Geometry & Topology

Volume 29 (2025)

Heegaard Floer homology and integer surgeries on links

Ciprian Manolescu
Peter Ozsváth

msp

Geometry & Topology **29**:6 (2025) 2783–3062
DOI: 10.2140/gt.2025.29.2783
Published: 22 September 2025

Heegaard Floer homology and integer surgeries on links

CIPRIAN MANOLESCU
PETER OZSVÁTH

Let L be a link in an integral homology three-sphere. We give a description of the Heegaard Floer homology of integral surgeries on L in terms of some data associated to L, which we call a complete system of hyperboxes for L. Roughly, a complete system of hyperboxes consists of chain complexes for (some versions of) the link Floer homology of L and all its sublinks, together with several chain maps between these complexes. Further, we introduce a way of presenting closed four-manifolds with $b_2^+ \geq 2$ by four-colored framed links in the three-sphere. Given a link presentation of this kind for a four-manifold X, we then describe the Ozsváth–Szabó mixed invariants of X in terms of a complete system of hyperboxes for the link. Finally, we explain how a grid diagram produces a particular complete system of hyperboxes for the corresponding link.

57R58; 57R57

1 Introduction

Heegaard Floer homology is a tool for studying low-dimensional manifolds, introduced by Zoltán Szabó and the second author. In the case of closed three-manifolds, a genus g Heegaard diagram naturally endows the g-fold symmetric product of the Heegaard surface with a pair of half-dimensional tori; and a suitable adaptation of Lagrangian Floer homology in the symmetric product (relative to the Heegaard tori) turns out to depend only on the underlying three-manifold. This three-manifold invariant is constructed in [Ozsváth and Szabó 2004b; 2004c]; related invariants of four-dimensional cobordisms are constructed in [Ozsváth and Szabó 2006; Zemke 2015]; invariants for knots and links in three-manifolds are developed in [Ozsváth and Szabó 2004a; 2008a; Rasmussen 2003]. Of particular interest to us here are the mixed invariants of closed four-manifolds defined in [Ozsváth and Szabó 2006]: they can detect exotic smooth structures and, in fact, are conjecturally identical to the Seiberg–Witten invariants [Witten 1994].

The knot Floer homology groups from [Ozsváth and Szabó 2004a; Rasmussen 2003] are closely related to the Heegaard Floer homology groups of (closed) three-manifolds obtained as surgeries on the knot. Indeed, both these works showed that the filtered knot Floer complex contains enough information to recover the Heegaard Floer homologies of all sufficiently large surgeries on the respective knot. This is used as a stepping stone to reconstruct the Heegaard Floer homology of arbitrary surgeries on a knot in [Ozsváth and Szabó 2008b; 2011].

Since every closed three-manifold can be obtained by surgery on a link in the three-sphere, a natural question is whether the results from [Ozsváth and Szabó 2008b] admit a generalization for links. The goal of this paper is to present such a generalization.

Let $K \subset Y$ be a knot in an integral homology three-sphere. Recall that the knot Floer homology of K is constructed starting from a Heegaard diagram for Y that has two basepoints w and z, which specify the knot. One can build Floer homology groups by counting pseudoholomorphic curves in the symmetric product of the Heegaard surface in various ways. For example, one can require the support of the curves to avoid z, and at the same time keep track of the intersections with w by powers of a U variable: this gives rise to knot Floer homology. Alternatively, one can define complexes A_s^+ in which one keeps track the intersection number of the curve with both w and z, in a way that depends on the value of an auxiliary parameter $s \in \mathbb{Z}$. When $s \gg 0$, the complex A_s^+ corresponds to ignoring z completely and keeping track of w via a U variable. When $s \ll 0$, we have the reverse: A_s^+ corresponds to ignoring w completely and keeping track of z. Note that, whenever $|s| \gg 0$, the homology $H_*(A_s^+)$ is the Heegaard Floer homology of Y, regardless of the sign of s (and, in particular, they are independent of the knot $K \subset Y$). The intermediate complexes A_s^+ for $s \in \mathbb{Z}$, however, contain nontrivial information about the knot. Indeed, according to [Ozsváth and Szabó 2004a; Rasmussen 2003], the complexes A_s^+ (as s varies) capture the Floer homology groups of the three-manifolds obtained by sufficiently large surgeries on K. Moreover, these complexes are basic building blocks of the constructions from [Ozsváth and Szabó 2008b; 2011].

Consider now an oriented link $\vec{L}$ in an integral homology three-sphere Y. The analogue of knot Floer homology was defined in [Ozsváth and Szabó 2008a] and is called *link Floer homology*. The construction starts with a Heegaard diagram $\mathcal{H}^L$ for Y that has several w and z basepoints, specifying the link. Let $L_1, \dots, L_\ell$ be the components of L. Following [loc. cit.], we consider the affine lattice $\mathbb{H}(L)$ over $\mathbb{Z}^\ell$ defined by

$$\mathbb{H}(L) = \bigtimes_{i=1}^{\ell} \mathbb{H}(L)_i, \quad \mathbb{H}(L)_i = \tfrac{1}{2}\,\mathrm{lk}(L_i, L - L_i) + \mathbb{Z},$$

where lk denotes linking number. By keeping track of the basepoints in various ways, we can define generalized Floer chain complexes $\boldsymbol{A}^-(\mathcal{H}^L, \boldsymbol{s})$ for $\boldsymbol{s} \in \mathbb{H}(L)$, which are the analogues of the groups A_s^+ for knots with two basepoints. In the case where the diagram $\mathcal{H}^L$ is link-minimal (that is, it has only one w and one z basepoint on each link component), the complexes $\boldsymbol{A}^-(\mathcal{H}^L, \boldsymbol{s})$ are made of free modules; in general, this is not the case, and we will have to construct some resolutions $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ of $\boldsymbol{A}^-(\mathcal{H}^L, \boldsymbol{s})$. We will then use the groups $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ to reconstruct the Heegaard Floer homology of integer surgeries on L. Note that, to keep in line with the conventions in [loc. cit.], we will phrase our construction in terms of the HF^- rather than the HF^+ version of Heegaard Floer homology. We also ignore sign issues and work over the field $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$.

Further, for technical reasons, we find it useful to use a slightly different variant of HF^- than the one defined in [Ozsváth and Szabó 2004b]: we complete the groups with respect to the U variables, so that they become modules over the power series ring $\mathbb{F}[\![U]\!]$; compare also [Kronheimer and Mrowka 2007]. We denote this completed version by $\boldsymbol{HF}^-$; it has the following technical advantage over HF^-. The uncompleted version HF^- is functorial under cobordisms equipped with Spin^c structures, but it is not functorial under cobordisms per se, whereas the completed version $\boldsymbol{HF}^-$ is. In particular, $\boldsymbol{HF}^-$ satisfies surgery exact triangles (analogous to those in [Ozsváth and Szabó 2004c]) just like HF^+.

Fix a framing Λ for the link $\vec{L}$. For a component L_i of L, we let Λ_i be its induced framing, thought of as an element in $H_1(Y - L)$. The latter group can be identified with $\mathbb{Z}^\ell$ via the basis of oriented meridians for $\vec{L}$. Given a sublink $M \subseteq L$, we let $\Omega(M)$ be the set of all possible orientations on M. For $\vec{M} \in \Omega(M)$, we let $I_-(\vec{L}, \vec{M})$ denote the set of indices i such that the component L_i is in M and its orientation induced from $\vec{M}$ is opposite to the one induced from $\vec{L}$. Set

$$\Lambda_{\vec{L},\vec{M}} = \sum_{i \in I_-(\vec{L},\vec{M})} \Lambda_i \in H_1(Y - L) \cong \mathbb{Z}^\ell.$$

Let $Y_\Lambda(L)$ be the three-manifold obtained from Y by surgery on the framed link (L, Λ). The input that we use to reconstruct $\boldsymbol{HF}^-(Y_\Lambda(L))$ is called a *complete system of hyperboxes* for the link $\vec{L}$. The precise definition is given in Section 8.6. Roughly, a complete system $\mathcal{H}$ consists of Heegaard diagrams $\mathcal{H}^{L'}$ representing all possible sublinks $L' \subseteq L$, together with some additional data that produces maps

$$\Phi_{\boldsymbol{s}}^{\vec{M}} : \mathfrak{A}^-(\mathcal{H}^{L'}, \boldsymbol{s}) \to \mathfrak{A}^-(\mathcal{H}^{L'-M}, \psi^{\vec{M}}(\boldsymbol{s}))$$

for any $M \subseteq L' \subseteq L$, $\boldsymbol{s} \in \mathbb{H}(L')$ and $\vec{M} \in \Omega(M)$. Here, $\psi^{\vec{M}} : \mathbb{H}(L') \to \mathbb{H}(L' - M)$ are natural reduction maps. The orientation $\vec{M}$ comes into play as follows: starting with $\mathcal{H}^{L'}$, we delete the w basepoints corresponding to components $L_i \subseteq M$ with $i \in I_-(\vec{L}, \vec{M})$, and delete the z basepoints corresponding to the remaining components $L_i \subseteq M$. The resulting diagram represents the link $L - M$, and the complete system gives us a sequence of steps that relate it to $\mathcal{H}^{L-M}$, a diagram that also represents $L - M$. The map $\Phi^{\vec{M}}_{\boldsymbol{s}}$ is constructed by following that sequence of steps, and counting pseudoholomorphic polygons. (In the case where the diagram is not link-minimal, the definition of $\Phi^{\vec{M}}_{\boldsymbol{s}}$ also involves certain transition maps, which will be described in Section 13.) In particular, when $M = \varnothing$, the map $\Phi^{\vec{M}}_{\boldsymbol{s}}$ is the differential on the complex $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$.

Define

$$\mathcal{C}^-(\mathcal{H}, \Lambda) = \bigoplus_{M \subseteq L} \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s})), \tag{1}$$

where ψ^M simply means $\psi^{\vec{M}}$ with $\vec{M}$ being the orientation induced from the one on $\vec{L}$. Equip $\mathcal{C}^-(\mathcal{H}, \Lambda)$ with a boundary operator as follows. For $\boldsymbol{s} \in \mathbb{H}(L)$ and $\boldsymbol{x} \in \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$, set

$$\begin{aligned} \mathcal{D}^-(\boldsymbol{s}, \boldsymbol{x}) &= \sum_{N \subseteq L-M} \sum_{\vec{N} \in \Omega(N)} (\boldsymbol{s} + \Lambda_{\vec{L}, \vec{N}}, \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}(\boldsymbol{x})) \\ &\in \bigoplus_{N \subseteq L-M} \bigoplus_{\vec{N} \in \Omega(N)} \mathfrak{A}^-(\mathcal{H}^{L-M-N}, \psi^{M \cup \vec{N}}(\boldsymbol{s})) \subseteq \mathcal{C}^-(\mathcal{H}, \Lambda). \end{aligned} \tag{2}$$

Our main result is:

Theorem 1.1 *Fix a complete system of hyperboxes $\mathcal{H}$ for an oriented, ℓ-component link $\vec{L}$ in an integral homology three-sphere Y, and fix a framing Λ of L. There is an isomorphism of homology groups*

$$H_*(\mathcal{C}^-(\mathcal{H}, \Lambda)) \cong \boldsymbol{HF}^-_*(Y_\Lambda(L)), \tag{3}$$

where $\boldsymbol{HF}^-$ is the completed version of Heegaard Floer homology over the power series ring $\mathbb{F}[[U]]$.

Similar results hold for other variants of Heegaard Floer homology: $\widehat{HF}$, HF^+ and $\boldsymbol{HF}^\infty$ (the last being a completed version of HF^∞).

We refer to the isomorphism (3) as the *link surgery formula*. A more detailed description of this formula (but still without many technicalities) is given in Section 4. We focus there on knots and two-component links, and go over the example of surgeries on the Hopf link. The complete statement of the link surgery formula, in the particular case of complete systems that are link-minimal (that is, the Heegaard diagrams in the system are link-minimal), is found in Section 9. The statement in full generality, for arbitrary complete systems, is given in Section 13.

The proof of Theorem 1.1 also gives a way of describing the maps from $\boldsymbol{HF}^-(Y)$ to $\boldsymbol{HF}^-(Y_\Lambda(L))$ induced by the surgery cobordism. More generally, let W be a cobordism between two connected three-manifolds Y_1 and Y_2 such that W consists of two-handle additions only. We can then find a link $\vec{L} \subset S^3$ with a framing Λ and a sublink $L' \subseteq L$ such that surgery on L' (with framing specified in Λ) produces Y_1, and the framed link specified by $L - L'$, thought of as a subset of Y_1, exactly corresponds to the cobordism W going to Y_2. Using this setup, we can describe the map on Heegaard Floer homology induced by the cobordism W in terms of a complete system of hyperboxes for L.

Refining this idea, we arrive at a similar description of a nontrivial smooth, closed four-manifold invariant. Let X be a closed, oriented four-manifold with $b_2^+(X) \geq 2$. The constructions from [Ozsváth and Szabó 2006] associated to every Spin^c structure $\mathfrak{s}$ on X an invariant called the *mixed invariant* $\Phi_{X,\mathfrak{s}} \in \mathbb{F}$. (The original definition was over $\mathbb{Z}$ and involved a homology action, but we ignore this extra structure in the present paper.) The manifold X can be presented in terms of a link as follows. Delete two four-balls from X to obtain a cobordism W from S^3 to S^3. Then split this cobordism into four parts,

$$W = W_1 \cup_{Y_1} W_2 \cup_{Y_2} W_3 \cup_{Y_3} W_4,$$

such that W_1 consists of one-handles only, W_2 and W_3 of two-handles, and W_4 of three-handles; further, we arrange so that Y_2 is an admissible cut in the sense of [ibid., Definition 8.3]. Next, find a framed link $(\vec{L} \subset S^3, \Lambda)$ that splits as a disjoint union $L_1 \cup L_2 \cup L_3$ such that surgery on L_1 produces Y_1, and surgery on L_2 and L_3 corresponds to the cobordisms W_2 and W_3, respectively. We refer to the data $(\vec{L} = L_1 \cup L_2 \cup L_3, \Lambda)$ as a *cut link presentation* for *X*.

Theorem 1.2 *Let X be a closed four-manifold with $b_2^+(X) \geq 2$, together with a cut link presentation $(\vec{L} = L_1 \cup L_2 \cup L_3, \Lambda)$. One can describe the mixed invariants $\Phi_{X,\mathfrak{s}}, \mathfrak{s} \in \mathrm{Spin}^c(X)$ with coefficients in $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$ in terms of the framing Λ and a complete system of hyperboxes for $\vec{L}$.*

The advantage of presenting the three- and four-manifold invariants in terms of link Floer complexes is that the latter are understood better. Indeed, there exist several combinatorial descriptions of knot Floer homology; see [Manolescu et al. 2009a; Sarkar and Wang 2010; Ozsváth and Szabó 2009]. We focus on the description in [Manolescu et al. 2009a], in terms of grid diagrams, which has the advantage that it extends to all versions of link Floer homology for links in S^3. It turns out that a grid diagram for a link gives rise to a complete system of hypercubes for that link (provided that the grid has at least one free marking; see Section 15 for the exact condition). Hence, we can apply Theorems 1.1 and 1.2 to obtain a description of the Heegaard Floer invariants for three- and four-manifolds in terms of counts of holomorphic curves on symmetric products of grid diagrams. In [Manolescu et al. 2009b], we use this result to describe the Heegaard Floer invariants in purely combinatorial terms.

It would be interesting to see whether the procedure of constructing three- and four-manifold invariants from data associated to links, as presented here, can be applied to other settings. Indeed, there are several homological invariants for links in S^3 that have much in common with Floer homology, but for which it is unknown whether they admit extensions to three- and four-manifolds. We are referring in particular to the link homologies constructed in [Khovanov 2000; Khovanov and Rozansky 2008a; 2008b].

In a different direction, the maps induced by surgery cobordisms can be used to construct another link invariant, a *link surgery spectral sequence* (see [Ozsváth and Szabó 2005, Theorem 4.1; Baldwin 2011, Theorem 5.2]). The methods giving Theorem 1.1 can also be used to give a description of this spectral sequence in terms of complete systems of hyperboxes; see Theorem 14.12 below for the precise statement.

The organization of this paper is as follows. In Section 2 we explain the difference between the completed theories $\boldsymbol{HF}^-$ and $\boldsymbol{HF}^\infty$ used in this paper, and the versions HF^- and HF^∞ originally defined in [Ozsváth and Szabó 2004b]. In Section 3 we define the generalized Floer chain complexes $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$, and explain their invariance properties. We also define maps between generalized Floer complexes by counting J-holomorphic polygons. Section 4 contains an overview of the link surgery formula, including a comparison with the knot surgery formula from [Ozsváth and Szabó 2008b], and an explicit computation for the case of the Hopf link. Sections 2–4 are sufficient for the reader who wants to reach a working understanding of the main result, without going into the proof or even into the full details of the statement.

In Section 5 we discuss some homological algebra that is needed throughout the rest of the paper: we introduce the notion of a hyperbox of chain complexes, and describe several operations on hyperboxes. Section 6 contains some analytical results about the behavior of holomorphic polygon maps under a move called quasistabilization. Section 7 contains similar analytical results for a move called index zero/three link stabilization. Section 8 is devoted to building up the definition of a complete system of hyperboxes for a link. In Section 9 we then give a more precise statement of Theorem 1.1, in the case of link-minimal complete systems. The proof of Theorem 1.1 for link-minimal complete systems occupies Section 12, and is based on a truncation procedure explained in Section 10 together with a surgery long exact sequence discussed in Section 11. We then turn to general complete systems: in Section 13 we give both the statement and the proof of Theorem 1.1 for arbitrary complete systems.

In Section 14 we present the extensions of Theorem 1.1 to the other versions of Heegaard Floer homology, and to the invariants associated to cobordisms; we also prove Theorem 1.2 and discuss the link surgeries spectral sequence. Finally, in Section 15, we explain how certain grid diagrams for links in S^3 give rise to complete systems of hyperboxes. Many of the diagrams in the resulting complete systems are quasistabilized, and therefore the results from Section 6 can be used to simplify the description of the surgery complex in terms of grids.

Acknowledgements We are thankful to Dylan Thurston for useful suggestions regarding the material in Section 5, to András Juhász, Matthew Hedden, Sucharit Sarkar and Zoltán Szabó for helpful conversations, and to John Baldwin, Damek Davis, Haofei Fan, Jonathan Hales, Tye Lidman, Yajing Liu, Yi Ni, Faramarz Vafaee and Ian Zemke for comments on previous versions of the paper. We are particularly grateful to the referee for many useful suggestions, and to Ian Zemke for pointing out the need for trajectory α-slides in Section 8.9. Manolescu would also like to thank the Mathematics Department at the University of Cambridge for its hospitality during his stay there in the spring of 2009.

Manolescu was supported by NSF grants DMS-0852439, DMS-1104406 and DMS-1402914 and a Royal Society University Research Fellowship. Ozsváth was supported by NSF grants DMS-0804121 and DMS-1258274 and a Guggenheim Fellowship.

2 Conventions

Throughout this paper we work with Floer homology groups with base field $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$.

Let Y be a closed, connected, oriented three-manifold. We consider the Heegaard Floer homology groups $\widehat{HF}$, HF^+, HF^- and HF^∞ defined in [Ozsváth and Szabó 2004b; 2004c]. These are modules over the polynomial ring $\mathbb{F}[U]$. (In the case of $\widehat{HF}$, the action of U is trivial.)

Let $\boldsymbol{HF}^-$ and $\boldsymbol{HF}^\infty$ denote the completions of HF^- and HF^∞ with respect to the maximal ideal (U) in the ring $\mathbb{F}[U]$. Since completion is an exact functor, we can alternatively think of $\boldsymbol{HF}^-$ as the homology of the complex $\boldsymbol{CF}^-$ with the same generators as CF^-, but whose coefficient ring is the formal power series ring $\mathbb{F}[\![U]\!]$ rather than $\mathbb{F}[U]$. Similarly, $\boldsymbol{HF}^\infty$ is the homology of this same complex, whose base ring is now the field ring of semi-infinite Laurent polynomials $\mathbb{F}[\![U, U^{-1}]$ (rather than $\mathbb{F}[U, U^{-1}]$, as in the construction of the usual HF^∞).

When $\mathfrak{s}$ is a torsion Spin^c structure, recall that $HF^-(Y, \mathfrak{s})$ is equipped with an absolute $\mathbb{Q}$-grading such that U drops grading by 2; see [Ozsváth and Szabó 2003]. We define the i^{th} graded piece of $\boldsymbol{HF}^-(Y, \mathfrak{s})$ to be the same as that of $HF^-(Y, \mathfrak{s})$. This does not quite induce a grading on $\boldsymbol{HF}^-(Y, \mathfrak{s})$ in the usual sense, because $\boldsymbol{HF}^-(Y, \mathfrak{s})$ is not the direct sum of its graded parts, but rather the completion of the direct sum. Still, by a slight abuse of terminology, we will refer to this structure as an absolute $\mathbb{Q}$-grading on $\boldsymbol{HF}^-(Y, \mathfrak{s})$. The same goes for $\boldsymbol{HF}^\infty(Y, \mathfrak{s})$.

When $\mathfrak{s}$ is nontorsion, $HF^-(Y, \mathfrak{s})$ and $HF^\infty(Y, \mathfrak{s})$ admit relative $\mathbb{Z}/2k\mathbb{Z}$-gradings for suitable k (depending on $\mathfrak{s}$). In this case, they induce true relative $\mathbb{Z}/2k\mathbb{Z}$-gradings on $\boldsymbol{HF}^-(Y, \mathfrak{s})$ and $\boldsymbol{HF}^\infty(Y, \mathfrak{s})$, characterized by the fact that each generator (ie intersection point between totally real tori) has the same grading as it does when it is thought of as a generator of $CF^-(Y, \mathfrak{s})$, and, further, that multiplication by U drops the grading by 2.

One can define cobordism maps and mixed invariants as in [Ozsváth and Szabó 2006], using $\boldsymbol{HF}^-$ and $\boldsymbol{HF}^\infty$ rather than HF^- and HF^∞. This new setting is parallel to the one developed in [Kronheimer and Mrowka 2007] in the context of gauge theory. It has the advantage that $\boldsymbol{HF}^-$ and $\boldsymbol{HF}^\infty$ now admit exact triangles analogous to those for $\widehat{HF}$ and HF^+ from [Ozsváth and Szabó 2004c, Section 9]. Further, whereas for the definition of HF^- and HF^∞ one needs to use strongly admissible Heegaard diagrams as in [Ozsváth and Szabó 2004b, Definition 4.10], in order to define $\boldsymbol{HF}^-$ and $\boldsymbol{HF}^\infty$ it suffices to consider weakly admissible diagrams. Indeed, Lemma 4.13 in [Ozsváth and Szabó 2004b] shows that the differentials of $\boldsymbol{CF}^-$ and $\boldsymbol{CF}^\infty$ are finite whenever the respective Heegaard diagrams are weakly admissible.

More generally, whenever we discuss versions of Heegaard and link Floer homology that were defined originally over polynomial rings $\mathbb{F}[U_1, \dots, U_\ell]$, in this paper we use their completions, which are modules over formal power series rings $\mathbb{F}[\![U_1, \dots, U_\ell]\!]$.

One could also define versions of $\widehat{HF}$ and HF^+ using the completed ring; but, since those are generated by complexes on which the action of multiplication by U is nilpotent on each generator, the resulting invariants coincide with the versions defined over $\mathbb{F}[U]$. In particular, in the completed context, we have an exact sequence for any closed, oriented three-manifold Y,

$$\cdots \to \boldsymbol{HF}^-(Y) \to \boldsymbol{HF}^\infty(Y) \to HF^+(Y) \to \cdots, \tag{4}$$

where HF^+ is the Heegaard Floer homology group from [Ozsváth and Szabó 2004b].

Note that, when Y is a three-manifold and $\mathfrak{s}$ is a torsion Spin^c structure on Y, the groups $HF^-(Y, \mathfrak{s})$ and $HF^\infty(Y, \mathfrak{s})$ are determined by $\boldsymbol{HF}^-$ and $\boldsymbol{HF}^\infty$, respectively. Indeed, $\boldsymbol{HF}^-_{\geq i}(Y, \mathfrak{s}) \cong (HF^-)_{\geq i}(Y, \mathfrak{s})$ for any given degree $i \in \mathbb{Q}$, and, since the groups (and their module structure) are determined by their truncations, the claim follows. Similar remarks apply to $\boldsymbol{HF}^\infty$ and HF^∞.

For nontorsion Spin^c structures $\mathfrak{s}$, there is some loss of information when passing from $HF^-(Y, \mathfrak{s})$ and $HF^\infty(Y, \mathfrak{s})$ to their completed analogues. For example, when $Y = S^1 \times S^2$, let h be a generator of $H^2(Y; \mathfrak{s}) \cong \mathbb{Z}$, and $\mathfrak{s}_k$ the Spin^c structure with $c_1(\mathfrak{s}) = 2kh$. Then

$$HF^-(S^1 \times S^2, \mathfrak{s}_k) \cong \mathbb{F}[U]/(U^k - 1), \quad HF^\infty(S^1 \times S^2, \mathfrak{s}_k) \cong \mathbb{F}[U, U^{-1}]/(U^k - 1),$$

so

$$\boldsymbol{HF}^-(S^1 \times S^2, \mathfrak{s}_k) \cong \boldsymbol{HF}^\infty(S^1 \times S^2, \mathfrak{s}_k) = 0.$$

In general, for any nontorsion Spin^c structure $\mathfrak{s}$,

$$\boldsymbol{HF}^\infty(Y, \mathfrak{s}) = 0. \tag{5}$$

Indeed, Lemma 2.3 in [Ozsváth and Szabó 2004d] says that $(1 - U^N) HF^\infty(Y, \mathfrak{u}) = 0$ for some $N \geq 1$. Since $1 - U^N$ is invertible as a power series, after taking the completion we get (5). Consequently, equation (5) combined with exactness in the sequence (4) gives a (grading-preserving) isomorphism

$$\boldsymbol{HF}^-(Y, \mathfrak{s}) \cong HF^+(Y, \mathfrak{s}) \tag{6}$$

for any nontorsion Spin^c structure $\mathfrak{s}$.

3 Generalized Heegaard Floer complexes for links

We define here some complexes associated to a Heegaard diagram for a link. As we shall see in Section 12.1, these are the complexes which govern large surgeries on links.

3.1 Heegaard diagrams

Combining the constructions of [Ozsváth and Szabó 2008a; Manolescu et al. 2009a], define a *multipointed Heegaard diagram* to be data of the form $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$, where:

- Σ is a closed, oriented surface of genus g.
- $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_{g+k-1}\}$ is a collection of disjoint, simple closed curves on Σ which span a g-dimensional lattice of $H_1(\Sigma; \mathbb{Z})$, and hence specify a handlebody U_α; the same goes for $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_{g+k-1}\}$, which specify a handlebody U_β.
- $\boldsymbol{w} = \{w_1, \dots, w_k\}$ and $\boldsymbol{z} = \{z_1, \dots, z_m\}$ (with $k \geq m$) are collections of points on Σ with the following property. Let $\{A_i\}_{i=1}^k$ be the connected components of $\Sigma - \alpha_1 - \dots - \alpha_{g+k-1}$ and $\{B_i\}_{i=1}^k$ be the connected components of $\Sigma - \beta_1 - \dots - \beta_{g+k-1}$. Then there is a permutation σ of $\{1, \dots, m\}$ such that $w_i \in A_i \cap B_i$ for $i = 1, \dots, k$, and $z_i \in A_i \cap B_{\sigma(i)}$ for $i = 1, \dots, m$.

We do not take the orderings of the curves and basepoints to be part of the data of the Heegaard diagram; rather, we just ask for such orderings to exist, so that the conditions above are satisfied.

A Heegaard diagram $\mathcal{H}$ describes a closed, connected, oriented three-manifold $Y = U_\alpha \cup_\Sigma U_\beta$, and an oriented link $\vec{L} \subset Y$ (with $\ell \leq m$ components), obtained as follows. For $i = 1, \dots, m$, we join w_i to z_i inside A_i by an arc which we then push by an isotopy into the handlebody U_α; then we join z_i to $w_{\sigma(i)}$ inside B_i by an arc which we then push into U_β. The union of these arcs (with the induced orientation) is the link $\vec{L}$. We then say that $\mathcal{H}$ is a multipointed Heegaard diagram representing $\vec{L} \subset Y$. Note that the definition we work with here is more general than the notion of a multipointed Heegaard diagram from [Ozsváth and Szabó 2008a], as we allow here for more than two basepoints per link component; moreover, we are allowing for extra basepoints of type w which are not thought of as belonging to a link component. We refer to $w_{m+1}, \dots, w_k$ as *free basepoints*. The set of free basepoints is denoted by $\boldsymbol{w}^{\text{free}}$.

In order to define the chain complexes associated to a Heegaard diagram $\mathcal{H}$ (as below, Section 3.2), we need to require that it is *generic*, ie the alpha and beta curves intersect each other transversely. Further, we should require that it is *admissible* in the sense of [Manolescu et al. 2009a, Definition 2.2]. More precisely:

Definition 3.1 Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ be a multipointed Heegaard diagram.

(a) A *region* in $\mathcal{H}$ is the closure of a connected component of $\Sigma - (\alpha_1 \cup \dots \cup \alpha_{g+k-1} \cup \beta_1 \cup \dots \cup \beta_{g+k-1})$.

(b) A *periodic domain* in $\mathcal{H}$ is a two-chain ϕ on Σ, obtained as a linear combination of regions (with integer coefficients), such that the boundary of ϕ is a linear combination of alpha and beta curves, and the local multiplicity of ϕ at every $w_i \in \boldsymbol{w}$ is zero.

(c) The diagram $\mathcal{H}$ is called *admissible* if every nontrivial periodic domain has some positive local multiplicities and some negative local multiplicities.

Admissibility, as defined in part (c), is equivalent to the requirement that the underlying diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w})$ representing Y be weakly admissible; see [Ozsváth and Szabó 2004b, Definition 4.10]. As mentioned in Section 2, since we will use coefficients in power series rings, there is no need to impose the strong admissibility condition from [loc. cit.].

It is helpful to introduce some terminology for more restrictive classes of Heegaard diagrams:

Definition 3.2 A Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ is called *link-minimal* if $m = \ell$, that is, each link component has only two basepoints.

Definition 3.3 A Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ for a nonempty link is called *minimally pointed* if $k = m = \ell$, that is, each link component has only two basepoints and there are no free basepoints. If the diagram represents an empty link, we call it *minimally pointed* if $k = 1$ and $m = 0$.

Clearly, a minimally pointed diagram is link-minimal. Minimally pointed diagrams were the ones originally studied in [Ozsváth and Szabó 2004b; 2008a].

We can impose an even more restrictive condition:

Definition 3.4 A Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ is called *basic* if it is minimally pointed and, further, for each $i = 1, \dots, \ell$, the basepoints w_i and z_i (which determine one of the link components) lie on each side of a beta curve β_i, and are not separated by any alpha curves.

Basic Heegaard diagrams for knots were constructed in [Ozsváth and Szabó 2004a, Section 2.2], as the doubly pointed diagrams associated to marked Heegaard diagrams. To see that every link admits a basic Heegaard diagram, we can start with a minimally pointed diagram; then, for each $i = 1, \dots, \ell$, we add a 1-handle with one foot near w_i and one near z_i, and choose a path between these feet that does not intersect any alpha curves. Further, we add alpha curves along the chosen paths and the cores of the handles, and beta curves along the cocores of the handles. See Figures 3 and 4 for an example of how this is done in the case of a diagram for the Hopf link.

More generally, we will also need to consider *multipointed Heegaard multidiagrams*: data of the form $(\Sigma, \{\boldsymbol{\eta}^i\}_{i=0}^l, \boldsymbol{w}, \boldsymbol{z})$ such that each $(\Sigma, \boldsymbol{\eta}^i, \boldsymbol{\eta}^j, \boldsymbol{w}, \boldsymbol{z})$ (for $0 \le i \neq j \le l$) is a multipointed Heegaard diagram. A multidiagram is called *generic* if all the curves in it intersect transversely, and no three curves intersect at one point. Finally, we have the following admissibility condition; compare [Ozsváth and Szabó 2004b, Definition 8.8 and Section 8.4.2]:

Definition 3.5 A *multiperiodic domain* ϕ on a multipointed Heegaard multidiagram $(\Sigma, \{\boldsymbol{\eta}^i\}_{i=0}^l, \boldsymbol{w}, \boldsymbol{z})$ is a two-chain on Σ that is a linear combination of the connected components of $\Sigma \setminus \left(\bigcup_i \boldsymbol{\eta}^i\right)$, with integer coefficients, such that $\partial\phi$ is a linear combination of eta curves, and the local multiplicity of ϕ at every $w_i \in \boldsymbol{w}$ is zero. The Heegaard multidiagram is called *admissible* if every multiperiodic domain has some positive local multiplicities and some negative local multiplicities.

3.2 Generalized Floer complexes

Let $\vec{L}$ be an oriented ℓ-component link in an integral homology sphere Y. We denote its components by $L_1, \dots, L_\ell$. Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ be an admissible, generic, multipointed Heegaard diagram for $\vec{L}$. As in the previous subsection, k denotes the number of basepoints in $\boldsymbol{w}$, and m denotes the number of basepoints in $\boldsymbol{z}$, so that $\ell \le m \le k$.

In the introduction we defined sets

$$\mathbb{H}(L)_i = \tfrac{1}{2}\,\mathrm{lk}(L_i, L - L_i) + \mathbb{Z} \subset \mathbb{Q}, \quad \mathbb{H}(L) = \bigtimes_{i=1}^{\ell} \mathbb{H}(L)_i,$$

where lk denotes linking number. Let us also set

$$\overline{\mathbb{H}}(L)_i = \mathbb{H}(L)_i \cup \{-\infty, +\infty\}, \quad \overline{\mathbb{H}}(L) = \bigtimes_{i=1}^{\ell} \overline{\mathbb{H}}(L)_i.$$

Remark 3.6 More invariantly, we could think of $\mathbb{H}(L)$ as an affine lattice over $H_1(Y - L; \mathbb{Z})$; see [Ozsváth and Szabó 2008a]. The latter group is identified with $\mathbb{Z}^\ell$ using the oriented meridians of each component. Furthermore, $\mathbb{H}(L)$ itself is canonically identified with the space of Spin^c structures on Y relative to the link $\vec{L}$; see [ibid., Section 8.1].

The Heegaard diagram determines tori

$$\mathbb{T}_\alpha = \alpha_1 \times \cdots \times \alpha_{g+k-1}, \quad \mathbb{T}_\beta = \beta_1 \times \cdots \times \beta_{g+k-1} \subset \mathrm{Sym}^{g+k-1}(\Sigma).$$

For $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we let $\pi_2(\boldsymbol{x}, \boldsymbol{y})$ be the set of homology classes of Whitney disks from $\boldsymbol{x}$ to $\boldsymbol{y}$ relative to $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$, as in [Ozsváth and Szabó 2004b]. For each homology class of disks $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$, we denote by $n_{w_j}(\phi)$ and $n_{z_j}(\phi) \in \mathbb{Z}$ the multiplicity of w_j and z_j, respectively, in the domain of ϕ. Further, we let $\mu(\phi)$ be the Maslov index of ϕ.

Any intersection point $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ has a Maslov grading $M(\boldsymbol{x}) \in \mathbb{Z}$ and an Alexander multigrading given by

$$A_i(\boldsymbol{x}) \in \mathbb{H}(L)_i, \quad i \in \{1, \dots, \ell\}.$$

Let $\mathbb{W}_i$ and $\mathbb{Z}_i$ be the set of indices for the w's and z's, respectively, belonging to the i^{th} component of the link. We then have

$$A_i(\boldsymbol{x}) - A_i(\boldsymbol{y}) = \sum_{j \in \mathbb{Z}_i} n_{z_j}(\phi) - \sum_{j \in \mathbb{W}_i} n_{w_j}(\phi),$$

where ϕ is any class in $\pi_2(\boldsymbol{x}, \boldsymbol{y})$. Note that this formula determines the Alexander grading up to an overall integer shift. The absolute Alexander grading can be pinned down either by relating intersection points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ to relative Spin^c structures on $(Y, \vec{L})$ via vector fields, or (more simply) by using the symmetry of link Floer homology. See [Ozsváth and Szabó 2008a, Sections 4 and 8; Manolescu et al. 2009a, Section 2.1] for details.

Following [Ozsváth and Szabó 2008a; Manolescu et al. 2009a], we define the *link Floer complex* $\boldsymbol{CFL}^-(\mathcal{H})$ as follows. We let $\boldsymbol{CFL}^-(\mathcal{H})$ be the free module over $\mathcal{R}(\mathcal{H}) = \mathbb{F}[\![U_1, \dots, U_k]\!]$ generated by $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and equipped with the differential

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{n_{w_1}(\phi)} \cdots U_k^{n_{w_k}(\phi)} \boldsymbol{y}. \tag{7}$$

Here, $\mathcal{M}(\phi)$ is the moduli space of pseudoholomorphic curves (solutions to Floer's equation) in the class ϕ, and $\mathbb{R}$ acts on $\mathcal{M}(\phi)$ by translations. Note that $\mathcal{M}(\phi)$ depends on the choice of a suitable path of almost complex structures on the symmetric product. We suppress the almost complex structures from notation for simplicity.

The Maslov grading M produces the homological grading on $\boldsymbol{CFL}^-(\mathcal{H})$, with each U_j decreasing M by two. Furthermore, each Alexander grading A_i defines a filtration on $\boldsymbol{CFL}^-(\mathcal{H})$, with U_j decreasing the filtration level A_i by one when $j \in \mathbb{W}_i$, and leaving A_j constant otherwise.

Of course, if we ignore the Alexander filtrations, then $\boldsymbol{CFL}^-(\mathcal{H})$ is simply a Heegaard Floer complex representing the underlying three-manifold Y, and we could write it as $\boldsymbol{CF}^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w})$ or $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$. However, we use the notation $\boldsymbol{CFL}^-(\mathcal{H})$ when we want to view it as an $\mathbb{H}(L)$-filtered complex.

Definition 3.7 Given $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \mathbb{H}(L)$, we define the *generalized Heegaard Floer complex*

$$\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) = \boldsymbol{A}^-(\mathcal{H}, s_1, \dots, s_\ell) = \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s})$$

to be the subcomplex of $\boldsymbol{CFL}^-(\mathcal{H})$ consisting of the elements with filtration degrees $A_i \le s_i$ for all $i = 1, \dots, \ell$.

Thus, $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ is generated as a product over $\mathbb{F}$ by the elements $U_1^{n_1} \cdots U_k^{n_k} \boldsymbol{x}$ where $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $n_1, \dots, n_k \in \mathbb{Z}$ are such that

$$A_i(U_1^{n_1} \cdots U_k^{n_k} \boldsymbol{x}) = A_i(\boldsymbol{x}) - \sum_{j \in \mathbb{W}_i} n_i \le s_i$$

for all $i = 1, \dots, \ell$.

Remark 3.8 When $L = K$ is a knot and the diagram $\mathcal{H}$ has only two basepoints (one w and one z), the complex $\boldsymbol{CFL}^-(\mathcal{H})$ is the completion of the subcomplex $C\{i \le 0\}$ of the knot Floer complex $CFK^\infty(Y, K)$, in the notation of [Ozsváth and Szabó 2004a]. Further, the generalized Floer complex $\boldsymbol{A}^-(\mathcal{H}, s)$ is the completion of the subcomplex $A_s^- = C\{\max(i, j-s) \le 0\}$. Note that, in [Ozsváth and Szabó 2008b], the formula for integral surgeries along a knot was stated in terms of HF^+. It involved the similar complex $A_s^+ = C\{\max(i, j-s) \ge 0\}$.

Remark 3.9 The restriction to the underlying three-manifold Y being an integral homology sphere is not essential. Indeed, with minor modifications, one can define generalized Heegaard Floer complexes for arbitrary null-homologous links in a three-manifold. Let us explain the only other version of the construction relevant to this paper, namely when $\mathcal{H}$ represents an unlink $\vec{L}$ (of m components) inside a ball in $Y = \#^n(S^1 \times S^2)$. The generators $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are then grouped in equivalence classes according to the Spin^c structures on Y. We focus attention on only those $\boldsymbol{x}$ that lie in the torsion Spin^c structure $\mathfrak{s}_0$. They have an Alexander grading $A(\boldsymbol{x}) \in \mathbb{Z}^m$, and generate a link Floer complex $\boldsymbol{CFL}^-(\mathcal{H}; \mathfrak{s}_0)$. For $\boldsymbol{s} \in \mathbb{Z}^m$, we then define $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ by the same procedure as above, but using only generators in $\mathfrak{s}_0$.

So far we have defined $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ for $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \mathbb{H}(L)$. We can extend this definition to all $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$; that is, we can let some of the s_i be $\pm\infty$. Indeed, if we allow some s_i to be $+\infty$, we can use the same definition as before, and the condition $A_i(U_1^{n_1} \cdots U_k^{n_k} \boldsymbol{x}) \le s_i$ is vacuous. In particular, when all s_i are $+\infty$, the complex $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ is just the completed Heegaard Floer complex

$$\boldsymbol{CF}^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}) = \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}) = \boldsymbol{CF}^-(Y)$$

representing the underlying three-manifold Y.

If we allow some of the s_i to be $-\infty$, we cannot use the same definition, since $A_i(U_1^{n_1} \cdots U_k^{n_k} \boldsymbol{x}) \le -\infty$ is impossible. Instead, we let $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ be generated by those elements with $A_i(U_1^{n_1} \cdots U_k^{n_k} \boldsymbol{x}) \le s_i$ for the values of i with $s_i \neq -\infty$, whereas for the values of i that are $-\infty$, we do not put any condition on A_i, but rather replace the coefficient $U_j^{n_{w_j}(\phi)}$ in (7) with $U_j^{n_{z(w_j)}(\phi)}$, where $z(w_j)$ is the z basepoint that follows w as we go along that link component (with its given orientation). In particular, when all s_i are $-\infty$, the complex $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ is just the completed Heegaard Floer complex

$$\boldsymbol{CF}^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{z} \cup \boldsymbol{w}^{\text{free}}) = \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{z} \cup \boldsymbol{w}^{\text{free}}),$$

again representing the three-manifold Y.

3.3 Heegaard moves

Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ and $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{w}', \boldsymbol{z}')$ be two multipointed Heegaard diagrams representing the same link $\vec{L} \subset Y$. If this is the case, we say that $\mathcal{H}$ and $\mathcal{H}'$ are *equivalent*. It is important to note that we are fixing the three-manifold Y, not just its diffeomorphism class, and we are also fixing the link $\vec{L}$, not just its isotopy class. Similarly, the Heegaard surfaces Σ and Σ' are thought of as embedded in Y.

We say that $\mathcal{H}$ and $\mathcal{H}'$ are related by

(i) a *three-manifold isotopy* if there is a self-diffeomorphism $\phi\colon Y \to Y$ isotopic to the identity such that $\phi(L) = L$, $\phi(\Sigma) = \Sigma'$, and ϕ takes all the curves and basepoints on Σ into the corresponding ones on Σ';

(ii) an *α-curve isotopy* if $\Sigma = \Sigma'$, $\boldsymbol{w} = \boldsymbol{w}'$, $\boldsymbol{z} = \boldsymbol{z}'$, $\boldsymbol{\beta} = \boldsymbol{\beta}'$, and the curves in $\boldsymbol{\alpha}'$ are obtained from those in $\boldsymbol{\alpha}$ by an ambient isotopy of Σ supported away from all the basepoints; a *β-curve isotopy* is similar, with the roles of $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ switched with those of $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$;

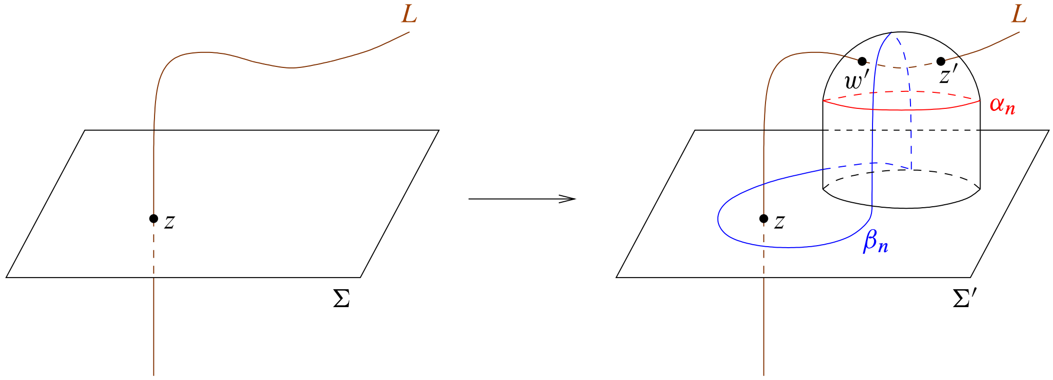

Figure 1: An index zero/three link stabilization. The surface Σ' is obtained from Σ by deleting a disk and adding a cap, ie taking the connected sum with a sphere inside Y. This is the same picture as [Manolescu et al. 2009a, Figure 5], but we emphasize the fact that the construction is done inside the fixed three-manifold Y, with a fixed link L.

(iii) an *α-handleslide* if $\Sigma = \Sigma'$, $\boldsymbol{w} = \boldsymbol{w}'$, $\boldsymbol{z} = \boldsymbol{z}'$, $\boldsymbol{\beta} = \boldsymbol{\beta}'$, a curve $\alpha_i' \in \boldsymbol{\alpha}'$ is obtained by handlesliding the respective curve $\alpha_i \in \boldsymbol{\alpha}$ over another curve in $\boldsymbol{\alpha}$ (with the handleslide being done away from the basepoints), and the other curves in $\boldsymbol{\alpha}'$ are isotopic (by an isotopy away from the basepoints) to the respective curves in $\boldsymbol{\alpha}$; a *β-handleslide* is similar, with the roles of $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ switched with those of $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$;

(iv) an *index one/two stabilization* if Σ' is obtained by taking the connected sum (inside Y) of Σ and a genus one surface with one alpha and beta curve intersecting transversely at a point; the connected sum is done away from the link L and the basepoints, so that $\boldsymbol{w} = \boldsymbol{w}'$, $\boldsymbol{z} = \boldsymbol{z}'$;

(v) a *free index zero/three stabilization* if $\Sigma' = \Sigma$, $\boldsymbol{\alpha}' = \alpha \cup \{\alpha_n\}$, $\boldsymbol{\beta}' = \boldsymbol{\beta} \cup \{\beta'\}$, $\boldsymbol{w}' = \boldsymbol{w} \cup \{w'\}$, $\boldsymbol{z}' = \boldsymbol{z}$, where the new curves α' and β' intersect each other in two points, do not intersect any of the other curves, and both bound disks containing the new free basepoint w'; this kind of stabilization was called simple in [Ozsváth and Szabó 2008a, Section 6.1] (see [ibid., Figure 3] for a picture);

(vi) an *index zero/three link stabilization* if Σ' is obtained from Σ by taking the connected sum (inside Y) with a sphere, in the neighborhood of a z basepoint, as shown in Figure 1; this introduces the new curves α' and β', as well as an additional (w, z) pair, denoted by (w', z') in the picture.

The moves above and their inverses are called *Heegaard moves*. (The inverses of the stabilization moves are called *destabilizations*.) An α-curve isotopy or an α-handleslide (resp. a β-curve isotopy or β-handleslide) is called *admissible* if the corresponding Heegaard triple diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\alpha}', \boldsymbol{w}, \boldsymbol{z})$ (resp. $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \boldsymbol{w}, \boldsymbol{z})$) is admissible. Admissible curve isotopies, admissible handleslides, together with all other moves (i) and (iv)–(vi) and their inverses are called *admissible Heegaard moves*.

Combining the arguments in [Ozsváth and Szabó 2004b; 2008a, Proposition 3.9; Manolescu et al. 2009a, Lemma 2.4], we obtain the following:

Proposition 3.10 *Any two equivalent, generic, admissible, multipointed Heegaard diagrams for $\vec{L}$ can be related by a sequence of admissible Heegaard moves.*

Remark 3.11 In fact, one can check that three-manifold isotopies are redundant in the list (i)–(vi) above, that is, they can be obtained as combinations of the other moves. For example, a three-manifold isotopy away from all the curves and the basepoints can be viewed as an index one/two stabilization, followed by an index one/two destabilization. Nevertheless, we keep (i) in the list because we will refer to this move in some of our future discussions, and it makes the exposition clearer.

We have the following refinement of Proposition 3.10:

Proposition 3.12 *Let $\mathcal{H}$ and $\mathcal{H}'$ be two equivalent, generic, admissible, multipointed Heegaard diagrams representing the same link $\vec{L}$ with ℓ components. Suppose the diagram $\mathcal{H}$ is link-minimal, and that $\mathcal{H}'$ has m basepoints of type z. For $\mathcal{H}'$, let us choose one basepoint of type w on each component of L. The remaining $m-\ell$ basepoints of type w on $\mathcal{H}'$ that are not free are called* ***subsidiary****.*

Then the diagram $\mathcal{H}'$ can be obtained from $\mathcal{H}$ by a sequence of admissible Heegaard moves that includes exactly $m-\ell$ index zero/three link stabilizations, with each of these stabilizations introducing a subsidiary basepoint.

Proof This follows from the arguments in [Manolescu et al. 2009a, proof of Lemma 2.4]. □

For future reference, let us introduce some more terminology. If two Heegaard diagrams $\mathcal{H}$ and $\mathcal{H}'$ have the same underlying Heegaard surface Σ, and their collections of curves $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$ are related by isotopies and handleslides only (supported away from the basepoints), we say that $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$ are *strongly equivalent*.

Definition 3.13 Consider multipointed Heegaard diagrams $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{w},\boldsymbol{z})$ and $\mathcal{H}'=(\Sigma',\boldsymbol{\alpha}',\boldsymbol{\beta}',\boldsymbol{w}',\boldsymbol{z}')$ representing the same link $\vec{L}\subset Y$.

(a) The diagrams $\mathcal{H}$ and $\mathcal{H}'$ are called *strongly equivalent* if $\Sigma=\Sigma'$, $\boldsymbol{w}=\boldsymbol{w}'$, $\boldsymbol{z}=\boldsymbol{z}'$, the curve collections $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ are strongly equivalent, and $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$ are strongly equivalent as well. In other words, $\mathcal{H}$ and $\mathcal{H}'$ should differ by the Heegaard moves (ii) and (iii) only.

(b) We say that $\mathcal{H}$ and $\mathcal{H}'$ differ by a *surface isotopy* if $\Sigma=\Sigma'$ and there is a self-diffeomorphism $\phi\colon\Sigma\to\Sigma$ isotopic to the identity and supported away from the link $\vec{L}$, and such that ϕ takes the curves and basepoints in $\mathcal{H}$ to the corresponding curves and basepoints in $\mathcal{H}'$. If $\mathcal{H}$ and $\mathcal{H}'$ are surface isotopic, we write $\mathcal{H}\cong\mathcal{H}'$.

Note that any surface isotopy can be embedded into an ambient three-manifold isotopy. For this reason, we did not include it in the list of Heegaard moves (i)–(vi). However, surface isotopies will play an essential role in the definition of complete systems of hyperboxes in Section 8.6. Observe also that, by definition, surface isotopies keep the basepoints on $\vec{L}$ fixed. If a surface isotopy kept all the basepoints fixed, it could simply be viewed as a composition of curve isotopies. Therefore, the main role of surface isotopies is to allow a way of moving the free basepoints on Σ, while keeping Σ fixed.

3.4 Invariance

Just as in the usual case of Heegaard Floer homology (or link Floer homology), there is an invariance statement for generalized Floer complexes for links. Before stating it, we need some preliminaries.

We start with a well-known algebraic lemma. We will use this lemma and variants of it throughout the paper, so we include its proof for completeness.

Lemma 3.14 *Let R be any $\mathbb{F}$-algebra, and let C be a free complex over $R[\![U_1]\!]$. Let C_- and C_+ be copies of C. Consider the mapping cone complex*

$$C' = \mathrm{Cone}(C_-[\![U_2]\!] \xrightarrow{U_1-U_2} C_+[\![U_2]\!]).$$

Then the complexes C and C' are chain homotopy equivalent over $R[\![U_1]\!]$.

Proof Let $\boldsymbol{S}$ be a set of free generators for C, over $R[\![U_1]\!]$. For each $\boldsymbol{x} \in \boldsymbol{S}$, we denote by $\boldsymbol{x}_-$ and $\boldsymbol{x}_+$ the corresponding generators of C_- and C_+. Then C' is free over $R[\![U_1, U_2]\!]$, with generators $\boldsymbol{x}_-$ and $\boldsymbol{x}_+$ for each $\boldsymbol{x} \in \boldsymbol{S}$.

Consider the morphisms over $R[\![U_1]\!]$ given by

$$\iota\colon C \to C', \quad \iota(\boldsymbol{x}) = \boldsymbol{x}_+,$$

and

$$\rho\colon C' \to C, \qquad \rho(U_2^n \boldsymbol{x}_-) = 0, \quad \rho(U_2^n \boldsymbol{x}_+) = U_1^n \boldsymbol{x}.$$

We claim that ι and ρ are homotopy inverses. Indeed, $\rho \circ \iota$ is the identity, whereas $\iota \circ \rho$ is chain homotopic to the identity via the homotopy

$$H : C' \to C', \qquad H(U_2^n \boldsymbol{x}_-) = 0, \quad H(U_2^n \boldsymbol{x}_+) = \frac{U_1^n - U_2^n}{U_1 - U_2} \boldsymbol{x}_-.$$

Here, by the fraction $(U_1^n - U_2^n)/(U_1 - U_2)$ we meant the polynomial expression

$$U_1^{n-1} + U_1^{n-2} U_2 + \cdots + U_1 U_2^{n-2} + U_2^{n-1}. \qquad \square$$

Remark 3.15 The map ι is a morphism over $R[\![U_1, U_2]\!]$, which induces an isomorphism on homology. Thus, if we let U_2 act on C just as U_1 does, we find that C and C' are quasi-isomorphic over $R[\![U_1, U_2]\!]$. However, they are not normally chain homotopy equivalent over this bigger ring.

Next, following [Ozsváth and Szabó 2008a, Section 2], if C_* and C'_* are complexes with a filtration by a $\mathbb{Z}^\ell$-affine space such as $\mathbb{H}(L)$, we say that C_* and C'_* are filtered chain homotopy equivalent if there are filtered chain maps $f : C_* \to C'_*$ and $g\colon C'_* \to C_*$ and filtered chain homotopies $H_1 : C_* \to C'_{*+1}$ and $H_2\colon C'_* \to C_{*+1}$ such that $g \circ f - \mathrm{id} = \partial \circ H_1 + H_1 \circ \partial$ and $f \circ g - \mathrm{id} = \partial \circ H_2 + H_2 \circ \partial$.

Recall that we defined $\boldsymbol{CFL}^-(\mathcal{H})$ and $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ as complexes over the ring $\mathcal{R}(\mathcal{H}) = \mathbb{F}[\![U_1, \dots, U_k]\!]$, with one variable for each w basepoint. Assuming L is nonempty, let us also consider the ring

$$\mathcal{R}(L) = \mathbb{F}[\![\widetilde{U}_1, \dots, \widetilde{U}_\ell]\!],$$

where each variable $\widetilde{U}_i$ is associated to a link component $L_i \subseteq L$. We define a *natural inclusion* of $\mathcal{R}(L)$ into $\mathcal{R}(\mathcal{H})$ to be a homomorphism (of unitary rings) that takes each $\widetilde{U}_i$ to one of the U_j variables such that the basepoint w_j is on the link component L_i. Note that when the diagram $\mathcal{H}$ is link-minimal, then there is a unique natural inclusion $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H})$, and that inclusion is in fact an isomorphism.

If L is empty, we let $\mathcal{R}(L) = \mathbb{F}[\![\widetilde{U}]\!]$, and a natural inclusion is a homomorphism that takes $\widetilde{U}$ to any of the U_j variables.

Given a natural inclusion $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H})$, we can view $\boldsymbol{CFL}^-(\mathcal{H})$ and $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ as complexes over $\mathcal{R}(L)$, by restriction of scalars.

We are now ready to state the invariance theorem.

Theorem 3.16 *Let $\mathcal{H}$ and $\mathcal{H}'$ be two generic, admissible, multipointed Heegaard diagram for the same oriented link $\vec{L} \subset Y$. Pick $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$. Pick also natural inclusions $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H})$ and $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H}')$, and view $\boldsymbol{CFL}^-(\mathcal{H})$, $\boldsymbol{CFL}^-(\mathcal{H}')$, $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ and $\boldsymbol{A}^-(\mathcal{H}', \boldsymbol{s})$ as complexes over $\mathcal{R}(L)$, according to those inclusions. Then*:

(a) *The complexes $\boldsymbol{CFL}^-(\mathcal{H})$ and $\boldsymbol{CFL}^-(\mathcal{H}')$ are filtered chain homotopy equivalent over $\mathcal{R}(L)$.*

(b) *The complexes $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ and $\boldsymbol{A}^-(\mathcal{H}', \boldsymbol{s})$ are chain homotopy equivalent over $\mathcal{R}(L)$.*

Note that part (b) is a simple algebraic consequence of part (a). Furthermore, part (a) follows by combining the arguments from [Ozsváth and Szabó 2008a, Theorem 4.4; Manolescu et al. 2009a, Proposition 2.5]. However, some aspects of this proof will be of particular importance for the present paper. We give below an outline of the proof with special emphasis on those aspects.

Beginning of the proof of Theorem 3.16 We need to verify invariance under each of the Heegaard moves (i)–(vi). Let us discuss these moves briefly in turn.

The proof of invariance under curve isotopies and index one/two stabilizations is completely similar to the one in [Ozsváth and Szabó 2004b].

To check invariance under a three-manifold isotopy, one needs to pull back the almost complex structure on the symmetric product under the isotopy, and then use the contractibility of the space of compatible almost complex structures. (This is the same argument as the one used in [ibid., Theorem 6.1].)

Invariance under handleslides can also be proved along the lines in [ibid.], but one has to define suitable (filtered) polygon maps between the link Floer complexes: the details of their construction are explained in Section 3.5.

Let us now discuss the remaining two Heegaard moves (v)–(vi), which change the number of basepoints.

Suppose that we change a diagram $\mathcal{H}$ (with k basepoints of type w) by a free index zero/three stabilization introducing an additional free basepoint w'. Let $\mathcal{H}'$ be the new diagram. By the same argument as in

[Ozsváth and Szabó 2008a, proof of Proposition 6.5], the new link Floer complex $\boldsymbol{CFL}^-(\mathcal{H}')$ is the mapping cone

$$\boldsymbol{CFL}^-(\mathcal{H})[\![U_{k+1}]\!] \xrightarrow{U_j-U_{k+1}} \boldsymbol{CFL}^-(\mathcal{H})[\![U_{k+1}]\!] \tag{8}$$

for some $j \in \{1,\dots,k\}$. The Alexander filtrations on both the domain and the target of (258) are induced from those on $\boldsymbol{CFL}^-(\mathcal{H})$, with U_{k+1} keeping the filtration level fixed (since it corresponds to a free basepoint). By Lemma 3.14, $\boldsymbol{CFL}^-(\mathcal{H}')$ is chain homotopy equivalent to $\boldsymbol{CFL}^-(\mathcal{H})$ over the previous ring $\mathcal{R}(\mathcal{H})$ and hence over $\mathcal{R}(L)$, assuming that the natural inclusion $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H}')$ factors through the natural inclusion into $\mathcal{R}(\mathcal{H})$. Moreover, the maps constructed in the proof of Lemma 3.14 are filtered, so $\boldsymbol{CFL}^-(\mathcal{H})$ and $\boldsymbol{CFL}^-(\mathcal{H}')$ are filtered chain homotopy equivalent.

Given this, $\boldsymbol{A}^-(\mathcal{H}',\boldsymbol{s})$ and $\boldsymbol{A}^-(\mathcal{H},\boldsymbol{s})$ are chain homotopy equivalent. In fact, $\boldsymbol{A}^-(\mathcal{H}',\boldsymbol{s})$ is the mapping cone

$$\boldsymbol{A}^-(\mathcal{H},\boldsymbol{s})[\![U_{k+1}]\!] \xrightarrow{U_j-U_{k+1}} \boldsymbol{A}^-(\mathcal{H},\boldsymbol{s})[\![U_{k+1}]\!]. \tag{9}$$

Finally, suppose that we change a diagram $\mathcal{H}$ by an index zero/three link stabilization introducing an additional (w,z) pair on a link component $L_i \subseteq L$. By the argument in [Manolescu et al. 2009a, Proposition 2.3], the new link Floer complex $\boldsymbol{CFL}^-(\mathcal{H}')$ is also of the form (8), albeit in this case the A_i filtration on the domain of (8) is shifted by -1 from the original one on $\boldsymbol{CFL}^-(\mathcal{H})$. Nevertheless, since U_{k+1} also decreases the A_i filtration level by one, $\boldsymbol{CFL}^-(\mathcal{H}')$ is still filtered chain homotopy equivalent to $\boldsymbol{CFL}^-(\mathcal{H})$. This implies that $\boldsymbol{A}^-(\mathcal{H}',\boldsymbol{s})$ and $\boldsymbol{A}^-(\mathcal{H},\boldsymbol{s})$ are chain homotopy equivalent, although we should point out that in this case $\boldsymbol{A}^-(\mathcal{H}',\boldsymbol{s})$ is not of the form (9). Once again, the chain homotopy equivalence is over $\mathcal{R}(L)$, and we need to assume that there is a factorization of natural inclusions $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H}) \hookrightarrow \mathcal{R}(\mathcal{H}')$.

This completes the proof of invariance over $\mathcal{R}(L)$, with respect to *some* choices of natural inclusions $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H})$ and $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H}')$. To see that invariance holds with respect to any such choices, we can make use of Proposition 3.12; this allows us to choose the basepoints that are introduced when we do the index zero/three link stabilizations. □

We are grateful to András Juhász and Sucharit Sarkar for conversations about the following point:

Remark 3.17 Part (b) of Theorem 3.16 implies that the homology $H_*(\boldsymbol{A}^-(\mathcal{H},\boldsymbol{s}))$ (for basic diagrams $\mathcal{H}$) is an invariant of $\vec{L} \subset Y$ and $\boldsymbol{s}$, up to isomorphism. However, this is not so up to canonical isomorphism. The problem is that, when relating a Heegaard diagram to itself via equivalence moves that include three-manifold isotopies, the basepoints may trace homotopically nontrivial loops during those isotopies. Nevertheless, the generalized Floer homology $H_*(\boldsymbol{A}^-(\mathcal{H},\boldsymbol{s}))$ (for basic diagrams) is an invariant of $\vec{L} \subset Y$, the level $\boldsymbol{s}$ and the basepoints $\boldsymbol{w},\boldsymbol{z} \subset L$, up to canonical isomorphism. This can be proved along the lines of [Ozsváth and Stipsicz 2010, Section 6].

3.5 Polygon maps

In this section we define certain maps between generalized Floer complexes that count holomorphic polygons. For the sake of clarity, we start with the case of triangles, which (along with quadrilaterals) are used in the part of the proof of Theorem 3.16 that deals with invariance under handleslides.

Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ be a generic, admissible, multipointed Heegaard diagram representing a link $\vec{L} \subset Y$, as in Section 3.1. Let $\boldsymbol{\gamma}$ be a new set of attaching curves for Σ, which is strongly equivalent to $\boldsymbol{\beta}$. We assume that the Heegaard multidiagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{z})$ is generic and admissible. Note that the diagram $\mathcal{H}'' = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{z})$ represents $\vec{L} \subset Y$, whereas $\mathcal{H}' = (\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{z})$ represents an unlink in the connected sum of several $S^1 \times S^2$. More precisely, with g, m and k as in Section 3.1, $\mathcal{H}'$ represents the unlink of m components inside $\#^g(S^1 \times S^2)$.

We have link Floer complexes $\boldsymbol{CFL}^-(\mathcal{H})$ and $\boldsymbol{CFL}^-(\mathcal{H}'')$ filtered by $\mathbb{H}(L)$, and (according to Remark 3.9) a complex $\boldsymbol{CFL}^-(\mathcal{H}'; \mathfrak{s}_0)$ filtered by $\mathbb{Z}^m$. We also have generalized Floer complexes $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ and $\boldsymbol{A}^-(\mathcal{H}'', \boldsymbol{s})$ for $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, and $\boldsymbol{A}^-(\mathcal{H}', \boldsymbol{s}')$ for $\boldsymbol{s}' \in \mathbb{Z}^m$.

Note that there is a projection $\mathbb{Z}^m \to \mathbb{Z}^\ell$, which takes $\boldsymbol{s}' = (s'_1, \dots, s'_m)$ to $\bar{\boldsymbol{s}}' = (\bar{s}'_1, \dots, \bar{s}'_\ell)$, where

$$\bar{s}'_i = \sum_{j \in \mathbb{W}_i} s'_j.$$

Since $\overline{\mathbb{H}}(L)$ is an affine space over $\mathbb{Z}^\ell$, for any $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$ and $\boldsymbol{s}' \in \mathbb{Z}^m$ we can make sense of the expression $\boldsymbol{s} + \bar{\boldsymbol{s}}' \in \overline{\mathbb{H}}(L)$.

For $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta, \boldsymbol{x}' \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma, \boldsymbol{y} \in \mathbb{T}_\gamma \cap \mathbb{T}_\alpha$, we let $\pi_2(\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{y})$ be the space of homotopy classes of Whitney triangles connecting $\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{y}$, as in [Ozsváth and Szabó 2004b, Section 8.1.2]. We define a triangle map

$$f_{\alpha\beta\gamma} \colon \boldsymbol{CFL}^-(\mathcal{H}) \otimes \boldsymbol{CFL}^-(\mathcal{H}'; \mathfrak{s}_0) \to \boldsymbol{CFL}^-(\mathcal{H}'') \tag{10}$$

by

$$f_{\alpha\beta\gamma}(\boldsymbol{x} \otimes \boldsymbol{x}') = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{x}', \boldsymbol{y}) \\ \mu(\phi)=0}} \#(\mathcal{M}(\phi)) \cdot U_1^{n_{w_1}(\phi)} \cdots U_k^{n_{w_k}(\phi)} \boldsymbol{y}.$$

Here, $\mathcal{M}(\phi)$ is the moduli space of pseudoholomorphic triangles in the class ϕ, which is required to have Maslov index $\mu(\phi) = 0$. We define a filtration by $\mathbb{H}(L)$ on the domain of $f_{\alpha\beta\gamma}$ by adding up the filtrations on $\boldsymbol{CFL}^-(\mathcal{H})$ and $\boldsymbol{CFL}^-(\mathcal{H}'; \mathfrak{s}_0)$ via $(\boldsymbol{s}, \boldsymbol{s}') \to \boldsymbol{s} + \bar{\boldsymbol{s}}'$. One can check that $f_{\alpha\beta\gamma}$ is a filtered chain map, using arguments as in [Ozsváth and Szabó 2004a, proof of Proposition 8.1].

Hence, for $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \overline{\mathbb{H}}(L)$ and $\boldsymbol{s}' = (s'_1, \dots, s'_\ell) \in \mathbb{Z}^\ell$, we obtain triangle maps

$$\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) \otimes \boldsymbol{A}^-(\mathcal{H}', \boldsymbol{s}') \to \boldsymbol{A}^-(\mathcal{H}'', \boldsymbol{s} + \bar{\boldsymbol{s}}'),$$

which we still denote by $f_{\alpha\beta\gamma}$.

This construction of triangle maps can be extended to more general polygon maps; compare [Ozsváth and Szabó 2005, Section 4.2]. Let $(\Sigma, \{\boldsymbol{\eta}^j\}_{j=0}^r, \boldsymbol{w}, \boldsymbol{z})$ be an admissible, generic, multipointed Heegaard multidiagram such that each $\boldsymbol{\eta}^j$ is a $(g+k-1)$-tuple of attaching circles. For simplicity, we assume that the curve collections $\boldsymbol{\eta}^j$ come in two equivalence classes such that the collections in the same class are strongly equivalent, and the diagrams formed by curve collections in different classes represent a link $\vec{L}$ in an integral homology sphere Y. We can then define Floer groups $\boldsymbol{A}^-(\mathbb{T}_{\eta^i}, \mathbb{T}_{\eta^j}, \boldsymbol{s})$, where each $\boldsymbol{s}$ is either in $\mathbb{Z}^m$ or in $\overline{\mathbb{H}}(L)$. In the former case, we have a well-defined projection $\bar{\boldsymbol{s}} \in \mathbb{Z}^\ell$; in the latter, we use the notation $\bar{\boldsymbol{s}}$ to simply denote $\boldsymbol{s}$.

Suppose $r > 1$, and that $\boldsymbol{\eta}^0$ and $\boldsymbol{\eta}^r$ are in different equivalence classes, so that $(\Sigma, \boldsymbol{\eta}^0, \boldsymbol{\eta}^r, \boldsymbol{w}, \boldsymbol{z})$ represents the link in the integral homology sphere. Then we can define linear maps

$$\begin{gathered} f_{\eta^0,\dots,\eta^r} : \bigotimes_{j=1}^r \boldsymbol{A}^-(\mathbb{T}_{\eta^{j-1}}, \mathbb{T}_{\eta^j}, \boldsymbol{s}_j) \to \boldsymbol{A}^-(\mathbb{T}_{\eta^0}, \mathbb{T}_{\eta^r}, \bar{\boldsymbol{s}}_1 + \dots + \bar{\boldsymbol{s}}_r), \\ f_{\eta^0,\dots,\eta^r}(\boldsymbol{x}_1 \otimes \dots \otimes \boldsymbol{x}_r) = \sum_{\boldsymbol{y} \in \mathbb{T}_{\eta^0} \cap \mathbb{T}_{\eta^r}} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}_1,\dots,\boldsymbol{x}_r,\boldsymbol{y}) \\ \mu(\phi) = 2-r}} \#(\mathcal{M}(\phi)) \cdot U_1^{n_{w_1}(\phi)} \cdots U_k^{n_{w_k}(\phi)} \boldsymbol{y}, \end{gathered} \tag{11}$$

given by counting isolated pseudoholomorphic $(r+1)$-gons in $\mathrm{Sym}^{g+k-1}(\Sigma)$, with edges on $\mathbb{T}_{\eta^0}, \dots, \mathbb{T}_{\eta^r}$, and with specified vertices. Here, the Maslov index $\mu(\phi)$ denotes the expected dimension of the space of pseudoholomorphic polygons in the class ϕ, where the domain (a disk with $r+1$ marked points) has a fixed conformal structure. Since the moduli space of conformal structures on the domain has dimension $r-2$, the Maslov index equals the expected dimension of $\mathcal{M}(\phi)$ minus $r-2$, where $\mathcal{M}(\phi)$ is the space of all pseudoholomorphic $(r+1)$-gons in the class ϕ. We warn the reader that this definition of $\mu(\phi)$ is different from the one in [Ozsváth and Szabó 2005, Section 4.2], where $\mu(\phi)$ was simply the expected dimension of $\mathcal{M}(\phi)$. Our definition of $\mu(\phi)$ coincides with that used in [Sarkar 2011, Section 4]. It has the advantage that it makes the Maslov index additive under the natural juxtaposition maps.

When $r = 1$, by f_{η^0,η^1} we simply mean the differential ∂ for a generalized Floer complex $\boldsymbol{A}^-(\mathbb{T}_{\eta^0}, \mathbb{T}_{\eta^1}, \boldsymbol{s})$.

The maps $f_{\eta^0,\dots,\eta^r}$ can also be defined when all the curve collections $\boldsymbol{\eta}^0, \dots, \boldsymbol{\eta}^r$ are strongly equivalent. The definition is completely analogous, except there is no need for the bars on the values $\boldsymbol{s}$; the image should be in $\boldsymbol{A}^-(\mathbb{T}_{\eta^0}, \mathbb{T}_{\eta^r}, \boldsymbol{s}_1 + \dots + \boldsymbol{s}_r)$, where $\boldsymbol{s}_1 + \dots + \boldsymbol{s}_r \in \mathbb{Z}^m$.

For simplicity, we will ignore the subscripts on the maps $f_{\eta^0,\dots,\eta^r}$, and denote them all by f. The maps f satisfy a generalized associativity property, which can be written as

$$\sum_{0 \le i < j \le r} f(\boldsymbol{x}_1, \dots, \boldsymbol{x}_i, f(\boldsymbol{x}_{i+1}, \dots, \boldsymbol{x}_j), \boldsymbol{x}_{j+1}, \dots, \boldsymbol{x}_r) = 0 \tag{12}$$

for any $\boldsymbol{x}_i \in \boldsymbol{A}^-(\mathbb{T}_{\eta^{j-1}}, \mathbb{T}_{\eta^j}, \boldsymbol{s}_j)$, $j = 1, \dots, r$. Compare [Ozsváth and Szabó 2005, equation (9)].

Completion of the proof of Theorem 3.16 As mentioned in Section 3.3, the triangle and quadrilateral maps defined here are used to prove the part of Theorem 3.16 that deals with invariance under handleslides.

Indeed, one can follow the arguments of [Ozsváth and Szabó 2004b, Section 9] almost verbatim. The only difference is that we use link Floer complexes, and we have to make sure that our maps are filtered. In particular, if $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are curve collections that differ from each other by an elementary handleslide as in [ibid., Figure 9], instead of $CF^-(\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathfrak{s}_0)$ we use $\boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_\gamma, \boldsymbol{0})$, where $\boldsymbol{0} = (0, \dots, 0)$ denotes the zero vector. Since all the 2^{g+k-1} generators are in Alexander grading zero, it is easy to see that $\boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_\gamma, \boldsymbol{0})$ has trivial differential. Therefore, its homology is that of a $(g+k-1)$-dimensional torus, over the base ring $\mathcal{R}(\mathcal{H}) = \mathbb{F}[[U_1, \dots, U_k]]$. When constructing the triangle maps corresponding to this handleslide, we use the top-degree generator $\Theta \in \boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_\gamma, \boldsymbol{0})$. Observe that, when $\boldsymbol{s}' = \boldsymbol{0}$, the map $f_{\alpha\beta\gamma}$ from (10) produces a filtered map

$$\boldsymbol{CFL}^-(\mathcal{H}) \to \boldsymbol{CFL}^-(\mathcal{H}''), \quad \boldsymbol{x} \mapsto f_{\alpha\beta\gamma}(\boldsymbol{x} \otimes \Theta),$$

and hence chain maps from $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ to $\boldsymbol{A}^-(\mathcal{H}'', \boldsymbol{s})$. These can be shown to be chain homotopy equivalences by using (filtered) quadrilateral maps, along the lines of [ibid., Section 9]. It should be noted that in [loc. cit.] it was only proved that the maps induced by handleslides are quasi-isomorphisms. However, the stronger statement that they are chain homotopy equivalences follows from the existence of a chain homotopy between the triangle map given by a small isotopy, and a continuation map; see [Lipshitz 2006, proof of Proposition 11.4; Ozsváth and Stipsicz 2010, proof of Theorem 6.6] and the proof of Lemma 8.14 below. □

3.6 An alternative description

In general, the generalized Floer complex $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ is not free as an $\mathcal{R}(\mathcal{H})$-module. In Section 13 we will construct a resolution of $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$, denoted by $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$. For now, we will focus on the special case when the Heegaard diagram $\mathcal{H}$ is link-minimal (see Definition 3.2). In that case, $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ admits an alternative description as a free module $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$, as follows.

Since we assume that $\mathcal{H}$ is link-minimal, we have $\mathbb{W}_i = \{w_i\}$ and $\mathbb{Z}_i = \{z_i\}$ for each $i = 1, \dots, \ell$. We may also have some free basepoints $w_{\ell+1}, \dots, w_k$. For $i \in \{1, \dots, \ell\}$, $s \in \overline{\mathbb{H}}(L)_i$, and $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$, we set

$$E^i_s(\phi) = \begin{cases} n_{w_i}(\phi) & \text{if } A_i(\boldsymbol{x}) \le s,\ A_i(\boldsymbol{y}) \le s, \\ s - A_i(\boldsymbol{x}) + n_{z_i}(\phi) & \text{if } A_i(\boldsymbol{x}) \le s,\ A_i(\boldsymbol{y}) \ge s, \\ n_{z_i}(\phi) & \text{if } A_i(\boldsymbol{x}) \ge s,\ A_i(\boldsymbol{y}) \ge s, \\ A_i(\boldsymbol{x}) - s + n_{w_i}(\phi) & \text{if } A_i(\boldsymbol{x}) \ge s,\ A_i(\boldsymbol{y}) \le s. \end{cases} \tag{13}$$

Equivalently, we can write

$$\begin{aligned} E^i_s(\phi) &= \max(s - A_i(\boldsymbol{x}), 0) - \max(s - A_i(\boldsymbol{y}), 0) + n_{z_i}(\phi) \\ &= \max(A_i(\boldsymbol{x}) - s, 0) - \max(A_i(\boldsymbol{y}) - s, 0) + n_{w_i}(\phi). \end{aligned} \tag{14}$$

In particular, $E^i_{-\infty}(\phi) = n_{z_i}(\phi)$ and $E^i_{+\infty}(\phi) = n_{w_i}(\phi)$.

Given $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \overline{\mathbb{H}}(L)$, we define the complex $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ as the free module over $\mathcal{R}(\mathcal{H}) = \mathbb{F}[\![U_1, \dots, U_k]\!]$ generated by $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and equipped with the differential

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{E^1_{s_1}(\phi)} \cdots U_\ell^{E^\ell_{s_\ell}(\phi)} \cdot U_{\ell+1}^{n_{w_{\ell+1}}(\phi)} \cdots U_k^{n_{w_k}(\phi)} \boldsymbol{y}. \tag{15}$$

We also put a $\mathbb{Z}$-grading $\mu_{\boldsymbol{s}}$ on $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ by setting

$$\mu_{\boldsymbol{s}}(\boldsymbol{x}) = M(\boldsymbol{x}) - 2 \sum_{i=1}^{\ell} \max(A_i(\boldsymbol{x}) - s_i, 0), \tag{16}$$

at least when none of the values s_i is $-\infty$. When some of the values s_i are $-\infty$, we replace the corresponding expressions $\max(A_i(\boldsymbol{x}) - s_i, 0)$ by $A_i(\boldsymbol{x})$ in (16).

Lemma 3.18 *The complexes $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ and $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ are isomorphic over $\mathcal{R}(\mathcal{H})$.*

Proof When none of the values s_i is $-\infty$, the isomorphism $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) \xrightarrow{\cong} \mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ is given by

$$\boldsymbol{x} \mapsto U_1^{\max(A_1(\boldsymbol{x}) - s_1, 0)} \cdots U_\ell^{\max(A_\ell(\boldsymbol{x}) - s_\ell, 0)} \boldsymbol{x}. \tag{17}$$

When some of the values s_i are $-\infty$, we replace $\max(A_i(\boldsymbol{x}) - s_i, 0)$ by $A_i(\boldsymbol{x})$ in the above formula. □

Remark 3.19 Although we could identify the complexes $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ and $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ using the isomorphism from Lemma 3.18, in practice we will use the notation $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ when we think of the subcomplex of $\boldsymbol{CFL}^-(\mathcal{H})$, and $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ when we think of the alternative description above, as a free complex.

Remark 3.20 When the Heegaard diagram $\mathcal{H}$ is not link-minimal, we have k variables U_j, but only ℓ Alexander gradings. If we set the variables U_j corresponding to the same link component to be equal to each other, then one can define a complex by a formula similar to (15). However, this complex would not be well behaved with Heegaard move (vi) from Section 3.3, the index zero/three link stabilization. Indeed, under this move, the generalized Floer complex $\mathfrak{A}^-_*(\mathcal{H}, \boldsymbol{s})$ would turn into the direct sum $\mathfrak{A}^-_{*+1}(\mathcal{H}, \boldsymbol{s}') \oplus \mathfrak{A}^-_*(\mathcal{H}, \boldsymbol{s})$, where $\boldsymbol{s}'$ differs from $\boldsymbol{s}$ by -1 in the spot corresponding to the link component that is being stabilized.

Recall from Remark 3.9 that the complexes $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ can also be defined when $\mathcal{H}$ represents an unlink in $\#^n(S^1 \times S^2)$. When $\mathcal{H}$ is also link-minimal, we can construct isomorphic free complexes $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ in that context as well. Furthermore, the polygon maps from (11) can be rewritten in terms of the free complexes as

$$\begin{gathered} f_{\eta^0, \dots, \eta^r} \colon \bigotimes_{j=1}^{r} \mathfrak{A}^-(\mathbb{T}_{\eta^{j-1}}, \mathbb{T}_{\eta^j}, \boldsymbol{s}_j) \to \mathfrak{A}^-(\mathbb{T}_{\eta^0}, \mathbb{T}_{\eta^r}, \boldsymbol{s}_1 + \cdots + \boldsymbol{s}_r), \\ f_{\eta^0, \dots, \eta^r}(\boldsymbol{x}_1 \otimes \cdots \otimes \boldsymbol{x}_r) = \sum_{\boldsymbol{y} \in \mathbb{T}_{\eta^0} \cap \mathbb{T}_{\eta^r}} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}_1, \dots, \boldsymbol{x}_r, \boldsymbol{y}) \\ \mu(\phi) = 2 - r}} \#(\mathcal{M}(\phi)) \cdot \boldsymbol{U}^{E_{\boldsymbol{s}_1, \dots, \boldsymbol{s}_r}(\phi)} \boldsymbol{y}, \end{gathered} \tag{18}$$

where we use the notation

$$\boldsymbol{U}^{E_{\boldsymbol{s}_1, \dots, \boldsymbol{s}_r}}(\phi) = U_1^{E^1_{s_{1,1}, \dots, s_{r,1}}(\phi)} \cdots U_\ell^{E^\ell_{s_{1,\ell}, \dots, s_{r,\ell}}(\phi)} \cdot U_{\ell+1}^{n_{w_{\ell+1}}(\phi)} \cdots U_k^{n_{w_k}(\phi)},$$

with

$$\boldsymbol{s}_j = (s_{j,1}, \dots, s_{j,\ell}), \quad j = 1, \dots, r,$$

and

$$\begin{aligned} E^i_{s_{1,i},\dots,s_{r,i}}(\phi) &= \sum_{j=1}^{r} \max(s_{j,i} - A_i(\boldsymbol{x}_j), 0) - \max\Bigl(\sum_{j=1}^{r} s_{j,i} - A_i(\boldsymbol{y}), 0\Bigr) + n_{z_i}(\phi) \\ &= \sum_{j=1}^{r} \max(A_i(\boldsymbol{x}_j) - s_{j,i}, 0) - \max(A_i(\boldsymbol{y}) - \sum_{j=1}^{r} s_{j,i}, 0) + n_{w_i}(\phi). \end{aligned}$$

3.7 Reduction

Suppose that M is a sublink of $L = L_1 \amalg \cdots \amalg L_\ell$. We choose an orientation on M (possibly different from the one induced from $\vec{L}$), and denote the corresponding oriented link by $\vec{M}$. We let $I_+(\vec{L}, \vec{M})$ (resp. $I_-(\vec{L}, \vec{M})$) be the set of indices i such that the component L_i is in M and its orientation induced from $\vec{L}$ is the same as (resp. opposite to) the one induced from $\vec{M}$.

For $i \in \{1, \dots, \ell\}$, we define a projection map $p^{\vec{M}}_i : \overline{\mathbb{H}}(L)_i \to \overline{\mathbb{H}}(L)_i$ by

$$p^{\vec{M}}_i(s) = \begin{cases} +\infty & \text{if } i \in I_+(\vec{L}, \vec{M}), \\ -\infty & \text{if } i \in I_-(\vec{L}, \vec{M}), \\ s & \text{otherwise.} \end{cases}$$

Then, for $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \overline{\mathbb{H}}(L)$, we set

$$p^{\vec{M}}(\boldsymbol{s}) = (p^{\vec{M}}_1(s_1), \dots, p^{\vec{M}}_\ell(s_\ell)).$$

Set $N = L - M$. We define a map

$$\psi^{\vec{M}} : \overline{\mathbb{H}}(L) \to \overline{\mathbb{H}}(N) \tag{19}$$

as follows. The map $\psi^{\vec{M}}$ depends only on the summands $\overline{\mathbb{H}}(L)_i$ of $\overline{\mathbb{H}}(L)$ corresponding to $L_i \subseteq N$. Each of these L_i's appears in N with a (possibly different) index j_i, so there is a corresponding summand $\overline{\mathbb{H}}(N)_{j_i}$ of $\overline{\mathbb{H}}(N)$. We then set

$$\psi^{\vec{M}}_i : \overline{\mathbb{H}}(L)_i \to \overline{\mathbb{H}}(N)_{j_i}, \quad s_i \to s_i - \tfrac{1}{2}\operatorname{lk}(L_i, \vec{M}), \tag{20}$$

where L_i is considered with the orientation induced from L, while $\vec{M}$ is with its own orientation. We then define $\psi^{\vec{M}}$ to be the direct sum of the maps $\psi^{\vec{M}}_i$, precomposed with projection to the relevant factors.

Remark 3.21 If we view $\overline{\mathbb{H}}(L)$ as a lattice over $H_1(S^3 - L)$ — see Remark 3.6 — we can describe the map $\psi^{\vec{M}}$ as

$$\psi^{\vec{M}}(\boldsymbol{s}) = \boldsymbol{s} - \tfrac{1}{2}[\vec{M}],$$

where we denote elements in $H_1(S^3 - L)$ the same as their inclusions into $H_1(S^3 - N)$.

Definition 3.22 Let $\mathcal{H}$ be a multipointed Heegaard diagram representing a link $\vec{L} \subset Y$. Let $M \subseteq L$ be a sublink, with an orientation $\vec{M}$ (not necessarily the one induced from $\vec{L}$). The *reduction* of $\mathcal{H}$ at $\vec{M}$, denoted by $r_{\vec{M}}\mathcal{H}$, is the Heegaard diagram for $\vec{L} - M$ obtained from $\mathcal{H}$ as follows: first, we delete the basepoints z from all components of $M \subseteq L$ oriented the same way in $\vec{L}$ as in $\vec{M}$; second, we delete the basepoints w, and relabel the basepoints z as w, from all components of M oriented the opposite way in $\vec{L}$ as in $\vec{M}$.

Remark 3.23 The reduction of a link-minimal Heegaard diagram is link-minimal. By contrast, the reduction of a minimally pointed diagram is usually not minimally pointed (because we introduce new free basepoints).

