

Naturality and functoriality in involutive Heegaard Floer homology

Kristen Hendricks, Jennifer Hom, Matthew Stoffregen, and Ian Zemke

Abstract. We prove first-order naturality of involutive Heegaard Floer homology, and furthermore, construct well-defined maps on involutive Heegaard Floer homology associated to cobordisms between three-manifolds. We also prove analogous naturality and functoriality results for involutive Floer theory for knots and links. The proof relies on the doubling model for the involution, as well as several variations.

Contents

1. Introduction	2
2. Algebraic preliminaries	10
3. Involutive Heegaard Floer homology	17
4. On naturality in Heegaard Floer theory	23
5. Maps for elementary handleslide equivalences	25
6. Holomorphic polygons and stabilizations	28
7. An expanded model of the involution	33
8. Relating the expanded and ordinary models of the involution	42
9. 1- and 3-handles	59
10. 2-handles	72
11. Naturality maps for stabilizations	84
12. Handleswaps	89
13. Completing the proofs of naturality and functoriality	97
14. The cobordism map for $S^2 \times S^2$	101
15. Knots and links	105
16. 2-handle maps and the knot surgery formula	114
17. The $SO(3)$ -twist	123
References	142

Mathematics Subject Classification 2020: 57K41 (primary); 57K18 (secondary).

Keywords: 3-manifolds, cobordism, homology cobordism, Heegaard Floer homology, involutive Heegaard Floer homology.

1. Introduction

Heegaard Floer homology, defined by Ozsváth and Szabó in the early 2000s [22, 23], is a powerful suite of invariants of 3-manifolds, knots and links inside of them, and 4-dimensional cobordisms between them. Given a basepointed 3-manifold (Y, w) together with a choice of Spin^c structure \mathfrak{s} on Y , the initial construction of the invariant goes by choosing a Heegaard diagram \mathcal{H} for (Y, w) and associating to it a free finitely-generated chain complex $CF^-(\mathcal{H}, \mathfrak{s})$ over the ring $\mathbb{F}[U]$, where U is a variable of degree -2 . If \mathcal{H} and \mathcal{H}' are two Heegaard diagrams representing (Y, w) , the Reidemeister–Singer theorem implies that there is some sequence of Heegaard moves that may be applied to \mathcal{H} to produce a diagram related to \mathcal{H}' by a basepoint-preserving diffeomorphism isotopic to the identity in Y . Ozsváth and Szabó [23, Section 9.2] proved that these sequences of moves induce transition maps, that is, chain homotopy equivalences

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

showing that the isomorphism class of the homology group

$$HF^-(\mathcal{H}) = \bigoplus_{\mathfrak{s} \in \text{Spin}^c(Y)} HF^-(\mathcal{H}, \mathfrak{s})$$

is an invariant of (Y, w) .

Juhász, Thurston, and the fourth author [9] proved that this data is *first-order natural* in the sense that the maps $\Psi_{\mathcal{H} \rightarrow \mathcal{H}'}$ are independent up to chain homotopy of the choice of Heegaard moves and diffeomorphisms isotopic to the identity relating the two diagrams. (Here, “first-order” means that the chain homotopies relating the transition maps associated to two possible sequences of Heegaard moves are not known to be independent of the choices involved in their definition.) It follows that Heegaard Floer homology associates to (Y, w) a specific module $HF^-(Y, w)$ rather than an isomorphism class of $\mathbb{F}[U]$ -modules.

In 2015, the first author and Manolescu [5] put additional structure on the Heegaard Floer package in the form of a homotopy involution ι , leading to involutive Heegaard Floer homology. The construction of the invariant goes as follows: given a basepointed 3-manifold (Y, w) together with a conjugation-invariant Spin^c structure \mathfrak{s} on Y , there is a chain isomorphism

$$\eta: CF^-(\mathcal{H}, \mathfrak{s}) \rightarrow CF^-(\bar{\mathcal{H}}, \mathfrak{s}),$$

where $\bar{\mathcal{H}}$ denotes the conjugate Heegaard diagram. One may then compose with the transition map

$$\Psi_{\bar{\mathcal{H}} \rightarrow \mathcal{H}}: CF^-(\bar{\mathcal{H}}, \mathfrak{s}) \rightarrow CF^-(\mathcal{H}, \mathfrak{s})$$

obtaining the map ι . This data is variously packaged either as the pair $(CF^-(Y), \iota)$, called an ι -complex, or as the mapping cone

$$CFI^-(\mathcal{H}, \mathfrak{s}) = \text{Cone}(CF^-(\mathcal{H}, \mathfrak{s}) \xrightarrow{Q(1+\iota)} Q \cdot CF^-(\mathcal{H}, \mathfrak{s})[-1]).$$

Here, Q is a formal variable of degree -1 , and we view the cone as a complex over $\mathbb{F}[U, Q]/Q^2$. Juhász–Thurston–Zemke naturality implies that the equivariant chain homotopy equivalence class of the pair $(CF^-(\mathcal{H}, \mathfrak{s}), \iota)$, or equivalently $\mathbb{F}[U, Q]/Q^2$ chain homotopy equivalence class of $CFI^-(\mathcal{H}, \mathfrak{s})$, is an invariant of Y [5, Section 2.1]. However, this does not suffice to show that involutive Heegaard Floer homology is itself natural.

In this paper, we prove that the involutive Heegaard Floer homology is a natural invariant of the basepointed 3-manifold (Y, w) together with a choice of framing $\xi = (\xi_1, \xi_2, \xi_3)$ of the oriented normal bundle to the basepoint; equivalently, that transition maps between the involutive complexes induced by Heegaard moves and basepointed diffeomorphisms isotopic to the identity in Y which preserve ξ are unique up to homotopy. Our proof relies on equipping a Heegaard diagram \mathcal{H} for (Y, w, ξ) with a set of *doubling data* which is used to give a tractable model for the involution. We write \mathfrak{D} for a Heegaard diagram \mathcal{H} equipped with such a collection of data, and we refer to such a \mathfrak{D} as a *doubling enhanced Heegaard diagram* for (Y, w, ξ) . We refer to the involutive Heegaard Floer homology of (Y, w, ξ) as defined using the Heegaard diagram \mathcal{H} with this model for the involution as $CFI^-(\mathfrak{D})$. (For more information on the doubling model for the involution, see Section 1.2.)

We also describe a set of moves which relate any two choices of doubling enhanced Heegaard diagrams. Given a sequence of such moves relating \mathfrak{D} and \mathfrak{D}' , we define a transition map

$$\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} : CFI^-(\mathfrak{D}) \rightarrow CFI^-(\mathfrak{D}')$$

and show they are independent of the sequence chosen, as follows.

Theorem 1.1. *The transition map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ is independent up to $\mathbb{F}[U, Q]/Q^2$ -equivariant chain homotopy from the sequence of moves between doubling enhanced Heegaard diagrams. Furthermore, $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}} \simeq \text{id}$, and $\Psi_{\mathfrak{D}' \rightarrow \mathfrak{D}''} \circ \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} \simeq \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}''}$.*

Remark 1.2. The involutive transition map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ above contains data equivalent to a pair $(\Psi_{\mathcal{H} \rightarrow \mathcal{H}'}, H_{\mathfrak{D} \rightarrow \mathfrak{D}'})$, where $\Psi_{\mathcal{H} \rightarrow \mathcal{H}'}$ is the non-involutive transition map from $CF^-(\mathcal{H})$ to $CF^-(\mathcal{H}')$, and $H_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ is a distinguished chain-homotopy between $\Psi_{\mathcal{H} \rightarrow \mathcal{H}'} \iota_{\mathfrak{D}}$ and $\iota_{\mathfrak{D}'} \Psi_{\mathcal{H} \rightarrow \mathcal{H}'}$. In this manner, Theorem 1.1 implies that the chain homotopy $H_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ is also well defined, up to a suitable notion of further chain-homotopy.

Given a 4-dimensional cobordism W with

$$\partial W = -Y_1 \coprod Y_2$$

and a Spin^c structure \mathfrak{s} on W , Ozsváth and Szabó [23] also constructed cobordism maps

$$CF(W, \mathfrak{s}) : CF^-(Y_1, w_1, \mathfrak{s}|_{Y_1}) \rightarrow CF^-(Y_2, w_2, \mathfrak{s}|_{Y_2}).$$

These maps depend on a choice of path γ connecting the basepoints in Y_1 and Y_2 but are otherwise independent of choices made in their definition [23, Sections 8 and 9], cf. also [39, Corollary F].

In [5], the first author and Manolescu gave a construction of cobordism maps on involutive Heegaard Floer homology, but did not show their construction was invariant of a choice of handle decomposition of the cobordism. In this paper, we refine the construction to give maps

$$CFI(W, \xi, \mathfrak{s}) : CFI^-(Y_1, w_1, \mathfrak{s}|_{Y_1}, \xi|_{Y_1}) \rightarrow CFI^-(Y_2, w_2, \mathfrak{s}|_{Y_2}, \xi|_{Y_2})$$

associated to a cobordism W from Y_1 to Y_2 , a conjugation-invariant Spin^c structure on W , and a choice of framing ξ of the normal bundle to a choice of path γ which induces the framings $\xi|_{Y_i}$ of the normal bundle to the basepoint w_i in Y_i . We prove these maps are well defined in the following sense.

Theorem 1.3. *The cobordism map $CFI(W, \xi, \mathfrak{s})$ is well defined; that is, it depends only on the choice of W , \mathfrak{s} , γ , and ξ .*

Our theory has some subtle differences from the Ozsváth–Szabó cobordism maps; in particular, the role of duality is not presently obvious. See Remark 14.2 for further discussion.

Heegaard Floer homology have a counterpart for knots and links in 3-manifolds, introduced by Ozsváth and Szabó [21] and independently in the case of knots by J. Rasmussen [28]. In its modern form, this theory associates to a Heegaard diagram \mathcal{H} for a link L in a 3-manifold Y , with two basepoints w_i and z_i on each component, a chain complex $\mathcal{CF}\mathcal{L}(\mathcal{H})$ over the ring $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$. This again decomposes over Spin^c structures on Y such that

$$\mathcal{CF}\mathcal{L}(\mathcal{H}) = \bigoplus_{\mathfrak{s} \in \text{Spin}^c(Y)} \mathcal{CF}\mathcal{L}(\mathcal{H}, \mathfrak{s}).$$

As previously, there are maps associated to sequence of Heegaard moves between Heegaard diagrams \mathcal{H} and \mathcal{H}' for $(Y, L, \mathbf{w}, \mathbf{z})$:

$$\Psi_{\mathcal{H} \rightarrow \mathcal{H}'} : \mathcal{CF}\mathcal{L}(\mathcal{H}, \mathfrak{s}) \rightarrow \mathcal{CF}\mathcal{L}(\mathcal{H}', \mathfrak{s}).$$

These maps are again well defined up to chain homotopy [9, Theorem 1.8].

The first author and Manolescu [5] defined an endomorphism ι_L on $\mathcal{CF}\mathcal{L}(\mathcal{H})$, called the link involution, in analogy with the procedure for defining the 3-manifold

involution ι , and showed that the equivariant chain homotopy equivalence class of the pair

$$\mathcal{CFLI}(\mathcal{H}) = (\mathcal{CFL}(\mathcal{H}), \iota_L)$$

is an invariant of (Y, L) . (We note that [5] focused on the case of knots, rather than links, but the same construction works for links.) The doubling model for the involution ι on 3-manifolds may be adapted to a model for the endomorphism ι_L . We use this model to prove naturality and functoriality of the link variant of Floer homology in the following sense. Once again, we let \mathfrak{D} denote an appropriate set of doubling data for a Heegaard diagram \mathcal{H} for $(Y, L, \mathbf{w}, \mathbf{z})$, and let $\mathcal{CFLI}(\mathfrak{D})$ refer to the ι_L -complex defined using this data. This data need *not* include a framing of either basepoint; see the discussion in Section 17.7 for more on this issue. Given a sequence of Heegaard moves relating two doubling-enhanced Heegaard link diagrams, we define a transition map

$$\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} : \mathcal{CFLI}(\mathfrak{D}) \rightarrow \mathcal{CFLI}(\mathfrak{D}')$$

which are enhanced ι_L -homomorphisms. (Types of morphisms between ι_L complexes are reviewed in Section 2.5.) We then prove the following.

Theorem 1.4. *The transition maps $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ are well defined up to homotopies of morphisms of ι_L -complexes.*

The fourth author [35], following work of Juhasz [7], defined maps associated to decorated link cobordisms $(W, \Sigma, \mathfrak{s})$ between $(Y_1, L_1, \mathbf{w}_1, \mathbf{z}_1, \mathfrak{s}|_{Y_1})$ and $(Y_2, L_2, \mathbf{w}_2, \mathbf{z}_2, \mathfrak{s}|_{Y_2})$ as follows:

$$\mathcal{CFL}(W, \Sigma, \mathfrak{s}) : \mathcal{CFL}(Y_1, L_1, \mathbf{w}_1, \mathbf{z}_1, \mathfrak{s}|_{Y_1}) \rightarrow \mathcal{CFL}(Y_2, L_2, \mathbf{w}_2, \mathbf{z}_2, \mathfrak{s}|_{Y_2}),$$

which are again invariant of the choices involved in their definitions. Here, the pair (W, Σ) consists of a compact 4-manifold W with embedded surface Σ such that W is a cobordism from Y_1 to Y_2 and

$$\partial \Sigma = -L_1 \coprod L_2,$$

together with some decorations on Σ . We will be interested in the case that Σ is a collection of annuli, each with one boundary component in Y_1 and one in Y_2 , and each decorated with two parallel longitudinal arcs. Maps associated to knot concordances or knot cobordisms have also appeared in [1, 8, 31].

Given a link cobordism (W, Σ) from $(Y_1, L_1, \mathbf{w}, \mathbf{z})$ to $(Y_2, L_2, \mathbf{w}, \mathbf{z})$ consisting of a set of annuli each with two parallel longitudinal arcs as described above, and a Spin^c structure \mathfrak{s} on W with the property that $\bar{\mathfrak{s}} = \mathfrak{s} + PD([\Sigma])$, we construct cobordism maps

$$\mathcal{CFLI}(W, \Sigma, \mathfrak{s}) : \mathcal{CFLI}(Y_1, L_1, \mathbf{w}, \mathbf{z}, \mathfrak{s}|_{Y_1}) \rightarrow \mathcal{CFLI}(Y_2, L_2, \mathbf{w}, \mathbf{z}, \mathfrak{s}|_{Y_2}),$$

which are again enhanced homomorphisms of ι_L -complexes. We prove the following functoriality theorem.

Theorem 1.5. *The cobordism maps $\mathcal{CFI}(W, \Sigma, \mathfrak{s})$ are well defined up to homotopies of ι_L -complexes.*

1.1. Computations and examples

In our paper, we also perform several example computations. The first example we present is for the cobordism $W = S^2 \times S^2$, with its unique self-conjugate Spin^c structure \mathfrak{s} . If we remove two 4-balls from W , we obtain a cobordism from S^3 to S^3 , and hence a map

$$CFI(W, \xi, \mathfrak{s}) : \mathbb{F}[U, Q]/Q^2 \rightarrow \mathbb{F}[U, Q]/Q^2.$$

In Section 14, we show via a direct diagrammatic computation that $CFI(W, \xi, \mathfrak{s})$ is multiplication by Q . (The choice of framings does not affect the result in this case.) This agrees with the prediction from $\text{Pin}(2)$ -equivariant monopole Floer homology; see the proof of [13, Theorem 5].

As a second example, we compute certain 2-handle cobordisms using the knot surgery formula from [4]. We recall that to a knot $K \subseteq S^3$, Ozsváth and Szabó [27, Theorem 1.1] described a complex $\mathbb{X}_n(K)$ which computes $CF^-(S_n^3(K))$. Therein, they also described a way of computing the non-involutive cobordism map $CF(W_n(K), \mathfrak{s})$ in terms of $\mathcal{CFK}(K)$, where $W_n(K)$ denotes the natural 2-handle cobordism from S^3 to $S_n^3(K)$, for each $\mathfrak{s} \in \text{Spin}^c(W_n(K))$. In [4], we described an involutive refinement of Ozsváth and Szabó's mapping cone formula, which we denoted by $\mathbb{XI}_n(K)$, and which computes $CFI(S_n^3(K))$. In Section 16, we describe a refinement of our work from [4] which computes certain cobordism maps. From the description in [4], it turns out to be most convenient to consider the cobordism $W'_n(K)$ from $S_n^3(K)$ to S^3 obtained by reversing the orientation of $W_n(K)$. This manifold possesses a self-conjugate Spin^c structure if and only if n is even. In Theorem 16.1, we describe an algebraic formula for the involutive cobordism map $CFI(W'_{2n}(K), \xi, \mathfrak{s})$ when \mathfrak{s} is the unique self-conjugate Spin^c structure on $W'_{2n}(K)$. Again, the choice of framings does not affect the computation.

1.2. The doubling model of the involution and several variations

The constructions of this paper rely on a model for the involution ι , used in our previous paper [4], based on the procedure of *doubling* a Heegaard diagram. We now briefly recall this model. If Σ is a Heegaard splitting of Y , containing a basepoint w , we construct another Heegaard splitting $D(\Sigma)$ of Y , with $D(\Sigma) \cong \Sigma \# \bar{\Sigma}$, which is embedded as the boundary of a regular neighborhood of $\Sigma \setminus N(w)$. A schematic

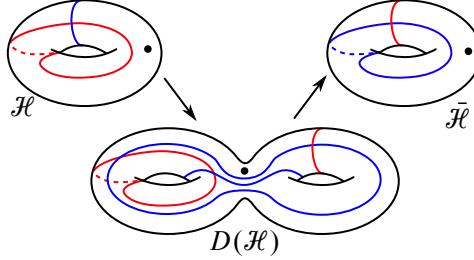


Figure 1.1. Realizing the involution on S^3 by doubling.

appears in Figure 1.1. The attaching curves for doubled diagrams are described in Section 3.2. An important property of the doubling operation is that if \mathcal{H} is a diagram for Y , and $D(\mathcal{H})$ is a double, then the maps relating $CF^-(\mathcal{H})$ and $CF^-(D(\mathcal{H}))$ have a conceptually simple form, and similarly for the maps from $CF^-(D(\mathcal{H}))$ to $CF^-(\bar{\mathcal{H}})$; see Section 3.3. (Doubled Heegaard diagrams had previously been considered in [10, 11, 36].)

The proofs in the present paper require additional variations of this model for the involution, based on adding extra basepoints or extra tubes between the two copies of Σ and $\bar{\Sigma}$. In the 3-manifold case, we begin by describing transition maps between the involutive complexes associated to Heegaard diagrams for a 3-manifold which are related by *elementary equivalences* in terms of maps on the complexes associated to the doubled diagrams, and define maps for general handleslide equivalences as compositions of the maps for elementary equivalences. However, it is somewhat difficult to see that this construction fails to depend on the sequence of elementary equivalences chosen. To show this, we consider an expanded version of this model, in which an additional basepoint and pair of curves are added in the doubling region. We introduce this model in detail in Section 7. Using this model we define transition maps for handleslide equivalences which are more obviously invariant of the choices involved, and then use this to prove invariance of the transition maps for the standard doubling model. In the case of knots and links, our situation is somewhat more technically complex, so for simplicity we in fact work solely with the analog of the expanded model for the involution. Throughout, the curve-counting arguments in our computations rely heavily on the cylindrical formulation of [14].

Remark 1.6. We emphasize here that there are two possible doubling models one could use to construct the map ι and our cobordism maps. The model described above and pictured in Figure 1.1 is called the *beta-doubling* model of the involution; the analogous construction with the roles of the alpha and beta curves reversed is the *alpha-doubling* model. This distinction is examined more closely in Section 3.5. Our techniques also show naturality and functoriality for involutive Heegaard Floer

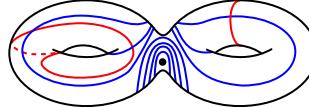


Figure 1.2. The effect of changing the framing on the doubled diagram from Figure 1.1.

homology constructed using the alpha-doubling model. We do not, however, attempt to relate these two models.

1.3. The role of framings

Theorem 1.1 differs from Juhász–Thurston–Zemke naturality in the incorporation of a framing of the normal bundle to the basepoint w in Y . In particular, $CFI(Y, w)$ is acted on not by $\text{Diff}^+(Y, w)$, the orientation-preserving diffeomorphisms of (Y, w) , but by $\text{Diff}^f(Y, w)$, the diffeomorphisms of (Y, w) which preserving the framing ξ at the basepoint. The Dehn twist tw in a neighborhood of w is an example of a class in $\pi_0(\text{Diff}(Y, w, \xi))$ which is trivial in $\pi_0(\text{Diff}(Y, w))$, but is not *a priori* trivial in $\pi_0(\text{Diff}(Y, w, \xi))$. We calculate the following.

Theorem 1.7. *The element $\text{tw} \in \pi_0(\text{Diff}^f(Y, w))$ acts on $CFI(Y, w)$ by $\text{Id} + Q\Phi$.*

Here, Φ is a formal derivative of the differential with respect to U ; for a review of this map, see Section 2.3. The action of tw on $CFI(Y, w)$ can also be interpreted as an action of loops in the space of framings of (Y, w) , as in Section 17. The space of such framings is a copy of $SO(3)$, and therefore has fundamental group $\pi_1(SO(3)) \cong \mathbb{Z}/2$, whose nontrivial element acts as tw above.

In the doubling model described above, the choice of framing affects the choice of curves in the doubled diagram; concretely, it corresponds to a Dehn twist around a meridian of the connect sum tube in $\Sigma \# \bar{\Sigma}$, as in Figure 1.2.

The reader will note that we do not require a framing of the basepoint for the naturality and functoriality results for links. Heuristically, this follows from the fact that the connect sum region of (any variation of) the doubled diagram for a link does not contain a basepoint; for more details, see Section 17.7.

We expect that the dependence on the framings and the appearance of the map Φ in this computation are related to a conjectural $\text{Pin}(2)$ -equivariant structure on Heegaard Floer homology, in analogy with Manolescu’s $\text{Pin}(2)$ -Seiberg Witten Floer homology and Lin’s $\text{Pin}(2)$ -monopole Floer homology, and hope to investigate this connection in future work. We further note that the appearance of the framings in our theorems implies that any infinite order naturality statement for Heegaard Floer homology would necessarily also take framings of the basepoint into account.

1.4. Organization

This paper is organized as follows. In Section 2, we discuss some algebraic preliminaries and their Floer-theoretic counterparts; in particular, we discuss hypercubes of chain complexes and hypercubes of attaching curves for a Heegaard surface. We additionally review the algebraic formalism of ι -complexes and their relationship to hypercubes of chain complexes. In Section 3, we review some background on involutive Heegaard Floer homology, and in particular recall the operation of doubling a Heegaard diagram and the resulting model for the involution on Heegaard Floer homology from [4]. In Section 4, we recall the structure of Juhász, Thurston, and Zemke’s proof of naturality of ordinary Heegaard Floer homology and describe how to adapt it to our case. In Section 5, we define transition maps for Heegaard diagrams related by elementary equivalences, and define general transition maps for Heegaard diagrams related by handleslide equivalences as compositions of maps for elementary equivalences. We prove in Proposition 5.4 that these maps satisfy the continuity axiom, in analogy with [9, Proposition 9.27]. In Section 6, we prove some necessary technical results concerning stabilizations and holomorphic polygons, generalizing the results from our previous paper [4]. In Section 7, we introduce the basepoint expanded doubling diagrams and the basepoint expanded model of the involution, and construct transition maps for handleslide equivalences using this model. In Section 8, we define a chain homotopy equivalence between involutive chain complexes associated to the ordinary and basepoint expanded doubles of a Heegaard diagram, and prove that these maps commute with the transitions maps associated to handleslide equivalences. We further show that the transition maps defined in Section 7 for generalized handleslide equivalences are unique up to chain homotopy equivalence, which proves the result for the non-expanded model. We then turn our attention to constructing the maps between involutive complexes associated to cobordism; in Section 9, we construct the maps associated to one- and three-handles and in Section 10, we construct the maps associated to two-handles. In Section 11, we define construct naturality maps for stabilizations as compositions of cobordism maps associated to cancelling handles and prove they commute with the transition maps associated to handleslide equivalences and with each other. In Section 12, we prove invariance of our transition maps under handleswaps, in analogy with [9, Section 9.3]. Finally, in Section 13, we complete the proofs of Theorem 1.1, following the structure of the proof in [9, Section 9.4], and of Theorem 1.3. As an example, in Section 14, we present the calculation of the cobordism map for a twice punctured copy of $S^2 \times S^2$, decomposed as 2-handle cobordisms $S^3 \rightarrow S^1 \times S^2$ and $S^1 \times S^2 \rightarrow S^3$. In Section 15, we describe how to adapt the naturality and functoriality results of the previous sections to the case of knots and links, and prove Theorems 1.4 and 1.5. We use this in Section 16 to provide a computation of the involutive cobordism map

associated to the cobordism $W'_n(K)$ from $S_n^3(K)$ to S^3 for n even in terms of the involutive mapping cone formula. We conclude by analyzing the effect of changing the framing in Section 17, giving a proof of Theorem 1.7.

2. Algebraic preliminaries

2.1. Hypercubes and hyperboxes

In this section, we recall the hypercube formalism of Manolescu and Ozsváth [18]. We write \mathbb{E}_n for the set of points $\{0, 1\}^n \subseteq \mathbb{Z}^n$. Similarly, if $\mathbf{d} = (d_1, \dots, d_n)$ is a tuple of n positive integers, we write

$$\mathbb{E}(\mathbf{d}) = \{0, \dots, d_1\} \times \dots \times \{0, \dots, d_n\}.$$

If $\varepsilon, \varepsilon' \in \mathbb{R}^n$, we say that $\varepsilon \leq \varepsilon'$ if the inequality holds for each component of ε and ε' .

Definition 2.1. An n -dimensional hypercube of chain complexes $(C^\varepsilon, D^{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ is a collection of groups C^ε , ranging over $\varepsilon \in \mathbb{E}_n$, together with a collection of maps $D^{\varepsilon, \varepsilon'}: C^\varepsilon \rightarrow C^{\varepsilon'}$, whenever $\varepsilon \leq \varepsilon'$. Furthermore, we assume the following compatibility condition is satisfied whenever $\varepsilon \leq \varepsilon''$:

$$\sum_{\varepsilon': \varepsilon \leq \varepsilon' \leq \varepsilon''} D^{\varepsilon', \varepsilon''} \circ D^{\varepsilon, \varepsilon'} = 0. \quad (2.1)$$

We usually call $D^{\varepsilon, \varepsilon}$ the internal differential of C^ε . The hypercube structure relation in equation (2.1) is equivalent to $(\bigoplus_{\varepsilon \in \mathbb{E}_n} C^\varepsilon, \sum_{\varepsilon \leq \varepsilon'} D^{\varepsilon, \varepsilon'})$ being a chain complex.

Definition 2.2. A hyperbox of chain complexes of size \mathbf{d} consists of a collection of groups $(C^\varepsilon)_{\varepsilon \in \mathbb{E}(\mathbf{d})}$, together with a choice of map $D^{\varepsilon, \varepsilon'}: C^\varepsilon \rightarrow C^{\varepsilon'}$ whenever $|\varepsilon' - \varepsilon|_{L^\infty} \leq 1$. We assume equation (2.1) holds whenever $|\varepsilon'' - \varepsilon|_{L^\infty} \leq 1$.

An important operation involving hyperboxes is *compression* [18, Section 5]. This operation takes a hyperbox of size $\mathbf{d} = (d_1, \dots, d_n)$ and returns an n -dimensional hypercube of chain complexes. We illustrate with an example. Consider a 1-dimensional hyperbox

$$C_0 \xrightarrow{f_{0,1}} C_1 \xrightarrow{f_{1,2}} \dots \xrightarrow{f_{n-1,n}} C_n.$$

The compression of the above hyperbox is the hypercube

$$C_0 \xrightarrow{f_{n-1,n} \circ \dots \circ f_{0,1}} C_n.$$

More generally, the same description also works for an n -dimensional hyperbox \mathcal{C} of size $(1, \dots, 1, d)$. That is, we may view \mathcal{C} as a 1-dimensional hyperbox of size

(d), which we call \mathcal{C}' , where the complex at each point ε_0 of $\mathbb{E}(d)$ is the $(n-1)$ -dimensional hypercube $\mathcal{C}'_{\varepsilon_0} = (\bigoplus_{\varepsilon \in \mathbb{E}_{n-1}} C^{\varepsilon \varepsilon_0}, \sum_{\varepsilon \leq \varepsilon'} D^{\varepsilon \varepsilon_0, \varepsilon' \varepsilon_0})$, viewed as a chain complex. Schematically, \mathcal{C}' is represented as a diagram

$$\mathcal{C}'_0 \xrightarrow{f_{0,1}} \mathcal{C}'_1 \xrightarrow{f_{1,2}} \cdots \xrightarrow{f_{d-1,d}} \mathcal{C}'_d,$$

where the maps $f_{i,i+1}$ are constructed as a sum of the maps $D^{\varepsilon_i, \varepsilon'_i + 1}$ for $\varepsilon \leq \varepsilon' \in \mathbb{E}_{n-1}$.

The definition given above for compression of a 1-dimensional hyperbox of chain complexes then applies to construct a 1-dimensional hypercube:

$$\mathcal{C}'_0 \xrightarrow{f_{d-1,d} \circ \cdots \circ f_{0,1}} \mathcal{C}'_n.$$

For $\varepsilon \leq \varepsilon' \in \mathbb{E}_n$ with $\varepsilon_n \neq \varepsilon'_n$, the map $D^{\varepsilon, \varepsilon'}$ of the compression of \mathcal{C} is the component of $f_{d-1,d} \circ \cdots \circ f_{0,1}$ from \mathcal{C}^ε to $\mathcal{C}^{\varepsilon'}$.

For a hyperbox C of size (d_1, \dots, d_n) , the above description may be iterated to give the compression, as follows. This description depends on a choice of ordering of the axes of the cube. (The motivated reader may verify that re-ordering the axes results in a homotopic hypercube.) Firstly, we may view C as a collection of d_1, \dots, d_{n-1} hyperboxes of size $(1, \dots, 1, d_n)$. We compress each of these hyperboxes using the function-composition description above. Stacking these hyperboxes results in a hyperbox of size $(d_1, \dots, d_{n-1}, 1)$. We may view this as a collection of d_1, \dots, d_{n-2} hypercubes of size $(1, \dots, 1, d_{n-1}, 1)$, which we compress using the function-composition description. Stacking the resulting hyperboxes gives a hyperbox of size $(d_1, \dots, d_{n-2}, 1, 1)$. Repeating this procedure gives the compression of C . It is possible to relate this description to the description in terms of the algebra of songs in [18], cf. [17, Section 4.1.2].

2.2. Hypercubes of attaching curves

In this section, we describe the notion of a hypercube in the Fukaya category. Such objects are described by Manolescu and Ozsváth [18], and are referred to as *hyperboxes of Heegaard diagrams*. The reader may also compare this construction to Lipshitz, Ozsváth and Thurston's notion of a *chain complex of attaching circles* [15]. The construction is also described in [4, Section 5.5].

We begin with a preliminary definition about Spin^c structures. Firstly, we say $\mathcal{L}_\beta = (\beta^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$ is an *empty hypercube of beta (or alpha)-attaching curves* if it is a collection of attaching curves on a surface Σ , indexed by points of \mathbb{E}_n . Let $Y_{\varepsilon_1, \varepsilon_2}$ be the 3-manifold determined by the two sets of attaching curves $(\beta^{\varepsilon_1}, \beta^{\varepsilon_2})$ and $X_{\varepsilon_1, \dots, \varepsilon_m}$ be the smooth 4-manifold associated to the collection of sets of attaching curves $(\beta^{\varepsilon_1}, \dots, \beta^{\varepsilon_m})$ in the usual way [23, Section 8.1].

Definition 2.3. (1) Suppose that $\mathcal{L}_\beta = (\beta^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$ is an empty hypercube of beta-curves. A *hypercube of Spin^c-structures* for \mathcal{L}_β consists of a collection of Spin^c structures

$$\mathfrak{S}_{\varepsilon_1, \dots, \varepsilon_m} \subseteq \text{Spin}^c(X_{\varepsilon_1, \dots, \varepsilon_m})$$

for each sequence $\varepsilon_1 < \dots < \varepsilon_m$ in \mathbb{E}_n , satisfying $m > 1$ and the following compatibility relations. Firstly, if $1 \leq i < j \leq m$, then $\mathfrak{S}_{\varepsilon_1, \dots, \varepsilon_m}$ is closed under the action of $\delta^1 H^1(Y_{\varepsilon_i, \varepsilon_j})$. Secondly if $\varepsilon_1 < \dots < \varepsilon_m$ is a sequence, and $\varepsilon_{i_1} < \dots < \varepsilon_{i_j}$ is a subsequence, where $1 \leq i_1 < \dots < i_j \leq m$, then $\mathfrak{S}_{\varepsilon_{i_1}, \dots, \varepsilon_{i_j}}$ is the image of $\mathfrak{S}_{\varepsilon_1, \dots, \varepsilon_m}$ under the natural restriction map

$$\text{Spin}^c(X_{\varepsilon_1, \dots, \varepsilon_m}) \rightarrow \text{Spin}^c(X_{\varepsilon_{i_1}, \dots, \varepsilon_{i_j}}).$$

(2) If $\mathcal{L}_\alpha = (\alpha^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$ and $\mathcal{L}_\beta = (\beta^\varepsilon)_{\varepsilon \in \mathbb{E}_m}$ are two empty hypercubes of alpha and beta attaching curves, then a hypercube of Spin^c structures for the pair $(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ consists of the following. For every pair of sequences $(\nu_1 < \dots < \nu_k)$ in \mathbb{E}_n and $(\varepsilon_1 < \dots < \varepsilon_\ell)$ in \mathbb{E}_m such that $k + \ell > 1$, a set of Spin^c structures $\mathfrak{S}_{\nu_k, \dots, \nu_1, \varepsilon_1, \dots, \varepsilon_\ell} \subseteq \text{Spin}^c(X_{\nu_k, \dots, \nu_1, \varepsilon_1, \dots, \varepsilon_\ell})$ that are closed under δ^1 -orbits and compatible with restriction, similar to above. The case that one of $(\nu_1 < \dots < \nu_k)$ and $(\varepsilon_1 < \dots < \varepsilon_\ell)$ is the empty list is allowed.

Note that a hypercube of Spin^c structures for the pair $(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ induces a hypercube of Spin^c structures for each of \mathcal{L}_α and \mathcal{L}_β .

Definition 2.4. An n -dimensional hypercube of beta attaching curves \mathcal{L}_β on a pointed surface (Σ, w) consists of an empty hypercube of beta attaching curves $(\beta^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$, a hypercube of Spin^c structures on \mathcal{L}_β , together with a choice of Floer chain $\Theta_{\varepsilon, \varepsilon'} \in CF^-(\Sigma, \beta^\varepsilon, \beta^{\varepsilon'}, w)$ whenever $\varepsilon < \varepsilon'$. Furthermore, the chains are required to satisfy the following compatibility condition, whenever $\varepsilon < \varepsilon'$:

$$\sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_n = \varepsilon'} f_{\beta^{\varepsilon_1}, \beta^{\varepsilon_2}, \dots, \beta^{\varepsilon_n}}(\Theta_{\varepsilon_1, \varepsilon_2}, \dots, \Theta_{\varepsilon_{n-1}, \varepsilon_n}) = 0. \quad (2.2)$$

Here, $f_{\beta^{\varepsilon_1}, \beta^{\varepsilon_2}, \dots, \beta^{\varepsilon_n}}$ is polygon-counting map associated to $(\beta^{\varepsilon_1}, \beta^{\varepsilon_2}, \dots, \beta^{\varepsilon_n})$ as in [23, Section 8.1]. The sum is taken over Spin^c structures according to the hypercube of Spin^c structures.

We will also write the Floer chain $\Theta_{\varepsilon, \varepsilon'}$ as $\Theta_{\beta^\varepsilon, \beta^{\varepsilon'}}$ when it seems clearest to refer to the sets of curves involved. In the above, we assume that each Heegaard multi-diagram is weakly admissible, and the appropriate finiteness of counts so that the sum makes sense. Assuming weak admissibility, one may always obtain a sensible expression by working over the power series ring $\mathbb{F}[[U]]$. Hypercubes of *alpha* attaching curves are defined by a notational modification of the above definition.

We may *pair* hypercubes of attaching curves, as follows. If $\mathcal{L}_\alpha = (\alpha^\nu, \Theta_{\nu', \nu})_{\nu \in \mathbb{E}_m}$ and $\mathcal{L}_\beta = (\beta^\varepsilon, \Theta_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ are hypercubes of attaching curves for some hypercubes of Spin^c structures, and we have a hypercube of Spin^c structures \mathfrak{S} for the pair $(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ which extends the hypercubes for \mathcal{L}_α and \mathcal{L}_β , then we may form an $(n + m)$ -dimensional hypercube of chain complexes denoted by $CF^-(\mathcal{L}_\alpha, \mathcal{L}_\beta, \mathfrak{S})$. The group at index (ν, ε) is the ordinary Floer complex $CF^-(\Sigma, \alpha^\nu, \beta^\varepsilon, \mathfrak{S}_{\nu, \varepsilon})$. The morphism from (ν, ε) to (ν', ε') is the map

$$D_{(\nu, \varepsilon), (\nu', \varepsilon')}(x) = \sum_{\substack{\nu = \nu_1 < \dots < \nu_k = \nu' \\ \varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'}} f_{\nu_k, \dots, \nu_1, \varepsilon_1, \dots, \varepsilon_j, \mathfrak{S}_{\nu_k, \dots, \nu_1, \varepsilon_1, \dots, \varepsilon_j}} \times (\Theta_{\nu_k, \nu_{k-1}}, \dots, \Theta_{\nu_2, \nu_1}, x, \Theta_{\varepsilon_1, \varepsilon_2}, \dots, \Theta_{\varepsilon_{j-1}, \varepsilon_j}).$$

It is straightforward to use the compatibility condition in (2.2) to see that $CF^-(\mathcal{L}_\alpha, \mathcal{L}_\beta, \mathfrak{S})$ is a hypercube of chain complexes. In specific examples, we will also use the abbreviated notation

$$f_{\alpha^{\nu_1} \rightarrow \dots \rightarrow \alpha^{\nu_k}}^{\beta^{\varepsilon_1} \rightarrow \dots \rightarrow \beta^{\varepsilon_j}}(x) = f_{\nu_k, \dots, \nu_1, \varepsilon_1, \dots, \varepsilon_j, \mathfrak{S}_{\nu_k, \dots, \nu_1, \varepsilon_1, \dots, \varepsilon_j}}(\Theta_{\nu_k, \nu_{k-1}}, \dots, \Theta_{\nu_2, \nu_1}, x, \Theta_{\varepsilon_1, \varepsilon_2}, \dots, \Theta_{\varepsilon_{j-1}, \varepsilon_j}).$$

If $\nu = \nu'$, then we will usually write $f_{\alpha^\nu}^{\beta^\varepsilon \rightarrow \beta^{\varepsilon'}}$, and similarly if $\varepsilon = \varepsilon'$. With respect to this notation the hypercube map becomes

$$D_{(\nu, \varepsilon), (\nu', \varepsilon')}(x) = \sum_{\substack{\nu = \nu_1 < \dots < \nu_k = \nu' \\ \varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'}} f_{\alpha^{\nu_1} \rightarrow \dots \rightarrow \alpha^{\nu_k}}^{\beta^{\varepsilon_1} \rightarrow \dots \rightarrow \beta^{\varepsilon_j}}.$$

We will also frequently substitute the letter h for f in the specific case of a map counting pseudoholomorphic rectangles.

2.3. ι -complexes

We recall the following definition from [5].

Definition 2.5. An ι -complex is a chain complex (C, ∂) which is free and finitely generated over $\mathbb{F}[U]$ and equipped with an endomorphism ι . Furthermore, the following hold.

- (1) C is equipped with a \mathbb{Z} -grading, such that U has grading -2 . We call this grading the *Maslov* or *homological* grading.
- (2) There is a grading preserving isomorphism $U^{-1}H_*(C) \cong \mathbb{F}[U, U^{-1}]$.
- (3) ι is a grading preserving chain map and $\iota^2 \simeq \text{id}$.

Definition 2.6. If (C, ι) and (C', ι') are ι -complexes, we define the group of *enhanced ι -morphisms* to be

$$\underline{\text{Mor}}((C, \iota), (C', \iota')) := \text{Hom}_{\mathbb{F}[U]}(C, C') \oplus \text{Hom}_{\mathbb{F}[U]}(C, C')[1],$$

where $\text{Hom}_{\mathbb{F}[U]}(C, C')$ denotes homogeneous (but not necessarily grading-preserving) morphisms and $[1]$ denotes a grading shift. We define

$$\partial_{\underline{\text{Mor}}}(F, h) = (\partial' F + F\partial, F\iota + \iota' F + \partial'h + h\partial),$$

which makes the category of ι -complexes into a dg-category. An enhanced ι -homomorphism is an enhanced ι -morphism (F, h) satisfying $\partial_{\underline{\text{Mor}}}(F, h) = 0$. An *enhanced ι -homotopy equivalence* is an enhanced ι -homomorphism (F, h) such that there exists an enhanced ι -homomorphism (G, k) with the property that $(G, k) \circ (F, h)$ and $(F, h) \circ (G, k)$ both differ from $(\text{id}, 0)$, the identity enhanced ι -morphism, by a boundary in $\underline{\text{Mor}}$. Composition is given by $(F, h) \circ (F', h') = (F \circ F', F'h + h'F)$.

If (C, ∂) is a free, finitely generated chain complex over $\mathbb{F}[U]$, we now describe the chain map

$$\Phi: C \rightarrow C.$$

Write ∂ as a matrix with respect to a free $\mathbb{F}[U]$ -basis of C . We then define Φ to be the endomorphism obtained by differentiating each entry of this matrix with respect to U and extend $\mathbb{F}[U]$ -linearly. This map appears naturally in considering the basepoint action on Heegaard Floer homology [9, 37, 39]. The map Φ is independent of the choice of basis up to $\mathbb{F}[U]$ -equivariant chain homotopy. In the special case that the basis is the set of generators associated to a particular Heegaard diagram for (Y, w) , such that powers of the variable U in the differential record intersections with the divisor $\{w\} \times \text{Sym}^{g-1}(\Sigma)$, we may write this map as Φ_w .

2.4. Hypercubes and ι -complexes

In this section, we describe a relation between enhanced ι -morphisms and hypercubes. An enhanced ι -morphism $(F, h): (C, \iota) \rightarrow (C', \iota')$ may be encoded into a diagram of the following form:

$$\begin{array}{ccc} C & \xrightarrow{F} & C \\ \downarrow \iota & \searrow h & \downarrow \iota' \\ C & \xrightarrow{F} & C' \end{array} \tag{2.3}$$

The pair (F, h) is an enhanced ι -homomorphism if and only if the above diagram is a hypercube of chain complexes. Composition of ι -morphisms is encoded by stacking and compressing hypercubes.

Remark 2.7. The diagram in equation (2.3) is algebraically equivalent to the following diagram:

$$\begin{array}{ccccc}
 C & \xrightarrow{F} & C' & & \\
 \downarrow Q(\iota + \text{id}) & \nearrow Q \cdot h & \downarrow Q(\iota' + \text{id}) & & \\
 Q \cdot C & \xrightarrow{F} & Q \cdot C' & &
 \end{array}$$

In particular, we may view enhanced iota morphisms between iota complexes as containing equivalent information to an $\mathbb{F}[U, Q]/Q^2$ -equivariant map between the associated complexes over $\mathbb{F}[U, Q]/Q^2$. We will use the two perspectives interchangeably.

Relations between compositions of ι -morphisms are also naturally encoded into hypercubes.

Lemma 2.8. Suppose that (A, ι_A) , (B, ι_B) , (C, ι_C) and (D, ι_D) are ι -complexes, and (F, f) , (G, g) , (H, h) , (I, i) , and (J, j) are enhanced ι -morphisms which fit into the following diagram:

$$\begin{array}{ccccc}
 A & \xrightarrow{I} & B & & \\
 \downarrow \iota_A & \nearrow G & \downarrow \iota_B & \nearrow H & \\
 C & \xrightarrow{F} & D & & \\
 \downarrow \iota_C & \nearrow i & \downarrow \iota_D & \nearrow j & \\
 A & \xrightarrow{I} & B & \xrightarrow{f} & D \\
 \downarrow \iota_A & \nearrow G & \downarrow \iota_C & \nearrow H & \downarrow \iota_D \\
 C & \xrightarrow{F} & D & &
 \end{array}$$

Then, the hypercube relations for the above diagram are equivalent to each of (F, f) , (G, g) , (H, h) and (I, i) being enhanced ι -homomorphisms, as well as the relation

$$(F, f) \circ (G, g) + (H, h) \circ (I, i) = \partial_{\text{Mor}}(J, j).$$

Proof. The length 1 relations imply that F , G , H , and I are chain maps. The length 2 relations along the left, right, front, and back faces are easily seen to be equivalent to (F, f) , (G, g) , (H, h) and (I, i) being enhanced ι -homomorphisms. The length 2 relation along the top and bottom faces is equivalent to

$$FG + HI = [\partial, J].$$

The length 3 relation is equivalent to

$$Fg + fG + Hi + hI = \iota_D J + J\iota_A + [\partial, j].$$

On the other hand, by definition,

$$(F, f) \circ (G, g) + (H, h) \circ (I, i) = (FG + HI, Fg + fG + Hi + hI),$$

while

$$\partial_{\underline{\text{Mor}}}(J, j) = ([\partial, J], \iota_D J + J\iota_A + [\partial, j]).$$

The main claim follows immediately. \blacksquare

2.5. ι_K -complexes and ι_L -complexes

We now review the refinement of the analog of an ι -complex for knots, called an ι_K -complex, and its extension to links. First, we recall some algebraic background. Let (C_K, ∂) be a free, finitely generated complex over the ring $\mathbb{F}[\mathcal{U}, \mathcal{V}]$. There are two naturally associated maps

$$\Phi, \Psi: C_K \rightarrow C_K,$$

constructed as follows. Write ∂ as a matrix with respect to a free $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -basis of C_K . We then define Φ to be the endomorphism obtained by differentiating each entry of this matrix with respect to \mathcal{U} . We define Ψ to be the endomorphism obtained by differentiating each entry with respect to \mathcal{V} . These maps appear naturally in the study of knot Floer homology, see [29, 32, 35]. The maps Φ and Ψ are independent of the choice of basis up to $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -equivariant chain homotopy [33, Corollary 2.9].

An \mathbb{F} -linear map $F: C_K \rightarrow C'_K$ is *skew- $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -equivariant* if

$$F \circ \mathcal{V} = \mathcal{U} \circ F \quad \text{and} \quad F \circ \mathcal{U} = \mathcal{V} \circ F.$$

We recall the following definition from [4]. (Compare also [5, Definition 6.2] and [33, Definition 2.2].)

Definition 2.9. An ι_K -complex (C_K, ∂, ι_K) is a finitely generated, free chain complex (C_K, ∂) over $\mathbb{F}[\mathcal{U}, \mathcal{V}]$, equipped with a skew-equivariant endomorphism ι_K satisfying

$$\iota_K^2 \simeq \text{id} + \Phi\Psi.$$

Remark 2.10. If (C_K, ∂, ι_K) is an ι_K -complex, then ι_K commutes with $U = \mathcal{U}\mathcal{V}$, and hence we can view C_K as an (infinitely generated) complex over $\mathbb{F}[U]$ with an $\mathbb{F}[U]$ -equivariant endomorphism ι_K .

We have the following notions of morphism of ι_K -complexes.

Definition 2.11. Suppose that $\mathcal{C} = (C_K, \partial, \iota_K)$ and $\mathcal{C}' = (C'_K, \partial', \iota'_K)$ are ι_K -complexes.

- (1) An ι_K -homomorphism from \mathcal{C} to \mathcal{C}' consists of an $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -equivariant chain map $F: C_K \rightarrow C'_K$, which satisfies $\iota'_K F + F \iota_K \simeq 0$ (where \simeq denotes skew-equivariantly chain homotopy equivalence).
- (2) The group of *enhanced ι_K -morphisms* is

$$\underline{\text{Mor}}(\mathcal{C}, \mathcal{C}') := \text{Hom}_{\mathbb{F}[\mathcal{U}, \mathcal{V}]}(C_K, C'_K) \oplus \overline{\text{Hom}}_{\mathbb{F}[\mathcal{U}, \mathcal{V}]}(C_K, C'_K)[1, 1],$$

where $\overline{\text{Hom}}_{\mathbb{F}[\mathcal{U}, \mathcal{V}]}(C_K, C'_K)$ denotes the group of $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -skew-equivariant maps. The differential on $\underline{\text{Mor}}(\mathcal{C}_K, \mathcal{C}'_K)$ is given by

$$\partial_{\underline{\text{Mor}}}(F, g) = (F\partial + \partial'F, F\iota_K + \iota'_K F + \partial'g + g\partial).$$

We say (F, g) is an enhanced ι_K -homomorphism if

$$\partial_{\underline{\text{Mor}}}(F, g) = 0.$$

The two enhanced ι_K -morphisms are ι_K -homotopic if their sum is a boundary in $\underline{\text{Mor}}(\mathcal{C}_K, \mathcal{C}'_K)$.

We will also be interested in the case of free, finitely-generated complexes (C_L, ∂) over the ring $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$, which arise when we study links L with multiple components. In this case, there are maps Φ_i and Ψ_i for $1 \leq i \leq \ell$ arising from differentiating with respect to each variable. An \mathbb{F} -linear map $F: C_L \rightarrow C'_L$ is said to be skew-equivariant if it exchanges \mathcal{U}_i and \mathcal{V}_i for each i . We may then consider ι_L -complexes (C_L, ∂, ι_L) , where the map ι_L satisfies the formula

$$\iota_L^2 \simeq (\text{id} + \Phi_\ell \Psi_\ell) \circ \dots \circ (\text{id} + \Phi_1 \Psi_1).$$

This is the diffeomorphism map for performing a Dehn twist on each link component. See [32, Theorem B] for further discussion. The definitions of ι_K -morphisms and their variations now extend straightforwardly to ι_L -complexes.

As in Section 2.4, an enhanced ι_K -homomorphism is equivalent to a hypercubes of chain complexes, compositions of ι_K -morphisms are encoded as hypercubes in the same manner as in Lemma 2.8.

3. Involutive Heegaard Floer homology

In this section, we review Hendricks and Manolescu's construction of involutive Heegaard Floer homology [5], and also define the doubling models which feature in our statements of naturality and functoriality.

3.1. The involutive Floer complexes

We presently recall Hendricks and Manolescu's original construction [5]. Suppose that $\mathcal{H} = (\Sigma, \alpha, \beta, w)$ is a weakly admissible Heegaard diagram for Y . Suppose that \mathfrak{s} is a self-conjugate Spin^c structure. Write $\bar{\mathcal{H}} = (\bar{\Sigma}, \bar{\beta}, \bar{\alpha}, w)$ for the diagram obtained from \mathcal{H} by reversing the orientation of Σ and reversing the roles of α and β . There is a canonical chain isomorphism

$$\eta: CF^-(\bar{\mathcal{H}}, \mathfrak{s}) \rightarrow CF^-(\mathcal{H}, \mathfrak{s}).$$

Hendricks and Manolescu consider the map

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

given by the formula

$$\iota := \eta \circ \Psi_{\mathcal{H} \rightarrow \bar{\mathcal{H}}},$$

where $\Psi_{\mathcal{H} \rightarrow \bar{\mathcal{H}}}$ is the map from naturality as in [9]. They define the involutive Heegaard Floer complex $CFI(Y, \mathfrak{s})$ to be the $\mathbb{F}[U, Q]/(Q^2)$ -chain complex whose underlying $\mathbb{F}[U, Q]/(Q^2)$ -module is

$$CF^-(Y, \mathfrak{s}) \otimes_{\mathbb{F}[U]} \mathbb{F}[U, Q]/(Q^2) = CF^-(Y, \mathfrak{s}) \oplus QCF^-(Y, \mathfrak{s})$$

with differential $\partial_{CFI(Y, \mathfrak{s})} = \partial_{CF^-(Y, \mathfrak{s})} \otimes \text{id} + (\text{id} + \iota) \otimes Q$. That is, the involutive Heegaard Floer complex is the mapping cone

$$\text{Cone}(Q(\text{id} + \iota): CF^-(Y, \mathfrak{s}) \rightarrow QCF^-(Y, \mathfrak{s}))$$

with the evident action of the $\mathbb{F}[U, Q]/(Q^2)$. Hendricks and Manolescu prove that $CFI^-(Y, \mathfrak{s})$ is well-defined up to chain homotopy equivalence.

Hendricks and Manolescu also define a refinement for knots, which we recall presently along with its extension to links. Beginning with the knot case, suppose that K is a null-homologous, oriented knot in a 3-manifold Y , and $\mathfrak{s} \in \text{Spin}^c(Y)$ is self-conjugate. We recall the following standard definition.

Definition 3.1. A Heegaard diagram for a pair (Y, K) consists of a tuple $(\Sigma, \alpha, \beta, w, z)$, such that (Σ, α, β) is a Heegaard diagram for Y , such that the knot K intersects Σ in the points w and z . Furthermore, if U_α and U_β are the closures of the two components of $Y \setminus \Sigma$, then K intersects each of U_α and U_β in a boundary-parallel arc. The standard convention is that K intersects Σ positively at z , and negatively at w .

We will work over the version of knot Floer homology [21, 28] which is freely generated over a 2-variable polynomial ring $\mathbb{F}[\mathcal{U}, \mathcal{V}]$. We denote this chain complex by $\mathcal{CFK}(K)$. See, e.g., [33, Section 3] for background on this version of knot Floer homology.

Let $\phi^+ : (Y, K, w, z) \rightarrow (Y, K, z, w)$ be the diffeomorphism of tuples supported in a neighborhood of K corresponding to a positive half-twist along K , and similarly for ϕ^- and the negative half-twist. There is a canonical chain isomorphism

$$\eta_K : \mathcal{CFK}(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}, z, w) \rightarrow \mathcal{CFK}(\Sigma, \alpha, \beta, w, z),$$

which satisfies $\eta_K(\mathcal{U} \cdot x) = \mathcal{V} \cdot \eta_K(x)$ and $\eta_K(\mathcal{V} \cdot x) = \mathcal{U} \cdot \eta_K(x)$. We define

$$\iota_{K,+} := \eta_K \circ \Psi_{\phi^+(\mathcal{H}) \rightarrow \bar{\mathcal{H}}} \circ \phi_*^+,$$

and we define $\iota_{K,-}$ similarly. Here, ϕ_*^+ denotes the tautological diffeomorphism map, and $\Psi_{\phi^+(\mathcal{H}) \rightarrow \bar{\mathcal{H}}}$ denotes the map from naturality obtained by picking a sequence of Heegaard moves relating $\phi^+(\mathcal{H})$ and $\bar{\mathcal{H}}$. There is a map ι_K^- defined similarly. Hendricks and Manolescu define the involutive knot Floer complex to be the data consisting of $\mathcal{CFK}(K)$ equipped with the endomorphism $\iota_{K,\pm}$, which gives an ι_K -complex.

Remark 3.2. In [5], Hendricks and Manolescu considered only the involution we call $\iota_{K,+}$, which is written simply as ι_K . It is unknown whether there are any knots K where $(\mathcal{CFK}(K), \iota_{K,+})$ is inequivalent to $(\mathcal{CFK}(K), \iota_{K,-})$. The distinction between $\iota_{K,+}$ and $\iota_{K,-}$ has previously been observed in [2].

The case of links is similar. For a link L of $\ell = |L|$ components, we consider Heegaard diagrams $(\Sigma, \alpha, \beta, w, z)$ such that $w = \{w_1, \dots, w_n\}$ and $z = \{z_1, \dots, z_\ell\}$ and the i th component K_i of L intersects Σ positively at z_i and negatively at w_i . The resulting chain complex $\mathcal{CF}\mathcal{L}(L)$ is freely-generated over the polynomial ring $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$. In principle, there are 2^ℓ choices of orientation which may be used to define the involution; however, for simplicity, we assume that an orientation is fixed, and we consider only the two orientations which are either coherent, or opposite to our preferred orientation. It is now straightforward to adapt the construction of $\iota_{K,\pm}$ given above to models for $\iota_{L,\pm}$ for these versions of the link involution.

3.2. Doubled Heegaard diagrams

Suppose that $\mathcal{H} = (\Sigma, \alpha, \beta, w)$ is a pointed Heegaard diagram for Y . We view a regular neighborhood of Σ as $\Sigma \times [0, 1]$. Pick a small disk $D \subseteq \Sigma$, which contains w along its boundary. The submanifold $U_0 = (\Sigma \setminus D) \times [0, 1]$ is a handlebody of genus $2g$, where g denotes $g(\Sigma)$. The boundary of this submanifold is canonically identified with $\Sigma \# \bar{\Sigma}$, where the connected sum occurs at w . Write U_1 for the closure of the complement of U_0 . Clearly, U_1 is also a handlebody of genus $2g$. We write $D(\Sigma) = \Sigma \# \bar{\Sigma}$ for this Heegaard surface.

We may naturally equip $D(\Sigma)$ with attaching curves, as follows. For U_1 , we may use the curves $\alpha \cup \bar{\beta}$, where $\alpha \subseteq \Sigma$ and $\bar{\beta} \subseteq \bar{\Sigma}$. For U_0 , we pick compressing disks by

picking a collection $\delta_1, \dots, \delta_{2g}$ of properly embedded arcs on $\Sigma \setminus N(w)$, called the doubling arcs. We add a basepoint to $\partial N(w)$, for which we also write w . We assume that the arcs $\delta_1, \dots, \delta_{2g}$ are disjoint from w , and also form a basis of $H_1(\Sigma \setminus D, \partial D \setminus \{w\})$. We form attaching curves $\Delta \subseteq \Sigma \# \bar{\Sigma}$ by doubling the curves $\delta_1, \dots, \delta_{2g}$ onto $\Sigma \# \bar{\Sigma}$. We write $D(\mathcal{H}) = (\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \Delta, w)$ for the resulting diagram.

If $\xi = (\xi_1, \xi_2, \xi_3)$ is a framing of $w \in Y$, then the framing is used in the construction as follows. We assume the Heegaard surface is tangent to the 2-plane field $(\xi_1, \xi_2) \subseteq T_w Y$. Additionally, in the tube connecting $\Sigma \# \bar{\Sigma}$, we must place the basepoint w . We place it in the direction corresponding to ξ_1 .

Remark 3.3. The diagram $D(\mathcal{H})$ may also be described by gluing two bordered Heegaard diagrams for genus g handlebodies [16].

3.3. Doubling and the involution

We use doubled Heegaard diagrams to give a natural model of the involution. This model appears in [4, 36]. In this section, we describe the construction.

