

A cylindrical reformulation of Heegaard Floer homology

ROBERT LIPSHITZ

We reformulate Heegaard Floer homology in terms of holomorphic curves in the cylindrical manifold $\Sigma \times [0, 1] \times R$, where Σ is the Heegaard surface, instead of $\text{Sym}^g(\Sigma)$. We then show that the entire invariance proof can be carried out in our setting. In the process, we derive a new formula for the index of the $\bar{\partial}$ -operator in Heegaard Floer homology, and shorten several proofs. After proving invariance, we show that our construction is equivalent to the original construction of Ozsváth–Szabó. We conclude with a discussion of elaborations of Heegaard Floer homology suggested by our construction, as well as a brief discussion of the relation with a program of C Taubes.

57R17; 57R58, 57M27

In [21], P Ozsváth and Z Szabó associated to a three-manifold Y and a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on Y a collection of abelian groups, known together as Heegaard Floer homology. These groups, which are believed to be isomorphic to certain Seiberg–Witten Floer homology groups (Ozsváth–Szabó [20] and Kronheimer–Mannolescu [12]), fit into the framework of a $(3 + 1)$ -dimensional topological quantum field theory. Since its discovery around the turn of the millennium, Heegaard Floer homology has been applied by Ozsváth, Rasmussen and Szabó to the study of knots and surgery [19; 25; 18], contact structures [23] and symplectic structures [22], and is strong enough to reprove most results about smooth four-manifolds originally proved by gauge theory [17]. In this paper we give an alternate definition of the Heegaard Floer homology groups.

Rather than being associated directly to a three-manifold Y , the Heegaard Floer homology groups defined in [21] and in this paper are associated to a Heegaard diagram for Y , as well as a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} and some additional structure. A *Heegaard diagram* is a closed, orientable surface Σ of genus g , together with two g -tuples of pairwise disjoint, homologically linearly independent, simple closed curves $\vec{\alpha} = \{\alpha_1, \dots, \alpha_g\}$ and $\vec{\beta} = \{\beta_1, \dots, \beta_g\}$ in Σ . A Heegaard diagram specifies a three-manifold as follows. Thicken Σ to $\Sigma \times [0, 1]$. Glue thickened disks along the $\alpha_i \times \{0\}$ and along the $\beta_j \times \{1\}$. The resulting space has two boundary components, each homeomorphic to S^2 . Cap each with a three-ball. The result is the three-manifold specified by $(\Sigma, \vec{\alpha}, \vec{\beta})$.

Different Heegaard diagrams can specify the same three–manifold. Two different Heegaard diagrams specify the same three–manifold if and only if they agree after a sequence of moves of the following three kinds:

- Isotopies of the α – or β –circles.
- Handleslides among the α – or β –circles. These correspond to pulling one α – (or β –) circle over another.
- Stabilization, which corresponds to taking the connect sum of the Heegaard diagram with the standard genus–one Heegaard diagram for S^3 .

See Gompf and Stipsicz [9, Sections 4.3 and 5.1] or Ozsváth and Szabó [21, Section 2] for more details.

So, after associating the Heegaard Floer homology groups to a Heegaard diagram, one must prove they are unchanged by these three kinds of Heegaard moves (as well as deforming the additional structure involved in their definition). Doing so comprises most of [21] for the original definition. Similarly, for our definition, most of this paper is involved in proving:

Theorem 1 *The Heegaard Floer homology groups*

$$HF^\infty(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{s}), \quad HF^+(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{s}), \quad HF^-(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{s}) \quad \text{and} \quad \widehat{HF}(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{s})$$

associated to a Heegaard diagram $(\Sigma, \vec{\alpha}, \vec{\beta})$ and $\text{Spin}^{\mathbb{C}}$ –structure \mathfrak{s} are in fact invariants of the pair (Y, \mathfrak{s}) .

We are also able to prove:

Theorem 2 *The Heegaard Floer homology groups defined in this paper are isomorphic to the corresponding groups defined in [21].*

Theorem 2 is proved in Section 13. The proof does not rely on the invariance results proved in this paper; it could be carried out immediately after Section 8. (We defer the proof to the end to avoid interrupting the narrative flow.) Clearly, Theorem 2 implies Theorem 1. However, one key goal of this paper is to demonstrate that the entire invariance proof can be carried out in our setting, and to develop the tools necessary to do so.

The only essentially new results in this paper are in Section 4, where we give a nice formula for the index of the $\bar{\partial}$ operator in our setup, and hence also the Maslov index in the traditional setting, and in the discussion of elaborations of Heegaard Floer homology

in the last section (Section 14). The casual reader might also be interested in looking at the elaboration and speculation in Section 14.

Although this paper is essentially self contained, it is probably most useful to read it in parallel with [21]. To facilitate this, the paper is organized similarly to [21], and throughout there are precise references to corresponding results in their original forms. In addition, the last appendix is a table cross referencing most of the results in this paper with those of [21].

A more technical discussion of the difference between our setup and that of [21] follows.

The original definition of Heegaard Floer homology involves holomorphic disks in $\text{Sym}^g(\Sigma)$. In this paper, we consider holomorphic curves in $\Sigma \times [0, 1] \times \mathbb{R}$. For instance, for us the chain complex \widehat{CF} is generated by g -tuples of Reeb chords $\{x_i \times [0, 1] \mid x_i \in \alpha_i \cap \beta_{\sigma(i)}\}$. For an appropriate almost complex structure J on $\Sigma \times [0, 1] \times \mathbb{R}$, the coefficient of $\{y_i \times [0, 1]\}$ in $\partial(\{x_i \times [0, 1]\})$ is given by counting holomorphic curves in $\Sigma \times [0, 1] \times \mathbb{R}$ asymptotic to $\{x_i \times [0, 1]\}$ at $-\infty$ and to $\{y_i \times [0, 1]\}$ at ∞ , with boundary mapped to the Lagrangian cylinders $\{\alpha_j \times \{1\} \times \mathbb{R}\}$ and $\{\beta_j \times \{0\} \times \mathbb{R}\}$. (We impose a few further technical conditions on the curves that we count; see Section 1.)

If J is the split complex structure $j_\Sigma \times j_\mathbb{D}$ then a holomorphic curve in $\Sigma \times [0, 1] \times \mathbb{R}$ is just a surface S and a pair of holomorphic maps $u_\Sigma : (S, \partial S) \rightarrow (\Sigma, \alpha_1 \cup \dots \cup \alpha_g \cup \beta_1 \cup \dots \cup \beta_g)$ and $u_\mathbb{D} : (S, \partial S) \rightarrow (\mathbb{D}, \partial \mathbb{D})$. If the map $u_\mathbb{D}$ is a g -fold branched covering then this data specifies a map $\mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ as follows. For $p \in \mathbb{D}$, let p_1, \dots, p_g be the preimages of p under $\pi_\mathbb{D} \circ u$, listed with multiplicity. Then the map $\mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ sends p to $\{u_\Sigma(p_1), \dots, u_\Sigma(p_g)\}$.

Note that the idea of viewing a map to $\text{Sym}^g(\Sigma)$ as a pair

$$(\text{a } g\text{-fold covering } S \rightarrow \mathbb{D}, \text{ a map } S \rightarrow \Sigma)$$

is already implicit in [21], although they use this idea mainly for calculations in special cases.

Working in $\Sigma \times [0, 1] \times \mathbb{R}$ has several advantages. A main advantage is that, unlike a g -fold symmetric product, one can actually visualize $\Sigma \times [0, 1] \times \mathbb{R}$. A second advantage is that a number of the technical details become somewhat simpler. The main disadvantage is that we must now consider higher genus holomorphic curves, not just disks. Another difficulty is that our setup requires compactness for holomorphic curves in manifolds with cylindrical ends, proved by Bourgeois et al in [2]. I also borrow from the language of symplectic field theory. Fortunately, much of the subtle machinery of symplectic field theory, like virtual cycles or the operator formalism, is unnecessary for this paper.

The paper is organized as follows. The first two sections are devoted to basic definitions and notation, and certain algebro–topological considerations. The third section proves transversality results necessary for the rest of the paper. These results should be standard, but I am unaware of a reference that applies to our setting.

The fourth section discusses the index of the $\bar{\partial}$ –operator in our context. We prove this index is the same as the Maslov index in the traditional setting, and obtain a combinatorial formula for it. The fifth section discusses so–called admissibility criteria necessary for the case $b_1(Y) > 0$. The definitions and results are completely analogous with [21]. The sixth section discusses coherent orientations of the moduli spaces. Again, our treatment is close to [21].

The seventh section rules out undesirable codimension–one degenerations of our holomorphic curves. After doing so, we are finally ready to define the Heegaard Floer chain complexes in Section 8, and turn to the invariance proof. The ninth section proves isotopy independence. Before proving handleslide independence, we introduce triangle maps in Section 10. (As in [21], to a *Heegaard triple–diagram* $(\Sigma, \alpha_1, \dots, \alpha_g, \beta_1, \dots, \beta_g, \gamma_1, \dots, \gamma_g)$ is associated maps

$$HF(\Sigma, \vec{\alpha}, \vec{\beta}) \otimes HF(\Sigma, \vec{\beta}, \vec{\gamma}) \rightarrow HF(\Sigma, \vec{\alpha}, \vec{\gamma}),$$

for various decorations of HF .) Using these triangle maps and a model computation, we prove handleslide invariance in Section 11.

Finally, in section twelve we prove stabilization invariance, completing the invariance proof. After this, we devote a section to proving equivalence with traditional Heegaard Floer homology and a section to elaborations and speculation.

There are also two appendices. The first is devoted to the gluing results used throughout the paper. The second cross references our results with those in [21].

For technical results about holomorphic curves, this paper sometimes cites recent sources when older ones would suffice. This generally reflects either that the newer results are more broadly applicable or that I found the newer exposition significantly clearer.

I thank Ya Eliashberg, who is responsible for communicating to me most of the ideas in this paper. I also thank Z Szabó for a helpful conversation about the index; P Ozsváth for a helpful conversation about annoying curves (see Section 8 below) and pointing out a serious omission in Section 13; P Melvin for a stimulating conversation about the index; M Hutchings for a discussion clarifying the relation between the $H_1(Y)$ –action and twisted coefficients; M Hedden and C Wendl for pointing out errors, both typographical and otherwise, in a previous version; and W Hsiang, C Manolescu, L

Ng and B Parker for comments that have improved the exposition. Finally, I thank the referees for finding several errors and making many helpful suggestions.

This work was partially supported by the NSF Graduate Research Fellowship Program, and partly by the NSF Focused Research Group grant DMS-0244663.

1 Basic definitions and notation

By a *pointed Heegaard diagram* we mean a Heegaard diagram (as discussed in the introduction) together with a chosen point z of the Heegaard surface in the complement of the α - and β -circles. Fix a pointed Heegaard diagram $\mathcal{H} = (\Sigma_g, \vec{\alpha} = \{\alpha_1, \dots, \alpha_g\}, \vec{\beta} = \{\beta_1, \dots, \beta_g\}, z)$. Let $\alpha = \alpha_1 \cup \dots \cup \alpha_g \subset \Sigma$ and $\beta = \beta_1 \cup \dots \cup \beta_g \subset \Sigma$. Consider the manifold $W = \Sigma \times [0, 1] \times \mathbb{R}$. We let (p, s, t) denote a point in W (so $p \in \Sigma$, $s \in [0, 1]$ and $t \in \mathbb{R}$). Let $\pi_{\mathbb{D}}: W \rightarrow [0, 1] \times \mathbb{R}$, $\pi_{\mathbb{R}}: W \rightarrow \mathbb{R}$ and $\pi_{\Sigma}: W \rightarrow \Sigma_g$ denote the obvious projections. Consider the cylinders $C_{\alpha} = \alpha \times \{1\} \times \mathbb{R}$ and $C_{\beta} = \beta \times \{0\} \times \mathbb{R}$. We will obtain Heegaard Floer homology by constructing a boundary map counting holomorphic curves with boundary on $C_{\alpha} \cup C_{\beta}$ and appropriate asymptotics at $\pm\infty$.

We shall always assume $g > 1$, as the $g = 1$ case is slightly different technically. Since we can stabilize any Heegaard diagram, this does not restrict the class of manifolds under consideration.

Fix a point z_i in each component D_i of $\Sigma_g \setminus (\alpha \cup \beta)$. Let dA be an area form on Σ , and j_{Σ} a complex structure on Σ tamed by dA . Let $\omega = ds \wedge dt + dA$, a split symplectic form on W . Let J be an almost complex structure on W such that

- (J1) J is tamed by ω .
- (J2) In a cylindrical neighborhood $U_{\{z_i\}}$ of $\{z_i\} \times [0, 1] \times \mathbb{R}$, $J = j_{\Sigma} \times j_{\mathbb{D}}$ is split. (Here, $U_{\{z_i\}}$ is small enough that its closure does not intersect $(\alpha \cup \beta) \times [0, 1] \times \mathbb{R}$.)
- (J3) J is translation invariant in the \mathbb{R} -factor.
- (J4) $J(\partial/\partial t) = \partial/\partial s$
- (J5) J preserves $T(\Sigma \times \{(s, t)\}) \subset TW$ for all $(s, t) \in [0, 1] \times \mathbb{R}$.

The first requirement is in order to obtain compactness of the moduli spaces. The second is for “positivity of domains” (see twelve paragraphs below). The third and fourth make W *cylindrical* as defined in [2, Section 2.1]. The fifth ensures that our complex structure is *symmetric* and *adjusted to ω* in the sense of [2, Section 2.1 and Section 2.2]. (Note that W is Levi-flat as defined there. The vector field \mathbf{R} introduced there is $\partial/\partial s$. The form λ is just ds .)

Note that we can view J as a path J_s of complex structures on Σ . Also notice that C_α and C_β are Lagrangian with respect to ω .

At one point later – the proof of 8.2 – we need to consider almost complex structures which, instead of satisfying **(J5)**, satisfy the slightly less restrictive condition

(J5') there is a 2–plane distribution ξ on $\Sigma \times [0, 1]$ such that the restriction of ω to ξ is non–degenerate, J preserves ξ , and the restriction of J to ξ is compatible with ω . We further assume that ξ is tangent to Σ near $(\alpha \cup \beta) \times [0, 1]$ and near $\Sigma \times (\partial[0, 1])$.

This still guarantees that J is symmetric and adjusted to ω .

By an *intersection point* we mean a set of g distinct points $\vec{x} = \{x_1, \dots, x_g\}$ in $\alpha \cap \beta$ such that exactly one x_i lies on each α_j and exactly one x_i lies on each β_k . (This corresponds to an intersection point of the α – and β –tori in [21].)

Observe that the characteristic foliation on $\Sigma \times [0, 1]$ induced by ω has leaves $\{p\} \times [0, 1] \times \{t\}$. So, an intersection point \vec{x} specifies a g –tuple of distinct “Reeb chords” (with respect to the characteristic foliation on $\Sigma \times [0, 1]$ induced by ω) in $\Sigma \times [0, 1]$ with boundaries on $\alpha \times \{1\} \cup \beta \times \{0\}$. (The collection of Reeb chords is just $\{x_i\} \times [0, 1]$.) We will call a g –tuple of Reeb chords at $\pm\infty$ specified by an intersection point an *I–chord collection*. (I stands for “intersection.”) We will abuse notation and also use \vec{x} to denote the I–chord collection specified by \vec{x} .

Let \mathcal{M} denote the moduli space of Riemann surfaces S with boundary, g “negative” punctures $\vec{p} = \{p_1, \dots, p_g\}$ and g “positive” punctures $\vec{q} = \{q_1, \dots, q_g\}$, all on the boundary of S , and such that S is compact away from the punctures.

For J satisfying **(J1)**–**(J5)**, we will consider J –holomorphic maps $u: S \rightarrow W$ such that

- (M0)** The source S is smooth.
- (M1)** $u(\partial(S)) \subset C_\alpha \cup C_\beta$.
- (M2)** There are no components of S on which $\pi_{\mathbb{D}} \circ u$ is constant.
- (M3)** For each i , $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ and $u^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ each consist of exactly one component of $\partial S \setminus \{p_1, \dots, p_g, q_1, \dots, q_g\}$.
- (M4)** $\lim_{w \rightarrow p_i} \pi_{\mathbb{R}} \circ u(w) = -\infty$ and $\lim_{w \rightarrow q_i} \pi_{\mathbb{R}} \circ u(w) = \infty$.
- (M5)** The *energy* of u , as defined in [2, Section 5.3], is finite. (For the moduli spaces defined later in the paper we shall always assume this technical condition is satisfied, but shall not usually state it.)




Figure 1: A curve in W we might consider. Note that our curves can also be disconnected.

(M6) u is an embedding.

Note that condition **(M3)** implies that $\partial S \setminus \{p_1, \dots, p_g, q_1, \dots, q_g\}$ consist of exactly $2g$ components, none of them compact. Also note that we allow holomorphic curves to be disconnected.

It follows from [2, Proposition 5.8] that near each negative puncture (respectively positive puncture), a holomorphic curve satisfying **(M0)**–**(M6)** converges exponentially (in t) to an I–chord collection \vec{x} (respectively \vec{y}) at $-\infty$ (respectively ∞). We say the holomorphic curve connects \vec{x} to \vec{y} . It follows from this asymptotic convergence to Reeb chords that $\pi_{\mathbb{D}} \circ u$ is a g –fold branched covering map.

Consider the space $\overline{W} = \Sigma \times [0, 1] \times [-1, 1]$ as a compactification of W . Let $\overline{C}_\alpha, \overline{C}_\beta$ denote the closures of the images of C_α and C_β in \overline{W} . Let \overline{S} denote the surface obtained by blowing up S at the punctures. Then, the asymptotic convergence to Reeb orbits mentioned earlier implies that u can be extended to a continuous map $\overline{u}: \overline{S} \rightarrow \overline{W}$. (Compare, for example, [2, Proposition 6.2].)

Let $\pi_2(\vec{x}, \vec{y})$ denote the set of homology classes of continuous maps $(S, \partial S) \rightarrow (W, C_\alpha \cup C_\beta)$ which converge to \vec{x} (respectively \vec{y}) near the negative (respectively

positive) punctures of S . That is, two such maps are equivalent if they induce the same element in $H_2(\overline{W}, \overline{C_\alpha} \cup \overline{C_\beta} \cup (\{x_i\} \times [0, 1] \times \{-1\}) \cup (\{y_i\} \times [0, 1] \times \{1\}))$. (The notation is chosen to be consistent with [21], where the notation π_2 makes sense.)

Each holomorphic curve connecting \vec{x} to \vec{y} represents an element of $\pi_2(x, y)$. For $A \in \pi_2(x, y)$, we denote by \mathcal{M}^A the space of holomorphic curves connecting \vec{x} and \vec{y} in the homology class A . (We always mod out by automorphism of the source S .) Since we are considering cylindrical complex structures, \mathbb{R} acts on \mathcal{M}^A by translation. Let $\widehat{\mathcal{M}}^A = \mathcal{M}^A / \mathbb{R}$. We denote by $\widehat{\widehat{\mathcal{M}}}^A$ the compactification, as in [2, Section 7], of $\widehat{\mathcal{M}}^A$.

Given a homology class $A \in \pi_2(\vec{x}, \vec{y})$, let $n_3(A)$ denote the intersection number of A with $\{3\} \times [0, 1] \times \mathbb{R}$. Define $n_{3_i}(A)$ similarly. If u is a curve in the homology class A we will sometimes write $n_3(u)$ or $n_{3_i}(u)$ for $n_3(A)$ or $n_{3_i}(A)$. We say that a homology class A is *positive* if $n_{3_i}(A) \geq 0$ for all i . Notice that if A has a holomorphic representative (with respect to any complex structure satisfying **(J2)**) then A is positive; this is the *positivity of domains* mentioned twelve paragraphs above. We shall let $\widehat{\pi}_2(\vec{x}, \vec{y}) = \{A \in \pi_2(\vec{x}, \vec{y}) \mid n_3(A) = 0\}$. Elements of $\widehat{\pi}_2(\vec{x}, \vec{x})$ are called *periodic classes*.

Remark In fact, even without **(J2)**, positivity of domains would still hold by Micallef and White [15, Theorem 7.1]. (See also Lemma 3.1.) On the other hand, by requiring **(J2)**, which is easy to obtain, we can avoid invoking here this hard analytic result.

Given a homology class A , we define the *domain of A* to be the formal linear combination $\sum n_{3_i}(A) D_i$. If u represents A then we define the domain of u to be the domain of A . The domains of periodic classes are called *periodic domains*.

As in [21], concatenation makes $\pi_2(\vec{x}, \vec{y})$ into a $\pi_2(\vec{x}, \vec{x})$ -torseur. We shall sometimes write concatenation with a $+$ and sometimes with a $*$, depending on whether we are thinking of domains or maps.

2 Homotopy preliminaries

These issues are substantially simplified from [21] because we need only deal with homology, not homotopy. This is reasonable: by analogy to the Dold–Thom theorem, the low-dimensional homotopy theory of $\text{Sym}^g(\Sigma)$ should agree with the homology theory of Σ .

Given an intersection point \vec{x} , observe that projection from W gives rise to an isomorphism from $\pi_2(\vec{x}, \vec{x})$ to $H_2(\Sigma_g \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$. Given intersection points

\vec{x}, \vec{y} , either $\pi_2(\vec{x}, \vec{y})$ is empty or $\pi_2(\vec{x}, \vec{y}) \cong H_2(\Sigma_g \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$. The isomorphism is not canonical; it is given by fixing an element of $\pi_2(\vec{x}, \vec{y})$ and then subtracting the homology class it represents from all other elements of $\pi_2(\vec{x}, \vec{y})$. We calculate $H_2(\Sigma_g \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$:

Lemma 2.1 (Compare [21, Proposition 2.15]) *There is a natural short exact sequence*

$$0 \rightarrow \mathbb{Z} \rightarrow H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}) \rightarrow H_2(Y) \rightarrow 0.$$

The choice of basepoint z gives a splitting $n_3: H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}) \rightarrow \mathbb{Z}$ of this sequence.

Proof The long exact sequence for the pair $(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$ gives

$$0 \rightarrow H_2(\Sigma \times [0, 1]) \rightarrow H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}) \rightarrow H_1(\alpha \times \{1\} \cup \beta \times \{0\}).$$

The image of the last map is isomorphic to $H_1(\alpha) \cap H_1(\beta)$, viewed as a submodule of $H_1(\Sigma)$.

Let $Y = U_1 \cup_\Sigma U_2$ be the Heegaard splitting. The Mayer–Vietoris sequence gives

$$H_2(U_1) \oplus H_2(U_2) \rightarrow H_2(Y) \rightarrow H_1(\Sigma) \rightarrow H_1(U_1) \oplus H_1(U_2).$$

Here, the kernel of the last map is $H_1(\alpha) \cap H_1(\beta)$. The groups $H_2(U_1)$ and $H_2(U_2)$ are both trivial, so $H_2(Y) \cong H_1(\alpha) \cap H_1(\beta)$. Combining this with the first sequence and using the fact that $H_2(\Sigma \times [0, 1]) \cong \mathbb{Z}$ gives the first part of the claim. With n_3 defined as in Section 1 the second part of the claim is obvious. \square

If we identify $\Sigma \times [0, 1]$ with $f^{-1}[3/2 - \epsilon, 3/2 + \epsilon]$ for some self-indexing Morse function f on Y then the map $H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}) \rightarrow H_2(Y)$ is simply given by “capping off” a cycle with the ascending / descending disks from the index 1 and 2 critical points of f . Also, notice that a homology class $A \in \pi_2(\vec{x}, \vec{y})$ specifies and is specified by its domain. (The domain need not, however, specify uniquely the intersection points \vec{x} and \vec{y} which it connects.)

Following [21, Section 2.6], we observe that a choice of basepoint z and intersection point \vec{x} specify a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on Y as follows. Choose a metric $\langle \cdot, \cdot \rangle$ and a self-indexing Morse function f which specify the Heegaard diagram \mathcal{H} . Then \vec{x} specifies a g -tuple of flows of ∇f from the index 1 critical points of f to the index 2 critical points of f . The point z lies on a flow from the index 0 critical point of f to the index 3 critical point of f . Choose small ball neighborhoods of (the closure of) each of these flow lines. Call the union of these neighborhoods B . Then, in the complement of B , ∇f is nonvanishing. One can extend ∇f to a nonvanishing

vector field v on all of Y . The vector field v reduces the structure group of TY from $SO(3)$ to $SO(2) \oplus SO(1) \cong U(1) \oplus 1 \subset U(2) = \text{Spin}^{\mathbb{C}}(3)$, and thus determines a $\text{Spin}^{\mathbb{C}}$ -structure on Y . We have, thus, defined a map \mathfrak{s}_3 from the set of intersection points in \mathcal{H} to the set of $\text{Spin}^{\mathbb{C}}$ -structures on Y .

It is clear that the $\text{Spin}^{\mathbb{C}}$ -structure $\mathfrak{s}_3(\vec{x})$ is independent of the metric and particular Morse function used to define it.

Given a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on Y , we shall often suppress 3 and write $\vec{x} \in \mathfrak{s}$ to mean $\mathfrak{s}_3(\vec{x}) = \mathfrak{s}$.

Note that by the previous construction, any nonvanishing vector field on a 3-manifold Y gives rise to a $\text{Spin}^{\mathbb{C}}$ -structure. It is not hard to show that two nonvanishing vector fields give rise to the same $\text{Spin}^{\mathbb{C}}$ -structure if and only if they are *homologous*, ie, homotopic through nonvanishing vector fields in the complement of some 3-ball; see [27]. We will use the analogous construction in the case of 4-manifolds in Subsubsection 10.1.2.

Our reason for introducing $\text{Spin}^{\mathbb{C}}$ -structures will become clear in a moment. First, one more definition. Fix a pair of intersection points \vec{x} and \vec{y} , as well as a Morse function f and Riemannian metric $\langle \cdot, \cdot \rangle$ which realize the Heegaard diagram. This data specifies a homology class $\epsilon(\vec{x}, \vec{y})$ as follows. Regard each of \vec{x} and \vec{y} as (the closure of) a g -tuple of gradient flow trajectories in Y from the g index 1 critical points to the g index 2 critical points. Then, $\vec{x} - \vec{y}$ is a 1-cycle in Y . We define $\epsilon(\vec{x}, \vec{y})$ to be the homology class in $H_1(Y)$ of the 1-cycle $\vec{x} - \vec{y}$.

The element $\epsilon(\vec{x}, \vec{y})$ can be calculated entirely in \mathcal{H} by the following equivalent definition. Let γ_α (respectively γ_β) be a 1-cycle in α (respectively β) such that $\partial\gamma_\alpha = \partial\gamma_\beta = \vec{x} - \vec{y}$. Then $\gamma_\alpha - \gamma_\beta$ is a 1-cycle in Σ . Define $\epsilon(\vec{x}, \vec{y})$ to be the image of $\gamma_\alpha - \gamma_\beta$ under the map

$$H_1(\Sigma) \rightarrow \frac{H_1(\Sigma)}{H_1(\alpha) + H_1(\beta)} \cong H_1(Y).$$

The equivalence of the two definitions is easy: in the notation used just above, $\vec{x} + \gamma_\alpha - \gamma_\beta$ is homologous, rel endpoints, to \vec{y} . It is obvious that the second definition is independent of the choices of γ_α and γ_β .

The following lemma justifies our introduction of ϵ and of $\text{Spin}^{\mathbb{C}}$ -structures:

Lemma 2.2 (Compare [21, Proposition 2.15, Lemma 2.19]) *Given a pointed Heegaard diagram \mathcal{H} and intersection points \vec{x} and \vec{y} , the following are equivalent:*

- (1) $\pi_2(\vec{x}, \vec{y})$ is nonempty

- (2) $\epsilon(\vec{x}, \vec{y}) = 0$
- (3) $\mathfrak{s}_3(\vec{x}) = \mathfrak{s}_3(\vec{y})$.

Proof

(1) \Rightarrow (2) Let $A \in \pi_2(\vec{x}, \vec{y})$. View A as a domain in Σ , ie, a chain in Σ . Then we can use ∂A to define $\epsilon(\vec{x}, \vec{y})$, which is thus zero in homology.

(2) \Rightarrow (1) Suppose that $\epsilon(\vec{x}, \vec{y}) = 0$. Then, using the same notation as just before the lemma, for an appropriate choice of γ_α and γ_β , $\gamma_\alpha - \gamma_\beta$ is null-homologous in $H_1(\Sigma)$. We can assume that γ_α and γ_β are cellular 1-chains in the cellulation of Σ induced by the Heegaard diagram. Then, there is a cellular 2-chain A with boundary $\gamma_\alpha - \gamma_\beta$, and A is the domain of an element of $\pi_2(\vec{x}, \vec{y})$.

(2) \Leftrightarrow (3) Let $v_{\vec{x}}$ and $v_{\vec{y}}$ denote the vector fields used to define $\mathfrak{s}_3(\vec{x})$ and $\mathfrak{s}_3(\vec{y})$, respectively. Let $v_{\vec{y}} = Av_{\vec{x}}$ where $A: Y \rightarrow SO(3)$. Let $Fr(v_{\vec{x}}^\perp)$ and $Fr(v_{\vec{y}}^\perp)$ denote the principal $SO(2) = U(1)$ -bundles of frames of $v_{\vec{x}}^\perp$ and $v_{\vec{y}}^\perp$. Then the principal $Spin^C$ bundles induced by $v_{\vec{x}}$ and $v_{\vec{y}}$ are $\mathfrak{s}_z(\vec{x}): Fr(v_{\vec{x}}^\perp) \times_{U(1)} U(2) \rightarrow Fr(TY)$ and $\mathfrak{s}_z(\vec{y}): Fr(v_{\vec{y}}^\perp) \times_{U(1)} U(2) \rightarrow Fr(TY)$.

Note that $\mathfrak{s}_3(\vec{x})$ and $\mathfrak{s}_3(\vec{y})$ are equivalent if and only if A is homotopic to a map $Y \rightarrow SO(2)$. So, the two $Spin^C$ -structures are equivalent if and only if the composition $h: Y \rightarrow SO(3) \rightarrow SO(3)/SO(2) = S^2$ is null homotopic.

Now, homotopy classes of maps from a 3-manifold Y to S^2 correspond to elements of $H^2(Y)$. The Poincaré dual to such a map is the homology class of the preimage of a regular value in S^2 .

For a generic choice of the two Morse functions and metrics used to define them, the flows through \vec{x} and \vec{y} glue together to a disjoint collection of circles γ . Let γ' be a smoothing of γ . The map h is homotopic to a Thom collapse map of a neighborhood of γ' . It follows that the preimage of a regular value is homologous to γ .

So, $\mathfrak{s}_3(\vec{x}) = \mathfrak{s}_3(\vec{y})$ if and only if γ is null-homologous. But γ is a cycle defining $\epsilon(\vec{x}, \vec{y})$, so the result follows. □

Note that the previous proof in fact shows that the map \mathfrak{s}_3 from intersection points to $Spin^C$ -structures is a map of $H^2(Y) = H_1(Y)$ -torseurs.

The following result (part of [21, Lemma 2.19]) is nice to know, but will not be used explicitly in this paper. The reader can imitate the proof of the previous proposition to prove it, or see [21].

Lemma 2.3 *Let \vec{x} be an intersection point of a Heegaard diagram \mathcal{H} , and $\mathfrak{z}_1, \mathfrak{z}_2$ two different basepoints for \mathcal{H} . Suppose that \mathfrak{z}_1 can be joined to \mathfrak{z}_2 by a path \mathfrak{z}_t disjoint from the β circles and such that $\#(\mathfrak{z}_t \cap \alpha_i) = \delta_{i,j}$ (Kronecker delta). Let γ be a loop in Σ such that $\gamma \cdot \alpha_i = \delta_{i,j}$. Then, $\mathfrak{s}_{\mathfrak{z}_2}(\vec{x}) - \mathfrak{s}_{\mathfrak{z}_1}(\vec{x}) = PD(\gamma)$, the Poincaré dual to γ .*

3 Transversality

We need to check that we can achieve transversality for the generalized Cauchy–Riemann equations within the class of almost complex structures satisfying (J1)–(J5). The argument is relatively standard, and is almost the same as the one found in [14, Chapter 3]. This section is somewhat technical, and the reader might want to skip most of it on a first reading.

Before proving our transversality result we need a few lemmas about the geometry of holomorphic curves in W .

Lemma 3.1 *Let $\pi: E \rightarrow B$ be a smooth fiber bundle, with $\dim(E) = 4$, $\dim(B) = 2$. Let J be an almost complex structure on E with respect to which the fibers are holomorphic. Let $u: S \rightarrow E$ be a J –holomorphic map, S connected, with $\pi \circ u$ not constant. Let $p \in S$ be a critical point of $\pi \circ u$, $q = \pi \circ u(p)$. Then there are neighborhoods $U \ni p$ and $V \ni q$, and C^2 coordinate charts $z: U \rightarrow \mathbb{C}$, $w: V \rightarrow \mathbb{C}$ such that $w \circ (\pi \circ u)(z) = z^k$, for some $k > 0$.*

Proof This follows immediately from [15, Theorem 7.1] applied to the intersection of u with the fiber of π over q . \square

Corollary 3.2 *Let $\pi: E \rightarrow B$ be a smooth fiber bundle, with $\dim(E) = 4$, $\dim(B) = 2$. Let J be an almost complex structure on E with respect to which the fibers are holomorphic. Let $u: S \rightarrow E$ be a J –holomorphic map, S connected, with $\pi \circ u$ not constant. Then the Riemann–Hurwitz formula applies to $\pi \circ u$. That is, if S is closed then*

$$\chi(S) = \chi(\pi \circ u(S)) - \sum_{p \in S} (e_{\pi \circ u}(p) - 1)$$

where $e_{\pi \circ u}(p)$ is the ramification index of p . If S has boundary and punctures then the same formula holds with Euler measure in place of Euler characteristic.

(The Euler measure of a surface S with boundary and punctures is $1/2\pi$ times the integral over S of the curvature of a metric on S for which ∂S is geodesic and the punctures of S are right angles. See Section 4 for further discussion of Euler measure.)

Lemma 3.3 Let $\pi: \Sigma \times [0, 1] \times \mathbb{R} \rightarrow \Sigma \times [0, 1]$ denote projection. Let u be a holomorphic curve in $\Sigma \times [0, 1] \times \mathbb{R}$ (with respect to some almost complex structure satisfying **(J1)**–**(J5)**). Let S' be a component of S on which u is not a trivial disk and $\pi_{\mathbb{D}} \circ u$ is not constant. Then there is a nonempty, open subset U of S' on which $\pi \circ u$ is injective and $\pi \circ u(U) \cap \pi \circ u(S \setminus U) = \emptyset$. Further, we can require that $u(U)$ be disjoint from $\overline{U_{3i}}$ and that $\pi_{\Sigma} \circ du$ and $\pi_{\mathbb{D}} \circ du$ be nonsingular on U .

(By a *trivial disk* we mean a component of S mapped diffeomorphically by u to $\{x\} \times [0, 1] \times \mathbb{R}$ for some $x \in \Sigma$.)

Proof Let x be such that $u|_{S'}$ is asymptotic to the Reeb chord $\{x\} \times [0, 1]$ at infinity. Let \overline{S} denote the surface obtained by blowing-up S at its punctures. As discussed earlier, we can extend u to a continuous map $\overline{S} \rightarrow \overline{W} = \Sigma \times [0, 1] \times [-1, 1]$. Let $\overline{\pi}: \overline{W} \rightarrow \Sigma \times [0, 1]$ denote projection. Let E denote the set of points $(x, s) \in \Sigma \times [0, 1]$ such that either $(\overline{\pi} \circ \overline{u})^{-1}(x, s)$ has cardinality larger than 1 or contains the image of a critical point of $\pi_{\Sigma} \circ u$ or $\pi_{\mathbb{D}} \circ u$. Then E is closed.

By the preceding corollary, there are only finitely many critical points of $\pi_{\Sigma} \circ u$ or $\pi_{\mathbb{D}} \circ u$. Further, “positivity of intersections” (eg, [15, Theorem 7.1]), applied to u and $\{x\} \times [0, 1] \times \mathbb{R}$, implies that there are only finitely many points in $\pi \circ u^{-1}(\{x\} \times [0, 1])$. So, there are only finitely many points in $E \cap \{x\} \times [0, 1]$.

However, $\{x\} \times [0, 1]$ is contained in the image of $\overline{\pi} \circ \overline{u}$. Choose $s \in [0, 1]$ such that $(x, s) \in \{x\} \times [0, 1] \setminus E$. Let V be an open neighborhood of (x, s) disjoint from E . Then $(\pi \circ u)^{-1}(V)$ has the desired properties. \square

To prove transversality we need to specify precisely the spaces under consideration. Fix $p > 2$, $k \geq 0$ ($k \in \mathbb{Z}$) and $d > 0$.

Definition 3.4 For a Riemannian manifold $(M, \partial M)$, a function $f: M \rightarrow \mathbb{R}$ lies in $L_k^p(M)$ if f has k weak derivatives in L^p . The L_k^p -norm of f is

$$\|f\|_{L_k^p} = \|f\|_{L^p} + \|f'\|_{L^p} + \cdots + \|f^{(k)}\|_{L^p}.$$

A function $f: M \rightarrow \mathbb{R}^n$ lies in L_k^p if each coordinate of f does, and its L_k^p -norm is the sum of the L_k^p -norms of its coordinate functions.

Fix a map $u: (S, \partial S) \rightarrow (W, C_{\alpha} \cup C_{\beta})$. Fix a Riemannian metric on S ; the particular choice is unimportant. Let $\{p_i^-\}$ denote the negative punctures of S and $\{p_i^+\}$ the positive punctures. Suppose u is asymptotic to $x_i^{\pm} \times [0, 1]$ at p_i^{\pm} . Identify a neighborhood U_i^- of each p_i^- with $[0, 1] \times (-\infty, 0]$ and a neighborhood U_i^+ of each p_i^+

with $[0, 1] \times [0, \infty)$. Let $(\sigma_i^\pm, \tau_i^\pm)$ denote the coordinates near p_i^\pm induced by this identification. Fix also a smooth embedding of Σ in \mathbb{R}^{N-2} for some N . This induces an embedding of $W = \Sigma \times [0, 1] \times \mathbb{R}$ in \mathbb{R}^N in an obvious way. For the following definition we identify W with its image in \mathbb{R}^N .

Definition 3.5 We say that u lies in $W_k^{p,d}((S, \partial S); (W, C_\alpha \cup C_\beta))$ if for some choice of constants $\{t_{0,i}^\pm\} \in \mathbb{R}$,

- the restriction of u to $S \setminus (U_1 \cup \dots \cup U_g \cup V_1 \cup \dots \cup V_g)$ lies in L_k^p (as a function to \mathbb{R}^N) and
- on each U_i^\pm the functions $e^{d|\tau_i^\pm|} (s \circ u(\sigma_i^\pm, \tau_i^\pm) - \sigma_i^\pm)$, $e^{d|\tau_i^\pm|} (t \circ u(\sigma_i^\pm, \tau_i^\pm) - \tau_i^\pm - t_{0,i}^\pm)$ and $e^{d|\tau_i^\pm|} (u(\sigma_i^\pm, \tau_i^\pm) - x_i^\pm)$ from $[0, \infty) \times \mathbb{R}$ or $(-\infty, 0]$ to \mathbb{R} lie in L_k^p .

For d small enough, all finite energy holomorphic curves (in the sense of [2, Section 5.3]) in $(W, C_\alpha \cup C_\beta)$ lie in $W_k^{p,d}$; see for instance [1, Chapter 3], particularly Propositions 3.5 and 3.6. Conversely, all maps in $W_k^{p,d}$ have finite energy.

Choose a homology class A of maps to W and a surface S . Let $\mathcal{X}_k^{p,d}$ denote the collection of maps $u \in W_\delta^{k,p}((S, \partial S); (W, C_\alpha \cup C_\beta))$ in class A .

Definition 3.6 Let E be a Riemannian vector bundle over S . Let f be a section of E . Then the $L_k^{p,d}$ -norm of f is

$$\|f\|_{L_k^{p,d}} = \|f\|_{S \setminus (U_1^- \cup \dots \cup U_g^+)} \|L_k^p + \sum_{i=1}^g \left(\|e^{d|\tau_i^+|} f|_{U_i^+} \|_{L_k^p} + \|e^{d|\tau_i^-|} f|_{U_i^-} \|_{L_k^p} \right)$$

Let $L_k^{p,d}(E)$ denote the Banach space of all sections of E with finite $L_k^{p,d}$ -norm.

Note that the tangent space at u to $\mathcal{X}_k^{p,d}$ is $\mathbb{R}^{2g} \oplus L_k^{p,d}(u^*TW, \partial)$ where $L_k^{p,d}(u^*TW, \partial)$ is the subspace of $L_k^{p,d}(u^*TW)$ of sections which lie in $u^*T(C_\alpha \cup C_\beta)$ over ∂S . The \mathbb{R}^{2g} factor corresponds to varying the $2g$ constants $t_{0,i}^\pm$ in Definition 3.5. Choosing $2g$ smooth vector fields v_i^\pm given by $\frac{\partial}{\partial t}$ on a neighborhood of p_i^\pm and zero near the other punctures p_j^\pm , we can include the \mathbb{R}^{2g} into $\Gamma(u^*TW)$ as $\text{Span}(\{v_i^\pm\})$.

Let \mathcal{J}^ℓ denote the space of C^ℓ almost complex structures on W which satisfy **(J1)**–**(J5)**. Let $\mathcal{J}^\ell(S)$ denote the space of C^ℓ almost complex structures on S . Let $\mathcal{M}^\ell = \{(u, j, J_s) \in \mathcal{X}_\delta^{k,p} \times \mathcal{J}^\ell(S) \times \mathcal{J}^\ell | \bar{\partial}_j J_s u = 0\}$.

Let $\text{End}(TS, j)$ denote the bundle whose fiber at $p \in S$ is the space of linear $Y: T_p S \rightarrow T_p S$ such that $Yj + jY = 0$. Then the tangent space at j to $\mathcal{J}^\ell(S)$ is the space of C^ℓ sections of $\text{End}(TS, j)$. Similarly, let $\text{End}(TW, J_s)$ denote the space of C^ℓ paths Y_s of linear maps $T\Sigma \rightarrow T\Sigma$ such that $Y_s J_s + J_s Y_s = 0$.

By convention, if we omit the superscripts k, ℓ , and p then we are referring to smooth objects.

By an *annoying curve* we mean a curve $u: S \rightarrow W$ such that there is a nonempty open subset of S on which $\pi_{\mathbb{D}} \circ u$ is constant.

Proposition 3.7 *For $\ell \geq 1$ the space \mathcal{M}^ℓ is a smooth Banach manifold away from annoying curves.*

Proof (This proof is a slight modification of [14, Proposition 3.4.1, page 34]. The reader is referred there for a less terse exposition.)

Let \mathcal{E}_{k-1}^p be the bundle over $\mathcal{X}_k^{p,d} \times \mathcal{J}^\ell(S) \times \mathcal{J}^\ell$ whose fiber over a point (u, j, J_s) is $L_{k-1}^{p,d}(\Lambda^{0,1} T^* S \otimes_{j, J_s} u^* T W)$.

We view $\bar{\partial}$ as a section of \mathcal{E}_{k-1}^p , and want to show that it is transverse to the zero section. At the zero section, the tangent space to \mathcal{E} splits as

$$T(\mathcal{X}_k^{p,d} \times \mathcal{J}^\ell(S) \times \mathcal{J}^\ell) \oplus L_{k-1}^{p,d}(\Lambda^{0,1} T^* S \otimes_{j, J_s} u^* T W).$$

Let

$$\begin{aligned} D\bar{\partial}(u, j, J_s): \mathbb{R}^{2g} \times L_k^{p,d}(u^* T W, \partial) \times C^\ell(\text{End}(TS, j)) \times C^\ell(\text{End}(TW, J_s)) \\ \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1} T^* S \otimes_{j, J_s} u^* T W) \end{aligned}$$

denote projection of the differential of $\bar{\partial}$ onto the vertical component of the tangent space to \mathcal{E} at a zero (u, j, J_s) of $\bar{\partial}$. We must show that $D\bar{\partial}$ is surjective.

The restriction of $D\bar{\partial}$ to any trivial disk is surjective by [10, Theorem 2]. So, for the rest of the proof we consider only the components of S which are not trivial disks. For these components we will, in fact, show that the restriction of $D\bar{\partial}$ to $0 \times L_k^{p,d}(u^* T W, \partial) \times C^\ell(\text{End}(TS, j)) \times C^\ell(\text{End}(TW, J_s))$ is surjective, and will focus on this restriction from now on.

The differential $D\bar{\partial}$ is given by

$$D\bar{\partial}(u, j, J_s)(\xi, Y, Y_s) = D_u \xi + \frac{1}{2} Y_s(u) \circ du \circ j + \frac{1}{2} J_s \circ du \circ Y$$

where $D_u \xi$ denotes the differential holding j and J_s fixed.

The operator $D\bar{\partial}(u, j, J_s)$ has closed range since D_u is Fredholm, and we only need to show that its range is dense. First, take $k = 1$. If the range is not dense then there exists $\eta \in L^{q,d}(\Lambda^{0,1}T^*S \otimes_{j,J_s} u^*TW)$ (where $1/p + 1/q = 1$) which annihilates the range of $D\bar{\partial}$. So, for any choice of (ξ, Y, Y_s) , we have

$$(1) \quad \int_S \langle \eta, D_u \xi \rangle = 0$$

$$(2) \quad \int_S \langle \eta, Y_s \circ du \circ j \rangle = 0$$

$$(3) \quad \int_S \langle \eta, J_s \circ du \circ Y \rangle = 0$$

The first equation says that η is a weak solution of $D_u^* \eta = 0$. So, by elliptic regularity, $\eta \in L^{r,d}_{\ell+1}$ for any $r > 0$. Further, it suffices to show that η vanishes on some open set to show that η vanishes identically.

Let U be as in the previous lemma and $z_0 \in U$. Choose coordinates (x, y) on S near z_0 with respect to which j is represented by the matrix $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Choose coordinates (x_1, y_1, x_2, y_2) near $u(z_0)$ preserving the splitting $TW = T\Sigma \oplus T\mathbb{D}$ and with respect to which the complex structure J_s on W has the form

$$\begin{pmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

The map η has the form $\begin{pmatrix} a & b \\ b & -a \\ c & d \\ d & -c \end{pmatrix}$. If $Y = \begin{pmatrix} \alpha & \beta \\ \beta & -\alpha \end{pmatrix}$ and $Y_s = \begin{pmatrix} \gamma & \delta \\ \delta & -\gamma \end{pmatrix}$. Let

$u_1 = x_1 \circ u, u_2 = x_2 \circ u$. Then, we have

$$J_s \circ du \circ Y = \begin{pmatrix} -\alpha \frac{\partial u_1}{\partial y} + \beta \frac{\partial u_1}{\partial x} & -\beta \frac{\partial u_1}{\partial y} - \alpha \frac{\partial u_1}{\partial x} \\ -\alpha \frac{\partial u_1}{\partial x} - \beta \frac{\partial u_1}{\partial y} & -\beta \frac{\partial u_1}{\partial x} + \alpha \frac{\partial u_1}{\partial y} \\ -\alpha \frac{\partial u_2}{\partial y} + \beta \frac{\partial u_2}{\partial x} & -\beta \frac{\partial u_2}{\partial y} - \alpha \frac{\partial u_2}{\partial x} \\ -\alpha \frac{\partial u_2}{\partial x} - \beta \frac{\partial u_2}{\partial y} & -\beta \frac{\partial u_2}{\partial x} + \alpha \frac{\partial u_2}{\partial y} \end{pmatrix}$$

and

$$Y_s \circ du \circ j = \begin{pmatrix} -\gamma \frac{\partial u_1}{\partial y} - \delta \frac{\partial u_1}{\partial x} & \gamma \frac{\partial u_1}{\partial x} - \delta \frac{\partial u_1}{\partial y} \\ -\delta \frac{\partial u_1}{\partial y} + \gamma \frac{\partial u_1}{\partial x} & \delta \frac{\partial u_1}{\partial x} + \gamma \frac{\partial u_1}{\partial y} \\ 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

By choosing γ and δ appropriately we can force $a = b = 0$ near z_0 . (This uses the injectivity established in the lemma and the nonvanishing of $\pi_\Sigma \circ du(z_0)$.) Then, choosing α and β appropriately we can force $c = d = 0$ near z_0 . (This uses the nonvanishing of $\pi_{\mathbb{D}} \circ du(z_0)$.) This establishes the surjectivity of $D\bar{\partial}$ and hence the $k = 1$ case.

For general k , suppose $\eta \in L_{k-1}^{p,d}(\Lambda^{0,1}T^*S \otimes_{j,J_s} u^*TW)$. From the $k = 1$ case, choose a triple $\xi \in L_1^p(u^*TW)$, $Y \in C^\ell(\text{End}(TS, j))$, and $Y_s \in C^\ell(\text{End}(TW, J_s))$ such that $D\bar{\partial}(u, j, J_s)(\xi, Y, Y_s) = \eta$. Then, elliptic regularity implies $\xi \in L_k^{p,d}$, so $D\bar{\partial}$ is surjective. Since $D\bar{\partial}$ is Fredholm, it follows from the infinite-dimensional implicit function theorem that \mathcal{M}^ℓ is a Banach manifold. \square

Proposition 3.8 *For a dense set J_{reg} of C^∞ paths of smooth complex structures on Σ , the moduli space of holomorphic curves satisfying (M1), (M2), (M4) and (M5), and without multiply covered components, is a smooth manifold.*

Proof Observing that (M2) implies the absence of annoying curve components, this follows easily from the previous result. The set J_{reg} is exactly the set of regular values for the projection of \mathcal{M} onto \mathcal{J} . For \mathcal{J}^ℓ it is immediate from Smale’s infinite-dimensional version of Sard’s theorem that J_{reg}^ℓ is dense. For the C^∞ statement a short approximation argument is required. We refer the reader to [14, page 36]; our case is just the same as theirs. \square

Remark Note that (M6) implies that u has no multiply covered components.

We will often say a complex structure J achieves transversality to mean $J \in J_{reg}$.

There is a second way that we can sometimes achieve transversality, which is more convenient for computations: by keeping the complex structure on W split and perturbing the α and β circles. Specifically we have:

Proposition 3.9 *Suppose that a homology class $A \in \pi_2(\vec{x}, \vec{y})$ with $\text{ind}(A) = 1$ is such that any $j_\Sigma \times j_{\mathbb{D}}$ -holomorphic curve $u: S \rightarrow W$ in the homology class A must have $\pi_\Sigma \circ u|_{\partial S}$ somewhere injective. Then for a generic perturbation of the α - and β -circles, for any u in the homology class A the linearization $D\bar{\partial}$, computed with respect to the complex structure $j_\Sigma \times j_{\mathbb{D}}$ on W , is surjective.*

The proof of this proposition is analogous to the argument in [16]. It is also a corollary of [21, Proposition 3.9], so we omit the proof.

We shall refer to the condition in the preceding proposition as *boundary injectivity*. One obvious time when boundary injectivity holds is the following:

Lemma 3.10 *Suppose the homology class $A \in \pi_2(\vec{x}, \vec{y})$ is represented by a domain $D = \sum_i n_i D_i$ such that for some i and j , $n_i = 1$, $n_j = 0$, and $\partial D_i \cap \partial D_j \neq \emptyset$. Then A satisfies the boundary injectivity hypothesis.*

For computing the homologies defined in Section 8, if the boundary injectivity criterion is met by every domain with index 1 it will suffice to take a generic perturbation of the boundary conditions and the split complex structure $j_\Sigma \times j_\mathbb{D}$ rather than a generic path J_s of complex structures. (The only time this is relevant in this paper is Section 11, but in practice it is necessary for most direct computations.)

4 Index

In this section we compute the index of the linearized $\bar{\partial}$ -operator $D\bar{\partial}$ at a holomorphic map $u: (S, j) \rightarrow (W, J)$. We start by reducing to a result discussed in [1] via a doubling argument similar to the one found in [10]. We then reinterpret this index several times, obtaining the Chern class formula for the index of periodic domains (Corollary 4.12, which is [21, Theorem 4.9]), J Rasmussen's formula ([25, Theorem 9.1], proved here in Proposition 4.8) and a combinatorial formula for the index near an embedded curve (Corollary 4.3) and, consequently, the Maslov index in traditional Heegaard Floer homology (Corollary 4.10).

4.1 First formulas for the index

We may assume that J is split, since deformations of J will not change the index. Also, we assume that the α and β curves meet in right angles.

Let a_1, \dots, a_g be the components of the boundary of S (in the complement of the punctures) mapped to α -cylinders and b_1, \dots, b_g the components mapped to β -cylinders. We define the quadruple of S , denoted $4 \times S$, by gluing four copies of S , denoted S_1, S_2, S_3 , and S_4 , as follows. Glue each a_i in S_1 to a_i in S_2 and each a_i in S_3 to a_i in S_4 . Similarly, glue each b_i in S_1 to b_i in S_3 and each b_i in S_2 to b_i in S_4 . Define a complex structure on $4 \times S$ by taking the complex structure j on S_1 and S_4 and its conjugate \bar{j} on S_2 and S_3 . Notice that these complex structures glue together correctly.

The complex vector bundle (u^*TW, u^*J) extends to a vector bundle over $4 \times S$, which we denote $(\overline{u^*T\overline{W}}, \overline{J})$, in an obvious way, and $D\overline{\partial}$ extends to an operator $4 \times D\overline{\partial}$ on the sections of this vector bundle. We restrict $4 \times D\overline{\partial}$ to the space of sections which approach zero near each puncture, as we require fixed asymptotics.

By [1, Corollary 5.4, page 53], the index of $4 \times D\overline{\partial}$ is

$$(4) \quad -\chi(4 \times S) + 2c_1(A).$$

Here, $c_1(A)$ is defined as follows. Choose a small disk near each point in $\alpha \cap \beta$. Trivialize $(T\Sigma, J)$ over these disks. This gives a trivialization of u^*TW in a neighborhood of the punctures in $4 \times S$ which extends to a trivialization of $\overline{u^*T\overline{W}}$ over the surface $\overline{4 \times S}$ obtained by filling in the punctures. Then, $c_1(A)$ is the pairing of the first Chern class of $\overline{u^*T\overline{W}}$ with the fundamental class of $\overline{4 \times S}$. Note that since $T([0, 1] \times \mathbb{R})$ is trivial, to compute c_1 we need only look at the Σ factor. Also, it is necessary to observe that F. Bourgeois’s calculations in [1, Section 5] are all done in the pullback bundle, so the fact that our index problem does not correspond to a genuine map is a nonissue.

We convert Formula (4) into one not involving the quadruple of S . First we compute that $\chi(4 \times S) = 4\chi(S) - 4g$. Indeed, after doubling along the α -arcs the Euler characteristic is $2\chi(S) - g$. Doubling again we obtain $\chi(4 \times S) = 2(2\chi(S) - g) - 2g$. (The last summand of $-2g$ comes from the $2g$ punctures in $4 \times S$.)


Second, $c_1(A)$ can be computed from Maslov-type indices as follows. Choose the trivializations of TW over the neighborhoods \mathcal{V} of $\alpha \cap \beta$ above so that for $p \in \alpha \cap \beta$, $T_p\beta = \mathbb{R} \subset \mathbb{C}$ and $T_p\alpha = i\mathbb{R} \subset \mathbb{C}$. Trivialize $(\pi_\Sigma \circ u)^*T\Sigma$ over S so that this trivialization agrees with the specified trivialization of TW over the neighborhoods \mathcal{V} . Then, each boundary arc a_i or b_i gives a loop of lines in \mathbb{C} , and so has a well-defined Maslov index $\mu(a_i)$ or $\mu(b_i)$. It is not hard to see that $c_1(A) = 2(\sum_{i=1}^g \mu(a_i) - \mu(b_i))$. This is independent of the choice of trivialization subject to the specified criteria.

Fix a map $u: (S, \partial S) \rightarrow (W, C_\alpha \cup C_\beta)$. An argument almost exactly like the one in [10] shows that

$$(5) \quad \text{ind}(D\overline{\partial}) = \frac{1}{4}\text{ind}(4 \times D\overline{\partial}) = g - \chi(S) + \sum_{i=1}^g \mu(a_i) - \sum_{i=1}^g \mu(b_i).$$

The factor of $1/4$ comes from the “matching conditions” on the boundary of S .

We again reinterpret the Maslov indices. Given a domain D , we define the Euler measure of D as follows. Suppose first that D is a surface with boundary and corners. Choose a metric on D such that ∂D is geodesic and such that the corners of D are

Figure 2: S is the shaded region.

right angles. Then the Euler measure $e(D)$ is defined to be $\frac{1}{2\pi}$ times the integral over D of the curvature of the metric. (This is normalized so that the Euler measure of a sphere is 2, agreeing with its Euler characteristic.) From this definition it is clear that the Euler measure is additive under disjoint unions and gluing of components along boundaries, and so the definition extends naturally to domains (linear combinations of regions in Σ).

It follows from the Gauss–Bonnet theorem that the Euler measure of a surface S with k acute right–angled corners (see Figure 2) and ℓ obtuse right–angled corners is $\chi(S) - k/4 + \ell/4$. As with the previous formulation of Euler measure, this formula is additive, so the Euler measure of a domain $D = \sum_i D_i$ is $e(D) = \sum_i e(D_i)$.

