



Geometry & Topology

Volume 30 (2026)

Graph cobordisms and Heegaard Floer homology

IAN ZEMKE

Graph cobordisms and Heegaard Floer homology

IAN ZEMKE

We construct a graph TQFT for the minus flavor of Heegaard Floer homology. Our graph TQFT extends Ozsváth and Szabó’s TQFT for closed and connected 3-manifolds, and assigns maps to cobordisms with disconnected ends. As an application, we give an explicit formula for the chain homotopy type of the π_1 -action on Heegaard Floer homology. We show that on homology the π_1 -action is trivial on the plus, minus and infinity flavors, but give examples where it is nontrivial on the hat flavor.

1. Introduction	389
2. Categorical preliminaries	396
3. Heegaard Floer homology	398
4. Relative homology action	415
5. Free-stabilization maps	427
6. Graph action map	449
7. 1- and 3-handles	464
8. 2-handles	475
9. Constructing the graph TQFT, I	483
10. Constructing the graph TQFT, II: adding punctures	488
11. The composition law	490
12. Path cobordisms and the normalization axiom	491
13. Basepoint-moving maps and the normalization axiom	492
Acknowledgements	526
References	527

1 Introduction

Heegaard Floer homology, introduced by Ozsváth and Szabó [2004b], associates functorial invariants to 3- and 4-manifolds. To an oriented, closed and connected 3-manifold Y , with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, they construct groups

$$\text{HF}^-(Y, \mathfrak{s}), \quad \text{HF}^\infty(Y, \mathfrak{s}), \quad \text{HF}^+(Y, \mathfrak{s}) \quad \text{and} \quad \widehat{\text{HF}}(Y, \mathfrak{s}),$$

which are modules over the ring $\mathbb{F}_2[U]$. These modules fit into the framework of a (3+1)-dimensional TQFT:

MSC2010: 57M27, 57R58.

© 2026 The Author, under license to MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

Theorem 1.1 [Ozsváth and Szabó 2006] *Suppose W is an oriented, compact 4-manifold with boundary $\partial W = -Y_0 \sqcup Y_1$, and W , Y_0 and Y_1 are all nonempty and connected. If $\mathfrak{s} \in \text{Spin}^c(W)$ and $\circ \in \{+, -, \infty, \wedge\}$, then there is a functorial map*

$$F_{W,\mathfrak{s}}: \text{HF}^\circ(Y_0, \mathfrak{s}|_{Y_0}) \rightarrow \text{HF}^\circ(Y_1, \mathfrak{s}|_{Y_1}),$$

which depends on a choice of path in W , connecting basepoints in Y_0 and Y_1 .

The dependence on a choice of path was not explicitly stated in Ozsváth and Szabó's original work, though their construction implicitly depends on a choice of path. Indeed, we will show that the maps appearing in Theorem 1.1 are not invariants without a choice of path; see Corollary E, below.

The goal of this paper is two-fold. Firstly, we determine precisely the dependence of the maps in Theorem 1.1 on a choice of path. Secondly, and more broadly, we provide a TQFT framework for Heegaard Floer homology which incorporates the basepoints in a natural way and extends Ozsváth and Szabó's framework to disconnected 3- and 4-manifolds.

Throughout the paper we work over the field $\mathbb{F}_2 := \mathbb{Z}/2\mathbb{Z}$.

1.1 The graph TQFT

Ozsváth and Szabó [2008], extending their [2004b] original construction for 3-manifolds, described Heegaard Floer complexes $\text{CF}^-(Y, \mathbf{w}, \mathfrak{s})$ for 3-manifolds with collections of basepoints $\mathbf{w} \subseteq Y$. For the purposes of the introduction, the complex $\text{CF}^-(Y, \mathbf{w}, \mathfrak{s})$ is a free, finitely generated chain complex over the polynomial ring $\mathbb{F}_2[U]$. Later, we will consider an algebraic generalization over a more general polynomial ring.

In this paper, we define a notion of cobordism between manifolds with collections of basepoints. A *ribbon graph cobordism* (W, Γ) from (Y_0, \mathbf{w}_0) to (Y_1, \mathbf{w}_1) is a cobordism W from Y_0 to Y_1 which contains an embedded graph Γ , whose intersection with Y_i is \mathbf{w}_i . Furthermore, Γ is decorated with a *formal ribbon structure*, i.e., a choice of cyclic ordering of the edges adjacent to each vertex of Γ .

Our present paper centers on proving the following:

Theorem A *Suppose $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a ribbon graph cobordism and $\mathfrak{s} \in \text{Spin}^c(W)$. The construction in this paper yields two chain maps,*

$$F_{W,\Gamma,\mathfrak{s}}^A, F_{W,\Gamma,\mathfrak{s}}^B: \text{CF}^-(Y_0, \mathbf{w}_0, \mathfrak{s}|_{Y_0}) \rightarrow \text{CF}^-(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}),$$

which are diffeomorphism invariants of (W, Γ) , up to $\mathbb{F}_2[U]$ -equivariant chain homotopy.

Theorem A also applies to CF^∞ , CF^+ and $\widehat{\text{CF}}$, since they can be recovered via tensor products with CF^- ; see Section 3.2.

The type- A and type- B cobordism maps satisfy the symmetry

$$(1-1) \quad F_{W,\Gamma,\mathfrak{s}}^A \simeq F_{W,\bar{\Gamma},\mathfrak{s}}^B,$$

where $\bar{\Gamma}$ denotes Γ with cyclic orders reversed [Hendricks et al. 2018, Lemma 5.9]. Despite this symmetry, it is natural to describe the A and B versions separately.

Unlike Ozsváth and Szabó’s original construction, Theorem A applies to cobordisms which are disconnected, or which have disconnected 3-manifolds appearing in their ends, as long as each component contains a basepoint.

Diffeomorphism-invariance in Theorem A amounts to the following. Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ and $(W', \Gamma'): (Y'_0, \mathbf{w}'_0) \rightarrow (Y'_1, \mathbf{w}'_1)$ are two graph cobordisms, and $D: (W, \Gamma) \rightarrow (W', \Gamma')$ is an orientation-preserving diffeomorphism. The map D restricts to give diffeomorphisms $d_0: Y_0 \rightarrow Y'_0$ and $d_1: Y_1 \rightarrow Y'_1$. Let $\mathfrak{s} \in \text{Spin}^c(W)$ and $\mathfrak{s}' := D_*(\mathfrak{s})$. The following diagram commutes up to chain homotopy:

$$(1-2) \quad \begin{array}{ccc} \text{CF}^-(Y_0, \mathbf{w}_0, \mathfrak{s}|_{Y_0}) & \xrightarrow{(d_0)_*} & \text{CF}^-(Y'_0, \mathbf{w}'_0, \mathfrak{s}'|_{Y'_0}) \\ \downarrow F_{W, \Gamma, \mathfrak{s}}^A & & \downarrow F_{W', \Gamma', \mathfrak{s}'}^A \\ \text{CF}^-(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) & \xrightarrow{(d_1)_*} & \text{CF}^-(Y'_1, \mathbf{w}'_1, \mathfrak{s}'|_{Y'_1}) \end{array}$$

Concerning graph cobordisms where the graph consists of a collection of paths from Y_0 to Y_1 , we prove the following:

Theorem B *Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a graph cobordism, and Γ consists of a collection of paths, each connecting \mathbf{w}_0 to \mathbf{w}_1 .*

(1) *The A and B versions of the maps coincide:*

$$F_{W, \Gamma, \mathfrak{s}}^A \simeq F_{W, \Gamma, \mathfrak{s}}^B.$$

(2) *If $\phi: (Y, \mathbf{w}) \rightarrow (Y, \mathbf{w})$ is an orientation-preserving diffeomorphism, let $W(\phi)$ denote the **mapping cylinder** (i.e., $[0, 1] \times Y$, with $\{0\} \times Y$ identified with Y via id_Y and $\{1\} \times Y$ identified with Y via ϕ). Then*

$$F_{W(\phi), [0, 1] \times \mathbf{w}, \mathfrak{s}}^A \simeq F_{W(\phi), [0, 1] \times \mathbf{w}, \mathfrak{s}}^B \simeq (\phi_*: \text{CF}^-(Y, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(Y, \mathbf{w}, \phi_*\mathfrak{s})).$$

(3) *Suppose $(W, \gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a cobordism such that W, Y_0 and Y_1 are nonempty and connected, and γ is a path from \mathbf{w}_0 to \mathbf{w}_1 . Then $F_{W, \gamma, \mathfrak{s}}^A \simeq F_{W, \gamma, \mathfrak{s}}^B$, and both maps coincide with the map defined by Ozsváth and Szabó.*

Analogous to the cobordism maps defined by Ozsváth and Szabó, the graph cobordism maps satisfy a Spin^c composition law:

Theorem C *Suppose that (W, Γ) is ribbon graph cobordism which decomposes as a composition $(W, \Gamma) = (W_2, \Gamma_2) \cup (W_1, \Gamma_1)$. If \mathfrak{s}_1 and \mathfrak{s}_2 are Spin^c structures on W_1 and W_2 , then*

$$F_{W_2, \Gamma_2, \mathfrak{s}_2}^A \circ F_{W_1, \Gamma_1, \mathfrak{s}_1}^A \simeq \sum_{\substack{\mathfrak{s} \in \text{Spin}^c(W) \\ \mathfrak{s}|_{W_2} = \mathfrak{s}_2 \\ \mathfrak{s}|_{W_1} = \mathfrak{s}_1}} F_{W, \Gamma, \mathfrak{s}}^A.$$

The same relation holds for the type-B graph cobordism maps.




Figure 1: The broken path graph cobordism for the endomorphism Φ_w . The underlying 4-manifold is $[0, 1] \times Y$. When w is not the only basepoint, there are additional, unbroken strands from the bottom to the top.

1.2 Moving basepoints and the π_1 -action

If $w \in Y$, there is a fibration

$$(1-3) \quad \text{Diff}(Y, w) \rightarrow \text{Diff}(Y) \xrightarrow{\text{ev}_w} Y.$$

The long exact sequence of homotopy groups for the fibration in (1-3) gives a homomorphism

$$\pi_1(Y, w) \rightarrow \text{MCG}(Y, w),$$

where $\text{MCG}(Y, w)$ denotes the *based mapping class group* of orientation-preserving diffeomorphisms modulo smooth isotopies which fix the point w .

If $\gamma \in \pi_1(Y, w)$, we write

$$\gamma_* : \text{CF}^-(Y, w, \mathfrak{s}) \rightarrow \text{CF}^-(Y, w, \mathfrak{s})$$

for the induced diffeomorphism map.

If $w \in \mathbf{w}$ is a chosen of basepoint, there is a +1-graded endomorphism

$$\Phi_w : \text{CF}^-(Y, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(Y, \mathbf{w}, \mathfrak{s}).$$

One interpretation of the map Φ_w is as the map for the *broken path cobordism*, shown in Figure 1. Alternatively, Φ_w can be interpreted as the map for a count of holomorphic disks on a Heegaard diagram; see (13-30).

An analog of the map Φ_w for link Floer homology was discovered by Sarkar [2015] in the context of basepoint-moving diffeomorphisms on link Floer homology.

Theorem D Suppose (Y, \mathbf{w}) is a multipointed 3-manifold, $w \in \mathbf{w}$ and $\gamma \in \pi_1(Y, w)$. Then the diffeomorphism map γ_* on $\text{CF}^-(Y, \mathbf{w}, \mathfrak{s})$ satisfies

$$\gamma_* \simeq \text{id} + \Phi_w \circ A_\gamma,$$

where Φ_w is the broken path cobordism map and A_γ denotes the action of the class $[\gamma] \in H_1(Y, \mathbb{Z})/\text{Tors}$.

Using Theorem D, we exhibit 3-manifolds Y where the induced map

$$\gamma_*: \widehat{\text{HF}}(Y, w, \mathfrak{s}) \rightarrow \widehat{\text{HF}}(Y, w, \mathfrak{s})$$

is not the identity:

Corollary E *Let Y be a 3-manifold and $\mathfrak{s} \in \text{Spin}^c(Y)$.*

(1) *If \mathfrak{s} is torsion and there is an $x \in \text{HF}^+(Y, w, \mathfrak{s})$ such that*

$$U \cdot x = 0 \quad \text{and} \quad x \notin U \cdot \text{HF}^+(Y, w, \mathfrak{s}),$$

then $\pi_1(Y \# S^1 \times S^2, w)$ acts nontrivially on

$$\widehat{\text{HF}}(Y \# S^1 \times S^2, w).$$

(2) *Suppose $[\gamma] \in H_1(Y; \mathbb{Z})$ is a class whose action on $\text{HF}^+(Y, w, \mathfrak{s})$ does not vanish. If $|\mathbf{w}| \geq 2$, then the diffeomorphism map γ_* acts nontrivially on the $\mathbb{F}_2[U]$ -module*

$$\text{HF}^-(Y, \mathbf{w}, \mathfrak{s}).$$

Note that $Y = \Sigma(2, 3, 7)$ satisfies the conditions of part (1) of Corollary E. The manifold $Y = S^1 \times S^2$ satisfies part (2).

When we restrict Theorem D to 3-manifolds with a single basepoint w , the formula defining Φ_w has the following algebraic interpretation. View $\text{CF}^-(Y, w, \mathfrak{s})$ as a free module over $\mathbb{F}_2[U]$ with basis equal to the set of intersection points. The differential can be written as a matrix over this basis, with entries in $\mathbb{F}_2[U]$. The endomorphism Φ_w is obtained by differentiating each entry of this matrix with respect to U . Using this basis, we can also define a derivative map d/dU as an endomorphism of $\text{CF}^-(Y, w, \mathfrak{s})$, which is not U -equivariant. The Leibniz rule implies

$$(1-4) \quad \Phi_w = \partial \circ \frac{d}{dU} + \frac{d}{dU} \circ \partial.$$

As a consequence, we prove the following folklore result known to experts:

Corollary F *If (Y, w) is a singly based 3-manifold and $\gamma \in \pi_1(Y, w)$, then the induced map γ_* is equal to the identity on the homology groups $\text{HF}^-(Y, w, \mathfrak{s})$, $\text{HF}^\infty(Y, w, \mathfrak{s})$ and $\text{HF}^+(Y, w, \mathfrak{s})$, but not necessarily $\widehat{\text{HF}}(Y, w, \mathfrak{s})$. Consequently, if (W, γ) is path cobordism between two singly based 3-manifolds (Y_0, w_0) and (Y_1, w_1) , then the cobordism map*

$$F_{W, \gamma, \mathfrak{s}}: \text{HF}^\circ(Y_0, w_0, \mathfrak{s}|_{Y_0}) \rightarrow \text{HF}^\circ(Y_1, w_1, \mathfrak{s}|_{Y_1})$$

is independent of γ if $\circ \in \{-, \infty, +\}$.

The chain homotopy $H = d/dU$ appearing in (1-4) is not $\mathbb{F}_2[U]$ -equivariant, and hence does not induce a chain homotopy on $\widehat{\text{CF}}$.

In another direction, if λ is a path between two basepoints $w_1, w_2 \in \mathbf{w} \subseteq Y$, there is a diffeomorphism $\text{Sw}_\lambda: (Y, \mathbf{w}) \rightarrow (Y, \mathbf{w})$ which swaps the two basepoints w_1 and w_2 along the path λ (this is well defined,




Figure 2: Graph cobordisms for the free-stabilization maps. The underlying 4-manifold is $[0, 1] \times Y$.

up to isotopy, since Y is 3-dimensional). Using the graph TQFT, we prove in Proposition 13.24 that

$$S_{w\lambda} \simeq \Phi_{w_1} A_\lambda + A_\lambda \Phi_{w_2} \simeq A_\lambda \Phi_{w_1} + \Phi_{w_2} A_\lambda.$$

1.3 Outline of the construction of the graph TQFT

We now describe the main ingredients of our graph TQFT. The first ingredients are the 4-dimensional handle attachment maps, similar to the maps defined by Ozsváth and Szabó [2006], with the exception that we define maps for 0-handles and 4-handles, and our construction of 1-handle maps and 3-handle maps is more flexible than their construction. There are two additional, novel ingredients of the graph TQFT: the *free-stabilization maps* and the *relative homology maps*.

If $w \notin \mathbf{w}$, we define two free-stabilization maps

$$S_w^+ : CF^-(Y, \mathbf{w}, \mathfrak{s}) \rightarrow CF^-(Y, \mathbf{w} \cup \{w\}, \mathfrak{s}) \quad \text{and} \quad S_w^- : CF^-(Y, \mathbf{w} \cup \{w\}, \mathfrak{s}) \rightarrow CF^-(Y, \mathbf{w}, \mathfrak{s}).$$

The maps S_w^+ and S_w^- are induced by the graph cobordisms shown in Figure 2.

Another new ingredient of our graph TQFT is the *relative homology action*. To a path λ between two basepoints $w_1, w_2 \in \mathbf{w}$ we construct an endomorphism

$$A_\lambda : CF^-(Y, \mathbf{w}, \mathfrak{s}) \rightarrow CF^-(Y, \mathbf{w}, \mathfrak{s}).$$

The definition of the map A_λ is asymmetrical with respect to the α and β curves. By switching the roles of α and β in the construction, we obtain another map B_λ , with the same domain and codomain, which is also a chain map.

The maps A_λ and B_λ turn out to be the graph cobordism maps for H -shaped graphs in $[0, 1] \times Y$ for different ribbon structures, though we will not make use of this fact in this paper; see [Zemke 2019, Lemma 14.11].

As an intermediate step towards defining the graph cobordism maps, we construct maps for graphs embedded in a fixed 3-manifold Y . We say $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is an embedded *flow-graph* in Y if $\Gamma \subseteq Y$ is an embedded ribbon graph, each point of \mathbf{w}_0 and \mathbf{w}_1 has valence 1 in Γ , and Γ has no valence 0 vertices.

As a key step towards the full construction of the graph TQFT, we construct $\mathbb{F}_2[U]$ -equivariant chain maps

$$\mathfrak{A}_{\mathcal{G}}, \mathfrak{B}_{\mathcal{G}} : CF^-(Y, \mathbf{w}_0, \mathfrak{s}) \rightarrow CF^-(Y, \mathbf{w}_1, \mathfrak{s}),$$

which we call the *type-A* and *type-B graph action maps*. The two maps $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ are related by the same symmetry as the cobordism maps in (1-1).

If $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is a ribbon flow-graph in Y , then, by definition,

$$(1-5) \quad \mathfrak{A}_{\mathcal{G}} \simeq F_{[0,1] \times Y, \Gamma', s}^A \quad \text{and} \quad \mathfrak{B}_{\mathcal{G}} \simeq F_{[0,1] \times Y, \Gamma', s}^B,$$

where $\Gamma' \subseteq [0, 1] \times Y$ is a ribbon graph which projects to Γ and is such that $\Gamma' \cap (\{0\} \times Y) = \mathbf{w}_0$ and $\Gamma' \cap (\{1\} \times Y) = \mathbf{w}_1$.

1.4 Further developments

This paper focuses on defining the graph cobordism maps, $F_{W, \Gamma, s}^A$ and $F_{W, \Gamma, s}^B$, and proving invariance. In subsequent papers, further properties and applications have been explored, which we briefly summarize.

Hendricks, Manolescu and the author [Hendricks et al. 2018, Proposition 5.2] showed that pair-of-pants graph cobordisms containing a trivalent graph induce chain homotopy equivalences between $\text{CF}^-(Y_1 \# Y_2)$ and $\text{CF}^-(Y_1) \otimes \text{CF}^-(Y_2)$, giving a cobordism perspective on Ozsváth and Szabó’s connected sum formula [2004a, Theorem 1.5]. We used this fact to prove a connected sum formula for involutive Heegaard Floer homology [Hendricks et al. 2018, Theorem 1.1].

In another subsequent paper, we prove that the graph cobordism maps for the *trace cobordism*

$$([0, 1] \times Y, [0, 1] \times \{w\}): (-Y, w) \sqcup (Y, w) \rightarrow \emptyset$$

are chain homotopic to the canonical trace pairing between $\text{CF}^-(Y)$ and $\text{CF}^-(-Y)$ [Zemke 2021a, Theorem 1.6]. A similar result holds for the *cotrace cobordism*, obtained by turning around the trace cobordism. As an application, the author [ibid., Theorem 1.1] computed the Heegaard Floer mixed invariants of mapping tori in terms of Lefschetz numbers on $\text{HF}_{\text{red}}^+(Y)$.

In another direction, the author described a TQFT for link Floer homology [Zemke 2019]. It turns out that the link Floer TQFT recovers the graph TQFT in a natural way; see [ibid., Theorem C] for a precise statement. One interpretation of the dependence on cyclic orders from this paper can be seen as an artifact of the relation with graphs embedded on surfaces.

The graph cobordism maps are invariant under a more general equivalence than ambient diffeomorphism of the 4-manifold, as stated in (1-2). The cobordism maps are also invariant under homotopies of the graph inside W which are only required to be smooth on each edge, and hence may not extend to a smooth isotopy of W ; see Definition 2.4. Using the aforementioned connection with link cobordisms, it turns out that the graph cobordisms are invariant under modifications of the graph which preserve the *ribbon equivalence class* of the ribbon graph; see [ibid., Corollary D] for a precise statement.

1.5 Further commentary

It is worthwhile to compare the cobordism maps in Heegaard Floer homology to the cobordism maps in Kronheimer and Mrowka’s construction [2007] of monopole Floer homology. It follows from work of Kutluhan, Lee and Taubes [Kutluhan et al. 2020a; 2020b; 2020c; 2020d; 2020e], and, independently,

Colin, Ghiggini and Honda [Colin et al. 2025; 2024b; 2024c; 2024a], that certain versions of Heegaard Floer homology, monopole Floer homology, and embedded contact homology are isomorphic. A proof that these isomorphisms extend to the level of 4-manifold invariants has not yet appeared in the literature.

In monopole Floer homology, as well as embedded contact homology, the chain complexes are defined independently of a basepoint, but the action of U requires a choice of basepoint, and a path can be used to construct a chain homotopy between the two U maps; see [Hutchings and Taubes 2009, Section 2.5].

In a different direction, we note that an early version of Theorem D appeared in [Zemke 2021b], for \widehat{HF} . The proof in [loc. cit.] used Juhász's sutured TQFT [2006; 2016] as well as the contact gluing map of Honda, Kazez and Matić [Honda et al. 2008], and hence had a different flavor than the one we explore in this paper. It is an interesting question whether the graph cobordism maps described in this paper have an interpretation in terms of the contact gluing map, perhaps using a limiting construction as in [Etnyre et al. 2017; Golla 2015] or using a minus version of sutured Floer homology described by Alishahi and Eftekhary [2015].

Outline of the paper

In Section 2, we define the category of ribbon graph cobordisms, and some related cobordism categories which appear in our paper. In Section 3, we describe some background material on Heegaard Floer homology, and prove several results about admissibility for multipointed Heegaard Floer diagrams. In Section 4, we construct the relative homology action. In Section 5, we construct the free-stabilization maps, for adding or removing a basepoint. In Section 6, we combine the free-stabilization maps with the relative homology action to construct a restricted version of our TQFT, which we call the graph action map. In Sections 7 and 8, we describe maps for 4-dimensional handles. In Sections 9 and 10, we define the graph cobordism maps and prove invariance (Theorem A). In Section 11, we prove the composition law (Theorem C). In Section 12, we give a summary of the proof of the normalization axiom (Theorem B). In Section 13, we perform several technical, holomorphic curve arguments to finish the proof of the normalization axiom, and also give several formulas for basepoint-moving diffeomorphism maps (in particular, Theorem D and Corollaries E and F).

2 Categorical preliminaries

In this section, we define a category of graph cobordisms GrCob , and a related category $\text{FlGr}(Y)$ of immersed flow-graphs in a fixed 3-manifold Y .

2.1 The graph cobordism category

The objects of GrCob are the following:

Definition 2.1 A *multipointed 3-manifold* is a pair (Y, \mathbf{w}) consisting of a closed, oriented 3-manifold Y (not necessarily connected), together with a finite collection of basepoints $\mathbf{w} \subseteq Y$, such that each component of Y contains at least one basepoint.

Morphisms in GrCob are the following:

Definition 2.2 A *ribbon graph cobordism* from (Y_0, \mathbf{w}_0) to (Y_1, \mathbf{w}_1) is a pair (W, Γ) satisfying the following:

- (1) W is a cobordism from Y_0 to Y_1 .
- (2) Γ is an embedded graph in W such that $\Gamma \cap Y_i = \mathbf{w}_i$. Furthermore, each point of \mathbf{w}_i has valence 1 in Γ .
- (3) Γ has finitely many edges and vertices, and no vertices of valence 0.
- (4) The embedding of Γ is smooth on each edge.
- (5) Γ is decorated with a *formal ribbon structure*, i.e., a formal choice of cyclic ordering of the edges adjacent to each vertex.

Remark 2.3 In GrCob , the embedding of a graph need not respect the formal ribbon structure (e.g., the embedding of Γ near a vertex need not map the adjacent edges into a 2-plane centered at the vertex in a way which respects the cyclic ordering).

The identity graph cobordism from (Y, \mathbf{w}) to itself is $([0, 1] \times Y, [0, 1] \times \mathbf{w})$.

As always in cobordism categories, for GrCob to form an honest category, we must also include parametrizing diffeomorphisms of the boundary as data of a morphism, and we must also quotient by diffeomorphisms which respect these parametrizations. (Otherwise, there is no identity morphism.) To simplify exposition, we suppress this fact.

Some care is required when defining equivalences of graph cobordisms, since vertices of valence greater than 1 are inherently nonsmooth. In the morphism sets of GrCob , we quotient by diffeomorphisms of the ambient 4-manifold, as well as the following notion of graph isotopy:

Definition 2.4 Suppose that W is a 4-manifold and Γ is an abstract graph. We say that a continuous map

$$i: [0, 1] \times \Gamma \rightarrow W$$

is a *smooth isotopy of embedded graphs* if there is a finite subset $\mathbf{v} \subseteq \Gamma$, consisting of valence 1 vertices, such that the following are satisfied:

- (1) For each t , the map $i|_{\{t\} \times \Gamma}$ is a topological embedding.
- (2) $i^{-1}(\partial W) = [0, 1] \times \mathbf{v}$, and i is constant on \mathbf{v} .
- (3) If $e \subseteq \Gamma$ is an edge, then the restriction $i|_{[0, 1] \times e}$ is smooth.

2.2 The flow-graph category

There is a related category of interest to us, which we call the *flow-graph* category of Y . The objects of $\text{FlGr}(Y)$ are collections of basepoints $\mathbf{w} \subseteq Y$ such that each component of Y contains at least one basepoint. The morphisms in FlGr are the following:

Definition 2.5 Suppose Y is a closed, oriented 3-manifold and \mathbf{w}_0 and \mathbf{w}_1 are two collections of basepoints in Y . We say a tuple $\mathcal{G} = (\Gamma, i, \mathbf{v}_0, \mathbf{v}_1)$ is an immersed *ribbon flow-graph* from \mathbf{w}_0 to \mathbf{w}_1 if the following are satisfied:

- (1) Γ is an abstract ribbon graph. Furthermore, $i: \Gamma \rightarrow Y$ is an immersion.
- (2) Each vertex of Γ has valence at least 1.
- (3) \mathbf{v}_0 and \mathbf{v}_1 disjoint collections of valence 1 vertices in Γ . Furthermore, i maps \mathbf{v}_0 bijectively to \mathbf{w}_0 and i maps \mathbf{v}_1 bijectively to \mathbf{w}_1 .

We identify two ribbon flow-graphs $(\Gamma, i, \mathbf{v}_0, \mathbf{v}_1)$ and $(\Gamma', i', \mathbf{v}'_0, \mathbf{v}'_1)$ if there is a homeomorphism $h: \Gamma \rightarrow \Gamma'$ such that $i' = h \circ i$.

If $\mathcal{G} = (\Gamma, i, \mathbf{v}_0, \mathbf{v}_1)$ is a ribbon flow-graph, we say that \mathbf{v}_0 and \mathbf{v}_1 are the *boundary* vertices of Γ . We say all other vertices are *interior* vertices. We will usually omit the immersion i from the notation, and write simply $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ for a flow-graph. For convenience, we will usually assume that our flow-graphs are embedded.

We use the following notion of equivalence of immersed flow-graphs:

Definition 2.6 Suppose that Γ is an abstract graph with pairwise disjoint sets of valence 1 vertices \mathbf{v}_0 and \mathbf{v}_1 . If $\mathbf{w}_0, \mathbf{w}_1$ are two collections of basepoints in Y , we say a continuous map

$$h: [0, 1] \times \Gamma \rightarrow Y$$

is a *smooth homotopy of immersed flow-graphs* from \mathbf{w}_0 to \mathbf{w}_1 if the following are satisfied:

- (1) For each t , the map $h|_{\{t\} \times \Gamma}$ is locally an embedding.
- (2) $h(\{t\} \times \mathbf{v}_0) = \mathbf{w}_0$ and $h(\{t\} \times \mathbf{v}_1) = \mathbf{w}_1$ for all t .
- (3) If $e \subseteq \Gamma$ is an edge, then $h|_{[0,1] \times e}$ is smooth.

The identity flow-graph from \mathbf{w} to itself consists of the pair (Γ, i) where Γ is the disjoint union of $|\mathbf{w}|$ copies of the interval $[0, 1]$ and i is a small perturbation of the map which sends a copy of $[0, 1]$ to the corresponding basepoint in \mathbf{w} .

There is a functor from $\text{FlGr}(Y)$ to GrCob which sends an immersed flow-graph $(\Gamma, i, \mathbf{v}_0, \mathbf{v}_1)$ to a graph cobordism $([0, 1] \times Y, \Gamma'): (Y, i(\mathbf{v}_0)) \rightarrow (Y, i(\mathbf{v}_1))$, where Γ' is an embedded graph in $[0, 1] \times Y$ which projects to Γ (up to ambient isotopy).

3 Heegaard Floer homology

In this section, we recall Ozsváth and Szabó's construction of the Heegaard Floer complexes for multi-pointed 3-manifolds.

3.1 Multipointed Heegaard diagrams and the Heegaard Floer complexes

Definition 3.1 Suppose (Y, \mathbf{w}) is a multipointed 3-manifold. A *multipointed Heegaard diagram* $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ for (Y, \mathbf{w}) is a tuple satisfying the following:

- (1) Σ is a closed, oriented surface, embedded in Y , such that $\mathbf{w} \subseteq \Sigma \setminus (\alpha \cup \beta)$. Furthermore, Σ splits Y into two handlebodies, U_α and U_β , oriented so that $\Sigma = \partial U_\alpha = -\partial U_\beta$.
- (2) $\alpha = \{\alpha_1, \dots, \alpha_n\}$ is a collection of $n = g(\Sigma) + |\mathbf{w}| - 1$ pairwise disjoint, simple, closed curves on Σ , bounding pairwise disjoint compressing disks in U_α . Each component of $\Sigma \setminus \alpha$ is planar and contains a single basepoint.
- (3) $\beta = \{\beta_1, \dots, \beta_n\}$ is a collection of pairwise disjoint, simple, closed curves on Σ bounding pairwise disjoint compressing disks in U_β . Each component of $\Sigma \setminus \beta$ is planar and contains a single basepoint.

We will also require Heegaard diagrams to satisfy an admissibility requirement; see Section 3.7.

3.2 The Heegaard Floer chain complexes

In this section, we describe Ozsváth and Szabó’s Heegaard Floer chain complexes.

If $\mathbf{w} = \{w_1, \dots, w_n\}$ is a set of basepoints in Y , we define the free polynomial ring

$$\mathbb{F}_2[U_{\mathbf{w}}] := \mathbb{F}_2[U_{w_1}, \dots, U_{w_n}].$$

Write $\mathbb{F}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}]$ for the ring obtained by formally inverting each of the variables U_{w_i} .

If $\mathbf{k} = (k_1, \dots, k_n)$ is an n -tuple, write

$$U_{\mathbf{w}}^{\mathbf{k}} := U_{w_1}^{k_1} \dots U_{w_n}^{k_n}.$$

Suppose (Y, \mathbf{w}) is a connected, multipointed 3-manifold and $\mathfrak{s} \in \text{Spin}^c(Y)$. Pick a diagram $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ of (Y, \mathbf{w}) , and consider the two tori

$$\mathbb{T}_\alpha := \alpha_1 \times \dots \times \alpha_n \quad \text{and} \quad \mathbb{T}_\beta := \beta_1 \times \dots \times \beta_n,$$

inside the symmetric product $\text{Sym}^n(\Sigma)$, where $n = g(\Sigma) + |\mathbf{w}| - 1$.

Ozsváth and Szabó [2004b, Section 2.6] construct a map

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

As an $\mathbb{F}_2[U_{\mathbf{w}}]$ -module, $\text{CF}^-(\mathcal{H}, \mathfrak{s})$ is defined to be the free $\mathbb{F}_2[U_{\mathbf{w}}]$ -module generated by intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with

$$\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}.$$

We define

$$(3-1) \quad \text{CF}^\infty(\mathcal{H}, \mathfrak{s}) := \text{CF}^-(\mathcal{H}, \mathfrak{s}) \otimes_{\mathbb{F}_2[U_{\mathbf{w}}]} \mathbb{F}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}] \quad \text{and} \quad \text{CF}^+(\mathcal{H}, \mathfrak{s}) := \text{CF}^\infty(\mathcal{H}, \mathfrak{s}) / \text{CF}^-(\mathcal{H}, \mathfrak{s}).$$

To equip $\text{CF}^-(\mathcal{H}, \mathfrak{s})$ with a differential, we pick an auxiliary path $(J_s)_{s \in [0,1]}$ of almost complex structures on $\text{Sym}^{g+|\mathbf{w}|-1}(\Sigma)$. We write $\mathcal{M}_{J_s}(\phi)$ for the moduli space of J_s -holomorphic maps $u: [0, 1] \times \mathbb{R} \rightarrow \text{Sym}^{g+|\mathbf{w}|-1}(\Sigma)$ which represent the class ϕ . The moduli space $\mathcal{M}_{J_s}(\phi)$ has a natural action of \mathbb{R} ,

corresponding to reparametrization of the source. We write

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

For a sufficiently generic J_s , we define the differential

$$\partial_{J_s} : \text{CF}^-(\mathcal{H}, \mathfrak{s}) \rightarrow \text{CF}^-(\mathcal{H}, \mathfrak{s})$$

via the formula

$$(3-2) \quad \partial_{J_s}(\mathbf{x}) = \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} \# \widehat{\mathcal{M}}_{J_s}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y},$$

extended linearly over $\mathbb{F}_2[U_{\mathbf{w}}]$.

The endomorphism ∂_{J_s} satisfies

$$\partial_{J_s} \circ \partial_{J_s} = 0.$$

We refer the reader to [Ozsváth and Szabó 2008, Lemma 4.3] for a proof.

We write $\text{CF}_{J_s}^-(\mathcal{H}, \mathfrak{s})$ for the $\mathbb{F}_2[U_{\mathbf{w}}]$ -module $\text{CF}^-(\mathcal{H}, \mathfrak{s})$ equipped with differential ∂_{J_s} . When there is no ambiguity, we will usually drop the J_s subscripts from both the chain complexes and the differential.

If (Y, \mathbf{w}) is a disconnected, multipointed 3-manifold, then a Heegaard diagram for (Y, \mathbf{w}) is a disjoint union of Heegaard diagrams for each component. The Heegaard Floer complex for such a diagram is the tensor product over \mathbb{F}_2 of the complexes for each component.

3.3 Coloring the Heegaard Floer complexes

In this section, we describe a way of algebraically modifying the Heegaard Floer complexes to achieve functoriality of the Heegaard Floer complexes under graph cobordisms.

Definition 3.2 If X is a topological space, a *coloring* of X is a function

$$\sigma : \pi_0(X) \rightarrow \mathbf{P},$$

where $\pi_0(X)$ denotes the set of connected components of X , and \mathbf{P} is a finite set.

When X is a finite set (such as a set of basepoints), we view a coloring as a map from X to \mathbf{P} .

If $\mathbf{P} = \{p_1, \dots, p_n\}$ is a finite set, write $\mathcal{R}_{\mathbf{P}}$ for the n -variable polynomial ring

$$\mathcal{R}_{\mathbf{P}} := \mathbb{F}_2[U_{p_1}, \dots, U_{p_n}].$$

If $\sigma : \mathbf{w} \rightarrow \mathbf{P}$ is a coloring, then σ induces an action of $\mathbb{F}_2[U_{\mathbf{w}}]$ on $\mathcal{R}_{\mathbf{P}}$, defined via the formula

$$U_{\mathbf{w}} \cdot x := U_{\sigma(\mathbf{w})} \cdot x$$

for $x \in \mathcal{R}_{\mathbf{P}}$. Hence, if M is an $\mathbb{F}_2[U_{\mathbf{w}}]$ -module, we can form an $\mathcal{R}_{\mathbf{P}}$ -module M^σ as the tensor product

$$(3-3) \quad M^\sigma := M \otimes_{\mathbb{F}_2[U_{\mathbf{w}}]} \mathcal{R}_{\mathbf{P}},$$

which we think of as being obtained by formally identifying U_{w_i} with $U_{\sigma(w_i)}$.

If (Y, \mathbf{w}) is a multipointed 3-manifold and $\sigma: \mathbf{w} \rightarrow \mathbf{P}$ is a coloring, then we write $\text{CF}^-(\mathcal{H}, \mathfrak{s}, \sigma)$ for the complex

$$(3-4) \quad \text{CF}^-(\mathcal{H}, \sigma, \mathfrak{s}) := \text{CF}^-(\mathcal{H}, \mathfrak{s})^\sigma = \text{CF}^-(\mathcal{H}, \mathfrak{s}) \otimes_{\mathbb{F}_2[U_{\mathbf{w}}]} \mathcal{R}_{\mathbf{P}}.$$

We briefly highlight the behavior of the complexes under disjoint unions:

Remark 3.3 Suppose $(Y, \{w_1, w_2\})$ is the disjoint union of (Y_1, w_1) and (Y_2, w_2) . If \mathcal{H}_i is a diagram for (Y_i, w_i) , then $\mathcal{H}_1 \sqcup \mathcal{H}_2$ is a diagram for $(Y, \{w_1, w_2\})$. By definition,

$$\text{CF}^-(\mathcal{H}_1 \sqcup \mathcal{H}_2, \mathfrak{s}_1 \sqcup \mathfrak{s}_2) := \text{CF}^-(\mathcal{H}_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2} \text{CF}^-(\mathcal{H}_2, \mathfrak{s}_2),$$

which is a module over $\mathbb{F}_2[U_{w_1}, U_{w_2}]$.

Let \mathbf{P} be the singleton $\{p\}$, and let $\sigma: \{w_1, w_2\} \rightarrow \mathbf{P}$ denote the unique map. Let σ_1 and σ_2 denote the restrictions of σ to $\{w_1\}$ and $\{w_2\}$. Then

$$\text{CF}^-(\mathcal{H}_1 \sqcup \mathcal{H}_2, \sigma, \mathfrak{s}_1 \sqcup \mathfrak{s}_2) \cong \text{CF}^-(\mathcal{H}_1, \sigma_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2[U_p]} \text{CF}^-(\mathcal{H}_2, \sigma_2, \mathfrak{s}_2),$$

which is a module over the single variable polynomial ring $\mathcal{R}_{\mathbf{P}} = \mathbb{F}_2[U_p]$.

3.4 Lipshitz’s cylindrical reformulation

For many technical arguments in this paper, it will be convenient to work with Lipshitz’s *cylindrical reformulation* [2006] of Heegaard Floer homology. If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w)$ is a Heegaard diagram for the singly pointed 3-manifold (Y, w) , Lipshitz shows that instead of counting holomorphic disks in $\text{Sym}^g(\Sigma)$, one can instead count holomorphic curves which map into the 4-manifold $\Sigma \times [0, 1] \times \mathbb{R}$. The equivalence is (morally) due to a *tautological correspondence* between holomorphic disks mapping into $\text{Sym}^g(\Sigma)$ and holomorphic curves (of higher genus) mapping into $\Sigma \times [0, 1] \times \mathbb{R}$ whose projection to $[0, 1] \times \mathbb{R}$ are g -fold branched covers. Lipshitz’s cylindrical setting has a similar extension for multipointed Heegaard diagrams; see [Ozsváth and Szabó 2008, Section 5.2].

We now describe the almost complex structures and moduli spaces which we consider in this paper (mostly taken directly from or slightly modified from [Lipshitz 2006]).

If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a multipointed Heegaard diagram, we focus on almost complex structures J on $\Sigma \times [0, 1] \times \mathbb{R}$ which satisfy the following axioms (from [loc. cit.]):

- (J1) J is tamed by the product symplectic form.
- (J2) There is a finite collection of points $P \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ such that J is split (i.e., equal to $j_\Sigma \times j_{\mathbb{D}}$) on a cylindrical neighborhood of $P \times [0, 1] \times \mathbb{R}$.
- (J3) J is translation-invariant in the \mathbb{R} factor.
- (J4) $J(\partial/\partial s) = \partial/\partial t$.
- (J5) J preserves the 2-planes $T(\Sigma \times \{(s, t)\})$ for all $(s, t) \in [0, 1] \times \mathbb{R}$.

Sometimes we will need to consider slightly more harshly perturbed almost complex structures, satisfying the following alternative to axiom (J5):

- (J5') There is a 2-plane distribution ξ on $\Sigma \times [0, 1]$ such that the restriction of ω to ξ is nondegenerate, J preserves ξ , and the restriction of J to ξ is compatible with ω . We further assume that ξ is tangent to $\Sigma \times \{\text{pt}\}$ near $(\alpha \cup \beta) \times [0, 1] \times \mathbb{R}$ and near $\Sigma \times \{0, 1\} \times \mathbb{R}$.

Following [loc. cit.], we will consider holomorphic curves $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ such that S is a Riemann surface with boundary, as well as positive punctures p_1, \dots, p_n and negative punctures q_1, \dots, q_n (where $n = g(\Sigma) + |\mathbf{w}| - 1$), satisfying the following:

- (M1) S is a smooth (not nodal) Riemann surface.
 (M2) $u(\partial S) \subseteq (\alpha \times \{1\} \times \mathbb{R}) \cup (\beta \times \{0\} \times \mathbb{R})$.
 (M3) $\lim_{z \rightarrow p_i} (\pi_{\mathbb{R}} \circ u)(z) = -\infty$ and $\lim_{z \rightarrow q_i} (\pi_{\mathbb{R}} \circ u)(z) = \infty$.
 (M4) u has finite energy.
 (M5) $\pi_{[0,1] \times \mathbb{R}} \circ u$ is locally nonconstant.
 (M6) u is an embedding.

When considering almost complex structures which satisfy (J5') instead of (J5), we will need to consider the following weaker version of the (M5) axiom:

- (M5') There is no nonempty open subset $U \subseteq S$ such that $\pi_{[0,1] \times \mathbb{R}} \circ u|_U$ is constant and takes value near $\{0, 1\} \times \mathbb{R}$ (in the sense of (J5')).

There is a similar description of holomorphic triangle maps in Lipshitz's cylindrical setting. Let Δ denote an unbounded, triangular region in the complex plane with three cylindrical ends, each identified with $[0, 1] \times [0, \infty)$. Following [ibid., Section 10.2], we consider almost complex structures on $\Sigma \times \Delta$ satisfying the following:

- (J'1) J is tamed by the split symplectic form on $\Sigma \times \Delta$.
 (J'2) There is a finite collection of points $P \subseteq \Sigma \setminus (\alpha \cup \beta \cup \gamma)$ with at least one point in each component of $\Sigma \setminus (\alpha \cup \beta \cup \gamma)$ such that J is split on a product neighborhood of $P \times \Delta$.
 (J'3) In the cylindrical ends of Δ , J is equal to a cylindrical almost complex structure satisfying (J1)–(J5).
 (J'4) The projection map $\pi_{\Delta}: \Sigma \times \Delta \rightarrow \Delta$ is holomorphic and the tangent space of each fiber of π_{Σ} is a complex line.

At times, we will need to deal with more harshly perturbed almost complex structures on $\Sigma \times \Delta$, which satisfy the following alternatives to axioms (J'3) and (J'4):

- (J'3') In the cylindrical ends of Δ , J agrees with cylindrical almost complex structures satisfying (J1)–(J4) and (J5'), above.
 (J'4') The 2-planes of $T(\{p\} \times \Delta)$ are complex lines of J for all $p \in \Sigma$.

(J5') The 2-planes $T(\Sigma \times \{d\})$, for $d \in \Delta$, are complex lines for J near $(\alpha \cup \beta \cup \gamma) \times \Delta$ and on $\Sigma \times U$ for an open subset $U \subseteq \Delta$ containing the three components of $\partial\Delta$.

3.5 Expected dimensions and transversality

For some technical arguments, we will need compute the expected dimension of certain moduli spaces, and also know when the expected dimension is generically correct. To give precise results, our expected dimensions must take into account the source curve S of a holomorphic map $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$, and not just the homology class. If S is a topological source and ϕ is a homology class, we define

$$\mathcal{M}(S, \phi)$$

to be the set of holomorphic curves $u: S \times [0, 1] \times \mathbb{R}$ which satisfy (M1)–(M5) (but possibly not (M6)). Near any curve where $D\bar{\partial}$ achieves transversality, the set $\mathcal{M}(S, \phi)$ will be a smooth manifold of dimension equal to the Fredholm index of $D\bar{\partial}$ at u .

Lipshitz [2014, Proposition 4.2'] proved that if $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ is a holomorphic curve which satisfies (M1)–(M6) (in particular, u is embedded), then the Fredholm index is equal to the Maslov index. More generally, at curves which only satisfy (M1)–(M5), the Fredholm index satisfies

$$\text{ind}(u) = \mu(\phi) - 2 \text{Sing}(u),$$

where $\text{Sing}(u)$ denotes the number of double points of u in an equivalent singularity.

If $X \subseteq \text{Sym}^n([0, 1] \times \mathbb{R})$ is a smooth submanifold, $p \in \Sigma \setminus (\alpha \cup \beta)$ is a point, and ϕ is a homology class with $n_p(\phi) = n$, there is a *matched moduli space*

$$\mathcal{M}(S, \phi, X) := \{u \in \mathcal{M}(S, \phi) : \rho^p(u) \in X\},$$

where and $\rho^p: \mathcal{M}(S, \phi) \rightarrow \text{Sym}^n([0, 1] \times \mathbb{R})$ is the map

$$(3-5) \quad \rho^p(u) := (u \circ \pi_{[0,1] \times \mathbb{R}})((u \circ \pi_\Sigma)^{-1}(p)).$$

The analysis becomes simpler if we restrict to submanifolds $X \subseteq \text{Sym}^n([0, 1] \times \mathbb{R})$ which avoid the *fat diagonal*, i.e., the codimension 2 subset of points with at least one repeated entry.

We state the following transversality result:

Proposition 3.4 *Suppose J is a generic almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ satisfying (J1)–(J5). Near any holomorphic curve $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ satisfying (M1)–(M5), the moduli space $\mathcal{M}(S, \phi)$ is transversely cut out and of dimension*

$$\text{ind}(u) = \mu(\phi) - 2 \text{Sing}(u).$$

If $X \subseteq \text{Sym}^n([0, 1] \times \mathbb{R})$ is a submanifold which avoids the fat diagonal, then, near any curve $u \in \mathcal{M}(S, \phi, X)$ satisfying (M1)–(M5), the space $\mathcal{M}(S, \phi, X)$ is transversely cut out and of dimension

$$\mu(\phi) - 2 \text{Sing}(u) - \text{codim}(X).$$

If J is a generic almost complex structure satisfying (J1)–(J4) and (J5'), then the statements hold at holomorphic curves which satisfy (M1)–(M4) and (M5') with no multiply covered closed components, and with no components S_0 such that $\pi_{[0,1] \times \mathbb{R}} \circ u|_{S_0}$ is constant and takes on a value near $\{0, 1\} \times \mathbb{R}$ (in the sense of (J5')).

See [Lipshitz 2006, Sections 3 and 4] for proofs of the statements about the unmatched moduli spaces near embedded curves, and see [Lipshitz 2014] for the statements about the unmatched moduli spaces near nonembedded curves. See [Juhász et al. 2021, Section 9.3] for an account of the proof involving the matched moduli spaces.

An analog of Proposition 3.4 holds for the moduli spaces of holomorphic triangles satisfying the natural analogs of (M1)–(M5) for almost complex structures satisfying (J'1)–(J'4). Furthermore, there is also a version for more harshly perturbed almost complex structures satisfying (J'1)–(J'2) and (J'3')–(J'5') using the moduli space axioms obtained by modifying (M1)–(M4) and (M5') in the obvious ways; see [Juhász et al. 2021, Section 9.3].

3.6 Naturality of Heegaard Floer homology

For functoriality, we need to understand the relation between the chain complexes obtained from different Heegaard diagrams of the same 3-manifold.

The following is standard:

Lemma 3.5 *Any two Heegaard diagrams for a multipointed 3-manifold can be connected by a sequence of the following moves:*

- (1) isotopies of the α or β curves not passing over the w basepoints;
- (2) handleslides of the α curves across each other or handleslides of the β curves across each other;
- (3) simple stabilizations or destabilizations;
- (4) pushing forward the diagram $(\Sigma, \alpha, \beta, w)$ under an automorphism ϕ of (Y, w) which is isotopic to id_Y , relative to w .

We explain move (3), appearing in Lemma 3.5. Suppose $B^3 \subseteq Y$ is a closed, embedded 3-ball which intersects Σ in a disk and is disjoint from $\alpha \cup \beta \cup w$. We say a diagram $(\Sigma', \alpha', \beta', w)$ is a *simple stabilization* of $(\Sigma, \alpha, \beta, w)$ if Σ agrees with Σ' outside of B^3 , $\alpha' = \alpha \cup \{\alpha'\}$, $\beta' = \beta \cup \{\beta'\}$, and $\Sigma' \cap B^3$ consists of a once-punctured torus. Furthermore, α' and β' intersect in a single, transverse intersection point, and are contained in the region $\Sigma' \cap B^3$. A *simple destabilization* is the inverse of a simple stabilization.

To each of the moves appearing in Lemma 3.5, Ozsváth and Szabó [2004b] associated a chain map between the corresponding Heegaard Floer complexes.

We now state the following naturality result:

Proposition 3.6 *Suppose that (Y, w) is a multipointed 3-manifold. To each pair (\mathcal{H}, J) and (\mathcal{H}', J') , there is a map*

$$\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')} : \text{CF}_J^-(\mathcal{H}, \mathfrak{s}) \rightarrow \text{CF}_{J'}^-(\mathcal{H}', \mathfrak{s}),$$

which is well defined up to $\mathbb{F}_2[U]$ -equivariant chain homotopy. Furthermore, the following are satisfied:

(1) If (\mathcal{H}, J) , (\mathcal{H}', J') and (\mathcal{H}'', J'') are three diagrams with almost complex structures, then

$$\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}'', J'')} \simeq \Psi_{(\mathcal{H}', J') \rightarrow (\mathcal{H}'', J'')} \circ \Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}.$$

(2) $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}, J)} \simeq \text{id}_{\text{CF}^-_{\mathcal{J}}(\mathcal{H}, s)}$.

Proof The details of the proof are due to many authors. Ozsváth and Szabó [2004b] constructed the transition maps $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$ and proved they were quasi-isomorphisms (establishing invariance of the isomorphism type of the homology groups). They also proved [2006, Theorem 2.1] most of the Floer-theoretic results necessary to show well-definedness of $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$. Lipshitz [2006, Proposition 11.4] showed that the transition maps are chain homotopy equivalences, as opposed to quasi-isomorphisms. Juhász and Thurston pointed out that [Ozsváth and Szabó 2006, Theorem 2.1] contains a gap, since the space of isotopies taking one embedded Heegaard surface to another is not connected. The proof of invariance of the transition maps is completed in [Juhász et al. 2021], by a careful topological analysis of the space of Heegaard splittings, followed by verification that the Heegaard Floer transition maps have no monodromy around a special loop of Heegaard diagrams (the *simple handleswap* loop [ibid., Definition 2.31]). □

We now describe the maps associated to the moves in Lemma 3.5, as well as changes of the almost complex structure.

If $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) , and J and J' are two almost complex structures on $\Sigma \times [0, 1] \times \mathbb{R}$, then the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}, J')}$ is defined by picking a noncylindrical almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$ which agrees with J on $\Sigma \times [0, 1] \times (-\infty, -1]$ and J' on $\Sigma \times [0, 1] \times [1, \infty)$, and counting Maslov index 0 \tilde{J} -holomorphic curves in $\Sigma \times [0, 1] \times \mathbb{R}$.

Next, we consider the maps associated to handleslides and isotopies of α and β . In this case, the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$ can be computed by a sequence of holomorphic triangle maps. If α' is obtained from α by a handleslide or isotopy, and the triple $(\Sigma, \alpha', \alpha, \beta, \mathbf{w})$ satisfies an admissibility condition (see Definition 3.12), the map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$ can be computed by counting holomorphic triangles via the formula

$$(3-6) \quad \Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}(-) = F_{\alpha', \alpha, \beta}(\Theta_{\alpha', \alpha}^+ \otimes -).$$

In (3-6), $\Theta_{\alpha', \alpha}^+$ denotes a cycle in $\text{CF}^-(\Sigma, \alpha', \alpha, \mathbf{w}, \mathfrak{s}_0)$ which represents the top-graded homogeneous element of $\text{HF}^-(\Sigma, \alpha', \alpha, \mathbf{w}, \mathfrak{s}_0)$, and \mathfrak{s}_0 denotes the torsion Spin^c structure on $(S^1 \times S^2)^{\#g(\Sigma)}$.

If the triple $(\Sigma, \alpha', \alpha, \beta, \mathbf{w})$ is not admissible, the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$ is obtained by a composition of triangle maps as in (3-6). We will often write $\Psi_{\alpha \rightarrow \alpha'}^\beta$ for the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$, in this situation.

Moves of the β curves can be computed using the holomorphic triangle maps, in an analogous fashion. If β' is obtained from β via a sequence of handleslides or isotopies, we write $\Psi_{\alpha}^{\beta \rightarrow \beta'}$ for the corresponding transition map.

If $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a Heegaard diagram for (Y, \mathbf{w}) , and $\mathcal{H}' = (\Sigma', \alpha \cup \{\alpha'\}, \beta \cup \{\beta'\}, \mathbf{w})$ is a simple stabilization, then we consider the map $\sigma: \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(\Sigma', \alpha', \beta', \mathbf{w}, \mathfrak{s})$ defined by the formula

$$\sigma(\mathbf{x}) := \mathbf{x} \times c,$$

extended $\mathbb{F}_2[U_{\mathbf{w}}]$ -equivariantly, where $\{c\} = \alpha' \cap \beta'$.

If J is an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ which is split in a neighborhood of the connected sum point, and $T > 0$, then one can construct an almost complex structure $J(T)$ on $\Sigma' \times [0, 1] \times \mathbb{R}$ which has a connected sum neck of length T inserted. According to the proof of [Lipshitz 2006, Proposition 12.5], if T is sufficiently large, the map σ will satisfy

$$\sigma \circ \partial_J = \partial_{J(T)} \circ \sigma.$$

For appropriately large T , the transition map from (\mathcal{H}, J) to $(\mathcal{H}', J(T))$ is defined to be σ . The meaning of “appropriately large” can be made precise: we say a neck length $T > 0$ satisfies *stabilization condition* (SC1) (and can be used to compute the simple stabilization map) if, for all $T' \geq T$, there is an almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$ interpolating $J(T)$ and $J(T')$ such that

$$(SC1) \quad \Psi_{\tilde{J}}(\mathbf{x} \times c) = \mathbf{x} \times c$$

for all $\mathbf{x} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$. We will consider a similar stabilizing condition when we define the free-stabilization, 1-handle and 3-handle maps; see conditions (SC2) and (SC3).

If J' is a general almost complex structure on $\Sigma' \times [0, 1] \times \mathbb{R}$ (possibly not satisfying condition (SC1)) then the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$ is defined as the composition

$$\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')} := \Psi_{(\mathcal{H}', J(T)) \rightarrow (\mathcal{H}', J')} \circ \Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J(T))} = \Psi_{(\mathcal{H}', J(T)) \rightarrow (\mathcal{H}', J')} \circ \sigma$$

for a T which satisfies condition (SC1).

Finally, if (\mathcal{H}', J') is obtained by pushing forward \mathcal{H} under a diffeomorphism ϕ of (Y, \mathbf{w}) which is isotopic to the identity, relative to \mathbf{w} , then the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}', J')}$ is defined to be the tautological map ϕ_* induced by ϕ .

If (Y, \mathbf{w}) is a multipointed 3-manifold and $\sigma: \mathbf{w} \rightarrow \mathbf{P}$ is a coloring, then the $\mathcal{R}_{\mathbf{P}}$ -modules $\text{CF}^-(\mathcal{H}, \sigma, \mathfrak{s})$ from (3-4) form a transitive system of chain complexes, for which we write $\text{CF}^-(Y, \mathbf{w}^{\sigma}, \mathfrak{s})$.

3.7 Admissibility of Heegaard diagrams

In order to achieve finite counts of holomorphic disks and triangles for CF^- , Ozsváth and Szabó [2004b, Sections 4 and 8] define several admissibility conditions (*weak admissibility* and *strong \mathfrak{s} -admissibility*) for singly pointed Heegaard diagrams, triples and quadruples. For multipointed diagrams of integer homology spheres, they also described [2008, Section 3.4] a weaker version of admissibility (*weak admissibility*), though this is not sufficient for our purposes. In this section, we extend their work on strong \mathfrak{s} -admissibility to multipointed diagrams of arbitrary 3-manifolds (see Definition 3.9).

Definition 3.7 If $(\Sigma, \alpha, \beta, \mathbf{w})$ is a multipointed Heegaard diagram, a *periodic domain* is an integral 2-chain P on Σ with boundary equal to a linear combination of the α and β curves, with $n_w(P) = 0$ for all $w \in \mathbf{w}$. A *periodic class* is a homology class $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ for some $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ such that $n_w(\phi) = 0$ for all $w \in \mathbf{w}$.

The domain of a periodic class is a periodic domain, though we usually will not make a distinction between the two. If R is a ring, we can also consider the set of R -valued periodic domains, for which we write Π_R .

There is a natural map

$$H: \Pi_{\mathbb{Z}} \rightarrow H_2(Y \setminus \mathbf{w}; \mathbb{Z}),$$

obtained by capping ∂P with an integral combination of compressing disks for the α and β curves. The construction of H works for other rings R , as well.

If $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is a homology class of disks, we write $D(\phi)$ for the domain of ϕ , viewed as a 2-chain on Σ . If $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ is a periodic class, define

$$H(\phi) := H(D(\phi)).$$

If $P \in \Pi_{\mathbb{Z}}$ is a periodic domain and λ is an integral 1-chain on Σ , with boundary equal to an integer sum of \mathbf{w} basepoints, then we will write $a(\lambda, P)$ for the integer obtained by summing each local difference of the class λ across each α curve as one traverses λ . We define an integer $b(\lambda, P)$ analogously, by summing differences of the domain P across the β curves as one traverses λ . See Section 4.1 for more details on $a(\lambda, P)$ and $b(\lambda, P)$.

If λ is an integral 1-chain on Σ , with boundary equal to a linear combination of the \mathbf{w} basepoints, then, by pushing the interior of γ into the α -handlebody, we obtain the formula

$$(3-7) \quad a(\lambda, P) = \#(\lambda \cap H(P)).$$

Similarly, if γ is an integral 1-cycle on Σ , then

$$(3-8) \quad a(\gamma, P) = \#(\gamma \cap H(P)) = \langle \text{PD}[\gamma], H(P) \rangle.$$

A helpful topological fact is the following:

Lemma 3.8 *The map $H: \Pi_R \rightarrow H_2(Y \setminus \mathbf{w}; R)$ is an isomorphism for any ring R .*

Proof Since $\Pi_{\mathbb{Z}}$ and $H_2(Y \setminus \mathbf{w}; \mathbb{Z})$ are free \mathbb{Z} -modules, it is sufficient to show the claim for $R = \mathbb{Z}$.

To see that H is a surjection, note that $Y \setminus N(\mathbf{w})$ is obtained by attaching 2-handles to $[0, 1] \times (\Sigma \setminus N(\mathbf{w}))$ along $\{0\} \times \alpha$ and $\{1\} \times \beta$. By putting a closed 2-cycle into general position, we can ensure that it intersects the cocores of the 2-handles transversely. After a homotopy, it becomes homologous to a 2-chain in $\Sigma \setminus N(\mathbf{w})$ with boundary equal to a linear combination of the α and β curves, together with some linear combination of the cores of the 2-handles. Hence H is surjective.

For injectivity, suppose $H(P) = 0 \in H_2(Y \setminus \mathbf{w}; \mathbb{Z})$. By (3-7), $a(\lambda, P) = 0$ for any λ which is a closed curve or a path connecting two basepoints in \mathbf{w} . Similarly, $b(\lambda, P) = 0$ for any such λ . By considering λ

arcs or curves which are dual to α and β curves, it follows that $H(P)$ has no changes across any of the α and β curves. Since P also has zero multiplicity at the basepoints, P must be zero everywhere. \square

By abuse of notation, we will also write $H(P)$ or $H(\phi)$ for the induced class in $H_2(Y; \mathbb{Z})$. The capping map H satisfies

$$(3-9) \quad \mu(\phi) = \langle c_1(\mathfrak{s}), H(\phi) \rangle + 2 \sum_{w \in \mathbf{w}} n_w(\phi).$$

This can be proven by a simple modification of Ozsváth and Szabó’s proof [2004a, Proposition 7.5] for singly pointed diagrams.

We make the following definition, extending Ozsváth and Szabó’s *strong \mathfrak{s} -admissibility* condition [2004b, Section 4] to multipointed diagrams:

Definition 3.9 If $\mathfrak{s} \in \text{Spin}^c(Y)$, we say a Heegaard diagram $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is *strongly \mathfrak{s} -admissible* if for each $N > 0$ and each nontrivial periodic domain $P \in \Pi_{\mathbb{Z}}$, the inequality

$$\langle c_1(\mathfrak{s}), H(P) \rangle = 2N \geq 0$$

implies that P has some multiplicity strictly greater than N .

Strong \mathfrak{s} -admissibility ensures finiteness of the differential on CF^- :

Lemma 3.10 *If \mathcal{H} is strongly \mathfrak{s} -admissible and j is fixed, then there are only finitely many homology classes $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = j$ and $D(\phi) \geq 0$.*

Proof The proof is a modification of Ozsváth and Szabó’s proof [2004b, Lemma 4.14] for singly pointed diagrams.