Firstly, we define the following map as a composition of 1-handle maps

$$F_1^{\bar{\beta}, \bar{\beta}}: CF^-(\Sigma, \alpha, \beta, w) \rightarrow CF^-(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \beta \bar{\beta}, w).$$

Next, we note that the diagram $(\Sigma \# \bar{\Sigma}, \beta \bar{\beta}, \Delta, w)$ is a doubled diagram for $(S^1 \times S^2)^{\# g}$, and hence we may pick a cycle $\Theta_{\beta \bar{\beta}, \Delta}$ generating the top degree of homology. The cycle $\Theta_{\beta \bar{\beta}, \Delta}$ is not unique, but any two choices are related by a boundary. We define the holomorphic triangle map

$$f_{\alpha \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}: CF^-(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \beta \bar{\beta}, w) \rightarrow CF^-(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \Delta, w)$$

by the formula

$$f_{\alpha \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}(x) = f_{\alpha \bar{\beta}, \beta \bar{\beta}, \Delta}(x, \Theta_{\beta \bar{\beta}, \Delta}).$$

Analogously, there is another holomorphic triangle map

$$f_{\alpha \bar{\beta}}^{\Delta \rightarrow \alpha \bar{\alpha}}: CF^-(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \Delta, w) \rightarrow CF^-(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \alpha \bar{\alpha}, w)$$

as well as a 3-handle map

$$F_3^{\alpha, \alpha}: CF^-(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \alpha \bar{\alpha}, w) \rightarrow CF^-(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}, w).$$

Proposition 3.4 ([36, Propositions 7.2 and 7.8]). *The composition $f_{\alpha \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta} \circ F_1^{\bar{\beta}, \bar{\beta}}$ is chain homotopic to the map $\Psi_{\mathcal{H} \rightarrow D(\mathcal{H})}$ from naturality. Dually, the composition $F_3^{\alpha, \alpha} \circ f_{\alpha \bar{\beta}}^{\Delta \rightarrow \alpha \bar{\alpha}}$ is chain homotopic to the map $\Psi_{D(\mathcal{H}) \rightarrow \bar{\mathcal{H}}}$.*

The idea of the proof of the above result is that the composition is topologically equivalent to a composition of the maps for canceling 1-handles and 2-handles, or canceling 2-handles and 3-handles.

Corollary 3.5. *The involution satisfies*

$$\iota \simeq \eta \circ F_3^{\alpha, \alpha} \circ f_{\alpha \bar{\beta}}^{\Delta \rightarrow \alpha \bar{\alpha}} \circ f_{\alpha \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta} \circ F_1^{\bar{\beta}, \bar{\beta}},$$

where η is the canonical chain isomorphism

$$\eta: CF^-(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}, w) \rightarrow CF^-(\Sigma, \alpha, \beta, w).$$

Definition 3.6. Suppose (Y, w, ξ) is a based manifold with a framing $\xi = (\xi_1, \xi_2, \xi_3)$ of $T_w Y$. A *doubling enhanced Heegaard diagram* \mathfrak{D} consists of the following data.

- (1) A Heegaard diagram $\mathcal{H} = (\Sigma, \alpha, \beta, w)$. Furthermore, Σ is positively tangent to the 2-plane spanned by ξ_1 and ξ_2 .
- (2) A set of attaching curves Δ on $\Sigma \# \bar{\Sigma}$, constructed by doubling a basis of arcs $\delta_1, \dots, \delta_{2g}$. Furthermore, we view the endpoint of each δ_i arc as determining an oriented tangent space (oriented to point into δ_i). We assume that none of these oriented tangent spaces coincide with the span of ξ_1 .
- (3) Choices of almost complex structures to compute the four Floer complexes $CF(\Sigma, \alpha, \beta)$, $CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \beta \bar{\beta})$, $CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \Delta)$ and $CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \alpha \bar{\alpha})$, as well as almost complex structures to compute the triangle maps $f_{\alpha \bar{\beta}}^{\Delta \rightarrow \alpha \bar{\alpha}}$ and $f_{\alpha \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}$.

We say that \mathfrak{D} is a *doubling enhancement* of \mathcal{H} . We say that \mathfrak{D} is *admissible* if each of the Heegaard diagrams and triples are weakly admissible.

3.4. Doubling and the knot and link involution

The knot involution of [5] also admits a convenient doubling model, which we introduced in [4] (see also [11, Section 9.4]). In this section we recall this model. In our proof of naturality, we will also need to use several stabilized models, which we introduce in Section 15.

The doubling model for knots has a similar description to Corollary 3.5, but more care will be needed to define the maps. There is furthermore one additional choice to be made, which will correspond to a choice of which direction along K to perform a half-twist.

We begin by describing doubled knot diagrams. If $(\Sigma, \alpha, \beta, w, z)$ is a diagram for (Y, K) , we consider the following two diagrams:

$$D^+(\mathcal{H}) = (\Sigma \#_z \bar{\Sigma}, \alpha \bar{\beta}, \Delta, w, \bar{w}) \quad \text{and} \quad D^-(\mathcal{H}) = (\Sigma \#_w \bar{\Sigma}, \alpha \bar{\beta}, \Delta, \bar{z}, z).$$

We start by analyzing the diagram $\mathcal{D}^+(\mathcal{H})$. There is a map

$$F_{1,+}^{\bar{\beta},\bar{\beta}} : \mathcal{CFK}(\Sigma, \alpha, \beta, w, z) \rightarrow \mathcal{CFK}(\Sigma \#_z \bar{\Sigma}, \alpha \bar{\beta}, \beta \bar{\beta}, w, \bar{w})$$

obtained as the composition of a 1-handle map, as well as a diffeomorphism map for moving z to \bar{w} . As in the setting of closed 3-manifolds, we have holomorphic triangle maps $f_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta}$ and $f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}}$. Finally, there is a map

$$F_{3,+}^{\alpha,\alpha} : \mathcal{CFK}(\Sigma \#_z \bar{\Sigma}, \alpha \bar{\beta}, \alpha \bar{\alpha}, w, \bar{w}) \rightarrow \mathcal{CFK}(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}, \bar{z}, \bar{w}),$$

obtained by moving w to the position of \bar{z} , and then applying the 3-handle maps to remove the α curves. Similar logic to Proposition 3.4 then implies that

$$\iota_{K,+} \simeq \eta_K \circ F_{3,+}^{\alpha,\alpha} \circ f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} \circ f_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} \circ F_{1,+}^{\bar{\beta},\bar{\beta}}.$$

There are analogous maps $F_{1,-}^{\bar{\beta},\bar{\beta}}$ and $F_{3,-}^{\alpha,\alpha}$ which instead involve the diagram $D^-(\mathcal{H})$, such that

$$\iota_{K,-} \simeq \eta_K \circ F_{3,-}^{\alpha,\alpha} \circ f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} \circ f_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} \circ F_{1,-}^{\bar{\beta},\bar{\beta}}.$$

The same procedure with only notational changes may be used to give doubling models for the link involutions $\iota_{L,\pm}$ for links with more than one component. See Section 15 for more details.

3.5. Alpha versus beta doubling

The reader may note that a somewhat arbitrary choice has been made in the definition of ι . Indeed, there is another diagram

$$\bar{D}(\mathcal{H}) = (\Sigma \# \bar{\Sigma}, \Delta, \beta \bar{\alpha}, w),$$

conjugate to the diagram $D(\mathcal{H})$ considered above. Similar to Proposition 3.4 and using the same notation from that proposition, $f_{\alpha\bar{\alpha}}^{\beta\bar{\alpha} \rightarrow \Delta} \circ F_1^{\bar{\alpha},\bar{\alpha}}$ is also chain homotopic to $\Psi_{\mathcal{H} \rightarrow \bar{D}(\mathcal{H})}$. Dually, $F_3^{\beta,\beta} \circ f_{\Delta \rightarrow \beta\bar{\beta}}^{\beta\bar{\alpha}}$ is homotopic to $\Psi_{\bar{D}(\mathcal{H}) \rightarrow \mathcal{H}}$. As a consequence, parallel to Corollary 3.5,

$$\iota \simeq \eta \circ F_3^{\beta,\beta} \circ f_{\Delta \rightarrow \beta\bar{\beta}}^{\beta\bar{\alpha}} \circ f_{\alpha\bar{\alpha}}^{\beta\bar{\alpha} \rightarrow \Delta} \circ F_1^{\bar{\alpha},\bar{\alpha}}. \quad (3.1)$$

We will refer to the composite on the right-hand side of equation (3.1) as the *alpha-doubling model* for ι , and we refer to the model $\eta \circ F_3^{\alpha,\alpha} \circ f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} \circ f_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} \circ F_1^{\beta,\bar{\beta}}$ from Corollary 3.5 as the *beta-doubling model*. In spite of the fact that the alpha-doubling and beta-doubling models are homotopic by naturality, we do not claim there is a canonical homotopy between them. In this paper, we arbitrarily work over the beta-doubling model exclusively.

4. On naturality in Heegaard Floer theory

In this section, we discuss naturality in Heegaard Floer theory and how it adapts to our setting. Our goal in this section is to state a version of the naturality theorem which covers our use of framed basepoints in the involutive Heegaard Floer complex.

We recall that the proof of [9] centers on proving naturality for invariants of sutured 3-manifolds. We recall that if Y is equipped with a basepoint w , then we may obtain a sutured manifold by removing a regular neighborhood of w and adding a single suture (closed loop) to the boundary. The reader should think of the suture as corresponding to the intersection of a Heegaard surface for Y with the boundary.

We cannot apply the work of [9] without any change, because our construction of the involutive Floer complex requires a choice of framing at the basepoint. This is incompatible with the *continuity* axiom of [9], which requires any diffeomorphism of a sutured manifold $\phi \in \text{Diff}^+(M, \gamma)$ which is isotopic to the identity through sutured diffeomorphisms to act by the identity on the complex. In our present context, this is too strong of an axiom to impose, since such diffeomorphisms may correspond to diffeomorphisms of (Y, w) which act non-trivially on the framing of the basepoint. Instead, we wish to weaken the continuity axiom to only require that diffeomorphisms which are isotopic to the identity *relative to the boundary* to act by the identity on our invariant. In this section, we explain how the techniques of [9] adapt to our present situation.

Remark 4.1. Note that in the original setting [9], this was of little importance, since isotopies of sutured manifolds which are supported in a small neighborhood of the boundary do not change the position of the alpha or beta curves.

We begin by recalling some notation from [9]. We first recall the graph $\mathcal{G}_{(M, \gamma)}$ defined therein. The vertices consist of embedded Heegaard diagrams for (M, γ) , with attaching curves taken up to isotopy on the Heegaard surface. The edges of $\mathcal{G}_{(M, \gamma)}$ consist of the following.

- (1) Alpha equivalences on a fixed Heegaard surface.
- (2) Beta equivalences on a fixed Heegaard surface.
- (3) Index (1, 2)-stabilizations.
- (4) Diffeomorphisms of (M, γ) which are isotopic to the identity as sutured diffeomorphisms.

We now define a subgraph $\mathcal{G}_{(M, \gamma)}^\partial \subseteq \mathcal{G}_{(M, \gamma)}$, which we call the graph of *boundary framed Heegaard moves*. This graph has the same vertices as $\mathcal{G}_{(M, \gamma)}$. We add all of the same edges, except that diffeomorphisms are required to be the identity on the boundary of M , and also are required to be isotopic to the identity rel boundary. We consider the following loops of Heegaard diagrams (thought of as 2-cells). These are

a restriction of the distinguished rectangles from [9, Definition 2.32]. These consist of the following loops.

- (1) Commutation of alpha and beta equivalences.
- (2) Commutations of two index $(1, 2)$ -stabilizations.
- (3) Commutations of an index $(1, 2)$ -stabilization and an alpha/beta equivalence.
- (4) Simple handleswap loops.
- (5) A sequence of diffeomorphism edges in $\mathcal{G}_{(M, \gamma)}^\partial$, starting and ending at a fixed Heegaard diagram \mathcal{H} , such that the composition of the diffeomorphisms is a contractible loop in $\text{Diff}(M, \partial M)$.

Proposition 4.2. *After attaching 2-cells to the graph $\mathcal{G}_{(M, \gamma)}^\partial$ corresponding to the loops (1)–(5), we obtain a simply connected space.*

Proof. The result follows from essentially the same argument as [9], with some extra attention paid to the behavior of diffeomorphisms on ∂M . Note that $\mathcal{G}_{(M, \gamma)}^\partial$ is a subcomplex of $\mathcal{G}_{(M, \gamma)}$. One may go through the proof described in [9, Section 8]. We consider a loop of Heegaard moves in $\mathcal{G}_{(M, \gamma)}^\partial$ of k -steps:

$$\mathcal{H}_0 \xrightarrow{e_1} \mathcal{H}_1 \xrightarrow{e_1} \mathcal{H}_2 \xrightarrow{e_2} \cdots \xrightarrow{e_k} \mathcal{H}_k = \mathcal{H}_0.$$

We now build a map $S^1 \rightarrow \mathcal{F}\mathcal{V}_{\geq 1}(M, \gamma)$, denoting the space of pairs (f, v) of sutured functions and gradient like vector fields which lie in codimension 0 or 1. (See [9] for this notation). This is constructed the same as therein. Namely, for each diagram \mathcal{H}_i , we pick a compatible pair $(f, v) \in \mathcal{G}\mathcal{V}_0(M, \gamma)$ inducing this pair. For each move of Heegaard diagrams, we pick a compatible 1-parameter family of pairs $(f_t, v_t) \in \mathcal{G}\mathcal{V}_{\leq 1}(M, \gamma)$. Accordingly, one constructs a polyhedral decomposition of D^2 , such that each edge contains at most one codimension 1 singularity, and each 2-cell contains at most one codimension 2 singularity. From here, one argues that after subdividing it is sufficient to consider only 2-cells whose edges compose to give one of the loops (1)–(5).

The difference between our situation and [9] is that we must argue that it is sufficient to consider only sutured isotopies which fix the boundary pointwise. One must do this for both edges and for 2-cells. Consider first a 1-cell:

$$\sigma_1: D^1 \rightarrow \mathcal{F}\mathcal{V}_0(M, \gamma).$$

Given such a path, one may construct a family of embedded Heegaard surfaces. To see that these are related by an isotopy, one uses a version of the isotopy extension lemma, as stated in [9, Lemma 6.19]. Note that this lemma actually constructs an isotopy which is the identity on ∂M , because it is obtained by integrating a time dependent family of vector fields on (M, γ) which is supported on the interior of (M, γ) .

Given a map $\sigma_2: D^2 \rightarrow \mathcal{F}\mathcal{V}_0(M, \gamma)$, the above construction, applied to $S^1 = \partial D^2$, defines a loop of diffeomorphisms $\phi_\theta: S^1 \rightarrow \text{Diff}(M, \partial M)$. Since the corresponding path in $\mathcal{F}\mathcal{V}_0(M, \gamma)$ is contractible, we may extend this loop ϕ_θ over D^2 . We observe that [9, Lemma 6.20] constructs a null-homotopy of the corresponding loop of diffeomorphisms. We observe that this null-homotopy is again constructed by integrating along a time dependent family of vector fields which is supported in $\text{int}(M)$, and hence the corresponding family of diffeomorphisms are the identity on ∂M . Hence, the boundary framed loops in (1)–(5) are sufficient to contract an arbitrary loop. ■

5. Maps for elementary handleslide equivalences

In this section, we define the transition maps for changes of the doubling data which fix the Heegaard surface Σ . These correspond to changes of the α , β and Δ curves by handleslides and isotopies. Furthermore, in this section, we consider only a restricted set of alpha and beta equivalences, as follows.

Definition 5.1. (1) Suppose that (Σ, \mathbf{w}) is a Heegaard surface with $n \geq 1$ basepoints and γ and γ' are attaching curves on Σ . We say that γ and γ' are *handleslide equivalent* if they may be related by a sequence of handleslides and isotopies.

(2) Suppose Σ is a Heegaard surface and γ and γ' are handleslide equivalent attaching curves on Σ . We say that γ and γ' are related by an *elementary equivalence* if $(\Sigma, \gamma', \gamma, \mathbf{w})$ is a weakly admissible diagram, and $\mathbb{T}_{\gamma'} \cap \mathbb{T}_\gamma$ consists of exactly $2^{|\gamma|-1}$ generators.

Suppose (Σ, α, β) is a Heegaard diagram, and α' and β' are handleslide equivalent to α and β , respectively. Suppose further that Δ and Δ' are choices of doubling curves. Let \mathfrak{D} denote the data consisting of (Σ, α, β) , the curves Δ , and appropriate choices of almost complex structures to compute the involution. Let \mathfrak{D}' denote the analogous data with α' , β' and Δ' . We say that \mathfrak{D}' is obtained from \mathfrak{D} by an *elementary equivalence* if α' and β' are obtained from α and β by elementary equivalences, and also if the tuple

$$(\Sigma \# \bar{\Sigma}, \alpha' \bar{\alpha}', \alpha \bar{\alpha}, \beta \bar{\beta}, \beta' \bar{\beta}', \Delta, \Delta', \mathbf{w})$$

is weakly admissible.

In the case that \mathfrak{D}' is obtained from \mathfrak{D} by an elementary equivalence, we define the transition map

$$\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}: CFI(\mathfrak{D}) \rightarrow CFI(\mathfrak{D}')$$

as the compression of the hyperbox of chain complexes shown in Figure 5.1. Therein, the rows with the 1-handle maps are constructed similarly to [4, Section 14], where

$$\begin{array}{ccccc}
CF(\alpha, \beta) & \longrightarrow & CF(\alpha', \beta) & \longrightarrow & CF(\alpha', \beta') \\
\downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}', \bar{\beta}} & & \downarrow F_1^{\bar{\beta}', \bar{\beta}'} \\
CF(\alpha \bar{\beta}, \beta \bar{\beta}) & \longrightarrow & CF(\alpha' \bar{\beta}', \beta \bar{\beta}) & \longrightarrow & CF(\alpha' \bar{\beta}', \beta' \bar{\beta}') \\
\downarrow & \searrow & \downarrow & \searrow & \downarrow \\
CF(\alpha \bar{\beta}, \Delta) & \longrightarrow & CF(\alpha' \bar{\beta}', \Delta) & \longrightarrow & CF(\alpha' \bar{\beta}', \Delta') \\
\downarrow & \searrow & \downarrow & \searrow & \downarrow \\
CF(\alpha \bar{\beta}, \alpha \bar{\alpha}) & \longrightarrow & CF(\alpha' \bar{\beta}', \alpha \bar{\alpha}) & \longrightarrow & CF(\alpha' \bar{\beta}', \alpha' \bar{\alpha}') \\
\downarrow F_3^{\alpha, \alpha} & & \downarrow F_3^{\alpha', \alpha} & & \downarrow F_3^{\alpha', \alpha'} \\
CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}', \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}', \bar{\alpha}') \\
\downarrow \text{id} & \searrow & & & \downarrow \text{id} \\
CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}', \bar{\alpha}') & \longrightarrow & CF(\bar{\beta}', \bar{\alpha}')$$

Figure 5.1. The transition map for an elementary equivalence. An additional row involving the canonical map $\eta: CF(\bar{\beta}, \bar{\alpha}) \rightarrow CF(\alpha, \beta)$ is omitted.

they are called *hypercubes of stabilization*. The construction here requires that the α' be obtained from α by an elementary equivalence, and similar for β and β' , so that the degeneration argument from [4, Section 10] applies. The rows with 3-handles are similar. The remaining subcubes are obtained by pairing hypercubes of attaching curves.

Note that in Figure 5.1, we have omitted a final level involving the tautological map $\eta: CF(\bar{\beta}, \bar{\alpha}) \rightarrow CF(\alpha, \beta)$. Throughout the paper, we will usually omit this row.

If \mathfrak{D} and \mathfrak{D}' are handleslide equivalent (but perhaps not related by an elementary equivalence), we pick an arbitrary sequence

$$\mathfrak{D} = \mathfrak{D}_1, \dots, \mathfrak{D}_n = \mathfrak{D}'$$

such that \mathfrak{D}_{i+1} and \mathfrak{D}_i are related by an elementary equivalence for all i . We define

$$\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} := \Psi_{\mathfrak{D}_{n-1} \rightarrow \mathfrak{D}_n} \circ \dots \circ \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_2}. \quad (5.1)$$

A key step in our proof of naturality will be proving the following theorem.

Theorem 5.2. *Suppose that \mathfrak{D} , \mathfrak{D}' and \mathfrak{D}'' are doubling enhanced Heegaard diagrams which all have the same underlying Heegaard surface Σ and are handleslide equivalent.*

- (1) *The map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ defined in (5.1) is independent of intermediate sequence of doubling enhanced Heegaard diagrams, up to $\mathbb{F}[U, Q]/Q^2$ -equivariant chain homotopy.*
- (2) $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}} \simeq \text{id}_{CFI(\mathfrak{D})}$.
- (3) $\Psi_{\mathfrak{D}' \rightarrow \mathfrak{D}''} \circ \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} \simeq \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}''}$.

Part (2) of Theorem 5.2 is proven in the subsequent Section 5.1. Parts (1) and (3) are more technical and are proven in Sections 7 and 8; in particular, they follow from Proposition 8.1.

Remark 5.3. We are only able to define the map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ using Figure 5.1 when \mathfrak{D} and \mathfrak{D}' are related by an elementary equivalence because we are not able to construct the 1-handle hypercubes if \mathfrak{D} and \mathfrak{D}' are related by an arbitrary handleslide equivalence. To show the map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ is independent of the choice of intermediate diagrams, we will define a more complicated transition map $\tilde{\Psi}_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ in Section 7.3.

5.1. Continuity

In this section, we prove the continuity axiom from [9, Definition 2.32] for our maps associated to a handleslide equivalence. In our present context, the continuity axiom amounts to the following.

Proposition 5.4. *Let \mathfrak{D} denote an admissible doubling enhanced Heegaard diagram for (Y, w, ξ) . Let $g: (\Sigma, w) \rightarrow (\Sigma, w)$ be a pointed diffeomorphism which is smoothly isotopic to the identity through diffeomorphisms which are the identity on a neighborhood of w . Let $g(\mathfrak{D})$ be the push-forward of the diagram \mathfrak{D} under g , and let*

$$CFI(g): CFI(\mathfrak{D}) \rightarrow CFI(g(\mathfrak{D}))$$

denote the tautological map. Then

$$\Psi_{\mathfrak{D} \rightarrow g(\mathfrak{D})} \simeq CFI(g).$$

Proof. We adapt the argument of [9, Proposition 9.27] to the involutive case by also using the small translate theorems for triangles and quadrilaterals from our previous paper [4]. Write $(\Sigma, \alpha, \beta, w)$ for the underlying Heegaard diagram of \mathfrak{D} , and let Δ be the doubling curves. As a first step, pick a C^∞ small diffeomorphism g_0 of (Σ, w) and write $\alpha' = g_0(\alpha)$, and define β' and Δ' similarly. Assume that $|\alpha'_i \cap \alpha_j| = 2\delta_{ij}$, where δ_{ij} denotes the Kronecker delta, and assume the analogous statement holds for

β' , β , Δ' and Δ . It follows from the small translate theorems [4, Propositions 11.1 and 11.5] that if g_0 is chosen suitably small, then we may pick an almost complex structure on $\Sigma \times \Delta$ which interpolates between a fixed J on the (Σ, α, β) cylindrical end and $(g_0)_*(J)$ on the (Σ, α', β) cylindrical end such that the map $f_{\alpha', \alpha, \beta}(\Theta_{\alpha', \alpha}, x)$ counts only small triangles, so that $f_{\alpha', \alpha, \beta}(\Theta_{\alpha', \alpha}, x) = x_{np}$ for all $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Here, x_{np} denotes the nearest point of $\mathbb{T}_{\alpha'} \cap \mathbb{T}_\beta$ to x . In a similar manner, the small translate theorem for quadrilaterals implies that all of the length 2 maps counted by $\Psi_{\mathcal{D} \rightarrow g_0(\mathcal{D})}$ vanish (as no holomorphic curves are counted). In particular, $\Psi_{\mathcal{D} \rightarrow g_0(\mathcal{D})}$ coincides with the nearest point map.

Similarly, if g'_0 is C^∞ close to g_0 , then the same argument as the small translate theorems also implies that $\Psi_{\mathcal{D} \rightarrow g'_0(\mathcal{D})}$ and $\Psi_{g'_0(\mathcal{D}) \rightarrow \mathcal{D}}$ coincide with the nearest point maps.

More generally, if g is C^∞ small relative to g_0 , then we can decompose

$$\Psi_{\mathcal{D} \rightarrow g(\mathcal{D})} \simeq \Psi_{g_0(\mathcal{D}) \rightarrow g(\mathcal{D})} \circ \Psi_{\mathcal{D} \rightarrow g_0(\mathcal{D})}$$

and apply the nearest point results to the latter two maps to identify $\Psi_{\mathcal{D} \rightarrow g(\mathcal{D})}$ with $CFI(g)$. This establishes the claim when g is C^∞ -small. Since both sides of the desired equality are functorial under composition, and any diffeomorphism $g: (\Sigma, w) \rightarrow (\Sigma, w)$ which is smoothly isotopic to the identity may be decomposed into a composition of C^∞ -small diffeomorphisms, this proves the claim in general. ■

6. Holomorphic polygons and stabilizations

In this section, we describe some technical results concerning stabilizations of Heegaard diagrams and holomorphic polygons. The results we cover are a generalization of those considered in [4, Section 9].

Definition 6.1. Suppose that $\mathcal{D}_0 = (\Sigma_0, \gamma_1, \dots, \gamma_n, w)$ is a weakly admissible Heegaard tuple and \mathfrak{S} is a set of Spin^c structures on $X_{\gamma_1, \dots, \gamma_n}$, which is closed under the action of $\delta H^1(Y_{\gamma_i, \gamma_j})$ for $i < j$.

- (1) We say that $(\mathcal{D}_0, \mathfrak{S})$ is a *multi-stabilizing diagram* if the following are satisfied. All elements of \mathfrak{S} restrict to a single element of $\text{Spin}^c(\partial X_{\gamma_1, \dots, \gamma_n})$ which is furthermore torsion. Additionally, we require that if $\mathfrak{s}_1, \mathfrak{s}_2 \in \mathfrak{S}$, then $d(\mathfrak{s}_1) = d(\mathfrak{s}_2)$, where $d(\mathfrak{s}) = (c_1^2(\mathfrak{s}) - 2\chi - 3\sigma)/4$.
- (2) We say $(\mathcal{D}_0, \mathfrak{S})$ is *algebraically rigid* if $\partial = 0$ on $\widehat{CF}(\Sigma_0, \gamma_i, \gamma_{i+1}, \mathfrak{s}_{i,i+1}, w)$, where $\mathfrak{s}_{i,i+1}$ is the restriction to $Y_{\gamma_i, \gamma_{i+1}}$ of the elements of \mathfrak{S} .

Example 6.2. Let \mathcal{D}_0 be a Heegaard triple where all γ_i are pairwise related by small Hamiltonian translates. Let $\mathfrak{S} = \{\mathfrak{s}_0\}$, where \mathfrak{s}_0 is the unique Spin^c structure which

has torsion restriction to the boundary. Then, $(\mathcal{D}_0, \mathfrak{S})$ is a multi-stabilizing diagram. For suitable choices of translates, \mathcal{D}_0 is algebraically rigid.

Remark 6.3. When the set \mathfrak{S} is understood from context (such as the case when there is a unique Spin^c structure on $X_{\gamma_1, \dots, \gamma_n}$ which has torsion restriction to $\partial X_{\gamma_1, \dots, \gamma_n}$), we will usually refer to the diagram \mathcal{D}_0 as being a multi-stabilizing diagram.

In this section, we state a helpful result about stabilizations and holomorphic curves. Before we state our result, we introduce some notation.

As a first step, we recall some notation and basic facts about Stasheff's associahedron, K_n . We view K_n as a convex polytope which models the compactification of $n + 1$ marked points on the boundary of a disk. It is well known that K_n admits an embedding into a Euclidean space and has the homology of a point. There is a well-known cell structure on K_n , giving a convenient model for its homology, which we denote $C_*^{\text{cell}}(K_n)$. The cells of this decomposition are in bijective correspondence with planar trees with n inputs, one output, and no internal vertices of valence less than 3. If T is such a tree, then the degree of the corresponding cell is

$$n - 1 - \#(\text{internal vertices}).$$

The differential of a tree T is the sum of all ways of breaking an internal vertex into two vertices, both with valence at least 3. The degree 0 cells correspond to trees with only valence 3 internal vertices. There is a codimension 0 cell of degree $n - 2$, which has just one internal vertex.

Suppose that $\mathcal{D} = (\Sigma, \gamma_1, \dots, \gamma_n, w)$ is a weakly admissible Heegaard tuple and $J = (J_x)_{x \in K_{n-1}}$ is a stratified family of almost complex structures on $\Sigma \times D_n$ for computing holomorphic n -gons. Suppose also that $\mathfrak{S} \subseteq \text{Spin}^c(X_{\gamma_1, \dots, \gamma_n})$ is a set which is closed under the actions of $\delta^1 H^1(Y_{\gamma_i, \gamma_j})$ for all $i < j$, and assume for simplicity that all elements of \mathfrak{S} restrict to a single element of $\text{Spin}^c(\partial X_{\gamma_1, \dots, \gamma_n})$. Given a tree $T \in C_*^{\text{cell}}(K_{n-1})$, there is a map

$$\hat{f}_{\mathcal{D}, \mathfrak{S}; J, T} : \widehat{CF}(\gamma_1, \gamma_2, \mathfrak{s}_{1,2}) \otimes \cdots \otimes \widehat{CF}(\gamma_{n-1}, \gamma_n, \mathfrak{s}_{n-1,n}) \rightarrow \widehat{CF}(\gamma_1, \gamma_n, \mathfrak{s}_{1,n}),$$

obtained by composing the polygon maps corresponding to each internal vertex of T . If T has only valence 3 internal vertices, then there is a unique point of K_{n-1} corresponding to T , which is in the strata of maximal codimension. We write J_T for the corresponding almost complex structure, and otherwise suppress the tree from the notation in this case. The map $\hat{f}_{\mathcal{D}, \mathfrak{S}; J, T}$ is obtained by successively composing holomorphic triangle maps according to the tree T .

More generally, we can view the holomorphic polygon maps as fitting together compatibly to give a chain map

$$\hat{f}_{\mathcal{D}, \mathfrak{S}; J} : C_*^{\text{cell}}(K_{n-1}) \otimes \widehat{CF}(\gamma_1, \gamma_2, \mathfrak{s}_{1,2}) \otimes \cdots \otimes \widehat{CF}(\gamma_{n-1}, \gamma_n, \mathfrak{s}_{n-1,n}) \rightarrow \widehat{CF}(\gamma_1, \gamma_n, \mathfrak{s}_{1,n}).$$

The same statement holds for CF^- as long as we are more careful about finiteness of curve counts, either by restricting Spin^c structures or working over $\mathbb{F}[[U]]$. We call this version of the map $f_{\mathcal{D}, \mathfrak{S}; J}$ and its specialization to a particular tree $f_{\mathcal{D}, \mathfrak{S}; J, T}$ analogously.

Lemma 6.4. *Suppose that $\mathcal{D}_0 = (\Sigma_0, \gamma_1, \dots, \gamma_n, \mathbf{w})$ is a diagram and \mathfrak{S} is a set of Spin^c structures such that $(\mathcal{D}_0, \mathfrak{S})$ is algebraically rigid and multi-stabilizing. Let T_1 and T_2 be two degree 0 trees. Then,*

$$\hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_{T_1}}(\mathbf{x}_1, \dots, \mathbf{x}_{n-1}) = \hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_{T_2}}(\mathbf{x}_1, \dots, \mathbf{x}_{n-1})$$

for any $\mathbf{x}_i \in \mathbb{T}_{\gamma_i} \cap \mathbb{T}_{\gamma_{i+1}}$.

Proof. The trees T_1 and T_2 are homologous as elements of $C_*^{\text{cell}}(K_{n-1})$, since

$$H_*^{\text{cell}}(K_{n-1}) \cong \mathbb{Z}.$$

Hence, the maps

$$\hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_{T_1}}, \hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_{T_2}}: \widehat{CF}(\gamma_1, \gamma_2, \mathfrak{s}_{1,2}) \otimes \dots \otimes \widehat{CF}(\gamma_{n-1}, \gamma_n, \mathfrak{s}_{n-1,n}) \rightarrow \widehat{CF}(\gamma_1, \gamma_n, \mathfrak{s}_{1,n})$$

are chain homotopic. The differentials on domain and codomain of the two maps vanish, so any two chain homotopic maps between them are equal. \blacksquare

Proposition 6.5. *Let $\mathcal{D} = (\Sigma, \delta_1, \dots, \delta_n, \mathbf{w})$ and $\mathcal{D}_0 = (\Sigma_0, \gamma_1, \dots, \gamma_n, \mathbf{w}_0)$ be Heegaard n -diagrams, where $n > 2$. Assume $\mathfrak{s} \in \text{Spin}^c(X_{\delta_1, \dots, \delta_n})$ and $\mathfrak{S} \subseteq \text{Spin}^c(X_{\gamma_1, \dots, \gamma_n})$ and assume that $(\mathcal{D}_0, \mathfrak{S})$ is an algebraically rigid multi-stabilizing Heegaard diagram. From the connected sum $\mathcal{D} \# \mathcal{D}_0$ by adding a tube from \mathcal{D}_0 at some $w_0 \in \mathbf{w}_0$ to a point on \mathcal{D} . Let T be a degree 0 tree representing the generator of $C_*^{\text{cell}}(K_{n-1})$. Suppose $J = (J_x)_{x \in K_{n-1}}$ and $J^0 = (J_x^0)_{x \in K_{n-1}}$ are stratified families of almost complex structures on $\Sigma \times D_n$ and $\Sigma_0 \times D_n$, respectively, for counting holomorphic n -gons. Let $\theta_1, \dots, \theta_{n-1}$ be homogeneously graded elements (necessarily cycles) of $\widehat{CF}(\Sigma_0, \gamma_1, \gamma_2, \mathfrak{s}_{1,2}), \dots, \widehat{CF}(\Sigma_0, \gamma_{n-1}, \gamma_n, \mathfrak{s}_{n-1,n})$, respectively, and suppose that*

$$\mathbf{y} = \hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_T^0}(\theta_1, \dots, \theta_{n-1})$$

is non-zero. Then, for any tree T'

$$\begin{aligned} f_{\mathcal{D} \# \mathcal{D}_0, \mathfrak{s} \# \mathfrak{S}; (J \wedge J^0), T'}(\mathbf{x}_1 \times \theta_1, \dots, \mathbf{x}_{n-1} \times \theta_{n-1}) \\ = f_{\mathcal{D}, \mathfrak{S}; J, T'}(\mathbf{x}_1, \dots, \mathbf{x}_{n-1}) \otimes \mathbf{y} + \sum_{\substack{\mathbf{z} \in \mathbb{T}_{\gamma_1} \cap \mathbb{T}_{\gamma_n} \\ \text{gr}(\mathbf{z}) > \text{gr}(\mathbf{y})}} \mathbf{q}_{\mathbf{z}} \otimes \mathbf{z}, \end{aligned}$$

where $\mathbf{q}_{\mathbf{z}}$ are elements of $CF^-(\Sigma, \delta_1, \delta_n)$.

Proof. The proof follows from similar reasoning to [4, Propositions 10.2 and 10.6], as we now sketch. See also [25, Section 6.1]. Since we are using almost complex structures which are singular at the connected sum point, the relevant moduli spaces are fibered products over $D_n^k \times K_{n-1}$, where k is the multiplicity of a class at the connected sum point. Consider a class $\psi \wedge \psi_0$, where $\psi_0 \in \pi_2(\theta_1, \dots, \theta_{n-1}, y_0)$, for some intersection point y_0 , and such that $s_{w_0}(\psi_0) \in \mathfrak{S}$. The assumption that $\hat{f}_{\gamma_1, \dots, \gamma_n, \mathfrak{S}; J_T^0}(\theta_1, \dots, \theta_{n-1}) = y$ is non-zero implies that there exists a class of triangles $\psi' \in \pi_2(\theta_1, \dots, \theta_{n-1}, y)$ which has $s_{w_0}(\psi') \in \mathfrak{S}$, $n_{w_0}(\psi') = 0$ and $\mu(\psi') = 0$. Since all elements of \mathfrak{S} have the same degree, we conclude that

$$\text{gr}(y_0) - 2n_{w_0}(\psi_0) + \mu(\psi_0) = \text{gr}(y).$$

Using Sarkar's formula for the Maslov index, we see that

$$\mu(\psi \wedge \psi_0) = \mu(\psi) + \mu(\psi_0) - 2n_{w_0}(\psi_0).$$

(The only difference between the formulas for $\mu(\psi \wedge \psi_0)$ and $\mu(\psi) + \mu(\psi_0)$ is in the Euler measure, which is corrected on the right-hand side by $2n_{w_0}(\psi_0)$). Hence, combining the above two formulas, we obtain

$$\mu(\psi \wedge \psi_0) = \mu(\psi) + \text{gr}(y) - \text{gr}(y_0) + n_{w_0 \setminus \{w_0\}}(\psi_0).$$

The main claim concerns only classes, where $\text{gr}(y) - \text{gr}(y_0) \geq 0$. As the map

$$f_{\mathcal{D}^{\#} \mathcal{D}_0, \mathfrak{s} \# \mathfrak{s}_0; J \wedge J^0, T'}$$

counts index $3 - n$ curves, we assume that

$$\mu(\psi) + \text{gr}(y) - \text{gr}(y_0) + n_{w_0 \setminus \{w_0\}}(\psi_0) = 3 - n.$$

By transversality for curves representing ψ , we may assume that $\mu(\psi) \geq 3 - n$. Hence, if $\text{gr}(y) - \text{gr}(y_0) \geq 0$, then we must have $\mu(\psi) = 3 - n$, $\text{gr}(y) = \text{gr}(y_0) = 0$ and $n_{w_0 \setminus \{w_0\}}(\psi_0) = 0$. In particular, the unconstrained moduli space $\mathcal{M}(\psi)$ is 0-dimensional. It suffices to understand the moduli spaces on the Σ_0 side.

Let $k \geq 0$ be a fixed integer, and assume that $\text{gr}(y) = \text{gr}(y_0)$. Write $\mathcal{M}_k(\theta_1, \dots, \theta_{n-1}, y_0)$ for the moduli space of curves representing a class in $\pi_2(\theta_1, \dots, \theta_{n-1}, y_0)$ with multiplicity k at w_0 . We now claim that

$$\text{ev}: \mathcal{M}_k(\theta_1, \dots, \theta_{n-1}, y_0) \rightarrow D_n^k \times K_{n-1}$$

is odd degree if y_0 is a summand of y in $\widehat{CF}(\gamma_1, \gamma_n)$, and is even degree if y_0 is not a summand of y . This clearly implies the claim.

Let $(d, x) \in D_n^k \times K_{n-1}$ be generic, and pick a path $\gamma: [0, \infty) \rightarrow D_n^k \times K_{n-1}$, such that $\gamma(0) = (d, x)$ and $\gamma(t)$ has the following behavior as $t \rightarrow \infty$. As $t \rightarrow \infty$,

we assume that $x(t)$ approaches the tree T , viewed as a point in ∂K_{n-1} , while the k -tuple of points $\mathbf{d}(t)$ travel towards one of the boundary punctures of D_n^k . Furthermore, we assume that under the identification of this end as a half cylinder $[0, 1] \times [0, \infty)$, they approach some fixed $\mathbf{d}' \subseteq ((0, 1) \times \mathbb{R})^k$, modulo the \mathbb{R} -action on $[0, \infty)$. We consider the 1-dimensional moduli space of curves which have $\text{ev}(u) \in \text{im } \gamma$. The generic degenerations of this moduli space at finite t correspond exactly to index 1 disks breaking off of the cylindrical ends. These cancel modulo 2, since the differentials vanish on each $\widehat{CF}(\mathbf{y}_i, \mathbf{y}_{i+1}, \mathbf{s}_{i,i+1})$. The ends appearing as $t \rightarrow \infty$ correspond to points in the following set:

$$\left(\coprod_{\substack{\phi \in \pi_2(\theta_1, \theta_1) \\ \mu(\phi) = 2k}} \mathcal{M}(\phi, \mathbf{d}') \right) \times \left(\coprod_{\substack{\psi_0 \in \pi_2(\theta_1, \dots, \theta_{n-1}, \mathbf{y}_0) \\ \mu(\psi_0) = 0 \\ n_{\mathbf{w}_0}(\psi_0) = 0}} \mathcal{M}_{J_T}(\psi_0) \right).$$

The left-hand factor has odd cardinality by [36, equation (31)]. The count of the right-hand side is exactly the \mathbf{y}_0 coefficient of $\hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_T^0}(\theta_1, \dots, \theta_{n-1})$, completing the proof. \blacksquare

6.1. Remarks and special cases

Proposition 6.5 generalizes [4, Propositions 10.2 and 10.6], which concern triangles and quadrilaterals. In this section we describe some special cases.

Remark 6.6. The results of [4, Propositions 10.2 and 10.6] are stated in terms of an *index tuple*, which does not appear in our Proposition 6.5. The stabilization formulas from [4] are stated only in the case that one of the index tuple coordinates was 0. The entry of the index tuple being zero implied a specific formula for $\mathbf{y} = \hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_T^0}(\theta_1, \dots, \theta_n)$. See [4, Lemma 10.4]. In our present paper, we also need the stabilization formulas in cases when the entry of the index tuple is non-zero, but where we know the value of $\hat{f}_{\mathcal{D}_0, \mathfrak{S}; J_T^0}(\theta_1, \dots, \theta_n)$.

Example 6.7. Suppose that \mathcal{D} and \mathcal{D}_0 are as in the statement of Proposition 6.5, and suppose further that \mathcal{D}_0 has the property that each \mathbf{y}_i is obtained from \mathbf{y}_1 by a small Hamiltonian translation of each of the curves. Then, Proposition 6.5 implies that

$$f_{\mathcal{D} \# \mathcal{D}_0, \mathfrak{s} \# \mathfrak{s}_0; J \wedge J^0, T'}(\mathbf{x}_1 \times \theta_{1,2}^+, \dots, \mathbf{x}_{n-1} \times \theta_{n-1,n}^+) = f_{\mathcal{D}, \mathfrak{s}; J, T'}(\mathbf{x}_1, \dots, \mathbf{x}_{n-1}) \otimes \Theta_{1,n}^+.$$

Similarly,

$$\begin{aligned} & f_{\mathcal{D} \# \mathcal{D}_0, \mathfrak{s} \# \mathfrak{s}_0; J \wedge J^0, T'}(\mathbf{x}_1 \times \theta_{1,2}^-, \mathbf{x}_2 \times \theta_{2,3}^+, \dots, \mathbf{x}_{n-1} \times \theta_{n-1,n}^+) \\ &= f_{\mathcal{D}, \mathfrak{s}; J, T'}(\mathbf{x}_1, \dots, \mathbf{x}_{n-1}) \otimes \theta_{1,n}^- + \sum_{\substack{\Theta \in \mathbb{T}_{\gamma_1} \cap \mathbb{T}_{\gamma_n} \\ \text{gr}(\Theta) > \text{gr}(\theta_{1,n}^-)}} \mathbf{q}_\Theta \otimes \Theta. \end{aligned}$$

7. An expanded model of the involution

In this section, we describe a *basepoint-expanded* doubling model of the involution, which will allow us to understand the transition maps for non-elementary handleslide equivalences. An outline of this section is as follows. In Section 7.1, we define the notion of a *basepoint expanded, doubling enhanced Heegaard diagram* $\tilde{\mathfrak{D}}$, which can be used to construct an $\mathbb{F}[U, Q]/Q^2$ -complex $CFI(\tilde{\mathfrak{D}})$. In Section 7.3, we construct transition maps

$$\tilde{\Psi}_{\tilde{\mathfrak{D}}_1 \rightarrow \tilde{\mathfrak{D}}_2} : CFI(\tilde{\mathfrak{D}}_1) \rightarrow CFI(\tilde{\mathfrak{D}}_2)$$

in the case that $\tilde{\mathfrak{D}}_1$ and $\tilde{\mathfrak{D}}_2$ are two basepoint expanded, doubling enhanced Heegaard diagrams which have the same underlying Heegaard surface, and which satisfy a weak admissibility condition. Unlike in Section 5, we do not require $\tilde{\mathfrak{D}}_1$ and $\tilde{\mathfrak{D}}_2$ to be related by an elementary handleslide. In the subsequent Section 8, we relate the complexes $CFI(\tilde{\mathfrak{D}})$ and the transition maps $\tilde{\Psi}_{\tilde{\mathfrak{D}}_1 \rightarrow \tilde{\mathfrak{D}}_2}$ to the non-expanded models described earlier. The key motivation for considering the maps $\tilde{\Psi}_{\tilde{\mathfrak{D}}_1 \rightarrow \tilde{\mathfrak{D}}_2}$ is that they may be defined for non-elementary handleslide equivalences in a manner which clearly is independent of additional choices.

7.1. Doubling with extra basepoints

We now describe our expanded model of the involution, which uses an extra basepoint. The presence of the extra basepoint simplifies some of the gluing arguments.

Definition 7.1. A *basepoint expanded, doubling enhanced Heegaard diagram* $\tilde{\mathfrak{D}}$ consists of a singly pointed Heegaard diagram

$$\mathcal{H} = (\Sigma, \alpha, \beta, w),$$

a small disk $D \subseteq \Sigma$, containing w along its boundary, a choice of point w' , as well as a collection of $2g + 2$ attaching curves $\Delta \subseteq \Sigma \# \bar{\Sigma}$, as follows. The curves Δ are constructed by doubling a basis of pairwise disjoint, properly embedded arcs $\delta_1, \dots, \delta_{2g+2}$ on $\Sigma \setminus D$, which avoid w and w' and which form a basis of $H_1(\Sigma \setminus D, \partial D \setminus \{w, w'\})$. Additionally, $\tilde{\mathfrak{D}}$ contains the choice of almost complex structures used to compute the holomorphic triangle maps.

Using a basepoint expanded diagram $\tilde{\mathfrak{D}}$, as above, we obtain a model of the involution by modifying the formula in Corollary 3.5, as follows. We let c and c' be small perturbations of the circle ∂D , as in Figure 7.1.

We define a 1-handle map

$$F_1^{c\bar{\beta}, c\bar{\beta}} : CF(\Sigma, \alpha, \beta, w) \rightarrow CF(\Sigma \# \bar{\Sigma}, \alpha c \bar{\beta}, \beta c \bar{\beta}, w, w')$$

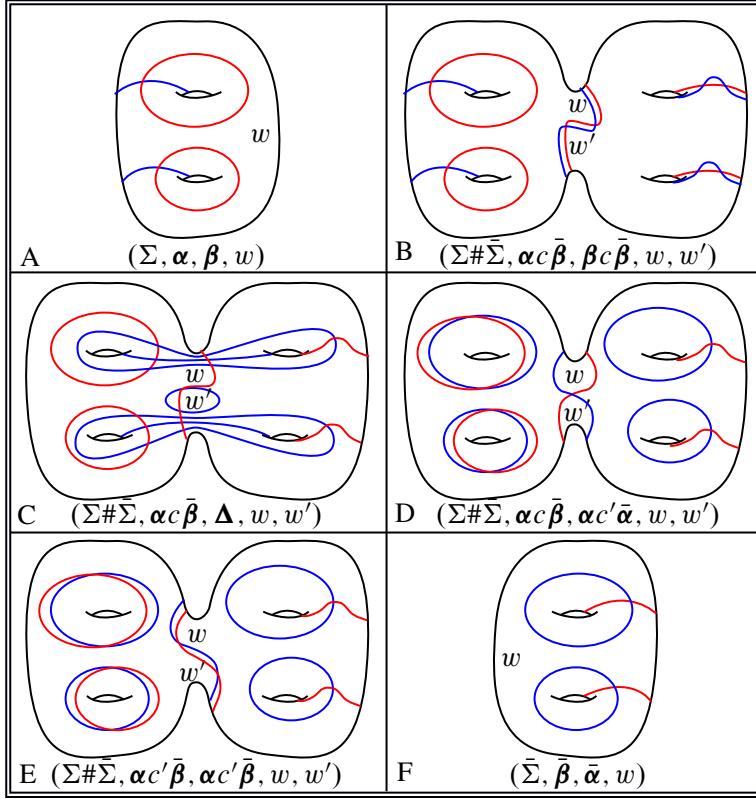


Figure 7.1. The basepoint expanded doubling model of the involution. (Frames are read in alphabetical order).

by tensoring with the top degree generator. This map adds the basepoint w' . Similarly, we define a map

$$F_3^{\alpha c', \alpha c'} : CF(\Sigma \# \bar{\Sigma}, \alpha c' \bar{\beta}, \alpha c' \bar{\alpha}, w, w') \rightarrow CF(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}, w)$$

by using the same formula as the standard 3-handle map.

We define the basepoint-expanded model for the involution via the formula

$$\iota := \eta \circ F_3^{\alpha c', \alpha c'} \circ f_{\alpha c \bar{\beta} \rightarrow \alpha c' \bar{\beta}}^{\alpha c' \bar{\alpha}} \circ A_\lambda \circ f_{\alpha c \bar{\beta}}^{\Delta \rightarrow \alpha c' \bar{\alpha}} \circ f_{\alpha c \bar{\beta}}^{\beta c \bar{\beta} \rightarrow \Delta} \circ F_1^{\beta c \bar{\beta}, c \bar{\beta}}.$$

In the above equation, A_λ denotes the relative homology map for an arc λ connecting w and w' in the connected sum region. This map is discussed further in the subsequent Section 7.2.

Remark 7.2. Similar to the non-expanded model, the expanded model also naturally depends on our framing $\xi = (\xi_1, \xi_2, \xi_3)$ of the basepoint. We assume that the original

Heegaard surface Σ is positively tangent to (ξ_1, ξ_2) , that w is in the direction of ξ_1 and w' is in the direction of ξ_2 , in the tube region. Furthermore, we assume that λ is chosen so that if we identify $\text{Span}(\xi_1, \xi_2)$ with \mathbb{C} , where $\xi_1 = 1$ and $\xi_2 = i$, and we identify the intersection with the tube region of $\Sigma \# \bar{\Sigma}$ with $S^1 \subseteq \mathbb{C}$, then λ corresponds to e^{it} for $t \in [0, \pi/2]$.

7.2. Homology actions on hypercubes

In this section, we describe how to construct a homology action on a hypercube of Floer chain complexes. We can construct both a version for closed curves on Σ , as well as a relative one for arcs on Σ with boundary in \mathbf{w} . We focus on the relative case.

We recall the construction from [39, Section 5]. If $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a multi-pointed Heegaard diagram, and λ is a path on Σ connecting $w_1, w_2 \in \mathbf{w}$, there is an endomorphism

$$A_\lambda: CF(\mathcal{H}, \mathbf{s}) \rightarrow CF(\mathcal{H}, \mathbf{s}), \quad (7.1)$$

which satisfies

$$[\partial, A_\lambda] = U_{w_1} + U_{w_2}. \quad (7.2)$$

In our present paper, we work with just one U -variable, so A_λ becomes a chain map. The map A_λ is defined as follows. If $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is a class of disks, we may define a quantity $a(\lambda, \phi) \in \mathbb{Z}$ by summing the changes of the multiplicities of ϕ across the alpha curves as one travels along λ . (Compare [19].) Sometimes, it is helpful to write $a(\lambda, \phi) = (\partial_\alpha D(\phi)) \cdot \lambda$, where $\partial_\alpha D(\phi)$ is the alpha boundary of the domain of ϕ . The map A_λ is defined via the formula

$$A_\lambda(\mathbf{x}) := \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} a(\lambda, \phi) \#(\mathcal{M}(\phi)/\mathbb{R}) U^{n_{w_1}(\phi) + \dots + n_{w_n}(\phi)} \cdot \mathbf{y}.$$

The map A_λ is sometimes a helpful tool when adding and removing basepoints. We now describe a version of A_λ for hypercubes. Suppose we are given two hypercubes of attaching curves \mathcal{L}_α and \mathcal{L}_β on a multi-pointed Heegaard diagram, each consisting of handleslide equivalent attaching curves, such that the length 1 morphisms are all the top degree generators of $CF^-(\#^n(S^1 \times S^2))$. We now suppose that λ is an arc which connects two basepoints, w_1 and w_2 . Consider the pairing $CF(\mathcal{L}_\alpha, \mathcal{L}_\beta)$, where we identify the variables of the basepoints w_1 and w_2 to a single U .

We now construct a morphism of hypercubes

$$\mathcal{A}_\lambda: CF(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow CF(\mathcal{L}_\alpha, \mathcal{L}_\beta).$$

Recall that a morphism of hypercubes can itself be thought of as a hypercube of dimension 1 larger than $\dim CF(\mathcal{L}_\alpha, \mathcal{L}_\beta)$. Our morphism of hypercubes \mathcal{A}_λ will

have the property that the length one morphisms coincide with the chain maps A_λ from (7.1). We will show that \mathcal{A}_λ is well defined up to chain homotopy of hypercube morphisms, and is natural with respect to restriction to sub-cubes of $CF(\mathcal{L}_\alpha, \mathcal{L}_\beta)$.

We note that we may formally view λ as an input for a holomorphic polygon counting map by setting

$$f_{\gamma_1, \dots, \gamma_j, \gamma_j, \dots, \gamma_m}(\mathbf{x}_{1,2}, \dots, \mathbf{x}_{j-1,j}, \lambda, \mathbf{x}_{j,j+1}, \dots, \mathbf{x}_{m-1,m}) \\ := \sum_{\substack{\psi \in \pi_2(\mathbf{x}_{0,1}, \dots, \mathbf{x}_{m-1,m}, \mathbf{z}) \\ \mu(\psi) = 3-m}} (\partial_{\gamma_j} D(\psi) \cdot \lambda) \# \mathcal{M}(\psi) U^{n_{w_1}(\psi) + \dots + n_{w_m}(\psi)} \cdot \mathbf{z},$$

where $\mathcal{M}(\psi)$ denotes the moduli space of holomorphic m -gons (in the usual sense).

As particular examples,

$$f_{\alpha, \alpha, \beta}(\lambda, \mathbf{x}) = A_\lambda(\mathbf{x}) \quad \text{and} \quad f_{\alpha, \beta, \beta}(\mathbf{x}, \lambda) = B_\lambda(\mathbf{x}),$$

where A_λ counts holomorphic disks weighted by $\#(\partial_\alpha(\phi) \cap \lambda)$, while B_λ counts disks weighted by $\#(\partial_\beta(\phi) \cap \lambda)$.

Lemma 7.3. *The holomorphic polygon counting maps, defined with an extra input λ , satisfy the standard associativity rule as long as one uses the convention that $\partial\lambda = U_{w_1} + U_{w_2}$, together with conventions that the polygon maps are U_{w_i} -equivariant and strictly unital (i.e., they vanish if they have 1 as an input when there are more than 2 total inputs). In particular, if we work over a single U variable, they satisfy the associativity relations with the convention that $\partial\lambda = 0$.*

Proof. The proof differs based on whether $\ell = 2$ or $\ell > 2$. If $\ell = 2$, then the result follows from equation (7.2). For the case that $\ell > 2$, one counts the ends of index $(3 - \ell)$ moduli spaces weighted by $\partial_{\gamma_i} D(\psi) \cdot \lambda$ and quickly obtains the result. ■

The construction of \mathcal{A}_λ is as follows. We will formally construct a cube-shaped diagram $\mathcal{L}_\alpha^\lambda$ of dimension $\dim \mathcal{L}_\alpha + 1$. To begin, we take two copies of \mathcal{L}_α , for which we write $\mathcal{L}_\alpha \times \{0\}$ and $\mathcal{L}_\alpha \times \{1\}$, and formally adjoin additional length 1 arrows from $\mathcal{L}_\alpha \times \{0\}$ to $\mathcal{L}_\alpha \times \{1\}$, which we label by the character λ .

We begin by constructing the length 2 chains of $\mathcal{L}_\alpha^\lambda$, via the following argument. We will show that the 2-dimensional faces can be constructed to have the following form:

$$\begin{array}{ccc} \alpha & \xrightarrow{\Theta_{\alpha', \alpha}} & \alpha' \\ \downarrow \lambda & \searrow U\eta_{\alpha', \alpha} & \downarrow \lambda \\ \alpha & \xrightarrow{\Theta_{\alpha', \alpha}} & \alpha' \end{array} \tag{7.3}$$

(We have not yet defined $\eta_{\alpha',\alpha}$.) The desired relation is that

$$A_\lambda(\Theta_{\alpha',\alpha}) + A'_\lambda(\Theta_{\alpha',\alpha}) = \partial U \eta_{\alpha',\alpha}. \quad (7.4)$$

Here, A_λ counts changes of a disk class across α , while A'_λ counts changes across α' . Note that by [6, Section 3.5], we have that

$$A_\lambda(\Theta_{\alpha',\alpha}) + A'_\lambda(\Theta_{\alpha',\alpha}) = U(\Phi_{w_1} + \Phi_{w_2})(\Theta_{\alpha',\alpha}), \quad (7.5)$$

where Φ_{w_i} denotes the map

$$\Phi_{w_i}(x) = U^{-1} \sum_{\substack{\phi \in \pi_2(x,y) \\ \mu(\phi)=1}} n_{w_i}(\phi) \#(\mathcal{M}(\phi)/\mathbb{R}) U^{(n_{w_1} + \dots + n_{w_n})(\phi)},$$

extended $\mathbb{F}[U]$ -equivariantly.

We note that $\Phi_{w_i}(\Theta_{\alpha',\alpha})$ is a cycle of grading 1 higher than $\Theta_{\alpha',\alpha}$, so it is null-homologous, since $\Theta_{\alpha',\alpha}$ is the top degree generator. Let $\eta_{\alpha',\alpha}$ be any chain of homogeneous grading one higher than $\Theta_{\alpha',\alpha}$ so that

$$\partial \eta_{\alpha',\alpha} = (\Phi_{w_1} + \Phi_{w_2})(\Theta_{\alpha',\alpha}).$$

We perform this construction to build all of the length 2 chains in $\mathcal{L}_\alpha^\lambda$. (If we did not quotient by $U_{w_1} + U_{w_2}$, then we would instead obtain a length 2 morphism of the form $U_1 \eta_{\alpha',\alpha}^1 + U_2 \eta_{\alpha',\alpha}^2$.)

We now construct the higher length chains in the cube $\mathcal{L}_\alpha^\lambda$. Their construction is similar to the construction of the length 2 chains. We assume, by induction, that each length $m > 1$ chain is of the form $U \eta_{\alpha',\alpha}$. (This is only necessary when $m \in \{2, 3\}$, but is satisfiable for all $m > 1$).

