

Holomorphic disks, link invariants and the multi-variable Alexander polynomial

PETER OZSVÁTH
ZOLTÁN SZABÓ

The knot Floer homology is an invariant of knots in S^3 whose Euler characteristic is the Alexander polynomial of the knot. In this paper we generalize this to links in S^3 giving an invariant whose Euler characteristic is the multi-variable Alexander polynomial. We study basic properties of this invariant, and give some calculations.

57M27; 57M25

1 Introduction

The knot Floer homology defined by the authors [20] and by Rasmussen [24] is an invariant for knots in S^3 whose Euler characteristic is the Alexander polynomial of the knot. Our aim here is to give a suitable generalization of this invariant to links in S^3 , giving rise to an invariant whose Euler characteristic is the multi-variable Alexander polynomial.

Specifically, let $L \subset S^3$ be a link with ℓ components. Let $H = H_1(S^3 - L; \mathbb{Z})$. Let $\mathbb{Z}[H]$ denote the group-ring of H , written as sums

$$\sum_{h \in H} a_h \cdot e^h,$$

where $a_h \in \mathbb{Z}$ is zero for all but finitely many $h \in H$. Note that $H \cong \mathbb{Z}^\ell$ is generated by the meridians $\{\mu_i\}_{i=1}^\ell$ for the components K_i of L . Thus, an orientation for L , denoted by \vec{L} , induces an identification between $\mathbb{Z}[H]$ and the ring of Laurent polynomials in ℓ variables (corresponding to the components of L). Consider now the affine lattice $\mathbb{H} = \mathbb{H}(L)$ over H , given by elements

$$\sum_{i=1}^{\ell} a_i \cdot [\mu_i],$$

where $a_i \in \mathbb{Q}$ satisfies the property that $2a_i + \text{lk}(K_i, L - K_i)$ is an even integer, where here lk denotes linking number.

We define a link invariant $\widehat{\text{HFL}}(\vec{L})$, which has the structure of a vector space over the field $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$ equipped with a splitting into direct summands indexed by pairs consisting of an integer (the “homological grading”) and an element of \mathbb{H}

$$\widehat{\text{HFL}}(\vec{L}) \cong \bigoplus_{d \in \mathbb{Z}, h \in \mathbb{H}} \widehat{\text{HFL}}_d(\vec{L}, h).$$

A few remarks are in order about these gradings. First of all, the ranks of the groups $\widehat{\text{HFL}}(\vec{L}, h)$ are independent of the orientation of L , but their homological gradings depend on this data. Its graded Euler characteristic is the Alexander polynomial in the following sense.

Recall that an ℓ -component link L has a symmetric multi-variable Alexander polynomial Δ_L . The link invariant is related to the multi-variable Alexander polynomial by the relation

$$(1) \quad \sum_{h \in \mathbb{H}} \chi(\widehat{\text{HFL}}_*(\vec{L}, h)) \cdot e^h = \begin{cases} \left(\prod_{i=1}^{\ell} (T_i^{\frac{1}{2}} - T_i^{-\frac{1}{2}}) \right) \cdot \Delta_L & \text{if } \ell > 1, \\ \Delta_L & \text{if } \ell = 1. \end{cases}$$

The symmetry of the Alexander polynomial has the following manifestation in link Floer homology: There is an identification

$$\widehat{\text{HFL}}_*(\vec{L}, h) \cong \widehat{\text{HFL}}_{*-2\delta(h)}(\vec{L}, -h),$$

where here

$$\delta \left(\sum_{i=1}^{\ell} a_i \cdot [\mu_i] \right) = \sum_{i=1}^{\ell} a_i.$$

This invariant should be compared with the link invariant described in [20]. Specifically, that paper gives an invariant for oriented links $\vec{L} \subset S^3$

$$\widehat{\text{HFK}}(\vec{L}) \cong \bigoplus_{d, s \in \mathbb{Z}} \widehat{\text{HFK}}_d(\vec{L}, s),$$

using the observation that an oriented ℓ -component link in S^3 naturally gives rise to a null-homologous knot in $\#^{\ell-1}(S^2 \times S^1)$, to which one can in turn apply the knot Floer homology functor, obtaining a bigraded theory associated to this link. One of the factors of the bigrading comes from the \mathbb{Z} -grading by $s \in \mathbb{Z}$ as above, and the second comes from the internal homological grading of the Heegaard Floer homology of $\#^{\ell-1}(S^2 \times S^1)$ (whose grading takes values in $\mathbb{Z} + \frac{\ell-1}{2}$). Taking the Euler characteristic of this knot Floer homology in a suitable sense gives a normalized

version of the Alexander-Conway polynomial of the oriented link. An orientation \vec{L} of a link L gives rise to a homomorphism $o: H \rightarrow \mathbb{Z}$, which extends to a map of \mathbb{H} to the integers. Under the map o , the multi-variable Alexander polynomial is carried to the Alexander-Conway polynomial. This fact admits the following generalization on the homological level:

Theorem 1.1 *Let $L \subset S^3$ be a link endowed with an orientation, denoted \vec{L} . Then, we have an identification*

$$\widehat{\text{HFK}}_{*+(\frac{\ell-1}{2})}(\vec{L}, s) \cong \sum_{\{h \in \mathbb{H} \mid o(h)=s\}} \widehat{\text{HFL}}_*(\vec{L}, h),$$

where $o: \mathbb{H} \rightarrow \mathbb{Z}$ is the natural homomorphism induced by the orientation.

Knot Floer homology can be viewed as the homology of the graded object associated to a filtered chain complex whose total homology is \mathbb{Z} ; and indeed, the filtered chain homotopy type of this complex is a knot invariant, cf Ozsváth–Szabó [19] and Rasmussen [24] (see also Lee [13], Rasmussen [25] and Gornik [6] for corresponding results in Khovanov’s homology, Khovanov [9] and Khovanov–Rozansky [11]). Thus, there is a spectral sequence starting with knot Floer homology, and converging to an E^∞ term which has rank one. We have the following generalization of this fact to link Floer homology:

Theorem 1.2 *There is a spectral sequence whose E_1 term is $\widehat{\text{HFL}}(\vec{L})$, and whose E^∞ term is isomorphic to the exterior algebra $\Lambda^* V$, where V is a vector space of rank $\ell - 1$. In fact, the spectral sequence is an invariant of the link L .*

The proof is given in Section 6, but we pause here for a few remarks on the construction of this spectral sequence and its meaning.

An orientation for L gives a basis (m_1, \dots, m_ℓ) for H (given by the oriented meridians of the various components of L), and hence a partial ordering on H , defined by $h_1 \geq h_2$ if

$$h_1 - h_2 = \sum_i t_i \cdot m_i,$$

where all the t_i are non-negative integers. The above theorem is proved by constructing a chain complex $\widehat{\text{CFL}}(\vec{L})$ (cf Definition 4.9) which admits both a \mathbb{Z} -grading and an \mathbb{H} -filtration, i.e. the group underlying this chain complex splits as a group

$$\widehat{\text{CFL}}(\vec{L}) \cong \bigoplus_{d \in \mathbb{Z}, h \in \mathbb{H}} \widehat{\text{CFL}}_d(\vec{L}, h),$$

and its differential $\widehat{\partial}: \widehat{CFL}(\vec{L}) \rightarrow \widehat{CFL}(\vec{L})$ carries $\widehat{CFL}_d(\vec{L}, h)$ into

$$\bigoplus_{h' \leq h} \widehat{CFL}_{d-1}(\vec{L}, h').$$

Moreover, there is an induced differential

$$\widehat{\partial}: \widehat{CFL}_d(\vec{L}, h) \rightarrow \widehat{CFL}_{d-1}(\vec{L}, h),$$

obtained by post-composing the differential $\widehat{\partial}$ with the projection map

$$\bigoplus_{h' \leq h} \widehat{CFL}(\vec{L}, h') \rightarrow \widehat{CFL}(\vec{L}, h).$$

The homology of this graded object is identified with $\widehat{HFL}_d(\vec{L}, h)$. Moreover, the filtered chain homotopy type of $\widehat{CFL}(\vec{L})$ is a link invariant, and its total homology is identified with Λ^*V .

Succinctly, this equips $\widehat{HFL}(\vec{L})$ with an additional differential mapping $\widehat{HFL}_d(\vec{L}, h)$ into $\bigoplus_{h' < h} \widehat{HFL}_{d-1}(\vec{L}, h')$, in such a way that the total homology is Λ^*V . Moreover, the absolute gradings on $\widehat{HFL}(\vec{L})$ are fixed so that the top-dimensional class in Λ^*V is supported in degree zero.

In [18] we showed that the knot Floer homology of an alternating knot is determined explicitly in terms of the Alexander polynomial and signature of the knot. Using this result, together with Equation (1) and Theorem 1.1, we obtain the following:

Theorem 1.3 *Let L be an $\ell > 1$ -component oriented link with connected, alternating projection. Letting $\Delta_L(T_1, \dots, T_\ell)$ denote the multi-variable Alexander polynomial of L , write*

$$\left(\prod_{i=1}^{\ell} \left(T_i^{\frac{1}{2}} - T_i^{-\frac{1}{2}} \right) \right) \cdot \Delta_L(T_1, \dots, T_\ell) = \sum_{h \in \mathbb{H}} a_h \cdot e^h.$$

Then,

$$\widehat{HFL}(\vec{L}, h) = \mathbb{F}^{|a_h|}_{\binom{o(h) + \frac{\sigma - \ell + 1}{2}}{d}}$$

where here σ denotes the signature of the oriented link L , and $\mathbb{F}^n_{\binom{n}{d}}$ denotes the n -dimensional graded \mathbb{F} -vector space supported entirely in grading d .

Link Floer homology satisfies a Künneth principle for connected sums. Specifically, let L_1 and L_2 be a pair of oriented links, and distinguish components $K_1 \in L_1$ and

$K_2 \in L_2$. Let $L_1 \# L_2$ denote the link obtained from the disjoint union of L_1 and L_2 , via a connected sum joining K_1 and K_2 . There is a natural map

$$\mathbb{H}(L_1) \oplus \mathbb{H}(L_2) \longrightarrow \mathbb{H}(L_1 \# L_2),$$

written $h_1, h_2 \mapsto h_1 \# h_2$. This is the map which sends both the meridian for K_1 and the meridian for K_2 (in L_1 and L_2) to the meridian for the connected sum $K_1 \# K_2$.

Theorem 1.4 *There is an isomorphism:*

$$\widehat{\text{HFL}}_*(\vec{L}_1 \# \vec{L}_2, h) \cong \sum_{\{h_1 \in \mathbb{H}(L_1), h_2 \in \mathbb{H}(L_2) \mid h_1 \# h_2 = h\}} \widehat{\text{HFL}}_*(\vec{L}_1, h_1) \otimes \widehat{\text{HFL}}_*(\vec{L}_2, h_2)$$

The above theorem is proved in Section 11, along with some of its natural generalizations.

1.1 A further variant

There are several variations of link Floer homology. In Section 4, we construct a chain complex $\text{CFL}^-(\vec{L}, h)$, which is a \mathbb{Z} -graded and H -filtered chain complex of free modules over the ring $\mathbb{F}[U_1, \dots, U_\ell]$. Multiplication by U_i lowers homological degree by two and it lowers the filtration level by the basis element $m_i \in H$. As shown in Theorem 4.4, the filtered chain homotopy type of this complex is an invariant of the link. The relationship between this and the earlier construction is encoded in the fact that $\widehat{\text{CFL}}(\vec{L})$ is gotten from $\text{CFL}^-(\vec{L})$ by setting each $U_i = 0$.

The homology of the associated graded object is an oriented link invariant $\text{HFL}^-(\vec{L})$ which is a module over the ring $\mathbb{F}[U_1, \dots, U_\ell]$, endowed with a \mathbb{Z} -grading (inherited from a \mathbb{Z} -grading on $\text{CFL}^-(\vec{L})$) and an additional grading by elements of \mathbb{H} (induced from the filtration). We denote this multi-grading

$$\text{HFL}^-(\vec{L}) \cong \bigoplus_{d \in \mathbb{Z}, h \in \mathbb{H}} \text{HFL}_d^-(\vec{L}, h).$$

Calculating $\text{HFL}^-(\vec{L}, h)$ is more challenging than calculating $\widehat{\text{HFL}}(\vec{L}, h)$, as its differential counts more holomorphic disks.

The Euler characteristic in this case is given by the formula

$$(2) \quad \sum_{h \in \mathbb{H}} \chi(\text{HFL}_*^-(\vec{L}, h)) \cdot e^h \doteq \begin{cases} \Delta_L & \text{if } \ell > 1, \\ \frac{\Delta_L}{(1-T)} & \text{if } \ell = 1, \end{cases}$$

where here $f \doteq g$ means that two polynomials differ by multiplication by units. More succinctly, Equation (2) says that the Euler characteristic of $\text{HFL}^-(\vec{L}, h)$ is the Milnor torsion of $L \subset S^3$ [17].

The fact that $\widehat{\text{CFL}}(\vec{L}) = \text{CFL}^-(\vec{L})/\{U_i = 0\}_{i=1}^\ell$ as chain complexes has the following manifestation on the level of homology, proved in Section 7:

Theorem 1.5 *For each fixed $d \in \mathbb{Z}$ and $h \in \mathbb{H}$, the \mathbb{F} -module $\text{HFL}_d^-(\vec{L}, h)$ is the homology of a filtered chain complex whose E_1 term in dimension d is given by*

$$\bigoplus_{(a_1, \dots, a_\ell) \geq 0} U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \widehat{\text{HFL}}_{d+2a_1+\dots+2a_\ell}(\vec{L}, h + a_1 \cdot m_1 + \cdots + a_\ell \cdot m_\ell).$$

For HFL^- , we have the following analogue of Theorem 1.2, proved in Section 10:

Theorem 1.6 *There is a spectral sequence, which is a link invariant, whose E_1 term is $\text{HFL}^-(\vec{L})$, and whose E^∞ term is isomorphic to the $\mathbb{F}[U_1, \dots, U_\ell]$ -module $\mathbb{F}[U]$, where each U_i acts as multiplication by U .*

1.2 About the construction

Link Floer homology is constructed using suitable multiply-pointed Heegaard diagrams for links. More precisely, if L is a link with ℓ components, we consider a Heegaard decomposition of S^3 as $U_\alpha \cup U_\beta$, with the property that $L \cap U_\alpha$ and $L \cap U_\beta$ consists of ℓ unknotted arcs. The link L can now be encoded in a genus g Heegaard diagram for S^3 , with $g + \ell - 1$ attaching circles $\alpha_1, \dots, \alpha_{g+\ell-1}$ for the index one handles, and $g + \ell - 1$ attaching circles $\beta_1, \dots, \beta_{g+\ell-1}$ for the index two attaching circles, and also 2ℓ points $w_1, z_1, \dots, w_\ell, z_\ell$ where the link crosses the mid-level. $\widehat{\text{HFL}}$ is a variant of Lagrangian Floer homology in the $g + \ell - 1$ -fold symmetric product of Σ punctured in the basepoints w_i and z_i . These topological considerations lead naturally to the notion of *balanced* Heegaard diagrams, which are Heegaard diagrams for a (closed, oriented) three-manifold with ℓ zero- and ℓ three-handles. Theorems 1.2 and 1.6 are obtained from extending Heegaard Floer homology to the case of such balanced Heegaard diagrams.

This paper is organized as follows. In Section 2, we review some of the algebraic terms used throughout this paper. In Section 3, we discuss balanced Heegaard diagrams associated to links, and also the topological data which can be extracted from them. We also address admissibility issues which will be required to define the Heegaard Floer complexes. In Section 4, we describe the Heegaard Floer homology complexes associated to balanced Heegaard diagrams. In this case, the proof that $\partial^2 = 0$ is

slightly more subtle than the usual case considered in [21]: Specifically, it is now no longer true that the total count of *boundary degenerations*, disks with boundary lying entirely in \mathbb{T}_α (or \mathbb{T}_β), is zero. These issues are addressed in Section 5, where the analytical preliminaries are set up. For certain technical reasons, we find it also convenient to adopt the “cylindrical” approach to Heegaard Floer homology developed by Lipshitz [15], where one considers pseudo-holomorphic multi-section of the trivial Σ -bundle over a disk, rather than disks in the symmetric product of Σ .

With the help of this, in Section 6, it is established that $\partial^2 = 0$, and indeed, the Heegaard Floer homology for balanced Heegaard diagrams is identified with the usual Heegaard Floer homology.

With this background in place, the invariants for links are easy to construct, and their invariance properties are readily verified in Section 7. In particular, Theorems 1.2, 1.5, and 1.6 are quick consequences of the constructions.

With the link invariants in hand, we turn to some of their basic properties. In Section 8 we establish certain symmetry properties, which parallel the usual symmetry of the Alexander polynomial.

In Section 9, we turn to the Euler characteristic statements, verifying Equations (1) and (2).

In Section 10 we relate the present form of link homology with the earlier form derived from knot Floer homology in [20], establishing Theorem 1.1.

The Künneth principle for connected sums (Theorem 1.4 above, and also some more general statements for CFL^-) is established in Section 11.

In Section 12, we establish Theorem 1.3. We give also some principles which help computing the spectral sequence from Theorem 1.2. These principles allow one to determine the spectral sequence for all two-bridge links from the signature and the multi-variable Alexander polynomial.

We illustrate these principles in some particular examples, giving also some calculations for the two non-alternating, seven-crossing links, as well.

1.3 Further remarks

We have set up here link Floer homology $\widehat{\text{HFL}}$ as the homology of a graded object associated to a filtration of a chain complex for $\#^{\ell-1}(S^2 \times S^1)$. As such, it gets the extra differentials promised in Theorem 1.2. If one is not interested in this extra structure, but only $\widehat{\text{HFL}}$ as a graded group, then its construction is somewhat more

elementary than the constructions described here. Properties of this invariant, and further computations, are given in [22].

To some degree, link Floer homology can be viewed as a categorification of the multi-variable Alexander polynomial. It is interesting to compare this with the recent categorification of the HOMFLY polynomial, Khovanov and Rozansky [10], see also Khovanov [9], Khovanov and Rozansky [11] and Dunfield, Gukov and Rasmussen [2].

Acknowledgements

ZSz was supported by NSF grant number DMS 0406155. PSO was supported by NSF grant number DMS 0234311. We would like to thank Robert Lipshitz, Jacob Rasmussen, and András Stipsicz for useful conversations during the course of this work. We would also like to thank Jiajun Wang for his input on an early version of this manuscript.

2 Algebraic preliminaries

We begin by fixing some terminology from homological algebra which will be used throughout this paper. We give the set \mathbb{Z}^ℓ the partial ordering $a = (a_1, \dots, a_\ell) \leq b = (b_1, \dots, b_\ell)$ if each $a_i \leq b_i$. Let \mathfrak{S} be an affine space for \mathbb{Z}^ℓ . Natural examples include the affine space \mathbb{H} introduced in the introduction, where we think of \mathbb{Z}^ℓ as identified with the first homology of a link complement (where the identification induced by orientations on the link). Another is the set of relative Spin^c structures over the link complement, cf Section 3.2 below. The ordering on \mathbb{Z}^ℓ induces an ordering on \mathfrak{S} . An \mathfrak{S} -filtered module over \mathbb{F} is an R -module M equipped with an exhausting family of sub-modules $\mathcal{F}(M, a) \subseteq M$ indexed by $a \in \mathfrak{S}$, with the containment relation $\mathcal{F}(M, a) \subseteq \mathcal{F}(M, b)$ if $a \leq b$. A module homomorphism $\phi: M \rightarrow M'$ with the property that for all $a \in \mathbb{Z}^\ell$, $\phi(\mathcal{F}(M, a)) \subseteq \mathcal{F}(M', a)$ is called a *morphism of \mathfrak{S} -filtered modules*. A *filtered \mathfrak{S} -complex* is a chain complex for which the differential ∂ is a morphism of \mathfrak{S} -filtered modules. The homology of an \mathfrak{S} -filtered chain complex inherits a natural \mathbb{Z}^ℓ filtration.

Two \mathfrak{S} -filtered chain maps $\phi_1, \phi_2: A \rightarrow B$ between \mathfrak{S} -filtered chain complexes A and B are said to be *filtered chain homotopic* if there is a morphism of \mathfrak{S} -filtered modules $H: A \rightarrow B$ with $\partial \circ H - H \circ \partial = \phi_1 - \phi_2$. Two \mathfrak{S} -filtered chain complexes A and B are *\mathfrak{S} -filtered chain homotopy equivalent* if there are filtered chain maps $f: A \rightarrow B$ and $g: B \rightarrow A$ with the property that both $f \circ g$ and $g \circ f$ are filtered chain homotopic to the corresponding identity maps. If C and C' are filtered chain homotopy equivalent, we write $C \simeq C'$. Clearly, this forms an equivalence relation on

the set of \mathfrak{S} filtered chain complexes, and the induced equivalence class of a given \mathfrak{S} -filtered chain complex is called its \mathfrak{S} -filtered chain homotopy type.

Given a \mathfrak{S} -filtered complex, we can form the associated graded object

$$\text{gr}(C) = \bigoplus_{a \in \mathfrak{S}} \text{gr}(C, a),$$

where

$$\text{gr}(C, a) = \text{Coker} \left(\bigoplus_{b < a} \mathcal{F}(C, b) \longrightarrow \mathcal{F}(C, a) \right).$$

Clearly the homology of the associated graded object of a \mathbb{Z}^ℓ -filtered chain complex C depends on only the filtered chain homotopy type of C .

Given any set $\mathfrak{T} \subset \mathfrak{S}$ with the property that for all $a \in \mathfrak{T}$ if $b \leq a$, then $b \in \mathfrak{T}$, we can form the subcomplex $C(\mathfrak{T}) \subset C$.

We will consider modules over the ring $R = \mathbb{F}[U_1, \dots, U_\ell]$. An \mathfrak{S} filtered R -module is an R -module whose underlying \mathbb{F} -module (gotten by forgetting the action of U_i) is \mathfrak{S} -filtered, and has the additional property that

$$U^{a_1} \cdot \dots \cdot U^{a_\ell} \cdot \mathcal{F}(M, b) \subseteq \mathcal{F}(M, b - a),$$

where $a = (a_1, \dots, a_\ell)$. The notions of morphisms, chain complexes, homotopies, and homotopy type extend in a straightforward manner: we consider maps which are simultaneously \mathfrak{S} filtered and which are also R -modules. If C is a chain \mathfrak{S} -filtered chain complex of R -modules, the chain complex $\widehat{C} = C \otimes_{\mathbb{F}[U_1, \dots, U_\ell]} \mathbb{F}$ gotten by setting each $U_i = 0$ is also a \mathbb{Z}^ℓ -filtered chain complex (whose filtered chain homotopy type depends on C only up to its filtered chain homotopy type).

A free \mathbb{Z}^n -filtered chain complex of R -modules is one which admits a homogeneous generating which freely generates the underlying complex over $\mathbb{F}[U_1, \dots, U_\ell]$. In particular, as a \mathbb{F} -module, C splits as a direct sum

$$C = \bigoplus_{a \in \mathfrak{S}} C\{a\}.$$

2.1 Operations on \mathfrak{S} -filtered chain complexes.

If C is a \mathbb{Z}^n -filtered chain complex, and $a \in \mathbb{Z}^n$, then we can form the \mathbb{Z}^n filtered complex $C[a]$ whose underlying chain complex agrees with C , but whose \mathbb{Z}^n -filtration is shifted by a ; ie the filtration $\mathcal{F}(C[a], b) = \mathcal{F}(C, a + b)$.

If C is free \mathbb{Z}^n -filtered complex, and $i \in \{1, \dots, n\}$, we can split the differential into components $D_a^b: C\{a\} \rightarrow C\{b\}$ with $b \leq a$; we can form a \mathbb{Z}^{n-1} filtered chain complex $C^{(1)}$ by “taking homology in the first component.” Specifically, write $D = D^1 + D'$, where D^1 consists of all the components D_a^b where a and b agree on all but the first place. It is easy to see that D^1 is a differential, and hence we can form the \mathbb{Z}^{n-1} -filtered chain complex

$$\mathcal{F}(C^{(1)}, b) = H_* \left(\bigcup_{\{a \mid (a_2, \dots, a_\ell) \leq b\}} \mathcal{F}(C, a), D^1 \right)$$

endowed with the differential induced from D' . Note that there remains an extra action of U_1 on this chain complex (which does not change the filtration level).

Of course, this notion admits a straightforward adaptation to taking the homology in the i^{th} components for any $i \in \{1, \dots, \ell\}$.

3 Heegaard diagrams

We discuss here basic topological aspects of multiply-pointed Heegaard diagrams, which are relevant for the study of links in three-manifolds. The material here is mostly a straightforward generalization of the singly-pointed case, which was studied in [21], and the doubly-pointed case from [20].

3.1 Heegaard diagrams for three-manifolds

Definition 3.1 A *balanced ℓ -pointed Heegaard diagram* is a quadruple of data

$$(\Sigma, \boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_{g+\ell-1}\}, \boldsymbol{\beta} = \{\beta_1, \dots, \beta_{g+\ell-1}\}, \mathbf{w} = \{w_1, \dots, w_\ell\}),$$

where ℓ is a positive integer, Σ is an oriented surface of genus g ,

$$\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_{g+\ell-1}\}$$

is a $g + \ell - 1$ -tuple of disjoint, simple closed curves which span a g -dimensional sublattice of $H_1(\Sigma; \mathbb{Z})$ (and hence they specify a handlebody U_α which is bounded by Σ), $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_{g+\ell-1}\}$ is a $g + \ell - 1$ -tuple of disjoint, simple closed curves which span another g -dimensional sublattice of $H_1(\Sigma; \mathbb{Z})$ (specifying another handlebody U_β), and \mathbf{w} is a collection of points in Σ chosen as follows. Let $\{A_i\}_{i=1}^\ell$ denote the connected components of $\Sigma - \alpha_1 - \dots - \alpha_{g+\ell-1}$; let $\{B_i\}_{i=1}^\ell$ denote the connected components of $\Sigma - \beta_1 - \dots - \beta_{g+\ell-1}$. The points $w_i \in \Sigma$ are constrained so that $w_i \in A_i \cap B_i$.

A balanced ℓ -pointed Heegaard diagram specifies a closed, oriented three-manifold $Y = U_\alpha \cup_\Sigma U_\beta$, endowing it with a cellular decomposition whose zero-cells correspond to the components $\{A_i\}_{i=1}^\ell$, its one-cells correspond to the circles $\alpha_1, \dots, \alpha_{g+\ell-1}$, its two-cells correspond to $\beta_1, \dots, \beta_{g+\ell-1}$, and its three-cells correspond to the $\{B_i\}_{i=1}^\ell$.

Definition 3.2 A balanced ℓ -pointed Heegaard diagram is called *generic* if the circles α_i and β_j meet transversally, for all $i, j \in \{1, \dots, g + \ell - 1\}$.

Fix a connected, oriented three-manifold Y , and a generic self-indexing Morse function on Y which has the same number ℓ of index zero and three critical points. Fix also generic metric g , together with a choice of ℓ gradient flowlines connecting each of the index zero and three critical points. Then, there is an associated generic balanced ℓ -pointed Heegaard diagram for Y whose surface Σ is the mid-level of the Morse function; α_i is the locus of points on Σ where the gradient flowlines leaving the i^{th} index one critical point meets Σ ; similarly, β_i is the locus of points on Σ which flow into the i^{th} index two critical point. Finally, for $i = 1, \dots, \ell$, w_i is the point on Σ which lies on the distinguished gradient flow-line connecting the i^{th} index zero and index three critical point. If $(\Sigma, \alpha, \beta, \mathbf{w})$ is obtained in this manner from a Morse function f , we call f a Morse function *compatible* with the balanced Heegaard diagram for Y .

Given a generic ℓ -pointed balanced Heegaard diagram for Y , it is easy to construct a compatible Morse function f .

