

Heegaard–Floer homology and string links

LAWRENCE ROBERTS

We extend knot Floer homology to string links in $D^2 \times I$ and to d -based links in arbitrary three manifolds. As with knot Floer homology we obtain a description of the Euler characteristic of the resulting homology groups (in $D^2 \times I$) in terms of the torsion of the string link. Additionally, a state summation approach is described using the equivalent of Kauffman states. Furthermore, we examine the situation for braids, prove that for alternating string links the Euler characteristic determines the homology, and develop similar composition formulas and long exact sequences as in knot Floer homology.

57M27; 57M25

1 Introduction

In [13] P Ozsváth and Z Szabó use the technology of Heegaard–Floer homology to refine the Alexander–Conway polynomial of a marked knot in S^3 . In particular, they define knot-Floer homology groups for relative Spin^c structures that correspond to the terms in the polynomial: the Euler characteristic of the homology with rational coefficients corresponding to i gives the coefficient of t^i . They have since shown that the non-vanishing of these groups characterizes the genus of the knot, [14]. In [10] they employ Kauffman’s state summation approach to the Alexander–Conway polynomial to give a concrete realization of the graded generators of the chain complex for each i (though not, unfortunately, for the differential). Furthermore, the techniques extend to *null-homologous* knots in an arbitrary three manifold, Y , where the knot may also be interpreted as giving a filtration of the Heegaard–Floer chain groups for Y that is also an invariant of the isotopy class of the knot. They also refine the one variable Alexander–Conway polynomial of an m -component link in S^3 by converting the link, in a way preserving isotopy classes, to a null-homologous knot in $\#^{m-1} S^1 \times S^2$.

In this paper we simultaneously generalize the preceding picture in two ways: first, by removing the restriction on the homology class of the embedding, and second, by defining chain complexes for a “string link,” thereby obtaining the Heegaard–Floer analog of the *multi-variable* torsion of the string link for its universal Abelian covering

space. There is a relationship between the torsion of a string link and the multi-variable Alexander polynomial of a simple link closure of the string link. Recently, P Ozsváth and Z Szabó have announced a version of Heegaard–Floer homology for links which enhances the multi-variable Alexander polynomial, [16], as the knot Floer homology enhanced the single-variable version. While there should be a relationship between the string link homology and the link homology, it should be noted that the string link homology appears to be a different beast; it does not share the symmetry under change of the components’ orientation, for example. On the other hand, we will show that in most respects the string link homology is a natural generalization of the knot Floer homology.

In general, we start with a basing of a link in a three-manifold, Y : we require an oriented disc, D , embedded in Y so that the link components intersect the disc once, positively. This configuration is known as a d –base, following the work of N Habegger and X S Lin, [4]. There are many d –basings of the same link, and our invariant will be sensitive to these. In S^3 there is a more perspicuous description of the configuration, called a string link.

Definition 1.1 Choose k points p_1, \dots, p_k in D^2 . A k –stranded string link in $D^2 \times I$ is a proper embedding, $\coprod_{i=1}^k f_i$ of $\coprod_{i=1}^k I_i$ into $D^2 \times I$, where $f_i: I_i \rightarrow D^2 \times I$, such that $f_i(0) = p_j \times 0$ and $f_i(1) = p_s \times 1$. The string link is called “pure” if $j = s$ for each interval.

A neighborhood of a d –base is a copy of $D^2 \times I$, and its complement in S^3 is also a copy of $D^2 \times I$. In the latter $D^2 \times I$, the d –based link appears as k copies of I extending from one end to the other. Thus, for general three-manifolds we can present our d –based link as a string link in $D^2 \times I$ along with a framed link diagram representing Y .

For a d –based link, $\mathbb{L} \cup D$, embedded in a general three-manifold, Y , we measure the homology of the components of the embedding by a lattice, Λ , in \mathbb{Z}^k consisting of all vectors of the form $([h] \cap L_1, \dots, [h] \cap L_k)$, where $[h] \in H_2(Y; \mathbb{Z})$ and the L_i are the components of the associated link. We can now state our main theorem.

Theorem 1.2 Let $\mathbb{L} \cup D$ be a d –based link in Y . Then, for each Spin^c structure, \mathfrak{s} , on Y , there is a relatively \mathbb{Z}^k / Λ –indexed Abelian group $\widehat{HF}(Y, \Gamma; \mathfrak{s})$. This group can be decomposed into isotopy invariant factors corresponding to elements of an affine space isomorphic to \mathbb{Z}^k / Λ .

When $\Lambda \equiv 0$, most of the results for knots transfer straightforwardly. In particular, the presence of the d –based link imposes a filtration upon the Heegaard–Floer chain

complexes of the ambient manifold, that is filtered chain homotopy invariant up to isotopy of the based link. The case where $\Lambda \neq 0$ occurs naturally when trying to generalize this picture. Furthermore, the theory can be extended to the other flavors of Heegaard–Floer homology.

Note In [13] the knot-Floer homology is denoted by $\widehat{HFK}(Y, K; \cdot)$. We will assume that the presence of $\Gamma = \mathbb{L} \cup D$ or K implies the use of the data determined by that object. Thus, we will use $\widehat{HF}(Y, K; \cdot)$ for the knot-Floer homology. When we wish to refer to the Heegaard–Floer homology of the ambient three-manifold (ignoring the information provided by K) we will simply omit reference to the knot or bouquet. However, these are not relative homology groups according to the classical axioms, nor do they solely depend upon the complement of the knot or link.

When $\mathbb{L} \cup D \subset S^3$, we may restate the theorem in terms of the associated string link, S .

Corollary 1.3 *For each $\bar{v} \in \mathbb{Z}^k$ there is an isotopy invariant $\widehat{HF}(S; \bar{v})$ of the string link S .*

In Section 8, we realize the generators of the chain complex for a string link S in $D^2 \times I$ from a projection of S . They are identified with a sub-set of maximal forests – satisfying specific constraints imposed by the meridians – in a planar graph constructed from the projection of S . This description generalizes the description of generators, their indices, and their gradings given by P Ozsváth and Z Szabó in [10].

Lemma 1.4 *There are vector weights assigned to crossings so that for each tree, adding the weights calculates the index, \bar{v} , for the corresponding generator. Furthermore, there are weights assigned to crossings which likewise calculate the grading for each generator.*

For the specific weights see Figure 6. This lemma requires a generalization of L Kauffman’s “Clock Lemma,” [5], to maximal forests that describes the connectivity of the set of maximal forests in a planar graph under two natural operations.

The weights and gradings are enough to form the Euler characteristic of the homology groups with rational coefficients, which is related to the Alexander–Conway polynomials of the link components in Section 8.4. However, the Euler characteristic can also be interpreted as a polynomial arising from the first homology of a covering space. We let $X = D^2 \times I - \text{int}(N(S))$ and $E = \partial X - D^2 \times \{0\}$. Consider the \mathbb{Z}^k -covering space, \widetilde{X} , determined by the Hurewicz map $\pi_1(X) \rightarrow H_1(X; \mathbb{Z}) \cong \mathbb{Z}^k$. Let \widetilde{E} be the

pre-image of E under the covering map. Then there is a presentation, (Litherland [9]), $\mathbb{Z}^k \xrightarrow{M} \mathbb{Z}^k \rightarrow H_1(\tilde{X}, \tilde{E}; \mathbb{Z}) \rightarrow 0$ of $\mathbb{Z}[t_1^{\pm 1}, \dots, t_k^{\pm 1}]$ -modules whose 0th elementary ideal is generated by $\det M$, a polynomial called the *torsion* of the string link (Kirk–Livingston–Wang [7]).

Theorem 1.5 *Let S be a string link. The Euler characteristic of $\widehat{HF}(S; \bar{v}; \mathbb{Q})$ is the coefficient of $t_1^{v_1} \cdots t_k^{v_k}$ in a polynomial $p(t_1, \dots, t_k)$ describing the torsion of the string link, $\tau(S)$.*

R Litherland appears to have originated the study of the module $H_1(\tilde{X}, \tilde{E}; \mathbb{Z})$ as a source of Alexander polynomials, [9]. He used it to study generalized θ -graphs, which, once we pick an edge, correspond to the string links above.

Many results follow from trying to replicate known properties of the torsion. Braids are a special sub-class of string links, for which it is known that the torsion is always trivial. Likewise, we can show that the string link homology is trivial.

Lemma 1.6 *If the string link S is isotopic to a braid, then $\widehat{HF}(S) \cong \mathbb{Z}_{(0)}$ where $\widehat{HF}(S) \cong \bigoplus_{\bar{v}} \widehat{HF}(S; \bar{v})$.*

The subscript in $\mathbb{Z}_{(0)}$ designates the grading. This result should be likened to the analogous result for 1-stranded string links, ie marked knots, that are also braids: the unknot has trivial knot Floer homology. While string links are usually considered up to isotopy fixing their endpoints on $D^2 \times \{0, 1\}$, this result has the implication that our invariant will be unchanged if we move the ends of the strings on $D^2 \times \{0, 1\}$.

Furthermore, as in Ozsváth–Szabó [10], alternating string link projections possess an important property.

Lemma 1.7 *Alternating string links in $D^2 \times I$ have trivial differential in each index, \bar{v} , for the Heegaard decomposition arising from an alternating projection.*

The proof may be found in [Section 11](#).

As P Ozsváth and Z Szabó can extend the knot Floer homology to links, we may extend the constructions for string links to a sub-class of colored tangles in $D^2 \times I$. For a tangle, T , we allow closed components, in addition to open components requiring that the open components independently form a string link. To each open or closed component we assign a color $\{1, \dots, k\}$ which corresponds to the variable t_i used for that component. We require that each color be applied to one and only one open component. We may

then use the colors to construct a string link, $S(T)$, in a second manifold, $\# S^1 \times S^2$, where n is the number of closed components in T . The isotopy class of $S(T)$ is determined by that of T in $D^2 \times I$, allowing us to consider $\widehat{HF}(Y, S(T); \mathfrak{s}_0, i)$ as an isotopy invariant of T . With this definition, we may extend the skein exact sequence of [13] to crossings where each strand has the same color.

Finally, we analyze how the homology changes for three types of string link compositions. Each has the form of connect sum in its Heegaard diagram, and the proofs roughly follow the approach for connect sums taken by P Ozsváth and Z Szabó. We picture our three-manifolds as given by surgery on framed links in $D^2 \times I$ with an additional string link, S , with k components. Given such diagrams for S_1 in Y_1 and S_2 in Y_2 we may 1) place them side by side to create a string link, $S_1 + S_2$ with $k_1 + k_2$ components, 2) when $k_1 = k_2$ we may stack one diagram on top of the other (as with composition of braids) to obtain the string link $S_1 \cdot S_2$, and 3) we may replace a tubular neighborhood, ie a copy of $D^2 \times I$, of the i th strand in S_1 with the entire picture for S_2 to obtain a string satellite, $S_1(i, S_2)$. The analysis of the second type proceeds differently than in [13]: we consider it as a closure of $S_1 + S_2$ found by joining the ends of S_1 on $D^2 \times \{1\}$ with the ends of S_2 on $D^2 \times \{0\}$ in a particular way. We prove the following formulas for the homologies, where $\mathfrak{s} = \mathfrak{s}_0 \# \mathfrak{s}_1$:

$$\begin{aligned} & \widehat{HF}(Y, S_1 + S_2; \mathfrak{s}, [\bar{j}_0] \oplus [\bar{j}_1]) \\ & \cong H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\bar{j}_0]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{j}_2])) \\ & \widehat{HF}(Y, S_1 \cdot S_2; \mathfrak{s}, [\bar{k}]) \\ & \cong \bigoplus_{[\bar{k}_1] + [\bar{k}_2] = [\bar{k}] \pmod{\Lambda}} H_*(\widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{k}_1]) \otimes \widehat{CF}(Y_2, S_2; \mathfrak{s}_2, [\bar{k}_2])) \\ & \widehat{HF}(Y, S_1(i, S_2); \mathfrak{s}, [(l_1, \dots, l_{k_1+k_2-1})]) \\ & \cong \bigoplus_{[\bar{v}'] + [\bar{w}'] = [\bar{l}] \pmod{\Lambda'}} H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\bar{v}]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{w}])), \end{aligned}$$

where $\bar{v}' = (v_1, \dots, v_{i-1}, v_i, \dots, v_i, v_{i+1}, \dots, v_{k_1})$, repeating v_i a total of k_2 -times, and $\bar{w}' = (0, \dots, 0, w_1, \dots, w_{k_2}, 0, \dots, 0)$ with zero entries except for places $i, \dots, i + k_2 - 1$. and $\Lambda' = \Lambda + \bar{0} \oplus \Lambda_1 \oplus \bar{0}$. Although the isotopy class of $S_1(i, S_2)$ depends upon a framing to specify the gluing, when considered using the more flexible isotopy described above, this framing becomes irrelevant.

Note We do not address issues of orientation of moduli spaces in the paper. However, nothing we say will alter the existence of the coherent orientations. It is merely

convenient to suppress this information. As usual we may work with $\mathbb{Z}/2\mathbb{Z}$ -coefficients to avoid these issues.

2 String links and Heegaard diagrams

2.1 String links in general three-manifolds

In this section, we describe the topological input for the paper. Let Y denote a closed, oriented three-manifold.

Definition 2.1 A d -base for an oriented link \mathbb{L} is an oriented disc, D , embedded in Y whose interior intersects each component of \mathbb{L} exactly once, positively.

By thickening the disc D we obtain an embedded ball with the structure of a cylinder, $D^2 \times I$. The complement of this ball together with its intersection with \mathbb{L} will be called a “string link” for Y . We give several examples that will be important in the remainder of the paper.

Example 1 For knots in S^3 a d -basing amounts to picking a point on the knot, the intersection point with D . This is a marked knot as used in [13]. The string link associated to a knot comes from dividing the knot at that point and pulling the ends apart to obtain a knotted strand in $D^2 \times I$.

Example 2 A d -base for a link in S^3 corresponds to the usual definition of a string link.

Definition 2.2 Choose k points p_1, \dots, p_k in D^2 . A k -stranded string link in $D^2 \times I$ is a proper, tame embedding, $\coprod_{i=1}^k f_i$ of $\coprod_{i=1}^k I_i$ into $D^2 \times I$, where $f_i: I_i \rightarrow D^2 \times I$, such that $f_i(0) = p_j \times 0$ and $f_i(1) = p_s \times 1$. The string link is *pure* if $j = s$ for each interval. We orient the strands “down” from 1 to 0.

To each strand of a string link, we may associate a knot: ignore the other strands and join the two ends of our strand with an unknotted arc in the complement of $D^2 \times I$. Likewise we may associate a link to a pure string link by using k unknotted, unlinked arcs lying in the projection plane, as for the closure of a braid. Furthermore, by retaining $D^2 \times \{0\}$, oriented as the boundary of $D^2 \times I$, we have an embedded disc which intersects the strands of the link in one point with $L_i \cap D = +1$. This is a d -base in the language of Habegger and Lin, [4].

Example 3 To draw the Heegaard diagram for a string link in Y , we think of $D \cup \mathbb{L}$ embedded in a framed link diagram for Y in S^3 . By inverting a neighborhood of the disc, we can present the string link in Y by a string link in a framed link diagram in $D^2 \times I$. Handleslides of the framed components, or blowing up/down will preserve the isotopy class of this string link.

When we refer to a string link, we will mean a configuration of strands and framed knots in $D^2 \times I$, usually as in [Example 2](#) with the strands ending on prescribed points. String links are typically considered up to isotopies which pointwise fix the set $D^2 \times \{0, 1\}$. In this sense, the set of braids is a subset of the set of string links in $D^2 \times I$. However, we will consider string links up to a different equivalence.

Definition 2.3 A flexible isotopy of a string link, is an isotopy which setwise preserves the components $D^2 \times \{0, 1\}$.

As an example of the difference, any braid is flexibly isotopic to the trivial braid with the same number of strands. The groups we define later will be invariant under this less rigid form of isotopy. Likewise, we can consider d -based links up to isotopy of $D \cup \mathbb{L}$ as a complex, which retains the intersection data, or more flexibly by allowing the strands of \mathbb{L} not to return to the same point, and moving their ends independently on, but not between, the two sides of D .

2.2 Heegaard diagrams for d -based links

We describe the general approach to Heegaard diagrams subordinate to a d -based link, $D \cup \mathbb{L}$, and the Heegaard equivalences between them. In particular, we will be concerned with retaining the data provided by D .

Let Y be a three-manifold. Assume we have a Heegaard decomposition, $Y = H_\alpha \cup_{\Sigma^g} H_\beta$, where $\partial H_\alpha = \Sigma = -\partial H_\beta$ and the gradient flow corresponds to the outward-pointing normal to H_α . If we choose decompositions of each handlebody into a zero cell and g one handles, then we can specify this decomposition by $\{\alpha_i\}_{i=1}^g$, the co-cores of the one handles for a handle decomposition in H_α sitting in Σ , and likewise for $\{\beta_i\}_{i=1}^g$ and H_β . This configuration of curves in Σ is called a Heegaard diagram. We will denote the data of the Heegaard diagram by $(\Sigma, \{\alpha_i\}_{i=1}^g, \{\beta_i\}_{i=1}^g)$ or $\Sigma_{\alpha\beta}$ and the underlying three-manifold by $M(\Sigma, \{\alpha_i\}_{i=1}^g, \{\beta_i\}_{i=1}^g)$.

To make our diagram reflect $\mathbb{L} \cup D$, we additionally require that our Heegaard data satisfies the following.

- (1) $M(\Sigma, \{\alpha_i\}_{i=k+1}^g, \{\beta_i\}_{i=1}^g)$ is $Y^3 - N(\mathbb{L} \cup D)$.

- (2) The complement of $M(\Sigma, \{\alpha_i\}_{i \neq j}, \{\beta_i\}_{i=1}^g)$ for $j \leq k$, is homeomorphic to a tubular neighborhood of L_j in Y with α_j as an oriented meridian for this tubular neighborhood.
- (3) There is a smooth embedding of the closed unit disc with image $D' \subset \Sigma$ such that $D' \cap \bigcup_{i=1}^g \beta_i = \emptyset$ while $D' \cap \bigcup_{i=1}^g \alpha_i$ consists of k disjoint closed intervals, $I_j \subset \Sigma$, $j = 1, \dots, k$ with $I_j \subset \alpha_j$. The string link formed from D' glued to the attaching discs for $\alpha_1, \dots, \alpha_k$ and the components of \mathbb{L} from (1) is isotopic to $\mathbb{L} \cup D$.

The first two conditions require that the diagram be subordinate to the link. The third requires that there be a disc in the Heegaard data that, when attached to the compression discs for α_i , $i \leq k$ and pushed into H_α , produces a d -basis for \mathbb{L} , isotopic to the original one. The decomposition associated to such a diagram is said to be *subordinate* to $\mathbb{L} \cup D$. Furthermore, a diagram with such a choice of α and a disc D' determines a string link as in item (3). Note that D' must be oriented opposite to Σ , and that D' will usually not be a full component of $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$.

We may relate diagrams subordinate to an embedding of $\mathbb{L} \cup D$ in Y by the following lemma.

Lemma 2.4 (Ozsváth–Szabó [15, Lemma 4.5]) *Let Y be a closed, oriented three-manifold. Let $\mathbb{L} \cup D \subset Y$ be an embedded d -base. Then there is a Heegaard diagram subordinate to $\mathbb{L} \cup D$ and any two such subordinate diagrams may be connected by a sequence of the following moves:*

- (1) *handleslides and isotopies among the elements of $\{\alpha_i\}_{i=k+1}^g$,*
- (2) *handleslides and isotopies of $\{\beta_i\}_{i=1}^g$,*
- (3) *stabilization introducing $\alpha_{g+1}, \beta_{g+1}$ intersecting in a single point,*
- (4) *isotopies of $\{\alpha_i\}_{i=1}^k$ and handleslides of them over elements of $\{\alpha_i\}_{i=k+1}^g$,*

where we disallow isotopies and handleslides of any attaching circles resulting in a curve intersecting the disc D' .

Proof Let $\Gamma = \mathbb{L} \cup D$. Then $N(\Gamma)$ is a handlebody where the disc D is thickened to a copy of $D \times I$ to which handles are added for each component of \mathbb{L} . One component of $D \times \partial I$ is oriented identically with the boundary of the handlebody, the other is oriented opposite, D'' . Inside D'' , there are k smaller discs where the handles for each component of \mathbb{L} are attached. Choose a point in the D'' away from these attaching discs and pick k disjoint arcs from this point to the boundary of each of the attaching

discs. D' is then a closed disc neighborhood that is a thickening of the point and the arcs whose boundary has k –arcs, one on the boundary of each of the attaching discs. If we take the boundary of each attaching discs to be one of the α_i with $i = 1, \dots, k$, we can then extend to a Heegaard decomposition subordinate to $D \cup \mathbb{L}$. The additional α and β may be chosen to avoid D' as D' is contractible in Σ . Thus, a diagram for Γ always exists.

Following the proof of Ozsváth–Szabó [11, Proposition 7.1], we see that any β isotopy whose support intersects D' , and which starts and ends with D' not intersecting any β –curve, may be accomplished by handleslides not intersecting D' . Since $D' \subset \Sigma$ is a disc, it suffices to consider isotopies taking an arc of β_i so that the entire disc moves to the other side of β_i from which it started (locally). If we surger all the β –curves except β_i we obtain a torus with β_i and D' still present, and a collection of disjoint discs, pairs of which represent each of the other β –curves, and which are disjoint from D' and β_i . If we do the same to the configuration of β –curves in Σ at the end of the isotopy we obtain another configuration in the torus, with D' and the other discs unchange, and β_i now on the other side of the disc. There are two ways to move between these configurations in the torus: 1) isotope β_i across D' , corresponding to the original isotopy in Σ , or 2) isotope β_i around the torus opposite to the isotopy direction, until it comes to the end configuration. In the latter case, each isotopy across one of the surgery discs corresponds in Σ to a handleslide whose support does not intersect D' in Σ . As a result the choice made in pushing the β off D' does not affect the equivalence class of the diagram.

Given the diagrams for two isotopic embeddings of Γ , we must see that they can be related by the moves described in the lemma. These moves preserve the region D' in the original diagram. D' and the attaching discs for the meridians act as the disk for the d –base. On the other hand, the isotopy carries a neighborhood of the disc into a neighborhood of the new disc; these neighborhoods are all homeomorphic and may be used as a 3–handle for each of the Heegaard decompositions. That the isotopy preserves \mathbb{L} outside this ball allows us to keep the meridians constant and consider only the additional handles describing $Y - N(\Gamma)$.

If we consider $\partial N(\Gamma)$ as $\partial_+(Y - N(\Gamma))$ then the existence of the Heegaard diagram follows from the existence of a relative Morse function that is equal to 1 on the boundary and that the isotopy class of $\{\alpha_i\}_{i=1}^k$ is determined by their being meridians of the knots determined by the strands of \mathbb{L} . As usual, we may cancel 1–handles with 0–handles until there is only one 0–handle. Similarly we may cancel off 3–handles until there are none. The relative version of Cerf’s theorem states that any two such diagrams can be linked through the first three moves and the introduction of new index 0/1 canceling pairs or new index 2/3 canceling pairs.

However, we would like to ensure that the path can be chosen through diagrams with only one index 0 handle and no index 3 handles. As we introduce a new 3–handle, we also introduce a canceling 2–handle. The new 2–handle will have one end of its co-core on $\partial N(\Gamma)$, since there are no other 3–handles. In a diagram for Y with the prescribed meridians, this 2–handle has a core that is a homological linear combination of the α . If we cut Σ along $\{\alpha_{k+1}, \dots, \alpha_g\}$, and cap the new boundaries with discs, the image of the core will be null-homotopic: it will be homotopic to the boundary of the B^2 at the end of the co-core lying on $\partial N(\Gamma)$. Since it is null-homotopic, the core cannot have a non-zero coefficient for a meridian for its homology class. Thus, it is linear combination of $\{\alpha_{k+1}, \dots, \alpha_g\}$. According to [11, Lemma 2.3], the core curve can be obtained as the image of a $\{\alpha_{k+1}, \dots, \alpha_g\}$ under handleslides. Thus, any diagram obtained from the diagram after a 2/3 pair is added, could be obtained from the old diagram, as handleslides over the new core can be given by handleslides over α , not using any meridians. We may bypass 2/3–handle pairs. The same argument applies for 0–handles, as there are fewer restrictions on the β .

Once again allowing the meridians to move, the last equivalence follows from considering the handle decomposition after surgering out $\{\alpha_i\}_{i=k+1}^g$. The meridians are determined up to isotopy, and as those isotopies cross surgery discs there are corresponding handleslides. Likewise, two curves for the same meridian, abutting the same disc D' at a specified point, will be isotopic. This takes care of all our choices, so any two diagrams subordinate to the same isotopy class of string link can be related with the moves in the lemma. \square

Note If we allow equivalences that intersect D' , the moves above would still preserve the handlebody neighborhood of $\mathbb{L} \cup D$ since we cannot slide over the meridional discs. This is weaker than preserving the isotopy class of $\mathbb{L} \cup D$. The additional restriction imposed by D' prevents an isotopy from twisting the ends of the same strand at the intersection point with D . We may still disjoin then ends of the link components and braid the ends on one side of the disc independently of the other. However the data from D prevents the ends from moving between the different sides of D .

2.3 Specific Heegaard diagrams for string links

2.3.1 Heegaard diagrams from projections We describe a general procedure for drawing the Heegaard diagram of a string link presented in $D^2 \times I$, along with a framed link specifying the three-manifold Y . We start with the situation in [Example 2](#) above.

Consider $D^2 \times I$ embedded inside S^3 . We isotope the string link so that projection onto the plane formed from a diameter for D^2 times I is generic. Such a projection

provides the data for a Heegaard decomposition of S^3 . For a string link whose strands are oriented downwards, we may draw a Heegaard diagram as in Figure 1 according to the process described presently.

