

Holomorphic disks and topological invariants for closed three-manifolds

By PETER OZSVÁTH and ZOLTÁN SZABÓ*

Abstract

The aim of this article is to introduce certain topological invariants for closed, oriented three-manifolds Y , equipped with a Spin^c structure. Given a Heegaard splitting of $Y = U_0 \cup_{\Sigma} U_1$, these theories are variants of the Lagrangian Floer homology for the g -fold symmetric product of Σ relative to certain totally real subspaces associated to U_0 and U_1 .

1. Introduction

Let Y be a connected, closed, oriented three-manifold, equipped with a Spin^c structure \mathfrak{s} . Our aim in this paper is to define certain Floer homology groups $\widehat{HF}(Y, \mathfrak{s})$, $HF^+(Y, \mathfrak{s})$, $HF^-(Y, \mathfrak{s})$, $HF^\infty(Y, \mathfrak{s})$, and $HF_{\text{red}}(Y, \mathfrak{s})$ using Heegaard splittings of Y . For calculations and applications of these invariants, we refer the reader to the sequel, [28].

Recall that a Heegaard splitting of Y is a decomposition $Y = U_0 \cup_{\Sigma} U_1$, where U_0 and U_1 are handlebodies joined along their boundary Σ . The splitting is determined by specifying a connected, closed, oriented two-manifold Σ of genus g and two collections $\{\alpha_1, \dots, \alpha_g\}$ and $\{\beta_1, \dots, \beta_g\}$ of simple, closed curves in Σ .

The invariants are defined by studying the g -fold symmetric product of the Riemann surface Σ , a space which we denote by $\text{Sym}^g(\Sigma)$: i.e. this is the quotient of the g -fold product of Σ , which we denote by $\Sigma^{\times g}$, by the action of the symmetric group on g letters. There is a quotient map

$$\pi: \Sigma^{\times g} \longrightarrow \text{Sym}^g(\Sigma).$$

$\text{Sym}^g(\Sigma)$ is a smooth manifold; in fact, a complex structure on Σ naturally gives rise to a complex structure on $\text{Sym}^g(\Sigma)$, for which π is a holomorphic map.

*PSO was supported by NSF grant number DMS-9971950 and a Sloan Research Fellowship. ZSz was supported by NSF grant number DMS-9704359, a Sloan Research Fellowship, and a Packard Fellowship.

In [7], Floer considers a homology theory defined for a symplectic manifold and a pair of Lagrangian submanifolds, whose generators correspond to intersection points of the Lagrangian submanifolds (when the Lagrangians are in sufficiently general position), and whose boundary maps count pseudo-holomorphic disks with appropriate boundary conditions. We spell out a similar theory, where the ambient manifold is $\text{Sym}^g(\Sigma)$ and the submanifolds playing the role of the Lagrangians are tori $\mathbb{T}_\alpha = \alpha_1 \times \cdots \times \alpha_g$ and $\mathbb{T}_\beta = \beta_1 \times \cdots \times \beta_g$. These tori are half-dimensional totally real submanifolds with respect to any complex structure on the symmetric product induced from a complex structure on Σ . These tori are transverse to one another when all the α_i are transverse to the β_j . To bring Spin^c structures into the picture, we fix a point $z \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$. We show in Section 2.6 that the choice of z induces a natural map from the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ to the set of Spin^c structures over Y .

While the submanifolds \mathbb{T}_α and \mathbb{T}_β in $\text{Sym}^g(\Sigma)$ are not *a priori* Lagrangian, we show that certain constructions from Floer's theory can still be applied, to define a chain complex $CF^\infty(Y, \mathfrak{s})$. This complex is freely generated by pairs consisting of an intersection point of the tori (which represents the given Spin^c structure) and an integer which keeps track of the intersection number of the holomorphic disks with the subvariety $\{z\} \times \text{Sym}^{g-1}(\Sigma)$; and its differential counts pseudo-holomorphic disks in $\text{Sym}^g(\Sigma)$ satisfying appropriate boundary conditions. Indeed, a natural filtration on the complex gives rise to an auxiliary collection of complexes $CF^-(Y, \mathfrak{s})$, $CF^+(Y, \mathfrak{s})$, and $\widehat{CF}(Y, \mathfrak{s})$. We let HF^- , HF^∞ , HF^+ , and \widehat{HF} denote the homology groups of the corresponding complexes.

These homology groups are relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{s})\mathbb{Z}$ -graded Abelian groups, where $\mathfrak{d}(\mathfrak{s})$ is the integer given by

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

where $c_1(\mathfrak{s})$ denotes the first Chern class of the Spin^c structure. In particular, when $c_1(\mathfrak{s})$ is a torsion class (which is guaranteed, for example, if $b_1(Y) = 0$), then the groups are relatively \mathbb{Z} -graded.

Moreover, we define actions

$$U: HF^\infty(Y, \mathfrak{s}) \longrightarrow HF^\infty(Y, \mathfrak{s})$$

and

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

which decrease the relative degree in $HF^\infty(Y, \mathfrak{s})$ by two and one respectively. These induce actions on \widehat{HF} , HF^+ , and HF^- (although the induced U -action on \widehat{HF} is trivial), endowing the homology groups with the structure of a module over $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(H_1(Y; \mathbb{Z})/\text{Tors})$. We show in Section 4 that the quotient

$HF^+(Y, \mathfrak{s})/U^d HF^+(Y, \mathfrak{s})$ stabilizes for all sufficiently large exponents d , and we let $HF_{\text{red}}(Y, \mathfrak{s})$ denote the group so obtained. After defining the groups, we turn to their topological invariance:

THEOREM 1.1. *The invariants $\widehat{HF}(Y, \mathfrak{s})$, $HF^-(Y, \mathfrak{s})$, $HF^\infty(Y, \mathfrak{s})$, $HF^+(Y, \mathfrak{s})$, and $HF_{\text{red}}(Y, \mathfrak{s})$, thought of as modules over*

$$\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^*(H_1(Y; \mathbb{Z})/\text{Tors}),$$

are topological invariants of Y and \mathfrak{s} , in the sense that they are independent of the Heegaard splitting, the choice of attaching circles, the basepoint z , and the complex structures used in their definition.

See also Theorem 11.1 for a more precise statement. The proof of the above theorem consists of many steps, and indeed, they take up the rest of the present paper.

In Section 2, we recall the topological preliminaries on Heegaard splittings and symmetric products used throughout the paper. In Section 3, we describe the modifications to the usual Lagrangian set-up which are necessary to define the totally real Floer homologies for the Heegaard splittings. In Subsection 3.3, we address the issue of smoothness for the moduli spaces of disks. In Subsection 3.4, we prove *a priori* energy estimates for pseudo-holomorphic disks which are essential for proving compactness results for the moduli spaces.

With these pieces in place, we define the Floer homology groups in Section 4. We begin with the technically simpler case of three-manifolds with $b_1(Y) = 0$, in Subsection 4.1. We then turn to the case where $b_1(Y) > 0$ in Section 4.2. In this case, we must work with a special class of Heegaard diagrams (so-called *admissible* diagrams) to obtain groups which are independent of the isotopy class of Heegaard diagram. The precise type of Heegaard diagram needed depends on the Spin^c structure in question, and the variant of $HF(Y, \mathfrak{s})$ which one wishes to consider. We define the types of Heegaard diagrams in Subsection 4.2.2, and discuss some of the additional algebraic structures on the homology theories when $b_1(Y) > 0$ in Subsection 4.2.5. With these definitions in hand, we turn to the construction of admissible Heegaard diagrams required when $b_1(Y) > 0$ in Section 5.

After defining the groups, we show that they are independent of initial analytical choices (complex structures) which go into their definition. This is established in Section 6, by use of chain homotopies which follow familiar constructions in Lagrangian Floer homology. Thus, the groups now depend on the Heegaard diagram.

In Section 7, we turn to the question of topological invariance. To show that we have a topological invariant for three-manifolds, we must show that the groups are invariant under the three basic Heegaard moves: isotopies of the attaching circles, handleslides among the attaching circles, and stabilizations of

the Heegaard diagram. Isotopy invariance is established in Subsection 7.3, and its proof is closely modeled on the invariance of Lagrangian Floer homology under exact Hamiltonian isotopies.

To establish handleslide invariance, we show that a handleslide induces a natural chain homotopy between the corresponding chain complexes. With a view towards this application, we describe in Section 8 the chain maps induced by counting holomorphic triangles, which are associated to three g -tuples of attaching circles. Indeed, we start with the four-dimensional topological preliminaries of this construction in Subsection 8.1, and turn to the Floer homological construction in later subsections. In fact, we set up this theory in more generality than is needed for handleslide invariance, to make our job easier in the sequel [28].

With the requisite naturality in hand, we turn to the proof of handleslide invariance in Section 9. This starts with a model calculation in $\#^g(S^1 \times S^2)$ (cf. Subsection 9.1), which we transfer to an arbitrary three-manifold in Subsection 9.2.

In Section 10, we prove stabilization invariance. In the case of \widehat{HF} , the result is quite straightforward, while for the others, we must establish certain gluing results for holomorphic disks.

In Section 11 we assemble the various components of the proof of Theorem 1.1.

1.1. On the Floer homology package. Before delving into the constructions, we pause for a moment to justify the profusion of Floer homology groups. Suppose for simplicity that $b_1(Y) = 0$.

Given a Heegaard diagram for Y , the complex underlying $CF^\infty(Y, \mathfrak{s})$ can be thought of as a variant of Lagrangian Floer homology in $\text{Sym}^g(\Sigma)$ relative to the subsets \mathbb{T}_α and \mathbb{T}_β , and with coefficients in the ring of Laurent polynomials $\mathbb{Z}[U, U^{-1}]$ to keep track of the homotopy classes of connecting disks. This complex in itself is independent of the choice of basepoint in the Heegaard diagram (and hence gives a homology theory which is independent of the choice of Spin^c structure on Y). Indeed (especially when $b_1(Y) = 0$) the homology groups of this complex turn out to be uninteresting (cf. Section 10 of [28]).

However, the choice of basepoint z gives rise to a \mathbb{Z} -filtration on $CF^\infty(Y, \mathfrak{s})$ which respects the action of the polynomial subalgebra $\mathbb{Z}[U] \subset \mathbb{Z}[U, U^{-1}]$. Indeed, the filtration has the following form: there is a $\mathbb{Z}[U]$ -subcomplex $CF^-(Y, \mathfrak{s}) \subset CF^\infty(Y, \mathfrak{s})$, and for $k \in \mathbb{Z}$, the k^{th} term in the filtration is given by $U^k CF^-(Y, \mathfrak{s}) \subset CF^\infty(Y, \mathfrak{s})$. It is now the chain homotopy type of CF^∞ as a *filtered* complex which gives an interesting three-manifold invariant. To detect this object, we consider the invariants HF^- , HF^+ , \widehat{HF} , and HF^∞ which are the homology groups of

$$CF^-, \quad \frac{CF^\infty}{CF^-}, \quad \frac{U^{-1} \cdot CF^-(Y, \mathfrak{s})}{CF^-(Y, \mathfrak{s})}, \quad \text{and} \quad CF^\infty$$

respectively. From their construction, it is clear that there are relationships between these various homology groups including, in particular, a long exact sequence relating HF^- , HF^∞ , and HF^+ . So, although HF^∞ in itself contains no interesting information, we claim that its subcomplex, quotient complex, and indeed the connecting maps all do.

1.2. *Further developments.* We give more motivation for these invariants, and their relationship with gauge theory, in the introduction to the sequel, [28]. Indeed, first computations and applications of these Floer homology groups are given in that paper. See also [29] where a corresponding smooth four-manifold invariant is constructed, and [27] where we endow the Floer homology groups with an absolute grading, and give topological applications of this extra structure.


1.3. *Acknowledgements.* We would like to thank Stefan Bauer, John Morgan, Tom Mrowka, Rob Kirby, and András Stipsicz for helpful discussions during the course of the writing of this paper.

2. Topological preliminaries

In this section, we 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 relationships with Heegaard diagrams.

2.1. *Heegaard diagrams.* A genus g Heegaard splitting of a connected, closed, oriented three-manifold Y is a decomposition of $Y = U_0 \cup_\Sigma U_1$ where Σ is an oriented, connected, closed 2-manifold with genus g , and U_0 and U_1 are handlebodies with $\partial U_0 = \Sigma = -\partial U_1$. Every closed, oriented three-manifold admits a Heegaard decomposition. For modern surveys on the theory of Heegaard splittings, see [34] and [41].

A handlebody U bounding Σ can be described using Kirby calculus. U is obtained from Σ by first attaching g two-handles along g disjoint, simple closed curves $\{\gamma_1, \dots, \gamma_g\}$ which are linearly independent in $H_1(\Sigma; \mathbb{Z})$, and then one three-handle. The curves $\gamma_1, \dots, \gamma_g$ are called *attaching circles for U* . Since the three-handle is unique, U is determined by the attaching circles. Note that the attaching circles are not uniquely determined by U . For example, they can be moved by isotopies. But more importantly, if $\gamma_1, \dots, \gamma_g$ are attaching circles for U , then so are $\gamma'_1, \gamma_2, \dots, \gamma_g$, where γ'_1 is obtained by “sliding” the handle of γ_1 over another handle, say, γ_2 ; i.e. γ'_1 is any simple, closed curve which is disjoint from the $\gamma_1, \dots, \gamma_g$ with the property that γ'_1, γ_1 and γ_2 bound an embedded pair of pants in $\Sigma - \gamma_3 - \dots - \gamma_g$ (see Figure 1 for an illustration in the $g = 2$ case).

Figure 1: Handlesliding γ_1 over γ_2

In view of these remarks, one can concretely think of a genus g Heegaard splitting of a closed three-manifold $Y = U_0 \cup_{\Sigma} U_1$ as specified by a genus g surface Σ , and a pair of g -tuples of curves in Σ , $\alpha = \{\alpha_1, \dots, \alpha_g\}$ and $\beta = \{\beta_1, \dots, \beta_g\}$, which are g -tuples of attaching circles for the U_0 - and U_1 -handlebodies respectively. The triple (Σ, α, β) is called a *Heegaard diagram*.

Note that Heegaard diagrams have a Morse-theoretic interpretation as follows (see for instance [13]). If $f: Y \rightarrow [0, 3]$ is a self-indexing Morse function on Y with one minimum and one maximum, then f induces a Heegaard decomposition with surface $\Sigma = f^{-1}(3/2)$, $U_0 = f^{-1}[0, 3/2]$, $U_1 = f^{-1}[3/2, 3]$. The attaching circles α and β are the intersections of Σ with the ascending and descending manifolds for the index one and two critical points respectively (with respect to some choice of Riemannian metric over Y). We will call such a Morse function on Y *compatible* with the Heegaard diagram (Σ, α, β) .

Definition 2.1. Let (Σ, α, β) and $(\Sigma', \alpha', \beta')$ be a pair of Heegaard diagrams. We say that the Heegaard diagrams are *isotopic* if $\Sigma = \Sigma'$ and there are two one-parameter families α_t and β_t of g -tuples of curves, moving by isotopies so that for each t , both the α_t and the β_t are g -tuples of smoothly embedded, pairwise disjoint curves. We say that $(\Sigma', \alpha', \beta')$ is obtained from (Σ, α, β) by *handleslides* if $\Sigma = \Sigma'$ and α' are obtained by handleslides amongst the α , and β' is obtained by handleslides amongst the β . Finally, we say that $(\Sigma', \alpha', \beta')$ is obtained from (Σ, α, β) by *stabilization*, if $\Sigma' \cong \Sigma \# E$, and $\alpha' = \{\alpha_1, \dots, \alpha_g, \alpha_{g+1}\}$, $\beta' = \{\beta_1, \dots, \beta_g, \beta_{g+1}\}$, where E is a two-torus, and $\alpha_{g+1}, \beta_{g+1}$ are a pair of curves in E which meet transversally in a single point. Conversely, in this case, we say that (Σ, α, β) is obtained from $(\Sigma', \alpha', \beta')$ by *destabilization*. Collectively, we will call isotopies, handleslides, stabilizations, and destabilizations of Heegaard diagrams *Heegaard moves*.

Recall the following basic result (compare [31] and [35]):

PROPOSITION 2.2. *Any two Heegaard diagrams (Σ, α, β) and $(\Sigma', \alpha', \beta')$ which specify the same three-manifold are diffeomorphic after a finite sequence of Heegaard moves.*

For the above statement, two Heegaard diagrams (Σ, α, β) and $(\Sigma', \alpha', \beta')$ are said to be diffeomorphic if there is an orientation-preserving diffeomorphism of Σ to Σ' which carries α to α' and β to β' .

Most of Proposition 2.2 follows from the usual handle calculus (as described, for example, in [13]). Introducing a canceling pair of index one and two critical points increases the genus of the Heegaard surface by one. After possible isotopies and handleslides, this corresponds to the stablization procedure described above. *A priori*, we might have to introduce canceling pairs of critical points with indices as zero and one, or two and three. (The two and three case is dual to the index zero and one case, so that we can consider only the latter.) To consider new index zero critical points, we have to relax the notion of attaching circles: any set $\{\alpha_1, \dots, \alpha_d\}$ of pairwise disjoint, embedded circles in Σ which bound disjoint, embedded disks in U and span the image of the boundary homomorphism $\partial: H_2(U, \Sigma; \mathbb{Z}) \rightarrow H_1(\Sigma, \mathbb{Z})$ is called an *extended set of attaching circles for U* (i.e., here we have $d \geq g$). Introducing a canceling zero and one pair corresponds to preserving Σ but introducing a new attaching circle (which cancels with the index zero critical point). Pair cancellations correspond to deleting an attaching circle which can be homologically expressed in terms of the other attaching circles. Proposition 2.2 is established once we see that handleslides using these additional attaching circles can be expressed in terms of handleslides amongst a minimal set of attaching circles. To this end, we have the following lemmas:

LEMMA 2.3. *Let $\{\alpha_1, \dots, \alpha_g\}$ be a set of attaching circles in Σ for U . Suppose that γ is a simple, closed curve which is disjoint from $\{\alpha_1, \dots, \alpha_g\}$. Then, either γ is null-homologous or there is some α_i with the property that γ is isotopic to a curve obtained by handlesliding α_i across some collection of the α_j for $j \neq i$.*

Proof. If we surger out the $\alpha_1, \dots, \alpha_g$, we replace Σ by the two-sphere S^2 , with $2g$ marked points $\{p_1, q_1, \dots, p_g, q_g\}$ (i.e. the pair $\{p_i, q_i\}$ corresponds to the zero-sphere which replaced the circle α_i in Σ). Now, γ induces a Jordan curve γ' in this two-sphere. If γ' does not separate any of the p_i from the corresponding q_i , then it is easy to see that the original curve γ had to be null-homologous. On the other hand, if p_i is separated from q_i , then it is easy to see that γ is obtained by handlesliding α_i across some collection of the α_j for $j \neq i$. \square

LEMMA 2.4. *Let $\{\alpha_1, \dots, \alpha_d\}$ be an extended set of attaching circles in Σ for U . Then, any two g -tuples of these circles which form a set of attaching circles for U are related by a series of isotopies and handleslides.*

Proof. This is proved by induction on g . The case $g = 1$ is obvious: if two embedded curves in the torus represent the same generator in homology, they are isotopic.

Next, if the two subsets have some element, say α_1 , in common, then we can reduce the genus, by surgering out α_1 . This gives a new Riemann surface Σ' of genus $g - 1$ with two marked points. Each isotopy of a curve in Σ' which crosses one of the marked points corresponds to a handleslide in Σ across α_1 . Thus, by the inductive hypothesis, the two subsets are related by isotopies and handleslides.

Consider then the case where the two subsets are disjoint, labeled $\{\alpha_1, \dots, \alpha_g\}$ and $\{\alpha'_1, \dots, \alpha'_g\}$. Obviously, α'_1 is not null-homologous, so, according to Lemma 2.3, after renumbering, we can obtain α'_1 by handlesliding α_1 across some collection of the α_i ($i = 2, \dots, g$). Thus, we have reduced to the case where the two subsets are not disjoint. \square

Proof of Proposition 2.2. Given any two Heegaard diagrams of Y , we connect corresponding compatible Morse functions through a generic family f_t of functions, and equip Y with a generic metric. The genericity ensures that the gradient flow-lines for each of the f_t never flow from higher- to lower-index critical points. In particular, at all but finitely many t (where there is cancellation of index one and two critical points), we get induced Heegaard diagrams for Y , whose extended sets of attaching circles undergo only handleslides and pair creations and cancellations.

Suppose, now that two sets of attaching circles $\{\alpha_1, \dots, \alpha_g\}$ and $\{\alpha'_1, \dots, \alpha'_g\}$ for U can be extended to sets of attaching circles $\{\alpha_1, \dots, \alpha_d\}$ and $\{\alpha'_1, \dots, \alpha'_d\}$ for U , which are related by isotopies and handleslides. We claim that the original sets $\{\alpha_1, \dots, \alpha_g\}$ and $\{\alpha'_1, \dots, \alpha'_g\}$ are related by isotopies and handleslides, as well. To see this, suppose that α'_i (for some fixed $i \in \{1, \dots, d\}$) is obtained by handle-sliding α_i over some α_j (for $j = 1, \dots, d$), then since α'_i can be made disjoint from all the other α -curves, we can view the extended subset $\{\alpha_1, \dots, \alpha_d, \alpha'_i\}$ as a set of attaching circles for U . Thus, Lemma 2.4 applies, proving the claim for a single handleslide amongst the $\{\alpha_1, \dots, \alpha_d\}$, and hence also for arbitrary many handleslides. The proposition then follows. \square

In light of Proposition 2.2, we see that any quantity associated to Heegaard diagrams which is unchanged by isotopies, handleslides, and stabilization is actually a topological invariant of the underlying three-manifold. Indeed, we will need a slight refinement of Proposition 2.2. To this end, we will find it convenient to fix an additional reference point $z \in \Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$.

Definition 2.5. The collection $(\Sigma, \alpha, \beta, z)$ is called a *pointed Heegaard diagram*. Heegaard moves which are supported in a complement of z — i.e. during

the isotopies, the curves never cross the basepoint z , and for handleslides, the pair of pants does not contain z — are called *pointed Heegaard moves*.

2.2. Symmetric products. In this section, we review the topology of symmetric products. For more details, see [22].

The *diagonal* D in $\text{Sym}^g(\Sigma)$ consists of those g -tuples of points in Σ , where at least two entries coincide.

LEMMA 2.6. *Let Σ be a genus g surface. Then*

$$\pi_1(\text{Sym}^g(\Sigma)) \cong H_1(\text{Sym}^g(\Sigma)) \cong H_1(\Sigma).$$

Proof. We begin by proving the isomorphism on the level of homology. There is an obvious map

$$H_1(\Sigma) \rightarrow H_1(\text{Sym}^g(\Sigma))$$

induced from the inclusion $\Sigma \times \{x\} \times \dots \times \{x\} \subset \text{Sym}^g(\Sigma)$. To invert this, note that a curve (in general position) in $\text{Sym}^g(\Sigma)$ corresponds to a map of a g -fold cover of S^1 to Σ , giving us a homology class in $H_1(\Sigma)$. This gives a well-defined map $H_1(\text{Sym}^g(\Sigma)) \rightarrow H_1(\Sigma)$, since a cobordism Z in $\text{Sym}^g(\Sigma)$, which meets the diagonal transversally gives rise to a branched cover \tilde{Z} which maps to Σ . It is easy to see that these two maps are inverses of each other.

To see that $\pi_1(\text{Sym}^g(\Sigma))$ is Abelian, consider a null-homologous curve $\gamma: S^1 \rightarrow \text{Sym}^g(\Sigma)$, which misses the diagonal. As above, this corresponds to a map $\hat{\gamma}$ of a g -fold cover of the circle into Σ , which is null-homologous; i.e. there is a map of a two-manifold-with-boundary F into Σ , $i: F \rightarrow \Sigma$, with $i|\partial F = \hat{\gamma}$. By increasing the genus of F if necessary, we can extend the g -fold covering of the circle to a branched g -fold covering of the disk $\pi: F \rightarrow D$. Then, the map sending $z \in D$ to the image of $\pi^{-1}(z)$ under i induces the requisite null-homotopy of γ . \square

The isomorphism above is Poincaré dual to the one induced from the Abel-Jacobi map

$$\Theta: \text{Sym}^g(\Sigma) \rightarrow \text{Pic}^g(\Sigma)$$

which associates to each divisor the corresponding (isomorphism class of) line bundle. Here, $\text{Pic}^g(\Sigma)$ is the set of isomorphism classes of degree g line bundles over Σ , which in turn is isomorphic to the torus

$$\frac{H^1(\Sigma, \mathbb{R})}{H^1(\Sigma, \mathbb{Z})} \cong T^{2g}.$$

Since, $H_1(\text{Pic}^g(\Sigma)) = H^1(\Sigma, \mathbb{Z})$, we obtain an isomorphism

$$\mu: H_1(\Sigma; \mathbb{Z}) \rightarrow H^1(\text{Sym}^g(\Sigma); \mathbb{Z}).$$

The cohomology of $\mathrm{Sym}^g(\Sigma)$ was studied in [22]. It is proved there that the cohomology ring is generated by the image of the above map μ , and one additional two-dimensional cohomology class, which we denote by U , which is Poincaré dual to the submanifold

$$\{x\} \times \mathrm{Sym}^{g-1}(\Sigma) \subset \mathrm{Sym}^g(\Sigma),$$

where x is any fixed point in Σ .

As is implicit in the above discussion, a holomorphic structure j on Σ naturally endows the symmetric product $\mathrm{Sym}^g(\Sigma)$ with a holomorphic structure, denoted $\mathrm{Sym}^g(j)$. This structure $\mathrm{Sym}^g(j)$ is specified by the property that the natural quotient map

$$\pi: \Sigma^{\times g} \longrightarrow \mathrm{Sym}^g(\Sigma)$$

is holomorphic (where the product space is endowed with a product holomorphic structure). Indeed, this complex structure can be Kähler: any Riemann surface has a projective embedding, inducing naturally a projective embedding on the g -fold product $\Sigma^{\times g}$, so that elementary geometric invariant theory (as explained in Chapter 10 of [15]) endows $\mathrm{Sym}^g(\Sigma)$, its quotient by the symmetric group on g letters (a finite group acting holomorphically), with the structure of a projective algebraic variety.

As is usual in the study of Gromov invariants and Lagrangian Floer theory, we must understand the holomorphic spheres in our manifold $\mathrm{Sym}^g(\Sigma)$. To this end, we study how the first Chern class c_1 (of the tangent bundle $T\mathrm{Sym}^g(\Sigma)$) evaluates on homology classes which are representable by spheres. First, we identify these homology classes. To this end, we introduce a little notation. If X is a connected space endowed with a basepoint $x \in X$, let $\pi'_2(X)$ denote the quotient of $\pi_2(X, x)$ by the action of $\pi_1(X, x)$. Note that this group is independent of the choice of basepoint x , and also that the natural Hurewicz homomorphism from $\pi_2(X, x)$ to $H_2(X; \mathbb{Z})$ factors through $\pi'_2(X)$.

PROPOSITION 2.7. *Let Σ be a Riemann surface of genus $g > 1$, then*

$$\pi'_2(\mathrm{Sym}^g(\Sigma)) \cong \mathbb{Z}.$$

Furthermore, if $\{A_i, B_i\}$ is a symplectic basis for $H_1(\Sigma)$, then there is a generator of $\pi'_2(\mathrm{Sym}^g(\Sigma))$, denoted S , whose image under the Hurewicz homomorphism is Poincaré dual to

$$(1 - g)U^{g-1} + \sum_{i=1}^g \mu(A_i)\mu(B_i)U^{g-2}.$$

In the case where $g > 2$, $\pi_1(\mathrm{Sym}^g(\Sigma))$ acts trivially on $\pi_2(\mathrm{Sym}^g(\Sigma))$ and thus

$$\pi_2(\mathrm{Sym}^g(\Sigma)) \cong \mathbb{Z}.$$

Proof. The isomorphism $\pi'_2(\text{Sym}^g(\Sigma)) \cong \mathbb{Z}$ is given by the intersection number with the submanifold $x \times \text{Sym}^{g-1}(\Sigma)$, for generic x . Specifically, if we take a hyperelliptic structure on Σ , the hyperelliptic involution gives rise to a sphere $S_0 \subset \text{Sym}^2(\Sigma)$, which we can then use to construct a sphere $S = S_0 \times x_3 \times \dots \times x_g \subset \text{Sym}^g(\Sigma)$. Clearly, S maps to 1 under this isomorphism.

Consider a sphere Z in the kernel of this map. By moving Z into general position, we can arrange that Z meets $x \times \text{Sym}^{g-1}(\Sigma)$ transversally in finitely many points. By splicing in homotopic translates of S (with appropriate signs) at the intersection points, we can find a new sphere Z' homotopic to Z which misses $x \times \text{Sym}^{g-1}(\Sigma)$; i.e. we can think of Z' as a sphere in $\text{Sym}^g(\Sigma - x)$. Note that since this splicing construction makes no reference to a basepoint, this operation is taking place in $\pi'_2(\text{Sym}^g(\Sigma))$. We claim that $\pi_2(\text{Sym}^g(\Sigma - x)) = 0$, for $g > 2$.

One way to see that $\pi_2(\text{Sym}^g(\Sigma - x)) = 0$ is to observe that $\Sigma - x$ is homotopy equivalent to the wedge of $2g$ circles or, equivalently, the complement in \mathbb{C} of $2g$ points $\{z_1, \dots, z_{2g}\}$. Now, $\text{Sym}^g(\mathbb{C} - \{z_1, \dots, z_{2g}\})$ can be thought of as the space of monic degree g polynomials p in one variable, with $p(z_i) \neq 0$ for $i = 1, \dots, 2g$. When we consider the coefficients of p , this is equivalent to considering \mathbb{C}^g minus $2g$ generic hyperplanes. A theorem of Hattori [17] states that the homology groups of the universal covering space of this complement are trivial except in dimension zero or g . This proves that $\pi_2(\text{Sym}^g(\Sigma - x)) = 0$ and hence completes the proof that $\pi'_2(\text{Sym}^g(\Sigma)) = \mathbb{Z}$ for $g > 2$.

In the case where $g = 2$ it is easy to see that $\text{Sym}^2(\Sigma)$ is diffeomorphic to the blowup of T^4 (indeed, the Abel-Jacobi map gives the map to the torus, and the exceptional sphere is the sphere $S_0 \subset \text{Sym}^2(\Sigma)$ induced from the hyperelliptic involution on the genus two Riemann surface). In this case, the calculation of π'_2 is straightforward.

To verify the second claim, note that the Poincaré dual of S is characterized by the fact that:

$$\text{PD}[S] \cup U = \text{PD}[1] \text{ and } \text{PD}[S] \cup \mu(A_i) \cup \mu(B_j) = 0,$$

where the latter equation holds for all $i, j = 1, \dots, g$. It is easy to see that $(1 - g)U^{g-1} + \sum_{i=1}^g \mu(A_i)\mu(B_i)U^{g-2}$ satisfies these properties, as claimed.

Finally, in the case where $g > 2$, we verify that the action of $\pi_1(\text{Sym}^g(\Sigma), \{x \times \dots \times x\})$ is trivial. Fix maps

$$\gamma: S^1 \longrightarrow \text{Sym}^g(\Sigma)$$

and

$$\sigma: S^2 \longrightarrow \text{Sym}^g(\Sigma).$$

According to Lemma 2.6, we can arrange after a homotopy that $\gamma = \gamma_1 \times \{x, \dots, x\}$ where $\gamma_1: S^1 \longrightarrow \Sigma$, and according to our calculation of π'_2 , we can

arrange that σ has the form $\sigma = \{x\} \times \sigma_1$, where $\sigma_1: S^2 \longrightarrow \text{Sym}^{g-1}(\Sigma)$. Now, the map

$$\gamma \vee \sigma: S^1 \vee S^2 \longrightarrow \text{Sym}^g(\Sigma)$$

admits an obvious extension

$$\gamma_1 \times \sigma_1: S^1 \times S^2 \longrightarrow \text{Sym}^g(\Sigma).$$

Since the action of $\pi_1(S^1 \times S^2)$ on $\pi_2(S^1 \times S^2)$ is trivial, the claim now follows immediately. \square

The evaluation of the first Chern class on the generator S is given in the following:

LEMMA 2.8. *The first Chern class of $\text{Sym}^g(\Sigma_g)$ is given by*

$$c_1 = U - \sum_{i=1}^g \mu(A_i)\mu(B_i).$$

In particular, $\langle c_1, [S] \rangle = 1$.

Proof. See [22] for the calculation of c_1 . The rest follows from this, together with Proposition 2.7. \square

2.3. *Totally real tori.* Fix a Heegaard diagram (Σ, α, β) . There is a naturally induced pair of smoothly embedded, g -dimensional tori

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

in $\text{Sym}^g(\Sigma)$. More precisely \mathbb{T}_α consists of those g -tuples of points $\{x_1, \dots, x_g\}$ for which $x_i \in \alpha_i$ for $i = 1, \dots, g$.

These tori enjoy a certain compatibility with any complex structure on $\text{Sym}^g(\Sigma)$ induced (as in Section 2.2) from Σ .

Definition 2.9. Let (Z, J) be a complex manifold, and $L \subset Z$ be a submanifold. Then, L is called *totally real* if none of its tangent spaces contains a J -complex line, i.e. $T_\lambda L \cap JT_\lambda L = (0)$ for each $\lambda \in L$.

LEMMA 2.10. *Let $\mathbb{T}_\alpha \subset \text{Sym}^g(\Sigma)$ be the torus induced from a set of attaching circles α . Then, \mathbb{T}_α is a totally real submanifold of $\text{Sym}^g(\Sigma)$ (for any complex structure induced from Σ).*

Proof. Note that the projection map $\pi: \Sigma^{\times g} \longrightarrow \text{Sym}^g(\Sigma)$ is a holomorphic local diffeomorphism away from the diagonal subspaces (consisting of those g -tuples for which at least two of the points coincide). Since $\mathbb{T}_\alpha \subset \text{Sym}^g(\Sigma)$ misses the diagonal, the claims about \mathbb{T}_α follow immediately from the fact that $\alpha_1 \times \cdots \times \alpha_g \subset \Sigma^{\times g}$ is a totally real submanifold (for the product complex structure), which follows easily from the definitions. \square

Note also that if all the α_i curves meet all the β_j curves transversally, then the tori \mathbb{T}_α and \mathbb{T}_β meet transversally. We will make these transversality assumptions as needed.

2.4. Intersection points and disks. Let $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ be a pair of intersection points. Choose a pair of paths $a: [0, 1] \rightarrow \mathbb{T}_\alpha$, $b: [0, 1] \rightarrow \mathbb{T}_\beta$ from \mathbf{x} to \mathbf{y} in \mathbb{T}_α and \mathbb{T}_β respectively. The difference $a - b$, then, gives a loop in $\text{Sym}^g(\Sigma)$.

Definition 2.11. Let $\varepsilon(\mathbf{x}, \mathbf{y})$ denote the image of $a - b$ under the map

$$\frac{H_1(\text{Sym}^g(\Sigma))}{H_1(\mathbb{T}_\alpha) \oplus H_1(\mathbb{T}_\beta)} \cong \frac{H_1(\Sigma)}{[\alpha_1], \dots, [\alpha_g], [\beta_1], \dots, [\beta_g]} \cong H_1(Y; \mathbb{Z}).$$

Of course, $\varepsilon(\mathbf{x}, \mathbf{y})$ is independent of the choice of the paths a and b .

It is worth emphasizing that ε can be calculated in Σ , using the identification between $\pi_1(\text{Sym}^g(\Sigma))$ and $H_1(\Sigma)$ described in Lemma 2.6. Specifically, writing $\mathbf{x} = \{x_1, \dots, x_g\}$ and $\mathbf{y} = \{y_1, \dots, y_g\}$, we can think of the path $a: [0, 1] \rightarrow \mathbb{T}_\alpha$ as a collection of arcs in $\alpha_1 \cup \dots \cup \alpha_g \subset \Sigma$, whose boundary (thought of as a zero-chain in Σ) is given by $\partial a = y_1 + \dots + y_g - x_1 - \dots - x_g$; similarly, we think of the path $b: [0, 1] \rightarrow \mathbb{T}_\beta$ as a collection of arcs in $\beta_1 \cup \dots \cup \beta_g \subset \Sigma$, whose boundary is given by $\partial b = y_1 + \dots + y_g - x_1 - \dots - x_g$. Thus, the difference $a - b$ is a closed one-cycle in Σ , whose image in $H_1(Y; \mathbb{Z})$ is the difference $\varepsilon(\mathbf{x}, \mathbf{y})$ defined above.

Clearly ε is additive, in the sense that

$$\varepsilon(\mathbf{x}, \mathbf{y}) + \varepsilon(\mathbf{y}, \mathbf{z}) = \varepsilon(\mathbf{x}, \mathbf{z}),$$

so that ε allows us to partition the intersection points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ into equivalence classes, where $\mathbf{x} \sim \mathbf{y}$ if $\varepsilon(\mathbf{x}, \mathbf{y}) = 0$.

We will study holomorphic disk connecting \mathbf{x} and \mathbf{y} . These can be naturally partitioned into homotopy classes of disks with certain boundary conditions. To describe this, we consider the unit disk \mathbb{D} in \mathbb{C} , and let $e_1 \subset \partial\mathbb{D}$ denote the arc where $\text{Re}(z) \geq 0$, and $e_2 \subset \partial\mathbb{D}$ denote the arc where $\text{Re}(z) \leq 0$. In the case where $g > 2$, let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the space of homotopy classes of maps

$$\left\{ u: \mathbb{D} \rightarrow \text{Sym}^g(\Sigma) \left| \begin{array}{l} u(-i) = \mathbf{x}, u(i) = \mathbf{y} \\ u(e_1) \subset \mathbb{T}_\alpha, u(e_2) \subset \mathbb{T}_\beta \end{array} \right. \right\}.$$

In the case where $g = 2$, we let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the quotient of this set by the natural action of $\pi_1(\text{Sym}^g(\Sigma))$. In general, $\pi_2(\mathbf{x}, \mathbf{y})$ is empty if $\varepsilon(\mathbf{x}, \mathbf{y}) \neq 0$.

The set $\pi_2(\mathbf{x}, \mathbf{y})$ is equipped with certain algebraic structure. Note that there is a natural splicing action

$$\pi'_2(\text{Sym}^g(\Sigma)) * \pi_2(\mathbf{x}, \mathbf{y}) \rightarrow \pi_2(\mathbf{x}, \mathbf{y}).$$

Also, if we take a Whitney disk 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

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

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 2.12. Let A be a collection of functions

$$\{A_{\mathbf{x}, \mathbf{y}}: \pi_2(\mathbf{x}, \mathbf{y}) \longrightarrow \mathbb{Z}\}_{\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta},$$

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*.

For example, for each fixed basepoint $z \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$, the map which sends a Whitney disk u to the algebraic intersection number

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

descends to homotopy classes, to give an additive assignment

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

This assignment can be used to define the domain belonging to a Whitney disk:

Definition 2.13. Let $\mathcal{D}_1, \dots, \mathcal{D}_m$ denote the closures of the components of $\Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$. Given a Whitney disk $u: \mathbb{D} \longrightarrow \text{Sym}^g(\Sigma)$, 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 . If all the coefficients $n_{z_i}(u) \geq 0$, then we write $\mathcal{D}(u) \geq 0$.

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 2.14. For a pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$, 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 2.15. *Suppose $g > 1$. 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 $\varepsilon(\mathbf{x}, \mathbf{y}) \neq 0$, 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 $\pi'_2(\text{Sym}^g(\Sigma)) \times \Pi_{\mathbf{x}}(z)$ 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 (when $g > 1$).

Proof. Suppose that $g > 2$. 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 $\text{Sym}^g(\Sigma)$ joining \mathbb{T}_α to \mathbb{T}_β , 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 $\text{Sym}^g(\Sigma)$):

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

whose associated homotopy long exact sequence gives:

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

But under the identification $\pi_1(\text{Sym}^g(\Sigma)) \cong H^1(\Sigma; \mathbb{Z})$, 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 $n_z: \pi_2(\mathbf{x}, \mathbf{x}) \longrightarrow \mathbb{Z}$ provides a splitting for the sequence. The proposition in the case where $g > 2$ follows. The case where $g = 2$ follows similarly, only now one must divide by the action of $\pi_1(\text{Sym}^g(\Sigma))$.

In the case where $\mathbf{x} \neq \mathbf{y}$ and $\varepsilon(\mathbf{x}, \mathbf{y}) = 0$, then $\pi_2(\mathbf{x}, \mathbf{y})$ is nonempty, so the above reasoning applies. \square

Remark 2.16. The above result, of course, fails when $g = 1$. However, it is still clear that $\pi_2(\mathbf{x}, \mathbf{y}) \rightarrow \mathbb{Z} \oplus H^1(Y; \mathbb{Z})$ is injective, and that is the only part of this result which is required for the Floer homology constructions described below to work. (Note also that the only three-manifolds which admit genus one Heegaard diagrams are lens spaces and $S^2 \times S^1$.)

2.5. Periodic domains and surfaces in Y . Given a periodic domain \mathcal{P} , there is a map from a surface-with-boundary $\Phi: F \rightarrow \Sigma$ representing \mathcal{P} , in the sense that $\Phi_*[F] = \mathcal{P}$ as chains (where here $[F]$ 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 2.17. *Consider a chain $\mathcal{P} + \ell[\Sigma]$ with ℓ sufficiently large so that $n_{z'}(\mathcal{P} + \ell[\Sigma]) \geq 0$ for all $z' \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$. Then there are an oriented two-manifold with boundary F and a map $\Phi: F \rightarrow \Sigma$ 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 $j = 1, \dots, m(\mathcal{D}_1(a))$, we identify

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

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 $j = 1, \dots, m(\mathcal{D}_1(b))$, we identify

$$\left(x^{(j)} \in \mathcal{D}_1^{(j)}(b)\right) \sim \left(x^{(j)} \in \mathcal{D}_2^{(j)}(b)\right).$$

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 $\Phi: F \longrightarrow \Sigma$ be a representative for a periodic domain $\mathcal{P} + \ell[\Sigma]$ as constructed in Lemma 2.17. Φ can be extended to a map into the three-manifold:

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

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}_α meets \mathbb{T}_β , is Poincaré dual to the isomorphism of Proposition 2.15.

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 by requiring that the linking number of γ with α_i in U_0 is $+1$.

LEMMA 2.18. *Let $z_1, z_2 \in \Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$ 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 $\alpha_2, \dots, \alpha_g$, 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 zero to the index three 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

2.6. *Spin^c structures.* Fix a point $z \in \Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$. In this section we define a natural map

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

To construct this, it is convenient to use Turaev's formulation of Spin^c structures in terms of homology classes of vector fields (see [38]; see also [19]). Fix a Riemannian metric g over a closed, oriented three-manifold Y . Following [38], we say that two unit vector fields v_1, v_2 are said to be *homologous* if they are homotopic in the complement of a three-ball in Y (or, equivalently, in the complement of finitely many disjoint three-balls in Y). Denote the space of homology classes of unit vector fields over Y by $\text{Spin}^c(Y)$. When we fix an ortho-normal trivialization τ of the tangent bundle TY , there is a natural one-to-one correspondence between vector fields over Y and maps from Y to S^2 , which descends to homology classes (where we say that two maps $f_0, f_1: Y \rightarrow S^2$ are *homologous* if they are homotopic in the complement of a three-ball). Fixing a generator μ for $H^2(S^2; \mathbb{Z})$, it follows from elementary obstruction theory that the assignment which associates to a map from Y to S^2 the pull-back of μ induces an identification between the space of homology classes of maps from Y to S^2 and the cohomology group $H^2(Y; \mathbb{Z})$. Hence, we obtain a one-to-one correspondence, depending on the trivialization τ :

$$\delta^\tau: \text{Spin}^c(Y) \rightarrow H^2(Y; \mathbb{Z}).$$

More canonically, if v_1 and v_2 are a pair of nowhere vanishing vector fields over Y , then the difference

$$\delta(v_1, v_2) = \delta^\tau(v_1) - \delta^\tau(v_2) \in H^2(Y; \mathbb{Z})$$

is independent of the trivialization τ , since any two trivializations τ and τ' differ by the action of a map $g: Y \rightarrow \text{SO}(3)$, and, as is elementary to check,

$$\delta^{g \cdot \tau}(v) - \delta^\tau(v) = g^*(w),$$

where w is the generator of $H^2(\text{SO}(3); \mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z}$. Moreover, since (for any fixed $v \in \text{Spin}^c(Y)$) the map $\delta(v, \cdot)$ defines a one-to-one correspondence between $\text{Spin}^c(Y)$ and $H^2(Y; \mathbb{Z})$, and $\delta(v_1, v_2) + \delta(v_2, v_3) = \delta(v_1, v_3)$, the space $\text{Spin}^c(Y)$ is naturally an affine space for $H^2(Y; \mathbb{Z})$. It is convenient to write the action additively, so that if $a \in H^2(Y; \mathbb{Z})$ and $v \in \text{Spin}^c(Y)$, then $a + v \in \text{Spin}^c(Y)$ is characterized by the property that $\delta(a + v, v) = a$. Moreover, given $v_1, v_2 \in \text{Spin}^c(Y)$, we let $v_1 - v_2$ denote $\delta(v_1, v_2)$.