Fix any $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$. If $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is another class, then we can uniquely write

$$D(\phi) = D(\psi) + P + A,$$

where $P \in \Pi_{\mathbb{Z}}$ and A is a \mathbb{Z} -linear combination of the components of $\Sigma \setminus \alpha$. If $D(\phi) \geq 0$, then

$$(3-10) \quad -D(\psi) \leq P + A.$$

Suppose, to the contrary of the main statement, that there is an infinite sequence ϕ_n of distinct classes in $\pi_2(\mathbf{x}, \mathbf{y})$, with $\mu(\phi_n) = j$ and $D(\phi_n) \geq 0$. We obtain an infinite sequence of pairwise distinct pairs (P_n, A_n) which satisfy (3-10).

Since there are only finitely many components of $\Sigma \setminus (\alpha \cup \beta)$,

$$\|P_n + A_n\|_{\infty} \rightarrow \infty,$$

where $\|D\|_{\infty}$ denotes the maximum absolute value of the components of a domain D .

The coefficients of $P_n + A_n$ are bounded below by the coefficients of $-D(\psi)$. Similarly, since P_n satisfies $n_w(P_n) = 0$ for all n , it follows that $n_w(A_n) = n_w(\phi_n) - n_w(\psi)$. Since $D(\phi_n) \geq 0$, the coefficients of A_n are also bounded from below.

Since P_n is zero on \mathbf{w} , and A_n is determined by its values on \mathbf{w} , we have

$$(3-11) \quad \|A_n\|_\infty \leq \|P_n + A_n\|_\infty.$$

By the triangle inequality,

$$(3-12) \quad \|P_n\|_\infty \leq \|A_n\|_\infty + \|P_n + A_n\|_\infty.$$

Inequalities (3-11) and (3-12) imply $A_n/\|P_n + A_n\|_\infty$ and $P_n/\|A_n + P_n\|_\infty$ are both bounded, and hence admit subsequences which converge to real domains A_∞ and P_∞ , respectively.

Since $\mu(\phi_n) = \mu(\psi) = j$, we have

$$(3-13) \quad \mu(P_n^{\mathbf{x}} + A_n^{\mathbf{x}}) = 0,$$

where $P_n^{\mathbf{x}}, A_n^{\mathbf{x}} \in \pi_2(\mathbf{x}, \mathbf{x})$ denote the periodic classes corresponding to P_n and A_n . Combining (3-9) and (3-13), as well as the fact that $H(A_n) = 0 \in H_2(Y; \mathbb{Z})$ and $n_{\mathbf{w}}(P_n) = 0$, we obtain

$$(3-14) \quad 0 = \mu(P_n^{\mathbf{x}} + A_n^{\mathbf{x}}) = \langle c_1(\mathfrak{s}), H(P_n) \rangle + \sum_{\mathbf{w} \in \mathbf{w}} 2n_{\mathbf{w}}(A_n).$$

Taking limits in (3-14), we obtain

$$(3-15) \quad \langle c_1(\mathfrak{s}), H(P_\infty) \rangle + \sum_{\mathbf{w} \in \mathbf{w}} 2n_{\mathbf{w}}(A_\infty) = 0.$$

Since the multiplicities of A_n and $A_n + P_n$ are bounded below, and $\|A_n + P_n\|_\infty \rightarrow \infty$, we conclude

$$(3-16) \quad A_\infty \geq 0 \quad \text{and} \quad P_\infty + A_\infty \geq 0.$$

If there is a real pair (P_∞, A_∞) satisfying (3-15) and (3-16), then it is not hard to see that there is also a nearby rational domain which also satisfies (3-15) and (3-16). By clearing denominators, we can find a pair (P', A') of integral domains satisfying the same relations.

By (3-15) and (3-16),

$$\langle c_1(\mathfrak{s}), H(-P') \rangle = 2N \geq 0,$$

where $N = \sum_{\mathbf{w} \in \mathbf{w}} n_{\mathbf{w}}(A')$. Hence, $-P'$ has a multiplicity which is greater than N by strong admissibility. Since $A' \geq 0$, this contradicts the fact that $A' + P' \geq 0$. \square

We now prove that all multipointed Heegaard diagrams can be made admissible by performing an isotopy:

Proposition 3.11 *If $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a multipointed Heegaard diagram for (Y, \mathbf{w}) , and \mathfrak{s} is a fixed Spin^c structure on Y , then \mathcal{H} is isotopic to a strongly \mathfrak{s} -admissible diagram.*

Proof Our proof is a modification of Ozsváth and Szabó’s procedure [2004b, Lemma 5.4] for achieving strong \mathfrak{s} -admissibility for singly pointed diagrams, and their procedure [2008, Proposition 3.6] for achieving weak admissibility for multipointed diagrams of integer homology spheres.

Pick a collection of closed curves $\gamma_1, \dots, \gamma_k$ and arcs $\lambda_1, \dots, \lambda_n$ on Σ satisfying the following:

- (1) $\gamma_1, \dots, \gamma_k$ are pairwise disjoint, simple, closed curves which span $H_1(Y; \mathbb{Z})$.
- (2) Each λ_i is an embedded arc with boundary equal to two basepoints of \mathbf{w} .
- (3) The interiors of λ_i and λ_j are disjoint if $i \neq j$, and λ_i is disjoint from γ_j for all i and j .
- (4) Each basepoint of \mathbf{w} is in the boundary of at least one λ_i , and $\lambda_1 \cup \dots \cup \lambda_n$ is connected.

The above conditions imply that $\gamma_1, \dots, \gamma_k, \lambda_1, \dots, \lambda_n$ span $H_1(Y, \mathbf{w}; \mathbb{Z})$.

Such a collection $\gamma_1, \dots, \gamma_k, \lambda_1, \dots, \lambda_n$ can be constructed as follows. Let α_0 be any tuple of $g(\Sigma)$ attaching curves on Σ , which are disjoint from \mathbf{w} and which bound compressing disks in the handlebody U_α , such that $\Sigma \setminus \alpha_0$ is a connected, planar surface. The curves $\gamma_1, \dots, \gamma_k$ can be chosen to be duals of the curves of α_0 . The complement of α_0 is planar, connected, and contains all of the \mathbf{w} basepoints. The λ_i can be chosen to form an embedded tree such that each λ_i connects one \mathbf{w} basepoint to a chosen central basepoint of \mathbf{w} .

Write $\Pi'_\mathbb{Q}$ for the set of rational 2-chains on Σ of the form

$$P - \frac{1}{2} \langle c_1(\mathfrak{s}), H(P) \rangle \cdot [\Sigma] \quad \text{for } P \in \Pi_\mathbb{Q}.$$

Following Ozsváth and Szabó’s terminology, we call such domains \mathfrak{s} -renormalized periodic domains. The groups $\Pi_\mathbb{Q}$ and $\Pi'_\mathbb{Q}$ are canonically isomorphic, so there is a well-defined capping map

$$H: \Pi'_\mathbb{Q} \rightarrow H_2(Y \setminus \mathbf{w}; \mathbb{Q}),$$

which is an isomorphism by Lemma 3.8.

Let R_1, \dots, R_n be a collection of pairwise disjoint, embedded rectangles in Σ such that the following hold:

- (1) $\lambda_i \cap R_i$ is a connected arc.
- (2) $\lambda_i \cap R_j = \emptyset$ if $i \neq j$, and $\gamma_i \cap R_j = \emptyset$ for all i and j .
- (3) If τ is an attaching curve in α or β , then $\tau \cap \lambda_i \subseteq R_i$.

Fix $N > 0$. We perform the following two winding moves to α to construct a diagram $\mathcal{H}_N = (\Sigma, \alpha_N, \beta)$:

- (W1) Let γ_i^+ and γ_i^- denote two small, parallel pushoffs of γ_i , which are disjoint from each other and γ_i . Wind the α curves positively N times around γ_i^+ , and negatively N times around γ_i^- , as shown in Figure 3
- (W2) Along each λ_i , we perform a zigzag move to the α curves, as shown in Figure 4. The zigzag move is supported in the rectangle R_i .

We write $\Pi'_{\mathbb{Q},N}$ for the set of \mathfrak{s} -renormalized periodic domains on \mathcal{H}_N .

The groups $\Pi'_{\mathbb{Q},N}$ and $\Pi'_\mathbb{Q}$ are canonically isomorphic: indeed, Lemma 3.8 implies both are canonically isomorphic to $H_2(Y \setminus \mathbf{w}; \mathbb{Q})$. Write

$$\mathcal{W}_N: \Pi'_\mathbb{Q} \rightarrow \Pi'_{\mathbb{Q},N}$$

for this isomorphism.

Subclaim 3.11.1 For sufficiently large N , any nonzero \mathfrak{s} -renormalized periodic domain on \mathcal{H}_N has both positive and negative multiplicities.

It is straightforward to see that Subclaim 3.11.1 implies that, for sufficiently large N , \mathcal{H}_N is strongly \mathfrak{s} -admissible.

We now prove Subclaim 3.11.1. Define an L^∞ -norm on $H_2(Y \setminus \mathbf{w}; \mathbb{Q})$ via the formula

$$\|\sigma\|_\infty^{Y \setminus \mathbf{w}} = \max\{|\#(\lambda_1 \cap \sigma)|, \dots, |\#(\lambda_n \cap \sigma)|, |\#(\gamma_1 \cap \sigma)|, \dots, |\#(\gamma_k \cap \sigma)|\},$$

where $\#(\lambda_i \cap \sigma)$ denotes the algebraic intersection number. Similarly, define an L^∞ -norm on $H_2(Y; \mathbb{Q})$ via the formula

$$\|\sigma\|_\infty^Y := \max\{|\#(\gamma_1 \cap \sigma)|, \dots, |\#(\gamma_k \cap \sigma)|\}.$$

Let $S'_\mathbb{Q}$ denote the unit sphere in $\Pi'_\mathbb{Q}$ with respect to the $\|\cdot\|_\infty^{Y \setminus \mathbf{w}}$ -norm, i.e.,

$$S'_\mathbb{Q} := \{P \in \Pi'_\mathbb{Q} : \|H(P)\|_\infty^{Y \setminus \mathbf{w}} = 1\}.$$

It is sufficient to show Subclaim 3.11.1 for \mathfrak{s} -renormalized periodic domains P such that $H(P) \in S'_\mathbb{Q}$, since scaling by \mathbb{Q} preserves the property of having positive and negative coefficients.

Since $\gamma_1, \dots, \gamma_k$ span $H_1(Y; \mathbb{Z})$, we can write

$$(3-17) \quad \text{PD}[c_1(\mathfrak{s})] = a_1[\gamma_1] + \dots + a_k[\gamma_k]$$

for some $a_i \in \mathbb{Z}$. Let M denote the quantity

$$M := \sum_{i=1}^k |a_i|.$$

Pick a real number ϵ so that

$$(3-18) \quad 0 < \epsilon < \frac{1}{2} \quad \text{and} \quad M\epsilon < \frac{1}{2}.$$

Let C be the maximum absolute value of any multiplicity of any domain in $S'_\mathbb{Q}$, and let $N > 0$ be an integer which satisfies

$$(3-19) \quad N \cdot \epsilon > C.$$

We will show that, if N satisfies (3-19), then any nonzero \mathfrak{s} -renormalized domain on \mathcal{H}_N has both positive and negative multiplicities.

Suppose that $P \in S'_\mathbb{Q}$. Note that by definition \mathcal{W}_N does not change the induced class in $H_2(Y \setminus \mathbf{w}; \mathbb{Z})$, and hence \mathcal{W}_N does not change either $\|\cdot\|_\infty^{Y \setminus \mathbf{w}}$ or $\|\cdot\|_\infty^Y$. We break the argument into two cases:

- (1) $\|H(P)\|_\infty^Y > \epsilon$.
- (2) $\|H(P)\|_\infty^Y \leq \epsilon$.

In case (1), let γ_i be a curve such that $|\#(\gamma_i \cap H(P))| > \epsilon$. By (3-7), $|a(\gamma_i, P)| > \epsilon$. Pick $x \in \gamma_i \setminus (\alpha \cup \beta)$ and let x^+ and x^- be nearby points on $\gamma_i^+ \setminus (\alpha \cup \beta)$ and $\gamma_i^- \setminus (\alpha \cup \beta)$.

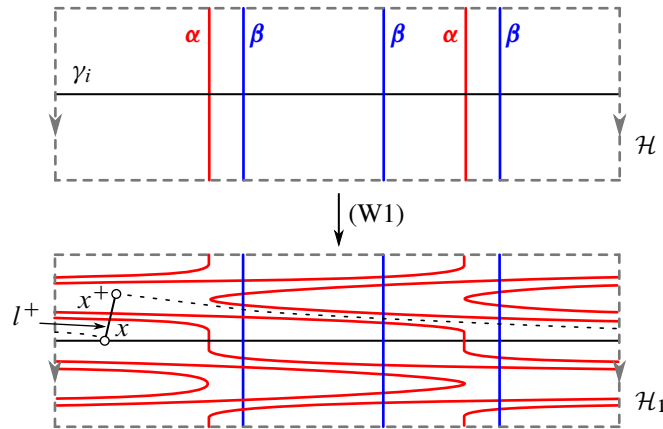


Figure 3: Winding move (W1) near a closed curve γ_i on Σ . The arc l^+ is shown. The left and right sides of each rectangle are identified to form an annulus on Σ .

Let l^+ be an arc from x to x^+ which is contained in a neighborhood of γ_i and does not intersect α or β . On \mathcal{H}_N , l^+ intersects a curve in α_N with N times the geometric multiplicity that γ_i intersects the corresponding curve of α . Furthermore, the arc l^+ can be concatenated with an arc which connects x and x^+ and is disjoint from α_N such that the concatenation is homologous to $N \cdot \gamma$ on Σ (see Figure 3, bottom). Writing l^- for a similar path from x to x^- , we have

$$(3-20) \quad a(l^+, \mathcal{W}_N(P)) = N \cdot a(\gamma_i, P) \quad \text{and} \quad a(l^-, \mathcal{W}_N(P)) = -N \cdot a(\gamma_i, P).$$

Observing that l^+ and l^- intersect no β curves, it follows from (3-20) that

$$(3-21) \quad n_{x^+}(\mathcal{W}_N(P)) = n_x(P) + N \cdot a(\gamma_i, P) \quad \text{and} \quad n_{x^-}(\mathcal{W}_N(P)) = n_x(P) - N \cdot a(\gamma_i, P).$$

Together, equations (3-19) and (3-21) imply that $\mathcal{W}_N(P)$ has both positive and negative multiplicities, proving Subclaim 3.11.1 in case (1).




Figure 4: The zigzag move (W2) near an arc λ_i with $\partial\lambda_i = \{w_1, w_2\}$. The arcs l_i and l'_i near w_1 and w_2 are shown. The move is supported in the rectangle R_i , whose boundary is the dashed box.

Next, we consider case (2). Combining (3-17) and (3-18) and the triangle inequality, we obtain

$$(3-22) \quad |\langle c_1(\mathfrak{s}), H(P) \rangle| = \left| \sum_{i=1}^k a_i \langle \text{PD}[\gamma_i], H(P) \rangle \right| \leq \sum_{i=1}^k |a_i| \cdot |\#(\gamma_i \cap H(P))| \leq M \cdot \epsilon < \frac{1}{2}.$$

Define $m := -\langle c_1(\mathfrak{s}), H(P) \rangle$, which is the multiplicity of $\mathcal{W}_N(P)$ at each basepoint $w \in \mathbf{w}$. By (3-22), $m \in (-\frac{1}{2}, \frac{1}{2})$. Furthermore, since $\|H(P)\|_\infty^{Y \setminus \mathbf{w}} = 1$, there is a λ_i such that $a(\lambda_i, P) = 1$. We consider the two arcs l_i and l'_i on Σ , shown in Figure 4, which intersect no curves on Σ , and intersect only the α curves on \mathcal{H}_N . Note that

$$a(\lambda_i, \mathcal{W}_N(P)) = a(l_i, \mathcal{W}_N(P)) = a(l'_i, \mathcal{W}_N(P)),$$

since l_i and l'_i can both be completed to curves on Σ which are isotopic to λ_i , and whose only intersections with α occur along l_i or l'_i (see the dashed lines on the bottom of Figure 4). Since l_i and l'_i do not intersect the β curves, the difference in multiplicity of $\mathcal{W}_N(P)$ between the two points of ∂l_i is $a(\lambda_i, \mathcal{W}_N(P)) = 1$, and similarly for the two points of $\partial l'_i$. Hence, $\mathcal{W}_N(P)$ has multiplicities of $m + 1$ and $m - 1$. Since $|m| < \frac{1}{2}$, it follows that $m - 1 < 0$ and $m + 1 > 0$. Subclaim 3.11.1 follows in case (2). \square

3.8 Admissibility for Heegaard triples and quadruples

We also need to consider admissibility of Heegaard triples and multidagrams.

If $(\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a Heegaard triple, Ozsváth and Szabó [2004b, Section 8.1] construct the 4-manifold

$$(3-23) \quad X_{\alpha, \beta, \gamma} := (\Delta \times \Sigma) \cup (e_\alpha \times U_\alpha) \cup (e_\beta \times U_\beta) \cup (e_\gamma \times U_\gamma),$$

where Δ is a triangle with boundary edges e_α, e_β and e_γ , and U_α, U_β and U_γ are standard handlebodies with boundary identified with Σ . The manifold $X_{\alpha, \beta, \gamma}$ satisfies

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

Ozsváth and Szabó [ibid., Section 8.1.4] construct a map

$$(3-24) \quad \mathfrak{s}_{\mathbf{w}}: \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \rightarrow \text{Spin}^c(X_{\alpha, \beta, \gamma}).$$

Definition 3.12 If $(\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a multipointed Heegaard triple, we say that an integral 2-chain P on Σ is a *triply periodic domain* if ∂P is a linear combination of the α, β and γ curves, and $n_w(P) = 0$ for all $w \in \mathbf{w}$.

If $\mathcal{Q} = (\Sigma, \alpha, \beta, \gamma, \delta, \mathbf{w})$ is a multipointed Heegaard quadruple, a *quadruply periodic domain* on \mathcal{Q} is defined analogously.

Ozsváth and Szabó [ibid., Definition 8.8] define a notion of strong admissibility for Heegaard triples, which extends to the multipointed setting:

Definition 3.13 Suppose $\mathcal{T} = (\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a multipointed Heegaard triple. If $\mathfrak{s} \in \text{Spin}^c(X_{\alpha, \beta, \gamma})$, we say \mathcal{T} is *strongly \mathfrak{s} -admissible* if each nonzero triply periodic domain P which can be written as a

sum of doubly periodic domains

$$P = P_{\alpha,\beta} + P_{\beta,\gamma} + P_{\alpha,\gamma} \quad \text{with} \quad \langle c_1(\mathfrak{s}), H(P_{\alpha,\beta}) \rangle + \langle c_1(\mathfrak{s}), H(P_{\beta,\gamma}) \rangle + \langle c_1(\mathfrak{s}), H(P_{\alpha,\gamma}) \rangle = 2N \geq 0$$

has a multiplicity which is strictly greater than N .

A straightforward extension of Lemma 3.10 and Proposition 3.11 gives the following (compare [ibid., Lemmas 8.10 and 8.11]):

Lemma 3.14 *Suppose that $\mathcal{T} = (\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a Heegaard triple.*

- (1) *If \mathcal{T} is strongly \mathfrak{s} -admissible, \mathbf{x} , \mathbf{y} and \mathbf{z} are fixed intersection points on \mathcal{T} , and j is a fixed integer, then there are only finitely many $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, $\mu(\psi) = j$ and $D(\psi) \geq 0$.*
- (2) *If $\mathfrak{s} \in \text{Spin}^c(X_{\alpha,\beta,\gamma})$, then \mathcal{T} is isotopic to a strongly \mathfrak{s} -admissible Heegaard triple.*

Given a Heegaard quadruple $(\Sigma, \alpha, \beta, \gamma, \delta, \mathbf{w})$, there is a four-ended cobordism $X_{\alpha,\beta,\gamma,\delta}$, and a map

$$\mathfrak{s}_{\mathbf{w}}: \pi_2(\mathbf{w}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \rightarrow \text{Spin}^c(X_{\alpha,\beta,\gamma,\delta}).$$

There are two natural decompositions of the 4-manifold $X_{\alpha,\beta,\gamma,\delta}$:

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

For the purposes of ensuring finiteness of holomorphic rectangle counts, we need finiteness of homotopy classes in an entire $\delta H^1(Y_{\beta,\delta}) + \delta H^1(Y_{\alpha,\gamma})$ -orbit \mathfrak{S} in $\text{Spin}^c(X_{\alpha,\beta,\gamma,\delta})$.

Definition 3.15 *Suppose $\mathcal{Q} = (\Sigma, \alpha, \beta, \gamma, \delta, \mathbf{w})$ is a multipointed Heegaard quadruple and \mathfrak{S} is a $\delta H^1(Y_{\beta,\delta}) + \delta H^1(Y_{\alpha,\gamma})$ -orbit of Spin^c structures on $X_{\alpha,\beta,\gamma,\delta}$. We say that \mathcal{Q} is *strongly \mathfrak{S} -admissible* if, whenever $\mathfrak{s} \in \mathfrak{S}$ and P is a nonzero periodic class which can be written as a sum of doubly periodic domains*

$$P = \sum_{\{\xi,\eta\} \subseteq \{\alpha,\beta,\gamma,\delta\}} P_{\xi,\eta}$$

such that

$$\sum_{\{\xi,\eta\} \subseteq \{\alpha,\beta,\gamma,\delta\}} \langle c_1(\mathfrak{s}|_{Y_{\xi,\eta}}), H(P_{\xi,\eta}) \rangle = 2N \geq 0,$$

then P has a multiplicity strictly greater than N .

If \mathcal{Q} is strongly \mathfrak{S} -admissible Heegaard quadruple, then a finiteness result for positive rectangle classes of a fixed Maslov index, similar to Lemma 3.14 can be proven by adapting Lemma 3.10.

If \mathcal{Q} is an arbitrary Heegaard quadruple and \mathfrak{S} is an $\delta H^1(Y_{\beta,\delta}) + \delta H^1(Y_{\alpha,\gamma})$ -orbit, then strong \mathfrak{S} -admissibility can be achieved by modifying the winding and zigzag procedure from Proposition 3.11 as long as the quantity

$$\sum_{\{\xi,\eta\} \subseteq \{\alpha,\beta,\gamma,\delta\}} \langle c_1(\mathfrak{s}|_{Y_{\xi,\eta}}), H(P_{\xi,\eta}) \rangle$$

is independent of the choice of $\mathfrak{s} \in \mathfrak{S}$. In particular, strong admissibility can be achieved by winding the curves α , β , γ and δ as long as

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

4 Relative homology action

Ozsváth and Szabó [2004b, Section 4.2.5] constructed an action of $\Lambda^*(H_1(Y; \mathbb{Z})/\text{Tors})$ on the singly pointed Heegaard Floer groups $\text{HF}^\circ(Y, w, \mathfrak{s})$. In this section, we describe similar maps on the multipointed Floer complexes for closed loops in Y , as well as paths between pairs of basepoints. We call these maps the *relative homology action*. Our construction is similar to an action of relative homology on sutured Floer homology constructed by Ni [2014].

If (Y, \mathbf{w}) is a multipointed 3-manifold and λ is a path between two basepoints $w_1, w_2 \in \mathbf{w}$, we construct two maps

$$A_\lambda, B_\lambda: \text{CF}^-(Y, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(Y, \mathbf{w}, \mathfrak{s}),$$

which we call the *type-A* and *type-B* relative homology maps, respectively.

We will show that A_λ and B_λ satisfy

$$\partial A_\lambda + A_\lambda \partial = \partial B_\lambda + B_\lambda \partial = U_{w_1} + U_{w_2}.$$

See Lemma 4.1.

If γ is a closed loop in Y , there are similar chain maps A_γ and B_γ . For a closed curve γ , the maps A_γ and B_γ coincide, and both agree with the familiar action of $H_1(Y; \mathbb{Z})/\text{Tors}$. See Lemma 4.7.

4.1 Construction of the relative homology action

Suppose $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a multipointed Heegaard diagram for (Y, \mathbf{w}) . If λ is a path in Y from w_1 to w_2 , we can homotope λ so that it has image in Σ and is an immersion. Furthermore, we can assume that λ intersects the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves transversely, and is disjoint from all intersections $\alpha_i \cap \beta_j$. Let a_1, \dots, a_m denote the points of intersection between λ and the $\boldsymbol{\alpha}$ curves. Let b_1, \dots, b_k denote the points of intersection between λ and the $\boldsymbol{\beta}$ curves. Given a homology class of disks $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we let $d_i^{\alpha, \lambda}(\phi)$ denote the difference between the multiplicities of ϕ on the two sides of the point a_i . Similarly, we let $d_i^{\beta, \lambda}(\phi)$ denote the difference between the multiplicities of ϕ on the two sides of b_i . We define

$$(4-1) \quad a(\lambda, \phi) = \sum_{i=1}^m d_i^{\alpha, \lambda}(\phi) \quad \text{and} \quad b(\lambda, \phi) = \sum_{i=1}^k d_i^{\beta, \lambda}(\phi),$$

which we view as elements of \mathbb{F}_2 . An orientation of λ allows us to lift the quantities $a(\lambda, \phi)$ and $b(\lambda, \phi)$ to \mathbb{Z} , which will occasionally be useful.

We define the endomorphism A_λ via the formula

$$(4-2) \quad A_\lambda(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} a(\lambda, \phi) \cdot \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}.$$

The map B_λ similarly to (4-2), by replacing the factor of $a(\lambda, \phi)$ with $b(\lambda, \phi)$.

If γ is a closed curve in Y , we can similarly modify (4-2) to define homology actions A_γ and B_γ on $\text{CF}^-(Y, \mathbf{w}, \mathfrak{s})$.

4.2 Properties of the relative homology action

In this section, we prove some basic properties of the relative homology action.

Lemma 4.1 *Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, w)$ is a diagram for (Y, w) , and λ is an immersed path in Σ which connects two basepoints w_1 and w_2 . Then*

$$A_\lambda \partial + \partial A_\lambda = B_\lambda \partial + \partial B_\lambda = U_{w_1} + U_{w_2}.$$

If γ is a closed, immersed curve on Σ , then

$$A_\gamma \partial + \partial A_\gamma = B_\gamma \partial + \partial B_\gamma = 0.$$

Proof We focus on the claim about the map A_λ , when λ connects two basepoints w_1 and w_2 . The other claims are straightforward modifications.

Our argument proceeds by counting the ends of index 2 moduli spaces. Suppose that $\phi \in \pi_2(x, z)$ is a homology class of disks with $\mu(\phi) = 2$. The 1-dimensional space $\widehat{\mathcal{M}}(\phi) := \mathcal{M}(\phi)/\mathbb{R}$ admits a compactification $\overline{\widehat{\mathcal{M}}}(\phi)$, whose ends correspond to strip breaking and Maslov index 2 boundary degenerations. Since compact 1-manifolds have an even number of ends, we have $\#(\partial \overline{\widehat{\mathcal{M}}}(\phi)) = 0$, and hence

$$(4-3) \quad a(\lambda, \phi) \cdot \#(\partial \overline{\widehat{\mathcal{M}}}(\phi)) = 0.$$

If $\phi_1 \in \pi_2(x, y)$ and $\phi_2 \in \pi_2(y, z)$ are two homology classes, then

$$(4-4) \quad a(\lambda, \phi_2 * \phi_1) = a(\lambda, \phi_2) + a(\lambda, \phi_1).$$

We consider separately the cases that $x \neq z$ or $x = z$. Consider first the case that $x \neq z$. We write $\pi_2^\alpha(x)$ and $\pi_2^\beta(x)$ for the groups of homology classes of α and β degenerations on $(\Sigma, \alpha, \beta, w)$. We can view $\pi_2^\alpha(x)$ (resp. $\pi_2^\beta(x)$) as the set of integral 2-chains on Σ with boundary equal to a linear combination of the α curves (resp. β curves).

If $w \in w$ and $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, then there is a unique class $A_w^x \in \pi_2^\alpha(x)$ whose domain has multiplicity 1 at w and 0 on the other components of $\Sigma \setminus \alpha$. A class $B_w^x \in \pi_2^\beta(x)$ is similarly specified. Furthermore, any homology class $A^x \in \pi_2^\alpha(x)$ decomposes as

$$A^x = \sum_{w \in w} n_w(A) \cdot A_w^x.$$

Using Lipshitz’s formula [2006, equation 8] for the Maslov index, we compute

$$\mu(A_w^x) = \mu(B_w^x) = 2.$$

Consequently, if $A^x \in \pi_2^\alpha(x)$ and $B^x \in \pi_2^\beta(x)$, then

$$(4-5) \quad \mu(A^x) = 2 \sum_{w \in w} n_w(A^x) \quad \text{and} \quad \mu(B^x) = 2 \sum_{w \in w} n_w(B^x).$$

It follows that if ϕ is a Maslov index 2 class and a broken holomorphic curve appears in $\overline{\partial\widehat{\mathcal{M}}(\phi)}$ which contains a nontrivial boundary degeneration, then the remaining curves must have Maslov index 0, and hence must represent the constant class, by transversality. Hence, if $\mathbf{x} \neq \mathbf{z}$, then boundary degenerations cannot occur in the ends of $\widehat{\mathcal{M}}(\phi)$, so (4-3) implies that, for each index 2 class $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$,

$$\begin{aligned} 0 &= a(\lambda, \phi) \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \phi_2 \in \pi_2(\mathbf{y}, \mathbf{z}) \\ \mu(\phi_1) = \mu(\phi_2) = 1 \\ \phi_1 * \phi_2 = \phi}} \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1) + n_{\mathbf{w}}(\phi_2)} \cdot \mathbf{z} \\ &= \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \phi_2 \in \pi_2(\mathbf{y}, \mathbf{z}) \\ \mu(\phi_1) = \mu(\phi_2) = 1 \\ \phi_1 * \phi_2 = \phi}} (a(\lambda, \phi_1) + a(\lambda, \phi_2)) \cdot \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1) + n_{\mathbf{w}}(\phi_2)} \cdot \mathbf{z}. \end{aligned}$$

Summing over all $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ with $\mu(\phi) = 2$, we get that the \mathbf{z} -component of $(A_\lambda \partial + \partial A_\lambda)(\mathbf{x})$ is zero.

If $\mathbf{x} = \mathbf{z}$, there may be ends of $\widehat{\mathcal{M}}(\phi)$ corresponding to boundary degenerations. If ϕ is one of the classes $A_w^{\mathbf{x}}$, we write $\widehat{\mathcal{N}}^\alpha(\phi)$ for the moduli space of cylindrical α boundary degenerations

$$u: S \rightarrow \Sigma \times [0, \infty) \times \mathbb{R}$$

representing ϕ , modulo conformal automorphisms of $[0, \infty) \times \mathbb{R}$.

Hence, the \mathbf{x} -component of $(\partial A_\lambda + A_\lambda \partial)(\mathbf{x})$ is equal to

$$\sum_{w \in \mathbf{w}} a(\lambda, A_w^{\mathbf{x}}) \cdot \#\widehat{\mathcal{N}}^\alpha(A_w^{\mathbf{x}}) U_w \cdot \mathbf{x} + \sum_{w \in \mathbf{w}} a(\lambda, B_w^{\mathbf{x}}) \cdot \#\widehat{\mathcal{N}}^\beta(B_w^{\mathbf{x}}) U_w \cdot \mathbf{x}.$$

We note that $a(\lambda, A_w^{\mathbf{x}}) = 0$ unless $w \in \{w_1, w_2\}$, in which case

$$a(\lambda, A_{w_1}^{\mathbf{x}}) = a(\lambda, A_{w_2}^{\mathbf{x}}) = 1.$$

Furthermore, $a(\lambda, B_w^{\mathbf{x}}) = 0$ for all $w \in \mathbf{w}$. The counts of the moduli spaces of boundary degenerations were proven by Ozsváth and Szabó [2008, Theorem 5.5]: for generic almost complex structure,

$$(4-6) \quad \#\widehat{\mathcal{N}}(A_w^{\mathbf{x}}) \equiv \begin{cases} 1 \pmod{2} & \text{if } |\mathbf{w}| > 1, \\ 0 & \text{if } |\mathbf{w}| = 1. \end{cases}$$

Hence, the \mathbf{x} -component of

$$(A_\lambda \partial + \partial A_\lambda + U_{w_1} + U_{w_2})(\mathbf{x})$$

is zero. Combining this with the computation in the case that $\mathbf{x} \neq \mathbf{z}$, the proof is complete. □

We now consider the interaction between the holomorphic triangle maps and the relative homology maps. If $(\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a multipointed Heegaard triple and λ is an immersed path between two basepoints $w_1, w_2 \in \mathbf{w}$, as an extension of (4-1), let $a(\lambda, \phi)$, $b(\lambda, \phi)$ and $c(\lambda, \phi)$ denote the sums of differences

of the multiplicities of ϕ across the α , β or γ curves, respectively. Let A_λ , B_λ and C_λ denote the endomorphisms of $\text{CF}^-(\Sigma, \alpha, \beta, w)$, $\text{CF}^-(\Sigma, \beta, \gamma, w)$ and $\text{CF}^-(\Sigma, \alpha, \gamma, w)$, defined by modifying (4-2). Note that all three endomorphisms A_λ , B_λ and C_λ are defined on all three complexes.

Lemma 4.2 *Suppose that $(\Sigma, \alpha, \beta, \gamma, w)$ is a multipointed Heegaard triple and $\mathfrak{s} \in \text{Spin}^c(X_{\alpha, \beta, \gamma})$. Then*

$$\begin{aligned} F_{\alpha, \beta, \gamma, \mathfrak{s}}(A_\lambda \otimes \text{id}) &\simeq A_\lambda \circ F_{\alpha, \beta, \gamma, \mathfrak{s}}(\text{id} \otimes \text{id}), \\ F_{\alpha, \beta, \gamma, \mathfrak{s}}(B_\lambda \otimes \text{id}) &\simeq F_{\alpha, \beta, \gamma, \mathfrak{s}}(\text{id} \otimes B_\lambda), \\ F_{\alpha, \beta, \gamma, \mathfrak{s}}(\text{id} \otimes C_\lambda) &\simeq C_\lambda \circ F_{\alpha, \beta, \gamma, \mathfrak{s}}(\text{id} \otimes \text{id}), \end{aligned}$$

as maps from $\text{CF}^-(\Sigma, \alpha, \beta, w, \mathfrak{s}|_{Y_{\alpha, \beta}}) \otimes_{\mathbb{F}_2[U_w]} \text{CF}^-(\Sigma, \beta, \gamma, w, \mathfrak{s}|_{Y_{\beta, \gamma}})$ to $\text{CF}^-(\Sigma, \alpha, \gamma, w, \mathfrak{s}|_{Y_{\alpha, \gamma}})$.

Proof Consider the first relation, involving A_λ . The subsequent two relations involving B_λ and C_λ can be proven mutatis mutandis. We prove the relation by counting the ends of index 1 moduli spaces of triangles. Suppose that $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is a homology class with $\mathfrak{s}_w(\psi) = \mathfrak{s}$, with $\mu(\psi) = 1$. The moduli space $\mathcal{M}(\psi)$ can be compactified into a compact 1-manifold $\overline{\mathcal{M}(\psi)}$ whose ends consist of pairs consisting of an index 1 holomorphic strip and an index 0 holomorphic triangle. Since compact 1-manifolds have an even number of ends, we have

$$\sum_{\substack{\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \\ \mu(\psi) = 1 \\ \mathfrak{s}_w(\psi) = \mathfrak{s}}} a(\lambda, \psi) \cdot \#(\partial \overline{\mathcal{M}(\psi)}) \cdot U_w^{n_w(\psi)} = 0.$$

If ϕ is a homology class of disks and ψ is a homology class of triangles, then, similar to (4-4),

$$a(\lambda, \psi * \phi) = a(\lambda, \psi) + a(\lambda, \phi).$$

It follows that

$$\begin{aligned} (4-7) \quad A_\lambda \circ F_{\alpha, \beta, \gamma, \mathfrak{s}}(\text{id} \otimes \text{id}) + F_{\alpha, \beta, \gamma, \mathfrak{s}}(A_\lambda \otimes \text{id}) + F_{\alpha, \beta, \gamma, \mathfrak{s}}(\text{id} \otimes A_\lambda) \\ = H_{\alpha, \beta, \gamma, \mathfrak{s}}^{A, \lambda} \circ (\partial_{\alpha, \beta} \otimes \text{id} + \text{id} \otimes \partial_{\beta, \gamma}) + \partial_{\alpha, \gamma} \circ H_{\alpha, \beta, \gamma, \mathfrak{s}}^{A, \lambda}, \end{aligned}$$

where $H_{\alpha, \beta, \gamma, \mathfrak{s}}^{A, \lambda}$ is the map defined on intersection points by the formula

$$H_{\alpha, \beta, \gamma, \mathfrak{s}}^{A, \lambda}(\mathbf{x} \otimes \mathbf{y}) = \sum_{\substack{\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma \\ \psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \\ \mu(\psi) = 0}} a(\lambda, \psi) \cdot \#\mathcal{M}(\psi) U_w^{n_w(\psi)} \cdot \mathbf{z},$$

and extended $\mathbb{F}_2[U_w]$ -equivariantly.

Finally, we note that if ϕ is a homology class of disks on the diagram $(\Sigma, \beta, \gamma, w)$, then the quantities $a(\lambda, \phi)$ vanish, since ϕ has no changes across the α curves. Hence the map

$$A_\lambda: \text{CF}^-(\Sigma, \beta, \gamma, w, \mathfrak{s}|_{\beta\gamma}) \rightarrow \text{CF}^-(\Sigma, \beta, \gamma, w, \mathfrak{s}|_{\beta\gamma})$$

vanishes. Combining this fact with (4-7), the main statement follows. □

Lemma 4.3 (1) Suppose that λ is a path on Σ connecting a pair of basepoints on \mathcal{H} which can be written as a concatenation $\lambda_2 * \lambda_1$ of two paths which connect pairs of basepoints. Then

$$A_{\lambda_2 * \lambda_1} = A_{\lambda_2} + A_{\lambda_1}.$$

(2) Suppose that λ is a path on Σ connecting a pair of basepoints, and γ is a closed loop on Σ , which has nontrivial intersection with γ . If $\lambda * \gamma$ denotes the path obtained by splicing γ into λ , then

$$A_{\gamma * \lambda} = A_{\gamma} + A_{\lambda}.$$

(3) If γ and γ' are two closed curves on Σ , then

$$A_{\gamma' * \gamma} = A_{\gamma'} + A_{\gamma}.$$

The same relations hold for the type-B relative homology maps.

Proof The first claim follows immediately from (4-2) since

$$a(\lambda_2 * \lambda_1, \phi) = a(\lambda_2, \phi) + a(\lambda_1, \phi)$$

for any homology class of disks ϕ . The second and third claims are proven similarly. □

We now compute the commutator of the relative homology maps:

Lemma 4.4 Suppose λ_1 and λ_2 are two paths connecting pairs of basepoints in \mathbf{w} . Then

$$A_{\lambda_1} A_{\lambda_2} + A_{\lambda_2} A_{\lambda_1} = \sum_{w \in \partial \lambda_1 \cap \partial \lambda_2} U_w.$$

Proof Our proof proceeds by counting the ends of $\widehat{\mathcal{M}}(\phi)$ for classes $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ with Maslov index 2. As in Lemma 4.1, there are two cases to consider: $\mathbf{x} \neq \mathbf{z}$ and $\mathbf{x} = \mathbf{z}$.

If $\mathbf{x} \neq \mathbf{z}$ and $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ has Maslov index 2, then the ends of $\widehat{\mathcal{M}}(\phi)$ all correspond to strip breaking. Summing over ϕ , we have

$$(4-8) \quad 0 = \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{z}) \\ \mu(\phi)=2}} a(\lambda_1, \phi) a(\lambda_2, \phi) \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \phi_2 \in \pi_2(\mathbf{y}, \mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1 \\ \phi_1 + \phi_2 = \phi}} \# \widehat{\mathcal{M}}(\phi_1) \cdot \# \widehat{\mathcal{M}}(\phi_2) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{z}.$$

Noting that

$$\begin{aligned} & a(\lambda_1, \phi_1 + \phi_2) a(\lambda_2, \phi_1 + \phi_2) \\ &= (a(\lambda_1, \phi_1) a(\lambda_2, \phi_2) + a(\lambda_1, \phi_2) a(\lambda_2, \phi_1)) + (a(\lambda_1, \phi_1) a(\lambda_2, \phi_1) + a(\lambda_1, \phi_2) a(\lambda_2, \phi_2)), \end{aligned}$$

we can rewrite (4-8) to obtain

$$(4-9) \quad 0 = \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(x, y) \\ \phi_2 \in \pi_2(y, z) \\ \mu(\phi_1) = \mu(\phi_2) = 1}} (a(\lambda_1, \phi_1)a(\lambda_2, \phi_2) + a(\lambda_1, \phi_2)a(\lambda_2, \phi_1)) \cdot \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_w^{n_w(\phi_1) + n_w(\phi_2)} \cdot z$$

$$+ \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(x, y) \\ \phi_2 \in \pi_2(y, z) \\ \mu(\phi_1) = \mu(\phi_2) = 1}} (a(\lambda_1, \phi_1)a(\lambda_2, \phi_1) + a(\lambda_1, \phi_2)a(\lambda_2, \phi_2)) \cdot \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_w^{n_w(\phi_1) + n_w(\phi_2)} \cdot z.$$

The right side of (4-9) is the z -component of

$$(A_{\lambda_1} A_{\lambda_2} + A_{\lambda_2} A_{\lambda_1} + \partial H_{\lambda_1 \lambda_2} + H_{\lambda_1 \lambda_2} \partial)(x),$$

where $H_{\lambda_1 \lambda_2}$ is the map

$$H_{\lambda_1 \lambda_2}(x) = \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi \in \pi_2(x, y) \\ \mu(\phi) = 1}} a(\lambda_1, \phi)a(\lambda_2, \phi) \cdot \#\widehat{\mathcal{M}}(\phi) U_w^{n_w(\phi)} \cdot y.$$

If $x = z$, then boundary degenerations may appear in the ends of $\widehat{\mathcal{M}}(\phi)$. These make an additional contribution to (4-8) of

$$(4-10) \quad \sum_{w \in \mathbf{w}} a(\lambda_1, A_w^x) a(\lambda_2, A_w^x) \cdot \#\widehat{\mathcal{N}}^\alpha(A_w^x) U_w \cdot x + \sum_{w \in \mathbf{w}} a(\lambda_1, B_w^x) a(\lambda_2, B_w^x) \cdot \#\widehat{\mathcal{N}}^\beta(B_w^x) U_w \cdot x.$$

We note that $a(\lambda, B_w^x) = 0$ for all w , and

$$(4-11) \quad a(\lambda, A_w^x) = \begin{cases} 1 & \text{if } w \in \partial\lambda, \\ 0 & \text{if } w \notin \partial\lambda. \end{cases}$$

The stated formula now follows from (4-10), (4-11) and the count of boundary degenerations from (4-6). \square

In the context of a single basepoint in Heegaard Floer homology, the homology action squares to zero (see [Ozsváth and Szabó 2004b, Proposition 4.17; Lipshitz 2006, Proposition 8.6]). We prove a similar result:

Lemma 4.5 *Suppose λ is an immersed path in Σ from w_1 to w_2 . Then*

$$A_\lambda \circ A_\lambda \simeq U_{w_1} \simeq U_{w_2}.$$

If γ is an immersed closed curve in Σ , then

$$A_\gamma \circ A_\gamma \simeq 0.$$

Proof Note that $U_{w_1} \simeq U_{w_2}$ by Lemma 4.1. We focus on the claim that $A_\lambda \circ A_\lambda \simeq U_{w_1}$. The proof of the relation $A_\gamma \circ A_\gamma \simeq 0$ follows with only minor modification. Pick an orientation of λ , which gives a lift of $a(\lambda, \phi)$ from \mathbb{F}_2 to \mathbb{Z} .

Note that both $a(\lambda, A_{w_1}^x)$ and $a(\lambda, A_{w_2}^x)$ are ± 1 , and, in fact, the two quantities have opposite signs. Orient λ so that

$$(4-12) \quad a(\lambda, A_{w_1}^x) = 1 \quad \text{and} \quad a(\lambda, A_{w_2}^x) = -1.$$

Define the map

$$H_\lambda(x) = \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi \in \pi_2(x, y) \\ \mu(\phi) = 1}} \frac{1}{2} a(\lambda, \phi) (a(\lambda, \phi) + 1) \cdot \#\widehat{\mathcal{M}}(\phi) U_w^{n_w(\phi)} \cdot y.$$

We will show that

$$(4-13) \quad A_\lambda^2(x) = (\partial H_\lambda + H_\lambda \partial + U_{w_1})(x).$$

If $z \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, the z -component of $(\partial H_\lambda + H_\lambda \partial + A_\lambda^2)(x)$ is

$$(4-14) \quad \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(x, y) \\ \phi_2 \in \pi_2(y, z) \\ \mu(\phi_1) = \mu(\phi_2) = 1}} \left(\frac{1}{2} a(\lambda, \phi_1) (a(\lambda, \phi_1) + 1) + \frac{1}{2} a(\lambda, \phi_2) (a(\lambda, \phi_2) + 1) \right) \cdot \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_w^{n_w(\phi_1 + \phi_2)}.$$

Rearranging, (4-14) becomes

$$(4-15) \quad \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(x, y) \\ \phi_2 \in \pi_2(y, z) \\ \mu(\phi_1) = \mu(\phi_2) = 1}} \frac{1}{2} a(\lambda, \phi_1 + \phi_2) (a(\lambda, \phi_1 + \phi_2) + 1) \cdot \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_w^{n_w(\phi_1 + \phi_2)} \\ + \sum_{\substack{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(x, y) \\ \phi_2 \in \pi_2(y, z) \\ \mu(\phi_1) = \mu(\phi_2) = 1}} a(\lambda, \phi_1) a(\lambda, \phi_2) \cdot \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_2) U_w^{n_w(\phi_1 + \phi_2)}.$$

The second summand of (4-15) is the z coefficient of $(A_\lambda \circ A_\lambda)(x)$. If $x \neq z$, then the first summand of (4-15) is equal to

$$(4-16) \quad \sum_{\substack{\phi \in \pi_2(x, z) \\ \mu(\phi) = 2}} \frac{1}{2} a(\lambda, \phi) (a(\lambda, \phi) + 1) \cdot \#\overline{\widehat{\mathcal{M}}(\phi)} U_w^{n_w(\phi)} = 0.$$

Hence, the z -component of $(H_\lambda \partial + dH_\lambda + A_\lambda^2)(x)$ is zero, when $x \neq z$.

We now consider the case that $z = x$. The expression in (4-16) is equal to $(\partial H_\lambda + H_\lambda \partial + A_\lambda^2)(x)$ plus the following contribution, due to boundary degenerations:

$$\sum_{w \in \mathbf{w}} \frac{1}{2} a(\lambda, A_w^x) (a(\lambda, A_w^x) + 1) \cdot \#\widehat{\mathcal{N}}^\alpha(A_w^x) U_w \cdot x + \sum_{w \in \mathbf{w}} \frac{1}{2} a(\lambda, B_w^x) (a(\lambda, B_w^x) + 1) \cdot \#\widehat{\mathcal{N}}^\beta(B_w^x) U_w \cdot x.$$

Finally, we note that $a(\lambda, B_w^{\mathbf{x}}) = 0$ for all $w \in \mathbf{w}$, while $a(\lambda, A_w^{\mathbf{x}}) = 0$ unless $w \in \{w_1, w_2\}$. Furthermore, from (4-12),

$$\frac{1}{2}a(\lambda, A_{w_1}^{\mathbf{x}})(a(\lambda, A_{w_1}^{\mathbf{x}}) + 1) = 1 \quad \text{and} \quad \frac{1}{2}a(\lambda, A_{w_2}^{\mathbf{x}})(a(\lambda, A_{w_2}^{\mathbf{x}}) + 1) = 0.$$

It follows that the \mathbf{x} -component of $(A_\lambda^2 + \partial H_\lambda + H_\lambda \partial)(\mathbf{x})$ is $U_{w_1} \cdot \mathbf{x}$, completing the proof. □

Since the Heegaard surface Σ is embedded in Y , there is a natural map $i_*: H_1(\Sigma; \mathbb{Z}) \rightarrow H_1(Y; \mathbb{Z})$. Since Y can be built by attaching 2-handles and 3-handles to $[0, 1] \times \Sigma$, the map i_* is surjective, and $\ker(i_*) = \text{Span}\{[\alpha_1], \dots, [\alpha_n], [\beta_1], \dots, [\beta_n]\}$. Hence,

$$(4-17) \quad H_1(Y; \mathbb{Z}) \cong \frac{H_1(\Sigma; \mathbb{Z})}{\text{Span}\{[\alpha_1], \dots, [\alpha_n], [\beta_1], \dots, [\beta_n]\}}.$$

Lemma 4.6 *If γ is a closed loop on Σ such that*

$$i_*([\gamma]) = 0 \in H_1(Y; \mathbb{Z})/\text{Tors},$$

then

$$A_\gamma \simeq 0.$$

Proof Our proof is a modification of Ni’s proof [2014, Lemma 2.4] of a closely related result. Suppose that γ is an integral 1-cycle on Σ such that $i_*(k \cdot \gamma) = 0$, for some integer $k \neq 0$.

By orienting γ , we obtain a lift of the quantities $a(\gamma, \phi)$ to \mathbb{Z} . We can assume that γ is immersed, intersects α and β transversely, and is disjoint from any intersections of the α and β curves. From the isomorphism in (4-17), it follows that the class $k \cdot [\gamma] \in H_1(\Sigma; \mathbb{Z})$ can be written as an integral combination of the α and β curves.

Consequently, there is an integral 2-chain S on Σ such that

$$\partial S = k \cdot \gamma + C,$$

where C is an integral 1-cycle on Σ which consists of an integral combination of small pushoffs of the α and β curves on Σ .

Let $n_S: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \mathbb{Z}$ denote the function

$$n_S(\mathbf{x}) = \left(\sum_{x \in \mathbf{x}} n_x(S) \right).$$

We claim that if $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, then

$$(4-18) \quad -a(k \cdot \gamma + C, \phi) = n_S(\mathbf{y}) - n_S(\mathbf{x}).$$

To establish (4-18), we note that the homology class ϕ determines a collection of paths of $g + |\mathbf{w}| - 1$ arcs $\mathbf{a}(\phi)$, which are contained in the α curves and run from the points of \mathbf{x} to the points of \mathbf{y} . The quantity $a(k \cdot \gamma + C, \phi)$ can be reinterpreted as the oriented intersection number

$$\#(\mathbf{a}(\phi) \cap (k \cdot \gamma + C)).$$

On the other hand, $\mathbf{a}(\phi) \cap S$ is a compact 1-manifold, and hence the algebraic count of its ends is zero. By the Leibniz rule,

$$0 = \#\partial(\mathbf{a}(\phi) \cap S) = \#(\partial\mathbf{a}(\phi) \cap S) + \#(\mathbf{a}(\phi) \cap \partial S),$$

from which (4-18) follows. Using additivity of (γ, ϕ) with respect to γ , from (4-18) we obtain

$$(4-19) \quad -k \cdot a(\gamma, \phi) - a(C, \phi) = n_S(\mathbf{y}) - n_S(\mathbf{x}).$$

If α' is a small pushoff of an α curve, then $a(\alpha', \phi) = 0$, since there are no changes of ϕ across any α curves as one traverses α' . Similarly, if β' is a small pushoff of a β curve, then $a(\beta', \phi) = 0$, since β' intersects no β curves, so the sum of differences of ϕ across the α curves as one traverses β' telescopes, and is zero. With this observation, equation (4-19) now reads

$$(4-20) \quad -k \cdot a(\gamma, \phi) = n_S(\mathbf{y}) - n_S(\mathbf{x}).$$

The right-hand side of (4-20) does not depend on ϕ . It follows that, modulo k , the expression $n_S(\mathbf{x})$ is independent of \mathbf{x} for \mathbf{x} representing a fixed Spin^c structure. By adding copies of $[\Sigma]$ to S , we can ensure that $n_S(\mathbf{x})$ is divisible by k for each $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ representing \mathfrak{s} . Hence,

$$(4-21) \quad -a(\gamma, \phi) = \frac{n_S(\mathbf{y})}{k} - \frac{n_S(\mathbf{x})}{k}.$$

We define the map

$$H_S: \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s})$$

via the formula

$$H_S(\mathbf{x}) = \frac{n_S(\mathbf{x})}{k} \cdot \mathbf{x},$$

extended equivariantly over $\mathbb{F}_2[U_{\mathbf{w}}]$.

After projecting to \mathbb{F}_2 , equation (4-21) implies that

$$A_\gamma(\mathbf{x}) = (\partial H_S + H_S \partial)(\mathbf{x}). \quad \square$$

We now describe a simple relation between the maps A_λ and B_λ . To describe their relation, we must introduce a new map. If $w \in \mathbf{w}$ is a basepoint, we define $\Phi_w: \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s})$ via the formula

$$(4-22) \quad \Phi_w(\mathbf{x}) := U_w^{-1} \cdot \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} n_w(\phi) \cdot \#\widehat{\mathcal{M}}(\phi) U_w^{n_w(\phi)} \cdot \mathbf{y},$$

extended $\mathbb{F}_2[U_{\mathbf{w}}]$ -equivariantly. Despite the initial factor of U_w^{-1} , the map Φ_w maps CF^- into CF^- . Furthermore, Φ_w is a chain map; see Lemma 13.13 below.

Lemma 4.7 *Suppose that $(\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) , λ is an immersed curve on Σ with endpoints w_1 and w_2 , and γ is an immersed closed curve on Σ . Then*

- (1) $A_\gamma = B_\gamma$, and
- (2) $A_\lambda + B_\lambda = U_{w_1} \Phi_{w_1} + U_{w_2} \Phi_{w_2}$.

Proof The map A_γ counts holomorphic strips weighted by the factor $a(\gamma, \phi)$, while B_γ counts holomorphic disks weighted by $b(\gamma, \phi)$. The sum $a(\gamma, \phi) + b(\gamma, \phi)$ is the total change in multiplicity of ϕ across all curves (either α or β); however, this is zero since γ is a closed curve. Part (1) follows.

For an arc λ connecting basepoints w_1 and w_2 , we instead have

$$a(\lambda, \phi) + b(\lambda, \phi) = n_{w_1}(\phi) - n_{w_2}(\phi),$$

from which part (2) follows. □

4.3 Naturality of the relative homology action

In this section, we prove the following:

Proposition 4.8 *Suppose that $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) , and λ and λ' are two immersed paths in Σ from w_1 to w_2 which represent the same element in $H_1(Y, \{w_1, w_2\}; \mathbb{Z})/\text{Tors}$. Then*

$$A_\lambda \simeq A_{\lambda'}.$$

Furthermore, if \mathcal{H} and \mathcal{H}' are two diagrams for (Y, \mathbf{w}) , then

$$A_\lambda \circ \Psi_{\mathcal{H} \rightarrow \mathcal{H}'} \simeq \Psi_{\mathcal{H} \rightarrow \mathcal{H}'} \circ A_\lambda.$$

The same holds for the map B_λ , as well as the homology action associated to closed loops in Y .

As a first step towards Proposition 4.8, we prove that the maps A_λ commute with the change of almost complex structure maps:

Lemma 4.9 *Suppose J and J' are two cylindrical almost complex structures on $\Sigma \times [0, 1] \times \mathbb{R}$ which satisfy axioms (J1)–(J5). Then*

$$A_\lambda \circ \Psi_{J \rightarrow J'} \simeq \Psi_{J \rightarrow J'} \circ A_\lambda,$$

where $\Psi_{J \rightarrow J'}$ is the change of almost complex structures map. The same relation holds for the type-B maps.

Proof To compute $\Psi_{J \rightarrow J'}$, one first picks a noncylindrical almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$ which agrees with J on $\Sigma \times [0, 1] \times (-\infty, -1]$ and agrees with J' on $\Sigma \times [0, 1] \times [1, \infty)$. The map $\Psi_{J \rightarrow J'}$ is defined via the formula

$$\Psi_{J \rightarrow J'}(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=0}} \# \mathcal{M}_{\tilde{J}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y},$$

extended linearly over $\mathbb{F}_2[U_{\mathbf{w}}]$.

We define the map $H_{\tilde{J}, \lambda}^A : \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s}) \rightarrow \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s})$ via the formula

$$H_{\tilde{J}, \lambda}^A(\mathbf{x}) := \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=0}} a(\lambda, \phi) \cdot \# \mathcal{M}_{\tilde{J}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y},$$

extended linearly over $\mathbb{F}_2[U_{\mathbf{w}}]$.

If $\phi \in \pi_2(x, y)$ is a Maslov index 1 homology class, then the 1-dimensional moduli space $\mathcal{M}_{\tilde{J}}(\phi)$ can be compactified into a compact 1-manifold, whose ends correspond to pairs of index 0 \tilde{J} -holomorphic curves, and index 1 J - or J' -holomorphic curves.

Using (4-4) (additivity of the quantity $a(\lambda, \phi)$ with respect to ϕ), it follows that

$$A_\lambda \circ \Psi_{J \rightarrow J'} + \Psi_{J \rightarrow J'} \circ A_\lambda + \partial_{J'} \circ H_{\tilde{J}, \lambda}^A + H_{\tilde{J}, \lambda}^A \circ \partial_J = 0.$$

The relation for the type- B homology actions is proved analogously. □

Towards proving that A_λ commutes with changes of the α curves, we prove the following:

Lemma 4.10 *Suppose that $(\Sigma, \alpha', \alpha, \beta, w)$ is a Heegaard triple such that $(\Sigma, \alpha', \alpha, w)$ is a diagram for $(S^1 \times S^2)^{\#k}$ for some k . Suppose that λ is an immersed curve in Σ , with endpoints w_1 and w_2 , or λ is an immersed closed curve. Let A'_λ, A_λ and B_λ denote the relative homology maps defined by counting changes over the α', α or β curves, respectively. Suppose that $\Theta_{\alpha', \alpha}^+ \in \text{CF}^-(\Sigma, \alpha', \alpha, w, \mathfrak{s}_0)$ is a cycle which represents the top-degree generator of homology, and $\mathfrak{s} \in \text{Spin}^c(X_{\alpha', \alpha, \beta})$ restricts to \mathfrak{s}_0 on $Y_{\alpha', \alpha}$. Then:*

- (1) $F_{\alpha', \alpha, \beta, \mathfrak{s}}(\Theta_{\alpha', \alpha}^+ \otimes A_\lambda(-)) \simeq A'_\lambda \circ F_{\alpha', \alpha, \beta, \mathfrak{s}}(\Theta_{\alpha', \alpha}^+ \otimes -)$.
- (2) $F_{\alpha', \alpha, \beta, \mathfrak{s}}(\Theta_{\alpha', \alpha}^+ \otimes B_\lambda(-)) \simeq B_\lambda \circ F_{\alpha', \alpha, \beta, \mathfrak{s}}(\Theta_{\alpha', \alpha}^+ \otimes -)$.

An analogous statement holds for triples $(\Sigma, \alpha, \beta, \beta', w)$ where $(\Sigma, \beta, \beta', w)$ is a diagram for $(S^1 \times S^2)^{\#k}$.

Proof We focus on the claim when λ is an immersed path connecting w_1 and w_2 . The claim when λ is an immersed closed curve is a simple modification.

Part (2) follows immediately from Lemma 4.2, though part (1) does not follow from a symmetric argument.

By using Lemmas 4.2 and 4.7, we see

$$\begin{aligned} (4-23) \quad F_{\alpha', \alpha, \beta, \mathfrak{s}}(\Theta_{\alpha', \alpha}^+ \otimes A_\lambda(-)) + A'_\lambda \circ F_{\alpha', \alpha, \beta, \mathfrak{s}}(\Theta_{\alpha', \alpha}^+ \otimes -) \\ \simeq F_{\alpha', \alpha, \beta, \mathfrak{s}}((A'_\lambda + A_\lambda)(\Theta_{\alpha', \alpha}^+) \otimes -) \\ \simeq F_{\alpha', \alpha, \beta, \mathfrak{s}}((U_{w_1} \Phi_{w_1} + U_{w_2} \Phi_{w_2})(\Theta_{\alpha', \alpha}^+) \otimes -). \end{aligned}$$

However, from (4-22) we see that the maps Φ_{w_i} are $+1$ -graded chain maps. Since $[\Theta_{\alpha', \alpha}^+]$ is the highest-graded nonzero element of $\text{HF}^-(\Sigma, \alpha', \alpha, w, \mathfrak{s}_0)$, we must have

$$[\Phi_{w_1}(\Theta_{\alpha', \alpha}^+)] = [\Phi_{w_2}(\Theta_{\alpha', \alpha}^+)] = 0 \in \text{HF}^-(\Sigma, \alpha', \alpha, w, \mathfrak{s}_0).$$

The associativity relations for holomorphic triangles imply that the map $F_{\alpha', \alpha, \beta, \mathfrak{s}}(\partial\eta, -)$ is chain homotopic to the zero map for any $\eta \in \text{CF}^-(\Sigma, \alpha, \alpha', w, \mathfrak{s}_0)$. Hence,

$$(4-24) \quad F_{\alpha', \alpha, \beta, \mathfrak{s}}((U_{w_1} \Phi_{w_1} + U_{w_2} \Phi_{w_2})(\Theta_{\alpha', \alpha}^+) \otimes -) \simeq 0.$$

Equations (4-23) and (4-24) imply part (1). □

Corollary 4.11 *Suppose Σ is a Heegaard surface for (Y, \mathbf{w}) , and $\lambda \subseteq \Sigma$ is either an immersed, closed curve, or an immersed path connecting two basepoints. If α and α' are attaching curves for the α -handlebody, and β and β' are attaching curves for the β -handlebody, then*

- (1) $\Psi_{\alpha \rightarrow \alpha'}^{\beta} \circ A_{\lambda} \simeq A_{\lambda} \circ \Psi_{\alpha \rightarrow \alpha'}^{\beta}$, and
- (2) $\Psi_{\alpha}^{\beta \rightarrow \beta'} \circ A_{\lambda} \simeq A_{\lambda} \circ \Psi_{\alpha}^{\beta \rightarrow \beta'}$.

The same holds for the type-B maps.

Proof The transition maps $\Psi_{\alpha \rightarrow \alpha'}^{\beta}$ and $\Psi_{\alpha}^{\beta \rightarrow \beta'}$ can both be computed via a sequence of holomorphic triangle maps, so the result follows from Lemma 4.10. □

Next we consider the transition maps associated to simple stabilizations of the Heegaard surface.