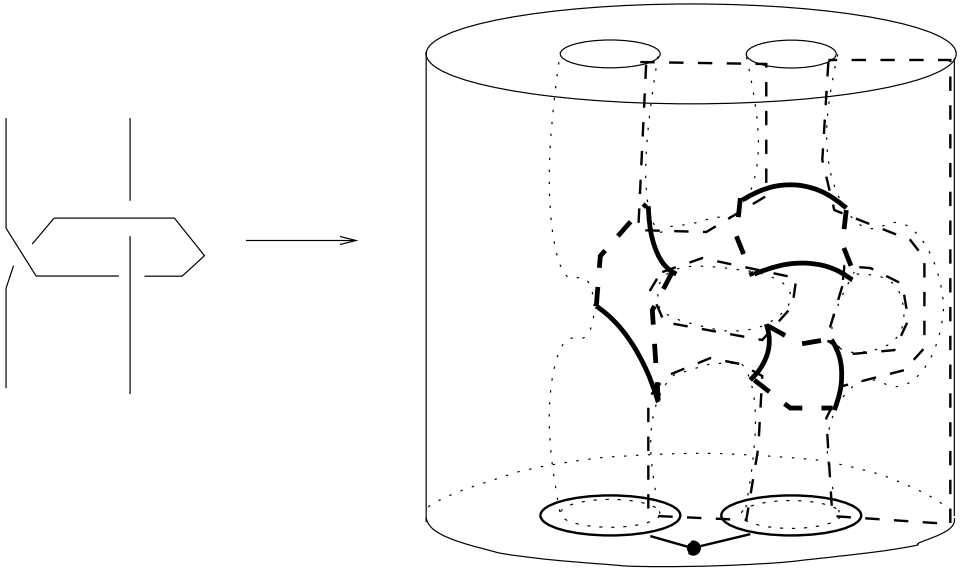


Figure 1: The Heegaard Diagram for a string link in $D^2 \times I$. We place the meridians at the bottom of the diagram and α curves, depicted as thick black curves, at each of the crossings, reflecting the type of crossing. The β -curves, which are dashed, correspond to faces of the projection. The graph in $D^2 \times \{0\}$ illustrates that used in constructing marked diagrams.

Take a small tubular neighborhood of each strand in $D^2 \times I$ and glue it to the three ball that is $S^3 \setminus (D^2 \times I)$. This is a handlebody to which we attach 1 handles at each of the crossings. These handles occur along the axis of the projection, between the two strands; when we compress the handle to obtain the tubular neighborhood of an “X”, the attaching circles appear as in Figure 1, one for each type of crossing, cf [10]. The attaching circles for the handles from the strands are placed in $D^2 \times \{0\}$ as meridians for each strand.

The complement of this in S^3 is also a handlebody. It can be described by taking two 0–handles in $D^2 \times I$ above and below the plane we projected onto (thought of as cutting $D^2 \times I$ through the middle). We then attach handles through the faces of the image of the projection: a copy of the D^2 factor in the handle should be the face. We use the 1–handle for the face on the leftmost side of the diagram, called U , to cancel

one of the 0–handles. As these two handlebodies have the same boundary, they must have the same number of handles, and the decomposition is a Heegaard decomposition.

To extend to the case including framed closed components, we first find a generic projection of both the strands and the framed components. We choose paths from each of the framed components to one of the strands, placed generically. We can then find a Heegaard diagram as above. However, on each of the closed components we will use the framing, not the meridian, to place the α curves, oriented opposite the longitudes from S^3 . We need the choice of paths in order to have a Heegaard decomposition, but if a component crosses one of the strands, we can use the path along the projection direction. Paths only need to be specified for framed components that are disjoint from the string link in the projection.

Lemma 2.5 *Let S be a string link for a three-manifold specified by a collection of framed knots, thought of being in $D^2 \times I$. Then the Heegaard diagram after*

- (1) *sliding a strand over a framed component, or*
- (2) *sliding a framed component over another component, or*
- (3) *adding a unlinked ± 1 framed unknot,*

is equivalent to the Heegaard diagram before.

Proof Sliding a strand in the string link over a framed curve produces a new string link which is isotopic to the first in Y . Thus the two Heegaard diagrams are related by Heegaard equivalences. Likewise, handleslides of framed components over each other can be effected by a bouquet with a path joining the two components and adding ± 1 framed unlinked, unknots (blowing up/down) can be effected using the reduced Heegaard equivalences.

Finally, how the framed components are joined to the strands does not change the equivalence class of the diagram. We describe specifically how to use the equivalences to see this conclusion. The argument comes from [15], for the same result for a bouquet, with minor alteration. Consider two distinct arcs, s_1 and s_2 , joining a framed component to Γ , and form a regular neighborhood of the graph provided by $\Gamma \cup s_1 \cup s_2$. We extend this to a diagram for S^3 (by adding handles for crossings, etc). We draw a subordinate diagram using s_1 to attach the framed component by placing α_{k+1} as a meridian to s_2 . To obtain a diagram subordinate to the second choice of paths we erase α_{k+1} and replace it with α'_{k+1} , a meridian to s_1 . We surger out all the α_i for $i > k + 1$ to obtain a genus $k + 1$ surface. We wish to see that α_{k+1} and α'_{k+1} are isotopic; if they are we may move the corresponding curves in the original diagram through isotopies

and handleslides one into the other. The two curves, along with some non-meridional α , bound a punctured torus coming from the framed component. After surgering, the other boundaries are filled, and surgering the framing attaching circle transforms the torus into a cylinder. Therefore, the α_{k+1} and α'_{k+1} now bound a cylinder which does not involve the disc D' . Each time the isotopy of curves determined by the cylinder crosses a disc coming from the surgered handles, there is a corresponding handleslide in the original picture. This provides a sequence of Heegaard equivalences that are allowed under the reduced equivalences of d -based links. Hence, using two different framed link descriptions for Y will not change the equivalence of the Heegaard diagrams. \square

Note In such diagrams, the meridians will intersect at most two β curves. This happens for all but one of the meridians in the construction above. By sliding one of these β curves over the other we can ensure that in our Heegaard diagrams the meridians each intersect only one β -curve, provided the β curves are distinct. However, we can introduce crossings of strands using the second Reidemeister move to ensure that they the β -curves are distinct. We will often assume that this simplification has been made.

2.3.2 Marked diagrams It is cumbersome to retain the disc D' in our Heegaard diagrams. Furthermore, since the restricted Heegaard equivalences eliminate handleslides over the meridians they often impede the simplification of the Heegaard diagram. We now give a different way to encode a d -based link in a Heegaard diagram that is more tractable.

Definition 2.6 A multi-pointed Heegaard diagram for Y is a Heegaard diagram $(\Sigma, \{\alpha_i\}_{i=1}^g, \{\beta_i\}_{i=1}^g)$ which specifies Y , and a collection of points w, z_1, \dots, z_k in $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$.

We consider multi-pointed diagrams up to the usual Heegaard equivalences with one exception. The supports of isotopies of curves in either $\{\alpha_i\}_{i=1}^g$ or $\{\beta_i\}_{i=1}^g$ are not allowed to contain any marked points. We then have the following relationship with flexible isotopy.

Proposition 2.7 A marked point diagram for Y specifies a Heegaard diagram for a d -based link embedded in Y , considered up to flexible isotopy. Furthermore, a Heegaard diagram for a d -based link specifies a multi-pointed diagram, and these two constructions are inverses up to Heegaard equivalence. The equivalences of Heegaard diagrams subordinate to a d -based link correspond to equivalences for a marked-point diagram.

Proof Using the the disc D' we may introduce marked points into Σ . We choose w to be in the interior of D' and z_i to be on the other side of α_i in the region of $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$ abutting the same segment as D' .

Conversely, for any Heegaard diagram with a choice of marked points w, z_1, \dots, z_k in $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$ we can construct a Heegaard diagram subordinate to a string link. First, choose paths from each z_i to w crossing only α . Then take neighborhoods of the gradient flow lines in H_β joining the index 0 critical point and the marked points in Σ , remove these neighborhoods from H_β and add them to H_α . The complement in H_β is still a handlebody, since we can see this as removing the neighborhoods of $k + 1$ radii in a closed three ball and then add the other handles to the boundary. Adding the neighborhoods to H_α creates a new handlebody $H_{\alpha'}$. The new α are meridians of the gradient flow lines, and the new β are loops following the flow line from w to the critical point, then to z_i and back along the path we chose in Σ , crossing only α . To find the region D' , take the portion of $\partial H_{\alpha'}$ coming from the newly added handles. This is a sphere with $k + 1$ discs removed. Find a point in this disc and join it by k arcs in this sphere to the meridians of the flow lines through the z_i . Now take a small enough neighborhood of this graph to recover D' . Adding the attaching discs produces the disc D regardless of the paths chosen in defining D' . Thus, a multi-pointed diagram gives us a diagram subordinate to a string link through the preferred disc D . Furthermore, if we use the preferred disc, D , to produce a multi-pointed diagram, after some handleslides of the α over the new meridians and stabilizations, we can de-stabilize the new α and β to obtain the original multi-pointed diagram.

This construction depends upon the choice of the new β -paths and the graph joining to the new α_i . However, if we surger all the β attaching circles in the multi-pointed diagram, the chosen paths become a star in S^2 joining z_i to w for all i . If two such stars are isotopic in the complement of the marked points, the resulting diagrams are equivalent. Each time the isotopy of a segment crosses a disc introduced by the surgery, we should think of our new β being slid over an old β . Braiding of the marked points in S^2 , carrying along the star, will not, in general, produce a diagram isotopic to the one with the star before braiding. However, these diagrams are both stabilizations of the same diagram and will thus still be equivalent. This discrepancy accounts for the use of flexible isotopy in the statement of the proposition.

As to the relationship between the equivalences in the two types of diagram, when going from a multi-pointed diagram to a string link diagram, performing an illicit isotopy over a marked point z_i corresponds to an illicit handleslide over the meridian of the strand corresponding to z_i , according to the previous construction. Were we to surger out the meridians, the point z_i would correspond to one of the two discs used to

replace that meridian (the other would be close to w inside the region D). Isotoping across it would be the same as a handleslide across a meridian. \square

2.3.3 Admissibility and d -based links In Heegaard–Floer homology, when $H_2(Y; \mathbb{Z}) \neq 0$, we must use diagrams submitting to certain admissibility requirements, [11]. We argue here that by presenting a d -based link as a string link in $D^2 \times I$, with an additional framed link defining Y , we can make the diagram admissible without disrupting the disc D or the structure of $D^2 \times I$. We make use of the lemmas in [11, Section 3]. Those readers willing to accept this proposition can safely skip this technical result.

Lemma 2.8 *Let $\mathbb{L} \cup D \subset Y$ be a d -based link. Let \mathfrak{s} be a Spin^c structure on Y . Then there is a strongly/weakly admissible diagram for $(Y, \mathfrak{s}, \mathbb{L} \cup D)$ presented as a framed link in the complement of a string link in $D^2 \times I$.*

Proof Suppose Y is presented as surgery on a link in S^3 , and Γ is $\mathbb{L} \cup D$ in this diagram. We adjust this, as above, to be a framed link diagram in $D^2 \times I$ with a string link. Recall that we must join the framed components to $\mathbb{L} \cup D$ by paths which do not touch $D^2 \times \{0\}$. With the framing curves as α -curves, this provides a diagram for Y . However, it need not be admissible; we may need to wind the attaching circles to make it so. We must ensure that the winding paths do not affect the discs $D^2 \times \{0\}$ or $D^2 \times \{1\}$. We make two observations.

- (1) First, any doubly periodic domain must have at least one boundary containing multiples of a framing curve. Otherwise, by replacing framing curves with the meridians of the link components, we would obtain a periodic domain in a diagram for S^3 . Furthermore, two periodic domains may not produce the same linear combination of framing curves in their boundaries.
- (2) Second, a meridian of a framed component may be chosen to intersect the framing curve which replaces it, once and only once, and intersect no other α . Each will, however, intersect at least one β curve in the projection. By Gompf–Stipsicz [3, Proposition 5.3.11], these meridians generate all of $H_1(Y; \mathbb{Z})$. By winding along them we may obtain intersection points representing any Spin^c structure; we have an intersection point which employs the framing curves intersecting the same β as the meridian.

These are the conditions necessary to draw the conclusion of Lemmas 5.2, 5.4, and 5.6 of [11]. These lemmas guarantee the results in the proposition. \square

3 Multi-Pointed Heegaard–Floer homology

3.1 Background and notation

The reader should consult [11] for the notation used below.

Let $(\Sigma, \{\alpha_i\}_{i=1}^g, \{\beta_i\}_{i=1}^g)$ and a choice of $w, z_1, \dots, z_k \in \Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$ be a multi-pointed strongly/weakly admissible Heegaard diagram for Y . We will denote the additional marked point data by Γ . Then $\{\alpha_i\}_{i=1}^g$ is a set of g disjoint, simple, closed curves in Σ whose images are linearly independent in $H_1(\Sigma; \mathbb{Z})$ and which bound compression discs in H_α . $\{\beta_i\}_{i=1}^g$ has analogous properties for H_β . We assume the curves in $\{\alpha_i\}_{i=1}^g$ and $\{\beta_i\}_{i=1}^g$ are in general position. choose a path of generic nearly symmetric almost-complex structures, J_s , on $\text{Sym}^g(\Sigma)$, in accordance with the restrictions in [11]. Furthermore, choose an equivalence class of intersection points, \mathfrak{s} , for Y and a coherent system of orientations for the equivalence class, [11].

For now we will consider the chain group $\widehat{CF}(Y; \mathfrak{s})$ as in [11]. When $c_1(\mathfrak{s}_w(\mathbf{x}))$ is not torsion we will employ the relative $\mathbb{Z}/\delta(\mathfrak{s})\mathbb{Z}$ -grading on the chain complexes, where

$$\delta(\mathfrak{s}) = \gcd_{\xi \in H_2(Y; \mathbb{Z})} \langle c_1(\mathfrak{s}), \xi \rangle.$$

If \mathfrak{s} is a torsion Spin^c structure, the chain complexes has an absolute \mathbb{Q} -grading, $\text{gr}_{\mathbb{Q}}$ defined in [15].

Finally, we introduce some extra notation: let $\mathcal{I}(\mathfrak{s})$ be the set of all intersection points, $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, which represent the Spin^c -structure, \mathfrak{s} , for the basepoint w .

3.2 Filtration indices

According to our conventions, the orientation on each α curve representing a meridian is the one induced from the attaching disk, oriented to intersect L_i positively: $L_i \cap D_i = +1$. If we choose a periodic domain, \mathcal{P} , representing the homology class $h \in H_2$ then $\partial\mathcal{P}$ may only contain multiples of the meridians, not segments on the meridians. With this orientation convention $(n_w - n_{z_i})(\mathcal{P}) = -L_i \cap h$. Thus, the quantity $n_w - n_{z_i}$ measures the number of times the i th meridian occurs in $\partial\mathcal{P}$.

Definition 3.1 $\Lambda_{(Y, \Gamma)} \subset \mathbb{Z}^k$ is the lattice consisting of the vectors

$$\{(n_{z_1}(\mathcal{P}), \dots, n_{z_k}(\mathcal{P})) \mid \mathcal{P} \text{ is a periodic domain}\}.$$

We will usually suppress the subscript, (Y, Γ) , when the context is clear. The additional marked points z_1, \dots, z_k then provide an additional index for the generators of $\widehat{CF}(Y, \Gamma; \mathfrak{s})$.

Definition 3.2 A filtration index, \overline{F} , for \mathfrak{s} , is a map $\overline{\mathcal{F}}: \mathcal{I}(\mathfrak{s}) \rightarrow \mathbb{Z}^k/\Lambda$ with the following properties.

- (1) Let $n_{\overline{z}}(\phi) = (n_{z_1}(\phi), \dots, n_{z_k}(\phi))$ and $n_{\overline{w}}(\phi) = (n_w(\phi), \dots, n_w(\phi))$ for $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ then

$$\overline{\mathcal{F}}(\mathbf{y}) - \overline{\mathcal{F}}(\mathbf{x}) = (n_{\overline{w}} - n_{\overline{z}})(\phi) \pmod{\Lambda}.$$

- (2) If z_i and z_j are in the same component of $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$ then the i th and j th coordinate maps, \mathcal{F}_i and \mathcal{F}_j , are equal.
- (3) If z_i is in the same component as w then $\mathcal{F}_i \equiv 0$, a constant map.

Since for a periodic domain $n_w(\mathcal{P}) = 0$ and $n_{\overline{z}}(\mathcal{P}) \in \Lambda$, adding a periodic domain \mathcal{P} to ϕ does not change the relation on the the right hand side of 1). Furthermore, adding or subtracting $[S]$ will change $n_{\overline{w}}(\phi)$ and $n_{\overline{z}}(\phi)$ by the same vector. Thus, the filtration difference between generators is well-defined by the relation above. However, we can obtain a new index by taking $\overline{\mathcal{F}} + v$ where $v \in \mathbb{Z}^k/\Lambda$ is a vector not disrupting properties 2) and 3). Thus we obtain only a relative \mathbb{Z}^k/Λ -index on $\widehat{CF}(Y, \mathfrak{s})$.

Example For null-homologous knots and torsion \mathfrak{s} there is a canonical choice of filtration index found from the first Chern class, [13]. Suppose all the knots in Y found by closing strands in S are null-homologous, and that we have a Heegaard diagram where there is a unique intersection point between each meridian and $\{\beta_i\}_{i=1}^g$. As we have noted previously, this can always be arranged. Let λ_i be a longitude for the closure of the i th strand, L_i . This curve can be realized in the Heegaard diagram for the string link as a curve crossing only one α -curve, the i th meridian. Interchanging the meridian with this longitude gives a Heegaard diagram for the manifold found by performing 0-framed surgery on L_i . To each intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ we can associate an intersection point, \mathbf{x}' , for the new Heegaard diagram, cf [13]. A Seifert surface for the closure of the i th strand becomes a doubly periodic domain, \mathcal{P}_i , in the new diagram. Following the argument in [13] shows that we may choose

$$\mathcal{F}_i(\mathbf{x}) = \frac{1}{2} \langle c_1(\mathfrak{s}_w(\mathbf{x}'), [\mathcal{P}_i]) \rangle$$

for our filtration index. This provides a canonical choice over different Spin^c -structures on Y , which the axiomatic description lacks, although it still depends on the choice of Seifert surface if Y is not a rational homology sphere.

3.3 The differential and the homology theory

According to [11] there is a differential, $\widehat{\partial}$ on $\widehat{CF}(Y; \mathfrak{s})$ defined by the linear extension of

$$\widehat{\partial}[\mathbf{x}] = \sum_{\mathbf{y} \in \mathcal{I}(\mathfrak{s})} \sum_{\phi \in D(\mathbf{x}, \mathbf{y})} \# \widehat{\mathcal{M}}(\phi)[\mathbf{y}]$$

where

$$D(\mathbf{x}, \mathbf{y}) = \{ \phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi) = 1, n_w(\phi) = 0 \}$$

and the signed count is made with respect to the choice of a coherent orientation system.

We can define a new chain complex $\widehat{CF}(Y, \Gamma; \mathfrak{s})$ by using

$$D'(\mathbf{x}, \mathbf{y}) = \{ \phi \in D(\mathbf{x}, \mathbf{y}) \mid n_{\bar{z}}(\phi) \equiv 0 \}$$

instead of D' . This still defines a differential since $n_{\bar{z}}(\phi_1 * \phi_2) \equiv 0$ implies, for ϕ_1, ϕ_2 with holomorphic representatives, that $n_{\bar{z}}(\phi_i) \equiv 0, i = 1, 2$, as holomorphicity requires positive multiplicity in every domain.

Since $\bar{\mathcal{F}}(\mathbf{y}) \equiv \bar{\mathcal{F}}(\mathbf{x}) - n_{\bar{z}}(\phi) \pmod{\Lambda}$ for every $\phi \in D(\mathbf{x}, \mathbf{y})$, a choice of $\bar{\mathcal{F}}$ splits the complex $\widehat{CF}(Y, \Gamma; \mathfrak{s})$ into a direct sum of complexes

$$\widehat{CF}(Y, \Gamma; \mathfrak{s}) \cong \bigoplus_{\bar{\mathbf{v}} \in \mathbb{Z}^k / \Lambda} \widehat{CF}(Y, \Gamma; \mathfrak{s}, \bar{\mathbf{v}})$$

where $\widehat{CF}(Y, \Gamma; \mathfrak{s}, \bar{\mathbf{v}})$ is generated by those $[\mathbf{x}]$ with

$$\bar{\mathcal{F}}(\mathbf{x}) \equiv (v_1, v_2, \dots, v_k) \pmod{\Lambda}.$$

In addition, when the lattice $\Lambda_{(Y, \Gamma)} \equiv 0$, a choice of index $\bar{\mathcal{F}}$ not only splits $\widehat{CF}(Y, \Gamma; \mathfrak{s})$, but it also \mathbb{Z}^k -filters $\widehat{CF}(Y, \mathfrak{s})$ by the relation

$$(j_1, \dots, j_k) < (j'_1, \dots, j'_k)$$

when $j_l \leq j'_l$ for all l with strict inequality for some j_l . We have $\bar{\mathcal{F}}(\widehat{\partial}[\mathbf{x}]) \leq \bar{\mathcal{F}}([\mathbf{x}])$ since $n_{z_i}(\phi) \geq 0$ on classes represented by a holomorphic disc. (The partial ordering on non-zero linear combinations of generators, $\sum \mathbf{y}_i \leq \sum \mathbf{x}_j$, occurs when every generator $\mathbf{y}_i \leq \mathbf{x}_j$ for each j). $\widehat{CF}(Y, \Gamma; \mathfrak{s})$ is then analogous to defining an E^0 page of a spectral sequence for this filtration.

The same argument applies to any z_j such that $n_{z_j}(\mathcal{P}) = 0$ for every periodic domain, \mathcal{P} . Then \mathcal{F}_j , the j th coordinate of $\bar{\mathcal{F}}$, which takes values in \mathbb{Z} , filters $\widehat{CF}(Y, \Gamma \setminus \{z_j\}; \mathfrak{s})$. Thus, for any collection, Z , of the z_j , which together will produce a lattice equivalent to $\bar{0}$, we obtain, using their filtration index coordinates, a filtration on $\widehat{CF}(Y, \Gamma \setminus Z; \mathfrak{s})$.

As usual define $\widehat{HF}(Y, \Gamma; \mathfrak{s})$ to be the homology of $\widehat{CF}(Y, \Gamma; \mathfrak{s})$ and $\widehat{HF}(Y, \Gamma; \mathfrak{s}, [v])$ to be the homology of $\widehat{CF}(Y, \Gamma; \mathfrak{s}, [v])$ for a specific choice of filtration index. We then have

$$\widehat{HF}(Y, \Gamma; \mathfrak{s}) \cong \bigoplus_{\bar{v} \in \mathbb{Z}^k / \Lambda} \widehat{HF}(Y, \Gamma; \mathfrak{s}, \bar{v}).$$

The action of $H_1(Y, \mathbb{Z})/\text{Tors}$ on the Heegaard–Floer homology, [11], extends to an action on $\widehat{HF}(Y, \Gamma; \mathfrak{s})$. Let $\gamma \subset \Sigma$ be a simple, closed curve representing the non-torsion class $h \in H_1$ and missing every intersection point between an α -curve and a β -curve. Let $a(\gamma, \phi)$ be the intersection number in \mathbb{T}_α of $\gamma \times \text{Sym}^{g-1}(\Sigma) \cap \mathbb{T}_\alpha$ and $u(1 + it)$ where u represents ϕ . This induces a map $\zeta \in Z^1(\Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta), \mathbb{Z})$. The action of such a co-cycle is defined by the formula:

$$A_\zeta([\mathbf{x}]) = \sum_{\mathbf{y} \in \mathcal{I}(\mathfrak{s})} \sum_{\phi \in D'(\mathbf{x}, \mathbf{y})} \zeta(\phi) \cdot (\#\widehat{\mathcal{M}}(\phi))[\mathbf{y}].$$

A_ζ preserves the summands $\widehat{CF}(Y, \Gamma; \mathfrak{s}, [v])$. In addition, if $\Lambda \equiv 0$ then we can replace $D'(\mathbf{x}, \mathbf{y})$ by $D(\mathbf{x}, \mathbf{y})$ in the map above see that usual action on $\widehat{CF}(Y, \mathfrak{s})$ is a filtered chain morphism whose top term is the action on $\widehat{CF}(Y, \Gamma; \mathfrak{s})$.

Following the argument presented above that $\partial^2 = 0$, we can verify that this is a chain map. We can then proved the analog of [11, Proposition 4.17].

Theorem 3.3 *There is a natural action of the exterior algebra, $\bigwedge^*(H_1(Y, \mathbb{Z})/\text{Tors})$ on the homology $\widehat{HF}(Y, \Gamma; \mathfrak{s})$, where $\zeta \in H_1(Y, \mathbb{Z})/\text{Tors}$ lowers degree by 1 and induces a filtered morphism of the chain complex when $\Lambda_Y \equiv 0$. Furthermore, this action respects the splitting into the subgroups $\widehat{HF}(Y, \Gamma; \mathfrak{s}, [v])$.*

Proof To see that this is a chain map note that the formula in [11, Lemma 4.17] for the coefficients of $\partial A_\zeta \pm A_\zeta \partial$ still applies as it only depends upon the ϕ and not upon the filtration indices. That the relevant homotopy classes have positive multiplicity allows us to use the argument above. As for the differential, any ϕ used in $\pi_2(\mathbf{x}, \mathbf{z})$ with $\mu(\phi) = 2$ will give the same set of indices for \mathbf{z} . The same observation applies to [11, Lemmas 4.18 and 4.19]. \square

In the next few sections, we will prove the following theorem, an extension of the main theorem in [13] which provided the statement for null-homologous knots.