Let $\varepsilon < \varepsilon'$ be points \mathbb{E}_d , where

$$d = \dim \mathcal{L}_\alpha,$$

and suppose that all chains of length less than $|\varepsilon' - \varepsilon|_{L^1}$ have already been defined. Set

$$\begin{aligned} C_{\varepsilon',\varepsilon} := & \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} \sum_{i=0}^j f_{\alpha^{\varepsilon_j}, \dots, \alpha^{\varepsilon_1}} (\Theta_{\alpha^{\varepsilon_j}, \alpha^{\varepsilon_{j-1}}, \dots, \Theta_{\alpha^{\varepsilon_{i+1}}, \alpha^{\varepsilon_i}}, \lambda, \\ & \quad \Theta_{\alpha^{\varepsilon_i}, \alpha^{\varepsilon_{i-1}}, \dots, \Theta_{\alpha^{\varepsilon_2}, \alpha^{\varepsilon_1}}}) \\ & + \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} \sum_{i=0}^j f_{\alpha^{\varepsilon_j}, \dots, \alpha^{\varepsilon_1}} (\Theta_{\alpha^{\varepsilon_j}, \alpha^{\varepsilon_{j-1}}, \dots, \Theta_{\alpha^{\varepsilon_{i+1}}, \alpha^{\varepsilon_i}}, U \eta_{\alpha^{\varepsilon_i}, \alpha^{\varepsilon_{i-1}}, \\ & \quad \Theta_{\alpha^{\varepsilon_{i-1}}, \alpha^{\varepsilon_{i-2}}, \dots, \Theta_{\alpha^{\varepsilon_2}, \alpha^{\varepsilon_1}}}). \end{aligned} \quad (7.6)$$

However, by an entirely analogous argument to (7.5), we have

$$\begin{aligned}
C_{\varepsilon',\varepsilon} = & \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} f_{\alpha^{\varepsilon_n}, \dots, \alpha^{\varepsilon_1}}^{n w_1 + n w_2} (\Theta_{\alpha^{\varepsilon_j}, \alpha^{\varepsilon_{j-1}}, \dots, \Theta_{\alpha^{\varepsilon_2}, \alpha^{\varepsilon_1}}}) \\
& + \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} \sum_{i=0}^j f_{\alpha^{\varepsilon_j}, \dots, \alpha^{\varepsilon_1}} (\Theta_{\alpha^{\varepsilon_j}, \alpha^{\varepsilon_{j-1}}, \dots, \Theta_{\alpha^{\varepsilon_{i+1}}, \alpha^{\varepsilon_i}}, U \eta_{\alpha^{\varepsilon_i}, \alpha^{\varepsilon_{i-1}}, \dots, \Theta_{\alpha^{\varepsilon_{i-2}}, \alpha^{\varepsilon_1}}}),
\end{aligned} \tag{7.7}$$

where $f_{\alpha^{\varepsilon_n}, \dots, \alpha^{\varepsilon_1}}^{n w_1 + n w_2}$ counts holomorphic polygons with a factor of $n w_1(\psi) + n w_2(\psi)$. In particular, $C_{\varepsilon',\varepsilon} = U C'_{\varepsilon',\varepsilon}$, for some chain $C'_{\varepsilon',\varepsilon}$. Since the hypercube relations are satisfied for all proper faces, the element $C'_{\varepsilon',\varepsilon}$ is a cycle. Since the complexes are free $\mathbb{F}[U]$ -modules, this implies that $C'_{\varepsilon',\varepsilon}$ is also a cycle. However, $C_{\varepsilon',\varepsilon}$ has grading $m-2$ higher than the top degree generator, where $m = |\varepsilon' - \varepsilon|_{L^1}$. Hence, $C'_{\varepsilon',\varepsilon}$ has grading m higher than the top degree generator. As long as $m \geq 1$, we conclude that $C'_{\varepsilon',\varepsilon}$ is a boundary. We set $\eta_{\alpha^{\varepsilon'}, \alpha^{\varepsilon}}$ to be any chain of homogeneous grading $m+1$ such that

$$\partial \eta_{\alpha^{\varepsilon'}, \alpha^{\varepsilon}} = C'_{\varepsilon',\varepsilon}.$$

Finally, we define the hypercube corresponding to the morphism \mathcal{A}_λ to be the pairing of $CF(\mathcal{L}_\alpha^\lambda, \mathcal{L}_\beta)$, where the hypercube maps are defined by using the normal formulas for the maps in a pairing, while allowing the chains λ to be inputs.

Lemma 7.4. (1) *The diagram $\mathcal{L}_\alpha^\lambda$ satisfies the hypercube relations (interpreted in the way described above).*

- (2) *The map \mathcal{A}_λ satisfies $\partial_{\text{Mor}}(\mathcal{A}_\lambda) = 0$ (i.e., \mathcal{A}_λ is a cycle).*
- (3) *The map \mathcal{A}_λ is independent, up to chain-homotopy of hypercube morphisms, from the choices of chains used in its construction.*

Proof. The first claim follows immediately from the construction. The second claim follows quickly from Gromov compactness interpreted via Lemma 7.3.

We consider the final claim. This is proven as follows. Suppose we have two models of $\mathcal{L}_\alpha^\lambda$, which are both constructed using the above procedure. We consider a partially constructed hypercube of dimension $d+2$, given by the following diagram:

$$\begin{array}{ccc}
\mathcal{L}_\alpha & \xrightarrow{\lambda_1} & \mathcal{L}_\alpha \\
\downarrow \text{id} & \searrow * & \downarrow \text{id} \\
\mathcal{L}_\alpha & \xrightarrow{\lambda_2} & \mathcal{L}_\alpha
\end{array}$$

Here, λ_1 denotes the chains in one model of $\mathcal{L}_\alpha^\lambda$, while λ_2 denotes the chains constructed in the other model, and the $*$ arrows are yet to be defined. The hypercube relations are satisfied on the four d -dimensional faces of the above diagram. We wish to define chains which increment both of the displayed coordinates, such that the hypercube relations are satisfied. The construction of such arrows (corresponding to the $*$ arrow above) follows nearly verbatim from the construction of the cube $\mathcal{L}_\alpha^\lambda$ itself. Pairing $\mathcal{L}_\alpha^{\lambda_1 \rightarrow \lambda_2}$ with \mathcal{L}_β gives the homotopy between the two models of \mathcal{A}_λ , completing the proof. \blacksquare

Remark 7.5. (1) If one works over the ring $\mathbb{F}[U_{w_1}, \dots, U_{w_n}]$, one obtains instead that $\partial_{\text{Mor}}(\mathcal{A}_\lambda) = (U_{w_1} + U_{w_2}) \cdot \text{id}$, where $\partial\lambda = U_{w_1} + U_{w_2}$.

(2) The simplest case of the map \mathcal{A}_λ is when \mathcal{L}_α is 0-dimensional. Then, $\mathcal{L}_\alpha^\lambda$ is just the formal hypercube

$$\alpha \xrightarrow{\lambda} \alpha.$$

If \mathcal{L}_β is a beta hypercube, then the map

$$\mathcal{A}_\lambda: CF(\alpha, \mathcal{L}_\beta) \rightarrow CF(\alpha, \mathcal{L}_\beta)$$

requires no additional choices.

(3) Naturally, one may also build a hypercube $\mathcal{L}_\beta^\lambda$, similar to above, and define

$$\mathcal{B}_\lambda: CF(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow CF(\mathcal{L}_\alpha, \mathcal{L}_\beta)$$

as the hypercube $CF(\mathcal{L}_\alpha, \mathcal{L}_\beta^\lambda)$.

7.3. Maps for general equivalences on the expanded model

We now define the naturality maps for a general alpha or beta equivalence in terms of the expanded model of the involution. Suppose that $\tilde{\mathcal{D}}$ and $\tilde{\mathcal{D}'}$ are two basepoint expanded doubling enhanced Heegaard diagrams, which share the same underlying Heegaard surface Σ . Write $\mathcal{H} = (\Sigma, \alpha, \beta, w)$ and $\mathcal{H}' = (\Sigma, \alpha', \beta', w)$ for the underlying non-involutive Heegaard diagrams of $\tilde{\mathcal{D}}$ and $\tilde{\mathcal{D}'}$, respectively. Write Δ and Δ' for the choices of doubling curves. Furthermore, we assume that the diagram

$$(\Sigma \# \bar{\Sigma}, \alpha', \bar{\alpha}', \alpha, \bar{\alpha}, \beta, \bar{\beta}, \beta', \bar{\beta}', c, c', \Delta, \Delta', w, w')$$

is weakly admissible. (This may be achieved by winding α, α', β and β' sufficiently).

Remark 7.6. Note that the above is not a valid Heegaard diagram since none of the attaching curve sets have $2g(\Sigma) + 1$ curves, however, it still makes sense to require weak-admissibility. Indeed, the definition of weak-admissibility is that if P is a non-zero integral 2-chain with boundary equal to a linear combination of curves on the

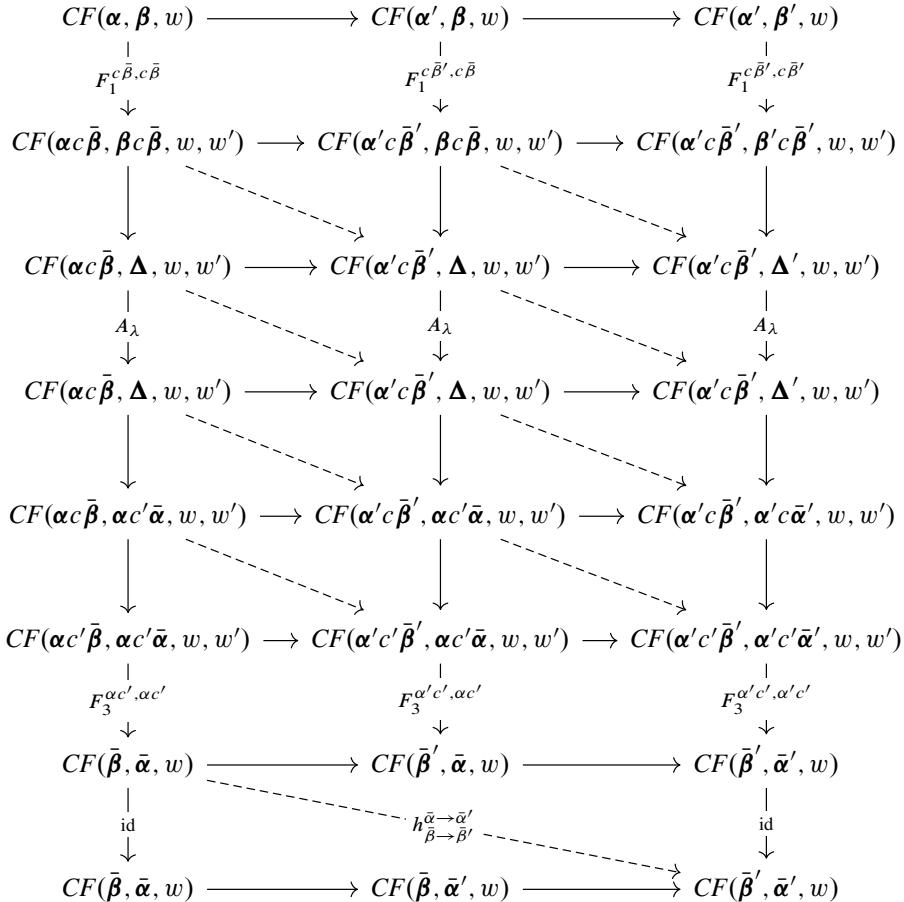


Figure 7.2. The hyperbox whose compression is the basepoint expanded transition map $\tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}}'}$.

diagram and $n_w(P) = n_{w'}(P) = 0$, then P has both positive and negative multiplicities.

Our transition map $\tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}}'}$ is defined by the hyperbox in Figure 7.2. We now describe aspects of the construction in more detail. The second and third levels are obtained by pairing hyperboxes of attaching curves in the obvious manner. The fourth level is obtained from the argument in Section 7.2, where we take λ to be a path in the neck region which connects the two basepoints w and w' .

We now consider the levels involving the 1-handle and 3-handle maps. The map $F_1^{c \beta', c \bar{\beta}}$ is different than the other 1-handle maps we have seen so far, since we do not

require the diagram $(\Sigma, \bar{\beta}, \bar{\beta}')$ to be a standard diagram for $(S^1 \times S^2)^{\#g}$. Instead, we only require the diagram to be admissible. We define the top level as the compression of the following hyperbox:

$$\begin{array}{ccccc}
CF(\alpha, \beta, w) & \longrightarrow & CF(\alpha', \beta, w) & \longrightarrow & CF(\alpha', \beta', w) \\
\downarrow \otimes \Theta_{\bar{\beta}, \bar{\beta}} & & \downarrow \otimes \Theta_{\bar{\beta}', \bar{\beta}} & & \downarrow \otimes \Theta_{\bar{\beta}', \bar{\beta}'} \\
CF(\alpha, \beta, w) \otimes CF(\bar{\beta}, \bar{\beta}, w') & \rightarrow & CF(\alpha', \beta, w) \otimes CF(\bar{\beta}', \bar{\beta}, w') & \rightarrow & CF(\alpha', \beta', w) \otimes CF(\bar{\beta}', \bar{\beta}', w') \\
\downarrow F_1^{c,c} & & \downarrow F_1^{c,c} & & \downarrow F_1^{c,c} \\
CF(\alpha c \bar{\beta}, \beta c \bar{\beta}, w, w') & \longrightarrow & CF(\alpha' c \bar{\beta}', \beta c \bar{\beta}, w, w') & \longrightarrow & CF(\alpha' c \bar{\beta}', \beta' c \bar{\beta}', w, w')
\end{array} \tag{7.8}$$

The horizontal arrows in equation (7.8) are holomorphic triangle maps. In the middle, they are the tensor product of the holomorphic triangle maps on the two relevant Heegaard triples. (Or equivalently, the map which counts holomorphic triangles in the symplectic manifold $(\Sigma \sqcup \bar{\Sigma}) \times D_3$, where D_3 denotes a disk with three boundary punctures.)

Lemma 7.7. *The diagram in equation (7.8) satisfies the hyperbox relations.*

Proof. The fact that the maps $\otimes \Theta_{\bar{\beta}, \bar{\beta}}$ (and so forth) are chain maps follows from the fact that $\Theta_{\bar{\beta}, \bar{\beta}}$ is a cycle. The maps $F_1^{c,c}$ are chain maps by [39, Proposition 8.5]. The fact that the top left and right squares commute is a consequence of the fact that

$$f_{\bar{\beta}', \bar{\beta}, \bar{\beta}}(\Theta_{\bar{\beta}', \bar{\beta}}, \Theta_{\bar{\beta}, \bar{\beta}}) = \Theta_{\bar{\beta}', \bar{\beta}} \quad \text{and} \quad f_{\bar{\beta}', \bar{\beta}, \bar{\beta}'}(\Theta_{\bar{\beta}', \bar{\beta}}, \Theta_{\bar{\beta}, \bar{\beta}'}) = \Theta_{\bar{\beta}', \bar{\beta}'}.$$

The left relation follows from a small triangle argument [4, Proposition 11.1], while the right one follows from the grading considerations. Commutativity of the bottom two squares follows from the stabilization results for triangles [39, Theorem 8.8]. ■

Remark 7.8. We remark that our main motivation for using the expanded model of involution is to simplify the construction of the 1-handle and 3-handle hypercubes in the transition map. Indeed, the 1-handle hypercube at the top level of equation (5.1) is challenging to construct for general β and β' . The challenge is that in Heegaard Floer theory, degenerating the holomorphic ℓ -gon maps along a connected sum neck in the Heegaard diagram yields moduli spaces which are fibered products over an evaluation map to $\text{Sym}^n(D_\ell) \times K_{\ell-1}$, where $K_{\ell-1}$ denotes an associahedron. Here, n denotes the multiplicity at the connected sum point. The factor of $\text{Sym}^n(D_\ell)$ records asymptotics of holomorphic curves at the connected sum point. The factor of $K_{\ell-1}$ records a choice of almost complex structure on D_ℓ . The construction of the 1-handle

hypercube amounts to a problem of deforming the diagonal in

$$\text{Sym}^n(D_\ell) \times \text{Sym}^n(D_\ell) \times K_{\ell-1} \times K_{\ell-1}.$$

However, in the expanded model of the involution, we can use stabilization results like Proposition 6.5 to reduce the problem to considering diagonals in $K_{\ell-1} \times K_{\ell-1}$, which is a simpler problem and is sufficient for our purposes.

8. Relating the expanded and ordinary models of the involution

In this section, we relate the expanded model of the involution from Section 7 with the ordinary model from Section 3.2. If \mathfrak{D} is a doubling enhanced Heegaard diagram, and $\tilde{\mathfrak{D}}$ is a basepoint expanded doubling diagram, such that the underlying (non-involutive) Heegaard diagrams for \mathfrak{D} and $\tilde{\mathfrak{D}}$ coincide, we will construct a chain homotopy equivalence

$$F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}: \text{CFI}(\mathfrak{D}) \rightarrow \text{CFI}(\tilde{\mathfrak{D}}).$$

Proposition 8.1. (1) *If \mathfrak{D} and \mathfrak{D}' are ordinary doubling diagrams which differ by an elementary equivalence, then*

$$\tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}}'} \circ F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}} + F_{\mathfrak{D}' \rightarrow \tilde{\mathfrak{D}}} \circ \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} \simeq 0.$$

(2) *If $\tilde{\mathfrak{D}}$, $\tilde{\mathfrak{D}'}$, and $\tilde{\mathfrak{D}''}$ are three basepoint expanded doubling enhancements of three handleslide equivalent diagrams, then*

$$\tilde{\Psi}_{\tilde{\mathfrak{D}'} \rightarrow \tilde{\mathfrak{D}''}} \circ \tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}'}} \simeq \tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}''}}.$$

The definition of the map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$ is in Section 8.2. The proof of part (1) of Proposition 8.1 is in Section 8.3. Finally, we prove part (2) in Section 8.4.

8.1. Preliminaries

In this section, we review a few maps that will appear in the construction of $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$, and also construct an important hypercube.

Firstly, we recall the *free-stabilization* maps from [39, Section 6]. If $\mathcal{H} = (\Sigma, \alpha, \beta, w)$ is a Heegaard diagram, and w' is a point in $\Sigma \setminus (\alpha \cup \beta \cup w)$, we may form a new diagram $\mathcal{H}^+ = (\Sigma, \alpha \cup \alpha_0, \beta \cup \beta_0, w \cup \{w'\})$ by adding attaching curves α_0 and β_0 which are contained in a small ball centered at w' . We assume that α_0 and β_0 intersect in two points, and furthermore, are disjoint from $\alpha \cup \beta \cup w$.

In the above setting, there are chain maps

$$S_w^+: CF^-(\mathcal{H}, \mathfrak{s}) \rightarrow CF^-(\mathcal{H}^+, \mathfrak{s}) \quad \text{and} \quad S_{w'}^-: CF^-(\mathcal{H}^+, \mathfrak{s}) \rightarrow CF^-(\mathcal{H}, \mathfrak{s})$$

called the free-stabilization maps. They are given by the formulas

$$S_{w'}^+(x) = x \times \theta^+, \quad S_{w'}^-(x \times \theta^+) = 0, \quad S_{w'}^-(x \times \theta^-) = x,$$

extended equivariantly over $\mathbb{F}[U]$.

If λ is an arc which connects w' to another basepoint in w , then

$$S_{w'}^- A_\lambda S_{w'}^+ = \text{id}, \quad (8.1)$$

for appropriately chosen almost complex structure. This is proven in [39, Lemma 7.10] by a model computation.

We now generalize equation (8.1) to the setting of hypercubes. Suppose that $\mathcal{H} = CF(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ is a hypercube formed by pairing hypercubes of attaching curves \mathcal{L}_α and \mathcal{L}_β . We form new hypercubes, \mathcal{L}_α^+ and \mathcal{L}_β^+ , by adding small translates of α_0 and β_0 to each attaching curve of \mathcal{L}_α and \mathcal{L}_β . We form the morphisms of \mathcal{L}_α^+ by tensoring each morphism of \mathcal{L}_α with $\theta_{\alpha_0, \alpha_0}^+$. We construct \mathcal{L}_β^+ similarly. The hypercube relations for \mathcal{L}_α^+ and \mathcal{L}_β^+ follow from Proposition 6.5. Finally, we define a new hypercube of chain complexes, $\mathcal{H}^+ = CF(\mathcal{L}_\alpha^+, \mathcal{L}_\beta^+)$.

We assume, additionally, that \mathcal{L}_α is a hypercube of handleslide equivalent alpha attaching curves, and that the length 1 morphisms are all cycles representing the top degree generator of homology. We construct the morphism

$$\mathcal{A}_\lambda: CF(\mathcal{L}_\alpha^+, \mathcal{L}_\beta^+) \rightarrow CF(\mathcal{L}_\alpha^+, \mathcal{L}_\beta^+)$$

as in Section 7.2.

We consider the following diagram:

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{\text{id}} & \mathcal{H} \\ \downarrow & & \downarrow S_{w'}^+ \\ & & \mathcal{H}^+ \\ \downarrow \text{id} & & \downarrow \mathcal{A}_\lambda \\ & & \mathcal{H}^+ \\ \downarrow & & \downarrow S_{w'}^- \\ \mathcal{H} & \xrightarrow{\text{id}} & \mathcal{H} \end{array} \quad (8.2)$$

The maps $S_{w'}^+$ and $S_{w'}^-$ denote hypercube versions of the free-stabilization maps; that is, they are hypercube morphisms with only length 1 maps, which are the ordinary free-stabilization maps on the Floer complexes. In the language of [4], these are *hypercubes of stabilization*.

If we compress the right-hand side of (8.2) we obtain a cube-shaped diagram. In the following, we verify that the resulting diagram is a hypercube of chain complexes.

Proposition 8.2. *Let \mathcal{H} and \mathcal{H}^+ be as above. Then, \mathcal{A}_λ may be chosen so that the diagram in (8.2) becomes a hypercube of chain complexes once we compress the right-hand side. In the hypercubes along the right-hand side, we use almost complex structures which are nodal near the free-stabilization region.*

We begin with the following lemma.

Lemma 8.3. *Let \mathcal{H} and \mathcal{H}^+ be as above. The hypercube $(\mathcal{L}_\alpha^+)^{\lambda}$ may be chosen so that each chain of length at least 2 in the λ direction is of the form $U\eta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+$.*

Proof. The proof is to follow the steps of the construction in Section 7.2 and verify that each of the chains in $\mathcal{L}_\alpha^\lambda$ of length 2 or more may be taken to be of the stated form. Suppose $\varepsilon < \varepsilon'$ and $|\varepsilon' - \varepsilon|_{L^1} = 1$. The chain of \mathcal{L}_α^+ from α_ε to $\alpha_{\varepsilon'}$ is of the form $\Theta_{\alpha_{\varepsilon'},\alpha_\varepsilon} \otimes \theta_{\alpha_0,\alpha_0}^+$ by definition. As in equation (7.4), the desired length 2 hypercube relation is

$$(A_\lambda + A'_\lambda)(\Theta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+) = \partial(U\eta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+). \quad (8.3)$$

From equation (7.5), we know that

$$(A_\lambda + A'_\lambda)(\Theta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+) = U(\Phi_w + \Phi_{w'})(\Theta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+).$$

Using the differential computation from [39, Proposition 6.5] we see that $\Phi_{w'}(\Theta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+) = 0$ and

$$\Phi_w(\Theta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+) = \Phi_w^-(\Theta_{\alpha',\alpha}) \otimes \theta_{\alpha_0,\alpha_0}^+,$$

where Φ_w^- denotes the basepoint action on the unstabilized complexes. In particular, we see that $\Phi_w^-(\Theta_{\alpha',\alpha}) = \partial\eta_{\alpha',\alpha}$, for some $\eta_{\alpha',\alpha}$, since Φ_w^- increases grading by 1, and $\Theta_{\alpha',\alpha}$ is in the maximal grading of homology. Using the fact that S_w^+ is a chain map [39, Proposition 6.5], we also have

$$\partial(\eta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+) = \partial(\eta_{\alpha',\alpha}) \otimes \theta_{\alpha_0,\alpha_0}^+.$$

In particular, the chain $\eta_{\alpha',\alpha}$, constructed as above, satisfies equation (8.3).

The higher length chains of $(\mathcal{L}_\alpha^+)^{\lambda}$ are analyzed similarly. In our present context, by using the stabilization result for holomorphic polygons in Proposition 6.5, we see that the right-hand side of equation (7.6) can be written as $C_{\varepsilon',\varepsilon} \otimes \theta_{\alpha_0,\alpha_0}^+$, where $C_{\varepsilon',\varepsilon}$ is a chain on the unstabilized complex. By the exact same reasoning as before, $C_{\varepsilon',\varepsilon}$ is a multiple of U , and is a cycle, so we may pick a primitive of the form $U\eta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+$, completing the proof. ■

Proof of Proposition 8.2. The hypercube relations for the right-hand side follow from the fact that each of the maps $\mathcal{S}_{w'}^+$, \mathcal{A}_λ and $\mathcal{S}_{w'}^-$ is a homomorphism of hypercubes (i.e., a chain map).

The hypercube relations for the entire cube, after compressing the right-hand side of the diagram in equation (8.2) are equivalent to the relation

$$\mathcal{S}_{w'}^- \mathcal{A}_\lambda \mathcal{S}_{w'}^+ = \text{id},$$

as hypercube morphisms. First, we use Lemma 8.3 to construct the hypercube $(\mathcal{L}_\alpha^+)^{\lambda}$ so that each chain of length at least 2 which has non-trivial λ direction is of the form $U\eta_{\alpha',\alpha} \otimes \theta_{\alpha_0,\alpha_0}^+$. In the composition $\mathcal{A}_\lambda \mathcal{S}_{w'}^+$, each arrow with length greater than 1 will be obtained by a holomorphic polygon count which only has inputs of the form $x \otimes \theta_{\alpha_0,\alpha_0}^+$. By Proposition 6.5, the output of such a polygon count will also have a tensor factor of $\theta_{\alpha_0,\alpha_0}^+$. Such elements lie in the kernel of $\mathcal{S}_{w'}^-$. Hence, the composition $\mathcal{S}_{w'}^- \mathcal{A}_\lambda \mathcal{S}_{w'}^+$ has no arrows of length greater than 1. The arrows of length 1 correspond to the composition $\mathcal{S}_{w'}^- \mathcal{A}_\lambda \mathcal{S}_{w'}^+$, which is the identity, by [39, Lemma 7.10]. ■

8.2. The map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$

In this section, we define the map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$. The map will correspond to a hypercube relating the expanded model of the involution from Section 7 and the non-expanded model from Section 3.2.

The construction proceeds in two steps. We first relate the singly-pointed involution to a model which adds a basepoint in a very simple way. We call this the *trivially expanded* involution. As a second step, we relate the trivially expanded model to the expanded model defined in Section 7.

Let $(\Sigma, \alpha, \beta, w)$ denote a Heegaard diagram, and let w' be a new basepoint, chosen near w . Let α_0 and β_0 be two new alpha and beta circles, centered at w' . We will let s denote either of the curves α_0 or β_0 (or small translates thereof).

In Figure 8.1, each square commutes, and no length 2 maps are necessary. As our definition, we take the compression of the right-hand side of Figure 8.1 as the trivially expanded model of the involution.

We now describe how to relate the trivially expanded model in Figure 8.1 to the expanded model described in Section 7. To do this, we use the hyperbox shown in Figure 8.2.

We now explain the maps appearing in equation Figure 8.2. Except for the top- and bottom-most cubes, all of the faces either commute on the nose, or are obtained by pairing hypercubes of attaching curves. The hypercubes are all of handleslide equivalent attaching curves, or handleslide equivalent attaching curves with an extra direction corresponding to λ . Hence, the construction follows by the technique of

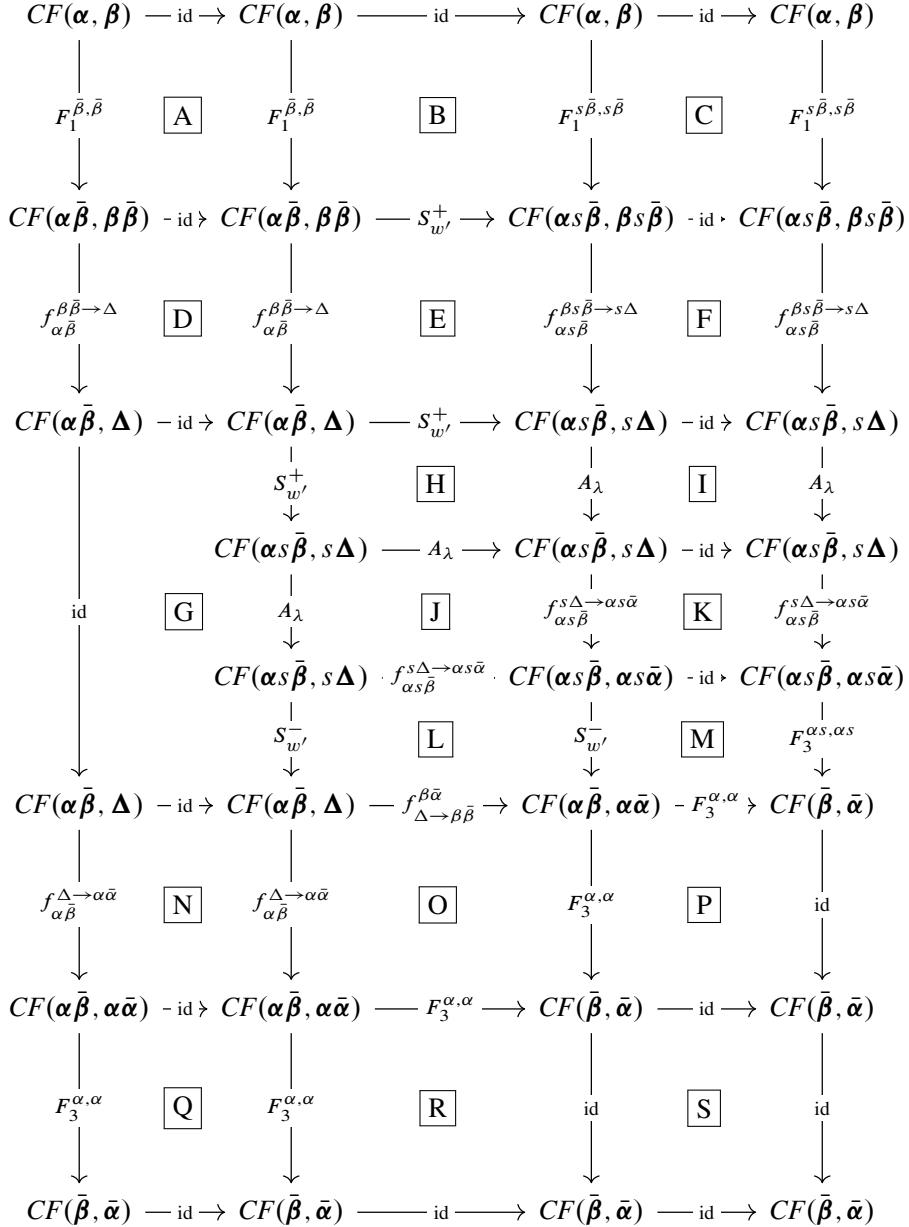


Figure 8.1. The first hyperbox used to construct the map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$. None of the faces have length 2 maps. The boxed letter will be referred to in the proof of part (1) of Proposition 8.1.

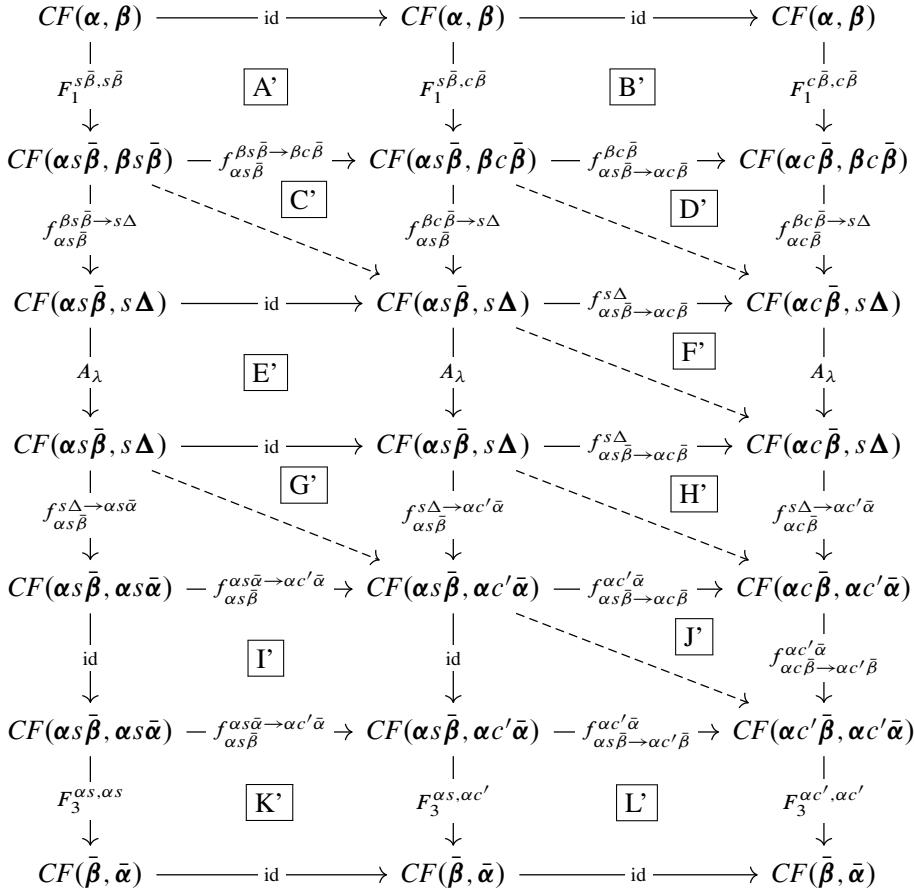


Figure 8.2. The second hyperbox used to construct the map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$. The boxed letters label the faces (in particular, they do not label maps in the faces).

filling empty hypercubes of handleslide equivalent attaching curves [18, Lemma 8.6], or Section 7.2.

We now consider the top-most and bottom-most hypercubes. We claim that the hypercube relations may be verified by appealing to Proposition 6.5. To this end, we state the following lemma.

Lemma 8.4. (1) *With respect to the unique Spin^c structures which have torsion restriction to the boundaries of the corresponding 4-manifolds, the Heegaard triples*

$$(\bar{\Sigma}, s\bar{\beta}, s\bar{\beta}, c\bar{\beta}, w, w') \quad \text{and} \quad (\bar{\Sigma}, c\bar{\beta}, s\bar{\beta}, c\bar{\beta}, w, w')$$

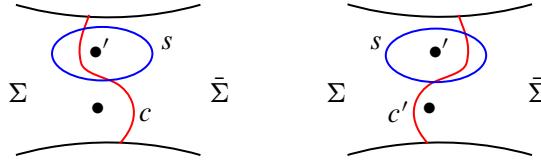


Figure 8.3. The curves s , c and c' in the connected sum neck region. The w basepoint is indicated by \bullet and the w' is indicated by \bullet' .

are weakly admissible multi-stabilizing Heegaard triples. Furthermore,

$$\begin{aligned}\hat{f}_{s\bar{\beta},s\bar{\beta},c\bar{\beta}}(\Theta_{s\bar{\beta},s\bar{\beta}}^+, \Theta_{s\bar{\beta},c\bar{\beta}}^+) &= \Theta_{s\bar{\beta},c\bar{\beta}}^+, \\ \hat{f}_{c\bar{\beta},s\bar{\beta},c\bar{\beta}}(\Theta_{c\bar{\beta},s\bar{\beta}}^+, \Theta_{s\bar{\beta},c\bar{\beta}}^+) &= \Theta_{c\bar{\beta},c\bar{\beta}}^+.\end{aligned}$$

Here, $\hat{f}_{s\bar{\beta},s\bar{\beta},c\bar{\beta}}$ denotes the map on \widehat{CF} , and similarly for the other holomorphic triangle map.

(2) Similarly, with respect to the unique Spin^c structures which have torsion restriction to the boundary,

$$(\Sigma, \alpha s, \alpha s, \alpha c', w, w') \quad \text{and} \quad (\Sigma, \alpha c', \alpha s, \alpha c', w, w')$$

are weakly admissible multi-stabilizing diagrams, and

$$\begin{aligned}\hat{f}_{\alpha s, \alpha s, \alpha c'}(\Theta_{\alpha s, \alpha s}^-, \Theta_{\alpha s, \alpha c'}^+) &= \Theta_{\alpha s, \alpha c'}^-, \\ \hat{f}_{\alpha c', \alpha s, \alpha c'}(\Theta_{\alpha c', \alpha s}^+, \Theta_{\alpha s, \alpha c'}^-) &= \Theta_{\alpha c', \alpha c'}^-.\end{aligned}$$

Proof. The claims about admissibility are proven by reducing to the genus zero triples (S^2, s, s, c, w, w') , and so forth, by removing pairwise isotopic triples of curves from $\bar{\beta}$ or α . See Figure 8.3 for the model case. It is straightforward to verify admissibility for the genus 0 diagrams. Clearly the other conditions of a multi-stabilizing diagram are satisfied.

We now prove the claim about the holomorphic triangle maps. Both triangle maps may be interpreted as naturality maps for a sequence of handleslides. Since each of the Floer complexes appearing in the statement have a unique top degree intersection point, and a unique bottom degree intersection point, and also since the claims only involve the maps on \widehat{CF} , the claims follow from grading considerations. ■

8.3. Proof of part (1) of Proposition 8.1

We now show that the map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$, defined in the previous section, commutes up to chain homotopy for the transition maps for elementary equivalences of the alpha and beta curves. This is part (1) of Proposition 8.1.

Our proof is to build a 3-dimensional hypercube which realizes the homotopy commutation

$$F_{\mathfrak{D}' \rightarrow \tilde{\mathfrak{D}'}} \circ \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'} \simeq \tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}'}} \circ F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}},$$

whenever \mathfrak{D} and \mathfrak{D}' differ by an elementary equivalence. We build this hypercube as follows. For each of the faces labeled (A)–(S) in Figure 8.1, and each of the faces labeled (A')–(L') in Figure 8.2, we build a hyperbox of size $(1, 1, 2)$, which extends the corresponding sub-cube of the map $F_{\mathfrak{D} \rightarrow \tilde{\mathfrak{D}}}$ in the direction of the maps $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ and $\tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}'}}$. We can view each of these hyperboxes as a sequence of two hypercubes. The third coordinate of the first hypercube is an alpha equivalence, and the third coordinate of the second hypercube is a beta equivalence. (This pattern is parallel to the construction of the map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$; see Figure 5.1). Write α' , β' , and Δ' for the attaching curves of \mathfrak{D}' .

The construction of the size $(1, 1, 2)$ hyperbox over the facet labeled B in Figure 8.1 is shown below.

$$\begin{array}{ccccccc}
CF(\alpha, \beta) & \xrightarrow{f_{\alpha \rightarrow \alpha'}^\beta} & CF(\alpha', \beta) & \xrightarrow{f_{\alpha' \rightarrow \beta'}^\beta} & CF(\alpha', \beta') & \xrightarrow{f_{\alpha' \rightarrow \beta'}^{\beta' \rightarrow \beta'}} & CF(\alpha', \beta') \\
\downarrow id & & \downarrow id & & \downarrow id & & \downarrow id \\
CF(\alpha, \beta) & \xrightarrow{f_{\alpha \rightarrow \alpha'}^\beta} & CF(\alpha', \beta) & \xrightarrow{f_{\alpha' \rightarrow \beta'}^\beta} & CF(\alpha', \beta') & \xrightarrow{f_{\alpha' \rightarrow \beta'}^{\beta' \rightarrow \beta'}} & CF(\alpha', \beta') \\
\downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}', \bar{\beta}} & & \downarrow F_1^{\bar{\beta}'', \bar{\beta}'} & & \downarrow F_1^{\bar{\beta}'', \bar{\beta}'} \\
\boxed{B} & & & & & & \\
\downarrow F_1^{s\bar{\beta}', s\bar{\beta}} & & \downarrow F_1^{s\bar{\beta}'', s\bar{\beta}} & & \downarrow F_1^{s\bar{\beta}'', s\bar{\beta}'} & & \downarrow F_1^{s\bar{\beta}'', s\bar{\beta}'} \\
CF(\alpha\bar{\beta}, \beta\bar{\beta}) & \xrightarrow{f_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}'}^\beta} & CF(\alpha'\bar{\beta}', \beta\bar{\beta}) & \xrightarrow{f_{\alpha'\bar{\beta}' \rightarrow \beta'\bar{\beta}'}^{\beta'' \rightarrow \beta'}} & CF(\alpha'\bar{\beta}', \beta'\bar{\beta}') & \xrightarrow{f_{\alpha'\bar{\beta}' \rightarrow \beta'\bar{\beta}'}^{\beta'' \rightarrow \beta' \bar{\beta}'}} & CF(\alpha'\bar{\beta}', \beta'\bar{\beta}') \\
\downarrow S_{w'}^+ & & \downarrow S_{w'}^+ & & \downarrow S_{w'}^+ & & \downarrow S_{w'}^+ \\
CF(\alpha s\bar{\beta}, \beta s\bar{\beta}) & \xrightarrow{f_{\alpha s\bar{\beta} \rightarrow \alpha' s\bar{\beta}'}^\beta} & CF(\alpha' s\bar{\beta}', \beta s\bar{\beta}) & \xrightarrow{f_{\alpha' s\bar{\beta}' \rightarrow \beta' s\bar{\beta}'}^{\beta'' \rightarrow \beta' s\bar{\beta}'}} & CF(\alpha' s\bar{\beta}', \beta' s\bar{\beta}') & \xrightarrow{f_{\alpha' s\bar{\beta}' \rightarrow \beta' s\bar{\beta}'}^{\beta'' \rightarrow \beta' s\bar{\beta}'}} & CF(\alpha' s\bar{\beta}', \beta' s\bar{\beta}') \\
\end{array} \tag{8.4}$$

The hypercube relations for the diagram in (8.4) follow by using Proposition 6.5 to destabilize the holomorphic triangle counts.

The hyperbox over the face labeled (G) in Figure 8.1 follows from the construction from Proposition 8.2. The size $(1, 1, 2)$ hyperboxes over the remaining faces (A)–(S) are similar to the construction in equation (8.4). We leave the details to the reader.

Next, we move on to extending the hyperboxes from Figure 8.2. The extensions of the faces labeled (C')–(J') are obtained by pairing hypercubes of handleslide equivalent attaching curves (with an extra dimension corresponding to λ in the case of E' and F'). In particular, the extensions are constructed using techniques which are now standard, so we leave them to the reader.

We investigate the hypercubes labeled A', B', K' and L'. All four are constructed similarly, so we focus our attention on the face labeled A'. The extension is shown in

the following diagram:

$$\begin{array}{ccccccc}
CF(\alpha, \beta) & \xrightarrow{f_{\alpha \rightarrow \alpha'}^\beta} & CF(\alpha', \beta) & \xrightarrow{f_{\alpha'}^{\beta \rightarrow \beta'}} & CF(\alpha', \beta') & \xrightarrow{id} & CF(\alpha', \beta') \\
\downarrow id & & \downarrow id & & \downarrow id & & \downarrow id \\
CF(\alpha, \beta) & \xrightarrow{f_{\alpha \rightarrow \alpha'}^\beta} & CF(\alpha', \beta) & \xrightarrow{f_{\alpha'}^{\beta \rightarrow \beta'}} & CF(\alpha', \beta') & \xrightarrow{id} & CF(\alpha', \beta') \\
\downarrow F_1^{s\bar{\beta}, s\bar{\beta}} & & \downarrow F_1^{s\bar{\beta}', s\bar{\beta}} & & \downarrow F_1^{\bar{\beta}', s\bar{\beta}'} & & \downarrow F_1^{s\bar{\beta}', c\bar{\beta}'} \\
\boxed{A} & & & & & & \\
\downarrow F_1^{s\bar{\beta}, c\bar{\beta}} & & \downarrow F_1^{s\bar{\beta}', c\bar{\beta}} & & \downarrow F_1^{s\bar{\beta}', c\bar{\beta}} & & \downarrow F_1^{s\bar{\beta}', c\bar{\beta}'} \\
CF(\alpha s\bar{\beta}, \beta s\bar{\beta}) & \xrightarrow{-f_{\alpha s\bar{\beta} \rightarrow \alpha' s\bar{\beta}'}^{\beta s\bar{\beta}}} & CF(\alpha' s\bar{\beta}', \beta s\bar{\beta}) & \xrightarrow{-f_{\alpha' s\bar{\beta}' \rightarrow \beta' s\bar{\beta}'}^{\beta s\bar{\beta}'}} & CF(\alpha' s\bar{\beta}', \beta' s\bar{\beta}') & \xrightarrow{f_{\alpha' s\bar{\beta}' \rightarrow \beta' s\bar{\beta}'}^{\beta s\bar{\beta} \rightarrow \beta' s\bar{\beta}'}} & CF(\alpha' s\bar{\beta}', \beta' s\bar{\beta}') \\
\downarrow f_{\alpha s\bar{\beta} \rightarrow \beta c\bar{\beta}}^{\beta s\bar{\beta} \rightarrow \beta c\bar{\beta}} & & \downarrow f_{\alpha' s\bar{\beta}' \rightarrow \beta c\bar{\beta}'}^{\beta s\bar{\beta}' \rightarrow \beta c\bar{\beta}'} & & \downarrow f_{\alpha' s\bar{\beta}' \rightarrow \beta' c\bar{\beta}'}^{\beta s\bar{\beta}' \rightarrow \beta' c\bar{\beta}'} & & \downarrow f_{\alpha' s\bar{\beta}' \rightarrow \beta' c\bar{\beta}'}^{\beta' s\bar{\beta}' \rightarrow \beta' c\bar{\beta}'} \\
CF(\alpha s\bar{\beta}, \beta c\bar{\beta}) & \xrightarrow{f_{\alpha s\bar{\beta} \rightarrow \alpha' s\bar{\beta}'}^{\beta c\bar{\beta}}} & CF(\alpha' s\bar{\beta}', \beta c\bar{\beta}) & \xrightarrow{f_{\alpha' s\bar{\beta}' \rightarrow \beta' c\bar{\beta}'}^{\beta c\bar{\beta}' \rightarrow \beta' c\bar{\beta}'}} & CF(\alpha' s\bar{\beta}', \beta' c\bar{\beta}') & \xrightarrow{id} & CF(\alpha' s\bar{\beta}', \beta' c\bar{\beta}') \\
\end{array} \tag{8.5}$$

The hypercube relations follow from the stabilization results of Proposition 6.5, together with Lemma 8.4.

8.4. The composition law for expanded transition maps

In this section, we prove part (2) of Proposition 8.1. As in the hypotheses of the proposition, suppose that we have three handleslide equivalent Heegaard diagrams

$$(\Sigma, \alpha, \beta, w), \quad (\Sigma, \alpha', \beta', w), \quad (\Sigma, \alpha'', \beta'', w),$$

all with underlying surface Σ . We also pick three choices of basepoint expanded doubling curves Δ , Δ' and Δ'' , and let c and c' denote the curves in the connected sum region. We further assume that the tuple

$$(\Sigma \# \bar{\Sigma}, \alpha'', \bar{\alpha}'', \alpha', \bar{\alpha}', \alpha, \bar{\alpha}, \beta, \bar{\beta}, \beta', \bar{\beta}', \beta'', \bar{\beta}'', \Delta, \Delta', \Delta'', c, c', w)$$

is weakly admissible (cf. Remark 7.6). Write $\tilde{\mathfrak{D}}$, $\tilde{\mathfrak{D}'}$ and $\tilde{\mathfrak{D}''}$ for the associated basepoint expanded, doubling enhanced Heegaard diagrams. We will show that

$$\tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}''}} \simeq \tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}'}} \circ \tilde{\Psi}_{\tilde{\mathfrak{D}'} \rightarrow \tilde{\mathfrak{D}''}}.$$

Our proof goes by constructing a hyperbox which realizes the stated relation. For each vertical level in the construction of $\tilde{\Psi}_{\tilde{\mathfrak{D}} \rightarrow \tilde{\mathfrak{D}'}}$, the hyperbox we construct will have a level which takes the form of a hyperbox of size $(2, 2, 1)$.

We describe the 1-handle hyperbox first. See Figure 8.4 for the overall shape of the 1-handle hyperbox. As a first step, we construct the hypercube shown in the following

equation:

$$\begin{array}{ccccc}
 \mathbb{F}[U] & \xrightarrow{\text{id}} & \mathbb{F}[U] & & \\
 \downarrow \Theta_{\bar{\beta}'}, \bar{\beta} & \searrow \text{id} & \downarrow \Theta_{\bar{\beta}'}, \bar{\beta}' & \searrow \text{id} & \downarrow \Theta_{\bar{\beta}'}, \bar{\beta}' \\
 & \lambda_{\bar{\beta}''}, \bar{\beta} & & \omega_{\bar{\beta}''}, \bar{\beta}' & \\
 & \downarrow & \Theta_{\bar{\beta}''}, \bar{\beta} & \downarrow & \Theta_{\bar{\beta}''}, \bar{\beta}' \\
 CF(\bar{\beta}', \bar{\beta}) & \xrightarrow{f_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}'}} & CF(\bar{\beta}', \bar{\beta}') & \xrightarrow{f_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}''}} & CF(\bar{\beta}'', \bar{\beta}'') \\
 \downarrow f_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}''} & \searrow h_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}''} & \downarrow f_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}''} & \searrow f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'} & \downarrow f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'} \\
 CF(\bar{\beta}'', \bar{\beta}) & \xrightarrow{f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}} & CF(\bar{\beta}'', \bar{\beta}') & \xrightarrow{f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}''}} & CF(\bar{\beta}'', \bar{\beta}'')
 \end{array} \tag{8.6}$$

The hypercube in equation (8.6) is constructed similarly to the procedure for *filling* a hypercube of handleslide equivalent attaching curves, described by Manolescu and Ozsváth [18, Lemma 8.6]. In equation (8.6), an arrow labeled by an element $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ means the map from $\mathbb{F}[U]$ to $CF(\alpha, \beta)$ given by

$$1 \mapsto x,$$

extended equivariantly over U . The existence of $\lambda_{\bar{\beta}''}, \bar{\beta}$, $\lambda_{\bar{\beta}''}, \bar{\beta}'$ and $\omega_{\bar{\beta}''}, \bar{\beta}'$ so that the diagram in (8.6) satisfies the hypercube relations follows from the fact that $\bar{\beta}$, $\bar{\beta}'$ and $\bar{\beta}''$ are handleslide equivalent.

It is helpful to concretely understand the hypercube relations in (8.6). The length 1 relations amount to each chain marked with Θ being a cycle. The length 2 relations are as follows:

$$\begin{aligned}
 f_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}'}(\Theta_{\bar{\beta}'}, \bar{\beta}) &= \Theta_{\bar{\beta}'}, \bar{\beta}', \\
 f_{\bar{\beta}'}^{\bar{\beta}'}(\Theta_{\bar{\beta}'}, \bar{\beta}') &= \Theta_{\bar{\beta}'}, \bar{\beta}', \\
 \Theta_{\bar{\beta}'}, \bar{\beta} + f_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta}''}(\Theta_{\bar{\beta}'}, \bar{\beta}) &= \partial \lambda_{\bar{\beta}''}, \bar{\beta}, \\
 f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\Theta_{\bar{\beta}'}, \bar{\beta}) + \Theta_{\bar{\beta}'}, \bar{\beta}' &= \partial \lambda_{\bar{\beta}''}, \bar{\beta}', \\
 f_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta}'} \circ f_{\bar{\beta}'}^{\bar{\beta} \rightarrow \bar{\beta}'} + f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'} \circ f_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta}'} &= [\partial, h_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}]. \tag{8.7}
 \end{aligned}$$

Note that the first relation is forced by grading considerations, and the second is implied by a nearest point map computation. The third and fourth relations are clearly satisfiable, and the final relation is the associativity relations for quadrilaterals.

Finally, the length 3 relation is

$$f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\lambda_{\bar{\beta}'', \bar{\beta}}) + \lambda_{\bar{\beta}'', \bar{\beta}'} + h_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\Theta_{\bar{\beta}'', \bar{\beta}}) = \partial \omega_{\bar{\beta}'', \bar{\beta}'}.$$

Note that the hypercube in (8.6) is a new type of hypercube. It is not naturally realized by pairing alpha and beta hypercubes together. We can think of the left-hand side as determining the following hypercube of *alpha* attaching curves:

$$\begin{array}{ccc} \bar{\beta} & \xrightarrow{1} & \bar{\beta} \\ | & \searrow & | \\ \Theta_{\bar{\beta}', \bar{\beta}} & \lambda_{\bar{\beta}'', \bar{\beta}} & \Theta_{\bar{\beta}'', \bar{\beta}} \\ \downarrow & \searrow & \downarrow \\ \bar{\beta}' & - \Theta_{\bar{\beta}'', \bar{\beta}'} \rightarrow & \bar{\beta}'' \end{array} \quad (8.8)$$

On the other hand, the front side is more naturally viewed as corresponding to the following hypercube of *beta* attaching curves:

$$\begin{array}{ccc} \bar{\beta}'' & \xrightarrow{1} & \bar{\beta}'' \\ | & \searrow & | \\ \Theta_{\bar{\beta}'', \bar{\beta}} & \lambda_{\bar{\beta}'', \bar{\beta}'} & \Theta_{\bar{\beta}'', \bar{\beta}'} \\ \downarrow & \searrow & \downarrow \\ \bar{\beta} & - \Theta_{\bar{\beta}, \bar{\beta}'} \rightarrow & \bar{\beta}' \end{array} \quad (8.9)$$

The bottom face of (8.6) is the pairing of hypercubes of attaching curves.

Next, we pick hypercubes of alpha attaching curves and beta attaching curves (respectively):

$$\begin{array}{ccc} \alpha & - \Theta_{\alpha', \alpha} \rightarrow & \alpha' \\ | & \searrow & | \\ 1 & \lambda_{\alpha'', \alpha} & \Theta_{\alpha'', \alpha'} \\ \downarrow & \searrow & \downarrow \\ \alpha & - \Theta_{\alpha'', \alpha} \rightarrow & \alpha'' \end{array} \quad \text{and} \quad \begin{array}{ccc} \beta & - \Theta_{\beta, \beta'} \rightarrow & \beta' \\ | & \searrow & | \\ 1 & \lambda_{\beta, \beta''} & \Theta_{\beta, \beta''} \\ \downarrow & \searrow & \downarrow \\ \beta & - \Theta_{\beta, \beta''} \rightarrow & \beta'' \end{array} \quad (8.10)$$

As usual, we assume that the length 1 morphisms represent top degree cycles. Note that we do not assume that the chains in (8.10) have any relation with the chains in (8.8) or (8.9).

We now consider the diagram in Figure 8.4. We will presently define the maps appearing therein, and then subsequently prove the hypercube relations. All the maps

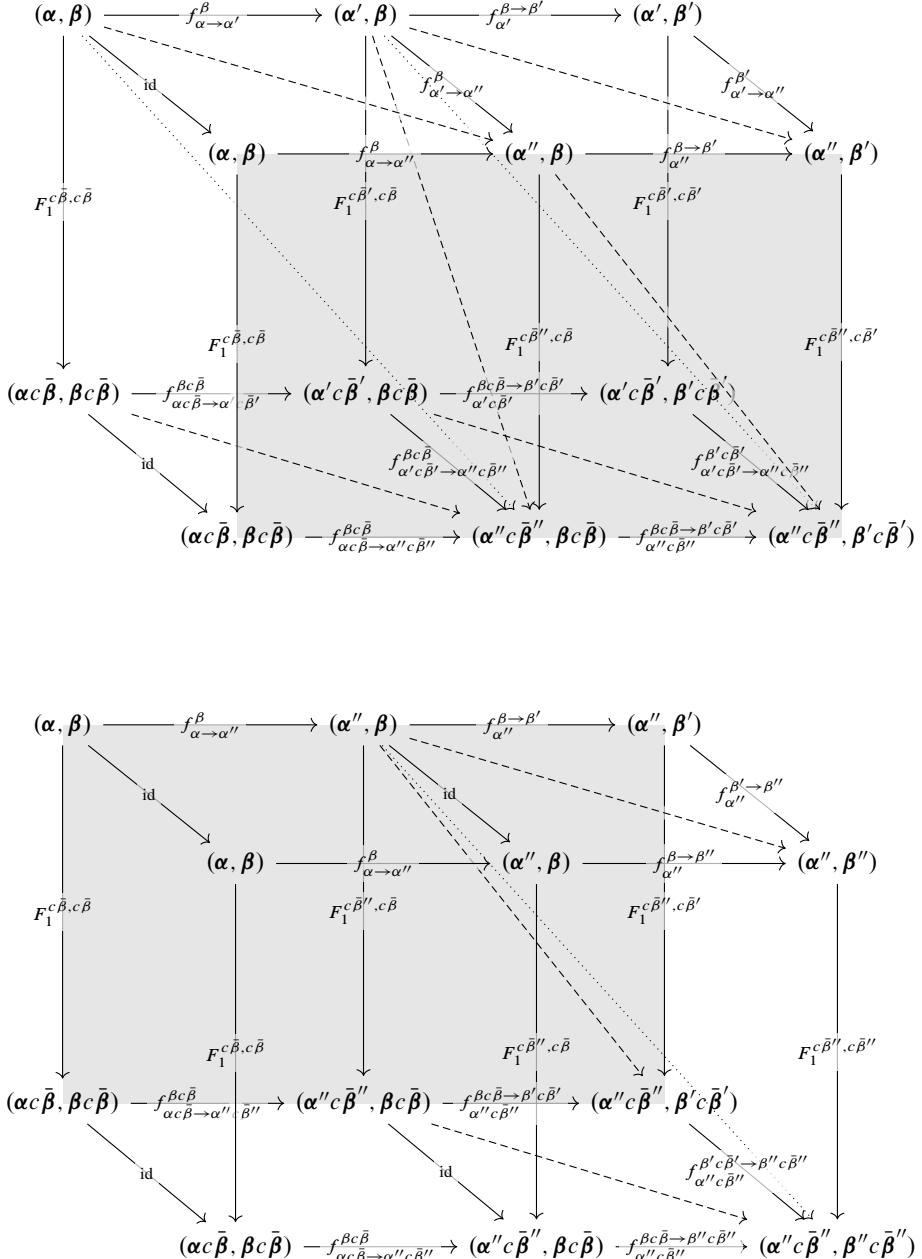


Figure 8.4. Four hypercubes. The length 2 maps are sums of quadrilateral and triangle maps (obtained by pairing hypercubes). The two boxes are stacked along the gray faces.

are built from the hypercubes in (8.6) and (8.10). We note that the hypercube relations are slightly subtle, since they do not arise from simply pairing hypercubes of attaching curves together.

We focus first on the top left cube of Figure 8.4. The maps appearing there are as follows.

(1) The length 2 map along the top is one associated to pairing the 0-dimensional beta hypercube β with the alpha hypercube

$$\begin{array}{ccc} \alpha & \xrightarrow{\Theta_{\alpha',\alpha}} & \alpha' \\ | & \searrow \lambda_{\alpha'',\alpha} & | \\ 1 & & \Theta_{\alpha'',\alpha'} \\ \downarrow & & \searrow \\ \alpha & \xrightarrow{\Theta_{\alpha'',\alpha}} & \alpha'' \end{array}$$

(2) The length 2 map along the bottom is the one obtained by pairing the 0-dimensional beta hypercube $\beta c \bar{\beta}$ with the hypercube

$$\begin{array}{ccc} \alpha c \bar{\beta} & \xrightarrow{\Theta_{\alpha',\alpha}|\theta^+|\Theta_{\bar{\beta}',\bar{\beta}}} & \alpha' c \bar{\beta}' \\ | & \searrow \lambda & | \\ 1 & & \Theta_{\alpha'',\alpha'}|\theta^+|\Theta_{\bar{\beta}'',\bar{\beta}'} \\ \downarrow & & \searrow \\ \alpha c \bar{\beta} & \xrightarrow{\Theta_{\alpha'',\alpha}|\theta^+|\Theta_{\bar{\beta}'',\bar{\beta}}} & \alpha'' c \bar{\beta}'' \end{array}$$

where

$$\lambda = \lambda_{\alpha'',\alpha}|\theta^+|f_{\bar{\beta}'',\bar{\beta}',\bar{\beta}}(\Theta_{\bar{\beta}'',\bar{\beta}'}, \Theta_{\bar{\beta}',\bar{\beta}}) + \Theta_{\alpha'',\alpha}|\theta^+|\lambda_{\bar{\beta}'',\bar{\beta}}.$$

(3) The length 2 map along the right-hand side is

$$x \mapsto f_{\alpha' \rightarrow \alpha''}^{\beta}(x)|\theta^+|\lambda_{\bar{\beta}'',\bar{\beta}}.$$

(4) The length 3 map is

$$x \mapsto h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(x)|\theta^+|\lambda_{\bar{\beta}'',\bar{\beta}} + f_{\alpha'',\alpha,\beta}(\lambda_{\alpha'',\alpha}, x)|\theta^+|\lambda_{\bar{\beta}'',\bar{\beta}}.$$

Lemma 8.5. *The maps described above for the top left subcube of the diagram in Figure 8.4 satisfy the hypercube relations.*

Proof. The length 1 relations are immediate, as are the length 2 relations for all but the right face of the cube. For the right face of the cube, using the stabilization result from Proposition 6.5, applied to holomorphic triangles, we see that the hypercube

relations amount to showing that if $\mathbf{x} \in \mathbb{T}_{\alpha''} \cap \mathbb{T}_{\beta}$, then

$$\begin{aligned} & f_{\alpha' \rightarrow \alpha''}^{\beta}(\mathbf{x})|\theta^+|(\Theta_{\bar{\beta}'', \bar{\beta}} + f_{\bar{\beta}'', \bar{\beta}'', \bar{\beta}}(\Theta_{\bar{\beta}'', \bar{\beta}'}, \Theta_{\bar{\beta}'', \bar{\beta}})) \\ &= \partial(f_{\alpha' \rightarrow \alpha''}^{\beta}(\mathbf{x})|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}}) + f_{\alpha' \rightarrow \alpha''}^{\beta}(\partial \mathbf{x})|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}}, \end{aligned}$$

which follows from the form of the differential on the Heegaard diagram obtained by attaching a 1-handle. The key point is that the right-hand side of the above equation simplifies to just $f_{\alpha' \rightarrow \alpha''}^{\beta}(\mathbf{x})|\theta^+|\partial \lambda_{\bar{\beta}'', \bar{\beta}}$, which coincides with the left-hand side by equation (8.7).