From the Gauss–Bonnet theorem, we also know that if we endow D with a flat metric such that all corners are right angles then the Euler measure $e(D)$ is $\frac{1}{2\pi}$ times the geodesic curvature of ∂D . It is then clear that for D the domain corresponding to u , $\sum_{i=1}^g \mu(a_i) - \mu(b_i) = 2e(D)$. So, we can recast the index formula as

$$(6) \quad \text{ind}(D\bar{\partial}) = g - \chi(S) + 2e(D).$$

4.2 Determining S from A

Note that the formulas for the index derived so far depend not only on the homology class but also on the topological type of the source. This is as it should be. However, as we will show presently, for *embedded* holomorphic curves the Euler characteristic of the source is determined by the homology class. (Actually, we prove this more generally for any curves satisfying certain hypotheses described in Lemma 4.1 below, not just holomorphic ones.) In fact, we can give an explicit formula for the Euler characteristic, allowing us to give a combinatorial formula for the index. We will see




Figure 3: A degenerate corner

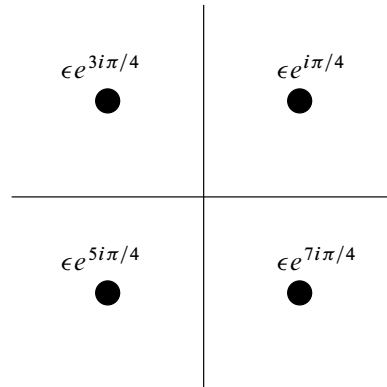


Figure 4: $n_p(A) = \frac{1}{4} (n_{\epsilon e^{i\pi/4}}(A) + n_{\epsilon e^{3i\pi/4}}(A) + n_{\epsilon e^{5i\pi/4}}(A) + n_{\epsilon e^{7i\pi/4}}(A))$

in Subsection 4.3 that this formula calculates the Maslov index in the setup of [21] as well. Before proving this claim we introduce some more terminology and notation.

Given an intersection point \vec{x} we call each $x_i \in \vec{x}$ a *corner* of \vec{x} . Following Rasmussen, we define a corner x_i of \vec{x} to be *degenerate* for a homology class $A \in \pi_2(\vec{x}, \vec{y})$ if $x_i = y_j$ for some $y_j \in \vec{y}$. This definition will be convenient presently.

Let $p \in \alpha_i \cap \beta_j$. For a homology class $A \in \pi_2(\vec{x}, \vec{y})$, define $n_p(A)$ to be the average of the coefficients of A of the four cells with corners at p . More precisely, choose coordinates identifying a neighborhood of p in Σ with the unit disk in \mathbb{C} , α_i with the real axis, and β_j with the imaginary axis. Then $n_p(A) = \frac{1}{4} (n_{\epsilon e^{i\pi/4}}(A) + n_{\epsilon e^{3i\pi/4}}(A) + n_{\epsilon e^{5i\pi/4}}(A) + n_{\epsilon e^{7i\pi/4}}(A))$ for some $\epsilon < 1$. See Figure 4. Define $n_{\vec{x}}(A)$ to be $\sum_{x_i \in \vec{x}} n_{x_i}(A)$, and $n_{\vec{y}}(A) = \sum_{y_i \in \vec{y}} n_{y_i}(A)$. (See [20, page 1202].)

We need a lemma about representability of homology classes with positive coefficients:

Lemma 4.1 *Suppose $A \in \pi_2(\vec{x}, \vec{y})$ is a positive homology class. Then there is a Riemann surface with boundary and corners \overline{S} and smooth map $u: S \rightarrow W$ (where S denotes the complement in \overline{S} of the corners of \overline{S}) in the homology class A such that:*

- (1) $u^{-1}(C_\alpha \cup C_\beta) = \partial S$.
- (2) For each i , $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ and $u^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ each consists of one arc in ∂S .
- (3) The map u is J -holomorphic in a neighborhood of $(\pi_\Sigma \circ u)^{-1}(\alpha \cup \beta)$ for some J satisfying **(J1)**–**(J5)** (in fact, for $j_\Sigma \times j_{\mathbb{D}}$).
- (4) For each component of S , either
 - The component is a disk with two boundary punctures and the map is a diffeomorphism to $\{x_i\} \times [0, 1] \times \mathbb{R}$ for some $x_i \in \alpha \cap \beta$ (such a component is a degenerate disk) or
 - The map $\pi_\Sigma \circ u$ extends to a branched covering map $\overline{\pi_\Sigma \circ u}$, none of whose branch points map to points in $\alpha \cap \beta$.
- (5) All the corners of S are acute (see Figure 2).
- (6) The map u is an embedding.

(Note that it follows from the conditions in the lemma that the map to Σ is orientation-preserving. Also, observe that a generic holomorphic representative of the homology class satisfies all of the properties of the lemma.)

Proof (For a similar construction of a u with slightly different properties, see Rasmussen, [25, Lemma 9.3]. His construction is slightly more subtle than we need. Another inspiring construction can be found in [21, Lemma 2.17].)

Let D_1, \dots, D_N denote the closures of the components of $\Sigma \setminus (\alpha \cup \beta)$, enumerated so that $\mathfrak{z}_i \in D_i$. Form a surface \overline{S}_0 by gluing together $n_{\mathfrak{z}_i}(A)$ copies of D_i ($i = 1, \dots, N$) pairwise along common boundaries maximally. There is then an obvious orientation-preserving map $p_{\Sigma,0}: \overline{S}_0 \rightarrow \Sigma$ which covers \mathfrak{z}_i $n_{\mathfrak{z}_i}(A)$ times.

It is possible to perform the specified gluing so that the only corners of \overline{S}_0 correspond one-to-one with non-degenerate corners of the domain. This is not automatic; see Figure 5. One way to achieve this is to first glue maximally along α -arcs. Then glue along the β -arcs as much as possible without introducing any obtuse corners. After




Figure 5: Two possible gluings \bar{S}_0 for a specified domain. The latter leads to extra corners of \bar{S}_0 .

doing so, any remaining corners must correspond to non-degenerate corners of the domain.

The surface \bar{S}_0 lacks corners at degenerate corners of the domain. Let \mathbb{D}_2^2 denote a disk with two punctures on the boundary. Let \bar{S}_1 denote the disjoint union of \bar{S}_0 with a copy of \mathbb{D}_2^2 for each degenerate corner of the domain. Extend $p_{\Sigma,0}$ to a map $p_{\Sigma,1}$ from \mathbb{D}_2^2 by mapping one copy of \mathbb{D}_2^2 to each degenerate corner of the domain.

Now, S_0 inherits a complex structure from Σ . Extend this complex structure arbitrarily over the new disks in S_1 to obtain a complex structure on S_1 . It is easy to choose a map $p_{\mathbb{D},1}: S_1 \rightarrow [0, 1] \times \mathbb{R}$ such that the map $(p_{\Sigma,1} \times p_{\mathbb{D},1}): S_1 \rightarrow W$ satisfies all of the properties specified in the statement of the lemma except perhaps numbers (2) and (6). Perturbing $p_{\mathbb{D},1}$ we may assume that $(p_{\Sigma,1} \times p_{\mathbb{D},1})$ is an embedding except for a collection of transverse double points.

For each α_i (respectively β_j), there will be exactly one arc in $\partial\Sigma_1$ mapped by $p_{\Sigma,1}$ to α_i (respectively β_j), and possibly some circles in $\partial\Sigma_1$ mapped by $p_{\Sigma,1}$ to α_i (respectively β_j). The map $p_{\mathbb{D},1}$ must be constant near each closed component of

∂S_1 , and the image of the arc under $(p_{\Sigma,1} \times p_{\mathbb{D},1})$ must intersect the image of each closed component of ∂S_1 mapped to α_i (respectively β_j) exactly once.

Modifying S_1 and $p_{\Sigma,1} \times p_{\mathbb{D},1}$ near the double points of $p_{\Sigma,1} \times p_{\mathbb{D},1}$ we can obtain a new map $u: S \rightarrow W$ satisfying all of the stated properties: in the process of deforming away the double points, we necessarily achieve property (2) as well. \square

Proposition 4.2 *Let $u: S \rightarrow W$ be a map satisfying the conditions enumerated in the previous lemma, representing a homology class A . Then the Euler characteristic $\chi(S)$ is given by*

$$(7) \quad \chi(S) = g - n_{\bar{x}}(A) - n_{\bar{y}}(A) + e(A).$$

Proof Applying the Riemann–Hurwitz formula to $\pi_{\Sigma} \circ u$, we only need to calculate the degree of branching of $\pi_{\Sigma} \circ u$.

To calculate the number of branch points of $\pi_{\Sigma} \circ u$ we reinterpret this number as a self–intersection number. We will assume all branch points of $\pi_{\Sigma} \circ u$ have order 2; we can clearly arrange this. Observe that since S has no obtuse corners, by the Riemann–Hurwitz formula,

$$\begin{aligned} \chi(S) = e(S) + g/2 = e(A) - (\text{number of branch points}) \\ + \frac{1}{2}(\text{number of trivial disks}) + g/2. \end{aligned}$$

(Branch points on ∂S should each be counted as half of a branch point.)

Assume for the time being that u contains no trivial disks, and in fact has no degenerate corners.

Notice that the number of branch points of $\pi_{\Sigma} \circ u$ is equal to the number of times the vector field $\frac{\partial}{\partial t}$ is tangent to u . (Tangencies on ∂S should each be counted as half of a tangency.) Let u' denote the curve obtained from u by translating a distance R in the \mathbb{R} –direction. Then, for small R , the number of branch points of $\pi_{\Sigma} \circ u$ is equal to the intersection number of u and u' . (Intersections on ∂S should each be counted as half of an intersection.)

This intersection number is invariant under isotopies of u' such that all intersection points of u and u' remain in a compact subset of W . (The only thing to check is that when an intersection point in the interior of W hits the boundary it gives rise to a pair of intersection points on the boundary. It is not hard to check this using a doubling argument in a neighborhood of the boundary.) We will calculate the intersection number by translating u' far in the \mathbb{R} –direction of W .

Translate u' by some $R \gg 0$ in the \mathbb{R} -factor of W . All intersection points between u and u' stay in a compact subset of W , so the intersection number $\#u \cap u'$ is unchanged.

We can modify u' so that near each negative puncture (corresponding to some x_i) u' agrees with the trivial disk $x_i \times [0, 1] \times \mathbb{R}$. Further, we can do this modification so that all intersection points between u and u' stay within some compact subset of W . (This follows from the simple asymptotic behavior of u' near $-\infty$.)

Similarly, we can modify u so that near the positive punctures of S , u agrees with the trivial disks $\{y_i\} \times [0, 1] \times \mathbb{R}$ ensuring in the process that all intersection points between u and u' stay within some compact subset of W .

Finally, for R large enough, we can assume that after the two modifications every intersection point between u and u' corresponds to an intersection point between u and $\{x_i\} \times [0, 1] \times \mathbb{R}$ or between $\{y_j\} \times [0, 1] \times \mathbb{R}$ and u' .

Now, for each corner $c_{k,\ell}$ of each component E_k of $S \setminus (\pi_\Sigma \circ u)^{-1}(\alpha \cup \beta)$, one of the following four phenomena occurs:

- (1) The corner $c_{k,\ell}$ is mapped by $\pi_\Sigma \circ u$ somewhere other than x_i . That is, $\lim_{p \rightarrow c_{k,\ell}} \pi_\Sigma \circ u(p) \neq x_i$.
- (2) The corner $c_{k,\ell}$ is mapped by $\pi_\Sigma \circ u$ to x_i , but at $-\infty$. That is, $\lim_{p \rightarrow c_{k,\ell}} \pi_\Sigma \circ u(p) = x_i$ but $\lim_{p \rightarrow c_{k,\ell}} \pi_{\mathbb{R}} \circ u(p) = -\infty$.
- (3) The corner $c_{k,\ell}$ is mapped by $\pi_\Sigma \circ u$ to x_i , and is mapped by u to the boundary of W . That is, $\pi_\Sigma \circ u(c_{k,\ell}) = x_i$ and $\pi_{\mathbb{D}} \circ u(c_{k,\ell}) \in \partial[0, 1] \times \mathbb{R}$.
- (4) The corner $c_{k,\ell}$ is mapped by $\pi_\Sigma \circ u$ to x_i , and is mapped by u to the interior of W . That is, $\pi_\Sigma \circ u(c_{k,\ell}) = x_i$ and $\pi_{\mathbb{D}} \circ u(c_{k,\ell}) \in (0, 1) \times \mathbb{R}$.

(Compare Figure 6.)

In the first two cases, E_k does not contribute to $\#(u \cap x_i \times [0, 1] \times \mathbb{R})$. In the last two, E_k contributes $1/4$ to the intersection number. Notice that the second case occurs exactly once, since S has no obtuse corners. There are a total of $4n_{\bar{x}}$ corners satisfying one of conditions (2)–(4) for some x_i . Exactly g of them satisfy condition (2). So, $\#(u \cap \{x_i\} \times [0, 1] \times \mathbb{R}) = n_{\bar{x}} - g/4$.

A similar analysis works for the intersection points between u' and $\{y_j\} \times [0, 1] \times \mathbb{R}$. So, it follows that the intersection number between u and u' is

$$\#(u \cap u') = n_{\bar{x}}(A) + n_{\bar{y}}(A) - g/2.$$

It follows that $\chi(S) = e(A) - n_{\bar{x}}(A) - n_{\bar{y}}(A) + g$.

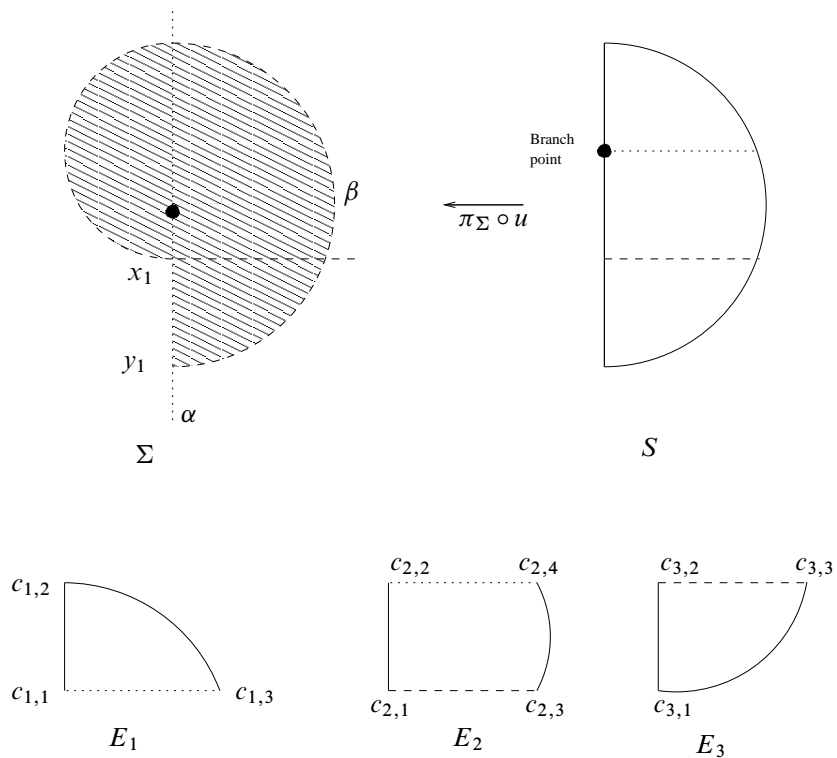


Figure 6: With respect to x_1 , corner $c_{1,2}$ has type (2), corners $c_{2,1}$ and $c_{3,2}$ have type (3), and all others have type (1).

In the proof so far we assumed that there were no trivial disks. Suppose u contains trivial disks corresponding to the intersection points x_{i_1}, \dots, x_{i_k} . Since we are considering only embedded curves, $n_{x_{i_j}}(A) = 0$ for $j = 1, \dots, k$. By the argument above, after ignoring the trivial disks, we find $\#(u \cap u') = n_{\bar{x}}(A) + n_{\bar{y}}(A) - g/2 + k/2$. So, we have the same formula for $\chi(S)$ as before.

Finally, we deal with degenerate corners which are not trivial disks. Since we are assuming S has only acute corners, and acute degenerate corners have exactly one shared boundary component under $\pi_\Sigma \circ u$, it is easy to see that after translating in the \mathbb{R} -direction, there will be one intersection point along the boundary near the puncture, so the extra corner mapped to $\pm\infty$ contributes $\frac{1}{2}$ to the intersection number, as one would expect from our formula. This concludes the proof. \square

Note that if $u: S \rightarrow W$ is an embedded holomorphic curve (with respect to any complex structure J on W satisfying **(J1)**–**(J5)**) then, after slitting S near any obtuse corners

and perturbing u slightly, $u: S \rightarrow W$ satisfies the conditions of Lemma 4.1. It follows that Proposition 4.2 calculates the Euler characteristic of S .

Corollary 4.3 For A a positive homology class and $u: S \rightarrow W$ a representative for A satisfying the conditions of Lemma 4.1, the index of the $D\bar{\partial}$ operator near u is given by

$$(8) \quad \text{ind}(D\bar{\partial}) = e(A) + n_{\vec{x}}(A) + n_{\vec{y}}(A)$$

Proof This is immediate from formula (6) and Proposition 4.2. \square

Definition 4.4 Given a positive homology class A define the index $\text{ind}(A)$ of A to be the index of the $D\bar{\partial}$ operator near any curve satisfying the conditions of Lemma 4.1.

Corollary 4.5 If A and $A + k[\Sigma]$ are both positive then $\text{ind}(A + k[\Sigma]) = \text{ind}(A) + 2k$.

Proof By Corollary 4.3,

$$\begin{aligned} \text{ind}(A + k[\Sigma]) &= e(A + k[\Sigma]) + n_{\vec{x}}(A + k[\Sigma]) + n_{\vec{y}}(A + k[\Sigma]) \\ &= e(A) + (2 - 2g)k + n_{\vec{x}}(A) + gk + n_{\vec{y}}(A) + gk \\ &= e(A) + n_{\vec{x}}(A) + n_{\vec{y}}(A) + 2k \\ &= \text{ind}(A) + 2k. \end{aligned}$$

\square

Definition 4.6 For any homology class A define the index $\text{ind}(A)$ of A to be $\text{ind}(A + k[\Sigma]) - 2k$ where k is chosen large enough that $A + k[\Sigma]$ is positive.

Corollary 4.7 Suppose that $A \in \pi_2(\vec{x}, \vec{y})$ and $A' \in \pi_2(\vec{y}, \vec{z})$. Then $\text{ind}(A + A') = \text{ind}(A) + \text{ind}(A')$.

Proof We may clearly assume that A and A' are both positive. Let $u: S \rightarrow W$ and $u': S' \rightarrow W$ be maps satisfying the conditions Lemma 4.1 representing A and A' respectively. Then we can glue u and u' to a map $u \natural u': S \natural S' \rightarrow W$ representing $A + A'$. It follows from general gluing results for the index that $\text{ind}(D\bar{\partial})(u \natural u') = \text{ind}(D\bar{\partial})(u) + \text{ind}(D\bar{\partial})(u')$. (Alternately, it follows from the additivity of Formula (6) under gluing.) \square

Remark Formula (8) was suggested to me by Z Szabó. Specifically, he suggested that it seemed the Maslov index in [21] can be calculated by this formula. In particular, in the special case when $A \in \pi_2(\vec{x}, \vec{x})$, Ozsváth and Szabó proved ([21, Theorem 4.9] and [20, Proposition 7.5]) that Formula (8) does compute the Maslov index. Note that it is not even clear *a priori* that Formula (8) is additive. In fact, I do not know a more direct proof than the one we used to obtain Corollary 4.7.

4.3 Comparison with classical Heegaard Floer homology

In this subsection we assume familiarity with [21].

By considering domains, for example, there is a natural identification of our $\pi_2(\vec{x}, \vec{y})$ with $\pi_2(\vec{x}, \vec{y})$ as defined in [21, Section 2.4]. For $A \in \pi_2(\vec{x}, \vec{y})$, let $\mu(A)$ denote the Maslov index of A , viewed as a homotopy class of maps disks in $(\text{Sym}^g(\Sigma), T_\alpha \cup T_\beta)$. The goal of this subsection is to prove the following

Proposition 4.8 *For $A \in \pi_2(\vec{x}, \vec{y})$ we have $\text{ind}(A) = \mu(A)$.*

Proof It is possible to give a direct proof (see [25, proof of Theorem 9.1]), but instead of doing so we will show our formula agrees with the one given by Rasmussen in [25, Theorem 9.1]. He proves that at a disk $\phi: (\mathbb{D}, \partial\mathbb{D}) \rightarrow (\text{Sym}^g(\Sigma), T_\alpha \cup T_\beta)$,

$$(9) \quad \mu(\phi) = \Delta \cdot \phi + 2e(\phi).$$

Here, e is the Euler measure defined in Subsection 4.1 and $\Delta \cdot \phi$ is the algebraic intersection number of ϕ with the diagonal in $\text{Sym}^g(\Sigma)$. (The diagonal is an algebraic subvariety of $\text{Sym}^g(\Sigma)$ of real codimension 2 so the intersection number is well-defined.)

To compare his result with ours, we need a slight strengthening of Lemma 4.1:

Lemma 4.9 *Suppose A is a positive homology class. Then we can represent $A + [\Sigma]$ by a map $u: S \rightarrow W$ satisfying all the conditions of Lemma 4.1 and such that, additionally:*

- *The map $\pi_{\mathbb{D}} \circ u$ is a g -fold branched covering map with all its branch points of order 2.*
- *The map u is holomorphic near the preimages of the branch points of $\pi_{\mathbb{D}} \circ u$.*

Proof Construct a map $u_1: S_1 \rightarrow W$ representing A as in Lemma 4.1. We would like to say that we can then choose a branched cover $p_{\mathbb{D},1}: S_1 \rightarrow [0, 1] \times \mathbb{R}$ (mapping arcs on the boundary appropriately). This may not, however, be the case: suppose, for instance, that $g = 2$ and S_1 were the disjoint union of a disk and a surface of genus one with one boundary component.

However, note that $[\Sigma] \in \pi_2(\vec{y}, \vec{y})$ can be represented by a map with connected source. Specifically, let S_Σ be obtained by making small slits in Σ along α_i and β_i starting at $y_i \in \vec{y}$ for $i = 1, \dots, g$. There is an obvious map $S_\Sigma \rightarrow \Sigma$.

Gluing the negative corners of S_Σ to the positive corners of S_1 we obtain a connected surface S_2 and map $p_{\Sigma,2}: S_2 \rightarrow \Sigma$. Since S_2 is connected it is possible to choose a

branched covering map $p_{\mathbb{D},2}: S_2 \rightarrow \mathbb{D}$ with appropriate boundary behavior. Perturbing this map we can assume all of its branch points have order 2. Finally, deforming away the double points of $p_{\Sigma,2} \times p_{\mathbb{D},2}$ and perturbing it to be holomorphic in appropriate places, we obtain an embedding satisfying the specified conditions. \square

Now, fix a positive homology class A in $\pi_2(\vec{x}, \vec{y})$ and a map u representing A as in the previous lemma. The map u induces a map $\phi: \mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ as follows. For $a \in \mathbb{D}$ let $(\pi_{\mathbb{D}} \circ u)^{-1}(a) = \{a_1, \dots, a_g\}$. Then define $\phi(a) = \{\pi_{\Sigma} \circ u(a_1), \dots, \pi_{\Sigma} \circ u(a_g)\}$.

There is a one-to-one correspondence between order 2 branch points of $\pi_{\mathbb{D}} \circ u$ and transverse intersections of ϕ with the top-dimensional stratum of the diagonal. By the Riemann–Hurwitz formula, $\chi(S) = g\chi(\mathbb{D}^2) - \phi \cdot \Delta = g - \phi \cdot \Delta$, so $\text{ind}(D\bar{\partial})(u) = g - \chi(S) + 2e(A + [\Sigma]) = \phi \cdot \Delta + 2e(A + [\Sigma])$.

This is exactly Rasmussen’s formula for the Maslov index. Thus, we have shown that $\text{ind}(A + [\Sigma]) = \mu(A + [\Sigma])$ for A positive. But both ind and μ are additive, and have $\text{ind}([\Sigma]) = \mu([\Sigma]) = 2$. Thus, it follows that $\text{ind}(A) = \mu(A)$ for all A .

Corollary 4.10 *In Heegaard Floer homology, the Maslov index of a domain D is given by*

$$\mu(D) = n_{\vec{x}}(D) + n_{\vec{y}}(D) + e(D).$$

Remark There are no assumptions on the domain.

4.4 Index for $A \in \pi_2(\vec{x}, \vec{x})$

The following result, which we will use below, is proved by Ozsváth and Szabó in [20, Proposition 7.5] by direct geometrical argument.

Lemma 4.11 *If $\vec{x} \in \mathfrak{s}$ and $A \in \hat{\pi}_2(\vec{x}, \vec{x})$ then $\langle c_1(\mathfrak{s}), A \rangle = e(A) + 2n_{\vec{x}}(A)$. Here $\langle \cdot, \cdot \rangle$ denotes the natural pairing between homology and cohomology, $c_1(\mathfrak{s})$ the first Chern class of \mathfrak{s} , and A is viewed as an element of $H_2(Y)$.*

The following is completely analogous to [21, Theorem 4.9].

Corollary 4.12 *Let P be a homology class in $\pi_2(\vec{x}, \vec{x})$. Then*

$$\text{ind}(P) = \langle c_1(\mathfrak{s}), P \rangle + 2n_{\vec{x}}(P).$$

5 Admissibility criteria

In order to define the differential in our chain complexes it will be important that for any intersection points \vec{x} and \vec{y} , only finitely many homology classes $A \in \pi_2(\vec{x}, \vec{y})$ with $\text{ind}(A) = 1$ support holomorphic curves. For a rational homology sphere, this is automatic: there are only finitely many homology classes in $\pi_2(\vec{x}, \vec{y})$. In general, following [21], we use special Heegaard diagrams and “positivity of domains” to ensure that only finitely many homology classes support holomorphic curves. Our definitions are the same as theirs. For the reader’s amusement, we provide slightly different proofs of two of the fundamental lemmas about admissibility.

Definition 5.1 (Compare [21, Definition 4.10]) The pointed Heegaard diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ is called *weakly admissible* for the $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} if every nontrivial periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 0$ has both positive and negative coefficients.

Definition 5.2 (Compare [21, Definition 4.10]) The pointed Heegaard diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ is called *strongly admissible* for the $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} if every nontrivial periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 2n > 0$ has $n_{\mathfrak{z}_i}(P) > n$ for some \mathfrak{z}_i .

Remark Notice that for any $\text{Spin}^{\mathbb{C}}$ -structure, $c_1(\mathfrak{s})$ is an even cohomology class: $c_1(\mathfrak{s})$ is the first Chern class of v^{\perp} for some nonvanishing vector field v . Then $c_1(\mathfrak{s}) \equiv w_2(v^{\perp}) \pmod{2}$. Since TM is trivial and the line field determined by v is obviously trivial, $1 = (1 + w_1(v^{\perp}) + w_2(v^{\perp}))$, so $w_2(v^{\perp}) = 0$.

We now need two kinds of result. The first is the finiteness mentioned just above in the case of weak / strong admissibility. The second is that the admissibility criteria can be achieved, and that any two admissible Heegaard diagrams can be connected by a sequence of Heegaard moves through admissible diagrams.

First, a few simple observations. A $\text{Spin}^{\mathbb{C}}$ -structure is called “torsion” if $c_1(\mathfrak{s})$ is a torsion homology class. For a torsion $\text{Spin}^{\mathbb{C}}$ -structure, $\langle c_1(\mathfrak{s}), P \rangle = 0$ for any periodic class P . So, the two definitions of admissibility agree. Further, if a Heegaard diagram is weakly (or equivalently strongly) admissible for some torsion $\text{Spin}^{\mathbb{C}}$ -structure then it is weakly admissible for every $\text{Spin}^{\mathbb{C}}$ -structure. This point is useful for computations.

Both admissibility criteria are, obviously, vacuous for a rational homology sphere.

It will be useful to have equivalent definitions of weak / strong admissibility:

Lemma 5.3 Fix a pointed Heegaard diagram $\mathcal{H} = (\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ and a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} .

- (Compare [21, Lemma 4.12]) *The diagram \mathcal{H} is weakly admissible for \mathfrak{s} if and only if there is an area form on Σ with respect to which every periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 0$ has zero signed area.*
- *The diagram \mathcal{H} is strongly admissible for \mathfrak{s} if there is an area form on Σ with respect to which every periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 2n$ has signed area equal to n , and with respect to which Σ has area 1.*

Proof The proofs of the two statements are very similar, and the proof of the first statement is in [21, Lemma 4.12]. We give here only the proof of the second statement.

Let $\{D_i\}$, $i = 1, \dots, N$ denote the components of $\Sigma \setminus (\alpha \cup \beta)$. We can view the space of periodic domains as a linear subspace V of $\mathbb{Z}^N \subset \mathbb{R}^N$. Suppose an area form assigns the area a_i to D_i . Then the area assigned to P is $P \cdot (a_i)$, the dot product of the vector $P \in \mathbb{R}^N$ and the vector (a_i) . Since this is the only way the area form enters the discussion, we will refer to the vector (a_i) as the area form.

Suppose there is an area form (a_i) on Σ with respect to which every periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 2n > 0$ has signed area equal to n and Σ has area 1. Suppose $\langle c_1(\mathfrak{s}), P \rangle = 2n$. Then by assumption $P \cdot (a_i) = n$. So, $(P - n[\Sigma]) \cdot (a_i) = 0$. Hence, $P - n[\Sigma]$ must have some positive coefficient. Hence, P must have some coefficient greater than n .

The converse is slightly more involved. Note that since $area(-P) = -area(P)$, it suffices to construct an area form with the desired property for periodic domains with $\langle c_1(\mathfrak{s}), P \rangle \geq 0$.

By Lemma 4.11, the function which assigns to a periodic domain P the number $\langle c_1(\mathfrak{s}), P \rangle$ extends to an \mathbb{R} -linear functional ℓ on V . The map $v \mapsto v - (\ell(v)/2)[\Sigma]$ gives a linear projection map $p: V \rightarrow \ker(\ell)$. Let $V' = p(V)$.

Now, we want to choose $a = (a_i)$ orthogonal to V' so that $a_i > 0$ for all i . We will show that one can choose such an a presently; for now, assume that such an a has been chosen. Multiplying a by some positive real number, we can assume that $a \cdot [\Sigma] = 1$. Now, for $v \in V$,

$$a \cdot v = a \cdot p(v) + (\ell(v)/2) a \cdot [\Sigma] = \ell(v)/2 = \langle c_1(\mathfrak{s}), P \rangle / 2$$

as desired.

Finally, we need to show such an a exists. The linear space V is spanned by the periodic domains P with $\langle c_1(\mathfrak{s}), P \rangle \geq 0$, so V' is spanned by their images under p . Every periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 2n \geq 0$ has a coefficient bigger than $n = \ell(P)/2$, so every $p(P)$ has a positive coefficient. It is also true that every $p(P)$

has a negative coefficient: if $\ell(P) = 0$ this follows by applying the hypothesis to $-P$; if $\ell(P) > 0$ this follows from the fact that $n_3(P) = 0$.

Now, we are reduced to showing the following: let V' be a subspace of \mathbb{R}^N such that every nonzero vector in V' has both positive and negative coefficients. Then there is a vector orthogonal to V' with all its entries positive. The proof of this claim is a linear algebra exercise; see [21, Lemma 4.12]. \square

Now we get to the two lemmas justifying the introduction of our admissibility criteria. The following is [21, Lemma 4.13].

Lemma 5.4 *If $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ is weakly admissible for \mathfrak{s} then for each $\vec{x}, \vec{y} \in \mathfrak{s}$ and $j, k \in \mathbb{Z}$ there are only finitely many positive homology classes $A \in \pi_2(\vec{x}, \vec{y})$ with $\text{ind}(A) = j$ and $n_3(A) = k$.*

Proof If $A, B \in \pi_2(\vec{x}, \vec{y})$ and $\text{ind}(A) = \text{ind}(B) = j$, $n_3(A) = n_3(B) = k$ then $P = A - B$ is a periodic domain with $\langle c_1(\mathfrak{s}), P \rangle = 0$. So, we must show that there are only finitely many periodic domains P with $\langle c_1(\mathfrak{s}), P \rangle = 0$ such that $A + P$ is positive.

Choose an area form on Σ so that the signed area of any periodic domain P with $\langle c_1(\mathfrak{s}), P \rangle = 0$ is zero. The condition that $A + P$ be positive obviously gives a lower bound for every coefficient of P . This and the condition that the signed area of P is zero gives an upper bound for every coefficient of P . The coefficients are all integers, so the result is immediate. \square

The following is [21, Lemma 4.14].

Lemma 5.5 *If $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ is strongly admissible for \mathfrak{s} then for each $\vec{x}, \vec{y} \in \mathfrak{s}$ and $j \in \mathbb{Z}$ there are only finitely many positive homology classes $A \in \pi_2(\vec{x}, \vec{y})$ with $\text{ind}(A) = j$.*

Proof Fix a homology class $A \in \pi_2(\vec{x}, \vec{y})$ with $n_3(A) = 0$ and $\text{ind}(A) = j_0$. Then any other homology class $B \in \pi_2(\vec{x}, \vec{y})$ can be written as $A + P + k[\Sigma]$ for some integer k and periodic domain P . We have $\text{ind}(B) = j_0 + k + \langle c_1(\mathfrak{s}), P \rangle$. If we assume that $\text{ind}(B) = j$ then $\langle c_1(\mathfrak{s}), P \rangle = j - j_0 - 2k$.

Fix an area form such that the area of Σ is 1 and the area of any periodic domain P is $\frac{1}{2}\langle c_1(\mathfrak{s}), P \rangle$. Then, the area of P is $\frac{j-j_0}{2} - k$.

If we impose the condition that B be positive then we automatically get lower bounds for every coefficient of P (which are independent of k). Note that $k \geq 0$, since

$k = n_{3i}(B)$ for some i . The condition that the area of P be $\frac{j-j_0}{2} - k \leq \frac{j-j_0}{2}$ and the lower bound for the coefficients of P gives an upper bound for the coefficients of P , independent of k . This completes the proof. \square

The following is [21, Lemma 5.8 and Proposition 7.2]. We refer the reader there for its (somewhat involved but essentially elementary) proof.

Proposition 5.6 Fix a 3-manifold Y and $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on Y .

- (1) There is a weakly (respectively strongly) admissible Heegaard diagram for \mathfrak{s} .
- (2) Suppose that $\mathcal{H}_1 = (\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ and $\mathcal{H}_2 = (\Sigma', \vec{\alpha}', \vec{\beta}', \mathfrak{z}')$ are weakly (respectively strongly) admissible Heegaard diagrams for \mathfrak{s} . Then there is a sequence of pointed Heegaard moves (ie, Heegaard moves supported in the complement of \mathfrak{z}) connecting \mathcal{H}_1 to \mathcal{H}_2 such that each intermediate Heegaard diagram is weakly (respectively strongly) admissible for \mathfrak{s} .

6 Orientations

In order to be able to work with \mathbb{Z} coefficients, we need to be able to choose orientations for the moduli spaces $\widehat{\mathcal{M}}^A$ in a coherent way. First we need to know that each $\widehat{\mathcal{M}}^A$ (or equivalently, each \mathcal{M}^A) is orientable. Then we will discuss what we mean by a coherent orientation and why such orientations exist. This is all somewhat technical, and we will sometimes supply references rather than details.

Suppose that we have chosen an almost complex structure J that achieves transversality for the moduli space \mathcal{M}^A . Then for $u \in \mathcal{M}^A$ the tangent space $T_u\mathcal{M}^A$ is naturally identified with the kernel $\ker(D_u\bar{\partial})$ of the linearized $\bar{\partial}$ operator at u . In fact, the spaces $\ker(D_u\bar{\partial})$ fit together to form a vector bundle over \mathcal{M}^A naturally isomorphic to $T\mathcal{M}^A$. So, orienting \mathcal{M}^A is the same as trivializing the top exterior power of the vector bundle $\ker(D\bar{\partial})$ over \mathcal{M}^A .

Rather than working with $\ker(D\bar{\partial})$ it is better to work with the line bundle $\mathcal{L} = \det(D\bar{\partial})$ which is defined to be the tensor product of the top exterior power of $\ker(D\bar{\partial})$ with the dual of the top exterior power of $\text{coker}(D\bar{\partial})$. (This is the “determinant line bundle of the virtual index bundle of the $\bar{\partial}$ -operator.”) Note that if J achieves transversality at a curve u then \mathcal{L}_u is just the top exterior power of $\ker(D_u\bar{\partial})$.

To keep the exposition clean we will assume that our sources are stable, ie, have no components which are twice-punctured disks. Let \mathcal{T} denote the *Teichmüller configuration space* of pairs (j, u) where j is a complex structure on S (ie, an

element of the Teichmüller space of Riemann surfaces) and $u : S \rightarrow W$ is a map satisfying **(M1)**, **(M3)** and **(M4)** and which is asymptotic to the planes $\{x_i\} \times [0, 1] \times \mathbb{R}$ or $\{y_i\} \times [0, 1] \times \mathbb{R}$ at the appropriate punctures. Say (j, u) and (j', u') are equivalent if there is an isomorphism of Riemann surfaces $\phi : (S, j) \rightarrow (S', j')$ so that

$$\begin{array}{ccc}
 S' & \xrightarrow{\phi} & S \\
 & \searrow u & \swarrow u \\
 & & W
 \end{array}$$

commutes. Let \mathcal{B} be the quotient of \mathcal{T} by this equivalence relation, so \mathcal{B} is the *moduli configuration space of maps to W* . The reason to work with \mathcal{L} is that the bundle \mathcal{L} is defined over all of \mathcal{B} .

Different topological types of S correspond to different components of \mathcal{B} . So from now on we restrict attention to the subspace \mathcal{B}^S corresponding to maps from a single topological type of source S .

Note in particular that we can talk about choosing an orientation over the homology class A even if \mathcal{M}^A is empty. This will be useful when we discuss coherence.

The determinant line bundle \mathcal{L} is, in our case, always trivial. To prove this we combine constructions from [5] and [13].

Under our stability assumption, we have a fiber bundle

$$\begin{array}{ccc}
 \text{Map}(S, W) & \longrightarrow & \mathcal{B}^S \\
 & & \downarrow \\
 & & \mathcal{M}_S
 \end{array}$$

where $\text{Map}(S, W)$ consists of maps $S \rightarrow W$ satisfying **(M1)**, **(M3)** and **(M4)** with appropriate asymptotics, and \mathcal{M}_S corresponds to the moduli space of conformal structures on S .

Call a space X *homotopy discrete* if every connected component of X is contractible. The following proposition is somewhat stronger than we need. It is, however, of some independent interest, and will be mentioned again in Section 14.

Proposition 6.1 *The space $\text{Map}(S, W)$ is homotopy discrete.*

Proof The space $\text{Map}(S, W)$ is the product $\text{Map}(S, \Sigma) \times \text{Map}(S, [0, 1] \times \mathbb{R})$. The space $\text{Map}(S, [0, 1] \times \mathbb{R})$ is convex, so it suffices to prove that $\text{Map}(S, \Sigma)$ is homotopy

discrete. (Here, both $\text{Map}(S, [0, 1] \times \mathbb{R})$ and $\text{Map}(S, \Sigma)$ refer to spaces of maps with certain obvious boundary and asymptotic conditions.)

There is a fibration

$$\text{Map}(S, \Sigma) \rightarrow \Omega(\alpha_1) \times \cdots \times \Omega(\alpha_g) \times \Omega(\beta_1) \times \cdots \times \Omega(\beta_g)$$

given by restricting a map to the boundary and identifying the space of paths in α_i (or β_j) with endpoints x_i and y_i with $\Omega(\alpha_i)$ (or $\Omega(\beta_j)$). (Here, Ω denotes the based loop space.) Since $\Omega(\alpha_1) \times \cdots \times \Omega(\beta_g)$ is homotopy discrete it suffices to prove that each fiber of the fibration is homotopy discrete. Let $\text{Map}(S, \Sigma; \partial)$ denote a fiber of the fibration.

Let $\text{Map}'(S, \Sigma)$ denote the space of all maps $S \rightarrow \Sigma$ in the homotopy class of $\text{Map}(S, \Sigma; \partial)$ with no boundary conditions. There is a fibration

$$\begin{array}{ccc} \text{Map}(S, \Sigma; \partial) & \longrightarrow & \text{Map}'(S, \Sigma) \\ & & \downarrow \\ & & \text{Map}(\partial S, \Sigma) \end{array}$$

Different fibers of this fibration are homotopy equivalent. So, we can replace $\text{Map}(S, \Sigma; \partial)$ with the space $\text{Map}'(S, \Sigma; \partial)$ of maps with all $2g$ punctures of S mapped to a single point $p \in \Sigma$ and the boundary arcs mapped to a fixed list C_1, \dots, C_{2g} of circles.

There is a fibration

$$\begin{array}{ccc} \text{Map}'(S, \Sigma; \partial) & \longrightarrow & \text{Map}'_*(S, \Sigma) \\ & & \downarrow \\ & & (\Omega \Sigma)^{2g} \end{array}$$

where $\text{Map}'_*(S, \Sigma)$ denotes the component of the space of based maps $S \rightarrow \Sigma$ containing $\text{Map}'(S, \Sigma; \partial)$. Since $\Omega \Sigma$ is homotopy discrete, it suffices to prove that $\text{Map}'_*(S, \Sigma)$ is homotopy discrete. This follows from the following lemma.

Lemma 6.2 *Let K be a $K(\pi, 1)$ and X any finite, connected CW complex. Fix basepoints in K and X . Then the space of based maps from X to K is homotopy discrete.*

Proof We may assume the zero-skeleton $X^{(0)}$ of X consists of just the basepoint, so the one-skeleton $X^{(1)}$ consists of a bouquet of circles. Let $\text{Map}_*(X, K)$ denote based maps from X to K . There is a fibration $\text{Map}_*(X, K) \rightarrow \text{Map}_*(X^{(1)}, K) = (\Omega K)^N$

(where N is the number of 1-cells in X). The base is homotopy discrete since K is a $K(\pi, 1)$. So, it suffices to prove that the nonempty fibers are homotopy discrete. Each fiber F_1 consists of maps of X to K fixed on $X^{(1)}$.

The proof now proceeds by induction over the skeleta of X . There is a fibration

$$F_1 \rightarrow \prod_{2\text{-cells}} \text{Map}_*(\mathbb{D}^2, K; \partial)$$

where $\text{Map}_*(\mathbb{D}^2, K; \partial)$ denotes (based) maps of a disk with the boundary mapped as specified by the map of the one-skeleton. The base is homotopy discrete, as one deduces from the fibration

$$\begin{array}{ccc} \text{Map}_*(\mathbb{D}^2, K; \partial) & \longrightarrow & \text{Map}_*(\mathbb{D}^2, K) \\ & & \downarrow \\ & & \Omega(K) \end{array}$$

in which the total space is contractible and the base space is homotopy discrete. Thus, it suffices to show that any fiber F_2 of $F_1 \rightarrow \prod_{2\text{-cells}} \text{Map}_*(\mathbb{D}^2, K; \partial)$ is contractible. Notice that each F_2 consists of maps of X specified on the two-skeleton.

Proceeding as before, we have a fibration $F_3 \rightarrow \prod_{3\text{-cells}} \text{Map}_*(\mathbb{D}^3, K; \partial)$, now with contractible base. It suffices to prove the fiber F_4 is contractible, and so on. Since X is finite, the process terminates at F_n for $n = \dim(X)$. □

Since Σ is a $K(\pi, 1)$, this completes the proof of the proposition.

If the space \mathcal{M}_S is contractible, or even just has trivial H^1 , then we are finished: we then have $H^1(\mathcal{B}^S)$ trivial and hence all line bundles over \mathcal{B}^S are orientable. This is, in fact, the case if S is a union of disks with boundary punctures. In general, however, \mathcal{M}_S can have interesting topology.

Remark It is not hard to show, by an argument similar to but simpler than the proof of Proposition 6.1, that the configuration spaces of disks in the original construction of Heegaard Floer homology [21] have trivial π_1 , proving orientability of the moduli spaces in that setting.

Proposition 6.3 *The determinant line bundle \mathcal{L} over \mathcal{B}^S is trivial.*

Proof By Proposition 6.1 for a given source S and homology class A , the space of possible maps of ∂S is contractible. Then, it follows from the argument of [13, Lemma

6.37] or [5, Lemma 3.8], which we sketch briefly, that the line bundle \mathcal{L} is trivial. Fix a map u_0 of a collar neighborhood C of ∂S . (By a collar neighborhood, we mean that ∂C is the union of ∂S and a collection of circles, and each component of C is an annulus with punctures on one of its boundary circles.) We restrict to maps $u : S \rightarrow W$ agreeing with u_0 on C . (That it suffices to consider such maps is the only place we use anything special about our situation. If we choose u_0 to be an embedding, say, it makes sense to say u agrees with u_0 on a collar neighborhood of the boundary, even though \mathcal{B}^S is a fiber bundle and not a product of \mathcal{M}_S and $\text{Map}(S, W)$.)

Let $\partial_{in}(C)$ denote the “interior boundary” of C , ie, the circles without punctures.

As in Section 3, we are interested in the kernel and cokernel of

$$D\bar{\partial} : L_1^{p,d}(u^*TW, \partial) \oplus \mathbb{R}^{2g} \rightarrow L^{p,d}(\Lambda^{0,1}u^*TW).$$

Let S' denote the surface obtained from S by collapsing $\partial_{in}C$. The operator $D\bar{\partial}$ induces an operator $D'\bar{\partial}$ between bundles over S' , and it is not hard to see that the determinant lines of $D\bar{\partial}$ and $D'\bar{\partial}$ are naturally identified. Let C' be the image of C in S' , so S' consists of C' and some closed components. The determinant line of $D'\bar{\partial}$ is the tensor product of the determinant lines of the restriction of $D'\bar{\partial}$ to each component of S' .

Over the closed components of S' the determinant line has a natural “complex” orientation. Fix once and for all trivializations of the determinant lines of $D'\bar{\partial}$ over the components of C' . (Since $u|_C = u_0$, the restriction of $D'\bar{\partial}$ to C' does not depend on u .) Then, these trivializations induce an orientation of the determinant of $D'\bar{\partial}$, and hence of $D\bar{\partial}$, independent of u . This proves orientability of \mathcal{L} . \square

Remark For comparison, in [21], Ozsváth and Szabó only use the fact that the tangent bundles to the Lagrangian submanifolds in question are trivial. In [5], although their setup is not quite the same as standard symplectic field theory, T. Ekholm, J. Etnyre and M. Sullivan prove orientability under the assumption that the Lagrangian manifolds in question are Spin.

The coherence we want for our orientation system is the following. Suppose that we have maps $u : S \rightarrow W$ and $u' : S' \rightarrow W$ in homology classes $A \in \pi_2(\vec{x}, \vec{x}')$ and $A' \in \pi_2(\vec{x}', \vec{x}'')$, respectively. Then by gluing the positive corners of u to the negative corners of u' (in some fixed, concrete way) one can construct a 1-parameter family of curves $u \natural_r u' : S \natural_r S' \rightarrow W$. One obtains, thus, an inclusion $\text{Map}_A(T, W) \times \text{Map}_{A'}(T, W) \times (R, \infty) \rightarrow \text{Map}_{A+A'}(T, W)$. This inclusion is covered by a map of determinant line bundles. Coherence means that this map is orientation-preserving.

A few more details can be found in the proof of Proposition 13.7, and a complete discussion in any of [6, Section 1.8], [3] or [5].

There are strong general results about the existence of coherent orientations (see [6], [3] or [5]). In our case, however, we can construct them quite concretely:

As in [21, Section 3.6], given a Heegaard diagram, a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} , and an intersection point $\vec{x}_0 \in \mathfrak{s}$, a *complete set of paths for \mathfrak{s}* is a choice of homology class $T_i \in \pi_2(\vec{x}_0, \vec{x}_i)$ for every other $\vec{x}_i \in \mathfrak{s}$. Fix $\vec{x}_0 \in \mathfrak{s}$, a representative P_j of each homology class in some basis for $\hat{\pi}_2(\vec{x}_0, \vec{x}_0) = H_2(Y)$, and a complete set of paths for \mathfrak{s} . By the gluing properties of coherent orientations described two paragraphs earlier, specifying a particular coherent orientation of the moduli space is equivalent to specifying an orientation over each T_i , each P_j , and over $[\Sigma] \in \pi_2(\vec{x}_0, \vec{x}_0)$.

Over $[\Sigma]$ there is a canonical orientation, given by viewing the map $\Sigma \coprod (\coprod_{k=1}^g \mathbb{D}^2) \rightarrow \Sigma \times \mathbb{D}$ given by $id \times \{0\}$ on Σ and $\{x_{0,i}\} \times id$ on the i^{th} \mathbb{D}^2 as positively oriented. We shall always use this orientation over $[\Sigma]$. Over the T_i and P_j orientations can be chosen arbitrarily. This is exactly as in [21].

Following [21], we shall let $\sigma(A)$ denote the choice of orientation over the homology class A .

A canonical choice of coherent orientation system is specified in [20]. An analogous construction presumably works in our case as well. This is not, however, necessary for the current paper.

7 Bubbling

By a *holomorphic building* in W we mean a list v_1, v_2, \dots, v_k of holomorphic curves in W , defined up to translation in \mathbb{R} , such that the asymptotics of v_i at $+\infty$ agree with the asymptotics of v_{i+1} at $-\infty$. We call k the *height* of the holomorphic building, the v_i the *stories* of the building, and i the *level* of v_i . It is proved in [2, Theorem 10.1] that any sequence $\{u_i\}$ of holomorphic curves in W (possibly disconnected) in a given homology class has a subsequence converging to some (possibly nodal) holomorphic building.

The meaning of convergence is, as usual for holomorphic curves, somewhat involved. Roughly, convergence to a several story holomorphic building means that some parts of u_i go to infinity (in the \mathbb{R} -factor) with respect to other parts. In the source, such *level splitting* corresponds to the degeneration of the complex structure along a collection of disjoint arcs or, in principle, circles. It is also possible for the source to degenerate

along a collection of circles or arcs without level splitting occurring; this corresponds to formation of nodes as in traditional Deligne–Mumford theory. In principle, disks or spheres can also bubble off. See [2, Section 7] for a precise definition of convergence taking into account all of these possibilities.

For the definition of the Floer homology in subsequent sections to work, we need the following:

Proposition 7.1 *Fix an almost complex structure J on W satisfying (J1)–(J5) and achieving transversality. Let $u_i: S \rightarrow W$ be a sequence of holomorphic curves in a homology class A satisfying (M0)–(M6) and converging to some holomorphic building v . Assume that $\text{ind}(A) \leq 2$ and that $A \neq [\Sigma]$. Then, each story v_j of v satisfies (M0)–(M6).*

Here, by $[\Sigma]$ we mean the homology class with $n_{3i} = 1$ for every i .

Proof First we check that Deligne–Mumford type degenerations which do not correspond to level splitting are impossible. Bubbling of spheres or disks is impossible because $\pi_2(W)$ and $\pi_2(W, C_\alpha \cup C_\beta)$ vanish. This leaves us to rule out Deligne–Mumford type degenerations along non–contractible curves and arcs.

Given transversality, Deligne–Mumford type degenerations in the interior of S are prohibited for a generic choice of J because they have codimension 2 in the space of all holomorphic curves. That is, these degenerations have codimension 2 after quotienting by translation. Hence, they only occur if $\text{ind}(A) \geq 3$. This rules out all but one kind of degeneration along the interior: bubbling of a copy of $\Sigma \times (s_0, t_0)$. (Recall that our complex structure does not achieve transversality for maps u with $\pi_{\mathbb{D}} \circ u$ constant.) However, Corollary 4.3 shows that if a copy of Σ were to bubble off then the remainder of v would have to be a collection of trivial disks, and so $n_{3i}(A) = 1$ for every i . Bubbling more than one copy of Σ or a multiply covered copy of Σ is prohibited by Corollary 4.3.

So, we need to check that cusp degenerations (ie, degenerations along arcs with boundary on ∂S) are impossible.

Suppose that a cusp degeneration occurs. Let A denote an arc in S which collapses. Since different components of ∂S are mapped by u to different cylinders in W , both endpoints of A must lie on the same boundary component C of S . Without loss of generality, suppose that C is mapped by the u_i to $\alpha_1 \times \{1\} \times \mathbb{R}$. Let S' denote the nodal surface obtained from S by collapsing the arcs along which the complex structure degenerates. In the limit curve v there will be more than one boundary component

mapped to $\alpha_1 \times \{1\} \times \mathbb{R}$. Let C' denote all the boundary components mapped to $\alpha_1 \times \{1\} \times \mathbb{R}$. Consider the map $\pi_{\mathbb{D}} \circ v|_{C'}: C' \rightarrow i\mathbb{R}$. One of the boundary components in C' will have a local (in fact, global) maximum. But this and the open mapping theorem imply that the map $\pi_{\mathbb{D}} \circ v$ must be constant near that boundary component, hence constant on a connected component of S' . In particular, this implies that all of the boundary of S' is mapped to the α -circles.

However, the α -circles are non-separating, so the component of S' on which $\pi_{\mathbb{D}} \circ v$ is constant must be mapped diffeomorphically onto Σ . As above, it then follows from our index computation that the rest of v must consist of g trivial disks, and $A = [\Sigma]$.

Thus, we have proved the each u_i satisfies **(M0)**. Condition **(M1)** is automatically satisfied.

The only way that $\pi_{\mathbb{D}} \circ v$ could be constant on some component of S would be as a result of bubbling, which we showed is prohibited. This implies **(M2)** for the limit curves.

Condition **(M4)** and the condition that there are exactly $2g$ punctures on each story also follow from the maximum modulus, or open mapping, theorem: applying the open mapping theorem to $\pi_{\mathbb{D}} \circ u_j$, one sees that the restriction of $\pi_{\mathbb{R}} \circ u_j$ to any component of $\partial S \setminus \{p_1, \dots, p_g, q_1, \dots, q_g\}$ must be monotone, so no new Reeb chords could form as $j \rightarrow \infty$.

Since projection onto $[0, 1] \times \mathbb{R}$ is holomorphic, the open mapping theorem prohibits any new boundary components from forming. The maximum modulus theorem prohibits boundary components from disappearing, so **(M3)** is satisfied by the limit curves.

That condition **(M5)** is satisfied by the limit curves is part of the statement of the compactness theorem [2, Theorem 10.1].

Let S_j denote the source of v_j . Suppose that v has height ℓ . Then since all degenerations in S correspond to level splitting, $\chi(S) = (1 - \ell)g + \sum_{j=1}^{\ell} \chi(S_j)$. Let D_j be the domain in Σ corresponding to v_j . Then we have $e(D) = \sum_{j=1}^{\ell} e(D_j)$. So, $\text{ind}(u_i) = g - \chi(S) + 2e(D) = \ell g - \sum_{j=1}^{\ell} (\chi(S_j) + 2e(D_j)) = \sum_{j=1}^{\ell} \text{ind}(v_j)$. That is, the index formula adds over levels.

Finally, near any immersed curve with the equivalent of k double points and with respect to which the complex structure achieves transversality there is a $2k$ -dimensional family of embedded holomorphic curves. This shows that the v_j must all be embedded. \square

For an almost complex structure J which achieves transversality, it follows from the previous proposition, Proposition A.1 and the compactness Theorem 10.1 from [2] that:

Corollary 7.2 *The space of holomorphic curves satisfying (M0)–(M6) in a given homology class A (modulo translation in the \mathbb{R} -factor of W) with $\text{ind}(A) \leq 2$ and $A \neq [\Sigma]$ forms a smooth manifold $\widehat{\mathcal{M}}^A$ of dimension $\text{ind}(A) - 1$ which is the interior of a compact manifold with boundary $\widehat{\mathcal{M}}^A$. The boundary of that manifold $\partial \widehat{\mathcal{M}}^A$ consists of all multi-story holomorphic buildings each component of which satisfies (M0)–(M6) in the homology class A .*

8 Chain complexes

Here we define the four basic chain complexes used by Ozsváth–Szabó – \widehat{CF} , CF^∞ , CF^- and CF^+ . Once these are defined CF_{red}^\pm are defined in exactly the same way as in [21]. Generalizing our definitions to include the twisted theories of [20, Section 8] is straightforward.

Fix: a pointed Heegaard diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$, an almost complex structure J satisfying (J1)–(J5) and achieving transversality, and a coherent orientation system. We shall assume that the Heegaard diagram satisfies the weak or strong admissibility criterion as necessary.

Fix a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} .

First we define the chain complex \widehat{CF} . Suppose that our Heegaard diagram is weakly admissible for \mathfrak{s} . Recall that, by Lemma 5.4, for $\vec{x}, \vec{y} \in \mathfrak{s}$, there are only finitely many $A \in \widehat{\pi}_2(\vec{x}, \vec{y})$ and $n_{z_i}(A) \geq 0$ for all i and $\text{ind}(A) = 1$. Let $\widehat{CF}(Y, \mathfrak{s})$ be the free Abelian group generated by the intersection points \vec{x} .

Define

$$\widehat{\mathcal{M}}(\vec{x}, \vec{y}) = \bigcup_{A \in \widehat{\pi}_2(\vec{x}, \vec{y}), \text{ind}(A)=1} \widehat{\mathcal{M}}^A$$

We define a boundary operator on \widehat{CF} by

$$\partial \vec{x} = \sum_{\vec{y}} \left(\# \widehat{\mathcal{M}}(\vec{x}, \vec{y}) \right) \vec{y}.$$

This sum is finite by the admissibility criterion and Corollary 7.2.