Theorem 3.4 *For $\Gamma \subset Y$ coming from a d -based link, the homology $\widehat{HF}(Y, \Gamma; \mathfrak{s})$ is a relatively \mathbb{Z}^k / Λ indexed invariant of $\mathbb{L} \cup D$, considered up to isotopy, which is a direct sum $\bigoplus \widehat{HF}(Y, \Gamma; \mathfrak{s}, [v])$ of sub-groups that are individually invariant. Furthermore, the*

relative \mathbb{Z}^k/Λ -structure is invariant, although the precise assignment to cosets depends upon the choice of $\overline{\mathcal{F}}$. There is a natural action of $\bigwedge^*(H_1(Y, \mathbb{Z})/\text{Tors})$ on each of the summands.

Theorem 3.5 *Furthermore, suppose the link L can be divided into two disjoint sets of components, L_1 and L_2 , where L_1 consists of only of null-homologous components. As above, L_1 induces a $\mathbb{Z}^{|\mathbb{L}_1|}$ -filtration on the complex $\widehat{CF}(Y, D \cup \mathbb{L}_2; \mathfrak{s})$ which is invariant up to filtered chain homotopy. The H_1 -action will then act as $\mathbb{Z}^{|\mathbb{L}_1|}$ -filtered chain map on $\widehat{CF}(Y, D \cup \mathbb{L}_2; \mathfrak{s})$.*

Comment Although it seems plausible, when $\Lambda \neq 0$, to consider those ϕ with $n_{\overline{z}}(\phi) \in \Lambda$, this does not prescribe a differential. That $\phi = \phi_1 * \phi_2$ does not imply that $n_{\overline{z}}(\phi_j) \in \Lambda$. Hence the terms in ∂^2 may give rise to complementary boundaries one of which does not arise from the definition of the differential.

4 Chain maps and the proof of invariance

4.1 Some chain maps of multi-pointed complexes

4.1.1 Push-forward filtration indices Let $(\Sigma, \{\alpha_i\}_{i=1}^g, \{\beta_i\}_{i=1}^g, \{\gamma_i\}_{i=1}^g)$ be a Heegaard triple equipped with marked points w, z_1, \dots, z_k in the complement of the attaching curves. We assume that $M(\Sigma, \{\gamma_i\}_{i=1}^g, \{\alpha_i\}_{i=1}^g) \cong \#^l S^1 \times S^2$. Let Θ^+ be a closed generator, representing the generator of $\widehat{HF}(\Sigma_{\gamma\alpha}, \mathfrak{s}_0)$ as an H_1 -module. In analogy with the three dimensional case, let $\Lambda_{\alpha\beta\gamma}$ be the lattice in \mathbb{Z}^k of vectors

$$(n_{z_1}(\mathcal{T}), \dots, n_{z_k}(\mathcal{T}))$$

where \mathcal{T} is any triply periodic domain (including doubly periodic domains).

Given filtration indices on $Y_{\alpha\beta}$ and $Y_{\gamma\alpha}$, the latter assigning Θ^+ to the index 0, we can define a push-forward filtration index on $Y_{\gamma\beta}$ by

$$\overline{\mathcal{G}}(\mathbf{z}) = \overline{\mathcal{F}}_{\alpha\beta}(\mathbf{x}) + (n_{\overline{w}} - n_{\overline{z}})(\psi) \pmod{\Lambda_{\alpha\beta\gamma}}$$

where $\psi \in \pi_2(\mathbf{x}, \Theta^+, \mathbf{z})$. The expression on the right does not change under the addition of any triply periodic domains or the class $[S]$. Thus it does not depend upon the specific homotopy class of triangles joining the three intersection points. Nor does it change if we use a triangle abutting \mathbf{z} but with a different initial point. Let $\phi_1 \in \pi_2(\mathbf{x}', \mathbf{x})$, then

$$\overline{\mathcal{F}}_{\alpha\beta}(\mathbf{x}) - \overline{\mathcal{F}}_{\alpha\beta}(\mathbf{x}') = (n_{\overline{w}} - n_{\overline{z}})(\phi_1) \pmod{\Lambda_{\alpha\beta}}$$

when $\psi' = \psi * \phi_1$, we find

$$\overline{\mathcal{G}}(\mathbf{z}) = \overline{\mathcal{F}}_{\alpha\beta}(\mathbf{x}') + (n_{\overline{w}} - n_{\overline{z}})(\psi') \pmod{\Lambda_{\alpha\beta\gamma}}$$

since $\Lambda_{\alpha\beta} \subset \Lambda_{\alpha\beta\gamma}$.

On the other hand, $\overline{\mathcal{G}}$ is (almost) a filtration index for the Spin^c -structure on the $\gamma\beta$ boundary. Let $\phi \in \pi_2(\mathbf{z}, \mathbf{z}')$ and $\psi \in \pi_2(\mathbf{x}, \Theta^+, \mathbf{z})$, then for $\psi' = \phi * \psi$,

$$\overline{\mathcal{G}}(\mathbf{z}') - \overline{\mathcal{G}}(\mathbf{z}) = (n_{\overline{w}} - n_{\overline{z}})(\phi) \pmod{\Lambda_{\alpha\beta\gamma}}.$$

This fails to be a filtration index only because the relation holds modulo a larger lattice than we would like.

4.1.2 Chain maps We continue in the setting of the previous section. As in [11] and [15] we will now define chain maps for multi-pointed Heegaard diagrams. Let u be a Spin^c -structure on $X_{\alpha\beta\gamma}$, restricting to $\Sigma_{\gamma\alpha}$ as the unique Spin^c -structure with torsion first Chern class. We may define a multi-pointed chain map

$$F_u: \widehat{CF}(Y_{\alpha\beta}, \Gamma; u|_{Y_{\alpha\beta}}, \overline{\mathcal{F}}) \longrightarrow \widehat{CF}(Y_{\gamma\beta}, \Gamma; u|_{Y_{\gamma\beta}}, \overline{\mathcal{G}})$$

by

$$F_u([\mathbf{x}]) = \sum_{\mathbf{z}} \sum_{\psi} \#\mathcal{M}(\psi)[\mathbf{z}]$$

where $\mathbf{z} \in \mathcal{I}(u|_{Y_{\gamma\beta}})$, and ψ is a homotopy class representing u with $\mu(\psi) = 0$, $n_w(\psi) = 0$ and $n_{\overline{z}}(\psi) \equiv 0$. That this is a chain map follows from the usual arguments by examining ends of moduli spaces with $\mu(\psi') = 1$. Note that the map has image in the homology using the push-forward filtration, $\overline{\mathcal{G}}$. The latter is just taking the direct sum of the homology groups associated to cosets of $Z^k/\Lambda_{Y_{\gamma\beta}}$, which map to the same image when we additionally take the quotient under the image of $\Lambda_{Y_{\alpha\beta}}$.

If there is a $j \leq k$ such that $n_{z_j}(\mathcal{T}) = 0$ for all triply and doubly periodic domains we can adjust F_u to be

$$F_u([\mathbf{x}], \mathcal{F}_j(\mathbf{x})) = \sum_{\mathbf{z}} \sum_{\psi} \#\mathcal{M}(\psi)[\mathbf{z}, \mathcal{F}_j(\mathbf{x}) - n_{z_j}(\psi)]$$

and let $n_{z_j}(\psi) \geq 0$. This will then define a chain map for the chain groups where we allow terms in the differential with $n_{z_j} \geq 0$. This latter chain group is filtered by the value of \mathcal{F}_j , and the chain map above will be a filtered chain map relative to the j th coordinate of the push-forward filtration index.

Example ($\Lambda_{\alpha\beta\gamma} \equiv 0.$) For example, we have a string link in S^3 and the cobordism is generated by surgeries on curves which are algebraically split from S . In this case, the push-forward filtration satisfies

$$\overline{\mathcal{G}}(\mathbf{y}) = \overline{\mathcal{F}}(\mathbf{x}) + (n_{\overline{w}} - n_{\overline{z}})(\psi)$$

for every ψ representing a Spin^c -structure on the cobordism restricting in a specified way to the ends, and for any choice of a complete set of paths. The filtrations on the ends are \mathbb{Z}^k filtrations and the chain map \widehat{F} on the usual Heegaard–Floer homology, \widehat{HF} , is a filtered chain map. This situation occurs in the long exact skein sequence of [13]. In S^3 , a filtration index on each component can be calculated using the first Chern class of a Spin^c structure on the manifold obtained by performing 0 surgery on the knot. In [13], P Ozsváth and Z Szabó show that, in the case under consideration, the push forward of this filtration is the one determined by the first Chern class calculation on Y_1 and the formula above corresponds to their identity for c_1 .

4.2 Invariance

To prove Theorem 3.4, we need to show that performing any of the following moves will produce a chain homotopy equivalent complex.

- Handleslides and isotopies of $\{\beta_i\}_{i=1}^g$.
- Handleslides and isotopies of $\{\alpha_i\}_{i=k+1}^g$.
- Stabilization.
- Isotopies of $\{\alpha_i\}_{i=1}^k$ and handleslides of them over element of $\{\alpha_i\}_{i=k+1}^g$.

Furthermore, we are not allowed to isotope or slide over any portion of the disc D' . We can, however, arrange for a β curve to isotope past the entire disc. The resulting diagram can be achieved by allowable handleslides in $\{\beta_i\}_{i=1}^g$ because the disc is contractible, [11].

We develop the proof of invariance through the multi-pointed diagrams; it precisely mimics the proof for Heegaard–Floer homology, [11], and uses the same technical results.

4.2.1 Invariance of the choice of $\overline{\mathcal{F}}$ If we choose $\overline{\mathcal{F}}': \mathcal{I}(\mathfrak{s}) \rightarrow \mathbb{Z}^k/\Lambda$ to assign intersection points to cosets, then for any \mathbf{x} we have an element $v \in \mathbb{Z}^k/\Lambda$ such that $\overline{\mathcal{F}}'(\mathbf{x}) = \overline{\mathcal{F}}(\mathbf{x}) + v$. This permutes the summands in the decomposition of $\widehat{CF}(Y, \Gamma; \mathfrak{s})$ but does not alter the differential on each summand nor the differences $\overline{\mathcal{F}}(\mathbf{y}) - \overline{\mathcal{F}}(\mathbf{x})$. Thus the relative \mathbb{Z}^k/Λ -structure remains intact. Furthermore, if $n_{z_i}(\mathcal{P}) = 0$ for

all periodic domains, the vector v will shift the i th coordinates but not change their relative position. Thus upon letting L_i induce a filtration on the chain complex where we allow ϕ into the differential with $n_{z_i}(\phi) \geq 0$, we will have a filtrated chain map induced by the permutation, so the filtered chain type will also be invariant under change of the filtration index.

4.2.2 Results on admissibility Strong/Weak admissibility can be achieved for all our diagrams without disrupting the assumptions coming from Γ . We have seen the existence of such diagrams already. In [11, Section 5] P Ozsváth and Z Szabó show how to ensure that isotopies, handleslides and stabilizations can be realized through such diagrams. In each case this is achieved by finding a set $\{\gamma_i\}_{i=1}^g \subset \Sigma$ of disjoint, simple closed curves with the property that $\#(\beta_i \cap \gamma_j) = \delta_{i,j}$ and that $\mathbb{T}_\alpha \cap \mathbb{T}_\gamma \neq \emptyset$ (or the same but with the roles of α and β switched).

We convert, through stabilization, our multi-pointed diagram into a diagram with an embedded disc. We only need require that w and the entire disc do not intersect the winding region. However, we may always choose our γ to lie in the disc’s complement. If we wish to wind the α we can do the same, requiring only that each γ_i that intersects a meridian does so away from the segment in $\partial D'$. With this arrangement of $\{\gamma_i\}_{i=1}^g$ the proofs of Lemmas 5.4, 5.6, and 5.7 of [11] carry through.

4.2.3 Invariance of complex structure and isotopy invariance In these cases, we rewrite the chain maps defined in [11] to incorporate the new indices. For example, P Ozsváth and Z Szabó define a chain map for a homotopy of paths of almost complex structures which we adjust to be, cf [11]:

$$\Phi_{J_{s,t}}[\mathbf{x}] = \sum_{\mathbf{y}} \sum_{\phi} \# \mathcal{M}_{J_{s,t}}(\phi)[\mathbf{y}]$$

where the sum is over all ϕ with $\mu(\phi) = 0$, $n_w(\phi) = 0$, and $n_{\bar{z}}(\phi) = 0$. The moduli space consists of holomorphic representatives of ϕ satisfying a Cauchy–Riemann like equation depending on $J_{s,t}$.

That the filtration index relation holds for all homotopy classes of discs, ϕ , ensures this is still a chain map and that it preserves the splitting of $\widehat{CF}(Y, \Gamma; \mathfrak{s})$ into a direct sum of complexes, given a filtration index. Furthermore, when $n_{z_i}(\mathcal{P}) = 0$ for all periodic domains, the above map may be adjusted by allowing all homotopy classes where $n_{z_i}(\phi) \geq 0$. This new map is a filtered chain morphism, by the positivity of n_{z_i} , for the filtration provided by the i th coordinate of the filtration index.

We adjust the chain homotopy in [11] by considering homotopy classes such that $n_w(\phi) = 0$, $n_{\bar{z}}(\phi) = \bar{0}$. It is then a consequence of Gromov compactness that $\Phi_{J_{s,t}}$

has an inverse chain map up to homotopy, and thus the map above is an isomorphism on homology. Furthermore, if $n_{z_i}(\mathcal{P}) = 0$ for all periodic domains, we can extend the homotopy inverse in the same way we extended Φ above to obtain a filtered chain map which is an inverse up to filtered chain homotopy. The invariance of the action of $H_1(Y, \mathbb{Z})/\text{Tors}$ follows as in [11], adjusting the maps similarly in regard to z_1, \dots, z_k .

For the isotopy invariance, the same argument applies (as the proofs are roughly parallel). We write the chain map coming from the introduction/removal of a pair of intersection points as:

$$\Gamma^\Psi[\mathbf{x}] = \sum_y \sum_{\phi \in \pi_2^\Psi(\mathbf{x}, y)} \#\mathcal{M}^\Psi(\phi)[y]$$

where we count holomorphic representatives with moving boundary, [11].

When we have a pair creation, we get new pairs of intersection points \mathbf{q}_+ and \mathbf{q}_- along with a holomorphic disc for a class in $\pi_2(\mathbf{q}_+, \mathbf{q}_-)$. The homotopy classes of discs joining intersection points from the original diagram do not change in this process, so their relative indices do not change. We can compute the index for \mathbf{q}_+ from any disc, and then note that \mathbf{q}_- receives the same value in \mathbb{Z}^k/Λ , using the newly introduced holomorphic disc. Otherwise, the argument is the same as for invariance under alteration of the path of almost complex structures. If z_i introduces a filtration, as above, then we can allow Γ^Ψ to include discs with $n_{z_i} \geq 0$ to obtain a filtered chain map with up to homotopy filtered inverse.

4.2.4 Invariance under handleslides The standard proof for handleslide invariance applies in this context, using chain maps induced by a Heegaard triple as modified below. We show below that the element corresponding to Θ^+ is still closed and explain which additive assignment for homotopy classes of triangles will work.

We will describe this solely for the α . Away from the curves involved in the handleslide, the resulting boundary $\Sigma_{\gamma\alpha}$ is the connected sum of genus 1 diagrams for $S^1 \times S^2$. Because we have not moved a curve across a marked point, the corresponding multi-pointed diagram has all the basepoints, w and z_i , in the same domain D of $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\gamma_i\}_{i=1}^g)$.

If we calculate the multi-pointed filtration indices for $\Sigma_{\gamma\alpha}$ we see that 1) $\Lambda \equiv 0$ and 2) all 2^g representatives of the torsion Spin^c structure have filtration index 0. None of the holomorphic discs have domain containing any marked points, so the homology with marked points is isomorphic the standard homology for the connected sum of the $S^1 \times S^2$. As usual, we will use the canonical generator Θ^+ , the maximally graded generator for $\widehat{HF}(\#^g S^1 \times S^2; \mathfrak{s}_0)$. We will use strongly admissible diagrams for the

Spin^c structures on the ends of the cobordism. The cobordism, $X_{\alpha\beta\gamma}$ is then $Y \times I$, after filling in the $Y_{\gamma\alpha}$ –boundary, so each triangle will represent $\mathfrak{s} \times I$.

As in the proof of handleslide invariance in Heegaard–Floer homology, [11], our diagrams may be drawn so that each intersection point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ can be joined to an intersection point $\mathbf{x}' \in T_\gamma \cap T_\beta$ by a unique holomorphic triangle $\psi_{\mathbf{x}} \in \pi_2(\mathbf{x}, \Theta^+, \mathbf{x}')$ with domain contained in the sum of the periodic domains from $\Sigma_{\gamma\alpha}$. Because the handleslide does not cross a meridian we can ensure that for these special homotopy classes of triangles $n_w(\psi_{\mathbf{x}}) = n_{z_i}(\psi_{\mathbf{x}}) = 0$. If we identify $\pi_2(\mathbf{x}, \mathbf{x})$ with $\pi_2(\mathbf{x}', \mathbf{x}')$ using $\psi_{\mathbf{x}}$ then we implicitly have an identification $H_2(Y_0) \cong H_2(Y \times I) \cong H_2(Y_1)$, an identification $\Lambda_{Y_0} = \Lambda_{Y \times I} = \Lambda_{Y_1}$, and a push-forward index

$$\bar{\mathcal{G}}(\mathbf{x}') = \bar{\mathcal{F}}(\mathbf{x}) + (n_{\bar{w}} - n_{\bar{z}})(\psi_{\mathbf{x}}) = \bar{\mathcal{F}}(\mathbf{x}) \pmod{\Lambda_Y}.$$

We can now repeat the argument from [11] using the chain maps described in the previous section, for the filtration index and its push-forward. Furthermore, associativity, up to chain homotopy, can be established using moduli spaces holomorphic quadrilaterals, and the push-forward filtration indices are compatible with the associativity relation, since $n_w = n_{z_j} = 0$ for all relevant homotopy classes of quadrilaterals. Lastly, in the associativity relation, there will be can unique holomorphic representatives of homotopy classes of triangles for $X_{\beta\gamma\delta}$ with $n_w = n_{z_j}$ and joining the maximal generators for the torsion Spin^c–structures in these boundaries. This follow again from the fact that we do not allow handleslides with marked points in their supports. This is enough to complete the argument in [11].

When $n_{z_j}(\mathcal{P}) = 0$ for every doubly periodic domain, \mathcal{P} , we can extend the chain map to a filtered chain map for the filtration induced by z_j on the chain group allowing $n_{z_j}(\phi) \geq 0$. The associativity argument also extends to this case, allowing us to conclude that the filtration is invariant under the handleslide.

4.2.5 Stabilization invariance Stabilization changes the surface Σ to $\Sigma' = \Sigma \# \mathbb{T}^2$ and adds an additional α and an additional β curve, intersecting in a single point c . Stabilization does not alter H_2 nor does it affect the structures of $\pi_2(\mathbf{x}, \mathbf{y})$. We make the necessary alterations for the gluing result, [11, Theorem 10.4], to hold. This theorem provides the invariance under stabilization in the standard case. In addition, in [11], P Ozsváth and Z Szabó show that the action of $H_1(Y, \mathbb{Z})/\text{Tors}$ is invariant under stabilization. That this also applies to multi-pointed diagrams follows analogously to the case of the differential.

5 Additional chain complexes

We can use the multi-pointed data to more directly mimic the development of knot Floer homology.

Definition 5.1 For $\mathfrak{s} \in \text{Spin}^c(Y)$, Let $CF_{\Gamma}^{\infty}(Y; \mathfrak{s})$ be the $\mathbb{Z}[U]$ -module

$$\text{Span}_{\mathbb{Z}}\{[\mathbf{x}, i, \bar{v}] | \mathbf{x} \in \mathcal{I}(\mathfrak{s}), i \in \mathbb{Z}, \bar{v} \in \mathbb{Z}^k\}$$

with U -action

$$U([\mathbf{x}, i, \bar{v}]) = [\mathbf{x}, i - 1, v_1 - 1, \dots, v_k - 1].$$

Let $CF_{\Gamma}^{-}(Y; \mathfrak{s})$ be the submodule where $i < 0$ and let $CF_{\Gamma}^{+}(Y; \mathfrak{s})$ be

$$CF_{\Gamma}^{\infty}(Y; \mathfrak{s}) / CF_{\Gamma}^{-}(Y; \mathfrak{s})$$

(with $i \geq 0$).

We pick an additive assignment.

Definition 5.2 An additive assignment for \mathfrak{s} is a map

$$\mathcal{A}: \pi_2(\mathbf{x}, \mathbf{y}) \rightarrow H_2(Y; \mathbb{Z})$$

for all $\mathbf{x}, \mathbf{y} \in \mathcal{I}(\mathfrak{s})$, such that

- (1) $\mathcal{A}(\phi_1 * \phi_2) = \mathcal{A}(\phi_1) + \mathcal{A}(\phi_2)$,
- (2) $\mathcal{A}(\Sigma) = 0$.

We will also pick a basepoint $\mathbf{x}_0 \in \mathcal{I}(\mathfrak{s})$ and interpret $n_{\bar{z}}(\mathcal{A}(\phi))$ as $n_{\bar{z}}(\mathcal{D})$ where \mathcal{D} is the periodic domain corresponding to $\mathcal{A}(\phi)$ in $\pi_2(\mathbf{x}_0, \mathbf{x}_0)$.

Given an additive assignment and a basepoint, we can redefine filtration indices to have image in \mathbb{Z}^k .

Definition 5.3 A filtration index, \bar{F} , for \mathfrak{s} , is a map $\bar{F}: \mathcal{I}(\mathfrak{s}) \rightarrow \mathbb{Z}^k$ with the following properties.

- (1) Let $n_{\bar{z}}(\phi) = (n_{z_1}(\phi), \dots, n_{z_k}(\phi))$ and $n_{\bar{w}}(\phi) = (n_w(\phi), \dots, n_w(\phi))$

$$\bar{F}(\mathbf{y}) - \bar{F}(\mathbf{x}) = (n_{\bar{w}} - n_{\bar{z}})(\phi) + n_{\bar{z}}(\mathcal{A}(\phi)).$$

- (2) If z_i and z_j are in the same component of $\Sigma \setminus (\{\alpha_i\}_{i=1}^g \cup \{\beta_i\}_{i=1}^g)$ then the i th and j th coordinate maps, \mathcal{F}_i and \mathcal{F}_j , are equal.
- (3) If z_i is in the same component as w then $\mathcal{F}_i \equiv 0$, a constant map.

$$(4) \quad \overline{\mathcal{F}}(\mathbf{x}_0) \equiv 0.$$

For any \mathbf{x} there is a unique $\phi_{\mathbf{x}} \in \pi_2(\mathbf{x}_0, \mathbf{x})$ such that $\mathcal{A}(\phi) = 0$ and $n_w(\phi) = 0$. For a periodic domain $\mathcal{P} \in \pi_2(\mathbf{x}, \mathbf{x})$, we let $\mathcal{P}' = \phi_{\mathbf{x}}^{-1} * \mathcal{P} * \phi_{\mathbf{x}}$. Then $(\mathcal{A})(\mathcal{P}') = \mathcal{A}(\mathcal{P})$. When we add a periodic domain \mathcal{P} to ϕ , the right hand side of 1) will then change by $-n_{\overline{z}}(\mathcal{P}) + n_{\overline{z}}(\mathcal{P}) = 0$. Thus the filtration index is well-defined. Choosing a different \mathcal{A} alters the value assigned to each ϕ by an element of Λ . Changing \mathbf{x}_0 changes $\overline{\mathcal{F}}$ to $\overline{\mathcal{F}} + v$. Thus, modulo Λ , the filtration index relation defines a relative \mathbb{Z}^k/Λ -grading.

Given a choice of $\mathbf{x}_0 \in \mathcal{I}(\mathfrak{s})$ and \mathcal{A} . There is a differential, ∂^∞ on $CF_\Gamma^\infty(Y; \mathfrak{s})$, defined by the linear extension of

$$\partial[\mathbf{x}, i, \overline{v}] = \sum_{\mathbf{y} \in \mathcal{I}(\mathfrak{s})} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} \# \widehat{\mathcal{M}}(\phi)[\mathbf{y}, i - n_w(\phi), \overline{v} - n_{\overline{z}}(\phi) + n_{\overline{z}}(\mathcal{A}(\phi))].$$

The differential is a $\mathbb{Z}[U]$ -module map which also specifies a differential on CF_Γ^\pm .

Definition 5.4 Let $CF^\infty(Y, \Gamma; \mathfrak{s}, \mathcal{A}, \mathbf{x}_0)$ be the subgroup of $CF_\Gamma^\infty(Y; \mathfrak{s})$ generated by those $[\mathbf{x}, i, \overline{v}]$ with

$$(v_1, v_2, \dots, v_k) - (i, i, \dots, i) = \mathcal{F}(\mathbf{x}).$$

It is then possible to find a more convenient complex.

Proposition 5.5 $CF^\infty(Y, \Gamma; \mathfrak{s}, \mathcal{A}, \mathbf{x}_0)$ is a subcomplex of $CF_\Gamma^\infty(Y; \mathfrak{s})$ as a $\mathbb{Z}[U]$ -modules.

Sketch of proof We must check three requirements.