Using the interpretation of $\mathbb{H}(L)$ as a space of relative Spin^c structures (see Remark 3.6) it follows from [Ozsváth and Szabó 2008a, Section 3.7] that there is an identification

$$A^-(\mathcal{H}, p^{\vec{M}}(s)) \xrightarrow{\cong} A^-(r_{\vec{M}}(\mathcal{H}), \psi^{\vec{M}}(s)). \tag{21}$$

When the diagram $\mathcal{H}$ is link-minimal, we can write this identification in terms of the free complexes from Section 3.6 as

$$\mathfrak{A}^-(\mathcal{H}, p^{\vec{M}}(s)) \xrightarrow{\cong} \mathfrak{A}^-(r_{\vec{M}}(\mathcal{H}), \psi^{\vec{M}}(s)). \tag{22}$$

3.8 Inclusion maps

Let $\vec{L} \subset Y$ and $M \subseteq L$ be as in the previous subsection. Let also $\mathcal{H}$ be a generic, admissible, multipointed Heegaard diagram for $\vec{L}$. If M is equipped with the orientation induced from $\vec{L}$, then, for $s \in \overline{\mathbb{H}}(L)$, we have a natural inclusion map between subcomplexes of $CFL^-(\mathcal{H})$,

$$\mathcal{I}^M_s : A^-(\mathcal{H}, s) \to A^-(\mathcal{H}, p^M(s)).$$

If $\mathcal{H}$ is link-minimal, we can write $\mathcal{I}^M_s$ in terms of the alternative description of generalized Floer complexes from Section 3.6. In that setting, we have

$$\mathcal{I}^M_s : \mathfrak{A}^-(\mathcal{H}, s) \to \mathfrak{A}^-(\mathcal{H}, p^M(s)), \quad \mathcal{I}^M_s x = \prod_{\substack{i \\ L_i \subseteq L - M}} U_i^{\max(A_i(x) - s_i, 0)} \cdot x. \tag{23}$$

Now suppose that M is equipped with an arbitrary orientation $\vec{M}$, not necessarily the one from $\vec{L}$. We use the notation $I_\pm(\vec{L}, \vec{M})$ from the previous subsection. If $\mathcal{H}$ is link-minimal then, by analogy with (23), we define an inclusion map

$$\mathcal{I}^{\vec{M}}_s : \mathfrak{A}^-(\mathcal{H}, s) \to \mathfrak{A}^-(\mathcal{H}, p^{\vec{M}}(s))$$

by

$$\mathcal{I}^{\vec{M}}_s x = \prod_{i \in I_+(\vec{L}, \vec{M})} U_i^{\max(A_i(x) - s_i, 0)} \cdot \prod_{i \in I_-(\vec{L}, \vec{M})} U_i^{\max(s_i - A_i(x), 0)} \cdot x. \tag{24}$$

Alternatively, in terms of the original description, we have

$$(25) \qquad \mathcal{I}_{\boldsymbol{s}}^{\vec{M}} : \boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathcal{H}, p^{\vec{M}}(\boldsymbol{s})), \quad \mathcal{I}_{\boldsymbol{s}}^{\vec{M}} \boldsymbol{x} = \prod_{i \in I_-(\vec{L}, \vec{M})} U_i^{s_i - A_i(\boldsymbol{x})} \cdot \boldsymbol{x}.$$

Note that the expressions (23), (24) and (25) are well defined only when the exponents of the U_i variables involved there are finite. That is, we need to require that $s_i \neq -\infty$ for all $i \in I_+(\vec{L}, \vec{M})$, and $s_i \neq +\infty$ for all $i \in I_-(\vec{L}, \vec{M})$. These conditions will always be satisfied when we consider inclusion maps in this paper.

Observe that $\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}$ is a chain map, and it shifts the grading (16) by a definite amount:

$$(26) \qquad \mu_{p^{\vec{M}}(\boldsymbol{s})}(\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}(\boldsymbol{x})) = \mu_{\boldsymbol{s}}(\boldsymbol{x}) - 2 \sum_{i \in I_-(\vec{L}, \vec{M})} s_i.$$

Remark 3.24 When $L = K$ is a knot, the maps $\mathcal{I}_s^{\vec{K}}$ and $\mathcal{I}_s^{-\vec{K}}$ correspond to the inclusions v_s^- and h_s^- of A_s^- into the subcomplexes $C(\{i \leq 0\})$ and $C(\{j \leq s\})$, respectively; compare Remark 3.8. In [Ozsváth and Szabó 2008b], there are analogous maps $v_s^+, h_s^+ : A_s^+ \to B^+ \cong CF^+(Y)$.

Remark 3.25 When $\mathcal{H}$ is not link-minimal, there does not seem to be a natural way of defining maps $\mathcal{I}_{\boldsymbol{s}}^{\vec{M}} : \boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathcal{H}, p^{\vec{M}}(\boldsymbol{s}))$ for arbitrary orientations $\vec{M}$. This is one of the reasons why, in Section 13, we will construct resolutions $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ of $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$; the resolutions will come with maps $\mathcal{I}_{\boldsymbol{s}}^{\vec{M}} : \mathfrak{A}^-(\mathcal{H}, \boldsymbol{s}) \to \mathfrak{A}^-(\mathcal{H}, p^{\vec{M}}(\boldsymbol{s}))$ that generalize (24).

4 An overview of the link surgery formula

According to the link surgery formula (Theorem 1.1), the Heegaard Floer homology groups $\boldsymbol{HF}_*^-$ of a link surgery $Y_\Lambda(L)$ can be computed from the surgery complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$. Constructing the surgery complex is rather technical, and the full details will be postponed until Section 9. Until then, the purpose of the present section is to give the reader a basic understanding of the link surgery formula, and to show how it can be applied to compute some examples.

4.1 Preliminaries

As mentioned in the introduction, the input data for the surgery complex is called a complete system of hyperboxes $\mathcal{H}$ for the link $\vec{L}$. In this section, we will focus on the simplest type of complete system, called *basic*. Complete systems of hyperboxes will be defined in Section 8, and basic ones in Section 8.7. For now, it suffices to note that a basic system exists for every link, and is determined by the choice of a single Heegaard diagram $\mathcal{H}^L$ for L. The diagram $\mathcal{H}^L$ has to be basic in the sense of Definition 3.4; that is, it should have no free basepoints, and every link component should intersect the Heegaard surface in exactly two basepoints (one w_i and one z_i), situated on each side of a beta curve.

Apart from $\mathscr{H}^L$, the other diagrams appearing in a basic complete system $\mathscr{H}$ are the reductions

$$\mathscr{H}^{L-M} := r_M(\mathscr{H}^L),$$

obtained from $\mathscr{H}^L$ by deleting the z basepoints on the sublink M; see Section 3.7. (In our case, M is given the orientation induced from $\vec{L}$.) Note that all the diagrams $\mathscr{H}^{L-M}$ are link-minimal.

To the diagrams $\mathscr{H}^{L-M}$ we associate generalized Floer complexes $\mathfrak{A}^-(\mathscr{H}^{L-M}, \boldsymbol{s})$, as in Section 3.6. These are modules over the ground ring

$$\mathscr{R} := \mathscr{R}(\mathscr{H}^L) = \mathbb{F}[\![U_1, \dots, U_\ell]\!].$$

With this in mind, recall from (1) that, as an $\mathscr{R}$-module, the surgery complex is the infinite direct product

$$\mathscr{C}^-(\mathscr{H}, \Lambda) = \bigoplus_{M \subseteq L} \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s})). \tag{27}$$

Here, ψ^M is the map from (19), with M again having the orientation induced from $\vec{L}$.

To simplify the notation somewhat, we denote a typical term in the chain complex (27) by

$$\mathscr{C}^{\varepsilon}_{\boldsymbol{s}} = \mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s})), \tag{28}$$

where $\varepsilon = \varepsilon(M) = (\varepsilon_1, \dots, \varepsilon_\ell) \in \{0,1\}^\ell$ is such that $L_i \subseteq M$ if and only if $\varepsilon_i = 1$.

Furthermore, as noted in (2), the differential on the complex $\mathscr{C}^-(\mathscr{H}, \Lambda)$ is given by

$$\mathscr{D}^-(\boldsymbol{s}, \boldsymbol{x}) = \sum_{N \subseteq L-M} \sum_{\vec{N} \in \Omega(N)} (\boldsymbol{s} + \Lambda_{\vec{L}, \vec{N}}, \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}(\boldsymbol{x}))$$

for $\boldsymbol{s} \in \mathbb{H}(L)$ and $\boldsymbol{x} \in \mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$. In this formula, the maps

$$\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})} : \mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s})) \to \mathfrak{A}^-(\mathscr{H}^{L-M-N}, \psi^{M \cup N}(\boldsymbol{s}))$$

are constructed from polygon maps of the type considered in Section 3.5; see also (18). We shall not give their exact definition here, but we note that when $N = \varnothing$, the map $\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}$ is just the usual differential on $\mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$, counting holomorphic disks.

Let us denote a typical summand in the differential $\mathscr{D}^-$ by

$$\mathscr{D}^{\varepsilon, \varepsilon'}_{\varepsilon^0, \boldsymbol{s}} = \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})} : \mathscr{C}^{\varepsilon^0}_{\boldsymbol{s}} \to \mathscr{C}^{\varepsilon^0 + \varepsilon}_{\boldsymbol{s} + \varepsilon' \cdot \Lambda}, \tag{29}$$

where $\varepsilon^0 = \varepsilon(M)$, $\varepsilon = \varepsilon(N)$, and $\varepsilon' \in \{0,1\}^\ell$ is such that $i \in I_-(\vec{L}, \vec{N})$ if and only if $\varepsilon'_i = 1$. The dot product $\varepsilon' \cdot \Lambda$ denotes the vector $\sum \varepsilon'_i \Lambda_i$. Furthermore, as in Section 3.7, $I_-(\vec{L}, \vec{N})$ is the set of indices i such that the orientation of $L_i \subset L - M$ is different in $\vec{N}$ from the one in $\vec{L}$.

Whenever we drop a subscript or superscript from our notation, we will mean the direct product (or sum, as the case may be) over all possible values of that subscript or superscript. For example, $\mathscr{C}^{\varepsilon} = \prod_{\boldsymbol{s}} \mathscr{C}^{\varepsilon}_{\boldsymbol{s}}$ and $\mathscr{C} = \bigoplus_{\varepsilon} \mathscr{C}^{\varepsilon} = \mathscr{C}^-(\mathscr{H}, \Lambda)$.

In practice, if one wants to compute the Heegaard Floer homology groups $\boldsymbol{HF}^-(Y_\Lambda(L),\mathfrak{u})$ using Theorem 1.1, it is helpful to replace the infinite direct product from (27) with a finite one. We refer to this procedure as *horizontal truncation*. It can be done along the same lines as the corresponding argument in the case of surgery on knots, from [Ozsváth and Szabó 2008b, Section 4.1]. There are many ways of doing horizontal truncation. A general procedure will be given in Section 10.1, but, until then, we will explain a few ad hoc variants in the examples below.

The organization of the rest of Section 4 is as follows. In Section 4.2 we specialize the link surgery formula to the case of knots, and explain the relation to the already-existing knot surgery formula from [ibid.]. In Section 4.3 we explain how horizontal truncation works in the simple case of $+1$ surgery on the unknot, and we also justify our use of direct products rather than direct sums. In Section 4.4, we turn our attention to two-component links, and describe the surgery complex in more detail in that case. Finally, in Section 4.5, we show how the truncation technique can be used to arrive at an explicit computation in the case of surgeries on the Hopf link.

4.2 Knots

Let $\vec{L}=\vec{K}$ be an oriented knot in a homology three-sphere Y. In this case, only two Heegaard diagrams appear in the surgery complex: a basic diagram $\mathcal{H}^K$ for K (with one w basepoint and one z basepoint), and its reduction $\mathcal{H}^\varnothing$, a diagram for Y that is obtained from $\mathcal{H}^K$ by deleting z. The lattice $\mathbb{H}(L)$ is just $\mathbb{Z}$, and $\mathbb{H}(\varnothing)$ consists of the single element 0. Further, let $U=U_1$ be the unique variable in the ground ring $\mathcal{R}=\mathbb{F}[[U]]$, and let $m\in\mathbb{Z}$ be the surgery coefficient that specifies the framing Λ. As an $\mathcal{R}$-module, the surgery complex $\mathcal{C}=\mathcal{C}^-(\mathcal{H},m)$ takes the form

$$\mathcal{C}=\prod_{s\in\mathbb{Z}}\mathcal{C}^0_s\oplus\prod_{s\in\mathbb{Z}}\mathcal{C}^1_s.$$

Here, $\mathcal{C}^0_s$ is the generalized Floer complex $\mathfrak{A}^-(\mathcal{H}^K,s)$, and each $\mathcal{C}^1_s$ is a copy of the complex

$$\mathfrak{A}^-(\mathcal{H}^\varnothing,0)=\boldsymbol{CF}^-(\mathcal{H}^\varnothing),$$

whose homology is $\boldsymbol{HF}^-(Y)$.

With regard to the differential, this is composed of the following terms: $\mathcal{D}^{0,0}_{0,s}$ (which is the usual differential on $\mathcal{C}^0_s$), $\mathcal{D}^{0,0}_{1,s}$ (which is the usual differential on $\mathcal{C}^1_s$), as well as

$$\mathcal{D}^{1,0}_{0,s}=\Phi^K_s\colon\mathcal{C}^0_s\to\mathcal{C}^1_s\quad\text{and}\quad\mathcal{D}^{1,1}_{0,s}=\Phi^{-K}_s\colon\mathcal{C}^0_s\to\mathcal{C}^1_{s+m}.$$

Thus, the complex $\mathcal{C}$ can be viewed as the mapping cone of the map

$$\prod_{s\in\mathbb{Z}}\mathcal{C}^0_s\to\prod_{s\in\mathbb{Z}}\mathcal{C}^1_s,\quad (s,\boldsymbol{x})\mapsto(s,\Phi^K_s)+(s+m,\Phi^{-K}_s). \tag{30}$$

Let us be more specific about what the maps Φ^K_s and Φ^{-K}_s are in this case. The map Φ^K_s is just the inclusion $\mathcal{I}^K_s$ defined in Section 3.8, whose codomain $\mathfrak{A}^-(\mathcal{H}^K,+\infty)$ is identified with $\mathfrak{A}^-(\mathcal{H}^\varnothing,0)$ by (22).

The map Φ_s^{-K} is also constructed from the corresponding inclusion $\mathcal{I}_s^{-K}$, but note that the codomain of $\mathcal{I}_s^{-K}$ is the complex $\mathfrak{A}^-(\mathcal{H}^K, -\infty) \cong \mathfrak{A}^-(r_{-K}(\mathcal{H}^K), 0)$. The diagram $r_{-K}(\mathcal{H}^K)$ is obtained from $\mathcal{H}^K$ by deleting w and relabeling z as the new w. This diagram represents the underlying three-manifold Y, just like $\mathcal{H}^\varnothing$, so it can be related to the latter by a sequence of Heegaard moves. In fact, in a basic system, we use a specific set of Heegaard moves, which will be described in Section 8.7. The Heegaard moves induce a chain homotopy equivalence from $\mathfrak{A}^-(r_{-K}(\mathcal{H}^K))$ to $\mathfrak{A}^-(\mathcal{H}^\varnothing, 0)$. We define Φ_s^{-K} to be the composition of $\mathcal{I}_s^{-K}$ with this homotopy equivalence.

The mapping cone (30) is strongly reminiscent of the mapping cone

$$\bigoplus_{s\in\mathbb{Z}} A_s^+ \to \bigoplus_{s\in\mathbb{Z}} B_s^+, \quad (s, \boldsymbol{x}) \mapsto (s, v_s^+(\boldsymbol{x})) + (s+m, h_s^+(\boldsymbol{x})),$$

that appears in the knot surgery formula from [Ozsváth and Szabó 2008b]. The main difference is that our complex $\mathcal{C}$ computes the completed, minus flavor of Heegaard Floer homology, $\boldsymbol{HF}^-(Y_m(K))$, whereas the one in [loc. cit.] computed the plus flavor $HF^+(Y_m(K))$. As noted in Remark 3.8, the complexes $\mathcal{C}_s^0$ are the completions of $A_s^- = C\{\max(i, j-s) \le 0\}$, in the notation of [loc. cit.]. Similarly, the complexes $\mathcal{C}_s^1 = \boldsymbol{CF}^-(Y)$ are the analogues of $B_s^+ = CF^+(Y)$. Moreover, as explained in Remark 3.24, the maps Φ_s^K and Φ_s^{-K} are the analogues of v_s^+ and h_s^+ from [loc. cit.].

4.3 Remarks on direct products and the unknot

The reader who is familiar with the knot surgery formula from [Ozsváth and Szabó 2008b] may wonder why our definition of the complex $\mathcal{C}$ in (27) involved a direct product rather than a direct sum. The results in [loc. cit.] were phrased in terms of direct sums, but they only applied to HF^+ and $\widehat{HF}$. In the case of $\boldsymbol{HF}^-$, direct sums do not give the right answer. This can be seen even in the simple case of $+1$ surgery on the unknot $\vec{U}$ in S^3.

Specifically, let us consider a genus one Heegaard diagram for the unknot, with one alpha curve and one beta curve intersecting transversely at a single point, and two basepoints w and z. This is a basic Heegaard diagram, to which we can associate a basic complete system of hyperboxes. The complex $\mathcal{C}$ for $+1$ surgery is

$$\mathcal{C} = \bigoplus_{\varepsilon\in\{0,1\}} \prod_{s\in\mathbb{Z}} \mathcal{C}_s^\varepsilon. \tag{31}$$

Here, if we use the alternative description of the complexes $\mathcal{C}_s^0 = \mathfrak{A}^-(\mathcal{H}^K, s)$ from Section 3.6, then each such complex is freely generated by the single intersection point, so it is a copy of $\mathbb{F}[\![U]\!]$. The complexes $\mathcal{C}_s^1$ are also copies of $\mathbb{F}[\![U]\!]$. Let $a_s \in \mathcal{C}_s^0$ and $b_s \in \mathcal{C}_s^1$ be the generators of each piece.

The differential $\mathcal{D}$ is obtained as the sum of maps

$$\Phi_s^{\vec{U}} : \mathcal{C}_s^0 \to \mathcal{C}_s^1 \quad \text{and} \quad \Phi_s^{-\vec{U}} : \mathcal{C}_s^0 \to \mathcal{C}_{s+1}^1.$$

The maps $\Phi_s^{\vec{U}}$ are inclusion maps $\mathcal{I}_s^{\vec{U}}$, which are given by multiplication by a certain power of U; see (24). The maps $\Phi_s^{-\vec{U}}$ are inclusion maps $\mathcal{I}_s^{-\vec{U}}$ postcomposed with a homotopy equivalence, which in this case is the identity. Thus, using (24), we obtain

$$\Phi_s^{\vec{U}} = \begin{cases} 1 & \text{if } s \geq 0, \\ U^{-s} & \text{if } s \leq 0, \end{cases} \qquad \Phi_s^{-\vec{U}} = \begin{cases} U^s & \text{if } s \geq 0, \\ 1 & \text{if } s \leq 0. \end{cases} \tag{32}$$

(Compare [Ozsváth and Szabó 2008b, Section 2.6].) The homology of the complex $\mathcal{C}$ is then isomorphic to $\mathbb{F}[\![U]\!]$, being freely generated by the element

$$\sum_{s \in \mathbb{Z}} U^{|s|(|s|-1)/2} a_s.$$

The Heegaard Floer homology of $+1$ surgery on the unknot is $\boldsymbol{HF}^-(S^3) \cong \mathbb{F}[\![U]\!]$, so Theorem 1.1 gives the right answer. However, if instead of the direct product in (31) we had used a direct sum, the homology of the resulting complex would have been a more complicated $\mathbb{F}[\![U]\!]$-module. Indeed, the map $\mathcal{D}$ would then have nontrivial cokernel in the new $\mathcal{C}^1 = \bigoplus \mathcal{C}_s^1$. The cokernel would be generated as an $\mathbb{F}[\![U]\!]$-module by classes $[b_i]$ for $i \in \mathbb{Z}$, subject to the relations

$$[b_0] = [b_1] = U[b_{-1}] = U[b_2] = U^3[b_{-2}] = U^3[b_3] = \cdots.$$

The basic reason why direct sums are not suitable for the surgery theorem for $\boldsymbol{HF}^-$ is that they do not behave as well with respect to filtrations as direct products do. As we shall see soon, filtrations play an important role in the truncation procedure, which in turn is essential for the proof of Theorem 1.1.

Let us fix some terminology. For us, a *filtration* $\mathcal{F}$ on an $\mathcal{R}$-module $\mathcal{A}$ is a collection of $\mathcal{R}$-submodules $\{\mathcal{F}^i(\mathcal{A}) \mid i \in \mathbb{Z}\}$ of $\mathcal{A}$ such that $\mathcal{F}^i(\mathcal{A}) \subseteq \mathcal{F}^j(\mathcal{A})$ for all $i \leq j$. A filtration is called *bounded above* if $\mathcal{F}^i(\mathcal{A}) = \mathcal{A}$ for $i \gg 0$, and *bounded below* if $\mathcal{F}^i(\mathcal{A}) = 0$ for $i \ll 0$. A filtration is called *bounded* if it is both bounded above and bounded below.

If $\mathcal{A}$ is equipped with a differential ∂ that turns it into a chain complex, we say that the chain complex $(\mathcal{A}, \partial)$ is filtered by $\mathcal{F}$ if ∂ preserves each submodule $\mathcal{F}^i(\mathcal{A})$. The *associated graded* complex $\mathrm{gr}_{\mathcal{F}} \mathcal{A}$ is defined as

$$\mathrm{gr}_{\mathcal{F}}(\mathcal{A}) = \bigoplus_{i \in \mathbb{Z}} (\mathcal{F}^i(\mathcal{A})/\mathcal{F}^{i-1}(\mathcal{A})), \tag{33}$$

equipped with the differential induced from $\mathcal{F}$.

If $\mathcal{F}$ is a bounded filtration on a chain complex $(\mathcal{A}, \partial)$, a standard result from homological algebra says that, if $\mathrm{gr}_{\mathcal{F}}(\mathcal{A})$ is acyclic, then $\mathcal{A}$ is acyclic as well. (Note that this can fail for filtrations that are not bounded.)

A standard way to construct bounded filtrations is as follows. If $\mathcal{A}$ is freely generated over $\mathcal{R}$ by a collection of generators G, a bounded map $\mathcal{F}\colon G \to \mathbb{Z}$ defines a bounded filtration on $\mathcal{A}$ by letting $\mathcal{F}^i(\mathcal{A})$ be the submodule generated by the elements $g \in G$ with $\mathcal{F}(g) \leq i$.

Suppose now that we have a direct product of $\mathscr{R}$-modules

$$\mathscr{A} = \prod_{s \in S} \mathscr{A}_s,$$

indexed over a countable set S. Suppose further that each $\mathscr{A}_s$ is a free module over $\mathscr{R}$ with a set of generators G_s. Assume that $\mathscr{A}$ is equipped with a differential ∂. (Typically, in our examples, each G_s is finite, and the differential ∂ is defined on each term so that it is locally finite, as in Definition 9.5.)

In this situation, an assignment

$$\mathscr{F}\colon \bigcup_{s \in S} G_s \to \mathbb{Z}$$

specifies bounded filtrations on each $\mathscr{A}_s$. Together these produce a filtration on $\mathscr{A}$ given by

$$\mathscr{F}^i(\mathscr{A}) = \prod_{s \in S} \mathscr{F}^i(\mathscr{A}_s).$$

This filtration on $\mathscr{A}$ is generally neither bounded above nor bounded below. It is bounded above provided that there exists $i \gg 0$ such that $\mathscr{F}^i(\mathscr{A}_s) = \mathscr{A}_s$ for all s, that is, if $\mathscr{F}(g) \le i$ for all $g \in G_s$ and $s \in S$. We have the following:

Lemma 4.1 *Consider a module $\mathscr{A} = \prod_{s \in S} \mathscr{A}_s$, where each $\mathscr{A}_s$ is freely generated over $\mathscr{R}$ by a set of generators G_s. Suppose $\mathscr{F}\colon \bigcup_{s \in S} G_s \to \mathbb{Z}$ defines a filtration on $\mathscr{A}$ that is bounded above. Further, suppose $\mathscr{A}$ is equipped with a differential ∂, and that the associated graded complex $\mathrm{gr}_{\mathscr{F}}(\mathscr{A})$ is acyclic. Then $\mathscr{A}$ itself is acyclic.*

Proof Consider the spectral sequence associated to $\mathscr{F}$, whose E^1 term is $H_*(\mathrm{gr}_{\mathscr{F}}(\mathscr{A})) = 0$. According to [Weibel 1994, Theorem 5.5.10], the spectral sequence converges to $H_*(\mathscr{A})$ if the filtration is complete, exhaustive, and the spectral sequence is regular. Completeness of the filtration means that $\mathscr{A} = \varprojlim \mathscr{A}/\mathscr{F}^i\mathscr{A}$, which is true because $\mathscr{A}$ is constructed as a direct product. (Note that for a complete filtration we have $\bigcap_i \mathscr{F}^i\mathscr{A} = 0$.) Exhaustiveness is a weaker condition than being bounded above. Regularity (as defined in [ibid., Definition 5.2.10]) is automatic when the E^1 term is zero, because the higher differentials have to be zero as well. □

One can apply Lemma 4.1 to calculate $H_*(\mathscr{C}^-(\vec{U}, 1))$, where $\mathscr{C} = \mathscr{C}^-(\vec{U}, 1)$ is the surgery complex for $+1$ surgery on the unknot, considered above. (Of course, this calculation can also be done directly.) As a module, the complex $\mathscr{C}$ is the direct product of all $\mathscr{C}_s^\varepsilon$ over $\varepsilon \in \{0, 1\}$ and $s \in \mathbb{Z}$. Each term in the direct product is freely generated by an element a_s or b_s. We have two filtrations $\mathscr{F}_0$ and $\mathscr{F}_1$ on $\mathscr{C}$ defined by

$$\mathscr{F}_0(a_s) = \mathscr{F}_0(b_s) = -s, \quad \mathscr{F}_1(a_s) = \mathscr{F}_1(b_{s+1}) = s.$$

The map $\Phi_s^{\vec{U}}$ preserves the $\mathscr{F}_0$-grading and decreases $\mathscr{F}_1$ by one, whereas $\Phi_s^{-\vec{U}}$ decreases $\mathscr{F}_0$ by one and preserves $\mathscr{F}_1$. Neither $\mathscr{F}_0$ nor $\mathscr{F}_1$ are bounded above. However, we can consider the subcomplexes of $\mathscr{C}$

$$\mathscr{C}_{>0} = \prod_{s>0} (\mathscr{C}_s^0 \oplus \mathscr{C}_s^1) \quad \text{and} \quad \mathscr{C}_{<0} = \prod_{s<0} (\mathscr{C}_s^0 \oplus \mathscr{C}_{s+1}^1).$$

Then the restriction of $\mathcal{F}_0$ to $\mathcal{C}_{>0}$ and the restriction of $\mathcal{F}_1$ to $\mathcal{C}_{<0}$ are both bounded above. Further, in the associated graded of these subcomplexes the differential cancels out all the terms in pairs. Applying Lemma 4.1 gives

$$H_*(\mathcal{C}_{>0}) \cong H_*(\mathcal{C}_{<0}) \cong 0.$$

The quotient complex $\mathcal{C}/(\mathcal{C}_{>0} \oplus \mathcal{C}_{<0})$ is simply $\mathcal{C}_0^0 \cong \mathbb{F}[\![U]\!]$ with a trivial differential. From the corresponding long exact sequence, we get

$$H_*(\mathcal{C}) \cong H_*(\mathcal{C}_0^0) \cong \mathbb{F}[\![U]\!],$$

as expected.

Although Lemma 4.1 can be successfully used to compute the homology of the complex associated to $+1$ surgery on the unknot, if we try to do -1 surgery instead, we will need to work with filtrations that are bounded below, rather than above. As the following example shows, the exact analogue of Lemma 4.1 fails for bounded below filtrations:

Example 4.2 Let $\mathcal{A} = \prod_{s \geq 0} \mathcal{A}_s$, where each $\mathcal{A}_s$ is a free module over $\mathbb{F}[\![U]\!]$ of rank two, with generators a_s and b_s. We equip $\mathcal{A}$ with the locally finite differential ∂ given by $\partial a_0 = b_0$, $\partial a_s = b_s + b_{s-1}$ for $s > 0$, and $\partial b_s = 0$ for all s. Let $\mathcal{F}$ be the filtration on $\mathcal{A}$ defined by $\mathcal{F}(a_s) = \mathcal{F}(b_s) = s$. Then $\mathcal{F}$ is bounded below but not bounded above, and $\operatorname{gr}_{\mathcal{F}} \mathcal{A}$ is acyclic. However, $\mathcal{A}$ itself is not acyclic, because $\sum_s a_s$ is in the kernel of its differential.

Nevertheless, we can obtain the desired result by imposing an additional condition on our filtration:

Definition 4.3 Suppose we have a module $\mathcal{A} = \prod_{s \in S} \mathcal{A}_s$, where each $\mathcal{A}_s$ is freely generated over $\mathcal{R} = \mathbb{F}[\![U_1, \dots, U_\ell]\!]$ by a finite set of generators G_s. Suppose $\mathcal{A}$ is equipped with a locally finite differential ∂, and with a filtration induced by an assignment $\mathcal{F}\colon \bigcup_{s \in S} G_s \to \mathbb{Z}$. For $i \in \mathbb{Z}$, let $\mathcal{A}^{[i]} \cong \mathcal{F}^i(\mathcal{A})/\mathcal{F}^{i-1}(\mathcal{A})$ be the free $\mathcal{R}$-module generated by those $g \in \bigcup_s G_s$ with $\mathcal{F}(g) = i$. Then we say that the filtration $\mathcal{F}$ is *U-tame* if there exists $k \in \mathbb{Z}$ such that, in the restriction of ∂ to $\mathcal{A}/\mathcal{F}^k\mathcal{A}$, all terms that drop the filtration level contain at least a nontrivial power of U_j for some j, that is, for all $i > k$,

$$\partial(\mathcal{A}^{[i]}) \subset \mathcal{A}^{[i]} \oplus ((U_1, \dots, U_\ell) \cdot \mathcal{F}^{i-1}(\mathcal{A})).$$

Lemma 4.4 *Let $\mathcal{A}$ be a chain complex as in Definition 4.3, with a U-tame filtration $\mathcal{F}$. Suppose that the associated graded complex $\operatorname{gr}_{\mathcal{F}}(\mathcal{A})$ is acyclic. Then $\mathcal{A}$ itself is acyclic.*

Proof Let $k \in \mathbb{Z}$ be as in Definition 4.3. The complex $\mathcal{F}^k(\mathcal{A})$ (with the restriction of the filtration $\mathcal{F}$) satisfies the hypotheses of Lemma 4.1, so it is acyclic. Thus, it suffices to prove that the quotient complex $\mathcal{A}' = \mathcal{A}/\mathcal{F}^k(\mathcal{A})$ is acyclic.

If we view $\mathcal{A}'$ as a complex over $\mathbb{F}$ (rather than over $\mathcal{R}$), we can equip it with a second filtration $\mathcal{U}$, given by the total degree of the U_j variables. Precisely, if $g \in \bigcup_s G_s$ is a generator with $\mathcal{F}(g_s) > k$, we set

$$\mathcal{U}(U_1^{m_1} \cdots U_\ell^{m_\ell} \cdot g) = -(m_1 + \cdots + m_\ell).$$

The original $\mathscr{R}$-structure on $\mathscr{A}'$ implies that $\mathscr{U}$ is indeed a filtration on $\mathscr{A}'$ (viewed as an $\mathbb{F}$-complex). The U-tameness assumption implies that the differential on the associated graded $\operatorname{gr}_{\mathscr{U}} \mathscr{A}'$ preserves each $\mathscr{F}$-filtration level $\mathscr{A}^{[i]}$ for $i > k$. Thus, $\operatorname{gr}_{\mathscr{U}} \mathscr{A}'$ splits as a direct sum of complexes of the form $\operatorname{gr}_{\mathscr{U}} \mathscr{A}^{[i]}$ for $i > k$.

We claim that each $\operatorname{gr}_{\mathscr{U}} \mathscr{A}^{[i]}$ is acyclic for $i > k$. We know that $\mathscr{A}^{[i]}$ itself (being a summand of $\operatorname{gr}_{\mathscr{F}}(\mathscr{A})$) is acyclic. The short exact sequence

$$0 \to \mathscr{A}^{[i]} \xrightarrow{U_1} \mathscr{A}^{[i]} \to \mathscr{A}^{[i]}/U_1\mathscr{A}^{[i]} \to 0$$

induces a long exact sequence in homology, showing that $\mathscr{A}^{[i]}/U_1\mathscr{A}^{[i]}$ is acyclic. Iterating this argument shows $\mathscr{A}^{[i]}/(U_1, \dots, U_\ell)\mathscr{A}^{[i]}$ is acyclic. Since $\operatorname{gr}_{\mathscr{U}} \mathscr{A}^{[i]}$ is just a direct sum of complexes isomorphic to $\mathscr{A}^{[i]}/(U_1, \dots, U_\ell)\mathscr{A}^{[i]}$, it must be acyclic as well. We deduce from this that $\operatorname{gr}_{\mathscr{U}} \mathscr{A}'$ is acyclic.

It now follows that $\mathscr{A}'$ is acyclic, by Lemma 4.1 applied to the filtration $\mathscr{U}$. (Note that $\mathscr{U}$ is bounded above by 0.) Consequently $\mathscr{A}$ is also acyclic. $\square$

As an application of Lemma 4.4, let us compute the homology of the complex $\mathscr{C} = \mathscr{C}^-(\vec{U}, -1)$, associated to -1 surgery on the unknot. Just as in the case of $+1$ surgery, our new complex $\mathscr{C}$ is the direct product (as a module) of all $\mathscr{C}_s^\varepsilon$, over $\varepsilon \in \{0, 1\}$ and $s \in \mathbb{Z}$. Each term in the direct product is freely generated by one element, denoted by a_s (for $\mathscr{C}_s^0$) and b_s (for $\mathscr{C}_s^1$). The maps $\Phi_s^{\vec{U}}$ and $\Phi_s^{-\vec{U}}$ are still given by (32), but now $\Phi_s^{-\vec{U}}$ takes $\mathscr{C}_s^0$ to $\mathscr{C}_{s-1}^1$. Explicitly, we have

$$\partial a_s = \begin{cases} b_s + U^s b_{s-1} & \text{if } s \geq 0, \\ U^{-s} b_s + b_{s-1} & \text{if } s \leq 0. \end{cases}, \qquad \partial b_s = 0.$$

Consider the quotient complexes of $\mathscr{C}$

$$\mathscr{C}_{>0} = \prod_{s>0} (\mathscr{C}_s^0 \oplus \mathscr{C}_s^1) \quad \text{and} \quad \mathscr{C}_{<0} = \prod_{s<0} (\mathscr{C}_{s+1}^0 \oplus \mathscr{C}_s^1).$$

On $\mathscr{C}_{>0}$, we define a filtration $\mathscr{F}_0$ by

$$\mathscr{F}_0(a_s) = \mathscr{F}_0(b_s) = s.$$

This filtration is not bounded above, but it is U-tame. Since its associated graded is acyclic, $\mathscr{C}_{>0}$ is acyclic by Lemma 4.4.

Similarly, on $\mathscr{C}_{<0}$ we have a U-tame filtration $\mathscr{F}_1$ given by the assignment

$$\mathscr{F}_1(a_s) = \mathscr{F}_1(b_{s-1}) = -s.$$

Its associated graded is acyclic, so $\mathscr{C}_{<0}$ is acyclic as well. We deduce that $\mathscr{C}$ is quasi-isomorphic to its subcomplex $\mathscr{C}_0^1$, which has homology $\mathbb{F}[\![U]\!]$.

The calculations above (for $+1$ and -1 surgery on the unknot) are the simplest examples of horizontal truncation, and serve as models for the general case. In Section 4.5 we will explain horizontal truncation in a more complicated example, that of the Hopf link.

4.4 Two-component links

We now consider the case of a two-component link $\vec{L} = L_1 \cup L_2 \subset Y$. We are interested in surgery on L with surgery coefficients p_1 on L_1 and p_2 on L_2. The framing Λ is given by the matrix

$$\Lambda = \begin{pmatrix} p_1 & c \\ c & p_2 \end{pmatrix},$$

where c is the linking number between L_1 and L_2. We denote by Λ_1 and Λ_2 the two rows of Λ.

The affine lattice $\mathbb{H}(L)$ associated to L is $\mathbb{Z} \times \mathbb{Z}$ if c is even, and $\left(\frac{1}{2} + \mathbb{Z}\right) \times \left(\frac{1}{2} + \mathbb{Z}\right)$ if c is odd. The lattices $\mathbb{H}(L_1)$ and $\mathbb{H}(L_2)$ are copies of $\mathbb{Z}$, and $\mathbb{H}(\varnothing) = 0$.

A basic complete system $\mathbb{H}$ for $\vec{L}$ involves a basic diagram $\mathscr{H}^L$ for L (with four basepoints: w_1 and z_1 on L_1, and w_2 and z_2 on L_2), as well as its reductions $\mathscr{H}^{L_1}$ (obtained from $\mathbb{H}^L$ by deleting z_1), $\mathscr{H}^{L_2}$ (obtained from $\mathscr{H}^L$ by deleting z_2), and $\mathscr{H}^{\varnothing}$ (obtained by deleting both z_1 and z_2).

The surgery complex $\mathscr{C}$ is a "double mapping cone" of the form

$$\begin{tikzcd} \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathscr{C}_{\boldsymbol{s}}^{00} \arrow[r] \arrow[d] \arrow[dr] & \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathscr{C}_{\boldsymbol{s}}^{10} \arrow[d] \\ \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathscr{C}_{\boldsymbol{s}}^{01} \arrow[r] & \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathscr{C}_{\boldsymbol{s}}^{11} \end{tikzcd} \tag{34}$$

Here, the horizontal arrows consist of maps of the form $\Phi_{\boldsymbol{s}}^{\pm L_1}$, the vertical arrows of maps $\Phi_{\boldsymbol{s}}^{\pm L_2}$, and the diagonal of maps $\Phi_{\boldsymbol{s}}^{\pm L_1 \cup \pm L_2}$ and $\Phi_{\boldsymbol{s}}^{\pm L_1 \cup \mp L_2}$. Here, $+L_i$ (resp. $-L_i$) means L_i with the orientation induced from (resp. opposite to) $\vec{L}$. There are also maps $\Phi_{\boldsymbol{s}}^{\varnothing}$ (not shown in the diagram above), which represent the differentials on each complex $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon}$.

The values of $\boldsymbol{s}$ in the targets of each map are shifted by an amount depending on the type of map and the framing Λ. Precisely, whenever we have a negatively oriented component $-\vec{L}_i$ in the superscript of a map $\Phi_{\boldsymbol{s}}$, we add the vector $\Lambda_i \in \mathbb{Z} \times \mathbb{Z}$ to $\boldsymbol{s}$. Note that all the maps preserve the equivalence class of $\boldsymbol{s}$ in the quotient

$$\mathbb{H}(L)/\langle \Lambda_1, \Lambda_2 \rangle.$$

This quotient can be naturally identified with the space of Spin^c structures on the surgered manifold $Y_\Lambda(L)$; see Remark 3.6 and [Ozsváth and Szabó 2008a, Section 3.7]. Therefore, the complex $\mathscr{C}$ splits as a direct sum, according to these Spin^c structures. In fact, there is a more refined version of Theorem 1.1, which establishes an isomorphism between the homology of each of these summands and the corresponding $\boldsymbol{HF}^-(Y_\Lambda(L), [\boldsymbol{s}])$; see Theorem 9.6.

Let us now say a few words about the maps in (34). Those of the form $\Phi_{\boldsymbol{s}}^{L_i}$ are given by inclusions $\mathscr{I}_{\boldsymbol{s}}^{L_i}$, as in Section 3.8. Those of the form $\Phi_{\boldsymbol{s}}^{-L_i}$ are the compositions of inclusions $\mathscr{I}_{\boldsymbol{s}}^{-L_i}$ with certain "descent maps" $D_{p^{L_i}(\boldsymbol{s})}^{L_i}$. The general definition of descent maps will be given in (99), but in the case at hand they are chain homotopy equivalences induced by Heegaard moves between the respective Heegaard diagrams (just as in the case of knots, discussed in Section 4.2).

Finally, the diagonal maps in (34) are chain homotopies between the corresponding compositions of vertical and horizontal maps. For example, $\Phi_{\boldsymbol{s}}^{L_1 \cup -L_2}$ is a chain homotopy between $\Phi^{L_1}_{\psi^{-L_2}(\boldsymbol{s})} \circ \Phi_{\boldsymbol{s}}^{-L_2}$ and $\Phi^{-L_2}_{\psi^{L_1}(\boldsymbol{s})} \circ \Phi_{\boldsymbol{s}}^{L_1}$. These chain homotopies are constructed as compositions of inclusion and descent maps. In the case of a basic system for two component links, such as the one considered here, it turns out that most of the descent maps, and hence the corresponding diagonal maps in (34), are zero.[1] Specifically, the only diagonal map that may be nontrivial is $\Phi_{\boldsymbol{s}}^{-L_1 \cup -L_2}$. This is the composition of the inclusion $\mathcal{I}_{\boldsymbol{s}}^{-L_1 \cup -L_2}$ with a descent map $D_{\boldsymbol{s}}^{-L_1 \cup -L_2}$.

The descent map $D_{\boldsymbol{s}}^{-L_1 \cup -L_2}$ is a chain homotopy between two maps induced by Heegaard moves between the diagrams $r_{-L}(\mathcal{H}^L)$ and $\mathcal{H}^{\varnothing}$ (that is, between the diagrams obtained from $\mathcal{H}^L$ by deleting the w basepoints or deleting the z basepoints). While the maps induced by Heegaard moves usually involve counts of holomorphic triangles, the chain homotopy $D_{\boldsymbol{s}}^{-L_1 \cup -L_2}$ involves also counts of holomorphic quadrilaterals. In practice, these counts are harder to get a handle on, but in some cases (such as the example discussed in the next section) we can simplify the complex $\mathcal{C}$ so that we do not have to compute $D_{\boldsymbol{s}}^{-L_1 \cup -L_2}$ at all.

4.5 Surgeries on the Hopf link

Let $\vec{L} = L_1 \cup L_2 \subset S^3$ be the positive Hopf link shown in Figure 2. Its link Floer homology was computed in [Ozsváth and Szabó 2008a, Section 12]. Let p_1 and p_2 be two integers with $p_1 p_2 \neq \pm 1$. Then (p_1, p_2) surgery on $\vec{L}$ produces the lens space $L(p_1 p_2 - 1, p_1)$, which is a rational homology sphere admitting $p_1 p_2 - 1$ different Spin^c structures. For any such Spin^c structure $\mathfrak{s}$, the corresponding Heegaard Floer homology group is $\boldsymbol{HF}^-(S^3_{p_1,p_2}(\vec{L}), \mathfrak{s}) \cong \mathbb{F}[\![U]\!]$ as a relatively $\mathbb{Z}$-graded module; see [Ozsváth and Szabó 2004c, Proposition 3.1]. In this section we show how Theorem 1.1 can be used to recover this calculation. We assume for simplicity that $p_1, p_2 \geq 2$.

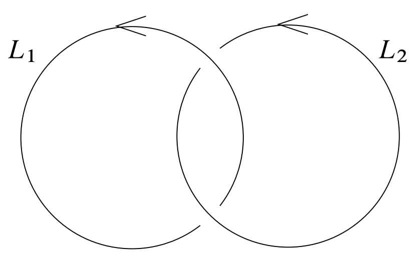

Figure 2: The positive Hopf link.

[1]The reason why those descent maps vanish is because, in the terminology of Section 5 below, they are obtained by compressing hyperboxes of size $\boldsymbol{d} = (d_1, d_2)$, where at least one of the d_i is zero. See (41) and (100).

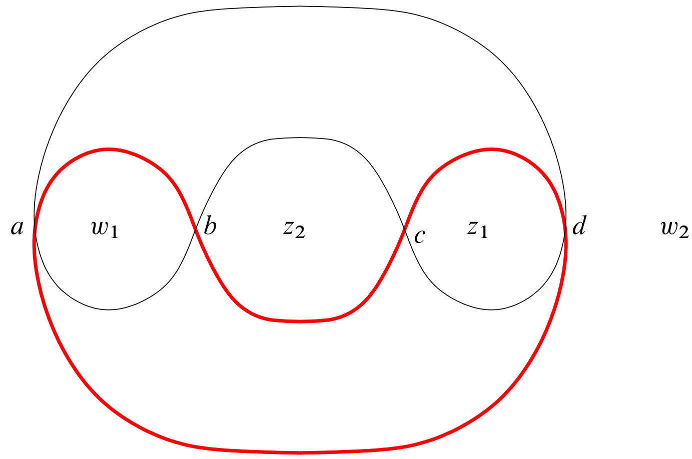

Figure 3: A Heegaard diagram for the Hopf link. The thicker (red) curve is α, while the thinner (black) curve is β.

We use a basic system of hyperboxes for the Hopf link $\vec{L}$. To construct it, we start with the genus zero Heegaard diagram for $\vec{L}$ from Figure 3. We then stabilize it (and do some handleslides) to end up with the basic Heegaard diagram pictured in Figure 4, in which w_i and z_i for $i = 1, 2$ are separated by the new beta curves β_i. The old alpha and beta curves are denoted by α_3 and β_3, respectively.

In the genus zero Heegaard diagram in Figure 3, the intersection $\alpha \cap \beta$ consists of four points, denoted by a, b, c and d. The index one holomorphic disks correspond to bigons. There are twelve such bigons:

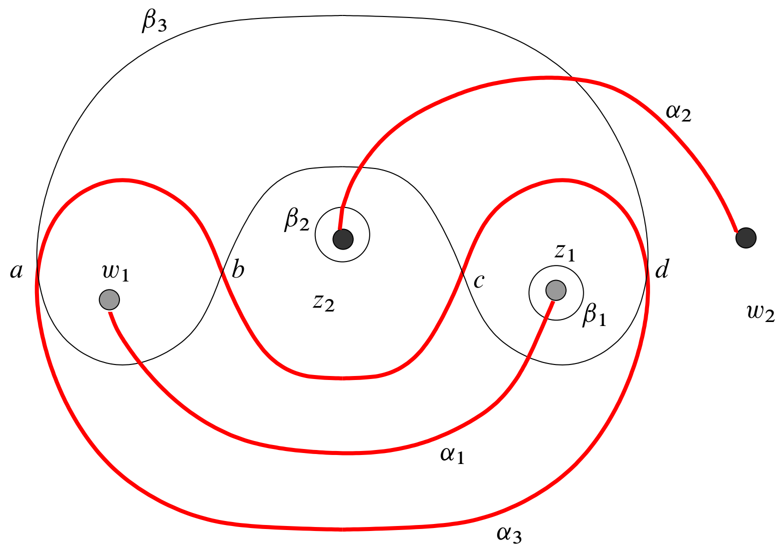

Figure 4: Another Heegaard diagram for the Hopf link. This picture is obtained from Figure 3 by stabilizing twice and doing some handleslides. It has the advantage that the basepoints come in pairs (w_1, z_1) and (w_2, z_2), with w_i and z_i on each side of the curve β_i.

one from b to a containing w_1, one from b to c containing z_2, one from d to c containing z_1, on from d to a containing w_2; two other bigons go from a to b and contain z_1, but they cancel each other in the Floer complex, so for all our purposes they can be ignored. Similarly, there are two bigons from c to d containing w_1, two bigons from a to d containing z_2, and two bigons from c to b containing w_2; all of these can be ignored.

In the genus two α–β Heegaard diagram from Figure 4 (which is the diagram denoted by $\mathscr{H}^L$ as part of the basic system), the tori $\mathbb{T}_\alpha = \alpha_1 \times \alpha_2 \times \alpha_3$ and $\mathbb{T}_\beta = \beta_1 \times \beta_2 \times \beta_3$ again intersect each other in four points. Indeed, since β_1 intersects a single alpha curve, namely α_1, and that intersection consists of a single point, that point must contribute to any intersection in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Similar remarks apply to $\alpha_1 \cap \beta_2$. Therefore, the intersection points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are determined by their component from $\alpha_3 \cap \beta_3$. We denote them still by a, b, c and d, using the obvious correspondence with the intersections in Figure 3.

The index one holomorphic disks between $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ in Figure 3 are also in one-to-one correspondence with those in Figure 4. Indeed, each bigon from Figure 3 corresponds to an annular domain in Figure 4. These annular domains are of the same kind as those considered in [Ozsváth and Szabó 2004c, proof of Lemma 3.4], where it is proved that they support exactly one holomorphic representative (modulo 2). Therefore, when building α–β Floer chain complexes from Figure 4, we may just as well look at the simpler Figure 3 and count the corresponding bigons.

The chain complexes we build from $\mathscr{H}^L$ are $\mathfrak{A}^-(\mathscr{H}^L, \boldsymbol{s})$ for $\boldsymbol{s} \in \mathbb{H}(L)$. We use the alternative definition for these complexes given in Section 3.6. Note that

$$\mathbb{H}(L) = \left(\tfrac{1}{2} + \mathbb{Z}\right) \times \left(\tfrac{1}{2} + \mathbb{Z}\right).$$

The Alexander gradings (A_1, A_2) of a, b, c and d are $\left(\frac{1}{2}, \frac{1}{2}\right)$, $\left(-\frac{1}{2}, \frac{1}{2}\right)$, $\left(-\frac{1}{2}, -\frac{1}{2}\right)$ and $\left(\frac{1}{2}, -\frac{1}{2}\right)$, respectively. Therefore, the formulas for the exponents $E^i_{s_i}(\phi)$ that appear in the definitions of $\mathfrak{A}^-(\mathscr{H}^L, \boldsymbol{s})$ depend only on the signs of s_i.

More precisely, for each $\boldsymbol{s} = (s_1, s_2) \in \mathbb{H}(L)$, the complex $\mathfrak{A}^-(\mathscr{H}^L, \boldsymbol{s})$ is freely generated over $\mathbb{F}[\![U_1, U_2]\!]$ by a, b, c and d. When $s_1, s_2 > 0$, the differential ∂ on the respective complex counts powers of U according to the multiplicities of w_1 and w_2 and ignores z_1 and z_2. We get the complex

$$\mathfrak{A}^{++}: \qquad \partial a = \partial c = 0, \quad \partial b = U_1 a + c, \quad \partial d = U_2 a + c.$$

When $s_1 > 0$ and $s_2 < 0$, we use w_1 and z_2, and ignore z_1 and w_2. We get the complex

$$\mathfrak{A}^{+-}: \qquad \partial a = \partial c = 0, \quad \partial b = U_1 a + U_2 c, \quad \partial d = a + c.$$

When $s_1 < 0$ and $s_2 > 0$, we use z_1 and w_2 and obtain

$$\mathfrak{A}^{-+}: \qquad \partial a = \partial c = 0, \quad \partial b = a + c, \quad \partial d = U_2 a + U_1 c.$$

Finally, when $s_1 < 0$ and $s_2 < 0$, we use z_1 and z_2 and obtain

$$\mathfrak{A}^{--}: \qquad \partial a = \partial c = 0, \quad \partial b = a + U_2 c, \quad \partial d = a + U_1 c.$$

We now turn our attention to the other three Heegaard diagrams in the basic system $\mathscr{H}$, namely $\mathscr{H}^{L_1}$, $\mathscr{H}^{L_2}$ and $\mathscr{H}^{\varnothing}$. Note that

$$\mathbb{H}(L_1) = \mathbb{H}(L_2) = \mathbb{Z}, \quad \mathbb{H}(\varnothing) = 0$$

and we have

$$\psi^{\pm L_1} \colon \mathbb{H}(L) \to \mathbb{H}(L_2), \quad \psi^{\pm L_1}(s_1, s_2) = s_2 \mp \tfrac{1}{2},$$
$$\psi^{\pm L_2} \colon \mathbb{H}(L) \to \mathbb{H}(L_1), \quad \psi^{\pm L_1}(s_1, s_2) = s_1 \mp \tfrac{1}{2}.$$

The diagram $\mathscr{H}^{L_1}$ is obtained from $\mathscr{H}^{L}$ by deleting z_2. Let us study $\mathfrak{A}^-(\mathscr{H}^{L_1}, s)$ for $s \in \mathbb{Z}$. The four generators a, b, c and d have Alexander gradings 0, -1, -1 and 0, respectively. Thus, the complex $\mathfrak{A}^-(\mathscr{H}^{L_1}, s)$ is isomorphic to $\mathfrak{A}^{++}$ for $s \geq 0$ and to $\mathfrak{A}^{-+}$ for $s < 0$.

Similarly, $\mathscr{H}^{L_2}$ is obtained from $\mathscr{H}^{L}$ by deleting z_1. The complex $\mathfrak{A}^-(\mathscr{H}^{L_2}, s)$ is isomorphic to $\mathfrak{A}^{++}$ for $s \geq 0$ and to $\mathfrak{A}^{+-}$ for $s < 0$.

Lastly, $\mathscr{H}^{\varnothing}$ is obtained from $\mathscr{H}^{L}$ by deleting both z_1 and z_2, and the corresponding complex $\mathfrak{A}^-(\mathscr{H}^{\varnothing}, 0)$ is a copy of $\mathfrak{A}^{++}$.

The surgery coefficients p_1, p_2 on the two components of L describe a framing Λ of L. Let $(1, 0)$ and $(0, 1)$ be the generators of $H_1(S^3 - L) \cong \mathbb{Z}^2$ corresponding to the meridians of L_1 and L_2, respectively. Since the linking number between L_1 and L_2 is 1, the framings of the components are $\Lambda_1 = (p_1, 1)$ and $\Lambda_2 = (1, p_2)$.

Let us now describe the full complex $\mathscr{C}^-(\mathscr{H}, \Lambda)$, whose homology is presumed to produce

$$\boldsymbol{HF}^-(S^3_\Lambda(\vec{L})) \cong \mathbb{F}[\![U]\!]^{\oplus(p_1 p_2 - 1)}.$$

As an $\mathscr{R}$-module, $\mathscr{C} = \mathscr{C}^-(\mathscr{H}, \Lambda)$ is the direct product of complexes $\mathscr{C}_{\boldsymbol{s}}$ over $\boldsymbol{s} = (s_1, s_2) \in \mathbb{H}(L) = \left(\mathbb{Z} + \tfrac{1}{2}\right)^2$, where

$$\mathscr{C}_{\boldsymbol{s}} = \mathfrak{A}^-(\mathscr{H}^L, \boldsymbol{s}) \oplus \mathfrak{A}^-\big(\mathscr{H}^{L_1}, s_1 - \tfrac{1}{2}\big) \oplus \mathfrak{A}^-\big(\mathscr{H}^{L_2}, s_2 - \tfrac{1}{2}\big) \oplus \mathfrak{A}^-(\mathscr{H}^{\varnothing}, 0).$$

Following (28), we denote the four terms in the direct sum above by $\mathscr{C}^{00}_{\boldsymbol{s}}$, $\mathscr{C}^{01}_{\boldsymbol{s}}$, $\mathscr{C}^{10}_{\boldsymbol{s}}$, $\mathscr{C}^{11}_{\boldsymbol{s}}$ in this order; for simplicity, we write the superscript as 00 rather than $(0, 0)$, etc. Recall that each of the terms is freely generated over $\mathbb{F}[\![U_1, U_2]\!]$ by four generators a, b, c and d. In order to be able to tell these generators apart, we denote them by $a^{\varepsilon_1\varepsilon_2}_{\boldsymbol{s}}$, $b^{\varepsilon_1\varepsilon_2}_{\boldsymbol{s}}$, etc when they live in $\mathscr{C}^{\varepsilon_1\varepsilon_2}_{\boldsymbol{s}}$.

The differential $\mathscr{D}^- \colon \mathscr{C}^-(\mathscr{H}, \Lambda) \to \mathscr{C}^-(\mathscr{H}, \Lambda)$ also splits as a sum of four terms

$$\mathscr{D}^- = \mathscr{D}^{00} + \mathscr{D}^{01} + \mathscr{D}^{10} + \mathscr{D}^{11},$$

where $\mathscr{D}^{\varepsilon'_1\varepsilon'_2}$ maps $\mathscr{C}_{\boldsymbol{s}}$ to $\mathscr{C}_{\boldsymbol{s}+\varepsilon'_1\Lambda_1+\varepsilon'_2\Lambda_2}$. We have chosen here to drop ε^0, ε and $\boldsymbol{s}$ from (29). A graphical representation of the differentials $\mathscr{D}^{\varepsilon'_1\varepsilon'_2}$ is given in Figure 5. Note that the equivalence relation on $\mathbb{H}(L)$ generated by

$$(s_1, s_2) \sim (s_1 + p_1, s_2 + 1), \quad (s_1, s_2) \sim (s_1 + 1, s_2 + p_2)$$

breaks $\mathbb{H}(L)$ into $p_1 p_2 - 1$ equivalence classes, corresponding to the Spin^c structures on the surgered manifold $S^3_\Lambda(\vec{L})$.

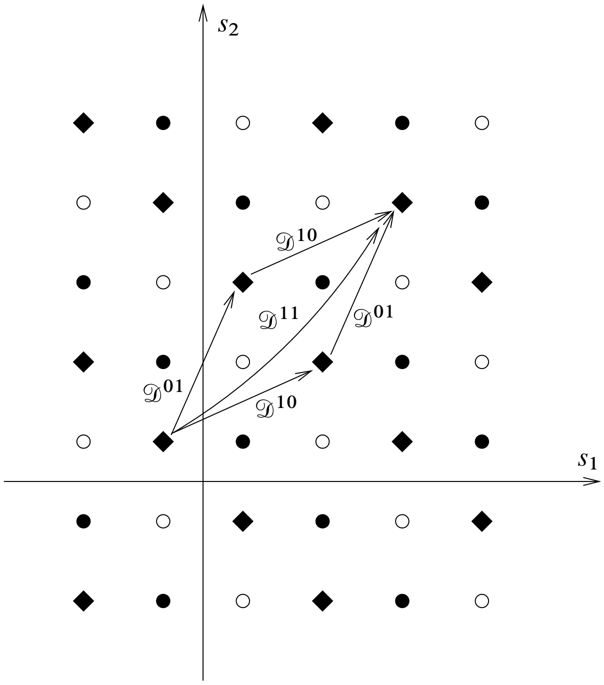

Figure 5: The complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$ for $p_1 = p_2 = 2$. We show here the lattice $\mathbb{H}(L)$, as a union of various icons in the plane: black dots, white dots and black diamonds. Each type of icon corresponds to a particular Spin^c structure on the surgered manifold. We also show how various parts of the differential $\mathfrak{D}^-$ act on the lattice. Not shown is $\mathfrak{D}^{00}$, which simply preserves each icon. Note that all parts of $\mathfrak{D}^-$ preserve the type of the icon.

More precisely, $\mathfrak{D}^{00}$ consists of the differentials $\partial = \Phi_{\boldsymbol{s}}^{\varnothing}$ on the chain complexes $\mathcal{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ themselves, plus the cross terms

$$\Phi_{\boldsymbol{s}}^{L_1}\colon \mathcal{C}_{\boldsymbol{s}}^{00} \to \mathcal{C}_{\boldsymbol{s}}^{10}, \quad \Phi_{\boldsymbol{s}}^{L_2}\colon \mathcal{C}_{\boldsymbol{s}}^{00} \to \mathcal{C}_{\boldsymbol{s}}^{01}, \quad \Phi_{\psi^{L_2}(\boldsymbol{s})}^{L_1}\colon \mathcal{C}_{\boldsymbol{s}}^{01} \to \mathcal{C}_{\boldsymbol{s}}^{11}, \quad \Phi_{\psi^{L_1}(\boldsymbol{s})}^{L_2}\colon \mathcal{C}_{\boldsymbol{s}}^{10} \to \mathcal{C}_{\boldsymbol{s}}^{11}, \quad \Phi_{\boldsymbol{s}}^{L}\colon \mathcal{C}_{\boldsymbol{s}}^{00} \to \mathcal{C}_{\boldsymbol{s}}^{11}.$$

Note that $\Phi_{\boldsymbol{s}}^{L}$ is a chain homotopy between $\Phi_{\psi^{L_2}(\boldsymbol{s})}^{L_1} \circ \Phi_{\boldsymbol{s}}^{L_2}$ and $\Phi_{\psi^{L_1}(\boldsymbol{s})}^{L_2} \circ \Phi_{\boldsymbol{s}}^{L_1}$. In fact, as noted in Section 4.4, the map $\Phi_{\boldsymbol{s}}^{L}$ vanishes in the case of a basic system.

The term $\mathfrak{D}^{10}$ is simpler. It consists of the maps

$$\begin{aligned} \Phi_{(s_1,s_2)}^{-L_1}\colon \mathcal{C}_{(s_1,s_2)}^{00} &\to \mathcal{C}_{(s_1+p_1,s_2+1)}^{10}, \\ \Phi_{\psi^{L_2}(s_1,s_2)}^{-L_1}\colon \mathcal{C}_{(s_1,s_2)}^{01} &\to \mathcal{C}_{(s_1+p_1,s_2+1)}^{11}, \\ \Phi_{(s_1,s_2)}^{(-L_1)\cup L_2}\colon \mathcal{C}_{(s_1,s_2)}^{00} &\to \mathcal{C}_{(s_1+p_1,s_2+1)}^{11}. \end{aligned}$$

The chain homotopy $\Phi_{(s_1,s_2)}^{(-L_1)\cup L_2}$ is again zero. , A similar description applies to $\mathfrak{D}^{01}$. Finally, the term $\mathfrak{D}^{11}$ is the simplest of all, consisting only of maps of the type

$$\Phi_{(s_1,s_2)}^{-L}\colon \mathcal{C}_{(s_1,s_2)}^{00} \to \mathcal{C}_{(s_1+p_1+1,s_2+p_2+1)}^{11}.$$

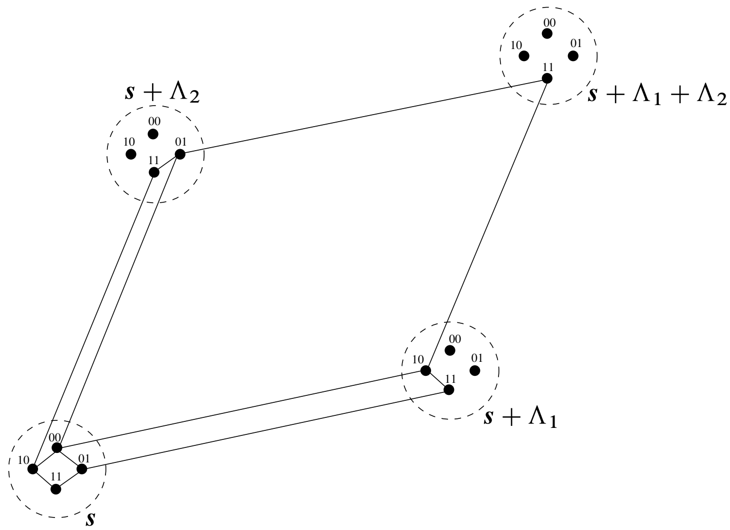

Figure 6: Four locally defined filtrations. This is a more in-depth look at Figure 5. Each dashed circle corresponds to one of the icons from Figure 5, and the four bullets inside a dashed circle are the four summands $\mathscr{C}_{s}^{\varepsilon}$. The values of ε are noted near each bullet, and the values of s near each dashed circle. Each of the four parallelograms represents a summand (namely, the one containing $\mathscr{C}_{s}^{00}$) in the associated graded complex of one of the four filtrations $\mathscr{F}_{00}$, $\mathscr{F}_{10}$, $\mathscr{F}_{01}$ and $\mathscr{F}_{11}$.

Let us introduce four locally defined filtrations

$$\mathscr{F}_{\omega_1\omega_2}, \quad \omega_1, \omega_2 \in \{0, 1\},$$

on the complex $\mathscr{C}^-(\mathscr{H}, \Lambda)$. These filtrations will play an important role in calculating the homology. The differential on the associated graded of $\mathscr{F}_{\omega_1\omega_2}$ will be denoted by $\widetilde{\mathscr{D}}^{\omega_1\omega_2}$. These associated graded complexes are shown schematically in Figure 6.

Recall that the generators of $\mathscr{C}^-(\mathscr{H}, \Lambda)$ (as a direct product) are of the form

$$g_{(s_1,s_2)}^{\varepsilon_1,\varepsilon_2}, \quad g \in \{a, b, c, d\},\ \varepsilon_1, \varepsilon_2 \in \{0, 1\},\ s_1, s_2 \in \mathbb{Z} + \tfrac{1}{2}.$$

The first filtration $\mathscr{F}_{00}$ is defined on generators by

$$\mathscr{F}_{00}(g_{(s_1,s_2)}^{\varepsilon_1,\varepsilon_2}) = \min\{-s_1, -s_2\}.$$

The differential $\mathscr{D}^-$ either preserves or decreases the filtration level. In the associated graded, the only visible part of $\mathscr{D}^-$ is $\widetilde{\mathscr{D}}^{00} = \mathscr{D}^{00}$. Consequently, the associated graded splits as a direct product of terms of the form $(\mathscr{C}_{s}, \mathscr{D}^{00})$.

The next filtration $\mathscr{F}_{10}$ is defined on generators by

$$\mathscr{F}_{10}(g_{(s_1,s_2)}^{\varepsilon_1,\varepsilon_2}) = s_1 - (p_1 - 1)s_2 - \varepsilon_1.$$

Again, the differential $\mathcal{D}^-$ either preserves or decreases the filtration level. In the associated graded, the only visible part of $\mathcal{D}^-$ is $\widetilde{\mathcal{D}}^{10}$, which is the sum of $\mathcal{D}^{10}$ and the parts of $\mathcal{D}^{00}$ that preserve ε_1. Consequently, the associated graded splits as a direct product of terms of the form

$$\begin{tikzcd} \mathcal{C}^{00}_{(s_1,s_2)} \arrow[r] \arrow[d] \arrow[dr] & \mathcal{C}^{10}_{(s_1+p_1,s_2+1)} \arrow[d] \\ \mathcal{C}^{01}_{(s_1,s_2)} \arrow[r] & \mathcal{C}^{11}_{(s_1+p_1,s_2+1)} \end{tikzcd} \tag{35}$$

We similarly have a filtration $\mathcal{F}_{01}$ given by

$$\mathcal{F}_{01}(g^{\varepsilon_1,\varepsilon_2}_{(s_1,s_2)}) = s_2-(p_2-1)s_1-\varepsilon_2.$$

The differential $\widetilde{\mathcal{D}}^{01}$ on its associated graded consists of $\mathcal{D}^{01}$ and the parts of $\mathcal{D}^{00}$ that preserve ε_2. There is a direct product splitting of the associated graded analogous to (35):

$$\begin{tikzcd} \mathcal{C}^{00}_{(s_1,s_2)} \arrow[r] \arrow[d] \arrow[dr] & \mathcal{C}^{10}_{(s_1,s_2)} \arrow[d] \\ \mathcal{C}^{01}_{(s_1+1,s_2+p_2)} \arrow[r] & \mathcal{C}^{11}_{(s_1+1,s_2+p_2)} \end{tikzcd} \tag{36}$$

The last filtration $\mathcal{F}_{11}$ is defined on generators by

$$\mathcal{F}_{11}(g^{\varepsilon_1,\varepsilon_2}_{(s_1,s_2)}) = \min\{s_1-p_1\varepsilon_1-\varepsilon_2,\, s_2-\varepsilon_1-p_2\varepsilon_2\}.$$

The corresponding differential $\widetilde{\mathcal{D}}^{11}$ on the associated graded is the sum of the following terms: the differentials $\partial=\Phi^{\varnothing}_{\boldsymbol{s}}$ on the chain complexes $\mathcal{C}^{\varepsilon_1\varepsilon_2}_{\boldsymbol{s}}$ themselves, plus the cross terms

$$\Phi^{-L_1}_{(s_1,s_2)}\colon \mathcal{C}^{00}_{(s_1,s_2)} \to \mathcal{C}^{10}_{(s_1+p_1,s_2+1)}, \quad \Phi^{-L_1}_{\psi^{L_2}(s_1,s_2)}\colon \mathcal{C}^{01}_{(s_1,s_2)} \to \mathcal{C}^{11}_{(s_1+p_1,s_2+1)},$$
$$\Phi^{-L_2}_{(s_1,s_2)}\colon \mathcal{C}^{00}_{(s_1,s_2)} \to \mathcal{C}^{01}_{(s_1+1,s_2+p_2)}, \quad \Phi^{-L_2}_{\psi^{L_1}(s_1,s_2)}\colon \mathcal{C}^{10}_{(s_1,s_2)} \to \mathcal{C}^{11}_{(s_1+1,s_2+p_2)},$$

and

$$\mathcal{D}^{11}=\Phi^{-L}_{(s_1,s_2)}\colon \mathcal{C}^{00}_{(s_1,s_2)} \to \mathcal{C}^{11}_{(s_1+p_1+1,s_2+p_2+1)}.$$

Consequently, the associated graded of $\mathcal{F}_{11}$ splits as a direct product of terms of the form

$$\begin{tikzcd} \mathcal{C}^{00}_{(s_1,s_2)} \arrow[r] \arrow[d] \arrow[dr] & \mathcal{C}^{10}_{(s_1+p_1,s_2+1)} \arrow[d] \\ \mathcal{C}^{01}_{(s_1+1,s_2+p_2)} \arrow[r] & \mathcal{C}^{11}_{(s_1+p_1+1,s_2+p_2+1)} \end{tikzcd} \tag{37}$$

Let us now turn to the computation of the homology of $\mathcal{C}^-(\mathcal{H},\Lambda)$. This complex has a subcomplex

$$\mathcal{C}_{\geq 0} = \prod_{\substack{(s_1,s_2)\\ \max(s_1,s_2)>0}} \mathcal{C}_{(s_1,s_2)}.$$

Lemma 4.5 $H_*(\mathcal{C}_{\geq 0})=0.$

Proof We use the restriction of the filtration $\mathscr{F}_{00}$ to $\mathscr{C}_{\geq 0}$. Note that this restriction is bounded above. In light of Lemma 4.1, it suffices to show that the homology of the associated graded groups $H_*(\mathscr{C}_{(s_1,s_2)}, \mathscr{D}^{00})$ vanishes, whenever $s_1 > 0$ or $s_2 > 0$.

Let us first consider the case $s_1, s_2 > 0$. Then $\mathscr{C}_{(s_1,s_2)}$ consists of four copies of the same complex, $\mathfrak{A}^{++}$, related by the five cross terms in the description of $\mathscr{D}^{00}$ above. Among these, $\Phi_{\boldsymbol{s}}^{L_1}$, $\Phi_{\boldsymbol{s}}^{L_2}$, $\Phi_{\psi^{L_2}(\boldsymbol{s})}^{L_1}$ and $\Phi_{\psi^{L_1}(\boldsymbol{s})}^{L_2}$ are all isomorphisms, corresponding to the identity on the complex $\mathfrak{A}^{++}$, while the chain homotopy $\Phi_{\boldsymbol{s}}^{L}$ is trivial. Thus, the complex $\mathscr{C}_{\boldsymbol{s}}$ can be described as

$$\begin{tikzcd} \mathfrak{A}^{++} \arrow[r, "\cong"] \arrow[d, "\cong"'] & \mathfrak{A}^{++} \arrow[d, "\cong"] \\ \mathfrak{A}^{++} \arrow[r, "\cong"] & \mathfrak{A}^{++} \end{tikzcd}$$

This is clearly acyclic.

Next, let us consider the case $s_1 > 0$ and $s_2 < 0$. Then $\mathscr{C}_{(s_1,s_2)}$ consists of two copies of $\mathfrak{A}^{+-}$, namely $\mathscr{C}^{00}_{(s_1,s_2)}$ and $\mathscr{C}^{10}_{(s_1,s_2)}$, and two copies of $\mathfrak{A}^{++}$, namely $\mathscr{C}^{01}_{(s_1,s_2)}$ and $\mathscr{C}^{11}_{(s_1,s_2)}$. Thus $\mathscr{C}_{\boldsymbol{s}}$ can be described as

$$\begin{tikzcd} \mathfrak{A}^{+-} \arrow[r, "\cong"] \arrow[d] & \mathfrak{A}^{+-} \arrow[d] \\ \mathfrak{A}^{++} \arrow[r, "\cong"] & \mathfrak{A}^{++} \end{tikzcd}$$

Even though the vertical maps are not isomorphisms, the horizontal ones are identities. This suffices to show that $\mathscr{C}_{\boldsymbol{s}}$ is acyclic. Indeed, the filtration $\mathscr{F}_{01}$ restricts to a filtration on $\mathscr{C}_{\boldsymbol{s}}$, whose associated graded differential consists of the two horizontal maps above. Therefore, this associated graded is acyclic, and so is $\mathscr{C}_{\boldsymbol{s}}$.

Similar remarks apply to the case $s_1 < 0$ and $s_2 > 0$. This shows that $\mathscr{C}_{\geq 0}$ is acyclic. $\square$

Lemma 4.5 implies that the homology of $\mathscr{C}^-(H, \Lambda)$ is the same as that of its quotient complex

$$\mathscr{C}_{<0} = \prod_{s_1,s_2<0} \mathscr{C}_{(s_1,s_2)}.$$

Next, we show that a large part of the complex $\mathscr{C}_{<0}$ is also acyclic. Pick small $\zeta_1, \zeta_2 > 0$ linearly independent over $\mathbb{Q}$. Consider the parallelogram $P_{\mathbb{R}}$ in the plane with vertices

$$(-\zeta_1, -\zeta_2), \quad (-\zeta_1 - p_1, -\zeta_2 - 1), \quad (-\zeta_1 - 1, -\zeta_2 - p_2), \quad (-\zeta_1 - p_1 - 1, -\zeta_2 - p_2 - 1).$$

If ζ_1 and ζ_2 are sufficiently small, the parallelogram $P_{\mathbb{R}}$ contains a unique representative from each equivalence class in $\mathbb{H}(L)$, ie from each Spin^c structure on the surgered manifold. Set $P = P_{\mathbb{R}} \cap \mathbb{H}(L)$.

For $\omega_1, \omega_2 \in \{0, 1\}$, let $Q_{\omega_1\omega_2}$ be the quadrant in the plane given by

$$Q_{\omega_1\omega_2} = \{(s_1, s_2) \in \mathbb{H}(L) \mid s_1 < -\zeta_1 + (\omega_1 - 1)p_1 + (\omega_2 - 1), \quad s_2 < -\zeta_2 + (\omega_1 - 1) + (\omega_2 - 1)p_2\}.$$

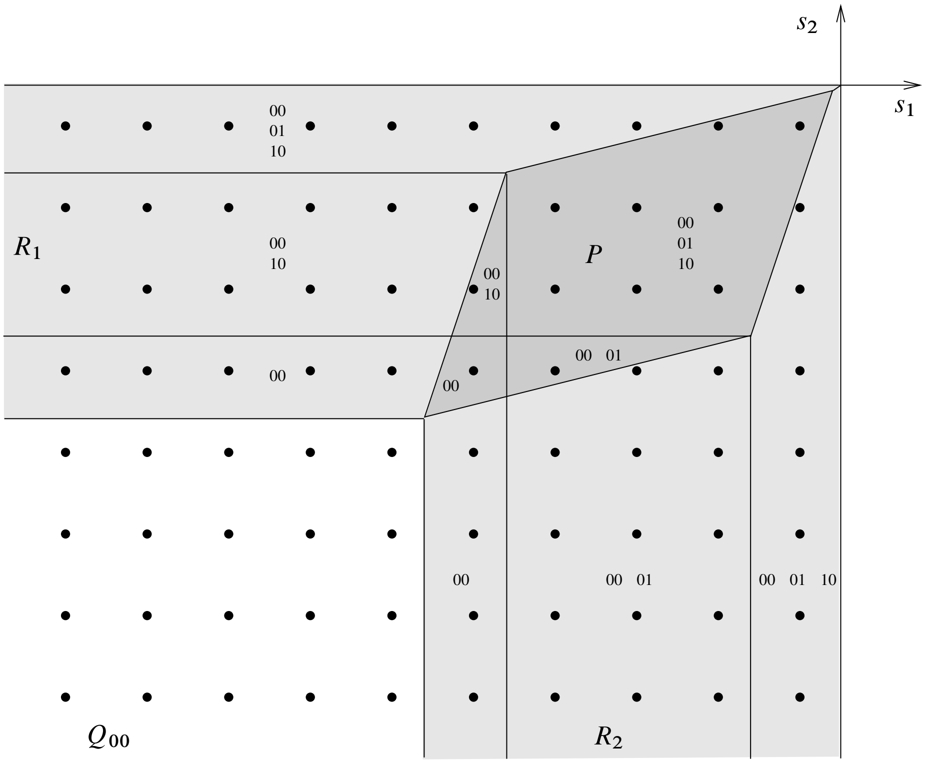

Figure 7: The complex $\mathscr{C}_{<0}$ for $p_1 = 4$ and $p_2 = 3$. The parallelogram P is darkly shaded, and the two regions R_1 and R_2 are more lightly shaded. Each dot represents an element of $Q_{11} \subset \mathbb{H}(L)$. The boundaries of Q_{00}, Q_{01}, Q_{10}, Q_{11} and P split the lower left quadrant in eleven regions. In each of these regions we mark the values $\varepsilon_1\varepsilon_2$ for which the respective groups $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ are part of the complex $\mathscr{C}_R = \mathscr{C}_{<0}/\mathscr{C}_{<R}$.

The complement $R = Q_{11} \setminus Q_{00}$ consists of the union of P and two other regions: one to the left of P, which we denote by R_1, and one below P, which we denote by R_2. These regions and the parallelogram P are shown in Figure 7.

Consider the submodule of $\mathscr{C}_{<0}$

$$\mathscr{C}_{<R} = \prod_{\boldsymbol{s} \in Q_{00}} (\mathscr{C}_{\boldsymbol{s}}^{00} \oplus \mathscr{C}_{\boldsymbol{s}+\Lambda_1}^{10} \oplus \mathscr{C}_{\boldsymbol{s}+\Lambda_2}^{01} \oplus \mathscr{C}_{\boldsymbol{s}+\Lambda_1+\Lambda_2}^{11}).$$

It is straightforward to check that $\mathscr{C}_{<R}$ is a subcomplex of $\mathscr{C}_{<0}$. We denote the corresponding quotient complex by $\mathscr{C}_R$.

Lemma 4.6 $$H_*(\mathscr{C}_{<R}) = 0.$$

Proof We use the restriction of the filtration $\mathscr{F}_{11}$ to the subcomplex $\mathscr{C}_{<R}$. Since this restriction is bounded above, by Lemma 4.1 it suffices to prove that the associated graded groups are acyclic.

Recall that the differential of the associated graded of $\mathscr{F}_{11}$ is denoted by $\widetilde{\mathfrak{D}}^{11}$; see the description of the terms of $\widetilde{\mathfrak{D}}^{11}$ before the diagram (37). As discussed in Section 4.4, the maps Φ appearing in $\widetilde{\mathfrak{D}}^{11}$ are compositions of descent maps D with inclusion maps $\mathscr{I}$. Since $s_1, s_2 \leq -\frac{1}{2}$ and the generators all have Alexander gradings at least $-\frac{1}{2}$, from the description (24) of the inclusion maps we see that in our case the inclusion maps are the identity.

Turning our attention to the descent maps, one could in principle compute them explicitly by counting holomorphic polygons. However, it is not necessary to do so. All we need to know is that the descent maps at one component only (be it $-L_1$ or $-L_2$), given by counting holomorphic triangles, induce isomorphisms on homology. This is true because they correspond to the natural triangle maps between strongly equivalent Heegaard diagrams.

The associated graded of $\mathscr{C}_{<R}$ with respect to $\mathscr{F}_{11}$ breaks down into a direct product of factors of the form (37), one for each $\boldsymbol{s} \in Q_{00}$. The fact that the vertical and horizontal maps in (37) are quasi-isomorphisms implies that the respective factors are acyclic. The claim follows. □

We denote the quotient complex of $\mathscr{C}_{<R} \subset \mathscr{C}_{<0}$ by C_R. Lemma 4.6 implies that

$$H_*(\mathscr{C}_{<0}) = H_*(\mathscr{C}_R).$$

Let $\mathscr{C}_P \subset \mathscr{C}_R$ be the submodule

$$\mathscr{C}_P = \prod_{\boldsymbol{s} \in P} \mathscr{C}^{00}_{\boldsymbol{s}}.$$

This is a quotient complex of $\mathscr{C}_R$. The respective subcomplex $\mathscr{C}_{R_1 \cup R_2}$ splits as a direct sum of two complexes

$$\mathscr{C}_{R_1} = \Big(\prod_{\boldsymbol{s} \in R_1} \mathscr{C}_{\boldsymbol{s}} \oplus \prod_{\boldsymbol{s} \in P} \mathscr{C}^{10}_{\boldsymbol{s}} \Big) \cap \mathscr{C}_R, \quad \mathscr{C}_{R_2} = \Big(\prod_{\boldsymbol{s} \in R_2} \mathscr{C}_{\boldsymbol{s}} \oplus \prod_{\boldsymbol{s} \in P} \mathscr{C}^{01}_{\boldsymbol{s}} \Big) \cap \mathscr{C}_R.$$

Lemma 4.7

$$H_*(\mathscr{C}_{R_1}) = H_*(\mathscr{C}_{R_2}) = 0.$$

Proof Consider the restriction of the filtration $\mathscr{F}_{10}$ to the complex $\mathscr{C}_{R_1}$. This restriction is bounded above. The respective associated graded splits as a direct product of mapping cone complexes of two possible kinds:

$$\mathfrak{A}^{--} \cong \mathscr{C}^{00}_{\boldsymbol{s}} \xrightarrow{\Phi^{-L_1}_{\boldsymbol{s}}} \mathscr{C}^{10}_{\boldsymbol{s}+\Lambda_1} \cong \mathfrak{A}^{+-}$$

for $\boldsymbol{s} \in Q_{01} \cap R_1$, and

$$\mathfrak{A}^{--} \cong \mathscr{C}^{00}_{\boldsymbol{s}} \xrightarrow{\Phi^{L_2}_{\boldsymbol{s}}} \mathscr{C}^{01}_{\boldsymbol{s}} \cong \mathfrak{A}^{-+}$$

for $\boldsymbol{s} = \big(s_1, -\frac{1}{2}\big) \in R_1$.

Mapping cone complexes of the first kind are acyclic by the same reasoning as in Lemma 4.6, because $\Phi^{-L_1}_{\boldsymbol{s}}$ is a quasi-isomorphism.

Let us study a mapping complex of the second kind. To compute $\Phi_{\boldsymbol{s}}^{L_2}$, note that it is the composition of a descent map D and an inclusion map $\mathcal{I}$; the former is the identity and the latter is multiplication by suitable powers of U_2. More precisely, in terms of the generators a, b, c and d of $\mathfrak{A}^{--}$ and $\mathfrak{A}^{-+}$, the map $\Phi_{\boldsymbol{s}}^{L_2}$ is given by

$$a \mapsto U_2 a, \quad b \mapsto U_2 b, \quad c \mapsto c, \quad d \mapsto d.$$

This induces an isomorphism on homology, which implies that the respective mapping cone complex is acyclic. Hence $H_*(\mathscr{C}_{R_1}) = 0$.

The proof that $\mathscr{C}_{R_2}$ is acyclic is similar, but uses the filtration $\mathscr{F}_{01}$ instead of $\mathscr{F}_{10}$. $\square$

Putting together Lemmas 4.5, 4.6 and 4.7 we obtain that the homology of the full complex $\mathscr{C}^-(\mathscr{H}, \Lambda)$ is the same as the homology of the complex

$$\mathscr{C}_P = \prod_{\boldsymbol{s} \in P} \mathscr{C}_{\boldsymbol{s}}^{00}.$$

For $s_1, s_2 < 0$, we have $\mathscr{C}_{\boldsymbol{s}}^{00} \cong \mathfrak{A}^{--}$, whose homology is easily seen to be isomorphic to

$$\mathbb{F}[\![U_1, U_2]\!]/(U_1 - U_2) \cong \mathbb{F}[\![U]\!].$$

Since there are $p_1 p_2 - 1$ lattice points from $\mathbb{H}(L)$ inside the parallelogram P, we obtain

$$H_*(\mathscr{C}^-(\mathscr{H}, \Lambda)) \cong \mathbb{F}[\![U]\!]^{\oplus (p_1 p_2 - 1)},$$

as expected.

4.6 Remarks about the general case

For an ℓ-component link $\vec{L} \subset Y$, the surgery complex (27) takes the form of an ℓ-dimensional hypercube, similar to the double complex (34). At the vertices of the hypercube we have complexes $\mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$ which are resolutions of the generalized Floer complexes $\boldsymbol{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$. When the diagrams $\mathscr{H}^{L-M}$ are link-minimal (that is, have only two basepoints on each link component), the complexes $\mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$ were constructed in Section 3.6, and are isomorphic to $\boldsymbol{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$. In the general case, the resolutions $\mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(\boldsymbol{s}))$ will be defined in Section 13.

Describing the maps $\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}$ that are the building blocks of the differential requires additional work, even in the link-minimal case (such as for a basic system). In the link-minimal case, the maps $\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}$ are the composition of inclusion maps $\mathcal{I}$ (of the type defined in Section 3.8), and certain descent maps D. The simplest descent maps (those that appear along the edges of the hypercube) are equivalences induced by Heegaard moves. The descent maps that are needed for the higher-dimensional faces of the hypercube are chain homotopies between the edge maps, and higher homotopies between these. They will be defined by counting holomorphic polygons. Note that, for the equivalences induced by Heegaard moves, we need to compose the maps corresponding to a single move. Similarly, for the higher descent maps, we need to compose higher homotopies from hypercubes. The homological algebra needed to describe this kind of compositions will be developed in Section 5.

The general link surgery formula is stated not just for a basic system, but for a general complete system of hyperboxes. (We need this more general formulation in the context of grid diagrams, for example; see Section 15.) Complete systems will be defined in Section 8. In a complete system, the diagram $\mathcal{H}^M$ associated to $M \subset L$ is not necessarily the reduction of $\mathcal{H}^L$ at M. Rather, for any orientation $\vec{M}$ of M, as part of the complete system we specify a sequence of Heegaard moves that relate $r_{\vec{M}}(\mathcal{H}^L)$ to $\mathcal{H}^M$. These Heegaard moves are then used to construct the descent maps. Furthermore, in the case where the diagrams are not link-minimal, to get the descent maps we will compose the maps induced by Heegaard moves with certain other transition maps; we refer to Section 13 for more details.

5 Hyperboxes of chain complexes and compression

In this section we develop some homological algebra that is essential for the statement of the surgery theorem. All the vector spaces we consider are over $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$.

When f is a function, we denote its n^{th} iterate by $f^{\circ n}$, ie $f^{\circ 0} = \mathrm{id}$, $f^{\circ 1} = f$ and $f^{\circ(n+1)} = f^{\circ n} \circ f$.

5.1 Hyperboxes of chain complexes

An n-dimensional *hyperbox* is a subset of $\mathbb{R}^n$ of the form $[0, d_1] \times \cdots \times [0, d_n]$, where $d_i \geq 0$ for $i = 1, \ldots, n$. We will assume that $\boldsymbol{d} = (d_1, \ldots, d_n) \in (\mathbb{Z}_{\geq 0})^n$ is a collection of nonnegative integers. We then let $\mathbb{E}(\boldsymbol{d})$ be the set of points in the corresponding hyperbox with integer coordinates, ie

$$\mathbb{E}(\boldsymbol{d}) = \{\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n) \mid \varepsilon_i \in \{0, 1, \ldots, d_i\}, i = 1, \ldots, n\}.$$

In particular, $\mathbb{E}_n = \mathbb{E}(1, \ldots, 1) = \{0, 1\}^n$ is the set of vertices of the n-dimensional unit hypercube.

For $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n) \in \mathbb{E}(\boldsymbol{d})$, we set

$$\|\varepsilon\| = \varepsilon_1 + \cdots + \varepsilon_n.$$

We can view the elements of $\mathbb{E}(\boldsymbol{d})$ as vectors in $\mathbb{R}^n$. There is a partial ordering on $\mathbb{E}(\boldsymbol{d})$, given by $\varepsilon' \leq \varepsilon$ if and only if $\varepsilon'_i \leq \varepsilon_i$ for all i. We write $\varepsilon' < \varepsilon$ if $\varepsilon' \leq \varepsilon$ and $\varepsilon' \neq \varepsilon$.