Next, we consider the length 3 relation of the cube. The desired relation is that if $\mathbf{x} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$, then

$$\begin{aligned} & h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(\mathbf{x})|\theta^+|\Theta_{\bar{\beta}'', \bar{\beta}} \\ &+ f_{\alpha'', \alpha, \beta}(\lambda_{\alpha'', \alpha}, \mathbf{x})|\theta^+|\Theta_{\bar{\beta}'', \bar{\beta}} \\ &+ h_{\alpha'' c \bar{\beta}'', \alpha' c \bar{\beta}'', \alpha c \bar{\beta}, \beta c \bar{\beta}}(\Theta_{\alpha'', \alpha'}|\theta^+|\Theta_{\bar{\beta}'', \bar{\beta}'}, \Theta_{\alpha', \alpha}|\theta^+|\Theta_{\bar{\beta}'', \bar{\beta}'}, \mathbf{x}|\theta^+|\Theta_{\bar{\beta}, \bar{\beta}}) \\ &+ (f_{\alpha' \rightarrow \alpha''}^{\beta} \circ f_{\alpha \rightarrow \alpha'}^{\beta})(\mathbf{x})|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}} \\ &+ f_{\alpha'' c \bar{\beta}'', \alpha c \bar{\beta}, \beta c \bar{\beta}}(\lambda_{\alpha'', \alpha}|\theta^+|f_{\bar{\beta}'', \bar{\beta}'', \bar{\beta}}(\Theta_{\bar{\beta}'', \bar{\beta}'}, \Theta_{\bar{\beta}'', \bar{\beta}}), \mathbf{x}|\theta^+|\Theta_{\bar{\beta}, \bar{\beta}}^+) \\ &+ f_{\alpha'' c \bar{\beta}'', \alpha c \bar{\beta}, \beta c \bar{\beta}}(\Theta_{\alpha'', \alpha}|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}}, \mathbf{x}|\theta^+|\Theta_{\bar{\beta}, \bar{\beta}}) \\ &= \left[h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(-)|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}} + f_{\alpha'', \alpha, \beta}(\lambda_{\alpha'', \alpha}, -)|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}}, \partial \right](\mathbf{x}). \end{aligned} \tag{8.11}$$

We note that

$$\begin{aligned} & \left[h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(-)|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}}, \partial \right](\mathbf{x}) \\ &= f_{\alpha'', \alpha, \beta}(f_{\alpha'', \alpha', \alpha}(\Theta_{\alpha'', \alpha'}, \Theta_{\alpha', \alpha}), \mathbf{x})|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}} \\ &+ (f_{\alpha' \rightarrow \alpha''}^{\beta} \circ f_{\alpha \rightarrow \alpha'}^{\beta})(\mathbf{x})|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}} \\ &+ h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(\mathbf{x})|\theta^+|\Theta_{\bar{\beta}'', \bar{\beta}} \\ &+ h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(\mathbf{x})|\theta^+|f_{\bar{\beta}'', \bar{\beta}'', \bar{\beta}}(\Theta_{\bar{\beta}'', \bar{\beta}'}, \Theta_{\bar{\beta}'', \bar{\beta}}). \end{aligned} \tag{8.12}$$

Also,

$$\begin{aligned} & \left[f_{\alpha'', \alpha, \beta}(\lambda_{\alpha'', \alpha}, -)|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}}, \partial \right](\mathbf{x}) \\ &= f_{\alpha'', \alpha, \beta}(\Theta_{\alpha'', \alpha} + f_{\alpha'', \alpha', \alpha}(\Theta_{\alpha'', \alpha'}, \Theta_{\alpha', \alpha}), \mathbf{x})|\theta^+|\lambda_{\bar{\beta}'', \bar{\beta}} \\ &+ f_{\alpha'', \alpha, \beta}(\lambda_{\alpha'', \alpha}, \mathbf{x})|\theta^+|(\Theta_{\bar{\beta}'', \bar{\beta}} + f_{\bar{\beta}'', \bar{\beta}'', \bar{\beta}}(\Theta_{\bar{\beta}'', \bar{\beta}'}, \Theta_{\bar{\beta}'', \bar{\beta}})). \end{aligned} \tag{8.13}$$

Next, we note that the stabilization result from Proposition 6.5 implies that

$$\begin{aligned}
& f_{\alpha''c\bar{\beta}'',\alpha c\bar{\beta},\beta c\bar{\beta}}(\lambda_{\alpha'',\alpha}|\theta^+|f_{\bar{\beta}'',\bar{\beta}',\bar{\beta}}(\Theta_{\bar{\beta}'',\bar{\beta}'},\Theta_{\bar{\beta}',\bar{\beta}}),x|\theta^+|\Theta_{\bar{\beta},\bar{\beta}}) \\
& = f_{\alpha'',\alpha,\beta}(\lambda_{\alpha'',\alpha},x)|\theta^+|f_{\bar{\beta}'',\bar{\beta}',\bar{\beta}}(\Theta_{\bar{\beta}'',\bar{\beta}'},\Theta_{\bar{\beta}',\bar{\beta}}), \\
& f_{\alpha''c\bar{\beta}'',\alpha c\bar{\beta},\beta c\bar{\beta}}(\Theta_{\alpha'',\alpha}|\theta^+|\lambda_{\bar{\beta}'',\bar{\beta}},x|\theta^+|\Theta_{\bar{\beta},\bar{\beta}}) \\
& = f_{\alpha'',\alpha,\beta}(\Theta_{\alpha'',\alpha},x)|\theta^+|\lambda_{\bar{\beta}'',\bar{\beta}}. \tag{8.14}
\end{aligned}$$

Using (8.12), (8.13), and (8.14) to simplify (8.11), we see that the length 3 hypercube relation becomes equivalent to

$$\begin{aligned}
& h_{\alpha''c\bar{\beta}'',\alpha'c\bar{\beta}',\alpha c\bar{\beta},\beta c\bar{\beta}}(\Theta_{\alpha'',\alpha'}|\theta^+|\Theta_{\bar{\beta}'',\bar{\beta}'},\Theta_{\alpha',\alpha}|\theta^+|\Theta_{\bar{\beta}',\bar{\beta}},x|\theta^+|\Theta_{\bar{\beta},\bar{\beta}}) \\
& = h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(x)|\theta^+|f_{\bar{\beta}'',\bar{\beta}',\bar{\beta}}(\Theta_{\bar{\beta}'',\bar{\beta}'},\Theta_{\bar{\beta}',\bar{\beta}}). \tag{8.15}
\end{aligned}$$

We now prove (8.15). Suppose that ψ and ψ_0 are classes of rectangles on $(\Sigma, \alpha'', \alpha', \alpha, \beta)$ and $(\bar{\Sigma}, \bar{\beta}'', \bar{\beta}', \bar{\beta}, \bar{\beta})$ such that $\mu(\psi) + \mu(\psi_0) = -1$. Let $C(\psi, \psi_0)$ denote the set of all homology classes which coincide with ψ and ψ_0 outside the 1-handle region, and have θ^+ as their input and output in the 1-handle region. (Such rectangles necessarily have index -1 .) Applying Proposition 6.5 to the 1-handle region allows us to remove the 1-handle region connecting Σ and $\bar{\Sigma}$. We obtain

$$\sum_{\phi \in C(\psi, \psi_0)} \# \mathcal{M}(\phi) \equiv \#(\mathcal{M}(\psi) \times_{\text{ev} \square} \mathcal{M}(\psi_0)) \pmod{2}. \tag{8.16}$$

Here, we are also using the nodal almost complex structures around the 1-handle region on the left-hand side. Furthermore, any class of rectangles which is not in some $C(\psi, \psi_0)$ for a pair (ψ, ψ_0) satisfying $\mu(\psi) + \mu(\psi_0) = -1$ has no holomorphic representatives. Also, $\times_{\text{ev} \square}$ denotes the fibered product over the evaluation map to the parameter space $\mathbb{R} \cong \mathcal{M}(\square)$.

We write h^0 for the map from $CF^-(\alpha, \beta)$ to $CF^-(\alpha''c\bar{\beta}'', \beta c\bar{\beta})$ which is the composition

$$h^0 = h_{\alpha c \bar{\beta} \rightarrow \alpha' c \bar{\beta}' \rightarrow \alpha'' c \bar{\beta}''}^{\beta c \bar{\beta}} \circ F_1^{c \bar{\beta}, c \bar{\beta}}.$$

By equation (8.16), we may identify h^0 as the map which counts holomorphic rectangles on the disjoint union of two diagrams, viewed as a fibered product over $\text{ev} \square$. We may consider a map h^t for any generic $t \in \mathbb{R}$, which counts curve pairs (u, u_0) on the disjoint union consisting of holomorphic quadrilaterals such that $\text{ev} \square(u) = \text{ev} \square(u_0) + t$.

We claim that $h^t = h^{t'}$ for generic t and t' in \mathbb{R} . Firstly, we note that since $\bar{\beta}$ appears twice (once as a small Hamiltonian isotopy), the nearest point argument of [4, Proposition 11.5] (cf. [15, Section 3.4]) implies that for suitably small choices of translations, the only rectangles on $\bar{\Sigma}$ with representatives have Maslov index at

least 0. In particular, all holomorphic rectangles counted by h^t for generic t have $\mu(\psi) = -1$ and $\mu(\psi_0) = 0$.

We consider the ends of the moduli spaces counted by h^t , ranging over $t \in [0, \infty)$. No degenerations are possible on the Σ -side, since all rectangles there have index -1 . On the $\bar{\Sigma}$ -side, rectangles have index 0, so the possible degenerations are a holomorphic disk splitting off at some finite t , as well as the ends which appear as $t \rightarrow \infty$. The ends appearing at $t \rightarrow \infty$ correspond exactly to the map

$$h_{\alpha \rightarrow \alpha' \rightarrow \alpha''}^{\beta}(x) | \theta^+ | f_{\bar{\beta}'', \bar{\beta}', \bar{\beta}}(\Theta_{\bar{\beta}'', \bar{\beta}'}, f_{\bar{\beta}', \bar{\beta}, \bar{\beta}}(\Theta_{\bar{\beta}'', \bar{\beta}'}, \Theta_{\bar{\beta}, \bar{\beta}})).$$

The ends at finite t are slightly more subtle. First, one end which can appear is an end where a disk breaks off into the $\bar{\beta}$ - $\bar{\beta}$ end. These cancel modulo two since $\Theta_{\bar{\beta}, \bar{\beta}}$ is a cycle. Holomorphic disks breaking off into the other ends are prohibited, since they would leave an index -1 holomorphic rectangle, with $\Theta_{\bar{\beta}, \bar{\beta}}$ as input. (The most important end where this is prohibited is the outgoing $\bar{\beta}''$ - $\bar{\beta}$ end, where a more general diagram might have degenerations.) In particular, we obtain (8.15), completing the proof. ■

The hypercube in the bottom right of Figure 8.4 is constructed by simply interchanging the roles of the alpha and beta curves in the above procedure.

We now construct the maps in the top right hypercube of Figure 8.4. As a first step, we pick chains to give hypercubes of attaching curves, as follows.

(1) The top face is obtained by pairing the two hypercubes

$$\alpha' \xrightarrow{\Theta_{\alpha'', \alpha'}} \alpha'' \quad \text{and} \quad \beta \xrightarrow{\Theta_{\beta, \beta'}} \beta'.$$

(2) The bottom face is obtained by pairing the two hypercubes

$$\alpha' c \bar{\beta}' \xrightarrow{\Theta_{\alpha'', \alpha'} | \theta^+ | \Theta_{\bar{\beta}'', \bar{\beta}'}} \alpha'' c \bar{\beta}'' \quad \text{and} \quad \beta c \bar{\beta} \xrightarrow{\Theta_{\beta, \beta'} | c | \Theta_{\bar{\beta}, \bar{\beta}'}} \beta' c \bar{\beta}'.$$

(3) The length 2 map along the left face is

$$x \mapsto f_{\alpha' \rightarrow \alpha''}^{\beta}(x) | \theta^+ | \lambda_{\bar{\beta}'', \bar{\beta}}.$$

(4) The length 2 map along the front face is

$$x \mapsto f_{\alpha''}^{\beta \rightarrow \beta'}(x) | \theta^+ | \lambda_{\bar{\beta}'', \bar{\beta}'}.$$

(5) The length 3 map is the sum of the following two maps:

$$\begin{aligned} x &\mapsto (f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(x) | \theta^+ | \omega_{\bar{\beta}'', \bar{\beta}'}, \\ x &\mapsto h^{t < t_0}(x). \end{aligned}$$

The map $h^{t < t_0}$ is a new map, as follows. It counts pairs (u, u_0) , where u and u_0 are holomorphic rectangles on $(\Sigma, \alpha'', \alpha', \beta, \beta')$ and $(\bar{\Sigma}, \bar{\beta}'', \bar{\beta}', \bar{\beta}, \bar{\beta}')$, respectively, both of index -1 , such that

$$\text{ev}_{\square}(u) < \text{ev}_{\square}(u_0).$$

If y and z are the outputs of u and u_0 , respectively, then the $h^{t < t_0}(x)$ has a summand of $y \times \theta^+ \times z$ (weighted by the appropriate U -power).

Lemma 8.6. *With the maps described above, the top right subcube of Figure 8.4 is a hypercube of chain complexes.*

Proof. The proof is similar to the proof of Lemma 8.5. The length 1 and 2 hypercube relations are similar to the previous case, so we focus on the length 3 relation. If $x \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_{\beta}$, the desired relation is

$$\begin{aligned} h_{\alpha' \rightarrow \alpha''}^{\beta \rightarrow \beta'}(x) | \theta^+ | \Theta_{\bar{\beta}'', \bar{\beta}'} &+ f_{\alpha'' c \bar{\beta}'', \beta c \bar{\beta}, \beta' c \bar{\beta}'}(f_{\alpha' \rightarrow \alpha''}^{\beta}(x) | \theta^+ | \lambda_{\bar{\beta}'', \bar{\beta}}, \Theta_{\beta, \beta'} | \theta^+ | \Theta_{\bar{\beta}, \bar{\beta}'}) \\ &+ (f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(x) | \theta^+ | \lambda_{\bar{\beta}'', \bar{\beta}'} \\ &+ h_{\alpha' c \bar{\beta}' \rightarrow \alpha'' c \bar{\beta}''}^{\beta c \bar{\beta} \rightarrow \beta' c \bar{\beta}'}(x | \theta^+ | \Theta_{\bar{\beta}', \bar{\beta}}) \\ &= [(f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(-) | \theta^+ | \omega_{\bar{\beta}'', \bar{\beta}'} + h^{t < t_0}, \partial](x). \end{aligned} \quad (8.17)$$

By excising the special curves in the connected sum region using the results about holomorphic triangles and 1-handles from [39, Theorem 8.8] (compare Proposition 6.5 above), we obtain

$$\begin{aligned} f_{\alpha'' c \bar{\beta}'', \beta c \bar{\beta}, \beta' c \bar{\beta}'}(f_{\alpha' \rightarrow \alpha''}^{\beta}(x) | \theta^+ | \lambda_{\bar{\beta}'', \bar{\beta}}, \Theta_{\beta, \beta'} | \theta^+ | \Theta_{\bar{\beta}, \bar{\beta}'}) \\ = (f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(x) | \theta^+ | f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\lambda_{\bar{\beta}'', \bar{\beta}}). \end{aligned} \quad (8.18)$$

We now analyze the quadrilateral map

$$h_{\alpha' c \bar{\beta}' \rightarrow \alpha'' c \bar{\beta}''}^{\beta c \bar{\beta} \rightarrow \beta' c \bar{\beta}'}(x | \theta^+ | \Theta_{\bar{\beta}', \bar{\beta}}).$$

As with the previous case, we can destabilize the 1-handle region and identify the expression with the count of index -1 holomorphic quadrilaterals on $(\Sigma \sqcup \bar{\Sigma}) \times \square$. We can view this moduli space as the fibered product over the evaluation map to $\mathcal{M}(\square) \cong \mathbb{R}$. If (ψ, ψ_0) is a homology class on $(\Sigma \sqcup \bar{\Sigma}) \times \square$ which admits a holomorphic representative, then there are generically two possibilities:

- (1) $\mu(\psi) = 0$ and $\mu(\psi_0) = -1$;
- (2) $\mu(\psi) = -1$ and $\mu(\psi_0) = 0$.

As before we will write h^0 for the map $h_{\alpha'c\bar{\beta}' \rightarrow \alpha''c\bar{\beta}''}^{\beta c\bar{\beta} \rightarrow \beta'c\bar{\beta}'}(x|\theta^+|\Theta_{\bar{\beta}',\bar{\beta}})$. For $t \in [0, \infty)$, we write h^t for a deformation, which counts curves such that

$$\text{ev}_{\square}(u_0) - \text{ev}_{\square}(u) = t.$$

We count the ends of the moduli spaces defining the maps h^t (ranging over $t \in [0, \infty)$). We obtain

$$\begin{aligned} [\partial, h^{t < t_0}](x) &= h_{\alpha'c\bar{\beta}' \rightarrow \alpha''c\bar{\beta}''}^{\beta c\bar{\beta} \rightarrow \beta'c\bar{\beta}'}(x|\theta^+|\Theta_{\bar{\beta}',\bar{\beta}}) \\ &+ (f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(x)|\theta^+|h_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\Theta_{\bar{\beta}',\bar{\beta}}) + h_{\alpha' \rightarrow \alpha''}^{\beta \rightarrow \beta'}(x)|\theta^+|\Theta_{\bar{\beta}'',\bar{\beta}'}. \end{aligned} \quad (8.19)$$

Note that instead of $\Theta_{\bar{\beta}'',\bar{\beta}'}$ in the final factor, what more naturally appears is

$$f_{\bar{\beta}'',\bar{\beta}',\bar{\beta}'}(\Theta_{\bar{\beta}'',\bar{\beta}'}, f_{\bar{\beta}',\bar{\beta},\bar{\beta}'}(\Theta_{\bar{\beta}',\bar{\beta}}, \Theta_{\bar{\beta},\bar{\beta}'})).$$

However, this coincides with $\Theta_{\bar{\beta}'',\bar{\beta}'}$, by a nearest point computation, cf. [4, Proposition 11.1].

Finally,

$$\begin{aligned} &[(f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(-)]|\theta^+|\omega_{\bar{\beta}'',\bar{\beta}'}, \partial](x) \\ &= (f_{\alpha''}^{\beta \rightarrow \beta'} \circ f_{\alpha' \rightarrow \alpha''}^{\beta})(x)|\theta^+|(f_{\bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\lambda_{\bar{\beta}'',\bar{\beta}}) + \lambda_{\bar{\beta}'',\bar{\beta}'} + h_{\bar{\beta}' \rightarrow \bar{\beta}''}^{\bar{\beta} \rightarrow \bar{\beta}'}(\Theta_{\bar{\beta}',\bar{\beta}})). \end{aligned} \quad (8.20)$$

By combining (8.18), (8.19), and (8.20), we quickly obtain (8.17), completing the proof. \blacksquare

9. 1- and 3-handles

In this section, we define our cobordism maps for 1-handle and 3-handles. Since for the maps in this section there will be an obvious choice of Spin^c structure and framed path, we omit these from the notation throughout.

9.1. The construction

We now construct the map for a 4-dimensional 1-handle. Suppose that $\mathbb{S}_0 = \{p, p'\}$ is a (framed) 0-sphere in Y . We pick a Heegaard surface (Σ, α, β) so that $\mathbb{S}_0 \subseteq \Sigma \setminus (\alpha \cup \beta)$. We form a Heegaard surface for the surgered manifold $Y(\mathbb{S}_0)$ by cutting out small disks centered at p and p' , and gluing in an annular diagram (A, γ_0, γ_0) , where γ_0 and γ_0' are meridians of the annulus which intersect in two points. (Here and onwards, we abuse notation and view one copy of γ_0 as a Hamiltonian translate of the other). Write $(\Sigma', \alpha \gamma_0, \beta \gamma_0, w)$ for the diagram obtained by this procedure.

We now define the 4-dimensional 1-handle map. To do this, we first pick paths λ and λ' from p and p' to w . We pick doubling arcs δ for Σ which are disjoint from λ and λ' . Furthermore, we choose two more doubling arcs δ_0 on Σ' , which are contained in the union of a neighborhood of λ and λ' , as well as the annular 1-handle region. We assume $\delta \cup \delta_0$ is a valid collection of doubling arcs for Σ' . Write Δ and Δ_0 for the curves obtained by doubling δ and δ_0 , respectively.

Additionally, the involutive 1-handle map requires a choice of closed curve $\tau \subseteq \Sigma' \# \bar{\Sigma}'$, such that the path τ is obtained by doubling a properly embedded arc on $\Sigma' \setminus N(w)$ which intersects γ_0 transversely at a single point.

We define the 1-handle map to be the diagram obtained by compressing the following hyperbox:

$$\begin{array}{ccccc}
CF(\alpha, \beta) & \xrightarrow{\quad} & F_1^{\gamma_0, \bar{\gamma}_0} & \xrightarrow{\quad} & CF(\alpha\gamma_0, \beta\gamma_0) \\
\downarrow F_1^{\bar{\beta}, \bar{\beta}} & & & & \downarrow F_1^{\bar{\beta}\bar{\gamma}_0, \bar{\beta}\bar{\gamma}_0} \\
CF(\alpha\bar{\beta}, \beta\bar{\beta}) & - F_1^{\gamma_0\bar{\gamma}_0, \bar{\gamma}_0\bar{\gamma}_0} \rightarrow & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) & \xlongequal{\quad} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) \\
\downarrow f_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} & & \downarrow f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\beta\bar{\beta}\gamma_0\bar{\gamma}_0 \rightarrow \Delta\gamma_0\bar{\gamma}_0} & \dashrightarrow & \downarrow f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\beta\bar{\beta}\gamma_0\bar{\gamma}_0 \rightarrow \Delta\Delta_0} \\
CF(\alpha\bar{\beta}, \Delta) & - F_1^{\gamma_0\bar{\gamma}_0, \bar{\gamma}_0\bar{\gamma}_0} \rightarrow & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \Delta\gamma_0\bar{\gamma}_0) & - f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \Delta\Delta_0} \rightarrow & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \Delta\Delta_0) \\
\downarrow f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} & & \downarrow f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} & \dashrightarrow & \downarrow f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\Delta_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} \\
CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & - F_1^{\gamma_0\bar{\gamma}_0, \bar{\gamma}_0\bar{\gamma}_0} \rightarrow & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \xrightarrow{B_\tau} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) \\
\downarrow F_3^{\alpha, \alpha} & & & & \downarrow F_3^{\alpha\gamma_0, \alpha\gamma_0} \\
CF(\bar{\beta}, \bar{\alpha}) & \xrightarrow{\quad} & F_1^{\bar{\gamma}_0, \bar{\gamma}_0} & \xrightarrow{\quad} & CF(\bar{\beta}\bar{\gamma}_0, \bar{\alpha}\bar{\gamma}_0)
\end{array} \tag{9.1}$$

It will be helpful to observe that the bottom-most level of (9.1) has the following slightly expanded form:

$$\begin{array}{ccccc}
CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & - F_1^{\bar{\gamma}_0, \bar{\gamma}_0} \rightarrow & CF(\alpha\bar{\beta}\bar{\gamma}_0, \alpha\bar{\alpha}\bar{\gamma}_0) & - F_1^{\gamma_0, \bar{\gamma}_0} \rightarrow & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) \\
\parallel & & \parallel & & \downarrow F_3^{\gamma_0, \bar{\gamma}_0} \\
CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & - F_1^{\bar{\gamma}_0, \bar{\gamma}_0} \rightarrow & CF(\alpha\bar{\beta}\bar{\gamma}_0, \alpha\bar{\alpha}\bar{\gamma}_0) & \xlongequal{\quad} & CF(\alpha\bar{\beta}\bar{\gamma}_0, \alpha\bar{\alpha}\bar{\gamma}_0) \\
\downarrow F_3^{\alpha, \alpha} & & & & \downarrow F_3^{\alpha, \alpha} \\
CF(\bar{\beta}, \bar{\alpha}) & \xrightarrow{\quad} & F_1^{\bar{\gamma}_0, \bar{\gamma}_0} & \xrightarrow{\quad} & CF(\bar{\beta}\bar{\gamma}_0, \bar{\alpha}\bar{\gamma}_0)
\end{array} \tag{9.2}$$

We now describe the construction of (9.1). The top-most face which has a diagonal map is obtained by pairing a hypercube of beta attaching curves with the hypercube of alpha attaching curves consisting of $\alpha\bar{\beta}\gamma_0\bar{\gamma}_0$.

We now consider the lowest face of (9.1) which has a diagonal map. This map is obtained by pairing the 0-dimensional alpha hypercube $\alpha\bar{\beta}\gamma_0\bar{\gamma}_0$ with the diagram below. The map B_τ is the H_1 -action for τ , defined by counting changes across the beta curves, $\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0$.

$$\begin{array}{ccc}
 \Delta\gamma_0\bar{\gamma}_0 & \xrightarrow{\Theta_{\Delta\gamma_0\bar{\gamma}_0, \Delta\Delta_0}} & \Delta\Delta_0 \\
 \downarrow & \searrow \eta & \downarrow \\
 \Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} & & \Theta_{\Delta\Delta_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} \\
 \downarrow & & \downarrow \\
 \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0 & \xrightarrow{\tau} & \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0
 \end{array} \tag{9.3}$$

In equation (9.3), the hypercube relations are to be interpreted using the formalism of Section 7.2. In particular, the hypercube relations amount to the equation

$$\partial\eta = B_\tau(\Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}) + f_{\Delta\gamma_0\bar{\gamma}_0, \Delta\Delta_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}(\Theta_{\Delta\gamma_0\bar{\gamma}_0, \Delta\Delta_0}, \Theta_{\Delta\Delta_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}). \tag{9.4}$$

Remark 9.1. There are many choices of morphisms in the diagram in equation (9.3). Any two choices of top-degree generators $\Theta_{\Delta\gamma_0\bar{\gamma}_0, \Delta\Delta_0}$ are homologous. Similarly any two choices of $\Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}$ and $\Theta_{\Delta\Delta_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}$ are homologous. Unlike these morphisms, it is not the case that any two choices of diagonal morphism $\eta = \eta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}$ are homologous. Indeed, this chain lies in the top degree of $CF^-(\#^{g+1} S^1 \times S^2)$, and hence, there two choices which are not homologous. Similarly, it is not the case that any two choices for the curve τ are homologous. Nonetheless, we will show in Lemma 9.5 that resulting involutive 1-handle map obtained by compressing the diagram in equation (9.2) is independent of these choices, up to chain homotopy.

We now show that this is satisfiable.

Lemma 9.2. *There is a homogeneously graded $\eta \in CF(\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0)$ satisfying (9.4).*

Proof. Any two choices of $\Delta\Delta_0$ are related by a sequence of handleslides and isotopies. Via an associativity argument, it is sufficient to show the claim for any convenient choice of Δ and Δ_0 . In particular, we may assume that the curves of Δ_0 are disjoint from α and $\bar{\alpha}$. In particular, the relevant triple diagram

$$(\Sigma' \# \bar{\Sigma}', \Delta\gamma_0\bar{\gamma}_0, \Delta\Delta_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0)$$

may be decomposed as a connected sum

$$(\Sigma \# \bar{\Sigma}, \Delta, \Delta, \alpha\bar{\alpha}) \# (\mathbb{T}^2 \# \mathbb{T}^2, \gamma_0\bar{\gamma}_0, \Delta_0, \gamma_0\bar{\gamma}_0).$$

By the stabilization result in Proposition 6.5, it suffices to show the claim for the genus 2 triple $(\mathbb{T}^2 \# \mathbb{T}^2, \gamma_0 \bar{\gamma}_0, \Delta_0, \gamma_0 \bar{\gamma}_0)$.

There are two ways to verify the claim in this situation. The first strategy is to argue topologically via the following reasoning.

- (1) Topologically the triple $(\mathbb{T}^2 \# \mathbb{T}^2, \gamma_0 \bar{\gamma}_0, \Delta_0, \gamma_0 \bar{\gamma}_0)$ is for 0-surgery on an unknot in $S^1 \times S^2$.
- (2) By picking Δ_0 appropriately, we may assume that τ runs parallel to a curve $\Delta \in \Delta_0$. Furthermore, Δ is the meridian of the knot on which we are performing surgery, and in particular defines the dual of the surgery knot.
- (3) In general, if U is a 0-framed unknot in Y , and μ is a meridian, then

$$CF(W(U)) = A_\mu \circ F_1,$$

where F_1 is the 1-handle map. This is easily verified using a genus 1 diagram and a stabilization result for triangles.

- (4) By combining the ideas above, the proof is complete.

Alternately, one may explicitly perform the computation by counting holomorphic triangles. Write θ^\pm (resp., $\bar{\theta}^\pm$) for the intersection points of γ_0 with its translate (resp., $\bar{\gamma}_0$ with its translate). In Figure 9.1, we show the four index 0 triangle classes which can contribute. Two of the classes are in $\pi_2(\Theta^+, \Theta^+, \theta^+ | \bar{\theta}^-)$ and the other two classes are in $\pi_2(\Theta^+, \Theta^+, \theta^- | \bar{\theta}^+)$. If ψ_1 and ψ_2 denote the two classes in $\pi_2(\Theta^+, \Theta^+, \theta^+ | \bar{\theta}^-)$, we claim

$$\#\mathcal{M}(\psi_1) + \#\mathcal{M}(\psi_2) \equiv 1 \pmod{2}, \quad (9.5)$$

and similarly for the two classes in $\pi_2(\Theta^+, \Theta^+, \theta^- | \bar{\theta}^+)$. This claim is verified via a Gromov compactness argument which is illustrated and explained in Figure 9.2. Hence, we conclude that $f_{\gamma_0 \bar{\gamma}_0, \Delta_0, \gamma_0 \bar{\gamma}_0}(\Theta^+, \Theta^+) = \theta^+ | \bar{\theta}^- + \theta^- | \bar{\theta}^+$, which is also easily seen to be the action τ on $\theta^+ | \bar{\theta}^-$. ■

Finally, we verify that the bottom-most face of (9.1) is a hypercube. Note that this face may be expanded into the diagram shown in (9.2). In light of this, it suffices to show the following relation (note the similarity to equation (8.2)).

Lemma 9.3. *As maps from $CF(\alpha \bar{\beta} \bar{\gamma}_0, \alpha \bar{\alpha} \bar{\gamma}_0)$ to $CF(\alpha \bar{\beta} \bar{\gamma}_0, \alpha \bar{\alpha} \bar{\gamma}_0)$, there is an equality*

$$F_3^{\gamma_0, \bar{\gamma}_0} \circ B_\tau \circ F_1^{\gamma_0, \bar{\gamma}_0} = \text{id}.$$

Proof. Compare [39, Lemma 7.10]. The relation follows from the form of the differential after attaching a 1-handle. Namely, it follows from [39, Proposition 8.5] that the differential satisfies

$$\partial(\mathbf{x} \times \theta_{\gamma_0, \bar{\gamma}_0}^+) = \partial(\mathbf{x}) \otimes \theta_{\gamma_0, \bar{\gamma}_0}^+.$$

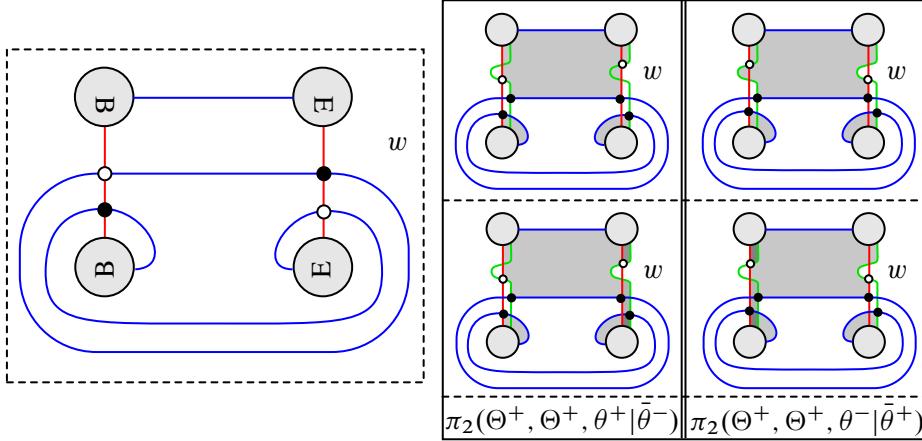


Figure 9.1. Verifying equation (9.4) by counting triangles. Left: The diagram $(T^2 \# T^2, \gamma_0 \bar{\gamma}_0, \Delta_0)$. Right: The four classes of index 0 holomorphic triangles which contribute to the triangle map $f_{\gamma_0 \bar{\gamma}_0, \Delta_0, \gamma_0 \bar{\gamma}_0}(\Theta^+, \Theta^+)$.

Furthermore, the proof shows that the holomorphic curves which go from $x \times \theta_{\gamma_0, \gamma_0}^+$ to $y \times \theta_{\gamma_0, \gamma_0}^-$ for some x and y all have $x = y$ and have domain consisting of one of the two bigons contained in the 1-handle region. The 1-chain τ intersects the boundary of exactly one of these bigons, and hence,

$$B_\tau(x \times \theta_{\gamma_0, \gamma_0}^+) = x \times \theta_{\gamma_0, \gamma_0}^- + z \otimes \theta_{\gamma_0, \gamma_0}^+$$

for some $z \in CF(\alpha\bar{\beta}, \alpha\bar{\alpha})$. The stated claim follows immediately. ■

9.2. Well-definedness of the 1-handle maps

In this section, we prove that the stabilization maps are compatible with the maps for changing the alpha and beta curves via handleslides or isotopies, or changing the choice of doubling curves Δ and Δ_0 . We also prove that the involutive 1-handle maps are independent of the choice of the curve $\tau \subseteq \Sigma' \# \bar{\Sigma}'$, as well as the choice of chains used to construct the hypercubes in (9.1).

As a first step, we verify independence of the curve τ .

Lemma 9.4. *Suppose that τ and τ' are two closed curves on $\Sigma' \# \bar{\Sigma}'$ obtained by doubling properly embedded arcs on $\Sigma' \setminus N(w)$ which intersect γ_0 transversely at a single point. Then, there are choices of the additional auxiliary data in the construction so that the models of the 1-handle maps in equation (9.1) are chain homotopic.*

Proof. In the compression of equation (9.1) the choice of τ only affects the diagonal map. There are two ways in which it appears. The first is via the choice of η , and the

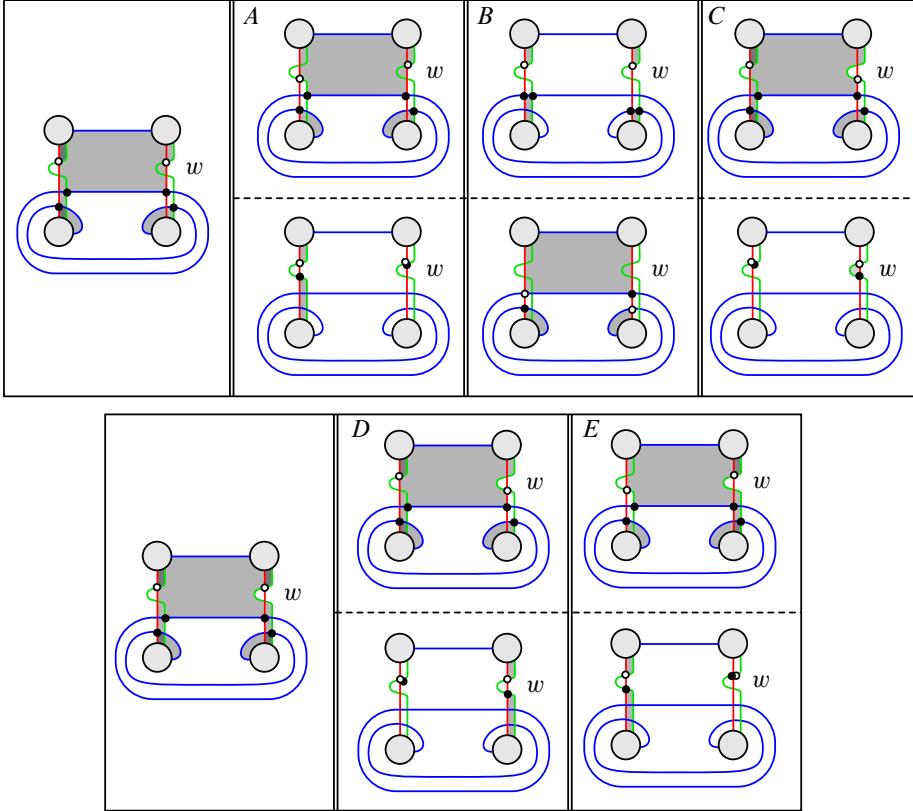


Figure 9.2. A Gromov compactness argument which shows that the total count of holomorphic triangles of index 0 in $\pi_2(\Theta^+, \Theta^+, \theta^+ | \bar{\theta}^-)$ is 1. In the left-most columns are index 1 classes of triangles. On the other columns of a given row, we show all decompositions of that class into an index 0 triangle class and an index 1 disk class (such that both classes have nonnegative multiplicities). Configurations A and E have the same total count as the $\theta^+ | \bar{\theta}^-$ coefficient of our map. Configuration B has total count 1. Configurations C and D have the same count. Gromov compactness gives $\#A + \#B + \#C \equiv 0$ and $\#D + \#E \equiv 0$, which quickly gives equation (9.5).

second is a summand which has a factor of

$$F_3^{\alpha\gamma_0, \alpha\gamma_0} \circ f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} \circ \tau^{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} \circ F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}. \quad (9.6)$$

We consider this latter summand first. The map in equation (9.6) counts holomorphic triangles weighted by

$$\#\tau \cap \partial_{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}(\psi).$$

If $\psi \in \pi_2(x, y, z)$ is a class of triangles counted in the above composition, then both x and y (the incoming intersection points) have $\theta_{\gamma_0, \gamma_0}^+$ as factors. By the stabilization

result of Proposition 6.5, if such a class has a representative, then z must have a factor of $\theta_{\gamma_0, \gamma_0}^+$ as well. However, such a term would evaluate trivially under $F_3^{\alpha\gamma_0, \alpha\gamma_0}$. Hence, equation (9.6) vanishes.

We now consider the first summand, which involves the choice of η . Let τ and τ' be two choices of curves, as above, and let η_τ be some choice of diagonal chain for the hypercube from equation (9.3), when τ is used. It is straightforward to see that $[\tau] - [\tau'] \in H_1(\Sigma' \# \bar{\Sigma}')$ lies in the span of the Δ curves. Let C be a 2-chain on Σ such that $\partial\Sigma$ is $\tau - \tau'$ plus a linear combination of small parallel pushoffs of the Δ curves.

We define the 2-chain

$$\eta_{\tau'} := \eta_\tau + H_C(\Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}).$$

In the above, H_C denotes the $\mathbb{F}[U]$ -linear map whose value on an intersection point $x \in \mathbb{T}_{\Delta\gamma_0\bar{\gamma}_0} \cap \mathbb{T}_{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}$ is the sum of the multiplicities of C at each factor of x . It is straightforward to see that $\eta_{\tau'}$ is a valid choice of diagonal for the hypercube in equation (9.3) which uses τ' . Since each summand of $H_C(\Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0})$ has $\theta_{\gamma_0, \gamma_0}^+$ as a factor, the same argument as above shows that

$$F_3^{\gamma_0, \gamma_0} \circ f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}(F_1^{\gamma_0, \gamma_0}(x), H_C(\Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0})) = 0,$$

so the diagonal maps computed using η_τ and $\eta_{\tau'}$ coincide. \blacksquare

Lemma 9.5. *Suppose that $\mathcal{H} = (\Sigma, \alpha, \beta)$ and $\mathcal{H}' = (\Sigma, \alpha', \beta')$ are two Heegaard diagrams for Y with the same underlying Heegaard surface, such that α' is obtained from α by an elementary handleslide or isotopy. We assume β' is obtained similarly from β . Let \mathfrak{D} and \mathfrak{D}' denote doubling enhancements of \mathcal{H} and \mathcal{H}' with choices of almost complex structures and doubling curves Δ and Δ' . Let \mathbb{S}_0 denote a 0-sphere which is embedded in $\Sigma \setminus (\alpha \cup \beta \cup \alpha' \cup \beta')$ and let $\Sigma(\mathbb{S}_0)$ denote the result of attaching a 1-handle along \mathbb{S}_0 to Σ . Let Δ_0 and Δ'_0 denote extra choices of doubling arcs, as in the construction of the 1-handle map, and let $\mathfrak{D}(\mathbb{S}_0)$ and $\mathfrak{D}'(\mathbb{S}_0)$ denote the doubling enhancements of $(\Sigma(\mathbb{S}_0), \alpha\gamma_0, \beta\gamma_0)$ and $(\Sigma(\mathbb{S}_0), \alpha'\gamma_0, \beta'\gamma_0)$ by using the new doubling curves Δ_0 and Δ'_0 , respectively. Let τ be a single choice of symmetric 1-cycle on $\Sigma(\mathbb{S}_0) \# \bar{\Sigma}(\mathbb{S}_0)$ which intersects γ_0 transversely at a single point. Finally, let $CFI(W(\mathbb{S}_0))$ and $CFI(W(\mathbb{S}_0))'$ denote the involutive 1-handle maps, computed with these two choices of data. Then*

$$\Psi_{\mathfrak{D}(\mathbb{S}_0) \rightarrow \mathfrak{D}'(\mathbb{S}_0)} \circ CFI(W(\mathbb{S}_0)) \simeq CFI(W(\mathbb{S}_0))' \circ \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}. \quad (9.7)$$

Proof. The proof is to enlarge the diagram in (9.1) into a 3-dimensional hyperbox by adding an extra direction coming out of the page. One face will correspond to the map $CFI(W(\mathbb{S}_0))$ and an opposite face will correspond to the map $CFI(W(\mathbb{S}_0))'$. It will

be straightforward to unpack the hypercube relations to obtain (9.7). We presently describe the construction of the hyperbox.

We extend the top-most hypercube of (9.1) to the hyperbox shown in (9.8). The unlabeled arrows are holomorphic triangle maps. There are no morphisms of length 2 or 3. The hypercube relations follow from the stabilization formula of Proposition 6.5.

$$\begin{array}{ccccc}
 CF(\alpha, \beta) & \xrightarrow{F_1^{\gamma_0, \gamma_0}} & CF(\alpha\gamma_0, \beta\gamma_0) & & \\
 \downarrow F_1^{\bar{\beta}, \bar{\beta}} & \searrow & \downarrow F_1^{\gamma_0, \gamma_0} & \searrow & \\
 CF(\alpha', \beta) & & CF(\alpha'\gamma_0, \beta\gamma_0) & & \\
 \downarrow F_1^{\bar{\beta}', \bar{\beta}} & & \downarrow F_1^{\bar{\beta}\bar{\gamma}_0, \bar{\beta}\bar{\gamma}_0} & & \\
 CF(\alpha\bar{\beta}, \beta\bar{\beta}) & \xrightarrow{F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) & & \\
 \downarrow F_1^{\bar{\beta}', \bar{\beta}} & \searrow & \downarrow F_1^{\bar{\beta}'\bar{\gamma}_0, \bar{\beta}'\bar{\gamma}_0} & & \\
 CF(\alpha'\bar{\beta}', \beta\bar{\beta}) & \xrightarrow{F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) & & \\
 & & & & \\
 & & CF(\alpha', \beta) & \xrightarrow{F_1^{\gamma_0, \gamma_0}} & CF(\alpha'\gamma_0, \beta\gamma_0) \\
 \downarrow F_1^{\bar{\beta}', \bar{\beta}} & \searrow & \downarrow F_1^{\gamma_0, \gamma_0} & \searrow & \\
 CF(\alpha', \beta') & & CF(\alpha'\gamma_0, \beta'\gamma_0) & & \\
 \downarrow F_1^{\bar{\beta}', \bar{\beta}'} & & \downarrow F_1^{\bar{\beta}'\bar{\gamma}_0, \bar{\beta}'\bar{\gamma}_0} & & \\
 CF(\alpha'\bar{\beta}', \beta\bar{\beta}) & \xrightarrow{F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) & & \\
 \downarrow F_1^{\bar{\beta}', \bar{\beta}'} & \searrow & \downarrow F_1^{\bar{\beta}'\bar{\gamma}_0, \bar{\beta}'\bar{\gamma}_0} & & \\
 CF(\alpha'\bar{\beta}', \beta'\bar{\beta}') & \xrightarrow{F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \beta'\bar{\beta}'\gamma_0\bar{\gamma}_0) & &
 \end{array} \tag{9.8}$$

The hypercubes obtained by extending the two small left-most cubes of (9.1) are similarly defined to be hyperboxes of stabilizations.

We now consider the right-most cubes of (9.1) (i.e., the two cubes which have length 2-maps). The top-most cube is easily extended by building hypercubes of

attaching curves where each length 1 chain is a cycle generating the top degree of HF^- . We leave the details to the reader.

We now consider the bottom-most cube of (9.1) which has a length 2 map, and we describe its extension to a 3-dimensional hyperbox. (Recall that this hypercube involved more choices than usual, such as the choice of the chain η and the curve τ ; see Remark 9.1.) The construction is obtained by pairing hypercubes of attaching curves. The main technical challenge in extending the 2-dimensional cube in (9.3) is building a 3-dimensional hypercube of attaching curves with the following form:

$$\begin{array}{ccccc}
 \Delta \gamma_0 \bar{\gamma}_0 & \xrightarrow{\quad} & \Delta \Delta_0 & & \\
 \downarrow & \searrow & \downarrow & \searrow & \\
 \Delta' \gamma_0 \bar{\gamma}_0 & \xrightarrow{\quad} & \Delta' \Delta'_0 & & \\
 \downarrow \eta & \downarrow & \downarrow \eta' & & \\
 \alpha \bar{\alpha} \gamma_0 \bar{\gamma}_0 & \xrightarrow{\quad \tau \quad} & \alpha \bar{\alpha} \gamma_0 \bar{\gamma}_0 & & \\
 \downarrow & \searrow & \downarrow h & \searrow & \\
 \alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0 & \xrightarrow{\quad \tau \quad} & \alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0 & &
 \end{array} \tag{9.9}$$

In equation (9.9), the chains η and η' are the ones used to construct $CFI(W(\mathbb{S}_0))$ and $CFI(W(\mathbb{S}_0))'$, as constructed in Lemma 9.2. In particular, these morphisms are fixed, and similarly the top length 2 morphism has also already been chosen in the previous step, so we assume that chain is fixed. Note that the two morphisms from $\alpha \bar{\alpha} \gamma_0 \bar{\gamma}_0$ to $\alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0$ are forced to coincide, since α' is obtained from α by an elementary handleslide, so there is a single top degree generator of $CF(\alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0, \alpha \bar{\alpha} \gamma_0 \bar{\gamma}_0)$, which is

$$\Theta_{\alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0} \otimes \theta_{\gamma_0 \bar{\gamma}_0}^+.$$

Assuming for the moment the existence of the hypercube in (9.9), one extends the bottom-most cube of (9.1) by constructing two hypercubes. The first hypercube is obtained by pairing the back face of (9.9) with the hypercube

$$\alpha \bar{\beta} \xrightarrow{\Theta_{\alpha \bar{\beta}, \alpha' \bar{\beta}'}} \alpha' \bar{\beta}'.$$

The second hypercube is obtained by pairing the entirety of equation (9.9) with the 0-dimensional hypercube $\alpha' \bar{\beta}'$.

We now prove the existence of the hypercube in (9.9). The existence of length 2 chains in equation (9.9) along the top, left and right faces is straightforward.

We now consider the bottom face of (9.9). The desired relation is

$$A_\tau(\Theta_{\alpha\bar{\alpha}\bar{\gamma}_0,\alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0}) + B_\tau(\Theta_{\alpha\bar{\alpha}\bar{\gamma}_0,\alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0}) = \delta(h),$$

where h denotes the diagonal chain. Since τ is a closed curve, we have an equality $A_\tau = B_\tau$, the above equation is satisfied by $h = 0$. (Compare equation (7.5) for arcs instead of closed curves). As we will see, setting $h = 0$ might not allow the length 3 hypercube relation to be satisfied, so instead, we use the chain

$$h = \varepsilon \cdot \Theta_{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0, \alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0}, \quad (9.10)$$

where $\varepsilon \in \mathbb{F}$ is to be determined.

Finally, we claim that we may make a choice of the constant ε in equation (9.10), and also pick a length 3 chain for the cube in (9.9) so that the hypercube relations are satisfied. To see this, let $C \in CF(\Delta \gamma_0 \bar{\gamma}_0, \alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0)$ denote the sum of all compositions featuring in the length 3 hypercube relations, except for the length 3 map itself. Here, we momentarily set $\varepsilon = 0$ in (9.10). The hypercube relations on the proper faces force C to be a cycle. We note however that C has the same grading as the top degree element generator of

$$HF^-(\Delta \gamma_0 \bar{\gamma}_0, \alpha' \bar{\alpha}' \gamma_0 \bar{\gamma}_0).$$

In particular, it either represents the top degree element or is null-homologous. If it is null-homologous, we set $\varepsilon = 0$ in (9.9) and pick any primitive of C in the appropriate grading to be the length 3 map in our hypercube. If C represents the top degree generator, then we pick $\varepsilon = 1$, and let the length 3 chain be any primitive of

$$f_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0, \alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0}(\Theta_{\Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}, \Theta_{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0, \alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0}) + C.$$

The hypercube relations are clearly satisfied.

Finally, it remains to extend the bottom-most cube of Figure 9.1. To extend this hypercube, it is convenient to consider the expanded description in (9.2). In fact, the only hypercube which requires explanation to extend is the cube

We define our 3-dimensional extension of this cube to have trivial length 3 map. The hypercube relations for the 3-dimensional extension of this cube are proven somewhat similarly to the proof of Proposition 8.2. The only non-trivial length 2 map in the

above diagram is in the top subface. This subface has the following form:

$$\begin{array}{ccccccc}
 CF(\alpha\bar{\beta}\bar{\gamma}_0, \alpha\bar{\alpha}\bar{\gamma}_0) & \xrightarrow{F_1^{\gamma_0, \gamma_0}} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \xrightarrow{B_\tau} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \downarrow & \\
 \downarrow & & \downarrow & \searrow & \downarrow & & \\
 CF(\alpha'\bar{\beta}'\bar{\gamma}_0, \alpha\bar{\alpha}\bar{\gamma}_0) & \xrightarrow{F_1^{\gamma_0, \gamma_0}} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \xrightarrow{B_\tau} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \downarrow & \\
 \downarrow & & \downarrow & \searrow & \downarrow & & \\
 CF(\alpha'\bar{\beta}'\bar{\gamma}_0, \alpha'\bar{\alpha}'\bar{\gamma}_0) & \xrightarrow{F_1^{\gamma_0, \gamma_0}} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0) & \xrightarrow{B_\tau} & CF(\alpha'\bar{\beta}'\gamma_0\bar{\gamma}_0, \alpha'\bar{\alpha}'\gamma_0\bar{\gamma}_0) & \downarrow &
 \end{array}$$

It is sufficient to show that the compression of this diagram composes trivially with the map $F_3^{\gamma_0, \gamma_0}$. The key tool we will use is that the triangle maps, when applied to elements of the form $(x \times \theta_{\gamma_0, \gamma_0}^+) \otimes (y \times \theta_{\gamma_0, \gamma_0}^+)$ have only summands of the form $z \times \theta_{\gamma_0, \gamma_0}^+$. This fact follows from Proposition 6.5. The top diagonal map counts holomorphic triangles which are weighted by the quantity $\#(\partial_{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}(\psi) \cap \tau) \in \mathbb{F}$. Similarly, the bottom triangle map is a sum of holomorphic triangle maps with certain weights. The special inputs for the length 1 triangle maps all have tensor factors of $\theta_{\gamma_0, \gamma_0}^+$, as does the special length 2 element used for the bottom diagonal map by equation (9.10). Hence, in compression of the diagram, the only holomorphic triangles which are counted have only tensor factors of $\theta_{\gamma_0, \gamma_0}^+$, and hence so does their output. In particular, the output vanishes when composed with $F_3^{\gamma_0, \gamma_0}$. ■

9.3. 3-handles

We define the 3-handle map by dualizing the construction of the 1-handle map. Concretely, this amounts to compressing the diagram shown in the following equation:

$$\begin{array}{ccccc}
 CF(\alpha\gamma_0, \beta\gamma_0) & \xrightarrow{F_3^{\gamma_0, \gamma_0}} & CF(\alpha, \beta) & & \\
 \downarrow F_1^{\beta\bar{\gamma}_0, \beta\bar{\gamma}_0} & & \downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \\
 CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) & \xrightarrow{B_\tau} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \beta\bar{\beta}\gamma_0\bar{\gamma}_0) & \xrightarrow{F_3^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha\bar{\beta}, \beta\bar{\beta}) \\
 \downarrow f_{\alpha\beta\gamma_0\bar{\gamma}_0}^{\beta\bar{\beta}\gamma_0\bar{\gamma}_0 \rightarrow \Delta\Delta_0} & \searrow & \downarrow f_{\alpha\beta\gamma_0\bar{\gamma}_0}^{\beta\bar{\beta}\gamma_0\bar{\gamma}_0 \rightarrow \Delta\gamma_0\bar{\gamma}_0} & \searrow & \downarrow f_{\alpha\beta}^{\beta\bar{\beta} \rightarrow \Delta} \\
 CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \Delta\Delta_0) & \xrightarrow{f_{\alpha\bar{\gamma}_0\bar{\gamma}_0}^{\Delta\Delta_0 \rightarrow \gamma_0\bar{\gamma}_0}} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \Delta\gamma_0\bar{\gamma}_0) & \xrightarrow{F_3^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha\bar{\beta}, \Delta) \\
 \downarrow f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} & \searrow & \downarrow f_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} & \searrow & \downarrow f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} \\
 CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \xlongequal{F_3^{\alpha\gamma_0, \alpha\gamma_0}} & CF(\alpha\bar{\beta}\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0) & \xrightarrow{F_3^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}} & CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) \\
 \downarrow & & \downarrow & & \downarrow F_3^{\alpha\bar{\alpha}, \alpha\bar{\alpha}} \\
 CF(\bar{\beta}\bar{\gamma}_0, \bar{\alpha}\bar{\gamma}_0) & \xrightarrow{F_3^{\bar{\gamma}_0, \bar{\gamma}_0}} & & & CF(\bar{\beta}, \bar{\alpha})
 \end{array}$$

9.4. Attaching 1-handles near the basepoint

We may obtain a slightly simpler formula for the 1-handle map when we perform the attachment near the basepoint w .

Lemma 9.6. *Suppose $\mathbb{S}_0 \subseteq Y$ is a 0-sphere and Y is connected. We have an isomorphism of groups*

$$CFI^-(Y(\mathbb{S}_0)) \cong CFI^-(Y) \otimes_{\mathbb{F}} \langle \theta^+, \theta^- \rangle.$$

With respect to this decomposition, the involutive 1-handle map takes the form

$$CFI(W(\mathbb{S}_0))(x) = x \times \theta^+,$$

extended equivariantly over $\mathbb{F}[U, Q]/Q^2$. The 3-handle map takes the form $x \times \theta^+ \mapsto 0$ and $x \times \theta^- \mapsto x$, extended $\mathbb{F}[U, Q]/Q^2$ -equivariantly.

Proof. We focus on the 1-handle map. We pick our diagram so that, before doubling, there is a special genus 1 region near the basepoint, and so, the doubled arcs have the configuration shown in Figure 9.1 in the double of this region. We form the compression of (9.1) by first compression horizontally, then vertically.

We first show that for suitable almost complex structures, the diagonal maps make trivial contribution to compression of the diagram. We consider the top-most diagonal map. This map counts holomorphic quadrilaterals on the connected sum

$$(\alpha\bar{\beta}, \beta\bar{\beta}, \Delta, \Delta) \# (\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0, \Delta_0).$$

Note that each Heegaard subdiagram of the right-handle quadruple has vanishing differential (since it has the minimal number of generators possible; see Figure 9.1). Hence, the stabilization results for holomorphic quadrilaterals from Proposition 6.5 imply that

$$h_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\beta\bar{\beta}\gamma_0\bar{\gamma}_0 \rightarrow \Delta\gamma_0\bar{\gamma}_0 \rightarrow \Delta\Delta_0} \circ F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0} = F_1^{\gamma_0\bar{\gamma}_0, \Delta_0} \circ h_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta \rightarrow \Delta}.$$

By the small translate theorem for holomorphic quadrilaterals [4, Proposition 11.5], the map $h_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta \rightarrow \Delta}$ vanishes.

We now consider the bottom-most diagonal map in (9.1). Because of our choice of diagram, the chain η in (9.3) may be chosen to vanish, by the same model computation we used to verify Lemma 9.2. Hence, the dashed length 2 map is the sum

$$h_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \Delta\Delta_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} + h_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0 \xrightarrow{\tau} \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}. \quad (9.11)$$

Here, $h_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0 \xrightarrow{\tau} \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}$ denotes the map which counts holomorphic triangles on the diagram

$$(\Sigma' \# \bar{\Sigma}', \alpha\bar{\beta}\gamma_0, \bar{\gamma}_0, \Delta\gamma_0\bar{\gamma}_0, \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0),$$

which are weighted by an extra factor of $\partial_{\alpha\bar{\alpha}\gamma_0\bar{\gamma}_0} D(\psi) \cdot \tau$.