- (1) That $CF^\infty(Y, \Gamma; \mathfrak{s}, \mathcal{A}, \mathbf{x}_0)$ is preserved by $\mathbb{Z}[U]$. This is a byproduct of the filtration index relation.
- (2) If $[\mathbf{x}, i, v] \in CF^\infty(Y, \Gamma; \mathfrak{s}, \mathcal{A}, \mathbf{x}_0)$ and $\langle \partial^\infty[\mathbf{x}, i, v], [\mathbf{y}, j, w] \rangle \neq 0$ then $[\mathbf{y}, j, w] \in CF^\infty(Y, \Gamma; \mathfrak{s}, \mathcal{A}, \mathbf{x}_0)$. This follows from the additivity of \mathcal{A} and the filtration index relation applied to calculating $\overline{\mathcal{F}}(\mathbf{x})$ using $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $i - j = n_w(\phi)$ and $v - w = n_{\overline{z}}(\phi)$.
- (3) That $(\partial^\infty)^2 = 0$. It could be that a $\mu = 2$ class ψ has holomorphic moduli space with boundaries $\phi_1 * \phi_2$ and $\phi'_1 * \phi'_2$, but such that the v -index changes differently in the calculation for the two different boundaries. However, since \mathcal{A} , n_w , and $n_{\overline{z}}$ are all additive, this does not happen.

Otherwise, we follow the proof in [11]. □

The action of $H_1(Y, \mathbb{Z})/\text{Tors}$ on the Heegaard–Floer homology, [11], also extends to an action on the homology of the $CF^\infty(Y, \Gamma; \mathfrak{s}, \mathcal{A}, \mathbf{x}_0)$ by

$$A_\zeta([\mathbf{x}, i, \bar{v}]) = \sum_y \sum_{\{\phi: \mu(\phi)=1\}} \zeta(\phi) \cdot (\#\widehat{\mathcal{M}}(\phi))[y, i - n_w(\phi), \bar{v} - n_z(\phi - \mathcal{A}(\phi))].$$

Combining the techniques of [11] for the H_1 -action with the previous sections one can then prove the following theorem.

Theorem 5.6 *For Γ coming from a d -based link. $HF^\pm(Y, \Gamma; \mathfrak{s})$ and $HF^\infty(Y, \Gamma; \mathfrak{s})$ each have a relative \mathbb{Z}^k/Λ grading which is an invariant of the isotopy class of the d -based link. U acts non-trivially on this additional grading. Furthermore, there is a natural action of the exterior algebra, $\bigwedge^*(H_1(Y, \mathbb{Z})/\text{Tors})$ on the homology $HF^\infty(Y, \Gamma; \mathfrak{s})$, where $\zeta \in H_1(Y, \mathbb{Z})/\text{Tors}$ lowers degree by 1.*

The details are omitted, except to say that one adjusts all the chain maps, relative to an additive assignment on a Heegaard triple, used in the usual proof of invariance in the same manner as the differential was adjusted above, relative to an additive assignment on the Heegaard diagram.

6 Basic properties of $HF(Y, S; \mathfrak{s})$

Note For the remainder of the paper, we will think in terms of a string link, S , in $D^2 \times I$, possibly with some framed link specifying a three-manifold other than S^3 , but without reference to a d -base.

6.1 An example

In Figure 2 we examine the homology of a knot in $S^1 \times S^2$ which intersects $[S^2]$ precisely once. Regardless of the knot, K , we may find a diagram as in the figure. The intersection points, after handlesliding, giving generators in the chain complex are precisely $\Theta^\pm \times \mathbf{x}$, where \mathbf{x} is a generator for the Knot Floer homology of K in S^3 . However, in the complex, only one of the two homotopy classes from Θ^+ to Θ^- does not cross the marked point, z . It is straightforward to verify that $\widehat{\partial}\Theta^+ \times \mathbf{x} = \Theta^- \times \mathbf{x} + \Theta^+ \times \widehat{\partial}_K \mathbf{x}$. In each Spin^c structure, the filtration index collapses; however, the differential as above also produces trivial homology.

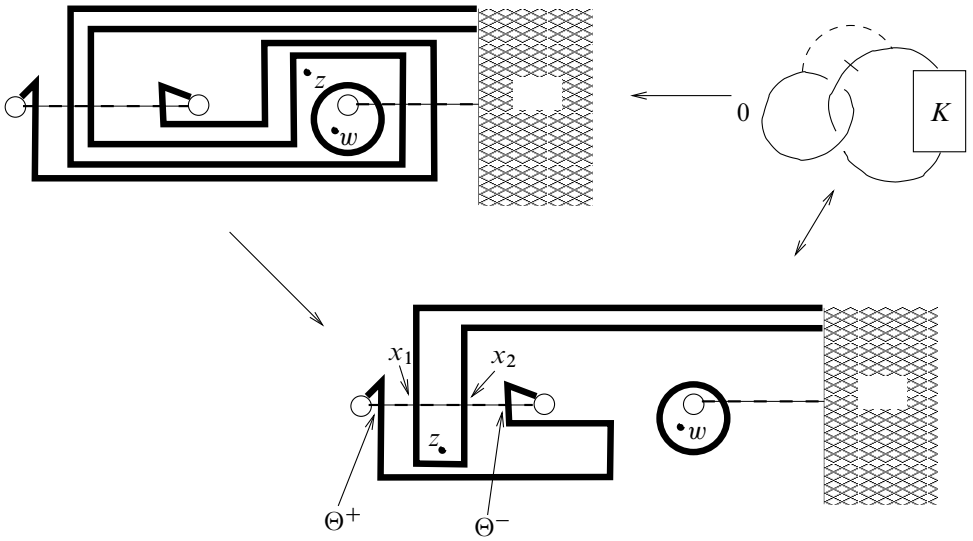


Figure 2: 0–surgery on an unknot linked with K . The d –base follows the dashed arc without twisting. The hatched part of the diagram includes the effect of K on the α (heavy) and β (dashed) curves. The portion shown in detail describes the linking and the surgery. Note that placing a meridian instead of the framing curve gives a diagram for K . The intersection points never use x_1 or x_2 , hence come from intersection points for the knot diagram. Destabilizing the meridian for K produces a diagram where the only holomorphic discs are those from the complex for K and for $S^1 \times S^2$.

6.2 Subtracting a strand

Suppose we remove a component strand of $S \subset (Y - B^3)$ to obtain a new string link, S' . We may use any of the complexes defined previously from a Heegaard diagram for S to compute find the complex for S' by ignoring z_k . The diagram without z_k is a Heegaard diagram subordinate to S' , and the differential incorporates the same holomorphic discs. When L_k does not algebraically intersect any homology class in Y , we can view the last coordinate, j_k , as a filtration on the complex $\widehat{CF}(Y, S; \mathfrak{s})$ and use the associated spectral sequence with E^1 term $\bigoplus_r \widehat{HF}(Y, S; \mathfrak{s}, [j_1, \dots, j_{k-1}], r)$ to calculate $\widehat{HF}(Y, S'; \mathfrak{s}, [j_1, \dots, j_{k-1}])$. It collapses in finitely many steps.

Adding or subtracting an unknotted, unlinked, null-homologous component corresponds to adding or subtracting an index which behaves like i . The Heegaard diagram corresponds to stabilizing in the region containing w and placing a new point z_{k+1} at

the end of an arc which crosses the new α once geometrically and misses the new β . As z_{k+1} will still be in the same domain as w , so there is no change in homology.

6.3 Mirror string links

Let $S \subset Y - B^3$ be a string link in standard form, lying in the plane which defined the projection of our framed link diagram except in neighborhoods of the crossings. Let \mathfrak{s} denote a Spin^c structure on this manifold. Let S' be the string link found through reflection in this plane, reflecting the framed components as well and switching the sign of their framing. Then S' is the string link induced by S in $-Y$ under orientation reversal. Drawing a Heegaard diagram from a projection of (Y, S) , we may use the same β -curves, and change the α -curves for each crossing and framing to obtain a diagram for a projection of $(-Y, S')$. Keeping the orientation of the mirror plane fixed, we can compare portion of the three-manifolds on either side of the mirror plane. The meridians intersect both sides; however, the marked points for S lie on one side of the plane whereas their counterparts for S' , z'_i , lie on the other side of the mirror plane, since they too must be reflected. The intersection points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ from the original diagram are in bijection with those of the new diagram. Each homotopy class ϕ is carried to a new homotopy class ϕ' , but to join the same intersection points it must map in with reversed multiplicities. All this implies that we may calculate $\widehat{HF}_*(-Y, S'; \mathfrak{s}')$ by looking at the intersection points for (Y, S) and the differential for the complex using $-\Sigma$. As in Heegaard–Floer homology, this new complex calculates the co-homology $\widehat{CF}^*(Y, S; \mathfrak{s})$; there is thus an isomorphism $\widehat{HF}_*(Y, S; \mathfrak{s}) \rightarrow \widehat{HF}^*(-Y, S'; \mathfrak{s}')$. (This isomorphism maps absolute degrees as $d \rightarrow -d$ if \mathfrak{s} is torsion). Using the same marked points, but the image of the basepoint and paths in the complete set of paths we find that $-\overline{\mathcal{F}}$ will be a filtration index for S' when $\overline{\mathcal{F}}$ is for S ($\Lambda_{(-Y, S')} = -\Lambda_{(Y, S)} = \Lambda_{(Y, S)}$). In particular, since each intersection point has fixed image on each meridian, the boundary of a class ϕ must contain whole multiples of the meridian and so $n_{z_i}(\phi) = -n_{z'_i}(\phi')$ since the multiplicities reversed, but $n_{z'_i}(\phi') = n_{z_i}(\phi)$. In summary, there is an isomorphism (including the absolute grading when present):

$$\widehat{HF}_*^{(-d)}(Y, S; \mathfrak{s}, [\overline{j}]) \rightarrow \widehat{HF}_{(d)}^*(-Y, S'; \mathfrak{s}', [-\overline{j}]).$$

When $Y = S^3$ and S is a normal string link, then S' is the mirror image of S found by switching all the crossings. The change in indices corresponds to the alteration $t_i \rightarrow t_i^{-1}$ in the Alexander polynomial (see Section 7).

7 Alexander invariants for string links in $D^2 \times I$

Let S be a string link in $D^2 \times I$. In this section we relate the Euler characteristic of $\widehat{CF}(S; \bar{v}) \otimes \mathbb{Q}$ to classical Alexander invariants built from coverings of $D^2 \times I - S$.

7.1 Alexander invariants for string links

Let $S \subset D^2 \times I$ be a string link; let $X = D^2 \times I - \text{int } N(S)$ be the complement of its tubular neighborhood. Let $E_0 = X \cap (D^2 \times \{0\})$ and $E_1 = \partial X - (X \cap \text{int } D^2 \times \{0\})$. Both are planar Riemann surfaces. By the Meyer–Vietoris sequence we have that $H_1(X; \mathbb{Z}) \cong \mathbb{Z}^k$, generated by the meridians of the strands in E_0 . Denote by \mathfrak{H} the ring $\mathbb{Z}[t_1^{\pm 1}, \dots, t_k^{\pm 1}]$ and let $\pi: \pi_1(X, x_0) \rightarrow H_1(X, \mathbb{Z})$ be the Hurewicz map for some basepoint, x_0 , made more specific below.

Definition 7.1 The torsion of a string link in $D^2 \times I$, denoted $\tau(S)$, is the Reidemeister torsion of the \mathfrak{H} -module $H_1(\tilde{X}, \tilde{E}_1; \mathbb{Z})$, where \tilde{X} is the cover of X determined by the Hurewicz map and \tilde{E}_1 is the pre-image of E_1 under the covering map. Note that this is defined only up to multiplication by a unit of \mathfrak{H} .

This Abelian invariant first appeared in a paper of R Litherland, [9], on the Alexander invariants of θ_{k+1} -graphs. String links and d -based links can be seen through the lens of θ_{k+1} -graphs, and the definition above is then equivalent to that in [9]. P Kirk, C Livingston and Z Wang, [7], were the first to use torsion properties in studying string links. Their definition is slightly different, using the fundamental group, but they relate it to the Reidemeister torsion of the based, acyclic co-chain complex $C^*(X, E_0; F)$ where $F = \mathbb{Q}(h_1, \dots, h_k)$ and the coefficients are twisted by the Hurewicz map. This is essentially the same as our definition.

To compute $\tau(S)$ we will construct a relative cell decomposition for (X, E_1) . We start with a relative cell complex for $(E_0, E_0 \cap E_1)$. Think of E_0 as the unit disc D^2 minus non-overlapping, equal radii open discs centered at $(i/k, 0)$ for $i = 1 - k, 3 - k, \dots, k - 3, k - 1$. We use k one-cells, the first joining $(1, 0)$ to the boundary of the disc centered at $(k - 1)/k$ and lying on the x -axis, and each successive one cell joining the boundaries of the discs centered at j/k and $(j - 2)/k$ and lying on the x -axis. The complement of the ∂E_0 union these one cells is an open 2-cell. This relative cell decomposition may be extended to the entirety of X by attaching one-cells at each of the crossings in a projection of S , along the axis of the projection, and then adding two-cells for each face in the projection with the exception of the leftmost one, which we have called U . By the definition of a cell complex, there are two-cells, F_1, \dots, F_k , arising from faces, that intersect E_0 in each of the one-cells in

the decomposition of E_0 . The complement in X of ∂X union the one and two cells is the interior of a single three-cell, which completes the decomposition.

We may collapse the two-cell in E_0 into the union of E_1 and the other two-cells by contracting the three-cell. Likewise we may collapse the one-cells in E_0 into the union of E_1 and the other one-cells by contracting F_1, \dots, F_k respectively. This leaves a relative cell complex with an equal number of 1- and 2-cells. We call this cell complex Y . The homotopy and homology properties of the pair (X, E_1) are determined by this complex. In particular, the chain complex for \tilde{Y} , the lift of Y under the covering map corresponding to the Hurewicz map, as a relative cell complex has a boundary map

$$0 \longrightarrow C_2(\tilde{Y}) \xrightarrow{\tilde{\partial}} C_1(\tilde{Y}) \longrightarrow 0.$$

Thus, $H_1(\tilde{X}, \tilde{E}_1) \cong \text{coker } \tilde{\partial}$, and $\tilde{\partial}$ defines a square presentation matrix, P , for $H_1(\tilde{X}, \tilde{E}_1)$. Let $\epsilon: \mathfrak{H} \rightarrow \mathbb{Z}$ be the homomorphism defined by substituting 1 for each variable, and $\epsilon(P)$ be the matrix found from P by applying ϵ to each of its entries. The $\epsilon(P)$ is the boundary map for the relative chain complex for Y . Since $H_1(X, E_1; \mathbb{Z}) \cong 0$, as E_1 also contains all the meridians, $\epsilon(P)$ is invertible and P has non-zero determinant, $\tau(S)$.

It is easier to calculate P using Fox calculus applied to the fundamental group of the complement of the string link, see Kauffman [6] or Burde–Zieschang [1]. We present the fundamental group of $X = D^2 \times I - S$ by choosing a basepoint in $D^2 \times I - S$ and loops through faces of the projection as generators. The crossings then provide relations in the usual manner, and we obtain a presentation of the fundamental group, similar to the Dehn presentation for a knot group, with k more generators than relations. If we denote the generators by s_i and the relations by R_j then the matrix $\left(\frac{\partial R_j}{\partial s_i}\right)$ found by first applying free differentials $\frac{\partial}{\partial s_i}$ to R_j , then taking the image of each element in $H_1(X; \mathbb{Z})$ is presentation matrix for $H_1(\tilde{X}, \tilde{x}_0)$ as an \mathfrak{H} -module. This presentation matrix has k more columns than rows, corresponding to the faces F_1, \dots, F_k . Ignoring these columns corresponds to collapsing the faces, and the resulting square matrix gives a presentation for $H_1(\tilde{Y}, \tilde{E}_1)$.

7.2 Heegaard splittings and Fox calculus

Proposition 7.2 *Let $\tau(S)$ be the torsion of $H_1(\tilde{X}, \tilde{E}_1)$ as in the previous subsection. Then $\tau(S)$ equals, up to multiplication by a unit,*

$$\sum_{\bar{v} \in \mathbb{Z}^k} \chi(\widehat{HF}(S, \bar{v}; \mathbb{Q})) t_1^{v_1} \cdots t_k^{v_k}$$

where $\widehat{HF}(S, \bar{v}; \mathbb{Q})$ is the homology of $\widehat{CF}(S^3, S; \bar{v}) \otimes \mathbb{Q}$ and the Euler characteristic is taken with respect to the canonical $\mathbb{Z}/2\mathbb{Z}$ -grading, [11].

We will prove this using Fox calculus. Various other authors have used much the same argument in different settings; J Rasmussen provides a very similar argument for Heegaard diagrams for three-manifolds in [17].

Proof Let S be a string link in $D^2 \times I$. We consider the standard Heegaard decomposition induced from a projection described in section 1. Let H_α be the handlebody determined by $\{\alpha_i\}_{i=1}^g$, and H_β be the handlebody determined by $\{\beta_i\}_{i=1}^g$. We assume that our meridians lie in $D^2 \times \{0\}$. In such a diagram, each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ has only one choice for its representative point on each meridian (an intersection with the attaching curve for one of the faces F_1, \dots, F_k). This corresponds to collapsing these faces in the relative cell complex.

We take as our basepoint, p_0 , for $\pi_1(X)$ the 0-cell in H_β . For each of the faces, we choose a path f_i , the gradient flow line oriented from the basepoint to the critical point corresponding to β_i which links the core positively in S^3 . The other gradient line oriented from the index 1 critical point to the 0-cell will be called \bar{f}_i . The loops $b_i = \bar{f}_i \circ f_i$ generate $\pi_1(X, p_0)$.

The α , not including the meridians, induce the relations for a presentation of π_1 corresponding to the Dehn presentation of the fundamental group, [6]. We choose an intersection point $\mathbf{u} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ which corresponds to points $u_i \in \alpha_{\sigma(i)} \cap \beta_i$ in Σ for some permutation $\sigma \in S_g$. Note that the choice along the meridians is prescribed for each such intersection point: there are k meridians and $k + 1$ faces intersecting them, but we cast one aside. This arrangement implies that our only choices occur on the non-meridional α . For a non-meridional α , let $[\alpha_{\sigma(i)}]$ be the path from the basepoint, along f_i , through the attaching disc for β_i to u_i , and around $\alpha_{\sigma(i)}$ with the its orientation, and then back the same way to the basepoint. Each time $[\alpha_j]$ crosses β_i positively, we append a b_i to the relation; each time it intersects negatively we append a b_i^{-1} . The word so obtained is called a_i . We derive this principle by looking at the segments α_j^s into which the β cut α_i and flowing them forwards along the gradient flow. The interior of each segment flows to the basepoint, while the endpoints flow to critical points in the attaching discs for the β . Thus, the path from one endpoint of the segment, to the critical point corresponding to that β_s , then along some f_s^{-1} or \bar{f}_s , and back along one of f_t or \bar{f}_t^{-1} , then to the other end of the segment, and back along the segment, is null homotopic. This allows us to break the α up into the various β it crosses.

Then $\left(\frac{\partial a_j}{\partial b_i}\right)$, ignoring the columns corresponding to the faces abutting $D^2 \times \{0\}$, is a presentation matrix for the \mathfrak{H} -module $H_1(\tilde{X}, \tilde{E}_1)$, and hence its determinant will provide the Alexander invariant.

If we consider the free derivative of a_j with respect to a b_i we find terms which correspond to each intersection point of α_j with β_i . The term possesses a minus sign when the two intersect negatively, otherwise it possesses a positive sign. The terms correspond to paths from the basepoint through $f_{\sigma^{-1}(j)}$ to $\beta_{\sigma^{-1}(j)}$, through the attaching disc to $u_{\sigma^{-1}(j)}$, along α_j to the intersection point with β_i and then back along f_i^{-1} . This can be rewritten as a word in the b_i . Summing over all intersection points with β_i equals $\partial_{b_i}(a_j)$.

Let μ be the Hurewicz map from the fundamental group to the first homology group. According to the Fox calculus, the matrix $[\mu(\partial_{b_i} a_j)]$, is a presentation matrix for the homology of the cover as an \mathfrak{H} -module. Again, we ignore the b_j corresponding to the faces abutting $D^2 \times \{0\}$. We calculate the Alexander invariant by computing the determinant of this matrix. Each term in this determinant has the form $\text{sgn}(\sigma)(-1)^{\#} h_1^{\rho_1} \cdots h_k^{\rho_k}$, where ρ_i is the sum of the powers of h_i over the terms in the determinant multiplying to this monomial; we do not allow any cancellation of terms. This monomial corresponds to a specific intersection point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ found from the pairing of rows and columns in the matrix. Likewise $\#$ is the number of negative intersections $\alpha_{\sigma(i)} \cap \beta_i$ in the g -tuple corresponding to this term.

Let \mathbf{x} and \mathbf{y} be two intersection points. We will consider the differences

$$\#_{\mathbf{y}} - \#_{\mathbf{x}} \quad \rho_i(\mathbf{y}) - \rho_i(\mathbf{x}).$$

Since we are considering points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ for a diagram of S^3 , there is a homotopy class of discs $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$.

We place marked points in the diagram corresponding to the strands and according to our conventions. We may measure how many times a 2-chain in X , representing a homotopy class of discs, ϕ , intersects the link components by evaluating $(n_w - n_{z_i})(\phi)$. We wish to show that

$$\rho_i(\mathbf{y}) - \rho_i(\mathbf{x}) = (n_w - n_{z_i})(\phi).$$

The right hand side counts the number of times that the boundary of $\mathcal{D}(\phi)$ winds around the i th meridian. We need only show that the same is true of the left, or, equivalently, that μ_i , the i th coordinate of the boundary, equals the left hand side.

In the boundary of the disc we have the α oriented from the points in \mathbf{x} to those in \mathbf{y} . We can take segments starting at u_i and traveling along $\alpha_{\sigma(i)}$ to x_i and y_i so that their difference is the oriented segment of the boundary of ϕ in α_i . We join this to the

basepoint by using paths in the attaching discs for the β and the preferred paths f_j or f_j^{-1} at each endpoint. Breaking this up as before, we can convert this path into a word of b_i and their inverses. If we look at one of the β_i boundary segments in ϕ , we see that the concatenation of the words for the α segments corresponding to the intersection points with β_i homotopes into the α boundary and the β boundary of ϕ .

Thus μ_i of the concatenation equals μ_i of the boundary of $\mathcal{D}(\phi)$. Furthermore, μ_i applied to each word of the concatenation tells us how many more times the segment in one α corresponding to \mathbf{y} , converted into a word of generators, wraps around the i th meridian than does the segment corresponding to \mathbf{x} . Taking μ_i of the concatenation gives the sum of these differences, or $\rho_i(\mathbf{y}) - \rho_i(\mathbf{x})$.

We now consider the difference in $\#$ between the two intersection points. The intersection point determines a permutation $\sigma_{\mathbf{x}}$ where $x_i \in \alpha_{\sigma(i)} \cap \beta_i$. We orient \mathbb{T}_{α} by the projection $\alpha_1 \times \cdots \times \alpha_g \rightarrow \mathbb{T}_{\alpha}$, and likewise for \mathbb{T}_{β} . The orientation of $\text{Sym}^g(\Sigma)$ is given by the orientation of $T_{x_1}\Sigma \oplus \cdots \oplus T_{x_g}\Sigma$. Then $\mathbb{T}_{\alpha} \cap_{\mathbf{x}} \mathbb{T}_{\beta}$ has local sign

$$\text{sgn}(\sigma_{\mathbf{x}})(-1)^{\frac{g(g-1)}{2}} \cdot (\alpha_{\sigma(1)} \cap_{x_1} \beta_1) \times \cdots \times (\alpha_{\sigma(g)} \cap_{x_g} \beta_g)$$

or

$$\text{sgn}(\sigma_{\mathbf{x}})(-1)^{\frac{g(g-1)}{2}} (-1)^{\#(\mathbf{x})}.$$

The difference in sign between \mathbf{y} and \mathbf{x} is then multiplication by $\text{sgn}(\sigma_{\mathbf{y}})\text{sgn}(\sigma_{\mathbf{x}}) \cdot (-1)^{\#_{\mathbf{y}} - \#_{\mathbf{x}}}$. This is also the difference in sign between terms in the determinant, and corresponds to the $\mathbb{Z}/2\mathbb{Z}$ grading in Ozsváth and Szabó [12, Section 10.4].

Thus, if we consider those intersection points with $\rho(\mathbf{x}) = \bar{v}$, for a given vector \bar{v} , we recover the intersection points for a given filtration index since ρ satisfies the index relation. In addition, these each correspond to the term $h_1^{v_1} \cdots h_k^{v_k}$ and occur with sign given by the $\mathbb{Z}/2\mathbb{Z}$ grading of the Heegaard–Floer homology, which is also the sign of the corresponding term in the determinant. For rational coefficients, the sum of these generators with sign is the Euler characteristic of the homology group corresponding to \bar{v} for this filtration index. □

8 State summation for Alexander invariants of string links in \mathcal{S}^3

For the Heegaard diagram derived from a generic projection of a string link S in $D^2 \times I$ we can give a more precise description of the generators of $\widehat{CF}(S^3, S; \mathfrak{s}_0)$ in terms of the combinatorics of a planar graph. The intersection points of $\mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$ will correspond to a subset of the maximal spanning forests of this graph, subject to

certain constraints imposed by the meridians. As in [10; 5], we will then use these forests to compute the coordinate functions, \mathcal{F}_i , for a filtration index, and a function G , which will yield the grading of the corresponding generator of \overline{CF} . We do this by showing that certain sums of weights assigned to each crossing of the projection satisfy the same relations, relative a homotopy classes ϕ , as the indices and gradings of the chain complex. As we have already seen that the Euler characteristic of the chain complex, suitably interpreted, gives $\tau(S)$, we will also obtain a combinatorial way to compute $\tau(S)$ from a state summation. Once we have adapted the spanning tree construction to apply to string links, the argument is an analog of the argument in [10], albeit with more cases.

8.1 Planar graph preliminaries

We consider planar graphs in the unit square, I^2 . Choose a number, k , and place $\frac{k}{2}$ vertices, marked by $*$, along the bottom edge when k is even ($\frac{k+1}{2}$ when k is odd). Place additional vertices, labeled by \bullet , along the top edge until there are k vertices in total. Let Δ be a *connected*, planar graph in I^2 which includes these vertices, but whose other vertices and all its edges are in the interior of the square. We let F be a maximal spanning forest for Δ , with a tree component for each $*$ on the boundary, rooted at $*$, and oriented away from its root.