For $i = 1, \ldots, n$, let $\tau_i \in \mathbb{E}_n$ be the n-tuple formed of $n-1$ zeros and a single one, where the one is in position i. Then, for any $\varepsilon \in \mathbb{E}(\boldsymbol{d})$,

$$\varepsilon = \varepsilon_1 \tau_1 + \cdots + \varepsilon_n \tau_n.$$

Definition 5.1 An *n-dimensional hyperbox of chain complexes* of size $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$ consists of a collection of $\mathbb{Z}$-graded vector spaces

$$(C^\varepsilon)_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}, \quad C^\varepsilon = \bigoplus_{* \in \mathbb{Z}} C^\varepsilon_*,$$

together with a collection of linear maps

$$D^\varepsilon_{\varepsilon^0} : C^{\varepsilon^0}_* \to C^{\varepsilon^0 + \varepsilon}_{*-1+\|\varepsilon\|},$$

one for each $\varepsilon^0 \in \mathbb{E}(\boldsymbol{d})$ and $\varepsilon \in \mathbb{E}_n$ such that $\varepsilon^0 + \varepsilon \in \mathbb{E}(\boldsymbol{d})$. The maps are required to satisfy the relations

$$\sum_{\varepsilon' \le \varepsilon} D^{\varepsilon - \varepsilon'}_{\varepsilon^0 + \varepsilon'} \circ D^{\varepsilon'}_{\varepsilon^0} = 0 \tag{38}$$

for all $\varepsilon^0 \in \mathbb{E}(\boldsymbol{d})$ and $\varepsilon \in \mathbb{E}_n$ such that $\varepsilon^0 + \varepsilon \in \mathbb{E}(\boldsymbol{d})$.

Given a hyperbox of chain complexes as above, we write

$$C = \bigoplus_{\varepsilon \in \mathbb{E}(\boldsymbol{d})} C^{\varepsilon}$$

and define linear maps $D^{\varepsilon} \colon C \to C$, by setting them on generators to

$$D^{\varepsilon}(x) = \begin{cases} D^{\varepsilon}_{\varepsilon^0}(x) & \text{for } x \in C^{\varepsilon^0} \text{ with } \varepsilon^0 + \varepsilon \in \mathbb{E}(\boldsymbol{d}), \\ 0 & \text{for } x \in C^{\varepsilon^0} \text{ with } \varepsilon^0 + \varepsilon \notin \mathbb{E}(\boldsymbol{d}). \end{cases}$$

We denote a typical hyperbox of chain complexes by $H = ((C^{\varepsilon})_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}, (D^{\varepsilon})_{\varepsilon \in \mathbb{E}_n})$;[2] the maps $D^{\varepsilon}_{\varepsilon^0}$ are implicitly taken into account in the direct sums D^{ε}. Sometimes, by abuse of notation, we let D^{ε} stand for any of its terms $D^{\varepsilon}_{\varepsilon^0}$. If $\boldsymbol{d} = (1, \dots, 1)$, we say that H is a *hypercube of chain complexes*.

Observe that a zero-dimensional hyperbox of chain complexes is simply a chain complex, while a one-dimensional hyperbox with $\boldsymbol{d} = (d)$ consists of chain complexes $C^{(i)}$ for $i = 0, \dots, d$ together with a string of chain maps

$$C^{(0)} \xrightarrow{D^{(1)}} C^{(1)} \xrightarrow{D^{(1)}} \cdots \xrightarrow{D^{(1)}} C^{(d)}. \tag{39}$$

To give another example, a two-dimensional hypercube is a diagram of complexes and chain maps

$$\begin{array}{ccc} C^{(0,0)} & \xrightarrow{D^{(1,0)}} & C^{(1,0)} \\ {\scriptstyle D^{(0,1)}}\downarrow & & \downarrow{\scriptstyle D^{(0,1)}} \\ C^{(0,1)} & \xrightarrow{D^{(1,0)}} & C^{(1,1)} \end{array}$$

together with a chain homotopy

$$D^{(1,1)} \colon C^{(0,0)} \to C^{(1,1)} \tag{40}$$

between $D^{(1,0)} \circ D^{(0,1)}$ and $D^{(0,1)} \circ D^{(1,0)}$.

In general, if $((C^{\varepsilon})_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}, (D^{\varepsilon})_{\varepsilon \in \mathbb{E}_n})$ is an n-dimensional hyperbox, then $(C^{\varepsilon}, D^{(0,\dots,0)})$ are chain complexes. Along the edges of the hyperbox we see strings of chain maps D^{τ_i} for $i = 1, \dots, n$. In fact, let us imagine the hyperbox $[0, d_1] \times \cdots \times [0, d_n]$ to be split into $d_1 d_2 \cdots d_n$ unit hypercubes. Then, along each edge of one of these hypercubes, we see a chain map. Along the two-dimensional faces of the unit hypercubes we have chain homotopies, and along higher-dimensional faces we have higher homotopies.

[2]The same notation ε is used for both elements of $\mathbb{E}(\boldsymbol{d})$ and elements of $\mathbb{E}_n$. Of course, $\mathbb{E}_n$ can be viewed a subset of $\mathbb{E}(\boldsymbol{d})$.

Observe that, when $H = (C^\varepsilon, D^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$ is a hypercube, we can also form a total complex $H_{\text{tot}} = (C_*, D)$, where the grading on C is given by

$$C_* = \bigoplus_{\varepsilon \in \mathbb{E}(\boldsymbol{d})} C^\varepsilon_{*+\|\varepsilon\|}$$

and the boundary map is the sum $D = \sum D^\varepsilon$.

5.2 Compression

Let $H = ((C^\varepsilon)_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}, (D^\varepsilon)_{\varepsilon \in \mathbb{E}_n})$ be an n-dimensional hyperbox of chain complexes. We will explain how to construct from H an n-dimensional hypercube $\hat{H} = (\hat{C}^\varepsilon, \hat{D}^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$. The process of turning H into $\hat{H}$ will be called *compression*.

The simplest example of compression is when $n = 1$, and H is a string of chain complexes and chain maps as in (39). Then compression is composing the maps. Precisely, the compressed hypercube $\hat{H}$ consists of the complexes $\hat{C}^{(0)} = C^{(0)}$ and $\hat{C}^{(1)} = C^{(d)}$, linked by the chain map

$$(D^{(1)})^{\circ d} = D^{(1)} \circ \cdots \circ D^{(1)} \colon C^{(0)} \to C^{(d)}.$$

For general n and $\boldsymbol{d} = (d_1, \dots, d_n)$, the compressed hypercube $\hat{H}$ has at its vertices the same complexes as those at the vertices of the original hyperbox H:

$$\hat{C}^{(\varepsilon_1, \dots, \varepsilon_n)} = C^{(\varepsilon_1 d_1, \dots, \varepsilon_n d_n)}, \quad \varepsilon = (\varepsilon_1, \dots, \varepsilon_n) \in \mathbb{E}_n.$$

Further, along each edge of $\hat{H}$ we should see the composition of the respective edge maps in H, ie

$$\hat{D}^{\tau_i} = (D^{\tau_i})^{\circ d_i}.$$

The construction of the maps corresponding to the higher-dimensional faces of $\hat{H}$ is rather involved, and will occupy Sections 5.3–5.6. For now, to give a flavor of the respective formulas, let us explain the simplest nontrivial case, namely $n = 2$.

When $n = 2$ and $\boldsymbol{d} = (d_1, d_2)$, the hyperbox H is a rectangle split into $d_1 d_2$ unit squares. Along the horizontal edges we have chain maps denoted by $f_1 = D^{(1,0)}$ and along the vertical edges we have chain maps denoted by $f_2 = D^{(0,1)}$. Further, each unit square carries a chain homotopy $f_{\{1,2\}} = D^{(1,1)}$ between $f_1 \circ f_2$ and $f_2 \circ f_1$. Then, on the edges of the compressed hypercube $\hat{H}$, we have maps

$$\hat{D}^{(1,0)} = f_1^{\circ d_1}, \quad \hat{D}^{(0,1)} = f_2^{\circ d_2}.$$

For the diagonal map $\hat{D}^{(1,1)}$ we choose

$$\hat{D}^{(1,1)} = \sum_{j_1=1}^{d_1} \sum_{j_2=1}^{d_2} f_1^{\circ(j_1-1)} \circ f_2^{\circ(j_2-1)} \circ f_{\{1,2\}} \circ f_2^{\circ(d_2-j_2)} \circ f_1^{\circ(d_1-j_1)}. \tag{41}$$

It is easy to check that $\hat{D}^{(1,1)}$ is a chain homotopy between $\hat{D}^{(1,0)} \circ \hat{D}^{(0,1)} = f_1^{\circ d_1} \circ f_2^{\circ d_2}$ and $\hat{D}^{(0,1)} \circ \hat{D}^{(1,0)} = f_2^{\circ d_2} \circ f_1^{\circ d_1}$. See Figure 8 for a pictorial interpretation.

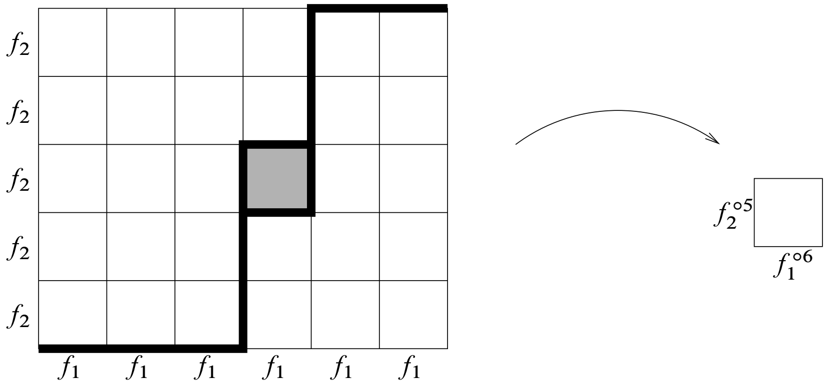

Figure 8: Compression of a rectangle into a square. This is the graphical representation of the term of the form $f_1^{\circ 2} \circ f_2^{\circ 2} \circ f_{\{1,2\}} \circ f_2^{\circ 2} \circ f_1^{\circ 3}$ appearing in the sum (41), with $d_1 = 6$, $d_2 = 5$ and $j_1 = j_2 = 3$. Each unit segment which is part of the thick line corresponds to a map f_1 or f_2, while the shaded square is the chain homotopy $f_{\{1,2\}}$. Taking the sum of all the terms in (41) corresponds to filling up the whole rectangle with 30 unit squares, and represents a chain homotopy between $f_1^{\circ 6} \circ f_2^{\circ 5}$ and $f_2^{\circ 5} \circ f_1^{\circ 6}$.

5.3 The algebra of songs

Let X be a finite set. (Typically, X will be a subset of the set of nonnegative integers.)

Definition 5.2 An *X-valued song* is a finite ordered list of *items*, where each item can be either a *note*, ie an element of X, or a *harmony*, ie a subset of X.

For example, if $X = \{1, 2, 3\}$, a typical X-valued song is written as

$$s = (213\{2,3\}2\{\}12\{3\}).$$

The song s has nine items, six being notes and three being harmonies. Note that we allow the empty harmony $\{\}$, and that we distinguish between the note 3 and the one-element harmony $\{3\}$. Also, our convention is to write songs between parentheses. For example, if $A \subseteq X$ is a harmony, by (A) we mean the one-item song made of that harmony.

Let $\widetilde{S}(X)$ be the $\mathbb{F}$-vector space freely generated by all X-valued songs. Since songs form a monoid under concatenation (with the empty song as the identity element), this induces the structure of a noncommutative algebra on $\widetilde{S}(X)$.

Definition 5.3 The *algebra of X-valued songs*, denoted by $S(X)$, is the quotient of $\widetilde{S}(X)$ by the ideal $I(X)$ generated by the following relations:

- For any note $x \in X$,

$$(x\{\}) = (\{\}x) \tag{42}$$

and

$$(x\{x\}) = (\{x\}x). \tag{43}$$

- For any note $x \in X$ and song $s \in \widetilde{S}(X)$,

$$(x\{x\}sx) + (xs\{x\}x) = (xs) + (sx). \tag{44}$$

- For any harmony $A \subseteq X$,

$$\sum_{B \subseteq A} (B)(A \setminus B) = 0. \tag{45}$$

Remark 5.4 The reader may wonder why we defined the algebra $S(X)$ this way. The motivation behind our choice of the relations (42)–(45) is that they do not affect the playing of songs, as defined in Section 5.5 below. See Lemma 5.11 for the relevant result.

Let $Y = X \cup \{y\}$ be the set obtained from X by adding a new note y. We define an operation on songs

$$\psi_y \colon \widetilde{S}(X) \to \widetilde{S}(Y),$$

as follows. For a note $x \in X$, we set

$$\psi_y(x) = (xy\{x, y\}yx).$$

For a harmony $A \subseteq X$, let $\Pi(A)$ be the set of all ordered decompositions $(A_1, \dots, A_k)$ of A into a disjoint union

$$A = A_1 \amalg A_2 \amalg \cdots \amalg A_k,$$

where $k \geq 0$ and all A_i's are nonempty. For $A \neq \{\}$, we set

$$\psi_y(A) = \sum_{(A_1,\dots,A_k) \in \Pi(A)} (y(A_1 \cup \{y\})y(A_2 \cup \{y\})y \cdots y(A_k \cup \{y\})y),$$

while for $A = \{\}$ we set $\psi_y(A) = (y)$.

So far we defined ψ_y only on one-item songs, consisting of either one note or one harmony. We extend it to arbitrary songs by requiring it to act as a derivation:

$$\psi_y(s_1 s_2) = \psi_y(s_1)s_2 + s_1\psi_y(s_2).$$

For example,

$$\psi_3(2\{1,2\}) = (23\{2,3\}32\{1,2\}) + (23\{1,2,3\}3) + (23\{1,3\}3\{2,3\}3) + (23\{2,3\}3\{1,3\}3).$$

Finally, we extend ψ_y to all of $\widetilde{S}(X)$ by requiring it to be linear.

Lemma 5.5 *The operation ψ_y descends to a linear map between $S(X)$ and $S(Y)$.*

Proof We need to check that when we apply ψ_y to the relations (42)–(45) from Definition 5.3, we obtain relations that hold true in $S(Y)$, ie lie in the ideal $I(Y)$.

Let us first look at the relation (42). The claim is that $\psi_y(x\{\}) = \psi_y(\{\}x)$. Indeed,

$$
\begin{aligned}
\psi_y(x\{\}) + \psi_y(\{\}x) &= (xy\{x,y\}yx\{\}) + (xy) + (\{\}xy\{x,y\}yx) + (yx) \\
&= (xy(\{x,y\}\{\} + \{\}\{x,y\})yx) + (xy) + (yx) \\
&= (xy(\{x\}\{y\} + \{y\}\{x\})yx) + (xy) + (yx) \\
&= (x(y\{x\}\{y\}y + y\{y\}\{x\}y)x) + (xy) + (yx) \\
&= (x(y\{x\} + \{x\}y)x) + (xy) + (yx) \\
&= (xy\{x\}x) + (x\{x\}yx) + (xy) + (yx) \\
&= 0.
\end{aligned} \tag{46}
$$

To get the second equality in (46) we used (42), namely the fact that $\{\}$ commutes with x and y. To get the third equality we applied (45) for $A = \{x, y\}$, while to get the fifth and seventh equalities we applied (44).

The similar result for (43) is simpler. It suffices to apply (43) twice and(44) once:

$$
\begin{aligned}
\psi_y(x\{x\}) + \psi_y(\{x\}x) &= (xy\{x,y\}yx\{x\}) + (\{x\}xy\{x,y\}yx) + (xy\{x,y\}y) + (y\{x,y\}yx) \\
&= (xy\{x,y\}y\{x\}x) + (x\{x\}y\{x,y\}yx) + (xy\{x,y\}y) + (y\{x,y\}yx) \\
&= 0.
\end{aligned} \tag{47}
$$

Here is the analogous result for (44):

$$
\begin{aligned}
&\psi_y(x\{x\}sx) + \psi_y(xs\{x\}x) + \psi_y(xs) + \psi_y(sx) \\
&\quad = (xy\{x,y\}yx\{x\}sx) + (xy\{x,y\}ysx) + (x\{x\}\psi_y(s)x) + (x\{x\}sxy\{x,y\}yx) \\
&\qquad + (xy\{x,y\}yxs\{x\}x) + (xsy\{x,y\}yx) + (x\psi_y(s)\{x\}x) + (xs\{x\}xy\{x,y\}yx) \\
&\qquad + (xy\{x,y\}yxs) + (x\psi_y(s)) + (\psi_y(s)x) + (sxy\{x,y\}yx) \\
&\quad = \big((xy\{x,y\}yx\{x\}sx) + (xy\{x,y\}yxs\{x\}x) + (xy\{x,y\}ysx) + (xy\{x,y\}yxs)\big) \\
&\qquad + \big((x\{x\}sxy\{x,y\}yx) + (xs\{x\}xy\{x,y\}yx) + (xsy\{x,y\}yx) + (sxy\{x,y\}yx)\big) \\
&\qquad + \big((x\{x\}\psi_y(s)x) + (x\psi_y(s)\{x\}x) + (x\psi_y(s)) + (\psi_y(s)x)\big) \\
&\quad = 0.
\end{aligned} \tag{48}
$$

In the last step, the four terms within each of the large parentheses cancel each other out by applying (44).

Lastly, we prove that (45) holds true after applying ψ_y. Let us introduce the following notational shortcut: if A is a subset of X, we write $\widetilde{A} = A \cup \{y\} \subseteq Y$.

For any $A \subseteq X$, we have

$$
\begin{aligned}
(49)\quad & \psi_y\Big(\sum_{B\subseteq A}(B)(A\setminus B)\Big)\\
&= \sum_{B\subseteq A}\psi_y(B)\cdot(A\setminus B)+(A\setminus B)\cdot\psi_y(B)\\
&= \sum_{B\subseteq A}\sum_{(B_1,\dots,B_k)\in\Pi(B)}\big((y\widetilde{B}_1y\cdots y\widetilde{B}_ky(A\setminus B))+((A\setminus B)y\widetilde{B}_1y\cdots y\widetilde{B}_ky)\big)\\
&= \sum_{(A_1,\dots,A_k)\in\Pi(A)}((y\widetilde{A}_1y\cdots y\widetilde{A}_{k-1}yA_k)+(A_1y\widetilde{A}_2y\cdots y\widetilde{A}_ky))\\
&\qquad\qquad + \sum_{(A_1,\dots,A_k)\in\Pi(A)}\big((y\widetilde{A}_1y\cdots y\widetilde{A}_ky\{\})+(\{\}y\widetilde{A}_1y\cdots y\widetilde{A}_ky)\big).
\end{aligned}
$$

The last equality was obtained by splitting the summation on the second line into terms with $B \neq A$ and $B = A$.

The final expression in (49) is a sum of two terms, where each term is a summation over the elements of $\Pi(A)$. Using the fact that $\{\}$ and y commute, the second summation (the one appearing on the very last line of (49)) is seen to equal

$$
\sum_{(A_1,\dots,A_k)\in\Pi(A)}\sum_{i=1}^{k}\big((y\widetilde{A}_1y\cdots y\widetilde{A}_{i-1}y)\cdot\big((\{\}\widetilde{A}_i)+(\widetilde{A}_i\{\})\big)\cdot(y\widetilde{A}_{i+1}y\cdots y\widetilde{A}_ky)\big). \tag{50}
$$

Applying (45) to $\widetilde{A}_i$, we see that the middle factor of the summand in (50) is

$$
(\{\}\widetilde{A}_i)+(\widetilde{A}_i\{\}) = \sum_{\substack{B\subseteq\widetilde{A}_i\\ B\neq\varnothing,\widetilde{A}_i}}(B)(\widetilde{A}_i\setminus B) = (A_i\{y\})+(\{y\}A_i)+\sum_{\substack{B\subseteq A_i\\ B\neq\varnothing,A_i}}((\widetilde{B})(A_i\setminus B)+(B)(\widetilde{A}_i\setminus B)).
$$

Plugging this back into (50), we obtain that (50) equals

$$
\begin{aligned}
(51)\quad & \sum_{(A_1,\dots,A_k)\in\Pi(A)}\sum_{i=1}^{k}\big((y\widetilde{A}_1y\cdots y\widetilde{A}_{i-1})\cdot\big((y\{y\}A_iy)+(yA_i\{y\}y)\big)\cdot(\widetilde{A}_{i+1}y\cdots y\widetilde{A}_ky)\big)\\
&\qquad + \sum_{(A_1,\dots,A_k)\in\Pi(A)}\sum_{i=1}^{k-1}(y\widetilde{A}_1y\cdots y\widetilde{A}_{i-1}A_iy\cdots y\widetilde{A}_ky)\\
&\qquad\qquad + \sum_{(A_1,\dots,A_k)\in\Pi(A)}\sum_{i=1}^{k-1}(y\widetilde{A}_1y\cdots yA_{i-1}\widetilde{A}_iy\cdots y\widetilde{A}_ky).
\end{aligned}
$$

Here, in the last two summations, we changed notation so that $(A_1,\dots,A_{i-1},B,A_i\setminus B,A_{i+1},\dots,A_k)$ is renamed $(A_1,\dots,A_k)$.

Now, applying (44) to the middle factor $(y\{y\}A_iy)+(yA_i\{y\}y)$ in the first summand in (51), we can replace it with $(yA_i)+(A_iy)$. Consequently, most of the terms in that first summation cancel out with terms in the second and third summations in (51). The only remaining terms are some corresponding to $i=1$ and $i=k$. More precisely, we get that (51) equals

$$\sum_{(A_1,\dots,A_k)\in\Pi(A)} ((y\widetilde{A}_1 y\cdots y\widetilde{A}_{k-1} yA_k)+(A_1 y\widetilde{A}_2 y\cdots y\widetilde{A}_k y)).$$

This exactly corresponds to the first summation in the final expression in (49). Hence, we obtain

$$\psi_y\biggl(\sum_{B\subseteq A}(B)(A\setminus B)\biggr)=0,$$

as desired. □

5.4 Symphonies

Definition 5.6 Let X be a finite, totally ordered set, and $m(X)$ the maximal element in X. The *symphony* $\alpha(X)\in\widetilde{S}(X)$ *on the set* X is defined, recursively, by

$$\alpha(\varnothing)=(\{\}),\quad \alpha(X)=\psi_{m(X)}\bigl(\alpha(X\setminus\{m(X)\})\bigr).$$

We call $\alpha_n=\alpha(\{1,2,\dots,n\})$ the n^{th} *standard symphony*.

For example,

$$\begin{aligned}
\alpha_1&=\psi_1(\{\})=(1),\\
\alpha_2&=\psi_2(1)=(12\{1,2\}21),\\
\alpha_3&=\psi_3(12\{1,2\}21)\\
&=(123\{1,2,3\}321)+(123\{2,3\}32\{1,2\}21)+(12\{1,2\}23\{2,3\}321)+(12\{1,2\}213\{1,3\}31)\\
&\qquad+(13\{1,3\}312\{1,2\}21)+(123\{1,3\}3\{2,3\}321)+(123\{2,3\}3\{1,3\}321).
\end{aligned}$$

Computer experimentation shows that α_4 is a linear combination of 97 different songs, and α_5 a linear combination of 2051 different songs. In general, a song s that appears with nonzero multiplicity in α_n is easily seen to satisfy the following two conditions. Let s consist of k notes and l harmonies, and let $h_1,\dots,h_l$ be the cardinalities of each harmony. Then

$$k=2n+l-1,\quad \sum_{i=1}^{l}h_i=n+l-1. \tag{52}$$

Of course, not every song that satisfies (52) appears in the formula for α_n.

Lemma 5.7 *For any finite, totally ordered set X, we have the relation, in* $S(X)$,

$$\sum_{Y\subseteq X}\alpha(Y)\alpha(X\setminus Y)=0. \tag{53}$$

Proof We proceed by induction on the cardinality n of X. For $n=0$ the corresponding relation $\{\}\{\}=0$ is (45) for $A=\{\}$, while for $n=1$ the corresponding relation $(x\{\})=(\{\}x)$ is (42).

Let X be a set of cardinality $n\geq 2$, and write $m=m(X)$. Suppose that (53) is true for all sets of cardinality $<n$, and, in particular, for $X'=X\setminus\{m\}$. Then

$$\begin{aligned}\sum_{Y\subseteq X}\alpha(Y)\alpha(X\setminus Y)&=\sum_{Y\subseteq X'}\big(\alpha(Y)\alpha(X'\cup\{m\}\setminus Y)+\alpha(Y\cup\{m\})\alpha(X'\setminus Y)\big)\\&=\sum_{Y\subseteq X'}\big(\alpha(Y)\psi_m(\alpha(X'\setminus Y))+\big(\psi_m(\alpha(Y))\big)\alpha(X'\setminus Y)\big)\\&=\psi_m\bigg(\sum_{Y\subseteq X'}\alpha(Y)\alpha(X'\setminus Y)\bigg)\\&=\psi_m(0)=0.\end{aligned}$$

□

5.5 Hypercubical collections

Let X be a finite set and $(A,+,*)$ be a (possibly noncommutative) algebra over $\mathbb{F}$. Given an element $\mathscr{A}\in A$, we denote by $\mathscr{A}^{*j}=\mathscr{A}*\cdots*\mathscr{A}$ its j^{th} power. In particular, $\mathscr{A}^{*0}=1$ is the unit.

Definition 5.8 An n-dimensional *hypercubical collection* in the algebra A, modeled on X, is a collection $\mathscr{A}$ composed of elements $\mathscr{A}_Z\in A$, one for each $Z\subseteq X$, which are required to satisfy the relations

$$\sum_{Z'\subseteq Z}\mathscr{A}_{Z'}*\mathscr{A}_{Z\setminus Z'}=0 \tag{54}$$

for any $Z\subseteq X$.

Example 5.9 Let $H=(C^{\varepsilon},D^{\varepsilon})$ be a hyperbox of chain complexes as in Section 5.1. Choose $X=\{1,2,\ldots,n\}$. For any $Z\subseteq X$, we can define an element $\zeta(Z)=(\zeta(Z)_1,\ldots,\zeta(Z)_n)\in\mathbb{E}_n=\{0,1\}^n$ by

$$\zeta(Z)_i=\begin{cases}1 & \text{if } i\in Z,\\ 0 & \text{otherwise.}\end{cases} \tag{55}$$

Then the

$$\mathscr{A}_Z=D^{\zeta(Z)}$$

form a hypercubical collection in the algebra $\operatorname{End}(C)$, modeled on X.

Definition 5.10 Let X be a finite set and let $\boldsymbol{d}=(d_x)_{x\in X}$ be a collection of positive integers indexed by X. Let $\mathscr{A}=\{A_Z\}_{Z\subseteq X}$ be a hypercubical collection in an algebra A. Let also s be an X-valued song, spelled out as a sequence of items

$$s=(x_1^1\cdots x_{r_1}^1\{y_1^1,\ldots,y_{t_1}^1\}x_1^2\cdots x_{r_2}^2\{y_1^2,\ldots,y_{t_2}^2\}\cdots x_1^l\cdots x_{r_l}^l\{y_1^l,\ldots,y_{t_l}^l\}x_1^{l+1}\cdots x_{r_{l+1}}^{l+1}),$$

where some of the r_i's can be zero.

The result of *playing the song s to the hypercubical collection $\mathscr{A}$, in the register $\boldsymbol{d}$*, is the algebra element $\mathrm{pl}^{\boldsymbol{d}}_{\mathscr{A}}(s) \in A$ defined by

$$\text{(56)} \quad \mathrm{pl}^{\boldsymbol{d}}_{\mathscr{A}}(s) = \sum \mathscr{A}^{*j^1_1}_{\{x^1_1\}} * \cdots * \mathscr{A}^{*j^1_{r_1}}_{\{x^1_{r_1}\}} * \mathscr{A}_{\{y^1_1,\dots,y^1_{t_1}\}} * \mathscr{A}^{*j^2_1}_{\{x^2_1\}} * \cdots * \mathscr{A}^{*j^2_{r_2}}_{\{x^2_{r_2}\}} * \mathscr{A}_{\{y^2_1,\dots,y^2_{t_2}\}} * \cdots * \mathscr{A}^{*j^l_1}_{\{x^l_1\}} * \cdots * \mathscr{A}^{*j^l_{r_l}}_{\{x^l_{r_l}\}} * \mathscr{A}_{\{y^l_1,\dots,y^l_{t_l}\}} * \mathscr{A}^{*j^{l+1}_1}_{\{x^{l+1}_1\}} * \cdots * \mathscr{A}^{*j^{l+1}_{r_{l+1}}}_{\{x^{l+1}_{r_{l+1}}\}},$$

where the sum is over all possible powers $j^\sigma_u \geq 0$ satisfying, for each $x \in X$,

$$\sum_{\sigma=1}^{l+1} \sum_{u=1}^{r_\sigma} \delta_{x,x^\sigma_u} j^\sigma_u + \sum_{\sigma=1}^{l} \sum_{u=1}^{t_\sigma} \delta_{x,y^\sigma_u} = d_x.$$

Here $\delta_{i,j}$ is the Kronecker delta symbol.

It is worth giving some examples of playing. First, note that, unless every $x \in X$ appears at least once in the song s (either as a note or as part of a harmony), we have $\mathrm{pl}^{\boldsymbol{d}}_{\mathscr{A}}(s) = 0$.

If $X = \{1\}$ and $\boldsymbol{d} = (d)$, then $\mathscr{A}$ consists of two algebra elements $\mathscr{A}_\varnothing$ and $\mathscr{A}_{\{1\}}$. Playing the song (1) to $\mathscr{A}$ yields the d^{th} power $\mathscr{A}^{*d}_{\{1\}}$. On the other hand, playing the song $(\{1\})$ to $\mathscr{A}$ yields $\mathscr{A}_{\{1\}}$ when $d = 1$ and 0 otherwise.

If $X = \{1, 2\}$ and $\boldsymbol{d} = (d_1, d_2)$, then $\mathscr{A}$ consists of four algebra elements $\mathscr{A}_\varnothing$, $\mathscr{A}_{\{1\}}$, $\mathscr{A}_{\{2\}}$ and $\mathscr{A}_{\{1,2\}}$. Playing the second standard symphony $\alpha_2 = (12\{1,2\}21)$ to $\mathscr{A}$ yields

$$\mathrm{pl}^{(d_1,d_2)}_{\mathscr{A}}(\alpha_2) = \sum_{\substack{j^1_1+j^2_2=d_1-1 \\ j^1_2+j^2_1=d_2-1}} \mathscr{A}^{*j^1_1}_{\{1\}} * \mathscr{A}^{*j^1_2}_{\{2\}} * \mathscr{A}_{\{1,2\}} * \mathscr{A}^{*j^2_1}_{\{2\}} * \mathscr{A}^{*j^2_2}_{\{1\}},$$

or, equivalently,

$$\mathrm{pl}^{(d_1,d_2)}_{\mathscr{A}}(\alpha_2) = \sum_{j_1=1}^{d_1} \sum_{j_2=1}^{d_2} \mathscr{A}^{*(j_1-1)}_{\{1\}} * \mathscr{A}^{*(j_2-1)}_{\{2\}} * \mathscr{A}_{\{1,2\}} * \mathscr{A}^{*(d_2-j_2)}_{\{2\}} * \mathscr{A}^{*(d_1-j_1)}_{\{1\}}.$$

(Compare (41) and Figure 8.)

Lemma 5.11 *Let $\mathscr{A}$ be a hypercubical collection in an algebra A, modeled on a set X, and pick a series of nonnegative integers $\boldsymbol{d} = (d_x)_{x \in X}$. Then the operation of playing songs to $\mathscr{A}$ in the register $\boldsymbol{d}$ descends to a linear map $\mathrm{pl}^{\boldsymbol{d}}_{\mathscr{A}} \colon S(X) \to A$.*

Proof We extend the playing of songs linearly to a map $\widetilde{S}(X) \to A$. In order to show that it descends to $S(X)$, we need to check that the relations (42)–(45) hold true after playing them. In fact, since playing is not multiplicative, one needs to check that these relations, when multiplied on the left and right with arbitrary songs, still hold true after playing.

For example, let us look at the relation (42), namely $(x\{\}) = (\{\}x)$ for any $x \in X$. The claim is that the maps $\mathrm{pl}_{(s_1 x\{\} s_2)}$ and $\mathrm{pl}_{(s_1\{\}x s_2)}$ are equal, for any songs s_1 and s_2. This is true because $x\{\}$ corresponds to taking a power of $\mathscr{A}_{\{x\}}$ in the big summation in (56), followed by the factor $\mathscr{A}_\varnothing$. Equation (54) for $Z = \{x\}$ implies that $\mathscr{A}_\varnothing$ commutes with $\mathscr{A}_{\{x\}}$, so it also commutes with its power.

The relation (43) holds true after playing because both $(x\{x\})$ and $(\{x\}x)$ correspond to taking an arbitrary, but nontrivial, power of $\mathscr{A}_{\{x\}}$ in the big summation in (56).

The relation (44) holds true after playing because the left-hand side is a sum of two terms, both roughly of the form xsx, except that in the first we impose the condition that the exponent of $\mathscr{A}_{\{x\}}$ terms is nonzero on the left of s, and in the second that it is nonzero on the right of s. Therefore, most of the terms obtained after playing cancel in pairs. The only remaining ones are those obtained by playing either sx or xs (with no $\mathscr{A}_{\{x\}}$ powers on the left and right, respectively).

Finally, the fact that (45) holds true after playing boils down to (54) in the definition of a hypercubical collection. □

Let $\mathscr{A}$ be a hypercubical collection modeled on a set X. For any $X' \subseteq X$, the subcollection composed of $\mathscr{A}_Z$ for $Z \subseteq X'$ is a hypercubical collection modeled on X', which we denote by $\mathscr{A}|X'$. Further, if $\boldsymbol{d} = (d_x)_{x\in X}$ is a series of nonnegative integers, by picking only the terms d_x for $x \in X'$ we obtain a new series, denoted by $\boldsymbol{d}|X'$.

The following is a straightforward consequence of the definition of playing taking into account Lemma 5.11:

Corollary 5.12 *Let $\mathscr{A}$ be a hypercubical collection in an algebra A, modeled on a set X. Let $X = X' \amalg X''$ be a decomposition of X as a disjoint union. Then*

$$\mathrm{pl}^{\boldsymbol{d}|X'}_{\mathscr{A}|X'}(s') * \mathrm{pl}^{\boldsymbol{d}|X''}_{\mathscr{A}|X''}(s'') = \mathrm{pl}^{\boldsymbol{d}}_{\mathscr{A}}(s's'')$$

for any $s' \in S(X')$ and $s'' \in S(X'')$.

Suppose now that X is a finite, totally ordered set. We equip all $Z \subseteq X$ with the induced total ordering.

Let $\mathscr{A}$ be a hypercubical collection in an algebra A, modeled on X. Pick $\boldsymbol{d} = (d_x)_{x\in X}$ with $d_x > 0$ and, for any $Z \subseteq X$, define

$$\mathscr{A}^{\boldsymbol{d}}_Z = \mathrm{pl}^{\boldsymbol{d}|Z}_{\mathscr{A}|Z}(\alpha(Z)),$$

where $\alpha(Z)$ is the symphony on Z from Definition 5.6.

Lemma 5.13 *The elements $\{\mathscr{A}^{\boldsymbol{d}}_Z\}_{Z\subseteq X}$ form a new hypercubical collection $\mathscr{A}^{\boldsymbol{d}}$ in A.*

Proof We need to check that, for any $Z \subseteq X$,

$$\sum_{Z'\amalg Z''=Z} \mathscr{A}^{\boldsymbol{d}}_{Z'} * \mathscr{A}^{\boldsymbol{d}}_{Z''} = 0. \tag{57}$$

Indeed, the left-hand side in (57) equals

$$\sum_{Z'\amalg Z''=Z} \mathrm{pl}^{\boldsymbol{d}|Z'}_{\mathscr{A}|Z'}(\alpha(Z')) * \mathrm{pl}^{\boldsymbol{d}|Z''}_{\mathscr{A}|Z''}(\alpha(Z'')) = \sum_{Z'\amalg Z''=Z} \mathrm{pl}^{\boldsymbol{d}|Z}_{\mathscr{A}|Z}(\alpha(Z')\alpha(Z'')) \tag{58}$$
$$= \mathrm{pl}^{\boldsymbol{d}|Z}_{\mathscr{A}|Z}\Biggl(\sum_{Z'\amalg Z''=Z} \alpha(Z')\alpha(Z'')\Biggr) = \mathrm{pl}^{\boldsymbol{d}|Z}_{\mathscr{A}|Z}(0) = 0.$$

The first equality in (58) is a consequence of Corollary 5.12, the second of linearity (Lemma 5.11), and the third of Lemma 5.7. $\square$

5.6 Back to compression

Let $H = ((C^\varepsilon)_{\varepsilon\in\mathbb{E}(\boldsymbol{d})}, (D^\varepsilon)_{\varepsilon\in\mathbb{E}_n})$ be an n-dimensional hyperbox of chain complexes as in Section 5.1. In Section 5.2 we advertised the construction of a compressed hypercube $\hat{H} = (\hat{C}^\varepsilon, \hat{D}^\varepsilon)_{\varepsilon\in\mathbb{E}_n}$, with

$$\hat{C}^{(\varepsilon_1,\dots,\varepsilon_n)} = C^{(\varepsilon_1 d_1,\dots,\varepsilon_n d_n)}.$$

We are now ready to explain the exact construction of the maps $\hat{D}^\varepsilon$. Let A be the algebra $\mathrm{End}(C)$ under composition, where $C = \bigoplus_{\varepsilon\in\mathbb{E}(\boldsymbol{d})} C^\varepsilon$. As mentioned in Example 5.9, the maps $D^\varepsilon = D^{\zeta(Z)} = \mathscr{A}_Z$ form a hypercubical collection $\mathscr{A}$ in A, modeled on $X = \{1,\dots,n\}$. (Note that every $\varepsilon \in \mathbb{E}_n$ can be written as $\zeta(Z)$ for a unique $Z \subseteq X$.)

For $Z \subseteq \{1,\dots,n\}$, set

$$\hat{D}^{\zeta(Z)} = \mathscr{A}^{\boldsymbol{d}}_Z = \mathrm{pl}^{\boldsymbol{d}|Z}_{\mathscr{A}|Z}(\alpha(Z)). \tag{59}$$

For example, when $Z = \{i\}$ and $\varepsilon = \tau_i$ for some $i \in \{1,\dots,n\}$, the map along the corresponding edge of the hypercube is $\hat{D}^{\tau_i} = \mathrm{pl}^{(d_i)}_{\mathscr{A}|\{i\}}(i) = (D^{\tau_i})^{\circ d_i}$, as noted in Section 5.2. For $n = 2$, by playing the symphony on a set of two elements, we recover formula (41) for $\hat{D}^{(1,1)}$.

Proposition 5.14 *$\hat{H} = (\hat{C}^\varepsilon, \hat{D}^\varepsilon)_{\varepsilon\in\mathbb{E}_n}$ is a hypercube of chain complexes.*

Proof The relations (38) are a direct consequence of Lemma 5.13. The fact that $\hat{D}^\varepsilon$ changes grading by $\|\varepsilon\| - 1$ (as required in the definition of a hypercube) follows from the similar property for the maps D^ε, together with the second relation in (52), which is satisfied by all the terms appearing in a symphony on a set of size n. $\square$

5.7 Chain maps and homotopies

Let

$${}^0H = (({}^0C^\varepsilon)_{\varepsilon\in\mathbb{E}(\boldsymbol{d})}, ({}^0D^\varepsilon)_{\varepsilon\in\mathbb{E}_n}), \quad {}^1H = (({}^1C^\varepsilon)_{\varepsilon\in\mathbb{E}(\boldsymbol{d})}, ({}^1D^\varepsilon)_{\varepsilon\in\mathbb{E}_n})$$

be two hyperboxes of chain complexes, having the same size $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$. Let $(\boldsymbol{d}, 1) \in (\mathbb{Z}_{\geq 0})^{n+1}$ be the sequence obtained from $\boldsymbol{d}$ by adding 1 at the end.

Definition 5.15 A *chain map* $F: {}^0H \to {}^1H$ is a collection of linear maps

$$F^{\varepsilon}_{\varepsilon^0}: {}^0C^{\varepsilon^0}_{*} \to {}^1C^{\varepsilon^0+\varepsilon}_{*+\|\varepsilon\|}$$

satisfying

$$\sum_{\varepsilon' \le \varepsilon} (D^{\varepsilon-\varepsilon'}_{\varepsilon^0+\varepsilon'} \circ F^{\varepsilon'}_{\varepsilon^0} + F^{\varepsilon-\varepsilon'}_{\varepsilon^0+\varepsilon'} \circ D^{\varepsilon'}_{\varepsilon^0}) = 0 \tag{60}$$

for all $\varepsilon^0 \in \mathbb{E}(\boldsymbol{d})$ and $\varepsilon \in \mathbb{E}_n$ such that $\varepsilon^0 + \varepsilon \in \mathbb{E}(\boldsymbol{d})$.

In other words, a chain map between the hyperboxes ${}^0H$ and ${}^1H$ is an $(n+1)$-dimensional hyperbox of chain complexes, of size $(\boldsymbol{d}, 1)$, such that the subhyperbox corresponding to $\varepsilon_{n+1} = 0$ is ${}^0H$ and the one corresponding to $\varepsilon_{n+1} = 1$ is ${}^1H$. The maps F are those maps D in the new hyperbox that increase ε_{n+1} by 1.

Note that, in the particular case when $\boldsymbol{d} = (0, \dots, 0)$, so that ${}^0H$ and ${}^1H$ are ordinary chain complexes, the notion of chain map coincides with the usual one. Also note that, when $\boldsymbol{d} = (1, \dots, 1)$, so that ${}^0H$ and ${}^1H$ are hypercubes, a chain map F induces an ordinary chain map F_{tot} between the corresponding total complexes ${}^0C_{\mathrm{tot}}$ and ${}^1C_{\mathrm{tot}}$.

The identity chain map $\mathrm{Id}: H \to H$ is defined to consist of the identity maps $\mathrm{Id}^{\varepsilon}_{\varepsilon^0}$ when $\varepsilon = (0, \dots, 0)$, and zero for other ε.

Lemma 5.16 *A chain map F between hyperboxes ${}^0H$ and ${}^1H$ induces a natural chain map $\hat{F}$ between the compressed hypercubes ${}^0\hat{H}$ and ${}^1\hat{H}$.*

Proof As mentioned above, the information in F can be used to build a new hyperbox of size $(\boldsymbol{d}, 1)$ composed of ${}^0H$ and ${}^1H$. Compressing this bigger hyperbox gives the required map $\hat{F}$. □

Definition 5.17 Let $F: {}^0H \to {}^1H$ and $G: {}^1H \to {}^2H$ be chain maps between hyperboxes of the same size. Their composite $G \circ F$ is defined to consist of the maps

$$(G \circ F)^{\varepsilon}_{\varepsilon^0} = \sum_{\varepsilon' \le \varepsilon} G^{\varepsilon-\varepsilon'}_{\varepsilon^0+\varepsilon'} \circ F^{\varepsilon'}_{\varepsilon^0}.$$

Definition 5.18 Let F and G be two chain maps between hyperboxes ${}^0H$ and ${}^1H$. A *chain homotopy* between F and G is a collection Ψ of linear maps

$$\Psi^{\varepsilon}_{\varepsilon^0}: {}^0C^{\varepsilon^0}_{*} \to {}^1C^{\varepsilon^0+\varepsilon}_{*+1+\|\varepsilon\|}$$

satisfying

$$F^{\varepsilon}_{\varepsilon^0} - G^{\varepsilon}_{\varepsilon^0} = \sum_{\varepsilon' \le \varepsilon} (D^{\varepsilon-\varepsilon'}_{\varepsilon^0+\varepsilon'} \circ \Psi^{\varepsilon'}_{\varepsilon^0} + \Psi^{\varepsilon-\varepsilon'}_{\varepsilon^0+\varepsilon'} \circ D^{\varepsilon'}_{\varepsilon^0})$$

for $\varepsilon^0 \in \mathbb{E}(\boldsymbol{d})$ and $\varepsilon \in \mathbb{E}_n$ such that $\varepsilon^0 + \varepsilon \in \mathbb{E}(\boldsymbol{d})$.

Note that we can also interpret a chain homotopy as a bigger hyperbox, namely an $(n+2)$-dimensional one of size $(\boldsymbol{d}, 1, 1)$, where $\varepsilon_{n+1} = 0$ and $\varepsilon_{n+1} = 1$ are the $(n+1)$-dimensional hyperboxes corresponding to F and G, respectively, and the maps in the new direction (from $\varepsilon_{n+2} = 0$ to $\varepsilon_{n+2} = 1$) are the identities (preserving ε_{n+1}) and Ψ (increasing ε_{n+1} by one).

We define a *chain homotopy equivalence* between hyperboxes as a chain map that has an inverse up to chain homotopy. The following follows from the same kind of argument as Lemma 5.16:

Lemma 5.19 *If $F : {}^0H \to {}^1H$ is a chain homotopy equivalence, then the compressed map $\widehat{F}$ is also a chain homotopy equivalence.*

Observe that, if $F : {}^0H \to {}^1H$ is a chain homotopy equivalence between hypercubes, the map F_{tot} between the respective total complexes is an ordinary chain homotopy equivalence.

Definition 5.20 A chain map $F : {}^0H \to {}^1H$ is called a *quasi-isomorphism of hyperboxes* if its components

$$F^{(0,\dots,0)}_{\varepsilon^0} : {}^0C^{\varepsilon^0} \to {}^1C^{\varepsilon^0}$$

induce isomorphisms on homology for all $\varepsilon^0 \in \mathbb{E}(\boldsymbol{d})$.

We note that a chain homotopy equivalence of hyperboxes is a quasi-isomorphism. Further, if $F : {}^0H \to {}^1H$ is a quasi-isomorphism between hypercubes, the total map F_{tot} is an ordinary quasi-isomorphism.

5.8 Elementary enlargements

For future reference, we introduce here a simple operation on hyperboxes, called *elementary enlargement*.

Let $H = ((C^{\varepsilon})_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}, (D^{\varepsilon})_{\varepsilon \in \mathbb{E}_n})$ be a hyperbox of chain complexes, of size $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$. Pick $k \in \{1, \dots, n\}$ and $j \in \{0, 1, \dots, d_k\}$. Define $\boldsymbol{d}^+ = (d_1^+, \dots, d_n^+) \in (\mathbb{Z}_{\geq 0})^n$ by $\boldsymbol{d}^+ = \boldsymbol{d} + \tau_k$, ie

$$d_i^+ = \begin{cases} d_i & \text{if } i \neq k, \\ d_i + 1 & \text{if } i = k. \end{cases}$$

We construct a new hyperbox

$$H^+ = ((C^{+,\varepsilon})_{\varepsilon \in \mathbb{E}(\boldsymbol{d}^+)}, (D^{+,\varepsilon})_{\varepsilon \in \mathbb{E}_n})$$

by replicating (ie introducing a new copy of) the complexes in positions with $\varepsilon_k = j$. The new copy will be in position $\varepsilon_k = j+1$, and everything with higher ε_k is shifted one step to the right. The two identical copies (which can be thought of as subhyperboxes) are linked by the identity chain map. Precisely, we set

$$C^{+,\varepsilon} = \begin{cases} C^{\varepsilon} & \text{if } \varepsilon_k \leq j, \\ C^{\varepsilon - \tau_k} & \text{if } \varepsilon_k \geq j+1, \end{cases} \qquad \text{and} \qquad D^{+,\varepsilon}_{\varepsilon^0} = \begin{cases} D^{\varepsilon}_{\varepsilon^0} & \text{if } \varepsilon^0_k + \varepsilon_k \leq j, \\ \mathrm{Id} & \text{if } \varepsilon^0_k = j,\ \varepsilon = \tau_k \\ 0 & \text{if } \varepsilon^0_k = j,\ \varepsilon_k = 1,\ \|\varepsilon\| \geq 2, \\ D^{\varepsilon}_{\varepsilon^0 - \tau_k} & \text{if } \varepsilon^0_k > j. \end{cases}$$

We then say that H^+ is obtained from H by an elementary enlargement at position (j, k). The following is easy to check from the definitions:

Lemma 5.21 *Let H^+ be an elementary enlargement of a hyperbox of chain complexes H. Then the compressed hypercubes $\widehat{H}$ and $\widehat{H}^+$ are identical.*

5.9 Canonical inclusions

We describe here yet another construction that will be useful to us later. We restrict to the case of hypercubes (since this is all we need), but everything can also be done in the more general context of hyperboxes.

Definition 5.22 Let (K_*, ∂) be a chain complex. The n-dimensional *canonical hypercube* $H(K, n)$ associated to (K_*, ∂) is defined to consist of the vector spaces

$$C_*^{\varepsilon} = K_*, \quad \varepsilon \in \mathbb{E}_n,$$

together with the maps

$$D_{\varepsilon^0}^{\varepsilon} = \begin{cases} \partial & \text{if } \|\varepsilon\| = 0, \\ \mathrm{Id} & \text{if } \|\varepsilon\| = 1, \\ 0 & \text{if } \|\varepsilon\| \geq 2. \end{cases}$$

Now suppose that $H = (C^{\varepsilon}, D^{\varepsilon})_{\varepsilon \in \mathbb{E}_n}$ is an arbitrary n-dimensional hypercube of chain complexes. Our aim is to construct a chain map

$$F_H^{\mathrm{can}} \colon H(C^{(0,\dots,0)}, n) \to H,$$

which will be called the *canonical inclusion*.

When $n = 0$ this is simply the identity map. When $n = 1$, the hypercube H consists of a single chain map $f = D^{(1)} \colon C^{(0)} \to C^{(1)}$ between two chain complexes. The canonical inclusion is then the square

$$\begin{tikzcd} C^{(0)} \arrow[r, "\mathrm{Id}"] \arrow[d, "\mathrm{Id}"'] & C^{(0)} \arrow[d, "f"] \\ C^{(0)} \arrow[r, "f"] & C^{(1)} \end{tikzcd}$$

with the diagonal map being trivial.

For general n, we construct the canonical inclusion as a composition of n different chain maps as follows. For $i \in \{0, \dots, n\}$ and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n) \in \mathbb{E}_n$, we let $\varepsilon[\leq i] \in \mathbb{E}_n$ be the multi-index obtained from ε by changing all entries above i to be zero. In other words,

$$\varepsilon[\leq i]_j = \begin{cases} \varepsilon_j & \text{if } j \leq i, \\ 0 & \text{if } j > i. \end{cases}$$

Similarly, we let $\varepsilon[> i]$ to be obtained from ε by setting all entries less than or equal to i to be zero.

We define a hypercube $H[i]$ to consist of the chain groups

$$C[i]^{\varepsilon} = C^{\varepsilon[\leq i]}, \quad \varepsilon \in \mathbb{E}_n,$$

and the maps

$$D[i]_{\varepsilon}^{\varepsilon'-\varepsilon} = \begin{cases} D_{\varepsilon[\le i]}^{(\varepsilon'-\varepsilon)[\le i]} & \text{if } \varepsilon[> i] = \varepsilon'[> i], \\ \mathrm{Id} & \text{if } \varepsilon[\le i] = \varepsilon'[\le i] \text{ and } \|\varepsilon'[> i] - \varepsilon[> i]\| = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $H[0]$ is the canonical hypercube $H(C^{(0,\dots,0)}, n)$, while $H[n] = H$.

For $i = 1, \dots, n$, we define chain maps

$$F[i] \colon H[i-1] \to H[i]$$

to consist of

$$F[i]_{\varepsilon}^{\varepsilon'-\varepsilon} = \begin{cases} D_{\varepsilon[\le (i-1)]}^{\varepsilon'[\le i]-\varepsilon[\le (i-1)]} & \text{if } \varepsilon_i = 1,\ \varepsilon[> i] = \varepsilon'[> i], \\ \mathrm{Id} & \text{if } \varepsilon = \varepsilon' \text{ and } \varepsilon_i = 0, \\ 0 & \text{otherwise.} \end{cases}$$

The canonical inclusion is then

$$F_H^{\mathrm{can}} = F[n] \circ F[n-1] \circ \cdots \circ F[1].$$

6 Quasistabilizations

In this section we introduce a new move that relates certain equivalent, multipointed Heegaard diagrams, called quasistabilization. Basically, a quasistabilization is the composition of a free index zero/three stabilization and some handleslides.

Our goal is to study how the polygon maps on Heegaard Floer complexes behave under this move. There are two motivations for this. First, the behavior of polygon maps under ordinary (free) index zero/three stabilizations (which can be viewed as particular examples of quasistabilizations) is one of the inputs in the construction of complete systems of hyperboxes in Section 8, as well as in the proof of the surgery theorem, Theorem 9.6. Second, the more general quasistabilizations are needed in Section 15, where they appear in the context of grid diagrams.

For concreteness, we will first describe the results in the case of triangles (for general quasistabilizations). Then we will explain how similar arguments can be used to study higher polygon maps. At the end we will specialize to the case of ordinary index zero/three stabilizations.

6.1 The setup

Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ be a multipointed Heegaard diagram, as in Section 3.1. Suppose $\mathcal{H}$ represents a link $\vec{L}$ in an integral homology sphere Y. Fix $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, so that we have a well-defined generalized Floer complex $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) = \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s})$.

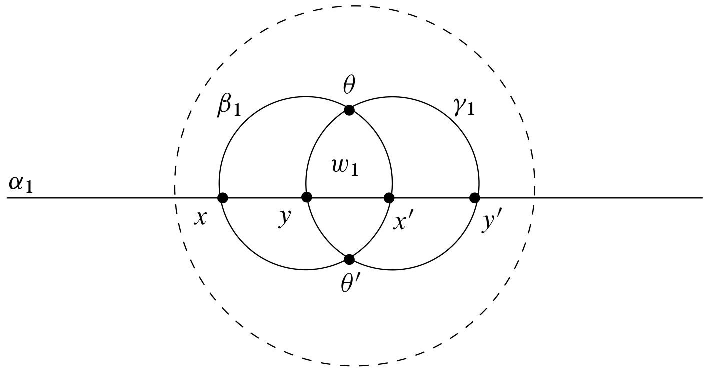

Figure 9: Quasistabilization. We show here a part of the triple Heegaard diagram described in Section 6.1. In Section 6.5, we will stretch the complex structure along the dashed curve.

Let g be the genus of Σ and d the number of alpha (or beta) curves. We assume that $\beta_1 \in \boldsymbol{\beta}$ bounds a disk containing a free basepoint w_1, and that the only alpha curve intersecting β_1 is α_1, which does so at two points x and x'. On the other hand, α_1 can intersect other beta curves. We also require that the homology class $[\alpha_1]$ lies in the span of the other $[\alpha_i]$ in $H_1(\Sigma)$.

Let $\overline{\mathscr{H}} = (\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\boldsymbol{w}}, \boldsymbol{z})$ be the diagram obtained from $\mathscr{H}$ by deleting α_1, β_1 and w_1. We then say that $\overline{\mathscr{H}}$ is obtained from $\mathscr{H}$ by *quasidestabilization*. The reverse process is called *quasistabilization*.

Remark 6.1 By handlesliding α_1 over other alpha curves, we can arrange that it does not intersect any beta curve except β_1. The resulting diagram is then a usual free index zero/three stabilization of $\overline{\mathscr{H}}$; see [Ozsváth and Szabó 2008a].

Consider now an extra collection of d attaching curves $\boldsymbol{\gamma}$ on Σ such that $\gamma_1 \in \boldsymbol{\gamma}$ has the same properties as β_1: it bounds a disk containing w_1, and the only alpha curve that it intersects is α_1, with the respective intersection consisting of two points y and y'. Furthermore, we assume that γ_1 is a small Hamiltonian translate of β_1, and intersects β_1 in two points θ and θ', as in Figure 9. We assume that the relative positions of x, x', y, y', θ and θ' are exactly as in the figure.

Let $\bar{\boldsymbol{\gamma}}$ be the collection of curves obtained from $\boldsymbol{\gamma}$ by removing γ_1. Then $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{z})$ is a quasistabilization of $(\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\gamma}}, \bar{\boldsymbol{w}}, \boldsymbol{z})$. For any $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, the intersection $\boldsymbol{x} \cap \alpha_1$ is either x or x'. We denote by $\bar{\boldsymbol{x}} \in \mathbb{T}_{\bar{\alpha}} \cap \mathbb{T}_{\bar{\beta}}$ the generator obtained from $\boldsymbol{x}$ by deleting the point in $\boldsymbol{x} \cap \alpha_1$. Similarly, for $\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, there is a corresponding generator $\bar{\boldsymbol{y}}$ in $\mathbb{T}_{\bar{\alpha}} \cap \mathbb{T}_{\bar{\gamma}}$, obtained by deleting y or y'.

Pick an intersection point $\boldsymbol{\theta}$ in $\mathbb{T}_\beta \cap \mathbb{T}_\gamma$ such that $\theta \in \boldsymbol{\theta}$. We have a map

$$F \colon \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{s}), \quad F(\boldsymbol{x}) = f(\boldsymbol{x} \otimes \boldsymbol{\theta}),$$

which counts index zero pseudoholomorphic triangles with one vertex at $\boldsymbol{\theta}$, as in Section 3.5.

Set $\bar{\boldsymbol{\theta}} = \boldsymbol{\theta} - \{\theta\} \in \mathbb{T}_{\bar{\beta}} \cap \mathbb{T}_{\bar{\gamma}}$. There is a corresponding triangle map in the quasidestabilized diagram

$$\bar{F} \colon \boldsymbol{A}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\beta}}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\gamma}}, \boldsymbol{s}), \quad \bar{F}(\bar{\boldsymbol{x}}) = f(\bar{\boldsymbol{x}} \otimes \bar{\boldsymbol{\theta}}).$$

Next, we define a map $G\colon \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{s})$ as follows. Writing the coefficients of $\bar{F}$ as $n_{\bar{\boldsymbol{x}},\bar{y}}$, so that, for $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$,

$$\bar{F}(\bar{\boldsymbol{x}}) = \sum_{\bar{\boldsymbol{y}} \in \mathbb{T}_{\bar{\alpha}} \cap \mathbb{T}_{\bar{\gamma}}} n_{\bar{\boldsymbol{x}},\bar{\boldsymbol{y}}} \bar{\boldsymbol{y}},$$

we set

$$G(\boldsymbol{x}) = \sum_{\bar{\boldsymbol{y}} \in \mathbb{T}_{\bar{\alpha}} \cap \mathbb{T}_{\bar{\gamma}}} n_{\bar{\boldsymbol{x}},\bar{\boldsymbol{y}}} (\bar{\boldsymbol{y}} \cup (\boldsymbol{x} \cap \alpha_1)).$$

Proposition 6.2 *For a quasistabilized triple Heegaard diagram* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{z})$ *as above, and suitable almost complex structures on the symmetric products, the maps* F *and* G *coincide.*

The proof of Proposition 6.2 will occupy Sections 6.2–6.5.

6.2 Cylindrical formulations

We recall Lipshitz's cylindrical formulation [2006] of Heegaard Floer homology; see also [Ozsváth and Szabó 2008a, Section 5.2]. Instead of holomorphic strips $[0,1] \times \mathbb{R} \to \mathrm{Sym}^d(\Sigma)$ with boundaries on $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$, Lipshitz considers pseudoholomorphic maps from a Riemann surface S (with boundary) to the target

$$W = \Sigma \times [0,1] \times \mathbb{R}.$$

The four-manifold W admits two natural projection maps

$$\pi_\Sigma\colon W \to \Sigma \quad \text{and} \quad \pi_{\mathbb{D}}\colon W \to [0,1] \times \mathbb{R}.$$

The notation $\pi_{\mathbb{D}}$ refers to the unit disk $\mathbb{D} \subset \mathbb{C}$, which can be viewed as the conformal compactification of $[0,1] \times \mathbb{R}$, obtained by adding the points $\pm i$.

We equip W with an almost complex structure J translation invariant in the $\mathbb{R}$-factor, and such that $\pi_{\mathbb{D}}$ is a pseudoholomorphic map. Further, we ask for J to be tamed by a natural split symplectic form on W. Typically, we choose J to be a small perturbation of a split complex structure $j_\Sigma \times j_{\mathbb{D}}$, where j_Σ and $j_{\mathbb{D}}$ are complex structures on Σ and $[0,1] \times \mathbb{R}$, respectively. Sometimes (for example, to ensure positivity of intersections) it will be convenient to require J to be split on U, that is, split on $U \times [0,1] \times \mathbb{R}$, where $U \subset \Sigma$ is an open subset.

Definition 6.3 An *annoying curve* is a pseudoholomorphic curve in W contained in a fiber of $\pi_{\mathbb{D}}$.

To define the differential on the cylindrical Heegaard Floer complex, Lipshitz uses pseudoholomorphic maps

$$u\colon S \to W = \Sigma \times [0,1] \times \mathbb{R}$$

with the following properties:

- S is a Riemann surface with boundary and $2d$ punctures on its boundary, of two types: d "positive" punctures $\{p_1, \dots, p_d\}$ and d "negative" punctures $\{q_1, \dots, q_d\}$.
- u is a smooth embedding.

- $u(\partial S) \subset (\boldsymbol{\alpha} \times \{1\} \times \mathbb{R}) \cup (\boldsymbol{\beta} \times \{0\} \times \mathbb{R})$.
- u has finite energy (where the energy is defined as in [Bourgeois et al. 2003, Section 6.3]).
- For each i, $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ and $u^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ consist of exactly one component of $\partial S - \{p_1, \dots, p_d, q_1, \dots, q_d\}$.
- No components of the image $u(S)$ are annoying curves.
- Any sequence of points in S converging to q_i (resp. p_i) is mapped under $\pi_{\mathbb{D}}$ to a sequence of points whose second coordinate converges to $-\infty$ (resp. $+\infty$).

Curves of this kind are called *cylindrical flow lines*. Any cylindrical flow line u can be extended to a map $\bar{u}$ from the closure of S to the compactification $\Sigma \times \mathbb{D}$. The image of this extension contains the points $x \times \{-i\}$ and $y \times \{i\}$ for $x \in \boldsymbol{x}$ and $y \in \boldsymbol{y}$, where $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \subset \mathrm{Sym}^d(\Sigma)$. We then say that u connects $\boldsymbol{x}$ to $\boldsymbol{y}$.

To every cylindrical flow line $u\colon S \to W$ one can associate a strip $\tilde{u}\colon [0,1] \times \mathbb{R} \to \mathrm{Sym}^d(\Sigma)$ with boundaries on $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$, by setting $\tilde{u}(z) = \pi_\Sigma((\pi_{\mathbb{D}} \circ u)^{-1}(z))$. Thus, cylindrical flow lines can be organized according to moduli spaces $\mathcal{M}(\phi)$, indexed by homology classes $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ for the corresponding Whitney disks. Moreover, Lipshitz [2006, Appendix A] identifies the moduli spaces of cylindrical flow lines in a class ϕ with the respective moduli spaces of pseudoholomorphic strips (ordinary flow lines), for suitable almost complex structures, in the case when the Maslov index $\mu(\phi)$ is one. It follows that the Heegaard Floer complex can be defined just as well by counting cylindrical flow lines instead of pseudoholomorphic strips.

When studying degenerations of cylindrical flow lines (for example, in the proof that $\partial^2 = 0$ in the cylindrical setting), we also encounter maps of the following kind:

Definition 6.4 Consider a Riemann surface S with boundary and d punctures $\{p_1, \dots, p_d\}$ on its boundary. A (cylindrical) *boundary degeneration* is a pseudoholomorphic map $u\colon S \to \Sigma \times (-\infty, 1] \times \mathbb{R}$ which has finite energy, is a smooth embedding, sends ∂S into $\boldsymbol{\alpha} \times \{1\} \times \mathbb{R}$, contains no component in the fiber of the projection to $(-\infty, 1] \times \mathbb{R}$, and has the property that each component of $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ consists of exactly one component of $\partial S \setminus \{p_1, \dots, p_d\}$. A similar definition can be made with β playing the role of α and using the interval $[0, \infty)$ instead of $(-\infty, 1]$.

Note that, for a boundary degeneration u, the points at infinity must be mapped to a fixed $\boldsymbol{x} \in \mathbb{T}_\alpha$. The boundary degenerations with endpoint $\boldsymbol{x}$ can be organized into moduli spaces $\mathcal{N}(\psi)$ according to homology classes $\psi \in \pi_2^\alpha(\boldsymbol{x}) \cong H_2(\Sigma, \boldsymbol{\alpha})$.

Next let us recall from Section 3.5 that, when one has three collections of curves $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ on a fixed Heegaard surface with marked basepoints $(\Sigma, \boldsymbol{w})$, one can define a map

$$f = f_{\alpha\beta\gamma}\colon \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s}) \otimes \boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_\gamma, \boldsymbol{0}) \to \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{s})$$

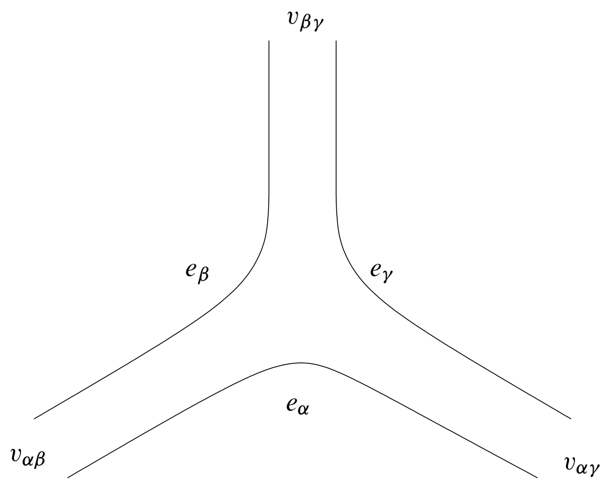

Figure 10: The triangular region Δ. This region is conformally equivalent to a triangle with punctures at the vertices.

by counting index zero pseudoholomorphic triangles in $\mathrm{Sym}^d(\Sigma)$, with boundaries on $\mathbb{T}_\alpha$, $\mathbb{T}_\beta$ and $\mathbb{T}_\gamma$. These maps admit a cylindrical formulation, too; see [Lipshitz 2006, Section 10]. Indeed, consider a contractible subset $\Delta \subset \mathbb{C}$ as in Figure 10, with three boundary components e_α, e_β and e_γ, and three infinite ends $v_{\alpha\beta}$, $v_{\beta\gamma}$ and $v_{\alpha\gamma}$, all diffeomorphic to $[0,1] \times (0,\infty)$. Setting

$$W_\Delta = \Sigma \times \Delta,$$

note that there are natural projections π_Σ and π_Δ to the two factors. We equip W_Δ with an almost complex structure having properties analogous to those of the almost complex structure on W.

We then consider pseudoholomorphic maps

$$u \colon S \to W_\Delta$$

with the following properties:

- S is a Riemann surface with boundary and $3d$ punctures $p_i^{\alpha\beta}$, $p_i^{\beta\gamma}$ and $p_i^{\alpha\gamma}$ for $i \in \{1, \dots, d\}$ on the boundary ∂S.
- u is a smooth embedding.
- $u(\partial S) \subset (\boldsymbol{\alpha} \times e_\alpha) \cup (\boldsymbol{\beta} \times e_\beta) \cup (\boldsymbol{\gamma} \times e_\gamma)$.
- u has finite energy.
- For each $i = 1, \dots, d$ and $\sigma \in \{\alpha, \beta, \gamma\}$, the preimage $u^{-1}(\sigma_i \times e_\sigma)$ consists of exactly one component of the punctured boundary of S.
- No components of the image $u(S)$ are annoying curves.
- Any sequence of points in S converging to $p_i^{\alpha\beta}$ (resp. $p_i^{\beta\gamma}$, $p_i^{\alpha\gamma}$) is mapped under π_Δ to a sequence of points converges towards infinity in the strip-like end $v_{\alpha\beta}$ (resp. $v_{\beta\gamma}$, $v_{\alpha\gamma}$).

Maps u like this are called *cylindrical triangles*. They can be organized into moduli spaces $\mathcal{M}(\phi)$ according to homology classes $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ for $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\boldsymbol{y} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$ and $\boldsymbol{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$. The moduli space of such maps in classes ϕ with $\mu(\phi) = 0$ can be identified with the moduli space of ordinary pseudoholomorphic triangles in $\mathrm{Sym}^d(\Sigma)$, as used in the definition of the map $f = f_{\alpha\beta\gamma}$. (This can be proved by the same methods as in [Lipshitz 2006, Section 13].) It follows that f can be described in terms of counts of cylindrical triangles.

Similar descriptions can be given to the higher polygon maps from Section 3.5.

6.3 Domains

Let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ be a homology class of Whitney disks in a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w})$. The curves $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ split the Heegaard surface into several connected components $R_1, \dots, R_r$, which we call *regions*. A *domain* $\mathcal{D}$ on the Heegaard diagram is by definition a linear combination of regions, with integer coefficients. The class ϕ has an associated *domain* $\mathcal{D}(\phi)$ — see [Ozsváth and Szabó 2004b, Section 3.5] —

$$\mathcal{D}(\phi) = \sum_{i=1}^{r} (\phi \cdot [\{z_i\} \times \mathrm{Sym}^{d-1}(\Sigma)]), R_i,$$

where z_i is a point chosen in the interior of the region R_i and $\cdot$ denotes intersection product.

Let $\mathcal{D}$ be a domain and $x \in \alpha_i \cap \beta_j$ an intersection point for some i and j. A neighborhood of x is split by α_i and β_j into four quadrants. Two of the four quadrants have the property that, as we move counterclockwise around x, we first see α_i on their boundary and then β_j; for the other two quadrants, we first see β_j and then α_i. Let $m^{\alpha\beta}(\mathcal{D}, x)$ be the sum of the multiplicities of $\mathcal{D}$ in the two quadrants of the first type, and $m^{\beta\alpha}(\mathcal{D}, x)$ the sum of multiplicities in the other two quadrants.

Given a point $x \in \alpha_i \cap \beta_j$ and a d-tuple $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we set

$$\delta(\boldsymbol{x}, x) = \begin{cases} 1 & \text{if } x \in \boldsymbol{x}, \\ 0 & \text{otherwise.} \end{cases}$$

Definition 6.5 Let $\mathcal{D}$ be a domain on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The domain $\mathcal{D}$ is said to be *acceptable for the pair* $(\boldsymbol{x}, \boldsymbol{y})$ if, for every $i, j = 1, \dots, d$ and $x \in \alpha_i \cap \beta_j$, we have

$$m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = \delta(\boldsymbol{x}, x) - \delta(\boldsymbol{y}, x). \tag{61}$$

The proof of the following lemma is straightforward:

Lemma 6.6 *A domain $\mathcal{D}$ is acceptable for the pair $(\boldsymbol{x}, \boldsymbol{y})$ if and only if it is of the form $\mathcal{D}(\phi)$ for some $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$.*

We now turn to the Maslov index $\mu(\phi)$, which is the expected dimension of the moduli space of pseudoholomorphic representatives of $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$. The Maslov index can be calculated in terms of the domain $\mathcal{D} = \mathcal{D}(\phi)$ using the following formula, due to [Lipshitz 2006, Corollary 4.3]:

$$\mu(\phi) = e(\mathcal{D}) + \sum_{x \in \boldsymbol{x}} n_x(\mathcal{D}) + \sum_{y \in \boldsymbol{y}} n_y(\mathcal{D}). \tag{62}$$

Here, $n_p(\mathcal{D})$ denotes the average multiplicity of $\mathcal{D}$ in the four quadrants around a point p, while $e(\mathcal{D})$ is the Euler measure of the domain, as defined in [ibid.].

Now consider a boundary degeneration class $\psi \in \pi_2^\alpha(\boldsymbol{x})$. Its domain is then an *α-periodic domain* $\mathcal{P} = \mathcal{P}(\psi) \in H_2(\Sigma, \boldsymbol{\alpha})$, ie a linear combination of components of $\Sigma - \boldsymbol{\alpha}$. In fact, there is a one-to-one correspondence between periodic domains and classes in $\pi_2^\alpha(\boldsymbol{x})$. The respective Maslov index is given by

$$\mu(\psi) = e(\mathcal{P}) + 2\sum_{x \in \boldsymbol{x}} n_x(\mathcal{P}). \tag{63}$$

We also have an alternative characterization; see [Ozsváth and Szabó 2008a, Lemma 5.4]. Recall that we have a basepoint w_i in each component of $\Sigma - \boldsymbol{\alpha}$. Then

$$\mu(\psi) = 2\sum_{i=1}^{d-g+1} n_{w_i}(\mathcal{P}). \tag{64}$$

Next, we turn to homology classes of triangles. Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{w}, \boldsymbol{z})$ be a triple Heegaard diagram, with each curve collection consisting of d curves. By *regions* we now mean the connected components of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma})$. Given a homology class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ for $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\boldsymbol{y} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$ and $\boldsymbol{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, its domain $\mathcal{D} = \mathcal{D}(\phi)$ is defined as before. We have analogues of Definition 6.5 and Lemma 6.6:

Lemma 6.7 *The necessary and sufficient condition for a domain $\mathcal{D}$ to be of the form $\mathcal{D}(\phi)$ for some $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ is that $\mathcal{D}$ be* ***acceptable*** *for the triple $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$, that is, it should satisfy*

$$m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = \delta(\boldsymbol{x}, x) \quad \text{for } x \in \alpha_i \cap \beta_j, \tag{65}$$

$$m^{\beta\gamma}(\mathcal{D}, y) - m^{\gamma\beta}(\mathcal{D}, y) = \delta(\boldsymbol{y}, y) \quad \text{for } y \in \beta_i \cap \gamma_j, \tag{66}$$

$$m^{\gamma\alpha}(\mathcal{D}, z) - m^{\alpha\gamma}(\mathcal{D}, z) = \delta(\boldsymbol{z}, z) \quad \text{for } z \in \gamma_i \cap \alpha_j. \tag{67}$$

The vertex multiplicities of a domain are defined as in the case of bigons. We can similarly define the Euler measure. Further, we let $a(\mathcal{D})$ denote the intersection $\partial\mathcal{D} \cap \boldsymbol{\alpha}$, viewed as a one-chain on Σ, supported on $\boldsymbol{\alpha}$. Similarly, we define $b(\mathcal{D}) = \partial\mathcal{D} \cap \boldsymbol{\beta}$ and $c(\mathcal{D}) = \partial\mathcal{D} \cap \boldsymbol{\gamma}$. We let $a(\mathcal{D}).c(\mathcal{D})$ denote the average of the four algebraic intersection numbers between $a'(\mathcal{D})$ and $c(\mathcal{D})$, where $a'(\mathcal{D})$ is a small translate of $a(\mathcal{D})$ in any of the four "diagonal" directions off $\boldsymbol{\alpha}$, such that no endpoint of $a(\mathcal{D})$ lies on $\boldsymbol{\gamma}$, and no endpoint of $c(\mathcal{D})$ lies on $\boldsymbol{\alpha}$. We could similarly define $b(\mathcal{D}).a(\mathcal{D})$ or $c(\mathcal{D}).b(\mathcal{D})$. Sarkar [2011, Theorem 4.1] proved the following formula for the index of holomorphic triangles:[3]

$$\mu(\phi) = e(\mathcal{D}) + \sum_{x \in \boldsymbol{x}} n_x(\mathcal{D}) + \sum_{y \in \boldsymbol{y}} n_y(\mathcal{D}) - a(\mathcal{D}).c(\mathcal{D}) - \tfrac{1}{2}d. \tag{68}$$

[3]Our conventions are different from those in [Sarkar 2011]. There, the sides of a holomorphic triangle were on α, β and γ in counterclockwise order, whereas we have them clockwise. Thus, Sarkar's expression $a(\mathcal{D}).c(\mathcal{D})$ differs from ours by a negative sign. The formula (68) above is written with our conventions.

Sarkar also gave a generalization of this formula to higher polygons. Suppose we have curve collections $\boldsymbol{\eta}^i$ for $i = 0, \dots, l$ on a pointed surface $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$ such that each collection consists of d curves. We then consider a homotopy class of $(l+1)$-gons $\phi \in \pi_2(\boldsymbol{x}^0, \dots, \boldsymbol{x}^l)$, where $\boldsymbol{x}^i \in \mathbb{T}_{\eta^i} \cap \mathbb{T}_{\eta^{i+1}}$ for $i < l$ and $\boldsymbol{x}^l \in \mathbb{T}_{\eta^0} \cap \mathbb{T}_{\eta^l}$. We can define $\mathcal{D}(\phi)$ as before, and we have acceptability conditions similar to (65)–(67). We let $a^i(\mathcal{D}) = \partial\mathcal{D} \cap \boldsymbol{\eta}^i$. With our conventions, Theorem 4.1 in [Sarkar 2011] says

$$(69)\quad \mu(\phi) = e(\mathcal{D}) + \sum_{x \in \boldsymbol{x}^0} n_x(\mathcal{D}) + \sum_{y \in \boldsymbol{x}^1} n_y(\mathcal{D}) - a^0(\mathcal{D}). \sum_{j=2}^{l} a^j(\mathcal{D}) - \sum_{j > k > 1} a^j(\mathcal{D}).a^k(\mathcal{D}) - \tfrac{1}{2} d(l-1).$$

Here, the Maslov index $\mu(\phi)$ is as defined in Section 3.5.

6.4 Convergence and gluing for the moduli spaces of triangles

Let $(\Sigma^1, \boldsymbol{\alpha}^1, \boldsymbol{\beta}^1, \boldsymbol{\gamma}^1)$ and $(\Sigma^2, \boldsymbol{\alpha}^2, \boldsymbol{\beta}^2, \boldsymbol{\gamma}^2)$ be two triple Heegaard diagrams. (For the purposes of this subsection, we can ignore the basepoints.) For $i = 1, 2$, we let d_i be the number of curves in the collection α^i (or β^i, or γ^i), and $g_i \le d_i$ the genus of Σ^i.

Consider an extra simple closed curve α_s^1 on Σ^1 that is disjoint from the other curves in $\boldsymbol{\alpha}^1$ and lies in their homological span. Set

$$\boldsymbol{\alpha}^{1+} = \boldsymbol{\alpha}^1 \cup \{\alpha_s^1\}.$$

Pick also one of the curves in $\boldsymbol{\alpha}^2$, say α_1^2, and call it α_s^2. The subscript s stands for "special".

Pick points $p_i \in \alpha_s^i$ for $i = 1, 2$ that do not lie on any of the beta or gamma curves. We form the connected sum $\Sigma = \Sigma_1 \# \Sigma_2$ at p_1 and p_2, of genus $g = g_1 + g_2$. By joining each of the two ends of α_s^1 at p_1 with the respective end of α_s^2 at p_2, we obtain a new curve $\alpha_s = \alpha_s^1 \# \alpha_s^2$ on Σ. We set

$$\boldsymbol{\alpha} = \boldsymbol{\alpha}^1 \cup (\boldsymbol{\alpha}^2 - \{\alpha_s^2\}) \cup \{\alpha_s\}.$$

This is a collection of $d = d_1 + d_2$ attaching curves on Σ. We can also form collections $\boldsymbol{\beta} = \boldsymbol{\beta}^1 \cup \boldsymbol{\beta}^2$ and $\boldsymbol{\gamma} = \boldsymbol{\gamma}^1 \cup \boldsymbol{\gamma}^3$. Together, they turn Σ into a triple Heegaard diagram, which we call the *special connected sum* of $(\Sigma^1, \boldsymbol{\alpha}^1, \boldsymbol{\beta}^1, \boldsymbol{\gamma}^1)$ and $(\Sigma^2, \boldsymbol{\alpha}^2, \boldsymbol{\beta}^2, \boldsymbol{\gamma}^2)$.

Example 6.8 A triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ as in Section 6.1 can be viewed as the special connected sum of the diagram $(\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\gamma}})$ with the genus zero diagram $(\mathcal{S}, \alpha_1, \beta_1, \gamma_1)$ shown in Figure 11. The notation is as in Section 6.1.

Given a special connected sum of triple Heegaard diagrams, note that intersection points $\boldsymbol{x}^1 \in \mathbb{T}_{\alpha^1} \cap \mathbb{T}_{\beta^1}$ and $\boldsymbol{x}^2 \in \mathbb{T}_{\alpha^2} \cap \mathbb{T}_{\beta^2}$ give rise to an intersection point $\boldsymbol{x}^1 \times \boldsymbol{x}^2 \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Conversely, any $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is of this form, because none of the points on $\alpha_s^1 \cap \beta_i^1$ for $i = 1, \dots, d_1$ can be part of $\boldsymbol{x}$. Similar remarks apply to generators of the form $\boldsymbol{y} = \boldsymbol{y}_1 \times \boldsymbol{y}_2 \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$ and $\boldsymbol{z} = \boldsymbol{z}_1 \times \boldsymbol{z}_2 \in \mathbb{T}_\gamma \cap \mathbb{T}_\alpha$.

Given a homology class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ on the special connected sum, we denote by $m_1 = m_1(\mathcal{D})$ and $m_2 = m_2(\mathcal{D})$ the multiplicities of $\mathcal{D}$ on each side of the curve α_s, near the connected sum neck. (See Figure 11 for an example.)

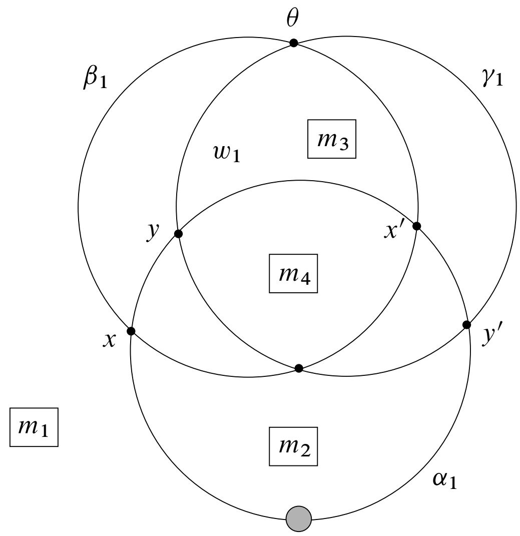

Figure 11: The sphere $\mathcal{S}$. This is the one-point compactification of the disk bounded by the dashed curve in Figure 9. It could be viewed as a triple Heegaard diagram, except it is missing a basepoint. The compactification point p^2 is shown by a gray dot. The four boxes show the multiplicities of a triangular domain inside the corresponding regions.

Lemma 6.9 *For a triple Heegaard diagram obtained as a special connected sum as above, pick* $\boldsymbol{x} = \boldsymbol{x}^1 \times \boldsymbol{x}^2 \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, $\boldsymbol{y} = \boldsymbol{y}_1 \times \boldsymbol{y}_2 \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$ *and* $\boldsymbol{z} = \boldsymbol{z}_1 \times \boldsymbol{z}_2 \in \mathbb{T}_\gamma \cap \mathbb{T}_\alpha$. *Then there is a natural surjective map*

$$p\colon \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}) \to \pi_2(\boldsymbol{x}^2, \boldsymbol{y}^2, \boldsymbol{z}^2). \tag{70}$$

Furthermore, a choice of a domain $\mathcal{P} \in H_2(\Sigma, \boldsymbol{\alpha}^{1+})$ *whose boundary contains* α_s^1 *with multiplicity one gives rise to a map*

$$\sigma\colon \pi_2(\boldsymbol{x}^2, \boldsymbol{y}^2, \boldsymbol{z}^2) \to \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$$

such that $p \circ \sigma = \mathrm{id}$.

Proof We use the identification between homology classes $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ and acceptable domains $\mathcal{D} = \mathcal{D}(\phi)$; see Lemma 6.7. An acceptable domain $\mathcal{D}$ on the special connected sum gives rise to an acceptable domain $\mathcal{D}_2$ on Σ_2 with respect to $(\boldsymbol{x}^2, \boldsymbol{y}^2, \boldsymbol{z}^2)$, by restriction. This produces the map p.

Given $\mathcal{P}$ is as in the statement of the lemma, the desired map s takes an acceptable domain $\mathcal{D}_2$ on Σ_2 to the acceptable domain $\mathcal{D} = \mathcal{D}_2 + (m_1(\mathcal{D}_2) - m_2(\mathcal{D}_2))\mathcal{P}$ on the special connected sum. □

Let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ have a domain $\mathcal{D}$. We define an equivalence relation on pairs $(\phi^1, \mathcal{P})$, where $\phi^1 \in \pi_2(\boldsymbol{x}^1, \boldsymbol{y}^1, \boldsymbol{z}^1)$ and $\mathcal{P} \in H_2(\Sigma, \boldsymbol{\alpha}^{1+})$ has α_s^1 with multiplicity $m_1(\mathcal{D}) - m_2(\mathcal{D})$ on its boundary. Two pairs $(\phi_1^1, \mathcal{P}_1)$ and $(\phi_2^1, \mathcal{P}_2)$ are set to be equivalent if $\phi_1^1 + \mathcal{P}_1 = \phi_2^1 + \mathcal{P}_2$, as two-chains on Σ^1.

By identifying homology classes with acceptable domains (as in the proof of Lemma 6.9), we see that every ϕ determines a unique such equivalence class ϕ^{1+}. We set

$$\mu(\phi^{1+}) = \mu(\phi^1) + \mu(\mathcal{P})$$

for any $(\phi^1, \mathcal{P}) \in \phi^{1+}$.

Lemma 6.10 *Let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ be a homology class of triangles in a triple Heegaard diagram obtained by special connected sum, as above. Let $\phi^2 \in \pi_2(\boldsymbol{x}^2, \boldsymbol{y}^2, \boldsymbol{z}^2)$ be its restriction to Σ^2, and ϕ^{1+} the corresponding equivalence class of pairs on Σ^1. Then*

$$\mu(\phi) = \mu(\phi^{1+}) + \mu(\phi^2) - m_1(\mathcal{D}) - m_2(\mathcal{D}),$$

where $\mathcal{D} = \mathcal{D}(\phi)$ is the domain of ϕ.

Proof In Sarkar's formula (68), all terms except $e(\mathcal{D})$ are additive under the special connected sum. When adding up the Euler measures, we have to subtract $m_1(\mathcal{D}) + m_2(\mathcal{D})$ because doing the special connected sum involves deleting two disks, each made of two bigons. Two of these four bigons have multiplicity $m_1(\mathcal{D})$, the other two $m_2(\mathcal{D})$, and the Euler measure of a bigon is $\frac{1}{2}$. □

We now proceed to study holomorphic triangles on a special connected sum. We will use Lipshitz's cylindrical formulation from Section 6.2.

Note that, if a homology class ϕ (of cylindrical flow lines, boundary degenerations, triangles, etc) admits pseudoholomorphic representatives, the principle of positivity of intersections implies that the domain $\mathcal{D}(\phi)$ is a linear combination of regions with only nonnegative coefficients:

$$\mathcal{D}(\phi) \geq 0;$$

see [Ozsváth and Szabó 2004b, Lemma 3.2].

In addition to the cylindrical flow lines, boundary degenerations and triangles from Section 6.2, when discussing special connected sums we will also need to study some new objects:

Definition 6.11 Consider a Heegaard surface Σ and a collection of attaching circles $\boldsymbol{\alpha}$ on Σ. An *annoying α-degeneration* is a holomorphic curve $u\colon S \to W_\Delta = \Sigma \times \Delta$ such that S is a connected Riemann surface with boundary and punctures on the boundary, and there exists an unpunctured component $\partial_0 S$ of the boundary ∂S satisfying $u(\partial_0 S) \subset \boldsymbol{\alpha} \times e_\alpha$. Here, Δ is as in Figure 10.