Proposition 3.3 Any two generic balanced ℓ -pointed Heegaard diagrams for Y can be connected by a sequence of the following moves:

- (i) *Isotopies and handleslides of the α supported in the complement of \mathbf{w} ;*
- (ii) *Isotopies and handleslides of the β supported in the complement of \mathbf{w} ;*
- (iii) *Index one/two stabilizations (and their inverses): Forming the connected sum of $(\Sigma, \alpha, \beta, \mathbf{w})$ with a torus equipped with a new pair of curves α_g and β_g which meet transversally in a single point;*
- (iv) *Index zero/three stabilizations (and their inverses): Introducing a new pair of homotopic curves $\alpha_{g+\ell}$ (disjoint from the α_i for $1 \leq i \leq g + \ell - 1$) and $\beta_{g+\ell}$ (disjoint from the other β_i) and a new basepoint $w_{\ell+1}$ in such a manner that $\alpha_{g+\ell}$ and $\beta_{g+\ell}$ are homotopic in $\Sigma - w_1 - \dots - w_\ell - w_{\ell+1}$, and each component of $\Sigma - \alpha_1 - \dots - \alpha_{g+\ell}$ and $\Sigma - \beta_1 - \dots - \beta_{g+\ell}$ contains some w_i .*

Proof It follows from standard Morse theory that any two (unpointed) Heegaard diagrams can be connected by moves of Types (i)–(iii). The fact that this can be done in the complement of a single basepoint (the case $\ell = 1$) can be established by trading an isotopy across the basepoint for a sequence of handleslides in the opposite direction [21, Proposition 7.1].

The proof in general is established by showing that we can use the above moves to reduce the number of basepoints, and hence reducing to the case of a singly-pointed Heegaard diagram. This is done as follows. Let A be the two-chain with boundary a combination of $\alpha_1, \dots, \alpha_{g+\ell}$ which has $n_{w_{\ell+1}}(A) = 1$ and $n_{w_i}(A) = 0$ for all $i \leq \ell$. After a sequence of handleslides among the α , we can arrange for A to have only one boundary component, which we label $\alpha_{g+\ell}$. Indeed, another sequence of handleslides can be done to arrange furthermore for the genus of A to be zero. Let B be the corresponding two-chain with boundary amongst the $\beta_1, \dots, \beta_{g+\ell}$, with $n_{w_{\ell+1}}(B) = 1$ and $n_{w_i}(B) = 0$ for all $i \leq \ell$. Performing a sequence of handleslides among the β , we can reduce to the case where B is a disk bounded by $\beta_{g+\ell}$. We can now form an index zero/three de-stabilization to delete $\alpha_{g+\ell}$, $\beta_{g+\ell}$ and $w_{g+\ell}$. The proof then follows by induction. \square

3.2 Relative Spin^c structures

We pause our discussion on Heegaard diagrams to recall Turaev's interpretation of Spin^c structures on three-manifolds, see [28], compare also [12].

Let Y be a closed, oriented three-manifold. We say that two nowhere vanishing vector fields v and v' are *homologous* if there is a ball $B \subset Y$ with the property that v and v' are homotopic (through nowhere vanishing vector fields) on the complement of B . The set of equivalence classes of such vector fields can be naturally identified with the space $\text{Spin}^c(Y)$ of Spin^c structures over Y . In particular it is an affine space for $H^2(Y; \mathbb{Z})$.

This notion has a straightforward generalization to the case of three-manifolds with toroidal boundary. Specifically, let $(M, \partial M)$ be a three-manifold with boundary consisting of a disjoint union of tori $T_1 \cup \dots \cup T_\ell$. The tangent bundle to the two-torus has a canonical nowhere vanishing vector field, which is unique up to homotopy (through nowhere vanishing vector fields). Consider now nowhere vector fields v on Y whose restriction to ∂M are identified with the canonical nowhere vanishing vector field on the boundary tori (in particular, the vector field has no normal component at the boundary). Two such vector fields v and v' are declared homologous if there is a ball $B \subset M - \partial M$ with the property that the restrictions of v and v' to $M - B$ are homotopic. The set of homology classes of such vector fields is called the set of

relative Spin^c structures, and it is an affine space for $H^2(M, \partial M; \mathbb{Z})$. We denote this set by $\underline{\text{Spin}}^c(M, \partial M)$.

Multiplying vector fields by -1 induces an involution on the space of relative Spin^c structures

$$J: \underline{\text{Spin}}^c(M, \partial M) \longrightarrow \underline{\text{Spin}}^c(M, \partial M).$$

If \vec{v} is a vector field on a three-manifold M with toroidal boundaries, whose restriction to each bounding torus gives the canonical trivialization of the torus' tangent bundle, then we can consider the oriented two-plane field \vec{v}^\perp of vectors orthogonal to \vec{v} . Along ∂M , this two-plane field has a canonical trivialization, by outward pointing vectors. Hence, there is a well-defined notion of a relative Chern class of this line field relative to its trivialization, thought of as an element of $H^2(M, \partial M; \mathbb{Z})$. This descends to a well-defined assignment

$$c_1: \underline{\text{Spin}}^c(M, \partial M) \longrightarrow H^2(M, \partial M; \mathbb{Z}).$$

Clearly,

$$c_1(J \cdot \underline{s}) = -c_1(\underline{s}).$$

The reader familiar with [20] should be warned that in that case, we were considering null-homologous knots, rather than links, and hence there is a well-defined notion of a zero-surgery. In [20], relative Spin^c structures were thought of as absolute Spin^c structures on this zero-surgery. This is a slightly different point of view than the one taken here (where we no longer have the luxury of referring to a zero-surgery).

3.3 Intersection points and Spin^c structures

Fix a generic ℓ -pointed balanced Heegaard diagram for Y , and let f be a compatible Morse function.

Consider the $g + \ell - 1$ -fold symmetric product of Σ , $\text{Sym}^{g+\ell-1}(\Sigma)$, and let

$$\mathbb{T}_\alpha = \alpha_1 \times \dots \times \alpha_{g+\ell-1} \text{ and } \mathbb{T}_\beta = \beta_1 \times \dots \times \beta_{g+\ell-1}.$$

Clearly, an intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ corresponds to a $g + \ell - 1$ -tuple of gradient flow-lines which connect all the index one and two critical points.

Let $\gamma_{\mathbf{x}}$ be the union of gradient flowlines passing through each $x_i \in \mathbf{x}$, and $\gamma_{\mathbf{w}}$ be the union of gradient flowlines passing through each $w_i \in \mathbf{w}$. The closure of $\gamma_{\mathbf{x}} \cup \gamma_{\mathbf{w}}$ is a collection of arcs whose boundaries consist of all the critical points of f . Moreover, each component contains a pair of critical points whose indices have opposite parities. Thus, we can modify the gradient vector field in an arbitrarily small neighborhood of $\gamma_{\mathbf{x}} \cup \gamma_{\mathbf{w}}$ to obtain a new vector field which vanishes nowhere in Y . Taking the homology

class of this vector field in the sense of Turaev [28], we obtain a map from intersection points to Spin^c structures over Y

$$\mathfrak{s}_{\mathbf{w}}: \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta} \longrightarrow \text{Spin}^c(Y).$$

It is easy to see that this is well-defined, i.e. independent of the choice of compatible Morse function f and modification of the vector field (compare [21, Section 2.6]).

3.4 Whitney disks and admissibility

Fix a generic ℓ -pointed Heegaard diagram for Y , $(\Sigma, \alpha, \beta, \mathbf{w})$. Given $\mathbf{x}, \mathbf{y} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$, we can consider Whitney disks from \mathbf{x} to \mathbf{y} relative to \mathbb{T}_{α} and \mathbb{T}_{β} . Each such Whitney disk gives rise to a two-chain on Σ ; specifically,

$$\Sigma - \alpha_1 - \cdots - \alpha_{g+\ell-1} - \beta_1 - \cdots - \beta_{g+\ell-1}$$

consists of a collection of regions $\{\Omega_i\}_{i=1}^m$. Fix a reference point $p_i \in \Omega_i$, and let $\mathcal{D}(\phi)$ denote the two-chain

$$\sum_{i=1}^m n_{p_i}(\phi)[\Omega_i],$$

where here $n_p(\phi)$ denotes the algebraic intersection number of ϕ with the subvariety $\{p\} \times \text{Sym}^{g+\ell-2}(\Sigma)$. The element $\mathcal{D}(\phi)$ specifies the relative homology class induced from the Whitney disk ϕ .

Let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the space of homology classes of Whitney disks. Let

$$n_{\mathbf{w}}: \pi_2(\mathbf{x}, \mathbf{y}) \longrightarrow \mathbb{Z}^{\ell}$$

be the map which sends ϕ to $(n_{w_1}(\phi), \dots, n_{w_{\ell}}(\phi))$

Let $\pi_2(\alpha)$ denote the space of homology classes of disks with boundary in \mathbb{T}_{α} . Clearly, we have isomorphisms

$$(3) \quad n_{\mathbf{w}}: \pi_2(\alpha) \xrightarrow{\cong} \mathbb{Z}^{\ell} \quad \text{and} \quad n_{\mathbf{w}}: \pi_2(\beta) \xrightarrow{\cong} \mathbb{Z}^{\ell}.$$

There is an exact sequence

$$(4) \quad 0 \longrightarrow \mathbb{Z} \longrightarrow \pi_2(\alpha) \oplus \pi_2(\beta) \longrightarrow \pi_2(\mathbf{x}, \mathbf{y}) \longrightarrow H^1(Y; \mathbb{Z}) \longrightarrow 0.$$

Moreover, we have an exact sequence

$$(5) \quad 0 \longrightarrow \mathcal{P} \longrightarrow \pi_2(\mathbf{x}, \mathbf{y}) \xrightarrow{n_{\mathbf{w}}} \mathbb{Z}^{\ell} \longrightarrow 0.$$

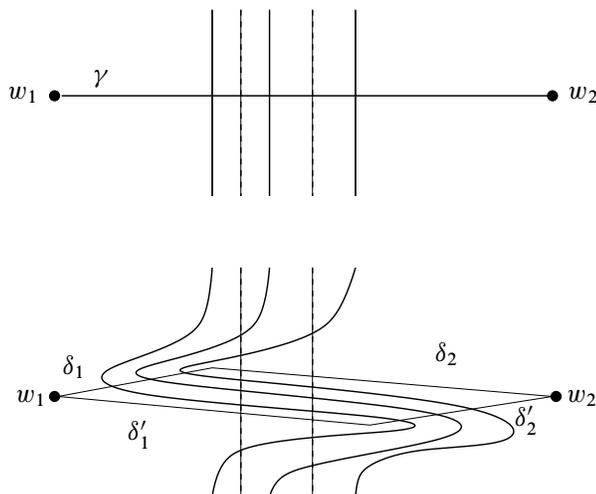


Figure 1: **Winding to achieve admissibility** This is an illustration of the procedure described in Proposition 3.6. Here, the dotted lines denote the possible α curves, and the solid lines represent β curves.

Definition 3.4 The group \mathcal{P} is called the group of *periodic domains*.

In the case where $H^1(Y; \mathbb{Z}) = 0$, we have that $\mathcal{P} \cong \mathbb{Z}^{\ell-1}$.

As in Section 3.4 of [21], we need further restrictions on the Heegaard diagram to obtain a reasonable chain complex (whose homology is the Heegaard Floer homology) defined.

Definition 3.5 A generic, balanced ℓ -pointed Heegaard diagram is called *weakly admissible* if for any non-trivial homology class $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ with $n_{\mathbf{w}}(\phi) = 0$ (meaning that $n_{w_i}(\phi) = 0$ for each $w_i \in \mathbf{w}$), the domain $\mathcal{D}(\phi)$ has both positive and negative local multiplicities.

Proposition 3.6 Let Y be a three-manifold with $H^1(Y; \mathbb{Z}) = 0$. Any balanced ℓ -pointed Heegaard diagram for Y is isotopic to a weakly admissible balanced ℓ -pointed Heegaard diagram.

Proof We embed in Σ a tree Γ whose ℓ vertices are the points w_i . We claim that admissibility can be achieved by isotoping some of the β -curves in a regular neighborhood of Γ . Specifically, if γ is an arc connecting w_1 to w_2 in Γ , then perform an isotopy of the β circles in a regular neighborhood of γ in such a manner that there is a pair of arcs δ_1 and δ_2 so that $\delta_1 \cup \delta_2$ is isotopic to γ as an arc from

w_1 to w_2 , but δ_1 is disjoint from the β circles while δ_2 is disjoint from the α circles. Moreover, we find another pair of arcs δ'_1 and δ'_2 so that $\delta'_1 \cup \delta'_2$ is isotopic to γ , only now δ'_1 is disjoint from the α circles while δ'_2 is disjoint from the β circles (see Figure 1 for an illustration). Isotoping the β circles in a regular neighborhood of all the edges in Γ as above, we obtain a Heegaard diagram that we claim is weakly admissible.

According to Equation (4), any $P \in \mathcal{P}$ can be decomposed as $P = A + B$, with $A \in \pi_2(\alpha)$ and $B \in \pi_2(\beta)$. The condition that P is a periodic domain ensures that $n_{\mathbf{w}}(A) + n_{\mathbf{w}}(B) = 0$. According to Equation (3), P is uniquely determined by $n_{\mathbf{w}}(A)$, modulo addition of Σ .

Suppose now that P has the property that the oriented intersection number of ∂A with γ is non-zero. Then, at the intermediate endpoint of δ_1 , we see that $A + B$ has local multiplicity given by $\partial A \cap \gamma$, while at the intermediate endpoint of δ'_1 , $A + B$ has local multiplicity given by $\partial B \cap \gamma = -\partial A \cap \gamma$. Thus, if for some edge γ in Γ , $\partial A \cap \gamma \neq 0$, then $P = A + B$ has both positive and negative coefficients. However, if $P = A + B$, where $A \in \pi_2(\alpha)$ and $B \in \pi_2(\beta)$, and ∂A has algebraic intersection number equal to zero with each edge in Γ , then, after subtracting off some number ($n_{w_1}(A)$) of copies of Σ , we can write $P = A' + B'$, where $A' \in \pi_2(\alpha)$, $B' \in \pi_2(\beta)$ and $n_{w_i}(A) = n_{w_i}(B) = 0$. According to Equation (3), then $A' = 0$, and hence $P = 0$. \square

3.5 Heegaard diagrams and links

Fix an ℓ -pointed Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w})$ for a three-manifold Y , and choose also an additional ℓ -tuple of basepoints $\mathbf{z} = \{z_1, \dots, z_\ell\}$, with the property that for each $i = 1, \dots, \ell$, both w_i and z_i are contained in the same component

$$A_i \subset \Sigma - \alpha_1 - \dots - \alpha_{g+\ell-1} \quad \text{and} \quad B_i \subset \Sigma - \beta_1 - \dots - \beta_{g+\ell-1}.$$

This data gives rise to an oriented, ℓ -component link L in $Y = U_\alpha \cup_\Sigma U_\beta$.

Definition 3.7 The diagram $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ as above is said to be a 2ℓ -pointed Heegaard diagram for the oriented link \vec{L} in Y .

Conversely, given an oriented, ℓ -component link, one can find a self-indexing Morse function $f: Y \rightarrow \mathbb{R}$ with ℓ index zero and three critical points, and $g + \ell - 1$ index one and two critical points, with the additional property that there are two ℓ -tuples of flowlines $\gamma_{\mathbf{w}}$ and $\gamma_{\mathbf{z}}$ connecting all the index three and index zero critical points, so that our oriented link L can be realized as the difference $\gamma_{\mathbf{z}} - \gamma_{\mathbf{w}}$. Such a Morse function gives rise to a 2ℓ -pointed Heegaard diagram for \vec{L} in Y .

Definition 3.8 A 2ℓ -pointed Heegaard diagram for a link $\vec{L} \subset Y$ is *weakly admissible* if the underlying ℓ -pointed Heegaard diagram for Y (gotten by disregarding the \mathbf{z}) is weakly admissible.

Proposition 3.9 If \vec{L} is an oriented ℓ -component link in Y , then there is a corresponding (weakly admissible) 2ℓ -pointed Heegaard diagram. Any two (weakly admissible) 2ℓ -pointed Heegaard diagrams for the same oriented link $\vec{L} \subset Y$ can be connected by a sequence of moves of the following types:

- Isotopies and handleslides of the α supported in the complement of \mathbf{w} and \mathbf{z} ;
- Isotopies and handleslides of the β supported in the complement of \mathbf{w} and \mathbf{z} ;
- Index one/two stabilizations (and their inverses), forming the connected sum of $(\Sigma, \alpha, \beta, \mathbf{w})$ with a torus equipped with a new pair of curves α_g and β_g which meet transversally in a single point.

Moreover, if we start and end with weakly admissible Heegaard diagrams, then we can assume that all the intermediate Heegaard diagrams are also weakly admissible.

Proof Without the admissibility hypothesis, the above result follows from Morse theory in the usual manner. Admissibility can be achieved as in the proof of Proposition 3.6. □

3.6 Intersection points and link diagrams

Given a 2ℓ -pointed Heegaard diagram for a link, consider the tori

$$\mathbb{T}_\alpha = \alpha_1 \times \dots \times \alpha_{g+\ell-1} \quad \text{and} \quad \mathbb{T}_\beta = \beta_1 \times \dots \times \beta_{g+\ell-1}.$$

Given a pair $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ of intersection points, we can find paths

$$a: [0, 1] \longrightarrow \mathbb{T}_\alpha \quad \text{and} \quad b: [0, 1] \longrightarrow \mathbb{T}_\beta,$$

with $\partial a = \partial b = \mathbf{x} - \mathbf{y}$. Viewing these paths as one-chains in Σ , supported away from the reference points $\{w_i\}_{i=1}^\ell$ and $\{z_i\}_{i=1}^\ell$, we obtain a one-cycle $\epsilon(\mathbf{x}, \mathbf{y})$ in the complement $Y - L$. This assignment clearly descends to give a well-defined map

$$\epsilon_{\mathbf{w}, \mathbf{z}}: (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \longrightarrow H_1(Y - L; \mathbb{Z}).$$

Lemma 3.10 *An oriented link in Y gives rise to a map*

$$\Pi: H_1(Y - L) \longrightarrow \mathbb{Z}^\ell,$$

where $\Pi_i(\gamma)$ is the linking number of γ with the component $L_i \subset L$, which is an isomorphism in the case where $H^1(Y; \mathbb{Z}) = 0$. Given an oriented link, and $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is any homology class, then

$$\Pi(\underline{\epsilon}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}, \mathbf{y})) = n_{\mathbf{z}}(\phi) - n_{\mathbf{w}}(\phi).$$

Proof The homology class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ gives rise to a null-homology of $\underline{\epsilon}(\mathbf{x}, \mathbf{y})$ inside Y . This null-homology meets the i^{th} component of the oriented link \vec{L} with intersection number $n_{\mathbf{z}}(\phi) - n_{\mathbf{w}}(\phi)$. The lemma follows at once. \square

We can lift $\underline{\epsilon}$ to a map from intersection points to relative Spin^c structures, generalizing the map from Section 3.3.

Let $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ be a pointed Heegaard diagram for an oriented link \vec{L} . We define the map

$$(6) \quad \underline{\epsilon}_{\mathbf{w}, \mathbf{z}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \longrightarrow \underline{\text{Spin}}^c(Y, \vec{L})$$

as follows. Note that we write $\underline{\text{Spin}}^c(Y, \vec{L})$ for the space of relative Spin^c structures on $Y - \text{nd}(L)$, cf Section 3.2.

For this map, we fix a choice as follows. Let γ be a gradient flowline connecting an index zero and index three critical point, and let $N(\gamma)$ denote a neighborhood of this flowline. One can construct a nowhere vanishing vector field \vec{v} over $N(\gamma)$, which has an integral flowline P which enters $N(\gamma)$ from its boundary, contains γ as a subset, and then exits $N(\gamma)$.

Let f be an orientation-preserving involution of $N(\gamma)$. We can arrange for $-\vec{v}|_{\partial N(\gamma)}$ to agree with $f^*(\vec{v})|_{\partial(N(\gamma) \cup P)}$. Indeed, we can construct \vec{v} in such a manner that the difference $\vec{v} - f^*(\vec{v})$ is the Poincaré dual of a meridian for γ , thought of as an element of $H_1(N(\gamma) - P)$.

This vector field is illustrated in Figure 2.

Armed with this vector field, we define the map promised in Equation (6). Fix a Morse function f compatible with the Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$. Given $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, consider the flowlines $\gamma_{\mathbf{x}}$, $\gamma_{\mathbf{w}}$, and $\gamma_{\mathbf{z}}$. We replace the gradient vector field in a neighborhood of $\gamma_{\mathbf{x}}$ so as not to vanish there. Similarly, we replace the gradient vector field in a neighborhood of $\gamma_{\mathbf{w}}$ using \vec{v} so that it does not vanish there. In fact, arranging for P to consist of arcs on $\gamma_{\mathbf{z}} \cup \gamma_{\mathbf{w}}$, we obtain in this manner a vector

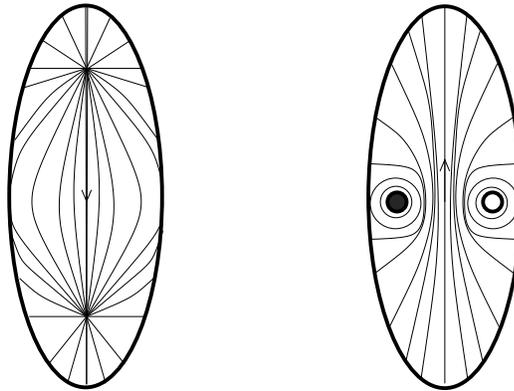


Figure 2: **Non-vanishing vector field** On the left, we represent the gradient vector field in a standard neighborhood of a gradient flowline γ connecting an index zero and three critical point, by sketching various integral curves. The vector field on the three-ball is obtained by rotating this vector field through the axis through the center of the picture. Replace this vector field with the nowhere vanishing vector field \vec{v} represented on the right. Specifically, consider the vector field represented here (after rotating through the axis in the center), which vanishes on the oriented meridian represented here by the pair of circles. Modify the vector field to point into the page at the empty circle and out of the page at the dark circle. This gives the nowhere vanishing vector field \vec{v} , with a closed orbit which coincides with an oriented meridian for $\gamma \subset P$.

field on Y which contains \vec{L} as a closed orbit. It is easy to see that this is equivalent to a vector field on $Y - \text{nd}(\vec{L})$ which is a standard non-vanishing vector field on the boundary tori.

Lemma 3.11 *We have that*

$$\underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x}) - \underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{y}) = \text{PD}[\underline{\epsilon}(\mathbf{x}, \mathbf{y})].$$

where here PD denotes the Poincaré duality map

$$\text{PD}: H_1(Y - L) \longrightarrow H^2(Y, L).$$

Indeed, given $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we have that

$$(7) \quad \underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x}) - \underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{y}) = \sum_{i=1}^{\ell} (n_{z_i}(\phi) - n_{w_i}(\phi)) \text{PD}[\mu_i],$$

where here μ_i is the meridian for the i^{th} component of L , with its induced orientation from the orientation of L .

Proof The vector fields $\underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x})$ and $\underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{y})$ differ in a neighborhood of $\gamma_x - \gamma_y$. It is now a local calculation to see that $\underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x}) - \underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{y}) = \text{PD}[\gamma_x - \gamma_y]$ (compare [21, Lemma 2.9] for the corresponding statement for Spin^c structures). It is easy to see that $\gamma_x - \gamma_y$ is homologous to $\underline{e}(\mathbf{x}, \mathbf{y})$.

The second remark follows immediately from Lemma 3.10. \square

The following fact will be useful in studying the symmetry properties of link Floer homology.

Lemma 3.12 *Let $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ be a multiply pointed Heegaard diagram representing an oriented link $\bar{L} \subset Y$, and let $\underline{s}_{\mathbf{w},\mathbf{z}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \underline{\text{Spin}}^c(Y, L)$ denote the induced map to relative Spin^c structures. Then, $(-\Sigma, \beta, \alpha, \mathbf{w}, \mathbf{z})$ specifies the same link, endowed with the opposite orientation. Let $\underline{s}'_{\mathbf{w},\mathbf{z}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \underline{\text{Spin}}^c(Y, L)$ denote its induced map. Then, we have that*

$$(8) \quad \underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x}) = J \circ \underline{s}'_{\mathbf{w},\mathbf{z}}(\mathbf{x}).$$

Also, $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ is another diagram representing the link with the opposite orientation; let $\underline{s}_{\mathbf{z},\mathbf{w}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \underline{\text{Spin}}^c(Y, L)$ denote its corresponding map. Then,

$$(9) \quad \underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x}) = \underline{s}_{\mathbf{z},\mathbf{w}}(\mathbf{x}) + \sum_{i=1}^{\ell} \text{PD}[\mu_i].$$

Proof For the first remark, note that if f is a Morse function compatible with $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$, then $-f$ is a Morse function compatible with $(-\Sigma, \beta, \alpha, \mathbf{w}, \mathbf{z})$. It is now a straightforward consequence of its definition that if $\underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x})$ is represented by \vec{v} , then $-\vec{v}$ represents $\underline{s}'_{\mathbf{w},\mathbf{z}}(\mathbf{x})$. Thus, Equation (8) follows.

For the second observation, note that the vector fields \vec{v} and \vec{w} representing $\underline{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x})$ and $\underline{s}_{\mathbf{z},\mathbf{w}}(\mathbf{x})$ respectively differ only in a collar neighborhood of the boundary. In fact, one can see that in this neighborhood, \vec{w} and \vec{v} can be made isotopic away from a neighborhood of the meridian, where they point in opposite directions. In this way, Equation (9) follows. \square

3.7 Filling relative Spin^c structures

Let $\vec{L} \subset Y$ be an oriented link. Fix a component K_1 of the underlying link. We have a natural “filling map”

$$G_{K_1}: \underline{\text{Spin}}^c(Y, \vec{L}) \longrightarrow \underline{\text{Spin}}^c(Y, \vec{L} - K_1).$$

This is gotten by simply viewing the relative Spin^c structure on Y relative to \vec{L} as one over Y relative to $\vec{L} - K_1$. More specifically, if we think of $\underline{\text{Spin}}^c(Y, \vec{L})$ as generated by vector fields which have a closed orbits consisting of the components of L (traversed with their given orientations), then we can view these also as relative Spin^c structures which have a closed orbits consisting of the components of $L - K_1$ (traversed with their given orientations).

Lemma 3.13 *We have the following*

$$c_1(G_{K_1}(\underline{s})) = c_1(\underline{s}) + \text{PD}[K_1].$$

Proof To verify this, let F_2 be an oriented surface in $Y - \text{nd}(L - K_1)$ with boundary on $\partial \text{nd}(L - K_1)$, representing some relative homology class $H_2(Y, L - K_1)$. We can assume that F_2 meets K_1 transversally. We can consider the surface F'_2 in $Y - \text{nd}(L)$ with boundary on $\text{nd}(L)$ gotten by deleting $\text{nd}(K_1) \cap F_2$ from F_2 . The homology class of F'_2 represents the natural map $H_2(Y, L - K_1) \longrightarrow H_2(Y, L)$.

If $\#(K_1 \cap F_2)$ denotes the algebraic intersection number of F_2 with K_1 (endowed with the orientation it inherits from \vec{L}) we have that

$$\langle c_1(G_{K_1}(\underline{s})), [F_2] \rangle = \langle c_1(\underline{s}), [F'_2] \rangle + \#(K_1 \cap F_2),$$

which is equivalent to the claim in the lemma. □

4 Definition of Heegaard Floer homology for multi-pointed Heegaard diagrams

4.1 Heegaard diagrams for three-manifolds

For convenience, we work always with Floer homology with coefficients in $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$. We also consider the case where the ambient manifold is a rational homology three-sphere, and hence Proposition 3.6 applies.

Let $(\Sigma, \alpha, \beta, \mathbf{w})$ be an ℓ -pointed balanced weakly admissible Heegaard diagram for a rational homology three-sphere Y . We define $CF^-(\Sigma, \alpha, \beta, \mathbf{w})$ to be the free

module over the polynomial algebra $\mathbb{F}[U_1, \dots, U_\ell]$ generated by the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ inside $\text{Sym}^{g+\ell-1}(\Sigma)$. There is a module homomorphism

$$\partial^-: CF^-(\Sigma, \alpha, \beta, \mathbf{w}) \longrightarrow CF^-(\Sigma, \alpha, \beta, \mathbf{w})$$

defined by

$$(10) \quad \partial^- \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi) = 1\}} \# \widehat{\mathcal{M}}(\phi) \cdot U_1^{n_{w_1}(\phi)} \cdots U_\ell^{n_{w_\ell}(\phi)} \otimes \mathbf{y}.$$

Here, as usual, $\mathcal{M}(\phi)$ denotes the moduli space of pseudo-holomorphic representatives of the given homology class of Whitney disks, and $\widehat{\mathcal{M}}(\phi)$ denotes the quotient of this moduli space by the action of \mathbb{R} . Also, $\mu(\phi)$ denotes the expected dimension of $\mathcal{M}(\phi)$ (i.e. the Maslov index of ϕ). (Indeed, we will find it useful later to use Lipshitz's cylindrical formulation [15] instead; we return to this point in Section 5.)

