ISSN 1364-0380 (on line) 1465-3060 (printed) **2013**

Geometry & Topology
Volume 9 (2005) 2013–2078
Published: 26 October 2005

Contact homology and one parameter families of Legendrian knots

Tamás Kálmán

Department of Mathematics, University of Southern California
Los Angeles, CA 90089, USA

Email: `tkalman@usc.edu`

Abstract

We consider S^1–families of Legendrian knots in the standard contact $\mathbf{R}^3$. We define the monodromy of such a loop, which is an automorphism of the Chekanov–Eliashberg contact homology of the starting (and ending) point. We prove this monodromy is a homotopy invariant of the loop (Theorem 1.1). We also establish techniques to address the issue of Reidemeister moves of Lagrangian projections of Legendrian links. As an application, we exhibit a loop of right-handed Legendrian torus knots which is non-contractible in the space $\mathrm{Leg}(S^1, \mathbf{R}^3)$ of Legendrian knots, although it is contractible in the space $\mathrm{Emb}(S^1, \mathbf{R}^3)$ of smooth knots. For this result, we also compute the contact homology of what we call the Legendrian closure of a positive braid (Definition 6.1) and construct an augmentation for each such link diagram.

AMS Classification numbers Primary: 53D40

Secondary: 57M25

Keywords: Legendrian contact homology, monodromy, Reidemeister moves, braid positive knots, torus knots

Proposed: Yasha Eliashberg Received: 3 October 2004
Seconded: Peter Ozsváth, Tomasz Mrowka Revised: 24 July 2005

1 Introduction

A Legendrian knot is an embedding $S^1 \hookrightarrow \mathbf{R}^3_{xyz}$ which is everywhere tangent to the kernel of the standard contact 1–form $\alpha = \mathrm{d}z - y\mathrm{d}x$. Two such knots are equivalent if they are homotopic through Legendrian knots. This is a refinement of the obvious subdivision according to classical knot type; just how much finer is a topic of current research. (We may refer to this as the π_0–problem of the space $\mathrm{Leg}(S^1, \mathbf{R}^3)$ of Legendrian knots). The integer-valued invariants tb (Thurston–Bennequin number) and r (rotation number) guarantee that every classical knot type contains infinitely many Legendrian equivalence classes. It was a major break-through when Chekanov [6] introduced a much more complicated, Floer-theoretical invariant which was able to distinguish between Legendrian knots of the same tb, r, and classical type (since then, Etnyre and Honda [11, 12, 13] have been able to prove similar results with different methods). The form in which we'll consider Chekanov's invariant is that of a system of homology groups, which we will call contact homology.

In this paper we use contact homology to tackle another, so called π_1–problem; namely, we use monodromy induced on this homology to provide the first examples that Legendrian and classical knots are different at the level of 1–parameter families.

1.1 Statements of results

In its most general formulation, Chekanov's invariant of Legendrian knots [6] takes the form of a stable tame isomorphism class of differential graded algebras, or DGA's for short, $(\mathcal{A}_*, \partial)$. What's more important for us is that in particular, the homology $\ker\partial / \operatorname{im}\partial$ is an invariant of the Legendrian knot type. This homology, which is actually also a graded algebra over $\mathbf{Z}_2$, will be referred to as contact homology and will be denoted by $H(L)$, where L is a (representative of a) Legendrian knot type. In [14], Etnyre, Ng, and Sabloff defined a natural extension of Chekanov's work, namely a DGA and corresponding contact homology with $\mathbf{Z}[t, t^{-1}]$–coefficients. We will carry out our proofs in this more general setting, even though in all of the applications known to the author the original $\mathbf{Z}_2$–version (which will be used in sections 6–10) would be sufficient.

The main result of the paper is the following:

Theorem 1.1 *For every Legendrian knot type $\mathcal{L} \subset \mathrm{Leg}(S^1, \mathbf{R}^3)$ and generic $L \in \mathcal{L}$, there exists a multiplicative homomorphism*

$$\mu\colon \pi_1(\mathcal{L}, L) \to \mathrm{Aut}(H(L)),$$

defined by continuation on the Chekanov–Eliashberg contact homology H. This is true with either $\mathbf{Z}_2$ or $\mathbf{Z}[t,t^{-1}]$–coefficients.

In particular, monodromy calculations are carried out in the homology and neither in its linearized [6] or abelianized [14] versions, nor in the DGA itself. Augmentations (which are the objects also needed to define linearized homology) do however play an important role in proving that μ is non-trivial. This will be demonstrated by the following example. Let $\mathcal{L}$ be the space of positive Legendrian (p,q) torus knots with maximal Thurston–Bennequin number (where $p,q \geq 2$ are relatively prime). $\mathcal{L}$ is connected by a result of Etnyre and Honda [11]. Consider the loop $\Omega_{p,q} \subset \mathcal{L}$ of Figure 1 (a more elaborate definition is in section 8). We claim the following:

Theorem 1.2 *$[\Omega_{p,q}] \in \pi_1(\mathcal{L})$ is either an element of infinite order or its order is divisible by $p+q$.*

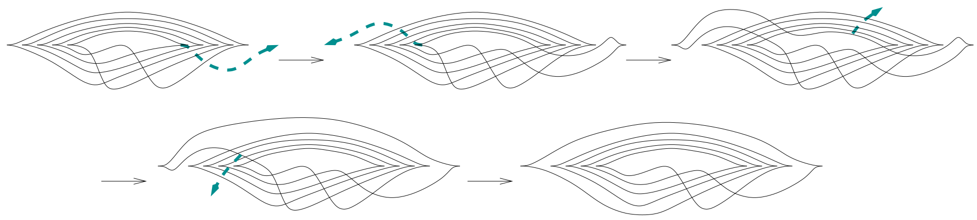

Figure 1: The loop $\Omega_{p,q}$ when $p=2$ and $q=5$, shown in the front projection

By Theorem 1.1, this follows immediately from the following:

Theorem 1.3 *The restricted monodromy $\mu_0(\Omega_{p,q}) = \mu(\Omega_{p,q})\big|_{H_0(L)}$ of the loop $\Omega_{p,q}$, where $H_0(L)$ is the index 0 part of the $\mathbf{Z}_2$–coefficient contact homology $H(L)$, has order $p+q$.*

In particular, $\mu(\Omega_{p,q})$ is different from the identity automorphism, and because $p+q$ doesn't divide $2p$, so is $(\mu(\Omega_{p,q}))^{2p}$. Thus by Theorem 1.1, $[\Omega_{p,q}]^{2p} \in \pi_1(\mathcal{L})$ is non-trivial. On the other hand, if $\mathcal{K}$ is the space of smooth positive (p,q) torus knots, then $[\Omega_{p,q}]^{2p} \in \pi_1(\mathcal{K})$ *is* trivial[1], because $\Omega_{p,q}$ is homotopic to a rotation by $2\pi/p$ radians and $\pi_1(SO(3)) = \mathbf{Z}_2$. So Theorem 1.2 had to be demonstrated via a specifically "Legendrian" technique, which is the monodromy invariant μ. We also have

[1]In fact, an elementary application of Hatcher's work [20] shows that if $p+q$ is odd, then $\Omega_{p,q}$ is itself contractible in $\mathcal{K}$, and if $p+q$ is even, then $[\Omega_{p,q}] \in \pi_1(\mathcal{K})$ has order 2.

Corollary 1.4 *There exist Legendrian knot types $\mathcal{L}$ so that for the corresponding smooth knot type $\mathcal{K} \supset \mathcal{L}$, the homomorphism $\pi_1(\mathcal{L}) \to \pi_1(\mathcal{K})$ induced by the inclusion is not injective.*

It is an open problem whether this homomorphism, for fixed $\mathcal{K}$ *and* $\mathcal{L}$, is always surjective. It is also worthwhile to compare our results to those of Benham, Lin, and Miller in [2]. They prove that if $\mathcal{W}$ is the set of knots in $\mathcal{K}$ with writhe $tb(\mathcal{L})$ (so that $\mathcal{L} \subset \mathcal{W} \subset \mathcal{K}$) then the inclusion $\mathcal{W} \hookrightarrow \mathcal{K}$ is a weak homotopy equivalence. For results about the homotopy type of $\mathcal{K}$, see Hatcher [20].

There are two widely used ways to represent Legendrian knots in the standard contact $\mathbf{R}^3$. In Figure 1, we used the front projection. For most of the paper, however, we will use Lagrangian projections, for the simple reason that H is naturally defined in terms of those. (Ng [23] gives a definition in terms of fronts.) Lagrangian projections are also more symmetric, but on the other hand less robust than fronts. By this I mean that certain linear constraints need to be satisfied both for the existence of an actual Legendrian lift of the diagram and for doing Reidemeister moves as well. While [6, section 11] contains a thorough description of the former phenomenon, our Theorem 4.1 appears to be the first account in the literature on the latter problem. We formulate a solution in terms of techniques and language of linear programming.

The paper is organized as follows. In section 1.2 we give some motivation and describe related work. We review some basic material in section 2. Section 3 contains the definition of monodromy and the proof of Theorem 1.1. Section 4 concludes the general discussion of one parameter families of Legendrian links by addressing Reidemeister moves.

Then we turn to examples; section 5 contains the simplest non-trivial one known to the author, a loop of Legendrian trefoils. In section 8 we include it in a much more general class of examples. Those loops consist of so called Legendrian closures of positive braids. Section 6 examines their base points; in particular, we compute their DGA. In section 7 we construct an augmentation for that DGA. Finally, we specialize our discussion again to the case when the Legendrian closure is a positive torus knot and use the last two sections to prove Theorem 1.3.

Acknowledgments This paper is almost identical to the author's PhD thesis that was submitted in 2004 at the University of California at Berkeley. It is my pleasure to thank the institution, and in particular my advisors, Rob Kirby and Michael Hutchings, for all the generous support that they provided. It was in fact Mike Hutchings who gave me the idea of this project. I would also like to thank Lenny Ng and Josh Sabloff for checking earlier versions of the manuscript, as well as them, Tobias Ekholm, Michael Sullivan and András Szűcs for stimulating conversations. I was supported by NSF grant DMS-0244558 during part of the research.

1.2 Floer theory

Loosely speaking, the idea of Floer theory is to formulate a Morse-type homology theory of certain real-valued smooth functionals on certain infinite-dimensional manifolds. The manifold and the functional are typically associated to an object such as a pair of Lagrangians, a symplectomorphism, a contact structure, a Legendrian submanifold etc. Some special cases of this project have had great success (alas, analytical details are sometimes missing) and we refer to these as "Floer theories." We should mention at least one case when the analogy with finite-dimensional Morse theory starts breaking down: in symplectic field theory [9], of which contact homology is a special case [14], generators of Floer homology are indeed critical points of an action functional, but flow lines (which are pseudoholomorphic curves in an appropriate sense) starting from a critical point may split and end at a finite collection of critical points (as opposed to just one endpoint). The natural way to include this phenomenon in our algebraic formulation of the theory is to write the multiple endpoints as a product. This is how the chain complex in this case takes the form of a differential graded *algebra* (DGA).

Besides finding new Floer theories, it is also an important and promising direction of research to generalize more and more features of finite-dimensional Morse theory to these new settings. One such feature can be the so called *continuation map*. Suppose (f_t, g_t), $t \in [0,1]$ is a 1–parameter family of smooth functions and Riemannian metrics on the finite-dimensional smooth manifold F. Suppose that the family is generic. This in particular means that the pairs (f_t, g_t) are Morse–Smale (ie, they have associated Morse homologies $H^{(t)}$) with finitely many exceptions $0 < t_1 < \ldots < t_k < 1$. Note that we can *not* speak of homologies $H^{(t_i)}$. Yet we may define a vector field on $[0,1] \times F$ by $V = t(1-t)\partial_t + V_t$, where V_t is the negative gradient of f_t with respect to g_t, and after observing that $\mathrm{Crit}_i(V) = (\{0\} \times \mathrm{Crit}_{i-1}(V_0)) \cup (\{1\} \times \mathrm{Crit}_i(V_1))$, we may count flow lines of V to define a chain map from the Morse complex of (f_0, g_0) to that of (f_1, g_1). The fact that it induces an isomorphism (which we call the continuation map) $H^{(0)} \to H^{(1)}$ can be established by a certain construction of chain homotopies [16], similar in flavor to the above[2]. This provides an alter-

[2]What has just been outlined is what I call the 'analytic approach.' Alternatively,

native proof that the isomorphism class of Morse homology depends only on the manifold, thus making Morse theory a self-contained homology theory.

But the usefulness of the continuation map doesn't stop here. Consider a fibration $\pi\colon Z \to S^1$ with smooth manifold fiber F, along with a smooth function $f\colon Z \to \mathbf{R}$ and a Riemannian metric g on Z. Suppose that for $1 \in S^1$, the pair $(f\big|_{\pi^{-1}(1)}, g\big|_{\pi^{-1}(1)})$ is Morse–Smale and the family $(f\big|_{\pi^{-1}(t)}, g\big|_{\pi^{-1}(t)})$ is generic. Then by continuation, we obtain an automorphism (called *monodromy*) $H^{(1)} \to H^{(1)}$. This is naturally equivalent to the automorphism of the singular homology of $\pi^{-1}(1)$ induced by the gluing map of the fibration (which is well defined up to homotopy, therefore its action is well defined). In particular, a non-trivial monodromy implies the non-triviality of the fibration.

As already demonstrated by Seidel [26] and Bourgeois [4], the continuation (monodromy) map can prove non-triviality theorems, in a way analogous to the above, about the fundamental group of the space of objects ("object" means a symplectomorphism in [26] and a contact structure in [4]), to which a version of Floer theory is associated[3]. To get such conclusions, it is of course a key point to prove that monodromy is *invariant* under homotopies of the loop of objects. In this paper we obtain such a result about the space of Legendrian knots $\mathrm{Leg}(S^1, \mathbf{R}^3)$.

There is also a generalized invariant of families of objects parametrized by an arbitrary finite-dimensional manifold, recently discovered by Hutchings [21]. The value of this is a spectral sequence. [21] contains its description in the case of finite-dimensional Morse theory, where the spectral sequence associated to a generic 4–tuple (Z, B, f, g) (as above, except that S^1 is replaced by the arbitrary finite-dimensional manifold B) turns out to be isomorphic to the Leray–Serre spectral sequence of the fibration. A paper [22] detailing the case of the Floer theory of symplectomorphisms, thus generalizing Seidel's results, is in preparation. In other Floer theories, apart from Bourgeois's results (where $B = S^n$) and the present paper (where $B = S^1$), the spectral sequence invariant

one may study the *bifurcations* that take place at $t = t_1, \ldots, t_k$ and define 'combinatorial' continuation maps $H^{(t_i-\varepsilon)} \to H^{(t_i+\varepsilon)}$, as well as chain homotopies to guarantee that these are isomorphisms. The approach taken in this paper will in fact be a generalization of the latter.

[3]Strictly speaking, Floer homology is associated only to generic objects. Once we have proven that its isomorphism class is invariant within a path-component of the objects, this distinction becomes immaterial, provided that we only wish to use Floer theory to separate such path-components (ie, we do π_0–theory). But the presence of non-generic objects is potentially an issue again when we investigate the topology of the path-components (like their π_1, as in this paper).

is yet to be either established (see the next paragraph for more on this when $B = S^1$) or exploited.

The motivation behind the group of results concerning one-parameter families is that Floer homology should not just be viewed as an invariant of certain objects, but should also be treated as a twisted coefficient system over the space of those objects; its non-trivial twisting is then capable of detecting the fundamental group of the moduli space. The existence of this twisted coefficient system, modulo an issue of fixing the signs in the continuation maps (see [21] for more), is an immediate consequence of what's already in the literature for many versions of Floer theory. These are the versions where they prove topological invariance by constructing continuation maps and chain homotopies using the analytic methods introduced in Floer's seminal paper [16] (as opposed to a study of bifurcations, which also appears in earlier work of Floer). In these cases the existence of a monodromy invariant, which is equivalent to the spectral sequence when $B = S^1$, is automatic too.

However, the invariance of contact homology was established in a different way, namely through a certain combinatorial "bifurcation analysis" [6]. Therefore we can use the analytic considerations mentioned so far only as a source of motivation that the monodromy invariant (and the spectral sequence invariant) exists in this context, even though one suspects that once monodromy is defined analytically, it will agree with our combinatorially defined version[4]. Following this "hint" then, we streamline Chekanov's original proof in section 3.1 (without changing its combinatorial nature) to emphasize the role of continuation more, and then in section 3.2, we continue the combinatorial study of bifurcations to provide the chain homotopies needed to prove that monodromy is a well-defined invariant.

2 Preliminaries

For a more thorough treatment and proofs of what follows see [6, 14, 10].

Throughout this paper, we work in the standard contact $\mathbf{R}^3$; we write the 1–form, whose kernel field is the contact distribution, as $\alpha = \mathrm{d}z - y\mathrm{d}x$. A curve $\lambda\colon S^1 \to \mathbf{R}^3$ is called a *Legendrian knot* if it is embedded (in particular,

[4]It *is* proven in [14] that Chekanov's combinatorially defined homology is isomorphic to the relative contact homology whose analytical definition was outlined in [9] and completed in [7]. For proof of invariance, however, the authors use bifurcation analysis in [7].

it has an everywhere non-zero tangent) and $\lambda^*(\alpha) = 0$. The space of such knots, modulo reparametrization, is denoted by $\mathrm{Leg}(S^1, \mathbf{R}^3)$. We will refer to connected components of $\mathrm{Leg}(S^1, \mathbf{R}^3)$ (with respect to the quotient of the C^∞ topology) as *Legendrian knot types*[5]. When talking about *classical* or *smooth knots*, we simply mean that the Legendrian assumption is not made. Modulo reparametrization, these form the moduli space $\mathrm{Emb}(S^1, \mathbf{R}^3)$, whose connected components are called *classical* or *smooth knot types*.

A (Legendrian or classical) *link* is a finite disjoint union of knots.

2.1 Link projections

The *Lagrangian projection* is $\pi\colon (x, y, z) \mapsto (x, y)$. If L is a Legendrian link, it admits an xy–projection that is an immersion with the additional properties described below in Proposition 2.2. We call L and $\gamma = \pi(L)$ *generic* if γ satisfies the usual assumptions for classical link diagrams, ie, it has no singularities other than finitely many transverse self-crossings.

The *front projection* is $(x, y, z) \mapsto (x, z)$. The image of a generic Legendrian link under this is a smooth curve with finitely many semicubical cusp points, no self-tangencies and no vertical tangents. As $y = \mathrm{d}z/\mathrm{d}x$, at every crossing the strand with smaller slope is the overcrossing one. Generic homotopies are composed of isotopies and a list of Reidemeister moves (see Figure 2). Unlike for Lagrangian projections, these moves can always be carried out without further restrictions. For example, the definition of $\Omega_{p,q}$ in Figure 1 (with $2(q-1)$ Reidemeister II moves and $2(q-1)$ Reidemeister II^{-1} moves) is sound. For the rest of the paper, we won't need to use fronts, except occasionally for motivation.

Fix an orientation of the Legendrian knot L. Let a be a simple transverse self-intersection (crossing) in its Lagrangian projection $\gamma = \pi(L)$ and denote the difference in z–coordinate between the two preimages of a by $h(a)$. We will refer to the positive number $h(a)$ as the *height* of a. Starting at the undercrossing, follow γ until we reach a again, this time on the upper strand. This path is called the *capping path* of a and we denote it by γ_a.

We use capping paths to define the *grading*, or *index* of a crossing, as follows. Assume that all crossings happen at a right angle. Then the rotation $r(\gamma_a)$ of the tangent to the capping path γ_a (with respect to the orientation $\mathrm{d}x \wedge \mathrm{d}y$

[5] We will use λ for actual maps and L for their equivalence classes modulo reparametrization, while $\mathcal{L}$ will be used to denote Legendrian knot types.

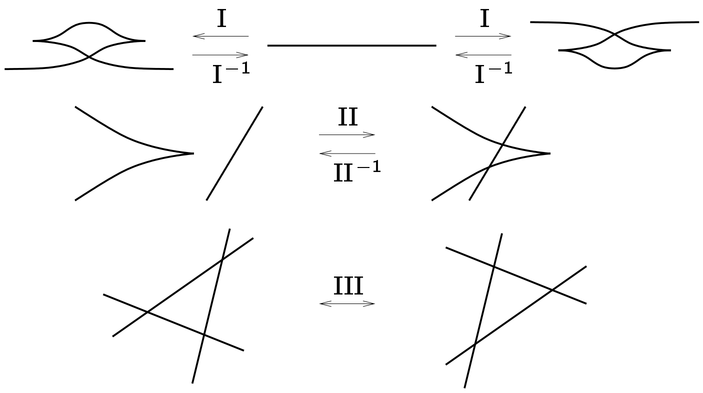

Figure 2: Reidemeister moves in the front projection; the list is completed by three more moves, which are just the reflections of the middle move in the vertical, horizontal, and in both, axes.

of the plane) is an odd multiple of $1/4$. The grading of the crossing a is the integer

$$|a| = -2r(\gamma_a) - \frac{1}{2}. \tag{1}$$

The *rotation number* r of a Legendrian knot is the Whitney index of its (oriented) Lagrangian projection. Its sign does depend on the choice of orientation. (However the modulo $2r$ residue of the grading of a crossing does not depend on the orientation. This residue is the value of Chekanov's original grading.) The *Thurston–Bennequin number* tb of a Legendrian knot is the writhe of its Lagrangian projection, ie, the number of crossings with even grading minus the number of those with odd grading[6]. It doesn't depend on the orientation of the knot. The parities of r and tb are opposite. The list of the three *classical invariants* of a Legendrian knot is completed by its *classical knot type.*

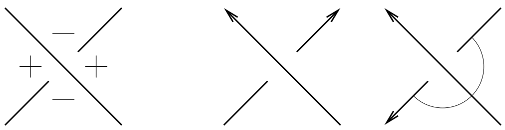

Figure 3: Reeb signs and orientation signs

The *Reeb signs* of the four quadrants surrounding a crossing are defined according to the scheme on the left of Figure 3. For the definition of the DGA with

[6]It is easy to verify that $(-1)^{|a|}$ coincides with the sign of the crossing a in classical knot theory. This fact is assumed in Figure 3, too.

$\mathbf{Z}[t,t^{-1}]$ coefficients in the next subsection, we will also need the concept of the so called *orientation sign*. This is $+1$ for all quadrants surrounding an oddly graded crossing and at evenly graded crossings, the two marked quadrants of Figure 3 have orientation sign -1, the other two have $+1$. The only property of the orientation signs that matters for the purposes of defining a homology theory is the following:

Lemma 2.1 [14] *Around the crossing a, the product of the orientation signs of any two opposite quadrants is $-(-1)^{|a|}$.*

Let $a_1, a_2, \dots, a_n$ be the crossings of a *classical* link diagram γ. Let us denote the bounded components of the complement by $U_1, U_2, \dots, U_{n+1}$. Let $S_\gamma = \bigoplus_{j=1}^{n} \mathbf{R}a_j$ and $V_\gamma = \bigoplus_{i=1}^{n+1} \mathbf{R}U_i$ be vector spaces with labeled bases and let S_γ^+ and V_γ^+ be their respective (open) positive cones. Chekanov defined the linear map $\Psi_\gamma \colon S_\gamma \to V_\gamma$ by the matrix $[E_{ij}]$, where $E_{ij} \in \{-2,-1,0,1,2\}$ is the number of quadrants at a_j that face U_i, counted with Reeb signs.