Lemma 6.12 *An annoying α-degeneration $u\colon S \to W_\Delta$ is an annoying curve in the sense of Definition 6.3, that is, all of S is mapped to a fixed point $p \in e_\alpha$ under $\pi_\Delta \circ u$.*

Proof Let $D(S)$ be the double of S taken along the component $\partial_0 S \subseteq \partial S$, and $D(\Delta)$ the double of Δ along e_α. We can extend $\pi_\Delta \circ u$ to a holomorphic map $f\colon D(S) \to D(\Delta)$ using Schwartz reflection. Since $\partial_0 S$ is compact and e_α is not, there exists some $z_0 \in \partial_0 S$ with $f'(z_0) = 0$. If f were not constant, it would have a branch point of order $k \geq 2$ at z_0. This contradicts the local model near z_0, which is that of a holomorphic function f mapped to $\mathbb{C}$ such that $\operatorname{Re} f(z) > 0$ for $\operatorname{Re} z > 0$. □

Annoying α-degenerations can be organized according to their domains, which are relative homology classes $\mathcal{P} \in H_2(\Sigma, \boldsymbol{\alpha})$. The domain $\mathcal{P}$ must be nonnegative. Even though there are no transversality results for annoying curves, compare [Lipshitz 2006, Section 3], one can still define the Maslov index $\mu(\mathcal{P})$ according to (63), by treating $\mathcal{P}$ as in the case of usual boundary degenerations.

We will mostly be interested in annoying α-degenerations on the first surface Σ^1 that is part of the special connected sum. The degenerations will be taken with respect to the collection of curves $\boldsymbol{\alpha}^{1+}$ that includes α_s^1. For an annoying α-degeneration of this form, with domain $\mathcal{P}$, and any $\boldsymbol{x} \in \mathbb{T}_{\alpha^1}$, we have (63). However, in (63) there are no vertex multiplicity contributions from the curve α_s^1, so (64) needs to be modified accordingly. Precisely, if we place basepoints $w_1, \dots, w_{d^1-g^1}$ in all components of $\Sigma^1 - \boldsymbol{\alpha}^{1+}$ except the two that have α_s^1 on their boundary, we obtain

$$\mu(\mathcal{P}) = m_1(\mathcal{P}) + m_2(\mathcal{P}) + 2 \sum_{i=1}^{d^1-g^1} n_{w_i}(\mathcal{P}). \tag{71}$$

This has the following consequence:

Lemma 6.13 *Let u be an annoying α-degeneration in $(\Sigma^1, \boldsymbol{\alpha}^{1+})$, with domain $\mathcal{P}$. Then $\mu(\mathcal{P}) \geq 0$, with equality if and only if $\mathcal{P} = 0$.*

We now turn to studying how cylindrical triangles in $W_\Delta = \Sigma \times \Delta$ relate to those in $W_\Delta^1 = \Sigma^1 \times \Delta$ and $W_\Delta^2 = \Sigma^2 \times \Delta$ when we do a special connected sum. Pick almost complex structures J^1 and J^2 on W_Δ^1 and W_Δ^2 and disk neighborhoods D^1 of p^1 in Σ^1 and D^2 of p^2 in Σ^2. We assume that J^1 and J^2 are split near D^1 and D^2. For $T > 0$, we form the connected sum

$$\Sigma(T) = (\Sigma^1 - D^1) \cup ([-T-1, T+1] \times S^1) \cup (\Sigma^2 - D^2)$$

by inserting a long cylinder, using the identifications $\partial D^1 \cong \{-T\} \times S^1$ and $\partial D^2 = \{T\} \times S^1$. We construct an almost complex structure $J(T)$ on

$$W_\Delta(T) = \Sigma(T) \times \Delta$$

by extending J^1 and J^2 on the two sides, and using a split complex structure on the cylinder. The quantity T is called the *neck-length*.

By a *broken triangle* in a homology class ϕ, we mean the juxtaposition of a cylindrical triangle with some cylindrical flow lines and ordinary boundary degenerations such that the sum of all their homology classes (as a two-chain on the Heegaard surface) is ϕ.

We then have the following convergence result:

Proposition 6.14 *Suppose we have a special connected sum, with the notation above. Consider a homology class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$. Let $\phi^2 \in \pi_2(\boldsymbol{x}^2, \boldsymbol{y}^2, \boldsymbol{z}^2)$ be its restriction to Σ^2, and ϕ^{1+} the corresponding* *equivalence class of pairs on Σ^1. Suppose that the moduli space $\mathcal{M}(\phi)$ of cylindrical triangles is nonempty for a sequence of almost complex structures $J(T_i)$ with $T_i \to \infty$. Then the moduli space of broken holomorphic triangles in the class ϕ^2 is nonempty. Further, there exists a representative $(\phi^1, \mathcal{P})$ of the equivalence class ϕ^{1+} such that the moduli space of broken holomorphic triangles in the class ϕ^1 is nonempty and there exist some annoying α-degenerations with domains that sum to $\mathcal{P}$.*

Proof The proof is similar to that of the second part of Theorem 5.1 in [Ozsváth and Szabó 2008a], and is based on Gromov compactness; compare also [Lipshitz 2006, Sections 7, 8 and 10]. In the limit $T_i \to \infty$, the sequence of holomorphic triangles must have a subsequence converging to some holomorphic objects on Σ^1 and Σ^2. On Σ^2 the object is a broken triangle u^2. The only new twist is that, when $\pi_\Sigma \circ u^2$ maps a point of the boundary of the domain to the connected sum point $p^2 \in \Sigma^2$, on the other side (ie on Σ^1) an annoying α-degeneration must appear in the limit. In the end on Σ^1 we obtain a union of a broken triangle and some annoying α-degenerations. □

There is also a gluing result:

Proposition 6.15 *Consider a homology class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ on a triple Heegaard diagram obtained by special connected sum. Let $\phi^2 \in \pi_2(\boldsymbol{x}^2, \boldsymbol{y}^2, \boldsymbol{z}^2)$ be the restriction of ϕ to Σ^2, and ϕ^{1+} the equivalence class of pairs obtained by restricting ϕ to Σ^1. Suppose that ϕ^{1+} contains a representative of the form $(\phi^1, 0)$, with $\phi^1 \in \pi_2(\boldsymbol{x}^1, \boldsymbol{y}^1, \boldsymbol{z}^1)$. Further, suppose that $d_2 > g_2$, $\mu(\phi^1) = 0$, $\mu(\phi^2) = 2m$, and the domain $\mathcal{D}$ of ϕ has $m_1(\mathcal{D}) = m_2(\mathcal{D}) = m$, so that $\mu(\phi) = 0$; see Lemma 6.10. Consider the maps*

$$\rho^1 : \mathcal{M}(\phi^1) \to \mathrm{Sym}^m(\Delta) \quad \textit{and} \quad \rho^2 : \mathcal{M}(\phi^2) \to \mathrm{Sym}^m(\mathrm{int}(\Delta) \cup e_\alpha),$$

where

$$\rho^i(u) = \pi_\Delta((\pi_\Sigma \circ u^i)^{-1}(\{p^i\})).$$

If the fibered product

$$\mathcal{M}(\phi^1) \times_{\mathrm{Sym}^m(\Delta)} \mathcal{M}(\phi^2) = \{u^1 \times u^2 \in \mathcal{M}(\phi^1) \times \mathcal{M}(\phi^2) \mid \rho^1(u^1) = \rho^2(u^2)\}$$

is a smooth manifold, then this fibered product can be identified with the moduli space $\mathcal{M}(\phi)$ for sufficiently large neck-length.

Proof We claim that, for sufficiently large neck-length, if $u : S \to \Sigma \times \Delta$ is a holomorphic representative of ϕ, no point of ∂S is mapped to p^2 under $\pi_\Sigma \circ u$. Indeed, if such points existed, in the limit when $T \to \infty$ we would get a broken triangle in a class ψ^1 and one (or more) annoying α-degenerations on Σ^1, summing up to a class $\mathcal{P} \in H_2(\Sigma^1, \boldsymbol{\alpha}^{1+})$. We must have $\mu(\psi^1) \geq 0$ because of the existence of a holomorphic representative, $\mu(\mathcal{P}) > 0$ by Lemma 6.13, and $\mu(\psi^1) + \mu(\mathcal{P}) = 0$ because $(\psi^1, \mathcal{P}) \sim (\phi^1, 0)$. This is a contradiction, so our claim was true.

By the definition of the fibered product, if a holomorphic triangle $u^2 \in \mathcal{M}(\phi^2)$ is such that $\rho^2(u^2) \not\subset \mathrm{Sym}^m(\mathrm{int}(\Delta))$, that triangle cannot appear in the fibered product. With these observations in mind, the rest of the proof is completely analogous to that of the third part of Theorem 5.1 in [Ozsváth and Szabó 2008a]. Basically, the index conditions forbid the presence of flow lines and boundary degenerations as part of broken triangles in the relevant moduli space. Further, the hypothesis $d_2 > g_2$ is used to exclude the presence of sphere bubbles on the Σ^2 side. One can then use the gluing arguments from [Lipshitz 2006], applied to triangles. □

6.5 A degeneration argument

We now return to the setting of Section 6.1. We view the triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ as the special connected sum of the diagram $(\Sigma, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}})$ with the genus zero diagram $(\mathcal{S}, \alpha_1, \beta_1, \gamma_1)$ from Figure 11; see Example 6.8.

Lemma 6.16 *Let $\psi \in \pi_2(a, \theta, b)$ be a homology class of triangles in $\mathcal{S}$, with $a \in \{x, x'\}$ and $b \in \{y, y'\}$. Let m_1, m_2, m_3 and m_4 be the local multiplicities of the domain of ψ in the regions marked as such in Figure 11. Then*

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

Proof The claimed equality is true when ψ is the index zero triangle in $\pi_2(x, \theta, y)$. Any other class ψ is related to this by the juxtaposition of a linear combination of embedded index one α–β or α–γ bigons (ie homology classes of flow lines) and index two disks (ie homology classes of boundary degenerations on α, β or γ). It is straightforward to check the equality for these disks and bigons. □

Proof of Proposition 6.2 We seek to understand the moduli space of triangles $\mathcal{M}(\phi)$ for $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{\theta}, \boldsymbol{y})$ with $\mu(\phi) = 0$.

Suppose $\mathcal{M}(\phi) \neq \varnothing$ for any sufficiently large neck-length T. Let $\psi = \phi^2 \in \pi_2(a, \theta, b)$ be the restriction of ϕ to $\mathcal{S}$, where $a \in \{x, x'\}$ and $b \in \{y, y'\}$. Let also ϕ^{1+} be the equivalence class of pairs which is the restriction of ϕ to $(\Sigma, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}})$. Using Proposition 6.14, there must be a pair $(\overline{\phi}, \mathcal{P}) \in \phi^{1+}$ that admits holomorphic representatives. Hence, $\overline{\phi}$ and $\mathcal{P}$ are nonnegative domains, and $\mu(\phi^{1+}) \geq 0$.

On the other hand, by Lemmas 6.10 and 6.16, we have

$$0 = \mu(\phi) = \mu(\phi^{1+}) + \mu(\psi) - m_1 - m_2 = \mu(\phi^{1+}) + m_3 + m_4.$$

Since all the terms on the right-hand side are nonnegative, we deduce that $\mu(\phi^{1+}) = m_3 = m_4 = 0$. The fact that $\mu(\phi^{1+}) = 0$, together with Lemma 6.13, implies that there can be no annoying α-degenerations: $\mathcal{P} = 0$, $\mu(\overline{\phi}) = 0$ and $m_1 = m_2$. Denote by m the common value $m_1 = m_2$. We are now able to apply Proposition 6.15 to obtain an identification

$$\mathcal{M}(\phi) \cong \mathcal{M}(\overline{\phi}) \times_{\mathrm{Sym}^m(\Delta)} \mathcal{M}(\psi). \tag{72}$$

The fact that $m_1 = m_2$ and $m_3 = m_4 = 0$ implies that ψ must be a class in either $\pi_2(x, \theta, y)$ or $\pi_2(x', \theta, y')$. Without loss of generality, let us consider $\psi \in \pi_2(x, \theta, y)$. We have $\mu(\psi) = 2m$. From the proof of Proposition 6.15, we know that, for any $u \in \mathcal{M}(\psi)$,

$$\rho(u) = \pi_\Delta((\pi_\Sigma \circ u)^{-1}(\{p^2\}))$$

lies in $\mathrm{Sym}^m(\mathrm{int}(\Delta))$, that is, it does not contain any points on e_α. Given $\boldsymbol{p} \in \mathrm{Sym}^m(\Delta)$, set

$$\mathcal{M}(\psi, \boldsymbol{p}) = \{u \in \mathcal{M}(\psi) \mid \rho(u) = \boldsymbol{p}\}.$$

Define

$$M(\boldsymbol{p}) = \sum_{\substack{\psi \in \pi_2(x,\theta,y) \\ m_1(\psi)=m_2(\psi)=m \\ m_3(\psi)=m_4(\psi)=0}} \#\mathcal{M}(\psi, \boldsymbol{p}).$$

A Gromov compactness argument shows that $M(\boldsymbol{p})$ is independent of $\boldsymbol{p}$, modulo 2; compare [Ozsváth and Szabó 2008a, Lemma 6.4]. By taking the limit as $\boldsymbol{p}$ consists of m distinct points, all approaching the edge e_β of Δ with spacing at least T between them, with $T \to \infty$, we obtain that the respective contributions $M(\boldsymbol{p})$ are splicings of the index zero triangle class $\psi_0 \in \pi_2(x, \theta, y)$ and m β-boundary degenerations of index two; compare [loc. cit.]. There is a unique possible class of β-boundary degenerations of index two with $m_3 = m_4 = 0$, namely the exterior of the curve β_1 from Figure 11. For this class the count of pseudoholomorphic representatives (modulo the two-dimensional group of automorphisms) is 1 (mod 2); see [ibid., Theorem 5.5]. Moreover, the index zero triangle class ψ_0 has a unique pseudoholomorphic representative. It follows that

$$M(\boldsymbol{p}) \equiv 1 \pmod 2.$$

A similar equality holds for the sum of contributions from classes $\psi \in \pi_2(x', \theta, y')$ with $m_1(\psi) = m_2(\psi) = m$ and $m_3(\psi) = m_4(\psi) = 0$. Combining these observations with (72), we deduce that

$$\sum_{\substack{\phi \in \pi_2(\boldsymbol{x},\boldsymbol{\theta},\boldsymbol{y}) \\ \mu(\phi)=0}} \mathcal{M}(\phi) \equiv \sum_{\substack{\bar{\phi} \in \pi_2(\bar{\boldsymbol{x}},\bar{\boldsymbol{\theta}},\bar{\boldsymbol{y}}) \\ \mu(\bar{\phi})=0}} \mathcal{M}(\bar{\phi}) \pmod 2$$

for sufficiently large neck-length. This implies that the triangle maps F and G are the same. □

6.6 Higher polygons

We now turn to a generalization of Proposition 6.2. Suppose we have $l \geq 2$ collections of attaching curves $\boldsymbol{\alpha}, \boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, \dots, \boldsymbol{\beta}^{(l)}$ on a multipointed Heegaard surface $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$ such that each diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}^{(i)}, \boldsymbol{w}, \boldsymbol{z})$ is the quasistabilization of a diagram $(\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}^{(i)}, \bar{\boldsymbol{w}}, \boldsymbol{z})$, obtained by adding curves α_1 and $\beta_1^{(i)}$ and the basepoint w_1. We also assume that, for every $i \neq j$, the curves $\beta_1^{(i)}$ and $\boldsymbol{\beta}_1^{(j)}$ differ by a small Hamiltonian isotopy, and intersect each other in two points. See Figure 12.

Let $\{x, x'\} = \alpha_1 \cap \beta_1^{(1)}$ and $\{y, y'\} = \alpha_1 \cap \beta_1^{(l)}$, with x and y to the left of x' and y', as in Figure 12. Let also $\theta^{(i)} \in \beta_1^{(i)} \cap \beta_1^{(i+1)}$ be the upper intersection point between the two curves.

There is a unique index zero $(l+1)$-gon class ψ_0 on the sphere $\mathcal{S}$ with vertices at $x, \theta^{(1)}, \dots, \theta^{(l-1)}, y$ such that ψ_0 has a positive domain; see Figure 12. The moduli space of holomorphic representatives for ψ_0 is $(l-2)$-dimensional, corresponding to all possible lengths of the cuts at $\theta^{(1)}, \dots, \theta^{(l-2)}$.

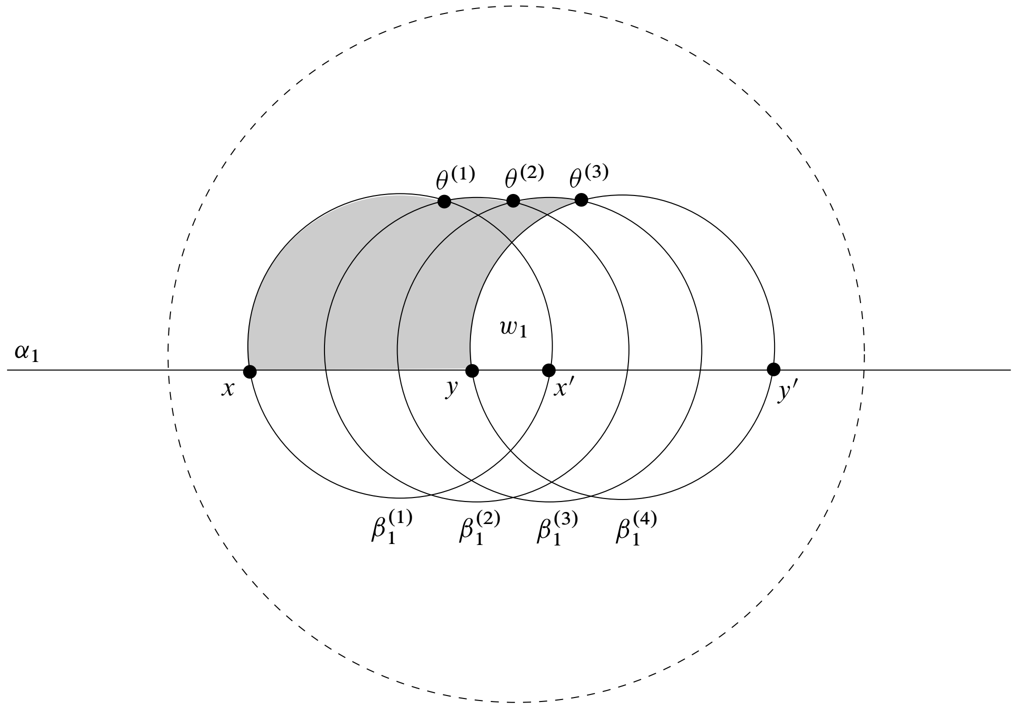

Figure 12: Several curve collections. This is the analogue of Figure 9 for higher polygons. We show here the case $l = 4$. The shaded domain is a pentagon of index zero.

Lemma 6.17 *For a fixed generic conformal structure on the domain* (*which is a disk with* $l+1$ *marked points on its boundary*), *the number of pseudoholomorphic representatives of* ψ_0 *is* 1 (mod 2).

Proof We use induction on l. The case $l = 3$ was treated in [Ozsváth and Szabó 2005, proof of Theorem 4.7]. For $l \geq 4$, let J_0 be a generic conformal structure on the domain. Take a generic, smooth path of conformal structures $(J_t)_{t\in[0,1)}$, which starts at J_0 and limits (as $t \to 1$) to a degenerate conformal structure, corresponding to taking the length of one of the cuts starting at $\theta^{(l-2)}$ to infinity. Thus, in the limit $t \to 1$ the domain degenerates into the union of an l-sided polygon and a triangle. Let $\mathcal{M}_t = \mathcal{M}_t(\psi_0)$ be the moduli space of pseudoholomorphic representatives of ψ_0 with the conformal structure J_t on the domain. By Gromov compactness and generic transversality, the union

$$\mathcal{M} = \bigcup_{t\in[0,1]} \mathcal{M}_t$$

is a one-dimensional compact manifold with boundary $\mathcal{M}_0 \cup \mathcal{M}_1$. By the inductive hypothesis, the cardinality of $\mathcal{M}_1$ is odd; hence, the same must be true for $\mathcal{M}_0$. $\square$

For $i = 1, \dots, l-1$, pick $\bar{\boldsymbol{\theta}}^{(i)} \in \mathbb{T}_{\bar{\beta}^{(i)}} \cap \mathbb{T}_{\bar{\beta}^{(i+1)}}$ and set

$$\boldsymbol{\theta}^{(i)} = \bar{\boldsymbol{\theta}}^{(i)} \cup \{\theta^{(i)}\} \in \mathbb{T}_{\beta^{(i)}} \cap \mathbb{T}_{\beta^{(i+1)}}.$$

Just as in Section 6, there is a map

$$F\colon \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta^{(1)}}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta^{(l)}}, \boldsymbol{s}), \quad F(\boldsymbol{x}) = f(\boldsymbol{x} \otimes \boldsymbol{\theta}^{(1)} \otimes \cdots \otimes \boldsymbol{\theta}^{(l-1)}),$$

this time given by counting pseudoholomorphic $(l+1)$-gons of index $2-l$; see Section 3.5.

For any $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta^{(1)}}$ and $\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta^{(l)}}$, we can eliminate their intersections with α_1 to obtain generators $\bar{\boldsymbol{x}} \in \mathbb{T}_{\overline{\alpha}} \cap \mathbb{T}_{\overline{\beta}^{(1)}}$ and $\boldsymbol{y} \in \mathbb{T}_{\overline{\alpha}} \cap \mathbb{T}_{\overline{\beta}^{(l)}}$. We can define a map $\overline{F}$ by counting $(l+1)$-gons in the destabilized diagram with $l-1$ fixed vertices at $\bar{\boldsymbol{\theta}}^{(1)}, \dots, \bar{\boldsymbol{\theta}}^{(l-1)}$. If $\overline{F}(\bar{\boldsymbol{x}}) = \sum_{\bar{\boldsymbol{y}} \in \mathbb{T}_{\overline{\alpha}} \cap \mathbb{T}_{\overline{\beta}^{(l)}}} n_{\bar{\boldsymbol{x}},\bar{\boldsymbol{y}}} \bar{\boldsymbol{y}}$, we set

$$G(\boldsymbol{x}) = \sum_{\bar{\boldsymbol{y}} \in \mathbb{T}_{\overline{\alpha}} \cap \mathbb{T}_{\overline{\beta}^{(l)}}} n_{\bar{\boldsymbol{x}},\bar{\boldsymbol{y}}} (\bar{\boldsymbol{y}} \times (\boldsymbol{x} \cap \alpha_1)).$$

Proposition 6.18 *For $\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(l)}, \boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(l-1)}$ as above and suitable almost complex structures on the symmetric products, the maps F and G coincide.*

Proof The arguments are completely analogous to the ones for $l = 2$. We insert a long neck along the dashed curve from Figure 12. We thus view Σ as the special connected sum of the sphere $\mathcal{S}$ and the quasidestabilized diagram, except now each has $l+1$ collections of attaching curves. Let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(l-1)}, \boldsymbol{y})$ be a homology class of $(l+1)$-gons, of index $2-l$, that admits pseudoholomorphic representatives. In the limit when the neck-length $T \to \infty$, the class ϕ splits into homology classes of broken $(l+1)$-gons ψ on the sphere $\mathcal{S}$, $\overline{\phi}$ on the quasidestabilized diagram, and a class $\mathcal{P}$ of annoying α-degenerations. Lemma 6.10, equation (71) and Lemma 6.16 still hold true, and therefore

$$2-l = \mu(\phi) = \mu(\overline{\phi}) + \mu(\mathcal{P}) + m_3 + m_4 \geq (2-l) + 0 + 0 + 0. \tag{73}$$

Hence, $\mathcal{P} = 0$ (so there are no annoying α-degenerations), $m_3 = m_4 = 0$, and we end up with a fibered product description of $\mathcal{M}(\phi)$ analogous to (72). On the $\mathcal{S}$ side, we can use a limiting process to ensure that the holomorphic representatives of ψ are splicings of an index zero $(l+1)$-gon and several $\beta^{(1)}$-boundary degenerations. One possibility for the index zero $(l+1)$-gon is that it lies in the class ψ_0 analyzed in Lemma 6.17. In the fibered product description the conformal structure of the domain of a pseudoholomorphic $(l+1)$-gon in ψ_0 is specified by the $(l+1)$-gon on the quasidestabilized diagram. By Lemma 6.17, the number of the holomorphic representatives of ψ_0 with this constraint is 1 (mod 2). A similar discussion applies to the index zero $(l+1)$-gon class with vertices at $x', \theta^{(1)}, \dots, \theta^{(l-1)}, y'$. The fibered product description then implies the identification of the two maps F and G. □

The reader may wonder what happens in the case $l = 1$, that is, how are flow lines in a diagram $\overline{\mathcal{H}}$ related to flow lines in its quasistabilization $\mathcal{H}$. This question is more difficult, because if we try to degenerate along the special connected sum neck, we can no longer avoid the presence of annoying α-degenerations. Indeed, in (73) the inequality $\mu(\overline{\phi}) \geq 2-l = 1$ does not hold true, due to the existence of index zero flow lines on $\overline{\mathcal{H}}$ (namely, trivial ones). This leaves open the possibility that $\mu(\mathcal{P}) = 1$.

Nevertheless, we make the following:

Conjecture 6.19 *Let $\overline{\mathcal{H}} = (\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\boldsymbol{w}}, z)$ be a Heegaard diagram, and $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, z)$ be its quasistabilization, as in Section 6.1. Suppose w_2 is the second basepoint (apart from w_1) in the component of $\Sigma - \bar{\boldsymbol{\alpha}}$ that contains the curve α_1. Suppose the variables corresponding to w_1 and w_2 are U_1 and U_2. Then, for suitable almost complex structures, there is an identification between the Floer complex $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ and the mapping cone complex*

$$\boldsymbol{A}^-(\overline{\mathcal{H}}, \boldsymbol{s})[\![U_1]\!] \xrightarrow{U_1 - U_2} \boldsymbol{A}^-(\overline{\mathcal{H}}, \boldsymbol{s})[\![U_1]\!].$$

Note that this is a generalization of [Ozsváth and Szabó 2008a, Proposition 6.5], which dealt with ordinary index zero/three stabilizations. The difficulty in proving Conjecture 6.19 is the lack of available transversality and gluing results for annoying curves. In Section 15.4, we will give a proof of the conjecture for a particular class of Heegaard diagrams, using rather ad hoc arguments.

6.7 Stabilizations

We now specialize to the case of ordinary (free) index zero/three stabilizations.

Let $(\Sigma, \boldsymbol{\alpha}, \{\boldsymbol{\beta}^{(i)}\}_{i=1}^{l}, \boldsymbol{w}, z)$ be as in the previous section, and

$$F\colon \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta^{(1)}}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta^{(l)}}, \boldsymbol{s}), \quad F(\boldsymbol{x}) = f(\boldsymbol{x} \otimes \boldsymbol{\theta}^{(1)} \otimes \cdots \otimes \boldsymbol{\theta}^{(l-1)}),$$

the corresponding polygon map. Assume further that the curves $\beta_1^{(i)}$ is a small Hamiltonian translate of α_1, intersecting it in two points. Thus, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}^{(i)}, \boldsymbol{w}, z)$ is obtained from $(\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}^{(i)}, \bar{\boldsymbol{w}}, z)$ by an index zero/three stabilization, and

$$\bar{F}\colon \boldsymbol{A}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\beta}^{(1)}}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\beta}^{(l)}}, \boldsymbol{s}), \quad \bar{F}(\bar{\boldsymbol{x}}) = f(\bar{\boldsymbol{x}} \otimes \bar{\boldsymbol{\theta}}^{(1)} \otimes \cdots \otimes \bar{\boldsymbol{\theta}}^{(l-1)}).$$

We keep the notation from Conjecture 6.19, with basepoints w_1, w_2 and variables U_1, U_2. Associate one U variable to each w basepoint, and let $\mathcal{R}$ be the power series ring generated over $\mathbb{F}$ by all U variables except U_1.

From the proofs of Theorem 3.16 and Lemma 3.14, we know that there are destabilization maps

$$\mathrm{Destab}^i\colon \boldsymbol{A}^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}^{(i)}, \boldsymbol{s}) \to \boldsymbol{A}^-(\Sigma, \bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}^{(i)}, \boldsymbol{s}),$$

defined by

$$U_1^n(\bar{\boldsymbol{x}} \times x') \mapsto 0, \quad U_1^n(\bar{\boldsymbol{x}} \times x) \mapsto U_2^n \bar{\boldsymbol{x}},$$

and these maps are equivalences over the ring $\mathcal{R}$.

Proposition 6.20 *The equivalences induced by destabilization commute with the polygon maps; that is,*

$$\mathrm{Destab}^l \circ F = \bar{F} \circ \mathrm{Destab}^1.$$

Proof Use Proposition 6.18. ◻

Versions of Proposition 6.20 are true for ordinary Heegaard Floer complexes, as well. In that case, we can consider an admissible multidiagram $(\Sigma, \boldsymbol{\alpha}, \{\boldsymbol{\beta}^{(i)}\}_{i=1}^{l}, \boldsymbol{w})$, where again the curves $\beta_1^{(i)}$ are small isotopic translates of α_1, so that we can form the index zero/three destabilizations, and corresponding maps

$$F_o\colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta^{(1)}}, \boldsymbol{w}) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta^{(l)}}, \boldsymbol{w}), \quad \overline{F}_o\colon \boldsymbol{CF}^-(\mathbb{T}_{\overline{\alpha}}, \mathbb{T}_{\overline{\beta}^{(1)}}, \overline{\boldsymbol{w}}) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\overline{\beta}^{(l)}}, \overline{\boldsymbol{w}}).$$

(Note that the Heegaard diagrams appearing here can represent arbitrary three-manifolds.)

We have analogous destabilization maps Destab_o.

Proposition 6.21 *In the case of ordinary Heegaard Floer complexes, the equivalences induced by destabilization commute with the polygon maps; that is,*

$$\mathrm{Destab}_o^l \circ F_o = \overline{F}_o \circ \mathrm{Destab}_o^1.$$

Proof This is analogous to the proof of Proposition 6.20. □

7 Index zero/three link stabilizations

In this section we study one of the Heegaard moves that appeared in Section 3.3, namely index zero/three link stabilization. The behavior of holomorphic disks under this move was discussed in [Manolescu et al. 2009a; Zemke 2015]. We will review those results, and then (in Section 7.4) we extend them to the case of more general holomorphic polygons. Furthermore, in Section 7.5, we will study what happens to holomorphic polygons under a type of strong equivalence that relates different kinds of stabilizations.

The analytical input for this section comes from [Zemke 2015].

7.1 Algebraic preliminaries

In Section 3.4 we established Lemma 3.14, which said that a free complex C over $R[\![U_1]\!]$ is chain homotopy equivalent to the mapping cone

$$\mathrm{Cone}(C_-[\![U_2]\!] \xrightarrow{U_1-U_2} C_+[\![U_2]\!]). \tag{74}$$

That cone appeared naturally in the context of (both free and link) index zero/three stabilizations; see the proof of Theorem 3.16.

To better understand link index zero/three stabilizations, we will also need the following variant of Lemma 3.14:

Lemma 7.1 *Let R be an $\mathbb{F}$-algebra, and let C be a free complex over $R[\![U_1]\!]$. Let $C_\pm^{U_1\to U_2}$ be two copies of the free complex over $R[\![U_2]\!]$ obtained from C by replacing the variable U_1 with U_2. Then the complexes C and*

$$C' = \mathrm{Cone}(C_+^{U_1\to U_2}[\![U_1]\!] \xrightarrow{U_1-U_2} C_-^{U_1\to U_2}[\![U_1]\!])$$

are chain homotopy equivalent over $R[\![U_1]\!]$.

Proof Let us fix a set of free generators $\boldsymbol{S}$ for C. We write the differential on C as

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \boldsymbol{S}} c(\boldsymbol{x}, \boldsymbol{y}) U_1^{n(\boldsymbol{x}, \boldsymbol{y})} \boldsymbol{y},$$

where $c(\boldsymbol{x}, \boldsymbol{y}) \in \{0, 1\}$ and $n(\boldsymbol{x}, \boldsymbol{y})$ are nonnegative integers. Further, given a generator $\boldsymbol{x} \in \boldsymbol{S}$, we denote by $\boldsymbol{x}_+$ and $\boldsymbol{x}_-$ the corresponding generators in $C_+^{U_1 \to U_2}$ and $C_-^{U_1 \to U_2}$.

We construct chain maps between C and C' in both directions such that they are module maps over $R[\![U_1]\!]$ and they are homotopy inverses.

We let $\rho\colon C' \to C$ be given by

$$\rho(U_2^m \boldsymbol{x}_-) = U_1^m \boldsymbol{x}, \quad \rho(U_2^m \boldsymbol{x}_+) = 0.$$

In the opposite direction, we define $\iota\colon C \to C'$ by

$$\iota(\boldsymbol{x}) = \boldsymbol{x}_- + \sum_{\boldsymbol{y} \in \boldsymbol{S}} c(\boldsymbol{x}, \boldsymbol{y}) \frac{U_1^{n(\boldsymbol{x}, \boldsymbol{y})} - U_2^{n(\boldsymbol{x}, \boldsymbol{y})}}{U_1 - U_2} \boldsymbol{y}_+.$$

One can check that ρ and ι commute with the differentials, and $\rho \circ \iota$ is the identity. Moreover, if we define

$$H\colon C' \to C', \quad H(U_2^n \boldsymbol{x}_+) = 0, \quad H(U_2^n \boldsymbol{x}_-) = \frac{U_1^n - U_2^n}{U_1 - U_2} \boldsymbol{x}_+,$$

then we find that $\iota \circ \rho$ is chain homotopic to the identity, via the homotopy H. □

7.2 Holomorphic disks

As defined in Section 3.3, an index zero/three link stabilization is a local move on a Heegaard diagram, in the neighborhood of an existing basepoint of type z. The move introduces a new alpha curve, a new beta curve, and a basepoint of each type; see Figure 1. For convenience, we include Figure 1 again here, as Figure 13, with the curves α_n and β_n relabeled as α_1 and β_1, and the basepoints z, w' and z' relabeled as z_1, w_2 and z_2. (Not shown is, for example, the basepoint w_1, which is separated from z_1 by a number of beta curves.)

We will denote the initial diagram by $\overline{\mathcal{H}} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{w}}, \overline{\boldsymbol{z}})$ and the stabilized diagram by $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$. We let $L_1 \subseteq L$ be the link component on which the stabilization is performed.

We are interested in the relation between the holomorphic disks in $\mathcal{H}$ and those in $\overline{\mathcal{H}}$. If we ignore the z basepoints, our picture is just that of a free index zero/three stabilization, and an analysis of the holomorphic disks was done in [Ozsváth and Szabó 2008a, Section 6]; see also [Manolescu et al. 2009a, Section 2]. The analysis is based on viewing $\mathcal{H}$ as the connected sum of $\overline{\mathcal{H}}$ (with z_1 deleted) and a sphere $\mathcal{S}$ containing the new curves (ie containing what is shown in Figure 13), and then degenerating the Heegaard surface along the curve c in Figure 13, by taking the length of the connected sum neck to infinity. Further, the connected sum point p is then taken to be close to one of the circles α_1 or β_1 in Figure 13.

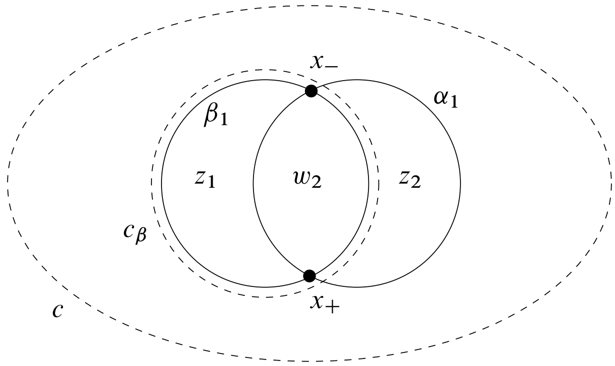

Figure 13: An index zero/three link stabilization.

The analysis in [Ozsváth and Szabó 2008a; Manolescu et al. 2009a] did not distinguish between the disks that cross z_1 and those that cross z_2. A more complete characterization of the holomorphic disks was given in [Zemke 2015, Lemmas 14.3 and 14.4]. This is based on a different degeneration. We still stretch the neck along the dashed curve c in Figure 13, but then, instead of taking the connected sum point close to a curve, we stretch along the dashed curve c_β. Note that stretching along c_β is very similar to the degeneration used for quasistabilizations in Section 6.

We state the results from [loc. cit.] here, with slightly different conventions: z_1 and z_2 play the roles of w' and w in [ibid., Lemma 14.3], and α_1 and β_1 are switched.

Proposition 7.2 [Zemke 2015] *For generic almost complex structures associated to sufficiently large neck stretching along c and c_β, with regard to the counts of holomorphic disks (modulo 2) in $\mathcal{H}$:*

- *The only disks of Maslov index one in a class $\phi \in \pi_2(\boldsymbol{x} \times x_-, \boldsymbol{y} \times x_-)$ from $\mathcal{H}$ appear when $n_{z_1}(\phi) = n_{w_2}(\phi) = 0$, in which case they are in one-to-one correspondence with the holomorphic disks of Maslov index one in the corresponding class $\overline{\phi} \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ in $\overline{\mathcal{H}}$. Here, ϕ is obtained from $\overline{\phi}$ by taking connected sum (along c) with some copies of the complement of the disk bounded by β_1 in Figure 13, so that $n_{z_2}(\phi) = n_{z_1}(\overline{\phi})$.*
- *The same description applies to the holomorphic disks of Maslov index one in classes $\phi \in \pi_2(\boldsymbol{x} \times x_+, \boldsymbol{y} \times x_+)$.*
- *The only holomorphic disks of Maslov index one in a class $\phi \in \pi_2(\boldsymbol{x} \times x_-, \boldsymbol{y} \times x_+)$ from $\mathcal{H}$ appear when $\boldsymbol{x} = \boldsymbol{y}$ and the domain of ϕ is either the bigon in Figure 13 containing w_1 or the bigon containing w_2; moreover, in each of these cases, there is only one such disk.*
- *The only holomorphic disks of Maslov index one in a class $\phi \in \pi_2(\boldsymbol{x} \times x_+, \boldsymbol{y} \times x_-)$ from $\mathcal{H}$ appear when $\boldsymbol{x} = \boldsymbol{y}$ and $n_{z_1}(\phi) = n_{z_2}(\phi) = 0$. Moreover, for each $\boldsymbol{x}$, there are exactly two classes ϕ such that the total count of holomorphic disks in ϕ is $1 \pmod 2$: one class is the one whose domain is the bigon in Figure 13 containing w_2, and the other is a class ϕ that has $n_{w_1}(\phi) = 1$ and goes over no other basepoints.*

Proposition 7.2 allows us to relate the various Heegaard Floer complexes associated to $\mathcal{H}$ and $\overline{\mathcal{H}}$. For example, if we ignore the z basepoints, we are in the setting of a free index zero/three stabilization. Consider the complex $C = \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\beta}}, \bar{\boldsymbol{w}})$. Then the corresponding complex $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$ in the stabilized diagram is identified with

$$C_-[\![U_2]\!] \xrightarrow{U_2-U_1} C_+[\![U_2]\!], \tag{75}$$

where we denoted by $C_\pm$ the copies of C that are obtained by including the basepoint $x_\pm$ in $\mathcal{S}$. This cone appeared in (74), and explains why Lemma 3.14 is relevant to our situation. In particular, Lemma 3.14 implies the following result (which we have already seen as part of the proof of Theorem 3.16):

Corollary 7.3 *In the setting of Figure 13, the destabilization map*

$$\rho\colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}) \to \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\beta}}, \bar{\boldsymbol{w}})$$

given by

$$\rho(U_2^m(\boldsymbol{x} \times x_+)) = U_1^m \boldsymbol{x}, \quad \rho(U_2^m(\boldsymbol{x} \times x_-)) = 0$$

is a chain homotopy equivalence over $\overline{\mathcal{R}}$ (the ring with one U variable for each w basepoint in the diagram $\overline{\mathcal{H}}$).

Furthermore, the equivalence ρ is filtered with respect to the Alexander filtrations, and therefore gives rise to equivalences

$$\rho\colon \boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s}) \to \boldsymbol{A}^-(\overline{\mathcal{H}}, \boldsymbol{s}).$$

Note that $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ does not admit a mapping cone description similar to (75). However, such descriptions do apply to, for example, generalized link Floer complexes that involve only the w basepoints on the component L_1, but may involve both w's and z's on the other components.

Also relevant to link stabilizations is Lemma 7.1. Indeed, let us ignore the w basepoints and look at the complexes constructed with z basepoints. We relabel the z basepoints on L_1 as w's, so that our notation is consistent with that from Section 3.1. Then Figure 13 becomes Figure 14. The move in Figure 14 consists in introducing two new curves and a new basepoint w_2 in the neighborhood of w_1. This is a variant of a free index zero/three stabilization: it differs from that by a strong equivalence (see Section 7.5 below), and in fact it could replace the ordinary free index zero/three stabilizations in the list of Heegaard moves in Section 3.3.

If we now write $C = \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}}, \mathbb{T}_{\bar{\beta}}, \bar{\boldsymbol{w}})$, then Proposition 7.2 implies that the corresponding complex $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$ in $\mathcal{H}$ is

$$C_+^{U_1\to U_2}[\![U_1]\!] \xrightarrow{U_1-U_2} C_-^{U_1\to U_2}[\![U_1]\!]. \tag{76}$$

Lemma 7.1 has the following:

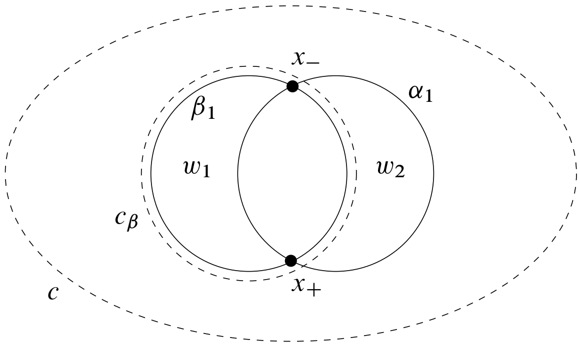

Figure 14: A variant of the free index zero/three stabilization.

Corollary 7.4 *For the move shown in Figure 14, the destabilization map*

(77) $\rho\colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}) \to \boldsymbol{CF}^-(\mathbb{T}_{\bar\alpha}, \mathbb{T}_{\bar\beta}, \bar{\boldsymbol{w}}), \qquad \rho(U_2^m(\boldsymbol{x} \times x_-)) = U_1^m \boldsymbol{x}, \quad \rho(U_2^m(\boldsymbol{x} \times x_+)) = 0,$

is a chain homotopy equivalence over $\overline{\mathcal{R}}$.

Remark 7.5 The result of Corollary 7.4 also holds if instead of $\boldsymbol{CF}^-$ we work with generalized link Floer complexes that involve only the z basepoints on the component L_1 (now relabeled as w's), but may involve both w's and z's on the other components.

7.3 Changes in the almost complex structures

In the previous subsection we worked with a generic family (call it J_β) of almost complex structures associated to neck stretching along the curves c and c_β. We could just as well stretch along c and a curve c_α around α_1; let us denote the corresponding family of almost complex structures by J_α.

With J_α, the analogue of Proposition 7.2 holds, with the difference that the disks in a class $\phi \in \pi_2(\boldsymbol{x} \times x_-, \boldsymbol{y} \times x_-)$ or $\pi_2(\boldsymbol{x} \times x_+, \boldsymbol{y} \times x_+)$ have $n_{z_2}(\phi) = 0$ instead of $n_{z_1}(\phi) = 0$; and the classes ϕ are obtained from $\bar\phi$ by taking connected sum with some copies of the complement of the disk bounded by α_1. Note that the description of $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$ as the cone (75) is still true, and so is Corollary 7.3.

On the other hand, once we ignore the w basepoints and relabel z's as w's, the description we have of $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$ from (76) changes. In the situation depicted on the right-hand side of Figure 15, with J_α, the complex $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$ is simply the cone

$$C_+[\![U_2]\!] \xrightarrow{U_2 - U_1} C_-[\![U_2]\!]. \tag{78}$$

Corollary 7.4 still holds, with the ρ map defined by the same formula (77).

One can interpolate between J_β and J_α by a family of almost complex structures, all with sufficiently large connected sum neck along c. This induces a map on Floer complexes

$$\Phi_{J_\beta \to J_\alpha}\colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}, J_\beta) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}, J_\alpha). \tag{79}$$

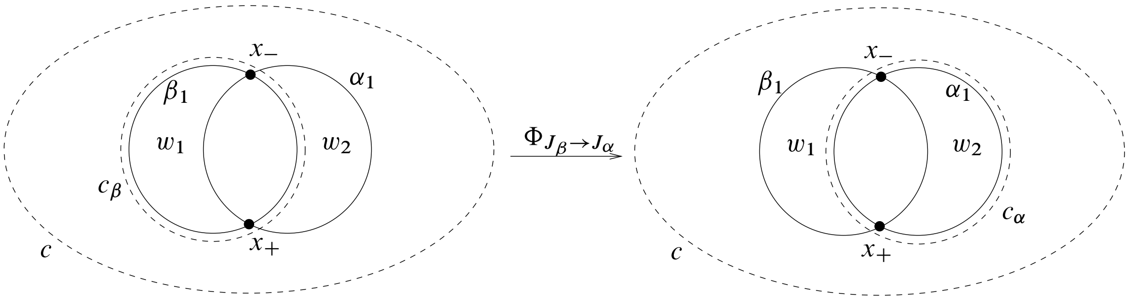

Figure 15: A change of almost complex structures from Figure 14.

See Figure 15. The change of almost complex structure map was studied in [Zemke 2015, Lemma 14.5]. We state the result below with our conventions.

Here, and later in the paper, we will write maps between mapping cones as 2×2 matrices; each cone will be viewed (as a vector space) as the direct sum of its domain and target, in this order.

Proposition 7.6 [Zemke 2015] *For sufficiently stretched almost complex structures, the map* (79) *can be written as*

$$\Phi_{J_\beta \to J_\alpha} = \begin{pmatrix} \mathrm{id} & * \\ 0 & \mathrm{id} \end{pmatrix}$$

or, in more expanded form,

$$\begin{tikzcd} C_+^{U_1 \to U_2}[\![U_1]\!] \arrow[r, "\mathrm{id}"] \arrow[d, "U_1 - U_2"'] & C_+[\![U_2]\!] \arrow[d, "U_1 - U_2"] \\ C_-^{U_1 \to U_2}[\![U_1]\!] \arrow[r, "\mathrm{id}"] \arrow[ur, "*"] & C_-[\![U_2]\!] \end{tikzcd}$$

Remark 7.7 The upper right entry $*$ in $\Phi_{J_\beta \to J_\alpha}$ can be computed explicitly; see Remark 14.6 of [Zemke 2015]. However, we do not need to know it for our purposes.

Corollary 7.8 *The map* (79) *is related to the projections ρ of the form* (77) *by*

$$\rho \circ \Phi_{J_\beta \to J_\alpha} = \rho.$$

Proof This follows immediately from Proposition 7.6. Indeed, recall that the maps ρ take the domains of the cones to zero, and act on the targets by taking both U_1 and U_2 to U_1. □

We also need to consider the change of almost complex structures map for a free index zero/three stabilization done in the neighborhood of a basepoint w_1; see Figure 16. This will be useful to us in Section 7.5 below. To be in agreement with the notation there, we will denote the unstabilized diagram by $(\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\gamma}}, \overline{\boldsymbol{w}})$, and the stabilized one by $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w})$, with $\boldsymbol{\alpha} = \overline{\boldsymbol{\alpha}} \cup \{\alpha_1\}$, $\boldsymbol{\gamma} = \overline{\boldsymbol{\gamma}} \cup \{\gamma_1\}$ and $\boldsymbol{w} = \overline{\boldsymbol{w}} \cup \{w_2\}$. We consider almost complex structures J_α stretched along c and c_α, and J_γ stretched along c and c_γ. We

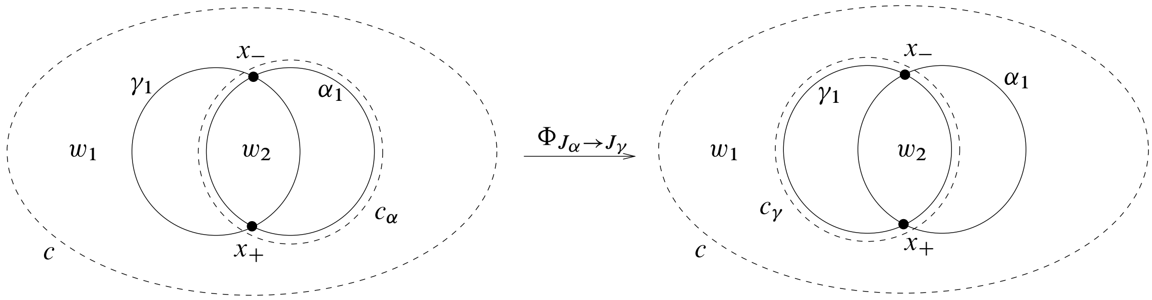

Figure 16: A change of almost complex structures for free index zero/three stabilizations.

then interpolate between them using a family of almost complex structures sufficiently stretched along c. This induces a map

$$\Phi_{J_\alpha \to J_\gamma} \colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w}, J_\alpha) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w}, J_\gamma). \tag{80}$$

Note that, if we write $C = \boldsymbol{CF}^-(\mathbb{T}_{\bar\alpha}, \mathbb{T}_{\bar\gamma}, \bar{\boldsymbol{w}})$, then both the domain and the target of $\Phi_{J_\alpha \to J_\gamma}$ can be identified with the cone

$$C_-[\![U_2]\!] \xrightarrow{U_2 - U_1} C_+[\![U_2]\!]. \tag{81}$$

The following result comes from [Zemke 2015, Proposition 14.22]:

Proposition 7.9 [Zemke 2015] *For sufficiently stretched almost complex structures, the map* (80) *can be written as*

$$\Phi_{J_\alpha \to J_\gamma} = \begin{pmatrix} \mathrm{id} & 0 \\ * & \mathrm{id} \end{pmatrix}$$

or, in more expanded form,

$$\begin{tikzcd} C_-[\![U_2]\!] \arrow[r, "\mathrm{id}"] \arrow[d, "U_1 - U_2"'] \arrow[dr, "*"] & C_-[\![U_2]\!] \arrow[d, "U_1 - U_2"] \\ C_+[\![U_2]\!] \arrow[r, "\mathrm{id}"] & C_+[\![U_2]\!] \end{tikzcd}$$

Once again, the lower right term $*$ is computed explicitly in [Zemke 2015, Proposition 14.22], but we do not need to know that formula here.

We are interested in the relation of the map (80) with the projections ρ.

Corollary 7.10 *There is a chain homotopy*

$$\rho \circ \Phi_{J_\alpha \to J_\gamma} \simeq \rho.$$

Proof Note that, since now the lower right term is nonzero, we do not have $\rho \circ \Phi_{J_\alpha \to J_\gamma} = \rho$ on the nose, as in Corollary 7.8. On the other hand, we can look at the homotopy inverses to ρ. From the proof of Lemma 3.14, such homotopy inverses are given by

$$\iota \colon \boldsymbol{CF}^-(\mathbb{T}_{\bar\alpha}, \mathbb{T}_{\bar\gamma}, \bar{\boldsymbol{w}}) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w}), \quad \iota(\boldsymbol{x}) = \boldsymbol{x} \times x_+.$$

Proposition 7.9 implies that $\iota = \Phi_{J_\alpha \to J_\gamma} \circ \iota$, on the nose. Since $\iota \circ \rho \simeq \mathrm{id}$ and $\rho \circ \iota = \mathrm{id}$, we have

$$\rho \circ \Phi_{J_\alpha \to J_\gamma} \simeq \rho \circ \Phi_{J_\alpha \to J_\gamma} \circ \iota \circ \rho = \rho \circ \iota \circ \rho = \rho. \qquad \square$$

7.4 Holomorphic polygons

Our next goal is to study the behavior of holomorphic polygons under an index zero/three link stabilization. The discussion will be modeled on Section 6, where we did a similar study for quasistabilizations.

Consider a Heegaard multidiagram

$$(\overline{\Sigma}, \overline{\boldsymbol{\alpha}}^{(1)}, \dots, \overline{\boldsymbol{\alpha}}^{(k)}, \overline{\boldsymbol{\beta}}^{(1)}, \dots, \overline{\boldsymbol{\beta}}^{(l)}, \overline{\boldsymbol{w}}, \bar{\boldsymbol{z}}),$$

and let

$$(\Sigma, \boldsymbol{\alpha}^{(1)}, \dots, \boldsymbol{\alpha}^{(k)}, \boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(l)}, \boldsymbol{w}, \boldsymbol{z})$$

be its index zero/three link stabilization, in the sense that each pair $(\boldsymbol{\alpha}^{(i)}, \boldsymbol{\beta}^{(j)})$ is related to $(\overline{\boldsymbol{\alpha}}^{(i)}, \overline{\boldsymbol{\beta}}^{(j)})$ by an index zero/three link stabilization, introducing two new basepoints w_2 and z_2 near z_1, and two new curves $\alpha_1^{(i)}$ and $\beta_1^{(j)}$. Furthermore, we want the new curves $\alpha_1^{(i)}$ to be Hamiltonian translates of one another, intersecting at two points each; and the same for the new curves $\beta_1^{(j)}$. See Figure 17 for an example, with $k = 2$ and $l = 3$.

Mark the top-degree intersection points $\theta_\alpha^{(i)} \in \alpha_1^{(i)} \cap \alpha_1^{(i+1)}$ for $i = 1, \dots, k-1$, and $\theta_\beta^{(j)} \in \beta_1^{(j)} \cap \beta_1^{(j+1)}$ for $j = 1, \dots, l-1$. Pick $\boldsymbol{\theta}_\alpha^{(i)} \in \mathbb{T}_{\alpha^{(i)}} \cap \mathbb{T}_{\alpha^{(i+1)}}$ containing $\theta_\alpha^{(i)}$, and pick $\boldsymbol{\theta}_\beta^{(j)} \in \mathbb{T}_{\beta^{(j)}} \cap \mathbb{T}_{\beta^{(j+1)}}$ containing $\theta_\beta^{(j)}$. Let $\bar{\boldsymbol{\theta}}_\alpha^{(i)} = \boldsymbol{\theta}_\alpha^{(i)} - \{\theta_\alpha^{(i)}\}$ and $\bar{\boldsymbol{\theta}}_\beta^{(j)} = \boldsymbol{\theta}_\beta^{(j)} - \{\theta_\beta^{(j)}\}$ be the corresponding intersection points in the destabilized diagram.

Let x_+ and x_- be the two points of intersection between $\alpha_1^{(1)}$ and $\beta_1^{(1)}$ such that there is a bigon from x_+ to x_- going over z_1 (and also one going over z_2). Similarly, let y_+ and y_- be the two points of intersection between $\alpha_1^{(k)}$ and $\beta_1^{(l)}$ such that there is a bigon from y_+ to y_- going over z_1.

We want to study holomorphic $(k+l)$-gons with boundaries on $\boldsymbol{\alpha}^{(k)}, \dots, \boldsymbol{\alpha}^{(1)}, \boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(l)}$, in this clockwise order, and with vertices at $\boldsymbol{\theta}_\alpha^{(i)}$ and $\boldsymbol{\theta}_\beta^{(j)}$ as well as arbitrary $\boldsymbol{x} \in \mathbb{T}_{\alpha^{(1)}} \cap \mathbb{T}_{\beta^{(1)}}$ and $\boldsymbol{y} \in \mathbb{T}_{\alpha^{(k)}} \cap \mathbb{T}_{\beta^{(l)}}$. A class ϕ of such holomorphic $(k+l)$-gons can be viewed as the connected sum of a class $\overline{\phi}$ on the destabilized diagram, and a class ψ of polygons on the sphere $\mathcal{S}$. Furthermore, given $\overline{\phi}$ and intersection points $r_1 \in \alpha_1^{(1)} \cap \beta_1^{(1)}$ and $r_2 \in \alpha_1^{(k)} \cap \beta_1^{(l)}$, we define $\Phi(\overline{\phi}, r_1, r_2)$ as the set of classes $\phi = (\overline{\phi}, \psi)$ satisfying $n_{w_2}(\phi) = 0$, and with ψ going from r_1 to r_2. This set is nonempty only when $(r_1, r_2) = (x_+, y_+)$ or $(r_1, r_2) = (x_-, y_-)$, in which case it consists in splicing of a standard polygon ψ_0 on $\mathcal{S}$ (an example is the shaded region in Figure 17) and some beta boundary degenerations.

Proposition 7.11 *Choose generic almost complex structures J_β associated to a sufficiently large connected sum neck along c, and also sufficiently stretched along a curve c_β enclosing the $\beta_1^{(j)}$ curves, as in Figure 17. Then*:

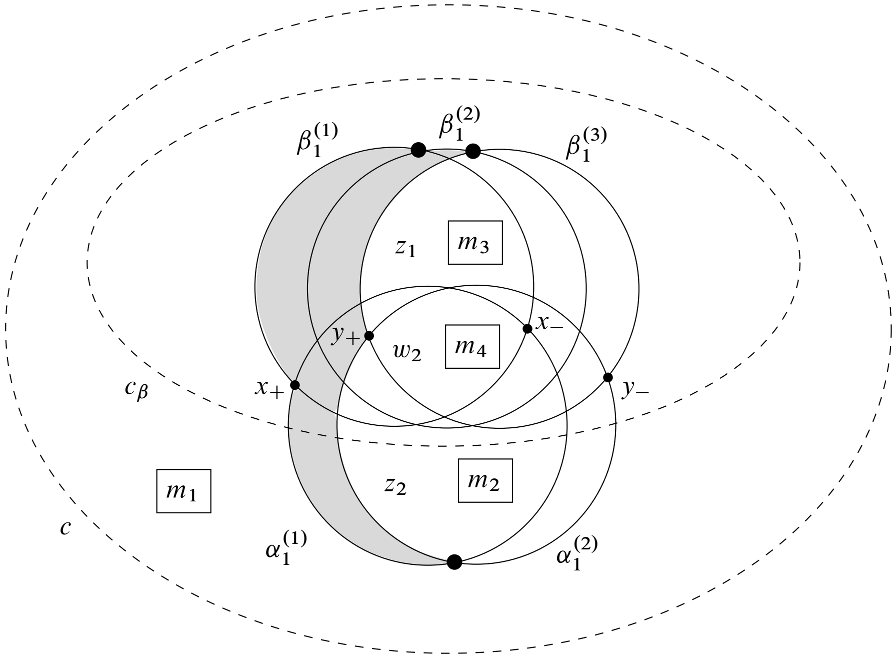

Figure 17: General polygons under an index zero/three stabilization. The black dots represent the θ intersection points. The shaded region is the domain of a pentagon class ψ_0.

(a) *When the components of $(\boldsymbol{x}, \boldsymbol{y})$ on $\mathcal{S}$ are the pair (x_+, y_-), the counts of holomorphic $(k+l)$-gons in classes*

$$\phi \in \pi_2(\boldsymbol{\theta}_\alpha^{(k-1)}, \dots, \boldsymbol{\theta}_\alpha^{(1)}, \boldsymbol{x}, \boldsymbol{\theta}_\beta^{(1)}, \dots, \boldsymbol{\theta}_\beta^{(l-1)}, \boldsymbol{y})$$

are always zero.

(b) *Suppose $\boldsymbol{x} = \bar{\boldsymbol{x}} \times r_1$ and $\boldsymbol{y} = \bar{\boldsymbol{y}} \times r_2$ with $(r_1, r_2) = (x_+, y_+)$ or (x_-, y_-). Then the counts of holomorphic $(k+l)$-gons in classes*

$$\phi \in \pi_2(\boldsymbol{\theta}_\alpha^{(k-1)}, \dots, \boldsymbol{\theta}_\alpha^{(1)}, \boldsymbol{x}, \boldsymbol{\theta}_\beta^{(1)}, \dots, \boldsymbol{\theta}_\beta^{(l-1)}, \boldsymbol{y})$$

can be nonzero only when $n_{w_2}(\phi) = 0$. Further, for

$$\bar{\phi} \in \pi_2(\bar{\boldsymbol{\theta}}_\alpha^{(k-1)}, \dots, \bar{\boldsymbol{\theta}}_\alpha^{(1)}, \bar{\boldsymbol{x}}, \bar{\boldsymbol{\theta}}_\beta^{(1)}, \dots, \bar{\boldsymbol{\theta}}_\beta^{(l-1)}, \bar{\boldsymbol{y}}),$$

if we define $\Phi(\bar{\phi}, r_1, r_2)$ as the set of classes $\phi = (\bar{\phi}, \psi)$ satisfying $n_{w_2}(\phi) = 0$, and with ψ going from r_1 to r_2, then

$$\#\mathcal{M}(\bar{\phi}) = \sum_{\phi \in \Phi(\bar{\phi}, r_1, r_2)} \#\mathcal{M}(\phi) \pmod 2. \tag{82}$$

Proof We use similar arguments to those in the proof of Proposition 6.18.

In the limit as we stretch the connected sum neck along c, holomorphic polygons become splicings of (possibly broken) holomorphic polygons on the two sides. The homology classes of polygons split also, as $\phi = (\overline{\phi}, \psi)$, where $\overline{\phi}$ is on $\overline{\Sigma}$ and ψ on $\mathcal{S}$.

Unlike for quasistabilizations, now the curve α_1 does not go through the point p, and hence cannot serve as boundary for the domains on $\overline{\Sigma}$. Hence, $\overline{\phi}$ is now an ordinary class of polygons (rather than a pair of a polygon and an α boundary degeneration), so the analysis becomes somewhat simpler.

Consider a few domain multiplicities for ϕ, shown in Figure 17:

$$m_1 = n_p(\phi), \quad m_2 = n_{z_2}(\phi), \quad m_3 = n_{z_1}(\phi), \quad m_4 = n_{w_2}(\phi).$$

The analogue of Lemma 6.10 gives the index of the class ϕ:

$$\mu(\phi) = \mu(\overline{\phi}) + \mu(\psi) - 2m_1.$$

There is also an analogue of Lemma 6.13, which gives the index of the component ψ:

$$\mu(\psi) = (3-k-l) + m_1 + m_2 + m_3 + m_4.$$

Hence,

$$\mu(\phi) = \mu(\overline{\phi}) + m_2 + m_3 + m_4 - m_1. \tag{83}$$

Further, by the relations for a domain to be acceptable (see Lemma 6.7), we have

$$m_3 - m_4 = m_1 - m_2 + \delta,$$

with

$$\delta = \begin{cases} 1 & \text{if } (r_1, r_2) = (x_+, y_-), \\ 0 & \text{if } (r_1, r_2) = (x_+, y_+) \text{ or } (x_-, y_-), \\ -1 & \text{if } (r_1, r_2) = (x_-, y_+). \end{cases}$$

From here and (83), we get

$$\mu(\phi) = \mu(\overline{\phi}) + 2m_4 + \delta. \tag{84}$$

We are interested in classes ϕ with $\mu(\phi) = 3-k-l$, so that we can count rigid holomorphic $(k+l)$-gons. If such a class ϕ contains holomorphic representatives, we must have $\mu(\overline{\phi}) \geq 3-k-l$. The relation (84) implies that $2m_4 + \delta \leq 0$. Since $m_4 \geq 0$, this disallows the case $\delta = 1$; that is, we cannot have $(r_1, r_2) = (x_+, y_-)$. Part (a) is proved.

Part (b) deals with the case $\delta = 0$. Then $\mu(\overline{\phi}) \geq 3-k-l = \mu(\phi)$ and (84) imply that

$$\mu(\overline{\phi}) = 0, \quad m_4 = 0.$$

This shows that $\phi \in \Phi(\overline{\phi}, r_1, r_2)$. Observe that the classes in $\Phi(\overline{\phi}, r_1, r_2)$ are splicings of a standard polygon class ψ_0 on $\mathcal{S}$ (an example is the shaded region in Figure 17) and some disk classes (namely, exteriors of the curves $\alpha_1^{(i)}$ or $\beta_1^{(j)}$ in Figure 17).

We have convergence and gluing results entirely similar to those in the proof of Proposition 6.18; compare Propositions 6.14 and 6.15. From there we obtain the desired relation (82) between the holomorphic polygons in the class $\overline{\phi}$ and those in the possible classes ϕ. □

Remark 7.12 If we had a single set of alpha curves (and several betas), then the result of Proposition 7.11 would follow more directly from the analysis in Section 6.6; see also [Zemke 2015, Lemma 14.25]. In that case, we would actually get a stronger conclusion; for example, the holomorphic polygon counts would be zero for $(r_1, r_2) = (x_-, y_+)$ as well.

Remark 7.13 In the proof of Proposition 7.11 we actually did not use that the almost complex structures are stretched along c_β, only along c. However, we need to stretch along c_β in order to identify holomorphic disks — and thus express the domain and target of polygon maps in terms of mapping cones — for Propositions 7.16 and 7.14 below.

We present below two consequences of Proposition 7.11. They are similar in spirit to Propositions 6.21 and 6.20.

First, we consider the polygon maps on ordinary Floer complexes $\boldsymbol{CF}^-$, and study their behavior with respect to the variant of index zero/three stabilization shown in Figure 14. The relevant picture is obtained from Figure 17 by deleting w_2 and relabeling z_1 and z_2 as w_1 and w_2. In the stabilized diagram, we have a map

$$F : \boldsymbol{CF}^-(\mathbb{T}_{\alpha^{(1)}}, \mathbb{T}_{\beta^{(1)}}) \to \boldsymbol{CF}^-(\mathbb{T}_{\alpha^{(k)}}, \mathbb{T}_{\beta^{(l)}}), \quad F(\boldsymbol{x}) = f(\boldsymbol{\theta}_\alpha^{(k-1)} \otimes \cdots \otimes \boldsymbol{\theta}_\alpha^{(1)} \otimes \boldsymbol{x} \otimes \boldsymbol{\theta}_\beta^{(1)} \otimes \cdots \otimes \boldsymbol{\theta}_\beta^{(l-1)}),$$

whereas, in the original (destabilized) diagram, we have

$$\bar{F} : \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}^{(1)}}, \mathbb{T}_{\bar{\beta}^{(1)}}) \to \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}^{(k)}}, \mathbb{T}_{\bar{\beta}^{(l)}}), \quad F(\bar{\boldsymbol{x}}) = f(\bar{\boldsymbol{\theta}}_\alpha^{(k-1)} \otimes \cdots \otimes \bar{\boldsymbol{\theta}}_\alpha^{(1)} \otimes \bar{\boldsymbol{x}} \otimes \bar{\boldsymbol{\theta}}_\beta^{(1)} \otimes \cdots \otimes \bar{\boldsymbol{\theta}}_\beta^{(l-1)}).$$

Proposition 7.14 *In the setting of Figure 14, the polygon maps induced on $\boldsymbol{CF}^-$ commute* (*on the nose*) *with the projections ρ from Corollary 7.4*; *that is*,

$$\rho \circ F = \overline{F} \circ \rho.$$

Proof The map ρ takes generators containing x_+ or y_+ to zero, and otherwise sends both U_1 and U_2 to U_1. In view of this, the result follows readily from Proposition 7.11: if $\bar{\boldsymbol{y}}$ appears with some coefficient U_1^n in the polygon map (on the destabilized diagram) applied to some $\bar{\boldsymbol{x}}$, then $\bar{\boldsymbol{y}} \times y_-$ appears with a coefficient $U_1^{n_1} U_2^{n_2}$ in the polygon map applied to $\bar{\boldsymbol{x}} \times x_-$ such that $n_1 + n_2 = n$. This implies commutation with ρ. □

Remark 7.15 Proposition 7.14 easily extends to maps between generalized link Floer complexes, where these complexes may use both types of basepoints on the link components that are not involved in the stabilization move.

Secondly, we can look at an index zero/three stabilization as in Figure 17, and consider the polygon maps induced on generalized link Floer complexes

$$\begin{aligned} &F\colon \boldsymbol{A}^-(\mathbb{T}_{\alpha^{(1)}}, \mathbb{T}_{\beta^{(1)}}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_{\alpha^{(k)}}, \mathbb{T}_{\beta^{(l)}}, \boldsymbol{s}), \\ &F(\boldsymbol{x}) = f(\boldsymbol{\theta}_\alpha^{(k-1)} \otimes \cdots \otimes \boldsymbol{\theta}_\alpha^{(1)} \otimes \boldsymbol{x} \otimes \boldsymbol{\theta}_\beta^{(1)} \otimes \cdots \otimes \boldsymbol{\theta}_\beta^{(l-1)}) \end{aligned} \tag{85}$$

and

$$\begin{aligned} &\bar{F}\colon \boldsymbol{A}^-(\mathbb{T}_{\bar{\alpha}^{(1)}}, \mathbb{T}_{\bar{\beta}^{(1)}}, \boldsymbol{s}) \to \boldsymbol{A}^-(\mathbb{T}_{\bar{\alpha}^{(k)}}, \mathbb{T}_{\bar{\beta}^{(l)}}, \boldsymbol{s}), \\ &F(\bar{\boldsymbol{x}}) = f(\bar{\boldsymbol{\theta}}_\alpha^{(k-1)} \otimes \cdots \otimes \bar{\boldsymbol{\theta}}_\alpha^{(1)} \otimes \bar{\boldsymbol{x}} \otimes \bar{\boldsymbol{\theta}}_\beta^{(1)} \otimes \cdots \otimes \bar{\boldsymbol{\theta}}_\beta^{(l-1)}). \end{aligned} \tag{86}$$

We are interested in whether these maps commute with the projections ρ. In Corollary 7.3, the maps ρ on $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$ are induced from those on $\boldsymbol{CF}^-$, using the Alexander filtration. Further, homotopy inverses ι to the ρ maps on $\boldsymbol{CF}^-$ can be found from the proof of Lemma 3.14: they are given by

$$\iota(\boldsymbol{x}) = \boldsymbol{x} \times x_+, \quad \iota(\boldsymbol{y}) = \boldsymbol{y} \times y_+.$$

The same formulas must give homotopy inverses for the filtered versions of ρ, on $\boldsymbol{A}^-(\mathcal{H}, \boldsymbol{s})$.

Proposition 7.11 now shows that the polygon maps commute with ι, on the nose:

$$F \circ \iota = \iota \circ \bar{F}. \tag{87}$$

In the case when F and $\bar{F}$ are triangle maps (ie $k = 1$ and $l = 2$, or $k = 2$ and $l = 1$), the same argument as in the proof of Corollary 7.10 shows that these triangle maps commute with the homotopy inverses to ι, which are the maps ρ:

$$\rho \circ F \simeq \bar{F} \circ \rho. \tag{88}$$

For $k + l \geq 4$, however, the polygon maps F and $\bar{F}$ are not chain maps, so a chain homotopy as in (88) would not make sense. Nevertheless, we can state a similar result by considering hypercubes of chain complexes, as in Section 5. Suppose that

$$C = (C^\varepsilon, D^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$$

is a hypercube, where the C^ε are generalized link Floer complexes of the form $\boldsymbol{A}^-(\mathbb{T}_{\alpha(\varepsilon)}, \mathbb{T}_{\beta(\varepsilon)}, \boldsymbol{s})$ for some collections of curves $\boldsymbol{\alpha}(\varepsilon)$ and $\boldsymbol{\beta}(\varepsilon)$ depending on ε, and the maps D^ε are polygon maps of the form (85) above, with $k + l = \|\varepsilon\| + 2$. (Such hypercubes will appear throughout this paper; see Section 8.3 below.) Furthermore, suppose that all $\boldsymbol{\alpha}(\varepsilon)$ and $\boldsymbol{\beta}(\varepsilon)$ are obtained from collections $\bar{\boldsymbol{\alpha}}(\varepsilon)$ and $\bar{\boldsymbol{\beta}}(\varepsilon)$ by index zero/three stabilizations in the same spot. We have a hypercube

$$\bar{C} = (\bar{C}^\varepsilon, \bar{D}^\varepsilon)_{\varepsilon \in \mathbb{E}_n}$$

in the destabilized diagram, where the maps $\bar{D}^\varepsilon$ are of the form (86).

The analogue of (88) in this general situation is the following:

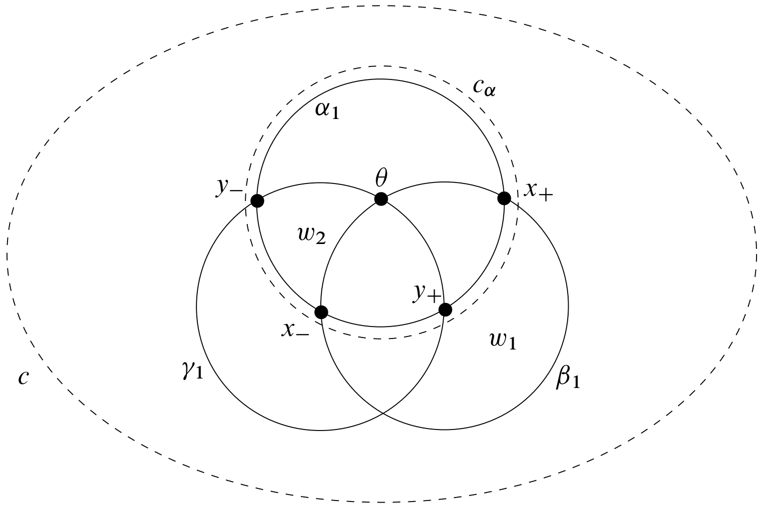

Figure 18: A strong equivalence between different kinds of stabilizations.

Proposition 7.16 *Consider two hypercubes of chain complexes $\overline{C}$ and C as above, related by an index zero/three link stabilization, with almost complex structures chosen as in Proposition 7.11. Then there exists a chain map (as in Definition 5.15)*

$$\Psi\colon C \to \overline{C}$$

whose components $\Psi^{\mathbf{0}}_{\varepsilon}\colon C^{\varepsilon} \to \overline{C}^{\varepsilon}$ are the projections ρ from Corollary 7.3.

Proof For the components of Ψ that increase the value of $\|\varepsilon\|$ by one, we choose the chain homotopies in (88); then (88) says that the condition for Ψ to be a chain map, equation (60), is satisfied along two-dimensional faces.

The higher components of Ψ are constructed inductively in the dimension of the faces. This is done using (87) and arguments similar to those in the proof of Corollary 7.10, but applied to chain maps between hypercubes. □

7.5 A strong equivalence

We now turn to discussing a move that will appear naturally in the context of index zero/three link stabilizations for complete systems of hyperboxes, in Section 8.9. This move is pictured in Figure 18, and consists in replacing the curve β_1 by γ_1 in that diagram, as well as replacing all the other beta curves (not shown) with small isotopic translates of themselves, intersecting them in two points. We denote by $\mathscr{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w})$ the initial diagram, by $\mathscr{H}' = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w})$ the final one, and by $\overline{\mathscr{H}} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{w}})$ and $\overline{\mathscr{H}}' = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\gamma}}, \overline{\boldsymbol{w}})$ the destabilized diagrams, with a single w in place of Figure 18. Note that $\overline{\mathscr{H}}$ and $\overline{\mathscr{H}}'$ are strongly equivalent, and in fact surface isotopic, in the sense of Definition 3.13. As usual, we view Σ as the connected sum of $\overline{\Sigma}$ and a sphere $\mathscr{S}$ that contains Figure 18.

Observe that $\mathcal{H}$ is exactly the variant of free index zero/three stabilization pictured in Figure 14, whereas $\mathcal{H}'$ is an ordinary free index zero/three stabilization of $\overline{\mathcal{H}}'$. Our move from $\mathcal{H}$ to $\mathcal{H}'$ is a strong equivalence — see Definition 3.13(a) — and it can be viewed as handlesliding β_1 over several other beta curves — those on the boundary of the component of $\overline{\Sigma} \setminus \boldsymbol{\beta}$ that contains w_1.

For each curve β_i not shown in the figure, let γ_i be its translate, and let $\theta_i \in \beta_i \cap \gamma_i$ be the intersection point that gives higher homological grading. We denote by $\boldsymbol{\theta} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$ the generator formed by all θ_i together with the point θ in Figure 18. Let $\bar{\boldsymbol{\theta}} = \boldsymbol{\theta} - \{\theta\} \in \mathbb{T}_{\bar{\beta}} \cap \mathbb{T}_{\bar{\gamma}}$ be the corresponding generator in the destabilized diagram.

To compute Floer complexes, we will use almost complex structures for sufficiently large neck stretching along the curves c and c_α in Figure 18. With these almost complex structures, recall from Section 7.3 that the Heegaard Floer complexes $\boldsymbol{CF}^-(\mathcal{H}, J_\alpha)$ and $\boldsymbol{CF}^-(\mathcal{H}', J_\alpha)$ can be described as the mapping cones (78) and (81), respectively. Thus, they are chain homotopy equivalent to $\boldsymbol{CF}^-(\overline{\mathcal{H}})$ and $\boldsymbol{CF}^-(\overline{\mathcal{H}}')$, respectively, via the projections ρ.

At the level of Heegaard Floer complexes, the strong equivalence from $\overline{\mathcal{H}}$ to $\overline{\mathcal{H}}'$ induces a map

$$\bar{F} : \boldsymbol{CF}^-(\overline{\mathcal{H}}) \to \boldsymbol{CF}^-(\overline{\mathcal{H}}')$$

given by counting holomorphic triangles with one vertex at $\bar{\boldsymbol{\theta}}$.

Similarly, the strong equivalence from $\mathcal{H}$ to $\mathcal{H}'$ induces a map

$$F_\alpha : \boldsymbol{CF}^-(\mathcal{H}, J_\alpha) \to \boldsymbol{CF}^-(\mathcal{H}', J_\alpha)$$

given by counting holomorphic triangles with one vertex at $\boldsymbol{\theta}$.

The map F_α is a version of the transition map computed in [Zemke 2015, Proposition 14.8]. Translated into our setting, his result reads:

Proposition 7.17 [Zemke 2015] *For generic almost complex structures J_α associated to a sufficiently large neck along the curves c and c_α, the map F_α is given by the 2×2 matrix*

$$F_\alpha = \begin{pmatrix} \bar{F} & 0 \\ 0 & \bar{F} \end{pmatrix}, \tag{89}$$

or, in a more expanded form,

$$\begin{tikzcd} C_+[\![U_2]\!] \arrow[r, "\bar{F}"] \arrow[d, "U_1-U_2"'] & C'_-[\![U_2]\!] \arrow[d, "U_1-U_2"] \\ C_-[\![U_2]\!] \arrow[r, "\bar{F}"] & C'_+[\![U_2]\!] \end{tikzcd}$$

where $C = \boldsymbol{CF}^-(\overline{\mathcal{H}})$ and $C' = \boldsymbol{CF}^-(\overline{\mathcal{H}}')$.

Remark 7.18 Since the $\bar{\boldsymbol{\beta}}$ and $\bar{\boldsymbol{\gamma}}$ curves are obtained from each other by small isotopies, the triangle map $\bar{F}$ is chain homotopic to the nearest-point map; see [Ozsváth and Stipsicz 2010, proof of Theorem 6.6]. On homology, we can think of the map induced by $\bar{F}$ as the identity.

Corollary 7.19 *Under the hypotheses of Proposition 7.17, we have a commutative diagram*

$$\begin{tikzcd} \boldsymbol{CF}^-(\mathcal{H}, J_\alpha) \arrow[r, "F_\alpha"] \arrow[d, "\rho"'] & \boldsymbol{CF}^-(\mathcal{H}', J_\alpha) \arrow[d, "\rho"] \\ \boldsymbol{CF}^-(\overline{\mathcal{H}}) \arrow[r, "\overline{F}"] & \boldsymbol{CF}^-(\overline{\mathcal{H}}') \end{tikzcd}$$

Proof This is an immediate consequence of Proposition 7.17. □

While in Proposition 7.17 we worked with almost complex structures J_α stretched along c and c_α, in this paper we will mostly work in the setting of Section 7.2; that is, with almost complex structures J_β, stretched along c and c_β (or J_γ, stretched along c and c_γ). Recall from (76) that $\boldsymbol{CF}^-(\mathcal{H}, J_\beta)$ can be identified with the cone

$$C_+^{U_1 \to U_2}[\![U_1]\!] \xrightarrow{U_1 - U_2} C_-^{U_1 \to U_2}[\![U_1]\!].$$

Further, $\boldsymbol{CF}^-(\overline{\mathcal{H}}, J_\gamma)$ can be identified with the cone from (75), namely

$$C_-[\![U_2]\!] \xrightarrow{U_1 - U_2} C_+[\![U_2]\!].$$

We can go from $\boldsymbol{CF}^-(\mathcal{H}, J_\beta)$ to $\boldsymbol{CF}^-(\mathcal{H}', J_\gamma)$ by first changing the almost complex structures from J_β to J_α as in Proposition 7.6, then applying the triangle map F_α from Proposition 7.17, and then changing the almost complex structure from J_α to J_γ as in Proposition 7.9; that is, we define

$$F : \boldsymbol{CF}^-(\mathcal{H}, J_\beta) \to \boldsymbol{CF}^-(\mathcal{H}', J_\gamma), \quad F = \Phi_{J_\alpha \to J_\gamma} \circ F_\alpha \circ \Phi_{J_\beta \to J_\alpha}.$$

Our main interest is the relation of F with the destabilization maps ρ from Corollaries 7.3 and 7.4.

Corollary 7.20 *The following diagram commutes up to a chain homotopy (represented by the diagonal map):*

$$\begin{tikzcd} \boldsymbol{CF}^-(\mathcal{H}, J_\alpha) \arrow[r, "F"] \arrow[d, "\rho"'] \arrow[dr] & \boldsymbol{CF}^-(\mathcal{H}', J_\gamma) \arrow[d, "\rho"] \\ \boldsymbol{CF}^-(\overline{\mathcal{H}}) \arrow[r, "\overline{F}"] & \boldsymbol{CF}^-(\overline{\mathcal{H}}') \end{tikzcd}$$

Proof Put together the results of Corollaries 7.8, 7.19 and 7.10. □

We now turn to discussing how higher polygon maps interact with the move in Figure 18.

Specifically, in Figure 18, we will introduce several isotopic translates of each curve. Suppose we have a Heegaard multidiagram

$$(\overline{\Sigma}, \overline{\boldsymbol{\alpha}}^{(1)}, \dots, \overline{\boldsymbol{\alpha}}^{(k)}, \overline{\boldsymbol{\beta}}^{(1)}, \dots, \overline{\boldsymbol{\beta}}^{(l)}, \overline{\boldsymbol{\gamma}}^{(1)}, \dots, \overline{\boldsymbol{\gamma}}^{(m)}, \overline{\boldsymbol{w}}, \boldsymbol{z}),$$

and let

$$(\Sigma, \boldsymbol{\alpha}^{(1)}, \dots, \boldsymbol{\alpha}^{(k)}, \boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(l)}, \boldsymbol{\gamma}^{(1)}, \dots, \boldsymbol{\gamma}^{(m)}, \boldsymbol{w}, \boldsymbol{z})$$

be obtained from it by adding a new basepoint w_2 and new curves

$$\alpha_1^{(i)}, \quad i = 1, \dots, k; \qquad \beta_1^{(i)}, \quad i = 1, \dots, l; \qquad \gamma_1^{(i)}, \quad i = 1, \dots, m$$

that are translates of α_1, β_1 and γ_1 from Figure 18. Furthermore, we want the new curves $\alpha_1^{(i)}$ to intersect one another at two points each, and similarly for the new beta and gamma curves.

We mark the top-degree intersection points $\theta_\alpha^{(i)} \in \alpha_1^{(i)} \cap \alpha_1^{(i+1)}$ for $i = 1, \dots, k-1$, and similarly $\theta_\beta^{(i)}$ and $\theta_\gamma^{(i)}$. Pick $\boldsymbol{\theta}_\alpha^{(i)} \in \mathbb{T}_{\alpha^{(i)}} \cap \mathbb{T}_{\alpha^{(i+1)}}$ containing $\theta_\alpha^{(i)}$, and let $\bar{\boldsymbol{\theta}}_\alpha^{(i)} = \boldsymbol{\theta}_\alpha^{(i)} - \{\theta_\alpha^{(i)}\}$. We pick $\boldsymbol{\theta}_\beta^{(i)}$ containing $\theta_\beta^{(i)}$ and $\boldsymbol{\theta}_\gamma^{(i)}$ containing $\theta_\gamma^{(i)}$, and define $\bar{\boldsymbol{\theta}}_\beta^{(i)}$ and $\bar{\boldsymbol{\theta}}_\gamma^{(i)}$ similarly.

Furthermore, we assume that the curves $\bar{\boldsymbol{\beta}}^{(l)}$ and $\bar{\boldsymbol{\gamma}}^{(1)}$ are small isotopic translates of each other, intersecting in two points. We let $\bar{\boldsymbol{\theta}}$ be the top intersection point in $\mathbb{T}_{\bar{\beta}^{(l)}} \cap \mathbb{T}_{\bar{\gamma}^{(1)}}$, and then set $\boldsymbol{\theta} = \bar{\boldsymbol{\theta}} \times \theta$, where $\theta \in \beta_1^{(l)} \cap \gamma_1^{(1)}$ is as in Figure 18.

We now count holomorphic $(k+l+m)$-gons with boundaries, in clockwise order, on $\boldsymbol{\alpha}^{(k)}, \dots, \boldsymbol{\alpha}^{(1)}$, $\boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(l)}$, $\boldsymbol{\gamma}^{(1)}, \dots, \boldsymbol{\gamma}^{(m)}$ and with vertices at $\boldsymbol{\theta}_\alpha^{(i)}$, $\boldsymbol{\theta}_\beta^{(i)}$, $\boldsymbol{\theta}$ and $\boldsymbol{\theta}_\gamma^{(i)}$ as well as arbitrary $\boldsymbol{x} \in \mathbb{T}_{\alpha^{(1)}} \cap \mathbb{T}_{\beta^{(1)}}$ and $\boldsymbol{y} \in \mathbb{T}_{\alpha^{(k)}} \cap \mathbb{T}_{\gamma^{(m)}}$. If we use almost complex structures J_α stretched along c and c_α (where c_α encloses the curves $\alpha_1^{(i)}$), the holomorphic polygon counts give rise to a map

$$G_\alpha \colon \boldsymbol{CF}^-(\mathbb{T}_{\alpha^{(1)}}, \mathbb{T}_{\beta^{(1)}}, J_\alpha) \to \boldsymbol{CF}^-(\mathbb{T}_{\alpha^{(k)}}, \mathbb{T}_{\gamma^{(m)}}, J_\alpha).$$

There is a corresponding map in the destabilized diagram

$$\overline{G} \colon \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}^{(1)}}, \mathbb{T}_{\bar{\beta}^{(1)}}) \to \boldsymbol{CF}^-(\mathbb{T}_{\bar{\alpha}^{(k)}}, \mathbb{T}_{\bar{\gamma}^{(m)}}).$$

We can also consider almost complex structures J_β stretched along c and c_β (where c_β encloses $\beta_1^{(1)}$), and J_γ stretched along c and c_γ (where c_γ encloses $\gamma_1^{(m)}$). By pre- and post-composing with change of almost complex structure maps, we can define

$$G = \Phi_{J_\alpha \to J_\gamma} \circ G_\alpha \circ \Phi_{J_\beta \to J_\alpha} \colon \boldsymbol{CF}^-(\mathbb{T}_{\alpha^{(1)}}, \mathbb{T}_{\beta^{(1)}}, J_\beta) \to \boldsymbol{CF}^-(\mathbb{T}_{\alpha^{(k)}}, \mathbb{T}_{\gamma^{(m)}}, J_\gamma).$$

The analysis done for holomorphic triangles in [Zemke 2015, Proposition 14.8] extends to the case of higher polygons (see also the proof of Proposition 7.11). We obtain the following result:

Proposition 7.21 *The higher polygon maps G_α commute with the projections ρ:*

$$\rho \circ G_\alpha = \overline{G} \circ \rho.$$

Observe that Proposition 7.21 implies an analogous result for the maps G instead of G_α, but with the commutation only up to homotopy, and assuming that the maps G are chain maps:

$$\rho \circ G \simeq \overline{G} \circ \rho.$$

Of course, in general, the polygon maps G are not chain maps, but rather fit into hypercubes of chain complexes similar to those in the setting of Proposition 7.16. The "commutation up to homotopy" can then be phrased in terms of the existence of a chain map between hypercubes, which has the maps ρ along its edges.