We will establish the following in Section 5.3:

Proposition 4.1 *For an ℓ -pointed, balanced, Heegaard diagram for a rational homology three-sphere Y , for any $\phi, \phi' \in \pi_2(\mathbf{x}, \mathbf{y})$, we have that*

$$(11) \quad \mu(\phi) - \mu(\phi') = 2 \sum_{i=1}^{\ell} (n_{w_i}(\phi) - n_{w_i}(\phi')).$$

Lemma 4.2 *For an ℓ -pointed balanced weakly-admissible Heegaard diagram for a three-manifold Y and any $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, the right hand side of Equation (10) consists of only finitely many different non-zero terms.*

Proof First, we prove that the coefficient of $U_1^{a_1} \cdots U_\ell^{a_\ell}$ in \mathbf{y} is finite, for any given $\mathbf{a} = a_1 \times \dots \times a_\ell \in \mathbb{Z}^\ell$. But this follows readily from the fact that for any two homology classes $\phi, \phi' \in \pi_2(\mathbf{x}, \mathbf{y})$ with $n_{\mathbf{w}}(\phi) = n_{\mathbf{w}}(\phi')$, $\phi - \phi'$ is a periodic domain, which hence must have both positive and negative coefficients. It follows from this that there are at most finitely many homology classes ϕ with $n_{\mathbf{w}}(\phi) = \mathbf{a}$ and $\mathcal{D}(\phi) \geq 0$, and hence which can support holomorphic representatives.

If ϕ admits a holomorphic representative, then $n_{w_i}(\phi) \geq 0$ for all i ; from Proposition 4.1, it follows at once that for $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = 1$, the quantity $\sum_i n_{w_i}(\phi)$ depends only on \mathbf{x} and \mathbf{y} . Thus, given \mathbf{x} and \mathbf{y} it follows that there are only finitely many different possibilities for $U_1^{a_1} \cdots U_\ell^{a_\ell}$ with non-zero coefficient on the right hand side of Equation (10). \square

Lemma 4.3 *The map ∂^- is a differential on $CF^-(\Sigma, \alpha, \beta, \mathbf{w})$.*

The above lemma is proved in Section 6, where we also prove that the homology groups are identified with the usual Heegaard Floer homology of [21]:

Theorem 4.4 *The chain homotopy type of the complex $CF^-(\Sigma, \alpha, \beta, \mathbf{w})$ over the polynomial ring $\mathbb{F}[U_1]$ is a three-manifold invariant; indeed, its homology coincides with the Heegaard Floer homology group $HF^-(Y)$.*

A more precise version of the above theorem can be stated, which respects the splitting of $HF^-(Y)$ according to its various Spin^c structures. We do not belabour this point now, as the primary example we have in mind here is the case where $Y = S^3$ (which has a unique Spin^c structure).

There is a simpler variant of the above construction, where we set $U_1 = 0$, formally – ie we consider the free module over the polynomial algebra $\mathbb{F}[U_2, \dots, U_\ell]$ generated by $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, endowed with the differential:

$$\partial \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid n_{w_1}(\phi) = 0, \mu(\phi) = 1\}} \# \widehat{\mathcal{M}}(\phi) \cdot U_2^{n_{w_2}(\phi)} \cdots U_\ell^{n_{w_\ell}(\phi)} \otimes \mathbf{y}$$

It is easy to see that the arguments from Theorem 4.4 also show that the homology of this chain complex calculates $\widehat{HF}(Y)$.

The simplest variant counts holomorphic disks which are disjoint from all the w_i , to obtain a complex $\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{w})$ ie specializing $CF^-(\Sigma, \alpha, \beta, \mathbf{w})$ to the case where $U_i = 0$ for all i . Explicitly, this is the chain complex of \mathbb{F} -vector spaces spanned by the intersection points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and equipped with the differential

$$\partial \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid n_{\mathbf{w}}(\phi) = 0, \mu(\phi) = 1\}} \# \widehat{\mathcal{M}}(\phi) \cdot \mathbf{y}.$$

For this complex, we have the following:

Theorem 4.5 *The complex $\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{w})$ calculates $\widehat{HF}(Y \# (\#^{\ell-1}(S^2 \times S^1)))$.*

4.2 Links and multi-filtrations of \widehat{HF}

The multi-pointed Heegaard diagram for an oriented link \vec{L} endows the chain complex $CF^-(\Sigma, \alpha, \beta, \mathbf{w})$ with a relative \mathbb{Z}^ℓ -filtration, as follows.

Let $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ be the Heegaard diagram for an oriented link in a three-manifold with $H^1(Y; \mathbb{Z}) = 0$. In Section 3.6, we defined a map

$$\underline{\mathfrak{s}}_{\mathbf{w}, \mathbf{z}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \longrightarrow \underline{\text{Spin}}^c(Y, L).$$

We extend the above-defined function to generators of the chain complex

$$CF^-(\Sigma, \alpha, \beta, \mathbf{w})$$

of the form $U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{x}$, using the formula

$$(12) \quad \underline{s}_{\mathbf{w}, \mathbf{z}}(U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{x}) = \underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) - a_1 \cdot \text{PD}[\mu_1] - \cdots - a_\ell \cdot \text{PD}[\mu_\ell],$$

where here $\{\mu_i\}_{i=1}^\ell$ are meridians for the link (compatible with its given orientation).

Lemma 4.6 *Let $\vec{L} \subset Y$ be an oriented link in a three-manifold Y with $H_1(Y; \mathbb{Z}) = 0$, and consider the corresponding identification $\mathbb{Z}^\ell \cong H^2(Y, L)$, given by*

$$(a_1, \dots, a_\ell) \mapsto \sum_{i=1}^{\ell} a_i \cdot \text{PD}[\mu_i],$$

making $\text{Spin}^c(Y, L)$ into an affine space for \mathbb{Z}^ℓ . Then, the function \underline{s} (as defined in Section 3.2, and extended in Equation (12) above) induces a $\text{Spin}^c(Y, L)$ -filtration on the chain complex $CF^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ (endowed with the differential from Equation (10)).

Proof Suppose $U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{y}$ appears in $\partial \mathbf{x}$ with non-zero coefficient. We must prove that then $\underline{s}(\mathbf{x}) \geq \underline{s}(U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{y})$. But in this case, there is a $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with a holomorphic representative, and hence $n_{\mathbf{w}}(\phi), n_{\mathbf{z}}(\phi) \geq 0$. On the other hand, $a_i = n_{w_i}(\phi)$, and hence

$$\underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) - \underline{s}_{\mathbf{w}, \mathbf{z}}(U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{y}) = \underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) - \underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{y}) + n_{\mathbf{w}}(\phi) = n_{\mathbf{z}}(\phi),$$

in view of Equation (7). □

The following will be verified in Section 7:

Theorem 4.7 *The $\text{Spin}^c(Y, L)$ -filtered chain homotopy type of the chain complex $CF^-(\alpha, \beta, \mathbf{w}, \mathbf{z})$ of $\mathbb{Z}[U_1, \dots, U_\ell]$ -modules is an invariant of the underlying oriented link.*

Definition 4.8 The $\text{Spin}^c(Y, L)$ -filtered chain homotopy type of the chain complex associated to a link will be denoted $\text{CFL}^-(Y, \vec{L})$. The homology of the associated graded object is denoted

$$\text{HFL}^-(Y, \vec{L}) = \bigoplus_{\underline{s} \in \text{Spin}^c(Y, L)} \text{HFL}^-(Y, \vec{L}, \underline{s}).$$

Explicitly, the associated graded object is generated as a $\mathbb{F}[U_1, \dots, U_\ell]$ -module by intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, endowed with a differential akin to Equation (10), only now we sum over those $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = 1$ and also $n_z(\phi) = 0$. Indeed, the summand in grading \underline{s} is generated by symbols $U^{a_1} \cdots U^{a_\ell} \otimes \mathbf{x}$, where a_1, \dots, a_ℓ are non-negative integers and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, satisfying the constraint that

$$\underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) - a_1 \text{PD}[\mu_1] - \cdots - a_\ell \text{PD}[\mu_\ell] = \underline{s}.$$

The homology of this complex is the group $\text{HFL}^-(Y, \vec{L}, \underline{s})$.

We can also set $U_1 = \cdots = U_\ell = 0$, to obtain a filtration of a chain complex which, according to Theorem 4.5, calculates \widehat{HF} of $Y \# (\#^{\ell-1}(S^2 \times S^1))$. More concretely, we consider the chain complex $\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{w})$ generated over \mathbb{F} by intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, endowed with the module homomorphism

$$(13) \quad \partial \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi) = 1, n_w(\phi) = 0\}} \# \widehat{\mathcal{M}}(\phi) \cdot \mathbf{y}.$$

As in the case of CFL^- , the function $\underline{s}_{\mathbf{w}, \mathbf{z}}$ endows $\widehat{CFL}(Y, \vec{L})$ with a filtration.

Definition 4.9 The $\text{Spin}^c(Y, L)$ -filtered chain homotopy type of the chain complex associated to a link will be denoted $\widehat{CFL}(Y, \vec{L})$. The homology of the associated graded complex is the link invariant $\widehat{\text{HFL}}(Y, \vec{L})$.

More precisely, for $\underline{s} \in \text{Spin}^c(Y, \vec{L})$, $\widehat{\text{HFL}}(Y, \vec{L}, \underline{s})$ is the homology of the chain complex generated by $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $\underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) = \underline{s}$, endowed with the differential

$$(14) \quad \partial \mathbf{x} = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi) = 1, n_w(\phi) = n_z(\phi) = 0\}} \# \widehat{\mathcal{M}}(\phi) \cdot \mathbf{y}.$$

In the introduction, no mention was made of relative Spin^c structures; rather, link Floer homology was described as a group graded by elements of \mathbb{H} . For the case of links in S^3 , the equivalence of these two points of view is given as follows. Given an element $h = \sum a_i[\mu_i] \in \mathbb{H}$, there is a unique relative Spin^c structure \underline{s} with the property that

$$c_1(\underline{s}) - \sum_{i=1}^{\ell} \text{PD}[\mu_i] = 2\text{PD}[h].$$

The group $\widehat{\text{HFL}}(L, h)$ from the introduction, then, is the group $\widehat{\text{HFL}}(L, \underline{s})$. This convention is quite natural, as we shall see in Section 8.1 below.

In practice, it can be taxing to calculate relative Spin^c structures. It is much simpler, rather, to define the link filtration on relative terms, declaring that

$$(15) \quad \mathbf{x} > \mathbf{y} \quad \text{if and only if} \quad n_{\mathbf{z}}(\phi) \geq n_{\mathbf{w}}(\phi),$$

where here ϕ is any element in $\pi_2(\mathbf{x}, \mathbf{y})$, for example in the case where $H_1(Y; \mathbb{Z}) = 0$. (The equivalence of this with our earlier point of view is a direct consequence of Equation (7).) This determines the filtration only up to an overall shift (by an element of $H^2(Y, L; \mathbb{Z})$), but this indeterminacy can be removed using the symmetry properties of the link invariant, cf Section 8 below.

5 Analytic input

We will be concerned in this section with gluing results for pseudo-holomorphic disks. Consider an ℓ -pointed Heegaard diagram $(\Sigma, \alpha, \beta, \{w_1, \dots, w_\ell\})$, and let $d = g + \ell - 1$. Given $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and a homology class of Whitney disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we can form the moduli space $\mathcal{M}(\phi)$.

Recall that these moduli spaces have Gromov compactifications, cf [7; 16; 4; 5]. For a given moduli space of pseudo-holomorphic Whitney disks, these Gromov compactifications include possibly moduli spaces of pseudo-holomorphic Whitney disks connecting other intersection points, moduli spaces of pseudo-holomorphic spheres, and finally also moduli spaces of further degenerate disks called *boundary degenerations*. More formally, given a point $\mathbf{x} \in \mathbb{T}_\alpha$, we let $\pi_2^\beta(\mathbf{x})$ denote the space of homology classes of maps:

$$\left\{ u: [0, \infty) \times \mathbb{R} \longrightarrow \text{Sym}^g(\Sigma) \left| \begin{array}{l} u(\{0\} \times \mathbb{R}) \subset \mathbb{T}_\beta \\ \lim_{z \rightarrow \infty} u(z) = \mathbf{x} \end{array} \right. \right\}$$

Such a map is called a β -boundary degeneration. We let $\pi_2^\alpha(\mathbf{x})$ denote the set of α -boundary degenerations, defined analogously.

Loosely speaking, Gromov's compactness theorem states that a sequence of pseudo-holomorphic curves representing ϕ , has a subsequence which converges locally to a "broken flow-line," consisting of collection of pseudo-holomorphic flow-lines $\{\phi_i\}$, a collection of α - and β -boundary degenerations $\{\psi_j\}$, and finally a collection of pseudo-holomorphic spheres $\{S_k\}$ with

$$\sum \mathcal{D}(\phi_i) + \sum \mathcal{D}(\psi_j) + \sum \mathcal{D}(S_k) = \mathcal{D}(\phi).$$

We will also find it necessary at several future points to study ends of moduli spaces as the Heegaard surface is degenerated. We turn to this more formally as follows.

5.1 Gluing moduli spaces

Let $(\Sigma_1, \alpha_1, \beta_1, z_1)$ and $(\Sigma_2, \alpha_2, \beta_2, z_2)$ be a pair of Heegaard diagrams, where here α_i and β_i are d_i -tuples of attaching circles in Σ_i . We can form their connected sum at $z_1 \in \Sigma_1$ and $z_2 \in \Sigma_2$ to obtain a new surface $\Sigma = \Sigma_1 \# \Sigma_2$, endowed with sets of attaching circles $\alpha_1 \cup \alpha_2$ and $\beta_1 \cup \beta_2$. We will assume that $d_i \geq g_i$, and write $d = d_1 + d_2$ and $g = g_1 + g_2$.

We will need here descriptions of the moduli spaces of flowlines in the connected sum diagram, in terms of moduli spaces for the two pieces.

Fix points $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{T}_{\alpha_i} \cap \mathbb{T}_{\beta_i} \subset \text{Sym}^{d_i}(\Sigma_i)$, which in turn give rise to points $\mathbf{x}_1 \times \mathbf{x}_2, \mathbf{y}_1 \times \mathbf{y}_2$ in $\text{Sym}^{d_1+d_2}(\Sigma_1 \# \Sigma_2)$. Fix a pair of homology classes of Whitney disks $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{y}_i)$ with $n_{z_1}(\phi_1) = n_{z_2}(\phi_2)$. These can be combined naturally to form a homology class $\phi_1 \# \phi_2 \in \pi_2(\mathbf{x}_1 \times \mathbf{x}_2, \mathbf{y}_1 \times \mathbf{y}_2)$. Specifically, the local multiplicities of $\phi_1 \# \phi_2$ at each domain $\mathcal{D} \subset \Sigma_1 \# \Sigma_2$ is the local multiplicity of ϕ_i at the corresponding $\mathcal{D} \subset \Sigma_1$ or Σ_2 .

Conversely, each homology class $\phi \in \pi_2(\mathbf{x}_1 \times \mathbf{y}_1, \mathbf{x}_2 \times \mathbf{y}_2)$ can be uniquely decomposed as $\phi = \phi_1 \# \phi_2$ for some pair of $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{y}_i)$ with $n_{z_1}(\phi_1) = n_{z_2}(\phi_2)$.

Moreover, given complex structures j_1 and j_2 on Σ_1 and Σ_2 , we can form a complex structure $J(T)$ with neck-length T as follows: Find conformal disks D_1 and D_2 about z_1 and z_2 , and form the connected sum

$$\Sigma(T) = (\Sigma_1 - D_1) \# \left([-T - 1, T + 1] \times S^1 \right) \# (\Sigma_2 - D_2),$$

under identifications $\partial D_1 \cong \{-T\} \times S^1$ and $\{T\} \times S^1 \cong \partial D_2$. As $T \mapsto \infty$, the conformal structure on $\Sigma(T)$ converges to the nodal curve $\Sigma_1 \vee \Sigma_2$.

Theorem 5.1 *Fix diagrams $(\Sigma_i, \alpha_i, \beta_i, z_i)$ for $i = 1, 2$ as above, where here α_i and β_i are d_i -tuples of attaching circles. Given a homology class $\phi = \phi_1 \# \phi_2$ for the connected sum of the two diagrams, we have that*

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2k,$$

where $k = n_{z_1}(\phi) = n_{z_2}(\phi)$. Suppose that $\mathcal{M}(\phi) \neq \emptyset$ for a sequence of almost-complex structures $J(T_i)$ with $T_i \mapsto \infty$. Then, the moduli spaces of broken pseudo-holomorphic flowlines representing ϕ_1 and ϕ_2 (ie the Gromov compactifications of these two moduli spaces) are non-empty. Finally, suppose that $\mu(\phi_1) = 1, \mu(\phi_2) = 2k$, and also that $d_2 > g_2$; and consider the maps

$$\rho_1: \mathcal{M}(\phi_1) \longrightarrow \text{Sym}^k(\mathbb{D}) \quad \text{and} \quad \rho_2: \mathcal{M}(\phi_2) \longrightarrow \text{Sym}^k(\mathbb{D}),$$

where here $\rho_i(u) = u^{-1}(\{z_i\} \times \text{Sym}^{d_i-1}(\Sigma_i))$.

If the fibered product of $\mathcal{M}(\phi_1)$ and $\mathcal{M}(\phi_2)$

$$\mathcal{M}(\phi_1) \times_{\text{Sym}^k(\mathbb{D})} \mathcal{M}(\phi_2) = \{u_1 \times u_2 \in \mathcal{M}(\phi_1) \times \mathcal{M}(\phi_2) \mid \rho_1(u_1) = \rho_2(u_2)\}$$

is a smooth manifold, then there is an identification of this moduli space with the moduli space $\mathcal{M}(\phi)$, for sufficiently long connected sum length.

There are three assertions in the above theorem: One concerns the Maslov index, one the existence of weak limits, and the third is a gluing result. They are arranged in order of difficulty; the third requires the most work. One could approach this problem from the point of view of degenerating $\text{Sym}^d(\Sigma_1 \# \Sigma_2)$ as the connected sum degenerates into $\Sigma_1 \vee \Sigma_2$, following the approach to stabilization invariance from [21]. In this case, the limiting symplectic space is $\text{Sym}^d(\Sigma_1 \vee \Sigma_2)$, which has a fairly complicated singular set (consisting of those d -tuples where one element is the singular point $p \in \Sigma_1 \vee \Sigma_2$). In particular, the singular set is, in itself, not a symplectic manifold but rather a singular space whose singularities consist of those d -tuples where at least two elements are the point p . Under the present circumstances, holomorphic disks we wish to resolve – whose boundary lies from in the top stratum $\text{Sym}^{d_1}(\Sigma_1) \times \text{Sym}^{d_2}(\Sigma_2) \subset \text{Sym}^d(\Sigma_1 \vee \Sigma_2)$ – meet the singular set (consisting of those tuples where at least one point is the singular point $z \in \Sigma_1 \vee \Sigma_2$) in a complex codimension one subset, ie where at least two coordinates agree with this singular point.

A much simpler approach can be given using Lipshitz's cylindrical reformulation of Heegaard Floer homology [15]. With this reformulation, then, the gluing problem takes place in a four-manifold, along a singular set which is a manifold, (placing it on roughly an equal footing with the proof of stabilization invariance for cylindrical reformulation, cf [15]). This kind of degeneration has been extensively studied in the literature, cf Ionel and Parker [8], Li and Ruan [14], Eliashberg, Givental and Hofer [3], Bourgeois, Eliashberg, Hofer, Wysocki and Zehnder [1], and of course [15]. We turn to this approach, first recalling the basic set-up of Lipshitz's picture.

5.2 Lipshitz's cylindrical formulation of Heegaard Floer homology

The starting point of Lipshitz's formulation is that a holomorphic disk $u: \mathbb{D} \rightarrow \text{Sym}^d(\Sigma)$ corresponds to a holomorphic curve in $\mathbb{D} \times \Sigma$, which, for any $p \in \mathbb{D}$, meets the fiber $\Sigma \times \{p\}$ in the set of d points $u(p)$. Thus, one could reformulate the chain complex defining Heegaard Floer homology as counting certain pseudo-holomorphic surfaces in $\mathbb{D} \times \Sigma$.

This can be made more precisely as follows. Consider the four-manifold

$$W = \Sigma \times [0, 1] \times \mathbb{R},$$

equipped with two projection maps

$$\pi_\Sigma: W \longrightarrow \Sigma \quad \text{and} \quad \pi_{\mathbb{D}}: W \longrightarrow [0, 1] \times \mathbb{R}.$$

(As usual, we think of the unit disk in the complex plane \mathbb{D} as the conformal compactification of the infinite strip $[0, 1] \times \mathbb{R}$, obtained by adding points $\pm i$ at infinity.) Endow W with an almost-complex structure J tamed by a natural split symplectic form on W , which is translation invariant in the \mathbb{R} -factor, and for which the projection $\pi_{\mathbb{D}}$ is a pseudo-holomorphic map. For example, the product complex structure – which will be called a *split* complex structure – satisfies these conditions, but it is often useful to perturb this. However, to ensure positivity, it is convenient to choose points $z \in \Sigma$, and require that J is split in a neighborhood of z (in Σ) times $[0, 1] \times \mathbb{R}$. Such an almost-complex structure J is called *split near z* .

Consider next a Riemann surface S with boundary, d “positive” punctures

$$\{p_1, \dots, p_d\}$$

and d “negative” punctures $\{q_1, \dots, q_d\}$ on its boundary.

Lipshitz considers pseudo-holomorphic maps $\tilde{u}: S \longrightarrow W$, satisfying the following conditions:

- \tilde{u} is smooth.
- $\tilde{u}(\partial S) \subset (\alpha \times \{1\} \times \mathbb{R}) \cup (\beta \times \{0\} \times \mathbb{R})$
- No component in the image of $\tilde{u}(S)$ is contained in a fiber of $\pi_{\mathbb{D}}$.
- For each i , $\tilde{u}^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ and $\tilde{u}^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ consist of exactly one component of $\partial S - \{p_1, \dots, p_d, q_1, \dots, q_d\}$.
- The energy of \tilde{u} is finite.
- \tilde{u} is an embedding.
- Any sequence of points in S converging to q_i resp. p_i is mapped under $\pi_{\mathbb{D}}$ to a sequence of points whose second coordinate converges to $\{-\infty\}$ resp $\{+\infty\}$.

We call holomorphic curves of this type *cylindrical flow-lines*. Thinking of the complex disk \mathbb{D} as a compactification of $[0, 1] \times \mathbb{R}$, a map \tilde{u} as above can be extended continuously to a map of the closure of S into $\overline{W} = \Sigma \times \mathbb{D}$. We say that \tilde{u} connects \mathbf{x} to \mathbf{y} if the image of this extension meets $\Sigma \times \{-i\}$ in the points $\mathbf{x} \times \{-i\}$ and it meets $\Sigma \times \{i\}$ in the points $\mathbf{y} \times \{i\}$.

Projecting such a map \tilde{u} onto Σ , we obtain a relative two-chain in Σ relative to $\alpha \cup \beta$, whose local multiplicity at some point $z \in \Sigma$ is given by the intersection number

$$\tilde{n}_z(\tilde{u}) = \#(\tilde{u} \cap (\{z\} \times [0, 1] \times \mathbb{R})).$$

Conversely, given $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, let $\widetilde{\mathcal{M}}(\phi)$ denote the moduli space of cylindrical flow-lines \widetilde{u} which induce the same two-chain as ϕ .

It is sometimes also useful to consider the analogue of boundary degenerations in this cylindrical context.

Definition 5.2 Consider a Riemann surface S with boundary and d punctures

$$\{p_1, \dots, p_d\}$$

on its boundary. Consider now pseudo-holomorphic maps $\widetilde{u}: S \rightarrow \Sigma \times [0, \infty) \times \mathbb{R}$ which are finite energy, smooth embeddings, sending the boundary of Σ into β , containing no component in the fiber of the projection to $[0, \infty) \times \mathbb{R}$, so that each component of $\widetilde{u}^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ consists of exactly one component of $\partial S - \{p_1, \dots, p_d\}$. Such a map is called a *cylindrical boundary degeneration*. For such a map, the point at infinity is mapped into a fixed $\mathbf{x} \in \mathbb{T}_\beta$. These maps can be organized into moduli spaces of $\widetilde{\mathcal{N}}(\psi^\beta)$ indexed by homology classes $\psi \in H_2(\Sigma, \beta)$. A corresponding definition can also be made with α playing the role of β .

In [15], Lipshitz sets up a theory analogous to Heegaard Floer homology (in the case where $d = g$), counting elements of $\widetilde{\mathcal{M}}(\widetilde{\phi})$. In setting up this theory, he establishes the necessary transversality properties: For a generic choice of almost-complex structure J over W as above, all the moduli spaces $\widetilde{\mathcal{M}}(\widetilde{\phi})$ with $\mu(\widetilde{\phi}) < 0$ are empty; non-empty moduli spaces $\widetilde{\mathcal{M}}(\widetilde{\phi})$ with $\mu(\widetilde{\phi}) = 0$ consist of constant flowlines, and moduli spaces with $\mu(\widetilde{\phi}) = 1$ are smooth one-manifolds. He also shows that these moduli spaces have the necessary Gromov compactifications analogous to those for the moduli spaces of holomorphic disks in a more traditional Lagrangian set-up. Indeed, in [15, Appendix A], Lipshitz establishes identifications $\widetilde{\mathcal{M}}(\widetilde{\phi}) \cong \mathcal{M}(\phi)$ for suitably generic choices of almost-complex structures, in cases where $\mu(\phi) = 1$. In particular, if we consider the map defined as in Equation (10), only using moduli spaces $\widetilde{\mathcal{M}}(\widetilde{\phi})/\mathbb{R}$ in place of $\mathcal{M}(\phi)/\mathbb{R}$, then the two maps actually agree. Although Lipshitz considers the case where $d = g$, the logic here applies immediately to prove the corresponding result in the case where $d > g$, as well.

With this cylindrical formulation, now the denegeration considered in Theorem 5.1 becomes more transparent. Specifically, the degeneration of Σ to $\Sigma_1 \vee \Sigma_2$ corresponds to a generation of W into $W_1 = \Sigma_1 \times [0, 1] \times \mathbb{R}$ and $W_2 = \Sigma_2 \times [0, 1] \times \mathbb{R}$, two symplectic manifolds which meet along the hypersurface $[0, 1] \times \mathbb{R}$, joined along $\{p_1\} \times [0, 1] \times \mathbb{R} \subset W_1$ and $\{p_2\} \times [0, 1] \times \mathbb{R} \subset W_2$. (Compare also [8; 14; 3; 1]. Finally, observe that this is precisely the set-up which Lipshitz uses in his proof of stabilization invariance for the cylindrical theory, see [15]).

More precisely, we start with almost-complex structures J_1 and J_2 on W_1 and W_2 and neighborhoods D_1 and D_2 of z_1 and z_2 . We assume that J_1 and J_2 which are split on $D_1 \times [0, 1] \times \mathbb{R}$ and $D_2 \times [0, 1] \times \mathbb{R}$ respectively. From this, we construct a complex structure $J(T)$ on

$$W(T) = \Sigma(T) \times [0, 1] \times \mathbb{R},$$

which agrees with J_1 near $(\Sigma_1 - D_1) \times [0, 1] \times \mathbb{R}$, and J_2 near $(\Sigma_2 - D_2) \times [0, 1] \times \mathbb{R}$, and which is split over $S^1 \times [-T - 1, T + 1] \times [0, 1] \times \mathbb{R}$.

Definition 5.3 A *pre-glued flowline representing the homology class* $\phi = \phi_1 \# \phi_2 \in \pi_2(\mathbf{x}_1 \times \mathbf{y}_1, \mathbf{x}_2 \times \mathbf{y}_2)$ is a pair of cylindrical flow-lines $\tilde{u}_1 \in \widetilde{\mathcal{M}}(\phi_1)$ and $\tilde{u}_2 \in \mathcal{M}(\phi_2)$ satisfying the matching condition

$$(\pi_{\mathbb{D}} \circ \tilde{u}_1)(\pi_{\Sigma_1} \circ \tilde{u}_1)^{-1}(\{z_1\}) = (\pi_{\mathbb{D}} \circ \tilde{u}_2)(\pi_{\Sigma_2} \circ \tilde{u}_2)^{-1}(\{z_2\}).$$

Similarly, a *pre-glued α -boundary degeneration* is a pair of α -boundary degenerations $\tilde{v}_1 \in \mathcal{M}(\phi_1)$ and $\tilde{v}_2 \in \mathcal{M}(\phi_2)$ satisfying the analogous matching condition $(\pi_{\mathbb{D}} \circ \tilde{u})(\pi_{\Sigma_1} \circ \tilde{v}_1)^{-1}(\{z_1\}) = (\pi_{\mathbb{D}} \circ \tilde{u})(\pi_{\Sigma_2} \circ \tilde{v}_2)^{-1}(\{z_2\})$. A similar definition can also be made for β -boundary degenerations.