We may define a dual for Δ by taking its planar dual inside I^2 , Δ^* , and placing the vertices that correspond to faces of $I^2 - \Delta$ touching ∂I^2 on ∂I^2 . Since Δ is connected, this choice of arc on the boundary is unambiguous. There is one vertex which corresponds to the left side of the square. We replace it with an $*$ and continue counter-clockwise, changing boundary vertices to $*$ until we have altered $\frac{k}{2} + 1$ (or $\frac{k+1}{2}$, k odd). This graph must also be connected.

We say that F admits a dual forest if the edges in Δ^* corresponding to edges of $\Delta - F$ form a maximal spanning forest, F^* , with each component rooted at a single $*$. In that case, we orient the forest away from its roots. Not every F admits a dual forest: a component of F^* may contain two $*$. We consider the set $\overline{\mathcal{F}}$ of forests in Δ that are part of a dual pair (F, F^*) . We will encode F^* in the diagram for Δ by inscribing the edges in $\Delta - F$ with a transverse arrow which concurs with the orientation of F^* .

Now consider a string link, S in $D^2 \times I$ and a generic projection of S into I^2 . We decompose $S = S_1 \cup \dots \cup S_l$, where each S_j consists of a maximal string link with connected projection, ie one whose projection into I^2 forms a connected graph.

Lemma 8.1 *For the Heegaard decomposition of S^3 defined by the connected projection of a string link, S , there is a one-to-one correspondence between the generators of $\overline{CF}(S)$ and the set of dual pairs, $\overline{\mathcal{F}}$, for a planar graph $\Delta \subset I^2$ as above.*

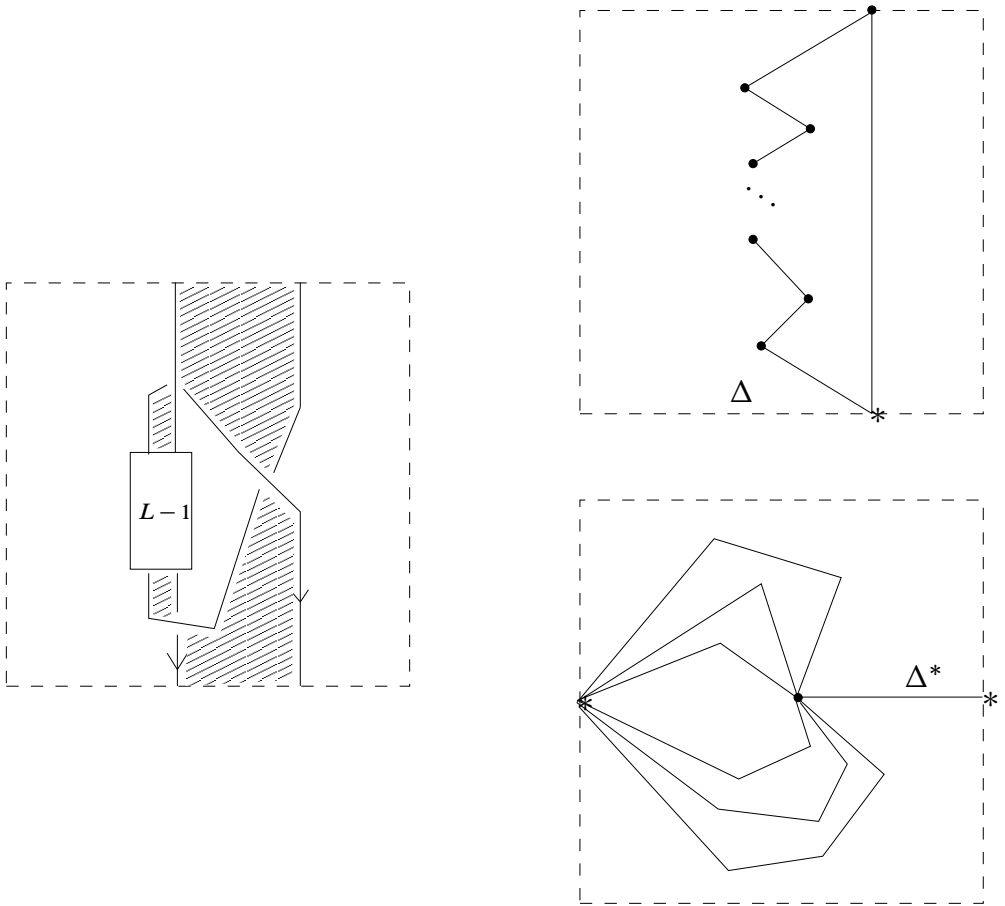


Figure 3: An example of the ancillary graph for a simple string link.

Proof The regions in $I^2 - p(S)$ can be colored with 2 colors as the projected graph is 4-valent. We label the leftmost region with the letter “U” and color it white. We then alternate between black and white across the edges of the projection. By using vertices corresponding to the black regions and edges corresponding to crossings where two (not necessarily distinct) black regions abut, we may form an second planar graph. For those regions touching the border of I^2 , the vertex should be place on the boundary.

We replace the vertex of each black region abutting the bottom edge of I^2 with an $*$. Thus, we have a graph embedded in D^2 with k vertices on the boundary, $\frac{k+1}{2}$ of which are $*$ when k is odd and $\frac{k}{2}$ when k is even. An example is given in Figure 3. In particular, Δ is connected: starting at a vertex of Δ inside a black region of $p(S)$ we can take a path to $p(S)$ and follow $p(S)$ to a point in the boundary of any other

black region. Since any edge of $p(S)$ touches a black region, we may then perturb the path into the black regions and find a corresponding path in the graph of black regions. If we place the U in all the white regions abutting the bottom edge of I^2 , then the graph of white regions is connected and, replacing the U with $*$ is the dual graph, Δ^* , from above.

Suppose we use the Heegaard decomposition of S^3 arising from the diagram for S . Following [10], we describe an intersection point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ by local data at the vertices of $p(S)$. For each non-meridional α there are four intersection points with $\{\beta_i\}_{i=1}^g$, corresponding to the four regions in the projection abutting the crossing defined by α . For each meridian there are one or two intersection points depending upon whether it intersects the region U . However, there can be only one choice along all the meridians which assigns each meridian to a distinct β . We will place a \bullet in the quadrant corresponding to the intersection point at each vertex.

Every intersection point corresponds to a pair of dual maximal spanning forests in the black and white graphs of the projection. The unique choice along the meridians corresponds to the rooting of the forests. We then choose the edges in the black graph which join two regions through a quadrant marked with a \bullet . As each black region contains a \bullet , this produces a subgraph, F , containing all the vertices of Δ . We can perform the same operation in the white graph to obtain a second sub-graph.

Furthermore, all the components of these sub-graphs are trees. A cycle in F would bound a disc in S^2 not containing a region labeled U . Rounding the crossings of $p(S)$ along the cycle, we find a 4-valent planar graph with B_{in} crossings and $B_{in} + 1$ faces not touching the cycle. The original intersection point must form a 1-1 correspondence between these faces and crossings as all the surrounding faces were consumed by the cycle. There can be no such identification and thus no cycle in the black graph. Similarly, if the intersection point does not produce a forest in the white graph there is a contradiction. Thus we have two maximal spanning forests.

Every component must contain precisely one $*$. It cannot contain more as there is a \bullet for each edge in the tree and for each root. In order for the number of edges plus roots (the α) to equal the number of β there must be precisely one root. Thus the two sub-graphs are a dual pair of maximal spanning forests for the graphs of black and white regions.

Conversely, the arrows on the edges of Δ found from a dual pair (F, F^*) tell us how to complete the assignment of the α to β from the unique assignment along the meridians: for each non-meridional α_i we choose the intersection point in $\alpha_i \cap \beta_{\sigma(i)}$ pointed to by the arrow on the edge corresponding to the crossing defined by α_i . The existence of

F^* ensures that no arrow contradicts the assignment along the meridians by pointing into a region labeled with U . □

8.2 A variant of the clock theorem

We will now examine the structure of the set of dual maximal spanning forests. As in Gilmer–Litherland [2] we consider two moves performed on the decorations a dual pair inscribes on Δ : the clock and counter-clock moves. These are moves that interchange the two pictures in Figure 4. There should be a face – not labeled with a U – of $I^2 - \Delta$ abutting these two edges at their common vertex. Note that this allows a portion of Δ to be wholly contained in the interior of the face.

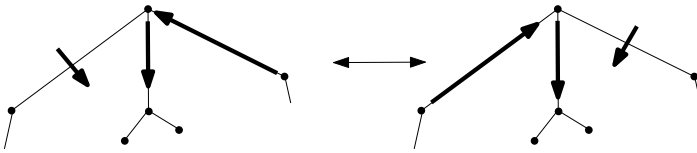


Figure 4: The clock (\rightarrow) and counter-clock (\leftarrow) moves on maximal forests at a vertex. The central sub-tree should be wholly contained in the face around which the move is made.

A clock move performed on a forest, F , in Δ corresponds to a clock move performed on F^* in Δ^* . These moves take dual maximal forest pairs into dual maximal forest pairs. If the edge with the transverse arrow joins two distinct components of F , then the new oriented sub-graph, F' , of Δ after the clock move is still a forest. The portion of F beyond the vertex where the move occurs is a tree which does not contain the vertex on the other component. The clock move merely prunes this portion of F and glues it to the other component. If the edge with transverse arrow joins vertices in the same component of F , it is conceivable that a cycle could form. However, this can only happen if the original transverse arrow points out of the disc bounded by this cycle, and thus a root of the dual graph must be contained in the cycle. Since those roots lie on ∂I^2 , this cannot happen.

For a connected, finite planar graph, Δ , with only one root, the structure of maximal spanning trees is already understood, [2; 5]. We require that the root be in the boundary of U , the unbounded component of $\mathbb{R}^2 - \Delta$. Pick one of the trees, T . Each additional edge in Δ , when adjoined to the tree, divides the plane into a bounded and an unbounded component. Draw an arrow pointing into the bounded component along each of these edges. This is the decoration inscribed by the dual tree as before.

In [2], Gilmer and Litherland reprove and strengthen Kauffman’s clock theorem, [5].

Theorem 8.2 (The Clock Theorem) *The set \mathcal{T} of maximal, spanning trees is a graded, distributive lattice under the partial order defined by $T \geq T'$ if we can move from T to T' solely by using clock-moves.*

We will only need that any $T \in \mathcal{T}$ can be obtained from any other T' by making clock and counter-clock moves. It is shown in [2] that only a finite number of clock (or counter-clock) moves can be made successively before we reach a tree not admitting another such move. Furthermore, this tree is unique for the type of move. Therefore, we can go from any tree to any other by continually making clockwise moves until we reach the unique un-clocked tree and then make counter-clock moves to get to the other tree.

8.2.1 The clock theorem for forests Our analog of the clock theorem is the following lemma.

Lemma 8.3 *The dual pairs (F, F^*) for the graph of black regions, Δ , found for a connected projection of a string link may be converted, one into another, by clock and counter-clock moves performed on the decorations coming from the rooting of the string link.*

Proof The vertices ∂I^2 divide the boundary into arcs. We draw an arrow into the the regions labeled by U , across the corresponding arcs. Place arrows pointing out along the other edges. A dual pair (F, F^*) for the string link S extends these arrows in the sense that each face has exactly one arrow pointing into it.

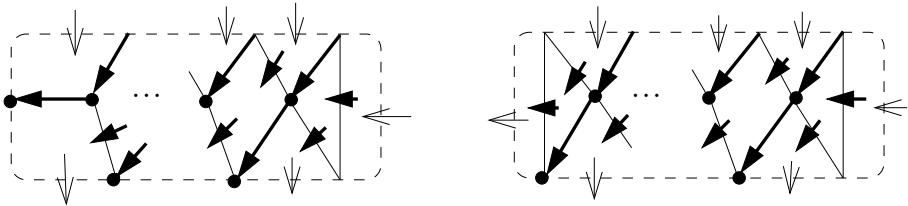


Figure 5: Graphs in I^2 admitting a unique pair of maximal spanning forests. In particular, no counter-clock moves can be performed in them. Furthermore, embedded appropriately in a planar graph with a maximal tree inscribing the same decorations, no counter-clock move in the planar graph can alter the decorations in this region. We use the one on the left when k is even, the one on the right when k is odd.

We use the arrows on ∂I^2 to extend Δ to a planar graph with a single root, so that each dual pair (F, F^*) corresponds to a unique maximal spanning *tree* in the resulting

graph. To do this we consider a new square with the reverse of the decorations (“out” goes to “in” and vice-versa) on the boundary of the old square, with the exception that the decorations on the right edge and the first root from right to left remain the same. We can then extend these decorations by a graph, Δ' , and a dual pair as in Figure 5. By inspection, these are the unique decorations providing a dual pair of forests in this graph and extending the boundary conditions. (No two arrows may point into the same white region, if the dual is to be a forest).

We can glue this decorated graph to the one from S to obtain a planar graph where there is only one U , corresponding to the leftmost edge of S , one $*$, and the pair (F, F^*) becomes the decorations from a maximal spanning tree as in the clock theorem. Furthermore, a maximal spanning tree in the glued graph inscribing the decorations on the Δ' -portion as in Figure 5 corresponds to a maximal pair (F, F^*) in the graph of black regions for S .

Now we perform counter-clock moves until we reach the maximally clocked tree. At no time do these moves disrupt the decorations in the Δ' -portion. No such move can occur on a face in Δ' as there is no vertex with the requisite arrangement of tree edge and transverse arrow. Furthermore, the Δ' region can be disrupted from outside only when a counter-clock move occurs on a face abutting Δ . Noting that the arrows point out of the vertices on the bottom, and into the vertices on the top, inspection shows that no counter-clock move can occur on such a face at a vertex from Δ' . This is *not* true for clock moves, which can occur on the top left of Δ' . Finally, since the transverse arrows point into the faces that were formerly labeled by U , these arrows are never altered. Consequently, no counter-clock move ever involves a face formerly labeled with U . However, as any forest pair for S may be extended to a tree for the new graph and counter-clocked to the maximal clocked tree, there is always a sequence of counter-clock and clock moves, not involving Δ' , which connect any two pairs for S . □

8.3 State summation

Following [10], we will prescribe weights at the crossings of the string link as in Figure 6, and extend those weights to apply to more than one component. For each intersection point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we consider the associated dual pair (F, F^*) and locally place the \bullet in the quadrant at each crossing pointed to by the arrow on the edge of Δ corresponding to that crossing. We sum the weights from the marked region at each crossing and meridian. Likewise, there are weights to calculate the grading for each intersection point, see Figure 7. For each intersection point we add the weights over all the crossings without considering which strands appear. The meridians do not contribute to the grading.

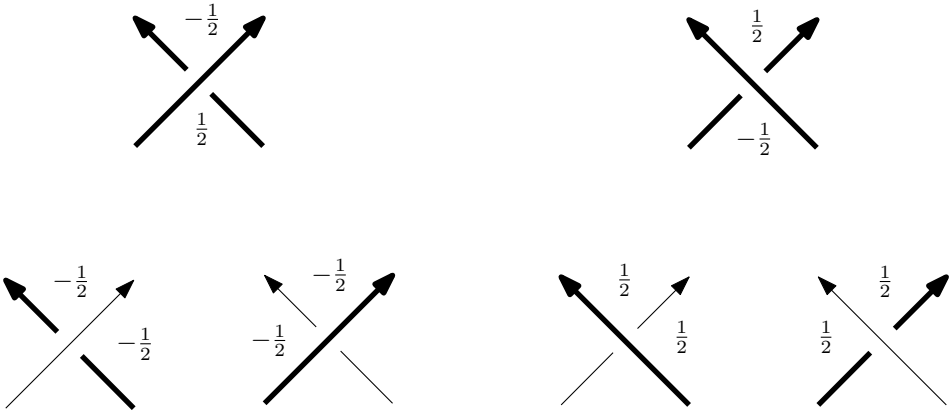


Figure 6: Filtration weights depicted for the thick strand. If this strand is the i th strand then the weights contribute to a sum computing \mathcal{F}_i . The meridians do not contribute to any of the weights; nor do crossings not involving the thick strand contribute to the computation of \mathcal{F}_i .



Figure 7: Weights for the absolute grading.

Proposition 8.4 Let (F, F^*) be a dual pair of forests for the graph Δ , arising from a projection of the string link $S \subset D^2 \times I$. At each crossing, c , the dual pair assigns an arrow pointing into one of the quadrants. Assign a value $\epsilon_i(c)$ given by the weight in Figure 6, with the thick strand being the i th strand, in the quadrant determined by (F, F^*) . Assign a value $g(c)$, by taking the weight in Figure 7. Let \mathbf{x} be the generator of $\widehat{CF}(S)$ corresponding to (F, F^*) . Then

$$\mathcal{F}_i(\mathbf{x}) = \sum_{c \in C} \epsilon_i(c)$$

where C is the set of crossings, defines the i th coordinate in a filtration index for S , and

$$G(\mathbf{x}) = \sum_{c \in C} g(c)$$

is a relative \mathbb{Z} –Maslov grading for \mathbf{x} .

Note For a single knot in S^3 , a string link consists solely in the choice of a point on the knot. This gives a preferred position for the meridian and the two points w, z used in [13] to calculate the knot Floer homology. The weights then agree with those of [10].

Proof When we change from the intersection point obtained from one tree to that of another tree, differing by a clock or counter-clock move, we have shown that the change in the exponent of t_i in $\tau(S)$ is given by $(n_w - n_{z_i})(\phi)$ for any ϕ in $\pi_2(\mathbf{x}, \mathbf{y})$ that joins these intersection points. Once we verify that this also equals the change in the weights for \mathcal{F}_i , we will have the result. As the Heegaard diagram describes S^3 , an integer homology sphere, such a class ϕ must exist. Since the dual pair (F, F^*) coming from an intersection point can be obtained from the dual pair for any other intersection point by clock and counter-clock moves, we may then compute the difference in exponents for any two intersection points by looking at the difference in the overall weights. The same argument will apply to comparing G with the Maslov grading.

We must now show that the difference in grading and filtration values between one maximal forest and the forest that results after a clock or counter-clock move equals $(n_w - n_{z_i})(\phi)$ for some ϕ . We denote the local contribution to each intersection point by placing a \bullet or \circ at each crossing. We will assume that \bullet and \circ are identical for crossings that are not depicted. Following the definitions of maximal trees and clock and counter-clock moves given above, we can verify that the moves from \circ to \bullet fulfill our requirements and exhaust all possible moves. We break the argument up into cases.

Case I Figure 8 shows the cases where a counter-clock move joins intersection points at two crossings (not meridians). These correspond to unique discs, namely squares, “atop” the Heegaard surface. As the squares do not cross any of the multi-points, there will be no change in exponents corresponding to any of the three strands. This equals the change in the weights. On the other hand, squares always have a one dimensional space of holomorphic representatives, so the intersection point \circ has grading 1 greater than \bullet in Heegaard–Floer homology. This equals the change in the grading weights.

Case II Figure 9 shows the same alteration but with a different configuration of under and over-crossings. The homotopy class we choose is now a “square” with punctures and handles added to it. In particular, the disc travels off the end of the

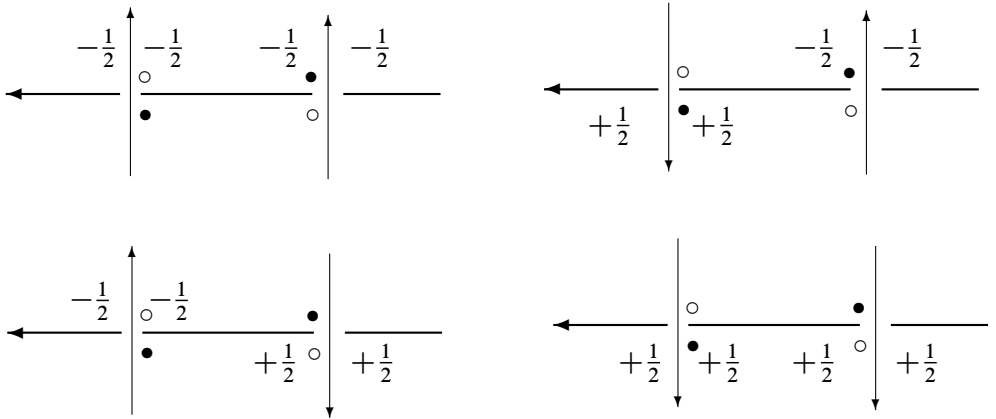


Figure 8: Case I: Weights, depicted for the horizontal strand, do not change under the alteration from \circ to \bullet . The thin strands do not need to come from the same component. The thin component on the left receives weight 0 from these configurations. The thin strand on the right receives the same weight from \circ and \bullet . The grading change occurs along the middle strand, as inspection of the crossings shows.

figure in the direction of the knot picking up punctures at crossings and joins punctures into handles if it happens to go through the same crossing twice. It terminates on the meridian corresponding to the horizontal strand. Thus, the filtrations remain unchanged except in the i th component. But the “disc” passes over z_i once, so $\mathcal{F}_i(\bullet) - \mathcal{F}_i(\circ) = (n_w - n_{z_i})(\phi) = -1$. In [10], P Ozsváth and Z Szabó show that $\mu(\phi) = 1$ for such a class, so the grading change equals the change in grading weights.

Case III For the other cases with three distinct strands, the strand on the right should go under the horizontal strand. However, if we rotate the figures in cases (1) and (2) 180° using the horizontal strand as an axis, we get precisely those cases. The disc also rotates and occurs “beneath” the Heegaard diagram. This disc represents a counter-clock move; However, we will still calculate the difference for a clock move from \circ to \bullet . For this we use ϕ^{-1} with $\mathcal{D}(\phi^{-1}) = -\mathcal{D}(\phi)$. The weights on the horizontal strands reflect across the horizontal strand as do the weights for the grading. The weights for the thin strands remain the same. Thus $\mathcal{F}_i^{III}(\bullet) = \mathcal{F}_i^{II}(\circ)$ and $\mathcal{F}_i^{III}(\bullet) - \mathcal{F}_i^{III}(\circ) = -(\mathcal{F}_i^{II}(\bullet) - \mathcal{F}_i^{II}(\circ)) = +1$ for those configurations in case II. But $(n_w - n_{z_i})(\phi_{III}^{-1}) = -(n_w - n_{z_i})(\phi_{III}) = -(n_w - n_{z_i})(\phi_{II}) = +1$ as ϕ_{II} and ϕ_{III} include the i th meridian the same number of times.

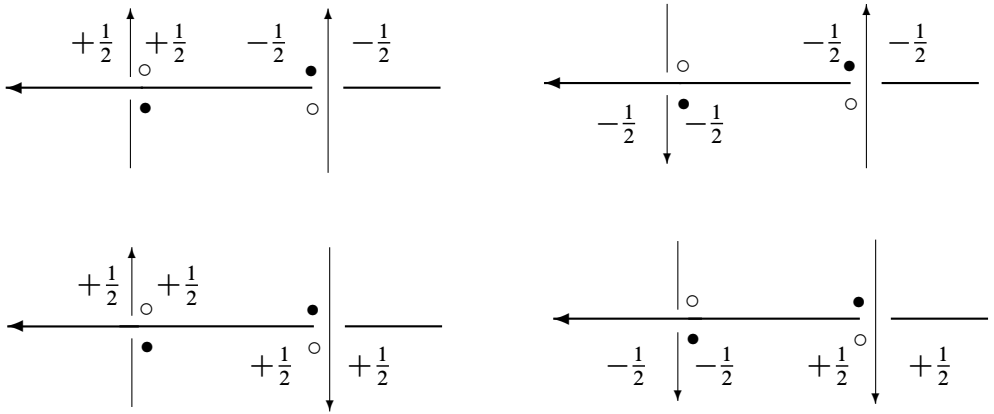


Figure 9: Case II: Weights, depicted for the horizontal strand, reduce by 1 under the alteration from \circ to \bullet . The same comments as in the caption for case I apply to the thin strands and the grading.

Case IV In the cases where two or more of the above strands are in the same component, we employ the following observations: 1) if the two thin strands belong to the same link component, then nothing changes, and 2) if the horizontal strand corresponds to the same component as one of the thin strands (or both), the sum of the weights in each quadrant differs by the same amount from that quadrant’s weight as a self-crossing. Thus the difference between intersection points of the sum is the same as the difference of the self-intersection weights. The grading computations don’t change. Inspecting the values in Figure 8 and Figure 9, we see that the filtration difference still equals the difference in the weights for the horizontal strand.

Finally, we have implicitly assumed that the horizontal strand between the two intersections is locally unknotted. Local knotting alters the topology of the domain $\mathcal{D}(\phi)$ above. Take the square in case I. If we knot the horizontal strand, there is still a class ϕ , with positive domain, joining the two intersection points, but it is a punctured disc with the same four points on its outer boundary, and one point that is both \bullet and \circ on each of its other boundaries. These new boundaries come from the faces in the projection of the local knot, and consist entirely of β curves. The α curve at each intersection point on such a boundary joins that boundary to another boundary β , possibly from the original square. According to R Lipshitz, in [8, Proposition 4.8], we can compute $\mu(\phi)$ as $e(\mathcal{D}(\phi)) + n_{\bar{x}}(\phi) + n_{\bar{y}}(\phi)$. Here $e(\mathcal{D}(\phi)) = (1 - n) - 1$, where $1 - n$ is the Euler characteristic of the domain of ϕ , and we subtract 1 due to the contributions of the acute angles from the square. Furthermore, $n_{\bar{x}}(\phi) = n_{\bar{y}}(\phi) = \frac{1}{2} + \frac{n}{2}$. The first summand, $\frac{1}{2}$ comes from the two points on the quadrilateral boundary, whereas the

other $\frac{n}{2}$ comes from the point on each of the other boundaries. Adding these shows that $\mu(\phi) = 1$. \square

Proposition 8.5

$$G(\mathbf{x}) = \sum_{c \in C} g(c)$$

is the absolute \mathbb{Q} -grading for \mathbf{x} as a generator of $\widehat{CF}(S^3)$.