Lemma 8.1 *With this boundary operator, \widehat{CF} is a chain complex.*

Proof Fix \vec{x}, \vec{z} . We will show that the coefficient of \vec{z} in $\partial^2 \vec{x}$ is zero. Consider the space

$$\widehat{\mathcal{M}}_2(\vec{x}, \vec{z}) = \bigcup_{A \in \widehat{\pi}_2(\vec{x}, \vec{z}), \text{ind}(A)=2} \widehat{\mathcal{M}}^A$$

with compactification $\overline{\widehat{\mathcal{M}}_2(\vec{x}, \vec{z})}$. (Note that by Lemma 5.4 there are only finitely many classes A with $\text{ind}(A) = 2$ and $\widehat{\mathcal{M}}^A \neq \emptyset$, so $\overline{\widehat{\mathcal{M}}_2(\vec{x}, \vec{z})}$ is compact.) From Corollary 7.2, $\overline{\widehat{\mathcal{M}}_2(\vec{x}, \vec{z})}$ is a 1-manifold with boundary and $\partial(\overline{\widehat{\mathcal{M}}_2(\vec{x}, \vec{z})})$ consists of broken trajectories connecting \vec{x} to \vec{z} . Thus, $0 = \#(\partial \overline{\widehat{\mathcal{M}}_2(\vec{x}, \vec{z})})$ is the coefficient of \vec{z} in $\partial^2 \vec{x}$. Note that the definition of a coherent orientation system is chosen exactly to make this argument work. \square

Next we define CF^∞ . Assume the strong admissibility criterion is satisfied. The chain group of $CF^\infty(Y; \mathfrak{s})$ is freely generated by pairs $[\vec{x}, n]$ where $\vec{x} \in \mathfrak{s}$ and $n \in \mathbb{Z}$.

The boundary operator is given by

$$\partial[\vec{x}, n] = \sum_{\vec{y}} \sum_{\substack{A \in \pi_2(\vec{x}, \vec{y}) \\ \text{ind}(A)=1}} (\#\widehat{\mathcal{M}}^A) [\vec{y}, n - n_3(A)]$$

The coefficient of each $[\vec{y}, m]$ is finite by Lemma 5.5 and Corollary 7.2.

Lemma 8.2 *Suppose that $J|_{s=0}$ and $J|_{s=1}$ have been chosen appropriately (in a sense to be specified in the proof) and that J achieves transversality for all holomorphic curves of index ≤ 1 . Then CF^∞ is a chain complex.*

Proof The proof is almost exactly the same as for \widehat{CF} . The only nuance is that the homology class $[\Sigma]$ was an exception to Corollary 7.2. One resolution of this difficulty is the following:

Recall that an annoying curve is a curve $u: S \rightarrow W$ with a component on which $\pi_{\mathbb{D}} \circ u$ is constant. The difficulty with Corollary 7.2 for the homology class $[\Sigma]$ is the possibility of a curve degenerating to a collection of trivial disks together with a copy of Σ mapped to a constant by $\pi_{\mathbb{D}} \circ u$.

Suppose that instead of considering almost complex structures satisfying (J1)–(J5) we considered the broader class of almost complex structures satisfying (J1)–(J4) and (J5'). All the results proved so far would still hold. (The only potential issue is

the proof of Proposition 7.1, but the requirement that ξ be horizontal near $\alpha \cup \beta$ is sufficient for that proof.) The transversality result would then hold for all holomorphic curves except for annoying curves mapped entirely into ∂W . Thus, annoying curves mapped into the interior of W , which are non-generic, would cease to exist.

To eliminate annoying curves mapped into ∂W we adopt an idea from [21]. Suppose that a sequence of holomorphic curves $\{u_j\}$ converges to a holomorphic curve u containing an annoying curve, which is mapped by u to $p \in \alpha \times \{1\} \times \mathbb{R}$, say. Then, rescaling near p we obtain from $\lim_{j \rightarrow \infty} \pi_{\mathbb{D}} \circ u_j$ a g -fold branched covering map $(S_0, \partial S_0) \rightarrow (\mathbb{H}, \mathbb{R})$ (where \mathbb{H} denotes the upper half plane). Here, S_0 is a surface obtained by cutting Σ along the α -circles.

Suppose that the complex structure J is given by $j'_\Sigma \times j_{\mathbb{D}}$ at $s = 1$. We will show that for appropriate choice of j'_Σ there is no map $(S_0, \partial S_0) \rightarrow (\mathbb{H}, \mathbb{R})$. Choose curves on Σ not intersecting the α -circles whose complement in Σ is a disjoint union of punctured tori (with one α -circle contained in each). Let $\{j'_{\Sigma,n}\}$ be a sequence of complex structures on Σ obtained by stretching Σ along the chosen curves. So, as $n \rightarrow \infty$, Σ degenerates to a wedge sum of tori.

Suppose that for all n there were sequences of holomorphic curves converging to annoying curves. Then for each n we obtain a map $\pi_n: (S_{0,n}, \partial S_{0,n}) \rightarrow (\mathbb{H}, \mathbb{R})$, where $S_{0,n}$ is obtained from $(\Sigma, j'_{\Sigma,n})$ by cutting along the α -circles. Choosing a convergent subsequence of the π_n , in the limit we obtain a g -fold covering map from a disjoint union of g punctured tori to \mathbb{H} . Such a map clearly does not exist.

So, for large enough n , if J agrees with $j'_{\Sigma,n} \times j_{\mathbb{D}}$ for $s = 1$ then there are no annoying curves mapped by $\pi_{\mathbb{D}}$ to $\{1\} \times \mathbb{R}$.

A similar argument shows that if we choose $J|_{s=0}$ appropriately then there are no annoying curves mapped by $\pi_{\mathbb{D}}$ to $\{0\} \times \mathbb{R}$.

It follows that, with respect to a generic complex structure extending the specified $J|_{s=0}$ and $J|_{s=1}$ and satisfying **(J1)**–**(J4)** and **(J5')**, $\partial^2 = 0$, by the same argument as for \widehat{CF} above.

Now, let J be an almost complex structure on W satisfying **(J1)**–**(J5)**, extending the specified $J|_{s=0}$ and $J|_{s=1}$ and achieving transversality for holomorphic curves of index 1. Let $\{J_n\}$ be a sequence of almost complex structures satisfying **(J1)**–**(J4)** and **(J5')** and achieving transversality which converges to J .

Let ∂_J denote the boundary map in CF^∞ computed with respect to J , ∂_{J_n} the boundary map computed with respect to J_n . Given a finite collection of homology classes $\{A_j \in \pi_2(\vec{x}_j, \vec{y}_j)\}$ such that $\text{ind}(A_j) = 1$ for all j , there is some N such that for $n > N$, $\widehat{\mathcal{M}}_{J_n}^{A_j} \cong \widehat{\mathcal{M}}_J^{A_j}$. So, since $\partial_{J_n}^2 = 0$ for all n , $\partial_J^2 = 0$. \square

Note that in Section 9 we will show that for any two choices of (generic) almost complex structure on W the pairs (CF^∞, ∂) are chain homotopy equivalent. This implies in particular that the restriction on $J|_{s=0}$ and $J|_{s=1}$ are unnecessary. Until then, we shall assume that $J|_{s=0}$ and $J|_{s=1}$ have been chosen so that the preceding proof works.

Since $n_3(A) \geq 0$ for any A supporting a holomorphic disk, CF^∞ has a subcomplex CF^- generated by the $[\vec{x}, n]$ with $n < 0$. The quotient complex we denote CF^+ . The homologies of CF^∞ , CF^+ , CF^- and \widehat{CF} we denote by HF^∞ , HF^+ , HF^- and \widehat{HF} respectively.

For computations, it is helpful to observe that CF^+ is defined even if one only assumes weak admissibility: the sum in the definition of $\partial[\vec{x}, n]$ only involves the $[\vec{x}, n - j]$ for $0 \leq j \leq n$. The weak admissibility criterion implies that for each of these finite collection of j 's, there are at most finitely many homology classes for which the moduli space is nonempty.

There is a natural action $U: CF^\infty(Y; \mathfrak{s}) \rightarrow CF^\infty(Y; \mathfrak{s})$ given by $U[\vec{x}, i] = [\vec{x}, i - 1]$. This action obviously descends to HF^∞ , making $HF^\infty(Y; \mathfrak{s})$ into a module over $\mathbb{Z}[U, U^{-1}]$. Further, the action of U preserves CF^- , so HF^- and HF^+ are modules over $\mathbb{Z}[U]$.

There are relative gradings on all four homology theories. On \widehat{CF} , define $gr(\vec{x}, \vec{y}) = \text{ind}(A)$ for any $A \in \widehat{\pi}_2(\vec{x}, \vec{y})$. It follows from Corollary 4.12 that this relative grading is defined mod n , where $n = \text{gcd}\{c_1(\mathfrak{s}), A\}$ for $A \in H_2(Y)$. Obviously the boundary map lowers the relative grading by 1, so the relative grading descends to $\widehat{HF}(Y; \mathfrak{s})$.

Similarly for $CF^\infty(Y, \mathfrak{s})$, $CF^-(Y, \mathfrak{s})$ and $CF^+(Y, \mathfrak{s})$ there are relative mod n gradings, $n = \text{gcd}\{c_1(\mathfrak{s}), A\}$, given by $gr([\vec{x}, i], [\vec{y}, j]) = \text{ind}(A) + 2(i - j)$, where $A \in \widehat{\pi}_2(\vec{x}, \vec{y})$. As before, this relative grading descends to $HF^\infty(Y, \mathfrak{s})$, $HF^+(Y, \mathfrak{s})$ and $HF^-(Y, \mathfrak{s})$. Note that the action of U lowers the relative grading by 2.

The chain complexes we have defined depend on the choice of coherent orientation system (see Section 6). It turns out, however, that some orientation systems give isomorphic chain complexes. Recall that the orientation system was given by specifying orientations over a complete set of paths T_i and a basis P_j for $\widehat{\pi}_2(\vec{x}_0, \vec{x}_0)$ (for some choice of $\vec{x}_0 \in \mathfrak{s}$).

Proposition 8.3 *Different choices of orientation over the T_i yield isomorphic chain complexes.*

Proof Let \mathfrak{o} and \mathfrak{o}' be two choices of orientation system which agree over P_j for all j . Define σ_i to be 1 if \mathfrak{o} agrees with \mathfrak{o}' over T_i and -1 otherwise. (Define σ_0 to

be 1.) It is easy to check that the map sending $[\vec{x}_i, k]$ to $\sigma_i[\vec{x}_i, k]$ (respectively \vec{x}_i to $\sigma_i\vec{x}_i$) induces isomorphisms on CF^∞ , CF^+ , and CF^- (respectively an isomorphism on \widehat{CF}). \square

It follows that there are “only” $2^{b_2(Y)}$ genuinely different choices of coherent orientation system.

Although the homologies depend on the choice of orientation system, we shall usually suppress it from the notation. Similarly, we shall often suppress the Spin^C -structure \mathfrak{s} from the notation.

8.1 Action of $H_1(Y, \mathbb{Z})/Tors$

We now define an action of $H_1(Y, \mathbb{Z})/Tors$ on the Floer homologies; cf [21, Section 4.2.5]. We will give the details only for HF^∞ . The corresponding results for HF^\pm are immediate consequences, and the results for \widehat{HF} require only slight modifications of the proofs.

Recall that for any \vec{x}_i , $\pi_2(\vec{x}_i, \vec{x}_i)$ is identified with $H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$. From Lemma 2.1, choosing a basepoint \mathfrak{z} gives an isomorphism $\Upsilon: H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}), \mathbb{Z} \xrightarrow{\sim} \mathbb{Z} \oplus H_2(Y, \mathbb{Z})$ and hence $Hom(H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}), \mathbb{Z}) \cong \mathbb{Z} \oplus Hom(H_2(Y), \mathbb{Z}) \cong \mathbb{Z} \oplus H_1(Y)/Tors$.

Choose a complete set of paths $\{T_i \in \pi_2(\vec{x}_0, \vec{x}_i)\}$. For each pair of intersection points \vec{x}_i, \vec{x}_j this gives an isomorphism $\Xi_{\{T_i\}}: \pi_2(\vec{x}_i, \vec{x}_j) \xrightarrow{\sim} H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$ via $\Xi_{\{T_i\}}(A) = \Upsilon(T_i + A - T_j)$.

Let $\zeta \in Hom(\mathbb{Z} \oplus H_2(Y), \mathbb{Z})$. Define $A_{\zeta, \{T_i\}}: CF^\infty(Y, \mathfrak{s}) \rightarrow CF^\infty(Y, \mathfrak{s})$ by

$$A_{\zeta, \{T_i\}}([\vec{x}, i]) = \sum_{\vec{y} \in \mathfrak{s}} \sum_{\substack{A \in \pi_2(\vec{x}, \vec{y}) \\ \text{ind}(A)=1}} \zeta(\Xi_{\{T_i\}}(A)) \cdot (\#\widehat{\mathcal{M}}^A)[\vec{y}, i - n_z(A)].$$

Notice that our definition is superficially different from the one used in [21, Section 4.2.5], although their definition makes sense in our language, too.

Lemma 8.4 $A_{\zeta, \{T_i\}}$ is a chain map.

Proof The proof is the same as for [21, Lemma 4.18]. Notice that $\zeta(\Xi_{\{T_i\}}(A + B)) = \zeta(\Xi_{\{T_i\}}(A)) + \zeta(\Xi_{\{T_i\}}(B))$. Suppose $C \in \pi_2(\vec{x}, \vec{y})$, $\text{ind}(C) = 2$. Then

$$0 = \left(\# \partial \widehat{\mathcal{M}}^C \right), \text{ so}$$

$$\begin{aligned} 0 &= \sum_{\substack{C \in \pi_2(\vec{x}, \vec{z}) \\ \text{ind}(C)=2 \\ n_z(C)=k}} \zeta(\Xi_{\{T_i\}}(C)) \left(\# \left(\partial \widehat{\mathcal{M}}^C \right) \right) \\ &= \sum_{\substack{C \in \pi_2(\vec{x}, \vec{z}) \\ \text{ind}(C)=2 \\ n_z(C)=k}} \sum_{\substack{A+B=C \\ \text{ind}(A)=\text{ind}(B)=1}} \left(\zeta(\Xi_{\{T_i\}}(A)) + \zeta(\Xi_{\{T_i\}}(B)) \right) \left(\# \widehat{\mathcal{M}}^A \right) \left(\# \widehat{\mathcal{M}}^B \right). \end{aligned}$$

This is the coefficient of $[\vec{z}, i - k]$ in $(\partial \circ A_{\zeta, \{T_i\}} + A_{\zeta, \{T_i\}} \circ \partial)([\vec{x}, i])$, proving the result. \square

The following lemma is analogous to [21, Lemma 4.19].

Lemma 8.5 *If T'_i is another complete set of paths then $A_{\zeta, \{T_i\}}$ and $A_{\zeta, \{T'_i\}}$ are chain homotopic.*

Proof Let $P_i = T_i - T'_i$. Consider the map $H: CF^\infty(Y, \mathfrak{s}) \rightarrow CF^\infty(Y, \mathfrak{s})$ given by $H([\vec{x}_i, j]) = \zeta(\Upsilon(P_i))[\vec{x}_i, j]$. Then

$$\begin{aligned} &(A_{\zeta, \{T_i\}} - A_{\zeta, \{T'_i\}})([\vec{x}_i, j]) \\ &= \sum_{\substack{\vec{x}_k \in \mathfrak{s} \\ \text{ind}(A)=1}} \sum_{A \in \pi_2(\vec{x}_i, \vec{x}_k)} \zeta(\Xi_{\{T_i\}}(A) - \Xi_{\{T'_i\}}(A)) \cdot \left(\# \widehat{\mathcal{M}}^A \right) [\vec{x}_k, j - n_3(A)] \\ &= \sum_{\substack{\vec{x}_k \in \mathfrak{s} \\ \text{ind}(A)=1}} \sum_{A \in \pi_2(\vec{x}_i, \vec{x}_k)} \zeta(\Upsilon(A + T_i - T_k) - \Upsilon(A + T'_i - T'_k)) \\ &\quad \cdot \left(\# \widehat{\mathcal{M}}^A \right) [\vec{x}_k, j - n_3(A)] \\ &= \sum_{\substack{\vec{x}_k \in \mathfrak{s} \\ \text{ind}(A)=1}} \sum_{A \in \pi_2(\vec{x}_i, \vec{x}_k)} \zeta(\Upsilon(T_i - T'_i) - \Upsilon(T_k - T'_k)) \cdot \left(\# \widehat{\mathcal{M}}^A \right) [\vec{x}_k, j - n_3(A)] \\ &= \sum_{\substack{\vec{x}_k \in \mathfrak{s} \\ \text{ind}(A)=1}} \sum_{A \in \pi_2(\vec{x}_i, \vec{x}_k)} \zeta(\Upsilon(P_i) - \Upsilon(P_k)) \cdot \left(\# \widehat{\mathcal{M}}^A \right) [\vec{x}_k, j - n_3(A)] \\ &= \partial \circ H + H \circ \partial \end{aligned}$$

\square

We are now justified in denoting the map on HF^∞ induced by $A_{\xi, \{T_i\}}$ (for any complete set of paths $\{T_i\}$) by simply A_ξ .

The following is [21, Proposition 4.17].

Proposition 8.6 *There is a natural action of the group $\mathbb{Z} \oplus \text{Hom}(H_2(Y), \mathbb{Z})$ on $HF^\infty(Y, \mathfrak{s})$, $HF^+(Y, \mathfrak{s})$, $HF^-(Y, \mathfrak{s})$ and $\widehat{HF}(Y, \mathfrak{s})$ lowering degree by one. This induces an action of the exterior algebra $\bigwedge^*(H_1(Y)/\text{Tors})$ on each group.*

Proof The action of ζ is given by A_ζ . Obviously A_ζ lowers the grading by one and $A_{\zeta+\zeta'} = A_\zeta + A_{\zeta'}$. We need to check that $A_\zeta \circ A_\zeta = 0$.

Choose a curve K in $f^{-1}(3/2 - \epsilon, 3/2 + \epsilon) = \Sigma \times [0, 1]$ representing the homology class ζ in $H_1(Y)/\text{Tors}$. That is, choose K so that $A_\zeta(A)$ is the intersection number of A with K for $A \in \pi_2(\vec{x}, \vec{x})$. Perturb K so that $K \times \mathbb{R}$ meets transversely every holomorphic curve u with $\text{ind}(u) = 1$, and is transverse to families of holomorphic curves with $\text{ind} = 2$.

Let \mathcal{M}_1 denote the moduli space of Riemann surfaces with one marked point p . Let \mathcal{M}_K^A denote the space of holomorphic curves $u: (S, p) \rightarrow (W, K \times \mathbb{R})$ with $(S, p) \in \mathcal{M}_1$ which, after forgetting the marked point, represent the homology class A . Let $\widehat{\mathcal{M}}_K^A = \mathcal{M}_K^A/\mathbb{R}$. Then, assuming the appropriate transversality result, which is left to the reader, we have

$$A_\zeta([\vec{x}, i]) = \sum_{\substack{A \in \pi_2(\vec{x}, \vec{y}) \\ \text{ind}(A)=1}} \left(\#\widehat{\mathcal{M}}_K(A) \right) [\vec{y}, i - n_3(A)].$$

To prove the proposition, consider the space \mathcal{M}_2 of Riemann surfaces with 2 marked points $\{p_1, p_2\}$. Let $\mathcal{M}_{K,2}^A$ denote the space of holomorphic maps $u: (S, \{p_1, p_2\}) \rightarrow (W, K \times \mathbb{R})$, for $(S, \{p_1, p_2\}) \in \mathcal{M}_2$. Let $\widehat{\mathcal{M}}_{K,2}^A/\mathbb{R} = \mathcal{M}_{K,2}^A/\mathbb{R}$.

For a generic choice of J_s , and a homology class A with $\text{ind}(A) = 2$, $\widehat{\mathcal{M}}_{K,2}^A(A)$ is a smooth 1-manifold. The manifold $\widehat{\mathcal{M}}_{K,2}^A(A)$ has four kinds of ends:

- (1) Ends corresponding to $\pi_{\mathbb{R}} \circ u(p_1) - \pi_{\mathbb{R}} \circ u(p_2) \rightarrow \infty$. These correspond to $A_\zeta \circ A_\zeta$.
- (2) Ends corresponding to $\pi_{\mathbb{R}} \circ u(p_1) - \pi_{\mathbb{R}} \circ u(p_2) \rightarrow -\infty$. These also correspond to $A_\zeta \circ A_\zeta$.
- (3) Ends where there is a q in S with $\pi_{\mathbb{R}} \circ u(q) - \pi_{\mathbb{R}} \circ u(p_1) \rightarrow \infty$ but $\pi_{\mathbb{R}} \circ u(p_1) - \pi_{\mathbb{R}} \circ u(p_2)$ stays bounded. These correspond to $\partial \circ A_\zeta$.

- (4) Ends where there is a q in S with $\pi_{\mathbb{R}} \circ u(q) - \pi_{\mathbb{R}} \circ u(p_1) \rightarrow -\infty$ but $\pi_{\mathbb{R}} \circ u(p_1) - \pi_{\mathbb{R}} \circ u(p_2)$ stays bounded. These correspond to $A_{\xi} \circ \partial$.

There is also a free action of $\mathbb{Z}/2\mathbb{Z}$ on $\widehat{\mathcal{M}}_{K,2}^A$ exchanging the labeling of p_1 and p_2 . Counting the ends of $\widehat{\mathcal{M}}_{K,2}^A/(\mathbb{Z}/2\mathbb{Z})$ and summing over A then shows that $A_{\xi} \circ A_{\xi}([\vec{x}, i])$ is chain homotopic to 0.

The result for HF^{∞} is immediate. The constructions of the lemmas and this proof preserve CF^{-} , so the results for HF^{\pm} follow. The proofs for \widehat{HF} are analogous. \square

9 Isotopy invariance

In this section we prove that the homologies defined in Section 8 are independent of deformations of the almost complex structure and isotopies of the α - and β -circles not crossing the basepoint \mathfrak{z} . Two parts of this story are slightly nonstandard. One is extending the coherent orientation system through isotopies which introduce new intersection points. The other is that we want to allow seemingly non-Hamiltonian isotopies of the α and β curves. The rest of the proof is analogous to the discussion in [6, Section 1.9].

We discuss how to extend the orientation system first. On a first reading the reader might want to skip the next three paragraphs.

Suppose that we have two pointed Heegaard diagrams $D_1 = (\Sigma, j, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ and $D_2 = (\Sigma, j', \vec{\alpha}', \vec{\beta}', \mathfrak{z})$ which differ by pointed isotopies (ie, isotopies during which none of the curves cross \mathfrak{z}) and deformations of the complex structure. The only interesting case is when the $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} , viewed as a collection of intersection points, is not empty in either Heegaard diagram. (If in one of the diagrams \mathfrak{s} contains no intersection points then any choice of orientation system for the other will be fine.) By choosing an appropriate isotopy and then breaking it into a sequences of isotopies, we can assume that some intersection point $\vec{x}_0 \in \mathfrak{s}$ exists in both Heegaard diagrams. For convenience, we fix a parametrization of Σ such that the α -circles stay fixed during the isotopy. Then, with this parametrization, the β -circles move and the complex structure deforms during the isotopy.

We can identify $\pi_2(\vec{x}_0, \vec{x}_0)$ in D_1 with $\pi_2(\vec{x}_0, \vec{x}_0)$ in D_2 as follows. Denote the isotopy by $I: \beta \times [0, 1] \rightarrow \Sigma$. Let $\phi: (C, \partial C) \rightarrow (W, \alpha \cup \beta)$ be any singular 2-chain in $\pi_2^{D_1}(\vec{x}_0, \vec{x}_0)$, C a simplicial complex. Let $\partial_{\beta} C = \partial C \cap I^{-1}(\beta)$. Define $\psi: \partial_{\beta} C \times [0, 1] \rightarrow W$ by $\psi(x, t) = (I(\pi_{\Sigma} \circ \phi(x), t), \pi_{\mathbb{D}} \circ \phi(x))$. Then identify ϕ with the chain in $\pi_2^{D_2}(\vec{x}_0, \vec{x}_0)$ given by $\phi + \psi$.

Thus, an orientation system over a complete set of paths in D_2 and an orientation system for D_1 determine an orientation system for D_2 . We assume that we are computing the homologies of D_2 with respect to an orientation system determined from the orientation system of D_1 in this way; as observed in the section on chain complexes, different orientation systems over a complete set of paths lead to isomorphic homologies. So, which particular one we choose is unimportant. Note also that our choices determine, for \vec{x}_i^k an intersection point in D_k , an isomorphism $\pi_2(\vec{x}_i^1, \vec{x}_j^2) \cong \mathbb{Z} \oplus H_2(Y)$.

By a *basic isotopy* we mean an isotopy (α_t, β_t) with one of the following two properties.

- (1) For all times t , α_t intersects β_t transversally. (These are *basic isotopies of the first type*)
- (2) The isotopy introduces one pair of transverse intersections between α_t and β_t by a Hamiltonian deformation (“finger move”) of the α -circles. (These are *basic isotopies of the second type*.)

We only consider isotopies which are sequences of basic isotopies. Call such an isotopy *strongly admissible* (respectively *weakly admissible*) if before and after each basic isotopy the Heegaard diagram is strongly (respectively weakly) admissible (for \mathfrak{s}).

It is clear that if two Heegaard diagrams are isotopic then they are isotopic through a sequence of basic isotopies. In fact, by Proposition 5.6, any two isotopic strongly admissible Heegaard diagrams are isotopic through a sequence of strongly admissible basic isotopies, and any two isotopic weakly admissible Heegaard diagrams are isotopic through a sequence of weakly admissible basic isotopy.

Following [21, Section 7.3], we use the fact that basic isotopies of the first type are equivalent to deformations of the complex structures on Σ and W . We make this precise. Suppose that D_2 differs from D_1 by a basic isotopy of the first type. Then, there is an orientation-preserving diffeomorphism ψ of Σ taking D_1 to D_2 and mapping $U_{z_i}^{D_1}$ onto $U_{z_i}^{D_2}$. It follows that computing the homologies of D_2 with respect to the complex structures j_Σ on Σ and J_s on W is the same as computing the homologies of D_1 with respect to $\psi^* j_\Sigma$ on Σ and $(\psi \times Id)^* J_s$ on W . Note that if J_s satisfies **(J1)**–**(J5)** with respect to j_Σ then so does $(\psi \times Id)^* J_s$ with respect to $\psi^* j_\Sigma$. Consequently, independence of the homologies with respect to isotopies preserving transversality of $\alpha \cap \beta$ will follow from independence with respect to complex structure.

The other point we need to check is that basic isotopies of the second type do not change the homologies, either. Suppose that D_2 is obtained from D_1 by a basic

isotopy of the second type. Then we can find a collection of Lagrangian cylinders in $\Sigma \times [0, 1] \times \mathbb{R}$ which agrees with $(\alpha \times \{1\} \times \mathbb{R}) \cup (\beta \times \{0\} \times \mathbb{R})$ near $t = -\infty$ and with $(\alpha' \times \{1\} \times \mathbb{R}) \cup (\beta' \times \{0\} \times \mathbb{R})$ near $t = \infty$. Call the collection of these Lagrangian cylinders C .

We combine invariance under both types of basic isotopy into one:

Proposition 9.1 *Suppose that either D_1 differs from D_2 only by an isotopy preserving transversality of the α and β circles, or that they differ by such an isotopy and a pair creation, and that both D_1 and D_2 are strongly (or, in the case of CF^+ and \widehat{CF} , weakly) admissible Heegaard diagrams. Suppose J_1 (respectively J_2) satisfies **(J1)**–**(J5)** and achieves transversality for D_1 (respectively D_2), with respect to j_1 (respectively j_2) on Σ . Then the chain complexes $CF_{D_1}^\infty$ and $CF_{D_2}^\infty$ (respectively $CF_{D_1}^-$ and $CF_{D_2}^-$; $CF_{D_1}^+$ and $CF_{D_2}^+$; and \widehat{CF}_{D_1} and \widehat{CF}_{D_2}) are chain homotopy equivalent. Further, the isomorphisms induced on homologies respect the $H_1(Y)/\text{Tor}s$ -module and, where appropriate, the $\mathbb{Z}[U, U^{-1}]$ - or $\mathbb{Z}[U]$ -module structures.*

By the discussion preceding the proposition, proving this proposition proves the independence of the Floer homologies under isotopies. Note that it also implies that the restrictions we imposed in Section 8 on $J|_{s=0}$ and $J|_{s=1}$ for CF^∞ and CF^\pm are not needed for ∂^2 to be 0.

Fix $T > 0$. Choose an almost complex structure J on $\Sigma \times [0, 1] \times \mathbb{R}$ which satisfies **(J1)**, **(J2)** and **(J4)**, agrees with J_1 on $(-\infty, -T)$ and with J_2 on $[T, \infty)$, and achieves transversality for holomorphic curves of the form $u: (S, \partial S) \rightarrow (\Sigma \times [0, 1] \times \mathbb{R}, C)$. This is possible by essentially the same argument as in Section 3.

We define a chain map from each chain complex defined on D_1 to the corresponding chain complex on D_2 by counting J -holomorphic curves in W . We carry out the details for CF^∞ ; the results for CF^+ and CF^- will follow, and the proof for \widehat{CF} is similar. Let (CF^∞, ∂_i) be the chain complex defined on D_i , $i = 1, 2$.

For \vec{x}^1 (respectively \vec{x}^2) an intersection point in D_1 (respectively D_2), let $\mathcal{M}_0(\vec{x}^1, \vec{x}^2)$ (respectively $\mathcal{M}_1(\vec{x}^1, \vec{x}^2)$) denote the space of all holomorphic curves u in (W, J) connecting \vec{x}^1 to \vec{x}^2 in homology classes A with $\text{ind}(A) = 0$ (respectively $\text{ind}(A) = 1$), and satisfying **(M0)**–**(M6)** (with respect to the new Lagrangian cylinders C).

Consider the map $\Phi: CF_1^\infty \rightarrow CF_2^\infty$ defined by

$$\Phi([\vec{x}^1, i]) = \sum_{\vec{x}^2 \in s} \sum_{u \in \mathcal{M}_0(\vec{x}^1, \vec{x}^2)} [\vec{x}^2, i - n_z(u)].$$

We need to check that the coefficient of $[\vec{x}^2, j]$ is a finite sum for each \vec{x}^2 and j :

Lemma 9.2 *Given $j \in \mathbb{Z}$ there are at most finitely many homology classes A connecting \vec{x}^1 to \vec{x}^2 with $\text{ind}(A) = j$ which admit a holomorphic curve.*

Proof This follows from the strong admissibility criterion; our proof is essentially the same as [21, Lemma 7.4].

Choose a point z_i in each component of $\Sigma \setminus (\alpha^1 \cup \beta^1)$ in such a way that none of the α - or β -circles cross any of the z_i during the isotopy. If A supports a J -holomorphic curve then $n_{z_i}(A) \geq 0$ for all i . Choose any homology class in $B \in \pi_2^{D_2}(\vec{x}^2, \vec{x}^1)$. We can view $A + B$ as an element of $\pi_2^{D_1}(\vec{x}^1, \vec{x}^1)$, and $n_{z_i}(A + B) \geq n_{z_i}(B)$ for all i . Now the argument used in Lemma 5.5 gives bounds for the coefficients of $A + B$, and hence bounds for the coefficients of A . This proves the result. \square

We shall sometimes denote Φ by Φ_{12} to emphasize that Φ is induced from a bordism from D_1 to D_2 .

Lemma 9.3 *Φ is a chain map.*

Proof Let ∂_i denote the boundary map on the chain complex for D_i .

We consider the compactified moduli space $\overline{\mathcal{M}}_1(\vec{x}^1, \vec{y}^2, k)$ of index 1 holomorphic curves in homology classes A with $n_3(A) = k$. There is still no bubbling, so this is a compact one-manifold whose boundary consists of height two holomorphic buildings one story of which lies in (W, J) and the other of which lies in either (W, J_1) or (W, J_2) . Hence,

$$0 = \# \left(\partial \overline{\mathcal{M}}_1(\vec{x}^1, \vec{y}^2, k) \right).$$

But this is the coefficient of $[\vec{y}^2, i - k]$ in $\Phi \circ \partial_1([\vec{x}^1, i]) + \partial_2 \circ \Phi([\vec{x}^1, i])$. \square

Lemma 9.4 *Given two different choices of complex structure J and J' connecting J_1 to J_2 and two different choices C and C' for the cylinders connecting $\alpha \times \{1\} \cup \beta \times \{0\}$ to $\alpha' \times \{1\} \cup \beta' \times \{0\}$ the maps Φ and Φ' are chain homotopic.*

Proof We outline the proof; further details are left to the reader. Choose a generic path J_t from J to J' and a Lagrangian isotopy C_t from C to C' . Let $\mathcal{M}_{-1,t}(\vec{x}^1, \vec{x}^2)$ denote the moduli space of holomorphic curves u connecting \vec{x}^1 to \vec{x}^2 in homology classes of index -1 . For a finite collection of t_i 's, $0 < t_1 < \dots < t_k < 1$, J_{t_i} is degenerate in such a way that $\mathcal{M}_{-1,t}(\vec{x}^1, \vec{x}^2)$ is nonempty. (As usual, the finiteness uses the admissibility hypothesis.) Then we define a chain homotopy Δ by

$$\Delta([\vec{x}^1, i]) = \sum_t \sum_{\vec{x}^2} \sum_{u \in \mathcal{M}_{-1,t}} [\vec{x}^2, i - n_3(u)].$$

Similarly, define $\mathcal{M}_{0,t}(\vec{x}^1, \vec{y}^2)$ to be the moduli space of holomorphic curves u connecting \vec{x}^1 to \vec{y}^2 in homology classes of index 0. Then $\cup_t \mathcal{M}_{0,t}(\vec{x}^1, \vec{y}^2)$ is a 1-manifold with boundary with four types of ends:

- (1) Ends corresponding to height two holomorphic buildings, with a J_{t_k} -holomorphic curve of index -1 and a J_1 -holomorphic curve of index 1 . These correspond to $\Delta \circ \partial_1$.
- (2) Ends corresponding to height two holomorphic buildings, with a J_{t_k} -holomorphic curve of index -1 and a J_2 -holomorphic curve of index 1 . These correspond to $\partial_2 \circ \Delta$.
- (3) Ends corresponding to $t = 0$ and a J -holomorphic curve. These correspond to $-\Phi$.
- (4) Ends corresponding to $t = 1$ and a J' -holomorphic curve. These correspond to Φ' .

So, counting the ends gives the result. \square

Lemma 9.5 *If we have a third diagram D_3 and a bordism from D_2 to D_3 then on homology $\Phi_{23} \circ \Phi_{12} = \Phi_{13}$.*

Note that since different choices of bordism give chain homotopic Φ_{ij} the exact choices of Φ_{ij} are unimportant here.

Proof This is immediate from compactness and Proposition A.1. \square

Lemma 9.6 *The maps on the homologies induced by Φ preserve the $H_1(Y)/Tors$ -module structure.*

Proof Fix $\zeta \in H_1(Y)/Tors$. Let K be a knot in $\Sigma \times [0, 1]$ representing ζ , as in the proof of Proposition 8.6.

Let \mathcal{M}_1 denote the moduli space of Riemann surfaces with one marked point p . For $i = 1, 2$, let $\mathcal{M}_{K,i}^A$ denote the space of J_i -holomorphic curves $u: (S, p) \rightarrow (W, K \times \mathbb{R})$ with $(S, p) \in \mathcal{M}_1$ which, after forgetting the marked point, represent the homology class A . Let $\widehat{\mathcal{M}}_{K,i}^A = \mathcal{M}_{K,i}^A / \mathbb{R}$. Assume K (or J_1 and J_2) is chosen so that we have transversality. Then, on the chain level the action of $H_1(Y)/Tors$ on $HF_{D_i}^\infty$ is given by

$$A_{\zeta,i}([\vec{x}^i, j]) = \sum_{\substack{A \in \pi_2(\vec{x}^i, \vec{y}^i) \\ \text{ind}(A)=1}} \left(\# \widehat{\mathcal{M}}_{K,i}^A \right) [\vec{y}^i, j - n_z(A)].$$

Let \mathcal{M}_K^A denote the space of J -holomorphic curves $u: (S, p) \rightarrow (W, K \times \mathbb{R})$ with $(S, p) \in \mathcal{M}_1$ which, after forgetting the marked point, represent the homology class A . Assume K, J_1, J_2 and J are chosen so that we have transversality. Consider the ends of

$$\mathcal{M}_K(\vec{x}^1, \vec{y}^2, k) = \bigcup_{\substack{A \in \pi_2(\vec{x}^1, \vec{y}^2) \\ \text{ind}(A)=1 \\ n_3(A)=k}} \mathcal{M}_K(A).$$

There are two kinds of ends:

- (1) Ends where $\pi_{\mathbb{D}} \circ u(p) \rightarrow \infty$. These ends correspond to the coefficient of $[\vec{y}, i - k]$ in $A_{\xi, 2} \circ \Phi([\vec{x}, i])$
- (2) Ends where $\pi_{\mathbb{R}} \circ u(p) \rightarrow -\infty$. These ends correspond to the coefficient of $[\vec{y}, i - k]$ in $\Phi \circ A_{\xi, 1}([\vec{x}, i])$.
- (3) Ends where $\pi_{\mathbb{R}} \circ u(p)$ stays bounded and an index 1 curve splits-off at $+\infty$. These correspond to the coefficient of $[\vec{y}, i - k]$ in $\partial \circ \Phi([\vec{x}^1, i - k])$.
- (4) Ends where $\pi_{\mathbb{R}} \circ u(p)$ stays bounded and an index 1 curve splits-off at $-\infty$. These correspond to the coefficient of $[\vec{y}, i - k]$ in $\Phi \circ \partial([\vec{x}^1, i - k])$.

Counting the ends, and using the fact that $\Phi \circ \partial + \partial \circ \Phi = 0$, gives the result. □

Proof of Proposition 9.1.

The proposition follows immediately from these four lemmas. From Lemma 9.3 we have chain maps Φ_{12} and Φ_{21} whose induced maps preserve the $H_1(Y)/Tors$ -structure by Lemma 9.6. From Lemma 9.4 and Lemma 9.5, $\Phi_{12} \circ \Phi_{21}$ is chain homotopic to the identity map, as is $\Phi_{21} \circ \Phi_{12}$. All the maps are obviously maps of $\mathbb{Z}[U, U^{-1}]$ -modules. This proves Proposition 9.1 for HF^∞ . To conclude the result for HF^+ and HF^- it is only necessary to make the trivial observation that all of the maps used preserve HF^- .

The proof for \widehat{HF} is completely analogous – one simply restricts in each case to holomorphic curves with $n_3 = 0$.

10 Triangles

By a pointed Heegaard triple-diagram we mean a Riemann surface Σ together with three g -tuples of pairwise disjoint, homologically linearly independent simple closed curves $\vec{\alpha} = \{\alpha_1, \dots, \alpha_g\}$, $\vec{\beta} = \{\beta_1, \dots, \beta_g\}$ and $\vec{\gamma} = \{\gamma_1, \dots, \gamma_g\}$, together with a basepoint $z \in \Sigma \setminus (\alpha \cup \beta \cup \gamma)$. A Heegaard triple-diagram specifies three 3-manifolds $Y_{\alpha, \beta}, Y_{\beta, \gamma}$

and $Y_{\alpha,\gamma}$. It also specifies a 4-manifold $X_{\alpha,\beta,\gamma}$ with boundary $Y_{\alpha,\beta} \cup Y_{\beta,\gamma} \cup -Y_{\alpha,\gamma}$ as follows. The curves $\vec{\alpha}$ (respectively $\vec{\beta}$, $\vec{\gamma}$) specify a handlebody U_α (respectively U_β , U_γ) with boundary Σ . Let \overline{T} be a triangle with edges e_1 , e_2 , and e_3 . Define $X_{\alpha,\beta,\gamma}$ to be the manifold obtained by gluing $U_\alpha \times [0, 1]$, $U_\beta \times [0, 1]$, and $U_\gamma \times [0, 1]$ to $\Sigma \times \overline{T}$ along $\Sigma \times e_1$, $\Sigma \times e_2$, and $\Sigma \times e_3$ by identifying $\alpha \in \partial U_\alpha \times \{p\}$ with $\alpha \in \Sigma \times \{p\}$ ($p \in e_1 = [0, 1]$) and similarly for β and γ .

In this (rather long) section we associate to a Heegaard triple-diagram maps between the Floer homologies of $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$ and $Y_{\alpha,\gamma}$. Specifically

Construct 10.1 *To an admissible (see Subsection 10.4) pointed Heegaard triple-diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \mathfrak{z})$ and a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{t} on $X_{\alpha,\beta,\gamma}$, as well as a coherent orientation system, as described in Subsection 10.3, we associate U -equivariant homomorphisms*

$$\begin{aligned} \widehat{F}_{\alpha,\beta,\gamma} &: \widehat{HF}(Y_{\alpha,\beta}; \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes_{\mathbb{Z}} \widehat{HF}(Y_{\beta,\gamma}; \mathfrak{t}|_{Y_{\beta,\gamma}}) \rightarrow \widehat{HF}(Y_{\alpha,\gamma}; \mathfrak{t}|_{Y_{\alpha,\gamma}}) \\ F_{\alpha,\beta,\gamma}^{\infty} &: HF^{\infty}(Y_{\alpha,\beta}; \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes_{\mathbb{Z}[U]} HF^{\infty}(Y_{\beta,\gamma}; \mathfrak{t}|_{Y_{\beta,\gamma}}) \rightarrow HF^{\infty}(Y_{\alpha,\gamma}; \mathfrak{t}|_{Y_{\alpha,\gamma}}) \\ F_{\alpha,\beta,\gamma}^{+} &: HF^{+}(Y_{\alpha,\beta}; \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes_{\mathbb{Z}[U]} HF^{\leq 0}(Y_{\beta,\gamma}; \mathfrak{t}|_{Y_{\beta,\gamma}}) \rightarrow HF^{+}(Y_{\alpha,\gamma}; \mathfrak{t}|_{Y_{\alpha,\gamma}}) \\ F_{\alpha,\beta,\gamma}^{-} &: HF^{\leq 0}(Y_{\alpha,\beta}; \mathfrak{t}|_{Y_{\alpha,\beta}}) \otimes_{\mathbb{Z}[U]} HF^{\leq 0}(Y_{\beta,\gamma}; \mathfrak{t}|_{Y_{\beta,\gamma}}) \rightarrow HF^{\leq 0}(Y_{\alpha,\gamma}; \mathfrak{t}|_{Y_{\alpha,\gamma}}). \end{aligned}$$

These maps satisfy an associativity property stated in Proposition 10.29.

Of course, the zero map satisfies these conditions, but our maps are usually more interesting. In particular, an instance of the construction will be used to prove handleslide invariance in the next section. (They are also a key computational tool, have been used to define 4-manifold invariants, contact invariants, and so on.) They will be produced by counting holomorphic curves in the product of Σ and a disk with three boundary punctures, ie, a triangle.

The outline of this section is as follows. In Subsection 10.1 we discuss basic topological properties of maps to $\Sigma \times T$, T a “triangle”. In particular, we discuss when such maps exist, how many homology classes of them there are, and how such homology classes specify $\text{Spin}^{\mathbb{C}}$ -structures. In Subsection 10.2 we discuss some basic prerequisites for construction of the triangle maps (the complex structures we consider, the index) and then define the triangle maps, conditional on certain technicalities to be addressed in the following two sections. In Subsection 10.3 we address the first of these technicalities: coherent orientations of the moduli spaces of triangles. In Subsection 10.4 we deal with the second technicality: admissibility criteria for Heegaard triple-diagrams necessary for the triangle maps to be defined. In Subsection 10.5 we return to the definition of the triangle maps, proving that the maps are chain maps, and are independent of the




Figure 7: The triangle T

complex structure on $\Sigma \times T$ and isotopies of the α -, β - and γ -circles. Finally, in Subsection 10.6 we prove an associativity property of triangle maps.

10.1 Topological preliminaries on triangles

Fix a pointed Heegaard triple-diagram $\mathcal{H}^3 = (\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \mathfrak{z})$. Let $\mathcal{H}_{\alpha,\beta}$, $\mathcal{H}_{\beta,\gamma}$ and $\mathcal{H}_{\alpha,\gamma}$ denote the pointed Heegaard diagrams for $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$ and $Y_{\alpha,\gamma}$ specified by \mathcal{H}^3 .

10.1.1 Homological preliminaries on triangles From now on, by the triangle T we mean a Y-shaped region in \mathbb{C} with three cylindrical ends, as shown in Figure 7. Note that T is conformally equivalent to a (in fact, any) triangle with punctures at the corners. (We will occasionally use a closed triangle, which we continue to denote \bar{T} and think of as the closure of T .) Let e_1 , e_2 , and e_3 denote the three boundary components of T , ordered clockwise, and v_{12} , v_{23} and v_{13} the ends between e_1 and e_2 , e_2 and e_3 , and e_1 and e_3 respectively. Let $W_{\alpha,\beta,\gamma} = \Sigma \times T$. For I-chord collections \vec{x} of \mathcal{H}^1 , \vec{y} of \mathcal{H}^2 and \vec{z} of \mathcal{H}^3 , let $\pi_2(\vec{x}, \vec{y}, \vec{z})$ denote the collection of homology classes of maps $(S, \partial S) \rightarrow (W_{\alpha,\beta,\gamma}, \alpha \times e_1 \cup \beta \times e_2 \cup \gamma \times e_3)$ (S a Riemann surface with boundary and punctures on the boundary) which are asymptotic to \vec{x} at v_{12} , \vec{y} at v_{23} and \vec{z} at v_{13} . As before, there is a map $n_3: \pi_2(\vec{x}, \vec{y}, \vec{z}) \rightarrow \mathbb{Z}$. Let $\hat{\pi}_2(\vec{x}, \vec{y}, \vec{z}) = \{A \in \pi_2(\vec{x}, \vec{y}, \vec{z}) \mid n_3(A) = 0\}$.

Let $W_{\alpha,\beta}$, $W_{\beta,\gamma}$ and $W_{\alpha,\gamma}$ denote the three cylindrical manifolds which are the ends of $W_{\alpha,\beta,\gamma}$.

Note that there are concatenation maps $*$: $\pi_2(\vec{x}, \vec{x}') \times \pi_2(\vec{x}, \vec{y}, \vec{z}) \rightarrow \pi_2(\vec{x}', \vec{y}, \vec{z})$ (and similarly for \vec{y} and \vec{z}).

Suppose that $\pi_2(\vec{x}, \vec{y}, \vec{z})$ is nonempty. Fix an element $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$. Let G denote the tree with one vertex of valence three and three vertices of valence 1 (so G looks like a figure Y). Let v_1, v_2 , and v_3 denote the three valence–one vertices in G . Then A determines a map

$$\pi_2(\vec{x}, \vec{y}, \vec{z}) \rightarrow H_2(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\})$$

by subtracting from each element of $\pi_2(\vec{x}, \vec{y}, \vec{z})$ a representative for A and pushing forward via the retract suggested by Figure 7. It is easy to see that this map is bijective.

Recall that $\pi_2(\vec{x}, \vec{x})$ is canonically identified with $H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\})$. Viewing $[0, 1]$ as the path from v_2 through the trivalent vertex of G to v_1 , we obtain an inclusion $H_2(\Sigma \times [0, 1], \alpha \times \{1\} \cup \beta \times \{0\}) \rightarrow H_2(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\})$. Under these identifications, concatenation $\pi_2(\vec{x}, \vec{x}) \times \pi_2(\vec{x}, \vec{y}, \vec{z}) \rightarrow \pi_2(\vec{x}, \vec{y}, \vec{z})$ corresponds to addition in $H_2(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\})$. Similar remarks apply to \vec{y} and \vec{z} , of course.

We have:

Lemma 10.2 (Compare [21, Proposition 8.2 and Proposition 8.3]) *There is a natural short exact sequence*

$$0 \rightarrow \mathbb{Z} \rightarrow H_2(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\}) \rightarrow H_2(X_{\alpha, \beta, \gamma}) \rightarrow 0.$$

The basepoint $3 \in \Sigma \setminus (\alpha \cup \beta \cup \gamma)$ determines a splitting

$$n_3: H_2(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\}) \rightarrow \mathbb{Z}$$

of this sequence.

Proof From the long exact sequence for the pair $(X_{\alpha, \beta, \gamma}, \Sigma \times \bar{T})$ we have:

$$H_2(\Sigma \times T) \rightarrow H_2(X_{\alpha, \beta, \gamma}) \rightarrow H_2(X_{\alpha, \beta, \gamma}, \Sigma \times \bar{T}) \rightarrow H_1(\Sigma \times \bar{T}).$$

The first map is trivial since Σ is null–homologous in $X_{\alpha, \beta, \gamma}$ (it bounds in U_α , for example). Since the boundary map in the long exact sequence for the pair (U_α, Σ) takes $H_2(U_\alpha, \Sigma)$ one–to–one onto $H_1(\alpha) \subset H_1(\Sigma)$, and similarly for β and γ , the kernel of the last map is the same as $\ker(H_1(\alpha) \oplus H_1(\beta) \oplus H_1(\gamma) \rightarrow H_1(\Sigma))$. (Here, the map is induced from including α, β , and γ in Σ .) Thus, $H_2(X_{\alpha, \beta, \gamma})$ is isomorphic to this kernel.

From the long exact sequence for the pair $(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\})$ we have

$$0 \rightarrow H_2(\Sigma \times G) \rightarrow H_2(\Sigma \times G, \alpha \times \{v_1\} \cup \beta \times \{v_2\} \cup \gamma \times \{v_3\}) \rightarrow H_1(\alpha) \oplus H_1(\beta) \oplus H_1(\gamma) \rightarrow H_1(\Sigma \times G)$$

The kernel of the last map is $\ker(H_1(\alpha) \oplus H_1(\beta) \oplus H_1(\gamma) \rightarrow H_1(\Sigma)) \cong H_2(X_{\alpha,\beta,\gamma})$. $H_2(\Sigma \times G) \cong \mathbb{Z}$. The statement about the splitting is clear. \square

The previous lemma tells us what $\pi_2(\vec{x}, \vec{y}, \vec{z})$ is if nonempty. It is worth knowing when $\pi_2(\vec{x}, \vec{y}, \vec{z})$ is in fact empty. Define $\epsilon(\vec{x}, \vec{y}, \vec{z})$ as follows. Choose a chain p_α (respectively p_β, p_γ) in α (respectively β, γ) with $\partial p_\alpha = \vec{x} - \vec{y}$ (respectively $\partial p_\beta = \vec{y} - \vec{z}, \partial p_\gamma = \vec{z} - \vec{x}$). Then $\epsilon(\vec{x}, \vec{y}, \vec{z})$ is the image of $p_\alpha + p_\beta + p_\gamma$ under the map

$$H_1(\Sigma) \rightarrow \frac{H_1(\Sigma)}{H_1(\alpha) + H_1(\beta) + H_1(\gamma)} \cong H_1(X_{\alpha,\beta,\gamma}).$$

Lemma 10.3 (Compare [21, Proposition 8.3]) *The set $\pi_2(\vec{x}, \vec{y}, \vec{z})$ is nonempty if and only if $\epsilon(\vec{x}, \vec{y}, \vec{z}) = 0$.*

Proof If $\pi_2(\vec{x}, \vec{y}, \vec{z})$ is nonempty, choose an element $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$. View A as a chain in Σ . Then the boundary of A is a chain which defines $\epsilon(\vec{x}, \vec{y}, \vec{z})$, and is obviously zero in homology.

Conversely, if $\epsilon(\vec{x}, \vec{y}, \vec{z})$ is zero then we can choose p_α, p_β and p_γ to be cellular chains (with respect to the cellulation of Σ induced by α, β and γ) so that $p_\alpha + p_\beta + p_\gamma$ is null-homologous. Any chain with boundary $p_\alpha + p_\beta + p_\gamma$ is an element of $\pi_2(\vec{x}, \vec{y}, \vec{z})$. \square

10.1.2 $\text{Spin}^{\mathbb{C}}$ -structures and triangles While each intersection point in a Heegaard diagram specifies a $\text{Spin}^{\mathbb{C}}$ -structure on the underlying 3-manifold, for a Heegaard triple-diagram it is the elements of $\pi_2(\vec{x}, \vec{y}, \vec{z})$ that specify $\text{Spin}^{\mathbb{C}}$ -structures on the corresponding 4-manifold. This correspondence is somewhat more complicated than in the 3-manifold case. Our exposition will be very close to that in [21, Section 8.1.4], but the reader may find some points clearer in one treatment or the other.

Recall (Section 2) that the definition of a $\text{Spin}^{\mathbb{C}}$ -structure on a 3-manifold which we have used is a “homology class of nonvanishing vector fields.” On a 4-manifold, the analogous definition is:

Definition 10.4 Fix a connected 4–manifold M . Suppose that J_1 and J_2 are almost complex structures on M , defined in the complement of some 4–ball in M . (We will say that J_1 and J_2 are *almost defined*.) We say that J_1 and J_2 are *homologous* if J_1 and J_2 are isotopic in the complement of some (larger) 4–ball in M . We define a $\text{Spin}^{\mathbb{C}}$ –structure on M to be a homology class of almost defined almost complex structures.

We sketch the identification with the standard definition of $\text{Spin}^{\mathbb{C}}$ –structures. Suppose we are given an almost complex J structure defined in the complement of some 4–ball B . On the complement of B , the almost complex structure J determines canonically a $\text{Spin}^{\mathbb{C}}$ lifting of the bundle of frames. The obstruction to extending the $\text{Spin}^{\mathbb{C}}$ –structure over B lies in $H^3(B, \partial B) = 0$ and the collection of distinct extensions correspond to $H^2(B, \partial B) = 0$.

Conversely, a $\text{Spin}^{\mathbb{C}}$ lifting of the bundle of frames determines complex positive and negative spinor bundles. Choosing a section s of the positive spinor bundle vanishing at a finite number of points, Clifford multiplication by s gives an isomorphism of the negative spinor bundle with TM away from a finite number of points. Choose a ball B containing these points. Then the identification of TM with the negative spinor bundle determines an almost complex structure on $M \setminus B$.

Now, fix a homology class $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$. Between here and Lemma 10.5 we associate a $\text{Spin}^{\mathbb{C}}$ –structure on $X_{\alpha, \beta, \gamma}$ to A and some extra data. (It will turn out that the $\text{Spin}^{\mathbb{C}}$ –structure does not depend on the extra data.)

Fix a height function f_{α} (respectively f_{β} , f_{γ}) on the handlebody U_{α} (respectively U_{β} , U_{γ}) with one index 0 critical point and g index 1 critical points, such that $f_{\alpha}|_{\partial U_{\alpha}}$ (respectively $f_{\beta}|_{\partial U_{\beta}}$, $f_{\gamma}|_{\partial U_{\gamma}}$) is constant.

Choose a smooth immersion $\phi: S \rightarrow W_{\alpha, \beta, \gamma}$ representing A . We will place some requirements on ϕ presently.

Let F denote the surface obtained by capping off $\phi(S) \cup \{3\} \times T$ with the downward gradient flows of f_{α} , f_{β} and f_{γ} . So, F is an immersed surface with boundary on the $3g + 3$ critical lines in $X_{\alpha, \beta, \gamma}$. (By a critical line we mean a line of the form (critical point of f_i) $\times e_i$.)

Let \mathcal{L} be the 2–plane field on $X_{\alpha, \beta, \gamma} \setminus F$ given by

- $\mathcal{L}(p) = T(\{p_1\} \times T) \subset T_{(p_1, p_2)}\Sigma \times T$ for $p = (p_1, p_2) \in \Sigma \times T$
- $\mathcal{L}(p) = \ker df_{\alpha}(p)$ (respectively $\ker df_{\beta}(p)$, $\ker df_{\gamma}(p)$) for $p \in U_{\alpha} \times e_1$ (respectively $p \in U_{\beta} \times e_2$, $p \in U_{\gamma} \times e_3$).




Figure 8: Foliation of T

To use \mathcal{L} to define a $\text{Spin}^{\mathbb{C}}$ -structure, we need to extend it further. Fix a point $x \in T$ and line segments a , b and c from x to the edges e_1 , e_2 and e_3 of \mathcal{L} , respectively. Let $v_{ij} = e_i \cap e_j$. Let $\ell_{\alpha,\beta}(t)$, $\ell_{\beta,\gamma}(t)$, and $\ell_{\alpha,\gamma}(t)$ denote the foliations of $T \setminus (a \cup b \cup c)$, parametrized by $(0, 1)$, so that as $t \rightarrow 0$, $\ell_{\alpha,\beta}(t)$ degenerates to the corner v_{12} and as $t \rightarrow 1$ $\ell_{\alpha,\beta}(t)$ degenerates to $a \cup b$. See Figure 8. The map π_T extends in an obvious way to a map $\pi_T: X_{\alpha,\beta,\gamma} \rightarrow T$. Let $\tilde{\ell}_{\alpha,\beta}(t) = \pi_T^{-1}(\ell_{\alpha,\beta}(t)) \subset X_{\alpha,\beta,\gamma}$, $\tilde{\ell}_{\beta,\gamma}(t) = \pi_T^{-1}(\ell_{\beta,\gamma}(t)) \subset X_{\alpha,\beta,\gamma}$, and $\tilde{\ell}_{\alpha,\gamma}(t) = \pi_T^{-1}(\ell_{\alpha,\gamma}(t)) \subset X_{\alpha,\beta,\gamma}$.