We claim that both of these maps in equation (9.11) vanish when precomposed with $F_1^{\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0}$ and post-composed with $F_3^{\alpha\gamma_0, \alpha\gamma_0}$. Consider the first quadrilateral map. It counts curves on the diagram

$$(\alpha\bar{\beta}, \Delta, \Delta, \alpha\bar{\alpha}) \# (\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0, \Delta_0, \gamma_0\bar{\gamma}_0).$$

We can use Proposition 6.5 to destabilize the quadruple, noting that the evaluation of the right-hand triple on a degree 0 tree of $C_*^{\text{cell}}(K_4)$ is

$$\theta^+|\bar{\theta}^- + \theta^-|\bar{\theta}^+$$

by Lemma 9.2. Hence, $\theta_{\gamma_0, \gamma_0}^-$ factor of the composition is

$$h_{\alpha\bar{\beta}}^{\Delta \rightarrow \Delta \rightarrow \alpha\bar{\alpha}}(x) \otimes \theta_{\gamma_0, \bar{\gamma}_0}^+.$$

This expression vanishes by the small translate theorem for quadrilaterals [4, Proposition 11.5], since Δ appears twice.

Finally, we argue that the triangle counting map $h_{\alpha\bar{\beta}\gamma_0\bar{\gamma}_0}^{\Delta\gamma_0\bar{\gamma}_0 \rightarrow \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0 \xrightarrow{\tau} \alpha\bar{\alpha}\gamma_0\bar{\gamma}_0}$ also makes trivial contribution. In this case, the triangle count occurs on

$$(\alpha\bar{\beta}, \Delta, \alpha\bar{\alpha}) \# (\gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0, \gamma_0\bar{\gamma}_0).$$

Proposition 6.5 implies there are no triangles with $\theta_{\gamma_0, \gamma_0}^-$ as output when only terms involving $\theta_{\gamma_0, \gamma_0}^+$ are the input, so this triangle map makes trivial contribution to the compressed diagram.

The argument for 3-handles follows because the 3-handle map is defined dually to the 1-handle map, so the above argument immediately translates after taking duals. ■

9.5. Commutation of 1-handles

Lemma 9.7. *Suppose that \mathbb{S} and \mathbb{S}' are two embedded and disjoint spheres in Y , whose dimensions lie in $\{0, 2\}$. Then*

$$CFI(W(Y(\mathbb{S}), \mathbb{S}')) \circ CFI(W(Y, \mathbb{S})) \simeq CFI(W(Y(\mathbb{S}'), \mathbb{S})) \circ CFI(W(Y, \mathbb{S}')).$$

Lemma 9.7 may be proven by using the model computations from Lemma 9.6. Alternatively, it may be proven by building a 3-dimensional hyperbox which has two faces corresponding to the definition of the 1-handle maps from (9.1).

10. 2-handles

In this section, we describe the map for a framed link $\mathbb{L} \subseteq Y$, with a self-conjugate Spin^c structure \mathbf{s} on $W(Y, \mathbb{L})$ (or more generally, a family $\mathfrak{S} \subseteq \text{Spin}^c(W(Y, \mathbb{L}))$ which is closed under conjugation). As in Section 9, we may omit the path and framing from our notation.

The existence of a map for 2-handles was described by Hendricks and Manolescu [5], though they were not able to prove invariance of the construction. The construction we present in this section should be thought of as a systematic version of the construction therein using the doubling operation.

10.1. The construction

Suppose \mathbb{L} is a framed link in Y . We assume that the framing is Morse, so that there is a corresponding 2-handle cobordism $W(Y, \mathbb{L})$. We pick a bouquet for \mathbb{L} , in the sense of [24, Definition 4.1], as well as a Heegaard triple $(\Sigma, \alpha', \alpha, \beta)$ subordinate to this bouquet [24, Definition 4.2].

Remark 10.1. Like the 2-handle maps from [24], our 2-handle maps also depend on a bouquet for the link \mathbb{L} . Independence from the bouquet will be proven in Lemma 13.1.

We first construct several auxiliary hypercubes. We firstly have two hypercubes of alpha attaching circles

$$\alpha \xrightarrow{\Theta_{\alpha', \alpha}} \alpha' \quad \text{and} \quad \alpha \bar{\beta} \xrightarrow{\Theta_{\alpha' \bar{\beta}, \alpha \bar{\beta}}} \alpha' \bar{\beta} \quad (10.1)$$

and a hypercube of beta attaching circles

$$\beta \bar{\beta} \xrightarrow{\Theta_{\beta \bar{\beta}, \Delta}} \Delta. \quad (10.2)$$

There is a unique choice of top degree cycle in the hypercubes in equation (10.1), while the cycle $\Theta_{\beta \bar{\beta}, \Delta}$ in equation (10.2) is only well defined up to addition of a boundary.

Next, we construct a more complicated hypercube of beta attaching curves:

$$\begin{array}{ccccc} \Delta & \xrightarrow{\Theta_{\Delta, \alpha' \bar{\alpha}'}} & \alpha' \bar{\alpha}' & & \\ \downarrow \Theta_{\Delta, \alpha \bar{\alpha}} & \searrow \lambda_{\Delta, \alpha \bar{\alpha}'} & \downarrow \Theta_{\alpha' \bar{\alpha}', \alpha \bar{\alpha}'} & & \\ \alpha \bar{\alpha} & \xrightarrow{\Theta_{\alpha \bar{\alpha}, \alpha \bar{\alpha}'}} & \alpha \bar{\alpha}' & & \end{array} \quad (10.3)$$

We now prove that the hypercube in (10.3) may be constructed.

Lemma 10.2. *If $\Theta_{\Delta, \alpha'\bar{\alpha}'}, \Theta_{\alpha'\bar{\alpha}', \alpha\bar{\alpha}'}, \Theta_{\Delta, \alpha\bar{\alpha}}$ and $\Theta_{\alpha\bar{\alpha}, \alpha\bar{\alpha}'}$ are cycles which represent the top degree elements of their respective HF^- groups, then*

$$[f_{\Delta, \alpha'\bar{\alpha}', \alpha\bar{\alpha}'}(\Theta_{\Delta, \alpha'\bar{\alpha}'}, \Theta_{\alpha'\bar{\alpha}', \alpha\bar{\alpha}'})] = [f_{\Delta, \alpha\bar{\alpha}, \alpha\bar{\alpha}'}(\Theta_{\Delta, \alpha\bar{\alpha}}, \Theta_{\alpha\bar{\alpha}, \alpha\bar{\alpha}'})],$$

where the brackets denote the induced element in homology. In particular, $\lambda_{\Delta, \alpha\bar{\alpha}'} \in CF(\Delta, \alpha\bar{\alpha}')$ may be chosen so that (10.3) is a hypercube of attaching curves.

Proof. Both triples $(\Sigma \# \bar{\Sigma}, \Delta, \alpha'\bar{\alpha}', \alpha\bar{\alpha}')$ and $(\Sigma \# \bar{\Sigma}, \Delta, \alpha\bar{\alpha}, \alpha\bar{\alpha}')$ represent the same cobordism, topologically. Namely, they represent a cobordism from $(S^1 \times S^2)^{\# g}$ to $(S^1 \times S^2)^{g-|\mathbb{L}|}$ which is surgery on a $|\mathbb{L}|$ component link, each component of which is an S^1 -fiber of a $S^1 \times S^2$ -summand. \blacksquare

Since the element $[f_{\Delta, \alpha'\bar{\alpha}', \alpha\bar{\alpha}'}(\Theta_{\Delta, \alpha'\bar{\alpha}'}, \Theta_{\alpha'\bar{\alpha}', \alpha\bar{\alpha}'})]$ represents a top-degree class of $CF(S^1 \times S^2)^{g-|\mathbb{L}|}$, we have that any two choices of $\lambda_{\Delta, \alpha\bar{\alpha}'}$ satisfying (10.3) are homologous.

Now, suppose that $\mathfrak{S} \subseteq \text{Spin}^c(W(Y, \mathbb{L}))$ is a finite set which is closed under conjugation. We now define the 2-handle map for \mathbb{L} , which denote by $CFI(W(Y, \mathbb{L}), \mathfrak{S})$, using the hypercubes in (10.1), (10.2) and (10.3). We define the map $CFI(W(Y, \mathbb{L}), \mathfrak{S})$ to be the compression of the hyperbox shown in (10.4). A priori, $CFI(W(Y, \mathbb{L}), \mathfrak{S})$ depends on the choices in its construction; for well-definedness see Section 10.2.

$$\begin{array}{ccccc}
CF(\alpha, \beta) & \xrightarrow{f_{\alpha \rightarrow \alpha'}^\beta} & CF(\alpha', \beta) & & \\
\downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \\
CF(\alpha\bar{\beta}, \beta\bar{\beta}) & \xrightarrow{f_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^{\beta\bar{\beta}}} & CF(\alpha'\bar{\beta}, \beta\bar{\beta}) & & \\
\downarrow f_{\alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} & \searrow h_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} & \downarrow f_{\alpha'\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} & & \\
CF(\alpha\bar{\beta}, \Delta) & \xrightarrow{f_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^\Delta} & CF(\alpha'\bar{\beta}, \Delta) & & \\
\parallel & \searrow h_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^{\Delta \rightarrow \alpha'\bar{\alpha}'} & \downarrow f_{\alpha'\bar{\beta}}^{\Delta \rightarrow \alpha'\bar{\alpha}'} & & \\
CF(\alpha\bar{\beta}, \Delta) & \xrightarrow{f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha'\bar{\alpha}'}} & CF(\alpha\bar{\beta}, \alpha'\bar{\alpha}') & \xrightarrow{f_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^{\alpha'\bar{\alpha}'}} & CF(\alpha'\bar{\beta}, \alpha'\bar{\alpha}') \\
\downarrow f_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}} & \searrow H_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}'} & \downarrow f_{\alpha\bar{\beta}}^{\alpha'\bar{\alpha}' \rightarrow \alpha\bar{\alpha}'} & & \downarrow F_3^{\alpha', \alpha} \\
CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & \xrightarrow{f_{\alpha\bar{\beta}}^{\alpha\bar{\alpha} \rightarrow \alpha\bar{\alpha}'}} & CF(\alpha\bar{\beta}, \alpha\bar{\alpha}') & \xrightarrow{F_3^{\alpha, \alpha}} & CF(\bar{\beta}, \bar{\alpha}') \\
\downarrow F_3^{\alpha, \alpha} & & \downarrow F_3^{\alpha, \alpha} & & \parallel \\
CF(\bar{\beta}, \bar{\alpha}) & \xrightarrow{f_{\bar{\beta}}^{\bar{\alpha} \rightarrow \bar{\alpha}'}} & CF(\bar{\beta}, \bar{\alpha}') & = & CF(\bar{\beta}, \bar{\alpha}') \\
& & & & (10.4)
\end{array}$$

Lemma 10.3. *If we compress each horizontal level in (10.4), then the resulting diagram is a hyperbox of chain complexes.*

Proof. It suffices to show that each face is a hypercube of chain complexes. The subcubes with length 2 maps are all obtained by pairing hypercubes of attaching curves, so the claim is immediate for these faces. The top-most and bottom-most subcubes are hypercubes of stabilization, so the claim is immediate. The only remaining cube of interest is the second-bottom-most cube on the right. The hypercube relations amount to the relation

$$F_3^{\alpha',\alpha'} \circ f_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^{\alpha'\bar{\alpha}'} = F_3^{\alpha,\alpha} \circ f_{\alpha\bar{\beta}}^{\alpha'\bar{\alpha}' \rightarrow \alpha\bar{\alpha}'}.$$
 (10.5)

This is proven as follows. We observe that the Heegaard triples for these triangle maps are $(\Sigma, \alpha', \alpha, \alpha') \# (\bar{\Sigma}, \bar{\beta}, \bar{\beta}, \bar{\alpha}')$ and $(\Sigma, \alpha, \alpha', \alpha) \# (\bar{\Sigma}, \bar{\beta}, \bar{\beta}, \bar{\alpha}')$. The first triple is a stabilization of the small isotopy triple $(\bar{\Sigma}, \bar{\beta}, \bar{\beta}, \bar{\alpha}')$, and similarly the second triple is also a stabilization of a Heegaard triple for a small isotopy. Using the stabilization results of Proposition 6.5 and the small triangle theorem [4, Proposition 11.1], we obtain (10.5), completing the proof. ■

Remark 10.4. If $\mathfrak{S} \subseteq \text{Spin}^c(W(Y, \mathbb{L}))$ is a set of Spin^c structures which is closed under the conjugation action, then a cobordism map may be defined by summing the maps above for all $s \in \mathfrak{S}$. If \mathfrak{S} is infinite, one may need to use coefficients in $\mathbb{F}[[U]]$ to obtain a well-defined map.

10.2. Independence of simple handleslides

We now prove that the construction above commutes with the maps for simple handleslides and changes of the doubling arcs. We phrase our result in enough generality that they also imply that the 2-handle maps are independent of the choice of bouquet for \mathbb{L} , and are independent of handleslides amongst the components of \mathbb{L} .

Suppose that $(\Sigma, \alpha', \alpha, \beta)$ is a Heegaard triple subordinate to a bouquet of $\mathbb{L} \subseteq Y$. Suppose that $(\Sigma, \alpha'_H, \alpha_H, \beta_H)$ is another Heegaard triple, such that α'_H, α_H and β_H are obtained from α', α and β (respectively) by elementary handleslide equivalences, in the sense of Section 5. We assume, furthermore, that each of the following sets admits a unique top degree intersection point, which is a cycle in its respective Floer complex:

$$\mathbb{T}_{\alpha_H} \cap \mathbb{T}_\alpha, \quad \mathbb{T}_{\alpha'_H} \cap \mathbb{T}_{\alpha'}, \quad \mathbb{T}_{\alpha'_H} \cap \mathbb{T}_\alpha, \quad \text{and} \quad \mathbb{T}_\beta \cap \mathbb{T}_{\beta_H}.$$

We also suppose that we pick two sets of doubling arcs, Δ and Δ_H . We require no particular relation between Δ and Δ_H , except that the diagram $(\Sigma \# \bar{\Sigma}, \Delta, \Delta_H, w)$ be admissible.

Proposition 10.5. Suppose that $(\Sigma, \alpha', \alpha, \beta)$ and $(\Sigma, \alpha'_H, \alpha_H, \beta_H)$ are as above. Let $CFI(W(\mathbb{L}))$ and $CFI(W(\mathbb{L}))^H$ denote the 2-handle maps, computed with the two triples $(\Sigma, \alpha', \alpha, \beta)$ and $(\Sigma, \alpha'_H, \alpha_H, \beta_H)$, and the doubling arcs Δ and Δ_H , respectively. Let \mathfrak{D} denote the doubling data for Y determined by (Σ, α, β) , Δ , and some choices of almost-complex structures. Let \mathfrak{D}_H denote the doubling data determined by $(\Sigma, \alpha_H, \beta_H)$, Δ_H , and some choices of almost-complex structures. Let \mathfrak{D}' and \mathfrak{D}'_H denote the analogous data for $Y(\mathbb{L})$. Then

$$\Psi_{\mathfrak{D}' \rightarrow \mathfrak{D}'_H} \circ CFI(W(\mathbb{L})) \simeq CFI(W(\mathbb{L}))^H \circ \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}_H}.$$

Here, $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}_H}$ and $\Psi_{\mathfrak{D}' \rightarrow \mathfrak{D}'_H}$ denote the transition maps for elementary handleslides and changes of the doubling data, and \simeq denotes $\mathbb{F}[U, Q]/Q^2$ -equivariant chain homotopy.

Proof. Mirroring the construction of the maps $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}_H}$, we construct for each subcube of equation (10.4) a hyperbox of size $(2, 1, 1)$. The final coordinate corresponds to the involution. If we collapse this axis of the cube, we obtain a hyperbox of size $(2, 1)$. If we compress this hyperbox, we obtain the following hypercube, which realizes the homotopy commutation in the statement.

$$\begin{array}{ccccc} CFI(\mathfrak{D}) & \xrightarrow{F_{W(\mathbb{L})}} & CFI(\mathfrak{D}') & & \\ \searrow \Psi_{\mathfrak{D} \rightarrow \mathfrak{D}_H} & & & \swarrow \Psi_{\mathfrak{D}' \rightarrow \mathfrak{D}'_H} & \\ & \nearrow & & \searrow & \\ CFI(\mathfrak{D}_H) & \xrightarrow{F_{W(\mathbb{L})}^H} & CFI(\mathfrak{D}'_H) & & \end{array}$$

The top-most hyperbox is obtained by extending the top-most hypercube of equation (10.4) out of the page, and is obtained by stacking two hypercubes of stabilization. We leave the details to the reader.

We now consider the hypercube in (10.4) which is second from the top. The corresponding hyperbox is obtained by pairing hypercubes, as follows. We think of a hyperbox of size $(2, 1, 1)$ as being obtained by stacking two hypercubes. The first hypercube is obtained by constructing and pairing the following hypercubes of attaching curves, where the length-1 edges are obtained using the unique top-degree generators specified earlier; the 2-faces are filled as in Lemma 10.2:

$$\begin{array}{ccc} \alpha \bar{\beta} & \xrightarrow{\quad} & \alpha' \bar{\beta} \\ \searrow & \nearrow & \searrow \\ & \alpha_H \bar{\beta}_H & \xrightarrow{\quad} \alpha'_H \bar{\beta}_H \end{array} \quad \text{and} \quad \begin{array}{c} \beta \bar{\beta} \\ \downarrow \\ \Delta \end{array}.$$

The next hypercube is obtained by constructing and pairing the following hypercubes of attaching curves:

$$\alpha_H \bar{\beta}_H \longrightarrow \alpha'_H \bar{\beta}_H$$

and

$$\begin{array}{ccc} \beta \bar{\beta} & & \\ \downarrow & \searrow & \\ \Delta & & \beta_H \bar{\beta}_H \\ & \swarrow & \downarrow \\ & & \Delta_H \end{array}$$

The remaining hypercubes which have a length 2 map are constructed by pairing hypercubes of attaching curves in a straightforward manner. We remark that the central-most hypercube of (10.4) is obtained by a small rearrangement of the pairing of an alpha hypercube with a beta hypercube. The corresponding 3-dimensional hyperbox is obtained by performing the analogous rearrangement to two hypercubes obtained by pairing.

Finally, we consider the three bottom-most hypercubes of (10.4) which do not have a length-2 map. We consider the right hypercube in the second to bottom-most level.

$$\begin{array}{ccccc} CF(\alpha \bar{\beta}, \alpha' \bar{\alpha}') & \longrightarrow & CF(\alpha' \bar{\beta}, \alpha' \bar{\alpha}') & & \\ \downarrow & \searrow & \downarrow & \searrow & \\ CF(\alpha_H \bar{\beta}_H, \alpha' \bar{\alpha}') & & CF(\alpha'_H \bar{\beta}_H, \alpha' \bar{\alpha}') & & \\ & \downarrow F_3^{\alpha', \alpha'} & & \downarrow F_3^{\alpha'_H, \alpha'} & \\ CF(\alpha \bar{\beta}, \alpha \bar{\alpha}') & \xrightarrow{F_3^{\alpha, \alpha'}} & CF(\bar{\beta}, \bar{\alpha}) & & \\ \downarrow & \searrow & \downarrow & \searrow & \\ CF(\alpha_H \bar{\beta}_H, \alpha \bar{\alpha}') & \xrightarrow{F_3^{\alpha_H, \alpha}} & CF(\bar{\beta}_H, \bar{\alpha}') & & \end{array}$$

$$\begin{array}{ccccc}
CF(\alpha_H \bar{\beta}_H, \alpha' \bar{\alpha}') & \longrightarrow & CF(\alpha'_H \bar{\beta}_H, \alpha' \bar{\alpha}') & & \\
\downarrow & \searrow & \downarrow & \searrow & \\
CF(\alpha_H \bar{\beta}_H, \alpha'_H \bar{\alpha}'_H) & & & & CF(\alpha'_H \bar{\beta}_H, \alpha'_H \bar{\alpha}'_H) \\
& & F_3^{\alpha'_H, \alpha'} & & \\
\downarrow & \searrow & \downarrow & & \downarrow F_3^{\alpha'_H, \alpha'_H} \\
CF(\alpha_H \bar{\beta}_H, \alpha \bar{\alpha}') & \xrightarrow{F_3^{\alpha_H, \alpha}} & CF(\bar{\beta}_H, \bar{\alpha}') & & \\
\downarrow & \searrow & \downarrow & & \downarrow \\
CF(\alpha_H \bar{\beta}_H, \alpha_H \bar{\alpha}'_H) & \xrightarrow{F_3^{\alpha_H, \alpha_H}} & CF(\bar{\beta}_H, \bar{\alpha}'_H) & &
\end{array} \tag{10.6}$$

In (10.6), all of the length 1 and 2 maps which are unlabeled are holomorphic triangle maps or quadrilateral maps. There are no length 3 maps.

The length 2 hypercube relations for the diagrams in (10.6) are verified as follows. For the faces which have length 2 maps, the hypercube relations are obvious. For the other faces, the length 2 relations are proven similarly to the proof of (10.5).

It remains to prove the length 3 relations for the diagrams in (10.6). To verify these, we note that each summand in the length 3 relation is a composition of a sum of quadrilateral maps followed by a 3-handle map. We claim that each summand vanishes identically. To see this, consider the length 2 map along the top face of the top cube of (10.6). This map is the sum of two quadrilateral maps, which take place on the diagrams $(\Sigma, \alpha'_H, \alpha', \alpha, \alpha') \# (\bar{\Sigma}, \bar{\beta}_H, \bar{\beta}, \bar{\beta}, \bar{\alpha}')$ and $(\Sigma, \alpha'_H, \alpha_H, \alpha, \alpha') \# (\bar{\Sigma}, \bar{\beta}_H, \bar{\beta}_H, \bar{\beta}, \bar{\alpha}')$. In particular, both diagrams are algebraically rigid and admissible multi-stabilizations. By our result on stabilizations in Proposition 6.5, we see that

$$F_3^{\alpha'_H, \alpha_H} \circ h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta} \rightarrow \alpha'_H \bar{\beta}_H}^{\alpha' \bar{\alpha}'} = h_{\bar{\beta} \rightarrow \bar{\beta} \rightarrow \bar{\beta}_H}^{\bar{\alpha}'} \circ F_3^{\alpha, \alpha'},$$

however, $h_{\bar{\beta} \rightarrow \bar{\beta} \rightarrow \bar{\beta}_H}^{\bar{\alpha}'} = 0$ by the small translate theorem for quadrilaterals [4, Proposition 11.5], since $\bar{\beta}$ is repeated twice in the diagram. The same argument works for all of the other summands, verifying the hypercube relations for the diagrams in (10.6).

One now constructs hypercubes for the bottom two faces of (10.4). The construction of these hypercubes is similar to the ones constructed above, and we leave the details to the reader.

There is one final hyperbox to construct, corresponding to the lowest level of Figure 5.1. This hyperbox is constructed easily by pairing hypercubes of attaching curves and rearranging the result slightly. We leave the details to the reader. ■

10.3. The composition law

In this section, we prove the composition law for cobordisms with only 2-handles.

Proposition 10.6. *Suppose that \mathbb{L}_1 and \mathbb{L}_2 are two framed links in Y , and $\mathfrak{s}_1 \in \text{Spin}^c(W(Y, \mathbb{L}_1))$ and $\mathfrak{s}_2 \in \text{Spin}^c(W(Y(\mathbb{L}_1), \mathbb{L}_2))$ are two self-conjugate Spin^c structures. Then*

$$CFI(W(Y(\mathbb{L}_1), \mathbb{L}_2), \mathfrak{s}_2) \circ CFI(W(Y, \mathbb{L}_1), \mathfrak{s}_1) \simeq CFI(W(Y, \mathbb{L}_1 \cup \mathbb{L}_2), \mathfrak{S}(\mathfrak{s}_1, \mathfrak{s}_2)),$$

where $\mathfrak{S}(\mathfrak{s}_1, \mathfrak{s}_2)$ is the set of Spin^c structures on $W(Y, \mathbb{L}_1 \cup \mathbb{L}_2)$ which restrict to \mathfrak{s}_1 and \mathfrak{s}_2 .

Proof. The proof of Proposition 10.6 is to construct a 3-dimensional hypercube with the following properties.

- (1) One direction corresponds to the involution.
- (2) One face corresponds to the cobordism map $CFI(W(Y, \mathbb{L}_1), \mathfrak{s}_1)$.
- (3) One face corresponds to the cobordism map $CFI(W(Y(\mathbb{L}_1), \mathbb{L}_2), \mathfrak{s}_2)$.
- (4) One face corresponds to the identity map from $CFI(Y)$ to itself.
- (5) One face corresponds to the map $CFI(W(Y, \mathbb{L}_1 \cup \mathbb{L}_2), \mathfrak{S}(\mathfrak{s}_1, \mathfrak{s}_2))$.

If we view the 3-dimensional hypercube as forming a 2-dimensional hypercube involving the CFI complexes, then the corresponding shape is as follows:

$$\begin{array}{ccccc} CFI(Y) & \xrightarrow{F_{W(Y, \mathbb{L}_1), \mathfrak{s}_1}} & CFI(Y(\mathbb{L}_1)) & & \\ & \searrow F_{W(Y, \mathbb{L}_1 \cup \mathbb{L}_2), \mathfrak{S}(\mathfrak{s}_1, \mathfrak{s}_2)} & & \swarrow F_{W(Y(\mathbb{L}_1), \mathbb{L}_2), \mathfrak{s}_2} & \\ & & CFI(Y(\mathbb{L}_1, \mathbb{L}_2)) & \xrightarrow{\text{id}} & CFI(Y(\mathbb{L}_1, \mathbb{L}_2)) \end{array}$$

As a first step, pick a Heegaard quadruple $(\Sigma, \alpha'', \alpha', \alpha, \beta)$ satisfying the following properties.

- (1) $(\Sigma, \alpha', \alpha, \beta)$ is subordinate to a bouquet of the link $\mathbb{L}_1 \subseteq Y$.
- (2) $(\Sigma, \alpha'', \alpha', \beta)$ is subordinate to a bouquet of the link $\mathbb{L}_2 \subseteq Y(\mathbb{L}_1)$.
- (3) $(\Sigma, \alpha'', \alpha, \beta)$ is subordinate to a bouquet of the link $\mathbb{L}_1 \cup \mathbb{L}_2 \subseteq Y$.

Let $\Theta_{\alpha', \alpha}$, $\Theta_{\alpha'', \alpha'}$ and $\Theta_{\alpha'', \alpha}$ be the canonical generators. As in Ozsváth and Szabó's original proof for HF^- , an easy model computation implies that the following diagram is a hypercube of alpha attaching curves:

$$\begin{array}{ccccc} \alpha & \xrightarrow{\Theta_{\alpha', \alpha}} & \alpha' & & \\ \downarrow \Theta_{\alpha'', \alpha} & & \downarrow \Theta_{\alpha'', \alpha'} & & \\ \alpha'' & \xrightarrow{1} & \alpha'' & & \end{array} \tag{10.7}$$

Pairing with the 0-dimensional hypercube β and restricting to the Spin^c structures in the statement gives the non-involutive version of Proposition 10.6. Note that we can stabilize (10.7) to obtain the following hypercube of attaching curves on $\Sigma \# \bar{\Sigma}$:

$$\begin{array}{ccc}
 \alpha \bar{\beta} & \xrightarrow{\Theta_{\alpha' \bar{\beta}, \alpha \bar{\beta}}} & \alpha' \bar{\beta} \\
 | & & | \\
 \Theta_{\alpha'' \bar{\beta}, \alpha \bar{\beta}} & \downarrow & \Theta_{\alpha'' \bar{\beta}, \alpha' \bar{\beta}} \\
 \alpha'' \bar{\beta} & \xrightarrow{1} & \alpha'' \bar{\beta}
 \end{array} \tag{10.8}$$

Our hyperbox for Proposition 10.6 will be of size $(1, 1, 6)$, which we organize into 6 hypercubes $\mathcal{C}_1, \dots, \mathcal{C}_6$, which we stack to obtain a hyperbox. The cubes \mathcal{C}_1 and \mathcal{C}_2 are shown in Figure 10.1.

The second hypercube in Figure 10.1 also has a length 3 map which is not drawn. In the figure, the top hypercube is a hypercube of stabilization, and the bottom hypercube is obtained by pairing the hypercube in (10.8) with the 1-dimensional hypercube

$$\beta \bar{\beta} \xrightarrow{\Theta_{\beta \bar{\beta}, \Delta}} \Delta.$$

We now move to the construction of the hypercube \mathcal{C}_3 and \mathcal{C}_4 , starting by constructing some auxiliary hypercubes. In Figure 10.2, we construct the 3-dimensional analog of the hypercube of beta-attaching curves in (10.3). We next consider the hypercube of chain complexes obtained by pairing the cube in (10.8) with the 1-dimensional cube

$$\Delta \xrightarrow{\Theta_{\Delta, \alpha'' \bar{\alpha}''}} \alpha'' \bar{\alpha}''.$$

We reorganize the resulting cube slightly, to obtain the diagram shown in Figure 10.3. Although the diagram in the figure is not technically a hypercube, since the bottom face is subdivided, if we compress the bottom face, we obtain a hypercube of chain complexes, which we call \mathcal{C}_3 .

Next, we pair the top face of Figure 10.2 with the cube

$$\alpha \bar{\beta} \xrightarrow{\Theta_{\alpha \bar{\beta}, \alpha' \bar{\beta}}} \alpha' \bar{\beta}$$

and modify the resulting hypercube to obtain the back-most portion of the diagram shown on the top of Figure 10.4. This diagram is not technically a hypercube, since some of the faces are subdivided. Nonetheless, we can form a hypercube \mathcal{C}_4 as follows. We view the diagram as being obtained by stacking two box-like diagrams (one on the back, and one on the front). Some of the faces of these boxes are subdivided, so we compress these faces. This gives two cube-diagrams. The hypercube relations for these cubes are straightforward to verify. For the back-most cube, they follow from the hypercube relations for the original cube described above. For the front cube, the

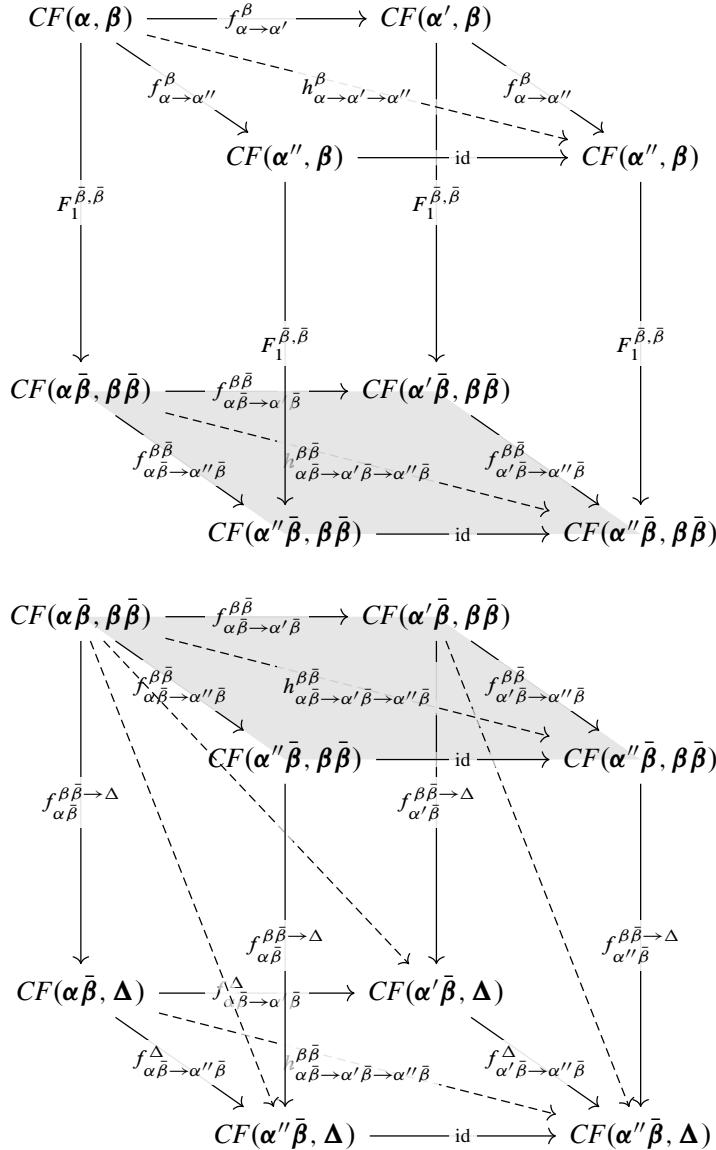


Figure 10.1. The hypercubes \mathcal{C}_1 and \mathcal{C}_2 . The gray faces are stacked.

hypercube relations are straightforward to verify directly, since all of the triangle maps involve stabilizations of Heegaard triples for small isotopies of the attaching curves. Hence, Proposition 6.5 may be used to destabilize the relevant counts, and then the small translate theorem for triangles [4, Proposition 11.1] may be used to identify the destabilized maps with nearest point maps.

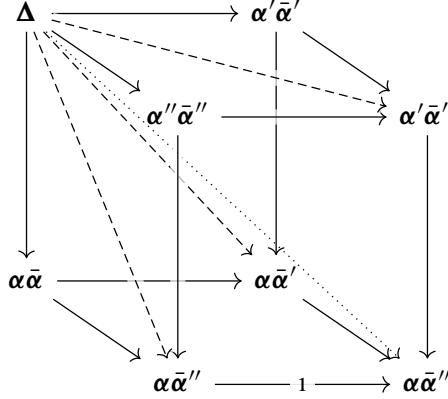


Figure 10.2. An auxiliary hypercube of β -attaching curves used in the construction of \mathcal{C}_4 .

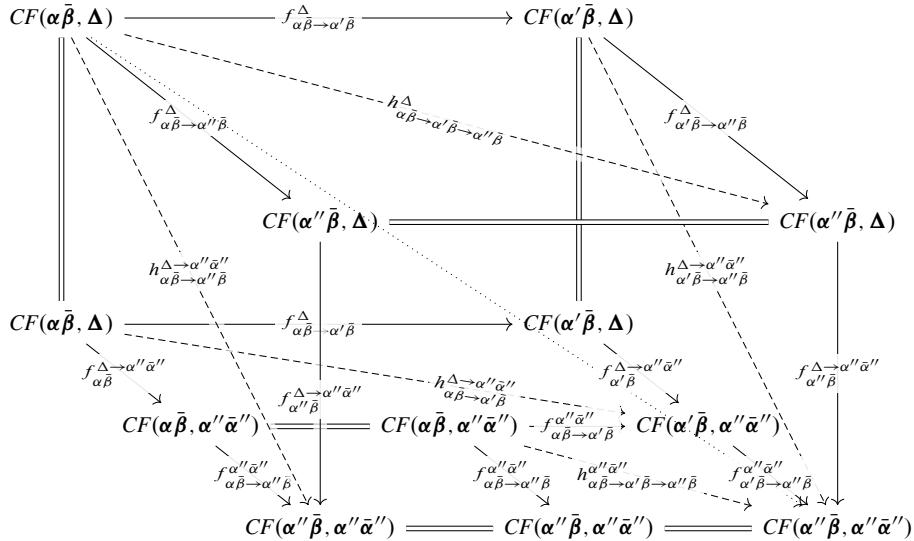


Figure 10.3. The diagram \mathcal{C}_3 .

We now describe the construction of the hypercube \mathcal{C}_5 . We first pair the 0-dimensional hypercube $\alpha\bar{\beta}$ with the hypercube in Figure 10.2. Then, we rearrange, and obtain the diagram shown in the middle of Figure 10.4.

There is a final hyperbox involving 3-handles, which we compress to obtain the hypercube \mathcal{C}_6 , shown on the bottom of Figure 10.4. Stacking and compressing the hypercubes $\mathcal{C}_1, \dots, \mathcal{C}_6$ proves the statement. ■

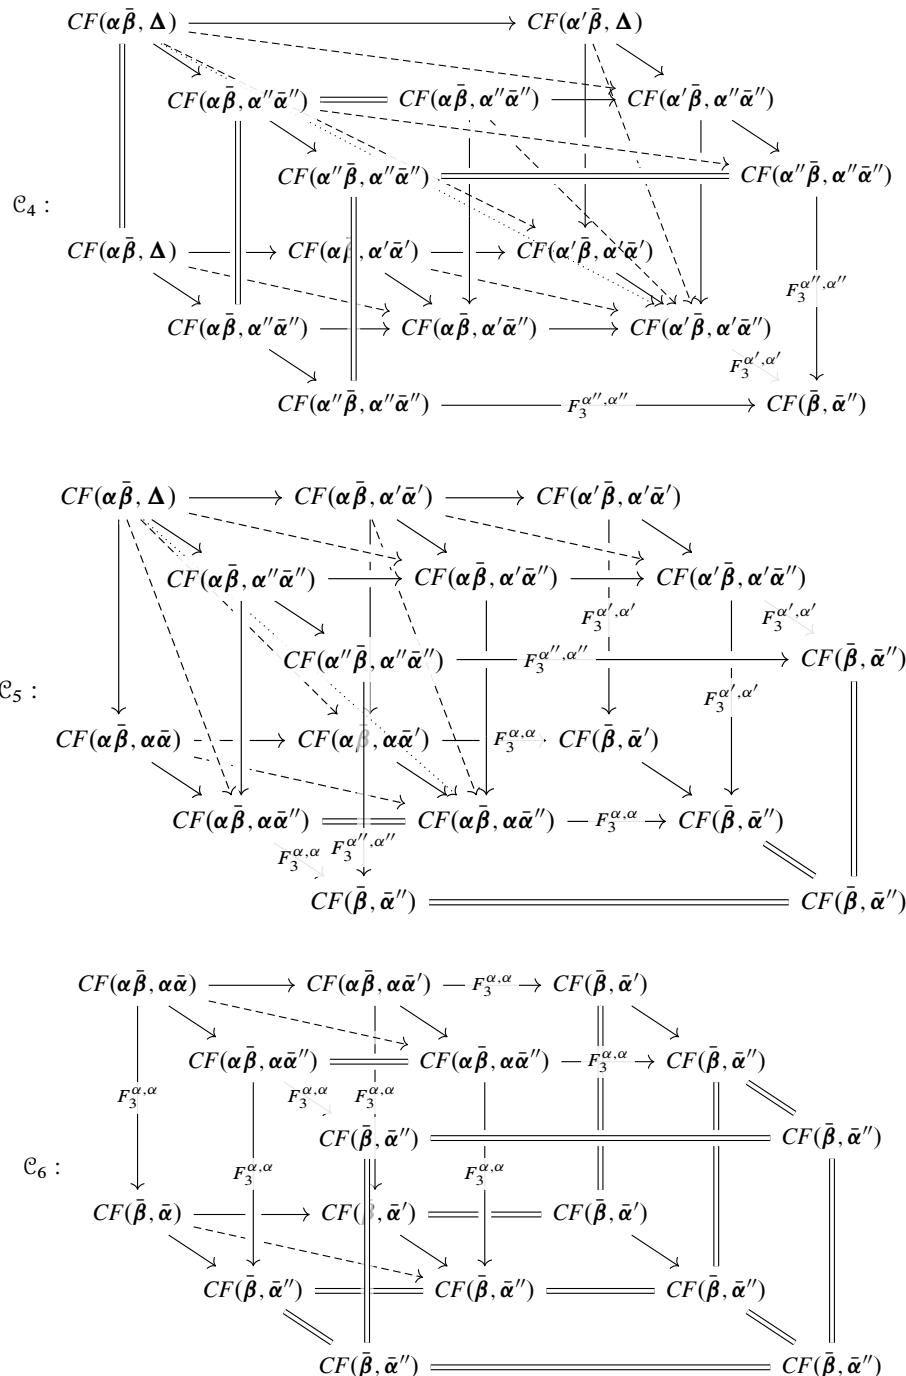


Figure 10.4. The diagrams used to construct \mathcal{C}_4 , \mathcal{C}_5 , and \mathcal{C}_6 .

10.4. More on commutations of handles

In this section, we prove that 1-handle and 3-handle maps may be commuted with 2-handle maps, when the corresponding topological handle manipulation is also possible.

Proposition 10.7. *Suppose that \mathbb{L} is a framed 1-dimensional link in Y , and \mathbb{S} is a 0- or 2-dimensional sphere in $Y \setminus \mathbb{L}$. Let $\mathfrak{S} \subseteq \text{Spin}^c(W(Y, \mathbb{L}))$ be a set of Spin^c structures which is closed under conjugation, and let*

$$\mathfrak{S}' \subseteq \text{Spin}^c(W(Y(\mathbb{S}), \mathbb{L}))$$

be the corresponding set. Then

$$CFI(W(Y(\mathbb{L}), \mathbb{S})) \circ CFI(W(Y, \mathbb{L}), \mathfrak{S}) \simeq CFI(W(Y(\mathbb{S}), \mathbb{L}), \mathfrak{S}') \circ CFI(W(Y, \mathbb{S})).$$

Proof. The proof is similar to the proof of Proposition 10.6. We will build a 3-dimensional hypercube, four of whose faces correspond to the four maps

$$\begin{aligned} CFI(W(Y(\mathbb{L}), \mathbb{S})), & \quad CFI(W(Y, \mathbb{L}), \mathfrak{S}), \\ CFI(W(Y(\mathbb{S}), \mathbb{L}), \mathfrak{S}'), & \quad CFI(W(Y, \mathbb{S})). \end{aligned}$$

The hypercube relations for the compression of this hyperbox translate exactly to the chain homotopy for the commutation described in the statement.

The hypercube realizing the commutation will be constructed as the compression of a hyperbox of size $(1, 1, 5)$, which we view as five hypercubes, stacked on top of each other. We write $\mathcal{B}_1, \dots, \mathcal{B}_5$ for these hypercubes.

The top-most hypercubes \mathcal{B}_1 and \mathcal{B}_2 , are shown in Figure 10.5. They are constructed by stacking hypercubes constructed by pairing hypercubes of attaching curves, as well as hypercubes of stabilization.

The diagram used to construct \mathcal{B}_3 is shown in Figure 10.6.

The hypercube \mathcal{B}_4 is displayed in Figure 10.7. The construction is as follows. The back-left cube is a hypercube of stabilization for the map $F_1^{\gamma_0 \bar{\gamma}_0, \gamma_0 \bar{\gamma}_0}$. The back-right cube is constructed by pairing the 0-dimensional hypercube $\alpha \bar{\beta} \gamma_0 \bar{\gamma}_0$ with a 3-dimensional generalized hypercube of beta attaching curves which has τ as one of the length 1 morphisms. The construction is similar to the definition of the 1-handle map, and the proof of well definedness is similar to Lemma 9.5. The front portion is verified to satisfy the hypercube relations (after compressing the top face) by observing that the length 2 map composes trivially with the 3-handle map $F_3^{\alpha' \gamma_0, \alpha' \gamma_0}$ by applying our stabilization result from Proposition 6.5 and applying the small translate theorems for triangles and quadrilaterals [4, Propositions 11.1 and 11.5].

Our final hypercube \mathcal{B}_5 is shown in Figure 10.8. Stacking and compressing these hypercubes completes the proof. ■

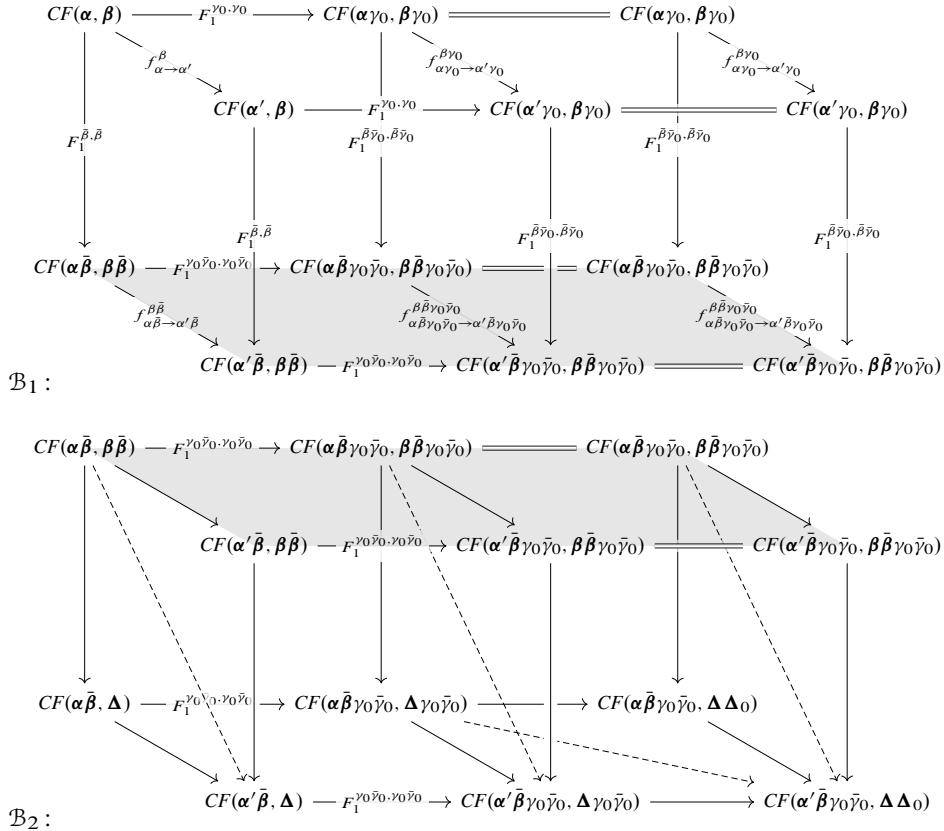


Figure 10.5. The hypercubes \mathcal{B}_1 (top) and \mathcal{B}_2 (bottom) are obtained by compressing these two diagrams. The unlabeled arrows are all holomorphic triangle or quadrilateral maps. There is also a length 3 arrow in the right subcube for \mathcal{B}_2 which is not drawn.

11. Naturality maps for stabilizations

In this section, we construct our naturality maps for stabilizations of the Heegaard surface, and prove some basic properties. Our definition of the naturality map for a stabilization is simple: we define the map for a stabilization as the composition of a 1-handle cobordism map with the 2-handle cobordism map for a 2-handle which topologically cancels the 1-handle.

The definition gives some diagrammatic flexibility for computing the stabilization map. In particular, for a stabilization we require two points p_0 and p_1 to be chosen on $\Sigma \setminus (\alpha \cup \beta)$. We then add a tube to Σ with feet at p_0 and p_1 . We add one new alpha curve and one beta curve. The beta curve is a meridian of the tube, and the

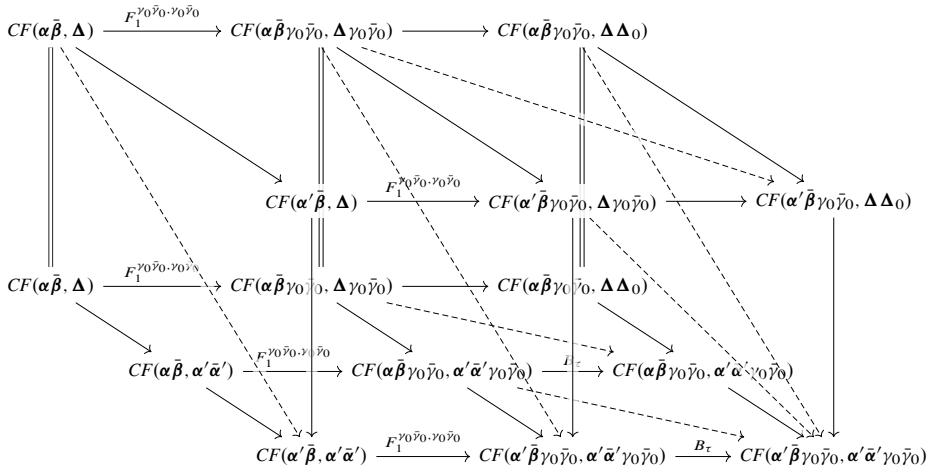


Figure 10.6. The diagram used to construct the hypercube B_3 .

alpha curve is obtained by concatenating a longitude of the curve with an embedded curve on $\Sigma \setminus \alpha$. If Σ is embedded in Y , then we may view the construction as being determined by picking an embedded path on $\Sigma \setminus \alpha$ which connects p_0 and p_1 , and having the tube run parallel to this path. The new alpha circle runs parallel to this arc. We define the destabilization map as the composition of a 2-handle map for a 0-framed unknot, followed by the 3-handle cobordism map.

Note that this stabilization operation naturally determines a 3-ball in Y which intersects Σ in a disk. We say two stabilizations are *disjoint* if there are two choices of corresponding 3-balls in Y which are disjoint.

We write σ for the stabilization map, and τ for the destabilization map. The main results of this section are summarized in the following proposition.

Proposition 11.1. (1) *The stabilization and destabilization maps commute with the transition map for elementary equivalences of doubling-enhanced Heegaard diagram.*

(2) If σ and σ' are the maps for two disjoint stabilizations, then

$$\sigma \circ \sigma' \cong \sigma' \circ \sigma.$$

(3) If σ and τ are stabilizations and destabilizations for the same 3-ball, then

$$\sigma \circ \tau \cong \text{id} \quad \text{and} \quad \tau \circ \sigma \cong \text{id}.$$

Our proof of parts (2) and (3) goes by way of computing basepoint-adjacent stabilizations and destabilizations (note that by handlesliding the original alpha and beta curves, such a configuration may always be achieved).

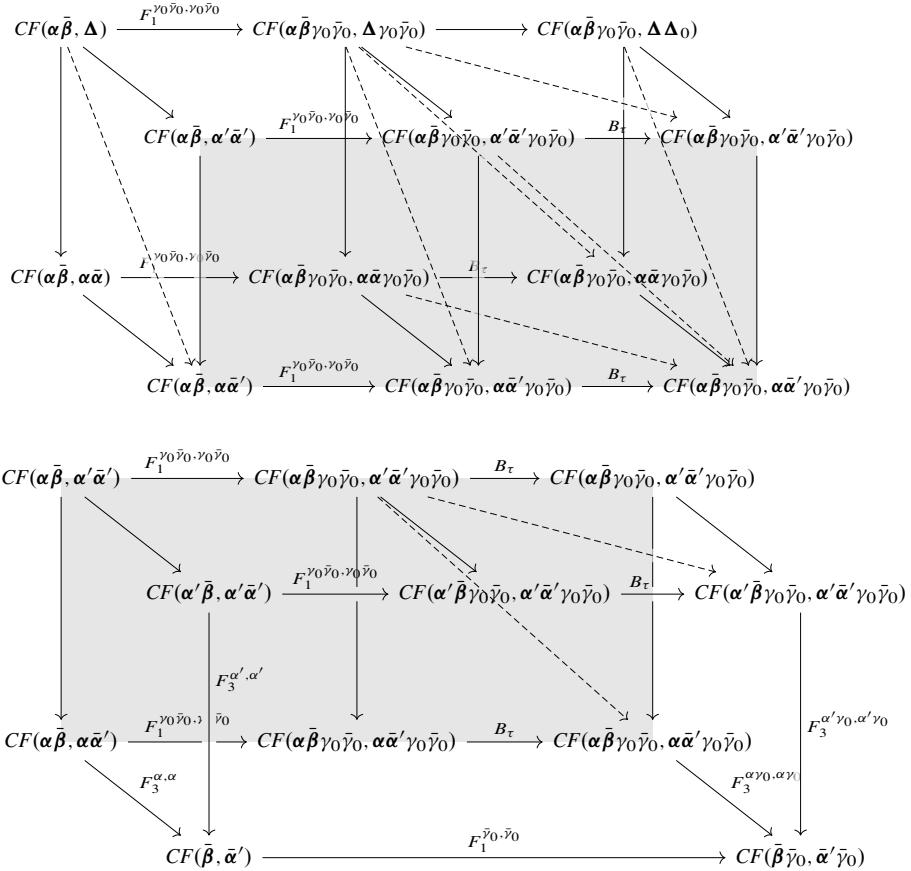


Figure 10.7. The diagram used to build the hypercube \mathcal{B}_4 . We stack the diagrams along the gray face. There is a length 3 map in the right portion of top-most diagram which is not drawn.

Lemma 11.2. Suppose that p_1 and p_0 are two points on $\Sigma \setminus (\alpha \cup \beta)$ which both lie in the same component as the basepoint w . For a suitable choice of doubling curves on the stabilized diagram and suitably degenerated choices of almost complex structures, we have

$$\sigma \simeq (x \mapsto x \times c) \quad \text{and} \quad \tau \simeq (x \times c \mapsto x),$$

where $\{c\} = \alpha_0 \cap \beta_0$, and $(x \mapsto x \times c)$ denotes the $\mathbb{F}[U, Q]/Q^2$ -equivariant map which sends x to $x \times c$, and similarly for $(x \times c \mapsto x)$.

Proof. Consider σ (the proof for τ is essentially the same). By definition, σ is the composition of the 1-handle cobordism map, followed by the map for a canceling 2-handle. By Lemma 9.6, for a suitable choice of almost complex structure, the 1-handle

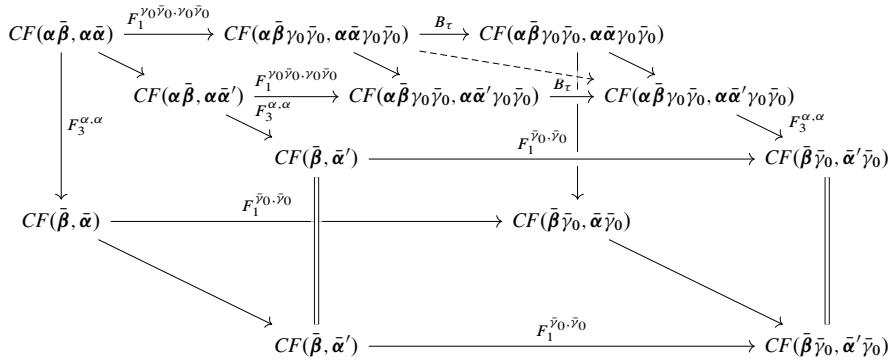


Figure 10.8. The diagram used to construct the hypercube \mathcal{B}_5 .

map takes the form $x \mapsto x \times \theta^+$, extended $\mathbb{F}[U, Q]/Q^2$ -equivariantly. We will show that the 2-handle cobordism map is chain homotopic to a map which sends $x \times \theta^+$ to $x \times c$.

We recall that the 2-handle map is defined in equation (10.4). We claim that the diagonal map of the compression is trivial. To see this, we make the following sub-claims about the diagram in (10.4) in the present context.

- (1) The top-most length 2 map vanishes on elements of the form $x \times \theta_{\beta_0, \beta_0}^+ \times \Theta_{\bar{\beta}, \bar{\beta}}^+ \times \theta_{\bar{\beta}_0, \bar{\beta}_0}^+$, for $x \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$.
- (2) The middle length 2 map composes trivially with the 3-handle map $F_3^{\alpha\beta_0, \alpha\beta_0}$ (labeled $F_3^{\alpha', \alpha'}$ in equation (10.4)), and the bottom-most length 2 map composes trivially with the map $F_3^{\alpha\alpha_0, \alpha\alpha_0}$ (labeled $F_3^{\alpha, \alpha}$ in equation (10.4)).

We consider claim (1) first. There is a genus 2 portion of the Heegaard quadruple corresponding to the stabilization, which takes the following form:

$$\mathcal{Q}_0 = (\mathbb{T}^2 \# \mathbb{T}^2, \alpha_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0, \Delta_0).$$

We observe that

$$\hat{f}_{\alpha_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0, \Delta_0} \left(\hat{f}_{\alpha_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0} \left(\Theta_{\alpha_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0}^+, \Theta_{\beta_0 \bar{\beta}_0, \beta_0 \bar{\beta}_0}^+ \right), \Theta_{\beta_0 \bar{\beta}_0, \Delta_0}^+ \right) = \Theta_{\alpha_0 \bar{\beta}_0, \Delta_0}^+.$$

This may be seen since the inner triangle map corresponds to the naturality map for a small perturbation of the curves $\beta_0 \bar{\beta}_0$, while the outer triangle map corresponds to a factor of the doubling model of the involution for the diagram $(\mathbb{T}^2, \alpha_0, \beta_0)$. Hence, by applying Proposition 6.5, we obtain the equality

$$h_{\alpha\beta_0 \bar{\beta}_0 \rightarrow \alpha\alpha_0 \bar{\beta} \bar{\beta}_0}^{\beta\beta_0 \bar{\beta} \bar{\beta}_0 \rightarrow \Delta\Delta_0} \left(x \times \theta_{\beta_0, \beta_0}^+ \times \Theta_{\bar{\beta}, \bar{\beta}}^+ \times \theta_{\bar{\beta}_0, \bar{\beta}_0}^+ \right) = h_{\alpha\bar{\beta} \rightarrow \alpha\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta} \left(x \times \theta_{\beta_0, \beta_0}^+ \right) \otimes \Theta_{\alpha_0 \bar{\beta}_0, \Delta_0}^+,$$

which vanishes by the small translate theorem for holomorphic quadrilaterals [4, Proposition 11.5].

Next, we consider claim (2), from above. There are two quadrilateral maps which appear in this claim. The proof for both maps is essentially the same, so we focus on the map

$$h_{\alpha\beta_0\bar{\beta}\bar{\beta}_0 \rightarrow \alpha\alpha_0\bar{\beta}\bar{\beta}_0}^{\Delta\Delta_0 \rightarrow \alpha\alpha_0\bar{\alpha}\bar{\alpha}_0}.$$

This is the middle length 2 map in (10.4). We wish to show that this map composes trivially with $F_3^{\alpha',\alpha'}$. To see this, we consider the genus 2 Heegaard quadruple

$$Q_1 = (\mathbb{T}^2 \# \mathbb{T}^2, \alpha_0\bar{\alpha}_0, \Delta_0, \beta_0\bar{\beta}_0, \alpha_0\bar{\beta}_0).$$

This is an admissible, algebraically rigid multi-stabilizing quadruple. Furthermore, we claim that

$$\hat{f}_{\alpha_0\bar{\alpha}_0, \beta_0\bar{\beta}_0, \alpha_0\bar{\beta}_0}(\hat{f}_{\alpha_0\bar{\alpha}_0, \Delta_0, \beta_0\bar{\beta}_0}(\Theta_{\alpha_0\bar{\alpha}_0, \beta_0\bar{\beta}_0}^+, \Theta_{\Delta_0, \beta_0\bar{\beta}_0}^+, \Theta_{\beta_0\bar{\beta}_0, \alpha_0\bar{\beta}_0}^+)) = \Theta_{\alpha_0\bar{\alpha}_0, \alpha_0\bar{\beta}_0}^-. \quad (11.1)$$

The above equation is verified by noting that the inner triangle map may be interpreted as the cobordism map for a 2-handle attachment along a fiber of $S^1 \times S^2$, while the outer triangle map may be interpreted as a 2-handle attached along a 0-framed unknot. The fact that $\Theta_{\alpha_0\bar{\alpha}_0, \alpha_0\bar{\beta}_0}^-$ is the output may be viewed as a consequence of the above topological description.

Combining (11.1) with Proposition 6.5, we conclude that

$$\begin{aligned} h_{\alpha\beta_0\bar{\beta}\bar{\beta}_0 \rightarrow \alpha\alpha_0\bar{\beta}\bar{\beta}_0}^{\Delta\Delta_0 \rightarrow \alpha\alpha_0\bar{\alpha}\bar{\alpha}_0}(x \times \Theta_{\beta_0\bar{\beta}_0, \Delta}^+) \\ = h_{\alpha\bar{\beta} \rightarrow \alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}}(x) \times \Theta_{\alpha_0\bar{\beta}_0, \alpha_0\bar{\alpha}_0}^- + \sum_{z \in \mathbb{T}_{\alpha\bar{\beta}} \cap \mathbb{T}_{\alpha\bar{\alpha}}} C_z \cdot z \times \Theta_{\alpha_0\bar{\beta}_0, \alpha_0\bar{\alpha}_0}^+, \end{aligned} \quad (11.2)$$

for some $C_z \in \mathbb{F}[U]$. Equation (11.2) vanishes once composed with $F_3^{\alpha\alpha_0, \alpha\alpha_0}$, by the small translate theorem for holomorphic quadrilaterals [4, Proposition 11.5].