The proof occupies the the remainder of this section. The argument in [10] adapts readily to show that $\text{gr}(\mathbf{y}) - \text{gr}(\mathbf{x})$ is the same as the difference in weights. However, we will show that there is an intersection point for which the sum of the weights above is the absolute grading obtained from the Heegaard–Floer theory. Thus the sum of the grading weights will equal the grading for every intersection point since the difference between these is equal for distinct intersection points. We will need a lemma before we proceed.

Lemma 8.6 *Let $G \subset \mathbb{R}^2$ be a finite graph and let H_α be the handlebody that is its regular neighborhood in S^3 . Then $S^3 - H_\alpha$ is a handlebody with the co-cores of its one-handles corresponding to the bounded faces of $\mathbb{R}^2 - G$. We choose these co-cores to be the $\{\beta_i\}_{i=1}^g$ of a Heegaard diagram for S^3 . Suppose that $\{\alpha_i\}_{i=1}^g$ contains α which intersect at most 2 β , each geometrically once. Furthermore, assume that each α links exactly one edge of G . Then there is only one point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$.*

Proof Suppose \mathbf{x} and \mathbf{y} are points in $\widehat{\mathbb{T}}_\alpha \cap \widehat{\mathbb{T}}_\beta$ and $\mathbf{x} \neq \mathbf{y}$. Then $\mathbf{x} = \{x_1, \dots, x_g\}$ and $\mathbf{y} = \{y_1, \dots, y_g\}$ where $x_i \in a_{\sigma(i)} \cap \beta_i$ and $y_i \in a_{\psi(i)} \cap \beta_i$ with $\sigma, \psi \in S_n$. If $\mathbf{x} \neq \mathbf{y}$ then $\psi^{-1} \circ \sigma$ is not the identity. It must therefore have a decomposition into cycles with at least one of length greater than or equal to 2. In the planar graph formed by placing a vertex in each bounded face of Δ , and an edge between each pair of faces abutted by an α curve, each non-trivial cycle in the cycle decomposition corresponds to a cycle in the graph. Each cycle in the graph implies the existence of a collection of α that are null-homologous in Σ , contradicting the Heegaard assumption. Thus at most one intersection point exists. Since $\widehat{HF}(S^3) \simeq \mathbb{Z}_{(0)}$, there is at least one intersection point, \mathbf{x}_0 . \square

Starting with a string link in S^3 , we use handleslides to construct a Heegaard diagram as in the lemma. The intersection point in this Heegaard diagram will correspond to one in the diagram for the string link. We will ensure that the handleslides do not alter the absolute grading. The intersection point for the string link must then have absolute grading equal to 0. This will also be the value of $G(\mathbf{x})$ for that intersection

point. Since we know that $G(\mathbf{y}) - G(\mathbf{x}) = \text{gr}(\mathbf{y}) - \text{gr}(\mathbf{x})$, the weights above will give the absolute grading of each intersection point.

To obtain an acceptable diagram, it suffices that the new α link an edge entering the crossing defined by the old α before the handlesliding as the weights for the grading only occur in the quadrant between the two edges exiting a crossing. The unique intersection point will then correspond to an old intersection point with $G(\mathbf{x}) = 0$. We ignore the points $\{z_1, \dots, z_k\}$ to find a pointed diagram for S^3 .

We order the crossings in the string link projection by the following conditions. Each crossing of L_k is larger than any crossing of strands L_i and L_j with $i, j < k$. For any k , the crossings with L_i for $i \leq k$ are enumerated from largest to smallest by the first time they are encountered while travelling backwards along L_k from the meridian. We adjust each crossing in increasing order by starting at that crossing (either for L_j with itself or between L_i and L_j with $i \leq j$) and isotoping and handlesliding the α -curve in the direction of L_j , going over all the α along that route. At self-crossings we choose to follow the edge which exits the crossing and arrives at a meridian without returning to the crossing. When we arrive at the meridian we handleslide across it, and then repeat the procedure in reverse. This produces an α linking the penultimate edge through the original crossing. Furthermore, for each crossing the ordering implies that there is a path to the meridian along the orientation of one or other strand, along which the crossing does not recur, and such that none of the α encountered have been previously altered.

The handleslides do not take the α across w . By inspecting the standard handleslide diagram, [11], we can see that the new intersection point is in the image of one of the original intersection points under the handleslide map, the one with marking in the same quadrant at each crossing, under the composition of the handleslide maps. The cobordism induced by the handleslides is $S^3 \times I$, so the formula calculating the change in absolute grading implies that the grading does not change. Moreover, since the triangle does not cross w we know that the image is $[\mathbf{x}_0, 0]$ which has absolute grading zero. This equals the absolute grading assigned by the weights.

8.4 Euler characteristic calculations

We have seen that $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ corresponds to an assignment of a quadrant to each crossing in a projection of S . Let $G(\mathbf{x})$ be the sum of the grading weights over the quadrant corresponding to \mathbf{x} at each crossing. Likewise let $\mathcal{F}_i(\mathbf{x})$ be the sums of weights assigned to the i th strand over each crossing. To each intersection point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ assign the monomial

$$(-1)^{G(\mathbf{x})} h_1^{\mathcal{F}_1(\mathbf{x})} \dots h_k^{\mathcal{F}_k(\mathbf{x})}$$

in $\mathbb{Z}[h_1^{\pm 1}, \dots, h_k^{\pm 1}]$. From the previous section we know that the sum of these monomials will equal

$$\sum_{\bar{v} \in \mathbb{Z}^k} \chi(\widehat{HF}(S, \bar{v}; \mathbb{Q})) t_1^{v_1} \dots t_k^{v_k}$$

up to multiplication by a unit, given the indeterminacy of the filtration indices. Therefore, the Laurent polynomial

$$\nabla_S(h_1, \dots, h_k) = \sum_{\mathbf{x}} (-1)^{G(\mathbf{x})} h_1^{\mathcal{F}_1(\mathbf{x})} \dots h_k^{\mathcal{F}_k(\mathbf{x})}$$

is the torsion of the string link, defined up to multiplication by a unit. Furthermore, it satisfies the Alexander–Conway skein relation in h_i at self-crossings of L_i . The proof is a comparison of the weights assigned to forests in the projections for the positive, negative, and resolved crossings. In fact, the three Reidemeister moves also preserve this summation.

We now show how the particular choice of weights actually normalizes the torsion.

Proposition 8.7 *Let L_i be the knot which corresponds to ignoring all the strands except the i th. The polynomial $\nabla_S(h_1, \dots, h_k)$ evaluates to $\Delta_{L_i}(t_i)$, the Alexander–Conway polynomial of L_i , upon setting $h_j = 1$ for $j \neq i$.*

It is shown in [10] that for 1–stranded string links (marked knots) ∇ and Δ are identical. The proof follows from two observations: 1) from Kauffman, that the polynomial formed by the weights at crossings is the Alexander–Conway polynomial, and 2) that the polynomial formed by using the first Chern class as the filtration index is symmetric due to the symmetries of Spin^c –structures on the three manifold found from 0–surgery on the knot. Since both schemes assign values to intersection points that satisfy the filtration relation, and both produce symmetric polynomials under $h \rightarrow h^{-1}$, they must be equal. For string links, we do not have an analog of the first observation, and thus will resort to model calculations. In particular, it is not sufficient to simply ignore all the strands but one (ie using only z_i and ignoring the other marked points) and appeal to the case for knots. The argument in [10] applies to a specific type of Heegaard decomposition for a knot, and while we may ignore all the z_i but one, the projection of a string link is still not such a Heegaard decomposition, due to the other strands.

Proof We first prove a supporting lemma.

Lemma 8.8 *Suppose we may interchange*

- (1) *crossings of L_i with itself,*

- (2) crossings of L_j with L_k when $j, k \neq i$. This includes self-crossings of L_j , $j \neq i$.

Then the string link S may be put in the form of a braid as found in [Figure 10](#).

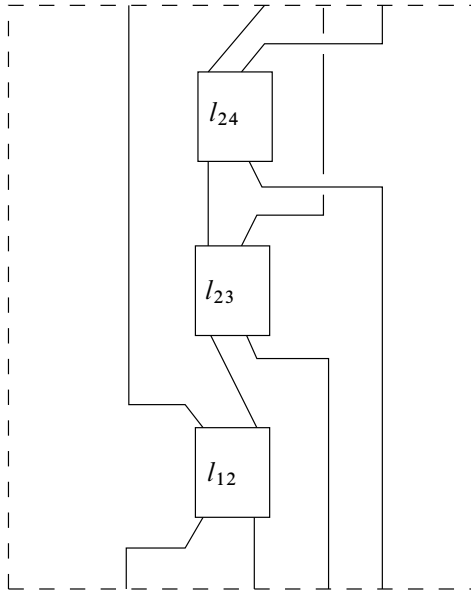


Figure 10: Reduced Form for a String Link when $i = 2$. The numbered boxes indicate the number of full twists between the two strands. For general i , we require that the other strands have no self-crossings, can link L_i some number of times representable as a box, and otherwise if $j > k$ then L_j always crosses over L_k . Furthermore, the box for L_k should be lower than that for L_j . Notice that these string links are actually braids.

Proof By interchanging self-crossings of L_i we may arrange for this strand to be unknotted. We may then isotope so that it is a vertical strand. Consider $D^2 \times I$ to be $I^2 \times I$ with the i th strand given as $(\frac{1}{2}, \frac{3}{4}) \times I$. By isotoping the other strands vertically, switching crossings where necessary, we can ensure that each is contained in a narrow band $I^2 \times (a_j - \epsilon, a_j + \epsilon)$ except when coming from or going to the boundary, when we require that they are vertical. If we look in the band $I^2 \times (a_j - \epsilon, a_j + \epsilon)$ for the j th strand, we see that all the other strands are vertical segments and the j th strand is some strand in a string link. We may isotope L_j past all the vertical segments except that for L_i . By effecting self-crossing changes in L_j we can make it unknotted. It can

then be isotoped so that it reaches the level $I^2 \times a_j$ from the bottom, winds around L_i in that plane exactly $lk(L_i, L_j)$ times, and then proceeds vertically to the top. We un-spool this winding vertically so that the result is a braid, where all the clasps are once again in $I^2 \times (a_j - \epsilon, a_j + \epsilon)$. Doing this to each of the strands produces a string link with projection as in [Figure 10](#) when we choose $a_j < a_k$ if $j < k$ for $j, k \neq i$. \square

Consider the closure, L_i formed by taking the i th strand and joining the two ends in S^3 by a simple, unknotted arc. The intersection points in the Heegaard diagram give rise to Spin^c -structures on the three-manifold formed by taking 0-surgery on this knot, [\[13\]](#). Let \mathcal{P}_i be the periodic domain corresponding to the Seifert surface in the Heegaard diagram for the 0-surgery manifold. We already know that the sum of the filtration weights assigned to \mathbf{x} by L_i differs from $\frac{1}{2}\langle c_1(\mathfrak{s}(\mathbf{x}), \mathcal{P}_i) \rangle$ by a constant, independent of \mathbf{x} , but which may depend upon S . Each of these quantities satisfies the same difference relation for a class ϕ joining two distinct intersection points. The crossings of L_j with L_k do not contribute to either calculation as the periodic domain \mathcal{P} does not change topology or multiplicities when we interchange such crossings. It consists of punctured cylinders arising from the linking of L_i with L_j or L_k and terminating on the meridian for that strand. However, the punctures and multiplicities remain the same regardless of the type of crossing between L_j and L_k . Since only the topology and multiplicities contribute to the calculation of the first Chern class for the intersection points, and the intersection points correspond under crossing changes, the value of the first Chern class does not change. Nor do the weights change as such crossings are assigned a weight of 0. As a result, interchanging crossings of L_i and L_j does not affect the value of the constant.

It requires more effort to see that interchanging self-crossings of L_i will not alter the constant. But presuming that, we see that the polynomial $\nabla_S(h_1, 1, \dots, 1)$ must be $h_i^{C_i(S)} \nabla_{L_i}$, since by [\[10\]](#) the polynomial determined by the first Chern class is the Alexander–Conway polynomial of the knot. Furthermore, we may calculate $C_i(S)$ by finding the polynomial assigned to the reduced form above, since it does not change under the moves of [Lemma 8.8](#). The i th strand is then the unknot, with Alexander–Conway polynomial equal to 1, and the value of $C_i(S)$ is 0 from our calculation for braids at the beginning of the next section.

That $C_i(S)$ does not change when interchanging self-crossings of L_i follows by seeing that the first Chern class calculation changes in the same way as the filtration weights. Notice that we may do the calculation for the multiplicities shown in [Figure 11](#).

By adding multiples of $[\Sigma]$ and \mathcal{P} we may realize a periodic region with the multiplicities for the crossing as shown in the figure. We fix the intersection point to be considered. On the left we have a local contribution to $n_{\mathbf{x}}$ of $+2$ on top, $+1$ on

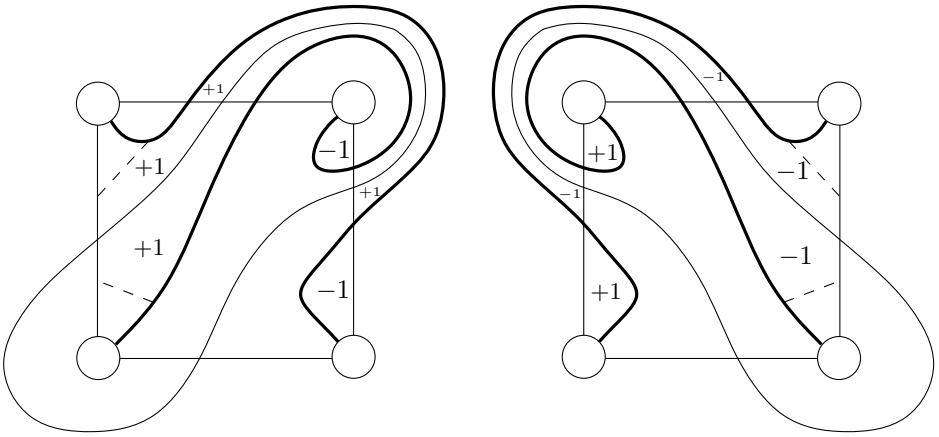


Figure 11: Local Heegaard diagrams for self-crossings of L_i . The thick curves are the α -curves. The thickest is the longitude, used as a surgery curve; while the middle thickness depicts the crossing. There is a positive crossing on the left, and a negative crossing on the right. The multiplicities are for the band added at the crossing according to Seifert’s algorithm.

left and right, and 0 on bottom. We may always pick an intersection point where the longitude pairs with a β intersecting the meridian, and thus does not incur an additional contribution from any crossing. On the right, the contribution is -2 , -1 , and 0 , respectively. All other local contributions are equal as is the value of n_w on the periodic region. The only other variation that can occur will be in the Euler measure, $\hat{\chi}(\mathcal{P})$. Outside of the depicted region, \mathcal{P} does not change. Indeed, if we divide \mathcal{P} by cutting the corners of the large $+1$ or -1 region along the dashed lines, we have a disc of multiplicity ± 1 and the remainder, R , of \mathcal{P} , which is the same in both diagrams. Then on the left $\hat{\chi}(\mathcal{P}) = \hat{\chi}(R) + 1 - 2$ while on the right $\hat{\chi}(\mathcal{P}') = \hat{\chi}(R) - 1 + 2$. The sign changes occur because of the Euler measure; in the first we add a $+1$ disc joined along two segments to a $+1$ region and at two points to a -1 region, which contributes nothing. In the second case the multiplicities are reversed, and the Euler measure must be calculated differently. Taking the difference of these, and adding the difference of the contribution from each quadrant gives $+2 - (-2) + (-1 - 1) = 2$ on top, $+1 - (-1) + (-1 - 1) = 0$ on left and right, and $0 - 0 + (-1 - 1) = -2$ on bottom. This is -2 times the difference in the weights for these quadrants, but that is precisely the factor we divide into the first Chern class to get a filtration index. Hence, the weights for any intersection point before and after a crossing change differ from the first Chern class calculation for the corresponding intersection point by the same amount. In particular, $C_i(S)$ does not change. \square

Implicit in this discussion is the following corollary.

Corollary 8.9 *For the filtration indices, \mathcal{F}_i , calculated from the weights*

$$\langle c_1(\mathfrak{s}(\mathbf{x})), \mathcal{P}_i \rangle = 2\mathcal{F}_i(\mathbf{x}).$$

9 Triviality for braids

In [Section 7](#), we saw that the Euler characteristic of the chain groups for a string link in S^3 produces the torsion of the string link $\tau(S_L)$ defined in [\[7\]](#). The torsion is known to be trivial if the string link is a pure braid. We now prove the analog for the homology groups.

Lemma 9.1 *$\widehat{HF}(S_L) \cong \mathbb{Z}_0$ when S_L is a pure braid. Using the weights from the previous section, the homology is non-trivial only for the index $(0, \dots, 0, 0)$.*

Proof There are two ways to prove this statement. The first analyzes the combinatorics of Kauffman states in the diagram. There is only one state: The meridians consume the first row of β -curves. Each crossing then has three regions already claimed, so it must use the fourth. This consumes the second set, and we proceed up the diagram. However, a more conceptual explanation may also be given. We know that the invariant we have defined does not depend upon how the strands move about on $D^2 \times \{1\}$. We may simply undo the braid to obtain the trivial string link.

The statement about filtration indices follows from the weights defined in the previous section and the observation that the unique state assigns its local contributions to the quadrant between two edges pointing into the crossing as the strands are oriented down the page. \square

10 Three operations on string links

Given two string links, S_0 and S_1 , in Y_0 and Y_1 , there are three simple operations we can perform to construct a new string link in a new manifold. We picture our three-manifolds as given by surgery on framed links in $D^2 \times I$ equipped with a string link, S , with k components. Given such diagrams for S_0 in Y_1 and S_1 in Y_2 we may 1) place them side by side to create a string link, $S_0 + S_1$ with $k_0 + k_1$ components, 2) when $k_0 = k_1$ we may stack one diagram on top of the other (as with composition of braids) to obtain the string link $S_0 \cdot S_1$, and 3) we may replace a tubular neighborhood, ie a copy of $D^2 \times I$, of the i th strand in S_0 with the entire picture for S_1 to obtain a

string satellite, $S_0(i, S_1)$. See Figure 12 for examples of each operation. The satellite operation depends on a choice of how to glue the vertical boundary of one $D^2 \times I$ into the boundary of the tubular neighborhood of the i th strand. We will designate this by a framing, which we will generally suppress as it will make no difference when we start calculating the Heegaard–Floer invariants, since it only adds braiding to the companion string link. The analysis of the second type proceeds differently than in [13]: we consider it as a closure of $S_0 + S_1$ found by joining the ends of S_0 on $D^2 \times \{1\}$ with the ends of S_1 on $D^2 \times \{0\}$ in a particular way. We always assume that the strands are oriented downwards. We will now analyze the effects these operations have on the Floer homology.

10.0.1 $S_0 + S_1$ We will show the following.

Proposition 10.1

$$\widehat{HF}(Y, S_1 + S_2; \mathfrak{s}, [\bar{j}_0] \oplus [\bar{j}_1]) \cong H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\bar{j}_0]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{j}_2])).$$

Proof We assume that we have (Y_0, S_0) and (Y_1, S_1) presented as string links in framed surgery diagrams in $D^2 \times I$. We assume that these have been put in standard form. This means we arrange that all the meridians, at the bottom of each diagram, intersect at most two β . However, there is only one choice possible for every g -tuple of intersection points due to the presence of U . Amalgamating the second string link does not affect this property for the meridians. Alternately, we can wind in the complement of $D^2 \times \{0\}$ and the amalgamation region as their union is contractible in Σ .

Topologically, the amalgamation is a connect sum of Y_0 and Y_1 where the sums occur for balls removed outside the region depicted as $D^2 \times I$. Thus for two Spin^c -structures, \mathfrak{s}_0 and \mathfrak{s}_1 , there is a unique $\mathfrak{s} = \mathfrak{s}_0 \# \mathfrak{s}_1$ on the amalgamated picture. Furthermore, $H_2 \cong H_2(Y_0; \mathbb{Z}) \oplus H_2(Y_1; \mathbb{Z})$. If the first string link has k_0 strands and the second k_1 strands, the amalgamation has filtration index taking values in $\mathbb{Z}^{k_0} / \Lambda_{Y_0} \oplus \mathbb{Z}^{k_1} / \Lambda_{Y_1}$.

Counting α and β from the portion of Y coming from Y_0 demonstrates that for an intersection point we must have an α from Y_0 pairing with a β from Y_0 and likewise for Y_1 . Even if some β extend from the Y_0 region to the Y_1 region (which we can avoid if we like), this remains true. In particular, the diagrams drawn from projections in $D^2 \times I$ will have one such β . Therefore, the generators for the new chain group are the product of generators for the two previous groups: as groups $\widehat{CF}(Y, S_0 + S_1; \mathfrak{s}) \cong \widehat{CF}(Y_0, S_0; \mathfrak{s}_0) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1)$. We may choose filtration indices for both links, choosing basepoints and complete sets of paths. The amalgamation will have $(\overline{\mathcal{F}}_0, \overline{\mathcal{F}}_1)$ as a filtration index for the complete set of paths found by using the product of the two

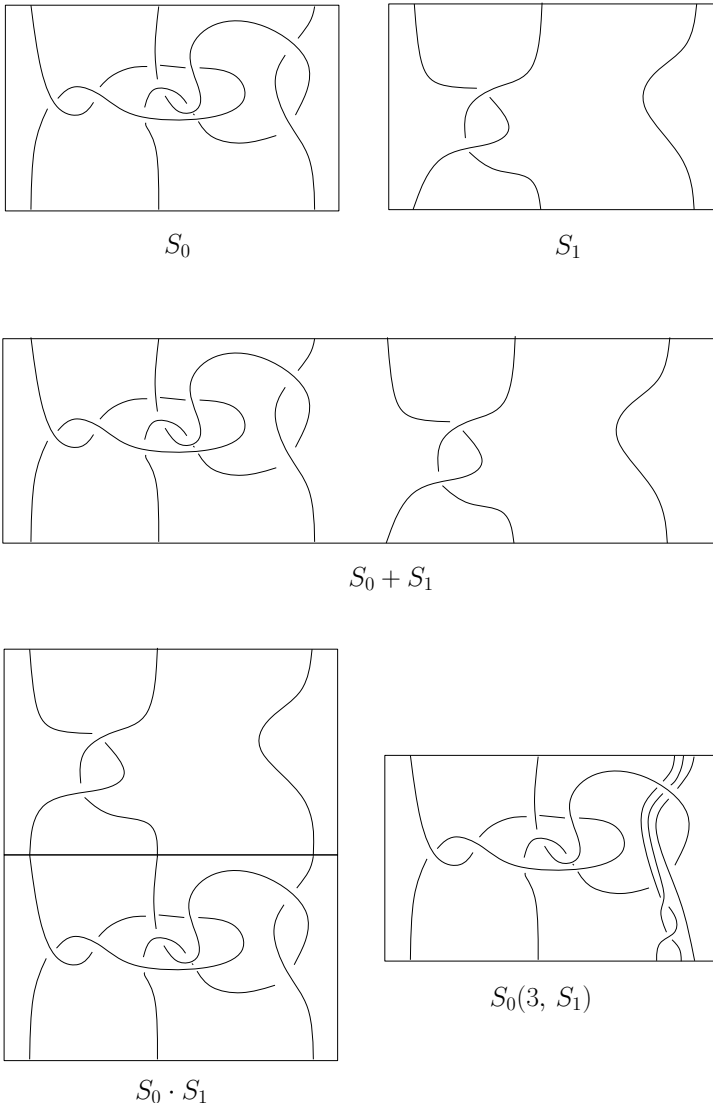


Figure 12: Examples of the three operations on string links.

basepoints and the paths from complete sets for Y_0 and Y_1 . That the domain containing w corresponds to the outer boundary of $D^2 \times I - (S_0 + S_1)$, and that the same is true for each string link individually, ensures that the paths in each complete set need not be altered. Furthermore, this region separates the domains for homotopy classes, ϕ , used in the differentials in the two complexes; thus $\widehat{\partial}_{S_0+S_1} = (\widehat{\partial}_{S_0} \otimes I) \oplus (I \otimes \widehat{\partial}_{S_1})$.

We have verified that

$$\widehat{HF}(Y, S_1 + S_2; \mathfrak{s}, [\bar{j}_0] \oplus [\bar{j}_1]) \cong H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\bar{j}_0]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{j}_2]))$$

up to gradings. The grading calculation follows as for connected sums [12]. In particular, for two torsion Spin^c -structures on Y_0 and Y_1 , the absolute grading satisfies $\text{gr}(\mathbf{x} \otimes \mathbf{y}) = \text{gr}_{S_0}(\mathbf{x}) + \text{gr}_{S_1}(\mathbf{y})$, which is all we require for string links in S^3 . We may also establish this relation by using the Maslov index calculation for Y_0 or Y_1 , presented as surgery on a link in S^3 , found in the absolute gradings section of [15]. Our assumptions include such presentations for Y_0 and Y_1 and the connect sum provides one for Y . We may use triangles with $n_w = 0$ for Y_0 and Y_1 . This allows use to use a product triangle in the calculation for Y . The first Chern class for the associated Spin^c -structure will be the sum of those for Y_0 and Y_1 . Since the intersection form splits and the Euler characteristics and signatures of the cobordisms add, we see that the gradings for the complexes for Y_0 and Y_1 add to give that for the complex on Y . \square

10.0.2 $S_1 \cdot S_2$ The second operation is called the composition of string links, the analog of composition for braids. The torsion of the composite is the product of the torsions of the two factors, [7]. We may prove the analogous result for the homologies of the string links. For $n = 1$ stranded string links composition corresponds to the connect sum of knots.