Choose ϕ so that:

- (1) The intersection of F with each $\tilde{\ell}_{\alpha,\beta}(t)$, $\tilde{\ell}_{\beta,\gamma}(t)$ or $\tilde{\ell}_{\alpha,\gamma}(t)$ is a finite disjoint union of contractible 1-complexes.
- (2) For all but finitely many t , $F \cap \tilde{\ell}_{\alpha,\beta}(t)$ consists of $g + 1$ disjoint embedded arcs.
- (3) In some small neighborhood of the corner v_{12} (respectively v_{23} , v_{13}) of T , ϕ agrees with $\vec{x} \times T \subset \Sigma \times T$ (respectively $\vec{y} \times T \subset \Sigma \times T$, $\vec{z} \times T \subset \Sigma \times T$).
- (4) The intersections $F \cap \pi_T^{-1}(a)$, $F \cap \pi_T^{-1}(b)$ and $F \cap \pi_T^{-1}(c)$ each consist of $g + 1$ disjoint embedded arcs.
- (5) The preimage under ϕ of $\alpha_i \times e_1$, $\beta_i \times e_2$ and $\gamma_i \times e_3$ is a connected arc for each i .

Such ϕ exist.

Fix a t such that $F \cap \tilde{\ell}_{\alpha,\beta}(t)$ consists of $g + 1$ disjoint embedded arcs. Observe that $\tilde{\ell}_{\alpha,\beta}(t)$ is diffeomorphic to $Y_{\alpha,\beta}$. For an appropriate choice of the height functions f_α and f_β , they and a parameter for the interval $\ell_{\alpha,\beta}(t)$ determine a Morse function f_t on $Y_{\alpha,\beta}$. For an appropriate choice of metric, the 2-plane field \mathcal{L} is the orthogonal complement of ∇f_t . As when we associated $\text{Spin}^{\mathbb{C}}$ -structures on 3-manifolds to intersection points, one can then use the $(g + 1)$ -tuple of paths $F \cap \tilde{\ell}_{\alpha,\beta}(t)$ to replace \mathcal{L} with a 2-plane field defined on all of $\tilde{\ell}_{\alpha,\beta}(t)$.

Doing this construction uniformly in t , we can extend \mathcal{L} over all of $\tilde{\ell}_{\alpha,\beta}(t)$ for all t such that $F \cap \tilde{\ell}_{\alpha,\beta}(t)$ consists of $g + 1$ disjoint arcs. The same construction obviously works for $\tilde{\ell}_{\beta,\gamma}(t)$ and $\tilde{\ell}_{\alpha,\gamma}(t)$.

Note that we can also extend \mathcal{L} over the boundary $Y_{\alpha,\beta} \cup Y_{\beta,\gamma} \cup Y_{\alpha,\gamma}$, by exactly the same method.

Now we have defined \mathcal{L} except on the intersection of F with

- $\pi_T^{-1}(a \cup b \cup c)$ and
- $\tilde{\ell}_{\alpha,\beta}(t_i)$, $\tilde{\ell}_{\beta,\gamma}(t'_i)$ and $\tilde{\ell}_{\alpha,\gamma}(t''_i)$ for some finite collection of t_i , t'_i and t''_i .

Thus the region to which we have not extended \mathcal{L} consists of a collection of disjoint contractible 1-complexes. So, we can find an open ball B in $X_{\alpha,\beta,\gamma}$ such that \mathcal{L} is defined on $X_{\alpha,\beta,\gamma} \setminus B$.

Choose a metric on $X_{\alpha,\beta,\gamma}$. Then the metric, orientation, and 2-plane field \mathcal{L} define an almost complex structure on $X_{\alpha,\beta,\gamma} \setminus B$, and hence a $\text{Spin}^{\mathbb{C}}$ -structure on $X_{\alpha,\beta,\gamma}$.

The first question we address is how this construction depends on ϕ .

Lemma 10.5 *The $\text{Spin}^{\mathbb{C}}$ -structure just constructed depends only on the restriction of ϕ to the boundary of S .*

Proof Observe that in the construction, the restriction of ϕ to the boundary of S determined the $\text{Spin}^{\mathbb{C}}$ -structure on $X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T}$. For a manifold M , let $\text{Spin}^{\mathbb{C}}(M)$ denote the collection of $\text{Spin}^{\mathbb{C}}$ -structures on M . We check that the restriction map $\text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma}) \rightarrow \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T})$ is injective; this then proves the result.

Recall that $\text{Spin}^{\mathbb{C}}(M)$ is an affine copy of $H^2(M; \mathbb{Z})$. Further, if $N \subset M$ then the restriction map from $\text{Spin}^{\mathbb{C}}(M)$ to $\text{Spin}^{\mathbb{C}}(N)$ commutes with the H^2 -action. That is, let $\xi \in \text{Spin}^{\mathbb{C}}(M)$, $\xi|_N$ its restriction to $\text{Spin}^{\mathbb{C}}(N)$, $a \in H^2(M; \mathbb{Z})$ and $a|_N$ its restriction to $H^2(N; \mathbb{Z})$. Then, $(a \cdot \xi)|_N = a|_N \cdot \xi|_N$. This can be proved, for example, by thinking about Čech cocycles.

So, it suffices to prove that the restriction map $H^2(X_{\alpha,\beta,\gamma}; \mathbb{Z}) \rightarrow H^2(X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T}; \mathbb{Z})$ is injective. Consider the following commutative diagram:

$$\begin{array}{ccccc}
 H^2(X_{\alpha,\beta,\gamma}, X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T}; \mathbb{Z}) & \longrightarrow & H^2(X_{\alpha,\beta,\gamma}; \mathbb{Z}) & \longrightarrow & H^2(X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T}; \mathbb{Z}) . \\
 \downarrow \cong & & \uparrow & & \\
 H^2(\Sigma \times \overline{\mathbb{D}}, \Sigma \times \partial \overline{\mathbb{D}}; \mathbb{Z}) & \xrightarrow{0} & H^2(\Sigma \times \overline{\mathbb{D}}; \mathbb{Z}) & \hookrightarrow & H^2(\Sigma \times \partial \overline{\mathbb{D}}; \mathbb{Z})
 \end{array}$$

(The top row is from the long exact sequence for the pair $(X_{\alpha,\beta,\gamma}, X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T})$. The bottom is from the long exact sequence for the pair $(\Sigma \times \overline{T}, \Sigma \times T)$.) It follows from the diagram that the map $H^2(X_{\alpha,\beta,\gamma}; \mathbb{Z}) \rightarrow H^2(X_{\alpha,\beta,\gamma} \setminus \Sigma \times \overline{T}; \mathbb{Z})$ is an injection. \square

It would be nice to have a better way of presenting the construction of a $\text{Spin}^{\mathbb{C}}$ -structure on $X_{\alpha,\beta,\gamma}$ from an element of $\pi_2(\vec{x}, \vec{y}, \vec{z})$. Unfortunately I do not know one.

The following is [21, Proposition 8.4]:

Lemma 10.6 *The assignment described above induces a well-defined map*

$$\mathfrak{s}_3: \pi_2(\vec{x}, \vec{y}, \vec{z}) \rightarrow \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma}).$$

Proof By the previous lemma, the construction depends only on the restriction of ϕ to ∂S . This restriction is defined up to isotopy by the element in $\pi_2(\vec{x}, \vec{y}, \vec{z})$. It is clear from the construction that an isotopy of ϕ does not change the $\text{Spin}^{\mathbb{C}}$ -structure constructed. \square

For $\mathfrak{s} \in \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma})$ we shall often write $A \in \mathfrak{s}$ to mean that $\mathfrak{s}_3(A) = \mathfrak{s}$.

Definition 10.7 Given two triples of intersection points $(\vec{x}, \vec{y}, \vec{z})$ and $(\vec{x}', \vec{y}', \vec{z}')$, and $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, $A' \in \pi_2(\vec{x}', \vec{y}', \vec{z}')$ define ψ and ψ' to be $\text{Spin}^{\mathbb{C}}$ -equivalent if there exist elements $B_{\alpha,\beta} \in \pi_2(\vec{x}, \vec{x}')$, $B_{\beta,\gamma} \in \pi_2(\vec{y}, \vec{y}')$, and $B_{\alpha,\gamma} \in \pi_2(\vec{z}, \vec{z}')$ such that $A + B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma} = A'$. We let $S_{\alpha,\beta,\gamma}$ denote the set of $\text{Spin}^{\mathbb{C}}$ -equivalence classes of triangles.

The following is [21, Proposition 8.5]:

Lemma 10.8 *The assignment $\pi_2(\vec{x}, \vec{y}, \vec{z}) \rightarrow \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma})$ defined above descends to a map $\mathfrak{s}_3: S_{\alpha,\beta,\gamma} \rightarrow \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma})$. This new map is injective. The image of \mathfrak{s}_3 consists of all $\text{Spin}^{\mathbb{C}}$ -structures whose restrictions to $\partial X_{\alpha,\beta,\gamma}$ are represented by intersection points.*

Proof Our proof is the same as in [21].

The fact that \mathfrak{s}_3 descends to $S_{\alpha,\beta,\gamma}$ follows from the fact that the restriction of the 2–plane field used to define $\mathfrak{s}_3(A)$ to $Y_{\alpha,\beta}$ is homologous to the 2–plane field used to define $\mathfrak{s}_3(\vec{x})$ (and similarly for the restrictions to $Y_{\beta,\gamma}$ and $Y_{\alpha,\gamma}$). That is, suppose A and A' are $\text{Spin}^{\mathbb{C}}$ –equivalent elements of $\pi_2(\vec{x}, \vec{y}, \vec{z})$, $A' = A + B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$. Let ξ_A and $\xi_{A'}$ denote the 2–plane fields constructed above to define $\mathfrak{s}_3(A)$ and $\mathfrak{s}_3(A')$ respectively. We can assume that ξ_A and $\xi_{A'}$ agree outside some collar neighborhood U of $\partial W_{\alpha,\beta,\gamma}$. Further, on $Y_{\alpha,\beta} \times [0, 1) \subset U$ (respectively $Y_{\beta,\gamma} \times [0, 1) \subset U$, $Y_{\alpha,\gamma} \times [0, 1) \subset U$) we can assume that ξ_A and $\xi_{A'}$ are given by $\nabla f_{\alpha,\beta}^{\perp}$ (respectively $\nabla f_{\beta,\gamma}^{\perp}$, $\nabla f_{\alpha,\gamma}^{\perp}$) for some Morse function $f_{\alpha,\beta}$ on $Y_{\alpha,\beta}$ (respectively $f_{\beta,\gamma}$ on $Y_{\beta,\gamma}$, $f_{\alpha,\gamma}$ on $Y_{\alpha,\gamma}$) outside some ball neighborhoods of $(\vec{x} \cup \mathfrak{z}) \times [0, 1)$ (respectively $(\vec{y} \cup \mathfrak{z}) \times [0, 1)$, $(\vec{z} \cup \mathfrak{z}) \times [0, 1)$). But it is then immediate from the definition that ξ_A and $\xi_{A'}$ define the same $\text{Spin}^{\mathbb{C}}$ –structure on $X_{\alpha,\beta,\gamma}$.

Recall that the restriction map $\text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma}) \rightarrow \text{Spin}^{\mathbb{C}}(\partial X_{\alpha,\beta,\gamma})$ commutes with the (transitive) actions of $H^2(X_{\alpha,\beta,\gamma})$ and $H^2(\partial X_{\alpha,\beta,\gamma})$ respectively. It follows that the cokernel $\text{Spin}^{\mathbb{C}}(\partial X_{\alpha,\beta,\gamma}) / \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma})$ is naturally identified with the cokernel of the restriction map $H^2(X_{\alpha,\beta,\gamma}; \mathbb{Z}) \rightarrow H^2(\partial X_{\alpha,\beta,\gamma}; \mathbb{Z})$. This in turn is identified with the image of the connecting homomorphism $\delta: H^2(\partial X_{\alpha,\beta,\gamma}; \mathbb{Z}) \rightarrow H^3(X_{\alpha,\beta,\gamma}, \partial X_{\alpha,\beta,\gamma}; \mathbb{Z})$.

Summarizing, we have a map $\epsilon': \text{Spin}^{\mathbb{C}}(\partial X_{\alpha,\beta,\gamma}) \rightarrow H^3(X_{\alpha,\beta,\gamma}, \partial X_{\alpha,\beta,\gamma}; \mathbb{Z})$ given by the composition of the coboundary map

$$\delta: H^2(\partial X_{\alpha,\beta,\gamma}; \mathbb{Z}) / H^2(X_{\alpha,\beta,\gamma}; \mathbb{Z}) \rightarrow H^3(X_{\alpha,\beta,\gamma}, \partial X_{\alpha,\beta,\gamma}; \mathbb{Z})$$

with the projection

$$\text{Spin}^{\mathbb{C}}(\partial X_{\alpha,\beta,\gamma}) \rightarrow \text{Spin}^{\mathbb{C}}(\partial X_{\alpha,\beta,\gamma}) / \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma}) = H^2(\partial X_{\alpha,\beta,\gamma}) / H^2(X_{\alpha,\beta,\gamma}).$$

The element $\epsilon'(\mathfrak{s})$ vanishes if and only if \mathfrak{s} can be extended to all of $X_{\alpha,\beta,\gamma}$.

Recall that we defined earlier in this section the obstruction $\epsilon(\vec{x}, \vec{y}, \vec{z}) \in H_1(X_{\alpha,\beta,\gamma})$ to the existence of elements in $\pi_2(\vec{x}, \vec{y}, \vec{z})$. Let $\mathfrak{s}_3(\vec{x}, \vec{y}, \vec{z})$ denote the $\text{Spin}^{\mathbb{C}}$ –structure induced by \vec{x} , \vec{y} and \vec{z} on $\partial X_{\alpha,\beta,\gamma}$. We next check that $\epsilon(\vec{x}, \vec{y}, \vec{z}) = PD(\epsilon'(\mathfrak{s}_3(\vec{x}, \vec{y}, \vec{z})))$, where PD denotes the Poincaré dual.

To show this, isotope the α –, β – and γ –circles so that there are intersection points \vec{x}' , \vec{y}' and \vec{z}' with $\pi_2(\vec{x}', \vec{y}', \vec{z}') \neq \emptyset$ (this is easy). Then, $\epsilon(\vec{x}', \vec{y}', \vec{z}') = 0$. We already showed that the $\text{Spin}^{\mathbb{C}}$ –structure $\mathfrak{s}_3(\vec{x}', \vec{y}', \vec{z}')$ extends to all of $X_{\alpha,\beta,\gamma}$, so $\epsilon'(\mathfrak{s}_3(\vec{x}', \vec{y}', \vec{z}')) = 0$. It is obvious from the definitions that, up to a universal sign, $\epsilon(\vec{x}, \vec{y}, \vec{z}) - \epsilon(\vec{x}', \vec{y}', \vec{z}') = i(\epsilon(\vec{x}, \vec{x}')) + i(\epsilon(\vec{y}, \vec{y}')) + i(\epsilon(\vec{z}, \vec{z}'))$. (Here, $i: H_1(\partial X_{\alpha,\beta,\gamma})$

→ $H_1(X_{\alpha,\beta,\gamma})$ is the map induced by inclusion.) Naturality of Poincaré duality gives the commutative diagram

$$\begin{array}{ccc} H_1(\partial X_{\alpha,\beta,\gamma}) & \xrightarrow{i} & H_1(X_{\alpha,\beta,\gamma}) \\ \uparrow PD & & \uparrow PD \\ H^2(\partial X_{\alpha,\beta,\gamma}) & \xrightarrow{\delta} & H^3(X_{\alpha,\beta,\gamma}, \partial X_{\alpha,\beta,\gamma}) \end{array}$$

which implies that

$$PD(\epsilon(\vec{x}, \vec{y}, \vec{z})) - PD(\epsilon(\vec{x}', \vec{y}', \vec{z}')) = \delta(PD(\epsilon(\vec{x}, \vec{x}') \oplus \epsilon(\vec{y}, \vec{y}') \oplus \epsilon(\vec{z}, \vec{z}'))).$$

From its definition,

$$\begin{aligned} \epsilon'(\mathfrak{s}_3(\vec{x}, \vec{y}, \vec{z})) - \epsilon'(\mathfrak{s}_3(\vec{x}', \vec{y}', \vec{z}')) &= \delta((\mathfrak{s}_3(\vec{x}) - \mathfrak{s}_3(\vec{x}')) \oplus (\mathfrak{s}_3(\vec{y}) - \mathfrak{s}_3(\vec{y}')) \\ &\quad \oplus (\mathfrak{s}_3(\vec{z}) - \mathfrak{s}_3(\vec{z}'))). \end{aligned}$$

Now, $PD(\epsilon(\vec{x}, \vec{x}')) = \mathfrak{s}_3(\vec{x}) - \mathfrak{s}_3(\vec{x}')$ (and similarly for \vec{y} and \vec{z}). So, since

$$PD(\epsilon(\vec{x}', \vec{y}', \vec{z}')) = \epsilon'(\mathfrak{s}_3(\vec{x}', \vec{y}', \vec{z}')),$$

it follows that $PD(\epsilon(\vec{x}, \vec{y}, \vec{z})) = \epsilon'(\vec{x}, \vec{y}, \vec{z})$.

Now suppose we had a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on $X_{\alpha,\beta,\gamma}$ whose restriction to the boundary is realized by intersection points \vec{x} , \vec{y} , and \vec{z} . Then $0 = \epsilon'(\vec{x}, \vec{y}, \vec{z}) = \epsilon(\vec{x}, \vec{y}, \vec{z})$. It follows that \mathfrak{s} is in the image of \mathfrak{s}_3 . The converse is obvious.

Finally, it remains to show injectivity. Fix a homology class $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$. The following diagram commutes:

$$\begin{array}{ccccccc} \hat{\pi}_2(\vec{x}, \vec{x}) \oplus \hat{\pi}_2(\vec{y}, \vec{y}) \oplus \hat{\pi}_2(\vec{z}, \vec{z}) & \xrightarrow{+A} & \hat{\pi}_2(\vec{x}, \vec{y}, \vec{z}) & \longrightarrow & S_{\alpha,\beta,\gamma} & \longrightarrow & 0 \\ \downarrow \cong & & \downarrow -A \cong & & \downarrow [-A] & \searrow \mathfrak{s}_3 & \searrow \\ & & & & & & \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma}) \\ H_2(\partial X_{\alpha,\beta,\gamma}) & \longrightarrow & H_2(X_{\alpha,\beta,\gamma}) & \longrightarrow & H_2(X_{\alpha,\beta,\gamma}, \partial X_{\alpha,\beta,\gamma}) & \xrightarrow{-\mathfrak{s}_3(A)} & \cong \end{array}$$

Injectivity of \mathfrak{s}_3 is immediate. □

10.2 Definitions of the moduli spaces and maps

We now deal with the analysis involved in defining the triangle maps. We start by discussing the almost complex structures with which we will work. Fix a point \mathfrak{z}_i in each

component of $\Sigma \setminus (\alpha \cup \beta \cup \gamma)$. Fix an almost complex structure $J_{\alpha, \beta}$ (respectively $J_{\beta, \gamma}$, $J_{\alpha, \gamma}$) on $W_{\alpha, \beta}$ (respectively $W_{\beta, \gamma}$, $W_{\alpha, \gamma}$) satisfying **(J1)**–**(J5)** and which achieves transversality for holomorphic curves of index ≤ 1 . We will work with complex structures J on $W_{\alpha, \beta, \gamma}$ such that:

- (J'1)** J is tamed by η , the split symplectic form on $\Sigma \times T$.
- (J'2)** In a neighborhood $U_{\{3i\}}$ of $\{3i\} \times T$, $J = j_\Sigma \times j_T$ is split. (Here, $U_{\{3i\}}$ is small enough that its closure does not intersect $(\alpha \cup \beta \cup \gamma) \times [0, 1] \times T$).
- (J'3)** Near $\Sigma \times \{v_1\}$, J agrees with $J_{\alpha, \beta}$. Similarly, J agrees with $J_{\beta, \gamma}$ near $\Sigma \times \{v_2\}$ and with $J_{\alpha, \gamma}$ near $\Sigma \times \{v_3\}$.
- (J'4)** Projection π_T onto T is holomorphic and each fiber of π_Σ is holomorphic.

For now, fix an almost complex structure J satisfying **(J'1)**–**(J'4)**. For $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, let \mathcal{M}^A denote the moduli space of embedded holomorphic curves $u: (S, \partial S) \rightarrow (W_{\alpha, \beta, \gamma}, \alpha \times e_1 \cup \beta \times e_2 \cup \gamma \times e_3)$ asymptotic to \vec{x} , \vec{y} and \vec{z} at the three ends of $\Sigma \times T$ and in the homology class A . We require that u map exactly one component of ∂S to each of the 3g cylinders $\alpha_i \times e_1$, $\beta_i \times e_2$ and $\gamma_i \times e_3$. We also require that there be no components of S on which $\pi_T \circ u$ is constant.

We digress briefly to discuss the index of the $\bar{\partial}$ -operator for triangles. Fix a homology class $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$. Suppose that $u: S \rightarrow W_{\alpha, \beta, \gamma}$ is a map in the homology class A . Then an argument similar to the one given in the beginning of Section 4 shows that

$$\text{ind}(D\bar{\partial})(u) = \frac{1}{2}g - \chi(S) + 2e(A).$$

Here, $e(A)$ denotes the Euler measure of A , as described in Section 4. The $\frac{1}{2}g$ looks strange, but for g odd it is easy to see that $2e(A)$ is a half-integer. When deriving this formula, one should keep in mind an extra $-g$, not appearing in Section 4, coming from the Maslov index of $\pi_T \circ u$.

Analogously to Subsection 4.2 we have:

Proposition 10.9 *The Euler characteristic of an embedded holomorphic curve $u: S \rightarrow W_{\alpha, \beta, \gamma}$ is determined by the homology class of u .*

Proof The only part of the index which does not depend *a priori* only on A is $\chi(S)$. As in the proof of Proposition 4.2 we will re-interpret the Euler characteristic as an intersection number.

Let $u: S \rightarrow W_{\alpha, \beta, \gamma}$ be a holomorphic map in the homology class A . By the Riemann–Hurwitz formula, the degree of branching of $\pi_\Sigma \circ u$ determines $\chi(S)$, and vice versa.

Choose a diffeomorphism $T \rightarrow \overline{\mathbb{D}^2} \setminus \{z^3 = 1\}$ from T to the unit disk with three boundary punctures. Let ξ_T be the vector field on T induced by the vector field $r \frac{\partial}{\partial \theta}$ on \mathbb{D}^2 . Here, we choose the diffeomorphism so that the preimage of 0 is not a branch point of $\pi_T \circ u$. Let ξ be the vector field $0 \times \xi_T$ on $W_{\alpha, \beta, \gamma}$. Then, for ϵ small, the degree of branching of $\pi_\Sigma \circ u$ is given by $\#(u \cap \exp_{\epsilon \xi}(u)) - g$.

Fix embedded holomorphic curves $u: S \rightarrow W_{\alpha, \beta, \gamma}$, $u': S' \rightarrow W_{\alpha, \beta, \gamma}$ in the homology class A . Because we are in a low dimension (two), we can choose a (proper) bordism $v: R \rightarrow W_{\alpha, \beta, \gamma}$ from u to u' . Choose a Morse function $f: R \rightarrow \mathbb{R}$ with $\partial R = f^{-1}(\{0, 1\})$, and so that v restricts to u on $f^{-1}(0)$ and to u' on $f^{-1}(1)$. We will think of f as a coordinate on R , and write R_a for $f^{-1}(a)$ and $v_a = v|_{R_a}$.

For an appropriate choice of R , v and f there is a partition $0 = t_0 < s_1 < t_1 < s_2 < t_2 < \dots < s_k = 1$ of $[0, 1]$ such that

- The function f has no critical values in $[s_i, t_i]$, $i = 1, \dots, k$ (and so $f^{-1}([s_i, t_i])$ is a product).
- For each point s_i (respectively t_i), v_{s_i} (respectively v_{t_i}) is an embedding in the homology class A which projects as a branched cover to T .
- The map $v|_{f^{-1}((t_i, s_{i+1}))}$ is constant near infinity (as a function of $f(p)$).

In words, we have chosen R , v and f so that we can subdivide $[0, 1]$ into subintervals over which either the map doesn't change near the punctures or the topology of S doesn't change near the punctures (and so that each interval starts and ends with maps for which the degree of branching of $\pi_\Sigma \circ v_a$ determines $\chi(R_a)$).

For each i , $\chi(R_{s_i}) = \chi(R_{t_i})$, by the first property.

Also, if $a, a', a'' \in (t_i, s_{i+1})$ are regular values then $\#(v_a \cap v_{a'}) = \#(v_a \cap v_{a''})$. It follows from the Riemann–Hurwitz formula that $\chi(R_{t_i}) = \chi(R_{t_{i+1}})$.

But this implies that $\chi(S) = \chi(R_{t_0}) = \chi(R_{s_k}) = \chi(S')$, completing the proof. \square

For $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, let $\text{ind}(A)$ denote the index of the $D\bar{\partial}$ problem for embedded curves in the homology class A , if such an embedded curve exists. Note that the index is additive in the sense that for $B_{\alpha, \beta} \in \pi_2(\vec{x}, \vec{x}')$, $\text{ind}(A + B_{\alpha, \beta}) = \text{ind}(A) + \text{ind}(B_{\alpha, \beta})$ (and similarly for the other ends), if both sides are defined. We shall omit the words “if defined” from all subsequent discussion of the index – that we are only discussing homology classes representable by holomorphic curves will be implicit throughout.

Remark We could have proved Proposition 10.9 in more generality by using an analog of Lemma 4.1. Then, $\text{ind}(A)$ would have a natural meaning for any homology

class $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$. The only cases in which we are interested in the index, however, are when there is an embedded holomorphic curve.

We now must check that, for generic J satisfying **(J'1)**–**(J'4)**, the moduli spaces of J –holomorphic curves \mathcal{M}^A are reasonably well behaved. Again, we achieve transversality by varying J . The argument to show that one can achieve transversality among J satisfying **(J'1)**–**(J'4)** is analogous to the one given in Section 3.

Note that bubbling and Deligne–Mumford type degenerations in moduli spaces with $\text{ind} \leq 1$ are prohibited by the argument used in Section 7. So, by the compactness theorem [2, Theorem 10.2], the compactification of \mathcal{M}^A consists of holomorphic buildings with one story in $W_{\alpha,\beta,\gamma}$ and all of their other stories in $W_{\alpha,\beta}$, $W_{\beta,\gamma}$ or $W_{\alpha,\gamma}$. (As is standard in symplectic field theory, the stories in the cylindrical bordisms $W_{\alpha,\beta}$, $W_{\beta,\gamma}$, and $W_{\alpha,\gamma}$ are only defined up to translation.)

For the rest of this section, fix a $\text{Spin}^{\mathbb{C}}$ equivalence class of triangles. Denote it $\mathfrak{s}_{\alpha,\beta,\gamma}$. Fix a complex structure satisfying **(J'1)**–**(J'4)** and achieving transversality.

There are still some technical details to address before we can define the triangle maps. However, we will give the definitions now, asking the reader to trust that all the symbols make sense and sums are finite. We will justify this trust presently.

Given the choice of $\mathfrak{s}_{\alpha,\beta,\gamma} \in \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma})$ we will define a map

$$f_{\alpha,\beta,\gamma}^{\infty}: CF^{\infty}(Y_{\alpha,\beta}) \otimes_{\mathbb{Z}[U]} CF^{\infty}(Y_{\beta,\gamma}) \rightarrow CF^{\infty}(Y_{\alpha,\gamma})$$

by

$$f_{\alpha,\beta,\gamma}^{\infty}([\vec{x}, i] \otimes [\vec{y}, j]) = \sum_{\vec{z}} \sum_{\substack{A \in \pi_2(\vec{x}, \vec{y}, \vec{z}) \cap \mathfrak{s}_{\alpha,\beta,\gamma} \\ \text{ind}(A)=0}} (\#\mathcal{M}^A)[\vec{z}, i + j - n_z(A)].$$

Remark For the complexes CF^{∞} , CF^{-} , CF^{+} and $CF^{\leq 0}$ our tensor products shall always be over $\mathbb{Z}[U]$. For \widehat{CF} they shall be over \mathbb{Z} . In the case of CF^{∞} it would be equivalent to take the tensor product over $\mathbb{Z}[U, U^{-1}]$. It is not equivalent, and quickly leads to nonsense, to take all the tensor products over \mathbb{Z} . The corresponding remark also applies to the Hom functor, if one wanted to obtain cohomology theories; cf [20, Section 2].

There are two obvious issues that need to be addressed. Firstly, since we have been working with \mathbb{Z} –coefficients, the symbol “#” implies that we have chosen orientations for the \mathcal{M}^A , which should presumably be consistent with the orientations for the moduli spaces for $\mathcal{H}_{\alpha,\beta}$, $\mathcal{H}_{\beta,\gamma}$ and $\mathcal{H}_{\alpha,\gamma}$. We will address this issue in Subsection

10.3. Secondly, we need to know that the coefficient of $[\bar{z}, k]$ in $f_{\alpha,\beta,\gamma}^\infty([\bar{x}, i], [\bar{y}, j])$ is a finite sum. This will require that we impose an admissibility condition on the Heegaard triple–diagram, as we will discuss in Subsection 10.4. Note, however, that if we work with $\mathbb{Z}/2$ –coefficients and if $H_2(X_{\alpha,\beta,\gamma})$ is finite then the formula defining $f_{\alpha,\beta,\gamma}^\infty$ already makes perfect sense.

Before addressing the issues of orientations and admissibility, we define the rest of the maps that will appear.

We will define a map

$$\hat{f}_{\alpha,\beta,\gamma}: \widehat{CF}(Y_{\alpha,\beta}) \otimes_{\mathbb{Z}} \widehat{CF}(Y_{\beta,\gamma}) \rightarrow \widehat{CF}(Y_{\alpha,\gamma})$$

by

$$\hat{f}_{\alpha,\beta,\gamma}(\bar{x} \otimes \bar{y}) = \sum_{\bar{z}} \sum_{\substack{A \in \hat{\pi}_2(\bar{x}, \bar{y}, \bar{z}) \cap \mathfrak{s}_{\alpha,\beta,\gamma} \\ \text{ind}(A)=0}} (\#\mathcal{M}^A) \bar{z}.$$

There is a subcomplex $CF^{\leq 0}$ of CF^∞ generated by all $[\bar{x}, i]$ with $i \leq 0$. The homology $HF^{\leq 0}$ of $CF^{\leq 0}$ is naturally isomorphic to HF^- , but the results on triangles are phrased most simply in terms of $CF^{\leq 0}$. The map $f_{\alpha,\beta,\gamma}^\infty$ restricts to a map $f_{\alpha,\beta,\gamma}^{\leq 0}: CF^{\leq 0}(Y_{\alpha,\beta}) \otimes CF^{\leq 0}(Y_{\beta,\gamma}) \rightarrow CF^{\leq 0}(Y_{\alpha,\gamma})$. Hence it also induces a map

$$f_{\alpha,\beta,\gamma}^+: CF^+(Y_{\alpha,\beta}) \otimes CF^{\leq 0}(Y_{\beta,\gamma}; M_{\beta,\gamma}) \rightarrow CF^+(Y_{\alpha,\gamma}).$$

Note that $f_{\alpha,\beta,\gamma}^\infty$ is a map of $\mathbb{Z}[U]$ –modules (in fact, $\mathbb{Z}[U, U^{-1}]$ –modules), so $f_{\alpha,\beta,\gamma}^{\leq}$ and $f_{\alpha,\beta,\gamma}^+$ are also maps of $\mathbb{Z}[U]$ –modules. Further, since $CF^{\leq 0}(Y_{\beta,\gamma})$ is a subcomplex of $CF^\infty(Y_{\beta,\gamma})$, the map $f_{\alpha,\beta,\gamma}^\infty$ restricts to a map $CF^\infty(Y_{\alpha,\beta}) \otimes_{\mathbb{Z}[U]} CF^{\leq 0}(Y_{\beta,\gamma}) \rightarrow CF^\infty(Y_{\alpha,\gamma})$, which we also denote $f_{\alpha,\beta,\gamma}^\infty$.

The next step is to show that the maps just defined are chain maps. The proof is completely standard, but before giving it we digress to deal with orientations and admissibility.

10.3 Orientations

By the same arguments as in Section 6, it follows that the moduli space $\mathcal{M}^A(X_{\alpha,\beta,\gamma})$ are orientable, and by general arguments we can find orientations for all the $\mathcal{M}^A(X_{\alpha,\beta,\gamma})$, $\mathcal{M}^{B_{\alpha,\beta}}(Y_{\alpha,\beta})$, $\mathcal{M}^{B_{\beta,\gamma}}(Y_{\beta,\gamma})$ and $\mathcal{M}^{B_{\alpha,\gamma}}(Y_{\alpha,\gamma})$ (or rather, for the determinant line bundles over the corresponding configuration spaces) which are consistent with gluings. However, we would like somewhat more. Specifically:

Lemma 10.10 (Compare [21, Lemma 8.7]) *Given coherent orientation systems $\mathfrak{o}_{\alpha,\beta}(B_{\alpha,\beta})$ and $\mathfrak{o}_{\beta,\gamma}(B_{\beta,\gamma})$ (for all $B_{\alpha,\beta}$ and $B_{\beta,\gamma}$) there are orientation systems $\mathfrak{o}_{\alpha,\gamma}$ and $\mathfrak{o}_{\alpha,\beta,\gamma}$ consistent with $\mathfrak{o}_{\alpha,\beta}$ and $\mathfrak{o}_{\beta,\gamma}$.*

Note that we have not claimed that $\mathfrak{o}_{\alpha,\gamma}$ and $\mathfrak{o}_{\alpha,\beta,\gamma}$ are unique. The indeterminacy will be clear from the proof.

Proof of Lemma 10.10 Our proof is the same as in [21, Section 8.2].

Fix a $\text{Spin}^{\mathbb{C}}$ -structure $\mathfrak{s}_{\alpha,\beta,\gamma}$ on $X_{\alpha,\beta,\gamma}$ and intersection points $\vec{x}_0 \in \mathfrak{s}_{\alpha,\beta,\gamma}|_{Y_{\alpha,\beta}}$, $\vec{y}_0 \in \mathfrak{s}_{\alpha,\beta,\gamma}|_{Y_{\beta,\gamma}}$ and $\vec{z}_0 \in \mathfrak{s}_{\alpha,\beta,\gamma}|_{Y_{\alpha,\gamma}}$. Fix $A_0 \in \pi_2(\vec{x}_0, \vec{y}_0, \vec{z}_0)$. Choose any orientation over A_0 .

Let

$$K = \{B_{\alpha,\gamma} \in \hat{\pi}_2(\vec{z}_0, \vec{z}_0) \mid \exists B_{\alpha,\beta} \in \hat{\pi}_2(\vec{x}_0, \vec{x}_0), B_{\beta,\gamma} \in \hat{\pi}_2(\vec{y}_0, \vec{y}_0) \\ \text{such that } A_0 + B_{\alpha,\gamma} = A_0 + B_{\beta,\gamma} + B_{\alpha,\beta}\}.$$

Sublemma 10.11 (1) $\hat{\pi}_2(\vec{z}_0, \vec{z}_0) \cong K \oplus \mathbb{Z}^N$ for some N .

(2) Given $B_{\alpha,\gamma} \in K$ there is only one pair $B_{\beta,\gamma}, B_{\alpha,\beta}$ such that $A_0 + B_{\alpha,\gamma} = A_0 + B_{\beta,\gamma} + B_{\alpha,\beta}$.

Assuming the sublemma, the lemma is almost immediate. By Part 2 of the sublemma, $\mathfrak{o}_{\alpha,\beta,\gamma}(A_0)$, $\mathfrak{o}_{\alpha,\beta}$, and $\mathfrak{o}_{\beta,\gamma}$ determine $\mathfrak{o}_{\alpha,\gamma}$ over K . Choosing $\mathfrak{o}_{\alpha,\gamma}$ arbitrarily over a basis of \mathbb{Z}^N determines $\mathfrak{o}_{\alpha,\gamma}$ over $\hat{\pi}_2(\vec{z}_0, \vec{z}_0)$. The orientation over $[\Sigma]$ is determined as in Section 6. Choosing a homology class $B_{\alpha,\gamma,\vec{z}} \in \pi_2(\vec{z}_0, \vec{z})$ for each intersection point \vec{z} (ie, a complete set of paths for $Y_{\alpha,\gamma}$ in the sense of Section 6) and then choosing orientations arbitrarily over the $B_{\alpha,\gamma,\vec{z}}$ determines $\mathfrak{o}_{\alpha,\gamma}$ over all of $\mathfrak{s}_{\alpha,\beta,\gamma}|_{Y_{\alpha,\gamma}}$. The orientation over A_0 and $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma}$ together determine $\mathfrak{o}_{\alpha,\beta,\gamma}$ over all of $\mathfrak{s}_{\alpha,\beta,\gamma}$. This completes the proof, except for the

Proof of Sublemma 10.11 The subgroup K is canonically identified with the intersection of $H_2(Y_{\alpha,\beta} \cup Y_{\beta,\gamma})$ and $H_2(Y_{\alpha,\gamma})$ in $H_2(X_{\alpha,\beta,\gamma})$. From the fragment

$$H_3(X_{\alpha,\beta,\gamma}, Y_{\alpha,\beta} \cup Y_{\beta,\gamma}) \rightarrow H_2(Y_{\alpha,\beta} \cup Y_{\beta,\gamma}) \xrightarrow{j} H_2(X_{\alpha,\beta,\gamma}) \xrightarrow{p} H^2(X_{\alpha,\beta,\gamma}, Y_{\alpha,\beta} \cup Y_{\beta,\gamma})$$

of the long exact sequence for the pair $(X_{\alpha,\beta,\gamma}, Y_{\alpha,\beta} \cup Y_{\beta,\gamma})$, one sees $\hat{\pi}_2(\vec{z}_0, \vec{z}_0)/K \cong p \circ i(H_2(Y_{\alpha,\gamma}))$. By excision, $H_*(X_{\alpha,\beta,\gamma}, Y_{\alpha,\beta} \cup Y_{\beta,\gamma}) \cong H_*(U_\beta \times [0, 1], U_\beta \times \partial[0, 1])$.

So, $H_3(X_{\alpha,\beta,\gamma}, Y_{\alpha,\beta} \cup Y_{\beta,\gamma}) = 0$, so j is injective, implying part (2) of the sublemma. Further,

$$\begin{aligned} H_2(X_{\alpha,\beta,\gamma}, Y_{\alpha,\beta} \cup Y_{\beta,\gamma}) &\cong H_2(U_\beta \times [0, 1], U_\beta \times \partial[0, 1]) \\ &\cong H_2\left(\left(\bigvee_{i=1}^g S^1\right) \times S^1, \left(\bigvee_{i=1}^g S^1\right) \times \{\text{pt}\}\right) \\ &\cong \mathbb{Z}^g \end{aligned}$$

(from the long exact sequence for the pair), so $\hat{\pi}_2(\vec{z}_0, \vec{z}_0)/K$ is free Abelian. It follows that the sequence

$$0 \rightarrow K \rightarrow \hat{\pi}_2(\vec{z}_0, \vec{z}_0) \rightarrow \hat{\pi}_2(\vec{z}_0, \vec{z}_0)/K \rightarrow 0$$

splits, yielding the result. □

In the rest of this section, we shall always assume that coherent orientations have been chosen for the moduli spaces under consideration, but shall suppress the orientation systems from the notation.

10.4 Admissibility

As when we defined the chain complexes, we will need the Heegaard triple–diagram to satisfy certain admissibility criteria in order to ensure finiteness when we define maps between Floer homologies.

Definition 10.12 The pointed Heegaard triple–diagram \mathcal{H}^3 is *weakly admissible* if the following condition is met. For any $B_{\alpha,\beta} \in \hat{\pi}_2(\vec{x}, \vec{x})$, $B_{\beta,\gamma} \in \hat{\pi}_2(\vec{y}, \vec{y})$, and $B_{\alpha,\gamma} \in \hat{\pi}_2(\vec{z}, \vec{z})$ we require that $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ have both positive and negative coefficients (or be identically zero).

Note that the definition given in Section 5 of weak admissibility here corresponds to the definition of weak admissibility for all $\text{Spin}^{\mathbb{C}}$ –structures.

Definition 10.13 Fix a $\text{Spin}^{\mathbb{C}}$ –structure \mathfrak{s} on $X_{\alpha,\beta,\gamma}$ and let $\mathfrak{s}_{\alpha,\beta}$, $\mathfrak{s}_{\beta,\gamma}$ and $\mathfrak{s}_{\alpha,\gamma}$ be the restrictions of $\mathfrak{s}_{\alpha,\beta,\gamma}$ to $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$ and $Y_{\alpha,\gamma}$ respectively. We say that \mathcal{H}^3 is *strongly admissible for $\mathfrak{s}_{\alpha,\beta,\gamma}$* if for any $\vec{x} \in \mathfrak{s}_{\alpha,\beta}$, $\vec{y} \in \mathfrak{s}_{\beta,\gamma}$, $\vec{z} \in \mathfrak{s}_{\alpha,\gamma}$, $B_{\alpha,\beta} \in \hat{\pi}_2(\vec{x}, \vec{x})$, $B_{\beta,\gamma} \in \hat{\pi}_2(\vec{y}, \vec{y})$ and $B_{\alpha,\gamma} \in \hat{\pi}_2(\vec{z}, \vec{z})$ with

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

and $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ not identically zero there is some coefficient of $A + B + C$ strictly greater than n .

Note that weak (respectively strong) admissibility for \mathcal{H}^3 implies weak (strong) admissibility for each of $\mathcal{H}_{\alpha,\beta}$, $\mathcal{H}_{\beta,\gamma}$ and $\mathcal{H}_{\alpha,\gamma}$.

The proof of the following alternate characterization of weak admissibility is the same as the proof of Lemma 5.3.

Lemma 10.14 *With notation as above:*

- *The diagram \mathcal{H}^3 is weakly admissible if and only if there is an area form on Σ with respect to which for any $B_{\alpha,\beta}$, $B_{\beta,\gamma}$ and $B_{\alpha,\gamma}$ as in the definition of weak admissibility, the domain $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ has zero signed area.*
- *The diagram \mathcal{H}^3 is strongly admissible for \mathfrak{s} if there is an area form on Σ with respect to which for any $B_{\alpha,\beta}$, $B_{\beta,\gamma}$ and $B_{\alpha,\gamma}$ as in the definition of strong admissibility, $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ has signed area equal to n , and with respect to which Σ has area 1.*

Recall that we call a homology class A positive if the corresponding domain has no negative coefficients.

The following is [21, Lemma 8.9]

Lemma 10.15 *Suppose \mathcal{H}^3 is weakly admissible. Fix intersection points \vec{x} , \vec{y} , and \vec{z} and a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on $X_{\alpha,\beta,\gamma}$. Then for each $j, k \in \mathbb{Z}$ there are only finitely many positive $A \in \mathfrak{s}_{\alpha,\beta,\gamma} \cap \pi_2(\vec{x}, \vec{y}, \vec{z})$ such that*

- $\text{ind}(A) = j$,
- $n_3(A) = k$,

Proof Suppose that $A, A' \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, $n_3(A) = n_3(A') = k$, and $A, A' \in \mathfrak{s}_{\alpha,\beta,\gamma}$. Then by Lemma 10.8, A and A' are $\text{Spin}^{\mathbb{C}}$ -equivalent, so $A' = A + B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ where $B_{\alpha,\beta} \in \hat{\pi}_2(\vec{x}, \vec{x})$, $B_{\beta,\gamma} \in \hat{\pi}_2(\vec{y}, \vec{y})$, and $B_{\alpha,\gamma} \in \hat{\pi}_2(\vec{z}, \vec{z})$. By the previous lemma, we can choose an area form on Σ so that $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ has zero signed area. The result then follows as in Lemma 5.4. \square

It follows from this lemma and compactness that if \mathcal{H}^3 satisfies the weak admissibility criterion then the sums defining $\hat{f}_{\alpha,\beta,\gamma}$ and $f_{\alpha,\beta,\gamma}^+$ are finite.

The following is [21, Lemma 8.10].

Lemma 10.16 *Fix $j \in \mathbb{Z}$, intersection points \vec{x} , \vec{y} , \vec{z} , and a $\text{Spin}^{\mathbb{C}}$ -structure $\mathfrak{s}_{\alpha,\beta,\gamma}$. Suppose \mathcal{H}^3 is strongly admissible for \mathfrak{s} . Then there are only finitely many positive $A \in \pi_2(\vec{x}, \vec{y}, \vec{z}) \cap \mathfrak{s}_{\alpha,\beta,\gamma}$ such that $\text{ind}(A) = j$.*

Proof As in the previous proof, given $A, A' \in \pi_2(\vec{x}, \vec{y}, \vec{z})$ satisfying the hypotheses, $A - A' = B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\alpha,\gamma}$ for some $B_{\alpha,\beta} \in \pi_2(\vec{x}, \vec{x}')$, $B_{\beta,\gamma} \in \pi_2(\vec{y}, \vec{y}')$ and $B_{\alpha,\gamma} \in \pi_2(\vec{z}, \vec{z}')$. The proof then follows as in Lemma 5.5. \square

It follows from this lemma and compactness that if \mathcal{H}^3 satisfies the strong admissibility criterion then the sums defining $f_{\alpha,\beta,\gamma}^\infty$ and $f_{\alpha,\beta,\gamma}^{\leq 0}$ are finite.

The following is [21, Lemma 8.11]. We refer the reader there for its proof.

Proposition 10.17 *Given any pointed Heegaard triple–diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \mathfrak{z})$ there is an isotopic weakly admissible Heegaard triple–diagram. Given any pointed Heegaard triple–diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \mathfrak{z})$ and a $\text{Spin}^{\mathbb{C}}$ –structure $\mathfrak{s}_{\alpha,\beta,\gamma}$ on $X_{\alpha,\beta,\gamma}$ there is an isotopic pointed Heegaard triple–diagram which is strongly admissible for $\mathfrak{s}_{\alpha,\beta,\gamma}$.*

10.5 Moduli spaces and maps, part 2

If we wish to make a statement about all of $f_{\alpha,\beta,\gamma}^\infty$, $f_{\alpha,\beta,\gamma}^{\leq 0}$, $f_{\alpha,\beta,\gamma}^+$, or $\hat{f}_{\alpha,\beta,\gamma}$ at once we will simply write $f_{\alpha,\beta,\gamma}$. For example

Lemma 10.18 *The maps $f_{\alpha,\beta,\gamma}$ are chain maps.*

Proof This follows by considering the 1–dimensional moduli spaces $\mathcal{M}^{A_{\alpha,\beta,\gamma}}$ where $A_{\alpha,\beta,\gamma} \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, $\text{ind}(A_{\alpha,\beta,\gamma}) = 1$. The proof of Proposition 7.1 still works, so the boundary of $\mathcal{M}^{A_{\alpha,\beta,\gamma}}$ consists of height two holomorphic buildings in the homology class $A_{\alpha,\beta,\gamma}$. Each of these height two holomorphic buildings consists of

- (1) a curve of index 0 in $\Sigma \times T$ and
- (2) a curve of index 1 (defined up to translation) in one of $W_{\alpha,\beta}$, $W_{\beta,\gamma}$ or $W_{\gamma,\alpha}$,

and every such building is in $\partial\mathcal{M}^{A_{\alpha,\beta,\gamma}}$ for some $A_{\alpha,\beta,\gamma}$ of index 1. This follows from [2, Theorem 10.2] and Proposition A.1. Hence,

$$\begin{aligned} 0 = \#(\partial\mathcal{M}^C) &= \sum_{\substack{B_{\alpha,\beta} \in \pi_2(\vec{x}, \vec{x}') \\ \text{ind}(B_{\alpha,\beta})=1}} \left(\#\widehat{\mathcal{M}}^{B_{\alpha,\beta}} \right) \left(\#\mathcal{M}^{A_{\alpha,\beta,\gamma} - B_{\alpha,\beta}} \right) \\ &+ \sum_{\substack{B_{\beta,\gamma} \in \pi_2(\vec{y}, \vec{y}') \\ \text{ind}(B_{\beta,\gamma})=1}} \left(\#\widehat{\mathcal{M}}^{B_{\beta,\gamma}} \right) \left(\#\mathcal{M}^{A_{\alpha,\beta,\gamma} - B_{\beta,\gamma}} \right) \\ &+ \sum_{\substack{B_{\alpha,\gamma} \in \pi_2(\vec{z}', \vec{z}) \\ \text{ind}(B_{\alpha,\gamma})=1}} \left(\#\widehat{\mathcal{M}}_{\alpha,\gamma}^B \right) \left(\#\mathcal{M}^{A_{\alpha,\beta,\gamma} - B_{\alpha,\gamma}} \right). \end{aligned}$$

But summing this over $A_{\alpha,\beta,\gamma}$ with $n_3(A_{\alpha,\beta,\gamma}) = k$ gives the coefficient of $[\vec{z}, i + j - k]$ in $\partial \circ f_{\alpha,\beta,\gamma}^\infty + f_{\alpha,\beta,\gamma}^\infty \circ \partial$, proving the result for $f_{\alpha,\beta,\gamma}^\infty$.

The results for $f_{\alpha,\beta,\gamma}^{\leq 0}$ and $f_{\alpha,\beta,\gamma}^+$ follow immediately. The proof for $\widehat{f}_{\alpha,\beta,\gamma}$ is analogous, restricting to curves with $n_3 = 0$; we leave the details of this case to the reader. \square

We use $F_{\alpha,\beta,\gamma}$ (appropriately decorated) to denote the maps on homology induced by $f_{\alpha,\beta,\gamma}$.

Lemma 10.19 *The maps $F_{\alpha,\beta,\gamma}$ just defined are independent of the choice of complex structure J on $W_{\alpha,\beta,\gamma}$ satisfying (J'1)–(J'4).*

Proof Suppose J and J' are two complex structures on $W_{\alpha,\beta,\gamma}$ satisfying (J'1)–(J'4). Note in particular that J and J' agree on the ends of $W_{\alpha,\beta,\gamma}$. Let $f_{\alpha,\beta,\gamma,J}^\infty$ and $f_{\alpha,\beta,\gamma,J'}^\infty$ denote the maps defined above, computed with respect to J and J' respectively.

Choose a generic path J_t connecting J to J' , which is fixed on the ends of $W_{\alpha,\beta,\gamma}$. Then for any $k \in \mathbb{Z}$ there are a finite collection of $t \in (0, 1)$ such that $\mathcal{M}^{A_{\alpha,\beta,\gamma}} \neq \emptyset$ for some $A_{\alpha,\beta,\gamma}$ with $\text{ind}(A_{\alpha,\beta,\gamma}) = -1$ and $n_3(A_{\alpha,\beta,\gamma}) \leq k$. (This uses the admissibility hypothesis.) Define a map $\Phi: CF^\infty(Y_{\alpha,\beta}) \otimes CF^\infty(Y_{\beta,\gamma}) \rightarrow CF^\infty(Y_{\alpha,\gamma})$ by

$$\Phi([\vec{x}, i], [\vec{y}, j]) = \sum_{\vec{z}} \sum_{\substack{(A_{\alpha,\beta,\gamma,t}) \in \pi_2(\vec{x}, \vec{y}, \vec{z}) \times (0,1) \\ \text{ind}(A_{\alpha,\beta,\gamma,t}) = -1}} \#\mathcal{M}^{A_{\alpha,\beta,\gamma}}[\vec{z}, i + j - n_3(A_{\alpha,\beta,\gamma})].$$

The coefficient of each $[\vec{z}, k]$ is a finite sum. By exactly the same argument as used in Lemma 9.4, Φ is a chain homotopy from $f_{\alpha,\beta,\gamma,J}^\infty$ to $f_{\alpha,\beta,\gamma,J'}^\infty$.

The results for $F_{\alpha,\beta,\gamma}^\infty$, $F_{\alpha,\beta,\gamma}^{\leq 0}$ and $F_{\alpha,\beta,\gamma}^+$ follow. The result for $\widehat{F}_{\alpha,\beta,\gamma}$ is proved in an analogous way; as has become our habit we leave the details of this case to the reader. \square

Lemma 10.20 *The maps $F_{\alpha,\beta,\gamma}$ are independent of the choices of complex structures $J_{\alpha,\beta}$, $J_{\beta,\gamma}$ and $J_{\alpha,\gamma}$ satisfying (J1)–(J5) and achieving transversality, and are independent of isotopies of the α , β and γ preserving the admissibility hypotheses and not crossing \mathfrak{z} .*

Proof Let α' , β' and γ' be isotopic to α , β , and γ . As in Section 9, we may assume the isotopy introduces or cancels only one pair of intersection points, and can be realized by Lagrangian cylinders. Let $J_{\alpha',\beta'}$, $J_{\beta',\gamma'}$ and $J_{\alpha',\gamma'}$ be a complex structures on




Figure 9

$W_{\alpha,\beta}$, $W_{\beta,\gamma}$ and $W_{\alpha,\gamma}$ respectively satisfying **(J1)**–**(J5)** and achieving transversality. Choose a path of complex structures $J_{t,\alpha,\beta}$ (respectively $J_{t,\beta,\gamma}$, $J_{t,\alpha,\gamma}$), $t \in [0, 1]$, interpolating between $J_{\alpha,\beta}$ and $J_{\alpha',\beta'}$ (respectively $J_{\beta,\gamma}$ and $J_{\beta',\gamma'}$, $J_{\alpha,\gamma}$ and $J_{\alpha',\gamma'}$), as in Section 9.

Divide $W_{\alpha',\beta',\gamma'}$ into seven regions: the three ends $E_{\alpha,\beta}$, $E_{\beta,\gamma}$ and $E_{\alpha,\gamma}$; three interpolation regions $I_{\alpha,\beta}$, $I_{\beta,\gamma}$, and $I_{\gamma,\alpha}$ just below the ends; and the rest of $W_{\alpha,\beta,\gamma}$, which we call the heart, as in Figure 9. Let J be a complex structure on $W_{\alpha,\beta,\gamma}$ which

- (1) satisfies **(J'1)** and **(J'4)**.
- (2) agrees with $J_{\alpha',\beta'}$ over $E_{\alpha,\beta}$, with $J_{\beta',\gamma'}$ over $E_{\beta,\gamma}$, and with $J_{\alpha',\gamma'}$ over $E_{\alpha,\gamma}$;
- (3) agrees with $J_{t,\alpha,\beta}$ over $I_{\alpha,\beta}$, with $J_{t,\beta,\gamma}$ over $I_{\beta,\gamma}$, and with $J_{t,\alpha,\gamma}$ over $I_{\alpha,\gamma}$;
- (4) is standard in a neighborhood of $\{z_i\} \times T$ for all but at most one i ;
- (5) achieves transversality.

Let J_s be a the complex structure obtained from J by inserting a neck of length s just between each interpolation region and the heart. (See Subsection A.2 for a precise definition of this process.)

Let C_α (respectively C_β , C_γ) be Lagrangian cylinders which interpolate between the α and α' (respectively β and β' , γ and γ') curves in the regions $I_{\alpha,\beta}$ and $I_{\alpha,\gamma}$ (respectively $I_{\alpha,\beta}$ and $I_{\beta,\gamma}$, $I_{\alpha,\gamma}$ and $I_{\beta,\gamma}$).

Counting J_s -holomorphic curves with boundary on the C_α we obtain maps $f_{s,\alpha',\beta',\gamma'}$. By the same proof as the previous proposition, the maps $f_{s,\alpha',\beta',\gamma'}$ are all chain homotopic to the map $f_{\alpha',\beta',\gamma'}$.

Taking $s \rightarrow \infty$ and using the compactness result [2, Theorem 10.3] and Proposition A.2 we find that $f_{\alpha',\beta',\gamma'}$ is chain homotopic to $\Phi_{\alpha,\gamma} \circ f_{\alpha,\beta,\gamma} \circ (\Phi_{\alpha,\beta} \otimes \Phi_{\beta,\gamma})$. Here, $\Phi_{\alpha,\beta}$ (respectively $\Phi_{\beta,\gamma}$, $\Phi_{\alpha,\gamma}$) is the chain map defined in Section 9, for the isotopy between $(\vec{\alpha}, \vec{\beta})$ and $(\vec{\alpha}', \vec{\beta}')$ (respectively between $(\vec{\beta}, \vec{\gamma})$ and $(\vec{\beta}', \vec{\gamma}')$, between $(\vec{\alpha}, \vec{\gamma})$ and $(\vec{\alpha}', \vec{\gamma}')$).

Orientation systems, which have been implicit in the discussion, are extended as discussed in Section 9. \square

10.6 Associativity of triangle maps

Next we show that the maps $F_{\alpha,\beta,\gamma}$ satisfy an associativity property. Before stating it, however, we need some basic properties of Heegaard quadruple-diagrams.

10.6.1 Heegaard quadruple-diagrams As the reader can presumably guess, a pointed Heegaard quadruple-diagram consists of a genus g surface Σ , four g -tuples of pairwise disjoint homologically linearly independent simple closed curves $\vec{\alpha}$, $\vec{\beta}$, $\vec{\gamma}$ and $\vec{\delta}$, and a distinguished point $z \in \Sigma \setminus (\alpha \cup \beta \cup \gamma \cup \delta)$. In this section we will state the analogs for Heegaard quadruple-diagrams of the basic definitions and lemmas stated earlier for Heegaard triple diagrams. Except as noted, the proofs are the same as for Heegaard triple-diagrams, and hence are omitted.