Proposition 2.2 *Suppose γ is the Lagrangian projection of a generic Legendrian link L. For the areas f_i of U_i, heights h_j of a_j, and incidence coefficients E_{ij} as described above, we have*

$$f_i = \sum_{j=1}^{n} E_{ij} h_j. \tag{2}$$

Proof First, it is an elementary fact that any arc $g = (x,y)\colon [0,1] \to \mathbf{R}^2_{xy}$ has Legendrian lifts $\lambda(t) = (g(t), z(t))$ in $\mathbf{R}^3$: let $z(0) = 0$ and define $z(t) = \int_0^t y(t)\mathrm{d}x(t)$. Further, any two such lifts differ by a translation in the z direction.

For each i, consider a piecewise continuously differentiable closed path $\kappa = (x,y)\colon [0,1] \to \mathbf{R}^2_{xy}$ that traverses the boundary of U_i once counterclockwise. Suppose it meets the crossings $b_1, \dots, b_s$ in this order (repetition in the list is allowed). Construct a Legendrian lift $\lambda(t) = (\kappa(t), z(t))$ of $\kappa(t)$. This is not a closed curve; in fact by Stokes' theorem, $z(1) - z(0) = \int_0^1 y(t)\mathrm{d}x(t) = \int_{U_i} \mathrm{d}y \wedge \mathrm{d}x = -f_i$. On the other hand, the same λ can also be constructed by gluing together z–translates of pieces of L. Inspecting Figure 3, it is easy to see that the total change of z along the latter is exactly $\pm h(b_1) \pm \cdots \pm h(b_s)$, where the counting happens with the opposite of the Reeb sign. □

We will also need the following observation from [6] later:

Lemma 2.3 *In the situation of Proposition 2.2,*

$$\sum_{j=1}^{n} E_{ij}|a_j| \equiv 2 - \#(\textit{positive corners of } U_i) \pmod{2r}.$$

Remark 2.4 The equations (2) imply the well known fact that a Lagrangian projection bounds zero area. Indeed, at each crossing, the sum of the indices with respect to γ of either pair of opposite regions (sharing the same Reeb sign) is the same.

Remark 2.5 By equation (2), the *height vector* $h = (h_1, h_2, \ldots, h_n) \in S_\gamma^+$ of a generic Lagrangian projection determines the *area vector* $f = (f_1, f_2, \ldots, f_{n+1})$. Moreover, h is such that $f = \Psi_\gamma h$ lies in V_γ^+. In addition to $h_i > 0$ for all i, this forces $n+1$ homogeneous linear inequalities for h. By *satisfying the linear constraints* we either mean that h is such that $h \in S_\gamma^+$ and $\Psi_\gamma h \in V_\gamma^+$, or that the pair $(h, f) \in S_\gamma^+ \times V_\gamma^+$ is such that (2) holds.

Definition 2.6 Let γ be a classical link diagram and consider the open convex cone

$$\mathcal{C}_\gamma = \{\, h \in S_\gamma^+ \mid \Psi_\gamma h \in V_\gamma^+ \,\}.$$

We say that γ is a *Lagrangian diagram* if $\mathcal{C}_\gamma$ is non-empty.

By the previous remark, Lagrangian projections of generic Legendrian links are Lagrangian diagrams. Lagrangian diagrams are the objects represented in our figures; in particular, we do not draw areas of regions to scale. All Lagrangian diagrams can, though, be *isotoped* (in the smooth sense) to actual Lagrangian projections (or, equivalently, we can use orientation-preserving diffeomorphisms of $\mathbf{R}^2$ to take one diagram to another). Moreover, *any* element of $\mathcal{C}_\gamma$ can be realized this way and because $\mathcal{C}_\gamma$ is convex, it is easy to see that all Lagrangian projections with the same underlying Lagrangian diagram can be isotoped into each other (this time in the Legendrian sense, ie, keeping the linear constraints intact throughout). In particular, a Lagrangian diagram represents a well-defined Legendrian link.

Example 2.7 Consider the knot diagram γ on the left side of Figure 4. We claim that $\mathcal{C}_\gamma$ is non-empty, thus γ is the Lagrangian diagram of a Legendrian knot L. The equations (2) take the following form (note the Reeb signs in

Figure 4):

$$\begin{array}{rclcccccccccc}
f_1 & = & & & h(a_2) & & & & & & \\
f_2 & = & -2h(a_1) & + & h(a_2) & + & h(b_1) & & & + & h(b_3) \\
f_3 & = & h(a_1) & & & & & & & & \\
f_4 & = & h(a_1) & & & - & h(b_1) & - & h(b_2) & - & h(b_3) \\
f_5 & = & & & & & h(b_1) & + & h(b_2) & & \\
f_6 & = & & & & & & & h(b_2) & + & h(b_3).
\end{array}$$

It is easy to check that the choice of $h(a_1) = 4$, $h(a_2) = 7$, $h(b_1) = h(b_2) = h(b_3) = 1$ yields the values $f_1 = 7$, $f_2 = 1$, $f_3 = 4$, $f_4 = 1$, and $f_5 = f_6 = 2$. As these are all positive, our linear constraints can indeed be satisfied.

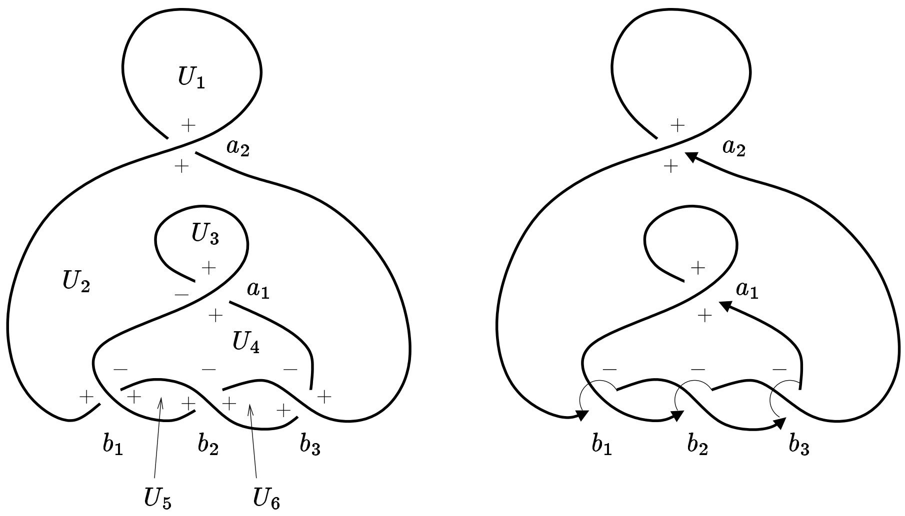

Figure 4: Right-handed Legendrian trefoil knot

The claim can be proven in a less direct way, too: the diagram γ is the result of the resolution [23] of a certain front diagram, analogous to the base point of the loop of Figure 1. See section 6 for more.

The rotation number of this knot (with either orientation, of course) is $r(L) = 0$. The gradings of the crossings (with either orientation, because $r = 0$) are $|a_1| = |a_2| = 1$ and $|b_1| = |b_2| = |b_3| = 0$. Hence the Thurston–Bennequin number is $tb(L) = 1$. Finally, the classical knot type of L is the right-handed trefoil knot.

Remark 2.8 Let p and q be arbitrary points on the same component C of the Legendrian link L. Then $\pi(p)$ and $\pi(q)$ divide $\gamma = \pi(C)$ into two arcs, γ_1 and γ_2. If γ_1' is another immersed arc with the same endpoints and starting and ending tangent vectors as γ_1, and so that $\gamma_1' \cup (-\gamma_1)$ bounds zero area, then $\gamma_1' \cup \gamma_2$ possesses a Legendrian lift (in general, an immersion) which coincides with C over γ_2. This diagrammatical method can be used to construct modifications of Legendrian links that only affect a small segment of the link, although possibly quite dramatically. We will use this idea in sections 4, 6, and 8.

2.2 Knot DGA and homology

Definition 2.9 Let R be a commutative ring. A *differential graded algebra*, or *DGA*, is a graded associative algebra $\mathcal{A} = \bigoplus_{i=-\infty}^{\infty} \mathcal{A}_i$ over R with identity $1 \in \mathcal{A}_0$, together with a differential $\partial\colon \mathcal{A} \to \mathcal{A}$. We will also refer to the grading as the *index*, which is multiplicative and takes values in a cyclic group Γ. We require that ∂ be an index -1, R–linear map so that $\partial(ab) = \partial(a)b + (-1)^{|a|}a\partial(b)$ for all $b \in \mathcal{A}$ and $a \in \mathcal{A}$ with pure index $|a|$ (consequently, $\partial(1) = 0$), and finally that $\partial^2 = 0$.

A DGA $\mathcal{A}$ is called *semi-free* (with generators $a_1, \dots, a_n$) if its underlying algebra is the ring (tensor algebra) $T(a_1, \dots, a_n)$ of non-commutative polynomials in a finite collection of pure-index elements $a_1, \dots, a_n \in \mathcal{A}$. We will use the term *degree* to refer to the degree of polynomials in this situation.

Throughout the paper, we will use the term polynomial to mean non-commutative polynomial.

Let L be an oriented, generic Legendrian knot in $\mathbf{R}^3$ with Lagrangian projection $\gamma \subset \mathbf{R}^2_{xy}$. Let the crossings of γ be $a_1, a_2, \dots, a_n$. We associate a semi-free DGA to L as follows.

The graded algebra Let $\mathcal{A}$ be the non-commutative, associative, unital polynomial algebra over $R = \mathbf{Z}[t, t^{-1}]$ generated by the symbols $a_1, a_2, \dots, a_n$. Grade each generator using the value $|a_i|$ defined by (1) and let $|t| = 2r$ (this value, ie twice the rotation number, is often called the *Maslov number* of the knot). Extend the grading to monomials multiplicatively, obtaining a $\mathbf{Z}$–grading so that we may write $\mathcal{A}_*$. Note that this grading, also called index, is different from the grading given by the degree of the polynomials.

For each positive integer k, fix an oriented disc Π_k (thus the boundary $\partial\Pi_k$ is also oriented) with k points $x_0^k, x_1^k, \dots, x_{k-1}^k$ marked on its boundary in the given cyclic order. A continuous map $f\colon \Pi_k \to \mathbf{R}^2_{xy}$ is *admissible with respect to*

γ (sometimes we'll say f is an *admissible disc in* γ) if it is an immersion away from the marked points, it preserves the orientation of the disc and moreover it takes $\partial\Pi_k$ to γ so that each marked point is mapped to a crossing in such a way that locally, the image of the disc forms an angle less than 180°. We also require that the quadrant occupied near the crossing $f(x_0^k)$ (the so called *positive corner*) has Reeb sign $+1$ and all other quadrants, at the so called *negative corners* $f(x_1^k),\dots,f(x_{k-1}^k)$, have Reeb sign -1. We'll also say that the disc *turns* at its corners $f(x_0^k), f(x_1^k),\dots,f(x_{k-1}^k)$, or that it turns at certain quadrants located at those points.

Note that we didn't require that the orientations of the curves γ and $f(\partial\Pi_k)$ match. Let us call an admissible immersion f *compatible with the orientation of* γ (or a *compatible disc*) if those two orientations either agree at all points of $\partial\Pi_k$ or disagree at all such points. We will need the following lemma in section 6.

Lemma 2.10 *Let γ be an oriented Lagrangian diagram. An admissible immersion f is compatible with the orientation of γ if and only if its positive corner has odd index and all of its negative corners have even indices.*

Proof An immersion is compatible if and only if the orientation requirement is satisfied in arbitrarily small neighborhoods of the punctures. Then the statement follows very easily from the observation (also mentioned in footnote 6) that $(-1)^{|a|}$ coincides with the sign of the crossing a in classical knot theory. Indeed, the middle diagram of Figure 3 shows an oddly graded crossing and we see (comparing with the diagram on the left) that a compatible disc may only turn at its positive quadrants. The diagram on the right shows an evenly graded crossing and a similar examination yields that compatible discs may only turn at its negative quadrants. □

The following lemma is a consequence of Proposition 2.2.

Lemma 2.11 *Let $f\colon \Pi_k \to \mathbf{R}^2_{xy}$ be an admissible immersion. Then*

$$h(f(x_0^k)) = \int_{\Pi_k} f^*(\mathrm{d}x\wedge \mathrm{d}y) + \sum_{i=1}^{k-1} h(f(x_i^k)).$$

In particular, $h(f(x_0^k)) > h(f(x_i^k))$ for all $i=1,\dots,k-1$.

The differential To compute $\partial(a)$ for a given crossing a, one considers all admissible immersions, modulo reparametrization, with positive corner a. A typical one described above contributes the monomial

$$t^{-l} \cdot \varepsilon_0 \varepsilon_1 \dots \varepsilon_{k-1} f(x_1^k) \dots f(x_{k-1}^k),$$

where $\varepsilon_0, \varepsilon_1, \dots, \varepsilon_{k-1}$ are the orientation signs (see Figure 3) of the quadrants occupied by $f(\Pi_k)$ at its corners, and l is defined as follows. Take the union of the oriented segments of $f(\partial \Pi_k)$ (the orientation is induced from that of $\partial \Pi_k$) with the capping path $\gamma_{f(x_0^k)} = \gamma_a$ and the reversed capping paths $-\gamma_{f(x_1^k)}, \dots, -\gamma_{f(x_{k-1}^k)}$. This is a cycle whose homology class in $H(L; \mathbf{Z}) \cong \mathbf{Z}$ is identified with the integer l via the chosen orientation of L.

$\partial(a)$ is the sum of these contributions over all admissible immersions. We extend ∂ to $\mathcal{A}$ using $\mathbf{Z}[t, t^{-1}]$–linearity and the signed Leibniz rule $\partial(ab) = \partial(a)b + (-1)^{|a|} a \partial(b)$.

Theorem 2.12 (Chekanov–Etnyre–Ng–Sabloff) *The sum taken when computing $\partial(a)$ for any crossing a is finite, hence $\partial\colon \mathcal{A} \to \mathcal{A}$ is well defined. It has index -1 and satisfies $\partial^2 = 0$.*

Definition 2.13 The semi-free DGA $(\mathcal{A}_*, \partial)$ is called the *Chekanov–Eliashberg differential graded algebra* associated to the generic Legendrian knot L. Its homology $H(L) = \ker(\partial)/\operatorname{im}(\partial)$ is the *Chekanov–Eliashberg contact homology*, or simply *contact homology* of L.

Chekanov's main theorem is that $H(L)$ (and further, a certain equivalence class of $(\mathcal{A}_*, \partial)$, called its stable tame isomorphism class) is unchanged when one passes to another generic Legendrian knot L' that is homotopic to L through Legendrian knots. (In this case we also say that L' is *Legendrian isotopic* to L.) We will state this result in Theorem 3.2.

Example 2.14 The trefoil DGA: please refer to Figure 4 and Example 2.7. We choose the orientation as on the right side of the figure. Quadrants with negative orientation sign are marked. Note that $|t| = 0$ and thus for index reasons, $\partial(b_1) = \partial(b_2) = \partial(b_3) = 0$. On the other hand,

$$\partial(a_1) = 1 - b_1 - b_3 - t b_1 b_2 b_3. \tag{3}$$

The admissible discs contributing these terms are U_3, $U_4 \cup U_5$, $U_4 \cup U_6$, and U_4, respectively. Similarly,

$$\partial(a_2) = 1 - b_2 + t^{-1} + b_2 b_3 + b_1 b_2 + t b_2 b_3 b_1 b_2, \tag{4}$$

with terms contributed by[7] U_1, $U_2 \cup U_3 \cup U_4$, $U_2 \cup 2U_4 \cup U_5 \cup U_6$, $U_2 \cup 2U_4 \cup U_5$, $U_2 \cup 2U_4 \cup U_6$, and $U_2 \cup 2U_4$, respectively. Note that all admissible immersions in γ are compatible with the orientation of the diagram. The following observation will be useful later in the proof of Theorem 1.2 for $p = 3,\ q = 2$:

$$\begin{aligned}\partial(a_2b_3 + (t^{-1} + b_2b_3)a_1) &= b_3 - b_2b_3 + t^{-1}b_3 + b_2b_3^2 + b_1b_2b_3 + tb_2b_3b_1b_2b_3\\ &\quad + t^{-1} - t^{-1}b_1 - t^{-1}b_3 - b_1b_2b_3\\ &\quad + b_2b_3 - b_2b_3b_1 - b_2b_3^2 - tb_2b_3b_1b_2b_3\\ &= b_3 + t^{-1} - t^{-1}b_1 - b_2b_3b_1.\end{aligned} \tag{5}$$

In other words, $b_3 + t^{-1} - t^{-1}b_1 - b_2b_3b_1 = 0$ in the contact homology.

We may recover Chekanov's original DGA by keeping the same generators and admissible discs, but putting $t = 1$, reducing coefficients modulo 2 and the grading modulo $2r$. For example, in this theory, (3) and (4) reduce to $\partial(a_1) = 1 + b_1 + b_3 + b_1b_2b_3$ and $\partial(a_2) = b_2 + b_2b_3 + b_1b_2 + b_2b_3b_1b_2$, respectively.

2.3 DGA maps and the product structure

Chain complexes of homology theories are primarily Abelian groups with a differential. The presence of the (non-commutative) product structure in Chekanov–Eliashberg theory is an 'extra' feature. In this section we point out how this does not actually complicate matters; namely, how we can concentrate on generators only and then let the algebra take care of itself. For more on general DGA theory, see [15].

Definition 2.15 Let $(\mathcal{A}, \partial)$ and $(\mathcal{B}, \partial')$ be differential graded algebras (over the same commutative ring R and graded by the same cyclic group Γ). A linear map $\varphi\colon \mathcal{A} \to \mathcal{B}$ is a *chain map* if it is index-preserving and intertwines the differentials: $\varphi(\partial(x)) = \partial'(\varphi(x))$ for all $x \in \mathcal{A}$. If such a φ is also an algebra homomorphism, then we call it a *DGA morphism.*

Definition 2.16 Let $(\mathcal{A}, \partial)$ and $(\mathcal{B}, \partial')$ be differential graded algebras (over the same R, graded by the same Γ), and $\varphi, \psi\colon \mathcal{A} \to \mathcal{B}$ chain maps between

[7] This way of describing the actual immersions works fine in this case. In general however it is possible that two different admissible discs cover the same regions with the same multiplicities. As to whether we indeed listed all admissible discs, the reader can check that as an exercise or can refer to section 6, where we shall determine all admissible discs in a certain class of Lagrangian diagrams generalizing γ.

them. An index r map $S\colon \mathcal{A} \to \mathcal{B}$ is called a *(φ,ψ)–derivation of index r* if for all $a \in \mathcal{A}$ of pure index $|a|$ and for all $b \in \mathcal{A}$, we have

$$S(ab) = S(a)\psi(b) + (-1)^{r|a|}\varphi(a)S(b).$$

If $(\mathcal{A},\partial) = (\mathcal{B},\partial')$ and $\varphi = \psi = \mathrm{id}_{\mathcal{A}}$, then we simply speak of a derivation of index r.

In this paper we consider three types of DGA mappings. They are all R–linear, but each has a different relation to the product structure of the DGA. Namely, by definition, the *differential* itself is a derivation of index -1. All *chain maps* that we consider are going to be DGA morphisms (in particular, they fix constants), and all *chain homotopies* connecting chain maps φ, ψ will be (φ,ψ)–derivations of index 1 (from which it follows that they take constants to 0). The following two lemmas show that these self-imposed restrictions are very natural and useful.

Lemma 2.17 *Let $(\mathcal{A},\partial)$ and $(\mathcal{B},\partial')$ be semi-free differential graded algebras over a commutative ring R, graded by Γ. If an index 0 algebra homomorphism $\eta\colon \mathcal{A} \to \mathcal{B}$ is such that $\eta(\partial(a)) = \partial'(\eta(a))$ for all generators of $\mathcal{A}$, then the same holds for all elements $a \in \mathcal{A}$ (ie, η is a chain map).*

Moreover, any assignment of values $\eta(a) \in \mathcal{B}$ to the generators a of $\mathcal{A}$ which satisfies $|\eta(a)| = |a|$ for all a can be uniquely extended to an index 0 algebra homomorphism.

Proof To complete a trivial induction proof, we need to show that if a and b are homogeneous elements of $\mathcal{A}$ that satisfy the desired property, then such is ab, too:

$$\begin{aligned}\eta(\partial(ab)) &= \eta\left(\partial(a)b + (-1)^{|a|}a\partial(b)\right) = \eta(\partial(a))\eta(b) + (-1)^{|a|}\eta(a)\eta(\partial(b))\\ &= \partial'(\eta(a))\eta(b) + (-1)^{|\eta(a)|}\eta(a)\partial'(\eta(b)) = \partial'(\eta(a)\eta(b)) = \partial'(\eta(ab)).\end{aligned}$$

In the last assertion, the existence of an extension is obvious, since $\mathcal{A}$ is semi-free with the given generators and $\mathcal{B}$ is associative. We just need to check that the unique extension is also of index 0, which is established with the same induction argument based on noting that if $|\eta(a)| = |a|$ and $|\eta(b)| = |b|$, then $|\eta(ab)| = |\eta(a)\eta(b)| = |\eta(a)| + |\eta(b)| = |a| + |b| = |ab|$. □

Lemma 2.18 *Let $(\mathcal{A},\partial)$ and $(\mathcal{B},\partial')$ be semi-free differential graded algebras over the commutative ring R, graded by Γ, and let $\varphi,\psi\colon \mathcal{A} \to \mathcal{B}$ be DGA* *morphisms. If $K\colon \mathcal{A} \to \mathcal{B}$ is a (φ,ψ)–derivation of index 1 so that $K(\partial(a)) + \partial'(K(a)) = \varphi(a) - \psi(a)$ for all generators of $\mathcal{A}$, then the same holds for all elements $a \in \mathcal{A}$ (ie, K is a chain homotopy between φ and ψ).*

Moreover, any assignment of values $K(a) \in \mathcal{B}$ to the generators a of $\mathcal{A}$ which satisfies $|K(a)| = |a| + 1$ for all a can be uniquely extended to a (φ,ψ)–derivation of index 1.