Remark 7.22 The discussion in this section was for Heegaard Floer complexes, using only w basepoints. However, if there are some z basepoints (away from the sphere $\mathscr{S}$, so that w_1 and w_2 are still free), the same results apply to generalized link Floer complexes.

8 Hyperboxes of Heegaard diagrams

In this section we define the notion of a complete system of hyperboxes for a link. As advertised in the introduction, this notion is the basic input for the surgery theorem.

Let $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{g+k-1})$ and $\boldsymbol{\beta}' = (\beta'_1, \dots, \beta'_{g+k-1})$ be two collections of curves on $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$. We will need the following terminology:

Definition 8.1 Suppose that, for any i, the curve β'_i is obtained from β_i by an isotopy such that β_i and β'_i intersect each other transversely, in exactly two points, and do not intersect any of the other curves in the diagrams. If this is the case, we write $\boldsymbol{\beta} \approx \boldsymbol{\beta}'$, and we say that $\boldsymbol{\beta}'$ *approximates* $\boldsymbol{\beta}$. Replacing $\boldsymbol{\beta}$ by $\boldsymbol{\beta}'$ is called an *approximation*.

Fix a multipointed, generic, admissible Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ for a link L, and $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$. If $\boldsymbol{\beta}'$ approximates $\boldsymbol{\beta}$ sufficiently closely, then there is a nearest-point map from $\boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s})$ to $\boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta'}, \boldsymbol{s})$ taking each intersection point to the corresponding nearest intersection point.

Lemma 8.2 *If $\boldsymbol{\beta}'$ approximates $\boldsymbol{\beta}$ sufficiently closely, then the nearest-point map is an isomorphism of chain complexes.*

Proof When the approximation is sufficiently small, the nearest-point map, which is clearly an isomorphism of modules, coincides with a continuation map obtained from varying the almost complex structure. The latter is always a chain map; see [Ozsváth and Szabó 2004b, Section 7.3]. □

8.1 $\boldsymbol{\beta}$-hyperboxes

Let $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$ be collections of curves on $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$ that are strongly equivalent. Note that the Heegaard diagram $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}', \boldsymbol{w}, \boldsymbol{z})$ is link-minimal, and represents an unlink inside the connected sum of several copies of $S^1 \times S^2$. If the pair $(\boldsymbol{\beta}, \boldsymbol{\beta}')$ is generic and admissible, we can define a generalized Floer chain complex $\boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \boldsymbol{0})$ as in Section 3.2; see in particular Remark 3.9. (Here, $\boldsymbol{0}$ is the zero vector.) We will use the alternative description of that complex from Section 3.6, and thus denote it by $\mathfrak{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \boldsymbol{0})$. The resulting Floer homology $H_*(\mathfrak{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \boldsymbol{0}))$ is the generalized Heegaard Floer homology of an unlink inside the connected sum of several copies of $S^1 \times S^2$, hence it equals the homology of a torus; see the completion of the proof of Theorem 3.16 at the end of Section 3.5. As such, there is a well-defined maximal-degree element $\theta_{\beta,\beta'} \in H_*(\mathfrak{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \boldsymbol{0}))$. Observe that, if the surface Σ has genus g, then this maximal degree is $\mu = \frac{1}{2}g$. In the particular case when $\boldsymbol{\beta} \approx \boldsymbol{\beta}'$, there is also a canonical cycle (intersection point) $\Theta^{\text{can}}_{\beta,\beta'}$ representing $\theta_{\beta,\beta'}$.

Recall the notation from Section 5.1; in particular, pick $\boldsymbol{d} = (d_1, \dots, d_n) \in (\mathbb{Z}_{\geq 0})^n$, $n \geq 0$, and consider the set of multi-indices $\mathbb{E}(\boldsymbol{d})$. We say that two multi-indices $\varepsilon, \varepsilon' \in \mathbb{E}(\boldsymbol{d})$ with $\varepsilon \leq \varepsilon'$ are *neighbors* if $\varepsilon' - \varepsilon \in \mathbb{E}_n$. Note that, in the definition of a hyperbox of chain complexes (Definition 5.1), we only have linear maps $D^{\varepsilon'-\varepsilon}\colon C^{\varepsilon} \to C^{\varepsilon'}$ in the case when ε and ε' are neighbors.

The hyperbox $\mathbb{E}(\boldsymbol{d})$ can be viewed as a union of several unit hypercubes: if all d_i's are nonzero, there are $d_1 d_2 \cdots d_n$ unit hypercubes of dimension n, whereas, if some d_i's are zero, we get a union of unit hypercubes of smaller dimension. More precisely, let $\boldsymbol{d}^{\circ} = (d_1^{\circ}, \dots, d_n^{\circ})$ consist of the values $d_i^{\circ} = \max(d_i - 1, 0)$, and let n° be the number of nonzero d_i's. Then $\mathbb{E}(\boldsymbol{d})$ is the union of the unit hypercubes $\varepsilon + \mathbb{E}_{n^{\circ}}$ for $\varepsilon \in \boldsymbol{d}^{\circ}$.

Definition 8.3 An *empty β-hyperbox* of size $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$ on a fixed surface with marked points $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$ consists of a collection $\{\boldsymbol{\beta}^{\varepsilon}\}_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}$ of strongly equivalent sets of attaching beta curves $\boldsymbol{\beta}^{\varepsilon}$, indexed by $\varepsilon \in \mathbb{E}(\boldsymbol{d})$. Further, we require that, for each unit hypercube of the form $\varepsilon + \mathbb{E}_{n^{\circ}}$ with $\varepsilon \in \boldsymbol{d}^{\circ}$, the corresponding Heegaard multidiagram $(\Sigma, \{\boldsymbol{\beta}^{\varepsilon'}\}_{\varepsilon' \in (\varepsilon + \mathbb{E}_{n^{\circ}})}, \boldsymbol{w}, \boldsymbol{z})$ is generic and admissible.

Definition 8.4 Let $\{\beta^{\varepsilon}\}_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}$ be an empty β-hyperbox, consisting of diagrams on a surface of genus g. Set $\mu = \frac{1}{2} g$. A *filling* Θ of the hyperbox consists of chain elements

$$\Theta_{\varepsilon,\varepsilon'} \in \mathfrak{A}^{-}_{\mu + \|\varepsilon' - \varepsilon\| - 1}(\mathbb{T}_{\beta^{\varepsilon}}, \mathbb{T}_{\beta^{\varepsilon'}}, \mathbf{0}),$$

one for each pair $(\varepsilon, \varepsilon')$ such that $\varepsilon < \varepsilon'$ and ε and ε' are neighbors. The chains $\Theta_{\varepsilon,\varepsilon'}$ are required to satisfy the following conditions:

- When $\varepsilon < \varepsilon'$ and $\|\varepsilon' - \varepsilon\| = 1$ (ie ε and ε' are the endpoints of an edge in the hyperbox), $\Theta_{\varepsilon,\varepsilon'}$ is a cycle representing the maximal-degree element $\theta_{\beta^{\varepsilon}, \beta^{\varepsilon'}}$ in Floer homology.
- For any $\varepsilon < \varepsilon'$ such that ε and ε' are neighbors,

$$\sum_{l=1}^{\|\varepsilon' - \varepsilon\|} \sum_{\{\varepsilon = \varepsilon^0 < \cdots < \varepsilon^l = \varepsilon'\}} f(\Theta_{\varepsilon^0, \varepsilon^1} \otimes \cdots \otimes \Theta_{\varepsilon^{l-1}, \varepsilon^l}) = 0. \tag{90}$$

The data consisting of an empty β-hyperbox $\{\beta^{\varepsilon}\}_{\varepsilon \in \mathbb{E}(\boldsymbol{d})}$ and a filling by $\{\Theta_{\varepsilon,\varepsilon'}\}$ is simply called a *β-hyperbox*.

Remark 8.5 The simplest kind of β-hyperbox is a β-hypercube, ie one with $\boldsymbol{d} = (1, \dots, 1)$, so that $\mathbb{E}(\boldsymbol{d}) = \mathbb{E}_n$. Then the Heegaard multidiagram consisting of all 2^n curve collections has to be admissible. Further, any two $\varepsilon, \varepsilon' \in \mathbb{E}_n$ with $\varepsilon \leq \varepsilon'$ are neighbors, so, for any such ε and ε', the pair $(\boldsymbol{\beta}^{\varepsilon}, \boldsymbol{\beta}^{\varepsilon'})$ needs to come equipped with a chain $\Theta_{\varepsilon,\varepsilon'}$ such that these chains satisfy (90).

This kind of β-hypercube is quite natural, and has already appeared implicitly in the Heegaard Floer literature, eg in [Ozsváth and Szabó 2005]. There are two reasons why we need the more general kind of hyperbox. The first is that we want to allow for some of the beta pairs in the hyperbox (those for which ε and ε' are not neighbors) to form nonadmissible diagrams. These nonadmissible diagrams make an appearance in the construction of basic systems; see Section 8.7 below. The second reason is that general hyperboxes appear naturally in the context of grid diagrams; see Section 15.6 below.

Lemma 8.6 *Any empty β-hyperbox admits a filling. Moreover, if a filling is partially defined on the m-skeleton of a β-hyperbox* (*ie the elements $\Theta_{\varepsilon,\varepsilon'}$ are defined only for neighbors ε and ε' with $\|\varepsilon'-\varepsilon\| \le m$, and satisfy the required conditions*), *it can be extended to a filling on the whole hyperbox.*

Proof We construct the chain elements $\Theta_{\varepsilon,\varepsilon'}$ inductively on $\|\varepsilon'-\varepsilon\|$.

When $\|\varepsilon'-\varepsilon\| = 1$, we choose arbitrary cycles representing the maximal-degree elements in homology.

When $\|\varepsilon'-\varepsilon\| > 1$, suppose we want to define $\Theta_{\varepsilon,\varepsilon'}$ and we have defined elements $\Theta_{\gamma,\gamma'}$, whenever $\|\gamma'-\gamma\| < \|\varepsilon'-\varepsilon\|$, satisfying (90). Set

$$c = \sum_{l=2}^{\|\varepsilon'-\varepsilon\|} \sum_{\{\varepsilon=\varepsilon^0<\cdots<\varepsilon^l=\varepsilon'\}} f(\Theta_{\varepsilon^0,\varepsilon^1} \otimes \cdots \otimes \Theta_{\varepsilon^{l-1},\varepsilon^l}).$$

Using (12), we obtain

$$\begin{aligned}
\partial c &= \sum_{l=2}^{\|\varepsilon'-\varepsilon\|} \sum_{\{\varepsilon=\varepsilon^0<\cdots<\varepsilon^l=\varepsilon'\}} f(f(\Theta_{\varepsilon^0,\varepsilon^1} \otimes \cdots \otimes \Theta_{\varepsilon^{l-1},\varepsilon^l})) \\
&= \sum_{l=2}^{\|\varepsilon'-\varepsilon\|} \sum_{\{\varepsilon=\varepsilon^0<\cdots<\varepsilon^l=\varepsilon'\}} \sum_{\substack{0\le i<j\le l \\ (i,j)\neq(0,l)}} f(\Theta_{\varepsilon^0,\varepsilon^1} \otimes \cdots \otimes f(\Theta_{\varepsilon^i,\varepsilon^{i+1}} \otimes \cdots \otimes \Theta_{\varepsilon^{j-1},\varepsilon^j}) \otimes \cdots \otimes \Theta_{\varepsilon^{l-1},\varepsilon^l}) \\
&= 0.
\end{aligned}$$

Here, for the last equality we applied (90) to the pair $(\varepsilon^i, \varepsilon^j)$.

Thus, c is a cycle in $\mathfrak{A}^-_{\mu+\|\varepsilon'-\varepsilon\|-2}(\mathbb{T}_{\beta^\varepsilon}, \mathbb{T}_{\beta^{\varepsilon'}}, \mathbf{0})$. If $\|\varepsilon'-\varepsilon\| = 2$, then c is the sum of two terms of the form $f(\Theta_{\varepsilon,\varepsilon^1} \otimes \Theta_{\varepsilon_1,\varepsilon'})$. Both of these terms represent the nonzero homology class of maximal degree (compare [Ozsváth and Szabó 2005, proof of Lemma 4.5]), and therefore their sum c is a boundary. If $\|\varepsilon'-\varepsilon\| > 2$, then the respective homology group $\mathfrak{A}^-_{\mu+\|\varepsilon'-\varepsilon\|-2}(\mathbb{T}_{\beta^\varepsilon}, \mathbb{T}_{\beta^{\varepsilon'}}, \mathbf{0})$ is zero (being beyond the maximal degree μ), so c is again a boundary. Thus, in any case we can choose $\Theta_{\varepsilon,\varepsilon'}$ such that $\partial\Theta_{\varepsilon,\varepsilon'} = c$. Then (90) is satisfied for the pair $(\varepsilon, \varepsilon')$. □

8.2 Hyperboxes of strongly equivalent Heegaard diagrams

We define an *α-hyperbox* on $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$ to be the same as a β-hyperbox, except we denote the collections of curves by $\boldsymbol{\alpha}$'s and, for any neighbors $\varepsilon < \varepsilon'$, we are given elements $\Theta_{\varepsilon',\varepsilon} \in \mathfrak{A}^-(\mathbb{T}_{\alpha^{\varepsilon'}}, \mathbb{T}_{\alpha^\varepsilon}, \mathbf{0})$ rather than in $\mathfrak{A}^-(\mathbb{T}_{\alpha^\varepsilon}, \mathbb{T}_{\alpha^{\varepsilon'}}, \mathbf{0})$. The compatibility relation (90) has to be modified accordingly:

$$\sum_{l=1}^{\|\varepsilon'-\varepsilon\|} \sum_{\{\varepsilon'=\varepsilon^0>\cdots>\varepsilon^l=\varepsilon\}} f(\Theta_{\varepsilon^0,\varepsilon^1} \otimes \cdots \otimes \Theta_{\varepsilon^{l-1},\varepsilon^l}) = 0. \tag{91}$$

Now suppose that we have $\boldsymbol{d} = (d_1, \dots, d_n) \in (\mathbb{Z}_{\geq 0})^n$ for some $n \geq 0$, and we are given maps

$$r_i : \{1, \dots, d_i\} \to \{\alpha, \beta\}, \quad i = 1, \dots, n.$$

We can then assign to each edge $(\varepsilon, \varepsilon')$ of the hyperbox $\mathbb{E}(\boldsymbol{d})$ a symbol $r(\varepsilon, \varepsilon') \in \{\alpha, \beta\}$ as follows: if the edge is parallel to the i^{th} axis and its projection to that axis is the segment $[j-1, j]$, we choose $r(\varepsilon, \varepsilon') = r_i(j)$. We refer to $\{r_i\}_{i=1}^n$ as *bipartition maps*.

Given bipartition maps as above and $\varepsilon \in \mathbb{E}(\boldsymbol{d})$, we set

$$\varepsilon_i^\alpha = \#(r_i^{-1}(\alpha) \cap \{1, \dots, \varepsilon_i\}), \quad \varepsilon_i^\beta = \#(r_i^{-1}(\beta) \cap \{1, \dots, \varepsilon_i\})$$

for $i = 1, \dots, n$. These define n-tuples $\varepsilon^\alpha \in \mathbb{E}(\boldsymbol{d}^\alpha)$ and $\varepsilon^\beta \in \mathbb{E}(\boldsymbol{d}^\beta)$, where $\boldsymbol{d}^\alpha = (d_i^\alpha)_{i=1}^n$ and $\boldsymbol{d}^\beta = (d_i^\beta)_{i=1}^n$ (where, of course, $d_i^\alpha = \# r_i^{-1}(\alpha)$ and $d_i^\beta = \# r_i^{-1}(\beta)$). Note that $\varepsilon_i^\alpha + \varepsilon_i^\beta = \varepsilon_i$ for any i.

Definition 8.7 Choose $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$ and bipartition maps r_i for $i = 1, \dots, n$. A *hyperbox $\mathscr{H}$ of strongly equivalent Heegaard diagrams* consists of an α-hyperbox of size $\boldsymbol{d}^\alpha \in (\mathbb{Z}_{\geq 0})^n$ and a β-hyperbox of size $\boldsymbol{d}^\beta \in (\mathbb{Z}_{\geq 0})^n$, both on the same multipointed surface $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$. These are required to satisfy the following condition: For each multi-index $\varepsilon \in \mathbb{E}(\boldsymbol{d})$, we can consider the Heegaard diagram

$$\mathscr{H}_\varepsilon = (\Sigma, \boldsymbol{\alpha}^{\varepsilon^\alpha}, \boldsymbol{\beta}^{\varepsilon^\beta}, \boldsymbol{w}, \boldsymbol{z}).$$

For each unit hypercube $\varepsilon + \mathbb{E}_{n^\circ} \subseteq \mathbb{E}(\boldsymbol{d})$, the curve collections appearing in the diagrams $\mathscr{H}_{\varepsilon'}$, for all $\varepsilon' \in \varepsilon + \mathbb{E}_{n^\circ}$, are required to form a generic, admissible Heegaard multidiagram.

We can view some of the information in $\mathscr{H}$ in the following way. The hyperbox $\mathscr{H}$ has at each of its vertices a Heegaard diagram $\mathscr{H}_\varepsilon$ such that all of these are strongly equivalent, and each edge in the hyperbox corresponds to changing either the alpha or the beta curves. Further, we have elements $\Theta^\alpha_{\varepsilon',\varepsilon} = \Theta_{\varepsilon',\varepsilon}$ for $\varepsilon < \varepsilon'$ neighbors in $\mathbb{E}(\boldsymbol{d}^\alpha)$, and $\Theta^\beta_{\varepsilon,\varepsilon'} = \Theta_{\varepsilon,\varepsilon'}$ for $\varepsilon < \varepsilon'$ neighbors in $\mathbb{E}(\boldsymbol{d}^\beta)$. We usually refer to a hyperbox $\mathscr{H}$ as going between the Heegaard diagrams $\mathscr{H}_{(0,\dots,0)}$ and $\mathscr{H}_{(d_1,\dots,d_n)}$.

Remark 8.8 Suppose we have a hyperbox $\mathscr{H}$ of strongly equivalent Heegaard diagrams representing a link $\vec{L} \subset Y$. Then the hyperbox $\mathscr{H}$ naturally gives rise to a *reduced hyperbox* $r_{\vec{M}}(\mathscr{H})$, with the same size and bipartition maps, made of the reduced Heegaard diagrams $(r_{\vec{M}}(\mathscr{H}))_\varepsilon = r_{\vec{M}}(\mathscr{H}_\varepsilon)$; see Definition 3.22. Indeed, in $r_{\vec{M}}(\mathscr{H})$ we can take the images of the Θ-chain elements from $\mathscr{H}$ under the corresponding inclusion maps defined in Section 3.8.

8.3 Hyperboxes of Floer complexes

Let $\mathscr{H}$ be a hyperbox of strongly equivalent Heegaard diagrams, of size $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$ and with partition maps r_i for $i = 1, \dots, n$. At each vertex we have an admissible Heegaard diagram $\mathscr{H}_\varepsilon$. Let us assume that $\mathscr{H}_\varepsilon$ represents a link $\vec{L}$ in an integral homology sphere Y. Fix $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$. We will construct an associated hyperbox of chain complexes as in Section 5.1. To each vertex $\varepsilon \in \mathbb{E}(\boldsymbol{d})$ we assign the generalized Floer chain complex

$$C_*^{\varepsilon} = \boldsymbol{A}_*^-(\mathcal{H}_{\varepsilon}, \boldsymbol{s}) = \boldsymbol{A}_*^-(\mathbb{T}_{\alpha^{\varepsilon^\alpha}}, \mathbb{T}_{\beta^{\varepsilon^\beta}}, \boldsymbol{s});$$

compare Section 3.2.

We define the linear maps $D_{\varepsilon}^{\varepsilon'-\varepsilon} \colon C^{\varepsilon} \to C^{\varepsilon'}$ by

$$D_{\varepsilon}^{\varepsilon'-\varepsilon}(\boldsymbol{x}) = \sum_{l,p} \sum_{\{\varepsilon'^{\alpha}=\gamma^0>\cdots>\gamma^l=\varepsilon^{\alpha}\}} \sum_{\{\varepsilon^{\beta}=\zeta^0<\cdots<\zeta^p=\varepsilon'^{\beta}\}} f(\Theta^{\alpha}_{\gamma^0,\gamma^1}\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{l-1},\gamma^l}\otimes\boldsymbol{x}\otimes\Theta^{\beta}_{\zeta^0,\zeta^1}\otimes\cdots\otimes\Theta^{\beta}_{\zeta^{p-1},\zeta^p}).$$

Example 8.9 When $\varepsilon = \varepsilon'$, the corresponding map $D^{\boldsymbol{0}}_{\varepsilon}$ is just the differential $f = \partial$ on the generalized Floer complex $\boldsymbol{A}_*^-(\mathcal{H}_{\varepsilon}, \boldsymbol{s})$.

Example 8.10 When $\|\varepsilon' - \varepsilon\| = 1$, ie $(\varepsilon, \varepsilon')$ is an edge in the hypercube, $D_{\varepsilon}^{\varepsilon'-\varepsilon}$ is given by counting holomorphic triangles with one specified vertex. Since that vertex is a Θ-cycle representing the maximal-degree element in homology, the chain map $D_{\varepsilon}^{\varepsilon'-\varepsilon}$ is a chain homotopy equivalence.

Example 8.11 When $\|\varepsilon' - \varepsilon\| = 2$, we distinguish two cases, according to whether the edges of the square between ε and ε' are marked with only one or both of the symbols α and β. In the first case, assuming they are marked with β, the map $D_{\varepsilon}^{\varepsilon'-\varepsilon}$ is a sum of three different polygon maps: one counting triangles with one vertex at $\Theta^{\beta}_{\varepsilon^{\beta},\varepsilon'^{\beta}}$ and two counting quadrilaterals with two vertices at $\Theta^{\beta}_{\varepsilon^{\beta},\zeta}$ and $\Theta^{\beta}_{\varepsilon^{\beta},\zeta}$, where ζ is an intermediate multi-index between ε^{β} and ε'^{β}. (There are two possibilities for ζ.) In the second case, when ε and ε' are the vertices of a square with two edges marked with α and two with β, the map $D_{\varepsilon}^{\varepsilon'-\varepsilon}$ is given by counting quadrilaterals with two specified vertices at $\Theta^{\beta}_{\varepsilon^{\beta},\varepsilon'^{\beta}}$ and $\Theta^{\alpha}_{\varepsilon'^{\alpha},\varepsilon^{\alpha}}$.

Lemma 8.12 $A^-(\mathcal{H}, \boldsymbol{s}) = ((C^{\varepsilon})_{\varepsilon\in\mathbb{E}(\boldsymbol{d})}, (D^{\varepsilon})_{\varepsilon\in\mathbb{E}_n})$ *is a hyperbox of chain complexes.*

Proof We need to check (38), ie that, for any $\varepsilon, \varepsilon'' \in \mathbb{E}_n$,

$$\sum_{\varepsilon<\varepsilon'<\varepsilon''} D_{\varepsilon'}^{\varepsilon''-\varepsilon'} \circ D_{\varepsilon}^{\varepsilon'-\varepsilon} = 0. \tag{92}$$

Indeed, the summation in (92) equals

$$\sum f\big(\Theta^{\alpha}_{\gamma^0,\gamma^1}\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{i-1},\gamma^i}\otimes f(\Theta^{\alpha}_{\gamma^i,\gamma^{i+1}}\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{l-1},\gamma^l}\otimes\boldsymbol{x}\otimes\Theta^{\beta}_{\zeta^0,\zeta^1}\otimes\cdots\otimes\Theta^{\beta}_{\zeta^{j-1},\zeta^j}) \otimes\Theta^{\beta}_{\zeta^j,\zeta^{j+1}}\otimes\cdots\otimes\Theta^{\beta}_{\zeta^{p-1},\zeta^p}\big),$$

where the sum is taken over all possible l, p, i, j and multi-indices $\varepsilon'^{\alpha} = \gamma^0 > \cdots > \gamma^l = \varepsilon^{\alpha}$ and $\varepsilon^{\beta} = \zeta^0 < \cdots < \zeta^p = \varepsilon'^{\beta}$. Applying (12), we find that this sum further equals

$$\sum f(\Theta^{\alpha}_{\gamma^0,\gamma^1}\otimes\cdots\otimes f(\Theta^{\alpha}_{\gamma^{i-1},\gamma^i}\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{j-1},\gamma^j})\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{l-1},\gamma^l}\otimes\boldsymbol{x}\otimes\Theta^{\beta}_{\zeta^0,\zeta^1}\otimes\cdots\otimes\Theta^{\beta}_{\zeta^{p-1},\zeta^p})$$
$$+\sum f(\Theta^{\alpha}_{\gamma^0,\gamma^1}\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{l-1},\gamma^l}\otimes\boldsymbol{x}\otimes\Theta^{\beta}_{\zeta^0,\zeta^1}\otimes\cdots\otimes f(\Theta^{\beta}_{\zeta^{i-1},\zeta^i}\otimes\cdots\otimes\Theta^{\beta}_{\zeta^{j-1},\zeta^j})\otimes\cdots\otimes\Theta^{\beta}_{\zeta^{p-1},\zeta^p}).$$

Both of these sums vanish. Indeed, let us fix i, $l-j$, p and all γ's and ζ's in the first sum, except for $\gamma^i,\dots,\gamma^{j-1}$. Then the corresponding sum of $f(\Theta^{\alpha}_{\gamma^{i-1},\gamma^i}\otimes\cdots\otimes\Theta^{\alpha}_{\gamma^{j-1},\gamma^j})$ over $\gamma^i,\dots,\gamma^{j-1}$ is zero by (91). Similarly, by applying (90) we find that the second sum is also zero. □

In the case where the hyperbox $\mathscr{H}$ consists of link-minimal diagrams, we can use the alternative description of $\boldsymbol{A}^-(\mathscr{H}_\varepsilon,\boldsymbol{s})$ as free complexes $\mathfrak{A}^-_*(\mathscr{H}_\varepsilon,\boldsymbol{s})$, as in Section 3.6. We then denote the resulting hyperbox of chain complexes by $\mathfrak{A}^-(\mathscr{H},\boldsymbol{s})$.

8.4 Moves on hyperboxes

In this section we describe a series of moves on hyperboxes of strongly equivalent Heegaard diagrams. In light of Lemma 8.6, we see that hyperboxes are easy to construct, and thus are rather flexible objects.

Recall the list of Heegaard moves between multipointed Heegaard diagrams from Section 3.3. Suppose now that $\mathscr{H}$ is a hyperbox of strongly equivalent Heegaard diagrams. We have a similar list of *hyperbox Heegaard moves* on $\mathscr{H}$:

(i) A *three-manifold isotopy* of $\mathscr{H}$ consists of applying the same three-manifold isotopy to all the Heegaard diagrams $\mathscr{H}^\varepsilon$ simultaneously, keeping the same partition maps and taking the Θ-chain elements (which are linear combinations of collections of intersection points) to the corresponding linear combinations in the isotopic diagrams.

(ii) An *index one/two stabilization* of a hyperbox $\mathscr{H}$ consists of a simultaneous index one/two stabilization of all diagrams $\mathscr{H}_\varepsilon$, in the same position. Note that if two collections of attaching curves are strongly equivalent, they remain so after the stabilization. With regard to the fillings, we need to pair each Θ-chain element with the unique intersection point between the new alpha curve and new beta curve. The fact that the new Θ-chain elements still satisfy the relations (90) follows from the argument of [Ozsváth and Szabó 2006, Lemma 4.7] generalized to polygons.

(iii) A *free index zero/three stabilization* of a hyperbox $\mathscr{H}$ consists of a simultaneous free index zero/three stabilization of all diagrams $\mathscr{H}_\varepsilon$, in the same position. We again note that if two collections of attaching curves are strongly equivalent, they remain so after the stabilization. With regard to the fillings, we pair each Θ-chain element with the maximal-degree intersection point between the new alpha and beta curves. The fact that the new Θ-chain elements still satisfy the relations (90) follows from the first part of Proposition 6.20.

(iv) An *index zero/three link stabilization* of a hyperbox $\mathscr{H}$ consists of a simultaneous index zero/three link stabilization of all diagrams $\mathscr{H}_\varepsilon$, in the same position. The fillings are constructed as in (iv), by pairing with the maximal-degree intersection point between the new curves. The relations (90) are then a consequence of Proposition 7.16.

(v) Instead of curve isotopies and handleslides, we now have global shifts, defined as follows. Consider two hyperboxes $\mathscr{H}$ and $\mathscr{H}'$ having the same size $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$ and the same partition maps r_i. A *global shift* S from $\mathscr{H}$ to $\mathscr{H}'$ is an $(n+1)$-dimensional hyperbox of size $(\boldsymbol{d}, 1) \in (\mathbb{Z}_{\geq 0})^{n+1}$ such that its subhyperbox corresponding to $\varepsilon_{n+1} = 0$ is $\mathscr{H}$ and its subhyperbox corresponding to $\varepsilon_{n+1} = 1$ is $\mathscr{H}'$. Note that there are two kinds of global shifts, α and β, according to the value of the map $r_{n+1} \colon \{1\} \to \{\alpha, \beta\}$. Observe that, by definition, a global shift does not change the Heegaard surface with basepoints $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$.

(vi) An *elementary enlargement* will be defined later (see Definition 8.17).

(vii) Inverses to the above. In particular, the inverse process to a stabilization is called *destabilization*, and the inverse of an elementary enlargement is called a *contraction*.

Note that the moves (i)–(v) (and their inverses) preserve the size of the respective hyperbox.

As a particular kind of three-manifold isotopy, we define a *surface isotopy* of $\mathscr{H}$ to consist in applying the same surface isotopy to all the Heegaard diagrams simultaneously, compare Definition 3.13(b). If $\mathscr{H}$ and $\mathscr{H}'$ are surface isotopic hyperboxes, we write $\mathscr{H} \cong \mathscr{H}'$.

Let us now focus on understanding global shifts further. Note that a global shift induces a chain map between the respective hyperboxes of generalized Floer chain complexes; compare Definition 5.15. It turns out that this map is always a chain homotopy equivalence. Before proving this fact, we need some preliminaries.

Definition 8.13 Let $\mathscr{H}$ and $\mathscr{H}'$ be two hyperboxes of strongly equivalent Heegaard diagrams such that either

(a) the corresponding alpha curves on $\mathscr{H}$ and $\mathscr{H}'$ coincide, whereas the corresponding beta curve collections approximate each other ($\beta^\varepsilon \approx {\beta'}^\varepsilon$) in the sense of Definition 8.1; or

(b) the corresponding beta curves coincide, and the corresponding alpha curve collections approximate each other.

Further, suppose that the approximations are sufficiently small (for Lemma 8.2 to hold) and suppose the Θ-chains correspond to each other under the respective nearest-point maps. We then say that $\mathscr{H}'$ *approximates* $\mathscr{H}$, and write $\mathscr{H} \approx \mathscr{H}'$. If we are in case (a), we call this an approximation of type β, and if we are in case (b), we call it of type α.

Lemma 8.14 *Suppose $\mathscr{H} \approx \mathscr{H}'$, where the approximation is of type β. Then there exists a canonical global shift between $\mathscr{H}$ and $\mathscr{H}'$ such that along the new edges we see the canonical elements* (*mentioned in Section 8.1*)

$$\Theta^{\mathrm{can}}_\varepsilon = \Theta^{\mathrm{can}}_{\boldsymbol{\beta}^\varepsilon, \boldsymbol{\beta}'^\varepsilon} \in \mathbb{T}_{\beta^\varepsilon} \cap \mathbb{T}_{\beta'^\varepsilon}.$$

Proof For simplicity, we explain the construction in the case when the hyperboxes are one-dimensional of length one. Thus, $\mathscr{H}$ consists of two curve collections $\boldsymbol{\beta}^0$ and $\boldsymbol{\beta}^1$ and a chain element Θ relating them.

Similarly, $\mathscr{H}'$ consists of curve collections $\boldsymbol{\beta}'^0$ and $\boldsymbol{\beta}'^1$ and a chain element Θ' that is the image of Θ under the nearest-point map. Along the new edges we are required to place the two intersection points Θ_0^{can} and Θ_1^{can}.

We also have a cycle

$$\Theta'' \in \mathfrak{A}^-(\mathbb{T}_{\beta^0}, \mathbb{T}_{\beta^{1'}}, \mathbf{0})$$

that is the image of Θ under the nearest-point map changing $\boldsymbol{\beta}^1$ into $\boldsymbol{\beta}'^1$. (This is a cycle because of Lemma 8.2.) Recall that there is a canonical homotopy between the continuation map $\mathfrak{A}^-(\mathbb{T}_{\beta^{0'}}, \mathbb{T}_{\beta^{1'}}, \mathbf{0}) \to \mathfrak{A}^-(\mathbb{T}_{\beta^0}, \mathbb{T}_{\beta^{1'}}, \mathbf{0})$ and the triangle map $\boldsymbol{x} \mapsto f(\Theta_0^{\text{can}} \otimes \boldsymbol{x})$; see [Lipshitz 2006, proof of Proposition 11.4] or [Ozsváth and Stipsicz 2010, proof of Theorem 6.6]. This chain homotopy is defined by counting holomorphic bigons with a varying almost complex structure and stretching of the neck. Thus,

$$f(\Theta_0^{\text{can}} \otimes \Theta') = \Theta'' + dH_0,$$

where H_0 is the image of Θ' under the canonical chain homotopy. Similarly,

$$f(\Theta \otimes \Theta_1^{\text{can}}) = \Theta'' + dH_1,$$

where H_1 is the image of a map counting stretched bigons with one vertex at Θ. We then place the chain element $H_0 + H_1$ on the two-dimensional face of our global shift. Relation (90) is satisfied.

If $\mathscr{H}$ and $\mathscr{H}'$ are one-dimensional (of arbitrary length), we place chain elements as above on all their two-dimensional faces. Further, the construction generalizes to hyperboxes of higher dimension, by considering continuation maps with several stretched necks. □

The global shift constructed in Lemma 8.14 is called the *identity shift*. There exists a similar identity shift between hyperboxes that approximate each other through an approximation of type α.

Now suppose $\mathscr{H} \approx \mathscr{H}'$ and fix $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$. Let $(C^\varepsilon, D^\varepsilon)$ and $(C'^\varepsilon, D'^\varepsilon)$ be the hyperboxes of generalized Floer complexes (for the value $\boldsymbol{s}$) associated to $\mathscr{H}$ and $\mathscr{H}'$, respectively. Note that there is a nearest-point map from C^ε to C'^ε, defined by taking each generator (intersection point) to its image under the approximation. This is a chain map provided that the almost complex structures are chosen in a compatible way. In fact, by adding zeros on the higher-dimensional faces, we obtain a chain map between the respective hyperboxes of generalized Floer complexes. We refer to this as the nearest-point map, too.

Lemma 8.15 *An identity shift induces a map on the hyperboxes of Floer complexes which is chain homotopic to the nearest-point map.*

Proof The triangle map induced by an approximation is chain homotopic to the nearest-point map; see [Ozsváth and Stipsicz 2010, proof of Theorem 6.6]. The argument in [loc. cit.] can also be applied to higher polygons, with the result that the chain homotopies between triangle and nearest-point maps lift to chain homotopies between the respective hyperboxes. □

Lemma 8.16 *The chain map $F(S)$ induced by any global shift S on hyperboxes of Floer complexes is a chain homotopy equivalence.*

Proof Let S be a global shift between two n-dimensional hyperboxes $\mathcal{H}, \mathcal{H}'$ of size $\boldsymbol{d}$, so that S has size $(\boldsymbol{d}, 1)$. Let $\mathcal{H}'', \mathcal{H}'''$ be two hyperboxes that approximate $\mathcal{H}$ and each other.

We construct an $(n+2)$-dimensional hyperbox $\widetilde{S}$ of size $(\boldsymbol{d}, 1, 1)$ as follows. Its subhyperbox corresponding to $\varepsilon_{n+2} = 0$ is S, so that $\varepsilon_{n+1} = \varepsilon_{n+2} = 0$ corresponds to $\mathcal{H}$, and $\varepsilon_{n+1} = 1$ and $\varepsilon_{n+2} = 0$ is $\mathcal{H}'$. Its subhyperbox corresponding to $\varepsilon_{n+1} = 1$ is any global shift between $\mathcal{H}'$ and $\mathcal{H}''$; for example, the reverse S^r of S composed with the identity shift from $\mathcal{H}$ to $\mathcal{H}''$. Thus $\varepsilon_{n+1} = \varepsilon_{n+2} = 1$ corresponds to $\mathcal{H}''$. For $\varepsilon_{n+1} = 0$ and $\varepsilon_{n+2} = 1$, we take the other hyperbox $\mathcal{H}'''$ that approximates $\mathcal{H}$. For $\varepsilon_{n+1} = 0$ we choose the identity shift between $\mathcal{H}$ and $\mathcal{H}'''$, and for $\varepsilon_{n+2} = 1$ the identity shift between $\mathcal{H}'''$ and $\mathcal{H}''$. We then fill in the remaining faces of the hyperbox $\widetilde{S}$ (those corresponding to increasing both ε_{n+1} and ε_{n+2} by 1) with Θ-chain elements in an arbitrary way.

On the level of Floer complexes, the hyperbox $\widetilde{S}$ then produces a chain homotopy between $F(S^r) \circ F(S)$ and the identity. □

Definition 8.17 We define an *elementary enlargement* of a hyperbox $\mathcal{H}$ of strongly equivalent Heegaard diagrams by analogy with the corresponding concept from Section 5.8. Let $\boldsymbol{d} \in (\mathbb{Z}_{\geq 0})^n$ be the size of $\mathcal{H}$. Pick $i_0 \in \{1, \dots, n\}$ and $j_0 \in \{0, 1, \dots, d_k\}$. Define $\boldsymbol{d}^+ = \boldsymbol{d} + \tau_k$ as in Section 5.8. We also choose a symbol $\kappa = \alpha$ or β.

We construct a new hyperbox $\mathcal{H}^+$ of size $\boldsymbol{d}^+$, with bipartition maps $r_i^+ = r_i$ for $i \neq i_0$, and

$$r_{i_0}^+(j) = \begin{cases} r_{i_0}(j) & \text{if } j \leq j_0, \\ \kappa & \text{if } j = j_0 + 1, \\ r_{i_0}(j-1) & \text{if } j \geq j_0 + 2. \end{cases}$$

The hyperbox $\mathcal{H}^+$ is obtained from $\mathcal{H}$ by splitting the latter into two halves along the hyperplane $\varepsilon_{i_0} = j_0$, adding a new subhyperbox at $\varepsilon_{i_0} = j_0 + 1$ that approximates the one at $\varepsilon_{i_0} = j_0$, and inserting the identity global shift between them. We say that H^+ is obtained from H by an elementary enlargement. The reverse process is called *elementary contraction*.

In Section 5.8 we defined a similar notion of elementary enlargement for hyperboxes of chain complexes. If $\mathcal{H}^+$ is the elementary enlargement of a hyperbox of Heegaard diagrams $\mathcal{H}$, let $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$ and $\mathfrak{A}^-(\mathcal{H}^+, \boldsymbol{s})$ be the respective hyperboxes of generalized Floer chain complexes, as in Section 8.3. Let also $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})^+$ be the corresponding elementary enlargement of the hyperbox of chain complexes $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$. By applying Lemma 8.15, we see that $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})^+$ is chain homotopy equivalent to $\mathfrak{A}^-(\mathcal{H}^+, \boldsymbol{s})$.

Lemma 8.18 *Let $\mathcal{H}$ and $\mathcal{H}'$ be two hyperboxes of strongly equivalent Heegaard diagrams representing the same oriented link $\vec{L} \subset Y$ and having the same dimension. Then*:

(a) *The hyperboxes $\mathcal{H}$ and $\mathcal{H}'$ can be related by a sequence of hyperbox Heegaard moves.*

(b) *Suppose further that the diagrams in $\mathcal{H}$ are link-minimal, and the diagrams in $\mathcal{H}'$ have m basepoints of type z, and therefore $m-\ell$ subsidiary w basepoints, as in Proposition 3.12. (We choose the same basepoints to be subsidiary in all the diagrams in $\mathcal{H}'$.) Then $\mathcal{H}'$ can be obtained from $\mathcal{H}$ using a sequence of hyperbox Heegaard moves that includes exactly $m-\ell$ index zero/three link stabilizations, with each of these stabilizations introducing a subsidiary basepoint.*

Proof For part (a), use three-manifold isotopies, stabilizations, and destabilizations to transform the hyperboxes into two new ones that have the same underlying Heegaard surface $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$. After a few more index one/two stabilizations as in [Manolescu et al. 2009a, Lemma 2.4], we can also arrange so that the alpha curves in $\mathcal{H}$ are strongly equivalent to the alpha curves in $\mathcal{H}'$, and the beta curves in $\mathcal{H}$ are strongly equivalent to the beta curves in $\mathcal{H}'$. We can then use elementary enlargements to arrange so that the hyperboxes $\mathcal{H}$ and $\mathcal{H}'$ have the same size and bipartition maps. Using the winding procedure from [Ozsváth and Szabó 2004b, Section 5], we can find a sequence of empty hyperboxes (of the same size and with the same bipartition maps) that interpolate between the two: $\mathcal{H}=\mathcal{H}^0, \mathcal{H}^1, \dots, \mathcal{H}^p=\mathcal{H}'$ such that each pair $(\mathcal{H}^{j-1}, \mathcal{H}^j)$ forms an empty hyperbox of one dimension bigger (satisfying the required admissibility conditions). We then choose arbitrary fillings of the intermediate hyperboxes $\mathcal{H}^j$ (see Lemma 8.6) as well as arbitrary global shifts between $\mathcal{H}^{j-1}$ and $\mathcal{H}^j$ for $j=1,\dots,p$. The result is a sequence of global shifts relating $\mathcal{H}$ and $\mathcal{H}'$.

Part (b) follows by combining the strategy above with Proposition 3.12. □

Recall the invariance statement about generalized Floer complexes for links (Theorem 3.16). We have an analogous result for hyperboxes of strongly equivalent Heegaard diagrams representing a link:

Proposition 8.19 *Let $\mathcal{H}, \mathcal{H}'$ be two hyperboxes of strongly equivalent Heegaard diagrams that represent the same oriented link $\vec{L} \subset Y$, and have the same dimension. Let $A^-(\mathcal{H}, \boldsymbol{s})$ and $A^-(\mathcal{H}', \boldsymbol{s})$ be the hyperboxes of generalized Floer chain complexes associated to $\mathcal{H}$ and $\mathcal{H}'$, respectively, for some $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$. Pick also natural inclusions $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H})$ and $\mathcal{R}(L) \hookrightarrow \mathcal{R}(\mathcal{H}')$, as in Theorem 3.16. Then the hypercubes obtained from $A^-(\mathcal{H}, \boldsymbol{s})$ and $A^-(\mathcal{H}', \boldsymbol{s})$ by compression (see Section 5.2) are chain homotopy equivalent over $\mathcal{R}(L)$, in the sense of Section 5.7.*

Proof By Lemma 8.18, it suffices to investigate the effect of hyperbox Heegaard moves on the corresponding (compressed) hyperboxes of generalized Floer complexes. Index one/two stabilizations and three-manifold isotopies produce chain homotopy equivalences between the respective hyperboxes; see [Ozsváth and Szabó 2004b; 2006]. The same goes for global shifts, according to Lemma 8.16. The chain homotopy equivalences then descend to the compressed hyperboxes, according to Lemma 5.19. Elementary enlargements also produce quasi-isomorphisms between the respective compressions; see Lemma 5.21. Free index zero/three stabilizations result in equivalences; see [Ozsváth and Szabó 2008a, Proposition 6.5] and Proposition 6.20. Index zero/three link stabilizations also result in equivalences, as a consequence of Corollary 7.3 and Proposition 7.16. □

8.5 Sublinks and hyperboxes

Let $\vec{L} \subset Y$ be an oriented link, and $M \subseteq L$ a sublink with m components. We choose an orientation $\vec{M}$ of M, not necessarily the one induced from $\vec{L}$. We denote by $\vec{L} - M$ the sublink $L - M$ with the orientation induced from $\vec{L}$.

Definition 8.20 A *hyperbox of Heegaard diagrams for the pair* $(\vec{L}, \vec{M})$ is an m-dimensional hyperbox $\mathcal{H}$ of strongly equivalent Heegaard diagrams representing the link $\vec{L} - M$, together with an ordering $M_1, \dots, M_m$ of the components of M.

Remark 8.21 In particular, a hyperbox of Heegaard diagrams for a pair $(\vec{L}, \varnothing)$ is simply a Heegaard diagram for $\vec{L}$.

Definition 8.20 may appear mysterious at first. The intuition behind it is that a hyperbox $\mathcal{H}$ for a pair $(\vec{L}, \vec{M})$ corresponds to subtracting the link $\vec{M}$ from $\vec{L}$. Indeed, although this is not required by the definition, the initial diagram $\mathcal{H}_{(0,\dots,0)}$ of $\mathcal{H}$ will always be obtained from a Heegaard diagram for $\vec{L}$ by deleting some of its basepoints. The final diagram will represent $\vec{L} - M$, and the ordering of the components is telling us the parts of the hyperbox associated to subtracting the respective components M_i from $\vec{L}$. Indeed, we can think of the ordering as a one-to-one correspondence between the coordinate axes of the hyperbox and the components of M, where the i^{th} coordinate corresponds to M_i.

To make this more precise, let $\mathcal{H}^{\vec{L},\vec{M}}$ be a hyperbox for the pair $(\vec{L}, \vec{M})$. Then, for each $M' \subseteq M$, we denote by $\mathcal{H}^{\vec{L},\vec{M}}(M')$ the Heegaard diagram $\mathcal{H}^{\vec{L},\vec{M}}_{\varepsilon(M')}$, where $\varepsilon(M')$ is the multi-index with components $\varepsilon(M')_i$ for $i = 1, \dots, m$, given by

$$\varepsilon(M')_i = \begin{cases} d_i & \text{if } M_i \subseteq M', \\ 0 & \text{otherwise.} \end{cases}$$

Informally, we think of $\mathcal{H}^{\vec{L},\vec{M}}(M')$ as the intermediate step in the hyperbox obtained after subtracting M' from $\vec{L}$. In particular, the initial and final vertex of the hyperbox $\mathcal{H}^{\vec{L},\vec{M}}$ are $\mathcal{H}^{\vec{L},\vec{M}}(\varnothing)$ and $\mathcal{H}^{\vec{L},\vec{M}}(M)$, respectively.

For every $M' \subseteq M$, there is a subhyperbox of $\mathcal{H}^{\vec{L},\vec{M}}$ going from $\mathcal{H}^{\vec{L},\vec{M}}(M')$ to $\mathcal{H}^{\vec{L},\vec{M}}(M)$. We denote it by $\mathcal{H}^{\vec{L},\vec{M}}(M', M)$. Note that $\mathcal{H}^{\vec{L},\vec{M}}(M', M)$ is a hyperbox associated to the pair $(\vec{L} - M', \vec{M} - M')$. We also denote by $\mathcal{H}^{\vec{L},\vec{M}}(\varnothing, M')$ the subhyperbox of $\mathcal{H}^{\vec{L},\vec{M}}$ going from $\mathcal{H}^{\vec{L},\vec{M}}(\varnothing)$ to $\mathcal{H}^{\vec{L},\vec{M}}(M')$, which is complementary to $\mathcal{H}^{\vec{L},\vec{M}}(M', M)$.

We say that two hyperboxes $\mathcal{H}^{\vec{L},\vec{M}}$ and $\widetilde{\mathcal{H}}^{\vec{L},\vec{M}}$ for the same pair $(\vec{L}, \vec{M})$ are *isotopic* if they have the same ordering of the components of M, and the underlying hyperboxes of strongly equivalent Heegaard diagrams are isotopic. If this is the case, we write $\mathcal{H}^{\vec{L},\vec{M}} \cong \widetilde{\mathcal{H}}^{\vec{L},\vec{M}}$.

Up to now, the orientation for M has not played any role in the definition of a hyperbox for a pair $(\vec{L}, \vec{M})$. However, it played a role when we discussed reduction (see Definition 3.22 and Remark 8.8). Observe that if $\mathcal{H}^{\vec{L},\vec{M}}$ is a hyperbox for a pair $(\vec{L}, \vec{M})$, its reduction at an oriented sublink $\vec{N} \subseteq L - M$, denoted by $r_{\vec{N}}(\mathcal{H}^{\vec{L},\vec{M}})$, is a hyperbox for the pair $(\vec{L} - N, \vec{M})$.

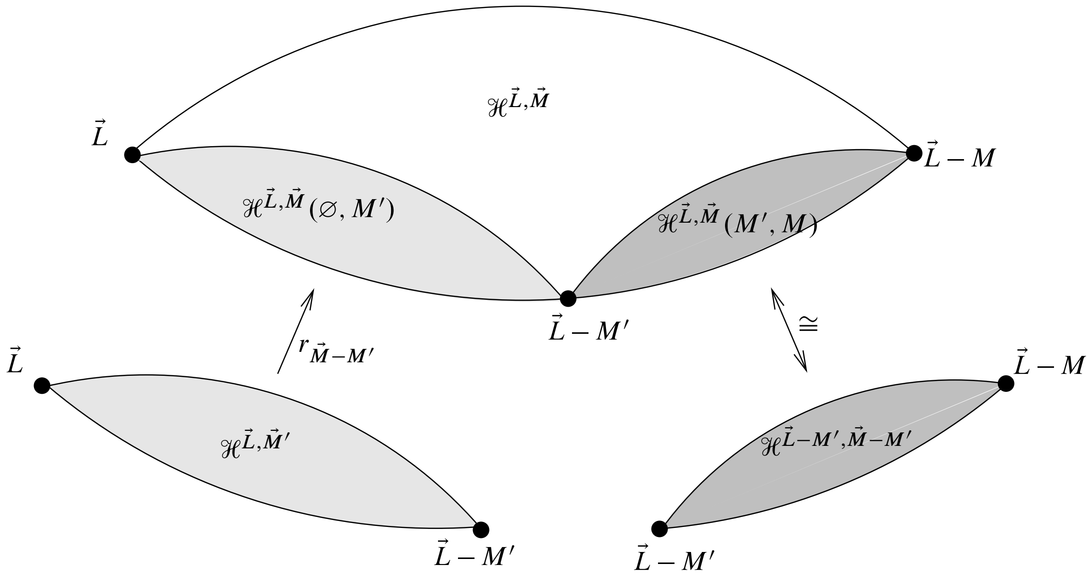

Figure 19: Compatibility between hyperboxes. We represent here hyperboxes schematically by lenses. On top we have a hyperbox $\mathcal{H}^{\vec{L}',\vec{M}}$ and two subhyperboxes, shown as sublenses with different shadings. The compatibility condition (a) in Definition 8.22 is the equality between the two lightly shaded lenses. Condition (b) is the equality between the two darkly shaded lenses.

We are now ready to mention the following notions of compatibility for hyperboxes (see Figure 19):

Definition 8.22 Let $\mathcal{H}^{\vec{L},\vec{M}}$ be a hyperbox for a pair $(\vec{L}, \vec{M})$, and $\mathcal{H}^{\vec{L},\vec{M}'}$ be a hyperbox for the pair $(\vec{L}, \vec{M}')$, where M' is a sublink of M with the orientation $\vec{M}'$ induced from $\vec{M}$. Let also $\mathcal{H}^{\vec{L}-M',\vec{M}-M'}$ be a hyperbox for the pair $(\vec{L}-M', \vec{M}-M')$.

(a) We say that the hyperboxes $\mathcal{H}^{\vec{L},\vec{M}}$ and $\mathcal{H}^{\vec{L},\vec{M}'}$ are *compatible* if

$$r_{\vec{M}-M'}(\mathcal{H}^{\vec{L},\vec{M}'}) = \mathcal{H}^{\vec{L},\vec{M}}(\varnothing, M'). \tag{93}$$

(b) We say that the hyperboxes $\mathcal{H}^{\vec{L},\vec{M}}$ and $\mathcal{H}^{\vec{L}-M',\vec{M}-M'}$ are *compatible* if there is a surface isotopy

$$\mathcal{H}^{\vec{L}-M',\vec{M}-M'} \cong \mathcal{H}^{\vec{L},\vec{M}}(M', M). \tag{94}$$

8.6 Complete systems of hyperboxes

In the following definition (and in the rest of this subsection), all sublinks of a link L that are denoted by L' will come with the orientation $\vec{L}'$ induced from $\vec{L}$ and, similarly, all sublinks in a link M that are denoted by M' will come with the orientation $\vec{M}'$ induced from $\vec{M}$.

Definition 8.23 A *complete presystem of hyperboxes* $\mathcal{H}$ representing the link $\vec{L}$ consists of a collection of hyperboxes subject to certain compatibility conditions, as follows. For each pair of subsets $M \subseteq L' \subseteq L$, and each orientation $\vec{M} \in \Omega(M)$, the complete presystem assigns a hyperbox $\mathcal{H}^{\vec{L}',\vec{M}}$ for the pair $(\vec{L}', \vec{M})$. Moreover, the hyperbox $\mathcal{H}^{\vec{L}',\vec{M}}$ is required to be compatible with both $\mathcal{H}^{\vec{L}',\vec{M}'}$ and $\mathcal{H}^{\vec{L}'-M',\vec{M}-M'}$.

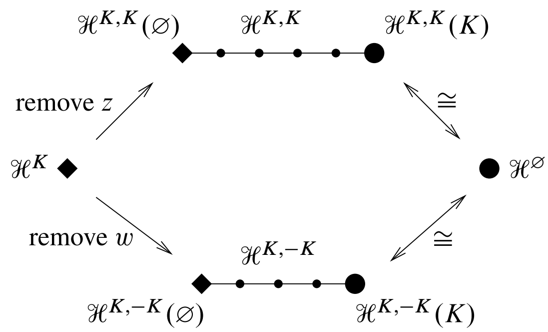

Figure 20: A complete presystem of hyperboxes for a knot $\vec{K}$. For simplicity, we assume here that $\mathcal{H}^K$ has no free basepoints, and exactly two linked basepoints (w and z). The symbols $\cong$ indicate surface isotopies.

In particular, note that a complete presystem contains hyperboxes of the form $\mathcal{H}^{\vec{L}',\varnothing}$, which are zero-dimensional; in other words, they consist of a single Heegaard diagram, which we denote by $\mathcal{H}^{L'}$. The diagram $\mathcal{H}^{L'}$ represents the sublink $L' \subseteq Y$.

As previously mentioned in Section 8.5, we think of the hyperbox $\mathcal{H}^{\vec{L}',\vec{M}}$ as a way of destabilizing $\vec{L}'$ at the components of $\vec{M}$. Indeed, $\mathcal{H}^{\vec{L}',\vec{M}}$ goes from $r_{\vec{M}}(\mathcal{H}^{L'})$ (ie $\mathcal{H}^{L'}$ with half of the basepoints on M deleted, according to the orientation $\vec{M}$) to $\mathcal{H}^{L'-M}$.

Example 8.24 Let $\vec{L} = \vec{K} \subset Y$ be a knot. Then a complete presystem of hyperboxes for $\vec{K}$ consists of two multipointed Heegaard diagrams, $\mathcal{H}^K$ for K and $\mathcal{H}^\varnothing$ for Y itself, together with two one-dimensional hyperboxes of strongly equivalent Heegaard diagrams: one, $\mathcal{H}^{K,K}$, going from a diagram $\mathcal{H}^{K,K}(\varnothing)$ (which is just $\mathcal{H}^K$ with the z points on K removed) to some diagram $\mathcal{H}^{K,K}(K)$ (surface isotopic to $\mathcal{H}^\varnothing$); and another one, $\mathcal{H}^{K,-K}$, going from the diagram $\mathcal{H}^{K,-K}(\varnothing)$ (which is just $\mathcal{H}^K$ with the w points removed), down to a diagram $\mathcal{H}^{K,-K}(K)$ (surface isotopic to $\mathcal{H}^\varnothing$). This is illustrated in Figure 20.

In short, a complete presystem of hyperboxes for a knot K produces a multipointed diagram $\mathcal{H}^K$ for the knot, together with a way of relating the diagram $\mathcal{H}^K$ with the z points on K removed to the diagram $\mathcal{H}^K$ with the w points on K removed, via a certain sequence of strongly equivalent Heegaard diagrams and Θ-cycles, plus some surface isotopies in the middle (moving the w basepoints into the z basepoints). Note that both of the diagrams that we relate in this fashion represent Y itself. Observe also that the sequence of Θ-cycles (and surface isotopies) induces a corresponding sequence of chain maps on Floer complexes as in Section 8.3. These chain maps are all chain homotopy equivalences; see Example 8.10. By composing them we obtain a chain homotopy equivalence between the initial and the final Floer complex. This is exactly the kind of structure that was used in [Ozsváth and Szabó 2008b], in the context of describing the Heegaard Floer homology of integer surgeries on knots. One should view complete presystems as a generalization of this structure to the case of links.

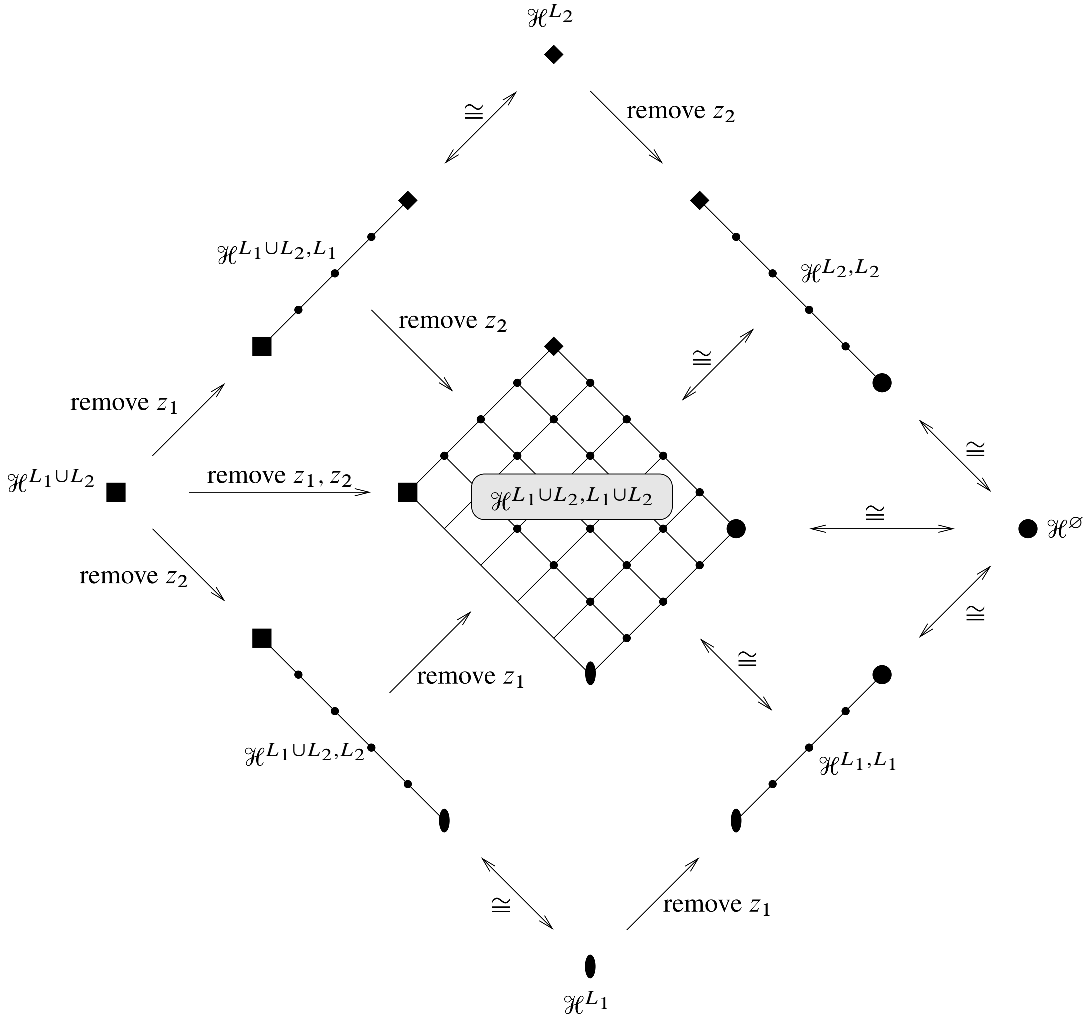

Figure 21: A part of a complete presystem of hyperboxes for a link $\vec{L} = L_1 \cup L_2$. For simplicity, we assume that $\mathcal{H}^L$ has exactly two basepoints w_1 and z_1 on L_1, and two basepoints w_2 and z_2 on L_2. We show here one quarter of the complete presystem, consisting of the two-dimensional hyperbox $\mathcal{H}^{L_1 \cup L_2, L_1 \cup L_2}$, and all smaller hyperboxes related to it by compatibility conditions.

Example 8.25 When $\vec{L} = L_1 \cup L_2$ is a link of two components, a complete presystem of hyperboxes for $\vec{L}$ consists of four zero-dimensional hyperboxes

$$\mathcal{H}^{L_1 \cup L_2}, \mathcal{H}^{L_1}, \mathcal{H}^{L_2}, \mathcal{H}^{\varnothing},$$

eight one-dimensional hyperboxes

$$\mathcal{H}^{L_1 \cup L_2, L_1}, \mathcal{H}^{L_1 \cup L_2, -L_1}, \mathcal{H}^{L_1 \cup L_2, L_2}, \mathcal{H}^{L_1 \cup L_2, -L_2}, \mathcal{H}^{L_1, L_1}, \mathcal{H}^{L_1, -L_1}, \mathcal{H}^{L_2, L_2}, \mathcal{H}^{L_2, -L_2},$$

and four two-dimensional hyperboxes

$$\mathcal{H}^{L_1 \cup L_2, L_1 \cup L_2}, \mathcal{H}^{L_1 \cup L_2, -L_1 \cup L_2}, \mathcal{H}^{L_1 \cup L_2, L_1 \cup -L_2}, \mathcal{H}^{L_1 \cup L_2, -L_1 \cup -L_2}.$$

These hyperboxes are related by various compatibility conditions. Some of these conditions are illustrated in Figure 21.

In the case of a knot $K \subset S^3$, the complete presystem only contains data for quasi-isomorphisms (in fact, chain homotopy equivalences) between various Heegaard Floer complexes of $Y = S^3$. At the level of homology, there is a unique $\mathbb{F}[\![U]\!]$-module isomorphism from $\boldsymbol{HF}^-(Y) = \mathbb{F}[\![U]\!]$ to itself, namely the identity. However, in the case of links (or even for knots, but in more general homology spheres Y), complete presystems give rise to quasi-isomorphisms between generalized Floer complexes of sublinks, whose homology can be complicated. Recall from Remark 3.17 that, in general, the resulting isomorphisms on generalized Floer homology are not canonical. To ensure naturality, we need to control the paths traced by basepoints under the surface isotopies involved in the compatibility relations (94).

We start with the following:

Definition 8.26 Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ be a multipointed Heegaard diagram for an oriented link $\vec{L} \subset Y$, with a corresponding Heegaard splitting $Y = U_\alpha \cup_\Sigma U_\beta$. For each basepoint $w \in \boldsymbol{w} \cap L$, we denote by $z(w) \in \boldsymbol{z}$ the *successor of* w *on* $\vec{L}$, that is, the basepoint on the same component $\vec{L}_j$ of $\vec{L}$ as w which appears just after w as we go around $\vec{L}_j$ according to its orientation. A *good set of trajectories*

$$\boldsymbol{c} = \{c_w \mid w \in \boldsymbol{w} \cap L\}$$

for $\mathcal{H}$ consists of disjoint, smoothly embedded paths

$$c_w \colon [0,1] \to \Sigma, \qquad c_w(0) = w, \quad c_w(1) = z(w),$$

such that each path c_w is homotopic (inside the handlebody U_α) to the minimal oriented segment l_w on $\vec{L}$ going from w to $z(w)$, by a homotopy whose interior avoids the link L. Further, we require that the image of each c_w is disjoint from the free basepoints (ie the basepoints in $\boldsymbol{w} \setminus L$).

Observe that the alpha and the beta curves did not play an essential role in the definition above; we could have made the same definition by starting with a Heegaard splitting (and suitable basepoints) instead of a multipointed Heegaard diagram.

Definition 8.27 Let $\vec{L} \subset Y$ be an oriented link. A *set of surface data* $(\Sigma, \boldsymbol{w}, \boldsymbol{z}, \boldsymbol{c})$ for $\vec{L}$ consists of

- an (oriented) embedded surface $\Sigma \subset Y$ that produces a Heegaard splitting $Y = U_\alpha \cup_\Sigma U_\beta$ with $\partial U_\alpha = -\partial U_\beta = \Sigma$; we assume Σ to be transverse to L;
- a decomposition of $L \cap \Sigma$ into alternating w and z basepoints; $\boldsymbol{z}$ is the set of z basepoints, and $\boldsymbol{w}$ is the union of those basepoints and possibly some additional ones on $\Sigma \setminus L$;
- a good set of trajectories $\boldsymbol{c}$ as in Definition 8.26.

Consider now a complete presystem of hyperboxes $\mathscr{H}$ for $\vec{L} \subset Y$. Note that all the hyperboxes $\mathscr{H}^{\vec{L}',\vec{M}}$ in $\mathscr{H}$ contain Heegaard diagrams with the same underlying Heegaard surface Σ. Let $\boldsymbol{w}^{\vec{L}',\vec{M}}$ (resp. $\boldsymbol{z}^{\vec{L}',\vec{M}}$) be the set of basepoints of type w (resp. z) on the diagrams in the hyperbox $\mathscr{H}^{\vec{L}',\vec{M}}$. In particular, we let $\boldsymbol{w}^{L'} = \boldsymbol{w}^{\vec{L}',\varnothing}$ and $\boldsymbol{z}^{L'} = \boldsymbol{z}^{\vec{L}',\varnothing}$. We also set $\boldsymbol{w} = \boldsymbol{w}^{L}$ and $\boldsymbol{z} = \boldsymbol{z}^{L}$.

Given a sublink $M \subseteq L$ with an orientation $\vec{M}$, we let $M = M_+ \amalg M_-$, where M_+ (resp. M_-) consists of those components of M that have the same (resp. opposite) orientation in $\vec{L}$ as in $\vec{M}$. Suppose M is a sublink of some $L' \subseteq L$ (with the orientation $\vec{L}'$ coming from $\vec{L}$). Because of the compatibility condition (93) (applied to $\vec{L}'$ instead of $\vec{L}$, and taking $M' = \varnothing$), we must have

$$\boldsymbol{w}^{\vec{L}',\vec{M}} = (\boldsymbol{w}^{L'} \setminus M_-) \cup (\boldsymbol{z}^{L'} \cap M_-) \quad \text{and} \quad \boldsymbol{z}^{\vec{L}',\vec{M}} = \boldsymbol{z}^{L'} \cap (L' - M).$$

Thus, the basepoints on all hyperboxes are determined by those on the diagrams $\mathscr{H}^{L'}$ for $L' \subseteq L$. To get further control on the basepoints (and the way they move under surface isotopies), consider a good set of trajectories $\boldsymbol{c} = \{c_w\}$ for the initial diagram $\mathscr{H}^{L}$ (which represents $\vec{L}$). We say that the complete presystem $\mathscr{H}$ is *dependent* on the trajectory set $\boldsymbol{c}$ if the following conditions are satisfied:

- For any $\vec{L}' \subseteq \vec{L}$, we have $\boldsymbol{w}^{L'} = \boldsymbol{w}$ and $\boldsymbol{z}^{L'} = \boldsymbol{z} \cap L'$, so that, for any $M \subset L'$ with an orientation $\vec{M}$,
$$\boldsymbol{w}^{\vec{L}',\vec{M}} = (\boldsymbol{w} \setminus M_-) \cup (\boldsymbol{z} \cap M_-), \quad \boldsymbol{z}^{\vec{L}',\vec{M}} = \boldsymbol{z} \cap (L' - M).$$
- For any $M' \subseteq M \subseteq L' \subseteq L$, and any orientation $\vec{M}$ of M, in the respective compatibility relation (94), which reads
$$\mathscr{H}^{\vec{L}'-M',\vec{M}-M'} \cong \mathscr{H}^{\vec{L}',\vec{M}}(M', M), \tag{95}$$
we use a surface isotopy that moves each basepoint $w \in \boldsymbol{w} \cap M'_-$ (appearing in the diagram on the left-hand side of (95)) to its successor $z(w) \in \boldsymbol{z} \cap M'_-$ (appearing in the diagram on the right-hand side of (95)), exactly tracing the path c_w. Moreover, we require the surface isotopy to be supported in a small neighborhood of the path c_w, and in particular to fix all the basepoints in
$$(\boldsymbol{w} \setminus M_-) \cup (\boldsymbol{z} \cap (M - M')_-) \cup (\boldsymbol{z} \cap (L' - M)),$$
which appear on both sides of (95).

Definition 8.28 A *complete system of hyperboxes* $(\mathscr{H}, \boldsymbol{c})$ representing a link $\vec{L} \subset Y$ consists of a complete presystem of hyperboxes $\mathscr{H}$ for $\vec{L}$ together with a good set of trajectories $\boldsymbol{c}$ for $\mathscr{H}^{L}$ such that $\mathscr{H}$ is dependent on $\boldsymbol{c}$. (We usually drop $\boldsymbol{c}$ from the notation, and refer to $\mathscr{H}$ as a complete system of hyperboxes.)

If $\mathscr{H}$ is a complete system of hyperboxes for $\vec{L}$ with a good set of trajectories $\boldsymbol{c}$, and if $(\Sigma, \boldsymbol{w}, \boldsymbol{z})$ is the multipointed Heegaard surface of the initial diagram $\mathscr{H}^{L}$, we refer to $(\Sigma, \boldsymbol{w}, \boldsymbol{z}, \boldsymbol{c})$ as the *underlying surface data* for $\mathscr{H}$. (Compare Definition 8.27.)

Definition 8.29 A complete system of hyperboxes $(\mathscr{H}, \boldsymbol{c})$ for $\vec{L} \subset Y$ is called *link-minimal* if $\mathscr{H}^{L}$ (and hence all the Heegaard diagrams in $\mathscr{H}$) are link-minimal, in the sense of Definition 3.2.

Definition 8.30 A complete system of hyperboxes $(\mathcal{H}, \boldsymbol{c})$ for $\vec{L} \subset Y$ is called *minimally pointed* if the diagram $\mathcal{H}^{L}$ is minimally pointed, in the sense of Definition 3.3. (Note that this does not imply that the other diagrams in the system are minimally pointed.)

Definition 8.31 A complete system of hyperboxes $(\mathcal{H}, \boldsymbol{c})$ for $\vec{L} \subset Y$ is called *of pure β-type* if all its hyperboxes are pure β-hyperboxes, ie the corresponding bipartition maps r_i take all the indices to β. (Hence, in such a system, the alpha curves always stay fixed, and we have $\varepsilon = \varepsilon^{\beta}$ for any multi-index ε.)

8.7 Basic systems

Let $\vec{L} \subset Y$ be a link. In this section we describe a special kind of minimally pointed complete system of pure β-type for $\vec{L}$, which we call *basic*.

In a basic complete system, all hyperboxes of the form $\mathcal{H}^{\vec{L}',\vec{M}}$ will be trivial, ie of size $(0,\dots,0)$, when $\vec{M}$ has the orientation induced from $\vec{L}$ (or, equivalently, from $\vec{L}'$). In particular, all diagrams $\mathcal{H}^{L'} = \mathcal{H}^{L',\varnothing}$ will simply be obtained from an initial diagram, $\mathcal{H}^{L}$, by deleting the z basepoints on the components of $L - L'$.

Let ℓ be the number of components of L. We choose the Heegaard diagram $\mathcal{H}^{L}$ to be basic (see Definition 3.4), that is, of genus g, with $g+\ell-1$ alpha curves and $g+\ell-1$ beta curves, ℓ basepoints marked w, ℓ basepoints marked z, and such that the basepoints w_i and z_i for $i = 1\dots,\ell$ lie on each side of a beta curve β_i, and are not separated by any alpha curves. Thus, the beta curves split the surface Σ into ℓ components $\Sigma_1,\dots,\Sigma_\ell$, numbered such that Σ_i contains the basepoints w_i and z_i, and has both sides of β_i as parts of its boundary. We denote by L_i the component of L on which w_i and z_i lie. We construct a good set of trajectories $\boldsymbol{c} = \{c_{w_i}\}$ for $\mathcal{H}^{L}$ to consist of small paths c_{w_i} from w_i to z_i that intersect β_i once and do not intersect any of the other curves. We will take $(\Sigma, \{w_i\}, \{z_i\}, \boldsymbol{c})$ to be the underlying surface data for our complete system.

Let β_i'' be the curve obtained from β_i by a small isotopy pushing w_i into z_i, following the path c_{w_i}, as in Figure 22. Note that, if we ignore w_i, then β_i'' can also be obtained from β_i by handleslides supported in Σ_i away from z_i, namely by handlesliding β_i over all the other curves on the boundary of Σ_i; compare [Ozsváth and Szabó 2004b, proof of Proposition 7.1].

We would like to construct a complete system using the curves β_i and β_i''. Of course, if the isotopy from w_i to z_i is supported just in a small neighborhood of c_{w_i}, then the curves β_i and β_i'' are not transverse (they have a whole arc in common away from that neighborhood), so cannot be part of the same Heegaard diagram. This issue persists even if we slightly isotope β_i'' further so that its intersection with β_i is transverse and consists of two points, as in Figure 22. Indeed, let B_i^1 be the bigon between β_i and β_i'' containing z_i and B_i^2 the other bigon. Then $\Sigma_i - B_i^1 + B_i^2$ is a periodic domain in Σ with only nonnegative multiplicities, and with zero multiplicity at the basepoint z_i. This is a source of problems if one tries to construct a Heegaard diagram using the curves β_i and β_i'' and the basepoint z_i, because such diagrams

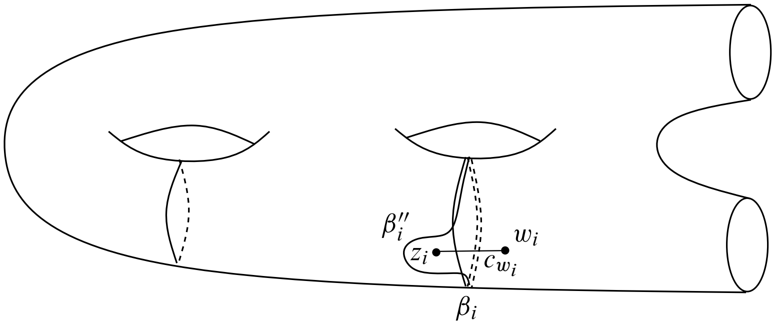

Figure 22: The curves β_i and β_i''. We show here a component Σ_i of the complement of the beta curves in Σ.

are not admissible. In order to fix this problem, we introduce an intermediate curve β_i' as in Figure 23. Then there are no periodic domains as above between β_i and β_i', nor between β_i' and β_i''.

Let us now describe the hyperbox $\mathscr{H}^{\vec{L},-\vec{L}}$, which is the biggest hyperbox in our basic system. As mentioned above, the hyperbox $\mathscr{H}^{\vec{L},-\vec{L}}$ is a pure β-hyperbox. Its punctures are z_i for $i = 1, \dots, \ell$, but playing the role of w's in the definition of a multipointed Heegaard diagram. The hyperbox has dimension ℓ and size $\boldsymbol{d} = (2, 2, \dots, 2)$. For $\varepsilon \in \mathbb{E}(\boldsymbol{d}) = \{0, 1, 2\}^{\ell}$, the collection of curves $\boldsymbol{\beta}^{\varepsilon}$ is given by an approximation

$$\beta_i^{\varepsilon} \approx \begin{cases} \beta_i & \text{if } \varepsilon_i = 0, \\ \beta_i' & \text{if } \varepsilon_i = 1, \\ \beta_i'' & \text{if } \varepsilon_i = 2, \end{cases}$$

for $i \le \ell$, and $\beta_i^{\varepsilon} \approx \beta_i$ for $i > \ell$. We also arrange so that $\beta_i^{\varepsilon} \approx \beta_i^{\varepsilon'}$ for any $\varepsilon \ne \varepsilon'$ with $\varepsilon_i = \varepsilon_i'$.

Note that with this choice of $\boldsymbol{\beta}^{\varepsilon}$, when two multi-indices ε and ε' are neighbors, we never see a pair of curves that approximate β_i and β_i'' (for the same i) in the Heegaard diagram $(\Sigma, \boldsymbol{\beta}^{\varepsilon}, \boldsymbol{\beta}^{\varepsilon'}, \boldsymbol{w}, \boldsymbol{z})$. Hence, the admissibility hypothesis in the definition of the β-hyperbox is satisfied. Further, there is a natural choice for the chain $\Theta_{\varepsilon,\varepsilon'}$ when $(\varepsilon, \varepsilon')$ is an edge in the hyperbox: namely, the respective intersection points of maximal degree. For example, if in the Heegaard diagram $(\Sigma, \boldsymbol{\beta}^{\varepsilon}, \boldsymbol{\beta}^{\varepsilon'}, \boldsymbol{w}, \boldsymbol{z})$ we see the curves β_i and β_i', then the point marked θ_i in Figure 23 would be part of the corresponding Θ-chain. When $\varepsilon < \varepsilon'$ are neighbors but $\|\varepsilon' - \varepsilon\| \ge 2$, we set $\Theta_{\varepsilon,\varepsilon'} = 0$.

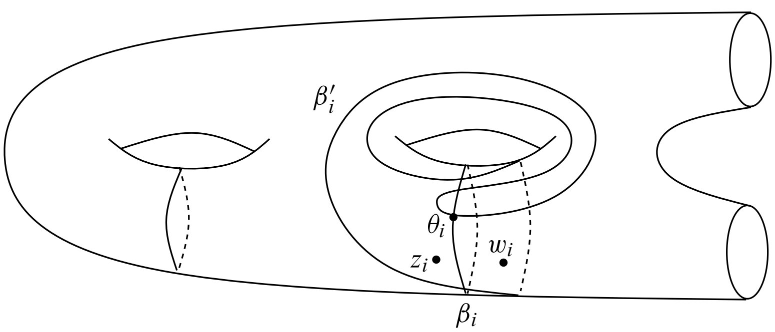

Figure 23: The curves β_i and β_i'. The diagram formed by them is admissible.

The fact that the Θ-chains satisfy the compatibility relations (90) has a proof similar to those of [Ozsváth and Szabó 2004b, Lemma 9.7; 2005, Lemma 4.3].

For future reference, we denote the beta curves in the $\mathscr{H}^{\vec{L},-\vec{L}}$ hyperbox by

$$\widetilde{\beta}_i^{\varepsilon} = \beta_i^{\varepsilon}, \qquad \varepsilon \in \{0,1,2\}^{\ell}, \quad i \in \{1,\dots,g+\ell-1\}.$$

Let us now describe an arbitrary hyperbox $\mathscr{H}^{\vec{L}',\vec{M}}$ in the basic system. Let $M = M_+ \amalg M_-$ as in Section 8.6, where M_+ (resp. M_-) consists of those components of M that have the same (resp. opposite) orientation in $\vec{L}$ as in $\vec{M}$. Denote by m (resp. m_+, m_-) the number of components in M (resp. M_+, M_-). Order the components of M according to their ordering as components of L:

$$M = L_{i_1} \cup \cdots \cup L_{i_m}, \quad i_1 < \cdots < i_m.$$

The hyperbox $\mathscr{H}^{\vec{L}',\vec{M}}$ has size $\boldsymbol{d}^{\vec{M}} = (d_1^{\vec{M}}, \dots, d_m^{\vec{M}})$, where

$$d_j^{\vec{M}} = \begin{cases} 0 & \text{if } L_{i_j} \subseteq M_+, \\ 2 & \text{if } L_{i_j} \subseteq M_-. \end{cases}$$

Its diagrams all lie on the surface Σ with punctures $\boldsymbol{w}^{\vec{L}',\vec{M}}$ and $\boldsymbol{z}^{\vec{L}',\vec{M}}$ given by

$$\boldsymbol{w}^{\vec{L}',\vec{M}} = \{w_i \mid L_i \subseteq L - M_-\} \cup \{z_i \mid L_i \subseteq M_-\} \quad \text{and} \quad \boldsymbol{z}^{\vec{L}',\vec{M}} = \{z_i \mid L_i \subseteq L' - M\}.$$

Consider the injective map

$$\lambda^{\vec{M}} : \mathbb{E}(\boldsymbol{d}^{\vec{M}}) \to \mathbb{E}(\boldsymbol{d}^{-\vec{L}}) = \{0,1,2\}^{\ell},$$

given by

$$(\lambda^{\vec{M}}(\varepsilon_1,\dots,\varepsilon_m))_i = \begin{cases} \varepsilon_j & \text{if } i = i_j \text{ for some } j, \\ 0 & \text{otherwise.} \end{cases}$$

For the beta curves β_i^{ε} in the hyperbox $\mathscr{H}^{\vec{L}',\vec{M}}$, we choose

$$\beta_i^{\varepsilon} = \widetilde{\beta}_i^{\lambda^{\vec{M}}(\varepsilon)}, \quad \varepsilon \in \mathbb{E}(\boldsymbol{d}^{\vec{M}}),\ i \in \{1,\dots,g+\ell-1\}. \tag{96}$$

Finally, we let the Θ-cycles on the edges be the respective intersection points of maximal degree, and we let the rest of the Θ-chain elements be zero.

This completes the description of the hyperboxes $\mathscr{H}^{\vec{L}',\vec{M}}$. (Note that, when $M_- = \varnothing$, we indeed get a trivial hyperbox consisting of the initial Heegaard diagram $\mathscr{H}^{L}$ with the points z_i on $(L - L') \cup M$ removed.) The verification of the conditions in the definition of a complete system is an easy exercise.

Definition 8.32 A complete system of hyperboxes $\mathscr{H}$ representing the link $\vec{L}$ is called *basic* if it is constructed as above. If this is the case, we refer to $\mathscr{H}$ more simply as a *basic system* for the link $\vec{L}$.

To review, every basic system for $\vec{L}$ is associated to a particular 2ℓ-pointed Heegaard diagram

$$(\Sigma, \alpha_1, \dots, \alpha_{g+\ell-1}, \beta_1, \dots, \beta_{g+\ell-1}, w_1, \dots, w_{\ell}; z_1, \dots, z_{\ell})$$

with the property that, for each $i = 1, \dots, \ell$, the basepoints w_i and z_i lie on each side of the curve β_i and are not separated by any alpha curves. In addition, the construction of the basic system requires several choices (the curves β_i' and β_i'' and their approximations). For simplicity, we typically forget those choices and just say that the basic system is associated to a special Heegaard diagram as above.

Remark 8.33 Another complete system of hyperboxes for a link is described in Section 15.6, using grid diagrams.

8.8 A more general construction

While Section 8.7 was focused on basic systems, the same ideas can be used to construct complete systems from any Heegaard diagram and good set of trajectories.

Proposition 8.34 *Let $\mathscr{H}^L = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ be a multipointed Heegaard diagram for a link $\vec{L} \subset Y$, and let $\boldsymbol{c}$ be a good set of trajectories for $\mathscr{H}^L$. Then there exists a complete system of hyperboxes $\mathscr{H}$ which is of pure β-type, is dependent on $\boldsymbol{c}$, and has $\mathscr{H}^L$ as its initial diagram.*

Proof We will construct $\mathscr{H}$ so that its hyperboxes $\mathscr{H}^{\vec{L}',\vec{M}}$ are of the same size as the ones defined for a basic system in Section 8.7. Also, since $\mathscr{H}$ is meant to be of pure β-type, the collections of alpha curves will be fixed to be the same collection $\boldsymbol{\alpha}$ from $\mathscr{H}^L$, in all the hyperboxes.

We begin by describing the biggest hyperbox $\mathscr{H}^{\vec{L},-\vec{L}}$, which is of size $\boldsymbol{d} = (2, 2, \dots, 2)$. Recall that each trajectory c_w in $\boldsymbol{c}$ gives a surface isotopy supported in a neighborhood of c_w, taking w to its successor $z(w)$. For $\varepsilon \in \mathbb{E}(\boldsymbol{d}) = \{0, 1, 2\}^\ell$, let τ_ε be the composition of all surface isotopies corresponding to c_w for basepoints w on link components L_i with $\varepsilon_i \neq 0$.

If none of the values ε_i are 1, we choose the collection of beta curves $\boldsymbol{\beta}^\varepsilon$ to be

$$\boldsymbol{\beta}^\varepsilon := \tau_\varepsilon(\boldsymbol{\beta}).$$

Thus, when some value ε_i changes from 0 to 2, the beta curves change by an isotopy moving some basepoints, similarly to how β_i'' was obtained from β_i in Section 8.7.

We construct the remaining collections $\boldsymbol{\beta}^\varepsilon$ by induction on the number of values 1 among the ε_i. Let $\bar{\varepsilon} \in \{0, 2\}^\ell$ be obtained from ε by changing all values 1 among ε_i to 2. Then we define $\boldsymbol{\beta}^\varepsilon$ by winding the curves $\boldsymbol{\beta}^{\bar{\varepsilon}} = \tau_\varepsilon(\boldsymbol{\beta})$ sufficiently to achieve admissibility with the already defined collections of beta curves, in each hypercube. This is done as in [Ozsváth and Szabó 2004b, Sections 5 and 8.4.2], and is similar to how we obtained the curve β_i' from β_i'' in Section 8.7.

After the beta curves are constructed, we fill in the hyperboxes in $\mathscr{H}^{\vec{L},-\vec{L}'}$ with Θ-chain elements using Lemma 8.6. We then define the other hyperboxes $\mathscr{H}^{\vec{L}',\vec{M}}$ based on what we have in $\mathscr{H}^{\vec{L},-\vec{L}'}$, using the same formula (96) for the beta curves as in the basic case. We also let the Θ-chain elements to be induced from those in $\mathscr{H}^{\vec{L},-\vec{L}'}$: the element for the pair $(\varepsilon, \varepsilon')$ in $\mathscr{H}^{\vec{L}',\vec{M}}$ should be the same as the one for $(\lambda^{\vec{M}}(\varepsilon), \lambda^{\vec{M}}(\varepsilon'))$ in $\mathscr{H}^{\vec{L},-\vec{L}'}$. □

8.9 Moves on complete systems of hyperboxes

In Section 8.4 we defined several moves on hyperboxes of strongly equivalent Heegaard diagrams. We can define an analogous list of *system Heegaard moves* between complete systems of hyperboxes.