Lemma 4.12 *Suppose that \mathcal{H}' is a simple stabilization of \mathcal{H} , and let σ denote the transition map from $CF_J^-(\mathcal{H}, \mathfrak{s})$ to $CF_{J(T)}^-(\mathcal{H}', \mathfrak{s})$. If $\lambda \subseteq \Sigma$ is an immersed, closed loop or path connecting two basepoints, then*

$$A_{\lambda} \circ \sigma = \sigma \circ A_{\lambda}.$$

Proof The proof follows from the same count of holomorphic curves used to prove stabilization-invariance [Lipshitz 2006, Proposition 12.5]. If k is an integer, let $\phi_k \in \pi_2(c, c)$ denote the class which has multiplicity k on the single domain of $(\mathbb{T}^2, \alpha_0, \beta_0)$. The class ϕ_k has Maslov index $2k$. If $\phi \in \pi_2(x, y)$ is a homology class of disks on \mathcal{H} , with multiplicity k at the connected sum point, then using Lipshitz’s formula [ibid., equation 8] for the Maslov index together with the fact that a disk has Euler measure 1, we obtain

$$\mu(\phi \# \phi_k) = \mu(\phi).$$

Furthermore, any homology class of disks on \mathcal{H}' can be written as such a connected sum.

According to the proof of [ibid., Proposition 12.5], for sufficiently large T , one has

$$(4-25) \quad \#\widehat{\mathcal{M}}_J(\phi) = \#\widehat{\mathcal{M}}_{J(T)}(\phi \# \phi_k).$$

A straightforward computation shows that

$$(4-26) \quad a(\lambda, \phi) = a(\lambda, \phi \# \phi_k).$$

The claim follows immediately from (4-25) and (4-26). □

We now prove well-definedness of the relative homology actions:

Proof of Proposition 4.8 Suppose that λ and λ' are paths from w_1 to w_2 in Y . The claim that $A_{\lambda} \simeq A_{\lambda'}$ if λ and λ' represent homologous elements of $H_1(Y, \{w_1, w_2\}; \mathbb{Z})/\text{Tors}$ is proven in Lemma 4.6.

To show that the map A_{λ} commutes with the transition maps up to chain homotopy, it is sufficient to show that A_{λ} commutes with the transition maps associated to changes of the almost complex structure, as well as each elementary Heegaard move from Lemma 3.5. Commutation of the relative homology maps with the transition maps associated to changing the almost complex structure is proven in Lemma 4.9.

Commutation with the maps associated to isotopies and handleslides of the α and β curves is proven in Corollary 4.11. Commutation with the simple stabilization maps is proven in Lemma 4.12. Commutation with the maps induced by isotopies of the Heegaard surface inside of Y is tautological. \square

5 Free-stabilization maps

In this section we describe maps for adding or removing a basepoint, which we call the *free-stabilization* maps. Suppose (Y, \mathbf{w}) is a multipointed 3-manifold, $w \notin \mathbf{w}$, and

$$\sigma: \mathbf{w} \rightarrow P \quad \text{and} \quad \sigma': \mathbf{w} \cup \{w\} \rightarrow P$$

are colorings satisfying $\sigma'|_{\mathbf{w}} = \sigma$. In this section, we describe homomorphisms of \mathcal{R}_P -modules

$$S_w^+: CF^-(Y, \mathbf{w}^\sigma, \mathfrak{s}) \rightarrow CF^-(Y, (\mathbf{w} \cup \{w\})^{\sigma'}, \mathfrak{s}) \quad \text{and} \quad S_w^-: CF^-(Y, (\mathbf{w} \cup \{w\})^{\sigma'}, \mathfrak{s}) \rightarrow CF^-(Y, \mathbf{w}^\sigma, \mathfrak{s}).$$

Since the maps S_w^+ and S_w^- are \mathcal{R}_P -equivariant, they induce maps on the $+$, ∞ and \wedge flavors as well, by tensoring with the identity map (see (3-1)).

We now state the formula defining the free-stabilization maps. Suppose that $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) such that $w \in \Sigma \setminus (\alpha \cup \beta)$. Pick a small disk $D \subseteq \Sigma \setminus (\alpha \cup \beta)$ containing the point w . Pick two curves α_0 and β_0 , inside of D , such that

$$|\alpha_0 \cap \beta_0| = 2$$

and both α_0 and β_0 bound a disk containing w . The two intersection points of $\alpha_0 \cap \beta_0$ are distinguished by their relative Maslov index. Let θ^+ and θ^- denote the higher- and lower-graded intersection points, respectively. See Figure 5.

For appropriately chosen almost complex structures on $\Sigma \times [0, 1] \times \mathbb{R}$ (described precisely in Section 5.1), the maps S_w^+ and S_w^- are defined via the formulas

$$(5-1) \quad S_w^+(\mathbf{x}) = \mathbf{x} \times \theta^+, \quad S_w^-(\mathbf{x} \times \theta^-) = \mathbf{x} \quad \text{and} \quad S_w^-(\mathbf{x} \times \theta^+) = 0,$$

extended \mathcal{R}_P -equivariantly.




Figure 5: A free stabilization. We can think of the dashed circle as the connected sum tube. Alternatively, if we collapse the dashed circle to a point p_0 , we get a doubly pointed diagram on S^2 .

5.1 Gluing data for stretching the neck

In this section, we describe precisely which almost complex structures we use to define the free-stabilization maps.

It is convenient to view the free-stabilization operation as taking the connected sum of \mathcal{H} with the diagram $(S^2, \alpha_0, \beta_0, w_0, p_0)$, at the points w and p_0 .

Fix an embedded disk $D_0 \subseteq S^2 \setminus (\alpha_0 \cup \beta_0)$, centered at p_0 . We make the following definition to parametrize the construction of a stretched almost complex structure:

Definition 5.1 Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) and $w \in \Sigma \setminus (\mathbf{w} \cup \alpha \cup \beta)$. We call a tuple $\mathfrak{d} = (J^\mathfrak{d}, J_0^\mathfrak{d}, D, \iota)$ a *gluing datum for free-stabilizing at w* if the following hold:

- (1) $D \subseteq \Sigma \setminus (\alpha \cup \beta)$ is a closed disk containing w .
- (2) $J^\mathfrak{d}$ is an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ which is split on D .
- (3) $J_0^\mathfrak{d}$ is an almost complex structure on $S^2 \times [0, 1] \times \mathbb{R}$ which is split on D_0 .
- (4) $\iota: S^2 \setminus (\frac{1}{2} \cdot D_0) \rightarrow D$ is an embedding which maps w_0 to w . Furthermore, ι maps the annulus $D_0 \setminus (\frac{1}{2} \cdot D_0)$ conformally onto $D \setminus (\frac{1}{2} \cdot D)$. Here $\frac{1}{2} \cdot D_0$ denotes the subdisk of radius $\frac{1}{2}$ centered at p_0 , obtained from the unique (up to rotation) conformal identification of the pair (D_0, p_0) with $(\{z \in \mathbb{C} : |z| \leq 1\}, 0)$.

Given a gluing datum \mathfrak{d} for free-stabilizing at w , we can form a diagram

$$\mathcal{H}^+ := (\Sigma, \alpha \cup \iota(\alpha_0), \beta \cup \iota(\beta_0), \mathbf{w} \cup \{w\}),$$

which depends on the choice of \mathfrak{d} through the embedding ι (though we suppress this from the notation).

We define the stretched almost complex structure $J^\mathfrak{d}(T)$ on $\Sigma \times [0, 1] \times \mathbb{R}$, whenever $T > 0$. We begin by defining $J^\mathfrak{d}(T)$ when $T = 2$, by letting $J^\mathfrak{d}(2)$ coincide with $J^\mathfrak{d}$ and $J_0^\mathfrak{d}$ on $(\Sigma \setminus (\frac{1}{2} \cdot D)) \times [0, 1] \times \mathbb{R}$ and $(S^2 \setminus (\frac{1}{2} \cdot D_0)) \times [0, 1] \times \mathbb{R}$, respectively. For $T \geq 2$, we construct an almost complex structure $J^\mathfrak{d}(T)$ by replacing the almost complex structure in the annulus region $D \setminus (\frac{1}{2} \cdot D)$ with one conformally equivalent to the annulus $D \setminus \frac{1}{T} \cdot D$. We set $J^\mathfrak{d}(T) = J^\mathfrak{d}(2)$ if $2 \geq T > 0$.

Recall that if J and J' are two cylindrical almost complex structures on $\Sigma \times [0, 1] \times \mathbb{R}$, the transition map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}, J')}$ can be computed by picking a noncylindrical almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$ which agrees with J on $\Sigma \times [0, 1] \times (-\infty, -1]$ and with J' on $\Sigma \times [0, 1] \times [1, \infty)$. The transition map is defined by counting index 0 \tilde{J} -holomorphic curves via the formula

$$(5-2) \quad \Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}, J')}(\mathbf{x}) := \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 0}} \# \mathcal{M}_{\tilde{J}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}.$$

Write $\Psi_{\tilde{J}}$ for the map $\Psi_{(\mathcal{H}, J) \rightarrow (\mathcal{H}, J')}$ appearing in (5-2), computed using \tilde{J} .

Definition 5.2 Suppose \mathcal{H} is a Heegaard diagram and \mathfrak{d} is a gluing datum for free-stabilizing at w . We say that a real number $T > 0$ satisfies *stabilizing condition (SC2)* if for any two $T_1, T_2 \geq T$, there is a

noncylindrical almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$ interpolating $J^\partial(T_1)$ and $J^\partial(T_2)$, such that for all $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we have

$$(SC2) \quad \Psi_{\tilde{J}}(\mathbf{x} \times \theta^+) = \mathbf{x} \times \theta^+ \quad \text{and} \quad \Psi_{\tilde{J}}(\mathbf{x} \times \theta^-) = \mathbf{x} \times \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x}, \mathbf{y}} \cdot \mathbf{y} \times \theta^+$$

for some $C_{\mathbf{x}, \mathbf{y}} \in \mathbb{F}_2[U_w]$ (which may depend on ∂ , T_1 and T_2).

We define

$$S_w^+ : CF_{J^\partial}^-(\mathcal{H}, \sigma, \mathfrak{s}) \rightarrow CF_{J^\partial(T)}^-(\mathcal{H}^+, \sigma', \mathfrak{s}),$$

using (5-1) whenever T satisfies condition (SC2). The map S_w^- is also defined using (5-1), under the same assumption.

If J' is an arbitrary almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$, then the map S_w^+ from $CF_{J^\partial}^-(\mathcal{H}, \sigma, \mathfrak{s})$ to $CF_{J'}^-(\mathcal{H}^+, \sigma', \mathfrak{s})$ is defined as the composition of the map in (5-1) together with the transition map $\Psi_{J^\partial(T) \rightarrow J'}$.

Proposition 5.3 *If ∂ is a gluing datum for free-stabilizing at w , then there is a $T > 0$ which satisfies stabilizing condition (SC2).*

Before proving Proposition 5.3, we prove a Maslov index formula:

Lemma 5.4 *Let $(S^2, \alpha_0, \beta_0, w, p_0)$ denote the diagram in Figure 5. If $x, y \in \alpha_0 \cap \beta_0$ and $\phi_0 \in \pi_2(x, y)$ is a homology class of disks, then*

$$\mu(\phi_0) = 2n_w(\phi_0) + 2n_{p_0}(\phi_0) + \text{gr}(x, y),$$

where $\text{gr}(x, y)$ denotes the relative Maslov grading between x and y .

Furthermore, if $m_1(\phi_0), m_2(\phi_0), m_3(\phi_0)$ and $m_4(\phi_0)$ denote the multiplicities of ϕ_0 in the components of $S^2 \setminus (\alpha_0 \cup \beta_0)$, then

$$\mu(\phi_0) = m_1(\phi_0) + m_2(\phi_0) + m_3(\phi_0) + m_4(\phi_0).$$

Proof The first formula is equivalent to

$$\text{gr}(x, y) = \mu(\phi) - 2n_w(\phi) - 2n_{p_0}(\phi),$$

which is the definition of the relative Maslov grading $\text{gr}(x, y)$.

To prove the second formula, we verify it for a constant homology class $e_x \in \pi_2(x, x)$ (for which the claim is trivial), and then note that it respects slicing in bigons. Since any two classes on this diagram can be related by splicing in bigons, the formula follows in general. \square

Proof of Proposition 5.3 We focus on the claim that if T_1 and T_2 satisfy (SC2), then \tilde{J} can be chosen so that

$$\Psi_{\tilde{J}}(\mathbf{x} \times \theta^+) = \mathbf{x} \times \theta^+.$$

The claim about $\mathbf{x} \times \theta^-$ (which is the dual statement) follows by a simple modification.

We will write p for the point w , viewed as a point on \mathcal{H} , and write w for the new basepoint on \mathcal{H}^+ .

Suppose $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times y)$ is a homology class of disks with Maslov index 0 for $y \in \{\theta^+, \theta^-\}$. We will show that for T_1 and T_2 sufficiently large, \tilde{J} can be chosen so that if $\phi \# \phi_0$ has a \tilde{J} -holomorphic representative, then $\phi \# \phi_0$ is the constant class, $e_{\mathbf{x}} \times e_{\theta^+} \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{x} \times \theta^+)$. Furthermore, we will show that $e_{\mathbf{x}} \times e_{\theta^+}$ always has a unique representative, which will imply the statement.

Suppose that $T_{1,i}$ and $T_{2,i}$ are sequences of neck lengths which both approach $+\infty$. We can pick a sequence of interpolating almost complex structures \tilde{J}_i such that $(\Sigma \times [0, 1] \times \mathbb{R}, \tilde{J}_i)$ contains the almost complex manifold $((\Sigma \setminus N_i) \times [0, 1] \times \mathbb{R}, J)$, where N_i is a nested sequence of open balls on Σ whose intersection is $\{p\}$.

Suppose u_i is a sequence of \tilde{J}_i -holomorphic curves representing $\phi \# \phi_0$. By adapting [Lipshitz 2006, Proposition 12.4], we can extract a broken limiting curve on the punctured manifold $(\Sigma \setminus \{p\}) \times [0, 1] \times \mathbb{R}$. Such a holomorphic curve can be completed over $\{p\} \times [0, 1] \times \mathbb{R}$ to obtain a (potentially broken) representative \mathcal{U} of the homology class ϕ on $(\Sigma, \alpha, \beta, w)$.

We have

$$(5-3) \quad \mu(\phi \# \phi_0) = \mu(\phi) + \mu(\phi_0) - 2n_{p_0}(\phi_0) = \mu(\phi) + \text{gr}(\theta^+, y) + 2n_w(\phi_0).$$

The first equality of (5-3) is justified by Lipshitz’s formula for the Maslov index [ibid., equation 8], together with the fact that a disk has Euler measure 1. The second equality follows from Lemma 5.4.

Since ϕ admits a broken representative for J , we conclude that $\mu(\phi) \geq 0$ by transversality. Since the last line of (5-3) involves only nonnegative terms, and the sum is zero, we conclude that

$$\mu(\phi) = \text{gr}(\theta^+, y) = n_w(\phi_0) = 0.$$

Since $\mu(\phi) = 0$ and ϕ admits a broken J -holomorphic representative, and J is cylindrical, it follows that $\mathbf{x} = \mathbf{y}$ and ϕ is the constant class, $e_{\mathbf{x}}$, by transversality. Since $\text{gr}(\theta^+, y) = 0$, it follows that $y = \theta^+$. Since $n_{p_0}(\phi_0) = n_w(\phi_0) = 0$, as well, ϕ_0 is the constant class e_{θ^+} .

Conversely, $e_{\mathbf{x}} \times e_{\theta^+}$ admits a unique \tilde{J}_i -holomorphic representative for any i , since each \tilde{J}_i is cylindrical in a neighborhood of $(\mathbf{x} \times \{\theta^+\}) \times [0, 1] \times \mathbb{R}$. □

5.2 Free stabilization and the differential

In this section, we prove that the free-stabilization maps are chain maps. The argument is essentially the same as [Ozsváth and Szabó 2008, Proposition 6.5]. We repeat the argument since we will later prove several refinements and analogous holomorphic curve counts.

Proposition 5.5 *Suppose \mathcal{H} is a Heegaard diagram for (Y, \mathbf{w}) , \mathcal{H}^+ is its free stabilization at w , and \mathfrak{d} is a gluing datum for the free stabilization. Then, for all T which satisfy (SC2),*

$$(5-4) \quad \begin{aligned} \partial_{\mathcal{H}^+, J^\circ(T)}(\mathbf{x} \times \theta^+) &= \partial_{\mathcal{H}, J^\circ}(\mathbf{x}) \otimes \theta^+, \\ \partial_{\mathcal{H}^+, J^\circ(T)}(\mathbf{x} \times \theta^-) &= \partial_{\mathcal{H}, J^\circ}(\mathbf{x}) \otimes \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x}, \mathbf{y}} \cdot \mathbf{y} \otimes \theta^+ \end{aligned}$$

for $C_{\mathbf{x}, \mathbf{y}} \in \mathbb{F}_2[U_{\mathbf{w}}, U_w]$ (which depend on T and \mathfrak{d}).

Proof Since the transition maps $\Psi_{J^\circ(T_1) \rightarrow J^\circ(T_2)}$ are chain maps, it is easy to check that condition (SC2) algebraically implies that if (5-4) holds for some T which satisfies condition (SC2), then it also holds for all T which satisfy (SC2). Hence, it is sufficient to establish (5-4) for any sufficiently large T .

Equation (5-4) is implied by the following two subclaims:

- (d1) The $\mathbf{y} \times \theta^-$ coefficient of $\partial_{\mathcal{H}^+, J^\circ(T)}(\mathbf{x} \times \theta^+)$ vanishes whenever T is sufficiently large.
- (d2) If $\theta \in \{\theta^+, \theta^-\}$ and T is sufficiently large, then the $\mathbf{y} \times \theta$ coefficient of $\partial_{\mathcal{H}^+, J^\circ(T)}(\mathbf{x} \times \theta)$ (an element of $\mathbb{F}_2[U_w, U_w]$) is equal to the \mathbf{y} coefficient of $\partial_{\mathcal{H}, J^\circ}(\mathbf{x})$ (an element of $\mathbb{F}_2[U_w]$).

We will write p for the point w on Σ , viewed as a point on \mathcal{H} , and write w for the basepoint on the free-stabilized diagram \mathcal{H}^+ .

Let $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y)$ be a Maslov index 1 class. Pick a sequence of neck lengths T_i approaching $+\infty$, and consider a sequence u_i of $J^\circ(T_i)$ -holomorphic curves representing $\phi \# \phi_0$. For such a sequence, we can extract a broken limit consisting of collections \mathcal{U} , \mathcal{U}_m and \mathcal{U}_0 , where \mathcal{U} is a collection on $\Sigma \setminus \{p\} \times [0, 1] \times \mathbb{R}$ whose total class is ϕ , and \mathcal{U}_0 is a collection on $S^2 \setminus \{p_0\} \times [0, 1] \times \mathbb{R}$ whose total class is ϕ_0 . The collection \mathcal{U}_m consists of curves in the tube region $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$ (ultimately, we will rule out any nontrivial curves in \mathcal{U}_m , due to codimension considerations). The curves in \mathcal{U} and \mathcal{U}_0 may be completed over p and p_0 to obtain curves on the diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and (S^2, α_0, β_0) . The process of obtaining limiting curves is described in [Lipshitz 2006, Proposition 12.4].

Since ϕ has the broken holomorphic representative \mathcal{U} , it follows from Proposition 3.4 that

$$(5-5) \quad \mu(\phi) \geq 0.$$

On the other hand, equation (5-3) implies

$$(5-6) \quad \mu(\phi \# \phi_0) = \mu(\phi) + \text{gr}(x, y) + 2n_w(\phi_0).$$

We first consider subclaim (d1), when $x = \theta^+$ and $y = \theta^-$. In this case, we conclude from (5-6) that

$$\mu(\phi) = 0 \quad \text{and} \quad n_w(\phi_0) = 0.$$

By transversality, ϕ is the constant class e_x , and ϕ_0 is one of the two bigons in the free-stabilization region which have zero multiplicity over w . Both classes have unique representatives for any almost complex structure, and hence have canceling contribution to the differential. Subclaim (d1) is established.

We now consider subclaim (d2). In this case, equation (5-6) implies

$$\mu(\phi) = 1 \quad \text{and} \quad n_w(\phi_0) = 0.$$

For the moment, we trim off all ghost curves from \mathcal{U} , \mathcal{U}_m , and \mathcal{U}_0 , i.e., components of the limit which have constant image in $\Sigma \times [0, 1] \times \mathbb{R}$ or $S^2 \times [0, 1] \times \mathbb{R}$. (We will shortly prove that, generically, no ghost curves appear.)

Having trimmed off ghost curves, we claim that transversality is achieved at the remaining curves in \mathcal{U} . By Proposition 3.4, this amounts to showing that the limiting curves satisfy (M1)–(M5). The only axiom

which is nontrivial is (M5), i.e., that the limiting curves have no components v with $\pi_{[0,1] \times \mathbb{R}} \circ v$ constant. Having trimmed off ghost curves, we can assume that any curve v with $\pi_{[0,1] \times \mathbb{R}} \circ v$ constant has $\pi_\Sigma \circ v$ nonconstant. Such a curve v has Maslov index at least 2, since its domain must be a sum of connected components of $\Sigma \setminus \alpha$ and $\Sigma \setminus \beta$, each weighted with a nonnegative integer. If we delete from \mathcal{U} all such curves v , then we obtain a curve at which transversality is achieved, by Proposition 3.4. However, the Maslov index of the remaining components is at most -1 (since $\mu(\phi) = 1$ and we have removed curves whose total Maslov index is at least 2). There are no holomorphic curves with Maslov index -1 at which transversality is obtained, so such curves v are prohibited from appearing in \mathcal{U} .

It follows that \mathcal{U} (after trimming ghost curves) consists of a single curve $u: S \rightarrow \Sigma \times [0, 1] \times \mathbb{R}$ satisfying (M1)–(M5). Since ϕ has Maslov index 1, from Proposition 3.4 it also follows that u is embedded, and hence satisfies (M6).

Let

$$\rho^p: \mathcal{M}(\phi) \rightarrow \text{Sym}^{n_p(\phi)}([0, 1] \times \mathbb{R})$$

denote the map

$$\rho^p(u) := (u \circ \pi_{[0,1] \times \mathbb{R}})((u \circ \pi_\Sigma)^{-1}(p)).$$

Consider the 1-dimensional set

$$X(\phi) := \{\rho^p(u) : u \in \mathcal{M}(\phi)\} \subseteq \text{Sym}^{n_p(\phi)}([0, 1] \times \mathbb{R}).$$

By perturbing the almost complex structure slightly near p , we can assume that $X(\phi)$ is disjoint from the fat diagonal in $\text{Sym}^{n_p(\phi)}([0, 1] \times \mathbb{R})$, a codimension 2 subset.

We claim that \mathcal{U}_m consists of a union of $n_p(\phi)$ once-covered cylinders

$$u: S^1 \times \mathbb{R} \rightarrow S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R},$$

which each have constant projection to $[0, 1] \times \mathbb{R}$, together with some ghost curves. This follows since the maximum modulus principle implies that the projection to $[0, 1] \times \mathbb{R}$ of any holomorphic curve in $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$ must be constant. The asymptotics of \mathcal{U}_m must match those of u , and hence there must be exactly $n_p(\phi)$ once-covered cylinders which each project to a different point in $[0, 1] \times \mathbb{R}$. Any additional curves must have constant image in $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$. Write C for these cylinders.

There must be a component u_0 of \mathcal{U}_0 which has the same asymptotics (as a curve on $S^2 \setminus \{p_0\} \times [0, 1] \times \mathbb{R}$) at p_0 as the cylinders in \mathcal{U}_m (which we have already reasoned are the same as the asymptotics of u at p), i.e.,

$$\rho^p(u) = \rho^{p_0}(u_0).$$

Write ϕ'_0 for the homology class of u_0 . It is not hard to see that after trimming ghost curves, u_0 must satisfy (M1)–(M5) (embeddedness, (M6), is not yet clear). Write S_0 for the source curve of u_0 .

We now show that \mathcal{U}_0 contains no curves other than u_0 and possibly ghost curves, and that u_0 is embedded (we subsequently will rule out ghost curves).

By Proposition 3.4, for a generically chosen almost complex structure, near u_0 the set $\mathcal{M}(S_0, \phi'_0, X(\phi))$ is a manifold of dimension

$$(5-7) \quad \dim \mathcal{M}(S_0, \phi'_0, X(\phi)) = \mu(\phi'_0) - \text{codim}(X(\phi)) - 2 \text{Sing}(u_0).$$

Since $D(\phi'_0) \leq D(\phi_0)$, Lemma 5.4 implies that

$$(5-8) \quad \mu(\phi'_0) \leq \mu(\phi_0) \quad \text{and} \quad \mu(\phi_0) = 2n_p(\phi).$$

(The first inequality of (5-8) can also be proven using a transversality argument.)

Since $\text{codim}(X(\phi)) = 2n_p(\phi) - 1$, (5-8) implies

$$\mu(\phi'_0) \leq \text{codim}(X(\phi)) + 1.$$

Hence, (5-7) reduces to

$$\dim \mathcal{M}(S_0, \phi'_0, X(\phi)) \leq 1 - 2 \text{Sing}(u_0),$$

with equality if and only if $\mu(\phi'_0) = \mu(\phi_0)$. In particular, $\mathcal{M}(S_0, \phi'_0, X(\phi))$ is generically empty unless u_0 is embedded and $\phi'_0 = \phi_0$. It follows that \mathcal{U}_0 consists only of the unbroken curve u_0 , which satisfies (M1)–(M6), as well as possibly some ghost curves.

We now show that, generically, no ghost curves appear in $\mathcal{U}, \mathcal{U}_m$ or \mathcal{U}_0 . Our argument is essentially standard; see [Lipshitz et al. 2018, Lemma 5.57]. Let us write S_i for the source of u_i (a curve in our original sequence of $J^0(T_i)$ -holomorphic curves). As one stretches the neck, the sources S_i degenerate along a collection of boundary-to-boundary arcs and closed loops, as in the Deligne–Mumford compactification of stable curves. We can assume that all of the sources S_i are topologically identified with a fixed surface, which we denote by \hat{S} . The limiting curve is a nodal curve, with nodes on the boundary or interior, corresponding to where arcs and closed curves in \hat{S} collapse. Furthermore, we can assume that all of the curves which appear in the limit are *stable*, which implies that each ghost curve which is a disk or sphere has at least three nodes. Let G_1, \dots, G_n denote the ghost components.

Let us glue any nodes connecting two ghost components, and smooth any identified pair of nodes on a single ghost component. Abusing notation slightly, we write G_1, \dots, G_n for the resulting surfaces, which we call the *smoothed ghosts*. Note that gluing and smoothing ghost curves preserves stability.

We make the following claims about the smoothed ghosts:

- (g1) There are no ghosts which have image in $\{p\} \times [0, 1] \times \mathbb{R}$ or $\{p_0\} \times [0, 1] \times \mathbb{R}$.
- (g2) Each G_i has exactly one node to S, S_0 , or to a cylinder in C in the tube region.
- (g3) Each G_i has no boundary components with no nodes.
- (g4) Each G_i has genus at least 1.

Claim (g1) follows since the limiting collections \mathcal{U} and \mathcal{U}_0 are obtained by taking collections of curves in $\Sigma \setminus \{p\} \times [0, 1] \times \mathbb{R}$ and $S^2 \setminus \{p_0\} \times [0, 1] \times \mathbb{R}$, and completing over the points p and p_0 .

Claim (g2) is established as follows. That there is at least one node is easy to see: if there are no nodes on some G_i , then u_i must contain a closed, null-homologous component, which is prohibited by

axiom (M5). To see that there cannot be more than one node connecting a ghost G_i to S or S_i , we note that such a configuration would imply that u or u_0 had a double point, which are prohibited since u and u_0 are embedded.

We now prove claim (g3). If G_i had a boundary component with no nodes, then for large j the unbroken curve u_j would have a boundary component with no punctures, which must be mapped to $\Sigma \times \{0, 1\} \times \mathbb{R}$. However, such a configuration violates the maximum modulus principle applied to $\pi_{[0,1] \times \mathbb{R}} \circ u_j$.

Finally, claim (g4) is proven by noting that each G_i is stable, and has at most one node and boundary component by claims (g2) and (g3), so must have genus at least 1.

Having established claims (g1)–(g4), we now prove that, generically, ghosts do not appear. Lipshitz’s formula [2006, equation 6] for the Fredholm index implies

$$(5-9) \quad \text{ind}(S, \phi) = d - \chi(S) + 2e(D(\phi)) \quad \text{and} \quad \text{ind}(S_0, \phi_0) = 1 - \chi(S_0) + 2e(D(\phi_0)),$$

where $d = |\alpha| = |\beta|$ and $e(D(\phi))$ is the Euler measure of the domain of ϕ . We have already reasoned that the Fredholm indices in (5-9) are 1 and $2n_p(\phi)$, respectively, since ϕ and ϕ_0 have Maslov indices 1 and $2n_p(\phi)$, respectively, and u and u_0 are embeddings. Similarly,

$$(5-10) \quad \text{ind}(\hat{S}, \phi \# \phi_0) = (d + 1) - \chi(\hat{S}) + 2e(D(\phi \# \phi_0)),$$

which also must be 1, since the u_i are embedded and have Maslov index 1, by assumption. Noting that $e(D(\phi \# \phi_0)) = e(D(\phi)) + e(D(\phi_0)) - 2n_{p_0}(\phi_0)$, equations (5-9) and (5-10) combine to imply that

$$(5-11) \quad \chi(\hat{S}) = \chi(S) + \chi(S_0) - 2n_{p_0}(\phi_0).$$

Finally, equation (5-11) prohibits any ghost components in \mathcal{U} , \mathcal{U}_m or \mathcal{U}_0 : \hat{S} may be topologically constructed by gluing S and S_0 along the $n_{p_0}(\phi_0)$ punctures, as well as gluing any ghost components in. However, claims (g2), (g3) and (g4) imply that gluing in each G_i drops $\chi(\hat{S})$ by at least 2, contradicting (5-11). We conclude that, generically, no ghost curves appear in \mathcal{U} , \mathcal{U}_m or \mathcal{U}_0 .

Summarizing, any sequence u_i of $J^\partial(T_i)$ -holomorphic curves representing $\phi \# \phi_0$ has a subsequence which converges to a pair (u, u_0) which satisfies (M1)–(M6) and $\rho^p(u) = \rho^{p_0}(u_0)$.

We can view such a pair (u, u_0) as being a point in the compactification of the 1-dimensional moduli space

$$(5-12) \quad \bigcup_{T>0} \hat{\mathcal{M}}_{J^\partial(T)}(\phi \# \phi_0).$$

Standard gluing results for holomorphic curves (see [Lipshitz 2006, Proposition A.2; Bourgeois 2002, Section 5.3]) imply that for sufficiently large T , near (u, u_0) the compactification of the moduli space in (5-12) is modeled on a half-open interval times a fibered product,

$$(5-13) \quad [0, 1) \times ((\mathcal{M}_{J^\partial}(\phi) \times_\rho \mathcal{M}_{J^\partial}(\phi_0)) / \mathbb{R}).$$

Since $\mu(\phi) = 1$, the set $\widehat{\mathcal{M}}_{J^\circ}(\phi)$ is 0-dimensional. Hence, the characterization (5-13) implies that for sufficiently large T ,

$$(5-14) \quad \#\widehat{\mathcal{M}}_{J^\circ(T)}(\phi \# \phi_0) \equiv \sum_{u \in \widehat{\mathcal{M}}(\phi)} \#\widehat{\mathcal{M}}_{J_0^\circ}(\phi_0, \rho^P(u)) \pmod{2}.$$

According to [Ozsváth and Szabó 2008, Lemma 6.4], if $k > 0$ and $\mathbf{d} \in \text{Sym}^k([0, 1] \times \mathbb{R})$ is in the complement of the fat diagonal, then

$$(5-15) \quad \sum_{\substack{\phi_0 \in \pi_2(\theta, \theta) \\ n_{\rho_0}(\phi_0) = k \\ n_w(\phi_0) = 0}} \#\mathcal{M}_{J_0^\circ}(\phi_0, \mathbf{d}) \equiv 1 \pmod{2}.$$

Subclaim (d2) now follows by combining equations (5-14) and (5-15), together with our argument which showed that the only index 1 classes $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \theta, \mathbf{y} \times \theta)$ which have holomorphic representatives for large T have $\mu(\phi) = 1$ and $n_w(\phi_0) = 0$. □

Corollary 5.6 *Suppose (Y, \mathbf{w}) is a multipointed 3-manifold, $w \in Y \setminus \mathbf{w}$, and $\sigma: \mathbf{w} \rightarrow \mathbf{P}$ is a coloring which is extended by $\sigma': \mathbf{w} \cup \{w\} \rightarrow \mathbf{P}$. For sufficiently large T , the maps*

$$S_w^+: \text{CF}_{J^\circ}^-(\mathcal{H}, \sigma, \mathfrak{s}) \rightarrow \text{CF}_{J^\circ(T)}^-(\mathcal{H}^+, \sigma', \mathfrak{s}) \quad \text{and} \quad S_w^-: \text{CF}_{J^\circ(T)}^-(\mathcal{H}^+, \sigma', \mathfrak{s}) \rightarrow \text{CF}_{J^\circ}^-(\mathcal{H}, \sigma, \mathfrak{s})$$

are chain maps.

Proof The claim is a restatement of Proposition 5.5, using the definition of the free-stabilization maps in (5-1). □

5.3 Free stabilization and holomorphic triangles

In this section, we prove an analog of Proposition 5.5 for holomorphic triangles. If $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a multipointed Heegaard triple and $w \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w})$, consider the triple

$$\mathcal{T}^+ := (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w} \cup \{w\}),$$

where α_0, β_0 and γ_0 are new attaching curves contained in a small neighborhood of w on Σ . Furthermore, assume that the free-stabilization region has the configuration shown in Figure 6. In particular,

$$|\alpha_0 \cap \beta_0| = |\beta_0 \cap \gamma_0| = |\alpha_0 \cap \gamma_0| = 2.$$

If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple and $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w} \cup \{w\})$ is its free stabilization, then there is a canonical diffeomorphism

$$(5-16) \quad X_{\boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}} \cong X_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}},$$

where $X_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}}$ is the 4-manifold defined in (3-23). Equation (5-16) follows from the fact that neither the Heegaard surface Σ nor the handlebodies U_α, U_β or U_γ are changed by a free stabilization.




Figure 6: A free-stabilized Heegaard triple. If we collapse the dashed circle to the point p_0 , we get the doubly based Heegaard triple $(S^2, \alpha_0, \beta_0, \gamma_0, p_0, w)$.

It is straightforward to verify that the Spin^c structure map \mathfrak{s}_w from (3-24) is also unchanged, in the sense that if $\psi \# \psi_0 \in \pi_2(\mathbf{x} \times \mathbf{x}, \mathbf{y} \times \mathbf{y}, \mathbf{z} \times \mathbf{z})$ is a homology class of triangles on the free-stabilized diagram, then

$$\mathfrak{s}_w(\psi \# \psi_0) = \mathfrak{s}_w(\psi),$$

with respect to the diffeomorphism from (5-16).

Theorem 5.7 Suppose \mathcal{T} is a Heegaard triple, \mathcal{T}^+ is its free stabilization at the point w , and \mathfrak{d} is a gluing datum for free stabilization. Then, for sufficiently large T , we have

$$\begin{aligned} F_{\mathcal{T}^+, J^{\mathfrak{d}}(T), \mathfrak{s}}(S_w^+(-), S_w^+(-)) &= S_w^+ F_{\mathcal{T}, J^{\mathfrak{d}}, \mathfrak{s}}(-, -), \\ S_w^- F_{\mathcal{T}^+, J^{\mathfrak{d}}(T), \mathfrak{s}}(S_w^+(-), -) &= F_{\mathcal{T}, J^{\mathfrak{d}}, \mathfrak{s}}(-, S_w^-(-)), \\ S_w^- F_{\mathcal{T}^+, J^{\mathfrak{d}}(T), \mathfrak{s}}(-, S_w^+(-)) &= F_{\mathcal{T}, J^{\mathfrak{d}}, \mathfrak{s}}(S_w^-(-), -). \end{aligned}$$

Before we prove Theorem 5.7, we prove a useful Maslov index formula:

Lemma 5.8 Suppose that $\psi_0 \in \pi_2(x, y, z)$ is a homology class of triangles on the Heegaard triple $(S^2, \alpha_0, \beta_0, \gamma_0, w, p_0)$ in Figure 6. Then

$$\mu(\psi_0) = -\text{gr}(x^+, x) - \text{gr}(y^+, y) + \text{gr}(z^+, z) + 2n_w(\psi_0) + 2n_{p_0}(\psi_0).$$

If $m_1(\psi_0), \dots, m_8(\psi_0)$ denote the multiplicities of ψ_0 in each of the eight regions of $S^2 \setminus (\alpha_0 \cup \beta_0 \cup \gamma_0)$, then

$$\mu(\psi_0) = \frac{1}{2}(m_1(\psi_0) + \dots + m_8(\psi_0) - 1).$$

Proof Both formulas can be verified using the same strategy. First, both are easily checked for the triangle class in $\pi_2(x^+, y^+, z^+)$ whose domain consists of a single component of $S^2 \setminus (\alpha_0 \cup \beta_0 \cup \gamma_0)$ with multiplicity 1. Both formulas respect splicing in any of the 12 bigons on the diagram. Since any two triangle classes can be related by splicing in such bigons, both formulas follow in general. \square

Proof of Theorem 5.7 The proof is similar to the proof of Proposition 5.5.

Rephrasing the statement in terms of intersection points, it suffices to show that if T is sufficiently large, then

$$\begin{aligned}
 (5-17) \quad & F_{\mathcal{T}^+, J^\partial(T), s}(\mathbf{x} \times x^+, \mathbf{y} \times y^+) = F_{\mathcal{T}, J^\partial, s}(\mathbf{x}, \mathbf{y}) \otimes z^+, \\
 & F_{\mathcal{T}^+, J^\partial(T), s}(\mathbf{x} \times x^+, \mathbf{y} \times y^-) = F_{\mathcal{T}, J^\partial, s}(\mathbf{x}, \mathbf{y}) \otimes z^- + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} C_{\mathbf{x}, \mathbf{y}, z}^1 \cdot z \times z^+, \\
 & F_{\mathcal{T}^+, J^\partial(T), s}(\mathbf{x} \times x^-, \mathbf{y} \times y^+) = F_{\mathcal{T}, J^\partial, s}(\mathbf{x}, \mathbf{y}) \otimes z^- + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} C_{\mathbf{x}, \mathbf{y}, z}^2 \cdot z \times z^+
 \end{aligned}$$

for $C_{\mathbf{x}, \mathbf{y}, z}^1, C_{\mathbf{x}, \mathbf{y}, z}^2 \in \mathbb{F}_2[U_w, U_w]$ (which depend on T and ∂).

Suppose that $\psi \# \psi_0 \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y, z \times z)$ is a homology class of triangles with Maslov index 0. Define the quantity

$$\delta(x, y, z) := -\text{gr}(x^+, x) - \text{gr}(y^+, y) + \text{gr}(z^+, z).$$

It is easy check that (5-17) is implied by the following two subclaims:

- (f1) If $\delta(x, y, z) = 1$, then any Maslov index 0 class in $\pi_2(\mathbf{x} \times x, \mathbf{y} \times y, z \times z)$ has no $J^\partial(T)$ -holomorphic representatives if T is sufficiently large.
- (f2) If $\delta(x, y, z) = 0$, then the $z \times z$ coefficient of $F_{\mathcal{T}, J^\partial(T), s}(\mathbf{x} \times x, \mathbf{y} \times y)$ is equal to the z coefficient of $F_{\mathcal{T}, J^\partial, s}(\mathbf{x}, \mathbf{y})$.

The classes with $\delta(x, y, z) \in \{-2, -1\}$ do not contribute to the present theorem statement.

In our proof, we will write $p \in \Sigma$ for the point w , viewed as a point on the unstabilized Heegaard triple, and write w for the new basepoint on the free-stabilized Heegaard triple.

We compute

$$(5-18) \quad \mu(\psi \# \psi_0) = \mu(\psi) + \mu(\psi_0) - 2n_{p_0}(\psi_0) = \mu(\psi) + \delta(x, y, z) + 2n_w(\psi_0).$$

The first equality of (5-18) follows from Sarkar’s formula [2011] for the Maslov index, as well as the fact that a disk has Euler measure 1. The second equality follows from Lemma 5.8.

Suppose that $\psi \# \psi_0$ admits holomorphic representatives for arbitrarily large neck length parameter T . As in the proof of Proposition 5.5, we may find a subsequence which converges to broken collections of curves $\mathcal{U}, \mathcal{U}_m$ and \mathcal{U}_0 , on $\Sigma \times \Delta, S^1 \times \mathbb{R} \times \Delta$ and $S^2 \times \Delta$, and their cylindrical ends, $\Sigma \times [0, 1] \times \mathbb{R}, S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$ and $S^2 \times [0, 1] \times \mathbb{R}$.

Consider claim (f1). In particular, ψ and ψ_0 both admit broken representatives. Using (5-18) and the equalities $\mu(\psi \# \psi_0) = 0$ and $\delta(x, y, z) = 1$, we conclude that $\mu(\psi) = -1$. However, the triangular analog of Proposition 3.4 implies that, generically, $\mathcal{M}(\psi)$ is empty, contradicting the existence of the broken triangle \mathcal{U} . Claim (f1) follows.

We now consider claim (f2). Let $\psi \# \psi_0$ denote a class with $\delta(x, y, z) = 0$, and let $\mathcal{U}, \mathcal{U}_m$ and \mathcal{U}_0 denote broken holomorphic triangles arising in the limit. We conclude from (5-18) and transversality that $\mu(\psi) = 0$.

As a first step, trim off any ghost components of \mathcal{U} , \mathcal{U}_m and \mathcal{U}_0 (i.e., curves which have constant image in $\Sigma \times \Delta$ or one of its cylindrical ends). We will later see that ghost curves do not appear generically.

Next, we claim that \mathcal{U} cannot contain any curves v mapping into $\Sigma \times \Delta$ or one of its cylindrical ends, such that $\pi_\Delta \circ v$ (or $\pi_{[0,1] \times \mathbb{R}} \circ v$) is constant, but $\pi_\Sigma \circ v$ is nonconstant. Such a curve v must have Maslov index at least 2, since it has domain equal to a nonnegative integral combination of components of $\Sigma \setminus \alpha$, $\Sigma \setminus \beta$ and $\Sigma \setminus \gamma$. The presence of such a curve would force the remaining curves, on which transversality is obtained, to have Maslov index at most -2 , which is impossible, generically.

Consequently, all curves in \mathcal{U} (except the ghost components) satisfy the analogs of (M1)–(M5), so, by the analog of Proposition 3.4 for triangles, transversality is obtained at u . Since $\mu(\psi) = 0$, it follows from transversality that \mathcal{U} consists only of a single holomorphic triangle u satisfying the analogs of (M1)–(M6) for holomorphic triangles (as well as possibly ghost curves). In particular, u is embedded.

Arguing identically to the proof of Proposition 5.5, the curves of \mathcal{U}_m , generically, consist of $n^p(\phi)$ once-covered cylinders, which each map to a distinct point in Δ . (There may also be ghost components in \mathcal{U}_m , though we will shortly rule those out.)

There must be a curve u_0 in \mathcal{U}_0 whose asymptotics (before completing over $\{p\} \times \Delta$) match those of u , i.e.,

$$(5-19) \quad \rho^p(u) = \rho^{p_0}(u_0),$$

where $\rho^p: \mathcal{M}(\psi) \rightarrow \text{Sym}^{n_p(\psi)}(\Delta)$ is the map defined analogously to (3-5).

Let ψ'_0 denote the homology class of u_0 . Since $D(\psi'_0) \leq D(\psi_0)$, we conclude from Lemma 5.8 that

$$(5-20) \quad \mu(\psi'_0) \leq \mu(\psi_0) = 2n_{p_0}(\psi_0).$$

Define

$$X(\psi) := \{\rho^p(u) : u \in \mathcal{M}(\psi)\},$$

a 0-dimensional subset of $\text{Sym}^{n_p(\psi)}(\Delta)$.

Write S_0 for the topological source of u_0 . Using Proposition 3.4 and (5-20), for a generically chosen almost complex structure,

$$(5-21) \quad \begin{aligned} \dim \mathcal{M}(S_0, \psi'_0, X(\psi_0)) &= \mu(\psi'_0) - \text{Sing}(u_0) - \text{codim } X(\psi_0) \\ &= \mu(\psi'_0) - \text{Sing}(u_0) - 2n_{p_0}(u_0) \\ &\leq -\text{Sing}(u_0). \end{aligned}$$

Furthermore, equality holds if and only if $\mu(\psi'_0) = \mu(\psi_0) = 2n_{p_0}(u_0)$ and $\text{Sing}(u_0) = 0$. Consequently, if $\mathcal{M}(S_0, \psi'_0, X(\psi_0))$ is nonempty, we must generically have $\mu(\psi'_0) = \mu(\psi_0)$ and $\text{Sing}(u_0) = 0$. In particular, \mathcal{U}_0 cannot have any other components (other than ghost components) and u_0 must be embedded.

The appearance of ghost components in \mathcal{U} and \mathcal{U}_0 , which we trimmed off earlier, can be prohibited from appearing generically by using the same index argument as in the proof of Proposition 5.5.

Summarizing, if T_i is a sequence of neck lengths approaching $+\infty$, then any sequence of $J^\partial(T_i)$ -holomorphic triangles u_i representing $\psi \# \psi_0$ has limit equal to a pair (u, u_0) of holomorphic triangles satisfying (M1)–(M6) and $\rho^P(u) = \rho^{P_0}(u_0)$.

As in the proof of Proposition 5.5, for large T , gluing gives a bijection

$$(5-22) \quad \mathcal{M}_{J^\partial(T)}(\psi \# \psi_0) \cong \mathcal{M}_{J^\partial}(\psi) \times_\rho \mathcal{M}_{J_0^\partial}(\psi_0).$$

Since $\mathcal{M}_{J^\partial}(\psi)$ is 0-dimensional, counting $\mathcal{M}_{J^\partial(T)}(\psi \# \psi_0)$ is reduced to the problem of counting the elements in the matched moduli space $\mathcal{M}_{J_0^\partial}(\psi_0, \mathbf{d})$ for generic $\mathbf{d} \in \text{Sym}^n(\Delta)$. We will show that, for generic J_0^∂ ,

$$(5-23) \quad \sum_{\substack{\psi_0 \in \pi_2(x, y, z) \\ n_{w_0}(\psi_0) = 0 \\ n_{p_0}(\psi_0) = n}} \# \mathcal{M}_{J_0^\partial}(\psi_0, \mathbf{d}) \equiv 1 \pmod{2}$$

whenever $\delta(x, y, z) = 0$ and $\mathbf{d} \in \text{Sym}^n(\Delta)$ is not in the fat diagonal.

To establish (5-23), we use the following cobordism argument. Consider a path $(\mathbf{d}_t)_{t \in [1, \infty)}$ such that:

- (1) $\mathbf{d}_1 = \mathbf{d}$.
- (2) Each \mathbf{d}_t is not contained in the fat diagonal.
- (3) As $t \rightarrow \infty$, the points in \mathbf{d}_t all travel to ∞ in the α - β cylindrical end of Δ .
- (4) If we identify the α - β cylindrical end of Δ with $[0, 1] \times (-\infty, 0]$, then the points of \mathbf{d}_t are spaced at least t distance apart for large t .
- (5) The $[0, 1]$ -component of each point of \mathbf{d}_t approaches some fixed $s_0 \in (0, 1)$.

Write

$$\mathcal{D} := \{\mathbf{d}_t : t \in [1, \infty)\}.$$

Consider the ends of the 1-dimensional moduli space

$$\mathcal{M}_{(x, y, z)}(\mathcal{D}) := \coprod_{\substack{\psi_0 \in \pi_2(x, y, z) \\ n_{w_0}(\psi_0) = 0 \\ n_{p_0}(\psi_0) = n}} \mathcal{M}_{J_0^\partial}(\psi_0, \mathcal{D}).$$

There are three types of ends to consider:

- (e1) curves at $t = 1$;
- (e2) degenerations occurring at finite $t \in (1, \infty)$;
- (e3) limiting curves as $t \rightarrow \infty$.

Ends of $\mathcal{M}_{(x, y, z)}(\mathcal{D})$ of type (e1) correspond to the disjoint union of $\mathcal{M}_{J_0^\partial}(\psi_0, \mathbf{d})$ over all $\psi_0 \in \pi_2(x, y, z)$ with $n_w(\psi_0) = 0$.

Ends of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ of type (e2) correspond to sequences of holomorphic triangles in $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ degenerating into a broken holomorphic triangle at some finite $t \in (0, \infty)$. We wish to show that such a degeneration can consist only of a holomorphic strip and triangle, both satisfying (M1)–(M6) or their triangular analogs. Furthermore, the holomorphic strip must have Maslov index 1, and zero multiplicity on p_0 , and the triangle must have Maslov index $2n_{p_0}(\psi_0) - 1$, and must match some \mathbf{d}_t .

Let \mathcal{V} denote a broken holomorphic triangle appearing as the limit of a sequence of curves in $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ at some $t \in (1, \infty)$. There must be a holomorphic triangle v_0 of \mathcal{V} which matches \mathbf{d}_t .

Delete curves v in \mathcal{V} which do not map into $\Sigma \times \Delta$, or such that $\pi_\Delta \circ v$ is constant. Since $n_w(\psi_0) = 0$, there are no nonghost components v of \mathcal{V} such that $\pi_\Delta \circ v$ is constant and equals a point of \mathbf{d}_t . Consequently, after removing any components whose projection to Δ is constant, the remaining curve v_0 still matches \mathbf{d}_t and satisfies the triangular analogs of (M1)–(M5) (embeddedness, (M6), is not clear yet). Write ψ'_0 for the homology class of v_0 and S_0 for the source curve of v_0 . By the triangular analog of Proposition 3.4, the space $\mathcal{M}(S_0, \psi'_0, \mathcal{D})$ is transversely cut out and of dimension

$$(5-24) \quad \mu(\psi'_0) - 2 \text{Sing}(u_0) - \text{codim}(\mathcal{D}) = \mu(\psi'_0) - 2 \text{Sing}(v_0) - (2|\mathbf{d}| - 1).$$

The quantity $\text{Sing}(v_0)$ takes values in $\frac{1}{2} \cdot \mathbb{Z}^{\geq 0}$; a singularity along the interior contributes 1, while a singularity along the boundary contributes $\frac{1}{2}$. Note that $\mu(\psi_0) = 2|\mathbf{d}|$, and $\mu(\psi'_0) \leq \mu(\psi_0)$ by Lemma 5.8. Hence, equation (5-24) implies that there are two possibilities which are not prohibited by dimension counts: $\mu(\psi'_0) = 2|\mathbf{d}| - 1$ and $\text{Sing}(v_0) = 0$ (v_0 is embedded); or $\mu(\psi'_0) = 2|\mathbf{d}|$ and $\text{Sing}(v_0) = \frac{1}{2}$ (v_0 has a singular point along its boundary).

We now rule out a boundary double point. A boundary double point arises when a boundary-to-boundary arc c in the source curve collapses. Suppose $v_i: S \rightarrow S^2 \times \Delta$ is a sequence of curves in $\mathcal{M}_{(x,y,z)}(\mathcal{D})$, all sharing the same topological source, which realizes such a degeneration. Write $\partial_\alpha S$, $\partial_\beta S$ and $\partial_\gamma S$ for the subsets of ∂S which map to α_0 , β_0 and γ_0 , respectively. If q_1 and q_2 are the two ends of c , then both q_1 and q_2 must be contained in a single one of $\partial_\alpha S$, $\partial_\beta S$ or $\partial_\gamma S$. Condition (J'4) implies that $v_i|_{\partial_\tau S}$ is monotonic for $\tau \in \{\alpha, \beta, \gamma\}$. Consequently, if a boundary double point forms, the subarc region between q_1 and q_2 along the boundary of S must map constantly to Δ , in the limit. Consequently, the limiting curve must have a component which is a boundary degeneration. Since \mathcal{D} is bounded away from $\partial\Delta$, and all curves in $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ have zero multiplicity on w , boundary degenerations cannot form in the ends of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$, since any boundary degeneration on $(S^2, \alpha_0, \beta_0, \gamma_0)$, covers one of p_0 or w . It follows that $\text{Sing}(v_0) = 0$ and

$$\mu(\psi'_0) \in \{2|\mathbf{d}|, 2|\mathbf{d}| - 1\}.$$

If $\mu(\psi'_0) = 2|\mathbf{d}| = \mu(\psi_0)$, then the remaining curves in the broken triangle have Maslov index 0. Using transversality, it is not hard to see that any such curves must be constant holomorphic strips in the cylindrical ends. By definition, constant strips are not stable, so cannot appear in a broken triangle in the boundary of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$. Hence, no curves with $\mu(\psi'_0) = 2|\mathbf{d}|$ appear in the boundary of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ for $t \in (1, \infty)$.

Next, we consider the case that $\mu(\psi'_0) = 2|\mathbf{d}| - 1$. In this case, the remaining curves (call them \mathcal{V}_1) have total Maslov index 1. It is straightforward to see that this implies that \mathcal{V}_1 consists of a single holomorphic flowline v_1 on $S^2 \times [0, 1] \times \mathbb{R}$ on one of the three subdiagrams (S^2, α_0, β_0) , (S^2, β_0, γ_0) or $(S^2, \alpha_0, \gamma_0)$. The curve v_1 must have multiplicity zero on p_0 and w . Since the differential vanishes on all three complexes $\widehat{\text{CF}}(S^2, \alpha_0, \beta_0, p_0, w)$, $\widehat{\text{CF}}(S^2, \alpha_0, \gamma_0, p_0, w)$ and $\widehat{\text{CF}}(S^2, \beta_0, \gamma_0, p_0, w)$, ends of type (e2) come in canceling pairs.

Finally we consider the ends of type (e3). Using the Maslov index formulas from Lemmas 5.4 and 5.8, together with the transversality result from Proposition 3.4 (as well as its triangular analog for curves in $S^2 \times \Delta$), we can show that the limit of a sequence of curves with $t \rightarrow \infty$ must consist of $|\mathbf{d}|$ flowlines $v_1, \dots, v_{|\mathbf{d}|}$ on (S^2, α_0, β_0) , representing Maslov index 2 classes in $\pi_2(x, x)$ with $n_{p_0} = 0$, each of which matches some fixed point $d \in [0, 1] \times \mathbb{R}$, as well as a single index 0 holomorphic triangle v with Maslov index 0 and $n_w(v) = n_{p_0}(v) = 0$. Using [Lipshitz 2006, Proposition A.1] to glue, such ends of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ at $t = \infty$ correspond to the 0-dimensional space

$$(5-25) \quad \left(\coprod_{\substack{\psi_0^0 \in \pi_2(x,y,z) \\ n_w(\psi_0^0) = n_{p_0}(\psi_0^0) = 0}} \mathcal{M}(\psi_0^0) \right) \times \left(\coprod_{\substack{\phi \in \pi_2(x,x) \\ n_{p_0}(\phi) = 1 \\ n_w(\phi) = 0}} \mathcal{M}(\phi, d) \right)^{|\mathbf{d}|}.$$

It is easy to verify that there is only one nonnegative class $\psi_0^0 \in \pi_2(x, y, z)$ with Maslov index 0 and $n_w(\psi_0^0) = n_{p_0}(\psi_0^0) = 0$. Furthermore, ψ_0^0 has a unique representative by the Riemann mapping theorem.

Using Ozsváth and Szabó’s count from (5-23), the number of elements in the set in (5-25) is 1, modulo 2. In particular, counting up all ends of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$, equation (5-23) follows.

Combining (5-22) and (5-23) yields claim (f2), completing the proof. □

5.4 Independence from the gluing datum

In this section, we prove that the free-stabilization maps are independent of the choice of gluing datum:

Proposition 5.9 *Suppose \mathcal{H} is diagram for (Y, \mathbf{w}) , \mathfrak{d}_1 and \mathfrak{d}_2 are two gluing data for free-stabilizing at w , and T_1 and T_2 satisfy the stabilizing condition (SC2) for \mathfrak{d}_1 and \mathfrak{d}_2 , respectively. Write \mathcal{H}_1^+ and \mathcal{H}_2^+ for the two free-stabilized diagrams constructed from \mathcal{H} , \mathfrak{d}_1 and \mathfrak{d}_2 (note that \mathcal{H}_1^+ and \mathcal{H}_2^+ differ only in the embedding of α_0 and β_0 in the free-stabilization region). The following diagram commutes up to chain homotopy:*

$$(5-26) \quad \begin{array}{ccc} \text{CF}_{J^{\mathfrak{d}_1}}^-(\mathcal{H}, \mathfrak{s}) & \xrightarrow{\Psi_{J^{\mathfrak{d}_1} \rightarrow J^{\mathfrak{d}_2}}} & \text{CF}_{J^{\mathfrak{d}_2}}^-(\mathcal{H}, \mathfrak{s}) \\ \downarrow S_w^+ & & \downarrow S_w^+ \\ \text{CF}_{J^{\mathfrak{d}_1}(T_1)}^-(\mathcal{H}_1^+, \mathfrak{s}) & \xrightarrow{\Psi_{(\mathcal{H}_1^+, J^{\mathfrak{d}_1}(T_1)) \rightarrow (\mathcal{H}_2^+, J^{\mathfrak{d}_2}(T_2))}} & \text{CF}_{J^{\mathfrak{d}_2}(T_2)}^-(\mathcal{H}_2^+, \mathfrak{s}) \end{array}$$

An analogous relation holds for the maps S_w^- .

Proof Note that if the claim holds for any particular T_1 and T_2 which satisfy (SC2), then it holds for all T_1 and T_2 which satisfy (SC2).

Let us call two gluing data \mathfrak{d}_1 and \mathfrak{d}_2 *equivalent*, and write $\mathfrak{d}_1 \sim \mathfrak{d}_2$, if, for all sufficiently large T_1 and T_2 , the diagram in (5-26), as well as its analog for S_w^- , commutes up to chain homotopy. The relation of two gluing data being equivalent is transitive and reflexive. It is sufficient to show that all generic gluing data are equivalent.

The proof is an elaboration of the proofs of Propositions 5.3 and 5.5. We prove the main statement in three steps. Write

$$\mathfrak{d}_1 = (J_1, J_{0,1}, D_1, \iota_1) \quad \text{and} \quad \mathfrak{d}_2 = (J_2, J_{0,2}, D_2, \iota_2).$$

Step 1 *If $D_1 = D_2$, $J_1|_{D_i \times [0,1] \times \mathbb{R}} = J_2|_{D_i \times [0,1] \times \mathbb{R}}$, $J_{0,1} = J_{0,2}$ and $\iota_1 = \iota_2$, then $\mathfrak{d}_1 \sim \mathfrak{d}_2$.*

We can pick an almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$ which interpolates J_1 and J_2 and is cylindrical on $D_i \times [0, 1] \times \mathbb{R}$. Since $J_{0,1} = J_{0,2}$ and $\iota_1 = \iota_2$, using \tilde{J} we can construct almost complex structures $\tilde{J}(T)$ which interpolate $J^{\mathfrak{d}_1}(T)$ and $J^{\mathfrak{d}_2}(T)$ and agree with \tilde{J} on $(\Sigma \setminus \frac{1}{T} \cdot D) \times [0, 1] \times \mathbb{R}$.

We claim that if T is sufficiently large, then

$$(5-27) \quad \Psi_{\tilde{J}(T)}(\mathbf{x} \times \theta^+) = \Psi_{\tilde{J}}(\mathbf{x}) \times \theta^+ \quad \text{and} \quad \Psi_{\tilde{J}(T)}(\mathbf{x} \times \theta^-) = \Psi_{\tilde{J}}(\mathbf{x}) \times \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x}, \mathbf{y}} \cdot \mathbf{y} \times \theta^+$$

for $C_{\mathbf{x}, \mathbf{y}} \in \mathbb{F}_2[U_{\mathbf{w}}]$ (depending on T , \mathfrak{d}_1 and \mathfrak{d}_2). Note that (5-27) is equivalent to the statement that $\mathfrak{d}_1 \sim \mathfrak{d}_2$.

Equation (5-27) follows from a modification of the proof of Proposition 5.5. We briefly sketch the argument.

Suppose that $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \mathbf{x}, \mathbf{y} \times \mathbf{y})$ is a homology class of disks with Maslov index 0. If T_i is a sequence of neck lengths approaching $+\infty$, then, from a sequence u_i of $\tilde{J}(T_i)$ holomorphic curves representing $\phi \# \phi_0$, we may extract two broken limiting curves, \mathcal{U} and \mathcal{U}_0 , representing ϕ and ϕ_0 . From Lemma 5.4, we know

$$0 = \mu(\phi \# \phi_0) = \mu(\phi) + \text{gr}(x, y) + n_w(\phi_0).$$

Equation (5-27) involves only classes with $\text{gr}(x, y) \geq 0$. Since ϕ and ϕ_0 both have broken representatives, we conclude that $\mu(\phi) \geq 0$. Furthermore, if $\text{gr}(x, y) = 1$, then there are no representatives of $\phi \# \phi_0$. Hence, if $\phi \# \phi_0$ has a representative for large T and $\text{gr}(x, y) \geq 0$, then

$$0 = \mu(\phi) = \text{gr}(x, y) = n_w(\phi_0).$$

Exactly one story of \mathcal{U} will be \tilde{J} -holomorphic, while the others will be $J^{\mathfrak{d}_1}$ - or $J^{\mathfrak{d}_2}$ -holomorphic. Since $\mu(\phi) = 0$, \mathcal{U} can consist only of a single, embedded curve u , which is \tilde{J} -holomorphic. The collection \mathcal{U}_0 will have a matching component u_0 which satisfies $\rho^w(u) = \rho^p(u_0)$. Using a transversality argument, analogous to the one described in Proposition 5.5, for generically chosen J_0 , u_0 will be the only curve in \mathcal{U}_0 , and u_0 will be embedded. From here, the argument proceeds via a gluing argument, analogous to the proof of Proposition 5.5.

Step 2 *If $J_1 = J_2$ and $D_1 = D_2$, then $\mathfrak{d}_1 \sim \mathfrak{d}_2$.*

The transition map $\Psi_{(\mathcal{H}_1^+, J^{\mathfrak{d}_1}(T_1)) \rightarrow (\mathcal{H}_2^+, J^{\mathfrak{d}_2}(T_2))}$ can be written as a composition of the map induced by a diffeomorphism $\phi: \Sigma \rightarrow \Sigma$ which is supported in a small neighborhood of $\frac{1}{2} \cdot D$, which moves the image

of the α_0 and β_0 curves under ι_1 to their images under ι_2 , followed by a map $\Psi_{\tilde{J}}$ which counts index 0 \tilde{J} -holomorphic disks for an almost complex structure \tilde{J} which interpolates $\phi_* J^{\partial_1}(T_1)$ and $J^{\partial_2}(T_2)$.