Thus, one could simply define the space of Spin^c structures over Y to be the space of homology classes of vector fields. The correspondence with the more traditional definition of Spin^c structures is given by associating to the vector v the “canonical” Spin^c structure associated to the reduction of the structure group of TY to $\text{SO}(2)$ (for this, and other equivalent formulations, see [38]).

The natural map s_z is defined as follows. Let f be a Morse function on Y compatible with the attaching circles α, β ; see Section 2.1. Then each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ determines a g -tuple of trajectories for the gradient flow of f

connecting the index one critical points to index two critical points. Similarly z gives a trajectory connecting the index zero critical point with the index three critical point. Deleting tubular neighborhoods of these $g+1$ trajectories, we obtain a subset of Y where the gradient vector field $\vec{\nabla}f$ does not vanish. Since each trajectory connects critical points of different parities, the gradient vector field has index 0 on all the boundary spheres of the subset, so it can be extended as a nowhere vanishing vector field over Y . The homology class of the nowhere vanishing vector field obtained in this manner (after renormalizing, to make it a unit vector field) gives the Spin^c structure $s_z(\mathbf{x})$. Clearly $s_z(\mathbf{x})$ does not depend on the choice of the compatible Morse function f or the extension of the vector field $\vec{\nabla}f$ to the balls.

Now we investigate how $s_z(\mathbf{x}) \in \text{Spin}^c(Y)$ depends on \mathbf{x} and z .

LEMMA 2.19. *When $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$,*

$$(1) \quad s_z(\mathbf{y}) - s_z(\mathbf{x}) = \text{PD}[\varepsilon(\mathbf{x}, \mathbf{y})].$$

Furthermore if $z_1, z_2 \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$ can be connected in Σ by an arc z_t from z_1 to z_2 which is disjoint from the β , whose intersection number $\#(\alpha_i \cap z_t) = 1$, and $\#(\alpha_j \cap z_t)$ for $j \neq i$ vanishes, then for all $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$,

$$(2) \quad s_{z_2}(\mathbf{x}) - s_{z_1}(\mathbf{x}) = \alpha_i^*,$$

where $\alpha_i^ \in H^2(Y, \mathbb{Z})$ is Poincaré dual to the homology class in Y induced from a curve γ in Σ with $\alpha_i \cdot \gamma = 1$, and whose intersection number with all other α_j for $j \neq i$ vanishes.*

Proof. Given $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, let $\gamma_{\mathbf{x}}$ denote the g trajectories for $\vec{\nabla}f$ connecting the index one to the index two critical points which contain the g -tuple \mathbf{x} ; similarly, given $z \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$, let γ_z denote the corresponding trajectory from the index zero to the index three critical point.

Thus, if $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}}$ is a closed loop in Y . A representative for $s_z(\mathbf{x})$ is obtained by modifying the vector field $\vec{\nabla}f$ in a neighborhood of $\gamma_{\mathbf{x}} \cup \gamma_z$. It follows then that $s_z(\mathbf{x}) - s_z(\mathbf{y})$ can be represented by a cohomology class which is compactly supported in a neighborhood of $\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}}$ (we can use the same vector field to represent both Spin^c structures outside this neighborhood).

It follows that the difference $s_z(\mathbf{x}) - s_z(\mathbf{y})$ is some multiple of the Poincaré dual of $\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}}$ (at least if the curve is connected; though the following argument is easily seen to apply in the disconnected case as well). To find out which multiple, we fix a disk D_0 transverse to $\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}}$; to find such a disk take some $x_i \in \mathbf{x}$ so that $x_i \notin \mathbf{y}$ (if no such x_i can be found, then $\mathbf{x} = \mathbf{y}$, and Equation (1) is trivial), and let D_0 be a small neighborhood of x_i in Σ . Our representative $v_{\mathbf{x}}$ of $s_z(\mathbf{x})$ can be chosen to agree with $\vec{\nabla}f$ near ∂D_0 ; and the representative $v_{\mathbf{y}}$ for $s_z(\mathbf{y})$ can be chosen to agree with $\vec{\nabla}f$ over D_0 . With respect to any fixed trivialization of TY , the two maps from Y to S^2 corresponding to $v_{\mathbf{x}}$ and

$v_{\mathbf{y}}$ agree on ∂D_0 . It makes sense, then, to compare the difference between the degrees $\deg_{D_0}(v_{\mathbf{x}})$ and $\deg_{D_0}(v_{\mathbf{y}})$ (maps from the disk to the sphere, relative to their boundary). Indeed,

$$s_z(\mathbf{x}) - s_z(\mathbf{y}) = (\deg_{D_0}(v_{\mathbf{x}}) - \deg_{D_0}(v_{\mathbf{y}})) \text{PD}(\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}}).$$

To calculate this difference, take another disk D_1 with the same boundary as D_0 , so that $D_0 \cup D_1$ bounds a three-ball in Y containing the index one critical point corresponding to x_i (and no other critical point); thus we can assume that $v_{\mathbf{x}} \equiv \vec{\nabla} f$ over D_1 . Now, since $v_{\mathbf{x}}$ does not vanish inside this three-ball, we have:

$$0 = \deg_{D_0}(v_{\mathbf{x}}) + \deg_{D_1}(v_{\mathbf{x}}) = \deg_{D_0}(v_{\mathbf{x}}) + \deg_{D_1}(\vec{\nabla} f).$$

Thus,

$$\deg_{D_0}(v_{\mathbf{x}}) - \deg_{D_0}(v_{\mathbf{y}}) = -\deg_{D_1}(\vec{\nabla} f) - \deg_{D_0}(\vec{\nabla} f) = 1,$$

since $\vec{\nabla} f$ vanishes with winding number -1 around the index 1 critical points of f . It follows from this calculation that $v_{\mathbf{x}} - v_{\mathbf{y}} = \text{PD}(\gamma_{\mathbf{x}} - \gamma_{\mathbf{y}})$. Letting $a \subset \alpha_1 \cup \dots \cup \alpha_g$ be a collection of arcs with $\partial a = \mathbf{y} - \mathbf{x}$, and $b \subset \beta_1 \cup \dots \cup \beta_g$ be such a collection with $\partial b = \mathbf{y} - \mathbf{x}$, we know that $a - b$ represents $\varepsilon(\mathbf{x}, \mathbf{y})$. On the other hand, if $a_i \subset a$ is one of the arcs which connects x_i to y_i , then it is easy to see that a_i is homotopic relative to its boundary to the segment in U_0 formed by joining the two gradient trajectories connecting x_i and y_i to the index one critical point. It follows from this (and the analogous statement in U_1) that $a - b$ is homologous to $\gamma_{\mathbf{y}} - \gamma_{\mathbf{x}}$. Equation (1) follows.

Equation (2) follows from similar considerations. Note first that $s_{z_1}(\mathbf{x})$ agrees with $s_{z_2}(\mathbf{y})$ away from $\gamma_{z_1} - \gamma_{z_2}$. Letting now D_0 be a disk which meets γ_{z_1} transversally in a single positive point (and is disjoint from γ_{z_2}), and D_1 be a disk with the same boundary as D_0 so that $D_0 \cup D_1$ contains the index zero critical point, we have that

$$\begin{aligned} \deg_{D_0}(v_{z_1}) - \deg_{D_0}(v_{z_2}) &= -\deg_{D_1}(v_{z_1}) - \deg_{D_0}(v_{z_2}) \\ &= -\deg_{D_1}(\vec{\nabla} f) - \deg_{D_0}(\vec{\nabla} f) = -1 \end{aligned}$$

(note now that $\vec{\nabla} f$ vanishes with winding number $+1$ around the index zero critical point). It follows that $s_{z_1}(\mathbf{x}) - s_{z_2}(\mathbf{x}) = -\text{PD}(\gamma_{z_1} - \gamma_{z_2})$. Now, $\gamma_{z_1} - \gamma_{z_2}$ is easily seen to be Poincaré dual to α_i^* . \square

It is not difficult to generalize the above discussion to give a one-to-one correspondence between Spin^c structures and homotopy classes of paths of \mathbb{T}_{α} to \mathbb{T}_{β} (having fixed the base point z). This is closely related to Turaev's notion of "Euler systems" (see [38]).

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 $\mathfrak{s} \mapsto \bar{\mathfrak{s}}$. Sometimes, $\bar{\mathfrak{s}}$ is called the *conjugate* Spin^c structure to \mathfrak{s} .

There is also a natural map

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

the first Chern class. This is defined by $c_1(\mathfrak{s}) = \mathfrak{s} - \bar{\mathfrak{s}}$. 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 $c_1(\bar{\mathfrak{s}}) = -c_1(\mathfrak{s})$.

3. Analytical aspects

Lagrangian Floer homology (see [7]) is a homology theory associated to a pair L_0 and L_1 of Lagrangian submanifolds in a symplectic manifold. Its boundary map counts certain pseudo-holomorphic disks whose boundary is mapped into the union of L_0 and L_1 . Our set-up here differs slightly from Floer's: we are considering a pair of totally real submanifolds, \mathbb{T}_α and \mathbb{T}_β , in the symmetric product. It is the aim of this section to show that the essential analytical aspects — Fredholm theory, transversality, and compactness — carry over to this context. We then turn our attention to orientations. In the final subsection, we discuss certain disks, whose boundary lies entirely in either \mathbb{T}_α or \mathbb{T}_β .

3.1. Nearly symmetric almost-complex structures. We will be counting pseudo-holomorphic disks in $\text{Sym}^g(\Sigma)$, using a restricted class of almost-complex structures over $\text{Sym}^g(\Sigma)$ (which can be thought of as a suitable elaboration of the taming condition from symplectic geometry).

Recall that an almost-complex structure J over a symplectic manifold (M, ω) is said to *tame* ω if $\omega(\xi, J\xi) > 0$ for every nonzero tangent vector ξ to M . This is an open condition on J .

The quotient map

$$\pi : \Sigma^{\times g} \longrightarrow \text{Sym}^g(\Sigma)$$

induces a covering space of $\text{Sym}^g(\Sigma) - D$, where $D \subset \Sigma^{\times g}$ is the diagonal; see Subsection 2.2. Let η be a Kähler form over Σ , and $\omega_0 = \eta^{\times g}$. Clearly, ω_0 is invariant under the covering action, so it induces a Kähler form $\pi_*(\omega_0)$ over $\text{Sym}^g(\Sigma) - D$.

Definition 3.1. Fix a Kähler structure (j, η) over Σ , a finite collection of points

$$\{z_i\}_{i=1}^m \subset \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g,$$

and an open set V with

$$\left(\{z_i\}_{i=1}^m \times \text{Sym}^{g-1}(\Sigma) \bigcup D \right) \subset V \subset \text{Sym}^g(\Sigma)$$

and

$$\overline{V} \cap (\mathbb{T}_\alpha \cup \mathbb{T}_\beta) = \emptyset.$$

An almost-complex structure J on $\text{Sym}^g(\Sigma)$ is called (j, η, V) -*nearly symmetric* if

- J tames $\pi_*(\omega_0)$ over $\text{Sym}^g(\Sigma) - \overline{V}$
- J agrees with $\text{Sym}^g(j)$ over V .

The space of (j, η, V) -*nearly symmetric almost-complex structures* will be denoted $\mathcal{J}(j, \eta, V)$.

Note that since \mathbb{T}_α and \mathbb{T}_β are Lagrangian with respect to $\pi_*(\omega_0)$, and J tames $\pi_*(\omega_0)$, the tori \mathbb{T}_α and \mathbb{T}_β are totally real for J .

The space $\mathcal{J}(j, \eta, V)$ is a subset of the set of all almost-complex structures, and as such it can be endowed with Banach space topologies \mathcal{C}^ℓ for any ℓ . In fact, $\text{Sym}^g(j)$ is (j, η, V) -nearly symmetric for any choice of η and V ; and the space $\mathcal{J}(j, \eta, V)$ is an open neighborhood of $\text{Sym}^g(j)$ in the space of almost-complex structures which agree with $\text{Sym}^g(j)$ over V .

Unless otherwise specified, we choose the points $\{z_i\}_{i=1}^m$ so that there is some z_i in each connected component of $\Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$.

3.2. Fredholm theory. We recall the Fredholm theory for pseudo-holomorphic disks, with appropriate boundary conditions. For more details, we refer the reader to [9]; see also [26], [11], and [12].

To set this up we assume that \mathbb{T}_α and \mathbb{T}_β meet transversally, i.e. that each α_i meets each β_j transversally.

We consider the moduli space of holomorphic strips connecting \mathbf{x} to \mathbf{y} , suitably generalized as follows. Let $\mathbb{D} = [0, 1] \times i\mathbb{R} \subset \mathbb{C}$ be the strip in the complex plane. Fix a path J_s of almost-complex structures over $\text{Sym}^g(\Sigma)$. Let $\mathcal{M}_{J_s}(\mathbf{x}, \mathbf{y})$ be the set of maps satisfying the following conditions:

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

For $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, the space $\mathcal{M}_{J_s}(\phi)$ denotes the subset consisting of maps as above which represent the given homotopy class ϕ (or equivalence class, when $g = 2$). The translation action on \mathbb{D} endows this moduli space with an \mathbb{R} action. The space of *unparametrized 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).

We will be considering moduli space $\mathcal{M}_{J_s}(\mathbf{x}, \mathbf{y})$, where J_s is a one-parameter family of nearly symmetric almost-complex structures: i.e. where we have some fixed (j, η, V) for which each J_s is (j, η, V) -nearly symmetric (see Definition 3.1) for each $s \in [0, 1]$.

In the definition of nearly-symmetric almost-complex structure, the almost-complex structure in a neighborhood of D is fixed to help prove the required energy bound, cf. Subsection 3.4. Moreover, the complex structure in a neighborhood of the $\{z_i\}_{i=1}^m \times \text{Sym}^{g-1}(\Sigma)$ is fixed to establish the following:

LEMMA 3.2. *If $u \in \mathcal{M}_{J_s}(\phi)$ is any J_s -holomorphic disk, then $\mathcal{D}(u) \geq 0$.*

Proof. In a neighborhood of $\{z_i\}_{i=1}^m \times \text{Sym}^{g-1}(\Sigma)$, we are using an integrable complex structure, so the disk u must either be contained in the subvariety (which is excluded by the boundary conditions) or it must meet it nonnegatively. \square

Let E be a vector bundle over $[0, 1] \times \mathbb{R}$ equipped with a metric and compatible connection, p, δ be positive real numbers, and k be a nonnegative integer. The δ -weighted Sobolev space of sections of E , written $L_{k,\delta}^p([0, 1] \times \mathbb{R}, E)$, is the space of sections σ for which the norm

$$\|\sigma\|_{L_{k,\delta}^p(E)} = \sum_{\ell=0}^k \int_{[0,1] \times \mathbb{R}} |\nabla^{(\ell)} \sigma(s+it)|^p e^{\delta\tau(t)} ds \wedge dt$$

is finite. Here, $\tau: \mathbb{R} \rightarrow \mathbb{R}$ is a smooth function with $\tau(t) = |t|$ provided that $|t| \geq 1$.

Fix some $p > 2$. Let $\mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$ denote the space of maps

$$u: [0, 1] \times \mathbb{R} \rightarrow \text{Sym}^g(\Sigma)$$

in $L_{1,\text{loc}}^p$, satisfying the boundary conditions

$$u(\{1\} \times \mathbb{R}) \subset \mathbb{T}_\alpha, \text{ and } u(\{0\} \times \mathbb{R}) \subset \mathbb{T}_\beta,$$

which are asymptotic to \mathbf{x} and \mathbf{y} as $t \mapsto -\infty$ and $+\infty$, in the following sense. There is a real number $T > 0$ and sections

$$\xi_- \in L_{1,\delta}^p\left([0, 1] \times (-\infty, -T], T_{\mathbf{x}}\text{Sym}^g(\Sigma)\right)$$

and

$$\xi_+ \in L_{1,\delta}^p\left([0, 1] \times [T, \infty), T_{\mathbf{y}}\text{Sym}^g(\Sigma)\right)$$

with the property that

$$u(s+it) = \exp_{\mathbf{x}}(\xi_-(s+it)) \text{ and } u(s+it) = \exp_{\mathbf{y}}(\xi_+(s+it)),$$

for all $t < -T$ and $t > T$ respectively. Here, \exp denotes the usual exponential map for some Riemannian metric on $\text{Sym}^g(\Sigma)$. Note that $\mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$ can be naturally given the structure of a Banach manifold, whose tangent space at any $u \in \mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$ is given by

$$L_{1,\delta}^p(u) := \left\{ \xi \in L_{1,\delta}^p([0, 1] \times \mathbb{R}, u^*(T\text{Sym}^g(\Sigma))) \mid \begin{array}{l} \xi(1, t) \in T_{u(1+it)}(\mathbb{T}_\alpha), \forall t \in \mathbb{R} \\ \xi(0, t) \in T_{u(0+it)}(\mathbb{T}_\beta), \forall t \in \mathbb{R} \end{array} \right\}.$$

Moreover, at each $u \in \mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$, we denote the space of sections

$$L_\delta^p(\Lambda^{0,1}u) := L_\delta^p([0, 1] \times \mathbb{R}, u^*(T\text{Sym}^g(\Sigma))).$$

These Banach spaces fit together to form a bundle \mathcal{L}_δ^p over $\mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$. At each $u \in \mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$, $\bar{\partial}_{J_s}u = \frac{d}{ds} + J(s)\frac{d}{dt}$ lies in the space $L_\delta^p(\Lambda^{0,1}(u))$ and is zero exactly when u is a J_s -holomorphic map. (Note that our definition of $\bar{\partial}_{J_s}$ implicitly uses the natural trivialization of the the bundle $\Lambda^{0,1}$ over \mathbb{D} , which is why the bundle does not appear in the definition of $L_\delta^p(\Lambda^{0,1}u)$, but does appear in its notation.) This assignment fits together over $\mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$ to induce a Fredholm section of \mathcal{L}_δ^p . The linearization of this section is denoted

$$D_u : L_{1,\delta}^p(u) \longrightarrow L_\delta^p(\Lambda^{0,1}u),$$

and it is given by the formula

$$D_u(\nu) = \frac{d\nu}{ds} + J(s)\frac{d\nu}{dt} + (\nabla_\nu J(s))\frac{du}{dt}.$$

Since the intersection of \mathbb{T}_α and \mathbb{T}_β is transverse, this linear map is Fredholm for all sufficiently small nonnegative δ . Indeed, there is some $\delta_0 > 0$ with the property that any map $u \in \mathcal{M}(\mathbf{x}, \mathbf{y})$ lies in $\mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$, for all $0 \leq \delta < \delta_0$.

The components of $\mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$ can be partitioned according to homotopy classes $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. The index of D_u , acting on the unweighted space ($\delta = 0$) descends to a function on $\pi_2(\mathbf{x}, \mathbf{y})$. Indeed, the index is calculated by the Maslov index μ of the map u (see [8], [33], [32], [39]). We conclude the subsection with a result about the Maslov index which will be of relevance to us later:

LEMMA 3.3. *Let $S \in \pi_2'(\text{Sym}^g(\Sigma))$ be the positive generator. Then for any $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$,*

$$\mu(\phi + k[S]) = \mu(\phi) + 2k.$$

In particular, if $O_{\mathbf{x}} \in \pi_2(\mathbf{x}, \mathbf{x})$ denotes the class of the constant map, then

$$\mu(O_{\mathbf{x}} + kS) = 2k.$$

Proof. It follows from the excision principle for the index that attaching a topological sphere Z to a disk changes the Maslov index by $2\langle c_1, [Z] \rangle$ (see [8], [23]). On the other hand for the positive generator we have $\langle c_1, [S] \rangle = 1$ according to Lemma 2.8. \square

3.3. Transversality. Given a Heegaard diagram (Σ, α, β) for which all the α_i meet the β_j transversally, the tori \mathbb{T}_α and \mathbb{T}_β meet transversally, so the holomorphic disks connecting \mathbb{T}_α with \mathbb{T}_β are naturally endowed with a Fredholm deformation theory.

Indeed, the usual arguments from Floer theory (see [9], [26] and [11]) can be modified to prove the following result:

THEOREM 3.4. *Fix a Heegaard diagram (Σ, α, β) with the property that each α_i meets β_j transversally, and fix (j, η, V) as in Definition 3.1. Then, for a dense set of paths J_s of (j, η, V) -nearly-symmetric almost-complex structures, the moduli spaces $\mathcal{M}_{J_s}(\mathbf{x}, \mathbf{y})$ are all smoothly cut out by the defining equations.*

In the above statement, “dense” is meant in the C^∞ topology on the path-space of $\mathcal{J}(j, \eta, V)$.

Proof of Theorem 3.4. This is a modification of the usual proof of transversality, see [9], [26] and [11].

Recall (see for instance Theorem 5.1 of [26]) that if u is any nonconstant holomorphic disk, then there is a dense set of points $(s, t) \in [0, 1] \times \mathbb{R}$ satisfying the two conditions that $du_{(s,t)} \neq 0$ and $u(s, t) \cap u(s, \mathbb{R} - \{t\}) = \emptyset$. By restricting to an open neighborhood of the boundary of \mathbb{D} (note that we have assumed that \bar{V} is disjoint from \mathbb{T}_α and \mathbb{T}_β), we can find such an (s, t) with $u(s, t) \notin \bar{V}$. When we vary the path J_s in a neighborhood of $u(s, t)$, the usual arguments show that u is a smooth point for the parametrized moduli space \mathfrak{M} , consisting of pairs (J_s, u) for which $\bar{\partial}_{J_s} u = 0$. The result then follows from the Sard-Smale theorem, applied to the Fredholm projection from \mathfrak{M} to the space of paths of nearly-symmetric almost-complex structures. \square

Under certain topological hypotheses, one can achieve transversality by placing the curves α and β in general position, but leaving the almost-complex structure fixed: indeed, letting J_s be the constant path $\text{Sym}^g(j)$. We return to this in Proposition 3.9, after setting up more of the theory of holomorphic disks in $\text{Sym}^g(\Sigma)$.

3.4. Energy bounds. Let Ω be a domain in \mathbb{C} . Recall that the energy of a map $u: \Omega \rightarrow X$ to a Riemannian manifold (X, g) is given by

$$E(u) = \frac{1}{2} \int_{\Omega} |du|^2.$$

Fix $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. In order to get the usual compactness results for holomorphic disks representing ϕ , we need an *a priori* energy bound for any holomorphic representative u for ϕ .

Such a bound exists in the symplectic context. Suppose that (X, ω) is a compact symplectic manifold, with a tame almost-complex structure J , then

there is a constant C for which

$$E(u) \leq C \int_{\Omega} u^*(\omega),$$

for each J -holomorphic map u . When the u has Lagrangian boundary conditions, the integral on the right-hand side depends only on the homotopy class of the map. This principle holds in our context as well, according to the following lemma.

LEMMA 3.5. *Fix a path J_s in the space of nearly-symmetric almost-complex structures. Then, for each pair of intersection points $\mathbf{x}, \mathbf{y} \in \text{Sym}^g(\Sigma)$, and $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, there is an upper bound on the energy of any holomorphic representative of ϕ .*

Proof. Given

$$u: (\mathbb{D}, \partial\mathbb{D}) \longrightarrow (\text{Sym}^g(\Sigma), \mathbb{T}_{\alpha} \cup \mathbb{T}_{\beta}),$$

we consider the lift

$$\tilde{u}: (F, \partial F) \longrightarrow (\Sigma^{\times g}, \pi^{-1}(\mathbb{T}_{\alpha} \cup \mathbb{T}_{\beta}))$$

obtained by pulling back the branched covering space $\pi: \Sigma^{\times g} \longrightarrow \text{Sym}^g(\Sigma)$. (That is to say, F is defined to be the covering space of the image of u away from the diagonal $D \subset \text{Sym}^g(\Sigma)$, and in a neighborhood of D , F is defined as a subvariety of $\Sigma^{\times g}$ — it is here that we are using the fact that each of the J_s agree with the standard complex structure near D .)

We break the energy integral into two regions:

$$E(u) = \int_{u^{-1}(\text{Sym}^g(\Sigma) - V)} |du|^2 + \int_{u^{-1}(V)} |du|^2.$$

To estimate the integral on $\text{Sym}^g(\Sigma) - V$, we use the fact that each J_s tames $\pi_*(\omega_0)$, from which it follows that there is a constant C_1 for which

$$\begin{aligned} (3) \quad E(u|_{\text{Sym}^g(\Sigma) - V}) &\leq C_1 \int_{u^{-1}(\text{Sym}^g(\Sigma) - V)} u^*(\pi_*(\omega_0)) \\ &= \frac{C_1}{g!} \int_{\tilde{u}^{-1}(\Sigma^{\times g} - \tilde{V})} \tilde{u}^*(\omega_0), \end{aligned}$$

where $\tilde{V} = \pi^{-1}(\overline{V})$.

To estimate the other integrand, choose a Kähler form ω over $\text{Sym}^g(\Sigma)$. Over V all the J_s agree with $\text{Sym}^g(j)$, so that u is $\text{Sym}^g(j)$ -holomorphic in that region, and there is some constant C_2 with the property that

$$(4) \quad E(u|_V) \leq C_2 \int_{u^{-1}(V)} u^*(\omega)$$

(the constant C_2 depends on the Riemannian metric used over $\text{Sym}^g(\Sigma)$ and the choice of Kähler form ω). Moreover, the right-hand side can be calculated using \tilde{u} according to the following formula:

$$(5) \quad \int_{u^{-1}(V)} u^*(\omega) = \frac{1}{g!} \int_{\tilde{u}^{-1}(\tilde{V})} \tilde{u}^*(\pi^*(\omega)).$$

Now, fix any two-form ω_1 over $\Sigma^{\times g}$. Then there is a constant C_3 with the following property. Let

$$\tilde{u}: F \longrightarrow \Sigma^{\times g}$$

be any map which is $j^{\times g}$ -holomorphic on $\tilde{u}^{-1}(\tilde{V})$. Then there is the inequality

$$(6) \quad \int_{\tilde{u}^{-1}(\tilde{V})} \tilde{u}^*(\omega_1) \leq C_3 \int_{\tilde{u}^{-1}(\tilde{V})} \tilde{u}^*(\omega_0).$$

This holds for the constant with the property that for each tangent vector ξ to $\Sigma^{\times g}$,

$$\omega_1(\xi, J\xi) \leq C_3 \omega_0(\xi, J\xi),$$

where $J = j^{\times g}$. Such a constant can be found since $\Sigma^{\times g}$ is compact and $\omega_0(\cdot, J\cdot)$ determines a nondegenerate quadratic form on each tangent space $T\Sigma^{\times g}$.

Applying Inequality (6) for the form $\omega_1 = \pi^*(\omega)$, and combining this with Inequality (3), we find a constant C_0 with the property that

$$(7) \quad E(u) \leq C_0 \int_F \tilde{u}^*(\omega_0).$$

Moreover, with respect to the symplectic form ω_0 , the pre-image under π of \mathbb{T}_α and \mathbb{T}_β are both Lagrangian. This gives a topological interpretation to the right-hand side of Equation (7):

$$(8) \quad \int_F \tilde{u}^*(\omega_0) = \langle \omega_0, [F, \partial F] \rangle,$$

which makes sense since ω_0 defines a relative cohomology class in

$$H^2(\Sigma^{\times g}, \pi^{-1}(\mathbb{T}_\alpha \cup \mathbb{T}_\beta)).$$

Note that the correspondence $u \mapsto \tilde{u}$ induces a right inverse, up to a multiplicative constant, to the map on homology

$$\pi_*: H_2(\Sigma^{\times g}, \pi^{-1}(\mathbb{T}_\alpha \cup \mathbb{T}_\beta)) \longrightarrow H_2(\text{Sym}^g(\Sigma), \mathbb{T}_\alpha \cup \mathbb{T}_\beta);$$

thus, the homology class $[F, \partial F]$ depends only on the relative homology class of u , thought of as a class in $H_2(\text{Sym}^g(\Sigma), \mathbb{T}_\alpha \cup \mathbb{T}_\beta)$ — in particular, it depends only on the equivalence class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ of u .

Thus, given a class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, this gives us an *a priori* bound on the ω_0 -energy of the (branched) lift of any holomorphic disk $u \in \mathcal{M}_{J_s}(\phi)$. Combining Inequality (4), Equation (5), Inequality (7), and Equation (8), we get that

$$(9) \quad E(u) \leq C_0 \langle \omega_0, [F, \partial F] \rangle,$$

(for some constant C_0 independent of the class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$). \square

3.5. *Holomorphic disks in the symmetric product.* Suppose that the path J_s is constant, and is given by $\text{Sym}^g(\mathbf{j})$ for some complex structure \mathbf{j} over Σ . Then, the space of holomorphic disks connecting \mathbf{x}, \mathbf{y} can be given an alternate description, using only maps between one-dimensional complex manifolds.

LEMMA 3.6. *Given any holomorphic disk $u \in \mathcal{M}(\mathbf{x}, \mathbf{y})$, there is a branched g -fold covering space $p: \widehat{\mathbb{D}} \rightarrow \mathbb{D}$ and a holomorphic map $\widehat{u}: \widehat{\mathbb{D}} \rightarrow \Sigma$, with the property that for each $z \in \mathbb{D}$, $u(z)$ is the image under \widehat{u} of the pre-image $p^{-1}(z)$.*

Proof. Given a holomorphic map $u: \mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ which does not lie in the diagonal, we can find a branched $g!$ -fold cover $p: \widetilde{\mathbb{D}} \rightarrow \mathbb{D}$ pulling back the canonical $g!$ -fold cover $\pi: \Sigma^{\times g} \rightarrow \text{Sym}^g(\Sigma)$, i.e. making the following diagram commutative:

$$\begin{array}{ccc} \widetilde{\mathbb{D}} & \xrightarrow{\widetilde{u}} & \Sigma^{\times g} \\ p \downarrow & & \downarrow \pi \\ \mathbb{D} & \xrightarrow{u} & \text{Sym}^g(\Sigma). \end{array}$$

Indeed, $\widetilde{\mathbb{D}}$ inherits an action by the symmetric group on g letters S_g , and \widetilde{u} is equivariant for the action (and its quotient is u). Let $\Pi_1: \Sigma^{\times g} \rightarrow \Sigma$ denote projection onto the first factor. Then, the composite map

$$\Pi_1 \circ \widetilde{u}: \widetilde{\mathbb{D}} \rightarrow \Sigma$$

is invariant under the action of $S_{g-1} \subset S_g$ consisting of permutations which fix the first letter. Then, we let $\widehat{\mathbb{D}} = \widetilde{\mathbb{D}}/S_{g-1}$, and \widehat{u} be the induced map from $\widehat{\mathbb{D}}$ to Σ . It is easy to verify that \widehat{u} has the desired properties. \square

Remark 3.7. It is straightforward to find appropriate topological conditions on $\widehat{u}|\partial\widehat{\mathbb{D}}$ to give a one-to-one correspondence between flows in $\mathcal{M}(\mathbf{x}, \mathbf{y})$ and certain pairs $(p: \widehat{\mathbb{D}} \rightarrow \mathbb{D}, \widehat{u}: \widehat{\mathbb{D}} \rightarrow \Sigma)$.

Let $\mathcal{D}_1, \dots, \mathcal{D}_m$ denote the connected components of $\Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$. Fix a basepoint z_i inside each \mathcal{D}_i . Then for any $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ we define the *domain* associated to ϕ , as a formal linear combination of components:

$$\mathcal{D}(\phi) = \sum_{i=1}^m n_{z_i}(\phi) \cdot \mathcal{D}_i.$$

Similarly the *area* of ϕ is given by

$$\mathcal{A}(\phi) = \sum_{i=1}^m n_{z_i}(\phi) \cdot \text{Area}_\eta(\mathcal{D}_i),$$

where η is the Kähler form on Σ . This area gives us a concrete way to understand the energy bound from the previous section since, as is easy to verify,

$$\int_F \tilde{u}^*(\omega_0) = (g!) \mathcal{A}(\phi).$$

As an application of Lemma 3.6, we observe that for certain special homotopy classes of maps in $\pi_2(\mathbf{x}, \mathbf{y})$, transversality can also be achieved by moving the curves α and β , following the approach of Oh [25]. (This observation will prove helpful in the explicit calculations of Section 9 and also Section 3 of [28].)

To state it, we need the following:

Definition 3.8. A domain $\mathcal{D}(\phi)$ is called α -injective if all of its multiplicities are 0 or 1, if its interior (i.e. the interior of the region with multiplicity 1) is disjoint from each α_i for $i = 1, \dots, g$, and its boundary contains intervals in each α_i .

PROPOSITION 3.9. Let $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ be an α -injective homotopy class, and fix a complex structure j over Σ . Then, for generic perturbations of the α , the moduli space $\mathcal{M}(\phi)$ of $\text{Sym}^g(j)$ -holomorphic disks is smoothly cut out by its defining equation.

Proof. The hypotheses ensure that for all $t \in \mathbb{R}$, we have that $u(1 + it) = \{a_1, \dots, a_g\} \in \mathbb{T}_\alpha$ where $a_i \notin u(1 + it')$ for any $t' \neq t$. This is true because the α -injectivity hypothesis ensures that the corresponding map $\hat{u}: F \rightarrow \Sigma$, coming from Lemma 3.6, is injective (with injective linearization, by elementary complex analysis) on the region mapping to the α -curves $p^{-1}(\{1\} \times \mathbb{R})$. Thus, following [25], by varying the α_i in a neighborhood of the a_i , one can see that the map u is a smooth point in a parametrized moduli space (parametrized now by variations in the curves). Thus, according to the Sard-Smale theorem, for generic small variations in the α , the corresponding moduli spaces are smooth. \square

3.6. Orientability. In this subsection, we show that the moduli spaces of flows $\mathcal{M}(\mathbf{x}, \mathbf{y})$ are orientable. As is usual in the gauge-theoretic set-up, this is done by proving triviality of the determinant line bundle of the linearization of the equations (the $\bar{\partial}$ -equation) which cut out the moduli spaces. A thorough treatment of orientability in general can be found in [12].

For some fixed $p > 2$ and some real $\delta > 0$ (both of which we suppress from the notation), consider the space $\mathcal{B}(\mathbf{x}, \mathbf{y}) = \mathcal{B}_\delta(\mathbf{x}, \mathbf{y})$ of maps discussed in Subsection 3.2. The moduli spaces of holomorphic disks are finite-dimensional subspaces of this Banach manifold.

Recall that for a family F_x of Fredholm operators parametrized by an auxiliary space X , the virtual vector spaces $\ker F_x - \text{Coker } F_x$ naturally fit together

to give rise to an element in the K -theory of X (see [2]), the *virtual index bundle*. The determinant of this is a real line bundle over X , the *determinant line bundle* of the family F_x .

PROPOSITION 3.10. *There is a trivial line bundle over $\mathcal{B}(\mathbf{x}, \mathbf{y})$ whose restriction to the moduli space $\mathcal{M}_{J_s}(\mathbf{x}, \mathbf{y}) \subset \mathcal{B}(\mathbf{x}, \mathbf{y})$ is naturally identified with the determinant line bundle for the linearization $\det(D_u)$, where J_s is any path of (j, η, V) -nearly-symmetric almost-complex structures.*

As we shall see, the main ingredient in the above proposition is the fact that the totally real subspaces \mathbb{T}_α and \mathbb{T}_β have trivial tangent bundles. We shall give the proof after a preliminary discussion.

Let $L_0(t)$ and $L_1(t)$ be a pair of paths of totally real subspaces of \mathbb{C}^n , indexed by $t \in \mathbb{R}$, which are asymptotically constant as $t \mapsto \pm\infty$, i.e. there are totally real subspaces L_0^-, L_1^-, L_0^+ , and L_1^+ with the property that

$$\lim_{t \rightarrow \pm\infty} L_0(t) = L_0^\pm, \quad \lim_{t \rightarrow \pm\infty} L_1(t) = L_1^\pm.$$

Suppose moreover that L_0^- and L_1^- are transverse, and similarly L_0^+ and L_1^+ are transverse, too. Then, the $\bar{\partial}$ on \mathbb{C}^g -valued functions on the strip, satisfying boundary conditions specified by the paths $L_0(t)$ and $L_1(t)$

$$\bar{\partial}: \left\{ f \in L_1^p([0, 1] \times \mathbb{R}; \mathbb{C}^g) \left| \begin{array}{l} f(1 + it) \in L_0(t), \\ f(0 + it) \in L_1(t), \\ \bar{\partial}f = 0 \end{array} \right. \right\} \longrightarrow L^p([0, 1] \times \mathbb{R}; \mathbb{C}^g)$$

is Fredholm. Thus, the $\bar{\partial}$ operator induces a family of Fredholm operators indexed by the space

$$\mathcal{P} = \left\{ L_0, L_1: [0, 1] \longrightarrow G\mathbb{R}(g) \left| \begin{array}{l} L_0(0) = L_0^-, \\ L_1(0) = L_1^-, \\ L_0(1) = L_0^+, \\ L_1(1) = L_1^+ \end{array} \right. \right\}$$

(after reparametrizing the paths in \mathcal{P} to be indexed by $\mathbb{R} \cup \{\pm\infty\}$ rather than $[0, 1]$), where $G\mathbb{R}(g)$ denotes the Grassmannian of totally real g -dimensional subspaces of \mathbb{C}^g .

The index of the linearization D_u of the $\bar{\partial}_{J_s}$ operator on maps of the disk into $\text{Sym}^g(\Sigma)$ can be related to the index of the $\bar{\partial}$ -operators over \mathcal{P} , as follows. First observe that D_u depends on a path J_s of almost-complex structures. However, we can connect the family to the constant path $\text{Sym}^g(j)$, without changing the index bundle. Next, fix a contraction of the unit disk to $-i$. Together with a connection over $T\text{Sym}^g(\Sigma)$, this induces a trivialization for any $u \in \mathcal{B}(\mathbf{x}, \mathbf{y})$ of the pull-back of the complex tangent bundle of $\text{Sym}^g(\Sigma)$ (induced from $\text{Sym}^g(j)$). Via these trivializations, the one-parameter family of

totally real subspaces

$$\{t \mapsto T_{u(1+it)}\mathbb{T}_\alpha \subset T_{u(1+it)}\mathrm{Sym}^g(\Sigma)\}, \{t \mapsto T_{u(0+it)}\mathbb{T}_\beta \subset T_{u(0+it)}\mathrm{Sym}^g(\Sigma)\}$$

induce one-parameter families $L_0(t)$ and $L_1(t)$ of totally real subspaces of \mathbb{C}^g . Indeed, if we use a connection over $T\mathrm{Sym}^g(\Sigma)$ which is trivial along \mathbb{T}_β , and we choose the contraction of our disk to preserve the left arc, both $t = 0$ and $t = 1$ endpoints of the families can be viewed as fixed (i.e. independent of the particular choice of u). Thus, we have a map

$$\Psi: \mathcal{B}(\mathbf{x}, \mathbf{y}) \longrightarrow \mathcal{P},$$

together with an identification between the pull-back of the (virtual) index bundle for $\bar{\partial}$ and the (virtual) index bundle for D_u (over the moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y}) \subset \mathcal{B}(\mathbf{x}, \mathbf{y})$).

We wish to study the index bundle over \mathcal{P} . There is a “difference” map

$$\delta: G\mathbb{R}(g) \times G\mathbb{R}(g) \longrightarrow \frac{\mathrm{Gl}_g(\mathbb{C})}{\mathrm{Gl}_g(\mathbb{R})},$$

where $\delta(L_0, L_1)$ is the equivalence class of any matrix $A \in \mathrm{Gl}_g(\mathbb{C})$ with the property that $AL_0 = L_1$. (By taking the difference, with $\mathbb{R}^g \subset \mathbb{C}^g$, we obtain a diffeomorphism between $G\mathbb{R}(g)$ and the homogeneous space $\frac{\mathrm{Gl}_g(\mathbb{C})}{\mathrm{Gl}_g(\mathbb{R})}$.) In this space, we have a Maslov cycle

$$Z_\mu \subset \frac{\mathrm{Gl}_g(\mathbb{C})}{\mathrm{Gl}_g(\mathbb{R})} = \{[A] \mid \mathbb{R}^g + A(\mathbb{R}^g) \neq \mathbb{C}^n\}.$$

Of course, L_0 and L_1 meet transversally if and only if their difference $\delta(L_0, L_1)$ does not lie in the Maslov cycle.

Let $[a_0] = \delta(L_0^-, L_1^-)$, $[a_1] = \delta(L_0^+, L_1^+)$. The difference map gives us a map

$$\Phi: \mathcal{P} \longrightarrow \mathcal{Q} = \{A: [0, 1] \longrightarrow \frac{\mathrm{Gl}_g(\mathbb{C})}{\mathrm{Gl}_g(\mathbb{R})} \mid A(0) = [a_0], A(1) = [a_1]\}.$$

In this notation, then, the numerical index of the $\bar{\partial}$ operator associated to a pair of paths $L_0(t)$ and $L_1(t)$ in \mathcal{P} is calculated by the intersection number of the difference with the Maslov cycle:

$$\mathrm{ind}(\bar{\partial}(L_0(t), L_1(t))) = \delta(L_0(t), L_1(t)) \cap Z_\mu.$$

Moreover, we could work entirely over \mathcal{Q} : \mathcal{Q} is identified with the subspace of \mathcal{P} where $L_0(t) \equiv \mathbb{R}^g$, so there is an index bundle over \mathcal{Q} , and the index bundle for $\bar{\partial}$ over \mathcal{P} is easily seen to be the pull-back of this index bundle over \mathcal{Q} (since the index bundle over \mathcal{P} is trivial over the fiber of Φ).

Clearly,

$$\pi_2 \left(\frac{\mathrm{Gl}_g(\mathbb{C})}{\mathrm{Gl}_g(\mathbb{R})} \right) \cong \pi_1(\mathcal{Q}).$$

Now, if $g \geq 2$, it is easy to see that

$$\pi_2 \left(\frac{\mathrm{GL}_g(\mathbb{C})}{\mathrm{GL}_g(\mathbb{R})} \right) \cong \mathbb{Z}/2\mathbb{Z}.$$

Thus, there is no *a priori* reason for the determinant line bundle $\det(\bar{\partial}) \longrightarrow \mathcal{Q}$ to be trivial: its first Stiefel-Whitney class may evaluate nontrivially on the nontrivial homotopy class $\mathbb{Z}/2\mathbb{Z}$. Proposition 3.10 is established by giving a suitable lift of the composite map $\Phi \circ \Psi$.