We will work with (Y_0, S_0) and (Y_1, S_1) with the assumption that S_0 and S_1 have the same number, n , of strands going from top to bottom. In addition, since we require each component of the string links to have one boundary on the top and one boundary on the bottom of the $D^2 \times I$ region in their respective manifolds, no closed component is formed by the stacking operation.

The affect on the homology groups can be calculated.

Proposition 10.2

$$\widehat{HF}(Y, S_0 \cdot S_1; \mathfrak{s}, [\bar{k}]) = \bigoplus_{[\bar{k}_0] + [\bar{k}_1] = [\bar{k}] \text{ mod } \Lambda} H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\bar{k}_0]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{k}_2]))$$

Strictly speaking, this formulation only applies to pure string links. The additional difficulty in the general case arises from the labeling of the strands. Relabeling variables to make the subscripts for each strand correspond in the composition solves the problem.

Proof Let $\Sigma_{\alpha^0 \beta^0}$ be a weakly admissible Heegaard diagram for (Y_0, S_0) with marked points w, z_1, \dots, z_n and $\Sigma_{\alpha^1 \beta^1}$ be a weakly admissible Heegaard diagram for (Y_0, S_0)

with marked points w', z'_1, \dots, z'_n . In each case, we use a diagram arising from a projection. Thus, the region containing w (or w') includes all of $\partial(D^2 \times I)$ minus the strands, and each meridian intersects at most two β -curves. The diagram with data $\alpha^0 \cup \alpha^1, \beta^0 \cup \beta^1$ is formed as in Figure 13, by joining Σ_0 and Σ_1 with a one handle attached at the gray discs. Once we add the handle, the w and w' are in the same domain, so we will consider only w' . We let γ_i^1 be a small Hamiltonian isotope of β_i^1 for all attaching curves, except for the meridians in S_1 . Each meridian for S_1 we replace with curves which traverse the tube and loop around the i th strand in each diagram, as depicted for the light gray curves in Figure 13. This effects the gluing of the different diagrams into one diagram for $S_0 \cdot S_1$. Note that after replacing meridians with the light gray curves each z'_i winds up in the same domain as w' . Furthermore, we require that the domain containing z'_i abut Θ^- .

We will analyze the cobordism generated by the triple $\alpha^0 \cup \alpha^1, \beta^0 \cup \beta^1$, and $\alpha^0 \cup \gamma^1$, where the occurrence of repeated sets of curves in a diagram indicates using Hamiltonian isotopes in the standard way. Each of the ends of this cobordism has $2n$ marked points, so we will have filtration indices taking values in \mathbb{Z}^{2n} , modulo some lattice. First, we describe the various boundary components.

Boundary I $\alpha^0 \cup \alpha^1, \beta^0 \cup \beta^1$. Topologically, this is a connect sum for Y_0 and Y_1 whose Heegaard diagram is drawn in the standard way. Each Spin^c -structure is therefore of the form $\mathfrak{s}_0 \# \mathfrak{s}_1$. Inclusion of the marked points puts us into the previous construction: amalgamation. Thus, the generators of the complex for \mathfrak{s} are products of generators from \mathfrak{s}_0 and \mathfrak{s}_1 , which we denote $\mathbf{x} \otimes \mathbf{y}$. As in the amalgamation case, we find that $\Lambda_I \equiv \Lambda_{(Y_0, S_0)} \oplus \Lambda_{(Y_1, S_1)}$ with filtration index $(\overline{\mathcal{F}}_0, \overline{\mathcal{F}}_1)$. We use the previous result to identify

$$\widehat{CF}(Y_0 \# Y_1, S_0 + S_1; \mathfrak{s}, [\overline{j}_0] \oplus [\overline{j}_1]) \cong \widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\overline{j}_0]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\overline{j}_1]).$$

Boundary II $\alpha^0 \cup \alpha^1, \alpha^0 \cup \gamma^1$. Topologically, this boundary is a $g_0 + g_1$ fold connect sums of $S^1 \times S^2$, which may be seen by isotoping the new α^1 components down the strands of S_0 , across the meridians found there, and back up the strands. However, with the additional marked points, it is unclear if the candidate for Θ_{std}^+ is closed in this diagram. We can choose Θ_{std}^+ as our basepoint, and use products of topological discs in Σ from each connect sum component as our complete set of paths. If we denote by e_i , the i th basis vector in \mathbb{Z}^{2n} , the lattice for this component will be $\Lambda_{II} \equiv \text{Span}\{e_i - e_{n+i}\}$. We now argue that Θ_{std}^+ is indeed closed for the differential missing *all* marked points.

There are precisely $2^{g_0+g_1}$ intersection possible for this diagram. We may use our complete set of topological discs to see that Θ_{std}^+ has maximal grading. By the

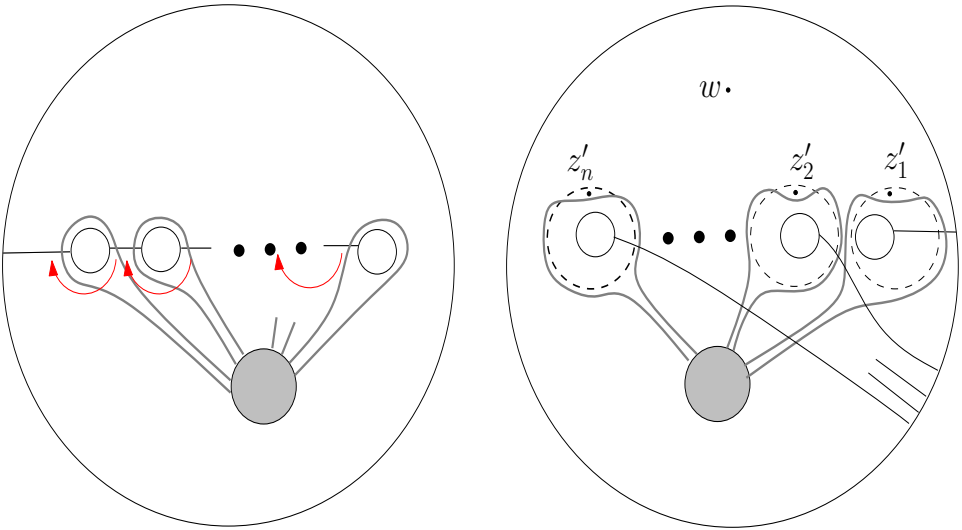


Figure 13: The left side depicts the curves in $D^2 \times \{1\}$ for S_0 and the right depicts the curves in $D^2 \times \{0\}$. The thin black curves are the β curves in the respective diagrams. The dashed curves are the meridians for S_1 . The marked points are as indicated. The process described in the text consists of attaching a 1–handle to the two gray discs and then using the gray curves to define the curves γ_1 through γ_n . These replace the meridians. The red arrows in the left hand picture show a sequence of handleslides, starting on the left, whose result applied in a mirror fashion in the diagram for S_1 is already depicted on the right. Note that each meridian intersects only one β –curve. Note also that although the γ curves can intersect more than one β , γ_n only intersects only one β , while γ_{n-1} intersects two, one of which already intersects γ_n and so forth. In particular there is a unique way to define the intersection point on these γ –curves (see the text for why the picture on the left plays no role). There is a small holomorphic triangle identifying the intersection point for the meridians with the β –curves and the γ –curves with the β –curves. Finally, note that in the $\beta^1 \gamma^1$ picture, all the z'_i are in the same domain as w .

Heegaard–Floer homology theory, we must have that the generator Θ_{std}^+ is closed for the differential only missing w . In particular, a holomorphic disc contributing to this differential cancels with some other holomorphic disc. Suppose we have two such homotopy classes of discs, ϕ and ϕ' . Splicing the inverse of one to the other, $\phi^{-1} * \phi'$ must produce a periodic domain. This periodic domain must evaluate to an element of Λ_{II} under the application of $n_{\bar{z}}$. However, $\sum n_{z_i}(\mathcal{P}) = 0$ for every periodic domain.

Since classes with holomorphic representatives must have non-negative multiplicities, it must be the case that when $n_{\bar{z}}(\phi) = 0$ so too $n_{\bar{z}}(\phi') = 0$. As the differential missing all marked points arises from a subset of the moduli spaces in the Heegaard–Floer differential and adding the extra marked points does not eliminate one disk in a cancelling pair without eliminating the other, so Θ_{std}^+ is still closed

Boundary III $\alpha^0 \cup \gamma^1, \beta^0 \cup \beta^1$. This diagram represents the result of composition. Topologically, each of the new γ^1 curves may be slid down a component of S_0 until it reaches a meridian. After sliding across the meridian, and back up the diagram, we have the connect sum of the diagrams for Y_0 and Y_1 . Again, each Spin^c -structure on Y is the sum of structures, \mathfrak{s}_i , from Y_0 and Y_1 . Additionally, $H_2(Y; \mathbb{Z}) \cong H_2(Y_0; \mathbb{Z}) \oplus H_2(Y_1; \mathbb{Z})$. However, the lattices now combine as $\Lambda_{III} \equiv (\Lambda_{Y_0} + \Lambda_{Y_1}) \oplus \bar{0}$, the span of the two original lattices, and z'_i is in the same domain as w . We required that our original diagrams be weakly admissible for our Spin^c -structures. We will see below how to extend periodic domains so that they continue to have positive and negative multiplicities in the diagram for Y . Thus, the new diagram will be weakly admissible.

In the diagram for S_0 there are g_0 α curves and g_0 β curves. As we have not changed these, the intersection point from a β_j^0 curve that contributes to a generator must come from an intersection with a curve from α^0 , even if β_j^0 intersects a curve γ_i . Hence, the intersection point on any curve γ_j must come from an intersection with a curve in β^1 . For the γ curves crossing the tube, these intersections have precisely the same form as intersections with the meridians in the original diagram for Y_1 . This allows us to establish a one-to-one correspondence between the product of generators for Y_0 and Y_1 and those of Y . The chain complex, as an abelian group, is the product of the original chain complexes. These generators we denote \mathbf{xy} .

Consider a class ϕ with $n_w(\phi) = 0$ in either of the original diagrams. If ϕ is in the lower diagram, we may use ϕ in the diagram for Y as the condition on n_w implies that the domain of the disc does not extend into the upper diagram: a small region at the top of each strand lies in the domain containing w in the diagram for Y_0 . If ϕ occurs in the upper diagram, its domain may cross $n_{z'_i}$ and include copies of the meridians for S_1 in its boundary. In the diagram for Y , we may extend this disc by following the strand down to a meridian from S_0 . At crossings, the disc gains a boundary component and an intersection point, or a copy of a framed component. However, the extension, ϕ' , will still have $n_w(\phi') = 0$, and $n_{z_i}(\phi')$ will be the same as $n_{z'_i}(\phi)$ (note the implicit relabelling if the string links are not pure). Periodic domains will continue to have the same multiplicities in the regions coming from their respective diagrams.

Given a class $\phi \in \pi_2(\mathbf{x}, \mathbf{x}')$ with $n_w(\phi) = 0$ in (Y_0, S_0) we may use it to calculate the difference in filtration index from \mathbf{xy} to $\mathbf{x}'\mathbf{y}$ to be $\bar{\mathcal{F}}_0(\mathbf{x}') - \bar{\mathcal{F}}_0(\mathbf{x})$. This holds for

any \mathbf{y} , an intersection point for (Y_1, S_1) . Likewise, if $\phi \in \pi_2(\mathbf{y}, \mathbf{y}')$ with $n_w(\phi) = 0$ in (Y_1, S_1) , we may extend to ϕ' as in the previous paragraph. This class may be used to calculate the filtration difference from \mathbf{xy} to \mathbf{xy}' . The result will be the same as $\overline{\mathcal{F}}_1(\mathbf{y}') - \overline{\mathcal{F}}_1(\mathbf{y})$ since the extension has $n_{z_i}(\phi') = n_{z_i}(\phi)$ and $n_w(\phi) = 0$. Thus, given filtration indices on the two diagrams, we can construct a filtration index on the composite which agrees with the vector sum in the first k components: $(\overline{\mathcal{F}}_0 + \overline{\mathcal{F}}_1, 0)$. This is taken modulo Λ_{III} .

We now return to the triple $\alpha^0 \cup \alpha^1, \beta^0 \cup \beta^1$, and $\alpha^0 \cup \gamma^1$. We call the induced four manifold X . We choose on X the Spin^c -structure \mathfrak{u} that is $\mathfrak{s} \times I$ and restricts to the torsion Spin^c -structure on the $\alpha^0 \cup \alpha^1, \alpha^0 \cup \gamma^1$ -boundary. We then have $\Lambda_X \equiv \Lambda_0 \oplus \Lambda_1 + \Lambda_{II}$. We use a homotopy class of triangles to join $\mathbf{x}_0 \otimes \mathbf{y}_0, \Theta_{\text{std}}^+$, and $\mathbf{x}_0 \mathbf{y}_0$; a choice specified as the unique local holomorphic class in the argument below. As we assign Θ_{std}^+ filtration index $\overline{0}$ and this local class has $n_w(\psi) = n_{\overline{z}}(\psi) = 0$, we have the following relation for the filtration indices on generators and for some $\lambda \in \Lambda_X$:

$$\overline{\mathcal{F}}(\mathbf{xy}) = \overline{\mathcal{F}}_0(\mathbf{x}) \oplus \overline{0} + \overline{0} \oplus \overline{\mathcal{F}}_1(\mathbf{y}) + \lambda_X.$$

Topologically, the cobordism, once we fill in the second boundary, is $(Y_0 \# Y_1) \times I$. If we take the quotient mod Λ_X , we recover the filtration index on Y as $\mathbb{Z}^{2n} / \Lambda_X \cong \mathbb{Z}^n / \Lambda_{III}$ and the filtrations will add correctly. Since z'_i is in the same domain as w in the diagram for Y , there is a chain isomorphism preserving filtrations which drops their entries in the filtration index. Thus, we recover $\widehat{CF}(Y, S_0 \cdot S_1)$ as a relatively indexed complex (and not, as initially could happen, a quotient of its index group).

The Heegaard triple will be weakly admissible for the doubly periodic domains, so we may choose an area form on Σ assigning the periodic domains signed area equal to zero. As it stands, this may assign large portions of the diagram small areas because the periodic domains abutting the old meridians from S_1 in the $\alpha^0 \cup \alpha^1, \alpha^0 \cup \gamma^1$ -boundary are quite substantial. We may address this difficulty by handlesliding the portion of the γ running over the tube down the diagram for S_0 until they are close to the meridians for S_0 . By doing this, we will have introduced new intersections between individual α and β curves; however, none of these may occur in a generator. Were we to use one of them, there would be too few γ remaining in the diagram for S_1 to pair with the β found there, and no means to ameliorate this deficiency with β from the bottom region. Furthermore, nothing in the previous analysis will be changed by this alteration.

In this new diagram, there are obvious holomorphic triangles abutting each intersection point $\mathbf{x} \otimes \mathbf{y}$ and Θ_{std}^+ . These consist of $g_0 + g_1$ disjoint topological triangles embedded in Σ whose domains are contained in the support of the periodic regions from the $\alpha^0 \cup \alpha^1, \alpha^0 \cup \gamma^1$ -boundary. The triangles near the meridians for S_1 are shown in

Figure 13. None of these triangles intersect a marked point We may make those periodic domains arbitrarily small in unsigned area, forcing our local triangles to have area smaller than ϵ . Without the adjustment in the previous paragraph, we would not be able to ensure that only the triangles identified above give rise to ϵ -“small” homotopy classes. Using the induced area filtrations, the chain map decomposes into a “small” portion, which is an isomorphism, and a “large” portion:

$$F((\mathbf{x} \otimes \mathbf{y}) \otimes \Theta_{\text{std}}^+) = \pm \mathbf{xy} + \text{lower order.}$$

We see then that the chain map found by counting triangles not crossing any marked points induces an injection of $\widehat{CF}(Y_0 \# Y_1, S_0 + S_1; \mathfrak{s}, [\bar{j}_0] \oplus [\bar{j}_1])$ into $\widehat{CF}(Y, S_0 + S_1; \mathfrak{s}, [\bar{j}_0 + \bar{j}_1])$ and that the map is a chain isomorphism on $\oplus \widehat{CF}(Y_0 \# Y_1, S_0 + S_1; \mathfrak{s}, [\bar{j}'_0] \oplus [\bar{j}'_1])$ where $[\bar{j}_0 + \bar{j}_1] = [\bar{j}'_0 + \bar{j}'_1] \text{ mod } \Lambda_{III}$. Together with our analysis of boundary I, this proves the result.

Finally, as the small triangles used in the argument each have $n_w = 0$ and $\mu = 0$, and the cobordism induces the torsion Spin^c -structure, the absolute grading for the image will be the sum of the absolute gradings for the original intersection points, when \mathfrak{s}_i are torsion. Since there are handleslides in the $\alpha_0 \alpha'_1, \beta_0 \beta_1$ diagram taking the curves replacing meridians in α'_1 back to the meridians, and the “small” triangles in each handleslide map link the corresponding generators, the absolute grading for the generators in this diagram are the same as for Y_0 . □

10.0.3 $S_0(i, S_1)$ The third operation is a form of string satellite to a string link. This can be formulated using the Heegaard diagram shown in [Figure 14](#). Note that this diagram applies regardless of the framing chosen for gluing in the companion string link. That the framing does not matter at the level of Heegaard diagrams can also be seen by noting that changing the framing, which can be seen as twists around the central axis of the companion string link, is equivalent to composing S_1 with a braid and then taking a satellite with a different framing. As we have seen the portion of the Heegaard diagram corresponding to the braid can be simplified to that for the trivial string link by the reduced Heegaard equivalences. For any framing we have

$$\widehat{HF}(Y, S_1(i, S_2); \mathfrak{s}, [l_1, \dots, l_{k_1+k_2-1}]) \cong \bigoplus_{[\bar{v}'] + [\bar{w}'] = [\bar{l}] \text{ mod } \Lambda'} H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\bar{v}]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\bar{w}])).$$

The notation is explained below.

Proof For string links in S^3 , there is only one way to pair meridians with β curves to achieve an intersection point. Indeed, if we draw the α as vertices of a graph and

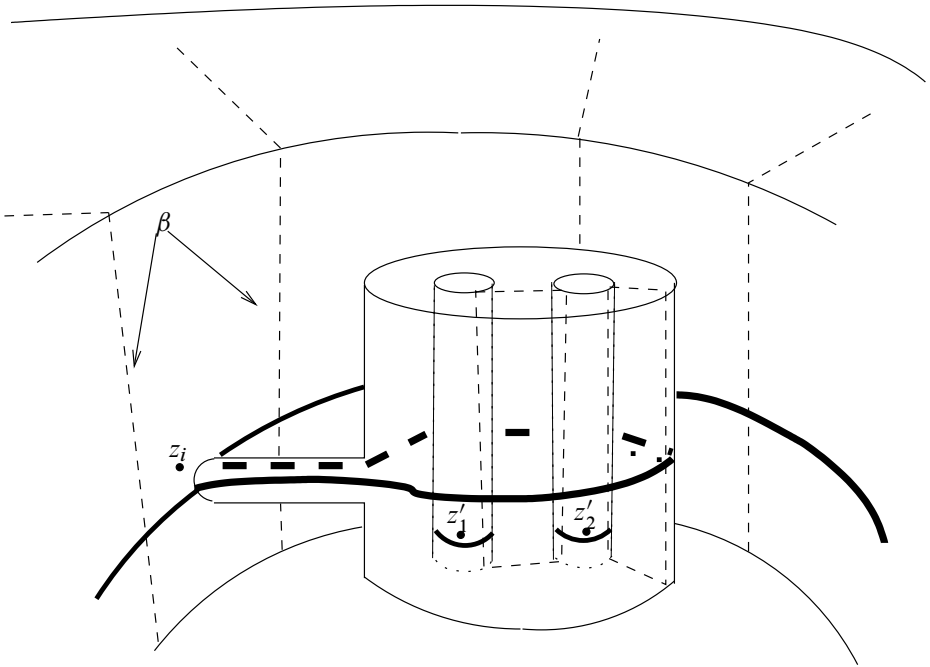


Figure 14: The Heegaard Diagram for a String Satellite. This is a cutaway view of a diagram for $S_0(i, S_1)$ with the cut through the i th strand of S_0 . All other attaching circles which have not been shown occur in one part or the other of this diagram, but do not cross from the inner cylinder to the outer one.

the β as edges, the meridians and the β that intersect them form a tree with one edge not possessing a vertex on one end. There is only one way to pair edges to end points in such a graph. The remaining β in the diagram can only intersect α according to the intersections in the original diagrams. However, the construction still applies to string links in more general manifolds, presented as surgery on framed links in $D^2 \times I$. A count of the α and β shows that generators for this new diagram occur as products of generators from the old diagrams, even when we have wound to achieve some admissibility and possibly increased the number of intersections at each meridian. Alternately, we may use the standard form to obtain diagrams to which the argument from S^3 still applies. Once again, the construction is a connect sum of two three-manifolds, and once again the Spin^c structures, etc transfer as expected.

Thus the generators of the chain complex correspond to the products of generators from the chain complexes for the constituent string links. The filtration indices, however,

differ from before. To ease the argument, we note that we may think of such a string link as the multiplication of one strand in (Y_0, S_0) n_1 times, followed by a composition with (Y_1, S_1) amalgamated with a trivial string link on n_0 strands. We already know the result of composition, hence we need only understand the string satellite where the inner constituent is an n_1 stranded trivial string link in S^3 . This has only one intersection point, hence the chain complex, as a group, is the same as that for (Y_0, S_0) for each Spin^c -structure.

Any class ϕ joining two generators, with $n_w(\phi) = 0$ can be extended to the new string link. It includes the new α and thus goes up the inner string link to the top and back down to the new meridians. For each time ϕ crosses z_i , each of the new meridians, m'_1, \dots, m'_{n_2} will be in the boundary of the new disc. In the trivial string link picture there is one generator: $u_1 \times \dots \times u_{n_1}$ with one intersection on each meridian. In particular, the new Λ in $\mathbb{Z}^{n_0+n_1-1}$ is spanned by vectors

$$v' = (\lambda_1, \dots, \lambda_i, \lambda_i, \dots, \lambda_i, \dots, \lambda_{n_1})$$

where λ_i is repeated n_1 times and $(\lambda_1, \dots, \lambda_{n_0})$ is a vector in Λ_{S_0} . We choose the extension of ϕ_x to $\phi_x \times u_1 \times \dots \times u_{n_1}$ to give our complete set of paths. The filtration index is now measured by

$$(\mathcal{F}_1, \dots, \mathcal{F}_{i-1}, \mathcal{F}_i, \dots, \mathcal{F}_i, \mathcal{F}_{i+1}, \dots, \mathcal{F}_{n_1})$$

with n_2 copies of \mathcal{F}_i .

As a consequence, discs with $n_{z_i}(\phi) = 0$ in the original diagram extend as themselves to the new diagram. In addition, any disc with $n_{z'_i}(\phi') = 0$, $i = 1, \dots, n_2$ corresponds to a disc in the original diagram with $n_{z_i} = 0$. The differentials $\widehat{\partial}$ and $\widehat{\partial}'$ must be the same, and count only classes of discs which do not need to be extended.

Putting all this together, if we denote the string satellite found by substituting S_1 in the i th strand of S_0 by $S_0(i, S_1)$ then

$$\widehat{HF}(Y, S_0(i, S_1); \mathfrak{s}, [(l_1, \dots, l_{n_1+n_2-1})]) \cong$$

$$\bigoplus_{[\vec{j}'] + [\vec{k}'] = [\vec{l}] \text{ mod } \Lambda'} H_*(\widehat{CF}(Y_0, S_0; \mathfrak{s}_0, [\vec{j}]) \otimes \widehat{CF}(Y_1, S_1; \mathfrak{s}_1, [\vec{k}]))$$

where

$$\vec{j}' = (j_1, \dots, j_{i-1}, j_i, \dots, j_i, j_{i+1}, \dots, j_{n_1})$$

and

$$\vec{k}' = (0, \dots, 0, k_1, \dots, k_{n_2}, 0, \dots, 0)$$

and

$$\Lambda' = \Lambda + \bar{0} \oplus \Lambda_1 \oplus \bar{0}.$$

We need also to calculate the absolute grading, when appropriate. When we have inserted the trivial string link into S_0 and a torsion Spin^c -structure on Y_0 we may handleslide the new α across the new meridians to arrive at a picture for a standard connect sum. At each handleslide, there is a small $\mu = 0$ homotopy class of triangles with $n_w = 0$ and admitting holomorphic representative joining each intersection point to the corresponding point in the new diagram (the product decomposition of generators is unchanged). In the connect sum picture, the gradings add – the grading of the product generator is the same as the grading of the generator from S_0 . In the cobordism induced by the handleslides, the grading does not change: $\text{gr}(\mathbf{x} \times u_1 \times \cdots \times u_{n_1}) = \text{gr}_{Y_0}(\mathbf{x})$. \square

11 Vanishing differential for alternating string links

We call a string link, S , *alternating* if there is a projection of S where proceeding along any strand in S from $D^2 \times \{0\}$ to $D^2 \times \{1\}$ encounters alternating over and under-crossings. Our goal is to prove that nothing contributes meaningfully to the differential.

Theorem 11.1 *Let S be a string link with alternating projection. Then the chain complex, $\widehat{CF}(S, \bar{j})$, arising from the projection Heegaard splitting has trivial differential for every index \bar{j} . In fact, all the generators have the same grading.*

This generalizes the result in [10] for alternating knots. We will need the knot case for the general result. The result in [10] is somewhat stronger (by identifying the grading).