Fix a pointed Heegaard quadruple-diagram $\mathcal{H}^4 = (\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \vec{\delta}, z)$. The diagram \mathcal{H}^4 specifies a 4-manifold $X_{\alpha,\beta,\gamma,\delta}$ with boundary $\partial X_{\alpha,\beta,\gamma,\delta} = Y_{\alpha,\beta} \cup Y_{\beta,\gamma} \cup Y_{\gamma,\delta} \cup (-Y_{\alpha,\delta})$ by gluing the elongated handlebodies $U_\alpha \times [0, 1]$, $U_\beta \times [0, 1]$, $U_\gamma \times [0, 1]$ and $U_\delta \times [0, 1]$ to the product of Σ and a square.

Let $W_{\alpha,\beta,\gamma,\delta}$ be the product of Σ and a disk R with four boundary punctures, thought of as a topological space without complex structure for the moment. Let e_1, \dots, e_4 denote the four boundary arcs of R , enumerated clockwise, v_{12}, v_{23}, v_{34} and v_{41} the four punctures, with v_{ij} between e_i and e_j ; see Figure 10. In $W_{\alpha,\beta,\gamma,\delta}$ we have $4g$ cylinders: $\alpha \times e_1$, $\beta \times e_2$, $\gamma \times e_3$ and $\delta \times e_4$. Let $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ denote the collection of homology classes of maps $(S, \partial S) \rightarrow (W_{\alpha,\beta,\gamma,\delta}, \alpha \times e_1 \cup \dots \cup \delta \times e_4)$ asymptotic to the I-chord collection \vec{x} for $\mathcal{H}_{\alpha,\beta}$ at v_1 , \vec{y} for $\mathcal{H}_{\beta,\gamma}$ at v_2 , and so on.

Given \vec{x}, \dots, \vec{w} there is an element $\epsilon(\vec{x}, \vec{y}, \vec{z}, \vec{w}) \in H_1(X_{\alpha,\beta,\gamma,\delta})$ defined as we defined $\epsilon(\vec{x}, \vec{y}, \vec{z})$ for triangles in Subsection 10.1.

- Lemma 10.21** (1) $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ is nonempty if and only if $\epsilon(\vec{x}, \vec{y}, \vec{z}, \vec{w}) = 0$.
- (2) If $\epsilon(\vec{x}, \vec{y}, \vec{z}, \vec{w}) = 0$ then $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w}) \cong \mathbb{Z} \oplus H_2(X_{\alpha,\beta,\gamma,\delta})$. The map $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w}) \rightarrow \mathbb{Z}$ is given by n_3 . The identification of $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})/\mathbb{Z}$ with $H_2(X_{\alpha,\beta,\gamma,\delta})$ is affine (but canonical up to translation).

Proof See Lemma 10.2 and Lemma 10.3. □

As with triangles, each element A of $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ specifies a $\text{Spin}^{\mathbb{C}}$ -structure $\mathfrak{s}_3(A)$ on $X_{\alpha,\beta,\gamma,\delta}$; the construction is the same as in Subsubsection 10.1.2. Also as before, there are obvious concatenation maps $\pi_2(\vec{x}', \vec{x}) \times \pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w}) \rightarrow \pi_2(\vec{x}', \vec{y}, \vec{z}, \vec{w})$, and similarly for \vec{y}, \vec{z} and \vec{w} . Again we say $A \in \pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ is $\text{Spin}^{\mathbb{C}}$ -equivalent to $A + B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\gamma,\delta} + B_{\alpha,\delta}$ for $B_{\alpha,\beta} \in \pi_2(\vec{x}, \vec{x}), \dots, B_{\alpha,\delta} \in \pi_2(\vec{w}, \vec{w})$, and let $S_{\alpha,\beta,\gamma,\delta}$ denote the collection of $\text{Spin}^{\mathbb{C}}$ -equivalence classes. Again we have:

Lemma 10.22 The map $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w}) \rightarrow \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma,\delta})$ descends to an injective map $S_{\alpha,\beta,\gamma,\delta} \hookrightarrow \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma,\delta})$ whose image consists of all those $\text{Spin}^{\mathbb{C}}$ -structures whose restrictions to $\partial X_{\alpha,\beta,\gamma,\delta}$ are realized to intersection points.

Proof See Lemma 10.8. □

Now, however, there is somewhat more structure. The manifold $X_{\alpha,\beta,\gamma,\delta}$ decomposes as $X_{\alpha,\beta,\gamma} \cup_{Y_{\alpha,\gamma}} X_{\alpha,\gamma,\delta}$ and as $X_{\alpha,\beta,\delta} \cup_{Y_{\beta,\delta}} X_{\beta,\gamma,\delta}$. Let $\delta_{\alpha,\gamma}$ (respectively $\delta_{\beta,\gamma}$) be the coboundary map for the Mayer–Vietoris sequence for the former (respectively the latter) decomposition. Working for the moment with the former decomposition, we have restriction maps $\text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma,\delta}) \xrightarrow{r} \text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma}) \times \text{Spin}^{\mathbb{C}}(X_{\alpha,\gamma,\delta})$. These maps commute with the H^2 -actions, and so by the Mayer–Vietoris theorem, the fibers of r are the orbits of the action of $\delta_{\alpha,\gamma} H^1(X_{\alpha,\gamma}) \subset H^2(X_{\alpha,\beta,\gamma,\delta})$ on $\text{Spin}^{\mathbb{C}}(X_{\alpha,\beta,\gamma,\delta})$. Corresponding statements hold for the decomposition $X_{\alpha,\beta,\gamma,\delta} \cong X_{\alpha,\beta,\delta} \cup_{Y_{\beta,\delta}} X_{\beta,\gamma,\delta}$.

Rather than fixing a single $\text{Spin}^{\mathbb{C}}$ -structure over $X_{\alpha,\beta,\gamma,\delta}$ we shall fix a $\delta_{\alpha,\gamma} H^1(Y_{\alpha,\gamma}) + \delta_{\beta,\delta} H^1(Y_{\beta,\delta})$ -orbit of $\text{Spin}^{\mathbb{C}}$ -structures. The reason is easier to see in terms of domains. Concatenation (addition) gives a well-defined map

$$\pi_2^{\alpha,\beta,\gamma}(\vec{x}, \vec{y}, \vec{a}) \times \pi_2^{\alpha,\gamma,\delta}(\vec{a}, \vec{z}, \vec{w}) \rightarrow \pi_2^{\alpha,\beta,\gamma,\delta}(\vec{x}, \vec{y}, \vec{z}, \vec{w}).$$

This map, however, does *not* descend to a map $S_{\alpha,\beta,\gamma} \times S_{\alpha,\gamma,\delta} \rightarrow S_{\alpha,\beta,\gamma,\delta}$: for $A_{\alpha,\beta,\gamma} \in \pi_2(\vec{x}, \vec{y}, \vec{a}), A_{\alpha,\gamma,\delta} \in \pi_2(\vec{a}, \vec{z}, \vec{w})$ and $B_{\alpha,\gamma} \in \pi_2(\vec{a}, \vec{a})$ the domains $A_{\alpha,\beta,\gamma}$ and $A_{\alpha,\beta,\gamma} + B_{\alpha,\gamma}$ define the same element of $S_{\alpha,\beta,\gamma}$, but $A_{\alpha,\beta,\gamma} + A_{\alpha,\gamma,\delta}$ and $A_{\alpha,\beta,\gamma} + B_{\alpha,\gamma} + A_{\alpha,\gamma,\delta}$ may not define the same element of $S_{\alpha,\beta,\gamma,\delta}$. So, the composition $F_{\alpha,\beta,\gamma}(F_{\alpha,\gamma,\delta}(\cdot, \cdot), \cdot)$ involves domains belonging to an entire $\pi_2^{\alpha,\gamma}(\vec{a}, \vec{a})$ -orbit

of elements of $S_{\alpha,\beta,\gamma,\delta}$. Since $\delta_{\alpha,\gamma}H^1(Y_{\alpha,\gamma})$ and $\delta_{\beta,\delta}H^1(Y_{\beta,\delta})$ may not coincide, the best we can expect to prove is associativity for certain sums of triangle maps. This is what we will (eventually) prove.

The next issue to address is admissibility.

Definition 10.23 The pointed Heegaard quadruple–diagram \mathcal{H}^4 is *weakly admissible* if given $B_{\alpha,\beta} \in \hat{\pi}_2(\vec{x}, \vec{x})$, $B_{\beta,\gamma} \in \hat{\pi}_2(\vec{y}, \vec{y})$, $B_{\gamma,\delta} \in \hat{\pi}_2(\vec{z}, \vec{z})$, $B_{\alpha,\delta} \in \hat{\pi}_2(\vec{w}, \vec{w})$ with $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\gamma,\delta} + B_{\alpha,\delta} \neq 0$, then $B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\gamma,\delta} + B_{\alpha,\delta}$ has both positive and negative coefficients.

Definition 10.24 A pointed Heegaard quadruple–diagram \mathcal{H}^4 is *strongly admissible* for a $\delta_{\alpha,\gamma}H^1(Y_{\alpha,\gamma}) + \delta_{\beta,\delta}H^1(Y_{\beta,\delta})$ –orbit of $\text{Spin}^{\mathbb{C}}$ –structures \mathfrak{S} if for any $\mathfrak{s}_{\alpha,\beta,\gamma,\delta} \in \mathfrak{S}$ and any six domains $B_{\xi,\eta} \in \hat{\pi}_2^{\xi,\eta}$, $\{\xi, \eta\} \subset \{\alpha, \beta, \gamma, \delta\}$ ($\xi \neq \eta$) such that

- $\sum_{\{\xi,\eta\}} B_{\xi,\eta} \neq 0$ and
- $\sum_{\{\xi,\eta\}} \langle c_1(\mathfrak{s}_{\alpha,\beta,\gamma}|_{Y_{\xi,\eta}}), B_{\xi,\eta} \rangle = 2n \geq 0$

then some coefficient of $\sum_{\{\xi,\eta\}} B_{\xi,\eta}$ is greater than n .

Lemma 10.25 (1) *Given any pointed Heegaard quadruple–diagram there is an isotopic weakly admissible pointed Heegaard quadruple–diagram.*

- (2) *Suppose that the pointed Heegaard quadruple–diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \vec{\delta}, \mathfrak{z})$ satisfies the conditions $\delta_{\beta,\delta}H^1(Y_{\beta,\delta})|_{Y_{\alpha,\gamma}} = 0$ and $\delta_{\alpha,\gamma}H^1(Y_{\alpha,\gamma})|_{Y_{\beta,\delta}} = 0$. Fix a $\delta_{\alpha,\gamma}H^1(Y_{\alpha,\gamma}) + \delta_{\beta,\delta}H^1(Y_{\beta,\delta})$ –orbit of $\text{Spin}^{\mathbb{C}}$ –structures \mathfrak{S} . Then there is an isotopic pointed Heegaard quadruple–diagram which is strongly admissible for \mathfrak{S} .*

For the argument, see [21, Section 8.4.2].

Lemma 10.26 *Suppose \mathcal{H}^4 is weakly admissible. Fix $j, k \in \mathbb{Z}$, intersection points $\vec{x}, \vec{y}, \vec{z}, \vec{w}$ and a $\delta H^1(Y_{\beta,\delta}) + \delta H^1(Y_{\alpha,\gamma})$ –orbit of $\text{Spin}^{\mathbb{C}}$ –structures \mathfrak{S} on $X_{\alpha,\beta,\gamma,\delta}$. Then there are only finitely many positive $A \in \pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ such that*

- $\text{ind}(A) = j$
- $n_3(A) = k$
- $\mathfrak{s}_3(A) \in \mathfrak{S}$

Proof See Lemma 10.15. □

Lemma 10.27 Suppose \mathcal{H}^4 is strongly admissible for a $\delta_{\alpha,\gamma}H^1(Y_{\alpha,\gamma}) + \delta_{\beta,\delta}H^1(Y_{\beta,\delta})$ -orbit of $\text{Spin}^{\mathbb{C}}$ -structures \mathfrak{S} on $X_{\alpha,\beta,\gamma,\delta}$. Fix $j \in \mathbb{Z}$ and intersection points $\vec{x}, \vec{y}, \vec{z}, \vec{w}$. Then there are only finitely many positive $A \in \pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ such that $\mathfrak{s}_3(A) \in \mathfrak{S}$ and $\text{ind}(A) = j$.

Proof See Lemma 10.16. □

10.6.2 Moduli spaces of squares Fix a pointed Heegaard quadruple-diagram $\mathcal{H}^4 = (\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \vec{\delta}, \mathfrak{z})$. Let R denote the unit disk with four punctures on its boundary. Let e_1, \dots, e_4 denote the four boundary components of R , enumerated clockwise, and v_{12}, v_{23}, v_{34} and $v_{4,1}$ the vertices of R , enumerated clockwise, with v_{12} between e_1 and e_2 . See Figure 10. Let $W_{\alpha,\beta,\gamma,\delta} = \Sigma \times R$.

The moduli space of conformal structures on a rectangle is parameterized by \mathbb{R} . Let $j_a, a \in \mathbb{R}$, sweep out this space. Do this in such a way that as $s \rightarrow -\infty$ an arc in R connecting e_1 to e_3 collapses, while as $s \rightarrow \infty$ an arc in R connecting e_2 to e_4 collapses; see Figure 10.

Fix a point \mathfrak{z}_i in each component of $\Sigma \setminus (\alpha \cup \beta \cup \gamma \cup \delta)$. Choose complex structures $J_{\xi,\eta}, \{\xi, \eta\} \subset \{\alpha, \beta, \gamma, \delta\}$ which satisfy **(J1)**–**(J5)** and achieve transversality.

Fix a path J_a of complex structures on $W_{\alpha,\beta,\gamma,\delta}$ such that

- (1) For every $a \in \mathbb{R}$, J_a is tamed by the split symplectic form on $\Sigma \times R$.
- (2) For every $a \in \mathbb{R}$, projection π_R onto R is (j_a, J_a) -holomorphic.
- (3) In a neighborhood $U_{\{\mathfrak{z}_i\}}$ of $\{\mathfrak{z}_i\} \times R$, $J = j_{\Sigma} \times j_T$ is split.
- (4) Near $\Sigma \times \{v_{12}\}$, J agrees with $J_{\alpha,\beta}$. Near $\Sigma \times \{v_{23}\}$, J agrees with $J_{\beta,\gamma}$. Near $\Sigma \times \{v_{34}\}$, J agrees with $J_{\gamma,\delta}$. Near $\Sigma \times \{v_{41}\}$, J agrees with $J_{\alpha,\delta}$.
- (5) As $a \rightarrow -\infty$, J_a degenerates to complex structures which satisfy **(J'1)**–**(J'4)** and achieve transversality for $W_{\alpha,\beta,\gamma}$ and $W_{\alpha,\gamma,\delta}$.
- (6) As $a \rightarrow \infty$, J_a degenerates to complex structures which satisfy **(J'1)**–**(J'4)** and achieve transversality for $W_{\delta,\alpha,\beta}$ and $W_{\gamma,\delta,\beta}$.
- (7) J_a achieves transversality (as a path of almost complex structures) for holomorphic curves with index ≤ 1 .

Checking that such a J_a exists is similar to the proof of transversality in Section 3.

Given I-chord collections $\vec{x}, \vec{y}, \vec{z}, \vec{w}$ for $Y_{\alpha,\beta}, Y_{\beta,\gamma}, Y_{\gamma,\delta}$ and $Y_{\alpha,\delta}$ let $\pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$ denote homology classes of maps to $W_{\alpha,\beta,\gamma,\delta}$ connecting $\vec{x}, \vec{y}, \vec{z}$, and \vec{w} . Given $A \in \pi_2(\vec{x}, \vec{y}, \vec{z}, \vec{w})$, let \mathcal{M}^A denote the union over $a \in \mathbb{R}$ of all embedded J_a -holomorphic




Figure 10: Degenerations of a rectangle

curves $u: (S, \partial S) \rightarrow (W_{\alpha, \beta, \gamma, \delta}, \alpha \times e_1 \cup \beta \times e_2 \cup \gamma \times e_3 \cup \delta \times e_4)$ (without closed components, and with one component of ∂S mapped to each Lagrangian cylinder) in the homology class A . Let $\text{ind}(A)$ denote the index of this $\bar{\partial}$ problem in the homology class A , so $\dim \mathcal{M}^A = \text{ind}(A) + 1$, if \mathcal{M}^A is nonempty. (The $+1$ appears because we consider a 1-parameter family of almost complex structures.)

10.6.3 Orienting squares The moduli spaces $\mathcal{M}^{A_{\alpha, \beta, \gamma, \delta}}$ are orientable for the same reason all the other moduli spaces considered so far have been. Again we want to choose orientations for the $\mathcal{M}^{A_{\alpha, \beta, \gamma, \delta}}$, $\mathcal{M}^{A_{\kappa, \eta, \xi}}$ and $\mathcal{M}^{B_{\eta, \xi}}$, $\{\kappa, \eta, \xi\} \subset \{\alpha, \beta, \gamma, \delta\}$ consistent with various gluings – all possible gluings of a 2-gon to a rectangle or triangle, and the two gluings $\mathcal{M}^{A_{\alpha, \beta, \gamma}} \times \mathcal{M}^{A_{\alpha, \gamma, \delta}} \times [R, \infty) \hookrightarrow \mathcal{M}^{A_{\alpha, \beta, \gamma} + A_{\alpha, \gamma, \delta}}$ and $\mathcal{M}^{A_{\alpha, \beta, \delta}} \times \mathcal{M}^{A_{\beta, \gamma, \delta}} \times [R, \infty) \hookrightarrow \mathcal{M}^{A_{\alpha, \beta, \delta} + A_{\beta, \gamma, \delta}}$. Again, it follows from standard arguments that there is some such coherent orientation system. And, again it is useful to have something slightly stronger:

Lemma 10.28 (Compare [21, Proposition 8.15]) *Suppose \mathcal{H}^4 is a Heegaard quadruple-diagram such that the image of $H_2(Y_{\beta, \delta})$ under the map $H_2(X_{\alpha, \beta, \gamma, \delta}) \rightarrow$*

$H_2(X_{\alpha,\beta,\gamma,\delta}, \partial X_{\alpha,\beta,\gamma,\delta})$ is zero. Fix a $\delta H^1(Y_{\alpha,\gamma})$ -orbit of $\text{Spin}^{\mathbb{C}}$ -structures \mathfrak{S} on $X_{\alpha,\beta,\gamma,\delta}$ and orientation systems $\mathfrak{o}_{\alpha,\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma,\delta}$ for $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \mathfrak{z}, \mathfrak{S}|_{X_{\alpha,\beta,\gamma}})$ and $(\Sigma, \vec{\alpha}, \vec{\gamma}, \vec{\delta}, \mathfrak{z}, \mathfrak{S}|_{X_{\alpha,\gamma,\delta}})$ inducing the same orientation over $Y_{\alpha,\gamma}$. Then there is a coherent orientation system for \mathcal{H}^4 extending $\mathfrak{o}_{\alpha,\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma,\delta}$.

Proof Our proof is the same as the one given in [21].

The orientation systems $\mathfrak{o}_{\alpha,\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma,\delta}$ determine orientations $\mathfrak{o}_{\alpha,\beta,\gamma,\delta}(A_{\alpha,\beta,\gamma,\delta})$ for all $A_{\alpha,\beta,\gamma,\delta}$ with $\mathfrak{s}_3(A_{\alpha,\beta,\gamma,\delta}) \in \mathfrak{S}$, which is well-defined since $\mathfrak{o}_{\alpha,\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma,\delta}$ induce the same orientation system on $Y_{\alpha,\gamma}$. The orientation systems $\mathfrak{o}_{\alpha,\beta,\gamma}$ and $\mathfrak{o}_{\alpha,\gamma,\delta}$ also determine orientation systems $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\gamma}$, $\mathfrak{o}_{\gamma,\delta}$, $\mathfrak{o}_{\alpha,\delta}$ and $\mathfrak{o}_{\alpha,\gamma}$ over $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$, $Y_{\gamma,\delta}$, $Y_{\alpha,\delta}$ and $Y_{\alpha,\gamma}$.

Choose intersection points $\vec{x} \in \mathfrak{S}|_{Y_{\alpha,\beta}}$, $\vec{y} \in \mathfrak{S}|_{Y_{\beta,\gamma}}$, $\vec{z} \in \mathfrak{S}|_{Y_{\gamma,\delta}}$, $\vec{w} \in \mathfrak{S}|_{Y_{\alpha,\delta}}$, $\vec{u} \in \mathfrak{S}|_{Y_{\alpha,\gamma}}$ and $\vec{v} \in \mathfrak{S}|_{Y_{\beta,\delta}}$. Choose any $A_{\alpha,\beta,\gamma,0} \in \pi_2(\vec{x}, \vec{y}, \vec{u})$ and $A_{\beta,\gamma,\delta,0} \in \pi_2(\vec{y}, \vec{z}, \vec{v})$. Choose an arbitrary orientation $\mathfrak{o}_{\alpha,\beta,\delta}(A_{\alpha,\beta,\delta,0})$ over $A_{\alpha,\beta,\delta,0}$. Then $\mathfrak{o}_{\alpha,\beta,\delta}(A_{\alpha,\beta,\delta,0})$ and $\mathfrak{o}_{\alpha,\beta,\gamma,\delta}(A_{\alpha,\beta,\delta,0} + A_{\beta,\gamma,\delta,0})$ determine an orientation $\mathfrak{o}_{\beta,\gamma,\delta}(A_{\beta,\gamma,\delta,0})$.

Now, we choose $\mathfrak{o}_{\beta,\delta}$ as follows. For $B_{\beta,\delta} \in \pi_2(\vec{v}, \vec{v})$, from the assumption that $H_2(Y_{\beta,\delta})$ is trivial inside $H_2(X_{\alpha,\beta,\gamma,\delta}, \partial X_{\alpha,\beta,\gamma,\delta})$, we can choose $B_{\alpha,\beta}$, $B_{\beta,\gamma}$, $B_{\gamma,\delta}$ and $B_{\alpha,\delta}$ so that

$$(10) \quad A_{\alpha,\beta,\delta,0} + A_{\beta,\gamma,\delta,0} + B_{\beta,\delta} = A_{\alpha,\beta,\delta,0} + A_{\beta,\gamma,\delta,0} + B_{\alpha,\beta} + B_{\beta,\gamma} + B_{\gamma,\delta} + B_{\alpha,\delta}$$

Choose $\mathfrak{o}_{\beta,\delta}(B_{\beta,\delta})$ so that the orientations induced by the two decompositions in Equation (10) agree. We have constructed $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\gamma}$, $\mathfrak{o}_{\gamma,\delta}$ and $\mathfrak{o}_{\alpha,\delta}$ so that this is independent of the choice of $B_{\alpha,\beta}, \dots, B_{\alpha,\delta}$. Choose arbitrary orientations over a complete set of paths for $Y_{\beta,\delta}$ to finish defining $\mathfrak{o}_{\beta,\delta}$.

Finally, $\mathfrak{o}_{\alpha,\beta,\delta}(A_{\alpha,\beta,\delta,0})$, $\mathfrak{o}_{\alpha,\beta}$, $\mathfrak{o}_{\beta,\delta}$ and $\mathfrak{o}_{\alpha,\delta}$ together specify $\mathfrak{o}_{\alpha,\beta,\delta}$ completely, and a similar remark applies to $\mathfrak{o}_{\beta,\gamma,\delta}$. We have, thus, finished constructing all the orientations. It is clear that they are coherent. \square

10.6.4 Proof of associativity

Proposition 10.29 Let $\mathcal{H}^4 = (\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \vec{\delta}, \mathfrak{z})$ be a pointed Heegaard quadruple-diagram. Fix a $\delta_{\alpha,\gamma} H^1(Y_{\alpha,\gamma}) + \delta_{\beta,\delta} H^1(Y_{\beta,\delta})$ -orbit of $\text{Spin}^{\mathbb{C}}$ -structures \mathfrak{S} on $X_{\alpha,\beta,\gamma,\delta}$. Assume that \mathcal{H}^4 is strongly admissible for \mathfrak{S} . Fix $*$ $\in \{\infty, +, -\}$. Then for any $\xi_{\alpha,\beta} \in HF^*(Y_{\alpha,\beta})$ and $\theta_{\beta,\gamma} \in HF^{\leq 0}(Y_{\beta,\gamma})$, $\theta_{\gamma,\delta} \in HF^{\leq 0}(Y_{\gamma,\delta})$ we have

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

An exactly similar statement holds with both HF^* and HF^{\leq} replaced by \widehat{HF} . For the case of \widehat{HF} , weak admissibility of the Heegaard quadruple–diagram is sufficient.

(In the proposition, we assume the maps are computed with respect to a coherent choice of orientations, as described in Subsubsection 10.6.3. Also note that every $\text{Spin}^{\mathbb{C}}$ –structure in \mathfrak{S} restricts to the same $\text{Spin}^{\mathbb{C}}$ –structure on $Y_{\alpha,\beta}$ (respectively $Y_{\beta,\gamma}$, $Y_{\gamma,\delta}$, $Y_{\alpha,\delta}$). When we speak of the Floer homology groups of $Y_{\alpha,\beta}$ (respectively $Y_{\beta,\gamma}$, $Y_{\gamma,\delta}$, $Y_{\alpha,\delta}$) we mean the Floer homology groups calculated with respect to this $\text{Spin}^{\mathbb{C}}$ –structure.)

Proof To make notation clearer, we will give the proof for $* = \infty$. The other cases are completely analogous. Define $h: CF^{\infty}(Y_{\alpha,\beta}) \otimes CF^{\leq 0}(Y_{\beta,\gamma}) \otimes CF^{\leq 0}(Y_{\gamma,\delta}) \rightarrow CF^{\infty}(Y_{\alpha,\delta})$ by

$$h([\vec{x}, i] \otimes [\vec{y}, j] \otimes [\vec{z}, k]) = \sum_{\vec{w}} \sum_{\substack{A \in \pi_2(\vec{x}, v\vec{y}, \vec{z}, \vec{w}) \\ s_3(A) \in \mathfrak{S} \\ \text{ind}(A) = -1}} (\#\mathcal{M}^A) [\vec{w}, i + j + k - n_3(A)].$$

The sum defining h makes sense by compactness and Lemma 10.27. Note that the map h has degree -1 .

Counting the ends of the space

$$\bigcup_{\substack{A \in \mathfrak{S} \\ \text{ind}(A) = 0}} \mathcal{M}^A$$

we find that

$$\begin{aligned} 0 &= \sum_{s \in \mathfrak{S}} f_{\alpha,\gamma,\delta}^{\infty} \left(f_{\alpha,\beta,\gamma}^{\infty}([\vec{x}, i] \otimes [\vec{y}, j]; \mathfrak{s}|_{X_{\alpha,\beta,\gamma}}) \otimes [\vec{z}, k]; \mathfrak{s}|_{X_{\alpha,\gamma,\delta}} \right) \\ &\quad - \sum_{s \in \mathfrak{S}} f_{\alpha,\beta,\delta}^* \left([\vec{x}, i] \otimes f_{\beta,\gamma,\delta}^{\leq 0}([\vec{y}, j] \otimes [\vec{z}, k]; \mathfrak{s}|_{X_{\beta,\gamma,\delta}}); \mathfrak{s}|_{X_{\alpha,\beta,\delta}} \right) \\ &\quad + h \circ \partial([\vec{x}, i] \otimes [\vec{y}, j] \otimes [\vec{z}, k]) + \partial \circ h([\vec{x}, i] \otimes [\vec{y}, j] \otimes [\vec{z}, k]). \end{aligned}$$

(The first two terms correspond to contributions from the degenerations shown in Figure 10. The last two terms correspond to the usual level splittings at the (four) cylindrical ends.)

The result is immediate. \square

11 Handleslides

In this section we will prove that the Floer homologies defined in Section 8 are unchanged by handleslides among the β - (and, symmetrically, the α -) circles. This follows from the associativity of the triangle maps defined in Section 10 and the calculation of a few specific moduli spaces.

Our proof is essentially the same as the proof given by Ozsváth and Szabó in [21, Section 9]. Indeed, their calculations are done in our language, and most of the rest of the proof follows formally from the results of Section 10. Our proof differs from theirs in the following ways. Firstly, we try to calculate the minimum amount necessary in order to prove handleslide invariance. We hope this will provide a slightly better understanding of the formal properties underlying handleslide invariance, and in turn lead to generalizations. Secondly, we give a more geometrical proof of Proposition 11.4 (roughly [21, Proposition 9.8]), using 1-gons. Thirdly, some details about the $H_1(Y)/Tors$ -actions look slightly different in our language.

We now state the main steps, and then return later to the proofs. First, notation. Let $\mathcal{H} = (\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ be a pointed Heegaard diagram. Let β'_i be a small Hamiltonian perturbation of β_i intersecting β_i transversally in two points and disjoint from β_j for $i \neq j$. Let β_1^H be a curve obtained by handlesliding β_1 over β_2 in the complement of \mathfrak{z} and then doing a small isotopy so that β_1^H intersects each of β_1 and β'_1 transversally in two points and is disjoint from β_i and β'_i for $i > 1$. Let β_i^H , $i > 1$, be a small Hamiltonian perturbation of β_i , intersecting each of β_i and β'_i transversally in two points and disjoint from β_j and β'_j for $j \neq i$. See Figure 11. In the notation of [21, Section 9], our β'_i is their δ_i and our β_i^H is their γ_i .

There are 2^g intersection points in $(\Sigma, \vec{\beta}, \vec{\beta}', \mathfrak{z})$. Let $\vec{\theta}_{\beta, \beta'} = \{\theta_1, \dots, \theta_g\}$ denote the unique intersection point of maximal grading. Similarly, in $(\Sigma, \vec{\beta}, \vec{\beta}^H, \mathfrak{z})$ (respectively $(\Sigma, \vec{\beta}, \vec{\beta}', \mathfrak{z})$) there are exactly 2^g intersection points. Let $\vec{\theta}_{\beta, \beta^H} = \{\theta_1^H, \dots, \theta_g^H\}$ (respectively $\vec{\theta}_{\beta', \beta^H} = \{\theta'_1, \dots, \theta'_g\}$) denote the unique intersection point of maximal grading. See Figure 11.

A few words about $\text{Spin}^{\mathbb{C}}$ -structures. We fix a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on $Y_{\alpha, \beta}$ for the rest of the section; when we refer to Floer homology groups of $Y_{\alpha, \beta}$ we mean the groups computed with respect to \mathfrak{s} . Now, note that all of the $3 \cdot 2^g$ intersection points in $(\Sigma, \vec{\beta}, \vec{\beta}')$, $(\Sigma, \vec{\beta}, \vec{\beta}^H)$ and $(\Sigma, \vec{\beta}', \vec{\beta}^H)$ represent the unique torsion $\text{Spin}^{\mathbb{C}}$ -structure on $\#^g S^1 \times S^2$. (This follows, for instance, from Lemma 4.11.) Throughout this section, when discussing Heegaard diagrams for $\#^g S^1 \times S^2$ we shall implicitly work with the torsion $\text{Spin}^{\mathbb{C}}$ -structure. Later in this section, we shall consider the Heegaard

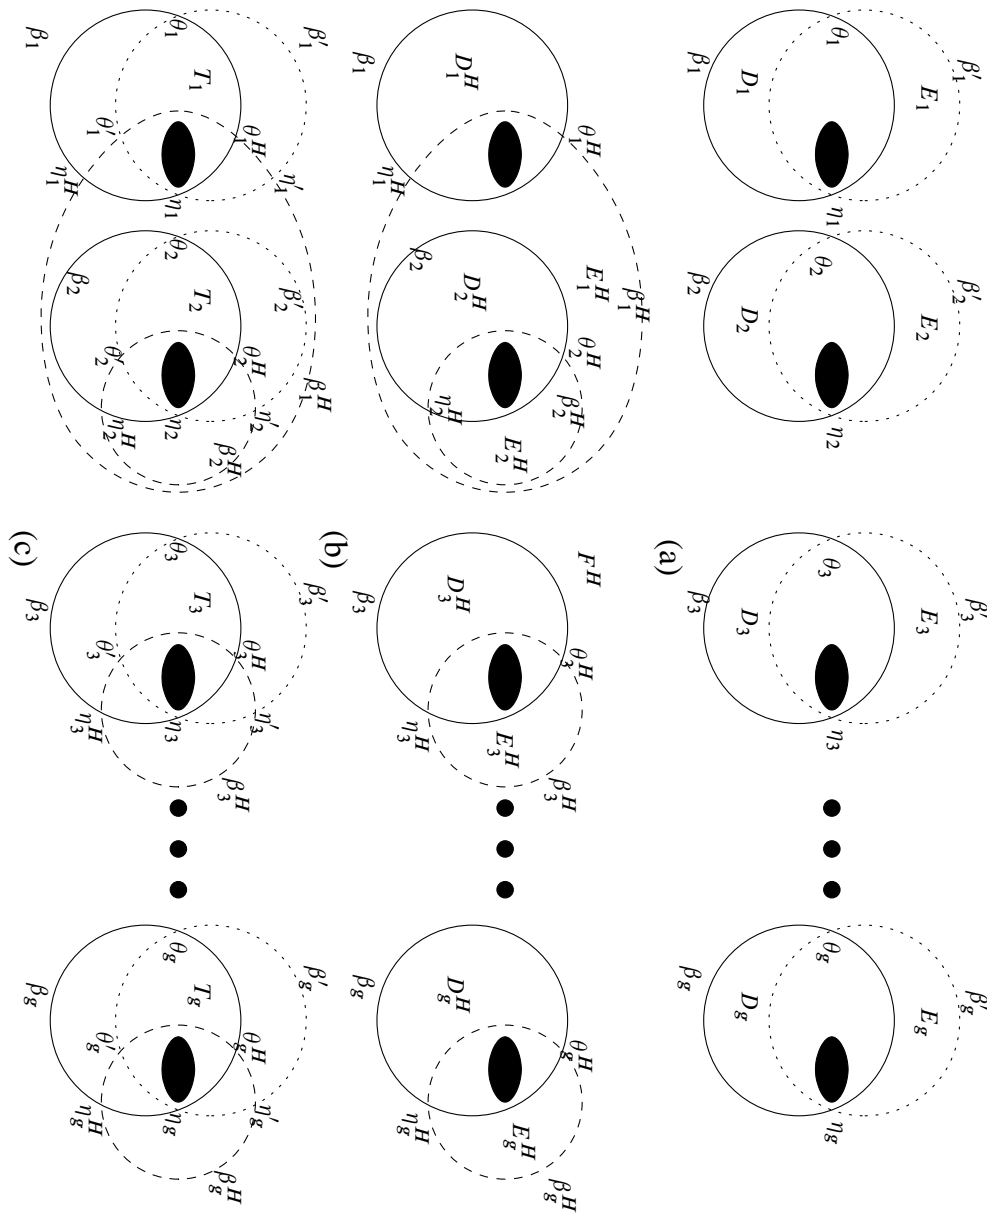


Figure 11: Heegaard diagrams for $\#^g S^1 \times S^2$

triple-diagrams $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\beta}^H)$, $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\beta}')$, and $(\Sigma, \vec{\beta}, \vec{\beta}^H, \vec{\beta}')$. It is easy to see that $X_{\alpha, \beta, \beta^H}$ and $X_{\alpha, \beta, \beta'}$ are each diffeomorphic to the complement of (a neighborhood of) a bouquet of g circles in $Y_{\alpha, \beta} \times [0, 1]$. (Compare [21, Example 8.1].) By considering cohomology groups, say, it follows that given a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on $Y_{\alpha, \beta}$ there is a

unique $\text{Spin}^{\mathbb{C}}$ -structure on X_{α,β,β^H} (respectively $X_{\alpha,\beta,\beta'}$) which restricts to \mathfrak{s} on $Y_{\alpha,\beta}$ and the torsion $\text{Spin}^{\mathbb{C}}$ -structure on Y_{β,β^H} (respectively $Y_{\beta,\beta'}$). It follows that there is a unique $\text{Spin}^{\mathbb{C}}$ -structure on $X_{\beta,\beta^H,\beta'}$ which restricts to the torsion $\text{Spin}^{\mathbb{C}}$ -structure on the three boundary components. When discussing triangle maps we shall always assume that they are computed with respect to these choices of $\text{Spin}^{\mathbb{C}}$ -structures on X_{α,β,β^H} , $X_{\alpha,\beta,\beta'}$ and $X_{\beta,\beta^H,\beta'}$. We shall, however, tend to suppress $\text{Spin}^{\mathbb{C}}$ -structures from the notation.

- Lemma 11.1** (1) *There is a coherent orientation system $\mathfrak{o}_{\beta,\beta'}$ for $(\Sigma, \vec{\beta}, \vec{\beta}', \mathfrak{z})$ with respect to which $\vec{\theta}_{\beta,\beta'}$ is a cycle in \widehat{CF} and $[\vec{\theta}_{\beta,\beta'}, 0]$ is a cycle in $CF^{\leq 0}$.*
- (2) *There is a coherent orientation system $\mathfrak{o}_{\beta,\beta^H}$ (respectively $\mathfrak{o}_{\beta',\beta^H}$) for $(\Sigma, \vec{\beta}, \vec{\beta}^H, \mathfrak{z})$ (respectively $(\Sigma, \vec{\beta}^H, \vec{\beta}', \mathfrak{z})$) with respect to which $\vec{\theta}_{\beta,\beta^H}$ (respectively $\vec{\theta}_{\beta^H,\beta'}$) is a cycle in \widehat{CF} and $[\vec{\theta}_{\beta,\beta^H}, 0]$ (respectively $[\vec{\theta}_{\beta^H,\beta'}, 0]$) is a cycle in $CF^{\leq 0}$.*
- (3) *The orientation systems $\mathfrak{o}_{\beta,\beta'}$, $\mathfrak{o}_{\beta,\beta^H}$ and $\mathfrak{o}_{\beta^H,\beta'}$ above can be chosen so that for some orientation system $\mathfrak{o}_{\beta,\beta^H,\beta'}$ for $(\Sigma, \vec{\beta}, \vec{\beta}^H, \vec{\beta}', \mathfrak{z})$, the orientation systems $\mathfrak{o}_{\beta,\beta^H,\beta'}$, $\mathfrak{o}_{\beta,\beta'}$, $\mathfrak{o}_{\beta,\beta^H}$ and $\mathfrak{o}_{\beta^H,\beta'}$ are coherent.*

The proof will be given in Subsection 11.1. We shall choose, once and for all, orientation systems as in the lemma.

The goal of this section is the following

Proposition 11.2 (Compare [21, Theorem 9.5]) *Fix a $\text{Spin}^{\mathbb{C}}$ -structure \mathfrak{s} on $Y_{\alpha,\beta} = Y_{\alpha,\beta^H}$. Then the map*

$$\widehat{F}_{\alpha,\beta,\beta^H}(\cdot \otimes \vec{\theta}_{\beta,\beta^H}): \widehat{HF}(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z}, \mathfrak{s}) \rightarrow \widehat{HF}(\Sigma, \vec{\alpha}, \vec{\beta}^H, \mathfrak{z}, \mathfrak{s})$$

is an isomorphism. The map

$$F_{\alpha,\beta,\beta^H}^*(\cdot \otimes [\vec{\theta}_{\beta,\beta^H}, 0]): HF^*(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z}, \mathfrak{s}) \rightarrow HF^*(\Sigma, \vec{\alpha}, \vec{\beta}^H, \mathfrak{z}, \mathfrak{s})$$

is an isomorphism for $$ $\in \{\infty, +, i\}$. These isomorphisms commute with the long exact sequences and the $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H^1(Y)/\text{Tors}$ -actions.*

Note we have suppressed some discussion of orientation systems from the statement of the proposition. See Lemma 11.11 below.

The essence of the proof of Proposition 11.2 is the following two propositions:

Proposition 11.3 (Compare [21, Lemma 9.7])

$$\begin{aligned}\widehat{F}_{\beta, \beta^H, \beta'}(\vec{\theta}_{\beta, \beta^H} \otimes \vec{\theta}_{\beta^H, \beta'}) &= \vec{\theta}_{\beta, \beta'} \\ F_{\beta, \beta^H, \beta'}^{\leq 0}([\vec{\theta}_{\beta, \beta^H}, 0] \otimes [\vec{\theta}_{\beta^H, \beta'}, 0]) &= [\vec{\theta}_{\beta, \beta'}, 0].\end{aligned}$$

(The proof is in Subsection 11.1.)

Proposition 11.4 (Compare [21, Proposition 9.8]) *The map $\widehat{F}_{\alpha, \beta, \beta'}(\cdot \otimes \vec{\theta}_{\beta, \beta'})$ (respectively the maps $F^*(\cdot \otimes [\vec{\theta}_{\beta, \beta'}, 0])$ for $* \in \{+, -, \infty\}$) is the isomorphism $\widehat{\Phi}$ (respectively Φ^* , $* \in \{+, -, \infty\}$) induced by the isotopy from $\vec{\beta}$ to $\vec{\beta}'$ as in Section 9.*

(The proof is in Subsection 11.2.)

We sketch the proof of Proposition 11.2 assuming Proposition 11.4 and Proposition 11.3.

Proof of Proposition 11.2 (sketch) Observe that

$$\begin{aligned}\widehat{F}_{\alpha, \beta^H, \beta'}\left(\widehat{F}_{\alpha, \beta, \beta^H}(\cdot \otimes \vec{\theta}_{\beta, \beta^H}) \otimes \vec{\theta}_{\beta^H, \beta'}\right) &= \widehat{F}_{\alpha, \beta, \beta'}\left(\cdot \otimes \widehat{F}_{\beta, \beta^H, \beta'}(\vec{\theta}_{\beta, \beta^H} \otimes \vec{\theta}_{\beta^H, \beta'})\right) \\ &= \widehat{F}_{\alpha, \beta, \beta'}(\cdot \otimes \vec{\theta}_{\beta, \beta'}) \\ &= \widehat{\Phi}_{\beta, \beta'}(\cdot)\end{aligned}$$

where $\widehat{\Phi}_{\beta, \beta'}$ is the isomorphism induced by the isotopy from $\vec{\beta}$ to $\vec{\beta}'$. (The first equality follows by the associativity of triangle maps (Proposition 10.29). The second follows from Proposition 11.3. The third follows from Proposition 11.4.)

It follows that $\widehat{F}_{\alpha, \beta^H, \beta'}$ is surjective and $\widehat{F}_{\alpha, \beta, \beta^H}$ is injective. The same argument with the roles of β and β^H exchanged shows that $\widehat{F}_{\alpha, \beta, \beta^H}$ is surjective. So, $\widehat{F}_{\alpha, \beta, \beta^H}$ is an isomorphism, proving that \widehat{HF} is invariant under handleslides. The proofs for HF^- , HF^+ and HF^∞ are just the same. \square

11.1 Proofs of Proposition 11.3 and Lemma 11.1

Proof of Proposition 11.3 With the notation of Figure 11, note that the domain $T_1 + \cdots + T_g$ achieves transversality, and admits a unique holomorphic representative, so $\#\mathcal{M}^{T_1 + \cdots + T_g} = \pm 1$.

For any intersection point $\xi \neq \vec{\theta}_{\beta, \beta'}$ of $(\Sigma, \vec{\beta}, \vec{\beta}', \mathfrak{z})$ and any $A \in \pi_2(\vec{\theta}_{\beta, \beta^H}, \vec{\theta}_{\beta^H, \beta'}, \xi)$, either $n_3(A) < 0$ or $\text{ind}(A) > 0$.

Any element of $\pi_2(\vec{\theta}_{\beta,\beta^H}, \vec{\theta}_{\beta^H,\beta'}, \vec{\theta}_{\beta,\beta'})$ other than $T_1 + \dots + T_g$ either has $\text{ind} > 0$ or some negative coefficient.

It follows that for any choice of coherent orientation system,

$$\begin{aligned} \widehat{f}_{\beta,\beta^H,\beta'}(\vec{\theta}_{\beta,\beta^H} \otimes \vec{\theta}_{\beta^H,\beta'}) &= \pm \vec{\theta}_{\beta,\beta'} \\ f_{\beta,\beta^H,\beta'}^{\leq 0}([\vec{\theta}_{\beta,\beta^H}, 0] \otimes [\vec{\theta}_{\beta^H,\beta'}, 0]) &= \pm [\vec{\theta}_{\beta,\beta'}, 0] \end{aligned} \quad \square$$

Proof of Lemma 11.1 Firstly, note that Part (2), Proposition 11.3, Lemma 10.10 and Lemma 10.18 imply Parts (1) and (3). So, it remains only to prove Part (2)

Label the components of $\Sigma \setminus (\beta \cup \beta^H)$ by $D_1^H, \dots, D_g^H, E_1^H, \dots, E_1^H$ and F^H , and the points in $\beta_i \cap \beta_i^H$ by θ_i^H and η_i^H as in Figure 11.

Observe that

$$gr(\vec{\theta}_{\beta,\beta^H}, \{\theta_{i_1}^H, \dots, \theta_{i_k}^H, \eta_{j_{k+1}}^H, \dots, \eta_{j_g}^H\}) = g - k.$$

It follows from this and positivity of domains that the only homology classes which could contribute to $\widehat{\partial}\vec{\theta}_{\beta,\beta^H}$ or $\partial^{\leq 0}[\vec{\theta}_{\beta,\beta^H}, 0]$ are $D_1^H, \dots, D_g^H, E_2^H, \dots, E_g^H, E_1^H + E_2^H$ and $E_1^H + D_2^H$.

By Proposition 3.9 and Lemma 3.10, a generic perturbation of the α - and β -circles achieves transversality for all of the moduli spaces under consideration. It follows from the Riemann mapping theorem that there is a unique equivalence class of holomorphic curves in each of $\widehat{\mathcal{M}}^{D_1^H}, \dots, \widehat{\mathcal{M}}^{D_g^H}$ and $\widehat{\mathcal{M}}^{E_2^H}, \dots, \widehat{\mathcal{M}}^{E_g^H}$.

We will show that for a generic perturbation of the α - and β -circles, one of $\widehat{\mathcal{M}}^{E_1^H + E_2^H}$ and $\widehat{\mathcal{M}}^{E_1^H + D_1^H}$ has one element and the other is empty. To this end, we use the following

Sublemma 11.5 *Let A be an annulus with boundary circles C_1 and C_2 and marked points x_i, y_i on C_i . Define the conformal angle between x_1 and y_1 (respectively x_2 and y_2) as follows. The annulus A is conformally equivalent to $\{1 \leq |z| \leq R\} \subset \mathbb{C}$ for some R . We can arrange that the equivalence takes x_1 (respectively x_2) to 1. Then the conformal angle between x_1 and y_1 (respectively x_2 and y_2) as the image of y_1 (respectively y_2) in S^1 under the equivalence. Then:*

- (1) *The conformal angle is well-defined.*
- (2) *There is a holomorphic involution of A exchanging x_1 and x_2 , and y_1 and y_2 , if and only if the conformal angle between x_1 and y_1 equals the conformal angle between x_2 and y_2 .*

Proof It is well-known that every conformal automorphism of the annulus $A' = \{1 \leq |z| \leq R\} \subset \mathbb{C}$ is either of the form $z \mapsto e^{i\theta}z$ or $z \mapsto \frac{Re^{i\theta}}{z}$. (Consider the universal cover.) Both parts of the claim are immediate from this observation. \square

Note that the condition in Part (2) of the sublemma is the same as the condition that there be a holomorphic 2-fold branched cover $c: A \rightarrow \mathbb{D}$ such that $c(x_1) = c(x_2) = -i$ and $c(y_1) = c(y_2) = i$.

Now, if $u: S \rightarrow W$ is an element of $\widehat{\mathcal{M}}^{E_1^H + E_2^H}$ then $\pi_\Sigma \circ u$ is an analytic isomorphism on the interior of S . It follows that S can be obtained by cutting $\overline{E_1^H + E_2^H}$ along β_2^H . Let x_1 denote the preimage of θ_1^H , y_1 the preimage of η_1^H and let x_2 and y_2 denote the two preimages of θ_2^H . (Which preimage should be labeled x_2 and which y_2 is clear.) Then elements of $\widehat{\mathcal{M}}^{E_1^H + E_2^H}$ correspond to choices of cuts of $\overline{E_1^H \cup E_2^H}$ along β_2^H starting at θ_2^H for which there are holomorphic involutions exchanging x_1 and x_2 , and y_1 and y_2 . A similar remark applies to $\widehat{\mathcal{M}}^{E_1^H + D_2^H}$ with β_2 in place of β_2^H .

For $p \in \beta_2^H$, let $a^H(p)$ denote the ratio of the conformal angle from x_2 to y_2 to the conformal angle from x_1 to y_1 , for the annulus obtained by cutting $\overline{E_1^H + E_2^H}$ along β_2^H to the point p . Let $a(p)$ denote the same with $\overline{E_1^H + D_2^H}$ in place of $\overline{E_1^H \cup E_2^H}$ and β_2 in place of β_2^H .

Observe that:

- (1) $a(p) = 1$ (respectively $a^H(p) = 1$) if and only if cutting to p gives an element of $\widehat{\mathcal{M}}^{E_1^H + D_2^H}$ (respectively $\widehat{\mathcal{M}}^{E_1^H + E_2^H}$).
- (2) $a(\theta_2^H) = a^H(\theta_2^H)$.
- (3) a and a^H are monotone as p travels from θ_2^H to η_2^H .
- (4) $a \rightarrow 0$ as $p \rightarrow \eta_2^H$. $a^H \rightarrow \infty$ as $p \rightarrow \eta_2^H$.

The first two claims are obvious. The third is clear. The fourth follows by considering the Gromov limit as $p \rightarrow \eta_2^H$.

For a generic choice of α - and β -circles, $a(\theta_2^H) \neq 1$. It follows that one of a or a^H assumes the value 1 once, and the other never assumes the value 1. Hence one of $\widehat{\mathcal{M}}^{E_1^H + D_2^H}$ or $\widehat{\mathcal{M}}^{E_1^H + E_2^H}$ has a unique element and the other is empty.

Now, specifying a coherent orientation system for $(\Sigma, \vec{\beta}, \vec{\beta}^H, \mathfrak{z})$ is equivalent to specifying orientations arbitrarily over $D_1^H, \dots, D_g^H, E_2^H, \dots, E_g^H$, and either $E_1^H + D_2^H$





Figure 12: Sketch of proof of Proposition 11.4

or $E_1^H + E_2^H$. Choose the orientations such that $\#\widehat{\mathcal{M}}^{D_i^H} = 1$ ($i = 1, \dots, g$), $\#\widehat{\mathcal{M}}^{E_i^H} = -1$ ($i = 2, \dots, g$) and $\#\widehat{\mathcal{M}}^{E_1^H + D_2^H} + \#\widehat{\mathcal{M}}^{E_1^H + E_2^H} = -1$. Then, θ_{β, β^H} is a cycle in \widehat{HF} and $[\theta_{\beta, \beta^H}, 0]$ is a cycle in $HF^{\leq 0}$.

11.2 Proof of Proposition 11.4

Our proof, illustrated schematically in Figure 12, proceeds by a neck-stretching argument. We will show that the moduli spaces used to define the chain map Φ corresponding to the isotopy from $\vec{\beta}$ to $\vec{\beta}'$ are the products of the moduli spaces used to define $F_{\alpha, \beta, \beta'}$ and a moduli space of 1-gons. Before showing this, we need to understand the moduli space of 1-gons.

Fix a Hamiltonian isotopy β_t from β to β' , agreeing with β for $t \ll 0$ and with β' for $t \gg 0$. Let \mathbb{H} denote the upper half plane in \mathbb{C} . The isotopy β_t defines a

Figure 13: Case $g = 2$

g -tuple of Lagrangian cylinders C_β in $\Sigma \times \partial\mathbb{H} \subset \Sigma \times \mathbb{H} =: W_\beta$. We will consider maps $(S, \partial S) \rightarrow (W_\beta, C_\beta)$, with one component of ∂S mapped to each cylinder in C_β . Such maps will converge to some I-chord collection $\vec{x} = \{x_i\} \subset \beta \cap \beta'$ at $\infty \in \partial\mathbb{H}$. Let $\pi_2(\vec{x})$ denote the collection of homology classes of such maps.

Lemma 11.6 *The map $n_3: \pi_2(\vec{x}) \rightarrow \mathbb{Z}$ given by $n_3(A) = \#(A \cap (\mathbb{H} \times \{3\}))$ is an isomorphism.*

Proof The β -circles are homologically linearly independent. \square

For a given almost complex structure $J_{\beta, \beta'}$ on $W_{\beta, \beta'}$ we work with complex structures J on W_β satisfying the obvious analogs of (J'1)–(J'4) from Subsection 10.2. For $A \in \pi_2(\vec{x})$, let \mathcal{M}^A denote the moduli space of embedded J -holomorphic curves in homology class A .

Lemma 11.7 *For $A \in \pi_2(\vec{x})$, $\text{ind}(A) = 2n_3(A)$.*

Proof Let $\gamma_1, \dots, \gamma_g$ be curves in Σ such that $(\Sigma, \vec{\beta}, \vec{\gamma})$ and $(\Sigma, \vec{\beta}', \vec{\gamma})$ are both the standard genus g Heegaard diagram for S^3 . Let \vec{y} be the intersection point between $\vec{\beta}$ and $\vec{\gamma}$, and \vec{z} the intersection point between $\vec{\beta}'$ and $\vec{\gamma}$. Let T_i be a small triangle connecting \vec{x}_i, \vec{y} and \vec{z} , as in Figure 13.

Then, $T_1 + \dots + T_g \in \pi_2(\vec{x}, \vec{y}, \vec{z})$ has $\text{ind}(T_1 + \dots + T_g) = 0$. Further, for $A \in \pi_2(\vec{x})$, $A + T_1 + \dots + T_g \in \pi_2(\vec{y}, \vec{z})$ has $\text{ind}(A + T_1 + \dots + T_g) = 2n_3(A)$. The result follows by additivity of the index under gluings. \square

Lemma 11.8 *There is a choice of C_β such that for any split complex structure on W_β , if $\vec{x} \neq \vec{\theta}_{\beta, \beta'}$ then for $A \in \pi_2(\vec{x})$ with $\text{ind}(A) = 2n_3(A) = 0$, $\mathcal{M}^A = \emptyset$.*




Figure 14: A Hamiltonian isotopy of β_i

Proof Note that if one understands symplectic field theory in the Morse–Bott case this is easy to prove by taking $\beta' \rightarrow \beta$. To avoid introducing this machinery we will instead give a somewhat perturbed argument.

Let t denote the first coordinate under the identification $\mathbb{H} = (-\infty, \infty) \times [0, \infty)$. We assume that for $t \in [-1/4, 1/4]$, $C_{\beta, \beta'}$ is the graph of an exact Hamiltonian isotopy of the following form. Fix δ with $0 < \delta < \pi$. Identify a neighborhood in Σ of each β_i with $S^1 \times (-2\delta, 2\delta)$. Let $\theta_i \in [0, 2\pi)$ and $r_i \in (-\delta, \delta)$ be coordinates on the i^{th} neighborhood. Fix a bump function $b(\theta)$ on the circle which is δ on $[\delta, \pi - \delta]$, 0 on the interval $[\pi + \delta, 2\pi - \delta]$, and monotone on $[2\pi - \delta, \delta]$ and $[\pi - \delta, \pi + \delta]$. For some fixed collection of constants C_i , $i = 1, \dots, g$, consider the Hamiltonian H given by $C_i + \epsilon(\sin(\theta_i) + b(\theta)r_i)$ on the i^{th} neighborhood, and extended arbitrarily outside the neighborhoods of the β_i . Here, we choose ϵ small enough that for each i the graph with respect to β_i of H is contained in the chosen neighborhood of β_i up to time 1. Then for each i , β'_i is the time 1 graph with respect to β_i of the Hamiltonian isotopy specified by H . Thus, the Hamiltonian isotopy of the β -curves looks like Figure 14.

Let A be as in the statement of the proposition and $u: S \rightarrow W_\beta$ be a holomorphic representative of A . Note that S is a disjoint union of g disks, and $\pi_\Sigma \circ u(\partial S)$ is a collection of g pairwise disjoint simple closed curves in $\bigcup_{\theta \in S^1} C_\theta$. Now, by considering orientations and the path taken by $\pi_\Sigma \circ u|_{\partial S}$, we see that u can not exist unless $\vec{x} = \vec{\theta}_{\beta, \beta'}$. See Figure 15. This proves the lemma. \square

Proof of Proposition 11.4 We will write the proof in the notation for HF^∞ , but the proofs for the other theories are just the same.




Figure 15: Two possible paths $\pi_\Sigma \circ u|_{\partial S}$ (shown in bold)

As usual, we let t denote the \mathbb{R} -coordinate on W . Let $C_\alpha = \alpha \times \{1\} \times \mathbb{R} \subset W$, and let $C_{\beta, \beta'}$ be a g -tuple of Lagrangian cylinders which agree with $\beta \times \{0\} \times \mathbb{R}$ for $t < -\frac{1}{4}$ and with $\beta' \times \{0\} \times \mathbb{R}$ for $t > \frac{1}{4}$.

For a given complex structure J on W , let J_R denote the complex structure obtained from J by inserting a neck of length R along the hypersurface $\Sigma \times \{|z| = 1/2, \Re(z) > 0\} \subset \Sigma \times [0, 1] \times \mathbb{R}$. (See Subsection A.2 for a discussion of the splitting process.)

Let $\Phi_R: CF^\infty(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}) \rightarrow CF^\infty(\Sigma, \vec{\alpha}, \vec{\delta}, \vec{\gamma})$ be the map defined in Section 9, with respect to the complex structure J_R . We showed in Section 9 that Φ_R is an isomorphism for each (J, R) such that J_R achieves transversality.

Taking the limit $R \rightarrow \infty$, W splits into two spaces. One is the space we denoted $W_{\alpha, \beta, \beta'}$ in Section 10. The other is a copy of W_β .