The composition of the bottom-most length 2 map from equation (10.4) with $F_3^{\alpha\beta_0, \alpha\beta_0}$ vanishes in our present setting by a nearly identical argument. ■

We can now prove Proposition 11.1.

Proof of Proposition 11.1. Claim (1) follows immediately from the fact that 1-handle, 2-handle and 3-handle cobordism maps commute (up to chain homotopy) with the transition maps for elementary handleslides.

Claims (2) and (3) follow from Lemma 11.2, since by (1), we may perform stabilizations near the basepoint, and then observe that the formulas given therein clearly satisfy the stated relations. ■

12. Handleswaps

In this section, we prove handleswap invariance.

12.1. Simple handleswaps

Of central importance to the approach to naturality in [9] are the following loops in the graph of Heegaard diagrams for a 3-manifold Y .

Definition 12.1. Suppose Y is a 3-manifold. A *simple handleswap loop* is a triangle of Heegaard diagrams

$$\begin{array}{ccc} & \mathcal{H}_1 & \\ g \nearrow & & \searrow e \\ \mathcal{H}_3 & \xleftarrow{f} & \mathcal{H}_2 \end{array}$$

which satisfies the following.

- (1) \mathcal{H}_i is an embedded Heegaard diagram for Y , which has the same Heegaard surface for each $i \in \{1, 2, 3\}$.
- (2) e is an α -equivalence, f is a β -equivalence, and g is a diffeomorphism.
- (3) The diagrams \mathcal{H}_i decompose as connected sums of a fixed Heegaard diagram $\mathcal{H} = (\Sigma, \alpha_u, \beta_u)$ with genus two diagrams $\mathcal{G}_i = (\Sigma_0, \alpha_i, \beta_i)$. Furthermore, in the punctured genus two surface $(\Sigma \# \Sigma_0) \setminus \Sigma$, the above triangle is diffeomorphic to the triangle in Figure 12.1.

In particular, in the above definition, α_1 consists of two closed curves α_1 and α_2 , while β_1 consists of two closed curves β_1 and β_2 . The arrow e corresponds to handlesliding α_1 over α_2 , while the arrow f corresponds to handlesliding β_1 over β_2 .

Recall that associated to handleslides and diffeomorphisms, we have defined morphisms, $CFI(e)$, $CFI(f)$ and $CFI(g)$, which are well defined up to $\mathbb{F}[U, Q]/Q^2$ -equivariant homotopy equivalence.

Theorem 12.2. $CFI(g) \circ CFI(f) \circ CFI(e) \simeq \text{Id}_{CFI(\mathcal{H}_1)}$.

The proof of Theorem 12.2 occupies the subsequent sections.

12.2. Reduction to basepoint-adjacent handleswaps

A simple handleswap need not occur in a region of the Heegaard diagram near a basepoint. However, we now show that for the purposes of naturality, it suffices to consider handleswaps where the special genus 2 region is adjacent to the basepoint. We note that this fact was also observed by Sungkyung Kang.

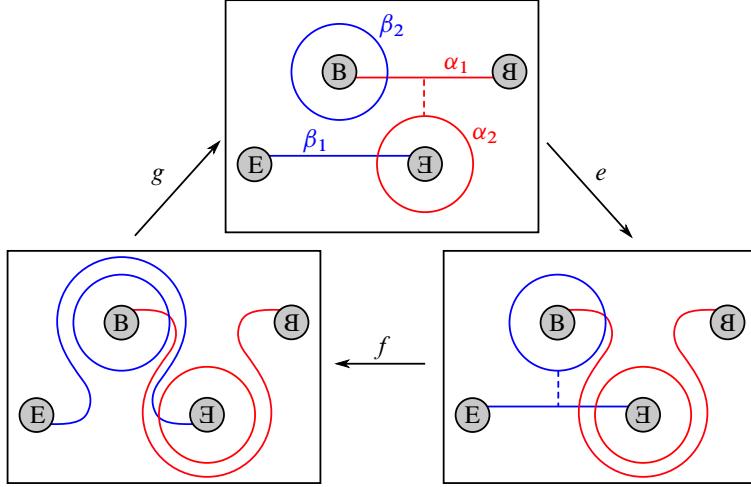
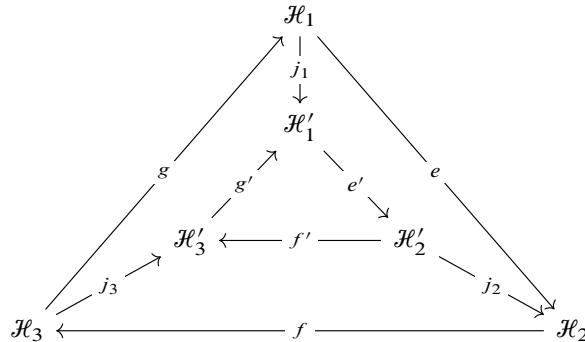


Figure 12.1. A simple handleswap.

Lemma 12.3. Suppose that F is a weak Heegaard invariant (in the sense of [9, Definition 2.24]) which satisfies all of the conditions of a strong Heegaard invariant [9, Definition 3.23] except possibly handleswap invariance. If F has no monodromy around basepoint-adjacent handleswaps, then it has no monodromy around arbitrary handleswaps.

Proof. Let $p \in \Sigma$ denote the connected sum point. Pick an arc λ which connects p to the basepoint w . We proceed by induction on the number of intersections of $(\alpha_u \cup \beta_u) \cap \lambda$. The base case is the hypothesis. If \mathcal{H} is arbitrary, there is an equivalence j between \mathcal{H} and a diagram \mathcal{H}' (obtained by handlesliding an alpha or beta curve over p) such that on \mathcal{H}' , there is one fewer intersection of $(\alpha_u \cup \beta_u) \cap \lambda$. If \mathcal{H}'_1 , \mathcal{H}'_2 and \mathcal{H}'_3 denote the diagrams obtained by connect summing the standard genus 2 diagrams from Figure 12.1, there is an analogous simple handleswap loop involving \mathcal{H}'_1 , \mathcal{H}'_2 and \mathcal{H}'_3 . Write e' , f' and g' for the analogous equivalences. We have a diagram of equivalences



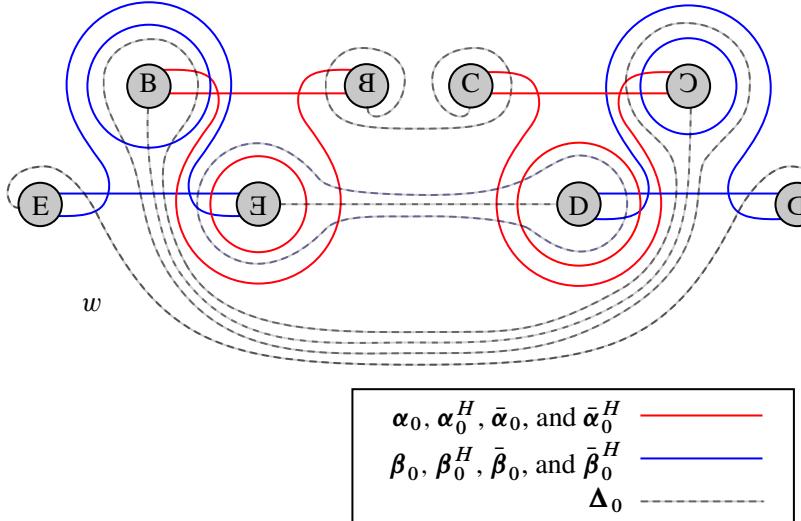


Figure 12.2. The double of the handleswap region. The left side is Σ , and the right is $\bar{\Sigma}$.

The three outer rectangles are obtained from distinguished squares by inverting some of the arrows, which are isomorphisms. By hypothesis, F commutes around these squares. Our inductive hypothesis is that F has no monodromy around the central triangle. This clearly implies that F has no monodromy around the outer triangle, completing the proof. ■

12.3. Proof of Theorem 12.2

Let us introduce some notation. For the curves in the handleswap region, we write α_0 , α_0^H , β_0 and β_0^H . We write α^H for $\alpha_u \cup \alpha_0^H$, and similarly for β_0^H . Hence, for the three diagrams appearing in the handleswap loop, we write

$$\mathcal{H}_1 = (\Sigma \# \Sigma_0, \alpha, \beta), \quad \mathcal{H}_2 = (\Sigma \# \Sigma_0, \alpha^H, \beta), \quad \mathcal{H}_3 = (\Sigma \# \Sigma_0, \alpha^H, \beta^H).$$

Similarly, let Δ_u be a collection of doubling curves on $\Sigma \# \bar{\Sigma}$. Let Δ_0 be choice of doubling curves on $\Sigma_0 \# \bar{\Sigma}_0$, as shown in Figure 12.2. Write Δ for $\Delta_u \cup \Delta_0$.

We write \mathcal{D} for \mathcal{H} , decorated with Δ_u . We pick a collection of four doubling curves Δ_0 on $\Sigma_0 \# \bar{\Sigma}_0$. We decorate \mathcal{H}_1 , \mathcal{H}_2 , and \mathcal{H}_3 with Δ to form the doubling enhanced diagrams \mathcal{D}_1 , \mathcal{D}_2 , and \mathcal{D}_3 , respectively.

We note that for $i \in \{1, 2, 3\}$, there are canonical isomorphisms

$$CFI(\mathcal{D}) \cong CFI(\mathcal{D}_i), \quad (12.1)$$

given by $x \mapsto x \times c_i$, extended $\mathbb{F}[U, Q]/Q^2$ -equivariantly. Here, we write c_i for the unique intersection points

$$c_1 \in \mathbb{T}_{\alpha_0} \cap \mathbb{T}_{\beta_0}, \quad c_2 \in \mathbb{T}_{\alpha_0^H} \cap \mathbb{T}_{\beta_0}, \quad \text{and} \quad c_3 \in \mathbb{T}_{\alpha_0^H} \cap \mathbb{T}_{\beta_0^H}.$$

In fact, we have the following lemma.

Lemma 12.4. *The genus 2 diagram shown in Figure 12.2 is weakly admissible. Furthermore, each of the following subdiagrams has vanishing differential:*

- (1) $(\Sigma_0 \# \bar{\Sigma}_0, \beta_0 \bar{\beta}_0, \Delta_0, w)$,
- (2) $(\Sigma_0 \# \bar{\Sigma}_0, \beta_0^H \bar{\beta}_0^H, \Delta_0, w)$,
- (3) $(\Sigma_0 \# \bar{\Sigma}_0, \Delta_0, \alpha_0 \bar{\alpha}_0, w)$,
- (4) $(\Sigma_0 \# \bar{\Sigma}_0, \Delta_0, \alpha_0^H \bar{\alpha}_0^H, w)$,
- (5) $(\Sigma_0 \# \bar{\Sigma}_0, \alpha_0 \bar{\beta}_0, \Delta_0, w)$,
- (6) $(\Sigma_0 \# \bar{\Sigma}_0, \alpha_0^H \bar{\beta}_0^H, \Delta_0, w)$.

Proof. Admissibility is a straightforward computation which we leave to the reader. To check that each of the stated complexes has vanishing differential, it suffices to check that each has the same number of generators as its homology. The first four subdiagrams represent a complex for $\#^2 S^1 \times S^2$ and also have exactly four generators. Similarly the final two subdiagrams represent S^3 , and each have a unique intersection point. ■

By definition,

$$CFI(e) = \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_2} \quad \text{and} \quad CFI(f) = \Psi_{\mathfrak{D}_2 \rightarrow \mathfrak{D}_3},$$

where $\Psi_{\mathfrak{D}_i \rightarrow \mathfrak{D}_{i+1}}$ is obtained by compressing a hyperbox as in Figure 5.1. Computing the composition $\Psi_{\mathfrak{D}_2 \rightarrow \mathfrak{D}_3} \circ \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_2}$ is somewhat inefficient, however, since the involutive transition maps naturally allow for changing both the alpha and the beta curves simultaneously. Instead, we will consider the map $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$ obtained by compressing the hyperbox in Figure 5.1. (See equation (12.4), below). By Theorem 5.2,

$$CFI(f) \circ CFI(e) = \Psi_{\mathfrak{D}_2 \rightarrow \mathfrak{D}_3} \circ \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_2} \simeq \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}. \quad (12.2)$$

To prove Theorem 12.2, we also need to understand the diffeomorphism map $CFI(g)$. We note that the diffeomorphism g does not fix Δ_0 . Hence, by definition, this map is the composition

$$CFI(g) := \Psi_{g(\mathfrak{D}_3) \rightarrow \mathfrak{D}_1} \circ T_g, \quad (12.3)$$

which we explain presently. The map T_g is the tautological map which sends an intersection point $x \in \mathbb{T}_{\alpha\alpha_0^H} \cap \mathbb{T}_{\beta\beta_0^H}$ to its image in $\mathbb{T}_{\alpha\alpha_0} \cap \mathbb{T}_{\beta\beta_0}$ under g , extended

$\mathbb{F}[U, Q]/Q^2$ -equivariantly. Also, $g(\mathfrak{D}_3)$ is obtained by enhancing Heegaard diagram \mathcal{H}_1 with the doubling curves $\Delta_u \cup \Delta'_0$, where $\Delta'_0 = g(\Delta_0)$. The map $\Psi_{g(\mathfrak{D}_3) \rightarrow \mathfrak{D}_1}$ is built by compressing a hyperbox similar to the one in Figure 5.1 which realizes the change of doubling curves.

With respect to the canonical isomorphisms of (12.1), the map T_g is clearly equal to the identity map. Theorem 12.2 follows from (12.2) and (12.3) and the following result.

Proposition 12.5. *With respect to the canonical isomorphisms in (12.1), the maps $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$ and $\Psi_{g(\mathfrak{D}_3) \rightarrow \mathfrak{D}_1}$ are chain homotopic to the identity map.*

Proof. We focus first on the computation of $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$. By definition, this map is obtained by compressing the following hyperbox:

$$\begin{array}{ccccc}
 CF(\alpha, \beta) & \longrightarrow & CF(\alpha^H, \beta) & \longrightarrow & CF(\alpha^H, \beta^H) \\
 \downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}^H, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}^H, \bar{\beta}^H} \\
 CF(\alpha\bar{\beta}, \beta\bar{\beta}) & \longrightarrow & CF(\alpha^H\bar{\beta}^H, \beta\bar{\beta}) & \longrightarrow & CF(\alpha^H\bar{\beta}^H, \beta^H\bar{\beta}^H) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 CF(\alpha\bar{\beta}, \Delta) & \longrightarrow & CF(\alpha^H\bar{\beta}^H, \Delta) & \longrightarrow & CF(\alpha^H\bar{\beta}^H, \Delta) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & \longrightarrow & CF(\alpha^H\bar{\beta}^H, \alpha\bar{\alpha}) & \longrightarrow & CF(\alpha^H\bar{\beta}^H, \alpha^H\bar{\alpha}^H) \\
 \downarrow F_3^{\alpha, \alpha} & & \downarrow F_3^{\alpha^H, \alpha} & & \downarrow F_3^{\alpha^H, \alpha^H} \\
 CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}^H, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}^H, \bar{\alpha}^H) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}^H, \bar{\alpha}^H) & \longrightarrow & CF(\bar{\beta}^H, \bar{\alpha}^H)
 \end{array} \tag{12.4}$$

We note that in the above diagram, each polygon counting map occurs on a connected sum. One summand is Σ or $\Sigma \# \bar{\Sigma}$ while the other summand is the genus 2 subsurface Σ_0 , or its double. We will argue by viewing the diagram as a multi-stabilization and applying Lemma 12.4 and Proposition 6.5.

There is a subtlety in that the almost complex structures used to construct the transition map equation (12.4) are not naturally adapted to the genus four stabilization region where the handleswap occurs. Instead, we must first argue that we can change

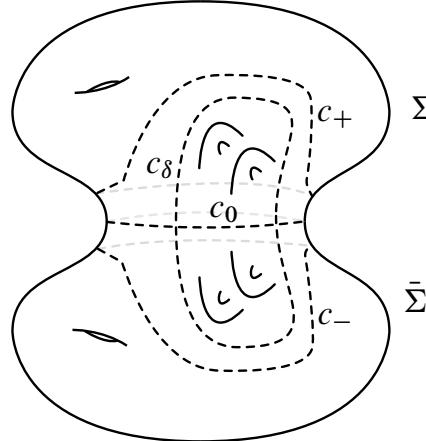


Figure 12.3. Curves along which we degenerate the almost complex structure.

the above composition to one which is naturally adapted to the handleswap. To do so, we consider the closed curves on $\Sigma \# \bar{\Sigma}$ shown in Figure 12.3. We label these as c_+ , c_- , c_δ and c_0 . These are as follows.

- (1) c_0 is a meridian of the connected sum tube used in the doubling operation.
- (2) c_+ is a closed curve on Σ which bounds the handleswap region and the connected sum point.
- (3) c_- is a symmetric closed curve on $\bar{\Sigma}$, and bounds the handleswap region as well as a the connected sum point.
- (4) c_δ is a closed curve on $\Sigma \# \bar{\Sigma}$ which bounds the union of both handleswap regions. Also, c_δ is disjoint from all Δ curves.

Write J_{c_0} for an almost complex structure or family of almost complex structures which is nodal along c_0 , and use similar notation to define almost complex structures degenerated on another curve or collections of pairwise disjoint curves on $\Sigma \# \bar{\Sigma}$.

We introduce some shorthand. Write $CF(L_1), \dots, CF(L_6)$ for the six horizontal levels of equation (12.4). (These are themselves 1-dimensional hyperboxes of chain complexes). By using change of almost complex structure hypercubes, similar to [4, Section 13], we may assume that the complexes and maps in equation (12.4) are computed with respect to almost complex structures which are *maximally pinched* along c_0 , c_+ and c_- , in the sense of [4, Section 8.3]. This gives the map $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$, as defined in the earlier sections of the present paper. On the other hand, using similar techniques, we may use change of almost complex structure hypercubes to relate this model to one which is maximally pinched along c_+ , c_- and c_δ . This gives a hyperbox

which we schematically indicate by the following diagram:

$$\begin{array}{ccc}
 CF_{J_{c_+}}(L_1) & \xrightarrow{\text{id}} & CF_{J_{c_+}}(L_1) \\
 \downarrow F_1 & & \downarrow F_1 \\
 CF_{J_{c_+,c_-,c_0}}(L_2) & \xrightarrow{\Psi_{J_{c_+,c_-,c_0} \rightarrow J_{c_+,c_-,c_\delta}}} & CF_{J_{c_+,c_-,c_\delta}}(L_2) \\
 \downarrow & \searrow \Psi_{J \rightarrow J_{c_\delta}} & \downarrow \\
 CF_J(L_3) & \xrightarrow{\Psi_{J \rightarrow J_{c_\delta}}} & CF_{J_{c_\delta}}(L_3) \\
 \downarrow & \searrow \Psi_{J_{c_+,c_-,c_0} \rightarrow J_{c_+,c_-,c_\delta}} & \downarrow \\
 CF_{J_{c_+,c_-,c_0}}(L_4) & \xrightarrow{\Psi_{J_{c_+,c_-,c_0} \rightarrow J_{c_+,c_-,c_\delta}}} & CF_{J_{c_+,c_-,c_\delta}}(L_4) \\
 \downarrow F_3 & & \downarrow F_3 \\
 CF_{J_{c_-}}(L_5) & \xrightarrow{\text{id}} & CF_{J_{c_-}}(L_5) \\
 \downarrow & \searrow \text{id} & \downarrow \\
 CF_{J_{c_-}}(L_6) & \xrightarrow{\text{id}} & CF_{J_{c_-}}(L_6)
 \end{array} \tag{12.5}$$

The hyperbox in equation (12.5) realizes a chain homotopy between the model of $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$, computed with respect to almost complex structures which are maximally pinched along c_+ , c_- and c_0 , and a model of $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$ which is computed using almost complex structures which are maximally pinched along c_+ , c_- and c_δ .

In the model of $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$ appearing along the right-hand side of equation (12.5), the holomorphic polygon counts all occur on a multi-stabilization of the diagram on Σ or $\Sigma \# \bar{\Sigma}$. The stabilization results for triangles imply that the two length 1 maps along the top of (12.4) compose to give a map which is intertwined with the identity under the canonical isomorphisms $CF(\mathcal{H}) \cong CF(\mathcal{H}_i)$ (similar to (12.1)). (This is the content of [9, Propositions 9.31 and 9.32]).

It remains to show that none of the diagonal maps make non-trivial contribution to the compression of (12.4). To this end, we make the following claims, which together imply that the compression has trivial diagonal map.

- (1) The compression of the top two levels of (12.4) has trivial diagonal map.
- (2) The compression of the third and fourth levels of (12.4) has trivial diagonal map.
- (3) The bottom-most level of (12.4) has trivial diagonal map.

Consider first the claim (3). The diagonal map is a quadrilateral counting map on the connected sum $(\Sigma, \alpha_u, \alpha_u, \beta_u, \beta_u) \# (\Sigma_0, \alpha_0^H, \alpha_0, \beta_0, \beta_0^H)$. Applying the stabilization result from Proposition 6.5 to the setting of holomorphic quadrilaterals, we may

destabilize the map to the diagram $(\Sigma, \alpha_u, \alpha_u, \beta_u, \beta_u)$. By the small translate theorem for quadrilaterals [4, Proposition 11.5], this map vanishes.

Next, consider (1). For this claim, we observe that the stabilization results for holomorphic triangles imply that the length 1 holomorphic triangle maps preserve the subspaces spanned by the top-degree generator (and furthermore, act tensorially on these subspaces). Furthermore, the same argument as for claim (3) implies that the holomorphic quadrilateral maps vanish on elements which have a factor which is the top degree generator. Since the top degree generator is the output of the 1-handle maps on the top row, this implies that the compression has trivial diagonal map. Claim (2) follows from an analogous argument. This completes the proof of the claim about $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$.

We now consider the claim about $\Psi_{g(\mathfrak{D}_3) \rightarrow \mathfrak{D}_1}$. This map may be computed by compressing the following diagram:

$$\begin{array}{ccccc}
 CF(\alpha, \beta) & \longrightarrow & CF(\alpha, \beta) & \longrightarrow & CF(\alpha, \beta) \\
 \downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}, \bar{\beta}} & & \downarrow F_1^{\bar{\beta}, \bar{\beta}} \\
 CF(\alpha\bar{\beta}, \beta\bar{\beta}) & \longrightarrow & CF(\alpha\bar{\beta}, \beta\bar{\beta}) & \longrightarrow & CF(\alpha\bar{\beta}, \beta\bar{\beta}) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 CF(\alpha\bar{\beta}, g(\Delta)) & \longrightarrow & CF(\alpha\bar{\beta}, g(\Delta)) & \longrightarrow & CF(\alpha\bar{\beta}, \Delta) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & \longrightarrow & CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) & \longrightarrow & CF(\alpha\bar{\beta}, \alpha\bar{\alpha}) \\
 \downarrow F_3^{\alpha, \alpha} & & \downarrow F_3^{\alpha, \alpha} & & \downarrow F_3^{\alpha, \alpha} \\
 CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}, \bar{\alpha}) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\beta}, \bar{\alpha})
 \end{array} \tag{12.6}$$

The description in (12.6) is not the most convenient for our purposes. Instead, we describe a sequence of doubling curves $g(\Delta_0) = \Delta_0^0, \dots, \Delta_0^6 = \Delta_0$ on $\Sigma_0 \# \bar{\Sigma}_0$, such that each Δ_0^i is obtained from Δ_0^{i-1} by a handleslide. Furthermore, we have the following properties for each $i \in \{0, \dots, 6\}$.

(1) The diagram

$$(\Sigma_0 \# \bar{\Sigma}_0, \alpha_0 \bar{\alpha}_0, \beta_0 \bar{\beta}_0, \Delta_0^i, \Delta_0^{i+1}, w)$$

is weakly admissible.

(2) Each of the following complexes $CF(\alpha_0 \bar{\beta}_0, \Delta_0^i)$, $CF(\alpha \bar{\alpha}, \Delta_0^i)$, $CF(\beta \bar{\beta}, \Delta_0^i)$, and $CF(\Delta_0^i, \Delta_0^{i+1})$ has vanishing differential.

The curves $\Delta_0^0, \dots, \Delta_0^6$ are shown in Figure 12.4. We leave it to the reader to verify the above claims. (Compare Lemma 12.4.)

We compute the map $\Psi_{g(\mathfrak{D}_3) \rightarrow \mathfrak{D}_1}$ as a composition of 6 maps, corresponding to the transition maps for changing Δ_0^i to Δ_0^{i+1} , each of which is computed by compressing a hyperbox similar to (12.6). The same argument as for $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$ implies that each of these maps is intertwined with the identity map by the canonical isomorphisms in (12.1). The proof is complete. ■

13. Completing the proofs of naturality and functoriality

In this section, we complete our proofs of naturality and functoriality.

13.1. Naturality

If \mathfrak{D} and \mathfrak{D}' are two arbitrary doubling-enhanced Heegaard diagrams (equipped with a framing at the basepoint), we define our transition map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ by composing the maps for an arbitrary sequence stabilizations, elementary equivalences of attaching curves, and pointed isotopies of the Heegaard diagram in (Y, w) .

We now prove Theorem 1.1, showing that this transition map $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}'}$ defined above is independent up to chain homotopy from the sequence of Heegaard moves.

Proof of Theorem 1.1. By [9, Theorem 2.38] and as in Section 4, it suffices to show that the transition maps we have defined satisfy the axioms of a strong Heegaard invariant [9, Definition 2.32]. We verify these axioms presently.

The *functoriality* axiom asserts two claims. Firstly, it asserts that we have well-defined morphisms associated to handleslide equivalences of alpha curves and morphisms associated to handleslide equivalences of beta curves (for a fixed Heegaard surface). This part of the axiom follows from Theorem 5.2. (We remind the reader that unlike in the setting of ordinary Heegaard Floer homology, it is not particularly natural to define an involutive transition map which changes just the alpha curves, or just the beta curves, since the transition map is defined by compressing hyperboxes as in (5.1). Such hyperboxes necessarily require changing both the alpha and beta curves simultaneously.)

The second part of the functoriality axiom asserts that the morphisms for stabilizations are inverse to the morphisms for destabilizations. This is proven in Proposition 11.1.

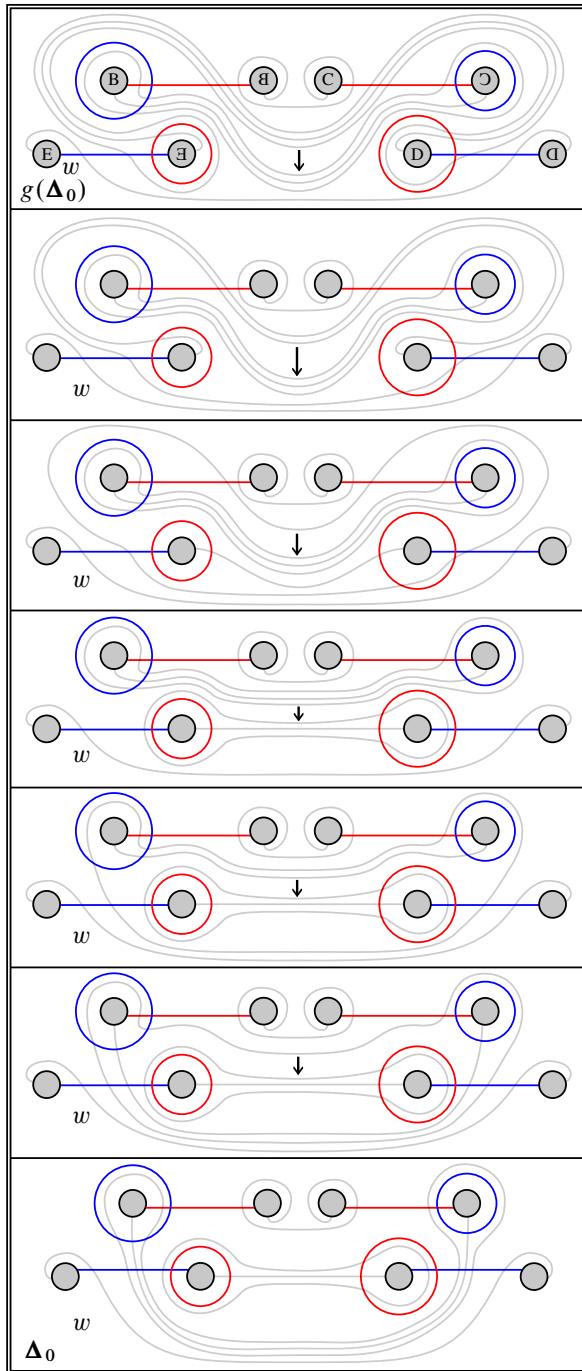


Figure 12.4. The curves $g(\Delta_0) = \Delta_0^0, \dots, \Delta_0^6 = \Delta_0$, which relate $g(\Delta_0)$ and Δ_0 by a sequence of handleslides. Each arrow indicates the subsequent handleslide.

The next axiom from [9, Definition 3.32] is the *commutativity* axiom, concerning commutativity of the distinguished rectangles from [9, Definition 2.29]. In our present case, we verify that there is no monodromy around the distinguished rectangles in Proposition 4.2. In our present setting, we may rephrase the distinguished rectangles from [9] as diagrams of embedded, doubling enhanced Heegaard diagrams with the following shape, satisfying one of five configurations:

$$\begin{array}{ccc} \mathfrak{D}_1 & \xrightarrow{e} & \mathfrak{D}_2 \\ \downarrow f & & \downarrow g \\ \mathfrak{D}_3 & \xrightarrow{h} & \mathfrak{D}_4 \end{array}$$

We presently enumerate the five configurations, and also prove that there is no monodromy.

- (1) e, f, g and h are handleslide equivalences. In our setting, we must consider changes of both of the attaching curves, as well as changes of the doubling datum. Note these are also referred to as *strong equivalences* in the literature, e.g., in [18]. The diagram commutes up to chain homotopy by Theorem 5.2. We note that in [9], the stated rectangle has the property that e and h are α -equivalences, while f and g are β -equivalences. However, the rectangle therein is also a rectangle of *isotopy diagrams*, so we must consider changes of both the alpha and beta curves along each edge.
- (2) e and h are handleslide equivalences of attaching curves and doubling curves, while f and g are both stabilizations. Commutativity follows from Proposition 11.1.
- (3) e and h are handleslide equivalences, while f and g are diffeomorphisms preserving the framing ξ . Commutativity around such rectangles is tautological.
- (4) e, f, g and h are all stabilizations. Furthermore, there are disjoint 3-balls B_1 and B_2 , such that e and h correspond to a stabilization in B_1 , while f and g correspond to a stabilization in B_2 . Commutativity around such rectangles follows from Proposition 11.1.
- (5) The maps e and h are stabilizations, while f and g are diffeomorphisms preserving the framing. Furthermore, the stabilization ball for h is the image of the stabilization ball for e under the diffeomorphism g . Commutativity around such rectangles is tautological.

The remaining axioms of [9, Definition 3.32], continuity and handleswap invariance, were verified in Proposition 5.4 and Theorem 12.2. ■

13.2. Functoriality

We now finish defining our cobordism maps and sketch the proof of Theorem 1.3, showing that these maps are well defined. Our proof is modeled on [24]. Since our approach is standard, we only provide an overview and sketch the points at which our construction deviates from [24].

We begin with the construction, and will shortly sketch a proof of invariance. Suppose W is a cobordism from Y_1 to Y_2 , and assume basepoints in Y_1 and Y_2 , as well as a framed path γ connecting them, have been chosen. We assume that W , Y_1 , Y_2 are connected. We pick a Morse function $f: W \rightarrow [0, 1]$. We may assume that the indices of the critical points of f are non-decreasing, and we assume all critical values are distinct and that f has critical points only of index 1, 2 and 3. Next, we pick a gradient like vector field v such that our path γ is a flow-line. A choice of gradient like vector field gives a diffeomorphism

$$W \cong W(\mathbb{S}_2) \cup W(\mathbb{S}_1) \cup W(\mathbb{S}_0),$$

where \mathbb{S}_i are framed i -dimensional links and $\mathbb{S}_0 \subseteq Y_1$, $\mathbb{S}_1 \subseteq Y_1(\mathbb{S}_0)$, and $\mathbb{S}_2 \subseteq Y_1(\mathbb{S}_0)(\mathbb{S}_1)$. If $\mathfrak{s} \in \text{Spin}^c(W)$ is self-conjugate, we define our cobordism map as the composition

$$CFI(W, \xi, \mathfrak{s}) = CFI(W(\mathbb{S}_2)) \circ CFI(W(\mathbb{S}_1), \mathfrak{s}|_{W(\mathbb{S}_1)}) \circ CFI(W(\mathbb{S}_0)).$$

We now prove Theorem 1.3. As a first step, we show that the 2-handle map is well defined.

Lemma 13.1. *Our 2-handle map $CFI(W(\mathbb{S}_1), \mathfrak{s}|_{W(\mathbb{S}_1)})$ is independent of the choice of diagram subordinate to a particular bouquet. Furthermore, it is also invariant from the choice of bouquet, as well as handleslides amongst the components of \mathbb{S}_1 .*

Proof. Ozsváth and Szabó [24, Lemma 4.5] described five moves which relate any two bouquets of a fixed link. Invariance under the first four of these moves follows immediately from invariance of our map from handleslides amongst the curves in our Heegaard triple, Proposition 10.5. The final move described by Ozsváth and Szabó is stabilization of the Heegaard triple. Invariance from this move is a formal consequence of our definition of the stabilization map. Indeed we defined the naturality map for stabilizations as the composition of the 1-handle map, followed by a canceling 2-handle map. Both the 1-handle map and map for a canceling 2-handle map commute with the 2-handle map $CFI(W(\mathbb{S}_1), \mathfrak{s}|_{W(\mathbb{S}_1)})$ by Propositions 10.6 and 10.7. Hence, our 2-handle map is independent of the choice of diagram subordinate to a fixed bouquet.

Next, we note that independence of the choice of bouquet, as well as independence of handleslides amongst the components of \mathbb{S}_1 both follow from independence of the

map from handleslides of the attaching curves, which we proved in Proposition 10.5. Compare [24, Lemma 4.8]. \blacksquare

Proof of Theorem 1.3. Firstly, we verify that the construction is independent of the choice of gradient-like vector field v for f . Subsequently, we sketch that the cobordism map is invariant from f .

The space of gradient like vector fields for a fixed Morse function is connected. Since we have already decomposed W along level sets which separate critical points of different indices, the codimension 1 singularities of a path of gradient like vector fields are handleslides amongst link components of the same indices. For critical points of index 1 or 3, invariance under these moves can instead be proven by reordering the attachment of the handles using the composition law for 1-handles and 3-handles, Proposition 9.7, since a handleslide is the same as an isotopy after attaching one of the handles. For 2-handles, we cannot add additional level sets because of the Spin^c decomposition, so instead handleslide invariance is proven directly in Lemma 13.1.

We now consider invariance under f . By standard Cerf theory [3, 12], any two Morse functions f_0 and f_1 with critical points ordered monotonically may be connected with a 1-parameter family $(f_t)_{t \in [0,1]}$ whose critical points are also ordered monotonically, except at finitely many points where a birth-death singularity occurs. Furthermore, if f_0 and f_1 have no index 0 or 4 critical points, then each f_t may also be chosen to have no index 0 or 4 critical points. We now show that the maps are invariant under index 1/2 handle cancelations, as well as index 2/3 handle cancelations. Invariance under such handle cancelation is essentially automatic from the definition of our naturality map for stabilizations. Indeed, consider the case that W has a Morse function with only index 2 critical points, whose descending manifolds intersect Y_0 in a framed link \mathbb{S}_1 . Let $\mathbb{S}_0 \subseteq Y_0$ be a 0-sphere, and let \mathbb{K} be a framed knot in $Y_0(\mathbb{S}_0)$ which intersects the co-core of \mathbb{S}_0 at exactly 1-point. Assume \mathbb{S}_1 is disjoint from \mathbb{S}_0 and \mathbb{K} . Using the composition law of Proposition 10.6, we obtain

$$\begin{aligned} CFI(W(\mathbb{S}_1 \cup \mathbb{K})) \circ CFI(W(\mathbb{S}_0)) \\ \simeq CFI(W(\mathbb{S}_1)) \circ CFI(W(\mathbb{K})) \circ CFI(W(\mathbb{S}_0)) := CFI(W(\mathbb{S}_1)) \circ \sigma, \end{aligned}$$

where σ is the naturality map for stabilization (and we are omitting Spin^c structures from the notation). Invariance under index 2/3 handle cancelations follows entirely analogously. The proof is complete. \blacksquare

14. The cobordism map for $S^2 \times S^2$

In this section we compute the cobordism map for $S^2 \times S^2$. We prove the following proposition.

Proposition 14.1. *Let W denote $S^2 \times S^2$, with two 4-balls removed. Let \mathfrak{s} denote the unique self-conjugate Spin^c structure on W . Then, the cobordism map*

$$CFI(W, \mathfrak{s}): \mathbb{F}[U, Q]/Q^2 \rightarrow \mathbb{F}[U, Q]/Q^2$$

is multiplication by Q .

Using the composition law, we may decompose W as two 2-handle cobordisms, $W = W_2 \circ W_1$. The first 2-handle is attached along a 0-framed unknot in S^3 to form W_1 . The second 2-handle is attached along an S^1 fiber of $S^1 \times S^2$ to form W_2 . Note that, in both cases, one end of the cobordism is a copy of S^3 , on which the map Φ is null-homotopic. This implies that the choice of framings of the basepoints of the three-manifolds and along choices of paths connecting them do not affect the final computation, and we therefore omit them.

We start with the standard Heegaard triple $(\Sigma, \alpha', \alpha, \beta)$ for W_1 . In Figure 14.1, we show the diagram $(\Sigma \# \bar{\Sigma}, \alpha' \bar{\beta}, \alpha \bar{\beta}, \beta \bar{\beta}, \Delta)$. Therein, the top half of the figure is Σ , and the bottom half is $\bar{\Sigma}$. This figure also encodes $(\Sigma, \alpha', \alpha, \beta)$.

The first quadrilateral counting map in equation (10.4) is $h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}$. The Heegaard quadruple $(\Sigma \# \bar{\Sigma}, \alpha' \bar{\beta}, \alpha \bar{\beta}, \beta \bar{\beta}, \Delta)$ is depicted in Figure 14.1. This diagram is weakly admissible. The input for $h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}$ from $CF(\beta \bar{\beta}, \Delta)$ is $\Theta_{\beta \bar{\beta}, \Delta}^+$, and the input from $CF(\alpha' \bar{\beta}, \alpha \bar{\beta})$ is $\Theta_{\alpha' \bar{\beta}, \alpha \bar{\beta}}^+$. We need only evaluate $h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}(\Theta_{\alpha \bar{\beta}, \beta \bar{\beta}}^+)$, since $F_1^{\bar{\beta}, \bar{\beta}}(\Theta_{\alpha, \beta}) = \Theta_{\alpha \bar{\beta}, \beta \bar{\beta}}^+$, and for grading reasons $h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}(\Theta_{\alpha \bar{\beta}, \beta \bar{\beta}}^+)$ is a multiple of $\Theta_{\alpha' \bar{\beta}, \Delta}^+$. Hence, we need only consider holomorphic quadrilaterals in

$$\pi_2(\Theta_{\alpha' \bar{\beta}, \alpha \bar{\beta}}^+, \Theta_{\alpha \bar{\beta}, \beta \bar{\beta}}^+, \Theta_{\beta \bar{\beta}, \Delta}^+, \Theta_{\alpha' \bar{\beta}, \Delta}^+).$$

There is an index -1 domain (shaded). This class has a holomorphic representative for a unique conformal class of rectangle. Indeed, for each slit length (along the alpha and beta arcs in the interior of P_2) there is a map of a varying conformal class of rectangles to the domain of P_2 . Meanwhile, there is a holomorphic rectangle mapping to P_1 for a unique conformal structure.

By inspection, there are no other positive domains, and so,

$$h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta}(\Theta_{\alpha \bar{\beta}, \beta \bar{\beta}}^+) = \Theta_{\alpha' \bar{\beta}, \Delta}^+.$$

The Heegaard quadruple for the second map $h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\Delta \rightarrow \alpha \bar{\alpha}'}$ is depicted in Figure 14.2. We need only consider summands in the image containing $\mathbf{x} = \{x_0^-, x_1^+\}$, since grading considerations imply that the only other generator that can appear in the image of $h_{\alpha \bar{\beta} \rightarrow \alpha' \bar{\beta}}^{\Delta \rightarrow \alpha' \bar{\alpha}'} \circ f_{\alpha \bar{\beta}}^{\beta \bar{\beta} \rightarrow \Delta} \circ F_1^{\bar{\beta}, \bar{\beta}}$ is $\{x_0^+, x_1^-\}$, which is annihilated by $F_3^{\alpha', \alpha'}$.

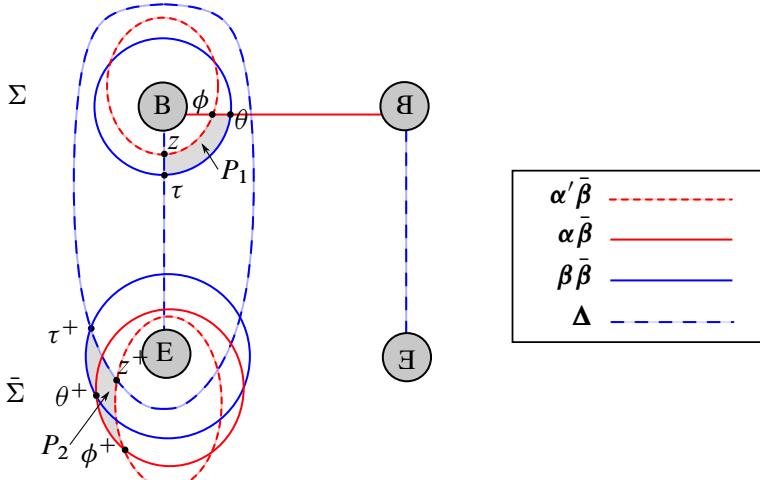


Figure 14.1. The Heegaard quadruple $(\Sigma \# \bar{\Sigma}, \alpha' \bar{\beta}, \alpha \bar{\beta}, \beta \bar{\beta}, \Delta)$. The generators of interest are $\Theta_{\alpha' \bar{\beta}, \alpha \bar{\beta}}^+ = \{\phi, \phi^+\}$, $\Theta_{\alpha \bar{\beta}, \beta \bar{\beta}}^+ = \{\theta, \theta^+\}$, $\Theta_{\beta \bar{\beta}, \Delta}^+ = \{\tau, \tau^+\}$ and $\Theta_{\alpha' \bar{\beta}, \Delta}^+ = \{z, z^+\}$.

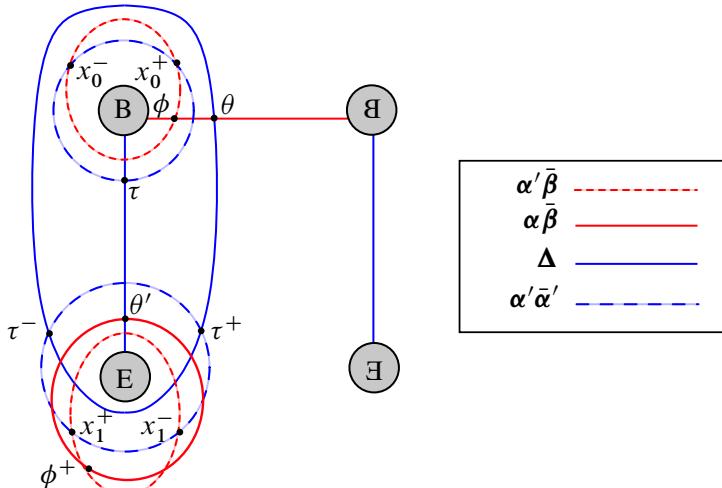


Figure 14.2. The Heegaard quadruple $(\Sigma, \alpha' \bar{\beta}, \alpha \bar{\beta}, \Delta, \alpha' \bar{\alpha}')$ in the definition of $h_{\alpha \beta \rightarrow \alpha' \bar{\beta}}^{\Delta \rightarrow \alpha' \bar{\alpha}'}$, with some intersection points labeled. The generators of interest are $\Theta_{\Delta, \alpha' \bar{\alpha}'}^+ = \{\tau^+, \tau\}$, $\Theta_{\alpha \bar{\beta}, \Delta}^+ = \{\theta, \theta'\}$, $\Theta_{\alpha' \bar{\beta}, \alpha \bar{\beta}}^+ = \{\phi^+, \phi\}$ and $x = \{x_0^-, x_1^+\}$.

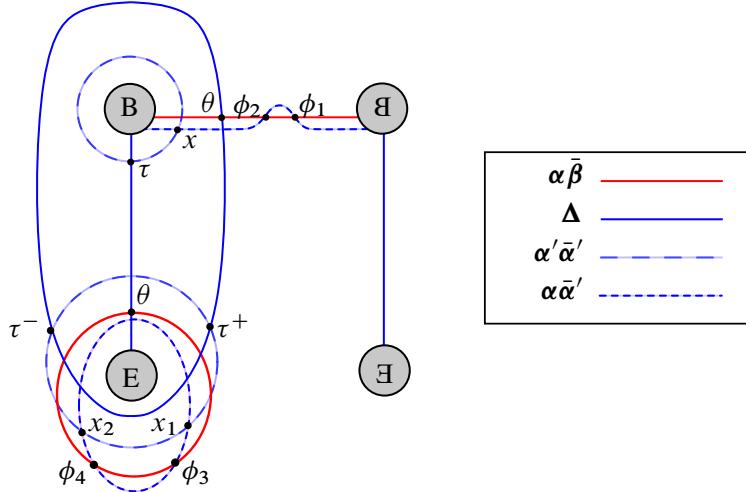


Figure 14.3. The Heegaard quadruple $(\Sigma, \alpha\bar{\beta}, \Delta, \alpha'\bar{\alpha}', \alpha\bar{\alpha}')$ in the definition of $H_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}'}$, with some intersection points labeled. The generators of interest are $\Theta_{\Delta, \alpha'\bar{\alpha}'}^+ = \{\tau^+, \tau\}$, $\Theta_{\alpha'\bar{\alpha}', \alpha\bar{\alpha}'}^+ = x_2 := \{x, x_2\}$, $y = \{\phi_2, \phi_3\}$ and $\Theta_{\alpha\bar{\beta}, \Delta} = \{\theta, \theta'\}$.

One checks directly that there are no positive domains among those generators for $h_{\alpha\bar{\beta} \rightarrow \alpha'\bar{\beta}}^{\Delta \rightarrow \alpha'\bar{\alpha}'}$ as well as $H_{\alpha\bar{\beta}}^{\Delta \rightarrow \alpha\bar{\alpha}'}$, the latter of which is shown in Figure 14.3.

It follows that the map $F_{W_1, \mathbf{s}|_{W_1}}$ takes the form

$$\left(CFI(S^3) \xrightarrow{F_{W_1, \mathbf{s}|_{W_1}}} CFI(S^1 \times S^2) \right)$$

$$\begin{array}{ccc}
 CF(S^3) & \xrightarrow{1 \mapsto \Theta^-} & CF(S^1 \times S^2) \\
 \simeq & \downarrow Q(1+\iota) & \downarrow Q(1+\iota) \\
 Q \cdot CF(S^3) & \xrightarrow{1 \mapsto \Theta^-} & Q \cdot CF(S^1 \times S^2)
 \end{array}$$

We now consider the 2-handle cobordism W from $S^1 \times S^2 \rightarrow S^3$. This is determined by the Heegaard triple $(\Sigma, \alpha'', \alpha', \beta)$, where α', β are as before, and α'' is a copy of α . Let $\theta_{\alpha', \beta}^-$ be the lower-degree generator of $CF^-(\Sigma, \alpha', \beta)$ and $\theta_{\bar{\beta}, \bar{\beta}}^+$ be the top-degree generator of $CF^-(\bar{\Sigma}, \bar{\beta}, \bar{\beta})$. We need to calculate $h_{\alpha'\bar{\beta} \rightarrow \alpha''\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta}(\theta_{\alpha', \beta}^-, \theta_{\bar{\beta}, \bar{\beta}}^+)$, $h_{\alpha'\bar{\beta} \rightarrow \alpha''\bar{\beta}}^{\Delta, \alpha''\bar{\alpha}''}(\Theta_{\alpha'\bar{\beta}, \Delta}^-)$, and $H_{\alpha'\bar{\beta}}^{\Delta \rightarrow \alpha''\bar{\alpha}''}(\Theta_{\alpha'\bar{\beta}, \Delta}^-)$.

We consider diagrams for each of those quadruples. Indeed, the first quadruple is represented in Figure 14.1 with the role of the $\alpha'\bar{\beta}$ and $\alpha\bar{\beta}$ curves exchanged.

A direct calculation shows there are no positive domains for $\pi_2(\Theta_{\alpha''\bar{\beta},\alpha'\bar{\beta}}^+, \theta_{\alpha',\beta}^- \times \theta_{\bar{\beta},\bar{\beta}}^+, \Theta_{\beta\bar{\beta},\Delta}^+, \Theta_{\alpha''\bar{\beta},\Delta}^+)$ so $h_{\alpha'\bar{\beta} \rightarrow \alpha''\bar{\beta}}^{\beta\bar{\beta} \rightarrow \Delta}(\theta_{\alpha',\beta}^- \times \theta_{\bar{\beta},\bar{\beta}}^+) = 0$. Entirely similar remarks apply for $h_{\alpha'\bar{\beta} \rightarrow \alpha''\bar{\beta}}^{\Delta, \alpha''\bar{\alpha}''}(\Theta_{\alpha'\bar{\beta},\Delta}^-)$, and $H_{\alpha'\bar{\beta}}^{\Delta \rightarrow \alpha'\bar{\alpha}''}(\Theta_{\alpha'\bar{\beta},\Delta}^-)$. Computing as above, we obtain that

$$\begin{array}{ccc} \left(CFI(S^1 \times S^2) \xrightarrow{F_{W_2, \mathfrak{s}|W_2}} CFI(S^3) \right) \\ \\ \begin{array}{ccc} CF(S^1 \times S^2) & \xrightarrow{\Theta^+ \mapsto 1, \Theta^- \mapsto 0} & CF(S^3) \\ \simeq & \downarrow Q(1+i) & \downarrow Q(1+i) \\ Q \cdot CF(S^1 \times S^2) & \xrightarrow{\Theta^+ \mapsto 1, \Theta^- \mapsto 0} & Q \cdot CF(S^3) \end{array} \end{array}$$

Composing the $CFI(W_2, \mathfrak{s}_2) \circ CFI(W_1, \mathfrak{s}_1)$, we obtain multiplication by Q .

Remark 14.2. The cobordism W_1 is obtained topologically by turning the cobordism W_2 around and reversing its orientation. The maps $CFI(W_1, \mathfrak{s}_1)$ and $CFI(W_2, \mathfrak{s}_2)$ become dual over $\mathbb{F}[U, Q]/Q^2$ after performing a change of basis, corresponding to a nontrivial automorphism on $CFI(S^1 \times S^2)$. Compare [24, Theorem 3.5]. Note that our definition of the 2-handle map in equation (10.4) does not naturally dualize to another 2-handle map. We hope to investigate duality more precisely in future work.

15. Knots and links

In this section, we describe how to adapt the naturality and functoriality results of the previous sections to the case of knots and links, and in particular prove Theorems 1.4 and 1.5.

If L is a link, in this section we write $\mathcal{CFI}(L)$ for the pair $(\mathcal{FL}(L), \iota_L)$. Analogously to the case of knots, our construction requires a choice of orientation on L , so there are potentially 2^ℓ different potential models for ι_L . To simplify the notation, we assume that an orientation is fixed, and we consider only the two orientations which are either coherent, or opposite to our preferred orientation. We write $\iota_{L,+}$ and $\iota_{L,-}$ for these models of the link involution.

In analogy to our construction in the case of closed 3-manifolds, we will describe several expanded models of the link involution $\tilde{\iota}_{L,\pm}$ in Section 15.1. To streamline the presentation, we only define transition maps with respect to the expanded models (unlike in the case of closed 3-manifolds, where we defined both non-expanded and expanded models). Since we only consider the expanded models of the involution

in this section, we will write \mathfrak{D} for the expanded, doubling enhanced link diagram we consider in Section 15.1. This notation departs from previous sections, where we wrote $\tilde{\mathfrak{D}}$ for expanded models of the involution. Similarly, we write $\Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_2}$ for the expanded models of the transition maps in this section, instead of $\tilde{\Psi}_{\tilde{\mathfrak{D}}_1 \rightarrow \tilde{\mathfrak{D}}_2}$.

In this section, we show how to adapt the techniques of the earlier sections to prove the following naturality theorem, which implies Theorem 1.4.

Theorem 15.1. *Suppose that L is a link in a 3-manifold.*

- (1) *The transition maps $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}}$ are well defined up to homotopies of morphisms of ι_L -complexes.*
- (2) $\Psi_{\mathfrak{D} \rightarrow \mathfrak{D}} \simeq \text{id}_{\mathcal{CFI}(\mathfrak{D})}$.
- (3) *If \mathfrak{D}_1 , \mathfrak{D}_2 and \mathfrak{D}_3 are three expanded, doubling enhanced Heegaard link diagrams for (Y, L) , then $\Psi_{\mathfrak{D}_2 \rightarrow \mathfrak{D}_3} \circ \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_2} \simeq \Psi_{\mathfrak{D}_1 \rightarrow \mathfrak{D}_3}$.*

Additionally, we prove the following functoriality theorem, which is an expanded version of Theorem 1.5.

Theorem 15.2. *Suppose that (W, Σ) is a link cobordism from (Y_1, L_1) to (Y_2, L_2) . Suppose further that Σ consists of a collection of annuli, each with one boundary component in Y_1 , and one boundary component in Y_2 . Suppose further that each annulus is decorated with two parallel longitudinal arcs, $\mathfrak{s} \in W$ is a Spin c structure such that $\bar{\mathfrak{s}} = \mathfrak{s} + PD[\Sigma]$. Then, the cobordism map $\mathcal{CFI}(W, \Sigma, \mathfrak{s})$ defined via the doubling model is an enhanced ι_L -homomorphism, and is well-defined up to enhanced ι_L -homotopy.*

15.1. Expanded models of the knot and link involutions

In this section, we describe our expanded model of the knot and link involution, similar to the expanded 3-manifold involution which appeared in Section 7. This is the model we use to prove Theorem 15.1.

The construction is as follows. Suppose that L is a link in Y , and $(\Sigma, \alpha, \beta, w, z)$ is a Heegaard link diagram for (Y, L) such that each component of L contains exactly two basepoints. We pick an auxiliary point $p \in \Sigma \setminus (\alpha \cup \beta \cup \{w, z\})$, at which we will form the connected sum of Σ and $\bar{\Sigma}$.

We write c for the curve which is a meridian of the connected sum tube, as in the 3-manifold case. There is a natural 1-handle map

$$F_1^{c\bar{\beta}, c\bar{\beta}} : \mathcal{CFI}(\Sigma, \alpha, \beta, w, z) \rightarrow \mathcal{CFI}(\Sigma \# \bar{\Sigma}, \alpha c \bar{\beta}, \beta c \bar{\beta}, w \cup \bar{z}, z \cup \bar{w}).$$

We pick a collection of properly embedded and pairwise disjoint arcs $\delta_1, \dots, \delta_{2g+2|L|-1}$ on $\Sigma \setminus N(p)$, such that after cutting $\Sigma \setminus N(p)$ along $\delta_1, \dots, \delta_{2g+2|L|-1}$,

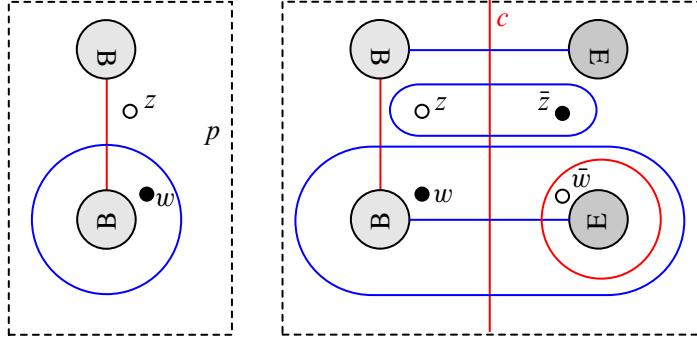


Figure 15.1. A doubly pointed Heegaard knot diagram with a special point p (left), and the double (right).

we are left with $2|L|$ punctured disks. We assume that each disk contains exactly one basepoint from $\mathbf{w} \cup \mathbf{z}$. Equivalently, we may assume that $\delta_1, \dots, \delta_{2g+2|L|-1}$ forms a basis of $H_1(\Sigma \setminus (\mathbf{w} \cup \mathbf{z}), p)$. We double the arcs $\delta_1, \dots, \delta_{2g+2|L|-1}$ to obtain the attaching curves Δ . See Figure 15.1 for an example.

The complex $\mathcal{CFL}(\Sigma \# \bar{\Sigma}, \beta c \bar{\beta}, \Delta, \mathbf{w} \cup \bar{z}, z \cup \bar{w})$ represents an $|L|$ -component unlink in $\#^g S^1 \times S^2$, where each component contains 4 basepoints. There are two canonical homology classes, represented by cycles $\Theta_{\beta c \bar{\beta}, \Delta}^{\mathbf{w}}$ and $\Theta_{\beta c \bar{\beta}, \Delta}^{\mathbf{z}}$. These are the generators of the top degree of homology with respect to the gradings $\text{gr}_{\mathbf{w}}$ and $\text{gr}_{\mathbf{z}}$, respectively. See [35, Lemma 3.7]. We define

$$f_{\alpha c \bar{\beta}}^{\beta c \bar{\beta} \rightarrow \Delta; \mathbf{z}}(-) = f_{\alpha c \bar{\beta}, \beta c \bar{\beta}, \Delta}(-, \Theta_{\beta c \bar{\beta}, \Delta}^{\mathbf{w}}),$$

and we similarly define $f_{\alpha c \bar{\beta}}^{\beta c \bar{\beta} \rightarrow \Delta; \mathbf{w}}$ to be the holomorphic triangle map with special input $\Theta_{\beta c \bar{\beta}, \Delta}^{\mathbf{z}}$. (The reason that the \mathbf{z} -map has the \mathbf{w} -generator as its input is due to the interpretation of the map in terms of decorated link cobordisms from [35].)

Similarly, there are two distinguished classes $\Theta_{\Delta, \alpha c \bar{\alpha}}^{\mathbf{w}}, \Theta_{\Delta, \alpha c \bar{\alpha}}^{\mathbf{z}}$ in $\mathcal{CFL}(\Sigma \# \bar{\Sigma}, \Delta, \alpha c \bar{\alpha}, \mathbf{w} \cup \bar{z}, z \cup \bar{w})$.