First, we collect a few observations about alternating string links that will be used in the proof. In a projection, we may place a small kink (formed using the first Reidemeister move) in any strand which does not initially participate in any crossing (self or otherwise) without changing whether the projection is alternating. We will assume that this is done, so every strand participates in some crossing.

We call the i th strand, s_i when counting from left to right on $D^2 \times \{0\}$ and s^i when counting on $D^2 \times \{1\}$. If we follow s_i from $D^2 \times \{0\}$ to $D^2 \times \{1\}$ find that it is $s^{\sigma(i)}$ for some permutation $\sigma \in S_k$.

At each end of $D^2 \times I$ we may label the strands as u or o denoting whether the strand is the over or under strand in the crossing immediately preceding (following) the end of the strand on $D^2 \times \{1\}$ ($D^2 \times \{0\}$). This assigns a k -tuple called the *trace* of the

string link at that end. We denote the trace at $D^2 \times \{i\}$ by $T_i(S)$. There is an inversion on such k -tuples found by interchanging the u and o , which we denote $v \rightarrow \bar{v}$. To compose alternating string links to have an alternating result requires $T_1(S_1) = \overline{T_0(S_2)}$.

We call Morse critical points of index 1 for the projection of s_i to the I factor a “cap”. Critical points of index 0 are called “cups”.

Lemma 11.2 *We may draw an alternating projection for S so that 1) Every crossing occurs with both strands oriented up, 2) No two crossings, caps, or cups occur in the same level in the I factor.*

Proof Rotate every crossing without the correct orientations so that both strands go up. This can be done in a small neighborhood of the crossing at the expense of introducing cups and caps. The second condition is achieved by a small perturbation of the diagram. \square

Given a diagram in this form, we proceed with a few combinatorial lemmas. These are devoted to showing that an alternating string link may be closed up to give an alternating knot with certain additional properties. At each time t in the parametrization, $s_i(t)$ of the i th strand, let $f_i(t)$ be the number of intersections between the string link and a horizontal line through that point on $s_i(t)$ and strictly to its the left.

Lemma 11.3 *Let t be a time when $s_i(t)$ is in a level (in I) not containing any caps or cups, and not at a crossing. The total number of crossings encountered along the strand s_i by the time t is $\equiv |f_i(t) - i + 1| \pmod{2}$ when $s_i(t)$ is oriented up and $\equiv |f_i(t) - i| \pmod{2}$ when $s_i(t)$ is oriented down.*

Proof The number of caps plus the number of cups encountered in the i th strand, by time t , is even when the strand is oriented up, and odd when the strand is oriented down. $f_i(t)$ changes value by 1 as $s_i(t)$ goes through a cap, cup or crossing. It changes value by 2 at levels where a cap or cup occurs to the left of $s_i(t)$. By reducing modulo 2 we eliminate the latter variation. Since there are $f_i(t)$ strands to the left, having started with $i - 1$ strands to the left, the number of cups, caps and crossings must be congruent to $|f_i(t) - i + 1|$ modulo 2. Removing the parity of the number of caps and cups gives the result. \square

Corollary 11.4 *The total number of crossings encountered by the i th strand is congruent modulo 2 to $|\sigma(i) - i|$.*

This lemma has the following consequence.

Lemma 11.5 *Suppose s_i and s_j cross somewhere in S . If $T_0(s_i) = T_0(s_j)$ then $i \equiv j \pmod 2$. If $T_0(s_i) \neq T_0(s_j)$ then $i \equiv j + 1 \pmod 2$.*

Proof Consider the first time that they cross in the ordering on s_i . Suppose that $i < j$, and that there are k strands to the left of the point in s_i just before the crossing. Then there must be $k \pm 1$ strands to the left of s_j . We label each point on the strands, except at crossings, by a u or an o depending upon whether an over, or under, crossing must occur next. The labels of the points on the two strands just before the crossing of s_i and s_j must be different. If s_i has encountered an even number of crossings prior to this point, it will have label $T_0(s_i)$, otherwise it has label $\overline{T_0(s_i)}$. The same will be true of s_j . We have assumed that both strands are oriented up just before the crossing. Thus, the parity of the number of crossings involving s_i is that of $|f_i(t_i) - i + 1| = |k - i + 1|$ and the same parity for s_j is $|f_j(t_j) - j + 1| \equiv |k - j|$. If $T_0(s_i) = T_0(s_j)$, then one strand must have experienced an even number of crossings, while the other experienced an odd number. This happens when $i \equiv j$. Otherwise, both strands must encounter the same parity of crossings and $i \equiv j + 1$. \square

We decompose $S = S_1 \cup \dots \cup S_l$, where S_j consists of a maximal string link with connected projection. We apply the following lemma to each S_j .

Lemma 11.6 *For a connected, alternating string link $T_0(S)$ must be either (u, o, u, \dots) with alternating entries, or (o, u, o, \dots) .*

Proof For s_i and s_j there is a sequence s_{i_0}, \dots, s_{i_r} with $s_{i_0} = s_i$ and $s_{i_r} = s_j$ and where consecutive entries cross one another. The result follows from induction using the conclusion of the previous lemma, which also proves the base case. \square

Since any strand in S_j divides the projection into two parts, we see that each S_j must consist of consecutive strands in the diagram (along both ends). We can show more, however.

Lemma 11.7 *For an alternating projection of S , $T_0(S) = \overline{T_1(S)}$.*

This lemma guarantees that the usual closure of the string link (join s^i to s_i for all i) is alternating.

Proof We divide $S = S_1 \cup \dots \cup S_l$ and apply the corollary above to each maximal sub-string link. By the preceding lemma, we have an alternating trace for the end $D^2 \times \{0\}$. By the corollary, the other end of s_i has the same label when $\sigma(i) - i$ is even, and different labels when $\sigma(i) - i$ is odd. The result follows directly. \square

In fact, a repetition of u or o in $T_0(S)$ implies that $S = S_1 \cup S_2$ for an alternating string link. S_1 consists of those strands including and to the left of the first u , and S_2 consists of those strands including and to the right of the second.

Note Recall that we add kinks by the first Reidemeister move to strands who don't cross any other strand (including themselves). This is how they get labeled.

We now turn to proving the main theorem of this section. The chain complex of a string link S which decomposes as a union of noncrossing components $S_1 \cup \dots \cup S_l$ for the index $(\bar{j}_1, \dots, \bar{j}_l)$ is $\widehat{CF}(S_1, \bar{j}_1) \otimes \dots \otimes \widehat{CF}(S_l, \bar{j}_l)$ with the standard tensor product differential. Thus, proving the theorem for connected, alternating string links will prove the general result. Our strategy is to compare the chain group from our projection to the chain group of an alternating knot.

We make a few observations about braids. First, their projection Heegaard diagrams possess only one generator. Second, given the projection of a braid, forgetting over and under crossings, there is precisely one set of crossing data with $T_0(B) = (u, o, u, o, \dots)$. Write the braid as a product of generators or their inverses. The traces picks out which (generator or inverse) must occur. Pushing up the diagram, the traces (u, o, u, o, \dots) repeats as $T_0(B')$ for the remainder of the braid.

We choose a braid representing a permutation τ such that $\tau \circ \sigma$ is a cyclic permutation taking 1 to k . As we have just seen, we may use $T_0(B) = \overline{T_1(S)}$ to choose a braid so that $S \# B$ is still an alternating string link. We complete the construction by closing the new string link as in one of the diagrams in [Figure 15](#). The trace at the bottom of $S \cdot B$ determines which to use. Due to our assumption about τ , the result is a knot, K .

By construction, the knot is alternating. We draw a Heegaard diagram for this knot by using m_1 from the string link diagram to give a meridian for the knot. This meridian intersects only one β . It is a basing for the knot in the Kauffman state picture of [\[10\]](#). We analyze the generators of this knot. In particular, we have seen that a generator \mathbf{x} of $\widehat{CF}(S, \bar{j})$ extends uniquely to a generator \mathbf{x}' for $S \cdot B$. We would like to extend \mathbf{x}' to a generator of $\widehat{CF}(K)$. However, when we forget the states on the other meridians, we have that the shaded regions of [Figure 16](#) receive an assignment, but the others do not. There is a unique way to complete the figure to a generator \mathbf{x}'' in the knot complex, and this corresponds to a Kauffman state.

Suppose that there is a $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\widehat{\mathcal{M}}(\phi) \neq \emptyset$ and with $n_w(\phi) = n_{z_i}(\phi) = 0$. The condition at the marked points, and that for a string link there is only one choice of assignment along the meridians, implies that the homotopy class ϕ gives a homotopy class $\phi'' \in \pi_2(\mathbf{x}'', \mathbf{y}'')$ in the diagram for K . That there is no variation on the meridians

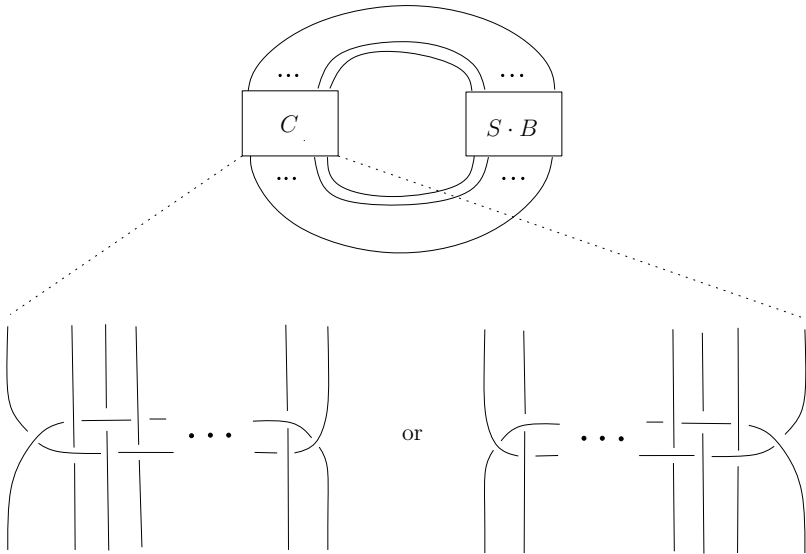


Figure 15: The closures used to construct an alternating knot from an alternating string link. The box with $S \cdot B$ indicates that the projection of the composition of the string link with the braid should lie in this box. The strands coming out of the box are assumed to match up to the strands on the left side of the diagram, without crossing in between. The projection in the box labeled C is one or other of the clasp diagrams, chosen to make the resulting projection alternating. The diagram only depicts the case when S has an even number of strands, but the case for an odd number of strands is similar. The basing (meridian) for the new knot is shown as a small line across the knot and intersecting U .

allows us to alter the generators as in the extension. The marked point conditions ensure that $\mathcal{D}(\phi)$ is wholly insulated from the additional braid and the closure construction. Furthermore, a neighborhood of each meridian is eliminated by the condition that $n_{z_i}(\phi) = 0$ as ϕ must contain whole multiples of the meridians in its boundary. There must then exist two intersection points \mathbf{x}'' and \mathbf{y}'' in the *same filtration index* for the knot (since ϕ doesn't cross the marked points) but which *differ in grading* by 1. This contradicts the statement in [10] that, for an alternating knot, the grading of every Kauffman state representing the same filtration index is the same (determined by the signature of the knot!). Hence, all moduli spaces contributing to the differential must be trivial, $\widehat{\partial}_S \equiv 0$. The same argument holds for ϕ as above with $\mu(\phi) \neq 0$, thus all the generators of the chain complex in a given filtration index in fact lie in the same grading.

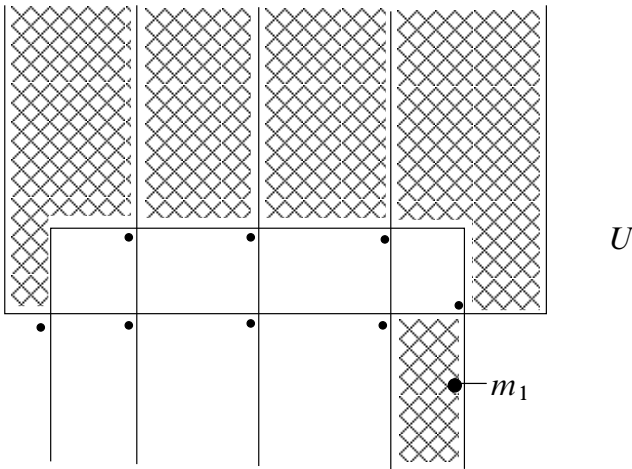


Figure 16: The Unique Extension of Generators. The diagram depicts the projection of the knot in the region of the clasp extension, but without any crossing data. The black dots indicate in which quadrant one should choose the intersection point for the α curve corresponding to that crossing. The short line m_1 should be taken to indicate the meridian of the knot. The shaded regions have already been assigned to a crossing by the generator in $S \cdot B$ (see the previous diagram) and thus are unavailable for the crossings in this portion of the knot projection. The extension of the generator depicted here corresponds to the extension of dual pairs used to prove the clock theorem for string links.

11.1 Example

In the graph of black regions for [Figure 3](#) we may form dual forests in two ways, see [Figure 17](#). First, we may have arrows pointing up the segments on the left side of the graph of black regions. Second, we may have some arrows pointing up that side, then a transverse arrow, then arrows pointing down the remainder of the segments on the left side. To join the vertices at top and bottom we have an arrow up the single segment on the right side.

When all the arrows go up the left side, they place the local intersection points on the opposite side of L_2 from that pointed to by the orientation of L_1 . Looking at the weights indicates that, regardless of the crossing type, such a tree contributes 0 to the filtration index for L_2 and 0 to the grading weight. On the other hand, when the crossings are positive all the segments on the left are assigned $-\frac{1}{2}$ for the weight on L_1 . The transverse arrow on the right will contribute nothing to either sum. A

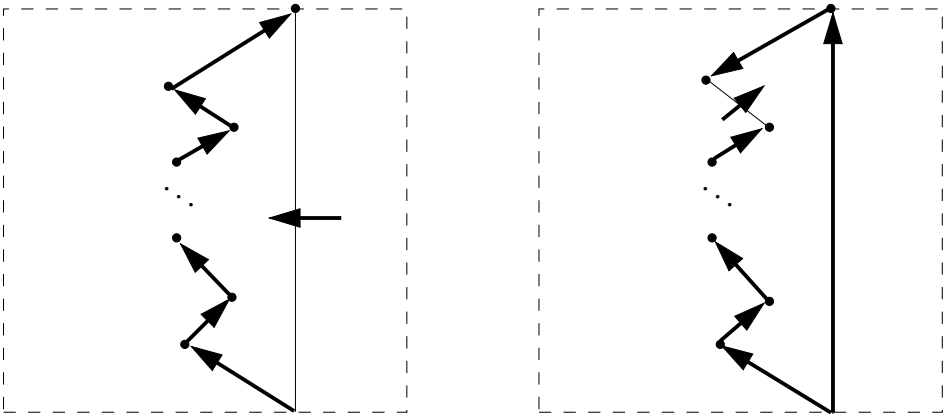


Figure 17: Examples of dual maximal forests for the graph Δ from Figure 3.

similar analysis for negative linking implies that this forest contributes $h_1^{-lk(L_1, L_2)}$ to the Euler characteristic.

For the second type of forest, let m be the number of arrows pointing up on the left side. Then m can vary between $m = 0$ and $m = 2lk(L_1, L_2) - 1 = 2L - 1$ when the linking number is positive. The monomial for this state is $h_1^{-s} h_2^{-L+s}$ when $m = 2s$ and $-h_1^{-s-1} h_2^{-L+s}$ when $m = 2s + 1$. The minus sign in the latter comes from the generator having grading -1 . Together these imply a polynomial of the form

$$(h_1^{-L} + h_1^{-L+1} h_2 + \cdots + h_2^{-L}) - (h_1^{-L} h_2^{-1} + \cdots + h_1^{-1} h_2^{-L}).$$

Since the string link is alternating, and the minus signs occur from grading -1 , we can determine the homology for the string link from the above polynomial (see the next section). It is $(v_1, v_2 \leq 0)$

$$\begin{aligned} \widehat{HF}(S, (v_1, v_2)) &\cong \mathbb{Z}_{(0)} & v_1 + v_2 &= -lk(L_1, L_2) \\ \widehat{HF}(S, (v_1, v_2)) &\cong \mathbb{Z}_{(-1)} & v_1 + v_2 &= -lk(L_1, L_2) - 1. \end{aligned}$$

12 Tangles and skein exact sequences

12.1 Tangles

Recall the definition of a tangle.

Definition 12.1 For each $k > 0$ choose a set, X_k , of k distinct points in D^2 . A (k, k') -tangle, τ , in $D^2 \times I$ is an oriented smooth 1-sub-manifold whose boundary equals $X_k \times \{0\} \cup X_{k'} \times \{1\}$.

Definition 12.2 Two tangles are flexibly isotopic if there exists a self-diffeomorphism of $D^2 \times I$, isotopic to the identity, setwise fixing the ends $D^2 \times \{0, 1\}$, and carrying one tangle into the other while preserving the orientations of the components.

Note that all braids are flexibly isotopic to the trivial string link.

Suppose we have a tangle with n -components. We will call it m -colored if there is a function, c , from the components of the tangle, τ_j , to $\{1, \dots, m\}$. We require our isotopies to preserve the value of c . We restrict to those tangles for which exactly one component in $c^{-1}(i)$ is open for each $i \in \{1, \dots, m\}$, and this component is oriented from $D^2 \times \{1\}$ to $D^2 \times \{0\}$. The collection of open components will then form a string link.

To construct a Heegaard diagram we convert the tangle, τ , to an associated string link, $S(\tau)$ in another three-manifold. First, connect the components with the same color by paths in the complement of the tangle so that, with the paths as edges and the components as vertices, we have a tree rooted at the open component for that color. These paths may be used to band sum the components together, using any number of half twists in the band which will match the orientations on components correctly. We then perform 0-surgery on an unknot linking the band once. The resulting manifold, Y , is a connect sum of s copies of $S^1 \times S^2$ where s is the number of closed components. The image of τ after performing the band sums is a multi-component ‘‘string link’’, $S(\tau)$, in the complement of a ball. We have a color map c on this string link which we use to index the components. The homology group for index \bar{j} for the colored tangle, τ , is defined to be $\widehat{HF}(Y, S(\tau); \mathfrak{s}_0, \bar{j})$ where \mathfrak{s}_0 is the torsion Spin^c structure. Underlying this construction is the following fundamental observation.

Lemma 12.3 *The flexible isotopy class of $S(\tau)$ is determined by that of τ .*

Proof This follows as for [13, Proposition 2.1]. It suffices to show that the choices made in performing the band sums do not affect the isotopy class of $S(\tau)$. Once we add the 0-framed handles, the choice of the bands no longer matters. In each band, we may remove full twists by using the belt trick to replace them with a self-crossing of the band. We may then isotope the 0-framed circle to the self-crossing and slide one of the strands in the band across the handle. Done appropriately this will undo a full twist in the band. Furthermore, by sliding across the 0-framed handles, we may

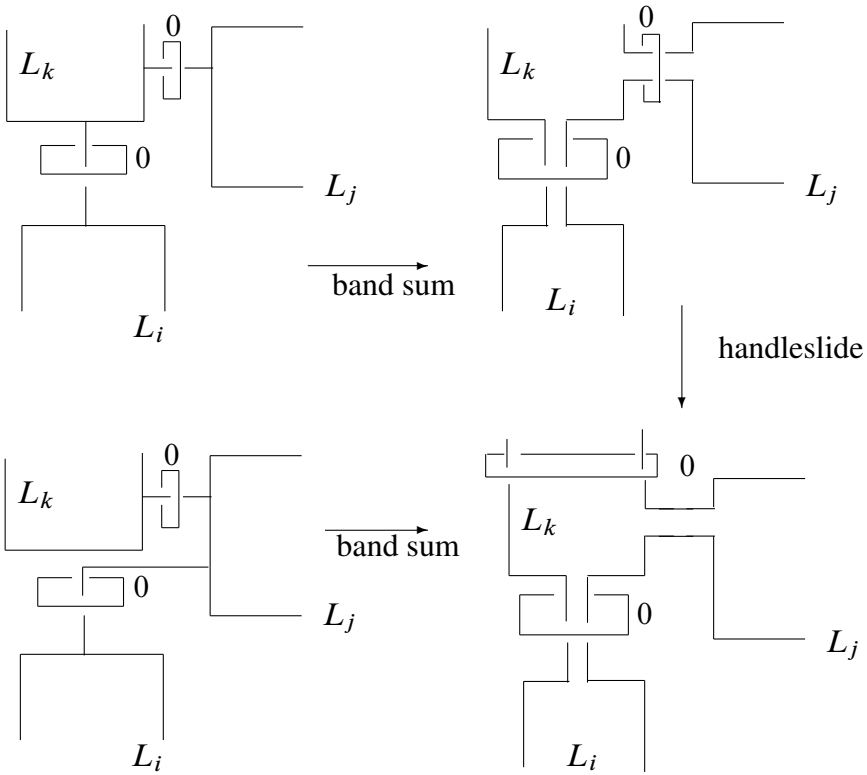


Figure 18: Kirby calculus interchange of connecting paths for a tangle.

move the bands past each other or any component in the string link. By using the trick from [13] illustrated in Figure 18, we may arrange that all the closed components of the same color are linked in a chain to a single open component. In addition, we may interchange any two components along the chain. The combination of these moves provides Heegaard equivalences, not involving the meridians, between any two ways of joining the closed components in a color to that with boundary, regardless of the paths for the band sums, or twists in the bands. \square

12.2 Long exact sequences

The tangle formulation where there is one open component for each color permits the introduction of the skein long exact sequence found in [13] where we may resolve crossings of components with the same color, but not crossings involving different colors. We should think of the resolved crossing, arising from 0–surgery on an unknot

in the long exact surgery sequence, in the context of tangles. We let L_-^c be a tangle with a negative self-crossing in color c , L_+^c be the tangle with a positive self-crossing instead, and L_0^c be the tangle resulting from resolving the crossing. As in [13] there are two sequences. If the crossing is a self-crossing of a component then

$$\rightarrow \widehat{HF}(L_-^c, \bar{j}) \rightarrow \widehat{HF}(L_0^c, \bar{j}) \rightarrow \widehat{HF}(L_+^c, \bar{j}) \rightarrow$$

whereas if the crossing occurs between different components of the same color we have

$$\rightarrow \widehat{HF}(L_-^c) \rightarrow \widehat{HF}(L_0^c) \otimes V \rightarrow \widehat{HF}(L_+^c) \rightarrow$$

where $V = V_{-1} \oplus V_0 \oplus V_{+1}$ and V_{-1} consists of a \mathbb{Z} in filtration index -1 for the color c and 0 for all others, V_0 consists of \mathbb{Z}^2 with filtration index 0 for all colors, and V_{+1} consists of a \mathbb{Z} in color c filtration index $+1$. The maps preserve the filtration indices with the tensor product index defined as the sum of those on the two factors. The proof is identical to that in [13].

Note Since the theory for tangles arises from thinking of them as string links in another manifold, the result for composition of string links extends to composition of this sub-class of tangles. The sub-class condition disallows the formation of new closed components when composing.

Acknowledgements The author would like to thank the referee, whose suggestions measurably improved the paper. The author was supported in part by NSF grant DMS-0353717 (RTG).

References

- [1] **G Burde, H Zieschang**, *Knots*, second edition, de Gruyter Studies in Mathematics 5, Berlin (2003) [MR1959408](#)
- [2] **P M Gilmer, R A Litherland**, *The duality conjecture in formal knot theory*, Osaka J. Math. 23 (1986) 229–247 [MR836713](#)
- [3] **R E Gompf, A I Stipsicz**, *4-manifolds and Kirby calculus*, Graduate Studies in Mathematics 20, Amer. Math. Soc. (1999) [MR1707327](#)
- [4] **N Habegger, X-S Lin**, *The classification of links up to link-homotopy*, J. Amer. Math. Soc. 3 (1990) 389–419 [MR1026062](#)
- [5] **L H Kauffman**, *Formal knot theory*, Mathematical Notes 30, Princeton University Press (1983) [MR712133](#)
- [6] **L H Kauffman**, *On knots*, Annals of Mathematics Studies 115, Princeton University Press (1987) [MR907872](#)

- [7] **P Kirk, C Livingston, Z Wang**, *The Gassner representation for string links*, Commun. Contemp. Math. 3 (2001) 87–136 [MR1820015](#)
- [8] **R Lipshitz**, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. 10 (2006) 955–1097 [MR2240908](#)
- [9] **R Litherland**, *The Alexander module of a knotted theta-curve*, Math. Proc. Cambridge Philos. Soc. 106 (1989) 95–106 [MR994083](#)
- [10] **P Ozsváth, Z Szabó**, *Heegaard Floer homology and alternating knots*, Geom. Topol. 7 (2003) 225–254 [MR1988285](#)
- [11] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) 159 (2004) 1027–1158 [MR2113019](#)
- [12] **P Ozsváth, Z Szabó**, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. (2) 159 (2004) 1159–1245 [MR2113020](#)
- [13] **P Ozsváth, Z Szabó**, *Holomorphic disks and knot invariants*, Adv. Math. 186 (2004) 58–116 [MR2065507](#)
- [14] **P Ozsváth, Z Szabó**, *Holomorphic disks and genus bounds*, Geom. Topol. 8 (2004) 311–334 [MR2023281](#)
- [15] **P Ozsváth, Z Szabó**, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math. 202 (2006) 326–400 [MR2222356](#)
- [16] **P Ozsváth, Z Szabó**, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, Algebr. Geom. Topol. 8 (2008) 615–692 [MR2443092](#)
- [17] **J Rasmussen**, *Floer Homology and Knot Complements*, PhD thesis, Harvard University (2003)

Department of Mathematics, Michigan State University
East Lansing, Michigan 48824, USA,

lawrence@math.msu.edu

<http://www.math.msu.edu/~lawrence>

Received: 17 January 2007 Revised: 18 June 2008