The curves $\alpha_1 \cup \dots \cup \alpha_{g+\ell-1}$ divide Σ into ℓ regions A_1, \dots, A_ℓ . Choose reference points w_i , one in each A_i .

It will be useful to have the following:

Lemma 5.4 Given $\tilde{\psi} \in \pi_2^\beta(\mathbf{x})$, we have that $\mu(\tilde{\psi}) = 2 \sum_{i=1}^\ell n_{w_i}(\tilde{\psi})$.

Proof This follows readily from the excision principle for the the linearized $\bar{\partial}$ operator, to reduce to the case of a disk. (See the proof of Theorem 5.5 for a more detailed discussion of a related problem.) □

We interrupt now our path to Theorem 5.1, paying off an earlier debt, supplying the following quick consequence of the above lemma:

Proof of Proposition 4.1 According to Equation (5), the homology classes of ϕ and ϕ' differ by the juxtaposition of a boundary degeneration with $\mathcal{D}(\psi) = \sum_i (n_{w_i}(\phi) - n_{w_i}(\phi'))$, whose index, according to Lemma 5.4, is given by $2 \sum_i (n_{w_i}(\phi) - n_{w_i}(\phi'))$. The result now follows from the additivity of the index under juxtaposition.

We give another consequence of Lemma 5.4. By arranging for the almost-complex structure on W to be split near neighborhoods of all the w_i , one can arrange for the usual positivity principle to hold: If a moduli space $\mathcal{N}(\psi^\beta)$ is non-empty, then all the $n_{w_i}(\psi^\beta) \geq 0$. It follows from this, together with Lemma 5.4, that if ψ is a homology class of boundary degenerations which contains a non-constant pseudo-holomorphic representative, then $\mu(\psi) \geq 2$. This principle is used in the following:

Proof of the cylindrical analogue of Theorem 5.1 We turn to Theorem 5.1, using moduli spaces of cylindrical flowlines in place of pseudo-holomorphic Whitney disks. In this context, the restriction maps are to be replaced by maps

$$\tilde{\rho}_i: \tilde{\mathcal{M}}(\tilde{\phi}_i) \longrightarrow \text{Sym}^k(\mathbb{D})$$

given by

$$\tilde{\rho}_i(\tilde{u}) = (\pi_{\mathbb{D}} \circ \tilde{u})((\pi_{\Sigma} \circ \tilde{u})^{-1}(\{z_i\})).$$

The formula for the Maslov index follows readily from the excision principle for the linearized $\bar{\partial}$ operator, using the cylindrical formulation. Consider a sequence of pseudo-holomorphic curves $\{v_t\}_{t \in \mathbb{Z}} \in \tilde{\mathcal{M}}_{J(t)}(\phi)$, where here the subscript $J(t)$ denotes the almost-complex structure induced on $W(t)$ with neck-length t as described earlier. Using Gromov's compactness theorem, after passing to a subsequence, v_t converges locally to a pseudo-holomorphic curve in the symplectic manifold

$$W(\infty) \cong \left(W_1 - \{z_1\} \times [0, 1] \times \mathbb{R} \right) \amalg \left(W_2 - \{z_2\} \times [0, 1] \times \mathbb{R} \right),$$

which in turn can be completed to a pseudo-holomorphic curve in W_1 and one in W_2 . More precisely, we obtain a broken flow-line whose components consist of pre-glued flowlines and boundary degenerations, finally also nodal curves supported entirely inside fibers $\Sigma_1 \vee \Sigma_2$. The representatives in the moduli spaces for ϕ_1 and ϕ_2 respectively are gotten by ignoring the matching conditions.

In the case where $\mu(\phi_1) = 1$, the limiting process generically gives rise to an unbroken, preglued flowline, according to the following dimension counts. Specifically, taking a Gromov compactification, we obtain a pseudo-holomorphic representative \tilde{u}_1 of ϕ_1 , and also a possibly broken flow-line representing ϕ_2 .

Assuming that this broken flow-line contains no components which are closed curves, there is some component of it u_2 with the property that u_1 and u_2 represents a pre-glued flowline in the sense of Definition 5.3. We claim that u_2 in fact represents ϕ_2 . If \tilde{u}_2 did not represent ϕ_2 , then it represents a homology class ϕ' which is a component in the Gromov compactification of ϕ' . If this compactification contains additional boundary degenerations, the Maslov index of ϕ' is at least 2 smaller than the Maslov

index of ϕ_2 (in view of Lemma 5.4). Moreover, if the compactification contains other flows, those will serve only to further decrease the Maslov index of ϕ' relative to that of ϕ . In sum, in the case where ϕ' does not agree with ϕ_2 , its Maslov index $\mu(\phi') < 2k$. But given $\Delta \in \text{Sym}^k(\mathbb{D})$, the moduli space

$$\rho_2^{-1}(\Delta) = \{\tilde{u} \in \widetilde{\mathcal{M}}(\tilde{\phi}') \mid (\pi_{\mathbb{D}} \circ \tilde{u}_2)(\pi_{\Sigma} \circ \tilde{u}_2)^{-1}(\{z_2\}) = \Delta\}$$

has expected dimension $\mu(\phi') - 2k < 0$. Thus, for a generic choice of Δ (gotten by $\rho_1(\phi_1)$), this space is empty.

It is easy to rule out also the case that the Gromov limit representing ϕ_2 cannot have any closed components. To this end, suppose it has some components which represent the homology class $m[\Sigma_2]$ for some $m > 2$. After deleting these components, we are left with a homology class ϕ' with $\mu(\phi'_2) = 2(d_2 - g_2 + 1)$ (cf Lemma 5.4). Some component u_2 of this Gromov compactification has $\rho_2(u_2) = \Delta'$, where Δ' is obtained from $\rho_1(\phi_1)$ by deleting m points. But the moduli space of such points has expected dimension given by

$$(16) \quad \mu(\phi'_2) - 2(k - m) = 2k - 2(k - m) - 2(d_2 + 1 - g_2)m = -2m(d_2 - g_2) < 0,$$

and hence it is empty (here, of course, is where we used our assumption that $d_2 > g_2$).

Thus, we have established that a sequence of holomorphic representatives for ϕ has a Gromov limit (as we stretch the neck) to a pre-glued flowline representing ϕ_1 and ϕ_2 . Conversely, given a pre-glued flowline, one can obtain a pseudo-holomorphic curve in $\widetilde{\mathcal{M}}(\phi)$ by gluing (cf [15] for further details on this gluing problem, and [1] for a discussion of gluing in a very general context).

5.3 Counting boundary degenerations

We let $\widetilde{\mathcal{N}}(\psi)$ denote the moduli space of pseudo-holomorphic boundary degenerations in the homology class of ψ . Note that $\mathbb{P}SL_2(\mathbb{R})$ acts on this moduli space, and we let $\widehat{\mathcal{N}}(\psi)$ denote the quotient by this action. The principles used in proving Theorem 5.1 can also be used to count boundary degenerations.

Theorem 5.5 *Consider Σ , a surface of genus g , equipped with a set of attaching circles $\alpha_1, \dots, \alpha_{g+\ell-1}$ for a handlebody. If $\mathcal{D}(\psi) \geq 0$ and $\mu(\psi) = 2$, then $\mathcal{D}(\psi) = A_i$ for some i ; and indeed in this case*

$$(17) \quad \#\widehat{\mathcal{N}}(\psi) \equiv \begin{cases} 0 \pmod{2} & \text{if } \ell = 1, \\ 1 \pmod{2} & \text{if } \ell > 1. \end{cases}$$

Proof It follows readily from Lemma 5.4 that if ψ is a non-zero homology class of boundary degenerations with $\mathcal{D}(\psi) \geq 0$, then $\mu(\psi) \geq 2$, with equality iff $\mathcal{D}(\psi) = A_i$. In the case where $\mathcal{D}(\phi) = A_i$, it remains to verify Equation (17). The case where $\ell = 1$ has already been established in [21, Theorem 3.15] for the usual Heegaard-Floer moduli spaces and [15] for the cylindrical version. Consider some region A_i , which we re-name simply A' . Recall that this is a Riemann surface with boundary, equipped with $m + p$ curves $\alpha_1, \dots, \alpha_{m+p}$, of which the first m comprise its boundary, and the rest are pairwise disjoint, embedded circles in the interior. We have fixed also $\mathbf{x} \in \alpha_1 \times \dots \times \alpha_{m+p}$. We reduce to the case where $p = 0$, by de-stabilizing. Next, we reduce to the case where the number of boundary circles is one. To this end, we can write $A' = A \# D$, where D here is a disk with boundary α_m , and A is a planar surface-with-boundary $\alpha_1, \dots, \alpha_{m-1}$. Denote the connected sum point in A by z_1 and the one in D by z_2 . Degenerating the connected sum tube, we obtain a fibered product description $\tilde{\mathcal{N}}(A') = \tilde{\mathcal{N}}(A) \times_{\text{Sym}^1(\mathbb{D})} \tilde{\mathcal{N}}(D)$, where the fibered product is taken over the maps

$$\rho_1: \tilde{\mathcal{N}}(A) \longrightarrow \text{Sym}^1(\mathbb{D}) \quad \text{and} \quad \rho_2: \tilde{\mathcal{N}}(D) \longrightarrow \text{Sym}^1(\mathbb{D})$$

defined as before. It is easy to see, though that $\tilde{\mathcal{N}}(D)$ is smooth, and the map ρ_2 is in fact a diffeomorphism. Thus, it follows that $\tilde{\mathcal{N}}(A') \cong \tilde{\mathcal{N}}(A)$. In this manner, we have reduced to the case where $p = 0$ and $m = 1$, which is, once again, the case where $A = D$. Now, since the map ρ_2 is $\mathbb{P}SL_2(\mathbb{R})$ -equivariant, we see that $\#\tilde{\mathcal{N}}(D)$ is precisely the degree of the map ρ_2 , which is one. \square

5.4 Notational remark

The notation of Theorem 5.1 suggests using moduli spaces of pseudo-holomorphic Whitney disks, while its proof uses cylindrical flow-lines. One could close the gap by either appealing to an identification between the two kinds of moduli spaces, cf [15], or simply adopting the cylindrical point of view instead; either approach has the same final outcome. In view of this remark, we henceforth drop the notational distinction between cylindrical and more traditional moduli spaces.

6 Heegaard Floer homology for multi-pointed Heegaard diagrams revisited

Having set up the analytical preliminaries, we now prove the invariance properties promised in Section 4.

Proof of Lemma 4.3 In the usual proof that $\partial^2 = 0$ from Floer homology (cf [21, Theorem 3.3] for a proof in the original case, and [15] for the cylindrical formulation), this is observed by counting ends of two-dimensional moduli spaces.

Specifically, fix intersection points \mathbf{x} and \mathbf{q} , and a vector $\mathbf{a} = (a_1, \dots, a_\ell) \in \mathbb{Z}^\ell$. We consider the ends of

$$\coprod_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{q}) \mid \mu(\phi) = 2, n_w(\phi) = \mathbf{a}\}} \widehat{\mathcal{M}}(\phi).$$

In the case where $\mathbf{x} \neq \mathbf{q}$, these ends cannot contain any boundary degenerations, since according to Lemma 5.4, these all carry Maslov index at least 2 (and hence if they appear in the boundary, the remaining configuration has Maslov index ≤ 0 , and hence, if it is non-empty, it must consist of the constant flowline alone). Thus, the ends in this case are modeled on

$$\coprod_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \coprod_{\{\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}), \phi_2 \in \pi_2(\mathbf{y}, \mathbf{q}) \mid \phi_1 * \phi_2 = \phi\}} \widehat{\mathcal{M}}(\phi_1) \times \widehat{\mathcal{M}}(\phi_2).$$

And the total count of these ends are given by

$$(18) \quad \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}), \phi_2 \in \pi_2(\mathbf{y}, \mathbf{q}) \mid n_w(\phi_1) + n_w(\phi_2) = \mathbf{a}\}} (\#\widehat{\mathcal{M}}(\phi_1)) \cdot (\#\widehat{\mathcal{M}}(\phi_2)),$$

which on the one hand must be even, on the other hand, it is easily seen to be the $U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{q}$ -component of $(\partial^-)^2(\mathbf{x})$. In the case where $\mathbf{q} = \mathbf{x}$, there are additional terms, which count boundary degenerations meeting constant flowlines, whose total signed count is

$$\sum_{\{\psi \in \pi_2^\alpha(\mathbf{x}) \mid n_w(\psi) = \mathbf{a}, \mu(\psi) = 2\}} \#\widehat{\mathcal{N}}^\alpha(\psi) + \sum_{\{\psi \in \pi_2^\beta(\mathbf{x}) \mid n_w(\psi) = \mathbf{a}, \mu(\psi) = 2\}} \#\widehat{\mathcal{N}}^\beta(\psi).$$

According to Theorem 5.5, this quantity vanishes. More precisely, $\#\widehat{\mathcal{N}}^\alpha(\psi) \equiv 0 \pmod{2}$ except in the case where $\ell > 1$ and $\mathcal{D}(\psi)$ is one of the components A_i of $\Sigma - \alpha_1 - \cdots - \alpha_g$, so that \mathbf{a} has the form that $a_i = 0$ for all but one value of i , where the component $a_i = 1$. In this case $\#\widehat{\mathcal{N}}^\alpha(\psi) \equiv 1 \pmod{2}$, but there is also a unique cancelling ψ' with $\mathcal{D}(\psi') = B_i$ and $\#\widehat{\mathcal{N}}^\alpha(\psi') \equiv 1 \pmod{2}$. Thus, we are left once again with a sum as in Equation (18) which can be interpreted as the $U_1^{a_1} \cdots U_\ell^{a_\ell} \cdot \mathbf{q}$ -component of $(\partial^-)^2(\mathbf{x})$.

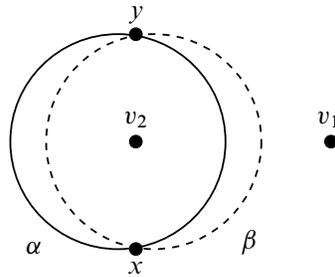


Figure 3: **Model doubly-pointed Heegaard diagram for S^3** We have illustrated here the doubly-pointed Heegaard diagram for S^3 $(S, \alpha, \beta, \{v_1, v_2\})$ considered above. The picture takes place on the two-sphere S .

6.1 Simple stabilization invariance

The aim of this subsection is to prove that the homology of the chain complex

$$CF^-(\Sigma, \alpha, \beta, \mathbf{w})$$

is invariant under a particularly simple index zero/three stabilizations (cf Proposition 3.3).

Specifically, recall that if we have a balanced Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w})$ with

$$\alpha = \{\alpha_1, \dots, \alpha_{g+l-1}\}, \quad \beta = \{\beta_1, \dots, \beta_{g+l-1}\}, \quad \mathbf{w} = \{w_1, \dots, w_\ell\},$$

we can construct a new balanced Heegaard diagram by introducing a new pair of separating curves α_{g+l} and β_{g+l} and a new basepoint $w_{\ell+1}$, so that α_{g+l} is isotopic to β_{g+l} in $\Sigma - w_1 - \dots - w_{\ell+1}$. Let

$$\alpha' = \{\alpha'_1, \dots, \alpha'_{g+l-1}, \alpha_{g+l}\}, \quad \beta' = \{\beta'_1, \dots, \beta'_{g+l-1}, \beta_{g+l}\},$$

where for $i = 1, \dots, g+l-1$, α'_i resp. β'_i is obtained from α_i resp. β_i by a small isotopic translate. As in Section 3, we say that this new diagram $(\Sigma, \alpha', \beta', \mathbf{w} \cup \{w_{\ell+1}\})$ is obtained from $(\Sigma, \alpha, \beta, \mathbf{w})$ by an index zero/three stabilization.

We call the stabilization *simple* if α_{g+l} bounds a disk in Σ whose closure is disjoint from the other α_i .

Let S be the two-sphere, and $\alpha \subset S$ be an embedded curve which divides S into two regions, each of which contains a basepoint v_1 or v_2 . Let β be a small Hamiltonian isotopic translate of α (in the complement of v_1 and v_2), meeting α in two points x and y , cf Figure 3. Of course, $(S, \alpha, \beta, \{v_1, v_2\})$ represents a balanced diagram for S^3 .

Lemma 6.1 *If $\phi \in \pi_2(x, x)$ or $\pi_2(y, y)$ with $\mathcal{D}(\phi) \geq 0$, then*

$$(19) \quad \mu(\phi) - 2n_{v_i}(\phi) \geq 0.$$

Moreover, the chain complex $CF^-(S, \alpha, \beta, \{v_1, v_2\})$ is given by

$$\mathbb{F}[U_1, U_2] \cdot x \xrightarrow{U_1 - U_2} \mathbb{F}[U_1, U_2] \cdot y.$$

Proof Inequality 19 follows easily from Lemma 5.4.

For the claim about the chain complex, it is easy to see that boundary maps from y to x come in cancelling pairs, whereas there are two holomorphic disks from x to y with Maslov index one, and they are bigons, one of which contains v_1 the other v_2 . \square

Definition 6.2 Let $(\Sigma_1, \xi, \eta, \{u_1, \dots, u_\ell\})$ and $(\Sigma_2, \gamma, \delta, \{v_1, \dots, v_m\})$ be a pair of balanced multi-pointed Heegaard diagrams, and choose basepoints $u_1 \in \mathbf{u}$, $v_1 \in \mathbf{v}$. We can form their *connected sum* $(\Sigma, \alpha, \beta, \mathbf{w})$, a balanced $\ell + m - 1$ -pointed Heegaard diagram whose underlying surface Σ is obtained from Σ_1 and Σ_2 by forming the connected sum at the points u_1 and v_1 . We let $\alpha = \xi \cup \gamma$, $\beta = \eta \cup \delta$, only now thought of as curves in Σ , and $\mathbf{w} = \{w_1, u_2, \dots, u_\ell, v_2, \dots, v_m\}$, where w_1 is some reference point on the connected sum neck. We denote this connected sum by $(\Sigma_1, \xi, \eta, \{u_1, \dots, u_\ell\}) \# (\Sigma_2, \gamma, \delta, \{v_1, \dots, v_m\})$.

In particular, given an arbitrary multi-pointed the Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w})$, a simple index zero/three stabilization can be thought of as the connected sum with $(S, \alpha, \beta, \{v_1, v_2\})$ as considered in Lemma 6.1.

Fix a Heegaard multi-diagram $(\Sigma, \alpha, \beta, \mathbf{w})$, choose $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $w_1 \in \mathbf{w}$. Correspondingly, let $\rho^{w_1}: \mathcal{M}(\phi) \rightarrow \text{Sym}^k(\mathbb{D})$, where $k = n_{w_1}(\phi)$, denote the map which assigns to the Whitney disk u the divisor

$$u^{-1}(w_1 \times \text{Sym}^{g+\ell-2}(\Sigma)),$$

or, in the cylindrical formulation, $\pi_{\mathbb{D}}(\pi_\Sigma \circ u)^{-1}(\{w_1\})$. Given $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ and $0 < t < 1$, we let $M(\phi, t)$ denote the moduli space of pseudo-holomorphic maps u representing ϕ with the additional constraint that $(t, 0) \in \rho^{w_1}(u)$.

Lemma 6.3 *Let $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ be a homotopy class of Whitney disks for a balanced Heegaard diagram*

$$(\Sigma, \{\alpha_1, \dots, \alpha_{g+\ell-1}\}, \{\beta_1, \dots, \beta_{g+\ell-1}\}, \{w_1, \dots, w_\ell\}).$$

If $\mu(\phi) = 2$, then $M(\phi, t)$ is generically a zero-dimension moduli space. Moreover, there is a number $\epsilon > 0$ with the property that for all $t \leq \epsilon$, the only possible non-empty

moduli spaces $M(\phi, t)$ with $\mu(\phi) = 2$ consist of moduli spaces for $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mathbf{x} = \mathbf{y}$, and indeed ϕ is obtained by splicing a $\mu(\psi) = 2$ boundary degeneration to the constant flowline. Indeed, for this moduli space,

$$\#M(\phi, t) = \begin{cases} 0 & \pmod{2} \text{ if } \ell = 1, \\ 1 & \pmod{2} \text{ if } \ell > 1. \end{cases}$$

Proof The dimension statement is clear. Note that the admissibility hypothesis ensures that there are at most finitely many $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ consisting of moduli spaces for with $\mu(\phi) = 2$ and $\mathcal{D}(\phi) \geq 0$, and hence only finitely many homotopy classes for which $M(\phi, t)$ could possibly be non-empty for some t . Consider one such homotopy class, and suppose that $M(\phi, t)$ is non-empty for a sequence of $t \mapsto 0$. Let $u_t \in M(\phi, t)$ be a corresponding sequence of pseudo-holomorphic curves. Taking their Gromov limit, we obtain a broken flow-line u representing the homology class ϕ . Since $\rho^{w_1}(u_t)$ contains points arbitrarily close to the line $\{0\} \times \mathbb{R}$ but w_1 does not lie in any of the β , we can conclude that the Gromov limit must contain a component which is a non-trivial boundary-degeneration ψ . According to Lemma 5.4, we can conclude that $\mu(\psi) \geq 2$. Thus, the remaining configuration $\phi - \psi$ has non-positive Maslov index, and it also has a pseudo-holomorphic representative. This forces it to be a constant flowline.

We have thus established that there is a real number $\epsilon > 0$ such that if $M(\phi, t)$ is non-empty for any $0 < t \leq \epsilon$, then ϕ is obtained by splicing a $\mu(\psi) = 2$ boundary degeneration to a constant flowline. The result now follows by gluing ψ to the constant flow-line, and applying Theorem 5.5. \square

We will also need a result about a suitable generalization of $M(\phi, t)$, but only for the Heegaard diagram $(S, \alpha, \beta, \{v_1, v_2\})$ introduced above.

Given a divisor $\Delta \in \text{Sym}^k(\mathbb{D})$, let $M(\phi, \Delta) = \{u \in \mathcal{M}(\phi) \mid \rho^{v_1}(\phi) = \Delta\}$.

Lemma 6.4 Consider the Heegaard diagram $(S, \alpha, \beta, \{v_1, v_2\})$ as above, with the two intersection points $x, y \in \alpha \cap \beta$. Fix also a generic $\Delta \in \text{Sym}^k(\mathbb{D})$ for some positive integer k . Then,

$$\sum_{\{\phi \in \pi_2(a, a) \mid \mu(\phi) = 2k, n_{v_2}(\phi) = 0\}} \#M(\phi, \Delta) \equiv 1 \pmod{2},$$

for $a = x$ or y .

Proof Consider the case where $a = x$. Let

$$S(x, \Delta) = \sum_{\{\phi \in \pi_2(x, x) \mid \mu(\phi) = 2k, n_{v_2}(\phi) = 0\}} \#M(\phi, \Delta).$$

We claim first that $S(x, \Delta)$ depends on Δ only through its total weight k . Specifically, if Δ_0, Δ_1 are two generic points in $\text{Sym}^k(\mathbb{D})$, then $S(x, \Delta_0) = S(x, \Delta_1)$. This follows from the following Gromov compactness argument. Let $\{\Delta_t\}_{t \in [0,1]}$ be a path in $\text{Sym}^k(\mathbb{D})$ connecting Δ_0 and Δ_1 . Consider the one-dimensional moduli space

$$\bigcup_{\{t \in [0,1], \phi \in \pi_2(x,x) \mid \mu(\phi) = 2k\}}$$

This has four types of ends. The first type appear in the expression

$$\bigcup_{\left\{ \begin{array}{l} \phi_1 \in \pi_2(x,y) \\ \phi_2 \in \pi_2(y,x) \\ t \in [0,1] \end{array} \mid \begin{array}{l} \mu(\phi_1) = 1 \\ n_{v_2}(\phi_1) = 0 = n_{v_2}(\phi_2) \\ \mu(\phi_2) = 2k - 1 \end{array} \right\}} \widehat{\mathcal{M}}(\phi_1) \times M(\phi_2, \Delta_t).$$

The total number of such ends is zero, since x is a cycle (in the chain complex for $(S, \alpha, \beta, \{v_1, v_2\})$). The second type of ends appear in

$$\bigcup_{\left\{ \begin{array}{l} \phi_1 \in \pi_2(x,y) \\ \phi_2 \in \pi_2(y,x) \\ t \in [0,1] \end{array} \mid \begin{array}{l} \mu(\phi_1) = 2k - 1, \\ \mu(\phi_2) = 1, \\ n_{v_2}(\phi_1) = 0 = n_{v_2}(\phi_2) \end{array} \right\}} M(\phi_1, \Delta_t) \times \widehat{\mathcal{M}}(\phi_2).$$

The total number of such ends is zero, since y is a cycle. Note that $n_{v_1}(\phi_1) = 0$ or $n_{v_2}(\phi_2) = 0$ for any such degeneration, since otherwise our divisor family Δ_t remains in a compact portion of the interior of the disk for all $t \in [0, 1]$.

The third and fourth types of ends appear in

$$\bigcup_{\{\phi \in \pi_2(x,x) \mid \mu(\phi) = 2k, n_{v_2}(\phi) = 0\}} M(\phi, \Delta_0),$$

and the related expression

$$\bigcup_{\{\phi \in \pi_2(x,x) \mid \mu(\phi) = 2k, n_{v_2}(\phi) = 0\}} M(\phi, \Delta_1).$$

Thus, taken together, the total number of ends is given by $S(x, \Delta_0) - S(x, \Delta_1) \pmod{2}$, which must therefore vanish.

Now let Δ_T consists of k points each of which have horizontal component $\leq \epsilon$ (chosen as in Lemma 6.3), and with vertical spacing of at least T between them. Taking a limit as $T \mapsto \infty$, we see that the ends of the parameterized moduli space

$$\bigcup_{\{T \in [1, \infty)\}} \bigcup_{\{\phi \in \pi_2(x,x) \mid \mu(\phi) = 2k\}} M(\phi, \Delta_T)$$

consist of $M(\phi, \Delta_1)$ (the end where $T = 1$), and a product $\prod_{i=1}^k M(\phi, t_i)$, in the notation of Lemma 6.3. Combining this observation with the result of that lemma, we see that $S(x, \Delta_1) \equiv 1 \pmod{2}$.

The case where $a = y$ follows similarly. \square

Proposition 6.5 *Suppose that $(\Sigma, \alpha', \beta', \mathbf{w}')$ is obtained from $(\Sigma, \alpha, \beta, \mathbf{w})$ by a simple index zero/three stabilization. Then, there is an identification of $\mathbb{F}[U_1, \dots, U_\ell]$ -modules*

$$HF^-(\Sigma, \alpha, \beta, \mathbf{w}) \cong HF^-(\Sigma, \alpha', \beta', \mathbf{w}').$$

Indeed, multiplication by $U_{\ell+1}$ is identified with multiplication by U_1 .

Proof We show that for suitable choices of almost-complex structures,

$$CF^-(\Sigma, \alpha', \beta', \mathbf{w}')$$

is identified with the mapping cone of a map

$$(20) \quad CF^-(\Sigma, \alpha, \beta, \mathbf{w})[U_{\ell+1}] \xrightarrow{U_{\ell+1}-U_1} CF^-(\Sigma, \alpha, \beta, \mathbf{w})[U_{\ell+1}],$$

where here $CF^-(\Sigma, \alpha, \beta, \mathbf{w})[U_{\ell+1}]$ denotes the chain complex

$$CF^-(\Sigma, \alpha, \beta, \mathbf{w}) \otimes_{\mathbb{F}[U_1, \dots, U_\ell]} \mathbb{F}[U_1, \dots, U_{\ell+1}].$$

We degenerate the Heegaard surface, to realize it as a connected sum

$$(\Sigma, \alpha, \beta, \mathbf{w}) \# (S, \alpha, \beta, v_1, v_2)$$

(with sufficiently large connected sum neck) where now w_1 is identified with v_1 ; or more precisely, w_1 and v_1 are the corresponding connected sum points, and we choose a new distinguished point w'_1 to lie in this connected sum region. Note that here $w_{\ell+1}$ is given by v_2 . We will use a complex structure on the connected sum surface with a very long connected sum length (and in fact, we will move the connected sum point v_1 close to the circle β in S , as explained below).

