HOLOMORPHIC DISKS AND THREE-MANIFOLD INVARIANTS: PROPERTIES AND APPLICATIONS

PETER OZSVÁTH AND ZOLTÁN SZABÓ

ABSTRACT. In [23], we introduced Floer homology theories $\widehat{HF}(Y,\mathfrak{s})$, $HF^+(Y,\mathfrak{s})$, $HF^-(Y,\mathfrak{s})$ and $HF_{\mathrm{red}}(Y,\mathfrak{s})$ associated to oriented rational homology 3-spheres Y and Spin^c structures $\mathfrak{s}\in\mathrm{Spin}^c(Y)$. In the first part of this paper we extend these constructions to all closed, oriented 3-manifolds. In the second part, we study the properties of these invariants. The properties include a relationship between the Euler characteristics of $HF^\pm$ and Turaev's torsion, a relationship with the minimal genus problem (Thurston norm), and surgery exact sequences. We also include some applications of these techniques to three-manifold topology.

1. Introduction

In [23], we defined topological invariants for closed, oriented three-manifolds Y with $b_1(Y)=0$. Starting with a Heegaard diagram for Y, with Riemann surface Σ and attaching circles $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, the invariants are defined by using variants of Lagrangian Floer homology in the g-fold symmetric product Σ, relative to the pair of totally real subspaces $\mathbb{T}_\alpha=\alpha_1\times...\times\alpha_g$ and $\mathbb{T}_\beta=\beta_1\times...\times\beta_g$, with the help of a reference point $z\in\Sigma-\alpha_1-\ldots-\alpha_g-\beta_1-\ldots-\beta_g$.

To put this theory into its proper context, we must extend the definitions to include all closed, oriented three-manifolds, a task which we undertake in the first part of the present paper. The theory is more complicated in the case where $b_1(Y)>0$, because the Maslov index of a disk connecting $\mathbf{x}$ and $\mathbf{y}$ in $\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ is no longer uniquely specified by the endpoints and its local intersection number with the hypersurface $\{z\}\times\mathrm{Sym}^{g-1}(\Sigma)\subset\mathrm{Sym}^g(\Sigma)$. This issue has two important consequences. First, given a Spin^c structure $\mathfrak{s}$ over Y, the relative grading for the homology theories for $\mathfrak{s}$ is now well-defined only modulo an indeterminacy

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

The other consequence is that the boundary map in the Floer theories makes sense only when we use certain special Heegaard diagrams, for which only finitely many relative homotopy classes of Maslov index one disks admit holomorphic representatives. The details are laid out in Section 4, where the theories are defined, and in Section 5, where it is shown that the necessary admissible Heegaard diagrams exist (and can be connected in a reasonable manner).

We summarize the basic properties of HF^+ in the following:

Theorem 1.1. *Let Y be an oriented three-manifold. Then, there is a topological invariant associated to Y and a Spin^c structure $\mathfrak{s}$ over Y which is a relatively $\mathbb{Z}/\delta(\mathfrak{s})\mathbb{Z}$-graded Abelian*

PSO was supported by NSF grant number DMS 9971950 and a Sloan Research Fellowship.
ZSz was supported by a Sloan Research Fellowship and a Packard Fellowship.

group, equipped with actions

$$U\colon HF^+(Y,\mathfrak{s}) \longrightarrow HF^+(Y,\mathfrak{s})$$

and

$$(H_1(Y,\mathbb{Z})/\mathrm{Tors}) \otimes HF^+(Y,\mathfrak{s}) \longrightarrow HF^+(Y,\mathfrak{s})$$

which decrease degree in $HF^+(Y,\mathfrak{s})$ by two and one respectively. These structures satisfy the following properties:

1. *given any $\xi \in HF^+(Y,\mathfrak{s})$, $U^d(\xi)=0$ for sufficiently large d,*
2. *there are only finitely many* Spin^c *structures for which $HF^+(Y,\mathfrak{s})\neq 0$,*
3. *if $c_1(\mathfrak{s})\in H^2(Y;\mathbb{Z})$ is not a torsion class, then $HF^+(Y,\mathfrak{s})$ is a finitely generated Abelian group*
4. *if $c_1(\mathfrak{s}_0)$ is a torsion class and $b_1(Y)\leq 2$, there is an identification between $HF^+(Y,\mathfrak{s})$ and $\mathbb{Z}[U^{-1}]\otimes_{\mathbb{Z}}\wedge^* H^1(Y,\mathbb{Z})$ in all sufficiently high degrees, which is compatible with the U- and $H_1(Y)/\mathrm{Tors}$-actions.*

As in [23], there are several related constructions giving rise to groups $\widehat{HF}(Y,\mathfrak{s})$ (which, for any Spin^c structure, is always a finitely generated, relatively-graded Abelian group), $HF^-(Y,\mathfrak{s})$, and $HF^\infty(Y,\mathfrak{s})$. When $b_1(Y)\leq 2$, the latter invariant is shown in Section 11 to be determined by $H_1(Y;\mathbb{Z})$, though it is still useful as a stepping-stone in the other theories; for example, this triviality result gives the structure of the invariant $HF^+(Y,\mathfrak{s}_0)$ for torsion Spin^c structures $\mathfrak{s}_0$ in sufficiently large degrees, as stated in Part 4 of Theorem 1.1 above.

After laying the foundations of these homology theories for three-manifolds with $b_1(Y)>0$, we turn to several of their more important properties. The first application illuminates their close connection with the the minimal genus problem in three dimensions (which could alternatively be stated in terms of Thurston's semi-norm, c.f. Section 8):

Theorem 1.2. *Let $Z\subset Y$ be an oriented, connected, embedded surface of genus $g(Z)>0$ in an oriented three-manifold with $b_1(Y)>0$. If $\mathfrak{s}$ is a* Spin^c *structure for which $HF^+(Y,\mathfrak{s})\neq 0$, then*

$$\big|\langle c_1(\mathfrak{s}),[Z]\rangle\big| \leq 2g(Z)-2.$$

The second application, discussed in Section 9, gives a non-triviality criterion for the homology groups on three-manifolds with $b_1(Y)>0$, relating their Euler characteristic with Turaev's torsion function (c.f. Theorem 9.1 in the case where $b_1(Y)=1$ and Theorem 9.10 when $b_1(Y)>1$):

Theorem 1.3. *Let Y be a three-manifold with $b_1(Y)>0$, and $\mathfrak{s}$ be a non-torsion* Spin^c *structure, then*

$$\chi(HF^+(Y,\mathfrak{s})) = \pm\tau(Y,\mathfrak{s}),$$

where $\tau\colon \mathrm{Spin}^c(Y)\longrightarrow \mathbb{Z}$ is Turaev's torsion function. In the case where $b_1(Y)=1$, $\tau(\mathfrak{s})$ is calculated in the "chamber" containing $c_1(\mathfrak{s})$.

For zero-surgery on a knot, there is a well-known formula for the Turaev torsion in terms of the Alexander polynomial, see [32]. With this, the above theorem has the following corollary (a more precise version of which is given in Proposition 11.13, where the signs are clarified):

Corollary 1.4. *Let Y_0 be the three-manifold obtained by zero-surgery on a knot $K \subset S^3$, and write its symmetrized Alexander polynomial as*

$$\Delta_K = a_0 + \sum_{i=1}^{d} a_i (T^i + T^{-i}).$$

Then, for each $i \neq 0$,

$$\chi(HF^+(Y_0, \mathfrak{s}_0 + iH)) = \pm \sum_{j=1}^{d} j a_{|i|+j},$$

where $\mathfrak{s}_0$ is the Spin^c *structure with trivial first Chern class, and H is a generator for $H^2(Y_0;\mathbb{Z})$.*

Indeed, a variant of Theorem 1.3 also holds in the case where the first Chern class is torsion, except that in this case, the homology must be appropriately truncated to obtain a finite Euler characteristic (see Theorem 11.16). Also, a similar result holds for $HF^-(Y,\mathfrak{s})$, see Section 11.6.

As one might expect, these homology theories contain more information than Turaev's torsion, as can be seen, for instance, in the following result:

Proposition 1.5. *Let Y be an oriented three-manifold equipped with a* Spin^c *structure $\mathfrak{s}$, and let $\mathfrak{s}_0$ denote the* Spin^c *structure on $S^2 \times S^1$ whose first Chern class vanishes, then there is an isomorphism*

$$HF^+(Y\#(S^2\times S^1), \mathfrak{s}\#\mathfrak{s}_0) \cong HF^+(Y,\mathfrak{s}) \otimes H_*(S^1).$$

In a similar vein, we have the following:

Theorem 1.6. *Let Y_1 and Y_2 be a pair of oriented three-manifolds, and $Y_1\#Y_2$ denote their connected sum. A* Spin^c *structure over $Y_1\#Y_2$ has non-trivial HF^+ if and only if it splits as a sum $\mathfrak{s}_1\#\mathfrak{s}_2$ with* Spin^c *structures $\mathfrak{s}_i$ over Y_i for $i=1,2$, with the property that both groups $HF^+(Y_i,\mathfrak{s}_i)$ are non-trivial.*

The following result provides a link between the theory for homology three-spheres developed in [23] and its generalization given here. Let Y be an integral homology three-sphere, $K \subset Y$ be a knot, Y_0 be the three-manifold obtained as 0-surgery on K, and Y_1 be obtained by $+1$-surgery on K. We let $HF^+(Y_0)$ denote the relatively $\mathbb{Z}/2\mathbb{Z}$-graded Abelian group

$$HF^+(Y_0) = \bigoplus_{\mathfrak{s}\in\mathrm{Spin}^c(Y_0)} HF^+(Y_0,\mathfrak{s}).$$

Theorem 1.7. *If Y is an integral homology three-sphere, then there is a U-equivariant exact sequence*

$$\ldots \longrightarrow HF^+(Y) \longrightarrow HF^+(Y_0) \longrightarrow HF^+(Y_1) \longrightarrow \ldots,$$

where we view $HF^+(Y)$ and $HF^+(Y_1)$ as $\mathbb{Z}/2\mathbb{Z}$-graded groups, as well.

A more general version of the above theorem is given in Section 10 which relates HF^+ for an oriented three-manifold Y and the three-manifolds obtained by surgery on a knot $K \subset Y$ with framing h, Y_h, and the three-manifold obtained by surgery along K with framing given by $h+m$ (where m is the meridian of K and $h\cdot m = 1$), c.f. Theorem 10.12. Other generalizations include: the case of $1/q$ surgeries (Subsection 10.3), the case of integer surgeries

(Subsection 10.5), a version using twisted coefficients – which can be thought of as $\mathbb{Z}$-lift of the usual HF^+ (for the definition, see Subsection 4.11, and Subsection 10.6 for the corresponding surgery exact sequence), and an analogous discussion on $\widehat{HF}$ (Subsection 10.4). More generalizations will appear in [26].

1.1. **First application: complexity of three-manifolds and surgeries.** There is a finite-dimensional theory which can be extracted from $HF^+(Y)$, given by

$$HF_{\mathrm{red}}(Y) = HF^+(Y)/\mathrm{Im}U^d,$$

where d is any sufficiently large integer. This can be used to define a numerical complexity of an integral homology three-sphere Y:

$$N(Y) = \mathrm{rk} HF_{\mathrm{red}}(Y).$$

An easy calculation shows that $N(S^3) = 0$ (see, for instance, Proposition 8.1 of [23]).

We define a complexity of the symmetrized Alexander polynomial $\Delta_K(T) = a_0 + \sum_{i=1}^d a_i(T^i + T^{-i})$ by the following formula:

$$\|\Delta_K\|_\circ = \max(0, -t_0) + 2\sum_{i=1}^d \big|t_i(K)\big|,$$

where

$$t_i(K) = \sum_{j=1}^d j a_{|i|+j}.$$

As an application of the theory outlined above, we have the following:

Theorem 1.8. *Let $K \subset Y$ be a knot in an integer homology three-sphere, and $n > 0$ be an integer, then*

$$n \cdot \big\|\Delta_K\big\|_\circ \leq N(Y) + N(Y_{1/n}),$$

where Δ_K is the Alexander polynomial of the knot, and $Y_{1/n}$ is the three-manifold obtained by $1/n$ surgery on Y along K.

This has the following immediate consequences:

Corollary 1.9. *If $N(Y) = 0$ (for example, if $Y \cong S^3$), and the symmetrized Alexander polynomial of K has degree greater than one, then $N(Y_{1/n}) > 0$; in particular, $Y_{1/n}$ is not homeomorphic to S^3.*

And also:

Corollary 1.10. *Let Y and Y' be a pair of integer homology three-spheres. Then there is a constant $C = C(Y, Y')$ with the property that if Y' can be obtained from Y by $\pm 1/n$-surgery on a knot $K \subset Y$ with $n > 0$, then $\|\Delta_K\|_\circ \leq C/n$.*

It is interesting to compare these results to analogous results obtained using Casson's invariant. Apart from the case where the degree of Δ_K is one, Corollary 1.9 applies to a wider class of knots. On the other hand, at present, $N(Y)$ does not give information about the fundamental group of Y.

There are generalizations of Theorem 1.8 (and its corollaries) using an absolute grading on the homology theories, and also which hold for other surgery coefficients, see [26].

Corollary 1.9 should be compared with the result of Gordon and Luecke which states that no non-trivial surgery on a non-trivial knot in the three-sphere can give back the three-sphere, see [12], [13], see also [5].

1.2. **Second application: bounding the number of gradient trajectories.** We give another application, to Morse theory over homology three-spheres.

Consider the following question. Let Y be a integral homology three-sphere. Equip Y with a self-indexing Morse function $f\colon Y \longrightarrow \mathbb{R}$ with only one index zero critical point and one index three critical point, and g index one and two critical points. Endowing Y with a metric μ, we then obtain a gradient flow equation over Y. Let $m(f,\mu)$ denote the number of g-tuples of disjoint gradient flowlines connecting the index one and two critical points (note that this is *not* a signed count). Let $M(Y)$ denote the minimum of $m(f,\mu)$, as f varies over all such Morse functions and μ varies over all Riemannian metrics. Of course, $M(Y)$ has an interpretation in terms of Heegaard diagrams: $M(Y)$ is the minimum number of intersection points between the tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ for any Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ or, more concretely, the minimum (again, over all Heegaard diagrams) of the quantity

$$m(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta}) = \sum_{\sigma\in S_g}\left(\prod_{i=1}^{g}\left|\alpha_i\cap\beta_{\sigma(i)}\right|\right),$$

where S_g is the symmetric group on g letters and $|\alpha\cap\beta|$ is the number of intersection points between curves α and β in Σ.

We call this quantity the *simultaneous trajectory number* of Y. It is easy to see that if $M(Y)=1$, then Y is the three-sphere. It is natural to consider the following

Problem: if Y is a three-manifold, find $M(Y)$.

Since the complex $\widehat{CF}(Y)$ calculating $\widehat{HF}(Y)$ is generated by intersection points between $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$, it is easy to see that we have the following:

Theorem 1.11. *If Y is an integral homology three-sphere, then*

$$\mathrm{rk}\widehat{HF}(Y)\leq M(Y).$$

Using this, the relationship between $HF^+(Y)$ and $\widehat{HF}(Y)$ (Proposition 7.1), and a surgery sequence for $\widehat{HF}$ analogous to Theorem 1.7 (Theorem 10.16), we obtain the following result, whose proof is given in Section 13:

Theorem 1.12. *Let $K\subset S^3$ be a knot, and let $Y_{1/n}$ be the three-manifold obtained by $+1/n$-surgery on K, then*

$$M(Y)\geq 4k+1,$$

where k is the number of positive integers i for which $t_i(K)$ is non-zero.

1.3. **Relationship with gauge theory.** The close relationship between this theory and Seiberg-Witten theory is clearly illuminated by some of the above results. For example, adjunction inequalities exist in both worlds (compare [1] and [17]). Also, the Euler characteristic calculation above has its natural analogue in Seiberg-Witten theory (see [21], [33]).

Two additional results presented in this paper – the surgery exact sequence and the algebraic structure of the Floer homology groups – have analogues in Floer's instanton homology, and conjectural analogues in Seiberg-Witten theory, with some partial results already established. For instance, a surgery exact sequence (analogous to Theorem 1.7) was established for instanton homology, see [7], [3]. Also, the algebraic structure of "Seiberg-Witten-Floer" homology for three-manifolds with positive first Betti number is still largely conjectural, but expected to match with that stated in Theorem 1.1 (see [15], [19], [24]); see also [2] for some corresponding results in instanton homology.

However, the geometric content of these homology theories, which gives rise to bounds on the number of gradient trajectories (Theorem 1.11 and Theorem 1.12) has, at present, no direct analogue in Seiberg-Witten theory; but it is interesting to compare it with Taubes' results connecting Seiberg-Witten theory over four-manifolds with the theory of pseudo-holomorphic curves, see [29]. For discussions on S^1-valued Morse theory and Seiberg-Witten invariants, see [30] and [14].

Gauge-theoretic invariants in three dimensions are closely related to smooth four-manifold topology: Floer's instanton homology is linked to Donaldson invariants, Seiberg-Witten-Floer homology should be the counterpart to Seiberg-Witten invariants for four-manifolds. In fact, there are four-manifold invariants related to the constructions presented in this paper. Manifestations of this four-dimensional picture can already be found in the discussion on holomorphic triangles (c.f. Sections 6 and 10). These invariants are presented in [26].

2. Topological preliminaries

In this section, we briefly recall some of the topological ingredients used in the definitions of the Floer homology theories: Heegaard diagrams, symmetric products, homotopy classes of connecting disks, Spin^c structures and their relationship with Heegaard diagrams. Most of this material is introduced at greater length in Section 2 of [23]. However, here the structure of homotopy classes is richer, linked as it is with surfaces in Y (c.f. Subsection 2.4).

2.1. **Heegaard diagrams.** Recall that any oriented three-manifold can be written as a union $Y = U_0 \cup_\Sigma U_1$ where U_0 and U_1 are handlebodies which meet along a common boundary, Σ, which is an oriented two-manifold. When the genus of Σ is g, such a decomposition is called a *genus g Heegaard decomposition of Y*.

A genus g Heegaard decomposition is specified by an oriented surface Σ of genus g and two g-tuples of embedded curves $\boldsymbol{\alpha} = \{\alpha_1, ..., \alpha_g\}$ and $\boldsymbol{\beta} = \{\beta_1, ...\beta_g\}$ in Σ, where for each $i \neq j$, $\alpha_i \cap \alpha_j = \beta_i \cap \beta_j = \emptyset$ and the homology classes $[\alpha_i]_{i=1}^g$ resp $[\beta_j]_{i=1}^g$ generate g-dimensional subspaces in $H_1(\Sigma;\mathbb{Z})$. The handlebody U_0 (resp. U_1) is obtained by first attaching disks along the $\boldsymbol{\alpha}$ (resp. $\boldsymbol{\beta}$), and then filling in the remaining two-sphere by a three-ball. The data $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is called *a Heegaard diagram*. There are several natural moves between various Heegaard diagrams of a given three-manifold: *isotopies* (of the $\boldsymbol{\alpha}$, or the $\boldsymbol{\beta}$, so that at each stage the curves retain the property that for each $i \neq j$, $\alpha_i \cap \alpha_j = \beta_i \cap \beta_j = \emptyset$), *handleslides* (amongst the $\boldsymbol{\alpha}$, or the $\boldsymbol{\beta}$), *stabilizations* (replacing Σ by its connected sum with a genus one surface, and augmenting the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ by a new pair α_{g+1} and β_{g+1} supported in the new torus, which meet in a single, transverse point). We call these moves, and the inverse to stabilization, *Heegaard moves*. It is a standard result that any two Heegaard diagram for a given three-manifold can be connected by a sequence of Heegaard moves (see [27] and [28]; see also Proposition 2.1 of [23]).

Sometimes, it is convenient to fix an additional reference point $z \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$. The collection data $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ is called a *pointed Heegaard diagram*. Heegaard moves which are supported in a complement of z – i.e. isotopies where the curves never cross the basepoint z, handleslides where we do not slide across z, and any stabilization – are called *pointed Heegaard moves*. It is shown in [23] that moving the basepoint can be realized as a sequence of pointed isotopies and handleslides. Thus, any two pointed Heegaard diagrams for the same three-manifold can be connected by pointed Heegaard moves (see Figure 8 of [23]).

2.2. **Symmetric products.** If Σ is a Riemann surface of genus g, its g^{th} symmetric product $\mathrm{Sym}^g(\Sigma)$ inherits a natural complex structure. For any fixed $z \in \Sigma$, the subset $\{z\} \times \mathrm{Sym}^{g-1}(\Sigma) \subset \mathrm{Sym}^g(\Sigma)$ is a complex submanifold.

There is a natural identification $\pi_1(\mathrm{Sym}^g(\Sigma)) \cong H_1(\Sigma;\mathbb{Z})$ (see [18]). Also, for $g > 1$, $\pi_2(\mathrm{Sym}^g(\Sigma)) \cong \mathbb{Z}$, which is generated by a sphere S whose intersection number with $\{z\} \times \mathrm{Sym}^{g-1}(\Sigma)$ is $+1$. The pairing between the first Chern class of the tangent bundle $T\mathrm{Sym}^g(\Sigma)$ and this sphere S is one.

Given a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, there is a pair of embedded submanifolds, $\mathbb{T}_\alpha, \mathbb{T}_\beta \subset \mathrm{Sym}^g(\Sigma)$ where

$$\mathbb{T}_\alpha = \alpha_1 \times \ldots \times \alpha_g \quad \text{and} \quad \mathbb{T}_\beta = \beta_1 \times \ldots \times \beta_g.$$

Giving $\mathrm{Sym}^g(\Sigma)$ any complex structure induced from Σ, the spaces $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ are totally real.

2.3. **Intersection points and disks.** We will be interested in intersection points between $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$, and the disks which connect them.

If $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are a pair of intersection points, then we let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the space of homotopy classes of disks which connect them. More precisely, let $\mathbb{D} \subset \mathbb{C}$ denote the unit disk in the complex plane, and divide its boundary into a union of arcs $\partial\mathbb{D} = e_1 \cup e_2$, where e_1 resp. e_2 consists of numbers whose real part is non-negative resp non-positive. Let

$$u\colon \mathbb{D} \longrightarrow \mathrm{Sym}^g(\Sigma)$$

be a map with

$$\begin{aligned} u(-i) &= \mathbf{x}, & u(i) &= \mathbf{y} \\ u(e_1) &\subset \mathbb{T}_\alpha, & u(e_2) &\subset \mathbb{T}_\beta. \end{aligned}$$

We will call such a map a *Whitney disk*. The set $\pi_2(\mathbf{x}, \mathbf{y})$, then, is the space of equivalence classes of such maps, where $u_0 \sim u_1$ if they are homotopic through Whitney disks.

There is a natural "difference" map

$$\epsilon\colon (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \longrightarrow H_1(Y; \mathbb{Z})$$

which vanishes if and only if $\pi_2(\mathbf{x}, \mathbf{y})$ is non-empty. The difference $\epsilon(\mathbf{x}, \mathbf{y})$ is defined by taking a path a in $\mathbb{T}_\alpha$ leaving $\mathbf{x}$ and arriving at $\mathbf{y}$, a similar path b in $\mathbb{T}_\beta$ leaving $\mathbf{x}$ and arriving at $\mathbf{y}$, and then considering the image of the difference $a - b$ (thought of as a cycle in Σ) in the homology group $H_1(Y; \mathbb{Z})$.

The set $\pi_2(\mathbf{x}, \mathbf{y})$ is equipped with certain structure. Note that $\pi_1(\mathrm{Sym}^g(\Sigma))$ acts trivially on $\pi_2(\mathrm{Sym}^g(\Sigma))$, and so there is a natural action

$$\pi_2(\mathrm{Sym}^g(\Sigma)) * \pi_2(\mathbf{x}, \mathbf{y}) \longrightarrow \pi_2(\mathbf{x}, \mathbf{y}).$$

Also, if we take a Whitney disks connecting $\mathbf{x}$ to $\mathbf{y}$, and one connecting $\mathbf{y}$ to $\mathbf{z}$, we can "splice" them, to get a Whitney disk connecting $\mathbf{x}$ to $\mathbf{z}$. This operation gives rise to a generalized multiplication

$$*\colon \pi_2(\mathbf{x}, \mathbf{y}) \times \pi_2(\mathbf{y}, \mathbf{z}) \longrightarrow \pi_2(\mathbf{x}, \mathbf{z}),$$

which is easily seen to be associative. As a special case, when $\mathbf{x} = \mathbf{y}$, we see that $\pi_2(\mathbf{x}, \mathbf{x})$ is a group.

Definition 2.1. *Let A be a collection of functions $\{A_{\mathbf{x},\mathbf{y}}\colon \pi_2(\mathbf{x}, \mathbf{y}) \longrightarrow \mathbb{Z}\}_{\mathbf{x},\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta}$, satisfying the property that*

$$A_{\mathbf{x},\mathbf{y}}(\phi) + A_{\mathbf{y},\mathbf{z}}(\psi) = A_{\mathbf{x},\mathbf{z}}(\phi * \psi),$$

for each $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, $\psi \in \pi_2(\mathbf{y}, \mathbf{z})$. Such a collection A is called an additive assignment.

One important additive assignment is obtained by choosing a base-point $z \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$. If u is a Whitney disk, we can consider the algebraic intersection number

$$n_z(u) = \#u^{-1}(\{z\} \times \mathrm{Sym}^{g-1}(\Sigma)).$$

This quantity descends to homotopy classes, to give an additive assignment

$$n_z\colon \pi_2(\mathbf{x}, \mathbf{y}) \longrightarrow \mathbb{Z}.$$

This assignment can be used to define the domain belonging to a Whitney disk. Let $\mathcal{D}_1, \ldots, \mathcal{D}_m$ denote the closures of the components of $\Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$. Given a Whitney disk $u\colon \mathbb{D} \longrightarrow \mathrm{Sym}^g(\Sigma)$, the *domain associated to* u is the formal linear combination of the domains $\{\mathcal{D}_i\}_{i=1}^m$:

$$\mathcal{D}(u) = \sum_{i=1}^m n_{z_i}(u)\mathcal{D}_i,$$

where $z_i \in \mathcal{D}_i$ are points in the interior of $\mathcal{D}_i$. This quantity is obviously independent of the choice of z_i, and indeed, $\mathcal{D}(u)$ depends only on the homotopy class of u.

Definition 2.2. *Fix a reference point* $z \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$. *A* periodic domain *is a two-chain* $\mathcal{P} = \sum_{i=1}^m a_i\mathcal{D}_i$ *whose boundary is a sum of* α- *and* β-*curves, and whose* $n_z(\mathcal{P}) = 0$. *For each* $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, *a class* $\phi \in \pi_2(\mathbf{x},\mathbf{x})$ *with* $n_z(\phi) = 0$ *is called a* periodic class. *The set* $\Pi_{\mathbf{x}}(z)$ *of periodic classes is naturally a subgroup of* $\pi_2(\mathbf{x},\mathbf{x})$. *The domain belonging to a periodic class is, of course, a periodic domain.*

The algebraic topology of the $\pi_2(\mathbf{x},\mathbf{y})$ is described in the following:

Proposition 2.3. *For all* $g > 1$, $\pi_2(\mathrm{Sym}^g(\Sigma)) \cong \mathbb{Z}$. *For all* $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, *there is an isomorphism*

$$\pi_2(\mathbf{x},\mathbf{x}) \cong \mathbb{Z} \oplus H^1(Y;\mathbb{Z});$$

which identifies the subgroup of periodic classes

$$\Pi_{\mathbf{x}}(z) \cong H^1(Y;\mathbb{Z}).$$

In general, for each $\mathbf{x},\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, *if* $\epsilon(\mathbf{x},\mathbf{y}) \neq 0$, *then* $\pi_2(\mathbf{x},\mathbf{y})$ *is empty; otherwise,*

$$\pi_2(\mathbf{x},\mathbf{y}) \cong \mathbb{Z} \oplus H^1(Y;\mathbb{Z})$$

as principal $\pi_2(\mathrm{Sym}^g(\Sigma)) \times \Pi_{\mathbf{x}}(z)$ *spaces.*

For each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, the above proposition shows that the natural map which associates to a periodic class in $\Pi_{\mathbf{x}}(z)$ its periodic domain is an isomorphism of groups.

Proof. The space $\pi_2(\mathbf{x},\mathbf{x})$ is naturally identified with the fundamental group of the space $\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)$ of paths in $\mathrm{Sym}^g(\Sigma)$ joining $\mathbb{T}_\alpha$ to $\mathbb{T}_\beta$, based at the constant ($\mathbf{x}$) path. Evaluation maps (at the two endpoints of the paths) induce a Serre fibration (with fiber the path-space of $\mathrm{Sym}^g(\Sigma)$)

$$\Omega\mathrm{Sym}^g(\Sigma) \longrightarrow \Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta) \longrightarrow \mathbb{T}_\alpha \times \mathbb{T}_\beta,$$

whose associated homotopy long exact sequence gives:

$$0 \longrightarrow \mathbb{Z} \cong \pi_2(\mathrm{Sym}^g(\Sigma)) \longrightarrow \pi_1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)) \longrightarrow \pi_1(\mathbb{T}_\alpha \times \mathbb{T}_\beta) \longrightarrow \pi_1(\mathrm{Sym}^g(\Sigma)).$$

But under the identification $\pi_1(\mathrm{Sym}^g(\Sigma)) \cong H^1(\Sigma;\mathbb{Z})$, the images of $\pi_1(\mathbb{T}_\alpha)$ and $\pi_1(\mathbb{T}_\beta)$ correspond to $H^1(U_0;\mathbb{Z})$ and $H^1(U_1;\mathbb{Z})$ respectively. Hence, after comparing with the cohomology long exact sequence for Y, we can reinterpret the above as a short exact sequence:

$$0 \longrightarrow \mathbb{Z} \longrightarrow \pi_2(\mathbf{x},\mathbf{x}) \longrightarrow H^1(Y;\mathbb{Z}) \longrightarrow 0.$$

The homomorphism $n_z\colon \pi_2(\mathbf{x},\mathbf{x}) \longrightarrow \mathbb{Z}$ provides a splitting for the sequence. The proposition follows. $\square$

Remark 2.4. *The above result, of course, fails when* $g=1$*. However, it is still clear that* $\pi_2(\mathbf{x},\mathbf{y}) \longrightarrow \mathbb{Z}\oplus H^1(Y;\mathbb{Z})$ *is injective, and that is the only part of this result which is required for the Floer homology constructions described below to work. (Another reason why we do not belabor this point is that the only manifold covered by this case, when* $b_1(Y)>0$*, is* $S^2\times S^1$*.)*

2.4. **Periodic domains and surfaces in** Y**.** Given a periodic domain $\mathcal{P}$, there is a map from a surface-with boundary

$$\Phi\colon F \longrightarrow \Sigma$$

representing $\mathcal{P}$, in the sense that

$$\Phi_*[F] = \mathcal{P}$$

as chains (where $[F]$ is a fundamental cycle of F). Typically, such representatives can be "inefficient": Φ need not be orientation preserving, so F can be quite complicated. However, for chains of the form $\mathcal{P}+\ell[\Sigma]$ with no negative coefficients, we can choose F in a special manner, according to the following.

Lemma 2.5. *Consider a chain* $\mathcal{P}+\ell[\Sigma]$ *with* ℓ *sufficiently large that* $n_{z'}(\mathcal{P}+\ell[\Sigma])\geq 0$ *for all* $z'\in\Sigma-\alpha_1-\ldots-\alpha_g-\beta_1-\ldots-\beta_g$*. Then there is an oriented two-manifold with boundary* F *and a map* $\Phi\colon F\longrightarrow\Sigma$ *with* $\Phi_*[F]=\mathcal{P}+\ell[\Sigma]$ *with the property that* Φ *is nowhere orientation reversing and the restriction of* Φ *to each boundary component of* F *is a diffeomorphism onto its image.*

Proof. Write

$$\mathcal{P}+\ell[\Sigma] = \sum_{i=1}^m n_i\mathcal{D}_i,$$

(where, by assumption, $n_i\geq 0$). If $\mathcal{D}$ is the domain $\mathcal{D}_i$, then we let $m(\mathcal{D})$ denote the coefficient n_i. The surface F is constructed as an identification space from

$$X = \coprod_{i=1}^m \coprod_{j=1}^{n_i} \mathcal{D}_i^{(j)},$$

where $\mathcal{D}_i^{(j)}$ is a diffeomorphic copy of the domain $\mathcal{D}_i$.

The α-curves are divided up by the β-curves into subsets, which we call α-arcs; and similarly, the β-curves are divided up by the α-curves into β-arcs. Each α or β-arc c is contained in two (not necessarily distinct) domains, $\mathcal{D}_1(c)$ and $\mathcal{D}_2(c)$. We order the domains so that

$$m(\mathcal{D}_1(c)) \leq m(\mathcal{D}_2(c)).$$

F is obtained from X by the following identifications. For each α-arc a, if $x\in a$, then for $j=1,\ldots,m(\mathcal{D}_1(a))$, we identify

$$\Big(x^{(j)}\in\mathcal{D}_1(a)\Big) \sim \Big(x^{(j+\delta_a)}\in\mathcal{D}_2(a)\Big),$$

where $\delta_a = m(\mathcal{D}_2(a))-m(\mathcal{D}_1(a))$. Similarly, for each β-arc b, if $x\in b$, then for $j=1,\ldots,m(\mathcal{D}_1(b))$, we identify

$$\Big(x^{(j)}\in\mathcal{D}_1(a)\Big) \sim \Big(x^{(j)}\in\mathcal{D}_2(a)\Big).$$

The map Φ, then, is induced from the natural projection map from X to Σ.

It is easy to verify that the space F is actually a manifold-with-boundary as claimed. $\square$

Let $\Phi\colon F \longrightarrow \Sigma$ be a representative for a periodic domain $\mathcal{P} + \ell[\Sigma]$ as constructed in Lemma 2.5 as above. Φ can be extended to a map into the three-manifold:

$$\widehat{\Phi}\colon \widehat{F} \longrightarrow Y$$

by gluing copies of the attaching disks for the index one and two critical points (with appropriate multiplicity) along the boundary of F. This gives us a concrete correspondence between periodic domains and homology classes in Y which, in the case where $\mathbb{T}_\alpha$ meets $\mathbb{T}_\beta$, is Poincaré dual to the isomorphism of Proposition 2.3.

One can also think of the intersection numbers n_z as taking place in Y. To set this up, note that each (oriented) attaching circle α_i naturally gives rise to a cohomology class $\alpha_i^* \in H^2(Y;\mathbb{Z})$. This class is, by definition, Poincaré dual to the closed curve $\gamma \subset U_0 \subset Y$ which is the difference between the two flow-lines connecting the corresponding index one critical point $a_i \in U_0 \subset Y$ with the index zero critical point. The sign of α_i^* is specified specified by requiring that the linking number of γ with α_i in U_0 is $+1$.

Lemma 2.6. *Let $z_1, z_2 \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$ be a pair of points which are separated by α_1, in the sense that there is a curve z_t from z_1 to z_2 which is disjoint from $\alpha_2, \ldots, \alpha_g$, and $\#(\alpha_1 \cap z_t) = +1$. Then, if $\mathcal{P}$ is a periodic domain (with respect to some possibly different base-point), then*

$$n_{z_1}(\mathcal{P}) - n_{z_2}(\mathcal{P}) = \langle H(\mathcal{P}), \alpha_1^* \rangle,$$

where $H(\mathcal{P}) \in H_2(Y;\mathbb{Z})$ is the homology class belonging to the periodic domain.

Proof. For $i = 1, 2$, let γ_i be the gradient flow line passing through z_i (connecting the index three to the index zero critical point). Clearly, $n_{z_i}(\mathcal{P}) = \#\gamma_i \cap \mathcal{P}$. Now the difference $\gamma_1 - \gamma_2$ is a closed loop in Y, which is clearly homologous to a loop in U_0 which meets the attaching disk for α_1 in a single transverse point (and is disjoint from the attaching disks for α_i for $i \neq 1$). The formula then follows. $\square$

2.5. Spin^c **structures.** In dimension three, the set of Spin^c structures is identical with the space of "homology classes" of nowhere vanishing vector fields, where two vector fields are called *homologous* if they are homotopic away from a three-ball (see [32]).

Given the base-point $z \in \Sigma$, there is a natural map

$$s_z\colon \mathbb{T}_\alpha \cap \mathbb{T}_\beta \longrightarrow \mathrm{Spin}^c(Y)$$

obtained as follows. We think of the Heegaard decomposition as arising from a self-indexing Morse function with one local maximum and one local minimum

$$f\colon Y \longrightarrow [0,3].$$

Then, $U_0 = f^{-1}([0,3/2])$, $U_1 = f^{-1}([3/2,3])$, $\Sigma = f^{-1}(3/2)$. For appropriate choices, we can arrange that the $\boldsymbol{\alpha}$ are the intersections with ascending manifolds of the index one critical points with the middle level Σ; similarly, the $\boldsymbol{\beta}$ are the intersections of the descending manifolds of the index two critical points with Σ.

Hence, an intersection point $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ corresponds to a g-tuple of connecting trajectories from index one to index two critical points. Then, any base-point $z\in\Sigma$ corresponds to a connecting trajectory between the index zero and index three critical points. We modify the gradient vector field $\vec{\nabla} f$ in a small neighborhood of the connecting trajectories to obtain a nowhere vanishing vector field over Y. The corresponding Spin^c structure is independent of the choice of extension.

The assignment respects the difference ϵ, in the sense that for each $\mathbf{x},\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$,

$$s_z(\mathbf{x})-s_z(\mathbf{y})=-\mathrm{PD}[\epsilon(\mathbf{x},\mathbf{y})].$$

Moreover, the dependence on the basepoint is reflected in the following. Suppose that z_1 and z_2 are two basepoints, and there is a path z_t for $t=1$ to 2 disjoint from $\boldsymbol{\beta}$ and $\{\alpha_2,...,\alpha_g\}$, and whose intersection number $\#\alpha_1\cap z_t=1$. Then,

$$s_{z_2}(\mathbf{x})-s_{z_1}(\mathbf{x})=\alpha_1^*,$$

(see Lemma 2.12 of [23]).

There is a natural involution on the space of Spin^c structures which carries the homology class of the vector field v to the homology class of $-v$. We denote this involution by the map $\mathfrak{s}\mapsto\overline{\mathfrak{s}}$.

There is also a natural map

$$c_1\colon \mathrm{Spin}^c(Y)\longrightarrow H^2(Y;\mathbb{Z}),$$

the first Chern class. This is defined by $c_1(\mathfrak{s})=\mathfrak{s}-\overline{\mathfrak{s}}$. Equivalently, if $\mathfrak{s}$ is represented by the vector field v, then $c_1(\mathfrak{s})$ is the first Chern class of the orthogonal complement of v, thought of as an oriented real two-plane (hence complex line) bundle over Y. It is clear that $c_1(\overline{\mathfrak{s}})=-c_1(\mathfrak{s})$.

3. Analytical preliminaries

Ultimately, the Floer homologies defined in this paper are obtained from counting holomorphic Whitney disks in $\mathrm{Sym}^g(\Sigma)$, where the notion of "holomorphic" must be suitably tailored for our purposes. We recall the necessary constructions briefly here, and refer the interested reader to [23], where they are built up in detail. (These constructions are modifications of Floer's original construction in [8].)

If one chooses a complex structure j over the two-manifold Σ, there is a naturally induced holomorphic structure on the g-fold symmetric power $\mathrm{Sym}^g(\Sigma)$, denoted $\mathrm{Sym}^g(j)$. We can interpret the local multiplicities $n_z(\phi)$, when $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ has a holomorphic representatives, as intersection numbers between complex subvarieties of a complex manifold. Thus, in such cases, the local multiplicities are non-negative.

Unfortunately, $\mathrm{Sym}^g(j)$ may not be generic enough for certain analytic constructions (gluing theory) to work. To remedy this, we consider instead moduli spaces of nearly-holomorphic strips connecting $\mathbf{x}$ to $\mathbf{y}$ (c.f. [8]). Letting $\mathbb{D} = [0,1] \times i\mathbb{R} \subset \mathbb{C}$ be the strip in the complex plane, and fixing a path J_s of almost-complex structures over $\mathrm{Sym}^g(\Sigma)$, we let $\mathcal{M}_{J_s}(\mathbf{x}, \mathbf{y})$ be the set of maps satisfying the following conditions:

$$\mathcal{M}_{J_s}(\mathbf{x},\mathbf{y}) = \left\{ u\colon \mathbb{D} \longrightarrow \mathrm{Sym}^g(\Sigma) \left| \begin{array}{l} u(\{1\}\times\mathbb{R}) \subset \mathbb{T}_\alpha \\ u(\{0\}\times\mathbb{R}) \subset \mathbb{T}_\beta \\ \lim_{t\mapsto -\infty} u(s+it) = \mathbf{x} \\ \lim_{t\mapsto +\infty} u(s+it) = \mathbf{y} \\ \frac{du}{ds} + J(s)\frac{du}{dt} = 0 \end{array} \right. \right\}.$$

The translation action on $\mathbb{D}$ endows this moduli space with an $\mathbb{R}$ action. The space of *unparameterized J_s-holomorphic disks* is the quotient

$$\widehat{\mathcal{M}}_{J_s}(\phi) = \frac{\mathcal{M}_{J_s}(\phi)}{\mathbb{R}}.$$

The word "disk" is used, in view of the holomorphic identification of the strip with the unit disk in the complex plane with two boundary points removed (and maps in the moduli space extend across these points, in view of the asymptotic conditions).

In an appropriate sense, made precise in Section 3 of [23], the path J_s is chosen to be a small perturbation of the constant path $\mathrm{Sym}^g(j)$. It is chosen to agree with the constant path near the "diagonal" of $\mathrm{Sym}^g(\Sigma)$ (which is useful for technical purposes), and we are also free to choose it to be a constant near any finite union of subsets of the form $\{z_i\} \times \mathrm{Sym}^{g-1}(\Sigma)$, where $z_i \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$. In view of this, we have the following result, whose proof (and precise statement) can be found in Lemma 3.2 and Theorem 3.16 of [23]:

Theorem 3.1. *There is an allowable family J_s with the following*

- **non-negativity property:** *If $\mathcal{M}_{J_s}(\phi)$ is non-empty, then $\mathcal{D}(\phi) \geq 0$; with equality implying that the only holomorphic map is constant.*

Indeed, there is a connected, open set $\mathcal{J}$ of such paths, whose generic element $J_s \in \mathcal{J}$ is sufficiently generic in the following sense:

- *If $\mu(\phi) < 1$ and $\mathcal{D}(\phi) \neq 0$, then $\mathcal{M}_{J_s}(\phi)$ is empty.*
- *If $\mu(\phi) = 1$, then the quotient $\widehat{\mathcal{M}}_{J_s}(\phi)$ is a compact zero-manifold.*

- *If $\mu(\phi)=2$, then $\widehat{\mathcal{M}}_{J_s}(\phi)$ has a Gromov compactification.*

As explained in the proof of Theorem 4.3 of [23], if we count ends of moduli spaces $\widehat{\mathcal{M}}_{J_s}(\phi)$ where $\phi\in\pi_2(\mathbf{x},\mathbf{w})$ with $\mu(\phi)=2$ and $\mathbf{x}\neq\mathbf{w}$, there are, in principle, many ends (in the Gromov compactification) which could contribute to the count. However, under generic circumstances, the only ends are "broken flow-lines"

$$\coprod_{\phi_1*\phi_2=\phi}\widehat{\mathcal{M}}(\phi_1)\times\widehat{\mathcal{M}}(\phi_2),$$

where the ϕ_1 and ϕ_2 range over all possible homotopy classes in $\pi_2(\mathbf{x},\mathbf{y})$ and $\pi_2(\mathbf{y},\mathbf{w})$ respectively with $\mu(\phi_1)=\mu(\phi_2)=1$, $\phi_1*\phi_2=\phi$, and $\mathbf{y}$ here is any possible intersection points in the same equivalence class as $\mathbf{x}$. In the case where $\mathbf{x}=\mathbf{w}$, there could also be "boundary degenerations", corresponding to holomorphic disks passing through $\mathbf{x}$, whose boundary lies entirely inside $\mathbb{T}_\alpha$ or $\mathbb{T}_\beta$. In fact, the algebraic contribution of these boundary degenerations turns out to vanish (see Theorem 3.13 of [23]).

3.1. **Orientations.** Strictly speaking, to make sense of the above algebraic count, we need orientations on the moduli spaces of flows unless, of course, one is content to use a theory with $\mathbb{Z}/2\mathbb{Z}$ coefficients. Indeed, most of the theory presented here works with $\mathbb{Z}/2\mathbb{Z}$ coefficients, and such coefficients are sufficient for all the applications given in the introduction, but we describe the signed refinement with future applications in mind. This signed refinement uses systems of coherent orientations (compare [10]).

Fix $\mathbf{x},\mathbf{y}$ representing some fixed Spin^c structure $\mathfrak{t}$ over Y. We let $\mathcal{B}(\phi)$ denote the space of maps from the strip into $\mathrm{Sym}^g(\Sigma)$ representing ϕ which live in some suitable Sobolev space (see Section 3 of [23]) satisfying the same boundary-value conditions and asymptotic conditions as $\mathcal{M}_{J_s}(\mathbf{x},\mathbf{y})$ (only without the holomorphicity condition). The moduli spaces $\mathcal{M}(\phi)$ are cut from the space of maps $\mathcal{B}(\phi)$ by a non-linear Fredholm operator, whose linearization D_u can be naturally extended for all $u\in\mathcal{B}(\phi)$. Thus, the determinant of the linearization is a real line bundle over $\mathcal{B}(\phi)$, whose restriction to $\mathcal{M}(\phi)$ is the top exterior power of its tangent bundle.

Recall the following result (see Proposition 3.10 of [23]):

Proposition 3.2. *The determinant line bundle* $\det(D_u)$ *is trivial over the connected space* $\mathcal{B}(\phi)$.

Note that splicing gives an identification

$$\det(u_1)\wedge\det(u_2)\cong\det(u_1*u_2),$$

where $u_1\in\pi_2(\mathbf{x},\mathbf{y})$ and $u_2\in\pi_2(\mathbf{y},\mathbf{w})$ are a pair of maps.

Definition 3.3. *A* coherent system of orientations for $\mathfrak{s}$, $\mathfrak{o}$, *is a choice of non-vanishing sections* $\mathfrak{o}(\phi)$ *of the determinant line bundle over each* $\phi\in\pi_2(\mathbf{x},\mathbf{y})$ *for each* $\mathbf{x},\mathbf{y}\in\mathcal{S}$ *and each* $\phi\in\pi_2(\mathbf{x},\mathbf{y})$, *which are compatible with gluing in the sense that*

$$\mathfrak{o}(\phi_1)\wedge\mathfrak{o}(\phi_2)=\mathfrak{o}(\phi_1*\phi_2),$$

under the identification coming from splicing, and

$$\mathfrak{o}(u*S)=\mathfrak{o}(u),$$

under the identification coming from the canonical orientation for the moduli space of holomorphic spheres.

To construct these it is useful to have the following:

Definition 3.4. *Let* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ *be a Heegaard diagram representing* Y*, and let* $\mathfrak{t}$ *be a* Spin^c *structure for* Y*. A* complete set of paths for $\mathfrak{t}$ *is an enumeration* $\{\mathbf{x}_0, ..., \mathbf{x}_m\} = \mathcal{S}$ *of all the intersection points of* $\mathbb{T}_\alpha$ *with* $\mathbb{T}_\beta$ *representing* $\mathfrak{t}$*, and a collection of homotopy classes* $\theta_i \in \pi_2(\mathbf{x}_0, \mathbf{x}_i)$ *for* $i = 1, ..., m$ *with* $n_z(\theta_i) = 0$*.*

Fix periodic classes $\phi_1, ..., \phi_b \in \pi_2(\mathbf{x}, \mathbf{x})$ representing a basis for $H^1(Y;\mathbb{Z})$, and non-vanishing sections of the determinant line bundle for bundle for the homotopy classes $\theta_1, ..., \theta_m$ and $\phi_1, ..., \phi_b$. These data uniquely determine coherent system of orientations by splicing, since any homotopy class $\phi \in \phi_2(\mathbf{x}_i, \mathbf{x}_j)$ can be uniquely written as

$$\phi = a_1\phi_1 + ... + a_b\phi_b - \theta_i + \theta_j.$$

4. Definition of Floer homologies

In this section, we extend the definitions of the homology theories from [23] to cover the case where $b_1(Y) > 0$. In the first few subsections we define: the relative grading between intersection points, the necessary admissibility hypotheses required to make the boundary map make sense, and then the chain complexes themselves. Next we define some additional algebraic structure: the action of U and the action of $H_1(Y;\mathbb{Z})/\mathrm{Tors}$. Finally, we sketch the topological invariance of all these structures, modeling the discussion on [23]. A key technical point concerning the admissibility hypotheses (needed for isotopy invariance) is postponed to Section 5, and the proof of handleslide invariance is postponed to Section 6.4. In the final subsections, we return to several more algebraic constructions one can make with the homology groups.

4.1. Grading.

Fix a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, and let $\mathcal{S}$ be the set of intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $s_z(\mathbf{x}) = \mathfrak{s}$. As in the case for rational homology three-spheres (see [23]), we define a relative grading between $\mathbf{x}$ and $\mathbf{y}$ in $\mathcal{S}$ by the formula

$$\mathrm{gr}(\mathbf{x},\mathbf{y}) = \mu(\phi) - 2n_z(\phi), \tag{1}$$

where $\phi \in \pi_2(\mathbf{x},\mathbf{y})$ is an arbitrary element. In the case where $b_1(Y) = 0$, this grading is well-defined since $\mu(\phi + S) = \mu(\phi) + 2$ (c.f. Proposition 2.6 of [23]). In the present case, however, this grading is well-defined only modulo an indeterminacy

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

in view of the following result, which is proved in Subsection 5.3:

Theorem 4.1. *Fix a* Spin^c *structure* $\mathfrak{s} \in \mathrm{Spin}^c(Y)$*. Then for each* $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ *with* $s_z(\mathbf{x}) = \mathfrak{s}$*, and for each periodic class* $\psi \in \Pi_{\mathbf{x}}$ *we have*

$$\mu(\psi) = \langle c_1(\mathfrak{s}), H(\psi) \rangle,$$

where $H(\psi) \in H_2(Y;\mathbb{Z})$ *is the homology class corresponding to the periodic class* ψ*.*

(For practical purposes, we give an explicit formula for this quantity in terms of data over the Heegaard diagram in Section 8, c.f. Proposition 8.5.)

4.2. Admissibility.

To ensure compactness of the index one moduli spaces connecting intersection points, we will use only certain special kinds of Heegaard diagrams, as follows. It turns out that these conditions are somewhat different for the various theories.

Definition 4.2. *A pointed Heegaard diagram is called* strongly admissible for the Spin^c structure $\mathfrak{s}$ *if for every non-trivial periodic domain* $\mathcal{D}$ *with*

$$\langle c_1(\mathfrak{s}), H(\mathcal{D}) \rangle = 2n \geq 0,$$

$\mathcal{D}$ *has some coefficient* $> n$*. A pointed Heegaard diagram is called* weakly admissible for $\mathfrak{s}$ *if for each non-trivial periodic domain* $\mathcal{D}$ *with*

$$\langle c_1(\mathfrak{s}), H(\mathcal{D}) \rangle = 0,$$

$\mathcal{D}$ *has both positive and negative coefficients.*

Remark 4.3. *Note that for a* Spin^c *structure with* $c_1(\mathfrak{s})$ *torsion, the weak and strong admissibility conditions coincide. Also note that if a Heegaard diagram is strongly admissible for any torsion* Spin^c *structure then in fact it is weakly admissible for all* Spin^c *structures.*

We have the following geometric reformulation of the weak admissibility condition (for all Spin^c structures):

Lemma 4.4. *A Heegaard diagram is weakly admissible for all* Spin^c *structures if and only if* Σ *can be endowed with a volume form for which each periodic domain has total signed area equal to zero.*

Proof. The existence of such a volume form obviously implies weak admissibility, since each non-trivial domain has positive area.

Assume, conversely, that each non-trivial periodic domain has both positive and negative coefficients. By changing the volume form, we are free to make each domain in $\Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$ have arbitrary positive area. Thus, the claim is equivalent to a linear algebra statement. We say that a vector subspace $V \subset \mathbb{R}^m$ is *balanced* if each of its non-zero vectors has both positive and negative components. The claim, then, follows form the fact that a vector subspace of $\mathbb{R}^m$ which is balanced admits an orthogonal vector each of whose coefficients is positive.

This fact is true by induction on the dimension of the ambient vector space (and it is vacuously true for $m = 1$). Now, suppose V is a balanced subspace of $\mathbb{R}^m$, and let $\Pi_i\colon \mathbb{R}^m \longrightarrow \mathbb{R}^{m-1}$ denote the projection map $\Pi_i(x_1, ..., x_m) = (x_1, ..., \widehat{x_i}, ..., x_m)$. Either $\Pi_i(V)$ is also balanced, or V contains a vector v whose i^{th} component is $+1$, all other components are non-positive, and at least one of them is negative. In this latter case, we construct the required positive orthogonal vector as follows. Apply the induction hypothesis to find a vector $\xi = (\xi_1, ..., \xi_{i-1}, 0, \xi_{i+1}, ..., \xi_m)$ with $\xi_j > 0$ for $i \neq j$, which is orthogonal to $V \cap \mathbb{R}^{m-1}$. The required vector, then, is $\xi - \langle v, \xi\rangle e_i$.