Proof The first part of the proof is a similar, but longer computation as in the previous argument:

$$
\begin{aligned}
(K\partial + \partial' K)(ab) &= K(\partial(ab)) + \partial'(K(ab)) \\
&= K\left(\partial(a)b + (-1)^{|a|}a\partial(b)\right) + \partial'\left(K(a)\psi(b) + (-1)^{|a|}\varphi(a)K(b)\right) \\
&= K(\partial(a))\psi(b) + (-1)^{|a|-1}\varphi(\partial(a))K(b) \\
&+ (-1)^{|a|}K(a)\psi(\partial(b)) + (-1)^{2|a|}\varphi(a)K(\partial(b)) \\
&+ \partial'(K(a))\psi(b) + (-1)^{|a|+1}K(a)\partial'(\psi(b)) \\
&+ (-1)^{|a|}\partial'(\varphi(a))K(b) + (-1)^{|a|+|\varphi(a)|}\varphi(a)\partial'(K(b)) \\
&= \big(K(\partial(a)) + \partial'(K(a))\big)\,\psi(b) + (-1)^{|a|}\big(-\varphi(\partial(a)) + \partial'(\varphi(a))\big)\,K(b) \\
&+ (-1)^{|a|}K(a)\,\big(\psi(\partial(b)) - \partial'(\psi(b))\big) + \varphi(a)\,\big(K(\partial(b)) + \partial'(K(b))\big) \\
&= (\varphi(a) - \psi(a))\psi(b) + 0 + 0 + \varphi(a)(\varphi(b) - \psi(b)) \\
&= \varphi(a)\varphi(b) - \psi(a)\psi(b) = \varphi(ab) - \psi(ab).
\end{aligned}
$$

In the second part, to show that the extension is unique, we need to establish that $K((ab)c) = K(a(bc))$, which is due to the following direct computation:

$$
\begin{aligned}
K((ab)c) &= K(ab)\psi(c) + (-1)^{|a|+|b|}\varphi(ab)K(c) \\
&= K(a)\psi(b)\psi(c) + (-1)^{|a|}\varphi(a)K(b)\psi(c) + (-1)^{|a|+|b|}\varphi(a)\varphi(b)K(c). \\
K(a(bc)) &= K(a)\psi(bc) + (-1)^{|a|}\varphi(a)K(bc) \\
&= K(a)\psi(b)\psi(c) + (-1)^{|a|}\varphi(a)K(b)\psi(c) + (-1)^{|a|}(-1)^{|b|}\varphi(a)\varphi(b)K(c).
\end{aligned}
$$

The extension is also index 1 by the already usual induction argument: If a and b are homogeneous with $|K(a)| = |a| + 1$ and $|K(b)| = |b| + 1$, then $|K(ab)| = |K(a)\psi(b) + (-1)^{|a|}\varphi(a)K(b)|$, where $|\psi(b)| = |b|$ and $|\varphi(a)| = |a|$, so $K(ab)$ is homogeneous of index $|a| + |b| + 1 = |ab| + 1$. □

Table 1 summarizes the relevant algebraic machinery. For all three rows, it is true that if the 'map' is defined on the generators so that the 'index' requirement is satisfied and then if it is extended to the whole DGA using linearity and the 'product identity,' then

map	property	index	product identity
differential ∂	$\partial^2 = 0$	-1	$\partial(ab) = \partial(a)b + (-1)^{\lvert a \rvert} a\partial(b)$
chain map η	$\eta\partial = \partial'\eta$	0	$\eta(ab) = \eta(a)\eta(b)$
chain homotopy K between chain maps φ, ψ	$K\partial + \partial' K$ $= \varphi - \psi$	1	$K(ab) = K(a)\psi(b) + (-1)^{\lvert a \rvert}\varphi(a)K(b)$

Table 1: Product identities

(1) such an extension is unique, it satisfies the 'index' requirement, and

(2) if the 'property' requirement holds for the generators, then it holds for all elements of the DGA.

3 The monodromy invariant

Let us first outline the contents of the section. Let $\mathcal{L} \subset \mathrm{Leg}(S^1, \mathbf{R}^3)$ be a Legendrian knot type. We call an element of $\mathcal{L}$ *generic*, as in section 2, if its Lagrangian projection is generic (this is equivalent to saying that the Legendrian knot doesn't have degenerate Reeb chords). These are the ones to which a graded chain complex $(\mathcal{A}_*, \partial)$, and a corresponding Floer homology H are actually associated. Non-generic objects form a codimension 1 discriminant $\mathcal{D} \subset \mathcal{L}$.

Generic homotopies, ie paths in $\mathcal{L}$, meet only the codimension 1 (codimension is always meant with respect to $\mathcal{L}$) stratum $\mathcal{D}_1 \subset \mathcal{D}$ and diagrammatically, this corresponds to Reidemeister II, II^{-1}, or III moves (see Figure 6). If L and L' are in adjacent chambers of $\mathcal{L} \setminus \mathcal{D}$, then following [6] (and [14] for the case of $\mathbf{Z}[t, t^{-1}]$–coefficients), one is able to write down a chain map $\varphi\colon \mathcal{A}(L) \to \mathcal{A}(L')$, corresponding to the Reidemeister move, that induces an isomorphism $\varphi_*\colon H(L) \to H(L')$. (In [6, 14] they use these chain maps to prove that the homology is independent of the actual projection.) Composing these chain maps/isomorphisms, every generic path $\Phi(t)$ that connects generic objects L_0 and L_1 has an induced isomorphism $\Phi_*\colon H(L_0) \to H(L_1)$, called the holonomy of the path. Details are found in section 3.1.

Now to prove that the holonomy only depends on the homotopy class (with fixed endpoints) of $\Phi(t)$, one has to consider two kinds of events: paths that are tangent to $\mathcal{D}_1$ (this is easy) and paths that meet $\mathcal{D}_2$. To deal with the latter, we take a small generic loop $\Psi(t)$, starting and ending at the generic element L, and linking the codimension 2 discriminant $\mathcal{D}_2 \subset \mathcal{D}$ once. While it is not necessarily the case that the induced *chain map* $\psi\colon \mathcal{A}_*(L) \to \mathcal{A}_*(L)$ is the identity, we will establish that there is always a *chain homotopy* $K\colon \mathcal{A}_*(L) \to \mathcal{A}_{*+1}(L)$ connecting ψ and $\mathrm{id}_{\mathcal{A}_*(L)}$, thereby proving that *on the homology level* $\psi_* = \Psi_* = \mathrm{id}_{H(L)}$, as desired.

In Figure 5, we represented a hypothetical case when two strands of the discriminant meet transversally. (As it is shown in [1, Figure 6] and as we'll see in section 3.2, the local description at other points of $\mathcal{D}_2$ may be slightly less or more complicated, but the approach is the same.) As we know that all arrows on the 'top level' represent isomorphisms, we may choose any two objects, which may even coincide, and direct all arrows from the first to the second. If this choice of directions can be accompanied with a chain homotopy on the 'middle level,' and in section 3.2 we prove that this is always the case, then we indeed get that the diagram of holonomies on the 'top level' commutes.

In section 3.3 we review our formulas in their simplified, $\mathbf{Z}_2$–coefficient versions.

3.1 Chain maps

Generic paths in $\mathrm{Leg}(S^1, \mathbf{R}^3)$ are isotopies with a finite number of Reidemeister moves. If we work with Lagrangian projections, then there are four different kinds of the latter, depicted in Figure 6. $\mathrm{III_a}$ and $\mathrm{III_b}$ moves are also called *triangle moves.* Move II is sometimes referred to as a *birth* and move II^{-1} as a *death* (of a pair of crossings).

Let us denote the DGA's of the diagrams in the upper row of Figure 6 by $(\mathcal{A}, \partial)$ and in the lower row, by $(\mathcal{A}', \partial')$. The crossings not affected by the moves are in an obvious one-to-one correspondence. If $x \in \mathcal{A}$ is one of these, then the corresponding crossing in $\mathcal{A}'$ will be denoted by x'.

Before stating Theorem 3.2, let us review Chekanov's analysis, although in slightly different terms, of the bifurcation that takes place on the right side of Figure 6. The re-phrasing is possible because we shall only be concerned with the homology, and not the stable tame isomorphism class of the DGA, as our invariant. Let the crossings that generate $\mathcal{A}$ be arranged by height as follows:

$$h(a_l) \geq \ldots \geq h(a_1) \geq h(a) > h(b) \geq h(b_1) \geq \ldots \geq h(b_m).$$

Let $\partial(a) = \varepsilon_a \varepsilon_b b + v$, where ε_a and ε_b are the orientation signs from Figure 6 (it is easy to prove that the exponent of t is 0 in the term coming from the

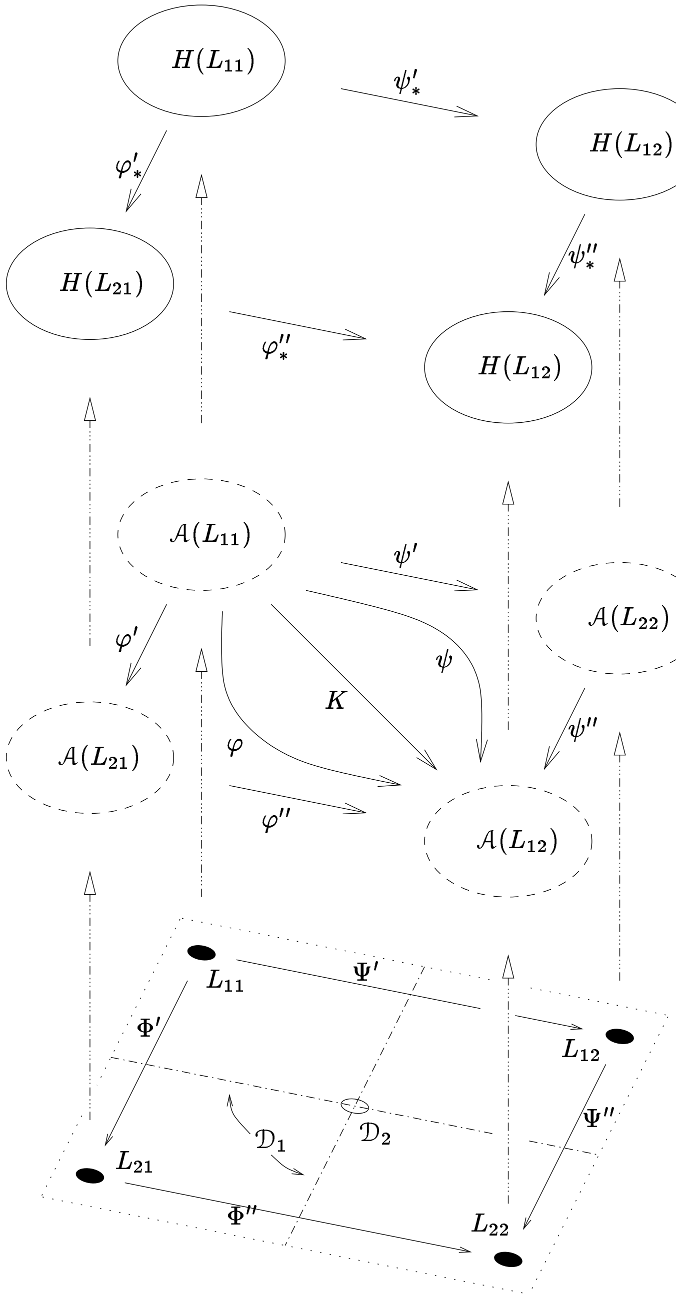

Figure 5: Chain complexes and homologies of different objects

vanishing bigon) and by Lemma 2.11, $v \in T(b_1, \dots, b_m)$. Define

$$\tau\colon \mathcal{A} \to \mathcal{A}'$$

by putting $\tau(a_i) = a'_i$ $(1 \le i \le l)$, $\tau(b_j) = b'_j$ $(1 \le j \le m)$, $\tau(a) = 0$, and $\tau(b) = \varepsilon'_a \varepsilon'_b v'$, where ε'_a and ε'_b are the other two orientation signs in Figure 6 and v' is obtained from v in the obvious way of putting primes on all the b_j's in the expression. Note that these assignments preserve the grading, so by the second statement of Lemma 2.17, we can extend τ to a well-defined index 0 algebra homomorphism. The geometric content of the bifurcation is captured by the following result of Chekanov:

Lemma 3.1 *$\tau\colon \mathcal{A} \to \mathcal{A}'$ is a chain map.*

One of the statements of Theorem 3.2 is that τ induces an isomorphism of homologies. Our goal now is to define what will be its homotopy inverse, that is another chain map[8] $\varphi\colon \mathcal{A}' \to \mathcal{A}$. While it will be the case that $\tau \circ \varphi = \mathrm{id}_{\mathcal{A}'}$, $\varphi \circ \tau$ and $\mathrm{id}_{\mathcal{A}}$ will only be chain homotopic via a certain $K\colon \mathcal{A} \to \mathcal{A}$. In fact, we will build φ and K with a somewhat subtle simultaneous construction[9], using the filtration that we will introduce next.

Define

$$\mathcal{A}_i = T(b_1, \dots, b_m, b, a, a_1, \dots, a_i)$$

and

$$\mathcal{A}'_i = T(b'_1, \dots, b'_m, a'_1, \dots, a'_i)$$

for all $i = 0, 1, \dots, l$. Note that by Lemma 2.11, $\partial(a_i) \in \mathcal{A}_{i-1}$ and $\partial(a'_i) \in \mathcal{A}'_{i-1}$. It follows that $\left(\mathcal{A}_i, \partial\big|_{\mathcal{A}_i}\right)$ and $\left(\mathcal{A}'_i, \partial'\big|_{\mathcal{A}'_i}\right)$ are DGA's themselves[10] for all $i = 0, 1, \dots, l$. Furthermore, for all such i, the restriction

$$\tau\colon \mathcal{A}_i \to \mathcal{A}'_i$$

is a DGA morphism. Define the index 0 algebra homomorphism

$$\varphi\colon \mathcal{A}'_0 \to \mathcal{A}_0$$

by $\varphi(b'_j) = b_j$ $(1 \le j \le m)$. By Lemma 2.17, this is also a chain map, for $\varphi(\partial'(b'_j)) = \varphi(\partial'(\tau(b_j))) = \varphi(\tau(\partial(b_j))) = \partial(b_j) = \partial(\varphi(b'_j))$. Here we used

[8] All of the considerations leading up from this point to Theorem 3.2 are purely algebraic. The chain map φ is less explicit in [6] and [14] than in what follows.

[9] In Remark 3.4, we express φ in even more explicit terms and without referring to K. That description is needed for applications, but would be impractical in the theoretical discussion.

[10] It is not true though that, for example, ∂ would map $\mathcal{A}_i$ to $\mathcal{A}_{i-1}$; $\partial(a_i^2)$ is usually not in $\mathcal{A}_{i-1}$.

Lemma 3.1 and that by Lemma 2.11, $\partial(b_j)$ can't contain a or b, and finally, the obvious fact that $\varphi\circ\tau\big|_{T(b_1,\dots,b_m)} = \mathrm{id}_{T(b_1,\dots,b_m)}$.

Now, $\varphi\circ\tau$ and $\mathrm{id}_{\mathcal{A}_0}$ are both DGA endomorphisms of $\mathcal{A}_0$. By the second statement of Lemma 2.18, the last three rows of the assignment

$$\begin{array}{rcll} a_i & \mapsto & 0 & \quad (i=1,\dots,l) \\ a & \mapsto & 0 & \\ b & \mapsto & \varepsilon'_a\varepsilon'_b a & \\ b_j & \mapsto & 0 & \quad (j=1,\dots,m) \end{array} \tag{6}$$

can be uniquely extended to a $(\varphi\circ\tau,\mathrm{id}_{\mathcal{A}_0})$–derivation

$$K\colon \mathcal{A}_0\to\mathcal{A}_0$$

of index 1; note that indeed, because ∂ is of index -1, we have $|a|=|b|+1$. Then we use the first part of Lemma 2.18 to show that K is a chain homotopy between $\varphi\circ\tau$ and $\mathrm{id}_{\mathcal{A}_0}$: For the generators $b_1,\dots,b_m$, the condition holds by the observation at the end of the previous paragraph, together with the fact, following from the definition, that the restriction of K to the DGA $T(b_1,\dots,b_m)$ is 0. For a and b, it is established below:

$$\begin{aligned} K(\partial(a))+\partial(K(a)) &= K(\varepsilon_a\varepsilon_b b+v)+0 \\ &= \varepsilon_a\varepsilon_b\varepsilon'_a\varepsilon'_b a+K(v) = (-1)^{|a|}(-1)^{|b|}a+0=-a \end{aligned}$$

by Lemma 2.1, and

$$K(\partial(b))+\partial(K(b)) = 0+\partial(\varepsilon'_a\varepsilon'_b a)=\varepsilon'_a\varepsilon'_b(\varepsilon_a\varepsilon_b b+v)=-b+\varepsilon'_a\varepsilon'_b v.$$

On the other hand,

$$\varphi(\tau(a))-a=\varphi(0)-a=-a \text{ and } \varphi(\tau(b))-b=\varphi(\varepsilon'_a\varepsilon'_b v')-b=\varepsilon'_a\varepsilon'_b v-b.$$

Next, we use a recursive process to define extensions of φ and K to the whole of $\mathcal{A}'$ and $\mathcal{A}$, respectively. Suppose that for some value $0\le i<l$, the DGA morphism $\varphi\colon\mathcal{A}'_i\to\mathcal{A}_i$ and the $(\varphi\circ\tau,\mathrm{id}_{\mathcal{A}_i})$–derivation $K\colon\mathcal{A}_i\to\mathcal{A}_i$ of index 1 are already defined so that K is a chain homotopy between $\varphi\circ\tau\big|_{\mathcal{A}_i}$ and $\mathrm{id}_{\mathcal{A}_i}$. Extend φ to $\mathcal{A}_{i+1}$ by letting

$$\varphi(a'_{i+1})=a_{i+1}+K(\partial(a_{i+1})). \tag{7}$$

As this preserves the grading of a'_{i+1}, by Lemma 2.17 (second part) there does exist an extension, which is an index 0 algebra homomorphism $\varphi\colon\mathcal{A}'_{i+1}\to\mathcal{A}_{i+1}$. To see that it's also a DGA morphism, we use Lemma 2.17 (first part), the inductive hypothesis, and the following computation:

$$\begin{aligned} \varphi(\partial'(a'_{i+1})) &= \varphi(\partial'(\tau(a_{i+1})))=\varphi(\tau(\partial(a_{i+1}))) \\ &= (K\circ\partial+\partial\circ K)(\partial(a_{i+1})+\partial(a_{i+1}) \\ &= \partial\left(K(\partial(a_{i+1}))+a_{i+1}\right)=\partial(\varphi(a'_{i+1})), \end{aligned}$$

which works because $\partial(a_{i+1}) \in \mathcal{A}_i$ and $\partial^2 = 0$.

Now we are in a position to use the DGA morphisms $\varphi \circ \tau \colon \mathcal{A}_{i+1} \to \mathcal{A}_{i+1}$ and $\mathrm{id}_{\mathcal{A}_{i+1}}$, along with the relevant values from (6) to define the $(\varphi \circ \tau, \mathrm{id}_{\mathcal{A}_{i+1}})$–derivation $K \colon \mathcal{A}_{i+1} \to \mathcal{A}_{i+1}$ of index 1, which is obviously an extension of the previous K. To make sure that K is also a chain homotopy between $\varphi \circ \tau$ and $\mathrm{id}_{\mathcal{A}_{i+1}}$, we still have to check that for a_{i+1}, but it's a tautology: $(K \circ \partial + \partial \circ K)(a_{i+1}) = K(\partial(a_{i+1})) = \varphi(a'_{i+1}) - a_{i+1} = \varphi(\tau(a_{i+1})) - a_{i+1}$.

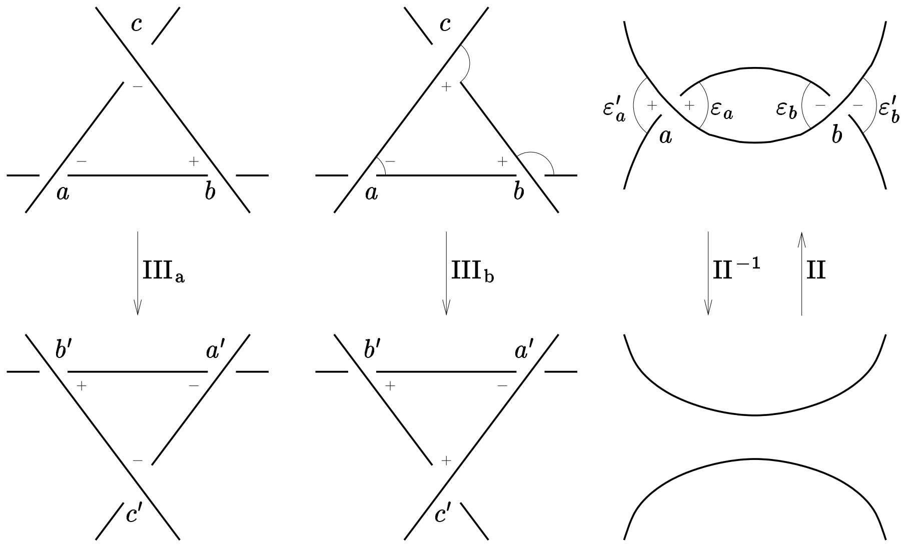

Figure 6: Legendrian Reidemeister moves in the Lagrangian projection. The $+$ and $-$ signs are Reeb signs and ε_a, ε'_a, ε_b, ε'_b refer to orientation signs.

Theorem 3.2 (Chekanov) *The following maps, when extended to $\mathcal{A}$ as algebra homomorphisms, are chain maps that induce isomorphisms (as graded algebras) of the corresponding Chekanov–Eliashberg contact homologies. (Recall that $(\mathcal{A}, \partial)$ is the DGA of the upper row, and $(\mathcal{A}', \partial')$ is the DGA of the lower row in Figure 6.)*

(1) *Move III$_\mathrm{a}$: Let $a \mapsto a'$, $b \mapsto b'$, $c \mapsto c'$, and $x \mapsto x'$, where x is any generator of $\mathcal{A}$ not shown in Figure 6.*

(2) *Move III$_\mathrm{b}$: Let*

$$a \mapsto a' - \varepsilon_a \varepsilon_b \varepsilon_c t^k c' b',$$

where ε_a, ε_b, and ε_c are the orientation signs of the quadrants indicated in Figure 6 and $k = k(a; b, c) \in \mathbf{Z}$ is the winding, with respect to the *chosen orientation of the knot, of the union of capping paths $-\gamma_a+\gamma_b+\gamma_c$ (thought of as starting and ending at the single point that the triangle shrinks to). Other generators are mapped trivially: $b \mapsto b'$, $c \mapsto c'$, and $x \mapsto x'$.*

(3) *Move II^{-1}: Let $\partial(a) = \pm b + v$ (recall that v does not contain neither a nor b). Define $x \mapsto x'$, which gives rise to the obvious re-labeling $v \mapsto v'$. Then let*

$$\begin{aligned} a &\mapsto 0 \\ b &\mapsto \mp v' \end{aligned}$$

(this is the map τ of the preceding discussion).

(4) *Move II: The map $\varphi\colon \mathcal{A}' \to \mathcal{A}$, constructed in the preceding paragraphs (for a more explicit description, see Remark 3.4).*

Proof We cite [6] and [14] that the maps associated to $\mathrm{III_a}$, $\mathrm{III_b}$, and II^{-1} moves are indeed chain maps (cf Lemma 3.1), and we have already seen that φ is one. In each of the four cases, the reverse of the move is again a Reidemeister move, with a corresponding chain map associated to it. It is enough to prove that the relevant pairwise compositions of these chain maps are chain homotopic to the identity. This is obvious for the pair of $\mathrm{III_a}$ moves on the left side of Figure 6. For the pair of $\mathrm{III_b}$ moves in the middle, note that the reverse move sends

$$a' \mapsto a - \varepsilon'_a\varepsilon'_b\varepsilon'_c t^k cb,$$

where ε'_a, ε'_b, and ε'_c are the orientation signs of the quadrants opposite to the previously considered ones, while the value of k is the same. So, since $\varepsilon_a\varepsilon_b\varepsilon_c\varepsilon'_a\varepsilon'_b\varepsilon'_c = [-(-1)^{|a|}][-(-1)^{|b|}][-(-1)^{|c|}] = -1$ by Lemmas 2.1 and 2.3, the chain maps themselves are inverses of each other.

Next, we claim that $\tau \circ \varphi = \mathrm{id}_{\mathcal{A}'}$. As both sides are DGA morphisms, it's enough to check for generators. The claim is obvious for the b'_j and follows from the definition for the a'_i: $\tau(\varphi(a'_i)) = \tau(a_i + K(\partial(a_i))) = \tau(a_i) = a'_i$, because $\tau \circ K = 0$. This last assertion is true because it's satisfied by all generators and $\tau \circ K$ and 0 are both $(\tau \circ \varphi \circ \tau, \tau)$–derivations of index 1.

Finally, it is only the composition $\varphi \circ \tau$ that requires a non-zero chain homotopy, but we have already constructed such a K before stating the theorem. □

Definition 3.3 We will refer to the graded algebra isomorphisms induced by the maps described in Theorem 3.2 as the *holonomies* of the corresponding Reidemeister moves. Sometimes the chain maps themselves will be called holonomies, too. Along generic paths, finitely many Reidemeister moves happen and the composition of the corresponding holonomies is called the holonomy of the path. When such a sequence of Reidemeister moves returns to the original starting knot, the holonomy is referred to as the *monodromy* of the loop.