Clearly, $\mathbb{T}'_\alpha \cap \mathbb{T}'_\beta = (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times \{x, y\}$; ie intersection points come in two types (containing x and y), and differentials can be of four types. Write C_x resp $C_y \subset CF^-(\Sigma', \alpha', \beta', \mathbf{w}')$ as the submodule generated by intersection points containing x resp $y \in \alpha \cap \beta$; ie we have a module splitting $C' = CF^-(\Sigma', \alpha', \beta', \mathbf{w}', \mathbf{z}') \cong C_x \oplus C_y$.

We begin by considering the C_x -component of the differential of a generator in C_x . Apply Theorem 5.1 to some homotopy class $\phi \in \pi_2(\mathbf{x} \times \{x\}, \mathbf{y} \times \{x\})$, in the case where $\ell > 1$. Writing $\phi = \phi_1 \# \phi_2$, we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2n_{v_1}(\phi).$$

If $\#\mathcal{M}(\phi) \neq 0$, then $\mathcal{D}(\phi_2) \geq 0$ (since ϕ_2 has a pseudo-holomorphic representative). As an easy consequence of Lemma 5.4, it follows that $\mu(\phi_2) - 2n_{v_1}(\phi) \geq 0$, with equality iff $n_{v_2}(\phi_2) = 0$. In the case where, $n_{v_2}(\phi_2) \neq 0$ we force $\mu(\phi_1) \leq -1$, and hence it is generically empty. Thus, we are left with the case of $\mu(\phi_2) - 2n_{v_1}(\phi) = 0$, and $n_{v_2}(\phi) = 0$. In this case, Theorem 5.1 shows that the $\mathbf{y} \times x$ -component of $\partial^-(\mathbf{x} \times x)$ is given by

$$\sum_{\{\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi_1) = 1\}} \sum_{u_1 \in \widehat{\mathcal{M}}(\phi_1)} \sum_{\phi_2 \in \pi_2(x, x)} \#\{u_2 \in \mathcal{M}(\phi_2) \mid \rho_1(u_1) = \rho_2(u_2)\}.$$

In this expression, we consider $u_1 \in \widehat{\mathcal{M}}(\phi_1)$ as an actual map (rather than only one modulo re-parameterization) by taking the representatives with the property that the projection of $\rho_1(u_1)$ onto the \mathbb{R} factor contains 0, and no positive real number. According to Lemma 6.4, then, for each $u_1 \in \widehat{\mathcal{M}}(\phi_1)$,

$$\sum_{\phi_2 \in \pi_2(x, x)} \#\{u_2 \in \mathcal{M}(\phi_2) \mid \rho_1(u_1) = \rho_2(u_2)\} \equiv 1 \pmod{2}.$$

Thus, the $\mathbf{y} \times x$ component of $\partial^-(\mathbf{x} \times x)$ is identified with the \mathbf{y} component of $\partial^-(\mathbf{x})$ for the original Heegaard diagram.

Thus we have identified the C_x component of the differential of an element in $C_x \subset C'$ with the differential coming from the obvious identification of C_x with the chain complex for the original diagram $CF^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$. In the same manner, the differential within C_y is identified with the differential of the original diagram.

We consider now the C_x -component of the differential of a generator in C_y . Again, we express $\mathcal{M}(\phi)$ as a fibered product over $\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\phi_2 \in \pi_2(y, x)$. Now, an application of Lemma 5.4 shows that if $\mathcal{D}(\phi_2) \geq 0$, then $\mu(\phi_2) - 2n_{v_1}(\phi_2) \geq 1$, with equality iff $n_{v_2}(\phi_2) = 0$. In the case of equality, we have that $\mu(\phi_1) = 0$, and hence it must be constant. This forces ϕ_2 to be one of the two flows from y to x with $\mu(\phi_2) = 1$ and $\mathcal{D}(\phi_2) \geq 0$. Each of these homotopy classes admits a unique holomorphic representative, and hence the differential cancels: The C_x component of the boundary of something in C_y is trivial.

Finally, we consider the C_y component of the differential of a generator in C_x . Splitting a homotopy class ϕ in the fibered sum description, we once again have

$$\mu(\phi_1) + \mu(\phi_2) - 2n_{v_1}(\phi_1) = 1,$$

and hence the condition that $\mu(\phi_1) \geq 0$ (which is needed for its corresponding moduli space to be non-empty) translates into the condition that $\mu(\phi_2) - 2n_{v_1}(\phi_2) \leq 1$.

Moreover, for $\phi_2 \in \pi_2(x, y)$ with $\mathcal{D}(\phi_2) \geq 0$, we have that $\mu(\phi_2) - 2n_{v_1}(\phi_2) \geq -1$. Parity considerations exclude the possibility that the quantity equals 0.

In the case where $\mu(\phi_2) - 2n_{v_1}(\phi_2) = 1$, we conclude that $\mu(\phi_1) = 0$, and hence it must be constant, forcing $n_{v_1}(\phi_2) = 0$. Thus, $\phi_2 \in \pi_2(x, y)$ must be the unique homotopy class with $\mu(\phi_2) = 1$, $\mathcal{D}(\phi_2) \geq 0$, $n_{v_1}(\phi_2) = 0$, and $n_{v_2}(\phi_2) = 1$. Thus, the corresponding component of $\partial(\mathbf{x} \times x)$ is $U_{\ell+1} \cdot \mathbf{x} \times y$ (note that in the connected sum diagram, the reference point v_2 corresponds to the new variable $U_{\ell+1}$).

In the remaining cases, $\mu(\phi_1) = 2$ and $\mu(\phi_2) - 2n_{v_1}(\phi_2) = -1$. Since $\mu(\phi_2) \geq 0$, we conclude that $n_{v_1}(\phi_2) > 0$, while the condition that $\mathcal{D}(\phi_2) \geq 0$ and $\mu(\phi_2) - 2n_{v_1}(\phi_2) = -1$ readily forces $n_{v_2}(\phi_2) = 0$.

Suppose that $n_{v_1}(\phi_2) = 1 = \mu(\phi_2)$. In this case, ϕ_2 is constrained to be a homotopy class with a unique holomorphic representative up to translation. Let u_2 be a holomorphic representative of ϕ_2 , and $\Delta_2 = u_2^{-1}(\{v_1\})$. Note that this consists of a single point up to translation in \mathbb{D} ; after suitable translation, we arrange for $\Delta_2 = (t, 0) \in \mathbb{D}$. We wish now to apply Theorem 5.1, with now Σ playing the role of Σ_2 in the statement of that theorem. For this, we need to assume that ℓ , the number of marked points \mathbf{w} , is greater than one (so that we are taking more than the g^{th} symmetric product of Σ , as required in the hypothesis of the last part of Theorem 5.1). We return to the case where $\ell = 1$ at the end of this proof.

Completion of the proof of the proposition when $\ell > 1$ According to Theorem 5.1, count of points in all Maslov index one moduli spaces $\phi_1 \# \phi_2$, where $\phi_2 \in \pi_2(x, y)$ has $n_{v_1}(\phi_2) = \mu(\phi_2) = 1$ is given by the map

$$\delta_1: C_x \longrightarrow C_y$$

defined by

$$\delta_1(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi_1) = 2, n_{w_1}(\phi_1) = 1\}} \left(\#M(\phi_1, t) \right) \cdot \left(\prod_{i=1}^{\ell} U_1^{n_{w_i}(\phi_1)} \right) \cdot \mathbf{y},$$

in the notation of Lemma 6.3. By choosing v_1 sufficiently close to β , we can arrange for t to be arbitrarily close to 0. According to Lemma 6.3 for suitable choice of t , this count is given by

$$(21) \quad \delta_1(\mathbf{x}) = U_1 \cdot \mathbf{x}$$

There are in principle other terms which count homotopy classes with $n_{v_1}(\phi_2) > 1$ (and corresponding to factorizations of ϕ as ϕ_1 and ϕ_2 with $k = n_{w_1}(\phi_1) > 1$, $\mu(\phi_1) = 2$).

Our claim is that for a sufficiently large parameter τ , these homotopy classes have trivial contribution. Here, τ parameterizes the choice of connected sum point $v_1(\tau)$ in S , with the limit $\lim_{\tau \rightarrow \infty} v_1(\tau)$ given by a point v_1^∞ on the curve β . Note that we have already taken τ large to apply Lemma 6.3 in establishing Equation (21).

Suppose now that for a sequence of τ going to infinity, the moduli spaces $\mathcal{M}_\tau(\phi)$ are non-empty for all choice of connected sum neck length. Then, for all sufficiently large τ , the fibered product $\mathcal{M}(\phi_1) \times_{\text{Sym}^k(\mathbb{D})} \mathcal{M}(\phi_2)$ is non-empty, where here $k = n_{v_1}(\phi_2)$.

Thus, we obtain a sequence of pseudo-holomorphic representatives $u_1^\tau \times u_2^\tau$ of the fibered product. Clearly, there are Gromov limits \bar{u}_1^∞ and \bar{u}_2^∞ as $\tau \mapsto \infty$ of the curves u_1^τ and u_2^τ . By dimension counts, since $\mu(u_1^\tau) = 2$, there are only three possible types of limit for \bar{u}_1^∞ : either it is a strong limit to a pseudo-holomorphic disk, or it is a weak limit to a singly-broken flowline, or it contains a boundary degeneration, in which case the remaining component must be a constant flowline. In this latter case, $k = 1$, and hence it has been covered earlier.

Suppose that the limit $\bar{u}_1^\infty = u_1^\infty$ is not a broken flowline, and let u_2^∞ denote the matching component of \bar{u}_2^∞ (ie \bar{u}_2^∞ could *a priori* be a broken flow-line, but it has some component u_2^∞ with the property that $\rho_1(u_1^\infty) = \rho_2(u_2^\infty)$). However, since $v_1(\tau)$ limits to β , we see that $\rho_2(u_2^\infty)$ contains some points on the β -boundary. To achieve this, we must have a sequence u_1^τ with arbitrarily large τ , with $\rho_1(u_1^\tau)$ containing points arbitrarily close to the line $\{0\} \times \mathbb{R}$. According to Lemma 6.3, this forces $k = 1$, a case considered earlier.

In the remaining case, the Gromov limit \bar{u}_1^∞ is given as a broken flow-line $\bar{u}_1^\infty = a \circ b$. Again, by simple dimension counts, we see that $\mu(a) = \mu(b) = 1$. There is also a corresponding $\bar{u}_2^\infty = a' \circ b'$. Since $\mu(\phi_2)$ is odd, we can conclude that so is either $\mu(a')$ or $\mu(b')$. Suppose it is $\mu(a')$ which is odd. It is easy to see that for $(a')^{-1}(v_1^\infty)$ contains points on the boundary $\{0\} \times \mathbb{R}$. Moreover, the same reasoning as before, with a and a' playing the roles of u_1^∞ and u_2^∞ respectively, shows that a is supported in a homotopy class which admits holomorphic representatives a^τ with $\rho_1(a^\tau)$ containing points arbitrarily close to $\{0\} \times \mathbb{R}$. But this is impossible, as w_1 is disjoint from β , and a^τ has to be one of finitely many holomorphic disks up to translation. The case where $\mu(b')$ is odd follows *mutas mutandis*.

Putting the above facts together, we obtain the desired identification of $CF^-(\Sigma, \alpha', \beta', \mathbf{w}')$ with the mapping cone of Equation (20), giving an expression

$$HF^-(\Sigma, \alpha', \beta', \mathbf{w}') \cong \frac{HF^-(\Sigma, \alpha, \beta, \mathbf{w})[U_{\ell+1}]}{U_{\ell+1} - U_1},$$

at least in the case where $\ell > 1$.

Proof of the proposition when $\ell = 1$ In this case, Theorem 5.1 cannot be applied directly to identify the C_y component of C_x .

Specifically, consider a homology class of $\phi = \phi_1 \# \phi_2$, with $\mu(\phi) = 1$, $\phi_2 \in \pi_2(x, y)$, and $n_{v_1}(\phi_2) = 1 = \mu(\phi_2)$. Take a Gromov limits of elements of $\mathcal{M}(\phi)$, as the connected sum is degenerated. Again, this limits to the unique (up to translation) holomorphic representative u_2 of ϕ_2 , but now we cannot assume that the Gromov limit \bar{u}_1 from the other side is a flow-line. More specifically, the dimension counts (cf Equation (16)) which ruled out the possibility of a closed component in the Gromov compactification no longer apply. Indeed, by pushing the connected sum point v_1 sufficiently close to β , we can use Lemma 6.3 to rule out the case where the component of u_1 in the Gromov compactification which matches with u_2 is an actual cylindrical flow-line. Rather, it must be a closed holomorphic curve representing the homology class $[\Sigma_1]$ (with multiplicity one). The argument from stabilization invariance as in [15] applies now to show that $\widehat{\mathcal{M}}(\phi) = 1$. More precisely, the Gromov limit u_1 equals a constant flow-line meeting a copy of Σ_1 , with $\rho_2(u_2) = p$. In this case, gluing can be used to show that $\#\mathcal{M}(\phi) = C \cdot \#\mathcal{M}(\phi_2)$, where here C is the count of representatives $\mathbb{D}\#\Sigma_1$ with a marked point q to $\Sigma_1 \times [0, 1] \times \mathbb{R}$ in the homology class $[\Sigma_1]$ which maps $u(p)$ to (w_1, p) . The fact that $C = 1$, in the case where $g_1 = 1$ is calculated in the proof of stabilization invariance of cylindrical Heegaard Floer homology (cf Appendix B of [15]). One can show that $C = 1$ for arbitrary g follows from this case by, for example, by realizing Σ_1 as a connected sum of g_1 copies of a genus one surface, and degenerating along all the necks. Now, $\mathcal{M}(\phi) = 1$ follows from the elementary calculation that $\#\mathcal{M}(\phi_2) = 1$.

Thus, we have computed that the counts of all Maslov index one moduli spaces of the form $\phi_1 \# \phi_2$, where $\phi_2(x, y)$ is the flowline with $n_{v_1}(\phi_2) = \mu(\phi_2) = 1$ is given by the same map δ_1 as in Equation (21). We can rule out the case where $\mu(\phi_1) = 2$ and $k > 1$ much as before. Specifically, by pushing the connected sum point close to β , Lemma 6.3 shows that the corresponding Gromov limit u_1 must contain at least one closed curve component. Removing this component, we are left with a broken flow-line with Maslov index zero, and hence, a constant flowline, hence showing that we needed to be in the case where $k = 1$. The argument is now completed as before. \square

6.2 Model calculations

The aim of the present subsection is to continue with the methods from the previous subsection to perform some model calculations which will be useful for establishing handleslide invariance.

Definition 6.6 Let C be a chain complex over the polynomial algebra $\mathbb{F}[U_1, \dots, U_\ell]$. We say that it is of $\mathbb{F}[U]$ -type if for $i \neq j$, there are chain homotopies $U_i \simeq U_j$ (thought of as endomorphisms of C).

Of course, if a chain complex C is of $\mathbb{F}[U]$ -type, then its homology is a module over the polynomial algebra $\mathbb{F}[U]$, where U acts by multiplication by any of the U_i .

Let $(\Sigma, \alpha, \beta, \mathbf{w})$ be a Heegaard diagram, and fix two preferred basepoints $w_1, w_2 \in \mathbf{w}$. Let Σ' be a surface obtained by attaching a one-handle in the neighborhood of w_1 and w_2 , and fix a pair of circles α_0 and β_0 which are supported inside the handle, each a small isotopic translate of one another, separating w_1 and w_2 . Equivalently, we form a double-connected sum of Σ along w_1 and w_2 with the sphere S as in Figure 3.

Proposition 6.7 Let $(\Sigma, \alpha, \beta, \mathbf{w})$ be a Heegaard diagram, and let $(\Sigma', \alpha', \beta', \mathbf{w}')$ be the Heegaard diagram obtained by attaching a one-handle in the above sense, so that if the original describes a three-manifold Y , then the second diagram describes $Y' = Y \# (S^2 \times S^1)$. Suppose that $HF^-(\Sigma, \alpha, \beta, \mathbf{w})$ of $\mathbb{F}[U]$ -type, then the same is true of $HF^-(\Sigma', \alpha', \beta', \mathbf{w}')$; and indeed

$$CF^-(\Sigma', \alpha', \beta', \mathbf{w}') \simeq CF^-(\Sigma, \alpha, \beta, \mathbf{w}) \oplus CF^-(\Sigma, \alpha, \beta, \mathbf{w}).$$

Proof We analyze as in Proposition 6.5, only this time stretching near both w_1 and w_2 . In this case, homotopy classes ϕ break as fibered products of homotopy classes ϕ_1 for $(\Sigma, \alpha, \beta, \mathbf{w})$, and ϕ_2 for $(S, \alpha, \beta, v_1, v_2)$. Now we have

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2n_{v_1}(\phi_2) - 2n_{v_2}(\phi_2).$$

The same dimension counts as in the proof of Proposition 6.5 express

$$CF^-(\Sigma', \alpha', \beta', \mathbf{w}')$$

as a mapping cone of a map

$$(22) \quad CF^-(\Sigma, \alpha, \beta, \mathbf{w}) \xrightarrow{U_1 - U_2 + \delta} CF^-(\Sigma, \alpha, \beta, \mathbf{w}).$$

As in the proof of that proposition, the first chain complex is identified with C_x and the second with C_y . The map δ counts those disks which have $n_{v_1}(\phi_2) + n_{v_2}(\phi_2) \geq 2$. By moving the connected sum points as in Proposition 6.5, we can arrange that these terms contribute trivially.

Thus, since $U_1 \simeq U_2$, the chain map from Equation (22) is null-homotopic, and it follows at once that $CF^-(\Sigma', \alpha', \beta', \mathbf{w}')$ is chain homotopic to the direct sum of two copies of $CF^-(\Sigma, \alpha, \beta, \mathbf{w})$. \square

Let Σ be an oriented two-manifold, and α be a collection of attaching circles, and let \mathbf{w} be a collection of basepoints, one in each component of $\Sigma - \alpha_1 - \cdots - \alpha_{g+\ell-1}$. Let α' be a set of attaching circles obtained as small exact Hamiltonian translates of α_i . Specifically α'_i is obtained as an exact Hamiltonian translate of α_i , so that $\alpha_i \cap \alpha'_j = \emptyset$ unless $i = j$, in which case the intersection consists of two points of transverse intersection. Moreover, the Hamiltonian isotopy never crosses any of the $w_i \in \mathbf{w}$. Clearly, $(\Sigma, \alpha, \alpha', \mathbf{w})$ represents $\#^g(S^2 \times S^1)$, where g is the genus of Σ .

Proposition 6.8 $CF^-(\Sigma, \alpha, \alpha', \mathbf{w})$ is of $\mathbb{F}[U]$ -type, and its homology is isomorphic to $\mathbb{F}[U] \otimes \Lambda^* V$, where V is a g -dimensional vector space.

Proof We can reduce to the case where $g = 0$ by repeatedly applying Proposition 6.7. In the case where $g = 0$, the proposition is proved after repeated applications of Proposition 6.5. \square

As an example, suppose that $(\Sigma, \alpha, \beta, \mathbf{w})$ is an admissible multi-pointed Heegaard diagram, and suppose that β' is obtained from β by a small perturbation as in Proposition 6.8. According to Proposition 6.8, there is an element $[\Theta_{\beta\beta'}] \in HF^-(\Sigma, \beta, \beta', \mathbf{w})$ of maximal degree.

We can define a corresponding map

$$\Phi_{\alpha\beta\beta'}: CF^-(\Sigma, \alpha, \beta, \mathbf{w}) \longrightarrow CF^-(\Sigma, \alpha, \beta', \mathbf{w})$$

$$\text{by: } \Phi_{\alpha\beta\beta'}(\xi) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta\beta'}, \mathbf{y})\}} \#\mathcal{M}(\psi) \left(\prod_{i=1}^{\ell} U_i^{n_{w_i}(\psi)} \right) \cdot \mathbf{y}$$

This map is a chain map (compare [21, Section 8]). (In this notation, we implicitly assume that $[\Theta_{\beta\beta'}]$ is represented by a single intersection point $\Theta_{\beta\beta'}$; more generally, our map is gotten by summing triangle maps over the various intersection points whose sum represents the homology class.)

Proposition 6.9 If the curves in β' are sufficiently close to those in β , chosen so that each β'_i meets β_i in precisely two intersection points, then the map $\Phi_{\alpha\beta\beta'}$ defined above induces an isomorphism in homology.

Proof This follows as in [21, Proposition 9.8]. The point is that for each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, there is a corresponding nearest point $\mathbf{x}' \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$, and also a corresponding small triangle $\psi \in \pi_2(\mathbf{x}, \Theta_{\beta\beta'}, \mathbf{y})$. Thus, the map obtained by counting only these smallest triangles induces an isomorphism of chain groups. Using the energy filtration, it follows that $\Phi_{\alpha\beta\beta'}$ is an isomorphism of chain complexes. \square

6.3 Handleslide invariance

We can now adapt the proof of handleslide invariance of Heegaard Floer homology as in [21, Section 9] to establish handleslide invariance of HF^- in the present context.

Specifically, start with a Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w})$. Let γ be obtained from β by a single handleslide.

Proposition 6.10 *There is an identification $HF^-(\Sigma, \alpha, \beta, \mathbf{w}) \cong HF^-(\Sigma, \alpha, \gamma, \mathbf{w})$.*

Proof As a first step, we claim that $HF^-(\Sigma, \beta, \gamma, \mathbf{w}) \cong \mathbb{F}[U] \otimes \Lambda^* V$, where V is a g -dimensional \mathbb{F} -vector space. To see this, de-stabilize using Propositions 6.5 and 6.7 to the case where there are exactly two β curves and two γ curves. There are three cases, according to whether $g = g(\Sigma)$ is 0, 1, or 2. The case where $g = 2$ is established in [21]; the other two cases are easily established by the same calculation.

With this said, there is a canonical top-dimensional generator $\Theta_{\beta\gamma}$ of

$$HF^-(\Sigma, \beta, \gamma, \mathbf{w}).$$

The handleslide map

$$\Psi_{\alpha\beta\gamma}: HF^-(\Sigma, \alpha, \beta, \mathbf{w}) \longrightarrow HF^-(\Sigma, \alpha, \gamma, \mathbf{w})$$

is defined by counting pseudo-holomorphic triangles: $\Psi_{\alpha\beta\gamma} = f_{\alpha\beta\gamma}(\cdot \otimes \Theta_{\beta\gamma})$; ie here

$$f_{\alpha\beta\gamma}(\mathbf{x} \otimes \Theta_{\beta\gamma}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta\gamma}, \mathbf{y}) \mid \mu(\psi) = 0\}} \#\mathcal{M}(\psi) \cdot \left(\prod_{i=1}^{\ell} U_i^{n_{w_i}(\psi)} \right) \cdot \mathbf{y}.$$

Now, according to associativity of the triangle maps, we see that the composite

$$\Psi_{\alpha\gamma\beta'} \circ \Psi_{\alpha\beta\gamma} = f_{\alpha\beta\beta'}(\cdot \otimes f_{\beta\gamma\beta'}(\Theta_{\beta\gamma} \otimes \Theta_{\gamma\beta'})).$$

Next, we verify that $f_{\beta\gamma\beta'}(\Theta_{\beta\gamma} \otimes \Theta_{\gamma\beta'}) = \Theta_{\beta\beta'}$ is the canonical top-dimensional generator of $HF^-(\Sigma, \beta, \beta', \mathbf{w})$, which is calculated in Proposition 6.8; ie we obtain the map $\Phi_{\alpha\beta\beta'}$ studied in Proposition 6.9. The map is an isomorphism now according to that proposition. \square

6.4 Invariance

We prove the invariance of HF^- as introduced in Section 4.

Proof of Theorem 4.4 We show that HF^- is invariant under the four Heegaard moves from Proposition 3.3, or, more specifically, their admissible versions as in Proposition 3.6.

Isotopy invariance follows exactly as in [21, Section 7].

Handleslide invariance was established in Proposition 6.10.

Invariance under index one and two stabilizations follows from exactly as in [21, Section 10].

Finally, it remains to consider invariance index zero/three stabilizations. In this case, we introduce a new pair of (isotopic) curves $\alpha_{g+\ell}$ and $\beta_{g+\ell}$, each of which separates Σ in two. In the case where the stabilization is simple in the sense of Section 6.1, ie $\alpha_{g+\ell}$ bounds a disk containing $w_{\ell+1}$ and none of the other α_i (or w_i), invariance is established in Proposition 6.5. It is easy to see that general index zero/three stabilizations can be achieved by simple index zero/three stabilizations, followed by a sequence of handleslides.

The case of \widehat{CF} can be handled more quickly:

Proof of Theorem 4.5 Let $(\Sigma, \alpha, \beta, \mathbf{w})$ be the multi-pointed Heegaard diagram, with $\mathbf{w} = \{w_1, \dots, w_\ell\}$. Let P be a planar surface with ℓ boundary components. Form the surface $\Sigma' = (\Sigma - \text{nd}(\mathbf{w})) \cup P$. Consider $(\Sigma', \alpha, \beta, w)$, where here w is chosen in the region $P \subset \Sigma'$. It is easy to see that $(\Sigma', \alpha, \beta, w)$ is a pointed Heegaard diagram for $Y\#(\#^{\ell-1}(S^2 \times S^1))$. Admissibility of the original diagram corresponds precisely to admissibility for this new diagram. Moreover, the intersection points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ in $\text{Sym}^{g+\ell-1}(\Sigma)$ are identical with those in $\text{Sym}^{g+\ell-1}(\Sigma')$. Finally, holomorphic curves in $\text{Sym}^{g+\ell-1}(\Sigma)$ which are disjoint from all the w_i are identical with those in $\text{Sym}^{g+\ell-1}(\Sigma')$ which are disjoint from w . The result now follows.

7 Invariance of link invariants

Invariance of the filtered chain homotopy type of the link filtration $\text{CFL}^-(S^3, \vec{L})$ is an easy consequence of the methods from Section 6.

Proof of Theorem 4.7 We must verify that $\text{CFL}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is invariant under isotopies and handleslides supported in the complement of \mathbf{w} and \mathbf{z} , and also index one/two stabilizations.

If $(\Sigma', \alpha', \beta', \mathbf{w}', \mathbf{z}')$ is obtained from an index one/two stabilization, then the methods from [21, Section 10] actually give an isomorphism of chain complexes

$$\text{CFL}^-(\Sigma', \alpha', \beta', \mathbf{w}', \mathbf{z}') \cong \text{CFL}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}).$$

Performing the stabilization away from \mathbf{w} and \mathbf{z} , we see that the assignment to relative Spin^c structures is unaffected.

Handleslides amongst the β are understood as follows. As before, if γ is obtained from β by a handleslide, then there is a chain map (cf Proposition 6.10)

$$\Psi_{\alpha\beta\gamma}: \text{CFL}^-(\Sigma, \alpha, \beta, \mathbf{w}) \longrightarrow \text{CFL}^-(\Sigma, \alpha, \gamma, \mathbf{w}),$$

defined by counting holomorphic triangles. Since all triply-periodic domains are disjoint from \mathbf{w} and \mathbf{z} , it follows easily that if $\mathbf{x} \leq \mathbf{y}$, then $\Psi_{\alpha,\beta,\gamma}(\mathbf{x}) \leq \Psi_{\alpha,\beta,\gamma}(\mathbf{y})$. We thus consider the induced map on the associated graded object, ie the map $\widehat{\Psi}_{\alpha\beta\gamma}$ induced by counting pseudo-holomorphic triangles which are disjoint from \mathbf{w} and \mathbf{z} alike. Note that

$$\underline{\mathfrak{s}}_{\mathbf{w},\mathbf{z}}(\mathbf{x}) = \underline{\mathfrak{s}}'_{\mathbf{w},\mathbf{z}}(\widehat{\Psi}_{\alpha\beta\gamma}(\mathbf{x})),$$

where here $\underline{\mathfrak{s}}'_{\mathbf{w},\mathbf{z}}$ is the map associated to the Heegaard diagram $(\Sigma, \alpha, \gamma, \mathbf{w}, \mathbf{z})$, since the triangles used in the definition of $\widehat{\Psi}$ this count are disjoint from \mathbf{w} and \mathbf{z} . Now, adapting the argument from Proposition 6.10, we have that the map

$$\widehat{\Psi}_{\alpha\gamma\beta'} \circ \widehat{\Psi}_{\alpha\beta\gamma} = \widehat{f}_{\alpha\beta\beta'}(\cdot \otimes \Theta_{\beta\beta'})$$

is an isomorphism. Specifically, here

$$\widehat{f}_{\alpha\beta\beta'}(\mathbf{x} \otimes \Theta_{\beta\beta'}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta\beta'}, \mathbf{y}) \mid \mu(\psi) = 0, n_{\mathbf{w}}(\psi) = n_{\mathbf{z}}(\psi) = 0\}} \#\mathcal{M}(\psi)\mathbf{y}.$$

The fact that this is an isomorphism follows readily as in Proposition 6.9, with the observation that the small triangles considered there do not contain any of the the basepoints $\{w_i, z_i\}_{i=1}^\ell$. Since $\widehat{\Psi}$ induces an isomorphism of $\text{Spin}^c(Y, L)$ -graded complexes, it follows formally that $\Psi_{\alpha\beta\gamma}$ induces an isomorphism of $\text{Spin}^c(Y, L)$ -filtered complexes.