Choose J so that:

- The complex structure $J_{\alpha, \beta, \beta'}$ induced on $W_{\alpha, \beta, \beta'}$ satisfies **(J'2)** (the other conditions in the triangles section are automatic).

- The complex structure $J_{\alpha,\beta,\beta'}$ and the complex structure J_β induced on W_β achieve transversality for curves of index 0 (and hence so does J_R for large R).

For such a J it follows from Proposition A.2 that, for R large,

$$\Phi_R^\infty([\vec{x}, i]) = \sum_{\vec{z}} \sum_{\substack{A \in \pi_2(\vec{x}) \\ B \in \pi_2(\vec{x}, \vec{y}, \vec{z}) \\ \text{ind}(A+B)=0}} \#(\mathcal{M}^A) \#(\mathcal{M}^B)[\vec{y}, i - n_3(A+B)]$$

where the first sum is over all intersection points between $\vec{\beta}$ and $\vec{\beta}'$.

By additivity of the index, $\#(\mathcal{M}^A) \#(\mathcal{M}^B) = 0$ unless both A and B have index 0. So, by Lemma 11.8, it follows that

$$\Phi_R^\infty = (\#\mathcal{M}(\vec{\theta}_{\beta,\beta'})) \cdot F_{\alpha,\beta,\delta}(\cdot \otimes \theta_{\beta,\beta'})$$

where $\mathcal{M}(\vec{\theta}_{\beta,\beta'})$ denotes the index 0 holomorphic maps in W_β asymptotic to $\vec{\theta}_{\beta,\beta'}$.

It only remains to understand $\mathcal{M}(\vec{\theta}_{\beta,\beta'})$. We do this by an indirect argument. Suppose that $(\Sigma, \vec{\alpha}, \vec{\beta})$ is the standard Heegaard diagram for S^3 . Then $\widehat{CF}(\Sigma, \vec{\alpha}, \vec{\beta}) = \widehat{CF}(\Sigma, \vec{\alpha}, \vec{\beta}') = \mathbb{Z}$ and the boundary maps are trivial. So, for Φ_R to be an isomorphism it must be multiplication by ± 1 . Further, by explicitly counting triangles (there is only one) we know that $F_{\alpha,\beta,\beta'}$ is the identity map. It follows that $\#\mathcal{M}(\vec{\theta}_{\beta,\beta'}) = \pm 1$. This finishes the proof. \square

11.3 Proof of Proposition 11.2

Lemma 11.9 *In the pointed Heegaard quadruple–diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\beta}^H, \vec{\beta}', \mathfrak{z})$, both $\delta H^1(Y_{\beta,\beta'})$ and $\delta H^1(Y_{\alpha,\beta^H})$ are the trivial subgroup of $H^2(X_{\alpha,\beta,\beta^H,\beta'}, \partial X_{\alpha,\beta,\beta^H,\beta'})$.*

Proof (Compare the proof of [21, Lemma 9.6])

The following diagram commutes:

$$\begin{array}{ccc}
 H^2(X_{\alpha,\beta,\beta^H,\beta'}, \partial X_{\alpha,\beta,\beta^H,\beta'}) & \xleftarrow[\cong]{P.D.} & H_2(X_{\alpha,\beta,\beta^H,\beta'}) \xleftarrow[\cong} \ker \left(\begin{array}{c} H_1(\alpha) \oplus H_1(\beta) \\ \oplus \\ H_1(\beta^H) \oplus H_1(\beta') \end{array} \rightarrow H_1(\Sigma) \right) \\
 \uparrow & & \uparrow \\
 H^1(\partial X_{\alpha,\beta,\beta^H,\beta'}) & \xleftarrow[\cong]{P.D.} & H_2(\partial X_{\alpha,\beta,\beta^H,\beta'}) \xleftarrow[\cong} \left(\begin{array}{c} \ker(H_1(\alpha) \oplus H_1(\beta) \rightarrow H_1(\Sigma)) \\ \oplus \\ \ker(H_1(\beta) \oplus H_1(\beta^H) \rightarrow H_1(\Sigma)) \\ \oplus \\ \ker(H_1(\beta^H) \oplus H_1(\beta') \rightarrow H_1(\Sigma)) \\ \oplus \\ \ker(H_1(\beta') \oplus H_1(\alpha) \rightarrow H_1(\Sigma)) \end{array} \right)
 \end{array}$$

It is clear that the right hand vertical map is surjective. So, from the long exact sequence of the pair $(X_{\alpha,\beta,\beta^H,\beta'}, \partial X_{\alpha,\beta,\beta^H,\beta'})$, the map $H_2(X_{\alpha,\beta,\beta^H,\beta'}) \rightarrow H_1(\partial X_{\alpha,\beta,\beta^H,\beta'}) \cong H^2(\partial X_{\alpha,\beta,\beta^H,\beta'})$ is injective. But the image of $\delta H^1(Y_{\beta,\beta'})$ or $\delta H^1(Y_{\alpha,\beta^H})$ in the latter group is obviously 0. \square

Corollary 11.10 (Compare [21, Lemma 9.6])

$$\widehat{F}_{\alpha,\beta^H,\beta'} \left(\widehat{F}_{\alpha,\beta,\beta^H}(\cdot \otimes \theta_{\beta,\beta^H}) \otimes \theta_{\beta^H,\beta'} \right) = \widehat{F}_{\alpha,\beta,\beta'} \left(\cdot \otimes \widehat{F}_{\beta,\beta^H,\beta'}(\theta_{\beta,\beta^H} \otimes \theta_{\beta^H,\beta'}) \right).$$

Similar statements hold for F^∞ , F^+ and F^- with $[\theta, 0]$ (appropriately subscripted) in place of θ .

Proof From the previous lemma, the $\delta H^1(Y_{\beta,\beta'}) + \delta H^1(Y_{\alpha,\beta^H})$ -orbit of $\text{Spin}^{\mathbb{C}}$ -structures which restrict to \mathfrak{s} has just one element. Further, strong admissibility for the quadruple is equivalent to strong admissibility of the six Heegaard diagrams $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z}), \dots, (\Sigma, \vec{\beta}^H, \vec{\beta}', \mathfrak{z})$, which in turn follows from admissibility for $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$. So, the result follows from Proposition 10.29. \square

Lemma 11.11 Let $\mathfrak{o}_{\alpha,\beta}$ be a coherent orientation system for $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$. Then there is a coherent orientation system for the Heegaard quadruple–diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \vec{\gamma}, \vec{\delta}, \mathfrak{z})$ extending $\mathfrak{o}_{\alpha,\beta}$ and the orientation systems constructed in Proposition 11.3.

(In fact, it is not hard to see that the coherent orientation system for the Heegaard quadruple–diagram is essentially unique (compare Proposition 8.3).)

Proof The orientation systems $\mathfrak{o}_{\alpha,\beta}$ and $\mathfrak{o}_{\beta,\beta^H}$ extend to orientation systems $\mathfrak{o}_{\alpha,\beta,\beta^H}$ and $\mathfrak{o}_{\alpha,\beta^H}$ by Lemma 10.10. For $\vec{\beta}'$ close to $\vec{\beta}$, $\mathfrak{o}_{\alpha,\beta}$ induces $\mathfrak{o}_{\alpha,\beta'}$, which extends to an orientation system $\mathfrak{o}_{\alpha,\beta',\beta^H}$. It is clear that we can perform the extensions so that

the induced orientations $\circ_{\beta',\beta H}$ and $\circ_{\beta,\beta H}$ are consistent with $\circ_{\beta,\beta H,\beta'}$. The result follows. \square

Lemma 11.12 *The map $F_{\alpha,\beta,\beta H}^\infty(\cdot \otimes [\theta_{\beta,\beta H}, 0])$ commutes with the actions of U and $H_1(Y_{\alpha,\beta})/Tors$. Similar statements hold for F^+ , F^- and \widehat{F} .*

Proof The fact that $F_{\alpha,\beta,\beta H}(\cdot \otimes [\vec{\theta}_{\beta,\beta H}, 0])$ commutes with the U -action is obvious. We check that it commutes with the $H_1/Tors$ -action.

Fix a homology class $\zeta \in H_1(Y)$. Choose a knot $K \hookrightarrow \Sigma$ representing ζ and meeting α transversely. Let \mathcal{M}_1 denote the space of Riemann surfaces with $2g$ boundary punctures (as usual) and one additional marked point p on the boundary. We give yet another definition of the map A_ζ for a given $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$. Given a homology class $B \in \pi_2(\vec{x}, \vec{y})$, let \mathcal{M}_K^B denote the moduli space consisting of holomorphic maps from surfaces $S \in \mathcal{M}_1$ to $W_{\alpha,\beta}$ in the homology class B mapping p to $K \times \{1\} \times \mathbb{R}$, and $\widehat{\mathcal{M}}_K^B = \mathcal{M}_K^B/\mathbb{R}$. (Note that this definition is slightly different in form from the one we used to prove Proposition 8.6 and Lemma 9.6. This definition is convenient here, but either would work.)

Now, for $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, let \mathcal{M}_K^A denote the space of holomorphic maps from surfaces $S \in \mathcal{M}_1$ to $W_{\alpha,\beta,\beta H}$ in the homology class A mapping p to $K \times e_1$. (Recall that e_1 is the edge of the triangle T corresponding to the α -circles.) Define a map $H: CF^\infty(Y_{\alpha,\beta}) \otimes CF^\infty(Y_{\beta,\beta H}) \rightarrow CF^\infty(Y_{\alpha,\beta H})$ by

$$H([\vec{x}, i] \otimes [\vec{y}, j]) = \sum_{\vec{z}} \sum_{\substack{A \in \pi_2(\vec{x}, \vec{y}, \vec{z}) \\ \text{ind}(A)=0}} \#(\mathcal{M}_K^A)[\vec{z}, i + j - n_3(A)].$$

For $A \in \pi_2(\vec{x}, \vec{y}, \vec{z})$, $\text{ind}(A) = 1$, consider the space \mathcal{M}_K^A . This is a one-dimensional manifold the ends of which are height two holomorphic buildings with an index 0 holomorphic curve in $W_{\alpha,\beta,\beta H}$ and an index 1 holomorphic curve (defined up to translation) in one of $W_{\alpha,\beta}$, $W_{\beta,\beta H}$ or $W_{\alpha,\beta H}$. For each of these buildings, either

- (1) p is mapped to $W_{\alpha,\beta,\beta H}$ or
- (2) p is mapped to $W_{\alpha,\beta}$ or
- (3) p is mapped to $W_{\alpha,\beta H}$.

Summing the number of ends over the different choices of A with $n_3(A) = k$ we obtain, in each case respectively,

- (1) the coefficient of $[\vec{z}, i + j - k]$ in $(\partial \circ H + H \circ \partial)([\vec{x}, i], [\vec{y}, j])$

- (2) the coefficient of $[\vec{z}, i + j - k]$ in $f_{\alpha, \beta, \beta^H}^\infty(A_\xi([\vec{x}, i]), [\vec{y}, j])$ or
 (3) the coefficient of $[\vec{z}, i + j - k]$ in $A_\xi(f_{\alpha, \beta, \beta^H}^\infty([\vec{x}, i], [\vec{y}, j]))$.

This proves that on the level of homology the $f_{\alpha, \beta, \beta^H}^\infty$ commute with A_ξ . As usual, the result for HF^+ and $HF^{\leq 0}$ follow; the result for \widehat{HF} is proved in an analogous way. \square

Proof of Proposition 11.2 Having proved Corollary 11.10, Proposition 11.4 and Proposition 11.3, we are completely justified in writing

$$\begin{aligned} \widehat{F}_{\alpha, \beta^H, \beta'} \left(\widehat{F}_{\alpha, \beta, \beta^H}(\cdot \otimes \vec{\theta}_{\beta, \beta^H}) \otimes \vec{\theta}_{\beta^H, \beta'} \right) &= \widehat{F}_{\alpha, \beta, \beta'} \left(\cdot \otimes \widehat{F}_{\beta, \beta^H, \beta'}(\vec{\theta}_{\beta, \beta^H} \otimes \vec{\theta}_{\beta^H, \beta'}) \right) \\ &= \widehat{F}_{\alpha, \beta, \beta'}(\cdot \otimes \vec{\theta}_{\beta, \beta'}) \\ &= \widehat{\Phi}_{\beta, \beta'}(\cdot) \end{aligned}$$

where $\widehat{\Phi}_{\beta, \beta'}$ is the isomorphism induced by the isotopy from $\vec{\beta}$ to $\vec{\beta}'$.

It follows that $\widehat{F}_{\alpha, \beta^H, \beta'}$ is surjective and $\widehat{F}_{\alpha, \beta, \beta^H}$ is injective. The same argument with the roles of β and β^H exchanged shows that $\widehat{F}_{\alpha, \beta, \beta^H}$ is surjective. It then follows from Lemma 11.12, that $\widehat{F}_{\alpha, \beta, \beta^H}$ is an isomorphism of $\mathbb{Z}[U] \otimes_{\mathbb{Z}} \Lambda^* H^1(Y) / \text{Tors}$ -modules, proving that \widehat{HF} is invariant under handleslides. The proofs for HF^- , HF^+ and HF^∞ are just the same. \square

12 Stabilization

We show that the homology groups defined in Section 8 are invariant under stabilization, or equivalently under taking connected sum of the Heegaard diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ with the standard Heegaard diagram for S^3 .

Because of handleslide invariance, it is enough to prove the result if we take the connected sum at the point \mathfrak{z} . Let $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ denote the original (pointed) Heegaard diagram, $(\mathcal{T}, \alpha_{g+1}, \beta_{g+1})$ the standard Heegaard diagram for S^3 , and $(\Sigma', \vec{\alpha}', \vec{\beta}', \mathfrak{z}')$ a Heegaard diagram obtained by taking a connect sum near \mathfrak{z} .

The result for the hat theories is quite easy. There is an identification between intersection points in Σ and intersection points in Σ' , identifying an intersection point $\vec{x} = \{x_1, \dots, x_g\}$ in Σ with the intersection point $\vec{x}' = \{x_1, \dots, x_g, \alpha_{g+1} \cap \beta_{g+1}\}$. There is then an obvious identification of $\pi_2^\Sigma(\vec{x}, \vec{y})$ with $\pi_2^{\Sigma'}(\vec{x}', \vec{y}')$. Fix a homology

class $A \in \pi_2^\Sigma(\vec{x}, \vec{y})$ with $n_z(A) = 0$. There is an inclusion map $\widehat{\mathcal{M}}_\Sigma^A \hookrightarrow \widehat{\mathcal{M}}_{\Sigma'}^A$, which takes any holomorphic curve $u \in \widehat{\mathcal{M}}_\Sigma^A$ to the disjoint union

$$u \coprod ((\alpha_{g+1} \cap \beta_{g+1}) \times [0, 1] \times \mathbb{R}).$$

This map is clearly onto, and hence a homeomorphism. With these identifications, the chain complex \widehat{CF} for the stabilized diagram is isomorphic to the chain complexes for the unstabilized one.

The other theories require more work.

Our strategy is to insert a longer and longer neck between the original Heegaard diagram Σ and the torus \mathcal{T} that has been spliced in. We will show that in the limit (neck length) $\rightarrow \infty$, the moduli spaces of curves we consider, correctly defined, are naturally identified with the moduli spaces for $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$. A gluing result then shows that the moduli spaces in the limit can be identified with those of large neck length. Since the Floer homologies are independent of complex structure, the invariance under stabilization will then follow.

Actually, although everything we say in this section is correct, it is somewhat dishonest: we leave most of the work for Proposition A.3 in the appendix. In particular, because of transversality issues, to prove the necessary gluing result we seem to need to work in a more harshly perturbed setting than we use in this section. Much of the work of the present section is redone in the proof of Proposition A.3 in the more general context. For example, we prove there a stronger version of the compactness result Proposition 12.4.

I have chosen to write this section in this slightly dishonest way for two reasons. Firstly, perhaps someone else will see a simpler way to correct the dishonesty (see the discussion at the beginning of the proof of Proposition 12.4). Secondly, as written this section is quite explicit, so one can actually see what is going on in the stabilization proof.

Returning to mathematics, observe that from the formula for the index in Section 4 we know that for a homology class $A \in \pi_2^\Sigma(\vec{x}, \vec{y})$ corresponding to $A' \in \pi_2^{\Sigma'}(\vec{x}, \vec{y})$, $\text{ind}(A) = \text{ind}(A')$. Since we only consider curves with index 1 when computing Floer homologies, we shall generally restrict to curves of index 1. This allows us to keep our definitions simpler and our theorems true.

We make precise what we mean by stretching the neck. Fix complex structures j_Σ on Σ and $j_{\mathcal{T}}$ on \mathcal{T} . Fix a point $\mathfrak{z}_0 \in \mathcal{T} \setminus (\alpha_{g+1} \cup \beta_{g+1})$. Choose small disks $D_\Sigma \supset D'_\Sigma \ni \mathfrak{z}$ and $D_{\mathcal{T}} \supset D'_{\mathcal{T}} \ni \mathfrak{z}_0$ so that $D_\Sigma \setminus D'_\Sigma$ is conformally identified with $S^1 \times [-1, 0]$ and

$D_{\mathcal{T}} \setminus D'_{\mathcal{T}}$ is conformally identified with $S^1 \times [0, 1]$. Here, $\partial D'_{\Sigma}$ is identified with $S^1 \times \{-1\}$ and $\partial D'_{\mathcal{T}}$ is identified with $S^1 \times \{1\}$. (The complex structure on a cylinder $S^1 \times [a, b]$ is given by identifying the cylinder with $\{e^{r+i\theta} \in \mathbb{C} \mid a \leq r \leq b\}$ under the identification $(\theta, t) \mapsto e^{t+i\theta}$.)

Let j_R be the complex structure on

$$\Sigma'_R = (\Sigma \setminus D'_{\Sigma}) \cup_{D_{\Sigma} \setminus D'_{\Sigma} \xrightarrow{\sim} S^1 \times [-R-1, -R]} S^1 \times [-R-1, R+1] \cup_{S^1 \times [R, R+1] \xrightarrow{\sim} \partial D_{\mathcal{T}}} (\mathcal{T} \setminus D_{\mathcal{T}})$$

induced by the complex structures on Σ , \mathcal{T} and $S^1 \times [-R-1, R+1]$, and ω_R the area form. We refer to $S^1 \times [-R-1, R+1]$ as the neck, and denote it N_R . Notice that in the limit $R \rightarrow \infty$, j_R degenerates to the complex structure $j_{\infty} = j_{\Sigma} \vee j_{\mathcal{T}}$ on $\Sigma \vee \mathcal{T}$.

Fix a complex structure J on $W'_0 = \Sigma'_0 \times [0, 1] \times \mathbb{R}$ satisfying **(J1)**–**(J5)**, which is split (ie, $j_0 \times j_{\mathbb{D}}$) near N_0 . Let J_R be the complex structure on $W'_R = \Sigma'_R \times [0, 1] \times \mathbb{R}$ which agrees with J outside $N_R \times [0, 1] \times \mathbb{R}$ and with $j_R \times j_{\mathbb{D}}$ on $N_R \times [0, 1] \times \mathbb{R}$. Note that J_R converges to a complex structure J_{∞} on $W'_{\infty} = (\Sigma \vee \mathcal{T}) \times [0, 1] \times \mathbb{R}$. The space $W = \Sigma \times [0, 1] \times \mathbb{R}$ lies inside W'_{∞} in an obvious way. We choose J so that

- (1) There is a collection of $R_i \rightarrow \infty$ such that J_{R_i} achieves transversality for holomorphic curves in W'_{R_i} satisfying **(M0)**–**(M6)**.
- (2) The restriction of J_{∞} to W achieves transversality for holomorphic curves in W satisfying **(M0)**–**(M6)**.

Since by choosing J appropriately we can end up with any J_{∞} which is split near \mathfrak{z} , we can always find such a J .

For a generic choice of \mathfrak{z} and for any holomorphic curve u with $\text{ind}(u) = 1$, $\pi_{\mathbb{D}}((\pi_{\Sigma} \circ u)^{-1}(\mathfrak{z}))$ consists of $n_{\mathfrak{z}}(u)$ distinct points. We choose \mathfrak{z} satisfying this condition.

The holomorphic curves that we consider in W'_{∞} are *holomorphic twin towers*. That is, fix a homology class $A \in \pi_2^{\Sigma}(\vec{x}, \vec{y}) \cong \pi_2^{\Sigma'}(\vec{x}', \vec{y}')$.

By a holomorphic twin tower u in the homology class A we mean a collection of holomorphic maps $(u_1, \dots, u_n, v_1, \dots, v_n)$ where $u_1, \dots, u_n: S_1, \dots, S_n \rightarrow W$ and $v_1, \dots, v_n: S'_1, \dots, S'_n \rightarrow T^2 \times [0, 1] \times \mathbb{R}$ such that:

- (1) The u_1, \dots, u_n are holomorphic curves in W satisfying **(M0)**–**(M6)**.
- (2) The S'_1, \dots, S'_n are all closed surfaces.

- (3) There is a sequence of intersection points $\vec{x}_1, \dots, \vec{x}_{n+1}$ such that u_i connects \vec{x}_i to \vec{x}_{i+1} .
- (4) $[u_1] + \dots + [u_n] = A$. Here, $[u_i]$ is the class in $\pi_2(\vec{x}_i, \vec{x}_{i+1})$ represented by u_i .
- (5) The v_1, \dots, v_n are holomorphic curves in $\mathcal{T} \times [0, 1] \times \mathbb{R}$ each connecting $\alpha_{g+1} \cap \beta_{g+1}$ to itself, satisfying **(M0)**, **(M1)**, **(M3)** and **(M5)**
- (6) As sets, $\pi_{\mathbb{D}}((\pi_{\Sigma} \circ u_i)^{-1}(\mathfrak{z})) = \pi_{\mathbb{D}}((\pi_{\Sigma} \circ v_i)^{-1}(\mathfrak{z}_0))$ for each i . (This is the matching condition on the horizontal levels.)

Remark If we had not restricted to holomorphic curves of index 1 and chosen \mathfrak{z} generically, we would need a slightly more complicated definition of holomorphic twin towers. In particular, we would need to allow pieces of the curves to live in the “horizontal” cylinders $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$.

Lemma 12.1 *If $\text{ind}(A) = 1$ then any holomorphic twin tower in the homology class A has height one (ie, in the previous definition, $n = 1$).*

Proof This is trivial: the index adds between stories. □

Lemma 12.2 *If (u_1, v_1) is a holomorphic twin tower in a homology class of index 1 then v_1 consists of a trivial cylinder $\alpha_{g+1} \cap \beta_{g+1} \times [0, 1] \times \mathbb{R}$ and $n_3(u_1)$ horizontal tori, $T^2 \times \pi_{\mathbb{D}}((\pi_{\Sigma} \circ u_1)^{-1}(\mathfrak{z}))$.*

Proof The restriction of $\pi_{\mathbb{D}} \circ v_1$ to the components on which it is nonconstant is a 1–fold covering of the disk. Hence, it must be a disk itself. It follows easily that the restriction of $\pi_{\mathcal{T}} \circ v_1$ to this component must be constant. Since the condition on z guarantees that the set $(\pi_{\mathbb{D}} \circ v_1)((\pi_{\mathcal{T}} \circ v_1)^{-1}(\mathfrak{z}_0))$ consists of $n_3([u_1]) = n_3([v_1])$ distinct points, each component on which $\pi_{\mathbb{D}} \circ v_1$ is constant must be a copy of T^2 , and the restriction of $\pi_{\mathcal{T}} \circ v_1$ to each component must be a diffeomorphism. □

Corollary 12.3 *The moduli space of holomorphic twin towers in a given homology class $A \in \pi_2^{\Sigma}(\vec{x}, \vec{y})$ with $\text{ind}(A) = 1$ is naturally identified with $\widehat{\mathcal{M}}_W^A$.*

Here, $\widehat{\mathcal{M}}_W^A$ denotes the moduli space of holomorphic curves in the homology class A in W .

We now need to identify the space of holomorphic twin towers with $\widehat{\mathcal{M}}_{W'_R}^A$ for R large enough. As usual, doing so requires two steps: compactness and gluing. To avoid work we torture the compactness argument slightly.




Figure 16: Interpolating Lagrangian cylinders

Proposition 12.4 Fix a sequence $\{u_R\}$ of J_R -holomorphic curves in W' with index 1. Then for \mathfrak{z} chosen generically, there is a subsequence of $\{u_R\}$ which converges to a holomorphic twin tower (of height one).

Proof The proof is in three stages. First we use the fact that $\pi_{\mathbb{D}} \circ u_R$ is holomorphic to extract the vertical level structure and conformal structure on the limit surface. Then we cut Σ into two overlapping regions, and view the u_R as maps into each of the two regions with Lagrangian boundary conditions. This allows us to extract a convergent subsequence of maps to W' .

By classical symplectic field theory [2, Theorem 10.1], we can replace u_R by a subsequence so that $\pi_{\mathbb{D}} \circ u_R$ converges to a holomorphic building, which we denote $\pi_{\mathbb{D}} \circ u_{\infty}$ (although u_{∞} does not yet make sense on its own). We can also assume that for all R the source of u_R is some fixed topological manifold S . Let S_{∞} denote the source of $\pi_{\mathbb{D}} \circ u_{\infty}$.

Let $C = S^1 \times \{0\} \subset \Sigma'_R$ denote the curve in Σ' along which we are splitting. For convenience, let us say that Σ “lies to the left of C ” while T “lies to the right of C .” Let $C_{\Sigma} = S^1 \times \{-R-1\} \subset \Sigma'_R$. Let $C_T = S^1 \times \{-R-1+\rho\} \subset \Sigma'_R$ for some $\rho < R+1$. Then, C_{Σ} and C_T lie to the left of C , and C_T lies to the right of C_{Σ} . Note that the complex structures J_R are split to the right of C_{Σ} . We will choose ρ large enough that C_T lies close to C in a sense we will specify soon. Let RC_{Σ} denote the region to the right of C_{Σ} and LC_T denote the region to the left of C_T . See Figure 16.

Let $\{C_k, A_{\ell}\}$ denote the collection of disjoint circles and arcs in S along which the complex structure degenerates.

Let $S_{RC_{\Sigma}, R} = (\pi_{\Sigma} \circ u_R)^{-1}(RC_{\Sigma})$. Since J_R is split over RC_{Σ} ,

$$\pi_{\Sigma} \circ u_R: (S_{RC_{\Sigma}, R}, \partial S_{RC_{\Sigma}, R}) \rightarrow (RC_{\Sigma}, C_{\Sigma} \cup \alpha_{g+1} \cup \beta_{g+1})$$

is a holomorphic map. So, again by classical symplectic field theory [2, Theorem 10.3], taking a further subsequence we can assume that $\pi_\Sigma \circ u_R|_{S_{RC_\Sigma, R}}$ converges to a holomorphic building. We denote this building by $\pi_\Sigma \circ u_\infty$, with the understanding that it is, so far, only defined over $S_{RC_\Sigma, \infty}$.

It now makes sense to talk about the circles in $\{C_k\}$ which correspond to the degeneration along C – ie, the level splitting of $\pi_\Sigma \circ u_R$. Let $\{p_j\}$ denote the corresponding points in S_∞ .

Since the arcs and circles $\{C_k, A_\ell\}$ are disjoint, we can choose C_T close enough to C (ie, ρ large enough) that $(\pi_\Sigma \circ u)^{-1}(C_T)$ is disjoint from all the C_k and A_ℓ . We can also choose ρ so that $\pi_\Sigma \circ u_\infty$ is transverse to C_T . We do so choose it. It follows that for R large enough $\pi_\Sigma \circ u_R$ is transverse to C_T .

Now, observe that each curve in $(\pi_{\mathbb{D}} \circ u_R)(\pi_\Sigma \circ u_R)^{-1}(C_T)$ converges in the C^∞ -topology as $R \rightarrow \infty$. Let $B_R = (\pi_{\mathbb{D}} \circ u_R)(\pi_\Sigma \circ u_R)^{-1}(C_T)$. Let $B_\infty = \lim_{R \rightarrow \infty} B_R$.

Let LC_T denote the portion of Σ' to the left of C_T . View $[0, 1]$ as lying inside $S^1 = [0, 2]/(0 \sim 2)$, say. We can, thus, consider $\pi_{\mathbb{D}} \circ u_R$ as a map to $S^1 \times \mathbb{R}$. Consider the symplectic 4-manifold with boundary $LC_T \times S^1 \times \mathbb{R}$, given the obvious (split) symplectic form. For each R , $C_T \times B_R$ is an immersed Lagrangian submanifold, and these submanifolds C^∞ converge to $C_T \times B_\infty$.

So, applying the compactness theorem [2, Theorem 10.1] for symplectic field theory to $u_R|_{S_{LC_T, R}}$, viewed as holomorphic curves with dynamic Lagrangian boundary conditions, we can extract a subsequence converging to a holomorphic building $u_\infty|_{S_{LC_T, \infty}}$ defined over $S_{LC_T, \infty}$.

Let $q_j = \pi_{\mathbb{D}} \circ u_\infty(p_j)$. There are $n_3([u_R])$ points q_j , all of them distinct. It follows that in the (horizontal) cylindrical regions $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$ connecting $\Sigma \times [0, 1] \times \mathbb{R}$ and $T^2 \times [0, 1] \times \mathbb{R}$, the building $u_\infty|_{S_{LC_T, \infty}}$ consists of trivial cylinders, so we can ignore these regions.

The holomorphic building $u_\infty|_{S_{LC_T, \infty}}$ must agree with $(\pi_\Sigma \circ u_\infty) \times (\pi_{\mathbb{D}} \circ u_\infty)$ where both are defined, so we can patch the two together to obtain a holomorphic twin tower to which (the subsequence of) the sequence u_R converges. \square

Proposition 12.5 *The Floer homologies HF^∞ , HF^+ and HF^- are invariant under stabilization.*

Proof This follows from the previous proposition and Proposition A.3. That is, for given intersection points the space of holomorphic twin towers (of height one, with index 1) connecting \vec{x}' to \vec{y}' is identified with $\widehat{\mathcal{M}}(\vec{x}, \vec{y})$. The admissibility criteria ensure

that only finitely many homology classes matter. Thus, by the previous proposition and Proposition A.3, we can use the space of holomorphic twin towers to compute the boundary maps in W' .

Using our original definition of the action of $H_1(Y)/Tors$ it is immediate that the two actions are the same. Obviously the U -actions correspond. \square

Proof of Theorem 1. By Proposition 5.6, there is an admissible Heegaard diagram for (Y, \mathfrak{s}) , and any two admissible Heegaard diagrams can be connected by a sequence of pointed Heegaard moves, ie, Heegaard moves supported in the complement of \mathfrak{z} . Thus, the theorem is immediate from Proposition 9.1, Proposition 11.2, and Proposition 12.5.

13 Comparison with Heegaard Floer homology

In this section we prove the equivalence of the theory described in this paper with Heegaard Floer homology as originally defined by Ozsváth and Szabó in [21]. For notational convenience, we will phrase the argument in terms of HF^∞ , but the same proof works for all four theories. This section assumes familiarity with [21].

By CF_{ours} we mean the chain complexes defined in this paper; by CF_{theirs} we mean the chain complex defined in [21]. We extend this notation functorially to $HF_{\text{ours/theirs}}$, $\pi_2^{\text{ours/theirs}}$, etc. When we have successfully identified two corresponding objects we drop the decorations “ours” or “theirs”.

Observe that there is an identification between our intersection points and those of [21], and so between the generators of CF_{ours}^∞ and the generators of $CF_{\text{theirs}}^\infty$. Similarly, for any intersection points \vec{x} and \vec{y} , $\pi_2^{\text{ours}}(\vec{x}, \vec{y})$ and $\pi_2^{\text{theirs}}(\vec{x}, \vec{y})$ are naturally identified, by considering domains for example.

Now we deal with a simple case. Suppose that the Heegaard diagram $(\Sigma, \vec{\alpha}, \vec{\beta}, \mathfrak{z})$ is such that the split complex structure $j_\Sigma \times j_\mathbb{D}$ achieves transversality for all homology classes of index 1 in our theory and such that $\text{Sym}^g(j_\Sigma)$ achieves transversality for domains of index 1 in their theory. Let $u: S \rightarrow W$ be a holomorphic curve with respect to $j_\Sigma \times j_\mathbb{D}$. Define a map $u': \mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ as follows. For $p \in \mathbb{D}$, let $(\pi_\mathbb{D} \circ u)^{-1}(p)$ denote the g preimages of p , listed with multiplicities. Then $u'(p) = \pi_\Sigma \circ u((\pi_\mathbb{D} \circ u)^{-1}(p))$ is a point in $\text{Sym}^g(\Sigma)$. It is easy to check that u being $(j_\Sigma \times j_\mathbb{D})$ -holomorphic implies that u' is holomorphic with respect to $\text{Sym}^g(j_\Sigma)$. So, for any $A \in \pi_2(\vec{x}, \vec{y})$ we have a map $\Phi: \widehat{\mathcal{M}}_{\text{ours}}^A \rightarrow \widehat{\mathcal{M}}_{\text{theirs}}^A$, which is clearly injective.

In [21, Lemma 3.6], Ozsváth and Szabó construct an inverse for Φ as follows. Fix a $\text{Sym}^g(j_\Sigma)$ -holomorphic map $u': \mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$. Let $(g-1)!u_\mathbb{D}: (g-1)!S \rightarrow \mathbb{D}$

and $(g - 1)!u_\Sigma: (g - 1)!S \rightarrow \Sigma^g$ denote the pullback via u of the branched cover $\Sigma^g \rightarrow \text{Sym}^g(\Sigma)$, as in the following diagram:

$$\begin{array}{ccc} (g - 1)!S & \xrightarrow{(g-1)!u_\Sigma} & \Sigma^g \\ \downarrow (g-1)!u_\mathbb{D} & & \downarrow p \\ \mathbb{D} & \xrightarrow{u'} & \text{Sym}^g(\Sigma) \end{array}$$

Since we are in the holomorphic category, pullbacks of branched covering maps are well-defined.

The symmetric group S_g acts on Σ^g by permuting the factors, and hence acts on $(g - 1)!S$. The maps $(g - 1)!u_\Sigma$ and $(g - 1)!u_\mathbb{D}$ are S_g -equivariant. Let S_{g-1} denote the permutations fixing the first factor in $\Sigma^g = \Sigma \times \Sigma^{g-1}$. Let $\pi: \Sigma^g \rightarrow \Sigma$ denote projection onto the first factor. Let $S = (g - 1)!S/S_{g-1}$. Let $u_\Sigma = (\pi \circ (g - 1)!u) / S_{g-1}: S \rightarrow \Sigma$ denote the map induced by $\pi \circ (g - 1)!u: (g - 1)!S \rightarrow \Sigma$, and $u_\mathbb{D} = ((g - 1)!u_\mathbb{D}) / S_{g-1}: S \rightarrow \mathbb{D}$ the map induced by $(g - 1)!u_\mathbb{D}$. Then, $u = u_\Sigma \times u_\mathbb{D}$ is a holomorphic map $S \rightarrow W$. Define $\Phi^{-1}(u') = u$.

Remark In the context of multivalued sections of Lefschetz pencils, the maps Φ and Φ^{-1} are called the *tautological correspondence* by M Usher [28], who attributes the term to I Smith. We will sometimes use this terminology below.

We check that $\Phi \circ \Phi^{-1}$ is the identity map. It suffices to show that $\Phi \circ \Phi^{-1}(u')$ agrees with u' away from the diagonal. Commutativity of the diagram

$$\begin{array}{ccccc} S & \xrightarrow{u_\Sigma} & \Sigma & & \\ & \searrow & \swarrow \pi & & \\ & & (g - 1)!S & \xrightarrow{(g-1)!u_\Sigma} & \Sigma^g \\ & \searrow u_\mathbb{D} & \downarrow (g-1)!u_\mathbb{D} & & \downarrow p \\ & & \mathbb{D} & \xrightarrow{u'} & \text{Sym}^g(\Sigma) \end{array}$$

shows that $\Phi \circ \Phi^{-1}(u')(x)$, viewed as a set of g distinct points in Σ , agrees with the set $\pi \circ p^{-1}(u'(x))$. (Here, it is important that elements of sets do *not* have multiplicities.) But this set is exactly $u'(x)$.

Observe that both Φ and Φ^{-1} are continuous. So, since Φ was injective, we have proved that Φ is a homeomorphism. This is enough to prove the equivalence of the two theories with $\mathbb{Z}/2$ -coefficients, under our assumption that $j_\Sigma \times j_\mathbb{D}$ achieves transversality. We will deal with the issue of orientations presently; first, we discuss the case when $j_\Sigma \times j_\mathbb{D}$ does not achieve transversality. The difficulty in the case when $j_\Sigma \times j_\mathbb{D}$ does not achieve transversality is that the nearly-symmetric almost complex structures used in [21] are required to be split near the diagonal, while compactness is proved in [2] for almost complex structures which are cylindrical. No non-split almost complex structure satisfies both conditions.

To address this problem, we define a class of complex structures on $\text{Sym}^g(\Sigma)$ which includes both (Sym^g of) the complex structures we consider and the complex structures considered in [21]. Specifically, fix an open neighborhood V_1 of $\{z_i\} \times \text{Sym}^{g-1}(\Sigma) \subset \text{Sym}^g(\Sigma)$ and an open neighborhood V_2 of the diagonal Δ in $\text{Sym}^g(\Sigma)$. Let $\pi: \text{Sym}^g(\Sigma) \rightarrow \Sigma$ denote projection, and $\omega_0 = (dA)^g$ the product symplectic form on Σ^g . As usual, we also fix a complex structure j_Σ on Σ .

Definition 13.1 By a *quasi-nearly-symmetric almost complex structure* on $\text{Sym}^g(\Sigma)$ we mean an almost complex structure \tilde{J} on $\text{Sym}^g(\Sigma)$ such that

- \tilde{J} is tamed by $\pi_*(\omega_0)$ on $\text{Sym}^g(\Sigma) \setminus V_2$.
- \tilde{J} agrees with $\text{Sym}^g(j_\Sigma)$ on V_1 .
- There is some complex structure j on Σ such that \tilde{J} agrees with $\text{Sym}^g(j)$ on V_2 .

Observe that the nearly-symmetric almost complex structures of [21, Section 3.1] are a special case of the preceding definition. So are complex structures of the form $\text{Sym}^g(j)$ for *any* complex structure j on Σ which is tamed by dA and agrees with j_Σ near the z_i . So, a path of nearly-symmetric almost complex structures is a path of quasi-nearly-symmetric almost complex structures. Also, an almost complex structure J on W satisfying **(J1)–(J5)** corresponds to a path J_s of complex structures on Σ , which in turn specifies a path of quasi-nearly-symmetric almost complex structures on $\text{Sym}^g(\Sigma)$, which we denote $\text{Sym}^g(J_s)$.

For a path \tilde{J}_t of quasi-nearly-symmetric almost complex structures, let j_t be the complex structure on Σ such that \tilde{J}_t agrees with $\text{Sym}^g(j_t)$ on V_2 .

Now, we prove the equivalence of our theory with that of [21] in three steps. First we show that the compactness proof of [21] extends to \tilde{J}_t -holomorphic curves in $\text{Sym}^g(\Sigma)$, where \tilde{J}_t is a path of quasi-nearly-symmetric almost complex structures. Then we observe that the class of paths of almost complex structures of the form

$\text{Sym}^g(J_s)$ is sufficient to achieve transversality for holomorphic disks in $\text{Sym}^g(\Sigma)$. Finally we show that our moduli spaces of J_s -holomorphic curves can be identified with the Ozsváth–Szabó moduli spaces of $\text{Sym}^g(J_s)$ -holomorphic curves.

After proving the following proposition, the first and third steps are essentially immediate.

Proposition 13.2 *Fix a path \tilde{J}_t of quasi-nearly-symmetric almost complex structures on $\text{Sym}^g(\Sigma)$. Fix a holomorphic disk $u': (\mathbb{D}, \partial\mathbb{D}) \rightarrow (\text{Sym}^g(\Sigma), T_\alpha \cup T_\beta)$. Then there is a $g!$ -fold branched covering $\tilde{u}_\mathbb{D}: \tilde{S} \rightarrow \mathbb{D}$ and a map $\tilde{u}_\Sigma: \tilde{S} \rightarrow \Sigma^g$ such that the following diagram commutes.*

$$\begin{array}{ccc}
 \tilde{S} & \xrightarrow{\tilde{u}_\Sigma} & \Sigma^g \\
 \downarrow \tilde{u}_\mathbb{D} & & \downarrow p \\
 \mathbb{D} & \xrightarrow{u'} & \text{Sym}^g(\Sigma)
 \end{array}
 .$$

The map \tilde{u}_Σ is holomorphic with respect to the path of almost complex structures on Σ^g induced by \tilde{J}_t , in the obvious sense, and is S_g -equivariant.

Proof The idea of the proof is that even though the complex structure is allowed to vary near the diagonal, since it varies in the class of split structures, locally near the diagonal we are still in the integrable case. The model for this argument is used to prove the following

Lemma 13.3 *Under the assumptions of Proposition 13.2, u' intersects the diagonal Δ in a discrete collection of points.*

Proof Suppose that u' intersects the diagonal Δ in a collection of points p_j with limit point p . Throwing out some of the points we may assume all of the p_j lie in the same stratum of Δ . For concreteness we will assume all of the p_j lie in the top dimensional stratum of Δ , but there is nothing special about this case.

Write $p = \{a_1, a_1, a_3, \dots, a_g\}$. Choose pairwise disjoint disk neighborhoods U_i of the a_i such that $U_1 \times U_1 \times U_3 \times \dots \times U_g$ is contained in V_2 . Choose three points on the boundary of each U_i . Then the Riemann mapping theorem gives a well-defined holomorphic identification of (U_i, j_t) and $(\mathbb{D}, j_\mathbb{D})$ for each t .

So in a neighborhood V of $u'^{-1}(p)$ we can view u' as a map to $(\text{Sym}^g(\mathbb{D}), \text{Sym}^g(j_\mathbb{D}))$. The diagonal $\Delta \subset \text{Sym}^g(\mathbb{D})$ is an analytic subvariety, so by elementary complex analysis, $u'|_V: V \rightarrow \Delta$.

Globally we know that the image of u' is not entirely contained in Δ . So, a standard open–closed argument gives a contradiction. \square

We return to the proof of Proposition 13.2. Away from the diagonal Δ , to obtain \tilde{S} , \tilde{u}_Σ and $\tilde{u}_\mathbb{D}$ we simply pullback the holomorphic covering space $\Sigma^g \rightarrow \text{Sym}^g(\Sigma)$. Near the diagonal, the argument used in the preceding lemma’s proof shows that we can pull back the branched covering (which locally looks like a piece of $((\mathbb{D}^2)^g, j_\mathbb{D}^g) \rightarrow (\text{Sym}^g(\mathbb{D}), \text{Sym}^g(j_\mathbb{D}))$) by u' . This proves the proposition.

Proposition 13.4 *The moduli spaces considered in [21], computed with respect to any quasi–nearly–symmetric almost complex structure, are compact.*

Proof The only place in their proof that Ozsváth and Szabó use the condition that their complex structures are standard near Δ is in the proof of their energy estimate [21, Lemma 3.5]. The only time they use it in that proof is to observe that the previous proposition holds. So, compactness is immediate from their work. \square

Proposition 13.5 *The class of paths of complex structures of the form $\text{Sym}^g(J_s)$ is sufficient to achieve transversality for disks $u: (\mathbb{D}, \partial\mathbb{D}) \rightarrow (\text{Sym}^g(\Sigma), T_\alpha \cup T_\beta)$.*

The proof, which we omit, is a simple adaptation of the one in Section 3.

Proposition 13.6 *Calculated with respect to J_t and $\text{Sym}^g(J_t)$ respectively, $\widehat{\mathcal{M}}_{\text{ours}}^A$ and $\widehat{\mathcal{M}}_{\text{theirs}}^A$ agree.*

The proof of this proposition is the same as the proof in the split case, using Proposition 13.2 where appropriate.

This is sufficient to prove that the two theories are equivalent with $\mathbb{Z}/2$ –coefficients. To prove the equivalence with \mathbb{Z} –coefficients, we need somewhat more. Specifically, we need to check that

Proposition 13.7 (1) *If u_1 and u_2 are two curves in the moduli space \mathcal{M}^A , $\text{ind}(A) = 1$, and $\text{sign}_{\text{ours}}(u_1) = \text{sign}_{\text{theirs}}(u_1)$ then $\text{sign}_{\text{ours}}(u_2) = \text{sign}_{\text{theirs}}(u_2)$. (Here, sign denotes the sign of a rigid curve induced by the coherent orientation.) That is, if the orientations agree at one curve then they agree at all curves in that moduli space.*

(2) *The coherence conditions for the two theories agree. That is, there is an identification of the determinant line bundles $\mathcal{L}_{\text{ours}}$ and $\mathcal{L}_{\text{theirs}}$ with respect to which any coherent orientation $\mathfrak{o}_{\text{ours}}$ of $\mathcal{L}_{\text{ours}}$ specifies a coherent orientation $\mathfrak{o}_{\text{theirs}}$ of $\mathcal{L}_{\text{theirs}}$.*

To keep the exposition clean, we will assume that the complex structure is split near the diagonal. As noted earlier, this assumption is quite restrictive. It can be removed by using, at appropriate times, a parametrized version of the Riemann mapping theorem to identify a neighborhood of u_Σ (branch points of $u_\mathbb{D}$) with a standard disk, as in the proof of Lemma 13.3. (Alternately, given a curve u , we can choose a family $J'_s : T\Sigma \rightarrow T\Sigma$ of almost complex structures on Σ constant near $\Phi(u) \cap \Delta$ and such that the (Σ, J'_s) are conformally isomorphic to the original family (Σ, J_s) , and then work with the family (Σ, J'_s) .) We leave further details to the interested reader.

From now on, when we want to discuss both our and their theories at once, we will drop the “ours” or “theirs” from the notation, even if we have not yet identified the corresponding objects.

Recall from Section 3 that the tangent space at u to $\mathcal{B}_{\text{ours}}$ is $\mathbb{R}^{2g} \oplus L_1^{p,d}(u^*TW, \partial)$, where the \mathbb{R}^{2g} includes into $\Gamma(u^*TW)$ as $\text{Span}(\{v_i^\pm\})$ where the v_i^\pm are some fixed smooth sections with v_i^\pm equal to $\frac{\partial}{\partial t}$ in a small neighborhood of the i^{th} positive or negative puncture and zero outside a slightly larger neighborhood of that puncture and $L_1^{p,d}(u^*TW, \partial)$ denotes sections tangent to $C_\alpha \cup C_\beta$ over ∂S . The tangent space at ϕ to $\mathcal{B}_{\text{theirs}}$ is just $L_1^{p,d}(\phi^*T\text{Sym}^g(\Sigma), \partial)$ of $L_1^{p,d}$ sections tangent to $T_\alpha \cup T_\beta$ over $\partial\mathbb{D}$ and we take as a model for \mathbb{D} the strip $[0, 1] \times \mathbb{R}$.

There is the minor complication that we are working in \mathbb{R} -invariant settings, so rigid curves have $\text{ind} = 1$. For convenience, we introduce the operator $P_{u,\text{ours}} : \mathbb{R}^{2g} \oplus L_k^{p,d}(u^*TW, \partial) \rightarrow \mathbb{R}$ defined by $P_{u,\text{ours}}(v_1, v_2) = \langle v_2, \partial/\partial t \rangle_{L^2}$. (Here, $\langle \cdot, \cdot \rangle_{L^2}$ denotes the L^2 inner product induced by some Riemannian metric on W and a canonical metric on the source S of u . Since $L_k^{p,d}$ is finer than L^2 , the operator $P_{u,\text{ours}}$ is continuous.) Then we replace the linearized $\bar{\partial}$ -map $D\bar{\partial}_{\text{ours}}$ by $\tilde{D}\bar{\partial}_{\text{ours}} = D\bar{\partial}_{\text{ours}} \oplus P_{\text{ours}}$. Similarly, for $\bar{\partial}_{\text{theirs}}$, let $P_{\phi,\text{theirs}}(v) = \langle v, d\phi(\partial/\partial t) \rangle_{L^2}$, where $\partial/\partial t$ generates the one-parameter group of automorphisms of $(\mathbb{D}^2, i, -i)$. Then, replace $D\bar{\partial}_{\text{theirs}}$ with $\tilde{D}\bar{\partial}_{\text{theirs}} = D\bar{\partial}_{\text{theirs}} \oplus P_{\text{theirs}}$. We retain the old meaning of ind , so the index of $\tilde{D}_u\bar{\partial}_{\text{ours/theirs}}$ is $\text{ind}(u) - 1$.

We tackle point (1) of the proposition first. We begin by recalling how the coherent orientation specifies signs of curves. Let u be a curve (not necessarily holomorphic), $\text{ind}(u) = 1$, at which $D\bar{\partial}$ (or equivalently $\tilde{D}\bar{\partial}$) is surjective. We can view the coherent orientation σ as a nonvanishing section of \mathcal{L} , defined up to multiplication by a positive scalar function. At u , there is also a canonically defined section

$$1 \otimes 1^* \in \mathbb{R} \otimes \mathbb{R}^* = (\Lambda^0 \mathbb{R}^0) \otimes (\Lambda^0 \mathbb{R}^0)^* = (\Lambda^{\text{top}} \ker(\tilde{D}_u\bar{\partial})) \otimes (\Lambda^{\text{top}} \text{coker}(\tilde{D}_u\bar{\partial}))^* = \mathcal{L}_u.$$

The sign $\text{sign}(u)$ is $+1$ (respectively -1) if $\sigma(u)$ is a positive (respectively negative) multiple of $1 \otimes 1^*$.

The section $1 \otimes 1^*$ does not extend continuously to the whole configuration space. Indeed, for a generic path $\{u_a\}$ (ie, one with only *normal crossings*) between curves u_0 and u_1 , with $D_{u_0} \bar{\partial}$ and $D_{u_1} \bar{\partial}$ surjective, the section $1 \otimes 1^*$ switches sign at each a for which $D_{u_a} \bar{\partial}$ is not surjective. So, if $c(\{u_a\})$ denotes the number of a for which $D_{u_a} \bar{\partial}$ is not surjective then $\text{sign}(u_0) = (-1)^{c(\{u_a\})} \text{sign}(u_1)$.

It follows that to prove part (1) of Proposition 13.7 it suffices to show that $\ker(\tilde{D} \bar{\partial}_{\text{ours}})$ and $\ker(\tilde{D} \bar{\partial}_{\text{theirs}})$ are nontrivial at the same curves. For this to make sense, we first need to identify the configuration spaces. Actually, we will define subspaces $B_{\text{ours}} \subset \mathcal{B}_{\text{ours}}$ and $B_{\text{theirs}} \subset \mathcal{B}_{\text{theirs}}$ such that

- $B_{\text{ours/theirs}} \supset \mathcal{M}_{\text{ours/theirs}}$ and
- $B_{\text{ours/theirs}}$ contains a path with only normal crossings between any two curves at which $D \bar{\partial}$ is surjective. (In particular, $B_{\text{ours/theirs}}$ is connected.)

Then, we will construct an identification Φ between B_{ours} and B_{theirs} extending the identification Φ between $\mathcal{M}_{\text{ours}}$ and $\mathcal{M}_{\text{theirs}}$. Finally, we will construct an identification of $\ker(\tilde{D} \bar{\partial}_{\text{ours}})$ and $\ker(\tilde{D} \bar{\partial}_{\Phi(u)} \bar{\partial}_{\text{theirs}})$. This suffices to prove part (1) of Proposition 13.7.

None of the steps involved are particularly intricate. The space B_{ours} consists of embedded curves u such that $u_{\mathbb{D}}$ is holomorphic with only order 2 branch points, near which u_{Σ} is holomorphic. The space B_{theirs} consists of curves ϕ intersecting the diagonal Δ only in the top-dimensional stratum, transverse to Δ , and holomorphic in a neighborhood of Δ . It is clear that B_{ours} and B_{theirs} have the requisite properties. The identification $\Phi: B_{\text{ours}} \rightarrow B_{\text{theirs}}$ is given in exactly the same way as $\Phi: \mathcal{M}_{\text{ours}} \rightarrow \mathcal{M}_{\text{theirs}}$ was defined at the beginning of the section.

The identification of kernels is somewhat more difficult. The key point is that near the branch points, the kernel of $\tilde{D} \bar{\partial}_{\text{ours}}$ itself consists of holomorphic sections of $T\Sigma$. Such sections correspond by the same tautological correspondence used to define Φ to holomorphic sections of $T \text{Sym}^g(\Sigma)$, which comprise the kernel of $\tilde{D} \bar{\partial}_{\text{theirs}}$ near the diagonal. More details follow.

Recall [14, page 28] that the linearized $\bar{\partial}$ -operator

$$(D_u \bar{\partial}): \mathbb{R}^{2g} \oplus L_1^{p,d} (u^* T W, \partial) \rightarrow L^{p,d} (\Lambda^{0,1} u^* T W)$$

at a curve $u: (S, \partial S) \rightarrow (W, C_\alpha \cup C_\beta)$, with fixed complex structure on S , is given by

$$D_u \bar{\partial}(\xi)(v) = \frac{1}{2} (\nabla_v \xi + J \nabla_{j(v)} \xi) - \frac{1}{2} J (\nabla_\xi J) \partial_J(u).$$

Here, ∇ is the metric connection on TW of the metric induced by ω and J , and $\xi \in \mathbb{R}^{2g} \oplus L_1^{p,d}(u^*TW, \partial) \subset \Gamma(u^*TW)$. The space $L_1^{p,d}(u^*TW, \partial)$ consists of sections of u^*TM tangent to $C_\alpha \cup C_\beta$ over ∂S with one derivative in $L^{p,d}$; see Section 3 for more details.

If we allow the complex structure on S to vary, then there is an additional term. The tangent space to Teichmüller space at S is a finite-dimensional subspace of $C^\infty(\text{End}(TS, j))$, where $\text{End}(TS, j)$ is the space of endomorphisms of TS anti-commuting with j . The linearized $\bar{\partial}$ -operator is then a restriction of the map

$$D_u \bar{\partial}: \mathbb{R}^{2g} \oplus L_1^{p,d}(u^*TW, \partial) \oplus C^\infty(\text{End}(TS, j)) \rightarrow L^{p,d}(\Lambda^{0,1}u^*TW)$$

defined by

$$D_u \bar{\partial}(\xi, Y)(v) = \frac{1}{2}(\nabla_v \xi + J \nabla_{j(v)} \xi) - \frac{1}{2}J(\nabla_\xi J) \partial_{J(v)}(u) + \frac{1}{2}J \circ du \circ Y(v).$$

For maps $\phi: \mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ the formulas are the same except that the \mathbb{R}^{2g} -factors are absent.

In the future, we will suppress the Lagrangians from the notation.

Fix $\epsilon > 0$. Let \bar{B}_{ours} denote the collection of maps u in B_{ours} for which

- $u_{\mathbb{D}}$ is holomorphic,
- $u_{\mathbb{D}}$ has only simple branch points p_1, \dots, p_ℓ ,
- $d(p_i, p_j) > 2\epsilon$ for $i \neq j$, and
- u_Σ has no branch points inside the $B_\epsilon(p_i)$, $i = 1, \dots, \ell$.

Let $B_{\text{ours}} \subset \bar{B}_{\text{ours}}$ denote the space of maps $u \in \bar{B}_{\text{ours}}$ for which u_Σ is holomorphic over $\cup_i u_{\mathbb{D}}^{-1}(B_\epsilon(p_i))$. For ϵ sufficiently small, J generic, and $\text{ind}(A) = 1$, $\mathcal{M}_{\text{ours}}^A \subset B_{\text{ours}}$. Note that B_{ours}^A is nonempty if and only if the intersection number $\Delta \cdot A$ of the diagonal in $\text{Sym}^g(\Sigma)$ with $A \in H_2(\text{Sym}^g(\Sigma), T_\alpha \cup T_\beta)$ is non-negative.

Let $u: S \rightarrow W$ be a map in \bar{B}_{ours} such that $u_{\mathbb{D}}$ has branch points $p_1, \dots, p_\ell \in (0, 1) \times \mathbb{R}$.

Lemma 13.8 *The tangent space to \bar{B}_{ours} at u is given by*

$$T_u \bar{B}_{\text{ours}} = L_k^{p,d}((\pi_\Sigma \circ u)^*T\Sigma, \partial) \oplus \left(\oplus_{i=1}^\ell \mathbb{C}\right).$$

Proof This is clear. A point in \bar{B}_{ours} is determined by a complex structure on S and a map from S to Σ . The complex structure is determined by the branch points p_1, \dots, p_ℓ . □

An inclusion of $T_u \overline{\mathcal{B}}_{\text{ours}}$ into $T_u \mathcal{B}_{\text{ours}}$ can be given as follows. The inclusion $L_1^{p,d}(u^* T \Sigma, \partial) \hookrightarrow L_1^{p,d}(u^* T W, \partial)$ is obvious. We include \mathbb{C}^ℓ into the space of infinitesimal deformations of the almost complex structure on the source. The \mathbb{C}^ℓ corresponds to moving the branch points p_1, \dots, p_g in \mathbb{D} . This, in turn, corresponds to deforming the almost complex structure $j_{\mathbb{D}}$ on \mathbb{D} . But any almost complex structure $j_{\mathbb{D}}$ on \mathbb{D} specifies an almost complex structure j_S on S via $j_S = (du_{\mathbb{D}})^{-1} \circ j_{\mathbb{D}} \circ du_{\mathbb{D}}$.