If, on the other hand, all i of the vector spaces $\Pi_i(V)$ are balanced, then by induction we can find vectors $\xi = (0, \xi_2, ..., \xi_m)$ and $\eta = (\eta_1, 0, \eta_3, ..., \eta_m)$ with $\xi_i > 0$ for $i \neq 1$, and $\eta_i > 0$ for $i \neq 2$. Then, $\xi + \eta$ is our required vector. ☐

The following two lemmas are, ultimately, the reasons for introducing the admissibility hypotheses.

Lemma 4.5. *Suppose that* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ *is weakly admissible for the* Spin^c *structure* $\mathfrak{s}$*, and fix integers* $j, k \in \mathbb{Z}$*. Then, for each* $\mathbf{x}, \mathbf{y} \in \mathcal{S}$*, there are only finitely many* $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ *for which* $\mu(\phi) = j$*,* $n_z(\phi) = k$*, and* $\mathcal{D}(\phi) \geq 0$*.*

Proof. Fix some class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = j$. Then, in view of Theorem 4.1, any other $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ has the form

$$\psi = \phi + \mathcal{P}_{\mathbf{x}} - \frac{\langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle}{2} S,$$

where $\mathcal{P}_x$ is some periodic class, $\mathcal{P}$ its associated periodic domain, and S is the positive generator of $\pi_2(\mathrm{Sym}^g(\Sigma))$. If $n_z(\psi) = n_z(\phi)$, this forces $\mathcal{D}(\psi) = \mathcal{D}(\phi) + \mathcal{P}$ for some periodic domain whose associated homology class is annihilated by $c_1(\mathfrak{s})$; moreover, if $\mathcal{M}(\psi)$ is non-empty, all of its coefficients must be non-negative, i.e.

$$\mathcal{P} \geq -\mathcal{D}(\phi).$$

Thus, the lemma follows from the observation that for each integer ℓ, there are only finitely many periodic domains $\mathcal{P}$ in the set

$$Q = \left\{\mathcal{P} \in \Pi_{\mathbf{x}} | \langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle = 0, \mathcal{P} \geq \ell[\Sigma]\right\}.$$

We see this as follows. Let m denote the total number of domains (components in $\Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$). We can think of Q as lattice points in the m-dimensional vector space generated by the domains $\mathcal{D}_i$. Given $p \in Q$, written as $p = \sum a_i \mathcal{D}_i$, we let $\|p\|$ denote its naturally induced Euclidean norm

$$\|p\| = \sqrt{\sum_{i=1}^{m} |a_i|^2}.$$

If Q had infinitely many elements, we could find a sequence of $\{p_j\}_{j=1}^{\infty} \subset Q$ with $\|p_j\| \mapsto \infty$. In particular, the sequence $\frac{p_j}{\|p_j\|}$ has a subsequence which converges to a unit vector in the vector space of periodic domains with real coefficients which annihilate $c_1(\mathfrak{s})$. We write the vector as $p = \sum b_i \mathcal{D}_i$. Since the coefficients of p_j are bounded below, but the lengths of the p_j diverge, it follows that all the coefficients of p are non-negative. Of course, if the polytope in $\mathrm{Ann}(c_1(\mathfrak{s}))$ consisting of periodic domains with only non-negative multiplicities has a non-trivial real vector, then it must also have a non-trivial rational vector. After clearing denominators, we obtain a periodic domain (with integer coefficients) annihilating $c_1(\mathfrak{s})$, with only non-negative coefficients. This contradicts the hypothesis of weak admissibility. □

Lemma 4.6. *For a strongly admissible pointed Heegaard diagram, and an integer j, there are only finitely many $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = j$ and $\mathcal{D}(\phi) \geq 0$.*

Proof. Fix a $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = j$. Then, as in the previous lemma, any other class $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ can be written as

$$\psi = \phi - \mathcal{P}_{\mathbf{x}} + \frac{\langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle}{2} S.$$

Thus, (as $\mathcal{D}(S) = \Sigma$),

$$-\mathcal{P} + \frac{\langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle}{2}[\Sigma] \geq -\mathcal{D}(\phi).$$

The finiteness then follows from the fact that for each $\ell \in \mathbb{Z}$, there are only finitely many periodic domains $\mathcal{P}$ for which

$$-\mathcal{P} + \frac{\langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle}{2}[\Sigma] \geq \ell[\Sigma].$$

This follows as in the proof of Lemma 4.5: an infinite number of such periodic domains would give rise to a a real periodic domain $\mathcal{P}$ for which

$$-\mathcal{P} + \frac{\langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle}{2}[\Sigma] \geq 0,$$

from which it is easy to see that there must be an integral periodic domain with the same property. But such a periodic domain would violate the strong admissibility hypothesis. □

4.3. **The chain complex.** To define the chain complexes $CF^\infty(Y,\mathfrak{s})$, $CF^-(Y,\mathfrak{s})$, $CF^+(Y,\mathfrak{s})$, and $\widehat{CF}(Y,\mathfrak{s})$, we need a pointed Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$. Indeed, when defining $CF^\infty(Y,\mathfrak{s})$ and $CF^-(Y,\mathfrak{s})$, we require the pointed Heegaard diagram to be strongly admissible for $\mathfrak{s}$, while to define $CF^+(Y,\mathfrak{s})$ and $\widehat{CF}(Y,\mathfrak{s})$, we require it to be only weakly admissible. Existence of such Heegaard diagrams will be established in Lemma 5.4 in the next section. Except in the case where the chain complex has $\mathbb{Z}/2\mathbb{Z}$ coefficients, we also need a coherent system of orientations $\mathfrak{o}$ for $\mathfrak{s}$ (in the sense of Definition 3.3).

Let $\mathcal{S}\subset\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ be the set of points $\mathbf{x}$ with $s_z(\mathbf{x})=\mathfrak{s}$. We define the relative grading function gr over $\mathcal{S}$ as in Equation (1). Let $CF^\infty(Y,\mathfrak{s})$ be the relatively $\mathbb{Z}/\delta(\mathfrak{s})$-graded chain complex (note that this is relatively $\mathbb{Z}$-graded when $c_1(\mathfrak{s})$ is torsion) which is freely generated by pairs $[\mathbf{x},i]\in\mathcal{S}\times\mathbb{Z}$, with relative grading

$$\mathrm{gr}([\mathbf{x},i],[\mathbf{y},j])=\mathrm{gr}(\mathbf{x},\mathbf{y})+2(i-j).$$

We define

$$c_k(\mathbf{x},\mathbf{y})=\sum_{\left\{\phi\in\pi_2(\mathbf{x},\mathbf{y})\middle|\mu(\phi)=1,n_z(\phi)=k\right\}}\#\widehat{\mathcal{M}}(\phi),$$

where the signed count is defined using the coherent system of orientations (except when we wish to use $\mathbb{Z}/2\mathbb{Z}$ coefficients, in which case $\#\mathcal{M}(\phi)$ is the number of points modulo two). It follows from Lemma 4.5 (and the weak admissibility hypothesis) together with and Theorem 3.1 that c_k is a finite sum. We then define

$$\partial^\infty[\mathbf{x},i]=\sum_{\{\mathbf{y}\in\mathcal{S}\}}\sum_{k=-\infty}^{\infty}c_k(\mathbf{x},\mathbf{y})[\mathbf{y},i-k].$$

It follows from Lemma 4.6 (and the strong admissibility hypothesis) that for each $[\mathbf{x},i]$, this is a finite sum.

The same arguments which apply in the $b_1(Y)=0$ case apply to give the following (see Theorem 4.3 of [23]):

Theorem 4.7. *The pair* $(CF^\infty,\partial^\infty)$ *is a chain complex, i.e.* $\partial^\infty\circ\partial^\infty=0$.

In view of the non-negativity property (Theorem 3.1), the subgroup $CF^-(Y,\mathfrak{s})\subset CF^\infty(Y,\mathfrak{s})$ generated by pairs $[\mathbf{x},i]$ with $i<0$ is a subcomplex. Its quotient complex is denoted $CF^+(Y,\mathfrak{s})$, which we can think of as being generated by $[\mathbf{x},i]$ with $i\geq 0$.

Similarly, we can define $\widehat{CF}$ to be the subcomplex of CF^+ generated elements of the form $[\mathbf{x},0]$, $\mathbf{x}\in\mathcal{S}$. Note that the chain complexes CF^+ and $\widehat{CF}$ are already well-defined under the weak admissibility hypothesis, according to Lemma 4.5.

4.4. **Additional algebra: the $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ and U-actions.** As in the case where $b_1(Y)=0$, the groups $HF^\infty(Y,\mathfrak{s})$, $HF^+(Y,\mathfrak{s})$, and $HF^-(Y,\mathfrak{s})$ come equipped with a natural action by a map U which lowers degree by two. It is induced by the map

$$U\colon CF^\infty(Y,\mathfrak{s}) \longrightarrow CF^\infty(Y,\mathfrak{s})$$

given by $U[\mathbf{x},i]=[\mathbf{x},i-1]$.

When $b_1(Y)>0$, there is a new algebraic object: an action of $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ on $HF^\infty(Y,\mathfrak{s})$, $HF^+(Y,\mathfrak{s})$, $HF^-(Y,\mathfrak{s})$, and $\widehat{HF}(Y,\mathfrak{s})$. Recall from the proof of Proposition 2.3 that the choice of basepoint gives an isomorphism

$$H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta);\mathbb{Z}) \cong \mathrm{Hom}(\pi_1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)),\mathbb{Z}) \cong \mathbb{Z}\oplus \mathrm{Hom}(H^1(Y,\mathbb{Z}),\mathbb{Z}).$$

Proposition 4.8. *There is a natural action of $H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta))$ lowering degree by one on $HF^\infty(Y,\mathfrak{s})$, $HF^+(Y,\mathfrak{s})$, $HF^-(Y,\mathfrak{s})$ and $\widehat{HF}(Y,\mathfrak{s})$. Furthermore, this induces actions of the exterior algebra $\Lambda^*(H_1(Y;\mathbb{Z})/\mathrm{Tors})\subset \Lambda^*(H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta));\mathbb{Z})$ on each group.*

To define this action, let $\zeta\in Z^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta);\mathbb{Z})$ be a one-cocycle in the space of paths connecting $\mathbb{T}_\alpha$ to $\mathbb{T}_\beta$. We define a map

$$A_\zeta\colon CF^\infty(Y,\mathfrak{s}) \longrightarrow CF^\infty(Y,\mathfrak{s})$$

which lowers degree by one, by the formula

$$A_\zeta([\mathbf{x},i]) = \sum_{\mathbf{y}\in\mathcal{S}} \sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y})|\mu(\phi)=1\}} \zeta(\phi)\cdot\left(\#\widehat{\mathcal{M}}(\phi)\right)[\mathbf{y}-n_z(\phi)].$$

By $\zeta(\phi)$, we mean the following. Choose any representative u for the homotopy class ϕ, and view it as an arc in $\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)$ which connects the constant paths $\mathbf{x}$ and $\mathbf{y}$. If we choose a different representative for the same homotopy class, then the corresponding paths will be homotopic (as arcs in $\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)$ connecting $\mathbf{x}$ to $\mathbf{y}$), so the evaluation of ζ is independent of the particular choice (since ζ is a cocycle).

We turn to the proof of Proposition 4.8, which we break into several lemmas.

Lemma 4.9. *A_ζ is a chain map.*

Proof. This is a variant on the usual proof that $\partial^2=0$. Suppose that $\phi\in\pi_2(\mathbf{x},\mathbf{w})$ satisfies $\mu(\phi)=2$, and let $k=n_z(\phi)$. Then, since $\zeta(\phi_1*\phi_2)=\zeta(\phi_1)+\zeta(\phi_2)$ (since ζ is a cocycle), we get that

$$\begin{aligned} 0 &= \zeta(\phi)\cdot\left(\#(\text{ends of }\widehat{\mathcal{M}}(\phi))\right) \\ &= \sum_{\{\phi_1,\phi_2|\phi=\phi_1*\phi_1,\mu(\phi_1)=\mu(\phi_2)=1\}} \left(\zeta(\phi_1)+\zeta(\phi_2)\right)\left(\#\widehat{\mathcal{M}}(\phi_1)\right)\cdot\left(\#\widehat{\mathcal{M}}(\phi_2)\right). \end{aligned}$$

(Note that boundary degenerations do not contribute to the above sum, as in the proof that $\partial^2=0$.) Summing over all $\phi\in\pi_2(\mathbf{x},\mathbf{w})$ with $n_z(\phi)=k$ and $\mu(\phi)=2$, we get the $[\mathbf{w},i-k]$-coefficient of $(\partial\circ A_\zeta + A_\zeta\circ\partial)[\mathbf{x},i]$. $\square$

Lemma 4.10. *If ζ is a coboundary, then A_ζ is chain homotopic to zero.*

Proof. If ζ is a coboundary, then there is a zero-cochain B (a possibly discontinuous map from $\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)$ to $\mathbb{Z}$) with the property that if γ is an arc in $\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)$ (a one-simplex), then $\zeta(\gamma)=B(\gamma(0))-B(\gamma(1))$. Let
$$H([\mathbf{x},i])=B(\mathbf{x})[\mathbf{x},i],$$
where the evaluation of B on $\mathbf{x}$ is performed by viewing the latter as a constant path from $\mathbb{T}_\alpha$ to $\mathbb{T}_\beta$. Then, it follows from the definitions that
$$A_\zeta=\partial\circ H+H\circ\partial.$$
□

Proof of Proposition 4.8. Together, Lemmas 4.9 and 4.10 show that the A_ζ descends to a well-defined action of $H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta))$ on HF^∞. To see that the action descends to the exterior algebra, we must verify that the composite $A_\zeta\circ A_\zeta=0$ in homology.

To see this, we think of A_ζ using codimension one constraints. Specifically, we begin with a map $f\colon\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)\longrightarrow S^1$ representing ζ. Given a generic point $p\in S^1$, and we let $V=f^{-1}(p)$, so that the action of ζ is given by
$$A_\zeta([\mathbf{x},i])=\sum_{\mathbf{y}}\sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y})\big|\mu(\phi)=1\}}a(\zeta,\phi)[\mathbf{y},i-n_z(\phi)],$$
where
$$a(\zeta,\phi)=\#\{u\in\mathcal{M}(\phi)\big|u(\{0\}\times[0,1])\in V\}.$$
Fix a homotopy class $\phi\in\pi_2(\mathbf{x},\mathbf{w})$ with $\mu(\phi)=2$. We consider the one-manifold
$$\Xi=\Big\{s\in[0,\infty),u\in\mathcal{M}(\phi)\Big|u(\{s\}\times[0,1])\in V,\ u(\{-s\}\times[0,1])\in V'\Big\}.$$
where V, V' are the preimages of p and p' under f. Choosing $p\neq p'$, the one-manifold Ξ has no boundary at $s=0$. The ends as $s\mapsto\infty$ (disregarding boundary degenerations, which do not contribute algebraically), are modeled on
$$\Big\{u_1\in\mathcal{M}(\phi_1)\Big|u_1(\{0\}\times[0,1])\in V\Big\}\times\Big\{u_2\in\mathcal{M}(\phi_2)\Big|u_2(\{0\}\times[0,1])\in V'\Big\},$$
where $\phi=\phi_1*\phi_2$. On the one hand, the number of points, counted with sign, must vanish; on the other hand, it is the $[\mathbf{w},i-n_z(\phi)]$ coefficient of $A_\zeta\circ A_\zeta$. It follows that the action of $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ on $HF^\infty(Y,\mathfrak{s})$ descends to an action of the exterior algebra

The chain map A_ζ, and the chain homotopy from Lemma 4.10 preserve $CF^-(Y,\mathfrak{s})$, so it is follows that A_ζ induces actions on HF^+ and HF^-. The action on $\widehat{HF}$ is defined in an analogous manner, as well. □

Note that in the statement of Theorem 1.1 in the introduction, we suppressed the action of the $\mathbb{Z}$ summand in $H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta);\mathbb{Z})$, belonging to the one-cycle associating to ϕ its intersection number $n_z(\phi)$. The reason for this is given by the following:

Proposition 4.11. *The induced action of $\mathbb{Z}\subset H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta))$ induced by the base-point $z\in\Sigma$ is trivial on $HF^\infty(Y,\mathfrak{s})$, $HF^+(Y,\mathfrak{s})$, $HF^-(Y,\mathfrak{s})$, and $\widehat{HF}(Y,\mathfrak{s})$.*

Proof. The generator ζ of the $\mathbb{Z}$ summand acts by

$$A_\zeta[\mathbf{x}, i] = \sum_{\mathbf{y}} \sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y}) | \mu(\phi)=1\}} (\#\mathcal{M}(\phi)) \cdot (n_z(\phi))\, [\mathbf{y}, i - n_z(\phi)].$$

This action is trivial on $\widehat{CF}(Y, \mathfrak{s})$. For the other theories, a null-homotopy is given by $H([\mathbf{x}, i]) = i \cdot [\mathbf{x}, i]$. □

Remark 4.12. *A geometric realization of the action of* $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ *can be given as follows. Let* $\gamma \in \Sigma$ *be a curve which misses the intersection points between the* α_i *and* β_j*, and let* $[\gamma]$ *be its induced homology class in* $H_1(Y;\mathbb{Z})$*. Then,*

$$A_{[\gamma]}([\mathbf{x}, i]) = \sum_{\mathbf{y}} \sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y}) | \mu(\phi)=1\}} a(\gamma, \phi)[\mathbf{y}, j - n_z(\phi)],$$

where

$$a(\gamma, \phi) = \#\{u \in \mathcal{M}(\phi) \big| u(0 \times 1) \in (\gamma \times \mathrm{Sym}^{g-1}(\Sigma)) \cap \mathbb{T}_\alpha\}$$

or, equivalently, $a(\gamma, \phi)$ *is the product of* $\#\widehat{\mathcal{M}}(\phi)$ *with the intersection number in* $\mathbb{T}_\alpha$ *between the codimension one submanifold* $\big(\gamma \cap \mathrm{Sym}^{g-1}(\Sigma)\big) \cap \mathbb{T}_\alpha$ *and the curve in* $\mathbb{T}_\alpha$ *obtained by restricting* u *to* $u(\mathbb{R} \times \{1\})$*, where* u *is any representative of* ϕ*.*

4.5. **Topological invariance: the statement.** The complexes CF^∞, CF^-, CF^+, and $\widehat{CF}$ depend on the pointed Heegaard diagram, the equivalence class of intersection points of the corresponding tori, and a system of coherent orientations.

Theorem 4.13. *The homology groups* $HF^\infty(Y, \mathfrak{s})$ $HF^-(Y, \mathfrak{s})$*,* $HF^+(Y, \mathfrak{s})$ *and* $\widehat{HF}(Y, \mathfrak{s})$ *(as modules over* $\wedge^*(H_1(Y;\mathbb{Z})/\mathrm{Tors})$ *and* $\mathbb{Z}[U]$*, where appropriate) are topological invariants. More precisely, if two* $\mathfrak{s}$*-strongly admissible pointed Heegaard diagrams represent the same oriented three-manifold, then the corresponding relatively graded modules are isomorphic.*

In [26], we prove a functorial version of the above result. In the case of $\widehat{HF}$ and HF^+ we have a stronger statement:

Theorem 4.14. *Let* Y *be an oriented three-manifold equipped with a* Spin^c *structure* $\mathfrak{s}$*. Any two* $\mathfrak{s}$*-weakly admissible pointed Heegaard diagrams induce isomorphic relatively graded groups* $\widehat{HF}(Y, \mathfrak{s})$ *and* $HF^+(Y, \mathfrak{s})$*.*

Strictly speaking, there are $2^{b_1(Y)}$ different candidates for these groups corresponding to different systems of coherent orientations.

Theorem 4.13 is a modification of the corresponding main theorem in [23]. The theorem is divided into four parts:

(1) independence of the complex structure j over Σ and associated path J_s of almost-complex structures used in the definition of the moduli spaces
(2) invariance under pointed isotopies of the Heegaard diagram
(3) invariance under pointed handleslides
(4) invariance under stabilization of the Heegaard diagram.

Special care must be taken to handle the admissibility hypotheses when $b_1(Y) > 0$. To this end, we have the following result, whose proof is relegated to Section 5.

Theorem 4.15. *Given any* Spin^c *structure* $\mathfrak{s} \in \mathrm{Spin}^c(Y)$*, there is a* $\mathfrak{s}$*-strongly admissible pointed Heegaard diagram for* Y*. Moreover, any two* $\mathfrak{s}$*-strongly admissible pointed Heegaard diagrams can be connected by Heegaard moves through* $\mathfrak{s}$*-strongly admissible pointed Heegaard diagrams.*

We address the four points presented above in the next few subsections, after which we return to the issue of orientations.

4.6. **Independence of complex structures.** To address the first point in topological invariance, we appeal to the proof of Theorem 4.9 of [23], which proves the analogous result when $b_1(Y)=0$.

Recall that first one fixes the complex structure j over Σ, and investigates how the chain complex depends on the variation of the family J_s. To this end, we consider one-parameter families of paths of almost complex structures $J_s(t)$, defining a chain map

$$\Phi^\infty_{J_{s,t}}\colon (CF^\infty(\mathbb{T}_\alpha,\mathbb{T}_\beta),\partial^\infty_{J_s(1)}) \longrightarrow (CF^\infty(\mathbb{T}_\alpha,\mathbb{T}_\beta),\partial^\infty_{J_s(0)})$$

by

$$\Phi^\infty_{J_{s,t}}[\mathbf{x},i] = \sum_{\mathbf{y}} \sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y})\big|\mu(\phi)=0\}} \#\left(\mathcal{M}_{J_{s,t}}(\phi)\right)[\mathbf{y}, i-n_z(\phi)],$$

where $\mathcal{M}_{J_{s,t}}(\phi)$ consists of maps

$$\left\{ u\colon \mathbb{D}\cong[0,1]\times\mathbb{R}\longrightarrow \mathrm{Sym}^g(\Sigma) \left| \begin{array}{l} \frac{du}{ds}+J(s,t)\frac{du}{dt}=0, \\ u(\{1\}\times\mathbb{R})\subset\mathbb{T}_\alpha, u(\{0\}\times\mathbb{R})\subset\mathbb{T}_\beta \\ \lim_{t\mapsto-\infty} u(s+it)=\mathbf{x}, \lim_{t\mapsto+\infty} u(s+it)=\mathbf{y} \end{array}\right.\right\},$$

which represent the homotopy class ϕ. Choosing the $J_s(t)$ to be constant at $\{z_i\}\times\mathrm{Sym}^{g-1}(\Sigma)$ (with one z_i in each domain of $\Sigma-\alpha_1-\ldots-\alpha_g-\beta_1-\ldots-\beta_g$), we have the non-negativity statement on $\mathcal{M}_{J_{s,t}}(\phi)$ from Theorem 3.1. This, together with admissibility and Lemma 4.6, ensure that the sum defining $\Phi^\infty_{J_{s,t}}$ is a finite sum.

The proof that it gives a chain map, with an inverse up to chain homotopy, works the same as in the case where $b_1(Y)=0$. To verify the invariance of the groups, thought of as $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ modules, it suffices to show that the map induced on homology by $\Phi^\infty_{J_{s,t}}$ is equivariant with respect to this action.

Lemma 4.16. *For any* $\zeta\in H_1(Y,\mathbb{Z})/\mathrm{Tors}$*,*

$$A_\zeta\circ(\Phi^\infty_{J_{z,t}}) = (\Phi^\infty_{J_{z,t}})\circ A_\zeta$$

as a map from $H_*(CF^\infty(\mathbb{T}_\alpha,\mathbb{T}_\beta),\partial^\infty_{J_s(1)}) \longrightarrow H_*(CF^\infty(\mathbb{T}_\alpha,\mathbb{T}_\beta),\partial^\infty_{J_s(0)})$

Proof. Let V be a codimension one constraint in $\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta)$ representing the class $\zeta\in H^1(\Omega(\mathbb{T}_\alpha,\mathbb{T}_\beta);\mathbb{Z})$, chosen to miss all the constant paths (corresponding to the intersection points $\mathbb{T}_\alpha\cap\mathbb{T}_\beta$).

Consider the map

$$h\colon CF^\infty(\mathbb{T}_\alpha,\mathbb{T}_\beta)\longrightarrow CF^\infty(\mathbb{T}_\alpha,\mathbb{T}_\beta),$$

defined by

$$h([\mathbf{x},i]) = \sum_{\mathbf{y}} \sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y})|\mu(\phi)=0\}} \#\{(r,u)\in\mathcal{M}_{J_{s,t}}(\phi)|u([0,1]\times r)\in V\}[\mathbf{y},i-n_z(\phi)].$$

We claim that

$$A_\zeta\circ\Phi^\infty_{J_{s,t}} - \Phi^\infty_{J_{s,t}}\circ A_\zeta = \partial_{J_s(0)}\circ h - h\circ\partial_{J_s(1)}. \tag{2}$$

This follows by considering the ends of the one-dimensional moduli spaces

$$\Xi = \{(r,u)\in\mathbb{R}\times\mathcal{M}_{J_{s,t}}(\psi)\big|u([0,1]\times\{r\})\in V\}$$

where $\mu(\psi)=1$. The ends where $r\mapsto\pm\infty$ correspond to the commutator of A_ζ and Φ^∞, while the ends where the maps $u\in\mathcal{M}_{J_{s,t}}$ bubble off correspond to the commutator of h with the corresponding boundary maps.

Equation (2), of course, says that A_ζ commutes with $\Phi^\infty_{J_{s,t}}$, on the level of homology. □

With this lemma in hand, the proof of Theorem 4.9 of [23] applies to prove the independence of the homology groups from the complex structures used. Incidentally, Lemma 4.16 is stated for CF^∞, but the same argument applies for CF^+, CF^-, and $\widehat{CF}$ as well, provided that the Heegaard diagram is strongly admissible. We return to the case of weakly admissible Heegaard diagrams (Theorem 4.14) in Subsection 4.10.

4.7. **Isotopy invariance.** Armed with Theorem 4.15, we approach isotopy invariance in the same manner as in [23]: an isotopy is thought of as composition of sequences of pair creations and annihilations, and changes of the complex structure Σ (the latter of which has already been handled).

Each pair creation is, once again, thought of as an exact Hamiltonian isotopy of, say, α_1, inducing a chain map with moving boundary conditions. In view of Theorem 4.15, we need only consider the case where the Heegaard diagrams are admissible before and after the pair creation.

Specifically, let $\pi_2^{\Psi_t}(\mathbf{x},\mathbf{y})$ denote the space of maps $u\colon\mathbb{D}\longrightarrow\mathrm{Sym}^g(\Sigma)$ satisfying

$$\left\{u\colon\mathbb{D}\longrightarrow\mathrm{Sym}^g(\Sigma)\left|\begin{array}{l}u(1+it)\in\Psi_t(\mathbb{T}_\alpha),\forall t\in\mathbb{R},\\ u(0+it)\in\mathbb{T}_\beta,\forall t\in\mathbb{R},\\ \lim_{t\mapsto-\infty}u(s+it)=\mathbf{x},\\ \lim_{t\mapsto+\infty}u(s+it)=\mathbf{y}\end{array}\right.,\right\}$$

where Ψ_t is an isotopy supported in a small neighborhood of $\mathbb{T}_\alpha\subset\mathrm{Sym}^g(\Sigma)$, with

$$\Psi_t(\mathbb{T}_\alpha)=\psi_t(\alpha_1)\times...\times\alpha_g,$$

and which is constant outside $t\in[0,1]$. We call such maps Whitney disks with dynamic boundary conditions. We have moduli spaces $\mathcal{M}^{\Psi_t}(\phi)$ of maps representing $\phi\in\pi_2^{\Psi_t}(\mathbf{x},\mathbf{y})$, which are also J_s-holomorphic. We then define the chain map

$$\Gamma^\infty_{\Psi_t}([\mathbf{x},i]) = \sum_{\{\mathbf{y}\in\mathfrak{S}'\}}\sum_{\{\phi\in\pi_2^{\Psi_t}(\mathbf{x},\mathbf{y})|\mu(\phi)=0\}} \#\left(\mathcal{M}^{\Psi_t}(\mathbf{x},\mathbf{y})\right)\cdot[\mathbf{y},i-n_z(\phi)],$$

where $\mu(\phi)$ is the expected dimension of the moduli space $\mathcal{M}_{\Psi_t}(\phi)$. We must verify that this sum is, in fact, a finite sum, for each given $[\mathbf{x},i]$.

In a single pair creation, the domains for the diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, and $(\Sigma, \{\psi(\alpha_1), \alpha_2, ..., \alpha_g\}, \boldsymbol{\beta})$ do not coincide: the latter has a new domain. Correspondingly a homotopy class $\pi_2^{\Psi}(\mathbf{x}, \mathbf{y})$ does not have a well-defined multiplicity at this new domain, since the $\Psi_t(\mathbb{T}_\alpha)$ crosses the subvariety $\{w\} \times \mathrm{Sym}^{g-1}(\Sigma) \subset \mathrm{Sym}^g(\Sigma)$, where w is any point in this new domain.

However, the multiplicities at the other domains are still well-defined; i.e. if $\mathcal{D}_i$ is any domain which exists before the pair-creation, and $w_i \in \mathcal{D}_i$ is a point in the interior of this domain, then the intersection number $u \cap \left(\{w_i\} \times \mathrm{Sym}^{g-1}(\Sigma)\right)$ (where u is any map representing $\phi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$) is independent of the choice of representative u and the point w_i (we choose the isotopy Ψ_t to be constant near $\{w_i\} \times \mathrm{Sym}^{g-1}(\Sigma)$). We call this collection of multiplicities the domain of ϕ.

Lemma 4.17. *Fix $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a strongly $\mathfrak{s}$-admissible pointed Heegaard diagram, and an isotopy Ψ_t is an isotopy as above. Then, for each pair of integers j, and for each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\mathbf{y} \in \mathbb{T}'_\alpha \cap \mathbb{T}_\beta$, there are only finitely many homotopy classes $\psi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ which support J_s-holomorphic representatives.*

Proof. Let $w_1, ..., w_m$ be points contained in the interiors of the domains before the pair-creation, and w_{m+1} be a point in the new domain. Let $\mathbb{T}'_\alpha$ be the torus $\psi_1(\alpha) \times ... \times \alpha_g$. As before, if $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the space of homotopy classes of Whitney disks for $\mathbb{T}_\alpha$, $\mathbb{T}_\beta$; if $\mathbf{y}' \in \mathbb{T}'_\alpha \cap \mathbb{T}_\beta$ we let $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$ denote the homotopy classes with moving boundary conditions defined above, and we let $\pi'_2(\mathbf{x}, \mathbf{y})$ denote the homotopy classes of Whitney disks for the pair $\mathbb{T}'_\alpha$ and $\mathbb{T}_\beta$, now thinking of $\mathbf{x}$ and $\mathbf{y}$ as intersections between those tori.

Fix $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $\mathbf{y} \in \mathbb{T}'_\alpha \cap \mathbb{T}_\beta$. It is easy to see that each homotopy class $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$ has a representative $u(s, t)$ which is constant for $t \leq 1$. As such, u can be thought of as representing a class $\pi'_2(\mathbf{x}, \mathbf{y})$. Indeed, this induces a one-to-one correspondence $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}) \cong \pi'_2(\mathbf{x}, \mathbf{y})$. In a similar manner, if $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we have identifications $\pi'_2(\mathbf{x}, \mathbf{x}) \cong \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{x}) \cong \pi_2(\mathbf{x}, \mathbf{x})$, which preserve all the local multiplicities n_{w_j} for all $j = 1, ..., m$.

Let $\{\psi_i\}$ be a sequence of homotopy classes in $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$ which support holomorphic representatives, and have a fixed Maslov index. Next, fix $\phi_0 \in \pi_2(\mathbf{y}, \mathbf{x})$. Since the ψ_i all support holomorphic representatives, the local multiplicities at the w_j for $j = 1, ..., m$ are non-negative; it follows that for $j = 1, ..., m$, $n_{w_j}(\psi_i * \phi_0) \geq n_{w_j}(\phi_0)$. But $\psi_i * \phi_0$ is a homotopy class connecting $\mathbf{x}$ with $\mathbf{x}$, which are intersection points which existed before the pair creation, so we can consider the corresponding element of $\pi_2(\mathbf{x}, \mathbf{x})$. From the above observations, the multiplicities at all w_i for $i = 1, ..., m$ are are bounded below, and the Maslov index is fixed, so there can be only finitely many such homotopy classes, according to Lemma 4.6. It follows that there are only finitely many distinct homotopy classes amongst the $\psi_i \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$. □

With these remarks in place, it follows that $\Gamma^\infty_{\Psi_t}$ is still finite sum (for each fixed $[\mathbf{x}, i]$) and the proof from [23] applies. Note that the Lemma 4.17 also holds for isotopies obtained by juxtaposing Ψ_t with Ψ_{1-t}.

Establishing its $H_1(Y; \mathbb{Z})/\mathrm{Tors}$-equivariance of the map Φ^∞ follows as in the proof of Lemma 4.16 above.

4.8. **Handleslides and Stabilizations.** Stabilization invariance is a direct consequence of the results from the gluing results of Section 6 of [23] (specifically, in that reference, use Theorem 6.1 for $\widehat{HF}$ and Theorem 6.2 for the others). The stabilization map is easily seen to be equivaraint under the action by $H_1(Y)/\mathrm{Tors}$: to see this, we represent the action by a codimension one constraint, and apply the gluing results with a constraint.

Handleslide invariance also follows the arguments from [23]. However, we will be at liberty to place the discussion there on a less *ad hoc* footing, once we develop the holomorphic triangle machinery a bit more. Hence, we will return to this point in Subsection 6.4.

4.9. **Orientations.** When the homology is calculated with $\mathbb{Z}$ coefficients, there is an additional choice in the definition: and that is the choice of coherent orientations. We investigate the dependence of the groups on this choice. To this effect, we have the following:

Let $\mathfrak{o}$ and $\mathfrak{o}'$ be a pair of systems of coherent orientations for $\mathfrak{t}$. Then, we define their difference $\delta = \delta(\mathfrak{o}, \mathfrak{o}') \in \mathrm{Hom}(H^1(Y;\mathbb{Z}), \mathbb{Z}/2\mathbb{Z})$ as follows. Let $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ be the periodic class representing some homology class $H \in H^1(Y;\mathbb{Z})$. Then, the section $\mathfrak{o}$ of the determinant line bundle over the component specified by ϕ is either a positive multiple of $\mathfrak{o}'$, in which case we let $\delta(H) = 0$, or it is a negative multiple of $\mathfrak{o}'$, in which case we let $\delta(H) = 1$. We say that two systems of coherent orientations are *equivalent* if their difference δ vanishes.

Now, if $\mathfrak{o}$ and $\mathfrak{o}'$ are equivalent systems in this sense, then it is easy to see that the chain complexes $CF^\infty(Y, \mathfrak{s}, \mathfrak{o})$ and $CF^\infty(Y, \mathfrak{s}, \mathfrak{o}')$ are chain homotopic. Moreover, it is easy to see that stabilizations and isotopies induce maps between coherent systems of orientations (well-defined up to this difference δ). We return to the case of handleslides in Subsection 6.4.

From this, it is clear that there are a priori $2^{b_1(Y)}$ different possible homologies, depending on the equivalence class of the orientation system. However, we usually drop the coherent system from the notation, and indeed in Section 11, we will give a canonical choice from these $2^{b_1(Y)}$ possibilities.

4.10. **Weakly admissible Heegaard diagrams.** Theorem 4.14 now follows from Thereom 4.13, together with the following observation.

Proposition 4.18. *Fix a Heegaard diagram for Y, and a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Then, for fixed complex structure j over Σ, the groups $\widehat{HF}(Y, \mathfrak{s})$ and $HF^+(Y, \mathfrak{s})$ are isomorphic to the corresponding groups calculated using some $\mathfrak{s}$-strongly admissible Heegaard diagram.*

Proof. It will be shown in Section 5 that a weakly admissible Heegaard diagram can be made strongly admissible by a large exact Hamiltonian isotopy through weakly admissible diagrams (c.f Lemma 5.7). Noting that the analogue of Lemma 4.17 also holds in the weakly admissible context (where $\mu(\psi)$ and $n_z(\psi)$ are both fixed), we construct chain homotopy equivalences $\Gamma^+_{\Psi_t}$ and $\widehat{\Gamma}_{\Psi_t}$ as in Subsection 4.7 to see that the groups are isomorphic. □

Proof of Theorem 4.14. This is now an immediate consequence of the above proposition, and Theorem 4.13. □

4.11. **Twisted coefficients.** There are variants of the above theory which work with a "twisted coefficient system". In particular, we can construct a homology theory with coefficients in the group-ring $\mathbb{Z}[H^1(Y;\mathbb{Z})]$. We write these homology theories $\underline{HF}(Y)$.

First, we need a surjective, additive map:

$$A\colon \pi_2(\mathbf{x},\mathbf{y}) \longrightarrow H^1(Y;\mathbb{Z}),$$

which is trivial on the action of $\pi_2(\mathrm{Sym}^g(\Sigma))$.

We can construct such a map as follows. A complete set of paths for $\mathfrak{t}$ in the sense of Definition 3.4 gives rise to identifications for any i, j:

$$\pi_2(\mathbf{x}_i,\mathbf{x}_j) \cong \pi_2(\mathbf{x}_0,\mathbf{x}_0),$$

by

$$\phi_i * \pi_2(\mathbf{x}_i,\mathbf{x}_j) \cong \pi_2(\mathbf{x}_0,\mathbf{x}_0) * \phi_j.$$

These isomorphisms fit together in an additive manner, thanks to the associativity of $*$. We then use the splitting $\pi_2(\mathbf{x}_0,\mathbf{x}_0) \cong \mathbb{Z}\times H^1(Y;\mathbb{Z})$ given by the basepoint, followed by the natural projection to the second factor.

We can then define

$$\underline{\partial}^\infty[\mathbf{x},i] = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \left(\sum_{\phi\in\pi_2(\mathbf{x},\mathbf{y})} \#\mathcal{M}(\phi) e^{A(\phi)}[\mathbf{y}, i-n_z(\phi)] \right),$$

which is a finite sum under the strong admissibility hypotheses.

Analogous constructions work for CF^+, CF^-, and $\widehat{CF}$, as well (with, once again, weak admissibility sufficing for CF^+ and $\widehat{CF}$).

Remark 4.19. *Note that there is a "universal" coefficient system for Lagrangian Floer homology, with coefficients in a group-ring over $\pi_1(\Omega(L_0,L_1))$. In fact, the construction we have here is a specialization of this: in our case, the fundamental group of the configuration space is $\mathbb{Z}\oplus H^1(Y,\mathbb{Z})$, but the $\mathbb{Z}$ summand is already implicit in our consideration of pairs $[\mathbf{x},i]\in(\mathbb{T}_\alpha\cap\mathbb{T}_\beta)\times\mathbb{Z}$.*

It is worth noting that, although the definition of the boundary map still depends on a coherent system of orientations $\mathfrak{o}$, the isomorphism class of the chain complex as a $\mathbb{Z}$-module does not: given a homomorphism $\mu\colon H^1(Y;\mathbb{Z})\longrightarrow \mathbb{Z}/2\mathbb{Z}$, the map

$$f(e^h[\mathbf{x},i]) = (-1)^{\mu(h)} e^h[\mathbf{x},i] \tag{3}$$

gives an isomorphism from the chain complex using $\mathfrak{o}$ to the chain complex using $\mathfrak{o}'$ with $\delta(\mathfrak{o},\mathfrak{o}')=\mu$.

Note that as $\mathbb{Z}$-modules, all of these chain complexes have a natural relative $\mathbb{Z}$ grading, which lifts the obvious relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{s})\mathbb{Z}$-grading. Specifically, given $g\otimes[\mathbf{x},i]$ and $h\otimes[\mathbf{y},j]$ with $g,h\in H^1(Y;\mathbb{Z})$, if we let ϕ be the class with $A(\phi)=g-h$ and $n_z(\phi)=i-j$ (this now uniquely specifies ϕ), we let the relative grading between $g\otimes[\mathbf{x},i]$ and $h\otimes[\mathbf{y},j]$ be given by the Maslov index of ϕ. In view of this, we can think of the corresponding homologies as analogues of a construction of Fintushel and Stern, for $\mathbb{Z}$ graded instanton homology (see [6]).

Theorem 4.20. *The groups* $\underline{HF}^\infty(Y,\mathfrak{s})$, $\underline{HF}^+(Y,\mathfrak{s})$, $\underline{HF}^-(Y,\mathfrak{s})$, *and* $\underline{\widehat{HF}}(Y,\mathfrak{s})$ *are topological invariants of* Y *with its* Spin^c *structure* $\mathfrak{s}$. *These groups are all modules over the group ring* $\mathbb{Z}[H^1(Y;\mathbb{Z})]$.

Proof. Independence of complex structure proceeds exactly as before. For isotopy invariance, observe that an isotopy Ψ_t as in Subsection 4.7 allows one to transfer an additive map A from $\pi_2(\mathbf{x},\mathbf{y})$ for $\mathbf{x},\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ to an additive map on $\pi_2(\mathbf{x}',\mathbf{y}')$ for $\mathbf{x}',\mathbf{y}'\in\Psi_1(\mathbb{T}_\alpha)\cap\mathbb{T}_\beta$. Stabilization is straightforward, and handleslide invariance in the present context is addressed specifically in Subsection 6.5.3. □

For any $\mathbb{Z}[H^1(Y;\mathbb{Z})]$-module M, we have homology groups

$$\underline{HF}(Y,\mathfrak{s};M)=H_*\left(\underline{CF}(Y,\mathfrak{s})\otimes_{Z[H^1(Y;\mathbb{Z})]}M\right)$$

(where $\underline{HF}$ can be any of $\underline{HF}^\infty$, $\underline{HF}^+$, $\underline{HF}^-$, or $\underline{\widehat{HF}}$). We can think of the homology groups with untwisted coefficients constructed earlier as special cases of this construction, thinking of $\mathbb{Z}$ as the trivial $\mathbb{Z}[H^1(Y;\mathbb{Z})]$-module. (In fact, the 2^g different choices of orientation systems over $\mathbb{Z}$ corresponding to the 2^g different ring homomorphisms $\mathbb{Z}[H^1(Y;\mathbb{Z})]\longrightarrow\mathbb{Z}$.)

Note also that the action of $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ on $CF^\infty(Y,\mathfrak{s})$ has an interpretation in this world: there is a natural pairing

$$\mathbb{Z}[H^1(Y;\mathbb{Z})]\otimes(H_1(Y;\mathbb{Z})/\mathrm{Tors})\longrightarrow\mathbb{Z}.$$

The action of ζ defined in Subsection 4.4 can be thought of as given by

$$\langle\underline{\partial}[\mathbf{x},i],\zeta\rangle.$$

4.12. **Further remarks.** There is another construction, which works even in the absence of admissibility hypotheses.

For this construction, we will work over the Novikov ring $\mathbb{A}$ consisting of formal power series $\sum_{r\geq 0}a_r e^r$, for which the support of the a_r (in r) is discrete, endowed with the multiplication law:

$$\left(\sum_{r\geq 0}a_r e^r\right)\cdot\left(\sum_{r\geq 0}b_r e^r\right)=\sum_{r\geq 0}\left(\sum_{s\geq 0}a_s b_{r-s}\right)e^r.$$

We define the boundary map by

$$\partial^+[\mathbf{x},i]=\sum_{\{\mathbf{y}\in\mathcal{S}\}}\sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y})\big|n_z(\phi)\leq i\}}e^{\mathcal{A}(\phi)}\left(\#\mathcal{M}(\phi)\right)\cdot[\mathbf{y},i-n_z(\phi)],$$

where $\mathcal{A}(\phi)$ denotes the area of the domain $\mathcal{D}(\phi)$. Observe that this construction depends on the choice of volume form for Σ through the induced areas of each periodic domain – a real valued function on $H_2(Y;\mathbb{R})$. That datum, in turn, can be thought of as a real two-dimensional cohomology class $\eta\in H^2(Y;\mathbb{R})$.

In this manner, we can obtain homology groups $HF^+_{\mathrm{Nov}}(Y,\mathfrak{s},\eta)$ which are invariants of the underlying topological data, and which require no admissibility hypotheses to define. We will have no further use for this construction in the present paper, though it may turn out to be useful in other applications. In particular, this construction is analogous to the Seiberg-Witten-Floer homology perturbed by a real two-dimensional cohomology class.

5. Special Heegaard moves and the Maslov index calculation

5.1. **Special Heegaard moves.** The purpose of this section is to show that strongly $\mathfrak{s}$-admissible Heegaard diagrams exist, and to study their isotopies, with the aim of proving Theorem 4.15 used in Section 4.

We will be considering certain special isotopies. Let γ be an oriented simple closed curve in Σ. By *winding along* γ we mean the diffeomorphism of Σ obtained by integrating a vector field X supported in a tubular neighborhood of δ, where it satisfies the property that $d\theta(X)>0$, with respect to a coordinate system $(t,\theta)\in(-\epsilon,\epsilon)\times S^1$ in the tubular neighborhood of $\gamma=\{0\}\times S^1$.

Choose a curve γ transverse to α_1, meeting it in a single transverse point, and which is disjoint from the other α_i for $i\neq 1$, and suppose that ϕ is some diffeomorphism which winds along γ. Suppose, moreover, that $\phi(\alpha_1)$ meets α_1 transversally in the neighborhood of γ, meeting it there in $2k$ points. Then, we say that ϕ winds α_1 along γ k times. See Figure 1.

We have the following notion:

Definition 5.1. *Fix a* Spin^c *structure* $\mathfrak{s}$ *over* Y. *A pointed Heegaard diagram*

$$(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$$

is called $\mathfrak{s}$*-realized if there is a point* $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ *with the property that* $s_z(\mathbf{x})=\mathfrak{s}$.

Lemma 5.2. *Fix* Y *and a* Spin^c *structure* $\mathfrak{s}$. *Then,* Y *admits an* $\mathfrak{s}$*-realized pointed Heegaard diagram.*

Proof. Begin with any Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ for Y and let $\boldsymbol{\gamma}$ be a collection of pairwise disjoint curves which are dual to the $\boldsymbol{\alpha}$, in the sense that for all i and j,

$$\#(\alpha_i\cap\gamma_j)=\delta_{i,j}$$

(the right hand side is Kronecker delta, and the left hand side denotes both the geometric and algebraic intersection numbers of the curves). By isotoping the $\boldsymbol{\beta}$ if necessary, we can arrange that $\mathbb{T}_\beta\cap\mathbb{T}_\gamma\neq\emptyset$. Choose a basepoint z distinct from $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ (indeed, choose z to be disjoint from the neighborhood of the $\boldsymbol{\gamma}$ where the winding is performed).

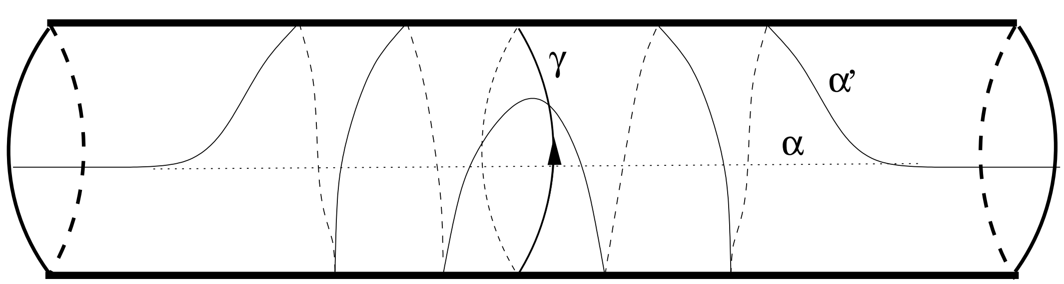

Figure 1. **Winding transverse to** α**.** We have pictured a cylindrical subregion of Σ, where α is the horizontal curve, which we wind twice along the vertical circle γ (in the direction indicated) to obtain α'.

Let $\mathbf{x} = \{x_1, ..., x_g\} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, labelled so that $x_i \in \beta_i \cap \gamma_i$ for $i = 1, ..., g$. Each time we wind α_i along γ_i. we create a new pair of intersection points near x_i between β_i and the new copy of α_i. Winding along each γ_i k times, then, we can label these intersection points $x_i^{\pm}(1), x_i^{\pm}(2), ..., x_i^{\pm}(k)$ (ordered in decreasing order of their distance to x_i, and with sign distinguishing which side of γ_i – in its tubular neighborhood – they lie in). Thus, we have induced intersection points

$$\mathbf{x}(i_1, ..., i_g) = \{x_1^+(i_1), ..., x_g^+(i_g)\} \in \mathbb{T}'_\alpha \cap \mathbb{T}_\beta$$

labeled by $i_1, ..., i_g \in 1, ..., k$.

No matter how many times we wind α_i along γ_i, the Spin^c structure of the farthest intersection point $\mathbf{x}(1, ..., 1)$ remains fixed (this is clear from the definition of $s_z(\mathbf{x})$: the winding isotopy induces an isotopy between the induced non-vanishing vector fields induced over Y). Moreover, we have that

$$\mathfrak{s}_z(\mathbf{x}(i_1, ..., i_g)) - \mathfrak{s}_z(\mathbf{x}(j_1, ..., j_g)) = \Big((i_1 - j_1)\mathrm{PD}[\gamma_1] + ... + (i_g - j_g)\mathrm{PD}[\gamma_g]\Big).$$

(This is a straightforward consequence of the definition of the difference map ϵ, and its compatibility with $\mathfrak{s}_z$; note that with our conventions, the short arc in α_i connecting $x_i(k)$ to $x_i(k+1)$, followed by the short arc in β_i with the same endpoints, is homologous to $-\gamma_i$ in Σ.)

Thus, we can find Heegaard diagrams which realize the Spin^c structures which differ from some fixed Spin^c structure $\mathfrak{s}_0$ by non-positive multiples of the $[\gamma_1]$,...,$[\gamma_g]$. Moreover, if we choose parallel copies $\{\gamma_1^-, ..., \gamma_g^-\}$ of the $\boldsymbol{\gamma}$, only with the opposite orientations, and wind along those in addition, we can realize all Spin^c structures which differ from $\mathfrak{s}_0$ by arbitrary multiples the $[\gamma_1]$,...,$[\gamma_g]$. Now, it is easy to see that the group $H^2(Y; \mathbb{Z})$ is generated by the Poincaré duals of the $\boldsymbol{\gamma}$. Hence, we can realize all Spin^c structures. ☐

Winding can be used also to arrange for strong admissibility.

It is useful to have the following:

Definition 5.3. *An* $\mathfrak{s}$-renormalized periodic domain *is a two-chain* $\mathcal{Q} = \sum a_i \mathcal{D}_i$ *in* Σ *whose boundary is a sum of the curves* $\boldsymbol{\alpha}$ *and* $\boldsymbol{\beta}$ *(with multiplicities), satisfying the additional property that*

$$n_z(\mathcal{Q}) = -\frac{\langle c_1(\mathfrak{s}), H(\mathcal{Q})\rangle}{2}.$$

Of course, the group of $\mathfrak{s}$-renormalized periodic domains is isomorphic to the group of periodic domains. (The periodic domain $\mathcal{P}$ gives rise to the renormalized periodic domain $\mathcal{P} - \frac{\langle c_1(\mathfrak{s}), H(\mathcal{P})\rangle}{2}[\Sigma]$.)

Lemma 5.4. *Fix* Y *and a* Spin^c *structure* $\mathfrak{s}$*. Then,* Y *admits a strongly* $\mathfrak{s}$*-admissible pointed Heegaard diagram.*

Proof. In view of Lemma 5.2, we can start with an $\mathfrak{s}$-realized Heegaard diagram. We will show that after winding the $\boldsymbol{\alpha}$ sufficiently many times along curves $\boldsymbol{\gamma}$ as in the proof of the previous lemma, we obtain a pointed Heegaard diagram for which each renormalized

$\mathfrak{s}$-periodic domain has both positive and negative coefficients. Such a Heegaard diagram is strongly $\mathfrak{s}$-admissible.

Write $b = b_1(Y)$, and choose a basis $\{\mathcal{Q}_1, ..., \mathcal{Q}_b\}$ for the group of renormalized periodic domains. Note that a renormalized periodic domain $\mathcal{Q}$ is uniquely determined by the corresponding vector in $\mathrm{Span}([\alpha_1], ..., [\alpha_g])$ which is the part of $\partial\mathcal{Q}$ – thus, we can think of of the space of renormalized periodic domains as a lattice in this g-dimensional $\mathbb{Z}$-module. After a change of basis of the $\{\mathcal{Q}_i\}$ and reordering the $\boldsymbol{\alpha}$, we can assume that for all $i = 1, ..., b$,

$$\partial\mathcal{Q}_i = \sum_{j=1}^{g} a_{i,j}\alpha_i + b_{i,j}\beta_j,$$

where $a_{i,j} = 0$ for $i > j$, and $a_{i,i} > 0$.