7.1 Link homology as an associated graded object

We consider now the link filtration.

Proof of Theorem 1.2 From its construction, we see that $\widehat{\text{HFL}}$ is the homology of the graded object associated to a \mathbb{Z}^ℓ -filtration of $\widehat{\text{CF}}(\Sigma, \alpha, \beta, \mathbf{w})$, which, according to Theorem 4.5, calculates $\widehat{\text{HF}}(\#^{\ell-1}(S^2 \times S^1))$. But $\widehat{\text{HF}}(\#^{\ell-1}(S^2 \times S^1))$ is identified with the exterior algebra of $H_1(\#^{\ell-1}(S^2 \times S^1); \mathbb{F})$ (cf [21, Lemma 9.1]). Invariance of the spectral sequence is a consequence of Theorem 4.7.

Proof of Theorem 1.5 Fix $\underline{s} \in \text{Spin}^c(Y, L)$. Recall that the chain complex computing $\text{HFL}^-(Y, \vec{L}, \underline{s})$ consists of symbols $U^{a_1} \cdots U^{a_\ell} \otimes \mathbf{x}$, where a_1, \dots, a_ℓ are non-negative integers and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, satisfying the constraint that

$$\underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) - a_1 \text{PD}[\mu_1] - \cdots - a_\ell \text{PD}[\mu_\ell] = \underline{s}.$$

This set is endowed with the differential as in Equation (10), except that now we sum over only those $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $n_{\mathbf{z}}(\phi) = 0$. The further filtration on $\widehat{\text{CFL}}(Y, \vec{L}, \underline{s})$ is given by the map $U^{a_1} \cdots U^{a_\ell} \otimes \mathbf{x} \mapsto (a_1, \dots, a_\ell)$. Its associated graded object is clearly the one stated in the theorem.

Proof of Theorem 1.6 The spectral sequence arises from the filtration of

$$\text{CFL}^-(S^3, \vec{L}).$$

The fact that the homology of this total complex is $\mathbb{Z}[U]$ follows from the fact that the underlying chain complex associated to $(\Sigma, \alpha, \beta, \mathbf{w})$ has homology $\text{HF}^-(S^3)$ (according to Theorem 4.4), which in turn is known to be $\mathbb{Z}[U]$ (cf [21]).

7.2 Forgetful functors

Let \vec{L} be an ℓ -component oriented link, which we assume to be in S^3 for simplicity. Distinguish a component $K_1 \subset L$, and consider the filling map

$$G_{K_1}: \text{Spin}^c(Y, L) \longrightarrow \text{Spin}^c(Y, L - K_1)$$

studied in Section 3.7. Under this map, a $\text{Spin}^c(Y, L)$ -filtered chain complex can be viewed as a $\text{Spin}^c(Y, L - K_1)$ -filtered chain complex.

Let M be the rank two graded vector space with one generator in grading 0 and another in grading -1 .

Proposition 7.1 *Let \vec{L} be an oriented, ℓ -component link in S^3 , and distinguish the first component K_1 . Consider the filtration $\widehat{\text{CFL}}(\vec{L})$, viewed as a $\text{Spin}^c(Y, L - K_1)$ -filtered chain complex, via the filling map using the distinguished component K_1 of L . The filtered chain homotopy type of this complex is identified with $\widehat{\text{CFL}}(\vec{L} - K_1) \otimes M$.*

In more elementary terms, a \mathbb{Z}^ℓ -filtered chain complex can be viewed as a $\mathbb{Z}^{\ell-1}$ -filtered one, by forgetting the first term in the relative filtration. The above proposition says that for the relatively \mathbb{Z}^ℓ -filtered chain complex $\widehat{\text{CFL}}(Y, \vec{L})$, if we forget the first term in the relative filtration, we obtain the relatively $\mathbb{Z}^{\ell-1}$ -filtered complex $\widehat{\text{CFL}}(\vec{L} - K_1) \otimes M$.

In particular, if we consider the higher differentials inducing a chain complex on $\widehat{\text{HFL}}(\vec{L})$ (from Theorem 1.2) and take their homology in the first component (in the sense of Section 2.1), we obtain $\widehat{\text{HFL}}(\vec{L}') \otimes M$ (up to some overall translation of the grading by relative Spin^c structures).

Proof of Proposition 7.1 We begin with the relative \mathbb{Z}^ℓ -filtered statement.

In view of Propositions 6.5 and 6.7, and keeping track of filtrations, we see that the $\mathbb{Z}^{\ell-1}$ -filtered chain complex obtained from $\text{CFL}^-(S^3, \vec{L})$ obtained by forgetting about the first factor is identified with the mapping cone of

$$\text{CFL}^-(S^3, \vec{L}') [U_1] \xrightarrow{U_1 - U_i} \text{CFL}^-(S^3, \vec{L}') [U_1],$$

where $\vec{L}' = \vec{L} - K_1$ and i is some integer > 1 . This identification is gotten by reducing to the case of simple stabilizations (in the sense of Section 6.1), after handlesliding across the basepoint z_1 .

Specializing to $U_j \equiv 0$ for all j , we obtain two copies of $\widehat{\text{CFL}}(S^3, \vec{L}')$, verifying the simplified statement of Proposition 7.1, using only the relative \mathbb{Z}^ℓ -filtered statement.

To obtain the version stated here, we must add the following observations. If

$$(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$$

is a multiply-pointed Heegaard diagram for \vec{L} , and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, then the map obtained by handlesliding across z_1 still preserves $G_{K_1}(\underline{\mathfrak{s}}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}))$. Moreover, in the identification of Proposition 6.5 between the complex obtained from applying the forgetful map to $\widehat{\text{CFL}}(S^3, L)$ and the complex $\widehat{\text{CFL}}(S^3, L - K_1) \otimes M$, we have an identification between the induced relative Spin^c structures.

8 Symmetry

We consider the symmetry properties of $\widehat{\text{HFL}}(\vec{L})$. The properties we give here are formally analogous to corresponding properties of knot Floer homology, cf [20; 24; 21].

Proposition 8.1 *Let \vec{L} be an oriented link. Then, we have an identification*

$$\widehat{\text{HFL}}_*(\vec{L}, \underline{\mathfrak{s}}) \cong \widehat{\text{HFL}}_*(-\vec{L}, J\underline{\mathfrak{s}}).$$

Proof Suppose $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is a pointed Heegaard diagram representing \vec{L} . Then, $(-\Sigma, \beta, \alpha, \mathbf{w}, \mathbf{z})$ is also a pointed Heegaard diagram for the link L , only now it represents the opposite orientation $-\vec{L}$. The two Heegaard diagrams give assignments

$$\underline{s}_{\mathbf{w}, \mathbf{z}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \longrightarrow \underline{\text{Spin}}^c(Y, L) \quad \text{and} \quad \underline{s}'_{\mathbf{w}, \mathbf{z}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \longrightarrow \underline{\text{Spin}}^c(Y, L),$$

respectively. Let \widehat{CF} denote the chain complex whose homology is \widehat{HF} , ie this is the associated graded object for the link filtration of \widehat{CF} . Explicitly, its generators are $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and it is endowed with the differential from Equation (14).

Of course, the generators for $\widehat{CF}(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ and $\widehat{CF}(-\Sigma, \beta, \alpha, \mathbf{w}, \mathbf{z})$ coincide. Indeed, by pre-composing with a reflection on the disk, we see that the differentials for the two chain complexes coincide.

However, $\underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) = J \underline{s}'_{\mathbf{w}, \mathbf{z}}(\mathbf{x})$, as in Lemma 3.12. The result follows. \square

Proposition 8.2 *Let \vec{L} be an oriented link in S^3 . Let $f: \underline{\text{Spin}}^c(Y, L) \longrightarrow \mathbb{Z}$ be the map defined as follows. Writing*

$$c_1(\underline{s}) = \sum_{i=1}^{\ell} a_i \cdot \text{PD}[\mu_i],$$

we have that $f(\underline{s}) = \sum_{i=1}^{\ell} a_i$. Then, we have an identification

$$\widehat{\text{HFL}}_d(\vec{L}, \underline{s}) \cong \widehat{\text{HFL}}_{d-f(\underline{s})+\ell}(\vec{L}, J(\underline{s}) + \sum_{i=1}^{\ell} \text{PD}[\mu_i]).$$

Proof Write $\mu = \sum_{i=1}^{\ell} \mu_i$. As in Lemma 3.12, we see that $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ and $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$ are Heegaard diagrams for \vec{L} and $-\vec{L}$ respectively. The generators of \widehat{CF} and the differentials coincide, but the two maps to relative Spin^c structures differ, with $\underline{s}_{\mathbf{w}, \mathbf{z}}(\mathbf{x}) = \underline{s}_{\mathbf{z}, \mathbf{w}}(\mathbf{x}) + \text{PD}[\mu]$. Thus, we see that there is a function $g: \underline{\text{Spin}}^c(Y, L) \longrightarrow \mathbb{Z}$ with the property that

$$\widehat{\text{HFL}}_d(\vec{L}, \underline{s}) \cong \widehat{\text{HFL}}_{d+g(\underline{s})}(-\vec{L}, \underline{s} - \mu).$$

Suppose that $f(\underline{s}_1) - f(\underline{s}_2) = k$, then there is a flowline $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\underline{s}(\mathbf{x}) = \underline{s}_1$ and $\underline{s}(\mathbf{y}) = \underline{s}_2$, $n_{\mathbf{w}}(\phi) = 0$, $n_{\mathbf{z}}(\phi) = k$. Now, in $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$, it follows that $\text{gr}(\mathbf{x}) - \text{gr}(\mathbf{y}) = \mu(\phi)$, whereas for $(\Sigma, \alpha, \beta, \mathbf{z}, \mathbf{w})$, $\text{gr}'(\mathbf{x}) - \text{gr}'(\mathbf{y}) = \mu(\phi) - 2k$ (where here we use gr' to denote the grading induced on the second diagram). It follows at once that there is some constant $c = c(\vec{L})$ (depending on the oriented link) with the property that for all $\underline{s} \in \underline{\text{Spin}}^c(Y, L)$, $g(\underline{s}) = -f(\underline{s}) + c(\vec{L})$.

Combining the above remarks with Proposition 8.1, we have that

$$(23) \quad \widehat{\text{HFL}}_d(\vec{L}, \underline{s}) = \widehat{\text{HFL}}_{d-f(\underline{s})+c(\vec{L})}(\vec{L}, J(\underline{s}) - \mu).$$

On the other hand, two applications of Equation (23) shows that

$$\widehat{\text{HFL}}_d(\vec{L}, \underline{s}) = \widehat{\text{HFL}}_{d-f(\underline{s})+c(\vec{L})-f(J(\underline{s}+\mu))+c(\vec{L})}(\vec{L}, \underline{s}),$$

and hence, since $f(J(\underline{s} + \mu)) = -f(\underline{s}) + 2n$ that $c(\vec{L}) = n$. □

There is a natural map

$$i: \text{Spin}^c(Y, L) \longrightarrow \text{Spin}^c(-Y, L),$$

since the notion of equivalence classes of nowhere vanishing vector fields does not refer to the orientation of the ambient manifold. Note, however, that $c_1(i(\underline{s})) = -c_1(J(\underline{s}))$, since the induced two-plane field on \vec{v}^\perp gets opposite orientations depending on the orientation of Y .

Proposition 8.3 *Let \vec{L} be an oriented link, and let $r(\vec{L})$ denote its mirror. Then,*

$$\widehat{\text{HFL}}_d(\vec{L}, \underline{s}) = \widehat{\text{HFL}}^{-d+\ell-1}(r(\vec{L}), i(\underline{s})).$$

Proof If $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ represents \vec{L} , we can think of $(-\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ as representing its mirror. It is easy to see that this latter chain complex is the dual complex for the former, ie

$$\widehat{CF}_d(-\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}) \cong \text{Hom}(\widehat{CF}_{-d}(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})),$$

so that the assignment from the intersection points to relative Spin^c structures is intertwined with the natural map i referred to above. The shift in grading follows from the fact that both total complexes for \widehat{CF} have homology isomorphic to $\Lambda^*\mathbb{F}^{\ell-1}$, graded so that its top-most term has grading 0 (cf Theorem 1.2). □

8.1 Notational remarks

Relative Spin^c structures are not very concrete objects, and hence it is sometimes awkward to describe link Floer homology as filtered by them. We could alternatively use relative two-dimensional homology, and indeed the above results suggest the following convention.

For $\kappa \in H^2(Y, L)$, we could define $\widehat{\text{HFL}}(\vec{L}, \kappa)$ to be $\widehat{\text{HFL}}(\vec{L}, \underline{\mathfrak{s}})$, where here $\underline{\mathfrak{s}}$ is the relative Spin^c structure determined by the formula

$$c_1(\underline{\mathfrak{s}}) + \sum_{i=1}^{\ell} \text{PD}[\mu_i] = \kappa.$$

Given $\kappa \in H^2(Y, L)$, write

$$\kappa = \sum_{i=1}^{\ell} a_i \cdot \text{PD}[\mu_i].$$

We let $|\kappa|$ be the quantity

$$|\kappa| = \sum_{i=1}^{\ell} a_i.$$

With these conventions, then, Proposition 8.2 reads

$$\widehat{\text{HFL}}_d(\vec{L}, \kappa) = \widehat{\text{HFL}}_{d-|\kappa|}(\vec{L}, -\kappa).$$

Also, according to Proposition 7.1,

$$\widehat{\text{HFL}}(\vec{L} - K_1, \kappa + \text{PD}[K_1])$$

is gotten by introducing extra differentials on $\bigoplus_{s \in \mathbb{Z}} \widehat{\text{HFL}}(\vec{L}, \kappa + 2s \cdot \text{PD}[\mu_1])$.

It is simpler yet to consider $\widehat{\text{HFL}}(\vec{L})$ as graded by elements of $\mathbb{H}(L)$; ie given $h \in \mathbb{H}(L)$, write $\widehat{\text{HFL}}(\vec{L}, h) = \widehat{\text{HFL}}(\vec{L}, \underline{\mathfrak{s}})$, where $\underline{\mathfrak{s}}$ is chosen so that

$$(24) \quad c_1(\underline{\mathfrak{s}}) + \sum_{i=1}^{\ell} \text{PD}[\mu_i] = 2 \cdot \text{PD}[h].$$

Now, Proposition 8.2 reads

$$(25) \quad \widehat{\text{HFL}}_d(\vec{L}, h) = \widehat{\text{HFL}}_{d-|h|}(\vec{L}, -h).$$

Note that we could have had the grading set be inside H , rather than \mathbb{H} , by dropping the factor of 2 in Equation (24). However, we have chosen the present formula so that $\widehat{\text{HFL}}(K, s \cdot \text{PD}[\mu]) = \widehat{\text{HFK}}(K, s)$, in the case where K is a (one-component) knot.

9 Euler characteristics

The aim of this section is to calculate the Euler characteristic of link homology, establishing Equations (1) and (2).

9.1 The case of $\widehat{\text{HFL}}$

Let X be a connected, m -dimensional CW complex, and let $H = H_1(X; \mathbb{Z})/\text{Tors}$. The Reidemeister–Franz torsion of X is defined as follows (see for example [17; 28]). Let \mathbb{K} denote the field of fractions of the group-ring $\mathbb{Z}[H]$. Let \widetilde{X} denote the covering space determined by the natural homomorphism of $\pi_1(X, x) \rightarrow H$. The action by H gives the CW complex of \widetilde{X} , $C_*(\widetilde{X}; \mathbb{Z})$ a complex of free $\mathbb{Z}[H]$ -modules. In particular, its differentials

$$\widetilde{\partial}_i: C_i(\widetilde{X}; \mathbb{Z}) \rightarrow C_{i-1}(\widetilde{X}; \mathbb{Z})$$

are all $\mathbb{Z}[U]$ -equivariant. A *fundamental family of cells* for \widetilde{X} is a collection of cells for \widetilde{X} , each of which projects to exactly one cell of X . The Reidemeister–Franz torsion is defined to vanish if $H_*(C_*(\widetilde{X}) \otimes_{\mathbb{Z}[H]} \mathbb{K}) \neq 0$. Otherwise, we proceed as follows. Suppose that c and c' are two ordered bases for some \mathbb{K} vector space. Then, their top exterior products differ by an element which we denote by $[c/c']$ (ie $\det(c) = [c/c'] \cdot \det(c')$). Let b_i be a sequence of vectors in $C_i(\widetilde{X}; \mathbb{Z})$ whose image under $\widetilde{\delta}_i$ forms a basis for this image. By our assumption, $\widetilde{\partial}_i(b_{i+1})b_i$ is a basis for $C_i(\widetilde{X}; \mathbb{Z})$; let c_i be a basis for $C_i(\widetilde{X}; \mathbb{Z})$ coming from a fundamental family of cells for \widetilde{X} . Then, the Reidemeister–Franz torsion $t(X) \in \frac{\mathbb{K}}{\pm H}$ is defined to be

$$t(X) = \prod_{i=0}^m \left[\frac{\widetilde{\partial}_i(b_{i+1})b_i}{c_i} \right]^{(-1)^{i+1}}.$$

In the case where X is the complement of an ℓ -component link $L \subset S^3$, $H \cong \mathbb{Z}^\ell$, and $t(X) \in \mathbb{Z}[H]$. Indeed, in this case the Reidemeister–Franz torsion, which we now denote by $t(L)$, is related to the (multi-variable) Alexander polynomial by the formula

$$t(L) = \begin{cases} \frac{\Delta_L(T)}{(T-1)} & \text{if } \ell = 1, \\ \Delta_L(T_1, \dots, T_\ell) & \text{if } \ell > 1. \end{cases}$$

Although this defines $t(L)$ only up to an overall sign and translation by elements of H , the latter indeterminacy can be resolved by taking the representative of $t(L)$ which satisfies the symmetry $t(L) = \bar{t}(L)$, where here $f \in \mathbb{Z}[H] \mapsto \bar{f}$ is the map induced by taking h to h^{-1} . (The existence of such a symmetric representative is standard in the theory of torsion, cf [17; 28].)

Proposition 9.1 *Given a link L , we have an identification*

$$(26) \quad \sum_{h \in \mathbb{H}} \chi(\widehat{\text{HFL}}_*(L, h)) \cdot e^h = \pm \left(\prod_{i=1}^{\ell} (T_i^{\frac{1}{2}} - T_i^{-\frac{1}{2}}) \right) \cdot t(L).$$

Here, we think of $T_i = e^{m_i}$, where m_i are meridians for L .

Proof Let $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ be a 2ℓ -pointed Heegaard diagram for S^3 subordinate to $L \subset S^3$.

By stabilizing if necessary, we can arrange for the following conditions to hold:

- The genus of Σ is g , with $\beta_1, \dots, \beta_\ell$ forming meridians for the various components of L .
- There are circles $\alpha_1, \dots, \alpha_\ell$, with the property that $\alpha_i \cap \beta_j$ for $i = 1, \dots, \ell$ and $j = 1, \dots, g$ is empty unless $i = j$, in which case α_i and β_i meet in a single point.
- w_i and z_i can be connected by an arc δ_i which is disjoint from all the α_j with $j = 1, \dots, g$, and it is also disjoint from all the β_j with $j \neq i$, meeting β_i transversally in a single point.

Let a_i be the one-handle in S^3 corresponding to the circle α_i , and b_j be the two-handle corresponding to the attaching circle β_j . The zero-handles A_1, \dots, A_ℓ , the one-handles $a_1, \dots, a_{g+\ell}$, and the two-handles $b_{\ell+1}, \dots, b_{\ell+g}$ together form a handle decomposition for $S^3 - L$. Let \tilde{A}_i, \tilde{a}_j , and \tilde{b}_k (with $i = 1, \dots, \ell$, $j = 1, \dots, g + \ell$, $k = \ell + 1, \dots, \ell + g$) denote a collection of lifts to the cover \tilde{M} . Note that we can label $\{\tilde{A}_i\}_{i=1}^\ell$ so that

$$(27) \quad \tilde{\partial} \tilde{a}_i = (1 - T_i) \tilde{A}_i.$$

for all $i = 1, \dots, \ell$.

Now,

$$t(\tilde{M}) = \left[\frac{\{\tilde{\partial} \tilde{a}_i\}_{i=1}^\ell}{\{\tilde{A}_i\}} \right]^{-1} \cdot \left[\frac{\{\tilde{\partial} \tilde{b}_k\}_{k=\ell+1}^{\ell+g} \{\tilde{a}_i\}_{i=1}^\ell}{\{\tilde{a}_j\}_{j=1}^{\ell+g}} \right] \cdot \left[\frac{\{\tilde{b}_k\}_{k=\ell+1}^{\ell+g}}{\{\tilde{b}_k\}_{k=\ell+1}^{\ell+g}} \right]^{-1}.$$

Of course,

$$\left[\frac{\{\tilde{b}_k\}_{k=\ell+1}^{\ell+g}}{\{\tilde{b}_k\}_{k=\ell+1}^{\ell+g}} \right] = 1.$$

Moreover, it is not difficult to see that

$$\left[\frac{\{\tilde{\partial} \tilde{b}_k\}_{k=\ell+1}^{\ell+g} \{\tilde{a}_i\}_{i=1}^\ell}{\{\tilde{a}_j\}_{j=1}^{\ell+g}} \right] = \left[\frac{\{\partial b_k\}_{k=\ell+1}^{\ell+g}}{\{a_k\}_{k=\ell+1}^{\ell+g}} \right] = \#(\mathbb{T}_\alpha \cap \mathbb{T}_\beta).$$

Finally, according to Equation (27),

$$\left[\frac{\{\widetilde{\partial a}_i\}_{i=1}^\ell}{\{A_i\}} \right] = \prod_{i=1}^\ell (1 - T_i).$$

Observe next that for any $\underline{s} \in \underline{\text{Spin}}^c(Y, L)$, we have that

$$\sum_{\underline{s} \in \underline{\text{Spin}}^c(Y, L)} \chi(\widehat{\text{HFL}}(L, \underline{s}_0 + h)) e^h \doteq \#(\widetilde{\mathbb{T}}_\alpha \cap \widetilde{\mathbb{T}}_\beta).$$

This establishes Equation (26), up to an overall translation by units in the Laurent polynomials in T_1, \dots, T_ℓ . The indeterminacy in translation by the various T_i is resolved by observing that both sides of the equation are symmetric under the involution of the ring of Laurent polynomials induced by $T_i \mapsto T_i^{-1}$ (for all i): The left hand side is invariant according to Proposition 8.2, while the right hand side is symmetric by basic properties of torsion [17]. \square

9.2 The case of HFL^-

Proposition 9.2 *We have an identification*

$$\sum_{h \in \mathbb{H}} \chi(\text{HFL}^-(L, h)) \cdot e^h = \pm t(L).$$

Proof Each generator $\mathbf{x} \in \widehat{\text{HFL}}(L)$ gives rise to infinitely many generators

$$T_1^{-a_1} \cdot \dots \cdot T_\ell^{-a_\ell} \cdot \mathbf{x}$$

for all $(a_1, \dots, a_\ell) \geq 0$, with $\mathcal{F}(e^h \cdot \mathbf{x}) = \mathcal{F}(\mathbf{x}) - h$. Thus, it follows that

$$\sum_{h \in \mathbb{H}} \chi(\text{HFL}^-(L, h)) \cdot e^h = \left(\sum_{\{(a_1, \dots, a_\ell) \in \mathbb{Z}^\ell \mid a \geq 0\}} T_1^{-a_1} \cdot \dots \cdot T_\ell^{-a_\ell} \right) \cdot \left(\sum_{h \in \mathbb{H}} \chi(\widehat{\text{HFL}}(L, h)) \cdot e^h \right).$$

Since

$$\sum_{\{(a_1, \dots, a_\ell) \in \mathbb{Z}^\ell \mid a \geq 0\}} T_1^{-a_1} \cdot \dots \cdot T_\ell^{-a_\ell} = \frac{1}{\prod_{i=1}^\ell (1 - T_i^{-1})},$$

the proposition now follows from Proposition 9.1. \square

10 Relationship with knot Floer homology

Our aim here is to show that certain aspects of the link Floer homology considered here can be extracted from the knot Floer homology of links considered in [20]. Specifically, we prove Theorem 1.1, which shows in particular that the total rank of $\widehat{\text{HFL}}$ agrees with the total rank of the knot Floer homology for an oriented link. From this, we quickly deduce the link Floer homology of alternating links, in Section 10.1.

First, we briefly recall the construction from [20]. Suppose that L is an ℓ -component oriented link in S^3 . Then, we can construct an oriented knot in $\#^{\ell-1}(S^2 \times S^1)$ as follows. Fix $2\ell - 2$ points $\{p_i, q_i\}_{i=1}^{\ell-1}$ in \vec{L} which are paired off in such a manner that if we formally identify each p_i with q_i in L , we obtain a connected graph. We view p_i and q_i as the feet of a one-handle to attach to S^3 . Attaching all $\ell - 1$ of these one-handles, we obtain a three-manifold $\kappa(S^3, \{p_i, q_i\})$ which is diffeomorphic to $\#^{\ell-1}(S^2 \times S^1)$. Inside each one-handle, we can find a band along which to perform a connected sum of the component of \vec{L} containing p_i with the component containing q_i . We choose the band so that the induced orientation of its boundary is compatible with the orientation of the link (moreover the band is chosen to be always transverse to the foliation of the one-handle by two-spheres). Our hypotheses on the number and distribution of the distinguished points ensures that the newly-constructed link gotten by performing all $\ell - 1$ of the connected sums is in fact a single-component knot. We denote this link by $\kappa(\vec{L}, \{p_i, q_i\})$ inside $\kappa(S^3, \{p_i, q_i\}) \cong \#^{\ell-1}(S^2 \times S^1)$. It is not difficult to see that the diffeomorphism type of the pair $(\kappa(S^3, \{p_i, q_i\}), \kappa(\vec{L}, \{p_i, q_i\}))$ depends on only the underlying oriented link \vec{L} [20, Proposition 2.1]; hence we denote this object simply by the knot $\kappa(\vec{L}) \subset \#^{\ell-1}(S^2 \times S^1)$.

With these preliminaries, the *knot Floer homology groups*

$$\widehat{\text{HFK}}(\vec{L}) = \bigoplus_{s \in \mathbb{Z}} \widehat{\text{HFK}}(\vec{L}, s)$$

of an oriented link are defined to be the knot Floer homology groups of $\kappa(\vec{L}) \subset \#^{\ell-1}(S^2 \times S^1)$:

$$\widehat{\text{HFK}}(\#^{\ell-1}(S^2 \times S^1), \kappa(\vec{L})) = \bigoplus_{s \in \mathbb{Z}} \widehat{\text{HFK}}(\#^{\ell-1}(S^2 \times S^1), \kappa(\vec{L}), s).$$

In our proof of Theorem 1.1, it will be useful to pass from a Heegaard diagram for a link $\vec{L} \subset S^3$ to a corresponding Heegaard diagram for $\kappa(\vec{L}) \subset \#^{\ell-1}(S^2 \times S^1)$. This is done as follows. Fix a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for an oriented, ℓ -component link $\vec{L} \subset S^3$. Write $\mathbf{w} = \{w_1, \dots, w_\ell\}$, $\mathbf{z} = \{z_1, \dots, z_\ell\}$. We attach $\ell - 1$ one-handles to Σ

to obtain a new surface Σ' as follows. The feet of the i^{th} one-handle is attached along a neighborhood of z_i and w_{i+1} . Letting $w = w_1$ and $z = z_\ell$, we get a two-pointed Heegaard diagram $(\Sigma', \alpha, \beta, w, z)$ for $\kappa(\vec{L}) \subset \#^{\ell-1}(S^2 \times S^1)$ (compare the proof of Theorem 4.5). Admissibility for $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ translates into admissibility for $(\Sigma', \alpha, \beta, w, z)$.