Proof of Proposition 3.10. Continuing the above notation, fix matrices $a_0, a_1 \in \mathrm{GL}_g(\mathbb{C})$, and consider the space

$$\tilde{\mathcal{Q}} = \{A: [0, 1] \longrightarrow \mathrm{GL}_g(\mathbb{C}) \mid A(0) = a_0, A(1) = a_1\}.$$

Since $\pi_2(\mathrm{GL}_g(\mathbb{C})) = 0$, we see that the index bundle of the $\bar{\partial}$ operator over $\tilde{\mathcal{Q}}$ is orientable. Thus, to establish orientability of the determinant line bundle over the moduli spaces of flows, we lift $\Phi \circ \Psi$ to a map

$$\tilde{\Phi}: \mathcal{B}(\mathbf{x}, \mathbf{y}) \longrightarrow \tilde{\mathcal{Q}}.$$

To define this lift, note that the tangent spaces to \mathbb{T}_α and \mathbb{T}_β respectively can be trivialized by ordering and orienting the attaching circles α and β . This in turn gives rise to a complex trivialization of the restrictions of $TSym^g(\Sigma)$ to \mathbb{T}_α and \mathbb{T}_β respectively (induced from the identifications $TSym^g(\Sigma)|_{\mathbb{T}_\alpha} \cong \mathbb{T}_\alpha \otimes_{\mathbb{R}} \mathbb{C}$, $TSym^g(\Sigma)|_{\mathbb{T}_\beta} \cong \mathbb{T}_\beta \otimes_{\mathbb{R}} \mathbb{C}$ arising from the corresponding totally real structures). Given a holomorphic disk u , then, we let $A(t)$ denote the matrix corresponding to the linear transformation from \mathbb{C}^g to itself given by parallel transporting the vector space $\mathbb{C}^g \cong T_{u(1+it)}\mathbb{T}_\alpha \otimes_{\mathbb{R}} \mathbb{C}$ to $T_{u(1)}\mathbb{T}_\alpha \otimes_{\mathbb{R}} \mathbb{C} \cong \mathbb{C}^g$, using the arc which is the image under u of the path prescribed by the fixed contraction of \mathbb{D} . Now, the composite $\Phi \circ \Psi$ factors through the projection from $\tilde{\mathcal{Q}}$ to \mathcal{Q} , so the pull-back of the determinant of the index bundle is trivial since it is trivial over $\tilde{\mathcal{Q}}$. \square

We would like to choose orientations for all moduli spaces in a consistent manner. To this end, we construct “coherent orientations” closely following [10]. 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 3.11. A *coherent system of orientations* for \mathfrak{s} , \mathfrak{o} is a choice of nonvanishing 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 3.12. Let $(\Sigma, \alpha, \beta, z)$ be a Heegaard diagram representing Y , and let \mathfrak{s} be a Spin^c structure for Y . A *complete set of paths for \mathfrak{s}* is an enumeration $\{\mathbf{x}_0, \dots, \mathbf{x}_m\} = \mathcal{S}$ of all the intersection points of \mathbb{T}_α with \mathbb{T}_β representing \mathfrak{s} , and a collection of homotopy classes $\theta_i \in \pi_2(\mathbf{x}_0, \mathbf{x}_i)$ for $i = 1, \dots, m$ with $n_z(\theta_i) = 0$.

Fix periodic classes $\phi_1, \dots, \phi_b \in \pi_2(\mathbf{x}, \mathbf{x})$ representing a basis for $H^1(Y; \mathbb{Z})$, and nonvanishing sections of the determinant line bundle for bundle for the homotopy classes $\theta_1, \dots, \theta_m$ and ϕ_1, \dots, ϕ_b . These data uniquely determine a coherent system of orientations by splicing, since any homotopy class $\phi \in \pi_2(\mathbf{x}_i, \mathbf{x}_j)$ can be uniquely written as

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

3.7. Degenerate disks. Fix a nearly symmetric almost-complex structure J over $\text{Sym}^g(\Sigma)$. For each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, the moduli space of α -degenerate disks is the set of maps

$$\mathcal{N}_J(\mathbf{x}) = \left\{ u: [0, \infty) \times \mathbb{R} \longrightarrow \text{Sym}^g(\Sigma) \left| \begin{array}{l} u(\{0\} \times \mathbb{R}) \subset \mathbb{T}_\alpha \\ \lim_{z \rightarrow \infty} u(z) = \mathbf{x} \\ \frac{du}{ds} + J \frac{du}{dt} = 0 \end{array} \right. \right\}.$$

Equivalently, we can think of $\mathcal{N}_J(\mathbf{x})$ as the moduli space of J -holomorphic maps of the unit disk D in \mathbb{C} to $\text{Sym}^g(\Sigma)$, which carry ∂D into \mathbb{T}_α , and i to \mathbf{x} . This also gives rise to a finite-dimensional moduli space, partitioned according to the homotopy classes of maps satisfying these boundary conditions, a set which we denote by $\pi_2(\mathbf{x})$ (again, when $g = 2$, we also divide out by the action of $\pi_1(\text{Sym}^2(\Sigma))$). Suppose that $g > 1$. Since the map $\pi_1(\mathbb{T}_\alpha) \longrightarrow \pi_1(\text{Sym}^g(\Sigma))$ is injective, $\pi_2(\mathbf{x}) \cong \mathbb{Z}$ under the map $n_z(u)$; equivalently, if $O_{\mathbf{x}} \in \pi_2(\mathbf{x})$ is the homotopy class of the constant, then any other is given by $O_{\mathbf{x}} + k[S]$ for $k \in \mathbb{Z}$. Of course, when $g = 1$, there is only one homotopy class, and that is $O_{\mathbf{x}}$.

Note that there is a two-dimensional automorphism group acting on $\mathcal{N}_J(\mathbf{x})$ (pre-composing u by either a purely imaginary translation or a real dilation), and we denote the quotient space by $\hat{\mathcal{N}}_J(\mathbf{x})$. If $\phi \in \pi_2(\mathbf{x})$ is a homotopy class, then we let $\mathcal{N}_J(\phi)$, resp. $\hat{\mathcal{N}}_J(\phi)$, denote its corresponding component in $\mathcal{N}_J(\mathbf{x})$, resp. $\hat{\mathcal{N}}_J(\mathbf{x})$.

In studying smoothness properties of $\mathcal{N}_J(\mathbf{x})$, it is useful to have the following result concerning the complex structures $\text{Sym}^g(\mathfrak{j})$:

LEMMA 3.13. *Given a finite collection of points $\{\mathbf{x}_i\}_{i=1}^n$ in $\text{Sym}^g(\Sigma)$, the set of complex structures \mathfrak{j} over Σ for which there is a $\text{Sym}^g(\mathfrak{j})$ -holomorphic sphere containing at least one of the \mathbf{x}_i has real codimension two.*

Proof. The spheres in $\text{Sym}^g(\Sigma)$ for the complex structure $\text{Sym}^g(\mathfrak{j})$ are all contained in the set of critical points for the Abel-Jacobi map

$$\Theta: \text{Sym}^g(\Sigma) \longrightarrow \text{Pic}^g(\Sigma) \cong H^1(\Sigma; S^1),$$

which is a degree one holomorphic map. Thus, the set of spheres is contained in a subset of real codimension two. \square

PROPOSITION 3.14. *Suppose $\mathbf{x} \in \mathbb{T}_\alpha$ is not contained in any $\text{Sym}^g(\mathfrak{j})$ -holomorphic sphere in $\text{Sym}^g(\Sigma)$. Then, there is a contractible neighborhood \mathcal{U} of $\text{Sym}^g(\mathfrak{j})$ in $\mathcal{J}(\mathfrak{j}, \eta, V)$ with the property that for generic $J \in \mathcal{U}$, the moduli space $\widehat{\mathcal{N}}_J(O_{\mathbf{x}} + [S])$ is a compact, formally zero-dimensional space which is smoothly cut out by its defining equations.*

Proof. To investigate compactness, note first that a sequence of elements in $\widehat{\mathcal{N}}_J(\mathbf{x})$ has a subsequence which either bubbles off spheres, or additional disks with boundaries lying in \mathbb{T}_α . However, it is impossible for a sequence to bubble off a null-homotopic disk with boundary lying in \mathbb{T}_α , since such disks must be constant, as they have no energy (according to the proof of Lemma 3.5, see Equation (9)). Moreover, sequences in $O_{\mathbf{x}} + [S]$ cannot bubble off homotopically nontrivial disks, because then one of the components in the decomposition would have negative ω -integral, and such homotopy classes have no holomorphic representatives.

This argument also rules out bubbling off spheres, except in the special case where the subsequence converges to a single sphere (more precisely, the constant disk mapping to \mathbf{x} , attached to some sphere). But this is ruled out by our hypothesis on \mathfrak{j} , which ensures that for any J sufficiently close to $\text{Sym}^g(\mathfrak{j})$, the J -holomorphic spheres are disjoint from \mathbf{x} .

To prove smoothness, note first that any holomorphic disk in $\widehat{\mathcal{N}}(\mathbf{x})$ for $\text{Sym}^g(\mathfrak{j})$ has a dense set of injective points. To see this, fix any point $z' \in \Sigma - \alpha_1 - \cdots - \alpha_g$. The intersection number of $\{z'\} \times \text{Sym}^{g-1}(\Sigma)$ with u is $+1$, and both are varieties; it follows that there is a single point of intersection, i.e. there is only one (s, t) for which $u(s, t) \in \{z'\} \times \text{Sym}^{g-1}(\Sigma)$. Thus, u is injective in a neighborhood of (s, t) . It follows that for any J sufficiently close to $\text{Sym}^g(\mathfrak{j})$, all the J -holomorphic degenerate disks are injective in a

neighborhood of $u(s, t)$. Thus, according to the usual proof of transversality, these pairs are all smooth points in the parametrized moduli space. Thus, the result follows from the Sard-Smale theorem. \square

Proceeding as earlier, we can orient the moduli spaces $\widehat{\mathcal{N}}_J$. Our aim now is to prove the following:

THEOREM 3.15. *Fix a finite collection $\{\mathbf{x}_i\}$ of points in $\text{Sym}^g(\Sigma)$, and an almost-complex structure \mathfrak{j} over Σ for which each $\text{Sym}^g(\mathfrak{j})$ -holomorphic sphere misses $\{\mathbf{x}_i\}$. Then, there is a contractible open neighborhood of $\text{Sym}^g(\mathfrak{j})$, \mathcal{U} , in the space of nearly-symmetric almost-complex structures, with the property that for generic $J \in \mathcal{U}$, the total signed number of points in $\widehat{\mathcal{N}}_J(O_{\mathbf{x}_i} + [S])$ is zero.*

Thinking of Σ as the connected sum of g tori, each of which contains exactly one α_i , we can endow Σ with a complex structure with long connected sum tubes.

PROPOSITION 3.16. *If \mathfrak{j} is sufficiently stretched out along the connected sum tubes, then the moduli space $\widehat{\mathcal{N}}_{\mathfrak{j}}(O_{\mathbf{x}} + [S])$ is empty for any $\mathbf{x} \in \mathbb{T}_{\alpha}$.*

Proof. Fix a genus one Riemann surface E . Let \mathfrak{j}_t denote the complex structure on Σ , thought of as the connected sum of g copies of E , connected along cylinders isometric to $S^1 \times [-t, t]$. As $t \mapsto \infty$, the Riemann surface degenerates to the wedge product of g copies of E , $\bigvee_{i=1}^g E_i$.

If for each \mathfrak{j}_t , the moduli space were nonempty, we could take the Gromov limit of a sequence u_t in $\widehat{\mathcal{N}}_{\mathfrak{j}_t}(O_{\mathbf{x}} + [S])$ to obtain a holomorphic map u_{∞} into $\text{Sym}^g(E_1 \vee \dots \vee E_g)$ (a linear chain of g tori meeting in $g - 1$ nodes). (In this argument, we have a one-parameter family of symmetric products, which we can embed into a fixed Kähler manifold, where we can apply the usual Gromov compactness theorem, see also Section 10.) The latter symmetric product decomposes into irreducible components

$$\bigcup_{\{k_1, \dots, k_g \in \mathbb{Z} \mid 0 \leq k_i \leq g, \ k_1 + \dots + k_g = g\}} \text{Sym}^{k_1}(E_1) \times \dots \times \text{Sym}^{k_g}(E_g).$$

These components meet along loci containing the connected sum points for the various E_i . Moreover, the torus \mathbb{T}_{α} can be viewed as a subset of the irreducible component $E_1 \times \dots \times E_g$ (corresponding to all $k_i = 1$). The Gromov limit u_{∞} then consists of a holomorphic disk v with boundary mapping into \mathbb{T}_{α} , and a possible collection of spheres bubbling off into the other irreducible components. But $\pi_2(E_1 \times \dots \times E_g, \alpha_1 \times \dots \times \alpha_g) = 0$, so it follows that v is constant, mapping to $\mathbf{x} \in \mathbb{T}_{\alpha}$ (which is disjoint from the connected sum

points). Since v misses the other components of the symmetric product, it cannot meet any of the spheres, so v is the Gromov limit of the u_t . But, $n_z(v) = 0$, while we have assumed that $n_z(u_t) = 1$. \square

LEMMA 3.17. *Let \mathbf{x} , \mathbf{j} , and \mathcal{U} be as in Proposition 3.14. Suppose that $J_1, J_2 \in \mathcal{U}$ are a pair of generic almost-complex structures, in the sense that $\widehat{\mathcal{N}}_{J_s}(\mathbf{x})$ is smooth for $s = 0$ and 1. Then, these moduli spaces are compactly cobordant.*

Proof. We connect J_1 and J_2 by a generic path $\{J_s\}$ in \mathcal{U} . As in the proof of Proposition 3.14, this gives rise to the required compact cobordism. Note that the possibility of bubbling off a sphere is ruled out, choosing \mathcal{U} small enough to ensure that \mathbf{x} is disjoint from all J -holomorphic spheres with $J \in \mathcal{U}$. \square

Proof of Theorem 3.15. Let \mathbf{j} be any complex structure over Σ for which the $\text{Sym}^g(\mathbf{j})$ -holomorphic spheres miss the $\{\mathbf{x}_i\}$. Let $\mathcal{U} \subset \bigcap_{i=1}^n \mathcal{U}_i$ be a contractible, open subset of the open subsets \mathcal{U}_i given to us by Proposition 3.14 for the points $\mathbf{x}_i \in \mathbb{T}_\alpha$. According to Lemma 3.17, the number of points $\#\widehat{\mathcal{N}}_J(O_{\mathbf{x}_i} + [S])$ is independent of J , i.e. it depends only on the complex structure \mathbf{j} over Σ . In fact, if J is a generic \mathbf{j} -nearly-symmetric almost-complex structure, and J' is a sufficiently close \mathbf{j}' -nearly-symmetric almost-complex structure, then the moduli spaces are identified. It follows that $\#\widehat{\mathcal{N}}_J(O_{\mathbf{x}_i} + [S])$ is a locally-constant function of the complex structure \mathbf{j} . Since the space of complex structures for which the $\text{Sym}^g(\mathbf{j})$ -holomorphic spheres miss $\{\mathbf{x}_i\}$ is connected, the theorem follows from Proposition 3.16. \square

3.8. Structure of moduli spaces.

THEOREM 3.18. *Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ be a Heegaard diagram with curves in general position. For a generic path J_s of nearly-symmetric almost-complex structures, the following holds: There is no nonconstant J_s -holomorphic disk u with $\mu(u) \leq 0$. Moreover for each $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = 1$, the quotient space*

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

is a compact, zero-dimensional manifold.

Proof. The first part follows directly from the Theorem 3.4.

Compactness follows from the usual compactification theorem for holomorphic curves (see [9] and also [14], [30], [40]), which hold thanks to the energy bound (Lemma 3.5).

Specifically, the compactness theorem says that a sequence of points in the moduli spaces converges to an ideal disk, with possible broken flowlines,

boundary degenerations, and bubblings of spheres. Broken flowlines are excluded by the additivity of the Maslov index, and the transversality result Theorem 3.4. Spheres and boundary degenerations both carry Maslov index at least two, so that these kinds of degenerations are excluded as well. \square

4. Definition of the Floer homology groups

We are now ready to define the Floer homology groups for three-manifolds. In Subsection 4.1, we consider the technically simpler case of three-manifolds with $b_1(Y) = 0$. We then turn our attention to the issues which arise when pushing this definition to the case of three-manifolds with $b_1(Y) > 0$: the cyclic gradings in Subsection 4.2.1, the “admissibility hypotheses” on the Heegaard diagrams required for topological invariance of the constructions in Subsection 4.2.2. (We will return to the construction of such Heegaard diagrams in Subsection 5.) With these technical pieces in place, we proceed as before to define the Floer homology groups when $b_1(Y) > 0$, in Subsection 4.2.3. These groups can be endowed the additional structure of the action by $H_1(Y; \mathbb{Z})$, which is constructed in Subsection 4.2.5. An additional modification — Floer homology with “twisted coefficients” — is introduced in Section 8 of [28].

4.1. *Floer homologies when $b_1(Y) = 0$.* Let $(\Sigma, \alpha, \beta, z)$ be a pointed Heegaard diagram with genus $g > 0$ for a rational homology three-sphere Y , where the α and β are in general position, and choose a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$. We let \mathfrak{S} be the set of intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $s_z(\mathbf{x}) = \mathfrak{s}$. We fix also the following auxiliary data:

- a coherent orientation system, in the sense of Definition 3.11 (note that this is not necessary when defining Floer homology groups with $\mathbb{Z}/2\mathbb{Z}$ coefficients),
- a generic complex structure j over Σ (generic in the sense that each intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is disjoint from the $\text{Sym}^g(j)$ -holomorphic spheres in $\text{Sym}^g(\Sigma)$ — see Lemma 3.13),
- a generic path of nearly symmetric almost-complex structure J_s over $\text{Sym}^g(\Sigma)$, contained in the open subset \mathcal{U} of Theorem 3.15 (associated to the subset $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and also V).

Let $\widehat{CF}(\alpha, \beta, \mathfrak{s})$ denote the free Abelian group generated by the points in $\mathfrak{S} \subset \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. This group can be endowed with a relative grading¹, defined by

$$(10) \quad \text{gr}(\mathbf{x}, \mathbf{y}) = \mu(\phi) - 2n_z(\phi),$$

¹A relatively graded Abelian group is one which is generated by elements partitioned into equivalence classes \mathfrak{S} , with a relative grading function $\text{gr}: \mathfrak{S} \times \mathfrak{S} \rightarrow \mathbb{Z}$, satisfying

where ϕ is any element $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, and μ is the Maslov index. In view of Proposition 2.15 and Lemma 3.3, this integer is independent of the choice of Whitney disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$.

Let

$$\partial: \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \longrightarrow \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$$

be the map defined by:

$$\partial \mathbf{x} = \sum_{\{\mathbf{y} \in \mathfrak{S} \mid \text{gr}(\mathbf{x}, \mathbf{y})=1\}} \# \left(\widehat{\mathcal{M}}_0(\mathbf{x}, \mathbf{y}) \right) \mathbf{y},$$

where $\widehat{\mathcal{M}}_0(\mathbf{x}, \mathbf{y}) = \widehat{\mathcal{M}}(\phi)$ for the element $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $n_z(\phi) = 0$ and $\mu(\phi) = 1$. Note that by Proposition 2.15 and Lemma 3.3, there is at most one such homotopy class. Also, counting these holomorphic disks in $\text{Sym}^g(\Sigma)$ is equivalent to counting holomorphic disks in $\text{Sym}^g(\Sigma - z)$, in view of Lemma 3.2. (We have suppressed the path J_s from the notation, but one should bear in mind that ∂ does depend on the path J_s . When it is important to call attention to this dependence, we write ∂_{J_s} ; see the proof of Theorem 6.1 below.)

The count appearing in the above boundary operator is, as usual, meant to signify a signed (oriented) count of points in the compact, zero-dimensional moduli spaces (see Theorem 3.18, and Subsection 3.6), and as such, it depends on a coherent orientation system as defined in Definition 3.11. As we shall see in Lemma 4.16 in the present case, such different choices give rise to isomorphic chain complexes, so we shall usually drop them from the notation.

THEOREM 4.1. *When $b_1(Y) = 0$, the pair $(\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}), \partial)$ is a chain complex; i.e. $\partial^2 = 0$.*

Proof. This follows in the usual manner from the compactifications of the one-dimensional moduli spaces $\widehat{\mathcal{M}}(\phi)$ with $\mu(\phi) = 2$ (together with the gluing descriptions of the neighborhoods of the ends). Note that if $\mathbf{x}' \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta - \mathfrak{S}$, then $\varepsilon(\mathbf{x}, \mathbf{x}') \neq 0$, so there are no flows connecting \mathbf{x} to \mathbf{x}' . We note also that there are no spheres in $\text{Sym}^g(\Sigma - z)$ or degenerate holomorphic disks (whose boundaries lie entirely in \mathbb{T}_α or \mathbb{T}_β), so the only boundary components in the compactification consist of broken flow-lines. \square

Definition 4.2. The Floer homology groups $\widehat{HF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ are the homology groups of the complex $(\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}), \partial)$.

Next, let $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ be the free Abelian group generated by pairs $[\mathbf{x}, i]$ where $\mathbf{x} \in \mathfrak{S}$, and $i \in \mathbb{Z}$ is an integer. We give the generators a relative grading

$\text{gr}(\mathbf{x}, \mathbf{y}) + \text{gr}(\mathbf{y}, \mathbf{w}) = \text{gr}(\mathbf{x}, \mathbf{w})$ for each $\mathbf{x}, \mathbf{y}, \mathbf{w} \in \mathfrak{S}$. When the corresponding theory for four-manifolds is developed, this relative \mathbb{Z} -grading can be lifted to an absolute \mathbb{Q} -grading; see [29].

defined by

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

Let

$$\partial^\infty: CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \longrightarrow CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$$

be the map defined by:

$$(11) \quad \partial^\infty[\mathbf{x}, i] = \sum_{\mathbf{y} \in \pi_2(\mathbf{x}, \mathbf{y})} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi)=1\}} \# \left(\widehat{\mathcal{M}}(\phi) \right) [\mathbf{y}, i - n_z(\phi)].$$

Although we have written the above expression as a double-sum, Proposition 2.15 and Lemma 3.3 ensure that for given \mathbf{x} and \mathbf{y} , there is at most one homotopy class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = 1$.

THEOREM 4.3. *When $b_1(Y) = 0$, the pair $(CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}), \partial^\infty)$ is a chain complex; i.e. $(\partial^\infty)^2 = 0$.*

Proof. As is usual in Floer's theory, one considers the ends of the moduli spaces $\widehat{\mathcal{M}}(\phi)$, where $\phi \in \pi_2(\mathbf{x}, \mathbf{w})$ satisfies $\mu(\phi) = 2$. This space has *a priori* three kinds of ends:

- (1) those corresponding to “broken flow-lines”, i.e. a pair $u \in \mathcal{M}(\mathbf{x}, \mathbf{y})$ and $v \in \mathcal{M}(\mathbf{y}, \mathbf{w})$ with $\mu(u) = \mu(v) = 1$,
- (2) those which correspond to a sphere bubbling off, i.e. another $v \in \mathcal{M}(\mathbf{x}, \mathbf{w})$ and a holomorphic sphere $S \in \mathrm{Sym}^g(\Sigma)$ which meets v ,
- (3) those which correspond to “boundary bubbling”; i.e., there are a $v \in \mathcal{M}(\mathbf{x}, \mathbf{w})$, and a holomorphic map u from the disk, whose boundary is mapped into \mathbb{T}_α or \mathbb{T}_β , which meet in a point on the boundary.

(In principle, several of the above degenerations could happen at once — multiple broken flows, spheres, and boundary degenerations, but these multiple degenerations are easily ruled out by dimension counts and the transversality theorem, Theorem 3.4.)

In the Cases (2) and (3), we argue that $[v] = \phi - \ell[S]$ (note that a disk whose boundary lies entirely inside \mathbb{T}_α or \mathbb{T}_β also has a corresponding domain $\mathcal{D}(u)$, which, in this case, must be a multiple of Σ ; if u is pseudo-holomorphic, then $\mathcal{D}(u) = \ell[\Sigma]$ for $\ell \geq 0$ according to Lemma 3.2, and if $\mathcal{D}(u) = 0$, the disk must be constant). Thus it follows from Lemma 3.3 that $\mu([v]) = \mu(\phi) - 2\ell$. From transversality (Theorem 3.4), it follows that $\ell = 1$ and v must be constant; so that, in particular, $\mathbf{x} = \mathbf{w}$. Now for generic \mathbf{j} , we know that the holomorphic spheres miss the intersection points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$; hence, the case of spheres bubbling off is excluded.

Thus, when $\mathbf{x} \neq \mathbf{w}$, boundary bubbles are excluded, so, counting the ends of the moduli space $\mathcal{M}(\phi)$, we get that

$$0 = \sum_{\mathbf{y}} \sum_{\{\psi \in \pi_2(\mathbf{x}, \mathbf{y}), \zeta \in \pi_2(\mathbf{y}, \mathbf{w}) \mid \psi * \zeta = \phi\}} \left(\# \mathcal{M}(\psi) \right) \cdot \left(\# \mathcal{M}(\zeta) \right).$$

When $\mathbf{x} = \mathbf{w}$, there are additional terms, corresponding to the boundary bubbles (and the gluing descriptions of the ends), giving a relation

$$0 = \# \hat{\mathcal{N}}^\alpha(\mathbf{x}) + \# \hat{\mathcal{N}}^\beta(\mathbf{x}) + \sum_{\mathbf{y}} \sum_{\{\psi \in \pi_2(\mathbf{x}, \mathbf{y}), \zeta \in \pi_2(\mathbf{y}, \mathbf{x}) \mid \psi * \zeta = \phi\}} \left(\# \mathcal{M}(\psi) \right) \cdot \left(\# \mathcal{M}(\zeta) \right);$$

see for example [12]. But the terms $\# \hat{\mathcal{N}}^\alpha(\mathbf{x})$ and $\# \hat{\mathcal{N}}^\beta(\mathbf{x})$ both vanish, according to Theorem 3.15.

From the additivity of n_z under juxtapositions of flow-lines, it follows that the double sums considered above are coefficients of $\partial^2[\mathbf{x}, i]$. \square

There is a chain map

$$U: CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \longrightarrow CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}),$$

which lowers the degree by two, defined by

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

Let $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ denote the subgroup of $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ which is freely generated by pairs $[\mathbf{x}, i]$, where $i < 0$. Let $CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ denote the quotient group $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})/CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$.

LEMMA 4.4. *The group $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ is a subcomplex of $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$, so there is a short exact sequence of chain complexes:*

$$0 \longrightarrow CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \xrightarrow{i} CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \xrightarrow{\pi} CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \longrightarrow 0.$$

Proof. The fact that $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ is a subcomplex is an easy consequence of Lemma 3.2. \square

Clearly, U restricts to an endomorphism of $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ (which lowers the degree by 2), and hence it also induces an endomorphism of the quotient $CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$. Sometimes, for clarity we denote the induced actions on these complexes (or their homologies) by U^- or U^+ . It is also easy to see that there is a short exact sequence

$$0 \longrightarrow \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \xrightarrow{\iota} CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \xrightarrow{U^+} CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}) \longrightarrow 0,$$

where $\iota(\mathbf{x}) = [\mathbf{x}, 0]$. In view of this, we declare the U action on $\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s})$ to be trivial.

Definition 4.5. Let $HF^\infty(\alpha, \beta, \mathfrak{s})$, $HF^-(\alpha, \beta, \mathfrak{s})$, and $HF^+(\alpha, \beta, \mathfrak{s})$ denote the homologies of the complexes $CF^-(\alpha, \beta, \mathfrak{s})$, $CF^\infty(\alpha, \beta, \mathfrak{s})$, and $CF^+(\alpha, \beta, \mathfrak{s})$ respectively, thought of as $\mathbb{Z}[U]$ modules.

Note that these constructions can be carried out with coefficients in any ring Λ : one defines the corresponding chain complexes as free Λ -modules, and the homology groups obtained in this way are modules over the polynomial algebra $\Lambda[U]$. We have no particular use for this construction, as \mathbb{Z} is a “universal” case, though we do point out that if we had chosen to use $\Lambda = \mathbb{Z}/2\mathbb{Z}$, then the issues of orientation (and choices of orientation systems) would become unnecessary.

In the interest of conciseness, we have suppressed additional data — notably, complex structures (and their perturbations) and orientation systems — from the notation of these homology groups. In fact, we will show that the homology groups are independent of these choices.

There is another algebraic construction which gives rise to a finitely-generated variant of Floer homology, which we define with the help of the following:

LEMMA 4.6. *If k is sufficiently large, then*

$$\mathrm{Im}(U^+)^k = \mathrm{Im}(\pi_*),$$

where $\pi_*: HF^\infty(\alpha, \beta, \mathfrak{s}) \longrightarrow HF^+(\alpha, \beta, \mathfrak{s})$ is the naturally induced map on homology. Similarly

$$\mathrm{Ker}(U^-)^k = \mathrm{Ker}(i_*),$$

for $i_*: HF^-(\alpha, \beta, \mathfrak{s}) \longrightarrow HF^\infty(\alpha, \beta, \mathfrak{s})$.

Proof. We begin with the claim about HF^- . It is clear that the kernel of i_* consists of those homology classes $\xi \in HF^-(\alpha, \beta, \mathfrak{s})$ for which $(U^-)^k \xi = 0$ for sufficiently large integers k . To see that k can be chosen independent of ξ , observe that CF^- is a finitely generated chain complex over the ring $\mathbb{Z}[U^-]$, and hence its homology is finitely generated over $\mathbb{Z}[U^-]$, as well. Thus, the sequence of submodules of $HF^-(\alpha, \beta, \mathfrak{s})$

$$\mathrm{Ker}(U^-) \subseteq \mathrm{Ker}(U^-)^2 \subseteq \mathrm{Ker}(U^-)^3 \subseteq \dots$$

must stabilize.

Now, we claim that for k large enough, $\mathrm{Ker}(U^-)^k = \mathrm{Ker}(i_*)$; it is also the case that $\mathrm{Im}(U^+)^k = \mathrm{Im}(\pi_*)$. First, since U is an automorphism of $HF^\infty(\alpha, \beta, \mathfrak{s})$, it is clear that $\mathrm{Im}(U^+)^k \supseteq \mathrm{Im}(\pi_*)$. Conversely, if $\xi \in \mathrm{Im}(U^+)^k$, then writing $\xi = (U^+)^k \eta$, we have that $\delta \xi = (U^-)^k \delta \eta = 0$ (since $i_* \circ \delta = 0$); thus $\xi \in \mathrm{Im}(\pi_*)$. \square

Definition 4.7. Let

$$HF_{\text{red}}^+(\alpha, \beta, \mathfrak{s}) = HF^+(\alpha, \beta, \mathfrak{s}) / \text{Im}(U^+)^k,$$

for sufficiently large k . Similarly, let

$$HF_{\text{red}}^-(\alpha, \beta, \mathfrak{s}) = \text{Ker}(U^-)^k \subset HF^-(\alpha, \beta, \mathfrak{s}).$$

PROPOSITION 4.8. *The boundary homomorphism of the long exact sequence induces a U -equivariant isomorphism*

$$HF_{\text{red}}^+(\alpha, \beta, \mathfrak{s}) \cong HF_{\text{red}}^-(\alpha, \beta, \mathfrak{s}).$$

Moreover, both are finitely generated \mathbb{Z} modules.

Proof. The isomorphism follows readily from Lemma 4.6 and the long exact sequence. Since $HF_{\text{red}}^-(\alpha, \beta, \mathfrak{s})$ is a finitely-generated $\mathbb{Z}[U^-]$ -module which is annihilated by $(U^-)^k$, it follows immediately that this module is also finitely generated over \mathbb{Z} . \square

In view of the above result, we will denote $HF_{\text{red}}^+(\alpha, \beta, \mathfrak{s}) \cong HF_{\text{red}}^-(\alpha, \beta, \mathfrak{s})$ simply by $HF_{\text{red}}(\alpha, \beta, \mathfrak{s})$.

4.2. Constructions when $b_1(Y) > 0$.

4.2.1. Grading. As before, we fix a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, and let \mathfrak{S} be the set of intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $s_z(\mathbf{x}) = \mathfrak{s}$. The expression used to define relative grading on \mathfrak{S} as in Equation (10) is now well-defined only modulo an indeterminacy, given by the Maslov indices of periodic classes. Indeed, this indeterminacy is given by the following more familiar quantity

$$(12) \quad \mathfrak{d}(\mathfrak{s}) = \gcd_{\xi \in H_2(Y; \mathbb{Z})} \langle c_1(\mathfrak{s}), \xi \rangle,$$

in view of the following result, which is proved in Subsection 9.3:

THEOREM 4.9. *Fix a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$. 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 $\psi \in \Pi_{\mathbf{x}}$,*

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

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

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

Definition 4.10. A pointed Heegaard diagram is called *strongly admissible* for the Spin^c structure \mathfrak{s} if for every nontrivial 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 nontrivial periodic domain \mathcal{D} with

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

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

Remark 4.11. 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 4.12. 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 nontrivial domain has positive area.

Assume, conversely, that each nontrivial periodic domain has both positive and negative coefficients. By changing the volume form, we are free to make each domain in $\Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$ have arbitrary positive area. Thus, the claim now reduces to some linear algebra. We say that a vector subspace $V \subset \mathbb{R}^m$ is *balanced* if each of its nonzero vectors has both positive and negative components. The claim, then, follows from 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 $\Pi_i: \mathbb{R}^m \rightarrow \mathbb{R}^{m-1}$ denote the projection map $\Pi_i(x_1, \dots, x_m) = (x_1, \dots, \widehat{x_i}, \dots, x_m)$. Either $\Pi_i(V)$ is also balanced, or V contains a vector v whose i^{th} component is $+1$, all other components are nonpositive, 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 $\xi = (\xi_1, \dots, \xi_{i-1}, 0, \xi_{i+1}, \dots, \xi_m)$ 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 $\xi = (0, \xi_2, \dots, \xi_m)$ and $\eta = (\eta_1, 0, \eta_3, \dots, \eta_m)$ with $\xi_i > 0$ for $i \neq 1$, and $\eta_i > 0$ for $i \neq 2$. Then, $\xi + \eta$ is our required vector. \square

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

LEMMA 4.13. *Suppose that $(\Sigma, \alpha, \beta, z)$ is weakly admissible for the Spin^c structure \mathfrak{s} , and fix integers $j, k \in \mathbb{Z}$. Then, for each $\mathbf{x}, \mathbf{y} \in \mathfrak{S}$, 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 initial $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$. Then, in view of Theorem 4.9, any other $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ has the form

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

where $\mathcal{P}_{\mathbf{x}}$ is some periodic class, \mathcal{P} its associated periodic domain, and S is the positive generator of $\pi_2(\text{Sym}^g(\Sigma))$. If $n_z(\psi) = n_z(\phi)$, this forces $\mathcal{D}(\phi) = \mathcal{D}(\psi) + \mathcal{P}$ for some periodic domain whose associated homology class is annihilated by $c_1(\mathfrak{s})$. If $\mathcal{D}(\phi) \geq 0$, then $\mathcal{P} \geq -\mathcal{D}(\psi)$.

Thus, the lemma follows from the observation that for any fixed $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$, there are only finitely many periodic domains \mathcal{P} in the set

$$Q = \{\mathcal{P} \in \Pi_{\mathbf{x}} \mid \langle c_1(\mathfrak{s}), H(\mathcal{P}) \rangle = 0, \mathcal{P} \geq -\mathcal{D}(\psi)\}.$$

We see this as follows. Let m denote the total number of domains (components in $\Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$). 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 nonnegative. Of course, if the polytope, the subspace of $H_2(Y; \mathbb{Z})$ annihilated by $c_1(\mathfrak{s})$, corresponding to periodic domains with only nonnegative multiplicities has a nontrivial real vector, then it must also have a nontrivial rational vector. After clearing denominators, we obtain a nonzero periodic domain (with integer coefficients) annihilating $c_1(\mathfrak{s})$, with only nonnegative coefficients. This contradicts the hypothesis of weak admissibility. \square

LEMMA 4.14. *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 reference $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$. Then, as in the previous lemma, any other class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ can be written as

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

Thus, each class ϕ with $\mathcal{D}(\phi) \geq 0$ corresponds to a periodic domain \mathcal{P} with

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

The lemma follows from the fact that (for fixed ψ) there are only finitely many periodic domains satisfying this inequality. This follows as in the proof of Lemma 4.13: an infinite number of such periodic domains would give rise to 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. \square

We will establish the existence of admissible Heegaard diagrams in Section 5.

4.2.3. *The chain complex.* Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a pointed Heegaard diagram for a three-manifold with $b_1(Y) > 0$, and fix a Spin^c structure \mathfrak{s} . Suppose moreover that the Heegaard diagram is strongly \mathfrak{s} -admissible. In this case, we define the groups $\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$, $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$ as in Subsection 4.1. Note that when $b_1(Y) > 0$, we include the coherent orientation system in the notation, since now the Floer homologies depend on this choice — we shall see in Lemma 4.16 that there are in principle $2^{b_1(Y)}$ different possible chain complexes corresponding to variations in this choice. Equation (10) now endows $\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$ with a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{s})$ -grading, where $\mathfrak{d}(\mathfrak{s})$ is given in Equation (12) (Subsection 4.2.1). We can define the subgroup $CF^-(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$ and quotient group $CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$ as before. We endow $CF^\infty(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$ with the differential from Equation (11), endowing $CF^+(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o})$ with the induced differential

$$\partial^+[\mathbf{x}, i] = \sum_{\mathbf{y} \in \pi_2(\mathbf{x}, \mathbf{y})} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi) = 1, i \geq n_z(\phi)\}} \# \left(\widehat{\mathcal{M}}(\phi) \right) [\mathbf{y}, i - n_z(\phi)].$$

THEOREM 4.15. *Let Y be a three-manifold equipped with a Spin^c structure \mathfrak{s} . Then,*

- *if $(\Sigma, \alpha, \beta, z)$ is strongly \mathfrak{s} -admissible, then $CF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o}, \partial^\infty)$ is a chain complex, with subcomplex CF^- and quotient complex CF^+ ;*
- *if $(\Sigma, \alpha, \beta, z)$ is weakly \mathfrak{s} -admissible, then $CF^+(\alpha, \beta, \mathfrak{s}, \mathfrak{o}, \partial^+)$ is a chain complex with subcomplex $\widehat{CF}(\alpha, \beta, \mathfrak{s}, \mathfrak{o}, \widehat{\partial})$.*

Proof. When $(\Sigma, \alpha, \beta, \mathfrak{s}, \mathfrak{o})$ is strongly \mathfrak{s} -admissible, the key point is to observe that the boundary operator defined in Equation (11) (Subsection 4.1) is actually a finite sum. This follows from Lemma 5.4. Similarly, when the diagram is only weakly admissible, Lemma 4.13 ensures that the differentials for \widehat{CF} and CF^+ are finite sums. With these remarks, the proof proceeds exactly as in the proof of Theorem 4.3. \square

4.2.4. *Coherent orientation systems.* Although the above chain complexes depend on the choice of orientation system, its isomorphism type depends on the orientation system only through its equivalence class, in the following sense. Let \mathfrak{o} and \mathfrak{o}' be a pair of systems of coherent orientations for \mathfrak{s} . Then, we define their difference $\delta = \delta(\mathfrak{o}, \mathfrak{o}') \in \text{Hom}(H^1(Y; \mathbb{Z}), \mathbb{Z}/2\mathbb{Z})$ as follows. Let $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ be the periodic class representing some cohomology 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. Clearly, there are $2^{b_1(Y)}$ inequivalent choices of orientation conventions.

LEMMA 4.16. *If \mathfrak{o} and \mathfrak{o}' are equivalent orientation systems in the above sense, then the chain complexes $CF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o})$ and $CF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o}')$ (and the corresponding CF^- , CF^+ , and \widehat{CF}) are isomorphic.*

Proof. Let \mathfrak{o} and \mathfrak{o}' be a pair of isomorphic coherent orientation systems. Fix a reference point $\mathbf{x}_0 \in \mathfrak{S}$. Given any other $\mathbf{x} \in \mathfrak{S}$ and path $\phi \in \pi_2(\mathbf{x}_0, \mathbf{x})$, there is a sign $\sigma(\mathbf{x}) \in \{\pm 1\}$ with the property that $\mathfrak{o}(\phi) = \sigma(\mathbf{x}) \cdot \mathfrak{o}'(\phi)$. As the notation suggests, if \mathfrak{o} and \mathfrak{o}' are isomorphic orientation systems, the number $\sigma(\mathbf{x})$ is independent of the choice of ϕ , so that we can define a map

$$f: CF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o}) \longrightarrow CF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o}')$$

by $f([\mathbf{x}, i]) = \sigma(\mathbf{x}) \cdot [\mathbf{x}, i]$. It is straightforward to verify that f induces an isomorphism of chain complexes. \square

4.2.5. *Additional algebra: the $H_1(Y; \mathbb{Z})/\text{Tors}$ and U -actions.* As in the case where $b_1(Y) = 0$, we define HF^- , HF^∞ , HF^+ , and \widehat{HF} to be the ho-

mologies of the corresponding chain complexes; and, as before, all of these homology groups come with the structure of a $\mathbb{Z}[U]$ module, where U lowers the relative grading by two. Similarly, we can form a reduced group HF_{red} using the procedure outlined earlier. Moreover, when $b_1(Y) > 0$, there is a new algebraic object: an action of $H_1(Y; \mathbb{Z})/\text{Tors}$, the Floer homology groups. Recall from the proof of Proposition 2.15 that the choice of basepoint gives an isomorphism

$$(13) \quad \begin{aligned} H^1(\Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta); \mathbb{Z}) &\cong \text{Hom}(\pi_1(\Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta)), \mathbb{Z}) \\ &\cong \pi_2(\text{Sym}^g(\Sigma)) \oplus \text{Hom}(H^1(Y, \mathbb{Z}), \mathbb{Z}). \end{aligned}$$

PROPOSITION 4.17. *There is a natural action of $H^1(\Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta))$ lowering the 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 $\Lambda^*(H_1(Y; \mathbb{Z})/\text{Tors}) \subset \Lambda^*(H^1(\Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta)); \mathbb{Z})$ 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}_α to \mathbb{T}_β . We define a map

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

which lowers the degree by one, by the formula

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

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 4.17, which we break into several lemmas.

LEMMA 4.18. *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

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

(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 4.19. *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}_α to \mathbb{T}_β . Then, it follows from the definitions that

$$A_\zeta = \partial \circ H - H \circ \partial. \quad \square$$

Proof of Proposition 4.17. Together, Lemmas 4.18 and 4.19 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 $f: \Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta) \rightarrow S^1$ representing ζ . Given a generic point $p \in S^1$, we let $V = f^{-1}(p)$, so that the action of ζ is given by

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

where

$$a(\zeta, \phi) = \#\{u \in \mathcal{M}(\phi) \mid u([0, 1] \times \{0\}) \in V\}.$$

Fix a homotopy class $\phi \in \pi_2(\mathbf{x}, \mathbf{w})$ with $\mu(\phi) = 2$. We consider the one-manifold

$$\Xi = \left\{ s \in [0, \infty), u \in \mathcal{M}(\phi) \mid u([0, 1] \times \{s\}) \in V, \quad u([0, 1] \times \{-s\}) \in V' \right\}.$$

where V, V' are the pre-images of p and p' under f . Choosing generic p and 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

$$\left\{ u_1 \in \mathcal{M}(\phi_1) \mid u_1([0, 1] \times \{0\}) \in V \right\} \times \left\{ u_2 \in \mathcal{M}(\phi_2) \mid u_2([0, 1] \times \{0\}) \in V' \right\},$$

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 $H_1(Y; \mathbb{Z})/\text{Tors}$ 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 4.19 preserve $CF^-(Y, \mathfrak{s})$, so it follows that A_ζ induces actions on HF^+ and HF^- . The action on \widehat{HF} is defined in an analogous manner, as well. \square

Although this gives an action of all of $H^1(\Omega(\mathbb{T}_\alpha, \mathbb{T}_\beta); \mathbb{Z})$ on the Floer homologies, the interesting part of the action is induced from $H_1(Y; \mathbb{Z})$ (cf. the isomorphism from Equation (13)): it is a straightforward verification that when $g > 2$, the additional \mathbb{Z} summand acts trivially on all the Floer homology groups.

Remark 4.20. A geometric realization of the action of $H_1(Y; \mathbb{Z})/\text{Tors}$ 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,

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

where

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

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

5. Special Heegaard moves

In Section 4, for three-manifolds with $b_1(Y) > 0$, we required that the Heegaard diagram satisfy additional admissibility hypotheses. The purpose of this section is to construct such Heegaard diagrams, and indeed to show that any admissible diagrams are isotopic through such diagrams.

To this end, 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 γ , where it satisfies the property that $d\theta(X) > 0$, with respect to a coordinate system $(t, \theta) \in (-\varepsilon, \varepsilon) \times S^1$ 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, 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 2.




Figure 2: 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' .

We have the following notion:

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

$$(\Sigma, \alpha, \beta, z)$$

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 (Σ, α, β) for Y and let γ be a collection of pairwise disjoint curves which are dual to the α , 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 β if necessary, we can arrange that $\mathbb{T}_\beta \cap \mathbb{T}_\gamma \neq \emptyset$. Choose a basepoint z distinct from α , β , and γ (indeed, choose z to be disjoint from the neighborhood of the γ where the winding is performed).

Let $\mathbf{x} = \{x_1, \dots, x_g\} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, labeled so that $x_i \in \beta_i \cap \gamma_i$ for $i = 1, \dots, g$. 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 $x_i^\pm(1), x_i^\pm(2), \dots, x_i^\pm(k)$ (ordered in decreasing order of their distance to x_i , and with the sign distinguishing which side of γ_i — in its tubular neighborhood — they lie in). Thus, we have induced intersection points

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

labeled by $i_1, \dots, i_g \in 1, \dots, k$. 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 γ_i in Σ .