The diagram

$$\begin{CD} \text{CF}_{\tilde{J}}^-(\mathcal{H}, \mathfrak{s}) @>\text{id}>> \text{CF}_{\tilde{J}}^-(\mathcal{H}, \mathfrak{s}) \\ @VV S_w^+ V @VV S_w^+ V \\ \text{CF}_{J^{\partial_1}(T_1)}^-(\mathcal{H}_1^+, \mathfrak{s}) @>\phi_*>> \text{CF}_{\phi_* J^{\partial_1}(T_1)}^-(\mathcal{H}_2^+, \mathfrak{s}) \end{CD}$$

clearly commutes, so it is sufficient to show that the analogous diagram involving $\Psi_{\phi_* J^{\partial_1}(T_1) \rightarrow J^{\partial_2}(T_2)}$ also commutes. It is sufficient to show that the diagram commutes whenever T_1 and T_2 are sufficiently large. Furthermore, using condition (SC2), we can assume $T_1 = T_2$.

Since ϕ_* is the identity outside of a neighborhood of $\frac{1}{2} \cdot D$, we can choose \tilde{J} so that it agrees with $J^{\partial_i}(T)$ (for both $i = 1, 2$) on $\Sigma \setminus \frac{1}{2} \cdot D$. However, equipped with the almost complex structure $J^{\partial_i}(T)$, $(\Sigma \setminus \frac{1}{2}) \times [0, 1] \times \mathbb{R}$ is holomorphically equivalent to $((\Sigma \setminus \frac{1}{T}) \cdot D) \times [0, 1] \times \mathbb{R}, J$.

For a sequence T_i approaching $+\infty$, we pick a sequence \tilde{J}_i of interpolating almost complex structures between $\phi_* J^{\partial_1}(T_i)$ and $J^{\partial_2}(T_i)$. Given \tilde{J}_i -holomorphic representatives u_i of an index 0 class $\phi \# \phi_0$, we extract a limit to broken curves \mathcal{U} and \mathcal{U}_0 . As in Step 1, this implies that there are no representatives of classes with $\text{gr}(x, y) = 1$, and the only remaining classes of interest have

$$\mu(\phi) = \text{gr}(x, y) = n_w(\phi_0) = 0.$$

However, \mathcal{U} is a broken J -holomorphic curve, which is a cylindrical almost complex structure. Since $\mu(\phi) = 0$, it follows from transversality that \mathcal{U} can consist only of a single story, consisting of a representative of the constant class e_x . Since $\phi \# \phi_0$ has Maslov index 0 and $\phi = e_x$, we clearly must have $\phi_0 = \{e_{\theta+}, e_{\theta-}\}$. On the other hand, it is straightforward to see that the classes $e_x \times e_x$ will always have \tilde{J}_i -holomorphic representatives. The stated formula follows.

Step 3 *If ∂_1 and ∂_2 are generic, then $\partial_1 \sim \partial_2$.*

Write J_i as $j_i \times j_{[0,1] \times \mathbb{R}}$ on D_i . By the uniformization theorem, there is a conformal diffeomorphism between (D_1, j_1) and (D_2, j_2) which fixes w , which we can assume extends to a diffeomorphism from Σ to itself and which fixes $\alpha \cup \beta \cup w \cup \{w\}$.

Writing ϕ also for the induced diffeomorphism of $\Sigma \times [0, 1] \times \mathbb{R}$, we tautologically have a commutative diagram

$$(5-28) \quad \begin{CD} \text{CF}_{J_1}^-(\mathcal{H}, \mathfrak{s}) @>\phi_*>> \text{CF}_{\phi_* J_1}^-(\mathcal{H}, \mathfrak{s}) \\ @VV S_w^+ V @VV S_w^+ V \\ \text{CF}_{J^{\partial_1}(T)}^-(\mathcal{H}_1^+, \mathfrak{s}) @>\phi_*>> \text{CF}_{J^{\phi_* \partial_1}(T)}^-(\mathcal{H}_2^+, \mathfrak{s}) \end{CD}$$

Since both horizontal arrows in (5-28) agree with the maps from naturality,

$$\partial_1 \sim \phi_*(\partial_1).$$

Define the gluing datum $\mathfrak{d}' := (\phi_*(J_1), J_{0,2}, D_2, \iota_2)$. Step 2 shows that

$$\phi_*(\mathfrak{d}_1) \sim \mathfrak{d}'.$$

Step 1 implies that

$$\mathfrak{d}' \sim \mathfrak{d}_2.$$

By transitivity of \sim , the proof is complete. □

5.5 Commuting simple stabilizations and free stabilizations

Lemma 5.10 *The transition maps for a simple stabilization commute with the free-stabilization maps.*

Proof We will focus on the positive free-stabilization map S_w^+ since S_w^- is the dual map. Write σ for the simple stabilization map. Suppose that the following hold:

- (1) $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) .
- (2) $\mathcal{H}^+ = (\Sigma, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \mathbf{w} \cup \{w\})$ is a free stabilization of \mathcal{H} .
- (3) $\mathcal{H}_\sigma = (\Sigma', \alpha \cup \{\alpha'\}, \beta \cup \{\beta'\}, \mathbf{w})$ is a simple stabilization of \mathcal{H} .
- (4) $\mathcal{H}_\sigma^+ = (\Sigma', \alpha \cup \{\alpha_0, \alpha'\}, \beta \cup \{\beta_0, \beta'\}, \mathbf{w} \cup \{w\})$ is obtained from \mathcal{H} by performing both stabilizations.

The formulas for σ and S_w^+ appear to commute; however, it is not clear that a single almost complex structure can be chosen to compute both maps, since σ requires stabilizing condition (SC1) to be satisfied, while S_w^+ requires condition (SC2) to be satisfied.

The main claim amounts to showing the following subclaim:

Subclaim 5.10.1 *If $T = (T, T')$ is a pair of neck lengths, write $J(T)$ for the almost complex structure on $\Sigma' \times [0, 1] \times \mathbb{R}$ which has had a neck of length T added in the free-stabilization region, and a neck of length T' inserted in the simple-stabilization region. If T_1 and T_2 are two pairs all of whose components are sufficiently large, then there is a noncylindrical almost complex structure \tilde{J} on $\Sigma' \times [0, 1] \times \mathbb{R}$, interpolating $J(T_1)$ and $J(T_2)$, such that*

$$\Psi_{\tilde{J}}(\mathbf{x} \times c \times \theta^+) = \mathbf{x} \times c \times \theta^+, \quad \Psi_{\tilde{J}}(\mathbf{x} \times c \times \theta^-) = \mathbf{x} \times c \times \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x}, \mathbf{y}} \cdot \mathbf{y} \times c \times \theta^+$$

for $C_{\mathbf{x}, \mathbf{y}} \in \mathbb{F}_2[U_w]$.

The proof of Subclaim 5.10.1 is a double neck-stretching argument similar to the proof of Proposition 5.3.

We can write any homology class on \mathcal{H}_σ^+ as $\phi \# \phi_0 \# \phi'$, where $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is a class on \mathcal{H} , $\phi_0 \in \pi_2(x, y)$ is a class on (S^2, α_0, β_0) , and $\phi' \in \pi_2(c, c)$ is a homology class on $(\mathbb{T}^2, \alpha', \beta')$. Adapting the proof of Lemma 5.4, we have

$$(5-29) \quad \mu(\phi \# \phi_0 \# \phi') = \mu(\phi) + \text{gr}(x, y) + 2n_w(\phi_0).$$

Suppose $T_{1,i}$ and $T_{2,i}$ are sequences of pairs of neck lengths all of whose components approach $+\infty$. We can pick a sequence of noncylindrical almost complex structures \tilde{J}_i interpolating $J(T_{1,i})$ and $J(T_{2,i})$ such that $(\Sigma' \times [0, 1] \times \mathbb{R}, \tilde{J}_i)$ contains an almost complex submanifold equivalent to $((\Sigma \setminus N_i) \times [0, 1] \times \mathbb{R}, J)$

for a nested sequence of open subsets $N_i \subseteq N_{i-1} \subseteq \Sigma$ such that $\bigcap_i N_i$ consists of exactly two points (the two connected sum points), and a cylindrical almost complex structure J on $\Sigma \times [0, 1] \times \mathbb{R}$.

If $\phi \# \phi_0 \# \phi'$ admits a sequence u_i of \tilde{J}_i -holomorphic representatives, then we can extract a broken J -holomorphic representative of ϕ on $\Sigma \times [0, 1] \times \mathbb{R}$. By examining (5-29), we conclude that if $\text{gr}(x, y) \geq 0$ (which are the only cases relevant to Subclaim 5.10.1), then

$$\mu(\phi) = n_w(\phi_0) = 0.$$

Since ϕ has a broken representative and $\mu(\phi) = 0$, we conclude that ϕ must be the constant class e_x , by transversality. Since $n_w(\phi_0) = 0$, this also implies that ϕ_0 is one of the constant classes e_{θ^+} or e_{θ^-} , and ϕ' is the constant class e_c . Conversely, the constant class $e_x \times e_{\theta^\pm} \times e_c$ always has a \tilde{J}_i -holomorphic representative. The subclaim follows, and so does the main claim. \square

5.6 Invariance of the free-stabilization maps

In this section, we put our previous results together and prove invariance of the free-stabilization maps:

Theorem 5.11 *The free-stabilization maps S_w^+ and S_w^- defined in Section 5.1 determine well-defined chain maps on the level of transitive systems of chain complexes.*

Before we prove Theorem 5.11, we need the following topological result about embeddings of Heegaard diagrams:

Lemma 5.12 *Suppose (Y, \mathbf{w}) is a multipointed 3-manifold, and $\mathbf{p} \subseteq Y \setminus \mathbf{w}$ is a finite collection of points. If $\mathcal{H}_1 = (\Sigma_1, \alpha_1, \beta_1, \mathbf{w})$ and $\mathcal{H}_2 = (\Sigma_2, \alpha_2, \beta_2, \mathbf{w})$ are two Heegaard diagrams such that $\mathbf{p} \subseteq \Sigma_i \setminus (\alpha_i \cup \beta_i)$ for $i = 1, 2$, then \mathcal{H}_1 and \mathcal{H}_2 can be connected by a sequence of the following moves:*

- (1) *handleslides and isotopies of the α and β curves (possibly passing over the points in \mathbf{p});*
- (2) *simple stabilizations, away from $\mathbf{w} \cup \mathbf{p}$;*
- (3) *isotopies $\phi_t: \Sigma \rightarrow Y$ of the Heegaard surface which are fixed on $\mathbf{w} \cup \mathbf{p}$ for all t .*

Proof Consider Morse functions with gradient-like vector fields (f_1, v_1) and (f_2, v_2) on Y which induce \mathcal{H}_1 and \mathcal{H}_2 (i.e., Σ_i is a middle level set of f_i , and α_i is the intersection of the ascending manifolds of the index 1 critical points of f_i with Σ , and β_i is the intersection of the descending manifolds of the index 2 critical points of f_i). We may pick a path $(f_t, v_t)_{t \in [1,2]}$ of functions with gradient-like vector fields, connecting (f_1, v_1) and (f_2, v_2) , such that the following hold:

- (1) The pair (f_t, v_t) is Morse–Smale at all but finitely many t .
- (2) At the finitely many t where f_t fails to be Morse, an index $\frac{1}{2}$ birth or death singularity occurs (corresponding to a simple stabilization of the Heegaard surface).
- (3) At the finitely many t where the pair (f_t, v_t) fails to be Smale, a handleslide between two α curves or β curves occurs.

This implies that the three stated moves suffice to connect \mathcal{H}_1 and \mathcal{H}_2 , except that we have not yet ensured that an isotopy appearing in an instance of move (3) can be taken to respect the points \mathbf{p} .

Generically, we can assume the following:

- (1) \mathbf{p} is disjoint from the flowlines of any v_t passing through \mathbf{w} .
- (2) \mathbf{p} is disjoint from the descending manifolds of any index 1 critical point of any f_t .
- (3) \mathbf{p} is disjoint from the ascending manifolds of any index 2 critical point of any f_t .

Let us write Σ_t for a middle level set of each f_t . We can thus assume that for each $p \in \mathbf{p}$ and $t \in [1, 2]$, there is a flowline $f_{t,p}$ which connects p with a point on $\Sigma_t \setminus \mathbf{w}$. By performing an isotopy of each Σ_t in a neighborhood of $f_{t,p}$, we can ensure that $\mathbf{p} \subseteq \Sigma_t$ for all t . Consequently, if ϕ_t is an isotopy appearing in an instance of move (3), we can assume that $\mathbf{p} \subseteq \phi_t(\Sigma)$.

Finally, if $\phi_t: \Sigma \rightarrow Y$ is an isotopy which fixes \mathbf{w} and is such that $\mathbf{p} \subseteq \phi_t(\Sigma)$ for all t , there will, generically, be finitely many t where the image of a point in \mathbf{p} intersects an α or β curve. Such an isotopy ϕ_t may thus be decomposed as a sequence of isotopies where \mathbf{p} is fixed, as well as isotopies of the type appearing in move (1), where the α and β curves are moved but the Heegaard surface is fixed. \square

Proof of Theorem 5.11 As S_w^- is dual to S_w^+ , we focus on the claim for S_w^+ . The proof amounts to showing that if $\mathcal{H}_1 = (\Sigma_1, \alpha_1, \beta_1, \mathbf{w})$ and $\mathcal{H}_2 = (\Sigma_2, \alpha_2, \beta_2, \mathbf{w})$ are two diagrams such that $w \in \Sigma_i \setminus (\alpha_i \cup \beta_i)$, and ∂_1 and ∂_2 are two gluing data, then the following diagram commutes up to chain homotopy:

$$\begin{array}{ccc}
 \text{CF}_{J^{\partial_1}}^-(\mathcal{H}_1, \mathfrak{s}) & \xrightarrow{\Psi_{J^{\partial_1} \rightarrow J^{\partial_2}}} & \text{CF}_{J^{\partial_2}}^-(\mathcal{H}_2, \mathfrak{s}) \\
 \downarrow S_w^+ & & \downarrow S_w^+ \\
 \text{CF}_{J^{\partial_1(T_1)}}^-(\mathcal{H}_1^+, \mathfrak{s}) & \xrightarrow{\Psi_{(\mathcal{H}_1^+, J^{\partial_1(T_1)}) \rightarrow (\mathcal{H}_2^+, J^{\partial_2(T_2))}}} & \text{CF}_{J^{\partial_2(T_2)}}^-(\mathcal{H}_2^+, \mathfrak{s})
 \end{array}
 \tag{5-30}$$

It suffices to show that the diagram in (5-30) commutes when \mathcal{H}_1 and \mathcal{H}_2 differ by one of the moves listed in Lemma 5.12 (for $\mathbf{p} = \{w\}$). First, we consider move (1), isotopies and handleslides of the α and β curves on a fixed Heegaard surface. The transition maps for handleslides and isotopies of the α and β curves can be computed using a triangle map. The triangle counts from Theorem 5.7 imply that the free-stabilization maps commute with such transition maps. Move (2), commutation of the free-stabilization maps with the simple stabilization maps, follows from Lemma 5.10. Finally, commutation with respect to move (3), isotopies of the Heegaard surface which fix $\mathbf{w} \cup \{w\}$, is a tautology. \square

5.7 Commuting free-stabilization and relative homology maps

Lemma 5.13 *Let (Y, \mathbf{w}) be a multipointed 3-manifold, and let $w \in Y \setminus \mathbf{w}$. Suppose that λ is either a closed path in Y or a path which connects two basepoints in \mathbf{w} . Then*

$$S_w^+ \circ A_\lambda + A_\lambda \circ S_w^+ \simeq 0 \quad \text{and} \quad S_w^- \circ A_\lambda + A_\lambda \circ S_w^- \simeq 0.$$

Proof For simplicity, let us consider the claim involving S_w^+ , since S_w^- is the dual map.

Let $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ be a diagram for (Y, \mathbf{w}) , and let \mathcal{H}_w^+ be its free stabilization at w . We will show that, for a sufficiently stretched almost complex structure,

$$S_w^+ \circ A_\lambda + A_\lambda \circ S_w^+ = 0.$$

Suppose that $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is an intersection point on \mathcal{H} . By definition,

$$(5-31) \quad (S_w^+ \circ A_\lambda)(\mathbf{x}) = \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} a(\lambda, \phi) \cdot \#\widehat{\mathcal{M}}_{J^\circ}(\phi) \cdot U_w^{n_w(\phi)} \cdot (\mathbf{y} \times \theta^+),$$

while

$$(5-32) \quad (A_\lambda \circ S_w^+)(\mathbf{x}) = \sum_{\substack{\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \mathbf{y}) \\ \mu(\phi \# \phi_0)=1}} a(\lambda, \phi \# \phi_0) \cdot \#\widehat{\mathcal{M}}_{J^\circ(T)}(\phi \# \phi_0) \cdot U_w^{n_w(\phi)} U_w^{n_w(\phi_0)} \cdot (\mathbf{y} \times \mathbf{y}).$$

We can assume that λ is disjoint from the free-stabilization region on Σ . Consequently,

$$(5-33) \quad a(\lambda, \phi \# \phi_0) = a(\lambda, \phi).$$

By following the proof of Proposition 5.5, we see that there are two types of index 1 classes in any $\pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \mathbf{y})$ which can have holomorphic representatives for sufficiently large T .

The first type of class which could have holomorphic representatives has $\mathbf{y} = \theta^-$, $\phi = e_{\mathbf{x}}$, and has ϕ_0 equal to a bigon supported in the free-stabilization region. For such a class, we have $a(\lambda, \phi \# \phi_0) = a(\lambda, e_{\mathbf{x}}) = 0$, by (5-33), so such classes make no contribution to (5-32).

The second type of class which could have holomorphic representatives has $\mu(\phi) = 1$, $\phi_0 \in \pi_2(\theta^+, \theta^+)$ and $n_w(\phi_0) = 0$. Using the counts of holomorphic disks representing classes in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^+)$ established in (5-14) and (5-15), together with (5-33), we conclude that (5-31) and (5-32) coincide. \square

5.8 Compositions of free stabilizations

We now prove that free-stabilization maps for different basepoints may be commuted.

Proposition 5.14 *For any $\circ_1, \circ_2 \in \{+, -\}$, the free-stabilization maps satisfy*

$$S_{w_1}^{\circ_1} S_{w_2}^{\circ_2} \simeq S_{w_2}^{\circ_2} S_{w_1}^{\circ_1},$$

as morphisms of transitive systems of chain complexes.

Proof The proof follows from a double neck-stretching argument similar to Lemma 5.10. We consider a diagram on which both free stabilizations have been performed, with underlying Heegaard surface Σ . Let J be a cylindrical almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$, viewed as a complex structure for the unstabilized diagram. If $T = (T_1, T_2)$ is a pair of neck lengths, let $J(T)$ denote an almost complex structure obtained by inserting two necks of length T_1 and T_2 .

We recall that the free-stabilization maps require stabilizing condition (SC2) to be satisfied; however, it is not immediately clear that condition (SC2) can be achieved on two necks simultaneously.

We prove the following:

Subclaim 5.14.1 Suppose that T and T' are two pairs of neck lengths for free-stabilizing at w_1 and w_2 . If all the components of T and T' are sufficiently large, then a noncylindrical almost complex structure \tilde{J} may be chosen which interpolates $J(T)$ and $J(T')$ and satisfies

$$\begin{aligned}
 \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^+ \times \theta_2^+) &= \mathbf{x} \times \theta_1^+ \times \theta_2^+, \\
 \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^+ \times \theta_2^-) &= \mathbf{x} \times \theta_1^+ \times \theta_2^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} (C_{\mathbf{x},\mathbf{y}}^1 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^+ + C_{\mathbf{x},\mathbf{y}}^2 \cdot \mathbf{y} \times \theta_1^- \times \theta_2^+), \\
 \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^- \times \theta_2^+) &= \mathbf{x} \times \theta_1^- \times \theta_2^+ + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} (C_{\mathbf{x},\mathbf{y}}^3 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^+ + C_{\mathbf{x},\mathbf{y}}^4 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^-), \\
 \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^- \times \theta_2^-) &= \mathbf{x} \times \theta_1^- \times \theta_2^- \\
 &\quad + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} (C_{\mathbf{x},\mathbf{y}}^5 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^- + C_{\mathbf{x},\mathbf{y}}^6 \cdot \mathbf{y} \times \theta_1^- \times \theta_2^+ + C_{\mathbf{x},\mathbf{y}}^7 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^+)
 \end{aligned}
 \tag{5-34}$$

for various $C_{\mathbf{x},\mathbf{y}}^i \in \mathbb{F}_2[U_w]$, which depend on \tilde{J} .

Before proving Subclaim 5.14.1, we briefly explain why it implies the main claim. Pick $T = (T_1, T_2)$ so that T_2 is large enough to compute the free-stabilization maps at w_2 , and then pick $T_1 \gg T_2$ which is large enough to compute the free-stabilization maps at w_1 (after having already performed the stabilization at w_2). Next, pick $T' = (T'_1, T'_2)$ so that $T'_2 \gg T'_1$, so that T'_1 is large enough to compute the free-stabilization maps at w_1 first, and then subsequently compute the free-stabilization maps at w'_2 . Hence, to compute $S_{w_2}^+ S_{w_1}^+$ we may use $J(T')$, while to compute $S_{w_1}^+ S_{w_2}^+$ we may use $J(T)$. To compare the two compositions on the level of morphisms of transitive systems of chain complexes, we must compose with the transition map $\Psi_{J(T) \rightarrow J(T')}$. With this in mind, the first line of (5-34) implies that $S_{w_2}^+ S_{w_1}^+ \simeq S_{w_1}^+ S_{w_2}^+$.

Similarly, the first and the second lines of (5-34) together imply that $S_{w_1}^+ S_{w_2}^- \simeq S_{w_2}^- S_{w_1}^+$. The first and third lines together imply that $S_{w_1}^- S_{w_2}^+ \simeq S_{w_2}^+ S_{w_1}^-$. Finally, all four lines together imply that $S_{w_1}^- S_{w_2}^- \simeq S_{w_2}^- S_{w_1}^-$.

We now proceed to prove Subclaim 5.14.1. The transition map $\Psi_{\tilde{J}}$ counts \tilde{J} -holomorphic representatives of Maslov index 0 homology classes of disks. If $\phi \# \phi_0^1 \# \phi_0^2 \in \pi_2(\mathbf{x} \times x_1 \times x_2, \mathbf{y} \times y_1 \times y_2)$ is a homology class on the doubly free-stabilized diagram, then adapting the proof of (5-3) gives

$$\mu(\phi \# \phi_0^1 \# \phi_0^2) = \mu(\phi) + \text{gr}(x_1, y_1) + \text{gr}(x_2, y_2) + 2n_{w_1}(\phi) + 2n_{w_2}(\phi_2).
 \tag{5-35}$$

To prove (5-34), it suffices to make the following three observations:

- (1) Assuming T and T' are sufficiently large, if $\mu(\phi) = 0$ and $\text{gr}(x_1, y_1) + \text{gr}(x_2, y_2) > 0$, then $\mathcal{M}_{\tilde{J}}(\phi)$ is empty.
- (2) Assuming T and T' are sufficiently large, if $\mu(\phi) = 0$, $\text{gr}(x_1, y_1) = \text{gr}(x_2, y_2) = 0$ and $\mathcal{M}_{\tilde{J}}(\phi)$ is nonempty, then $\phi = e_{\mathbf{x}} \times e_{x_1} \times e_{x_2}$.
- (3) If ϕ is a class with $\text{gr}(x_1, y_1) + \text{gr}(x_2, y_2) \leq 0$ but $\text{gr}(x_1, y_1)$ and $\text{gr}(x_2, y_2)$ are not both zero, then the count of $\mathcal{M}_{\tilde{J}}(\phi)$ is irrelevant to (5-34).

Observations (1) and (2) are proven similarly to the proof of Proposition 5.3. For sufficiently large T and T' , by picking \tilde{J} appropriately, we can ensure that ϕ has a broken representative for a cylindrical almost complex structure. In particular $\mu(\phi) \geq 0$. Consequently, if $\text{gr}(x_1, y_1) + \text{gr}(x_2, y_2) > 0$, then (5-35) can never be zero, so there are no such classes with representatives. If (5-35) is zero and $\text{gr}(x_1, y_1) = \text{gr}(x_2, y_2) = 0$, then also $\mu(\phi) = 0$, so ϕ is a constant homology class, by transversality. Using additionally the vertex multiplicities near x_1 and x_2 , we conclude that ϕ_0^1 and ϕ_0^2 must also be constant homology classes, so the second observation follows. The third observation is straightforward. This establishes Subclaim 5.14.1, and hence the main claim. \square

We now prove a simple relation involving the free-stabilization maps:

Lemma 5.15 *Suppose (Y, \mathbf{w}) is a multipointed 3-manifold and $w \in Y \setminus \mathbf{w}$. Then*

$$S_w^- S_w^+ \simeq 0.$$

Proof The relation follows immediately from the formulas for the free-stabilization maps in (5-1). \square

6 Graph action map

If $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is a flow-graph in Y , and σ is a coloring of Γ (see Section 2.2 and Definition 3.2 for the definitions of a flow-graph and coloring, respectively), in this section we construct two *graph action maps*

$$\mathfrak{A}_{\mathcal{G}}, \mathfrak{B}_{\mathcal{G}}: \text{CF}^-(Y, \mathbf{w}_0^{\sigma_0}, \mathfrak{s}) \rightarrow \text{CF}^-(Y, \mathbf{w}_1^{\sigma_1}, \mathfrak{s}),$$

where $\sigma_i := \sigma|_{\mathbf{w}_i}$.

Unlike in Section 2.2, where flow-graphs are allowed to be immersed, we focus on embedded flow-graphs in this section. Upgrading the results to immersed flow-graphs is straightforward (the only complication is that when $\mathbf{w}_0 \cap \mathbf{w}_1 \neq \emptyset$, one must define $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ as compositions of two graph action maps).

We think of the graph action map as a restricted version of the full graph TQFT. Indeed, we will later prove that the graph cobordism map for $([0, 1] \times Y, \Gamma)$ is equal to the graph action map for the graph obtained by projecting Γ into Y ; see Theorem 9.5(3), below.

To define the graph action map, we fix a decomposition of Γ into edges and vertices. We write $V(\Gamma)$ and $E(\Gamma)$ for the vertices and edges of Γ , respectively. Although we require such a decomposition for our construction, the resulting graph action maps turn out to be unchanged by subdivision; see Lemma 6.14. Our construction of $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ also requires picking a further decomposition of \mathcal{G} into elementary flow-graphs (see Definition 6.3), as well as a choice of absolute lift of the ribbon structure.

In this section, we prove the following:

Theorem 6.1 *Suppose $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is a ribbon flow-graph in Y . The maps $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ described in this section satisfy the following:*

- (a) *The maps $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ are independent from the choice of decomposition into elementary flow-graphs.*




Figure 7: Examples of elementary ribbon flow-graphs. The top and bottom dashed lines indicate the incoming and outgoing vertices.

(b) The maps $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ are functorial: if \mathcal{G} is a flow-graph from \mathbf{w}_0 to \mathbf{w}_1 and \mathcal{G}' is a flow-graph from \mathbf{w}_1 to \mathbf{w}_2 , then

$$\mathfrak{A}_{\mathcal{G}'} \circ \mathfrak{A}_{\mathcal{G}} = \mathfrak{A}_{\mathcal{G}' \cup \mathcal{G}} \quad \text{and} \quad \mathfrak{B}_{\mathcal{G}'} \circ \mathfrak{B}_{\mathcal{G}} = \mathfrak{B}_{\mathcal{G}' \cup \mathcal{G}}.$$

(c) The maps $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{B}_{\mathcal{G}}$ are independent from the choice of absolute lift of the ribbon structure used in their construction.

6.1 Constructing the graph action maps

Before defining the graph action maps, we establish some terminology.

Definition 6.2 A *strong* ribbon flow-graph is a ribbon flow-graph equipped with an absolute lift of the cyclic ordering at each vertex.

Definition 6.3 A ribbon flow-graph $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is *elementary* if it is one of the following three subtypes (see Figure 7):

- (ER1) Translation: $|\mathbf{w}_0| = |\mathbf{w}_1|$ and each edge of Γ connects a vertex of \mathbf{w}_0 to a vertex of \mathbf{w}_1 .
- (ER2) Interior vertex: There is a single vertex w_s of Γ which is not contained in $\mathbf{w}_0 \cup \mathbf{w}_1$, and all edges of Γ connect either \mathbf{w}_0 to \mathbf{w}_1 , or a point of $\mathbf{w}_0 \cup \mathbf{w}_1$ to w_s . We call w_s the *special vertex*.
- (ER3) Local extrema: $|\mathbf{w}_0| = |\mathbf{w}_1| \pm 2$, and all but one edge connect \mathbf{w}_0 to \mathbf{w}_1 . There is a single edge e_0 which connects two vertices of \mathbf{w}_0 to each other, or connects two distinct vertices of \mathbf{w}_1 together. We call e_0 the *special edge*. There are two subtypes:
 - (i) Local max: The special edge connects two vertices of \mathbf{w}_0 .
 - (ii) Local min: The special edge connects two vertices of \mathbf{w}_1 .

Definition 6.4 Suppose that $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is an elementary ribbon flow-graph of type (ER2), with special vertex w_s . If e is an edge which is adjacent to w_s , we say that e has *positive initial slope* if e connects w_s to \mathbf{w}_1 , and we say that e has *negative initial slope* if e connects w_s to \mathbf{w}_0 .

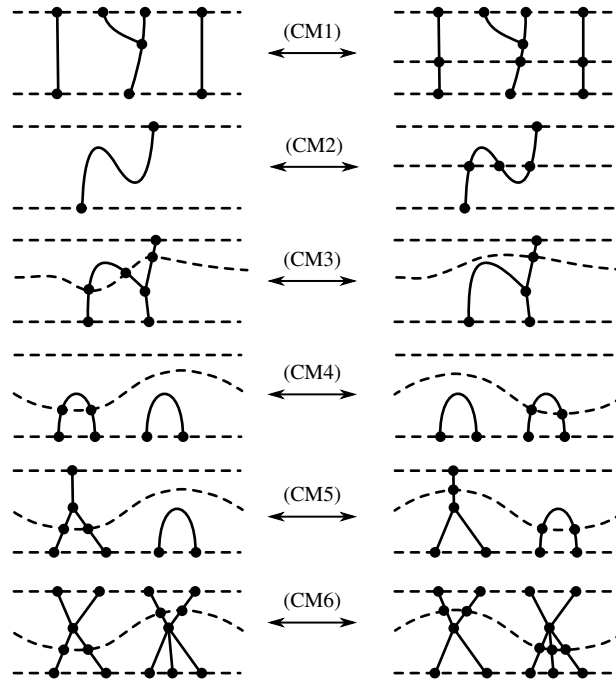


Figure 8: Examples of the Cerf moves from Definition 6.7.

Definition 6.5 (1) We say that a (strong) ribbon flow-graph $\mathcal{G}' = (\Gamma', \mathbf{w}_0, \mathbf{w}_1)$ is a *subdivision* of $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ if Γ' is obtained by adding extra vertices to the interior of edges of Γ . If \mathcal{G}' is a strong ribbon flow-graph, we allow the new vertices to be added with either absolute ordering.

(2) We say that two flow-graphs $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ and $\mathcal{G}' = (\Gamma', \mathbf{w}'_0, \mathbf{w}'_1)$ are *subdivision-equivalent* if \mathcal{G} and \mathcal{G}' have a common subdivision (in particular, $\Gamma = \Gamma'$ as subsets of Y).

We now define a notion of decomposition into elementary pieces for ribbon flow-graphs:

Definition 6.6 If $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is a ribbon flow-graph in Y , a *Cerf decomposition* of \mathcal{G} is a sequence of ribbon flow-graphs $\mathcal{G}_1, \dots, \mathcal{G}_n$, where $\mathcal{G}_i = (\Gamma_i, \mathbf{w}_{i,0}, \mathbf{w}_{i,1})$, such that the following hold:

- (1) Each \mathcal{G}_i is an elementary ribbon flow-graph.
- (2) $\mathbf{w}_{1,0} = \mathbf{w}_0$ and $\mathbf{w}_{n,1} = \mathbf{w}_1$.
- (3) $\mathbf{w}_{0,i+1} = \mathbf{w}_{1,i}$ for $i \in \{1, \dots, n-1\}$.
- (4) The concatenation $\mathcal{G}_n \circ \dots \circ \mathcal{G}_1$ is a subdivision of \mathcal{G} .
- (5) Each vertex of $V(\Gamma) \setminus (\mathbf{w}_0 \cup \mathbf{w}_1)$ is a special vertex of some \mathcal{G}_i . Furthermore, if w_s is a special vertex of any \mathcal{G}_i , then w_s is a vertex in $V(\Gamma)$.

We now define the following moves between Cerf decompositions of flow-graphs:

Definition 6.7 We say two Cerf decompositions of a flow-graph $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ in Y are *Cerf-equivalent* if they can be connected by a sequence of the following moves, or their inverses:

- (CM1) Level splitting: Replacing an elementary flow-graph \mathcal{G} with a subdivision-equivalent composition $\mathcal{G}_2 \circ \mathcal{G}_1$ such that one of \mathcal{G}_1 and \mathcal{G}_2 is of the same type as \mathcal{G} , and the other is of type (ER1).
- (CM2) Critical point birth/death: Replacing an elementary flow-graph \mathcal{G} of type (ER1) with a subdivision-equivalent composition $\mathcal{G}_2 \circ \mathcal{G}_1$ such that \mathcal{G}_2 and \mathcal{G}_1 are of type (ER3i) and (ER3ii), respectively.
- (CM3) Changing an initial slope: Replacing an elementary flow-graph \mathcal{G} of type (ER2) with a subdivision-equivalent composition $\mathcal{G}_2 \circ \mathcal{G}_1$ such that exactly one of \mathcal{G}_1 and \mathcal{G}_2 is of type (ER3), and the other is of type (ER2) (with the same special vertex as \mathcal{G}).
- (CM4) Critical value swap: Replacing the composition of two consecutive elementary flow-graphs $\mathcal{G}_2 \circ \mathcal{G}_1$ of type (ER3) with a subdivision-equivalent composition $\mathcal{G}'_2 \circ \mathcal{G}'_1$ such that both \mathcal{G}'_2 and \mathcal{G}'_1 are also of type (ER3). Furthermore, the special edges of \mathcal{G}_2 and \mathcal{G}_1 are disjoint, and the special edges of \mathcal{G}'_2 and \mathcal{G}'_1 are disjoint.
- (CM5) Vertex/critical value swap: Replacing a composition $\mathcal{G}_2 \circ \mathcal{G}_1$, where \mathcal{G}_2 is of type (ER2) and \mathcal{G}_1 of type (ER3), with a subdivision-equivalent composition $\mathcal{G}'_2 \circ \mathcal{G}'_1$ where \mathcal{G}'_2 is of type (ER3) and \mathcal{G}'_1 is of type (ER2). Furthermore, we assume that there is no path in $\mathcal{G}_2 \circ \mathcal{G}_1$ connecting the special vertex of \mathcal{G}_2 to the special edge of \mathcal{G}_1 .
- (CM6) Vertex value swap: Replacing a composition $\mathcal{G}_2 \circ \mathcal{G}_1$ of two elementary ribbon graphs of type (ER2) with a subdivision-equivalent composition $\mathcal{G}'_2 \circ \mathcal{G}'_1$, where \mathcal{G}'_2 and \mathcal{G}'_1 are also both of type (ER2). Furthermore, we assume that there is no path in $\mathcal{G}_2 \circ \mathcal{G}_1$ connecting the two special edges of \mathcal{G}_2 and \mathcal{G}_1 .

Examples of the moves from Definition 6.7 are shown in Figure 8.

We encourage the reader to compare the following to [Wehrheim and Woodward 2015, Theorem 5.9]:

Proposition 6.8 *Any two Cerf decompositions of a flow-graph \mathcal{G} are Cerf-equivalent (i.e., can be connected by the moves in Definition 6.7).*

Proof Suppose $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is a flow-graph in Y . We assume that each edge of Γ is smoothly immersed in Y (and that unless an edge has both endpoints on the same vertex, each edge is smoothly embedded). We say a function $f: \Gamma \rightarrow [0, 1]$ is *Morse* if the following hold:

- (1) $f^{-1}(0) = \mathbf{w}_0$ and $f^{-1}(1) = \mathbf{w}_1$.
- (2) If we write an edge e as the image of an immersion $\hat{e}: [0, 1] \rightarrow Y$, then $f \circ \hat{e}$ is Morse on $[0, 1]$ and 0 and 1 are not critical points of $f \circ \hat{e}$.

A Cerf decomposition can be obtained by picking a generic Morse function $f: \Gamma \rightarrow [0, 1]$, and picking a collection of values $0 = t_0 < t_1 < \dots < t_{n-1} < t_n = 1$ such that each $f^{-1}(t_i)$ contains no critical points or vertices, and each $f^{-1}([t_i, t_{i+1}])$ contains at most one critical point of f or vertex of $V(\Gamma) \setminus (\mathbf{w}_0 \cup \mathbf{w}_1)$. Furthermore, we assume that each $f^{-1}(t)$ is nonempty for all $t \in [0, 1]$. We can construct a Cerf decomposition $\mathcal{G}_n \circ \dots \circ \mathcal{G}_1$ of \mathcal{G} by setting

$$\mathcal{G}_i = (f^{-1}([t_{i-1}, t_i]), f^{-1}(t_{i-1}), f^{-1}(t_i)).$$

Conversely, it is not hard to see that any Cerf decomposition is induced by a Morse function, in the above sense.

Given two Cerf decompositions of \mathcal{G} , we let f_1 and f_2 be Morse functions on Γ which induce the two Cerf decompositions. We connect f_1 and f_2 as follows. First, we modify f_1 in a small neighborhood of each vertex so that it coincides with f_2 near $V(\Gamma)$. Write \tilde{f}_1 for the resulting Morse function. In particular, \tilde{f}_1 and f_2 have the same initial slopes along each edge (in the sense of Definition 6.4). The Cerf decompositions induced by f_1 to \tilde{f}_1 are related by repeated applications of move (CM3).

We take a generic perturbation of a linear homotopy between \tilde{f}_1 and f_2 , which is fixed in a neighborhood of the vertices $V(\Gamma)$. Write $(f_t)_{t \in [1,2]}$ for this path. Note that the critical set of f_t is bounded away from $V(\Gamma)$ since \tilde{f}_1 and f_2 coincide in a neighborhood of $V(\Gamma)$. By perturbing f_t slightly, the interval $[1, 2]$ can be subdivided by picking

$$1 = b_1 < b_2 < \dots < b_n = 2$$

so that, on each interval $[b_i, b_{i+1}]$, exactly one of the following holds:

- (1) f_t is Morse for all $t \in [b_i, b_{i+1}]$, and all vertices and critical points have distinct values.
- (2) f_t is Morse for all $t \in [b_i, b_{i+1}]$, and all vertices and critical points have distinct values, except at a single point $t_0 \in (b_i, b_{i+1})$, where two critical points or vertices exchange relative ordering.
- (3) f_t is Morse for all $t \in [b_i, b_{i+1}]$, except for at a single $t_0 \in (b_i, b_{i+1})$, where a critical point birth–death singularity occurs along the interior of an edge. All vertices and critical points have distinct values.

The above follows from standard Morse theory, as we now describe. If we view each edge e as a copy of $[0, 1]$, then the family f_t is fixed near 0 and 1 and hence extends to a Morse function on all of \mathbb{R} , which is fixed in the parameter t outside of the interior of e . The above moves are the generic codimension 1 singularities of smooth functions.

From these considerations, it is straightforward to see that \tilde{f}_1 and f_2 can be connected by a sequence of moves (CM1), (CM2), (CM4), (CM5) and (CM6).

Finally, we note one can ensure that all level sets of each function in the path f_t have nonempty level sets by first modifying f_1 and f_2 near w_0 , raising the value of a point near w_0 so that it is close to 1. This modification induces a sequence of moves (CM2), (CM4) and (CM5). □

We now define the graph action map. Suppose that $\mathcal{G} = (\Gamma, w_0, w_1)$ is a ribbon flow-graph in Y and $\mathfrak{s} \in \text{Spin}^c(Y)$. To define the map, we pick a strong ribbon structure on Γ which lifts the ribbon structure, as well as a Cerf decomposition of \mathcal{G} . We now define the graph action map for elementary strong ribbon flow-graphs.

If $\mathcal{G} = (\Gamma, w_0, w_1)$ is an elementary ribbon flow-graph of type (ER1) or (ER3), we define the graph action map by the formula

$$(6-1) \quad \mathfrak{A}_{\mathcal{G}} := \left(\prod_{w \in w_0} S_w^- \right) \circ \left(\prod_{e \in E(\Gamma)} A_e \right) \circ \left(\prod_{w \in w_1} S_w^+ \right).$$

The map $\mathfrak{B}_{\mathcal{G}}$ is defined using a similar formula, with the relative homology maps B_e in place of A_e .

In an elementary flow-graph of type (ER1) or (ER3), no edges share a common vertex. Consequently, the product of the relative homology maps appearing in (6-1) does not affect chain homotopy type of the composition, by Lemma 4.4. By Proposition 5.14, the ordering of the free-stabilization maps within the left and right factors of (6-1) also does not affect the composition.

Next, suppose that \mathcal{G} is an elementary strong ribbon flow-graph of type (ER2). Let w_s denote the special vertex, and let e_1, \dots, e_n denote the edges incident to w_s , ordered according to the strong ribbon structure. We define

$$(6-2) \quad \mathfrak{A}_{\mathcal{G}} := \left(\prod_{w \in \mathbf{w}_0 \cup \{w_s\}} S_w^- \right) \circ \left(\prod_{e \in E(\Gamma) \setminus \{e_1, \dots, e_n\}} A_e \right) \circ (A_{e_n} \circ \dots \circ A_{e_1}) \circ \left(\prod_{w \in \mathbf{w}_1 \cup \{w_s\}} S_w^+ \right).$$

The map $\mathfrak{B}_{\mathcal{G}}$ is defined using the maps B_e in place of A_e .

If $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is an arbitrary, strong ribbon flow-graph in Y , the graph action map is defined by picking a Cerf decomposition

$$\mathcal{G} = \mathcal{G}_n \circ \dots \circ \mathcal{G}_1$$

and setting

$$\mathfrak{A}_{\mathcal{G}} := \mathfrak{A}_{\mathcal{G}_n} \circ \dots \circ \mathfrak{A}_{\mathcal{G}_1}.$$

The type- B map $\mathfrak{B}_{\mathcal{G}}$ is defined similarly.

Remark 6.9 Since the graph action map $\mathfrak{A}_{\mathcal{G}}$ is defined using a relative homology map for each edge, the map $\mathfrak{A}_{\mathcal{G}}$ is unchanged by replacing an edge e of \mathcal{G} with another edge e' such that $\partial e = \partial e'$ and $e \cup e'$ is a null-homologous loop in Y .

6.2 Algebraic relations in the graph TQFT

In this section, we prove some algebraic relations involving the free-stabilization maps and the relative homology maps which will be useful in proving Theorem 6.1.

We call the following relation the *trivial strand relation* (the terminology is justified by Lemma 6.15, below):

Lemma 6.10 *Suppose that (Y, \mathbf{w}) is a multipointed 3-manifold, $w \in Y \setminus \mathbf{w}$ is a new basepoint, and $\lambda \subseteq Y$ is a path which connects w to a point $w_0 \in \mathbf{w}$. Suppose that $\sigma: \mathbf{w} \rightarrow \mathbf{P}$ and $\sigma': \mathbf{w} \cup \{w\} \rightarrow \mathbf{P}$ are colorings such that $\sigma'|_{\mathbf{w}} = \sigma$ and $\sigma'(w) = \sigma'(w_0)$. Then, with respect to the complexes which are algebraically colored using σ and σ' , we have*

$$S_w^- A_{\lambda} S_w^+ \simeq S_w^- B_{\lambda} S_w^+ \simeq \text{id}_{\text{CF}^-(Y, \mathbf{w}^{\sigma, s})}.$$

Proof We focus on the relation $S_w^- A_{\lambda} S_w^+ \simeq \text{id}_{\text{CF}^-(Y, \mathbf{w}^{\sigma, s})}$. The relation involving $S_w^- B_{\lambda} S_w^+$ is proven similarly.

A single free-stabilized diagram and stretched almost complex structure can be chosen to compute S_w^+ , A_{λ} and S_w^- . Such a diagram is shown in Figure 9.




Figure 9: One of the two bigons contributing to (6-4).

Using the definition of the free-stabilization maps in (5-1), it suffices to show that for a sufficiently stretched almost complex structure, the map A_λ satisfies

$$(6-3) \quad A_\lambda(x \times \theta^+) = x \times \theta^- + \sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{x,y} \cdot y \times \theta^+$$

for some $C_{x,y} \in \mathcal{RP}$. In (6-3), θ^+ and θ^- denote the two intersection points in the free-stabilized region.

Equation (6-3) follows from the proof of Proposition 5.5. More explicitly, for a sufficiently stretched almost complex structure, any Maslov index 1 class in $\pi_2(x \times \theta^+, y \times \theta^-)$ which supports holomorphic representatives has domain equal to one of the two bigons in the free-stabilization region. Writing ϕ_x^1 and ϕ_x^2 for these two classes, we note that both ϕ_x^1 and ϕ_x^2 have a unique holomorphic representative. An easy model computation (see Figure 9) shows that

$$(6-4) \quad a(\lambda, \phi_x^1) + a(\lambda, \phi_x^2) \equiv 1 \pmod{2}.$$

Equation (6-3) now follows from the definition of the map A_λ in (4-2). □

The following relation is related to subdivision-invariance of the graph TQFT:

Lemma 6.11 *Suppose that w_0, w_1 and w_2 are basepoints in Y , $w_1 \notin \{w_0, w_2\}$, and \mathbf{w} is a collection of basepoints containing w_0 and w_2 , but not w_1 . Suppose that λ_1 and λ_2 are paths in Y satisfying $\partial\lambda_i = \{w_{i-1}, w_i\}$, and write $\lambda_2 * \lambda_1$ for the concatenation. Then*

$$A_{\lambda_2 * \lambda_1} \simeq S_{w_1}^- A_{\lambda_2} A_{\lambda_1} S_{w_1}^+,$$

as endomorphisms of $CF^-(Y, \mathbf{w}^\sigma, \mathfrak{s})$, for a coloring σ which identifies U_{w_0}, U_{w_1} and U_{w_2} with a common variable, U .

Proof We compute

$$\begin{aligned} A_{\lambda_2 * \lambda_1} &\simeq A_{\lambda_2 * \lambda_1} (S_{w_1}^- A_{\lambda_1} S_{w_1}^+) && \text{(Lemma 6.10)} \\ &\simeq S_{w_1}^- A_{\lambda_2 * \lambda_1} A_{\lambda_1} S_{w_1}^+ && \text{(Lemma 5.13)} \\ &\simeq S_{w_1}^- (A_{\lambda_2} + A_{\lambda_1}) A_{\lambda_1} S_{w_1}^+ && \text{(Lemma 4.3)} \\ &\simeq S_{w_1}^- A_{\lambda_2} A_{\lambda_1} S_{w_1}^+ + S_{w_1}^- U S_{w_1}^+ && \text{(Lemma 4.5)} \\ &\simeq S_{w_1}^- A_{\lambda_2} A_{\lambda_1} S_{w_1}^+ && \text{(Lemma 5.15).} \end{aligned}$$

□




Figure 10: The configuration of vertices and edges in Lemma 6.12.

In the following lemma, and also henceforth, we write $S_{w_n w_{n-1} \dots w_1}^+$ for the composition

$$S_{w_n w_{n-1} \dots w_1}^+ := S_{w_n}^+ \circ S_{w_{n-1}}^+ \circ \dots \circ S_{w_1}^+,$$

and use similar notation for negative free stabilizations. By Proposition 5.14, the map $S_{w_n w_{n-1} \dots w_1}^+$ is independent of the ordering of the basepoints w_1, \dots, w_n .

Lemma 6.12 *Suppose that w_0, w_1, w_2 and w_3 are distinct points in Y , and λ_1, λ_2 and λ_3 are paths in Y such that $\partial\lambda_i = \{w_{i-1}, w_i\}$. See Figure 10. Furthermore, suppose that \mathbf{w} is collection of basepoints on Y , containing w_0 and w_3 , but not w_1 and w_2 . If $\tau: \{1, 2, 3\} \rightarrow \{1, 2, 3\}$ is a permutation, then*

$$S_{w_1 w_2}^- A_{\lambda_{\tau(3)}} A_{\lambda_{\tau(2)}} A_{\lambda_{\tau(1)}} S_{w_1 w_2}^+ \simeq S_{w_1 w_2}^- A_{\lambda_3} A_{\lambda_2} A_{\lambda_1} S_{w_1 w_2}^+,$$

as endomorphisms of $CF^-(Y, \mathbf{w}^\sigma, \mathfrak{s})$, where σ is a coloring which identifies $U_{w_0}, U_{w_1}, U_{w_2}$ and U_{w_3} with a common variable, U .

Proof It is sufficient to show that if $\{i, j, k\} = \{1, 2, 3\}$ as sets, then the map $S_{w_1 w_2}^- A_{\lambda_i} A_{\lambda_j} A_{\lambda_k} S_{w_1 w_2}^+$ is invariant (up to chain homotopy) under switching the order of λ_j and λ_k or switching the order of λ_i and λ_j . We will focus on the proof that the map is invariant under switching the order of λ_j and λ_k ; the proof for switching λ_i and λ_j is analogous.

There are two cases:

- (1) λ_j and λ_k are disjoint.
- (2) λ_j and λ_k share a vertex.

If λ_j and λ_k are disjoint, then $A_{\lambda_j} A_{\lambda_k} \simeq A_{\lambda_k} A_{\lambda_j}$ by Lemma 4.4. Hence, it remains to consider the case that λ_j and λ_k share a vertex $w_n \in \{w_1, w_2\}$. Let w_m denote the other vertex in $\{w_1, w_2\}$. We compute

$$\begin{aligned} S_{w_1 w_2}^- A_{\lambda_i} A_{\lambda_j} A_{\lambda_k} S_{w_1 w_2}^+ &\simeq S_{w_1 w_2}^- A_{\lambda_i} A_{\lambda_k} A_{\lambda_j} S_{w_1 w_2}^+ + U S_{w_1 w_2}^- A_{\lambda_i} S_{w_1 w_2}^+ \\ &\simeq S_{w_1 w_2}^- A_{\lambda_i} A_{\lambda_k} A_{\lambda_j} S_{w_1 w_2}^+ + U S_{w_m}^- S_{w_n}^- S_{w_n}^+ A_{\lambda_i} S_{w_m}^+ \\ &\simeq S_{w_1 w_2}^- A_{\lambda_i} A_{\lambda_k} A_{\lambda_j} S_{w_1 w_2}^+. \end{aligned}$$

The first chain homotopy is justified by Lemma 4.4. The second chain homotopy is justified by Lemma 5.13, noting that λ_i is disjoint from w_n . The final chain homotopy is justified by Lemma 5.15. \square

6.3 Independence from the Cerf decomposition

We now prove that the graph action map is independent of the choice of decomposition into elementary flow-graphs:




Figure 11: Subdividing an elementary flow-graph of type (ER1).

Proof of Theorem 6.1(a) We focus on $\mathfrak{A}_{\mathcal{G}}$, since the argument for $\mathfrak{B}_{\mathcal{G}}$ is no different. By Proposition 6.8, it is sufficient to show invariance under moves (CM1)–(CM6). Suppose that an absolute lift of the ribbon structure has been fixed.

We first consider move (CM1), splitting levels. Consider first when \mathcal{G} is of type (ER1) (translation flow-graph). Further restrict first to the case when \mathcal{G} consists of a single edge e , which goes from w_0 to w_1 . The map $\mathfrak{A}_{\mathcal{G}}$ is defined to be

$$(6-5) \quad \mathfrak{A}_{\mathcal{G}} := S_{w_0}^- A_e S_{w_1}^+.$$

Write e as the concatenation of two edges, e_1 and e_2 , such that e_1 goes from w_0 to w' , and e_2 goes from w' to w_1 . Let \mathcal{G}_1 denote the flow-graph (e_1, w_0, w') , and \mathcal{G}_2 denote the flow-graph (e_2, w', w_1) . See Figure 11.

By definition,

$$(6-6) \quad \mathfrak{A}_{\mathcal{G}_2} \circ \mathfrak{A}_{\mathcal{G}_1} := (S_{w'}^- A_{e_2} S_{w_1}^+) (S_{w_0}^- A_{e_1} S_{w'}^+).$$

Using Lemma 5.13 and Proposition 5.14, we see (6-6) is chain homotopic to

$$(6-7) \quad S_{w_0}^- S_{w'}^- A_{e_2} A_{e_1} S_{w'}^+ S_{w_1}^+.$$

Using Lemma 6.11, we conclude that (6-7) is chain homotopic to $S_{w_0}^- A_{e_2 * e_1} S_{w_1}^+$, which coincides with the expression for $\mathfrak{A}_{\mathcal{G}}$ in (6-5).

The general case of splitting an elementary flow-graph \mathcal{G} of type (ER1) with more than one strand is no different: one applies the above manipulation to each strand, noting that the terms associated with different strands of \mathcal{G} commute by Lemma 5.13 and Proposition 5.14.

The above subdivision technique also works when \mathcal{G} is of type (ER2) (interior vertex flow-graphs).

We now consider the case that \mathcal{G} is of type (ER3) (local extrema flow-graph). Let e be the special edge of \mathcal{G} , and suppose further that \mathcal{G} is of type (ER3i), i.e., e connects two incoming vertices, w_1 and w_2 .




Figure 12: Subdividing an elementary flow-graph of type (ER3i).




Figure 13: Applying move (CM2) to a flow-graph of type (ER1).

Let w_0 and w'_0 be two new vertices on the interior of e , and let e_1, e_2 and e_3 be the components of $e \setminus \{w_0, w'_0\}$, as shown in Figure 12.

We focus on the terms of $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{A}_{\mathcal{G}_2} \circ \mathfrak{A}_{\mathcal{G}_1}$ corresponding to e and its subdivision. The terms corresponding to the other strands can be subdivided using the procedure described above for flow-graphs of type (ER1).

The terms corresponding to e in $\mathfrak{A}_{\mathcal{G}}$ are

$$(6-8) \quad S_{w_1 w_2}^- A_e.$$

The corresponding terms of $\mathfrak{A}_{\mathcal{G}_2} \circ \mathfrak{A}_{\mathcal{G}_1}$ are

$$(6-9) \quad S_{w_0 w'_0}^- A_{e_2} S_{w_1 w_2}^- A_{e_1} A_{e_3} S_{w_0 w'_0}^+.$$

We compute

$$\begin{aligned} S_{w_1 w_2}^- A_e &\simeq S_{w_1 w_2 w'_0}^- A_{e_1 * e_2} A_{e_3} S_{w'_0}^+ && \text{(Lemma 6.11)} \\ &\simeq S_{w_1 w_2 w_0 w'_0}^- A_{e_2} A_{e_1} S_{w_0}^+ A_{e_3} S_{w'_0}^+ && \text{(Lemma 6.11)} \\ &\simeq S_{w_1 w_2 w_0 w'_0}^- A_{e_2} A_{e_1} A_{e_3} S_{w'_0 w_0}^+ && \text{(Lemma 5.13)} \\ &\simeq S_{w_0 w'_0}^- A_{e_2} S_{w_1 w_2}^- A_{e_1} A_{e_3} S_{w'_0 w_0}^+ && \text{(Lemma 5.13),} \end{aligned}$$

which coincides with (6-9). We conclude that level splitting does not change the homotopy type of the graph action map for elementary ribbon flow-graphs of type (ER3i). The argument for graphs of type (ER3ii) is a simple modification.

We now show independence from move (CM2), a critical point birth along the interior of an edge. By using move (CM1) we may assume that the new critical points occur inside a flow-graph of type (ER1). By using Lemma 5.13 and Proposition 5.14, it is sufficient to consider the case when \mathcal{G} consists of a single edge e , and we divide \mathcal{G} into a composition $\mathcal{G}_2 \circ \mathcal{G}_1$. Let e_1, e_2, e_3 and e_4 be the edges of $\mathcal{G}_2 \circ \mathcal{G}_1$, and let w_1, w_2, w_3, w_4 and w_5 be the vertices, as in Figure 13.

By definition,

$$(6-10) \quad \mathfrak{A}_{\mathcal{G}} := S_{w_1}^- A_e S_{w_5}^+,$$

while

$$(6-11) \quad \mathfrak{A}_{\mathcal{G}_2} \circ \mathfrak{A}_{\mathcal{G}_1} := S_{w_2 w_3 w_4}^- A_{e_4} A_{e_2} S_{w_5}^+ S_{w_1}^- A_{e_3} A_{e_1} S_{w_2 w_3 w_4}^+.$$




Figure 14: Applying move (CM3) to a flow-graph of type (ER2).

Using Lemma 5.13 and Proposition 5.14, equation (6-11) can be rearranged to

$$(6-12) \quad S_{w_1 w_2 w_4}^- A_{e_4} (S_{w_3}^- A_{e_2} A_{e_3} S_{w_3}^+) A_{e_1} S_{w_2 w_4 w_5}^+$$

Using Lemma 6.11, (6-12) becomes

$$S_{w_1 w_2 w_4}^- A_{e_4} A_{e_2 * e_3} A_{e_1} S_{w_2 w_4 w_5}^+$$

Proceeding similarly, using the aforementioned relations, we compute

$$\begin{aligned} S_{w_1 w_2 w_4}^- A_{e_4} A_{e_2 * e_3} A_{e_1} S_{w_2 w_4 w_5}^+ &\simeq S_{w_1 w_2}^- (S_{w_4}^- A_{e_4} A_{e_2 * e_3} S_{w_4}^+) A_{e_1} S_{w_2 w_5}^+ \\ &\simeq S_{w_1 w_2}^- A_{e_2 * e_3 * e_4} A_{e_1} S_{w_2 w_5}^+ \\ &\simeq S_{w_1}^- (S_{w_3}^- A_{e_2 * e_3 * e_4} A_{e_1} S_{w_2}^+) S_{w_5}^+ \\ &\simeq S_{w_1}^- A_e S_{w_5}^+. \end{aligned}$$

Hence, (6-10) and (6-11) are chain homotopic, and invariance under move (CM2) is established.

We now consider invariance under move (CM3) (changing an initial slope). Suppose \mathcal{G} is an elementary flow-graph of type (ER2) (an interior vertex flow-graph), with a special vertex w_s , and e_1, \dots, e_n are the edges incident to w_s , ordered according to the chosen absolute lift of the cyclic ordering. Suppose we wish to change the initial slope of e_i , and that e_i has downward initial slope in \mathcal{G} . We decompose \mathcal{G} as a composition $\mathcal{G}_2 \circ \mathcal{G}_1$, where \mathcal{G}_1 is of type (ER2), with special vertex w_s , and \mathcal{G}_2 is of type (ER3i) (a local max). We decompose e_i as the concatenation of e, e' and e'' , and let w' and w'' denote the new vertices, as in Figure 14.

Write w_0 for the incoming vertices of \mathcal{G} , and w_1 for the outgoing vertices. By the same argument as for invariance under move (CM1), we can reduce to the case that there are no components of \mathcal{G} which are disjoint from w_s . By definition,

$$(6-13) \quad \mathfrak{A}_{\mathcal{G}} := S_{w_0 \cup \{w_s\}}^- A_{e_n} \cdots A_{e_i} \cdots A_{e_1} S_{w_1 \cup \{w_s\}}^+$$

Using Lemmas 6.11 and 5.13, we have

$$(6-14) \quad A_{e_i} \simeq S_{w''}^- A_{e' * e} A_{e''} S_{w''}^+ \simeq S_{w'' w'}^- A_{e'} A_e S_{w''}^+ A_{e''} S_{w''}^+ \simeq S_{w'' w'}^- A_{e'} A_e A_{e''} S_{w'' w'}^+$$

Combining (6-13) and (6-14), we see that

$$\mathfrak{A}_{\mathcal{G}} \simeq S_{w_0 \cup \{w_s\}}^- A_{e_n} \cdots A_{e_{i+1}} (S_{w'' w'}^- A_{e'} A_e A_{e''} S_{w'' w'}^+) A_{e_{i-1}} \cdots A_{e_1} S_{w_1 \cup \{w_s\}}^+$$

Since e', e'', w'' and w' are disjoint from e_j when $j \neq i$, using Lemmas 4.4 and 5.13, we compute that

$$\begin{aligned} S_{w_0 \cup \{w_s\}}^- A_{e_n} \cdots A_{e_{i+1}} (S_{w''w'}^- A_{e'} A_e A_{e''} S_{w''w'}^+) A_{e_{i-1}} \cdots A_{e_1} S_{w_1 \cup \{w_s\}}^+ \\ \simeq (S_{w''w'}^- A_{e'}) S_{w_0 \cup \{w_s\}}^- (A_{e_n} \cdots A_{e_{i+1}} A_e A_{e_{i-1}} \cdots A_{e_1}) A_{e''} S_{w_1 \cup \{w'', w', w_s\}}^+ . \end{aligned}$$

The above expression is almost the expression for $A_{G_2} \circ A_{G_1}$, the only difference being that the remaining edges $e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_n$ with upward initial slope have not yet been subdivided. By using Lemma 6.11 to subdivide the remaining edges with upward initial slope, and then commuting terms exactly as in move (CM1), we arrive at the definition of $\mathfrak{A}_{G_2} \circ \mathfrak{A}_{G_1}$, completing the proof of invariance under move (CM3).

Invariance under moves (CM4), (CM5) and (CM6) can be proven by adapting the above arguments. For example, invariance under move (CM4) (critical value swaps) can be proven by using the manipulation from the proof of invariance of the maps for elementary flow-graphs of type (ER3) (local minima and maxima) under move (CM1) (level splitting) to subdivide one of the edges with a local extrema. Since the two components of the graph with local extrema are disjoint, all terms associated to one component can be commuted past the terms associated to the other component, using Lemmas 4.4 and 5.13 and Proposition 5.14. A similar strategy can be used for moves (CM5) and (CM6). □

Since the graph action is defined by taking a decomposition of \mathcal{G} into elementary flow-graphs and composing the maps for each piece, functoriality (Theorem 6.1(b)) is automatic.