Remark 3.4 For the purposes of computing holonomies (and monodromies), one should understand the map $\varphi\colon \mathcal{A}' \to \mathcal{A}$ in more concrete terms. Of course, we have

$$\varphi(b_j') = b_j \text{ for all } j = 1, \dots, m. \tag{8}$$

Next, recall that $\varphi(a_1') = a_1 + K(\partial(a_1))$. Let us write

$$\partial(a_1) = \sum \kappa B_1 b B_2 b B_3 b B_4 b \dots B_k b A,$$

where $k \geq 0$, $\kappa \in \mathbf{Z}[t, t^{-1}]$, $B_1, B_2, \dots, B_k \in T(b_1, \dots, b_m)$ are monomials, and in the monomial $A \in \mathcal{A}_0$, every b factor is preceded by an a factor (the point here is that the b factor right before A is the *last* b *before the first* a; if there is no a at all, then it's just the last b in the word). The sum is taken over all admissible discs with positive corner at a_1. Then it is not hard to compute, applying the Leibniz-type rule of Table 1 that we used to define K (recall also that $\partial(a) = \varepsilon_a \varepsilon_b b + v = -\varepsilon_a' \varepsilon_b' b + v$), that

$$\begin{aligned} \varphi(a_1') = a_1 + \sum \kappa \Big(& (-1)^{|B_1|} \varepsilon_a' \varepsilon_b' B_1 a B_2 b B_3 b B_4 b \dots B_k b A \\ & + (-1)^{|B_1|+|B_2|} (-1)^{|b|} B_1 v B_2 a B_3 b B_4 b \dots B_k b A \\ & + (-1)^{|B_1|+|B_2|+|B_3|} \varepsilon_a' \varepsilon_b' B_1 v B_2 v B_3 a B_4 b \dots B_k b A \\ & + (-1)^{|B_1|+|B_2|+|B_3|+|B_4|} (-1)^{|b|} B_1 v B_2 v B_3 v B_4 a \dots B_k b A \\ & + \dots \\ & + (-1)^{|B_1|+\dots+|B_k|} (-1)^{|b|} \big(\varepsilon_a' \varepsilon_b' (-1)^{|b|} \big)^k B_1 v B_2 v B_3 v B_4 v \dots B_k a A \Big). \end{aligned} \tag{9}$$

With this expression in hand, we may compute $\varphi(a_2') = a_2 + K(\partial(a_2))$. In fact, we find the following recursion: If $\partial(a_i) = \sum \kappa B_1 b B_2 b B_3 b B_4 b \dots B_k b A$, where $\kappa \in \mathbf{Z}[t, t^{-1}]$, $B_1, \dots, B_k \in T(b_1, \dots, b_m, a_1, \dots, a_{i-1})$, and A doesn't contain any copy of b which is not preceded by a copy of a, then

$$\begin{aligned}\varphi(a_i') = a_i + \sum \kappa\Big(&(-1)^{|B_1|}\varepsilon_a'\varepsilon_b'\bar{B}_1 a B_2 b B_3 b B_4 b \dots B_k b A\\ &+ (-1)^{|B_1|+|B_2|}(-1)^{|b|}\bar{B}_1 v \bar{B}_2 a B_3 b B_4 b \dots B_k b A\\ &+ (-1)^{|B_1|+|B_2|+|B_3|}\varepsilon_a'\varepsilon_b'\bar{B}_1 v \bar{B}_2 v \bar{B}_3 a B_4 b \dots B_k b A\\ &+ (-1)^{|B_1|+|B_2|+|B_3|+|B_4|}(-1)^{|b|}\bar{B}_1 v \bar{B}_2 v \bar{B}_3 v \bar{B}_4 a \dots B_k b A\\ &+ \dots\\ &+ (-1)^{|B_1|+\dots+|B_k|}(-1)^{|b|}\left(\varepsilon_a'\varepsilon_b'(-1)^{|b|}\right)^k \bar{B}_1 v \bar{B}_2 v \bar{B}_3 v \bar{B}_4 v \dots \bar{B}_k a A\Big),\end{aligned} \tag{10}$$

where $\bar{B}_1, \dots, \bar{B}_k$ are obtained from $B_1, \dots, B_k$ by replacing each of the symbols $a_1, \dots, a_{i-1}$ by the corresponding polynomial from the already constructed list $\varphi(a_1'), \dots, \varphi(a_{i-1}')$. Note that if $\partial(a_i)$ doesn't contain any of $a_1, \dots, a_{i-1}$, then (10) is analogous to (9).

The following is an easy consequence of (8) and either (7) or (10):

Proposition 3.5 *If $\varphi\colon \mathcal{A}' \to \mathcal{A}$ is the holonomy of a Reidemeister II move and $x' \in \mathcal{A}'$ is a generator so that for the corresponding $x \in \mathcal{A}$, we have $\partial(x) = 0$, then $\varphi(x') = x$.*

Example 3.6 Figure 7 shows the succession of four Reidemeister moves. Using Theorem 4.1 it will be very easy to check that all four of them are possible; see Example 4.2. Note that we dropped primed labels from our notation which we hope will not lead to confusion. Let us compute the images of the cycles b_1, b_2, and b_3 under the composition of the four holonomies. (We could apply the chain maps to a_1 and a_2 as well but because those are not cycles in the chain complex (see Example 2.14), the results in themselves are not very informative. Let us just mention as an illustration that under the first of the four moves, the image of a_2 is $a_2 + d + tb_2b_3d$.)

The first move is of type II. But because after the move, the diagram still only contains crossings with grading 0 or 1 ($|c_1| = 0$ and $|d| = 1$) and $|t| = 0$, the boundary of either one of the index 0 crossings b_1, b_2, and b_3 is still 0 after the move. Thus by Proposition 3.5, all of them are mapped trivially. The two type $\mathrm{III_b}$ moves that follow only affect d ($d \mapsto d + a_1c_1 \mapsto d - b_1a_2 + a_1c_1$) but this is not relevant for our purpose. Finally, the holonomy of the II^{-1} move takes b_1 to $-t^{-1} - b_3c_1$ (and d to 0).

Next, we prove that a non-trivial monodromy implies that the loop itself is non-trivial.

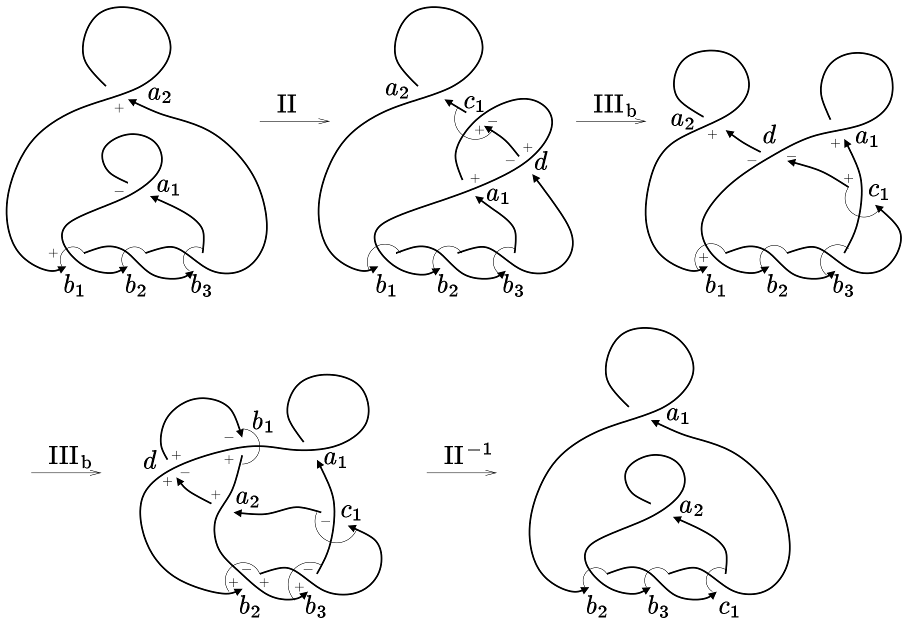

Figure 7: The loop $\Omega_{3,2}$ of trefoils, shown in the Lagrangian projection

3.2 Chain homotopies

It is a folk theorem of singularity theory (see for example [1, p. 18], where the more general situation of plane curves with semicubical cusp singularities is considered) that there are six different codimension 2 strata (bifurcations) of degenerate immersions of 1–manifolds into 2–manifolds. These are the three different possibilities of two distant Reidemeister moves happening at the same time, and three more types: quadruple points with pairwise transverse branches; degenerate triple points when exactly two branches are tangent; and cubic tangency of two branches. The neighborhoods (versal deformations) of points on each of these strata are also understood (see [1, Figure 6]).

For our purposes, each case has to be further separated according to the actual position of the branches in 3–space. Thus, we have six types of quadruple points corresponding to the six cyclic orderings of the branches. Degenerate triple points have three types distinguished by whether the third branch crosses over, under, or between the two whose projections are tangent. Cubic tangencies are separated into two obvious types. As to the cases of simultaneous Reidemeister moves, let us note that if two degenerate Reeb chords appear at the same time, their lengths, which we'll call the *height of the move*, can generically be assumed different.

Theorem 3.7 *Let $\mathcal{L} \subset \mathrm{Leg}(S^1, \mathbf{R}^3)$ be a Legendrian knot type. Let L_0 and L_1 be generic elements of $\mathcal{L}$ and $\Phi(t)$, $\Psi(t)$ two generic paths joining them, inducing the holonomies $\Phi_*, \Psi_* \colon H(L_0) \to H(L_1)$. If these paths are homotopic with their endpoints fixed, then $\Phi_* = \Psi_*$.*

Proof As $\Phi(t)$ and $\Psi(t)$ are homotopic, there is also a generic homotopy $\Phi_s(t)$ connecting them. This map meets the codimension 2 discriminant $\mathcal{D}_2$ transversely in finitely many points. If the number of such points is zero, then the only time the sequence of Reidemeister moves changes is those finitely many values of s when $\Phi_s(t)$ is tangent to $\mathcal{D}_1$ (for some value of t). At these instances, a pair of Reidemeister moves appears or disappears. Because these two moves are inverses of each other, their holonomies cancel each other out by what we have shown in the proof of Theorem 3.2. The rest of the proof consists of an analysis of the codimension 2 strata of the discriminant, as described in the introduction of the section.

Case 1 *Simultaneous Reidemeister moves where the two or three crossings affected by the moves form disjoint sets ("far away" moves)* Each such case fits the scheme (differentials noted for reference later)

$$\begin{tikzcd}
\text{before, before } (\partial) \arrow[r, "\varphi"] \arrow[d, "\psi"'] & \text{after, before } (\partial_2) \arrow[d, "\hat\psi"] \\
\text{before, after } (\partial_1) \arrow[r, "\hat\varphi"'] & \text{after, after } (\partial')
\end{tikzcd} \tag{11}$$

and we can arrange that if one or two pairs of Reidemeister II, II^{-1} moves are involved, then the arrows always point in the direction of the II^{-1} move. Then it suffices to check that in each case, the two compositions of chain maps from upper left to lower right agree. Therefore it will not be necessary to introduce a chain homotopy, or in other words, it can be chosen to be 0. This phenomenon will re-occur in each case except the very last one.

By our last assumption and the formulas of Theorem 3.2, the claim is clear for generators not directly affected by the moves. If any of the moves is a triangle move, it is straightforward to check for the vertices of the triangle, too. Suppose the horizontal arrows represent II^{-1} moves, with disappearing crossings a and b, so that $\partial(a) = \pm b + v$. The crossing a gets mapped to 0 under either composition. Let $\partial_1(a) = \pm b + w$. Then $\hat\varphi(\psi(b)) = \mp w$ and $\hat\psi(\varphi(b)) = \mp\hat\psi(v)$. But because $\psi(a) = a$ and hence $\psi(\partial(a)) = \partial_1(a)$, and

$\psi(b) = b$, we have $\psi(v) = w$. So to show that $\hat{\varphi}(\psi(b)) = \hat{\psi}(\varphi(b))$, we need to prove that[11] $\psi(v) = \hat{\psi}(v)$. This follows by a case-by-case analysis. If ψ and $\hat{\psi}$ are induced by triangle moves then they clearly act the same way on v. If it is a II^{-1} move with a bigger height than that of the one inducing φ, then $\psi(v) = \hat{\psi}(v) = v$. If the height is lower than that of the first move, then we have to make sure that the same expression computes the boundary of the higher-index vanishing crossing of the second move in each diagram of the upper row of (11), because this is used to define ψ and $\hat{\psi}$. But this is clear since φ is a chain map and the expressions in question can't contain neither a nor b.

Case 2 *Quadruple points* We are going to draw circles linking the singularity once and compute monodromies around these loops, starting and ending at the top left in each of the cases depicted in Figures 9 and 10. Such a loop meets the triple point discriminant eight times and in each case we will find that the composition of the corresponding eight chain maps is the identity.

Out of the six cases, the one on the bottom of Figure 9 is the most complicated, due to the fact that all eight moves are of type $\mathrm{III_b}$. Even in this case, only the crossings b, c, and e need checking, because others are always mapped trivially. On the left side of Figure 9, we indicated the orientation signs used in the following. In the second move, b's image is $b - \varepsilon'_b\varepsilon_d\varepsilon'_a t^{k(b:d,a)}ad$. The three generators in this expression are mapped trivially until the sixth move, when the image becomes $(b - \varepsilon_b\varepsilon'_d\varepsilon_a t^{k(b:d,a)}ad) - \varepsilon'_b\varepsilon_d\varepsilon'_a t^{k(b:d,a)}ad = b$ (this is by the same argument applied in the proof of Theorem 3.2). The generator e returns to itself through two completely analogous moves. Finally, we have the following for c:

$$
\begin{aligned}
&c \mapsto c - \varepsilon_c\varepsilon_d\varepsilon_e t^{k(c;d,e)}ed \mapsto \text{same} \mapsto c - \varepsilon_c\varepsilon_d\varepsilon_e t^{k(c;d,e)}(e - \varepsilon'_e\varepsilon_a\varepsilon_f t^{k(e;a,f)}fa)d\\
&\mapsto c - \varepsilon'_c\varepsilon_b\varepsilon_f t^{k(c;b,f)}fb - \varepsilon_c\varepsilon_d\varepsilon_e t^{k(c;d,e)}(e - \varepsilon'_e\varepsilon_a\varepsilon_f t^{k(e;a,f)}fa)d\\
&\mapsto c - \varepsilon'_c\varepsilon'_d\varepsilon'_e t^{k(c;d,e)}ed - \varepsilon'_c\varepsilon_b\varepsilon_f t^{k(c;b,f)}fb - \varepsilon_c\varepsilon_d\varepsilon_e t^{k(c;d,e)}(e - \varepsilon'_e\varepsilon_a\varepsilon_f t^{k(e;a,f)}fa)d\\
&= c - \varepsilon'_c\varepsilon_b\varepsilon_f t^{k(c;b,f)}fb - (-1)^{|e|}\varepsilon_c\varepsilon_d\varepsilon_a\varepsilon_f t^{k(c;d,e)+k(e;a,f)}fad\\
&\mapsto c - \varepsilon'_c\varepsilon_b\varepsilon_f t^{k(c;b,f)}f(b - \varepsilon_b\varepsilon'_d\varepsilon_a t^{k(b;d,a)}ad) - (-1)^{|e|}\varepsilon_c\varepsilon_d\varepsilon_a\varepsilon_f t^{k(c;d,e)+k(e;a,f)}fad\\
&= c - \varepsilon'_c\varepsilon_b\varepsilon_f t^{k(c;b,f)}fb \mapsto \text{same} \mapsto c - \varepsilon_c\varepsilon'_b\varepsilon'_f t^{k(c;b,f)}fb - \varepsilon'_c\varepsilon_b\varepsilon_f t^{k(c;b,f)}fb = c.
\end{aligned}
$$

[11]Note that, with a slight abuse of notation, we sometimes suppress the chain map relating DGA's of different diagrams when it simply replaces all generators in one expression by corresponding generators from the other diagram. Our "excuse" is that these new generators are denoted by the same symbols as the old ones. A more exact version of the claim would be that $\psi(v)$ and $\hat{\psi}(v)$ are expressed by identical polynomials in the two DGA's.

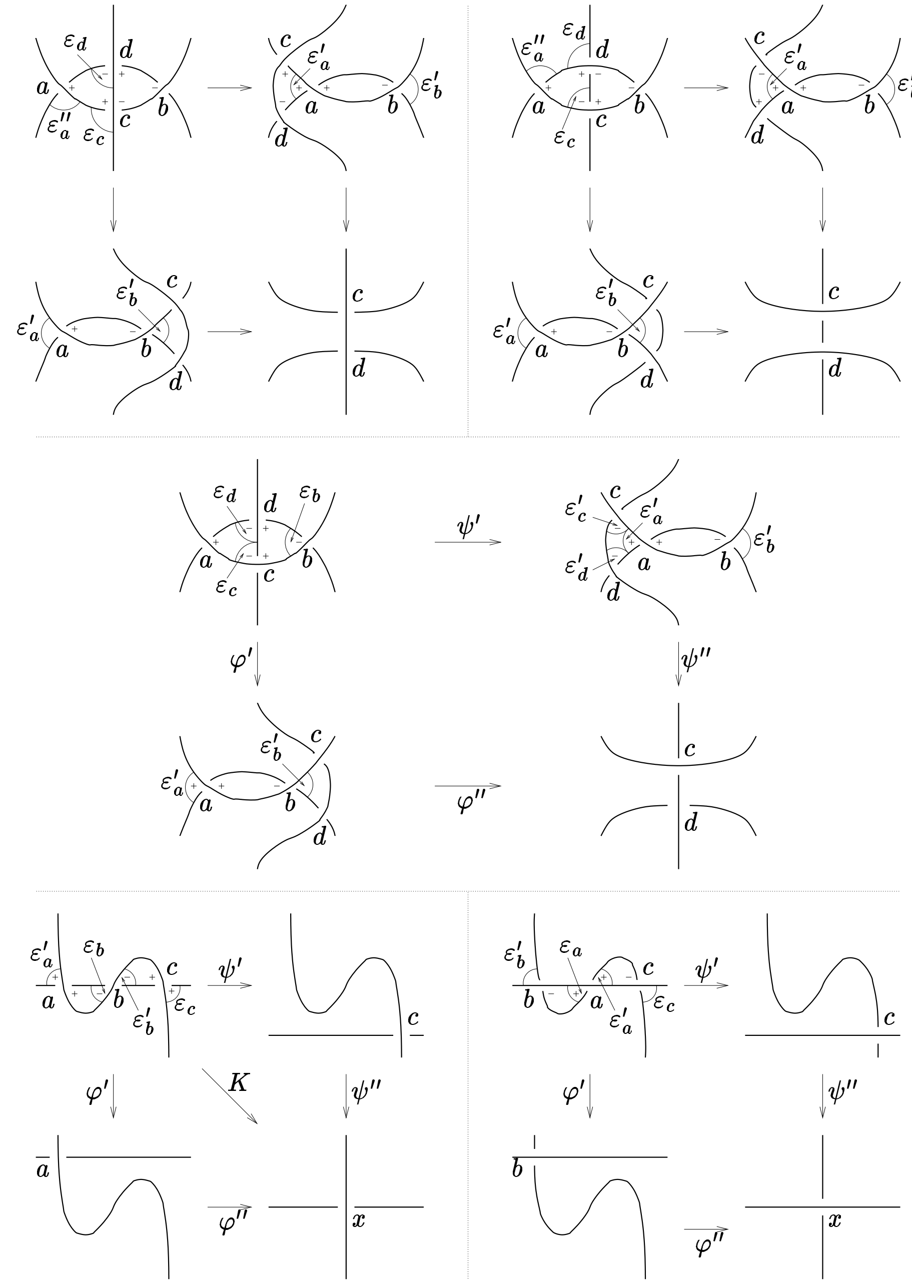

Figure 8: Degenerate triple points and cubic self-tangencies

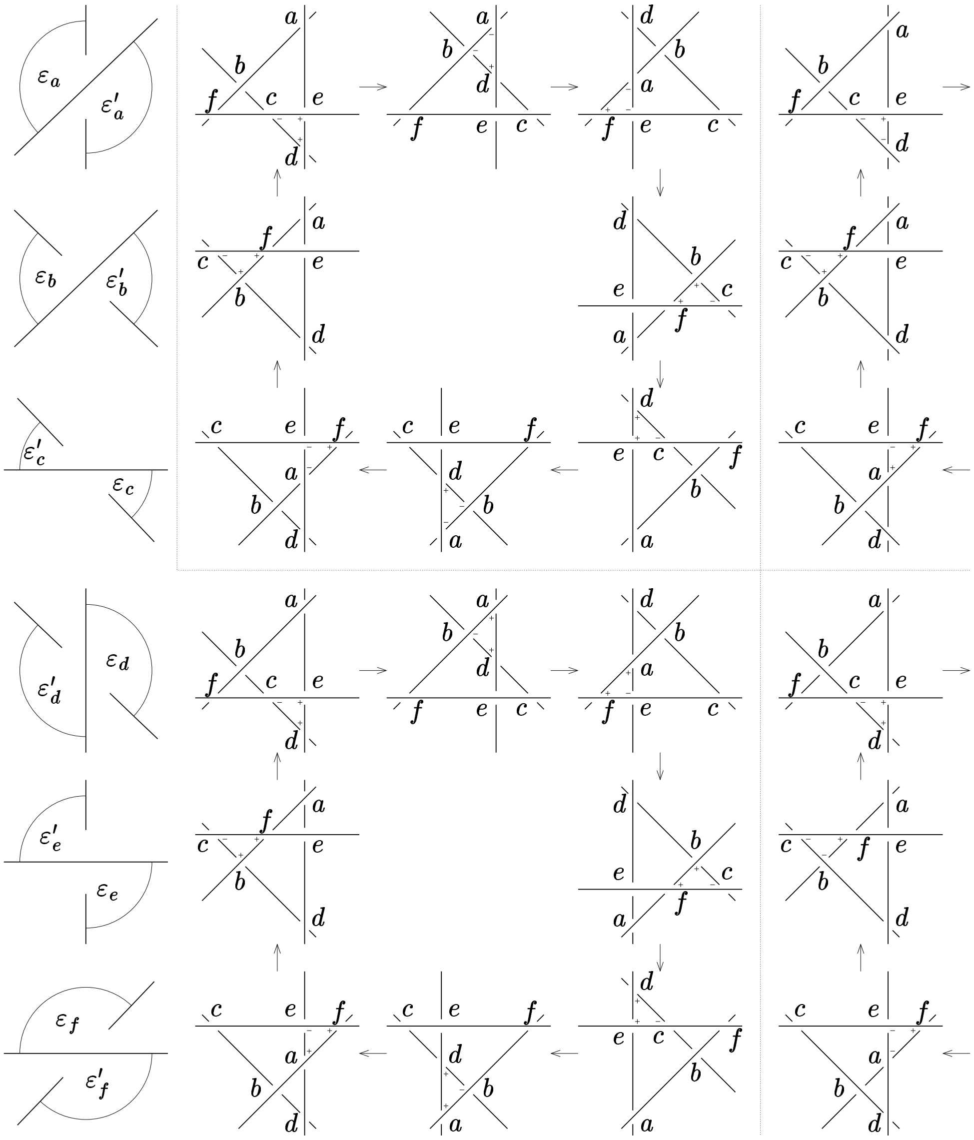

Figure 9: Quadruple points I

The equality stated in the fourth row, as well as the very last one, holds by the argument in the proof of Theorem 3.2. The other one in the sixth row is true because