No matter how many times we wind α_i along γ_i , the Spin^c structure of the farthest intersection point $\mathbf{x}(1, \dots, 1)$ remains fixed (this is clear from the definition of $s_z(\mathbf{x})$: the winding isotopy induces an isotopy between the induced nonvanishing vector fields induced over Y). Moreover, by the definition of the difference map ε introduced in Subsection 2.4, together with Lemma 2.19, we have that

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

Thus, we can find Heegaard diagrams which realize the Spin^c structures which differ from some fixed Spin^c structure \mathfrak{s}_0 by nonpositive multiples of the $[\gamma_1], \dots, [\gamma_g]$. Moreover, if we choose parallel copies $\{\gamma_1^-, \dots, \gamma_g^-\}$ of the γ , 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 $[\gamma_1], \dots, [\gamma_g]$. Now, it is easy to see that the group $H^2(Y; \mathbb{Z})$ is generated by the Poincaré duals of the γ . Hence, we can realize all Spin^c structures. \square

Winding can be used also to arrange for strong admissibility. For this, 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 α and β (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 α sufficiently many times along curves γ 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.

We find it convenient to use rational coefficients for our periodic domains. Write $b = b_1(Y)$, and choose a basis $\{\mathcal{Q}_1, \dots, \mathcal{Q}_b\}$ for the \mathbb{Q} -vector space of renormalized periodic domains. Note that a renormalized periodic domain \mathcal{Q} is uniquely determined by the part of its boundary which is spanned by $\{[\alpha_1], \dots, [\alpha_g]\}$; i.e. the map which associates to a periodic domain the α -coefficients of its boundary gives an injection of vector spaces. After a change of

basis of the $\{\mathcal{Q}_i\}$ and reordering the α , we can assume that for all $i = 1, \dots, b$,

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

where $a_{i,j}, b_{i,j} \in \mathbb{Q}$.

For each $i = 1, \dots, b$ choose points $w_i \in \gamma_i$ which are not contained in any of the α or β . Let

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

and then choose some integer N so that

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

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

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

In a similar manner, we see that

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

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. \square

Indeed, an elaboration of this argument gives the following refinement. But first, we give a definition.

Recall that a (generic) pointed isotopy between two Heegaard diagrams can be subdivided into a sequence of isotopies where, at each stage, there is one curve $\alpha_i \in \alpha$ and one curve $\beta_j \in \beta$ whose number of intersection points either increases by two (pair creation) or drops by two (pair annihilation), while all other curves α_k and β_ℓ when $(i, j) \neq (k, \ell)$ remain transverse throughout the isotopy.

Definition 5.5. Two strongly \mathfrak{s} -admissible pointed Heegaard diagrams are said to be *strongly \mathfrak{s} -isotopic* if all the intermediate Heegaard diagrams in the isotopy are also strongly \mathfrak{s} -admissible.

LEMMA 5.6. *Suppose that two strongly \mathfrak{s} -admissible pointed Heegaard diagrams are isotopic (via an isotopy which does not cross the basepoint z); then they are strongly \mathfrak{s} -isotopic.*

Proof. First, note that if $(\Sigma, \alpha, \beta, z)$ is strongly \mathfrak{s} -admissible, then if we choose curves along which to wind the α (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 coefficients, 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 β — taking β to β'), then if we wind their α -curves simultaneously along some collection of γ to obtain α' , then the pointed Heegaard diagrams $(\Sigma, \alpha', \beta, z)$ and $(\Sigma, \alpha', \beta', z)$ 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

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

where we think of $t \in [0, 1]$ as the parameter in some isotopy taking β to β' , 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. \square

Remark 5.7. Note that this lemma also proves that any two isotopic \mathfrak{s} -realized pointed Heegaard diagrams are isotopic through \mathfrak{s} -realized Heegaard diagrams.

Suppose that two weakly \mathfrak{s} -admissible pointed Heegaard diagrams are isotopic (via an isotopy which does not cross the basepoint z); then they are said to be *weakly \mathfrak{s} -isotopic*.

LEMMA 5.8. *Any weakly \mathfrak{s} -admissible Heegaard diagram is weakly \mathfrak{s} -isotopic to a strongly \mathfrak{s} -admissible Heegaard diagram.*

Proof. It is easy to see that if we take a weakly \mathfrak{s} -admissible Heegaard diagram, then Lemma 5.4 provides an isotopy to a strongly \mathfrak{s} -admissible Heegaard diagram, and that the given isotopy is an isotopy through weakly \mathfrak{s} -admissible Heegaard diagrams. \square

6. Independence of complex structures

In Section 4, we defined various chain complexes, whose definition required a Heegaard diagram (satisfying appropriate admissibility hypotheses), and analytical choices — a complex structure on Σ , and a one-parameter perturbation J . Our aim here is to prove that the homology groups of the chain complexes are independent of the latter choices, and hence define an invariant of pointed Heegaard diagrams. More precisely, we prove the following:

THEOREM 6.1. *Let Y be a closed, oriented three-manifold equipped with a Spin^c structure \mathfrak{s} , and let $(\Sigma, \alpha, \beta, z)$ be a strongly \mathfrak{s} -admissible Heegaard diagram, endowed with an equivalence class \mathfrak{o} of coherent orientation system. Then, the homology groups $\widehat{HF}(\alpha, \beta, \mathfrak{s}, \mathfrak{o})$, $HF^\pm(\alpha, \beta, \mathfrak{s}, \mathfrak{o})$ and $HF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o})$, thought of as modules over $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(Y; \mathbb{Z})/\text{Tors}$, are independent of the choice of complex structure j on Σ and the path J_s .*

We first prove Theorem 6.1 in the case where Y is a rational homology three-sphere:

Proof of Theorem 6.1 when $b_1(Y) = 0$. First, we argue that if we fix j over Σ , the Floer homology groups are independent of the choice of (η, j, V) -nearly symmetric path J_s (in \mathcal{U}). Suppose we have two paths $J_s(0)$ and $J_s(1)$ in \mathcal{U} . Since \mathcal{U} is simply connected, we can connect them by a two-parameter family $J: [0, 1] \times [0, 1] \rightarrow \mathcal{U}$, thought of as a one-parameter family of paths indexed by $t \in [0, 1]$, writing $J_s(t)$ for the path obtained by fixing t . In fact, we can arrange that $J_s(t)$ is independent of t for t near 0 and 1, so that $J_s(t)$ can be naturally extended to all $t \in \mathbb{R}$. Then (as is familiar in Floer theory) we can define an associated chain map

$$\Phi_{J_{s,t}}^\infty : (CF^\infty(\alpha, \beta, \mathfrak{s}), \partial_{J_{s(0)}}^\infty) \longrightarrow (CF^\infty(\alpha, \beta, \mathfrak{s}), \partial_{J_{s(1)}}^\infty)$$

by

$$(14) \quad \Phi_{J_{s,t}}^\infty[\mathbf{x}, i] = \sum_{\mathbf{y}} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi)=0\}} \#(\mathcal{M}_{J_{s,t}}(\phi)) \cdot [\mathbf{y}, i - n_z(\phi)],$$

where $\mathcal{M}_{J_{s,t}}(\phi)$ denotes holomorphic disks with a time-dependent complex structure on the target, i.e. $\mathcal{M}_{J_{s,t}}(\phi)$ consists of maps

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

which represent the homotopy class ϕ . Note that the argument of Lemma 3.5 still goes over to give an energy bound on this moduli space.

The usual arguments from Floer theory then apply to show that $\Phi_{J_{s,t}}^\infty$ is a chain map which induces an isomorphism in homology. We outline these briefly.

The transversality theorem (Theorem 3.4) shows that for a generic path $J_s(t)$, the zero-dimensional components of the moduli spaces $\mathcal{M}_{J_{s,t}}(\mathbf{x}, \mathbf{y})$ are smoothly cut out and compact, as in the proof of Theorem 3.18. Thus, the map $\Phi_{J_{s,t}}^\infty$ is well-defined. To show that it is a chain map, we consider the ends of the one-dimensional moduli spaces $\mathcal{M}_{J_{s,t}}(\psi)$ with $\mu(\psi) = 1$. We claim that the only ends of these moduli spaces correspond to products

$$\left(\coprod_{\phi * \phi' = \psi} \mathcal{M}_{J_{s,t}}(\phi) \times \widehat{\mathcal{M}}_{J_s(1)}(\phi') \right) \coprod \left(\coprod_{\phi' * \phi = \psi} \widehat{\mathcal{M}}_{J_s(0)}(\phi') \times \mathcal{M}_{J_{s,t}}(\phi) \right),$$

where the ϕ homotopy classes all have $\mu(\phi) = 0$ and the ϕ' have $\mu(\phi') = 1$. Counted with sign, these represent the coefficients of $\partial_{J_s(1)}^\infty \circ \Phi_{J_{s,t}}^\infty - \Phi_{J_{s,t}}^\infty \circ \partial_{J_s(0)}^\infty$. This follows from Gromov's compactness, together with the observation that there can be no spheres bubbling off or boundary degenerations, as both carry Maslov index at least two. Hence, $\Phi_{J_{s,t}}^\infty$ is a chain map.

To see that $\Phi_{J_{s,t}}^\infty$ induces an isomorphism in homology, we show that the composite $\Phi_{J_{s,t}}^\infty \circ \Phi_{J_{s,1-t}}^\infty$ is chain homotopic to the identity map. The chain homotopy is constructed using a homotopy $J_{s,t,\tau}$ between two two-parameter families of complex structures (once again, we let $J_{s,t}(\tau)$ denote the two-parameter family in s and t , with $\tau \in [0, 1]$ fixed); i.e. $J_{s,t}(0)$ is the family of complex structures obtained by juxtaposing $J_{s,t}$ with $J_{s,1-t}$, while $J_{s,t}(1) = J_s(0)$ is independent of t . We can define a moduli space

$$\mathcal{M}_{J_{s,t,\tau}}(\phi) = \bigcup_{\tau \in [0,1]} \mathcal{M}_{J_{s,t}(\tau)}(\phi),$$

for each fixed homotopy class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. For generic $J_{s,t,\tau}$, this is a manifold of dimension $\mu(\phi) + 1$. We define a map

$$H_{J_{s,t,\tau}}^\infty([\mathbf{x}, i]) = \sum_{\mathbf{y}} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) | \mu(\phi) = -1\}} \#(\mathcal{M}_{J_{s,t,\tau}}(\phi)) \cdot [\mathbf{y}, i - n_z(\phi)].$$

Note that $H_{J_{s,t,\tau}}^\infty$ has degree -1 . To see, then that $H_{J_{s,t,\tau}}^\infty$ is the chain homotopy between $\Phi_{J_t}^\infty \circ \Phi_{J_{1-t}}^\infty$ and the identity map, we consider ends of the moduli spaces $\mathcal{M}_{J_{s,t,\tau}}(\psi)$ with $\mu(\psi) = 0$. These have three kinds of ends: the $\tau = 0$ end, which corresponds to the composite $\Phi_{J_t}^\infty \circ \Phi_{J_{1-t}}^\infty$, the $\tau = 1$ end, corresponding to the identity map, and those ends which correspond to splittings

$$\left(\coprod_{\phi * \phi' = \psi} \mathcal{M}_{J_{s,t,\tau}}(\phi) \times \widehat{\mathcal{M}}_{J_s(0)}(\phi') \right) \coprod \left(\coprod_{\phi' * \phi = \psi} \widehat{\mathcal{M}}_{J_s(0)}(\phi') \times \mathcal{M}_{J_{s,t,\tau}}(\phi) \right),$$

where $\mu(\phi) = -1$, and $\mu(\phi') = 1$, i.e. this corresponds to $\partial_{J_s(0)}^\infty \circ H_{J_{s,t,\tau}}^\infty - H_{J_{s,t,\tau}}^\infty \circ \partial_{J_s(0)}^\infty$. Once again, there are no other possible bubbles generically.

Similar chain maps can be defined on \widehat{HF} , HF^+ and HF^- as well. For \widehat{HF} , the corresponding map $\widehat{\Phi}_{J_{s,t}}$ counts ϕ only if $n_z(\phi) = 0$. For HF^+ and HF^- , we can let $\Phi_{J_{s,t}}^+$ and $\Phi_{J_{s,t}}^-$ be the induced maps from $\Phi_{J_{s,t}}^\infty$, because ϕ admits no holomorphic representative if $n_z(\phi) < 0$. It is clear from their definition that Φ^∞ , Φ^+ , and Φ^- commute with the corresponding U -actions.

As a consequence, we see that the Floer homologies for fixed j are independent of the choice of generic (η, j, V) -nearly symmetric path. Thus, it follows also that the homology groups do not depend on the η and V .

Next, we see that the chain complex is independent of the complex structure j on the Riemann surface. To this end, we observe that the chain complexes remain unchanged under small perturbations of the path of almost-complex structures J_s , provided that we still have *a priori* energy bounds after the perturbation. Furthermore, we can approximate a j -nearly-symmetric path J_s by j' -nearly-symmetric paths J'_s , with j' close to j . This shows that the Floer homology is also independent of the choice of j , since the space of allowed complex structures over Σ is connected (see Lemma 3.13), as it is obtained from the space of all complex structures by removing a codimension two subset. \square

We turn attention to Theorem 6.1 in the case where $b_1(Y) > 0$. We assume that the pointed Heegaard diagram $(\Sigma, \alpha, \beta, z)$ is strongly \mathfrak{s} -admissible.

As in the proof when $b_1(Y) = 0$, consider a fixed complex structure j over Σ , and let $J_{s,t}$ be a one-parameter family of (η, j, V) -nearly symmetric paths in \mathcal{U} . Note that the nonnegativity result of Lemma 3.2 applied to the parametrized moduli spaces $\mathcal{M}_{J_{s,t}}(\phi)$, together with admissibility and Lemma 4.14, ensure that the sum defining $\Phi_{J_{s,t}}^\infty$ from Equation (14) is a finite sum. To see that the map respects the module structure, we have the following:

LEMMA 6.2. *For any $\zeta \in H_1(Y, \mathbb{Z})/\text{Tors}$,*

$$A_\zeta \circ (\Phi_{J_{s,t}}^\infty) = (\Phi_{J_{s,t}}^\infty) \circ A_\zeta$$

as a map from $H_(CF^\infty(\mathbb{T}_\alpha, \mathbb{T}_\beta), \partial_{J_{s(0)}}^\infty) \longrightarrow H_*(CF^\infty(\mathbb{T}_\alpha, \mathbb{T}_\beta), \partial_{J_{s(1)}}^\infty)$.*

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

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

defined by

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

We claim that

$$(15) \quad A_\zeta \circ \Phi_{J_{s,t}}^\infty - \Phi_{J_{s,t}}^\infty \circ A_\zeta = \partial_{J_s(1)} \circ h - h \circ \partial_{J_s(0)}.$$

This follows from consideration of the ends of the one-dimensional moduli spaces

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

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 (15), of course, says that A_ζ commutes with $\Phi_{J_{s,t}}^\infty$, on the level of homology. \square

Proof of Theorem 6.1 when $b_1(Y) > 0$. The proof proceeds exactly as in the case where $b_1(Y) = 0$. Lemma 6.2 is used to prove that the induced isomorphisms of Floer homologies are $\Lambda^*H_1(Y; \mathbb{Z})/\text{Tors}$ -module isomorphisms. \square


7. First steps towards topological invariance

7.1. Overview of topological invariance. According to Theorem 6.1, the Floer homology groups as defined in Section 4 are invariants of strongly \mathfrak{s} -admissible pointed Heegaard diagrams.

Loosely speaking, to show they are actually invariants of the underlying three-manifold, we must show that they are invariant under the three basic moves: isotopies, handleslides, and stabilizations. In fact, the invariants we have are invariants of *pointed* Heegaard diagrams, and we must restrict ourselves to moves which never cross the basepoint z (i.e. the isotopies do not cross z , and the pairs of pants involved in the handleslides do not contain z either). As we shall see in Proposition 7.1, an invariant of Heegaard diagrams which is invariant under these restricted moves still gives a topological invariant, since we can trade an isotopy which crosses the basepoint z for a sequence of handleslides which do not.

In fact, when $b_1(Y) > 0$, we allow ourselves only strongly \mathfrak{s} -admissible Heegaard diagrams for Y and allow ourselves an even more restricted set of Heegaard moves which connect such diagrams. By adapting the arguments from Section 5 (see Proposition 7.2 below), we see that a quantity associated to admissible Heegaard diagrams which is invariant under only these kinds of Heegaard moves still gives a topological invariant of the underlying three-manifold.

After establishing these topological preliminaries in Section 7.2, we establish isotopy invariance of the Floer homologies in Subsection 7.3 (and a corresponding version for the weakly \mathfrak{s} -admissible required to define \widehat{HF} and

Figure 3: Moving β_1 to $\bar{\beta}_1$ in $\Sigma - z$

HF^+ in Subsection 7.4), returning to handleslide invariance in Section 9 and stabilization invariance in Section 10.

7.2. Topological invariants and special Heegaard moves. A quantity associated to pointed Heegaard diagrams which is invariant under pointed Heegaard moves is a three-manifold invariant, according to the following:

PROPOSITION 7.1. *Any two Heegaard diagrams $(\Sigma, \alpha, \beta, z)$ and $(\Sigma', \alpha', \beta', z')$ which specify the same three-manifold are diffeomorphic after a finite sequence of pointed Heegaard moves (i.e. Heegaard moves supported in the complement of the basepoint).*

Proof. Given a sequence of Heegaard moves, we can clearly introduce isotopies as needed to arrange that no handleslides, only isotopies (of the α or the β), cross the basepoint z .

Since the roles of α, β are symmetric, it suffices to consider the case where the isotopy of β_i , say β_1 , crosses z . We denote the new isotopic curve by $\bar{\beta}_1$. We claim that β_1 can be moved by a series of handleslides and isotopies to $\bar{\beta}_1$ all of which are supported in $\Sigma - z$. This can be seen by first surgering out β_2, \dots, β_g to get a torus T^2 with $2g - 2$ marked points. Clearly, in $T^2 - z$ the curves induced by β_1 and $\bar{\beta}_1$ are isotopic. We can follow the isotopy by moves $\Sigma - z$ where isotopies across the marked points $T^2 - z$ are replaced by handleslides in $\Sigma - z$. See Figure 3 for the $g = 3$ case. \square

When $b_1(Y) > 0$, in view of the additional admissibility hypotheses on the Heegaard diagrams, we need the following refinement of Proposition 7.1:

PROPOSITION 7.2. *Let Y be a three-manifold, equipped with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$. Then, there is strongly \mathfrak{s} -admissible pointed Heegaard diagram for Y , and any two strongly \mathfrak{s} -admissible pointed Heegaard diagrams can*

be connected by a finite sequence of pointed Heegaard moves where, at each intermediate stage, the Heegaard diagrams are all strongly \mathfrak{s} -admissible pointed Heegaard diagrams.

Proof. Existence is ensured by Lemma 5.4. Given two strongly \mathfrak{s} -admissible Heegaard diagrams, Proposition 7.1 gives a sequence of pointed Heegaard moves which connect them. Now, by introducing additional isotopies as in Lemma 5.6, we can arrange for all the isotopies to go through strongly \mathfrak{s} -isotopic Heegaard diagrams. \square

7.3. Isotopy invariance. Theorem 6.1 implies that the Floer homologies remain unchanged under isotopies of the attaching circles which preserve the transversality of the α and β . To show isotopy invariance of the homology groups in general, we must have a larger class of isotopies which allows us to introduce new intersections between the attaching circles, in a controlled manner.

Such moves are provided by exact Hamiltonian motions of the attaching circles in Σ . Recall that on a symplectic manifold, a one-parameter family of real-valued functions H_t naturally gives rise to a unique one-parameter family of Hamiltonian vector fields X_t , specified by

$$\omega(X_t, \cdot) = dH_t$$

where the left-hand side denotes the contraction of the symplectic form ω with the vector field. A one-parameter family of diffeomorphisms Ψ_t is said to be an exact Hamiltonian isotopy if it is obtained by integrating a Hamiltonian vector field, i.e. if $\Psi_0 = \text{Id}$, and

$$\frac{d\Psi_t}{dt} = X_t.$$

By taking a positive bump function h supported in a neighborhood of a point which lies on α_1 , and letting $f: \mathbb{R} \rightarrow [0, 1]$ be a nonnegative smooth function whose support is $(0, 1)$ we can consider Hamiltonian $H_t = f(t)h$. The corresponding diffeomorphism moves the curve α_1 slightly (without moving any of the other α -curves). See Figure 4 for an illustration. (The picture takes place in a small Euclidean patch of Σ , which meets two curves α and β ; the isotopy is used to displace α , to give a new curve, α' .)

With these preliminaries in place we state the main result of this subsection:

THEOREM 7.3. *If Y is a three-manifold equipped with a Spin^c structure \mathfrak{s} , and endowed with two isotopic \mathfrak{s} -strongly admissible Heegaard diagrams, then there is an identification between equivalence classes of orientation systems for the two diagrams and corresponding identification between the homology groups*




Figure 4: Moving the curve α to α' by a Hamiltonian isotopy. In this manner, we can introduce a pair of canceling intersection points between α' and the curve β .

(thought of as $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(Y; \mathbb{Z}) / \text{Tors-modules}$)

$$\begin{aligned} \widehat{HF}(\Sigma, \alpha, \beta, \mathfrak{o}) &\cong \widehat{HF}(\Sigma', \alpha', \beta', \mathfrak{o}'), \\ HF^{\pm}(\Sigma, \alpha, \beta, \mathfrak{o}) &\cong HF^{\pm}(\Sigma', \alpha', \beta', \mathfrak{o}'), \\ HF^{\infty}(\Sigma, \alpha, \beta, \mathfrak{o}) &\cong HF^{\infty}(\Sigma', \alpha', \beta', \mathfrak{o}'). \end{aligned}$$

Proof of Theorem 7.3 when $b_1(Y) = 0$. First of all, isotopies which preserve the condition that the α are transverse to the β can be thought of as variations in the metric (or equivalently, the complex structure j) on Σ . Hence, the invariance of the Floer homology under such isotopies follows from Theorem 6.1.

It suffices then to show that the homology remains unchanged when a pair of canceling intersection points between, say, α_1 and β_1 is introduced. Such an isotopy can be realized by moving the α by an exact Hamiltonian diffeomorphism of Σ . We assume that the exact Hamiltonian is supported on $\Sigma - \alpha_2 - \cdots - \alpha_g - z$ (i.e. that the corresponding vector fields X_t are supported on the subset), and that the Hamiltonian is supported in $t \in [0, 1]$. The isotopy $\{\Psi_t\}$ induces an isotopy of \mathbb{T}_{α} . We must show that this isotopy of \mathbb{T}_{α} induces a map on Floer homology, by imitating the usual constructions from Lagrangian Floer theory.

The map is induced by counting points in the zero-dimensional components of the moduli spaces with a time-dependent constraint. To be precise, here, we will be using a fixed one-parameter family J of nearly-symmetric complex structure. Recall that for this notion, we needed to fix a collection of points $\{z_i\}_{i=1}^m \subset \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$. For the present argument, it suffices to choose $\{z_i\}_{i=1}^m = \{z\}$.

Now, we have homotopy classes of disks $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$, where $\mathbf{x} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$ and $\mathbf{y} \in \Psi_1(\mathbb{T}_{\alpha}) \cap \mathbb{T}_{\beta}$, which denote homotopy classes of maps $u: \mathbb{D} \longrightarrow \text{Sym}^g(\Sigma)$ satisfying

$$(16) \quad \left\{ u: \mathbb{D} \longrightarrow \text{Sym}^g(\Sigma) \left| \begin{array}{l} u(1+it) \in \Psi_t(\mathbb{T}_{\alpha}), \forall t \in \mathbb{R}, \\ u(0+it) \in \mathbb{T}_{\beta}, \forall t \in \mathbb{R}, \\ \lim_{t \rightarrow -\infty} u(s+it) = \mathbf{x}, \\ \lim_{t \rightarrow +\infty} u(s+it) = \mathbf{y} \end{array} \right. \right\}$$

(dividing out by the action of π_1 in the case where $g = 2$). We think of this set as a set of homotopy class of Whitney disks with “dynamic boundary conditions.” Fixing such a class $\phi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$, we have the moduli space of maps $\mathcal{M}^{\Psi_t}(\phi)$ satisfying the above boundary conditions and also $\bar{\partial}_{J_s} u = 0$. It is easy to see that if $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}) \neq \emptyset$, then $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $\mathbf{y} \in \Psi_1(\mathbb{T}_\alpha) \cap \mathbb{T}_\beta$ lie in equivalence classes \mathfrak{S} and \mathfrak{S}' corresponding to the same Spin^c structure \mathfrak{s} , for the fixed base-point z . Moreover, in this case $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}) \cong \mathbb{Z}$. Note that by juxtaposition, a single $\phi_0 \in \pi_2^{\Psi_t}(\mathbf{x}_0, \mathbf{y}_0)$ naturally gives rise to an identification between homotopy classes $\pi_2(\mathbf{x}_0, \mathbf{x}_0) \cong \pi_2(\mathbf{y}_0, \mathbf{y}_0)$ and a corresponding identification between isomorphism classes of coherent orientation systems for $(\Sigma, \alpha, \beta, \mathbf{z})$ and $(\Sigma, \alpha', \beta, \mathbf{z})$. (This point is more important, of course, in the case where $b_1(Y) > 0$.)

Now, we define the map associated to the isotopy

$$\Gamma_{\Psi_t}^\infty : CF^\infty(\alpha, \beta, z) \longrightarrow CF^\infty(\alpha', \beta, z)$$

by the formula

$$(17) \quad \Gamma_{\Psi_t}^\infty([\mathbf{x}, i]) = \sum_{\{\mathbf{y} \in \mathfrak{S}'\}} \sum_{\{\phi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}) \mid \mu(\phi)=0\}} \#(\mathcal{M}^{\Psi_t}(\mathbf{x}, \mathbf{y})) \cdot [\mathbf{y}, i - n_z(\phi)],$$

where $\mu(\phi)$ is the expected dimension of the moduli space $\mathcal{M}_{\Psi_t}(\phi)$, and the number $\#\mathcal{M}^{\Psi_t}(\phi)$ is the signed number of points in this zero-dimensional moduli space. Note that this is a finite sum for each given $[\mathbf{x}, i]$, since $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}) \cong \mathbb{Z}$. The map is well-defined up to one overall sign.

The important observation is that the moduli spaces considered have Gromov compactifications. This follows from the energy bounds on the moduli spaces of disks: given a class $\phi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$ with $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $\mathbf{y} \in \Psi_1(\mathbb{T}_\alpha) \cap \mathbb{T}_\beta$, we must show that there is a bound (depending only on ϕ) on the energy of any holomorphic disk $u \in \mathcal{M}^{\Psi_t}(\phi)$.

Specifically, suppose that u_0 is some fixed representative of the homotopy class ϕ . We will find that there is some fixed constant which bounds the energy of another such representative u_1 . Specifically, since u_0 and u_1 are homotopic holomorphic disks which connect \mathbf{x} to \mathbf{y} , i.e. there is a homotopy

$$U : \mathbb{D} \times [0, 1] \longrightarrow \text{Sym}^g(\Sigma),$$

with $U(s+it, 0) = u_0(s+it)$, $U(s+it, 1) = u_1(s+it)$, and $U(1+it, \tau) \in \Psi_t(\mathbb{T}_\alpha)$, $U(0+it, \tau) \in \mathbb{T}_\beta$. Then, there are corresponding lifts

$$\tilde{u}_0 : F_0 \longrightarrow \Sigma^{\times g} \text{ and } \tilde{u}_1 : F_1 \longrightarrow \Sigma^{\times g}$$

coming from pulling back the branched covering $\pi : \Sigma^{\times g} \longrightarrow \text{Sym}^g(\Sigma)$ as in Subsection 3.4. The homotopy of u_0 and u_1 lifts to give a homology between \tilde{u}_0 and \tilde{u}_1 , so that if ω_0 denotes the product form on $\Sigma^{\times g}$, the difference

$$\int_{F_1} \tilde{u}_1^* \omega_0 - \int_{F_0} \tilde{u}_0^* \omega_0$$

is calculated by integrating the pull-back of ω_0 over the cover (in $\Sigma^{\times g}$) of the restriction $U|(\partial\mathbb{D}) \times [0, 1]$. Indeed, the form ω_0 vanishes over the image of $\{0\} \times \mathbb{R} \subset \partial\mathbb{D}$, so that we need to bound the integral of the pull-back of ω_0 via the $g!$ maps

$$f: \mathbb{R} \times [0, 1] \longrightarrow \Sigma^{\times g},$$

given by

$$f(t, \tau) = \tilde{U}(1 + it, \tau)$$

(where $\mathbb{R} \times [0, 1]$ is one of the $g!$ components of the covering space of $\{1\} \times \mathbb{R} \times [0, 1] \subset \mathbb{D} \times [0, 1]$ induced from $\Sigma^{\times g}$). Thus, the map f satisfies $f(t, \tau) \in \Psi_t(\widetilde{\mathbb{T}}_\alpha)$ (where $\widetilde{T}a$ is the pre-image of \mathbb{T}_α in $\Sigma^{\times g}$) for all t, τ . Taking derivatives of this condition, we get that

$$\frac{df(t, \tau)}{dt} \equiv X_t f(t, \tau) \pmod{(\Psi_t)_*(T\widetilde{\mathbb{T}}_\alpha)}$$

and

$$\frac{df(t, \tau)}{d\tau} \equiv 0 \pmod{(\Psi_t)_*(T\widetilde{\mathbb{T}}_\alpha)}.$$

In view of the fact that ω_0 vanishes on all tangent spaces $(\Psi_t)_*(T\widetilde{\mathbb{T}}_\alpha)$,

$$\begin{aligned} \int f^* \omega_0 &= \int \omega_0 \left(\frac{df}{dt}, \frac{df}{d\tau} \right) dt \wedge d\tau \\ &= \int \omega_0 \left(X_t, \frac{df}{d\tau} \right) \\ &= \int \left\langle dH_t, \frac{df}{d\tau} \right\rangle dt \wedge d\tau \\ &= \int (H_t(f(t, 1)) - H_t(f(t, 0))) dt. \end{aligned}$$

Since H_t is identically zero outside $t \in [0, 1]$, the above integral is bounded by a constant

$$k = \sup_{(t, \mathbf{w}) \in \mathbb{R} \times \text{Sym}^g(\Sigma)} H_t(\mathbf{w}) - \inf_{(t, \mathbf{w}) \in \mathbb{R} \times \text{Sym}^g(\Sigma)} H_t(\mathbf{w})$$

which is independent of our choice of u_0 and u_1 . Thus, we see that

$$\int_{F_1} \tilde{u}_1^*(\omega_0) - \int_{F_0} \tilde{u}_0^*(\omega_0) \leq k.$$

But now, in view of the proof of Lemma 3.5, we see that if ω_1 denotes the symplectic form on $\text{Sym}^g(\Sigma)$ (which tames the product complex structure),

then

$$\begin{aligned} \int_{\mathbb{D}} u_1^*(\omega_1) &\leq C \int_{F_0} \tilde{u}_1(\omega_0) \\ &\leq C \left(\int_{F_0} \tilde{u}_0(\omega_0) + \left(\int_{F_1} \tilde{u}_1(\omega_0) - \int_{F_0} \tilde{u}_0(\omega_0) \right) \right) \\ &\leq C \left(\int_{F_0} \tilde{u}_0(\omega_0) + k \right) \end{aligned}$$

which in turn is independent of our choice of u_1 . (Indeed, for the case where $g = 2$, it is easy to see that the key property used for this bound is not that u_1 is homotopic to u_0 , but merely that the two maps induce the same relative homology class.)

With the energy bounds in place, now, we can show that Γ^∞ induces an isomorphism in HF^∞ .

As in the proof of Theorem 6.1 above, we verify that $\Gamma_{\Psi_t}^\infty$ is a chain map. Consider $\psi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = 1$. Then, $\mathcal{M}_{\Psi_t}(\psi)$ is a noncompact space, which can be compactified by:

$$\partial \mathcal{M}_{\Psi_t}(\psi) = \left(\coprod_{\phi * \phi' = \psi} \mathcal{M}_{\Psi_t}(\phi) \times \widehat{\mathcal{M}}(\phi') \right) \coprod \left(\coprod_{\phi' * \phi = \psi} \widehat{\mathcal{M}}(\phi') \times \mathcal{M}_{\Psi_t}(\phi) \right).$$

In the first decomposition, ϕ ranges over those elements of $\pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}')$ (where \mathbf{y}' is any point in $\Psi_1(\mathbb{T}_\alpha) \cap \mathbb{T}_\beta$) with $\mu(\phi) = 0$, while ϕ' ranges over those homotopy classes $\phi' \in \pi_2(\mathbf{y}', \mathbf{y})$ for the tori $\Psi_1(\mathbb{T}_\alpha)$ and \mathbb{T}_β , which satisfy $\mu(\phi') = 1$, and $\phi * \phi' = \psi$ (using the obvious juxtaposition operation). In the second, ϕ' ranges over those elements of $\pi_2(\mathbf{x}, \mathbf{x}')$ (where \mathbf{x}' is any point in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$) where $\mu(\phi') = 1$, while ϕ ranges over those elements of $\pi_2^{\Psi_t}(\mathbf{x}', \mathbf{y})$ with $\mu(\phi) = 0$, and $\phi' * \phi = \psi$. By counting points with sign, it follows that $\Gamma_{\Psi_t}^\infty$ is a chain map.

To see that $\Gamma_{\Psi_t}^\infty$ induces an isomorphism in homology, observe that the composite $\Gamma_{\Psi_t}^\infty \circ \Gamma_{\Psi_{1-t}}^\infty$ is chain homotopic to the identity map. The chain homotopy is constructed using a homotopy $\Phi_{t,\tau}$ between two one-parameter families of isotopies — thought of as $\tau \mapsto \Psi_t(\tau)$ — which connects the juxtaposition of Ψ_t with Ψ_{1-t} at $\tau = 0$, to the stationary identity isotopy at $\tau = 1$. Letting

$$\mathcal{M}^{\Phi_{t,\tau}}(\phi) = \bigcup_{\tau \in [0,1]} \mathcal{M}^{\Phi_t(\tau)}(\phi),$$

we define

$$H^\infty([\mathbf{x}, i]) = \sum_{\mathbf{y}} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid \mu(\phi) = -1\}} \#(\mathcal{M}^{\Phi_{t,\tau}}(\phi)) [\mathbf{y}, i - n_z(\phi)],$$

where we are implicitly using the identification $\pi_2(\mathbf{x}, \mathbf{y}) \cong \pi_2^{\Phi_t(\tau)}(\mathbf{x}, \mathbf{y})$ for each $\tau \in [0, 1]$. Note that H^∞ has degree -1 . To see, then that H^∞ is the chain homotopy between $\Gamma_{\Psi_t}^\infty \circ \Gamma_{\Psi_{1-t}}^\infty$ and the identity map, we consider ends of the moduli spaces $\mathcal{M}^{\Phi_t, \tau}(\psi)$, where ψ has $\mu(\psi) = 0$. These spaces have three kinds of ends: those at to $\tau = 0$, which correspond to the composite $\Gamma_{\Psi_t}^\infty \circ \Gamma_{\Psi_{1-t}}^\infty$, those at to $\tau = 1$, corresponding to the identity map, and those at the splittings

$$\left(\coprod_{\phi * \phi' = \psi} \mathcal{M}^{\Phi_t, \tau}(\phi) \times \widehat{\mathcal{M}}(\phi') \right) \coprod \left(\coprod_{\phi' * \phi = \psi} \widehat{\mathcal{M}}(\phi') \times \mathcal{M}^{\Phi_t, \tau}(\phi) \right);$$

where all ϕ satisfy $\mu(\phi) = 0$, and the ϕ' satisfy $\mu(\phi') = 1$. Counting these ends with sign, we obtain the relation

$$\Gamma_{\Psi_t}^\infty \circ \Gamma_{\Psi_{1-t}}^\infty = \text{Id} + \partial^\infty \circ H^\infty - H^\infty \circ \partial^\infty.$$

Switching the roles of Ψ_t and Ψ_{1-t} , it follows that Γ^∞ induces an isomorphism in homology.

This same technique applies to prove the result for \widehat{HF} , HF^+ , HF^- , and the U equivariance proves the result for HF_{red} . \square

The case where $b_1(Y) > 0$ proceeds much as before. However, special care must be taken to see that the sum appearing in the definition of the chain map induced by an isotopy of the α -curves, as given in Equation (17) is, in fact, a finite sum for each given $[\mathbf{x}, i]$. Again, this is done with the help of admissibility hypotheses.

In a single pair creation, the domains for (Σ, α, β) , and $(\Sigma, \{\psi(\alpha_1), \alpha_2, \dots, \alpha_g\}, \beta)$ 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 $\Psi_t(\mathbb{T}_\alpha)$ crosses the subvariety $\{w\} \times \text{Sym}^{g-1}(\Sigma) \subset \text{Sym}^g(\Sigma)$, 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 $u \cap (\{w_i\} \times \text{Sym}^{g-1}(\Sigma))$ (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 $\{w_i\} \times \text{Sym}^{g-1}(\Sigma)$). We call this collection of multiplicities the domain of ϕ . In fact, in our fixed one-parameter family of nearly-symmetric complex structures, we choose our basepoints $\{w_i\}_{i=1}^m \subset \Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$ with one in each of these domains \mathcal{D}_i .

LEMMA 7.4. *Fix $(\Sigma, \alpha, \beta, z)$ to be a strongly \mathfrak{s} -admissible pointed Heegaard diagram, and let Ψ_t be 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, \dots, 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}'_α be the torus $\psi_1(\alpha) \times \dots \times \alpha_g$. 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}'_α and \mathbb{T}_β , 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 $j = 1, \dots, m$.

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 $j = 1, \dots, m$ are nonnegative; it follows that for $j = 1, \dots, m$, $n_{w_j}(\psi_i * \phi_0) \geq n_{w_j}(\phi_0)$. 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 $i = 1, \dots, m$ are bounded below, and the Maslov index is fixed, so there can be only finitely many such homotopy classes, according to Lemma 4.14. The lemma now follows. \square

Proof of Theorem 7.3 when $b_1(Y) > 0$. In view of Lemma 7.4, it follows that $\Gamma_{\Psi_t}^\infty$ as defined above is a finite sum for each fixed $[\mathbf{x}, i]$, so that the earlier proof in the case where $b_1(Y) = 0$ applies. Note that the Lemma 7.4 also holds for isotopies obtained by juxtaposing Ψ_t with Ψ_{1-t} .

Establishing its $H_1(Y; \mathbb{Z})/\text{Tors}$ -equivariance of the map Φ^∞ follows as in the proof of Lemma 6.2 above. \square

7.4. Weakly admissible Heegaard diagrams.

THEOREM 7.5. *If Y is a three-manifold equipped with a Spin^c structure \mathfrak{s} , and endowed with a weakly \mathfrak{s} -admissible Heegaard diagram, then there is an isotopic strongly \mathfrak{s} -admissible Heegaard diagram, an identification between equivalence classes of orientation systems for the two diagrams and corresponding identification between the homology groups (thought of as $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(Y; \mathbb{Z})/\text{Tors-modules}$)*

$$\begin{aligned} \widehat{HF}(\Sigma, \alpha, \beta, \mathfrak{o}) &\cong \widehat{HF}(\Sigma', \alpha', \beta', \mathfrak{o}'), \\ HF^+(\Sigma, \alpha, \beta, \mathfrak{o}) &\cong HF^+(\Sigma', \alpha', \beta', \mathfrak{o}'). \end{aligned}$$

Proof. According to Lemma 5.8, we can find a weakly \mathfrak{s} -isotopic, strongly \mathfrak{s} -admissible Heegaard diagram. (Indeed, it is easy to give exact Hamiltonian isotopies of Σ which are suitable for the purposes of Lemma 5.4.) Note that the analogue of Lemma 7.4 also holds in the weakly admissible context (where now both $\mu(\psi)$ and $n_z(\psi)$ are fixed), so we can construct chain homotopy equivalences $\Gamma_{\Psi_t}^+$ and $\widehat{\Gamma}_{\Psi_t}$ by modifying Equation (17) (for example, in the definition of $\Gamma_{\Psi_t}^+$ we drop terms involving $[\mathbf{y}, j]$ with $j < 0$) showing that the corresponding groups are isomorphic. \square

8. 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. As we shall see in the sequel (cf. §9 of [28]), they are useful also when comparing the Floer homology groups of three-manifolds which differ by surgeries on a knot. More applications are also given in [29]. Thus, we allow ourselves now a lengthy digression into the properties of these maps.

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. In Subsection 8.2, we discuss issues of orientations. In Subsection 8.3, we discuss admissibility issues, and then set up the maps induced by holomorphic triangles, discussing various invariance properties of the maps. The maps enjoy a certain associativity property, which we will make use of in the proof of handleslide invariance, and in the sequel. This associativity is treated in Subsection 8.4; see also [12], [4].

8.1. Topological preliminaries on triangles. A *Heegaard triple-diagram of genus g* is an oriented two-manifold and three g -tuples α , β , and γ 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 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

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

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 8.1. Let (Σ, α, β) be a Heegaard diagram for Y , and let γ be a g -tuple of curves which are isotopic to β . Then the triple-diagram

$$(\Sigma, \alpha, \beta, \gamma)$$

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 $U_\beta \times \frac{1}{2}$.

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

PROPOSITION 8.2. *Let $\text{Span}([\alpha_i]_{i=1}^g) \subset H_1(\Sigma; \mathbb{Z})$ denote the lattice spanned by the one-dimensional homology classes induced by the α . Then, there are natural identifications*

(18)

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

or, equivalently,

(19)

$$H_2(X; \mathbb{Z}) \cong \text{Ker} \left(H_1(\mathbb{T}_\alpha; \mathbb{Z}) \oplus H_1(\mathbb{T}_\beta; \mathbb{Z}) \oplus H_1(\mathbb{T}_\gamma; \mathbb{Z}) \longrightarrow H_1(\text{Sym}^g(\Sigma); \mathbb{Z}) \right).$$

Similarly,

(20)

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

Proof. First, note that the boundary homomorphism $\partial: H_2(U_\alpha, \Sigma; \mathbb{Z}) \longrightarrow H_1(\Sigma; \mathbb{Z})$ is injective, and its image is $\text{Span}([\alpha_i]_{i=1}^g)$. 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 when 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 $H_1(\text{Sym}^g(\Sigma); \mathbb{Z}) \cong H_1(\Sigma; \mathbb{Z})$, the image of $H_1(\mathbb{T}_\alpha; \mathbb{Z})$ is identified with $\text{Span}([\alpha_i]_{i=1}^g)$.

The final isomorphism follows from the fact that

$$\begin{aligned} H_1(X, \Delta \times \Sigma) &\cong H_1(U_\alpha, \Sigma) \oplus H_1(U_\beta, \Sigma) \oplus H_1(U_\gamma, \Sigma) \\ &\cong H^2(U_\alpha) \oplus H^2(U_\beta) \oplus H^2(U_\gamma) = 0. \end{aligned}$$

□

Suppose $(a, b, c) \in \text{Span}([\alpha_i]_{i=1}^g) \oplus \text{Span}([\beta_i]_{i=1}^g) \oplus \text{Span}([\gamma_i]_{i=1}^g)$ 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 $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$ where we choose a reference point $z \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g - \gamma_1 - \cdots - \gamma_g$ is called a *pointed Heegaard triple-diagram*.

8.1.2. *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

$$u: \Delta \longrightarrow \text{Sym}^g(\Sigma)$$

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} . In the case where $g = 2$, this will actually denote equivalence classes of homotopy classes of Whitney triangles, where the equivalence relation identifies two homotopy classes which differ by the action of $\pi_1(\text{Sym}^g(\Sigma))$. (Note that one can alternately stabilize, as in Section 10 to reduce to the case where $g > 2$.)

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

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

which vanishes if $\pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ is nonempty. 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, $\varepsilon(\mathbf{x}, \mathbf{y}, \mathbf{w})$ is the equivalence class of the closed path $a + b + c$.

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

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

PROPOSITION 8.3. *When $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\mathbf{y} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, $\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, then $\pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ is nonempty if and only if $\varepsilon(\mathbf{x}, \mathbf{y}, \mathbf{w}) = 0$. Moreover, if $g > 1$ and $\varepsilon(\mathbf{x}, \mathbf{y}, \mathbf{w}) = 0$ then*

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

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

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

whose fiber is homotopy equivalent to the space of pointed maps from the sphere to $\text{Sym}^g(\Sigma)$ (the base space here is a product of path spaces). In the case where $g > 2$, this gives us an exact sequence

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

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

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

That image, in turn, is clearly identified with the kernel of the natural map $H_1(\mathbb{T}_\alpha) \oplus H_1(\mathbb{T}_\beta) \oplus H_1(\mathbb{T}_\gamma) \longrightarrow H_1(\text{Sym}^g(\Sigma))$ (we are using here the fact that $\pi_1(\text{Sym}^g(\Sigma))$ is Abelian). The proposition when $g > 2$ then follows from Proposition 8.2. The case where $g = 2$ follows as above, after dividing by the action of $\pi_1(\text{Sym}^g(\Sigma))$. \square

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).