For each $i = 1, ..., b$ choose points $w_i \in \gamma_i$ which are not contained in any of the $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$. Let

$$c_i = \max_{i=1,\ldots,b} |n_{w_i}(\mathcal{Q}_j)|,$$

and then choose some integer N with

$$N > b \cdot \left(\max_{i=1,\ldots,b} \frac{c_i}{a_{i,i}} \right).$$

Choose parallel copies γ_i^- of the γ_i for $i = 1, ..., b$, and let $\{\mathcal{Q}'_1, ..., \mathcal{Q}'_b\}$ be the new periodic domains, obtained after winding the curves $\{\alpha_1, ..., \alpha_b\}$ N times along the $\{\gamma_1, ..., \gamma_b\}$ and N times in the opposite direction along the $\{\gamma_1^-, ..., \gamma_b^-\}$. Note that

$$\begin{aligned} n_{w_i}(\mathcal{Q}'_i) &= n_{w_i}(\mathcal{Q}_i) + N a_{i,i} \\ &> n_{w_i}(\mathcal{Q}_i) + b c_i \\ &\geq (b-1) \max_{j=1,\ldots,i-1,i+1,\ldots,b} |n_{w_i}(\mathcal{Q}_j)| \\ &= (b-1) \max_{j=1,\ldots,i-1,i+1,\ldots,b} |n_{w_i}(\mathcal{Q}'_j)|. \end{aligned}$$

In a similar manner, we see that

$$n_{w_i^-}(\mathcal{Q}'_i) < -(b-1) \max_{j=1,\ldots,i-1,i+1,\ldots,b} |n_{w_i^-}(\mathcal{Q}'_j)|.$$

It is a straightforward matter, then, to verify that for any linear combination of the $\mathcal{Q}'_i$, one can find some point w for which n_w is positive, and another w' for which $n_{w'}$ is negative. □

Indeed, an elaboration of this argument gives the following:

Lemma 5.5. *Suppose that two strongly $\mathfrak{s}$-admissible pointed Heegaard diagrams are isotopic, then they are isotopic through strongly $\mathfrak{s}$-admissible pointed Heegaard diagrams.*

Proof. First, note that if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ is strongly $\mathfrak{s}$-admissible, then if we choose curves along which to wind the $\boldsymbol{\alpha}$ (disjoint from the basepoint z), then the winding gives an isotopy through strongly $\mathfrak{s}$-admissible pointed Heegaard diagrams. The reason for this is that, in the complement of a small neighborhood of the winding region, the various renormalized periodic domains remain unchanged; thus, if some renormalized periodic domain has positive coefficents, then it retains this property as it undergoes winding.

Thus, it suffices to show that if two Heegaard diagrams are isotopic (via an isotopy which we can assume without loss of generality takes place only among the $\boldsymbol{\beta}$ – taking $\boldsymbol{\beta}$ to $\{\beta'_1, ..., \beta'_g\}$), then if we wind their α-curves simultaneously along some collection of $\boldsymbol{\gamma}$ to obtain $\{\alpha'_1, ..., \alpha'_g\}$, then the pointed Heegaard diagrams $(\Sigma, \{\alpha'_1, ..., \alpha'_g\}, \boldsymbol{\beta}, z)$ and $(\Sigma, \{\alpha'_1, ..., \alpha'_g\}, \{\beta'_1, ..., \beta'_g\}, z)$ are isotopic through strongly $\mathfrak{s}$-admissible Heegaard diagrams. To see this, we choose γ_i curves and their translates γ_i^- as in the proof of Lemma 5.4. Now, we choose constants

$$c_i = \sup_{t\in[0,1]} \max_{i=1,...,b} |n_{w_i}(\mathcal{Q}_i(t))|,$$

where we think of $t \in [0,1]$ as the parameter in some isotopy taking $\boldsymbol{\beta}$ to $\{\beta'_1, ..., \beta'_g\}$, and $\mathcal{Q}_i(t)$ is the corresponding one-parameter family of renormalized periodic domains. (Strictly speaking, the point w_i generically lies on the translates of the β_i for finitely many t, so that for those values of t, the multiplicity $n_{w_i}(\mathcal{Q}_i(t))$ does not make sense as we have defined it; for those values of t, we use a small perturbation $w'_i \in \gamma_i$ of the basepoint w_i.) Using these constants c_i as in the proof of Lemma 5.4, the present lemma follows. □

Remark 5.6. *Note that this lemma also proves that any two isotopic $\mathfrak{s}$-realized pointed Heegaard diagrams are isotopic through $\mathfrak{s}$-realized Heegaard diagrams.*

Finally, we can prove Theorem 4.15.

Proof of Theorem 4.15. Lemma 5.4 gives the necessary existence statement. Any two Heegaard diagrams can be connected through a sequence of isotopies, handleslides, and stabilizations (see Proposition 2.1 of [23]). Lemma 5.5 shows that isotopic strongly $\mathfrak{s}$-admissible Heegaard diagrams are isotopic through strongly $\mathfrak{s}$-admissible Heegaard diagrams. In a similar manner, if two Heegaard diagrams can be connected through a handleslide amongst the $\boldsymbol{\beta}$, then by winding sufficiently many times normal to the $\boldsymbol{\alpha}$, we obtain isotopic strongly $\mathfrak{s}$-admissible Heegaard diagrams which are connected by a handleslide through strongly $\mathfrak{s}$-admissible Heegaard diagrams. Noting that the multiplicities of the periodic domains remain unchanged under stabilizations, it follows that stabilizations preserve the strongly $\mathfrak{s}$-admissible condition. The theorem then follows. □

5.2. **Weakly admissible Heegaard diagrams.**

We now justify the use of weakly admissible isotopies for $\widehat{HF}$ and HF^+.

Lemma 5.7. *Any $\mathfrak{s}$-weakly admissible pointed Heegaard diagram is isotopic through an exact Hamiltonian isotopy to some $\mathfrak{s}$-strongly admissible pointed Heegaard diagram*

Proof. This is proved in the same manner as Lemma 5.4, together with the following observation. Winding along γ is, of course, not an exact Hamiltonian isotopy. However, if we wind in the opposite direction along a parallel translate γ_- of γ at the same time, then such an isotopy can be realized by an exact Hamiltonian isotopy. Such an isotopy can be induced from a Hamiltonian function f supported inside an annular neighborhood A of γ. The result then follows. □

This was the remaining lemma used in establishing Theorem 4.14.

5.3. **The Maslov index of a periodic domain.** We can now prove Theorem 4.1, which was used in the definition of the relative gradings.

Proof of Theorem 4.1. Fix $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and consider the map $\mathrm{Hom}(H_2(Y;\mathbb{Z}), 2\mathbb{Z})$ which, given $c \in H_2(Y;\mathbb{Z})$, calculates $\mu(\psi(c))$, where $\psi(c) \in \pi_2(\mathbf{x},\mathbf{x})$ is the periodic class associated to $c \in H_2(Y;\mathbb{Z})$. Note that this is a homomorphism, since the Maslov index is additive. Indeed, this assignment depends on the point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ only through its induced Spin^c structure $\mathfrak{s} = s_z(\mathbf{x})$, by the additivity of the Maslov index. We denote the map by $m_\mathfrak{s} \in \mathrm{Hom}(H_2(Y;\mathbb{Z}),\mathbb{Z})$.

We argue that $m_\mathfrak{s}$ depends on Y alone, i.e. it is invariant under pointed isotopies, pointed handle-slides, and stabilization. In particular, it is also independent of the choice of base-point. To see stabilization invariance, it suffices to see how the Maslov index changes by adding $S \in \pi_2(\mathrm{Sym}^g(\Sigma))$, and thereby reducing to the case where the coefficient of the domain is zero on the two-torus. Handle-slide invariance follows from the holomorphic triangle construction, as in [23]. Specifically, let $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ be attaching circles, where $\boldsymbol{\gamma}$ are obtained from $\boldsymbol{\beta}$ by a handle slide and a small Hamiltonian isotopy. Let $\mathbf{x}'$ be the corresponding intersection of the $\boldsymbol{\alpha}$ with the $\boldsymbol{\gamma}$. There is a class in $\Delta \in \pi_2(\mathbf{x},\theta,\mathbf{x}')$ with $n_z = 0$ and $\mu(\Delta) = 0$. Now, there is an affine identification $\pi_2(\mathbf{x},\theta,\mathbf{x}') \cong \mathbb{Z} \oplus H^1(Y;\mathbb{Z}) \oplus H^1(\#^g(S^1 \times S^2);\mathbb{Z})$ (c.f. Propositions 6.3 and 6.2 in the next section). Hence if $p_\mathbf{x}$ is a periodic class for $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$, and $p'_\mathbf{x}$ is the corresponding periodic class for $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma})$, then there is a periodic class $\delta \in \pi_2(\theta,\theta)$ for $\mathbb{T}_\beta \cap \mathbb{T}_\gamma$ with the property that $p_\mathbf{x} + \Delta = p'_\mathbf{x} + \Delta + \delta$. Moreover, since the Maslov index on any such element δ vanishes (see [23]), it follows that $\mu(p_\mathbf{x}) = \mu(p_{\mathbf{x}'})$. Isotopy invariance is straightforward, except in the case where the isotopy cancels all intersection points belonging to the given Spin^c structure $\mathfrak{s}$. To avoid this we use only special isotopies, as in Lemma 5.2 and 5.5 (see Remark 5.6).

Now, we argue that if $\mathfrak{s}, \mathfrak{s}' \in \mathrm{Spin}^c(Y)$ are represented by intersection points, then we claim that
$$m_\mathfrak{s} = m_{\mathfrak{s}'} + 2c,$$
where $c \in H^2(Y;\mathbb{Z})$ is the class for which $\mathfrak{s}' = \mathfrak{s} - c$. To see this, it suffices to consider the effect of moving the base-point z across some fixed circle, say, α_1. Note then that $s_{z'} = s_z + \alpha_1^*$, according to Lemma 2.6. If ψ is the periodic class corresponding under the basepoint z to some $v \in H_2(Y;\mathbb{Z})$ then clearly $n_{z'}(\psi) = -\langle \alpha_1^*, v\rangle$. Moreover, the periodic class for $\psi(z',v) = \psi(z,v) - n_{z'}(\psi(z,v))[S]$. It follows that $m_\mathfrak{s} = m_{\mathfrak{s}'} + 2c$.

It follows that $m_\mathfrak{s} = c_1(\mathfrak{s}) + K$, in $\mathrm{Hom}(H_2(Y;\mathbb{Z}),\mathbb{Z})$ for some K which is independent of $\mathfrak{s}$. We wish to show that $K = 0$. To this end, we compare $m_\mathfrak{s}$ and $m_{\overline{\mathfrak{s}}}$. Switching the roles of $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ and reversing the orientation of Σ, we get a new Heegaard diagram describing Y, and an obvious identification of intersection points; letting $s'_z(\mathbf{x})$ be the Spin^c structure with respect to this new data, it is clear that $s'_z(\mathbf{x}) = \overline{s_z(\mathbf{x})}$. Note that switching the two tori and the orientation of Σ simultaneously leaves holomorphic data, such as the Maslov index of a given periodic domains, unchanged. In particular, if $\mathcal{P} = \sum a_i \mathcal{D}_i$ is a periodic domain, and $\mathcal{P}' = \sum a_i \mathcal{D}'_i$, where the $\mathcal{D}'_i$ have the opposite orientation to the $\mathcal{D}_i$, then $\langle m_\mathfrak{s}, H(\mathcal{P})\rangle = \langle m_{\overline{\mathfrak{s}}}, H(\mathcal{P}')\rangle$. However, $H(\mathcal{P}) = -H(\mathcal{P}')$. Thus, $m_{\overline{\mathfrak{s}}} = -m_\mathfrak{s}$. Since it is also true that $c_1(\mathfrak{s}) = -c_1(\overline{\mathfrak{s}})$, it follows that $K = 0$. □

6. Holomorphic triangles

Maps between Floer homologies can be constructed by counting pseudo-holomorphic triangles in a given equivalence class. This construction is fundamental to establishing the handleslide invariance of the Floer homologies considered here. They are also useful when comparing the Floer homology groups of three-manifolds which differ by surgeries on a knot (c.f. Section 10).

Since holomorphic triangles fit naturally into a four-dimensional framework, we begin the section by setting up the relevant (four-dimensional) topological preliminaries, including the map from homotopy classes of triangles to Spin^c structures over an associated four-manifold. Next, we study the holomorphic triangle construction itself, proving several basic properties of the maps. We then cast the handleslide invariance of the homology groups into this framework, building on the proof of handleslide invariance for the case where $b_1(Y) = 0$ (c.f. Section 5 of [23]). After that, we indicate the modifications necessary for holomorphic triangles to work in the presence of twisted coefficients (in the sense of Subsection 4.11).

6.1. **Topological preliminaries on triangles.** A *Heegaard triple-diagram of genus g* is an oriented two-manifold and three g-tuples $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ which are complete sets of attaching circles for handlebodies U_α, U_β, and U_γ respectively. Let $Y_{\alpha,\beta} = U_\alpha \cup U_\beta$, $Y_{\beta,\gamma} = U_\beta \cup U_\gamma$, and $Y_{\alpha,\gamma} = U_\alpha \cup U_\gamma$ denote the three induced three-manifolds. A Heegaard triple-diagram naturally specifies a cobordism $X_{\alpha,\beta,\gamma}$ between the these three-manifolds. The cobordism is constructed as follows.

Let Δ denote the two-simplex, with vertices $v_\alpha, v_\beta, v_\gamma$ labeled clockwise, and let e_i denote the edge v_j to v_k, where $\{i,j,k\} = \{\alpha,\beta,\gamma\}$. Then, we form the identification space

$$X_{\alpha,\beta,\gamma} = \frac{(\Delta\times\Sigma)\coprod(e_\alpha\times U_\alpha)\coprod(e_\beta\times U_\beta)\coprod(e_\gamma\times U_\gamma)}{(e_\alpha\times\Sigma)\sim(e_\alpha\times\partial U_\alpha),(e_\beta\times\Sigma)\sim(e_\beta\times\partial U_\beta),(e_\gamma\times\Sigma)\sim(e_\gamma\times\partial U_\gamma)}.$$

Over the vertices of Δ, this space has corners, which can be naturally smoothed out to obtain a smooth, oriented, four-dimensional cobordism between the three-manifolds $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, and $Y_{\alpha,\gamma}$ as claimed.

We will call the cobordism $X_{\alpha,\beta,\gamma}$ described above a *pair of pants connecting $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, and $Y_{\alpha,\gamma}$*. Note that

$$\partial X_{\alpha,\beta,\gamma} = -Y_{\alpha,\beta} - Y_{\beta,\gamma} + Y_{\alpha,\gamma},$$

with the obvious orientation.

Example 6.1. *Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ be a Heegaard diagram for Y, and let $\boldsymbol{\gamma}$ be a g-tuple of curves which are isotopic to $\boldsymbol{\beta}$. Then the triple-diagram*

$$(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma})$$

is a diagram for the cobordism between $-Y$, Y, and $\#^g(S^1\times S^2)$ obtained from $Y\times[0,1]$ by deleting a regular neighborhood of $U_\beta\times\frac{1}{2}$.

6.1.1. *Two-dimensional homology.* We can think of the two-dimensional homology of $X = X_{\alpha,\beta,\gamma}$ in terms of the $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ as follows:

Proposition 6.2. *Let* $\mathrm{Span}([\alpha_i]_{i=1}^g) \subset H_1(\Sigma;\mathbb{Z})$ *denote the lattice spanned by the one-dimensional homology classes induced by the* $\boldsymbol{\alpha}$*. Then, there are natural identifications*

$$H_2(X;\mathbb{Z}) \cong \mathrm{Ker}\Big(\mathrm{Span}([\alpha_i]_{i=1}^g) \oplus \mathrm{Span}([\beta_i]_{i=1}^g) \oplus \mathrm{Span}([\gamma_i]_{i=1}^g) \longrightarrow H_1(\Sigma;\mathbb{Z})\Big); \tag{4}$$

or, equivalently,

$$H_2(X;\mathbb{Z}) \cong \mathrm{Ker}\Big(H_1(\mathbb{T}_\alpha;\mathbb{Z}) \oplus H_1(\mathbb{T}_\beta;\mathbb{Z}) \oplus H_1(\mathbb{T}_\gamma;\mathbb{Z}) \longrightarrow H_1(\mathrm{Sym}^g(\Sigma);\mathbb{Z})\Big). \tag{5}$$

Similarly, we have

$$H_1(X;\mathbb{Z}) \cong \mathrm{Coker}\Big(\mathrm{Span}([\alpha_i]_{i=1}^g) \oplus \mathrm{Span}([\beta_i]_{i=1}^g) \oplus \mathrm{Span}([\gamma_i]_{i=1}^g) \longrightarrow H_1(\Sigma;\mathbb{Z})\Big); \tag{6}$$

Proof. First, note that the boundary homomorphism $\partial\colon H_2(U_\alpha,\Sigma;\mathbb{Z}) \longrightarrow H_1(\Sigma;\mathbb{Z})$ is injective, and its image is $\mathrm{Span}([\alpha_i]_{i=1}^g)$. The first isomorphism then follows from the long exact sequence in homology for the pair $(X, \Delta\times\Sigma)$, bearing in mind that

$$H_2(X,\Delta\times\Sigma) \cong H_2(U_\alpha,\Sigma) \oplus H_2(U_\beta,\Sigma) \oplus H_2(U_\gamma,\Sigma)$$

(by excision), and that the map $H_2(\Sigma) \longrightarrow H_2(X)$ is trivial: the Heegaard surface is obviously null-homologous in X.

The second isomorphism follows from the fact that under the natural identification $H_1(\mathrm{Sym}^g(\Sigma);\mathbb{Z}) \cong H_1(\Sigma;\mathbb{Z})$, the image of $H_1(\mathbb{T}_\alpha;\mathbb{Z})$ is identified with $\mathrm{Span}([\alpha_i]_{i=1}^g)$.

The final isomorphism follows from the fact that

$$H_1(X,\Delta\times\Sigma) \cong H_1(U_\alpha,\Sigma) \oplus H_1(U_\beta,\Sigma) \oplus H_1(U_\beta,\Sigma) \cong H^2(U_\alpha) \oplus H^2(U_\beta) \oplus H^2(U_\gamma) = 0.$$

□

Suppose $(a,b,c) \in \mathrm{Span}([\alpha_i]_{i=1}^g) \oplus \mathrm{Span}([\beta_i]_{i=1}^g) \oplus \mathrm{Span}([\gamma_i]_{i=1}^g)$ satisfies $a+b+c=0$. Then, of course, $a+b+c$ spans some two-chain in Σ. Two-chains of this type which also vanish at a given base-point z (lying outside the collection of attaching circles) are natural analogues of the periodic domains considered earlier. We call such two-chains *triply-periodic domains.* In keeping with earlier terminology, the data

$$(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$$

with $z \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g - \gamma_1 - \ldots - \gamma_g$ is called a *pointed Heegaard triple-diagram.*

6.1.2. *Homotopy classes of triangles.* Let $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\mathbf{y} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, $\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$. Consider the map

$$u\colon \Delta \longrightarrow \mathrm{Sym}^g(\Sigma)$$

with the boundary conditions that $u(v_\gamma) = \mathbf{x}$, $u(v_\alpha) = \mathbf{y}$, and $u(v_\beta) = \mathbf{w}$, and $u(e_\alpha) \subset \mathbb{T}_\alpha$, $u(e_\beta) \subset \mathbb{T}_\beta$, $u(e_\gamma) \subset \mathbb{T}_\gamma$. Such a map is called a *Whitney triangle connecting* $\mathbf{x}$*,* $\mathbf{y}$*, and* $\mathbf{w}$. Two Whitney triangles are homotopic if the maps are homotopic through maps which are all Whitney triangles. We let $\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$ denote the space of homotopy classes of Whitney triangles connecting $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{w}$.

As in the definition of Whitney disks, we have an obstruction

$$\epsilon\colon (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times (\mathbb{T}_\beta \cap \mathbb{T}_\gamma) \times (\mathbb{T}_\alpha \cap \mathbb{T}_\gamma) \longrightarrow \frac{H_1(\mathrm{Sym}^g(\Sigma))}{H_1(\mathbb{T}_\alpha) + H_1(\mathbb{T}_\beta) + H_1(\mathbb{T}_\gamma)} \cong H_1(X;\mathbb{Z})$$

which vanishes if $\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$ is non-empty. The obstruction is defined as follows. Choose an arc $a\subset\mathbb{T}_\beta$ from $\mathbf{x}$ to $\mathbf{y}$, $b\subset\mathbb{T}_\gamma$ from $\mathbf{y}$ to $\mathbf{w}$, and an arc $c\subset\mathbb{T}_\alpha$ from $\mathbf{w}$ to $\mathbf{x}$. Then, $\epsilon(\mathbf{x},\mathbf{y},\mathbf{w})$ is the equivalence class of the closed path $a+b+c$.

Using a base-point $z\in\Sigma-\alpha_1-\ldots-\alpha_g-\beta_1-\ldots-\beta_g-\gamma_1-\ldots-\gamma_g$, we obtain an intersection number

$$n_z\colon \pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\longrightarrow\mathbb{Z}.$$

Proposition 6.3. *Given $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, $\mathbf{y}\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma$, $\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma$, then $\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$ is non-empty if and only if $\epsilon(\mathbf{x},\mathbf{y},\mathbf{w})=0$. Moreover, if $g>1$ and $\epsilon(\mathbf{x},\mathbf{y},\mathbf{w})=0$ then*

$$\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\cong\mathbb{Z}\oplus H_2(X;\mathbb{Z}).$$

Proof. Let $\mathrm{Map}^W(\Delta,\mathrm{Sym}^g(\Sigma))$ denote the space of Whitney triangles connecting $\mathbf{x},\mathbf{y},\mathbf{w}$. Then, evaluation along the boundary gives a fibration

$$\mathrm{Map}^W(\Delta,\mathrm{Sym}^g(\Sigma))\longrightarrow\Omega_{\mathbb{T}_\alpha}(\mathbf{x},\mathbf{y})\times\Omega_{\mathbb{T}_\beta}(\mathbf{y},\mathbf{w})\times\Omega_{\mathbb{T}_\gamma}(\mathbf{x},\mathbf{w}),$$

whose fiber is homotopy equivalent to the space of pointed maps from the sphere to $\mathrm{Sym}^g(\Sigma)$ (the base space here is a product of path spaces). This gives us an exact sequence

$$0\longrightarrow\mathbb{Z}\longrightarrow\pi_0(\mathrm{Map}^W(\Delta,\mathrm{Sym}^g(\Sigma)))\longrightarrow H_1(\mathbb{T}_\alpha)\oplus H_1(\mathbb{T}_\beta)\oplus H_1(\mathbb{T}_\gamma).$$

By definition, $\pi_0(\mathrm{Map}^W(\Delta,\mathrm{Sym}^g(\Sigma)))\cong\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$. The evaluation n_z provides a splitting for the first inclusion, so that

$$\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\cong\mathbb{Z}\oplus\mathrm{Im}\Big(\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\longrightarrow H_1(\mathbb{T}_\alpha)\oplus H_1(\mathbb{T}_\beta)\oplus H_1(\mathbb{T}_\gamma)\Big).$$

That image, in turn, is clearly identified with the kernel of the natural map $H_1(\mathbb{T}_\alpha)\oplus H_1(\mathbb{T}_\beta)\oplus H_1(\mathbb{T}_\gamma)\longrightarrow H_1(\mathrm{Sym}^g(\Sigma))$ (we are using here the fact that $\pi_1(\mathrm{Sym}^g(\Sigma))$ is Abelian). The proposition then follows from Proposition 6.2. ☐

Note that the identification $\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\cong\mathbb{Z}\oplus H_2(X;\mathbb{Z})$ is not canonical, but rather it is affine. Specifically, if we fix a homotopy class: $\psi_0\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$, then any other homotopy class $\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$ differs from ψ_0 by an integer $n_z(\psi)-n_z(\psi_0)$, and a triply-periodic domain $\mathcal{D}(\psi)-\mathcal{D}(\psi_0)-(n_z(\psi)-n_z(\psi_0))[\Sigma]$ (which in turn can be thought of as a two-dimensional homology class in X).

6.1.3. Spin^c *structures.* There is a geometric interpretation of Spin^c structures in four dimensions, analogous to Turaev's interpretation of Spin^c structures in three-dimensions, compare [16] and [11].

Let X be a four-manifold. We consider pairs (J,P), where $P\subset X$ is a collection of finitely many points in X, and J is an almost-complex structure defined over $X-P$. We say that two pairs (J_1,P_1) and (J_2,P_2) are *homologous* if there is a compact one-manifold with boundary $C\subset X$ containing P_1 and P_2, with the property that $J_1|X-C$ is isotopic to $J_2|X-C$. We can think of a Spin^c structure on X as a homology class of such pairs (J,P).

The identification with a more traditional definition is as follows. Note that an almost-complex structure over $X-P$ has a canonical Spin^c structure, and that can be uniquely extended over the points P (the obstruction to extending lies in $H^3(X,X-P)=0$, and the indeterminacy in extending lies in $H^2(X,X-P)=0$). Conversely, given a Spin^c structure with spinor bundle W^+, a generic section $\Phi \in \Gamma(X, W^+)$ vanishes at finitely many points, away from which Clifford multiplication on Φ sets up an isomorphism between TX and W^-, hence endowing TX with a complex structure.

Given a pair (J, P), the first Chern class of the induced complex tangent bundle of $X - P$ canonically extends to give a two-dimensional cohomology class $c_1(J, P) \in H^2(X; \mathbb{Z})$. In fact, this agrees with the first Chern class $c_1(\mathfrak{s})$ of the spinor bundle W^+.

6.1.4. *Triangles and* Spin^c *structures.* The base-point $z \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g - \gamma_1 - ... - \gamma_g$ gives rise to a relationship between Spin^c structures on X and holomorphic triangles, analogous to the construction of the Spin^c structure on a three-manifold belonging to intersection point between $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ together with the basepoint z.

To describe this, fix "height functions" over the handlebodies $f_i \colon U_i \longrightarrow [0, 1]$ where $i = \alpha$, β, or γ with only g index one critical points and one index zero critical point, with $f_i(\partial U_i) = 1$.

Now, given a generic map $u \colon \Delta \longrightarrow \mathrm{Sym}^g(\Sigma)$ representing $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, there is an immersed surface-with-boundary $F = F_0 \cup F_1 \subset X_{\alpha,\beta,\gamma}$ constructed as follows. The intersection of the component F_0 with $U_\xi \times e_\xi$, is the product of e_ξ with the upward gradient connecting the index zero critical point with the point $z \in \Sigma$; its intersection with $\Delta \times \Sigma$ is simply $\Delta \times \{z\}$. The intersection of F_1 with $U_\xi \times e_\xi$ is given by the g-tuple of f_ξ gradient flow-lines connecting the various index one critical points with the g points over $(x, u(x))$ (where $x \in e_\xi$). Finally, in the inside region $\Delta \times \Sigma$, the subset F_1 consists of points (x, σ), where $\sigma \in u(x)$. Note that in the complement $X - (F_0 \cup F_1)$, there is a well-defined oriented two-plane field $\mathcal{L}$ which is tangent to Σ inside $\Delta \times \Sigma$, and agrees with the kernel of df_ξ in $TU_\xi \subset T(U_\xi \times e_\xi)$.

In fact, we extend the two-plane field further. Fix a central point $x \in \Delta$, and three paths a, b, and c from x to the edges e_α, e_β, and e_γ respectively. In the complement $\Delta - a \cup b \cup c$, there is a foliation by line segments which connect pairs of edges. For example, there is a family $\ell_{\alpha,\beta}(t)$ of paths connecting e_α to e_β which degenerates as $t \mapsto 0$ to the vertex v_γ, and as $t \mapsto 1$ it degenerates to $a \cup b$. There are analogous families of leaves $\ell_{\beta,\gamma}(t)$ and $\ell_{\alpha,\gamma}(t)$

There is a natural map $\pi \colon X \longrightarrow \Delta$. The preimage under π of $\ell_{\alpha,\beta}(t)$ for $t \in [0, 1)$, which we denote $\widetilde{\ell}_{\alpha,\beta}(t)$, is identified with $Y_{\alpha,\beta}$. For all but finitely many t in the open interval, the intersection of F with $\widetilde{\ell}_{\alpha,\beta}(t)$ consists of $g+1$ disjoint paths which connect the critical points of f_α in U_α to critical points of f_β in U_β. For t, we extend the oriented two-plane field in over a neighborhood of these $g+1$ paths (as in Subsection 2.5) in a continuous manner. In this way, we have extended $\mathcal{L}$ across the intersection of F with $\widetilde{\ell}_{\alpha,\beta}(t)$ for all but finitely many t.

We proceed in the analogous manner to extend over the $\widetilde{\ell}_{\beta,\gamma}(t)$ and $\widetilde{\ell}_{\alpha,\gamma}(t)$.

We have now extended $\mathcal{L}$ over X, except for the intersection of F with certain excluded leaves in the foliation of Δ. These excluded leaves fall into two categories. First, there is the singular leaf $a \cup b \cup c$; and then there are those leaves in Δ which contain a point x for which $\sigma(x)$ has either a repeated entry, or $\sigma(x)$ contains the basepoint $z \in \Sigma$. These are the points where the paths of F cross. One can see that generically the intersection of F with the preimages of these special leaves is a collection of contractible one-complexes; so its tubular neighborhood consists of a finite collection of disjoint four-balls embedded in X.

The two-plane field $\mathcal{L}$ and the orientation on X determine a complex structure over the complement of finitely many balls in X, and hence a Spin^c structure over X.

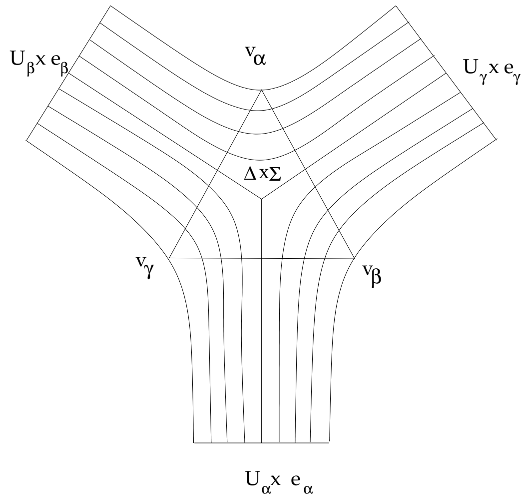

FIGURE 2. Schematic for the cobordism X. We have illustrated the foliation of the triangle by segments whose preimages are the three-manifolds $\widetilde{\ell}_{\xi,\eta}(t)$.

Proposition 6.4. *The above construction induces a map*

$$\mathfrak{s}_z\colon \pi_2(\mathbf{x},\mathbf{y},\mathbf{w}) \longrightarrow \mathrm{Spin}^c(X).$$

Proof. Recall that $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{w}$ determine the two-plane field on the boundary minus finitely many three-balls. Fix u and choose extensions over the three-balls (c.f. Section 2.5). This data specifies the two-plane field over $X_{\alpha,\beta,\gamma} - (\mathrm{int}(\Delta)\times\Sigma) - \mathrm{int}F_0 - \mathrm{int}F_1$. The above discussion shows that the Spin^c structure extends over this region, and, indeed, since the deleted region is topologically a $\Delta\times\Sigma$, it follows from a cohomology long exact sequence that the extension is unique. It is easy to see also that the induced Spin^c structure does not depend on the extension of the two-plane fields to the three-balls in the boundary.

Changing u by a homotopy moves F_1 by an isotopy, so it is easy to see that the induced Spin^c structure depends only on the homotopy class of u. □

Homotopy classes of Whitney triangles can be collected into Spin^c-equivalence classes, as follows. Let $\mathbf{x},\mathbf{x}'\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, $\mathbf{y},\mathbf{y}'\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma$, and $\mathbf{v},\mathbf{v}'\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma$. We say that two homotopy classes $\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{v})$ and $\psi'\in\pi_2(\mathbf{x}',\mathbf{y}',\mathbf{v}')$ are Spin^c-*equivalent*, or simply equivalent, if there are classes $\phi_1\in\pi_2(\mathbf{x},\mathbf{x}')$, $\phi_2\in\pi_2(\mathbf{y},\mathbf{y}')$, and $\phi_3\in\pi_2(\mathbf{v},\mathbf{v}')$ with

$$\psi' = \psi + \phi_1 + \phi_2 + \phi_3.$$

Let $S_{\alpha,\beta,\gamma}$ denote the space of homotopy classes of such triangles.

To justify the terminology, we claim that the Spin^c structure constructed above depends only on its Spin^c-equivalence class as follows:

Proposition 6.5. *The map from Proposition 6.4 descends to a map*

$$\mathfrak{s}_z\colon S_{\alpha,\beta,\gamma} \longrightarrow \mathrm{Spin}^c(X_{\alpha,\beta,\gamma})$$

which is one-to-one, with image consisting of those Spin^c*-structures whose restrictions to the boundary are realized by intersection points.*

Proof. First, we verify that we have characterized the image. Recall that for an arbitrary four-manifold-with-boundary (X,Y) there is a canonical map $\epsilon'\colon \mathrm{Spin}^c(Y) \longrightarrow H^3(X,Y;\mathbb{Z})$. which is defined as follows. Choose a Spin^c structure $\mathfrak{s}_0$ over X, and let

$$\epsilon'(\mathfrak{t}) = \delta(\mathfrak{t}-\mathfrak{s}_0|_Y),$$

where $\delta\colon H^2(Y;\mathbb{Z}) \longrightarrow H^3(X,Y;\mathbb{Z})$ is the coboundary map. It is easy to see that ϵ' is independent of the choice of $\mathfrak{s}_0$, and that it vanishes if and only if $\mathfrak{t}$ extends over X. Next, we argue that $\epsilon(\mathbf{x},\mathbf{y},\mathbf{w}) = \pm\epsilon'(\mathbf{x},\mathbf{y},\mathbf{w})$. To see this, isotope $\mathbb{T}_\alpha$, $\mathbb{T}_\beta$, and $\mathbb{T}_\gamma$ so that there are intersection points $\mathbf{x}'$, $\mathbf{y}'$, and $\mathbf{w}'$ for which $\epsilon(\mathbf{x}',\mathbf{y}',\mathbf{w}')=0$, so that there is a triangle connecting them. We have explicitly constructed the corresponding Spin^c structure, thus $\epsilon'(\mathbf{x}',\mathbf{y}',\mathbf{w}')=0$, as well. It is easy to see that

$$\epsilon(\mathbf{x},\mathbf{y},\mathbf{w}) - \epsilon(\mathbf{x}',\mathbf{y}',\mathbf{w}') = \pm\delta\Big(\mathrm{PD}(\epsilon(\mathbf{x},\mathbf{x}'))\oplus\mathrm{PD}(\epsilon(\mathbf{y},\mathbf{y}'))\oplus\mathrm{PD}(\epsilon(\mathbf{w},\mathbf{w}'))\Big).$$

Similarly,

$$\epsilon'(\mathbf{x},\mathbf{y},\mathbf{w}) - \epsilon'(\mathbf{x}',\mathbf{y}',\mathbf{w}') = \delta\Big(\big(\mathfrak{s}_z(\mathbf{x})-\mathfrak{s}_z(\mathbf{x}')\big)\oplus\big(\mathfrak{s}_z(\mathbf{y})-\mathfrak{s}_z(\mathbf{y}')\big)\oplus\big(\mathfrak{s}_z(\mathbf{w})-\mathfrak{s}_z(\mathbf{w}')\big)\Big).$$

It follows that $\epsilon(\mathbf{x},\mathbf{y},\mathbf{w}) = \pm\epsilon'(\mathbf{x},\mathbf{y},\mathbf{w})$: the obstructions to extending a Spin^c structure are the same as the obstruction to finding a Whitney triangle.

Suppose that u and v are a pair of triangles in $\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$ with $n_z(u)=n_z(v)$, so that their difference from Proposition 6.3 can be interpreted as the triply-periodic domain $\mathcal{D}(u)-\mathcal{D}(v)$. We claim that this triply-periodic domain is gives rise to a relative cohomology class in $H^2(X,\partial X;\mathbb{Z})$ whose image in $H^2(X)$ is the difference $\mathfrak{s}_z(u)-\mathfrak{s}_z(v)$. This is a local calculation since, as is easy to verify, the restriction map

$$H^2(X,\partial X)\longrightarrow H^2(U_\alpha\times(e_\alpha,\partial e_\alpha))\oplus H^2(U_\beta\times(e_\beta,\partial e_\beta))\oplus H^2(U_\gamma\times(e_\gamma,\partial e_\gamma))$$

is injective, and each of the latter groups is generated by the Poincaré duals to curves $[\xi_i^*]\times e_\xi$ (where $\xi=\alpha$, β, or γ, and $i=1,...,g$). On the one hand, the evaluation of a triply-periodic domain on, say, $\alpha_1^*\times[0,1]$ is easily seen to be simply the multiplicity of α_1 in the boundary of the triply-periodic domain. On the other hand, the pair of two-plane fields representing $\mathfrak{s}_z(u)$ and $\mathfrak{s}_z(v)$ differ over $\alpha_1^*\times e_\alpha$ only at those points where one of $u(e_\alpha)$ or $v(e_\alpha)$ contains $\alpha_1\cap\alpha_1^*$. The fact that the constant appearing here is one could be determined by calculating a model case (see [26]). ☐

6.1.5. *Higher polygons.* The above results for triangles admit straightforward generalizations to arbitrarily large collections of g-tuples, which call *Heegaard multi-diagrams* (or *pointed Heegaard multi-diagrams*, when they are equipped with a basepoint z in the complement of all the attaching circles). In fact, the only other case we will require in the present work is the case of squares. Specifically, an oriented two-manifold Σ and four g-tuples of attaching circles $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\delta}$ specify a four-manifold $X_{\alpha,\beta,\gamma,\delta}$ which provides a cobordism between $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, $Y_{\gamma,\delta}$ and $Y_{\alpha,\delta}$. It admits two obvious decompositions

$$X_{\alpha,\beta,\gamma,\delta} = X_{\alpha,\beta,\gamma} \cup_{Y_{\alpha,\gamma}} X_{\alpha,\gamma,\delta} = X_{\alpha,\beta,\delta} \cup_{Y_{\beta,\delta}} X_{\beta,\gamma,\delta}.$$

We can define homotopy classes of squares $\pi_2(\mathbf{x},\mathbf{y},\mathbf{v},\mathbf{w})$ in $\mathrm{Sym}^g(\Sigma)$, and equivalence classes of homotopy classes $S_{\alpha,\beta,\gamma,\delta}$ – i.e. two squares $\varphi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{v},\mathbf{w})$ and $\varphi' \in \pi_2(\mathbf{x}',\mathbf{y}',\mathbf{v}',\mathbf{w}')$ are equivalent if there are $\phi_1 \in \pi_2(\mathbf{x},\mathbf{x}')$, $\phi_2 \in \pi_2(\mathbf{y},\mathbf{y}')$, $\phi_3 \in \pi_2(\mathbf{v},\mathbf{v}')$, and $\phi_4 \in \pi_2(\mathbf{w},\mathbf{w}')$ with

$$\varphi + \phi_1 + \phi_2 + \phi_3 + \phi_4 = \varphi'.$$

Proposition 6.5 admits a straightforward generalization, giving a map from $S_{\alpha,\beta,\gamma,\delta}$ to the space of Spin^c structures over $X_{\alpha,\beta,\gamma,\delta}$.

6.2. **Orienting spaces of pseudo-holomorphic triangles.** We will be counting pseudo-holomorphic triangles. To achieve the required transversality, we allow J to be a function from Δ to the space of almost-complex structures over $\mathrm{Sym}^g(\Sigma)$ chosen to be compatible near the corners with the paths J_s used to define the notion of pseudo-holomorphic disk (see Section 5 of [23]). Moreover, we will use a class of perturbations of the constant complex structure for which the analogue of Theorem 3.1 still holds: if u is a J-holomorphic triangle, the domain associated to u is non-negative.

Now, we can collect the space of J-holomorphic Whitney triangles representing a fixed homotopy class into a moduli space, which we denote $\mathcal{M}(\psi)$. This moduli space has an expected dimension, which we will denote $\mu(\psi)$.

With the transversality in place, the modulo two count of $\mathcal{M}_J(\psi)$ is straightforward to define. When we wish to work over $\mathbb{Z}$, however, we must use a refined count. Again this can be done since the determinant line bundle of the tangent space admits an extension $\det(D_u)$ as a trivial line bundle over each component $\psi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$.

Let $\mathfrak{s}$ be a Spin^c structure over the four-manifold X specified by a pointed Heegaard triple $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},z)$, and let $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma}$ be coherent systems of orientations for the three bounding three-manifolds.

Definition 6.6. *A* coherent system of orientations for $\mathfrak{s}$ $\mathfrak{o}_{\alpha,\beta,\gamma}$, compatible with $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\gamma}$, and $\mathfrak{o}_{\alpha,\gamma}$ *is a collection of sections* $\mathfrak{o}_{\alpha,\beta,\gamma}$ *of the determinant line bundle* $\det(D_u)$ *for each homotopy class of triangle* ψ *representing the* Spin^c *structure* $\mathfrak{s}$*, which is compatible with splicing in the sense that if* $\psi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$*,* $\psi_1 \in \pi_2(\mathbf{x},\mathbf{x}')$*,* $\psi_2 \in \pi_2(\mathbf{y},\mathbf{y}')$*,* $\psi_3 \in \pi_2(\mathbf{w},\mathbf{w}')$ *are any three Whitney disks, then:*

$$\mathfrak{o}_{\alpha,\beta,\gamma}(\psi + \phi_1 + \phi_2 + \psi_3) = \mathfrak{o}_{\alpha,\beta,\gamma}(\psi) \wedge \mathfrak{o}_{\alpha,\beta}(\phi_1) \wedge \mathfrak{o}_{\beta,\gamma}(\phi_2) \wedge \mathfrak{o}_{\alpha,\gamma}(\phi_3),$$

under the identification coming from splicing.

Existence is ensured by the following:

Lemma 6.7. *Let* $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},z)$ *be a pointed Heegaard triple, and fix a* Spin^c *structure* $\mathfrak{s}$ *over* $X_{\alpha,\beta,\gamma}$ *whose restrictions* $\mathfrak{t}_{\alpha,\beta}$*,* $\mathfrak{t}_{\beta,\gamma}$ *and* $\mathfrak{t}_{\alpha,\gamma}$ *are all realized by intersection points. For coherent systems* $\mathfrak{o}_{\alpha,\beta}$ *and* $\mathfrak{o}_{\beta,\gamma}$ *for two of the boundary components, there always exists at least one system of coherent orientation system* $\mathfrak{o}_{\alpha,\gamma}$ *for the remaining boundary component, and a coherent system* $\mathfrak{o}_{\alpha,\beta,\gamma}$ *which is compatible with the* $\mathfrak{o}_{\alpha,\beta}$*,* $\mathfrak{o}_{\beta,\gamma}$*, and* $\mathfrak{o}_{\alpha,\gamma}$*.*

Proof. Let $\psi_0 \in \pi_2(\mathbf{x}_0, \mathbf{y}_0, \mathbf{w}_0)$ be a fixed homotopy class representing $\mathfrak{s}$. Fix an arbitrary orientation $\mathfrak{o}_{\alpha,\beta,\gamma}(\psi_0)$.

Next, we construct $\mathfrak{o}_{\alpha,\gamma}(\phi_{\alpha,\gamma})$, where $\phi_{\alpha,\gamma} \in \pi_2(\mathbf{w}_0, \mathbf{w}_0)$ are periodic classes. Observe that there is a subgroup K of periodic $\phi_3 \in \pi_2(\mathbf{w}_0, \mathbf{w}_0)$ which satisfy the property that

$$\psi_0 + \phi_3 = \psi_0 + \phi_1 + \phi_2$$

for some periodic domains ϕ_1 and ϕ_2 for $\pi_2(\mathbf{x}_0, \mathbf{x}_0)$ and $\pi_2(\mathbf{y}_0, \mathbf{y}_0)$ respectively. (Indeed, the ϕ_1 and ϕ_2 are uniquely specified, as $\mathcal{D}(\phi_1)$ is uniquely specified by its α-boundary, which should agree with the α-boundary of $\mathcal{D}(\phi_3)$, and $\mathcal{D}(\phi_2)$ is similarly determined by the γ-boundary of $\mathcal{D}(\phi_3)$.) It is easy to see that the quotient Q of $\pi_2(\mathbf{w}_0, \mathbf{w}_0)$ by the subgroup K has no torsion, so we have a splitting

$$\pi_2(\mathbf{w}_0, \mathbf{w}_0) \cong K \oplus Q.$$

For $\phi_3 \in K$, we define $\mathfrak{o}_{\alpha,\gamma}(\phi_3)$ so that

$$\mathfrak{o}_{\alpha,\beta,\gamma}(\psi_0) \wedge \mathfrak{o}_{\alpha,\gamma}(\phi_3) = \mathfrak{o}_{\alpha,\beta,\gamma}(\psi_0) \wedge \mathfrak{o}_{\alpha,\beta}(\phi_1) \wedge \mathfrak{o}_{\beta,\gamma}(\phi_2).$$

We then define $\mathfrak{o}_{\alpha,\gamma}(\phi)$ arbitrarily on a basis of generators for Q, and allow that to induce the orientation on all $\psi \in \pi_2(\mathbf{x}_0, \mathbf{y}_0, \mathbf{w}_0)$.

As a final step, we choose a complete set of paths $\{\theta_i\}_{i=1}^m$ for $Y_{\alpha,\gamma}$ over which we choose our orientations (for $\mathfrak{o}_{\alpha,\gamma}$) arbitrarily, and use them to define the orientation for all the remaining $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ in the given Spin^c-equivalence class. □

6.3. **Holomorphic triangles and maps between Floer homologies.** Our aim is to use these counts to define maps between Floer homologies. To do this, we will need our triple-diagram to satisfy some admissibility hypotheses, which are direct generalizations of the admissibility conditions from Subsection 4.2.

Definition 6.8. *A pointed Heegaard triple-diagram is called* weakly admissible *if each non-trivial triply-periodic domain which can be written as a sum of doubly-periodic domains has both positive and negative coefficients. A pointed triple-diagram is called* strongly admissible *for the* Spin^c *structure* $\mathfrak{s}$ *if for each triply-periodic domain* $\mathcal{D}$ *which can be written as a sum of doubly-periodic domains*

$$\mathcal{D} = \mathcal{D}_{\alpha,\beta} + \mathcal{D}_{\beta,\gamma} + \mathcal{D}_{\alpha,\gamma}$$

with the property that

$$\langle c_1(\mathfrak{s}_{\alpha,\beta}), H(\mathcal{D}_{\alpha,\beta})\rangle + \langle c_1(\mathfrak{s}_{\beta,\gamma}), H(\mathcal{D}_{\beta,\gamma})\rangle + \langle c_1(\mathfrak{s}_{\alpha,\gamma}), H(\mathcal{D}_{\alpha,\gamma})\rangle = 2n \geq 0,$$

there is some coefficient of $\mathcal{D} > n$. *(In the above expression, of course,* $\mathfrak{s}_{\xi,\eta}$ *is the restriction of* $\mathfrak{s}$ *to the boundary component* $Y_{\xi,\eta}$*).*

Note that the above notion of weak admissibility is independent of Spin^c structures – it corresponds to the notion of weak admissibility for any torsion Spin^c structure, for an ordinary pointed Heegaard diagram. (We could, of course, have given a slightly weaker formulation depending on the Spin^c structure, more parallel to the definition of weakly admissible for pointed Heegaard diagrams given earlier, but we have no particular use for this presently.)

The following are analogues of Lemmas 4.5 and 4.6:

Lemma 6.9. *Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$ be weakly admissible Heegaard triple, with underlying four-manifold X. Fix intersection points $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{w}$ and a* Spin^c *structure $\mathfrak{s}$ over X. Then, for each integer k, there are only finitely many homotopy classes $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ with $n_z(\psi) = k$ with $\mathfrak{s}_z(\psi) = \mathfrak{s}$, and which support holomorphic representatives.*

Proof. Given $\psi, \psi' \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ with $n_z(\psi) = n_z(\psi')$ and $\mathfrak{s}_z(\psi) = \mathfrak{s}_z(\psi')$, the difference $\mathcal{D}(\psi) - \mathcal{D}(\psi')$ is a triply-periodic domain which, in view of Proposition 6.5, can be written as a sum of doubly-periodic domains. Given this, finiteness follows as in the proof of Lemma 4.5. □

Lemma 6.10. *For a strongly admissible pointed Heegaard triple-diagram for a given* Spin^c *structure $\mathfrak{s}$, and an integer j, there are only finitely many $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ representing $\mathfrak{s}$ with $\mu(\psi) = j$ and which support holomorphic representatives.*

Proof. Suppose that $\psi, \psi' \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ satisfy $\mathfrak{s}_z(\psi) = \mathfrak{s}_z(\psi')$, and $\mu(\psi) = \mu(\psi')$. Then we can write $\psi' = \psi + \phi_1 + \phi_2 + \phi_2$; so by the additivity of the index, it follows that $\mu(\phi_1) + \mu(\phi_2) + \mu(\phi_3) = 0$ (which is identified with the first Chern class evaluation). The proof then follows from the proof of Lemma 4.6. □

Existence of admissible triples follows along the lines of Section 5.

Lemma 6.11. *Given a Heegaard triple-diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$, there is an isotopic weakly admissible Heegaard triple diagram. Moreover, given a* Spin^c *structure $\mathfrak{s}$ over X, there is an isotopic strongly $\mathfrak{s}$-admissible Heegaard triple diagram.*

Proof. This follows as in Lemma 5.4: we wind transverse to all of the $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ simultaneously. □

A Spin^c structure over X gives rise to a map

$$f^{\infty}(\,\cdot\,;\mathfrak{s})\colon CF^{\infty}(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}) \otimes CF^{\infty}(Y_{\beta,\gamma}, \mathfrak{s}_{\beta,\gamma}) \longrightarrow CF^{\infty}(Y_{\alpha,\gamma}, \mathfrak{s}_{\alpha,\gamma})$$

by the formula:

(7)

$$f^{\infty}_{\alpha,\beta,\gamma}([\mathbf{x}, i] \otimes [\mathbf{y}, j]; \mathfrak{s}) = \sum_{\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\{\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w}) \big| \mathfrak{s}_z(\psi) = \mathfrak{s}, \mu(\psi) = 0\}} \Big(\#\mathcal{M}(\psi)\Big) \cdot [\mathbf{w}, i + j - n_z(\psi)].$$

For each fixed $[\mathbf{x}, i]$ and $[\mathbf{y}, j]$ the above is a finite sum when the triple is strongly admissible for $\mathfrak{s}$.

In fact, for each fixed $[\mathbf{x}, i]$ and $[\mathbf{y}, j]$, the $[\mathbf{w}, k]$ coefficient is a sum of $\#\mathcal{M}(\psi)$, where ψ ranges over $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ with $\mathfrak{s}_z(\psi) = \mathfrak{s}$ and $n_z(\psi) = i + j - k$. Thus (according to Lemma 6.9), the $[\mathbf{w}, k]$ coefficient is given by a finite sum under the weak admissibility hypothesis.

Hence, if the triple is weakly admissible, the above sum induces a map

$$f^{+}_{\alpha,\beta,\gamma}(\cdot\,;\mathfrak{s})\colon CF^{+}(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}) \otimes CF^{\leq 0}(Y_{\beta,\gamma}, \mathfrak{s}_{\beta,\gamma}) \longrightarrow CF^{+}(Y_{\alpha,\gamma}, \mathfrak{s}_{\alpha,\gamma}),$$

where, $CF^{\leq 0}(Y,\mathfrak{s}) \subset CF^\infty(Y,\mathfrak{s})$ is the subcomplex generated by $[\mathbf{x},i]$ with $i \leq 0$. Of course, $CF^{\leq 0}(Y,\mathfrak{s})$ is isomorphic to $CF^-(Y,\mathfrak{s})$ as a chain complex (but the latter is generated by $[\mathbf{x},i]$ with $i<0$).

Similarly, we can define a map

$$\widehat{f}_{\alpha,\beta,\gamma}(\mathbf{x}\otimes\mathbf{y};\mathfrak{s}) = \sum_{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma} \sum_{\{\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\big|\mathfrak{s}_z(\psi)=\mathfrak{s},\mu(\psi)=0,n_z(\psi)=0\}} (\#\mathcal{M}(\psi))\,\mathbf{w}.$$

Again, this is a finite sum under the weak admissibility hypothesis.

Theorem 6.12. *Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},z)$ be a pointed Heegaard triple-diagram, which is strongly $\mathfrak{s}$-admissible for some Spin^c structure $\mathfrak{s}$ over the underlying four-manifold X. Then the sum on the right-hand-side of Equation (7) is finite, giving rise to a U-equivariant chain map which also induces maps on homology:*

$$\begin{aligned}
F^\infty_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon HF^\infty(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes HF^\infty(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma}) &\longrightarrow HF^\infty(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma})\\
F^{\leq 0}_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon HF^{\leq 0}(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes HF^{\leq 0}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma}) &\longrightarrow HF^{\leq 0}(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma}).
\end{aligned}$$

The induced (U-equivariant) chain map

$$f^+_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon CF^+(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes CF^{\leq 0}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma}) \longrightarrow CF^+(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma})$$

gives a well-defined chain map when the triple diagram is only weakly admissible, and the Heegaard diagram $(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},z)$ is strongly admissible for $\mathfrak{s}_{\beta,\gamma}$. In fact, the induced map

$$\widehat{f}_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \widehat{CF}(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes \widehat{CF}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma}) \longrightarrow \widehat{CF}(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma})$$

gives a well-defined chain map when the diagram is weakly admissible. There are induced maps on homology:

$$\begin{aligned}
\widehat{F}_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \widehat{HF}(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes \widehat{HF}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma}) &\longrightarrow \widehat{HF}(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma}),\\
F^+_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon HF^+(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes HF^{\leq 0}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma}) &\longrightarrow HF^+(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma}),
\end{aligned}$$

the latter of which is also U-equivariant.