6.4 Cyclic ordering

In our construction of the graph action map, we chose an absolute lift of the cyclic orderings. In this section, we show that the graph action map is independent of the choice of absolute lift. The following lemma implies Theorem 6.1(c):

Lemma 6.13 *Suppose that e_1, \dots, e_n are edges in Y such that w_s has valence 1 in each e_i , and $e_i \cap e_j = \{w_s\}$ for all i and j . For a coloring which identifies all of the variables U_{w_1}, \dots, U_{w_n} with a common variable U , we have*

$$S_{w_s}^- A_{e_n} A_{e_{n-1}} \cdots A_{e_1} S_{w_s}^+ \simeq S_{w_s}^- A_{e_{n-1}} \cdots A_{e_1} A_{e_n} S_{w_s}^+ .$$

Proof The proof is by induction. The $n = 1$ case is vacuous. The $n = 2$ case is easy: using Lemma 4.4, we compute

$$S_{w_s}^- A_{e_2} A_{e_1} S_{w_s}^+ \simeq S_{w_s}^- U S_{w_s}^+ + S_{w_s}^- A_{e_1} A_{e_2} S_{w_s}^+ \simeq S_{w_s}^- A_{e_1} A_{e_2} S_{w_s}^+ .$$

The $n = 3$ case is also relatively easy. We compute

$$S_{w_s}^- A_{e_3} A_{e_2} A_{e_1} S_{w_s}^+ \simeq S_{w_s}^- A_{e_2} A_{e_3} A_{e_1} S_{w_s}^+ + U S_{w_s}^- A_{e_1} S_{w_s}^+ \tag{Lemma 4.4}$$

$$\simeq S_{w_s}^- A_{e_2} A_{e_1} A_{e_3} S_{w_s}^+ + U S_{w_s}^- (A_{e_2} + A_{e_1}) S_{w_s}^+ \tag{Lemma 4.4}$$

$$\simeq S_{w_s}^- A_{e_2} A_{e_1} A_{e_3} S_{w_s}^+ + U S_{w_s}^- A_{e_2 * e_1} S_{w_s}^+ \tag{Lemma 4.3}$$

$$\simeq S_{w_s}^- A_{e_2} A_{e_1} A_{e_3} S_{w_s}^+ + UA_{e_2 * e_1} S_{w_s}^- S_{w_s}^+ \tag{Lemma 5.13}$$

$$\simeq S_{w_s}^- A_{e_2} A_{e_1} A_{e_3} S_{w_s}^+ \tag{Lemma 5.15}.$$

We now prove the statement when $n > 3$ by induction. Assume the lemma statement holds whenever w_s has valence $n - 1$. The idea is to replace $e_1 \cup e_2$ with a Y -shaped graph which has valence 1 at w_s . See Figure 15.

Let $w_0 \in Y$ be a point which is disjoint from all of the edges e_i . Pick a path e from w_s to w_0 , which is also disjoint from all e_i . We first prove the following:

$$(6-15) \quad A_{e_2} A_{e_1} \simeq S_{w_0}^- A_{e_2 * e} A_{e_1 * e} A_e S_{w_0}^+.$$

To establish (6-15), we compute

$$\begin{aligned} (6-16) \quad S_{w_0}^- A_{e_2 * e} A_{e_1 * e} A_e S_{w_0}^+ &\simeq S_{w_0}^- (A_{e_2} + A_e)(A_{e_1} + A_e) A_e S_{w_0}^+ \\ &\simeq S_{w_0}^- A_{e_2} A_{e_1} A_e S_{w_0}^+ + S_{w_0}^- (A_{e_2} A_e^2 + A_e A_{e_1} A_e + A_e^3) S_{w_0}^+ \\ &\simeq S_{w_0}^- A_{e_2} A_{e_1} A_e S_{w_0}^+ + S_{w_0}^- (A_{e_2} U + A_{e_1} U + UA_e + UA_e) S_{w_0}^+ \\ &\simeq A_{e_2} A_{e_1} S_{w_0}^- A_e S_{w_0}^+ + S_{w_0}^- (A_{e_2} U + A_{e_1} U + UA_e + UA_e) S_{w_0}^+ \\ &\simeq A_{e_2} A_{e_1} + S_{w_0}^- (A_{e_2} U + A_{e_1} U + UA_e + UA_e) S_{w_0}^+ \\ &\simeq A_{e_2} A_{e_1} + US_{w_0}^- A_{e_2 * e_1} S_{w_0}^+ \\ &\simeq A_{e_2} A_{e_1} + UA_{e_2 * e_1} S_{w_0}^- S_{w_0}^+ \\ &\simeq A_{e_2} A_{e_1}. \end{aligned}$$

Equation (6-16) is justified as follows. Line 1 is justified by additivity of the relative homology maps. Line 2 is obtained by algebra. Line 3 is justified by Lemmas 4.4 and 4.5. Line 4 is justified by Lemma 5.13. Line 5 is justified by Lemma 6.10. Line 6 follows by canceling the repeated UA_e terms, and using additivity of the relative homology maps. Line 7 follows from Lemma 5.13. The final line is justified by Lemma 6.10. Equation (6-15) is established.

Define new edges e'_1 and e'_2 as the concatenations

$$e'_1 := e_1 * e \quad \text{and} \quad e'_2 := e_2 * e.$$

See Figure 15.




Figure 15: The edges e_1 and e_2 , the vertices w_s and w_0 , and the new edges e , e'_1 and e'_2 .

Using (6-15), we see that

$$(6-17) \quad S_{w_s}^- A_{e_n} \cdots A_{e_2} A_{e_1} S_{w_s}^+ \simeq S_{w_s}^- A_{e_n} \cdots A_{e_3} (S_{w_0}^- A_{e_2}' A_{e_1}' A_e S_{w_0}^+) S_{w_s}^+.$$

Note that w_0, e_1' and e_2' are disjoint from e_2, \dots, e_n . Hence, Lemmas 4.4 and 5.13 and Proposition 5.14 imply that (6-17) is chain homotopic to

$$(6-18) \quad S_{w_0}^- A_{e_2}' A_{e_1}' (S_{w_s}^- A_{e_n} \cdots A_{e_3} A_e S_{w_s}^+) S_{w_0}^+.$$

By induction, we know that (6-18) is chain homotopic to

$$(6-19) \quad S_{w_0}^- A_{e_2}' A_{e_1}' (S_{w_s}^- A_{e_{n-1}} \cdots A_{e_3} A_e A_{e_n} S_{w_s}^+) S_{w_0}^+.$$

Commuting terms, using the same justification as above, we see that (6-19) is chain homotopic to

$$(6-20) \quad S_{w_s}^- A_{e_{n-1}} \cdots A_{e_3} (S_{w_0}^- A_{e_2}' A_{e_1}' A_e S_{w_0}^+) A_{e_n} S_{w_s}^+.$$

Applying (6-15), we can conclude that (6-20) is chain homotopic to

$$S_{w_s}^- A_{e_{n-1}} \cdots A_{e_3} A_{e_2} A_{e_1} A_{e_n} S_{w_s}^+. \quad \square$$

6.5 Subdivision-invariance and the trivial strand relation

In this section we prove two basic relations about the graph action map: subdivision-invariance and the trivial strand relation.

Lemma 6.14 *Suppose $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ and $\mathcal{G}' = (\Gamma', \mathbf{w}_0, \mathbf{w}_1)$ are two ribbon flow-graphs in Y such that Γ' is obtained from Γ by adding a vertex to the interior of an edge of Γ . Then*

$$\mathfrak{A}_{\mathcal{G}} \simeq \mathfrak{A}_{\mathcal{G}'} \quad \text{and} \quad \mathfrak{B}_{\mathcal{G}} \simeq \mathfrak{B}_{\mathcal{G}'}$$

Proof We focus on the maps $\mathfrak{A}_{\mathcal{G}}$ and $\mathfrak{A}_{\mathcal{G}'}$. The proof is essentially the same as the argument to show invariance under level splitting (move (CM1)). Using independence of the graph action map from the choice of Cerf decomposition (Theorem 6.1(a)), it is sufficient to show the claim when \mathcal{G} is an elementary flow-graph of type (ER1) (a translation flow-graph).

Suppose e is the edge of \mathcal{G} which we wish to subdivide, and let w_0 and w_1 be the incoming and outgoing vertices of e . Let w denote the vertex in the interior of e which is added to form \mathcal{G}' . Note that \mathcal{G} is of type (ER1), while \mathcal{G}' is of type (ER2), with valence 2 special vertex w . Let e_1 and e_2 denote the two components of $e \setminus \{w\}$.

Using Lemmas 4.4 and 5.13 and Proposition 5.14, it is sufficient to show the claim when $\mathcal{G} = (e, w_0, w_1)$, since the maps corresponding to other edges and vertices can be commuted past the maps for e .

By definition,

$$(6-21) \quad \mathfrak{A}_{\mathcal{G}} := S_{w_0}^- A_e S_{w_1}^+,$$




Figure 16: Adding a trivial strand. Left: an elementary flow-graph \mathcal{G} of type (ER1), with a single edge e . Right: a Cerf decomposition of a graph \mathcal{G}' obtained by adding a trivial strand to \mathcal{G} .

while

$$(6-22) \quad \mathfrak{A}_{\mathcal{G}'} := S_{w_0}^- S_w^- A_{e_2} A_{e_1} S_w^+ S_{w_1}^+.$$

Using Lemma 6.11, we conclude that (6-21) and (6-22) are chain homotopic. \square

Next, we prove the *trivial strand relation* for the graph action maps:

Lemma 6.15 *Suppose that $\mathcal{G} = (\Gamma, w_0, w_1)$ is a ribbon flow-graph in Y , and \mathcal{G}' is obtained by adjoining a new edge e_0 such that $e_0 \cap \Gamma$ consists of a single point which has valence 3 in $\Gamma \cup e_0$. (We allow the new valence 3 vertex to be given either cyclic order.) Then*

$$\mathfrak{A}_{\mathcal{G}} \simeq \mathfrak{A}_{\mathcal{G}'},$$

and similarly $\mathfrak{B}_{\mathcal{G}} \simeq \mathfrak{B}_{\mathcal{G}'}$.

Proof Using independence from the choice of Cerf decomposition (Theorem 6.1(a)), it is sufficient to show the claim when \mathcal{G} is an elementary flow-graph of type (ER1). As in the proof of Lemma 6.14, it is also sufficient to consider the case when \mathcal{G} has one connected component. Let us write e for the single edge of Γ , w for the incoming vertex of \mathcal{G} , and w' for the outgoing vertex. By definition,

$$(6-23) \quad \mathfrak{A}_{\mathcal{G}} = S_w^- A_e S_{w'}^+.$$

The graph \mathcal{G}' can be given a Cerf decomposition into two flow-graphs, $\mathcal{G}_2 \circ \mathcal{G}_1$, where \mathcal{G}_2 is an elementary flow-graph of type (ER2) with a valence 1 special vertex, and \mathcal{G}_1 is an elementary flow-graph of type (ER2) with a valence 3 special vertex. Let $e_1, e_2, e_3, e_4, e_5, w_1, w_2, w_3$ and w_4 denote the new vertices of \mathcal{G}' , as shown in Figure 16.

We assume that the edges incident to the new valence 3 vertex, w_1 , are ordered e_1, e_2 then e_3 . The argument for the opposite cyclic ordering is a simple modification, which we leave to the reader.

By definition,

$$\mathfrak{A}_{\mathcal{G}'} = S_{w_2 w_3 w_4}^- A_{e_4} A_{e_5} S_{w' w_4}^+ S_{w w_1}^- A_{e_3} A_{e_2} A_{e_1} S_{w_1 w_2 w_3}^+.$$

We manipulate the above expression for $\mathfrak{A}_{\mathcal{G}'}$, as follows:

$$\begin{aligned} S_{w_2 w_3 w_4}^- A_{e_4} A_{e_5} S_{w' w_4}^+ S_{w w_1}^- A_{e_3} A_{e_2} A_{e_1} S_{w_1 w_2 w_3}^+ \\ \simeq S_{w_2 w_3}^- A_{e_4} (S_{w_4}^- A_{e_5} S_{w_4}^+) S_{w'}^+ S_{w w_1}^- A_{e_3} A_{e_2} A_{e_1} S_{w_1 w_2 w_3}^+ \quad (\text{Lemma 5.13}) \end{aligned}$$

$$\begin{aligned}
 &\simeq S_{w_2 w_3}^- A_{e_4} S_{w'}^+ S_{w w_1}^- A_{e_3} A_{e_2} A_{e_1} S_{w_1 w_2 w_3}^+ && \text{(Lemma 6.10)} \\
 &\simeq S_{w w_1 w_3}^- (S_{w_2}^- A_{e_4} A_{e_3} S_{w_2}^+) A_{e_2} A_{e_1} S_{w' w_1 w_3}^+ && \text{(Lemma 5.13)} \\
 &\simeq S_{w w_1 w_3}^- A_{e_3 * e_4} A_{e_2} A_{e_1} S_{w' w_1 w_3}^+ && \text{(Lemma 6.11)} \\
 &\simeq S_{w w_1 w_3}^- A_{e_1} A_{e_3 * e_4} A_{e_2} S_{w' w_1 w_3}^+ && \text{(Lemma 6.13)} \\
 &\simeq S_{w w_1}^- A_{e_1} A_{e_3 * e_4} (S_{w_3}^- A_{e_2} S_{w_3}^+) S_{w' w_1}^+ && \text{(Lemma 5.13)} \\
 &\simeq S_{w w_1}^- A_{e_1} A_{e_3 * e_4} S_{w' w_1}^+ && \text{(Lemma 6.10)} \\
 &\simeq S_w^- A_{e_1 * e_3 * e_4} S_w^+ && \text{(Lemma 6.11),}
 \end{aligned}$$

which agrees with the expression for \mathfrak{A}_G in (6-23). □

7 1- and 3-handles

In this section, we define maps for 4-dimensional 1-handles and 3-handles, and prove invariance. Our construction is similar to the constructions of [Ozsváth and Szabó 2006; Juhász 2016], though there are some differences. Unlike the construction from [Ozsváth and Szabó 2006], our construction allows us to consider 1-handles which join two components of a 3-manifold, or 3-handles which separate a 3-manifold into two components. Unlike Juhász’s construction [2016], which works only for \widehat{CF} , our construction applies to CF^- , CF^+ and CF^∞ .

In Section 7.8, we show that our construction coincides with Ozsváth and Szabó’s for 1-handles with feet in the same component of the 3-manifold.

7.1 Definition of the 1-handle and 3-handle maps

Suppose that (Y, \mathbf{w}) is a multipointed 3-manifold and $\mathbb{S}^0 = \{p_1, p_2\}$ is a 0-sphere in $Y \setminus \mathbf{w}$. Pick a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ such that

$$\mathbb{S}^0 \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta}).$$

We construct a Heegaard surface $\widehat{\Sigma}$ for the surgered 3-manifold $Y(\mathbb{S}^0)$ by removing neighborhoods of p_1 and p_2 in Σ and gluing in an annulus contained in the 1-handle region of $Y(\mathbb{S}^0)$. In the annulus region, we add two new curves α_0 and β_0 which are homologically essential in the annulus and intersect transversely in a pair of points. The two points can be distinguished by the Maslov grading. Write θ^+ and θ^- for the two points of $\alpha_0 \cap \beta_0$, and $\widehat{\mathcal{H}}$ for the diagram $(\widehat{\Sigma}, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w})$.

There is a unique Spin^c structure $\widehat{\mathfrak{s}} \in \text{Spin}^c(Y(\mathbb{S}^0))$ which restricts to \mathfrak{s} on $Y \setminus N(\mathbb{S}^0)$ and evaluates trivially on the new 2-sphere in $Y(\mathbb{S}^0)$. There is a unique 4-dimensional Spin^c structure \mathfrak{t} on the 1-handle cobordism $W(Y, \mathbb{S}^0)$ which extends \mathfrak{s} . The Spin^c structure \mathfrak{t} restricts to $\widehat{\mathfrak{s}}$ on $Y(\mathbb{S}^0)$.

It is an easy exercise to see that if \mathcal{H} is strongly \mathfrak{s} -admissible then $\widehat{\mathcal{H}}$ is strongly $\widehat{\mathfrak{s}}$ -admissible.

We define the 1-handle map

$$F_{Y, \mathbb{S}^0, \mathfrak{t}}: CF^-(\mathcal{H}, \sigma, \mathfrak{s}) \rightarrow CF^-(\widehat{\mathcal{H}}, \sigma, \widehat{\mathfrak{s}})$$

using the formula

$$(7-1) \quad F_{Y, \mathbb{S}^0, t}(\mathbf{x}) = \mathbf{x} \times \theta^+,$$

extended equivariantly over $\mathbb{F}_2[U_{\mathbf{w}}]$. Like the free-stabilization maps, the 1-handle and 3-handle maps require the almost complex structure to be stretched. See Definition 7.2 for a precise definition of which almost complex structures can be chosen.

We now describe the 3-handle maps. Suppose that $\mathbb{S}^2 \subseteq Y \setminus \mathbf{w}$ is an embedded 2-sphere, and $\widehat{\mathcal{H}} = (\widehat{\Sigma}, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w})$ is a Heegaard diagram for (Y, \mathbf{w}) such that $\mathbb{S}^2 \cap \widehat{\Sigma}$ consists of a circle c which is disjoint from $\boldsymbol{\alpha} \cup \boldsymbol{\beta}$. Furthermore, assume that an annular neighborhood of c contains both α_0 and β_0 (which are homologically essential in this annulus) and $\alpha_0 \cap \beta_0 = \{\theta^+, \theta^-\}$. A diagram for $Y(\mathbb{S}^2)$ may be obtained by cutting out a neighborhood of c , removing α_0 and β_0 , and filling in the two boundary components with disks. Write \mathcal{H} for the resulting diagram.

For a sufficiently stretched almost complex structure, the 3-handle map

$$F_{Y, \mathbb{S}^2, t}: \text{CF}^-(\widehat{\mathcal{H}}, \sigma, \hat{\mathfrak{s}}) \rightarrow \text{CF}^-(\mathcal{H}, \sigma, \mathfrak{s})$$

is defined via the formulas

$$(7-2) \quad F_{Y, \mathbb{S}^2, t}(\mathbf{x} \times \theta^+) = 0 \quad \text{and} \quad F_{Y, \mathbb{S}^2, t}(\mathbf{x} \times \theta^-) = \mathbf{x},$$

extended $\mathbb{F}_2[U_{\mathbf{w}}]$ -equivariantly.

We note that the 1-handle maps and the 3-handle maps are algebraically dual, and hence any statement about the 1-handle maps has a corresponding statement about the 3-handle maps. To streamline the exposition, we focus mostly on the 1-handle maps.

7.2 Gluing data for 1- and 3-handles

We now describe precisely which almost complex structures can be used to compute the 1-handle maps. The technical details are similar to the free-stabilization maps, so we will be terse.

It is convenient to view the diagram $\widehat{\mathcal{H}}$ for $Y(\mathbb{S}^0)$, constructed in Section 7.1, as being obtained by connecting a diagram (S^2, α_0, β_0) to the diagram \mathcal{H} using two tubes attached to antipodal regions of $S^2 \setminus (\alpha_0 \cup \beta_0)$. Let p_1^0 and p_2^0 denote these two connected sum points on S^2 , and fix two disks $D_1^0, D_2^0 \subseteq S^2 \setminus (\alpha_0 \cup \beta_0)$, containing p_1^0 and p_2^0 , respectively. Also fix regular neighborhoods $N(p_1)$ and $N(p_2)$ in Y , to construct the surgered manifold $Y(\mathbb{S}^0)$.

Similar to Definition 5.1, we make the following definition:

Definition 7.1 Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for (Y, \mathbf{w}) and $p_1, p_2 \in \Sigma \setminus (\mathbf{w} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta})$. We call a tuple $\mathfrak{d} = (J^{\mathfrak{d}}, J_0^{\mathfrak{d}}, D_1, D_2, t)$ a *gluing datum for a 1-handle attached at p_1 and p_2* if the following hold:

- (1) $J_0^{\mathfrak{d}}$ is an almost complex structure on $S^2 \times [0, 1] \times \mathbb{R}$ which is split on D_1^0 and D_2^0 .
- (2) $D_i \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w})$ is a closed disk such that $\frac{1}{2} \cdot D_i$ contains $N(p_i) \cap \Sigma$ for $i \in \{1, 2\}$.

- (3) J^∂ is an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$ which is split on D_1 and D_2 .
- (4) $\iota: S^2 \setminus (\frac{1}{2} \cdot D_1^0 \cup \frac{1}{2} \cdot D_2^0) \rightarrow Y(S^0)$ is a smooth embedding such that the following hold:
 - (a) $\text{im}(\iota) \cap (Y \setminus (N(p_1) \cup N(p_2))) \subseteq \Sigma$.
 - (b) ι maps each annulus $D_i^0 \setminus \frac{1}{2} \cdot D_i^0$ conformally onto $D_i \setminus \frac{1}{2} \cdot D_i$.

If $T = (T_1, T_2)$ is a pair of positive real numbers, by adapting the construction for free-stabilized almost complex structures from Section 5.1, we can construct an almost complex structure $J^\partial(T)$ on $\widehat{\Sigma} \times [0, 1] \times \mathbb{R}$ with necks of length T_1 and T_2 .

Analogously to Definition 5.2, we make the following definition:

Definition 7.2 Suppose \mathcal{H} is a Heegaard diagram for (Y, \mathbf{w}) and ∂ is a gluing datum for attaching a 1-handle with feet $p_1, p_2 \in Y \setminus \mathbf{w}$. We say a pair of neck lengths T satisfies *stabilizing condition* (SC3) if for any two pairs of neck lengths T_1 and T_2 such that $T_1, T_2 \geq T$, componentwise, there is a noncylindrical almost complex structure \tilde{J} on $\Sigma \times [0, 1] \times \mathbb{R}$, interpolating $J^\partial(T_1)$ and $J^\partial(T_2)$, such that, for all $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we have

$$(SC3) \quad \Psi_{\tilde{J}}(\mathbf{x} \times \theta^+) = \mathbf{x} \times \theta^+ \quad \text{and} \quad \Psi_{\tilde{J}}(\mathbf{x} \times \theta^-) = \mathbf{x} \times \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x}, \mathbf{y}} \cdot \mathbf{y} \times \theta^+$$

for $C_{\mathbf{x}, \mathbf{y}} \in \mathbb{F}_2[U_{\mathbf{w}}]$ (which depend on ∂, T_1 and T_2).

Analogous to Proposition 5.3, we have the following:

Proposition 7.3 *If ∂ is a gluing datum for attaching a 1-handle at $\{p_1, p_2\}$, then there is a pair of neck lengths $T = (T_1, T_2)$ which satisfies stabilizing condition (SC3).*

We begin with the following Maslov index formula:

Lemma 7.4 *Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) and $\widehat{\mathcal{H}} = (\widehat{\Sigma}, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \mathbf{w})$ is a diagram for the surgered manifold $Y(S^0)$. If $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \mathbf{x}, \mathbf{y} \times \mathbf{y})$ is a homology class of disks on $\widehat{\mathcal{H}}$, where $\mathbf{x}, \mathbf{y} \in \{\theta^+, \theta^-\}$, then*

$$\mu(\phi \# \phi_0) = \mu(\phi) + \text{gr}(\mathbf{x}, \mathbf{y}).$$

Proof By Proposition 5.3, the index of ϕ_0 is

$$(7-3) \quad \mu(\phi_0) = 2n_{p_1^0}(\phi_0) + 2n_{p_2^0}(\phi_0) + \text{gr}(\mathbf{x}, \mathbf{y}).$$

Noting that the Euler measure of a disk is 1, Lipshitz’s formula [2006, equation 8] for the Maslov index implies that

$$(7-4) \quad \mu(\phi \# \phi_0) = \mu(\phi) + \mu(\phi_0) - 2n_{p_1^0}(\phi_0) - 2n_{p_2^0}(\phi_0).$$

Combining (7-3) and (7-4) implies the main statement. □

With the Maslov index formula from Lemma 7.4, Proposition 7.3 is proven by adapting the proof of Proposition 5.3 to handle stretching two necks instead of one. The main details of the argument are unchanged, so we leave them to the reader.

7.3 1-handles, 3-handles and the differential

Ozsváth and Szabó [2004a, Proposition 6.4] proved that their 1-handle and 3-handle maps are chain maps. We now prove that our version of the 1-handle and 3-handle maps are also chain maps.

Proposition 7.5 *Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) , and $\widehat{\mathcal{H}} = (\widehat{\Sigma}, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \mathbf{w})$ is the diagram constructed by attaching a 1-handle at $\{p_1, p_2\} \subseteq \Sigma \setminus (\alpha \cup \beta \cup \mathbf{w})$. If \mathfrak{d} is a gluing datum for this 1-handle, and \mathbf{T} is a pair of neck lengths satisfying stabilizing condition (SC3), then*

$$(7-5) \quad \partial_{\widehat{\mathcal{H}}, J^\circ(\mathbf{T})}(\mathbf{x} \times \theta^+) = \partial_{\mathcal{H}, J^\circ(\mathbf{T})}(\mathbf{x}) \otimes \theta^+, \quad \partial_{\widehat{\mathcal{H}}, J^\circ(\mathbf{T})}(\mathbf{x} \otimes \theta^-) = \partial_{\mathcal{H}, J^\circ(\mathbf{T})}(\mathbf{x}) \otimes \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x}, \mathbf{y}} \cdot \mathbf{y} \times \theta^+$$

for $C_{\mathbf{x}, \mathbf{y}} \in \mathbb{F}_2[U_{\mathbf{w}}]$.

Proof The proof is essentially the same as the proof of Proposition 5.5.

Suppose that $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \mathbf{x}, \mathbf{y} \times \mathbf{y})$ is a class with Maslov index 1. By Lemma 7.4,

$$(7-6) \quad \mu(\phi \# \phi_0) = \mu(\phi) + \text{gr}(\mathbf{x}, \mathbf{y}).$$

Note that (7-5) makes no claim about the counts of classes with $\text{gr}(\mathbf{x}, \mathbf{y}) = -1$ (i.e., where $\mathbf{x} = \theta^-$ and $\mathbf{y} = \theta^+$). These counts correspond to the $C_{\mathbf{x}, \mathbf{y}}$ in the statement.

As in the proof of Proposition 5.5, if $\mathcal{M}_{J^\circ(\mathbf{T})}(\phi \# \phi_0)$ is nonempty for arbitrarily large \mathbf{T} , then both ϕ and ϕ_0 must have broken representatives.

If $\text{gr}(\mathbf{x}, \mathbf{y}) = 1$, then (7-6) implies that $\mu(\phi) = 0$. Since ϕ has a broken representative, ϕ is the constant class $e_{\mathbf{x}} \in \pi_2(\mathbf{x}, \mathbf{x})$. In this case, $\phi \# \phi_0$ must have domain equal to one of the two bigons in the 1-handle region. These curves cancel, modulo 2, and hence make no contribution to the differential.

It remains to consider classes with $\text{gr}(\mathbf{x}, \mathbf{y}) = 0$, i.e., $\mathbf{x} = \mathbf{y} = \theta^+$ or $\mathbf{x} = \mathbf{y} = \theta^-$. For such classes, equation (7-6) implies that

$$\mu(\phi \# \phi_0) = \mu(\phi) = 1.$$

Write n_1 and n_2 for $n_{p_1}(\phi)$ and $n_{p_2}(\phi)$, respectively. Consider the map

$$\rho^{p_1, p_2}: \mathcal{M}_{J^\circ}(\phi) \rightarrow \text{Sym}^{n_1}([0, 1] \times \mathbb{R}) \times \text{Sym}^{n_2}([0, 1] \times \mathbb{R}).$$

Write $X(\phi)$ for the image $\rho^{p_1, p_2}(\mathcal{M}(\phi))$.

As in the proof of Proposition 5.5, for large \mathbf{T} , there is a fibered product description

$$(7-7) \quad \#\widehat{\mathcal{M}}_{J^\circ(\mathbf{T})}(\phi \# \phi_0) \equiv \sum_{u \in \widehat{\mathcal{M}}(\phi)} \#\mathcal{M}_{J^\circ}(\phi_0, \rho^{p_1, p_2}(u)).$$

Consequently, it is sufficient to show that, if $\theta \in \{\theta^+, \theta^-\}$ is fixed and

$$\mathbf{d}_1 \times \mathbf{d}_2 \in \text{Sym}^{n_1}([0, 1] \times \mathbb{R}) \times \text{Sym}^{n_2}([0, 1] \times \mathbb{R})$$

is a point with no repeated entries, then, for a generic almost complex structure J_0 on $S^2 \times [0, 1] \times \mathbb{R}$,

$$(7-8) \quad \sum_{\substack{\phi_0 \in \pi_2(\theta, \theta) \\ n_{p_1^0}(\phi_0) = n_1 \\ n_{p_2^0}(\phi_0) = n_2}} \# \mathcal{M}_{J_0}(\phi_0, \mathbf{d}_1 \times \mathbf{d}_2) \equiv 1 \pmod{2}.$$

If $\mathbf{d}_1 \times \mathbf{d}_2 \in \text{Sym}^{n_1}([0, 1] \times \mathbb{R}) \times \text{Sym}^{n_2}([0, 1] \times \mathbb{R})$ has no entries with the same $[0, 1]$ -component, consider the path $\mathbf{D}_T := \mathbf{d}_1^T \times \mathbf{d}_2$ obtained by translating \mathbf{d}_1 upwards by T units in the \mathbb{R} -direction. If $\mathbf{d}_1 \times \mathbf{d}_2$ is not in the fat diagonal, but two elements share the same $[0, 1]$ -component, a perturbation of this path may be chosen which avoids the fat diagonal. The compactification of the 1-dimensional space

$$\bigcup_{T \in [0, \infty)} \prod_{\phi_0 \in \pi_2(\theta, \theta)} \mathcal{M}_{J_0}(\phi_0, \mathbf{D}_T)$$

has ends in bijection with the Cartesian product

$$(7-9) \quad \left(\prod_{\substack{\phi_0 \in \pi_2(\theta, \theta) \\ n_{p_1^0}(\phi_0) = n_1 \\ n_{p_2^0}(\phi_0) = 0}} \mathcal{M}(\phi_0, \mathbf{d}_1) \right) \times \left(\prod_{\substack{\phi_0 \in \pi_2(\theta, \theta) \\ n_{p_1^0}(\phi_0) = 0 \\ n_{p_2^0}(\phi_0) = n_2}} \mathcal{M}(\phi_0, \mathbf{d}_2) \right).$$

Equation (5-15) implies that the count of the elements in (7-9) is 1, modulo 2. Equation (7-8) follows.

Combining (7-7) and (7-8), it follows that

$$\sum_{\substack{\phi_0 \in \pi_2(\theta, \theta) \\ n_{p_1^0}(\phi_0) = n_1 \\ n_{p_2^0}(\phi_0) = n_2}} \# \widehat{\mathcal{M}}(\phi \# \phi_0) \equiv \# \widehat{\mathcal{M}}(\phi) \pmod{2}.$$

The main claim now follows. □

Corollary 7.6 *The 1-handle maps $F_{Y, \mathbb{S}^0, t}$ and the 3-handle maps $F_{Y, \mathbb{S}^2, t}$ are chain maps.*

Proof The result is an immediate consequence of Proposition 7.5, together with the formulas for the 1-handle and 3-handle maps in (7-1) and (7-2). □

7.4 1-handles, 3-handles and triangle maps

Similar to Theorem 5.7, the 1-handle and 3-handle maps satisfy a useful relationship with the triangle maps.

Suppose $\mathcal{T} = (\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a multipointed Heegaard triple, and $p_1, p_2 \in \Sigma \setminus (\alpha \cup \beta \cup \gamma \cup \mathbf{w})$ are two points. We can form a new Heegaard triple $\widehat{\mathcal{T}} = (\widehat{\Sigma}, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}, \mathbf{w})$ by cutting out a neighborhood of the points p_1 and p_2 , gluing in an annulus to connect the new boundary components. We add three new attaching curves, α_0, β_0 and γ_0 , in the new annular region. We assume that α_0, β_0 and γ_0 satisfy the configuration shown in Figure 17.




Figure 17: Adding a 1-handle to a Heegaard triple. The dashed circles are the boundaries of the new annular region.

Lemma 7.7 *If $\mathcal{T} = (\Sigma, \alpha, \beta, \gamma, \mathbf{w})$ is a Heegaard triple, and $\widehat{\mathcal{T}} = (\widehat{\Sigma}, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}, \mathbf{w})$ is obtained by attaching a 1-handle, as above, then there is a canonical isomorphism*

$$(7-10) \quad \text{Spin}^c(X_{\alpha, \beta, \gamma}) \cong \text{Spin}^c(X_{\alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}}).$$

Proof Write Σ_0 for the surface $\Sigma \setminus (N(p_1) \cup N(p_2))$, and write $X_{\alpha, \beta, \gamma}^0$ for the 4-manifold obtained as the union

$$(\Delta \times \Sigma_0) \cup (e_\alpha \times U_\alpha^0) \cup (e_\beta \times U_\beta^0) \cup (e_\gamma \times U_\gamma^0),$$

where U_α^0 is the 3-manifold with boundary and corners obtained by gluing 2-handles to $[0, 1] \times \Sigma_0$ along $\alpha \times \{1\}$, and U_β^0 and U_γ^0 are defined similarly. There are two restriction maps

$$(7-11) \quad \text{Spin}^c(X_{\alpha, \beta, \gamma}) \rightarrow \text{Spin}^c(X_{\alpha, \beta, \gamma}^0) \quad \text{and} \quad \text{Spin}^c(X_{\alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}}) \rightarrow \text{Spin}^c(X_{\alpha, \beta, \gamma}^0).$$

We leave it as a straightforward exercise for the reader to use the Mayer–Vietoris long exact sequences on cohomology to verify that both maps in (7-11) are isomorphisms, leading to the isomorphism in (7-10). \square

Theorem 7.8 *Suppose that \mathcal{T} is a Heegaard triple, and $\widehat{\mathcal{T}}$ is obtained by attaching a 1-handle, as in Figure 17. Let \mathfrak{d} be a gluing datum for the 1-handle attachment. If T is a pair of neck lengths which are sufficiently large, then, with respect to the isomorphism of Spin^c structures in (7-10), we have*

$$\begin{aligned} F_{\widehat{\mathcal{T}}, J^{\mathfrak{d}}(\mathbf{T}), s}(\mathbf{x} \times x^+, \mathbf{y} \times y^+) &= F_{\mathcal{T}, J^{\mathfrak{d}}, s}(\mathbf{x}, \mathbf{y}) \otimes z^+, \\ F_{\widehat{\mathcal{T}}, J^{\mathfrak{d}}(\mathbf{T}), s}(\mathbf{x} \times x^+, \mathbf{y} \times y^-) &= F_{\mathcal{T}, J^{\mathfrak{d}}, s}(\mathbf{x}, \mathbf{y}) \otimes z^- + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} C_{\mathbf{x}, \mathbf{y}, z}^1 \cdot z \times z^+, \\ F_{\widehat{\mathcal{T}}, J^{\mathfrak{d}}(\mathbf{T}), s}(\mathbf{x} \times x^-, \mathbf{y} \times y^+) &= F_{\mathcal{T}, J^{\mathfrak{d}}, s}(\mathbf{x}, \mathbf{y}) \otimes z^- + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} C_{\mathbf{x}, \mathbf{y}, z}^2 \cdot z \times z^+ \end{aligned}$$

for $C_{\mathbf{x}, \mathbf{y}, z}^1, C_{\mathbf{x}, \mathbf{y}, z}^2 \in \mathbb{F}_2[U_{\mathbf{w}}]$ (which depend on $\mathcal{T}, \mathfrak{d}$ and T).

Proof The proof follows the same line of reasoning as the proof of Theorem 5.7.

Suppose that $\psi \# \psi_0 \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y, z \times z)$ is a homology class with Maslov index 0. Using Lemma 5.8, Sarkar’s formula [2011] for the Maslov index, as well as the fact that the Euler measure of a

disk is 1, we compute that

$$(7-12) \quad \mu(\psi \# \psi_0) = \mu(\psi) - \text{gr}(x^+, x) - \text{gr}(y^+, y) + \text{gr}(z^+, z).$$

Define

$$\delta(x, y, z) := -\text{gr}(x^+, x) - \text{gr}(y^+, y) + \text{gr}(z^+, z).$$

It is straightforward to check that the theorem statement follows from the following two subclaims:

- (h1) If $\delta(x, y, z) = 1$, then any Maslov index 0 class in $\pi_2(\mathbf{x} \times x, \mathbf{y} \times y, \mathbf{z} \times z)$ has no $J^\delta(\mathbf{T})$ -holomorphic representatives when \mathbf{T} is sufficiently large.
- (h2) If $\delta(x, y, z) = 0$, then the $\mathbf{z} \times z$ coefficient of $F_{\mathcal{T}, J^\delta(\mathbf{T}), s}(\mathbf{x} \times x, \mathbf{y} \times y)$ coincides with the \mathbf{z} coefficient of $F_{\widehat{\mathcal{T}}, J^\delta, s}(\mathbf{x}, \mathbf{y})$.

The counts of the moduli spaces with $\delta(x, y, z) \in \{-1, -2\}$ are not relevant to the theorem statement.

Claim (h1) is the easier of the two subclaims to verify, since it relies only on compactness and transversality, but does not require gluing. By (7-12), if $\delta(x, y, z) = 1$, then $\mu(\psi) = -1$. If $\psi \# \psi_0$ has representatives for arbitrarily large \mathbf{T} , then ψ must admit a broken representative. However, since $\mu(\psi) = -1$, there can be no broken representatives, by transversality. Hence, no index 0 classes with $\delta(x, y, z) = 1$ have nonempty moduli spaces for sufficiently large \mathbf{T} .

We now consider claim (h2). Assume $\delta(x, y, z) = 0$. In this case, we have $\mu(\psi) = 0$ by (7-12). As in the proof of Theorem 5.7, if \mathbf{T}_i is an increasing, unbounded sequence of pairs of neck lengths, then any sequence of $J^\delta(\mathbf{T}_i)$ -holomorphic representatives of $\psi \# \psi_0$ has a subsequence which converges to a pair (u, u_0) where $u \in \mathcal{M}_{J^\delta}(\psi)$ and $u_0 \in \mathcal{M}_{J^\delta}(\psi_0)$, and

$$(7-13) \quad \rho^{p_1, p_2}(u) = \rho^{p_1^0, p_2^0}(u_0).$$

In (7-13), $\rho^{p_1, p_2}(u) \in \text{Sym}^{n_{p_1}(\psi)}(\Delta) \times \text{Sym}^{n_{p_2}(\psi)}(\Delta)$ is the set

$$\rho^{p_1, p_2}(u) = ((\pi_\Delta \circ u)((\pi_\Sigma \circ u)^{-1}(p_1)), (\pi_\Delta \circ u)((\pi_\Sigma \circ u)^{-1}(p_2))),$$

and $\rho^{p_1^0, p_2^0}(u_0)$ is defined similarly.

Hence, using a gluing argument, for sufficiently large \mathbf{T} there is an identification

$$\mathcal{M}_{J^\delta(\mathbf{T})}(\psi \# \psi_0) \cong \mathcal{M}_{J^\delta}(\psi) \times_\rho \mathcal{M}_{J^\delta}(\psi_0).$$

Since $\mathcal{M}_{J^\delta}(\psi)$ is 0-dimensional, it is sufficient to show that if $\mathbf{d}_1 \times \mathbf{d}_2 \in \text{Sym}^{n_1}(\Delta) \times \text{Sym}^{n_2}(\Delta)$ is not in the fat diagonal, then

$$(7-14) \quad \sum_{\substack{\psi_0 \in \pi_2(x, y, z) \\ n_{p_1^0}(\psi_0) = n_1 \\ n_{p_2^0}(\psi_0) = n_2}} \#\mathcal{M}(\psi_0, \mathbf{d}_1 \times \mathbf{d}_2) \equiv 1$$

for a generic almost complex structure on $S^2 \times [0, 1] \times \mathbb{R}$.

Equation (7-14) is verified similarly to (5-23). Consider a path $(\mathbf{D}_t)_{t \in [1, \infty)}$ in $\text{Sym}^{n_1}(\Delta) \times \text{Sym}^{n_2}(\Delta)$, satisfying the following:

- (1) $\mathbf{D}_1 = \mathbf{d}_1 \times \mathbf{d}_2$.
- (2) The image of \mathbf{D}_t is disjoint from the fat diagonal.
- (3) All points of \mathbf{D}_t travel into the α - β cylindrical end of Δ .
- (4) For large t , the points of \mathbf{D}_t are spaced at least distance t apart (with respect to the identification of the α - β cylindrical end of Δ as $[0, 1] \times (-\infty, 0]$).
- (5) The $[0, 1]$ -components of all points in \mathbf{D}_t approach a fixed $s_0 \in (0, 1)$.

Write $\mathcal{D} := \{\mathbf{D}_t : t \in [1, \infty)\}$ and consider the matched moduli space

$$\mathcal{M}_{(x,y,z)}(\mathcal{D}) := \coprod_{\substack{\psi_0 \in \pi_2(x,y,z) \\ n_{p_1^0}(\psi_0) = n_1 \\ n_{p_2^0}(\psi_0) = n_2}} \bigcup_{t \in [1, \infty)} \mathcal{M}(\psi_0, \mathbf{D}_t).$$

As in the proof of (5-23), the space $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ has ends at $t = 1$, at $t \in (1, \infty)$ and at $t = \infty$. The ends at $t = 1$ correspond to the left-hand side of (7-14). The ends at $t \in (1, \infty)$ correspond to index 1 strips breaking off at finite t , which do not pass over p_1^0 or p_2^0 . The ends at $t = \infty$ correspond to the Cartesian product

$$(7-15) \quad \left(\coprod_{\substack{\psi_0^0 \in \pi_2(x,y,z) \\ n_{p_1^0}(\psi_0^0) = n_1 \\ n_{p_2^0}(\psi_0^0) = 0}} \mathcal{M}(\psi_0^0) \right) \times \left(\coprod_{\substack{\phi \in \pi_2(x,x) \\ n_{p_1^0}(\phi) = n_1 \\ n_{p_2^0}(\phi) = 0}} \mathcal{M}(\phi, d) \right)^{n_1} \times \left(\coprod_{\substack{\phi \in \pi_2(x,x) \\ n_{p_1^0}(\phi) = 0 \\ n_{p_2^0}(\phi) = n_2}} \mathcal{M}(\phi, d) \right)^{n_2},$$

where $d \in \{s_0\} \times \mathbb{R} \subseteq [0, 1] \times \mathbb{R}$ is a chosen point.

As in the proof of Theorem 5.7, the ends of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$ corresponding to strip breaking at finite t cancel modulo 2, since x is a cycle in the complex $\widehat{\text{CF}}(S^2, \alpha_0, \beta_0, p_1^0, p_2^0)$, and similarly y and z are cycles in their appropriate complexes.

The first factor of (7-15) has total count 1, modulo 2, since it is easy to check that when $\delta(x, y, z) = 0$, the only nonnegative triangle class which has zero multiplicity over p_1^0 and p_2^0 is a small triangle, which clearly has a unique representative. The latter two factors of (7-15) have total count 1, modulo 2, by [Ozsváth and Szabó 2008, Lemma 6.4].

Hence, by counting the ends of $\mathcal{M}_{(x,y,z)}(\mathcal{D})$, equation (7-14) follows. □

7.5 1-handles, 3-handles and gluing data

Analogous to Proposition 5.9, we have the following:

Proposition 7.9 *Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is diagram for (Y, \mathbf{w}) , $\mathbb{S}^0 = \{p_1, p_2\} \subseteq \Sigma \setminus (\alpha \cup \beta \cup \mathbf{w})$ is a 0-sphere, and ∂_1 and ∂_2 are two gluing data for attaching a 1-handle at \mathbb{S}^0 . Suppose further that T_1*

and T_2 are pairs of neck lengths which satisfy stabilizing condition (SC3) for ∂_1 and ∂_2 , respectively. Write $\widehat{\mathcal{H}}_1$ and $\widehat{\mathcal{H}}_2$ for the diagrams constructed by attaching a 1-handle to \mathcal{H} using ∂_1 and ∂_2 (note that $\widehat{\mathcal{H}}_1$ and $\widehat{\mathcal{H}}_2$ differ only by an isotopy in the 1-handle region). The following diagram commutes up to chain homotopy:

$$(7-16) \quad \begin{array}{ccc} \text{CF}_{J^{\partial_1}}^-(\mathcal{H}, \mathfrak{s}) & \xrightarrow{\Psi_{J^{\partial_1} \rightarrow J^{\partial_2}}} & \text{CF}_{J^{\partial_2}}^-(\mathcal{H}, \mathfrak{s}) \\ \downarrow F_{Y, \mathbb{S}^0, t} & & \downarrow F_{Y, \mathbb{S}^0, t} \\ \text{CF}_{J^{\partial_1}(T_1)}^-(\widehat{\mathcal{H}}_1, \widehat{\mathfrak{s}}) & \xrightarrow{\Psi_{(\widehat{\mathcal{H}}_1, J^{\partial_1}(T_1)) \rightarrow (\widehat{\mathcal{H}}_2, J^{\partial_2}(T_2))}} & \text{CF}_{J^{\partial_2}(T_2)}^-(\widehat{\mathcal{H}}_2, \widehat{\mathfrak{s}}) \end{array}$$

An analogous relation holds for the 3-handle maps.

Proposition 7.9 is proven by adapting our proof of the analogous result for the free-stabilization maps in Proposition 5.9.

7.6 Invariance of the 1- and 3-handle maps

In this section, we prove that the 1-handle and 3-handle maps are independent of the choices used in the construction.

Theorem 7.10 *The 1-handle and 3-handle maps determine well-defined chain maps of transitive systems of chain complexes.*

Proof We focus on the 1-handle maps, since the 3-handle maps are dual.

Suppose $\mathbb{S}^0 = \{p_1, p_2\}$ is an embedded 0-sphere in Y , $t \in \text{Spin}^c(W(Y, \mathbb{S}^0))$, and $\mathcal{H}_1 = (\Sigma_1, \alpha_1, \beta_1, \mathbf{w})$ and $\mathcal{H}_2 = (\Sigma_2, \alpha_2, \beta_2, \mathbf{w})$ are two Heegaard diagrams for (Y, \mathbf{w}) such that $\{p_1, p_2\} \subseteq \Sigma_i \setminus (\alpha_i \cup \beta_i \cup \mathbf{w})$ for $i \in \{1, 2\}$. Let ∂_1 and ∂_2 be two gluing data. It is sufficient to show that if T_1 and T_2 are two pairs of neck lengths satisfying condition (SC3), then the following diagram commutes up to chain homotopy:

$$(7-17) \quad \begin{array}{ccc} \text{CF}_{J^{\partial_1}}^-(\mathcal{H}_1, \mathfrak{s}) & \xrightarrow{\Psi_{(\mathcal{H}_1, J^{\partial_1}) \rightarrow (\mathcal{H}_2, J^{\partial_2})}} & \text{CF}_{J^{\partial_2}}^-(\mathcal{H}_2, \mathfrak{s}) \\ \downarrow F_{Y, \mathbb{S}, t} & & \downarrow F_{Y, \mathbb{S}, t} \\ \text{CF}_{J^{\partial_1}(T_1)}^-(\widehat{\mathcal{H}}_1, \widehat{\mathfrak{s}}) & \xrightarrow{\Psi_{(\widehat{\mathcal{H}}_1, J^{\partial_1}(T_1)) \rightarrow (\widehat{\mathcal{H}}_2, J^{\partial_2}(T_2))}} & \text{CF}_{J^{\partial_2}(T_2)}^-(\widehat{\mathcal{H}}_2, \widehat{\mathfrak{s}}) \end{array}$$

By Proposition 7.9, the 1-handle maps are independent of the choice of gluing data (i.e., diagram (7-17) commutes when $\mathcal{H}_1 = \mathcal{H}_2$).

By Lemma 5.12, we can connect \mathcal{H}_1 and \mathcal{H}_2 by a sequence of the following moves:

- (1) handleslides and isotopies of the α and β curve (possibly passing over p_1 and p_2);
- (2) simple stabilizations, away from $\mathbf{w} \cup \{p_1, p_2\}$;
- (3) changing the embedding of the Heegaard surface by an isotopy $\phi_t: \Sigma \rightarrow Y$ which is fixed on $\mathbf{w} \cup \{p_1, p_2\}$ for all t .

Since the transition maps for handleslides and isotopies can be computed using a composition of triangle maps, Theorem 7.8 implies invariance under move (1). Invariance under simple stabilizations away from $\mathbf{w} \cup \{p_1, p_2\}$ follows from a triple neck-stretching argument, similar to Lemma 5.10. Finally, invariance under move (3) is tautological. \square

7.7 Further properties of the 1- and 3-handle maps

In this section we prove several additional results about the 1-handle and 3-handle maps.

Lemma 7.11 *Suppose that λ is a path in Y , \mathbb{S} is an embedded 0-sphere or 2-sphere in $Y \setminus \mathbf{w}$ and $\mathfrak{t} \in \text{Spin}^c(W(Y, \mathbb{S}))$. Suppose that $\lambda \subseteq Y \setminus \mathbb{S}$ is a closed loop, or a path connecting two basepoints. Write λ also for the induced path in $Y(\mathbb{S})$. Then*

$$A_\lambda \circ F_{Y, \mathbb{S}, \mathfrak{t}} \simeq F_{Y, \mathbb{S}, \mathfrak{t}} \circ A_\lambda.$$

Proof The proof is similar to the proof of Lemma 5.13, and follows by examining the curves counted by the differential in Proposition 7.5, using the fact that the path λ does not enter the 1-handle region. We leave it to the reader to make the necessary notational modifications to the proof of Lemma 5.13. \square

Next, we consider commuting 1-handle and 3-handle maps amongst each other. If $\mathbb{S} \subseteq Y$ is a union of framed k -spheres, we write $W(Y, \mathbb{S})$ for the cobordism from Y to $Y(\mathbb{S})$ obtained by attaching $k + 1$ handles to $[0, 1] \times Y$ along $\{1\} \times \mathbb{S}$. Note that if \mathbb{S} and \mathbb{S}' are two disjoint, framed spheres in Y , then the 4-manifolds $W(Y, \mathbb{S} \cup \mathbb{S}')$ and $W(Y(\mathbb{S}), \mathbb{S}') \cup W(Y, \mathbb{S})$ and $W(Y(\mathbb{S}'), \mathbb{S}) \cup W(Y, \mathbb{S}')$ are diffeomorphic, via diffeomorphisms which are well defined up to isotopy.

Lemma 7.12 *Suppose that \mathbb{S} and \mathbb{S}' are two disjoint embedded spheres of dimension 0 or 2 in Y . Write \mathfrak{t} for a Spin^c structure on*

$$W(Y, \mathbb{S} \cup \mathbb{S}') \cong W(Y(\mathbb{S}), \mathbb{S}') \cup W(Y, \mathbb{S}) \cong W(Y(\mathbb{S}'), \mathbb{S}) \cup W(Y, \mathbb{S}').$$

Then

$$F_{Y(\mathbb{S}), \mathbb{S}', \mathfrak{t}|_{W(Y(\mathbb{S}), \mathbb{S}')}} \circ F_{Y, \mathbb{S}, \mathfrak{t}|_{W(Y, \mathbb{S})}} \simeq F_{Y(\mathbb{S}'), \mathbb{S}, \mathfrak{t}|_{W(Y(\mathbb{S}'), \mathbb{S})}} \circ F_{Y, \mathbb{S}', \mathfrak{t}|_{W(Y, \mathbb{S}')}}.$$

Proof The proof follows from a quadruple neck-stretching argument, similar to the double neck-stretching argument from Proposition 5.14 used to show that free-stabilization maps commute with each other. The present statement follows from the following subclaim:

Subclaim 7.12.1 *Suppose that \mathbf{T} and \mathbf{T}' are two 4-tuples of neck lengths for attaching two 1-handles. If all of the components of \mathbf{T} and \mathbf{T}' are sufficiently large, then a noncylindrical almost complex structure \tilde{J} interpolating $J(\mathbf{T})$ and $J(\mathbf{T}')$ may be chosen so that*

$$\begin{aligned} \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^+ \times \theta_2^+) &= \mathbf{x} \times \theta_1^+ \times \theta_2^+, \\ \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^+ \times \theta_2^-) &= \mathbf{x} \times \theta_1^+ \times \theta_2^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} (C_{\mathbf{x}, \mathbf{y}}^1 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^+ + C_{\mathbf{x}, \mathbf{y}}^2 \cdot \mathbf{y} \times \theta_1^- \times \theta_2^+), \end{aligned}$$

$$\Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^- \times \theta_2^+) = \mathbf{x} \times \theta_1^- \times \theta_2^+ + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} (C_{\mathbf{x},\mathbf{y}}^3 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^+ + C_{\mathbf{x},\mathbf{y}}^4 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^-),$$

$$\begin{aligned} \Psi_{\tilde{J}}(\mathbf{x} \times \theta_1^- \times \theta_2^-) &= \mathbf{x} \times \theta_1^- \times \theta_2^- \\ &+ \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} (C_{\mathbf{x},\mathbf{y}}^5 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^- + C_{\mathbf{x},\mathbf{y}}^6 \cdot \mathbf{y} \times \theta_1^- \times \theta_2^+ + C_{\mathbf{x},\mathbf{y}}^7 \cdot \mathbf{y} \times \theta_1^+ \times \theta_2^+) \end{aligned}$$

for various $C_{\mathbf{x},\mathbf{y}}^i \in \mathbb{F}_2[U_{\mathbf{w}}]$, which depend on \tilde{J} .

The proof of Subclaim 7.12.1 follows the same reasoning as the proof of Subclaim 5.14.1, the analogous subclaim of Proposition 5.14. The main difference is that we must use the index formula from Lemma 7.4 to compute the index of a homology class after adding two 1-handles. We leave it to the reader to verify that the argument carries over to our present context without major change. \square

Lemma 7.13 *Suppose that \mathbb{S} is an embedded 0- or 2-sphere in a multipointed 3-manifold (Y, \mathbf{w}) . If $w \in Y \setminus (\mathbf{w} \cup \mathbb{S})$ is a new basepoint, then*

$$S_w^+ \circ F_{Y,\mathbb{S},t} \simeq F_{Y,\mathbb{S},t} \circ S_w^+ \quad \text{and} \quad S_w^- \circ F_{Y,\mathbb{S},t} \simeq F_{Y,\mathbb{S},t} \circ S_w^-.$$

Proof The proof follows from a triple neck-stretching argument, similar to the ones we encountered in Proposition 5.14 and Lemma 7.12. We leave the necessary modifications to the reader. \square

7.8 Ozsváth and Szabó’s 1-handle and 3-handle maps

Ozsváth and Szabó originally defined the 1-handle and 3-handle maps by taking the connected sum of the Heegaard surface with a genus 1 Heegaard diagram for $S^1 \times S^2$ using the same formula as in (7-1) and (7-2). Morally, this amounts to picking a path between the two components of the attaching 0-sphere of a 1-handle. For showing invariance of the graph cobordism maps, it is convenient to show that our definition coincides with Ozsváth and Szabó’s original definition [2006, Section 4.3] when the feet of the 1-handle are in the same component of the 3-manifold. This amounts to the following change of almost complex structure computation:

Lemma 7.14 *Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ denote a Heegaard diagram and let $\mathcal{H}' = (\Sigma \# \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w})$ denote the diagram obtained by connect summing the diagram $(\mathbb{T}^2, \alpha_0, \beta_0)$ for $S^1 \times S^2$, as shown in Figure 18. Let c, c_1 and c_2 denote the three circles labeled therein. If $\mathbf{T} = (T, T_1, T_2)$ is a triple of positive real numbers, let $J(\mathbf{T})$ denote an almost complex structure which has been stretched along c, c_1 and c_2 , with neck lengths T, T_1 and T_2 . If all components of \mathbf{T} and \mathbf{T}' are sufficiently large, then there exists a noncylindrical almost complex structure \tilde{J} on $\Sigma \# \mathbb{T}^2 \times [0, 1] \times \mathbb{R}$, interpolating $J(\mathbf{T})$ and $J(\mathbf{T}')$, such that*

$$(7-18) \quad \Psi_{\tilde{J}}(\mathbf{x} \times \theta^+) = \mathbf{x} \times \theta^+ \quad \text{and} \quad \Psi_{\tilde{J}}(\mathbf{x} \times \theta^-) = \mathbf{x} \times \theta^- + \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} C_{\mathbf{x},\mathbf{y}} \cdot \mathbf{y} \times \theta^+$$

for some $C_{\mathbf{x},\mathbf{y}} \in \mathbb{F}_2[U_{\mathbf{w}}]$ (which depend on \tilde{J}).




Figure 18: The diagram $\mathcal{H}' = (\Sigma \# \mathbb{T}^2, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \mathbf{w})$ and the circles c, c_1 and c_2 in Lemma 7.14.

Proof First, a modification of Lemma 5.4 implies that if $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is a class of disks on $(\Sigma, \alpha, \beta, \mathbf{w})$ and $\phi_0 \in \pi_2(x, y)$ is a class of disks on $(\mathbb{T}^2, \alpha_0, \beta_0)$, then

$$(7-19) \quad \mu(\phi \# \phi_0) = \mu(\phi) + \text{gr}(x, y).$$

Note that the present lemma statement concerns only classes with $\text{gr}(x, y) \geq 0$.

Suppose that T_i and T'_i are two sequences of neck lengths, all of whose components approach $+\infty$. We choose noncylindrical almost complex structures \tilde{J}_i , interpolating $J(T_i)$ and $J(T'_i)$, such that $(\Sigma \# \mathbb{T}^2 \times [0, 1] \times \mathbb{R}, \tilde{J}_i)$ contains the almost complex submanifold $((\Sigma \setminus N_i) \times [0, 1] \times \mathbb{R}, J)$, where N_i is some nested sequence of open neighborhoods of the connected sum point $p \in \Sigma$, such that $\bigcap_{i \in \mathbb{N}} N_i = \{p\}$, and J is a fixed cylindrical almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$. Fix a Maslov index 0 class $\phi \# \phi_0$. As in the proof of Proposition 5.3, from a sequence u_i of \tilde{J}_i -holomorphic representatives of $\phi \# \phi_0$, we may extract a broken representative of ϕ .

If $\text{gr}(x, y) = 1$ and $\phi \# \phi_0$ is a Maslov index 0 class with \tilde{J}_i -holomorphic representatives for large i , then (7-19) implies that $\mu(\phi) = -1$. However, since ϕ admits a broken J -holomorphic representative and J is cylindrical, we must have $\mu(\phi) \geq 0$, so we obtain a contradiction. Hence, such classes $\phi \# \phi_0$ have no representatives for large i .

If $\text{gr}(x, y) = 0$ and $\phi \# \phi_0$ is a Maslov index 0 class with \tilde{J}_i -holomorphic representatives for large i , then (7-19) implies $\mu(\phi) = 0$. Since ϕ admits broken holomorphic representatives for a generic cylindrical almost complex structure, we conclude that ϕ is the constant class. There are no nonnegative Maslov index 0 classes on $(\mathbb{T}^2, \alpha_0, \beta_0)$ with zero multiplicity over the connected sum point. Hence, ϕ_0 must also be a constant class, e_θ for $\theta \in \{\theta^+, \theta^-\}$. On the other hand, the class $e_{\mathbf{x} \times \theta}$ always has \tilde{J}_i -holomorphic representatives. Equation (7-18) follows. \square

8 2-handles

In this section, we describe the cobordism maps for 2-handles. The maps we describe are essentially the same as those defined by Ozsváth and Szabó [2006]. They are also similar to the versions defined by Juhász [2016] in the setting of cobordisms of sutured manifolds.

8.1 Definition of the 2-handle maps

Definition 8.1 A framed link \mathbb{S}^1 in a 3-manifold Y is a collection of pairwise disjoint, embedded knots $K_1, \dots, K_m \subseteq Y$ together with a choice framing, i.e., a choice of homology classes $\ell_i \in H_1(\partial N(K_i); \mathbb{Z})$ satisfying $\mu_i \cdot \ell_i = 1$ for a meridian μ_i of K_i .

Adapting Ozsváth and Szabó’s definition [2006, Definition 4.1] in the singly pointed setting, we make the following definition:

Definition 8.2 Suppose (Y, \mathbf{w}) is a multipointed 3-manifold and \mathbb{S}^1 is a framed link in Y , with components K_1, \dots, K_m . A bouquet B of \mathbb{S}^1 is an embedded forest in Y (i.e., a collection of embedded, contractible graphs) such that each leaf of B is a point on K_i , or a point in \mathbf{w} . Furthermore, we assume the following:

- (1) For each $i \in \{1, \dots, m\}$, the set $B \cap K_i$ contains a single point.
- (2) Each connected component of B intersects exactly one point in \mathbf{w} , and furthermore that point is a leaf of B .

If B is a bouquet of the link \mathbb{S}^1 in (Y, \mathbf{w}) , then we form a manifold with boundary $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$, by removing a regular neighborhood of $B \cup \mathbb{S}^1$. We decorate the boundary of $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$ with a collection of sutures (i.e., a collection of oriented, simple closed curves which divide the boundary into two subsurfaces) by adding one contractible suture for each basepoint in \mathbf{w} .

Definition 8.3 Suppose that \mathbb{S}^1 is a framed link in Y , with components K_1, \dots, K_m . We say that a Heegaard triple

$$(\Sigma, \alpha, \beta, \beta', \mathbf{w}) = (\Sigma, \{\alpha_1, \dots, \alpha_n\}, \{\beta_1, \dots, \beta_n\}, \{\beta'_1, \dots, \beta'_n\}, \mathbf{w})$$

is subordinate to a bouquet B for the framed link \mathbb{S}^1 in Y if the following hold:

- (1) $(\Sigma, \alpha, \beta, \mathbf{w})$ is a Heegaard diagram for (Y, \mathbf{w}) .
- (2) If Σ_0 denotes the surface obtained by removing small neighborhoods of the basepoints \mathbf{w} , then

$$(\Sigma_0, \{\alpha_1, \dots, \alpha_n\}, \{\beta_{m+1}, \dots, \beta_n\})$$

is a sutured Heegaard diagram for the sutured manifold $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$.

- (3) The curves $\beta'_{m+1}, \dots, \beta'_n$ are small isotopies of the curves $\beta_{m+1}, \dots, \beta_n$. Furthermore,

$$|\beta_i \cap \beta'_j| = 2\delta_{ij}$$

whenever $m + 1 \leq i, j \leq n$.