8.1.3. 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 [19] and [13].

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 which 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^+ .

8.1.4. Triangles and Spin^c structures. The base-point $z \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g - \gamma_1 - \cdots - \gamma_g$ 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 points between $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ together with the basepoint z .

To describe this, fix “height functions” over the handlebodies $f_i: U_i \rightarrow [0, 1]$ where $i = \alpha, \beta$, 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 $u: \Delta \rightarrow \text{Sym}^g(\Sigma)$ representing $\phi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$, 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 straight 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 $\ell_{\alpha, \gamma}(t)$; see Figure 5.

There is a natural map $\pi: X \rightarrow \Delta$. The pre-image under π of $\ell_{\alpha, \beta}(t)$ for $t \in [0, 1]$, which we denote $\tilde{\ell}_{\alpha, \beta}(t)$, is identified with $Y_{\alpha, \beta}$. For all but finitely many t in the open interval, the intersection of F with $\tilde{\ell}_{\alpha, \beta}(t)$ 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 these generic t , we extend the oriented two-plane field over a neighborhood of these $g + 1$ paths (as in Subsection 2.6) in a continuous manner. In this way, we have extended \mathcal{L} across the intersection of F with $\tilde{\ell}_{\alpha, \beta}(t)$ for all but finitely many t .

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

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




Figure 5: Schematic for the cobordism X . We have illustrated the foliation of the triangle by segments whose pre-images are the three-manifolds $\tilde{\ell}_{\xi,\eta}(t)$.

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 pre-images 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 .

PROPOSITION 8.4. *The above construction induces a map*

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

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 (cf. Section 2.6). These data specify the two-plane field over $X_{\alpha,\beta,\gamma} - (\text{int}(\Delta) \times \Sigma) - \text{int}F_0 - \text{int}F_1$. 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 .

Indeed, it is easy to see from the above discussion that the restriction $u|_{\partial\Delta}$ determines the Spin^c structure. It follows that adding spheres to the homotopy class in $\pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ does not change the induced Spin^c structure. \square

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 8.5. *The map from Proposition 8.4 descends to a map*

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

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

Proof. Since addition of Whitney disks to a given homotopy class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w})$ changes the induced two-plane field in a neighborhood of the boundary, it follows that the above constructions descend to give a map

$$\mathfrak{s}_z : S_{\alpha, \beta, \gamma} \longrightarrow \text{Spin}^c(X_{\alpha, \beta, \gamma});$$

it remains to verify that we have identified the image, and also that \mathfrak{s}_z is one-to-one.

To see that we have characterized the image, recall that for an arbitrary four-manifold-with-boundary (X, Y) there is a canonical map $\varepsilon' : \text{Spin}^c(Y) \longrightarrow H^3(X, Y; \mathbb{Z})$ which is defined as follows. Choose a Spin^c structure \mathfrak{s}_0 over X , and let

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

where $\delta : H^2(Y; \mathbb{Z}) \longrightarrow H^3(X, Y; \mathbb{Z})$ 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 \mathbf{t} extends over X . Next, we argue that $\text{PD}[\varepsilon(\mathbf{x}, \mathbf{y}, \mathbf{w})] = \pm \varepsilon'(\mathbf{x}, \mathbf{y}, \mathbf{w})$. To see this, isotope \mathbb{T}_α , \mathbb{T}_β , and \mathbb{T}_γ so that there are intersection points \mathbf{x}' , \mathbf{y}' , and \mathbf{w}' for which $\varepsilon(\mathbf{x}', \mathbf{y}', \mathbf{w}') = 0$, so that there is a triangle connecting them. We have explicitly constructed the corresponding Spin^c structure, thus $\varepsilon'(\mathbf{x}', \mathbf{y}', \mathbf{w}') = 0$, as well. It is easy to see that

$$\varepsilon(\mathbf{x}, \mathbf{y}, \mathbf{w}) - \varepsilon(\mathbf{x}', \mathbf{y}', \mathbf{w}') = \pm \text{PD} \circ \delta \left(\text{PD}(\varepsilon(\mathbf{x}, \mathbf{x}')) \oplus \text{PD}(\varepsilon(\mathbf{y}, \mathbf{y}')) \oplus \text{PD}(\varepsilon(\mathbf{w}, \mathbf{w}')) \right).$$

Similarly,

$$\begin{aligned} \varepsilon'(\mathbf{x}, \mathbf{y}, \mathbf{w}) - \varepsilon'(\mathbf{x}', \mathbf{y}', \mathbf{w}') \\ = \delta \left((\mathfrak{s}_z(\mathbf{x}) - \mathfrak{s}_z(\mathbf{x}')) \oplus (\mathfrak{s}_z(\mathbf{y}) - \mathfrak{s}_z(\mathbf{y}')) \oplus (\mathfrak{s}_z(\mathbf{w}) - \mathfrak{s}_z(\mathbf{w}')) \right). \end{aligned}$$

It follows that $\text{PD}(\varepsilon(\mathbf{x}, \mathbf{y}, \mathbf{w})) = \pm \varepsilon'(\mathbf{x}, \mathbf{y}, \mathbf{w})$: 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 8.3 can be interpreted as the triply-periodic domain $\mathcal{D}(u) - \mathcal{D}(v)$. We claim that this triply-periodic domain corresponds to a relative cohomology class in $H^2(X, \partial X)$ whose image in $H^2(X)$ is the difference $\mathfrak{s}_z(u) - \mathfrak{s}_z(v)$, which in turn verifies that the map \mathfrak{s}_z is injective as claimed. But this now 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 the cylinders $[\xi_i^*] \times e_\xi$ (where $\xi = \alpha, \beta$, or γ , and $i = 1, \dots, g$). 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 [29]). \square

8.1.5. Higher polygons. The above results for triangles admit straightforward generalizations to arbitrarily large collections of g -tuples, which are called *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 α, β, γ , and δ 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 $\text{Sym}^g(\Sigma)$, and equivalence classes of homotopy classes $S_{\alpha, \beta, \gamma, \delta}$. That is, 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 8.5 admits a straightforward generalization, giving a map from $S_{\alpha, \beta, \gamma, \delta}$ to the space of Spin^c structures over $X_{\alpha, \beta, \gamma, \delta}$.

8.2. 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 (nearly-symmetric) almost-complex structures over $\text{Sym}^g(\Sigma)$ chosen to be compatible near the corners

with the paths J_s used to define the notion of pseudo-holomorphic disk. Moreover, we will use a class of perturbations of the constant complex structure for which the analogue of Lemma 3.2 still holds: if u is a J -holomorphic triangle, the domain associated to u is nonnegative.

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 8.6. 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:

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

under the identification coming from splicing.

Existence is ensured by the following:

LEMMA 8.7. 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 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 $\phi_{\alpha,\gamma} \in \Pi_{\mathbf{w}_0} \subset \pi_2(\mathbf{w}_0, \mathbf{w}_0)$ are periodic classes. To this end, let K denote the subgroup of $\Pi_{\mathbf{w}_0}$ consisting of elements $\phi_3 \in \Pi_{\mathbf{w}_0}$ which satisfy the property that

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

for some periodic classes ϕ_1 and ϕ_2 in $\Pi_{\mathbf{x}_0}$ and $\Pi_{\mathbf{y}_0}$ respectively. We claim that (1) given $\phi_3 \in K$, the corresponding ϕ_1 and ϕ_2 are uniquely determined,

and also that (2) the quotient Q of $\Pi_{\mathbf{w}_0}$ is a free Abelian group. According to Claim (2), we obtain a splitting

$$\Pi_{\mathbf{w}_0} \cong K \oplus Q.$$

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

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

According to Claim (1), this is a consistent definition. 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 $\psi \in \pi_2(\mathbf{x}_0, \mathbf{y}_0, \mathbf{w}_0)$. Both Claims (1) and (2) follow from the diagram below:

$$\begin{array}{ccccccc} & & & & H_2(Y_{\alpha,\gamma}) & & \\ & & & & \downarrow i & & \\ H_3(X, Y_{\alpha,\beta} \amalg Y_{\beta,\gamma}) & \rightarrow & H_2(Y_{\alpha,\beta} \amalg Y_{\beta,\gamma}) & \xrightarrow{j} & H_2(X) & \xrightarrow{q} & H_2(X, Y_{\alpha,\beta} \amalg Y_{\beta,\gamma}), \end{array}$$

where $X = X_{\alpha,\beta,\gamma}$. Interpreting the second homology as periodic domains, we see that the quotient group Q is the image of $q \circ i$. Now, by excision,

$$H_*(X, Y_{\alpha,\beta} \amalg Y_{\beta,\gamma}) \cong H_*(U_\beta \times (e_\beta, \partial e_\beta))$$

(where here U_β is a handlebody and e_β is an interval); in particular, j is injective, establishing Claim (1), and also the image of q is nontorsion, establishing Claim (2).

We have thus consistently oriented all triangles 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. \square

8.3. 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 4.2.2.

Definition 8.8. A pointed Heegaard triple-diagram is called *weakly admissible* if each nontrivial 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 4.13 and 4.14:

LEMMA 8.9. *Let $(\Sigma, \alpha, \beta, \gamma, z)$ be a 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 $\mathfrak{s}_z(\psi) = \mathfrak{s}$ 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 8.5, can be written as a sum of doubly-periodic domains. Given this, finiteness follows as in the proof of Lemma 4.13. \square

LEMMA 8.10. *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 $\psi' = \psi + \phi_1 + \phi_2 + \phi_3$; 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 4.14. \square

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

LEMMA 8.11. *Given a Heegaard triple diagram $(\Sigma, \alpha, \beta, \gamma, z)$, 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 α , β , and γ simultaneously. \square

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

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

by the formula:

$$(21) \quad f_{\alpha,\beta,\gamma}^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j]; \mathfrak{s}) \\ = \sum_{\mathbf{w} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\{\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{w}) \mid \mathfrak{s}_z(\psi) = \mathfrak{s}, \mu(\psi) = 0\}} \left(\# \mathcal{M}(\psi) \right) \cdot [\mathbf{w}, i + j - n_z(\psi)].$$

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 8.9), 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

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

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

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

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

THEOREM 8.12. *Let $(\Sigma, \alpha, \beta, \gamma, z)$ 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 (21) is finite, giving rise to a U -equivariant chain map which also induces maps on homology:*

$$F_{\alpha,\beta,\gamma}^\infty(\cdot; \mathfrak{s}_{\alpha,\beta,\gamma}): HF^\infty(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}) \otimes HF^\infty(Y_{\beta,\gamma}, \mathfrak{s}_{\beta,\gamma}) \longrightarrow HF^\infty(Y_{\alpha,\gamma}, \mathfrak{s}_{\alpha,\gamma}),$$

$$F_{\alpha,\beta,\gamma}^{\leq 0}(\cdot; \mathfrak{s}_{\alpha,\beta,\gamma}): HF^{\leq 0}(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}) \otimes HF^{\leq 0}(Y_{\beta,\gamma}, \mathfrak{s}_{\beta,\gamma}) \longrightarrow HF^{\leq 0}(Y_{\alpha,\gamma}, \mathfrak{s}_{\alpha,\gamma}).$$

The induced (U -equivariant) chain map

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

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

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

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

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

the latter of which is also U -equivariant.

Proof. The fact that $f_{\alpha,\beta,\gamma}^\infty$ is a chain map follows by counting ends of one-dimensional moduli spaces of holomorphic triangles (compare [23]). 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:

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

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

$$f_{\alpha,\beta,\gamma}^\infty(U([\mathbf{x}, i]) \otimes [\mathbf{y}, j]) = U \cdot f_{\alpha,\beta,\gamma}^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j])$$

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

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

PROPOSITION 8.13. *The maps on homology listed in Theorem 8.12 are independent of the choice of family J (and underlying complex structure j) over Σ used in their definition.*

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

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

defined by

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

Now, the ends of

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

count:

$$\begin{aligned} & f_{J_0}^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j]; \mathfrak{s}) - f_{J_1}^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j]; \mathfrak{s}) \\ & \quad + \partial \circ H^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j]) - H^\infty \circ \partial([\mathbf{x}, i] \otimes [\mathbf{y}, j]); \end{aligned}$$

i.e. $f_{J_0}^\infty$ and $f_{J_1}^\infty$ are chain homotopic.

The above argument shows that if we fix a map J into the space of almost-complex structures over $\text{Sym}^g(\Sigma)$, with specified behaviour near the vertices of the triangle, then the induced map is uniquely specified up to chain homotopy. Now, we investigate invariance as we change the one-parameter families of paths used at the corners. To this end, we claim that the following diagram,

$$\begin{array}{ccc} (CF^\infty(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}), \delta_{J_s(0)}^\infty) \otimes CF^\infty(Y_{\beta,\gamma}, \mathfrak{s}_{\alpha,\beta}) & \xrightarrow{f_{\alpha,\beta,\gamma}} & (CF^\infty(Y_{\alpha,\gamma}), \delta_{J_s(0)}^\infty) \\ \Phi_{\alpha,\beta}^\infty \otimes \text{Id} \downarrow & & \downarrow \Phi_{\alpha,\gamma}^\infty \\ (CF^\infty(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}), \delta_{J_s(1)}^\infty) \otimes CF^\infty(Y_{\beta,\gamma}, \mathfrak{s}_{\alpha,\beta}) & \xrightarrow{f_{\alpha,\beta,\gamma}} & (CF^\infty(Y_{\alpha,\gamma}), \delta_{J_s(1)}^\infty), \end{array}$$

commutes up to homotopy. Here, the maps Φ^∞ are the homomorphisms induced by variations of these paths (which we encountered in the proof of Theorem 6.1; note that we have now dropped the variation of path from the notation, and replaced it with an indication of the Heegaard diagram we are

working with). The homotopy now is obtained by counting triangles in a formally -1 -dimensional space, which are J_τ -holomorphic with respect to any J_τ for τ in a one-parameter family. We do not spell out the details here: a very similar argument is given in the next proposition. Similar remarks hold as we vary the boundary behaviour at the other vertices.

As a final remark, 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 j over Σ . \square

PROPOSITION 8.14. *The maps on homology listed in Theorem 8.12 are invariant under isotopies of the α , β , and γ preserving all the admissibility hypotheses.*

Proof. We begin with isotopies of the α . 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 α in Σ), and we consider moduli spaces with dynamic boundary conditions. Specifically, let $E_\alpha: \mathbb{R} \rightarrow \Delta$ be a parametrization of the edge e_α , with

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

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

$$\mathcal{M}_\tau = \left\{ u: \Delta \rightarrow \text{Sym}^g(\Sigma) \left| \begin{array}{l} u \circ E_\alpha(t) \in \Psi_{t+\tau}(\mathbb{T}_\alpha) \\ u(e_\beta) \subset \mathbb{T}_\beta, u(e_\gamma) \subset \mathbb{T}_\gamma \end{array} \right. \right\},$$

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

$$\begin{aligned} & H^\infty([\mathbf{x}, i] \otimes [\mathbf{y}, j]; \mathfrak{s}) \\ &= \sum_{\mathbf{w}} \sum_{\{\psi \in \pi_2^{\Psi_t}(\mathbf{x}, \mathbf{y}, \mathbf{w}) \mid \mu(\psi) = -1, \mathfrak{s}_z(\psi) = \mathfrak{s}\}} \left(\# \bigcup_{\tau \in \mathbb{R}} \mathcal{M}_\tau(\psi) \right) [\mathbf{w}, i + j - n_z(\psi)]. \end{aligned}$$

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 7.3), $\psi_{\alpha,\beta} \in \pi_2(\mathbf{x}', \mathbf{y}, \mathbf{w})$, and $\mu(\phi_{\alpha,\beta}) = \mu(\psi_{\alpha,\beta}) = 0$ (with analogous conditions on the second union). There are also ends of the form

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

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 7.3), $\psi_{\alpha,\beta} \in \pi_2^{\Psi_t}(\mathbf{x}', \mathbf{y}, \mathbf{w})$, 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

$$\Gamma_{\alpha,\alpha',\beta} : CF^\infty(Y_{\alpha,\beta}, \mathfrak{s}_{\alpha,\beta}) \longrightarrow CF(Y_{\alpha',\beta}, \mathfrak{s}_{\alpha',\beta})$$

and

$$\Gamma_{\alpha,\alpha',\gamma} : CF(Y_{\alpha,\gamma}, \mathfrak{s}_{\alpha,\gamma}) \longrightarrow CF(Y_{\alpha',\gamma}, \mathfrak{s}_{\alpha',\gamma})$$

are the chain maps induced by the isotopy Ψ_t , as constructed in Subsection 7.3 (note that here we have suppressed the isotopy Ψ_t from the notation).

Isotopies of the γ work the same way; we now set up isotopies of the β . Consider moduli spaces indexed by a real $\tau \in [0, \infty)$:

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

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 α , β , and γ . If $\mu(\psi) = -1$, then $\cup_{\tau \in [0, \infty)} \mathcal{M}_\tau(\psi)$ is generically a compact zero-dimensional manifold, so we can define

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

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

$$f_{\alpha,\beta',\gamma} \circ (\Gamma_{\alpha,\beta,\beta'} \otimes \Gamma_{\beta,\beta',\gamma}). \quad \square$$

8.4. Associativity of holomorphic triangles. The map induced by triangles satisfies an associativity property, on the level of homology. As is familiar in Lagrangian Floer homology, the chain homotopies required for the associativity are provided by a count of holomorphic squares; see [37], [12], [4].

Loosely speaking, this count is done as follows. Let \square denote the “rectangle”: the unit disk with four marked boundary points which we denote a_0 , a_1 , a_2 , and a_3 (labeled in clockwise order). Observe that the space of conformal structures on the rectangle, $\mathcal{M}(\square)$, is identified with \mathbb{R} under the map which is the logarithmic difference between the length and the width. We would like to count pseudo-holomorphic maps of the rectangle in $\text{Sym}^g(\Sigma)$ (without fixing the conformal structure of the domain). We wish to consider moduli spaces of pseudo-holomorphic Whitney rectangles with formal dimension one. Some ends of these moduli spaces are modeled on flowlines breaking off at the corners, but there is another type of end not encountered before in the counts of triangles, arising from the noncompactness of $\mathcal{M}(\square) \cong \mathbb{R}$. As this parameter goes to $\pm\infty$, the corresponding rectangle breaks up conformally into a pair of triangles meeting at a vertex (in two different ways, depending on which end we are considering), as illustrated in Figure 6. This is how a count of holomorphic squares induces a chain homotopy between two different compositions of holomorphic triangle counts..

The ingredients required in this construction are: the role of Spin^c structures in the construction, admissibility hypotheses required to make the holomorphic rectangle counts and finite compatibility with holomorphic triangle counts. With these components in place, the proof of associativity proceeds in the same way as it does in the usual Lagrangian Floer homology.




Figure 6: Ends of the moduli spaces holomorphic maps of rectangles. The two degenerations on the right correspond to degenerations of the conformal type of the underlying rectangle. (There are three additional degenerations like the one on the left, corresponding to the other three vertices.)

8.4.1. *Spin^c structures on rectangles.* Fix a pointed Heegaard quadruple $(\Sigma, \alpha, \beta, \gamma, \delta, z)$, and let $X_{\alpha, \beta, \gamma, \delta}$ be the corresponding cobordism. We have, of course, restriction maps:

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

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 $\text{Spin}^c(X_{\alpha, \beta, \gamma, \delta})$ are the fibers of this restriction map. Similarly, we have a restriction map

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

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.

8.4.2. *Admissibility for rectangles.* 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 $(\Sigma, \alpha, \beta, \gamma, \delta, z)$ 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:

$$(22) \quad \mathcal{P} = \sum_{\{\xi, \eta\} \subset \{\alpha, \beta, \gamma, \delta\}} \mathcal{D}_{\xi, \eta}$$

with the property that

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

(i.e. where $\mathfrak{s}_{\xi,\eta}$ 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 strongly \mathfrak{S} -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:

$$(23) \quad \delta H^1(Y_{\beta,\delta})|Y_{\alpha,\gamma} = 0 \text{ and } \delta H^1(Y_{\alpha,\gamma})|Y_{\beta,\delta} = 0.$$

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 (22) then

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

is a function of the periodic domain \mathcal{P} 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.

8.4.3. Compatibility with counts of squares. Having established the necessary admissibility requirements for defining counts of pseudo-holomorphic rectangles, we need two more ingredients: transversality and orientation.

We must set up the transverse perturbations for J -holomorphic rectangles in a manner which is compatible with the J_s paths used over strips and the “nearly-symmetric families” used on the triangles. To do this, we must fix a map

$$J: \mathcal{M}(\square) \times \square \longrightarrow \mathcal{U},$$

with certain properties. First, we arrange for \mathcal{U} to consist of (j, η, V) -nearly symmetric almost-complex structures, with V chosen to contain $\bigcup_{z_i} \{z_i\} \times \text{Sym}^{g-1}(\Sigma)$, where the $\{z_i\}$ are a finite collection of points, one from each domain in $\Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g - \gamma_1 - \dots - \gamma_g - \delta_1 - \dots - \delta_g$. Moreover, \overline{V} is chosen to be disjoint from all four tori \mathbb{T}_α , \mathbb{T}_β , \mathbb{T}_γ , and \mathbb{T}_δ . To arrange for compatibility with strips, we fix path $J_s^{(i)}$ for $i = 0, \dots, 3$, and assume that for each conformal structure in $\mathcal{M}(\square)$, the map J agrees with the $J_s^{(i)}$ path near the i^{th} vertex (under a conformal identification between a neighborhood of each corner and strip). To arrange for compatibility with triangles, we restrict the

behaviour of J as the parameter in $\mathcal{M}(\square)$ goes to infinity. As the parameter in $\mathcal{M}(\square)$ approaches one extreme, the complex structure over \square is identified with $[0, 1] \times [-T, T]$ for some large T ; when the parameter approaches the other extreme, the complex structure is identified with $[-T, T] \times [0, 1]$. For the first degeneration into triangles, we assume that $J|_{[0, 1] \times [-T, 0]}$ agrees, after translation by T on the second factor, with the restriction of some fixed admissible family on $[0, 1] \times [0, +\infty) \cong \Delta$, and similarly, that $J|_{[0, 1] \times [0, T]}$ agrees, after translation by $-T$ on the second factor, with the restriction of an admissible family on $[0, 1] \times (-\infty, 0] \cong \Delta$. We also assume analogous properties for the other degeneration.

Adapting the transversality proof, one can see that for generic $J: \mathcal{M}(\square) \times \square \rightarrow \mathcal{U}$ satisfying the compatibilities near infinity as above, the corresponding moduli spaces with formal dimension ≤ 1 are smooth.

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 $\text{Spin}^c(X_{\alpha, \beta, \gamma, \delta})$, which we denote \mathfrak{S} . A coherent system of orientations for \mathfrak{S} , then, is a collection of nonvanishing sections indexed by subsets $\{\xi_1, \dots, \xi_\ell\} \subset \{\alpha, \beta, \gamma, \delta\}$ with $\ell = 2, 3, 4$, $\mathfrak{o}_{\xi_1, \dots, \xi_\ell}(\phi_{\xi_1, \dots, \xi_\ell})$, for the determinant line bundle defined over the homotopy class of polygons $\phi_{\xi_1, \dots, \xi_\ell}$ (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 $X_{\xi_1, \dots, \xi_\ell}$ if $\ell = 3, 4$. These are required to be compatible with the gluings in the sense that

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

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

The following construction of coherent orientation systems will suffice for our purposes:

PROPOSITION 8.15. *Let $X_{\alpha, \beta, \gamma, \delta}$ belong to a Heegaard quadruple for which*

$$(24) \quad \left(\delta H^1(Y_{\beta, \delta}) \subseteq H^2(X_{\alpha, \beta, \gamma, \delta}, \partial X_{\alpha, \beta, \gamma, \delta}) \right) = 0.$$

Fix a $\delta H^1(Y_{\alpha, \gamma})$ -orbit \mathfrak{S} in the set of Spin^c structures over $X_{\alpha, \beta, \gamma, \delta}$. Fix also orientation systems $\mathfrak{o}_{\alpha, \beta, \gamma}$ and $\mathfrak{o}_{\alpha, \gamma, \delta}$ for the two corresponding Heegaard subtriples, which induce the same orientation over $Y_{\alpha, \gamma}$. Then, the given orientation systems can be completed compatibly to give an orientation system over the Heegaard quadruple $X_{\alpha, \beta, \gamma, \delta}$.

Proof. The given data already allow us to orient all squares representing \mathfrak{S} . Specifically, any such square can be decomposed as a sum of Whitney triangles $\varphi_{\alpha, \beta, \gamma, \delta} = \psi_{\alpha, \beta, \gamma} + \psi_{\alpha, \gamma, \delta}$, for which the ψ denote homotopy classes of Whitney

triangles for the subtriples indicated in the notation, and which represent the Spin^c -equivalence class induced from \mathfrak{S} . We then define the orientation for the given square by

$$\mathfrak{o}(\varphi_{\alpha,\beta,\gamma,\delta}) = \mathfrak{o}(\psi_{\alpha,\beta,\gamma}) \wedge \mathfrak{o}(\psi_{\alpha,\gamma,\delta})$$

(note that we have suppressed the subscripts from \mathfrak{o} indicating the diagram used for the polygon, as our notation for the homotopy classes already carries this information). Since the induced orientations over $Y_{\alpha,\gamma}$ are compatible, this definition is independent of the decomposition. In particular, it follows that for any four Whitney disks ϕ_1, \dots, ϕ_4 , for the Heegaard pairs on the boundary ($Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, $Y_{\gamma,\delta}$, and $Y_{\alpha,\delta}$ respectively), we already have that

$$(25) \quad \mathfrak{o}(\phi_{\alpha,\beta,\gamma,\delta} + \phi_1 + \dots + \phi_4) = \mathfrak{o}(\phi_{\alpha,\beta,\gamma,\delta}) \wedge \mathfrak{o}(\phi_1) \wedge \dots \wedge \mathfrak{o}(\phi_4).$$

Next, fix a pair of triangles $\psi_{\alpha,\beta,\delta} \in \pi_2(\mathbf{x}_0, \mathbf{u}_0, \mathbf{w}_0)$, $\psi_{\beta,\gamma,\delta} \in \pi_2(\mathbf{y}_0, \mathbf{v}_0, \mathbf{u}_0)$ for the corresponding Heegaard triples, and fix an arbitrary orientation $\mathfrak{o}(\psi_{\alpha,\beta,\delta})$, and then orient $\psi_{\beta,\gamma,\delta}$ so that

$$\mathfrak{o}(\psi_{\alpha,\beta,\delta}) \wedge \mathfrak{o}(\psi_{\beta,\gamma,\delta}) = \mathfrak{o}(\psi_{\alpha,\beta,\delta} + \psi_{\beta,\gamma,\delta}).$$

Given this, we can now orient all $\phi_{\beta,\delta} \in \pi_2(\mathbf{u}_0, \mathbf{u}_0)$: given any such $\phi_{\beta,\delta}$, in view of Equation (24) we can write

$$\psi_{\alpha,\beta,\delta} + \phi_{\beta,\delta} + \psi_{\alpha,\gamma,\delta} = \psi_{\alpha,\beta,\delta} + \psi_{\alpha,\gamma,\delta} + \phi_1 + \phi_2 + \phi_3 + \phi_4,$$

where $\phi_1 \in \pi_2(\mathbf{x}_0, \mathbf{x}_0)$, $\phi_2 \in \pi_2(\mathbf{y}_0, \mathbf{y}_0)$, $\phi_3 \in \pi_2(\mathbf{v}_0, \mathbf{v}_0)$, and $\phi_4 \in \pi_2(\mathbf{w}_0, \mathbf{w}_0)$. We now define $\mathfrak{o}(\phi_{\beta,\gamma})$ to be compatible with this decomposition, and observe that, according to Equation (25), the answer is independent of the choices of ϕ_1, \dots, ϕ_4 . We extend these orientations on $\pi_2(\mathbf{u}_0, \mathbf{u}_0)$ to an orientation system for $Y_{\beta,\delta}$ in the usual manner. (Note that the isomorphism class of this orientation system is uniquely determined.)

This orientation system of $Y_{\beta,\delta}$ and the fixed orientation $\mathfrak{o}(\psi_{\alpha,\beta,\delta})$ on the initial triangle $\psi_{\alpha,\beta,\delta}$ suffices to orient all the remaining triangles for the triples $X_{\alpha,\beta,\delta}$ in the Spin^c -equivalence class corresponding to the restriction of \mathfrak{S} ; any other such triangle $\psi'_{\alpha,\beta,\delta}$ can be written in the form

$$\psi'_{\alpha,\beta,\delta} = \psi_{\alpha,\beta,\delta} + \phi_{\alpha,\beta} + \phi_{\beta,\delta} + \phi_{\alpha,\delta},$$

and can be oriented correspondingly. The orientation obtained in this manner is well-defined, owing to Equation (25). The same remarks hold for the remaining subtriple $X_{\beta,\gamma,\delta}$, and since the induced orientation systems on the $Y_{\beta,\delta}$ agree, it is straightforward now to verify that the induced orientation system on the $X_{\alpha,\beta,\delta}$ and $X_{\beta,\gamma,\delta}$ are consistent with the orientation system on the squares in $X_{\alpha,\beta,\gamma,\delta}$. \square

8.4.4. *The associativity theorem.*

THEOREM 8.16. *Let $(\Sigma, \alpha, \beta, \gamma, \delta, z)$ 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 $\text{Spin}^c(X_{\alpha, \beta, \gamma, \delta})$. Then,*

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

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,

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

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

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

by

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

Note that the map above 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

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

and

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

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

9. Handleslide invariance

Our aim is to prove handleslide invariance of the Floer homology groups. In Subsection 9.1 we establish such a result for the three-manifold $\#^g(S^2 \times S^1)$ (equipped with a particular Spin^c structure) and a specific handleslide, by explicitly calculating the Floer homologies for both Heegaard diagrams. In Subsection 9.2, we use the holomorphic triangle construction to transfer the result for this specific three-manifold to a general three-manifold.

9.1. *The Floer homologies of $\#^g(S^2 \times S^1)$.* We give now a model calculation of a handleslide.

Let Σ be an oriented surface of genus g equipped with a basepoint z . Fix a g -tuple β of attaching circles for a handlebody bounding a genus g surface Σ (disjoint from z), and let γ be another g -tuple obtained by handlesliding β_1 over β_2 (in the complement of z). Fix also another g -tuple δ which is isotopic to the β . More precisely, let β'_1 be a curve obtained by handlesliding β_1 over β_2 . Move β'_1 by a small isotopy to γ_1 , so that it meets β_1 in a pair of transverse intersection points with opposite signs. Let γ_i for $i > 1$ be obtained by small isotopies of β_i . Furthermore let δ_i for $i = 1, \dots, g$ be obtained by small isotopies of β_i . We also require that for each $i = 1, \dots, g$ the pairwise intersections $\beta_i \cap \gamma_i$, $\beta_i \cap \delta_i$, $\gamma_i \cap \delta_i$ consist of two transverse points with opposite signs. We denote these points y_i^\pm , w_i^\pm , v_i^\pm respectively. All the above isotopies are taken to be small enough to be disjoint from the initial basepoint z . Moreover, we arrange that z is also outside of the pair of pants domain which bounds β_1, β_2 and β'_1 . (See the corresponding Figure 8 for $g = 2$, where the small circles are identified by their vertical pairs to give the genus two surface.)

It is easy to see that the three Heegaard diagrams $(\Sigma, \beta, \gamma, z)$, $(\Sigma, \gamma, \delta, z)$ all $(\Sigma, \beta, \gamma, z)$ represent the three-manifold $\#^g(S^2 \times S^1)$. Let \mathfrak{s}_0 denote the Spin^c structure on $\#^g(S^2 \times S^1)$ with $c_1(\mathfrak{s}_0) = 0$.

We will calculate the Floer homologies of $\#^g(S^2 \times S^1)$ in the Spin^c structure \mathfrak{s}_0 with respect to all three Heegaard diagrams. We start with the easiest case, $(\Sigma, \beta, \delta, z)$. Note that we can choose the curves β_i to be exact Hamiltonian isotopic to the δ_i , so that the following can be thought of as a natural analogue of a basic result of Floer.




Figure 7

LEMMA 9.1. *The pointed Heegaard diagram $(\Sigma, \beta, \delta, z)$ is admissible for the Spin^c structure \mathfrak{s}_0 . Moreover, there is a choice of orientation conventions $\mathfrak{o}_{\beta, \gamma}$ with the property that*

$$\begin{aligned}\widehat{HF}(\Sigma, \beta, \delta, \mathfrak{o}_{\beta, \delta}) &\cong H_*(T^g; \mathbb{Z}) \\ HF^-(\Sigma, \beta, \delta, \mathfrak{o}_{\beta, \delta}) &\cong \mathbb{Z}[U] \otimes H_*(T^g; \mathbb{Z}),\end{aligned}$$

where T^g denotes the g -dimensional torus.

Proof. There are altogether 2^g intersection points between $\mathbb{T}_\beta \cap \mathbb{T}_\gamma$. Fix some $i = 1, \dots, g$, and let $\mathbf{x} = \{x_1, \dots, x_g\}$ and $\mathbf{y} = \{y_1, \dots, y_g\}$ be intersection points with $x_i = w_i^+$, $y_i = w_i^-$ for all $j \neq i$ we have $x_j = y_j$. Note that β_i, δ_i bound two pairs of disks, which we denote by D_1 and D_2 , as in Figure 7. Let $\phi_1, \phi_2 \in \pi_2(\mathbf{x}, \mathbf{y})$ be homotopy classes given by $\mathcal{D}(\phi_1) = D_1$ and $\mathcal{D}(\phi_2) = D_2$ respectively. It is easy to see that $\mu(\phi_1) = \mu(\phi_2) = 1$. By juxtaposing the disks we get a periodic class $\phi_1 - \phi_2$ in $\pi_2(\mathbf{x}, \mathbf{x})$, with $\mu(\phi_1 - \phi_2) = 0$. It is easy to see that the subgroups of $\pi_2(\mathbf{x}, \mathbf{x})$ and $\pi_2(\mathbf{y}, \mathbf{y})$ consisting of homotopy classes with $n_z(\phi) = 0$ are generated by g classes, all of which can be realized in this way; so it follows that all 2^g intersection points represent the trivial Spin^c structure. Moreover, since each nontrivial combination of these periodic classes has both positive and negative coefficients, it follows that $(\Sigma, \beta, \delta, z)$ is admissible for \mathfrak{s}_0 (since $c_1(\mathfrak{s}_0) = 0$, the strong and weak admissibility notions coincide). It also follows that if $d(\mathbf{w})$ denotes the number of positive coordinates for $\mathbf{w} = \{w_1^\pm, \dots, w_g^\pm\}$, then the relative grading is given by $\text{gr}(\mathbf{w}, \mathbf{w}') = d(\mathbf{w}) - d(\mathbf{w}')$. We still have to show that all boundary maps are zero.

For the purpose of these calculations, we will use the path $J_t \equiv \text{Sym}^g(\mathbf{j})$, where \mathbf{j} is some fixed complex structure over Σ . In these calculations, transversality for the flow-lines considered is either immediate or it is achieved by moving the curves, as in Proposition 3.9.

First let \mathbf{x} and \mathbf{y} be as above. Then for all $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, with $n_z(\phi) = 0$ and $\phi \neq \phi_1, \phi_2$, the moduli space $\mathcal{M}(\phi)$ is empty, since $\mathcal{D}(\phi)$ has some negative

coefficients. Also $\widehat{\mathcal{M}}(\phi_1)$ contains a unique solution that maps the trivial g -fold cover of the disk to Σ , so that one of the sheets is mapped to D_1 , and other sheets map to x_j for $j \neq i$ respectively. Similarly $\widehat{\mathcal{M}}(\phi_2)$ also contains a unique smooth solution. Since the two domains differ by a periodic domain, we are free to choose an orientation system for which the signs corresponding to these solutions are different. For this system, the component of \mathbf{y} in $\widehat{\partial}\mathbf{x}$ is zero.

Now suppose that \mathbf{x} and \mathbf{y} differ in at least two coordinates, and $\text{gr}(\mathbf{x}, \mathbf{y}) = 1$. Then it is easy to see that for all $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, with $n_z(\phi) = 0$ the moduli space $\mathcal{M}(\phi)$ is empty, since $\mathcal{D}(\phi)$ has some negative coefficients. This shows that $\widehat{\partial} \equiv 0$, and consequently $\widehat{HF}(\boldsymbol{\beta}, \boldsymbol{\delta}) \cong H_*(T^g; \mathbb{Z})$. In fact, since the two disks D_1 and D_2 separately have a unique holomorphic representative, it follows that the $H_1(\#^g(S^2 \times S^1))$ -module structure is given by the identification $H_*(T^g) \cong \Lambda^* H_1(\#^g(S^2 \times S^1))$.

Let $\mathbf{y}^+ = w_1^+ \times \cdots \times w_g^+$ be the intersection point representing a top-dimensional homology class in $\widehat{CF}(\boldsymbol{\beta}, \boldsymbol{\delta}, z)$. We have shown that for all $\mathbf{y} \neq \mathbf{y}^+ \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, $\text{gr}(\mathbf{y}^+, \mathbf{y}) \geq 1$. It follows that for all $\phi \in \pi_2(\mathbf{y}^+, \mathbf{y})$ with $\mu(\phi) = 1$ we have $n_z(\phi) \leq 0$. By Lemma 3.2, when $n_z(\phi) < 0$, the moduli space $\mathcal{M}(\phi)$ is empty. The case $n_z(\phi) = 0$ corresponds to the boundary map in $\widehat{CF}(\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathfrak{s}_0)$ which was just shown to be trivial. It follows that $\partial^\infty([\mathbf{y}^+, i]) = 0$. The algebra action on $[\mathbf{y}^+, -1]$ gives then the map

$$\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H^1(\#^g(S^2 \times S^1); \mathbb{Z}) \longrightarrow CF^-(\boldsymbol{\beta}, \boldsymbol{\delta}, \mathfrak{s}_0),$$

which is easily seen to be an isomorphism by properties of the short exact sequence


$$0 \longrightarrow CF^-(\boldsymbol{\beta}, \boldsymbol{\delta}, \mathfrak{s}_0) \xrightarrow{U} CF^-(\boldsymbol{\beta}, \boldsymbol{\delta}, \mathfrak{s}_0) \longrightarrow \widehat{CF}(\boldsymbol{\beta}, \boldsymbol{\delta}, \mathfrak{s}_0) \longrightarrow 0.$$

□

Remark 9.2. In fact, the above proof shows that for any orientation system \mathfrak{o} , the chain complex for $\widehat{CF}(\boldsymbol{\beta}, \boldsymbol{\delta}, \mathfrak{s}_0, \mathfrak{o})$ is a g -fold tensor product of two-step complexes $\mathbb{Z} \longrightarrow \mathbb{Z}$, where the boundary map is multiplication by 0 or ± 2 . Of the 2^g possible (isomorphism classes of) complexes, the orientation of Lemma 9.1 is characterized as the only one with a nonzero, g -dimensional cycle.

In order to calculate the complexes associated to $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$, it is useful to have the following:

LEMMA 9.3. *Let A be an annulus, and fix arcs ξ_1 and ξ_2 on the outer and inner boundaries respectively. Then there is a holomorphic involution of A switching ξ_1 and ξ_2 if and only if the angles swept out by ξ_1 and ξ_2 agree. Moreover, if this condition is satisfied, the involution is unique.*

Figure 8: β_i , γ_i , δ_i in the genus 2 surface.

Proof. We can think of A as the set $\{z \in \mathbb{C} \mid r < |z| < 1/r\}$. By the Schwartz reflection principle, we can extend any involution of A to the Riemann sphere, so that it switches 0 and ∞ . Thus, this involution has the form $z \mapsto c/z$ for some $c \in \mathbb{C}^*$. Since the involution takes the set of complex numbers with modulus r to those with modulus $1/r$, it follows that $|c| = 1$, and the lemma follows. \square

LEMMA 9.4. *The pointed Heegaard diagram $(\Sigma, \beta, \gamma, z)$ is admissible for the Spin^c structure \mathfrak{s}_0 . Moreover, there is a choice of orientation conventions $\mathfrak{o}_{\beta, \gamma}$ with the property that*

$$\widehat{HF}(\Sigma, \beta, \gamma, \mathfrak{o}_{\beta, \gamma}) \cong H_*(T^g; \mathbb{Z}),$$

$$HF^-(\Sigma, \beta, \gamma, \mathfrak{o}_{\beta, \gamma}) \cong \mathbb{Z}[U] \otimes H_*(T^g; \mathbb{Z}),$$

where T^g denotes the g -dimensional torus.

Proof. We work out the $g = 2$ case, for notational convenience (the general case follows easily).

Let D_i for $1 \leq i \leq 5$ denote the connected components of $\Sigma_2 - \beta_1 - \beta_2 - \gamma_1 - \gamma_2$; see Figure 9. Let $\mathbf{x} = \{y_1^+, y_2^+\}$, and $\mathbf{y} = \{y_1^-, y_2^-\}$. Then there are exactly three classes, $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, with $n_z(\phi) = 0$ and $\mathcal{D}(\phi) \geq 0$, and these are given by $\mathcal{D}(\phi_1) = D_1$, $\mathcal{D}(\phi_2) = D_2 + D_3$, and $\mathcal{D}(\phi_3) = D_2 + D_4$. As before, $\mu(\phi_1) = 1$. We claim that $\mu(\phi_2) = \mu(\phi_3) = 1$. Since these cases are symmetric we deal with ϕ_2 .

According to Lemma 3.6, a holomorphic disk $u : D \rightarrow \text{Sym}^2(\Sigma)$ representing ϕ_2 gives rise to a branched double cover $\pi : F \rightarrow D$ and a map $\hat{u} : F \rightarrow \Sigma$. In our case \hat{u} has degree 1 in D_2 and D_3 and has degree 0 on the other regions. Here, F is an annulus, and the image of ∂F lies in the union of β_1 , β_2 ,




Figure 9




Figure 10

γ_1, γ_2 . In fact the part of the image that lies in γ_2 is an arc starting at y_2^+ , see Figure 10. More generally, for each $t \in [0, 1)$, we can consider the subset $B_t \subset \Sigma$, which is the interior of the region in $D_2 \cup D_3$ obtained by removing a length t subarc of γ_2 starting at y_2^+ , with arc-length normalized so that $t = 1$ corresponds to the endpoint y_2^- .

The region B_t is topologically an (open) annulus, and hence can be identified conformally with a standard (open) annulus $A_t^\circ = \{z \in \mathbb{C} \mid r_t < |z| < 1\}$, where, r_t is a nonzero real number depending on t . The identification is given by a map $\Phi_t: A_t \rightarrow B_t$ which extends to a continuous map from the closure A_t of A_t° to the closure of B_t , which maps the boundary of A_t into the union of $\beta_1, \beta_2, \gamma_1$, and γ_2 , so as to map onto the length t sub-arc of γ_2 . Let $\xi_1(t)$ and $\xi_2(t)$ denote the subsets of ∂A_t which map to β_1 and β_2 respectively, and let $f_1(t)$ and $f_2(t)$ denote the angles swept out by $\xi_1(t)$ and $\xi_2(t)$. When the map $\Phi_t: A_t \rightarrow B_t \subset \Sigma$ is induced from a holomorphic disk in the second symmetric product, the corresponding branched double-cover (of the disk, by A_t) induces an involution on A_t which switches $\xi_1(t)$ to $\xi_2(t)$. These assertions follow from simple modifications of the Riemann mapping theorem; see for example [1]. Indeed, it follows from these classical considerations that if the curves are smoothly embedded, then the objects r_t and Φ_t also depend smoothly on the parameter t .

In light of Lemmas 3.6 and 9.3, points in the moduli space $\widehat{\mathcal{M}}(\phi_2)$ are in one-to-one correspondence with $t \in (0, 1)$ with $f_1(t) = f_2(t)$. For certain (generic) choices of the curve β_1 , we can arrange that the graph of f_1 and f_2 have nonempty, transversal intersection; i.e. that there is a finite, nonempty collection of holomorphic disks. By a slight perturbation of the curves, it then follows that the formal dimension of $\widehat{\mathcal{M}}(\phi_2)$ is 0, so $\mu(\phi_2) = 1$.