Proof. The fact that $f^\infty_{\alpha,\beta,\gamma}$ is a chain map follows by counting ends of one-dimensional moduli spaces of holomorphic triangles (see Section 5 of [23]; compare [20]). Fix $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, $\mathbf{y}\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma$, $\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma$, and consider moduli spaces of holomorphic triangles $\mathcal{M}(\psi)$ where $\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})$, $\mathfrak{s}_z(\psi)=\mathfrak{s}$, and $\mu(\psi)=1$. The ends of this moduli space are modeled on:

$$\begin{gathered}
\left(\coprod_{\mathbf{x}'\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\coprod_{\phi_{\alpha,\beta}*\psi_{\alpha,\beta}=\psi}\widehat{\mathcal{M}}(\phi_{\alpha,\beta})\times\mathcal{M}(\psi_{\alpha,\beta})\right)\\
\coprod\\
\left(\coprod_{\mathbf{y}'\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma}\coprod_{\phi_{\beta,\gamma}*\psi_{\beta,\gamma}=\psi}\widehat{\mathcal{M}}(\phi_{\beta,\gamma})\times\mathcal{M}(\psi_{\beta,\gamma})\right)\\
\coprod\\
\left(\coprod_{\mathbf{w}'\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma}\coprod_{\phi_{\alpha,\gamma}*\psi_{\alpha,\gamma}=\psi}\widehat{\mathcal{M}}(\phi_{\alpha,\gamma})\times\mathcal{M}(\psi_{\alpha,\gamma})\right).
\end{gathered}$$

In the above expression, the pairs of homotopy classes $\phi_{\alpha,\beta}$ and $\psi_{\alpha,\beta}$ range over $\phi_{\alpha,\beta}\in\pi_2(\mathbf{x},\mathbf{x}')$ and $\psi_{\alpha,\beta}\in\pi_2(\mathbf{x}',\mathbf{y},\mathbf{w})$ with $\mu(\phi_{\alpha,\beta})=1$, $\mu(\psi_{\alpha,\beta})=0$, $\phi_{\alpha,\beta}*\psi_{\alpha,\beta}=\psi$ (with analogous conditions for the $\phi_{\beta,\gamma}\in\pi_2(\mathbf{y},\mathbf{y}')$ and $\phi_{\alpha,\gamma}\in\pi_2(\mathbf{w}',\mathbf{w})$). Counted with signs, the first two unions give the $[\mathbf{w},i+j-n_z(\psi)]$-coefficient of $f^\infty_{\alpha,\beta,\gamma}\circ\partial([\mathbf{x},i]\otimes[\mathbf{y},j])$ (using the natural differential on the tensor product), while the last gives the $[\mathbf{w},i+j-n_z(\psi)]$-coefficient of $\partial\circ f^\infty([\mathbf{x},i]\otimes[\mathbf{y},j])$.

Recall that if ψ has a holomorphic representative, then $n_z(\psi)\geq 0$. Thus, f^∞ maps the subcomplex

$$CF^{\leq 0}(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes CF^{\leq 0}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma})\subset CF^\infty(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes CF^\infty(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma})$$

into $CF^{\leq 0}(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma})$. Similarly, $f^+_{\alpha,\beta,\gamma}$ as above also gives a chain map.

The U-equivariance

$$f^\infty_{\alpha,\beta,\gamma}\left(U\left([\mathbf{x},i]\right)\otimes[\mathbf{y},j]\right)=U\cdot f^\infty_{\alpha,\beta,\gamma}\left([\mathbf{x},i]\otimes[\mathbf{y},j]\right)$$

(and indeed for the other induced maps, where stated) follows immediately from the definitions. □

Now familiar arguments can be used to establish invariance properties of these maps, as in the following:

Proposition 6.13. *The maps on homology listed in Theorem 6.12 are independent of the choice of family J (and underlying complex structure $\mathfrak{j}$ over Σ) used in its definition.*

Proof. Fix first the complex structure $\mathfrak{j}$ over Σ. Consider a one-parameter variation family of maps J_τ from Δ into the space of almost-complex structures over $\mathrm{Sym}^g(\Sigma)$, where τ is a real parameter $\tau\in[0,1]$ (which are perturbations of the symmetrized complex structure $\mathrm{Sym}^g(\mathfrak{j})$ over $\mathrm{Sym}^g(\Sigma)$). We write down the case of CF^∞; the other homology theories work the same way, with only notational changes. Consider the map

$$H^\infty\colon CF^\infty(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes CF^\infty(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma})\longrightarrow CF^\infty(Y_{\alpha,\gamma},\mathfrak{s}_{\alpha,\gamma})$$

defined by

$$H^\infty([\mathbf{x},i]\otimes[\mathbf{y},j],\mathfrak{s})=\sum_{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma}\ \sum_{\{\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\big|\mathfrak{s}_z(\psi)=\mathfrak{s},\mu(\psi)=-1\}}\#\left(\bigcup_{\tau\in[0,1]}\mathcal{M}_{J_\tau}(\psi)\right)[\mathbf{w},i+j-n_z(\psi)].$$

Now, the ends of

$$\coprod_{\{\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\big|\mathfrak{s}_z(\psi)=\mathfrak{s},\mu(\psi)=0\}}\left(\bigcup_{\tau\in[0,1]}\mathcal{M}_{J_\tau}(\psi)\right)$$

count

$$f^\infty_{J_0}([\mathbf{x},i]\otimes[\mathbf{y},j];\mathfrak{s})-f^\infty_{J_1}([\mathbf{x},i]\otimes[\mathbf{y},j];\mathfrak{s})+\partial\circ H^\infty([\mathbf{x},i]\otimes[\mathbf{y},j])+H^\infty\circ\partial([\mathbf{x},i]\otimes[\mathbf{y},j]);$$

i.e. $f^\infty_{J_0}$ and $f^\infty_{J_1}$ are chain homotopic.

Since the induced maps are invariant under small perturbations of the family J, it follows also that the induced map is independent under variations of the complex structure $\mathfrak{j}$ over Σ. $\square$

Proposition 6.14. *The maps on homology listed in Theorem 6.12 are invariant under isotopies of the $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ preserving all the admissibility hypotheses.*

Proof. We begin with isotopies of the $\boldsymbol{\alpha}$. As in the proof of isotopy invariance of Floer homologies, we let Ψ_τ be an isotopy (induced from an exact Hamiltonian isotopy of the $\boldsymbol{\alpha}$ in Σ), and we consider moduli spaces with dynamic boundary conditions. Specifically, let $E_\alpha\colon \mathbb{R} \longrightarrow \Delta$ be a parameterization of the edge e_α, with

$$\lim_{t\mapsto -\infty} E_\alpha(t) = v_\gamma \quad \text{and} \quad \lim_{t\mapsto +\infty} E_\alpha(t) = v_\beta$$

Consider moduli spaces indexed by a real parameter $\tau \in \mathbb{R}$:

$$\mathcal{M}_\tau = \left\{ u\colon \Delta \longrightarrow \mathrm{Sym}^g(\Sigma) \middle| \begin{array}{c} u\circ E_\alpha(t) \in \Psi_{t+\tau}(\mathbb{T}_\alpha) \\ u(e_\beta)\subset \mathbb{T}_\beta, u(e_\gamma)\subset\mathbb{T}_\gamma \end{array} \right\},$$

and divide them into homotopy classes $\pi_2^{\Psi_t}(\mathbf{x},\mathbf{y},\mathbf{w})$, with $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, $\mathbf{y}\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma$, $\mathbf{w}\in\mathbb{T}_\gamma\cap\Psi_1(\mathbb{T}_\alpha)$.

Note that if $\mu(\psi)=-1$, then $\bigcup_{\tau\in\mathbb{R}}\mathcal{M}_\tau(\psi)$ is generically a compact zero-dimensional manifold, so we can define

$$H^\infty([\mathbf{x},i]\otimes[\mathbf{y},j];\mathfrak{s}) = \sum_{\mathbf{w}} \sum_{\{\psi\in\Psi_t|\mu(\psi)=-1,\mathfrak{s}_z(\psi)=\mathfrak{s}\}} \left(\#\bigcup_{\tau\in\mathbb{R}}\mathcal{M}_\tau(\psi)\right)[\mathbf{w}, i+j-n_z(\psi)].$$

Fix, now, any homotopy class $\psi\in\pi_2^{\Psi_t}(\mathbf{x},\mathbf{y},\mathbf{w})$ with $\mathfrak{s}_z(\psi)=\mathfrak{s}$ and $\mu(\psi)=0$, and consider the one-manifold

$$\bigcup_{\tau\in\mathbb{R}}\mathcal{M}_\tau(\psi).$$

This has ends as $\tau\mapsto\pm\infty$, which are modeled on

$$\left(\coprod_{\phi_{\alpha,\beta}*\psi_{\alpha,\beta}=\psi}\mathcal{M}_{\Psi_t}(\phi_{\alpha,\beta})\times\mathcal{M}(\psi_{\alpha,\beta})\right)\coprod\left(\coprod_{\phi_{\alpha,\gamma}*\psi_{\alpha,\gamma}=\psi}\mathcal{M}_{\Psi_t}(\phi_{\alpha,\gamma})\times\mathcal{M}(\psi_{\alpha,\gamma})\right),$$

where the first union is over all $\mathbf{x}'\in\Psi_1(\mathbb{T}_\alpha)\cap\mathbb{T}_\beta$ with $\mathfrak{s}_z(\mathbf{x}')=\mathfrak{s}_{\alpha,\beta}$, $\phi_{\alpha,\beta}\in\pi_2^{\Psi_t}(\mathbf{x},\mathbf{x}')$ (in the sense of Subsection 4.7), $\psi_{\alpha,\beta}\in\pi_2(\mathbf{x}',\mathbf{y},\mathbf{w})$, and $\mu(\phi_{\alpha,\beta})=\mu(\psi_{\alpha,\beta})=0$ (with analogous conditions on the second union). There are also ends of the form

$$\begin{array}{c}\left(\coprod_{\phi_{\alpha,\beta}*\psi_{\alpha,\beta}=\psi}\widehat{\mathcal{M}}(\phi_{\alpha,\beta})\times\left(\bigcup_{\tau\in\mathbb{R}}\mathcal{M}_\tau(\psi_{\alpha,\beta})\right)\right)\\ \coprod \\ \left(\coprod_{\phi_{\beta,\gamma}*\psi_{\beta,\gamma}=\psi}\widehat{\mathcal{M}}(\phi_{\beta,\gamma})\times\left(\bigcup_{\tau\in\mathbb{R}}\mathcal{M}_\tau(\psi_{\beta,\gamma})\right)\right)\\ \coprod \\ \left(\coprod_{\phi_{\alpha,\gamma}*\psi_{\gamma,\alpha}=\psi}\widehat{\mathcal{M}}(\phi_{\alpha,\gamma})\times\left(\bigcup_{\tau\in\mathbb{R}}\mathcal{M}_\tau(\psi_{\alpha,\gamma})\right)\right)\end{array},$$

where the first union is over all $\mathbf{x}' \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ in the same equivalence class as $\mathbf{x}$, $\phi_{\alpha,\beta} \in \pi_2(\mathbf{x}, \mathbf{x}')$ (in the sense of Subsection 4.7), $\psi_{\alpha,\beta} \in \pi_2^{\Psi_t}(\mathbf{x}', \mathbf{y}, \mathbf{w})$, and $\mu(\phi_{\alpha,\beta}) = 1$ and $\mu(\phi_{\alpha,\beta}) = 1$ while $\mu(\psi_{\alpha,\beta}) = -1$ (with analogous conditions over the other two unions). Counting ends with sign, we get that

$$\Gamma_{\alpha,\alpha',\gamma} \circ f_{\alpha,\beta,\gamma} + f_{\alpha',\beta,\gamma} \circ \Gamma_{\alpha,\alpha'\beta} = \partial \circ H + H \circ \partial,$$

where

$$\Gamma_{\alpha,\alpha',\beta} \colon CF(\mathbb{T}_\alpha, \mathbb{T}_\beta) \longrightarrow CF(\mathbb{T}'_\alpha, \mathbb{T}_\beta) \quad \text{and} \quad \Gamma_{\alpha,\alpha',\gamma} \colon CF(\mathbb{T}_\alpha, \mathbb{T}_\gamma) \longrightarrow CF(\mathbb{T}'_\alpha, \mathbb{T}_\gamma)$$

are the chain maps induced by the isotopy Ψ_t, as constructed in Subsection 4.7 (note that here we have suppressed the isotopy Ψ_t from the notation.

Isotopies of the $\boldsymbol{\gamma}$ work the same way; we now set up isotopies of the $\boldsymbol{\beta}$. Consider moduli spaces indexed by a real $\tau \in [0, \infty)$

$$\mathcal{M}_\tau = \left\{ u \colon \Delta \longrightarrow \mathrm{Sym}^g(\Sigma) \middle| \begin{array}{c} u \circ E_\gamma(t) \in \Psi^{-1}_{t+\tau} \circ \Phi_{t-\tau}(\mathbb{T}_\gamma) \\ u(e_\alpha) \subset \mathbb{T}_\alpha, u(e_\gamma) \subset \mathbb{T}_\gamma \end{array} \right\}.$$

These moduli spaces partition according to homotopy classes $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ (with $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\mathbf{y} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, $\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$). Note that for $\tau = 0$, this is the usual moduli space for holomorphic triangles for $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$. Again, when $\mu(\psi) = 0$, the union $\bigcup_{\tau \in [0,\infty)} \mathcal{M}_\tau(\psi)$ is generically a compact, zero-dimensional manifold, and we can define Note that if $\mu(\psi) = -1$, then $\bigcup_{\tau \in [0,\infty)} \mathcal{M}_\tau(\psi)$ is generically a compact zero-dimensional manifold, so we can define

$$H^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j]; \mathfrak{s}) = \sum_{\mathbf{w}} \sum_{\{\psi \in \Psi_t | \mu(\psi) = -1, \mathfrak{s}_z(\psi) = \mathfrak{s}\}} \left(\# \bigcup_{\tau \in [0,\infty)} \mathcal{M}_\tau(\psi) \right) [\mathbf{w}, i + j - n_z(\psi)].$$

Fix a homotopy class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ with $\mathfrak{s}_z(\psi) = \mathfrak{s}$ and $\mu(\psi) = 0$, and consider the ends of

$$\bigcup_{\tau \in [0,\infty)} \mathcal{M}_\tau(\psi).$$

The ends as $\tau \mapsto \infty$ are modeled on

$$\bigcup_{\phi_{\alpha,\beta} * \psi * \phi_{\beta,\gamma}} \mathcal{M}_{\Psi_t}(\phi_{\alpha,\beta}) \times \mathcal{M}_{\Psi_t}(\phi_{\beta,\gamma}) \times \mathcal{M}(\psi_{\alpha,\beta',\gamma}),$$

where the union is over all $\mathbf{x}' \in \mathbb{T}_\alpha \cap \Psi_1(\mathbb{T}_\beta)$, $\mathbf{y}' \in \Psi_1(\mathbb{T}_\beta) \cap \mathbb{T}_\gamma$ with $\mathfrak{s}_z(\mathbf{y}') \in \mathfrak{s}_{\beta',\gamma}$, $\phi_{\alpha,\beta} \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{x}')$, $\phi_{\beta,\gamma} \in \pi_2^{\Psi_t}(\mathbf{y}, \mathbf{y}')$, and $\psi_{\alpha,\beta',\gamma} \in \pi_2(\mathbf{x}, \mathbf{y}', \mathbf{w})$ with $\mu(\phi_{\alpha,\beta}) = \mu(\phi_{\beta,\gamma}) = \mu(\psi_{\alpha,\beta,\gamma}) = 0$. Counting these ends with sign, we get a contribution of

$$f_{\alpha,\beta',\gamma} \circ (\Gamma_{\alpha,\beta,\beta'}([\mathbf{x}, i]) \otimes \Gamma_{\beta,\beta',\gamma}([\mathbf{y}, j])),$$

while the end as $\tau \mapsto \infty$ corresponds simply to $f_{\alpha,\beta,\gamma}([\mathbf{x}, i] \otimes [\mathbf{y}, j])$. There are other ends as before, whose contribution is

$$\partial \circ H^\infty + H^\infty \circ \partial.$$

Thus, we have exhibited a chain homotopy from $f_{\alpha,\beta,\gamma}$ with $f_{\alpha,\beta',\gamma} \circ (\Gamma_{\alpha,\beta,\beta'} \otimes \Gamma_{\beta,\beta',\gamma})$. $\square$

6.3.1. *Associativity.* The map induced by triangles satisfies an associativity property, using holomorphic squares.

Specifically, fix a pointed Heegaard quadruple

$$(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z),$$

and let $X_{\alpha,\beta,\gamma,\delta}$ be the corresponding cobordism. We have, of course, restriction maps:

$$\mathrm{Spin}^c(X_{\alpha,\beta,\gamma,\delta}) \longrightarrow \mathrm{Spin}^c(X_{\alpha,\beta,\gamma}) \times \mathrm{Spin}^c(X_{\alpha,\gamma,\delta}),$$

which correspond to splitting the cobordism along an embedded copy of $Y_{\alpha,\gamma}$. There is a subgroup $\delta H^1(Y_{\alpha,\gamma}) \subset H^2(X_{\alpha,\beta,\gamma,\delta})$, whose orbits on $\mathrm{Spin}^c(X_{\alpha,\beta,\gamma,\delta})$ are the fibers of this restriction map. Similarly, we have a restriction map

$$\mathrm{Spin}^c(X_{\alpha,\beta,\gamma,\delta}) \longrightarrow \mathrm{Spin}^c(X_{\alpha,\beta,\delta}) \times \mathrm{Spin}^c(X_{\beta,\gamma,\delta}),$$

which corresponds to splitting along $Y_{\beta,\delta}$. In view of this, we will find it convenient to fix not a single Spin^c structure over $X_{\alpha,\beta,\gamma,\delta}$, but rather a $\delta H^1(Y_{\beta,\delta}) + \delta H^1(Y_{\alpha,\gamma})$ orbit.

There are notions of admissibility for Heegaard quadruples (and, in general, multi-diagrams), which generalize the corresponding notions for triangles. For instance, a Heegaard quadruple $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z)$ is called *weakly admissible* if every periodic domain which can be written as sums of doubly-periodic domains for $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, $Y_{\gamma,\delta}$, and $Y_{\alpha,\delta}$ has both positive and negative coefficients. Existence is achieved by winding, as in Section 5.

Strong admissibility requires fixing a $\delta H^1(Y_{\beta,\delta}) + \delta H^1(Y_{\alpha,\gamma})$-orbit $\mathfrak{S}$ of a fixed Spin^c structure over $X_{\alpha,\beta,\gamma,\delta}$. We say that a Heegaard quadruple is *strongly admissible* for the orbit $\mathfrak{S}$ if for each Spin^c structure $\mathfrak{s} \in \mathfrak{S}$ and each quadruply-periodic domain which can be written as a sum of doubly-periodic domains:

$$\mathcal{P} = \sum_{\{\xi,\eta\}\subset\{\alpha,\beta,\gamma,\delta\}} \mathcal{D}_{\xi,\eta} \tag{8}$$

with the property that

$$\sum \langle c_1(\mathfrak{s}_{\xi,\eta}), H(\mathcal{D}_{\xi,\eta})\rangle = 2n \geq 0,$$

(i.e. where here $\mathfrak{s}_{\xi,\eta}$ denotes the corresponding restriction of $\mathfrak{s}$), it follows that some local multiplicity of $\mathcal{P}$ is strictly greater than n. Note that this notion involves the orbit $\mathfrak{S}$ only through its restrictions to the six three-manifolds $Y_{\xi,\eta}$, for all subsets $\{\xi,\eta\} \subset \{\alpha,\beta,\gamma,\delta\}$. Note also that if a Heegaard quadruple is $\mathfrak{S}$-strongly admissible, then the associated Heegaard diagrams for $Y_{\xi,\eta}$ (for all $\{\xi,\eta\} \subset \{\alpha,\beta,\gamma,\delta\}$) are automatically strongly $\mathfrak{s}_{\xi,\eta}$-admissible.

Since the orbit $\mathfrak{S}$ *a priori* may contain infinitely many Spin^c structures, it might be impossible to achieve strong admissibility. However, suppose that the Heegaard quadruple satisfies the hypothesis that:

$$\delta H^1(Y_{\beta,\delta})|Y_{\alpha,\gamma} = 0 \quad \text{and} \quad \delta H^1(Y_{\alpha,\gamma})|Y_{\beta,\delta} = 0. \tag{9}$$

In this case, if we fix any $\{\xi,\eta\} \subset \{\alpha,\beta,\gamma,\delta\}$, the restriction $\mathfrak{s}_{\xi,\eta}$ to $Y_{\xi,\eta}$ of any $\mathfrak{s} \in \mathfrak{S}$ is independent of the choice of $\mathfrak{s}$. Equivalently, if we choose any quadruply-periodic domain which can be written as a sum of doubly-periodic domains as in Equation (8) we have that

$$\sum_{\{\xi,\eta\}\subset\{\alpha,\beta,\gamma,\delta\}} \langle c_1(\mathfrak{s}_{\xi,\eta}), H(\mathcal{D}_{\xi,\eta})\rangle$$

is a function of the periodic domain $\mathfrak{Q}$ and the orbit $\mathfrak{S}$ (i.e. it is independent of the choice of Spin^c structure $\mathfrak{s}\in\mathfrak{S}$). Thus, the proof of Lemma 5.2 adapts immediately to show that $\mathfrak{S}$-strong admissibility can be achieved.

When working with quadruples, and $\mathbb{Z}$ coefficients, we need yet another generalization of the notion of coherent systems of orientations. We now fix a $\delta H^1(Y_{\beta,\delta})+\delta H^1(Y_{\alpha,\gamma})$ orbit in Spin$^c(X_{\alpha,\beta,\gamma,\delta})$, which we denote $\mathfrak{S}$. A coherent system of orientations for $\mathfrak{S}$, then, is a collection of non-vanishing sections indexed by subsets $\{\xi_1,...,\xi_\ell\}\subset\{\alpha,\beta,\gamma,\delta\}$ with $\ell=2,3,4$, $\mathfrak{o}_{\xi_1,...,\xi_\ell}(\phi_{\xi_1,...,\xi_\ell})$, for the determinant line bundle defined over the homotopy class of polygons $\phi_{\xi_1,...,\xi_\ell}$ (i.e. this can be a Whitney disk, triangle, or square) representing the restriction of some $\mathfrak{s}\in\mathfrak{S}$ to Y_{ξ_1,ξ_2} when $\ell=2$ or $X_{\xi_1,...,\xi_\ell}$ if $\ell=3,4$. These are required to be compatible with the gluings in the sense that

$$\mathfrak{o}_{\xi_1,...,\xi_\ell}(\phi_{\xi_1,...,\xi_\ell})\wedge\mathfrak{o}_{\eta_1,...,\eta_m}(\phi_{\eta_1,...,\eta_m})=\mathfrak{o}_{\xi_1,...,\xi_\ell,\eta_1,...,\eta_m}(\phi_{\xi_1,...,\xi_\ell}*\phi_{\eta_1,...,\eta_m}),$$

under gluing maps which are defined whenever we have subsets $\{\xi_1,...,\xi_\ell\}$ and $\{\eta_1,...,\eta_m\}$ with two elements, say ξ_1 and ξ_2, in common, and for which the polygons $\phi_{\xi_1,...,\xi_\ell}$ and $\phi_{\eta_1,...,\eta_m}$ meet in a single intersection point for $\mathbb{T}_{\eta_1}\cap\mathbb{T}_{\eta_2}$.

Following the lines of Lemma 6.7, one can build up such a coherent system. Observe first that since

$$\delta H^1(Y_{\beta,\delta})+\delta H^1(Y_{\alpha,\gamma})|\partial X_{\alpha,\beta,\gamma,\delta}\equiv 0,$$

for each given $\mathfrak{S}$, restriction to each boundary component uniquely determines Spin^c structures over these boundary components. Start with three orientation systems $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\gamma}$, $\mathfrak{o}_{\gamma,\delta}$ for three of these boundary components, and two systems $\mathfrak{o}_{\alpha,\beta,\gamma}$, $\mathfrak{o}_{\alpha,\gamma,\delta}$ for Spin^c structures obtained by restricting any given $\mathfrak{s}\in\mathfrak{S}$, which are compatible with the orientation conventions used of the three-manifold boundaries. Next, fix some arbitrary orientation for some square φ_0 representing some Spin^c structure in $\mathfrak{S}$. The compatibility conditions then impose some restrictions on the orientation conventions for $\mathfrak{o}_{\beta,\gamma,\delta}$ and $\mathfrak{o}_{\alpha,\gamma,\delta}$, but it is easy to see that these conditions are consistent.

The following is an elaboration of Lemma 5.11 of [23]:

Theorem 6.15. *Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z)$ be a pointed Heegaard quadruple which is strongly $\mathfrak{S}$-admissible, where $\mathfrak{S}$ is a $\delta H^1(Y_{\beta,\delta})+\delta H^1(Y_{\alpha,\gamma})$-orbit in* Spin$^c(X_{\alpha,\beta,\gamma,\delta})$. *Then, we have*

$$\begin{aligned}&\sum_{\mathfrak{s}\in\mathfrak{S}}F^*_{\alpha,\gamma,\delta}(F^*_{\alpha,\beta,\gamma}(\xi_{\alpha,\beta}\otimes\theta_{\beta,\gamma};\mathfrak{s}_{\alpha,\beta,\gamma})\otimes\theta_{\gamma,\delta};\mathfrak{s}_{\alpha,\gamma,\delta})\\ &\quad=\sum_{\mathfrak{s}\in\mathfrak{S}}F^*_{\alpha,\beta,\delta}(\xi_{\alpha,\beta}\otimes F^{\leq 0}_{\beta,\gamma,\delta}(\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta};\mathfrak{s}_{\beta,\gamma,\delta});\mathfrak{s}_{\alpha,\beta,\delta}),\end{aligned}$$

where $F^*=F^\infty$, F^+ or F^-, $\xi_{\alpha,\beta}\in HF^*(Y_{\alpha,\beta})$, $\theta_{\beta,\gamma}$ and $\theta_{\gamma,\delta}$ lie in $HF^{\leq 0}(Y_{\beta,\gamma})$ and $HF^{\leq 0}(Y_{\gamma,\delta})$ respectively; also,

$$\begin{aligned}&\sum_{\mathfrak{s}\in\mathfrak{S}}\widehat{F}_{\alpha,\gamma,\delta}(\widehat{F}_{\alpha,\beta,\gamma}(\xi_{\alpha,\beta}\otimes\xi_{\beta,\gamma};\mathfrak{s}_{\alpha,\beta,\gamma})\otimes\xi_{\gamma,\delta};\mathfrak{s}_{\alpha,\gamma,\delta})\\ &\quad=\sum_{\mathfrak{s}\in\mathfrak{S}}\widehat{F}_{\alpha,\beta,\delta}(\xi_{\alpha,\beta}\otimes\widehat{F}_{\beta,\gamma,\delta}(\xi_{\beta,\gamma}\otimes\xi_{\gamma,\delta};\mathfrak{s}_{\beta,\gamma,\delta});\mathfrak{s}_{\alpha,\beta,\delta}),\end{aligned}$$

where now $\xi_{\alpha,\beta}$, $\xi_{\beta,\gamma}$, and $\xi_{\alpha,\gamma}$ lie in $\widehat{HF}$ for the corresponding three-manifolds. When working over $\mathbb{Z}$, we assume a consistent family of orientations for all the Spin^c *structures in $\mathfrak{S}$, used in the definitions of the maps on triangles.*

Proof. We define a map

$$H^{\infty}(\ \cdot\ ,\mathfrak{S})\colon \bigoplus_{\mathfrak{s}\in\mathfrak{S}} CF^{\infty}(Y_{\alpha,\beta},\mathfrak{s}_{\alpha,\beta})\otimes CF^{\infty}(Y_{\beta,\gamma},\mathfrak{s}_{\beta,\gamma})\otimes CF^{\infty}(Y_{\gamma,\delta},\mathfrak{s}_{\gamma,\delta}) \longrightarrow \bigoplus_{\mathfrak{s}\in\mathfrak{S}} CF^{\infty}(Y_{\alpha,\delta})$$

by

$$H^{\infty}([\mathbf{x},i]\otimes[\mathbf{y},j]\otimes[\mathbf{w},k],\mathfrak{S}) = \sum_{\mathbf{p}\in\mathbb{T}_\alpha\cap\mathbb{T}_\delta} \sum_{\{\varphi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w},\mathbf{p})\big|\mathfrak{s}_z(\varphi)\in\mathfrak{S},\mu(\varphi)=0\}} \Big(\#\mathcal{M}(\varphi)\Big)[\mathbf{p},i+j+k-n_z(\varphi)].$$

Note that above map is a finite sum by the strong admissibility requirement on the Heegaard quadruple; indeed, it also implies that the sums appearing in the statement of the theorem are finite sums.

Counting ends of one-dimensional moduli spaces $\mathcal{M}(\varphi)$ with $\mu(\varphi) = 1$, we see that H induces a chain homotopy between the maps

$$\xi_{\alpha,\beta}\otimes\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta}\mapsto \sum_{\mathfrak{s}\in\mathfrak{S}} f_{\alpha,\gamma,\delta}(f_{\alpha,\beta,\gamma}(\xi_{\alpha,\beta}\otimes\theta_{\beta,\gamma},\mathfrak{s}_{\alpha,\beta,\gamma})\otimes\theta_{\gamma,\delta},\mathfrak{s}_{\alpha,\gamma,\delta})$$

and

$$\xi_{\alpha,\beta}\otimes\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta}\mapsto \sum_{\mathfrak{s}\in\mathfrak{S}} f_{\alpha,\beta,\delta}(\xi_{\alpha,\beta}\otimes f_{\beta,\gamma,\delta}(\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta},\mathfrak{s}_{\beta,\gamma,\delta}),\mathfrak{s}_{\alpha,\beta,\delta}).$$

Again, the other cases are established in the same manner.

□

6.4. **Handleslide invariance.** Here we describe the handleslide invariance of the homology groups. In fact, the proof when $b_1(Y) = 0$ actually applies in general (c.f. Section 5 of [23]). We can now, however, put the proof into some perspective.

Recall that to prove handleslide invariance, one starts with a pointed Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$, and considers additional circles $\boldsymbol{\gamma}$ obtained from $\boldsymbol{\beta}$ by a handleslide, and then $\boldsymbol{\delta}$, obtained from $\boldsymbol{\beta}$ by a small perturbation. Observe that the corresponding cobordism $X_{\alpha,\beta,\gamma}$ describes a cobordism between $-Y$, Y, and $\#^g(S^1\times S^2)$ (indeed, it is diffeomorphic to the cobordism of Example 6.1). One calculates the homologies $HF^{\infty}(\mathbb{T}_\beta,\mathbb{T}_\gamma)$, $HF^{\infty}(\mathbb{T}_\gamma,\mathbb{T}_\delta)$, and $HF^{\infty}(\mathbb{T}_\delta,\mathbb{T}_\beta)$, and shows that (for an appropriate choice of coherent orientations) all are isomorphic to $H_*(T^g;\mathbb{Z})\otimes\mathbb{Z}[U,U^{-1}]$ (Lemma 5.7 for $\widehat{HF}$, and Lemma 5.13 for CF^{∞}, both in [23]). As such, they all come with "top dimensional" generators $\theta_{\beta,\gamma}$, $\theta_{\gamma,\delta}$, and $\theta_{\beta,\delta}$ respectively. In fact, we will view them as elements in $HF^{\leq 0}$. These elements are related as follows. For the Spin^c structure $\mathfrak{s}_0$ over $(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z)$, $F^{\leq 0}(\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta}) = \theta_{\beta,\delta}$ (see Lemma 5.10 of [23]).

Now, (see Lemma 5.12 of [23]) the generator $\theta_{\beta,\delta}$ has the property that the map

$$HF(\mathbb{T}_\alpha,\mathbb{T}_\beta)\longrightarrow HF(\mathbb{T}_\alpha,\mathbb{T}_\delta)$$

given by tensoring with $\theta_{\beta,\delta}$ is an isomorphism: Recall that this is proved by isotoping the $\boldsymbol{\delta}$ to the $\boldsymbol{\beta}$ (thought of as a one-parameter family $\delta_i(\tau)$ with $\delta_i(0)=\beta_i$), and noticing that in the limit, the only contributing triangles are the canonical ones in $\psi_0\in\pi_2(\mathbf{x},\Theta_{\beta,\delta},\mathbf{x}')$ (where $\mathbf{x}'\in\mathbb{T}_\alpha\cap\mathbb{T}_\delta(\tau)$ is the point closest to $\mathbb{T}_\alpha\cap\mathbb{T}_\beta$). The others are excluded by dimension counts, which we must present slightly differently here (since now we have no absolute gradings at our disposal). Suppose that there is another family of triangles $\psi_\tau\in\pi_2(\mathbf{x},\Theta_{\beta,\delta},\mathbf{y}')$ (with $\mathbf{y}'\neq\mathbf{x}'$) with $\mu(\psi_\tau)=0$. Then there are corresponding $\phi_\tau\in\pi_2(\mathbf{x}',\mathbf{y}')$ with $\psi_\tau=\psi_0*\phi_t$. It follows that $\mu(\phi_\tau)=0$, so that ϕ_τ converge to a flow in $\phi_0\in\pi_2(\mathbf{x},\mathbf{y})$ with $\mu(\phi_0)=0$. But these moduli spaces are empty for generic choices. (Note that the Maslov index depends only on the underlying domain and the endpoints, and the domains of the ϕ_τ degenerate to the domain of ϕ_0.)

We make some remarks on the sign conventions. Observe that in the present case, $H^1(X_{\alpha,\beta,\gamma})$ maps isomorphically onto $H^1(Y_{\alpha,\gamma})$. Thus, from the proof of Lemma 6.7, we see that if we have fixed a coherent system of coherent orientations over $Y_{\alpha,\beta}$ (which we do arbitrarily), and another system of coherent orientations over $Y_{\beta,\delta}$ (which we select so that $\theta_{\beta,\delta}$ is a cycle), the coherent system of orientations over $Y_{\alpha,\delta}$ is uniquely defined up to equivalence (in the sense of Subsection 4.9). Indeed, this sets up a one-to-one correspondence between coherent orientation systems over $Y_{\alpha,\beta}$ and $Y_{\alpha,\delta}$ (up to isomorphism).

It follows from these observations, together with associativity that the map

$$HF(\mathbb{T}_\alpha,\mathbb{T}_\beta)\longrightarrow HF(\mathbb{T}_\alpha,\mathbb{T}_\gamma)$$

given by tensoring with $\theta_{\beta,\gamma}$ induces an isomorphism.

In this application of associativity, we observe that

$$\delta H^1(Y_{\beta,\delta})+\delta H^1(Y_{\alpha,\gamma})=0.$$

To see this, note that all quadruply-periodic domains for our Heegaard quadruple can be written as sums of doubly-periodic domains for the four bounding three-manifolds (this is an easy consequence of the fact that the spans in H_1 of the three g-tuples $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\delta}$ all coincide). Since the map from doubly-periodic domains to quadruply-periodic domains models the map $H^1(\partial X_{\alpha,\beta,\gamma,\delta})$ to $H^2(X_{\alpha,\beta,\gamma,\delta},\partial X_{\alpha,\beta,\gamma,\delta})$, it follows that the map $H^2(X_{\alpha,\beta,\gamma,\delta})\longrightarrow H^2(\partial X_{\alpha,\beta,\gamma,\delta})$ is injective. Since the restriction of $\delta H^1(Y_{\beta,\delta})+\delta^1(Y_{\alpha,\gamma})$ to the boundary is trivial, it follows that the subgroup itself is trivial.

Proposition 6.16. *The map*

$$HF^*(\mathbb{T}_\alpha,\mathbb{T}_\beta)\longrightarrow HF^*(\mathbb{T}_\alpha,\mathbb{T}_\gamma)$$

given by $F^*(\cdot\otimes\theta_{\beta,\gamma})$ commutes with the action by $H_1(Y;\mathbb{Z})$.

Proof. We represent the action by a codimension-one constraint $V\in\mathbb{T}_\alpha$. We then consider the moduli space

$$\mathcal{M}_V(\psi)=\bigcup_{\tau\in\mathbb{R}}\left\{u\in\mathcal{M}(\psi)\middle|u\circ E_\alpha(\tau)\in V\right\},$$

As usual, when $\mu(\psi)=-1$, this space is compact, and can be used to construct a chain homotopy

$$H([\mathbf{x},i]\otimes[\mathbf{y},j])=\sum_{\mathbf{w}}\sum_{\psi}\#\Big(\mathcal{M}_V(\psi)\Big)[\mathbf{w},i+j-n_z(\psi)]$$

Consider homotopy classes with $\mu(\psi)=0$. The ends as $\tau\mapsto\infty$ correspond to the commutator of F^* with the action of V; the other ends correspond to the commutator of H with the boundary maps.

□

6.5. **Twisted coefficients.** We discuss how to extend the previous results to the case of twisted coefficient systems. We use a refinement of the notion of Spin^c structures, that of *relative* Spin^c *structures*, which we define in a manner which will be most convenient for the applications.

6.5.1. *Relative* Spin^c *structures.* Suppose that $X_{\alpha,\beta,\gamma}$ is a cobordism between $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, and $Y_{\alpha,\gamma}$, and fix Spin^c structures $\mathfrak{t}_{\alpha,\beta}$, $\mathfrak{t}_{\beta,\gamma}$, $\mathfrak{t}_{\alpha,\gamma}$ over the three boundary components, with $\epsilon(\mathfrak{t}_{\alpha,\beta},\mathfrak{t}_{\beta,\gamma},\mathfrak{t}_{\alpha,\gamma})\neq 0$. Fixing complete sets of paths for each of these three Spin^c structures (in the sense of Definition 3.4). This gives us identifications

$$\pi_2(\mathbf{x}_0,\mathbf{y}_0,\mathbf{w}_0)=\pi_2(\mathbf{x},\mathbf{y},\mathbf{w}),$$

where $\mathbf{x}_0$ and $\mathbf{x}$ (resp. $\mathbf{y}_0$ and $\mathbf{y}$, resp. $\mathbf{w}_0$ and $\mathbf{w}$) both represent $\mathfrak{t}_{\alpha,\beta}$ (resp. $\mathfrak{t}_{\beta,\gamma}$ resp. $\mathfrak{t}_{\alpha,\gamma}$).

In effect, this allows us to think of $\pi_2(\mathbf{x}_0,\mathbf{y}_0,\mathbf{w}_0)$ as an affine space for $H^2(X,Y;\mathbb{Z})$ (c.f. Proposition 6.3), which maps onto the space of Spin^c structures extending $\mathfrak{t}_{\alpha,\beta}$, $\mathfrak{t}_{\beta,\gamma}$, $\mathfrak{t}_{\alpha,\gamma}$ (c.f. Proposition 6.5). When thinking of $\pi_2(\mathbf{x}_0,\mathbf{y}_0,\mathbf{w}_0)$ in this manner, we refer to it as a space of relative Spin^c structures, and denote it by $\underline{\mathrm{Spin}^c}(X_{\alpha,\beta,\gamma})$.

The fiber of a fixed Spin^c structure $\mathfrak{s}_{\alpha,\beta,\gamma}$ will be denoted $\mathrm{Spin}^c(X_{\alpha,\beta,\gamma};\mathfrak{s}_{\alpha,\beta,\gamma})$.

We will use this terminology for higher polygons, as well.

6.5.2. *The maps with twisted coefficients.* The space of relative Spin^c structures $\underline{\mathrm{Spin}^c}(X_{\alpha,\beta,\gamma};\mathfrak{s}_{\alpha,\beta,\gamma})$ (which induce a given Spin^c structure $\mathfrak{s}_{\alpha,\beta,\gamma}$ over $X_{\alpha,\beta,\gamma}$) is a space with a natural action of $H^1(Y_{\alpha,\beta};\mathbb{Z})\times H^1(Y_{\beta,\gamma};\mathbb{Z})\times H^1(Y_{\alpha,\gamma};\mathbb{Z})$. As such, it can be used to induce an $H^1(Y_{\alpha,\gamma};\mathbb{Z})$-module from a pair $M_{\alpha,\beta}$ and $M_{\beta,\gamma}$ of $H^1(Y_{\alpha,\beta};\mathbb{Z})$ and $H^1(Y_{\beta,\gamma};\mathbb{Z})$-modules:

$$\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}}=\frac{(m_{\alpha,\beta},m_{\beta,\gamma},\underline{\mathfrak{s}})\in M_{\alpha,\beta}\times M_{\beta,\gamma}\times\underline{\mathrm{Spin}^c}(X_{\alpha,\beta,\gamma},\mathfrak{s}_{\alpha,\beta,\gamma})}{(m_{\alpha,\beta},m_{\beta,\gamma},\underline{\mathfrak{s}})\sim(h_{\alpha,\beta}\cdot m_{\alpha,\beta},h_{\beta,\gamma}\cdot m_{\beta,\gamma},(h_{\alpha,\beta}\times h_{\beta,\gamma}\times 0)\cdot\underline{\mathfrak{s}})},$$

where $h_{\alpha,\beta}$ and $h_{\beta,\gamma}$ are arbitrary elements of $H^1(Y_{\alpha,\beta};\mathbb{Z})$ and $H^1(Y_{\beta,\gamma};\mathbb{Z})$ respectively.

Fix a Spin^c structure $\mathfrak{s}$ over $X_{\alpha,\beta,\gamma}$, whose restriction to $Y_{\alpha,\beta}$ and $Y_{\beta,\gamma}$ is $\mathfrak{t}_{\alpha,\beta}$ and $\mathfrak{t}_{\beta,\gamma}$ respectively. We can now define a map

$$\begin{aligned}\underline{f^{\infty}_{\alpha,\beta,\gamma}}(\ \cdot\ ,\mathfrak{s})\colon \underline{CF^{\infty}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta};M_{\alpha,\beta})\otimes\underline{CF^{\infty}}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma};M_{\beta,\gamma})&\\ \longrightarrow\underline{CF^{\infty}}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}}),&\end{aligned}$$

by the formula:

$$\underline{f^{\infty}_{\alpha,\beta,\gamma}}(m_{\alpha,\beta}[\mathbf{x},i]\otimes m_{\beta,\gamma}[\mathbf{y},j];\mathfrak{s}) =$$
$$\sum_{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\ \sum_{\{\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\big|\mathfrak{s}_z(\psi)=\mathfrak{s}\}} (\#\mathcal{M}(\psi))\,\{m_{\alpha,\beta}\otimes m_{\beta,\gamma}\otimes \underline{\mathfrak{s}}_z(\psi)\}\cdot[\mathbf{w},i+j-n_z(\psi)]. \tag{10}$$

The braces above indicate the natural map

$$\{\cdot\otimes\cdot\otimes\cdot\}:M_{\alpha,\beta}\otimes M_{\beta,\gamma}\otimes\underline{\mathrm{Spin}}^c(X_{\alpha,\beta,\gamma},\mathfrak{s})\longrightarrow\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}}.$$

The following analogue of Theorem 6.12 holds in the present context:

Theorem 6.17. *Let* $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},z)$ *be a pointed Heegaard triple-diagram, which is strongly* $\mathfrak{s}$*-admissible for some* Spin^c *structure* $\mathfrak{s}$ *over the underlying four-manifold* X*, and fix modules* $M_{\alpha,\beta}$ *and* $M_{\beta,\gamma}$ *for* $H^1(Y_{\alpha,\beta};\mathbb{Z})$ *and* $H^1(Y_{\beta,\gamma};\mathbb{Z})$ *respectively. Then the sum on the right-hand-side of Equation* (10) *is finite, giving rise to a chain map which also induces maps on homology:*

$$\begin{aligned}\underline{F^{\infty}_{\alpha,\beta,\gamma}}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \underline{HF^{\infty}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta};M_{\alpha,\beta})\otimes\underline{HF^{\infty}}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma};M_{\beta,\gamma})&\\ \longrightarrow\underline{HF^{\infty}}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}})&\\ \underline{F^{\leq 0}_{\alpha,\beta,\gamma}}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \underline{HF^{\leq 0}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta};M_{\alpha,\beta})\otimes\underline{HF^{\leq 0}}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma};M_{\beta,\gamma})&\\ \longrightarrow\underline{HF^{\leq 0}}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}}).&\end{aligned}$$

The induced chain map

$$\begin{aligned}\underline{f^{+}_{\alpha,\beta,\gamma}}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \underline{CF^{+}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta};M_{\alpha,\beta})\otimes\underline{CF^{\leq 0}}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma};M_{\beta,\gamma})&\\ \longrightarrow\underline{CF^{+}}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}})&\end{aligned}$$

gives a well-defined chain map when the triple diagram is only weakly admissible, and the Heegaard diagram $(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},z)$ *is strongly admissible for* $\mathfrak{t}_{\beta,\gamma}$*. In fact, the induced map*

$$\begin{aligned}\underline{\widehat{f}_{\alpha,\beta,\gamma}}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \underline{\widehat{CF}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta};M_{\alpha,\beta})\otimes\underline{\widehat{CF}}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma};M_{\beta,\gamma})&\\ \longrightarrow\underline{\widehat{CF}}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}})&\end{aligned}$$

gives a well-defined chain map when the diagram is weakly admissible. There are induced maps on homology:

$$\begin{aligned}\underline{\widehat{F}_{\alpha,\beta,\gamma}}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \underline{\widehat{HF}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta};M_{\alpha,\beta})\otimes\widehat{HF}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma};M_{\beta,\gamma})&\\ \longrightarrow\widehat{HF}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}})&\\ F^{+}_{\alpha,\beta,\gamma}(\cdot,\mathfrak{s}_{\alpha,\beta,\gamma})\colon \underline{HF^{+}}(Y_{\alpha,\beta},\mathfrak{t}_{\alpha,\beta})\otimes\underline{HF}^{\leq 0}(Y_{\beta,\gamma},\mathfrak{t}_{\beta,\gamma})&\\ \longrightarrow\underline{HF^{+}}(Y_{\alpha,\gamma},\mathfrak{t}_{\alpha,\gamma};\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}})).&\end{aligned}$$

Independence of complex structure (Proposition 6.13) and isotopy invariance (Proposition 6.14) proceed as before. Associativity, on the other hand, can be given a sharper statement.

Observe first that there is a canonical gluing

$$\underline{\mathrm{Spin}^c}(X_{\alpha,\beta,\gamma},\mathfrak{s}_{\alpha,\beta,\gamma})\times\underline{\mathrm{Spin}^c}(X_{\alpha,\gamma,\delta},\mathfrak{s}_{\alpha,\gamma,\delta})\longrightarrow\underline{\mathrm{Spin}^c}(X_{\alpha,\beta,\gamma,\delta})$$

which maps onto the set of all relative Spin^c structures over $X_{\alpha,\beta,\gamma,\delta}$ whose restrictions to $X_{\alpha,\beta,\gamma}$ and $X_{\alpha,\gamma,\delta}$ represent Spin^c structures $\mathfrak{s}_{\alpha,\beta,\gamma}$ and $\mathfrak{s}_{\alpha,\gamma,\delta}$ respectively. Thus, the set of Spin^c induced structures in $X_{\alpha,\beta,\gamma,\delta}$ under this map consists of a $\delta H^1(Y;\mathbb{Z})$-orbit. Using this gluing, we obtain an identification

$$\begin{aligned}&\{\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\mathfrak{s}_{\alpha,\beta,\gamma}}\otimes M_{\gamma,\delta}\}^{\mathfrak{s}_{\alpha,\beta,\delta}}\\ &\cong \coprod_{\{\mathfrak{s}\in\mathrm{Spin}^c(X_{\alpha,\beta,\gamma,\delta})\big|\mathfrak{s}|X_{\alpha,\beta,\gamma}=\mathfrak{s}_{\alpha,\beta,\gamma},\mathfrak{s}|X_{\alpha,\gamma,\delta}=\mathfrak{s}_{\alpha,\gamma,\delta}\}}\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\otimes M_{\gamma,\delta}\}^{\mathfrak{s}},\end{aligned}$$

where $\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\otimes M_{\gamma,\delta}\}^{\mathfrak{s}}$ denotes the $H^1(Y_{\alpha,\delta};\mathbb{Z})$-module induced from $M_{\alpha,\beta}$, $M_{\beta,\gamma}$, $M_{\gamma,\delta}$ and the set of relative Spin^c structures inducing the given Spin^c structure $\mathfrak{s}$ over the four-manifold $X_{\alpha,\beta,\gamma,\delta}$.

Theorem 6.18. *Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z)$ be a pointed Heegaard quadruple which is strongly $\mathfrak{S}$-admissible, where $\mathfrak{S}$ is a $\delta H^1(Y_{\beta,\delta})+\delta H^1(Y_{\alpha,\gamma})$-orbit in $\mathrm{Spin}^c(X_{\alpha,\beta,\gamma,\delta})$. Fix also modules $M_{\alpha,\beta}$, $M_{\beta,\gamma}$, and $M_{\gamma,\delta}$ for $H^1(Y_{\alpha,\beta};\mathbb{Z})$, $H^1(Y_{\beta,\gamma};\mathbb{Z})$, $H^1(Y_{\beta,\gamma};\mathbb{Z})$, and $H^1(Y_{\gamma,\delta};\mathbb{Z})$ respectively. Then,*

$$\begin{aligned}&\sum_{\mathfrak{s}\in\mathfrak{S}}\underline{F^*_{\alpha,\gamma,\delta}}(\underline{F^*_{\alpha,\beta,\gamma}}(\xi_{\alpha,\beta}\otimes\theta_{\beta,\gamma};\mathfrak{s}_{\alpha,\beta,\gamma})\otimes\theta_{\gamma,\delta};\mathfrak{s}_{\alpha,\gamma,\delta})\\ &\quad=\sum_{\mathfrak{s}\in\mathfrak{S}}\underline{F^*_{\alpha,\beta,\delta}}(\xi_{\alpha,\beta}\otimes\underline{F^{\leq 0}_{\beta,\gamma,\delta}}(\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta};\mathfrak{s}_{\beta,\gamma,\delta});\mathfrak{s}_{\alpha,\beta,\delta}),\end{aligned}$$

where $\underline{F^*}=\underline{F^\infty}$, $\underline{F^+}$ or $\underline{F^-}$; also,

$$\begin{aligned}&\sum_{\mathfrak{s}\in\mathfrak{S}}\underline{\widehat{F}_{\alpha,\gamma,\delta}}(\underline{\widehat{F}_{\alpha,\beta,\gamma}}(\xi_{\alpha,\beta}\otimes\theta_{\beta,\gamma};\mathfrak{s}_{\alpha,\beta,\gamma})\otimes\theta_{\gamma,\delta};\mathfrak{s}_{\alpha,\gamma,\delta})\\ &\quad=\sum_{\mathfrak{s}\in\mathfrak{S}}\underline{\widehat{F}_{\alpha,\beta,\delta}}(\xi_{\alpha,\beta}\otimes\underline{\widehat{F}_{\beta,\gamma,\delta}}(\theta_{\beta,\gamma}\otimes\theta_{\gamma,\delta};\mathfrak{s}_{\beta,\gamma,\delta});\mathfrak{s}_{\alpha,\beta,\delta}),\end{aligned}$$

where we are taking coefficients in coefficients in $\coprod_{\mathfrak{s}\in\mathfrak{S}}\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\otimes M_{\gamma,\delta}\}^{\mathfrak{s}}$ over $Y_{\alpha,\delta}$.

Proof. The proof is the same as the proof of Theorem 6.15, only keeping track now of the homotopy classes of the corresponding triangles. □

6.5.3. *Handleslide invariance.* To adapt the previous proof of handleslide invariance to the case of twisted coefficient systems, a few remarks must be made.

As in Subsection 6.4 we consider the cobordism $X_{\alpha,\beta,\gamma}$, where the $\boldsymbol{\gamma}$ are obtained by $\boldsymbol{\beta}$ by handleslides, so that $(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma})$ describes $Y_{\beta,\gamma}=\#^g(S^1\times S^2)$, and $Y\cong Y_{\alpha,\beta}\cong Y_{\alpha,\gamma}$. Note that the universal element $\theta_{\beta,\gamma}$ lives in $HF^{\leq 0}$ (or $\widehat{HF}$) of $\#^g(S^1\times S^2)$ with untwisted coefficients, or, equivalently, the homology with coefficients in a trivial module.

Note, however, that if we let $M_{\beta,\gamma}$ be the trivial $H^1(Y_{\beta,\gamma})$-module, then there is a canonical identification of $H^1(Y;\mathbb{Z})$-modules

$$M \cong \{M \otimes M_{\beta,\gamma}\},$$

where the pairing here uses the cobordism $X_{\alpha,\beta,\gamma}$. Hence, multiplying by $\theta_{\beta,\gamma}$ indeed gives rise to a map

$$\underline{HF}(\mathbb{T}_\alpha, \mathbb{T}_\beta; M) \longrightarrow \underline{HF}(\mathbb{T}_\alpha, \mathbb{T}_\gamma; M)$$

(with the analogous statement for multiplication by $\theta_{\gamma,\delta}$ and $\theta_{\beta,\delta}$).