- $k(c;b,f)+k(b;d,a)$ and $k(c;d,e)+k(e;a,f)$ both equal the total winding of $-\gamma_c+\gamma_d+\gamma_a+\gamma_f$, with each capping path thought of as starting and

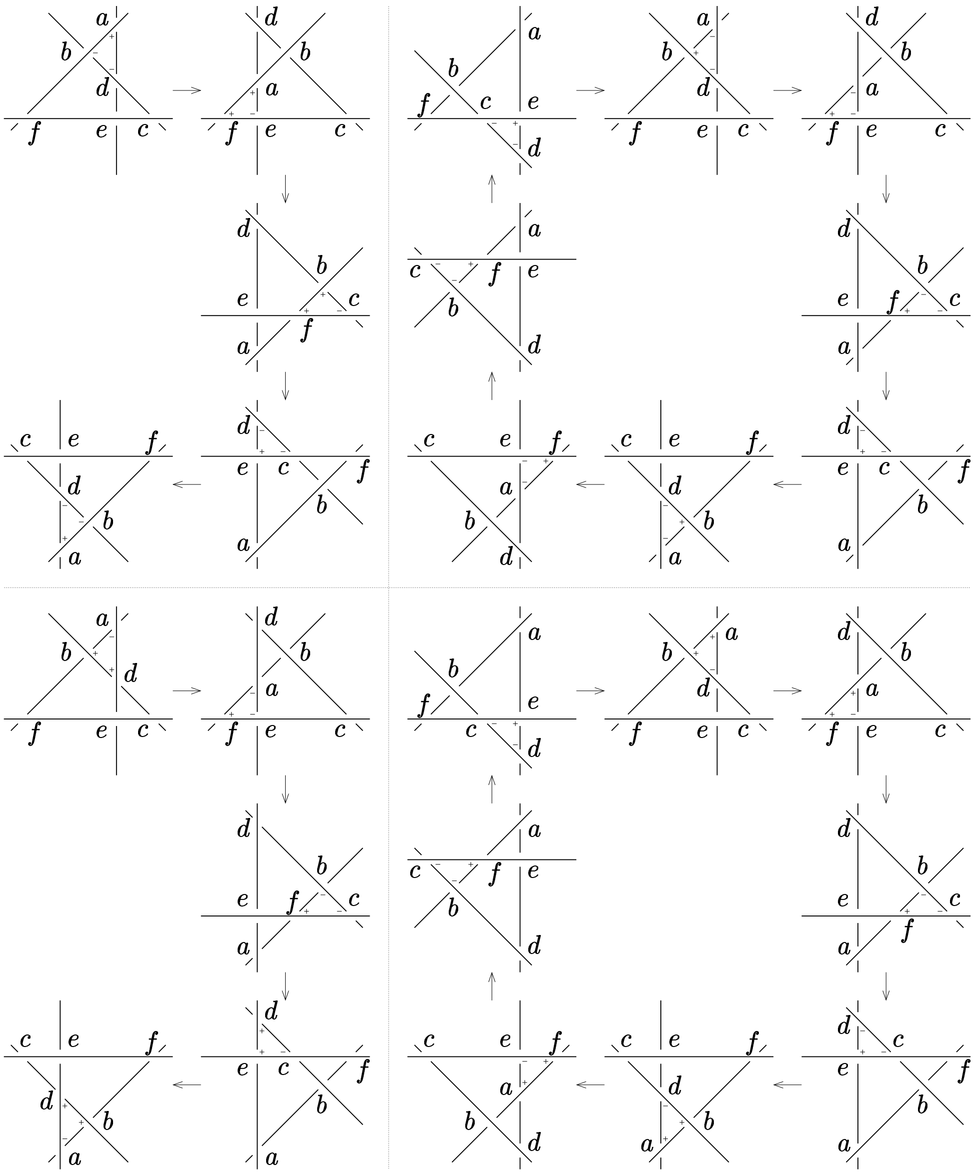

Figure 10: Quadruple points II

ending at the quadruple point.

- $\varepsilon'_c\varepsilon_f\varepsilon'_d\varepsilon_a(-1)^{|e|}\varepsilon_c\varepsilon_d\varepsilon_a\varepsilon_f = (-1)^{|e|}(-1)^{|c|}(-1)^{|d|} = 1$ by Lemma 2.3.

The other five quadruple point bifurcations can be handled with very similar, and simpler, computations.

Case 3 *Degenerate triple points* For the three bifurcations on the top of Figure 8, we claim that the chain maps corresponding to the two ways of moving from upper left to lower right in the diagrams induce the same chain maps. Please refer to Figure 8 for the positions of the orientation signs that appear in the calculations and note that ε and ε' always denote orientation signs of opposite quadrants. Consider the third case (middle of Figure 8). Denote the differential of the upper left diagram by ∂, of the lower left diagram by ∂_1 and that of the upper right diagram by ∂_2. Applying the chain maps of Theorem 3.2 to the generator b, we find the following:

$$b \xmapsto{\varphi'} b - \varepsilon_b\varepsilon_d\varepsilon_c t^{k(b;c,d)}cd \xmapsto{\varphi''} \varepsilon'_b\varepsilon'_a v - \varepsilon_b\varepsilon_d\varepsilon_c t^{k(b;c,d)}cd$$

and $$b \xmapsto{\psi'} b \xmapsto{\psi''} \varepsilon'_b\varepsilon'_a(\varepsilon'_a\varepsilon'_d\varepsilon'_c t^{k(a;d,c)}cd + w).$$

Here, v and w are defined by the equations $\partial_1(a) = \varepsilon_a\varepsilon_b b + v$ and $\partial_2(a) = \varepsilon_a\varepsilon_b b + \varepsilon'_a\varepsilon'_d\varepsilon'_c t^{k(a;d,c)}cd + w$. The coefficients of cd agree in the two images of b by Lemma 2.3 and the observation that the capping paths γ_a and γ_b used to compute the exponents of t coincide at the moment of self-tangency. Thus, we need to prove that $v = w$. For this, note that $(\varphi')^{-1}(a) = a$ and $\psi'(a) = a$, therefore $\partial(a) = (\varphi')^{-1}(\partial_1(a)) = \varepsilon_a\varepsilon_b(b - \varepsilon'_b\varepsilon'_d\varepsilon'_c t^{k(b;d,c)}cd) + v$, which is also the expression of $\partial_2(a) = \psi'(\partial(a))$. Now, $v = w$ follows because $(-\varepsilon_a\varepsilon_b\varepsilon'_b\varepsilon'_d\varepsilon'_c)(\varepsilon'_a\varepsilon'_d\varepsilon'_c) = -(-1)^{|a|+|b|} = 1$. The claim is clear for other generators of the upper left diagram.

In the other two cases (top row of Figure 8), the proof is very similar. It is important to note though that in the first diagram of the top row, $h(d) > h(a)$, and in the second, $h(c) > h(a)$ in a neighborhood of the codimension 2 discriminant. Due to this fact, the boundary of a is given by the same polynomial in both pairs of upper right and lower left diagrams.

Case 4 *Cubic self-tangencies* Note that on the two remaining diagrams (bottom row of Figure 8), we only have two Reidemeister moves; the chain maps φ'' and ψ'' are simple re-labelings. In the case depicted on the second diagram, we still don't need to introduce a non-zero chain homotopy. This is because $\varphi''(\varphi'(a)) = \psi''(\psi'(a)) = 0$, $\psi''(\psi'(b)) = \psi''(\varepsilon'_a\varepsilon'_b(\varepsilon'_a\varepsilon'_c c)) = x = \varphi''(\varphi'(b))$ (it is an easy fact that $\varepsilon'_b = \varepsilon'_c$), and a similar computation for c.

Finally, it is the situation on the bottom left of Figure 8 which does require a non-trivial chain homotopy K, mapping between the two indicated DGA's. Let $K(b) = \varepsilon_a\varepsilon_b x$ and let K map all other generators to 0. We extend K as in Lemma 2.18. Let $\varphi = \varphi'' \circ \varphi'$ and $\psi = \psi'' \circ \psi'$. Then $\varphi(a) - \psi(a) = x$, $\varphi(c) - \psi(c) = -x$ and $\varphi(b) - \psi(b) = \varepsilon_b\varepsilon_c v_c - \varepsilon'_a\varepsilon'_b v_a$, where the polynomials v_a

and v_c are defined in the upper left diagram by the equations $\partial(a) = \varepsilon_a\varepsilon_b b + v_a$ and $\partial(c) = \varepsilon'_c\varepsilon'_b b + v_c$. Finally, it is easy to check the following, too:

$$(K \circ \partial + \partial' \circ K)(a) = K(\varepsilon_a\varepsilon_b b + v_a) = x$$

(because v_a doesn't contain b);

$$(K \circ \partial + \partial' \circ K)(c) = K(\varepsilon'_c\varepsilon'_b b + v_c) = \varepsilon'_c\varepsilon'_b\varepsilon_a\varepsilon_b x = (-1)^{|a|+|b|}x = -x$$

(because $\varepsilon'_c = \varepsilon'_a$), and

$$(K \circ \partial + \partial' \circ K)(b) = \partial'(\varepsilon_a\varepsilon_b x) = \varepsilon_a\varepsilon_b v_a + \varepsilon_a\varepsilon_b v_c,$$

where, indeed, $\varepsilon_a = \varepsilon_c$ and $\varepsilon_a\varepsilon_b = -\varepsilon'_a\varepsilon'_b$. □

3.3 $\mathbf{Z}_2$–coefficients

In this subsection we briefly summarize the modifications needed to reduce our discussion to the original $\mathbf{Z}_2$–coefficient theory of Chekanov. (We will work in that context in sections 6 through 10.) Please note again that this means that we are going to describe the same geometry (Legendrian knots, their families, Reeb chords, and holomorphic discs) using slightly less sophisticated algebra: The generators are still the same Reeb chords, ie crossings, and the product is non-commutative, but we substitute $t = 1$, reduce integers modulo 2, and reduce the grading modulo $2r$, where $2r$ is the common Maslov number of each knot in the family. The only simplification we get in Table 1 is that now, $(-1)^{|a|} = 1$ (and $\varphi - \psi = \varphi + \psi$). Otherwise, we keep all notation introduced in this section and still refer to Figure 6 (orientation signs can now be ignored).

A list of maps that, when extended from the generators to the DGA as algebra morphisms, become the chain maps that are used to define holonomies and monodromies of (sequences of) Reidemeister moves, is as follows.

Move $\mathrm{III_a}$: Let $a \mapsto a'$, $b \mapsto b'$, $c \mapsto c'$, and $x \mapsto x'$, where x is any other crossing of the upper diagram.

Move $\mathrm{III_b}$: Let

$$a \mapsto a' + c'b',$$

while other generators are mapped trivially: $b \mapsto b'$, $c \mapsto c'$, and $x \mapsto x'$.

Move II^{-1}: Let $\partial(a) = b + v$. Define $x \mapsto x'$, which gives rise to the obvious re-labeling $v \mapsto v'$. Then let

$$\begin{aligned} a &\mapsto 0 \\ b &\mapsto v'. \end{aligned}$$

Move II: The map $\varphi\colon \mathcal{A}' \to \mathcal{A}$ still takes

$$b_j' \mapsto \varphi(b_j') = b_j$$

for all $j = 1, \ldots, m$. Formula (9) of Remark 3.4 becomes

$$\begin{aligned} a_1' \mapsto \varphi(a_1') = a_1 + \sum \Big(& B_1 a B_2 b B_3 b B_4 b \ldots B_k b A \\ & + B_1 v B_2 a B_3 b B_4 b \ldots B_k b A \\ & + B_1 v B_2 v B_3 a B_4 b \ldots B_k b A \\ & + B_1 v B_2 v B_3 v B_4 a \ldots B_k b A \\ & + \ldots \\ & + B_1 v B_2 v B_3 v B_4 v \ldots B_k a A \Big), \end{aligned}$$

and formula (10) takes the form

$$\begin{aligned} a_i' \mapsto \varphi(a_i') = a_i + \sum \Big(& \bar{B}_1 a B_2 b B_3 b B_4 b \ldots B_k b A \\ & + \bar{B}_1 v \bar{B}_2 a B_3 b B_4 b \ldots B_k b A \\ & + \bar{B}_1 v \bar{B}_2 v \bar{B}_3 a B_4 b \ldots B_k b A \\ & + \bar{B}_1 v \bar{B}_2 v \bar{B}_3 v \bar{B}_4 a \ldots B_k b A \\ & + \ldots \\ & + \bar{B}_1 v \bar{B}_2 v \bar{B}_3 v \bar{B}_4 v \ldots \bar{B}_k a A \Big). \end{aligned}$$

In the rest of the paper, each time we compute the holonomy of a Reidemeister II move, Proposition 3.5 applies. In other words, the detailed description of move II was included only for completeness.

4 Legendrian Reidemeister moves

In this section, we investigate a very basic question, which arises naturally in the theory of Legendrian knots and links. The question is this: if a Lagrangian diagram of a Legendrian link contains a part that is *topologically* fit for a Reidemeister II, II^{-1}, or III move, can we carry the move out in the Legendrian category? (If so, we will call the move *consistent.*) As it turns out, the answer is not always affirmative and a necessary and sufficient condition can be formulated in terms of a linear program.

We call two regions (of the complement of a knot diagram) *adjacent at a vertex* if they share the vertex (crossing) but don't share any edges.

Theorem 4.1 *Each Reidemeister move of Figure 6 is possible if and only if the Lagrangian diagram can be isotoped (in the Legendrian sense) so that a certain inequality is satisfied. The list is as follows:*

- *Move III_a: the area of the triangle is smaller than the area of any of the three regions adjacent to it at either of its vertices.*
- *Move III_b: the area of the triangle is smaller than the area of the region adjacent to it at the vertex a.*
- *Move II^{-1}: the area of the 2–gon is smaller than the sum of the areas of the two regions adjacent to it at its vertices.*
- *Move II: the sum of the heights of the crossings along the right side of the pinching region, counted with Reeb signs, is positive. (If the region to the right of the pinching, as it is depicted in Figure 6, is the unbounded one, then there is no obstruction to the move.)*

Recall from section 2 that every Lagrangian diagram γ has the non-empty cone $\mathcal{C}_\gamma$ associated to it. $\mathcal{C}_\gamma$ is defined by the positivity constraints on the heights associated with γ and finitely many homogeneous linear inequalities (coming from bounded regions of the complement) in terms of these heights (see Remark 2.5). Each inequality of Theorem 4.1 requires that a certain linear objective function[12] defined on $\mathcal{C}_\gamma$ take on a positive value. That is, in each case the necessary and sufficient condition can be re-phrased to require that a certain linear programming problem be unbounded. This can be checked by the simplex method or any other linear programming algorithm (for small diagrams even by hand, as the author can attest).

Proof The "only if" part of each statement is easy to justify. To see this for move II, use Proposition 2.2: at the moment of self-tangency of the projection, the sum in question is the area of the region pinching off on the right side plus the difference in z–coordinate between the two preimages of the self-tangency. As this is positive, our sum must have been positive for some period of time before the move happened, too. The other three moves involve the vanishing of the area of a region; clearly, for some time before that happens, the particular area is smallest among all others and the said inequalities hold.

Now to prove the "if" statement for move II, suppose that the sum S described in the statement is positive and extend a very narrow 'finger' that follows closely the right side of the region from the lower edge to the upper one. The area of

[12]The objective function is linear exactly because the areas of the regions are linear functions of the heights. These areas can also be thought of as slack variables.

the finger can be made arbitrarily small and therefore with a slight downward bump on the lower edge (see Remark 2.8) we can make sure that other parts of the knot (in particular the value of S) are not affected. As we are also able to arrange that the area that is pinching off be less than S, an application of Proposition 2.2 shows that at the moment the finger reaches the upper edge, it must do so so that it crosses underneath it. If the 'right side' that we used doesn't exist due to the unboundedness of the pinching region, then simply let the finger follow a path that encircles a large enough area (on the finger's left) so that it dips down sufficiently.

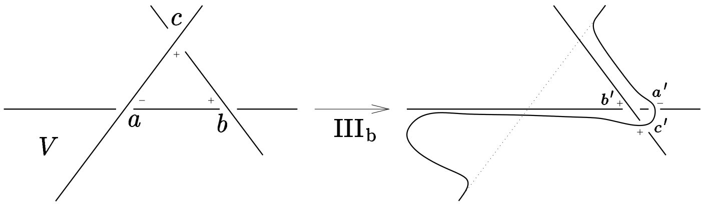

Figure 11: Reidemeister $\mathrm{III_b}$ move

In the other three cases, we again utilize the idea described in Remark 2.8. To carry out move $\mathrm{III_b}$, replace the uppermost strand as indicated in Figure 11. It is clear from the proof of Proposition 2.2 that if the new arc, starting from the old crossing c, follows the side cb close enough, then the new crossing c' has height arbitrarily close to $h(c)$, and $h(a')$ is arbitrarily close to $h(c)+h(b)$. In particular, the types of the crossings are as indicated. Then, if we followed the sides cb and ba close enough, we can use the excess area of the region V to compensate for the modification.

Three very similar arguments establish the claim for move $\mathrm{III_a}$, and Figure 12 is self-explanatory for the case of move II^{-1}. □

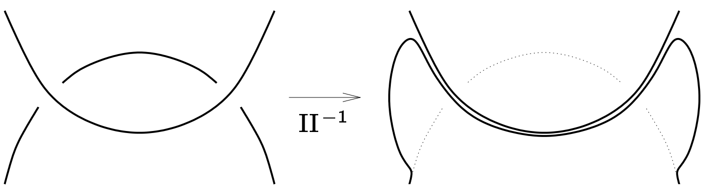

Figure 12: Reidemeister II^{-1} move

Example 4.2 We'll prove that the loop of trefoils in Example 3.6 is sound, ie, all four of the indicated Reidemeister moves can indeed be carried out.

Let us consider the first diagram in Figure 7. By adding "bulges" of equal area to the regions that were denoted by U_1 and U_2 in Figure 4, we may arrange that the areas of those, along with the height $h(a_2)$, be arbitrarily larger than any other area or height occurring in the projection. In particular, $h(a_2)+h(b_1)-h(a_1)$ can be made positive, which by Theorem 4.1 means that the first move is possible. The triangle move that follows is consistent because the region adjacent to the triangle at a' is the unbounded one. At this point, we may finish the argument by pointing out that the remaining two moves are mirror images of the first two, or we may continue as follows: For the other triangle move, enlarge the areas of the triangle c_1a_1a' and of the region capping off a_1 until the area of the triangle gets bigger than that of the triangle a_2b_1a' (this is meant in the third diagram of Figure 7). Finally, the same trick applied to the bounded regions adjacent to a_1 in the fourth diagram makes sure that the appropriate condition from Theorem 4.1 is satisfied and the last move of the loop is consistent, too.

5 First evidence of non-triviality

At this point we are ready to prove that the monodromy invariant of Theorem 1.1 is non-trivial. We will use essentially the same example to establish Corollary 1.4, too.

Note that the first and last diagrams of Figure 7 coincide, so in Examples 3.6 and 4.2 we investigated a loop of Legendrian trefoil knots. It is not hard to identify it with $\Omega_{3,2}$ (defined in the introduction). If we re-label the fifth diagram like the first one, we can write that the monodromy of the loop (which is the composition of the four holonomies of Example 3.6 and the re-labeling) acts on the cycles b_1, b_2, and b_3 as follows:

$$\begin{aligned} b_1 &\mapsto -t^{-1}-b_2b_3 \\ b_2 &\mapsto b_1 \\ b_3 &\mapsto b_2. \end{aligned}$$

Proposition 5.1 *The map μ_0 defined by the formulas above on the index 0 part of the contact homology (of the knot introduced in Example 2.7) has order 5. Thus Theorem 1.3 holds in the case $p=3$, $q=2$.*

Proof As $\mu_0(b_3)=b_2$ and $\mu_0^2(b_3)=b_1$, all three generators are on the same orbit. Then we compute that $\mu_0^4(b_3)=\mu_0(-t^{-1}-b_2b_3)=-t^{-1}-b_1b_2$ and that $\mu_0^5(b_3)=\mu_0(-t^{-1}-b_1b_2)=-t^{-1}-(-t^{-1}-b_2b_3)b_1=-t^{-1}+t^{-1}b_1+b_2b_3b_1=b_3$

by equation (5) of Example 2.14.

The argument for the order will be complete if we show that $b_3 \neq b_2$ in the homology, ie, that there isn't an element $a \in \mathcal{A}$ so that $\partial(a) = b_3 - b_2$. For this, notice that $\operatorname{im}\partial$ is contained in the two-sided ideal generated by $\partial(a_1)$ and $\partial(a_2)$. An examination of equations (3) and (4) shows that after substituting $t = -1$, they both contain an even number of terms that are powers of b_3 (namely, two and zero terms). Therefore the same property holds for the entire ideal. Since $b_3 - b_2$ is not such a polynomial, it cannot be in the image of the differential. $\square$

It is easy to see that if we fix the basepoint, but allow knots in our one parameter families that are not Legendrian, then the three-fold concatenation of $\Omega_{3,2}$ is homotopic to a continuous rotation by 2π radians. Therefore, as $\pi_1(SO(3)) = \mathbf{Z}_2$ is generated by exactly this loop, the six-fold concatenation of $\Omega_{3,2}$ is contractible in the space $\mathcal{K}$ of smooth right-handed trefoil knots. But because $\mu((\Omega_{3,2})^6) \neq \mathrm{id}$ (as $\mu_0((\Omega_{3,2})^6) = \mu_0(\Omega_{3,2}) \neq \mathrm{id}$), $(\Omega_{3,2})^6$ is non-contractible in the space $\mathcal{L}$ of Legendrian trefoil knots. This last implication, which establishes Corollary 1.4, holds by Theorem 3.7. (We remark that with a bit of extra work, it can be shown that $\Omega_{3,2}$ is itself contractible in $\mathcal{K}$, so taking its sixth power isn't really necessary in our argument.)

In sections 9 and 10, we will generalize the proof of Proposition 5.1 for an arbitrary (p,q) torus knot. (Except that to simplify the discussion, in those sections we will work with $\mathbf{Z}_2$–coefficients.) Before we do that, we use sections 6 and 8 to set up a more general picture. Section 7 contains a vital ingredient of the proof of Theorem 1.3.

6 Positive braid closures

Let us consider a positive braid β on q strands, as in Figure 15. Label the left and right endpoints of the strands from top to bottom with the first q whole numbers. The pair of left and right labels on each strand takes the form $(i, \sigma(i))$, where σ is the *underlying permutation* of β. Further, label the crossings of the braid with a pair of numbers, the first one the left label i of the overcrossing strand and the second one the right label j of the undercrossing one. (Note that not all pairs of numbers between 1 and q occur as labels of crossings; for example, $(i, \sigma(i))$ never does for any i.)

Definition 6.1 A positive braid β defines the front diagram of an oriented Legendrian link as on the upper half of Figure 13 (orient each strand of β from left to right). We call this construction the *Legendrian closure* of β and denote it by L_β.