Now let $\mathbf{x}' = \{y_1^+, y_2^-\}$, and let $\phi_4, \phi_5 \in \pi_2(\mathbf{x}, \mathbf{x}')$ be given by $\mathcal{D}(\phi_4) = D_3$, $\mathcal{D}(\phi_5) = D_4$. Clearly $\phi_4 - \phi_5$ and $\phi_1 - \phi_2$ generate the periodic classes of $\pi_2(\mathbf{x}, \mathbf{x}')$, and both of them has zero Maslov index. Clearly any nontrivial linear combination has positive and negative coefficients as well, so our diagram $(\Sigma, \beta, \gamma, \mathfrak{s}_0)$ is \mathfrak{s}_0 -admissible.

In order to compute the boundary maps, note that $\#\widehat{\mathcal{M}}(\phi_1) = \pm 1$, where the unique solution is again a product of a holomorphic disk and a constant map. We also claim that

$$\#\widehat{\mathcal{M}}(\phi_2) + \#\widehat{\mathcal{M}}(\phi_3) = \pm 1.$$

This computation follows easily from the previous discussion and some complex analysis. Let A denote the annulus given by the domain D_2 in Σ , and let ν_1 and ν_2 denote the conformal angles of the parts of ∂A that correspond to β_1 and β_2 respectively. By general position we can assume that $\nu_1 \neq \nu_2$.

Let f_1, f_2 be defined as above. Clearly $f_2(0) = 2\pi$, so for small enough t we have $f_1(t) < f_2(t)$. We claim that $\lim_{t \rightarrow 1} A_t = A$, $\lim_{t \rightarrow 1} f_1(t) = \nu_1$, $\lim_{t \rightarrow 1} f_2(t) = \nu_2$. This follows readily from Gromov's compactness (although in this case, it could also be proved using classical conformal analysis; see for example [1]). It follows that if $\nu_1 < \nu_2$, then $\#\widehat{\mathcal{M}}(\phi_2) = 0$, and if $\nu_1 > \nu_2$, then $\#\widehat{\mathcal{M}}(\phi_2) = \pm 1$.

Thus, we have calculated $\#\mathcal{M}(\phi_2)$. The computation for ϕ_3 is similar. Again we have an annulus with a cut of length $t \in [0, 1)$. Let $g(t)$ denote the conformal angle corresponding to β_1, β_2 respectively. In this case $g_2(0) = 0$, so for small enough t we have $g_1(t) > g_2(t)$. As above, as $t \rightarrow 1$, the conformal angle $g_i(t)$ converges to ν_i . It follows that if $\nu_1 < \nu_2$, then $\#\widehat{\mathcal{M}}(\phi_3) = \pm 1$, and if $\nu_1 > \nu_2$ then $\#\widehat{\mathcal{M}}(\phi_3) = 0$. This proves that $\#\widehat{\mathcal{M}}(\phi_2) + \#\widehat{\mathcal{M}}(\phi_3) = \pm 1$.

Again, we are free to choose an orientation system so that the sign of the flow obtained here cancels the sign in $\#\mathcal{M}(\phi_1)$ (since the domains differ by a periodic domain), so that the \mathbf{y} component of $\partial \mathbf{x}$ vanishes.

The same argument shows also that the \mathbf{y}' -component of $\widehat{\partial} \mathbf{x}' = 0$, where $\mathbf{y}' = \{y_1^-, y_2^-\}$. The remaining components of $\widehat{\partial}$ can be shown to vanish, by the arguments from Lemma 9.1. This shows that $\widehat{HF}(\beta, \gamma, z) \equiv H_*(T^2, \mathbb{Z})$. The corresponding statement about $HF^-(\beta, \gamma, z)$ also follows as before.

The case where $g > 2$ follows from the $g = 2$ case (where we have the same holomorphic maps multiplied with $g - 2$ constant maps). \square

9.2. *Naturality.* Continuing notation from the previous section, we let

$$\begin{aligned}\widehat{\Theta}_{\beta,\gamma} &\in \widehat{HF}(\beta, \gamma, \mathfrak{s}_0, \mathfrak{o}_{\beta,\gamma}), \\ \widehat{\Theta}_{\gamma,\delta} &\in \widehat{HF}(\gamma, \delta, \mathfrak{s}_0, \mathfrak{o}_{\gamma,\delta}), \\ \widehat{\Theta}_{\beta,\delta} &\in \widehat{HF}(\beta, \delta, \mathfrak{s}_0, \mathfrak{o}_{\beta,\delta})\end{aligned}$$

and $\Theta_{\beta,\gamma}^{\leq}$, $\Theta_{\gamma,\delta}^{\leq}$, and $\Theta_{\beta,\delta}^{\leq}$ denote the top-dimensional generators of $HF^{\leq 0}$ coming from Lemmas 9.1 and 9.4; i.e. $\widehat{\Theta}_{\beta,\gamma}$ is represented by the intersection point $\mathbf{y}^+ = \{y_1^+, \dots, y_g^+\}$, $\widehat{\Theta}_{\gamma,\delta}$ is represented by $\mathbf{v}^+ = \{v_1^+, \dots, v_g^+\}$, and $\widehat{\Theta}_{\beta,\delta}$ is represented by $\mathbf{w}^+ = \{w_1^+, \dots, w_g^+\}$.

Let Y be a three-manifold equipped with a Spin^c structure \mathfrak{s} , and $(\Sigma, \alpha, \beta, z)$ be a strongly \mathfrak{s} -admissible Heegaard diagram, and let γ be obtained from β by a handleslide as chosen in the isotopy class described in Subsection 9.1. Note that the Heegaard triple $(\Sigma, \alpha, \beta, \gamma, z)$ represents the cobordism $Y \times [0, 1]$, with a bouquet of g circles removed (cf. Example 8.1). It follows that for each Spin^c structure \mathfrak{s} over Y , there is a uniquely induced Spin^c structure on the Heegaard triple $X_{\alpha,\beta,\gamma}$ whose restriction to $Y = Y_{\alpha,\beta}$ is Y and $\#^g(S^2 \times S^1) = Y_{\beta,\gamma}$ has trivial first Chern class.

It is also easy to see that if we specify some arbitrary isomorphism classes of orientation system $\mathfrak{o}_{\alpha,\beta}$ for $(\Sigma, \alpha, \beta, z)$, the orientation system $\mathfrak{o}_{\beta,\gamma}$ given by Lemma 9.4, then these uniquely induce an orientation system $\mathfrak{o}_{\alpha,\gamma}$ on $(\Sigma, \alpha, \gamma, z)$.

THEOREM 9.5. *Let Y be a closed three-manifold equipped with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, and let $(\Sigma, \alpha, \beta, z)$ be a strongly \mathfrak{s} -admissible Heegaard diagram. Fix an arbitrary isomorphism class of orientation system $\mathfrak{o}_{\alpha,\beta}$ for $(\Sigma, \alpha, \beta, z)$, and let γ be the g -tuple of circles obtained from β by a handleslide as above. Then, for the induced orientation system $\mathfrak{o}_{\alpha,\gamma}$, the chain map $\xi \mapsto f^\infty(\xi \otimes \Theta_{\beta,\gamma}^{\leq})$ given by counting holomorphic triangles over the Heegaard triple $(\Sigma, \alpha, \beta, \gamma, z)$ induces an isomorphism of $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(Y; \mathbb{Z}) / \text{Tors-modules}$*

$$HF^\infty(\Sigma, \alpha, \beta, \mathfrak{s}) \longrightarrow HF^\infty(\Sigma, \alpha, \gamma, \mathfrak{s});$$

and indeed the triangle count induces isomorphisms on all the other homology groups HF^- , HF^+ , and \widehat{HF} .

LEMMA 9.6. *Let $(\Sigma, \alpha, \beta, \gamma, \delta, z)$ as before. Then, for any $\xi \in \widehat{HF}(\alpha, \beta, \mathfrak{s})$,*

$$\widehat{F}_{\alpha,\gamma,\delta}(\widehat{F}_{\alpha,\beta,\gamma}(\xi \otimes \widehat{\Theta}_{\beta,\gamma}) \otimes \widehat{\Theta}_{\gamma,\delta}) = \widehat{F}_{\alpha,\beta,\delta}(\xi \otimes \widehat{F}_{\beta,\gamma,\delta}(\widehat{\Theta}_{\beta,\gamma} \otimes \widehat{\Theta}_{\gamma,\delta})),$$

and for $\xi \in HF^\infty(\alpha, \beta, \mathfrak{s})$,

$$F_{\alpha,\gamma,\delta}^\infty(F_{\alpha,\beta,\gamma}^\infty(\xi \otimes \Theta_{\beta,\gamma}^{\leq}) \otimes \Theta_{\gamma,\delta}^{\leq}) = F_{\alpha,\beta,\delta}^\infty(\xi \otimes F_{\beta,\gamma,\delta}^\infty(\Theta_{\beta,\gamma}^{\leq} \otimes \Theta_{\gamma,\delta}^{\leq})).$$

Proof. This follows from associativity of the triangle construction.

First, we begin with some remarks on admissibility for the pointed Heegaard quadruple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, z)$. 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 $H^2(X_{\alpha, \beta, \gamma, \delta}) \rightarrow H^2(\partial X_{\alpha, \beta, \gamma, \delta})$ is injective. Since the restriction of $\delta H^1(Y_{\beta, \delta}) + \delta H^1(Y_{\alpha, \gamma})$ to the boundary is trivial, the subgroup itself is trivial. Thus, the orbits $\mathfrak{S} \subset \text{Spin}^c(X_{\alpha, \beta, \gamma, \delta})$ needed for admissibility (in the sense of Subsection 8.4.2) are trivial. Moreover, it is now easy to see that if we choose $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ all sufficiently close to one another, so the six Heegaard diagrams are strongly admissible, then the quadruple $X_{\alpha, \beta, \gamma, \delta}$ is strongly admissible.

The result is now a direct application of Theorem 8.16. \square

LEMMA 9.7. *The triple $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, z)$ is admissible in respect to z , and (for the orientation conventions from Lemmas 9.1 and 9.4),*

$$\begin{aligned} \widehat{F}_{\beta, \gamma, \delta}(\Theta_{\beta, \gamma} \otimes \Theta_{\gamma, \delta}) &= \widehat{\Theta}_{\beta, \delta}, \\ F_{\beta, \gamma, \delta}^{\leq}(\Theta_{\beta, \gamma}^{\leq} \otimes \Theta_{\gamma, \delta}^{\leq}) &= \Theta_{\beta, \delta}^{\leq}. \end{aligned}$$

Proof. The admissibility is easy to check. Now let Δ_i denote the triangular region in $\Sigma - \{\beta_i\} - \{\gamma_i\} - \{\delta_i\}$ with vertices y_i^+, v_i^+, w_i^+ , cf. Figure 8. Let $\psi \in \pi_2(\mathbf{y}^+, \mathbf{v}^+, \mathbf{w}^+)$ be given by $\mathcal{D}(\psi) = \sum_{i=1}^g \Delta_i$. Then all $\psi' \neq \psi$, with $n_z(\psi) = 0$, have some negative coefficients so that $\mathcal{M}(\psi') = \emptyset$. To study $\mathcal{M}(\psi)$, note that for each i there is a unique map from the two-simplex to Δ_i that satisfies the boundary conditions. The corresponding g -tuple of maps to Σ gives the unique solution in $\mathcal{M}(\psi)$, which is easily seen to be smooth. Now, the coefficient of \mathbf{w}^+ in $\widehat{f}(\mathbf{y}^+ \otimes \mathbf{v}^+)$ is ± 1 . For all $\mathbf{w} \neq \mathbf{w}^+$ we have $\text{gr}(\mathbf{y}^+, \mathbf{v}^+, \mathbf{w}) = \text{gr}(\mathbf{y}^+, \mathbf{v}^+, \mathbf{w}^+) + \text{gr}(\mathbf{w}^+, \mathbf{w}) > 0$. Thus, $\widehat{F}(\Theta_{\beta, \gamma} \otimes \Theta_{\gamma, \delta}) = \Theta_{\beta, \delta}$. More precisely, with orientation conventions on $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, z)$ and $(\Sigma, \boldsymbol{\gamma}, \boldsymbol{\delta}, z)$ as in Lemmas 9.1 and 9.4 (i.e. for which \mathbf{y}^+ and \mathbf{v}^+ are cycles), then for the induced orientation convention on $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\delta}, z)$, \mathbf{w}^+ is a cycle and so there, too, we are using the orientation convention of Lemma 9.1. The corresponding result for f^{\leq} follows similarly. \square

PROPOSITION 9.8. *If the δ are sufficiently close to the β , then the maps $f_{\alpha,\beta,\delta}(\cdot \otimes \Theta_{\beta,\delta})$ induce isomorphisms*

$$HF^\infty(\Sigma, \alpha, \beta, \mathfrak{s}) \cong HF^\infty(\Sigma, \alpha, \delta, \mathfrak{s}),$$

and indeed isomorphisms of all the other homology groups \widehat{HF} , HF^- , and HF^+ .

In fact, the above proof is slightly simpler in the case where $c_1(\mathfrak{s})$ is torsion, so we start with a proof in this case. Indeed, we will find it useful to introduce here an “energy filtration” on CF^∞ . To this end, we introduce some terminology.

Definition 9.9. A *filtered group* is a free Abelian group C which is freely generated by a distinguished set of generators \mathfrak{S} , and equipped with a map

$$\mathcal{F}: \mathfrak{S} \longrightarrow \mathbb{R}.$$

If ξ, η are any two elements in C , we say that $\xi < \eta$, if, writing $\xi = \sum_{\mathbf{x} \in \mathfrak{S}} \xi_{\mathbf{x}} \cdot \mathbf{x}$ and $\eta = \sum_{\mathbf{x} \in \mathfrak{S}} \eta_{\mathbf{x}} \cdot \mathbf{x}$ (where $\xi_{\mathbf{x}}, \eta_{\mathbf{x}} \in \Lambda$), we have that

$$\max\{\mathcal{F}(\mathbf{x}) \mid \xi_{\mathbf{x}} \neq 0\} < \max\{\mathcal{F}(\mathbf{x}) \mid \eta_{\mathbf{x}} \neq 0\}.$$

A filtration is *bounded below* if for each real number r , $(\mathcal{F}(\mathfrak{S})) \cap (-\infty, r]$ is a finite set.

There are, of course, more general notions where \mathbb{R} is replaced by an arbitrary partially ordered set, but we have no need for those presently.

The basic property is the following:

LEMMA 9.10. *Let $F: A \longrightarrow B$ be a map of filtered groups, which can be decomposed as a sum $F = F_0 + \ell$ where F_0 is a filtration-preserving isomorphism, and $\ell(\mathbf{x}) < F_0(\mathbf{x})$ for all $\mathbf{x} \in \mathfrak{S}$. Then, if the filtration on B is bounded below, we can conclude that F is an isomorphism of groups.*

Proof. Straightforward. □

When $c_1(\mathfrak{s})$ is torsion and $(\Sigma, \alpha, \beta, z)$ is strongly \mathfrak{s} -admissible, then that Heegaard diagram is, of course, weakly admissible for all Spin^c structures. Fix a reference point $\mathbf{x}_0 \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ which represents \mathfrak{s} , and equip Σ with a volume form for which the signed area of each periodic domain is zero. The *energy filtration* on $CF^\infty(\Sigma, \alpha, \beta, \mathfrak{s})$ is the filtration defined by $\mathcal{F}[\mathbf{x}, i] = -\mathcal{A}(\phi)$ where $\phi \in \pi_2(\mathbf{x}_0, \mathbf{x})$ is any homotopy class with $n_z(\phi) = -i$, and $\mathcal{A}(\phi)$ refers to its signed area.

Proof of Proposition 9.8 when $c_1(\mathfrak{s})$ is torsion. We endow $(\Sigma, \alpha, \beta, \mathfrak{s})$ with a volume form as above, and arrange for the δ_i to be exact Hamiltonian translates of the β_i , with respect to this volume form (so that all the β - δ periodic domains also have total signed area equal to zero). When all the δ_i are sufficiently close to their corresponding β_i , and all the α_j intersect the β_i transversally (as we always assume), then to each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, there is a uniquely associated point $\mathbf{x}' \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$ which is closest to \mathbf{x} . We let

$$\widehat{f}_0: \widehat{CF}(\alpha, \beta, \mathfrak{s}) \longrightarrow \widehat{CF}(\alpha, \delta, \mathfrak{s})$$

denote the group homomorphism induced by this closest point map. Clearly, this is an isomorphism of groups.

Let $\mathbf{w}^+ \in \mathbb{T}_\beta \cap \mathbb{T}_\delta$ denote the intersection point representing the class $\widehat{\Theta}_{\beta, \delta}$. Note that there is a unique “small triangle” homotopy class $\psi_0^\mathbf{x} \in \pi_2(\mathbf{x}, \mathbf{w}^+, \mathbf{x}')$, which is the only triangle satisfying the two properties that $\mathcal{D}(\psi_0^\mathbf{x}) \geq 0$, and $\mathcal{D}(\psi_0^\mathbf{x})$ is supported inside the support of the isotopy between β and δ . Indeed, $\mathcal{D}(\psi_0^\mathbf{x})$ has multiplicity one on the small triangles $\Delta(x_i, w_i^+, x'_i)$ connecting x_i , w_i^+ , and x'_i (for $i = 1, \dots, g$), and zero everywhere else. We claim that if J is a family which is a sufficiently small perturbation of the constant complex structure j , then $\#\mathcal{M}_J(\psi_0^\mathbf{x}) = 1$. To see this, we first consider complex structures of the form $J \equiv \text{Sym}^g(j)$. By elementary complex analysis we see that there is a unique holomorphic map f_i from the 2-simplex Δ to the triangle $\Delta(x_i, w_i^+, x'_i)$ that lines up the corresponding boundary segments. It follows that the moduli space $\mathcal{M}_{\text{Sym}^g(j)}(\psi)$ has a unique element, which is given as the product of f_1, \dots, f_g from the trivial g -fold cover of Δ to Σ (with respect to the constant almost-complex structure $\text{Sym}^g(j)$). It is easy to see that this is a smooth solution. Counted with the appropriate orientation, $\#(\mathcal{M}_{\text{Sym}^g(j)}(\psi)) = 1$. Thus, this assertion persists for all sufficiently small perturbations of the constant family.

It follows that we can decompose $\widehat{f}_{\alpha, \beta, \delta}(\cdot \otimes \widehat{\Theta}_{\beta, \delta})$ as

$$\widehat{f}_{\alpha, \beta, \delta}(\cdot \otimes \Theta_{\beta, \delta}) = \widehat{f}_0 + \widehat{\ell},$$

where $\widehat{\ell}$ counts holomorphic triangles ψ , with $n_z(\psi) = 0$, whose domains are not supported inside the small region bounded by β and δ . In fact, let ε denote the total (unsigned) area swept out by the isotopy between the β and the δ which, as its notation suggests, can be made arbitrarily small by varying δ . Let M_0 denote the minimal area of any (nonzero) domain in $\Sigma - \alpha - \beta$ (note that this is independent of δ). It is easy to see that the triangles counted by ℓ all have total area bounded below by $M_0 - \varepsilon$. Choose a reference point $\mathbf{x}_0 \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and consider the induced area filtration on $(\Sigma, \alpha, \delta, \mathfrak{s})$ induced from \mathbf{x}'_0 . Since for each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, the area of the canonical triangle $\psi_0^\mathbf{x}$ is bounded by ε , if we arrange that $\varepsilon < M_0/2$, then

$$\widehat{\ell}(\mathbf{x}) < \widehat{f}_0(\mathbf{x}),$$

with respect to the energy filtration induced on $\widehat{CF}(\alpha, \delta, \mathfrak{s})$. By Lemma 9.10, it follows that \widehat{f} , which we already know is a chain map, is actually an isomorphism of chain complexes.

By modifying the above constructions, we see that

$$f_{\alpha, \beta, \delta}^{\infty}(\cdot \otimes \Theta_{\beta, \delta}^{\leq}) = f_0^{\infty} + \text{lower order},$$

where

$$f_0^{\infty}[\mathbf{x}, i] = [\mathbf{x}', i].$$

Although now the filtration on $CF^{\infty}(\alpha, \beta, z)$ is no longer bounded below, it comes with a relative \mathbb{Z} -grading, and the map f^{∞} preserves this relative grading, in the sense that if we have $\xi, \xi' \in CF^{\infty}(\alpha, \beta, \mathfrak{s})$ with $\text{gr}(\xi, \xi') = 0$, then

$$\text{gr}(f_{\alpha, \beta, \delta}^{\infty}(\xi \otimes \Theta_{\beta, \delta}^{\leq}), f_{\alpha, \beta, \delta}^{\infty}(\xi' \otimes \Theta_{\beta, \delta}^{\leq})) = 0.$$

Now, the filtration induced on the subset of $CF^{\infty}(\alpha, \beta, \mathfrak{s})$ with fixed relative grading is easily seen to be bounded below (since, in fact, there are only finitely many generators representing any given degree). \square

When $c_1(\mathfrak{s})$ is nontorsion, we must use a refinement, since the relative grading on CF^{∞} is no longer a relative \mathbb{Z} -grading, but only a $\mathbb{Z}/\delta\mathbb{Z}$ -grading (where $\delta = \delta(\mathfrak{s})$ is the indeterminacy from Equation (12)), and the complex for each fixed degree might be an infinitely generated \mathbb{Z} -module.

So, 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 4.12. Given any $[\mathbf{x}, i]$ and $[\mathbf{y}, j]$ with the same relative grading, we can find some disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $n_z(\phi) = i - j$ and $\mu(\phi) = 0$. In each relative degree, then, there is a filtration (uniquely defined up to an additive constant) defined by

$$\mathcal{F}([\mathbf{x}, i]) - \mathcal{F}([\mathbf{y}, j]) = -\mathcal{A}(\mathcal{D}(\phi)).$$

Since any two possible choices of such disks ϕ, ϕ' differ by a renormalized periodic domain, it follows that the difference in filtration defined above is independent of the choice of disk.

Proof of Proposition 9.8. when $c_1(\mathfrak{s})$ is nontorsion. We consider CF^{-} , and use the notation from the earlier proof, writing

$$f_{\alpha, \beta, \delta}^{\infty} = f_0^{\infty} + \ell^{\infty},$$

where $f_0^{\infty}[\mathbf{x}, i] = [\mathbf{x}', i]$ as before. Clearly, for given $[\mathbf{x}, i]$, $f_0^{\infty}([\mathbf{x}, i])$ and $\ell^{\infty}([\mathbf{x}, i])$ have the same relative degree.

We claim that, just as before, if the δ are sufficiently close to the β , then ℓ^{∞} has lower order (than f_0^{∞}) with respect to the above refined energy filtration. To see this, let $[\mathbf{y}, j]$ be any element appearing with nonzero multiplicity

in $\ell^\infty([\mathbf{x}, i]) = [\mathbf{y}, j]$ (so in particular $\mathbf{y} \neq \mathbf{x}'$). Then, there is a pseudo-holomorphic triangle $\psi \in \pi_2(\mathbf{x}, \mathbf{w}^+, \mathbf{y})$ with $\mathcal{D}(\psi) \geq 0$ and $\mu(\psi) = 0$. It is easy to see that in this case, there is also a $\phi \in \pi_2(\mathbf{x}', \mathbf{y})$ with $\psi = \psi_0^{\mathbf{x}} * \phi$, so that $\mu(\phi) = 0$ and $n_z(\phi) = i - j$. Since $\mathcal{D}(\psi)$ is not supported inside the region between β and γ , we know that $\mathcal{A}(\psi) > M_0 - \varepsilon$. Thus, it follows that

$$\mathcal{A}(\phi) = \mathcal{A}(\psi) - \mathcal{A}(\psi_0) > M_0 - 2\varepsilon,$$

so, since $-\mathcal{A}(\phi)$ determines the filtration difference between $f_0^\infty[\mathbf{x}, i]$ and $\ell^\infty[\mathbf{x}, i]$, if we choose the exact Hamiltonian translates to be sufficiently close, $\ell^\infty([\mathbf{x}, i]) < f_0^\infty[\mathbf{x}, i]$.

We claim that the refined filtration is bounded below. To see this, observe that if δ denotes the divisibility of $c_1(\mathfrak{s})$, then for each $[\mathbf{x}, i]$ representing a fixed relative, we can find a Spin^c -renormalized periodic domain connecting $[\mathbf{x}, i]$ to $[\mathbf{x}, i + \delta]$, so these elements all have the same filtration. Thus, Lemma 9.10 applies, to prove that $f_{\alpha, \beta, \delta}^\infty$ induces an isomorphism of chain complexes. \square

Proof of Theorem 9.5. This follows easily from Lemma 9.6 and Proposition 9.8. A direct application of these results shows that the map induced from $\otimes_{\beta, \gamma}$ is injective, while the map induced from γ, δ is surjective. But the roles of these g -tuples is symmetric: we can introduce a fifth tuple of circles η which are small isotopic translates of the γ , and apply the above reasoning to see that the map induced from $\otimes_{\beta, \gamma}$ is surjective, as well.

To see that the map commutes with the $H_1(Y; \mathbb{Z})$ -module structure, we represent the action by a codimension-one constraint $V \in \mathbb{T}_\alpha$ as in Remark 4.20, and consider the moduli space

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

where $E_\alpha: \mathbb{R}$ is a parametrization of the α -edge as in the proof of Proposition 8.13. As usual, when $\mu(\psi) = -1$, this space is compact, and can be used to construct a chain homotopy

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

Consider homotopy classes of triangles $\pi_2(\mathbf{x}, \Theta_{\beta, \gamma}, \mathbf{y})$ with $\mu(\psi) = 0$. The ends as $\tau \mapsto \infty$ correspond to the commutator of $F^\infty(\cdot \otimes \Theta_{\beta, \gamma})$ with the action of V ; the other ends correspond to the commutator of H with the boundary maps. \square

9.3. The Maslov index of a periodic domain. We can now prove Theorem 4.9, which was used in the definition of the relative gradings.

Proof of Theorem 4.9. Fix $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and consider the map $\text{Hom}(H_2(Y; \mathbb{Z}), 2\mathbb{Z})$ 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 $m_{\mathfrak{s}} \in \text{Hom}(H_2(Y; \mathbb{Z}), \mathbb{Z})$.

We argue that $m_{\mathfrak{s}}$ depends on Y alone, i.e. it is invariant under pointed isotopies, pointed handleslides, and stabilization (cf. Proposition 7.1). To see stabilization invariance, it suffices to see how the Maslov index changes by adding $S \in \pi_2(\text{Sym}^g(\Sigma))$, and thereby reducing to the case where the coefficient of the domain is zero on the two-torus. Handleslide invariance follows from index calculations parallel to (but much simpler than) Theorem 9.5. Specifically, let α, β, γ be attaching circles, where γ are obtained from β by a handleslide and a small Hamiltonian isotopy. Indeed, if the pair of pants used for the handleslide is sufficiently small, we can associate to $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ a nearest intersection point $\mathbf{x}' \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, which is connected to \mathbf{x} by a small triangle $\Delta \in \pi_2(\mathbf{x}, \theta, \mathbf{x}')$ with $n_z = 0$ and $\mu(\Delta) = 0$ (here, θ is an intersection point representing the homology class $\hat{\Theta}_{\beta, \gamma}$ from Lemma 9.4). Now, in view of the affine identification $\pi_2(\mathbf{x}, \theta, \mathbf{x}') \cong \mathbb{Z} \oplus H^1(Y; \mathbb{Z})$ (cf. Propositions 8.3 and 8.2), if $p_{\mathbf{x}}$ is a periodic class for (Σ, α, β) , and $p'_{\mathbf{x}}$ is the corresponding periodic class for $(\Sigma, \alpha, \gamma, z)$, 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 (this is established in the course of the proof of Lemma 9.4), 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 5.6 (see Remark 5.7).

Now, we argue that if $\mathfrak{s}, \mathfrak{s}' \in \text{Spin}^c(Y)$ are represented by intersection points, then

$$m_{\mathfrak{s}} = m_{\mathfrak{s}'} + 2c,$$

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 2.19. 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 $m_{\mathfrak{s}} = m_{\mathfrak{s}'} + 2c$.

Now, $m_{\mathfrak{s}} = c_1(\mathfrak{s}) + K$, in $\text{Hom}(H_2(Y; \mathbb{Z}), \mathbb{Z})$ 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 $m_{\bar{\mathfrak{s}}}$. Switching the roles of \mathbb{T}_α and \mathbb{T}_β 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 these

new data, we see 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 domain, 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 $\langle m_{\mathfrak{s}}, H(\mathcal{P}) \rangle = \langle m_{\bar{\mathfrak{s}}}, H(\mathcal{P}') \rangle$. However, $H(\mathcal{P}) = -H(\mathcal{P}')$. Thus, $m_{\bar{\mathfrak{s}}} = -m_{\mathfrak{s}}$. Since it is also true that $c_1(\mathfrak{s}) = -c_1(\bar{\mathfrak{s}})$, it follows that $K = 0$. \square

10. Stabilization

The final step in establishing topological invariance of the Floer homologies is stabilization invariance. Fix a strongly \mathfrak{s} -admissible pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ for Y , where Σ is a surface of genus g and $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_g\}$, $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_g\}$. The stabilized diagram $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', z)$ is obtained by forming the connected sum $\Sigma' = \Sigma \# E$, which is the connected sum of Σ with E , a surface of genus g , and letting $\boldsymbol{\alpha}' = \{\alpha_1, \dots, \alpha_g, \alpha_{g+1}\}$ and $\boldsymbol{\beta}' = \{\beta_1, \dots, \beta_g, \beta_{g+1}\}$, where α_{g+1} and β_{g+1} are a pair of circles in the E summand which meet transversally in a single positive point c . For simplicity, we choose the point in Σ along which we form the connected sum to lie in the same path-component of $\Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$ as z . (For this choice, there is a one-to-one correspondence between periodic domains for the two diagrams, and hence the notions of admissibility coincide.)

We begin with the much simpler case of \widehat{CF} .

THEOREM 10.1. *Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a strongly \mathfrak{s} -admissible pointed Heegaard diagram for Y , and let $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', z)$ denote its stabilization. Then, for each orientation system \mathfrak{o} on the original Heegaard diagram, there are an induced orientation system \mathfrak{o}' on its stabilization, and a corresponding isomorphism of $\Lambda^* H_1(Y; \mathbb{Z})/\text{Tors}$ -modules:*

$$\widehat{HF}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathfrak{s}, \mathfrak{o}) \cong \widehat{HF}(\boldsymbol{\alpha}', \boldsymbol{\beta}', \mathfrak{s}, \mathfrak{o}').$$

Proof. It is easy to see that the intersection points correspond: $\mathbb{T}'_{\alpha} \cap \mathbb{T}'_{\beta} = (\mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}) \times \{c\}$. We choose a basepoint z in the same path-component of $\Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$ as $\sigma_1 \in \Sigma$, the point along which we perform the connected sum to obtain Σ' . Let z' denote the corresponding basepoint in Σ' . If $\mathbf{x} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$, and $\mathbf{x}' \in \mathbb{T}'_{\alpha} \cap \mathbb{T}'_{\beta}$ is the corresponding point $\mathbf{x} \times \{c\}$, then the induced Spin^c structures agree, $\mathfrak{s}_z(\mathbf{x}) = \mathfrak{s}_{z'}(\mathbf{x}')$, since the corresponding vector fields agree away from the three-ball where the stabilization occurs.

Let $\mathbf{x}, \mathbf{y} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$, and $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ be the class with coefficient $n_z(\phi) = 0$. Let $\mathbf{x}' = \mathbf{x} \times \{c\}$, $\mathbf{y}' = \mathbf{y} \times \{c\}$, and $\phi' \in \pi_2(\mathbf{x}', \mathbf{y}')$ be the class with $n_{z'}(\phi') = 0$. Then, we argue that for certain special paths of almost-complex structures, the moduli space $\mathcal{M}(\phi')$ is identified with $\mathcal{M}(\phi) \times \{c\}$ (together

with its deformation theory, including the determinant line bundles). Hence, the chain complexes are identical.

To set up the complex structure, we let σ_1 and σ_2 denote the connected sum points for Σ and E respectively. Recall that there is a holomorphic map

$$\mathrm{Sym}^g(\Sigma - B_{r_1}(\sigma_1)) \times (E - B_{r_2}(\sigma_2)) \longrightarrow \mathrm{Sym}^{g+1}(\Sigma'),$$

which is a diffeomorphism onto its image. Suppose that J_s is a (generic) family of almost-complex structures over $\mathrm{Sym}^g(\Sigma)$ which agrees with $\mathrm{Sym}^g(j)$ in a tubular neighborhood of $\{\sigma_1\} \times \mathrm{Sym}^{g-1}(\Sigma)$, a neighborhood which is holomorphically identified with $B_\varepsilon(\sigma_1) \times \mathrm{Sym}^{g-1}(\Sigma)$. Using the above product map, we can transfer $J_s \times j_E$ to an open subset of $\mathrm{Sym}^{g+1}(\Sigma')$, and extend it to all of $\mathrm{Sym}^{g+1}(\Sigma')$. It follows from our choice of J_s that if $n_z(\phi) = 0$, then any J_s -holomorphic representative for ϕ must have its image in $\mathrm{Sym}^g(\Sigma - B_\varepsilon)$, and hence its product with the constant map is $(J_s \times j_E)$ -holomorphic in $\mathrm{Sym}^{g+1}(\Sigma')$. Conversely, a $(J_s \times j_E)$ -holomorphic curve in $\mathrm{Sym}^{g+1}(\Sigma')$ with $n_{z'}(\phi') = 0$ must be contained in $\mathrm{Sym}^g(\Sigma) \times \{c\}$. The identification of deformation theories is straightforward.

Since the map we describe identifies individual moduli spaces it is clear (cf. Remark 4.20) that the map is equivariant under the action of $H_1(Y; \mathbb{Z})/\mathrm{Tors}$. \square

The corresponding fact for HF^\pm and HF^∞ is more subtle, and depends on a gluing theorem for holomorphic curves.

THEOREM 10.2. *Let $(\Sigma, \alpha, \beta, z)$ be a strongly \mathfrak{s} -admissible pointed Heegaard diagram for Y , and let $(\Sigma', \alpha', \beta', z)$ denote its stabilization. Then, for each orientation system \mathfrak{o} on the original Heegaard diagram, there are an induced orientation system \mathfrak{o}' on its stabilization, and corresponding $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(Y; \mathbb{Z})/\mathrm{Tors}$ -module isomorphisms*

$$HF^\infty(\alpha, \beta, \mathfrak{s}, \mathfrak{o}) \cong HF^\infty(\alpha', \beta', \mathfrak{s}, \mathfrak{o}') \text{ and } HF^\pm(\alpha, \beta, \mathfrak{s}, \mathfrak{o}) \cong HF^\pm(\alpha', \beta', \mathfrak{s}, \mathfrak{o}').$$

The proof of this theorem occupies the rest of this section. However, we give a brief outline of it here.

Sketch of proof. As in the proof of Theorem 10.1 above, the generators for the complexes are identified. We fix some complex structure j over Σ , and j_E over E . Let $j'(T)$ denote the complex structure on Σ' obtained by inserting a cylinder $[-T, T] \times S^1$ between the Σ and E . Correspondingly, the symmetric product $\mathrm{Sym}^{g+1}(\Sigma')$ is endowed with a complex structure $\mathrm{Sym}^{g+1}(j'(T))$, which admits an open subset which is holomorphically identified with

$$\mathrm{Sym}^g(\Sigma - B_{r_1}(\sigma_1)) \times \mathrm{Sym}^1(E - B_{r_2}(\sigma_2)),$$

given the complex structure obtained by restricting $\mathrm{Sym}^g(j) \times j_E$. Using this parametrization, we can transfer some fixed initial path J_s of nearly-

symmetric almost-complex structures over $\text{Sym}^g(\Sigma)$ to a path $J'_s(T)$ of $j'(T)$ -nearly-symmetric almost-complex structures which extend over $\text{Sym}^{g+1}(\Sigma')$.

For T sufficiently large, we show that for any pair $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, if $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is the homotopy class with $\mu(\phi) = 1$, then there is an identification $\mathcal{M}_{J_s}(\phi) \cong \mathcal{M}_{J'_s(T)}(\phi')$, where $\phi' \in \pi_2(\mathbf{x}', \mathbf{y}')$ is the corresponding class with $\mu(\phi') = 1$, and $\mathbf{x}' = \mathbf{x} \times \{c\}$, $\mathbf{y}' = \mathbf{y} \times \{c\}$ (see Theorem 10.4 below). It is an easy consequence of this that all the relevant corresponding chain complexes are identified:

$$\begin{aligned} (CF^\pm(\alpha, \beta, \mathfrak{s}), \partial_{J_s}^\pm) &\cong (CF^\pm(\alpha', \beta', \mathfrak{s}), \partial_{J_s(T)}^\pm), \\ (CF^\infty(\alpha, \beta, \mathfrak{s}), \partial_{J_s}^\infty) &\cong (CF^\infty(\alpha', \beta', \mathfrak{s}), \partial_{J_s(T)}^\infty). \end{aligned}$$

Since the homology groups of these chain complexes are independent of the paths of almost-complex structures, it follows that the groups are isomorphic; cf. the last paragraph in Subsection 10.1.

The identification between the moduli spaces is given by a now familiar gluing construction. Suppose that all flows in $\mathcal{M}_{J_s}(\phi)$ meet the subvariety $\{\sigma\} \times \text{Sym}^{g-1}(\Sigma)$ in general position; i.e., suppose that each Maslov index-one flow-line $u \in \mathcal{M}(\phi)$ in $\text{Sym}^g(\Sigma)$ meets the above subvariety corresponding to the connected sum point $\sigma \in \Sigma$ transversally (this can be easily arranged). Then, given such a flow $u \in \mathcal{M}(\phi)$, we have that the pre-image

$$u^{-1}(\{\sigma\} \times \text{Sym}^{g-1}(\Sigma)) = \{q_1, \dots, q_n\},$$

where $n = n_z(\phi)$. Taking the product of u with $\{c\}$, we obtain a map to $\text{Sym}^g(\Sigma) \times E$, which can be viewed as a subset of $\text{Sym}^{g+1}(\Sigma \vee E)$, which in turn can be thought of as a degenerated version of $\text{Sym}^{g+1}(\Sigma')$. Splicing in spheres in $\text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E)$ (which can also be thought of as a subset of the $(g+1)$ -fold symmetric product of the wedge), we obtain for any sufficiently large T , a nearly J'_s -holomorphic map into $\text{Sym}^{g+1}(\Sigma \#_T E)$ for all sufficiently large T . As the parameter $T \mapsto \infty$, this spliced map becomes closer to being $J'_s(T)$ -holomorphic (in an appropriate norm, as discussed below). Using the inverse function theorem in the usual manner, we obtain a nearby pseudo-holomorphic map $u' \in \mathcal{M}_{J'_s(T)}(\phi')$.

A Gromov compactness argument shows that for sufficiently large T , all of the moduli spaces $\mathcal{M}_{J'_s(T)}(\phi')$ lie in the domain of this gluing map. (Note also that in the case where $g = 1$ not every homotopy class $\phi' \in \pi_2(\mathbf{x}', \mathbf{y}')$ is obtained by stabilizing a class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, e.g. there may not be any homotopy class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = 1$. Nonetheless, the compactness argument applies, showing that then the moduli space $\mathcal{M}_{J'_s(T)}(\phi')$ is empty for sufficiently large T .)

Note that the $\mathbb{Z}[U]$ -equivariance of the identifications is a formal consequence of the fact that $n_z(\phi) = n_{z'}(\phi')$; the $H_1(Y; \mathbb{Z})/\text{Tors}$ -equivariance also follows similarly. \square

An important ingredient in the above story is a description of the spheres in $\text{Sym}^2(E)$. In particular, it is crucial that through each pair of points in E , there is a unique holomorphic sphere in $\text{Sym}^2(E)$ in the positive generator of $\pi_2(\text{Sym}^2(E))$. This follows from the next result:

LEMMA 10.3. *Let E be a genus one Riemann surface. The second symmetric product $\text{Sym}^2(E)$ is naturally a ruled surface over E .*

Proof. The map from $\text{Sym}^2(E)$ to the base E is the map sending the pair $\{x, y\}$ to the sum $x + y$. The fiber over a base point a is represented by pairs of the form $\{a + w, -w\}$ with $w \in E$. Two pairs $(a + w_1, -w_1)$ and $(a + w_2, -w_2)$ are equivalent if $-a - w_1 = w_2$. Now, the map $w \mapsto -a - w$ is an involution whose quotient is a projective line. \square

Ultimately, the proof of Theorem 10.2 outlined above is a modification of the usual picture for dealing with noncompactness in Gromov theory ([30], [23], [21]). The gluing here is analytically closely related to gluing problems which arose in the study of the Yang-Mills equations (see for example [36], [24]). Moreover, the degeneration of Σ' above is closely related to the degenerations studied by Ionel-Parker and Li-Ruan which concern degenerations of symplectic manifolds along codimension one symplectic submanifolds (see [18], [20]). The rest of this section is devoted to proving Theorem 10.2.

10.1. *Gluing: the statement.* In this subsection, we state the result which is the cornerstone of Theorem 10.2 above, allowing us to use flows in $\text{Sym}^g(\Sigma)$ to construct flows in $\text{Sym}^{g+1}(\Sigma \#_T E)$ for sufficiently large T .

For simplicity, we assume henceforth that the path J_s of almost-complex structures over $\text{Sym}^g(\Sigma)$ agrees with $\text{Sym}^g(j)$ for some complex structure j over Σ in a neighborhood of the subset $\text{Sym}^{g-1}(\Sigma) \times \{\sigma\} \subset \text{Sym}^g(\Sigma)$ (i.e. we are using (j, η, V) -nearly-symmetric almost-complex structures for a choice of V containing $\{\sigma\} \times \text{Sym}^{g-1}(\Sigma)$).

Moreover, as indicated earlier, the path J_s can be used to construct a one-parameter family $J'_s(T)$ of paths over $\text{Sym}^{g+1}(\Sigma')$ compatibly with J_s in the following sense. Recall that $\text{Sym}^{g+1}(\Sigma \#_T E)$ (endowed with the complex structure $\text{Sym}^{g+1}(j'(T))$, where $j'(T) = j \#_T j_E$) has an open subset holomorphically identified with:

$$\text{Sym}^g(\Sigma - B_{r_1}(\sigma_1)) \times \text{Sym}^1(E - B_{r_2}(\sigma_2)),$$

where r_1, r_2 are nonnegative real numbers. Fix another pair of real numbers $R_1 > r_1, R_2 > r_2$. We choose J'_s with the following properties:

- over $\text{Sym}^g(\Sigma - B_{R_1}(\sigma_1)) \times \text{Sym}^1(E - B_{R_2}(\sigma_2))$, the path J'_s agrees with $J_s \times j_E$,

- over $\text{Sym}^g(\Sigma - B_{r_1}(\sigma_1)) \times \text{Sym}^1(B_{R_2}(\sigma_2) - B_{r_2}(\sigma_2))$, as the normal parameter to σ_2 goes from R_2 to r_2 , J'_s connects from $J_s \times \mathbf{j}_E$ to $\text{Sym}^g(\mathbf{j}) \times \mathbf{j}_E$ (for each s), always splitting as a product of some almost-complex structure on the first factor with \mathbf{j}_E on the second,
- over the rest of $\text{Sym}^{g+1}(\Sigma')$, $J'_s \equiv \text{Sym}^{g+1}(\mathbf{j}'(T))$.

Note that this description fits together continuously, in light of our hypothesis that $J_s \equiv \text{Sym}^g(\mathbf{j})$ near $\text{Sym}^{g-1}(\Sigma) \times \{\sigma_1\}$.

Fix a class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, with $n_\sigma(\phi) = n$ for $\sigma \in \Sigma - \alpha_1 - \cdots - \alpha_g - \beta_1 - \cdots - \beta_g$. There is a naturally induced map

$$\rho_1: \mathcal{M}_{J_s}(\phi) \longrightarrow \text{Sym}^n(\text{Sym}^{g-1}(\Sigma) \times \text{Sym}^1(E))$$

given by

$$\rho_1(u) = u(u^{-1}(\text{Sym}^{g-1}(\Sigma) \times \{\sigma\})) \times \{c\}.$$

There is a complex codimension one subset in the range corresponding to the subvariety $\text{Sym}^{g-2}(\Sigma) \times \{\sigma\} \subset \text{Sym}^{g-1}(\Sigma)$. For general choice of σ , all one-dimensional moduli spaces of flows miss this special locus — i.e. the images of u in the symmetric product never contain the prospective connected sum point σ with multiplicity greater than one.