We define the following two expanded models of the knot involution:

$$\begin{aligned} \tilde{\iota}_{L,+} &= \eta_K \circ F_3^{\alpha c, \alpha c} \circ f_{\alpha c \bar{\alpha}}^{\Delta \rightarrow \alpha c \bar{\alpha}; \mathbf{w}} \circ f_{\alpha c \bar{\alpha}}^{\beta c \bar{\beta} \rightarrow \Delta; \mathbf{z}} \circ F_1^{c \bar{\beta}, c \bar{\beta}}, \\ \tilde{\iota}_{L,-} &= \eta_K \circ F_3^{\alpha c, \alpha c} \circ f_{\alpha c \bar{\alpha}}^{\Delta \rightarrow \alpha c \bar{\alpha}; \mathbf{z}} \circ f_{\alpha c \bar{\alpha}}^{\beta c \bar{\beta} \rightarrow \Delta; \mathbf{w}} \circ F_1^{c \bar{\beta}, c \bar{\beta}}. \end{aligned}$$

Lemma 15.3. *The maps $\tilde{\iota}_{L,-}$ and $\tilde{\iota}_{L,+}$ are homotopic to the maps $\iota_{L,-}$ and $\iota_{L,+}$, respectively, which were defined in Section 3.4.*

Proof. We focus on the case when L is a doubly pointed knot K , since the proof for links is not substantially different.

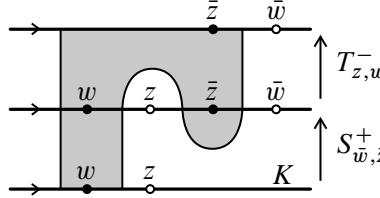


Figure 15.2. Decorations on a cylindrical link cobordism $[0, 1] \times K$ corresponding to a composition of a quasi-stabilization followed by a quasi-destabilization.

The proof is to identify each of the above compositions with a cobordism map from [35]. The map $F_1^{c\bar{\beta}, c\bar{\beta}}$ is equal to a composition of g 4-dimensional 1-handle maps, as well as the birth cobordism map for adding an unknot component. The map $F_3^{\alpha c, \alpha c}$ is similarly equal to the composition of g 3-handle maps, as well as a death cobordism map which deletes an unknot.

We now consider the triangle maps. The map $f_{\alpha c \bar{\alpha}}^{\beta c \bar{\beta} \rightarrow \Delta; z}$ is the composition of g 4-dimensional 2-handle maps (as in the ordinary doubling model of the involution), as well as the saddle cobordism map which attaches a band which has both of its ends in a component of $K \setminus \{w, z\}$ which is oriented from z to w (i.e., which lies in the beta handlebody). Furthermore, the band is contained in the type- z subsurface of the link cobordism. The map $f_{\alpha c \bar{\alpha}}^{\Delta \rightarrow \alpha c \bar{\alpha}; w}$ has a similar description, except that the band is instead in the type- w subsurface.

To make the identifications in the main statement, we argue as follows. By canceling the 1-handles and 2-handles, we may identify the composition $f_{\alpha c \bar{\alpha}}^{\beta c \bar{\beta} \rightarrow \Delta; z} \circ F_1^{c\bar{\beta}, c\bar{\beta}}$ with the composition of a birth cobordism, followed by a w -saddle map. By [35, Proposition 8.5], we may identify this composition with the quasi-stabilization map $S_{\bar{z}, \bar{w}}^+$, from [35, Section 4] and [32]. Analogously, we may identify the composition $F_3^{\alpha c, \alpha c} \circ f_{\alpha c \bar{\alpha}}^{\Delta \rightarrow \alpha c \bar{\alpha}; w}$ with $T_{z,w}^-$. The composition $T_{z,w}^- \circ S_{\bar{z}, \bar{w}}^+$ is homotopic to the basepoint moving map which sends w to \bar{z} and z to \bar{w} , both moving in the positive direction by [35, Section 4.4]. See Figure 15.2. ■

Remark 15.4. The compositions $\eta_K \circ F_3^{\alpha c, \alpha c} \circ f_{\alpha c \bar{\alpha}}^{\Delta \rightarrow \alpha c \bar{\alpha}; z} \circ f_{\alpha c \bar{\alpha}}^{\beta c \bar{\beta} \rightarrow \Delta; z} \circ F_1^{c\bar{\beta}, c\bar{\beta}}$ and $\eta_K \circ F_3^{\alpha c, \alpha c} \circ f_{\alpha c \bar{\alpha}}^{\Delta \rightarrow \alpha c \bar{\alpha}; w} \circ f_{\alpha c \bar{\alpha}}^{\beta c \bar{\beta} \rightarrow \Delta; w} \circ F_1^{c\bar{\beta}, c\bar{\beta}}$ are null-homotopic. Indeed, the same argument as above identifies the two maps with the link cobordism maps for a cylinder which contains a null-homotopic loop in the dividing set. Such a cobordism map induces the trivial map, since it may be identified with one of the compositions $S_{w,z}^- S_{w,z}^+$ or $T_{w,z}^- T_{w,z}^+$, which both vanish by [35, Lemma 4.13].

Remark 15.5. As previously mentioned, if L has ℓ -components, there are 2^ℓ natural models of the link involution, depending on how we twist each link component. The

expanded model described above can also be used to realize any of these models. We consider the diagram $(\Sigma \# \bar{\Sigma}, \beta c \bar{\beta}, \Delta, \mathbf{w} \cup \bar{z}, z \cup \bar{w})$. Suppose that \mathbb{O} is a map from the set of link components of L to the set $\{+, -\}$ (corresponding to a choice of whether to twist in the positive or negative direction). We let $\mathbb{Z}(\mathbb{O})$ consist of the basepoints from $\mathbf{w} \cup \bar{z}$, for components where $\mathbb{O} = -$, as well as the basepoints from $z \cup \bar{w}$ for components where $\mathbb{O} = +$. Let $\mathbb{W}(\mathbb{O})$ denote the complement. Then, each collection $\mathbb{Z}(\mathbb{O})$ and $\mathbb{W}(\mathbb{O})$ determine Maslov gradings on $\mathcal{CF}\mathcal{L}(\Sigma \# \bar{\Sigma}, \beta c \bar{\beta}, \Delta, \mathbf{w} \cup \bar{z}, z \cup \bar{w})$, and we let $\Theta^{\mathbb{Z}(\mathbb{O})}$ and $\Theta^{\mathbb{W}(\mathbb{O})}$ be the corresponding top-degree generators. A straightforward extension of Lemma 15.3 shows that the involution $\tilde{\iota}_{L, \mathbb{O}}$ is computed by using $\Theta^{\mathbb{Z}(\mathbb{O})}$ for the first triangle map, and $\Theta^{\mathbb{W}(\mathbb{O})}$ for the second.

15.2. Naturality for knots and links

Most of the ideas from the previous sections carry over to define transition maps and cobordism maps for the model of involutive link Floer homology described in Section 15.1.

The main new technical subtlety is that the model of involutive link Floer homology in Section 15.1 has an additional piece of data: the choice of connected sum point p . If p and p' are two choices of special points on Σ , a choice of path on Σ from p to p' can be used to relate the two models where the connected sum is taken at p or p' , by using a point-pushing diffeomorphism along this path. In principle, a transition map defined in this manner might depend on the choice of path. To handle this issue, we introduce further expansions of the doubling model, as follows.

Suppose that $\mathbf{p} = \{p_1, \dots, p_n\}$ is a collection of marked points on $\Sigma \setminus (\alpha \cup \beta \cup \mathbf{w} \cup \mathbf{z})$. We may form an expanded model of the involution, as follows. Firstly, we consider the surface $\Sigma \#_{\mathbf{p}} \bar{\Sigma}$, where we attach a connected sum tube at each point in \mathbf{p} . Let $\mathbf{c} = \{c_1, \dots, c_n\}$ denote a collection of meridians of the tubes. We pick properly embedded arcs $\delta_1, \dots, \delta_{2g+2|L|+|\mathbf{p}|-2}$, which form a basis of $H_1(\Sigma \setminus (\mathbf{w} \cup \mathbf{z}), \mathbf{p})$. It is straightforward to see that if Δ is constructed by doubling $\delta_1, \dots, \delta_{2g+2|L|+|\mathbf{p}|-2}$, then $(\Sigma \#_{\mathbf{p}} \bar{\Sigma}, \alpha c \bar{\beta}, \Delta, \mathbf{w} \cup \bar{z}, z \cup \bar{w})$ is a diagram for (Y, L) , where each component of L is given four basepoints.

Given a collection of points \mathbf{p} as well as doubling arcs, as above, the construction from Section 15.1 adapts verbatim to produce a model of the link involution. It remains to relate the models for different choices of \mathbf{p} .

Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is a Heegaard link diagram, and \mathbf{p} is a subset of $\Sigma \setminus (\alpha \cup \beta \cup \mathbf{w} \cup \mathbf{z})$, and q is a point on $\Sigma \setminus (\alpha \cup \beta \cup \mathbf{w} \cup \mathbf{z} \cup \mathbf{p})$. Suppose that $\mathfrak{D}_{\mathbf{p}}$ and $\mathfrak{D}_{\mathbf{p} \cup \{q\}}$ are doubling enhancements of \mathcal{H} , which use the connected sum points \mathbf{p} and $\mathbf{p} \cup \{q\}$, respectively. We now define *expansion* and *contraction* morphisms

$$\mathcal{E}_q: \mathcal{CFKI}(\mathfrak{D}_{\mathbf{p}}) \rightarrow \mathcal{CFKI}(\mathfrak{D}_{\mathbf{p} \cup \{q\}}), \quad \mathcal{C}_q: \mathcal{CFKI}(\mathfrak{D}_{\mathbf{p} \cup \{q\}}) \rightarrow \mathcal{CFKI}(\mathfrak{D}_{\mathbf{p}}).$$

We define the map \mathcal{E}_q in equation (15.1). Our definition of \mathcal{E}_q is similar to our definition of the 1-handle map, and \mathcal{C}_q is defined similarly to the 3-handle map. In particular, \mathcal{E}_q is given by the following hyperbox.

$$\begin{array}{ccccc}
\mathcal{CFK}(\alpha, \beta) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha, \beta) & & \\
\downarrow & & \downarrow & & \\
F_1^{cp\bar{\beta}, cp\bar{\beta}} & & & & F_1^{cp\bar{\beta}, cp\bar{\beta}, cp\bar{\beta}, cp\bar{\beta}} \\
\downarrow & & \downarrow & & \downarrow \\
\mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & \xrightarrow{F_1^{c_q, c_q}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) \\
\downarrow f_{\alpha c_p \bar{\beta}}^{\beta c_p \bar{\beta} \rightarrow \Delta; z} & & \downarrow f_{\alpha c_p c_q \bar{\beta}}^{\beta c_p c_q \bar{\beta} \rightarrow \Delta c_q; z} & \searrow & \downarrow f_{\alpha c_p c_q \bar{\beta}}^{\beta c_p c_q \bar{\beta} \rightarrow \Delta'; z} \\
\mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta) & \xrightarrow{F_1^{c_q, c_q}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \Delta c_q) & \xrightarrow{f_{\alpha c_p c_q \bar{\beta}}^{\Delta c_q \rightarrow \Delta'}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \Delta') \\
\downarrow f_{\alpha c_p \bar{\beta}}^{\Delta \rightarrow \alpha c_p \bar{\alpha}; w} & & \downarrow f_{\alpha c_p c_q \bar{\beta}}^{\Delta c_q \rightarrow \alpha c_p c_q \bar{\alpha}; w} & \searrow & \downarrow f_{\alpha c_p c_q \bar{\beta}}^{\Delta' \rightarrow \alpha c_p c_q \bar{\alpha}; w} \\
\mathcal{CFK}(\alpha c_p \bar{\beta}, \alpha c_p \bar{\alpha}) & \xrightarrow{F_1^{c_q, c_q}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \alpha c_p c_q \bar{\alpha}) & \xrightarrow{B_\tau} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \alpha c_p c_q \bar{\alpha}) \\
\downarrow F_3^{ac_p, ac_p} & & & & \downarrow F_3^{ac_p c_q, ac_p c_q} \\
\mathcal{CFK}(\bar{\beta}, \bar{\alpha}) & \xrightarrow{\text{id}} & \mathcal{CFK}(\bar{\beta}, \bar{\alpha}) & &
\end{array} \tag{15.1}$$

In the definition of \mathcal{E}_q , we write τ for a closed curve on $\Sigma \#_{p \cup \{q\}} \bar{\Sigma}$ which is symmetric with respect to the \mathbb{Z}_2 -symmetry, and has a single transverse intersection with c_q .

Lemma 15.6. (1) *The maps \mathcal{E}_q and \mathcal{C}_q commute with the transition maps for handle-slides and isotopies of the attaching curves. Furthermore, \mathcal{E}_q and \mathcal{C}_q are independent of the choice of τ , up to chain homotopy.*

- (2) *If $p \neq q$, then $[\mathcal{E}_p, \mathcal{C}_q] \simeq 0$, $[\mathcal{E}_p, \mathcal{E}_q] \simeq 0$, and $[\mathcal{C}_p, \mathcal{C}_q] \simeq 0$.*
- (3) *$\mathcal{E}_p \circ \mathcal{C}_p \simeq \text{id}$ and $\mathcal{C}_p \circ \mathcal{E}_p \simeq \text{id}$.*
- (4) *If ϕ is a surface isotopy of Σ which moves p to q and fixes all of the other connected sum points and basepoints, then $\phi_* \simeq \mathcal{C}_p \circ \mathcal{E}_q$.*

Proof. The proof of part (1) is essentially the same as for the 1-handle and 3-handle maps. See Lemmas 9.4 and 9.5.

The proofs of claims (2) and (3) amount to constructing five hyperboxes (one for each relation). All of these hyperboxes have a similar flavor, so we describe the hyperbox for the relation $\mathcal{C}_q \circ \mathcal{E}_q \simeq \text{id}$, and leave the remaining hyperboxes to the reader.

The hyperbox we build will correspond to the hypercube

$$\begin{array}{ccc}
 \mathcal{CFKI}(\mathfrak{D}_p) & \xrightarrow{-\varepsilon_q} & \mathcal{CFKI}(\mathfrak{D}_{p \cup \{q\}}) \\
 \downarrow \text{id} & \nearrow \text{id} & \downarrow \varepsilon_q \\
 \mathcal{CFKI}(\mathfrak{D}_p) & \xrightarrow{\text{id}} & \mathcal{CFKI}(\mathfrak{D}_p)
 \end{array}$$

We build four constituent hyperboxes, each corresponding to one of the four levels of equation (15.1). The highest level is the following diagram:

$$\begin{array}{ccccc}
 \mathcal{CFK}(\alpha, \beta) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha, \beta) & & \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \\
 \mathcal{CFK}(\alpha, \beta) & & \mathcal{CFK}(\alpha, \beta) & & \\
 \downarrow F_1^{cp\bar{\beta}, cp\bar{\beta}} & & \downarrow F_1^{cp c_q \bar{\beta}, cp c_q \bar{\beta}} & & \downarrow F_1^{cp\bar{\beta}, cp\bar{\beta}} \\
 \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & \xrightarrow{F_1^{\bar{\beta}, \bar{\beta}}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) & & \\
 \downarrow \text{id} & \downarrow \text{id} & \downarrow F_3^{c_q, c_q \circ B_\tau} & & \downarrow \text{id} \\
 \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & &
 \end{array}$$

The bottom-most level is very similar to the above, so we omit it.

We now consider the hyperbox corresponding to the second level from the top in equation (15.1):

$$\begin{array}{ccccc}
 \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & \xrightarrow{F_1^{c_q, c_q}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) & \xrightarrow{-\text{id}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} \\
 & & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) & & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) \\
 & & \downarrow B_\tau & & \downarrow B_\tau \\
 & & \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) \\
 & & \downarrow \text{id} & & \downarrow \text{id} \\
 & & \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p \bar{\beta}, \beta c_p \bar{\beta}) \\
 \downarrow \text{id} & & \downarrow \text{id} & & \downarrow \text{id} \\
 \mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta) & \xrightarrow{-F_1^{c_q, c_q}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \Delta c_q) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) \\
 \downarrow \text{id} & & \downarrow \text{id} & & \downarrow \text{id} \\
 & & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \Delta c_q) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p c_q \bar{\beta}, \beta c_p c_q \bar{\beta}) \\
 & & \downarrow B_\tau & & \downarrow B_\tau \\
 & & \mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta) \\
 & & \downarrow \text{id} & & \downarrow \text{id} \\
 & & \mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta) & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta) \\
 & & & \xrightarrow{\text{id}} & \mathcal{CFK}(\alpha c_p \bar{\beta}, \Delta)
 \end{array} \tag{15.2}$$

In equation (15.2), each unlabeled arrow is a triangle or quadrilateral map. The back right hypercube also has a length 3 map, which is not labeled. We now explain the

back-right hypercube. This hypercube is obtained by pairing $\alpha c_p c_q \bar{\beta}$ with the following hypercube of beta attaching curves:

$$\begin{array}{ccccc}
 \beta c_p c_q \bar{\beta} & \xrightarrow{1} & \beta c_p c_q \bar{\beta} & & \\
 \downarrow \Theta_{\beta c_p c_q \bar{\beta}, \Delta c_q}^z & \searrow \tau & \downarrow \Theta_{\beta c_p c_q \bar{\beta}, \Delta'}^z & \searrow \tau & \\
 \Delta c_q & \xrightarrow{\Theta_{\Delta c_q, \Delta'}^z} & \Delta' & \xrightarrow{\Theta_{\Delta', \Delta c_q}^z} & \Delta c_q \\
 \downarrow \Theta_{\Delta c_q, \Delta'}^z & \searrow \tau & \downarrow \eta & \searrow \Theta_{\Delta', \Delta c_q}^z & \\
 \Delta c_q & \xrightarrow{1} & \Delta c_q & \xrightarrow{1} & \Delta c_q
 \end{array} \tag{15.3}$$

The above hypercube of beta attaching curves is constructed as follows. Length 2 chains are constructed by adapting Lemma 9.2. Once these are chosen, a length 3 chain may be chosen by the standard cube-filling procedure. Namely, if C is the sum of all other terms of the length 3 relation, we know $\partial C = 0$ because of the hypercube relations on the other proper faces, and C is in $\mathcal{CFK}(\beta c_p c_q \bar{\beta}, \Delta c_q)$, a complex for an unlink in a connected sum of several copies of $S^1 \times S^2$, of gr_z -grading equal to that of the top degree generator of homology. By potentially adding a copy of the top degree generator to the 2-chain labeled η , in equation (15.3), the hypercube relations may be satisfied.

The hyperbox for the third level of equation (15.1) is constructed very similarly. Indeed, if we switch the roles of the alpha and beta curves, then the hypercube for the third level of equation (15.1) is obtained by dualizing equation (15.2). Stacking these hyperboxes and compressing gives the relation $\mathcal{C}_q \circ \mathcal{E}_q \simeq \text{id}$.

One important point to check is that the diagonal maps on the second and third hypercubes can be chosen to match. For example, to build the hypercube in equation (15.2), we may potentially have to add the top degree generator to η so that the hypercube relations are satisfiable. A similar issue occurs in constructing the length 2 map on the top face of the third hypercube as well. Nonetheless, as in the proof of Lemma 9.5, addition of the top degree generator to η in equation (15.3) has no effect on the diagonal map in the compression of equation (15.2). Hence, the compression of equation (15.2) may be stacked on top of the corresponding hypercube below it.

We now consider part (4). This follows from the first claims, as we now explain. Let λ be a path from p to q . Let $\phi_{q \rightarrow p}$ denote the surface isotopy which moves q to

p along λ , and is supported in a small neighborhood of λ . It is sufficient to show that

$$\phi_{q \rightarrow p} \circ \mathcal{C}_p \circ \mathcal{E}_q \simeq \text{id}.$$

Let x_0 be a point in the interior of λ . We compute as follows, using the previous relations:

$$\begin{aligned} \phi_{q \rightarrow p} \circ \mathcal{C}_p \circ \mathcal{E}_q &\simeq \phi_{q \rightarrow p} \circ \mathcal{C}_p \circ \mathcal{C}_{x_0} \circ \mathcal{E}_{x_0} \circ \mathcal{E}_q \\ &\simeq \phi_{q \rightarrow p} \circ \mathcal{C}_{x_0} \circ \mathcal{E}_q \circ \mathcal{C}_p \circ \mathcal{E}_{x_0} \\ &\simeq \mathcal{C}_{x_0} \circ \phi_{q \rightarrow p} \circ \mathcal{E}_q \circ \mathcal{C}_p \circ \mathcal{E}_{x_0} \\ &\simeq \mathcal{C}_{x_0} \circ \mathcal{E}_p \circ \mathcal{C}_p \circ \mathcal{E}_{x_0} \\ &\simeq \text{id}. \end{aligned}$$

Note that commuting $\phi_{q \rightarrow p}$ past \mathcal{C}_{x_0} in the second line is possible since the diffeomorphism $\phi_{q \rightarrow p}$ is isotopic to one which fixes x_0 . Next, $\phi_{q \rightarrow p} \circ \mathcal{E}_q \simeq \mathcal{E}_p$ by naturality of the expansion and contraction maps with respect to changes of attaching curves, which is part (1) of the statement. This establishes part (4). \blacksquare

15.3. Functoriality for knots and links

In this section, we describe several aspects of the proof of Theorem 15.2. Firstly, to construct cobordism maps for $(W, \Sigma): (Y_1, L_1) \rightarrow (Y_2, L_2)$ when Σ is a collection of annuli such that each component of Σ intersects Y_1 and Y_2 non-trivially, it is sufficient to describe cobordism maps on link Floer homology for 4-dimensional handles attached to the complement of a link.

Cobordism maps for 4-dimensional handles attached in the complement of a link are defined similarly to the maps for cobordisms between closed 3-manifolds (cf. [35, Section 5]). The 1-handle and 3-handle maps are defined as in Section 9, and the 2-handle maps are defined similarly to Section 10. One caveat is that there is an asymmetry in the definition of the 2-handle maps with respect to the basepoints, in that we choose the 2-handle maps to count holomorphic triangles which represent classes satisfying $\mathfrak{s}_w(\psi) = \mathfrak{s}$. Note that by [34, Lemma 3.8], we have

$$(\mathfrak{s}_w - \mathfrak{s}_z)(\psi) = PD[\Sigma].$$

In particular, the final hypercube we use in constructing the cobordism maps for our 2-handle map will have the following form:

$$\begin{array}{ccc} \mathcal{CFL}(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}) & \xrightarrow{f_{\bar{\beta}; \mathfrak{s}}^{\bar{\alpha} \rightarrow \bar{\alpha}'}} & \mathcal{CFL}(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}') \\ \downarrow \eta_L & & \downarrow \eta_L \\ \mathcal{CFL}(\Sigma, \alpha, \beta) & \xrightarrow{f_{\alpha \rightarrow \alpha', \mathfrak{s} + PD[\Sigma]}^{\beta}} & \mathcal{CFL}(\Sigma, \alpha', \beta) \end{array}$$

16. 2-handle maps and the knot surgery formula

In this section, we compute another family of involutive cobordism maps. If $K \subseteq S^3$, and $n \in \mathbb{Z}$, we will compute the map for the 2-handle cobordism from $S_{2n}^3(K)$ to S^3 . Our description will be in terms of the knot surgery formula of [4], and is in a similar spirit to the non-involutive description of the 2-handle map of Ozsváth and Szabó [27, Theorem 1.1].

We suppose that $K \subseteq S^3$ is a knot, and we consider the 2-handle cobordism $W'_m(K)$ from $S_m^3(K)$ to S^3 . This cobordism is Spin if and only if m is even. Furthermore, if $m = 2n$, then the unique self-conjugate Spin^c structure on $W'_{2n}(K)$ is the one which satisfies $\langle c_1(\mathfrak{s}), \Sigma \rangle = 0$, where Σ is obtained by capping K with the core of the 2-handle. This Spin^c structure restricts to $[n] \in \text{Spin}^c(S_{2n}^3(K))$ under the isomorphism $\text{Spin}^c(S_{2n}^3(K)) \cong \mathbb{Z}/2n$.

Our formula is phrased in terms of the mapping cone formula of Ozsváth and Szabó [27] and its involutive analog from [4]. We review these constructions briefly. We write $A_s(K)$ for the subcomplex of $CFK^\infty(K)$ generated by elements $\mathcal{U}^i \mathcal{V}^j \mathbf{x}$ where $A(\mathbf{x}) + j - i = 0$ and $j \geq -s$ and $i \geq 0$. We write $B_s(K)$ for the subspace generated by similar monomials, with no restriction on j . We write $\tilde{B}_s(K)$ for the subspace generated by monomials which satisfy $j \geq -s$, but with no restriction on i . We write $v_s: A_s \rightarrow B_s$ and $\tilde{v}_s: A_s \rightarrow \tilde{B}_s$ for the canonical inclusions. We recall from [27] that the mapping cone formula requires a choice of homotopy equivalence from $\tilde{B}_s(K)$ to $B_{s+m}(K)$. We write \mathfrak{F}_m for this map. We write $\mathbb{A}(K)$ for the direct product over s of $A_s(K)$, tensored with the power series ring $\mathbb{F}[[U]]$. Similarly $\mathbb{B}(K)$ denotes the direct product of the complexes $B_s(K)$, tensored with $\mathbb{F}[[U]]$.

We recall that Ozsváth and Szabó proved that if $K \subseteq S^3$ and $m \in \mathbb{Z}$, then

$$CF^-(S_m^3(K)) \cong \text{Cone}(v + h_m: \mathbb{A}(K) \rightarrow \mathbb{B}(K)),$$

where $h_m = \mathfrak{F}_m \tilde{v}$. In [4], we extended this to the involutive setting, and proved that

$$\begin{array}{ccc} \mathbb{A}(K) & \xrightarrow{v+h_m} & \mathbb{B}(K) \\ \mathfrak{CFI}^-(S_m^3(K)) \simeq & \begin{array}{c} \downarrow Q(\text{id} + \iota_{\mathbb{A}}) \\ \text{---} \\ \text{---} \end{array} & \begin{array}{c} \text{---} \\ \text{---} \\ \downarrow Q(\text{id} + \iota_{\mathbb{B}}) \end{array} \\ Q\mathbb{A}(K) & \xrightarrow{v+h_m} & Q\mathbb{B}(K) \end{array}$$

The map $\iota_{\mathbb{A}}: A_s \rightarrow A_{-s}$ is given by $\iota_{\mathbb{A}} = U^s \iota_K$ and the map $\iota_{\mathbb{B}}: B_s \rightarrow B_{-s+m}$ is given by $\mathfrak{F}_m U^s \iota_K$. The map H_m is defined by picking a map $H_m: \tilde{B}_s(K) \rightarrow B_{-s}(K)$ satisfying

$$U^s \iota_K + \mathfrak{F}_m U^s \iota_K \mathfrak{F}_m = [\partial, H_m]. \quad (16.1)$$

Indeed, any two choices of map H_m satisfying this equation are themselves related by appropriate homotopies. See [4, Section 3.5].

With the above notation in place, we now state our formula for the 2-handle map. Once again, because one end of the cobordism is the three-sphere, the formula is not affected by choices of framings, and we omit them from the notation.

Theorem 16.1. *Let \mathfrak{s} be the unique self-conjugate Spin^c structure on the 2-handle cobordism $W'_{2n}(K)$ from $S^3_{2n}(K)$ to S^3 . Write $BI_n(K) := \text{Cone}(Q(\text{id} + \iota_{\mathbb{B}})) : B_n(K) \rightarrow QB_n(K)$, and let*

$$J : \mathbb{X}\mathbb{I}_n(K) \rightarrow BI_n(K)$$

be the map $J := v_n \Pi_n^A + Q \Pi_n^B \iota_{\mathbb{X}}$, where Π_n^A is projection onto $A_s \oplus QA_s$ and Π_n^B is similar. Then, there is a hypercube of $\mathbb{F}[U, Q]/Q^2$ -equivariant maps

$$\begin{array}{ccccc} CFI^-(S^3_{2n}(K)) & \longrightarrow & CFI(W'_{2n}(K), \mathfrak{s}) & \longrightarrow & CFI^-(S^3) \\ \downarrow \cong & \searrow & & & \downarrow \cong \\ \mathbb{X}\mathbb{I}_{2n}(K) & \xrightarrow{\quad J \quad} & BI_n(K) & & \end{array}$$

Remark 16.2. In Theorem 16.1, $\iota_{\mathbb{X}}$ denotes the involution on the mapping cone, i.e., $\iota_{\mathbb{X}} = \iota_{\mathbb{A}} + \iota_{\mathbb{B}} + H_{2n}\tilde{v}$. Note $\iota_{\mathbb{A}}$ vanishes when composed with Π_n^B , and hence makes no contribution to the map J .

Example 16.3. Consider the case $K = U$ and $n = 0$, so that $W'_{2n}(K)$ is the cobordism W_2 from $S^1 \times S^2$ to S^3 considered in Section 14. We recall that $\mathbb{X}_0(K)$ may be taken to be the cone of $A_0(K) \xrightarrow{v+h_0} B_0(K)$ [27, Section 4.8]. In [4, Section 22.9] we show that $CFI^-(S^3_0(K)) \simeq \mathbb{X}\mathbb{I}_0(K)$, to wit the complex

$$\begin{array}{ccc} A_0(K) & \xrightarrow{v+h_0} & B_0(K) \\ \downarrow Q(\text{id} + \iota_{\mathbb{A}}) & \searrow H_0\tilde{v} & \downarrow Q(\text{id} + \iota_{\mathbb{B}}) \\ QA_0(K) & \xrightarrow{v+h_m} & QB_0(K). \end{array}$$

With this in mind, $BI_0(K) = B_0(K) \oplus QB_0(K)$ with appropriate gradings and differentials. We take $A_0(U)$ and $B_0(U)$ to be appropriately graded copies of $\mathbb{F}[U]$. For clarity, let a be the generator of $A_0(U)$ and b be the generator of $B_0(U)$. Choosing $H_0 \equiv 0$ satisfies the formula 16.1. With this in mind, the map J sends $J(a) = b$ and $J(b) = Qb$, and correspondingly, $J(Qa) = Qb$ and $J(Qb) = 0$. This differs from the map associated to the cobordism computed in Section 14 by a change of basis; setting $\Theta^+ = a$ and $\Theta^- = a + Qb$ recovers the map as described there.

We will focus on proving the case that $n \neq 0$. The case that $n = 0$ is similar.

We begin by considering the non-involutive analog of the above theorem, which is similar to [27, Theorem 1.1]. Ozsváth, Szabó constructed a map from $\mathbf{CF}^-(S_{2n}^3(K))$ to $\text{Cone}(\mathbf{CF}^-(S_{2n+m}^3(K)) \rightarrow \underline{\mathbf{CF}}^-(S^3))$, which we think of as a hypercube

$$\begin{array}{ccc} \mathbf{CF}^-(S_{2n}^3(K)) & & \\ \downarrow & \searrow^k & \\ \mathbf{CF}^-(S_{2n+m}^3(K)) & \longrightarrow & \underline{\mathbf{CF}}^-(S^3). \end{array} \quad (16.2)$$

In the above, $\underline{\mathbf{CF}}^-(S^3)$ denotes a version of twisted Floer homology which is isomorphic to $\mathbf{CF}^-(S^3) \otimes \mathbb{F}[T]/(T^m - 1)$.

We may consider instead a version of the above hypercube with twisted coefficients lying in $\mathbb{F}[T]$. We will write $\tilde{\mathbf{CF}}^-(S^3)$ for this version. We can also construct an analog of equation (16.2) over $\mathbb{F}[T]$, though the resulting hypercube has a different shape:

$$\begin{array}{ccc} \mathbf{CF}^-(S_{2n}^3(K)) & \xrightarrow{F} & \tilde{\mathbf{CF}}^-(S^3) \\ \downarrow & \searrow^k & \downarrow M \\ \mathbf{CF}^-(S_{2n+m}^3(K)) & \xrightarrow{G} & \tilde{\mathbf{CF}}^-(S^3). \end{array} \quad (16.3)$$

In equation (16.3), F is denoted as

$$F = \sum_{\mathfrak{s} \in \text{Spin}^c(W'_{2n}(K))} T^{A_{W'_{2n}, \Sigma}(\mathfrak{s})} \mathbf{CF}(W'_{2n}(K), \mathfrak{s}),$$

where

$$A_{W'_{2n}, \Sigma}(\mathfrak{s}) := \frac{\langle c_1(\mathfrak{s}), \widehat{\Sigma} \rangle + 2n}{2}.$$

Additionally,

$$M = \sum_{s \in 2\mathbb{Z}+1} U^{m(s^2-1)/8} T^{m(s+1)/2}.$$

Note that $M = 0$ if we set $T^m = 1$. Hence, the hypercube in equation (16.2) is obtained from equation (16.3) if we set $T^m = 1$ and remove the top right corner of the hypercube.

The formula for F is derived as follows. Ozsváth and Szabó proved (16.2) by considering a Heegaard quadruple $(\Sigma, \alpha_3, \alpha_2, \alpha_1, \beta, w, z)$, with a special genus 1 surgery region, as in [27, Section 3]. In the diagrams $(\Sigma, \alpha_1, \beta)$ and $(\Sigma, \alpha_2, \beta)$, the points w and z are immediately adjacent. The diagram $(\Sigma, \alpha_1, \beta, w)$ represents $S_{2n}^3(K)$. The diagram $(\Sigma, \alpha_2, \beta, w)$ represents $S_{2n+m}^3(K)$. The diagram $(\Sigma, \alpha_3, \beta, w, z)$ represents (S^3, K) . One obtains the hypercube in equation (16.3) by considering the degenerations of holomorphic quadrilaterals. The quantity $U^{m(s^2-1)/8} T^{m(s+1)/2}$ arises from

a model count of holomorphic triangles on the diagram $(\Sigma, \alpha_3, \alpha_2, \alpha_1, w, z)$ (see [27, Section 3]).

We now fix $\delta \gg 0$, and we set $U^\delta = 0$. Write CF^δ for CF^-/U^δ , and make similar notation for maps. We assume that $m \gg 0$, so the only terms of F with U -power less than δ have $s \in \{1, -1\}$. Hence,

$$F^\delta = (1 + T^m) \sum_{\mathfrak{s} \in \text{Spin}^c(W'_{2n}(K))} T^{A_{W'_{2n}, \Sigma}(\mathfrak{s})} CF^\delta(W'_{2n}(K), \mathfrak{s}).$$

On $D(-m, 1)$, Spin^c structures may be identified with $2\mathbb{Z} + 1$, where $s \in 2\mathbb{Z} + 1$ corresponds to the Spin^c structure \mathfrak{t} satisfying

$$\langle c_1(\mathfrak{t}), S^2 \rangle = s \cdot m.$$

We are most interested in the Spin^c structures corresponding to $s = \pm 1$, which are the Spin^c structures with maximal square. Note that the composition of the natural cobordisms from $S^3_{2n}(K)$ to $S^3_{2n+m}(K)$ (which has an extra boundary component $L(m, 1)$) and the cobordism from $S^3_{2n+m}(K)$ to S^3 is $W'_{2n}(K) \# D(-m, 1)$. We identify Spin^c structures on $W'_{2n}(K) \# D(-m, 1)$ as pairs $\mathfrak{s} \#[s]$ where $\mathfrak{s} \in \text{Spin}^c(W'_{2n}(K))$ and $s \in 2\mathbb{Z} + 1$.

We recall the notation of Ozsváth and Szabó for Spin^c structures on $W'_{2n}(K)$. They write \mathfrak{x}_s and \mathfrak{y}_s for the Spin^c structures which satisfy

$$\langle c_1(\mathfrak{x}_s), \Sigma \rangle + 2n = 2s \quad \text{and} \quad \langle c_1(\mathfrak{y}_s), \Sigma \rangle - 2n = 2s,$$

respectively.

There are also analogous Spin^c structures on $W'_{2n+m}(K)$, for which we write \mathfrak{X}_s and \mathfrak{Y}_s . The Spin^c structure \mathfrak{X}_s satisfies

$$\langle c_1(\mathfrak{X}_s), \Sigma \rangle + 2n + m = 2s,$$

and \mathfrak{Y}_s is similar. We write

$$G_{\mathfrak{X}_s}, G_{\mathfrak{Y}_s}: \mathbf{CF}^-(S^3_{2n+m}(K)) \rightarrow \mathbf{CF}^-(S^3)$$

for the corresponding cobordism maps.

If $s \in \mathbb{Z}$, we may define a projection map $\Pi_s: \tilde{CF}^\delta(S^3) \rightarrow CF^\delta(S^3)$ by reading off only the component T^s . Using this map, we may form a hypercube

$$\begin{array}{ccccc} CF^\delta(S^3_{2n}(K), [s]) & \xrightarrow{F} & \tilde{CF}^\delta(S^3) & \xrightarrow{\Pi_s M} & CF^\delta(S^3) \\ \downarrow & \searrow k & \downarrow M & & \downarrow \text{id} \\ CF^\delta(S^3_{2n+m}(K)) & \xrightarrow{G} & \tilde{CF}^\delta(S^3) & \xrightarrow{\Pi_s} & CF^\delta(S^3). \end{array}$$

If we compress the above hyperbox, we obtain the following diagram:

$$\begin{array}{ccc}
 CF^\delta(S_{2n}^3(K), [s]) & \xrightarrow{CF(W'_{2n}, \mathfrak{x}_s)} & CF^\delta(S^3) \\
 \downarrow & \searrow \Pi_s k & \downarrow \text{id} \\
 CF^\delta(S_{2n+m}^3(K)) & \xrightarrow{G_{\mathfrak{x}_s}} & CF^\delta(S^3)
 \end{array}$$

If we instead use Π_{s+2n+m} , we obtain a hypercube relating the maps $CF(W'_{2n}, \mathfrak{y}_s)$ and $G_{\mathfrak{y}_s}$.

The next step of the knot surgery formula is to construct a hypercube of the following form:

$$\begin{array}{ccc}
 CF^\delta(S_{2n+m}^3(K)) & \xrightarrow{G_{\mathfrak{x}_s} + G_{\mathfrak{y}_s}} & \underline{CF^\delta(S^3)} \\
 \downarrow \Gamma & \searrow j & \downarrow \theta_w \\
 \mathbb{A}^\delta(K) & \xrightarrow{v + h_{2n}} & \mathbb{B}^\delta(K)
 \end{array} \tag{16.4}$$

The map Γ is a holomorphic triangle counting map. The map θ_w is the map for trivializing the twisted coefficients defining $\underline{CF^\delta(S^3)}$.

The above hypercube is built by gluing together hypercubes with the following shape:

$$\begin{array}{ccc}
 CF^\delta(S_{2n+m}^3(K)) & \xrightarrow{G_{\mathfrak{x}_s}} & \underline{CF^\delta(S^3)} \\
 \downarrow \Gamma & \downarrow \theta_w & \text{and} \\
 \mathbb{A}^\delta(K) & \xrightarrow{v} & \mathbb{B}^\delta(K)
 \end{array} \quad
 \begin{array}{ccc}
 CF^\delta(S_{2n+m}^3(K)) & \xrightarrow{G_{\mathfrak{y}_s}} & \underline{CF^\delta(S^3)} \\
 \downarrow \Gamma & \searrow j & \downarrow \theta_w \\
 \mathbb{A}^\delta(K) & \xrightarrow{h_{2n}} & \mathbb{B}^\delta(K)
 \end{array}$$

We may combine the above hypercubes into a hyperbox of the following shape:

$$\begin{array}{ccccc}
 CF^\delta(S_{2n}^3(K), [s]) & \xrightarrow{\quad CF(W'_{2n}, \mathfrak{x}_s) \quad} & CF^\delta(S^3) & & \\
 \searrow \downarrow \Gamma & \searrow k & \searrow \downarrow \text{id} & & \\
 (CF^\delta(S_{2n+m}^3(K)) \cdot G \cdot \underline{CF^\delta(S^3)}) & \xrightarrow{\quad G_{\mathfrak{x}_s} \Pi_{[s]} \quad} & CF^\delta(S^3) & & \\
 \downarrow \Gamma & \searrow j & \downarrow \theta_w & \searrow \downarrow \theta_w & \\
 (\mathbb{A}^\delta(K) \xrightarrow{v + h_{2n}} \mathbb{B}^\delta(K)) & \xrightarrow{v \Pi_s^A} & B_n^\delta(K) & &
 \end{array} \tag{16.5}$$

In the above, $\Pi_{[s]}$ denotes projection of $CF^\delta(S_{2n+m}^3(K))$ onto the Spin^c structure identified with $[s] \in \mathbb{Z}/(2n+m)$. Also, we view each parenthesized group as being a single point in the set $\mathbb{E}(1, 2) \cong \{0, 1\} \times \{0, 1, 2\}$.

Symmetrically, there is another hypercube of the following form:

$$\begin{array}{ccccc}
 CF^\delta(S_{2n}^3(K), [s]) & \xrightarrow{\quad} & CF(W'_{2n}, \mathfrak{v}_s) & \xrightarrow{\quad} & CF^\delta(S^3) \\
 \downarrow CF(W_{2n,m}) & \searrow k & \downarrow k_{\mathfrak{v}_s \# [1]} & \searrow & \downarrow \text{id} \\
 (CF^\delta(S_{2n+m}^3(K)) \cdot G \cdot CF^\delta(S^3)) & \xrightarrow{\quad} & & \xrightarrow{\quad} & CF^\delta(S^3) \\
 \downarrow \Gamma & \searrow j & \downarrow \theta_w & \searrow G_{\mathfrak{v}_s \# [s]} \Pi_{s+2n+m} j & \downarrow \theta_w \\
 (\mathbb{A}^\delta(K) \xrightarrow{-v+h_{2n}} \mathbb{B}^\delta(K)) & \xrightarrow{\quad} & & \xrightarrow{\quad} & B_n^\delta(K) \\
 \downarrow h_{2n} \Pi_s^A & & & &
 \end{array} \tag{16.6}$$

Equations (16.5) and (16.6) can be used to compute the cobordism maps for $W'_{2n}(K)$, very similarly to Ozsváth and Szabó's [27, Theorem 1.1].

In particular, the fundamental objects to construct are the hypercubes in equations (16.2) and (16.3). The involutive extension of equation (16.2) is one of the main theorems of [4]. We now describe how to extend equation (16.3) to an involutive hypercube.

Proposition 16.4. *There is a hypercube of the form*

$$\begin{array}{ccc}
 CFI^-(S_{2n}^3(K)) & \xrightarrow{F} & \widetilde{CFI}^-(S^3) \\
 \downarrow & \searrow k & \downarrow M \\
 CFI^-(S_{2n+m}^3(K)) & \xrightarrow{G} & \widetilde{CFI}^-(S^3).
 \end{array}$$

Furthermore, G and F are the involutive cobordism maps summed over all Spin^c structures (as defined in this paper) and M is the map

$$\sum_{s \in 2\mathbb{Z}+1} U^{m(s^2-1)/8} T^{m(s+1)/2} \cdot \text{id} + EQ.$$

The map E decomposes into a sum

$$E = \sum_{s \in 2\mathbb{Z}+1} E_s,$$

where each E_s is of homogeneous grading $-m(s^2 - 1)/4 + 1$.

Proof. The proof is mostly by a straightforward reinterpretation of the techniques of [4]. There is one point where our description in [4] is not convenient for our present purpose, which is in the first subcube of the central hypercube. We recall the Heegaard quintuple $(\gamma'', \gamma', \gamma, \delta, \delta')$ used to define the first central hypercube. There is a natural hypercube, obtained by pairing two hypercubes (in [4], we referred to this hypercube

as the first *auxiliary* hypercube). Over $\mathbb{F}[T]$, this hypercube takes the following form:

$$\begin{array}{ccccc}
 \mathbf{CF}^-(\gamma, \delta) & \xrightarrow{f_{\gamma \rightarrow \gamma''}^\delta} & \tilde{\mathbf{CF}}^-(\gamma'', \delta) & & \\
 \downarrow & \searrow f_{\gamma \rightarrow \gamma'}^\delta & \downarrow h_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^\delta & \searrow M & \downarrow \\
 & \mathbf{CF}^-(\gamma', \delta) & \xrightarrow{f_{\gamma' \rightarrow \gamma''}^\delta} & \tilde{\mathbf{CF}}^-(\gamma'', \delta) & \\
 \downarrow f_{\gamma}^{\delta \rightarrow \delta'} & \downarrow h_{\gamma \rightarrow \gamma'}^{\delta \rightarrow \delta'} & \downarrow f_{\gamma''}^{\delta \rightarrow \delta'} & & \downarrow f_{\gamma''}^{\delta \rightarrow \delta'} \\
 & \mathbf{CF}^-(\gamma, \delta') & \xrightarrow{f_{\gamma \rightarrow \gamma''}^{\delta'}} & \tilde{\mathbf{CF}}^-(\gamma'', \delta') & \\
 \downarrow & \searrow f_{\gamma \rightarrow \gamma'}^{\delta'} & \downarrow h_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^{\delta'} & \searrow M & \downarrow \\
 & \mathbf{CF}^-(\gamma', \delta') & \xrightarrow{f_{\gamma' \rightarrow \gamma''}^{\delta'}} & \tilde{\mathbf{CF}}^-(\gamma'', \delta') & \\
 \downarrow & \searrow f_{\gamma \rightarrow \gamma'}^{\delta'} & \downarrow h_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^{\delta'} & \searrow M & \downarrow \\
 & \mathbf{CF}^-(\gamma', \delta') & \xrightarrow{f_{\gamma' \rightarrow \gamma''}^{\delta'}} & \tilde{\mathbf{CF}}^-(\gamma'', \delta') &
 \end{array}$$

The length 3 map (not shown) is the pentagon counting map $p_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^{\delta \rightarrow \delta'}$.

We recall that (γ', δ'') admits a 3-handle map to a diagram \mathcal{H}_0 , and that both (γ, δ') and (γ', δ') are genus 1 stabilizations of \mathcal{H}_0 . Furthermore, $f_{\gamma \rightarrow \gamma''}^\delta$ and $f_{\gamma' \rightarrow \gamma''}^\delta$ are the maps for surgering on a 0-framed unknot. Since the relevant Heegaard triples are both stabilizations, both of these maps admit left inverses, which are a 3-handle map followed by a stabilization. Write σ for the stabilization map which has (γ, δ') as its codomain, and write σ' for the stabilization map which has (γ', δ') as its codomain.

As a replacement for the diagram $\mathcal{C}_{\text{cen}}^{(1)}$ in [4], we build the following diagram:

$$\begin{array}{ccccc}
 \mathbf{CF}^-(\gamma, \delta) & \xrightarrow{f_{\gamma \rightarrow \gamma''}^\delta} & \tilde{\mathbf{CF}}^-(\gamma'', \delta) & & \\
 \downarrow \text{id} & \searrow f_{\gamma \rightarrow \gamma'}^\delta & \downarrow h_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^\delta & \searrow M & \downarrow \\
 & \mathbf{CF}^-(\gamma', \delta) & \xrightarrow{f_{\gamma' \rightarrow \gamma''}^\delta} & \tilde{\mathbf{CF}}^-(\gamma'', \delta) & \\
 \downarrow \sigma F_3 h_{\gamma \rightarrow \gamma''}^{\delta \rightarrow \delta'} & \downarrow \sigma F_3 f_{\gamma''}^{\delta \rightarrow \delta'} & \downarrow \sigma' F_3 h_{\gamma' \rightarrow \gamma''}^{\delta \rightarrow \delta'} & \downarrow \sigma' F_3 h_1 \sigma F_3 f_{\gamma''}^{\delta \rightarrow \delta'} & \downarrow \sigma' F_3 f_{\gamma''}^{\delta \rightarrow \delta'} \\
 & \mathbf{CF}^-(\gamma, \delta) & \xrightarrow{f_{\gamma \rightarrow \gamma'}^{\delta \rightarrow \delta'}} & \tilde{\mathbf{CF}}^-(\gamma', \delta') & \\
 \downarrow & \searrow f_{\gamma \rightarrow \gamma'}^{\delta \rightarrow \delta'} & \downarrow h_{\gamma \rightarrow \gamma'}^{\delta \rightarrow \delta'} & \searrow f_{\gamma' \rightarrow \gamma'}^{\delta \rightarrow \delta'} & \downarrow \\
 & \mathbf{CF}^-(\gamma', \delta) & \xrightarrow{f_{\gamma' \rightarrow \gamma''}^{\delta'}} & \tilde{\mathbf{CF}}^-(\gamma'', \delta') &
 \end{array}$$

The length 3 map (not shown) is

$$\sigma' F_3 p_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^{\delta \rightarrow \delta'} + \sigma' F_3 h_1 \sigma F_3 h_{\gamma \rightarrow \gamma''}^{\delta \rightarrow \delta'}.$$

Following the notation of [4], the map h_1 denotes $h_{\gamma \rightarrow \gamma' \rightarrow \gamma''}^{\delta'}$. The hypercube relations are straightforward to verify. Compare [4, Lemma 16.4].

The hypercubes $\mathcal{C}_{\text{cen}}^{(2)}$ and $\mathcal{C}_{\text{cen}}^{(3)}$ are more straightforward to modify, and we leave the details to the reader. Upon compressing and stacking these hypercubes, we get the diagram in the statement. It remains to verify the listed claims. The identification of J and G as cobordism maps is straightforward, and we leave the details to the reader. Compare [4, Lemma 16.14].

We now consider the map labeled M in the statement, the length one component of this map (i.e., the non- Q component) may be identified with

$$\sum_{s \in 2\mathbb{Z} + 1} U^{m(s^2-1)/8} T^{m(s+1)/2}$$

by the lattice point counting argument in [27, Section 3]. It remains to identify the Q -term. The Q -term is essentially the same as the map E considered in [4, Section 19]. Therein, we showed that E is null-homotopic over $\mathbb{F}[U, T]/(T^m - 1)$. Our argument used at several places that the coefficients were in $\mathbb{F}[U, T]/(T^m - 1)$ instead of $\mathbb{F}[U, T]$, so we give an alternate argument that is sufficient for our present purposes.

The first observation is that the map E decomposes over $\text{Spin}^c(D(-m, 1))$, which we have already identified with $2\mathbb{Z} + 1$. The grading change formula is also a routine consequence of this identification, and the standard grading change formulas for the cobordism maps in Heegaard Floer homology [20]. ■

Remark 16.5. We write $M = M_0 + QM_1$. Since $CF^-(S^3) \simeq \mathbb{F}[U]$ is supported only in even gradings, and M_1 may be viewed as a chain map from $CF^-(S^3)$ to itself which shifts the mod 2 grading by 1, we may conclude that $M_1 \simeq 0$. In particular, by adding a term to the length 3 map of the hypercube, we may assume that $M_1 = 0$. Furthermore, fixing δ and letting m be sufficiently large, we see that E^δ is homogeneously graded, so the null-homotopy may also be taken to be homogeneously graded.

In [4, Section 21], we also described how to extend the hypercube in (16.4) into the involutive setting. Using this, together with Proposition 16.4, we may construct the involutive analog of the hyperboxes in equations (16.5) and (16.6). We construct this hyperbox so that the top face has the hyperbox from equation (16.5), which involves the Spin^c structures \mathfrak{X}_s and \mathfrak{X}_s with $s = n$. There is an important subtlety in that the bottom face of this hyperbox will be the hypercube from equation (16.6), involving the Spin^c structures \mathfrak{Y}_{-n} and \mathfrak{Y}_{-n} . (Note that $\mathfrak{Y}_{-n} = \tilde{\mathfrak{X}}_n$.)

The front face of this hyperbox will compress to have the diagonal map $H_{2n}\tilde{v}$, where H_{2n} is the length 2 map from the involutive mapping cone of [4].

We add one more level to this hyperbox so that the top and bottom levels faces coincide. We add this level to the bottom of the hyperbox. We build this new level as a pair of hypercubes. The first hypercube has the following shape:

$$\begin{array}{ccccc}
 CF^\delta(S_{2n}^3(K), [n]) & \xrightarrow{CF(W'_{2n}, \mathfrak{y}_{-n})} & CF^\delta(S^3) & & \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} \\
 & \xrightarrow{CF(W_{2n, 2n+m})} & k & \xrightarrow{k_{\mathfrak{y}_{-n} \# [1]}} & \\
 & & \downarrow & & \\
 & (CF^\delta(S_{2n+m}(K)) \rightarrow \underline{CF^\delta(S^3)}) & & & \\
 & \downarrow \text{id} & & \downarrow \text{id} & \\
 & CF^\delta(S_{2n}^3(K), [n]) & \xrightarrow{CF(W'_{2n}, x_n)} & CF^\delta(S^3) & \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} \\
 & \xrightarrow{CF(W_{2n, 2n+m})} & k & \xrightarrow{k_{x_n \# [-1]}} & \\
 & & \downarrow & & \\
 & (CF^\delta(S_{2n+m}^3(K)) \rightarrow \underline{CF^\delta(S^3)}) & & & \\
 & & \downarrow G_{x_n} \Pi_{[n]} & &
 \end{array}$$

In the above diagram, Π_n denotes projection onto the B_n -summand of $\underline{CF^\delta}(S^3)$. The hypercube relations are straightforward.

We finally have one additional hypercube, as follows:

$$\begin{array}{ccccc}
 (CF^\delta(S_{2n+m}^3(K)) \rightarrow \underline{CF^\delta(S^3)}) & \xrightarrow{G_{\mathfrak{y}_{-n}} \Pi_{[-n]}} & CF^\delta(S^3) & & \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} \\
 & \xrightarrow{\Gamma} & j & \xrightarrow{\theta_w} & \\
 & & \downarrow \text{id} & & \\
 & (\mathbb{A}^\delta(K) \rightarrow \mathbb{B}^\delta(K)) & \xrightarrow{\Pi_n^B j} & & \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} \\
 & (CF^\delta(S_{2n+m}^3(K)) \rightarrow \underline{CF^\delta(S^3)}) & \xrightarrow{G_{x_n} \Pi_{[n]}} & CF^\delta(S^3) & \\
 \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} & \searrow \text{id} & \downarrow \text{id} \\
 & \xrightarrow{\Gamma} & j & \xrightarrow{\theta_w} & \\
 & & \downarrow \text{id} & & \\
 & (\mathbb{A}^\delta(K) \rightarrow \mathbb{B}^\delta(K)) & \xrightarrow{v\Pi_n^A} & & B_n^\delta(K)
 \end{array}$$

The above shape can be easily checked to be a hypercube of chain complexes, so the proof is complete.

17. The $SO(3)$ -twist

In this section, we show that the diffeomorphism map associated to $\text{tw} \in \pi_1(SO(3)) \cong \mathbb{Z}/2$ acts on $CFI(Y)$ by $\text{Id} + Q\Phi$, proving Theorem 1.7. Our argument will actually identify the map with $Q\Phi\iota$, though this is homotopic to $Q\Phi$ on $CFI(Y)$. In Section 17.7, we will further consider an analogous twisting diffeomorphism map on involutive knot Floer homology, and show that it is trivial.

Before beginning with the computation, we recall several topological perspectives about tw . The first is as a Dehn twist along a small S^2 enclosing the basepoint. A regular neighborhood of such a sphere is $S^2 \times I$. We recall that

$$\text{Diff}(S^2) \simeq O(3),$$

so the generator of $\pi_1(SO(3))$ gives a diffeomorphism of $S^2 \times I$ which is the identity on the boundary.

There is another perspective on tw , as follows. Write $\text{Diff}^+(Y, w)$ for the orientation preserving diffeomorphisms ϕ such that $\phi(w) = w$; likewise, write $\text{Diff}^f(Y, w)$ for the orientation preserving diffeomorphisms ϕ such that $\phi(w) = w$ and $d_w\phi = \text{id}$. If we fix a trivialization of $T_w Y$, we obtain a Serre fibration

$$\text{Diff}^f(Y, w) \rightarrow \text{Diff}^+(Y, w) \rightarrow SO(3),$$

and hence, a map

$$\pi_1(SO(3)) \rightarrow \pi_0(\text{Diff}^f(Y, w)).$$

17.1. Heegaard diagrams for the twist map

We first describe the Heegaard diagrams used to define tw_* . Recall that in non-involutive Heegaard Floer homology, if $\phi: \Sigma \rightarrow \Sigma$ is a diffeomorphism, then the action of ϕ_* on Heegaard Floer homology is induced by composing the maps in the following diagram:

$$CF(\Sigma, \alpha, \beta) \xrightarrow{f_{\alpha \rightarrow \phi(\alpha)}^{\beta \rightarrow \phi(\beta)}} CF(\Sigma, \phi^{-1}(\alpha), \phi^{-1}(\beta)) \xrightarrow{T_\phi} CF(\Sigma, \alpha, \beta),$$

where $f_{\alpha \rightarrow \phi(\alpha)}^{\beta \rightarrow \phi(\beta)}$ is a composition of holomorphic triangle maps and change of almost complex structure maps, and T_ϕ is the tautological map on intersection points. The diffeomorphism maps on involutive Heegaard Floer homology are defined similarly, except that we replace $CF(\Sigma, \alpha, \beta)$ with a one-dimensional hyperbox (whose maps of non-zero length correspond to the involution), and we replace $f_{\alpha \rightarrow \phi(\alpha)}^{\beta \rightarrow \phi(\beta)}$ and T_ϕ with morphisms of hyperboxes.

This has the following instantiation in our present setting. Let δ be a set of doubling arcs (with boundary on w) and $\tilde{\delta}$ a set obtained by applying a boundary Dehn

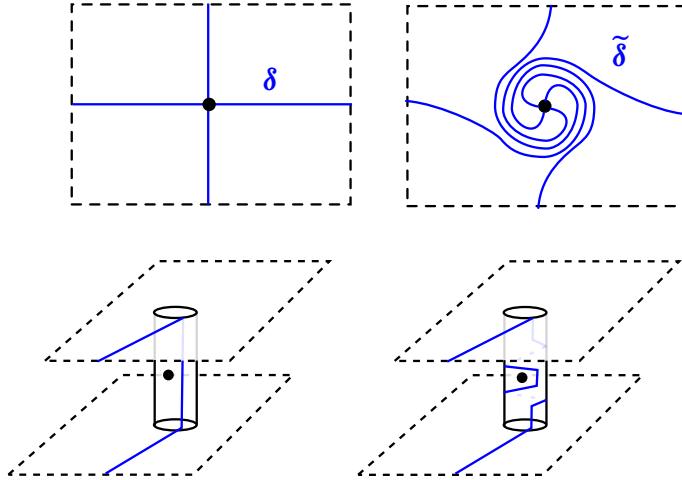


Figure 17.1. Top row: the doubling arcs δ (left) and $\tilde{\delta}$ (right). On the bottom row, we show the doubled diagrams and indicate how $\tilde{\Delta}$ is obtained from Δ by moving the basepoint in a loop around the connected sum tube.

twist in a loop which encircles w . Let Δ and $\tilde{\Delta}$ be their corresponding doubled curves on $\Sigma \# \bar{\Sigma}$. Then, tw_* is obtained by compressing the following diagram (where we are omitting an extra column corresponding to the tautological portion of the diffeomorphism map):

$$\begin{array}{ccc}
 CF(\Sigma, \alpha, \beta) & \longrightarrow & CF(\Sigma, \alpha, \beta) \\
 \downarrow & & \downarrow \\
 CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \beta \bar{\beta}) & \longrightarrow & CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \beta \bar{\beta}) \\
 \downarrow & \searrow & \downarrow \\
 CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \Delta) & \longrightarrow & CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \tilde{\Delta}) \\
 \downarrow & \searrow & \downarrow \\
 CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \alpha \bar{\alpha}) & \longrightarrow & CF(\Sigma \# \bar{\Sigma}, \alpha \bar{\beta}, \alpha \bar{\alpha}) \\
 \downarrow & & \downarrow \\
 CF(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}) & \longrightarrow & CF(\bar{\Sigma}, \bar{\beta}, \bar{\alpha})
 \end{array} \tag{17.1}$$

The relevant Heegaard diagrams are shown in Figure 17.1.