First note that *three-manifold isotopies*, *index one/two stabilizations*, *free index zero/three stabilizations*, *global shifts* (and all their inverses) have straightforward extensions to complete systems of hyperboxes. A certain move of one of these types on a complete system consists of applying that type of move to all the hyperboxes in the system, in a way compatible with restrictions. The good set of trajectories should be taken into the corresponding one by the respective moves. The (index one/two, or free index zero/three) stabilizations are required to be done away from the trajectories c_w in the good set, so that these trajectories are preserved.

The three other system Heegaard moves in the list are elementary enlargements / contractions, index zero/three link stabilizations (and destabilizations), and trajectory α-slides, which require more discussion.

An *elementary enlargement* of a complete system $\mathscr{H}$ consists of picking a component M_0 of L, with an orientation $\vec{M}_0$, and doing compatible elementary enlargements of those hyperboxes $\mathscr{H}^{\vec{L}',\vec{M}}$ with $\vec{M}_0 \subseteq \vec{M}$ (and the orientation on $\vec{M}_0$ is the one induced from $\vec{M}$); these elementary enlargements are all done along the coordinate axis corresponding to M_0. The good set of trajectories is unchanged. An *elementary contraction* is the inverse of an elementary enlargement.

An *index zero/three link stabilization* of a complete system $\mathscr{H}$ is as follows. We do an index zero/three link stabilization of the initial diagram $\mathscr{H}^L$ resulting in two additional basepoints w' and z' near an old basepoint z on a component $L_i \subseteq L$; see Figure 1. Let Σ' be the new Heegaard surface, and $\boldsymbol{w}'$ and $\boldsymbol{z}'$ the new collections of basepoints. We let $c_{w'}$ be a short trajectory (inside the disk bounded by α') joining w' to z'. Adding $c_{w'}$ to the good set of trajectories $\boldsymbol{c}$ for $\mathscr{H}$, we obtain a new trajectory set $\boldsymbol{c}'$. We take $(\Sigma, \boldsymbol{w}', \boldsymbol{z}', \boldsymbol{c}')$ to be the underlying surface data for the new stabilized complete system $\mathscr{H}'$, which we now describe. Let $M \subseteq L' \subseteq L$ be sublinks, with $M = M_+ \amalg M_-$ having an orientation $\vec{M}$ (that coincides with the orientation of $\vec{L}$ exactly on M_+). To get from the hyperboxes in $\mathscr{H}$ to the corresponding hyperboxes in $\mathscr{H}'$, we do the following.

If $L_i \subseteq L' - M_-$, we change the hyperbox $\mathscr{H}^{\vec{L}',\vec{M}}$ by an index zero/three link stabilization performed in the same place as in $\mathscr{H}^L$.

If $L_i \subseteq L - L'$, note that the z and z' basepoints should disappear. Therefore, we change $\mathscr{H}^{\vec{L}',\vec{M}}$ by a free index zero/three stabilization at the same location (followed by a three-manifold isotopy to get from Σ to Σ').

If $L_i \subseteq M_-$, note that the w' basepoint disappears, and we relabel z and z' as w and w', respectively. We first construct a hyperbox $\breve{\mathscr{H}}^{\vec{L}',\vec{M}}$ from $\mathscr{H}^{\vec{L}',\vec{M}}$ by taking (in each Heegaard diagram) the respective connected sum with a sphere, and adding the two new curves α' and β' encircling w and w', as in Figure 24. The corresponding hyperbox $(\mathscr{H}')^{\vec{L}',\vec{M}}$ in the new system $\mathscr{H}'$ is obtained from $\breve{\mathscr{H}}^{\vec{L}',\vec{M}}$ by increasing the

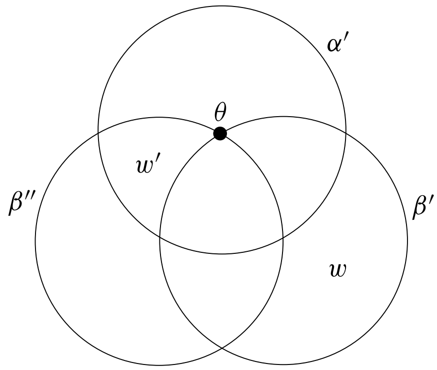

Figure 24: Index zero/three link stabilizations of complete systems. Compared to the right-hand side of Figure 1, we deleted w' and relabeled z and z' as w and w', respectively. The curve β'' can be obtained from β' by handlesliding it over other beta curves.

length of its side in the L_i direction by one, and adding on (at the end of each segment in the L_i direction) new diagrams in which the curve β' is replaced by β'' as in Figure 24. (Observe that replacing β' with β'' is the strong equivalence discussed in Section 7.5, with different notation.) The Θ-chain elements on the new edges in the L_i direction contain the point θ in the figure. After this modification, the subhyperbox $(\mathscr{H}')^{\vec{L}',\vec{M}}(L_i, M)$ is surface isotopic to the free index zero/three stabilization of $\mathscr{H}^{\vec{L}'-L_i,\vec{M}-L_i}$, as it should be according to the compatibility relation (94). Note that the change from $\breve{\mathscr{H}}^{\vec{L}',\vec{M}}$ to $(\mathscr{H}')^{\vec{L}',\vec{M}}$ can be realized as the composition of an elementary enlargement and a global shift (of hyperboxes). We do these moves in a compatible way on all hyperboxes corresponding to pairs $(\vec{L}', \vec{M})$ with $L_i \subseteq M_-$. This completes the description of the index zero/three link stabilization of $\mathscr{H}$.

Finally, we have one more system Heegaard move, the *trajectory α-slide*, which did not have an analogue in the list of moves on hyperboxes. This involves changing the good set of trajectories $\boldsymbol{c}$. It is a move that we only perform on complete systems of pure β-type, ie those where the alpha curves are fixed; see Definition 8.31. Suppose $\mathscr{H}$ is such a complete system, and let τ be an embedded arc connecting a basepoint z_i to an alpha curve α_j, and not intersecting any alpha curves in its interior. Observe that the trajectory c_{w_i} does not intersect alpha curves either — otherwise these would be moved by the isotopy along c_{w_i}, but we assume them to be fixed. On the other hand, both c_{w_i} and τ may intersect various beta curves, from different diagrams in the complete system. See Figure 25, top left.

The trajectory α-slide is a move that transforms $\mathscr{H}$ into a new complete system $\mathscr{H}'$, where the trajectory c_{w_i} is replaced by c'_{w_i}, the result of a handleslide of c_{w_i} over α_j along the arc τ, as in Figure 25, top right. To motivate the description of $\mathscr{H}'$, consider the relations

$$\mathscr{H}^{\vec{L}'-M',\vec{M}-M'} \cong \mathscr{H}^{\vec{L}',\vec{M}}(M', M)$$

from (94), where the surface isotopy goes along c_{w_i}; this happens when $w_i \in M'_-$. Let us denote this surface isotopy by ψ. Then, if on the left-hand side of the above equation we have a diagram with

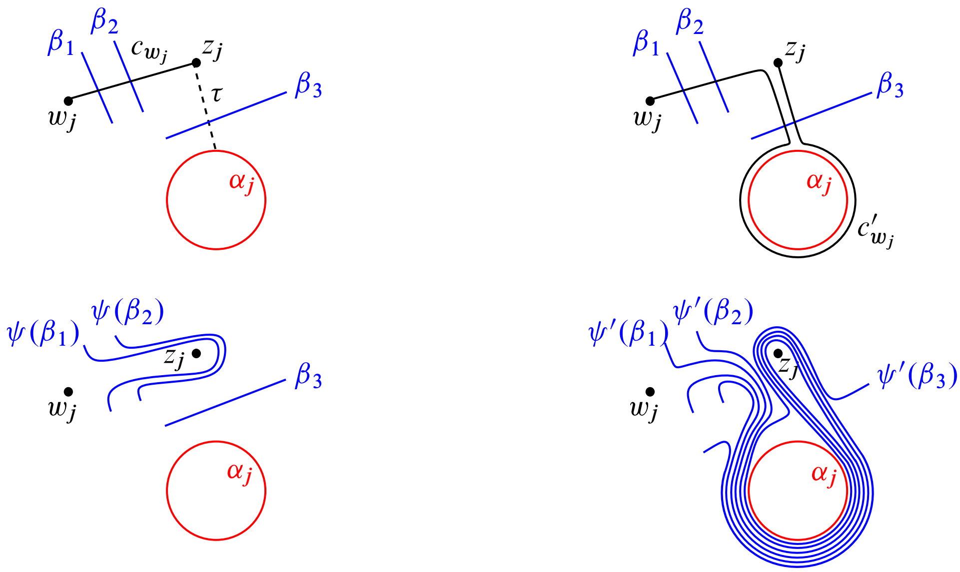

Figure 25: A trajectory α-slide.

curves β_k, on the right-hand side we will have the curves $\psi(\beta_k)$, as in Figure 25, bottom left. Now, in the new complete system $\mathscr{H}'$, we would like to have on the right-hand side the curves $\psi'(\beta_k)$ instead, where ψ' is the isotopy following c'_{w_i}, as in Figure 25, bottom right. To achieve this, given any hyperbox $\mathscr{H}^{\vec{L}',\vec{M}}$ with $w_i \in M_-$, we increase its side length corresponding to the component L_i (the one containing w_i) by 2, by letting $\psi(\beta_k)$ be followed by a new curve β'_k and then by $\psi'(\beta_k)$; the intermediate curves β'_k are chosen by twisting β_k sufficiently to ensure admissibility. We then fill in the new parts of the hyperboxes using the procedure from Lemma 8.6. The result is a new complete system of pure β-type, which is our desired $\mathscr{H}'$.

With the list of system Heegaard moves in place, we can now state the analogue of Lemma 8.18 in this context:

Proposition 8.35 *Let $\mathscr{H}$ and $\mathscr{H}'$ be two complete systems of hyperboxes representing the same oriented link $\vec{L} \subset Y$.*

(a) *The complete systems $\mathscr{H}$ and $\mathscr{H}'$ can be related by a sequence of system Heegaard moves.*

(b) *Suppose that $\mathscr{H}$ is a link-minimal system and that the diagram $(\mathscr{H}')^L$ in $\mathscr{H}'$ has m points of type z, and hence $m-\ell$ subsidiary w basepoints. Then $\mathscr{H}'$ can be obtained from $\mathscr{H}$ by a sequence of system Heegaard moves that includes exactly $m-\ell$ index zero/three link stabilizations, with each of these stabilizations introducing a subsidiary basepoint.*

Recall that each complete system $\mathscr{H}$ has some underlying surface data $S = (\Sigma, \boldsymbol{w}, \boldsymbol{z}, \boldsymbol{c})$; compare Definition 8.27. To prove Proposition 8.35, we will relate the complete systems by first relating their underlying surface data. With the exceptions of global shifts and elementary enlargements/contractions (which do not change the surface data), each system move has a corresponding move at the level of the surface data itself. Thus, we distinguish the following types of *moves of surface data*: three-manifold isotopies, index one/two stabilizations, index zero/three stabilizations (free or linked), and trajectory α-slides. (In the case of trajectory α-slides, to define these without reference to alpha curves, we simply require one of the trajectories in $\boldsymbol{c}$ to slide over a simple closed curve that bounds a disk in the handlebody U_α.)

Lemma 8.36 *Let S and S' be two sets of surface data for the same oriented link ℓ-component $\vec{L} \subset Y$.*

(a) *The sets S and S' can be related by a sequence of moves of surface data.*

(b) *Suppose that S is link-minimal (that is, it contains exactly ℓ basepoints of type z), and that S' has m points of type z, and hence $m-\ell$ subsidiary w basepoints. Then S' can be obtained from S by a sequence of moves of surface data that includes exactly $m-\ell$ index zero/three link stabilizations.*

Proof Let us first consider the case when both $S = (\Sigma, \boldsymbol{w}, \boldsymbol{z}, \boldsymbol{c})$ and $S' = (\Sigma', \boldsymbol{w}', \boldsymbol{z}', \boldsymbol{c}')$ are minimally pointed; that is, they each have only ℓ basepoints of type w and ℓ basepoints of type z.

In this case, after a three-manifold isotopy, we can assume S and S' have the same basepoints $\boldsymbol{w} = \boldsymbol{w}'$ and $\boldsymbol{z} = \boldsymbol{z}'$, and in fact that the corresponding surfaces Σ and Σ' agree in a neighborhood of each basepoint (intersecting that neighborhood in the same small disk). Moreover, we can make the trajectories c_w agree with c'_w in small neighborhoods of the basepoints.

After this, to relate S and S', we use a Morse-theoretic argument. We fix a Riemannian metric on Y. For each component L_i of L (with basepoints w_i and z_i), we also fix a standard self-indexing Morse function f_i in a small tubular neighborhood $\nu(L_i)$ of L_i. We ask for f_i to have exactly two critical points, both on L_i, one of index 0 and one of index 3, such that L_i consists of two gradient flow trajectories between these two critical points, and

$$f_i^{-1}\left(\tfrac{3}{2}\right) = \Sigma \cap \nu(L_i) = \Sigma' \cap \nu(L_i),$$

so that the alpha handlebody (from either S or S') intersects $\nu(L_i)$ in $f_i^{-1}\left[0, \tfrac{3}{2}\right]$, and the beta handlebody intersects $\nu(L_i)$ in $f_i^{-1}\left[\tfrac{3}{2}, 3\right]$.

Furthermore, we can assume that $c_{w_i} \cap \nu(L_i) = c'_{w_i} \cap \nu(L_i)$ consists of two small arcs on Σ: an arc a_i from w_i to some point w_i^* on $\partial\nu(L_i)$, and another arc b_i from some point z_i^* on $\partial\nu(L_i)$ to z_i. We let $c_i \subset c_{w_i}$ be the remaining arc from w_i^* to z_i^* along the trajectory c_{w_i}, after we remove a_i and b_i. We define $c'_i \subset c'_{w_i}$ similarly.

We choose a pushoff γ_i of the arc $L_i \cap f_i^{-1}\left[0, \tfrac{3}{2}\right]$, so that γ_i goes along the boundary $\partial\nu(L_i)$, from w_i^* to z_i^*. See Figure 26.

We extend the functions f_i to a single self-indexing Morse–Smale function f on all of Y, by introducing new index 1 and 2 critical points (but no additional index 0 or 3 critical points) so that $\Sigma = f^{-1}\left(\tfrac{3}{2}\right)$. We

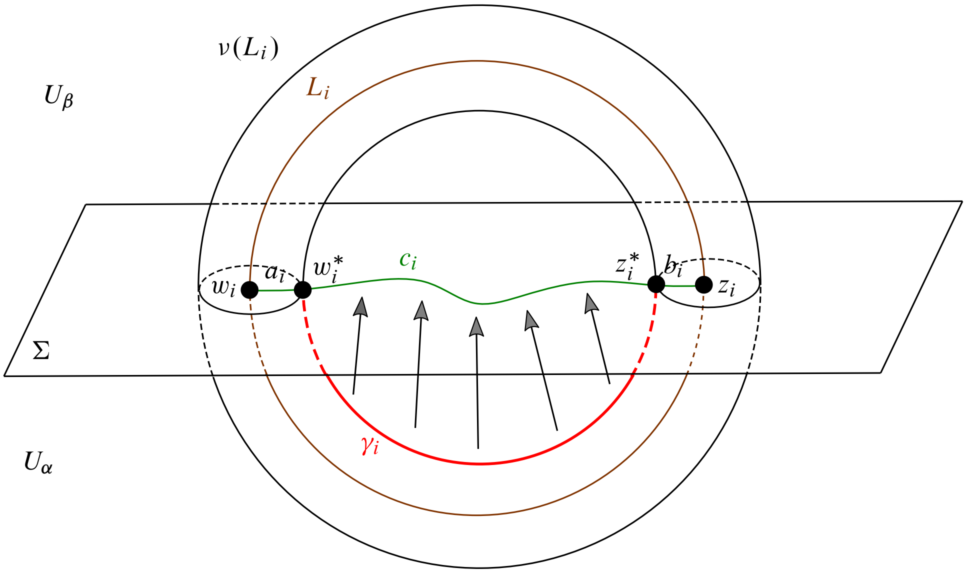

Figure 26: Constructing a Morse–Smale function. We show here a neighborhood of a link component L_i, together with the path c_i. The arrows represent gradient flow lines taking γ_i into c_i.

ask that the arc γ_i is taken by the forward gradient flow to the arc c_i on Σ, from w_i^* to z_i^*. To see that we can choose f this way, note that, by assumption, γ_i is homotopic to c_i inside the alpha handlebody, and therefore, by Dehn's lemma, the union of γ_i and c_i bounds an embedded disk D_i in that handlebody. We extend the functions f_i from $\nu(L_i)$ to a neighborhood of that disk, so that the flow takes γ_i to c_i, then extend them to the rest of the alpha handlebody — so that on Σ it takes the value $\frac{3}{2}$, and finally to the beta handlebody.

Next, we apply the same procedure to Σ', obtaining a Morse–Smale function g with $\Sigma' = g^{-1}\left(\frac{3}{2}\right)$ such that c_i' is the image of γ_i under the forward gradient flow.

We now relate the two functions f and g by a generic smooth homotopy $(F_t)_{0\le t\le 1}$, with $F_0 = f$, $F_1 = g$, keeping F_t fixed in the neighborhoods $\nu(L_i)$. Following the level set at $\frac{3}{2}$ along this homotopy, we can relate the surfaces Σ and Σ' by three-manifold isotopies and index 1/2 stabilizations and destabilizations; compare [Ozsváth and Szabó 2004b, Proposition 2.2] or [Juhász 2006, Proposition 2.15]. As for the trajectories c_{w_i} and c'_{w_i}, they remain fixed near the basepoints throughout the homotopy, while the arc c_i changes into c_i' by following the images of γ_i through the homotopy. In this generic one-parameter family of functions, at finitely many points it may happen that the forward flow takes γ_i into an index 1 critical point. Then c_i changes not by an isotopy on Σ, but rather by a trajectory α-slide (over the intersection of Σ with the ascending manifold of the index 1 critical point).

Thus, we have succeeded in relating any two minimally pointed sets of surface data. If we have arbitrary surface data S', we can find some minimally pointed surface data S and a sequence of three-manifold isotopies and index zero/three (free and linked) stabilizations, that take S into S'. Indeed, we can position these (de)stabilizations around the trajectories c_w and the free basepoints that we want to eliminate from S', until we get to the minimally pointed setting. This completes the proof of part (a). For part (b), note that in the above process we used the minimal number of index zero/three link stabilizations. $\square$

Proof of Proposition 8.35 First, let us consider the case when $\mathscr{H}$ and $\mathscr{H}'$ have the same underlying surface data $S = (\Sigma, \boldsymbol{w}, \boldsymbol{z}, \boldsymbol{c})$. If so, elementary enlargements and contractions can be used to make the sizes (and bipartition maps) of all corresponding hyperboxes to agree. Then we use global shifts to relate the resulting complete systems, as we did in Lemma 8.18 for hyperboxes.

In view of the above observation, instead of relating two complete systems $\mathscr{H}$ and $\mathscr{H}'$ by moves, it suffices to consider their surface data S and S', use Lemma 8.36 to relate these by a sequence of intermediate data

$$S = S_0,\ S_1,\ S_2,\ \dots,\ S_n = S' \tag{97}$$

and, for each $k = 0, \dots, n-1$, find *some* complete system $\mathscr{H}_k$ with data S_k, related by some system move to a complete system $\mathscr{H}'_k$ with data S_{k+1}. We would then be able to relate $\mathscr{H}$ to $\mathscr{H}_0$ (because they have surface data S_0), each $\mathscr{H}_k$ to $\mathscr{H}'_k$ and to $\mathscr{H}_{k+1}$ (as the latter two have the same surface data S_{k+1}), and finally $\mathscr{H}'_{n-1}$ with $\mathscr{H}'$. The combination of these system moves will enable us to get from $\mathscr{H}$ to $\mathscr{H}'$.

It remains to see that every move of surface data can be upgraded to a corresponding move of complete systems. We have already seen in Proposition 8.34 that for any surface data, we can find a complete system of pure β-type having that data. If our move is a three-manifold isotopy, we just follow the same isotopy to change the complete system. If our move is a stabilization (of any kind), let S be the data before the stabilization and S' the one after. Up to isotopy, we can assume that the stabilization happens in a neighborhood of some point p, and therefore we can choose a complete system $\mathscr{H}$, with underlying data S such that we can apply the same stabilization to it (ie the curves involved in the diagrams in $\mathscr{H}$ do not go through p); the result is a system $\mathscr{H}'$ for S'.

Finally, if our move is a trajectory α-slide over some curve α_0, we pick a complete system $\mathscr{H}$ of pure β-type that contains α_0 as one of its alpha curves. (This can be done because in the proof of Proposition 8.34, there was a lot of freedom in the choice of alpha curves: we just wanted them to represent the given handlebody, and we can choose such a set by starting with any curve bounding a disk in the handlebody.) Then we let the new complete system $\mathscr{H}'$ be obtained from $\mathscr{H}$ by the given trajectory α-slide. $\square$

9 Statement of the surgery theorem for link-minimal complete systems

We are now ready to state a version of Theorem 1.1 for the case of link-minimal complete systems. We let $\vec{L} = L_1 \amalg \cdots \amalg L_\ell$ be an oriented link in an integer homology sphere Y. We suppose we are given a link-minimal complete system of hyperboxes $\mathscr{H}$ for $\vec{L}$, as in Definition 8.29. In particular, $\mathscr{H}^L$ is a Heegaard diagram for $\vec{L}$. We let k be the number of w basepoints in $\mathscr{H}^L$ (or in any other diagram in $\mathscr{H}$), so that $\mathscr{H}^L$ has $k - \ell$ free basepoints.

Since all the diagrams we work with are link-minimal, we are free to use the description of generalized Floer complexes $\boldsymbol{A}^-(\cdot, \boldsymbol{s})$ as free complexes $\mathfrak{A}^-(\cdot, \boldsymbol{s})$; see Section 3.6.

9.1 Descent to a sublink

Let $M \subseteq L$ be a sublink, with an orientation $\vec{M}$, which can be different from the one coming from $\vec{L}$, as in Section 3.7.

Recall that in Section 3.7 we defined projection maps

$$p^{\vec{M}}: \overline{\mathbb{H}}(L) \to \overline{\mathbb{H}}(L),$$

as well as reduction maps

$$\psi^{\vec{M}}: \overline{\mathbb{H}}(L) \to \overline{\mathbb{H}}(L-M).$$

Fix $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$. Since $\psi^{\vec{M}}$ depends only on the components of L that are not on M, we have

$$\psi^{\vec{M}}(p^{\vec{M}}(\boldsymbol{s})) = \psi^{\vec{M}}(\boldsymbol{s}).$$

Consider the m-dimensional hyperbox $\mathscr{H}^{\vec{L},\vec{M}}$ from the complete system. This hyperbox consists of diagrams representing the link $L-M$. By the definition of the complete system, $\mathscr{H}^{\vec{L},\vec{M}}$ has to be compatible with $\mathscr{H}^{\vec{L},\varnothing} = \mathscr{H}^L$, ie the initial diagram in $\mathscr{H}^{\vec{L},\vec{M}}$ is obtained from $\mathscr{H}^L$ by deleting the basepoints z_i on components L_i with $i \in I_+(\vec{L}, \vec{M})$, deleting w_i on the components L_i with $i \in I_-(\vec{L}, \vec{M})$, and also relabeling z_i as w_i for the components L_i with $i \in I_-(\vec{L}, \vec{M})$. In view of (22), we have an identification

$$\mathfrak{A}^-(\mathscr{H}^L, p^{\vec{M}}(\boldsymbol{s})) \cong \mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}(\varnothing), \psi^{\vec{M}}(\boldsymbol{s})). \tag{98}$$

Let $\boldsymbol{d}^{\vec{M}} \in (\mathbb{Z}_{\geq 0})^m$ be the size of $\mathscr{H}^{\vec{L},\vec{M}}$. We have an associated hyperbox of generalized Floer complexes

$$\mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s})) = ((C^{\varepsilon}_{\boldsymbol{s}})_{\varepsilon \in \mathbb{E}(\boldsymbol{d}^{\vec{M}})}, (D^{\varepsilon}_{\boldsymbol{s}})_{\varepsilon \in \mathbb{E}_m})$$

as defined in Section 8.3. By compressing the hyperbox $(C^{\varepsilon}_{\boldsymbol{s}}, D^{\varepsilon}_{\boldsymbol{s}})$ we obtain a hypercube $(\hat{C}^{\varepsilon}_{\boldsymbol{s}}, \hat{D}^{\varepsilon}_{\boldsymbol{s}})_{\varepsilon \in \mathbb{E}_m}$; see Sections 5.2 and 5.6. We are only interested in the longest diagonal map in this hypercube, namely $\hat{D}^{(1,\dots,1)}_{\boldsymbol{s}}$, which we simply denote by $\hat{D}^{\vec{M}}_{\boldsymbol{s}}$. According to (59),

$$\hat{D}^{\vec{M}}_{\boldsymbol{s}} = \mathrm{pl}^{\boldsymbol{d}^{\vec{M}}}_{\{D^{\varepsilon}_{\boldsymbol{s}}\}}(\alpha_m),$$

where α_m is the m^{th} standard symphony and pl denotes the operation of playing songs; see Definitions 5.6 and 5.10.

Thus, for any $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, we have defined a *descent map*

$$\hat{D}^{\vec{M}}_{\boldsymbol{s}}: \mathfrak{A}^-(\mathscr{H}^L, p^{\vec{M}}(\boldsymbol{s})) \to \mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}(M), \psi^{\vec{M}}(\boldsymbol{s})),$$

as a sum over compositions of polygon maps associated to various subhyperboxes of $\mathscr{H}^{\vec{L},\vec{M}}$. We used here the identification (98). Since $\hat{D}^{\vec{M}}_{\boldsymbol{s}}$ depends only on the projection $p^{\vec{M}}(\boldsymbol{s})$, we will write $\hat{D}^{\vec{M}}_{\boldsymbol{s}}$ as $\hat{D}^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$.

$$
\begin{tikzcd}[column sep=tiny]
 & & & \mathfrak{A}^-(\mathcal{H}^{L_2},\psi^{L_1}(s)) \arrow[dr, "\mathcal{I}^{L_2}_{\psi^{L_1}(s)}"] & & & \\
 & & \mathfrak{A}^-(\mathcal{H}^{L,L_1}(L_1),\psi^{L_1}(s)) \arrow[ur, leftrightarrow, "\cong"] \arrow[dr, "\mathcal{I}^{L_2}_{\psi^{L_1}(s)}"] & & \mathfrak{A}^-\big(\mathcal{H}^{L_2},p^{L_2}(\psi^{L_1}(s))\big) \arrow[ddr, "\widehat{D}^{L_2}_{p^{L_2}(\psi^{L_1}(s))}"] & & \\
 & & & \mathfrak{A}^-\big(\mathcal{H}^{L,L_1}(L_1),p^{L_2}(\psi^{L_1}(s))\big) \arrow[ur, leftrightarrow, "\cong"] \arrow[ddr, "\widehat{D}^{L_2}_{p^{L_2}(\psi^{L_1}(s))}"] & & & \\
 & \mathfrak{A}^-(\mathcal{H}^{L},p^{L_1}(s)) \arrow[uur, "\widehat{D}^{L_1}_{p^{L_1}(s)}"] \arrow[dr, "\mathcal{I}^{L_2}_{p^{L_1}(s)}"] & & & & \mathfrak{A}^-(\mathcal{H}^{L_2,L_2}(L_2),\psi^{L}(s)) \arrow[dl, leftrightarrow, "\cong"] \arrow[dr, leftrightarrow, "\cong"] & \\
\mathfrak{A}^-(\mathcal{H}^{L},s) \arrow[ur, "\mathcal{I}^{L_1}_s"] \arrow[rr, "\mathcal{I}^{L_1\cup L_2}_s"] \arrow[dr, "\mathcal{I}^{L_2}_s"] & & \mathfrak{A}^-(\mathcal{H}^{L},p^{L}(s)) \arrow[uur, "\widehat{D}^{L_1}_{p^{L}(s)}"] \arrow[rr, "\widehat{D}^{L}_{p^{L}(s)}"] \arrow[ddr, "\widehat{D}^{L_2}_{p^{L}(s)}"] & & \mathfrak{A}^-(\mathcal{H}^{L,L}(L),\psi^{L}(s)) \arrow[rr, leftrightarrow, "\cong"] \arrow[dr, leftrightarrow, "\cong"] & & \mathfrak{A}^-(\mathcal{H}^{\varnothing},\psi^{L}(s)) \\
 & \mathfrak{A}^-(\mathcal{H}^{L},p^{L_2}(s)) \arrow[ur, "\mathcal{I}^{L_1}_{p^{L_2}(s)}"] \arrow[ddr, "\widehat{D}^{L_2}_{p^{L_2}(s)}"] & & & & \mathfrak{A}^-(\mathcal{H}^{L_1,L_1}(L_1),\psi^{L}(s)) \arrow[ur, leftrightarrow, "\cong"] & \\
 & & & \mathfrak{A}^-\big(\mathcal{H}^{L,L_2}(L_2),p^{L_1}(\psi^{L_2}(s))\big) \arrow[uur, "\widehat{D}^{L_1}_{p^{L_1}(\psi^{L_2}(s))}"] \arrow[dr, leftrightarrow, "\cong"] & & & \\
 & & \mathfrak{A}^-(\mathcal{H}^{L,L_2}(L_2),\psi^{L_2}(s)) \arrow[ur, "\mathcal{I}^{L_1}_{\psi^{L_2}(s)}"] \arrow[dr, leftrightarrow, "\cong"] & & \mathfrak{A}^-\big(\mathcal{H}^{L_1},p^{L_1}(\psi^{L_2}(s))\big) \arrow[uur, "\widehat{D}^{L_1}_{p^{L_1}(\psi^{L_2}(s))}"] & & \\
 & & & \mathfrak{A}^-(\mathcal{H}^{L_1},\psi^{L_2}(s)) \arrow[ur, "\mathcal{I}^{L_1}_{\psi^{L_2}(s)}"] & & &
\end{tikzcd}
$$

Figure 27: Chain maps coming from hyperboxes in a complete system. We consider here a complete system $\mathcal{H}$ for a link $\vec{L} = L_1 \cup L_2$. Recall that one quarter of such a complete system was illustrated in Figure 21. We show here the corresponding chain maps: to each reduction (deletion of basepoints) in the complete system we have an associated inclusion map $\mathcal{I}$, to each face of a hyperbox we have an associated descent map $\widehat{D}$, and to each surface isotopy we have an associated isomorphism denoted by the symbol $\cong$. For simplicity, all the sublinks are taken here with their induced orientation from $\vec{L}$, and we drop the arrows from the notation. There are similar maps corresponding to other hyperboxes in the complete system, in which some link components can appear with the opposite orientation.

Example 9.1 Suppose M is a single link component L_i. Each edge of the one-dimensional hyperbox $\mathcal{H}^{\vec{L},\vec{L}_i}$ comes equipped with a corresponding theta chain element, and this gives a triangle map between the Floer homology groups associated to the initial and final Heegaard diagrams for that edge. The map $\widehat{D}^{\vec{L}_i}_{p^{\vec{M}}(s)}$ is the composition of these triangle maps. Note that all the triangle maps, and hence also $\widehat{D}^{\vec{L}_i}_{p^{\vec{M}}(s)}$, are chain homotopy equivalences; see Example 8.10.

$$\begin{tikzcd}[column sep=small]
 & & \mathfrak{A}^-(\mathscr{H}^{L,L_1}(L_1),\psi^{L_1}(s)) \arrow[dr, "\mathscr{I}^{L_2}_{\psi^{L_1}(s)}"] & & \\
 & & & \mathfrak{A}^-\big(\mathscr{H}^{L_2},p^{L_2}(\psi^{L_1}(s))\big) \arrow[ddr, "D^{L_2}_{p^{L_2}(\psi^{L_1}(s))}"] & \\
 & \mathfrak{A}^-(\mathscr{H}^L,p^{L_1}(s)) \arrow[uur, "D^{L_1}_{p^{L_1}(s)}"] \arrow[dr, "\mathscr{I}^{L_2}_{p^{L_1}(s)}"] & & & \\
\mathfrak{A}^-(\mathscr{H}^L,s) \arrow[ur, "\mathscr{I}^{L_1}_s"] \arrow[rr, "\mathscr{I}^{L_1\cup L_2}_s"] \arrow[dr, "\mathscr{I}^{L_2}_s"'] & & \mathfrak{A}^-(\mathscr{H}^L,p^L(s)) \arrow[uur, "D^{K_1}_{p^L(s)}"] \arrow[rr, "D^L_{p^L(s)}"] \arrow[ddr, "D^{L_2}_{p^L(s)}"] & & \mathfrak{A}^-(\mathscr{H}^{\varnothing},\psi^L(s)) \\
 & \mathfrak{A}^-(\mathscr{H}^L,p^{L_2}(s)) \arrow[ur, "\mathscr{I}^{L_1}_{p^{L_2}(s)}"'] \arrow[ddr, "D^{L_2}_{p^{L_2}(s)}"'] & & & \\
 & & & \mathfrak{A}^-\big(\mathscr{H}^{L_1},p^{L_1}(\psi^{L_2}(s))\big) \arrow[uur, "D^{L_1}_{p^{L_1}(\psi^{L_2}(s))}"'] & \\
 & & \mathfrak{A}^-(\mathscr{H}^{L,L_2}(L_2),\psi^{L_2}(s)) \arrow[ur, "\mathscr{I}^{L_1}_{\psi^{L_2}(s)}"'] & &
\end{tikzcd}$$

Figure 28: More chain maps coming from hyperboxes in a complete system. This is the same as Figure 27, but we have composed the descent maps $\widehat{D}$ with chain isomorphisms (coming from surface isotopies) to obtain the maps denoted by D.

In a complete system of hyperboxes, we have an isotopy $\mathscr{H}^{\vec{L},\vec{M}}(M) \cong \mathscr{H}^{L-M}$, which induces an identification between the respective Floer complexes. When we change the range of $\widehat{D}^{\vec{M}}_{p^{\vec{M}}(s)}$ via this identification, we denote the resulting map by

$$D^{\vec{M}}_{p^{\vec{M}}(s)} \colon \mathfrak{A}^-(\mathscr{H}^L, p^{\vec{M}}(s)) \to \mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^{\vec{M}}(s)). \tag{99}$$

Example 9.2 In the case of a link with two components, some of the inclusion and descent maps are illustrated in Figures 27 and 28.

For any $s \in \mathbb{H}(L)$, we now define a map

$$\Phi^{\vec{M}}_s = D^{\vec{M}}_{p^{\vec{M}}(s)} \circ \mathscr{I}^{\vec{M}}_s \colon \mathfrak{A}^-(\mathscr{H}^L, s) \to \mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^{\vec{M}}(s)). \tag{100}$$

Note that we can define similar maps if we replace $\vec{L}$ by a sublink $\vec{L}'$. By abuse of notation, we will always denote the maps corresponding to inclusion and descent at $\vec{M}$ by $\mathscr{I}^{\vec{M}}_s$, $D^{\vec{M}}_{p^{\vec{M}}(s)}$ and $\Phi^{\vec{M}}_s$, even though their domains of definition may vary.

Lemma 9.3 *Let $M_1, M_2 \subseteq L$ be two disjoint sublinks, with orientations $\vec{M}_1$ and $\vec{M}_2$. For any $s \in \overline{\mathbb{H}}(L)$,*

$$\mathscr{I}^{\vec{M}_2}_{\psi^{\vec{M}_1}(s)} \circ D^{\vec{M}_1}_{p^{\vec{M}_1}(s)} = D^{\vec{M}_1}_{p^{\vec{M}_1 \cup \vec{M}_2}(s)} \circ \mathscr{I}^{\vec{M}_2}_{p^{\vec{M}_1}(s)}.$$

Proof For $s = (s_1, \dots, s_\ell)$, let us denote by s_i' the components of $s' = \psi^{\vec{M}_1}(s)$, obtained from s_i by subtracting half the linking number of L_i and M_1; compare (20).

Observe that both $D^{\vec{M}_1}_{p^{\vec{M}_1}(s)}$ and $D^{\vec{M}_1}_{p^{\vec{M}_1\cup\vec{M}_2}(s)}$ are sums of compositions of polygon maps. The same polygons get counted in both maps, but with different powers of U_i. More precisely, suppose we have a chain of polygons relating intersection points $\boldsymbol{x}$ from $\mathscr{H}^L$ and $\boldsymbol{y}$ from $\mathscr{H}^{L-M_1-M_2}$. If j is one of the colors, let e_j (resp. e'_j) be the exponent of U_j in the coefficient of $\boldsymbol{y}$ in $D^{\vec{M}_1}_{p^{\vec{M}_1}(s)}(\boldsymbol{x})$ (resp. $D^{\vec{M}_1}_{p^{\vec{M}_1\cup\vec{M}_2}(s)}(\boldsymbol{x})$). We have

$$e_j-e'_j=\sum_{\substack{i\in I_+(\vec{L},\vec{M}_2)\\ \tau_i=j}}\bigl(\max(A_i(\boldsymbol{x})-s_i,0)-\max(A_i(\boldsymbol{y})-s'_i,0)\bigr)+\sum_{\substack{i\in I_-(\vec{L},\vec{M}_2)\\ \tau_i=j}}\bigl(\max(s_i-A_i(\boldsymbol{x}),0)-\max(s'_i-A_i(\boldsymbol{y}),0)\bigr);$$

compare Section 3.5.

On the other hand, according to (24), the map $\mathscr{I}^{\vec{M}_2}_{p^{\vec{M}_1}(s)}$ contributes a power of U_j with exponent

$$\sum_{\substack{i\in I_+(\vec{L},\vec{M}_2)\\ \tau_i=j}}\max(A_i(\boldsymbol{x})-s_i,0)+\sum_{\substack{i\in I_-(\vec{L},\vec{M}_2)\\ \tau_i=j}}\max(s_i-A_i(\boldsymbol{x}),0)$$

and the map $\mathscr{I}^{\vec{M}_2}_{\psi^{\vec{M}_1}(s)}$ a power of U_j with exponent

$$\sum_{\substack{i\in I_+(\vec{L},\vec{M}_2)\\ \tau_i=j}}\max(A_i(\boldsymbol{y})-s'_i,0)+\sum_{\substack{i\in I_-(\vec{L},\vec{M}_2)\\ \tau_i=j}}\max(s'_i-A_i(\boldsymbol{y}),0).$$

These contributions exactly cancel out the difference between e'_j and e_j. □

Proposition 9.4 *Choose a sublink $M\subseteq L$, and endow it with an orientation $\vec{M}$. Then, for any $s\in\overline{\mathbb{H}}(L)$,*

$$\sum_{\vec{M}_1\amalg\vec{M}_2=\vec{M}}\Phi^{\vec{M}_2}_{\psi^{\vec{M}_1}(s)}\circ\Phi^{\vec{M}_1}_{s}=0, \tag{101}$$

where $\vec{M}_1$ and $\vec{M}_2$ are equipped with the orientations induced from $\vec{M}$.

Proof We have

$$\begin{aligned}\sum_{\vec{M}_1\amalg\vec{M}_2=\vec{M}}\Phi^{\vec{M}_2}_{\psi^{\vec{M}_1}(s)}\circ\Phi^{\vec{M}_1}_{s}&=\sum_{\vec{M}_1\amalg\vec{M}_2=\vec{M}}D^{\vec{M}_2}_{p^{\vec{M}_2}(\psi^{\vec{M}_1}(s))}\circ\mathscr{I}^{\vec{M}_2}_{\psi^{\vec{M}_1}(s)}\circ D^{\vec{M}_1}_{p^{\vec{M}_1}(s)}\circ\mathscr{I}^{\vec{M}_1}_{s}\\&=\sum_{\vec{M}_1\amalg\vec{M}_2=\vec{M}}D^{\vec{M}_2}_{p^{\vec{M}_2}(\psi^{\vec{M}_1}(s))}\circ D^{\vec{M}_1}_{p^{\vec{M}_2}(\psi^{\vec{M}_1}(s))}\circ\mathscr{I}^{\vec{M}_2}_{p^{\vec{M}_1}(s)}\circ\mathscr{I}^{\vec{M}_1}_{s}\\&=\Bigl(\sum_{\vec{M}_1\amalg\vec{M}_2=\vec{M}}D^{\vec{M}_2}_{p^{\vec{M}_2}(\psi^{\vec{M}_1}(s))}\circ D^{\vec{M}_1}_{p^{\vec{M}_2}(\psi^{\vec{M}_1}(s))}\Bigr)\circ\mathscr{I}^{\vec{M}_1\cup\vec{M}_2}_{s}\\&=0.\end{aligned}$$

Indeed, the second equality above follows from Lemma 9.3, and the last equality is a consequence of Proposition 5.14, together with the properties of a complete system of hyperboxes. □

9.2 The surgery theorem

Let us fix a framing Λ for the link $\vec{L}$. For a component L_i of L, we let Λ_i be its induced framing, thought of as an element in $H_1(Y-L)$.

Given a sublink $N \subseteq L$, we let $\Omega(N)$ be the set of all possible orientations on N, as in the introduction. For $\vec{N} \in \Omega(N)$, we let

$$\Lambda_{\vec{L},\vec{N}} = \sum_{i \in I_-(\vec{L},\vec{N})} \Lambda_i \in H_1(Y-L).$$

We view $\mathbb{H}(L) \subseteq H_1(Y-L;\mathbb{Q})$ as an affine lattice over $H_1(Y-L)$ as in Remark 3.6. Thus, if $\boldsymbol{s} \in \mathbb{H}(L)$, then $\boldsymbol{s} + \Lambda_{\vec{L},\vec{N}}$ is also in $\mathbb{H}(L)$.

Now consider the $\mathcal{R}$-module

$$\mathcal{C}^-(\mathcal{H},\Lambda) = \bigoplus_{M \subseteq L} \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s})), \tag{102}$$

where ψ^M simply means $\psi^{\vec{M}}$ with $\vec{M}$ being the orientation induced from the one on $\vec{L}$.

Note that the definition of $\mathcal{C}^-(\mathcal{H},\Lambda)$ involves direct products; in fact, the direct sum in (102) can equally be thought of as a direct product, since it is finite. It is worth saying a few words about how one can define maps between direct products:

Definition 9.5 Let S and T be countable index sets, and $\mathcal{A} = \prod_{s \in S} \mathcal{A}_s$ and $\mathcal{B} = \prod_{t \in T} \mathcal{B}_t$ direct products of modules over a commutative ring R. Suppose we are given module homomorphisms

$$F_{s,t} \colon \mathcal{A}_s \to \mathcal{B}_t$$

for each $s \in S$ and $t \in T$. The collection of maps $\{F_{s,t}\}$ is called *locally finite* if, for each $t \in T$, only finitely many $F_{s,t}$ are nonzero.

If $\{F_{s,t}\}$ is a locally finite collection of homomorphisms as in Definition 9.5, we can assemble them into a single homomorphism

$$F \colon \mathcal{A} \to \mathcal{B}, \quad F(\{a_s\}_{s \in S}) = \Bigl\{\sum_{s \in S} F_{s,t}(a_s)\Bigr\}_{t \in T}.$$

With this in mind, we equip the module $\mathcal{C}^-(\mathcal{H},\Lambda)$ with a boundary operator $\mathcal{D}^-$ as follows. For $\boldsymbol{s} \in \mathbb{H}(L)$ and $\boldsymbol{x} \in \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$, we set

$$\begin{aligned}\mathcal{D}^-(\boldsymbol{s},\boldsymbol{x}) &= \sum_{N \subseteq L-M} \sum_{\vec{N} \in \Omega(N)} (\boldsymbol{s} + \Lambda_{\vec{L},\vec{N}}, \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}(\boldsymbol{x})) \\ &\in \bigoplus_{N \subseteq L-M} \bigoplus_{\vec{N} \in \Omega(N)} \mathfrak{A}^-(\mathcal{H}^{L-M-N}, \psi^{M \cup \vec{N}}(\boldsymbol{s})) \subseteq \mathcal{C}^-(\mathcal{H},\Lambda).\end{aligned}$$

This defines a locally finite collection of maps between the modules $\mathfrak{A}^-(\mathscr{H}^{L-M}, \psi^M(s))$, producing a well-defined map

$$\mathscr{D}^-: \mathscr{C}^-(\mathscr{H}, \Lambda) \to \mathscr{C}^-(\mathscr{H}, \Lambda).$$

According to Proposition 9.4, $\mathscr{C}^-(\mathscr{H}, \Lambda)$ is a chain complex. Note that $\mathscr{C}^-(\mathscr{H}, \Lambda)$ naturally breaks into a direct product of terms $\mathscr{C}^-(\mathscr{H}, \Lambda, \mathfrak{u})$, according to equivalence classes $\mathfrak{u}$ of the values s. Here s_1 and s_2 are equivalent if they differ by an element in the (possibly degenerate) sublattice

$$H(L, \Lambda) \subseteq H_1(Y - L),$$

generated by all possible $\Lambda_{\vec{L},\vec{N}}$ or, equivalently, by the component framings $\Lambda_i \in H_1(Y - L)$. The space of equivalence classes is parametrized by the quotient

$$\mathbb{H}(L)/H(L, \Lambda),$$

which can be naturally identified with the space of Spin^c structures on the surgered manifold $Y_\Lambda(L)$; see Remark 3.6 and [Ozsváth and Szabó 2008a, Section 3.7].

Given a Spin^c structure $\mathfrak{u}$ on $Y_\Lambda(L)$, we set

$$\mathfrak{d}(\mathfrak{u}) = \gcd_{\xi \in H_2(Y_\Lambda(L);\mathbb{Z})} \langle c_1(\mathfrak{u}), \xi \rangle, \tag{103}$$

where $c_1(\mathfrak{u})$ is the first Chern class of the Spin^c structure. The Heegaard Floer homology $\boldsymbol{HF}^-_*(Y_\Lambda(L), \mathfrak{u})$ admits a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading; see [Ozsváth and Szabó 2004b]. In Section 9.3 we will construct a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading on the complex $\mathscr{C}^-(\mathscr{H}, \Lambda, \mathfrak{u})$ as well.

The surgery theorem (for link-minimal systems) then says:

Theorem 9.6 *Fix a link-minimal complete system of hyperboxes $\mathscr{H}$ for an oriented, ℓ-component link $\vec{L}$ in an integral homology three-sphere Y, and fix a framing Λ of L. Then, for any $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L, \Lambda)$, there is an isomorphism of relatively graded $\mathbb{F}[\![U]\!]$-modules*

$$H_*(\mathscr{C}^-(\mathscr{H}, \Lambda, \mathfrak{u}), \mathscr{D}^-) \cong \boldsymbol{HF}^-_*(Y_\Lambda(L), \mathfrak{u}). \tag{104}$$

Note that the left-hand side of (104) is a priori an $\mathscr{R}$-module, where $\mathscr{R} = \mathbb{F}[\![U_1, \dots, U_k]\!]$. However, part of the claim of the theorem is that all U_i's act the same, so we can think of it as an $\mathbb{F}[\![U]\!]$-module.

The proof of Theorem 9.6 will be given in Section 12.

Remark 9.7 In the case when the Spin^c structure $\mathfrak{u}$ is torsion, one should be able to use the same techniques as in [Ozsváth and Szabó 2008b] to obtain an isomorphism of absolutely graded groups, with a well-determined shift in grading between the two sides of (104). However, we will not pursue this direction here.

9.3 Gradings

As promised in the previous subsection, we proceed to construct a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading on the complex $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})$.

Let us identify $H_1(Y - L)$ with $\mathbb{Z}^\ell$ as in Remark 3.6. We view the framing Λ as an ℓ-by-ℓ symmetric matrix with columns Λ_i. The matrix element c_{ij} in Λ (for $i, j = 1, \dots, \ell$) is the linking number between L_i and L_j when $i \neq j$, and the surgery coefficient λ_i of L_i when $i = j$.

Let $H(L,\Lambda)^\perp \subset \mathbb{Z}^\ell$ be the orthogonal complement to $H(L,\Lambda)$, that is,

$$H(L,\Lambda)^\perp = \{\boldsymbol{v} \in \mathbb{Z}^\ell \mid \boldsymbol{v} \cdot \Lambda_i = 0 \text{ for all } i\} = \left\{(v_1, \dots, v_\ell) \in \mathbb{Z}^\ell \;\middle|\; \sum_i v_i \Lambda_i = 0\right\}.$$

There are natural identifications

$$H^2(Y_\Lambda(L)) \cong H_1(Y_\Lambda(L)) \cong \mathbb{Z}^\ell / H(L,\Lambda)$$

such that

$$c_1([\boldsymbol{s}]) = [2\boldsymbol{s}]$$

for any $\boldsymbol{s} \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$. This can be deduced from the formulas for the Chern class in [Ozsváth and Szabó 2008a, equation (24) and Lemma 3.13]; compare also Lemma 11.2 below.

Using Poincaré duality, we obtain a natural identification $H_2(Y_\Lambda(L)) \cong H(L,\Lambda)^\perp$. Hence,

$$\mathfrak{d}(\mathfrak{u}) = \gcd_{\boldsymbol{v} \in H(L,\Lambda)^\perp} \sum_i 2 s_i v_i, \tag{105}$$

where we wrote $\boldsymbol{v} = (v_1, \dots, v_\ell)$, and $\boldsymbol{s} = (s_1, \dots, s_\ell)$ is any element in the corresponding equivalence class $\mathfrak{u} \in \mathbb{H}(L)/H(L,\Lambda)$.

Remark 9.8 It is clear from (103) that $\mathfrak{d}(\mathfrak{u})$ is always even. One can also verify this using the description of $\mathfrak{d}(\mathfrak{u})$ given in (105). Indeed, let $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_\ell) = (c_{11}, \dots, c_{\ell\ell})$ be the diagonal vector of the framing matrix Λ. For $\boldsymbol{s} \in \mathbb{H}(L)$ and $\boldsymbol{v} \in H(L,\Lambda)^\perp$, we have $2\boldsymbol{s} \equiv \Lambda_1 + \cdots + \Lambda_\ell - \boldsymbol{\lambda} \pmod 2$ and $(\Lambda_1 + \cdots + \Lambda_\ell) \cdot \boldsymbol{v} = 0$, so

$$(2\boldsymbol{s}) \cdot \boldsymbol{v} \equiv \boldsymbol{\lambda} \cdot \boldsymbol{v} \equiv \sum_{i=1}^{\ell} c_{ii} v_i \equiv \sum_{i=1}^{\ell} c_{ii} v_i^2 \equiv \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} c_{ij} v_i v_j \equiv \boldsymbol{v}^T \Lambda \boldsymbol{v} \equiv 0 \pmod 2.$$

Lemma 9.9 *Fix an equivalence class $\mathfrak{u} \in \mathbb{H}(L)/H(L,\Lambda)$. There exists a function $\nu : \mathfrak{u} \to \mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$ with the property that*

$$\nu(\boldsymbol{s} + \Lambda_i) \equiv \nu(\boldsymbol{s}) + 2 s_i \tag{106}$$

for any $i = 1, \dots, \ell$ and $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \mathfrak{u}$.

Proof Pick some $\boldsymbol{s}^0 = (s_1^0, \dots, s_\ell^0) \in \mathfrak{u}$. Any other $\boldsymbol{s} \in \mathfrak{u}$ is of the form $\boldsymbol{s}^0 + \sum a_i \Lambda_i$, for some $a_i \in \mathbb{Z}$. Set

$$\nu\Bigl(\boldsymbol{s}^0 + \sum_{i=1}^{\ell} a_i \Lambda_i\Bigr) = \sum_{i=1}^{\ell} 2a_i s_i^0 + \sum_{i,j=1}^{\ell} a_i a_j c_{ij}.$$

There is an indeterminacy in expressing $\boldsymbol{s}$ as $\boldsymbol{s}^0 + \sum a_i \Lambda_i$, namely one can add an element in $H(L,\Lambda)^\perp$ to the vector $(a_1, \dots, a_\ell)$. It is easy to check that $\nu(\boldsymbol{s})$ is independent (modulo $\mathfrak{d}(\mathfrak{u})$) of how we express $\boldsymbol{s}$, and that (106) is satisfied. $\square$

Remark 9.10 The function ν from Lemma 9.9 is unique up to the addition of a constant.

Fix a function ν as in Lemma 9.9. Each factor $\mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$ appearing in the complex $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})$ admits a natural $\mathbb{Z}$-grading $\mu_{\boldsymbol{s}}^M = \mu_{\psi^M(\boldsymbol{s})}$ as in (16). We define a $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading μ on $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})$ as follows. For $\boldsymbol{s} \in \mathfrak{u}$ and $\boldsymbol{x} \in \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$, we set

$$\mu(\boldsymbol{s}, \boldsymbol{x}) = \mu_{\boldsymbol{s}}^M(\boldsymbol{x}) + \nu(\boldsymbol{s}) - |M|,$$

where $|M|$ denotes the number of components of M.

Lemma 9.11 *The differential $\mathcal{D}^-$ on $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})$ decreases μ by one modulo $\mathfrak{d}(\mathfrak{u})$.*

Proof Use (26), (100), and the fact that $D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$ (being the longest map in an $|M|$-dimensional hypercube of chain complexes) changes grading by $|M| - 1$. $\square$

10 Truncation

As noted near the end of Section 4.1, in order to calculate the Heegaard Floer homology groups $\boldsymbol{HF}^-(Y_\Lambda(L), \mathfrak{u})$ using Theorem 9.6, it is helpful to replace the infinite direct product from (102) with a finite one. This is called *horizontal truncation*, and we saw a few instances of it in Sections 4.3 and 4.5. In Section 10.1, we will describe a general way of doing horizontal truncation for surgery on an arbitrary link.

Even after horizontal truncation, the direct product in (102) is still an infinite-dimensional $\mathbb{F}$-vector space, due to the fact that each term is a free module over a ring of power series. However, in Section 10.2 we show that the power series ring can be replaced (essentially without any loss of information) by a finite-dimensional polynomial ring. This process is called *vertical truncation*, and is done by setting large powers of the U variables to zero. By combining horizontal and vertical truncation, we can replace the right-hand side of (102) with a finite-dimensional chain complex. In Section 10.3, we describe an alternative way of doing so, by applying a slightly different horizontal truncation to the vertically truncated complex. In Section 10.4, we alter the combined truncation procedure from Section 10.3 further, by adding certain "crossover maps". We call the result a *folded truncation* of the original complex. Folded truncation will play an important role in the proof of Theorem 9.6 presented in Section 12.

Before describing all these truncations, let us recall some notation from Section 4. We denote a typical term in the surgery complex (102) by

$$\mathcal{C}^{\varepsilon}_{\boldsymbol{s}} = \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s})),$$

where $\varepsilon = \varepsilon(M) \in \mathbb{E}_\ell = \{0,1\}^\ell$ is such that $L_i \subseteq M$ if and only if $\varepsilon_i = 1$. A typical summand in the differential $\mathcal{D}^-$ is denoted by

$$\mathcal{D}^{\varepsilon,\varepsilon'}_{\varepsilon^0,\boldsymbol{s}} = \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})} \colon \mathcal{C}^{\varepsilon^0}_{\boldsymbol{s}} \to \mathcal{C}^{\varepsilon^0+\varepsilon}_{\boldsymbol{s}+\varepsilon'\cdot\Lambda},$$

where $\varepsilon^0 = \varepsilon(M)$, $\varepsilon = \varepsilon(N)$, and $\varepsilon' \in \mathbb{E}_\ell$ is such that $i \in I_-(\vec{L}, \vec{N})$ if and only if $\varepsilon'_i = 1$. The dot product $\varepsilon' \cdot \Lambda$ denotes the vector $\sum \varepsilon'_i \Lambda_i$. Note that we always have $\varepsilon' \le \varepsilon$.

Dropping a subscript or superscript from the notation means considering the direct product over all possible values of that subscript or superscript. For example, $\mathcal{C}^\varepsilon = \prod_{\boldsymbol{s}} \mathcal{C}^\varepsilon_{\boldsymbol{s}}$, and $\mathcal{C} = \bigoplus_\varepsilon \mathcal{C}^\varepsilon = \mathcal{C}(\mathcal{H}, \Lambda)$. Observe that $(\mathcal{C}^\varepsilon, \mathcal{D}^\varepsilon)$ forms a hypercube of chain complexes as defined in Section 5.1 (except it may have only a $\mathbb{Z}/2\mathbb{Z}$-grading, rather than a $\mathbb{Z}$-grading) and $\mathcal{C}$ is the total complex of this hypercube.

10.1 Horizontal truncation

We now return to the setting of Section 9, where L is an arbitrary link, equipped with an orientation, a complete system of hyperboxes $\mathcal{H}$, and a framing Λ.

Lemma 10.1 *There exists a constant $b > 0$ such that, for any $i = 1, \dots, \ell$ and for any sublink $M \subseteq L$ not containing the component L_i, the chain map*

$$\Phi^{\vec{L}_i}_{\psi^{\vec{M}}(\boldsymbol{s})} \colon \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^{\vec{M}}(\boldsymbol{s})) \to \mathfrak{A}^-(\mathcal{H}^{L-M-L_i}, \psi^{\vec{M}\cup\vec{L}_i}(\boldsymbol{s}))$$

induces an isomorphism on homology provided that either

- *$\boldsymbol{s} \in \mathbb{H}(L)$ is such that $s_i > b$, and L_i is given the orientation induced from L; or*
- *$\boldsymbol{s} \in \mathbb{H}(L)$ is such that $s_i < -b$, and L_i is given the orientation opposite the one induced from L.*

Proof For $|s_i|$ sufficiently large, and L_i oriented as in the lemma (according to the sign of s_i), the inclusion map $\mathcal{I}^{\vec{L}_i}_{\psi^{\vec{M}}(\boldsymbol{s})}$ from (24) is the identity. Moreover, the descent map $D^{\vec{L}_i}_{p^{\vec{L}_i}(\psi^{\vec{M}}(\boldsymbol{s}))}$ is a composition of maps along the edges of the corresponding hyperbox $\mathcal{H}^{\vec{L}-M, L_i}$, and hence induces an isomorphism on homology; see Example 9.1. The conclusion then follows in light of (100). □

Lemma 10.1 is the key ingredient in truncating the complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$. Roughly, it allows the terms of this complex to cancel in pairs, whenever $\boldsymbol{s} \in \mathbb{H}(L)$ has at least one component s_i with $|s_i| > b$. The result is that the homology of the complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$ can be computed by restricting to some terms corresponding to $\boldsymbol{s}$ lying in a compact subset of $\mathbb{H}(L)$.

Of course, we need to be more explicit about how this is done. For simplicity, let us assume (for the moment) that the framing vectors $\Lambda_1, \dots, \Lambda_\ell$ are linearly independent in $\mathbb{R}^\ell$. This is equivalent to asking for $Y_\Lambda(L)$ to be a rational homology sphere.

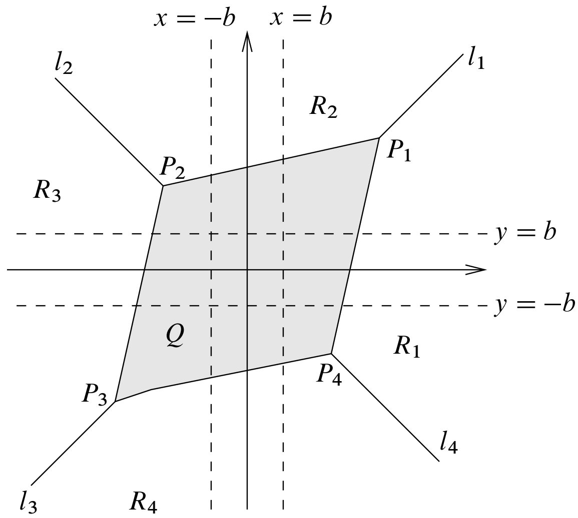

Figure 29: The parallelogram Q and the regions R_i. The parallelogram Q, which is (roughly) the support of the truncated complex $\mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle$, is shown shaded.

Let us first recall the case of surgery on knots; see [Ozsváth and Szabó 2008b, Section 4.1]. Then the framing coefficient Λ is a nonzero integer. Set

$$\mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle = \bigoplus_{-b\le s\le b} \mathfrak{A}^-(\mathcal{H}^L, s) \oplus \bigoplus_{-b+\Lambda\le s\le b} \mathfrak{A}^-(\mathcal{H}^\varnothing, \psi^L(s)) \subset \mathcal{C}^-(\mathcal{H},\Lambda). \tag{107}$$

It is easy to check that $\mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle$ is a subcomplex of $\mathcal{C}^-(\mathcal{H},\Lambda)$ for $\Lambda<0$, and a quotient complex of $\mathcal{C}^-(\mathcal{H},\Lambda)$ when $\Lambda>0$. In both cases, an application of Lemma 10.1 shows that $\mathcal{C}^-(\mathcal{H},\Lambda)$ and $\mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle$ are quasi-isomorphic.

Next, we turn to the case when the link L has two components. We denote by λ_1 and λ_2 the framings of L_1 and L_2 (as compared to the framings coming from Seifert surfaces for those components), and by c the linking number between L_1 and L_2. Thus, in terms of the standard basis of $H_1(Y-L)\cong\mathbb{Z}^2$,

$$\Lambda_1=(\lambda_1,c), \quad \Lambda_2=(c,\lambda_2).$$

Recall our assumption that Λ_1 and Λ_2 are linearly independent, ie $\lambda_1\lambda_2-c^2\neq 0$.

As before, $\mathcal{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ will denote the term $\mathfrak{A}^-(\mathcal{H}^{L-M},\psi^M(\boldsymbol{s}))$ appearing in (102), where $\varepsilon_i=1$ or 0 depending on whether or not $L_i\subseteq M$ for $i=1,2$. We say that $\mathcal{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ is supported at the point $\boldsymbol{s}=(s_1,s_2)\in\mathbb{H}(L)\subset\mathbb{R}^2$.

Let b be the constant from Lemma 10.1. We seek to define a chain complex $\mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle$ quasi-isomorphic to $\mathcal{C}^-(\mathcal{H},\Lambda)$, and composed of only finitely many of the terms $\mathcal{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$. For this purpose, we construct a parallelogram Q in the plane, with edges parallel to the vectors Λ_1 and Λ_2, and with vertices P_1, P_2, P_3 and P_4 as in Figure 29. We require the coordinates of P_1 to satisfy $x,y>b$, the coordinates of P_2

to satisfy $x < -b$, $y > b$, etc. Further, we choose half-lines l_1, l_2, l_3 and l_4 with l_i starting at P_i and staying in the i^{th} quadrant. For example, l_1 has to form an angle between 0 and $\frac{\pi}{2}$ with the x-axis, etc. The half-lines l_i split the complement $\mathbb{R}^2 - Q$ into four regions, denoted by R_1, R_2, R_3 and R_4; see Figure 29. We require that the edges of the parallelogram Q and the half-lines l_i miss the lattice $\mathbb{H}(L)$.

We require that the slopes of the lines l_1, l_2, l_3 and l_4, compared to those of the vectors Λ_1 and Λ_2, are as indicated in Figure 30. We distinguish there six cases, according to the values λ_1, λ_2 and c. (Note that these cases cover all the possibilities for $\lambda_1\lambda_2 - c^2 \neq 0$, with some overlap. When we are in an overlap situation, we are free to choose either setting.) For example, in Case I, we require both vectors Λ_1 and Λ_2 to point into R_2 when placed on the half-line l_2.

By Lemma 10.1, the maps Φ^{L_1} appearing in the definition of the differential for $\mathcal{C}^-(\mathcal{H}, \Lambda)$ induce isomorphisms on homology when restricted to terms supported in R_1. The same is true for the maps Φ^{L_2} supported in R_2, for the maps Φ^{-L_1} supported in R_3, and for the maps Φ^{-L_2} supported in R_4. (Here, as in Section 4.5, we let L_1 and L_2 have the orientation induced from L, and we add a minus sign when we wish to indicate the opposite orientation.) These facts allow us to truncate the complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$ to a obtain a quasi-isomorphic one $\mathcal{C}_Q = \mathcal{C}^-(\mathcal{H}, \Lambda)\langle b\rangle$, which is (roughly) supported in the parallelogram Q. This is obtained by taking successive subcomplexes and quotient complexes from $\mathcal{C}^-(\mathcal{H}, \Lambda)$, and canceling out complexes supported in R_1, R_2, R_3 and R_4. Some care has to be taken with what happens near the boundaries of these regions, so we proceed to do a case-by-case analysis. We present the first case in detail, and for the others we sketch the necessary modifications.

Case I $(\lambda_1, \lambda_2 > 0, \lambda_1\lambda_2 - c^2 > 0)$ Let $\mathcal{C}_{R_1 \cup R_2}$ be the subcomplex of $\mathcal{C}^-(\mathcal{H}, \Lambda)$ consisting of those terms supported in $R_1 \cup R_2$. There is a filtration $\mathcal{F}_{00}$ on $\mathcal{C}_{R_1\cup R_2}$ (analogous to the one used for the Hopf link in Section 4.5) such that $\mathcal{F}_{00}$ is bounded above, and in the associated graded we only see the differentials that preserve $\boldsymbol{s}$. The associated graded then splits into a direct sum $\mathcal{C}_{R_1} \oplus \mathcal{C}_{R_2}$, where $\mathcal{C}_{R_i}$ is the direct product of the terms supported in R_i. On $\mathcal{C}_{R_1}$ there is an additional two-step filtration such that the differential of its associated graded consists only of maps of the form Φ^{L_1}, which are isomorphisms on homology. Hence, by applying Lemma 4.1, we have $H_*(\mathcal{C}_{R_1}) = 0$. Similarly $H_*(\mathcal{C}_{R_2}) = 0$, using an associated graded that leaves only the maps Φ^{L_2} in the differential. Putting these together, and using Lemma 4.1 again, we deduce that $\mathcal{C}_{R_1\cup R_2}$ is acyclic.

The quotient complex of $\mathcal{C}_{R_1\cup R_2}$ is supported in $Q \cup R_3 \cup R_4$. Let us define a subcomplex of it, denoted by $\mathcal{C}_{R_3\cup R_4}$, to consist of those terms $\mathfrak{A}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ with the property that $\boldsymbol{s} - \varepsilon_1\Lambda_1 - \varepsilon_2\Lambda_2 \in R_3 \cup R_4$. This is roughly supported in $R_3 \cup R_4$, although some terms spill over into Q. We define $\mathcal{C}_Q$ to be its quotient complex.

We claim that $\mathcal{C}_Q$ is quasi-isomorphic to the original complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$. For this, we need to show that $\mathcal{C}_{R_3\cup R_4}$ is acyclic. In the region R_3, we would like to use the maps Φ^{-L_2} to cancel out terms in pairs. However, there exist a whole region of values $\boldsymbol{s} \in R_3$ such that $\boldsymbol{s} + \Lambda_2$ lands in R_2 rather than R_3. The direct product of $\mathcal{C}_{\boldsymbol{s}}^{00}$ and $\mathcal{C}_{\boldsymbol{s}}^{01}$ over the values $\boldsymbol{s}$ in that region forms a quotient complex of $\mathcal{C}_{R_3\cup R_4}$; this

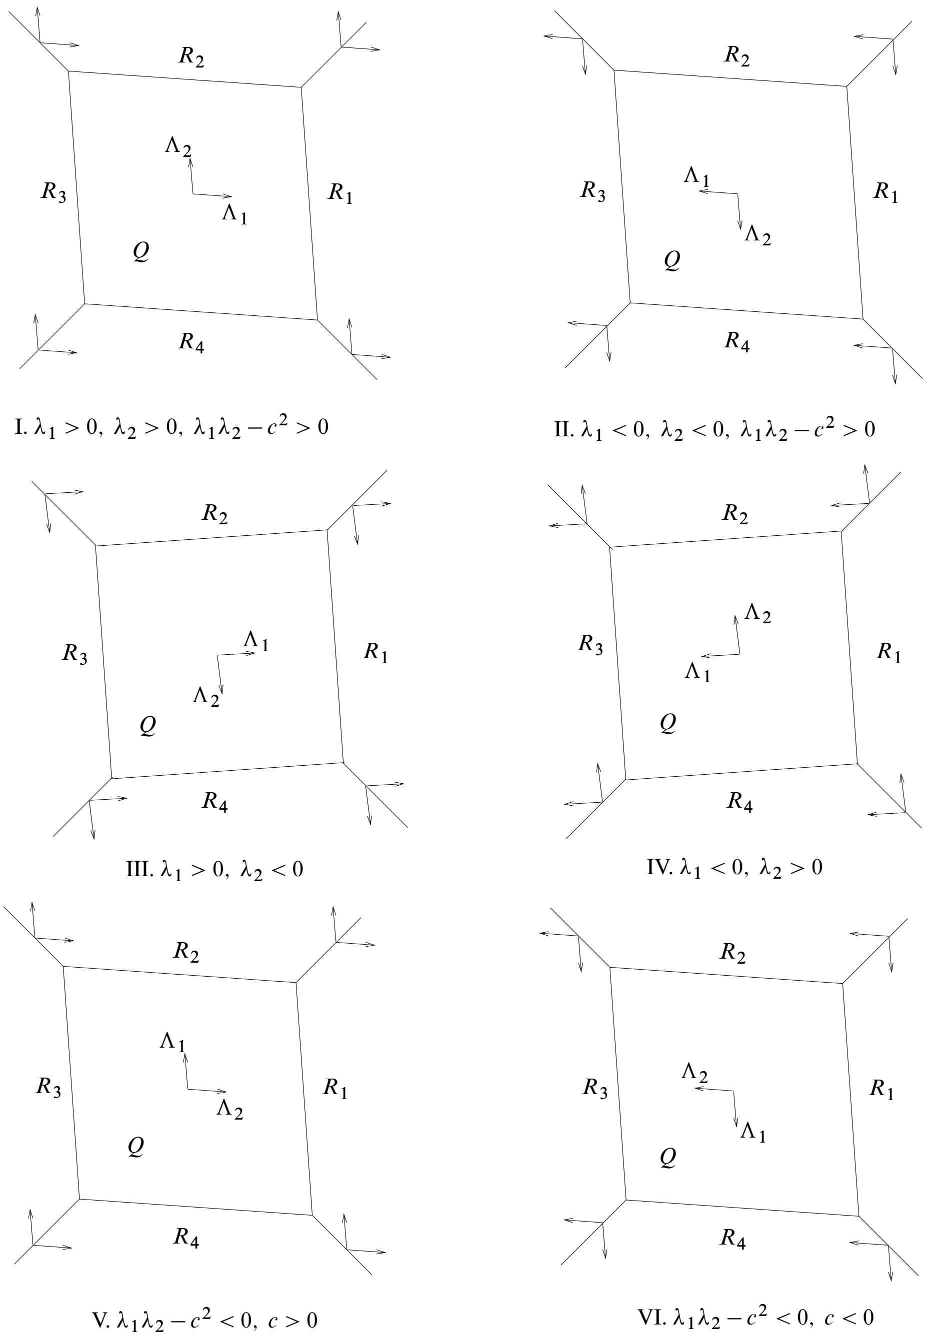

Figure 30: Horizontal truncation. We show here the required slopes for the half-lines l_i, in relation to the vectors Λ_1 and Λ_2. There are six cases.

quotient is acyclic, because the maps Φ^{L_2} (being close to the region R_2) make the terms $\mathscr{C}_{\boldsymbol{s}}^{00}$ and $\mathscr{C}_{\boldsymbol{s}}^{01}$ cancel out in pairs. The corresponding subcomplex $\mathscr{C}'_{R_3 \cup R_4}$ is quasi-isomorphic to $\mathscr{C}_{R_3 \cup R_4}$. Similarly, we can eliminate the terms $\mathscr{C}_{\boldsymbol{s}}^{00}$ and $\mathscr{C}_{\boldsymbol{s}}^{10}$ from $\mathscr{C}'_{R_3 \cup R_4}$ for those $\boldsymbol{s} \in R_4$ such that $\boldsymbol{s} + \Lambda_1 \in R_1$. The result is a quasi-isomorphic complex $\mathscr{C}''_{R_3\cup R_4}$. Consider the associated graded of this complex with respect to a filtration $\mathscr{F}_{11}$ (analogous to the one used for the Hopf link in Section 4.5) such that the remaining differentials preserve $\boldsymbol{s}-\varepsilon_1\Lambda_1-\varepsilon_2\Lambda_2$. The associated graded is acyclic, as it splits into a direct sum according to the regions R_3 and R_4, and those are acyclic by Lemma 10.1. We conclude that $H_*(\mathscr{C}_{R_3\cup R_4}) = H_*(\mathscr{C}''_{R_3\cup R_4}) = 0$.

Case II $(\lambda_1,\lambda_2<0,\ \lambda_1\lambda_2-c^2>0)$ This is similar to Case I, except quotient complexes and subcomplexes are swapped, and the unbounded filtrations are reversed. Specifically, $\mathscr{C}_{R_1\cup R_2}$ is a quotient complex, on which $-\mathscr{F}_{00}$ defines a filtration. Its associated graded splits as $\mathscr{C}_{R_1}\oplus\mathscr{C}_{R_2}$ and is acyclic (by using the same two-step filtration as in Case I). The filtration $-\mathscr{F}_{00}$ is no longer bounded above; still, it is U-tame in the sense of Definition 4.3, and we can apply Lemma 4.4 to deduce that $\mathscr{C}_{R_1\cup R_2}$ is acyclic. There is a subcomplex of the whole $\mathscr{C}^-(\mathscr{H},\Lambda)$ with $\mathscr{C}_{R_1\cup R_2}$ as the corresponding quotient; that subcomplex has a quotient complex $\mathscr{C}_{R_3\cup R_4}$ and a subcomplex $\mathscr{C}_Q$. We can prove that $\mathscr{C}_{R_3\cup R_4}$ is acyclic similarly to how we did so in Case I, this time relying on Lemma 4.4 when the filtrations are not bounded above. We conclude that $\mathscr{C}^-(\mathscr{H},\Lambda)$ is quasi-isomorphic to $\mathscr{C}_Q$.

Case III $(\lambda_1>0,\ \lambda_2<0)$ We define a subcomplex $\mathscr{C}_{R_1}$ composed of those terms $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ such that either $\boldsymbol{s}\in R_1$, or $\boldsymbol{s}\in R_4$, $\varepsilon_2=1$ and $\boldsymbol{s}-\Lambda_2\in R_1$. We also define another subcomplex $\mathscr{C}_{R_3}$ composed of those terms $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ such that either $\boldsymbol{s}-\varepsilon_1\Lambda_1\in R_3$, or $\boldsymbol{s}\in R_4$, $\varepsilon_2=1$ and $\boldsymbol{s}-\varepsilon_1\Lambda_1-\Lambda_2\in R_3$. Both $\mathscr{C}_{R_1}$ and $\mathscr{C}_{R_3}$ are acyclic. The corresponding quotient complex admits two further acyclic quotient complexes: one, $\mathscr{C}_{R_2}$, consisting of $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ with $\boldsymbol{s}\in R_2$ and the other, $\mathscr{C}_{R_4}$, consisting of $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon_1\varepsilon_2}$ such that $\boldsymbol{s}-\varepsilon_2\Lambda_2\in R_4$. We let $\mathscr{C}_Q$ be the resulting subcomplex.

Case IV $(\lambda_1<0,\lambda_2>0)$ This is similar to Case III, but now $\mathscr{C}_{R_1}$ and $\mathscr{C}_{R_3}$ are quotient complexes, while $\mathscr{C}_{R_2}$ and $\mathscr{C}_{R_4}$ are subcomplexes.

Case V $(c>0,\ \lambda_1\lambda_2-c^2<0)$ This is similar to Case I, except that in order to make the complex $\mathscr{C}_{R_3\cup R_4}$ acyclic, we need to slightly modify its definition near the corners P_2, P_4 of the parallelogram Q. Specifically, when we define $\mathscr{C}_{R_3\cup R_4}$, we still ask that $\boldsymbol{s}-\varepsilon_1\Lambda_1-\varepsilon_2\Lambda_2\in R_3\cup R_4$, but from such pairs $(\boldsymbol{s},\varepsilon)$ we eliminate those with

(a) $\boldsymbol{s}\in Q$, $\boldsymbol{s}+\Lambda_1\in R_2$, $\boldsymbol{s}-\Lambda_2\in R_3$ and $\varepsilon=(0,1)$; or

(b) $\boldsymbol{s}\in Q, \boldsymbol{s}+\Lambda_2\in R_1$, $\boldsymbol{s}-\Lambda_1\in R_4$ and $\varepsilon=(1,0)$.

Case VI $(c<0,\ \lambda_1\lambda_2-c^2<0)$ This is similar to Case II, except we need to make adjustments near the corners P_2 and P_4 as in Case V.

A similar construction can be done for surgery on links with an arbitrary number of components ℓ. Recall that we assume the framing vectors $\Lambda_1,\dots,\Lambda_\ell$ to be linearly independent in $\mathbb{R}^\ell$ (ie $Y_\Lambda(L)$ is a rational homology sphere). Let $x_1,\dots,x_\ell$ be the coordinates in $\mathbb{R}^\ell\cong H_1(Y-L;\mathbb{R})$. The coordinate hyperplanes $x_i=0$ split $\mathbb{R}^\ell$ into 2^ℓ "hyperquadrants". We construct a "skewed hyperbox" Q (analogous to the parallelogram Q in the case $\ell = 2$), with one vertex in each hyperquadrant, as follows. The vertices of Q are P_ε for $\varepsilon \in \mathbb{E}_\ell = \{0,1\}^\ell$, with coordinates satisfying

$$x_i(P_\varepsilon) > b \quad \text{if } \varepsilon_i = 0, \qquad x_i(P_\varepsilon) < -b \quad \text{if } \varepsilon_i = 1.$$

Thus, the skewed hyperbox Q contains the hypercube $[-b,b]^\ell$, where b is the constant in Lemma 10.1. Further, we require the edges of Q to be parallel to the vectors Λ_i.

Let

$$F_{i,\sigma}, \quad i \in \{1,\dots,\ell\},\ \sigma \in \{-1,1\},$$

be the hyperface of Q that lies completely in the half-space given by $\sigma x_i > 0$. In other words, $F_{i,\sigma}$ has as vertices all P_ε with $(-1)^{\varepsilon_i} = \sigma$.

We can truncate the complex $\mathcal{C}^-(\mathcal{H},\Lambda)$ to obtain a quasi-isomorphic one $\mathcal{C}_Q$, roughly supported in the skewed hyperbox Q. This truncated complex $\mathcal{C}_Q = \mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle$ is obtained from $\mathcal{C}^-(\mathcal{H},\Lambda)$ after canceling some acyclic subcomplexes and quotient complexes, one for each face $F_{i,\sigma}$. We use Lemma 10.1 to show acyclicity, along the same lines as in the case $\ell = 2$. Note that, near the faces $F_{i,+1}$, the truncation is done exactly along the boundaries of Q, while near the faces $F_{i,-1}$, we allow some terms to spill in or out of Q, ie instead of requiring (locally) that $\boldsymbol{s} \in Q$, we have a requirement of the form $\boldsymbol{s} - \varepsilon_i \Lambda_i \in Q$. This allows for the cancellation of the terms outside $\mathcal{C}_Q$. We leave the verification of the details to the interested reader.

Finally, let us turn to the case when the framing matrix Λ is degenerate, so that $b_1(Y_\Lambda(L)) > 0$. Then one can still truncate each complex $\mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u})$, corresponding to a specific Spin^c structure $\mathfrak{u}$ on $Y_\Lambda(L)$. The truncations is done in the same way as in the case $b_1(Y_\Lambda(L)) = 0$, but we cut only in $\ell - b_1(Y_\Lambda(L))$ directions.

In principle, the complex $\mathcal{C}^-(\mathcal{H},\Lambda)$ is a direct product over $\mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u})$, and there are infinitely many Spin^c structures $\mathfrak{u}$. Nevertheless, Theorem 2.3 in [Ozsváth and Szabó 2004c], together with (6), implies that there are only finitely many Spin^c structures $\mathfrak{u}$ for which $\boldsymbol{HF}^-(Y_\Lambda(L),\mathfrak{u}) \neq 0$. In addition, one can find an a priori bound (in terms of a suitable Heegaard diagram) to determine the range of possible $\mathfrak{u}$ with nonzero Floer homology. By only taking those particular $\mathfrak{u}$ in the direct product $\prod \mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u})$ we obtain a chain complex quasi-isomorphic to the original $\mathcal{C}^-(\mathcal{H},\Lambda)$, which we can then truncate to arrive at a quasi-isomorphic finite direct product. We denote the resulting chain complex with finite support by $\mathcal{C}^-(\mathcal{H},\Lambda)\langle b\rangle$, depending on $b \gg 0$.

10.2 Vertical truncation

This section is an analogue of [Ozsváth and Szabó 2008b, Section 2.7], with $\boldsymbol{HF}^-$ replacing HF^+, and with the use of possibly several U variables, as well as nontorsion Spin^c structures.

Let C be a chain complex over $\mathcal{R} = \mathbb{F}[[U_1, \dots, U_\ell]]$, with a relative $\mathbb{Z}/2N\mathbb{Z}$-grading, where N is a nonnegative integer and each U_i has degree -2. Let δ be a positive integer. Let $\mathcal{R}^\delta$ be the quotient of $\mathcal{R}$ by the ideal generated by $U_i^\delta, i = 1, \dots, \ell$. We then denote by C^δ the complex $C \otimes_{\mathcal{R}} \mathcal{R}^\delta$. Further, if $F: C \to D$ is a map between chain complexes over $\mathcal{R}$, we denote by F^δ the corresponding map between C^δ and D^δ. The procedure of replacing C by C^δ, or F by F^δ is referred to as *vertical truncation*.

If $\delta' \geq \delta$, note that there is a natural projection $\mathcal{R}^{\delta'} \to \mathcal{R}^\delta$, which gives a map $C^{\delta'} \to C^\delta$. Set

$$H_*^{\delta \leftarrow \delta'}(C) = \mathrm{Im}(H_*(C^{\delta'}) \to H_*(C^\delta)). \tag{108}$$

Definition 10.2 A chain complex C over $\mathcal{R} = \mathbb{F}[[U_1, \dots, U_\ell]]$ is said to be of *torsion* $\boldsymbol{CF}^-$ *type* if it admits an absolute $\mathbb{Q}$-grading and a relative $\mathbb{Z}$-grading, and it is quasi-isomorphic (over $\mathcal{R}$) to a finitely generated, free chain complex over $\mathcal{R}$.

Remark 10.3 The prototype of a chain complex of torsion $\boldsymbol{CF}^-$ type is the Heegaard Floer complex $\boldsymbol{CF}^-(Y, \mathfrak{u})$, where Y is a three-manifold and $\mathfrak{u}$ a torsion Spin^c over Y; see [Ozsváth and Szabó 2004b; 2004c; 2006]. Also, when $\mathfrak{u}$ is torsion, the complex $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})$ constructed in Section 9.2 is of $\boldsymbol{CF}^-$ type, because it is quasi-isomorphic to a (finitely generated, free) horizontally truncated complex as in Section 10.1.

Note that if C is of torsion $\boldsymbol{CF}^-$ type, then the homology $H_k(C)$ vanishes for $k \gg 0$. Further, as an $\mathcal{R}$-module, $H_*(C)$ admits a quotient module $H_{\geq k} = \bigoplus_{i \geq k} H_i(C)$ for any $k \in \mathbb{Z}$.

In the lemma below, the symbol $\cong$ denotes isomorphism of $\mathcal{R}$-modules.

Lemma 10.4 (a) *Let C be a complex of torsion $\boldsymbol{CF}^-$ type. Then, for any $k \in \mathbb{Q}$, there exists a constant d such that, for all integers $\delta \geq d$, we have $H_{\geq k}(C^\delta) \cong H_{\geq k}(C)$.*

(b) *If A and B are chain complexes of torsion $\boldsymbol{CF}^-$ type satisfying $H_*(A^\delta) \cong H_*(B^\delta)$ for all $\delta \gg 0$, then $H_*(A) \cong H_*(B)$.*

(c) *More generally, suppose we are given integers $a, b \geq 0$. If A and B are chain complexes of torsion $\boldsymbol{CF}^-$ type satisfying*

$$H_*(A^\delta) \otimes (H_{*+2\delta-1}(S^{2\delta-1}))^{\otimes a} \cong H_*(B^\delta) \otimes (H_{*+2\delta-1}(S^{2\delta-1}))^{\otimes b}$$

for all $\delta \gg 0$, then $H_*(A) \cong H_*(B)$. (The tensor products are taken over $\mathbb{F}$.)

Proof Part (a) is similar to [Ozsváth and Szabó 2008b, proof of Lemma 2.7]. We consider the short exact sequence

$$0 \to C \xrightarrow{U_1^\delta} C \to C/U_1^\delta C \to 0. \tag{109}$$

For a given k, we choose d such that $H_i(C) = 0$ for $i \geq k + 2d - 1$. The induced long exact sequence then gives an isomorphism $H_{\geq k}(C) \cong H_{\geq k}(C/U_1^\delta C)$ for $\delta \geq d$. Iterate this argument, replacing C with $C/U_1^\delta C$ and U_1 with U_2, then use U_3, etc.

Parts (b) and (c) follow from (a), combined with the fact that $H_{\geq k}(A) \cong H_{\geq k}(B)$ for all k implies $H_*(A) \cong H_*(B)$. For (c), we also need to note that $H_*(A^\delta)$ and $H_*(A^\delta) \otimes (H_{*+2\delta-1}(S^{2\delta-1}))^{\otimes a}$ have the same quotient modules corresponding to degrees $\geq k$ for any δ sufficiently large compared to k, and that a similar statement applies to B^δ. □

Remark 10.5 If C is finitely generated, we can estimate the value of d in Lemma 10.4(a) as follows: if m is the maximal degree of the generators of C_*, we can choose $d > \frac{1}{2}(m-k)$.

Definition 10.6 A chain complex C over $\mathcal{R} = \mathbb{F}[\![U_1, \dots, U_\ell]\!]$ is said to be of *nontorsion* $\boldsymbol{CF}^-$ *type* if it admits a relative $\mathbb{Z}/2N\mathbb{Z}$-grading for some $N > 0$ and it is quasi-isomorphic (over $\mathcal{R}$) to a finitely generated, free chain complex over $\mathcal{R}$.

Lemma 10.7 *Let Y be a three-manifold and $\mathfrak{u}$ a nontorsion* Spin^c *structure over Y. Then the Heegaard Floer complex $\boldsymbol{CF}^-(Y, \mathfrak{u})$ is of nontorsion $\boldsymbol{CF}^-$ type, and there exists $d \geq 0$ such that $U_i^d\, \boldsymbol{HF}^-(Y, \mathfrak{u}) = 0$ for all i.*

Proof The fact that $\boldsymbol{CF}^-(Y, \mathfrak{u})$ is of nontorsion $\boldsymbol{CF}^-$ type was established in [Ozsváth and Szabó 2004b].