Consider the moduli space of (unparametrized) holomorphic maps

$$\mathcal{M}\left(\bigcup_{i=1}^n \mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E)\right)$$

from n disjoint Riemann spheres, whose restriction to each component represents the generator of $\pi_2(\text{Sym}^2(E))$ (and hence is constant on the first factor).

There is a map

$$\rho_2: \mathcal{M}\left(\bigcup_{i=1}^n \mathbb{S} \rightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E)\right) \rightarrow \text{Sym}^n(\text{Sym}^{g-1}(\Sigma) \times \text{Sym}^1(E)),$$

given by

$$\rho_2(v) = v(v^{-1}(\text{Sym}^{g-1}(\Sigma) \times \{\sigma \times \text{Sym}^1(E)\})).$$

In the next section, we show how to splice the spheres coming from this moduli space to the disks in $\mathcal{M}(\phi) \times \{c\}$, giving rise to a map from the fibered product of $\rho_1 \times \rho_2$ to the space of maps connecting $\mathbf{x}' = \mathbf{x} \times \{c\}$ to $\mathbf{y}' = \mathbf{y} \times \{c\}$ which are nearly holomorphic. Indeed, this fibered product description is particularly simple in view of the fact that ρ_2 is a diffeomorphism (see Lemmas 10.7 and 10.11 below). In particular, each $u \in \mathcal{M}_{J_s}(\phi)$ has a corresponding unique $v \in \mathcal{M}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2))$ with $\rho_1(u) = \rho_2(v)$.

Our aim, then is to prove the following:

THEOREM 10.4. *Fix $\mathbf{x}, \mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and a class $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = 1$. For $\mathbf{x}' = \mathbf{x} \times \{c\}$ and $\mathbf{y}' = \mathbf{y} \times \{c\}$, let $\phi' \in \pi_2(\mathbf{x}', \mathbf{y}')$ be the class with the property that $n_\sigma(\phi) = n_{\sigma'}(\phi')$ where σ is any point in $\Sigma - \alpha_1 - \dots - \alpha_g - \beta_1 - \dots - \beta_g$, and σ' is the corresponding point in Σ' . Then, for all sufficiently large T , there is a diffeomorphism $\mathcal{M}_{J_s}(\phi) \cong \mathcal{M}_{J'_s(T)}(\phi')$.*

Remark 10.5. The diffeomorphism statement above is to be interpreted as an identification between deformation theories. In particular, $\mu(\phi) = \mu(\phi')$.

We argue that the $J'_s(T)$ are $j'(T)$ -nearly symmetric, and can be used to calculate the Floer homologies in $\text{Sym}^{g+1}(\Sigma')$. Fix a T so that Theorem 10.4 holds for all ϕ with $\mu(\phi) = 1$; then, fix a generic path I'_s lying in the open set \mathcal{U}' of Theorem 3.15. Connecting $J'_s(T)$ and I_s by a one-parameter family of $j(T)$ -nearly-symmetric paths of almost-complex structures (as in the proof of Theorem 6.1), we see that the Floer homology calculated using the path $J'_s(T)$ is the same as that calculated using I'_s .

10.2. Approximate gluing. We describe how to splice spheres to flows for $\text{Sym}^g(\Sigma)$ to obtain flows in $\text{Sym}^{g+1}(\Sigma')$. First, we must introduce some notation.

Write $W = \text{Sym}^{g-1}(\Sigma) \times \{\sigma_1\}$ for the subvariety of $V = \text{Sym}^g(\Sigma)$. We consider flows $u \in \mathcal{M}(\mathbf{x}, \mathbf{y})$ which meet W transversally in $n = n_{\sigma_1}[u]$ distinct points $\{q_1, \dots, q_n\}$. For such a flow u , it is easy to see that there are constants $0 < r_1 < R_1$ with the property that the intersection of the image of u with $B_{r_1}(\sigma_1) \times \text{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1))$ is contained in the subset

$$(26) \quad B_{r_1}(\sigma_1) \times \text{Sym}^{g-1}(\Sigma - B_{R_1}(\sigma_1)) \subset \text{Sym}^g(\Sigma).$$

We introduce the corresponding product-like metric on this region. Moreover, we fix a conformal identification

$$(27) \quad B_1(\sigma_1) - \sigma_1 \cong [0, \infty) \times S^1$$

(observe that all the J_s agree with the complex structure in this region). We can find an n -tuple of disjoint balls $B_\varepsilon(q_i)$ which are mapped into the subspace from Equation (26) by u . By elementary complex analysis, we can find an n -tuple $(w_1, t_1 + i\theta_1, \dots, w_n, t_n + i\theta_n)$, where $w_i \in \text{Sym}^{g-1}(\Sigma)$, $t_i \in \mathbb{R}$, $\theta_i \in [0, 2\pi)$, with the property that, with respect to a conformal identification with the unit disk with the ball around q_i

$$(28) \quad D \cong B_\varepsilon(q_i)$$

and the identification of the range with the product from Equation (26), the restriction of u to $B_\varepsilon(q_i)$ can be written as

$$(29) \quad z \mapsto (w_i, e^{t_i + i\theta_i} z) + O(|z|^2).$$

It is convenient to recast this in terms of cylindrical coordinates with respect to fixed conformal identifications

$$B_\varepsilon(q_i) - q_i \cong [0, \infty) \times S^1.$$

We use weighted Sobolev spaces with weight function $e^{\delta\tau_1}$, where

$$\tau_1: \Sigma - \{q_1, \dots, q_n\} \longrightarrow [0, \infty)$$

is a smooth function satisfying:

- τ_1 vanishes in the complement of $B_\varepsilon(q_i)$,
- $\tau_1(s + i\varphi) \equiv s$ for $s \geq 1$ over each cylinder.

In terms of the cylindrical coordinates, then, the Taylor expansion in Equation (29) gives the following. For each i , there is $w_i \in \text{Sym}^{g-1}(\Sigma)$ and $t_i + i\theta_i$, for which the restriction of u to $B_\varepsilon(q_i) - q_i \cong [0, \infty) \times S^1$ differs by an $L^p_{1,\delta}$ map (for some $\delta > 0$) from the smooth map

$$a_{t_i+i\theta_i, w_i}: [0, \infty) \times S^1 \longrightarrow \text{Sym}^{g-1}(\Sigma) \times [0, \infty) \times S^1 \subset \text{Sym}^g(\Sigma)$$

defined by

$$a_{t_i+i\theta_i, w_i}(s + i\varphi) = (w_i, (s + t_i) + i(\varphi + \theta_i)),$$

where we have used the conformal identification

$$\begin{aligned} & \text{Sym}^{g-1}(\Sigma - B_{R_1}(\sigma_1)) \times ([0, \infty) \times S^1) \\ & \cong \text{Sym}^{g-1}(\Sigma - B_{R_1}(\sigma_1)) \times (B_{r_1}(\sigma_1) - \sigma_1) \subset \text{Sym}^g(\Sigma) \end{aligned}$$

(strictly speaking, when $t_i < 0$, we must cut off $s + t_i$ in the region where $s < -t_i$).

We can use these asymptotics to “cut off” the pseudo-holomorphic map u to construct a nearly pseudo-holomorphic map into $\text{Sym}^g(\Sigma - \sigma_1)$, which agrees with a standard map between cylinders, as follows. Given a real number $T > 0$, let $X_1(T)$ denote the subset of $\mathbb{D} - \{q_1, \dots, q_n\}$ consisting of points with cylindrical coordinate $\leq T$, i.e. $X_1(T) = \tau_1^{-1}((0, T])$, and $X_1(\infty) \cong \mathbb{D} - \{q_1, \dots, q_n\}$. Consider the map

$$\tilde{u}_T: X_1(\infty) \longrightarrow \text{Sym}^g(\Sigma)$$

defined to agree with u away from the balls $B_\varepsilon(q_i)$ and over $B_\varepsilon(q_i) \cong [0, \infty) \times S^1$, defined by

$$\tilde{u}_T(s + i\varphi) = h(s - T)a_{t_i+i\theta_i, w_i}(s + i\varphi) + (1 - h(s - T))u(s + i\varphi),$$

where $h: \mathbb{R} \longrightarrow [0, 1]$ is a smooth, increasing cut-off function which is identically 0 for $t < 0$ and 1 for $t > 1$. In the latter formula, the convex combination is to be interpreted using the exponential map on the range, with respect to the product metric.

When considering δ -weighted spaces for functions defined over the entire $\mathbb{D} - \{q_i\}$, we use the weight function $e^{\delta\tau_1}$ in addition to the weights on \mathbb{D} described in Section 3.2. There, we chose a smooth function $\tau_0: [0, 1] \times \mathbb{R} \cong \mathbb{D} \longrightarrow \mathbb{R}$ with $\tau_0(s + it) = |t|$ for all sufficiently large t (and here we can assume that the support of τ_0 misses the balls $B_\varepsilon(q_i)$), using the weight function $e^{\delta_0\tau_0}$, for some δ_0 depending on the local geometry around the intersections $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Now, suppose

$$f: X_1(\infty) \longrightarrow \text{Sym}^g(\Sigma - \sigma_1)$$

is a map which sends the cylindrical ends of $X_1(\infty)$ to the locus $\sigma_1 \times \text{Sym}^{g-1}(\Sigma)$. Then, we define the $L_{k;\delta_0,\delta}^p$ norm of a section $\xi \in \Gamma(X_1(\infty), f^*(T\text{Sym}^g(\Sigma)))$ by

$$\|\xi\|_{L_{k;\delta_0,\delta}^p(f)} = \sum_{i=0}^p \int_{\mathbb{D}-\{q_i\}} |\nabla^{(i)}\xi|^p e^{\delta_0\tau_0} e^{\delta\tau_1} d\mu,$$

where μ is the volume form for $\mathbb{D} - \{q_i\}$ with the cylindrical-end metric (obtained using the cylindrical identification of Equation (28)), and the norm of $\nabla^{(i)}$ denotes the metric induced from the pull-back of the cylindrical metric on $T\text{Sym}^g(\Sigma - \sigma_1)$ (obtained from the identification (26) together with (27)). The Banach space $L_{k;\delta_0,\delta}^p(f)$ is the space of sections

$$L_{k;\delta_0,\delta}^p(f) := \left\{ \xi \in L_{k;\delta_0,\delta}^p(\mathbb{D} - \{q_i\}, f^*(T\text{Sym}^g(\Sigma))) \mid \begin{array}{l} \xi(0, t) \in T_{f(0+it)}(\mathbb{T}_\beta), \forall t \in \mathbb{R}, \\ \xi(1, t) \in T_{f(1+it)}(\mathbb{T}_\alpha), \forall t \in \mathbb{R} \end{array} \right\}.$$

Similarly, we let $L_{\delta_0,\delta,k-1}^p(\Lambda^{0,1}f)$ denote the Banach space of $L_{k-1;\delta_0,\delta}^p$ sections of $f^*(T\text{Sym}^g(\Sigma))$ with a doubly-weighted norm, without the boundary conditions (note that, as in Section 3, the notation is justified by the fact that the bundle $\Lambda^{0,1}$ is trivial over \mathbb{D}). For our purposes, it suffices to use $k = 1$, and fix some $p > 2$ once and for all. In fact, in the interest of notational expediency, we typically drop the δ_0 from the notation, simply writing $L_{1,\delta}^p(f)$ and $L_\delta^p(\Lambda^{0,1}f)$, with some δ_0 being understood as fixed. Moreover, the results below will hold for all $\delta > 0$ sufficiently small (indeed, any $\delta \in (0, p)$ will do).

LEMMA 10.6. *Fix u as above. There are constants $t, k > 0$, $S_0 > 0$, and $C > 0$ so that for each $S > S_0$, the map \tilde{u}_S constructed above satisfies*

$$(30) \quad \|\bar{\partial}_{J_s} \tilde{u}_S\|_{L_\delta^p(\Lambda^{0,1}\tilde{u}_S)} \leq C e^{-kS}.$$

Moreover, for each S , there is T_0 so that for all $T > T_0$, \tilde{u}_S maps $X_1(T)$ into

$$\text{Sym}^g(\Sigma - [T + t, \infty) \times S^1).$$

Proof. Choose holomorphic coordinates (z_1, \dots, z_g) centered at

$$w_i \times \sigma_1 \in \text{Sym}^{g-1}(\Sigma - B_{R_1}(\sigma_1)) \times B_{r_1}(\sigma_1),$$

with respect to which $z_g = 0$ corresponds to the intersection of the coordinate patch with the subvariety $\text{Sym}^{g-1}(\Sigma) \times \{\sigma_1\}$. (These are $\text{Sym}^g(j)$ -holomorphic coordinates, which are actually holomorphic for all the J_s , in the specified region.) With respect to these coordinates, we can write u as (u_1, \dots, u_{g-1}, f) . Next, fix a local coordinate function z around the point q_i , so that $f(0) = 0$. We can factor $f(z) = czg(z)$, where $c = f'(0)$ and $g(0) = 1$. Thus, in cylindrical coordinates, we can write this as the map sending $s + i\varphi \in [0, \infty) \times S^1$ to

$$-\log f(e^{-(s+i\varphi)}) = -\log c + s + i\varphi - \log g(e^{-(s+i\varphi)}).$$

Now $z \mapsto \log g(z)$ vanishes at the origin, so that there is a constant c_1 with the property that $|\log g(e^{-(s+i\varphi)})| \leq c_1 e^{-s}$. Applying the above cut-off construction forces us to consider the new function

$$\tilde{f}(e^{-(s+i\varphi)}) = -\log c + s + i\varphi + (h(s - S) - 1) \log g(e^{-(s+i\varphi)}).$$

Now, $\bar{\partial}$ of this is given by

$$\bar{\partial} \tilde{f} = (\bar{\partial} h(s - S)) \log g(e^{-(s+i\varphi)}),$$

so we have a constant c_2 with the property that

$$|\bar{\partial} \tilde{f}(s + i\varphi)| \leq c_3 e^{-s}$$

pointwise; moreover, this tensor is supported in the strip where s is constrained to lie in the interval $[S, S + 1]$. Thus there is a c_3 with

$$(31) \quad \int_{[0, \infty) \times S^1} |\bar{\partial} \tilde{f}|^p e^{\delta s} \leq c_3 e^{(-p+\delta)S}.$$

To handle the other components u_i , we proceed in a similar manner. Noting that $u_i(0) = 0$, the cutting off construction gives

$$\tilde{u}_i(e^{-(s+i\varphi)}) = (1 - h(s - S)) u_i(e^{-(s+i\varphi)}).$$

Applying $\bar{\partial}$ to this, we obtain

$$\bar{\partial} \tilde{u}_i(e^{-(s+i\varphi)}) = -(\bar{\partial} h(s - S)) u_i(e^{-(s+i\varphi)}).$$

We have a constant c_4 with $|u_i(z)| \leq c_4 z$, so that

$$|\bar{\partial} \tilde{u}_i(s + i\varphi)| \leq c_4 e^{-s}$$

pointwise. Once again, since the tensor is supported in a bounded strip, we have that

$$(32) \quad \int_{[0, \infty) \times S^1} |\bar{\partial} \tilde{u}_i|^p e^{\delta s} \leq c_3 e^{(-p+\delta)S}.$$

Together, Inequalities (31) and (32) give Inequality (30), with respect to cylindrical-end metrics on the range $\text{Sym}^g(\Sigma - \sigma_1)$ (which, over the subset $\text{Sym}^g(\Sigma - B_{r_1})$, can be made to be isometric to the corresponding subset of $\text{Sym}^g(\Sigma)$).

The second part of the lemma is straightforward. \square

Next, we turn our attention to spheres in $\text{Sym}^2(E)$. More precisely, consider holomorphic maps

$$v: \mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E)$$

which are constant on the first factor, and which represent the positive generator of $\pi_2(\text{Sym}^2(E))$ on the second (i.e. they satisfy $n_{\sigma_2}[v] = 1$). We denote the space of such maps, modulo holomorphic reparametrization, by

$$\mathcal{M}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2)).$$

LEMMA 10.7. *The map*

$$\mathcal{M}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2)) \longrightarrow \text{Sym}^{g-1}(\Sigma) \times E$$

which takes $[v] \in \mathcal{M}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2))$ to the unique pair (w, y) , with the property that $(w, \{y, \sigma_2\})$ is in the image of v , induces a one-to-one correspondence.

Proof. This is true because $\pi: \text{Sym}^2(E) \longrightarrow E$ is a ruled surface (see Lemma 10.3), the generator corresponds to holomorphic spheres representing the fiber class (according to the proof of that lemma), and the map $E \longrightarrow \text{Sym}^2(E)$ given by $\sigma \mapsto \{\sigma, \sigma_2\}$ is a section. Thus, the composite $\pi \circ v$ is a holomorphic map from the sphere to the torus, which must be constant. Hence, v maps to a fiber. Indeed, from our assumption on the homotopy class, it follows that v must map isomorphically to a fiber, and each fiber is determined by its intersection with the section. \square

Let v be as above. Thinking of \mathbb{S} as the Riemann sphere $\mathbb{S} \cong \mathbb{C} \cup \{\infty\}$, we normalize v so that $v(0) = w \times \{y, \sigma_2\}$ (this can be achieved by precomposing v with a Möbius transformation if necessary), and suppose moreover that $y \neq \sigma_2$. A neighborhood of the image $w \times \{y, \sigma_2\}$ can be identified with a product

$$\text{Sym}^{g-1}(\Sigma) \times B_{r_2}(\sigma_2) \times (E - B_{R_2}(\sigma_2)) \subset \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E),$$

for some $r_2 < R_2$. With respect to these coordinates, near 0, v then takes the form

$$z \mapsto (w, \sigma_2 + O(|z|), e^{t+i\theta}z) + O(|z|^2)$$

for some $t + i\theta$ (when the third coordinate = 0 corresponds to the point $y \in E$). In terms of cylindrical coordinates in $[0, \infty) \times S^1 \cong B_\varepsilon(0) - 0 \subset \mathbb{S}$ and $[0, \infty) \times S^1 \cong B_{r_2}(\sigma_2)$, this can be phrased as follows:

$$v: [0, \infty) \times S^1 \longrightarrow \text{Sym}^{g-1}(\Sigma) \times ([0, \infty) \times S^1) \times E$$

satisfies the property that there is a $w \in W$ and $t + i\theta$ so that v differs from the map

$$b_{(t+i\theta, w, y)}(s + i\varphi) = (w, (s + t) + i(\theta + \varphi), y)$$

by a map which is in $L_{1,\delta}^p$ for any $\delta > 0$. (Here, $L_{1,\delta}^p$ is defined with respect to a weight function $e^{\delta\tau_2}$, for $\tau_2: \mathbb{S} - 0 \longrightarrow \mathbb{R}^+$ defined in a manner analogous to τ_1 .)

In view of these asymptotics, we can form a “cutting off” construction as before. Given a real number $S > 0$, let $\mathbb{S}(S)$ denote the subspace of $x \in \mathbb{S}$ where $\tau_2(x) \leq S$. Next, consider the map

$$v_S: \mathbb{S} - \infty \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E - \sigma_2)$$

defined to agree with v over $\mathbb{S}(S)$, and given by the formula

$$v_S(s + i\varphi) = h(s - S)b_{t+i\theta} + (1 - h(s - S))v(s + i\varphi)$$

over the cylinder $[0, \infty) \times S^1 \cong B_{r_2}(0) - 0$.

The six-dimensional group $\mathbb{P}\text{SL}_2(\mathbb{C})$ acts on the space of parametrized holomorphic maps from \mathbb{S} . The normalization condition on v above cuts out two dimensions from this automorphism group. An additional two dimensions could be cut out by the following conditions. The holomorphic map v has an “energy measure” obtained by pulling back the symplectic form from $\text{Sym}^2(E)$. This measure in turn has a center of mass in the three-ball. Requiring the center of mass to lie on the z -axis (the axis connecting the points corresponding to 0 and ∞ in \mathbb{S}) cuts out an additional two-dimensions from the automorphism group, but we find it convenient to formulate this condition as follows. Let v_0 be a normalized holomorphic sphere with $v_0(0) = w \times \{c, \sigma_2\}$, whose center of mass lies on the z axis, and choose a point $\sigma'_2 \in E$ with the property that $v_0(\infty)$ lies in the submanifold $W' = \text{Sym}^{g-1}(\Sigma) \times \sigma'_2 \times E$ (there are two possible choices for σ'_2).

Definition 10.8. A holomorphic map

$$v: \mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E)$$

is called *centered* if the following two conditions are satisfied:

- $v(0) = w \times \{y, \sigma_2\}$ for some $w \in \text{Sym}^{g-1}(\Sigma)$ and $y \in E - \sigma_2$.
- $v(\infty) \in W'$.

We can collect the set of centered holomorphic maps into a moduli space, which we denote by

$$\mathcal{M}^{\text{cent}}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2)).$$

Since the space of holomorphic automorphisms which respect the centered condition is identified with \mathbb{C}^* , this moduli space has a \mathbb{C}^* action, and fibers over the region in $\mathcal{M}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2))$ corresponding to

$\text{Sym}^{g-1}(\Sigma) \times (E - B_{R_2}(\sigma_2))$. Now, we can view the assignment $v \mapsto (w, t + i\theta, y)$ as a map from the moduli space of centered maps

$$\rho_2: \mathcal{M}^{\text{cent}}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2)) \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \mathbb{R} \times S^1 \times E.$$

It is an easy consequence of Lemma 10.7 that ρ_2 is a \mathbb{C}^* -equivariant diffeomorphism. Given $(w, t, i\theta) \in \text{Sym}^{g-1}(\Sigma) \times \mathbb{R} \times S^1$, let

$$v_{(w,t,i\theta)}: \mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E)$$

denote the centered map with the property $\rho_2 v_{(w,t,i\theta)} = (w, t + i\theta, c)$.

Given the J_s -holomorphic disk u , we would like to splice in a number of centered spheres, to construct a nearly $J'_s(T)$ -holomorphic curve in $\Sigma'(T)$. For instance, at the marked point q_i , with value w_i we splice in a copy of $v_{(w_i, -t_i, \theta_i)}$.

More precisely, using the given conformal identifications of Equation (27) of the neighborhoods of the puncture points, the connected sum $\Sigma'(T)$ can be thought of as the space obtained from $\Sigma(2T)$ and $E(2T)$, by identifying the cylinders

$$[0, 2T] \times S^1 \subset \Sigma(2T) \text{ and } [0, 2T] \times S^1 \subset E(2T)$$

using the involution $(t, i\theta) \sim (2T - t, i\theta)$. Let

$$X_2(T) = \bigcup_{i=1}^n \mathbb{S}(T)_i,$$

and $X_1 \cup_T X_2$ denote the union of $X_1(T)$ with $X_2(T)$, glued along their common boundary. Now, given a holomorphic disk u and a pair S and T of real numbers with $0 < S < T - t$, let

$$\tilde{v}_S: X_2(\infty) \longrightarrow \text{Sym}^{g-1}(\Sigma) \times \text{Sym}^2(E - \sigma_2)$$

be the map whose restriction to $\mathbb{S}(T)_i \subset X_2(\infty)$ is the cut-off of $v_{(w_i, t_i + i\theta_i)}$. We can form a spliced map

$$\hat{\gamma} = \hat{\gamma}(u, S, T): X_1 \cup_T X_2 \longrightarrow \text{Sym}^{g+1}(\Sigma \#_T E),$$

which is defined to agree with $\tilde{u}_S \times \{c\}$ over $X_1(T)$, and to agree with the map \tilde{v}_S over $X_2(T)$. Note that $X_1 \cup_T X_2$ is conformally equivalent to \mathbb{D} , though its geometry is different. Reflecting these geometries, we find it convenient to introduce weighted Sobolev spaces over $X_1 \cup_T X_2$. Fix some $p > 2$, and let $\mathcal{Z}(T)$ denote the space of sections $\Gamma(X_1 \cup_T X_2, \Lambda^{0,1} \otimes \hat{\gamma}^*(T\text{Sym}^{g+1}(\Sigma')))$, endowed with the norm

$$\|\phi\|_{\mathcal{Z}_\delta(T)} = \left(\int_{X_1 \cup_T X_2} |\phi|^p e^{\delta_0 \tau_0} e^{\delta \tau} \right)^{1/p},$$

where τ_0 is the weight function used near the two ends of the strip

$$[0, 1] \times \infty \cong X_1 \cup_T X_2,$$

and $\tau: X_1 \cup_T X_2 \longrightarrow \mathbb{R}$ is a smooth function (depending, of course, on T) which agrees with τ_1 and τ_2 over $X_1(T-1)$ and $X_2(T-1)$, and satisfies a uniform C^∞ (i.e. independent of T) bound on $d\tau$.

LEMMA 10.9. *Given any sufficiently large S , there is a T_0 so that for all $T > T_0$, the map $\hat{\gamma}(u, S, T)$ determines a smooth map*

$$\hat{\gamma}(u, S, T): X_1 \cup_T X_2 \longrightarrow \text{Sym}^{g+1}(\Sigma \#_T E).$$

Moreover, there are positive constants k and C , so that for all $S > S_0$ and $T > T_0$, $\hat{\gamma}(u, S, T)$ satisfies

$$\|\bar{\partial}_{J'_s} \hat{\gamma}(u, S, T)\|_{\mathcal{Z}_s(T)} \leq C e^{-kS}.$$

Proof. This follows immediately from Lemma 10.6, and the corresponding estimate for \tilde{v}_S (which follows in the same manner). Note that the image of $\hat{\gamma}$ is contained in the region where the divisor on the E side is bounded away from σ_2 , so that the almost-complex structures J'_s split as $J_s \times j_E$. \square

10.3. *Parametrices.* To perform the gluing, we must construct an inverse for the linearization of the $\bar{\partial}$ -operator for the spliced map from $X_1 \cup_T X_2$, whose norm is bounded independent of T (see Proposition 10.12 below). To set this up, we describe suitable function spaces for the operators over the cylindrical-end versions:

$$D_{\tilde{u}}: \Gamma(X_1(\infty), E_1) \longrightarrow \Gamma(X_1(\infty), F_1)$$

and

$$D_{\tilde{v}}: \Gamma(X_2(\infty), E_2) \longrightarrow \Gamma(X_2(\infty), F_2),$$

where:

$$\begin{aligned} E_1 &= \tilde{u}^*(TSym^g(\Sigma - \sigma_1)) \\ F_1 &= \text{Hom}(TX_1(\infty), \tilde{u}^*(TSym^g(\Sigma - \sigma_1))), \\ E_2 &= \tilde{v}^*(T(\text{Sym}^{g-1}(\Sigma - \sigma_1) \times \text{Sym}^2(E - \sigma_2))) \\ F_2 &= \text{Hom}(TX_2(\infty), \tilde{v}^*(T(\text{Sym}^{g-1}(\Sigma - \sigma_1) \times \text{Sym}^2(E - \sigma_2)))). \end{aligned}$$

Note that, given the complex structures on X_1 and X_2 , there are canonical identifications

$$F_1 \cong \Lambda^{(0,1)} \otimes E_1 \text{ and } F_2 \cong \Lambda^{(0,1)} \otimes E_2,$$

but we have chosen to write F_1 and F_2 as above, as we will allow the complex structure on $X_1(\infty)$ to vary.

Indeed, to allow for this variation, we enlarge the domain of $D_{\tilde{u}}$, by $\mathbb{C}^n \cong \bigoplus_{i=1}^n T_{q_i} \mathbb{D}$. Let $\phi_\xi \in \text{Diff}(\mathbb{D})$ be a family of diffeomorphisms of \mathbb{D} indexed by $\xi \in \mathbb{C}^n$, which satisfies the following properties:

- $\phi_0 = \text{Id}$,
- $\phi_\xi|_{\mathbb{D} - \bigcup B_\varepsilon(q_i)} \equiv \text{Id}$,
- If ξ is sufficiently small, then $\phi_\xi|_{B_{\varepsilon/2}(q_i)} \longrightarrow \mathbb{D}$ is translation by ξ_i .

A family of diffeomorphisms ϕ_ξ as above is constructed, for example, by fixing a cut-off function h supported in $\bigcup B_\varepsilon(q_i)$ which is identically one on $B_{3\varepsilon/4}(q_i)$. Then, if $\xi \in \bigoplus T_{q_i}(\mathbb{D})$, we let $\tilde{\xi}$ be its locally constant extension to $\bigcup_{i=1}^n B_\varepsilon(q_i)$. Then, ϕ_ξ is obtained by exponentiating the compactly supported vector field $h\tilde{\xi}$. We then define $j_\xi = \phi_\xi^* j_0$, where j_0 is the standard complex structure on \mathbb{D} .

We then consider a modified Cauchy-Riemann operator

$$\bar{\partial}_{J_s}^{\text{ext}} : \text{Map}(X_1(\infty), \text{Sym}^g(\Sigma)) \times \mathbb{C}^n \longrightarrow \Gamma(X_1(\infty), F_1),$$

which assigns to $(u, \xi) \in \text{Map}(X_1(\infty), \text{Sym}^g(\Sigma)) \times \mathbb{C}^n$ the bundle map from $T\mathbb{D}$ to F_1 given by

$$\bar{\partial}_{J_s}^{\text{ext}}(u, \xi) = (u^* J_{\phi_\xi^{-1}}) \circ du - du \circ j_\xi,$$

where $J_{\phi_\xi^{-1}}$ is the family of almost-complex structures over \mathbb{D} given by

$$J_{\phi_\xi^{-1}}(s, t) = J(\phi_\xi^{-1}(s, t))$$

(where we view our original path J as a map on \mathbb{D} which is constant in the t directions). Clearly, $\bar{\partial}_{J_s}^{\text{ext}}(u, \xi) = 0$ if and only if the map $u \circ \phi_\xi^{-1}$ is J_s -holomorphic. The linearization of $\bar{\partial}_{J_s}^{\text{ext}}$ is the map

$$D_{(u, \xi)} : \Gamma(X_1(\infty), E_1) \oplus \mathbb{C}^n \longrightarrow \Gamma(X_1(\infty), F_1)$$

given by

$$(33) \quad D_{(u, \xi)}(\nu, x) = (u^* \nabla_\nu J_{\phi_\xi^{-1}}) \circ du + (u^* J_{\phi_\xi^{-1}}) \circ (\nabla \nu) \\ - (\nabla \nu) \circ j_\xi - (du) \circ (\mathcal{L}_x j_\xi) + (d_x(J_{\phi_\xi^{-1}})) \circ du.$$

The first term here involves a covariant derivative of the tensor $J_{\phi_\xi^{-1}(s, t)}$ (defined over $T\text{Sym}^g(\Sigma)$); the fourth term involves the Lie derivative of the complex structure j_ξ in the direction specified by the vector field corresponding to $x \in \mathbb{C}^n$, the final term involves a directional derivative of the map $J_{\phi_\xi^{-1}}$, thought of as a map from \mathbb{D} to the space of almost-complex structures, in the direction specified by x , thought of as a vector field along \mathbb{D} .

With these remarks in place, we turn attention to the Sobolev topologies to be used on the domain and range. For the operator $D_{\tilde{u}}$, we consider the Banach spaces

$$\mathcal{W}_1 = \left(L_{1; \delta_0, \delta}^p(\tilde{u}) + \mathcal{H} \right) \oplus \mathbb{C}^n,$$

and

$$\mathcal{Z}_1 = L_{\delta_0, \delta}^p(\Lambda^1 \otimes \tilde{u}),$$

where \mathcal{H} is the finite-dimensional vector space

$$\mathcal{H} = \bigoplus_{i=1}^n (T_{w_i} \text{Sym}^{g-1}(\Sigma) \oplus \mathbb{C}).$$

Recall that, by definition, $L_{1; \delta_0, \delta}^p(\tilde{u})$ and $L_{\delta_0, \delta}^p(\Lambda^1 \otimes \tilde{u}) \cong L_{\delta_0, \delta}^p(\Lambda^{(0,1)} \otimes \tilde{u})$ are spaces of sections of E_1 and F_1 respectively (satisfying the appropriate boundary and decay conditions, specified earlier).

To make sense of \mathcal{W}_1 , we view \mathcal{H} as a space of sections of E_1 , as follows. First, consider the tangent vectors $T_{w_i} \text{Sym}^{g-1}(\Sigma)$ as vector fields in $\tilde{u}^* T \text{Sym}^{g-1}(\Sigma)$ which are locally constant over the cylindrical ends. More precisely, a tangent vector $T_{w_i}(\text{Sym}^{g-1}(\Sigma))$ gives rise to a section of E_1 over the i^{th} cylindrical end. These vector fields are then to be multiplied by a fixed cut-off function (supported in the i^{th} end), to give rise to global sections of E_1 . The \mathbb{C} summand in \mathcal{H} corresponds to tangent vector fields to the cylinders which are constant, as follows. We identify the constant vector fields over $S^1 \times \mathbb{R}$ with \mathbb{C} . Again, using a smooth cut-off function, we can transfer the vector fields over $S^1 \times \mathbb{R}$ to $X_1(\infty)$ (to get a vector space of tangent vector fields, which are constant at the cylindrical ends). We can then use the derivative of \tilde{u} to push these forward to get sections of E_1 (which are holomorphic over the ends).

The vector space \mathbb{C}^n corresponds to variations in the complex structure over $X_1(\infty)$, corresponding to variations of the puncture points on the disks, as described above. Then, the differential operator $D_{(\tilde{u}, \xi)}$ extended as above induces a continuous linear map

$$D_{\tilde{u}}: \mathcal{W}_1 \longrightarrow \mathcal{Z}_1.$$

(Perhaps a more suggestive notation for \mathcal{W}_1 is to write it as:

$$L_{1; \delta_0, \delta}^p(\tilde{u}) + \bigoplus_{i=1}^n (T_{w_i} \text{Sym}^{g-1}(\Sigma) \oplus T_{q_i} \mathbb{D} \oplus N_{w_i} \text{Sym}^{g-1}(\Sigma)),$$

where $N_w \text{Sym}^{g-1}(\Sigma)$ is the normal vector space to $\text{Sym}^{g-1}(\Sigma) \times \{\sigma_1\}$ at the point $w \times \sigma_1$.) There is also a natural linear projection

$$\rho_1: \mathcal{W}_1 \longrightarrow \mathcal{H}.$$

Note that these operators (and indeed the Banach spaces on which they are defined) depend on the real number $S > 0$, which we have suppressed from the notation.

LEMMA 10.10. *For all sufficiently large $S > 0$, the operator $D_{\tilde{u}}$ has a one-dimensional kernel and no cokernel. Moreover, it admits a right inverse, with an inverse whose norm is bounded above independent of S .*

Proof. Consider the operator

$$D_{(u,\xi)}^{\text{cyl}}: \left(L_{1,\delta_0,\delta}^p(u) + \mathcal{H} \right) \oplus \mathbb{C}^n \longrightarrow L_{\delta_0,\delta}^p(\Lambda^1 \otimes u).$$

Here, u is viewed as a map between $X_1(\infty)$ and $\text{Sym}^g(\Sigma - \sigma_1)$, endowed with cylindrical metrics. By the removable singularities theorem, the kernel and cokernel are identified with kernel and cokernel of the flow-line in $\text{Sym}^g(\Sigma)$,

$$D_u: L_{1,\delta_0}^p(u) \longrightarrow L_{\delta_0}^p(\Lambda^1 \otimes u).$$

(See the corresponding discussion in [3].) By assumption, the cokernel vanishes and its kernel is one-dimensional, since we assumed that u was a generic flow-line. Moreover, the operators $D_{(\tilde{u}_S,\xi)}$ converge to $D_{(u,\xi)}^{\text{cyl}}$ as $S \mapsto \infty$. The lemma then follows. \square

An analogous discussion applies on the other side, as well. However, to compensate for the \mathbb{C}^n increase of the domain on the X_1 side, we must decrease the domain on the X_2 side. We define Banach spaces

$$\mathcal{W}_2^{\text{ext}} = L_{1,\delta}^p(X_2(\infty), E_2) + \mathcal{H}$$

and

$$\mathcal{Z}_2 = L_{\delta}^p(X_2(\infty), F_2)$$

(where \mathcal{H} is the finite-dimensional vector space defined earlier), and once again we have a map:

$$D_{\tilde{v}}^{\text{ext}}: \mathcal{W}_2^{\text{ext}} \longrightarrow \mathcal{Z}_2$$

(inherited by the differential operator $D_{\tilde{v}}$), and the linear projection

$$\rho_2: \mathcal{W}_2^{\text{ext}} \longrightarrow \mathcal{H}.$$

As it turns out, $D_{\tilde{v}}$ on $\mathcal{W}_2^{\text{ext}}$ has a nontrivial kernel. To get rid of the kernel, we consider a linearization of the centered condition described in the previous subsection. Specifically, if $x_1, \dots, x_n \in X_2(\infty)$ are the points corresponding to the origins under the identification $\bigcup_{i=1}^n \mathbb{C} \cong X_2(\infty)$, we assume that $\tilde{v}(x_i) \in W'$. Then, we can define

$$\mathcal{W}_2 = \{ \xi \in \mathcal{W}_2^{\text{ext}} \mid \xi(x_i) \in T_{v(x_i)} W' \quad \forall i = 1, \dots, n \}.$$

LEMMA 10.11. *For all sufficiently large $S > 0$, the map $D_{\tilde{v}}$ has no cokernel, and its kernel is identified under ρ_2 with \mathcal{H} .*

Proof. This follows as in the proof of Lemma 10.10. Note that (by a removable singularities theorem) the kernel and cokernel of

$$D_v^{\text{ext}}: \mathcal{W}_2^{\text{ext}} \longrightarrow \mathcal{Z}_2$$

can be identified with the corresponding spaces for the operator

$$D_v: L_1^p(v) \longrightarrow L^p(\Lambda^1 \otimes v),$$

where

$$L_1^p(v) = L_1^p\left(\bigcup_{i=1}^n \mathbb{S}, v^*(TSym^g(\Sigma))\right)$$

and

$$L_1^p(\Lambda^1 \otimes v) = L_1^p\left(\bigcup_{i=1}^n \mathbb{S}, \Lambda^1 \otimes v^*(TSym^g(\Sigma))\right).$$

In turn, these spaces are given by the deformation theory of (parametrized) holomorphic spheres, up to reparametrizations. The fact that the cokernel vanishes follows from the fact that the fiber class of a ruled surface forms a smooth moduli space, which in turn follows from an easy Leray-Serre spectral sequence argument, which we defer to the next paragraph. For each sphere, then, there is a complex three-dimensional family of holomorphic vector fields which fix the image curve. These generate a three-dimensional space of kernel elements (for each sphere). However, if we restrict the domain to \mathcal{W}_2 , we are considering spheres which are specified to lie in the subvariety $W \times W'$ at a fixed pair of points: hence, there is only a one-dimension space \mathbb{C}^* of automorphisms left. This kernel is easily seen to be captured by ρ_2 .

To see that the n -fold fiber class of $Sym^2(E)$ over E considered above has a smooth deformation theory, note that each n -tuple of fibers is given as the zero set of a section σ of a bundle \mathcal{L} , which is the pull-back of a Chern class n line bundle \mathcal{L}_0 over the base E . The cokernel of the deformation complex is identified with the one-dimensional cohomology of the quotient sheaf of \mathcal{L} by σ ,

$$H^1(Sym^2(E), \mathcal{L}/\sigma\mathcal{O}).$$

First, we prove the vanishing

$$(34) \quad H^1(Sym^2(E), \mathcal{L}) = 0.$$

The Leray-Serre spectral sequence for the fibration has an $E_2^{p,q}$ term with

$$H^p(E, R^q\pi_*\mathcal{L}) \Rightarrow H^{p+q}(Sym^2(E), \pi^*\mathcal{L}),$$

where $R^q\pi_*$ is the derived functor of the push-forward map π_* . It suffices to prove that both $H^1(E, \pi_*\mathcal{L}) = 0$ and $H^0(E, R^1\pi_*\mathcal{L}) = 0$. The first vanishing follows from Serre duality, since $\pi_*\mathcal{L} = \mathcal{L}_0$, a positive line bundle over an elliptic curve. The second vanishing statement follows from the projection formula $R^1\pi_*(\pi^*\mathcal{L}_0) \cong \mathcal{L}_0 \otimes R^1\pi_*\mathcal{O} = 0$ (see for example Chapter III.8 of [16]). This proves the vanishing in Equation (34). The vanishing of $H^1(Sym^2(E), \mathcal{L}/\sigma\mathcal{O})$, then, follows from Equation (34), the long exact sequence in cohomology associated to the short exact sequence of sheaves:

$$0 \longrightarrow \mathcal{O} \xrightarrow{\sigma} \mathcal{L} \longrightarrow \mathcal{L}/\sigma\mathcal{O} \longrightarrow 0,$$

and the fact that $H^2(Sym^2(E), \mathcal{O}) = 0$. □

We can consider norms on spaces of sections

$$\mathcal{X}(T)^{\text{ext}} = L_1^p(X_1 \cup_T X_2, \tilde{u} \#_T \tilde{v}^*(T(\text{Sym}^{g+1}(\Sigma')))) \oplus \mathbb{C}^n,$$

which reflect the neck stretching. These norms are obtained by using a partition of unity to transfer sections to $X_1(\infty)$ and $X_2(\infty)$ and use the norms of \mathcal{W}_1 and \mathcal{W}_2 on those spaces. More precisely, choose a partition of unity $\{\phi_1, \phi_2\}$ on $X_1 \cup_T X_2$ subordinate to its cover $\{X_1(T+1), X_2(T+1)\}$, constructed so that $d\phi_1$ is uniformly C^∞ bounded, independent of T . There is an injection

$$\iota: \mathcal{X}(T)^{\text{ext}} \longrightarrow \mathcal{W}_1 \oplus \mathcal{W}_2^{\text{ext}}$$

given by

$$\iota(a, \xi) = ((\phi_1 a + (1 - \phi_1)h, \xi), (\phi_2 a + (1 - \phi_2)h)),$$

where $h = \Pi_{\mathcal{H}}(a|_{S^1 \times \{0\}})$ over each cylinder, and where $\phi_1 a + (1 - \phi_1)h$ and $\phi_2 a + (1 - \phi_2)h$ are to be thought of as sections over $X_1(\infty)$ and $X_2(\infty)$ respectively in the obvious manner. The norm on $\mathcal{X}(T)^{\text{ext}}$ is defined to make ι an isometry onto its image; i.e.,

$$\|(a, \xi)\|_{\mathcal{X}(T)} = \|\phi_1 a + (1 - \phi_1)h\|_{L_{1,\delta}^p(X_1) + \mathcal{H}} + \|\phi_2 a + (1 - \phi_2)h\|_{L_{1,\delta}^p(X_2) + \mathcal{H}} + \|\xi\|_{\mathbb{C}^n}.$$

Let $\mathcal{X}(T) \subset \mathcal{X}(T)^{\text{ext}}$ denote the subset whose image under ι maps to $\mathcal{W}_1 \oplus \mathcal{W}_2 \subset \mathcal{W}_1 \oplus \mathcal{W}_2^{\text{ext}}$; i.e. where the vectors at the centers of X_2 are tangent to W' , $a(x_i) \in T_{v(x_i)}W'$ for $i = 1, \dots, n$. Similarly, we consider the Banach space

$$\mathcal{Z}(T) = L^p(X_1 \cup_T X_2, \Lambda^{0,1} \otimes \tilde{u} \#_T \tilde{v}^*(T\text{Sym}^n(\Sigma)))$$

with norm given by

$$\|b\|_{\mathcal{Z}(T)} = \|\phi_1 b\|_{\mathcal{Z}_1} + \|\phi_2 b\|_{\mathcal{Z}_2}.$$

It is easy to see that for any fixed T , the induced topology on $\mathcal{Z}(T)$ is the same as the ordinary unweighted Sobolev topologies on the corresponding bundle over $X_1 \#_T X_2$. These norms were chosen instead because they satisfy the following property:

PROPOSITION 10.12. *There is a constant C with the property that for all $T > T_0$, $D_T = D_{\tilde{u} \#_T \tilde{v}}$ has a continuous right inverse*

$$P_T: \mathcal{Z}(T) \longrightarrow \mathcal{X}(T),$$

whose operator norm is uniformly bounded above (independent of T).

The inverse is constructed by patching together right inverses over the cylindrical end versions $X_1(\infty)$ and $X_2(\infty)$. We describe these inverses, and prove Proposition 10.12 at the end of the subsection.