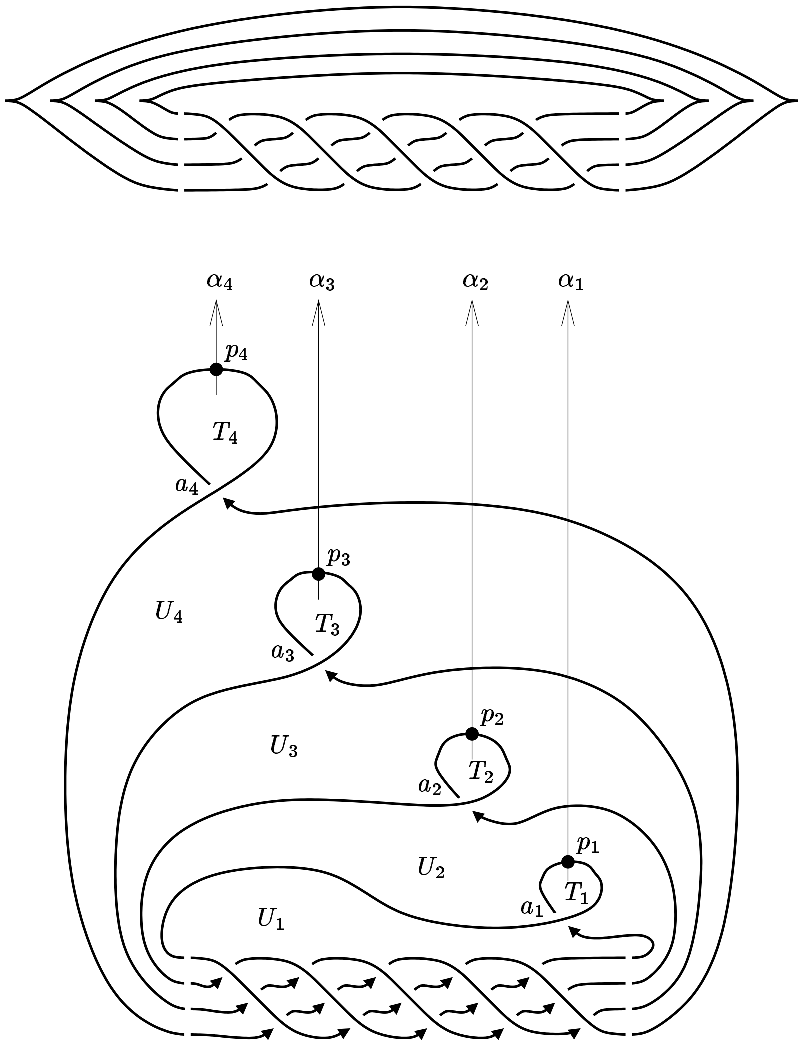

Figure 13: Front and Lagrangian diagrams of the closure of a positive braid

Note that the types of crossings are correctly determined by the slopes of the branches that meet there exactly because the braid is positive. Applying Ng's [23] construction of resolution to the front diagram, we obtain the Lagrangian diagram γ_β in the lower half of Figure 13.

In section 2, we omitted the modifications needed to define the invariants (such as rotation and Thurston–Bennequin numbers and contact homology itself) for oriented multi-component Legendrian links. We include an informal rundown here (see [23, section 2.5] for more). The rotation number is the sum of the rotations of the components. The Thurston–Bennequin number is still the writhe of the Lagrangian projection; note that now, as opposed to the case of a single component, the orientation matters in its definition.

From now on, we will always use $\mathbf{Z}_2$–coefficients in contact homology. For a multi-component link, there is not a single distinguished grading of its DGA, rather a family of so called admissible gradings. We consider those introduced in [23] and not the larger class of gradings described in [6, section 9.1]. In the $\mathbf{Z}_2$–coefficient theory, each of these is defined modulo the greatest common divisor of the Maslov numbers of the components (recall that the Maslov number is twice the rotation number of a knot). Self-crossings of individual components have the same well-defined index in any admissible grading. We will refer to such generators as *proper crossings*. Other crossings' indices have the same parity in each admissible grading. The parity of the index of any crossing coincides with its sign in classical knot theory.

The differential ∂ is of index -1 with respect to each admissible grading. (Furthermore, Lemma 2.3 holds for each admissible grading, too.) Contact homology is a well-defined invariant in the sense that if two diagrams are Legendrian isotopic, then there is a one-to-one correspondence between their sets of admissible gradings so that the corresponding contact homologies are isomorphic as graded algebras.

It is easy to calculate that the Legendrian closure of any positive braid has rotation number $r=0$ (in fact, every component has Maslov number 0). The Thurston–Bennequin number is $tb(L_\beta) = (\text{word length of } \beta) - q$. The positive Legendrian (p,q) torus knot obtained as a special case is the one with maximal Thurston–Bennequin number $p(q-1)-q$ (see [11] for a classification of Legendrian torus knots). The $(3,2)$ torus knot that we considered earlier arises from this construction, too. The unique Legendrian unknot with maximal Thurston–Bennequin number $tb=-1$ [3, 8] is the Legendrian closure of the trivial braid on a single strand.

In the Chekanov–Eliashberg DGA of γ_β, the grading (any admissible grading) is integer-valued. The crossings a_m $(m = 1, \dots, q)$ have index 1. The rest of the generators are the crossings of β and they have index 0. In the multi-component case we should say instead that the grading that assigns the index

0 to each is admissible; from now on, we will always work with this grading. These generators will be labeled by $b_{i,j,t}$, where the integers $1 \le i, j \le q$ are the ones defined at the beginning of the section. The third label t is used to distinguish between multiple intersections of strands, enumerating them from left to right, as shown in Figure 14.

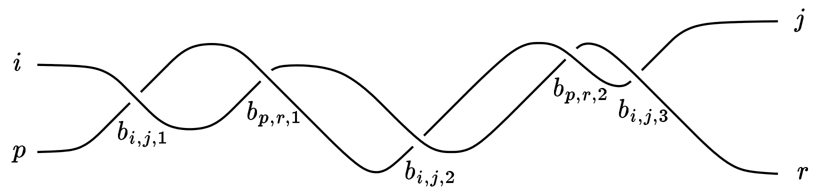

Figure 14: Labeling the crossings of a braid

Since ∂ lowers the index by 1,

$$\partial(b_{i,j,t}) = 0 \text{ for all } i, j = 1, \dots, q \text{ and } t.$$

In fact, there are no admissible discs whose positive corner is an index 0 crossing. The boundaries $\partial(a_m)$, $m = 1, \dots, q$, are polynomials in the non-commuting variables $b_{i,j,t}$ only. Our goal for the remainder of this section is to compute these polynomials.

The first observation is that for all $m = 1, \dots, q$, there is an admissible disc covering the teardrop-shaped region T_m right above a_m once. We'll call it the m^{th} *trivial disc.* It contributes 1 to $\partial(a_m)$. The m^{th} trivial disc is the only one that turns at the upward-facing positive quadrant at a_m. The rest of the contributions to $\partial(a_m)$ come from discs that turn at the positive quadrant facing down.

Suppose $f\colon \Pi_k \to \mathbf{R}^2_{xy}$ is an admissible immersion with respect to the projection $\gamma = \gamma_\beta$, with positive corner $f(x_0^k) = a_m$ and so that it is different from the m^{th} trivial disc. Fix points $p_1, \dots, p_q$ on γ as shown in Figure 13.

Lemma 6.2 *The curve $f(\partial\Pi_k)$ doesn't pass through the points $p_m, \dots, p_q$.*

Proof In fact, if we denote the region of the complement directly under a_i by U_i, then $f(\Pi_k)$ is disjoint from $U_{m+1}, \dots, U_q$, as well as from $T_m, \dots, T_q$. To see this, first shrink γ so that all areas and heights are smaller than ε^{q-m}, where $\varepsilon > 0$ is to be chosen later. Then add "bulges" of equal area 1 to T_q and U_q on the outside of the diagram (see Remark 2.8). Next, add bulges of area ε to T_{q-1} and U_{q-1} on their sides facing U_q. Continue all the way until adding bulges of area ε^{q-m-1} to T_{m+1} and U_{m+1} on their sides facing U_{m+2}. The result is a Lagrangian projection of the same link in which, after the choice of a small enough ε, the claim is obvious by Lemma 2.11, except for the case of T_m. But that follows from the easy observation that if any admissible disc passes through p_m so that locally, $f(\Pi_k)$ faces downward, then f is the m^{th} trivial disc. □

Recall that γ is oriented as shown in Figure 13. We have already noted that apart from its positive corner which is of index 1, f turns only at certain index 0 crossings $b_{i,j,t}$. By Lemma 2.10, this implies that f is compatible with the orientation of γ. Therefore, as the orientation of $f(\partial\Pi_k)$ agrees with that of γ near the positive corner a_m, these orientations agree at all points of $\partial\Pi_k$. This last observation for example implies that f can't turn at the downward facing negative quadrant at any of the $b_{i,j,t}$'s.

We may summarize our findings about f so far as follows. The curve $f(\partial\Pi_k)$ starts at a_m, follows γ until it reaches the braid at the left endpoint labeled m (while f extends this map toward U_m). Then it travels through the braid, possibly turning (left) at several crossings $b_{i,j,t}$ but always heading to the right, until it reaches a right endpoint labeled $i_1 \le m$. Then $f(\partial\Pi_k)$ climbs to a_{i_1} and, unless $i_1 = m$, beyond a_{i_1} (note that f can't have a positive corner at a_{i_1}) to p_{i_1}, at which point the extension f is toward U_{i_1+1}. Then $f(\partial\Pi_k)$ descends back to a_{i_1} so that it doesn't turn at the negative corner. So the process repeats with (m, i_1) replaced with (i_1, i_2) (where $i_2 \le m$), and so on until $i_{c+1} = m$ for some c.

Definition 6.3 A finite sequence of positive integers is called *admissible* if for all $s \ge 0$, between any two appearances of s in the sequence there is a number greater than s which appears between them. For $n \ge 1$, let us denote by D_n the set of all admissible sequences that are composed of the numbers $1, 2, \dots, n-1$.

For example, $D_1 = \{\varnothing\}$, $D_2 = \{\varnothing, \{1\}\}$, and

$$D_3 = \{\varnothing, \{1\}, \{2\}, \{1,2\}, \{2,1\}, \{1,2,1\}\}.$$

By induction on n and observing the position of the unique maximal term in the sequence, it is easy to prove that $|D_n| = |D_{n-1}|^2 + |D_{n-1}|$.

Proposition 6.4 *For the admissible disc f as above, the sequence of intermediate labels $\{i_1, i_2, \dots, i_c\}$ is an element of D_m.*

Proof There is an intuitive reason for the claim: if the index i was repeated without the boundary of the disc climbing higher between the two occurrences, then the disk would pinch off at a_i. We shall give a more rigorous proof using Blank's theorem (see [24], or the review in [17], from where we'll borrow our terminology).

Let $\tilde{f}$ be a small generic perturbation of f. This has the corners rounded so that $\tilde{f}|_{\partial\Pi_k}$ is an immersion; moreover, let us note that because $f(\partial\Pi_k)$ may cover parts of γ more than once, the complement of $\tilde{f}(\partial\Pi_k)$ typically contains more regions than that of $f(\partial\Pi_k)$. We are going to apply Blank's theorem to $\tilde{f}|_{\partial\Pi_k}$. We extend rays from each bounded component of $\mathbf{R}^2 \setminus \tilde{f}(\partial\Pi_k)$ to infinity. If the region is not one of those obtained from the T_i (let us denote these by T_i'), or one of the small ones resulting from perturbation close to the arc bounding T_i, this can be done so that γ only intersects the ray in a positive manner (ie, if the ray is oriented toward ∞, then γ crosses from the right side to the left side). By our observation on the compatibility of f, this implies that $\tilde{f}(\partial\Pi_k)$ also intersects those rays positively. Let us draw rays starting from the regions $T_1', \dots, T_q'$ as in Figure 13. Label all rays, in particular label these last q with the symbols $\alpha_1, \dots, \alpha_q$.

The *Blank word* of the disc $\tilde{f}$ is obtained by tracing $\tilde{f}(\partial\Pi_k)$, starting from, say, a_m, and writing down the labels of the rays we meet with exponents ± 1 according to positive and negative intersections (we'll call these *positive* and *negative symbols* or *letters*). A *grouping* of the Blank word is a set of properly nested disjoint unordered pairs of the form $\{\,\alpha, \alpha^{-1}\,\}$ so that each negative symbol is part of exactly one pair. Blank's theorem states that the set of groupings of the Blank word is in a one-to-one correspondence with non-equivalent extensions of the immersion $\tilde{f}|_{\partial\Pi_k}$ to immersions of Π_k.

As $\tilde{f}$ is such an extension, a grouping exists and we claim that this implies the Proposition. We delete labels different from the α_s from the Blank word of $\tilde{f}$ and concentrate on the remaining word, which also inherits a grouping (note that most of the deleted labels only appeared positively anyway, except for certain ones that belong to some regions that were the result of perturbation). We'll call this the *Blank word of* f and denote it by W_f. Note that if we only keep the negative letters α_s^{-1} from W_f, then the sequence of their indices is exactly $\{\,i_1, i_2, \dots, i_c\,\}$. In fact, W_f is decomposed into segments $S_1, \dots, S_c$

ended by these negative symbols, and a final segment S:

$$W_f = \overbrace{\alpha_1\alpha_2\dots\alpha_{i_1-1}\alpha_{i_1}^{-1}}^{S_1}\cdot\overbrace{\alpha_1\alpha_2\dots\alpha_{i_2-1}\alpha_{i_2}^{-1}}^{S_2}\cdot\dots \\ \cdot\underbrace{\alpha_1\alpha_2\dots\alpha_{i_c-1}\alpha_{i_c}^{-1}}_{S_c}\cdot\underbrace{\alpha_1\alpha_2\dots\alpha_{m-1}}_{S}.$$

We'll prove the following statements by induction on j:

(1) Each copy of α_{m-j}^{-1} is paired in the grouping with a copy of α_{m-j} which is located to the right of it.

(2) None of the pairs $\{\alpha_{m-j}^{-1},\alpha_{m-j}\}$ is nested in a pair $\{\alpha_{m-l}^{-1},\alpha_{m-l}\}$ for any $l>j$.

(3) Any two copies of α_{m-j}^{-1} are separated in W_f by a copy of α_{m-n}^{-1} for some $n<j$.

The last statement of course directly implies the Proposition.

The last letter of W_f is α_{m-1} and this is the only positive occurrence of this letter. Hence (1) and (2) are obvious for $j=1$. Also, since a second one wouldn't find a pair, there can be at most one copy of α_{m-1}^{-1} in W_f, thus (3) is vacuously true for $j=1$.

Assume that the statements hold for all $j'=1,\dots,j-1$. To prove (1), assume that a certain copy α^{-1} of α_{m-j}^{-1} forms the pair P with a copy α of α_{m-j}, which is to the left of it. Then, as α can't be part of the final segment S, it belongs to some S_b with last letter $\alpha_{i_b}^{-1}$, where $i_b>m-j$, ie $m-i_b<j$. This copy of $\alpha_{i_b}^{-1}=\alpha_{m-(m-i_b)}^{-1}$ is to the left of α^{-1}, hence it is part of a pair which is nested inside P, which contradicts the hypothesis (2) for $j'=m-i_b$.

To prove (2), suppose that a certain pair $\{\alpha_{m-j}^{-1},\alpha_{m-j}\}$ is nested in $P=\{\alpha,\alpha^{-1}\}$ where α has an index less than $m-j$. This copy of α_{m-j} can't be part of S because it has a symbol with lower index (namely, α or α^{-1}) to the right of it. It is clear then that the first negative letter after α_{m-j} has index $m-j'$, which is higher than $m-j$ (ie $j'<j$), and it is still nested in P. Then so is the pair containing it, which contradicts (2) for j'.

Finally for (3), assume the contrary again, namely that there are two copies of α_{m-j}^{-1} in W_f that are not separated by any higher index negative symbol. Then the pair of the first α_{m-j}^{-1} can not lie between them either, because it would be part of some S_b and then $\alpha_{i_b}^{-1}$ would separate. Thus, we have two pairs of the form $\{\alpha_{m-j}^{-1},\alpha_{m-j}\}$ nested in one another. But then the positive symbol of the inner pair would be followed by some $\alpha_{i_b}^{-1}$, which is part of a pair that is nested in the outer pair, and that contradicts (2) for $j' = m - i_b$. □

Definition 6.5 Let $1 \le i, j \le q$. The element $B_{i,j}$ of the DGA of γ_β is the sum of the following products. For each path composed of parts of the strands of the braid β that connects the left endpoint labeled i to the right endpoint labeled j so that it only turns at quadrants facing up, take the product of the crossings from left to right that it turns at.

For example, $B_{i,j}$ contains the constant term 1 if and only if $j = \sigma(i)$. We will need the following polynomials of the $B_{i,j}$:

Definition 6.6 Let $q \ge i > j \ge 1$ and let

$$C_{i,j} = \sum_{\{i_1,\dots,i_c,j\}\in D_i} B_{i,i_1}B_{i_1,i_2}B_{i_2,i_3}\dots B_{i_{c-1},i_c}B_{i_c,j}.$$

Similarly, for $1 \le m \le q$, let

$$C_{m,m} = \sum_{\{i_1,\dots,i_c\}\in D_m} B_{m,i_1}B_{i_1,i_2}B_{i_2,i_3}\dots B_{i_{c-1},i_c}B_{i_c,m}. \tag{12}$$

Finally, for any i and j, let

$$M_{i,j} = \sum_{\{i_1,\dots,i_c\}\in D_{\min\{i,j\}}} B_{i,i_1}B_{i_1,i_2}B_{i_2,i_3}\dots B_{i_{c-1},i_c}B_{i_c,j}.$$

In particular, for each summand in $C_{i,j}$, j is the last element of the admissible sequence but it may occur elsewhere, too. For example, $C_{1,1} = B_{1,1}$, $C_{2,1} = B_{2,1}$, and $C_{3,1} = B_{3,1} + B_{3,2}B_{2,1} + B_{3,1}B_{1,2}B_{2,1}$. Note also that $M_{1,j} = B_{1,j}$, $M_{i,1} = B_{i,1}$, $M_{m,m} = C_{m,m}$ and $M_{i,i-1} = C_{i,i-1}$, whenever these expressions are defined.

Theorem 6.7 $\partial(a_m) = 1 + C_{m,m}$. *Consequently, the index 0 part $H_0(L_\beta)$ of the contact homology $H(L_\beta)$ has a presentation where the generators are the crossings of β and the relations are $C_{m,m} = 1$ for $m = 1, \dots, q$.*

Proof The 1 in the formula comes from the m^{th} trivial disc. We claim that the rest of the contributions add up to $C_{m,m}$. From Proposition 6.4 and the paragraph preceding Definition 6.3, we know that there can't be any such terms other than the ones included in $C_{m,m}$. To see that all such monomials actually arise from admissible discs, we just need to find those discs. This can either be done by an inductive construction on c (for the inductive step, remove the smallest number from the sequence), or by applying Blank's theorem. □

It is not clear whether $H(L_\beta)$ contains any non-zero higher index part at all (except for the case of the unknot, when the single index 1 crossing is a non-nullhomologous cycle). This is mainly why we only work with $H_0(L_\beta)$ in this paper. This difficulty in handling contact homology also underlines the importance of the augmentation that we construct in the next section.

7 Augmentations of braid closures

Definition 7.1 An *augmentation* of the Lagrangian diagram γ of a Legendrian link L is a subset X of its crossings with the following properties.

- All elements of X are proper crossings of L (ie, intersections of different components are not allowed in X).
- The index of each element of X in any admissible grading is 0 (in fact, this requirement implies the previous one).
- For each generator a, the number of admissible discs with positive corner a and all negative corners in X is even.

The last requirement implies that the evaluation homomorphism (which is defined on the link DGA, and which is also called an augmentation) $\varepsilon_X\colon \mathcal{A} \to \mathbf{Z}_2$ that sends elements of X to 1 and other generators to 0, gives rise to an algebra homomorphism $(\varepsilon_X)_*\colon H(L) \to \mathbf{Z}_2$.

Example 7.2 Consider the right-handed Legendrian trefoil knot diagram of Figure 4. We claim that the set $\{\, b_3 \,\}$ is an augmentation. Indeed, the only two non-zero differentials (see Example 2.14 or the previous section) are $\partial(a_1) = 1 + b_1 + b_3 + b_1b_2b_3$ and $\partial(a_2) = b_2 + b_2b_3 + b_1b_2 + b_2b_3b_1b_2$, and even these vanish after mapping b_1 and b_2 to 0 and b_3 to 1.

Remark 7.3 Let us return to the front diagram in the upper half of Figure 13. Such a diagram always has an *admissible decomposition* (or *ruling*) in the sense of [5]: the only two values of the Maslov potential (even though it's $\mathbf{Z}$–valued) are 1 on the upper strands and 0 on the strands of the original braid β, thus all crossings are Maslov, and we may declare all of them switching (in the multi-component case, consider only proper crossings). This gives rise to a decomposition where the discs are nested in one another, so it's admissible.

The existence of a ruling implies that L_β is not a stabilized link type for any positive braid β [5]. Also, by a theorem of Fuchs [18], it implies that the diagram has an augmentation[13] (note this is by no means unique). In his proof, Fuchs constructs an augmentation of a diagram which is equivalent to the original but has a lot more crossings. This means that the original diagram can also be augmented: the pull-back of an augmentation by a DGA morphism, like the ones listed in section 3.3, is again an augmentation.

In the case of a braid closure, it would be impractical to pull back Fuchs' augmentation to the original diagram. Instead, we'll start from scratch and construct an augmentation of the Legendrian closure γ_β of an arbitrary positive braid β. When selecting crossings into X, it will suffice to work with the braid itself, as illustrated in Figure 15, for the index 0 proper crossings of γ_β are all crossings of β. (We will call these the *proper crossings of the braid* β.)

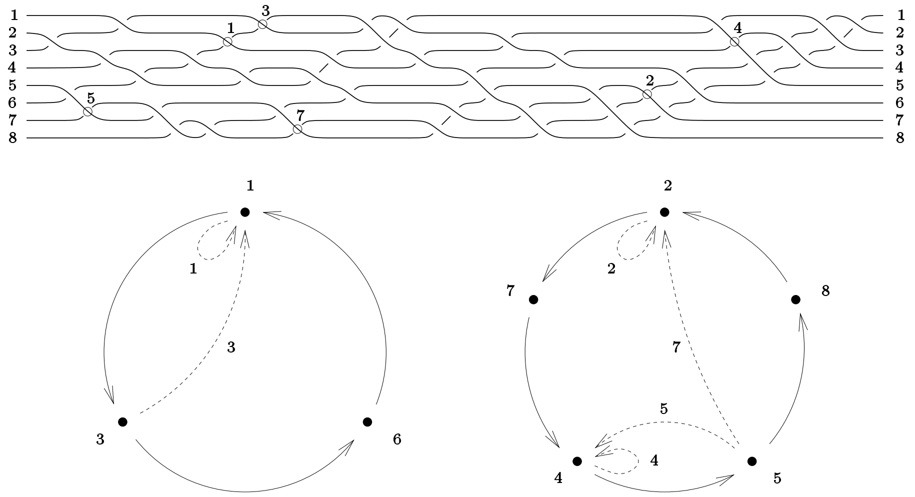

Figure 15: Constructing an augmentation for the closure of a braid

First, we associate an oriented graph to an arbitrary permutation $\sigma \in S_q$. Let o be one of the cycles of σ. Let us write the elements of o around the perimeter of a circle in the cyclic order suggested by σ, directing an edge from s to $\sigma(s)$ for all s. (See Figure 15 for an example. The permutation in the diagram is the one underlying the braid). If p is a non-maximal element of o, then follow the cycle in the forward direction starting from p until it hits the first number in o which is bigger then p. Let p_+ be the number immediately before that. In particular, $p_+ \leq p < \sigma(p_+)$. Do the same in the backward direction, resulting

[13]Fuchs and Ishkhanov [19] and independently Sabloff [25] have also proven that the existence of an augmentation implies the existence of a ruling.

in the number p_- such that $p_- \leq p < \sigma^{-1}(p_-)$. (For example, in Figure 15, 5 is an element of a 5–cycle o, with $5_+ = 5$ and $5_- = 4$.) Then, connect p_+ to p_- by a directed chord of the circle and label the chord by p. If $p_+ = p_- = p$, then instead, we attach a loop edge (labeled p) at p to the graph. This will also be called a chord. Draw this loop edge inside the circle, right next to the perimeter, on the side of p where the smaller of its two neighbors lies.