For the second statement, note that all U_i's act the same way on homology; see [Ozsváth and Szabó 2008a]. Let us denote their common action by U. We need to check that this action is nilpotent. Indeed, $\boldsymbol{HF}^\infty(Y, \mathfrak{u}) = 0$ by (5). Since $\boldsymbol{HF}^\infty$ is the ring of fractions of $\boldsymbol{HF}^-$ with respect to (U), the action of U on $\boldsymbol{HF}^-$ must be nilpotent. □

Lemma 10.8 (a) *Let C be a complex of nontorsion $\boldsymbol{CF}^-$ type. Suppose there exists $d \geq 0$ such that $U_i^d H_*(C) = 0$ for all i. Then, for all integers $\delta \geq d$ and $\delta' \geq \delta + d$, we have an isomorphism of relatively graded $\mathcal{R}$-modules*

$$H_*^{\delta \leftarrow \delta'}(C) \cong H_*(C). \tag{110}$$

(b) *Let A and B be relatively $\mathbb{Z}/2N\mathbb{Z}$-graded chain complexes of nontorsion $\boldsymbol{CF}^-$ type satisfying $H_*^{\delta \leftarrow \delta'}(A) \cong H_*^{\delta \leftarrow \delta'}(B)$ for all $\delta' \geq \delta \gg 0$. Suppose there exists $d \geq 0$ such that $U_i^d H_*(A) = 0$ for all i. Then $H_*(A) \cong H_*(B)$, as relatively graded $\mathcal{R}$-modules.*

Proof (a) Note that the short exact sequences (109) for δ' and δ fit into a commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & C \arrow[r, "U_1^{\delta'}"] \arrow[d, "U_1^{\delta'-\delta}"'] & C \arrow[r] \arrow[d, equal] & C/U_1^{\delta'}C \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & C \arrow[r, "U_1^{\delta}"] & C \arrow[r] & C/U_1^{\delta}C \arrow[r] & 0
\end{tikzcd}$$

At the level of homology, this produces the commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & H_*(C) \arrow[r] \arrow[d, equal] & H_*(C/U_1^{\delta'}C) \arrow[r] \arrow[d] & H_*(C) \arrow[r] \arrow[d, "U_1^{\delta'-\delta}"] & 0 \\
0 \arrow[r] & H_*(C) \arrow[r] & H_*(C/U_1^{\delta}C) \arrow[r] & H_*(C) \arrow[r] & 0
\end{tikzcd}$$

The third vertical arrow is zero by hypothesis. It follows that $\mathrm{Im}(H_*(C/U_1^{\delta'}C) \to H_*(C/U_1^{\delta}C)) \cong H_*(C)$. By iterating this argument ℓ times, we obtain (110).

(b) If U_i^d annihilates $H_*(A)$, the long exact sequence in homology associated to (109) implies that U_i^d also annihilates $H_*(A^\delta)$, for all $\delta \gg 0$. Since $H_*(A^\delta) \cong H_*(B^\delta)$ by hypothesis (where we chose $\delta' = \delta$), $U_i^d H_*(B^\delta) = 0$ for $\delta \gg 0$.

We claim that this implies $U_i^d H_*(B) = 0$ for all i. Let us explain the argument in the case $\ell = 1$. The long exact sequence in homology

$$\cdots \to H_*(B) \xrightarrow{U_1^\delta} H_*(B) \to H_*(B^\delta) \to \cdots$$

implies that $M = U_1^d H_*(B)$ is in $U_1^\delta H_*(B) = U_1^{\delta-d} M$ for all large δ. Since M is a finitely generated module over the local Noetherian ring $\mathcal{R} = \mathbb{F}[\![U_1, \dots, U_\ell]\!]$, Krull's theorem implies that $\bigcap_i U_1^i M = 0$; see [Atiyah and Macdonald 1969, Corollary 10.19]. Thus $U_1^d H_*(B) = 0$. Iterating this argument ℓ times produces the same conclusion for an arbitrary number of U_i variables.

The claim that $H_*(A) \cong H_*(B)$ now follows by applying (a) to both A and B. □

Lemma 10.9 *Let C be a complex of nontorsion $\boldsymbol{CF}^-$ type such that all U_i act the same on homology, and $U^d H_*(C) = 0$ for some $d \geq 0$, where U denotes the common U_i action. Suppose* $\dim_{\mathbb{F}}(H_*(C^{\delta+1})) = \dim_{\mathbb{F}}(H_*(C^\delta))$ *for some $\delta \geq 1$. Then $H_*^{\delta \leftarrow 2\delta}(C) \cong H_*(C)$.*

Proof Explicitly, the homology $H_*(C)$ is a direct sum $\bigoplus_j (\mathbb{F}[\![U]\!]/U^{k_j})$, so that $\dim_{\mathbb{F}}(H_*(C)) = \sum k_j$. The long exact sequence on homology associated to (109), iterated ℓ times, implies that $\dim_{\mathbb{F}}(H_*(C^\delta)) = 2^\ell \sum \min(k_j, \delta)$. If $\dim_{\mathbb{F}}(H_*(C^{\delta+1})) = \dim_{\mathbb{F}}(H_*(C^\delta))$, we must have $\delta \geq k_j$ for all j, which means that $U_i^\delta H_*(C) = 0$. The claim now follows from Lemma 10.8(a). □

Let us apply this discussion to the complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$ from (102). Suppose we understand the chain groups and differentials and we want to compute its homology. First, we decompose the complex into terms of the form $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})$, according to Spin^c structures $\mathfrak{u}$. Then we apply the horizontal truncation from Section 10.1 and get quasi-isomorphic complexes $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})\langle b\rangle$. If $\mathfrak{u}$ is torsion, the complex $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})\langle b\rangle$ is finite-dimensional in each given degree, so we can compute its homology. Alternatively, we could replace it by a quasi-isomorphic, finite-dimensional complex $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})\langle b\rangle^\delta$ (see Lemma 10.4), where δ can be estimated as in Remark 10.5. If $\mathfrak{u}$ is nontorsion, the complex $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})\langle b\rangle$ is of nontorsion $\boldsymbol{CF}^-$ type, and its homology is annihilated by a power of the U_i's; see Theorem 9.6 and Lemma 10.7. We then start computing the homology of the complexes $\mathcal{C}^-(\mathcal{H}, \Lambda, \mathfrak{u})\langle b\rangle^\delta$, and let δ increase by one until we find that the complexes for δ and $\delta+1$ have the same total rank. By Lemma 10.9, $H_*(\mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u})) \cong H_*^{\delta \leftarrow 2\delta}(\mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u}))$. The latter homology group can be computed from finite-dimensional complexes.

10.3 A combined truncation

Pick $\delta > 0$ and consider the vertically truncated complex $\mathcal{C}^\delta = \mathcal{C}^-(\mathcal{H},\Lambda)^\delta$. We could apply the horizontal truncation procedure from Section 10.1 to obtain a quasi-isomorphic, finite-dimensional complex $\mathcal{C}^-(\mathcal{H},\Lambda)^\delta\langle b\rangle$. In this section we describe a different way of doing horizontal truncation. This new procedure cannot be applied directly to $\mathcal{C}^-(\mathcal{H},\Lambda)$; it is essential to do the vertical truncation by δ first. The new horizontal truncation is an intermediate step towards constructing the folded truncation in Section 10.4.

We use the notation introduced at the beginning of Section 10, ie we denote by $\mathcal{C}_{\boldsymbol{s}}^{\varepsilon,\delta}$ the factors of $\mathcal{C}^\delta$, and by $\mathcal{D}_{\varepsilon^0,\boldsymbol{s}}^{\varepsilon,\varepsilon',\delta}$ the maps that form the differential. The property that distinguishes the vertically truncated complex $\mathcal{C}^\delta$ from $\mathcal{C}$ is the following:

Lemma 10.10 *Fix $\delta > 0$. Then there is a constant $b^\delta > 0$ such that, for any $i = 1,\dots,\ell$, the map $\mathcal{D}_{\varepsilon^0,\boldsymbol{s}}^{\varepsilon,\varepsilon',\delta}$ is a quasi-isomorphism provided that either*

- *$\varepsilon = \tau_i$ (ie $\varepsilon_i = 1$ and $\varepsilon_j = 0$ for $j \neq i$), $\varepsilon' = 0$, and $s_i > b^\delta$, or*
- *$\varepsilon = \varepsilon' = \tau_i$ and $s_i < -b^\delta$;*

and, further, $\mathcal{D}_{\varepsilon^0,\boldsymbol{s}}^{\varepsilon,\varepsilon',\delta} = 0$ provided that either

- *$\varepsilon_i = \varepsilon'_i = 1$ and $s_i > b^\delta$, or*
- *$\varepsilon_i = 1, \varepsilon'_i = 0$, and $s_i < -b^\delta$.*

Proof For the first part of the statement, it suffices to make sure that $b^\delta \geq b$, where b is the constant from Lemma 10.1. For the second statement (about the triviality of the respective maps), observe that, for example, $\varepsilon_i = \varepsilon'_i = 1$ means that $i \in I_-(\vec{L},\vec{N})$, and we are asked to show that $\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}$ is zero. This is true because by (100), one of the factors of $\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}$ is the map $\mathcal{I}^{\vec{N}}_{\psi^M(\boldsymbol{s})}$. This "inclusion" is the zero map because it contains a large power of U_i see (24), and that power is set to zero in the vertical truncation. The case $\varepsilon_i = 1, \varepsilon'_i = 0$ is similar. □

Let us fix some $\zeta = (\zeta_1,\dots,\zeta_\ell) \in \mathbb{R}^\ell$ such that the values ζ_i are very close to zero, and linearly independent over $\mathbb{Q}$. We let $P_{\mathbb{R}}(\Lambda) \subset \mathbb{R}^\ell$ be the hyperparallelepiped with vertices

$$\zeta + \tfrac{1}{2}(\pm\Lambda_1 \pm \Lambda_2 \pm \cdots \pm \Lambda_\ell)$$

for all possible choices of signs. This is a fundamental domain for $\mathbb{R}^\ell/H(L,\Lambda)$, where $H(L,\Lambda)$ is the lattice generated by the vectors Λ_i, as in Section 9.2. Let $P(\Lambda)$ be the collection of points in the lattice $\mathbb{H}(L)$ that also lie in $P_{\mathbb{R}}(\Lambda)$. Because of our choice of ζ, there are no lattice points on the boundary of $P_{\mathbb{R}}(\Lambda)$. Therefore,

$$P(\Lambda) \cong \mathbb{H}(L)/H(L,\Lambda). \tag{111}$$

In terms of the standard basis of $H_1(Y-L)$, we write

$$\Lambda_i = (c_{i1}, \dots, c_{i\ell}),$$

where $c_{ii} = \lambda_i$ is the framing coefficient on the component L_i, and c_{ij} is the linking number between $\vec{L}_i$ and $\vec{L}_j$ for $i \neq j$.

Pick $m_i \gg 0$, for $i = 1, \dots, \ell$, and let $\widetilde{\Lambda}$ be obtained from Λ by increasing the framing coefficients by m_i; that is,

$$\widetilde{\Lambda}_i = (\tilde{c}_{i1}, \dots, \tilde{c}_{i\ell}),$$

with

$$\tilde{c}_{ij} = \begin{cases} c_{ii} + m_i & \text{if } i = j, \\ c_{ij} & \text{if } i \neq j. \end{cases}$$

For each $\varepsilon \in \mathbb{E}_\ell$, consider the hyperparallelepiped $P(\widetilde{\Lambda}, \Lambda, \varepsilon)_{\mathbb{R}} \subset \mathbb{R}^\ell$ with vertices

$$\zeta + \frac{1}{2} \sum_{i=1}^{\ell} (\sigma_i \widetilde{\Lambda}_i + (1-\sigma_i)\varepsilon_i \Lambda_i)$$

over all possible choices of signs $\sigma_i \in \{\pm 1\}$. Set

$$P(\widetilde{\Lambda}, \Lambda, \varepsilon) = P(\widetilde{\Lambda}, \Lambda, \varepsilon)_{\mathbb{R}} \cap \mathbb{H}(L).$$

For $\varepsilon = (0, \dots, 0)$, we recover the old $P(\widetilde{\Lambda})$, while, for $\varepsilon = (1, \dots, 1)$, the hyperparallelepiped $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ is a rectangular hyperbox of size $(m_1, \dots, m_\ell)$.

Set

$$\mathcal{C}^\delta \langle\langle \widetilde{\Lambda} \rangle\rangle = \bigoplus_{\varepsilon \in \mathbb{E}_\ell} \bigoplus_{\boldsymbol{s} \in P(\widetilde{\Lambda}, \Lambda, \varepsilon)} \mathcal{C}_{\boldsymbol{s}}^{\varepsilon, \delta}. \tag{112}$$

See Figure 31 for an example.

Remark 10.11 The hyperparallelepipeds $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ behave nicely with respect to the maps $\psi^M \colon \mathbb{H}(L) \to \mathbb{H}(L-M)$ as defined in Section 3.7. Indeed, consider the link

$$M = \bigcup_{\varepsilon_i = 1} L_i$$

with the orientation induced from $\vec{L}$. Then ψ^M takes $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ exactly to the hyperparallelepiped $P(\widetilde{\Lambda}|_{L-M})$, and it does so in an $m(M)$-to-one fashion, where

$$m(M) = \prod_{L_i \subseteq M} m_i.$$

Therefore, we can rewrite (112) as

$$\mathcal{C}^\delta \langle\langle \widetilde{\Lambda} \rangle\rangle = \bigoplus_{M \subseteq L} \bigoplus_{\boldsymbol{s} \in P(\widetilde{\Lambda}|_{L-M})} \bigoplus^{m(M)} (\mathfrak{A}^{-,\delta}(\mathcal{H}^{L-M}, \boldsymbol{s})). \tag{113}$$

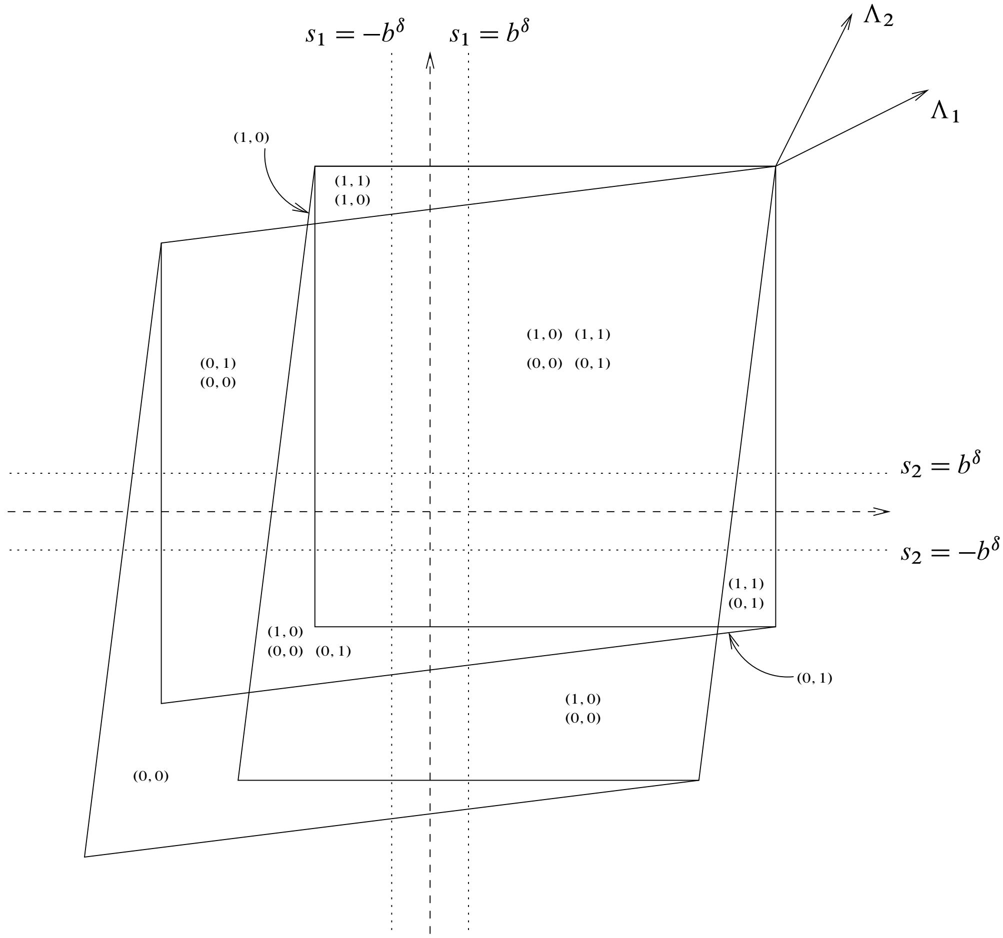

Figure 31: The complex $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$ for a two-component link. We show here the combined truncation corresponding to surgery on a framed link with, say, $\Lambda_1 = (200, 100)$, $\Lambda_2 = (100, 200)$, $m_1 = m_2 = 600$ and $b^\delta = 50$. The four parallelograms in the picture are $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ for $\ell = 2$, $\varepsilon \in \{0, 1\}^2$, and $\widetilde{\Lambda}$ and Λ fixed. In each region we mark the values of ε such that the corresponding $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ contains the region. This tells us which terms $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon,\delta}$ form the truncated complex $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$. The framing $\widetilde{\Lambda}$ is sufficiently large compared to b^δ, as explained in the proof of Proposition 10.12.

Proposition 10.12 *Fix $\delta > 0$. If we pick the values m_i sufficiently large, the direct sum $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$, equipped with the restriction of the differential $\mathscr{D}^{-,\delta}$, forms a chain complex quasi-isomorphic to $(\mathscr{C}^\delta, \mathscr{D}^{-,\delta})$.*

Proof Let $P_{\mathbb{R}}$ be a hyperparallelepiped in $\mathbb{R}^\ell$, with vertices V^σ, for $\sigma = (\sigma_1, \dots, \sigma_\ell) \in \{-1, 1\}^\ell$. We assume that $\sigma_i V_i^\sigma > 0$ for all σ and i; that is, each vertex lies in the hyperquadrant in $\mathbb{R}^\ell$ that corresponds to σ. For each

$$\omega = (\omega_1, \dots, \omega_\ell) \in \{-1, 0, 1\}^\ell,$$

we define a subset $P_{\mathbb{R}}[\omega] \subset \mathbb{R}^\ell$ as follows. First, define a completion of ω to be a vector $\sigma \in \{-1, 1\}^\ell$ such that $\sigma_i = \omega_i$ whenever $|\omega_i| = 1$. In other words, a completion of ω is a vector in which we

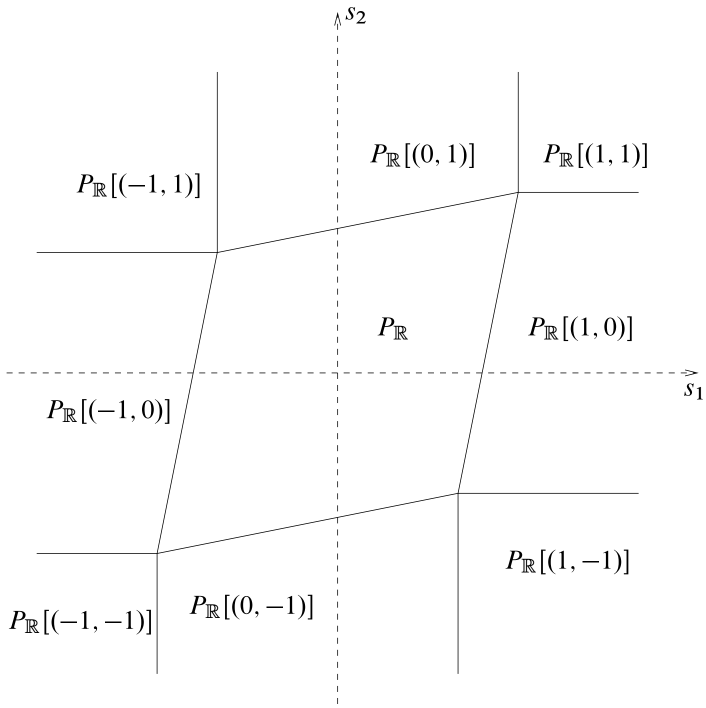

Figure 32: A decomposition of the plane into nine regions. We show here an example of the subsets $P_{\mathbb{R}}[\omega] \subset \mathbb{R}^2$ for $\omega \in \{-1,0,1\}^2$.

replace the zero entries of ω with $+1$ or -1. Let $P_{\mathbb{R}}^{\omega}$ be the subparallelepiped of $P_{\mathbb{R}}$ with vertices V^{σ}, where σ runs over all possible completions of ω. Further, given a vector with only nonnegative entries $\boldsymbol{t} = (t_1, \dots, t_\ell) \in [0,\infty)^\ell$, we define $\boldsymbol{t} * \omega$ to be the pointwise product

$$\boldsymbol{t} * \omega = (t_1\omega_1, \dots, t_\ell\omega_\ell) \in \mathbb{R}^\ell.$$

Set

$$P_{\mathbb{R}}[\omega] = \bigcup_{\boldsymbol{t} \in [0,\infty)^\ell} (P_{\mathbb{R}}^{\omega} + \boldsymbol{t} * \omega).$$

Note that $P_{\mathbb{R}}[(0,\dots,0)] = P_{\mathbb{R}}$. We have a decomposition

$$\mathbb{R}^\ell = \bigcup_{\omega \in \{-1,0,1\}^\ell} P_{\mathbb{R}}[\omega].$$

See Figure 32 for the case $\ell = 2$.

Suppose there are no points in the lattice $\mathbb{H}(L)$ on the boundary of any $P_{\mathbb{R}}[\omega]$. Then, letting $P[\omega] = P_{\mathbb{R}}[\omega] \cap \mathbb{H}(L)$, we get a decomposition as a disjoint union

$$\mathbb{H}(L) = \coprod_{\omega \in \{-1,0,1\}^\ell} P[\omega].$$

We can apply this to any of the hyperparallelepipeds $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ and obtain a decomposition of $\mathbb{H}(L)$ for each ε.

For $\omega \in \{-1,0,1\}^{\ell}$, consider the direct sum

$$\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega] = \bigoplus_{\varepsilon\in\mathbb{E}_\ell}\ \prod_{\boldsymbol{s}\in P(\widetilde{\Lambda},\Lambda,\varepsilon)[\omega]} \mathcal{C}_{\boldsymbol{s}}^{\varepsilon,\delta},$$

so that

$$\mathcal{C}^{\delta} = \bigoplus_{\omega\in\{-1,0,1\}^{\ell}} \mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]. \tag{114}$$

We choose the values m_i sufficiently large such that whenever V is a vertex of some $P(\widetilde{\Lambda},\Lambda,\varepsilon)$ with $\sigma_i V_i > 0$, we in fact have $\sigma_i V_i > b^{\delta}$. (Here b^{δ} is the constant in Lemma 10.10.)

Starting from here, the idea is to use Lemma 10.10 to show that the terms $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]$ in the direct sum (114) produce acyclic complexes for $\omega \neq (0,\dots,0)$. We eliminate them one by one, beginning with the ones with $|\omega| = \sum \omega_i = \ell$, then those with $|\omega| = \ell-1$, all the way to $|\omega| = 1$, after which we are left only with $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[(0,\dots,0)] = \mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle$.

We start with ω such that $|\omega| = \ell$, ie all values ω_i are 1 or -1. Let $\varpi = (\varpi_1,\dots,\varpi_\ell)$ be the vector consisting of the values $\varpi_i = \frac{1}{2}(1-\omega_i)$ for all i. Each $P(\widetilde{\Lambda},\Lambda,\varepsilon)[\omega]$ is obtained from the hyperquadrant determined by ω by translating by an amount that depends on ε. More precisely, if V is the unique vertex of $P(\widetilde{\Lambda},\Lambda,0)[\omega]$, then the vertex of $P(\widetilde{\Lambda},\Lambda,\varepsilon)[\omega]$ is $V + (\varpi * \varepsilon)\cdot\Lambda$.

The second part of Lemma 10.10 implies that the maps $D^{\varepsilon,\varepsilon',\delta}_{\varepsilon^0,\boldsymbol{s}}$ are trivial unless $\varepsilon' = \varpi * \varepsilon$. It follows from here that $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]$ is a direct summand of $\mathcal{C}^{\delta}$ as a chain complex, ie it is preserved by the differential. Moreover, since by definition $D^{\varepsilon,\varepsilon',\delta}_{\varepsilon^0,\boldsymbol{s}}$ maps $\mathcal{C}^{\varepsilon^0,\delta}_{\boldsymbol{s}}$ into $\mathcal{C}^{\varepsilon^0+\varepsilon,\delta}_{\boldsymbol{s}+\varepsilon'\cdot\Lambda}$, we deduce that the differential on $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]$ preserves the quantity

$$\boldsymbol{s} - (\varpi * \varepsilon)\cdot\Lambda.$$

Hence, the complex $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]$ splits into a direct sum of terms corresponding to values of $\boldsymbol{s} - (\varpi*\varepsilon)\cdot\Lambda$. Each such term is an ℓ-dimensional hypercube that has on its edges maps of the form $D^{\varepsilon,\varepsilon',\delta}_{\varepsilon^0,\boldsymbol{s}}$ with $\varepsilon = \tau_i$. By the first part of Lemma 10.10, all these edge maps are quasi-isomorphisms. Therefore, the respective hypercube complexes are acyclic, and so is $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]$.

We then proceed inductively on $\ell - |\omega|$. At each stage, we have a complex $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\leq|\omega|]$ obtained from $\mathcal{C}^{\delta}$ by canceling the terms with higher $|\omega|$. Each $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\omega]$, for the given $|\omega|$, splits as a direct sum of subcomplexes and quotient complexes of $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle[\leq|\omega|]$. They are acyclic by an application of Lemma 10.10. The claim follows. □

10.4 A folded truncation

In this section we describe the variant of truncation which will be useful to us in Section 12. We keep the notation from the previous subsection, with $\mathcal{C}^{\delta} = \mathcal{C}^{-}(\mathcal{H},\Lambda)^{\delta}$ and $\widetilde{\Lambda}$ as before. The folded truncation $\mathcal{C}^{\delta}\{\widetilde{\Lambda}\}$ that we are about to describe will be isomorphic to $\mathcal{C}^{\delta}\langle\!\langle\widetilde{\Lambda}\rangle\!\rangle$ as an $\mathcal{R}$-module, but its differential will contain some additional maps.

For $\varepsilon = \varepsilon(M) \in \mathbb{E}_\ell$ and $\boldsymbol{s} \in \mathbb{H}(L)$, we introduce the notation $[\boldsymbol{s}]_\varepsilon$ to denote the translate of $\boldsymbol{s}$ that lies in the hyperparallelepiped $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$, where we allow translations by multiples of $\widetilde{\Lambda}_i - \varepsilon_i \Lambda_i$ for any $i = 1, \dots, \ell$. In other words, we consider the tessellation of $\mathbb{R}^\ell$ by translates of $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$, and "fold" all the hyperparallelepipeds in this tessellation onto the fundamental domain $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ using translation. For simplicity, we write $\mathcal{C}^\varepsilon_{[\boldsymbol{s}]}$ for the complex $\mathcal{C}^\varepsilon_{[\boldsymbol{s}]_\varepsilon}$.

Observe that if $\boldsymbol{s} \in \mathbb{H}(L)$ is sufficiently close to a face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$ (say, at a distance of no more than the sum of the absolute values of the entries of Λ) and the values m_i were chosen sufficiently large, then $[\boldsymbol{s}]_\varepsilon$ is also close to a (possibly different) face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon)$. Further, in this case in $\mathcal{C}^\varepsilon_{\boldsymbol{s}} = \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$ all entries in $\psi^M(\boldsymbol{s})$ are either sufficiently large or sufficiently small that they can be replaced with $\pm\infty$, so $\mathcal{C}^\varepsilon_{\boldsymbol{s}}$ is naturally isomorphic to $\mathfrak{A}^-(\mathcal{H}^\varnothing)$, using the data in the complete system $\mathcal{H}$. The same is true of $\mathcal{C}^\varepsilon_{[\boldsymbol{s}]}$, so we have an identification between $\mathcal{C}^\varepsilon_{\boldsymbol{s}}$ and $\mathcal{C}^\varepsilon_{[\boldsymbol{s}]}$.

With this in mind, we define

$$\mathcal{C}\{\widetilde{\Lambda}\} = \bigoplus_{\varepsilon \in \mathbb{E}_\ell} \bigoplus_{\boldsymbol{s} \in P(\widetilde{\Lambda}, \Lambda, \varepsilon)} \mathcal{C}^\varepsilon_{\boldsymbol{s}},$$

with the differential $\mathcal{D}^-\{\widetilde{\Lambda}\}$ consisting of all maps

$$\mathcal{D}^{\varepsilon,\varepsilon'}_{\varepsilon^0,\boldsymbol{s}}\{\widetilde{\Lambda}\} = \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})} : \mathcal{C}^{\varepsilon^0}_{\boldsymbol{s}} \to \mathcal{C}^{\varepsilon^0+\varepsilon}_{[\boldsymbol{s}+\varepsilon'\cdot\Lambda]} \tag{115}$$

for $\varepsilon^0 = \varepsilon(M)$, $\varepsilon = \varepsilon(N)$ and $\varepsilon' \in \mathbb{E}_\ell$ such that $\varepsilon^0 + \varepsilon \in \mathbb{E}_\ell$ and $\varepsilon' \le \varepsilon$.

Of course, it may happen that the value $\boldsymbol{s}$ is in the hyperparallelepiped $P(\widetilde{\Lambda}, \Lambda, \varepsilon^0)$, but $\boldsymbol{s} + \varepsilon' \cdot \Lambda$ is not in $P(\widetilde{\Lambda}, \Lambda, \varepsilon^0 + \varepsilon)$. In such cases, note that $\boldsymbol{s}$ must be close enough to a face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon^0 + \varepsilon)$; hence, the target of the map $\mathcal{D}^{\varepsilon,\varepsilon'}_{\varepsilon^0,\boldsymbol{s}}$, which is a priori $\mathcal{C}^{\varepsilon^0+\varepsilon}_{\boldsymbol{s}+\varepsilon'\cdot\Lambda}$, is identified with $C^{\varepsilon^0+\varepsilon}_{[\boldsymbol{s}+\varepsilon'\cdot\Lambda]}$, as explained above. The maps $\mathcal{D}^{\varepsilon,\varepsilon'}_{\varepsilon^0,\boldsymbol{s}}$ of this form are called *crossover maps*.

We let $\mathcal{C}^\delta\{\widetilde{\Lambda}\}$ be the vertical truncation (by δ) of $\mathcal{C}\{\widetilde{\Lambda}\}$. Then $\mathcal{C}^\delta\{\widetilde{\Lambda}\}$ differs from the truncation $\mathcal{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$ only in that its differential contains the crossover maps.

Example 10.13 If K is a knot with framing coefficient $\Lambda \in \mathbb{Z}$, recall from (107) from Section 10.1 that we have a horizontal truncation

$$\mathcal{C}\langle b\rangle = \bigoplus_{-b \le s \le b} \mathcal{C}^0_s \oplus \bigoplus_{-b+\Lambda \le s \le b} \mathcal{C}^1_s \subset \mathcal{C}.$$

If we also do vertical truncation by δ, the resulting complex $\mathcal{C}^\delta\langle b\rangle$ is the same as the combined truncation $\mathcal{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$ from Section 10.3, with $\widetilde{\Lambda} = 2b$. When $\Lambda = 1$, the truncation $\mathcal{C}^\delta\langle b\rangle = \mathcal{C}^\delta\langle\langle 2b\rangle\rangle$ is shown in Figure 33. In principle, in the folded truncation $\mathcal{C}^\delta\{2b\}$ we have to add the crossover maps Φ^{+K}_{-b} and Φ^{-K}_{b} shown by dashed lines in Figure 33. (In our notation, for example, Φ^{+K}_{-b} is $\mathcal{D}^{\varepsilon,\varepsilon'}_{\varepsilon^0,\boldsymbol{s}}\{\widetilde{\Lambda}\}$ with $\varepsilon^0 = 0$, $\varepsilon = 1$, $\varepsilon' = 0$ and $\boldsymbol{s} = -b$, so that $[\boldsymbol{s} + \varepsilon' \cdot \Lambda]_{\varepsilon^0+\varepsilon} = [\boldsymbol{s}]_{\varepsilon^0+\varepsilon} = b$.) However, observe that, if b is chosen sufficiently large compared to b^δ, then the maps Φ^{+K}_{-b} and Φ^{-K}_{b} vanish, according to Lemma 10.10. Thus, in this case, the folded truncation is the same as the combined truncation. (In fact, the same is true for any Λ, that is, for any integer surgery on a knot.)

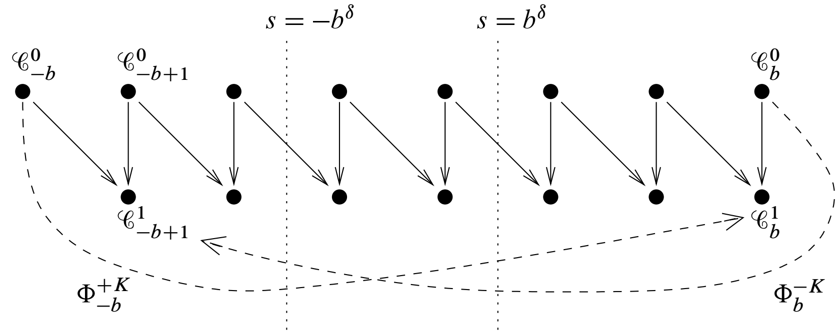

Figure 33: The folded truncation for surgery on a knot K. In this example, $\Lambda = 1$ and $b = 8$. Each black dot represents a complex $\mathscr{C}^\varepsilon_s$, each vertical arrow represents a map Φ^{+K}_s, and each diagonal arrow represents a map Φ^{-K}_s. The dashed arrows are crossover maps that appear in $\mathscr{C}\{2b\}$. They are set to zero in $\mathscr{C}^\delta\{2b\}$ because they originate outside the region between the two dotted lines; compare Lemma 10.10.

Example 10.14 When L is a link of two or more components, there exist crossover maps that remain nontrivial even when vertically truncating by δ. Examples are shown in Figure 34. For concreteness, consider the dashed arrow Φ^{+L_2} starting at the leftmost red triangle in the figure, and ending at the rightmost green triangle. In our notation, this is the crossover map $\mathscr{D}^{\varepsilon,\varepsilon'}_{\varepsilon^0,\boldsymbol{s}}\{\widetilde{\Lambda}\}$ for $\varepsilon^0 = (0,0)$, $\varepsilon = (0,1)$ and $\varepsilon' = (0,0)$, with $\boldsymbol{s}$ in the leftmost red triangle and $[\boldsymbol{s}]_{\varepsilon^0+\varepsilon}$ in the rightmost green triangle.

For future use, let us also write the definition of $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ in a slightly different way. For $\varepsilon = \varepsilon(M) \in \mathbb{E}_\ell$, let $\mathscr{T}^\varepsilon$ be the (finite) polynomial ring over $\mathbb{F}$ in variables T_i for those i such that $\varepsilon_i = 1$ (that is, $L_i \subseteq M$), with relations $T_i^{m_i} = 1$. If $M = L_{i_1} \cup \cdots \cup L_{i_p}$ and $\boldsymbol{k} = (k_1, \ldots, k_p) \in \mathbb{Z}^p$, we formally write

$$T^{\boldsymbol{k}} = \prod_{j=1}^{p} T_{i_j}^{k_j}.$$

Fix a reference value $\boldsymbol{s}^0 = (s^0_1, \ldots, s^0_\ell) \in \mathbb{H}(L)$ with all of its entries less than b^δ in absolute value. In view of Remark 10.11, we have an identification

$$P(\widetilde{\Lambda}, \Lambda, \varepsilon) \xrightarrow{\cong} P(\widetilde{\Lambda}|_{L-M}) \times \mathscr{T}^\varepsilon, \quad \boldsymbol{s} = (s_1, \ldots, s_\ell) \mapsto \Big(\psi^M(\boldsymbol{s}), \prod_{\varepsilon_i=1} T_i^{s_i - s_i^0}\Big).$$

Therefore, we can write

$$\mathscr{C}^\delta\{\widetilde{\Lambda}\} = \bigoplus_{M\subseteq L} \bigoplus_{\boldsymbol{s}\in P(\widetilde{\Lambda}|_{L-M})} \mathfrak{A}^{-,\delta}(\mathscr{H}^{L-M}, \boldsymbol{s}) \otimes \mathscr{T}^\varepsilon(M), \tag{116}$$

with the differential made of terms

$$\mathscr{D}^{\varepsilon,\varepsilon',\delta}_{\varepsilon^0,\boldsymbol{s}}\{\widetilde{\Lambda}\} = T^{(\Lambda|_N)_{\vec{L},\vec{N}}} \cdot \Phi^{\vec{N},\delta}_{\psi^M(\boldsymbol{s})}, \tag{117}$$

where $\varepsilon^0 = \varepsilon(M)$, $\varepsilon = \varepsilon(N)$, and the orientation $\vec{N}$ on N is given by $i \in I_-(\vec{L}, \vec{N})$ if and only if $\varepsilon'_i = 1$. By $(\Lambda|_N)_{\vec{L},\vec{N}}$ we mean the image of $\Lambda_{\vec{L},\vec{N}} = \sum_{i\in I_-(\vec{L},\vec{N})} \Lambda_i \in H_1(Y-L)$ under the natural map $H_1(Y-L) \to H_1(Y-N)$.

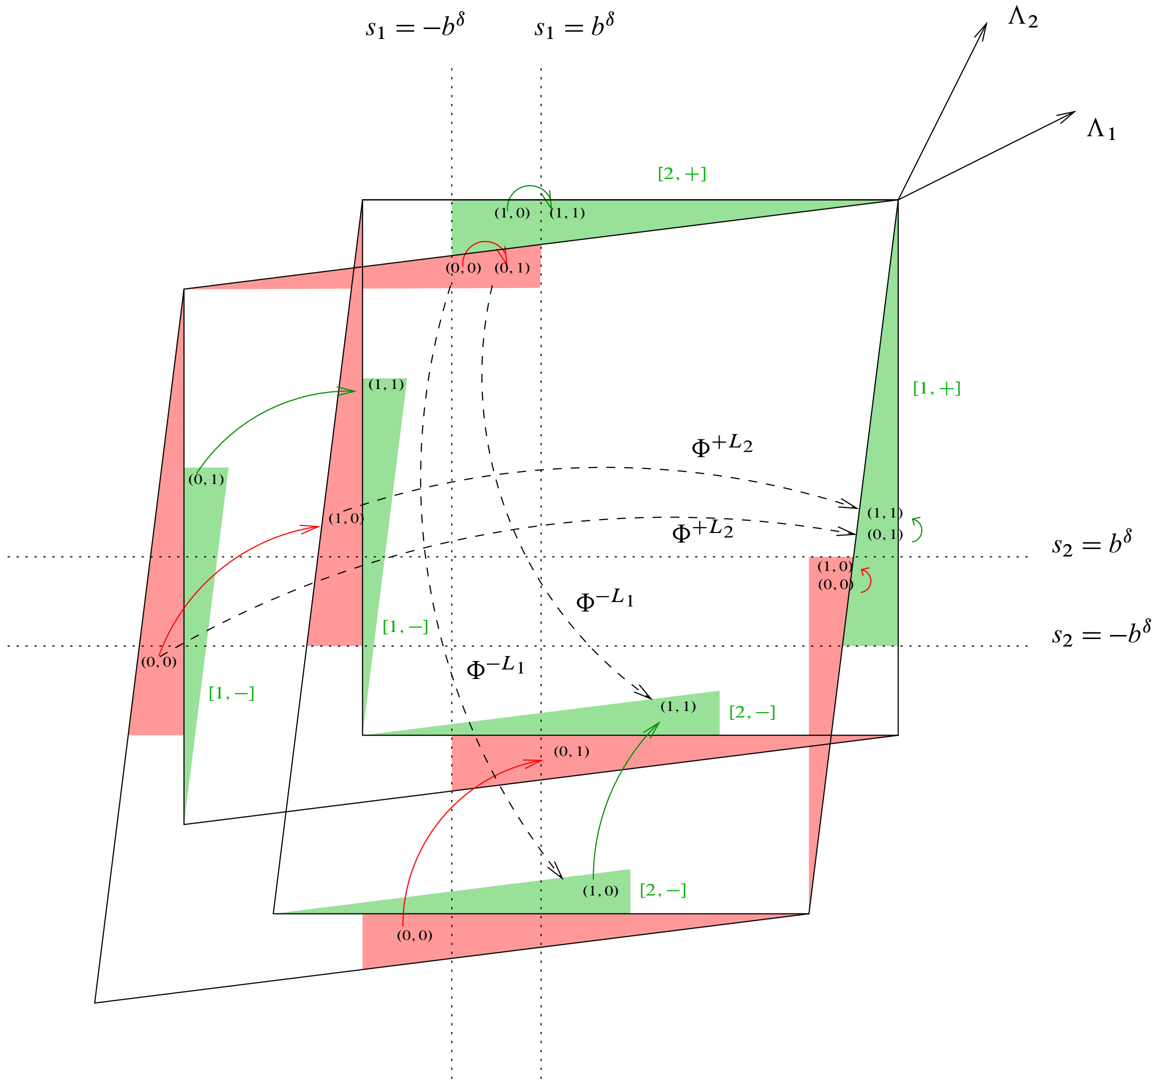

Figure 34: The folded truncation for surgery on a two-component link. We show here the same parallelograms as in Figure 31, corresponding to surgery on a link with $\Lambda_1 = (200, 100)$ and $\Lambda_2 = (100, 200)$. Examples of crossover maps are given by dashed lines. They go from the quotient complex $\mathscr{C}^\delta_{\mathrm{Out}}\{\widetilde{\Lambda}\}$ (indicated by the red regions) to the subcomplex $\mathscr{C}^\delta_{\mathrm{In}}\{\widetilde{\Lambda}\}$ (indicated by the green regions). In addition to the four crossover maps shown by dashed lines, there are four others: two of the form Φ^{+L_1} and two of the form Φ^{-L_2} (not shown here in order to avoid overcrowding the picture). In each green or red region we mark the values ε such that $\mathscr{C}^{\varepsilon,\delta}_{\boldsymbol{s}}$ is part of the respective quotient complex or subcomplex. Further, near each green region we mark (in green square brackets) the edge of the parallelogram that blocks the points in that region: $[1,+]$, $[1,-]$, $[2,+]$ or $[2,-]$; see the proof of Proposition 10.15 for an explanation. The subcomplex $\mathscr{C}^\delta_{\mathrm{In}}\{\widetilde{\Lambda}\}$ is acyclic because its differential contains quasi-isomorphisms (indicated by the green arrows) that cancel its terms in pairs. Similarly, $\mathscr{C}^\delta_{\mathrm{Out}}\{\widetilde{\Lambda}\}$ is acyclic because of the quasi-isomorphisms indicated by the red arrows.

Our main goal is to show:

Proposition 10.15 *Fix $\delta > 0$. If we pick the values m_i sufficiently large, then the complexes $\mathscr{C}^\delta$ and $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ are quasi-isomorphic.*

The proof of Proposition 10.15 will proceed along the following lines. Recall from Proposition 10.12 that $\mathscr{C}^\delta$ is quasi-isomorphic to the truncated complex $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$. Thus, it suffices to show that $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$ and $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ are quasi-isomorphic. The basic idea is that all nontrivial crossover maps go into an acyclic subcomplex of $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$; therefore, deleting them will not change the quasi-isomorphism type of the complex.

In order to make this argument precise, we first introduce the auxiliary complex we need, and show that it is acyclic.

Note that many of the crossover maps in $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ are automatically zero, by an application of Lemma 10.10. In particular, when L is a knot, all crossover maps vanish; compare Example 10.13 and Figure 33. On the other hand, in Figure 34, some crossover maps have to vanish and others do not: for example, near the top right corner of Figure 34, in principle we should have crossover maps of the form Φ^{-L_1} and Φ^{-L_2}; but these are zero, so we did not show them in the figure.

Roughly, the auxiliary complex $\mathscr{C}^\delta_{\mathrm{In}}\{\widetilde{\Lambda}\}$ (to be defined below) will be the target of all crossover maps that are not being automatically set to zero by applying Lemma 10.10. In Figure 34, this complex is indicated by the green regions. (When L is a knot, the auxiliary complex is zero.)

Precisely, we define an *incoming crossover point* to be a pair $(\varepsilon^0, \boldsymbol{s})$ with $\varepsilon^0 \in \mathbb{E}_\ell$ and $\boldsymbol{s} = (s_1, \dots, s_\ell) \in P(\widetilde{\Lambda}, \Lambda, \varepsilon^0)$ satisfying the following property: there exist $\varepsilon, \varepsilon' \in \mathbb{E}_\ell$ such that

- $\varepsilon^0 - \varepsilon \in \mathbb{E}_\ell$ and $\varepsilon' \le \varepsilon$;
- $\boldsymbol{s}' = (s'_1, \dots, s'_\ell) := \boldsymbol{s} - \varepsilon' \cdot \Lambda \notin P(\widetilde{\Lambda}, \Lambda, \varepsilon^0 - \varepsilon)$;
- there is no $i = 1, \dots, \ell$ with $\varepsilon_i = 1$, $\varepsilon'_i = 0$ and $s'_i < -b^\delta$;
- there is no $i = 1, \dots, \ell$ with $\varepsilon_i = 1$, $\varepsilon'_i = 1$ and $s'_i + \sum_{j \ne i} \mathrm{lk}(L_i, L_j)(1 - \varepsilon_j) > b^\delta$.

In other words, we ask for $(\varepsilon^0, \boldsymbol{s})$ to specify the target of a (not a priori trivial) crossover map

$$\mathscr{D}^{\varepsilon, \varepsilon'}_{\varepsilon^0 - \varepsilon', \boldsymbol{s}} : \mathscr{C}^{\varepsilon^0 - \varepsilon}_{\boldsymbol{s}'} \to \mathscr{C}^{\varepsilon^0}_{\boldsymbol{s}}.$$

We denote by $\mathrm{In}(\widetilde{\Lambda}, \Lambda)$ the set of all incoming crossover points. (In Figure 34, the set $\mathrm{In}(\widetilde{\Lambda}, \Lambda)$ is exactly the union of the green regions.) Define the *incoming crossover subcomplex* of $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ to be

$$\mathscr{C}^\delta_{\mathrm{In}}\{\widetilde{\Lambda}\} = \bigoplus_{(\varepsilon, \boldsymbol{s}) \in \mathrm{In}(\widetilde{\Lambda}, \Lambda)} \mathscr{C}^{\varepsilon, \delta}_{\boldsymbol{s}}.$$

Lemma 10.16 *The subcomplex $\mathscr{C}^\delta_{\mathrm{In}}\{\widetilde{\Lambda}\}$ is acyclic.*

Proof Recall that the crossover maps must end near the faces of the hyperparallelepipeds under consideration. Near each face Lemma 10.10 indicates the presence of a particular kind of quasi-isomorphism among the edge maps $\Phi_{\boldsymbol{s}}^{\pm L_i}$. We claim that these quasi-isomorphisms are part of the differential for $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$, and cause its terms to cancel out in pairs. (In Figure 34, the quasi-isomorphisms that produce the cancellation are indicated by the green arrows.)

To make this precise, let $(\varepsilon^{\mathbf{0}}, \boldsymbol{s})$ be an incoming crossover point. The condition $\boldsymbol{s} - \varepsilon' \cdot \Lambda \notin P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}} - \varepsilon)$ means that $\boldsymbol{s} - \varepsilon' \cdot \Lambda$ fails to satisfy at least one of the 2ℓ inequalities describing the hyperparallelepiped $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}} - \varepsilon)$. This inequality comes from one of the 2ℓ codimension one faces of the hyperparallelepiped, and we say that $(\varepsilon^{\mathbf{0}}, \boldsymbol{s})$ is an incoming crossover point *blocked* by that face. One way to distinguish one of the 2ℓ codimension one faces is to say which of the 2ℓ coordinate half-axes of $\mathbb{R}^{\ell}$ it intersects. In general, if P is a hyperparallelepiped like the ones in our construction, the face of P that intersects the positive s_i-axis will be called the $[i, +]$ face, and the face that intersects the negative s_i-axis will be called the $[i, -]$ face.

Suppose an inequality that fails for $\boldsymbol{s} - \varepsilon' \cdot \Lambda$ is the one defining the $[1, +]$ face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}} - \varepsilon)$. Since $\boldsymbol{s}$ does satisfy the inequality defining the $[1, +]$ face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}})$, it must be that $\boldsymbol{s}$ lies close to those faces — where by "close" we mean a distance comparable to the entries of Λ and much smaller than b^{δ} or the m_i's. (An example is when $\boldsymbol{s}$ is in the rightmost vertical green triangle in Figure 34, and $\varepsilon^{\mathbf{0}} = (0, 1)$ or $(1, 1)$; there, $\varepsilon = (0, 1)$ and $\varepsilon' = (0, 0)$, so $\Phi_{[\boldsymbol{s}']}^{+L_2}$ is the crossover map.) Suppose further that $\varepsilon_1^0 = 0$, and let $\hat{\varepsilon}^{\mathbf{0}} \in \mathbb{E}_{\ell}$ be obtained from $\varepsilon^{\mathbf{0}}$ by changing the entry ε_1^0 from 0 to 1. Changing $\varepsilon^{\mathbf{0}}$ into $\hat{\varepsilon}^{\mathbf{0}}$ does not change the $[1, +]$ face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}})$ at all. The same is true for the $[1, +]$ face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}} - \varepsilon)$. Therefore, $(\boldsymbol{s}, \varepsilon^{\mathbf{0}})$ is an incoming crossover point blocked by the $[1, +]$ face if and only if $(\boldsymbol{s}, \hat{\varepsilon}^{\mathbf{0}})$ is. Moreover, since we are near the $[1, +]$ face, the value s_1 is bigger than b^{δ}, so Lemma 10.10 says that

$$\Phi_{\boldsymbol{s}}^{+L_1} : \mathscr{C}_{\boldsymbol{s}}^{\varepsilon^{\mathbf{0}}, \delta} \to \mathscr{C}_{\boldsymbol{s}}^{\hat{\varepsilon}^{\mathbf{0}}, \delta}$$

is a quasi-isomorphism. (This is shown as one of the green arrows in Figure 34.) Consider the subcomplex of $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$ consisting of $\mathscr{C}_{\boldsymbol{s}}^{\varepsilon^{\mathbf{0}}, \delta}$ for all $(\varepsilon^{\mathbf{0}}, \boldsymbol{s}) \in \mathrm{In}(\widetilde{\Lambda}, \Lambda)$ blocked by the $[1, +]$ face. This subcomplex is acyclic, because it consists of two hypercubes of chain complexes (corresponding to $\varepsilon_1^0 = 0$ and $\varepsilon_1^0 = 1$, respectively), connected by a chain map which has quasi-isomorphisms $\Phi_{\boldsymbol{s}}^{+L_1}$ along the edges. We can kill off this acyclic subcomplex and obtain a new complex quasi-isomorphic to $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$, which does not involve any $(\boldsymbol{s}, \varepsilon^{\mathbf{0}})$ blocked by the $[1, +]$ face. Similarly, from this new complex we can kill off the terms blocked by the $[2, +]$ face, then the terms blocked by the $[3, +]$ face, and so on, all without changing the quasi-isomorphism type.

An analogous elimination procedure works for the terms blocked by faces of the type $[i, -]$. Indeed, suppose an inequality that fails for $\boldsymbol{s} - \varepsilon' \cdot \Lambda$ is the one defining the $[1, -]$ face of $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}} - \varepsilon)$. Since $\boldsymbol{s}$ does satisfy the corresponding inequality for $P(\widetilde{\Lambda}, \Lambda, \varepsilon^{\mathbf{0}})$, it must be that $\boldsymbol{s}$ lies close to those faces. (An example of this situation is when $\boldsymbol{s}$ is in one of the two leftmost vertical green triangles in Figure 34, and $\varepsilon^{\mathbf{0}} = (0, 1)$ or $(1, 1)$; there, $\varepsilon = \varepsilon' = (0, 1)$, so $\Phi_{[\boldsymbol{s}']}^{-L_2}$ is the crossover map.) Suppose further that

$\varepsilon_1^0 = 0$, and let $\hat{\varepsilon}^0 \in \mathbb{E}_\ell$ be obtained from ε^0 by changing the entry ε_1^0 from 0 to 1. This time, changing ε^0 into $\hat{\varepsilon}^0$ results in translating the $[1,-]$ face of $P(\widetilde{\Lambda},\Lambda,\varepsilon^0)$ by the vector Λ_1. The same happens to the $[1,-]$ face of $P(\widetilde{\Lambda},\Lambda,\varepsilon^0-\varepsilon)$. Therefore, $(\hat{\varepsilon}^0, \boldsymbol{s}+\Lambda_1)$ is also an incoming crossover point blocked by the $[1,-]$ face. Further, since $s_1 \ll -b^\delta$, Lemma 10.10 says that

$$\Phi_{\boldsymbol{s}}^{-L_1} : \mathscr{C}_{\boldsymbol{s}}^{\varepsilon^0,\delta} \to \mathscr{C}_{\boldsymbol{s}+\Lambda_1}^{\hat{\varepsilon}^0,\delta}$$

is a quasi-isomorphism. Consequently, the terms in $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$ that are blocked by the $[1,-]$ face cancel each other out in pairs (using $\Phi_{\boldsymbol{s}}^{-L_1}$). By a similar argument, we can kill off the remaining terms blocked by the $[2,-]$ face, and so forth. We conclude that the whole of $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$ is acyclic. □

It is worthwhile to note that we can similarly define an *outgoing crossover point* to be a pair $(\varepsilon^0, \boldsymbol{s})$ with $\varepsilon^0 \in \mathbb{E}_\ell$ and $\boldsymbol{s} = (s_1,\dots,s_\ell) \in P(\widetilde{\Lambda},\Lambda,\varepsilon^0)$ satisfying the following property: there exist $\varepsilon, \varepsilon' \in \mathbb{E}_\ell$ such that

- $\varepsilon^0 + \varepsilon \in \mathbb{E}_\ell$ and $\varepsilon' \le \varepsilon$;
- $\boldsymbol{s} + \varepsilon' \cdot \Lambda \notin P(\widetilde{\Lambda},\Lambda,\varepsilon^0+\varepsilon)$;
- there is no $i = 1,\dots,\ell$ with $\varepsilon_i = 1$, $\varepsilon_i' = 1$ and $s_i > b^\delta$;
- there is no $i = 1,\dots,\ell$ with $\varepsilon_i = 1$, $\varepsilon_i' = 0$ and $s_i + \sum_{j\neq i} \mathrm{lk}(L_i,L_j)(1-\varepsilon_j^0) < -b^\delta$.

We denote by $\mathrm{Out}(\widetilde{\Lambda},\Lambda)$ the set of all outgoing crossover points. (In Figure 34, the set $\mathrm{Out}(\widetilde{\Lambda},\Lambda)$ is the union of the red regions.) Define the *outgoing crossover quotient complex* of $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ to be

$$\mathscr{C}_{\mathrm{Out}}^{\delta}\{\widetilde{\Lambda}\} = \bigoplus_{(\varepsilon,\boldsymbol{s})\in \mathrm{Out}(\widetilde{\Lambda},\Lambda)} \mathscr{C}_{\boldsymbol{s}}^{\varepsilon,\delta},$$

with the differential induced from $\mathscr{C}_{\mathrm{Out}}^{\delta}\{\widetilde{\Lambda}\}$. Roughly, $\mathscr{C}_{\mathrm{Out}}^{\delta}\{\widetilde{\Lambda}\}$ is the domain of all crossover maps that are not being automatically set to zero by applying Lemma 10.10. One can show that $\mathscr{C}_{\mathrm{Out}}^{\delta}\{\widetilde{\Lambda}\}$ is acyclic by an argument similar to the one used for $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$. However, we will not need to use the complex $\mathscr{C}_{\mathrm{Out}}^{\delta}\{\widetilde{\Lambda}\}$ in the proof below.

Proof of Proposition 10.15 Let

$$\mathscr{C}^\delta\{\{\widetilde{\Lambda}\}\} = \bigoplus_{\substack{\varepsilon\in\mathbb{E}_\ell \\ \boldsymbol{s}\in P(\widetilde{\Lambda},\Lambda,\varepsilon) \\ (\varepsilon,\boldsymbol{s})\notin \mathrm{In}(\widetilde{\Lambda},\Lambda)}} \mathscr{C}_{\boldsymbol{s}}^{\varepsilon,\delta}$$

be the quotient complex of $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ obtained from by quotienting out the subcomplex $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$. Since $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$ are acyclic, it follows that $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ and $\mathscr{C}^\delta\{\{\widetilde{\Lambda}\}\}$ are quasi-isomorphic.

Equally well, we can think of $\mathscr{C}_{\mathrm{In}}^{\delta}\{\widetilde{\Lambda}\}$ as a subcomplex of the truncation $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$ (the one without crossover maps). Thus, $\mathscr{C}^\delta\{\{\widetilde{\Lambda}\}\}$ is a quotient of $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$, and quasi-isomorphic to it. We deduce that $\mathscr{C}^\delta\langle\langle\widetilde{\Lambda}\rangle\rangle$ and $\mathscr{C}^\delta\{\widetilde{\Lambda}\}$ are quasi-isomorphic. The conclusion then follows from Proposition 10.12. □

11 A general surgery exact sequence

Theorem 9.6 is a generalization of the corresponding result for knots [Ozsváth and Szabó 2008b, Theorem 1.1]. The key component of the proof of that result is the existence of a long exact sequence

$$\ldots \to \boldsymbol{HF}^-(Y_n(K)) \to \boldsymbol{HF}^-(Y_{m+n}(K)) \to \bigoplus^m \boldsymbol{HF}^-(Y) \to \cdots ; \tag{118}$$

see [ibid., Theorem 3.1], stated there with HF^+ instead of $\boldsymbol{HF}^-$ (see also [Ozsváth and Szabó 2011, Theorem 6.2], for a generalization to rationally null-homologous knots inside three-manifolds). Here $n, m \in \mathbb{Z}$ are surgery coefficients of the knot $K \subset Y$ with $m > 0$, and Y an integral homology sphere. The fact that we work with the completed version $\boldsymbol{HF}^-$ (see Section 2) allows the triangle to exist for $\boldsymbol{HF}^-$ just as for HF^+.

As a first step towards the proof of Theorem 9.6, in this section we will discuss a broader generalization of (118), in which K is an arbitrary knot in a closed, oriented three-manifold Y. In Sections 11.1 and 11.2 we establish a few useful facts about cobordism maps in Heegaard Floer homology. We then proceed to state the general exact sequence. Its proof in Section 11.3 is a straightforward adaptation of [Ozsváth and Szabó 2008b, Theorem 3.1], as long as one does not keep track of gradings and decompositions into Spin^c structures. We deal with these two last issues at length in Sections 11.5, 11.6 and 11.7 for the exact sequence with vertically truncated complexes.

11.1 Four-dimensional cobordisms

We spell out here a way of describing homology classes, cohomology classes, and Spin^c structures on four-dimensional cobordisms (given by two-handle additions) in terms of surgery. This will be useful to us when discussing the surgery long exact sequence later in this section, as well as when discussing surgery maps in Section 14.

Let $(\vec{L}, \Lambda)$ be a framed ℓ-component link inside an integral homology three-sphere Y. Recall that the space of Spin^c structures on $Y_\Lambda(L)$ is identified with the quotient $\mathbb{H}(L)/H(L,\Lambda)$ (see Section 9.2). Also, $H_1(Y_\Lambda(L)) \cong H^2(Y_\Lambda(L))$ is identified with $\mathbb{Z}^\ell/H(L,\Lambda)$ and $H_2(Y_\Lambda(L)) \cong H^1(Y_\Lambda(L))$ with $H(L,\Lambda)^\perp$ (see Section 9.3).

Let $L' \subseteq L$ be an ℓ'-component sublink, with the orientation induced from $\vec{L}$. Denote by

$$H(L, \Lambda|_{L'}) \subset H_1(Y - L) \cong \mathbb{Z}^\ell$$

the sublattice generated by the framings Λ_i, for $L_i \subseteq L'$.

Let $W_\Lambda(L', L)$ be the cobordism from $Y_{\Lambda|_{L'}}(L')$ to $Y_\Lambda(L)$ given by surgery on $L - L'$ (framed with the restriction of Λ). Further, let $W_\Lambda(L) = W_\Lambda(\varnothing, L)$, so that

$$W_\Lambda(L) = W_{\Lambda|_{L'}}(L') \cup W_\Lambda(L', L).$$

Lemma 11.1 (a) *There is a natural identification*

$$H_2(W_\Lambda(L',L)) \cong H(L,\Lambda|_{L'})^\perp = \{\boldsymbol{v} \in \mathbb{Z}^\ell \mid \boldsymbol{v} \cdot \Lambda_i = 0 \text{ for all } i \text{ with } L_i \subseteq L'\}, \tag{119}$$

under which the intersection product

$$H_2(W_\Lambda(L',L)) \otimes H_2(W_\Lambda(L',L)) \to \mathbb{Z}$$

is given by $\boldsymbol{v} \otimes \boldsymbol{v}' \to \boldsymbol{v}^t \Lambda \boldsymbol{v}'$.

(b) *There is a natural identification*

$$H^2(W_\Lambda(L',L)) \cong \mathbb{Z}^\ell / H(L,\Lambda|_{L'}), \tag{120}$$

under which the natural projection

$$\pi^{L,L'} \colon \mathbb{Z}^\ell / H(L,\Lambda|_{L'}) \to \mathbb{Z}^\ell / H(L,\Lambda)$$

corresponds to restriction to $H^2(Y_\Lambda(L))$, *while restriction to the coordinates corresponding to* L',

$$\mathbb{Z}^\ell / H(L,\Lambda|_{L'}) \to \mathbb{Z}^{\ell'} / H(L',\Lambda|_{L'}),$$

corresponds to restriction to $H^2(Y_{\Lambda|_{L'}}(L'))$.

(c) *Under the identifications* (119) *and* (120), *the evaluation map*

$$H^2(W_\Lambda(L',L)) \otimes H_2(W_\Lambda(L',L)) \to \mathbb{Z}$$

corresponds to the usual scalar multiplication of vectors in $\mathbb{Z}^\ell$. *Also, the composition of Poincaré duality with the natural map in the long exact sequence of a pair*

$$H_2(W_\Lambda(L',L)) \xrightarrow{\cong} H^2(W_\Lambda(L',L), \partial W_\Lambda(L',L)) \to H^2(W_\Lambda(L',L))$$

corresponds to the multiplication $\boldsymbol{v} \to \Lambda \boldsymbol{v}$.

Proof We start by proving the claims in the case when $L' = \varnothing$, ie $W_\Lambda(L',L) = W_\Lambda(L)$. Choose Seifert surfaces $F_i \subset Y$ for each link component L_i. (Of course, these surfaces may intersect each other.) Let $\widehat{F}_i$ be the surface obtained by capping off F_i in $W_\Lambda(L)$, using the core of the respective two-handle. Note that the homology class $[\widehat{F}_i]$ is independent of our choice of F_i. Further, since Y is an integral homology sphere, the classes $[\widehat{F}_i], i = 1, \dots, \ell$, form a basis of $H_2(W_{\widetilde{\Lambda}}(L))$, so we have the desired identification

$$H_2(W_\Lambda(L)) \cong \mathbb{Z}^\ell. \tag{121}$$

Note that the intersection form on $H_2(W_\Lambda(L))$ in this basis is the framing matrix Λ.

We also have another identification

$$H^2(W_\Lambda(L)) \cong \mathbb{Z}^\ell, \tag{122}$$

obtained by sending a cohomology class c to $(\langle c, [\widehat{F}_1]\rangle, \dots, \langle c, [\widehat{F}_\ell]\rangle)$. The claims in (c) are then easy to check for $W_\Lambda(L)$.

Let us now consider the general case of $L' \subseteq L$. For part (a), observe that

$$H_3(W_\Lambda(L), W_\Lambda(L', L)) \cong H_3(W_{\Lambda|_{L'}}(L'), Y_{\Lambda|_{L'}}(L')) = 0,$$

because $W_{\Lambda|_{L'}}(L')$ consists of two-handle additions only. Hence, the long exact sequence in homology for the pair $(W_\Lambda(L), W_\Lambda(L', L))$ reads

$$0 \to H_2(W_\Lambda(L', L)) \to H_2(W_\Lambda(L)) \xrightarrow{f} H_2(W_\Lambda(L), W_\Lambda(L', L)) \cong H^2(W_{\Lambda|_{L'}}(L')).$$

Using (121) and (122), we can view f as a map from $\mathbb{Z}^\ell$ to $\mathbb{Z}^{\ell'}$. From part (c) for $W_\Lambda(L)$, we get that f is given in matrix form by the restriction of Λ to the components of L'. The identification (119) follows.

For part (b), use the identifications (122) for $W_\Lambda(L)$ and $W_{\Lambda|_{L'}}(L')$, as well as the commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & H^2(W_\Lambda(L), W_\Lambda(L', L)) \arrow[r] \arrow[d, "\cong"] & H^2(W_\Lambda(L)) \arrow[r] \arrow[d] & H^2(W_\Lambda(L', L)) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & H^2(W_{\Lambda|_{L'}}(L'), Y_{\Lambda|_{L'}}(L')) \arrow[r] & H^2(W_{\Lambda|_{L'}}(L')) \arrow[r] & H^2(Y_{\Lambda|_{L'}}(L')) \arrow[r] & 0
\end{tikzcd}$$

Part (c) follows from the respective statements for $W_\Lambda(L)$. □

We also have a description of the space of Spin^c structures on the cobordism $W_\Lambda(L', L)$:

Lemma 11.2 *There is a natural identification*

$$\mathrm{Spin}^c(W_\Lambda(L', L)) \cong \mathbb{H}(L)/H(L, \Lambda|_{L'}) \tag{123}$$

under which the natural projection

$$\pi^{L,L'}\colon \mathbb{H}(L)/H(L, \Lambda|_{L'}) \to \mathbb{H}(L)/H(L, \Lambda)$$

corresponds to restricting the Spin^c *structures to* $Y_\Lambda(L)$*, while the map*

$$\psi^{L-L'}\colon \mathbb{H}(L)/H(L, \Lambda|_{L'}) \to \mathbb{H}(L')/H(L', \Lambda|_{L'})$$

corresponds to restricting them to $Y_{\Lambda|_{L'}}(L')$*. Further, the first Chern class map*

$$c_1\colon \mathrm{Spin}^c(W_\Lambda(L', L)) \to H^2(W_\Lambda(L', L)) \cong \mathbb{Z}^\ell/H(L, \Lambda|_{L'})$$

is given by

$$c_1([\boldsymbol{s}]) = [2\boldsymbol{s} - (\Lambda_1 + \cdots + \Lambda_\ell)]. \tag{124}$$

Proof The space of Spin^c structures on $W_\Lambda(L)$ is identified with the space of relative Spin^c structures on (Y, L), and hence with $\mathbb{H}(L)$ via the formula (124) (see Remark 3.6). Similarly, we have an identification $\mathrm{Spin}^c(W_{\Lambda|_{L'}}(L')) \cong \mathbb{H}(L')$. Moreover, there is a commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & \mathrm{Spin}^c(W_\Lambda(L), W_\Lambda(L', L)) \arrow[r] \arrow[d, "\cong"] & \mathrm{Spin}^c(W_\Lambda(L)) \arrow[r] \arrow[d] & \mathrm{Spin}^c(W_\Lambda(L', L)) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathrm{Spin}^c(W_{\Lambda|_{L'}}(L'), Y_{\Lambda|_{L'}}(L')) \arrow[r] & \mathrm{Spin}^c(W_{\Lambda|_{L'}}(L')) \arrow[r] & \mathrm{Spin}^c(Y_{\Lambda|_{L'}}(L')) \arrow[r] & 0
\end{tikzcd}$$

where the two horizontal rows are short exact sequences. The conclusion easily follows. □

11.2 Gradings, cobordism maps and twisted coefficients

We now discuss the general grading properties for Heegaard Floer complexes with twisted coefficients and cobordism maps between them. Let Y be a closed, oriented three-manifold with a Spin^c structure $\mathfrak{u}$ and a $\mathbb{F}[H^1(Y;\mathbb{Z})]$-module M. A pointed, admissible Heegaard diagram for Y (with a complete set of paths for the generators, as in [Ozsváth and Szabó 2004b, Definition 3.12]) gives rise to a twisted Heegaard Floer complex $\boldsymbol{CF}^-(Y,\mathfrak{u};M)$; see [Ozsváth and Szabó 2004c, Section 8.1]. (Here we use the completed version, and delete the usual underline from notation for simplicity.) This complex admits a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u},M)\mathbb{Z}$-grading, where

$$\mathfrak{d}(\mathfrak{u},M)=\gcd_{\substack{\xi\in H_2(Y;\mathbb{Z})\\ \mathrm{PD}(\xi)\cdot m=m \text{ for all } m\in M}}\langle c_1(\mathfrak{u}),\xi\rangle. \tag{125}$$

This is true because the ambiguity in the grading difference between two generators is the Maslov index of periodic domains for which the corresponding classes in $H^1(Y;\mathbb{Z})$ act trivially on M; compare [loc. cit.].

Next, we set up cobordism maps with twisted coefficients, following [ibid., Section 8.2], but in slightly more generality.

Consider a pointed, admissible triple Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},w)$ which represents a cobordism $X=X_{\alpha,\beta,\gamma}$ with boundaries $Y_{\alpha,\beta}$, $Y_{\beta,\gamma}$ and $Y_{\alpha,\gamma}$. Suppose we are given an $\mathbb{F}[H^1(Y_{\alpha,\beta};\mathbb{Z})]$-module $M_{\alpha,\beta}$ and an $\mathbb{F}[H^1(Y_{\beta,\gamma};\mathbb{Z})]$-module $M_{\beta,\gamma}$. Given a Spin^c structure $\boldsymbol{t}$ on X, we denote by $\underline{\mathrm{Spin}}^c(X;\boldsymbol{t})$ the space of relative Spin^c structures on X representing $\boldsymbol{t}$. Note that $\underline{\mathrm{Spin}}^c(X;\boldsymbol{t})$ has a natural action of $H^1(Y_{\alpha,\beta};\mathbb{Z})\times H^1(Y_{\beta,\gamma};\mathbb{Z})\times H^1(Y_{\alpha,\gamma};\mathbb{Z})$. The Spin^c structure $\boldsymbol{t}$ induces an $\mathbb{F}[H^1(Y_{\alpha,\gamma};\mathbb{Z})]$-module

$$\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\boldsymbol{t}}=\frac{(m_{\alpha,\beta},m_{\beta,\gamma},\underline{\boldsymbol{t}})\in M_{\alpha,\beta}\times M_{\beta,\gamma}\times\underline{\mathrm{Spin}}^c(X;\boldsymbol{t})}{(m_{\alpha,\beta},m_{\beta,\gamma},\underline{\boldsymbol{t}})\sim(h_{\alpha,\beta}\cdot m_{\alpha,\beta},h_{\beta,\gamma}\cdot m_{\beta,\gamma},(h_{\alpha,\beta}\times h_{\beta,\gamma}\times 0)\cdot\underline{\boldsymbol{t}})},$$

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

Suppose we are also given an $\mathbb{F}[H^1(Y_{\alpha,\gamma};\mathbb{Z})]$-module $M_{\alpha,\gamma}$ and a module homomorphism

$$\zeta\colon\{M_{\alpha,\beta}\otimes M_{\beta,\gamma}\}^{\boldsymbol{t}}\to M_{\alpha,\gamma}.$$

For simplicity, let us write M for the triple $(M_{\alpha,\beta},M_{\beta,\gamma},M_{\alpha,\gamma})$. We then have a cobordism map with twisted coefficients

$$f^-_{\alpha,\beta,\gamma;\boldsymbol{t},M,\zeta}\colon\boldsymbol{CF}^-(Y_{\alpha,\beta},\boldsymbol{t}|_{Y_{\alpha,\beta}};M_{\alpha,\beta})\otimes\boldsymbol{CF}^-(Y_{\beta,\gamma},\boldsymbol{t}|_{Y_{\beta,\gamma}};M_{\beta,\gamma})\to\boldsymbol{CF}^-(Y_{\alpha,\gamma},\boldsymbol{t}|_{Y_{\alpha,\gamma}};M_{\alpha,\gamma}),$$

given by

$$f^-_{\alpha,\beta,\gamma;\boldsymbol{t},M,\zeta}(m_{\alpha,\beta}\boldsymbol{x}\otimes m_{\beta\gamma}\boldsymbol{y})=\sum_{\boldsymbol{z}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma}\sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z})\\ \mu(\phi)=0\\ \boldsymbol{t}_w(\phi)=\boldsymbol{t}}}(\#\mathcal{M}(\phi))\cdot U^{n_w(\phi)}\zeta(m_{\alpha,\beta}\otimes m_{\beta\gamma}\otimes\bar{\boldsymbol{t}}_w(\phi))\boldsymbol{z},$$

with the usual notation in Heegaard Floer theory; compare [ibid., equation (10)].

Set

$$\mathfrak{d}(\boldsymbol{t}, M) = \gcd\big(\mathfrak{d}(\boldsymbol{t}|_{Y_{\alpha,\beta}}, M_{\alpha,\beta}), \mathfrak{d}(\boldsymbol{t}|_{Y_{\beta,\gamma}}, M_{\beta,\gamma}), \mathfrak{d}(\boldsymbol{t}|_{Y_{\alpha,\gamma}}, M_{\alpha,\gamma})\big). \tag{126}$$

Note that the Floer complexes

$$\boldsymbol{CF}^-(Y_{\alpha,\beta}, \boldsymbol{t}|_{Y_{\alpha,\beta}}; M_{\alpha,\beta}), \boldsymbol{CF}^-(Y_{\beta,\gamma}, \boldsymbol{t}|_{Y_{\beta,\gamma}}; M_{\beta,\gamma}), \boldsymbol{CF}^-(Y_{\alpha,\gamma}, \boldsymbol{t}|_{Y_{\alpha,\gamma}}; M_{\alpha,\gamma})$$

all admit relative $\mathbb{Z}/\mathfrak{d}(\boldsymbol{t}, M)\mathbb{Z}$-gradings.

Lemma 11.3 (a) *The map $f^-_{\alpha,\beta,\gamma;\boldsymbol{t},M,\zeta}$ preserves the relative $\mathbb{Z}/\mathfrak{d}(\boldsymbol{t}, M)\mathbb{Z}$-gradings.*

(b) *Let $\boldsymbol{t}'$ be a Spin^c structure on X with the same restrictions to ∂X as $\boldsymbol{t}$, and let $\zeta' : \{M_{\alpha,\beta} \otimes M_{\beta,\gamma}\}^{\boldsymbol{t}'} \to M_{\alpha,\gamma}$ be a module homomorphism. Set $\boldsymbol{t}' - \boldsymbol{t} = u \in H^2(X, \partial X; \mathbb{Z})$. Then, for every pair of homogeneous elements $\boldsymbol{x} \in \boldsymbol{CF}^-(Y_{\alpha,\beta}, \boldsymbol{t}|_{Y_{\alpha,\beta}}; M_{\alpha,\beta})$ and $\boldsymbol{y} \in \boldsymbol{CF}^-(Y_{\beta,\gamma}, \boldsymbol{t}|_{Y_{\beta,\gamma}}; M_{\beta,\gamma})$,*

$$\mathrm{gr}\, f^-_{\alpha,\beta,\gamma;\boldsymbol{t}',M,\zeta'}(\boldsymbol{x} \otimes \boldsymbol{y}) - \mathrm{gr}\, f^-_{\alpha,\beta,\gamma;\boldsymbol{t},M,\zeta}(\boldsymbol{x} \otimes \boldsymbol{y}) \equiv (c_1(\boldsymbol{t}) \smile u + u \smile u)[X] \pmod{\mathfrak{d}(\boldsymbol{t}, M)}. \tag{127}$$

Proof (a) The ambiguity in the grading shift comes from doubly periodic domains with trivial module actions. The contribution of these domains to the Maslov index is given by their pairing with $c_1(\boldsymbol{t})$.

(b) This follows from the formula for the Maslov index of a triply periodic domain in [Sarkar 2011, Section 5]. □

Of course, equation (125) and Lemma 11.3 apply equally well to truncated Floer complexes $\boldsymbol{CF}^{-,\delta}$ instead of $\boldsymbol{CF}^-$, and to triple Heegaard diagrams with several basepoints.

Remark 11.4 A particular example of a cobordism map is the untwisted one corresponding to a two-handle addition, as in [Ozsváth and Szabó 2006]. In this case, $Y_{\beta,\gamma}$ is a connected sum of $S^1 \times S^2$'s, and we consider only triangles with one vertex at a representative for the top-degree homology generator $\boldsymbol{y} = \Theta^{\mathrm{can}}_{\beta,\gamma}$. Supposing further that $c_1(\boldsymbol{t})$ has torsion restrictions to $Y_{\alpha,\beta}$ and $Y_{\alpha,\gamma}$, the respective Floer homology groups have absolute $\mathbb{Q}$-gradings compatible with their relative $\mathbb{Z}$-gradings. Let W be the cobordism from $Y_{\alpha,\beta}$ and $Y_{\alpha,\gamma}$ obtained from X by filling in the other boundary component with three-handles. It is shown in [loc. cit.] that the cobordism map shifts absolute grading by

$$\tfrac{1}{4}(c_1(\boldsymbol{t})^2 - 2\chi(W) - 3\sigma(W)), \tag{128}$$

where χ and σ denote Euler characteristic and signature, respectively. If we further suppose that $c_1(\boldsymbol{t}')$ has torsion restrictions to $Y_{\alpha,\beta}$ and $Y_{\alpha,\gamma}$, we can then view (127) as a simple consequence of (128), with $\mathfrak{d}(\boldsymbol{t}, M) = 0$.

11.3 A long exact sequence

In this section we sketch the construction of the surgery long exact sequence from [Ozsváth and Szabó 2008b, Section 3; 2011, Section 6.1]. It is stated there for rationally null-homologous knots, but the construction can be generalized to works for arbitrary knots inside three-manifolds. In the more general setting, one of the Floer complexes may appear with genuinely twisted coefficients; see Proposition 11.5

below, as well as its graded refinements Propositions 11.11, 11.15, 11.17, 11.24, 11.28 and 11.29. (A slightly different generalized surgery sequence with twisted coefficients was proved by Fink [2009].)

We work in the setting of Section 11.1, with $(\vec{L}, \Lambda)$ being a framed ℓ-component link inside an integral homology three-sphere Y. Let $\overline{\Lambda}$ be the framing on L obtained from Λ by adding $m_1 > 0$ to the surgery coefficient λ_1 of the first component L_1; that is, the corresponding framing vectors are $\overline{\Lambda}_1 = \Lambda_1 + m_1\tau_1$, where here τ_1 is a meridian for L_1, and $\overline{\Lambda}_i = \Lambda_i$ for $i \neq 1$. Also, we let $L' = L - L_1$ and denote by Λ' the restriction of the framing Λ to L', with framing vectors $\Lambda'_2, \dots, \Lambda'_\ell$. We can view $K = L_1$ as a knot inside the three-manifold $Y_{\Lambda'}(L')$, and $Y_\Lambda(L)$ as the result of surgery along that knot. (Note that, given any knot K inside a three-manifold M, we can find $L = L_1 \cup L' \subset Y = S^3$ such that $M = S^3_{\Lambda'}(L')$ and K corresponds to L_1.)

Suppose we have a multipointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, z_1)$ for $K \subset Y_{\Lambda'}(L')$, with possibly several free basepoints, but a single basepoint pair (w_1, z_1) on K. Let g be the genus of Σ and k the total number of w basepoints, as in Section 3.1. Moreover, we assume that w_1 and z_1 can be connected by a path which crosses β_{g+k-1} exactly once, and which is disjoint from all the other alpha and beta curves.

As in [Ozsváth and Szabó 2008b, proof of Theorem 3.1], we let γ_{g+k-1} be a simple, closed curve in Σ disjoint from $\beta_1, \dots, \beta_{g+k-2}$ which specifies the λ_1-framing of $L_1 \subset Y$. We complete this to a $(g+k-1)$-tuple $\boldsymbol{\gamma}$ of attaching curves on Σ by taking curves $\gamma_1, \dots, \gamma_{g+k-2}$ which approximate $\beta_1, \dots, \beta_{g+k-2}$ in the sense of Definition 8.1. We define another collection $\boldsymbol{\delta}$ similarly, only now δ_{g+k-1} specifies the framing $\lambda_1 + m_1$ on L_1. Thus, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta}, \boldsymbol{w})$ are Heegaard diagrams for $Y_\Lambda(L)$ and $Y_{\overline{\Lambda}}(L)$, respectively.

The Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\delta}, \boldsymbol{w})$ represents a four-manifold X_1 with three boundary components,

$$Y_\Lambda(L), \quad L(m_1, 1) \# \#^{g+k-2}(S^1 \times S^2) \quad \text{and} \quad Y_{\overline{\Lambda}}(L).$$

There is a canonical torsion Spin^c structure on the manifold $L(m_1, 1) \# \#^{g+k-2}(S^1 \times S^2)$; see [ibid., Definition 3.2]. We arrange that the Floer homology $\boldsymbol{HF}^-(\mathbb{T}_\gamma, \mathbb{T}_\delta, \boldsymbol{w})$ in that Spin^c structure, in the maximal degree with nonzero homology, is represented by a unique intersection point, which we call *canonical*. We define a map

$$f_1^- : \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w}) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\delta, \boldsymbol{w})$$

by counting holomorphic triangles with one vertex in the canonical intersection point.

Let us add $g + k - 2$ three-handles to X_1 to kill off the $S^1 \times S^2$ summands in the middle boundary component, then remove a neighborhood of a path between the first two boundary components. We thus obtain a cobordism W_1 from $Y_\Lambda(L) \# L(m_1, 1)$ to $Y_{\overline{\Lambda}}(L)$. It is easy to see that

$$f_1^-(\boldsymbol{x}) = f^-_{W_1}(\boldsymbol{x} \otimes \Theta^{\text{can}}_{\gamma,\delta}),$$

where $f^-_{W_1}$ is the map on Floer complexes induced by the cobordism W_1 (as in [Ozsváth and Szabó 2006]), and $\Theta^{\mathrm{can}}_{\gamma,\delta}$ is the top-degree generator for the Floer homology of $L(m_1,1)$ (in its canonical Spin^c structure $\mathfrak{u}^{\mathrm{can}}$).

Next, we look at the Heegaard triple $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\delta},\boldsymbol{\beta},\boldsymbol{w})$. This represents a cobordism X_2 with three boundary components,

$$Y_{\bar{\Lambda}}(L), \quad \#^{g+k-2}(S^1\times S^2) \quad \text{and} \quad Y_{\Lambda'}(L').$$

By filling in the middle component with three-handles, we obtain a cobordism W_2 from $Y_{\bar{\Lambda}}(L)$ to $Y_{\Lambda'}(L')$. This is simply the reverse of the cobordism $-W_{\bar{\Lambda}}(L',L)$ in the notation of Section 11.1. Here, by reverse of a cobordism we mean a reversal in direction, ie turning the cobordism around so that we view it as a cobordism from the final to the initial manifold. The minus sign denotes the additional reversal of orientation.

To X_2 we associate a cobordism map with twisted coefficients as in Section 11.2. (Compared to the notation there, we now have X_2 instead of X, $\boldsymbol{\delta}$ instead of $\boldsymbol{\beta}$, and $\boldsymbol{\beta}$ instead of $\boldsymbol{\gamma}$.) We let $M_{\alpha,\delta}$ and $M_{\delta,\beta}$ be the modules $\mathbb{F}$, with trivial action by the respective cohomology groups. Consider the ring

$$\mathscr{T}_1=\mathbb{F}[\mathbb{Z}/m_1\mathbb{Z}]=\mathbb{F}[T_1]/(T_1^{m_1}-1).$$

We make $\mathscr{T}_1$ into an $\mathbb{F}[H^1(Y_{\Lambda'}(L'))]$-module $M_{\alpha,\beta}$ by letting $h\in H^1(Y_{\Lambda'}(L'))$ act by multiplication by $T_1^{\langle h,[L_1]\rangle}$. In the corresponding Floer complex $\boldsymbol{CF}^-(Y_{\Lambda'}(L');\mathscr{T}_1)$ each isolated holomorphic strip gets counted in the differential with a coefficient $T_1^{n_{w_1}(\phi)-n_{z_1}(\phi)}$, where ϕ is the homotopy class of the strip.

Given $\boldsymbol{t}\in \mathrm{Spin}^c(W_2)$, we abuse notation slightly, letting $\boldsymbol{t}$ also denote the restriction of $\boldsymbol{t}$ from W_2 to X_2. We consider the module homomorphism

$$\zeta\colon \{M_{\alpha,\delta}\otimes M_{\delta,\beta}\}^{\boldsymbol{t}}\to M_{\alpha,\beta}$$

given by composing the projection

$$\{M_{\alpha,\delta}\otimes M_{\delta,\beta}\}^{\boldsymbol{t}}\to \frac{\underline{\mathrm{Spin}}^c(X_2;\boldsymbol{t})}{H^1(Y_{\bar{\Lambda}}(L))\times H^1\big(\#^{g+k-2}(S^1\times S^2)\big)}$$

with the map

$$\frac{\underline{\mathrm{Spin}}^c(X_2;\boldsymbol{t})}{H^1(Y_{\bar{\Lambda}}(L))\times H^1\big(\#^{g+k-2}(S^1\times S^2)\big)}\to\mathscr{T}_1,\quad [\psi]\mapsto T_1^{n_{w_1}(\psi)-n_{z_1}(\psi)}, \tag{129}$$

where $\psi\in\underline{\mathrm{Spin}}^c(X_2;\boldsymbol{t})$ is viewed as a homotopy class of triangles; compare [Ozsváth and Szabó 2004c, Section 8.2.1]. Note that the map (129) is well defined, because every doubly periodic domain ϕ for $Y_{\bar{\Lambda}}(L)$ or $\#^{g+k-2}(S^1\times S^2)$ has $n_{w_1}(\phi)=n_{z_1}(\phi)$.

We define

$$f_2^-\colon \boldsymbol{CF}^-(\mathbb{T}_\alpha,\mathbb{T}_\delta,\boldsymbol{w})\to\boldsymbol{CF}^-(\mathbb{T}_\alpha,\mathbb{T}_\beta,\boldsymbol{w};\mathscr{T}_1)$$

by

$$f_2^-(\boldsymbol{x}) = \sum_{\boldsymbol{t} \in \mathrm{Spin}^c(W_2)} f^-_{\alpha,\delta,\beta;\boldsymbol{t},M,\zeta}(\boldsymbol{x} \otimes \Theta^{\mathrm{can}}_{\delta,\beta}),$$

where $\Theta^{\mathrm{can}}_{\delta,\beta}$ is the respective canonical generator. In other words, the map f_2^- counts holomorphic triangles in a class ψ with a coefficient $T_1^{n_{w_1}(\psi)-n_{z_1}(\psi)}$. For future reference, since we will be interested in the grading properties of f_2^-, we note that

$$\mathfrak{d}(\boldsymbol{t}, \mathscr{T}_1) = \gcd\big(\mathfrak{d}(\boldsymbol{t}|_{Y_{\bar\Lambda}(L)}), \mathfrak{d}(\boldsymbol{t}|_{Y_{\Lambda'}(L')}, \mathscr{T}_1)\big), \tag{130}$$

where we wrote $\mathfrak{d}(\boldsymbol{t}, \mathscr{T}_1) = \mathfrak{d}(\boldsymbol{t}, M)$ for simplicity. Equation (130) follows from (126) together with the triviality of the modules $M_{\alpha,\delta}$ and $M_{\delta,\beta}$.