Proof of Theorem 1.1 There is a one-to-one correspondence between generators for $\widehat{CFL}(\vec{L})$ and $\widehat{CFK}(\kappa(\vec{L}))$.

Collapse the relative \mathbb{Z}^ℓ filtration on $\widehat{CFL}(S^3, \vec{L})$ to a relative \mathbb{Z} -filtration by

$$\mathcal{F}_o(\mathbf{x}) - \mathcal{F}_o(\mathbf{y}) = \sum_i (n_{z_i}(\phi) - n_{w_i}(\phi)),$$

where here $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is any Whitney disk in $\text{Sym}^{g+\ell-1}(\Sigma)$. Recall also that the \mathbb{Z} -filtration on $\widehat{CF}(\#^{\ell-1}(S^2 \times S^1), \kappa(\vec{L}))$ is given by

$$\mathcal{F}'(\mathbf{x}) - \mathcal{F}'(\mathbf{y}) = n_{z_\ell}(\phi') - n_{w_1}(\phi')$$

where here $\phi' \in \pi_2(\mathbf{x}, \mathbf{y})$ is any Whitney disk in $\text{Sym}^{g+\ell-1}(\Sigma')$.

Clearly, any Whitney disk $\phi' \in \pi_2(\mathbf{x}, \mathbf{y})$ in $\text{Sym}^{g+\ell-1}(\Sigma')$ gives rise naturally to a Whitney disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ in $\text{Sym}^{g+\ell-1}(\Sigma)$, having the property that $n_{z_i}(\phi) = n_{w_{i+1}}(\phi)$ for $i = 1, \dots, \ell - 1$. From these conditions, it is immediate to see that

$$\mathcal{F}_o(\mathbf{x}) = \mathcal{F}'(\mathbf{x}),$$

identifying one of the factors in the bigrading for both theories.

Moreover, it is also straightforward to see that $\mu(\phi) = \mu(\phi')$, and hence the relative homological gradings of both theories are identified. A shift in the absolute homological bigradings by $\frac{\ell-1}{2}$ can be identified in a model calculation, according to which the top-dimensional generator $\widehat{HF}(\#^{\ell-1}(S^2 \times S^1))$ generates $\widehat{HF}(S^3)$ in the corresponding multiply-pointed Heegaard diagram. We can also remove the indeterminacy in the relative \mathbb{Z} -filtration, using the fact that both theories are enjoy the symmetry property (Proposition 8.2).

Having established that both complexes have the same generators with the same bigradings, it remains to show that the differentials are identified; ie for suitable choices of complex structure on Σ' , the holomorphic disks correspond, as well.

By stabilizing and renumbering curves if necessary, we can arrange for the diagram $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ to have the following properties:

- For positive, odd i less than or equal to ℓ , w_i and z_i are connected by an arc δ_i which meets α_i in a single point but which is disjoint from all the other α_i and β_j .
- For positive, even i less than or equal to ℓ , w_i and z_i are connected by an arc δ_i which meets β_i in a single point but which is disjoint from all the other α_i and β_j .

Our aim is to exhibit a complex structure on Σ' with the property that for all $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and all $\phi' \in \pi_2(\mathbf{x}, \mathbf{y})$ with $n_{w_1}(\phi') = n_{z_\ell}(\phi') = 0$ and $\mu(\phi') = 1$ (ie these are the disks which contribute in the differential for $\widehat{\text{CFK}}(\kappa(\bar{L}))$), if $\mathcal{M}(\phi')$ is non-empty, then we must have $n_{z_i}(\phi') = n_{w_i}(\phi') = 0$ for all i , as well.

The complex structure is obtained by thinking of Σ' as obtained from Σ by adding one-handles with a long connected sum length attached connecting z_i to w_{i+1} (for $i = 1, \dots, \ell$). We choose the feet of these handles z_i and w_{i+1} in a family $\{z_i(t)\}_{t \in [0, \infty)}$ and $\{w_{i+1}(t)\}_{t \in [0, \infty)}$ starting out at $z_i(0)$ and $w_{i+1}(0)$ our original z_i and w_{i+1} respectively, $z_i(t), w_{i+1}(t) \in \Sigma - \alpha_1 - \dots - \alpha_{g+\ell} - \beta_1 - \dots - \beta_{g+\ell}$, for odd i less than ℓ ,

$$\lim_{t \rightarrow \infty} z_i(t) = z_i^\infty \in \alpha_i,$$

while for even i ,

$$\lim_{t \rightarrow \infty} w_i(t) = w_i^\infty \in \beta_i.$$

This can be achieved by constraining $w_i(t)$ and $z_i(t)$ to be subarcs of δ_i .

We have a two-parameter family of complex structures on Σ' , $\Sigma'(\sigma, \tau)$, where σ denotes the conformal structure on the annulus in the connected sum region, and τ parametrizes the placement of the feet of the one-handles $w_i(\tau)$ and $z_i(\tau)$.

Starting with our homotopy class $\phi' \in \pi_2(\mathbf{x}, \mathbf{y})$, thought of representing a cylindrical flow-line in Σ' , let $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ be the corresponding Whitney disk in $\text{Sym}^{g+\ell-1}(\Sigma)$. If for all τ , we can find sufficiently large σ for which $\mathcal{M}_{\sigma, \tau}(\phi')$ is non-empty, then we can take a Gromov limit (as $\sigma \mapsto \infty$) for any choice of τ . This gives a possibly broken pseudo-holomorphic cylindrical flow-line \bar{u}_τ representing ϕ .

Observe that since $n_{w_1}(\phi') = 0$, it is clear that \bar{u}_τ contains no components which contain all of Σ . Moreover, it can contain no boundary degenerations. Specifically, each boundary degeneration satisfies $n_{w_i}(\psi) = n_{z_i}(\psi)$ for all $i = 1, \dots, \ell$, while a boundary degeneration arising in this manner must also satisfy the additional conditions that $n_{w_1}(\psi) = n_{z_\ell}(\psi) = 0$ (as $n_{w_1}(\phi) = n_{z_\ell}(\phi) = 0$), forcing all $n_{w_i}(\psi) = n_{z_i}(\psi)$. This forces $\psi = 0$.

Thus, \bar{u}_τ represents a juxtaposition of cylindrical flow-lines in Σ . Moreover, each component u_τ in the Gromov compactification must satisfy the conditions that

$$\rho^{z_i(\tau)}(u_\tau) = \rho^{w_{i+1}(\tau)}(u_\tau)$$

for $i = 1, \dots, \ell - 1$, while $\rho^{w_1}(u_\tau) = \rho^{z_\ell}(u_\tau) = 0$. Indeed, the same remarks apply taking $\tau \rightarrow \infty$, as well: we obtain a possibly broken flowline \bar{u}_∞ , each of whose components consist of cylindrical flowlines satisfying

$$(28) \quad \rho^{z_i^\infty}(u_\infty) = \rho^{w_{i+1}^\infty}(u_\infty).$$

Choose i minimal with the property that $n_{z_i}(\phi') \neq 0$. If there is no such i , of course, then ϕ' represents a flowline with all $n_{z_i}(\phi') = n_{w_i}(\phi') = 0$ as desired. Otherwise, we assume that $n_{w_i}(\phi') = 0$ (returning to the other case later). Then, some component u_∞ of \bar{u}_∞ has $\rho^{z_i^\infty}(u_\infty) \neq \emptyset$. In fact, if i is odd, then $\rho^{z_i^\infty}(u_\infty)$ is contained in $\{1\} \times \mathbb{R}$. However, since w_{i+1}^∞ is not on one of the α -circles, $\rho^{w_{i+1}^\infty}(u_\infty)$ is disjoint from $\{1\} \times \mathbb{R}$, contradicting Equation (28). A symmetrical argument applies when i is even. In the other case, where $n_{w_i}(\phi') \neq 0$, we can instead find the minimal j for which $n_{w_j}(\phi') \neq 0$, and then apply similar reasoning to derive a contradiction.

10.1 Alternating links

A large class of calculations is provided by Theorem 1.3, which follows quickly from the material established thus far, combined with results from [18].

Proof of Theorem 1.3 In [18, Theorem 4.1], an analogue of Theorem 1.3 is established for the knot invariant of L and the (one-variable) Alexander-Conway polynomial; specifically, letting Δ_L be its Alexander-Conway polynomial, and writing

$$(T^{-1/2} - T^{1/2})^{n-1} \cdot \Delta_L(T) = a_0 + \sum_{s>0} a_s(T^s + T^{-s}),$$

we have that

$$\widehat{\text{HFK}}(S^3, L, s) \cong \mathbb{F}_{\binom{|a_s|}{s+\frac{n}{2}}}.$$

Combining this with Theorem 1.1, the result follows.

11 The Künneth formula for connected sums

In this section, we study the behaviour of the link invariant under connected sums, generalizing corresponding results for knot Floer homology [20; 24]. We give a variant for the filtration of $CF^-(S^3)$, and use it to conclude a corresponding result for $\widehat{\text{HFL}}$ (Theorem 1.4).

Let C be a \mathbb{Z}^m filtered chain complex over the ring $\mathbb{F}[U_1, \dots, U_m]$ and C' be a \mathbb{Z}^n -filtered chain complex over $\mathbb{F}[U'_1, \dots, U'_n]$, we can form their tensor product $C \otimes_{U_1 \sim U'_1} C'$, thought of as a chain complex over the polynomial ring

$$\mathbb{F}[U_1, \dots, U_m, U'_1, \dots, U'_n] / U_1 \sim U'_1.$$

This complex is naturally \mathbb{Z}^{m+n-1} -filtered, by the rule that the part of $C \otimes C'$ in filtration level $(a_1, a_2, \dots, a_m, b_2, \dots, b_n)$,

$$\mathcal{F}(C \otimes_{U_1 \sim U'_1} C', (a_1, a_2, \dots, a_m, b_2, \dots, b_n)),$$

is generated by the sum over all $a + b = a_1$ of the subcomplexes

$$\mathcal{F}(C, a, a_2, \dots, a_m) \otimes \mathcal{F}(C', (b, b_2, \dots, b_n)).$$

Consider two links L_1 and L_2 , and fix components $p \in L_1$, $p' \in L_2$. There is a natural map

$$\#: \underline{\text{Spin}}^c(L_1) \times \underline{\text{Spin}}^c(L_2) \longrightarrow \underline{\text{Spin}}^c(L_1 \#_{p \sim p'} L_2),$$

defined as follows. Represent relative Spin^c structures $\underline{\mathfrak{s}}_1$ and $\underline{\mathfrak{s}}_2$ as nowhere vanishing vector fields \vec{v}_1 and \vec{v}_2 which are tangent to L_1 and L_2 respectively. Deleting a sufficiently small ball B_1 around $p \in L_1 \subset S^3$, the vector field \vec{v}_1 points normal to the boundary $S^3 - B_1$ at exactly two points, where the corresponding component of L_1 meets ∂B_1 . Find a corresponding ball B_2 around $p' \in L_2$. We can then identify $v_1|_{\partial B_1}$ with $v_2|_{\partial B_2}$, and hence extend the vector field over $S^3 \# S^3$, to obtain a nowhere vanishing vector field containing $L_1 \#_{p \sim p'} L_2$ as a closed orbit. This determines the stated map. It is easy to see that

$$(\underline{\mathfrak{s}}_1 + \text{PD}[h_1]) \# (\underline{\mathfrak{s}}_2 + \text{PD}[h_2]) = \underline{\mathfrak{s}}_1 \# \underline{\mathfrak{s}}_2 + \iota_1(\text{PD}[h_1]) + \iota_2(\text{PD}[h_2]),$$

where here

$$\iota_i: H_1(S^3 - L_i) \longrightarrow H_1(S^3 - L_1 \# L_2)$$

are the natural maps.

We can refine the above tensor product as follows. If C is a $\underline{\text{Spin}}^c(L_1)$ -filtered complex over $\mathbb{F}[U_1, \dots, U_m]$ and C' is a $\underline{\text{Spin}}^c(L_2)$ -filtered complex over $\mathbb{F}[U'_1, \dots, U'_n]$, we can form the chain complex

$$C \otimes_{U_1 \sim U'_1} C',$$

equipped with the $\underline{\text{Spin}}^c(L_1 \# L_2)$ -filtration, where $\mathcal{F}(C \otimes_{U_1 \sim U'_1} C', \underline{\mathfrak{s}})$ is generated by the sum of $\mathcal{F}(C, \underline{\mathfrak{s}}_1) \otimes \mathcal{F}(C', \underline{\mathfrak{s}}_2)$ over all pairs $\underline{\mathfrak{s}}_1 \in \underline{\text{Spin}}^c(Y, L_1)$ and $\underline{\mathfrak{s}}_2 \in \underline{\text{Spin}}^c(L_2)$ with $\underline{\mathfrak{s}} = \underline{\mathfrak{s}}_1 \# \underline{\mathfrak{s}}_2$.

Theorem 11.1 Let L_1 and L_2 be oriented links with n_1 and n_2 components respectively. Fix also a point $p \in L_1$ and $p' \in L_2$, and form the connected sum $L_1 \# L_2$ (near the two distinguished points). Then, the chain complex $\text{CFL}^-(L_1 \#_p \sim_{p'} L_2)$ is isomorphic as a $\text{Spin}^c(L_1 \# L_2)$ -filtered chain complex over

$$\mathbb{F}[U_1, \dots, U_n, U'_1, \dots, U'_m] / U_1 \sim U'_1.$$

to the tensor product $\text{CFL}^-(L_1) \otimes_{U_1 \sim U'_1} \text{CFL}^-(L_2)$, where here U_1 and U'_1 are variables corresponding to the components of L_1 and L_2 distinguished by the points p and p' respectively.

Proposition 11.2 Let (Y_1, L_1) and (Y_2, L_2) are links with multi-pointed Heegaard diagrams

$$(\Sigma_1, \alpha_1, \beta_1, \{w_i^1\}_{i=1}^{n_1}, \{z_i^1\}_{i=1}^{n_1}) \quad (\Sigma_2, \alpha_2, \beta_2, \{w_i^2, z_i^2\}_{i=1}^{n_2}).$$

The connected sum $(Y_1 \# Y_2, L_1 \# L_2)$, where we connect the first component of L_1 with the first component of L_2 , can be given a Heegaard diagram

$$(\Sigma_1 \# \Sigma_2, \alpha_1 \cup \alpha_2, \beta_1 \cup \beta_2, \{w_i^1\}_{i=1}^{n_1} \cup \{w_i^2\}_{i=2}^{n_2}, \{z_i^1\}_{i=2}^{n_1-1} \cup \{z_i^2\}_{i=1}^{n_2}),$$

whose underlying surface $\Sigma_1 \# \Sigma_2$ is formed by the connected sum of Σ_1 and Σ_2 along the point $z_1^1 \in \Sigma_1$ and $w_1^2 \in \Sigma_2$. In this new diagram, the pair of reference points w_1^1 and z_1^2 represent the connected sum of the first component of L_1 with the first component of L_2 .

Proof This is straightforward. □

Proof of Theorem 11.1 Let α'_1, β'_1 , and α'_2 be small perturbations of the sets of curves α_1, β_1 , and α_2 respectively. Correspondingly, there are canonical top-dimensional intersection points $\Theta_1 \in \mathbb{T}_{\beta_1} \cap \mathbb{T}_{\beta'_1}$ and $\Theta_2 \in \mathbb{T}_{\alpha_2} \cap \mathbb{T}_{\alpha'_2}$.

We have a map

$$\begin{aligned} CF^-(\alpha_1, \beta_1) &\longrightarrow CF^-(\alpha_1 \cup \alpha_2, \beta_1 \cup \alpha'_2) \\ \mathbf{x}_1 &\mapsto \mathbf{x}_1 \times \Theta_2. \end{aligned}$$

Combining Propositions 6.5 and 6.7, we see that this is a chain map. There is a similar map

$$CF^-(\alpha'_2, \beta_2) \longrightarrow CF^-(\beta_1 \cup \alpha'_2, \beta'_1 \cup \beta_2)$$

defined by

$$\mathbf{x}_2 \mapsto \Theta_1 \times \mathbf{x}_2.$$

Finally, we have a chain map

$$CF^-(\alpha_1 \cup \alpha_2, \beta_1 \cup \alpha'_2) \otimes CF^-(\beta_1 \cup \alpha'_2, \beta'_1 \cup \beta_2) \longrightarrow CF^-(\alpha_1 \cup \alpha_2, \beta'_1 \cup \beta_2)$$

defined by counting pseudo-holomorphic triangles.

Putting these together, we obtain a composite chain map

$$\Phi: CF^-(\alpha_1, \beta_1) \otimes CF^-(\alpha'_2, \beta_2) \longrightarrow CF^-(\alpha_1 \cup \alpha'_2, \beta'_1 \cup \beta_2).$$

There is also a “closest point map” inducing an isomorphism of modules over the polynomial algebra $(\mathbb{F}[U_1, \dots, U_m, U'_1, \dots, U'_{n'}]/U_1 \sim U'_1)$

$$\iota: CF^-(\alpha_1, \beta_1) \otimes CF^-(\alpha'_2, \beta_2) \longrightarrow CF^-(\alpha_1 \cup \alpha'_2, \beta'_1 \cup \beta_2)$$

with $\iota(\mathbf{x}_1 \times \mathbf{x}_2) = \mathbf{x}'_1 \times \mathbf{x}'_2$. Here $\mathbf{x}'_1 \in \mathbb{T}_{\alpha_1} \cap \mathbb{T}_{\beta'_1}$ is the point nearest to $\mathbf{x}_1 \in \mathbb{T}_{\alpha_1} \cap \mathbb{T}_{\beta_1}$, while $\mathbf{x}'_2 \in \mathbb{T}_{\alpha'_2} \cap \mathbb{T}_{\beta_2}$ is the point closest to $\mathbf{x}_2 \in \mathbb{T}_{\alpha'_2} \cap \mathbb{T}_{\beta_2}$. The closest point map can alternatively be thought of as counting holomorphic triangles with minimal area. Thus, Φ is a map of the form ι plus lower order terms (provided that the curves α' resp. β' are sufficiently to α resp. β'). In particular, Θ induces an isomorphism of chain complexes over $\mathbb{F}[U_1, \dots, U_m, U'_1, \dots, U'_{n'}]/U_1 \sim U'_1$.

It is easy to see that

$$\underline{\mathfrak{s}}_{w_1, z_1}(\mathbf{x}_1) \# \underline{\mathfrak{s}}_{w_2, z_2}(\mathbf{x}_2) = \underline{\mathfrak{s}}_{w_1 \cup w_2, z_1 \cup z_2}(\mathbf{x}'_1 \times \mathbf{x}'_2).$$

Thus the map ι induces an isomorphism on the associated $\text{Spin}^c(L_1 \# L_2)$ -graded complex. The other terms in Φ have lower order than ι , and hence Φ induces a $\text{Spin}^c(L_1 \# L_2)$ -filtered isomorphism, as claimed.

Proof of Theorem 1.4 This theorem follows readily from Theorem 11.1, and the principle that an isomorphism of \mathbb{Z}^ℓ -filtered complexes induces an isomorphism on the homology of its associated graded object.

12 Examples

We turn to some sample calculations. Specifically, we will be concerned here with calculation the filtered chain homotopy type of $\overline{CFL}(\vec{L})$ for various oriented links. In Section 12.1, we give a theorem which can be used to compute this data explicitly for a two-bridge link. In Section 12.2, we give some calculations of this data for the first two non-alternating links.

12.1 Two-component alternating links

Although Theorem 1.3 concerns only the homology of the associated graded object $\widehat{\text{HFL}}(L)$, in many cases, one can also use it to deduce information about the filtered homotopy type of the full chain complex.

We study here the filtered chain homotopy type $\widehat{\text{CFL}}$ of alternating links. For simplicity, we consider here the case of two-component links.

Theorem 12.1 *Let \vec{L} be an oriented, two-component link which admits a connected, alternating projection. The filtered chain homotopy type of $\widehat{\text{CFL}}(S^3, \vec{L})$ is determined by the following data:*

- The multi-variable Alexander polynomial of \vec{L} ;
- The signature of \vec{L} ;
- The linking number of K_1 and K_2 ;
- The filtered chain homotopy type $\widehat{\text{CF}}(S^3, K_1)$ and $\widehat{\text{CF}}(S^3, K_2)$ of the two components K_1 and K_2 of L .

In the course of the proof, we show explicitly how $\widehat{\text{CFL}}(S^3, \vec{L})$ is determined by the data. First, we must set up some notions.

We say that a $\mathbb{Z} \oplus \mathbb{Z}$ -filtered chain complex is E_2 -collapsed if it has a splitting $C = \{C_{i,j}\}_{i,j \in \mathbb{Z}}$ so that its differential ∂ is written has the form $\partial = D^1 + D^2$, where here

$$D^1|_{C_{i,j}} = D^1_{i,j}: C_{i,j} \longrightarrow C_{i-1,j} \quad \text{and} \quad D^2|_{C_{i,j}} = D^2_{i,j}: C_{i,j} \longrightarrow C_{i,j-1}.$$

For example, suppose that C is a $\mathbb{Z} \oplus \mathbb{Z}$ filtered chain complex which also has an internal grading g (which the differential drops by one). Suppose moreover that the filtration (i, j) and the grading g are related by $i + j - g = c$ for some constant c . Then, C is E_2 collapsed. Theorem 1.3 ensures that the link Floer homology of an alternating link is E_2 -collapsed.

We give some examples of E_2 -collapsed chain complexes. Let $B_{(d)}$ denote the chain complex with

$$(B_{(d)})_{i,j} = \begin{cases} \mathbb{F}_{(d)} & \text{if } (i, j) = (0, 0) \\ \mathbb{F}_{(d+1)} & \text{if } (i, j) \in \{(0, 1), (1, 0)\} \\ \mathbb{F}_{(d+2)} & \text{if } (i, j) = (1, 1) \\ 0 & \text{otherwise} \end{cases}$$

and the property that $D_{1,1}^1$ and $D_{1,0}^1$ and $D_{1,1}^2$ and $D_{0,1}^2$ are field isomorphisms.

Fix an integer ℓ , and let $V_{(d)}^\ell$ be the chain complex with

$$\left(V_{(d)}^\ell\right)_{i,j} = \begin{cases} \mathbb{F}_{(d)} & \text{if } (i, j) = (-j, j) \text{ with } j = 0, \dots, \ell - 1 \\ \mathbb{F}_{(d-1)} & \text{if } (i, j) = (-j - 1, j) \text{ with } j = 0, \dots, \ell - 1 \\ 0 & \text{otherwise} \end{cases}$$

where $D = D^1 + D^2$ is a sum of maps, where $D_{-j,j}^1$ are field isomorphisms for all i , $D_{-j,j}^2$ is an isomorphism for all $1 \leq j \leq \ell - 1$.

Similarly, let $H_{(d)}^\ell$ be the chain complex with

$$\left(H_{(d)}^\ell\right)_{i,j} = \begin{cases} \mathbb{F}_{(d)} & \text{if } (i, j) = (i, -i) \text{ with } i = 0, \dots, \ell - 1 \\ \mathbb{F}_{(d-1)} & \text{if } (i, j) = (i, -i - 1) \text{ with } i = 0, \dots, \ell - 1 \\ 0 & \text{otherwise} \end{cases}$$

where $D = D^1 + D^2$ is a sum of maps, where $D_{i,-i}^1$ is an isomorphism for all $i = 1, \dots, \ell - 1$, while $D_{i,-i}^2$ is an isomorphism for all i .

Note that $H_*(B) \cong H_*(V_{(d)}^\ell) \cong H^*(H_{(d)}^\ell) = 0$.

There are two basic types of E_2 -collapsed chain complexes with non-trivial homology. Let $X_{(d)}^\ell$ be the complex given by

$$\left(X_{(d)}^\ell\right)_{i,j} = \begin{cases} \mathbb{F}_{(d)} & \text{if } i + j = \ell \text{ and } i, j \geq 0 \\ \mathbb{F}_{(d+1)} & \text{if } i + j = \ell + 1 \text{ if } i, j > 0 \\ 0 & \text{otherwise} \end{cases}$$

where $D = D^1 + D^2$ is a sum of maps, where $D_{i,j}^1$ and $D_{i,j}^2$ are isomorphisms when $i + j = \ell + 1$ and $i, j > 0$, and zero otherwise.

Also, let $Y_{(d)}^\ell$ be the complex determined by

$$\left(Y_{(d)}^\ell\right)_{i,j} = \begin{cases} \mathbb{F}_{(d)} & \text{if } i + j = \ell \text{ and } i, j \geq 0 \\ \mathbb{F}_{(d-1)} & \text{if } i + j = \ell - 1 \text{ and } i, j \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

where $D = D^1 + D^2$ is a sum of maps, where $D_{i,j}^1$ is an isomorphism when $i > 0$, $j \geq 0$, and $i + j = \ell$, while $D_{i,-i}^2$ is an isomorphism when $i \geq 0$, $j > 0$, and $i + j = \ell$, and all other maps are zero.

For these latter two complexes, we see that $H_*(X_{(d)}^\ell) \cong H_*(Y_{(d)}^\ell) \cong \mathbb{F}_{(d)}$.

It is straightforward to see that for each E_2 -collapsed chain complex C is filtered chain homotopy equivalent to a filtered chain complex which splits as a direct sum of copies of chain complexes of the type

$$B_{(d)}[i, j], \quad V_{(d)}^\ell[i, j], \quad H_{(d)}^\ell[i, j], \quad X_{(d)}^\ell[i, j], \quad \text{and} \quad Y_{(d)}^\ell[i, j],$$

allowing ℓ , i , and j to vary. (Here, as in Section 2, given A a $\mathbb{Z} \oplus \mathbb{Z}$ -filtered chain complex, we let $A[i, j]$ denote the $\mathbb{Z} \oplus \mathbb{Z}$ filtered chain complex obtained from A by shifting the filtration, so that $(A[i, j])_{x,y} = A_{x+i,y+j}$).

Proof of Theorem 12.1 In view of Theorem 1.3, $\widehat{CF}(S^3, \vec{L})$ is an E_2 -collapsed chain complex, and as such decomposes into summands of the above five types. Thus, the filtered chain homotopy type is determined by the number of summands of each type, and their various parameters (i , j , d , and in four cases, ℓ). Our goal is to show how the data assumed in the statement of Theorem 12.1 can be used to extract all of the needed information.

Let $E_{(d)}^\ell$ denote the \mathbb{Z} -filtered chain complex

$$\left(E_{(d)}^\ell\right)_i = \begin{cases} \mathbb{F}_{(d)} & \text{if } i = 0 \\ \mathbb{F}_{(d-1)} & \text{if } i = -\ell \\ 0 & \text{otherwise} \end{cases}$$

endowed with a differential which is an isomorphism from $\left(E_{(d)}^\ell\right)_0$ to $\left(E_{(d)}^\ell\right)_{-\ell}$. It is easy to see that each \mathbb{Z} -filtered chain complex splits as a sum of complexes of the form $E_{(d)}^\ell[i]$, and also homologically non-trivial complexes of the form \mathbb{F} supported in some fixed degree and filtration level.

Observe that taking the homology in the vertical direction (ie taking the homology with respect to D^2 to obtain a \mathbb{Z} -filtered chain complex, endowed with the differential induced from D^1) has the property that it annihilates $B[i, j]$, $H^\ell[i, j]$. Moreover, taking the vertical homology of $V_{(d)}^\ell[i, j]$, we obtain $E_{(d)}^\ell[i]$. The vertical homology of $X^\ell[i, j]$ is a copy of \mathbb{F} in filtration level i , while the horizontal homology gives \mathbb{F} in filtration level $j - \ell$; the horizontal homology of $Y^\ell[i, j]$ is a copy of \mathbb{F} in filtration level j while its vertical homology gives \mathbb{F} in filtration level $i + \ell$.