- (4) If $1 \leq i \leq m$, then β_i is a meridian of K_i .
- (5) If $1 \leq i \leq m$, then β'_i is a longitude of K_i , corresponding to the framing.
- (6) If $1 \leq i \leq m$, then β'_i is disjoint from the curves $\beta_{m+1}, \dots, \beta_n$. Furthermore,

$$|\beta'_i \cap \beta_j| = \delta_{ij}$$

whenever $1 \leq i, j \leq m$.

Since $Y_{\beta, \beta'}$ is a connected sum of $g(\Sigma) - m$ copies of $S^1 \times S^2$, a theorem of Laudenbach and Poénaru [1972] implies that the diffeomorphism type of the 4-manifold obtained by attaching 3- and 4-handles to $Y_{\beta, \beta'} \subseteq \partial X_{\alpha, \beta, \beta'}$ is unique. Similar to [Ozsváth and Szabó 2006, Proposition 4.3], we have the following simple description of the resulting 4-manifold:

Lemma 8.4 *Suppose $(\Sigma, \alpha, \beta, \beta', \mathbf{w})$ is subordinate to a bouquet for a framed link \mathbb{S}^1 in Y . After filling in the boundary component $Y_{\beta, \beta'} \subseteq \partial X_{\alpha, \beta, \beta'}$ with 3- and 4-handles, we obtain the handle cobordism $W(Y, \mathbb{S}^1)$.*

Proof Add an extra product layer and view $W(Y, \mathbb{S}^1)$ as the union

$$(8-1) \quad W(Y, \mathbb{S}^1) = ([1, 2] \times Y(\mathbb{S}^1)) \cup H_2 \cup ([0, 1] \times Y),$$

where H_2 is a union of copies of $D^2 \times D^2$ (the 2-handles). Write

$$Y = U_\alpha \cup ([0, 1] \times \Sigma) \cup U_\beta \quad \text{and} \quad Y(\mathbb{S}) = U_\alpha \cup ([0, 1] \times \Sigma) \cup U_{\beta'},$$

and delete

$$W_0 := ([1, 1 + \epsilon] \times U_{\beta'}) \cup H$$

from $W(Y, \mathbb{S}^1)$. We note that W_0 is a 4-dimensional handlebody. Using the description of $W(Y, \mathbb{S}^1)$ from (8-1), we can write

$$(8-2) \quad W \setminus \text{int}(W_0) = (U_\alpha \times [0, 2]) \cup (U_\beta \times [0, 1]) \cup (U_{\beta'} \times [1 + \epsilon, 2]) \cup ([0, 2] \times [0, 1] \times \Sigma).$$

Upon rounding corners and identifying $[0, 2] \times [0, 1]$ topologically with a triangle, this is identical to $X_{\alpha, \beta, \beta'}$, as defined in (3-23). □

We now define the cobordism maps for 2-handle cobordisms. Suppose $\mathbb{S}^1 \subseteq Y$ is a framed link in Y , and B is a bouquet. Let $(\Sigma, \alpha, \beta, \beta', \mathbf{w})$ be a Heegaard triple subordinate to B . There is a unique top-graded intersection point $\Theta_{\beta, \beta'}^+ \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}$. Furthermore, it is straightforward to see that $\Theta_{\beta, \beta'}^+$ is a cycle in the complex $\text{CF}^-(\Sigma, \beta, \beta', \mathbf{w}, \mathfrak{s}_0)$, where $\mathfrak{s}_0 \in \text{Spin}^c(Y_{\beta, \beta'})$ is the unique torsion Spin^c structure.

By Lemma 8.4, there is a canonical isomorphism

$$\text{Spin}^c(X_{\alpha, \beta, \beta'}) \cong \text{Spin}^c(W(Y, \mathbb{S}^1)).$$

Therefore, we will not distinguish between Spin^c structures on $X_{\alpha, \beta, \beta'}$ and $W(Y, \mathbb{S}^1)$.

If $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}^1))$, the 2-handle map

$$F_{Y, \mathbb{S}^1, \mathfrak{s}}: \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s}|_Y) \rightarrow \text{CF}^-(\Sigma, \beta, \beta', \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S}^1)})$$

is defined as the holomorphic triangle map

$$(8-3) \quad F_{Y, \mathbb{S}^1, \mathfrak{s}}(\mathbf{x}) := F_{\alpha, \beta, \beta', \mathfrak{s}}(\mathbf{x} \otimes \Theta_{\beta, \beta'}^+) = \sum_{\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}} \sum_{\substack{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta, \beta'}^+, \mathbf{z}) \\ \mu(\psi) = 0 \\ \mathfrak{s}_w(\psi) = \mathfrak{s}}} \# \mathcal{M}(\psi) U_{\mathbf{w}}^{n_w(\psi)} \cdot \mathbf{z},$$

extended $\mathbb{F}_2[U_{\mathbf{w}}]$ -equivariantly.




Figure 19: The triple $(\mathbb{T}^2, \delta_0, \epsilon_0, \epsilon'_0, p_0)$ used in a simple stabilization of a Heegaard triple. We view p_0 as being the outer dashed circle, collapsed to a point.

8.2 Simple stabilizations and triangle maps

In this section, we prove that the holomorphic triangle maps are invariant under simple stabilizations. For the original argument in the context of the symmetric product, we refer the reader to [Ozsváth and Szabó 2004b, Theorem 10.4; 2006, Theorem 2.14].

Definition 8.5 Let $(\mathbb{T}^2, \delta_0, \epsilon_0, \epsilon'_0, p_0)$ denote the genus 1 Heegaard triple shown in Figure 19. If $(\Sigma, \alpha, \beta, \gamma, w)$ is a Heegaard triple, we say the Heegaard triple

$$(\Sigma', \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}, w)$$

is a *simple stabilization* of $(\Sigma, \alpha, \beta, \gamma, w)$ if $\Sigma' = \Sigma \# \mathbb{T}^2$ and $\{\alpha_0, \beta_0, \gamma_0\} = \{\delta_0, \epsilon_0, \epsilon'_0\}$ setwise.

Lemma 8.6 If $(\Sigma, \alpha, \beta, \gamma, w)$ is a Heegaard triple and $(\Sigma \# \mathbb{T}_0, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}, w)$ is a simple stabilization, then there is a canonical isomorphism

$$(8-4) \quad \text{Spin}^c(X_{\alpha, \beta, \gamma}) \cong \text{Spin}^c(X_{\alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \gamma \cup \{\gamma_0\}}).$$

The proof of Lemma 8.6 is a Mayer–Vietoris argument similar to Lemma 7.7, which we leave to the reader.

Theorem 8.7 Suppose $\mathcal{T} = (\Sigma, \alpha, \beta, \gamma, w)$ is a Heegaard triple, $(\mathbb{T}^2, \delta_0, \epsilon_0, \epsilon'_0, p_0)$ is the Heegaard triple shown in Figure 19, and \mathfrak{d} is a gluing datum for stretching the neck. Let a, b, θ^+ and θ^- denote the intersection points shown in Figure 19. For sufficiently large T , we have

$$(8-5) \quad \begin{aligned} F_{\alpha \cup \{\epsilon_0\}, \beta \cup \{\epsilon'_0\}, \gamma \cup \{\delta_0\}, J^\mathfrak{d}(T), s}(\mathbf{x} \times \theta^+, \mathbf{y} \times a) &= F_{\alpha, \beta, \gamma, J^\mathfrak{d}, s}(\mathbf{x}, \mathbf{y}) \otimes b, \\ F_{\alpha \cup \{\delta_0\}, \beta \cup \{\epsilon_0\}, \gamma \cup \{\epsilon'_0\}, J^\mathfrak{d}(T), s}(\mathbf{x} \times b, \mathbf{y} \times \theta^+) &= F_{\alpha, \beta, \gamma, J^\mathfrak{d}, s}(\mathbf{x}, \mathbf{y}) \otimes a, \\ F_{\alpha \cup \{\epsilon_0\}, \beta \cup \{\delta_0\}, \gamma \cup \{\epsilon'_0\}, J^\mathfrak{d}(T), s}(\mathbf{x} \times b, \mathbf{y} \times a) &= F_{\alpha, \beta, \gamma, J^\mathfrak{d}, s}(\mathbf{x}, \mathbf{y}) \otimes \theta^- + \sum_{z \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} C_{\mathbf{x}, \mathbf{y}, z} \cdot z \times \theta^+ \end{aligned}$$

for some $C_{\mathbf{x}, \mathbf{y}, z} \in \mathbb{F}_2[U_w]$, which depend on \mathfrak{d} and T .

Remark 8.8 Theorem 8.7 can be restated as follows. Let σ_b denote the simple stabilization map defined on intersection points as $\sigma_b(\mathbf{x}) = \mathbf{x} \times b$, and extended equivariantly over $\mathbb{F}_2[U_w]$. Define σ_a similarly.

Let $F_1^{\epsilon_0, \epsilon'_0}$ denote the 1-handle map, defined as $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$, extended equivariantly over $\mathbb{F}_2[U_{\mathbf{w}}]$. Define the 3-handle map $F_3^{\epsilon_0, \epsilon'_0}$ via the formulas $\mathbf{x} \times \theta^+ \mapsto 0$ and $\mathbf{x} \times \theta^- \mapsto \mathbf{x}$. Equation (8-5) can be restated as

$$\begin{aligned} F_{\alpha \cup \{\epsilon_0\}, \beta \cup \{\epsilon'_0\}, \gamma \cup \{\delta_0\}, J^\circ(T), s}(F_1^{\epsilon_0, \epsilon'_0}(\mathbf{x}), \sigma_a(\mathbf{y})) &= \sigma_b F_{\alpha, \beta, \gamma, J^\circ, s}(\mathbf{x}, \mathbf{y}), \\ F_{\alpha \cup \{\delta_0\}, \beta \cup \{\epsilon_0\}, \gamma \cup \{\epsilon'_0\}, J^\circ(T), s}(\sigma_b(\mathbf{x}), F_1^{\epsilon_0, \epsilon'_0}(\mathbf{y})) &= \sigma_a F_{\alpha, \beta, \gamma, J^\circ, s}(\mathbf{x}, \mathbf{y}), \\ F_{\alpha \cup \{\epsilon_0\}, \beta \cup \{\delta_0\}, \gamma \cup \{\epsilon'_0\}, J^\circ(T), s}(\sigma_b(\mathbf{x}), \sigma_a(\mathbf{y})) &= F_1^{\epsilon_0, \epsilon'_0} F_{\alpha, \beta, \gamma, J^\circ, s}(\mathbf{x}, \mathbf{y}) \otimes \theta^-. \end{aligned}$$

Proof of Theorem 8.7 The proof we present is similar to the proofs of Theorems 5.7 and 7.8. We focus on the first formula in (8-5), since the second and third are proven similarly.

A Maslov index computation similar to Lemma 5.4 shows that if $\psi_0 \in \pi_2(\theta^+, a, b)$, then

$$\mu(\psi_0) = 2n_{p_0}(\psi_0) + \text{gr}(\theta^+, \theta).$$

Consequently, if $\psi \# \psi_0 \in \pi_2(\mathbf{x} \times \theta, \mathbf{y} \times a, \mathbf{z} \times b)$, then

$$(8-6) \quad \mu(\psi \# \psi_0) = \mu(\psi) + \text{gr}(\theta^+, \theta).$$

Next, to obtain transversality at curves which appear in our proof, we consider almost complex structures on $\Delta \times \Sigma$ satisfying (J'1), (J'2), (J'3'), (J'4') and (J'5'); see Proposition 3.4 for a precise statement of transversality. These are the same almost complex structures considered in the proof of handleswap-invariance [Juhász et al. 2021, Section 9.3]; we refer the reader there for a detailed account of a similar argument.

As in the proofs of Theorems 5.7 and 7.8, a sequence of $J^\circ(T_i)$ -holomorphic representatives of $\psi \# \psi_0$ will degenerate into a pair of broken holomorphic triangles \mathcal{U} and \mathcal{U}_0 , representing ψ and ψ_0 , respectively. Using (8-6), as well as a transversality, we conclude that $\mu(\psi) = 0$. Consequently, arguing as in Theorems 5.7 and 7.8, the broken triangles \mathcal{U} and \mathcal{U}_0 each consist of a single curve, u and u_0 , respectively, which satisfy the matching condition

$$\rho^p(u) = \rho^{p_0}(u_0),$$

where p and p_0 are the connected sum points. Via a gluing argument, as in the proofs of Theorems 5.7 and 7.8, it is sufficient to show that for a generic $\mathbf{d} \in \text{Sym}^n(\Delta)$, the following count holds for the matched moduli spaces on $(\mathbb{T}^2, \epsilon_0, \epsilon'_0, \delta_0, p_0)$:

$$(8-7) \quad \sum_{\substack{\psi_0 \in \pi_2(\theta^+, a, b) \\ n_{p_0}(\psi_0) = n}} \#\mathcal{M}(\psi_0, \mathbf{d}) \equiv 1 \pmod{2}.$$

The argument to establish (8-7) is formally similar to the argument for establishing (5-15), (5-23), (7-8) and (7-14).

Pick a path $(\mathbf{d}_t)_{t \in [0, \infty)}$ in $\text{Sym}^n(\Delta)$ satisfying the following:

- (1) $\mathbf{d}_0 = \mathbf{d}$.
- (2) \mathbf{d}_t is disjoint from the fat diagonal.

- (3) The points in \mathbf{d}_t all enter into the $\epsilon_0\text{-}\delta_0$ cylindrical end of $\Delta \times \Sigma$.
- (4) For large t , the points in \mathbf{d}_t are spaced at least distance t apart with respect to the Euclidean metric under the identification of the cylindrical ends as $[0, 1] \times [0, \infty)$.
- (5) The points of \mathbf{d}_t all approach a line $\{s_0\} \times [0, \infty)$ as $t \rightarrow \infty$.

Write $\mathcal{D} := \{\mathbf{d}_t : t \in [0, \infty)\}$, and consider the 1-dimensional moduli space

$$\mathcal{M}_{(\theta^+, a, b)}(\mathcal{D}) := \coprod_{\substack{\psi_0 \in \pi_2(\theta^+, a, b) \\ n_{p_0}(\psi) = n}} \bigcup_{t \in [0, \infty)} \mathcal{M}_{J^\circ(T)}(\psi_0, \mathbf{d}_t).$$

We count the ends of $\mathcal{M}(\mathcal{D})$. As in the proof of (5-23), generically, the ends correspond to the following:

- (e'1) $t = 0$;
- (e'2) index 1 holomorphic strips breaking off at $t \in (0, \infty)$, with zero multiplicity at p_0 ;
- (e'3) $t \rightarrow \infty$.

The count of the ends of the form (e'1) is equal to the left-hand side of (8-7).

The ends of the form (e'2) cancel modulo 2, since θ^+ is a cycle in $\widehat{\text{CF}}(\mathbb{T}^2, \epsilon_0, \epsilon'_0, p_0)$, and a and b are cycles in their appropriate complexes.

The ends of the form (e'3) correspond to the Cartesian product

$$(8-8) \quad \left(\coprod_{\substack{\psi_0^0 \in \pi_2(\theta^+, a, b) \\ n_{p_0}(\psi_0^0) = 0}} \mathcal{M}(\psi_0^0) \right) \times \left(\coprod_{\substack{\phi \in \pi_2(b, b) \\ n_{p_0}(\phi) = n}} \mathcal{M}(\phi, d) \right),$$

where d is a point on the line $\{s_0\} \times \mathbb{R}$.

We wish to show that the total count in (8-8) is 1. The total count of the left factor is 1, since there is only one class which contributes, and that class is a small triangle. It remains to show that if ϕ denotes the index 2 class in $\pi_2(b, b)$ with $n_{p_0}(\phi) = 1$, then

$$(8-9) \quad \#\mathcal{M}(\phi, d) \equiv 1 \pmod{2}.$$

Note that by axiom (J'3'), $\mathcal{M}(\phi, d)$ consists of holomorphic curves for an almost complex structure on $\mathbb{T}^2 \times [0, 1] \times \mathbb{R}$ which satisfies (J1)–(J4) and (J5'). These are precisely the almost complex structures Lipshitz used to prove stabilization-invariance of the Heegaard Floer complexes. With respect to these almost complex structures, Lipshitz [2006, Sublemma A.12] proves (8-9) while proving stabilization-invariance. We repeat Lipshitz's argument here, for the convenience of the reader.

Consider the genus 2 diagram for $S^1 \times S^2$ in Figure 20. Let $\phi_1 \# \phi_0$ denote the index 1 class shown. The class ϕ_1 is a bigon, and ϕ_0 denotes the index 2 class on \mathbb{T}^2 . By a neck-stretching argument (analogous to the one above), for large neck length,

$$\#\widehat{\mathcal{M}}(\phi_1 \# \phi_0) = \#\widehat{\mathcal{M}}(\phi_1) \cdot \#\widehat{\mathcal{M}}(\phi_0, d).$$




Figure 20: A once-stabilized diagram of $S^1 \times S^2$ used in the proof of Theorem 8.7. The shaded class is $\phi_1 \# \phi_0$. We stretch along the dashed line.

On the other hand, by invariance of Heegaard Floer homology, since the diagram represents $S^1 \times S^2$ and there is only one other nonnegative index 1 class (a bigon), we must have $\#\widehat{\mathcal{M}}(\phi_1 \# \phi_0) = 1$. Equation (8-9) follows.

The other two stated formulas from (8-5) follow mutatis mutandis. □

8.3 Invariance of the 2-handle maps

We now prove invariance of the 2-handle maps. Since the 2-handle maps which feature in the graph TQFT are essentially identical to those defined by Ozsváth and Szabó, our exposition will be terse.

Extending [Ozsváth and Szabó 2006, Lemma 4.5] to multipointed 3-manifolds, we have the following:

Lemma 8.9 *Suppose \mathbb{S}^1 is an m -component framed link in (Y, \mathbf{w}) , and $B \subseteq Y$ is a fixed bouquet of \mathbb{S}^1 . Then there exists a Heegaard triple $(\Sigma, \alpha, \beta, \beta', \mathbf{w})$ subordinate to B . Furthermore, two Heegaard triples subordinate to \mathbb{S}^1 can be connected by a sequence of the following moves:*

- (1) *an isotopy or handleslide amongst the α curves;*
- (2) *an isotopy or handleslide amongst the curves $\beta_{m+1}, \dots, \beta_n$, followed by the corresponding move applied to $\beta'_{m+1}, \dots, \beta'_n$;*
- (3) *a simple stabilization or destabilization (of the first or second types in Theorem 8.7), in the complement of $B \cup \mathbb{S}^1$;*
- (4) *for $i \in \{1, \dots, m\}$, an isotopy of β_i , or a handleslide of β_i across one of the $\beta_{m+1}, \dots, \beta_n$;*
- (5) *for $i \in \{1, \dots, m\}$, an isotopy of β'_i , or a handleslide of β'_i across one of the $\beta'_{m+1}, \dots, \beta'_n$;*
- (6) *an isotopy Σ_t of the Heegaard surface Σ , inside of Y , fixing \mathbf{w} , such that Σ_t intersects $B \cup \mathbb{S}^1$ only along \mathbf{w} for all t .*

Proof To construct a Heegaard triple subordinate to B , pick a Heegaard diagram for the sutured manifold $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$ (formed by removing a neighborhood of $B \cup \mathbb{S}^1$ and adding a contractible suture for each basepoint in \mathbf{w}). Let $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_{m+1}, \dots, \beta_n)$ denote this diagram. We then define $\beta'_{m+1}, \dots, \beta'_n$ to be small isotopies of $\beta_{m+1}, \dots, \beta_n$.

A simple closed curve in $\partial(Y_{B \cup \mathbb{S}^1, \mathbf{w}})$ which avoids the contractible regions bounded by the sutures can be projected onto Σ to yield a curve which is in the complement of the curves $\beta_{m+1}, \dots, \beta_n$. Note that the curve on Σ obtained by projecting is only well defined up to isotopies and handleslides across the $\beta_{m+1}, \dots, \beta_n$ curves. Let $\beta_1, \dots, \beta_n \subseteq \Sigma$ be obtained by projecting meridians of the components of \mathbb{S}^1 onto Σ . Let $\beta'_1, \dots, \beta'_n$ denote projections of the longitudes. These curves are well defined up to moves (4) and (5).

Hence, given a sutured Heegaard diagram for $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$ we can obtain a Heegaard triple subordinate to $B \cup \mathbb{S}^1$ by the above procedure. Furthermore, any two triples constructed from a given Heegaard diagram via the above procedure can be connected by the listed moves. It remains to connect two sutured Heegaard diagrams for $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$. Any two sutured Heegaard diagrams for $Y_{B \cup \mathbb{S}^1, \mathbf{w}}$ can be connected by a sequence of Heegaard moves by [Juhász 2006, Proposition 2.15], which induce moves (1), (2), (3) and (6) on the resulting Heegaard triple. \square

We now show that the maps defined in (8-3) are independent of the choice of bouquet or subordinate triple:

Lemma 8.10 *Suppose \mathbb{S}^1 is a framed link in (Y, \mathbf{w}) , B_1 and B_2 are two bouquets for \mathbb{S}^1 , and \mathcal{T}_1 and \mathcal{T}_2 are Heegaard triples subordinate to B_1 and B_2 , respectively. Write $F_{Y, \mathbb{S}^1, s, \mathcal{T}_1}$ and $F_{Y, \mathbb{S}^1, s, \mathcal{T}_2}$ for the 2-handle maps computed with \mathcal{T}_1 and \mathcal{T}_2 . Write \mathcal{H}_1 and \mathcal{H}_2 for the diagrams of Y induced by \mathcal{T}_1 and \mathcal{T}_2 , respectively, and write \mathcal{H}'_1 and \mathcal{H}'_2 for the induced diagrams of $Y(\mathbb{S}^1)$. The following diagram commutes up to homotopy:*

$$\begin{CD} \text{CF}^-(\mathcal{H}_1, \mathbf{w}, \mathfrak{s}|_Y) @>\Psi_{\mathcal{H}_1 \rightarrow \mathcal{H}_2}>> \text{CF}^-(\mathcal{H}_2, \mathbf{w}, \mathfrak{s}|_Y) \\ @V F_{Y, \mathbb{S}^1, s, \mathcal{T}_1} VV @VV F_{Y, \mathbb{S}^1, s, \mathcal{T}_2} V \\ \text{CF}^-(\mathcal{H}'_1, \mathbf{w}, \mathfrak{s}|_{Y'}) @>\Psi_{\mathcal{H}'_1 \rightarrow \mathcal{H}'_2}>> \text{CF}^-(\mathcal{H}'_2, \mathbf{w}, \mathfrak{s}|_{Y'}) \end{CD}$$

Proof Our proof is no different than the original proof given by Ozsváth and Szabó [2006, Theorem 4.4], so we will be terse. See also [Juhász 2016, Theorem 6.9].

First, we fix a bouquet B for \mathbb{S}^1 and show that $F_{Y, \mathbb{S}^1, s, \mathcal{T}}$ is independent of the triple subordinate to B . This amounts to proving independence from the moves in Lemma 8.9. Invariance under moves (1), (2), (4) and (5) all follow from a relatively straightforward argument using associativity of the triangle maps. Invariance under move (3) (simple stabilizations and destabilizations) follows from Theorem 8.7. Finally, invariance under move (6) (isotopies of the Heegaard surface) is tautological.

Next, one needs to prove independence from the bouquet. The idea is that if $m + 1 \leq i \leq n$ and $1 \leq j \leq m$, then we can handleslide β_i twice across β_j , and handleslide β'_i twice across β'_j to change the bouquet, one 1-cell at a time. However, the 3-manifold represented by $(\Sigma, \beta, \beta', \mathbf{w})$ is unchanged by a sequence of such handleslides, and hence the corresponding transition map will preserve the top-degree generator. Consequently, a simple associativity argument shows that the cobordism map $F_{Y, \mathbb{S}^1, s, \mathcal{T}}$ is unchanged. \square

8.4 The composition law for 2-handle maps

We now state the Spin^c composition law for the 2-handle maps. We omit the proof, as it is identical to the proof given by Ozsváth and Szabó [2006, Proposition 4.9] for 2-handle cobordisms between singly pointed 3-manifolds using associativity of the holomorphic triangle maps [Ozsváth and Szabó 2004b, Theorem 8.16]:

Lemma 8.11 *Suppose that \mathbb{S}_1 and \mathbb{S}_2 are two disjoint, framed, 1-dimensional links in (Y, \mathbf{w}) , and $\mathfrak{s}_1 \in \text{Spin}^c(W(Y, \mathbb{S}_1))$ and $\mathfrak{s}_2 \in \text{Spin}^c(W(Y(\mathbb{S}_1), \mathbb{S}_2))$. Then*

$$F_{Y(\mathbb{S}_1), \mathbb{S}_2, \mathfrak{s}_2} \circ F_{Y, \mathbb{S}_1, \mathfrak{s}_1} = \sum_{\substack{\mathfrak{s} \in \text{Spin}^c(W(Y(\mathbb{S}_1), \mathbb{S}_2)) \\ \mathfrak{s}|_{W(Y, \mathbb{S}_1)} = \mathfrak{s}_1 \\ \mathfrak{s}|_{W(Y(\mathbb{S}_1), \mathbb{S}_2)} = \mathfrak{s}_2}} F_{Y, \mathbb{S}_2 \cup \mathbb{S}_1, \mathfrak{s}}$$

9 Constructing the graph TQFT, I

In this section, we describe our maps for graph cobordisms which satisfy the following condition:

Definition 9.1 A cobordism $W: Y_0 \rightarrow Y_1$ has enough ends if each connected component of W intersects both Y_0 and Y_1 nontrivially.

In the subsequent Section 10, we define the maps cobordisms which may not have enough ends.

9.1 Cerf-theoretic preliminaries

We need the following notion of a parametrized decomposition of a cobordism from [Juhász 2016, Section 8.1]:

Definition 9.2 Suppose that $W: Y_0 \rightarrow Y_1$ is a cobordism with enough ends. A parametrized Kirby decomposition \mathcal{K} of W consists of the following data:

- (1) a decomposition

$$W = W_n \circ \dots \circ W_0,$$

where W_i is a cobordism from \mathcal{Y}_i to \mathcal{Y}_{i+1} , and $Y_0 = \mathcal{Y}_0$ and $Y_1 = \mathcal{Y}_{n+1}$;

- (2) for each $i \in \{0, \dots, n\}$, a framed link $\mathbb{S}_i \subseteq \mathcal{Y}_i$, all of whose components have the same dimension; the possibility $\mathbb{S}_i = \emptyset$ is not excluded;
- (3) for each $i \in \{0, \dots, n\}$, a diffeomorphism $\Phi_i: W(\mathcal{Y}_i, \mathbb{S}_i) \rightarrow W_i$, defined up to isotopy, such that $\Phi_i(0, y) = y$ for all $y \in \mathcal{Y}_i$.

Furthermore, the following are satisfied:

- (1) There is a $c \in \{0, \dots, n\}$ such that \mathbb{S}_c has dimension 1.
- (2) If $c > i$, then \mathbb{S}_i has dimension 0. If $i > c$, then \mathbb{S}_i has dimension 2. (Note we allow \emptyset to have any dimension.)

- (3) If $\dim(\mathbb{S}_i) = 0$, then \mathbb{S}_i is empty or has two components. If $\dim(\mathbb{S}_i) = 2$, \mathbb{S}_i is empty or has one component.

Proposition 9.3 *Any two parametrized Kirby decompositions can be connected by a sequence of the following moves and their inverses:*

- (KM1) *adding or removing levels with $\mathbb{S} = \emptyset$;*
- (KM2) *pushing \mathcal{K} forward under a diffeomorphism of W , which is the identity on ∂W and is isotopic to the identity relative to ∂W ;*
- (KM3) *exchanging the relative ordering of two framed 0-spheres which are in adjacent levels;*
- (KM4) *exchanging the relative ordering of two framed 2-spheres which are in adjacent levels;*
- (KM5) *handlesliding the components of \mathbb{S}_c (a framed 1-dimensional link) across each other;*
- (KM6) *canceling a framed 0-sphere \mathbb{S}^0 with a framed 1-sphere \mathbb{K}^1 in the subsequent level if the belt sphere of \mathbb{S}^0 intersects \mathbb{K}^1 transversely in a single point;*
- (KM7) *canceling a framed 2-sphere \mathbb{S}^2 with a framed 1-sphere \mathbb{K}^1 in the previous level if the belt sphere of \mathbb{K}^1 intersects \mathbb{S}^2 transversely in a single point.*

We omit the proof of Proposition 9.3, and instead refer the reader to [Juhász 2018, Section 2; 2016, Theorem 8.9] for a careful Morse-theoretic argument. We content ourselves with a more detailed topological description of the moves in Proposition 9.3.

Note that in moves (KM3)–(KM7), we have not explained how the parametrizing diffeomorphisms are related. To this end, we note that if \mathbb{S} and \mathbb{S}' are pairwise disjoint framed spheres in Y , then there are canonical diffeomorphisms

$$W(Y(\mathbb{S}), \mathbb{S}') \circ W(Y, \mathbb{S}) \cong W(Y, \mathbb{S} \cup \mathbb{S}') \cong W(Y(\mathbb{S}'), \mathbb{S}) \circ W(Y, \mathbb{S}').$$

For example, the diffeomorphism $W(Y(\mathbb{S}), \mathbb{S}') \circ W(Y, \mathbb{S}) \cong W(Y, \mathbb{S} \cup \mathbb{S}')$ is obtained by noting that $W(Y(\mathbb{S}), \mathbb{S}') \circ W(Y, \mathbb{S})$ can be constructed by inserting a product layer $[0, 1] \times Y(\mathbb{S})$ into $W(Y, \mathbb{S} \cup \mathbb{S}')$. Hence, a diffeomorphism is obtained by picking a collar neighborhood of $Y(\mathbb{S})$ in $W(Y, \mathbb{S})$, which is unique up to isotopy. This describes the change in parametrizing diffeomorphisms in moves (KM3) and (KM4).

The change in parametrizing diffeomorphisms from move (KM5), a handleslide, is specified similarly. Suppose \mathbb{K} and \mathbb{K}' are two framed knots in Y , and \mathbb{K}'' is obtained by handlesliding \mathbb{K}' across \mathbb{K} . By our previous argument, there is a diffeomorphism between $W(Y, \mathbb{K} \cup \mathbb{K}')$ and $W(Y(\mathbb{K}), \mathbb{K}') \circ W(Y, \mathbb{K})$, which is well defined, up to isotopy. Next, $W(Y(\mathbb{K}), \mathbb{K}')$ and $W(Y(\mathbb{K}), \mathbb{K}'')$ are canonically diffeomorphic (up to isotopy), since \mathbb{K}' and \mathbb{K}'' are isotopic in $Y(\mathbb{K})$. Finally, $W(Y(\mathbb{K}), \mathbb{K}'') \circ W(Y, \mathbb{K})$ and $W(Y, \mathbb{K} \cup \mathbb{K}'')$ are canonically diffeomorphic, up to isotopy, by our previous argument on reordering handles. Composing these diffeomorphisms gives a diffeomorphism between $W(Y, \mathbb{K} \cup \mathbb{K}')$ and $W(Y, \mathbb{K} \cup \mathbb{K}'')$ which is well defined up to isotopy.

For the change in parametrizing diffeomorphism after handle cancellations, moves (KM6) and (KM7), we refer the reader to [Juhász 2018, Definition 2.17].

9.2 Graphs in 4-space

We need the following transversality result concerning the intersection of graphs and the ascending and descending manifolds of a Morse function:

Lemma 9.4 *Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a graph cobordism and f is a Morse function on W , with gradient-like vector field v .*

- (1) *If Γ is disjoint from $\text{Crit}(f)$, then, for generic v , Γ is disjoint from the descending manifolds of the index 1 critical points, and the ascending manifolds of the index 3 critical points. Generically, Γ is disjoint from both the ascending and descending manifolds of the index 2 critical points.*
- (2) *Suppose Γ is a fixed, abstract graph, and $i_t: \Gamma \rightarrow W$ is a family of embeddings, whose intersection with ∂W is fixed for all t . For generic i_t , the image of Γ is disjoint from $\text{Crit}(f)$, and is disjoint from the descending manifolds of the index 1 critical points and the ascending manifolds of the index 3 critical points. Generically, there are finitely many t where $i_t(\Gamma)$ transversely intersects the ascending or descending manifold of an index 2 critical point of f . Generically, such intersections occur along the interior of an edge of Γ .*

Proof We begin with the first claim, concerning a fixed graph. The descending manifolds of index 1 critical points of f are 1-dimensional, so, generically, a graph will be disjoint, since W is 4-dimensional. The same argument works for the ascending manifolds of index 3 critical points, which are also 1-dimensional. The ascending and descending manifolds of index 2 critical points are 2-dimensional, so a graph will be disjoint, generically.

The second claim, concerning 1-parameter families of graphs, follows from the same reasoning. \square

9.3 Definition of the graph cobordism maps

Let $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ be a graph cobordism with enough ends. Let \mathcal{K} be a parametrized Kirby decomposition of W , which decomposes W as

$$W = W_n \circ \cdots \circ W_0.$$

Let $c \in \{0, \dots, n\}$ denote the index of the 2-handle cobordism, and write \mathcal{Y}_c and \mathcal{Y}_{c+1} for the incoming and outgoing ends of W_c .

Suppose Γ is an embedded graph in W . The parametrizing diffeomorphisms of \mathcal{K} naturally equip each W_i with a Morse function and gradient-like vector field (f_i, v_i) , well defined up to isotopy in the handle attachment regions (see [Juhász 2018, Lemma 2.15]). We can assume that the (f_i, v_i) glue together to form a Morse function with gradient-like vector field (f, v) on all of W .

We can assume, after a small perturbation of the Morse functions, that Γ is disjoint from $\text{Crit}(f)$. By Lemma 9.4, for generically chosen v , the graph Γ is disjoint from the descending manifolds of the index 1

critical points, the ascending manifolds of the index 3 critical points, as well as both the ascending and descending manifolds of the index 2 critical points. Flow each point of Γ along v until it hits \mathcal{Y}_c . Note that we flow $\Gamma \cap (W_n \circ \dots \circ W_c)$ backwards along v , and we flow $\Gamma \cap (W_{c-1} \circ \dots \circ W_0)$ forward along v . Write

$$\Gamma_c \subseteq \mathcal{Y}_c$$

for the resulting graph. By perturbing v slightly, we may assume that Γ_c is embedded in \mathcal{Y}_c .

Upgrade Γ_c to a flow-graph \mathcal{G}_c by letting the initial and terminal vertices of \mathcal{G}_c be the images of the basepoints $\mathbf{w}_0 \subseteq Y_0$ and $\mathbf{w}_1 \subseteq Y_1$, under the flow of v .

Let $\mathbb{S}_0, \dots, \mathbb{S}_n$ denote the framed links in $\mathcal{Y}_0, \dots, \mathcal{Y}_n$, associated to \mathcal{K} , and let $\Phi_i: W(\mathcal{Y}_i, \mathbb{S}_i) \rightarrow W_i$ denote the parametrizing diffeomorphisms. Let $\phi_i: \mathcal{Y}_i(\mathbb{S}_i) \rightarrow \mathcal{Y}_{i+1}$ denote the restriction of Φ_i to the outgoing boundary of $W(\mathcal{Y}_i, \mathbb{S}_i)$.

The type- A graph cobordism map is defined as

$$(9-1) \quad F_{W,\Gamma,\mathfrak{s}}^A := \left(\prod_{i=c+1}^n (\phi_i)_* \circ F_{\mathcal{Y}_i, \mathbb{S}_i, \mathfrak{s}_i} \right) \circ ((\phi_i)_* \circ F_{\mathcal{Y}_c, \mathbb{S}_c, \mathfrak{s}_c}) \circ \mathfrak{A}_{\mathcal{G}_c} \circ \left(\prod_{i=0}^{c-1} (\phi_i)_* \circ F_{\mathcal{Y}_i, \mathbb{S}_i, \mathfrak{s}_i} \right),$$

where $\mathfrak{s}_i := \Phi_i^*(\mathfrak{s}|_{W_i})$.

We define the type- B maps by replacing $\mathfrak{A}_{\mathcal{G}}$ with $\mathfrak{B}_{\mathcal{G}}$, in (9-1).

Theorem 9.5 *Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a graph cobordism with enough ends.*

- (1) *For a fixed parametrized Kirby decomposition \mathcal{K} , the maps $F_{W,\Gamma,\mathfrak{s}}^A$ and $F_{W,\Gamma,\mathfrak{s}}^B$ are unchanged by smooth isotopies of the graph Γ (in the sense of Definition 2.4).*
- (2) *The maps $F_{W,\Gamma,\mathfrak{s}}^A$ and $F_{W,\Gamma,\mathfrak{s}}^B$ are independent from the choice of parametrized Kirby decomposition of W .*
- (3) *If $W = [0, 1] \times Y$, then $F_{W,\Gamma,\mathfrak{s}}^A \simeq \mathfrak{A}_{\mathcal{G}}$, where \mathcal{G} is an immersed ribbon flow-graph which is homotopic to the projection of Γ into Y . Similarly, $F_{W,\Gamma,\mathfrak{s}}^B \simeq \mathfrak{B}_{\mathcal{G}}$.*

Proof We first verify claim (1), independence from isotopies of Γ , for fixed \mathcal{K} . Lemma 9.4 describes the codimension 1 configurations which occur for a generic smooth isotopy Γ_t . At all but finitely many t , the graph Γ_t will be disjoint from the descending manifolds of the index 2 critical points. At all but these finitely many t , we may flow Γ_t along a gradient-like vector field (either upwards or downwards, depending on which portion of \mathcal{K} the graph Γ_t lies in) until it intersects \mathcal{Y}_c , and obtain an immersed flow-graph \mathcal{G}_t in \mathcal{Y}_c . Homotopies of immersed flow-graphs, which fix boundary vertices, do not affect the graph action map; see Remark 6.9. There are finitely many points of time t where Γ_t crosses a descending manifold of an index 2 critical point. If t_0 is such a value of t , and ϵ is sufficiently small, the induced flow-graphs $\mathcal{G}_{t_0-\epsilon}$ and $\mathcal{G}_{t_0+\epsilon}$ are related by sliding an edge of $\mathcal{G}_{t_0-\epsilon}$ across one of the link components of the framed link \mathbb{S}_c .

Let \mathcal{G}^- denote $\mathcal{G}_{t_0-\epsilon}$ and let \mathcal{G}^+ denote $\mathcal{G}_{t_0+\epsilon}$. We wish to show that

$$(9-2) \quad F_{\mathcal{Y}_c, \mathbb{S}_c, \mathfrak{s}_c} \circ \mathfrak{A}_{\mathcal{G}^-} \simeq F_{\mathcal{Y}_c, \mathbb{S}_c, \mathfrak{s}_c} \circ \mathfrak{A}_{\mathcal{G}^+}.$$




Figure 21: A region of a Heegaard triple subordinate to a framed knot which cancels a 1-handle. The almost complex structure is stretched along the curve c .

By Theorem 5.7 and Lemma 4.10, the graph action map commutes with 2-handle maps. However, in $Y(\mathbb{S})$, the flow-graphs \mathcal{G}^- and \mathcal{G}^+ are isotopic. Consequently,

$$F_{\mathcal{Y}_c, \mathcal{S}_c, \mathcal{S}_c} \circ \mathfrak{A}_{\mathcal{G}^-} \simeq \mathfrak{A}_{\mathcal{G}^-} \circ F_{\mathcal{Y}_c, \mathcal{S}_c, \mathcal{S}_c} \simeq \mathfrak{A}_{\mathcal{G}^+} \circ F_{\mathcal{Y}_c, \mathcal{S}_c, \mathcal{S}_c} \simeq F_{\mathcal{Y}_c, \mathcal{S}_c, \mathcal{S}_c} \circ \mathfrak{A}_{\mathcal{G}^+},$$

establishing (9-2). It follows that, for fixed \mathcal{K} , the maps $F_{\mathcal{W}, \Gamma, \mathfrak{s}}^A$ and $F_{\mathcal{W}, \Gamma, \mathfrak{s}}^B$ are invariant under smooth isotopies of the graph Γ .

We now show that $F_{\mathcal{W}, \Gamma, \mathfrak{s}}^A$ is independent of the choice of parametrized Kirby decomposition. It is sufficient to show invariance under the moves (KM1)–(KM7) from Proposition 9.3.

Invariance under move (KM1) (adding identity layers) is a tautology.

Invariance under move (KM2) (pushing \mathcal{K} forward under an isotopy) is established as follows. Let $(\Phi_t)_{t \in [0,1]}$ be an isotopy of W which is fixed on ∂W and satisfies $\Phi_0 = \text{id}_W$. Let \mathcal{K} be a parametrized Kirby decomposition, and let \mathcal{K}' denote the pushforward of \mathcal{K} under Φ_1 . Note that, tautologically, the map $F_{\mathcal{W}, \Gamma, \mathfrak{s}}^A$, computed with \mathcal{K} , agrees with $F_{\mathcal{W}, \Phi_1(\Gamma), \mathfrak{s}}^A$, computed with \mathcal{K}' . By our previous argument, the graph cobordism maps (for fixed \mathcal{K}) are invariant under smooth isotopies of the graph, so the map $F_{\mathcal{W}, \Phi_1(\Gamma), \mathfrak{s}}^A$, computed with \mathcal{K}' , coincides with the map $F_{\mathcal{W}, \Gamma, \mathfrak{s}}^A$, computed with \mathcal{K}' .

Invariance under moves (KM3) and (KM4) (reordering disjoint 1-handles or 3-handles) follows from Lemma 7.12.

Invariance under move (KM5) (handleslides) follows from [Ozsváth and Szabó 2006, Lemma 4.14], and is a consequence of invariance under the Heegaard Floer groups under handleslides of the α and β , together with an associativity argument.

Our proof of invariance under move (KM6) (canceling 1- and 2-handles) is formally the same as Ozsváth and Szabó’s original proof [2006, Lemma 4.16]. For notational simplicity, we focus on the case when we are computing the 2-handle map of a single framed knot \mathbb{K}^1 in Y which cancels a 1-handle, attached along $\mathbb{S}^0 \subseteq Y$. The general case when \mathbb{K}^1 is replaced by a framed link with multiple components follows from essentially the same argument. We can view \mathbb{S}^0 and \mathbb{K}^1 as being contained in a 3-ball in the original manifold Y . Hence, we can find a Heegaard triple subordinate to \mathbb{K}^1 in $Y(\mathbb{S}^0)$ which locally looks like Figure 21. Fix a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$. Let $\hat{\mathfrak{s}}$ denote the unique Spin^c structure

on $W(Y, \mathbb{S}^0)$ extending \mathfrak{s} , and let $t_0 \in \text{Spin}^c(W(Y(\mathbb{S}^0), \mathbb{K}^1))$ denote the unique Spin^c structure which evaluates trivially on the framed 2-sphere introduced by \mathbb{S}^0 , and restricts to \mathfrak{s} on $Y \cong Y(\mathbb{S}^0)(\mathbb{K}^1)$.

By Lemma 7.14, we can find almost complex structures on $\mathbb{T}^2 \times [0, 1] \times \mathbb{R}$ and $\Sigma \times [0, 1] \times \mathbb{R}$ such that, if $J(T)$ denotes the almost complex structure on $(\Sigma \# \mathbb{T}^2) \times [0, 1] \times \mathbb{R}$ obtained by gluing the two together, with neck length T along the circle c in Figure 21, then the 1-handle map $F_{Y, \mathbb{S}^0, \hat{\mathfrak{s}}}$ satisfies

$$(9-3) \quad F_{Y, \mathbb{S}^0, \hat{\mathfrak{s}}}(\mathbf{x}) = \mathbf{x} \times \theta_{\alpha_0, \beta_0}^+$$

If $(\Sigma, \alpha, \beta, \mathbf{w})$ denotes the original Heegaard diagram of Y , let β' denote small isotopies of β . The Heegaard triple $(\Sigma \# \mathbb{T}^2, \alpha \cup \{\alpha_0\}, \beta \cup \{\beta_0\}, \beta' \cup \{\beta'_0\}, \mathbf{w})$ is subordinate to a bouquet for \mathbb{K}^1 . By Theorem 8.7 and (9-3),

$$(9-4) \quad F_{Y(\mathbb{S}^0), \mathbb{K}^1, t_0}(F_{Y, \mathbb{S}^0, \hat{\mathfrak{s}}}(\mathbf{x})) = F_{Y(\mathbb{S}^0), \mathbb{K}^1, t_0}(\mathbf{x} \times \theta_{\alpha_0, \beta_0}^+) = \sigma(\Psi_{\alpha}^{\beta \rightarrow \beta'}(\mathbf{x})),$$

where σ and $\Psi_{\alpha}^{\beta \rightarrow \beta'}$ denote the maps from naturality associated to a simple stabilization and to the small isotopy moving β to β' , respectively. Invariance under move (KM6) follows.

Invariance under move (KM7) follows by turning around the above argument for move (KM6).

We now verify claim (3). Assume $W = [0, 1] \times Y$. In this case, we can pick \mathcal{K} to have a single level, which has the empty framed sphere. The map $F_{W, \Gamma, \mathfrak{s}}^A$ is, by definition, obtained by projecting Γ into Y and computing the map $\mathfrak{A}_{\mathcal{G}}$. The same argument works for the type- B map. \square

10 Constructing the graph TQFT, II: adding punctures

In this section, we define the cobordism maps when a cobordism does not have enough ends (in the sense of Definition 9.1). The idea is to remove small 4-balls from W , and connect the new copies of S^3 to Γ by adding a new edge.

Combined with our proof of invariance for graph cobordisms with enough ends (Theorem 9.5), this section concludes the construction of the graph cobordism maps and the proof of their invariance (Theorem A).

For a fixed P , the chain complex $\text{CF}^-(\emptyset)$ is defined to be the ring $\mathcal{R}_P := \mathbb{F}_2[U_{p_1}, \dots, U_{p_n}]$, where $P = \{p_1, \dots, p_n\}$, with vanishing differential.

10.1 0- and 4-handle maps

Suppose that (Y, \mathbf{w}) is a multipointed 3-manifold, equipped with a coloring $\sigma: \mathbf{w} \rightarrow P$. Let (S^3, w_0) denote a new copy of S^3 , and let $\sigma': \mathbf{w} \cup \{w_0\} \rightarrow P$ be an extension of σ .

If $(\Sigma, \alpha, \beta, \mathbf{w})$ is a diagram for (Y, \mathbf{w}) , and (S^2, w_0) is a genus 0 diagram for S^3 with no α or β curves (alternatively, if one wants to avoid Heegaard diagrams with no curves, one can use a genus 1 diagram for S^3 , with a single α and β curve).

There is a canonical chain isomorphism

$$(10-1) \quad \text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}^\sigma, \mathfrak{s}) \cong \text{CF}^-(\Sigma \cup S^2, \alpha, \beta, (\mathbf{w} \cup \{w_0\})^{\sigma'}, \mathfrak{s} \cup \mathfrak{s}_0),$$

obtained by sending $U_{p_1}^{i_1} \cdots U_{p_n}^{i_n} \cdot \mathbf{x}$ to $U_{p_1}^{i_1} \cdots U_{p_n}^{i_n} \cdot \mathbf{x}$, where $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and U_{p_1}, \dots, U_{p_n} are the variables in the ring \mathcal{R}_P .

The 4-manifold corresponding to a 0-handle attachment is the disjoint union of $[0, 1] \times Y$ with B^4 .

Writing \mathfrak{s}_0 for the unique element of $\text{Spin}^c(S^3)$, we define the 0-handle map

$$F_{Y,0}: \text{CF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s}) \rightarrow \text{CF}^-(Y \cup S^3, (\mathbf{w} \cup \{w_0\})^{\sigma'}, \mathfrak{s} \cup \mathfrak{s}_0)$$

using the canonical isomorphism in (10-1).

There is also a 4-handle map

$$F_{Y \cup S^3,4}: \text{CF}^-(Y \cup S^3, (\mathbf{w} \cup \{w_0\})^{\sigma'}, \mathfrak{s} \cup \mathfrak{s}_0) \rightarrow \text{CF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s}),$$

also defined using the canonical isomorphism in (10-1).

10.2 Puncturing graph cobordisms

Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a graph cobordism without enough ends. Construct a new 4-manifold W' by removing a collection of 4-balls from W , which are disjoint from Γ . For each new copy of S^3 in $\partial W'$, we add a new edge to Γ which connects the new S^3 to a point along the interior of an edge of Γ . The new vertices may be labeled with any cyclic order. We designate the new copies of S^3 in the boundary as incoming or outgoing, in such a way that (W', Γ') has enough ends, in the sense of Definition 9.2. Let F_0 denote the 0-handle maps corresponding to the new incoming boundary components, and let F_4 denote the 4-handle maps corresponding to the new outgoing boundary components.

We now define

$$(10-2) \quad F_{W,\Gamma,\mathfrak{s}}^A := F_4 \circ F_{W',\Gamma',\mathfrak{s}|_{W'}}^A \circ F_0,$$

and define $F_{W,\Gamma,\mathfrak{s}}^B$ similarly.

It remains to show that (10-2) is well defined. Since any two puncturings of (W, Γ) can be related by a common puncturing (up to diffeomorphism), it is sufficient to show that the puncturing operation does not change the maps for cobordisms which already have enough boundary components:

Proposition 10.1 *Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a graph cobordism with enough ends, and (W', Γ') is obtained by removing a 4-ball from W and connecting the new boundary component to a point on the interior of an edge of Γ . If the new S^3 is designated as incoming, then*

$$F_{W,\Gamma,\mathfrak{s}}^A \simeq F_{W',\Gamma',\mathfrak{s}|_{W'}}^A \circ F_{Y_0,0}.$$

If the new S^3 is designated as outgoing, then

$$F_{W,\Gamma,\mathfrak{s}}^A \simeq F_{Y_1 \cup S^3,4} \circ F_{W',\Gamma',\mathfrak{s}|_{W'}}^A.$$

Proof We focus on the case when the new S^3 is designated as incoming. The argument for the other case is similar. A handle decomposition of W' is obtained from a handle decomposition of W by adding in a

1-handle which connects the new copy of S^3 to Y_0 . Let S^0 denote the corresponding 0-sphere in $Y \cup S^3$, and let $F_{Y_0 \cup S^3, S^0}$ denote the map for this 1-handle. Immediately from the definitions, we have

$$(10-3) \quad F_{Y_0 \cup S^3, S^0} \circ F_{Y_0, 0} = S_{w_0}^+,$$

where w_0 is the basepoint which was added by the 0-handle map.

Let \mathcal{Y}_c denote the level set in W corresponding to the incoming boundary of the 2-handle portion of W , and let $\mathcal{G}_c \subseteq \mathcal{Y}_c$ denote the flow-graph obtained by isotoping Γ along the flowlines of a gradient-like vector field on W . Let \mathcal{G}'_c denote the flow-graph obtained by isotoping Γ' into \mathcal{Y}_c using the flow of a gradient-like vector field on W' . Let w'_0 denote the image of w_0 in \mathcal{Y}_c . We can assume that \mathcal{G}'_c is obtained from \mathcal{G}_c by connecting w'_0 to the interior of an edge of \mathcal{G}_c with an arc.

Using (10-3) and the fact $S_{w_0}^+$ can be commuted with the 1-handle maps of W using Lemma 7.13, the main statement reduces to showing that

$$\mathfrak{A}_{\mathcal{G}_c} = \mathfrak{A}_{\mathcal{G}'_c} \circ S_{w'_0}^+.$$

We note that, by definition, the function $\mathfrak{A}_{\mathcal{G}'_c} \circ S_{w'_0}^+$ is equal to the graph action map of \mathcal{G}'_c , with w'_0 viewed as an interior vertex. By the trivial strand relation in Lemma 6.15, the induced map is chain homotopic to $\mathfrak{A}_{\mathcal{G}_c}$. □

11 The composition law

We now prove the composition law for graph cobordisms:

Theorem C *Suppose that a graph cobordism (W, Γ) can be decomposed as the composition*

$$(W, \Gamma) = (W_2, \Gamma_2) \circ (W_1, \Gamma_1).$$

If $\mathfrak{s}_1 \in \text{Spin}^c(W_1)$ and $\mathfrak{s}_2 \in \text{Spin}^c(W_2)$, then

$$F_{W_2, \Gamma_2, \mathfrak{s}_2}^A \circ F_{W_1, \Gamma_1, \mathfrak{s}_1}^A \simeq \sum_{\substack{\mathfrak{s} \in \text{Spin}^c(W) \\ \mathfrak{s}|_{W_1} = \mathfrak{s}_1 \\ \mathfrak{s}|_{W_2} = \mathfrak{s}_2}} F_{W, \Gamma, \mathfrak{s}}^A.$$

The same relation holds for the type-B maps.

Proof First, note that the 0-handle maps of $F_{W_2, \Gamma_2, \mathfrak{s}_2}^A$ can be commuted to the right of all of the factors of $F_{W_1, \Gamma_1, \mathfrak{s}_1}^A$. Similarly, the 4-handle maps of $F_{W_1, \Gamma_1, \mathfrak{s}_1}^A$ can be commuted to the left of all of the factors of $F_{W_2, \Gamma_2, \mathfrak{s}_2}^A$. Consequently, it is sufficient to prove the composition law for cobordisms with enough ends (in the sense of Definition 9.1).

Pick parametrized Kirby decompositions \mathcal{K}_1 and \mathcal{K}_2 , for W_1 and W_2 , respectively. We can stack \mathcal{K}_1 and \mathcal{K}_2 to get a decomposition of \mathcal{K} of W into elementary cobordisms, though \mathcal{K} will not in general be a parametrized Kirby decomposition, as the handles do not appear in the correct order, and there are two

levels with 2-handles. We will refer to a decomposition of W into parametrized cobordisms, with at most two levels with 2-handles, as a *semi-Kirby decomposition* of W .

Let \mathcal{G}_1 denote the flow-graph in W_1 obtained by isotoping Γ_1 along the flowlines of a gradient-like vector field until it lies in the incoming boundary of the level of W_1 containing the 2-handles. Let \mathcal{G}_2 denote the analogous flow-graph in a level of W_2 .

The ascending manifolds of the index 3 critical points of W_1 are (after generic perturbation) disjoint from the descending manifolds of the index 1 and 2 critical points of W_2 , as well as the flow-graph \mathcal{G}_2 . Consequently, we may topologically move all of the 3-handles of \mathcal{K}_1 to the left of the 1-handles, 2-handles and flow-graph of W_2 , to obtain a new semi-Kirby decomposition of W . Moving the 3-handles past the 1-handles and 2-handles of W_2 does not affect the composition by Theorem 7.8 and Lemma 7.12. Commuting 3-handle maps past a graph action map for a flow-graph which is disjoint from the attaching sphere amounts to commuting the free-stabilization and relative homology maps (which are used to build the graph action map) past the 3-handle map. The free-stabilization and relative homology maps commute with the 3-handle maps by Lemmas 7.11 and 7.13.

In an entirely analogous manner, the 1-handle maps of W_2 may be commuted to the right of the 3-handles, 2-handles and flow-graph of W_1 .

Generically, the ascending manifolds of the 2-handles in W_2 will be disjoint from \mathcal{G}_2 , and consequently we may topologically move \mathcal{G}_2 past the 2-handles of W_2 and obtain a flow-graph \mathcal{G}'_2 in the same level as \mathcal{G}_1 . Commuting the graph action map of \mathcal{G}_2 with the 2-handle map of W_1 amounts to commuting the free-stabilization and relative homology maps past the 2-handle map, which can be done using Lemma 4.10 and Theorem 5.7.

By Theorem 6.1(b), we have $\mathfrak{A}_{\mathcal{G}'_2} \circ \mathfrak{A}_{\mathcal{G}_1} = \mathfrak{A}_{\mathcal{G}'_2 \cup \mathcal{G}_1}$.

In summary, we have shown that the composition $F_{W_2, \Gamma_2, \mathfrak{s}_2}^A \circ F_{W_1, \Gamma_1, \mathfrak{s}_1}^A$ is equal to the composition of handle maps and graph action map of a semi-Kirby decomposition, which fails to be a parametrized Kirby decomposition only in that it has two levels with 2-handles. Since these two levels are consecutive, by using the composition law for 2-handle maps in Lemma 8.11, the general Spin^c composition law is obtained. \square

12 Path cobordisms and the normalization axiom

Suppose (Y, \mathbf{w}) is a multipointed 3-manifold. Using the composition law, if $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y, \mathbf{w})$ is a graph cobordism and $\mathfrak{s} \in \text{Spin}^c(W)$, then

$$F_{[0,1] \times Y, [0,1] \times \mathbf{w}, \mathfrak{s}}^A \circ F_{W, \Gamma, \mathfrak{s}}^A \simeq F_{W, \Gamma, \mathfrak{s}}^A.$$

This of course does not imply that $F_{[0,1] \times Y, [0,1] \times \mathbf{w}, \mathfrak{s}}^A$ is the identity map. In the study of TQFTs, the relation

$$F_{[0,1] \times Y, [0,1] \times \mathbf{w}, \mathfrak{t}}^A \simeq \text{id}_{\text{CF}^-(Y, \mathbf{w}, \mathfrak{t})}$$

is often referred to as the *normalization axiom* (see [Turaev 1994, Section 1.4]). In this section, we prove the following (modulo a technical result, which we prove in Section 13):

Theorem B Suppose that $(W, \Gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a graph cobordism, and Γ is a collection of paths, each connecting \mathbf{w}_0 to \mathbf{w}_1 .

(1) The *A* and *B* versions coincide:

$$F_{W,\Gamma,s}^A \simeq F_{W,\Gamma,s}^B.$$

(2) Suppose $\phi: (Y, \mathbf{w}) \rightarrow (Y, \mathbf{w})$ is an orientation-preserving diffeomorphism, and let $W(\phi)$ denote the **mapping cylinder** (i.e., $[0, 1] \times Y$, with $\{0\} \times Y$ identified with Y via id_Y and $\{1\} \times Y$ identified with Y via ϕ). Then

$$F_{W(\phi),[0,1] \times \mathbf{w},s}^A \simeq F_{W(\phi),[0,1] \times \mathbf{w},s}^B \simeq (\phi_*: \text{CF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s}) \rightarrow \text{CF}^-(Y, \mathbf{w}^{\phi_*\sigma}, \phi_*\mathfrak{s})).$$

(3) Suppose $(W, \gamma): (Y_0, \mathbf{w}_0) \rightarrow (Y_1, \mathbf{w}_1)$ is a cobordism such that W, Y_0 and Y_1 are nonempty and connected, and γ is a path from \mathbf{w}_0 to \mathbf{w}_1 . Then $F_{W,\gamma,s}^A \simeq F_{W,\gamma,s}^B$, and both maps coincide with the map defined by Ozsváth and Szabó.

Proof Using naturality of the Heegaard Floer groups, the first and second claims amount to showing that if $\mathcal{G} = (\Gamma, \mathbf{w}_0, \mathbf{w}_1)$ is a flow-graph in Y such that \mathbf{w}_0 and \mathbf{w}_1 are pairwise disjoint collections of points in Y , and Γ consists of arcs, each connecting a point of \mathbf{w}_0 to a point of \mathbf{w}_1 , then

$$(12-1) \quad \mathfrak{A}_{\mathcal{G}} \simeq \mathfrak{B}_{\mathcal{G}} \simeq (\phi_{\Gamma})_*,$$

where $\phi_{\Gamma}: (Y, \mathbf{w}_0) \rightarrow (Y, \mathbf{w}_1)$ is the diffeomorphism obtained by moving \mathbf{w}_0 to \mathbf{w}_1 along Γ .

Similarly, to prove the third claim, we note that our maps $F_{W,\gamma,s}^A$ and $F_{W,\gamma,s}^B$ are defined in (9-1) by writing W as a composition of 1-, 2- and 3-handles. The equivalence of our definition of the 1- and 3-handle maps with Ozsváth and Szabó’s construction is established in Lemma 7.14. Our definition of the 2-handle maps coincides exactly with their definition. The only remaining difference between our map and theirs is that our map includes a graph action map between the 1- and 2-handle maps, for a flow-graph that consists of a single arc. Equation (12-1) implies that this graph action map is chain homotopic to a basepoint-moving diffeomorphism map, from which the main claim follows.

It remains to establish (12-1). We leave this for the final section; see Theorem 13.11. □

13 Basepoint-moving maps and the normalization axiom

If λ is a path in Y from w to w' , in this section we prove that the diffeomorphism map λ_* , obtained by pushing w to w' along λ , satisfies

$$(13-1) \quad \lambda_* \simeq S_w^- A_{\lambda} S_{w'}^+.$$

See Theorem 13.11. Using (13-1), we finish our proof of Theorem B (the normalization axiom) and prove Corollary F (our formula for the π_1 -action).




Figure 22: The diagrams \mathcal{H}_1 and \mathcal{H}_2 considered in Theorem 13.1.

13.1 A transition map computation

A key ingredient in our proof of Theorem 13.11 is a computation of the transition map shown in Figure 22.

Suppose that $\mathcal{H}_0 = (\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{z\})$ is a multipointed Heegaard diagram and z is a distinguished basepoint. We consider the free-stabilized diagrams,

$$\mathcal{H}_1 = (\Sigma, \alpha \cup \{\alpha_1\}, \beta \cup \{\beta_1\}, \mathbf{w}_0 \cup \{w, w'\}) \quad \text{and} \quad \mathcal{H}_2 = (\Sigma, \alpha' \cup \{\alpha_2\}, \beta' \cup \{\beta_2\}, \mathbf{w}_0 \cup \{w, w'\}),$$

obtained by removing z from \mathcal{H}_0 , and inserting the diagrams shown in Figure 22 into the region which contained z . The curves α' and β' are small Hamiltonian isotopies of α and β . The case that $\mathbf{w}_0 = \emptyset$ is not excluded. A key step in our proof of (13-1) is to compute the transition map between \mathcal{H}_1 and \mathcal{H}_2 .

We introduce some notation. If C is a module over the ring $\mathbb{F}_2[U_z]$, we write $C^{U_z \rightarrow U_w}$ for the tensor product

$$C^{U_z \rightarrow U_w} := C \otimes_{\mathbb{F}_2[U_z]} \mathbb{F}_2[U_z, U_w] / (U_w - U_z),$$

which we think of as the module obtained by formally setting U_z to the variable U_w . If $F: C_1 \rightarrow C_2$ is a map of $\mathbb{F}_2[U_z]$ -modules, write $F^{U_z \rightarrow U_w}$ for the induced map

$$F^{U_z \rightarrow U_w} := F \otimes \text{id}_{\mathbb{F}_2[U_z, U_w] / (U_w - U_z)}.$$

If R is a ring of characteristic 2 and $F: C_1 \rightarrow C_2$ is a map of R -modules, write $(F)_{U_w}$ for the map of $R \otimes_{\mathbb{F}_2} \mathbb{F}_2[U_w]$ -modules

$$(13-2) \quad F_{U_w} := F \otimes \text{id}_{\mathbb{F}_2[U_w]}: C_1 \otimes_{\mathbb{F}_2} \mathbb{F}_2[U_w] \rightarrow C_2 \otimes_{\mathbb{F}_2} \mathbb{F}_2[U_w].$$

Finally, we introduce a convenient matrix notation. Let V denote the 2-dimensional vector space $\langle \theta^+, \theta^- \rangle$. If F is a homomorphism

$$F: C_1 \otimes_{\mathbb{F}_2} V \rightarrow C_2 \otimes_{\mathbb{F}_2} V,$$

then we will write F as a 2×2 block matrix. We always use the ordered basis (θ^+, θ^-) , so the first row and column of such a matrix correspond to θ^+ , and the second row and column to θ^- .