We choose a family of infinitesimal deformations parametrizing the \mathbb{C}^ℓ vanishing over the branch points p_1, \dots, p_ℓ , ie, such that the almost complex structure remains fixed near the branch points.

Lemma 13.9 $\ker(D\bar{\partial})|_{\overline{\mathcal{B}}_{\text{ours}}} = \ker(D(\bar{\partial}|_{\overline{\mathcal{B}}_{\text{ours}}}))$.

Proof Again, this is clear. \square

To discuss the linearized $\bar{\partial}$ -operator for maps to $\text{Sym}^g(\Sigma)$, we must first fix a connection on $T \text{Sym}^g(\Sigma)$. Away from the diagonal we choose the metric connection of the metric induced by the split symplectic form and our almost complex structure. We extend this connection arbitrarily over the diagonal. (By [24, Corollary 2], we could in fact extend the symplectic form over the diagonal, and work with the induced metric connection.) Since we will work with curves which are holomorphic near the diagonal, the choice of linearization near the diagonal is unimportant.

Let $\overline{\mathcal{B}}_{\text{theirs}}$ denote the collection of maps $\phi: \mathbb{D} \rightarrow \text{Sym}^g(\Sigma)$ (with boundary on the α - and β -tori) intersecting the diagonal Δ transversely and only in its top stratum. Notice that $\overline{\mathcal{B}}_{\text{theirs}}$ is an open subset of $\mathcal{B}_{\text{theirs}}$. Let $\mathcal{B}_{\text{theirs}} \subset \overline{\mathcal{B}}_{\text{theirs}}$ denote the subspace of maps which are holomorphic near the diagonal.

Lemma 13.10 *The tangent space at ϕ to $\overline{\mathcal{B}}_{\text{theirs}}$ is*

$$T_\phi \overline{\mathcal{B}}_{\text{theirs}} = L_1^{p,d}(\phi^* T \text{Sym}^g(\Sigma), \partial, \Delta) \oplus \mathbb{C}^\ell$$

where $L_1^{p,d}(\phi^* T \text{Sym}^g(\Sigma), \partial, \Delta)$ is the space of $L_1^{p,d}$ vector fields along ϕ which are tangent to the diagonal over $\phi^{-1}(\Delta)$ and to the Lagrangian tori T_α and T_β over $\partial\mathbb{D}$, and ℓ is the number of intersections of ϕ with Δ .

Proof The tangent space to the space of maps

$$(\mathbb{D}^2, \{p_1, \dots, p_\ell\}) \rightarrow (\text{Sym}^g(\Sigma), \Delta)$$

is $L_1^{p,d}(\phi^* T \text{Sym}^g(\Sigma), \partial, \Delta)$. The \mathbb{C}^ℓ corresponds to allowing the p_i to move. \square

Again, we can identify the \mathbb{C}^ℓ with the tangent space to an ℓ -dimensional family of deformations of $j_{\mathbb{D}}$. We take this to be the same ℓ -dimensional family used before. In particular, the family of almost complex structures is constant near the p_i .

There is a map $\Phi: (\overline{B}_{\text{ours}}, B_{\text{ours}}) \rightarrow (\overline{B}_{\text{theirs}}, B_{\text{theirs}})$ defined just as $\Phi: \mathcal{M}_{\text{ours}} \rightarrow \mathcal{M}_{\text{theirs}}$ was defined at the beginning of this section.

Now, suppose that $u \in B_{\text{ours}}$ and

$$(v, w) \in \ker(D_u \bar{\partial}|_{B_{\text{ours}}}) \subset L_1^{p,d}(u^* T\Sigma, \partial) \times \mathbb{C}^\ell$$

We construct from (v, w) an element of $\ker(D_{\Phi(u)} \bar{\partial}|_{B_{\text{theirs}}})$. First we construct a vector field $v' \in L_1^{p,d}(\Phi(u)^* T \text{Sym}^g(\Sigma), \partial)$. For $p \in \mathbb{D}$, $p \notin \cup_i B_\epsilon(p_i)$, let $u_{\mathbb{D}}^{-1}(p) = \{q_1, \dots, q_g\}$ and $\Phi(u)(p) = \{x_1, \dots, x_g\} = \{u_\Sigma(q_1), \dots, u_\Sigma(q_g)\}$. Then define $v'(p) = \{v(q_1), \dots, v(q_g)\} \in T_{\{x_1, \dots, x_g\}} \text{Sym}^g(\Sigma)$.

Since $\Phi(u)(p_i)$ lies in the top-dimensional stratum of the diagonal, near $\Phi(u)(p_i)$, $\text{Sym}^g(\Sigma)$ decomposes as $\text{Sym}^2(\Sigma) \times \Sigma^{g-2}$. The projection of v' to $T\Sigma^{g-2}$ is given by the previous construction. This leaves us to define the projection v'_1 of v' to $T \text{Sym}^2(\Sigma)$.

Inside the $B_\epsilon(p_i)$, w is constant so the term $J \circ du \circ Y$ in $D_u \bar{\partial}$ is zero. Let U_i be the component of $u_{\mathbb{D}}^{-1}(B_\epsilon(p_i))$ on which $du_{\mathbb{D}}$ is singular. Identify U_i and $V_i = u_\Sigma(U_i)$ holomorphically with \mathbb{D} . Then, $T V_i$ is identified with $V_i \times \mathbb{C}$ and $u_\Sigma^* T V_i$ with $U_i \times \mathbb{C}$. Further, v becomes a map $U_i \rightarrow \mathbb{C}$, and the statement that $D_u \bar{\partial} v = 0$ becomes the statement that $v: U_i \rightarrow \mathbb{C}$ is holomorphic.

Now, $\Phi(u)(B_\epsilon(p_i)) \subset \text{Sym}^2(V_i) \subset \text{Sym}^2(\mathbb{C}) \cong \mathbb{C}^2$. So, $\Phi(u)^* T \text{Sym}^2(\Sigma)$ is identified with $B_\epsilon(p_i) \times \mathbb{C}^2$. The holomorphic map $(u_{\mathbb{D}}, v): U_i \rightarrow B_\epsilon(p_i) \times \mathbb{C}^2$ specifies a holomorphic map $v'_1: B_\epsilon(p_i) \rightarrow \text{Sym}^2(\mathbb{C}) \cong \mathbb{C}^2$ by the same tautological correspondence used to define Φ . We view v'_1 as a section of $\Phi(u)^* T \text{Sym}^2(\Sigma)$. Observe that v'_1 and hence v_1 is tangent to the diagonal.

It is easy to check that the two definitions of v' agree over $\partial B_\epsilon(p_i)$.

Finally, by its definition, w corresponded to an infinitesimal deformation of $j_{\mathbb{D}}$.

Lemma 13.11 *The pair (v', w') lies in the kernel of $D \bar{\partial}_{\text{theirs}}$.*

Proof This is direct from the definitions. Away from the diagonal, this follows from the fact that the complex structure and symplectic form on $\text{Sym}^g(\Sigma)$ have the form $\text{Sym}^g(j_\Sigma)$ and $\text{Sym}^g(\omega_\Sigma)$ respectively. Near the diagonal, on the $\text{Sym}^2(\Sigma)$ factor the kernel of $D \bar{\partial}_{\text{theirs}}$ corresponds under the trivializations used above to holomorphic

maps $B_\epsilon(p_i) \rightarrow \text{Sym}^2(\mathbb{C})$. The map v' is holomorphic near the diagonal by the usual tautological correspondence. \square

This proves that $\ker(D\bar{\partial}_{\text{ours}}) \subset \ker(D\bar{\partial}_{\text{theirs}})$. The reverse inclusion can be proved similarly, using the opposite direction Φ^{-1} of the tautological correspondence.

This proves part (1) of Proposition 13.7. It remains to check part (2), ie, that the coherence conditions for the two theories agree. Before doing so, we prove another

Lemma 13.12 *Fix A with $\text{ind}(A) > 0$ and $A \cdot \Delta \geq 0$. Let σ_{ours}^A and σ_{theirs}^A be sections of $\mathcal{L}_{\text{ours}}$ and $\mathcal{L}_{\text{theirs}}$ respectively. Let $u_1, u_2 \in B_{\text{ours}}^A$ with $D_{u_i}\bar{\partial}$ surjective for $i = 1, 2$. At u_i we constructed an identification of $\mathcal{L}_{\text{ours}}$ and $\mathcal{L}_{\text{theirs}}$. Suppose that with respect to the identification of determinant liens at u_1 , $\sigma_{\text{ours}}^A(u_1)$ is a positive multiple of $\sigma_{\text{theirs}}^A(u_1)$. Then with respect to the identification of determinant lines at u_2 , $\sigma_{\text{ours}}^A(u_2)$ is a positive multiple of $\sigma_{\text{ours}}^A(u_2)$.*

Proof Choose a path u_a in B_{ours} from u_1 to u_2 . Choose also a family of subspaces H of $L^{p,d}(\Lambda^{0,1}u_a^*T\Sigma)$ such that the operator

$$\bar{D}_{u_a}\bar{\partial}: L^{p,d}(u_a^*T\Sigma, \partial) \oplus \mathbb{C}^\ell \oplus H \rightarrow L^{p,d}(\Lambda^{0,1}u_a^*T\Sigma)$$

given by

$$\bar{D}_{u_a}\bar{\partial}(v, w, x) = D_{u_a}\bar{\partial}(v, w) + x$$

is surjective for all a . Further, choose H so that all sections in H vanish near the branch points of $u_{\mathbb{D}}$. This is possible for a generic path u_a .

Choose an orientation $\sigma(H)$ of H , ie, a section of $\Lambda^{top}H$. Then there is a canonical isomorphism between the determinant lines of $D\bar{\partial}$ and $\bar{D}\bar{\partial}$.

Since the sections in H vanish near the branch points of $u_{\mathbb{D}}$, the space H specifies a subspace H' of $L^{p,d}(\Lambda^{0,1}\Phi(u_a)^*T\text{Sym}^g(\Sigma))$ as follows: for $x \in H$ and $y \in T_p\mathbb{D}$ let y_1, \dots, y_g be the preimages of y under $u_{\mathbb{D}}$. Then define

$$x(y) = \{x(y_1), \dots, x(y_g)\} \in T_{\Phi(u)(p)}\text{Sym}^g(\Sigma).$$

The vanishing of the sections in H near the branch points of $u_{\mathbb{D}}$ also means that we have an identification of $\ker(\bar{D}\bar{\partial}_{\text{ours}})$ and $\ker(\bar{D}\bar{\partial}_{\text{theirs}})$, defined in the same way as the identification of $\ker(\tilde{D}\bar{\partial}_{\text{ours}})$ and $\ker(\tilde{D}\bar{\partial}_{\text{theirs}})$ above. Further, the following diagram commutes:

$$\begin{CD} \text{Det}(D_{u_i}\bar{\partial}) @>\cong>> \text{Det}(D_{\Phi(u_i)}\bar{\partial}) \\ @V\cong\otimes\sigma(H)VV @VV\cong\otimes\sigma(H)V \\ \text{Det}(\bar{D}_{u_i}\bar{\partial}) @>\cong>> \text{Det}(\bar{D}_{\Phi(u_i)}\bar{\partial}) \end{CD}$$

for $i = 1, 2$. But this implies the result. □

Note that as a porism of Lemma 13.12 we also have an identification of determinant lines for A with $A \cdot \Delta \geq 0$ but $\text{ind}(A) \leq 0$.

Next we recall the definition of a coherent orientation. A coherent orientation consists of a choice of nonvanishing section $\sigma(A)$ of the determinant line bundle \mathcal{L} over the configuration space in the homotopy class A such that the following coherence condition is satisfied:

Let $u_1 \in \mathcal{B}^{A_1}(\vec{x}, \vec{y})$ and $u_2 \in \mathcal{B}^{A_2}(\vec{y}, \vec{z})$ be curves for which $D_{u_i} \bar{\partial}$ is surjective. In Appendix A we construct a family of preglued curves $u_1 \natural_r u_2 \in \mathcal{B}^{A_1+A_2}(\vec{x}, \vec{z})$. For r large, the kernel of $D_{u_1 \natural_r u_2} \bar{\partial}$ is identified with $\ker(D_{u_1} \bar{\partial}) \oplus \ker(D_{u_2} \bar{\partial})$. So, $\sigma^{A_1} \otimes \sigma^{A_2}$ specifies a section of $\mathcal{L}^{A_1+A_2}$. The *coherence condition* is that $\sigma^{A_1+A_2}$ be a positive multiple of $\sigma^{A_1} \otimes \sigma^{A_2}$. One must check that this condition is independent of u_1 and u_2 ; see for instance [3, Corollary 7]. (In the case that $D_{u_i} \bar{\partial}$ is not surjective, one stabilizes $D_{u_i} \bar{\partial}$ by a finite-dimensional oriented subspace of $L^{p,d}(\Lambda^{0,1} u_i^* T W)$, as in the proof of Lemma 13.12.)

(From now on, when we write a sum $A + B$ of homotopy classes of curves we implicitly assume that the asymptotics of A at ∞ agree with those of B at $-\infty$.)

Returning to our situation, fix a coherent orientation σ_{ours} . By the previous lemma, this specifies an orientation σ_{theirs}^A for each A with $\text{ind}(A) > 0$ and $A \cdot \Delta \geq 0$. We will refer to the collection of A with $\text{ind}(A) > 0$ and $A \cdot \Delta \geq 0$ as the *positive cone* and denote it C_+ . Note that $[\Sigma] \in C_+$ and for any A , $A + N[\Sigma] \in C_+$ for N sufficiently large.

Lemma 13.13 *The orientation σ_{theirs} is coherent over the positive cone. That is, for $A, B \in C_+$, with respect to the identification induced by gluing, $\sigma_{\text{theirs}}^A \otimes \sigma_{\text{theirs}}^B$ is a positive multiple of $\sigma_{\text{theirs}}^{A+B}$.*

Proof This follows from the commutativity of

$$\begin{array}{ccc}
 \mathcal{L}_{\text{ours}}^{A_1} \otimes \mathcal{L}_{\text{ours}}^{A_2} & \longrightarrow & \mathcal{L}_{\text{theirs}}^{A_1} \otimes \mathcal{L}_{\text{theirs}}^{A_2} \\
 \downarrow & & \downarrow \\
 \mathcal{L}_{\text{ours}}^{A_1+A_2} & \longrightarrow & \mathcal{L}_{\text{theirs}}^{A_1+A_2}
 \end{array}$$

up to positive scaling, which follows easily from the definitions of the various maps involved. □

Lemma 13.14 *Any coherent orientation σ over the positive cone can be extended uniquely to a coherent orientation over all configuration spaces.*

Proof Fix A , and N large enough that $A + N[\Sigma] \in C_+$. Then the map

$$\mathcal{L}^A \xrightarrow{\otimes_{\sigma}^{[N\Sigma]}} \mathcal{L}^{A+N[\Sigma]}$$

and the orientation $\sigma^{A+N[\Sigma]}$ specifies an orientation σ^A of \mathcal{L}^A . By the coherence of σ over the positive cone, the orientation σ^A is independent of N , and obviously agrees with σ over the positive cone. Commutativity of the following diagram, up to positive scaling, implies that σ is coherent:

$$\begin{array}{ccc} \mathcal{L}^A \otimes \mathcal{L}^B & \xrightarrow{\otimes_{\sigma}^{[N_1\Sigma]} \otimes_{\sigma}^{[N_2\Sigma]}} & \mathcal{L}^{A+N_1[\Sigma]} \otimes \mathcal{L}^{B+N_2[\Sigma]} \\ \downarrow & & \downarrow \\ \mathcal{L}^{A+B} & \xrightarrow{\otimes_{\sigma}^{[N_1\Sigma]} \otimes_{\sigma}^{[N_2\Sigma]}} & \mathcal{L}^{A+B+(N_1+N_2)[\Sigma]} \end{array}$$

(Here, the vertical maps are induced by gluing. This requires stabilizing the spaces involved, as discussed above.) \square

We have now proved part (2) of Proposition 13.7. This concludes the proof of Theorem 2.

14 Other Remarks

14.1 Elaborations of Heegaard Floer

The 4–dimensional approach that we have used suggests possible elaborations of the Floer homology groups considered in this paper. For the first one, we will describe briefly the \widehat{HF} –case, the only case that I am confident works. Our new chain complex \widehat{CF}_{big} is freely generated over $\mathbb{Z}[[t]]$ by the intersection points. The differential is $\partial_{big} = \partial_0 + t\partial_1 + t^2\partial_2 + \dots$. Here, ∂_0 is boundary operator for \widehat{HF} that appears throughout this paper. The linear term ∂_1 counts holomorphic curves with exactly 1 double point, in homology classes A with $\text{ind}(A) = 3$. (This is the appropriate condition for the resulting moduli space to be zero–dimensional.) In general, ∂_i counts holomorphic curves with singularity equivalent to i double points, in homology classes A with $\text{ind}(A) = 2i + 1$. (This idea is inspired by the so–called Taubes series described by Ionel–Parker in [11]. The analog for disks in $\text{Sym}^g(\Sigma)$ is to consider only disks with a prescribed number of tangencies to the diagonal.)




Figure 17: A handleslide in genus two, with notation

The resulting filtered chain complex, up to filtered chain homotopy, is an invariant of the three–manifold. The proof is similar to the proof of invariance of \widehat{HF} above. Undesirable codimension–one degenerations are ruled out just as in Proposition 7.1. The fact that $\partial_{big}^2 = 0$ now expresses the additional fact that when a holomorphic curve with k double points splits, i of them go to one level and $k - i$ to the other.

Isotopy invariance follows from the arguments of Section 9. Like ∂_{big} , the chain maps Φ constructed there now take the form $\Phi_{big} = \Phi_0 + t\Phi_1 + t^2\Phi_2 + \dots$, where Φ_i counts holomorphic curves with i double points or equivalent. Again, the statement that Φ_{big} is a chain map expresses the observation that when a holomorphic curve with double points splits, some double points go to one level and the rest to the other level. Otherwise, the argument is unchanged. Similarly, the triangle maps $F_{\alpha,\beta,\gamma}$ now have the form $\sum_{i \geq 0} F_{\alpha,\beta,\gamma,i}$ where $F_{\alpha,\beta,\gamma,i}$ counts curves with i double points.

Handleslide invariance is proved as in Section 11, except that some of the model calculations are more complicated. We give those computations here. Since we are considering \widehat{HF} , it suffices to study the case $g = 2$. Let $\vec{\beta}$, $\vec{\beta}'$ and $\vec{\beta}^H$ be the curves shown in Figure 17. (This notation agrees with the notation of Section 11.) The proofs that for an appropriate choice of orientation system the generators $\vec{\theta}_{\beta,\beta'} = \{\theta_1, \theta_2\}$,

$\vec{\theta}_{\beta^H, \beta'}$ = $\{\theta'_1, \theta'_2\}$ and $\vec{\theta}_{\beta, \beta^H}$ = $\{\theta_1^H, \theta_2^H\}$ shown in Figure 17 are cycles are the same as in Section 11, with the additional observation that immersed curves can only contribute to ∂_{big} if the grading difference is at least three.

However, the proof that $F_{\beta, \beta', \beta^H}(\vec{\theta}_{\beta, \beta^H} \otimes \vec{\theta}_{\beta^H, \beta'}) = \vec{\theta}_{\beta, \beta'}$ (Proposition 11.4) is somewhat more involved. Since triangle maps count curves in even-dimensional moduli spaces, it is conceivable that the coefficient of $t\{\eta_1, \eta_2\}$ is nonzero in $F_{\beta, \beta', \beta^H}(\vec{\theta}_{\beta, \beta^H} \otimes \vec{\theta}_{\beta^H, \beta'})$. Again, with notation as in Figure 17, we need to show that none of the following domains correspond to moduli spaces containing immersed curves: $D_1 + D_2 + D_3 + D_7 + D_8 + D_9$, $D_1 + D_2 + D_3 + D_9 + D_{10} + D_{11}$, $D_1 + D_2 + D_3 + D_7 + D_{11} + D_{12}$, $D_3 + D_4 + D_5 + D_7 + D_8 + D_9$, $D_3 + D_4 + D_5 + D_9 + D_{10} + D_{11}$, $D_3 + D_4 + D_5 + D_7 + D_{11} + D_{12}$, $D_1 + D_5 + D_6 + D_7 + D_8 + D_9 + D_{10} + D_{11} + D_{12}$, $D_1 + D_5 + D_6 + D_7 + D_{10} + 2D_{11} + 2D_{12}$, $D_1 + D_5 + D_6 + D_7 + 2D_8 + 2D_9 + D_{10}$, $D_1 + D_5 + D_6 + D_9 + 2D_{10} + 2D_{11} + D_{12}$, $D_1 + D_5 + D_6 + D_8 + 2D_9 + 2D_{10} + D_{11}$, $D_1 + D_5 + D_6 + 2D_7 + 2D_8 + D_9 + D_{12}$, and $D_1 + D_5 + D_6 + 2D_7 + D_8 + D_{11} + 2D_{12}$. (These are the only positive domains in $\pi_2(\vec{\theta}_{\beta, \beta^H}, \vec{\theta}_{\beta^H, \beta'}, \vec{\eta})$, where $\vec{\eta} = \{\eta_1, \eta_2\}$.)

In the first seven of these domains, there are no coefficients larger than one. It follows that there can be no immersed curve in the corresponding moduli space. The last six of these domains exhibit enough symmetry that it suffices to consider one of them. We will focus on $D = D_1 + D_5 + D_6 + D_7 + D_{10} + 2D_{11} + 2D_{12}$. It is shown shaded in Figure 18.

We compute the Euler characteristic of an embedded curve representing D . Observe that $D = (D_5 + D_{11} + D_{12} + D_7) + (D_1 + D_6 + D_{10} + D_{11} + D_{12}) = E + F$. Here, $E = (D_5 + D_{11} + D_{12} + D_7) \in \pi_2(\vec{\theta}^H, \{\eta'_1, \theta'_2\}, \vec{\eta})$ and $F = (D_1 + D_6 + D_{10} + D_{11} + D_{12}) \in \pi_2(\{\theta'_2, \eta'_1\}, \vec{\theta}')$. The domain E is represented by a pair of embedded disks. At least for some (nonempty, open set of) almost complex structures, the domain F is represented by an embedded annulus. Gluing representatives for E and F , it follows that the Euler characteristic for an embedded representative of D is 0.

It follows that the source of a curve with double points representing D must have $\chi \geq 2$.

It is clear from considering the boundary of D that any representative for D satisfying **(M0)**–**(M5)** must be connected. But there are no connected surfaces with boundary with $\chi \geq 2$. This proves the non-embedded moduli spaces are empty, so

$$F_{\beta, \beta', \beta^H}(\vec{\theta}_{\beta, \beta^H} \otimes \vec{\theta}_{\beta^H, \beta'}) = \vec{\theta}_{\beta, \beta'}$$

as desired.




Figure 18: The domain D

Plugging this computation into the proof of handleslide invariance in Section 11 proves handleslide invariance for \widehat{HF}_{big} . Stabilization invariance of \widehat{HF}_{big} follows in the same simple way as stabilization invariance of \widehat{HF} . Note that all of the maps we have used are maps of $\mathbb{Z}[[t]]$ -modules. It follows that the chain complex \widehat{CF}_{big} , up to chain homotopy equivalence over $\mathbb{Z}[[t]]$, is an invariant of Y .

Unfortunately, I have been unable to compute a single case in which \widehat{CF}_{big} is not homotopy equivalent to a complex in which all higher differentials vanish.

I suspect that one could similarly elaborate HF^∞ and HF^\pm , but have not done the computations necessary to establish handleslide invariance, and am mildly concerned that there may be subtleties in the proof of stabilization invariance. If evidence appears to suggest that these variants of HF^∞ , HF^\pm or, for that matter, \widehat{HF} would be new or interesting then they will become the subject of a future paper.

Here are several other elaborations, also inspired by Gromov–Witten theory. For convenience, I will formulate only the \widehat{HF} -analogs, but in these cases analogs of the

other theories present few added difficulties. (One does have to deal with annoying curves, but these can be addressed similarly to the way we did in Section 8.) To start, fix a homology class $[K] \in H_1(Y)$, a knot $K \hookrightarrow \Sigma$ representing $[K]$, and a point $s_0 \in [0, 1]$. Let $\widehat{CF}_{[K]}$ be freely generated over $\mathbb{Z}[t]$ by the intersection points. Define $\partial_{[K]} = \partial_0 + \frac{1}{1!}t\partial_1 + \frac{1}{2!}t^2\partial_2 + \cdots$, where the coefficient of \vec{y} in $\partial_i \vec{x}$ counts holomorphic curves with i marked points in homology classes A with $ind(A) = 1$ so that each marked point is mapped to $K \times \{s_0\} \times \mathbb{R}$. Then the standard proof shows that $\partial_{[K]}^2 = 0$.

The resulting chain complexes, up to chain homotopy equivalence over $\mathbb{Z}[t]$, are indeed invariants of $(Y, [K])$. However, as pointed out to me by M. Hutchings, the resulting chain complex can be reconstructed from the chain complex with “totally twisted” (group ring) coefficients. (This is not at all surprising; the number of times a holomorphic curve intersects $K \times \{s_0\} \times \mathbb{R}$ depends only on the homology class of the curve. Since this construction imitates pulling back a 2–dimensional cohomology class from W , this is somewhat analogous to the divisor equation in Gromov–Witten theory.)

One could try an analogous construction by forcing points to be mapped to $\{p_0\} \times \{s_0\} \times \mathbb{R}$ for some choice of $p_0 \in \Sigma$. The result is again an invariant, and independent of p_0 . However, taking $p_0 \in \Sigma \setminus (\alpha \cup \beta)$ and considering $s_0 = 0$, one sees that all of the higher differentials vanish.

This exhausts the obvious cohomology classes one could pull back, so the next idea is to try descendent classes. For example, let \widehat{CF}_{desc} be freely generated over $\mathbb{Z}[[t]]$ by the intersection points. Define $\partial_{desc} = \partial_0 + \frac{1}{1!}t\partial_1 + \frac{1}{2!}t^2\partial_2 + \cdots$ where ∂_i counts holomorphic curves u from Riemann surfaces S with i marked points p_1, \dots, p_i such that

- The homology class of u has index $2i + 1$.
- For each i , $u(p_i) \in K \times \{s_0\} \times \mathbb{R}$.
- For each i , $(\pi_\Sigma \circ u)'(p_i) = 0$.

Again, the proof that $\partial_{desc}^2 = 0$ is standard. That the chain complex up to chain homotopy equivalence (over $\mathbb{Z}[[t]]$) is an invariant of $(Y, [K])$ is almost identical to the proof of invariance of the $H_1/Tors$ –action on Heegaard Floer homology. There are no new computations that need to be done.

Of course, this is only one piece of a much bigger chain complex one could consider, where one keeps track of higher branching of $\pi_\Sigma \circ u$ and several different elements of $H_1(Y)$. One could also allow curves with branching of $\pi_\Sigma \circ u$ at prescribed s_0 and arbitrary point in Σ . Again, invariance of the bigger complex is free.

(One could also consider curves with prescribed branching of $\pi_{\mathbb{D}} \circ u$, but taking $s_0 = 0$ our proof that bubbling is impossible forces higher differentials of this kind to be trivial.)

All of these deformed complexes should have roughly the same formal properties as Heegaard Floer homology. Unfortunately, even with this apparent wealth of additional information I have been unable to find nontrivial examples. That is, while there are examples where higher differentials are nontrivial, I do not know examples which are not chain homotopy equivalent to complexes in which all higher differentials vanish. Hopefully this is for lack of creativity or perseverance on my part.

One might hope to construct further elaborations by pairing with cohomology classes of the space of maps $S \rightarrow W$, but by Proposition 6.1 this space does not have interesting topology. Finally, one might be able to obtain invariants by pulling back cohomology classes from the space of holomorphic maps to a disk. Doing so in a useful way, however, would require a better understanding of the cohomology of the moduli space of maps to a disk than I presently possess.

14.2 Relationship with Taubes' program

We conclude with a few remarks about a likely relationship between Heegaard Floer homology as formulated in this paper and Taubes' program to understand holomorphic curves in 4-manifolds with singular symplectic forms.

First, a one paragraph sketch of Taubes' idea. Any 4-manifold M with $b_2^+ > 0$ can be endowed with a closed two-form ω which is nondegenerate in the complement of a collection of circles, and degenerates in a controlled way near the circles; see [26] for further information and references. The program is then to fix a complex structure J in the complement of the singular circles, adjusted to ω , and obtain smooth invariants by studying J -holomorphic curves of finite ω -energy in M .

The Floer homology associated to Taubes' program would be structured as follows. For a three-manifold Y , one chooses a closed 2-form ω on $Y \times \mathbb{R}$ which is nondegenerate on the complement of certain lines $\{p_i\} \times \mathbb{R}$. One would then choose a translation invariant complex structure J on $Y \times \mathbb{R}$ adjusted to ω and study J -holomorphic curves with some specified asymptotics at the singular lines and at $Y \times \{\pm\infty\}$.

Fix a Morse function f on Y , and a metric on Y . Then, the form $df \wedge dt + \star df$ is a closed 2-form on $Y \times \mathbb{R}$, which is nondegenerate on the complement of $\{\text{critical points of } f\} \times \mathbb{R}$. (Here, \star means the Hodge star on Y , not on $Y \times \mathbb{R}$.) One possible compatible complex structure pairs dt with ∇f and is given by rotation by $\pi/2$ on $\ker(df)$. We will call this the complex structure induced by f .

But this setting bears a close resemblance to $(W, C_\alpha \cup C_\beta)$. Assume f was self-indexing. Then, view $\Sigma \times [0, 1]$ as the slice $f^{-1}(3/2 - \epsilon, 3/2 + \epsilon)$ of Y . Then the complex structure induced by f satisfies **(J1)**–**(J5)**. Further, in the limit $\epsilon \rightarrow 1/2$, the boundary conditions we impose become certain asymptotic conditions on the holomorphic curves at the singular lines. So, our setup fits quite nicely in Taubes' picture.

A serious difficulty in studying holomorphic curves in $Y \times \mathbb{R}$ is understanding the asymptotics at the singular lines. By studying only a middle slice of Y , Heegaard Floer homology neatly avoids this issue. Still, it would be nice to be able to work in $Y \times \mathbb{R}$; for instance, this would probably illuminate the proof of handleslide invariance.

Appendix A Gluing lemmas

A.1 Statement of results

As seems conventional in the subject, we relegate the misery called “gluing lemmas” to the appendix.

In the following, by “symplectic manifold” we mean a symplectic manifold with cylindrical ends, and by “complex structure” we mean a symmetric almost complex structure, as defined in [2, Section 2]. We will always assume that Reeb orbits and chords corresponding to the almost complex structure and Lagrangian submanifolds in question are isolated, as the Morse–Bott case requires extra work, and all Reeb chords in this paper are isolated. When we refer to a holomorphic curve we always mean ones with finite energy in the sense of [2, Section 6.1]. With these conventions, the gluing results used in this paper are

Proposition A.1 *Let (M_1, M_2) be a chain of symplectic manifolds (cf [6, Section 1.6]) with M_2 cylindrical. Let (u_1, u_2) be a height two holomorphic building in (M_1, M_2) . Assume that the complex structures on M_1 and M_2 are chosen so that the $\bar{\partial}$ operator is transverse to the zero-section at u_1 and u_2 . Then there is a neighborhood of (u_1, u_2) in the space of (height one or two) holomorphic buildings diffeomorphic to $\mathbb{R}^{\text{ind}(u_1)+\text{ind}(u_2)-1} \times (0, 1]$. If M_1 is also cylindrical then the same statement holds with $\mathbb{R}^{\text{ind}(u_1)+\text{ind}(u_2)-1}$ replaced by $\mathbb{R}^{\text{ind}(u_1)+\text{ind}(u_2)-2}$. Both statements remain true in the relative case.*

Proposition A.2 *Let (M_1, M_2) be a chain of symplectic manifolds (cf [6, Section 1.6]) obtained by splitting a symplectic manifold M along a co-oriented hypersurface H as in [2, Section 3.4]. Let (u_1, u_2) be a height two holomorphic building in*

(M_1, M_2) . Assume that the complex structure on M is chosen so that the $\bar{\partial}$ operator is transverse to the zero-section at u_1 and u_2 . Let J_R denote the complex structure on M obtained by inserting a neck of length R along H . Let \mathcal{M}_R denote the space of J_R -holomorphic curves in M . Then there is a neighborhood of (u_1, u_2) in the space $\bigcup_R \mathcal{M}_R$ diffeomorphic to $\mathbb{R}^{\text{ind}(u_1)+\text{ind}(u_2)} \times (0, 1]$. The statement remains true in the relative case.

Proposition A.3 *Proposition A.2 remains true in the case of height one holomorphic twin towers studied in Section 12. That is, with notation as in Section 12, for $\text{ind}(A) = 1$, we have $\#\widehat{\mathcal{M}}_{W'_R}^A = \#\widehat{\mathcal{M}}_W^A$.*

The special case of Proposition A.1 when M_1 is the symplectization of a contact manifold follows from work of Bourgeois [1, Corollary 5.7]. His argument, however, extends without essential changes to prove the more general results stated here. To demonstrate this, we imitate his argument to prove Proposition A.2 in the relative case. Proposition A.1 is similar but marginally less complicated. After proving Proposition A.2 we discuss how the proof needs to be modified to prove Proposition A.3.

A.2 The splitting process

In this subsection we describe the process of splitting a symplectic manifold along a hypersurface. More details can be found in [2, Section 3.4]. We describe the splitting process (tersely) here partly because the relative case is hardly discussed in [2] but mostly just to fix notation.

Let (M, ω) be a symplectic manifold, J a complex structure on M compatible with ω , and $H \subset M$ a co-oriented compact hypersurface. We also assume that we are given a vector field $\vec{R} \in \ker(\omega|_H)$ on H so that the associated cylindrical structure on $H \times \mathbb{R}$ is symmetric (see [2, page 802]). Let α be the 1-form $\omega(J\vec{R}, \cdot)$ on H . Then a neighborhood of H is symplectomorphic to $((-\epsilon, \epsilon) \times H, \pi_H^* \omega|_H + d(t\alpha))$ so that $J\vec{R}$ is identified with $\frac{\partial}{\partial t}|_{H=0}$ (cf [2, page 806]). We will assume that J is preserved by the flow in the t -direction. This can always be arranged by perturbing J near H . In this case, let M° denote the manifold $M \setminus (-\epsilon/2, \epsilon/2) \times H$ with $\partial M^\circ = H' \cup H''$. Form

$$M_R = M^\circ \bigcup_{\substack{H' = \{-R\} \times H \\ H'' = \{R\} \times H}} [-R, R] \times H.$$

The complex structure on M_R is given by J on M° and on $[-R, R] \times H$ agrees with J on $H = \{0\} \times H$ and is translation invariant. The symplectic form on M_R is given by $\frac{2R}{\epsilon} \omega$ on M° and by $\pi_H^* \omega|_H + d(t\alpha)$ on $[-R, R] \times H$.

Similarly, let

$$M_\infty = (-\infty, 0] \times H \bigcup_{\{0\} \times H = H'} M^\circ \bigcup_{H'' = \{0\} \times H} [0, \infty) \times H.$$

As before, the complex structures on the ends $(-\infty, 0] \times H$ and $[0, \infty) \times H$ are induced by translation invariance and $J|_H$; on M° the complex structure is given by J . On M° , the symplectic form is just ω . On $(-\infty, 0] \times H$ and $[0, \infty) \times H$ it is given by $\pi_H^* \omega|_H + d(t\alpha)$.

Suppose $L \subset M$ is a Lagrangian submanifold intersecting H transversally. Perturbing L slightly we can assume that in the tubular neighborhood $(-\epsilon, \epsilon) \times H$ used to perform the stretching, L has the form $(-\epsilon, \epsilon) \times (L \cap H)$. (In fact, this can be achieved by a Hamiltonian deformation of L .) We will call this property being “cylindrical” near H . Then, the splitting procedure just described gives in an obvious way Lagrangian submanifolds L_R in M_R and L_∞ in M_∞ .

Now, fix

- a symplectic manifold (M, ω) with cylindrical ends
- a symmetric, cylindrical almost complex structure J on M adjusted to ω
- a compact, co-oriented hypersurface $H \subset M$ and
- a closed Lagrangian submanifold $L \subset M$ which is cylindrical near H , such that in each end $[0, \infty) \times \tilde{E}$ (respectively $(-\infty, 0] \times \tilde{E}$) of M , L has the form $[0, \infty) \times \tilde{L}_E$ (respectively $(-\infty, 0] \times \tilde{L}_E$) with $T\tilde{L}_E \subset T\tilde{E} \cap JT\tilde{E}$.

Let $\tilde{E}_1, \dots, \tilde{E}_N$ denote the ends of M , so near each \tilde{E}_i , M is modeled on $I_i \times \tilde{E}_i =: E_i$, where I_i is a semi-infinite interval. The isomorphism $E_i = \tilde{E}_i \times I_i$ is considered fixed as part of the definition of a cylindrical complex structure. If I_i is $[r, \infty)$ we call the end E_i positive; otherwise we call E_i negative. Let σ_i be $+1$ if E_i is positive and -1 if E_i is negative.

Let M_R and M_∞ denote the manifolds obtained by splitting M along H . The ends of M_R are $\tilde{E}_1, \dots, \tilde{E}_N$. The ends of M_∞ are $\tilde{E}_0^+, \tilde{E}_0^-, \tilde{E}_1, \dots, \tilde{E}_N$, where \tilde{E}_0^\pm correspond to H . Note that $\tilde{E}_0^+ = \tilde{E}_0^-$, where this identification respects the image of L and also the Reeb field. We sometimes write \tilde{E}_0 for \tilde{E}_0^\pm .

For $i = 0, \dots, N$, let $L \cap E_i = \tilde{L}_i \times I_i$. Let τ_i denote the coordinate on I_i , and $\vec{R}_i = J\partial_{\tau_i}$. Let Γ_i^o denote the space of closed orbits of \vec{R}_i , Γ_i^c the space of \vec{R}_i -chords, ie, flows of \vec{R}_i starting and ending on \tilde{L}_i . Let $\Gamma_i = \Gamma_i^o \cup \Gamma_i^c$.

We will assume throughout that the Γ_i are discrete.

Fix Reeb chords/orbits $\gamma_{0,1}, \dots, \gamma_{0,n_0} \in \Gamma_0$, $\gamma_{i,1}^+, \dots, \gamma_{i,n_i}^+ \in \Gamma_i$ ($i = 1, \dots, N$), and $\gamma_{i,1}^-, \dots, \gamma_{i,n_i}^- \in \Gamma_i$ ($i = 1, \dots, N$). For notational convenience, let $\gamma_{0,j}^+ = \gamma_{0,j}^- = \gamma_{0,j}$. Assume $\gamma_{i,1}^\pm, \dots, \gamma_{i,m_i}^\pm \in \Gamma_i^o$ are closed Reeb orbits and $\gamma_{i,m_i^\pm+1}, \dots, \gamma_{0,n_i^\pm} \in \Gamma_i^c$ are Reeb chords. For convenience and because it is the only case relevant to this paper we will assume that all the $\gamma_{0,k}$ and $\gamma_{i,k}^\pm$ are simple Reeb chords / orbits.

Fix smooth surfaces Σ^\pm with boundary and labeled punctures

$$p_{0,1}^\pm, \dots, p_{0,n_0}^\pm, \dots, p_{N,1}^\pm, \dots, p_{N,n_N}^\pm.$$

Let \mathcal{T}^\pm be smooth families of conformal structures on Σ so that \mathcal{T}^\pm surjects onto an open set in the moduli space of Riemann surfaces. (That is, points in \mathcal{T}^\pm are honest surfaces with conformal structures and marked points, not equivalence classes of such. The map from \mathcal{T}^\pm to moduli space need not be injective.) Choose \mathcal{T} so that there are small neighborhoods of the punctures in Σ in which the conformal structure is constant. Fix $p > 2$, $k \geq 1$. For convenience later, we will also fix $S_0^\pm \in \mathcal{T}^\pm$.

For $S^\pm \in \mathcal{T}^\pm$, let $\mathcal{B}_{S^\pm}^\pm = \mathcal{B}_k^{p,d,\pm}$ denote the Banach manifold comprising W_k^p -maps $(S^\pm, \partial S^\pm) \rightarrow (M_\infty, L)$ converging to $\gamma_{i,k}^\pm$ at $p_{i,k}^\pm$ ($i = 0, \dots, N$, $k = 1, \dots, n_i$) in $W_k^{p,d}$. That is:

Any \vec{R} -chord or orbit γ has some period T_γ defined by $T_\gamma = \frac{1}{2\pi} \int_\gamma \omega(\partial_\tau, \cdot)$. If γ is a Reeb orbit choose a tubular neighborhood N_γ of γ invariant under the \vec{R} -flow and a diffeomorphism ϕ from a neighborhood of $S^1 \times \vec{0}$ in $S^1 \times \mathbb{R}^{2n-2}$ to N_γ such that

- (1) $\phi(S^1 \times \vec{0}) = \gamma$
- (2) the pushforward under ϕ of T_γ times the unit tangent vector to $S^1 \times \vec{0}$ is \vec{R} and
- (3) $\phi^* \omega|_{S^1 \times \vec{0}} = \omega_0|_{S^1 \times \vec{0}}$, where ω_0 denotes the standard symplectic form on \mathbb{R}^{2n-2} .

If γ is a Reeb chord choose a tubular neighborhood N_γ of γ invariant under the \vec{R} -flow and a diffeomorphism ϕ from a neighborhood of $[0, 1] \times \vec{0}$ in $[0, 1] \times \mathbb{R}^{2n-2}$ to N_γ such that

- (1) $\phi([0, 1] \times \vec{0}) = \gamma$
- (2) the pushforward under ϕ of T_γ times the unit tangent vector to $[0, 1] \times \vec{0}$ is \vec{R} and
- (3) $\phi^* \omega|_{[0,1] \times \vec{0}} = \omega_0|_{[0,1] \times \vec{0}}$, where ω_0 denotes the standard symplectic form on \mathbb{R}^{2n-2} .

(Compare [2, Lemma A.1], but note that our N_γ is not the same as their N .) Let θ_γ parametrize S^1 in the first case, or $[0, 1]$ in the second case, by arc length. Let $z = (x_1, \dots, x_{n-1}, y_1, \dots, y_{n-1})$ be the standard coordinates on \mathbb{R}^{2n-2} .

Via ϕ , θ_γ and z become coordinates on N_γ .

For $S^\pm \in \mathcal{T}^\pm$ we fix a holomorphic identification of a neighborhood of each $p_{i,j}^\pm$ with $[0, \infty) \times S^1$ if $p_{i,j}^\pm$ is in the interior of S^\pm and with $[0, \infty) \times [-\pi, \pi]$ if $p_{i,j}^\pm$ is in the boundary of S^\pm . Let (t, s) denote the coordinates near $p_{i,j}^\pm$ induced by this decomposition.

The statement that a map $u_\pm: S^\pm \rightarrow M_\infty$ is asymptotic to $\gamma_{i,j}^\pm$ at $p_{i,j}^\pm$ means that for some $\tau_{i,j,0}^\pm, \theta_{i,j,0}^\pm \in \mathbb{R}$, $z \circ u_\pm(t, s) \rightarrow 0$ as $t \rightarrow \infty$, $\tau \circ u_\pm(t, s) - \sigma_i T_{\gamma_{i,j}^\pm} t - \sigma_i \tau_{i,j,0}^\pm \rightarrow 0$ as $t \rightarrow \infty$ and $\theta \circ u_\pm(t, s) - s - \theta_{i,j,0}^\pm \rightarrow 0$ as $t \rightarrow \infty$. We say that the map u_\pm is $L_k^{p,d}$ near $p_{i,j}^\pm$ if $z \circ u(t, s)$, $\tau \circ u(t, s) - \sigma_i T_{\gamma_{i,j}^\pm} t - \sigma_i \tau_{i,j,0}^\pm$ and $\theta \circ u(t, s) - s - \theta_{i,j,0}^\pm$ are all in $L_k^{p,d} = \{f \mid f(t, s) e^{d|t|/p} \in L_k^p\}$. Note that if $\gamma_{i,j}$ is a Reeb chord then $\theta_{i,j,0}^\pm = 0$ while if $\gamma_{i,j}$ is a closed Reeb orbit then after replacing s with $s - \theta_{i,j,0}^\pm$ we can assume that $\theta_{i,j,0}^\pm = 0$. In the future we will assume that s has been chosen so that $\theta_{i,j,0}^\pm = 0$. Replacing some of the neighborhoods of the punctures with smaller ones if necessary we can assume that for each i the constants $\tau_{i,j,0}^\pm$ are all the same, $\tau_{i,j,0}^\pm = \tau_{i,0}$.

The compactness result [2, Theorem 10.2] and asymptotic convergence result [2, Proposition 6.2] imply that there exists a $d > 0$ such that any holomorphic curve of finite energy with the specified asymptotics is in $\mathcal{B}_k^{p,d}$. Fix such a d , and let $\mathcal{B}^\pm = \mathcal{B}_k^{p,d,\pm}$.

Similarly, for any Riemannian vector bundle E over S^\pm we can consider the space of $L_k^{p,d}$ -sections of E , $L_k^{p,d}(E)$. The complex structure J and symplectic form ω induce a metric on M_∞ , so it makes sense to talk about the space $L_{k-1}^{p,d}(\Lambda^{0,1} T^* S^\pm \otimes_J u^* T M_\infty)$ of $L_{k-1}^{p,d}$ $(0, 1)$ -forms on S^\pm with values in $T^* M$.

Now, the spaces $\mathcal{B}_{S^\pm}^\pm$ fit together into a fiber bundle \mathcal{B}^\pm over \mathcal{T}^\pm . In turn, the spaces $L_{k-1}^{p,d}(\Lambda^{0,1} T^* S^\pm \otimes_J u^* T M_\infty)$ fit together into a vector bundle \mathcal{E}^\pm over \mathcal{B}^\pm . The $\bar{\partial}$ -operator gives a section $\mathcal{B}^\pm \rightarrow \mathcal{E}^\pm$. We will assume that this $\bar{\partial}$ map is transversal to the 0-section. We let \mathcal{M}^\pm denote the intersection of $\bar{\partial}$ with the 0-section.

Recall that we fixed reference Riemann surfaces $S_0^\pm \in \mathcal{T}^\pm$. As all Riemann surfaces in \mathcal{T}^\pm have fixed smooth source Σ^\pm , there is an obvious projection map $\mathcal{B}^\pm \rightarrow \mathcal{B}_{S_0^\pm}^\pm$. Identify \mathcal{B}^\pm with $\mathcal{T} \times \mathcal{B}_{S_0^\pm}^\pm$.

Remark While we have not yet chosen metrics or norms on the spaces under consideration, when we do, a figment of the details of our definition is that the identification of \mathcal{B}^\pm with $\mathcal{T} \times \mathcal{B}_{S_0^\pm}^\pm$ will not induce isometries. However, by choosing \mathcal{T} small we can make the induced maps arbitrarily close to isometries.

Choosing a connection on M_∞ , for instance the Levi–Civita connection, we obtain also a linearized $\bar{\partial}$ –operator $D\bar{\partial}: T_u\mathcal{B}^\pm \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1}T^*S \otimes_J u^*TM) = \mathcal{E}_u$. (The space \mathcal{E}_u is identified with the vertical tangent space to \mathcal{E}_u . The map $D\bar{\partial}$ is the composition of projection onto the vertical tangent space – defined with the help of the connection – with the derivative of the section $\bar{\partial}$. See [14, Section 3.3].) We will also sometimes be interested in the restriction of the $\bar{\partial}$ –operator to maps with fixed source; we denote the $\bar{\partial}$ –operator on maps with source S by $\bar{\partial}_S$ and its linearization by $D\bar{\partial}_S$.

Let $L_k^{p,d}(u_\pm^*TM_\infty, \partial)$ denote those $L_k^{p,d}$ sections of $u_\pm^*TM_\infty$ which lie in TL over ∂S .

The tangent space at (S_0^\pm, j) to \mathcal{T}^\pm is a finite–dimensional space V^\pm . The tangent space to $\mathcal{B}_{S^\pm}^\pm$ at some map u_\pm is

$$\mathbb{C}^{\sum_{i=0}^N m_i^\pm} \oplus \mathbb{R}^{\sum_{i=0}^N (n_i^\pm - m_i^\pm)} \oplus L_k^{p,d}(u_\pm^*TM_\infty, \partial) = F^\pm \oplus L_k^{p,d}(u_\pm^*TM_\infty, \partial).$$

Fix a vector field $v_{\tau,i,j}^\pm$ (respectively $v_{\theta,i,j}^\pm$) which agrees with $\frac{\partial}{\partial \tau}$ (respectively $\frac{\partial}{\partial \theta}$) near $p_{i,j}^\pm$ and lies in $L_k^{p,d}(u_\pm^*TM_\infty, \partial)$ away from $p_{i,j}^\pm$ for each interior puncture $p_{i,j}^\pm$ of S_0^\pm . Fix a vector field $v_{\tau,i,j}^\pm$ which agrees with $\frac{\partial}{\partial \tau}$ near $p_{i,j}^\pm$ and lies in $L_k^{p,d}(u_\pm^*TM_\infty, \partial)$ away from $p_{i,j}^\pm$ for each boundary puncture $p_{i,j}^\pm$ of S_0^\pm . Then, F^\pm is $\text{Span}\{v_{\tau,i,j}^\pm, v_{\theta,i,j}^\pm\} \subset \Gamma u^*TM_\infty$. (Equivalently, F^\pm corresponds to the constants $\tau_{i,j,0}^\pm$ and $\theta_{i,j,0}^\pm$ varying.) Fix a norm on F^\pm and use this norm to extend the $L_k^{p,d}$ –norm to $F^\pm \oplus L_k^{p,d}(u^*TM_\infty, \partial)$.

From the decomposition $\mathcal{B}^\pm = \mathcal{T} \times \mathcal{B}_{S_0^\pm}^\pm$ we obtain an identification of $T_{u_\pm}\mathcal{B}^\pm$ with $V^\pm \oplus F^\pm \oplus L_k^{p,d}(u_\pm^*TM_\infty, \partial)$. Let $W_{u_\pm}^\pm = D\bar{\partial}(V^\pm \oplus \{0\} \oplus \{0\}) \subset \Gamma(\Lambda^{0,1}T^*S^\pm \otimes_J u_\pm^*TM_\infty)$. (We use Γ to denote C_0^∞ sections.) Later, we will assume that the map $D\bar{\partial}$ is surjective at u_\pm . This is equivalent to the statement that $W_{u_\pm}^\pm + D\bar{\partial}_S(F^\pm \oplus L_k^{p,d}(u_\pm^*TM_\infty, \partial))$ is all of

$$L_{k-1}^{p,d}(\Lambda^{0,1}T^*S^\pm \otimes_J u_\pm^*TM_\infty) = L_{k-1}^{p,d}(u_\pm^*TM_\infty).$$

(Compare [1, pages 55–56].)

As the particulars of our asymptotics are irrelevant to linear statements, the following lemma is completely standard; see, for instance, [7, Theorem 4, page 797].

Lemma A.4 *For appropriate (small) $d > 0$, the linearized $\bar{\partial}$ -operator $D\bar{\partial}$ is Fredholm.*

We will use a gluing construction to define spaces \mathcal{T}_R of glued surfaces; then \mathcal{B}^R , \mathcal{E}^R , and \mathcal{M}^R are defined similarly to \mathcal{B}^\pm , \mathcal{E}^\pm , and \mathcal{M}^\pm . As described later, it will be important that we modify the metric on \mathcal{B}^R and norm on \mathcal{E}^R by adding exponential weights in the necks, but the topologies on the spaces will not be affected. We define \mathcal{B} (respectively \mathcal{E} , \mathcal{M}) to be the union over $R \in [0, \infty]$ of \mathcal{M}^R (respectively \mathcal{E}^R , \mathcal{M}^R).

A.3 Notation

In summary, so far we have a symplectic manifold M , hypersurface $H \hookrightarrow M$, and Lagrangian submanifold $L \hookrightarrow M$. Splitting M along H we obtain manifolds M_R and M_∞ , each with Lagrangian submanifolds obtained from L , which we still denote L . The split manifold M_∞ has two new ends E_0^\pm corresponding to H .

We have families \mathcal{T}^\pm of complex curves and spaces \mathcal{B}^\pm of $W_k^{p,d}$ maps of curves in \mathcal{T}^\pm into M_∞ , asymptotic to Reeb orbits $\gamma_{0,1} = \gamma_{0,1}^+ = \gamma_{0,1}^- \cdots \gamma_{0,m_0} = \gamma_{0,m_0}^+ = \gamma_{0,m_0}^-$ and Reeb chords $\gamma_{0,m_0+1} = \gamma_{0,m_0+1}^+ = \gamma_{0,m_0+1}^- \cdots \gamma_{0,n_0} = \gamma_{0,n_0}^+ = \gamma_{0,n_0}^-$ in the ends E_0^\pm at punctures $p_{0,j}^\pm$ and various other fixed Reeb chords in the other ends. We have defined bundles \mathcal{E}^\pm over \mathcal{B}^\pm of which $\bar{\partial}$ is a section, and \mathcal{M}^\pm to be the intersection of $\bar{\partial}$ with the 0-section of \mathcal{E}^\pm . The linearized $\bar{\partial}$ -operator is denoted

$$D\bar{\partial} : W_{u_\pm}^\pm \oplus F^\pm \oplus L_k^{p,d}(u_\pm^* TM_\infty, \partial) \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1} u_\pm^* TM_\infty)$$

where $F^\pm \cong \mathbb{C}^{\sum_{i=0}^N m_i^\pm} \oplus \mathbb{R}^{\sum_{i=0}^N (n_i^\pm - m_i^\pm)}$ is generated by sections $v_{\tau,i,j}^\pm$ and $v_{\theta,i,j}^\pm$ which agree with $\partial/\partial\tau$ and $\partial/\partial\theta$ respectively near the puncture $p_{i,j}^\pm$. When restricting to a fixed source S^\pm those objects are replaced by $\mathcal{B}_{S^\pm}^\pm$, $\mathcal{E}_{S^\pm}^\pm$, $\mathcal{M}_{S^\pm}^\pm$, $\bar{\partial}_{S^\pm}$ and $D\bar{\partial}_{S^\pm}$ respectively.

To perform the gluing we will need cutoff functions. Fix a smooth function $\beta: \mathbb{R} \rightarrow [0, 1]$ such that

$$\begin{cases} \beta(t) = 0 & \text{if } t \leq 0 \\ \beta(t) = 1 & \text{if } t \geq 1 \\ 0 \leq \beta^{(\ell)}(t) \leq 2^\ell & \text{for all } t \in \mathbb{R} \text{ and } \ell \in \mathbb{N}. \end{cases}$$

For $R/2 > T_{\gamma_{0,i}^\pm}$, define the cutoff function $\beta_{i,R}(t) = \beta\left(\frac{T_{\gamma_{0,i}^\pm}(t-1)}{R/2 - T_{\gamma_{0,i}^\pm}}\right)$, so that

$$\begin{cases} \beta_{i,R}(t) = 0 & \text{for } t \leq 1 \\ \beta_{i,R}(t) = 1 & \text{for } t \geq \frac{R}{2T} \\ 0 \leq \beta_{i,R}^{(\ell)}(t) \leq \left(\frac{2T}{R/2 - T}\right)^\ell & \text{for all } t \in \mathbb{R} \text{ and } \ell \in \mathbb{N}. \end{cases}$$

Remark For our proofs, the explicit bounds on the higher derivatives will not be important. All that will matter is that all derivatives of $\beta_{i,R}$ are uniformly bounded in R , and that $\beta'_{i,R} \rightarrow 0$ as $R \rightarrow \infty$.

A.4 Gluing estimates

Choose elements $u^\pm: S^\pm \rightarrow M_\infty$ in \mathcal{M}^\pm . Let ind^\pm denote the index of $D\bar{\partial}$ at u^\pm . Since we are assuming transversality, there is a neighborhood of u^\pm in \mathcal{M}^\pm diffeomorphic to \mathbb{R}^{ind^\pm} . We want to show that there is a neighborhood of the two-story holomorphic building (u^-, u^+) in \mathcal{M} diffeomorphic to $\mathbb{R}^{ind^-} \times \mathbb{R}^{ind^+} \times (0, 1]$. To do this, we will use the

Lemma A.5 (Implicit Function Theorem) Let $f: E \rightarrow F$ be a smooth map of Banach spaces with a Taylor expansion

$$f(\xi) = f(0) + Df(0)\xi + N(\xi).$$

Assume $Df(0)$ has a finite-dimensional kernel and a right inverse Q satisfying

$$\|QN(\xi) - QN(\eta)\| \leq C(\|\xi\| + \|\eta\|)\|\xi - \eta\|,$$

for some constant C . Assume also that $\|Qf(0)\| \leq \frac{1}{8C}$. Then for $\delta = 1/(4C)$, $f^{-1}(0) \cap B_\delta(\xi)$ is a smooth manifold of dimension $\dim \ker Df(0)$. In fact, there is a smooth function $\pi: \ker Df(0) \cap B_\delta(\xi) \rightarrow Q(F)$ such that $f(\xi + \phi(\xi)) = 0$ and all zeroes of f in $B_\delta(\xi)$ have the form $\xi + \phi(\xi)$.

This result is [8, Proposition 24, page 25]. The proof is essentially the same as the finite-dimensional case. The result, which A. Floer refers to as Newton’s method, is called “Floer’s Picard Lemma” by some authors. In the literature, I have found at least five incorrect references to its location.