Definition 7.4 If p is non-maximal in its cycle o, then the oriented loop Γ_p that starts from p, goes along the circle of o to p_+, then goes to p_- on the chord labeled p, then follows the circle again back to p will be called the *loop of* p. If $\hat{p}$ is the largest number in o, then let the *loop of* $\hat{p}$ be the loop $\Gamma_{\hat{p}}$ that travels around the original circle once.

Lemma 7.5 *If $p < r$, then*

(a) *Γ_p doesn't contain r*

(b) *the discs bounded by Γ_p and Γ_r are either disjoint or the latter contains the former.*

In particular, the $|o| - 1$ (oriented) chords obtained in the construction are pairwise disjoint. They can't be parallel to the original edges and they differ from each other as well.

Proof Chords differ from edges because it's impossible that each end of an edge be smaller than the other end (a chord and an edge can be parallel in the non-oriented sense, as in Figure 15). The other statements follow from (a) and (b). Statement (a) is obvious from the construction, and so is (b) if r is maximal in o.

The directed arc of the circle stretching from p_- to p_+ only contains numbers less than r. If this arc is disjoint from Γ_r, then of course so is Γ_p. Otherwise, the whole arc must be contained in Γ_r. From this, statement (b) is clear, except if the chord labeled p is a loop edge and either $p = r_+$ or $p = r_-$. In the first case, $\sigma(p) > r$, but $\sigma^{-1}(p) \leq r$, so (b) holds by the construction of the loop edge. The other case is handled analogously. □

Definition 7.6 Let Γ_σ be the disjoint union of the graphs constructed above over all cycles o of $\sigma \in S_q$. This oriented planar graph, with vertices labeled by the numbers $1, \ldots, q$, is called the *augmented graph of the permutation* σ.

Lemma 7.7 *For all $p \in \{1, \ldots, q\}$, the loop of p is the unique directed loop in Γ_σ starting and ending at p so that the sequence of the vertices (apart from p) visited by it is in D_p.*

Proof It is enough to prove the statement for a connected component associated to a cycle o. The loop Γ_p has the required property for all p, because all the vertices visited by it are less than p and no repetition occurs.

We have to rule out the existence of other loops. Since $\sigma(p_+)$ and $\sigma^{-1}(p_-)$ are larger than p, by the disjointness statement in Lemma 7.5, all loops in question are trapped in the disc bounded by Γ_p (and this is obviously true when p is maximal in o). Then by statement (b) of Lemma 7.5, apart from edges of Γ_p, they may only contain chords labeled by numbers less than p. Now, if the chord labeled by s occurred in the loop and $s < p$ was the smallest such number, then because p is not on Γ_s, the sequence of vertices on Γ_s would appear as a subsequence of the original vertex sequence of the loop. In that case, s would be repeated in the sequence without the two occurrences separated by a larger number, which is a contradiction. □

The following "edge reversal lemma," which we will need in section 10, can be proven very similarly.

Lemma 7.8 *For any $1 \leq p < r \leq q$ so that $\sigma(r) = p$, if p is the second largest vertex (after r) along Γ_r, then there is a unique oriented path in Γ_σ from p to r so that the sequence of the intermediate vertices is in D_p. Otherwise, there is no such path.*

Definition 7.9 Let Y be a set of crossings of the positive braid β. By the *graph realized by* Y we mean the oriented graph with vertices $1, \ldots, q$ so that a directed edge connects i to j if and only if $B_{i,j}\big|_Y = 1$. Here, $B_{i,j}$ is as in Definition 6.5 and by $B_{i,j}\big|_Y$ we mean the element of $\mathbf{Z}_2$ obtained by substituting 1 for elements of Y and 0 for other generators in $B_{i,j}$.

Lemma 7.10 *Let σ be the underlying permutation of β and Y a set of proper crossings of β. If the graph realized by Y agrees with the augmented graph Γ_σ of σ, then Y is an augmentation of the Legendrian closure of β.*

Proof Assume the two graphs do agree. Then by Lemma 7.7, exactly one of the summands of $C_{m,m}$ (see equation (12)) contributes 1 to the sum $C_{m,m}\big|_Y$, namely the one that belongs to the sequence of vertices on Γ_m. Therefore by Theorem 6.7, Y is an augmentation. □

Next, based on Γ_σ, we construct a candidate X, and then we will use Lemma 7.10 to prove that it's an augmentation. Loosely speaking, the edges of Γ_σ

connecting s to $\sigma(s)$ are always realized, even by the empty set. To realize the chord labeled p, we'll select the crossing $b_{p_+,p_-,1}$ into X. We can do this because it always exists: $p_+ \le p < \sigma^{-1}(p_-)$ (for left labels) and $p_- \le p < \sigma(p_+)$ (for right labels), therefore the strand connecting p_+ to $\sigma(p_+)$ always meets the strand connecting $\sigma^{-1}(p_-)$ to p_-. If there were more than one points with the first two labels p_+, p_-, we could have selected any of them[14]; we used the third label 1 for concreteness and for ease in the proof of Theorem 10.1. In Figure 15, we marked the selected crossings and labeled them with the label of the chord that they realize.

Proposition 7.11 *The set*

$$X = \{\, b_{p_+,p_-,1} \mid p \in \{\,1,\ldots,q\,\} \text{ is not a maximal element of a cycle of } \sigma \,\}$$

is an augmentation of the Legendrian closure of the positive braid β with underlying permutation σ.

In particular, for a pure braid β, the empty set is an augmentation. In other words, the DGA of the Legendrian closure of a pure braid is augmented, ie, the boundary of each generator is a polynomial without a constant term. If β is not pure, then $X \neq \varnothing$, hence $\varepsilon_X \neq 0$, and it follows that $H(L_\beta) \neq 0$.

Proof It is clear from the construction that all the crossings in X are proper. By Lemma 7.10, it suffices to prove that the graph G realized by X is the graph Γ_σ. For this, the chief claim is that no two points of X are connected with a part of a strand so that it arrives at both points from above. In other words, the situation of Figure 16 can not arise: there is no pair of numbers p, r so that $\sigma(r_+) = p_-$. Indeed, then we'd have $r < \sigma(r_+) = p_- \le p$ and similarly, $p < \sigma^{-1}(p_-) = r_+ \le r$, which would be a contradiction.

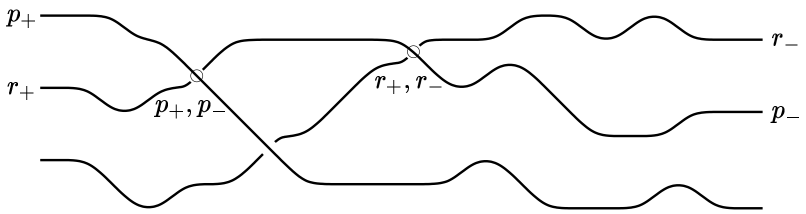

Figure 16: A situation that we have to rule out

[14]For instance, the crossing b_3 of Example 7.2 is denoted by $b_{1,1,2}$ in the general labeling system.

We know then that any path that is to contribute a non-zero summand to a certain $B_{i,j}\big|_X$ can have at most one corner, which of course has to be in X. Thus in the discussion before the Proposition we exhausted all such contributions: paths with no corners are responsible for the edges and paths with one corner are responsible for the chords of Γ_σ. So indeed, $G = \Gamma_\sigma$. □

Remark 7.12 Let β be the standard positive braid whose Legendrian closure is a positive (p,q) torus link (the braid used to produce Figure 13 is an example with $p=5$ and $q=4$). If we apply our construction to it, we find an interesting connection of the resulting augmentation to the Euclidean algorithm. We mention this here without proof; we will only need a small part of the statement which is hidden in the proof of Proposition 10.3.

Let us denote the quotients and residues in the Euclidean algorithm (with input p and q) as follows:

$$\begin{aligned}
p &= k_{-1}q + r_0 && (0 \le r_0 < q)\\
q &= k_0 r_0 + r_1 && (0 \le r_1 < r_0)\\
r_0 &= k_1 r_1 + r_2 && (0 \le r_2 < r_1)\\
r_1 &= k_2 r_2 + r_3 && (0 \le r_3 < r_2)\\
&\ \vdots\\
r_{l-3} &= k_{l-2} r_{l-2} + r_{l-1} && (0 \le r_{l-1} < r_{l-2})\\
r_{l-2} &= k_{l-1} r_{l-1} + r_l && (0 \le r_l < r_{l-1})\\
r_{l-1} &= k_l r_l + 0.
\end{aligned}$$

(Of course, $r_l = \gcd\{p,q\}$.) The points of the augmentation X are arranged in blocks of the following sizes: (k_0-1) blocks of size r_0; k_1 blocks of size r_1; k_2 blocks of size r_2 and so on until the last k_l blocks of size r_l. If we draw the diagram of the braid as in Figure 17, every block can be viewed as the diagonal of a square, and the squares can in turn be seen to be placed inside a $p\times q$ rectangular box so that they realize a 'graphic implementation' of the Euclidean algorithm.

8 A loop of positive links

Let L_β be the Legendrian closure of the positive braid β. Then there exists a natural closed loop in the connected component $\mathcal{L}_\beta$ of the space of Legendrian links that contains L_β, as follows. Let us write $\beta = \lambda_1 \dots \lambda_w$ as a product of

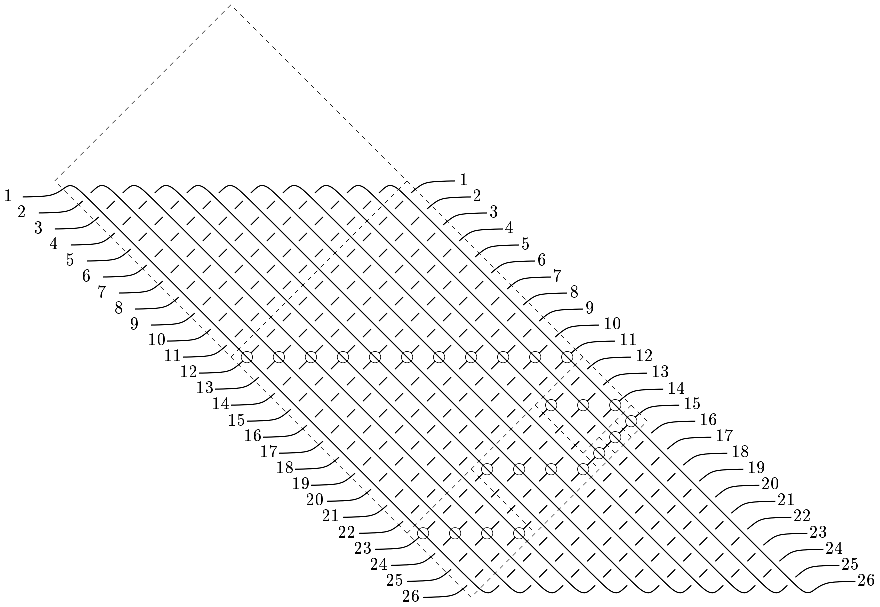

Figure 17: The augmentation of a (p,q) torus link implements the Euclidean algorithm; on the diagram, $p=11$ and $q=26$.

the braid group generators. In Figure 18, we show through an example how L_β can be changed into the Legendrian closure of the conjugate braid that results from moving the first factor λ_1 to the end of the word: if λ_1 is a half-twist of the m^{th} and $(m+1)^{\text{st}}$ strands of the braid, then one interchanges the m^{th} and $(m+1)^{\text{st}}$ strands *above* the braid. In Figure 19, the same path Φ_{λ_1} is shown, but in the Lagrangian projection. Note that the index 1 crossings a_m and a_{m+1} trade places. (The notation used on the diagram for the index 0 crossings is the one that we will introduce below for the special case of torus links.) The Lagrangian diagrams of the endpoints are clearly obtained by resolution of the corresponding fronts. However, we will not prove that the paths themselves agree, too (up to homotopy). Instead, we will content ourselves with checking (using Theorem 4.1) that the four Reidemeister moves in Figure 19 are consistent, and thereafter use the Lagrangian construction as our definition of Φ_{λ_1}. (There are no such consistency issues with Reidemeister moves of fronts, but we need the Lagrangian diagrams to compute holonomies.)

Theorem 8.1 *The sequence of Reidemeister moves in Figure 19, defining Φ_{λ_1}, is consistent.*

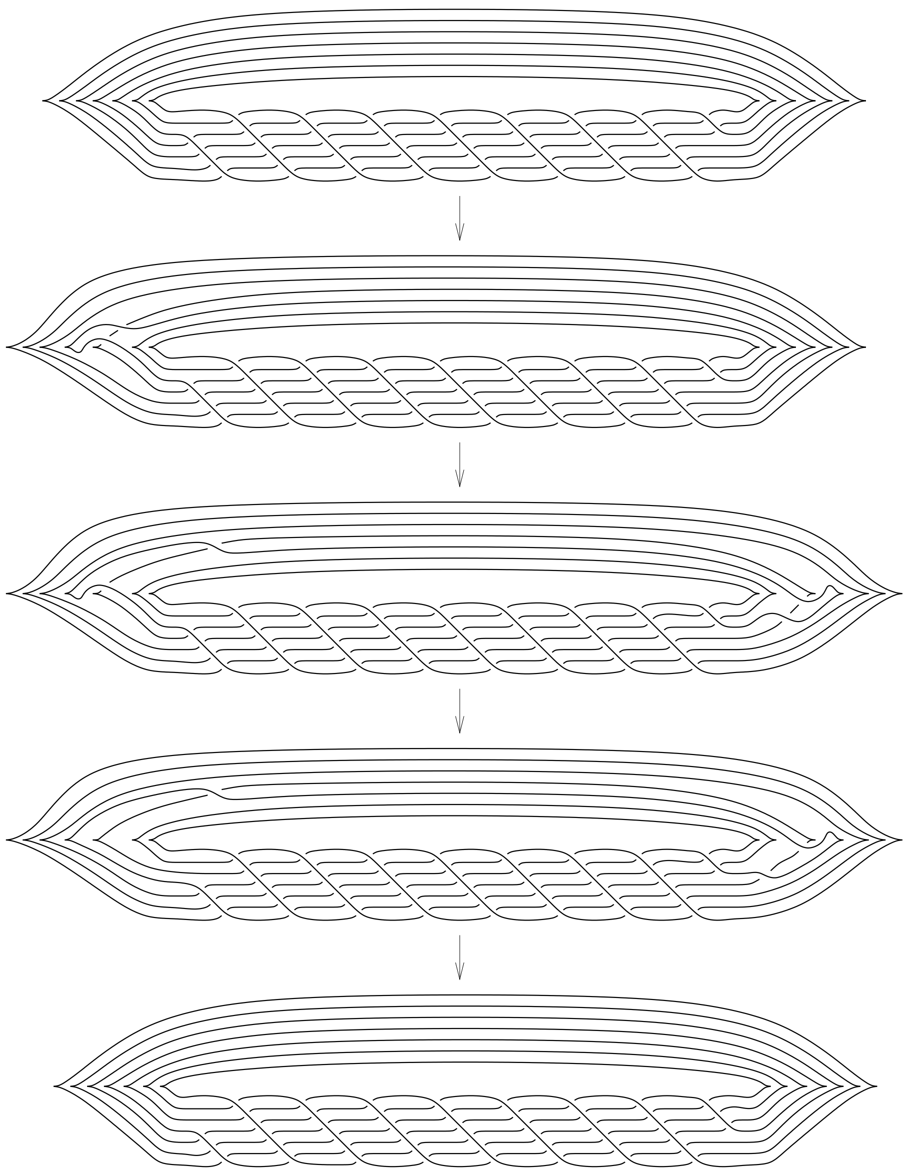

Figure 18: A path that corresponds to conjugating a braid, viewed in the front projection. The moves are similar to those in Figure 1.

Proof This is a generalization of Example 4.2. Isotope the diagram just like in the proof of Lemma 6.2. Then by choosing a small enough ε, the height $h(a_{m+1})$ will dominate the expression whose positivity is needed (by Theorem

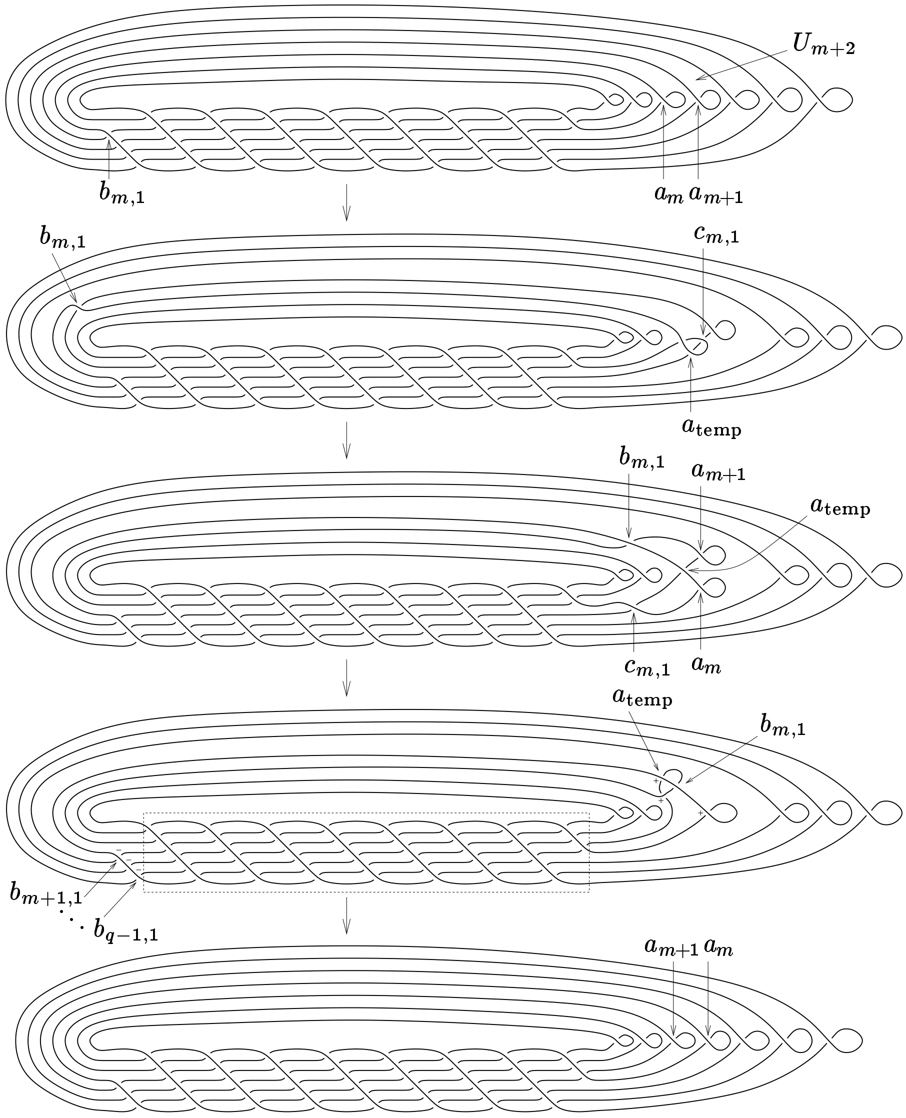

Figure 19: A path that corresponds to conjugating a braid, viewed in the Lagrangian projection. The individual Reidemeister moves *do not* correspond to those in Figure 18.

4.1) in order for the first move (which is a Reidemeister II move) to be consistent. To carry out the $\mathrm{III_b}$ move that follows, we need to isotope the second diagram from the top so that what remains from the region U_{m+2} after the first move has larger area than the vanishing triangle. This can be achieved by the same trick (this time, moving away the 'outer' $q-m-1$ strands). Next, the newborn triangle needs to be blown up so that it has larger area than the triangle which is due to vanish in the second $\mathrm{III_b}$ move (which is the third move altogether). For this, the same trick with the bulges still works: apply it to the $q-m-1$ outer strands and the one that crosses itself at a_m. Finally, the exact same argument guarantees that the fourth move, of type II^{-1}, is consistent too. □

Let us use the formulas of section 3.3 to compute the holonomy μ_1 of Φ_{λ_1} (cf Example 3.6). More precisely, we will compute the action of μ_1 on the index 0 crossings which generate the index 0 contact homology $H_0(L_\beta)$. In the first move, two new crossings appear; using Lemma 2.3 it is easy to show that their indices are 0 and 1. Let us denote the one with index 0 by[15] $c_{m,\sigma(m+1),1}$ and the one with index 1 by a_{temp}. The old index 0 crossings are not affected by this move (ie, the holonomy maps them trivially). This is true by Proposition 3.5: for index reasons, the boundary of any index 0 crossing is 0. It is easy to see that the following two triangle moves don't affect the old index 0 crossings, either.

In the fourth, Reidemeister II^{-1} move however the crossing $b_{m,\sigma(m+1),1}$ (together with a_{temp}) vanishes, and its image in the holonomy becomes the polynomial $M'_{m+1,m} = C'_{m+1,m}$. We used primed symbols to remind us that these are to be computed with respect to the conjugated braid. The fact that for each admissible disc that turns at the positive quadrant at a_{temp} that faces away from $b_{m,\sigma(m+1),1}$, the intermediate sequence of labels (of other index 1 generators through which the boundary of the admissible disk passes) doesn't contain any number larger than m can be shown just like in the proof of Lemma 6.2. The fact that the sequence is admissible follows by an argument very similar to the proof of Proposition 6.4. Finally, the fact that each such admissible sequence does contribute the said terms can be established as in the proof of Theorem 6.7. We have proven:

Proposition 8.2 *The holonomy μ_1 of Φ_{λ_1} maps each crossing of the braid trivially except the first one from the left, which is mapped to the polynomial $M'_{m+1,m} = C'_{m+1,m}$ (if the first crossing is between the m^{th} and $(m+1)^{\mathrm{st}}$ strands). This expression is to be computed as in Definition 6.6, with respect to the conjugate braid $\lambda_2\lambda_3\dots\lambda_w\lambda_1$.*

[15]In other words, let it inherit the labels of the crossing that is being moved to the other end of the braid. This is what we've done in Figure 19 too, except that there, a different notation is used for index 0 crossings.

Note also that at the end stage of Φ_{λ_1} each crossing of the (conjugate) braid has index 0, thus in the multi-component case our choice of preferred grading is justified.

Now, it is clear that the concatenation $\Omega_\beta = \Phi_{\lambda_1} \ldots \Phi_{\lambda_w}$ is a closed loop in $\mathcal{L}_\beta$ and by definition, its monodromy is the composition of holonomies $\mu_w \circ \ldots \circ \mu_1$, followed by a re-labeling to restore the original labels. Namely, each symbol c labeling an index 0 crossing needs to be changed back to b; in fact, a_i would have to be changed to $a_{\sigma(i)}$, but because we only concern ourselves with $H_0(L_\beta)$, this can be ignored.

9 $|\mu_0|$ divides $p+q$

The last two sections of the paper contain the proof of Theorem 1.3. The argument works for any p,q except when q divides p (the case of a pure braid), or p divides q. The reason why we don't claim Theorems 1.3 and 1.2 for multi-component torus links is that we only proved Theorems 3.2 and 3.7 (and thus Theorem 1.1) for knots. (However, the extension of those proofs should only be a matter of changing the formalism to that of link contact homology.)

Let us revisit the loop $\Omega_{p,q}$ of Legendrian (p,q) torus knots defined in the introduction. As the braid β is now composed of p periods, the general loop Ω_β described in section 8 is the p-fold concatenation of another, and the latter is easy to identify as homotopic[16] to $\Omega_{p,q}$. In particular, by Theorem 3.7, the 'full' monodromy takes the form μ^p, where μ is the monodromy of $\Omega_{p,q}$. From now on, we will concentrate on this map μ, and especially on its restriction μ_0 to the index 0 part of the contact homology of the standard torus link diagram γ (shown in Figure 13), representing the base point L.