Theorem 13.1 *Let \mathcal{H}_1 and \mathcal{H}_2 denote the free-stabilized diagrams in Figure 22. There are choices of almost complex structures J_1 and J_2 on \mathcal{H}_1 and \mathcal{H}_2 , respectively, such that J_1 can be used to compute*




Figure 23: The intermediate diagram $\mathcal{H}_{1.5}$.

S_w^+ and S_w^- , and J_2 can be used to compute $S_{w'}^+$ and $S_{w'}^-$, such that (for some choice of intermediate diagrams and almost complex structures)

$$\Psi_{(\mathcal{H}_1, J_1) \rightarrow (\mathcal{H}_2, J_2)} = \begin{pmatrix} (\Psi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_w}^{U_z \rightarrow U_{w'}} \circ (\Psi_{\alpha}^{\beta \rightarrow \beta'})_{U_{w'}}^{U_z \rightarrow U_w} & 0 \\ * & (\Psi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_w}^{U_z \rightarrow U_{w'}} \circ (\Psi_{\alpha}^{\beta \rightarrow \beta'})_{U_{w'}}^{U_z \rightarrow U_w} \end{pmatrix}.$$

Here $\Psi_{\alpha}^{\beta \rightarrow \beta'}$ denotes the transition map from $CF^-(\Sigma, \alpha, \beta, w_0 \cup \{z\}, \mathfrak{s})$ to $CF^-(\Sigma, \alpha, \beta', w_0 \cup \{z\}, \mathfrak{s})$, and $\Psi_{\alpha \rightarrow \alpha'}^{\beta'}$ is defined similarly.

We remark that the component marked $*$ can be computed explicitly. It is equal to

$$(13-3) \quad (\Phi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_w}^{U_z \rightarrow U_{w'}} \circ \left(\sum_{i, j \geq 0} U_w^i U_{w'}^j (\partial_{i+j+1})_{U_w, U_{w'}} \right) \circ (\Phi_{\alpha}^{\beta \rightarrow \beta'})_{U_{w'}}^{U_z \rightarrow U_w},$$

though we do not need this fact to prove Theorem 13.11. That $*$ equals (13-3) can be proven by a small modification of Proposition 13.22, below. In (13-3), the expression ∂_{i+j+1} is obtained by writing the differential on $CF^-(\mathcal{H}_0, \mathfrak{s})$ as

$$(\partial_{\mathcal{H}_0}) = \sum_{i=0}^{\infty} \partial_i U_z^i.$$

The proof of Theorem 13.1 involves computing several holomorphic triangle maps and several non-cylindrical holomorphic strip counts.

On $\mathcal{H}_1, \mathcal{H}_{1.5}$ (see Figure 23) and \mathcal{H}_2 we will write J_{α} for an almost complex structure which is stretched along the circles c and c_{α} in Figure 24. We write J_{β} for an almost complex structure which is stretched along the circles c and c_{β} in Figure 24.

To help simplify the statements of some of the results in this section, we prove the following:

Lemma 13.2 *Let \mathcal{H} be one of $\mathcal{H}_1, \mathcal{H}_{1.5}$ or \mathcal{H}_2 . If $T = (T, T')$, write $J_{\alpha}(T)$ for an almost complex structure which has been stretched along c and c_{α} , with neck lengths T and T' . There is a constant N such that if T_1 and T_2 are two pairs of neck lengths all of whose components are greater than N , then there is a noncylindrical almost complex structure \tilde{J} interpolating $J_{\alpha}(T_1)$ and $J_{\alpha}(T_2)$ satisfying*

$$\Psi_{J_{\alpha}(T_1) \rightarrow J_{\alpha}(T_2)} := \Psi_{\tilde{J}} = \begin{pmatrix} \text{id} & 0 \\ 0 & \text{id} \end{pmatrix}.$$




Figure 24: The almost complex structures J_α and J_β for $\mathcal{H}_{1.5}$.

Proof We focus on the case that $\mathcal{H} = \mathcal{H}_1$. The statement is proven similarly for $\mathcal{H}_{1.5}$ and \mathcal{H}_2 . The proof is a double neck-stretching argument. Suppose that $T_{1,i}$ and $T_{2,i}$ are two sequences of pairs of neck lengths, all of whose components approach $+\infty$. Write $T_{1,i} = (T_{1,i}, T'_{1,i})$ and $T_{2,i} = (T_{2,i}, T'_{2,i})$. Define

$$T_{\min,i} = \min(T_{1,i}, T_{2,i}) \quad \text{and} \quad T'_{\min,i} = \min(T'_{1,i}, T'_{2,i}).$$

We decompose neighborhoods of c and c_α as unions of three annuli, as shown in Figure 25:

$$N(c) = N_1 \cup N_2 \cup N_3 \quad \text{and} \quad N(c_\alpha) = N'_1 \cup N'_2 \cup N'_3.$$

Construct interpolating almost complex structures \tilde{J}_i between $J_\alpha(T_{1,i})$ and $J_\alpha(T_{2,i})$. We require that the almost complex structures \tilde{J}_i be split on each of N_1, N_3, N'_1 and N'_3 (in particular, \tilde{J}_i is cylindrical on these regions). Furthermore, we require that \tilde{J}_i be chosen so that N_1 and N_3 are both conformally equivalent to $S^1 \times [0, \frac{1}{3}T_{\min,i}]$, and so that N'_1 and N'_3 are conformally equivalent to $S^1 \times [0, \frac{1}{3}T'_{\min,i}]$. Further, we assume that \tilde{J}_i is only noncylindrical in the regions N_2 and N'_2 . We can pick \tilde{J}_i so that projection from $\Sigma \times [0, 1] \times \mathbb{R}$ to $[0, 1] \times \mathbb{R}$ is holomorphic.




Figure 25: Decomposing the stretching regions in Lemma 13.2.

Write

$$(13-4) \quad \Psi_{\tilde{J}_i} = \begin{pmatrix} A_i & B_i \\ C_i & D_i \end{pmatrix}.$$

If $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \mathbf{x}, \mathbf{y} \times \mathbf{y})$ is a Maslov index 0 class of disks, then (5-6) implies

$$(13-5) \quad \mu(\phi \# \phi_0) = \mu(\phi) + \text{gr}(x, y) + 2m_2(\phi_0).$$

Classes with $\text{gr}(x, y) = 1$ contribute to C_i , classes with $\text{gr}(x, y) = 0$ contribute to A_i or D_i , and classes with $\text{gr}(x, y) = -1$ contribute to B_i .

Given a sequence u_i of \tilde{J}_i -holomorphic curves representing $\phi \# \phi_0$, we may extract a subsequence which converges to broken holomorphic curves on the diagrams $(S^2, \alpha_1, \beta_1^l)$, (S^2, β_1^m) and (Σ, α, β) , whose total homology class is $\phi \# \phi_0$. The curves β_1^l and β_1^m are the result of cutting β_1 along its intersection with c_α , and then collapsing the ends to a point; see Figure 25. Furthermore, the broken limiting curves are pseudoholomorphic for cylindrical almost complex structures.

In somewhat more detail, to construct such a convergent subsequence, let D_0 denote the disk component of $\Sigma \setminus N(c_\alpha)$ containing α_1 , and let A_0 denote the annulus between $N(c)$ and $N(c_\alpha)$, as in Figure 25. Let Σ_0 denote the component of $\Sigma \setminus N(c)$ which is disjoint from w and w' .

If u_i is \tilde{J}_i -holomorphic, we take preimages of u_i into subregions of $\Sigma \times [0, 1] \times \mathbb{R}$ to construct holomorphic curves with additional boundary circles, u_i^l, u_i^m and u_i^r , which map into the 4-manifolds $(D_0 \cup N'_3) \times [0, 1] \times \mathbb{R}$, $(N'_1 \cup A_0 \cup N_3) \times [0, 1] \times \mathbb{R}$ and $(N_1 \cup \Sigma_0) \times [0, 1] \times \mathbb{R}$, respectively. These are holomorphic curves for cylindrical almost complex structures. We can view $((D_0 \cup N'_3) \times [0, 1] \times \mathbb{R}, \tilde{J}_i)$ as being contained in $((D_0 \cup N'_3) \times [0, 1] \times \mathbb{R}, \tilde{J}_k)$ whenever $T_{\min,i} \leq T_{\min,k}$. Consequently, given such a sequence u_i , we may find a subsequence such that u_i^l, u_i^m and u_i^r each converge to curves in the punctured manifolds $S^2 \setminus \{p_0\} \times [0, 1] \times \mathbb{R}$, $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$ and $\Sigma \setminus \{p\} \times [0, 1] \times \mathbb{R}$, where p_0 and p denote the connected sum points corresponding to the circles c_α and c , respectively.

Consequently, ϕ and ϕ_0 admit broken homomorphic representatives on (Σ, α, β) and $(S^2, \alpha_1, \beta_1^l)$. In particular, $\mu(\phi) \geq 0$, so (13-5) implies

$$m_2(\phi_0) = 0 \quad \text{and} \quad \text{gr}(x, y) \leq 0.$$

Hence, $B_i = 0$ for large i , since B_i counts curves with $\text{gr}(x, y) = +1$.

Next, we consider classes with $\text{gr}(x, y) = 0$ (which contribute to A_i and D_i). Since ϕ admits a broken representative and has index 0 by (13-5), we conclude that ϕ must represent the constant class e_x . It is straightforward to see that this also implies that ϕ_0 is the constant class e_x . Conversely, since \tilde{J}_i is cylindrical in a neighborhood of all the intersection points, the index 0 classes $e_x \times e_x$ have unique \tilde{J}_i -holomorphic representatives for all i . Consequently, $A_i = D_i = \text{id}$.

We now consider classes with $\text{gr}(x, y) = -1$, which contribute to C_i . For such classes, equation (13-5) and the inequality $\mu(\phi) \geq 0$ imply

$$\mu(\phi) = 1 \quad \text{and} \quad m_2(\phi_0) = 0.$$

Since $\mu(\phi) = 1$, transversality implies that the limiting curve on (Σ, α, β) is a nonbroken index 1 flowline. Since projection to $[0, 1] \times \mathbb{R}$ is \tilde{J}_i -holomorphic, the asymptotics of the curves on (S^2, β_1^m) must match the curve on (Σ, α, β) at the connected sum point. Since the limiting curve on (Σ, α, β) is a genuine flowline, its asymptotics at the connected sum point consist of $m_1(\phi_0)$ points in $[0, 1] \times \mathbb{R}$ (which we can assume are distinct, for a generically chosen almost complex structure). The only curve on the partial diagram (S^2, β_1^m) which could match $m_1(\phi_0)$ distinct points in $[0, 1] \times \mathbb{R}$ consists of $m_1(\phi_0)$ copies of S^2 which each map constantly to the $[0, 1] \times \mathbb{R}$ factor, as well as some collection of β_1^m boundary degenerations, which have $m_1 = 0$. The 2-spheres contribute equally to $m_1(\phi_0)$ and $n_2(\phi_0)$, while the boundary degenerations only contribute to $n_2(\phi_0)$. Consequently,

$$(13-6) \quad n_2(\phi_0) \geq m_1(\phi_0).$$

Since $\phi_0 \in \pi_2(\theta^-, \theta^+)$ and $m_2(\phi_0) = 0$, we must have

$$(13-7) \quad m_1(\phi_0) + 1 = n_1(\phi_0) + n_2(\phi_0),$$

which is satisfied for all classes $\phi_0 \in \pi_2(\theta^-, \theta^+)$ with $m_2(\phi_0) = 0$. Equation (13-7) contradicts (13-6), so C_i must be zero for large i . □

We now describe a refinement of the differential computation of Proposition 5.5. The following computation is similar to [Ozsváth and Szabó 2008, Proposition 6.5], though the placement of basepoints and choices of almost complex structures are different than they consider. Recall that $\mathcal{H}_{1.5}$ is the stabilization of \mathcal{H}_0 shown in Figure 23.

Lemma 13.3 *Let J_α denote an almost complex structure on $\mathcal{H}_{1.5}$ which is stretched along c and c_α . For sufficiently large neck lengths along c and c_α , we have*

$$\partial_{\mathcal{H}_{1.5}, J_\alpha} = \begin{pmatrix} (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_w} & U_w + U_{w'} \\ 0 & (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_w} \end{pmatrix}.$$

If J_β denotes an analogous almost complex structure stretched sufficiently along c and c_β , then

$$\partial_{\mathcal{H}_{1.5}, J_\beta} = \begin{pmatrix} (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_{w'}} & U_w + U_{w'} \\ 0 & (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_{w'}} \end{pmatrix}.$$

Proof We focus on computing $\partial_{\mathcal{H}_{1.5}, J_\alpha}$; the computation for J_β is a straightforward modification. Write

$$\partial_{\mathcal{H}_{1.5}, J_\alpha} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

By Lemma 13.2, if the claim holds any pair of arbitrarily large neck lengths, then it holds for all sufficiently large pairs of neck lengths (regardless of their relative lengths).

Let m_1, m_2, n_1 and n_2 denote the multiplicities of the regions of $(S^2, \alpha_1, \beta_2, w, w')$ shown in Figure 26.

Suppose $\phi \# \phi_0 \in \pi_2(x \times x, y \times y)$ is a homology class of disks on $\mathcal{H}_{1.5}$. Lemma 5.4 adapts to show

$$(13-8) \quad \mu(\phi \# \phi_0) = \mu(\phi) - \text{gr}(x, y) + 2m_2(\phi_0).$$




Figure 26: Stretching along c_α . On the left are multiplicities m_1, m_2, n_1 and n_2 . On the right we illustrate how the almost complex structure degenerates as we stretch J_α along c_α .

Classes with $\text{gr}(x, y) = 1$ contribute to C . Classes with $\text{gr}(x, y) = 0$ contribute to A or D . Classes with $\text{gr}(x, y) = -1$ contribute to B .

We begin by computing B . Suppose $\phi \# \phi_0 \in \pi_2(x \times \theta^-, y \times \theta^+)$ is a class counted by B . By stretching sufficiently along c , we may assume that ϕ has a broken representative, and consequently $\mu(\phi) \geq 0$. Equation (13-8) implies that $\mu(\phi) = m_2(\phi_0) = 0$. Since ϕ has a broken representative, ϕ is a constant class e_x . Since $m_2(\phi_0) = m_1(\phi_0) = 0$, we conclude that ϕ_0 must have domain equal to one of the two bigons going over w or w' . These classes also have a unique holomorphic representative. Consequently, the map B is multiplication by $U_w + U_{w'}$, as claimed.

We now compute the maps A and D , which count classes with $\text{gr}(x, y) = 0$. Equation (13-8) implies that $\mu(\phi) = 1$ and $m_2(\phi_0) = 0$. Let us fix a neck length along c_α . By the proof of Proposition 5.5, for sufficiently large neck length along c , there is an identification

$$(13-9) \quad \mathcal{M}(\phi \# \phi_0) \cong \{(u, u_0) \in \mathcal{M}(\phi) \times \mathcal{M}(\phi_0) : \rho^p(u) = \rho^{p_0}(u_0)\},$$

where p and p_0 denote the connected sum points on Σ and S^2 . Note that the required neck length along c for (13-9) to hold may depend on the fixed neck length along c_α .

By the proof of Proposition 5.5, we know that if $\theta \in \{\theta^+, \theta^-\}$ and $\mathbf{d} \in \text{Sym}^m([0, 1] \times \mathbb{R})$ is a fixed generic element, then

$$(13-10) \quad \sum_{\substack{\phi_0 \in \pi_2(x, x) \\ m_2(\phi_0) = 0 \\ m_1(\phi_0) = m}} \#\mathcal{M}(\phi_0, \mathbf{d}) \equiv 1 \pmod{2}.$$

Equations (13-9) and (13-10) together imply that if $\mu(\phi) = 1$ and $\theta \in \{\theta^+, \theta^-\}$, then

$$(13-11) \quad \#\widehat{\mathcal{M}}(\phi) \equiv \sum_{\substack{\phi_0 \in \pi_2(\theta, \theta) \\ m_2(\phi_0) = 0 \\ m_1(\phi_0) = m_1(\phi)}} \#\widehat{\mathcal{M}}(\phi \# \phi_0) \pmod{2}.$$

Note that (13-11) almost gives us the desired identification of A and D , except it does not inform us about the multiplicity on the basepoints w' and w . By considering the multiplicities in the four regions

around θ^+ , we know that if $\phi_0 \in \pi_2(\theta, \theta)$ has $m_2(\phi) = 0$, then

$$(13-12) \quad m_1(\phi_0) = n_1(\phi_0) + n_2(\phi_0).$$

The integer $m_1(\phi_0)$ is the power of U_z contributed by ϕ to $\partial_{\mathcal{H}_0}$, while $n_1(\phi_0)$ and $n_2(\phi_0)$ are the powers of U_w and $U_{w'}$, respectively, contributed by $\phi \# \phi_0$ to $\partial_{\mathcal{H}_{1.5}}$.

We finish our claim about the maps A and D , it is sufficient to show that if $\mathbf{d} \in \text{Sym}^m([0, 1] \times \mathbb{R})$ is a fixed point, then, for almost complex structure sufficiently stretched along c_α , the only classes $\phi_0 \in \pi_2(\theta, \theta)$ with nonempty $\mathcal{M}(\phi_0, \mathbf{d})$ have $n_2(\phi_0) = m_1(\phi_0)$ and $n_1(\phi_0) = 0$. This implies that any curve which makes nontrivial contribution to A or D is counted with a factor of $U_w^{m_1(\phi)}$ and no factor of $U_{w'}$. We prove this in the following subclaim:

Subclaim 13.3.1 *Suppose $\mathbf{d} \in \text{Sym}^m([0, 1] \times \mathbb{R})$ is not contained in the fat diagonal, and also does not contain any points of $\{0, 1\} \times \mathbb{R}$. Suppose $\phi_0 \in \pi_2(\theta, \theta)$ has $m_2(\phi_0) = 0$, $m_1(\phi_0) = m$ and $n_1(\phi_0) > 0$. If the almost complex structure on $S^2 \times [0, 1] \times \mathbb{R}$ is sufficiently stretched along c_α , then $\mathcal{M}(\phi_0, \mathbf{d})$ is empty.*

Proof The argument here is similar to our proof of Lemma 13.2. As we let the stretching parameter approach $+\infty$, we can extract broken limiting curves on $(S^2, \alpha_1, \beta_2^l)$ and the degenerate diagram (S^2, β_2^m) . See Figure 26. The curves β_2^l and β_2^m are obtained by cutting β_2 along c_α , and collapsing the endpoints.

Consider the limiting curves on (S^2, β_2^m) . Since these arose as the limit of curves which matched \mathbf{d} , the limiting curves must also match \mathbf{d} at the point p_0 (recall that p_0 corresponds to the circle c in Figure 26). There are no α curves on the degenerate diagram (S^2, β_2^m) . Consequently, any limiting curve either has no boundary, or has boundary which maps to β_2^m . Any curve which has boundary on β_2^m must map locally constantly to $\{0\} \times \mathbb{R} \subseteq [0, 1] \times \mathbb{R}$, by the maximum principle. Consequently, the curves which match \mathbf{d} can only be spheres which map to $S^2 \times \{d\} \subseteq S^2 \times [0, 1] \times \mathbb{R}$ for $d \in \mathbf{d}$, together with boundary degenerations which do not cover p_0 . Such a sphere contributes equally to m_1 and n_2 , while a boundary degeneration not covering p_0 only contributes to n_2 . Consequently, if $\mathcal{M}(\phi_0, \mathbf{d})$ is nonempty for a sufficiently stretched along complex structure, then

$$(13-13) \quad m_1(\phi_0) \leq n_2(\phi_0).$$

Combined with (13-12), since all multiplicities are nonnegative, inequality (13-13) implies that if $\mathcal{M}(\phi_0, \mathbf{d})$ is nonempty, then $m_1(\phi_0) = n_2(\phi_0)$ and $n_1(\phi_0) = 0$. \square

We now prove that C is zero (recall C counts representatives of classes in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$). For such classes, equation (13-8) constrains $m_2(\phi_0)$ to be in $\{0, 1\}$.

If $m_2(\phi_0) = 1$, then (13-8) implies that $\mu(\phi) = 0$, forcing ϕ to be a constant class, and ϕ_0 to be the bigon $m_2(\phi_0) = 1$. This class has a unique holomorphic representative. If $m_2(\phi_0) = 0$, then (13-8) implies that if the almost complex structure is sufficiently stretched along c , then $\mu(\phi) = 2$. It remains to count such classes and show that their total counts cancel the bigons with $m_2(\phi_0) = 1$.

We prove a helpful subclaim:

Subclaim 13.3.2 *Suppose that $d_i \in [0, 1] \times \mathbb{R}$ is a sequence of points approaching the line $\{0\} \times \mathbb{R}$, and suppose $\phi \in \pi_2(x, y)$ is a Maslov index 2 class on (Σ, α, β') . If the matched moduli space $\mathcal{M}(\phi, d_i)$ is nonempty for arbitrarily large i , then ϕ has domain equal to a connected component of $\Sigma \setminus \beta$.*

Proof Let u_i be a sequence of curves in $\mathcal{M}(\phi, d_i)$. Since p is contained in $\Sigma \setminus (\alpha \cup \beta)$, any curve u in the limit with $n_p(u) \neq 0$ must have $\pi_{[0,1] \times \mathbb{R}} \circ u$ constant, by the maximum principle. Consequently, u must have domain equal to $[\Sigma]$, or to a β boundary degeneration. Since such curves have Maslov index at least 2, and any other curves which appear will achieve transversality by Proposition 3.4, no additional curves may appear in the limit. \square

The only nonnegative class $\phi_0 \in \pi_2(\theta^+, \theta^-)$ with $m_1(\phi_0) = 1$ and $m_2(\phi_0) = 0$ is the bigon with $m_1(\phi_0) = 1$. Define the point

$$d(\phi_0) := (u \circ \pi_{[0,1] \times \mathbb{R}})((u \circ \pi_{S^2})^{-1}(p_0)) \in [0, 1] \times \mathbb{R},$$

where u is a representative of the bigon ϕ_0 (well defined up to \mathbb{R} -translation). Note that $d(\phi_0)$ depends on the choice of almost complex structure on $S^2 \times [0, 1] \times \mathbb{R}$. By stretching along c_α , we can make $d(\phi_0)$ arbitrarily close to $\{0\} \times \mathbb{R}$. Let us fix a neck length along c_α , so the conclusion of Subclaim 13.3.2 holds.

Having fixed a neck length along c_α , we now stretch along c . Suppose $\phi \# \phi_0$ admits a sequence of holomorphic representatives for a sequence of neck lengths along c which approach $+\infty$. We can extract broken limits. The curve appearing on $S^2 \times [0, 1] \times \mathbb{R}$ will be a representative of the bigon ϕ_0 . On $\Sigma \times [0, 1] \times \mathbb{R}$, there will be a curve in $\mathcal{M}(\phi', d(\phi_0))$ for some class ϕ' . We claim that ϕ' must actually be equal to the class ϕ . Using Proposition 3.4, it is straightforward to show that any curves appearing in the limit either achieve transversality, or have the homology class of a boundary degeneration or have domain $[\Sigma]$. Classes which have domain $[\Sigma]$ or which have the domain of a boundary degeneration have Maslov index at least 2 by (4-5). Hence, other curves are prohibited from appearing in the limit by a dimension count, since $\mu(\phi) = 2$. Consequently, $\phi = \phi'$.

Subclaim 13.3.2 implies that since $\mathcal{M}(\phi, d(\phi_0))$ is nonempty, ϕ must have domain equal to a connected component of $\Sigma \setminus \beta$.

Having restricted the homology classes which can contribute to C , we now count their holomorphic representatives. Let $\phi \# \phi_0 \in \pi_2(x \times \theta^+, x \times \theta^-)$ denote an index 1 homology class with $m_1(\phi_0) = 1$ and $m_2(\phi_0) = 0$ such that the domain of ϕ is a connected component of $\Sigma \setminus \beta$.

Let ϕ' denote the Maslov index 1 bigon going over w' . Splicing $\phi \# \phi_0$ and ϕ' together, we get the Maslov index 2 homology class $B'_{x \times \theta^+} = \phi' * (\phi \# \phi_0) \in \pi_2(x \times \theta^+, x \times \theta^+)$. The moduli space $\widehat{\mathcal{M}}(B'_{x \times \theta^+})$ is 1-dimensional. We count its ends. One end corresponds to the boundary degenerations in $\widehat{N}^\beta(B'_{x \times \theta^+})$, which have total count equal to 1 by [Ozsváth and Szabó 2008, Theorem 5.5]. The other ends correspond to representatives of $B'_{x \times \theta^+}$ breaking into two holomorphic strips. Our previous argument implies that the only other end of $\widehat{\mathcal{M}}(B)$ consists of the broken holomorphic strip consisting of a representative of $\phi \# \phi_0$ and a representative of the bigon ϕ' .




Figure 27: Multiplicities on \mathcal{H}_1 .

Summing over the ends of $\widehat{\mathcal{M}}(B'_{\mathbf{x} \times \theta^+})$, we conclude

$$\#\widehat{\mathcal{N}}^\beta(B'_{\mathbf{x} \times \theta^+}) + \#\widehat{\mathcal{M}}(\phi \# \phi_0) \cdot \#\widehat{\mathcal{M}}(\phi') = 0.$$

Since $\#\widehat{\mathcal{N}}^\beta(B'_{\mathbf{x} \times \theta^+}) = \#\widehat{\mathcal{M}}(\phi') = 1$, we conclude that $\#\widehat{\mathcal{M}}(\phi \# \phi_0) = 1$.

It follows that $C = 0$, completing the proof of Lemma 13.3. □

In a similar manner to Lemma 13.3, we may compute the differentials on \mathcal{H}_1 and \mathcal{H}_2 :

Lemma 13.4 *The differentials $\partial_{\mathcal{H}_1, J_\alpha}$ and $\partial_{\mathcal{H}_1, J_\beta}$ have the same form:*

$$\partial_{\mathcal{H}_1, J_\alpha} = \partial_{\mathcal{H}_1, J_\beta} = \begin{pmatrix} (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_w} & U_w + U_{w'} \\ 0 & (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_w} \end{pmatrix}.$$

Furthermore, using the multiplicities from Figure 27, the index 1 J_α -holomorphic curves counted by $\partial_{\mathcal{H}_1, J_\alpha}$ satisfy the following:

- (1) Any Maslov index 1 class in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^+)$ or $\pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^-)$ with holomorphic representatives has $n_1 = m_2 = 0$ and $n_2 = m_1$.
- (2) The index 1 classes in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$ with holomorphic representatives have domain equal to one of the two bigons with $n_1 = 1$ or $n_2 = 1$ (and all other multiplicities zero).
- (3) The index 1 classes in $\pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^+)$ with holomorphic representatives have one of two domains. One domain is a bigon on (S^2, α_1, β_1) with $m_2 = 1$ (and all other multiplicities zero). The other domain is the connected sum of the bigon on (S^2, α_1, β_1) with $m_1 = 1$, together with the domain on (Σ, α, β) consisting of the connected component of $\Sigma \setminus \beta$ which contains the connected sum point. Any class with either domain has one representative, modulo 2.

Similar statements hold for $\partial_{\mathcal{H}_1, J_\beta}$, $\partial_{\mathcal{H}_2, J_\alpha}$ and $\partial_{\mathcal{H}_2, J_\beta}$.

Proof The differential $\partial_{\mathcal{H}_1, J_\alpha}$ counts the same holomorphic curves as $\partial_{\mathcal{H}_{1.5}, J_\alpha}$, which is analyzed in Lemma 13.3. The present claim is proven by repeating the argument therein, while keeping track of the multiplicities over the basepoints, as they appear in \mathcal{H}_1 . □

We now compute the change of almost complex structure map $\Psi_{J_\alpha \rightarrow J_\beta}$ on $\text{CF}^-(\mathcal{H}_{1.5}, \mathfrak{s})$.

Lemma 13.5 Consider the almost complex structures J_α and J_β on $\mathcal{H}_{1,5}$, obtained by stretching along c and c_α , or c and c_β , respectively. Whenever the necks along c are sufficiently large,

$$(13-14) \quad \Psi_{J_\alpha \rightarrow J_\beta} = \begin{pmatrix} \text{id} & 0 \\ * & \text{id} \end{pmatrix}.$$

If all necks are sufficiently long, then the $*$ -component may be identified with

$$* = \sum_{i,j \geq 0} U_w^i U_{w'}^j (\partial_{i+j+1})_{U_w, U_{w'}}.$$

Proof Lemma 13.2 implies that the relative lengths along c and c_α for J_α do not affect the transition map, and similarly the relative lengths along c and c_β for J_β do not affect the transition map.

We fix a neck length along c_α for J_α , and a neck length along c_β for J_β . The computation of the three components not marked with a $*$ follows from a small adaptation to Proposition 5.3.

To compute the entry marked with a $*$, we stretch along all necks so that the differentials $\partial_{\mathcal{H}_{1.5, J_\alpha}}$ and $\partial_{\mathcal{H}_{1.5, J_\beta}}$ take the form described in Lemma 13.3. Write C for the entry labeled $*$ in (13-14). The transition map $\Psi_{J_\alpha \rightarrow J_\beta}$ is a chain map. We view the relation

$$\Psi_{J_\alpha \rightarrow J_\beta} \circ \partial_{\mathcal{H}_{1.5, J_\alpha}} + \partial_{\mathcal{H}_{1.5, J_\beta}} \circ \Psi_{J_\alpha \rightarrow J_\beta} = 0$$

as a matrix involving two-by-two matrices. The diagonal entries give

$$(\partial_{\mathcal{H}_0})^{U_z \rightarrow U_w} + (\partial_{\mathcal{H}_0})^{U_z \rightarrow U_{w'}} = (U_w + U_{w'}) \cdot C,$$

which algebraically implies the stated form of the map C . □

We now perform a triangle map computation for a triple which has been stabilized as in Figure 28.

Proposition 13.6 Suppose that $\mathcal{T} = (\Sigma, \alpha, \beta, \beta', \mathbf{w}_0 \cup \{z\})$ is a Heegaard triple with a distinguished basepoint z , and β' are small isotopies of β , satisfying $|\beta_i \cap \beta'_j| = 2\delta_{ij}$. Let

$$\widehat{\mathcal{T}} = (\Sigma, \alpha \cup \{\alpha_1\}, \beta \cup \{\beta_1\}, \beta' \cup \{\beta_2\}, \mathbf{w}_0 \cup \{w, w'\})$$

be the Heegaard triple obtained by replacing a neighborhood of z with the region shown in Figure 28. Let J be an almost complex structure on \mathcal{T} , and $J(T)$ an almost complex structure on $\widehat{\mathcal{T}}$ which is stretched along c_α . Using matrix notation, we have

$$F_{\widehat{\mathcal{T}}, J(T)}(\mathbf{x} \times -, \Theta_{\beta, \beta'}^+ \times \theta^+) = \begin{pmatrix} F_{\mathcal{T}, J}(\mathbf{x}, \Theta_{\beta, \beta'}^+)^{U_z \rightarrow U_w} & 0 \\ 0 & F_{\mathcal{T}, J}(\mathbf{x}, \Theta_{\beta, \beta'}^+)^{U_z \rightarrow U_w} \end{pmatrix}$$

for sufficiently large T .

The following lemma will be useful in the proof of Proposition 13.6:

Lemma 13.7 Suppose that $(\Sigma, \beta, \beta', \mathbf{w})$ is a diagram for $(S^1 \times S^2)^{\#g(\Sigma)}$ such that β' are small isotopies of β , and $|\beta_i \cap \beta'_j| = 2\delta_{ij}$. Let $\Theta_{\beta, \beta'}^+$ denote the top-graded intersection point of $\mathbb{T}_\beta \cap \mathbb{T}_{\beta'}$. If




Figure 28: The triple $\widehat{\mathcal{T}}$ in Proposition 13.6, a stabilization of \mathcal{T} . The shaded regions are two examples of small triangles, which might be counted.

$\phi \in \pi_2(\Theta_{\beta, \beta'}^+, \mathbf{y})$ is a nonnegative homology class, then $\mu(\phi) - n_{w_0}(\phi) \geq 0$ for all $w_0 \in \mathbf{w}$. Furthermore, $\mu(\phi) - n_{w_0}(\phi) = 0$ for some $w_0 \in \mathbf{w}$ if and only if ϕ is the constant homology class $e_{\Theta_{\beta, \beta'}^+}$.

Proof By the formula for the relative Maslov grading in [Ozsváth and Szabó 2004b], we have

$$(13-15) \quad \text{gr}(\Theta_{\beta, \beta'}^+, \mathbf{y}) = \mu(\phi) - 2 \sum_{w \in \mathbf{w}} n_w(\phi) \geq 0.$$

Hence, if ϕ is nonnegative,

$$(13-16) \quad \mu(\phi) - n_{w_0}(\phi) = \text{gr}(\Theta_{\beta, \beta'}^+, \mathbf{y}) + n_{w_0}(\phi) + 2 \sum_{w \in \mathbf{w} \setminus \{w_0\}} n_w(\phi) \geq n_{w_0}(\phi) + 2 \sum_{w \in \mathbf{w} \setminus \{w_0\}} n_w(\phi) \geq 0.$$

In particular, $\mu(\phi) - n_{w_0}(\phi) \geq 0$. Furthermore, equation (13-16) implies that if $\mu(\phi) - n_{w_0}(\phi) = 0$ for some w_0 , then $n_w(\phi) = 0$ for all $w \in \mathbf{w}$, and $\mu(\phi) = 0$. From (13-15), we see that $\text{gr}(\Theta_{\beta, \beta'}^+, \mathbf{y}) = 0$, so $\mathbf{y} = \Theta_{\beta, \beta'}^+$. By diagrammatic inspection, the only nonnegative class $\phi \in \pi_2(\Theta_{\beta, \beta'}^+, \Theta_{\beta, \beta'}^+)$ with $\mu(\phi) = 0$ and $n_w(\phi) = 0$ for all w is the constant class. \square

Proof of Proposition 13.6 We begin with a Maslov index computation:

Subclaim 13.7.1 Suppose $x \in \{x^+, x^-\}$ and $y \in \{y^+, y^-\}$, and $\psi_0 \in \pi_2(x, \theta^+, y)$ is a triangle class on the triple $(S^2, \alpha_1, \beta_1, \beta_2)$ shown in Figure 28. We claim

$$(13-17) \quad \mu(\psi_0) = (m_1 + m_2 + N_1 + N_2)(\psi_0).$$

Proof Equation (13-17) holds for both of the two small triangle classes shaded in Figure 28. Furthermore, equation (13-17) respects splicing in doubly periodic domains on (S^2, β_1, β_2) as well as bigons on (S^2, α_1, β_1) and (S^2, α_1, β_2) , implying the formula in general. \square




Figure 29: Decomposing $\widehat{\mathcal{T}}$ along c_α and the broken curves $\mathcal{U}_0, \mathcal{U}_m$ and \mathcal{U} .

Suppose $\psi \in \pi_2(x, \Theta_{\beta, \beta'}^+, y)$ is a class on \mathcal{T} , and $\psi_0 \in \pi_2(x, \theta^+, y)$ is a class on $(S^2, \alpha_1, \beta_1, \beta_2)$. We form the class $\psi \# \psi_0$ on $\widehat{\mathcal{T}}$ by taking the connected sum of ψ and ψ_0 , along c . We claim

$$(13-18) \quad \mu(\psi \# \psi_0) = \mu(\psi) + \mu(\psi_0) - 2m_1(\psi_0) = \mu(\psi) + (m_2 - m_1 + N_1 + N_2)(\psi_0).$$

The first equality of (13-18) follows from Sarkar’s formula [2011] for the Maslov index and the fact that a disk has Euler measure 1. The second equality follows from (13-17).

Let $p \in \Sigma$ and $p_0 \in S^2$ denote the points corresponding to c_α , which arise after we cut Σ along c_α , and collapse each of the resulting boundary components. Write

$$\Sigma_0 = \Sigma \setminus \{p\} \quad \text{and} \quad S_0^2 = S^2 \setminus \{p_0\}.$$

Let β_1^r and β_2^r denote the resulting arcs on Σ_0 .

We will also be interested in the tube region, for which we write $(S^1 \times \mathbb{R}, \beta_1^m, \beta_2^m)$. We write p_m^l and p_m^r for the two punctures of $S^1 \times \mathbb{R}$. See Figure 29 for a schematic.

If T_i is a sequence of neck lengths approaching ∞ , and $J(T_i)$ denotes an almost complex structure on $\Sigma \times \Delta$ with neck length T_i along c_α , then a sequence u_i of $J(T_i)$ -holomorphic triangles representing $\psi \# \psi_0$ has a subsequence which converges to three collections, $\mathcal{U}, \mathcal{U}_m$ and \mathcal{U}_0 , where \mathcal{U} is a broken representative of ψ on $(\Sigma_0, \alpha, \beta \cup \{\beta_1^r\}, \beta^r \cup \{\beta_2^r\})$, \mathcal{U}_0 is a broken representative of ψ_0 on $(S_0^2, \alpha_1, \beta_1^l, \beta_2^l)$, and \mathcal{U}_m is a collection of holomorphic curves in the tube region $(S^1 \times \mathbb{R}, \beta_1^m, \beta_2^m)$. The collections $\mathcal{U}, \mathcal{U}_0$ and \mathcal{U}_m may contain both holomorphic curves mapping into $\Sigma \times \Delta, S^2 \times \Delta$ or $S^1 \times \mathbb{R} \times \Delta$, and holomorphic curves mapping into $\Sigma \times [0, 1] \times \mathbb{R}, S^2 \times [0, 1] \times \mathbb{R}$ or $S^1 \times \mathbb{R} \times [0, 1] \times \mathbb{R}$.




Figure 30: Domains of troublesome curves which have asymptotics of type (a2) and (a3). We rule out such curves from appearing.

We now describe the possible asymptotic behavior of the curves appearing in \mathcal{U} , \mathcal{U}_m and \mathcal{U}_0 . For definiteness, let us focus on \mathcal{U} . View a neighborhood of the puncture on Σ_0 as $S^1 \times [0, \infty)$. If $u \in \mathcal{U}$, then the source of u may have punctures along the boundary, at which u is asymptotic to an intersection point on the Heegaard diagram. Additionally, a curve u in \mathcal{U} may have the following types of asymptotics:

- (a1) The source of u may have an interior puncture which is asymptotic to an orbit $S^1 \times \{p\} \subseteq S^1 \times \Delta$ for $p \in \Delta$ (multiple covers of such orbits are allowed). Similar asymptotics may occur for curves of \mathcal{U} mapping into $\Sigma \times [0, 1] \times \mathbb{R}$.
- (a2) The source u may have a boundary puncture which is asymptotic to a chord $a \times \{p\} \subseteq S^1 \times \Delta$, where p is a point in one of the three components of $\partial\Delta$, and a is a subarc of S^1 (thought of as the boundary of Σ with a point removed). The arc a connects β_1^r to β_1^r or β_2^r to β_2^r . The arc a may wind multiple times around S^1 . Similar asymptotics could also appear in curves in \mathcal{U} which map into $\Sigma \times [0, 1] \times \mathbb{R}$. (These asymptotics are studied in [Lipshitz et al. 2018].)
- (a3) The source of u may have a boundary puncture which is asymptotic to a chord which connects β_1^r and β_2^r . At this boundary puncture, $\pi_\Delta \circ u$ approaches ∞ in one of the cylindrical ends of Δ . Also, $\pi_\Sigma \circ u$ is asymptotic to an arc which connects β_1^r and β_2^r (perhaps winding many times). Similar asymptotics could also appear in curves in \mathcal{U} which map into $\Sigma \times [0, 1] \times \mathbb{R}$.

Examples of asymptotics (a2) and (a3) are shown in Figure 30. We will rule out asymptotics of type (a2) and (a3) from appearing generically in \mathcal{U} , \mathcal{U}_m or \mathcal{U}_0 .

We first consider \mathcal{U} , the curves on Σ_0 :

Subclaim 13.7.2 *If $u \in \mathcal{U}$, and S_0 is a connected component of the source of u , then $u|_{S_0}$ cannot have boundary on both α and β_1^r or β_2^r .*

Proof The claim follows from the maximum modulus principle (compare the proof of [Manolescu and Ozsváth 2025, Proposition 5.2]). Suppose $u: S \rightarrow \Sigma_0 \times \Delta$ is a holomorphic triangle with boundary on α and at least one of β_1^r or β_2^r . Let \bar{S} denote the surface obtained by compactifying S by adding its boundary punctures. There must be a connected component of $\partial\bar{S}$ which is mapped by $\pi_\Sigma \circ u$ to only β_1^r and β_2^r , since there are no intersection points between β_1^r or β_2^r and α , β or β' . By the maximum

modulus principle, $\pi_\Delta \circ u$ must be constant, which is a contradiction. A similar argument applies if u is a holomorphic strip on one of the subdiagrams $(\Sigma_0, \alpha, \beta \cup \{\beta_1^r\})$ or $(\Sigma_0, \alpha, \beta' \cup \{\beta_2^r\})$. \square

It follows from Subclaim 13.7.2 that \mathcal{U} may be arranged into the following collections:

- (1) a broken holomorphic triangle u_Σ on $(\Sigma, \alpha, \beta, \beta')$ representing a class $\psi_\Sigma \in \pi_2(x, \Theta', y)$;
- (2) a broken holomorphic strip $u_{\beta, \beta'}$ on $(\Sigma, \beta \cup \{\beta_1^r\}, \beta' \cup \{\beta_2^r\})$ representing a class $\phi_{\beta, \beta'}$ in $\pi_2(\Theta_{\beta, \beta'}^+ \times p, \Theta' \times p)$.

Subclaim 13.7.3 *The class $\phi_{\beta, \beta'}$ must be the constant class, and \mathcal{U} consists only of the curve u_Σ (and potentially some ghost curves). The class ψ_Σ is equal to ψ , and has Maslov index 0. The asymptotics of u_Σ at p consist only of orbits of type (a1). After perturbing the almost complex structure, the orbit asymptotics of u_Σ project to $m_1(\psi)$ distinct points in the interior of Δ . After completing over the puncture p , the curve u_Σ satisfies (M1)–(M6).*

Proof The proof is essentially combinatorial, and is based on obtaining convenient index formulas.

Let $\phi_{\beta, \beta'}^0$ denote the homology class in $\pi_2(\Theta_{\beta, \beta'}^+, \Theta')$ induced by $\phi_{\beta, \beta'}$ on (Σ, β, β') , obtained by removing a disk containing the β_1^r and β_2^r curves, and collapsing the resulting boundary component to a point.

Note that $\psi = \psi_\Sigma + \phi_{\beta, \beta'}^0$, since all multiplicities are represented. Hence,

$$(13-19) \quad \mu(\psi) = \mu(\psi_\Sigma) + \mu(\phi_{\beta, \beta'}^0).$$

We claim

$$(13-20) \quad (m_2 - m_1)(\psi_0) = (m_2 - m_1)(\psi_\Sigma) + (m_2 - m_1)(\phi_{\beta, \beta'}) = (m_2 - m_1)(\phi_{\beta, \beta'}).$$

The first equality of (13-20) follows since ψ_0 represents the entire homology class of the limiting curves on $(S_0^2, \alpha_1, \beta_1^l, \beta_2^l)$, while ψ_Σ and $\phi_{\beta, \beta'}$ represent the entire homology class of the limiting curves on $(\Sigma_0, \alpha, \beta \cup \{\beta_1^r\}, \beta' \cup \{\beta_2^r\})$. The second equality of (13-20) follows from the fact that $m_2(\psi_\Sigma) = m_1(\psi_\Sigma)$, since u_Σ has no boundary mapping to β_1 or β_2 .

Combining (13-18), (13-19) and (13-20), we obtain

$$(13-21) \quad \begin{aligned} \mu(\psi \# \psi_0) &= \mu(\psi) + (m_2 - m_1 + N_1 + N_2)(\psi_0) \\ &= \mu(\psi_\Sigma) + \mu(\phi_{\beta, \beta'}^0) + (m_2 - m_1 + N_1 + N_2)(\psi_0) \\ &= \mu(\psi_\Sigma) + (\mu(\phi_{\beta, \beta'}^0) - m_1(\phi_{\beta, \beta'})) + m_2(\phi_{\beta, \beta'}) + (N_1 + N_2)(\psi_0). \end{aligned}$$

By construction of the class $\phi_{\beta, \beta'}^0$, we have

$$(13-22) \quad m_1(\phi_{\beta, \beta'}) = m_1(\phi_{\beta, \beta'}^0).$$

Equation (13-22) and Lemma 13.7 imply

$$\mu(\phi_{\beta, \beta'}^0) - m_1(\phi_{\beta, \beta'}) = \mu(\phi_{\beta, \beta'}^0) - m_1(\phi_{\beta, \beta'}^0) \geq 0.$$

Consequently, equation (13-21) implies

$$(13-23) \quad 0 = \mu(\psi_\Sigma) = \mu(\phi_{\beta, \beta'}^0) - m_1(\phi_{\beta, \beta'}) = m_2(\phi_{\beta, \beta'}) = N_1(\psi_0) = N_2(\psi_0).$$

In particular, equation (13-23) implies that

$$\mu(\phi_{\beta, \beta'}^0) - m_1(\phi_{\beta, \beta'}) = \mu(\phi_{\beta, \beta'}^0) - m_1(\phi_{\beta, \beta'}) = 0,$$

so Lemma 13.7 implies that $\phi_{\beta, \beta'}^0$ must be a constant class.

Next, if $n \in \{n_1, n_2, m_1, m_2\}$, we note that

$$n(\phi_{\beta, \beta'}) + m_1(\psi_\Sigma) = n(\psi_0),$$

since $\phi_{\beta, \beta'}$, ψ_0 and ψ_Σ , when spliced together, represent the entire class $\psi \# \psi_0$. Since u_Σ has no boundary components on β_1^r or β_2^r ,

$$n_1(\psi_\Sigma) = n_2(\psi_\Sigma) = m_1(\psi_\Sigma) = m_2(\psi_\Sigma).$$

Consequently, the class $\phi_{\beta, \beta'}$ satisfies the same vertex relation at θ^- that the class ψ_0 satisfies:

$$(13-24) \quad (n_1 + n_2)(\phi_{\beta, \beta'}) = (m_1 + m_2)(\phi_{\beta, \beta'}).$$

Since $\phi_{\beta, \beta'}^0$ is the constant class, we know $m_1(\phi_{\beta, \beta'}^0) = 0$. Equation (13-22) implies that $m_1(\phi_{\beta, \beta'}) = 0$. From (13-23), we know that $m_2(\phi_{\beta, \beta'}) = 0$. Combined with (13-24), we conclude

$$n_1(\phi_{\beta, \beta'}) = n_2(\phi_{\beta, \beta'}) = m_1(\phi_{\beta, \beta'}) = m_2(\phi_{\beta, \beta'}) = 0,$$

so $\phi_{\beta, \beta'}$ is a constant class, as well.

Since $\phi_{\beta, \beta'}$ is the constant class, $\psi = \psi_\Sigma$.

Since u_Σ has no boundary mapping to β_1^r or β_2^r , it can only be asymptotic at p to orbits of type (a1). From (13-23) we know that $\mu(\psi) = 0$. Hence, for generic almost complex structure, after completing over p , u_Σ will satisfy (M1)–(M6).

Finally, by perturbing the almost complex structure slightly, we can ensure that the asymptotics of u_Σ at p consist of $m_1(\psi)$ once-covered orbits of type (a1), which each project to a different point of Δ . Such a perturbation can be done concretely, by perturbing the placement of the connected sum point so that it is disjoint from the image of the branch set of $\pi_\Sigma \circ v$, for any Maslov index 0 holomorphic triangle v (a codimension 2 subset of Σ). \square

We now analyze the curves in \mathcal{U}_0 and \mathcal{U}_m :

Subclaim 13.7.4 *The curves in \mathcal{U}_0 can only be asymptotic at the puncture p_0 to an orbit of type (a1). The curves in \mathcal{U}_m consist only of once-covered cylinders of the form $S^1 \times \mathbb{R} \times \{d\}$, where d is an interior point of Δ , as well as possibly some ghost curves.*

Proof A version of compactness in symplectic field theory is described in [Bourgeois et al. 2003], where a linear ordering of levels appears. Our present situation is more similar to the compactification via

holomorphic combs which appears in bordered Floer homology [Lipshitz et al. 2018, Section 5.4], since holomorphic curves can degenerate into the three cylindrical ends of Δ , and also into the tube region which results from cutting along c_α . Hence, the limiting curves may be arranged into a 2-component level structure. We refer to one component as the Δ -level, and the other as the Σ -level (compare [Lipshitz et al. 2018, Definition 5.20]). Furthermore, there is a single Δ -level in each Σ -level which consists of curves mapping into $S^2_0 \times \Delta$, $S^1 \times \mathbb{R} \times \Delta$ or $\Sigma_0 \times \Delta$. We refer to this level as the *central* Δ -level. All other curves map into one of the cylindrical 4-manifolds appearing in the ends.

We begin at the central Δ -level of \mathcal{U} , which consists only of the curve u_Σ by Subclaim 13.7.3 (and possibly some ghost curves, which we will later rule out). We proved that the asymptotics of u_Σ consisted of $m_1(\psi)$ orbits of type (a1), each of which projects to a distinct point in the interior of Δ .

There is a story \mathcal{V}_1 of \mathcal{U}_m which matches u_Σ . Since \mathcal{V}_1 matches $m_1(\psi)$ once-covered orbits of type (a1), which each project to an interior point of Δ , \mathcal{V}_1 must also be in the central Δ -level. Since there are no α curves on $(S^1 \times \mathbb{R}, \beta_1^m, \beta_2^m)$, any holomorphic curve with connected source which has such an orbit at p_m^r must project constantly to Δ . Consequently, such a curve must also be asymptotic to a once-covered orbit at p_m^l . According to Subclaim 13.7.3, we have

$$n_1(\psi_0) = n_2(\psi_0) = m_1(\psi_0) = m_2(\psi_0) = m_1(\psi).$$

Since \mathcal{V}_1 must contain $m_1(\psi)$ once-covered cylinders, which each project constantly to Δ , there can be no other holomorphic curves with nonconstant image, since they would raise the multiplicity in some region of Σ too high.

We now consider the level \mathcal{V}_2 of \mathcal{U}_m which matches the asymptotics of \mathcal{V}_1 . The previous argument implies that \mathcal{V}_2 consists only of $m_1(\psi)$ once-covered cylinders, and some ghost curves. We continue in this manner until we reach the central Δ -level of \mathcal{U}_0 . Since adjacent levels must have matching asymptotics, we conclude that the central Δ -level of \mathcal{U}_0 has asymptotics at p_0 consisting only of $m_1(\psi_0)$ once-covered orbits. Furthermore, the asymptotics match those of u_Σ . □

Subclaim 13.7.5 *The collection \mathcal{U}_0 consists of a single holomorphic triangle u_0 , satisfying (M1)–(M6), as well as possibly some ghost curves. Furthermore,*

$$\rho^p(u_\Sigma) = \rho^{p_0}(u_0).$$

Proof Subclaim 13.7.5 constrains the asymptotics at p_0 of the curves in \mathcal{U}_0 to satisfy (a1). The Maslov index of ψ_0 is $2m_1(\psi_0)$, since $\mu(\psi) = 0$ by the proof of Subclaim 13.7.3. The set $X(\psi) \subseteq \text{Sym}^{m_1(\psi_0)}(\Delta)$ defined by

$$X(\psi) := \{\rho^p(u) : u \in \mathcal{M}(\psi)\}$$

is finite. Consequently, as in the proof of Proposition 5.5, dimension counting using Proposition 3.4 implies that \mathcal{U}_0 consists of a single holomorphic triangle u_0 which satisfies (M1)–(M6). □

Having constrained the curves of \mathcal{U} , \mathcal{U}_m and \mathcal{U}_0 in Subclaims 13.7.3, 13.7.4 and 13.7.5, the index argument in the proof of Proposition 5.5 applies to show that, generically, no ghost curves appear.

The class ψ_0 is completely determined: since $N_1(\psi_0) = N_2(\psi_0) = 0$ by (13-23), ψ_0 has domain consisting of one of the two shaded small triangles in Figure 28 together with k copies of the component of $S^2 \setminus \alpha_0$ which has nonzero multiplicity on the region marked m_1 . In particular, this implies that $\psi_0 \in \pi_2(x^+, \theta^+, y^+)$ or $\psi_0 \in \pi_2(x^-, \theta^+, y^-)$, and hence the off-diagonal entries of the map $F_{\widehat{\mathcal{T}}, J(T)}(\mathbf{x} \times -, \Theta_{\beta, \beta'}^+ \times \theta^+)$ are zero. Let us write

$$\psi_0^{k,+} \in \pi_2(x^+, \theta^+, y^+) \quad \text{and} \quad \psi_0^{k,-} \in \pi_2(x^-, \theta^+, y^+)$$

for these two triangle classes.

It remains to count representatives of the class $\psi \# \psi_0$, when $\mu(\psi) = 0$ and $\psi_0 \in \{\psi_0^{m_1(\psi),+}, \psi_0^{m_1(\psi),-}\}$. If u_Σ denotes the holomorphic triangle in \mathcal{U} , and u_0 the triangle in \mathcal{U}_0 , then we can view the pair (u_Σ, u_0) as a point in the compactification of the space

$$\bigcup_{T>0} \mathcal{M}_{J(T)}(\psi \# \psi_0).$$

Gluing gives an identification of a neighborhood of the set of such pairs (u_Σ, u_0) with the Cartesian product

$$(13-25) \quad \{(u_\Sigma, u_0) \in \mathcal{M}(\psi) \times \mathcal{M}^{p_0}(\psi_0) : \rho^p(u_\Sigma) = \rho^{p_0}(u_0)\} \times [0, 1],$$

where $\mathcal{M}^{p_0}(\psi_0)$ denotes the set of holomorphic curves $u_0: S_0 \rightarrow S_0^2 \times [0, 1] \times \mathbb{R}$, representing the class ψ_0 on $(S_0^2, \alpha_1, \beta_1^l, \beta_2^l)$, such that $(\pi_{S^2} \circ u_0)^{-1}(p_0) \cap \partial S = \emptyset$ (i.e., the curves which only have asymptotics of type (a1) at p_0 , and none of type (a2) or (a3)).

Since $\mu(\psi) = 0$, it suffices to count the number of elements in the matched moduli space $\mathcal{M}^{p_0}(\psi_0, \mathbf{d})$, where $\mathbf{d} \in \text{Sym}^{m_1(\psi_0)}(\Delta)$ is a generic point. The strategy used by Ozsváth and Szabó [2008, Lemma 6.4] to prove (5-15) now readily adapts to our present situation. Pick a path of points $\mathbf{d}_t: [0, \infty) \rightarrow \text{Sym}^{m_1(\psi_0)}(\Delta)$, disjoint from the fat diagonal, consisting of $m_1(\psi_0)$ points which all travel into the α_1 - β_1 cylindrical end of Δ , such that the points of \mathbf{d}_t are spaced at least t distance apart (with respect to a metric obtained by embedding Δ conformally in the complex plane, so that each cylindrical end is identified with $[0, 1] \times [0, \infty)$).

Subclaim 13.7.6 *The only ends of the 1-dimensional moduli space $\bigcup_{t \in [0, \infty)} \mathcal{M}^{p_0}(\psi_0, \mathbf{d}_t)$ at finite t correspond to $\mathcal{M}^{p_0}(\psi_0, \mathbf{d}_0)$.*

Proof The proof is similar to our proof of (5-23). Degenerations may be analyzed by considering possible curves and arcs collapsing in the source curve. Such degenerations may be broadly classified into one of the following phenomena: holomorphic strips breaking off into one of the three cylindrical ends, boundary degenerations bubbling off, curves breaking off toward the puncture p_0 , or the source curve becoming nodal.

One example of a curve breaking off to p_0 would be a slit along β_1^l or β_2^l traveling out towards p_0 .

Strip breaking is prohibited as follows. Suppose $v: S \rightarrow S^2 \times [0, 1] \times \mathbb{R}$ is a hypothetical strip appearing in degeneration in $\mathcal{M}^{p_0}(\psi_0, \mathbf{d}_t)$ at some finite t . Since $N_1(v) = N_2(v) = 0$, the domain of v must

have nonzero multiplicity in one of the four regions adjacent to p_0 . Consequently, any holomorphic triangle also arising in the degeneration could not match any \mathbf{d}_t , since \mathbf{d}_t is bounded away from the three components of $\partial\Delta$.

Boundary degenerations are prohibited similarly, since the only possible class of the resulting boundary degeneration would have nonzero multiplicity around p_0 , prohibiting the triangular component in the limit from matching \mathbf{d}_t .

Curves breaking off in the direction of p_0 are also prohibited, since the resulting curve mapping into $S^1 \times \mathbb{R} \times \Delta$ must match \mathbf{d}_t at the puncture p_m^r . As argued in the proof of Subclaim 13.7.3, this constrains any curves appearing in $S^1 \times \mathbb{R} \times \Delta$ to consist only of cylinders, and possibly ghost curves. Ghost curves are prohibited from dimension counts, as in the proof of Proposition 5.5.

Finally, the formation of nodal singularities is prohibited similarly to the proof of (5-23). Double points appearing on the interior of the source are prohibited by dimension counts. Boundary double points are prohibited since they result in the formation of a boundary degeneration, which we have already prohibited. □

Consequently, the limiting curves which appear in the ends of $\mathcal{M}^{p_0}(\psi_0, \mathbf{d}_t)$ as $t \rightarrow \infty$ take the following form:

- (1) a single index 0 holomorphic triangle (with domain equal to one of the two shaded regions in Figure 28), together with
- (2) $m_1(\psi_0)$ Maslov index 2 holomorphic curves with domain A equal to the component of $S^2 \setminus \alpha_0$ with $m_1(A) = 1$; each curve matches a single point $d \in [0, 1] \times \mathbb{R}$.

If $x \in \{x^+, x^-\}$, write A_x for the Maslov index 2 homology class in $\pi_2(x, x)$ corresponding to the curves of type (2). It remains to count the matched moduli space $\mathcal{M}^{p_0}(A_x, d)$ for a generic point $d \in [0, 1] \times \mathbb{R}$. This is achieved by letting d approach $\{1\} \times \mathbb{R}$ (the α_1 boundary of $[0, 1] \times \mathbb{R}$). The limiting curve then consists of a single α_1 boundary degeneration, as well as a representative of the constant class, e_x . There are no other ends.

By [Ozsváth and Szabó 2008, Theorem 5.5], the count of α_1 boundary degenerations in the class A_x is 1, modulo 2. By the Riemann mapping theorem, each of the shaded triangles in Figure 28 has a unique representative. Consequently, gluing together these curves, we conclude that $\mathcal{M}^{p_0}(\psi_0, \mathbf{d})$ has a single element, modulo 2, for any generic $\mathbf{d} \in \text{Sym}^{m_1(\psi_0)}(\Delta)$.

Using (13-25), we conclude that, for large T ,

$$\#\mathcal{M}_J(\psi) = \#\mathcal{M}_{J(T)}(\psi \# \psi_0).$$

The class $\psi \# \psi_0$ has multiplicity $m_1(\psi)$ over w , and zero multiplicity on w' . Hence, any representative of $\psi \# \psi_0$ is counted with a factor of $U_{w'}^{n_z(\psi)}$ and no factor of U_w , completing the proof of Proposition 13.6. □

Proposition 13.8 *Suppose that $\mathcal{T} = (\Sigma, \alpha', \alpha, \beta', w_0 \cup \{z\})$ is a Heegaard triple with a distinguished basepoint z , and α' are small Hamiltonian isotopies of α satisfying $|\alpha'_i \cap \alpha_j| = 2\delta_{ij}$. Let $\widehat{\mathcal{T}} =$*




Figure 31: The Heegaard triple $\widehat{\mathcal{T}}$ in Proposition 13.8. The shaded regions are each examples of small triangles which might be counted.

$(\Sigma, \alpha' \cup \{\alpha_2\}, \alpha \cup \{\alpha_1\}, \beta' \cup \{\beta_2\}, \mathbf{w}_0 \cup \{w, w'\})$ be the Heegaard triple shown in Figure 31. Write J for an almost complex structure on $\Sigma \times \Delta$ for \mathcal{T} , and write $J(T)$ for an almost complex structure for $\widehat{\mathcal{T}}$, which has had a neck of length T inserted along c_β . For sufficiently large T ,

$$F_{\widehat{\mathcal{T}}, J(T)}(\Theta_{\alpha', \alpha}^+ \times \theta^+, \mathbf{x} \times -) = \begin{pmatrix} F_{\mathcal{T}, J}(\Theta_{\alpha', \alpha}^+, \mathbf{x})^{U_z \rightarrow U_{w'}} & 0 \\ 0 & F_{\mathcal{T}}(\Theta_{\alpha', \alpha}^+, \mathbf{x})^{U_z \rightarrow U_{w'}} \end{pmatrix}.$$

Proof The proof is identical to the proof of Proposition 13.6. □

Lemma 13.9 Almost complex structures J_α and J_β can be chosen on \mathcal{H}_1 , $\mathcal{H}_{1.5}$ and \mathcal{H}_2 so that the following hold:

- (1) Lemma 13.3 applies to compute the differentials on $\mathcal{H}_{1.5}$. Lemma 13.4 applies to compute the differential on \mathcal{H}_1 and \mathcal{H}_2 .
- (2) Proposition 13.6 applies to compute the transition map $\Psi_{(\mathcal{H}_1, J_\alpha) \rightarrow (\mathcal{H}_{1.5}, J_\alpha)}$ and Proposition 13.8 applies to compute $\Psi_{(\mathcal{H}_{1.5}, J_\beta) \rightarrow (\mathcal{H}_2, J_\beta)}$.
- (3) Lemma 13.5 applies to compute $\Psi_{(\mathcal{H}_{1.5}, J_\alpha) \rightarrow (\mathcal{H}_{1.5}, J_\beta)}$.
- (4) J_α can be used to compute S_w^+ on \mathcal{H}_1 , and J_β can be used to compute S_w^- on \mathcal{H}_2 .