Remark 17.1. Since tw_* can be interpreted in terms of moving the basepoint around the connected sum neck on $\Sigma \# \bar{\Sigma}$, one might expect that the formula for the hypercube

π_1 -action from [38, Section 13] would show that tw_* is null-homotopic as a morphism of 1-dimensional hypercubes. That is, if we view γ_* as a 2-dimensional diagram below, then there is a 3-dimensional hypercube taking the form shown on the right:

We observe however that the techniques of [38] do not guarantee the 3-dimensional hypercube so constructed has the same maps on the top and bottom faces of the hypercube. (This may seem surprising since the length one components of γ_* are the identity on the nose.) Concretely, the homotopy on the top face involves maps related to relative homology actions on stabilized diagrams, which are in general not preserved by the involution. In particular, such a hypercube does not give us an $\mathbb{F}[U, Q]/Q^2$ equivariant homotopy $\text{id} \cong \text{tw}_*$ on $CFI(Y)$.

17.2. Degenerating connected sums

Our argument depends on a somewhat complicated neck stretching argument. Our strategy is adapted from work of Ozsváth and Szabó in the setting of bordered knot Floer homology [26]. We have explored some of these ideas further in [4, 38].

We focus first on the case of two ordinary Heegaard diagrams

$$\mathcal{H} = (\Sigma, \alpha, \beta, w) \quad \text{and} \quad \mathcal{H}' = (\Sigma', \alpha', \beta', w'),$$

and we take the connected sum at w and w' (deleting one of the basepoints).

We first degenerate the connected sum tubes by letting neck length approach ∞ . We obtain maps which count *perfectly matched* moduli spaces. We write $CF_{\wedge}(\mathcal{H} \# \mathcal{H}')$ for the resulting chain complex.

There are homotopy equivalences

$$\Psi^{\Sigma < \Sigma'}, \Psi^{\Sigma' < \Sigma} : CF_{\wedge}(\mathcal{H} \# \mathcal{H}') \rightarrow CF(\mathcal{H}) \otimes_{\mathbb{F}[U]} CF(\mathcal{H}').$$

The map $\Psi^{\Sigma < \Sigma'}$ counts holomorphic disk pairs (u, u') of expected dimension 0 such that the following hold.

- (1) For $[0, 1] \times (-\infty, 0)$, the punctures of u and u' consist of perfectly matched pairs; that is, each puncture of u is paired with a puncture of u' which has identical projection to $[0, 1] \times \mathbb{R}$ (and vice-versa).
- (2) There is an even number of u and u' punctures on $[0, 1] \times \{0\}$, the *special line*. These punctures all have distinct projections to $[0, 1] \times \{0\}$, and alternate between u and u' . The left-most puncture is from u .
- (3) In $[0, 1] \times (0, \infty)$, u and u' may have any number of punctures, and there is no constraint.

The map $\Psi^{\Sigma' < \Sigma}$ is defined similarly, except that the left-most puncture on the special line is from u' .

Proposition 17.2. *The maps $\Psi^{\Sigma < \Sigma'}$ and $\Psi^{\Sigma' < \Sigma}$ are chain maps. Furthermore, they are homotopy equivalences, with homotopy inverses given by the maps which count curves in 0-dimensional moduli spaces with matching conditions in the regions above and below $[0, 1] \times \{0\}$ reversed.*

A proof, modeled on Ozsváth and Szabó's work [26], may be found in [38, Section 11.2]. An extension of the connected sum maps $\Psi^{\Sigma < \Sigma'}$ to the setting of connected sums of hypercubes (under certain topological assumptions) is also proven in [38, Proposition 11.8].

17.3. Setting up the computation

In Section 15.2, we considered several stabilized models of the involution where we used multiple connected sum tubes, but a single basepoint. To perform our computation, we will join Σ and $\bar{\Sigma}$ using a tube at w and a tube at another point p . Considering this stabilized model will allow us to localize the computation.

Write c_p for the meridian of the p -tube. Note that we can get a valid Heegaard diagram by replacing c_p with a meridian of the other tube. There are two natural choices, c_w and \bar{c}_w , depending on which side of w we put the meridian (see Figure 17.2).

We recall that the stabilized model of the involution from Section 15 takes the following form:

$$F_3^{\alpha c_p, \alpha c_p} \circ f_{\alpha c_p \bar{\beta}}^{\Delta \rightarrow (\alpha c_p \bar{\alpha})} \circ f_{\alpha c_p \bar{\beta}}^{(\beta c_p \bar{\beta}) \rightarrow \Delta} \circ F_1^{c_p \bar{\beta}, c_p \bar{\beta}}.$$

We now modify the sequence of curves we use in this Heegaard diagram, by breaking the sequence

$$\beta c_p \bar{\beta} \rightarrow \Delta \rightarrow \alpha c_p \bar{\alpha}$$

into a longer sequence. This will make some holomorphic curve counts easier.

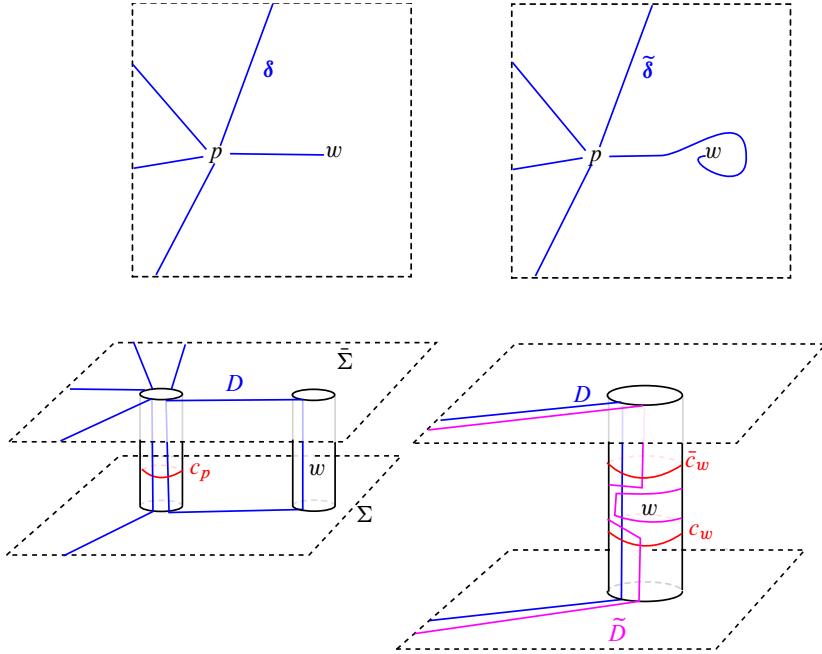


Figure 17.2. A once stabilized model of the involution. On the top, we show the doubling arcs on Σ . On the bottom left, we have an overview of the two tubes. On the bottom right, we zoom in on the w -tube region.

We pick sets of doubled curves $\Delta_1, \dots, \Delta_n$, containing $2g + 1$ curves each, on $\Sigma \# \bar{\Sigma}$ such that the following hold.

- (1) $\Delta_i = \Delta_i^0 \cup D$, where Δ_i^0 are doubled curves on the unstabilized diagram, obtained by doubling arcs with two endpoints at w . Here, D is a single curve which is obtained by doubling an arc which connects p to w .
- (2) $(\Sigma \# \bar{\Sigma}, \beta \bar{\beta}, \Delta_1^0)$ is algebraically rigid. (Compare [4, Definition 6.5].)
- (3) $(\Sigma \# \bar{\Sigma}, \alpha \bar{\alpha}, \Delta_n^0)$ is algebraically rigid.
- (4) Each Δ_{i+1}^0 is obtained from Δ_i^0 by a small isotopy or a simple arc slide of the doubling arcs on $\Sigma \setminus N(p)$, not crossing over the endpoint corresponding to D . (In particular, each $(\Sigma \# \bar{\Sigma}, \Delta_i^0, \Delta_{i+1}^0, w)$ is algebraically rigid.) Note that this implies that Δ_{i+1}^0 is obtained from Δ_i^0 by an elementary handleslide.

We will use the following model of the involution:

$$\begin{aligned}
 & F_3^{\alpha c_p, \alpha c_p} \circ \Psi_{(\alpha \bar{c}_w \bar{\beta}) \rightarrow (\alpha c_p \bar{\alpha})}^{(\alpha \bar{c}_w \bar{\alpha}) \rightarrow (\alpha c_p \bar{\alpha})} \circ f_{\alpha \bar{c}_w \bar{\beta}}^{\Delta_n \rightarrow (\alpha \bar{c}_w \bar{\alpha})} \\
 & \circ \dots \circ f_{\alpha \bar{c}_w \bar{\beta}}^{\Delta_1 \rightarrow \Delta_2} \circ f_{\alpha \bar{c}_w \bar{\beta}}^{(\beta \bar{c}_w \bar{\beta}) \rightarrow \Delta_1} \circ \Psi_{(\alpha c_p \bar{\beta}) \rightarrow (\beta \bar{c}_w \bar{\beta})}^{(\beta c_p \bar{\beta}) \rightarrow (\beta \bar{c}_w \bar{\beta})} \circ F_1^{c_p \bar{\beta}, c_p \bar{\beta}}
 \end{aligned}$$

The motivation for using this model is that most of the holomorphic curve counts in the above diagram occur on a connected sum of the ordinary doubled diagram for Y on $\Sigma \# \bar{\Sigma}$ and a genus 1 diagram where the diffeomorphism acts non-trivially. The genus 1 region is a neighborhood of $D \cup c_w$; we call it \mathbb{T} .

We now write $\tilde{\Delta}_i$ for $\text{tw}(\Delta_i)$. If we write $\Delta_i = \Delta_i^0 \cup D$, then $\tilde{\Delta}_i = \Delta_i^0 \cup \tilde{D}$, where \tilde{D} is obtained from D by applying the twist in the torus region.

Our argument will proceed as follows. We will view each of the holomorphic polygon counts as occurring on a connected sum of $\Sigma \# \bar{\Sigma}$ and \mathbb{T} . We will build a 3-dimensional hyperbox whose front and back faces are expanded versions of the hyperbox in equation (17.1). The new dimension corresponds to the map $\Psi^{\Sigma \# \bar{\Sigma} \times \mathbb{T}}$. The back face will consist of the perfectly matched complexes. The front face will consist of the trivially matched complexes with some extra terms. The extra terms correspond to the $Q\Phi$ summand of tw_* .

17.4. Initial and final hypercubes

In this section, we extend the following hypercubes (with holomorphic curves counted with perfect matching) into a third dimension corresponding to the map $\Psi^{\Sigma \# \bar{\Sigma} \times \mathbb{T}}$:

$$\begin{array}{ccccc}
 & & CF(\alpha \bar{c}_w \bar{\alpha}, \alpha \bar{c}_w \bar{\alpha}) & \xrightarrow{\text{id}} & CF(\alpha \bar{c}_w \bar{\alpha}, \alpha \bar{c}_w \bar{\alpha}) \\
 & & \downarrow & & \downarrow \\
 CF(\alpha, \beta) & \xrightarrow{\text{id}} & CF(\alpha, \beta) & & CF(\alpha c_p \bar{\alpha}, \alpha c_p \bar{\alpha}) \xrightarrow{\text{id}} CF(\alpha c_p \bar{\alpha}, \alpha c_p \bar{\alpha}) \\
 \downarrow F_1^{c_p \bar{\beta}, c_p \bar{\beta}} & & \downarrow F_1^{c_p \bar{\beta}, c_p \bar{\beta}} & & \downarrow F_3^{\alpha c_p, \alpha c_p} \\
 CF(\alpha c_p \bar{\beta}) & \xrightarrow{\text{id}} & CF(\alpha c_p \bar{\beta}) & & CF(\bar{\beta}, \bar{\alpha}) \xrightarrow{\text{id}} CF(\bar{\beta}, \bar{\alpha}).
 \end{array}$$

In the above, there is no diagonal term because the diffeomorphism tw fixes all of the attaching curves on the Heegaard diagram.

Lemma 17.3. *If we stretch the almost complex structure sufficiently on the connected sum tubes, we have the equality*

$$\Psi_{(\alpha c_p \bar{\beta}) \rightarrow (\alpha \bar{c}_w \bar{\beta})}^{(\beta c_p \bar{\beta}) \rightarrow (\beta \bar{c}_w \bar{\beta})} \circ F_1^{c_p \bar{\beta}, c_p \bar{\beta}} = F_1^{\bar{c}_w \bar{\beta}, \bar{c}_w \bar{\beta}}.$$

Proof. This follows from our stabilization result for holomorphic triangles in Proposition 6.5. Topologically, it is sufficient to observe that on $\bar{\Sigma}$, there is a sequence of handleslides of c_p over the other $\bar{\beta}$ curves (not crossing over w) which takes c_p to \bar{c}_w (i.e., handleslide c_p over each $\bar{\beta}$ curve twice). ■

We now investigate the next terms in the sequence. In Figure 17.3, we draw the diagram $(\mathbb{T}, \bar{c}_w, D, \tilde{D}, \bar{c}'_w, w)$.

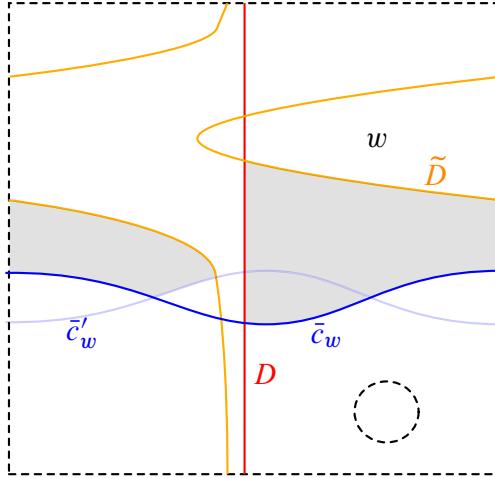


Figure 17.3. The special genus 1 region and the diagram $(\mathbb{T}^2, \bar{c}_w, D, \tilde{D}, \bar{c}'_w, w)$. The dashed circle is the connected sum point. Also shown is a holomorphic quadrilateral that will play a role in our argument.

Next, we observe

$$F_3^{\alpha c_p, \alpha c_p} = F_3^{c_p, c_p} \circ F_3^{\alpha, \alpha}.$$

We note that after surgering out the α curves, the curves c_p become isotopic to c_w . Hence, we identify

$$F_3^{c_p, c_p} \circ F_3^{\alpha, \alpha} = F_3^{c_w, c_w} \circ F_3^{\alpha, \alpha}.$$

In a similar manner to Lemma 17.3, we have the strict commutation

$$F_3^{\alpha, \alpha} \circ \Psi_{(\alpha \bar{c}_w \bar{\alpha}) \rightarrow (\alpha c_p \bar{\alpha})}^{(\alpha \bar{c}_w \bar{\beta}) \rightarrow (\alpha c_p \bar{\beta})} = \Psi_{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})}^{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})} \circ F_3^{\alpha, \alpha}. \quad (17.2)$$

We now prove a simple index bound for the diagram $(\mathbb{T}, \bar{c}_w, c_w)$. We view this diagram as having both a basepoint $w \in \mathbb{T}$, as well as a connected sum point $p \in \mathbb{T}$. Note that these do not coincide. See Figure 17.4.

Lemma 17.4. *Suppose that $\phi \in \pi_2(\theta^+, \theta)$ is a class of disks on $(\mathbb{T}, \bar{c}_w, c_w, w)$.*

- (1) *Then, $\mu(\phi) = 2n_w(\phi) + \text{gr}(\theta^+, \theta)$.*
- (2) *If $D(\phi) \geq 0$, then $\mu(\phi) \geq n_p(\phi)$.*

Proof. The first claim is the Maslov grading formula.

We now consider the second claim. Consider first the case that $\theta = \theta^+$. Then, any class $\phi = e_\theta + k \cdot P + n \cdot [\mathbb{T}]$, where P is the periodic domain class which has $n_w = 0$ and $n_p = 1$. We observe $\mu(\phi) = 2n$, and $n_p(\phi) = n + k$. Also, $n, n - k$ and

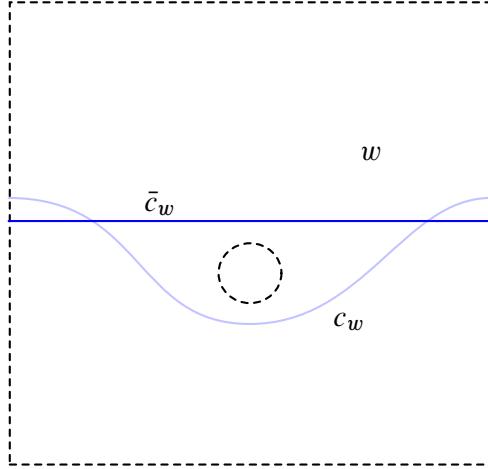


Figure 17.4. The diagram $(\mathbb{T}, \bar{c}_w, c_w, w)$. The dashed circle is the connected sum point p .

$n + k$ are multiplicities of this disk class, so all must be nonnegative. This implies, in particular, that $n \geq |k|$. Ergo,

$$\mu(\phi) = 2n \geq n + k = n_p(\phi),$$

as claimed.

We now consider the case that $\theta = \theta^-$. In this case, write ϕ_0 for the bigon which has $n_p = 1$ and write $\phi = \phi_0 + k \cdot P + n \cdot [\mathbb{T}]$ for a general class. In this case, n , $n - k$ and $n + k + 1$ are multiplicities on the diagram. Hence, $n \geq k$. We have

$$\mu(\phi) = 1 + 2n \geq 1 + n + k = n_p(\phi),$$

completing the proof. ■

We now perform a key computation.

Lemma 17.5. *We have the relation*

$$F_3^{c_w, c_w} \circ \Psi_{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})}^{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})} \simeq \left(\begin{array}{l} x \otimes \theta^- \mapsto x \\ x \otimes \theta^+ \mapsto \Phi_w^{\bar{\Sigma}}(x) \end{array} \right) \circ \Psi_{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}^{\bar{\Sigma} < \mathbb{T}},$$

where $\Phi_w^{\bar{\Sigma}}$ is the standard basepoint action on $CF(\bar{\Sigma}, \bar{\beta}, \bar{\alpha})$, usually denoted by Φ_w .

Proof. The map $\Psi_{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}^{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})}$ is a composition of two triangle maps. Figure 17.4 shows the diagram $(\Sigma, \bar{c}_w, c_w, w)$. The transition map may be viewed as handlesliding \bar{c}_w over the connected sum point.

We first claim that index bound which implies that the 3-handle map is well defined implies also that

$$F_3^{c_w, c_w} = F_3^{c_w, c_w} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}. \quad (17.3)$$

To see this, note that $\Psi^{\bar{\Sigma} < \mathbb{T}}$ counts holomorphic disk-pairs $(u_\Sigma, u_\mathbb{T})$ on

$$(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}) \# (\mathbb{T}, c_w, c_w),$$

which have a special line matching condition. Suppose that u_Σ has S marked points below the special line and M marked points along the special line. The matching condition implies that $u_\mathbb{T}$ has the same number of marked points below and on the special lines, respectively. If $\phi_\mathbb{T} \in \pi_2(\theta, \theta')$ is the class of $u_\mathbb{T}$, we observe that

$$\mu(\phi_\mathbb{T}) = 2n_w(\phi_\mathbb{T}) + \text{gr}(\theta, \theta') \geq 2M + 2S + \text{gr}(\theta, \theta').$$

In the composition $F_3^{c_w, c_w} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}$, only curves with $\theta' = \theta^-$ contribute. Therefore, the above shows that

$$\mu(\phi_\mathbb{T}) \geq 2M + 2S.$$

By assumption, $\Psi^{\bar{\Sigma} < \mathbb{T}}$ counts 0-dimensional moduli spaces, so we must have

$$0 = \mu(\phi_\Sigma) + \mu(\phi_\mathbb{T}) - 2S - 2M \geq \mu(\phi_\Sigma).$$

By transversality, if ϕ_Σ admits a holomorphic representative, then $\mu(\phi_\Sigma) \geq 0$ with equality if and only if ϕ_Σ is a constant class. Hence,

$$F_3^{c_w, c_w} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}$$

counts only curves which represent the constant class on Σ . It is easy to see by considering the diagram (\mathbb{T}, c_w, c_w) that this also implies that these classes are also constant on \mathbb{T} . Hence, equation (17.3) follows.

Therefore, we may replace $F_3^{c_w, c_w}$ in the left-hand side of the equation in the statement with $F_3^{c_w, c_w} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}$. We now consider the 1-parameter moduli spaces which would naturally be used to commute $\Psi^{\bar{\Sigma} < \mathbb{T}}$ past both of the triangle maps used to define $\Psi_{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})}^{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}$. These are the 1-dimensional moduli spaces of holomorphic triangles which have a special line which has an even number of marked points, which are $\bar{\Sigma} < \mathbb{T}$ matched. Here, we view the triangle Δ as $[0, 1] \times \mathbb{R}$ with the puncture $\{1\} \times \{0\}$ removed. The special lines are of the form $[0, 1] \times \{s\}$, $s \in \mathbb{R}$.

In these moduli spaces, there are ends corresponding to the composition

$$\Psi_{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})}^{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})} \circ \Psi^{\bar{\Sigma} < \mathbb{T}},$$

as well as ends corresponding to a chain homotopy. There is also the possibility of an additional end where a holomorphic disk forms at the same height as the special line,

and some the punctures along the special line may also degenerate into a holomorphic disk. We claim that these ends are constrained to a broken curve containing the following curves.

- (u-1) An index one disk on $(\mathbb{T}, \bar{c}_w, c_w)$ which has a single marked point, which is from \mathbb{T} .
- (u-2) A trivial disk on $(\bar{\Sigma}, \bar{\beta}, \bar{\beta})$.
- (u-3) A pair of triangles on $(\mathbb{T}, \bar{c}_w, c_w, c_w) \wedge (\bar{\Sigma}, \bar{\alpha}, \bar{\beta}, \bar{\beta})$ which have marked points at the same height as the $\bar{\beta}$ - $\bar{\beta}$ vertex of the triangle. Furthermore, $\bar{\Sigma}$ and \mathbb{T} marked points alternate along this line, and both the left-most and right-most punctures are from Σ .

See Figure 17.6. We say that triangle pair satisfying (u-3) is $(\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma})$ -matched.

To establish that these are the only ends, we consider a potential degeneration where the special line on the triangle Δ has the same height as the $\bar{\beta}$ - $\bar{\beta}$ puncture of Δ .

We consider a limiting curve, and write $\psi_{\bar{\Sigma}}^l, \psi_{\mathbb{T}}^l$ for the homology classes of the holomorphic triangles, and we write $\phi_{\bar{\Sigma}}^r$ and $\phi_{\mathbb{T}}^r$ for the classes of the holomorphic disks that break off. We suppose that there are $n_{\bar{\Sigma}}^r$ and $n_{\mathbb{T}}^r$ marked points of the disk classes along the special line, and $S^r = S_{\bar{\Sigma}}^r = S_{\mathbb{T}}^r$ marked points below the special line (i.e., perfectly matched). The expected dimension of the disk class with this matching condition is

$$\mu(\phi_{\bar{\Sigma}}^r) + \mu(\phi_{\mathbb{T}}^r) - n_{\bar{\Sigma}}^r - n_{\mathbb{T}}^r + 1 - 2S^r. \quad (17.4)$$

The above quantity must be 1 in a generic degeneration. On the other hand, since $\phi_{\bar{\Sigma}}^r$ is a disk class on $(\bar{\Sigma}, \bar{\beta}, \bar{\beta})$, and the connected sum point is the basepoint, we have

$$\mu(\phi_{\bar{\Sigma}}^r) \geq 2n_p(\phi_{\bar{\Sigma}}^r) \geq 2(n_{\bar{\Sigma}}^r + S^r)$$

by the absolute grading formula. Similarly, Lemma 17.4 implies that

$$\mu(\phi_{\mathbb{T}}^r) \geq n_{\mathbb{T}}^r + S^r.$$

We observe that the expected dimension in equation (17.4) is at least

$$n_{\bar{\Sigma}}^r + 1 + S^r.$$

Hence, $S^r = n_{\bar{\Sigma}}^r = 0$ for a generic degeneration. Since $|n_{\mathbb{T}}^r - n_{\bar{\Sigma}}^r| \leq 1$, we must have $n_{\mathbb{T}}^r \in \{0, 1\}$.

We therefore conclude that

$$\mu(\phi_{\bar{\Sigma}}^r) = n_p(\phi_{\bar{\Sigma}}^r) = 0 \quad \text{and} \quad \mu(\phi_{\mathbb{T}}^r) = 1,$$

and $n_{\mathbb{T}}^r = 1$, for a generic degeneration.

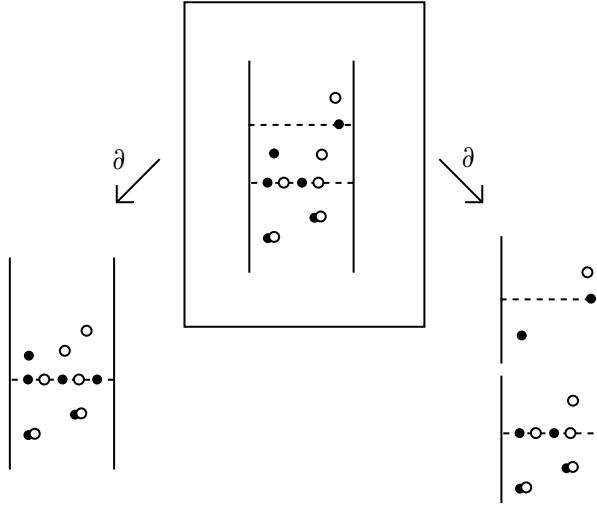


Figure 17.5. A schematic of the homotopy $\Psi^{\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma}} \simeq (\text{id} \otimes \Phi_w^{\bar{\Sigma}}) \circ \Psi^{\bar{\Sigma} < \mathbb{T}}$.

From this broken curve, we may extend the moduli space, by considering the $(\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma})$ -matched moduli space of triangles where the special line has height below the $\bar{\beta} \cdot \bar{\beta}$ boundary puncture of the triangle Δ (see Figure 17.5).

There is a natural map

$$\Psi^{\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma}}: CF_{\wedge}(\bar{\beta} c_w, \bar{\alpha} \bar{c}_w) \rightarrow CF_{\otimes}(\bar{\beta} c_w, \bar{\alpha} \bar{c}_w),$$

which counts curve pairs of expected dimension 0 with the following matchings.

- (1) Punctures along the special line are $(\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma})$ -matched.
- (2) Punctures below the special line are perfectly matched.
- (3) Punctures above the special line are trivially matched.

We claim that

$$\Psi^{\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma}} \simeq (\text{id} \otimes \Phi_w^{\bar{\Sigma}}) \circ \Psi^{\bar{\Sigma} < \mathbb{T}}.$$

This is proven by splitting the special line into two special lines. The higher special line contains a single $\bar{\Sigma}$ -marked point. The lower special line is $\bar{\Sigma} < \mathbb{T}$ matched. Other marked points are perfectly matched below the lower special line, and trivially matched above it. (i.e., there is no change in matching as we cross the $\bar{\Sigma}$ -matched line.)

We count the ends of these 1-parameter moduli spaces with two special lines. The only non-cancelling ends are as follows.

- (1) An index 1 disk breaking off while the special lines have finite, non-zero distance, giving a chain homotopy.

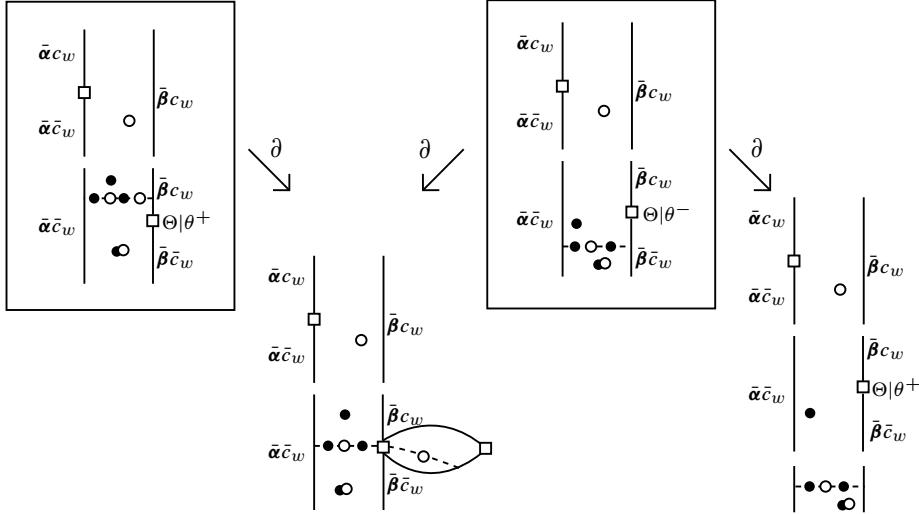


Figure 17.6. Canceling moduli space ends in Lemma 17.5. Squares denote special inputs of the holomorphic polygon maps. The curves in the square indicate 1-dimensional moduli spaces, and the arrows labeled with ∂ indicate codimension 1 degenerations. Solid dots are $\bar{\Sigma}$ -punctures and open dots are \mathbb{T} -punctures.

(2) The vertical distance between the two canceling ends approaches 0 or ∞ . The distance 0 end contributes $\Psi^{\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma}}$ while the distance ∞ end corresponds to the composition $(\text{id} \otimes \Phi_w^{\bar{\Sigma}}) \circ \Psi^{\bar{\Sigma} < \mathbb{T}}$.

Counting ends as above yields

$$\Psi^{\bar{\Sigma} < \mathbb{T} < \dots < \bar{\Sigma}} \simeq (\text{id} \otimes \Phi_w^{\bar{\Sigma}}) \circ \Psi^{\bar{\Sigma} < \mathbb{T}}.$$

Finally, we put all the pieces together. When counting triangles with trivial matching conditions, it is straightforward to see that

$$\Psi_{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}^{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})} = \text{id} \otimes \text{id},$$

on the level of intersection points.

Counting the ends of the above moduli spaces, we have ends which contribute the map

$$F_3^{c_w, c_w} \circ \Psi_{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}^{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}.$$

This is equal to

$$\begin{pmatrix} x \otimes \theta^- \mapsto x \\ x \otimes \theta^+ \mapsto 0 \end{pmatrix} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}.$$

Additionally, we have the contribution corresponding to extending the bigon degeneration. This gives an extra contribution of $\Psi^{\bar{\Sigma} \times \mathbb{T}}$ post-composed with

$$x \otimes \theta^+ \mapsto \Phi_w^{\bar{\Sigma}}(x).$$

Summing this with the previous contribution gives the main statement. \blacksquare

Remark 17.6. Note that the proof of equation (17.3) adapts easily to show that

$$F_3^{\alpha, \alpha} \circ \Psi^{\Sigma \# \bar{\Sigma} \times \mathbb{T}} = \Psi^{\bar{\Sigma} \times \mathbb{T}} \circ F_3^{\alpha, \alpha}.$$

17.5. Central hypercubes

In this section, we describe how to commute $\Psi^{\Sigma \# \bar{\Sigma} \times \mathbb{T}}$ past the levels of tw_* which involve the quadrilateral maps. The main result of this section is the following lemma.

Lemma 17.7. *The map $\Psi^{\Sigma \# \bar{\Sigma} \times \mathbb{T}}$ extends to a homotopy equivalence of hyperboxes from the perfectly matched hyperbox*

$$\begin{array}{ccc}
 CF(\alpha \bar{c}_w \bar{\beta}, \beta \bar{c}_w \bar{\beta}) & \longrightarrow & CF(\alpha \bar{c}_w \bar{\beta}, \beta \bar{c}_w \bar{\beta}) \\
 \downarrow & \searrow & \downarrow \\
 CF(\alpha \bar{c}_w \bar{\beta}, \beta D \bar{\beta}) & \longrightarrow & CF(\alpha \bar{c}_w \bar{\beta}, \beta \tilde{D} \bar{\beta}) \\
 \downarrow & \searrow & \downarrow \\
 CF(\alpha \bar{c}_w \bar{\beta}, \Delta_1) & \longrightarrow & CF(\alpha \bar{c}_w \bar{\beta}, \tilde{\Delta}_1) \\
 \downarrow & \searrow & \downarrow \\
 \vdots & \searrow & \vdots \\
 \downarrow & \searrow & \downarrow \\
 CF(\alpha \bar{c}_w \bar{\beta}, \Delta_n) & \longrightarrow & CF(\alpha \bar{c}_w \bar{\beta}, \tilde{\Delta}_n) \\
 \downarrow & \searrow & \downarrow \\
 CF(\alpha \bar{c}_w \bar{\beta}, \alpha D \bar{\alpha}) & \longrightarrow & CF(\alpha \bar{c}_w \bar{\beta}, \alpha \tilde{D} \bar{\alpha}) \\
 \downarrow & \searrow & \downarrow \\
 CF(\alpha \bar{c}_w \bar{\beta}, \alpha \bar{c}_w \bar{\alpha}) & \longrightarrow & CF(\alpha \bar{c}_w \bar{\beta}, \alpha \bar{c}_w \bar{\alpha})
 \end{array}$$

to its trivially matched counterpart.

Proof. We observe that all of the attaching curves which appear in the statement are disjoint from the connected sum tube which connects $\bar{\Sigma}$ and \mathbb{T}^2 . In fact, we may describe the entire hyperbox using a connected sum operation on hypercubes of

attaching curves. This connected sum operation on hypercubes of attaching curves is a small modification of the one considered by Lipshitz, Ozsváth and Thurston in [15, Section 3.5] (see [38, Section 9.5] for our present notation). If $\mathcal{L}_\beta = (\beta_\varepsilon, \Theta_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ and $\mathcal{L}_{\beta'} = (\beta'_{\nu}, \Theta_{\nu, \nu'})_{\nu \in \mathbb{E}_m}$ are two hypercubes of attaching curves on Heegaard surfaces Σ and Σ' , then $\mathcal{L}_\beta \# \mathcal{L}_{\beta'}$ is a hypercube of attaching curves of dimension $n + m$. The curves of $\mathcal{L}_\beta \# \mathcal{L}_{\beta'}$ are defined as

$$\mathcal{B}_{(\varepsilon, \nu)} = \beta_\varepsilon \cup \beta'_{\nu}$$

with small translations of curves taken to achieve admissibility. In our present setting, if the connected sum of Σ and Σ' is taken at either w or w' (deleting this basepoint), then the curves $\mathcal{B}_{(\varepsilon, \nu)}$ will also form an admissible diagram which is algebraically rigid, so we take the length 1 Floer chains of $\mathcal{L}_\beta \# \mathcal{L}_{\beta'}$ to be the top degree generators, and all higher length chains to vanish.

In the hyperbox in the statement of the lemma, each constituent hypercube of chain complexes is obtained by pairing the 0-dimensional hypercube $\alpha \bar{c}_w \bar{\beta}$ with a 2-dimensional hypercube of beta attaching curves. We view $\alpha \bar{c}_w \bar{\beta}$ as the connected sum of 0-dimensional hypercubes $\alpha \bar{\beta}$ with \bar{c}_w . In the top-most and bottom-most hypercubes, the beta hypercube of attaching curves are connected sums of a 0-dimensional hypercube on $\Sigma \# \bar{\Sigma}$ with a 2-dimensional hypercube on \mathbb{T}^2 . The remaining beta hypercubes are connected sums of a 1-dimensional hypercube on $\Sigma \# \bar{\Sigma}$ and a 1-dimensional hypercube on \mathbb{T}^2 .

The work of the last author [38, Section 11.3] constructs, under suitable hypotheses, a homotopy equivalence between the hyperbox in the lemma statement (i.e., the hyperbox obtained by counting curve pairs with perfect matching) and the hyperbox obtained by counting curve pairs with trivial matching. The precise statement may be found in [38, Proposition 11.8] and involves the following somewhat technical definitions.

- (1) The alpha and beta hypercubes of attaching curves must be *algebraically rigid* and consist of handleslide equivalent attaching curves.
- (2) A technical condition about the placement of the connected sum points and the basepoint must be satisfied. See [38, Definition 11.4]. We say that a hypercube of algebraically rigid, handleslide equivalent attaching curves is *graded* by the connected sum point p if for all nonnegative classes $\phi \in \pi_2(\Theta_{\varepsilon_1, \varepsilon_2}, \dots, \Theta_{\varepsilon_{n-1}, \varepsilon_n})$ (where $\Theta_{\varepsilon_i, \varepsilon_{i+1}}$ are chains from the hypercube), we have

$$\mu(\phi) \geq 2n_p(\phi).$$

We recall the statement of [38, Proposition 11.8]. Suppose that Σ and Σ' are Heegaard surfaces with distinguished points $p \in \Sigma$ and $p' \in \Sigma'$, and \mathcal{L}_α and \mathcal{L}_β are

algebraically rigid hypercubes on Σ and \mathcal{L}_α and \mathcal{L}_β are algebraically rigid hypercubes on Σ' . If \mathcal{L}_α is graded by p and \mathcal{L}_β is graded by p' , then the map

$$\Psi^{\Sigma' \times \Sigma}: CF_\wedge(\mathcal{L}_\alpha \# \mathcal{L}_{\alpha'}, \mathcal{L}_\beta \# \mathcal{L}_{\beta'}) \rightarrow CF_\otimes(\mathcal{L}_\alpha \# \mathcal{L}_{\alpha'}, \mathcal{L}_\beta \# \mathcal{L}_{\beta'})$$

is a homotopy equivalence. Here, CF_\wedge has differential counting perfectly matched curve pairs, while CF_\otimes counts trivially matched curve pairs.

We now return to the original setting of the present lemma. By construction, each hypercube is algebraically rigid. Furthermore, the alpha hypercubes $\alpha\bar{\beta}$ and \bar{c}_w are vacuously graded by the connected sum point since they are 0-dimensional. All of the hypercubes on $\Sigma \# \bar{\Sigma}$ are graded by the connected sum point since the connected sum point is adjacent to the basepoint and the hypercubes are algebraically rigid. The beta hypercubes on \mathbb{T}^2 are not all graded by the basepoint, but all of the alpha hypercubes which lie on \mathbb{T}^2 are. The only difference in our present situation and [38] is that the hypercubes of attaching curves do not all consist of handleslide equivalent attaching curves. For example, $\beta\bar{\beta}$ is not handleslide equivalent to Δ_1^0 . Nonetheless, the proof of [38, Proposition 11.8] still goes through since Heegaard diagram formed from a pair of curves in the diagram represents an algebraically rigid diagram for a connected sum of $S^1 \times S^2$'s and the grading assumption in (2) is satisfied. Hence, the same argument as in [38, Proposition 11.8] provides the homotopy equivalence in the statement. ■

Next, we view the above hyperbox as consisting of attaching curves on the disjoint union of $\Sigma \# \bar{\Sigma}$ and \mathbb{T} , instead of the connected sum of these two diagrams. Note that compression commutes with pairing up to homotopy, so we may instead compress the hyperbox of attaching curves:

$$\begin{array}{ccc}
 \beta\bar{c}_w\bar{\beta} & \longrightarrow & \beta\bar{c}_w\bar{\beta} \\
 \downarrow & & \downarrow \\
 \beta D\bar{\beta} & \longrightarrow & \beta\tilde{D}\bar{\beta} \\
 \downarrow & & \downarrow \\
 \Delta_1 & \longrightarrow & \tilde{\Delta}_1 \\
 \downarrow & & \downarrow \\
 \vdots & & \vdots \\
 \downarrow & & \downarrow \\
 \Delta_n & \longrightarrow & \tilde{\Delta}_n \\
 \downarrow & & \downarrow \\
 \alpha D\bar{\alpha} & \longrightarrow & \alpha\tilde{D}\bar{\alpha} \\
 \downarrow & & \downarrow \\
 \alpha\bar{c}_w\bar{\alpha} & \longrightarrow & \alpha\bar{c}_w\bar{\alpha}
 \end{array} \tag{17.5}$$

Note that we have not defined compression of hyperboxes of attaching curves. Nonetheless, the function-composition approach from Section 2.1 generalizes to this setting with minimal care. In our present case this amounts to viewing each of the $n + 2$ levels in the above diagram as a twisted complex in the Fukaya category. Each pair of adjacent levels in the diagram determines a morphism between 1-dimensional hypercubes of attaching curves. A pair of morphisms between hypercubes of attaching curves may be composed, so we compose iteratively all of these morphisms. Note that this is not the same as applying the ordinary A_∞ -composition map μ_{n+2} . Rather, it is the result of applying μ_2 repeatedly.

Lemma 17.8. *The compression of the diagram in equation (17.5) takes the following form:*

$$\begin{array}{ccc}
 \beta \bar{c}_w \bar{\beta} & \xrightarrow{1} & \beta \bar{c}_w \bar{\beta} \\
 \downarrow X|\theta^- & \searrow X'|\theta^+ & \downarrow X|\theta^- \\
 \alpha \bar{c}_w \bar{\alpha} & \xrightarrow{1} & \alpha \bar{c}_w \bar{\alpha}
 \end{array}$$

Here, X is the iterated composition of the vertical arrows on $\Sigma \# \bar{\Sigma}$, and θ^\pm are the intersection points of $\bar{c}_w \cap \bar{c}'_w$. Furthermore, the cycle X' is homologous to X .

Proof. We view the diagram as being obtained by composing $n + 1$ morphisms between 1-dimensional hypercubes of Lagrangians. Viewing this diagram as a sequence of morphisms $\theta_{n+1}, \dots, \theta_1$, we form the composition as

$$\mu_2(\theta_{n+1}, \dots, \mu_2(\theta_2, \theta_1)).$$

We first claim that except for the final composition, there are no diagonal terms in any partial composition. To see this, consider the first composition; the same argument will apply to all later compositions except for the final. There are three holomorphic quadrilateral counts which contribute to the diagonal term, corresponding to the following quadrilaterals:

- (1) $(\beta \bar{c}_w \bar{\beta}, \beta \bar{c}_w \bar{\beta}, \beta \tilde{D} \bar{\beta}, \tilde{\Delta}_1)$,
- (2) $(\beta \bar{c}_w \bar{\beta}, \beta D \bar{\beta}, \beta \tilde{D} \bar{\beta}, \tilde{\Delta}_1)$,
- (3) $(\beta \bar{c}_w \bar{\beta}, \beta D \bar{\beta}, \Delta_1, \tilde{\Delta}_1)$.

Each of these Heegaard quadruples decomposes as a connected sum where both factors have one set of attaching curves which is a small translation of an adjacent set. Since we are counting curves using trivial matchings, we use [4, Proposition 11.5] to see that no holomorphic quadrilaterals contribute. Similarly, using a model count on the genus 1 diagram, we observe that the two triangle compositions in the vertical direction coincide.

Therefore, the final composition is of a diagram of the following form:

$$\begin{array}{ccc}
 \beta \bar{c}_w \bar{\beta} & \longrightarrow & \beta \bar{c}_w \bar{\beta} \\
 \downarrow X|\theta & & \downarrow X|\theta \\
 \alpha D \bar{\alpha} & \longrightarrow & \alpha \tilde{D} \bar{\alpha} \\
 \downarrow & & \downarrow \\
 \alpha \bar{c}_w \bar{\alpha} & \longrightarrow & \alpha \bar{c}_w \bar{\alpha}
 \end{array} \tag{17.6}$$

where θ denotes the intersection point of $\bar{c}_w \cap D$ or $\bar{c}_w \cap \tilde{D}$.

As a first step, we consider the length 1 maps in the compression of (17.6). The length one chain on the left side of the compression is given by counting triangles on $(\Sigma \# \bar{\Sigma}, \beta \bar{\alpha}, \alpha \bar{\alpha}, \alpha \bar{\alpha})$ and $(\mathbb{T}^2, \bar{c}_w, D, \bar{c}_w)$. There is no matching condition, so we count triangles on the two diagrams separately and see that the triangle maps applied to $(X|\theta, \Theta^+|\theta)$ have output $X|\theta^-$. The computation on $\Sigma \# \bar{\Sigma}$ follows from nearest-point triangle counts, and the computation on \mathbb{T}^2 is a direct and straightforward triangle count.

We now consider the length 2 map of the compression. The small translate theorems [4, Propositions 11.1 and 11.5] prohibit holomorphic quadrilaterals except on

$$(\Sigma \# \bar{\Sigma}, \beta \bar{\beta}, \alpha \bar{\alpha}, \alpha \bar{\alpha}, \alpha \bar{\alpha}) \# (\mathbb{T}^2, \bar{c}_w, D, \tilde{D}, \bar{c}_w).$$

Since the diagram on $\Sigma \# \bar{\Sigma}$ has several adjacent sets of attaching curves which are small translates of one another, the only curves counted consist of pairs (u, u') where u is an index 0 quadrilateral on $\Sigma \# \bar{\Sigma}$ and u' is an index -1 quadrilateral on \mathbb{T}^2 . Furthermore, the two quadrilaterals share the same almost complex structure parameter (viewing the moduli space of rectangles as \mathbb{R}). There is a unique such curve u' , shown in Figure 17.3. Write $X'|\theta^-$ for the output.

To see that X' is homologous to X we deform the matching, by instead requiring the quadrilaterals u and u' to have almost $\text{ev}_\square(u) - \text{ev}_\square(u') = t$, where ev_\square denotes the evaluation map to the moduli space of complex rectangles, identified with \mathbb{R} . We then send $t \rightarrow \infty$. The effect is to change X' by a boundary. This shows that X' is homologous to the composition

$$\mu_2(\mu_2(X, \Theta_{\alpha \bar{\alpha}, \alpha \bar{\alpha}}), \Theta_{\alpha \bar{\alpha}, \alpha \bar{\alpha}}),$$

which we identify with X via the nearest point map. This completes the proof. \blacksquare

17.6. Proof of Theorem 1.7

We now complete our proof of the computation of the map tw_* .

Proof of Theorem 1.7. It remains to put the pieces together to finish the proof. We build a 3-dimensional hyperbox to realize the homotopy. The back face is the map tw_* , computed using perfect matching between $\Sigma \# \bar{\Sigma}$ and \mathbb{T}^2 .

The front face is defined very similarly, using the trivial (i.e., tensor product) matchings, with a modification of one cube. In place of

$$\begin{array}{ccc}
 CF_{\wedge}(\bar{c}_w \bar{\beta}, \bar{c}_w \bar{\alpha}) & \xrightarrow{\text{id}} & CF_{\wedge}(\bar{c}_w \bar{\beta}, \bar{c}_w \bar{\alpha}) \\
 \downarrow F_3^{cw, cw} \Psi_{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}^{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})} & & \downarrow F_3^{cw, cw} \Psi_{(\bar{c}_w \bar{\beta}) \rightarrow (c_w \bar{\beta})}^{(\bar{c}_w \bar{\alpha}) \rightarrow (c_w \bar{\alpha})} \\
 CF(\bar{\beta}, \bar{\alpha}) & \xrightarrow{\text{id}} & CF(\bar{\beta}, \bar{\alpha})
 \end{array}$$

we have

$$\begin{array}{ccc}
 CF_{\otimes}(\bar{c}_w \bar{\beta}, \bar{c}_w \bar{\alpha}) & \xrightarrow{\text{id}} & CF_{\otimes}(\bar{c}_w \bar{\beta}, \bar{c}_w \bar{\alpha}) \\
 \downarrow \Psi' & & \downarrow \Psi' \\
 CF(\bar{\beta}, \bar{\alpha}) & \xrightarrow{\text{id}} & CF(\bar{\beta}, \bar{\alpha})
 \end{array}$$

where

$$\Psi' = \begin{pmatrix} x \otimes \theta^- \mapsto x \\ x \otimes \theta^+ \mapsto \Phi_w(x) \end{pmatrix} \circ \Psi^{\bar{\Sigma} < \mathbb{T}}.$$

Lemmas 17.3, 17.7 and equation (17.3) show that the hyperbox relations are satisfied for appropriate choices of length 3 maps on the interior of the cube.

Next, we compress the front face of the hypercube described above. We pair the hypercube from Lemma 17.8 with $\alpha \bar{c}_w \bar{\beta}$ and then stack with the 2-dimensional hypercube

$$\begin{array}{ccc}
 CF(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}) & \xrightarrow{\text{id}} & CF(\bar{\Sigma}, \bar{\beta}, \bar{\alpha}) \\
 \downarrow \eta & & \downarrow \eta \\
 CF(\Sigma, \alpha, \beta) & \xrightarrow{\text{id}} & CF(\Sigma, \alpha, \beta)
 \end{array}$$

The so-constructed cube has the form

$$\begin{array}{ccc}
 CF(\alpha, \beta) & \xrightarrow{\text{id}} & CF(\alpha, \beta) \\
 \downarrow \iota & \searrow \Phi_w \iota & \downarrow \iota \\
 CF(\alpha, \beta) & \xrightarrow{\text{id}} & CF(\alpha, \beta),
 \end{array}$$

completing the proof. ■

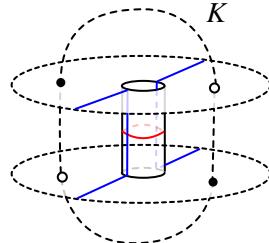


Figure 17.7. A doubled knot diagram (for the unknot). Meridional twisting preserves the diagram.

17.7. Knots and links

If K is a knot, then *a priori* there is an action tw_*^μ corresponding to rotating a neighborhood of K in the meridional direction (see Figure 17.7). (Rotation in the longitudinal direction is well known to induce an interesting automorphism by work of Sarkar [30].) Despite the interesting potential for a new map, we have the following proposition.

Proposition 17.9. *The map tw_*^μ is chain homotopic to id on $\mathcal{CFKI}(K)$. The same holds for links.*

Proof. We consider the stabilized model of the involution. For knots and links, the doubling arcs δ have endpoints at a point p which is not one of the basepoints of the knot. Hence, applying tw^μ to each Heegaard diagram has no effect on any of the attaching curves. ■

Remark 17.10. We can also recover the above result on the non-stabilized doubling models of the knot involution, considered in Section 3.4 and [4, Section 2.2], though the argument is more subtle. Here, we double at one of the basepoints (say $w \in \Sigma$), with the δ curves having endpoints at w . In this model, after attaching 1-handles, we move w to the position of \bar{z} . The map tw^μ still changes the curves Δ , but since there is no basepoint in the tube region (unlike in the case of closed 3-manifolds), the effect on the level of diagrams is to act by a surface diffeomorphism which is isotopic to the identity relative to the basepoints on $\Sigma \# \bar{\Sigma}$.

Acknowledgments. The authors are grateful to András Juhász, Ciprian Manolescu, and Dylan Thurston for helpful conversations (over many years). The authors are further grateful to the anonymous referee for helpful comments. Parts of this work were carried out while the first and second authors were in residence at the Simons Laufer Mathematical Institute/Mathematical Sciences Research Institute in Fall 2022; the authors are grateful to the institute for its hospitality.

Funding. The first author was partially supported by NSF grant DMS-2019396 and a Sloan Research Fellowship. The second author was partially supported by NSF grants DMS-1552285 and DMS-2104144. Parts of this work were carried out while the first and second authors were in residence at the Simons Laufer Mathematical Institute/Mathematical Sciences Research Institute in Fall 2022, and therefore supported by NSF Grant DMS-1440140. The third author was partially supported by NSF grant DMS-1702532. The fourth author was partially supported by NSF grants DMS-1703685 and DMS-2204375.

References

- [1] A. Alishahi and E. Eftekhary, [Tangle Floer homology and cobordisms between tangles](#). *J. Topol.* **13** (2020), no. 4, 1582–1657 Zbl 1511.57019 MR 4186139
- [2] F. Binns, S. Kang, J. Simone, and P. Truöl, [On the nonorientable four-ball genus of torus knots](#). 2021, arXiv:2109.09187
- [3] J. Cerf, [La stratification naturelle des espaces de fonctions différentiables réelles et le théorème de la pseudo-isotopie](#). *Inst. Hautes Études Sci. Publ. Math.* (1970), no. 39, 5–173 Zbl 0213.25202 MR 0292089
- [4] K. Hendricks, J. Hom, M. Stoffregen, and I. Zemke, [Surgery exact triangles and involutive Heegaard Floer homology](#). 2020, arXiv:2011.00113
- [5] K. Hendricks and C. Manolescu, [Involutive Heegaard Floer homology](#). *Duke Math. J.* **166** (2017), no. 7, 1211–1299 Zbl 1383.57036 MR 3649355
- [6] K. Hendricks, C. Manolescu, and I. Zemke, [A connected sum formula for involutive Heegaard Floer homology](#). *Selecta Math. (N.S.)* **24** (2018), no. 2, 1183–1245 Zbl 1407.57025 MR 3782421
- [7] A. Juhász, [Cobordisms of sutured manifolds and the functoriality of link Floer homology](#). *Adv. Math.* **299** (2016), 940–1038 Zbl 1358.57021 MR 3519484
- [8] A. Juhász and M. Marengon, [Concordance maps in knot Floer homology](#). *Geom. Topol.* **20** (2016), no. 6, 3623–3673 Zbl 1364.57013 MR 3590358
- [9] A. Juhász, D. Thurston, and I. Zemke, [Naturality and mapping class groups in Heegaard Floer homology](#). *Mem. Amer. Math. Soc.* **273** (2021), no. 1338, v+174 Zbl 1510.57001 MR 4337438
- [10] A. Juhász and I. Zemke, [Contact handles, duality, and sutured Floer homology](#). *Geom. Topol.* **24** (2020), no. 1, 179–307 Zbl 1467.57023 MR 4080483
- [11] A. Juhász and I. Zemke, [Stabilization distance bounds from link Floer homology](#). [v1] 2018, [v3] 2021, arXiv:1810.09158v3
- [12] R. Kirby, [A calculus for framed links in \$S^3\$](#) . *Invent. Math.* **45** (1978), no. 1, 35–56 Zbl 0377.55001 MR 0467753
- [13] F. Lin, [The surgery exact triangle in Pin\(2\)-monopole Floer homology](#). *Algebr. Geom. Topol.* **17** (2017), no. 5, 2915–2960 Zbl 1387.57027 MR 3704248
- [14] R. Lipshitz, [A cylindrical reformulation of Heegaard Floer homology](#). *Geom. Topol.* **10** (2006), 955–1096 Zbl 1130.57035 MR 2240908

- [15] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, [Bordered Floer homology and the spectral sequence of a branched double cover II: the spectral sequences agree](#). *J. Topol.* **9** (2016), no. 2, 607–686 Zbl 1345.57023 MR 3509974
- [16] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, [Bordered Heegaard Floer homology](#). *Mem. Amer. Math. Soc.* **254** (2018), no. 1216, viii+279 Zbl 1422.57080 MR 3827056
- [17] Y. Liu, [Heegaard Floer homology of surgeries on two-bridge links](#). *Algebr. Geom. Topol.* **22** (2022), no. 2, 473–557 Zbl 1505.57020 MR 4464459
- [18] C. Manolescu and P. Ozsváth, Heegaard Floer homology and integer surgeries on links. [v1] 2010, [v5] 2022, arXiv:1011.1317v5
- [19] Y. Ni, [Homological actions on sutured Floer homology](#). *Math. Res. Lett.* **21** (2014), no. 5, 1177–1197 Zbl 1310.57028 MR 3294567
- [20] P. Ozsváth and Z. Szabó, [Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary](#). *Adv. Math.* **173** (2003), no. 2, 179–261 Zbl 1025.57016 MR 1957829
- [21] P. Ozsváth and Z. Szabó, [Holomorphic disks and knot invariants](#). *Adv. Math.* **186** (2004), no. 1, 58–116 Zbl 1062.57019 MR 2065507
- [22] P. Ozsváth and Z. Szabó, [Holomorphic disks and three-manifold invariants: properties and applications](#). *Ann. of Math. (2)* **159** (2004), no. 3, 1159–1245 Zbl 1081.57013 MR 2113020
- [23] P. Ozsváth and Z. Szabó, [Holomorphic disks and topological invariants for closed three-manifolds](#). *Ann. of Math. (2)* **159** (2004), no. 3, 1027–1158 Zbl 1073.57009 MR 2113019
- [24] P. Ozsváth and Z. Szabó, [Holomorphic triangles and invariants for smooth four-manifolds](#). *Adv. Math.* **202** (2006), no. 2, 326–400 Zbl 1099.53058 MR 2222356
- [25] P. Ozsváth and Z. Szabó, [Holomorphic disks, link invariants and the multi-variable Alexander polynomial](#). *Algebr. Geom. Topol.* **8** (2008), no. 2, 615–692 Zbl 1144.57011 MR 2443092
- [26] P. Ozsváth and Z. Szabó, [Algebras with matchings and knot Floer homology](#). 2019, arXiv:1912.01657
- [27] P. S. Ozsváth and Z. Szabó, [Knot Floer homology and integer surgeries](#). *Algebr. Geom. Topol.* **8** (2008), no. 1, 101–153 Zbl 1181.57018 MR 2377279
- [28] J. A. Rasmussen, [Floer homology and knot complements](#). Ph.D. thesis, Harvard University, Ann Arbor, MI, 2003 Zbl 2704683
- [29] S. Sarkar, [Maslov index formulas for Whitney \$n\$ -gons](#). *J. Symplectic Geom.* **9** (2011), no. 2, 251–270 Zbl 1231.53068 MR 2811652
- [30] S. Sarkar, [Moving basepoints and the induced automorphisms of link Floer homology](#). *Algebr. Geom. Topol.* **15** (2015), no. 5, 2479–2515 Zbl 1331.57015 MR 3426686
- [31] L. Tsvetkova-Nelip, [Concordance maps in \$\text{HFK}^-\$](#) . *J. Knot Theory Ramifications* **31** (2022), no. 5, article no. 2250031 Zbl 1496.57014 MR 4450929
- [32] I. Zemke, [Quasistabilization and basepoint moving maps in link Floer homology](#). *Algebr. Geom. Topol.* **17** (2017), no. 6, 3461–3518 Zbl 1387.57020 MR 3709653
- [33] I. Zemke, [Connected sums and involutive knot Floer homology](#). *Proc. Lond. Math. Soc. (3)* **119** (2019), no. 1, 214–265 Zbl 1473.57043 MR 3957835

- [34] I. Zemke, [Link cobordisms and absolute gradings on link Floer homology](#). *Quantum Topol.* **10** (2019), no. 2, 207–323 Zbl [1478.57035](#) MR [3950650](#)
- [35] I. Zemke, [Link cobordisms and functoriality in link Floer homology](#). *J. Topol.* **12** (2019), no. 1, 94–220 Zbl [1455.57020](#) MR [3905679](#)
- [36] I. Zemke, [Duality and mapping tori in Heegaard Floer homology](#). *J. Topol.* **14** (2021), no. 3, 1027–1112 Zbl [1504.57022](#) MR [4503956](#)
- [37] I. Zemke, [A graph TQFT for hat Heegaard Floer homology](#). *Quantum Topol.* **12** (2021), no. 3, 439–460 Zbl [1479.57033](#) MR [4321211](#)
- [38] I. Zemke, Bordered manifolds with torus boundary and the link surgery formula. [v1] 2021, [v5] 2023, arXiv:[2109.11520v5](#)
- [39] I. Zemke, Graph cobordisms and Heegaard Floer homology. [v1] 2015, [v3] 2019, arXiv:[1512.01184v3](#)

Received 10 September 2022.

Kristen Hendricks

Department of Mathematics, Rutgers University New Brunswick, 110 Freylinghuysen Rd, New Brunswick, 08854, USA; kristen.hendricks@rutgers.edu

Jennifer Hom

Department of Mathematics, Georgia Institute of Technology, 686 Cherry Street, Atlanta, 30332, USA; hom@math.gatech.edu

Matthew Stoffregen

Department of Mathematics, Michigan State University, 619 Red Cedar Road, East Lansing, 48824, USA; stoffre1@msu.edu

Ian Zemke

Department of Mathematics, Princeton University, Fine Hall, Washington Road, Princeton, 08540, USA; izemke@math.princeton.edu