7. Basic properties

We collect here some properties of $\widehat{HF}$, HF^+, HF^-, and HF^∞ which follow easily from the definitions. Then, we turn to several of $\widehat{HF}$, which lead to Theorem 1.6 from the introduction. We also include a simple example, $S^2 \times S^1$ (for all four variants), which serves to illustrate some of the issues from the last section. In the final subsection, we describe the effect of connected sums with $S^2 \times S^1$ on HF^+, establishing Proposition 1.5.

7.1. **General properties.** Note that $\widehat{HF}$ and HF^+ distinguish certain Spin^c structures on Y – those for which the groups do not vanish.

Proposition 7.1. *For an oriented three-manifold Y with Spin^c structure $\mathfrak{s}$, $\widehat{HF}(Y,\mathfrak{s})$ is non-trivial if and only if $HF^+(Y,\mathfrak{s})$ is non-trivial (for the same orientation system).*

Proof. This follows from the natural long exact sequence:

$$... \longrightarrow \widehat{HF}(Y,\mathfrak{s}) \longrightarrow HF^+(Y,\mathfrak{s}) \stackrel{U}{\longrightarrow} HF^+(Y,\mathfrak{s}) \longrightarrow ...$$

induced from the short exact sequence of chain complexes

$$0 \longrightarrow \widehat{CF}(Y,\mathfrak{s}) \longrightarrow CF^+(Y,\mathfrak{s}) \stackrel{U}{\longrightarrow} CF^+(Y,\mathfrak{s}) \longrightarrow 0.$$

Now, observe that U is an isomorphism on $HF^+(Y,\mathfrak{s})$ if and only if the latter group is trivial, since each element of $HF^+(Y,\mathfrak{s})$ is annihilated by a sufficiently large power of U. □

Remark 7.2. *Indeed, the above proposition holds when we use an arbitrary coefficient ring. In particular, the rank of $HF^+(Y,\mathfrak{s})$ is non-zero if and only if the rank of $\widehat{HF}(Y,\mathfrak{s})$ is non-zero.*

Moreover, there are finitely many such Spin^c structures:

Theorem 7.3. *There are finitely many Spin^c structures $\mathfrak{s}$ for which $HF^+(Y,\mathfrak{s})$ is non-zero. The same holds for $\widehat{HF}(Y,\mathfrak{s})$.*

Proof. Consider a Heegaard diagram which is weakly $\mathfrak{s}$-admissible for all Spin^c structures (i.e. a diagram which is $\mathfrak{s}_0$-admissible Heegaard diagram, where $\mathfrak{s}_0$ is any torsion Spin^c structure, c.f. Remark 4.3 and, of course, Lemma 5.4). This diagram can be used to calculate HF^+ and $\widehat{HF}$ for all Spin^c-structures simultaneously. But the tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ have only finitely many intersection points, so there are only finitely many Spin^c structures for which the chain complexes $CF^+(Y,\mathfrak{s})$ and $\widehat{CF}(Y,\mathfrak{s})$ are non-zero. □

Theorem 7.4. *The Floer homologies HF^+, HF^-, HF^∞, and $\widehat{HF}$ are invariant under the involution $J\colon \mathrm{Spin}^c(Y) \longrightarrow \mathrm{Spin}^c(Y)$.*

Proof. Observe that if a Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$ is admissible for $\mathfrak{s}$, then the diagram $(-\Sigma,\boldsymbol{\beta},\boldsymbol{\alpha},z)$ is admissible for $J\mathfrak{s}$. The theorem then follows as in the $b_1(Y)=0$ case (c.f. Theorem 7.1 in [23]), noting again that the moduli spaces of flows are identified, although the assignment $\mathfrak{s}_z$ changes by the involution J. □

7.2. **On $\widehat{HF}$.** The following observation is clear from the definition of $\widehat{HF}(Y,\mathfrak{s})$ (c.f. Theorem 1.11 of Section 1):

Proposition 7.5. *Let Y be a rational homology three-sphere, then*

$$\mathrm{rk}\widehat{HF}(Y) = \sum_{\mathfrak{s}\in\mathrm{Spin}^c(Y)} \mathrm{rk}\widehat{HF}(Y,\mathfrak{s})$$

is a lower bound for the simultaneous trajectory number $M(Y)$ introduced in Subsection 1.2.

Proof. Clearly, the rank of the homology of $\widehat{CF}(Y)=\bigoplus_{\mathfrak{s}\in\mathrm{Spin}^c(Y)}\widehat{CF}(Y,\mathfrak{s})$ is a lower bound for the rank of $\widehat{CF}(Y)$, which in turn is the total number intersection points between $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$. □

Remark 7.6. *Again, this proposition works with arbitrary coefficient ring.*

Proposition 7.7. *The Euler characteristic of $\widehat{HF}$ is given by*

$$\chi(\widehat{HF}(Y,\mathfrak{s})) = \begin{cases} 1 & \text{if } b_1(Y)=0 \\ 0 & \text{if } b_1(Y)>0 \end{cases}.$$

Proof. Both cases follow from the observation that $\chi(\widehat{HF}(Y,\mathfrak{s}))$ is independent of the Spin^c structure $\mathfrak{s}$. To see this, note that for any β_j, we can wind normal to the $\boldsymbol{\alpha}$ so that $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$ and $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z')$ are both weakly $\mathfrak{s}$-admissible, where z and z' are two choices of basepoint which can be connected by an arc which meets only β_j. Now, both $\widehat{HF}(Y,\mathfrak{s})$ and $\widehat{HF}(Y,\mathfrak{s}+\mathrm{PD}[\beta_j^*])$ are calculated by the same equivalence class of intersection points, using the basepoint z in the first case and z' in the second. This changes only the boundary map, but leaves the (finitely generated) chain groups unchanged, hence leaving the Euler characteristic unchanged.

The result for $b_1(Y)>0$ then follows from this observation, together with Theorem 7.3.

For the case where $b_1(Y)=0$, recall that the Heegaard decomposition gives Y a chain complex with g one-dimensional generators corresponding to the $\boldsymbol{\alpha}$ (each of which is a cycle), and g two-dimensional generators corresponding to the $\boldsymbol{\beta}$. On the one hand, the determinant of the boundary map is the order of the finite group $H_1(Y;\mathbb{Z})$ (which, in turn, is the number of distinct Spin^c structures over Y); on the other hand, it is easily seen to agree with the intersection number $\#(\mathbb{T}_\alpha\cap\mathbb{T}_\beta)=\sum_{\mathfrak{s}\in\mathrm{Spin}^c(Y)}\chi(\widehat{HF}(Y,\mathfrak{s}))$. The result follows from this, together with $\mathfrak{s}$-independence of $\chi(\widehat{HF}(Y,\mathfrak{s}))$. □

Thus, we have the following:

Proposition 7.8. *Suppose that $b_1(Y)>0$. Then*

$$\mathrm{rk}\widehat{HF}(Y) \geq 2\#\{\mathfrak{s}\in\mathrm{Spin}^c(Y)\big|\mathrm{rk}HF^+(Y,\mathfrak{s})\neq 0\}.$$

Proof. This is a combination of Propositions 7.1 and 7.7. □

The behavior of $\widehat{HF}$ under connected sums is easy to understand. Recall that the connected sum of Spin^c structures over Y_1 and Y_2 is a Spin^c structure over $Y_1\#Y_2$, in fact giving rise to an identification:

$$\mathrm{Spin}^c(Y_1)\times\mathrm{Spin}^c(Y_2)\cong\mathrm{Spin}^c(Y_1\#Y_2).$$

Proposition 7.9. *Let Y_1 and Y_2 be a pair of oriented three-manifolds. Then $\widehat{HF}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)$ is the homology of the product of complexes calculating $\widehat{HF}(Y_1,\mathfrak{s}_1)$ and $\widehat{HF}(Y_2,\mathfrak{s}_2)$.*

Proof. Endow Y_1 and Y_2 with Heegaard diagrams which are weakly admissible for all Spin^c structures. Observe, then, that the connected sum diagram for $Y_1\#Y_2$ is weakly admissible for all Spin^c structures, too (when we take the connected sum points to be sufficiently close to the two base-points). With this observation in place, the proof proceeds exactly as the proof of Proposition 7.2 of [23]. □

Proof of Theorem 1.6. This is now a direct consequence of Propositions 7.1, 7.9, and the Künneth formula from homological algebra. □

Observe that, by contrast, if Y_1 and Y_2 are a pair of oriented three-manifolds with positive first Betti number, then the Alexander polynomial (or Turaev's torsion) of their connected sum $Y_1\#Y_2$ vanishes.

As we have pointed out, the results stated above hold over arbitrary coefficient rings. In particular, we can use homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$ (in which case, of course, the rank of the relevant homology group is its dimension as a vector space over $\mathbb{Z}/2\mathbb{Z}$). This allows us to bypass issues of orientation for the three-dimensional applications described in the introduction.

7.3. **Remarks on algebraic notation.** In the following example, and indeed, throughout this paper, we will describe modules over the graded ring $\mathbb{Z}[U]\otimes_{\mathbb{Z}}\Lambda^*(H_1(Y;\mathbb{Z})/\mathrm{Tors})$, graded so that U has grading two and non-zero elements of $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ have grading one. For example, $HF^+(Y,\mathfrak{s})$, $HF^-(Y,\mathfrak{s})$, and $HF^\infty(Y,\mathfrak{s})$ are all modules over this ring.

There are several other obvious graded modules over this ring, including:

$$\mathbb{Z}[U^{-1},U]\otimes_{\mathbb{Z}}\Lambda^*H^1(Y;\mathbb{Z}),$$

The U action is given by multiplying on the Laurent factor, and the action of $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ is given by contracting with the exterior product factor (noting that the Kronecker pairing identifies $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ with the dual of $H^1(Y;\mathbb{Z})$). This has a quotient module, which we write

$$\mathbb{Z}[U^{-1}]\otimes_{\mathbb{Z}}\Lambda^*H^1(Y;\mathbb{Z});$$

i.e. here we have polynomials in U^{-1}. The U-action again multiplies these polynomials, now dropping the new terms involving positive powers of U.

7.4. **A simple example: $S^1\times S^2$.** Consider the torus Σ with a homotopically non-trivial embedded curve α, and an isotopic translate β. The data (Σ,α,β) gives a Heegaard diagram for $S^1\times S^2$. (Actually, we met this example already in Section 5 of [23].)

We can choose the curves disjoint, dividing Σ into a pair of annuli. If the basepoint z lies in one annulus, the other annulus $\mathcal{P}$ is a periodic domain. Since there are no intersection points, one might be tempted to think that the homology groups are trivial; but this is not the case, as the Heegaard diagram is not weakly admissible for $\mathfrak{s}_0$, and also not strongly admissible for any Spin^c structure.

To make the diagram weakly admissible for the torsion Spin^c structure $\mathfrak{s}_0$, the periodic domain must have coefficients with both signs. This can be arranged by introducing canceling pairs of intersection points between α in β. The simplest such case occurs when there is only one pair of intersection points x^+ and x^-. There is now a pair of (non-homotopic) holomorphic disks connecting x^+ and x^- (both with Maslov index one), showing at once that

$$\widehat{HF}(S^1\times S^2,\mathfrak{s}_0)\cong H_*(S^1), \qquad HF^\infty(S^1\times S^2,\mathfrak{s}_0)\cong H_*(S^1)\otimes_{\mathbb{Z}}\mathbb{Z}[U,U^{-1}],$$
$$HF^+(S^1\times S^2,\mathfrak{s}_0)\cong H_*(S^1)\otimes_{\mathbb{Z}}\mathbb{Z}[U^{-1}] \qquad HF^-(S^1\times S^2,\mathfrak{s}_0)\cong H_*(S^1)\otimes_{\mathbb{Z}}\mathbb{Z}[U].$$

(We choose here the orientation system so that the two disks algebraically cancel, c.f. Subsection 4.9; but there are in fact two equivalence classes orientation systems giving two different Floer homology groups, just as there are two locally constant $\mathbb{Z}$ coefficient systems over S^1 giving two possible homology groups.) Since the described Heegaard decomposition is weakly admissible for all Spin^c structures, and both intersection points represent $\mathfrak{s}_0$, it follows that

$$\widehat{HF}(S^1\times S^2,\mathfrak{s})=HF^+(S^1\times S^2,\mathfrak{s})=0$$

if $\mathfrak{s}\neq\mathfrak{s}_0$.

To calculate the other homologies in non-torsion Spin^c structures, we must wind transverse to α, and then push the basepoint z across α some number of times, to achieve strong admissibility. Indeed, it is straightforward to verify that if $h\in H^2(S^1\times S^2)$ is a generator, then for $\mathfrak{s}=\mathfrak{s}_0+n\cdot h$ with $n>0$,

$$\partial^\infty[x^+,i]=[x^-,i]-[x^-,i-n];$$

in particular,

$$HF^-(S^2\times S^1,\mathfrak{s}_0+nh)\cong HF^\infty(S^2\times S^1,\mathfrak{s}_0+nh)\cong\mathbb{Z}[U]/(U^n-1).$$

7.5. **Connected sums with $S^1\times S^2$.** We can determine the effect of connected sums with $S^1\times S^2$ on HF^+, using the gluing theory which was used to establish the stabilization invariance of the homology groups (c.f. Section 6 of [23]).

Proposition 7.10. *Let $\mathfrak{s}_0$ be the* Spin^c *structure on $S^2\times S^1$ with $c_1(\mathfrak{s}_0)=0$, and let Y be an oriented three-manifold, equipped with a* Spin^c *structure $\mathfrak{s}$. There is a $\Lambda^*(H_1(Y\#(S^2\times S^1)/\mathrm{Tors}))$-equivariant isomorphism:*

$$HF^+(Y\#(S^2\times S^1),\mathfrak{s}\#\mathfrak{s}_0)\cong HF^+(Y,\mathfrak{s})\otimes\wedge^*H^1(S^2\times S^1).$$

For all other Spin^c *structures on $Y\#(S^2\times S^1)$, HF^+ vanishes.*

Proof. We consider first Spin^c structures on $Y\#(S^2\times S^1)$ of the form $\mathfrak{s}\#\mathfrak{s}_0$. Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},z)$ be a Heegaard diagram describing Y, which is weakly admissible for all Spin^c structures. We form the connected sum with E, a surface of genus one, with a pair of curves α_{g+1} and β_{g+1} which are isotopic (through an isotopy which does not cross the connected sum point), meeting in a pair x^+ and x^- of intersection points. As noted earlier, there is a pair of homotopy classes $\phi_1,\phi_2\in\pi_2(x^+,x^-)$ which contain holomorphic representatives, indeed both containing a unique smooth, holomorphic representative (for any constant complex structure on E). Of course, $(\Sigma\#E, \{\alpha_1, ..., \alpha_{g+1}\}, \{\beta_1, ..., \beta_{g+1}\}, z)$ describes $Y\#(S^2\times S^1)$, and it, too is weakly admissible for all Spin^c structures. Since the first Chern class of the Spin^c structure $(\mathfrak{s}\#\mathfrak{s}_0)$ evaluates trivially on the periodic domain associated to $S^2\subset S^2\times S^1$, we can choose our basepoint to lie in the Σ summand in $\Sigma\#E$, close to the connected sum point. Of course $\mathbb{T}'_\alpha\cap\mathbb{T}'_\beta = (\mathbb{T}_\alpha\cap\mathbb{T}_\beta)\times\{x^+, x^-\}$; thus $CF^+(Y_0, \mathfrak{s}\#\mathfrak{s}_0)$ is generated by $[\mathbf{x}, i]\otimes\{x^\pm\}$, where $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, and $\mathrm{gr}([\mathbf{x}, i]\otimes\{x^+\}, [\mathbf{x}, i]\otimes\{x^-\}) = 1$, i.e. $CF^+(Y\#(S^2\times S^1), \mathfrak{s}\#\mathfrak{s}_0)\cong CF^+(Y, \mathfrak{s})\oplus CF^+(Y, \mathfrak{s})$ (where the second factor is shifted in grading by one). We claim that when the neck is sufficiently long, the differential respects this splitting.

Fix $\mathbf{x}, \mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$. First, we claim that for sufficiently long neck lengths, the only homotopy classes $\phi'\in\pi_2(\mathbf{x}\times\{x^+\}, \mathbf{y}\times\{x^+\})$ with non-trivial holomorphic representatives are the ones which are constant on x^+. This follows from the following weak limit argument. Suppose there is a homotopy class $\phi'\in\pi_2(\{\mathbf{x}, x^+\}, \{\mathbf{y}, x^+\})$ with $\mu(\phi)\neq 0$ for which the moduli space is non-empty for arbitrarily large connected sum neck-length. Then, there is a limiting holomorphic disk in $\mathrm{Sym}^g(\Sigma)\times E$. On the E factor, the disk must be constant, since $\pi_2(x^+, x^+)\cong\mathbb{Z}$ (here we are in the first symmetric product of the genus one surface), and all non-constant homotopy classes have domains with positive and negative coefficients. Thus, the limiting flow has the form $\phi\times\{x^+\}$ for some $\phi\in\pi_2(\mathbf{x}, \mathbf{y})$ (in $\mathrm{Sym}^g(\Sigma)$). Theorem 6.4 of [23] applies then to give an identification $\mathcal{M}(\phi\times\{x^+\})\cong\mathcal{M}(\phi')$. Indeed, we have the same statement with x^- replacing x^+.

Next, we claim that (for generic choices) if $\phi'\in\pi_2(\mathbf{x}\times\{x^+\}, \mathbf{y}\times\{x^-\})$ is any homotopy class with $\mu(\phi') = 1$, which contains a holomorphic representative for arbitrarily long neck-lengths, then it must be the case that $\mathbf{x} = \mathbf{y}$, and $\phi' = \{\mathbf{x}\}\times\phi_1$ or $\phi' = \{\mathbf{x}\}\times\phi_2$. Again, this follows from weak limits. If it were not the case, we would be able to extract a sequence which converges to a holomorphic disk in $\mathrm{Sym}^g(\Sigma)\times E$, which has the form $\phi\times\phi_1$ or $\phi\times\phi_2$. Now, it is easy to see that $\phi\times\{x^+\} * (\{\mathbf{y}\}\times\phi_i) = \phi'$ for $i = 1$ or 2 (by, say, looking at domains); hence, $\mu(\phi\times\{x^+\}) = 0$. It follows that as a flow in $\mathrm{Sym}^g(\Sigma)$, $\mu(\phi) = 0$. Thus, there are generically no non-trivial holomorphic representatives, unless ϕ is constant. Observe, of course, that $\#\widehat{\mathcal{M}}(\{\mathbf{x}\}\times\phi_1) = \#\widehat{\mathcal{M}}(\{\mathbf{x}\}\times\phi_2) = 1$, and also $n_z(\{\mathbf{x}\}\times\phi_1) = n_z(\{\mathbf{x}\}\times\phi_2)$. With the appropriate orientation system, these flows cancel in the differential.

Putting these facts together, we have established that

$$\partial'([\mathbf{x}, i]\times\{x^\pm\}) = (\partial[\mathbf{x}, i])\times\{x^\pm\}$$

(where ∂' is the differential on $CF^+(Y\#(S^2\times S^1), \mathfrak{s}\#\mathfrak{s}_0)$, and ∂ is the differential on $CF^+(Y, \mathfrak{s})$. Indeed, it is easy to see the action of the one-dimensional homology generator coming from $S^2\times S^1$ annihilates $[\mathbf{x}, i]\times\{x^-\}$, and sends $[\mathbf{x}, i]\times\{x^+\}$ to $[\mathbf{x}, i]\times\{x^-\}$.

An adaptation of the above arguments can be used to show that the map

$$CF^+(Y)\otimes H_*(S^1)\longrightarrow CF^+(Y\#(S^2\times S^1))$$

defined by $[\mathbf{x}, i]\times x^\pm\mapsto[\mathbf{x}\times\{x^\pm\}, i]$ – which, of course, induces the isomorphism in homology (for sufficiently long connected sum necks) – is in fact equivariant under the action of $H_1(Y;\mathbb{Z})/\mathrm{Tors}\oplus H_1(S^2\times S^1)$. For instance, generators coming from $H_1(Y;\mathbb{Z})/\mathrm{Tors}$, can be represented by a constraint $\gamma\subset\Sigma$ (c.f. Remark 4.12). Then modifying the above arguments, we obtain an identification (again, for sufficiently long connected sum necks) of cut-down moduli spaces

$$\left\{u \in \mathcal{M}(\phi) \middle| u(0 \times 1) \in (\gamma \times \mathrm{Sym}^{g-1}(\Sigma)) \cap \mathbb{T}_\alpha\right\} \cong \\ \left\{u \in \mathcal{M}(\phi'_\pm) \middle| u'(0 \times 1) \in (\gamma \times \mathrm{Sym}^{g}(\Sigma \# E)) \cap (\mathbb{T}_\alpha \times \{\alpha_{g+1}\})\right\},$$

where $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ has $\mu(\phi) = 1$, and $\phi'_\pm$ is the corresponding element in $\pi_2(\mathbf{x} \times \{x^\pm\})$. This establishes the $H_1(Y;\mathbb{Z})/\mathrm{Tors}$-equivariance of the induced map on homology. The equivariance for the other factor $(H_1(S^2 \times S^1;\mathbb{Z}))$ follows similarly.

When the first Chern class of the Spin^c structure evaluates non-trivially on the $S^2 \times S^1$ factor, we can make α_{g+1} and β_{g+1} disjoint, and have a Heegaard diagram which is still weakly admissible for this Spin^c structure. Since there are no intersection points, it follows that HF^+ in this case is trivial. □

8. Adjunction Inequalities

As a first application of the Floer homologies constructed above, we relate them to the minimal genus problem in dimension three.

Theorem 8.1. *Let $Z \subset Y$ be a connected embedded two-manifold of genus $g(Z) > 0$ in an oriented three-manifold with $b_1(Y) > 0$. If $\mathfrak{s}$ is a Spin^c structure for which $HF^+(Y, \mathfrak{s}) \neq 0$, then*

$$\left|\langle c_1(\mathfrak{s}), [Z]\rangle\right| \leq 2g(Z) - 2.$$

We can reformulate this result using Thurston's semi-norm, see [31]. If $Z = \bigcup_{i=1}^{k} Z_i$ is a closed surface with k connected components, let

$$\chi_-(Z) = \sum_{i=1}^{k} \max(0, -\chi(Z_i)).$$

The Thurston semi-norm of a homology class $\xi \in H_2(Y;\mathbb{Z})$ is then defined by

$$\Theta(\xi) = \inf\{\chi_-(Z) \big| Z \subset Y, [Z] = \xi\}.$$

In this language, Theorem 8.1 says the following:

Corollary 8.2. *If $HF^+(Y, \mathfrak{s}) \neq 0$, then $\left|\langle c_1(\mathfrak{s}), \xi\rangle\right| \leq \Theta(\xi)$ for all $\xi \in H_2(Y;\mathbb{Z})$.*

Proof. First observe that if Z is an embedded sphere in Y, then for each $\mathfrak{s}$ for which $HF^+(Y, \mathfrak{s}) \neq 0$, we have that $\langle c_1(\mathfrak{s}), [Z]\rangle = 0$. This is a direct consequence of Theorem 8.1: attach a handle to Z to get a homologous torus Z' and apply the theorem.

Now, let $\cup_{i=1}^{k} Z_i$ be a representative of ξ whose χ_- is minimal, labeled so that Z_i for $i = 1, ..., \ell$ are the components with genus zero. Then,

$$|\langle c_1(\mathfrak{s}), \xi\rangle| \leq \sum_{i=1}^{k} |\langle c_1(\mathfrak{s}), Z_i\rangle| \leq \sum_{i=\ell+1}^{k} (2g(Z_i) - 2) = \Theta(\xi).$$

□

Theorem 8.1 is proved by constructing a special Heegaard diagram for Y, containing a periodic domain representative for Z with a particular form. The theorem then follows from a formula which calculates the evaluation of $c_1(\mathfrak{s})$ on Z.

The following lemma, which is proved at the end of this subsection, provides the required Heegaard diagram for Y.

Lemma 8.3. *Suppose $Z \subset Y$ is a homologically non-trivial, embedded two-manifold of genus $h = g(Z)$, then Y admits a genus g Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, with $g > 2h$, containing a periodic domain $\mathcal{P} \subset \Sigma$ representing $[Z]$, all of whose multiplicities are one or zero. Moreover, $\mathcal{P}$ is a connected surface whose Euler characteristic is equal to $-2h$, and $\mathcal{P}$ is bounded by β_1 and α_{2h+1}.*

Moreover, we have the following result, which follows from a more general theorem proved in Subsection 8.1:

Proposition 8.4. *If* $\mathbf{x} = \{x_1, ..., x_g\}$ *is an intersection point, and* z *is chosen in the complement of the periodic domain* $\mathcal{P}$ *of Lemma 8.3, then*

$$\langle c_1(\mathfrak{s}_z(\mathbf{x})), H(\mathcal{P})\rangle = 2 - 2h + 2\#(x_i \textit{ in the interior of } \mathcal{P}).$$

Proof of Theorem 8.1 If $\langle c_1(\mathfrak{s}), [Z]\rangle = 0$, then the inequality is obviously true.

We assume that $\langle c_1(\mathfrak{s}), [Z]\rangle$ is non-zero. If $Z \subset Y$ is an embedded surface of genus $g(Z) = h$, then we consider a special Heegaard decomposition constructed in Lemma 8.3. Suppose that $b_1(Y) = 1$. Then this Heegaard decomposition is weakly admissible for any non-torsion Spin^c structure $\mathfrak{s}$: there are no non-trivial periodic domains $\mathcal{D}$ with $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 0$. Clearly, of all $x_i \in \mathbf{x}$, exactly two must lie on the boundary. According to Proposition 8.4, then,

$$\langle c_1(\mathfrak{s}), \mathcal{P}\rangle = 2 - 2h + 2\#(x_i \in \mathrm{int}\mathcal{P});$$

i.e.

$$2 - 2h \leq \langle c_1(\mathfrak{s}), [Z]\rangle.$$

If we consider the same inequality for $-Z$ (or using the J invariance), we get the stated bounds.

In the case where $b_1(Y) > 1$, we must wind transverse to the $\alpha_1, ..., \widehat{\alpha_{2h+1}}, ..., \alpha_g$ to achieve weak admissibility. We choose our transverse curves to be of course disjoint from one another (and α_{2h+1}). In winding along these curves, we leave the periodic domain $\mathcal{P}$ representing S unchanged. Moreover, each periodic domain $\mathcal{Q}$ which evaluates trivially on $c_1(\mathfrak{s})$ must contain some α_j with $j \neq 2h+1$ on its boundary; thus, by twisting sufficiently along the γ-curves, we can arrange that the Heegaard decomposition is weakly admissible. The previous argument when $b_1(Y) = 1$ then applies. □

We now return to the proof of Lemma 8.3.

Proof of Lemma 8.3. The tubular neighborhood of Z, identified with $Z \times [-1, 1]$, has a handle decomposition with one zero-handle, $2h$ one-handles, and one two-handle; i.e. the tubular neighborhood admits a Morse function f with one index zero critical point p, $2h$ index one critical points $\{a_1, \ldots, a_{2h}\}$, and one index two critical point b_1. Hence, we have a genus $2h$ handlebody V_{2h}, with an embedded circle on its boundary $\beta_1 \subset \partial V_{2h} = \Sigma_{2h}$ (the descending manifold of b_1). The circle β_1 separates Σ_{2h}, and attaching a two-handle to V_{2h} along β_1 gives us the tubular neighborhood of Z. Choose a component of the complement of β_1, and denote its closure by $F_{2h} \subset \Sigma_{2h}$. Attaching the descending manifold of b_1 along $\partial F_{2h} = \beta_1$, we obtain a representative of $[Z]$ in this neighborhood.

We claim that the Morse function f can be extended to all of Y, so that the extension has one index three critical point and no additional index zero critical points. To see this, extend f to a Morse function $\widetilde{f}$, and first cancel off all new index zero critical points. This is a familiar argument from Morse theory (see for instance [22]): given another index zero critical point p', there is some index one critical point a which admits a unique flow to p' (if there no such index one critical points, then p' would generate a $\mathbb{Z}$ in the Morse complex for Y, which persists in $H_0(Y)$; but also, the sum of the other index zero critical points would not lie in the image of ∂, so it, too, would persist in homology, violating the connectedness hypothesis of Y). Thus, we can cancel p' and the critical point a.

Next, we argue that the extension $\widetilde{f}$ need contain only one index three critical point, as well. If there were two, call them q and q', we show that one of them can necessarily be canceled with an index two critical point other than b_1. If this could not be done, then both q and q' would have a unique flow-line to b_1. Thus, both q and q' would represent non-zero elements in $H_3(Y,Z) \cong H^0(Y-Z)$. But this is impossible since the complement $Y-Z$ is connected, thanks to our homological assumption on Z (which ensures that Z admits a dual circle which hits it algebraically a non-zero number of times). In fact, the extension generically contains no flows between index i and index j critical points with $j \geq i$, hence giving us a Heegaard decomposition of Y.

Thus, Y has a handlebody decomposition $Y = U_0 \cup_{\Sigma_g} U_1$, where U_0 is obtained from V_{2h} by attaching a sequence of one-handles. The attaching regions for each of these one-handles consists of two disjoint disks in Σ_{2h}, which are disjoint from β_1. At least one of them has one component inside F_{2h} and one outside. This follows from the fact that β_1 is homologically trivial in Σ_{2h}, but homologically non-trivial in the final Heegaard surface Σ. Let α_{2h+1} be the attaching circle for this one-handle. After handleslides across α_{2h+1}, we can arrange that all the other additional one-handles were attached in the complement of F_{2h}. The domain in F_{2h} between and α_{2h+1} and β_1 represents Z.

□

8.1. **The first Chern class formula.** Next, we give a proof Proposition 8.4. Indeed, we prove a more general result. But first, we introduce some data associated to periodic domains.

A periodic domain $\mathcal{P}$ is represented by an oriented two-manifold with boundary

$$\Phi\colon F \longrightarrow \Sigma,$$

whose boundary maps under Φ into $\boldsymbol{\alpha} \cup \boldsymbol{\beta}$. We consider the pull-back bundle $\Phi^*(T\Sigma)$ over F. This bundle is canonically trivialized over the boundary: the velocity vectors of the attaching circles give rise to natural trivializations. We define the *Euler measure* of the periodic domain $\mathcal{P}$ by the formula:

$$\chi(\mathcal{P}) = \langle c_1(\Phi^* T\Sigma; \partial), F\rangle,$$

where $c_1(\Phi^* T\Sigma;\partial)$ is first Chern class of $\Phi^*T\Sigma$ relative to this boundary trivialization. (It is easy to verify that $\chi(\mathcal{P})$ is independent of the representative $\Phi\colon F \longrightarrow \Sigma$.)

For example, if $\mathcal{P} \subset \Sigma$ is a periodic domain all of whose coefficients are one or zero, with $\partial\mathcal{P} = \cup_{i=1}^m \gamma_i$ where the γ_i are chosen among the $\boldsymbol{\alpha}$ and the $\boldsymbol{\beta}$, then $\chi(\mathcal{P})$ agrees with the usual Euler characteristic of $\mathcal{P}$, thought of as a subset of Σ.

Given a reference point $x \in \Sigma$, there is another quantity associated to periodic domains, obtained from a natural generalization of the multiplicity $n_x(\mathcal{P})$ defined in Section 2. This quantity, which we denote $\overline{n}_x(\mathcal{P})$ is defined by:

$$\overline{n}_x\Big(\sum_i a_i \mathcal{D}_i\Big) = \sum_i a_i \left(\begin{array}{ll} 1 & \text{if } x \text{ lies in the interior of } \mathcal{D}_i \\ \frac{1}{2} & \text{if } x \text{ lies in the interior of some edge of } \mathcal{D}_i \\ & \text{or two vertices of } \mathcal{D}_i \text{ are identified with } x \\ \frac{1}{4} & \text{if one vertex of } \mathcal{D}_i \text{ is identified with } x \\ 0 & \text{if } x \notin \mathcal{D}_i \end{array}\right).$$

Of course, if x lies in $\Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$, then $\overline{n}_x(\mathcal{P}) = n_x(\mathcal{P})$. If $\mathcal{P}$ has all multiplicities one or zero, and x is contained in its boundary, then $\overline{n}_x(\mathcal{P}) = \frac{1}{2}$.

Proposition 8.5. *Fix a class $\xi \in H_2(Y;\mathbb{Z})$, a base point $z \in \Sigma - \alpha_1 - \ldots - \alpha_g - \beta_1 - \ldots - \beta_g$, and a point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Let $\mathcal{P}$ be the periodic domain associated to z and ξ, and let $\mathfrak{s}$ be the* Spin^c *structure $s_z(\mathbf{x})$. Then the evaluation of the first Chern class of $\mathfrak{s}$ on ξ is calculated by*

$$\langle c_1(\mathfrak{s}), \xi \rangle = \chi(\mathcal{P}) + 2 \sum_{x_i \in \mathbf{x}} \overline{n}_{x_i}(\mathcal{P}).$$

Of course, Proposition 8.4 is a special case of this result, since in that case, two of the x_i are in the boundary of $\mathcal{P}$, so they have $\overline{n}_{x_i} = \frac{1}{2}$.

To prove the proposition, we need an explicit understanding of the vector field belonging to the Spin^c structure $s_z(\mathbf{x})$. Specifically, consider the normalized gradient vector field $\frac{\vec{\nabla} f}{|\vec{\nabla} f|}$, restricted to the mid-level Σ of the Morse function f (compatible with the given Heegaard decomposition of Y). Clearly, the orthogonal complement of the vector field is canonically identified with the tangent bundle of Σ. Suppose, then, that γ is a connecting trajectory between an index one and an index two critical point (which passes through Σ). We can replace the gradient vector field by another vector field v which agrees with $\frac{\vec{\nabla} f}{|\vec{\nabla} f|}$ outside of a small three-ball neighborhood B, which meets Σ in a disk D. Let τ be a trivialization of the two-plane field $v^\perp|\partial D$ which extends as a trivialization of $T\Sigma|D$. There is a well-defined relative first Chern class $c_1(v,\tau) \in H^2(D,\partial D)$, which we can calculate as follows:

Lemma 8.6. *For D, v, and τ as above, the relative first Chern number is given by*

$$\langle c_1(v,\tau), [D,\partial D] \rangle = 2$$

(where we orient D in the same manner as $\Sigma = \partial U_0$).

Proof. Using an appropriate trivialization of the tangent bundle $TY|B$, can view the normalized gradient vector field $\frac{\vec{\nabla} f}{|\vec{\nabla} f|}$ as constant over D. Let $S = \partial B$ be the boundary, which is divided into two hemispheres $S = D_1 \cup D_2$, so that the sphere $D_1 \cup D$ contains the index one critical point and $D \cup D_2$ contains the index two critical point. We can replace $\frac{\vec{\nabla} f}{|\vec{\nabla} f|}$ by another vector field v which agrees with the normalized gradient over S, and vanishes nowhere in B (and hence can be viewed as a unit vector field). With respect to the trivialization of $TY|B$, we can think of the vector field as a map to the two-sphere; indeed the restriction

$$v\colon D \longrightarrow S^2,$$

is constant along the boundary, so it has a degree, which in the present case is one, since

$$\begin{aligned} -1 &= \deg_{D_1}\left(\frac{\vec{\nabla} f}{|\vec{\nabla} f|}\right) + \deg_D\left(\frac{\vec{\nabla} f}{|\vec{\nabla} f|}\right) \\ &= \deg_{D_1}(v) \end{aligned}$$

and

$$0 = \deg_{D_1}(v) + \deg_D(v).$$

The line bundle we are considering, $v^\perp$, then, is the pull-back of the tangent bundle to S^2, whose first Chern number is the Euler characteristic for the sphere. □

Proof of Proposition 8.5. We find it convenient to consider domains with only non-negative multiplicities; thus, we prove the following formula (for sufficiently large m):

$$\langle c_1(\mathfrak{s}),\xi\rangle = \chi(\mathcal{P}+m[\Sigma]) + 2\left(\sum_{x_i\in\mathbf{x}} \overline{n}_{x_i}(\mathcal{P}+m[\Sigma])\right) - 2n_z(\mathcal{P}+m[\Sigma]). \tag{11}$$

In fact, since

$$\begin{aligned}\chi(\mathcal{P}+m[\Sigma]) &= \chi(\mathcal{P})+m(2-2g),\\ \sum_{x_i\in\mathbf{x}}\overline{n}_{x_i}(\mathcal{P}+m[\Sigma]) &= mg+\sum_{x_i\in\mathbf{x}}\overline{n}_{x_i}(\mathcal{P})\\ n_z(\mathcal{P}+m[\Sigma]) &= m,\end{aligned}$$

Equation (11) for any specific value of m implies the formula stated in the proposition.

The reformulation has the advantage that for m sufficiently large, $\mathcal{P}+m[\Sigma]$ is represented by a map

$$\Phi\colon F\longrightarrow\Sigma$$

which is nowhere orientation-reversing, and whose restriction to each boundary component is a diffeomorphism onto its image (see Lemma 2.5).

Near each boundary component of F, we can identify a neighborhood in F with the half-open cylinder $[0,1)\times S^1$. Suppose that the image of the boundary component is an β curve. The β curve canonically bounds a disk in U_1: this disk D consists of points which flow (under $\vec{\nabla} f$) into the associated index two critical point. Of course, we can glue this disk to F along the boundary, and correspondingly extend Φ across the disk as a map into Y, but then the gradient $\vec{\nabla} f$ vanishes at some point of the extended map. To avoid this, we can back off from the boundary of F: we delete a small neighborhood $[0,\epsilon)\times S^1$ from F, to obtain a new manifold-with-boundary F^-. In these local coordinates, now, the boundary of F^- is a translate of the β curve $\{\epsilon\}\times S^1$. Now, we can attach a translate of the disk, D_-. Now, it is easy to see that (a smoothing of) the cap $\left([\epsilon,1)\times S^1\right)\cup D_-$ is transverse to the gradient flow $\vec{\nabla} f$. (See the illustration in Figure 3.)

We can perform the analogous construction at the α-components of the boundary of F, only now, the α curve bounds a disk D in U_0, which consists of points flowing out of the corresponding index two critical point. By cutting out a neighborhood of the boundary, and attaching a translate of the D, we once again obtain a cap which is transverse to the gradient flow $\vec{\nabla} f$.

Observe that if $x_i\in\mathrm{int}\mathcal{P}$, then (if we chose the above ϵ sufficiently small),

$$\overline{n}_{x_i}(\mathcal{P}) = \#\{x\in F^-\big|\Phi(x)=x_i\} \tag{12}$$

(with the same formula holding for z in place of x_i). Moreover, if $x_i\in\partial\mathcal{P}$, then

$$\overline{n}_{x_i}(\mathcal{P}) = \frac{1}{2}\#\{x\in\partial F|\Phi(x)=x_i\}+\#\{x\in F^-|\Phi(x)=x_i\}. \tag{13}$$

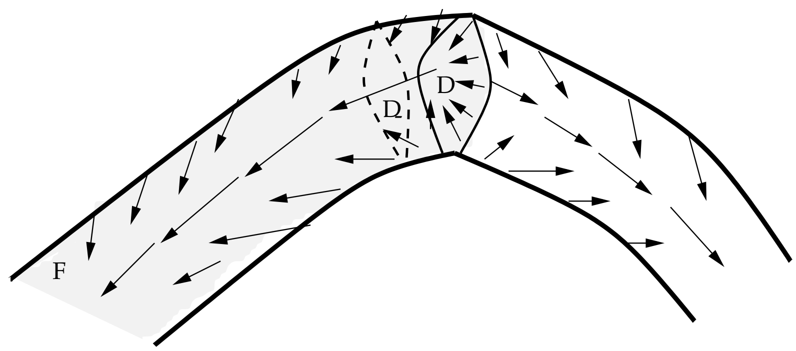

FIGURE 3. The gradient flow inside a one-handle. The shaded region on the boundary of the one-handle is a piece of F; the disk D (with solid boundary, in the center) goes through the index one critical point. Its translate D_- (with dotted boundary) does not, and the subregion of F terminating in the dotted circle, when capped off by D_-, is transverse to the gradient flow.

By adding the caps as above to F^-, we construct a closed, oriented two-manifold $\widehat{F}$ and a map

$$\widehat{\Phi}\colon \widehat{F} \longrightarrow Y,$$

which crosses the connecting trajectories between the index one and two critical points at each point $x \in F^-$ which maps under Φ to x_i, and similarly, $\widehat{\Phi}$ crosses the connecting trajectory belonging to z at those $x \in F^-$ which map under Φ to z.

Away from these points, we have a canonical identification

$$\widehat{\Phi}^*((\vec{\nabla} f)^\perp) \cong \Phi^*(v^\perp).$$

By the local calculation from Lemma 8.6, it follows that

(14)

$$\langle e\left(\widehat{\Phi}^*(v^\perp)\right), \widehat{F}\rangle = \langle e\left(\widehat{\Phi}^*(\vec{\nabla} f^\perp)\right), \widehat{F}\rangle + 2\#\{x \in F^- | \Phi(x) = x_i\} - 2\#\{x \in F^- | \Phi(x) = z\}.$$

(Note that the term involving z follows just as in the proof of Lemma 8.6, with the difference that now the index of the vector field v around the corresponding critical point in U_0 is $+1$ rather than -1, since the critical point has index zero rather than one.)

Moreover, the Euler number of $\widehat{\Phi}^*(\vec{\nabla} f^\perp)$ is $\chi(\mathcal{P})$ plus the number of disks which are attached to F^- to obtain the closed manifold $\widehat{F}$ (since each boundary disk is transverse to the gradient flow, so $\vec{\nabla} f^\perp$ is naturally identified with the tangent bundle of the disk, which has relative Euler number one relative to the trivialization it gets from the bounding circle). But the number of such disks is simply $\#\{x \in \partial F | \Phi(x) = x_i\}$. Combining this with Equations (12), (13), and (14), we obtain Equation (11), and hence proposition follows. □

9. The Euler characteristic of HF^+

In [32], Turaev defines a torsion function

$$\tau_Y \colon \mathrm{Spin}^c(Y) \longrightarrow \mathbb{Z},$$

which is a generalization of the Alexander polynomial. This function can be calculated from a Heegaard diagram of Y as follows. Fix integers i and j between 1 and g, and consider corresponding tori

$$\mathbb{T}^i_\alpha = \alpha_1 \times .. \times \widehat{\alpha_i} \times ... \times \alpha_g \quad \text{and} \quad \mathbb{T}^j_\beta = \beta_1 \times .. \times \widehat{\beta_j} \times ... \times \beta_g$$

in $\mathrm{Sym}^{g-1}(\Sigma)$ (where the hat denotes an omitted entry). There is a map σ from $\mathbb{T}^i_\alpha \cap \mathbb{T}^j_\beta$ to $\mathrm{Spin}^c(Y)$, which is given by thinking of each intersection point as a $(g-1)$-tuple of connecting trajectories from index one to index two critical points. Moreover, orienting α_i, there is a distinguished trajectory connecting the index zero critical point to the index one critical point a_i corresponding to α_i; similarly, orienting β_j, there is a distinguished trajectory connecting the critical point b_j corresponding to the circle β_j to the index index three critical point in Y. This $(g+1)$-tuple of trajectories then gives rise to a Spin^c structure in the usual manner (modifying the upward gradient flow in the neighborhoods of these trajectories). Thus, we can define

$$\Delta_{i,j}(\mathfrak{s}) = \pm \sum_{\{\mathbf{x}\in\mathbb{T}^i_\alpha\cap\mathbb{T}^j_\beta \big| \sigma(\mathbf{x})=\mathfrak{s}\}} \epsilon(\mathbf{x}),$$

where $\epsilon(\mathbf{x})$ is the local intersection number of $\mathbb{T}^i_\alpha$ and $\mathbb{T}^j_\beta$ at $\mathbf{x}$, and the overall sign depends on i, j and g. (It is straightforward to verify that this geometric interpretation is equivalent to the more algebraic definition of $\Delta_{i,j}$ given in [32], see for instance Section 7 from [25].)

Choose i and j so that both α_i^* and β_j^* have non-zero image in $H^2(Y;\mathbb{R})$. When $b_1(Y) > 1$, Turaev's torsion is characterized by the equation

$$\tau(\mathfrak{s}) - \tau(\mathfrak{s}+\alpha_i^*) - \tau(\mathfrak{s}+\beta_j^*) + \tau(\mathfrak{s}+\alpha_i^*+\beta_j^*) = \Delta_{i,j}(\mathfrak{s}), \tag{15}$$

and the property that it has finite support. (To define β_j^* here, let C be a curve in Σ with $\beta_i \cap C = \delta_{i,j}$, and let β_j^* be Poincaré dual to the induced homology class in Y.) When $b_1(Y) = 1$, we need a direction t in $H^2(Y;\mathbb{R})$ (which we can think of as a component of $H^2(Y;\mathbb{R}) - 0$). Then, τ_t is characterized by the above equation and the property that τ_t has finite support amongst Spin^c structures whose first Chern class lies in the component of t.

The relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{s})\mathbb{Z}$-grading on $CF^+(Y,\mathfrak{s})$ induces a natural relative $\mathbb{Z}/2\mathbb{Z}$-grading on all Spin^c structures. Alternatively, this relative $\mathbb{Z}/2\mathbb{Z}$ grading between $[\mathbf{x}, i]$ and $[\mathbf{y}, j]$ is calculated by orienting $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$, and letting the relative degree be given by the parity of the product of the local intersection numbers of $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ at $\mathbf{x}$ and $\mathbf{y}$. This relative $\mathbb{Z}/2\mathbb{Z}$-grading can be used to define an Euler characteristic $\chi(HF^+(Y,\mathfrak{s}))$ (when the homology groups are finitely generated), which is well-defined up to an overall sign.

In this section, we relate the Euler characteristics of $HF^+(Y,\mathfrak{s})$ with Turaev's torsion function, when $c_1(\mathfrak{s})$ is non-torsion. (The torsion case will be covered in Subsection 11.7, after more is known about HF^∞; related results also hold for HF^-, c.f. Subsection 11.6.)

The overall sign on $\chi(HF^+(Y,\mathfrak{s}))$ will be pinned down once we define an absolute $\mathbb{Z}/2\mathbb{Z}$ grading on $HF^+(Y,\mathfrak{s})$ in Subsection 11.5.

9.1. $\chi(HF^+(Y,\mathfrak{s}))$ **when** $b_1(Y)=1$ **and** $\mathfrak{s}$ **is non-torsion.** Our aim is to prove the following:

Theorem 9.1. *Suppose $b_1(Y)=1$. If $\mathfrak{s}$ is a non-torsion* Spin^c *structure, then $HF^+(Y,\mathfrak{s})$ is finitely generated, and indeed,*

$$\chi(HF^+(Y,\mathfrak{s})) = \pm\tau_t(Y,\mathfrak{s}),$$

where τ_t is Turaev's torsion function, with respect to the component t of $H^2(Y;\mathbb{R})-0$ containing $c_1(\mathfrak{s})$.

The proof of Theorem 9.1 occupies the rest of the present subsection.

Let $\mathfrak{s}$ be a non-torsion Spin^c structure on Y. Let H be the generator of $H_2(Y;\mathbb{Z})$ with the property that

$$\langle c_1(\mathfrak{s}), H\rangle < 0.$$

After handleslides, we can arrange that the periodic domain $\mathcal{P}$ corresponding to H contains α_1 with multiplicity one in its boundary.

Choose a curve γ transverse to α_1 and disjoint from all other α_i for $i>1$, oriented so that $\alpha_1\cap\gamma=+1$. (Note that $\mathrm{PD}[\gamma]=\alpha_1^*$.) This curve has the property, then, that

$$\langle PD[\gamma], H\rangle = -1.$$

Let $\mathbb{T}_\gamma = \gamma\times\alpha_2\times ...\times\alpha_g$. Winding α_1 n times along γ, we obtain a new α-torus, which we denote $\mathbb{T}_\alpha(n)$. For each intersection point $\mathbf{x}\in\mathbb{T}_\gamma\cap\mathbb{T}_\beta$ we obtain $2n$ intersection points in $\mathbb{T}_\alpha(n)\cap\mathbb{T}_\beta$

$$\mathbf{x}_1^\pm, ..., \mathbf{x}_n^\pm,$$

which we order with decreasing distance to γ, with a sign $\pm$ indicating which side of γ they lie on ($-$ indicates left, $+$ indicates right). We call the points in $\mathbb{T}_\alpha(n)\cap\mathbb{T}_\beta$ *γ-induced*: equivalently, a γ-induced intersection point between $\mathbb{T}_\alpha(n)$ and $\mathbb{T}_\beta$ is a g-tuple of points in Σ, one of which lies in the winding region about γ. It is easy to see that $\mathbf{x}_i^+$ and $\mathbf{x}_i^-$ lie in the same equivalence class: indeed, there is a canonical flow-line (with Maslov index 1) connecting each $\mathbf{x}_i^+$ to $\mathbf{x}_i^-$. Thus, (for any choice of base-point z),

$$\begin{aligned}\mathfrak{s}_z(\mathbf{x}_i^+) - \mathfrak{s}_z(\mathbf{x}_j^+) &= (i-j)\mathrm{PD}(\gamma),\\ \mathfrak{s}_z(\mathbf{x}_i^+) &= \mathfrak{s}_z(\mathbf{x}_i^-).\end{aligned}$$

Our twisting will always be done in a "sufficiently small" area, so that the area of each component of $\Sigma-\mathrm{nd}(\gamma)-\alpha_1-\alpha_2-...-\alpha_g-\beta_1-...-\beta_g$ is greater than n times the area of $\mathrm{nd}(\gamma)$.

We will place our base-point z to the right of γ, in the $\left(\frac{n}{2}\right)^{th}$ subregion of the winding region about γ. For this choice of basepoint, if $\mathbf{x}\in\mathbb{T}_\gamma\cap\mathbb{T}_\beta$ then the Spin^c structure induced by $\mathbf{x}_{n/2}^\pm$ is independent of n. Of course, the base-point is not uniquely determined by this requirement: this region is divided into components by the β-curves which intersect γ; but we fix any one such region, for the time being.

Lemma 9.2. *If we wind n times, and place the basepoint in the $\left(\frac{n}{2}\right)^{th}$ subregion, and let $\mathcal{P}_n$ denote the corresponding periodic domain, then there is a constant c with the property that we*

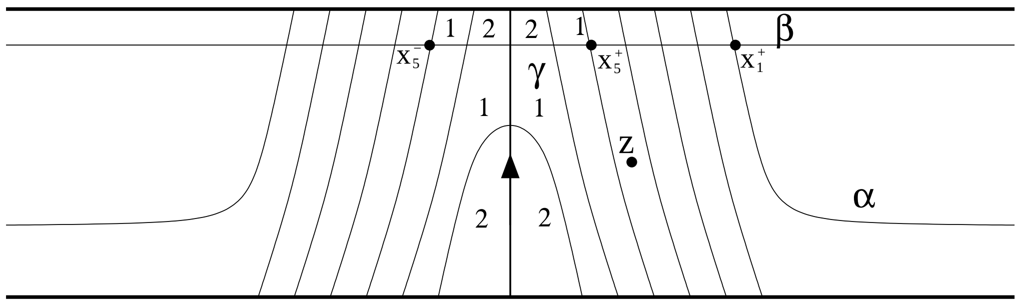

FIGURE 4. **Winding transverse to α.** We have pictured, once again, the cylindrical neighborhood of γ, and an α-curve obtained by winding six times transverse to γ. The basepoint z is placed in the third region, and intersection points corresponding to some β are labeled. The multiplicities correspond to the domain of a flow connecting x_5^+ to x_5^-.

can find basepoints w_1 and w_2 (near γ and away from γ respectively), so that

$$n_{w_1}(\mathcal{P}_n) \leq c - \frac{n}{2}, \quad \textit{and} \quad n_{w_2}(\mathcal{P}_n) \geq c + \frac{n}{2}.$$

Lemma 9.3. *Fix a* Spin^c *structure* $\mathfrak{s} \in Y$. *Then, if n is sufficiently large, the γ-induced intersection points of $\mathbb{T}_\alpha(n) \cap \mathbb{T}_\beta$ are the only ones which represent any of the* Spin^c *structures of the form* $\mathfrak{s} + k \cdot \mathrm{PD}[\gamma]$ *for* $k \geq 0$.

Proof. The intersection points between $\mathbb{T}_\alpha(n)$ and $\mathbb{T}_\beta$ which are not induced from γ correspond to the intersection points between the original $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$. So, suppose that $\mathbf{x}$ is an intersection point between $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ (there are, of course, finitely many such intersection points), and let z_0 be some basepoint outside the winding region. As we wind α_1 n times, and place the new basepoint z inside the winding region as above (so as not to cross any additional β-curves), we see that

$$s_z(\mathbf{x}) - s_{z_0}(\mathbf{x}) = -\frac{n}{2}\mathrm{PD}[\gamma],$$

where we think of $[\gamma]$ a one-dimension homology class in Y. The lemma then follows. □

Let $(\mathbb{T}_\alpha(n) \cap \mathbb{T}_\beta)^L \subset \mathcal{S}$ denote subset of γ-induced intersection points where the α_1 part lies to the "left" of γ, and $(\mathbb{T}_\alpha(n) \cap \mathbb{T}_\beta)^R$ denote subset of γ-induced intersection points where the α_1 part lies to the "right" of γ. (Note here that $\mathcal{S}$ denotes the subset of intersection points which induce the given Spin^c structure $\mathfrak{s}$ over Y.) There are corresponding subgroups L^+ and $R^+ \subset CF^+(Y)$; similarly we have L^∞ and $R^\infty \subset CF^\infty(Y)$.