Proof All of the results follow from the aforementioned results, together with Lemma 13.2, which shows that the transition map for changing between two choices of J_α with different relative neck lengths is the identity map, on the level of chain complexes. □

We now prove Theorem 13.1:

Proof of Theorem 13.1 Let J_1 and J_2 be the almost complex structures J_α and J_β on \mathcal{H}_1 and \mathcal{H}_2 , chosen so that the statements of Lemma 13.9 hold.

Decompose $\Psi_{(\mathcal{H}_1, J_\alpha) \rightarrow (\mathcal{H}_2, J_\beta)}$ as

$$(13-26) \quad \Psi_{(\mathcal{H}_1, J_\alpha) \rightarrow (\mathcal{H}_2, J_\beta)} = \Psi_{(\mathcal{H}_{1.5}, J_\beta) \rightarrow (\mathcal{H}_2, J_\beta)} \circ \Psi_{(\mathcal{H}_{1.5}, J_\alpha) \rightarrow (\mathcal{H}_{1.5}, J_\beta)} \circ \Psi_{(\mathcal{H}_1, J_\alpha) \rightarrow (\mathcal{H}_{1.5}, J_\alpha)}.$$




Figure 32: The path λ in the diagrams \mathcal{H}_1 and \mathcal{H}_2 .

Using Lemma 13.9, equation (13-26) implies that $\Psi_{(\mathcal{H}_1, J_\alpha) \rightarrow (\mathcal{H}_2, J_\beta)}$ is equal to, in matrix notation,

$$\begin{aligned} & \begin{pmatrix} (\Psi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_z \rightarrow U_{w'}} & 0 \\ 0 & (\Psi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_z \rightarrow U_{w'}} \end{pmatrix} \begin{pmatrix} \text{id} & 0 \\ * & \text{id} \end{pmatrix} \begin{pmatrix} (\Psi_{\alpha}^{\beta \rightarrow \beta'})_{U_z \rightarrow U_w} & 0 \\ 0 & (\Psi_{\alpha}^{\beta \rightarrow \beta'})_{U_z \rightarrow U_w} \end{pmatrix} \\ &= \begin{pmatrix} (\Psi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_z \rightarrow U_{w'}} \circ (\Psi_{\alpha}^{\beta \rightarrow \beta'})_{U_z \rightarrow U_w} & 0 \\ * & (\Psi_{\alpha}^{\beta \rightarrow \beta'})_{U_z \rightarrow U_w} \circ (\Psi_{\alpha \rightarrow \alpha'}^{\beta'})_{U_z \rightarrow U_w} \end{pmatrix}, \end{aligned}$$

as claimed. Furthermore, according to Lemma 13.9, the almost complex structure J_α can be used to compute $S_{w'}^+$ on \mathcal{H}_1 , and J_β can be used to compute S_w^- on \mathcal{H}_2 . \square

13.2 Basepoint-moving maps and the π_1 -action

In this section, we prove our proposed formulas for the basepoint-moving diffeomorphism maps.

We first need a computation of the relative homology map for the diagrams \mathcal{H}_1 and \mathcal{H}_2 considered in the previous section. In the diagrams \mathcal{H}_1 and \mathcal{H}_2 , let λ be the path shown in Figure 32.

Lemma 13.10 *On $(\mathcal{H}_1, J_\alpha)$, the relative homology map A_λ takes the form*

$$(A_\lambda)_{\mathcal{H}_1, J_\alpha} = \begin{pmatrix} U_w \cdot (\Phi_z)_{U_z \rightarrow U_w} & U_w \\ \text{id} & U_w \cdot (\Phi_z)_{U_z \rightarrow U_w} \end{pmatrix},$$

where Φ_z is the endomorphism of $\text{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{z\})$,

$$\Phi_z(x) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_z(\phi) \cdot \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}_0}^{n_{\mathbf{w}_0}(\phi)} U_z^{n_z(\phi) - 1} \cdot \mathbf{y}.$$

On (\mathcal{H}_2, J_β) , the map A_λ takes the form

$$(A_\lambda)_{\mathcal{H}_2, J_\beta} = \begin{pmatrix} 0 & U_w \\ \text{id} & 0 \end{pmatrix}.$$

Proof The statement follows from the holomorphic disk counts in Lemma 13.4, as we now explain.

We first consider the diagonal entries. The curves which contribute to the diagonal entries of $(A)_{\mathcal{H}_1, J_\alpha}$ represent classes of the form $\phi \# \phi_0$ in $\pi_2(\mathbf{x} \times \theta, \mathbf{y} \times \theta)$ for $\theta \in \{\theta^+, \theta^-\}$. Lemma 13.4 implies that these have no change across β_1 , and have a change of $n_z(\phi)$ across α_1 . The curves are counted by $(A_\lambda)_{\mathcal{H}_1, J_\alpha}$ with an additional factor of $a(\lambda, \phi \# \phi_0)$, which is thus $n_z(\phi)$. It follows that the diagonal entries are

equal to $U_w \cdot (\Phi_z)^{U_z \rightarrow U_w}$. We leave it to the reader to verify, using Lemma 13.4, that the off-diagonal entries are as claimed.

An entirely analogous argument works to compute $(A_\lambda)_{\mathcal{H}_1, J_\beta}$. □

Theorem 13.11 *Suppose λ is a path from w to w' , and $\mathbf{w}_0 \subseteq Y$ is a (possibly empty) collection of basepoints which contains neither w nor w' . The diffeomorphism map*

$$\lambda_*: \text{CF}^-(Y, (\mathbf{w}_0 \cup \{w\})^\sigma, \mathfrak{s}) \rightarrow \text{CF}^-(Y, (\mathbf{w}_0 \cup \{w'\})^{\sigma'}, \mathfrak{s})$$

satisfies

$$\lambda_* \simeq S_w^- A_\lambda S_{w'}^+$$

when σ' is the coloring obtained by pushing forward σ under λ .

Proof Assume that a diagram $(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w\})$ is chosen so that λ is embedded in Σ and is disjoint from α and β . Let $\phi_t: [0, 1] \times \Sigma \rightarrow \Sigma$ be an isotopy of Σ , supported in a small neighborhood of λ , such that $\phi_0 = \text{id}_\Sigma$ and $\phi_1(w) = w'$. Further, assume that ϕ_t is the identity outside a small neighborhood of λ . If J is an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$, then pushing forward under $\phi_1 \times \text{id}$ yields a tautological map

$$T: \text{CF}_J^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w\}) \rightarrow \text{CF}_{(\phi_1)_*J}^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w'\}).$$

By definition,

$$\lambda_*: \text{CF}_J^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w\}) \rightarrow \text{CF}_J^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w'\})$$

is the composition

$$(13-27) \quad \lambda_* = \Psi_{(\phi_1)_*J \rightarrow J} \circ T.$$

As a first step, we show that the expression for λ_* in (13-27) is equal to the tautological map on intersection points. The map $\Psi_{(\phi_1)_*J \rightarrow J}$ is obtained by counting Maslov index zero holomorphic disks for a dynamic almost complex structure which interpolates $(\phi_1)_*J$ and J . The isotopy ϕ_t induces an automorphism Φ of $\Sigma \times [0, 1] \times \mathbb{R}$,

$$\Phi(x, s, t) = (\phi_t(x), s, t).$$

We push J forward along Φ to get a dynamic almost complex structure, interpolating J and $(\phi_1)_*J$. However, $\Phi_*(J)$ -holomorphic disks on $\Sigma \times [0, 1] \times \mathbb{R}$ pull back under Φ to J -holomorphic disks on $\Sigma \times [0, 1] \times \mathbb{R}$. (The assumption that ϕ_t is fixed near α and β is necessary since it implies that the pullback of a $\Phi_*(J)$ -holomorphic disk has boundary mapping to the cylinders $\alpha \times \{0\} \times \mathbb{R}$ and $\beta \times \{0\} \times \mathbb{R}$.) By transversality for J , a Maslov index zero $\Phi_*(J)$ -holomorphic disk must be a constant disk. Hence, $\Psi_{(\phi_1)_*J \rightarrow J}$ is the tautological map on intersection points.

Hence, λ_* , being the composition of two tautological maps, is a tautological map.

We now consider the composition $S_w^- A_\lambda S_{w'}^+$. Let \mathcal{H}_1 and \mathcal{H}_2 be diagrams for $(Y, \mathbf{w}_0 \cup \{w, w'\})$ which are free stabilizations of $(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w\})$ or $(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w'\})$, respectively, as in Figure 32.

Similarly to Figure 24, let J_α and J_β be almost complex structures on \mathcal{H}_1 and \mathcal{H}_2 , obtained by stretching parallel to α_1 and β_2 , respectively.

The map $S_w^- A_\lambda S_{w'}^+$, viewed as a map from $CF_J^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w\})$ to $CF_J^-(\Sigma, \alpha, \beta, \mathbf{w}_0 \cup \{w'\})$, can be computed as the composition

$$(13-28) \quad S_w^- A_\lambda S_{w'}^+ = \Psi_{\alpha' \rightarrow \alpha}^{\beta' \rightarrow \beta} \circ S_w^- \circ (A_\lambda)_{\mathcal{H}_2, J_\beta} \circ \Psi_{(\mathcal{H}_1, J_\alpha) \rightarrow (\mathcal{H}_2, J_\beta)} \circ S_{w'}^+.$$

Using matrix notation, we can write

$$S_{w'}^+ = \begin{pmatrix} \text{id} \\ 0 \end{pmatrix} \quad \text{and} \quad S_w^- = (0 \text{ id}).$$

Using matrix notation, we rewrite (13-28) using Theorem 13.1 and Lemmas 13.5, and 13.10 as

$$(13-29) \quad S_w^- A_\lambda S_{w'}^+ \simeq \Psi_{\alpha' \rightarrow \alpha}^{\beta' \rightarrow \beta} (0 \text{ id}) \begin{pmatrix} 0 & U_w \\ \text{id} & 0 \end{pmatrix} \begin{pmatrix} \Psi_{\alpha \rightarrow \alpha'}^{\beta \rightarrow \beta'} & 0 \\ * & \Psi_{\alpha \rightarrow \alpha'}^{\beta \rightarrow \beta'} \end{pmatrix} \begin{pmatrix} \text{id} \\ 0 \end{pmatrix}.$$

Matrix multiplication reduces (13-29) to $\Psi_{\alpha' \rightarrow \alpha}^{\beta' \rightarrow \beta} \circ \Psi_{\alpha \rightarrow \alpha'}^{\beta \rightarrow \beta'}$, which is chain homotopic to the tautological map by naturality. Since we already showed that λ_* is the tautological map, the proof is complete. □

13.3 The Φ_w map

If $(\Sigma, \alpha, \beta, \mathbf{w})$ is a Heegaard diagram and $w \in \mathbf{w}$, then we write Φ_w for the endomorphism

$$(13-30) \quad \Phi_w(x) = U_w^{-1} \sum_{\substack{\phi \in \pi_2(x, y) \\ \mu(\phi)=1}} n_w(\phi) \cdot \# \widehat{\mathcal{M}}(\phi) U_w^{n_w(\phi)} \cdot y$$

of $CF^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathfrak{s})$.

Remark 13.12 The map Φ_w can be conveniently described algebraically as the *formal derivative* of the differential. Viewing $CF^-(\mathcal{H}, \mathfrak{s})$ as a free $\mathbb{F}_2[U_w]$ -module with basis $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we can write ∂ as a matrix with entries in $\mathbb{F}_2[U_w]$. The map Φ_w is obtained by differentiating each entry with respect to U_w .

We now prove several properties about the endomorphism Φ_w .

Lemma 13.13 *The map Φ_w is a chain map.*

Proof View the differential ∂ as a square matrix with entries in $\mathbb{F}_2[U_w]$. Differentiating $\partial \circ \partial = 0$ with respect to d/dU_w , using the Leibniz rule, yields

$$\partial \Phi_w + \Phi_w \partial = 0. \quad \square$$

Lemma 13.14 *The map Φ_w commutes with transition maps up to chain homotopy.*

Proof The transition maps satisfy

$$(13-31) \quad \partial \Psi_{\mathcal{H}_1 \rightarrow \mathcal{H}_2} + \Psi_{\mathcal{H}_1 \rightarrow \mathcal{H}_2} \partial = 0.$$

We apply d/dU_w to (13-31). The Leibniz rule implies

$$\Phi_w \Psi_{\mathcal{H}_1 \rightarrow \mathcal{H}_2} + \Psi_{\mathcal{H}_1 \rightarrow \mathcal{H}_2} \Phi_w + \partial \Psi'_{\mathcal{H}_1 \rightarrow \mathcal{H}_2} + \Psi'_{\mathcal{H}_1 \rightarrow \mathcal{H}_2} \partial = 0,$$

which says that Φ_w commutes with $\Psi_{\mathcal{H}_1 \rightarrow \mathcal{H}_2}$ up to chain homotopy. □

Lemma 13.15 *If $w \in \mathbf{w}$ and $|\mathbf{w}| > 1$ (so $S_w^+ S_w^-$ is defined), then*

$$\Phi_w \simeq S_w^+ S_w^-.$$

Proof Pick a diagram \mathcal{H} for (Y, \mathbf{w}) which is a free stabilization of a diagram \mathcal{H}_0 for $(Y, \mathbf{w} \setminus \{w\})$ at w . Furthermore, assume that the free stabilization is adjacent to another basepoint $w' \in \mathbf{w}$. Using Lemma 13.4, we may pick an almost complex structure so that

$$(13-32) \quad \partial_{\mathcal{H}} = \begin{pmatrix} (\partial_{\mathcal{H}_0})_{U_w} & U_w + U_{w'} \\ 0 & (\partial_{\mathcal{H}_0})_{U_w} \end{pmatrix},$$

where $(\partial_{\mathcal{H}_0})_{U_w}$ is the map obtained by extending $\partial_{\mathcal{H}_0}$ linearly over U_w , as in (13-2).

Differentiating (13-32) with respect to U_w yields

$$\Phi_w = \begin{pmatrix} 0 & \text{id} \\ 0 & 0 \end{pmatrix},$$

which is $S_w^+ S_w^-$. □

Lemma 13.16 *Suppose λ is a path between two distinct basepoints in \mathbf{w} . If $w \in \partial\lambda$, then*

$$A_\lambda \Phi_w + \Phi_w A_\lambda \simeq \text{id}.$$

If $w \notin \partial\lambda$, then

$$A_\lambda \Phi_w + \Phi_w A_\lambda \simeq 0.$$

Proof Assume first that $w \in \partial\lambda$. We apply d/dU_w to the equation

$$\partial A_\lambda + A_\lambda \partial = U_w + U_{w'}$$

from Lemma 4.1, obtaining the relation

$$A_\lambda \Phi_w + \Phi_w A_\lambda + A'_\lambda \partial + \partial A'_\lambda = \text{id}.$$

This proves the claim when $w \in \partial\lambda$. The claim when $w \notin \partial\lambda$ is proven similarly. □

Lemma 13.17 *If $w \neq w'$, then $\Phi_w S_{w'}^\circ + S_{w'}^\circ \Phi_w \simeq 0$ for $\circ \in \{+, -\}$.*

Proof Similar to Lemma 13.16, the claim is proven by differentiating the expression $\partial S_{w'}^\circ + S_{w'}^\circ \partial = 0$ with respect to U_w . □

Lemma 13.18 *The map Φ_w satisfies $(\Phi_w)^2 \simeq 0$ ($\mathbb{F}_2[U_w]$ -equivariantly).*

Proof Write $\partial = \sum_{n=0}^{\infty} \partial_n U_w^n$, where ∂_n does not involve any powers of U_w . Define

$$H := \sum_{n=2}^{\infty} \frac{1}{2}n(n-1)\partial_n U_w^{n-2};$$

the quantity $\frac{1}{2}n(n-1)$ is computed over \mathbb{Z} , then projected to \mathbb{F}_2 . The equality $\partial^2 = 0$ implies that

$$\sum_{i+j=k} \partial_i \partial_j = 0$$

for each $k \geq 0$. We now compute

$$\begin{aligned} \partial H + H \partial &= \sum_{i,j \geq 0} \left(\frac{1}{2}i(i-1) + \frac{1}{2}j(j-1) \right) \partial_i \partial_j U_w^{i+j-2} \\ &= \sum_{k=2}^{\infty} \left(\frac{1}{2}k(k-1) \sum_{i+j=k} \partial_i \partial_j \right) U_w^{k-2} + \sum_{i,j \geq 0} ij \partial_i \partial_j U_w^{i+j-2} \\ &= 0 + \Phi_w^2. \end{aligned}$$

□

13.4 The π_1 -action

We now prove our formula for the π_1 -action:

Theorem D *The action of $\gamma \in \pi_1(Y, w)$ satisfies*

$$\gamma_* \simeq \text{id} + \Phi_w \circ A_\gamma.$$

Proof Break γ into the concatenation of two paths, and write $\gamma = \lambda_2 * \lambda_1$, where λ_1 is a path from w to w' and λ_2 is a path from w' to w . Using Theorem 13.11, we have

$$(13-33) \quad (\gamma)_* \simeq (\lambda_2)_* \circ (\lambda_1)_* \simeq (S_{w'}^- A_{\lambda_2} S_w^+) (S_w^- A_{\lambda_1} S_{w'}^+).$$

We manipulate (13-33) as follows:

$$\begin{aligned} (\gamma)_* &\simeq S_{w'}^- A_{\lambda_2} \Phi_w A_{\lambda_1} S_{w'}^+ && \text{(Lemma 13.15)} \\ &\simeq S_{w'}^- (\Phi_w A_{\lambda_2} + \text{id}) A_{\lambda_1} S_{w'}^+ && \text{(Lemma 13.16)} \\ &\simeq S_{w'}^- \Phi_w A_{\lambda_2} A_{\lambda_1} S_{w'}^+ + S_{w'}^- A_{\lambda_1} S_{w'}^+ && \text{(algebra)} \\ &\simeq \Phi_w S_{w'}^- A_{\lambda_2} A_{\lambda_1} S_{w'}^+ + S_{w'}^- A_{\lambda_1} S_{w'}^+ && \text{(Lemma 13.17)} \\ &\simeq \Phi_w A_\gamma + S_{w'}^- A_{\lambda_1} S_{w'}^+ && \text{(Lemma 6.11)} \\ &\simeq \Phi_w A_\gamma + \text{id} && \text{(Lemma 6.10).} \end{aligned}$$

□

We now consider the triviality of the π_1 -action on homology. By Remark 13.12, the map Φ_w can be viewed as the derivative of the differential, with respect to U_w . Note that to apply the map d/dU_w to an element $\partial(x)$, we must use the Leibniz rule:

$$\frac{d}{dU_w}(\partial(x)) = \left(\frac{d}{dU_w} \partial \right)(x) + \partial \left(\frac{d}{dU_w} x \right).$$

Rearranging, we obtain

$$(13-34) \quad \Phi_w = \frac{d}{dU_w} \circ \partial + \partial \circ \frac{d}{dU_w}.$$

The map Φ_w is an $\mathbb{F}_2[U_w]$ -module homomorphism and hence induces an \mathcal{R}_P -module homomorphism after tensoring CF^- with the $\mathbb{F}_2[U_w]$ -module \mathcal{R}_P . The derivative d/dU_w is not an $\mathbb{F}_2[U_w]$ -module homomorphism and hence does not induce a chain null-homotopy on homology, unless the coloring is simple. Similarly, the map d/dU_w does not always persist to \widehat{CF} .

We summarize with the following:

Corollary 13.19 *The $\pi_1(Y, w)$ -action vanishes on $HF^\circ(Y, w^\sigma, \mathfrak{s})$ for $\circ \in \{+, -, \infty\}$ as long as w has a coloring distinct from all the other basepoints. In particular, it vanishes on the uncolored modules $HF^\circ(Y, w, \mathfrak{s})$ for $\circ \in \{+, -, \infty\}$.*

Corollary F is also a consequence:

Corollary F *If (W, γ) is a path cobordism between two singly based 3-manifolds (Y_0, w_0) and (Y_1, w_1) , then the cobordism map*

$$F_{W, \gamma, \mathfrak{s}}: HF^\circ(Y_0, w_0, \mathfrak{s}|_{Y_0}) \rightarrow HF^\circ(Y_1, w_1, \mathfrak{s}|_{Y_1})$$

is independent of γ if $\circ \in \{-, \infty, +\}$.

Proof Suppose that γ and γ' are two paths in W , connecting w_0 and w_1 . Decompose W as $W_3 \circ W_2 \circ W_1$, where W_i is obtained by attaching i -handles. Let \mathcal{Y} denote $W_2 \cap W_1$. Let $w'_0, w'_1 \in \mathcal{Y}$ denote the images of w_0 and w_1 , under the flow of gradient-like vector fields on $W_3 \circ W_2$, and W_1 . By flowing γ and γ' downwards in $W_3 \circ W_2$, and upwards in W_1 , we obtain two paths, λ and λ' , from w'_0 to w'_1 in \mathcal{Y} . We can decompose the two cobordism maps as

$$F_{W, \gamma, \mathfrak{s}} = F_3 \circ F_2 \circ \lambda_* \circ F_1 \quad \text{and} \quad F_{W, \gamma', \mathfrak{s}} = F_3 \circ F_2 \circ \lambda'_* \circ F_1,$$

where F_i denotes the i -handle map for W_i , and λ_* and λ'_* denote the two basepoint-moving diffeomorphism maps. That $\lambda_* = \lambda'_*$ as homomorphisms from $HF^\circ(\mathcal{Y}, w'_0, \mathfrak{s}|_{\mathcal{Y}})$ to $HF^\circ(\mathcal{Y}, w'_1, \mathfrak{s}|_{\mathcal{Y}})$ follows from Corollary 13.19. □

13.5 3-manifolds with nonvanishing π_1 -action

We are now equipped to prove Corollary E, and give examples where the π_1 -action is nonvanishing.

Corollary E *Let Y be a 3-manifold and $\mathfrak{s} \in \text{Spin}^c(Y)$.*

(1) *If \mathfrak{s} is torsion and there is an $x \in HF^+(Y, w, \mathfrak{s})$ such that*

$$U \cdot x = 0 \quad \text{and} \quad x \notin U \cdot HF^+(Y, w, \mathfrak{s}),$$

then $\pi_1(Y \# S^1 \times S^2, w)$ acts nontrivially on

$$\widehat{HF}(Y \# S^1 \times S^2, w).$$

(2) Suppose $[\gamma] \in H_1(Y; \mathbb{Z})$ is a class whose action on $\text{HF}^+(Y, w, \mathfrak{s})$ does not vanish. If $|\mathbf{w}| \geq 2$, then the diffeomorphism map γ_* acts nontrivially on the $\mathbb{F}_2[U]$ -module

$$\text{HF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s}),$$

where σ denotes the coloring which assigns all basepoints the variable U .

Proof We begin with claim (1). The classification theorem for finitely generated chain complexes over a PID (see, e.g., [Hendricks et al. 2018, Lemma 6.1]) implies that $\text{CF}^-(Y, w, \mathfrak{s})$ can be written as a direct sum of 1-step complexes (complexes with a single generator over $\mathbb{F}_2[U]$, and vanishing differential) and 2-step complexes $a \xrightarrow{p(U)} b$ (i.e., complexes with two generators over $\mathbb{F}_2[U]$, a and b , with $\partial(b) = 0$ and $\partial(a) = p(U) \cdot b$).

If $\mathfrak{s} \in \text{Spin}^c(Y)$ is torsion, then $\text{CF}^-(Y, w, \mathfrak{s})$ admits a relative \mathbb{Z} grading [Ozsváth and Szabó 2004b, Section 4.2]. Consequently, an algebraic argument (see [Hendricks et al. 2018, Lemma 6.2]) implies each of the 2-step complexes appearing must have $p(U) = U^i$ for some i .

The existence of an $x \in \text{HF}^+(Y, w, \mathfrak{s})$ such that $U \cdot x = 0$ and $x \notin U \cdot \text{HF}^+(Y, w, \mathfrak{s})$ implies that at least one summand of $\text{CF}^-(Y, w, \mathfrak{s})$ must be a 2-step complex of the form $a \xrightarrow{U} b$. The Φ_w map for such a complex satisfies $\Phi_w(a) = b$. In particular, on $\widehat{\text{HF}}(Y, w, \mathfrak{s})$ the map Φ_w is nonzero.

We view $\widehat{\text{HF}}(S^1 \times S^2, \mathfrak{s}_0)$ as the 2-dimensional vector space $V = \langle \theta^+, \theta^- \rangle$. If $\gamma \subseteq S^1 \times S^2$ is a curve which generates $H_1(S^1 \times S^2)$, then

$$A_\gamma(\theta^+) = \theta^- \quad \text{and} \quad A_\gamma(\theta^-) = 0.$$

Using the connected sum formula, $\widehat{\text{HF}}(Y \# S^1 \times S^2, w, \mathfrak{s} \# \mathfrak{s}_0)$ is isomorphic to $\widehat{\text{HF}}(Y, w, \mathfrak{s}) \otimes V$. From Theorem D and the connected sum formula, the map γ_* satisfies

$$\gamma_*(a \otimes \theta^+) = a \otimes \theta^+ + b \otimes \theta^- \quad \text{and} \quad \gamma_*(a \otimes \theta^-) = a \otimes \theta^-,$$

which is not equal to the identity map.

We now consider claim (2). Since $|\mathbf{w}| \geq 2$, pick a $w \in \mathbf{w}$ and write $\mathbf{w}_0 := \mathbf{w} \setminus \{w\}$. By Lemma 13.4, we obtain the decomposition

$$\text{HF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s}) \cong \text{HF}^-(Y, \mathbf{w}_0^\sigma, \mathfrak{s}) \otimes V.$$

Analyzing the curve counts appearing in Lemma 13.4, we conclude that the maps Φ_w and A_γ have the matrix descriptions

$$(13-35) \quad \Phi_w = \begin{pmatrix} 0 & \text{id} \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad A_\gamma = \begin{pmatrix} A_\gamma & 0 \\ 0 & A_\gamma \end{pmatrix}.$$

In (13-35), we are abusing notation slightly and writing A_γ for the induced map on both $\text{HF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s})$ and $\text{HF}^-(Y, \mathbf{w}_0^\sigma, \mathfrak{s})$.

Consequently,

$$(13-36) \quad \gamma_* = \text{id} + \Phi_w A_\gamma = \begin{pmatrix} \text{id} & A_\gamma \\ 0 & \text{id} \end{pmatrix},$$

which is not the identity if A_γ acts nontrivially on $\text{HF}^-(Y, \mathbf{w}_0^\sigma, \mathfrak{s})$. Using the formula from (13-35), we see that A_γ acts nontrivially on $\text{HF}^-(Y, \mathbf{w}_0^\sigma, \mathfrak{s})$ if and only if it acts nontrivially on the singly based complex. The main statement now follows. \square

13.6 Illustrating the π_1 -action on multipointed diagrams

In this section, we sketch an alternative proof of the formula for the π_1 -action, when there are at least two basepoints. In contrast to our proof of the normalization axiom (in particular Theorem 13.11), the proof we describe in this section is a direct holomorphic curve count, and does not use any functorial aspects of the graph TQFT.

Suppose \mathbf{w}_0 is a nonempty collection of basepoints in Y , and suppose $p \in Y \setminus \mathbf{w}_0$ is a new basepoint. If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0)$ is a Heegaard diagram for (Y, \mathbf{w}_0) , let

$$\mathcal{H}_p^+ = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w}_0 \cup \{p\})$$

denote the free stabilization of \mathcal{H} at p . As we observed in the proof of Corollary E (see (13-36)), the π_1 -action takes a simple form for free-stabilized diagrams. We now show that (13-36) can be proven using a direct holomorphic curve count:

Proposition 13.20 *Suppose \mathcal{H} is a diagram for (Y, \mathbf{w}_0) and \mathcal{H}_w^+ is obtained by free-stabilizing at w . If γ is a closed loop based at w , then*

$$(13-37) \quad \gamma_* \simeq \begin{pmatrix} \text{id} & A_\gamma \\ 0 & \text{id} \end{pmatrix},$$

with respect to the matrix decomposition $\text{CF}^-(\mathcal{H}_w^+, \mathfrak{s}) \cong \text{CF}^-(\mathcal{H}, \mathfrak{s}) \otimes_{\mathbb{F}_2} V$, where $V = \langle \theta^+, \theta^- \rangle$ is a 2-dimensional vector space.

If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram, and $p \in \Sigma \setminus \boldsymbol{\alpha} \cup \boldsymbol{\beta}$, we can define an endomorphism $\Omega_p: \text{CF}^-(\mathcal{H}, \mathfrak{s}) \rightarrow \text{CF}^-(\mathcal{H}, \mathfrak{s})$ via the formula

$$(13-38) \quad \Omega_p(\mathbf{x}) := \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} n_p(\phi) \cdot \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}_0}^{n_{\mathbf{w}_0}(\phi)} \cdot \mathbf{y}.$$

Remark 13.21 The map Ω_p appearing in (13-38) may not be a chain map. Instead Ω_p is a chain homotopy between the actions U_{w_α} and U_{w_β} , where $w_\alpha \in \mathbf{w}_0$ is the basepoint in the component of $\Sigma \setminus \boldsymbol{\alpha}$ containing p , and w_β is the basepoint in the component of $\Sigma \setminus \boldsymbol{\beta}$ containing p .

We need the following change of almost complex structure map computation:

Proposition 13.22 *Suppose \mathcal{H} is a Heegaard diagram for (Y, \mathbf{w}_0) , and \mathcal{H}_p^+ is the free stabilization of \mathcal{H} at the point p . Let J_α and J_β denote almost complex structures which have been stretched as in Figure 24.*




Figure 33: Multiplicities in the free-stabilization region.

For sufficiently long necks, the transition map $\Psi_{J_\beta \rightarrow J_\alpha}$ takes the form

$$\Psi_{J_\beta \rightarrow J_\alpha} = \begin{pmatrix} \text{id} & (\Omega_p)U_p \\ 0 & \text{id} \end{pmatrix},$$

where $\Omega_p: \text{CF}^-(\mathcal{H}, \mathfrak{s}) \rightarrow \text{CF}^-(\mathcal{H}, \mathfrak{s})$ is the map defined in (13-38).

Proof Similar to Lemma 13.5, we fix two neck lengths along c_α and c_β for J_α and J_β , respectively, and then stretch along c . Adapting the argument therein, for an appropriately chosen almost complex structure $\tilde{\mathcal{J}}$ on $\Sigma \times [0, 1] \times \mathbb{R}$, interpolating J_α and J_β with sufficiently large neck length along c , the transition map satisfies

$$(13-39) \quad \Psi_{J_\beta \rightarrow J_\alpha} := \Psi_{\tilde{\mathcal{J}}} = \begin{pmatrix} \text{id} & * \\ 0 & \text{id} \end{pmatrix}.$$

The $*$ appears in the upper right corner, whereas in Lemma 13.5 it appeared in the lower left corner, since the basepoint configuration is now different, and the Maslov grading of the points of $\alpha_0 \cap \beta_0$ is now reversed from Lemma 13.5.

We now show $* = \Omega_p$. This amounts to counting $\tilde{\mathcal{J}}$ -holomorphic representatives of index 0 classes $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^+)$. Let $n_1(\phi_0), n_2(\phi_0), m_1(\phi_0)$ and $n_p(\phi_0)$ denote the multiplicities shown in Figure 33.

For classes $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^+)$, equation (5-3) implies

$$(13-40) \quad \mu(\phi \# \phi_0) = \mu(\phi) + 2n_p(\phi_0) - 1.$$

If $\phi \# \phi_0$ is a class with $\tilde{\mathcal{J}}$ -holomorphic representatives with arbitrarily large neck length along c , then a limiting argument as in Lemma 13.2 shows that ϕ has a broken representative for a cylindrical almost complex structure. Hence, $\mu(\phi) \geq 0$, by transversality. Consequently, equation (13-40) implies that if $\phi \# \phi_0 \in \pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^+)$ has Maslov index 0 and has representatives for arbitrarily long neck length along c , then

$$\mu(\phi) = 1 \quad \text{and} \quad n_p(\phi_0) = 0.$$

If k is a fixed nonnegative integer, there are k nonnegative homology classes in $\pi_2(\theta^-, \theta^+)$ with $n_p = 0$ and $m_1 = k$. If k_1, k_2 are nonnegative integers with $k_1 + k_2 = k - 1$, we write $\phi_0^{k_1, k_2}$ for the class in $\pi_2(\theta^-, \theta^+)$ with

$$(13-41) \quad n_p(\phi_0^{k_1, k_2}) = 0, \quad m_1(\phi_0^{k_1, k_2}) = k, \quad n_1(\phi_0^{k_1, k_2}) = k_1 \quad \text{and} \quad n_2(\phi_0^{k_1, k_2}) = k_2.$$

Hence, to compute the component marked $*$ in (13-39), it is sufficient to compute

$$\#\mathcal{M}_{\tilde{\mathcal{J}}}(\phi \# \phi_0^{k_1, k_2})$$

when ϕ has Maslov index 1 and $k_1 + k_2 = m_1(\phi) - 1$. To this end, we will show that

$$(13-42) \quad \#\mathcal{M}_{\tilde{\mathcal{J}}}(\phi \# \phi_0^{k_1, k_2}) \equiv \#\widehat{\mathcal{M}}(\phi) \pmod{2}$$

for each nonnegative pair k_1, k_2 with $k_1 + k_2 = k - 1$. Note that the main statement follows quickly from (13-42).

Similar to the classes $\phi_0^{k_1, k_2}$, there are $k + 1$ nonnegative homology classes in $\pi_2(\theta^+, \theta^+)$ with $n_p = 0$ and $m_1 = k$. If $K_1 + K_2 = k$, write $\xi_0^{K_1, K_2}$ for the homology class in $\pi_2(\theta^+, \theta^+)$ defined similarly to (13-41). To establish (13-42), we count the ends of the 1-dimensional spaces $\mathcal{M}_{\tilde{\mathcal{J}}}(\phi \# \xi_0^{K_1, K_2})$.

As in the proof that $\Psi_{\tilde{\mathcal{J}}}$ is a chain map, the ends of $\mathcal{M}_{\tilde{\mathcal{J}}}(\phi \# \xi_0^{K_1, K_2})$ all result from strip breaking, i.e., correspond to pairs (\tilde{u}_1, u_2^β) or $(u_1^\alpha, \tilde{u}_2)$ where \tilde{u}_i denotes an index 0 $\tilde{\mathcal{J}}$ -holomorphic curve, and u_2^β and u_1^α denote index 1 J_β - or J_α -holomorphic curves. Furthermore, each end consists of one the following configurations:

- (e1) a pair $(\tilde{u}_1, u_2^\alpha)$, where \tilde{u}_1 represents an index 0 class in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{z} \times \theta^+)$ and u_2^α represents an index 1 class in $\pi_2(\mathbf{z} \times \theta^+, \mathbf{y} \times \theta^+)$;
- (e2) a pair $(\tilde{u}_1, u_2^\alpha)$ where \tilde{u}_1 represents an index 0 class in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{z} \times \theta^-)$ and u_2^α represents an index 1 class in $\pi_2(\mathbf{z} \times \theta^-, \mathbf{y} \times \theta^+)$;
- (e3) a pair (u_1^β, \tilde{u}_2) where u_1^β represents an index 1 class in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{z} \times \theta^+)$ and \tilde{u}_2 represents an index 0 class in $\pi_2(\mathbf{z} \times \theta^+, \mathbf{y} \times \theta^+)$;
- (e4) a pair (u_1^β, \tilde{u}_2) where u_1^β represents an index 1 class in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{z} \times \theta^-)$ and \tilde{u}_2 represents an index 0 class in $\pi_2(\mathbf{z} \times \theta^-, \mathbf{y} \times \theta^+)$.

In case (e1), the curve \tilde{u}_1 contributes to the first diagonal entry of $\Psi_{\tilde{\mathcal{J}}}$. The argument used to establish (13-39) implies that \tilde{u}_1 must represent a constant class $e_{\mathbf{x}} \times e_{\theta^+}$.

In case (e2), the curve \tilde{u}_1 contributes to the lower left entry of $\Psi_{\tilde{\mathcal{J}}}$. Our argument to establish (13-39) showed that no such curves \tilde{u}_1 exist, prohibiting the existence of these ends.

In case (e3), the curve \tilde{u}_2 contributes to the first diagonal entry of $\Psi_{\tilde{\mathcal{J}}}$. As with case (e1), the curve \tilde{u}_2 must represent the constant class $e_{\mathbf{y}} \times e_{\theta^+}$.

In case (e4) the curve u_1^β contributes to the lower left entry of the differential $\partial_{\mathcal{H}_p^+, J_\alpha}$, which vanishes. Hence, u_1^β represents one of two bigon classes in the free-stabilization region. See the proof of Lemma 13.3. The class \tilde{u}_2 contributes to the upper right component of the matrix for $\Psi_{\tilde{\mathcal{J}}}$ (which we are trying to compute). We write $B_{0,1}$ and $B_{1,0}$ for these bigons, chosen so that $n_1(B_{1,0}) = 1$ and $n_2(B_{0,1}) = 1$.

In summary, the ends of $\mathcal{M}_{\tilde{\mathcal{J}}}(\phi \# \xi_0^{K_1, K_2})$ are constrained to the following:

- (1) a $\tilde{\mathcal{J}}$ -holomorphic representative of $e_{\mathbf{x}} \times e_{\theta^+}$ and a J_α -holomorphic representative of $\phi \# \xi_0^{K_1, K_2}$;
- (2) a $\tilde{\mathcal{J}}$ -holomorphic representative of $e_{\mathbf{y}} \times e_{\theta^+}$ and a J_β -holomorphic representative of $\phi \# \xi_0^{K_1, K_2}$;




Figure 34: A Heegaard triple which has been free-stabilized. We stretch the almost complex structure along $c_{\beta, \beta'}$.

- (3) a \tilde{J} -holomorphic representative of $\phi \# \phi_0^{K_1-1, K_2}$ and a J_β -holomorphic representative of $B_{1,0}$;
- (4) a \tilde{J} -holomorphic representative of $\phi \# \phi_0^{K_1, K_2-1}$ and a J_β -holomorphic representative of $B_{0,1}$.

We can write out the ends as

$$(13-43) \quad \partial \bar{\mathcal{M}}_{\tilde{J}}(\phi \# \xi_0^{K_1, K_2}) = \widehat{\mathcal{M}}_{J_\alpha}(\phi \# \xi_0^{K_1, K_2}) \times \mathcal{M}_{\tilde{J}}(e_x \times e_{\theta^+}) \sqcup \mathcal{M}_{\tilde{J}}(e_y \times e_{\theta^+}) \times \widehat{\mathcal{M}}_{J_\beta}(\phi \# \xi_0^{K_1, K_2}) \\ \sqcup \widehat{\mathcal{M}}_{J_\beta}(B_{1,0}) \times \mathcal{M}_{\tilde{J}}(\phi \# \phi_0^{K_1-1, K_2}) \sqcup \widehat{\mathcal{M}}_{J_\beta}(B_{0,1}) \times \mathcal{M}_{\tilde{J}}(\phi \# \phi_0^{K_1, K_2-1}).$$

The bigons and constant classes appearing in (13-43) have unique holomorphic representatives. Since the total ends of a compact 1-manifold are 0, equation (13-43) implies

$$(13-44) \quad \#\widehat{\mathcal{M}}_{J_\alpha}(\phi \# \xi_0^{K_1, K_2}) + \#\widehat{\mathcal{M}}_{J_\beta}(\phi \# \xi_0^{K_1, K_2}) + \#\mathcal{M}_{\tilde{J}}(\phi \# \phi_0^{K_1-1, K_2}) + \#\mathcal{M}_{\tilde{J}}(\phi \# \phi_0^{K_1, K_2-1}) = 0.$$

The proof of Lemma 13.4 adapts immediately to show that

$$(13-45) \quad \#\widehat{\mathcal{M}}_{J_\beta}(\phi \# \xi_0^{K_1, K_2}) = \begin{cases} \#\widehat{\mathcal{M}}(\phi) & \text{if } (K_1, K_2) = (k, 0), \\ 0 & \text{otherwise,} \end{cases}$$

and

$$(13-46) \quad \#\widehat{\mathcal{M}}_{J_\alpha}(\phi \# \xi_0^{K_1, K_2}) = \begin{cases} \#\widehat{\mathcal{M}}(\phi) & \text{if } (K_1, K_2) = (0, k), \\ 0 & \text{otherwise.} \end{cases}$$

By starting at $(K_1, K_2) = (k, 0)$ and using (13-44), (13-45) and (13-46) to go through all pairs (K_1, K_2) with $K_1 + K_2 = k$, we obtain (13-42). □

We need the following triangle count:

Lemma 13.23 *Suppose $\mathcal{T} = (\Sigma, \alpha, \beta, \beta', w_0)$ is a Heegaard triple (we do not assume β' are small isotopies of β). Let \mathcal{T}_p^+ be the Heegaard triple obtained by free-stabilizing at p , as in Figure 34. Let J be*




Figure 35: Breaking the closed curve γ into a sequence of arcs $\lambda_1, \dots, \lambda_n$.

an almost complex structure on $\Sigma \times [0, 1] \times \mathbb{R}$, and let $J(T)$ denote an almost complex structure with a neck of length T added along $c_{\beta, \beta'}$. For sufficiently large T ,

$$F_{\mathcal{T}_p^+, J(T)}(\mathbf{x} \times -, \mathbf{y} \times \theta^+) = \begin{pmatrix} F_{\mathcal{T}, J}(\mathbf{x}, \mathbf{y}) & 0 \\ 0 & F_{\mathcal{T}, J}(\mathbf{x}, \mathbf{y}) \end{pmatrix}.$$

Proof The statement is a special case of [Manolescu and Ozsváth 2025, Proposition 5.2]. □

We now sketch our proof of the formula for the π_1 -action when there is more than one basepoint:

Proof of Proposition 13.20 Immerse γ in Σ , and break γ into a sequence of embedded arcs, $\lambda_1, \dots, \lambda_n$ such that each λ_i crosses a single α or β curve exactly once. See Figure 35.

Write $\partial\lambda_i = \{p_i, p_{i+1}\}$. Note that $p_1 = p_{n+1} = w$. Let $\mathcal{H}_{p_i}^+$ denote a free stabilization of \mathcal{H} at p_i .




Figure 36: Decomposing $(\lambda_i)_*$ into a triangle map followed by a tautological diffeomorphism map.

Let $J_{\alpha,i}$ denote an almost complex structure for $\mathcal{H}_{p_i}^+$ stretched along a circle c_α parallel to the α circle in the free-stabilization region, and let $J_{\beta,i}$ denote one stretched along a circle c_β parallel to the β circle.

Each arc λ_i induces a diffeomorphism map

$$(13-47) \quad (\lambda_i)_*: \text{CF}_{J_{\beta,i}}^-(\mathcal{H}_{p_i}^+, \mathfrak{s}) \rightarrow \text{CF}_{J_{\beta,i+1}}^-(\mathcal{H}_{p_{i+1}}^+, \mathfrak{s}).$$

First, we claim that if λ_i intersects a single β curve (and no α curves), then $(\lambda_i)_*$ is chain homotopic to

$$(13-48) \quad (\lambda_i)_* \simeq \begin{pmatrix} \text{id} & 0 \\ 0 & \text{id} \end{pmatrix}.$$

To establish (13-48), we decompose λ_i into two maps, a holomorphic triangle count to handleslide a β curve over p_i , followed by a tautological diffeomorphism map to move p_i to p_{i+1} . See Figure 36.

Equation (13-48) follows from Lemma 13.23, since we are using the almost complex structure $J_{\beta,i}$. Hence, the triangle counts on the free-stabilized Heegaard triple for the handleslide are obtained from the triangle counts on the unstabilized diagram. Once we compose with the tautological map for an isotopy, the diagonal entries are chain homotopic to the identity, by naturality.

Next, we claim that if λ_i intersects a single α curve (and no β curves), then $(\lambda_i)_*$ is chain homotopic to

$$(13-49) \quad (\lambda_i)_* \simeq \begin{pmatrix} \text{id} & \Omega_{p_i} + \Omega_{p_{i+1}} \\ 0 & \text{id} \end{pmatrix}.$$

Equation (13-49) is established as follows. By simply switching the roles of the α and β curves, our previous argument shows that the map

$$(\lambda_i)_*: \text{CF}_{J_{\alpha,i}}^-(\mathcal{H}_{p_i}^+, \mathfrak{s}) \rightarrow \text{CF}_{J_{\alpha,i}}^-(\mathcal{H}_{p_{i+1}}^+, \mathfrak{s})$$

takes the same form as in (13-48). To obtain the induced map $(\lambda_i)_*$ when the domain and codomain have almost complex structures as in (13-47), we must precompose with $\Psi_{J_{\beta,i} \rightarrow J_{\alpha,i}}$ and postcompose with $\Psi_{J_{\alpha,i+1} \rightarrow J_{\beta,i+1}}$. These two transition maps may be computed using Proposition 13.22. Upon multiplying the three matrices together, we obtain (13-49).

Finally, to compute

$$\gamma_*: \text{CF}_{J_{\beta,1}}^-(\mathcal{H}_w^+, \mathfrak{s}) \rightarrow \text{CF}_{J_{\beta,1}}^-(\mathcal{H}_w^+, \mathfrak{s}),$$

we write

$$\gamma_* = (\lambda_n)_* \circ \cdots \circ (\lambda_1)_*,$$

and use the expressions for λ_i from (13-48) and (13-49). Upon multiplying all the matrices out, we obtain that

$$(13-50) \quad \gamma_* = \begin{pmatrix} \text{id} & \sum_{i \in \{1, \dots, n\}, \lambda_i \cap \alpha \neq \emptyset} (\Omega_{p_i} + \Omega_{p_{i+1}}) \\ 0 & \text{id} \end{pmatrix}.$$

The map

$$\sum_{\substack{i \in \{1, \dots, n\} \\ \lambda_i \cap \alpha \neq \emptyset}} (\Omega_{p_i} + \Omega_{p_{i+1}})$$

is exactly equal to the map which counts index 1 holomorphic disks, with a factor equal to their total change across the α curves. This is exactly the homology action A_γ . Consequently, equation (13-50) coincides with (13-37). □

13.7 Basepoint-swapping diffeomorphism

We now compute an additional diffeomorphism map. Suppose $w_1, w_2 \in \mathbf{w}$ are two distinct basepoints, and λ is a path connecting w_1 and w_2 . Let

$$\text{Sw}_\lambda: (Y, \mathbf{w}) \rightarrow (Y, \mathbf{w})$$

denote the diffeomorphism corresponding to swapping w_1 and w_2 , along λ . The diffeomorphism Sw_λ is well defined up to isotopy, since Y is 3-dimensional.

The diffeomorphism Sw_λ is the map induced by the graph cobordism $([0, 1] \times Y, \Gamma)$, where Γ is the union of the arcs $[0, 1] \times \{w\}$ for $w \in \mathbf{w} \setminus \{w_1, w_2\}$, as well as two arcs, one connecting $\{w_1\} \times \{0\}$ to $\{w_2\} \times \{1\}$ and the other connecting $\{w_2\} \times \{0\}$ to $\{w_1\} \times \{1\}$, both of which project to λ , up to isotopy.

Note that Sw_λ only induces an \mathcal{R}_P -equivariant map on $\text{CF}^-(Y, \mathbf{w}^\sigma, \mathfrak{s})$ for colorings σ satisfying $\sigma(w_1) = \sigma(w_2)$.

Proposition 13.24 *Let $\text{Sw}_\lambda: (Y, \mathbf{w}) \rightarrow (Y, \mathbf{w})$ be the basepoint-swapping diffeomorphism. The induced map $(\text{Sw}_\lambda)_*$ satisfies*

$$(13-51) \quad (\text{Sw}_\lambda)_* \simeq \Phi_{w_1} A_\lambda + A_\lambda \Phi_{w_2} \simeq A_\lambda \Phi_{w_1} + \Phi_{w_2} A_\lambda.$$

Proof The second chain homotopy in (13-51) can be obtained from Lemma 13.16, so we focus on the first.

Let $w' \notin \mathbf{w}$ be a new basepoint, given the same color as w_1 and w_2 . Pick paths λ_1 and λ_2 from w_1 to w' and from w' to w_2 (respectively) such that the concatenation $\lambda_1 * \lambda_2$ is isotopic to λ . Break Sw_λ into the composition of the following three diffeomorphisms:

- (1) an isotopy of w_1 to w' , along λ_1 ;
- (2) an isotopy of w_2 to w_1 , along λ ;
- (3) an isotopy of w' to w_2 , along λ_2 .

Using Theorem 13.11, we obtain

$$(13-52) \quad (\text{Sw}_\lambda)_* \simeq (S_{w'}^- A_{\lambda_2} S_{w_2}^+) (S_{w_2}^- A_\lambda S_{w_1}^+) (S_{w_1}^- A_{\lambda_1} S_{w'}^+).$$

Using Lemmas 13.13–13.18, we manipulate (13-52) as follows:

$$\begin{aligned}
 (13-53) \quad (\text{Sw}_\lambda)_* & \\
 & \simeq S_{w'}^- A_{\lambda_2} \Phi_{w_2} A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ \\
 & \simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ + S_{w'}^- A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ \\
 & \simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda A_{\lambda_1} S_{w'}^+ \Phi_{w_1} + \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda S_{w'}^+ + S_{w'}^- A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ \\
 & \simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda A_{\lambda_1} S_{w'}^+ \Phi_{w_1} + \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda S_{w'}^+ + \Phi_{w_1} S_{w'}^- A_\lambda A_{\lambda_1} S_{w'}^+ + S_{w'}^- A_{\lambda_1} S_{w'}^+.
 \end{aligned}$$

We focus on the first term of the last line of (13-53). By Lemma 4.3, $A_\lambda = A_{\lambda_2} + A_{\lambda_1}$ when all are defined. Also $A_{\lambda_i}^2 \simeq U$ (where U denotes any of U_{w_1} , U_{w_2} or $U_{w'}$, which are identified by the coloring) by Lemma 4.5. Hence,

$$(13-54) \quad A_{\lambda_2} A_\lambda A_{\lambda_1} \simeq A_{\lambda_2} (A_{\lambda_1} + A_{\lambda_2}) A_{\lambda_1} \simeq A_{\lambda_2} A_{\lambda_1}^2 + A_{\lambda_2}^2 A_{\lambda_1} \simeq U (A_{\lambda_1} + A_{\lambda_2}) \simeq U A_\lambda.$$

Using (13-54), as well as Lemmas 5.13 and 5.15, we see

$$(13-55) \quad \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda A_{\lambda_1} S_{w'}^+ \Phi_{w_1} \simeq U \Phi_{w_2} S_{w'}^- A_\lambda S_{w'}^+ \Phi_{w_1} \simeq U \Phi_{w_2} A_\lambda S_{w'}^- S_{w'}^+ \Phi_{w_1} \simeq 0.$$

We now consider the second term of the last line of (13-53). Lemmas 5.13 and 6.10 imply

$$(13-56) \quad \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda S_{w'}^+ \simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} S_{w'}^+ A_\lambda \simeq \Phi_{w_2} A_\lambda.$$

Similarly, the third term of the last line of (13-53) becomes

$$(13-57) \quad \Phi_{w_1} S_{w'}^- A_\lambda A_{\lambda_1} S_{w'}^+ \simeq \Phi_{w_1} A_\lambda S_{w'}^- A_{\lambda_1} S_{w'}^+ \simeq \Phi_{w_1} A_\lambda.$$

The last term of the last line of (13-53) satisfies

$$(13-58) \quad S_{w'}^+ A_{\lambda_1} S_{w'}^+ \simeq \text{id},$$

by Lemma 6.10.

Combining (13-53), (13-55), (13-56), (13-57) and (13-58), we obtain

$$(13-59) \quad (\text{Sw}_\lambda)_* \simeq \Phi_{w_2} A_\lambda + \Phi_{w_1} A_\lambda + \text{id}.$$

Equation (13-59) is chain homotopic to the expression in the statement since $\Phi_{w_1} A_\lambda + A_\lambda \Phi_{w_1} \simeq \text{id}$ by Lemma 13.16. \square

Acknowledgements

I would like to thank Antonio Alfieri, András Juhász, Kristen Hendricks, Robert Lipshitz, Yajing Liu, Ciprian Manolescu, Mike Miller, and Matthew Stoffregen for helpful conversations. I would like to thank Ciprian Manolescu and Peter Ozsváth for helping me with the holomorphic geometry related to the 1- and 3-handle maps.

References

- [Alishahi and Eftekhary 2015] **A S Alishahi, E Eftekhary**, *A refinement of sutured Floer homology*, *J. Symplectic Geom.* 13:3 (2015) 609–743 MR
- [Bourgeois 2002] **F Bourgeois**, *A Morse–Bott approach to contact homology*, PhD thesis, Stanford University (2002) MR Available at <https://www.proquest.com/docview/305591502>
- [Bourgeois et al. 2003] **F Bourgeois, Y Eliashberg, H Hofer, K Wysocki, E Zehnder**, *Compactness results in symplectic field theory*, *Geom. Topol.* 7 (2003) 799–888 MR
- [Colin et al. 2024a] **V Colin, P Ghiggini, K Honda**, *The equivalence of Heegaard Floer homology and embedded contact homology, III: From hat to plus*, *Publ. Math. Inst. Hautes Études Sci.* 139 (2024) 349–385 MR
- [Colin et al. 2024b] **V Colin, P Ghiggini, K Honda**, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions, I*, *Publ. Math. Inst. Hautes Études Sci.* 139 (2024) 13–187 MR
- [Colin et al. 2024c] **V Colin, P Ghiggini, K Honda**, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions, II*, *Publ. Math. Inst. Hautes Études Sci.* 139 (2024) 189–348 MR
- [Colin et al. 2025] **V Colin, P Ghiggini, K Honda**, *Embedded contact homology and open book decompositions*, *Geom. Topol.* 29:7 (2025) 3345–3463 MR
- [Etnyre et al. 2017] **J B Etnyre, D S Vela-Vick, R Zarev**, *Sutured Floer homology and invariants of Legendrian and transverse knots*, *Geom. Topol.* 21:3 (2017) 1469–1582 MR
- [Golla 2015] **M Golla**, *Comparing invariants of Legendrian knots*, *Quantum Topol.* 6:3 (2015) 365–402 MR
- [Hendricks et al. 2018] **K Hendricks, C Manolescu, I Zemke**, *A connected sum formula for involutive Heegaard Floer homology*, *Selecta Math.* 24:2 (2018) 1183–1245 MR
- [Honda et al. 2008] **K Honda, W Kazez, G Matić**, *Contact structures, sutured Floer homology and TQFT*, preprint (2008) arXiv 0807.2431
- [Hutchings and Taubes 2009] **M Hutchings, C H Taubes**, *The Weinstein conjecture for stable Hamiltonian structures*, *Geom. Topol.* 13:2 (2009) 901–941 MR
- [Juhász 2006] **A Juhász**, *Holomorphic discs and sutured manifolds*, *Algebr. Geom. Topol.* 6 (2006) 1429–1457 MR
- [Juhász 2016] **A Juhász**, *Cobordisms of sutured manifolds and the functoriality of link Floer homology*, *Adv. Math.* 299 (2016) 940–1038 MR
- [Juhász 2018] **A Juhász**, *Defining and classifying TQFTs via surgery*, *Quantum Topol.* 9:2 (2018) 229–321 MR
- [Juhász et al. 2021] **A Juhász, D Thurston, I Zemke**, *Naturality and mapping class groups in Heegaard Floer homology*, *Mem. Amer. Math. Soc.* 1338, Amer. Math. Soc., Providence, RI (2021) MR
- [Kronheimer and Mrowka 2007] **P Kronheimer, T Mrowka**, *Monopoles and three-manifolds*, *New Mathematical Monographs* 10, Cambridge Univ. Press (2007) MR
- [Kutluhan et al. 2020a] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, I: Heegaard Floer homology and Seiberg–Witten Floer homology*, *Geom. Topol.* 24:6 (2020) 2829–2854 MR
- [Kutluhan et al. 2020b] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, II: Reeb orbits and holomorphic curves for the ech/Heegaard Floer correspondence*, *Geom. Topol.* 24:6 (2020) 2855–3012 MR
- [Kutluhan et al. 2020c] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, III: Holomorphic curves and the differential for the ech/Heegaard Floer correspondence*, *Geom. Topol.* 24:6 (2020) 3013–3218 MR
- [Kutluhan et al. 2020d] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, IV: The Seiberg–Witten Floer homology and ech correspondence*, *Geom. Topol.* 24:7 (2020) 3219–3469 MR
- [Kutluhan et al. 2020e] **Ç Kutluhan, Y-J Lee, C H Taubes**, *HF = HM, V: Seiberg–Witten Floer homology and handle additions*, *Geom. Topol.* 24:7 (2020) 3471–3748 MR
- [Laudenbach and Poénaru 1972] **F Laudenbach, V Poénaru**, *A note on 4-dimensional handlebodies*, *Bull. Soc. Math. France* 100 (1972) 337–344 MR
- [Lipshitz 2006] **R Lipshitz**, *A cylindrical reformulation of Heegaard Floer homology*, *Geom. Topol.* 10:2 (2006) 955–1096 MR

- [Lipshitz 2014] **R Lipshitz**, *Correction to the article “A cylindrical reformulation of Heegaard Floer homology”*, *Geom. Topol.* 18:1 (2014) 17–30 MR
- [Lipshitz et al. 2018] **R Lipshitz, P S Ozsvath, D P Thurston**, *Bordered Heegaard Floer homology*, *Mem. Amer. Math. Soc.* 1216, Amer. Math. Soc., Providence, RI (2018) MR
- [Manolescu and Ozsváth 2025] **C Manolescu, P Ozsváth**, *Heegaard Floer homology and integer surgeries on links*, *Geom. Topol.* 29:6 (2025) 2783–3062 MR
- [Ni 2014] **Y Ni**, *Homological actions on sutured Floer homology*, *Math. Res. Lett.* 21:5 (2014) 1177–1197 MR
- [Ozsváth and Szabó 2004a] **P Ozsváth, Z Szabó**, *Holomorphic disks and three-manifold invariants: properties and applications*, *Ann. of Math.* 159:3 (2004) 1159–1245 MR
- [Ozsváth and Szabó 2004b] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, *Ann. of Math.* 159:3 (2004) 1027–1158 MR
- [Ozsváth and Szabó 2006] **P Ozsváth, Z Szabó**, *Holomorphic triangles and invariants for smooth four-manifolds*, *Adv. Math.* 202:2 (2006) 326–400 MR
- [Ozsváth and Szabó 2008] **P Ozsváth, Z Szabó**, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, *Algebr. Geom. Topol.* 8:2 (2008) 615–692 MR
- [Sarkar 2011] **S Sarkar**, *Maslov index formulas for Whitney n -gons*, *J. Symplectic Geom.* 9:2 (2011) 251–270 MR
- [Sarkar 2015] **S Sarkar**, *Moving basepoints and the induced automorphisms of link Floer homology*, *Algebr. Geom. Topol.* 15:5 (2015) 2479–2515 MR
- [Turaev 1994] **V G Turaev**, *Quantum invariants of knots and 3-manifolds*, *De Gruyter Studies in Mathematics* 18, de Gruyter, Berlin (1994) MR
- [Wehrheim and Woodward 2015] **K Wehrheim, C Woodward**, *Floer field theory for tangles*, preprint (2015) arXiv 1503.07615
- [Zemke 2019] **I Zemke**, *Link cobordisms and functoriality in link Floer homology*, *J. Topol.* 12:1 (2019) 94–220 MR
- [Zemke 2021a] **I Zemke**, *Duality and mapping tori in Heegaard Floer homology*, *J. Topol.* 14:3 (2021) 1027–1112 MR
- [Zemke 2021b] **I Zemke**, *A graph TQFT for hat Heegaard Floer homology*, *Quantum Topol.* 12:3 (2021) 439–460 MR

IAN ZEMKE izemke@uoregon.edu

Department of Mathematics, Princeton University, Princeton, NJ, United States

Current address: *Department of Mathematics, University of Oregon, Eugene, OR, United States*

Proposed: Peter Ozsváth

Received: January 26, 2016

Seconded: Ciprian Manolescu, András I Stipsicz

Revised: June 6, 2019

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITORS

Robert Lipshitz University of Oregon
lipshitz@uoregon.edu
András I Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University abouzaid@stanford.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Dan Abramovich	Brown University dan_abramovich@brown.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Sándor Kovács	University of Washington skovacs@uw.edu
Arend Bayer	University of Edinburgh arend.bayer@ed.ac.uk	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Agnès Beaudry	University of Colorado Boulder agnes.beaudry@colorado.edu	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Mark Behrens	University of Notre Dame mbehren1@nd.edu	Jianfeng Lin	Tsinghua University linjian5477@mail.tsinghua.edu.cn
Mladen Bestvina	University of Utah bestvina@math.utah.edu	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Martin R Bridson	University of Oxford bridson@maths.ox.ac.uk	Haynes R Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Tobias H Colding	Massachusetts Institute of Technology colding@math.mit.edu	Aaron Naber	Institute for Advanced Studies anaber@ias.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Peter Ozsváth	Princeton University petero@math.princeton.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
David M Fisher	Rice University davidfisher@rice.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
Mike Freedman	Microsoft Research michaelf@microsoft.com	Roman Sauer	Karlsruhe Institute of Technology roman.sauer@kit.edu
David Gabai	Princeton University gabai@princeton.edu	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Natasa Sesum	Rutgers University natasas@math.rutgers.edu
Cameron Gordon	University of Texas gordon@math.utexas.edu	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHES and NYU, Courant Institute gromov@ihes.fr	Kirsten Wickelgren	Duke University kirsten.wickelgren@duke.edu
Mark Gross	University of Cambridge mgross@dpms.cam.ac.uk	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2026 is US \$905/year for the electronic version, and \$1275/year (+\$80, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

GT peer review and production are managed by EditFLOW[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2026 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 30 Issue 2 (pages 389–833) 2026

Graph cobordisms and Heegaard Floer homology	389
IAN ZEMKE	
Trivalent vertices and bordered knot Floer homology in the standard basis	529
ANDREW MANION	
Symplectomorphisms of some Weinstein 4-manifolds	645
PAUL HACKING and AILSA KEATING	
A Weiss–Williams theorem for spaces of embeddings and the homotopy type of spaces of long knots	701
SAMUEL MUÑOZ-ECHÁNIZ	
Cables of the figure-eight knot via real Frøyshov invariants	781
SUNGKYUNG KANG, JUNGHWAN PARK and MASAKI TANIGUCHI	