Let \mathcal{X}_∞ denote the fibered product of ρ_1 and ρ_2 ; i.e. it is the Banach space which fits into the short exact sequence

$$0 \longrightarrow \mathcal{X}_\infty \longrightarrow \mathcal{W}_1 \oplus \mathcal{W}_2 \xrightarrow{\rho_1 - \rho_2} \mathcal{H} \longrightarrow 0.$$

With $\mathcal{Z}_\infty = \mathcal{Z}_1 \oplus \mathcal{Z}_2$, the maps $D_{\tilde{u}}$ and $D_{\tilde{v}}$ induce a map

$$D_\infty: \mathcal{X}_\infty \longrightarrow \mathcal{Z}_\infty,$$

by

$$D_\infty(a, b) = (D_{\tilde{u}}a, D_{\tilde{v}}b).$$

Proof of Proposition 10.12. The restriction of ρ_2 to the kernel of $D_{\tilde{v}}$ in \mathcal{W}_2 is an isomorphism onto \mathcal{H} (according to Lemma 10.11). It follows from this, and Lemma 10.10 that $D_\infty: \mathcal{X}_\infty \longrightarrow \mathcal{Z}_\infty$ is invertible. Let $P_\infty = (P_\infty^{(1)}, P_\infty^{(2)})$ denote the inverse of D_∞ . A paramatrix Q_T for D_T can be defined by

$$Q_T(y) = \psi_{1;T} P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty}) + \psi_{2;T} P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty}),$$

where $\psi_{1;T}, \psi_{2;T}$ are functions defined on $X_1 \cup_T X_2$, which satisfy the property that $\psi_{1;T}^2 + \psi_{2;T}^2 \equiv 1$, and

$$\sup_{X_1 \cup_T X_2} |d\psi_{1;T}| \leq \frac{c}{T},$$

for some fixed constant c . Such a family can be obtained from some fixed $\{\psi_{1;1}, \psi_{2;1}\}$ supported in the tube by rescaling the tube by T .

The operator norms of the $Q_T: \mathcal{Z}(T) \longrightarrow \mathcal{X}(T)$ are uniformly bounded independent of T . Indeed,

$$\begin{aligned} \|Q_T y\|_{\mathcal{X}(T)} &\leq \|P_\infty(\psi_1 y, \psi_2 y)\|_{\mathcal{X}_\infty} \\ &\leq \|P_\infty\| \|(\psi_1 y, \psi_2 y)\|_{\mathcal{Z}_\infty} \\ &\leq \|P_\infty\| \|y\|_{\mathcal{Z}(T)}, \end{aligned}$$

where $\|P_\infty\|$ denotes the operator norm of

$$P_\infty: \mathcal{Z}_\infty \longrightarrow \mathcal{X}_\infty.$$

Moreover, Q_T is nearly an inverse for D_T , as

$$\begin{aligned} D_T \circ Q_T(y) &= D_T \psi_{1;T} P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty}) + \psi_{2;T} P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty}) \\ &= (d\psi_{1;T}) P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty}) + (d\psi_{2;T}) P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty}) \\ &\quad + \psi_{1;T} D_{\tilde{u}} P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty}) + \psi_{2;T} D_{\tilde{v}} P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty}) \\ &= (d\psi_{1;T}) P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty}) + (d\psi_{2;T}) P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty}) \\ &\quad + \psi_{1;T} \psi_{1;Ty} + \psi_{2;T} \psi_{2;Ty}; \end{aligned}$$

i.e., the operator $D_T \circ Q_T$ differs from the identity map (on $\mathcal{Z}(T)$) by the operator

$$y \mapsto (d\psi_{1;T}) P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty}) + (d\psi_{2;T}) P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty}).$$

But

$$\begin{aligned}
\|(d\psi_{1;T}P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty})\|_{\mathcal{X}(T)} &\leq \frac{c}{T}\|P_\infty^{(1)}(\psi_{1;Ty}, \psi_{2;Ty})\|_{\mathcal{W}_1} \\
&\leq \frac{c}{T}\|P_\infty(\psi_{1;Ty}, \psi_{2;Ty})\|_{\mathcal{X}_\infty} \\
&\leq \frac{c}{T}\|P_\infty\|\|(\psi_{1;Ty}, \psi_{2;Ty})\|_{\mathcal{Z}_\infty} \\
&\leq \frac{c}{T}\|P_\infty^{(1)}\|\|y\|_{\mathcal{Z}(T)},
\end{aligned}$$

where $\|P_\infty\|$ denotes the operator norm of

$$P_\infty: \mathcal{Z}_\infty \longrightarrow \mathcal{X}_\infty.$$

A similar estimate holds for $\|(d\psi_{2;T}P_\infty^{(2)}(\psi_{1;Ty}, \psi_{2;Ty})\|_{\mathcal{X}(T)}$. Thus, if T is sufficiently large relative to $\|P_\infty\|$, then we can arrange for, say,

$$\|D_T \circ Q_T - \text{Id}\| \leq \frac{1}{2},$$

so that the von Neumann expansion gives the inverse of $D_T \circ Q_T$, with $\|D_T \circ Q_T\| \leq 2$. Letting $P_T = Q_T \circ (D_T \circ Q_T)^{-1}$, it follows that P_T is a right inverse for D_T , with norm bounded by $2\|P_\infty\|$ (again, for sufficiently large T). \square

10.4. *Proof of the gluing result.* With the approximately holomorphic disks $\widehat{\gamma} = \widehat{\gamma}(u)$ constructed in Subsection 10.2 and the deformation theory from Subsection 10.3, the construction of the holomorphic disks in $\text{Sym}^{g+1}(\Sigma')$, as stated in Theorem 10.4, follows an application of the inverse function theorem.

We give a suitable modification of Floer's set-up (see [9]). There is an “exponential mapping”

$$\text{Exp}_{\widehat{\gamma}}: \mathcal{X}(T) \longrightarrow \text{Map}(X_1 \cup_T X_2, \text{Sym}^{g+1}(\Sigma')) \times \text{Diff}(\mathbb{D}),$$

which is obtained as follows. Let $(a, \xi) \in \mathcal{X}(T)$; i.e. a is a section of $\widehat{\gamma}_T^*(T\text{Sym}^{g+1}(\Sigma'))$, and $\xi \in \mathbb{C}^n \subset \Gamma(X_1 \cup_T X_2, T(X_1 \cup_T X_2))$. Fix a path \mathbf{g}_s of metrics for which \mathbb{T}'_β is \mathbf{g}_0 -totally geodesic, and \mathbb{T}'_α is \mathbf{g}_1 -totally geodesic; this family of metrics is also required to be cylindrical in the connected sum region (for the appropriate uniformity in T). In fact, we find it convenient to use an s -independent product Kähler metric on the region $\text{Sym}^{g+1}(\Sigma - B_{R_1}) \times \text{Sym}^1(E - B_{R_2}) \subset \text{Sym}^{g+1}(\Sigma')$. Then, we define

$$\text{Exp}(a, \xi)(s + it) = \exp_{\widehat{\gamma}_T(s+it)}^{\mathbf{g}_s}(a(s + it)) \times \exp_{s+it}(\xi),$$

where in the first factor

$$\exp_x^{\mathbf{g}}: T\text{Sym}^{g+1}(\Sigma') \longrightarrow \text{Sym}^{g+1}(\Sigma')$$

denotes the exponential map for the metric \mathbf{g} over $\text{Sym}^{g+1}(\Sigma')$ at the point $x \in \text{Sym}^{g+1}(\Sigma')$; while the second denotes simply the exponential map for \mathbb{D} .

The range of this exponential map consists of holomorphic disks with the appropriate boundary conditions.

We can think of the $\bar{\partial}$ -operator extended to $\text{Map}(\mathbb{D}, \text{Sym}^{g+1}(\Sigma')) \times \text{Diff}(\mathbb{D})$, which sends the pair (u, ϕ) to the section $\Gamma(\mathbb{D}, \text{End}(T\mathbb{D}, u^*(T\text{Sym}^{g+1}(\Sigma'))))$, defined by

$$\bar{\partial}_{J'_s}^{\text{ext}}(u, \xi) = (u^* J'_{\phi_s^{-1}}) \circ du - du \circ j_\xi,$$

Clearly, a section (u, ϕ) lies in the kernel of this operator if and only if $u \circ \phi$ is a J'_s -holomorphic map from \mathbb{D} to $\text{Sym}^{g+1}(\Sigma')$.

As in [9], this section has a second-order expansion

$$\bar{\partial}\text{Exp}_{\hat{\gamma}}(a, \xi) = \bar{\partial}\hat{\gamma} + D_{\hat{\gamma}}(a, \xi) + N_{\hat{\gamma}_T}(a, \xi),$$

for some operator

$$N_{\hat{\gamma}_T}: \mathcal{X}(T) \longrightarrow \mathcal{Z}(T).$$

LEMMA 10.13. *There are constants ε and c depending on u with the property that for all T large enough for $\hat{\gamma}_T$ to be defined, if $\|\alpha_1\|_{\mathcal{X}(T)}, \|\alpha_2\|_{\mathcal{X}(T)} < \varepsilon$, then*

$$\|N_{\hat{\gamma}_T}(\alpha_1) - N_{\hat{\gamma}_T}(\alpha_2)\|_{\mathcal{Z}(T)} \leq c\|\alpha_1 - \alpha_2\|_{\mathcal{X}(T)} (\|\alpha_1\|_{\mathcal{X}(T)} + \|\alpha_2\|_{\mathcal{X}(T)}).$$

Proof. We are free to use, metrics near $\{w_i\} \times \{c\} \in \text{Sym}^g(\Sigma - B_{r_1}) \times \text{Sym}^1(E - B_{r_2}) \subset \text{Sym}^{g+1}(\Sigma')$ which are Euclidean. In these neighborhoods, then, the nonlinear part of the second-order expansion vanishes. But these are the regions where the map depends on the neck-length T . \square

We can now apply Newton's iteration scheme (see for instance [5], or [21] in a more closely related context) to find the unique holomorphic map $\gamma = \gamma(u)$ in a sufficiently small neighborhood of $\hat{\gamma}$.

PROPOSITION 10.14. *There is an $\varepsilon > 0$ with the property that for all sufficiently large $T > 0$, there is a unique holomorphic curve $\gamma(u)$ which lies in an ε -neighborhood of $\hat{\gamma}(u)$, measured in the $\mathcal{X}(T)$ norm (for T sufficiently large relative to S).*

Proof. Given $\hat{\gamma}$, we find, for all sufficiently large T , a holomorphic

$$\gamma = \hat{\gamma} + P_T \eta,$$

by finding $\eta \in \mathcal{Z}(T)$ with the property that

$$\eta + N_{\hat{\gamma}}(P_T \eta) = \bar{\partial}_{J'_s(T)} \hat{\gamma}.$$

This can be done, since the map

$$\eta \mapsto N_{\hat{\gamma}} \circ P_T \eta$$

is uniformly quadratically contracting (this is a consequence of Lemmas 10.12 and 10.13), and the error term $\bar{\partial}_{J'_s(T)} \hat{\gamma}$ goes to zero in $\mathcal{Z}(T)$, according to Lemma 10.9. \square

In fact, from the construction of the right inverses, it is clear that the kernel of $D_{\hat{\gamma}}$ is identified with the kernel of the original D_u (see Lemma 10.10 and Proposition 10.12). Moreover, since D_{γ} is a small perturbation of $D_{\hat{\gamma}}$, their kernels are identified as well. Hence the deformation theory of u is identified with the deformation theory of $\gamma(u)$ (this was the claim made in Remark 10.5). This can also be used to identify the corresponding determinant line bundles, so that the signs appearing in the signed counts agree.

We wish to show that for sufficiently large $T > 0$, all the holomorphic curves in $\mathcal{M}(\mathbf{x}', \mathbf{y}')$ are contained in the domain of the gluing map constructed in Proposition 10.14. This can be seen after an application of the Gromov compactness theorem. To apply Gromov's theorem, it is important to set up a uniform version of the energy bounds from Subsection 3.4 which hold as $T \mapsto \infty$. To this end, we find it convenient to think of the degeneration (as $T \mapsto \infty$) as the formation of a node. Specifically, we consider a holomorphic one-parameter family of holomorphic curves parametrized by the complex disk

$$f: X \longrightarrow D,$$

whose fiber is singular at $t = 0$, and which contains the complex structures $j'(T)$ for all large T (i.e. X is an algebraic variety, together with a map to the disk, whose fiber X_t at over a point $t \in D$ is a curve of genus $g + 1$ whenever $t \neq 0$; and the limit $t \mapsto 0$ corresponds to $T \mapsto \infty$). Forming the fiberwise symmetric product, we obtain an algebraic variety

$$f: S(X) \longrightarrow D,$$

whose fiber over t is identified as $S(X)_t \cong \text{Sym}^{g+1}(X_t)$ (see also [6] for a symplectic construction of this object).

PROPOSITION 10.15. *Fix a class $\phi \in \pi_2(\mathbf{x}', \mathbf{y}')$. Then any sequence of holomorphic disks $u_t \in \mathcal{M}_{J_t}(\phi)$ with $t \mapsto 0$ has a Gromov limit which is a holomorphic disk in $\text{Sym}^{g+1}(\Sigma \vee E)$.*

Proof. We can work in an ambient manifold by embedding the algebraic variety $S(X)$ into a complex projective space, from which it inherits the Kähler form ω . As in Subsection 3.4, we must obtain an *a priori* bound on the ω -energy of any disk $u \in \mathcal{M}(\phi)$ (thought of as a disk in $S(X)$). We imitate the discussion from Subsection 3.4, only now in families.

The product form ω_0 gives a Kähler form on the product $X^{\times(g+1)}$. Restricting this to the subvariety $P(X) \subset X^{\times(g+1)}$, defined by

$$P(X) = \{(x_1, \dots, x_{g+1}) \in X^{\times(g+1)} \mid f(x_1) = \dots = f(x_{g+1})\},$$

we obtain a quadratic form $\mathfrak{g}_0(v) = \omega_0(v, Jv)$ on the tangent cone of $P(X)$. There is a holomorphic map $\pi: P(X) \longrightarrow S(X)$, under which we can pull back ω to the subvariety, to obtain another quadratic form $g(v) = \omega(v, Jv)$ on

the tangent cone. Note that \mathfrak{g}_0 is nowhere vanishing. Thus, we can form the ratio $\mathfrak{g}/\mathfrak{g}_0$, to obtain a continuous function on the projectivized tangent cone of $P(X)$. By compactness, this function is bounded above by some constant which we will denote by C_3 (since it is the constant C_3 appearing in Inequality (6)). It follows that if we have a $\text{Sym}^{g+1}(\mathfrak{j}'(T))$ -holomorphic map $u: \Omega \rightarrow P(X)$, then

$$\int_{\Omega} u^*(\omega) \leq C_3 \int_{\hat{\Omega}} \hat{u}^*(\omega_0).$$

But we must consider perturbations of these.

To control the energy integrand in the region where the complex structures are varying, consider first the (holomorphic) commutative diagram

$$\begin{array}{ccc} (\Sigma - B_{R_1}(\sigma_1))^{\times g} \times (E - B_{R_2}(\sigma_2)) \times D & \xrightarrow{\tilde{\iota}} & P(X) \\ \pi \downarrow & & \pi' \downarrow \\ \text{Sym}^g(\Sigma - B_{R_1}(\sigma_1)) \times (E - B_{R_2}(\sigma_2)) \times D & \xrightarrow{\iota} & S(X), \end{array}$$

where the vertical maps are induced by the quotients, and the maps all commute with projection to the base D . If we let V be some open subset of the diagonal in $\text{Sym}^g(\Sigma - B_{R_1}(\sigma_1))$, and \tilde{V} be the pre-image of its closure in $(\Sigma - B_{R_1}(\sigma_1))^{\times g}$, then the restriction of π to $\tilde{V} \times D$ is a covering space. Thus, we can consider the differential form $\pi_* \tilde{\iota}^*(\omega_0)$ over $V \times (E - B_{R_2}(\sigma_2)) \times D \subset \text{Sym}^g(\Sigma - B_{R_1}(\sigma_1)) \times (E - B_{R_2}(\sigma_2)) \times D$. Since the maps are all holomorphic local diffeomorphisms, we observe that the complex structure $\text{Sym}^g(\mathfrak{j}) \times \mathfrak{j}_E \times \mathfrak{j}_D$ tames the form $\pi_* \tilde{\iota}^*(\omega_0)$. It follows that if J'_s is a sufficiently small perturbation of the constant path $\text{Sym}^g(\mathfrak{j}) \times \mathfrak{j}_E$ (which is the restriction of $\text{Sym}^{g+1}(\mathfrak{j})$), then the taming condition is preserved. With respect to such a path J'_s , then, the energy of a J'_s -holomorphic map u from \mathbb{D} to any given t -fiber of $P(X)$ can be controlled as in Lemma 3.5 by the ω_0 -integral of its associated branched cover \hat{u} . On the other hand, for a given ϕ , this integral is controlled by some multiplicity (depending on the maximal multiplicity in $\mathcal{D}(\phi)$, which is, of course, a topological quantity) of the ω_0 -area a fiber of X . But the latter area is bounded in the family, since it is obtained from a smooth symplectic form which extends over the family X .

This gives us the uniform energy bound on the u_t required by Gromov's compactness theorem, and hence gives rise to a limiting cusp-curve which maps into $S(X)$. Indeed, by continuity of the limiting process, it follows that the cusp-curve maps to the $t = 0$ fiber, $\text{Sym}^{g+1}(\Sigma \vee E)$. \square

With respect to the cylindrical geometries on the domain and the range, this gives us the following:

PROPOSITION 10.16. *Fix $\mathbf{x}, \mathbf{y} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$, and let \mathbf{x}', \mathbf{y}' be their stabilizations. Fix a homotopy class $\phi' \in \pi_2(\mathbf{x}', \mathbf{y}')$ with $\mu(\phi') = 1$. There is a sequence*

of holomorphic curves $\{u_i\}_{m=1}^\infty$ with $u_m \in \mathcal{M}_{J'_s(T_m)}(\phi')$ for some unbounded sequence of real numbers $\{T_m\}$, with $\mu(u_m) = 1$. Then, for each $\varepsilon > 0$, after passing to a suitable subsequence of the $\{u_m\}$ if necessary, the following data are obtained:

- a collection of points $\{w_1, \dots, w_n\} \in \text{Sym}^g(\Sigma)$,
- a pair u and v , where $u \in \mathcal{M}(\phi)$, and

$$v \in \mathcal{M}^{\text{cent}}(\mathbb{S} \longrightarrow \text{Sym}^{g-1}(\Sigma_1) \times \text{Sym}^2(\Sigma_2)),$$
- a decomposition of the strip $\mathbb{D} = X_1 \cup_{T'_m} X_2$ (and, of course, a corresponding subset $\{q_1, \dots, q_n\} \subset \mathbb{D}$),
- a sequence $\{\xi_m\} \in \mathbb{C}^n$ converging to zero, with the property that $u_m \circ \phi_{\xi_m}(q_m) \in W'$ (recall that ϕ_{ξ_m} is some “translation” of the strip $\mathbb{D} = X_1 \cup_{T'_m} X_2$ which carries the center q_i of the i^{th} component of X_2 to some nearby point),

which satisfy the following properties:

- (1) Each $u_m \circ \phi_{\xi_m}$ maps X_1 into

$$\text{Sym}^g(\Sigma - B_{r_1}(\sigma_1)) \times \text{Sym}^1(E - B_{r_2}(\sigma_2)) \subset \text{Sym}^{g+1}(\Sigma').$$

- (2) Each $u_m \circ \phi_{\xi_m}$ maps the i^{th} component of X_2 into

$$B_\varepsilon(w_i) \times B_\varepsilon(c) \subset \text{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1)) \times \text{Sym}^1(E - B_{r_2}(\sigma_2)) \subset \text{Sym}^{g+1}(\Sigma').$$

- (3) For any $S > 0$, the restrictions $\{u_m \circ \phi_{\xi_m}|_{X_1(S)}\}_{m=1}^\infty$ and $\{u_m \circ \phi_{\xi_m}|_{X_2(S)}\}_{m=1}^\infty$ converge in the C^∞ topology to the restrictions $u \times \{c\}|_{X_1(S)}$ and $v|_{X_2(S)}$ respectively, where $u \times \{c\}$ and v are thought of as maps

$$u \times \{c\}: X_1(\infty) \longrightarrow \text{Sym}^g(\Sigma - \sigma_1) \times \{c\}$$

and

$$v: X_2(\infty) \longrightarrow \text{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1)) \times \text{Sym}^2(E - \sigma_2) \subset \text{Sym}^{g+1}(\Sigma')$$

respectively.

- (4) Under $u_m \circ \phi_{\xi_m}$, each cylinder $[-T'_m, T'_m] \times S_i^1 \subset X_1 \cup_{T'_m} X_2$ maps to the product

$$B_\varepsilon(w_i) \times ([-T_m, T_m] \times S^1) \times B_\varepsilon(c) \subset \text{Sym}^{g+1}(\Sigma'(T_m)).$$

Proof. Gromov's compactness theorem adapted to our situation (Proposition 10.15) gives a map

$$u_\infty: \mathbb{B} \longrightarrow \mathrm{Sym}^{g+1}(\Sigma \vee E),$$

which is the limit of the $\{u_m\}$, where \mathbb{B} here is a bubble-tree obtained by attaching n spheres \mathbb{S}_i to the disk D . For generic choice of σ_1 , its image lies in the union of $\mathrm{Sym}^g(\Sigma) \times \mathrm{Sym}^1(E)$ and $\mathrm{Sym}^{g-1}(\Sigma) \times \mathrm{Sym}^2(E)$. Since $\mathbb{T}'_\alpha \cup \mathbb{T}'_\beta$ is contained in the first set, it follows that the main component of the bubble-tree is mapped to the first set. Indeed, this gives rise to the $u \times \{c\}$. Moreover, by dimension counts, it follows that the spheres are mapped into the second subset, giving rise to the map v . (Note also that the limiting almost-complex structures J'_s all agree with the constant complex structure, over this latter subset.)

To view the domains of all the $\{u_m\}$ as $X_1 \cup_{T'_m} X_2$ for fixed X_1 , we use the diffeomorphisms generated by vector fields ξ_i ($i = 1, \dots, n$). Specifically, in the proof of the compactness theorem (see, for instance [30]) one rescales around a point where energy is accumulating, to construct the bubbles. These points where the energy concentrates form subsequences of n -tuples $\mathbf{q}_m = \{q_1^{(m)}, \dots, q_n^{(m)}\}$, which converge to \mathbf{q} . Now, after rescaling, we obtain a C^∞_{loc} convergent subsequence to spheres whose energy-measures are centered at the origin. This means that there is a sequence of points $\mathbf{q}'_m = \{q'^{(m)}_1, \dots, q'^{(m)}_n\}$ also converging to \mathbf{q} , with the property that $u_m(q'_i) \in W'$. We let ξ_m be the corresponding translation taking \mathbf{q} to \mathbf{q}'_m , so that $\phi_{\xi_m}(\mathbf{q}) = \mathbf{q}'_m$. Thus, the pairs $(u_m \circ \phi_{\xi_m}, \phi_{\xi_m}^{-1})$ satisfy the extended holomorphic condition on $\mathrm{Map}(X_1 \cup_{T'_m} X_2, \mathrm{Sym}^{g+1}(\Sigma')) \times \mathrm{Diff}(D)$ considered earlier. Moreover, the C^∞_{loc} convergence of the rescaled u_m ensures C^∞_{loc} convergence of the restrictions of $u_m \circ \phi_{\xi_m}|_{X_2(S)}$. The limit is a holomorphic map which is centered (in the sense of Definition 10.8).

The usual C^∞_{loc} on the main component of the bubble-tree gives us Property (1) and the corresponding part of Property (3), while the C^∞_{loc} convergence on the spheres gives us Property (2), and the rest of Property (3).

The C^0 convergence at the bubble points is equivalent to Property (4), in light of the following. Given $w \in \mathrm{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1))$, $\sigma \in \mathrm{Sym}^1(E - B_{r_2}(\sigma_2))$, and letting s be the wedge-point in $\Sigma \vee E$, we have a neighborhood of $w \times s \times \sigma \in \mathrm{Sym}^{g+1}(\Sigma \vee E)$ meeting the nearby fiber $\Sigma'(t)$ of the one-parameter family in a subset identified with the product

$$\begin{aligned} B_\varepsilon(w) \times ([-T, T] \times S^1) \times B_\varepsilon(\sigma_2) \subset \\ \mathrm{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1)) \times ([-T, T] \times S^1) \times \mathrm{Sym}^1(E - B_{r_2}(\sigma_2)) \subset \mathrm{Sym}^{g+1}(\Sigma'). \end{aligned}$$

□

We would like to prove that the properties arising from the compactness above, specifically the behaviour of the maps on the cylinders $u_m|[-T'_m, T'_m] \times S^1 \subset X_1 \cup_{T'_m} X_2$ above, imply stronger decay conditions. According to the proposition, we can view these as maps (which are now holomorphic in the usual sense) to the product space $B_\varepsilon(w_i) \times ([-T_m, T_m] \times S^1) \times B_\varepsilon(c)$. There is an *a priori* bound on the energy of these maps, where the energy is calculated with respect to the symplectic form on the target with collapsing connected sum region. But we wish to use the cylindrical geometry on the target instead, to stay in the geometric framework of Section 10.1. The $L^p_{1,\delta}$ convergence we are aiming for can be divided into two parts, then, corresponding to the two natural projections on the target, which we denote

$$\Pi_2: \text{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1)) \times ([-T, T] \times S^1) \times \text{Sym}^1(E - B_{r_2}(\sigma_2)) \longrightarrow \mathbb{R} \times S^1$$

and

$$\begin{aligned} \Pi_{1,3}: \text{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1)) \times ([-T, T] \times S^1) \times \text{Sym}^1(E - B_{r_2}(\sigma_2)) \\ \longrightarrow \text{Sym}^{g-1}(\Sigma - B_{r_1}(\sigma_1)) \times \text{Sym}^1(E - B_{r_2}(\sigma_2)). \end{aligned}$$

Our aim is to show that the projections $\Pi_{1,3} \circ u_m|[-T'_m, T'_m] \times S^1$ and $\Pi_2 \circ u_m|[-T'_m, T'_m] \times S^1$ all lie in a uniform $L^p_{1,\delta}$ ball. Note that, thanks to the C^∞ convergence on the compact pieces (Property (3)), the restrictions $\{u_m|_{\{-T'_m\} \times S^1}\}$ and $\{u_m|_{\{T'_m\} \times S^1}\}$ lie in C^∞ -bounded subsets, up to translations on the cylindrical factor on the range. In handling the $\Pi_{1,3}$ projection, we do not use this translation, and we can appeal to the following elementary result:

LEMMA 10.17. *Let T_i be an increasing sequence of real numbers, and suppose that $u_i: [-T_i, T_i] \times S^1 \longrightarrow \mathbb{C}$ is a sequence of holomorphic maps of the cylinders into the complex plane, with the property that*

$$u_i|_{([-T_i, -T_i+1] \cup [T_i-1, T_i]) \times S^1}$$

is uniformly C^∞ bounded. Then there is a subsequence $\{w_m\}_{m=1}^\infty \in D$ converging to some $w \in D$, with the property that $u_m - w_m$ is uniformly $L^p_{1,\delta}$ -bounded.

Proof. The point here is that the norm of a holomorphic function u on a cylinder is controlled by its behaviour near the boundary, provided that $\int_{\{T\} \times S^1} u = 0$. More precisely, for each integer $k \geq 0$, we have constants c_k with the property that if

$$u: [-T, T] \times S^1 \longrightarrow \mathbb{C}$$

is a holomorphic map with $\int_{\{-T\} \times S^1} u = 0$, then we have for each sufficiently small $\delta > 0$,

$$\|u\|_{L^2_{k,\delta}} \leq c_k \left(\|u\|_{L^2_{k,\delta}([-T, -T+1] \times S^1)} + \|u\|_{L^2_{k,\delta}([T-1, T] \times S^1)} \right).$$

This follows from an examination of the Fourier coefficients of u . We content ourselves here with the case where $k = 0$.

Write $u(t, \theta) = \sum_{n \neq 0} a_n e^{n(t+i\theta)}$. Then,

$$\begin{aligned} \|u\|_{L^2_\delta}^2 &= \sum_{n>0} |a_n|^2 \int_{-T}^T e^{(2n+\delta)t} + \sum_{n<0} |a_n|^2 \int_{-T}^T e^{(2n+\delta)t} \\ &= \sum_{n>0} |a_n|^2 \int_{-T}^T e^{(2n+\delta)t} + \sum_{n<0} |a_n|^2 \int_{-T}^T e^{(2n+\delta)t} \\ &\leq C \left(\sum_{n>0} |a_n|^2 \int_{T-1}^T e^{(2n+\delta)t} + \sum_{n<0} |a_n|^2 \int_{-T}^{-T+1} e^{(2n+\delta)t} \right) \\ &= C \left(\|u\|_{L^2_\delta([-T, -T+1] \times S^1)}^2 + \|u\|_{L^2_\delta([T-1, T] \times S^1)}^2 \right). \end{aligned}$$

The estimate for larger k follows in a similar manner. The corresponding estimate for $L^p_{1,\delta}$ follows from the uniform, T -independent inclusion (provided T is bounded below)

$$L^2_{2,\delta}([-T, T] \times S^1) \longrightarrow L^p_{1,\delta'}([-T, T] \times S^1),$$

where $\delta' = \delta p/2$, which in turn follows from the fact that the norm of the inclusion $L^2_2 \longrightarrow L^p_1$ is uniformly bounded, provided T is bounded below.

To apply this, let $w_i = \int_{\{-T_i\} \times S^1} u_i$. Clearly, the w_i have a convergent subsequence. The above argument gives a uniform $L^p_{k,\delta}$ -bound on $u_i - w_i$. \square

To handle the Π_2 projection, we must study a similar problem, except now the target is also a cylinder (rather than a disk). To set up notation, let

$$j_1: [0, 1] \times S^1 \longrightarrow \mathbb{R} \times S^1$$

denote the standard inclusion of an annulus into the cylinder.

LEMMA 10.18. *There is an $\varepsilon > 0$ with the property that if $u_m: [-T_m, T_m] \times S^1 \longrightarrow \mathbb{R} \times S^1$ is a sequence of maps with the property that, up to translations and rotations on the cylinder, both sequences $\{u_m|[-T_m, -T_m+1] \times S^1\}$ and $\{u_m|[T_m-1, T_m] \times S^1\}$ lie in an ε neighborhood of j_1 . Then, up to translations and rotations, the u_m lie in a uniform L^2_δ -neighborhood of the standard inclusion $[-T_m, T_m] \times S^1 \subset \mathbb{R} \times S^1$.*

Proof. Let

$$j_{T_m}: [-T_m, T_m] \times S^1 \longrightarrow \mathbb{R} \times S^1$$

denote the natural inclusion. Consider next the difference $j_{T_m} - u_m$. If the C^0 norm of the difference between $u_m|_{\{T_m\} \times S^1}$ and the standard inclusion of the

circle is sufficiently small (up to translation), then the two maps are homotopic. Hence, we can lift the difference $u_m - j_{T_m}$ to a sequence of holomorphic maps

$$u_m - j_{T_m} : [-T_m, T_m] \times S^1 \longrightarrow \mathbb{C},$$

where $\mathbb{C} = \mathbb{R} \times \mathbb{R}$ is thought of as the universal covering space of the annulus. We then apply the argument from Lemma 10.17. \square

Proof of Theorem 10.4. Given $\varepsilon > 0$, we can find a T_0 so that for all $T > T_0$, all holomorphic curves $\gamma \in \mathcal{M}(\phi')$ lie in the image of the map constructed in Proposition 10.14. This follows from Lemma 10.17 and Lemma 10.18. \square

11. Conclusion: topological invariance

We have established all the pieces now to conclude topological invariance of the homology groups, first stated in Theorem 1.1, which justifies our dropping the Heegaard diagram from the notation for the Floer homology groups. As a preliminary remark, recall that Theorem 6.1 shows that the homology groups are independent of the complex structure. We can now give the following more precise statement of Theorem 1.1:

THEOREM 11.1. *If (Y, \mathfrak{s}) is a three-manifold equipped with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, then the Floer homology groups are topological invariants of (Y, \mathfrak{s}) in the following sense. There is a strongly \mathfrak{s} -admissible Heegaard diagram, and if two different strongly \mathfrak{s} -admissible Heegaard diagrams $(\Sigma_1, \alpha_1, \beta_1, z_1)$ and $(\Sigma_2, \alpha_2, \beta_2, z_2)$ represent the same three-manifold, then there is a one-to-one correspondence between isomorphism classes of orientation conventions for the first and the second Heegaard diagram and corresponding $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H_1(Y; \mathbb{Z})/\text{Tors}$ -module isomorphisms (represented by the vertical arrows):*

$$(35) \quad \begin{array}{ccccccc} \cdots & \longrightarrow & HF^-(\Sigma_1, \alpha_1, \beta_1, \mathfrak{s}) & \longrightarrow & HF^\infty(\Sigma_1, \alpha_1, \beta_1, \mathfrak{s}) & \longrightarrow & HF^+(\Sigma_1, \alpha_1, \beta_1, \mathfrak{s}) \longrightarrow \cdots \\ & & \Phi^- \downarrow & & \Phi^\infty \downarrow & & \Phi^+ \downarrow \\ \cdots & \longrightarrow & HF^-(\Sigma_2, \alpha_2, \beta_2, \mathfrak{s}) & \longrightarrow & HF^\infty(\Sigma_2, \alpha_2, \beta_2, \mathfrak{s}) & \longrightarrow & HF^+(\Sigma_2, \alpha_2, \beta_2, \mathfrak{s}) \longrightarrow \cdots \end{array}$$

(where the groups in the first row can be calculated using any orientation system \mathfrak{o}_1 , and the second are calculated using the induced orientation system \mathfrak{o}_2). Indeed, if $(\Sigma_1, \alpha_1, \beta_1, z_1)$ and $(\Sigma_2, \alpha_2, \beta_2, z_2)$ are weakly \mathfrak{s} -admissible, there are isomorphisms

$$\begin{array}{ccccccc} \cdots & \longrightarrow & \widehat{HF}(\Sigma_1, \alpha_1, \beta_1, \mathfrak{s}) & \longrightarrow & HF^+(\Sigma_1, \alpha_1, \beta_1, \mathfrak{s}) & \longrightarrow & HF^+(\Sigma_1, \alpha_1, \beta_1, \mathfrak{s}) \xrightarrow{U} \cdots \\ & & \widehat{\Phi} \downarrow & & \Phi^+ \downarrow & & \Phi^+ \downarrow \\ \cdots & \longrightarrow & \widehat{HF}(\Sigma_2, \alpha_2, \beta_2, \mathfrak{s}) & \longrightarrow & HF^+(\Sigma_2, \alpha_2, \beta_2, \mathfrak{s}) & \xrightarrow{U} & HF^+(\Sigma_2, \alpha_2, \beta_2, \mathfrak{s}) \xrightarrow{U} \cdots \end{array}$$

Proof. The strongly \mathfrak{s} -admissible Heegaard diagrams exist according to Proposition 7.2, which also shows that they can be connected by a sequence of strongly \mathfrak{s} -admissible isotopies, handleslides, and stabilizations. Invariance under those isotopies in the above sense was established in Theorem 7.3. Invariance under handleslides was established in Theorem 9.5 for a special choice of isotopy type of the handleslide. Of course, any handleslide (which does not cross the basepoint) can be brought to this form after an isotopy of the attaching circles (which does not cross the basepoint), so that general handleslide invariance follows. Finally, stabilization invariance was established in Theorem 10.2 (and Theorem 10.1 for \widehat{HF}). When the Heegaard diagram is only weakly \mathfrak{s} -admissible, Lemma 5.8 gives us a strongly \mathfrak{s} -admissible diagram which is isotopic to the given diagram, and Theorem 7.5 gives the necessary identifications between \widehat{HF} and HF^+ . We then reduce to the strongly \mathfrak{s} -admissible case considered before. \square

Recall that $HF_{\text{red}}(Y, \mathfrak{t})$ can be determined from the long exact sequence in Equation (35), so it, too is a topological invariant.

When $b_1(Y) > 0$, we stress that there is still the auxiliary choice of an (isomorphism class of) orientation system, giving rise to $2^{b_1(Y)}$ different candidates for the “Floer homologies.” In fact, in Theorem 10.12 of [28], we show how to identify a canonical orientation system.

It is not difficult to establish naturality of the maps appearing in the above theorem; we return to this point in [29].

11.0.1. *Further remarks.* There are other variants of the Floer homologies defined in the present paper. For example, there are variants twisted with a $\mathbb{Z}[H^1(Y; \mathbb{Z})]$ -module M . We will describe this construction in Section 8 of [28].

There is another construction, which works even in the absence of admissibility hypotheses, which we sketch now. 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 emphasize that the symbol e here is a formal variable).

For a pointed Heegaard diagram $(\Sigma, \alpha, \beta, z)$ for Y , we consider a the chain complex CF_{Nov}^+ which is freely generated (over \mathbb{A}) by pairs $[\mathbf{x}, i] \in (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times \mathbb{Z}^{\geq 0}$, endowed with the boundary map by

$$\partial_{\text{Nov}}^+[\mathbf{x}, i] = \sum_{\{\mathbf{y} \in \mathcal{S}\}} \sum_{\{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \mid n_z(\phi) \leq i\}} e^{\mathcal{A}(\phi)} (\#\mathcal{M}(\phi)) \cdot [\mathbf{y}, i - n_z(\phi)],$$

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 $\eta \in H^2(Y; \mathbb{R})$.

Taking homologies, we obtain homology groups $HF_{\text{Nov}}^+(Y, \mathfrak{s}, \eta)$ which are invariants of the underlying topological data (and an orientation system, which we suppress), and which require no admissibility hypotheses to define. We will have no further use for this construction in [28], 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.

PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY
Current address: COLUMBIA UNIVERSITY, NEW YORK, NY
E-mail address: petero@math.columbia.edu

PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY
E-mail address: szabo@math.princeton.edu

REFERENCES

- [1] L. V. AHLFORS, *Conformal Invariants: Topics in Geometric Function Theory*, McGraw-Hill Book Co., New York, 1973.
- [2] M. F. ATIYAH, *K-theory* (Notes by D. W. Anderson), *Advanced Book Classics*, Addison-Wesley Publishing Company, 1989.
- [3] M. F. ATIYAH, V. K. PATODI, and I. M. SINGER, Spectral Asymmetry and Riemannian Geometry, I, *Math. Proc. Camb. Philos. Soc.* **77** (1975), 43–69.
- [4] V. DE SILVA, Products in the symplectic Floer homology of Lagrangian intersections, Ph.D. thesis, Oxford University, 1999.
- [5] S. K. DONALDSON and P. B. KRONHEIMER, *The Geometry of Four-Manifolds*, *Oxford Mathematical Monographs* Oxford University Press, 1990.
- [6] S. K. DONALDSON and I. SMITH, Lefschetz pencils and the canonical class for symplectic 4-manifolds, math.SG/0012067, 2000.
- [7] A. FLOER, Morse theory for Lagrangian intersections, *J. Differential Geom.* **28** (1988), 513–547.
- [8] ———, A relative Morse index for the symplectic action, *Comm. Pure Appl. Math.* **41** (1988), 393–407.
- [9] ———, The unregularized gradient flow of the symplectic action, *Comm. Pure Appl. Math.* **41** (1988), 775–813.
- [10] A. FLOER and H. HOFER, Coherent orientations for periodic orbit problems in symplectic geometry, *Math. Z.* **212** (1993), 13–38.
- [11] A. FLOER, H. HOFER, and D. SALAMON, Transversality in elliptic Morse theory for the symplectic action, *Duke Math. J.* **80** (1995), 251–329.
- [12] K. FUKAYA, Y-G. OH, K. ONO, and H. OHTA, Lagrangian intersection Floer theory— anomaly and obstruction, preprint, Kyoto University, 2000.

- [13] R. E. GOMPF and A. I. STIPSICZ, *4-Manifolds and Kirby Calculus, Graduate Studies in Mathematics* **20**, A.M.S., Providence, RI, 1999.
- [14] M. GROMOV, Pseudo holomorphic curves in symplectic manifolds, *Invent. Math.* **82** (1985), 307–347.
- [15] J. HARRIS, *Algebraic Geometry: A First Course, Graduate Texts in Mathematics* **133** Springer-Verlag, New York, 1995.
- [16] R. HARTSHORNE, *Algebraic Geometry, Graduate Texts in Mathematics* **52**, Springer-Verlag, New York, 1977.
- [17] A. HATTORI, Topology of \mathbb{C}^n minus a finite number of affine hyperplanes in general position, *J. Fac. Sci. Univ. Tokyo* **22** (1975), 205–219.
- [18] E.-N. IONEL and T. H. PARKER, Relative Gromov-Witten invariants, *Ann. of Math.* **157** (2003), 45–96.
- [19] P. B. KRONHEIMER and T. S. MROWKA, Monopoles and contact structures, *Invent. Math.* **130** (1997), 209–255.
- [20] A.-M. LI and Y. RUAN, Symplectic surgery and Gromov-Witten invariants of Calabi-Yau 3-folds, *Invent. Math.* **145** (2001), 151–218.
- [21] G. LIU, Associativity of quantum multiplication, *Comm. Math. Phys.* **191** (1998), 265–282.
- [22] I. G. MACDONALD, Symmetric products of an algebraic curve, *Topology* **1** (1962), 319–343.
- [23] D. MCDUFF and D. SALAMON, *J-Holomorphic Curves and Quantum Cohomology, University Lecture Series* **6**, A.M.S., Providence, RI, 1994.
- [24] T. S. MROWKA, A local Mayer-Vietoris principle for the Yang-Mills moduli space, Ph.D. Thesis, Univ. of California, Berkeley, 1989.
- [25] Y.-G. OH, Fredholm theory of holomorphic discs under the perturbation of boundary conditions, *Math. Z* **222** (1996), 505–520.
- [26] ———, On the structure of pseudo-holomorphic discs with totally real boundary conditions, *J. Geom. Anal.* **7** (1997), 305–327.
- [27] P. S. OZSVÁTH and Z. SZABÓ, Absolutely Graded Floer homologies and intersection forms for four-manifolds with boundary, *Adv. Math.* **173** (2003), 179–261.
- [28] ———, Holomorphic disks and three-manifold invariants: properties and applications, *Ann. of Math.* **159** (2004), 1159–1245.
- [29] ———, Holomorphic triangles and invariants for smooth four-manifolds, math/0110169, 2001.
- [30] T. H. PARKER and J. G. WOLFSON, Pseudo-holomorphic maps and bubble trees, *J. Geom. Anal.* **3** (1993), 63–98.
- [31] K. REIDEMEISTER, Zur dreidimensionalen Topologie, *Abh. Math. Sem. Univ. Hamburg* **9** (1933), 189–194.
- [32] J. ROBBIN and D. SALAMON, The spectral flow and the Maslov index, *Bull. London Math. Soc.* **27** (1995), 1–33.
- [33] D. SALAMON and E. ZEHNDER, Morse theory for periodic solutions of Hamiltonian systems and the Maslov index, *Comm. Pure Appl. Math.* **45** (1992), 1303–1360.
- [34] M. SCHARLEMANN, Heegaard splittings of compact 3-manifolds, in *Handbook of Geometric Topology*, 921–953, North-Holland, Amsterdam, 2002.
- [35] J. SINGER, Three-dimensional manifolds and their Heegaard diagrams, *Trans. Amer. Math. Soc.* **35** (1933), 88–111.

- [36] C. H. TAUBES, *Metrics, Connections and Gluing Theorems*, *CBMS Regional Conference Series in Mathematics*, A.M.S., Providence, RI, 1996.
- [37] G. TIAN, in *Current Developments in Mathematics*, 1995 (R. Bott, M. Hopkins, A. Jaffe, I. Singer, D. Stroock, and S-T. Yau, eds.), Internat. Press, Cambridge, MA, 1994.
- [38] V. TURAEV, Torsion invariants of Spin^c -Structures on 3-manifolds, *Math. Res. Lett.* **4** (1997), 679–695.
- [39] C. VITERBO, Intersection de sous-variétés lagrangiennes, fonctionnelles d'action et indice des systèmes hamiltoniens, *Bull. Soc. Math. France* **115** (1987), 361–390.
- [40] R. YE, Gromov's compactness theorem for pseudo holomorphic curves, *Trans. Amer. Math. Soc.* **342** (1994), 671–694.
- [41] H. ZIESCHANG, On Heegaard diagrams of 3-manifolds, in *On the Geometry of Differentiable Manifolds* (Rome, 1986), *Astérisque* **163–164** (1988), 247–280.

(Received October 26, 2000)

(Revised May 7, 2002)