Lemma 9.4. *Fix $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ and an integer n sufficiently large (in comparison with $\langle c_1(\mathfrak{s}), \mathcal{P}\rangle$). Then, for each γ-induced pair $\mathbf{x}^+$ and $\mathbf{y}^-$ inducing $\mathfrak{s}$, there are at most two homotopy classes $\phi^{\mathrm{in}}, \phi^{\mathrm{out}} \in \pi_2(\mathbf{x}^+, \mathbf{y}^-)$ with Maslov index one and with only non-negative multiplicities. Moreover, there are no such classes in $\pi_2(\mathbf{y}^-, \mathbf{x}^+)$.*

Proof. Assume $\mathrm{gr}(\mathbf{x}^+,\mathbf{y}^-)$ is odd, and let ϕ_n^{in} be the class with $\mu(\phi_n^{\mathrm{in}})=1$, and whose α_1 boundary lies entirely inside the tubular neighborhood of γ. We claim that $\mathcal{D}(\phi_{n+2}^{\mathrm{in}})$ is obtained from $\mathcal{D}(\phi_n^{\mathrm{in}})$ by winding only its α_1-boundary (and hence leaving the domain unchanged outside the winding region). This follows from the fact that the Maslov index is unchanged under totally real isotopies of the boundary. It follows then that the multiplicities of ϕ_n^{in} inside a neighborhood of γ grow like $n/2$. Recall that the multiplicities of $\mathcal{P}_n$ inside grow like $-n/2$, while outside they grow like $n/2$.

Now, the set of all $\mu=1$ homotopic classes connecting $\mathbf{x}^+$ to $\mathbf{y}^-$ is given by

$$\phi_n^{\mathrm{in}}+k\left(\mathcal{P}_n-\frac{\langle c_1(\mathfrak{s}),\mathcal{P}\rangle}{2}S\right).$$

If this class is to have non-negative multiplicities, we must have that $k=0$ or 1. This proves the assertion concerning classes from $\mathbf{x}^+$ to $\mathbf{y}^-$, letting $\phi_n^{\mathrm{out}}=\phi_n^{\mathrm{in}}+\left(\mathcal{P}_n-\frac{\langle c_1(\mathfrak{s}),\mathcal{P}\rangle}{2}S\right)$.

Considering classes from $\mathbf{y}^-$ to $\mathbf{x}^+$, note that all $\mu=1$ classes have the form

$$(S-\phi^{\mathrm{in}})+k\left(\mathcal{P}_n-\frac{\langle c_1(\mathfrak{s}),\mathcal{P}\rangle}{2}S\right).$$

When $k<0$, these classes have negative multiplicities outside γ. When $k\geq 0$, these have negative multiplicities inside the neighborhood of γ. $\square$

Proposition 9.5. *Given a* Spin^c *structure* $\mathfrak{s}$ *and an* n *sufficiently large, the subgroup* $L^\infty\subset CF^\infty(Y,\mathfrak{s})$ *is a subcomplex.*

Proof. This follows immediately from the previous lemma. $\square$

Of course, the above proposition allows us to think of R^∞ as a chain complex, as well, with differential induced from the quotient structure CF^∞/L^∞.

There is a natural map

$$\delta\colon R^\infty\longrightarrow L^\infty$$

given by taking the L^∞-component of the boundary of each element in R^∞. This induces the connecting homomorphism for the long exact sequence associated to the short exact sequence of complexes:

$$0\longrightarrow L^\infty\longrightarrow CF^\infty\longrightarrow R^\infty\longrightarrow 0.$$

To understand the homomorphism δ, let

$$f_1\colon R^\infty\longrightarrow L^\infty$$

be the homomorphism induced by $f_1([\mathbf{x}_i^+,j])=[\mathbf{x}_i^-,j-n_z(\phi)]$, where ϕ the disk connecting $\mathbf{x}_i^+$ to $\mathbf{x}_i^-$ which is supported in the tubular neighborhood of γ.

We can define an ordering on the γ-induced intersection points representing $\mathfrak{s}$ as follows. Let $[\mathbf{x},i],[\mathbf{y},j]\in\mathfrak{S}\times\mathbb{Z}$, then there is a unique $\phi\in\pi_2(\mathbf{x},\mathbf{y})$ with $n_z(\phi)=i-j$ and $\partial(\mathcal{D}(\phi))\cap\alpha_1$ supported inside the tubular neighborhood of γ. We denote the class ϕ by $\phi_{[\mathbf{x},i],[\mathbf{y},j]}$. We then say that

$$[\mathbf{x},i]>[\mathbf{y},j]$$

if

$$\mu(\phi_{[\mathbf{x},i],[\mathbf{y},j]}) > 0$$

or if

$$\mu(\phi_{[\mathbf{x},i],[\mathbf{y},j]}) = 0$$

and the area $\mathcal{A}(\mathcal{D}(\phi_{[\mathbf{x},i],[\mathbf{y},j]})) > 0$. Note that an ordering gives us a partial ordering for elements in $CF^\infty(Y,\mathfrak{s})$: fix $\xi,\eta \in CF^+(Y,\mathfrak{s})$, we say that $\xi < \eta$ if each $[\mathbf{x},i] \in \mathcal{S}\times\mathbb{Z}$ which appears with non-zero multiplicity in the expansion of ξ is smaller than each $[\mathbf{y},j] \in \mathcal{S}\times\mathbb{Z}$ which appears with non-zero multiplicity in the expansion of η.

In the following lemma, it is crucial to work with *negative* Spin^c structures, i.e. those for which $\langle c_1(\mathfrak{s}),\mathcal{P}\rangle < 0$.

Lemma 9.6. *If $\mathfrak{s}$ is a negative Spin^c structure, then the map*

$$\delta\colon R^\infty \longrightarrow L^\infty$$

can be written as

$$\delta = f_1 + f_2,$$

so that

$$f_2(g) < f_1(g)$$

for each $g = [\mathbf{x},i] \in R^\infty$.

Proof. Consider a pair of generators $[\mathbf{x}^+,i]$ and $[\mathbf{y}^-,j]$, for which the coefficient of δ is non-zero, i.e. that gives a homotopy class ψ for which $\mu(\psi)=1$ and $\mathcal{D}(\psi)\geq 0$. Thus, by Lemma 9.4, there are two possible cases, where $\psi = \phi^{\mathrm{in}}$ or $\psi=\phi^{\mathrm{out}}$ (for $\mathbf{x}^+$ and $\mathbf{y}^-$). Note also that $\phi^{\mathrm{in}} = \phi_{[\mathbf{x}^+,i][\mathbf{y}^-,j]}$.

The case where $\psi=\phi^{\mathrm{in}}$, has two subcases, according to whether or not $[\mathbf{y}^-,j] = f_1([\mathbf{x}^+,i])$. If $[\mathbf{y}^-,j] = f_1([\mathbf{x}^+,i])$, $\psi = \phi_{[\mathbf{x}^+,i]f_1([\mathbf{x}^+,i])}$, and it follows easily that $\#\mathcal{M}(\psi)=1$. Since the periodic domains have both positive and negative coefficients, the $[\mathbf{y}^-,j]$ coefficient of $f_2[\mathbf{x}^+,i]$ must vanish. If $[\mathbf{y}^-,j]\neq f_1([\mathbf{x}^+,i])$, then the domain of $\phi_{f_1([\mathbf{x}^+,i]),[\mathbf{y}^-,j]}$ must include some region outside the neighborhood of γ. Moreover, since

$$\phi_{[\mathbf{x}^+,i],f_1([\mathbf{x}^+,i])} + \phi_{f_1([\mathbf{x}^+,i]),[\mathbf{y}^-,j]} = \psi,$$

we have that $\mu(\phi_{f_1([\mathbf{x}^+,i]),[\mathbf{y}^-,j]})=0$; but since the support of the twisting region is sufficiently small, it follows that

$$\mathcal{A}(\phi_{f_1([\mathbf{x}^+,i]),[\mathbf{y}^-,j]})>0;$$

i.e. $f_1([\mathbf{x}^+,i]) > [\mathbf{y}^-,j]$.

When $\psi=\phi^{\mathrm{out}}$, it is easy to see that

$$\phi_{[\mathbf{x}^+,i],[\mathbf{y}^-,j]} = \phi^{\mathrm{out}} - \mathcal{P}.$$

It follows that $\mu(\phi_{[\mathbf{x}^+,i],[\mathbf{y}^-,j]}) = 1-\langle c_1(\mathfrak{s}),H(\mathcal{P})\rangle$. Moreover,

$$\phi_{[\mathbf{x}^+,i],f_1([\mathbf{x}^+,i])} + \phi_{f_1([\mathbf{x}^+,i]),[\mathbf{y}^-,j]} = \phi_{[x^+,i],[\mathbf{y}^-,j]},$$

so $\mu(\phi_{f_1([\mathbf{x}^+,i]),[\mathbf{y}^-,j]}) = -\langle c_1(\mathfrak{s}),H(\mathcal{P})\rangle > 0$, by our hypothesis on $\mathfrak{s}$, so that $f_1([\mathbf{x}^+,i]) > [\mathbf{y}^-,j]$. □

Proposition 9.7. *For negative* Spin^c *structures* $\mathfrak{s}$*, the map* $\delta^+\colon R^+ \longrightarrow L^+$ *is surjective, and its kernel is identified with the kernel of* f_1^+ *(as a* $\mathbb{Z}/\mathfrak{d}(\mathfrak{s})\mathbb{Z}$*-graded groups).*

Proof. This is an algebraic consequence of Lemma 9.6.

We can define a right inverse to f_1,

$$P_1[\mathbf{x}_i^-, j] = [\mathbf{x}_i^+, j + n_z(\phi)],$$

where ϕ is the disk connecting $\mathbf{x}_i^+$ to $\mathbf{x}_i^-$. Then, we define a map

$$P = \sum_{N=0}^{\infty} P_1 \circ (-f_2 \circ P_1)^{\circ N}.$$

Note that the right-hand-side makes sense, since the map $f_2 \circ P_1$ decreases the ordering (which is bounded below), so for any fixed $\xi \in R^+$, there is some N for which

$$(-f_2 \circ P_1)^{\circ N}(\xi) = 0.$$

It is easy to verify that P is a right inverse for δ^+.

The map sending $\xi \mapsto \xi - P \circ \delta^+(\xi)$ induces a map from $\mathrm{Ker} f_1$ to $\mathrm{Ker}\delta^+$, which is injective, since for any $\xi \in \mathrm{Ker} f_1$, we have that

$$P \circ \delta^+(\xi) = P \circ f_2(\xi) < \xi.$$

Similarly, the map $\xi \mapsto \xi - P_1 \circ f_1(\xi)$ supplies an injection $\mathrm{Ker}\delta^+ \longrightarrow \mathrm{Ker} f_1$. It follows that $\mathrm{Ker} f_1 \cong \mathrm{Ker}\delta^+$. □

Proposition 9.8. *For negative* Spin^c *structures, the rank* $HF^+(Y, \mathfrak{s})$ *is finite. Moreover, we have that* $\chi(H_*(\ker \delta_{\mathfrak{s}}^+)) = \chi(HF^+(Y, \mathfrak{s}))$*.*

Proof. According to Proposition 9.7 we have the short exact sequence

$$0 \longrightarrow \ker \delta^+ \longrightarrow R^+ \xrightarrow{\delta^+} L^+ \longrightarrow 0,$$

which we compare with the short exact sequence

$$0 \longrightarrow L^+ \longrightarrow CF^+ \longrightarrow R^+ \longrightarrow 0.$$

The result then follows by comparing the associated long exact sequences, and observing that the connecting homomorphism for the second sequence agrees with the map on homology induced by δ^+. □

Proposition 9.9. *Let* $\mathfrak{s}$ *be a negative* Spin^c *structure, then*

$$\chi(\mathrm{Ker} f_1(\mathfrak{s})) = \pm \tau_t(\mathfrak{s}),$$

where t *is the component of* $H^2(Y, \mathbb{Z})$ *containing* $c_1(\mathfrak{s})$*.*

Proof. The map f_1 depends on a base-point and an equivalence class of intersection point. However, according to Propositions 9.7 and 9.8, $\chi(\mathrm{Ker} f_1^+(\mathfrak{s}))$ depends on this data only through the underlying Spin^c structure $\mathfrak{s}$ (when the latter is negative). Let $\chi(\mathfrak{s})$ denote the Euler characteristic $\chi(\mathrm{Ker} f_1|_{\mathfrak{s}})$. We fix a basepoint z as before. We have a map

$$S_z\colon \mathbb{T}_\gamma \cap \mathbb{T}_\beta \longrightarrow \mathrm{Spin}^c(Y),$$

defined as follows. Given $\mathbf{x} \in \mathbb{T}_\gamma \cap \mathbb{T}_\beta$, we have

$$\mathfrak{s}_z(\mathbf{x}_1^+) + (n_z(\phi) - 1)\alpha_1^*,$$

where ϕ is the canonical homotopy class connecting $\mathbf{x}_1^+$ and $\mathbf{x}_1^-$, and $\alpha_1^* = \mathrm{PD}[\gamma]$. (In fact, it is easy to see that the above assignment is actually independent of the number of times we twist α_1 about γ.) There is a naturally induced function (depending on the basepoint)

$$a_z\colon \mathrm{Spin}^c(Y) \longrightarrow \mathbb{Z}$$

by

$$a_z(\mathfrak{s}) = \sum_{\left\{\mathbf{x}\in\mathbb{T}_\gamma\cap\mathbb{T}_\beta \,\middle|\, S_z(\mathbf{x})=\mathfrak{s}\right\}} \epsilon(\mathbf{x}),$$

where $\epsilon(\mathbf{x})$ is the local intersection number of $\mathbb{T}_\gamma \cap \mathbb{T}_\beta$ at $\mathbf{x}$. It is clear that

$$\chi(\mathfrak{s}) = \sum_{n=0}^{\infty} (n+1) \cdot a_z(\mathfrak{s} + n \cdot \alpha_1^*).$$

It follows that

$$\chi(\mathfrak{s}) - \chi(\mathfrak{s} + \alpha_1^*) = \sum_{n=0}^{\infty} a_z(\mathfrak{s} + n \cdot \alpha_1^*). \tag{16}$$

We investigate the dependence of a_z on the basepoint z. Note first that there must be some curve β_j which meets γ whose induced cohomology class β_j^* is not a torsion element in $H^2(Y;\mathbb{Z})$: indeed, any β_j appearing in the expression $\partial\mathcal{P}$ with non-zero multiplicity has this property. Suppose that z_1 and z_2 are a pair of possible base-points which can be connected by a path z_t disjoint from all the attaching circles except β_j, which it crosses transversally once, with $\#(\beta_j \cap z_t) = +1$. We have a corresponding intersection point $w \in \gamma \cap \beta_j$. We orient β_j so that this intersection number is negative (so that β_j points in the same direction as α_1).

Now, we have two classes of intersection points $\mathbf{x} \in \mathbb{T}_\gamma \cap \mathbb{T}_\beta$: those which contain w (each of these have the form $w \times \mathbb{T}_\alpha^1 \cap \mathbb{T}_\beta^j$), and those which do not. If $\mathbf{x}$ lies in the first set, then

$$S_{z_1}(\mathbf{x}) = S_{z_2}(\mathbf{x}) + \beta_j^* - \alpha_1^*;$$

if $\mathbf{x}$ lies in the second set, then

$$S_{z_1}(\mathbf{x}) = S_{z_2}(\mathbf{x}) + \beta_j^*.$$

Note that there is an assignment:

$$\sigma'\colon \mathbb{T}_\alpha^1 \cap \mathbb{T}_\beta^j \longrightarrow \mathrm{Spin}^c(Y)$$

obtained by restricting S_{z_2} to $w \times (\mathbb{T}_\alpha^1 \cap \mathbb{T}_\beta^j) \subset \mathbb{T}_\gamma \cap \mathbb{T}_\beta$, and hence a corresponding map

$$\Delta'\colon \mathrm{Spin}^c(Y) \longrightarrow \mathbb{Z}.$$

We have the relation that

$$a_{z_2}(\mathfrak{s}) - a_{z_1}(\mathfrak{s}+\beta_j^*) = \Delta'(\mathfrak{s}) - \Delta'(\mathfrak{s}+\alpha_1^*). \tag{17}$$

It follows from Equations (16) and (17) that

$$\begin{aligned} \chi(\mathfrak{s}) - \chi(\mathfrak{s}+\alpha_1^*) - \chi(\mathfrak{s}+\beta_j^*) + \chi(\mathfrak{s}+\alpha_1^*+\beta_j^*) &= \sum_{n=0}^{\infty} a_{z_2}(\mathfrak{s}+n\alpha_1^*) - a_{z_1}(\mathfrak{s}+n\alpha_1^*+\beta_j^*) \\ &= \sum_{n=0}^{\infty} \Delta'(\mathfrak{s}+n\alpha_1^*) - \Delta'(\mathfrak{s}+(n+1)\alpha_1^*) \\ &= \Delta'(\mathfrak{s}). \end{aligned}$$

(note that Δ' has finite support).

It is easy to see directly from the construction that Δ' and the term $\Delta_{1,j}$ from Equation (15) can differ at most by a sign and a translation with $C_1\alpha_1^*+C_2\beta_j^*$, where C_1 and C_2 are universal constants. Since $\tau(\mathfrak{s})$ and $\chi(HF^+(Y,\mathfrak{s}))$ are three-manifold invariants, by varying β_j^*, it follows that $C_2=0$. A simple calculation in $S^1\times S^2$ shows that $C_1=0$, too. It follows that $\tau(\mathfrak{s})$ must agree with $\pm\chi(HF^+(Y,\mathfrak{s}))$. □

Proof of Theorem 9.1. This is now a direct consequence of Propositions 9.7, 9.8 and 9.9. □

9.2. The Euler characteristic of $HF^+(Y,\mathfrak{s})$ when $b_1(Y)>1$, $\mathfrak{s}$ is non-torsion.

Theorem 9.10. *If $\mathfrak{s}$ is a non-torsion* Spin^c *structure, over an oriented three-manifold Y with $b_1(Y)>1$, then $HF^+(Y,\mathfrak{s})$ is finitely generated, and indeed,*

$$\chi(HF^+(Y,\mathfrak{s})) = \pm\tau(Y,\mathfrak{s}),$$

where τ is Turaev's torsion function.

The proof in subsection 9.1 applies, with the following modifications.

First of all, we use a Heegaard decomposition of Y for which there is a periodic domain $\mathcal{P}$ containing α_1 with multiplicity one in its boundary, and with the property that the induced real cohomology class $c_1(\mathfrak{s})$ is a non-zero multiple of $\mathrm{PD}[\alpha_1^*]$. (This can be arranged after handleslides amongst the α_i.) The subgroup $c_1(\mathfrak{s})^\perp$ of $H_2(Y;\mathbb{Z})$ which pairs trivially with $c_1(\mathfrak{s})$ corresponds to the set of periodic domains $\mathcal{P}$ whose boundary contains α_1 with multiplicity zero. Let $\mathcal{P}_2,...,\mathcal{P}_b$ be a basis for these domains. By winding normal to the $\alpha_2,...,\alpha_g$, we can arrange for all of these periodic domains to have both positive and negative coefficients with respect to any possible choice of base-point on γ. It follows that the Heegaard diagrams constructed above remain weakly admissible for any Spin^c structure. In the present case, the proof of Lemma 9.4 gives the following:

Lemma 9.11. *Fix $\mathfrak{s}$ and an n sufficiently large (in comparison with $\langle c_1(\mathfrak{s}),\mathcal{P}\rangle$). Then, for each γ-induced pair $\mathbf{x}^+$ and $\mathbf{y}^-$ inducing $\mathfrak{s}$, there are at most two homotopy classes modulo the action of $c_1(\mathfrak{s})^\perp$, $[\phi^{\mathrm{in}}],[\phi^{\mathrm{out}}]\in\pi_2(\mathbf{x}^+,\mathbf{y}^-)/c_1(\mathfrak{s})^\perp$ with Maslov index one and with only non-negative multiplicities. Moreover, there are no such classes in $\pi_2(\mathbf{y}^-,\mathbf{x}^+)$.*

Thus, Proposition 9.5 holds in the present context. In fact, the above lemma suffices to construct the ordering. Note that there is no longer a unique map connecting $\mathbf{x}$ to $\mathbf{y}$ with α_1-boundary near γ, with specified multiplicity at z (the map $\phi_{[\mathbf{x},i][\mathbf{y},j]}$ from before), but rather, any two such maps ϕ and ϕ' differ by the addition of periodic domains in $c_1(\mathfrak{s})^{\perp}$. Thus, in view of Theorem 4.1 the Maslov indices of ϕ and ϕ' agree. If we choose the volume form on Σ so that all of $\mathcal{P}_2, ..., \mathcal{P}_g$ have total signed area zero (c.f. Lemma 4.4), then the ordering defined by analogy with the previous subsection is independent of the choice of ϕ or ϕ'.

With these remarks in place, the proof of Theorem 9.1 applies, now proving that $\chi(\mathfrak{s}) = \pm\tau(\mathfrak{s})$, proving Theorem 9.10.

10. Surgery exact sequences

We investigate how surgeries on a three-manifold affect its invariants. We consider first the effect on HF^+ of $+1$ surgeries on integral homology three-spheres, then a generalization which holds for arbitrary (closed, oriented) three-manifolds, then we consider the case of fractional $1/q$-surgeries on an integral homology three-sphere. This latter case uses the homology theories with twisted coefficients. We then give analogous results for $\widehat{HF}$. After this, we present a surgery formula for integer surgeries. In the final subsection, we consider a $+1$ surgery formula with twisted coefficients.

10.1. $+1$ **surgeries on an integral homology three-sphere.** We start with the case of a homology three-sphere Y. Let $K \subset Y$ be a knot. Let Y_0 be the manifold obtained by 0-surgery on K, and Y_1 be obtained by $(+1)$-surgery. Let

$$HF^+(Y_0) \cong \bigoplus_{\mathfrak{s}\in\mathrm{Spin}^c(Y_0)} HF^+(Y_0,\mathfrak{s}),$$

viewed as a $\mathbb{Z}/2\mathbb{Z}$-relatively graded group. In fact, we will view the homology groups $HF^+(Y)$ and $HF^+(Y_1)$ as $\mathbb{Z}/2\mathbb{Z}$-graded, as well.

Theorem 10.1. *There is a U-equivariant exact sequence of relatively $\mathbb{Z}/2\mathbb{Z}$-graded complexes:*

$$... \longrightarrow HF^+(Y) \xrightarrow{F_1} HF^+(Y_0) \xrightarrow{F_2} HF^+(Y_1) \xrightarrow{F_3} ...$$

In fact, if we give $HF^+(Y)$ and $HF^+(Y_1)$ absolute $\mathbb{Z}/2\mathbb{Z}$-gradings so that $\chi(\widehat{HF}(Y)) = \chi(\widehat{HF}(Y_1)) = +1$, then F_3 preserves degree.

The maps in Theorem 10.1 are constructed with the help of holomorphic triangles. Thus, we must construct compatible Heegaard decompositions for all three manifolds Y, Y_0, and Y_1. Exactness is then proved using a filtration on the homology groups above, together with the homological-algebraic constructions used in establishing the surgery sequences for instanton Floer homology (see [9], [3]). The proof occupies the rest of the present subsection.

Lemma 10.2. *There is a pointed Heegaard multi-diagram*

$$(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z)$$

with the property that

(1) *the Heegaard diagrams $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$, $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma})$, and $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\delta})$ describe Y, Y_0, and Y_1 respectively,*
(2) *for each $i=1,...,g-1$, the curves β_i, γ_i, and δ_i are small isotopic translates of one another, each pairwise intersecting in a pair of canceling transverse intersection points (where the isotopies are supported in the complement of z),*
(3) *the curve γ_g is isotopic to the juxtaposition of δ_g and β_g (with appropriate orientations),*
(4) *every non-trivial multi-periodic domain has both positive and negative coefficients.*

Proof. Consider a Morse function on $Y - \mathrm{nd}(K)$ with one index zero critical point, g index one critical points and $g-1$ index two critical points. Let Σ be the 3/2-level of this function, $\boldsymbol{\alpha}$ be the curves where Σ meets the ascending manifolds of the index one critical points in Σ, and let $\beta_1, ..., \beta_{g-1}$ be the curves where Σ meets the descending manifolds of the index two critical points. By gluing in the solid torus in three possible ways, we get the manifolds Y, Y_0, Y_1. Extending the given Morse function to the glued in solid tori, (by introducing an additional index two and index three critical point), we obtain Heegaard decompositions for the manifolds Y, Y_0, and Y_1. We let γ_i and δ_i be small perturbations of β_i for $i = 1, ..., g-1$. In this manner, we have satisfied Properties (1)-(3).

To satisfy Property (4), we wind to achieve weak admissibility for all Spin^c structures for the Heegaard subdiagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, z)$: in fact, we can use a volume form over Σ for which all such doubly-periodic domains have zero signed area (c.f. Lemma 4.4). Then, for the $\{\beta_1, ...\beta_{g-1}\}$ and $\{\delta_1, ..., \delta_{g-1}\}$, we use small Hamiltonian translates of the $\{\gamma_1, ..., \gamma_{g-1}\}$ (ensuring that the corresponding new periodic domains each have zero energy). There is a triply-periodic domain which forms the homology between β_g, γ_g, and δ_g in a torus summand of Σ containing no other β_i or γ_i (for $i \neq g$). By adjusting the areas of the two triangles with non-zero area, we can arrange for the signed area of the triply-periodic domain to vanish.

□

For $i = 1, ..., g-1$, label

$$y_i^{\pm} = \beta_i \cap \gamma_i,\ v_i^{\pm} = \gamma_i \cap \delta_i,\ w_i^{\pm} = \beta_i \cap \delta_i,$$

where the sign indicates the sign of the intersection point. Also, let

$$y_g = \beta_g \cap \gamma_g,\ v_g = \gamma_g \cap \delta_g,\ w_g = \beta_g \cap \delta_g.$$

Then, let $\Theta_{\beta,\gamma} = \{y_1^+, ..., y_{g-1}^+, y_g\}$, $\Theta_{\gamma,\delta} = \{v_1^+, ..., v_{g-1}^+, v_g\}$, $\Theta_{\beta,\delta} = \{w_1^+, ..., w_{g-1}^+, w_g\}$ denote the corresponding intersection points between $\mathbb{T}_\beta \cap \mathbb{T}_\gamma$, $\mathbb{T}_\gamma \cap \mathbb{T}_\delta$ and $\mathbb{T}_\beta \cap \mathbb{T}_\delta$. (See Figure 6 for an illustration.)

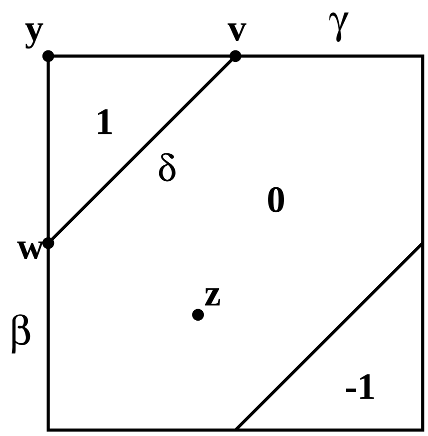

FIGURE 5. This picture takes place in the torus, with the usual edge identifications. The integers denote multiplicities for a triply-periodic domain.

Proposition 10.3. *The elements* $\theta_{\beta,\gamma} = [\Theta_{\beta,\gamma}, 0]$*,* $\theta_{\gamma,\delta} = [\Theta_{\gamma,\delta}, 0]$*,* $\theta_{\beta,\delta} = [\Theta_{\beta,\delta}, 0]$ *are cycles in* $CF^\infty(\mathbb{T}_\beta, \mathbb{T}_\gamma)$*,* $CF^\infty(\mathbb{T}_\gamma, \mathbb{T}_\delta)$ *and* $CF^\infty(\mathbb{T}_\beta, \mathbb{T}_\delta)$ *respectively.*

Proof. Note that the three-manifolds described here are $(g-1)$-fold connected sums of $S^1 \times S^2$, so the result follows from Proposition 7.10 (or, alternatively, see Section 5 of [23]). □

We can reduce the study of holomorphic triangles belonging to $X_{\beta,\gamma,\delta}$ to holomorphic triangles in the first symmetric product of the two-torus, with the help of the following analogue of the gluing theory used to establish stabilization invariance of the Floer homology groups.

Theorem 10.4. *Fix a pair of Heegaard diagrams*

$$(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, z) \quad \textit{and} \quad (E, \beta_0, \gamma_0, \delta_0, z_0),$$

where E *is a Riemann surface of genus one. We will form the connected sum* $\Sigma \# E$*, where the connected sum points are near the distinguished points* z *and* z_0 *respectively. Fix intersection points* $\mathbf{x}, \mathbf{y}, \mathbf{w}$ *for the first diagram and a class* $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$*, and intersection points* x_0*,* y_0*, and* w_0 *for the second, with a triangle* $\psi_0 \in \pi_2(x_0, y_0, w_0)$ *with* $\mu(\psi) = \mu(\psi_0) = 0$*. Suppose moreover that* $n_{z_0}(\psi_0) = 0$*. Then, for a suitable choice of complex structures and perturbations, we have a diffeomorphism of moduli spaces:*

$$\mathcal{M}(\psi') \cong \mathcal{M}(\psi) \times \mathcal{M}(\psi_0),$$

where $\psi' \in \pi_2(\mathbf{x} \times x_0, \mathbf{y} \times y_0, \mathbf{w} \times w_0)$ *is the triangle for* $\Sigma \# E$ *whose domain on the* Σ*-side agrees with* $\mathcal{D}(\psi)$*, and whose domain on the* E*-side agrees with* $\mathcal{D}(\psi_0) + n_z(\psi)[E]$*.*

Proof. The proof is obtained by suitably modifying Theorem 6.4 of [23].

Suppose that u and u_0 are holomorphic representatives of ψ and ψ_0 respectively. We obtain a nodal pseudo-holomorphic disk $u \vee u_0$ in the singular space $\mathrm{Sym}^{g+1}(\Sigma \vee E)$ specified as follows:

- At the stratum $\mathrm{Sym}^g(\Sigma) \times \mathrm{Sym}^1(E)$, $u \vee u_0$ is the product map $u \times u_0$.
- At the stratum $\mathrm{Sym}^{g-1}(\Sigma) \times \mathrm{Sym}^2(E)$, $u \vee u_0$ is given by $n_z(\psi)$ pseudo-holomorphic spheres which are constant on the first factor. More precisely, for each $p \in \Delta$ for which $u(p) = \{z, x_2, ..., x_g\}$ (where the $x_i \in \Sigma - \{z\}$ are arbitrary), there is a component of $u \vee u_0$ mapping into $\mathrm{Sym}^{g-1}(\Sigma) \times \mathrm{Sym}^2(E)$, consisting of the product of the constant map $\{x_2, ..., x_g\}$ with the sphere in $\mathrm{Sym}^2(E)$ which passes through $\{z\} \times u_0(p)$.
- The map $u \vee u_0$ misses all other strata of $\mathrm{Sym}^{g+1}(\Sigma \vee E)$.

As in Theorem 6.4 of [23], we can splice to obtain an approximately holomorphic disk $u \# u_0$ (a triangle) in $\mathrm{Sym}^{g+1}(\Sigma \# E)$. When the connected sum tube is sufficiently long, the the inverse function theorem can be used to find the nearby pseudo-holomorphic triangle. The domain belonging to $u \# u_0$ is clearly given by $\psi \# \psi_0$ described above. Conversely, by Gromov's compactness (see also Proposition 6.15 of [23]), any sequence of pseudo-holomorphic representatives $u_i \in \pi_2(\mathbf{x} \times x_0, \mathbf{y} \times y_0, \mathbf{w} \times w_0)$ for arbitrarily long connected sum neck must limit to a pseudo-holomorphic representative for $\psi' \# \psi_0'$, where $\mathcal{D}(\psi_0') - \mathcal{D}(\psi_0) = k[E]$ for some $0 \leq k \leq n_z(\psi)$. However, since $\pi_2(E) = 0$, it follows that $k = 0$. Thus, the gluing map covers the moduli space. □

Proposition 10.5. *There are homotopy classes of triangles $\{\psi_k^{\pm}\}_{k=1}^{\infty}$ in $\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ for the triple-diagram $(\Sigma,\boldsymbol{\beta},\boldsymbol{\gamma},\boldsymbol{\delta},z)$ satisfying the following properties:*

$$\begin{aligned}\mu(\psi_k^{\pm}) &= 0,\\ n_z(\psi_k^{\pm}) &= \frac{k(k-1)}{2}.\end{aligned}$$

Moreover, each triangle in $\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ is Spin^c equivalent to some $\psi_k^{\pm}$. Furthermore, there is a choice of perturbations and complex structure on Σ with the property that for each $\Psi\in\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\mathbf{x})$ (where $\mathbf{x}\in\mathbb{T}_\beta\cap\mathbb{T}_\delta$) with $\mu(\Psi)=0$, we have that

$$\#\mathcal{M}(\Psi)=\left\{\begin{array}{ll}\pm 1 & \text{if } \Psi\in\{\psi_k^{\pm}\}_{k=1}^{\infty}\\ 0 & \text{otherwise}\end{array}\right. .$$

Proof. First observe that the space of Spin^c structures over $X_{\beta,\gamma,\delta}$ extending a given one on the boundary is identified with $\mathbb{Z}$. In particular, modulo doubly-periodic domains for the three boundary three-manifolds, every triangle $\psi\in\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ can uniquely be written as $\psi_1+a[S]+b[\mathcal{P}]$ for some pair of integers a and b, where $\mathcal{P}$ is the generator of the space of triply-periodic domains: in fact, the integer a is determined by the intersection number n_z, and b can be determined by the signed number of times the arc in β_g obtained by restricting ψ to its boundary crosses some fixed $\tau\in\beta_g$. For the triangles $\{\psi_k^{\pm}\}$ this signed count can be any arbitrary integer, so these triangles represent all possible Spin^c-equivalence classes of triangles.

The other claims are straightforward in the case where $g=1$. In this case, the curves β, γ, δ lie in a surface of genus one, so the holomorphic triangle can be lifted to the complex plane. Hence, by standard complex analysis, it is smoothly cut out, and unique.

The fact that $\#\mathcal{M}(\psi_k^{\pm})=\pm 1$ for higher genus follows from induction, and the gluing result, Theorem 10.4. Specifically, if the result is known for genus g, then we can add a new torus E to Σ which contains three curves β_0, γ_0, δ_0 which are small Hamiltonian translates of one another (and the basepoint is chosen outside the support of the isotopy). The torus E contains a standard small triangle $\psi_0\in\pi_2(y_0^+,v_0^+,w_0^+)$, for which it is clear that $\#\mathcal{M}(\psi_0)=1$. Gluing this triangle to the $\{\psi_k^{\pm}\}$ in Σ, we obtain corresponding triangles in $\Sigma\#E$ satisfying all the above hypotheses.

The fact that $\#\mathcal{M}(\Psi)=0$ for $\Psi\notin\{\psi_k^{\pm}\}_{k=1}^{\infty}$ follows similarly, with the observation that the other moduli spaces of triangles on the torus are empty. □

We can define the map

$$F_1\colon HF^+(Y)\longrightarrow HF^+(Y_0)$$

by summing:

$$F_1(\xi)=\sum_{\mathfrak{s}\in\mathrm{Spin}^c(X_{\alpha,\beta,\gamma})}\pm F^+_{\alpha,\beta,\gamma}(\xi\otimes\theta_{\beta,\gamma},\mathfrak{s}).$$

On the chain level, F_1 is induced from a map:

$$f_1([\mathbf{x},i])=\sum_{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma}\;\sum_{\{\psi\in\pi_2(\mathbf{x},\Theta_{\beta,\gamma},\mathbf{w})\,\big|\,\mu(\psi)=0\}}(\#\mathcal{M}(\psi))\cdot[\mathbf{w},i-n_z(\psi)],$$

where $\#\mathcal{M}(\psi)$ is calculated with respect to a particular choice of coherent orientation system (see Proposition 10.6 below). It is important to note here that the sum on the right hand side will have only finitely many non-zero elements for each fixed $\xi \in CF^+(Y)$. The reason for this is that all the multi-periodic domains have both positive and negative coefficients. Similarly, we define

$$f_2([\mathbf{x}, i]) = \sum_{\{\psi \in \pi_2(\mathbf{x}, \Theta_{\gamma,\delta}, \mathbf{w}) \big| \mu(\psi)=0\}} (\#\mathcal{M}(\psi)) \cdot [\mathbf{w}, i - n_z(\psi)],$$

letting F_2 be the induced map on homology.

Observe that the maps f_1 and f_2 preserve the relative $\mathbb{Z}/2\mathbb{Z}$-grading. The reason for this is that the parity of the Maslov index of a triangle $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ depends only on the sign of the local intersection numbers of the $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\mathbb{T}_\beta \cap \mathbb{T}_\gamma$, and $\mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ at $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{w}$. (Each local intersection number is calculated by orienting the three tori consistently, but their product is independent of the choice of orientations.)

Proposition 10.6. *For any coherent system of orientations for Y_0, we can find coherent systems of orientations for the triangles defining f_1 and f_2 so that the composition $F_2 \circ F_1 = 0$.*

Proof. For any system of coherent orientations, associativity, together with Proposition 10.5, can be interpreted as saying that

$$\sum_{s_{\beta,\gamma,\delta} \in S_{\beta,\gamma,\delta}} f^{\leq 0}_{\beta,\gamma,\delta}(\theta_{\beta,\gamma} \otimes \theta_{\gamma,\delta}) = \sum_{k=1}^{\infty} \left[\Theta_{\beta,\delta}, -\frac{k(k-1)}{2}\right] \pm \left[\Theta_{\beta,\delta}, -\frac{k(k-1)}{2}\right]$$

(up to an overall sign), as a formal sum.

Of course, if we are using only $\mathbb{Z}/2\mathbb{Z}$ coefficients, the proof is complete.

Note that the orientation system for $Y_{\beta,\delta}$ is chosen so that $\Theta_{\beta,\delta}$ is a cycle. But this leaves the orientation system over $Y_{\alpha,\gamma}$ unconstrained, and any choice of such orientation system determines the choice over $X_{\alpha,\beta,\gamma}$ (up to an overall sign depending on the Spin^c structure used over $Y_{\alpha,\gamma}$). Now, the relative sign appearing above corresponds to the orientation of the triangles ψ_k^+ vs. the triangles ψ_k^- over $X_{\beta,\delta,\gamma}$, and each such pair of triangles belongs to different $\delta H^1(Y_{\alpha,\delta}) + \delta H^1(Y_{\beta,\delta})$-orbits for the square $X_{\alpha,\beta,\gamma,\delta}$. Thus, we can modify the relative sign at will. We choose it so that the terms pairwise cancel. □

We can choose a one-parameter family of γ-curves $\gamma_i(t)$ with the property that $\lim_{t \mapsto 0} \gamma_i(t) = \beta_i$ for $i = 1, ..., g-1$, and $\lim_{t \mapsto 0} \gamma_g(t) = \delta_g * \beta_g$ (juxtaposition of curves), and we choose our basepoint z to lie outside the support of the homotopies $\gamma_i(t)$. We choose another one-parameter family of δ-curves $\delta_i(t)$ for $i = 1, ..., g-1$ with $\lim_{t \mapsto 0} \delta_i(t) = \beta_i$. We assume that all α_i are disjoint from the $\beta_g \cap \delta_g$. Then, if t is sufficiently small, then there is a natural partitioning of $\mathbb{T}_\alpha \cap \mathbb{T}_{\gamma(t)}$ into two subsets, those which are nearest to points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and those which are nearest to points in $\mathbb{T}_\alpha \cap \mathbb{T}_\delta(t)$. (See Figure 6 for an illustration.) Correspondingly, we have a splitting

$$CF^+(Y_0) \cong CF^+(Y) \oplus CF^+(Y_1);$$

or, a short exact sequence of graded groups

$$0 \longrightarrow CF^+(Y) \xrightarrow{\iota} CF^+(Y_0) \xrightarrow{\pi} CF^+(Y_1) \longrightarrow 0$$

with splitting

$$R\colon CF^+(Y_1) \longrightarrow CF^+(Y_0),$$

where the maps ι, π, and R are not necessarily chain maps. Our goal is to construct a short exact sequence as above, which is compatible with the boundary maps.

Proposition 10.7. *The map f_1 is chain homotopic to a U-equivariant chain map g_1 with the property that*

$$0 \longrightarrow CF^+(Y) \xrightarrow{g_1} CF^+(Y_0) \xrightarrow{f_2} CF^+(Y_1) \longrightarrow 0.$$

is a short exact sequence of chain complexes.

Theorem 10.1 is a consequence of this proposition: the associated long exact sequence is the exact sequence of Theorem 10.1.

For the construction of g_1, we need the following ingredients:

- lower-bounded filtrations on the $CF^+(Y)$, $CF^+(Y_0)$, and $CF^+(Y_1)$, which are strictly decreasing for the boundary maps; i.e. each chain complex is generated by elements with $\partial\xi < \xi$.
- an injection ι and splitting map R as above, both of which respect the filtrations
- decompositions of $f_1 = \iota +$ lower order and $f_2 = \pi +$ lower order, where, here, lower order is with respect to the filtrations. More precisely $CF^+(Y)$ is generated by elements

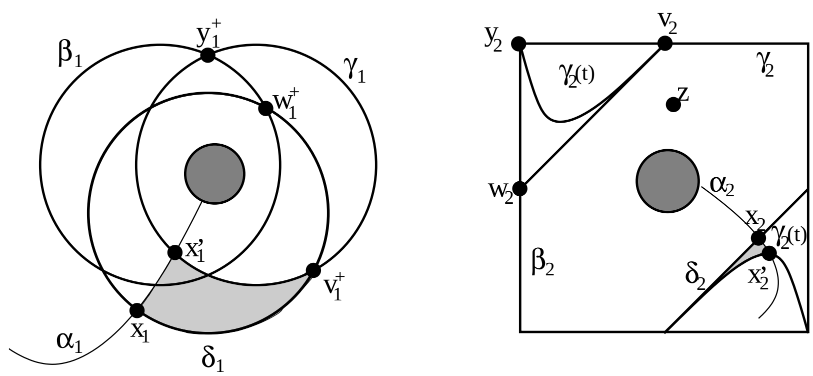

FIGURE 6. $+1$-surgery, $g = 2$. The left side takes place in an annulus, the right side in a torus minus a disk; both are pieces of our genus two surface Σ (the central disk missing from the annulus and the disk removed from the torus are both indicated by large grey circles). We have curves $\{\beta_1, \beta_2\}$, $\{\gamma_1, \gamma_2\}$ and $\{\delta_1, \delta_2\}$ as in Lemma 10.2, with intersection points $\Theta_{\beta,\gamma} = \{y_1^+, y_2\}$, $\Theta_{\gamma,\delta} = \{v_1^+, v_2\}$, and $\Theta_{\beta,\delta} = \{w_1^+, w_2\}$. The curve $\gamma_2(t)$ is isotopic to γ_2, but it approximates the juxtaposition of β_2 and δ_2. We have also pictured arcs in α_1 and α_2. There is an intersection point $\mathbf{x} = \{x_1, x_2\}$ for $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$, and its nearest point $\mathbb{T}_\alpha \cap \mathbb{T}_{\gamma(t)}$, $\{x_1', x_2'\} = \rho(\mathbf{x})$. Observe the two lightly shaded triangles: they correspond to the canonical triangle in $\pi_2(\rho(\mathbf{x}), \Theta_{\gamma,\delta}, \mathbf{x})$.

ξ with the property that $f_1(\xi) - \iota(\xi) < \iota(\xi)$, and $CF^+(Y_1)$ is generated by elements η with $\eta - f_2 \circ R(\eta) < \eta$.

- $f_2 \circ f_1$ is chain homotopic to zero by a U-equivariant homotopy

$$H\colon CF^+(Y) \longrightarrow CF^+(Y_1)$$

which decreases filtrations, in the sense that $R \circ H < \iota$.

Following Lemma 9 of [3], we define a right inverse R' for f_2 by

$$R' = R \circ \sum_{k=0}^{\infty} (\mathrm{Id} - f_2 \circ R)^{\circ k},$$

and let

$$g_1 = f_1 - (\partial(R' \circ H) + (R' \circ H)\partial);$$

so that our hypotheses ensure that $g_1 = \iota + \text{lower order}$. It follows that if L is the left inverse of ι induced from R, then $L \circ g_1$ is invertible, as $L \circ g_1(\xi) = \xi - N(\xi)$, where N decreases filtration (so we can define

$$(L \circ g_1)^{-1}(\xi) = \sum_{k=0}^{\infty} N^{\circ k}(\xi),$$

as the sum on the right contains only finitely many non-zero terms for each $\xi \in CF^+(Y)$); thus, $(L \circ g_1)^{-1} \circ L$ is a left inverse for g_1.

A similar argument shows surjectivity of f_2, and exactness at the middle stage (see [3]).

We will use an energy filtration on $CF^+(Y_0)$ defined presently. First, fix an $\mathbf{x}_0 \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. If $[\mathbf{y}, j] \in CF^+(Y_0)$, let $\psi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \mathbf{y})$ be a (homotopy class of) triangle, with $n_z(\psi) = -j$. We then define

$$\mathcal{F}_{Y_0}([\mathbf{y}, j]) = -\mathcal{A}(\psi).$$

(Note that $\pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \mathbf{y})$ is non-empty.) As in Lemma 4.4, the topological hypothesis from Lemma 10.2 allows us to use a volume form on Σ for which every periodic domain for Y_0 has zero area: every periodic domain for $(\mathbb{T}_\beta, \mathbb{T}_\delta)$, $(\mathbb{T}_\beta, \mathbb{T}_{\gamma(t)})$ and also the triply-periodic domain for $(\beta_g, \gamma_g(t), \delta_g)$ has area zero. (For example, we can start with the area form constructed in the proof of Lemma 10.2 for the initial $t = 0$ γ-curves, and then move those curves by a Hamiltonian isotopy.) The real-valued function $\mathcal{F}_{Y_0}$ on the generators of $CF^+(Y_0)$ gives the latter group an obvious partial ordering.

We will assume now that the $\gamma_g(t)$ is sufficiently close to the juxtaposition of β_g and δ_g, in the following sense. Let $\mathcal{P}$ be a triply-periodic domain between $\gamma_g(t)$, β_g, and δ_g which generates the group of such periodic domains (this is the domain pictured in Figure 5, before γ_g was isotoped); and for $i = 1, ..., g-1$, let $\mathcal{P}_i$ be the doubly-periodic domains with $\partial\mathcal{P}_i = \beta_i - \gamma_i(t)$. We let $\epsilon(t)$ be the sum of the absolute areas of all these periodic domains:

$$\epsilon(t) = \mathcal{A}\Big(|\mathcal{D}(\mathcal{P})|\Big) + \sum_{i=1}^{g-1} \mathcal{A}\Big(|\mathcal{D}(\mathcal{P}_i)|\Big),$$

where here the absolute signs denote the unsigned area. Note that $\lim_{t \mapsto 0} \epsilon(t) = 0$. Also, let M be the minimum of the area of any domain in $\Sigma - \alpha_1 - ... - \alpha_g - \beta_1 - ... - \beta_g - \delta_g$. We choose t small enough that $\epsilon(t) < M/2$. We assume that the absolute (unsigned) area of the periodic domain $\mathcal{Q}_i$ with $\partial(\mathcal{Q}_i) = \beta_i - \delta_i(t)$ agrees with the absolute area of $\mathcal{P}_i$.

Lemma 10.8. *For sufficiently small t, the function $\mathcal{F}_{Y_0}$ induces a filtration on $CF^+(Y_0)$. In particular,*

$$\partial[\mathbf{y}, j] < [\mathbf{y}, j].$$

Proof. It is important to observe that the area filtration defined above is indeed well-defined. The reason for this is that if ψ, ψ' are a pair of homotopy classes in $\pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \mathbf{y})$ with $n_z(\psi) = n_z(\psi')$, then $\mathcal{D}(\psi) - \mathcal{D}(\psi')$ is a triply-periodic domain. It follows from above that it must have total area zero.

Suppose that we have a pair of generators $[\mathbf{y}, j]$ and $[\mathbf{y}', j']$ which are connected by a flow ϕ. If $\psi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \mathbf{y})$ is a class with $n_z(\psi) = -j$, then, of course, $\psi + \phi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \mathbf{y}')$ is a class with $n_z(\psi + \phi) = -j'$; thus, $\mathcal{F}_{Y_0}([\mathbf{y}', j']) - \mathcal{F}_{Y_0}([\mathbf{y}, j]) = -\mathcal{A}(\phi)$; but $\mathcal{A}(\phi) > 0$, as all of its coefficients are non-negative (and at least one is positive). □

The filtration on CF^+, together with the data ι, π, and R, endow $CF^+(Y)$ and $CF^+(Y_1)$ with a filtration as well.

Lemma 10.9. *For t sufficiently small, the orderings induced on $CF^+(Y)$ and $CF^+(Y_1)$ give filtrations.*

Proof. There is a natural filtration on Y, defined by $\mathcal{F}_Y([\mathbf{x}, i]) = -\mathcal{A}(\phi)$, where $\phi \in \pi_2(\mathbf{x}_0, \mathbf{x})$ is the class with $n_z(\phi) = -i$. This is a filtration, in view of the usual positivity of holomorphic disks (see Theorem 3.1); indeed, the filtration decreases by at least M along flows.

The filtration induced by $\mathcal{F}_{Y_0}$ and the map ι, defined by $\mathcal{F}'_Y([\mathbf{x}, i]) = \mathcal{F}_{Y_0}(\iota[\mathbf{x}, i])$ very nearly agrees with this natural filtration, for sufficiently small t. To see this, note that there is a unique "small" triangle $\psi_0 \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \iota(\mathbf{x}))$ which has non-negative coefficients and is supported inside the support of $\mathcal{P} + \mathcal{P}_1 + ...\mathcal{P}_{g-1}$. Clearly, $\mathcal{A}(\psi_0) < \epsilon(t)$, and $n_z(\psi_0) = 0$. Now, if $\phi \in \pi_2(\mathbf{x}_0, \mathbf{x})$ is the class with $n_z(\phi) = -i$ the juxtaposition of $\psi_0 + \phi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \iota(\mathbf{x}))$ can be used to calculate the Y_0 filtration of $\iota(\mathbf{x})$; thus $|\mathcal{F}_Y([\mathbf{x}, i]) - \mathcal{F}'_Y([\mathbf{x}, i])| < \epsilon(t)$. In particular, since $\mathcal{F}_Y$ decreases by at least M along flowlines, $\mathcal{F}_{Y_0} \circ \iota$, too, must decrease along flows.

For Y_1, there is another filtration, this one induced by squares. Given $[\mathbf{y}, i] \in (\mathbb{T}_\alpha \cap \mathbb{T}_\delta) \times \mathbb{Z}^{\geq 0}$, consider $\varphi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \Theta_{\gamma,\delta}, \mathbf{y})$ with $n_z(\varphi) = -i$, and let

$$\mathcal{F}''_{Y_1}([\mathbf{y}, i]) = -\mathcal{A}(\mathcal{D}(\varphi)).$$

Indeed, if M' is the minimum area of any domain in $\Sigma - \alpha_1 - ... - \alpha_g - \delta_1(t) - ... - \delta_{g-1}(t) - \delta_g$, then $\mathcal{F}''_{Y_1}$ decreases by at least M' along each flowline. Note that $M' > M - \epsilon(t)$.