On the other hand, according to Proposition 7.1, if we take the vertical homology of $\widehat{CFL}(\vec{L})$, we obtain the knot filtration on $\widehat{CF}(S^3)$ coming from K_1 , tensored with $\mathbb{F} \oplus \mathbb{F}$ (supported in two consecutive dimensions) and shifted over (in its filtration) by n , and taking the horizontal homology, we obtain the knot filtration of K_2 again shifted by $\frac{n}{2}$ and tensored with $\mathbb{F} \oplus \mathbb{F}$. Specifically, these two projections are \mathbb{Z} -filtered chain complexes, and as such can be decomposed into one-step complexes \mathbb{F} (supported in

some filtration level), and two-step complexes $E_{(d)}^\ell[i]$, consisting of \mathbb{F} in filtration level i and $i - \ell$ (trivial otherwise), endowed with a differential which induces an isomorphism from the part in filtration level i to the one in filtration level $i - \ell$.

It follows from the above remarks that the summands in $\widehat{CFL}(\vec{L})$ of the form V are in two-to-one correspondence with the summands of type E in $\widehat{CF}(K_1)$: More precisely, if $E_{(d)}^\ell[i]$ appears in $\widehat{CF}(K_1)$, there are two summands,

$$V_{(d)}^\ell[i - \frac{n}{2}, d - i - \frac{n}{2}], \quad \text{and} \quad V_{(d-1)}^\ell[i - \frac{n}{2}, d - 1 - i - \frac{n}{2}].$$

Similarly, the summands in $\widehat{CFL}(\vec{L})$ of the form H are in two-to-one correspondence with the summands of type E in $\widehat{CF}(K_2)$: More precisely, if $E_{(d)}^\ell[i]$ appears in $\widehat{CF}(K_2)$, there are two summands,

$$H_{(d)}^\ell[i - \frac{n}{2}, d - i - \frac{n}{2}] \quad \text{and} \quad H_{(d-1)}^\ell[i - \frac{n}{2}, d - 1 - i - \frac{n}{2}].$$

Moreover, since $H_*(C) \cong \mathbb{F}_{(-1)} \oplus \mathbb{F}_{(0)}$, there can be at most two summands of type $X^\ell[i, j]$ or $Y^\ell[i, j]$. We claim that there are in fact only two possible cases: Either the two summands are $X_0^\ell[i, j]$ and $X_{-1}^{\ell-1}[i, j]$ or $Y_{-1}^\ell[i, j]$ and $Y_0^{\ell-1}[i - 1, j - 1]$. This follows from the constraints that $H_*(C) \cong \mathbb{F} \oplus \mathbb{F}$ are supported in two consecutive degrees (0 and -1), together with the constraint that if we take the horizontal resp. vertical homologies, we are to get a filtered chain complex whose homology $\mathbb{F} \oplus \mathbb{F}$ is supported in two consecutive gradings and the same filtration level. In fact, taking the vertical resp. horizontal homology gives the knot Floer homology of K_1 resp. K_2 , tensored with $\mathbb{F} \oplus \mathbb{F}$, supported in two consecutive dimensions and filtration level $\tau(K_1)$ resp. $\tau(K_2)$.

More specifically, writing $\tau_1 = \tau(K_1)$, $\tau_2 = \tau(K_2)$, and n for the linking number of K_1 with K_2 , we have two cases, according to the sign of

$$\ell = \tau_1 + \tau_2 + n + \frac{\sigma - 1}{2}.$$

If $\ell \geq 0$, $\widehat{CFL}(\vec{L})$ has summands

$$Y_{(0)}^\ell[\tau_2 + \frac{1 - \sigma - n}{2}, \tau_1 + \frac{1 - \sigma - n}{2}] \oplus Y_{(-1)}^{\ell-1}[\tau_2 + \frac{3 - \sigma - n}{2}, \tau_1 + \frac{3 - \sigma - n}{2}].$$

If $\ell \leq 0$, then $\widehat{CFL}(\vec{L})$ has summands

$$X_{(0)}^{|\ell|}[\tau_1 + \frac{n}{2}, \tau_2 + \frac{n}{2}] \oplus X_{(-1)}^{|\ell|-1}[\tau_1 + \frac{n}{2}, \tau_2 + \frac{n}{2}].$$

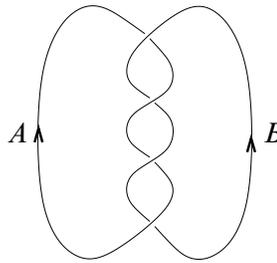


Figure 4: **The oriented link \vec{H}_n with $n = 2$**

The remaining summands of type B (and their precise placement) are determined by the Alexander polynomial of L .

For example, if \vec{L} is a two-component, two-bridge link, Theorem 12.1 applies immediately to express the knot filtration in terms of the Alexander polynomial and signatures of \vec{L} . Note in particular, that in this case, the two components K_1 and K_2 are individually unknotted; it follows in particular that there are no summands of type V or H . (Any two-bridge link can be represented by a four-pointed Heegaard diagram with genus zero, and hence the holomorphic curve counts defining the differential take place in the Riemann sphere. Thus, these counts are purely combinatorial, compare [23] and also [20, Section 6].)

Consider the Hopf link \vec{H} . There are two orientations for \vec{H} , distinguished by the signature. We denote the two cases by \vec{H}^\pm , with the convention that $\sigma(\vec{H}^\pm) = \pm(-1)$. Note also that $\Delta_{\vec{H}}(X, Y) = 1$, and hence $\widehat{\text{HFL}}(\vec{H})$ consists of four generators. It follows at once from the above considerations that

$$\widehat{\text{CFL}}(\vec{H}^+) \simeq Y_{(0)}^0[\frac{1}{2}, \frac{1}{2}] \oplus Y_{(-1)}^1[-\frac{1}{2}, -\frac{1}{2}]$$

while

$$\widehat{\text{CFL}}(\vec{H}^-) \simeq X_{(0)}^1[-\frac{1}{2}, -\frac{1}{2}] \oplus X_{(-1)}^0[-\frac{1}{2}, -\frac{1}{2}].$$

More generally, consider the link \vec{H}_n consisting of two unknotted circles which link each other both algebraically $n > 0$ and geometrically n times – ie this is the $(2, 2n)$ torus link with the specified orientation (see Figure 4). It is easy to see that

$$\Delta_{\vec{H}_n}(S, T) = S^{\frac{n-1}{2}} T^{\frac{1-n}{2}} \sum_{i=0}^{n-1} (S^{-1}T)^i$$

and $\sigma(\vec{H}_n) = -n$.

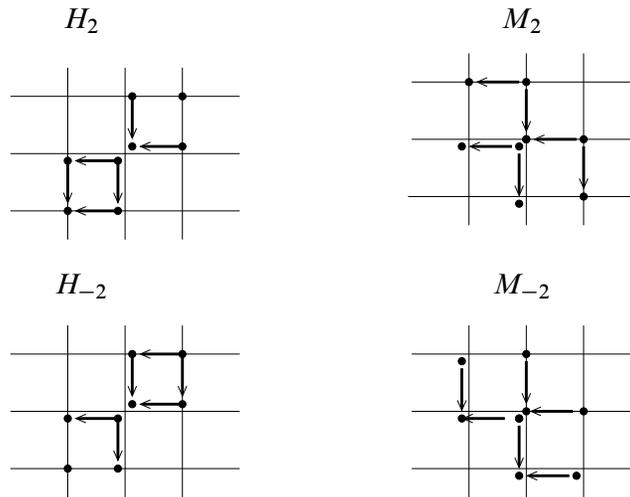


Figure 5: **Link Floer homologies of H_n with $n = 2$** We have illustrated the link Floer homologies of H_2 for the four different possible orientations. The upper left is \vec{H}_2 – the right-handed, positive version, as illustrated in Figure 4. The upper right-hand illustrates the link obtained by reversing the orientation of B . The second row illustrates the link Floer homology groups of the mirrors of the links in the first row.

It follows at once that

$$\widehat{CF}(\vec{H}_n) \simeq Y_{(0)}^0 \left[\frac{n}{2}, \frac{n}{2} \right] \oplus Y_{(-1)}^1 \left[\frac{n}{2} - 1, \frac{n}{2} - 1 \right] \bigoplus_{i=1}^{n-1} B_{(-2i)} \left[-i + \frac{n}{2}, -i + \frac{n}{2} \right].$$

Similarly, if we reverse the orientation of one of the components, we get a link denoted \vec{H}_{-n} , and we see that

$$\widehat{CF}(\vec{H}_{-n}) \simeq X_{(0)}^n \left[-\frac{n}{2}, -\frac{n}{2} \right] \oplus X_{(-1)}^{n-1} \left[-\frac{n}{2}, -\frac{n}{2} \right].$$

The complexes for \vec{H}_n and \vec{H}_{-n} , along with their mirrors, are illustrated in the case where $n = 2$ in Figure 5.

One can alternatively calculate this directly by looking at a genus zero Heegaard diagram, as illustrated in Figure 6.

12.2 Two non-alternating examples

We calculate $\widehat{CFL}(L)$ for the first two two-component, non-alternating links L_1 and L_2 . These are both seven-crossing links, both obtained as a union of a trefoil and

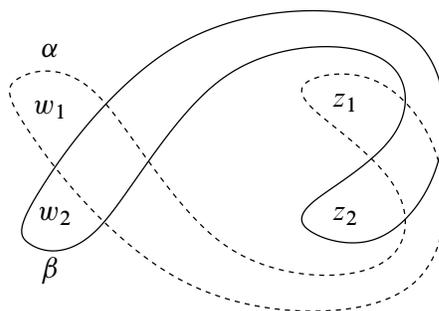


Figure 6: **Heegaard diagram for \vec{H}_n with $n = 2$** We have illustrated the four basepoints and the two curves (α and β). This picture takes place on the Riemann sphere. The eight intersection points are visible; and indeed each gives rise to a generator in homology, as illustrated in Figure 5. The various orientations are, of course, obtained by permuting the roles of w_1 with z_1 , while the mirrors can be found by reversing the roles of α and β .

an unknot, denoted 7^2_8 and 7^2_7 respectively in Rolfsen’s table [26], and $7n_2$ and $7n_1$ respectively in Thistlethwaite’s link table [27]. They are distinguishable immediately by the linking number of the two components. For 7^2_8 , the two components have linking number zero, while for 7^2_7 they have linking number ± 2 . Both links are illustrated in Figures 7 and 8 respectively.

12.2.1 The link L_1 (a.k.a. $7^2_8, 7n_2$) We claim that the link Floer homology groups of the link L_1 illustrated in Figure 7) have the form

$$(29) \quad \widehat{\text{HFL}}(7^2_8, i, j) = \begin{cases} \mathbb{F}^4_{(0)} & \text{if } (i, j) = (0, 0) \\ \mathbb{F}^2_{(i+j)} & \text{if } i, j \in \mathbb{Z}, |i| + |j| = 1 \\ \mathbb{F}_{(i+j)} & \text{if } i, j \in \{\pm 1\} \\ 0 & \text{otherwise.} \end{cases}$$

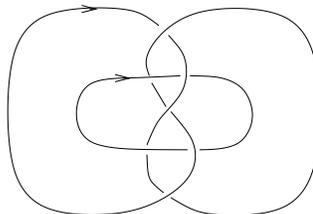


Figure 7: **A seven-crossing non-alternating link** This is the link we denote L_1 ; it is denoted 7^2_8 in Rolfsen’s notation; $7n_2$ in Thistlethwaite’s.

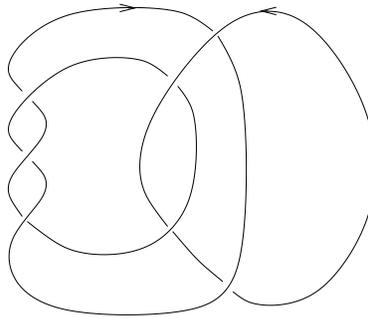


Figure 8: **The other seven-crossing non-alternating link** This is the link we denote L_2 ; it is denoted 7_7^2 in Rolfsen's notation; $7n_1$ in Thistlethwaite's.

In particular, this link, too, is E_2 -collapsed. Note that is not necessary to separate two possible orientations: Rotation through a vertical axis gives an identification between two possible orientations for L_1 .

Equation (29) can be seen by considering the (admissible) Heegaard diagram pictured in Figure 9, which takes place in a genus one surface, with attaching circles $\{\alpha_1, \alpha_2\}$ and $\{\beta_1, \beta_2\}$.

We label

$$\begin{aligned} \alpha_1 \cap \beta_1 &= \{x_1, x_2, x_3\} & \alpha_1 \cap \beta_2 &= \{b_1, b_2, b_3, b_4\} \\ \alpha_2 \cap \beta_1 &= \{y_1, y_2, y_3\} & \alpha_2 \cap \beta_2 &= \{a_1, a_2, a_3, a_4\} \end{aligned}$$

Any two intersection points between α_i and β_j can be connected by a sequence of consecutive embedded Whitney disks. Using these disks (each of which has $\mu(\phi) = \pm 1$, depending on its orientation), it is straightforward to calculate the relative gradings and filtrations of any two intersection points of the form $\{a_i, x_j\}$ (or any two intersection points of the form $\{b_i, y_j\}$). Finally, there is a square at the center of the diagram which represents a Whitney disk connecting $a_4 \times x_1$ to $b_4 \times y_3$, which allows one to complete the calculations all relative gradings and filtration levels of generators. Indeed, one readily sees that in all but four filtration levels (which in the normalization of Equation (29), are $(i, j) \in \{-1, -2\} \times \{0, -1\}$) the generators have the same relative gradings, and hence the ranks and degrees of the homology (up to an overall translation) are as given in Equation (29).

We consider the four remaining filtration levels $(-2, -1)$, $(-2, 0)$, $(-1, -1)$, $(-1, 0)$, represented by intersection points

$$\begin{aligned} \{a_4 \times x_1, b_4 \times y_3\}, & \quad \{a_2 \times x_1, b_2 \times y_3\}, \\ \{a_3 \times x_1, b_3 \times y_3, b_4 \times y_2\}, & \quad \{a_1 \times x_1, a_4 \times x_2, b_1 \times y_3, b_2 \times y_2\} \end{aligned}$$

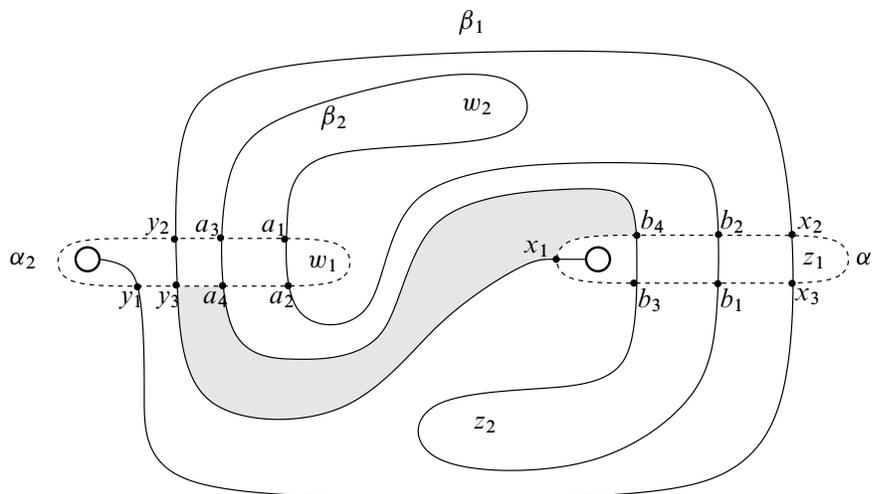


Figure 9: **Heegaard diagram for 7_8^2** Attach a one-handle to the plane at the two circles. This gives a Heegaard diagram for 7_8^2 , with the orientation as given in Figure 7. The shaded rectangle represents a flow-line from $a_4 \times x_1$ to $b_4 \times y_3$.

respectively. We need to show that the homology groups have ranks 0, 0, 1 and 2 respectively.

To handle the filtration level $(-2, -1)$, one inspects the Heegaard diagram in Figure 9, to find a rectangle connecting $a_4 \times x_1$ to $b_4 \times y_3$. It is easy to see that there are no other non-negative domains connecting these two disks, hence the differential annihilates this pair of generators. (For this, it is useful to observe that all other possible homology classes of disks which could contribute to the differential are obtained from this given

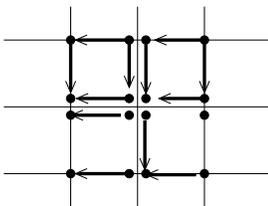


Figure 10: **Link Floer homology for L_1** We have illustrated the link Floer homology for the link L_1 as pictured in Figure 7. The first coordinate denotes the filtration induced by the trefoil component, while the second denotes the filtration by the unknot component.

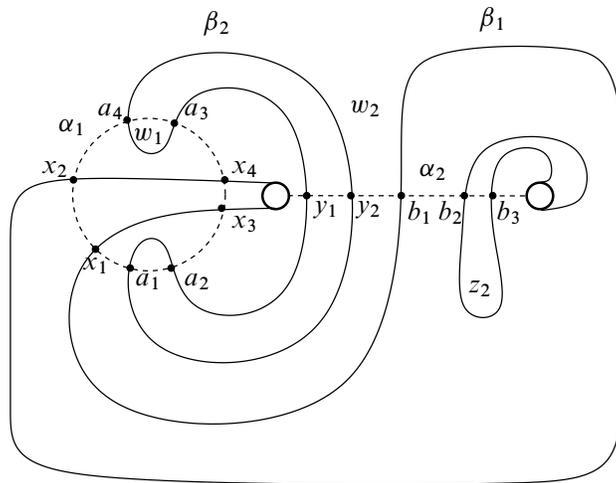


Figure 11: **Heegaard diagram for 7_7^2** Attach a one-handle to the plane at the two circles. This gives a Heegaard diagram for 7_7^2 , with the orientation as given in Figure 8.

square by the addition of a periodic domain. But any non-trivial periodic domain has both positive and negative local multiplicities in one of the regions adjoining b_1 , which is disjoint from our given square. Thus, there are no other non-negative homotopy classes.)

In the same manner, one can find a rectangle to show that the homology in the filtration level $(-2, 0)$ is trivial. Indeed, one can find also rectangles connecting generators $a_3 \times x_1$ to $b_4 \times y_2$ and $a_1 \times x_1$ to $b_2 \times y_2$, showing that the differentials in filtration levels $(-1, -1)$ and $(-1, 0)$ are non-trivial. It is straightforward then to conclude that the groups with their relative gradings are as given in Equation (29).

To verify the absolute gradings, recall that we need to orient the knot (and an orientation is implicit in the Heegaard diagram, via the choices of w_i and z_j). If we allow isotopies to cross z_2 and z_1 , it is easy to shrink β_2 in Figure 9 (canceling out intersection points $b_1, b_3, b_2, b_4, a_4, a_2$) and then perform a finger move on β_1 to cancel point x_2 and x_3 , to obtain an admissible doubly-pointed Heegaard diagram for S^3 with exactly two generators, $a_1 \times x_1$ and $a_3 \times x_1$. It is easy to see that $a_1 \times x_1$ represents a generator for $\widehat{HF}(S^3)$, and hence it is supported in degree zero. This completes the verification of Equation (29).

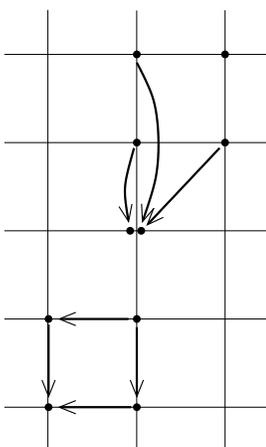


Figure 12: **Heegaard diagram for 7_7^2** Link Floer homology for L_2 with the orientation from Figure 8. The first coordinate represents the unknot factor, while the second represents the trefoil.

As in the alternating case, the higher differentials on $\widehat{\text{HFL}}(S^3, \vec{L})$ can be determined by the knot Floer homologies of the two components (the trefoil and the unknot). We illustrate the resulting complex in Figure 10.

12.2.2 The link L_2 (aka $7_7^2, 7n1$) The Floer homology of the link \vec{L}_2 is not E_2 -collapsed. In fact, we have:

$$(30) \quad \widehat{\text{HFL}}(\vec{L}_2) = \begin{cases} \mathbb{F}_{(i+j-3)} & \text{if } (i, j) \in \{0, 1\} \times \{1, 2\} \cup \{0, -1\} \times \{-1, -2\} \\ \mathbb{F}_{(-2)} \oplus \mathbb{F}_{(-3)} & \text{if } (i, j) = (0, 0) \\ 0 & \text{otherwise.} \end{cases}$$

To perform the calculation, we draw the Heegaard diagram in Figure 11.

This once again is a genus one Heegaard diagram with two pairs of attaching circles $\{\alpha_1, \alpha_2\}$ and $\{\beta_1, \beta_2\}$. Now we label

$$\begin{aligned} \alpha_1 \cap \beta_1 &= \{x_1, x_2, x_3, x_4\}, & \alpha_1 \cap \beta_2 &= \{a_1, a_2, a_3, a_4\}, \\ \alpha_2 \cap \beta_1 &= \{b_1, b_2, b_3\}, & \alpha_2 \cap \beta_2 &= \{y_1, y_2\}. \end{aligned}$$

Proceeding as before, it is easy to calculate filtration levels and gradings (up to an overall shift) of generators. There are three filtration levels, $(-1, 1)$, $(-1, 0)$, and $(1, 0)$, where there are exactly two generators, one, $(0, 1)$, where there are three, and finally one, $(0, 0)$ where there are four. In all the rest, there are one or zero generators, so in these other filtration levels, Equation (30) is immediately verified.

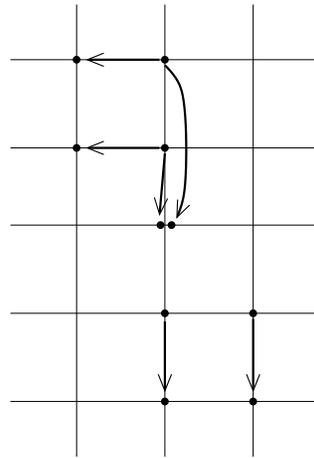


Figure 13: **Heegaard diagram for 7_7^2** Link Floer homology for 7_7^2 , reversing the orientation of with the orientation of the unknot from that depicted in Figure 8.

For $(i, j) = (-1, 1)$, there are two generators $\{a_3 \times b_3, x_4 \times y_1\}$, and a rectangle can be immediately found to show that the homology in this filtration level is trivial. The same remarks apply for $(i, j) = (-1, 0)$ and $(1, 0)$. Also, rectangles can be found connecting $a_4 \times b_3$ to $x_2 \times y_1$ (and also $a_4 \times b_3$ to $x_4 \times y_2$), forcing the differential to be non-trivial and hence the homology to be one-dimensional.

The remaining case of $(i, j) = (0, 0)$ can be either analyzed carefully in this way, or alternatively, one can argue that taking the horizontal homologies should give the knot Floer homology of the trefoil (as in Proposition 7.1) tensored with a two-dimensional vector space. In particular, the horizontal homology through the $j = 0$ line must be two-dimensional. But since the homology is trivial for all $i \neq 0$ and $j = 0$, it follows that this horizontal homology is identified with the homology in the filtration level $(i, j) = (0, 0)$.

Absolute degrees are now computed by considering isotopies which cross z_1 and z_2 . It is easy to find an isotopy of β_1 crossing z_2 which cancels $b_1, b_2, x_1, x_2, x_3, x_4$, and then an isotopy of β_2 which crosses z_1 cancelling a_1 and a_2 . This leaves a diagram with exactly two generators $a_3 \times b_1$ and $a_4 \times b_1$. In fact, $a_4 \times b_1$ is the generator which survives in $\widehat{HF}(S^3)$, and hence it must have absolute grading equal to zero.

This completes the verification of Equation (30).

To compute higher differentials, we can no longer argue that the link complex must be E_2 -collapsed. However, we still have a number of constraints: The two homological

projections give the homology of the trefoil and the unknot respectively, and the total homology must have rank two (in dimensions 0 and 1). These constraints suffice to determine the higher differentials uniquely. We have illustrated these in Figure 12.

Reverse the orientation of one of the two components. Again, we obtain a chain complex which is uniquely determined by these constraints. We have illustrated the unique solution in Figure 13.

References

- [1] **F Bourgeois, Y Eliashberg, H Hofer, K Wysocki, E Zehnder**, *Compactness results in symplectic field theory*, *Geom. Topol.* 7 (2003) 799–888 MR2026549
- [2] **N M Dunfield, S Gukov, J Rasmussen**, *The Superpolynomial for Knot Homologies*, *Experiment. Math* 15 (2006) 129–160 MR2253002
- [3] **Y Eliashberg, A Givental, H Hofer**, *Introduction to symplectic field theory*, from: “GAFA 2000 (Tel Aviv, 1999)”, *Geom. Funct. Anal. Special Volume, Part II* (2000) 560–673 MR1826267
- [4] **A Floer**, *Morse theory for Lagrangian intersections*, *J. Differential Geom.* 28 (1988) 513–547 MR965228
- [5] **K Fukaya, Y G Oh, H Ohta, K Ono**, *Lagrangian intersection Floer theory-anomaly and obstruction*, preprint 487 (2000) 488
- [6] **B Gornik**, *Note on Khovanov link cohomology* (2004) arXiv:math.QA/0402266
- [7] **M Gromov**, *Pseudoholomorphic curves in symplectic manifolds*, *Invent. Math.* 82 (1985) 307–347 MR809718
- [8] **E-N Ionel, T H Parker**, *Relative Gromov-Witten invariants*, *Ann. of Math. (2)* 157 (2003) 45–96 MR1954264
- [9] **M Khovanov**, *A categorification of the Jones polynomial*, *Duke Math. J.* 101 (2000) 359–426 MR1740682
- [10] **M Khovanov, L Rozansky**, *Matrix factorizations and link homology II* arXiv:QA/0505056
- [11] **M Khovanov, L Rozansky**, *Matrix factorizations and link homology* (2004) arXiv:math.QA/0401268
- [12] **P B Kronheimer, T S Mrowka**, *Monopoles and contact structures*, *Invent. Math.* 130 (1997) 209–255 MR1474156
- [13] **E S Lee**, *The support of the Khovanov’s invariants for alternating knots* (2002) arXiv:math.GT/0201105
- [14] **A-M Li, Y Ruan**, *Symplectic surgery and Gromov-Witten invariants of Calabi-Yau 3-folds*, *Invent. Math.* 145 (2001) 151–218 MR1839289

- [15] **R Lipshitz**, *A Cylindrical Reformulation of Heegaard Floer Homology* (2005) arXiv:math.SG/0502404
- [16] **D McDuff, D Salamon**, *J-holomorphic curves and quantum cohomology*, University Lecture Series 6, American Mathematical Society, Providence, RI (1994) MR1286255
- [17] **J Milnor**, *Whitehead torsion*, Bull. Amer. Math. Soc. 72 (1966) 358–426 MR0196736
- [18] **P Ozsváth, Z Szabó**, *Heegaard Floer homology and alternating knots*, Geom. Topol. 7 (2003) 225–254 MR1988285
- [19] **P Ozsváth, Z Szabó**, *Knot Floer homology and the four-ball genus*, Geom. Topol. 7 (2003) 615–639 MR2026543
- [20] **P Ozsváth, Z Szabó**, *Holomorphic disks and knot invariants*, Adv. Math. 186 (2004) 58–116 MR2065507
- [21] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) 159 (2004) 1027–1158 MR2113019
- [22] **P Ozsváth, Z Szabó**, *Link Floer homology and the Thurston norm* (2006) arXiv:math.GT/0601618
- [23] **J A Rasmussen**, *Floer homology of surgeries on two-bridge knots*, Algebr. Geom. Topol. 2 (2002) 757–789 MR1928176
- [24] **J A Rasmussen**, *Floer homology and knot complements*, PhD thesis, Harvard University (2003) arXiv:math.GT/0306378
- [25] **J A Rasmussen**, *Khovanov homology and the slice genus* (2004) arXiv:math.GT/0402131
- [26] **D Rolfsen**, *Knots and links*, Mathematics Lecture Series 7, Publish or Perish, Houston, TX (1990) MR1277811 Corrected reprint of the 1976 original
- [27] **M Thistlethwaite**, *Link table* Available at http://katlas.math.toronto.edu/wiki/The_Thistlethwaite_Link_Table
- [28] **V Turaev**, *Torsions of 3-manifolds*, from: “Invariants of knots and 3-manifolds (Kyoto, 2001)”, (T Ohtsuki, et al, editors), Geom. Topol. Monogr. 4 (2002) 295–302 MR2002617

Department of Mathematics, Columbia University
New York, NY 10027, USA

Department of Mathematics, Princeton University
New Jersey 08544, USA

petero@math.columbia.edu, szabo@math.princeton.edu

Received: 3 February 2003 Revised: 9 November 2007