We will apply the implicit function theorem to $\bar{\partial}: \mathcal{B} \rightarrow \mathcal{E}$. Recall that for $u: S \rightarrow M$, $T_u\mathcal{B} = V \oplus F \oplus L_k^{p,d}(u^*TM, \partial)$ for some finite-dimensional space $V \oplus F$. Choose

a norm on $V \oplus F$ arbitrarily; together with the $L_k^{p,d}$ -norm this makes $T\mathcal{B}$ into a Banach space.

We give two necessary general results before turning to the estimates needed by the implicit function theorem.

Lemma A.6 *At $(u: S \rightarrow M) \in W_k^{p,d}$ the remainder N in the Taylor expansion*

$$\bar{\partial}(u + \Delta u) = \bar{\partial}(u) + D\bar{\partial}(u)\Delta u + N(\Delta u) \in L_{k-1}^{p,d}(\Lambda^{0,1}T^*S \otimes_J u^*TM_\infty)$$

for the $\bar{\partial}$ -operator satisfies

$$(11) \quad \|N(\xi) - N(\eta)\| \leq C' (\|\xi\| + \|\eta\|) \|\xi - \eta\|$$

for ξ, η in $V \oplus F \oplus L_k^{p,d}(u^*TM, \partial)$ and some constant C (depending on $\|u\|_{W_{k-1}^{p,d}}$). (Here, we identify a neighborhood of the map u with a neighborhood of the 0-section in $V \oplus F \oplus L_k^{p,d}(u^*TM, \partial)$.)

Proof The argument is essentially the same as the one used by Floer to prove [7, Theorem 3a], and we omit it. Floer finds a relatively explicit formula [7, Lemma 3.2] for $N(\xi)$, using charts, and simple bounds on the terms in the formula. There are several differences between our setup and the setup of [7]. We claim the bound (11) for all $L_k^{p,d}$, while Floer only states it for $L_1^{p,d}$; however, Floer in fact proves the result for all $L_k^{p,d}$. The holomorphic curves in [7] all have source a strip, but this is again essentially irrelevant for his proof. The extra finite-dimensional spaces cause no additional complications. Finally, the asymptotics considered in [7] are somewhat different from the ones we consider, but as the estimate follows from uniform pointwise bounds, this is yet again irrelevant to the proof.

For slightly weaker estimates, still sufficient to apply the implicit function theorem, see [14, Section 3.3] □

Corollary A.7 *If Q is a bounded right inverse for $D\bar{\partial}(u)$ then there is a constant C , linear in $\|Q\|$, so that*

$$\|QN(\xi) - QN(\eta)\| \leq C (\|\xi\| + \|\eta\|) \|\xi - \eta\|$$

for $\xi, \eta \in V \oplus F \oplus L_k^{p,d}(u^*TM, \partial)$. Thus, the inverse function theorem applies to find nearby solutions of the $\bar{\partial}$ -equation.

Proof Take $C = \|Q\|C'$, for C' as in Lemma A.6. □

So, the strategy to prove the gluing lemma is to construct a family of pre-glued maps $(u_+ \natural_R u_-): (S^+ \natural_R S^-) \rightarrow M_R$ so that

- (1) the maps $u_+ \natural_R u_-$ converge to the height two building (u_+, u_-) and
- (2) at each $u_+ \natural_R u_-$ there is a right inverse Q_R to $D\bar{\partial}$ so that the Q_R are uniformly bounded (in R).

The first condition then implies that $D\bar{\partial}(u_+ \natural_R u_-) \rightarrow 0$ as $R \rightarrow \infty$. It follows from this and the second condition that for large enough R the implicit function theorem applies to give families of solutions of the $\bar{\partial}$ -equation near $u_+ \natural_R u_-$, and hence near (u_+, u_-) . Further, any other solution of the $\bar{\partial}$ -equation near (u_+, u_-) would lie in a small neighborhood of $u_+ \natural_R u_-$ for appropriate large R , and hence be one of the solutions given by the implicit function theorem. This then proves the gluing lemma.

Now we define the pre-glued maps $u_+ \natural_R u_-$. Choose holomorphic coordinates (t_i^\pm, s_i^\pm) ($s_i^\pm \in S^1$ for $i = 1, \dots, m_0$, $s_i^\pm \in [0, 1]$ for $i = m_0 + 1, \dots, n_0$, $t_i^\pm \in \mathbb{R}$) near the puncture $p_{0,i}^\pm$ so that $t_i^\pm \rightarrow \mp\infty$ as $p \rightarrow p_{0,i}^\pm$. For $i = 1, \dots, m_0$ we further require that $\lim_{t_i^+ \rightarrow -\infty} u_+(t_i^+, s_i^+) = \lim_{t_i^- \rightarrow \infty} u_-(t_i^-, s_i^-)$. For $i = m_0 + 1, \dots, n_0$ this is automatic. We will call such coordinates *cylindrical coordinates*.

Remark Our convention for which coordinate is denoted s and which is denoted t is exactly the opposite from [1], but agrees with the convention used in the rest of this paper.

Let τ^\pm denote the \mathbb{R} -coordinate on E_0^\pm chosen earlier in this section, let $\pi_{\tilde{E}_0}$ denote projection of E_0^\pm onto \tilde{E}_0 and fix some metric on \tilde{E}_0 so that $L \cap \tilde{E}_0$ is totally geodesic (this is possible by the Lagrangian neighborhood theorem). Near $p_{0,i}^\pm$ the map u_\pm has the form

$$\begin{aligned} \tau^\pm \circ u_\pm(t_i^\pm, s_i^\pm) &= T_{\gamma_{0,i}} t_i^\pm \pm \tau_{0,0}^\pm + \eta_i^\pm(t_i^\pm, s_i^\pm) \\ \pi_{\tilde{E}_0} \circ u &= \exp_{\gamma_{0,i}(s_i^\pm)}(U_i^\pm(t_i^\pm, s_i^\pm)) \end{aligned}$$

where η_i^\pm and U_i^\pm decay exponentially in t_i .

Choose R_0 large enough that for $R > R_0$ and $i = 1, \dots, n_0$, $(\mp \frac{R \pm \tau_{0,0}^\pm}{T_{\gamma_{0,i}}}, s_i^\pm)$ lies in the neighborhood of $p_{0,i}^\pm$ on which the coordinates in use are defined. From now on we will assume that $R > R_0$. Then $\tau^\pm \circ u_\pm(\mp \frac{R \pm \tau_{0,0}^\pm}{T_{\gamma_{0,i}}}, s_i^\pm) = \pm R$. So, define $S^+ \natural_R S^-$ to be the surface obtained by deleting $\left\{ \mp t_i^\pm > \frac{R \pm \tau_{0,0}^\pm}{T_{\gamma_{0,i}}} \right\}$ from S^\pm and gluing the

resulting surfaces along the newly created boundary in the obvious way. We will refer to the image of the neighborhoods of the punctures $p_{0,i}^\pm$ in $S^+ \natural_R S^-$ as the *necks*.

Define coordinates (t_i^0, s_i^0) on the necks by

$$\begin{cases} s_i^0 = s_i^+, & t_i^0 = t_i^+ + \frac{R+\tau_{0,0}}{T_{\gamma_{0,i}}} & \text{for } t_i^+ \geq -\frac{R+\tau_{0,0}}{T_{\gamma_{0,i}}} \\ s_i^0 = s_i^-, & t_i^0 = t_i^- - \frac{R-\tau_{0,0}}{T_{\gamma_{0,i}}} & \text{for } t_i^- \leq \frac{R-\tau_{0,0}}{T_{\gamma_{0,i}}} \end{cases} .$$

We define the pre-glued map $u_+ \natural_R u_-: S_+ \natural_R S_- \rightarrow M_R$ to agree with u_+ or u_- outside the necks. Recall that M_R has a neck $[-R, R] \times \tilde{E}_0$; let τ^R denote the $[-R, R]$ -coordinate on the neck of M_R , and $\pi_{\tilde{E}_0}$ projection of the neck of M_R onto \tilde{E}_0 . Then, on the i^{th} neck of $S_+ \natural_R S_-$, $u_+ \natural_R u_-$ is defined by

$$\begin{cases} \left. \begin{aligned} \tau^R \circ (u_+ \natural_R u_-)(s_i^0, t_i^0) &= T_{\gamma_{0,i}} t_i^0 + \beta(t_i^0 - 1) \eta_i^+(s_i^0, t_i^0) \\ \pi_{\tilde{E}_0} \circ (u_+ \natural_R u_-)(s_i^0, t_i^0) &= \exp_{\gamma_{0,i}(s_i)}(\beta(t_i^0 - 1) U_i^+(s_i^0, t_i^0)) \end{aligned} \right\} & \text{if } t_i^0 \geq 1 \\ \left. \begin{aligned} \tau^R \circ (u_+ \natural_R u_-)(s_i^0, t_i^0) &= T_{\gamma_{0,i}} t_i^0 \\ \pi_{\tilde{E}_0} \circ (u_+ \natural_R u_-)(s_i^0, t_i^0) &= \gamma_{0,i}(s_i) \end{aligned} \right\} & \text{if } -1 \leq t_i^0 \leq 1 \\ \left. \begin{aligned} \tau^R \circ (u_+ \natural_R u_-)(s_i^0, t_i^0) &= T_{\gamma_{0,i}} t_i^0 + \beta(-t_i^0 - 1) \eta_i^-(s_i^0, t_i^0) \\ \pi_{\tilde{E}_0} \circ (u_+ \natural_R u_-)(s_i^0, t_i^0) &= \exp_{\gamma_{0,i}(s_i)}(\beta(-t_i^0 - 1) U_i^-(s_i^0, t_i^0)) \end{aligned} \right\} & \text{if } t_i^0 \leq -1 \end{cases}$$

(Compare [1, page 54].) Note that since $L \cap \tilde{E}_0$ is totally geodesic and $U_i^\pm|_{\partial S^\pm}$ is tangent to $L \cap \tilde{E}_0$, this formula makes perfectly good sense in the relative setting.

We extend the gluing construction to a neighborhood of (u_+, u_-) in $\mathcal{B}^+ \times \mathcal{B}^-$. The details of this extension are unimportant, but for completeness we give them anyway. Extend the coordinates (t_i^\pm, s_i^\pm) smoothly to holomorphic coordinates on surfaces in a neighborhood of (S^+, S^-) . Let $\theta_1, \dots, \theta_{m_0}$ be in some interval around $0 \in S^1 = \mathbb{R}/\mathbb{Z}$ and r_1, \dots, r_{n_0} in \mathbb{R} , $R > R_0$. Let $\theta_{m_0+1} = \dots = \theta_{n_0} = 0$. Then, for $(u'_+: S'^+ \rightarrow M_\infty, u'_-: S'^- \rightarrow M_\infty)$, define $S'^+ \natural_{r_1, \dots, r_{n_0}, \theta_1, \dots, \theta_{m_0}} S'^-$ by deleting the disks $\left\{ \mp t_i^\pm > \frac{r_i \pm \tau_{0,0}}{T_{\gamma_{0,i}}} \right\}$ from S'^\pm and identifying the newly created boundaries by $s_i^+ \leftrightarrow s_i^- + \theta_i$. Define coordinates (t_i^0, s_i^0) on the necks of the glued surfaces by

$$\begin{cases} s_i^0 = s_i^+ - \beta(t_i^0) \theta_i & t_i^0 = \frac{R}{r_i} t_i^+ + \frac{r_i + \tau_{0,0}}{T_{\gamma_{0,i}}} & \text{for } t_i^+ \geq -\frac{r_i + \tau_{0,0}}{T_{\gamma_{0,i}}} \\ s_i^0 = s_i^-, & t_i^0 = t_i^- - \frac{R}{r_i} \frac{r_i - \tau_{0,0}}{T_{\gamma_{0,i}}} & \text{for } t_i^- \leq \frac{r_i - \tau_{0,0}}{T_{\gamma_{0,i}}} \end{cases}$$

Then, define $u'_+ \natural_{r_1, \dots, r_{n_0}, \theta_1, \dots, \theta_{m_0}} u'_-$ by the same formula used to define $u_+ \natural_R u_-$ above. Let \mathcal{T}_R denote the space of conformal structures given by the gluing construction

just explained. Note that \mathcal{T}_R projects onto an open neighborhood of $S^+ \natural_R S^-$ in moduli space. Let \mathcal{B}_R denote the space of $W_k^{p,d}$ maps from surfaces in \mathcal{T}_R with the specified boundary conditions and asymptotics. Note that the glued maps fill out a neighborhood of $u_+ \natural_R u_-$ in \mathcal{B}_R .

Recall that we defined spaces $W_{u^\pm}^\pm =: W^\pm$ as the image under $D\bar{\partial}$ of $V^\pm = T_{S^\pm} \mathcal{T}$. Since the conformal structures in \mathcal{T} agree near the punctures, we can view $V^+ \oplus V^-$ as a subspace of $\Gamma(\Lambda^{0,1} T^*(S^+ \natural_R S^-))$ and $W^+ \oplus W^-$ as a subspace of $(u_+ \natural_R u_-)^* TM_R$ for any R .

Recall also that we defined spaces $F^\pm = \mathbb{C}^{\sum_{i=0}^N m_i^\pm} \oplus \mathbb{R}^{\sum_{i=0}^N (n_i^\pm - m_i^\pm)}$ corresponding to the span of a fixed collection of $\sum_{i=0}^N m_i^\pm$ vector fields $v_{\tau,i,j}^\pm$ and $v_{\theta,i,j}^\pm$ given by $\frac{\partial}{\partial \tau}$ and $\frac{\partial}{\partial \theta}$ near the $\sum_{i=0}^N m_i^\pm$ interior punctures and a fixed collection of $\sum_{i=0}^N (n_i^\pm - m_i^\pm)$ vector fields $v_{\tau,i,j}^\pm$ given by $\frac{\partial}{\partial \tau}$ near the $\sum_{i=0}^N (n_i^\pm - m_i^\pm)$ boundary punctures of u^\pm . Let

$$F^0 = \text{Span}\{v_{\tau,0,j}, v_{\theta,0,j}\} = \mathbb{C}^{\sum_{i=1}^N (m_i^+ m_i^-)} \oplus \mathbb{R}^{\sum_{i=1}^N (n_i^+ + n_i^- - m_i^+ - m_i^-)}$$

be the sections corresponding to punctures which are not being glued.

The tangent space $T\mathcal{T}_R$ is $\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0 - m_0} \oplus V^+ \oplus V^-$. (The \mathbb{C} - and \mathbb{R} -summands correspond to the gluing parameters θ_i and r_i .) So, the tangent space to \mathcal{B}_R is

$$(12) \quad T_{u_+ \natural_R u_-} \mathcal{B}_R = \mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0 - m_0} \oplus V^+ \oplus V^- \oplus F^0 \oplus L_k^{p,d}((u_+ \natural_R u_-)^* TM_R, \partial).$$

The map $V^+ \oplus V^- \rightarrow W^+ \oplus W^-$ induced by $D\bar{\partial}$ is a surjective map of finite-dimensional vector spaces, and independent of R . It therefore has a uniformly bounded right inverse. It therefore suffices to construct a uniformly bounded right inverse to the map

$$\begin{aligned} \mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0 - m_0} \oplus W^+ \oplus W^- \oplus F^0 \oplus L_k^{p,d}((u_+ \natural_R u_-)^* TM_R, \partial) \\ \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1}(u_+ \natural_R u_-)^* TM_R) \end{aligned}$$

given by

$$(v_{\mathbb{C}}, v_{\mathbb{R}}, v_+, v_-, \xi_0, \xi) \mapsto v_+ + v_- + (D\bar{\partial})(v_{\mathbb{C}} + v_{\mathbb{R}}) + (D\bar{\partial}_{S^+ \natural_R S^-})(\xi_0 + \xi).$$

Here, $v_{\mathbb{C}}$ and $v_{\mathbb{R}}$ correspond to infinitesimal variations of the almost complex structure j_S , and $(D\bar{\partial})(v_{\mathbb{C}} + v_{\mathbb{R}})$ the image of the sum of these variations under $D\bar{\partial}$. If one views v_+ and v_- as sections of $\text{End}(TS, j)$ as defined in Section 3 then $D\bar{\partial}(v_{\mathbb{C}} + v_{\mathbb{R}})$ is given by $J \circ d(u_+ \natural_R u_-) \circ (v_{\mathbb{C}} + v_{\mathbb{R}})$.

Constructing the right inverse and proving its boundedness is what we shall do for most of the remainder of this section.

On spaces of sections over $u_+ \natural_R u_-$ instead of the Banach norm specified earlier in this section we use a Banach norm with additional weights of $e^{d(R/T-s_i^0)}$ on the necks. This paragraph and the next make this precise. Choose metrics $\langle \cdot, \cdot \rangle_R$ on $S^+ \natural_R S^-$ and diffeomorphisms $\phi_{R_0, R}: S^+ \natural_{R_0} S^- \rightarrow S^+ \natural_R S^-$ for so that $\langle \cdot, \cdot \rangle_{R_0} \leq \phi^* \langle \cdot, \cdot \rangle_R \leq \frac{2R}{R_0} \langle \cdot, \cdot \rangle_{R_0}$, pointwise. Integrals over S_R will be with respect to the volume forms induced by $\langle \cdot, \cdot \rangle_R$. Choose metrics $\langle \cdot, \cdot \rangle_R^M$ on M_R and diffeomorphisms $\phi_{R_0, R}^M: M_{R_0} \rightarrow M_R$ such that $\langle \cdot, \cdot \rangle_{R_0}^M \leq \phi^* \langle \cdot, \cdot \rangle_R^M \leq \frac{2R}{R_0} \langle \cdot, \cdot \rangle_{R_0}^M$, pointwise. The bundles $(u_+ \natural_R u_-)^* TM_R$ and $\Lambda^{0,1}(u_+ \natural_R u_-)^* TM_R$ inherit metrics from $\langle \cdot, \cdot \rangle_R$ and $\langle \cdot, \cdot \rangle_R^M$. Norms of elements of these vector bundles will be taken with respect to the induced metrics.

Choose cylindrical coordinates $(t_{i,j}^\pm, s_{i,j}^\pm)$ near the punctures $p_{i,j}^\pm$. Let ξ be a section of $(u_+ \natural_R u_-)^* TM_R$. For $1 \leq i \leq n_0$ define $\pi_{\tau,i}(\xi) = \int_{t_i^0=0} \left\langle \frac{\partial}{\partial \tau}, \xi \right\rangle$ and for $1 \leq i \leq m_0$ define $\pi_{\theta,i}(\xi) = \int_{t_i^0=0} \left\langle \frac{\partial}{\partial \theta}, \xi \right\rangle$. Then, define

$$\begin{aligned} \bar{\xi} = & \sum_{i=1}^{n_0} \pi_{\tau,i}(\xi) \left(1 - \beta_{i,R}(t_i^0)\right) \left(1 - \beta_{i,R}(-t_i^0)\right) \frac{\partial}{\partial \tau} \\ & + \sum_{i=1}^{m_0} \pi_{\theta,i}(\xi) \left(1 - \beta_{i,R}(t_i^0)\right) \left(1 - \beta_{i,R}(-t_i^0)\right) \frac{\partial}{\partial \theta}. \end{aligned}$$

One can think of $\bar{\xi}$ as an approximate projection of ξ to $\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0}$ in $F^+/(F^0 \cap F^+)$.

Then, the norm of a section ξ of $(u_+ \natural_R u_-)^* TM_R$ is given by the sum of the norm of the vector

$$(\pi_{\tau,1}(\xi), \dots, \pi_{\tau,n_0}(\xi), \pi_{\theta,1}(\xi), \dots, \pi_{\theta,m_0}(\xi)) \in \mathbb{R}^{n_0+m_0}$$

and

$$\sup_{\substack{|\alpha| \leq k \\ \left| \frac{\partial}{\partial \alpha} \right| = 1}} \left(\int_{S_+ \natural_R S_-} \left| \frac{\partial}{\partial \alpha} \left[\left(1 + \sum_{i=0}^{n_0} e^{d(R-|t_i^0|)/p} \beta(R-|t_i^0|) + \sum_{\pm} \sum_{i=1}^N \sum_{j=1}^{n_i} e^{d|t_{i,j}^\pm|/p} \beta(|t_{i,j}^\pm|) \right) (\xi - \bar{\xi}) \right] \right|^p dV \right)$$

where the sup is over all partial derivative of order at most k and norm 1.

The norm of a section η of $\Lambda^{0,1}(u_+ \natural_R u_-)^* TM_R$ is given by

$$\sup_{\substack{|\alpha| \leq k-1 \\ \left| \frac{\partial}{\partial \alpha} \right| = 1}} \left(\int_{S_+ \natural_R S_-} \left| \frac{\partial}{\partial \alpha} \left[\left(1 + \sum_{i=0}^{n_0} e^{d(R-|t_i^0|)/p} \beta(R-|t_i^0|) + \sum_{\pm} \sum_{i=1}^N \sum_{j=1}^{n_i} e^{d|t_{i,j}^{\pm}|/p} \beta(|t_{i,j}^{\pm}|) \right) \eta \right] \right|^p dV \right).$$

It is not necessary to split η into two parts.

Lemma A.8 (Compare [1, Lemma 5.5])

With this Banach structure on $\Gamma(\Lambda^{0,1}(u_+ \natural_R u_-))$, if u_{\pm} are holomorphic then

$$\lim_{R \rightarrow \infty} \|\bar{\partial} u_+ \natural_R u_-\|_R = 0.$$

Proof This is the same as in [1]. Since u_{\pm} are holomorphic, the section $\bar{\partial} u_+ \natural_R u_-$ is identically zero except in regions $-2 \leq t_i^0 \leq -1$ and $1 \leq t_i^0 \leq 2$ in the necks. In these regions we crudely bound $\|\bar{\partial} u_+ \natural_R u_-\|$ by

$$\sum_i \left(\left\| \chi \left[-\frac{R+\tau_{0,0}^-}{T}, -\frac{R+\tau_{0,0}^-}{T} + 2 \right] (\nabla U_i^-, \nabla \eta_i^-) \right\| + \left\| \chi \left[\frac{R-\tau_{0,0}^-}{T} - 2, \frac{R-\tau_{0,0}^+}{T} \right] (\nabla U_i^+, \nabla \eta_i^+) \right\| \right).$$

Recall that these norms are weighted by $e^{d(R/T-s_i^0)}$. However, since u_{\pm} are in $\mathcal{B}_k^{p,d}$, the sections $(U_i^{\pm}, \eta_i^{\pm})$ are in $L_k^{p,d}$ and so their derivatives are in $L_{k-1}^{p,d}$. It follows that the right hand side goes to zero as $R \rightarrow \infty$. \square

Proposition A.9 Suppose that $(D\bar{\partial})_{u_{\pm}}$, the linearized $\bar{\partial}$ -operator at u_{\pm} , is surjective. Then, for large enough R , the operator $(D\bar{\partial})_{u_+ \natural_R u_-}$ has a right inverse Q_R which is uniformly bounded in R .

Proof As with all of this section, I learned this proof mainly from Bourgeois [1]. Bourgeois in turn cites McDuff–Salamon [14], who say they adopted the argument from Donaldson–Kronheimer [4].

Once one has seen how it goes, the proof is not particularly hard. To prove the proposition we define linearized gluing and splitting maps

$$g_R: \Delta \oplus L_k^{p,d}(u_+^* TM_{\infty}, \partial) \oplus L_k^{p,d}(u_-^* TM_{\infty}, \partial) \rightarrow L_k^{p,d}((u_+ \natural_R u_-)^* TM_R, \partial)$$

and

$$s_R: L_{k-1}^{p,d}(\Lambda^{0,1}(u_+ \natural_R u_-)^* TM_R) \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1} u_+^* TM_{\infty}) \oplus L_{k-1}^{p,d}(\Lambda^{0,1} u_-^* TM_{\infty}).$$

where Δ is the diagonal in $\mathbb{C}^{m_0^+} \oplus \mathbb{R}^{n_0^+ - m_0^+} \oplus \mathbb{C}^{m_0^-} \oplus \mathbb{R}^{n_0^- - m_0^-} \subset F^+ \oplus F^-$ (so Δ corresponds to the sections at the punctures being glued which agree on the two sides).

Let $G_R = s_R \left(D\bar{\partial}(\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0}) \right)$ where $\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0}$ denotes the tangent space to the space of gluing parameters, as in Equation (12). (The map $s_R \circ D\bar{\partial}$ is an isomorphism when restricted to this $\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0}$, and so identifies G_R with $\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0}$. So, we will sometimes abuse notation and use G_R when we mean its preimage under $s_R \circ D\bar{\partial}$.)

We will check that

$$D\bar{\partial}_{(u_+, u_-)}: W^+ \oplus W^- \oplus \Delta \oplus F^0 \oplus G_R \oplus L_k^{p,d}(u_+^* TM_\infty, \partial) \oplus L_k^{p,d}(u_-^* TM_\infty, \partial) \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1} u_+^* TM_\infty) \oplus L_{k-1}^{p,d}(\Lambda^{0,1} u_-^* TM_\infty).$$

is surjective and has a uniformly bounded right inverse Q_∞ for large R .

We will then define an approximate right inverse \tilde{Q}_R for

$$D\bar{\partial}_{u_+ \natural_R u_-}: \mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0} \oplus W^+ \oplus W^- \oplus F^0 \oplus L_k^{p,d}((u_+ \natural_R u_-)^* TM_R, \partial) \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1}(u_+ \natural_R u_-)^* TM_R)$$

by the commutative diagram

$$\begin{array}{ccc} L_{k-1}^{p,d}(\Lambda^{0,1}(u_+ \natural_R u_-)^* TM_R) & \xrightarrow{\tilde{Q}_R} & W^+ \oplus W^- \oplus F^0 \oplus \mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0-m_0} \\ & & \oplus L_k^{p,d}((u_+ \natural_R u_-)^* TM_R, \partial) \\ \downarrow s_R & & \uparrow id \oplus id \oplus id \oplus id \oplus g_R \\ L_{k-1}^{p,d}(\Lambda^{0,1} u_+^* TM_\infty) \oplus L_{k-1}^{p,d}(\Lambda^{0,1} u_-^* TM_\infty) & \xrightarrow{Q_\infty} & W^+ \oplus W^- \oplus F^0 \oplus G^R \oplus \Delta \\ & & \oplus L_k^{p,d}(u_+^* TM_\infty, \partial) \oplus L_k^{p,d}(u_-^* TM_\infty, \partial). \end{array}$$

We will check that s_R and g_R are uniformly bounded, so that \tilde{Q}_R is, also. We will then show that $\|(D\bar{\partial})_{u_+ \natural_R u_-} \tilde{Q}_R - I\| \leq 1/2$, so that $(D\bar{\partial})_{u_+ \natural_R u_-} \tilde{Q}_R$ is invertible; the inverse of $(D\bar{\partial})_{u_+ \natural_R u_-}$ is then given by $\tilde{Q}_R \left((D\bar{\partial})_{u_+ \natural_R u_-} \tilde{Q}_R \right)^{-1}$, which is bounded by $2\|\tilde{Q}_R\|$.

For

$$\begin{aligned} & (\xi_+, \xi_-) \in \Delta \oplus L_k^{p,d}(u_+^* TM_\infty, \partial) \oplus L_k^{p,d}(u_-^* TM_\infty, \partial) \\ & \subset \mathbb{C}^{m_0^+} \oplus \mathbb{R}^{n_0^+-m_0^+} \oplus L_k^{p,d}(u_+^* TM_\infty, \partial) \oplus \mathbb{C}^{m_0^-} \oplus \mathbb{R}^{n_0^- - m_0^-} \oplus L_k^{p,d}(u_-^* TM_\infty, \partial) \end{aligned}$$

define $g_R(\xi_+, \xi_-)$ to agree with ξ_+ or ξ_- outside the necks and by

$$g_R(\xi_+, \xi_-) = \begin{cases} \xi_+(t_i^0, s_i^0) + (1 - \beta_{i,R}(t_i^0)) \xi_-(t_i^0, s_i^0) & \text{if } t_i^0 \geq 1 \\ \xi_+(t_i^0, s_i^0) + \xi_-(t_i^0, s_i^0) & \text{if } -1 \leq t_i^0 \leq 1 \\ (1 - \beta_{i,R}(-t_i^0)) \xi_+(t_i^0, s_i^0) + \xi_-(t_i^0, s_i^0) & \text{if } t_i^0 \leq -1 \end{cases}$$

in the necks. (Compare [1, pages 56–57].) Note that this formula makes as much sense in the relative case as in the closed case. Exactly the same formulas define a linearized gluing map

$$g_R: L_{k-1}^{p,d} \left(\Lambda^{0,1} u_+^* TM_\infty \right) \oplus L_{k-1}^{p,d} \left(\Lambda^{0,1} u_-^* TM_\infty \right) \rightarrow L_{k-1}^{p,d} \left(\Lambda^{0,1} (u_+ \natural_R u_-)^* TM_R \right),$$

which shall be useful when we estimate $(D\bar{\partial})\tilde{Q}_R - I$.

Define $s_R(\eta) = (\eta_+, \eta_-)$ where η_\pm agrees with η away from the punctures and near $p_{0,i}^\pm$ is given by

$$\begin{aligned} \eta_+(t_i^0, s_i^0) &= \beta(t_i^0)\eta(t_i^0, s_i^0) \\ \eta_-(t_i^0, s_i^0) &= (1 - \beta(t_i^0))\eta(t_i^0, s_i^0). \end{aligned}$$

(Compare [1, page 57].)

Lemma A.10 *The maps g_R and s_R are uniformly bounded in R .*

Proof Although for notational reasons it may appear involved, the proof is in fact straightforward. Let $(\xi_0, \xi_+, \xi_-) \in \Delta \oplus L_k^{p,d} (u_+^* TM_\infty, \partial) \oplus L_k^{p,d} (u_-^* TM_\infty, \partial)$. Recall that, by definition,

$$\|g_R(\xi_0, \xi_+, \xi_-)\| = \|\overline{g_R(\xi_0, \xi_+, \xi_-)}\| + \|g_R(\xi_0, \xi_+, \xi_-) - \overline{g_R(\xi_0, \xi_+, \xi_-)}\|_{L_k^{p,d}}$$

where $\overline{g_R(\xi_0, \xi_+, \xi_-)}$ is a certain projection of $g_R(\xi_0, \xi_+, \xi_-)$ to $\mathbb{C}^{m_0} \oplus \mathbb{R}^{n_0}$. Observe that as $R \rightarrow \infty$, $\overline{g_R(\xi_0, \xi_+, \xi_-)} \rightarrow \xi_0$, and hence $\|\overline{g_R(\xi_0, \xi_+, \xi_-)}\|$ is bounded. In fact, by the Sobolev inequalities,

$$\|\overline{g_R(\xi_0, \xi_+, \xi_-)}\| \leq C \left(\|\xi_0\| + \|\xi_+\|_{L_k^{p,d}} + \|\xi_-\|_{L_k^{p,d}} \right)$$

where C depends only on the cutoff function β , and not on R .

On the other hand,

$$\begin{aligned} \|g_R(\xi_0, \xi_+, \xi_-) - \overline{g_R(\xi_0, \xi_+, \xi_-)}\|_{L_k^{p,d}} &\leq \|g_R(0, \xi_+, \xi_-)\| + \|g_R(\xi_0, \xi_+, \xi_-) \\ &\quad - g_R(0, \xi_+, \xi_-) - \overline{g_R(\xi_0, \xi_+, \xi_-)}\|_{L_k^{p,d}} \\ &= \|g_R(0, \xi_+, \xi_-)\|_{L_k^{p,d}} + \|g_R(\xi_0, 0, 0) \\ &\quad - \overline{g_R(\xi_0, \xi_+, \xi_-)}\|_{L_k^{p,d}}. \end{aligned}$$

Now, on the i^{th} neck,

$$g_R(\xi_0, 0, 0) - \overline{g_R(\xi_0, \xi_+, \xi_-)} = (1 - \beta_{i,R}(t_i^0))(1 - \beta_{i,R}(-t_i^0)) \left[\xi_0 - \xi_0 - \int_{t_i^0=0} \left(\left\langle \frac{\partial}{\partial \theta}, \xi_+ + \xi_- \right\rangle \frac{\partial}{\partial \theta} + \left\langle \frac{\partial}{\partial \tau}, \xi_+ + \xi_- \right\rangle \frac{\partial}{\partial \tau} \right) \right].$$

So,

$$\|g_R(\xi_0, 0, 0) - \overline{g_R(\xi_0, \xi_+, \xi_-)}\|_{L_k^{p,d}} \leq C \left(\|\xi_+\|_{L_k^{p,d}} + \|\xi_-\|_{L_k^{p,d}} \right)$$

where C depends only on the cutoff function β , and not on R .

Finally,

$$\begin{aligned} \|g_R(0, \xi_+, \xi_-)\|_{L_k^{p,d}} &\leq \|\xi_+\|_{L_k^{p,d}} + \|\xi_-\|_{L_k^{p,d}} \\ &+ \sum_{i=1}^{n_0} \sum_{|\alpha| \leq k} \left(\int_{s_i^0 \in S^1} \int_{t_i^0 = -R}^{-1} \left| \frac{\partial}{\partial \alpha} \left(1 - \beta_{i,R}(-t_i^0) \xi_+(t_i^0, s_i^0) + \xi_-(t_i^0, s_i^0) \right) e^{d|t_i^0|/p} \right|^p dt_i^0 \right. \\ &+ \int_{s_i^0 \in S^1} \int_{t_i^0 = -1}^1 \left| \frac{\partial}{\partial \alpha} \left(\xi_+(t_i^0, s_i^0) + \xi_-(t_i^0, s_i^0) \right) e^{d|t_i^0|/p} \right|^p dt_i^0 \\ &\left. + \int_{s_i^0 \in S^1} \int_{t_i^0 = 1}^R \left| \frac{\partial}{\partial \alpha} \left(1 - \xi_+(t_i^0, s_i^0) + (1 - \beta_{i,R}(t_i^0)) \xi_-(t_i^0, s_i^0) \right) e^{d|t_i^0|/p} \right|^p dt_i^0 \right)^{1/p} \\ &\leq \|\xi_+\|_{L_k^{p,d}} + \|\xi_-\|_{L_k^{p,d}} + \sum_{i=1}^{n_0} \sum_{|\alpha| \leq k} C_\alpha \left(\|\xi_+\|_{L_k^{p,d}} + \|\xi_-\|_{L_k^{p,d}} \right) \\ &\leq C \left(\|\xi_+\|_{L_k^{p,d}} + \|\xi_-\|_{L_k^{p,d}} \right). \end{aligned}$$

Here, the terms $\|\xi_+\|_{L_k^{p,d}}$ and $\|\xi_-\|_{L_k^{p,d}}$ on the right hand side of the first inequality take care of the contribution of $S \setminus \{\text{the necks}\}$ to $\|g_R(0, \xi_+, \xi_-)\|_{L_k^{p,d}}$. The second inequality follows from the bound on the derivatives of $\beta_{i,R}$. This proves boundedness of g_R . The C_α are universal constants depending only on α and the cutoff function β , not on R .

Turning to s_R , we have

$$(13) \quad \|s_R(\eta)\|_{L_{k-1}^{p,d}} = \|\beta(t_i^0)\eta(t_i^0, s_i^0)\|_{L_{k-1}^{p,d}} + \left\| \left(1 - \beta(t_i^0) \right) \eta(t_i^0, s_i^0) \right\|_{L_{k-1}^{p,d}}.$$

Focussing on the first term on the right hand side, we have

$$\begin{aligned} & \|\beta(t_i^0)\eta(t_i^0, s_i^0)\|_{L_{k-1}^{p,d}} \\ & \leq \|\eta\|_{L_{k-1}^{p,d}} + \sum_{i=1}^{n_0} \sum_{|\alpha|\leq k-1} \left(\int_{s_i^+ \in S^1} \int_{t_i^+ = -\infty}^{-\frac{R+\tau_{0,0}}{T_{\gamma_{0,i}^+}}} \left| \frac{\partial}{\partial \alpha} \beta \left(t_i^+ + \frac{R+\tau_{0,0}}{T_{\gamma_{0,i}^+}} \right) \right. \right. \\ & \quad \left. \left. \times \eta \left(t_i^+ + \frac{R+\tau_{0,0}}{T_{\gamma_{0,i}^+}}, s_i^+ \right) e^{d|t_i^+|/p} \right|^p dt_i^+ \right)^{1/p} \\ & \leq \|\eta\|_{L_{k-1}^{p,d}} + \sum_{i=1}^{n_0} \sum_{|\alpha|\leq k-1} C'_\alpha \|\eta\|_{L_{k-1}^{p,d}} \\ & \leq C' \|\eta\|_{L_{k-1}^{p,d}} \end{aligned}$$

where the inequalities follow by the same reasoning as for $g_R(0, \xi_+, \xi_-)$. A similar argument applies to the second term on the right hand side of Equation (13), so s_R is uniformly bounded. \square

Lemma A.11 *The linearized $\bar{\partial}$ -operator*

$$\begin{aligned} D\bar{\partial}_{(u_+, u_-)}: W^+ \oplus W^- \oplus \Delta \oplus F^0 \oplus G_R \oplus L_k^{p,d}(u_+^* TM_\infty, \partial) \oplus L_k^{p,d}(u_-^* TM_\infty, \partial) \\ \rightarrow L_{k-1}^{p,d}(\Lambda^{0,1} u_+^* TM_\infty) \oplus L_{k-1}^{p,d}(\Lambda^{0,1} u_-^* TM_\infty). \end{aligned}$$

is surjective and has a uniformly bounded right inverse Q_∞ for large R .

Proof By assumption,

$$D\bar{\partial}_{u_\pm}: (W_{u_\pm}^\pm \oplus F^\pm \oplus L_k^{p,d}(u_\pm^* TM_\infty, \partial)) \rightarrow L_{k-1}^{p,d}(u_\pm^* TM_\infty)$$

is surjective. Recall that the inclusion of F^\pm in $\Gamma(u_\pm^* TM_\infty)$ required choosing particular sections $v_{\tau,i,j}^\pm$ and $v_{\theta,i,j}^\pm$, constant near the punctures, but the image of $D\bar{\partial}_{u_\pm}$ was independent of these choices. For appropriate choices of these sections, $D\bar{\partial}(F^0 \cap F^\pm) + D\bar{\partial}(G_R)$ is exactly F^\pm . The desired surjectivity follows. Uniform boundedness only requires that one observe that the ‘‘appropriate choices’’ converge, which is clear. (See also [1, Corollary 5.7].) \square

Now we estimate $(D\bar{\partial})\tilde{Q}_R - I$. Given a section (ξ_+, ξ_-) of $\Delta \oplus L_k^{p,d}(u_+^* TM_\infty, \partial) \oplus L_k^{p,d}(u_-^* TM_\infty, \partial)$, observe that

$$((D\bar{\partial})g_R - g_R(D\bar{\partial}))(\xi_+, \xi_-) = \begin{cases} 0 & \text{outside the necks} \\ ((D\bar{\partial})\beta_{i,R}(t_i^0))\xi_-(t_i^0, s_i^0) & \text{if } t_i^0 \geq 1 \\ 0 & \text{if } -1 \leq t_i^0 \leq 1 \\ -((D\bar{\partial})\beta_{i,R}(-t_i^0))\xi_+(t_i^0, s_i^0) & \text{if } t_i^0 \leq -1. \end{cases}$$

Now, $D\bar{\partial}$ is a pure first order differential operator, so $|(D\bar{\partial})\beta_{i,R}(t_i^0)| \leq C_1/(R/2 - T_{\gamma_{0,i}^\pm}) \leq C_2/R$ for R large and some constant C_2 . So, for large R ,

$$\|((D\bar{\partial})g_R - g_R(D\bar{\partial}))(\xi_+, \xi_-)\| \leq \frac{C_2}{R} \|(\xi_+, \xi_-)\|.$$

Also, outside the necks $g_R \circ s_R(\eta) = \eta$, while in the necks

$$\begin{aligned} g_R \circ s_R(\eta) &= g_R(\beta(t_i^0)\eta(t_i^0, s_i^0), (1 - \beta(t_i^0))\eta(t_i^0, s_i^0)) \\ &= \begin{cases} \beta(t_i^0)\eta(t_i^0, s_i^0) + (1 - \beta_{i,R}(t_i^0))(1 - \beta(t_i^0))\eta(t_i^0, s_i^0) & \text{if } t_i^0 \geq 1 \\ (\beta(t_i^0) + (1 - \beta(t_i^0)))\eta(t_i^0, s_i^0) = \eta(t_i^0, s_i^0) & \text{if } -1 \leq t_i^0 \leq 1 \\ (1 - \beta_{i,R}(-t_i^0))\beta(t_i^0)\eta(t_i^0, s_i^0) + (1 - \beta(t_i^0))\eta(t_i^0, s_i^0) & \text{if } t_i^0 \leq -1. \end{cases} \\ &= \eta(t_i^0, s_i^0). \end{aligned}$$

It follows that

$$\begin{aligned} \|D\bar{\partial}\tilde{Q}_R\eta - \eta\| &= \|(D\bar{\partial})(id \oplus id \oplus id \oplus id \oplus g_R)Q_\infty s_R\eta - \eta\| \\ &= \|(D\bar{\partial}(id \oplus id \oplus id \oplus id \oplus g_R) - (g_R)D\bar{\partial})Q_\infty s_R\eta \\ &\quad + g_R(D\bar{\partial})Q_\infty s_R\eta - \eta\| \\ &= \|(D\bar{\partial}(id \oplus id \oplus id \oplus id \oplus g_R) - (g_R)(D\bar{\partial}))Q_\infty s_R\eta + \eta - \eta\| \\ &\leq \frac{C_2}{R} \|Q_\infty s_R\eta\| \\ &\leq \frac{C_3}{R} \|\eta\| \\ &\rightarrow 0 \quad \text{as } R \rightarrow \infty. \end{aligned}$$

So, for R large enough we have $\|(D\bar{\partial})\tilde{Q}_R - I\| \leq 1/2$, so $(D\bar{\partial})\tilde{Q}_R$ is invertible for large enough R . Thus, $Q_R = \tilde{Q}_R \left((D\bar{\partial})\tilde{Q}_R \right)^{-1}$ is a uniformly bounded right inverse to $D\bar{\partial}$.

Proof of Proposition A.2 By Corollary A.7 and Proposition A.9, the implicit function theorem (Lemma A.5) proves the result. \square

Proof of Proposition A.3 The proof of Proposition A.2 all takes place in a small neighborhood of the split holomorphic curves in question. There are, therefore, no additional complications because the spaces W_∞ involved in Proposition A.3 have two kinds of ends: near each end the asymptotics are exactly the kind considered in the proof of Proposition A.2.

Therefore, the one point to check is that the split complex structure achieves transversality for maps $\mathbb{D}^2 \coprod \left(\coprod_{k=1}^\ell T^2 \right) \rightarrow T^2 \times [0, 1] \times \mathbb{R}$. Unfortunately, this isn't true. Presumably these maps do achieve transversality if one considers them as lying in the space of maps from a torus-with-boundary to $T^2 \times [0, 1] \times \mathbb{R}$. I do not even know how to formulate this statement properly, however. So, instead we use a rather indirect argument.

Fix ℓ distinct points w_1, \dots, w_ℓ in $(0, 1) \times \mathbb{R}$, and a point $z_0 \in T^2$. Fix an almost complex structure J on $T^2 \times [0, 1] \times \mathbb{R}$ satisfying (J1)–(J4) and (J5'), which is split near $\{z_0\} \times [0, 1] \times \mathbb{R}$. Let \mathcal{N}_J denote the space of maps u from $\mathbb{D}^2 \#_{k=1}^\ell T^2$ with ℓ marked points p_1, \dots, p_ℓ to $T^2 \times [0, 1] \times \mathbb{R}$ in the homology class $\ell[T^2]$, so that $u(p_i) = (z_0, w_i)$. For a generic choice of J , we have transversality of the $\bar{\partial}$ -operator, and the space \mathcal{N}_J is a smooth, compact, oriented 0-manifold. We want to show that $\#\mathcal{N}_J = 1$; then we will use \mathcal{N}_J to perform the gluing.

To show $\#\mathcal{N}_J = 1$, first consider the case $\ell = 1$. We show $\#\mathcal{N}_J = 1$ by using a special case of stabilization invariance. Specifically, consider the two Heegaard diagrams for $S^1 \times S^2$ shown in Figure 19. They differ by two handleslides of α_1 over α_2 . For an appropriate choice of coherent orientation system, in the first diagram (\mathcal{H}_1), it is easy to see that $\widehat{HF} = \mathbb{Z} \oplus \mathbb{Z}$. It follows from handleslide invariance that the same is true for the second diagram (\mathcal{H}_2). It follows that, for D the domain indicated by the numbers in the second diagram, $\#\mathcal{M}^D = \pm 1$.

Now, stretch the neck in the second diagram along the dark dotted circle C in Figure 19. In the limit, \mathcal{H}_2 degenerates to $T_{S^1 \times S^2}^2 \vee T_{S^3}^2$. Choose an almost complex structure on \mathcal{H}_2 so that, after stretching the neck, the corresponding almost complex structure on $T_{S^3}^2$ is J .

Sublemma A.12 *Compactness still holds in the current context. That is, let J_0 be an almost complex structure on \mathcal{H}_2 satisfying (J1)–(J4) and (J5'), which is split near C . Let J_R denote the almost complex structure on \mathcal{H}_2 obtained by inserting a neck of modulus R along C . Let $\{u_R: S_R \rightarrow \mathcal{H}_2 \times [0, 1] \times \mathbb{R}\}$ be a sequence of*




Figure 19: Model stabilization degeneration

J_R -holomorphic curves with $\text{ind}(u_R) = 1$. Then there is a subsequence of $\{u_R\}$ which converges to a holomorphic twin tower. (Cf Proposition 12.4.)

Proof We just sketch the proof.

Add enough marked points to each component of each S_R to stabilize it. Replacing the u_R by a subsequence if necessary, we may assume that the S_R converge. Now, consider the thick–thin decomposition of S_R . For given $\epsilon > 0$ and large R , $(\pi_\Sigma \circ u_R)^{-1}(C)$ lies in $\text{Thin}_\epsilon(S_R)$. For small ϵ and large R , the component of $\text{Thin}_\epsilon(S_R)$ containing C is an annulus in a neighborhood of which J_R is split.

On the thick part of S_R convergence of a subsequence follows from the Gromov–Schwarz lemma and Arzela–Ascoli theorem, in the standard way (cf [2, Section 10.2.2]). So, from now on assume the u_R converge in the thick part.

On all of the thin part of S_R except the components intersecting $(\pi_\Sigma \circ u_R)^{-1}(C)$, convergence of a subsequence also follows in the standard way (cf [2, Section 10.2.3]). So, from now on assume the u_R converge in these thin parts.

In components of $\text{Thin}_\epsilon(S_R)$ intersecting $(\pi_\Sigma \circ u_R)^{-1}(C)$, we extract convergence of subsequences of $\pi_\Sigma \circ u_R$ and then $\pi_{\mathbb{D}} \circ u_R$ by viewing them as sequences of holomorphic curves with converging Lagrangian boundary conditions, as in the proof of Proposition 12.4. That is, fix circles C_L and C_R in \mathcal{H}_2 so that the region between C_L and C_R is an annulus A containing the component of $\text{Thin}_\epsilon(\mathcal{H}_2)$ containing C , and so that J_R is split over A (for all R). Let $S_R^A = (\pi_\Sigma \circ u_R)^{-1}(A)$. The maps $(\pi_\Sigma \circ u_R)|_{S_R^A}: S_R^A \rightarrow A$ are a sequence of maps with Lagrangian boundary conditions, hence have a convergent subsequence. So, from now on assume that the $\pi_\Sigma \circ u_R|_{S_R^A}$ converge.

Finally, $(\pi_{\mathbb{D}} \circ u_R)((\pi_\Sigma \circ u_R)^{-1}(C_L \cup C_R))$ is a collection of circles in $[0, 1] \times \mathbb{R}$, which converges as $R \rightarrow \infty$ to a collection of circles in $[0, 1] \times \mathbb{R}$. Viewing $\pi_{\mathbb{D}} \circ u_R|_{S_R^A}$ as a family of maps with these Lagrangian boundary conditions, we obtain a convergent subsequence. Replacing the u_R with this subsequence, we finally have a convergent sequence of $u_R: S_R \rightarrow \mathcal{H}_2 \times [0, 1] \times \mathbb{R}$. \square

Let $z_{S^1 \times S^2}$ (respectively z_{S^3}) denote the wedge point in $T_{S^1 \times S^2}^2$ (respectively $T_{S^3}^2$). The domain D degenerates to a domain $D_{S^1 \times S^2}$ in $T_{S^1 \times S^2}^2$ (respectively D_{S^3} in $T_{S^3}^2$). Now, there is clearly a unique holomorphic curve $u_{S^1 \times S^2}$ in $\mathcal{M}^{D_{S^1 \times S^2}}$.

By choosing C appropriately, we can force z_{S^3} to be the point z_0 and

$$\pi_{\mathbb{D}} \circ u_{S^1 \times S^2} \left((\pi_\Sigma \circ u_{S^1 \times S^2})^{-1}(z_{S^1 \times S^2}) \right)$$

to be the point w_1 . Then, the claim that $\#\mathcal{N}_J = \pm 1$ follows from the fact that $\#\mathcal{M}^D = \pm 1$, the gluing result Proposition A.2, and the remark at the beginning of this proof. This deals with the case $\ell = 1$.

Now, the fact that $\#\mathcal{N}_J = \pm 1$ for general ℓ follows from the $\ell = 1$ case. Let u denote a holomorphic curve constructed for the $\ell = 1$ case. Consider a height ℓ holomorphic curve each story of which is u . Then, gluing these stories together and using the fact that changing the w_i gives bordant moduli spaces we see that $\#\mathcal{N}_J = \pm 1$ for general ℓ , z_0 , and $\{w_1, \dots, w_\ell\}$.

Since the $\bar{\partial}$ -operator is transversal for \mathcal{N}_J , stabilization invariance now follows from the gluing result A.2 and the remark at the beginning of this proof, just as described in Proposition 12.5. The present proposition is also immediate.

Appendix B Cross-references with [21]

For the reader's convenience I include a table indicating the correspondence between results in this paper and those in [21]. (The correspondence is also indicated in the text.) The third column indicates how similar the statements in the two places are, and the fourth how similar the proofs are. (Both are on a scale of 0, 1, 2 or ∞ , where ∞ indicates that I refer to [21] for the proof.) The comparison is clearly rather subjective.

Result of this paper	Result of [21]	Similarity of statement	Similarity of proof
Lemma 2.1 and Lemma 2.2	Proposition 2.15 and Lemma 2.19	2	1
Lemma 2.3	Lemma 2.19	2	∞
Proposition 3.9 and Lemma 3.10	Proposition 3.9	2	2
Corollary 4.5	Lemma 2.8	1	0
Lemma 4.11	Proposition 7.5 of [20]	2	∞
Corollary 4.12	Theorem 4.9 and Lemma 2.8	2	0
Lemma 5.3	Lemma 4.12	1	2
Lemma 5.4	Lemma 4.13	2	0
Lemma 5.5	Lemma 4.14	2	0
Proposition 5.6	Lemma 5.8 and Proposition 7.2	2	∞
Proposition 6.3	Proposition 3.10	2	1
Corollary 7.2	Theorem 3.18	1	0
Lemma 8.1 and Lemma 8.2	Theorem 4.15	2	1
Lemma 8.4	Proposition 4.18	2	2
Proposition 8.6	Proposition 4.17	2	1
Proposition 9.1	Theorem 7.3 and Theorem 6.1	2	1
Construct 10.1	Theorem 8.12	2	1
Lemma 10.2 and Lemma 10.3	Proposition 8.3	2	1
Lemma 10.6	Proposition 8.4	2	1
Lemma 10.8	Proposition 8.5	2	2
Lemma 10.10	Lemma 8.7	2	2
Lemma 10.15	Lemma 8.9	2	0
Lemma 10.16	Lemma 8.10	2	0
Proposition 10.17	Lemma 8.11	2	∞
Lemma 10.18	Theorem 8.12	1	1
Lemma 10.19	Proposition 8.13	2	1
Lemma 10.20	Proposition 8.14	2	1
Lemma 10.28	Proposition 8.15	2	2
Proposition 10.29	Theorem 8.16	2	1
Lemma 11.1	Lemma 9.1 and Lemma 9.4	1	2
Proposition 11.2	Theorem 9.5	2	1

Result of this paper	Result of [21]	Similarity of statement	Similarity of proof
Proposition 11.4	Proposition 9.8	1	0
Proposition 11.3	Lemma 9.7	2	2
Sublemma 11.5	Lemma 9.3	1	1
Lemma 11.9 and Corollary 11.10	Lemma 9.6	2	2
Proposition 12.5	Theorem 10.1	1	0

References

- [1] **F Bourgeois**, *A Morse–Bott approach to contact homology*, from: “Symplectic and contact topology: interactions and perspectives (Toronto, ON/Montreal, QC, 2001)”, Fields Inst. Commun. 35, Amer. Math. Soc., Providence, RI (2003) 55–77 MR1969267
- [2] **F Bourgeois**, **Y Eliashberg**, **H Hofer**, **K Wysocki**, **E Zehnder**, *Compactness results in symplectic field theory*, *Geom. Topol.* 7 (2003) 799–888 MR2026549
- [3] **F Bourgeois**, **K Mohnke**, *Coherent orientations in symplectic field theory*, *Math. Z.* 248 (2004) 123–146 MR2092725
- [4] **S K Donaldson**, **P B Kronheimer**, *The geometry of four-manifolds*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York (1990) MR1079726
- [5] **T Ekholm**, **J Etnyre**, **M Sullivan**, *Orientations in Legendrian contact homology and exact Lagrangian immersions*, *Internat. J. Math.* 16 (2005) 453–532 MR2141318
- [6] **Y Eliashberg**, **A Givental**, **H Hofer**, *Introduction to symplectic field theory*, *Geom. Funct. Anal.* (2000) 560–673 MR1826267
- [7] **A Floer**, *The unregularized gradient flow of the symplectic action*, *Comm. Pure Appl. Math.* 41 (1988) 775–813 MR948771
- [8] **A Floer**, *Monopoles on asymptotically flat manifolds*, from: “The Floer memorial volume”, *Progr. Math.* 133, Birkhäuser, Basel (1995) 3–41 MR1362821
- [9] **R E Gompf**, **A I Stipsicz**, *4-manifolds and Kirby calculus*, Graduate Studies in Mathematics 20, American Mathematical Society, Providence, RI (1999) MR1707327
- [10] **H Hofer**, **V Lizan**, **J-C Sikorav**, *On genericity for holomorphic curves in four-dimensional almost-complex manifolds*, *J. Geom. Anal.* 7 (1997) 149–159 MR1630789
- [11] **E-N Ionel**, **T H Parker**, *The Gromov invariants of Ruan-Tian and Taubes*, *Math. Res. Lett.* 4 (1997) 521–532 MR1470424
- [12] **P Kronheimer**, **C Manolescu**, *Periodic Floer pro-spectra from the Seiberg–Witten equations* arXiv:math.GT/0203243

- [13] **C-CM Liu**, *Moduli of J -Holomorphic Curves with Lagrangian Boundary Conditions and Open Gromov–Witten Invariants for an S^2 -Equivariant Pair* arXiv:math.SG/0210257
- [14] **D McDuff, D Salamon**, *J -holomorphic curves and quantum cohomology*, University Lecture Series 6, American Mathematical Society, Providence, RI (1994) MR1286255
- [15] **M J Micallef, B White**, *The structure of branch points in minimal surfaces and in pseudoholomorphic curves*, Ann. of Math. (2) 141 (1995) 35–85 MR1314031
- [16] **Y-G Oh**, *Fredholm theory of holomorphic discs under the perturbation of boundary conditions*, Math. Z. 222 (1996) 505–520 MR1400206
- [17] **P Ozsváth, Z Szabó**, *Holomorphic triangles and invariants for smooth four-manifolds* arXiv:math.SG/0110169
- [18] **P Ozsváth, Z Szabó**, *Holomorphic disks and genus bounds*, Geom. Topol. 8 (2004) 311–334 MR2023281
- [19] **P Ozsváth, Z Szabó**, *Holomorphic disks and knot invariants*, Adv. Math. 186 (2004) 58–116 MR2065507
- [20] **P Ozsváth, Z Szabó**, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. (2) 159 (2004) 1159–1245 MR2113020
- [21] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) 159 (2004) 1027–1158 MR2113019
- [22] **P Ozsváth, Z Szabó**, *Holomorphic triangle invariants and the topology of symplectic four-manifolds*, Duke Math. J. 121 (2004) 1–34 MR2031164
- [23] **P Ozsváth, Z Szabó**, *Heegaard Floer homology and contact structures*, Duke Math. J. 129 (2005) 39–61 MR2153455
- [24] **T Perutz**, *A remark on Kähler forms on symmetric products of Riemann surfaces* arXiv:math.SG/0501547
- [25] **J Rasmussen**, *Floer Homology and Knot Complements*, PhD thesis, Harvard University (2003) arXiv:math.GT/0306378
- [26] **CH Taubes**, *A compendium of pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$* , Geom. Topol. 6 (2002) 657–814 MR1943381
- [27] **V Turaev**, *Torsion invariants of Spin^c -structures on 3-manifolds*, Math. Res. Lett. 4 (1997) 679–695 MR1484699
- [28] **M Usher**, *The Gromov invariant and the Donaldson-Smith standard surface count*, Geom. Topol. 8 (2004) 565–610 MR2057774

Department of Mathematics, Stanford University, Stanford, CA 94305-2125, USA

lipshitz@math.stanford.edu

Proposed: Peter Ozsváth
 Seconded: John Morgan, Ronald Fintushel

Received: 14 May 2005
 Revised: 9 October 2005