Now, we claim that $\mathcal{F}''_{Y_1}$ nearly agrees with the filtration $\mathcal{F}'_{Y_1}$ induced by $\mathcal{F}_{Y_0}$ and the right inverse R: $\mathcal{F}'_{Y_1}([\mathbf{y}, j]) = \mathcal{F}_{Y_0}(R[\mathbf{y}, j])$. Again, if we let $\rho(\mathbf{y})$ denote the point in $\mathbb{T}_\alpha \cap \mathbb{T}\gamma(t)$ closest to $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$, there is a unique small triangle $\psi_0 \in \pi_2(\rho(\mathbf{y}), \Theta_{\gamma,\delta}, \mathbf{y})$. If $\psi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \rho(\mathbf{y}))$ is a triangle with $n_z(\psi) = -j$ (i.e. used to calculate $\mathcal{F}_{Y_0} \circ R$), then, the juxtaposition $\psi + \psi_0$ is a square which can be used to calculate $\mathcal{F}''_{Y_1}([\mathbf{y}, j])$. But $|\mathcal{A}(\psi + \psi_0) - \mathcal{A}(\psi)| \leq \epsilon(t)$, so since $\mathcal{F}''_{Y_1}$ decreases by at least M' for non-trivial flows, it follows that $\mathcal{F}_{Y_0} \circ R$, too, must decrease along flows. □

Lemma 10.10. *The maps f_1 and f_2 have the form:*

$$f_1 = \iota + \textit{lower order}, \qquad f_2|_{\mathrm{Im} R} = \pi + \textit{lower order}$$

Proof. The map $f_1([\mathbf{x}, i])$ counts the number of holomorphic triangles in homotopy classes with $\psi \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \mathbf{y})$, with $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\gamma(t)}$ and $\mu(\psi) = 0$. One of these triangles, of course is the canonical small triangle $\psi_0 \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \iota(\mathbf{x}))$. One can calculate that $\#\mathcal{M}(\psi_0) = 1$. This gives the ι component of f_1. Now, no other homotopy class $\psi \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \mathbf{y})$ with $\mathcal{D}(\psi) \geq 0$ has its domain $\mathcal{D}(\psi)$ contained inside the support of $\mathcal{P} + \mathcal{P}_1 + ... \mathcal{P}_{g-1}$; thus, if $\mathcal{M}(\psi)$ is non-empty, then $\mathcal{A}(\psi) > M - \epsilon(t) > M/2$. Moreover, in the proof of Lemma 10.9, we saw that if $\phi \in \pi_2(\mathbf{x}_0, \mathbf{x})$ is the homotopy class with $n_z(\phi) = -i$, then

$$|\mathcal{F}_{Y_0}(\iota([\mathbf{x}, i])) + \mathcal{A}(\phi)| < \epsilon(t).$$

But now $\psi + \phi$ can be used to calculate the filtration $\mathcal{F}_{Y_0}([\mathbf{y}, i - n_z(\psi)])$. Thus,

$$\mathcal{F}_{Y_0}([\mathbf{y}, i - n_z(\psi)]) - \mathcal{F}_{Y_0}(\iota[\mathbf{x}, i]) \leq -\mathcal{A}(\psi) + \epsilon(t) < 0.$$

Next, we consider f_2. As before, if $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$, we let $\rho(\mathbf{y}) \in \mathbb{T}_\alpha \cap \mathbb{T}_{\gamma(t)}$ denote the intersection point closest to $\mathbf{y}$. Suppose that $f_2([\rho(\mathbf{y}), i])$ has a non-zero component in $[\mathbf{w}, j]$ with $[\mathbf{y}, i] \neq [\mathbf{w}, j]$; thus, we have a $\psi \in \pi_2(\rho(\mathbf{y}), \Theta_{\gamma,\delta}, \mathbf{w})$ with $n_z(\psi) = i - j$, which supports a holomorphic triangle. Again, ψ cannot be supported inside the support of $\mathcal{P} + \mathcal{P}_1 + ... + \mathcal{P}_{g-1}$, so $\mathcal{A}(\psi) > M/2$. Fix $\psi_{\mathbf{w}} \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \rho(\mathbf{w}))$ (for $\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma$) with $n_z(\psi_{\mathbf{w}}) = -j$, and $\psi_{\mathbf{y}} \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \rho(\mathbf{y}))$ with $n_z(\psi_{\mathbf{y}}) = -i$. Clearly, the juxtaposition $\psi_{\mathbf{y}} + \psi \in \pi_2(\mathbf{x}_0, \Theta_{\beta,\gamma}, \Theta_{\gamma,\delta}, \mathbf{w})$ is a square whose area must agree with the square $\psi_{\mathbf{w}} + \psi_0$, where $\psi_0 \in \pi_2(\rho(\mathbf{w}), \Theta_{\gamma,\delta}, \mathbf{w})$ is the canonical small triangle, so

$$\mathcal{A}(\psi_{\mathbf{w}}) = \mathcal{A}(\psi_{\mathbf{y}}) - \mathcal{A}(\psi_0) + \mathcal{A}(\psi),$$

and hence $\mathcal{F}([\rho(\mathbf{y}), i]) > \mathcal{F}([\rho(\mathbf{w}), j])$. □

Lemma 10.11. *For sufficiently small t, there is a null-homotopy H of $f_2 \circ f_1$ satisfying $R \circ H < \iota$.*

Proof. Theorem 6.15 provides a null-homotopy H. The $[\mathbf{y}, j]$ coefficient of $H[\mathbf{x}, i]$ counts holomorphic squares $\varphi \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \Theta_{\gamma,\delta}, \mathbf{y})$ with $n_z(\varphi) = i - j$.

Our aim here is to prove that if the $[\mathbf{y}, j]$ component of $H[\mathbf{x}, i]$ is non-zero then $\iota[\mathbf{x}, i] > R[\mathbf{y}, j]$. Now, the filtration difference between $\iota([\mathbf{x}, i])$ and $R[\mathbf{y}, j]$ is calculated (to within $\epsilon(t)$) by $\mathcal{A}(\psi)$, where $\psi \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \rho(\mathbf{y}))$ has $n_z(\psi) = i - j$. Adding the smallest triangle in $\pi_2(\rho(\mathbf{y}), \Theta_{\gamma,\delta}, \mathbf{y})$ (and hence changing the area by no more than $\epsilon(t)$), we obtain another square $\varphi' \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma}, \Theta_{\gamma,\delta}, \mathbf{y})$ with $n_z(\varphi') = i - j$, whose area must agree with the area of φ. Now if t is sufficiently small ($\epsilon(t) < M/4$), it follows that the filtration difference between $\iota[\mathbf{x}, i]$ and $R[\mathbf{y}, j]$ is positive. □

Proof of Theorem 10.1. Theorem 10.1 is now a consequence of the long exact sequence associated to the short exact sequence from Proposition 10.7, with a few final observations regarding the $\mathbb{Z}/2\mathbb{Z}$ grading.

Orient the $\alpha_1, ..., \alpha_g$, the $\beta_1, ..., \beta_{g-1}$ arbitrarily (hence inducing orientations on the $\gamma_1, ..., \gamma_{g-1}$ and the $\delta_1, ..., \delta_{g-1}$). The orientation on β_g is then forced on us by the requirement that

$$1 = \chi(\widehat{HF}(Y)) = \#(\mathbb{T}_\alpha \cap \mathbb{T}_\beta),$$

where we orient the tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ in the obvious manner. Similarly, the orientation on δ_g is forced; indeed, so that

$$\delta_g = \beta_g \pm \gamma_g$$

We can orient γ_g so that the above sign is positive. It is then clear with these conventions (by looking at the small triangles) that F_1 preserves the absolute $\mathbb{Z}/2\mathbb{Z}$ grading, while F_2 reverses it. It follows then that F_3 preserves degree as claimed.

□

10.2. **A generalization.** Let Y be an oriented three-manifold, and let $K \subset Y$ be a knot. Let m be the meridian of K, and let $h \in H_1(\partial(Y - \mathrm{nd}(K)))$ be a homology class with $m \cdot h = 1$ (here, the torus is oriented as the boundary of the neighborhood of K). We let Y_h denote the three-manifold obtained by attaching a solid torus to $Y - \mathrm{nd}(K)$, with framing specified by h.

Fix a Spin^c structure $\mathfrak{s}_0$ over $Y - K$. We let

$$HF^+(Y_h, [\mathfrak{s}_0]) = \bigoplus_{\{\mathfrak{s} \big| \mathfrak{s}|_{Y-K} = \mathfrak{s}_0\}} HF^+(Y_h, \mathfrak{s}).$$

We define $HF^+(Y, [\mathfrak{s}_0])$ similarly.

The following is a direct generalization of Theorem 10.1 (the case where Y is an integer homology three-sphere, and h is the "longitude" of K):

Theorem 10.12. *For each* Spin^c *structure* $\mathfrak{s}_0$ *on* $Y - K$*, we have the* U*-equivariant exact sequence:*

$$... \longrightarrow HF^+(Y, [\mathfrak{s}_0]) \longrightarrow HF^+(Y_h, [\mathfrak{s}_0]) \longrightarrow HF^+(Y_{h+m}, [\mathfrak{s}_0]) \longrightarrow ...$$

Corollary 10.13. *Let* Y *be an integer homology three-sphere with a knot* $K \subset Y$*, and let* Y_n *be the three-manifold obtained by* n *surgery on* K *where* $n > 0$*, then there is a* U*-equivarant long exact sequence*

$$... \longrightarrow HF^+(Y) \longrightarrow HF^+(Y_n) \longrightarrow HF^+(Y_{n+1}) \longrightarrow ...$$

The proof given in the previous section adapts to this context, after a few observations.

Note first that the map from Y to Y_h defined by counting triangles is naturally partitioned into equivalence classes. To see the decomposition agrees with what we have stated, we observe the following. Let X be the pair-of-pants cobordism connecting Y, Y_h, and $\#^{g-1}(S^2 \times S^1)$. The four-manifold obtained by filling the last component with $\#^{g-1}(D^3 \times S^1)$ is the cobordism W_h from Y to Y_h obtained by attaching a two-handle to Y along K with framing h.

Now, Spin^c-equivalence classes of triangles for $\mathbb{T}_\alpha$, $\mathbb{T}_\beta$, $\mathbb{T}_\gamma$ agree with Spin^c structures on the cobordism W_h, since $\mathfrak{s}_z(\Theta_{\beta,\gamma})$ is a torsion Spin^c structure over $\#^{g-1}(S^2 \times S^1)$ (which extends uniquely over $\#^{g-1}(D^3 \times S^1)$). But two Spin^c structures on Y and Y_h extend over W_h if and only if they agree on the knot complement $Y - K$ (thought of as a subset of both Y and Y_h).

With this said, the maps f_1 and f_2 partition according to Spin^c structures on $Y - K$.

Next, we observe that there are in principle many periodic domains for the triple $(\mathbb{T}_\alpha, \mathbb{T}_\beta, \mathbb{T}_\gamma)$. By twisting normal to the $\boldsymbol{\alpha}$, however, we can arrange that the triple is admissible. By choosing the volume form on Σ appropriately, we can arrange that they all have zero signed area.

We can define the filtrations as before. Fix any $\mathbf{x}_0 \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ so that $\mathfrak{s}_z(\mathbf{x}_0)$ restricts to $\mathfrak{s}_0$ on $Y - K$. The triangle connecting $\mathbf{x}_0$, $\Theta_{\beta,\gamma}$ and any intersection point $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ with $\mathfrak{s}_z(\mathbf{y})|Y - K = \mathfrak{s}_0$ is guaranteed to exist, since the corresponding Spin^c structures extend over W_h. The area the of the domain of any such triangle can be used to define $\mathcal{F}_{Y_h}([\mathbf{y}, i])$. The proof given before, then, applies.

10.3. **Fractional Surgeries.** There are other directions to generalize Theorem 10.1. We consider presently the case of fractional $(1/q)$ surgeries on an integral homology three-sphere.

Let Y be an integer homology three-sphere, and $K \subset Y$ be a knot. Let Y_0 be the manifold obtained by zero-surgery on K, and let $Y_{1/q}$ be obtained by $1/q$ surgery on K, where q is a positive integer.

We fix a representation

$$H^1(Y;\mathbb{Z}) \longrightarrow \mathbb{Z}/q\mathbb{Z}$$

taking generators to generators, and let

$$\underline{HF^+}(Y_0, \mathbb{Z}/q\mathbb{Z}) \cong \bigoplus_{\mathfrak{s}\in\mathrm{Spin}^c(Y_0)} \underline{HF^+}(Y_0, \mathfrak{s})$$

denote the corresponding homology group with twisted coefficient ring (in the sense of Subsection 4.11).

Theorem 10.14. *Let Y be an integral homology three-sphere and let q be a positive integer. Then, there is a U-equivariant exact sequence*

$$\ldots \longrightarrow \underline{HF^+}(Y_0;\mathbb{Z}/q\mathbb{Z}) \longrightarrow HF^+(Y_{1/q}) \longrightarrow HF^+(Y) \longrightarrow \ldots$$

The proof of Lemma 10.2 in the present context gives us a generalized pointed Heegaard diagram

$$(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, z)$$

with the property that:

- the Heegaard diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma})$, and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta})$ describe Y, Y_0, and $Y_{1/q}$ respectively,
- for each $i = 1, ..., g-1$, the curves β_i, γ_i, and δ_i are small isotopic translates of one another, each pairwise intersecting in a pair of canceling transverse intersection points
- the curve δ_g is isotopic to the juxtaposition of β_g with the q-fold juxtaposition of γ_g.

We can think concretely about $\underline{CF^+}(Y_0;\mathbb{Z}/q\mathbb{Z})$ as follows. Let $\zeta = e^{\frac{2\pi i}{q}}$, and fix a reference point $\tau \in \gamma_g$, which we choose to be disjoint from all the other $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\delta}$. This gives rise to a codimension-one submanifold

$$V = \gamma_1 \times ... \times \times \gamma_{g-1} \times \{\tau\} \subset \mathbb{T}_\gamma.$$

Then, $\underline{CF^+}(Y_0;\mathbb{Z}/q\mathbb{Z})$ is generated over $\mathbb{Z}$ by the basis $[\mathbf{x}, i] \otimes \zeta^j$ where of course, $\mathbf{x}$ is an intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ in the appropriate equivalence class, i is a non-negative integer, and $j \in \mathbb{Z}/q\mathbb{Z}$. The boundary map then is given by

$$\partial\left([\mathbf{x},i]\otimes\zeta^j\right) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma}\ \sum_{\{\phi\in\pi_2(\mathbf{x},\mathbf{y})\big|\mu(\phi)=1\}} \left(\#\widehat{\mathcal{M}}(\phi)\right)\cdot[\mathbf{y}, i-n_z(\phi)]\otimes\zeta^{j+\#(V\cap\partial_\gamma(\phi))}. \tag{18}$$

The quantity $V\cap\partial_\gamma(\phi)$ is the intersection number between the codimension-one submanifold $V\subset\mathbb{T}_\gamma$ with the path in $\mathbb{T}_\gamma$ obtained by restricting ϕ to the appropriate edge.

Again, we let v_g be the intersection point between δ_g and γ_g. We now have q different intersection points between δ_g and β_g, of which we choose one, labelled w_g, in the following Proposition 10.15. We will have no need for the $q-1$ other intersection points. Let $\Theta_{\beta,\gamma}$, $\Theta_{\gamma,\delta}$, and $\Theta_{\beta,\delta}$ be as before.

As in Proposition 10.3, if we let $\theta_{\beta,\delta}=[\Theta_{\beta,\delta},0]$, then $\theta_{\beta,\delta}$ is a cycle in $CF^\infty(\mathbb{T}_\beta,\mathbb{T}_\delta)$. Note that the three-manifold described by the pair $(\Sigma,\boldsymbol{\beta},\boldsymbol{\delta})$ is now a sum $L(q,1)\#\left(\#_{i=1}^{g-1}(S^1\times S^2)\right)$ (where $L(q,1)$ is a lens space).

Proposition 10.15. *For an appropriate choice $w_g\in\beta_g\cap\delta_g$ for β_g with δ_g, there are homotopy classes of triangles $\{\psi_k^\pm\}_{k=1}^\infty\in\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ satisfying the following properties (for each k):*

$$\begin{aligned}\mu(\psi_k^\pm) &= 0,\\ n_z(\psi_k^+) &= n_z(\psi_k^-),\\ n_z(\psi_k^+) &< n_z(\psi_{k+1}^+),\end{aligned}$$

Moreover, each triangle in $\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ is Spin^c *equivalent to some $\psi_k^\pm$. Also, the congruence class modulo q of the intersection number $\#(V\cap\partial_\gamma(\psi))$ is independent of the choice of $\psi\in\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$. Furthermore, there is a choice of perturbations and complex structure with the property that for each $\Psi\in\pi_2(\mathbf{x},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ (where $\mathbf{x}\in\mathbb{T}_\beta\cap\mathbb{T}_\gamma$) with $\mu(\Psi)=0$, we have that*

$$\#\mathcal{M}(\Psi)=\begin{cases}\pm1 & \text{if } \Psi\in\{\psi_k^\pm\}_{k=1}^\infty\\ 0 & \text{otherwise}\end{cases}.$$

Proof. The proof follows along the lines of Proposition 10.5. In this case, letting $\mathcal{P}$ be the generating periodic domain in the torus, we have that

$$\partial\mathcal{P}=\beta_g+q\gamma_g-\delta_g.$$

We must choose w_g so that it is the β_g-δ_g corner point for the domain containing the basepoint z. Note that $\partial\mathcal{P}$ meets the reference point $\tau\in\gamma$ with multiplicity q. This proves the q independence of the intersection number $\#(V\cap\partial_\gamma(\psi))$ of the choice of $\psi\in\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$. (See Figure 7.)

□

Our choice of basepoint z and the intersection point $\Theta_{\beta,\delta}$, from the above proposition give us a Spin^c structure $\mathfrak{t}_{\beta,\delta}\in\text{Spin}^c(L(q,1)\#\left(\#_{i=1}^{g-1}(S^1\times S^2)\right))$.

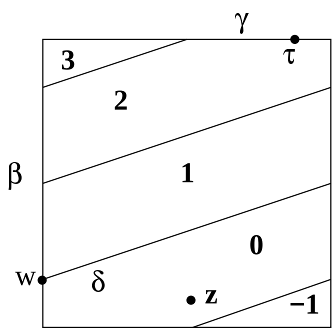

FIGURE 7. The triply-periodic domain in the torus relevant for $1/q$ surgery, with $q = 3$.

We consider the chain map

$$f_2\colon \underline{CF^+}(Y_0, \mathbb{Z}/q\mathbb{Z}) \longrightarrow CF^+(Y_{1/q})$$

defined by

$$f_2(\xi) = \sum_{\{\mathfrak{s}\in \mathrm{Spin}^c(X_{\alpha,\gamma,\delta})} \underline{f^+_{\alpha,\gamma,\delta}}(\xi\otimes\theta_{\gamma,\delta},\mathfrak{s}).$$

In the present context,

$$\underline{f^+_{\alpha,\gamma,\delta}}\Big([\mathbf{x},i]\otimes\zeta^k\otimes[\mathbf{y},j];\mathfrak{s}\Big) = \sum_{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\delta}\ \sum_{\{\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{w})\big|\#(V\cap\partial_\gamma\psi)=-k,\mathfrak{s}_z(\psi)=\mathfrak{s}\}} (\#\mathcal{M}(\psi))\cdot[\mathbf{w}, i+j-n_z(\psi)].$$

We define

$$f_3\colon CF^+(Y_{1/q}) \longrightarrow CF^+(Y)$$

by

$$f_3(\xi) = \sum_{\{\mathfrak{s}\in \mathrm{Spin}^c(X_{\alpha,\delta,\beta})\big|\mathfrak{s}|Y_{\beta,\delta}=\mathfrak{t}_{\beta,\delta}\}} f^+_{\alpha,\delta,\beta}(\xi\otimes\theta_{\beta,\delta},\mathfrak{s}).$$

This gives us maps:

$$\underline{CF^+}(Y_0,\mathbb{Z}/q\mathbb{Z}) \xrightarrow{f_2} CF^+(Y_{1/q}) \xrightarrow{f_3} CF^+(Y).$$

It follows, once again, from associativity, together with the Proposition 10.15 that the maps on homology $F_3\circ F_2 = 0$. Note that the chain homotopy evaluated on $\zeta^k\times[\mathbf{x},i]$ is constructed by counting squares in $\varphi\in\pi_2(\mathbf{x},\Theta_{\gamma,\delta},\Theta_{\delta,\beta},\mathbf{y})$ with $V\cap\partial_\gamma(\varphi) = -k$.

We homotope the δ-curve to the juxtaposition of the β_g with the q-fold juxtaposition of γ_g. This gives a short exact sequence of graded groups

$$0 \longrightarrow \underline{CF^+}(Y_0,\mathbb{Z}/q\mathbb{Z}) \xrightarrow{\iota} CF^+(Y_{1/q}) \xrightarrow{\pi} CF^+(Y) \longrightarrow 0.$$

To see the inclusion, note that each intersection point $\mathbf{x}$ of $\mathbb{T}_\alpha\cap\mathbb{T}_\gamma$ corresponds to q distinct intersection points between $\mathbb{T}_\alpha\cap\mathbb{T}_\delta$, labelled $(\mathbf{x}_1,...,\mathbf{x}_q)$. For each of these intersection points, there is a unique smallest triangle $u_1,...,u_q$, with $u_i\in\pi_2(\mathbf{x},,\Theta_{\gamma,\delta},\mathbf{x}_j)$. We claim that the q integers $\#(V \cap u_i)$ each lie in different congruence classes modulo q. This gives the inclusion. To see surjection, note that each intersection point of $\mathbf{x}' \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ gives rise to a unique intersection point $\rho(\mathbf{x}')$ between $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$, which can be joined by a small triangle in $\pi_2(\rho(\mathbf{x}'), \Theta_{\beta,\delta}, \mathbf{x}')$. (See Figure 8 for an illustration.)

With this said, then, the energy filtration is defined as before, calculating the energy of classes $\psi \in \pi_2(\mathbf{x}_0, \Theta_{\gamma,\delta}, \mathbf{y})$. Thus we obtain the required long exact sequence.

10.4. $\widehat{HF}$. Let Y be an oriented three-manifold, $K \subset Y$ be a knot, and $\mathfrak{s}_0$ be a fixed Spin^c structure over $Y - K$.

Theorem 10.16. *For each* Spin^c *structure* $\mathfrak{s}_0$ *on* $Y - K$*, we have the exact sequence:*

$$\ldots \longrightarrow \widehat{HF}(Y, [\mathfrak{s}_0]) \longrightarrow \widehat{HF}(Y_h, [\mathfrak{s}_0]) \longrightarrow \widehat{HF}(Y_{h+m}, [\mathfrak{s}_0]) \longrightarrow \ldots$$

Similarly, we have:

Theorem 10.17. *Let* Y *be an integral homology three-sphere and let* q *be a positive integer. Then, there is a* U*-equivariant exact sequence*

$$\ldots \longrightarrow \underline{\widehat{HF}}(Y_0; \mathbb{Z}/q\mathbb{Z}) \longrightarrow \widehat{HF}(Y_{1/q}) \longrightarrow \widehat{HF}(Y) \longrightarrow \ldots$$

For the proofs of these results, Proposition 10.15 (or Proposition 10.5, for the case of +1-surgeries) is replaced by the comparatively simpler:

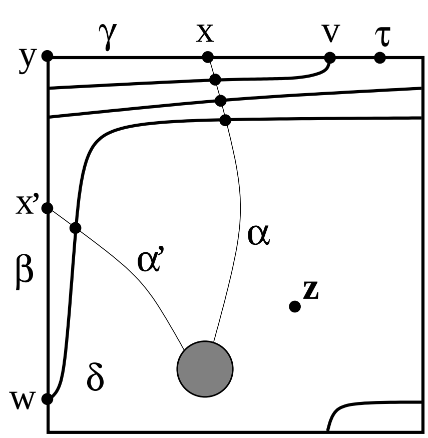

FIGURE 8. The analogue of Figure 6, only for $1/q$ surgery with $q = 3$. We have pictured here only the part of the surface taking place in the final torus summand, and correspondingly dropped the g subscripts. There are two α-curves crossing the region here, labelled α and α': the first of these meets γ at x, the second meets β at x'. Observe the three intersection points of $\alpha \cap \delta$ and the intersection point of $\alpha' \cap \delta$ corresponding to x and x' respectively.

Proposition 10.18. *There are two homotopy classes of triangles* ψ^+ *and* ψ^- *in* $\pi_2(\Theta_{\beta,\gamma},\Theta_{\gamma,\delta},\Theta_{\beta,\delta})$ *with*

$$\begin{aligned}\mu(\psi^{\pm}) &= 0,\\ n_z(\psi^{\pm}) &= 0,\\ \#(\partial_\gamma\psi^+) &= \#(\partial_\gamma\psi^-)+q.\end{aligned}$$

Indeed, these are the only two triangles with $\mathcal{D}(\psi)\geq 0$ *and* $n_z(\psi)=0$*. Also, each moduli space consists of a single, smooth isolated point.*

Proof. This now follows directly from the picture in the torus. In particular, in the present case, there is no need for Theorem 10.4. □

Proof of Theorems 10.16 and 10.17. The proofs here are now obtained by copying the earlier proofs for HF^+, with the obvious notational changes. □

10.5. **Integer surgeries.** Another generalization of Theorem 10.1 involves integer surgeries.

Let Y be an integer homology three-sphere, and $K\subset Y$ be a knot. Let Y_0 be the manifold obtained by zero-surgery on K, and let Y_p be obtained by $+p$ surgery on K, where p is a positive integer.

Theorem 10.19. *There is a surjective map* $Q\colon \mathrm{Spin}^c(Y_0)\longrightarrow \mathrm{Spin}^c(Y_p)$ *with the property that for each* Spin^c *structure* $\mathfrak{t}\in\mathrm{Spin}^c(Y_p)$*, we have a* U*-equivariant exact sequence*

$$\ldots \xrightarrow{F_1} HF^+(Y_0,[\mathfrak{t}]) \xrightarrow{F_2} HF^+(Y_p,\mathfrak{t}) \xrightarrow{F_3} HF^+(Y) \longrightarrow \ldots,$$

where

$$HF^+(Y_0,[\mathfrak{t}]) = \bigoplus_{\{\mathfrak{t}_0 \big| Q(\mathfrak{t}_0)=\mathfrak{t}\}} HF^+(Y_0,\mathfrak{t}_0).$$

Moreover, F_3 *preserves* $\mathbb{Z}/2\mathbb{Z}$ *degree, chosen so that*

$$\chi(\widehat{HF}(Y_p,\mathfrak{t}))=\chi(\widehat{HF}(Y))=1.$$

In particular, there is a U*-equivariant exact sequence*

$$\ldots \longrightarrow HF^+(Y_0) \longrightarrow HF^+(Y_p) \longrightarrow \bigoplus_{i=1}^p HF^+(Y) \longrightarrow \ldots,$$

Remark 10.20. *Indeed, a modification of the following proof can also be given to construct an exact sequence*

$$\ldots \xrightarrow{F_2} HF^+(Y) \xrightarrow{F_3} HF^+(Y_{-p},\mathfrak{t}) \xrightarrow{F_1} HF^+(Y_0,[\mathfrak{t}]) \longrightarrow \ldots,$$

where F_3 *preserves the* $\mathbb{Z}/2\mathbb{Z}$ *degree.*

Proof. This time, the curve δ_g is isotopic to the juxtaposition of the p-fold juxtaposition of β_g with the γ_g.

Now, we have p different intersection points between δ_g and γ_g. We choose one (so that the analogue of Proposition 10.15 holds, for our given choice of basepoint), and label it v_g. We will have no need for the remaining $p-1$ intersection points. Let w_g denote the intersection point between β_g and δ_g, and let $\Theta_{\beta,\gamma}$, $\Theta_{\gamma,\delta}$, and $\Theta_{\beta,\delta}$ be as before. We have a corresponding Spin^c structure $\mathfrak{t}_{\gamma,\delta}$ corresponding to $\Theta_{\gamma,\delta}$.

If $\mathfrak{t}' \in \mathrm{Spin}^c(Y_0)$, there is a unique Spin^c structure $\mathfrak{t} \in \mathrm{Spin}^c(Y_p)$ with the property that there is a Spin^c structure $\mathfrak{s}$ on $X_{\alpha,\gamma,\delta}$ with $\mathfrak{s}|Y_0 = \mathfrak{t}'$, $\mathfrak{s}|Y_{\gamma,\delta} = \mathfrak{t}_{\gamma,\delta}$, and $\mathfrak{s}|Y_{\alpha,\delta} = \mathfrak{t}$. We let $Q(\mathfrak{t}') = \mathfrak{t}$.

Fix a Spin^c structure $\mathfrak{t}$ over Y_p. We consider the chain map

$$f_2 \colon CF^+(Y_0) \longrightarrow CF^+(Y_p, \mathfrak{t})$$

defined by

$$f_2(\xi) = \sum_{\{\mathfrak{s}\in\mathrm{Spin}^c(X_{\alpha,\beta,\delta})\big|\mathfrak{s}|Y_{\alpha,\delta}=\mathfrak{t},\ \mathfrak{s}|_{\gamma,\delta}=\mathfrak{t}_{\gamma,\delta}\}} f^+_{\alpha,\gamma,\delta}(\xi\otimes\theta_{\gamma,\delta},\mathfrak{s}).$$

We define f_3 as follows. Consider

$$f_3(\xi) = \sum_{\{\mathfrak{s}\in\mathrm{Spin}^c(X_{\alpha,\delta,\beta})\big|\mathfrak{s}|Y_p=\mathfrak{t}\}} f^+_{\alpha,\delta,\beta}(\xi\otimes\theta_{\beta,\delta}).$$

This gives us maps:

$$CF^+(Y_0,[\mathfrak{t}]) \xrightarrow{f_2} CF^+(Y_p,\mathfrak{t}) \xrightarrow{f_3} CF^+(Y).$$

It follows once again from associativity, together with the analogue of Proposition 10.15, that $F_3 \circ F_2 = 0$.

We homotope the δ-curve to the juxtaposition of the p-fold multiple of β_g with γ_g. This gives a short exact sequence of graded groups

$$0 \longrightarrow CF^+(Y_0,[\mathfrak{t}]) \xrightarrow{\iota} CF^+(Y_p,\mathfrak{t}) \xrightarrow{\pi} CF^+(Y) \longrightarrow 0.$$

The inclusion follows as before: each intersection point $\mathbf{x}$ of $\mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ corresponds a unique intersection point between $\mathbb{T}_\alpha \cap \mathbb{T}_\delta$, which can be canonically connected by a small triangle. To see surjection, note that each intersection point of $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ gives rise to p different intersection points between $\mathbb{T}_\alpha\cap\mathbb{T}_\delta$, which we label $(\mathbf{y}_1, ..., \mathbf{y}_p)$. Note, however, that $\epsilon(\mathbf{y}_i,\mathbf{y}_j) = (i-j)\mathrm{PD}[\beta_g^*]$. Now, $\mathrm{PD}[\beta_g^*] \in H^2(Y_p)$ is a generator, so there will always be a unique induced intersection point representing the Spin^c structure $\mathfrak{t}$ over Y_p. The rest follows as before. □

10.6. $+1$ **surgeries for twisted coefficients.** There is also surgery exact sequence for $+1$ surgeries which uses twisted coefficients.

For simplicity, we state it in the case where we begin with a three-manifold Y which is an integer homology sphere. In that case, if we let T be a generator for $H^1(Y_0;\mathbb{Z})$, then we can think of $\mathbb{Z}[H^1(Y_0;\mathbb{Z})]$ as $\mathbb{Z}[T,T^{-1}]$. Given any $\mathbb{Z}[U]$ module M, let $M[T,T^{-1}]$ denote the induced module over $\mathbb{Z}[U,T,T^{-1}]$.

Theorem 10.21. *There is a $\mathbb{Z}[U,T,T^{-1}]$-equivariant long exact sequence:*

$$... \longrightarrow HF^+(Y)[T,T^{-1}] \xrightarrow{F_1^+} \underline{HF^+}(Y_0) \xrightarrow{F_2^+} HF^+(Y_1)[T,T^{-1}] \xrightarrow{F_3^+} ...$$

We will think of $\underline{HF^+}(Y_0)$ like we did in Subsection 10.3: we fix a reference point $\tau \in \gamma_g$, and let the boundary map record, in the power of T, the multiplicity with which ϕ meets τ along its boundary, as in Equation (18) (with the difference that now we use a formal variable T rather than a root of unity ζ).

We will similarly use a reference point $\tau' \in \delta_g$, again defining the boundary map for Y_1 which records the intersection with τ' in the power of T, to obtain a chain complex for Y_1, which we write as: $CF^+(Y_1, \mathbb{Z}[T, T^{-1}])$. Note that (by contrast with the case of Y_0) this has little effect on the homology. Indeed, it is easy to construct an isomorphism of chain complexes (over $\mathbb{Z}[U, T, T^{-1}]$):

$$CF^+(Y_1) \otimes_{\mathbb{Z}} \mathbb{Z}[T, T^{-1}] \cong CF^+(Y_1, \mathbb{Z}[T, T^{-1}]).$$

Moreover, it is clear that

$$H_*(CF^+(Y_1) \otimes_{\mathbb{Z}} \mathbb{Z}[T, T^{-1}]) \cong HF^+(Y_1) \otimes_{\mathbb{Z}} \mathbb{Z}[T, T^{-1}].$$

However, this device will be convenient in constructing the chain maps.

We choose τ' to lie on the boundary of ψ^- and τ to lie on the boundary of ψ^+ (where $\psi^\pm = \psi_1^\pm$ from Proposition 10.5), and let V, V' be the corresponding codimension one subsets of $\mathbb{T}_\gamma$ and $\mathbb{T}_\delta$ respectively. We then let

$$\underline{f_1^+}([\mathbf{x}, i]) = \sum_{\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta,\gamma} . \mathbf{w}) \big| \mu(\psi)=0\}} c(\mathbf{x}, \mathbf{w}, \psi) \cdot [\mathbf{w}, i - n_z(\psi)],$$

and

$$\underline{f_2^+}([\mathbf{x}, i]) = \sum_{\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta} \sum_{\{\psi \in \pi_2(\mathbf{x}, \Theta_{\gamma,\delta}, \mathbf{w}) \big| \mu(\psi)=0\}} c(\mathbf{x}, \mathbf{w}, \psi) \cdot [\mathbf{w}, i - n_z(\psi)];$$

where in both cases $c(\mathbf{x}, \mathbf{w}, \psi) \in \mathbb{Z}[T, T^{-1}]$ is given by

$$c(\mathbf{x}, \mathbf{w}, \psi) = (\#\mathcal{M}(\psi)) \cdot \left(T^{\#(\partial_\gamma \psi \cap V) + \#(\partial_\delta \psi \cap V')}\right).$$

We have the following analogue of Proposition 10.6:

Proposition 10.22. *The composition* $\underline{F_2^+} \circ \underline{F_1^+} = 0$.

Proof. Observe that for the homotopy classes $\{\psi_k^\pm\}_{k=1}^\infty$ from Proposition 10.5, we have that

$$\#(\partial_\beta \psi_k^+ \cap V) + \#(\partial_\delta \psi_k^+ \cap V') = \#(\partial_\beta \psi_k^- \cap V) + \#(\partial_\delta \psi_k^- \cap V') = 1$$

This implies that the formal sum

$$\begin{aligned} \sum_{\mathfrak{s}_{\beta,\gamma,\delta} \in S_{\beta,\gamma,\delta}} \underline{f_{\beta,\gamma,\delta}^{\leq 0}}(\theta_{\beta,\gamma} \otimes \theta_{\gamma,\delta}, \mathfrak{s}_{\beta,\gamma,\delta}) &= \sum_{k=1}^\infty T \otimes \left(\left[\Theta_{\beta,\delta}, -\frac{k(k-1)}{2}\right] - \left[\Theta_{\beta,\delta}, -\frac{k(k-1)}{2}\right]\right) \\ &= 0. \end{aligned}$$

Thus, the proof follows from associativity as before. □

Proof of Theorem 10.21. With Proposition 10.22 replacing Proposition 10.6, the proof proceeds as the proof of Theorem 10.1. □

We have also the generalization for integer surgeries:

Theorem 10.23. *Let Y be an integral homology three-sphere, let $K \subset Y$ be a knot in Y, and fix a positive integer p. For each* Spin^c *structure* $\mathfrak{t} \in \mathrm{Spin}^c(Y_p)$*, we have a* $\mathbb{Z}[U, T, T^{-1}]$*-equivariant exact sequence*

$$\dots \xrightarrow{F_1} \underline{HF^+}(Y_0, [\mathfrak{t}]) \xrightarrow{F_2} HF^+(Y_p, \mathfrak{t})[T, T^{-1}] \xrightarrow{F_3} HF^+(Y)[T, T^{-1}] \longrightarrow \dots,$$

where

$$\underline{HF^+}(Y_0, [\mathfrak{t}]) = \bigoplus_{\{\mathfrak{t}_0 \big| Q(\mathfrak{t}_0) = \mathfrak{t}\}} \underline{HF^+}(Y_0, \mathfrak{t}_0),$$

using the map $Q\colon \mathrm{Spin}^c(Y_0) \longrightarrow \mathrm{Spin}^c(Y_p)$ *be the map from Theorem 10.19.*

Proof. Combine the refinements from Theorem 10.19 with those of Theorem 10.21. □

11. Calculation of HF^∞

The main result of the present section is the complete calculation of $\underline{HF^\infty}(Y)$ purely in terms of the homological data of Y (we use the algebraic notation from Subsection 7.3). We also give the following similar calculation of $HF^\infty(Y)$ when $b_1(Y) \leq 2$.

Theorem 11.1. *Let Y be a closed, oriented three-manifold with $b_1(Y) \leq 2$. Then, there is an equivalence class of orientation system over Y with the following property. If $\mathfrak{s}_0$ is torsion, then*

$$HF^\infty(Y;\mathfrak{s}_0) \cong \mathbb{Z}[U,U^{-1}] \otimes_{\mathbb{Z}} \Lambda^* H^1(Y;\mathbb{Z})$$

as a $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(H_1(Y;\mathbb{Z})/\mathrm{Tors})$-module. Furthermore, if $\mathfrak{s}$ is not torsion,

$$HF^\infty(Y;\mathfrak{s}) \cong (\mathbb{Z}[U]/U^n - 1) \otimes_{\mathbb{Z}} \Lambda^* c_1(\mathfrak{s})^\perp,$$

where $c_1(\mathfrak{s})^\perp \subset H^1(Y;\mathbb{Z})$ is the subgroup pairing trivially with $c_1(\mathfrak{s})$, and $n = \mathfrak{d}(\mathfrak{s})/2$.

Remark 11.2. *Of course, in the above statement, we think of the usual cohomology $H^1(Y;\mathbb{Z})$ (with constant coefficients); but it will be apparent from the proof that for each choice of locally constant $\mathbb{Z}$ coefficient system, we obtain an orientation system for HF^∞ for which the analogous isomorphism holds: this gives an identification between locally constant $\mathbb{Z}$ coefficient systems over Y and equivalence classes of orientation system over Y.*

The proof in some important special cases is given in Subsection 11.1, and the general case is proved in Subsection 11.2. We give also a twisted analogue in Subsection 11.3 which holds for arbitrary $b_1(Y)$.

The theorem describes the module structure of $HF^+(Y,\mathfrak{s}_0)$ in sufficiently large degree, when $\mathfrak{s}_0$ is a torsion Spin^c structure and $b_1(Y) \leq 2$ (hence completing Part (4) of Theorem 1.1). It also allows us to pay off several other debts: first, it allows us to define an absolute $\mathbb{Z}/2\mathbb{Z}$ grading on the homology groups; then, combined with the discussion of Section 9, it allows us to relate $\chi(HF^-(Y,\mathfrak{s}))$ with Turaev's torsion in Subsection 11.6 (though an alternative calculation could also be given by modifying directly the discussion in Section 9). It also allows us to extend the Euler characteristic calculations for HF^+ to the case where the Spin^c structure is torsion, c.f. Subsection 11.7. Finally, the result allows us to identify a "standard" orientation system for Y: the one for which Theorem 11.1 holds, with the usual $H^1(Y;\mathbb{Z})$ on the right-hand-side. (This justifies our practice of dropping the coefficient system from the notation for HF^∞, and the other related groups.) Since the analogue of Theorem 11.1 in the twisted case (Theorem 11.12) holds without restriction on the Betti numbers of Y, it can be used to identify a canonical coherent system of orientations for any oriented three-manifold Y.

11.1. $HF^\infty(Y)$ when $H_1(Y;\mathbb{Z}) = 0$ or $\mathbb{Z}$.

Theorem 11.3. *Theorem 11.1 holds when Y is an integer homology three-sphere; i.e. over $\mathbb{Z}$, $HF^\infty(Y)$ is freely generated by generators y_i for $i \in \mathbb{Z}$, with $Uy_i = y_{i-1}$.*

Theorem 11.4. *Theorem 11.1 holds when the three-manifold in question Y_0 satisfies $H_1(Y_0) \cong \mathbb{Z}$. More concretely, let $H \in H^2(Y_0;\mathbb{Z})$ be a generator, and let $\mathfrak{s}_0$ denote the Spin^c structure with trivial first Chern class. Then if $\mathfrak{s} = \mathfrak{s}_0 \pm n \cdot H$ with $n > 0$, then $HF^\infty(Y_0,\mathfrak{s})$ is freely generated by generators x_i for $i \in 1,...,n$, with $Ux_i = x_{i-1}$, $Ux_1 = x_n$. Moreover, $HF^\infty(Y_0,\mathfrak{s}_0)$ is* *freely generated by generators* x_i, y_i *for* $i \in \mathbb{Z}$, *with* $Uy_i = y_{i-1}$, $Ux_i = x_{i-1}$ *and* $\mathrm{gr}(x_i, y_i) = 1$*; also,* $\mathrm{PD}[H] \cdot x_i = y_i$.

The main ingredient in the proof of the above results is the following:

Proposition 11.5. *Let* Y *be an integer homology three-sphere, and* $K \subset Y$ *be a knot, then there is an identification:*

$$HF^\infty(Y_0, \mathfrak{s}) \cong HF^\infty(Y, \mathfrak{s}_0)/(U^n - 1),$$

where Y_0 *is the three-manifold obtained by zero-surgery on* K*, and where the divisibility of* $c_1(\mathfrak{s})$ *is* $2n$.

This is proved in several steps.

We start with a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ describing Y_0, with the property that $(\Sigma, \{\alpha_2, ..., \alpha_g\}, \{\beta_1, ..., \beta_g\})$ describes the knot complement. Let γ be a curve which intersects α_1 once and is disjoint from $\{\alpha_2, ..., \alpha_g\}$, so that $(\Sigma, \{\gamma, \alpha_2, ..., \alpha_g\}, \{\beta_1, ..., \beta_g\})$ represents Y. Indeed, we let $\gamma_2, ..., \gamma_g$ be small isotopic translates of $\alpha_2, ...\alpha_g$, with $\gamma_i \cap \alpha_i$ for $i = 2, ..., g$ consisting of a canceling pair of points $w_i^\pm$. Such a diagram can always be found (compare Lemma 10.2). We twist α_1 along γ, and let $R^\infty(\mathfrak{s})$ resp. $L^\infty(\mathfrak{s})$ denote the subset of $CF^\infty(Y_0, \mathfrak{s})$, generated by the γ-induced intersection points to the right resp. the left of the curve γ. Recall that if we twist sufficiently, then $L^\infty(\mathfrak{s})$ is a subcomplex (c.f. Proposition 9.5).

We relate HF^∞ for Y with $H_*(R^\infty)$, as follows:

Lemma 11.6. *There is an isomorphism* $H_*(R^\infty) \cong HF^\infty(Y)$.

Proof. Let $\Theta_{\alpha,\gamma} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ be the intersection point $\{\gamma \cap \alpha_1, w_2^+, ..., w_g^+\}$. It follows as in the proof of Proposition 9.5 that there are no triangles $\psi \in \pi_2(\Theta_{\alpha,\gamma}, \mathbf{x}, \mathbf{y})$ with $\mathbf{x} \in L^\infty$, $\mathbf{y} \in \mathbb{T}_\gamma \cap \mathbb{T}_\beta$ and $\mathcal{D}(\psi) \geq 0$, and $\mu(\psi) = 0$. Hence, counting holomorphic triangles whose $\mathbb{T}_\alpha \cap \mathbb{T}_\gamma$-vertex is $\Theta_{\alpha,\gamma}$, we obtain a map $H_*(R^\infty) \longrightarrow HF^\infty(Y)$. On the chain level, this map has the form ι+lower order, where $\iota[\mathbf{x}, i] = [\mathbf{x}', i - n_z(\psi_{\mathbf{x}})]$ where $\mathbf{x}'$ is the intersection point on $\mathbb{T}_\gamma \cap \mathbb{T}_\beta$ closest to $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\psi_{\mathbf{x}}$ is the unique small triangle (supported in the neighborhood of γ and the support of the isotopies between γ_i and α_i with non-negative multiplicities) and lower order is taken with respect to the energy filtration on Y. Moreover, there is a relative $\mathbb{Z}$-grading on both complexes, given by the Maslov index (where we take an "in" domain for Y_0). The map preserves this grading. Moreover, there are only finitely many generators in each degree. It follows then that the induced map is an isomorphism. □

We have seen that the map $H_*(R^\infty) \longrightarrow H_*(L^\infty)$ naturally splits into two pieces, δ_1 and δ_2, where δ_1 uses the domains ϕ^{in} from Lemma 9.4.

Lemma 11.7. *The map* δ_1 *is an isomorphism.*

Proof. This follows from the fact that on the chain level, δ_1 has the form

$$\delta_1[\mathbf{x}^+, i] = [\mathbf{x}^-, i - n_z(\phi_{\mathbf{x}^+,\mathbf{x}^-})] + \text{lower order}.$$

(Lemma 9.6), together with the fact that δ_1 preserves the relative $\mathbb{Z}$ grading. □

Lemma 11.8. *The map δ_2 is an isomorphism.*

Proof. Fix an equivalence class of intersection points between $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, all of which are γ-induced. According to Section 9, if we wind sufficiently many times along γ and move the basepoint z sufficiently close to γ, then $\langle c_1(\mathfrak{s}), H\rangle$ can be made arbitrarily large. By moving the basepoint to change the Spin^c structure, we have that the complexes $L^+(\mathfrak{s})$ and $L^+(\mathfrak{s}')$ (resp. $R^+(\mathfrak{s})$ and $R^+(\mathfrak{s}')$) are identical. Moreover, if $\mathfrak{s}$ and $\mathfrak{s}'$ are sufficiently positive, then the map δ_2^+ is independent of the Spin^c structure.

Choose a degree i sufficiently large that $H_i(R^+) \cong H_i(R^\infty)$ and $H_i(L^+) \cong H_i(L^\infty)$, and note under these identifications, the map induced on homology

$$\delta_2^+ \colon H_i(R^+) \longrightarrow H_{i-1}(L^+)$$

agrees with δ_2. For fixed i and sufficiently large $\mathfrak{s}$, δ_1^+ on $H_i(R^+(\mathfrak{s}))$ vanishes. Since $HF^+(Y,\mathfrak{s})$ is zero for all sufficiently large $\mathfrak{s}$, it follows from the long exact sequence induced from

$$0 \longrightarrow L^+(\mathfrak{s}) \longrightarrow CF^+(Y,\mathfrak{s}) \longrightarrow R^+(\mathfrak{s}) \longrightarrow 0$$

that $\delta = \delta_1^+ + \delta_2^+ \colon H_*(R^+(\mathfrak{s})) \longrightarrow H_*(L^+(\mathfrak{s}))$ is an isomorphism. It follows that the kernel of δ_2^+ in degree i is trivial. From this, it follows in turn that the kernel of δ_2^+ is trivial in all larger degrees. Since δ_1^+ decreases degree more than δ_2^+, it is easy to see that the cokerenel of δ_2^+ in dimension i is trivial, as well. The lemma then follows. □

Proof of Proposition 11.5. Note that δ_1 and and δ_2 are both isomorphisms, and

$$\mathrm{gr}(\delta_1([\mathbf{x},i]), \delta_2([\mathbf{x},i])) = \pm 2n$$

for each generator $[\mathbf{x},i]$ for $CF^+(Y,\mathfrak{s})$. It follows that:

$$HF^\infty(Y_0,\mathfrak{s}) \cong H_*(R^\infty)/(U^n - 1).$$

Thus, the proposition follows from Lemma 11.6. □

Proof of Theorem 11.3. Since multiplication by U is an isomorphism on $HF^\infty(Y,\mathfrak{s}_0)$, Proposition 11.5 shows that $HF^\infty(Y) \cong HF^\infty(Y_1)$, where Y_1 denotes the $+1$ surgery on any knot $K \subset Y$. Since any two integer homology three-spheres can be connected by sequences of ± 1 surgeries, it follows that $HF^\infty(Y) \cong HF^\infty(S^3)$, which we know has the claimed form. □

Proof of Theorem 11.4. This is a direct consequence of Theorem 11.3 and Proposition 11.5 when $c_1(\mathfrak{s})$ is non-torsion. In the torsion case, the induced maps on homology satisfy either $\delta_1 = \delta_2$, or $\delta_1 = -\delta_2$, according to the two possible orientation conventions for Y. The two possibilities give two different homology groups (over $\mathbb{Z}$). We define the standard orientation convention to be the one for which $\delta_1 = -\delta_2$.

Finally, note that the action of $h \in H_1(Y_0;\mathbb{Z})$ is given by $\pm\delta_1$, as can be easily seen from the geometric representative for the circle action (see Remark 4.12). □

11.2. **The general case of Theorem 11.1.**

Definition 11.9. *Let Z be a compact three-manifold with $\partial Z = T^2$. The kernel of the map*

$$H_1(\partial Z) \longrightarrow H_1(Z)$$

is cyclic, generated by $d\ell$, where $\ell \subset T^2$ is a simple, closed curve. We call such a curve ℓ a longitude*, and d the* divisibility *of Z.*

Proposition 11.10. *Suppose that $b_1(Z) = 1$, and let h_1, h_2 be primitive homology classes in $H_1(T^2;\mathbb{Z})$ and with $h_1 \cdot \ell$ and $h_2 \cdot \ell$ positive with $h_1 \cdot h_2 = 1$. Then, if HF^∞ of Y_{h_1} and Y_{h_2} satisfy the property Theorem 11.1, then so does $Y_{h_1+h_2}$.*

Proof. Recall that the Floer homologies of a rational homology three-sphere have an absolute $\mathbb{Z}/2\mathbb{Z}$ grading, specified by

$$\chi(\widehat{HF}(Y)) = |H_1(Y;\mathbb{Z})|.$$

From the exact sequence of Theorem 10.12, we have that

$$.. \longrightarrow HF^+(Y_{h_1}) \xrightarrow{F_1} HF^+(Y_{h_2}) \xrightarrow{F_2} HF^+(Y_{h_1+h_2}) \longrightarrow ...$$

The hypothesis in the sign guarantees that the degree shift occurs at F_1 (using the absolute $\mathbb{Z}/2\mathbb{Z}$ grading on each group). It follows that $HF^\infty(Y_{h_1+h_2})$ vanishes in all odd degrees. Indeed, since this is true when we take coefficients in $\mathbb{Z}/p\mathbb{Z}$ for all p; hence, $HF^\infty(Y_{h_1+h_2})$ has no torsion in even degrees. Since $\chi(HF^\infty(Y,\mathfrak{s})/U) = 1$ for all rational homology three-spheres, the result follows. □

Proposition 11.11. *Suppose that Z be an oriented three-manifold with torus boundary. For each h with the property that $h \cdot \ell = 1$, we have an identification*

$$HF^\infty(Y_\ell, \mathfrak{s}) \cong HF^\infty(Y_h, \mathfrak{s}')/(U^n - 1)$$

where $\mathfrak{s}'$ is a torsion Spin^c *structure, $\mathfrak{s}_0|_Z = \mathfrak{s}|_Z$, and $\mathfrak{d}(\mathfrak{s}) = 2n$.*

Proof. We adapt the proof of Proposition 11.5. We start with $(\Sigma, \{\alpha_2, ..., \alpha_g\}, \{\beta_1, ..., \beta_g\})$ representing the knot complement Z, and then choose α_1 to represent ℓ and γ to represent h: i.e. $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ represents Y_ℓ and $(\Sigma, \{\gamma, \alpha_2, ..., \alpha_g\}, \{\beta_1, ..., \beta_g\})$ represents Y_h. There is an added feature now, since the divisibility d of Z could be greater than one. It is still the case that for sufficiently large winding, all the intersection points are represented from $R^\infty(\mathfrak{s})$ or $L^\infty(\mathfrak{s})$, and, as in Lemma 9.4, all homotopy classes of maps ϕ with $\mu(\phi) = 1$ admitting holomorphic representatives (connecting any two intersection points) satisfy that the property that $\partial_\alpha \phi$ uses the central point $p = \alpha_1 \cap \gamma$ either once or zero times. Recall δ_1 is the map defined using those homotopy classes which meet p once. Now, there is a difference map

$$\eta \colon (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \longrightarrow \mathbb{Z}/d\mathbb{Z},$$

which is defined by

$$\eta(\mathbf{x}, \mathbf{y}) = \#\left(\partial_{\alpha_1}\phi \cap p\right) \pmod{d}.$$

There are corresponding splittings

$$L^\infty(\mathfrak{s}) = L^\infty_1, ..., L^\infty_d \quad \text{and} \quad R^\infty(\mathfrak{s}) = R^\infty_1, ..., R^\infty_d.$$

labeled so that $\eta(\mathbf{x},\mathbf{y}) = 1$ if $\mathbf{x} \in R_i^\infty$ and $\mathbf{y} \in R_{i+1}^\infty$, and $\delta_1(R_i^\infty) \subset L_{i+1}^\infty$. and $\delta_2(R_i^\infty) \subset L_i^\infty$.

The proof of Lemma 11.6 gives us that $H_*(R_i^\infty) \cong HF^\infty(Y, \mathfrak{s}')$ (for $i = 1, ..., d$). Also, analogues of Lemmas 11.7 and 11.8 still hold: both δ_1 and δ_2 are isomorphisms. Now, the proposition easily follows as before. □

Proof of Theorem 11.1 We begin with the case where $b_1(Y) = 0$, and prove the claim by induction on $|H_1(Y;\mathbb{Z})|$. The base case is, of course, Theorem 11.3. For the inductive step, we choose a knot $K \subset Y$ which represents a non-trivial homology class. With appropriate orientation, we have that $m \cdot \ell > 0$. If $m \cdot \ell > 1$, the inductive step follows from Proposition 11.10, since m can be decomposed as $m = h_1 + h_2$ with $h_1 \cdot h_2 = 1$, $h_1 \cdot \ell, h_2 \cdot \ell > 1$. Note also that if $h \cdot \ell > 0$, then $|H_1(Y_h)|$ depends linearly on $h \cdot \ell$.