We will adjust our notation to this special situation. The crossings of β will be indexed with two integers (as opposed to three), namely $b_{m,n}$ ($m = 1, \ldots, q-1$, $n = 1, \ldots, p$) will denote the m^{th} crossing counted from the top in the n^{th} period of the braid. Note that in the definition of μ, after a full period of the braid has been moved from the left end to the right end, a re-labeling takes place, too: the second label of each crossing in the other $(p-1)$ periods is reduced by 1, and the labels $c_{m,1}$ in the now last period are changed to $b_{m,p}$.

[16]Because we'll omit the rigorous justification of this fact, the reader may treat the new description of $\Omega_{p,q}$ as the definition.

Proposition 9.1 *The monodromy μ of the loop $\Omega_{p,q}$ of Legendrian torus links acts on the index 0 generators as follows:*

$$\mu(b_{m,n}) = \begin{cases} b_{m,n-1} & \text{if } 2 \le n \le p \\ C_{q,m} & \text{if } n = 1 \end{cases} .$$

Proof The claim is clear for those crossings not in the first period: by Proposition 8.2, they are only affected, and in the described way, by the re-labeling. The rest of the statement will be proven by induction on $q-m$. When this value is 1, ie $m = q-1$, this is just the statement of Proposition 8.2 (the conjugate braid in this case is the original β again, and the re-labeling changes $C'_{q-1+1,q-1}$ into $C_{q,q-1}$). Assume the statement holds for $b_{q-1,1}, \dots, b_{m+1,1}$. Right after the conjugation that removes it from the left end of the braid, the image of $b_{m,1}$ is $M''_{m+1,m}$, computed with respect to the braid after this conjugation. This can be re-written (by grouping terms with respect to the first factor in the product) as

$$M''_{m+1,m} = b_{m+1,1}M'_{m+1,m} + b_{m+2,1}M'_{m+2,m} + \dots + b_{q-1,1}M'_{q-1,m} + M'_{q,m},$$

where the terms labeled M' on the right are to be computed in the braid indicated by the box in Figure 19. Note however that by Definition 6.6, the same expressions are obtained if we use the whole braid β (before re-labeling). If we apply the holonomies of the remaining $q-1-m$ conjugations and the re-labeling to this expression, we get (by the inductive hypothesis)

$$\mu(b_{m,1}) = C_{q,m+1}M_{m+1,m} + C_{q,m+2}M_{m+2,m} + \dots + C_{q,q-1}M_{q-1,m} + M_{q,m} = C_{q,m}.$$

The last equality is true because the middle expression is exactly what results if we group terms in $C_{q,m}$ with respect to the last label in the admissible sequence which is more than m. □

Proposition 9.2 *In the contact homology ring $H(L)$, we have:*

$$\mu(B_{i,j}) = \begin{cases} B_{i-1,j-1} + B_{i-1,q}b_{j-1,p} & \text{if } i,j \ge 2 \\ B_{i-1,q} & \text{if } i \ge 2 \text{ and } j = 1 \\ b_{j-1,p} & \text{if } i = 1 \text{ and } j \ge 2 \end{cases} , \tag{13}$$

$$\mu(C_{i,j}) = \begin{cases} C_{i-1,j-1} & \text{if } j \ge 2 \\ M_{i-1,q} & \text{if } j = 1 \end{cases} , \tag{14}$$

and

$$\mu(M_{i,j}) = M_{i-1,j-1}, \text{ whenever } i,j \ge 2. \tag{15}$$

We omitted $B_{1,1}$ and $C_{i,i}$ because they (and hence their images) are equal to 1 in the contact homology. Recall also that $M_{i,1} = B_{i,1}$ and $M_{1,j} = B_{1,j}$.

Proof When $i \geq 2$, none of the terms in $B_{i,j}$ contains any of $b_{1,1}, \dots, b_{q-1,1}$, so they only need to be re-labeled. This means that all crossings that the path through the braid which generated the term turned at, are shifted to the left by a unit. This operation changes the entry point from the one labeled i to the one labeled $i-1$. If $j=1$, then the shifted path can be completed by the overcrossing strand of the last period of the braid, which shows that the re-labeled expression is a summand in $B_{i-1,q}$. Moreover, all such summands are obtained in this way exactly once. When $j \geq 2$, the re-labeling results in a summand of $B_{i-1,j-1}$, but not all such are obtained: we miss contributions from paths that turn at the crossing (the last one on the strand with right endpoint $j-1$) $b_{j-1,p}$. Hence the correction term in the top row of (13) (note that the paths turning at $b_{j-1,p}$ are exactly those that would otherwise have arrived at q).

When $i=1$ and $j \geq 2$, we have

$$\begin{aligned}
\mu(B_{1,j}) &= \mu(b_{1,1}B_{2,j} + b_{2,1}B_{3,j} + \dots + b_{q-1,1}B_{q,j} + R) \\
&= C_{q,1}(B_{1,j-1} + B_{1,q}b_{j-1,p}) + C_{q,2}(B_{2,j-1} + B_{2,q}b_{j-1,p}) + \dots \\
&+ C_{q,q-1}(B_{q-1,j-1} + B_{q-1,q}b_{j-1,p}) + \mu(R) \\
&= C_{q,1}B_{1,j-1} + C_{q,2}B_{2,j-1} + \dots + C_{q,q-1}B_{q-1,j-1} \\
&+ (C_{q,1}B_{1,q} + C_{q,2}B_{2,q} + \dots + C_{q,q-1}B_{q-1,q})b_{j-1,p} \\
&+ B_{q,j-1} + B_{q,q}b_{j-1,p} \\
&= C_{q,j-1} + C_{q,j-1}C_{j-1,j-1} + C_{q,q}b_{j-1,p} = C_{q,j-1} + C_{q,j-1} + b_{j-1,p} \\
&= b_{j-1,p}.
\end{aligned}$$

Here, R is the sum of the contributions to $B_{1,j}$ that don't contain crossings of the first period. These terms only have to be re-labeled and that can be done just like in the argument above for $B_{i,j}$ when $i \geq 2$. In the sum of sums

$$C_{q,1}B_{1,j-1} + C_{q,2}B_{2,j-1} + \dots + C_{q,q-1}B_{q-1,j-1} + B_{q,j-1},$$

we re-grouped the terms; those with an admissible sequence of labels formed $C_{q,j-1}$, and the rest, where $j-1$ was repeated 'illegally,' formed $C_{q,j-1}C_{j-1,j-1}$.

Note that by the now proven (13), for all $i,j \geq 2$, $\mu(B_{i,j} + B_{i,1}B_{1,j}) = B_{i-1,j-1}$.

Therefore, when $i > j \geq 2$,

$$
\begin{aligned}
\mu(C_{i,j}) &= \mu\left(\sum_{\{i_1,\dots,i_c,j\}\in D_i} B_{i,i_1}B_{i_1,i_2}B_{i_2,i_3}\dots B_{i_{c-1},i_c}B_{i_c,j}\right)\\
&= \mu\left(\sum_{\substack{\{j_1,\dots,j_d,j\}\in D_i\\ j_1,\dots,j_d\geq 2}} \begin{matrix}(B_{i,j_1}+B_{i,1}B_{1,j_1})(B_{j_1,j_2}+B_{j_1,1}B_{1,j_2})\dots\\ (B_{j_{d-1},j_d}+B_{j_{d-1},1}B_{1,j_d})(B_{j_d,j}+B_{j_d,1}B_{1,j})\end{matrix}\right)\\
&= \sum_{\substack{\{j_1,\dots,j_d,j\}\in D_i\\ j_1,\dots,j_d\geq 2}} B_{i-1,j_1-1}B_{j_1-1,j_2-1}\dots B_{j_{d-1}-1,j_d-1}B_{j_d-1,j-1}\\
&= C_{i-1,j-1}.
\end{aligned}
$$

As a consequence of this and (13), for all $i \geq 2$,

$$\mu(C_{i,1}) = \mu\left(B_{i,1}+\sum_{j=2}^{i-1}C_{i,j}B_{j,1}\right) = B_{i-1,q}+\sum_{j=2}^{i-1}C_{i-1,j-1}B_{j-1,q} = M_{i-1,q}.$$

Finally, when $2 \leq i < j$,

$$
\begin{aligned}
\mu(M_{i,j}) &= \mu\left(B_{i,j}+\sum_{k=1}^{i-1}C_{i,k}B_{k,j}\right)\\
&= B_{i-1,j-1}+B_{i-1,q}b_{j-1,p}+M_{i-1,q}b_{j-1,p}\\
&\quad+\sum_{k=2}^{i-1}C_{i-1,k-1}(B_{k-1,j-1}+B_{k-1,q}b_{j-1,p})\\
&= M_{i-1,q}b_{j-1,p}+\left(B_{i-1,j-1}+\sum_{k=2}^{i-1}C_{i-1,k-1}B_{k-1,j-1}\right)\\
&\quad+\left(B_{i-1,q}+\sum_{k=2}^{i-1}C_{i-1,k-1}B_{k-1,q}\right)b_{j-1,p}\\
&= M_{i-1,q}b_{j-1,p}+M_{i-1,j-1}+M_{i-1,q}b_{j-1,p}\\
&= M_{i-1,j-1},
\end{aligned}
$$

and when $i > j$, the argument is very similar to the one we gave for $C_{i,j}$. □

Theorem 9.3 *The order of the (restricted) monodromy $\mu_0 = \mu\big|_{H_0(L)}$ of the loop $\Omega_{p,q}$ of Legendrian (p,q) torus knots divides $p+q$.*

Proof This is now a straightforward computation, generalizing the first paragraph of the proof of Proposition 5.1. Consider the generator $b_{m,p}$ ($m = 1, \ldots, q-1$). By Proposition 9.1, the first p iterations of μ act on it as follows:

$$\mu(b_{m,p}) = b_{m,p-1},\ \mu^2(b_{m,p}) = b_{m,p-2}, \ldots, \mu^{p-1}(b_{m,p}) = b_{m,1},\ \mu^p(b_{m,p}) = C_{q,m}.$$

Then by (14) of Proposition 9.2, the next m iterations are as follows:

$$\mu^{p+1}(b_{m,p}) = C_{q-1,m-1}, \ldots, \mu^{p+m-1}(b_{m,p}) = C_{q-m+1,1},\ \mu^{p+m}(b_{m,p}) = M_{q-m,q}.$$

Now by (15), the next $q-m-1$ iterations are

$$\mu^{p+m+1}(b_{m,p}) = M_{q-m-1,q-1}, \ldots, \mu^{p+q-1}(b_{m,p}) = M_{1,m+1}.$$

Finally, because $M_{1,m+1} = B_{1,m+1}$, (13) yields

$$\mu^{p+q}(b_{m,p}) = b_{m,p}.$$

Because $b_{m,n}$ is on the orbit of $b_{m,p}$ for all $n = 1, \ldots, p$, we see that μ^{p+q} is the identity on all of the degree 0 generators. □

10 $(p+q)$ divides $|\mu_0|$

Theorem 10.1 *The number $p+q$ divides the order of the monodromy μ.*

In the proof of Theorem 9.3, we described explicitly the $(p+q)$-element orbit of each index 0 generator $b_{m,n}$ (altogether $q-1$ orbits). Recall that all of those orbits contain a *b–sequence* $b_{m,p}, b_{m,p-1}, \ldots, b_{m,1}$ of length p, a *C–sequence* $C_{q,m}, C_{q-1,m-1}, \ldots, C_{q-m+1,1}$ of length m and an *M–sequence* $M_{q-m,q}, M_{q-m-1,q-1}, \ldots, M_{1,m+1}$ of length $q-m$. (These expressions are cycles in the chain complex $\mathcal{A}$, but we really mean the homology classes represented by them.) We will choose one of these orbits and show that it doesn't have a period shorter than $p+q$. For this, we will evaluate the augmentation $\varepsilon = \varepsilon_X$ (see Proposition 7.11) on elements of the orbit, and prove that the resulting sequence of 0's and 1's, which we will call the *0-1–sequence of the orbit*, has no such shorter period. In fact, we claim the following:

Proposition 10.2 *If $q > p$ but $p \nmid q$, the 0-1–sequence S of the orbit of $b_{p,p}$ consists of p consecutive 0's and q consecutive 1's. If $q < p$ but $q \nmid p$, then the same holds for the orbit of $b_{[p \pmod q)],p}$.*

In the latter case, we will denote the value $1 \leq [p \pmod q] \leq q-1$ by r_0. For the rest of the section, either this value r_0 or p, as the case may be, should be substituted for m in the formulas for the b–, C–, and M–sequences. Note that if $q \mid p$, the said orbit doesn't even exist (if $q > p$ and $p \mid q$, then its 0-1–sequence consists only of 1's).

Recall that X was constructed so that the graph realized by X was the augmented graph of the underlying permutation σ of the braid. Hence, this oriented graph Γ_σ has adjacency matrix $[\varepsilon_X(B_{i,j})]$ and therefore it contains all the information we need to evaluate the algebra homomorphism $\varepsilon = \varepsilon_X$ on the polynomial expressions of the C– and M–sequences. We will only need to refer to the actual braid in the case of the b–sequence. In our situation,

$$\sigma(i) = [(i-p) \pmod q], \; i = 1, \ldots, q.$$

When $q < p$, we could equivalently write $\sigma(i) = [(i - r_0) \pmod q]$. This explains why our choice of orbit in Proposition 10.2 is reasonable: both in the C–sequence and in the M–sequence the two lower indices are always the endpoints of an edge of Γ_σ, but *they are listed in the reverse order.* So when we evaluate ε on these polynomials, what we need to examine is whether the given edge of the graph can be "reversed," ie, if it is part of an oriented loop (and how many loops) with an admissible sequence of vertices.

We will re-state and prove Proposition 10.2 in a more detailed version.

Proposition 10.3 *The orbit specified in Proposition 10.2 contributes 0's and 1's to the sequence S as follows.*

(1) *If $2p \leq q$ but $p \nmid q$ (hence in fact $2p < q$), then we have*

$$\overbrace{b_{p,p}, \ldots, b_{p,1}}^{p \text{ copies of } 1}; \overbrace{C_{q,p}, \ldots, C_{q-p+1,1}}^{p \text{ copies of } 1};$$
$$\overbrace{M_{q-p,q}, \ldots, M_{q-2p+1,q-p+1}}^{p \text{ copies of } 0}, \overbrace{M_{q-2p,q-p}, \ldots, M_{1,p+1}}^{q-2p \text{ copies of } 1}.$$

(2) *If $2p > q$ but $q > p$, then the sequence is*

$$\overbrace{b_{p,p}, \ldots, b_{p,q-p+1}}^{2p-q \text{ copies of } 0}, \overbrace{b_{p,q-p}, \ldots, b_{p,1}}^{q-p \text{ copies of } 1}; \overbrace{C_{q,p}, \ldots, C_{q-p+1,1}}^{p \text{ copies of } 1}; \overbrace{M_{q-p,q}, \ldots, M_{1,p+1}}^{q-p \text{ copies of } 0}.$$

(3) *If $q < p$, $q \nmid p$, and $2r_0 < q$, then we get*

$$\overbrace{b_{r_0,p}, \ldots, b_{r_0,r_0+1}}^{p-r_0 \text{ copies of } 0}, \overbrace{b_{r_0,r_0}, \ldots, b_{r_0,1}}^{r_0 \text{ copies of } 1}; \overbrace{C_{q,r_0}, \ldots, C_{q-r_0+1,1}}^{r_0 \text{ copies of } 1};$$
$$\overbrace{M_{q-r_0,q}, \ldots, M_{r_0+1,2r_0+1}}^{q-2r_0 \text{ copies of } 1}, \overbrace{M_{r_0,2r_0}, \ldots, M_{1,r_0+1}}^{r_0 \text{ copies of } 0}.$$

(4) *Finally, if $q < p$, $q \nmid p$, and $2r_0 \geq q$, then we have*

$$\overbrace{b_{r_0,p},\ldots,b_{r_0,q-r_0+1}}^{p-q+r_0 \text{ copies of } 0},\overbrace{b_{r_0,q-r_0},\ldots,b_{r_0,1}}^{q-r_0 \text{ copies of } 1};$$

$$\overbrace{C_{q,r_0},\ldots,C_{q-r_0+1,1}}^{r_0 \text{ copies of } 1};\overbrace{M_{q-r_0,q},\ldots,M_{1,r_0+1}}^{q-r_0 \text{ copies of } 0}.$$

Proof ***b*–sequence** In case (1), all numbers $1 \leq j \leq p$ are so that $j < [(j+p) \pmod q] = j+p$ and $j < [(j-p) \pmod q]$. This means that $j_+ = j_- = j$, ie, that there is a loop edge attached to j in Γ_σ. It is easy to check that in the construction of X, the crossing that realizes this loop edge is exactly $b_{p,j}$.

In case (2), we similarly find loop edges but only attached to the numbers $1,\ldots,q-p$. These are realized by the crossings $b_{p,1},\ldots,b_{p,q-p}$. For the crossings $b_{p,q-p+1},\ldots,b_{p,p}$, we find that their second labels in the system that we used to label crossings of general braids in section 6 are also $q-p+1,\ldots,p$. These numbers can not be the endpoints of a chord because the numbers preceding them in the permutation (namely, $1,\ldots,2p-q$) are smaller than them. Therefore these crossings are indeed not selected into X.

In cases (3) and (4), first note that the crossings $b_{r_0,q+1},\ldots,b_{r_0,p}$ have third labels greater than 1 in the labeling system of section 6, so they never get selected into X. Neither do $b_{r_0,r_0+1},\ldots,b_{r_0,q}$, because their 'old' first labels are $r_0+1,\ldots,q$ and these are taken to the smaller values $1,\ldots,q-r_0$ by σ (ie, they'll never be the startpoint of a chord). After this, the rest of the b–sequence can be sorted out just like in the first two cases.

***C*–sequence** Here the claim is that it always contributes only 1's to S. This is because $i > \sigma^{-1}(i) = [(i+p) \pmod q]$ implies $\varepsilon(C_{i,[(i+p) \pmod q)]}) = 1$. Indeed, since it is preceded in the permutation by a smaller number, i can't be the endpoint of a chord, only of the single edge coming from $\sigma^{-1}(i)$. So there is a unique term in $C_{i,[(i+p) \pmod q)]}$ that contributes 1 to $\varepsilon(C_{i,[(i+p) \pmod q)]})$, namely the one which, when multiplied by $B_{[(i+p) \pmod q)],i}$ on the right, produces the term corresponding to the unique loop described in Lemma 7.7.

***M*–sequence** All the four claims in this case follow from Lemma 7.8. Note in particular that if we allowed $p \mid q$ in case (1), then $j-p$ would be the second largest vertex on Γ_j for all $j = p+1,\ldots,q$, and therefore the first p elements of the M–sequence wouldn't be mapped to 0 by ε. However if $p \nmid q$, then the first number x after the sequence $j, j-p, j-2p, \ldots$ 'wraps around' the circle $\mathbf{Z}_q$ is different from j. If x is also smaller than j (and this will be the case exactly when $j = q, q-1, \ldots, q-[q \pmod p]+1$), then it falls between

j and $j-p$, so $j-p$ is not second largest in Γ_j. If x is larger than j, then it is recognized as what we called $\sigma(j_+)$ in section 7. But then there exists a chord in Γ_σ, starting from j_+, and ending at the element j_- of Γ_j. This j_- is by construction such that $[(j_- + p) \pmod q] > j$. But because $p < q/2$, if we assume that $j > q-p$, then this is only possible if $j_- + p > j$, ie if $j_- > j-p$. This again means that in these cases, $j-p$ is not second largest on Γ_j. This proves the claim about the first part of the M–sequence. Finally, if $j \leq q-p$ (which implies $j_- = j$), then Γ_j only visits the positive elements of the arithmetic progression $j, j-p, j-2p, \ldots$ (ie, there is no wrapping around) and $j-p$ is obviously second largest among these. Cases (2), (3) and (4) can be handled similarly. □

Note that this was a generalization of the second paragraph in the proof of Proposition 5.1; there, we used the augmentation $X = \{b_3\}$.

References

[1] **V I Arnol'd**, *Invariants and perestroikas of fronts on a plane*, Trudy Mat. Inst. Steklov. 209 (1995) 14–64

[2] **C Benham**, **X-S Lin**, **D Miller**, *Subspaces of knot spaces*, Proc. Amer. Math. Soc. 129 (2001) 3121–3127

[3] **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

[4] **F Bourgeois**, *Contact homology and homotopy groups of the space of contact structures*, `arXiv:math.SG/0407531`

[5] **Y Chekanov**, *New invariants of Legendrian knots*, from: "European Congress of Mathematics, Vol. II (Barcelona, 2000)", Progr. Math. 202, Birkhäuser, Basel (2001) 525–534

[6] **Y Chekanov**, *Differential algebra of Legendrian links*, Invent. Math. 150 (2002) 441–483

[7] **T Ekholm**, **J Etnyre**, **M Sullivan**, *Legendrian submanifolds in $\mathbf{R}^{2n+1}$ and contact homology*, `arXiv:math.SG/0210124`

[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

[9] **Y Eliashberg**, **A Givental**, **H Hofer**, *Introduction to symplectic field theory*, Geom. Funct. Anal. (2000) 560–673

[10] **J B Etnyre**, *Legendrian and Transversal Knots*, `arXiv:math.SG/0306256`

[11] **J B Etnyre**, **K Honda**, *Knots and contact geometry. I. Torus knots and the figure eight knot*, J. Symplectic Geom. 1 (2001) 63–120

[12] **J B Etnyre**, **K Honda**, *On connected sums and Legendrian knots*, Adv. Math. 179 (2003) 59–74

[13] **J B Etnyre**, **K Honda**, *Cabling and transverse simplicity*, e-print, `arXiv:math.SG/0306330`

[14] **J B Etnyre**, **L L Ng**, **J M Sabloff**, *Invariants of Legendrian knots and coherent orientations*, J. Symplectic Geom. 1 (2002) 321–367

[15] **Y Félix**, **S Halperin**, **J-C Thomas**, *Differential graded algebras in topology*, from: "Handbook of algebraic topology", North-Holland, Amsterdam (1995) 829–865

[16] **A Floer**, *Symplectic fixed points and holomorphic spheres*, Comm. Math. Phys. 120 (1989) 575–611

[17] **G K Francis**, *Extensions to the disk of properly nested plane immersions of the circle*, Michigan Math. J. 17 (1970) 377–383

[18] **D Fuchs**, *Chekanov-Eliashberg invariant of Legendrian knots: existence of augmentations*, J. Geom. Phys. 47 (2003) 43–65

[19] **D Fuchs**, **T Ishkhanov**, *Invariants of Legendrian knots and decompositions of front diagrams*, Mosc. Math. J. 4 (2004) 707–717, 783

[20] **A Hatcher**, *Spaces of knots*, `arXiv:math.GT/9909095`

[21] **M Hutchings**, *Floer homology of families I* `arXiv:math.SG/0308115`

[22] **M Hutchings**, *Floer homology of families II: symplectomorphisms*, in preparation

[23] **L L Ng**, *Computable Legendrian invariants*, Topology 42 (2003) 55–82

[24] **V Poénaru**, *Extensions des immersions en codimension 1 (d'après Samuel Blank)*, Séminaire Bourbaki, Vol. 10, Exp. No. 342, Soc. Math. France, Paris (1995) 473–505

[25] **J M Sabloff**, *Augmentations and rulings of Legendrian knots*, Int. Math. Res. Not. (2005) 1157–1180

[26] **P Seidel**, π_1 *of symplectic automorphism groups and invertibles in quantum homology rings*, Geom. Funct. Anal. 7 (1997) 1046–1095