If $m \cdot \ell = 1$, then since K is homologically non-trivial, we must have that $d > 1$. Also, $|\mathrm{Tors} H_1(Y_\ell)| = \frac{1}{d}|\mathrm{Tors} H_1(Y)|$. Applying Proposition 11.11 along a different knot in Y_ℓ which represents a generator for $H_1(Y_\ell)/\mathrm{Tors}$, we see that

$$HF^\infty(Y_\ell, \mathfrak{s}) \cong HF^\infty(Y', \mathfrak{s}')/(U^n - 1),$$

where $|H_1(Y';\mathbb{Z})| < |H_1(Y;\mathbb{Z})|$. Applying the proposition again, and the induction hypothesis, we obtain that $HF^\infty(Y) \cong \mathbb{Z}[U, U^{-1}]$.

The proof for general $b_1(Y) = 1$ or 2 follows from an induction on $b_1(Y)$. Let Y be an oriented three-manifold with $b_1(Y) = 1$ or 2. Choose a knot $K \subset Y$ whose image in $H_1(Y;\mathbb{Z})/\mathrm{Tors}$ is primitive. (This implies that in $Y - K$, the divisibility $d = 1$.) If $\mathfrak{s}$ is a non-torsion Spin^c structure on Y_ℓ, then the result follows from Proposition 11.11. The other case follows from the fact that we have two maps δ_1 and δ_2 from $R^\infty(\mathfrak{s})$ to $L^\infty(\mathfrak{s})$, and both of these maps are isomorphisms of $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(H_1(Y_h;\mathbb{Z})/\mathrm{Tors})$-modules (between two modules are, in turn, isomorphic to $\mathbb{Z}[U^{-1}] \otimes_{\mathbb{Z}} \Lambda^* H^1(Y_h;\mathbb{Z})$). Now, observe that the automorphism of $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(H_1(Y_h;\mathbb{Z})/\mathrm{Tors})$-module $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(H_1(Y_h;\mathbb{Z})/\mathrm{Tors})$ is determined by its action on the determinant line $\Lambda^b(H_1(Y_h;\mathbb{Z})/\mathrm{Tors}) \cong \mathbb{Z}$, where it is either multiplication by $+1$ or -1. Thus, the maps δ_1 and δ_2 either cancel (for one orientation convention) or they do not (for the other one). The convention where $\delta_1 + \delta_2 = 0$ is the one for which the theorem follows; it is, in this case, the standard orientation convention for Y. □

11.3. **The twisted case.** We state a version of Theorem 11.1 which holds for arbitrary first Betti number.

Observe that the proof of Theorem 11.1 breaks down when $b_1(Y) \geq 3$, since now the module $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(\mathbb{Z}^{b-1})$ has non-trivial automorphisms, so that δ_1 and δ_2 do not necessarily cancel. Indeed, it is proved in [26] that

$$HF^\infty(T^3) \cong \mathbb{Z}[U, U^{-1}] \otimes_{\mathbb{Z}} \Big(H^1(T^3) \oplus H^2(T^3)\Big).$$

There is, however, a version which holds for twisted coefficient systems.

Observe first that the twisted homology group $\underline{HF^\infty}(Y, \mathfrak{s})$ is a module over the group-ring $\mathbb{Z}[H^1(Y;\mathbb{Z})] \otimes_{\mathbb{Z}} \mathbb{Z}[U, U^{-1}]$ (which can be thought of as a ring of Laurent polynomials in $b_1(Y) + 1$ variables). To make it the ring structure respect the relative grading, we give $\underline{HF^\infty}(Y, \mathfrak{s})$ a relative $\mathbb{Z}/2\mathbb{Z}$ grading.

Theorem 11.12. *Let Y be a closed, oriented three-manifold. Then, there is a unique equivalence class of orientation system for which we have a $\mathbb{Z}[U,U^{-1}]\otimes_{\mathbb{Z}}\mathbb{Z}[H^1(Y;\mathbb{Z})]$-module isomorphism*

$$\underline{HF^\infty}(Y,\mathfrak{s})\cong\mathbb{Z}[U,U^{-1}],$$

where the latter is endowed with a trivial action by $H^1(Y;\mathbb{Z})$.

Proof. The proof is obtained by modifying the above proof of Theorem 11.1, with minor modifications, which we outline presently.

For the case where $H_1(Y_0;\mathbb{Z})\cong\mathbb{Z}$, we adapt the proof of Theorem 11.4, thinking of $\mathbb{Z}[H^1(Y;\mathbb{Z})]$ as $\mathbb{Z}[T,T^{-1}]$. In this case, Lemma 11.6 is replaced by an isomorphism $H_*(R^\infty)\cong HF^\infty(Y)[T,T^{-1}]$ (with the same proof). Next, we observe that rather than having δ_1 and δ_2 cancel, as in the proof of Theorem 11.4, we have that $\delta_1=\pm\delta_2\cdot T$. In fact, for some choice of orientation convention, we can arrange for $\delta_1=-\delta_2$. The result then follows easily from the long exact sequence connecting $\underline{L}^\infty(Y,\mathfrak{s})$, $\underline{HF^\infty}(Y,\mathfrak{s})$, and $\underline{R}^\infty(\mathfrak{s})$ observing that the map

$$\mathbb{Z}[T,T^{-1}]\xrightarrow{1-T}\mathbb{Z}[T,T^{-1}]$$

injective, with cokernel $\mathbb{Z}$ (with trivial action by T).

The same modifications work to prove the general case (arbitrary $b_1(Y)$) as well.

We now turn to the uniqueness assertion on the orientation system. For the various equivalence classes of orientation systems, it is always true that $\underline{HF^\infty}(Y,\mathfrak{s})\cong\mathbb{Z}[U,U^{-1}]$ as a $\mathbb{Z}$ module. In fact, we saw (c.f. Equation (3)) that as a $\mathbb{Z}$ module, the isomorphism class of the chain complex $\underline{CF^\infty}(Y,\mathfrak{s})$ is independent of the choice of orientation system. Moreover, from Equation (3), it is clear that the $2^{b_1(Y)}$ different equivalence classes of coherent orientation system give rise to all $2^{b_1(Y)}$ different $\mathbb{Z}[H^1(Y;\mathbb{Z})]$-module structures on $\mathbb{Z}[U,U^{-1}]$ which correspond naturally to $\mathrm{Hom}(H^1(Y;\mathbb{Z}),\mathbb{Z}/2\mathbb{Z})$, with a distinguished module for which the action by $H^1(Y;\mathbb{Z})$ is trivial. ☐

11.4. **On the structure of HF^+.** We assemble now the pieces of Theorem 1.1 claimed in the introduction:

Proof of Theorem 1.1. The topological invariance of HF^+ was established in Theorem 4.13. The fact each $\xi\in HF^+(Y,\mathfrak{s})$ is annihilated by a sufficiently large power of U follows from the corresponding fact on the chain level, where it is obvious. The fact that the groups are finitely generated for non-torsion Spin^c structures was established in Theorem 9.10 (in fact, in Proposition 9.8). The module structures in sufficiently high degrees follows from Theorem 11.1, together with the fact that in high degrees, $HF^+(Y,\mathfrak{s})$ is identified with $HF^\infty(Y,\mathfrak{s})$. ☐

11.5. **Absolute $\mathbb{Z}/2\mathbb{Z}$ gradings.** With the help of Theorem 11.12, we can define an absolute $\mathbb{Z}/2\mathbb{Z}$ grading on $CF^\infty(Y,\mathfrak{s})$ (and hence all the other associated chain complexes), for all Spin^c structures, simultaneously.

We declare the non-zero generators of $\underline{HF^\infty}(Y,\mathfrak{s})$ to have even degree. Note that for a rational homology three-sphere, this orientation convention agrees with that used before, i.e. $\chi(\widehat{HF}(Y))=|H_1(Y;\mathbb{Z})|$. (In fact, if we orient $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ so that the intersection number

$\#(\mathbb{T}_\alpha \cap \mathbb{T}_\beta) = |H_1(Y;\mathbb{Z})|$, then the $\mathbb{Z}/2\mathbb{Z}$ grading at a generator $[\mathbf{x}, i]$ is $+1$ if and only if the local intersection number of $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ at $\mathbf{x}$ is $+1$.)

With this orientation convention, we have the following refinement of Corollary 1.4:

Proposition 11.13. *Let Y_0 be an oriented three-manifold with $b_1(Y_0) = 1$, and $\mathfrak{s}$ be a non-torsion* Spin^c *structure, then*

$$\chi(HF^+(Y_0, \mathfrak{s}_0 + nH)) = -\tau_t(Y_0, \mathfrak{s}),$$

where t is the component containing $c_1(\mathfrak{s})$, and the sign on $\tau_t(Y_0, \mathfrak{s})$ is specified by

$$\tau_{-t}(\mathfrak{s}) - \tau_t(\mathfrak{s}) = n.$$

In particular, if Y_0 is obtained by zero-surgery on a knot K in a homology three-sphere, whose symmetrized Alexander polynomial is

$$\Delta_K = a_0 + \sum_{i=1}^{d} a_i (T^i + T^{-i}),$$

then

$$\chi(HF^+(Y_0, \mathfrak{s}_0 + nH) = -\sum_{j=1}^{d} j a_{|n|+j}.$$

Proof. First observe that the sign comparing $\chi(HF^+(Y_0))$ and τ_t in Theorem 9.1 is universal, depending on the relative sign between $\Delta_{i,j}$ and $\Delta'_{i,j}$. Checking these signs for $S^1 \times S^2$, the Proposition follows. □

11.6. **The Euler characteristic of HF^-.** The following is an immediate consequence of Theorem 9.1, together with Theorem 11.4 (though a more direct proof can be given by modifying the discussion in Section 9):

Corollary 11.14. *Let Y be an oriented three-manifold with $b_1(Y) = 1$, and $\mathfrak{s} \in$* Spin$^c(Y)$ *be a non-torsion* Spin^c *structure. Then, $\chi(HF^-(Y, \mathfrak{s})) = \tau_{-t}(\mathfrak{s})$, where t is the component of $H^2(Y;\mathbb{Z}) - 0$ containing $c_1(\mathfrak{s})$*

Proof. The short exact sequence

$$0 \longrightarrow CF^-(Y, \mathfrak{s}) \longrightarrow CF^\infty(Y, \mathfrak{s}) \longrightarrow CF^+(Y, \mathfrak{s}) \longrightarrow 0.$$

induced a long exact sequence in homology

$$\longrightarrow HF^-(Y, \mathfrak{s}) \longrightarrow HF^\infty(Y, \mathfrak{s}) \longrightarrow HF^+(Y, \mathfrak{s}) \longrightarrow ...,$$

which shows that

$$\chi(HF^\infty(Y, \mathfrak{s})) = \chi(HF^+(Y, \mathfrak{s})) + \chi(HF^-(Y, \mathfrak{s})).$$

Moreover, Theorem 11.1 implies that

$$\chi(HF^\infty(Y, \mathfrak{s})) = n,$$

where $2n$ is the divisibility of $c_1(\mathfrak{s})$ in $H^2(Y,\mathfrak{s})/\mathrm{Tors}$. The result now follows from the "wall-crossing formula":

$$\tau_{-t}(Y,\mathfrak{s}) - \tau_t(Y,\mathfrak{s}) = n$$

for Turaev's torsion (see [32]). □

Corollary 11.15. *If Y is an oriented three-manifold with $b_1(Y)=1$ or 2 and $\mathfrak{s}\in\mathrm{Spin}^c(Y)$ is a non-torsion Spin^c structure, then $\chi\left(HF^-(Y,\mathfrak{s})\right)=\pm\tau(\mathfrak{s})$.*

Proof. This follows in the same manner as the previous corollary, except that now $c_1(\mathfrak{s})^\perp$ is a non-trivial vector space, so its exterior algebra has Euler characteristic zero: thus, $\chi(HF^\infty(Y,\mathfrak{s}))=0$. □

11.7. **The truncated Euler characteristic.** In Theorem 9.1, we worked with a non-torsion Spin^c structure. The reason for this, of course, is given Theorem 11.1: if $\mathfrak{s}_0$ is torsion and Y_0 is a three manifold with $0<b_1(Y)=b\leq 2$, then in all sufficiently large degrees i, $HF^+_i(Y_0,\mathfrak{s}_0)\cong HF^\infty_i(Y_0,\mathfrak{s}_0)\cong\mathbb{Z}^{2^{b_1(Y)-1}}$. This shows, however, that for all sufficiently large n, the Euler characteristic of the graded Abelian group $HF^+_{\leq n}(Y_0,\mathfrak{s}_0)$ takes on two possible values, depending on the parity of n (and the difference between the two values is $2^{b_1(Y)-1}$). In fact, we have the following:

Theorem 11.16. *Let Y be a three-manifold with $b_1(Y)=1$ or 2, equipped with a torsion Spin^c structure $\mathfrak{s}_0$. Then, when $b_1(Y)=1$, then for all sufficiently large n*

$$\chi(HF^+_{\leq n}(Y,\mathfrak{s}_0)) = \begin{cases} -\tau(Y) & \text{for odd } n \\ -\tau(Y)+1 & \text{for even } n \end{cases}$$

When $b_1(Y)=2$, then in all sufficiently large degrees,

$$\chi(HF^+_{\leq n}(Y,\mathfrak{s}_0)) = \pm\tau(Y)+(-1)^n.$$

Proof. As before, we have a short exact sequence

$$0 \longrightarrow L^+ \longrightarrow CF^+(Y_0,\mathfrak{s}_0) \longrightarrow R^+ \longrightarrow 0,$$

and hence a long exact sequence:

$$... \longrightarrow H_i(L^+) \longrightarrow HF^+_i(Y,\mathfrak{s}_0) \longrightarrow H_i(R^+) \xrightarrow{\delta} ...$$

Note that we are using a relative $\mathbb{Z}$ grading here, which we can do since $\mathfrak{s}_0$ is torsion. When i is sufficiently large, the coboundary map δ is zero, since on HF^∞, the map $H_*(L^\infty)\longrightarrow HF^\infty(Y)$ is an injection.

It follows that for all sufficiently large n,

(19) $$\chi(HF^+_{\leq n}(Y)) = \chi(H_{\leq n}(L^+)) + \chi(H_{\leq n}(R^+)).$$

On the other hand, we still have a short exact sequence:

$$0 \longrightarrow \ker f_1 \longrightarrow R^+ \xrightarrow{f_1} L^+ \longrightarrow 0,$$

inducing

$$\longrightarrow H_i(\ker f_1) \longrightarrow H_i(R^+) \xrightarrow{f_1} H_{i-1}(L^+) \longrightarrow \ldots$$

Note that with the earlier grading conventions, f_1 must decrease the grading by one. Of course, $\ker f_1$ is a finite-dimensional graded vector space, so the above gives the following relation for all sufficiently large n:

$$\chi(\ker f_1) = \chi(H_{\leq n}(R^+)) + \chi(H_{\leq n-1}(L^+)). \tag{20}$$

But from Proposition 9.9 applies in the present case, to identify $\chi(\ker f_1) = \tau(\mathfrak{s}_0)$. Note that the proof of the that proposition does not really require that $\mathfrak{s}$ be negative; it suffices to consider the case where $\mathfrak{s} + \alpha_1^*$, $\mathfrak{s} + \beta_j^*$ and $\mathfrak{s} + \alpha_1^* + \beta_j^*$ are negative, and $c_1(\mathfrak{s})$ is torsion. Combining this result, Equation (19), and Equation (20), we obtain that:

$$\chi(HF^+_{\leq n}(Y, \mathfrak{s}_0)) = -\tau(Y, \mathfrak{s}_0) + (-1)^n \mathrm{rk} H_n(L^+, \mathfrak{s}_0).$$

Suppose that $b_1(Y) = 1$. Then, (according to Theorem 11.1) for all sufficiently large n, $\mathrm{rk} H_n(L^+, \mathfrak{s}_0) = 1$ if n is even and 0 when n is odd. Similarly, when $b_1(Y) = 2$, we have

$$\mathrm{rk} H_n(L^+, \mathfrak{s}_0) = \mathrm{rk} HF^\infty_n(Y)/2 = 1.$$

□

11.8. **On the role of n_z.** The "triviality" of $HF^\infty(Y)$ – its dependence on the homological information of Y alone – underscores the importance of the quantity n_z in the construction of interesting Floer-homological invariants.

Another manifestation of this is the following. When Y is an integral homology three-sphere, we needed the base-point to define $\mathbb{Z}$-grading between intersection points. However, there is still a $\mathbb{Z}/2\mathbb{Z}$ graded-theory $CF'(Y)$, which is freely generated by the transverse intersection points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $\mathbb{Z}/2\mathbb{Z}$-graded by the local intersection number between $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$. The map

$$\partial \mathbf{x} = \sum_{\mathbf{y}} \sum_{\{\phi \in \pi_2(\mathbf{x},\mathbf{y}) | \mu(\phi)=1\}} \left(\#\widehat{\mathcal{M}}(\phi)\right) \mathbf{y}$$

gives a well-defined boundary map, and in fact, we can consider the homology group

$$HF'(Y) = H_*(CF'(Y), \partial).$$

However, it is a consequence of Theorem 11.3 that

$$HF'_*(Y) \cong \mathbb{Z} \oplus 0.$$

To see this, note that as a $\mathbb{Z}/2$-graded chain complex, $CF^\infty(Y)$ is naturally a (finitely generated, free) module over the ring of Laurent polynomials $\mathbb{Z}[U, U^{-1}]$. Moreover, its quotient by the action of U and U^{-1} is the complex $CF'(Y)$ defined above. More algebraically, we have that

$$CF'(Y) = CF^\infty(Y) \otimes_{\mathbb{Z}[U,U^{-1}]} \mathbb{Z},$$

where the homomorphism $\mathbb{Z}[U, U^{-1}] \longrightarrow \mathbb{Z}$ sends U to 1. Theorem 11.3 says that $HF^\infty(Y)$ is a free $\mathbb{Z}[U, U^{-1}]$-module of rank one. The claim about $HF'_*(Y)$ then follows immediately from the universal coefficients theorem spectral sequence (see, for instance [4]).

12. Connected sums

We have seen how $\widehat{HF}$ behaves under connected sum (Proposition 7.9), and this suffices to give a non-vanishing result for HF^+ under connected sums (Theorem 1.6). The purpose of the present subsection is to give a more precise description of the behaviour of HF^- under connected sum. (From this, and the calculation of HF^∞, it is easy to deduce the behaviour of HF^+ under this operation, as well.)

Note that $CF^-(Y,\mathfrak{s})$, viewed as a $\mathbb{Z}/2\mathbb{Z}$-graded chain complex, is finitely generated as a module over the ring $\mathbb{Z}[U]$.

Theorem 12.1. *Let Y_1 and Y_2 be a pair of oriented three-manifolds, equipped with* Spin^c *structures $\mathfrak{s}_1$ and $\mathfrak{s}_2$ respectively. Then,*

$$HF^-(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)\cong H_*\left(CF^-(Y_1,\mathfrak{s}_1)\otimes_{\mathbb{Z}[U]}CF^-(Y_2,\mathfrak{s}_2)\right).$$

Before proceeding with the proof of the above result, we give a consequence for rational homology three-spheres Y_1 and Y_2, using a field $\mathbb{F}$ instead of the base ring $\mathbb{Z}$. In this case, since $HF^-(Y,\mathfrak{s};\mathbb{F})$ is a finitely generated module over $\mathbb{F}[U]$, it splits as a direct sum of cyclic modules. Indeed, each cyclic summand is either isomorphic to $\mathbb{F}[U]$ or it has the form $\mathbb{F}[U]/U^n$ for some non-negative integer n, since if some polynomial in U, $f(U)$, acts trivially on any element $\xi\in HF^-(Y,\mathfrak{s})$, then clearly U must divide f. We call this exponent n the *order* of the corresponding generator, i.e. given a generator $\xi\in HF^-(Y,\mathfrak{s})$ as a $\mathbb{F}[U]$-module, we define its order

$$\mathrm{ord}(\xi)=\max\{i\in\mathbb{Z}^{\geq 0}\big|U^i\cdot\xi\neq 0\}.$$

Note that by the structure of $HF^\infty(Y,\mathfrak{s})$, in any set of generators for $HF^-(Y,\mathfrak{s})$ there is exactly one with infinite order.

Corollary 12.2. *Let $\mathbb{F}$ be a field, and fix rational homology spheres Y_1 and Y_2. Let ξ_i for $i=0,...,M$ resp. η_j for $j=0,...,N$ be generators of $HF^-(Y_1,\mathfrak{s}_1;\mathbb{F})$ resp. $HF^-(Y_2,\mathfrak{s}_2;\mathbb{F})$ as a $\mathbb{F}[U]$-module. We order these so that $\mathrm{ord}(\xi_0)=\mathrm{ord}(\eta_0)=+\infty$. Then, $HF^-(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2;\mathbb{F})$ is generated as a $\mathbb{F}[U]$-module by generators $\xi_i\otimes\eta_j$ with $(i,j)\in\{0,...,M\}\times\{0,...,N\}$ and also by generators $\xi_i*\eta_j$ for $(i,j)\in\{1,...,M\}\times\{1,...N\}$. Moreover, for all $(i,j)\in\{0,...M\}\times\{0,...,N\}$,*

$$\mathrm{ord}(\xi_i\otimes\eta_j)=\min(\mathrm{ord}(\xi_i),\mathrm{ord}(\eta_j))$$

and

$$\mathrm{gr}(\xi_i\otimes\eta_j)=\mathrm{gr}(\xi_i)+\mathrm{gr}(\eta_j);$$

while for all $(i,j)\in\{1,...,M\}\times\{1,...,N\}$, we have that

$$\mathrm{ord}(\xi_i*\eta_j)=\min(\mathrm{ord}(\xi_i),\mathrm{ord}(\eta_j))$$

and

$$\mathrm{gr}(\xi_i*\eta_j)=\mathrm{gr}(\xi_i)+\mathrm{gr}(\eta_j)-1.$$

In particular, we have that

$$\chi\left(HF^-_{red}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)\right)=\chi\left(HF^-_{red}(Y_1,\mathfrak{s}_1)\right)+\chi\left(HF^-_{red}(Y_2,\mathfrak{s}_2)\right).$$

Proof. This is an immediate application of Theorem 12.1 and the Künneth formula for chain complexes over the principal ideal domain $\mathbb{F}[U]$. Specifically, we have that

$$HF^-(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)\cong\left(HF^-(Y_1,\mathfrak{s}_1)\otimes_{\mathbb{F}[U]}HF^-(Y_1,\mathfrak{s}_2)\right)\oplus\left(HF^-(Y_1,\mathfrak{s}_1)*HF^-(Y_1,\mathfrak{s}_2)\right),$$

where $A*B$ denotes the Tor-complex, i.e.

$$(A*B)_k\cong\bigoplus_{i+j=k-1}\mathrm{Tor}_{\mathbb{F}[U]}(A_i,B_j).$$

It is easy to see then that for any pair of non-negative integers m and n,

$$(\mathbb{F}[U]/U^m)\otimes_{\mathbb{F}[U]}(\mathbb{F}[U]/U^n)\cong\mathbb{F}[U]/U^{\min(m,n)}\cong\mathrm{Tor}_{\mathbb{F}[U]}(\mathbb{F}[U]/U^m,\mathbb{F}[U]/U^n);$$

while for any $\mathbb{F}[U]$-module M, $\mathbb{F}[U]\otimes_{\mathbb{F}[U]}M\cong M$ and $\mathrm{Tor}_{\mathbb{F}[U]}(\mathbb{F}[U],M)=0$.

To see the Euler characteristic statement, we proceed as follows. First, observe that to calculate the Euler characteristic of the graded $\mathbb{Z}$-module $HF^-(Y,\mathfrak{s})$ is the same as the Euler characteristic of the $\mathbb{Q}$-vector space $HF^-(Y,\mathfrak{s};\mathbb{Q})$. From above, we have that $HF^-_{red}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2;\mathbb{Q})$ is freely generated over $\mathbb{Q}$ by $i,j\in\{0,...,M\}\times\{0,...,N\}-\{0,0\}$ with $U^m\xi_i\otimes\eta_j$ where $m\in 0,...,\mathrm{ord}(\xi_i\otimes\eta_j)$ (observe that all generators of the form $U^m(\xi_0\otimes\eta_0)$ inject into $HF^\infty(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2;\mathbb{F})$) and also generators $U^m(\xi_i*\eta_j)$ for $(i,j)\in\{1,...M\}\times\{1,...,N\}$ and $m\in\{0,...,\mathrm{ord}(\xi_i*\eta_j)\}$. Observe in particular that when i,j are both non-zero, $U^m(\xi_i\otimes\eta_j)$ has a corresponding element $U^m(\xi_i*\eta_j)$ whose degree differs by one, so these cancel in the Euler characteristic. The only remaining elements are those of the form $U^m(\xi_i\otimes\eta_0)$ with $i>0$ and $m\in 0,..,\mathrm{ord}(\xi_i)$, and also $U^n(\xi_0\otimes\eta_j)$ with $j>0$ and $n\in 0,...,\mathrm{ord}(\eta_j)$. These contribute $\chi(HF^-_{red}(Y_1,\mathfrak{s}_1))$ and $\chi(HF^-_{red}(Y_2,\mathfrak{s}_2))$ to the Euler characteristic $\chi(HF^-_{red}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2))$ respectively. □

Proof of Theorem 12.1 when $b_1(Y_1\#Y_2)=0$. First, we construct a chain map

$$\Gamma\colon CF^{\leq 0}(Y_1,\mathfrak{s}_1)\otimes_{\mathbb{Z}[U]}CF^{\leq 0}(Y_2,\mathfrak{s}_2)\longrightarrow CF^{\leq 0}(Y_1\#\mathfrak{s}_2,\mathfrak{s}_1\#\mathfrak{s}_2).$$

To this end, consider pointed Heegaard diagrams $(\Sigma_1,\boldsymbol{\alpha},\boldsymbol{\beta},z_1)$ and $(\Sigma_2,\boldsymbol{\xi},\boldsymbol{\eta},z_2)$ for Y_1 and Y_2 respectively. Then there is connected sum Heegaard triple $(\Sigma_1\#\Sigma_2,\boldsymbol{\alpha}\cup\boldsymbol{\xi},\boldsymbol{\beta}\cup\boldsymbol{\xi},\boldsymbol{\beta}\cup\boldsymbol{\eta},z)$. This triple describes a cobordism from $Y_1\#(\#^{g_2}(S^2\times S^1))\coprod(\#^{g_1}(S^2\times S^1))\#Y_2$ to $Y_1\#Y_2$ where where g_1 and g_2 are the genera of Σ_1 and Σ_2 respectively. In fact, we let $\boldsymbol{\beta}'$ and $\boldsymbol{\xi}'$ be exact Hamiltonian translates of the $\boldsymbol{\beta}$ and $\boldsymbol{\xi}$ respectively, so that the new triple

$$(\Sigma_1\#\Sigma_2,\boldsymbol{\alpha}\cup\boldsymbol{\xi}',\boldsymbol{\beta}\cup\boldsymbol{\xi},\boldsymbol{\beta}'\cup\boldsymbol{\eta},z),$$

is admissible. We let $\Theta_1\in\mathbb{T}_\beta\cap\mathbb{T}'_\beta$ and $\Theta_2\in\mathbb{T}_\xi\cap\mathbb{T}'_\xi$ denote the "top" intersection points in $\mathrm{Sym}^{g_1}(\Sigma_1)$ resp. $\mathrm{Sym}^{g_2}(\Sigma_2)$ between the tori corresponding to $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$ resp. $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$. Thus, the maps $[\mathbf{x},i]\mapsto[\mathbf{x}\times\Theta_2,i]$ and $[\mathbf{y},j]\mapsto[\Theta_1\times\mathbf{y},j]$ are chain maps from

$$\Phi_1\colon CF^{\leq 0}(Y_1,\mathfrak{s}_1)\longrightarrow CF^{\leq 0}(Y_1\#^{g_2}(S^2\times S^1),\mathfrak{s}_1\#\mathfrak{s}_0)$$

and

$$\Phi_2\colon CF^{\leq 0}(Y_2,\mathfrak{s}_2)\longrightarrow CF^{\leq 0}(\#^{g_1}(S^2\times S^1)Y_2,\mathfrak{s}_0\#\mathfrak{s}_2)$$

are the chain maps considered in Proposition 7.10. Now, we define Γ to be the the composite of $\Phi_1\otimes\Phi_2$ with the map

$$F\colon CF^{\leq 0}(Y_1\#^{g_2}(S^2\times S^1),\mathfrak{s}_1\#\mathfrak{s}_0)\otimes CF^{\leq 0}((\#^{g_1}(S^2\times S^1))\#Y_2,\mathfrak{s}_0\#\mathfrak{s}_2)\longrightarrow CF^{\leq 0}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)$$

defined by counting holomorphic triangles in the Heegaard triple considered above. Observe that $F([\mathbf{x},i-1]\otimes[\mathbf{y},j]) = F([\mathbf{x},i]\otimes[\mathbf{y},j-1])$, so that $F\circ(\Phi_1\otimes\Phi_2)$ is $\mathbb{Z}[U]$-bilinear, inducing the $\mathbb{Z}[U]$-equivariant chain map Γ.

Suppose that $\boldsymbol{\beta}'$ is sufficiently close to the $\boldsymbol{\beta}$. Then, for each intersection point $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, there is a unique closest intersection point $\mathbf{x}'\in\mathbb{T}_\alpha\cap\mathbb{T}'_\beta$; similarly, when $\boldsymbol{\xi}'$ is sufficiently close to $\boldsymbol{\xi}$, each intersection point $\mathbf{y}\in\mathbb{T}_\xi\cap\mathbb{T}_\eta$ corresponds to a unique closest intersection point $\mathbf{y}'\in\mathbb{T}'_\xi\cap\mathbb{T}_\eta$. In this case, there is an obvious map

$$\Gamma_0\colon CF^{\leq 0}(Y_1,\mathfrak{s}_1)\otimes_{\mathbb{Z}[U]}CF^{\leq 0}(Y_2,\mathfrak{s}_2)\longrightarrow CF^{\leq 0}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)$$

defined by

$$\Gamma_0([\mathbf{x},i]\otimes[\mathbf{y},j]) = [\mathbf{x}'\times\mathbf{y}', i+j].$$

The map ψ_0 is not necessarily a chain map, but it is clearly an isomorphism of relatively $\mathbb{Z}$-graded groups. Indeed, we claim that when the total unsigned area ϵ in the regions between the ξ_i and the corresponding ξ'_i (resp. β_i and corresponding β'_i) is sufficiently small, then, for the induced energy filtration on $CF^{\leq 0}(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)$, we have that

$$\Gamma = \Gamma_0 + \text{lower order}.$$

This is true because there is an obvious small holomorphic triangle ψ with $n_z(\psi)=0$, $\mu(\psi)=0$, and $\#\mathcal{M}(\psi)=1$ connecting $\mathbf{x}\times\Theta_2$, $\Theta_1\times\mathbf{y}$, and $\mathbf{x}'\times\mathbf{y}'$. The total area of this triangle is bounded by the total area ϵ (which we can arrange to be smaller than any other triangle $\psi'\in\pi_2(\mathbf{x}\times\Theta_2,\Theta_1\times\mathbf{y},\mathbf{w})$). Since the energy filtration is bounded below in each degree (where now we view the complexes as relatively $\mathbb{Z}$-graded modules over $\mathbb{Z}$), it follows that Φ also induces an isomorphism in each degree. It follows that Γ induces an isomorphism of $\mathbb{Z}$-modules

$$\gamma\colon H_*\left(CF^{\leq 0}(Y_1,\mathfrak{s}_1)\otimes_{\mathbb{Z}[U]}CF^{\leq 0}(Y_2,\mathfrak{s}_2)\right)\longrightarrow HF^{\leq 0}(Y_1\#\mathfrak{s}_2,\mathfrak{s}_1\#\mathfrak{s}_2).$$

□

The above proof actually holds provided that $\mathfrak{s}_1\#\mathfrak{s}_2$ is a torsion Spin^c structure, so that chain complex $CF^-(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)$ is finitely generated in each degree.

For non-torsion Spin^c structures, we must use a refined filtration. Suppose that Y is a three-manifold endowed with a Spin^c structure $\mathfrak{s}$ whose first Chern class is non-torsion. We will work with a special class of Heegaard diagrams: those for which each $\mathfrak{s}$-renormalized periodic domain (in the sense of Definition 5.3 has both positive and negative coefficients. These are the Heegaard diagrams constructed by winding as in Lemma 5.4. Next, we equip Σ with a volume form for which each $\mathfrak{s}$-renormalized periodic domain has total signed area zero. This can be arranged as in the proof of Lemma 4.4.

Now, given $[\mathbf{x},i]$ and $[\mathbf{y},j]$ with the same grading, we can find some disk $\phi\in\pi_2(\mathbf{x},\mathbf{y})$ with $n_z(\phi)=i-j$ and $\mu(\phi)=0$. We then define the filtration difference to be the area of the domain associated to ϕ:

$$\mathcal{F}([\mathbf{x},i],[\mathbf{y},j]) = -\mathcal{A}(\mathcal{D}(\phi)).$$

Since any possible choices of such disk ϕ, ϕ' differ by a renormalized periodic domain, it follows that the filtration defined above depends only on the disk.

Letting $\delta = \delta(\mathfrak{s})$ be the divisibility of $\mathfrak{s}$, we see that the subalgebra $\mathbb{Z}[U^\delta] \subset \mathbb{Z}[U]$ acts on $CF^-(Y,\mathfrak{s})$ preserving the relative $\mathbb{Z}/\delta(\mathfrak{s})\mathbb{Z}$-grading. We claim that $U^\delta[\mathbf{x},i]$ and $[\mathbf{x},i]$ can be connected by a map ϕ whose underlying domain is a renormalized periodic domain; thus, the filtration $\mathcal{F}$ we have defined above is a filtration of $CF^-(Y,\mathfrak{s})$ as a $\mathbb{Z}[U^\delta]$-module.

Proof of Theorem 12.1 when $b_1(Y_1\#Y_2)>0$

We argue first that the connected sum $Y_1\#Y_2$ can be endowed with a Heegaard diagram which is both special in the above sense (each $\mathfrak{s}_1\#\mathfrak{s}_2$-renormalied periodic domain has total area zero), and it also splits as a sum of Heegaard diagrams $(\Sigma_1\#\Sigma_2, \boldsymbol{\alpha}\cup\boldsymbol{\xi}, \boldsymbol{\beta}\cup\boldsymbol{\eta}, z)$. This is done by winding the $\boldsymbol{\alpha}$ within Σ_1, and the $\boldsymbol{\beta}$ within Σ_2. As in the proof of the theorem when $b_1(Y_1\#Y_2)=0$, we consider the Heegaard triple

$$(\Sigma_1\#\Sigma_2, \boldsymbol{\alpha}\cup\boldsymbol{\xi}', \boldsymbol{\beta}\cup\boldsymbol{\xi}, \boldsymbol{\beta}'\cup\boldsymbol{\eta}, z),$$

where $\boldsymbol{\xi}'$ and $\boldsymbol{\beta}'$ are obtained as sufficiently small Hamiltonian translates of the original $\boldsymbol{\xi}$ and $\boldsymbol{\beta}$, letting ϵ denote the total (unsigned) areas in the regions between the original curves and their Hamiltonian translates.

When $\mathfrak{s}_1\#\mathfrak{s}_2$ is a torsion Spin^c structure, the proof given under the assumption that $b_1(Y_1\#Y_2)=0$ adapts immediately in the present context. We claim that even when $\mathfrak{s}_1\#\mathfrak{s}_2$ is non-torsion, we can write

$$\Gamma = \Gamma_0 + \text{lower order}, \tag{21}$$

where now the lower order terms have lower order with respect to the filtration $\mathcal{F}$ defined right before this proof. To see this, suppose that ψ is a holomorphic triangle which contributes to Γ, i.e. $\psi\in\pi_2(\mathbf{x}\times\mathbf{y}, \Theta_1\times\Theta_2, \mathbf{p}\times\mathbf{q})$ satisfies $\mu(\psi)=0$ and $\mathcal{D}(\psi)>0$, while $\psi_0\in\pi_2(\mathbf{x}\times\mathbf{y},\Theta_1\times\Theta_2,\mathbf{x}'\times\mathbf{y}')$ is the canonical small triangle. Assuming that $\mathbf{x}'\times\mathbf{y}'\neq\mathbf{p}\times\mathbf{q}$, we argue that

$$\mathcal{F}([\mathbf{x}'\times\mathbf{y}',i],[\mathbf{p}\times\mathbf{q}, i-n_z(\psi)])<0.$$

To see this, find some $\phi\in\pi_2(\mathbf{x}'\times\mathbf{y}',\mathbf{p}\times\mathbf{q})$ with $\mu(\phi)=0$, so that both $\psi,\psi_0*\phi\in\pi_2(\mathbf{x}\times\mathbf{y},\Theta_1\times\Theta_2,\mathbf{p}\times\mathbf{q})$ have $\mu(\psi)=\mu(\psi_0+\phi)=0$. Now, we claim that

$$\mathcal{A}(\psi)=\mathcal{A}(\psi_0+\phi),$$

since the difference is a triply-periodic domain, while the $\boldsymbol{\xi}'$ and $\boldsymbol{\eta}'$ are obtained from $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ by exact Hamiltonian translation. Since $\mathcal{A}(\psi)>\epsilon$, while $\mathcal{A}(\psi_0)<\epsilon$, it follows that $\mathcal{A}(\phi)$ is positive.

Now, since $CF^-(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)$ is finitely generated as a chain complex over the polynomial algebra $\mathbb{Z}[U^\delta]$, and the the filtration is defined over this algebra, we can conclude from Equation (21) that Γ is an isomorphism of $\mathbb{Z}[U^\delta]$-chain complexes. The theorem then follows. $\square$

13. Applications

In this section, we prove the remaining results (Theorems 1.8 and 1.12) claimed in the introduction.

13.1. Complexity of three-manifolds.

The theorems in the introduction dealing with fractional surgeries are proved using surgery exact sequences with twisted theories (Theorems 10.14 and 10.17). Consequently, we will need the following analogue of Theorem 9.1 for the twisted theory:

Lemma 13.1. *Let Y_0 be a homology $S^1 \times S^2$, and choose a coefficient system corresponding to a representation*

$$H^1(Y_0;\mathbb{Z}) \longrightarrow \mathbb{Z}/n\mathbb{Z}.$$

Then, for each non-torsion Spin^c *structure over Y_0, we have that*

$$\chi(\underline{HF^+}(Y_0,\mathbb{Z}/n\mathbb{Z};\mathfrak{s})) = n\cdot\chi(HF^+(Y_0,\mathfrak{s})) = -n\cdot\tau_t(Y_0,\mathfrak{s})$$

(where on the left we are still taking the rank as a $\mathbb{Z}$-module, and t here is the component of $H^2(Y;\mathbb{Z})-0$ containing $c_1(\mathfrak{s})$). Similarly, for a torsion Spin^c *structure $\mathfrak{s}_0$, we have that*

$$\chi(HF^+_{\leq 2n+1}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z}) = -n\cdot\tau(Y_0,\mathfrak{s}_0).$$

Proof. The proof proceeds exactly as in the proof of Theorem 9.1 (with the sign pinned down in Proposition 11.13, and Theorem 11.16 in the case where the Spin^c structure is torsion), together with the observation that now $\chi(\mathrm{Ker} f_1)$ multiplies by n. □

We will also need the following result, which follows along the lines of Section 11.

Lemma 13.2. *Suppose that Y_0 is a homology $S^1\times S^2$, and choose a coefficient system corresponding to a map $H^1(Y_0;\mathbb{Z})\cong\mathbb{Z}\longrightarrow\mathbb{Z}/n\mathbb{Z}$ which maps generators to generators. Then, if $\mathfrak{s}_0$ is a torsion* Spin^c *structure, then $\underline{HF^\infty_i}(Y_0,\mathfrak{s}_0,\mathbb{Z}/n\mathbb{Z})\cong\mathbb{Z}$ in all degrees.*

Proof. We still have the long exact sequence

$$...\longrightarrow \underline{HF^\infty}(Y_0,\mathfrak{s}_0,\mathbb{Z}/n\mathbb{Z}) \longrightarrow H_*(\underline{R}^\infty,\mathbb{Z}/n\mathbb{Z}) \xrightarrow{\delta} H_*(\underline{L}^\infty,\mathbb{Z}/n\mathbb{Z}) \longrightarrow ...$$

We place a reference point p at the intersection of γ (the perturbing curve) with α_1. It is clear that $H_*(\underline{L}^\infty,\mathbb{Z}/n\mathbb{Z})\cong H_*(L^\infty)\otimes_{\mathbb{Z}}\mathbb{Z}[\mathbb{Z}/n\mathbb{Z}]$. Moreover, the coboundary splits as $\delta=\delta_1-\zeta\delta_2$, where ζ is is a primitive n^{th} root of unity, and δ_1 and δ_2 are the maps obtained from the δ_1 and δ_2 using $\mathbb{Z}$ coefficients, by a base-change to $\mathbb{Z}/n\mathbb{Z}$. In particular, both δ_1 and δ_2 are isomorphisms (Lemmas 11.7 and 11.8). Thus, in view of Theorem 11.1 (indeed, we're using here the special cases from Subsection 11.1), we have exactness for

$$0\longrightarrow\underline{HF^\infty_i}(Y_0,\mathfrak{s}_0,\mathbb{Z}/n\mathbb{Z})\longrightarrow\mathbb{Z}[\mathbb{Z}/n\mathbb{Z}]\xrightarrow{1-\zeta}\mathbb{Z}[\mathbb{Z}/n\mathbb{Z}]\longrightarrow\underline{HF^\infty_{i-1}}(Y_0,\mathfrak{s}_0,\mathbb{Z}/n\mathbb{Z})\longrightarrow 0$$

□

We can now prove Theorem 1.8.

Proof of Theorem 1.8. This is an application of the U-equivariant exact sequence of Theorem 10.14, which gives:

$$... \xrightarrow{F_1} \underline{HF^+}(Y_0;\mathbb{Z}/n\mathbb{Z}) \xrightarrow{F_2} HF^+(Y_{1/n}) \xrightarrow{F_3} HF^+(Y) \longrightarrow ...,.$$

Now, we claim that for all sufficiently large d, the map induced by F_2

$$\mathrm{Im} U^d \underline{HF^+}(Y_0,\mathbb{Z}/n\mathbb{Z}) \longrightarrow \mathrm{Im} U^d HF^+(Y_{1/n})$$

is surjective. It suffices to consider the $\mathfrak{s}_0$-summand of $\underline{HF^+}(Y_0,\mathbb{Z}/n\mathbb{Z})$, where $\mathfrak{s}_0$ is the torsion Spin^c structure. There, F_2 has a natural $\mathbb{Z}$-graded lift. For one parity, the corresponding $HF^\infty(Y_{1/n})$ vanishes (so the claim is obvious). For the other parity, in sufficiently high degree k, the image of F_1 is trivial, so, with the help of Lemma 13.2, our exact sequence reads:

$$0 \longrightarrow \underline{HF^+_k}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z}) \cong \underline{HF^\infty_k}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z}) \cong \mathbb{Z} \xrightarrow{F_2} HF^+_k(Y_{1/n}) \cong \mathbb{Z}.$$

Since $HF^\infty(Y)$ has no torsion, it easily follows that F_2 must surject onto the generator in $HF^+_k(Y_{1/n})$.

From this observation, together with the U-equivariant exact sequence, it follows that the map

$$\frac{HF^+(Y)}{U^d HF^+(Y)} \longrightarrow \frac{\underline{HF^+}(Y_0,\mathbb{Z}/n\mathbb{Z})}{U^d \underline{HF^+}(Y_0,\mathbb{Z}/n\mathbb{Z})} \longrightarrow \frac{HF^+(Y_{1/n})}{U^d HF^+(Y_{1/n})}.$$

is exact in the middle, and hence that

$$\mathrm{rk}\left(\underline{HF_{\mathrm{red}}}(Y_0,\mathbb{Z}/n\mathbb{Z})\right) \leq \mathrm{rk}\left(\underline{HF_{\mathrm{red}}}(Y)\right) + \mathrm{rk}\left(\underline{HF_{\mathrm{red}}}(Y_1)\right). \tag{22}$$

(Here, as in the case where $b_1 = 0$, $\underline{HF_{\mathrm{red}}}(Y_0,\mathbb{Z}/n\mathbb{Z})$ is defined to be the quotient of $\underline{HF^+}(Y_0,\mathbb{Z}/n\mathbb{Z})$ by the image of $\underline{HF^\infty}(Y_0,\mathbb{Z}/n\mathbb{Z})$.)

Now, observe that if $\mathfrak{s} \neq \mathfrak{s}_0$, $HF^+(Y_0,\mathfrak{s};\mathbb{Z}/n\mathbb{Z})$ is finitely generated, so that for sufficiently large d,

$$\underline{HF_{\mathrm{red}}}(Y_0,\mathfrak{s};\mathbb{Z}/n\mathbb{Z}) = \frac{\underline{HF^+}(Y_0,\mathfrak{s};\mathbb{Z}/n\mathbb{Z})}{U^d \underline{HF^+}(Y_0,\mathfrak{s},\mathbb{Z}/n\mathbb{Z})} = \underline{HF^+}(Y_0,\mathfrak{s};\mathbb{Z}/n\mathbb{Z}). \tag{23}$$

For $\mathfrak{s} = \mathfrak{s}_0$, we observe that

$$\max(0, -\chi(HF^+_{\leq 2n+1}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z}))) \leq \mathrm{rk} HF^+_{\leq 0}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z}). \tag{24}$$

The reason for this is that for all sufficiently large n, we have

$$\chi(\underline{HF^+}_{\leq 2n+1}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z})) = \\ \chi\left(\underline{HF_{\mathrm{red}}}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z})\right) + \chi\left(\underline{HF^+}_{\leq 2n+1}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z}) \cap \mathrm{Im}\underline{HF^\infty}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z})\right).$$

The second term above is negative: owing to the algebraic structure of $HF^\infty(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z})$ (the even-dimensional generators are the images of the odd-dimensional ones under an isomorphism), there are more odd-dimensional than even-dimensional generators coming from $U^d\underline{HF^+}(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z})$ in $\underline{HF^+}_{\leq 2n+1} f(Y_0,\mathfrak{s}_0;\mathbb{Z}/n\mathbb{Z})$.

The theorem is obtained by combining Inequality (22), Equation (23), Inequality (24), and Lemma 13.1. □

13.2. **Gradient trajectories.** We turn to Theorem 1.12.

Proof of Theorem 1.12. As a first step, observe that, since

$$\chi(\underline{HF^+}(Y_0, \mathfrak{s}_0 \pm iH; \mathbb{Z}/n\mathbb{Z})) = \pm n \cdot t_i(K),$$

it follows that the rank of $HF^+(Y_0, \mathbb{Z}/n\mathbb{Z}, \mathfrak{s})$ is non-zero for at least $2k$ distinct non-torsion Spin^c structures; thus the rank of $\underline{\widehat{HF}}(Y_0, \mathfrak{s}, \mathbb{Z}/n\mathbb{Z})$ is also non-zero in these Spin^c structures (c.f. Proposition 7.1). Moreover, from Lemma 13.2, it follows that the rank of $\underline{HF^+}(Y_0, \mathbb{Z}/n\mathbb{Z}, \mathfrak{s}_0)$ is non-zero, and hence so is the rank of $\underline{\widehat{HF}}(Y_0, \mathfrak{s}_0, \mathbb{Z}/n\mathbb{Z})$. Now, since for all Spin^c structures,

$$\chi(\underline{\widehat{HF}}(Y_0, \mathfrak{s}, \mathbb{Z}/n\mathbb{Z}) = 0$$

(again, using the twisted analogue of Proposition 7.7), the rank of $\widehat{HF}(Y_0, \mathbb{Z}/n\mathbb{Z})$ is at least $4k + 2$. The result then follows from the exact sequence of Theorem 10.17 and the analogue of Proposition 7.5 with twisted coefficients. □

References

[1] D. Auckly. The Thurston norm and three-dimensional Seiberg-Witten theory. *Osaka J. Math*, 33(3):737–750, 1996.

[2] D. M. Austin and P. J. Braam. Morse-Bott theory and equivariant cohomology. In H. Hofer, C. H. Taubes, A. Weinstein, and E. Zehnder, editors, *The Floer Memorial Volume*, number 133 in Progress in Mathematics, pages 123–183. Birkhäuser, 1995.

[3] P. Braam and S. K. Donaldson. Floer's work on instanton homology, knots, and surgery. In H. Hofer, C. H. Taubes, A. Weinstein, and E. Zehnder, editors, *The Floer Memorial Volume*, number 133 in Progress in Mathematics, pages 195–256. Birkhäuser, 1995.

[4] H. Cartan and S. Eilenberg. *Homological algebra*. Princeton University Press, 1956.

[5] M. Culler, C. McA. Gordon, J. Luecke, and P. B. Shalen. Dehn surgery on knots. *Ann. of Math*, 125(2):237–300, 1987.

[6] R. Fintushel and R. J. Stern. Integer graded instanton homology groups for homology three-spheres. *Topology*, 31(3):589–604, 1992.

[7] A. Floer. An instanton-invariant for 3-manifolds. *Comm. Math. Phys.*, 119:215–240, 1988.

[8] A. Floer. The unregularized gradient flow of the symplectic action. *Comm. Pure Appl. Math.*, 41(6):775–813, 1988.

[9] A. Floer. Instanton homology and Dehn surgery. In H. Hofer, C. H. Taubes, A. Weinstein, and E. Zehnder, editors, *The Floer Memorial Volume*, number 133 in Progress in Mathematics, pages 77–97. Birkhäuser, 1995.

[10] A. Floer and H. Hofer. Coherent orientations for periodic orbit problems in symplectic geometry. *Math. Z.*, 212(1):13–38, 1993.

[11] R. E. Gompf and A. I. Stipsicz. *4-manifolds and Kirby calculus*, volume 20 of *Graduate Studies in Mathematics*. American Mathematical Society, 1999.

[12] C. McA. Gordon and J. Luecke. Knots are determined by their complements. *J. Amer. Math. Soc.*, 2(2):371–415, 1989.

[13] C. McA. Gordon and J. Luecke. Knots are determined by their complements. *Bull. Amer. Math. Soc. (N.S.)*, 20(1):83–87, 1989.

[14] M. Hutchings and Y-J. Lee. Circle-valued Morse theory, Reidemeister torsion, and Seiberg-Witten invariants of 3-manifolds. *Topology*, 38(4):861–888, 1999.

[15] P. B. Kronheimer and T. S. Mrowka, 1996. I.P. Lecture.

[16] P. B. Kronheimer and T. S. Mrowka. Monopoles and contact structures. *Invent. Math.*, 130(2):209–255, 1997.

[17] P. B. Kronheimer and T. S. Mrowka. Scalar curvature and the Thurston norm. *Math. Res. Lett.*, (4):931–937, 1997.

[18] I. G. MacDonald. Symmetric products of an algebraic curve. *Topology*, 1:319–343, 1962.

[19] M. Marcolli and B-L. Wang. Equivariant Seiberg-Witten Floer homology. dg-ga/9606003.

[20] D. McDuff and D. Salamon. *J-holomorphic curves and quantum cohomology*. Number 6 in University Lecture Series. American Mathematical Society, 1994.

[21] G. Meng and C. H. Taubes. $\underline{\mathrm{SW}}$=Milnor torsion. *Math. Research Letters*, 3:661–674, 1996.

[22] J. Milnor. *Lectures on the h-cobordism theorem*. Princeton University Press, 1965. Notes by L. Siebenmann and J. Sondow.

[23] P. S. Ozsváth and Z. Szabó. Holomorphic disks and topological invariants for rational homology three-spheres. math.SG/0101206, 2000.

[24] P. S. Ozsváth and Z. Szabó. On embedding circle-bundles in four-manifolds. *Math. Res. Lett.*, 7(5-6):657–669, 2000.

[25] P. S. Ozsváth and Z. Szabó. The theta divisor and three-manifold invariants. math.GT/0006194, 2000.

[26] P. S. Ozsváth and Z. Szabó. Holomorphic triangles and invariants for smooth four-manifolds. 2001. In preparation.

[27] K. Reidemeister. Zur dreidimensionalen Topologie. *Abh. Math. Sem. Univ. Hamburg*, (9):189–194, 1933.

[28] J. Singer. Three-dimensional manifolds and their Heegaard diagrams. *Trans. Amer. Math. Soc.*, (35):88–111, 1933.

[29] C. H. Taubes. SW $\to$ Gr: from the Seiberg-Witten equations to pseudo-holomorphic curves. *J. Amer. Math. Soc.*, 9(3):845–918, 1996.

[30] C. H. Taubes. The geometry of the Seiberg-Witten invariants. In *Proceedings of the International Congress of Mathematicians, Vol. II*, pages 493–504, 1998.

[31] W. P. Thurston. *A norm for the homology of* 3*-manifolds*, volume 59 of *Mem. Amer. Math. Soc.*, pages i–vi and 99–130. 1986.

[32] V. Turaev. Torsion invariants of Spin^c-structures on 3-manifolds. *Math. Research Letters*, 4:679–695, 1997.

[33] V. Turaev. A combinatorial formulation for the Seiberg-Witten invariants of 3-manifolds. *Math. Research Letters*, 5(5):583–598, 1998.

Department of Mathematics, Princeton University, New Jersey 08540
petero@math.princeton.edu

Department of Mathematics, Princeton University, New Jersey 08540
szabo@math.princeton.edu