Finally, we consider the cobordism X_3 corresponding to the Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$. The associated filled-in cobordism W_3 from $Y_{\Lambda'}(L')$ to $Y_\Lambda(L)$ is simply $W_\Lambda(L', L)$, ie surgery on the framed knot (L_1, Λ_1). We set up a map

$$f_3^- \colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}; \mathscr{T}_1) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w})$$

as follows. (Compare the definition of the map f_3^+ in [Ozsváth and Szabó 2008b, Section 3].) Fix any triangle class $\psi \in \pi_2(\Theta^{\mathrm{can}}_{\gamma,\beta}, \Theta^{\mathrm{can}}_{\beta,\delta}, \Theta^{\mathrm{can}}_{\gamma,\delta})$, where $\Theta^{\mathrm{can}}_{i,j}$ are canonical generators as before. Set

$$c = n_{w_1}(\psi) - n_{z_1}(\psi).$$

By analyzing triply periodic domains, one can check that the residue of c modulo m_1 is independent of the choice of ψ. Set

$$f_3^-(T_1^{s_1} \cdot \boldsymbol{x}) = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \Theta^{\mathrm{can}}_{\beta,\gamma}, \boldsymbol{y}) \\ \mu(\phi)=0 \\ s_1 + n_{w_1}(\phi) - n_{z_1}(\phi) \equiv c \,(\mathrm{mod}\, m_1)}} (\#\mathcal{M}(\phi)) \cdot \boldsymbol{U}^{n_w(\phi)} \boldsymbol{y}, \tag{131}$$

where

$$\boldsymbol{U}^{n_w(\phi)} := U_1^{n_{w_1}(\phi)} \cdot U_k^{n_{w_k}(\phi)}.$$

Observe that f_3^- is not a usual cobordism map with twisted coefficients. Rather, its value depends on the exponent s_1 in the input $T_1^{s_1} \cdot \boldsymbol{x}$. When applied to a fixed $T_1^{s_1} \cdot \boldsymbol{x}$, the map f_3^- acts as a sum of several cobordism maps with twisted coefficients; these are associated to the cobordism $W_3 = W_\Lambda(L', L)$ and to the Spin^c structures $\boldsymbol{s} \in \mathrm{Spin}^c(W_3) \cong \mathbb{H}(L)/H(L, \Lambda|_{L'})$ with some fixed value of $\boldsymbol{s}$ modulo $(m_1, 0, \ldots, 0)$.

We denote by F_1^-, F_2^- and F_3^- the maps induced by f_1^-, f_2^- and f_3^- on homology.

Proposition 11.5 *For any framed link $(\vec{L}, \Lambda)$ inside an integral homology sphere, there is a long exact sequence*

$$\cdots \to \boldsymbol{HF}^-(Y_\Lambda(L)) \xrightarrow{F_1^-} \boldsymbol{HF}^-(Y_{\bar\Lambda}(L)) \xrightarrow{F_2^-} \boldsymbol{HF}^-(Y_{\Lambda'}(L'); \mathscr{T}_1) \xrightarrow{F_3^-} \cdots.$$

In fact, the complex $\boldsymbol{CF}^-(Y_\Lambda(L)) \cong \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w})$ is quasi-isomorphic to the mapping cone of f_2^-.

The proof of Proposition 11.5 is given in [Ozsváth and Szabó 2011, Section 6.1] for the case when L_1 is rationally null-homologous inside $Y_{\Lambda'}(L')$ (and for HF^+ instead of $\boldsymbol{HF}^-$), but it applies equally well to our situation, so we omit it.

We would like to have a refined statement of Proposition 11.5, in which we keep track of the decomposition of $\boldsymbol{CF}^-(Y_\Lambda(L))$ into Spin^c structures, as well as the respective relative gradings on it. However, keeping track of gradings is possible only if we work with vertically truncated complexes $\boldsymbol{CF}^{-,\delta}$ (as in Section 10.2) instead of $\boldsymbol{CF}^-$. We write f_i^δ (resp. H_i^δ) for the vertical truncation of f_i^- (resp. H_i^δ). The proof of Proposition 11.5 in fact gives the following more precise version:

Proposition 11.6 *Fix $\delta \geq 0$. For any framed link $(\vec{L}, \Lambda)$ inside an integral homology sphere, the complex $\boldsymbol{CF}^{-,\delta}(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w}) = \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L))$ is quasi-isomorphic to the mapping cone*

$$\boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L)) \xrightarrow{f_2^\delta} \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'); \mathscr{T}_1). \tag{132}$$

In fact, following [Ozsváth and Szabó 2008b; 2011], there are two natural quasi-isomorphisms that can be used to prove Proposition 11.5. First, we have the quasi-isomorphism

$$\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L)) \xrightarrow{\sim} \mathrm{Cone}\big(\boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L)) \xrightarrow{f_2^\delta} \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'); \mathscr{T}_1)\big) \tag{133}$$

given by a triangle-counting map f_1^δ to the first term in the mapping cone and a quadrilateral-counting map H_1^δ (a null-homotopy of $f_2^\delta \circ f_1^\delta$) to the second factor. Second, we have a quasi-isomorphism in the opposite direction

$$\mathrm{Cone}\big(\boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L)) \xrightarrow{f_2^\delta} \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'); \mathscr{T}_1)\big) \xrightarrow{\sim} \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L)), \tag{134}$$

given by a triangle-counting map f_3^δ from the second term of the mapping cone, and a quadrilateral-counting map H_2^δ (a null-homotopy of $f_3^\delta \circ f_2^\delta$) from the first term.

11.4 Refinements of Proposition 11.6

We will discuss below several refinements of Proposition 11.6, on a case-by-case basis, depending on Λ. By choosing m_1 judiciously (in particular, sufficiently large compared to δ), we describe cases where the quasi-isomorphism (133) or (134) has good Spin^c structure decompositions and good grading-preserving properties.

More precisely, we consider the following two cases:

- **Case I** Λ is nondegenerate.
- **Case II** Λ is degenerate, and furthermore $\Lambda_1 \in \mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)$.

(We do not consider the case where Λ is degenerate, but $\Lambda_1 \notin \mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)$. That case turns out not to be needed for our present applications.)

Case I is further subdivided as follows. Let $h > 0$ be the smallest integer such that the vector

$$\boldsymbol{a} = (a_1, \dots, a_\ell) = h\Lambda^{-1}(1, 0, \dots, 0)$$

has all integer coordinates. We choose $m_1 \gg 0$ such that the vector $m_1\tau_1 = (m_1, 0, \dots, 0)$ is in $H(L, \Lambda) = \mathrm{Span}(\Lambda_1, \dots, \Lambda_\ell)$.[4] Hence, the value $d = m_1/h$ is an integer, too. Choosing m_1 sufficiently large is the same as choosing d sufficiently large. In the case when $a_1 \neq 0$, we impose an additional constraint on our choices of sufficiently large m_1: namely, we ask for m_1 to be a multiple of $a_1 h$, ie for d to be a multiple of a_1.

It is easy to check that $a_1 = 0$ if and only if Λ' is degenerate, and when Λ' is nondegenerate, the sign of a_1 is the same as the sign of the restriction of Λ (viewed as a symmetric bilinear form) to the one-dimensional space $\mathrm{Span}_{\mathbb{Q}}(\Lambda_2, \dots, \Lambda_\ell)^{\perp} \subset \mathbb{Q}^\ell$. (Here, the orthogonal complement is taken with respect to the standard inner product.)

We thus distinguish three subcases, according to the sign of a_1:

Case I(a) $(a_1 > 0)$ In this subcase, the graded refinement of Proposition 11.6 is given in Proposition 11.11 below. The model example to keep in mind is that of positive surgery on a knot $K \subset Y$, eg $Y = S^3$, n is a positive integer, $\Lambda = (n)$, $\overline{\Lambda} = (n + m_1)$, $\Lambda' = \varnothing$, $a_1 = 1$ and $h = n$.

Case I(b) $(a_1 < 0)$ In this subcase, the refinement is given in Proposition 11.15 below. The model example to keep in mind is that of negative surgery on a knot $K \subset Y$, eg $Y = S^3$, n is a positive integer, $\Lambda = (-n)$, $\overline{\Lambda} = (-n + m_1)$, $\Lambda' = \varnothing$, $a_1 = -1$ and $h = n$.

Case I(c) $(a_1 = 0)$ In this subcase, the refinement is given in Proposition 11.17 below. The model example here is fractional surgery on a knot in S^3, as follows. Let $K \subset Y = S^3$ be a knot, which we can promote to a two-component link $L = L_1 \cup K$, where here L_1 is a meridian for K. In this case, write

$$\Lambda = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \overline{\Lambda} = \begin{pmatrix} m_1 & 1 \\ 1 & 0 \end{pmatrix}, \quad \Lambda' = (0), \quad a_1 = 0, \quad h = 1.$$

In this model example, $Y_\Lambda(L) = S^3$, $Y_{\overline{\Lambda}}(L) = S^3_{-1/m_1}(K)$ and $Y_{\Lambda'}(L') = S^3_0(K)$. Note now that $Y_{\Lambda'}(L')$ appears with twisted coefficients. Thus, Proposition 11.17 can be viewed as a generalized, negative-surgery version of the fractional surgery theorem [Ozsváth and Szabó 2004c, Theorem 9.14].

In Case II, there are two refinements, given in Propositions 11.24 and 11.28, corresponding to the two possible quasi-isomorphisms from (133) and (134), respectively. Another case, where the Floer complex for $Y_\Lambda(L)$ has twisted coefficients, is given in Proposition 11.29. The model example for Case II is that of 0 surgery on a knot in an integral homology three-sphere: $\Lambda = (0)$, $\overline{\Lambda} = (m_1)$, $\Lambda' = \varnothing$. In this example, Proposition 11.24 should be compared with the integer surgeries exact sequence [ibid., Theorem 9.19].

In all these cases, to understand the Spin^c decompositions, we find it useful to study further the cobordism W_1 from $Y_\Lambda(L) \# L(m_1, 1)$ to $Y_{\overline{\Lambda}}(L)$, in the manner of Section 11.1.

[4] By "span" we will mean the span of integral vectors over $\mathbb{Z}$, unless we explicitly refer to the "$\mathbb{Q}$-span".

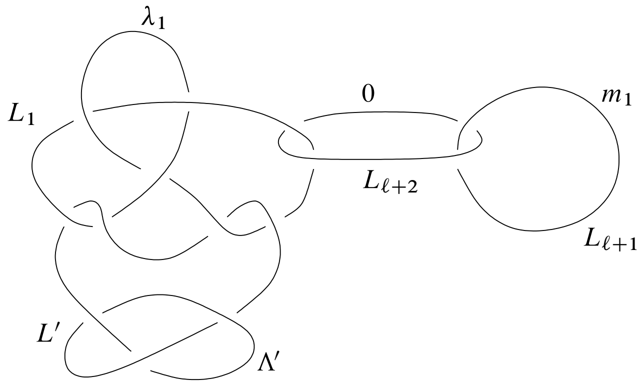

Figure 35: The cobordism W_1. By adding an unlinked, unknotted component $L_{\ell+1}$ with framing m to the original link L we obtain a surgery presentation for $Y_\Lambda(L)\,\#\,L(m_1,1)$. The cobordism W_1 is then given adding a two-handle along the unknot component L_{l+2} with framing 0, which has linking number one with both L_1 and $L_{\ell+1}$.

The cobordism W_1 consists of a single two-handle addition. A Kirby calculus picture for it is shown in Figure 35. If we denote by L^+ the $(\ell+2)$-component link $L \cup L_{\ell+1} \cup L_{\ell+2}$ from the figure, and by Λ^+ its given framing (also shown in the figure), we can express our cobordism as

$$W_1 = W_{\Lambda^+}(L^+ - L_{\ell+2}, L^+),$$

in the notation used in Section 11.1. In matrix form, the framing Λ^+ for L^+ is

$$\left(\begin{array}{cccc|cc} & & & & 0 & 1 \\ & \Lambda & & & 0 & 0 \\ & & & & \vdots & \vdots \\ & & & & 0 & 0 \\ \hline 0 & 0 & \cdots & 0 & m_1 & 1 \\ 1 & 0 & \cdots & 0 & 1 & 0 \end{array}\right) \tag{135}$$

We denote the rows of Λ^+ by Λ_i^+ for $i = 1, \dots, \ell+2$.

Definition 11.7 Two Spin^c structures $\mathfrak{u}$ on $Y_\Lambda(L)$ and $\bar{\mathfrak{u}}$ on $Y_{\bar\Lambda}(L)$ are called *linked* if there exists a Spin^c structure on W_1 which restricts to $\mathfrak{u}\,\#\,\mathfrak{u}^{\mathrm{can}}$ on $Y_\Lambda(L)\,\#\,L(m_1,1)$, and to $\bar{\mathfrak{u}}$ on $Y_{\bar\Lambda}(L)$.

Note that we can use Lemma 11.2 to describe Spin^c structures and restriction maps concretely. We have:

Lemma 11.8 *Two* Spin^c *structures* $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$ *and* $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\bar\Lambda}(L)) \cong \mathbb{H}(L)/H(L,\bar\Lambda)$ *are linked if and only if there exists* $\boldsymbol{s} \in \mathbb{H}(L)$ *and* $j \in \mathbb{Z}$ *such that* $\mathfrak{u} = \boldsymbol{s} + H(L,\Lambda)$ *and* $\bar{\mathfrak{u}} = \boldsymbol{s} + j m_1 \tau_1 + H(L,\bar\Lambda)$.

Proof Suppose we have a Spin^c structure

$$\boldsymbol{t} \in \mathrm{Spin}^c(W_1) \cong \mathbb{H}(L^+)/\mathrm{Span}(\Lambda_1^+, \dots, \Lambda_\ell^+, \Lambda_{\ell+1}^+),$$

and let $\boldsymbol{s}^+ = (s_1^+, \dots, s_\ell^+, s_{\ell+1}^+, s_{\ell+2}^+) \in \mathbb{H}(L^+)$ be a representative of $\boldsymbol{t}$. Using Lemma 11.2, the condition that the restriction of $\boldsymbol{t}$ to $Y_\Lambda(L) \# L(m_1, 1)$ is $\mathfrak{u} \# \mathfrak{u}^{\mathrm{can}}$ translates into

$$\mathfrak{u} = \boldsymbol{s} + H(L, \Lambda), \quad \text{where } \boldsymbol{s} = \left(s_1^+ - \tfrac{1}{2}, s_2^+, \dots, s_\ell^+\right) \in \mathbb{H}(L),$$

and

$$\psi^{L - L_{\ell+1}}(\boldsymbol{s}^+) = 0 \in \mathrm{Spin}^c(L(m_1, 1)) \cong \mathbb{Z}/m_1,$$

ie

$$s_{\ell+1}^+ = q m_1 + \tfrac{1}{2}$$

for some $q \in \mathbb{Z}$.

On the other hand, the restriction of $\boldsymbol{t}$ to $Y_{\overline{\Lambda}}(L)$ is

$$\boldsymbol{s}^+ + H(L^+, \Lambda^+) \in (\mathbb{H}(L^+)/H(L^+, \Lambda^+)).$$

Note that we have two different expressions for $\mathrm{Spin}^c(Y_{\overline{\Lambda}}(L))$ as a quotient, one coming from the cobordism W_1 and one from $W_{\overline{\Lambda}}(L)$. They are related by the isomorphism

$$\mathbb{H}(L^+)/H(L^+, \Lambda^+) \to \mathbb{H}(L)/H(L, \overline{\Lambda}), \quad \boldsymbol{s}^+ + H(L^+, \Lambda^+) \mapsto \bar{\boldsymbol{s}} + H(L, \overline{\Lambda}),$$

where

$$\bar{\boldsymbol{s}} = (\bar{s}_1, \dots, \bar{s}_\ell) = (s_1^+ - s_{\ell+1}^+ + m_1 s_{\ell+2}^+, s_2^+, \dots, s_\ell^+).$$

The conclusion follows by taking $j = s_{\ell+2}^+ - q$. $\square$

11.5 Refinements of Proposition 11.6, Case I: Λ is nondegenerate

Case I(a) $(a_1 > 0)$ (The model example to keep in mind is that of positive surgery on a knot $K \subset Y$, for which the grading issues have been studied in [Ozsváth and Szabó 2008b, Sections 4.5–4.6].)

In this situation $Y_\Lambda(L)$, $Y_{\overline{\Lambda}}(L)$, and $Y_{\Lambda'}(L')$ are all rational homology spheres, and therefore the corresponding Floer chain complexes all admit absolute $\mathbb{Q}$-gradings (and relative $\mathbb{Z}$-gradings) in each Spin^c structure; see [Ozsváth and Szabó 2006].

In fact, we can rephrase this subcase as being about positive surgery on the knot L_1 inside the rational homology sphere $Y_{\Lambda'}(L')$. As such, it was fully treated in [Ozsváth and Szabó 2011, Section 6.3] (see also [Ozsváth and Szabó 2008b, Sections 4.5–4.6]). However, we present a slightly different proof here, which will be easier to adapt to the other cases below, and which gives us the opportunity to establish some notation.

We aim to understand the quasi-isomorphism (f_1^δ, H_1^δ) from (133). We start by studying the map f_1^δ, which corresponds to the cobordism W_1 from Figure 35.

Lemma 11.9 *If $a_1 > 0$, then the cobordism W_1 is negative definite. In fact, a generator Σ_1 of $H_2(W_1; \mathbb{Z}) \cong \mathbb{Z}$ satisfies $\Sigma_1 \cdot \Sigma_1 = -m_1(a_1 d + 1)$.*

Proof Using the matrix form (135) for Λ^+ and the identification (119), the generator of $H_2(W_1;\mathbb{Z})$ is given in coordinates by

$$\Sigma_1 = (a_1 d, \dots, a_\ell d, 1, -m_1). \tag{136}$$

The conclusion then follows from Lemma 11.1(a). □

Using Lemma 11.8, we can specify the relation between the Spin^c structures on W_1 and its two boundary components. Precisely, since $m_1\tau_1 = (m_1, 0, \dots, 0) \in H(L, \Lambda)$, we deduce that Spin^c structures $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$ and $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L)) \cong \mathbb{H}(L)/H(L,\overline{\Lambda})$ are linked if and only if

$$\pi(\bar{\mathfrak{u}}) = \mathfrak{u},$$

where

$$\pi\colon (\mathbb{H}(L)/H(L,\overline{\Lambda})) \to (\mathbb{H}(L)/H(L,\Lambda))$$

is the natural projection. Note that the projection makes sense because

$$H(L,\overline{\Lambda}) = \mathrm{Span}(\Lambda_1 + m_1\tau_1, \Lambda_2, \dots, \Lambda_\ell)$$

is a subset of $H(L,\Lambda)$.

This observation implies that the map f_1^δ is a direct sum of the maps

$$f_{1,\mathfrak{u}}^\delta\colon \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u}) \to \bigoplus_{\substack{\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L)) \\ \pi(\bar{\mathfrak{u}})=\mathfrak{u}}} \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L),\bar{\mathfrak{u}})$$

for $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$. Further, the map $f_{1,\mathfrak{u}}^\delta$ is itself a sum of maps $f_{1,\mathfrak{u};\boldsymbol{t}}^\delta$, one for each Spin^c structure $\boldsymbol{t} \in \mathrm{Spin}^c(W_1)$ with $\psi^{L_{\ell+2}}(\boldsymbol{t}) = \mathfrak{u} \# \mathfrak{u}^{\mathrm{can}}$.

Let us now turn our attention to the map

$$f_2^\delta\colon \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L)) \to \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L');\mathscr{T}_1) \cong \bigoplus_{i=0}^{m_1-1} T_1^i \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L')),$$

whose mapping cone appears in (133). This is simply a twisted coefficient map associated to the cobordism W_2, the reverse of $-W_{\overline{\Lambda}}(L',L)$ in the notation of Section 11.1.

Lemma 11.10 *The cobordism W_2 is negative definite. Its second homology is generated by a class Σ_2 with $\Sigma_2 \cdot \Sigma_2 = -a_1 h(a_1 d + 1)$.*

Proof It is easier to think about the cobordism $W_{\overline{\Lambda}}(L',L)$, for which we can apply Lemma 11.1. Indeed, by part (a) of that lemma, the generator of the second homology is the vector $\boldsymbol{a}$, with $\boldsymbol{a}^t\,\overline{\Lambda}\boldsymbol{a} = a_1 h(a_1 d + 1)$. The change in sign in the final answer is due to the fact that in W_2, the orientation is reversed. □

Using Lemma 11.2, we can relate Spin^c structures on the two boundaries of W_2. By also keeping track of the powers of T_1, we obtain that f_2^δ is a sum of maps

$$f^\delta_{2,\bar{\mathfrak{u}};\boldsymbol{t}}\colon \boldsymbol{CF}^{-,\delta}(Y_{\bar{\Lambda}}(L),\bar{\mathfrak{u}}) \to T_1^{\boldsymbol{t}}\cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'),\psi^{L_1}(\boldsymbol{t})) \tag{137}$$

for $\boldsymbol{t}\in \mathrm{Spin}^c(W_2)\cong \mathbb{H}(L)/H(L,\Lambda|_{L'})$ restricting to $\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar{\Lambda}}(L))$. Here, $T_1^{\boldsymbol{t}}$ denotes $T_1^{\boldsymbol{a}\cdot(\boldsymbol{t}-\boldsymbol{t}_0)}$, where $\boldsymbol{t}_0$ is a fixed Spin^c structure on W_2. Note that the expression $\boldsymbol{a}\cdot(\boldsymbol{t}-\boldsymbol{t}_0)$ is well defined because $\boldsymbol{a}\cdot\boldsymbol{v}=0$ for all $\boldsymbol{v}\in H(L,\Lambda|_{L'})$. (In fact, as can be seen from the proof of Lemma 11.10, multiplication with $\boldsymbol{a}$ represents evaluation on the homology generator.)

Taking the sum of all maps $f^\delta_{2,\bar{\mathfrak{u}};\boldsymbol{t}}$ over all $\boldsymbol{t}$ (with $\bar{\mathfrak{u}}$ fixed), we obtain a map

$$f^\delta_{2,\bar{\mathfrak{u}}}\colon \boldsymbol{CF}^{-,\delta}(Y_{\bar{\Lambda}}(L),\bar{\mathfrak{u}}) \to \bigoplus_{\substack{\boldsymbol{t}\in\mathrm{Spin}^c(W_2)\\ \pi^{L,L'}(\boldsymbol{t})=\bar{\mathfrak{u}}}} T_1^{\boldsymbol{t}}\cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'),\psi^{L_1}(\boldsymbol{t})).$$

By a slight abuse of notation, for $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$, we set

$$f^\delta_{2,\mathfrak{u}} = \sum_{\substack{\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar{\Lambda}}(L))\\ \pi(\bar{\mathfrak{u}})=\mathfrak{u}}} f^\delta_{2,\bar{\mathfrak{u}}}.$$

Note that each term $T_1^i\cdot\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'),\mathfrak{u}')$ for $i=0,\dots,m_1-1$ and $\mathfrak{u}'\in\mathrm{Spin}^c(Y_{\Lambda'}(L'))$ appears in the target of exactly one of the maps $f^\delta_{2,\mathfrak{u}}$. Indeed, two Spin^c structures $\boldsymbol{t}$ and $\boldsymbol{t}'$ on W_2 have the same term in the target if and only if $\boldsymbol{t}-\boldsymbol{t}'\in\mathbb{Z}^\ell/H(L,\Lambda|_{L'})$ has a representative of the form $(dh/a_1,0,\dots,0)$. Since we have chosen d to be a multiple of a_1, and $(h,0,\dots,0)$ is in the span of $\Lambda_1,\dots,\Lambda_\ell$, we get that $\boldsymbol{t}$ and $\boldsymbol{t}'$ have the same reduction modulo $H(L,\Lambda)$, ie they correspond to the same $\mathfrak{u}$. We conclude that the map f_2^δ is the direct sum of all $f^\delta_{2,\mathfrak{u}}$, for $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$.

The last map that appears in (133) is the null-homotopy H_1^δ. This also splits as a direct sum of maps $H^\delta_{1,\mathfrak{u}}$ over $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$, where $H^\delta_{1,\mathfrak{u}}$ denotes the corresponding null-homotopy of $f^\delta_{2,\mathfrak{u}}\circ f^\delta_{1,\mathfrak{u}}$. Moreover, each $H^\delta_{1,\mathfrak{u}}$ is a sum of maps $H^\delta_{1,\mathfrak{u};\boldsymbol{t}}$, over Spin^c structures on $W_1\cup W_2$ restricting to $\mathfrak{u}\#\mathfrak{u}^{\mathrm{can}}$ on the boundary component $Y_\Lambda(L)\,\#\,L(m_1,1)$.

The quasi-isomorphism (133) can then be viewed as a direct sum itself. Precisely, the summand $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u})$ is quasi-isomorphic to the mapping cone of

$$\bigoplus_{\substack{\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar{\Lambda}}(L))\\ \pi(\bar{\mathfrak{u}})=\mathfrak{u}}} \boldsymbol{CF}^{-,\delta}(Y_{\bar{\Lambda}}(L),\bar{\mathfrak{u}}) \xrightarrow{f^\delta_{2,\mathfrak{u}}} \bigoplus_{\substack{\boldsymbol{t}\in\mathrm{Spin}^c(W_2)\\ \pi(\pi^{L,L'}(\boldsymbol{t}))=\mathfrak{u}}} T_1^{\boldsymbol{t}}\cdot\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'),\psi^{L_1}(\boldsymbol{t})) \tag{138}$$

via the map $(f^\delta_{1,\mathfrak{u}},H^\delta_{1,\mathfrak{u}})$.

Note that the direct summands appearing in (138) are absolutely $\mathbb{Q}$-graded and relatively $\mathbb{Z}$-graded (in a compatible way).

We seek to prove:

Proposition 11.11 *Fix $\delta > 0$. Suppose that Λ is nondegenerate, $a_1 > 0$, and m_1 (a multiple of $a_1 h$) is sufficiently large. Then, for every $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$, there is a relative $\mathbb{Z}$-grading on $\mathrm{Cone}(f^\delta_{2,\mathfrak{u}})$ such that the quasi-isomorphism*

$$(f^\delta_{1,\mathfrak{u}}, H^\delta_{1,\mathfrak{u}})\colon \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u}) \xrightarrow{\sim} \mathrm{Cone}(f^\delta_{2,\mathfrak{u}})$$

respects the relative $\mathbb{Z}$-gradings on the two sides.

For this, we need a few lemmas, for which we make use of the absolute $\mathbb{Q}$-gradings on our complexes:

Lemma 11.12 *Let $\vec{L} \subset Y$ and Λ be as above, and fix $\delta > 0$. Then there are constants C_1 and C_2 (depending only on $\vec{L}$, Λ and δ) such that, for all sufficiently large m_1 (chosen as specified above, ie a multiple of $a_1 h$) and for all $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L)) \cong \mathbb{H}(L)/H(L,\overline{\Lambda})$,*

(139) $$\max \mathrm{gr}\, \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L),\bar{\mathfrak{u}}) - \min \mathrm{gr}\, \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L),\bar{\mathfrak{u}}) \le C_1,$$

(140) $$\left| \max \mathrm{gr}\, \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L),\bar{\mathfrak{u}}) - \min_{\substack{\boldsymbol{s}\in\mathbb{H}(L)\\ [\boldsymbol{s}]=\bar{\mathfrak{u}}}} \frac{|(2\boldsymbol{s}-\overline{\Lambda}_1)\cdot \boldsymbol{a}|^2}{4a_1 h(a_1 d+1)} \right| \le C_2.$$

Proof This is an analogue of [Ozsváth and Szabó 2008b, Corollary 2.4], and has a similar proof, so we only sketch the argument. (See also [Ozsváth and Szabó 2011, Lemma 4.6 and the proof of Proposition 4.2].) The first inequality follows from the large surgeries theorem [ibid., Theorem 4.1], which gives a relatively graded identification of $\boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L),\bar{\mathfrak{u}})$ with a generalized Floer complex of L_1 inside $Y_{\Lambda'}(L')$. This latter complex is independent of m_1.

For the second inequality, note that if $\boldsymbol{t} \in \mathrm{Spin}^c(W_2) \cong \mathbb{H}(L)/H(L,\Lambda|_{L'})$, then, using Lemma 11.10 and the formula $\Sigma_2 = \boldsymbol{a}$ for the homology generator in standard coordinates, we get

$$\langle c_1(\boldsymbol{t}), [\Sigma_2]\rangle = (2\boldsymbol{s}-\overline{\Lambda}_1)\cdot \boldsymbol{a},$$

and

(141) $$c_1(\boldsymbol{t})^2 = -\frac{|(2\boldsymbol{s}-\overline{\Lambda}_1)\cdot \boldsymbol{a}|^2}{a_1 h(a_1 d+1)},$$

where $\boldsymbol{s} \in \mathbb{H}(L)$ is any representative of $\boldsymbol{t}$. Consider the map induced on (truncated) Floer homologies by the cobordism W_2 in the Spin^c structure $\boldsymbol{t}$ for which $c_1(\boldsymbol{t})^2$ is maximized along all $\boldsymbol{t}$ that restrict to $\bar{\mathfrak{u}}$. The large surgeries theorem [loc. cit.] also identifies this map with a standard inclusion map between generalized Floer complexes. Since the target graded group $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \psi^{L_1}(\boldsymbol{t}))$ is independent of m_1, inequality (140) follows from the formula (128) for absolute grading shifts. □

Lemma 11.13 *Fix $\vec{L}$, Λ and δ as above, and a constant $C_3 \in \mathbb{R}$. Then there is a constant b with the following property. For all sufficiently large m_1 (divisible by $a_1 h$), for any fixed $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L)) \cong \mathbb{H}(L)/H(L,\overline{\Lambda})$, there are at most two Spin^c structures $\boldsymbol{t}$ on W_2 whose restriction to $Y_{\overline{\Lambda}}(L)$ is $\bar{\mathfrak{u}}$, and with the property that*

(142) $$\max \mathrm{gr}\, \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L),\bar{\mathfrak{u}}) \ge C_3 - \tfrac{1}{4}c_1(\boldsymbol{t})^2;$$

these are the Spin^c *structures* $\boldsymbol{t} = \boldsymbol{t}^{\pm}_{\bar{\mathfrak{u}}}$ *with*

$$\langle c_1(\boldsymbol{t}^{\pm}_{\bar{\mathfrak{u}}}), [\Sigma_2]\rangle = (2\boldsymbol{s} \pm \bar{\Lambda}_1)\cdot \boldsymbol{a},$$

where $\boldsymbol{s} \in \mathbb{H}(L)$ *is any representative of* $\bar{\mathfrak{u}}$ *which satisfies*

$$-\tfrac{1}{2}h(a_1 d+1) \le \boldsymbol{s}\cdot\boldsymbol{a} < \tfrac{1}{2}h(a_1 d+1). \tag{143}$$

All other Spin^c *structures* $\boldsymbol{t}$ *restricting to* $\bar{\mathfrak{u}}$ *satisfy the inequality*

$$c_1(\boldsymbol{t})^2 \le -4m_1.$$

Moreover, if there is no representative $\boldsymbol{s}$ *of* $\bar{\mathfrak{u}}$ *satisfying* $|\boldsymbol{s}\cdot\boldsymbol{a}| < b$, *then there is a unique* Spin^c *structure* $\boldsymbol{t}$ *(restricting to* $\bar{\mathfrak{u}}$*) that satisfies* (142), *namely the one for which* $|\langle c_1(\boldsymbol{t}), [\Sigma_2]\rangle|$ *is minimal.*

Proof Since $\bar{\Lambda}_1\cdot\boldsymbol{a} = h(a_1 d+1)$ and $\Lambda_i\cdot\boldsymbol{a} = 0$ for $i > 1$, each $\bar{\mathfrak{u}} \in \mathbb{H}(L)/H(L,\bar{\Lambda})$ has indeed a representative $\boldsymbol{s}$ satisfying (143), and the value of $\boldsymbol{s}\cdot\boldsymbol{a}$ is independent of that representative.

We use formula (141) for $c_1(\boldsymbol{t})^2$, as well as the inequalities (139) and (140) to verify the statements of the lemma. For more details in a special case, see [Ozsváth and Szabó 2008b, Lemma 4.4]. □

Lemma 11.14 *Fix a constant* C_0. *For all sufficiently large* m_1 *(divisible by* $a_1 h$*), the following statement holds. Each* Spin^c *structure* $\bar{\mathfrak{u}}$ *over* $Y_{\bar{\Lambda}}(L)$ *has at most one extension* $\boldsymbol{t}$ *to* W_1 *whose restriction to* $Y_{\Lambda}(L) \# L(m_1, 1)$ *is* $\pi(\bar{\mathfrak{u}}) \# \mathfrak{u}^{\mathrm{can}}$ *and for which*

$$C_0 \le c_1(\boldsymbol{t})^2 + m_1. \tag{144}$$

Further, if such a Spin^c *structure* $\boldsymbol{t} = \boldsymbol{t}^{\circ}_{\bar{\mathfrak{u}}}$ *exists, then*

$$c_1(\boldsymbol{t}^{\circ}_{\bar{\mathfrak{u}}})^2 = -\frac{4d(\boldsymbol{s}\cdot\boldsymbol{a})^2}{h(a_1 d+1)}, \tag{145}$$

where $\boldsymbol{s} \in \mathbb{H}(L)$ *is the representative of* $[\boldsymbol{s}] = \bar{\mathfrak{u}} \in \mathbb{H}(L)/H(L,\bar{\Lambda})$ *for which the absolute value* $|\boldsymbol{s}\cdot\boldsymbol{a}|$ *is minimal.*

Proof (compare [Ozsváth and Szabó 2008b, Lemma 4.7; 2011, Lemma 6.7]) Since both boundaries of W_1 are rational homology spheres, for any $\boldsymbol{t} \in \mathrm{Spin}^c(W_1)$ we can write $c_1(\boldsymbol{t}) = \alpha\cdot\mathrm{PD}(\Sigma_1)$, for some $\alpha \in \mathbb{Q}$. Here, Σ_1 is the homology generator from Lemma 11.9. Using the computation $\Sigma_1\cdot\Sigma_1 = -m_1(a_1 d+1)$ from that lemma, inequality (144) becomes

$$|\alpha| \le \sqrt{\frac{m_1 - C_0}{m_1(a_1 d+1)}}. \tag{146}$$

(Recall that $d = m_1/h$.) Restriction to the boundaries determines $\boldsymbol{t}$ up to addition of $\mathrm{PD}(\Sigma_1)$, ie $\bar{\mathfrak{u}}$ determines α up to the addition of an even integer. Since the right-hand side of (146) becomes very small when d gets large, in each equivalence class mod $2\mathbb{Z}$ there is at most one α satisfying (146). Inequality (144) follows.

To establish (145), we investigate in more detail the possible values of $\langle c_1(\boldsymbol{t}), [\Sigma_1]\rangle$, over all $\boldsymbol{t}$ which have fixed restrictions $\bar{\mathfrak{u}}$ and $\pi(\bar{\mathfrak{u}}) \# \mathfrak{u}^{\mathrm{can}}$ to the boundaries. Given such a Spin^c structure

$$\boldsymbol{t} \in \mathrm{Spin}^c(W_1) \cong \mathbb{H}(L^+)/\mathrm{Span}(\Lambda_1^+, \dots, \Lambda_\ell^+, \Lambda_{\ell+1}^+),$$

let $\boldsymbol{s}^+ = (s_1^+, \dots, s_\ell^+, s_{\ell+1}^+, s_{\ell+2}^+) \in \mathbb{H}(L^+)$ be a representative of $\boldsymbol{t}$. Recall from the proof of Lemma 11.8 that

$$s_{\ell+1}^+ = m_1 + \tfrac{1}{2}, \tag{147}$$

and the restriction of $\boldsymbol{t}$ to $Y_{\bar{\Lambda}}(L)$ is $\bar{\mathfrak{u}} = [\boldsymbol{s}] \in \mathbb{H}(L)/H(L, \bar{\Lambda})$, where

$$\boldsymbol{s} = (s_1, \dots, s_\ell) = (s_1^+ - s_{\ell+1}^+ + m_1 s_{\ell+2}^+, s_2^+, \dots, s_\ell^+). \tag{148}$$

Using (147), (148) and the formula (136) for the homology generator Σ_1, we get

$$\begin{aligned}\langle c_1(\boldsymbol{t}), \Sigma_1\rangle &= (2\boldsymbol{s}^+ - \Lambda_{\ell+2}^+) \cdot (a_1 d, \dots, a_\ell d, 1, -m_1) \\ &= 2d\boldsymbol{s} \cdot \boldsymbol{a} + 2m_1(a_1 d + 1)(1 - s_{\ell+2}^+) \\ &= 2d(\boldsymbol{s} + \bar{\Lambda}_1(1 - s_{\ell+2}^+)) \cdot \boldsymbol{a}.\end{aligned}$$

Since the equivalence class $\bar{\mathfrak{u}} = [\boldsymbol{s}] \in \mathbb{H}(L)/H(L, \bar{\Lambda})$ is unchanged by the addition of a multiple of $\bar{\Lambda}_1$, by relabeling $\boldsymbol{t}$ we conclude that the possible values of $\langle c_1(\boldsymbol{t}), \Sigma_1\rangle$ (when $\bar{\mathfrak{u}}$ is fixed) are exactly given by

$$\langle c_1(\boldsymbol{t}), \Sigma_1\rangle = 2d\boldsymbol{s} \cdot \boldsymbol{a},$$

where $\boldsymbol{s} \in \mathbb{H}(L)$ is a representative of $\boldsymbol{t}$. Hence,

$$\alpha = \frac{2\boldsymbol{s} \cdot \boldsymbol{a}}{h(a_1 d + 1)} \in \mathbb{Q},$$

and

$$c_1(\boldsymbol{t})^2 = -\frac{4d(\boldsymbol{s} \cdot \boldsymbol{a})^2}{h(a_1 d + 1)}.$$

Thus, $|\alpha|$ is small if and only if $|\boldsymbol{s} \cdot \boldsymbol{a}|$ is small, and (145) follows. □

Proof of Proposition 11.11 We start by equipping the domain of the map $f_{2,\mathfrak{u}}^\delta$ with a relative $\mathbb{Z}$-grading such that $f_{1,\mathfrak{u}}^\delta$ respects the gradings. We do this as follows. Lemma 11.12 supplies a constant C_0 with the property that for all $\boldsymbol{t} \in \mathrm{Spin}^c(W_1)$ with $\psi^{L_{\ell+2}}(\boldsymbol{t}) = \mathfrak{u} \# \mathfrak{u}^{\mathrm{can}}$, the map $f_{1,\mathfrak{u};\boldsymbol{t}}^\delta$ is zero unless Inequality (144) is satisfied. Now, Lemma 11.14 shows that for each $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\bar{\Lambda}}(L))$ such that $\pi(\bar{\mathfrak{u}}) = \mathfrak{u}$, there is at most one extension $\boldsymbol{t} = \boldsymbol{t}_{\bar{\mathfrak{u}}}^\circ$ of $\bar{\mathfrak{u}}$ to W_1 satisfying (144). Now we choose a relative grading on the domain of $f_{2,\mathfrak{u}}^\delta$ so that each map $f_{1,\mathfrak{u};\boldsymbol{t}_{\bar{\mathfrak{u}}}^\circ}$ is grading-preserving.

We grade the range of $f_{2,\mathfrak{u}}^\delta$ so that this map (thought of as differential of the mapping cone) drops grading by one. To do this, we need to check that, for any given term

$$T_1^i \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}')$$

in the target of $f^{\delta}_{2,\mathrm{u}}$, the compositions $f^{\delta}_{2,\mathrm{u};\boldsymbol{t}_2} \circ f^{\delta}_{1,\mathrm{u};\boldsymbol{t}_1}$ which hit that summand induce the same shift in grading. Lemmas 11.13 and 11.14 imply that, if we fix the target, there are at most two pairs $(\boldsymbol{t}_1, \boldsymbol{t}_2)$ for which the corresponding compositions are nonzero. In fact, in many cases there is at most one such pair. If two pairs with nonzero compositions (with the same domain and target) exist, they are of the form

$$(\boldsymbol{t}_1, \boldsymbol{t}_2) = (\boldsymbol{t}^{\circ}_{\bar{\mathrm{u}}}, \boldsymbol{t}^{+}_{\bar{\mathrm{u}}}) \quad \text{and} \quad (\boldsymbol{t}'_1, \boldsymbol{t}'_2) = (\boldsymbol{t}^{\circ}_{\bar{\mathrm{u}}+\Lambda_1}, \boldsymbol{t}^{-}_{\bar{\mathrm{u}}+\Lambda_1}),$$

where $\bar{\mathrm{u}} \in \mathrm{Spin}^c(Y_{\bar{\Lambda}}(L))$ admits a representative $\boldsymbol{s}$ with $|\boldsymbol{s} \cdot \boldsymbol{a}| < b$.

Let $\boldsymbol{s}$ be the representative of $\bar{\mathrm{u}}$ with $|\boldsymbol{s} \cdot \boldsymbol{a}|$ minimal. Using (128), checking that $f^{\delta}_{2,\mathrm{u};\boldsymbol{t}_2} \circ f^{\delta}_{1,\mathrm{u};\boldsymbol{t}_1}$ and $f^{\delta}_{2,\mathrm{u};\boldsymbol{t}'_2} \circ f^{\delta}_{1,\mathrm{u};\boldsymbol{t}'_1}$ shift degree by the same amount is equivalent to showing that

$$c_1(\boldsymbol{t}_1)^2 + c_1(\boldsymbol{t}'_1)^2 = c_1(\boldsymbol{t}_2)^2 + c_1(\boldsymbol{t}_2)^2. \tag{149}$$

Indeed, using Lemma 11.14, we get

$$\begin{aligned} c_1(\boldsymbol{t}_1)^2 - c_1(\boldsymbol{t}'_1)^2 &= \frac{4d|(\boldsymbol{s}+\Lambda_1)\cdot \boldsymbol{a}|^2}{h(a_1 d+1)} - \frac{4d|\boldsymbol{s}\cdot\boldsymbol{a}|^2}{h(a_1 d+1)} \\ &= \frac{4d(\Lambda_1 \cdot \boldsymbol{a})((2\boldsymbol{s}+\Lambda_1)\cdot\boldsymbol{a})}{h(a_1 d+1)} \\ &= \frac{4d(2\boldsymbol{s}+\Lambda_1)\cdot\boldsymbol{a}}{a_1 d+1}. \end{aligned}$$

Using Lemma 11.13, we get

$$\begin{aligned} c_1(\boldsymbol{t}'_2)^2 - c_1(\boldsymbol{t}_2)^2 &= \frac{(2\boldsymbol{s}+\bar{\Lambda}_1)\cdot\boldsymbol{a}}{a_1 h(a_1 d+1)} - \frac{(2\boldsymbol{s}+2\Lambda_1-\bar{\Lambda}_1)\cdot\boldsymbol{a}}{a_1 h(a_1 d+1)} \\ &= \frac{4((\bar{\Lambda}_1-\Lambda_1)\cdot\boldsymbol{a})\cdot((2\boldsymbol{s}+\Lambda_1)\cdot\boldsymbol{a})}{a_1 h(a_1 d+1)} \\ &= \frac{4d(2\boldsymbol{s}+\Lambda_1)\cdot\boldsymbol{a}}{a_1 d+1}. \end{aligned}$$

Equation (149) is therefore satisfied, and we conclude that $f^{\delta}_{1,\mathrm{u}}$ preserves the relative grading. The other component of the quasi-isomorphism under consideration is the null-homotopy $H^{\delta}_{1,\mathrm{u}}$. To check that it is grading-preserving, note that it is a sum of terms $H^{\delta}_{1,\mathrm{u};\boldsymbol{t}}$ over certain Spin^c structures on $W_1 \cup W_2$. The inequalities (139), (140), (142) and (144) imply that the only nonzero terms correspond to Spin^c structures of the form $(\boldsymbol{t}_1, \boldsymbol{t}_2) = (\boldsymbol{t}^{\circ}_{\bar{\mathrm{u}}}, \boldsymbol{t}^{\pm}_{\bar{\mathrm{u}}})$; compare [Ozsváth and Szabó 2008b, Lemma 4.8]. The grading shifts are one less than to the corresponding shifts of the compositions $f^{\delta}_{2,\mathrm{u};\boldsymbol{t}_2} \circ f^{\delta}_{1,\mathrm{u};\boldsymbol{t}_1}$. Hence, since $f^{\delta}_{1,\mathrm{u}}$ preserves the relative grading, so does $H^{\delta}_{1,\mathrm{u}}$. □

Case I(b) $(a_1 < 0)$ The model example is that of negative surgery on a knot $K \subset Y$, which was discussed in [ibid., Section 4.7].

In this situation, again, $Y_{\Lambda}(L)$, $Y_{\bar{\Lambda}}(L)$ and $Y_{\Lambda'}(L')$ are rational homology spheres, so we can use absolute gradings and computations of grading shifts. The computations in Lemmas 11.9 and 11.10 still hold, with the important difference that now $\Sigma_1 \cdot \Sigma_1 = -m_1(a_1 d + 1)$ is positive, hence W_1 is positive definite. Since we need to work with maps induced by negative definite cobordisms, we consider instead the quasi-isomorphism in the other direction (ie (134) rather than (133)), and show that it preserves relative grading. To this end, instead of W_1 we use the negative definite cobordism $W_3 = W_\Lambda(L', L)$ from $Y_{\Lambda'}(L')$ to $Y_\Lambda(L)$, which gives rise to the map f_3^δ.

The quasi-isomorphism (134) is a direct sum of quasi-isomorphisms $(H^\delta_{2,\mathfrak{u}}, f^\delta_{3,\mathfrak{u}})$, one for each $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$.

Proposition 11.15 *Fix $\delta > 0$. If Λ is nondegenerate, $a_1 < 0$ and m_1 (a multiple of $a_1 h$) is sufficiently large, then, for every $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$, the quasi-isomorphism*

$$(H^\delta_{2,\mathfrak{u}}, f^\delta_{3,\mathfrak{u}})\colon \mathrm{Cone}(f^\delta_{2,\mathfrak{u}}) \xrightarrow{\sim} \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u})$$

respects the relative $\mathbb{Z}$-gradings on the two sides.

Sketch of the proof The quasi-isomorphism under consideration is a sum of several maps, corresponding to Spin^c structures on $W_2 \cup W_3$. Analogues of Lemmas 11.13 and 11.14 show that, if we fix the domain and the target of the maps, there at most two Spin^c structures on $W_2 \cup W_3$ which give nonzero maps, and a computation similar to that for (149) shows that those two maps shift the absolute gradings in the same way. Compare [Ozsváth and Szabó 2008b, Section 4.7; 2011, Section 6.4]. □

Case I(c) ($a_1 = 0$, ie Λ' is degenerate) (A model example to keep in mind is surgery on a link of two components $L = L_1 \cup L_2 \subset Y$, where the linking number of L_1 and L_2 is nonzero, and the framing coefficient of L_2 is zero.)

In this situation Lemma 11.9 still holds, we have $\Sigma_1 \cdot \Sigma_1 = -m_1$, so the cobordism W_1 is negative definite. Consequently, we choose to look at the quasi-isomorphism (133), just like in Case I(a). The main difference from that subcase is that now $Y_{\Lambda'}(L')$ has $b_1 = 1$, so the respective Floer complex (with twisted coefficients) does not have a relative $\mathbb{Z}$-grading.

Let us study the relationship between Spin^c structures on the boundaries of the cobordisms W_1 and W_2, just as we did in Case I(a). For W_1, it is still true that $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$ and $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\bar\Lambda}(L))$ are linked if and only if $\pi(\bar{\mathfrak{u}}) = \mathfrak{u}$, where

$$\pi\colon (\mathbb{H}(L)/H(L, \bar\Lambda)) \to (\mathbb{H}(L)/H(L, \Lambda))$$

is the natural projection. Moreover, because now $(m_1, 0, \dots, 0) = \bar\Lambda_1 - \Lambda_1$ is in the span of $\Lambda_2, \dots, \Lambda_\ell$, we have $H(L, \Lambda) = H(L, \bar\Lambda)$, and the projection π is actually a bijection. Thus, the triangle map f_1^δ decomposes as a direct sum of maps

$$f^\delta_{1,\mathfrak{u}}\colon \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u}) \to \boldsymbol{CF}^{-,\delta}(Y_{\bar\Lambda}(L), \pi^{-1}(\mathfrak{u}))$$

for $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$. Just as before, the map $f^\delta_{1,\mathfrak{u}}$ is itself a sum of maps $f^\delta_{1,\mathfrak{u};\boldsymbol{t}}$, over Spin^c structures $\boldsymbol{t} \in \mathrm{Spin}^c(W_1)$ with the given restriction to the boundary.

With regard to the cobordism W_2, we have:

Lemma 11.16 *If Λ is nondegenerate and $a_1 = 0$, then the natural restriction map*

$$\mathrm{Spin}^c(W_2) \to \mathrm{Spin}^c(\partial W_2) \cong \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L)) \times \mathrm{Spin}^c(Y_{\Lambda'}(L'))$$

is injective.

Proof It suffices to prove the analogous statement for second cohomology groups, which is easier because then the restriction map is a homomorphism and we can talk about its kernel. It is also easier to think of the cobordism $W_{\overline{\Lambda}}(L', L)$, which is the reverse of W_2 with the opposite orientation, and apply Lemma 11.1 to it. Suppose $\boldsymbol{s} \in \mathbb{Z}^\ell$ is such that

$$[\boldsymbol{s}] \in H^2(W_{\overline{\Lambda}}(L', L)) \cong \mathbb{Z}^\ell / \mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)$$

has trivial projections to $\mathbb{Z}^\ell / H(L, \overline{\Lambda})$ and to $\mathbb{Z}^{\ell-1}/H(L', \Lambda')$. We need to show that $[\boldsymbol{s}] = 0$.

By assumption, there exist $b_2, \dots, b_\ell$ such that $\boldsymbol{s} - b_2\Lambda_2 - \cdots - b_\ell\Lambda_\ell$ has only the first coordinate nonzero. Hence $[\boldsymbol{s}]$ has a representative $\boldsymbol{s}' = (s_1', 0, \dots, 0)$. Also by hypothesis, $\boldsymbol{s}'$ must be in $H(L, \overline{\Lambda}) = H(L, \Lambda) = \mathrm{Span}(\Lambda_1, \dots, \Lambda_\ell)$. But, since $(h, 0, \dots, 0)$ is in the span of $\Lambda_2, \dots, \Lambda_\ell$, and Λ is nondegenerate, we must have that $\boldsymbol{s}'$ is also in the span of $\Lambda_2, \dots, \Lambda_\ell$. Therefore, $[\boldsymbol{s}] = [\boldsymbol{s}'] = 0$. □

Recall that $\Lambda \cdot (0, a_2, \dots, a_\ell) = (h, 0, \dots, 0)$, so since Λ is nondegenerate, the kernel $H(L', \Lambda')^\perp$ of Λ' (which is identified with $H^1(Y_{\Lambda'}(L'))$; see Section 9.3) must be a copy of $\mathbb{Z}$, generated by $\boldsymbol{a}' = (a_2, \dots, a_\ell)$. If $\Lambda_1' = \Lambda_1|_{L'}$ is the vector of linking numbers between L_1 and the other components, we have $\Lambda_1' \cdot \boldsymbol{a}' = h$, ie the generator $\boldsymbol{a}'$ of $H^1(Y_{\Lambda'}(L'))$ evaluates to h on $[L_1]$. We define an $\mathbb{F}[H^1(Y_{\Lambda'}(L'))]$-module

$$\mathcal{V}_1 = \mathbb{F}[V_1]/(V_1^d - 1),$$

where $c \in H^1(Y_{\Lambda'}(L'))$ acts on $\mathcal{V}_1$ by multiplication by $V_1^{\langle c, [L_1]\rangle / h}$.

If we equate $T_1^h = V_1$, it is easy to check that for every $\mathfrak{u}' \in \mathrm{Spin}^c(Y_{\Lambda'}(L'))$ we have an isomorphism

$$\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathcal{T}_1) \cong \bigoplus_{i=0}^{h-1} T_1^i \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathcal{V}_1).$$

Lemma 11.16 implies that f_2^δ is a direct sum of maps

$$f^\delta_{2,\mathfrak{u},\mathfrak{u}'} \colon \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L), \pi^{-1}(\mathfrak{u})) \to T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathcal{V}_1),$$

each either zero or corresponding to the unique Spin^c structure on W_2 that extends $\pi^{-1}(\mathfrak{u}) \in \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L))$ and $\mathfrak{u}' \in \mathrm{Spin}^c(Y_{\Lambda'}(L'))$. Here, $\mathfrak{u}$ denotes a Spin^c structure on $Y_\Lambda(L)$, and $T_1^{\mathfrak{u}}$ is shorthand for $T_1^{(\mathfrak{u}-\mathfrak{u}_0)\cdot \boldsymbol{a}}$, where $\mathfrak{u}_0$ is a fixed Spin^c structure on $Y_\Lambda(L)$. (Observe that the dot product with $\boldsymbol{a}$ is well defined modulo m_1.) We let $f^\delta_{2,\mathfrak{u}}$ be the sum of $f^\delta_{2,\mathfrak{u},\mathfrak{u}'}$, over all possible $\mathfrak{u}'$ (fixing $\mathfrak{u}$).

For $\mathfrak{u}' = [(s_2, \dots, s_\ell)] \in \mathrm{Spin}^c(Y_{\Lambda'}(L')) \cong \mathbb{H}(L')/H(L', \Lambda')$, using (125) we get

$$\mathfrak{d}(\mathfrak{u}', \mathscr{T}_1) = \mathfrak{d}(\mathfrak{u}', \mathscr{V}_1) = \gcd_{\substack{\boldsymbol{v} \in H(L',\Lambda')^\perp \\ \Lambda'_1 \cdot \boldsymbol{v} \equiv 0 \,(\mathrm{mod}\, m_1)}} \sum_{i=2}^{\ell} 2 s_i v_i .$$

Note that $\Lambda'_1 \cdot \boldsymbol{v} \equiv 0 \,(\mathrm{mod}\, m_1)$, or equivalently $(\Lambda'_1/h) \cdot \boldsymbol{v} \equiv 0 \,(\mathrm{mod}\, d)$, is the condition on $\boldsymbol{v} \in H(L', \Lambda)^\perp \cong H^1(Y_{\Lambda'}(L'))$ to act trivially on the module $\mathscr{T}_1$ (or, equivalently, on $\mathscr{V}_1$).

Recall that $H(L', \Lambda')^\perp$ is a copy of $\mathbb{Z}$, generated by $\boldsymbol{a}' = (a_2, \dots, a_\ell)$. For $v \in \mathbb{Z}$, we have $\Lambda'_1 \cdot (v\boldsymbol{a}) = vh$, which is divisible by $m_1 = dh$ if and only if v is divisible by d. Hence,

$$\mathfrak{d}(\mathfrak{u}', \mathscr{V}_1) = d \cdot \left| \sum_{i=2}^{\ell} 2 s_i a_i \right| .$$

Note that this is always divisible by $2d$, because the expression inside the absolute value is even (see Remark 9.8). Therefore, in particular, the Floer complexes $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{V}_1)$ have relative $\mathbb{Z}/2d\mathbb{Z}$-gradings.

Further, for any $\boldsymbol{t} \in \mathrm{Spin}^c(W_2)$ whose restriction to $Y_{\Lambda'}(L')$ is $\mathfrak{u}'$, equation (130) says that $\mathfrak{d}(\boldsymbol{t}, \mathscr{V}_1) = \mathfrak{d}(\boldsymbol{t}, \mathscr{T}_1) = \mathfrak{d}(\mathfrak{u}', \mathscr{T}_1)$. Thus, according to Lemma 11.3(a), every map $f^{\delta}_{2,\mathfrak{u},\mathfrak{u}'}$ preserves the relative $\mathbb{Z}/2d\mathbb{Z}$-gradings.

We equip the mapping cone of $f^{\delta}_{2,\mathfrak{u}}$ with the relative $\mathbb{Z}/2d\mathbb{Z}$-grading which gets decreased by one under $f^{\delta}_{2,\mathfrak{u}}$. Note that the target of $f^{\delta}_{2,\mathfrak{u}}$ is a direct sum of terms corresponding to possible $\mathfrak{u}'$, which a priori have unrelated relative $\mathbb{Z}/2d\mathbb{Z}$-gradings. As part of the mapping cone, however, their direct sum becomes relatively $\mathbb{Z}/2d\mathbb{Z}$-graded.

The second ingredient in (133), the null-homotopy H^{δ}_1, is then a direct sum of maps $H^{\delta}_{1,\mathfrak{u}}$, which represent null-homotopies of $f^{\delta}_{2,\mathfrak{u}} \circ f^{\delta}_{1,\mathfrak{u}}$.

Proposition 11.17 *Fix $\delta > 0$. If Λ is nondegenerate and $a_1 = 0$, then, for every $m_1 = dh$ sufficiently large and for every $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$, the quasi-isomorphism*

$$(f^{\delta}_{1,\mathfrak{u}}, H^{\delta}_{1,\mathfrak{u}}) \colon \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u}) \xrightarrow{\sim} \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})$$

respects the relative $\mathbb{Z}/2d\mathbb{Z}$-gradings on the two sides.

Instead of Lemma 11.14, we now have the following:

Lemma 11.18 *Fix a constant C_0. For all sufficiently large $m_1 = dh$, the following statement holds. Each Spin^c structure $\bar{\mathfrak{u}}$ over $Y_{\overline{\Lambda}}(L)$ has at most two extensions $\boldsymbol{t}$ to W_1 whose restrictions to $Y_\Lambda(L) \# L(m_1, 1)$ are $\pi(\bar{\mathfrak{u}}) \# \mathfrak{u}^{\mathrm{can}}$ and for which*

$$C_0 \le c_1(\boldsymbol{t})^2 + m_1. \tag{150}$$

Further, if two such Spin^c structures $\boldsymbol{t}$ and $\boldsymbol{t}'$ exist, then

$$c_1(\boldsymbol{t})^2 = c_1(\boldsymbol{t}')^2. \tag{151}$$

Proof Just as in the proof of Lemma 11.14, we find that if $\boldsymbol{t}$ satisfies (150) and $c_1(\boldsymbol{t}) = \alpha \cdot \mathrm{PD}(\Sigma_1)$, then

$$|\alpha| \le \sqrt{\frac{m_1 - C_0}{m_1(a_1 d + 1)}} = \sqrt{1 - \frac{C_0}{m_1}}. \tag{152}$$

Also, $\alpha = 2(\boldsymbol{s} \cdot \boldsymbol{a})/h$ for some representative $\boldsymbol{s} \in \mathbb{H}(L)$ of $\boldsymbol{t}$, so α must live in the discrete subset $\frac{1}{h} \cdot \mathbb{Z} \subset \mathbb{Q}$. Hence, if we choose $m_1 > -C_0 h/(2h+1)$, the only way inequality (152) is satisfied is if

$$|\alpha| \le 1.$$

If we fix the restriction of $\boldsymbol{t}$ to ∂W_1, the value of α is determined up to the addition of even integers. Typically there is at most one $\boldsymbol{t}$ with $|\alpha| \le 1$. The only time there are two such Spin^c structures $\boldsymbol{t}$ and $\boldsymbol{t}'$, the respective values of α are ± 1, and we have

$$c_1(\boldsymbol{t})^2 = c_1(\boldsymbol{t}')^2 = -m_1. \qquad \square$$

Proof of Proposition 11.17 Using Lemma 11.18 and formula (128), we find that each $f^{\delta}_{1,\mathrm{u}}$ is the sum of at most two nonzero maps corresponding to Spin^c structures on W_1, and if there are two such maps, they shift absolute grading by the same amount. Thus, $f^{\delta}_{1,\mathrm{u}}$ preserves the relative $\mathbb{Z}$-gradings on the two sides (and hence their $\mathbb{Z}/2d\mathbb{Z}$ reductions). From our discussion of W_2 we also know that each $f^{\delta}_{2,\mathrm{u}}$ preserves relative grading. The Spin^c structures that give nontrivial contributions to $H^{\delta}_{1,\mathrm{u}}$ are subject to similar constraints, and shift grading by one degree less than the respective maps $f^{\delta}_{2,\mathrm{u}} \circ f^{\delta}_{1,\mathrm{u}}$. $\square$

11.6 Refinements of Proposition 11.6, Case II: $\Lambda_1 \in \mathrm{Span}_{\mathbb{Q}}(\Lambda_2, \dots, \Lambda_\ell)$

When the framing matrix Λ is degenerate, we will discuss refinements only in the case when Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$. (A model example for this is 0 surgery on a knot in the integral homology sphere Y.) Note that, in the present case, $b_1(Y_{\overline{\Lambda}}(L)) = b_1(Y_{\Lambda'}(L)) = b_1(Y_{\Lambda}(L)) - 1$.

In this situation we are free to choose whether to consider the quasi-isomorphism in (133), or the one in (134). We will focus on the quasi-isomorphism (133), which involves the cobordisms W_1 and W_2. We denote by h the smallest positive integer such that

$$h\Lambda_1 \in \mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell).$$

Lemma 11.19 *If Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$, then the vector $\tau_1 = (1, 0, \dots, 0)$ is not in the $\mathbb{Q}$-span of $\Lambda_1, \dots, \Lambda_\ell$.*

Proof Let us view the framing matrix Λ as a self-adjoint linear operator on $\mathbb{Q}^\ell$ (with the standard inner product). Since Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$, there exists a vector $\boldsymbol{v} = (v_1, v_2, \dots, v_\ell)$ in the kernel of Λ, with $v_1 \neq 0$. The kernel is orthogonal to the image of Λ, so, since $\boldsymbol{v} \cdot \tau_1 = v_1 \neq 0$, the image cannot contain τ_1. $\square$

Lemma 11.20 *The vector $j\tau_1$ is in $\mathrm{Span}(\overline{\Lambda}_1, \Lambda_2, \dots, \Lambda_\ell)$ if and only if j is divisible by $m_1 h$.*

Proof Writing

$$j\tau_1 = v_1(\Lambda_1 + m_1\tau_1) + v_2\Lambda_2 + \cdots + v_\ell\Lambda_\ell,$$

we have $(j - v_1 m_1)\tau_1 = v_1\Lambda_1 + \cdots + v_\ell\Lambda_\ell$. This can only happen if $j = v_1 m_1$ (by Lemma 11.19) and v_1 is divisible by h (by the definition of h). $\square$

Lemma 11.21 *Suppose Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$, and let $\Lambda'_1 = \Lambda_1|_{L'}$ be the vector of linking numbers between L_1 and the other components. Then $j\Lambda'_1 \in \mathrm{Span}(\Lambda'_2, \dots, \Lambda'_\ell)$ if and only if j is a multiple of h.*

Proof If i is a multiple of h, we already know that $i\Lambda_1 \in \mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)$. For the converse, suppose $j\Lambda'_1 = \sum v_i \Lambda'_2$ for some $v_i \in \mathbb{Z}, i = 2, \dots, \ell$. Then $j\Lambda_1 - \sum v_i\Lambda_i$ is a multiple of τ_1. From Lemma 11.19 we see that it must be zero, so j is a multiple of h. $\square$

Lemma 11.22 *If Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$, the natural restriction map*

$$\mathrm{Spin}^c(W_1) \to \mathrm{Spin}^c(\partial W_1) \cong \mathrm{Spin}^c(Y_\Lambda(L) \# L(m_1, 1)) \times \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L))$$

is injective.

Proof We prove the similar statement for second cohomology groups, by looking at the kernel of the corresponding homomorphism. Suppose $\boldsymbol{s} = (s_1, \dots, s_{\ell+2}) \in \mathbb{Z}^{\ell+2}$ is such that

$$[\boldsymbol{s}] \in H^2(W_1) \cong \mathbb{Z}^{\ell+2}/\mathrm{Span}(\Lambda_1^+, \dots, \Lambda_{\ell+1}^+)$$

has trivial projection to the boundary. We need to show that $[\boldsymbol{s}] = 0$.

By hypothesis, the vector $(s_1, \dots, s_\ell)$ is in the image of Λ, and $s_{\ell+1}$ is a multiple of m_1. By adding suitable multiples of $\Lambda_1^+, \dots, \Lambda_{\ell+1}^+$, we can assume that $\boldsymbol{s} = (0, \dots, 0, s_{\ell+2})$ without changing the equivalence class $[\boldsymbol{s}]$. Also by hypothesis, there exist integers b_i for $i = 1, \dots, \ell+2$ such that

$$\boldsymbol{s} = b_1\Lambda_1^+ + \cdots + b_{\ell+2}\Lambda_{\ell+2}^+.$$

Restricting attention to the first ℓ coordinates of $\boldsymbol{s}$, $(b_{\ell+2}, 0, \dots, 0)$ is in the span of $\Lambda_1, \dots, \Lambda_\ell$. From Lemma 11.19 we see that $b_{\ell+2} = 0$. Hence $\boldsymbol{s}$ is in the span on $\Lambda_1^+, \dots, \Lambda_{\ell+1}^+$, as desired. $\square$

Let us define an equivalence relation on Spin^c structures on $Y_\Lambda(L)$. For $\mathfrak{u}_1, \mathfrak{u}_2 \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L, \Lambda)$, we say $\mathfrak{u}_1$ is equivalent to $\mathfrak{u}_2$ if and only if there exist $\boldsymbol{s}_1, \boldsymbol{s}_2 \in \mathbb{H}(L)$ with

$$[\boldsymbol{s}_1] = \mathfrak{u}_1, \quad [\boldsymbol{s}_2] = \mathfrak{u}_2 \quad \text{and} \quad \boldsymbol{s}_1 - \boldsymbol{s}_2 = j m_1 \tau_1$$

for some $j \in \mathbb{Z}$.

We call a Spin^c structure $\mathfrak{u}$ on $Y_\Lambda(L)$ *small* if our complex $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L)) = \boldsymbol{CF}^{-,\delta}(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w})$ has at least one generator in that Spin^c structure. Clearly there are only finitely many small Spin^c structures.

Further, for $m_1 \gg 0$, using Lemma 11.19 we see that every equivalence class of Spin^c structures on $Y_\Lambda(L)$ contains at most one small structure. Let us pick one representative (in $\mathrm{Spin}^c(Y_\Lambda(L))$) from each equivalence class, in such a way that all small structures are picked. We call the chosen Spin^c structures *special*. By construction, for every special $\mathfrak{u}$,

$$\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u}) = \bigoplus_{\mathfrak{u}_1\sim\mathfrak{u}} \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u}). \tag{153}$$

Building up on Definition 11.7, we write $\mathfrak{u}\sim\bar{\mathfrak{u}}$ if $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$ and $\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar\Lambda}(L))$ are linked. Further, we say that two Spin^c structures $\bar{\mathfrak{u}}$ on $Y_{\bar\Lambda}(L)$ and $\mathfrak{u}'$ on $Y_{\Lambda'}(L')$ are *linked* (and write $\bar{\mathfrak{u}}\sim\mathfrak{u}'$) if there exists a Spin^c structure on W_2 interpolating between the two. We also say that $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$ is *linked* to $\mathfrak{u}'\in\mathrm{Spin}^c(Y_{\Lambda'}(L'))$ (and write $\mathfrak{u}\sim\mathfrak{u}'$) if there exists $\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar\Lambda}(L))$ such that $\mathfrak{u}\sim\bar{\mathfrak{u}}$ and $\bar{\mathfrak{u}}\sim\mathfrak{u}'$.

The following lemma describes how the Spin^c structures on the three manifolds $Y_\Lambda(L)$, $Y_{\bar\Lambda}$ and $Y_{\Lambda'}$ are linked to each other:

Lemma 11.23 *Suppose that Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2,\dots,\Lambda_\ell$, and $m_1>0$ is sufficiently large. Then*:

(a) *Every $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$ is linked to exactly h Spin^c structures on $Y_{\bar\Lambda}(L)$, while every $\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar\Lambda}(L))$ is linked to exactly one special Spin^c structure on $Y_\Lambda(L)$.*

(b) *Every $\bar{\mathfrak{u}}\in\mathrm{Spin}^c(Y_{\bar\Lambda}(L))$ is linked to exactly h Spin^c structures on $Y_{\Lambda'}(L')$, while every $\mathfrak{u}'\in\mathrm{Spin}^c(Y_{\Lambda'}(L'))$ is linked to exactly m_1h Spin^c structures on $Y_{\bar\Lambda}(L)$.*

(c) *Every $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$ is linked to exactly h Spin^c structures on $Y_{\Lambda'}(L')$, while every $\mathfrak{u}'\in\mathrm{Spin}^c(Y_{\Lambda'}(L'))$ is linked to exactly m_1 special Spin^c structures on $Y_\Lambda(L)$.*

Before proving the lemma, it is useful to illustrate its content in a particular example, shown graphically in Figure 36. We consider a framed link $L=L_1\cup L_2$ with framing matrix

$$\Lambda=\begin{pmatrix}1&3\\3&9\end{pmatrix},$$

so that $h=3$ and $\mathbb{H}(L)=\left(\mathbb{Z}+\frac{1}{2}\right)^2$. The set of Spin^c structures on $Y_\Lambda(L)$ is identified with $\mathbb{H}(L)/\langle(1,3)\rangle$. The ones that differ by $(m_1,0)$ are called equivalent. Thus, there are only $3m_1$ special Spin^c structures (one from each equivalence class). The set of special Spin^c structures on $Y_\Lambda(L)$ is the left rectangle in the picture (where $m_1=7$), with each structure being represented by a black dot. Going down one square in the rectangle (in a cyclical fashion) corresponds to adding the vector $(1,0)$. Going left corresponds to adding the vector $(0,1)$, and is also done in a cyclical fashion, except when we go from the third to the first column we also move one step up, according to the relation $(0,3)=(-1,0)$.

The second rectangle in Figure 36 represents the space of all $9m_1$ Spin^c structures on $Y_{\bar\Lambda}(L)$, where

$$\bar\Lambda=\begin{pmatrix}m_1+1&3\\3&9\end{pmatrix}.$$

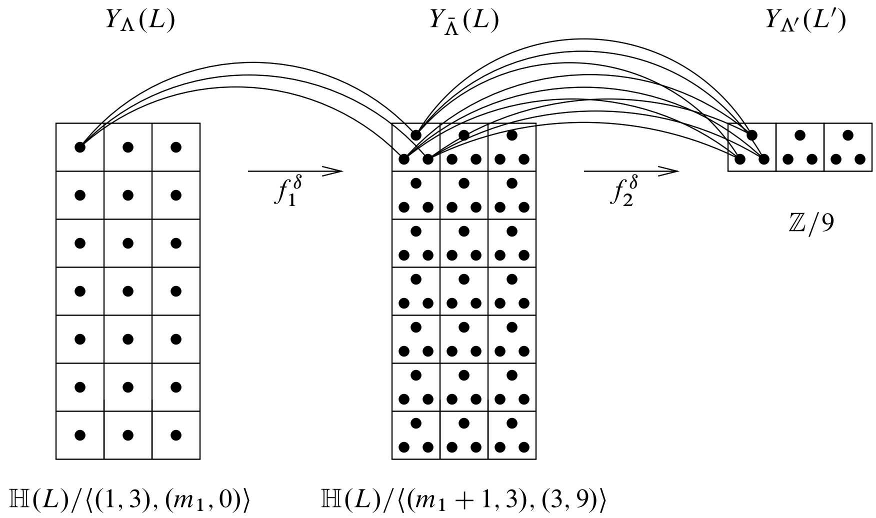

Figure 36: Linking of Spin^c structures. Each dot in the first rectangle corresponds to a special Spin^c structure on $Y_\Lambda(L)$, and is linked to the three dots (representing Spin^c structures on $Y_{\overline{\Lambda}}(L)$) in the corresponding square of the second rectangle. Furthermore, all the dots in the same column of the second rectangle are linked to the three dots (representing Spin^c structures on $Y_{\Lambda'}(L')$) in the corresponding column of the third rectangle.

The dots in the same square differ from each other by multiples of $(m_1, 0)$. Going down one square in the rectangle still means adding $(1, 0)$, and going left means adding $(0, 1)$. When we go horizontally in a cycle we move up one square as well, according to the relation $(0, 3) = (-m_1 - 1, 0)$.

Finally, the rectangle on the right of Figure 36 represents the set of Spin^c structures on $Y_{\Lambda'}(L')$, which is simply identified with $\mathbb{Z}/9$. The dots in the same square differ from each other by multiples of 3.

The linking of Spin^c structures between the three rectangles is as shown in Figure 36, and corresponds to the description in Lemma 11.23.

Proof of Lemma 11.23 (a) Consider two Spin^c structures $\bar{\mathfrak{u}}_1$ and $\bar{\mathfrak{u}}_2$ on $Y_{\overline{\Lambda}}(L)$. By Lemma 11.8, $\bar{\mathfrak{u}}_1$ and $\bar{\mathfrak{u}}_2$ are linked to a single $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$ if and only if the difference

$$\bar{\mathfrak{u}}_1 - \bar{\mathfrak{u}}_2 \in H^2(Y_{\overline{\Lambda}}(L)) \cong \mathbb{Z}^\ell / H(L, \overline{\Lambda})$$

can be represented (in $\mathbb{Z}^\ell$) by a multiple of $m_1\tau_1 = \overline{\Lambda}_1 - \Lambda_1$. The smallest such multiple that lies in the span of $\overline{\Lambda}_1, \dots, \Lambda_\ell$ (ie represents a trivial cohomology class) is $h(\overline{\Lambda}_1 - \Lambda_1)$.

In the other direction, consider two Spin^c structures $\mathfrak{u}_1$ and $\mathfrak{u}_2$ on $Y_\Lambda(L)$. Applying Lemma 11.8 again, we see that $\mathfrak{u}_1$ and $\mathfrak{u}_2$ can be linked to a single $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\overline{\Lambda}}(L))$ if and only if the difference

$$\mathfrak{u}_1 - \mathfrak{u}_2 \in H^2(Y_\Lambda(L)) \cong \mathbb{Z}^\ell / H(L, \Lambda)$$

can be represented (in $\mathbb{Z}^{\ell}$) by a multiple of $m_1\tau_1$, ie $\mathfrak{u}_1$ and $\mathfrak{u}_2$ are equivalent. There is a unique special structure in each equivalence class, by construction.

(b) Apply Lemma 11.1(b) to $W_{\overline{\Lambda}}(L', L)$, which is $-W_2$ turned upside down. Two structures $\mathfrak{u}'_1, \mathfrak{u}'_2 \in \mathrm{Spin}^c(Y_{\Lambda'}(L')) \cong \mathbb{H}(L')/H(L', \Lambda')$ are linked to a single one on $Y_{\overline{\Lambda}}(L)$ if and only if

$$\mathfrak{u}'_1 - \mathfrak{u}'_2 \in H^2(Y_{\Lambda'}(L')) \cong \mathbb{Z}^{\ell-1}/H(L', \Lambda')$$

can be represented (in $\mathbb{Z}^{\ell-1}$) by a multiple of $\overline{\Lambda}_1|_{L'} = \Lambda'_1$. From Lemma 11.21 we see that $j\Lambda'_1$ is zero in cohomology if and only if the respective factor j is divisible by h.

In the other direction, two Spin^c structures $\bar{\mathfrak{u}}_1, \bar{\mathfrak{u}}_2$ on $Y_{\overline{\Lambda}}(L)$ are linked to a single one on $Y_{\Lambda'}(L')$ if and only if $\bar{\mathfrak{u}}_1 - \bar{\mathfrak{u}}_2$ can be represented (in $\mathbb{Z}^{\ell}$) by a multiple of τ_1. By Lemma 11.20, the first such multiple that lies in $H(L, \overline{\Lambda})$ is $m_1 h \tau_1$.

(c) Use the descriptions of linking in (a) and (b). We get that two Spin^c structures $\mathfrak{u}'_1$ and $\mathfrak{u}'_2$ on $Y_{\Lambda'}(L')$ are linked to a single one on $Y_{\Lambda}(L)$ if and only if $\mathfrak{u}'_1 - \mathfrak{u}'_2$ can be represented by a multiple of Λ'_1, and two Spin^c structures $\mathfrak{u}_1$ and $\mathfrak{u}_2$ on $Y_{\Lambda}(L)$ are linked to a single one on $Y_{\Lambda'}(L')$ if and only if $\mathfrak{u}_1 - \mathfrak{u}_2$ can be represented by a multiple of τ_1. □

From (153) and Lemma 11.23(a), we see that the cobordism map f_1^{δ} splits into the direct sum of the maps

$$f_{1,\mathfrak{u}}^{\delta} \colon \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}) \to \bigoplus_{\bar{\mathfrak{u}} \sim \mathfrak{u}} \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L), \bar{\mathfrak{u}})$$

over special Spin^c structures $\mathfrak{u}$.

Turning our attention to the map f_2^{δ}, note that the action of $H^1(Y_{\Lambda'}(L')) \cong H(L', \Lambda')^{\perp}$ on the module $\mathscr{T}_1 = \mathbb{F}[T_1]/(T_1^{m_1} - 1)$ is trivial. Indeed, if $\boldsymbol{v}' \in \mathbb{Z}^{\ell-1}$ satisfies $\boldsymbol{v}' \cdot \Lambda'_i = 0$ for all $i = 2, \dots, \ell$, we must also have $\boldsymbol{v}' \cdot \Lambda'_1 = 0$ (because Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_{\ell}$).

Hence, for every $\mathfrak{u}' \in \mathrm{Spin}^c(Y_{\Lambda'}(L'))$, we have a decomposition

$$\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{T}_1) \cong \bigoplus_{i=0}^{m_1-1} T_1^i \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}').$$

Moreover, using the description of linking in Lemma 11.23(b) (see also Figure 37), f_2^{δ} splits as a direct sum of maps

$$f_{2,\mathfrak{u}}^{\delta} \colon \bigoplus_{\bar{\mathfrak{u}} \sim \mathfrak{u}} \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L), \bar{\mathfrak{u}}) \to \bigoplus_{\mathfrak{u}' \sim \mathfrak{u}} T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}')$$

over special Spin^c structures $\mathfrak{u}$ on $Y_{\Lambda}(L)$. Here, $T_1^{\mathfrak{u}}$ denotes $T_1^{(\mathfrak{u}-\mathfrak{u}_0)\cdot \boldsymbol{v}}$, where $\mathfrak{u}_0$ is a fixed Spin^c structure on $Y_{\Lambda}(L)$, and $\boldsymbol{v}$ is an arbitrary vector in $H_2(W_2) \cong H(L, \overline{\Lambda}|_{L'})^{\perp}$.

Using Lemma 11.23 again, we also obtain a similar decomposition of H_2^{δ} into maps $H_{2,\mathfrak{u}}^{\delta}$, one for each special Spin^c structure $\mathfrak{u}$ on $Y_{\Lambda}(L)$. Each $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u})$ is quasi-isomorphic to the mapping cone of the respective $f_{2,\mathfrak{u}}^{\delta}$, via the pair $(f_{1,\mathfrak{u}}^{\delta}, H_{1,\mathfrak{u}}^{\delta})$.

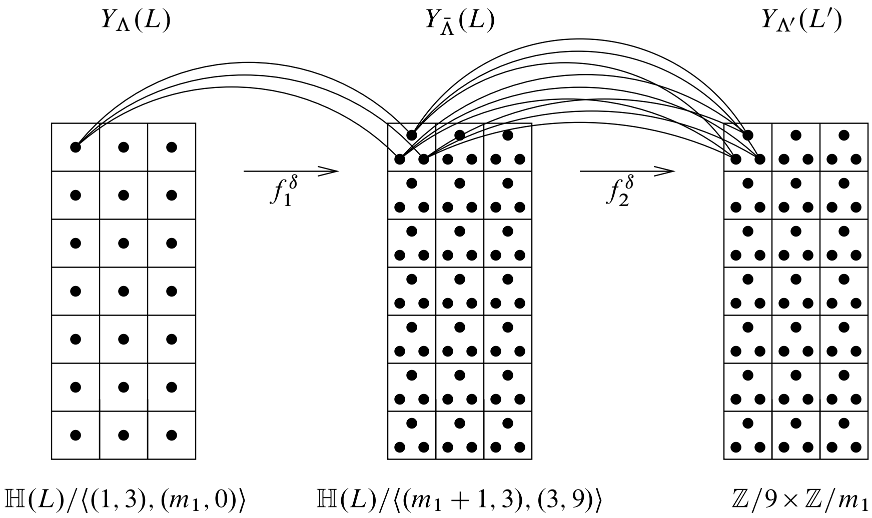

Figure 37: Linking of Spin^c structures, keeping track of the powers of T_1. The map f_2^δ decomposes as a sum according to powers of T_1. Each dot in the rectangle on the right represents a pair $(\mathfrak{u}', i)$, where $\mathfrak{u}' \in \mathrm{Spin}^c(Y_{\Lambda'}(L'))$ and $i \in \mathbb{Z}/_{m_1}$ is an exponent of T_1. Going down one square in the rectangle corresponds to multiplication by T_1. This way, a dot in the second rectangle is linked (that is, gives a contribution to f_2^δ) only with the three dots in the corresponding square of the third rectangle. Therefore, the pair of maps (f_1^δ, f_2^δ) decomposes as a direct sum according to dots in the first rectangle (ie special Spin^c structures $\mathfrak{u}$ on $Y_\Lambda(L)$).

Proposition 11.24 *Fix $\delta > 0$ and suppose that Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$. Then, for every $m_1 \gg 0$ and for every special* Spin^c *structure $\mathfrak{u}$ on $\mathrm{Spin}^c(Y_\Lambda(L))$ such that m_1 is divisible by $\mathfrak{d}(\mathfrak{u})$, the mapping cone* $\mathrm{Cone}(f_{2,\mathfrak{u}}^\delta)$ *admits a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading, and the quasi-isomorphism*

$$(f_{1,\mathfrak{u}}^\delta, H_{1,\mathfrak{u}}^\delta)\colon \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u}) \xrightarrow{\sim} \mathrm{Cone}(f_{2,\mathfrak{u}}^\delta)$$

respects the relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-gradings on the two sides.

Before proceeding with the proof, we establish a few lemmas. First, note that for $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$, equation (105) gives

$$\mathfrak{d}(\mathfrak{u}) = \gcd_{\substack{\boldsymbol{v}\in\mathbb{Z}^\ell \\ \Lambda\boldsymbol{v}=0}} 2\boldsymbol{v}\cdot\boldsymbol{s} \tag{154}$$

for $[\boldsymbol{s}] = \mathfrak{u}$.

Lemma 11.25 *Suppose $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$ and $\bar{\mathfrak{u}} \in \mathrm{Spin}^c(Y_{\bar\Lambda}(L))$ are linked. Then $\mathfrak{d}(\mathfrak{u})$ divides $\mathfrak{d}(\bar{\mathfrak{u}})$.*

Proof If $[\boldsymbol{s}] = \mathfrak{u}$, then, according to the proof of Lemma 11.23(a), $[\boldsymbol{s} + jm_1\tau_1] = \bar{\mathfrak{u}}$ for some $j \in \mathbb{Z}$. We have

$$\mathfrak{d}(\bar{\mathfrak{u}}) = \gcd_{\substack{\boldsymbol{v}=(v_1,\dots,v_\ell)\in\mathbb{Z}^\ell \\ \bar\Lambda\boldsymbol{v}=0}} 2(\boldsymbol{v}\cdot\boldsymbol{s} + jv_1m_1).$$

From Lemma 11.19 we see that $\bar{\Lambda}\boldsymbol{v} = m_1 v_1 \tau_1 + \Lambda\boldsymbol{v} = 0$ implies $v_1 = 0$ and $\Lambda\boldsymbol{v} = 0$. Hence, $2(\boldsymbol{v}\cdot\boldsymbol{s} + j v_1 m_1) = 2\boldsymbol{v}\cdot\boldsymbol{s}$ is one of the elements for which we take the greatest common divisor in (154). Hence, $\mathfrak{d}(\mathfrak{u})$ divides $\mathfrak{d}(\bar{\mathfrak{u}})$. □

Lemma 11.26 *Suppose* $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$ *and* $\mathfrak{u}' \in \mathrm{Spin}^c(Y_{\Lambda'}(L'))$ *are linked. Then* $\mathfrak{d}(\mathfrak{u})$ *divides* $\mathfrak{d}(\mathfrak{u}')$.

Proof If $\mathfrak{u} = [\boldsymbol{s}]$ with $\boldsymbol{s} = (s_1, \dots, s_\ell)$ and we let $\boldsymbol{s}' = (s_2, \dots, s_\ell)$, then, according to the proof of Lemma 11.23(c), $\mathfrak{u}' = [\boldsymbol{s}' + j\Lambda'_1]$ for some $j \in \mathbb{Z}$. We have

$$\mathfrak{d}(\mathfrak{u}') = \gcd_{\substack{\boldsymbol{v}'\in\mathbb{Z}^{\ell-1}\\ \Lambda'\boldsymbol{v}'=0}} 2\boldsymbol{v}'\cdot(\boldsymbol{s}' + j\Lambda'_1).$$

Note that $\Lambda'\boldsymbol{v}' = 0$ implies $\Lambda'_1\cdot\boldsymbol{v}' = 0$, so the expression in the gcd is $2\boldsymbol{v}'\cdot\boldsymbol{s}'$. Moreover, if we let $\boldsymbol{v} \in \mathbb{Z}^\ell$ be the vector with first coordinate zero and the others given by $\boldsymbol{v}'$, then $\Lambda\boldsymbol{v} = (u, 0, \dots, 0)$ for some $u \in \mathbb{Z}$. Lemma 11.19 implies that $u = 0$, so $\boldsymbol{v}$ gets counted in (154), and the conclusion follows. □

Proof of Proposition 11.24 Let $\mathfrak{u}$ be a special Spin^c structure on $Y_\Lambda(L)$ with m_1 divisible by $\mathfrak{d}(\mathfrak{u})$. Lemma 11.3(a), together with Lemmas 11.22 and 11.25 imply that we can equip

$$\bigoplus_{\bar{\mathfrak{u}}\sim\mathfrak{u}} \boldsymbol{CF}^{-,\delta}(Y_{\bar{\Lambda}}(L), \bar{\mathfrak{u}})$$

with a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading such that $f^\delta_{1,\mathfrak{u}}$ is grading-preserving. Furthermore, Lemma 11.26 shows that each direct summand $T_1^{\mathfrak{u}}\cdot\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}')$ in the target of $f^\delta_{2,\mathfrak{u}}$ is relatively $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-graded.

We would like to give the mapping cone of $f^\delta_{2,\mathfrak{u}}$ a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading such that $f^\delta_{1,\mathfrak{u}}$ is grading-preserving. To be able to do this, we need to check that for any two Spin^c structures $\boldsymbol{t}, \boldsymbol{t}+u$ on $W_1 \cup W_2$ which restrict to $\mathfrak{u}$ on $Y_\Lambda(L)$, and to the same Spin^c structure on $Y_{\Lambda'}(L')$, the contributions to the composition $f^\delta_{2,\mathfrak{u}} \circ f^\delta_{1,\mathfrak{u}}$ coming from $\boldsymbol{t}$ and $\boldsymbol{t}+u$ shift the relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-gradings by the same amount. Using Lemma 11.3(b), this is equivalent to showing that

$$\langle c_1(\boldsymbol{t}) \smile u + u \smile u, [W_1 \cup W_2]\rangle \equiv 0 \pmod{\mathfrak{d}(\mathfrak{u})}. \tag{155}$$

Here $u \in H^2(W_1 \cup W_2, \partial(W_1 \cup W_2); \mathbb{Z})$ is Poincaré dual to some class in $H_2(W_1 \cup W_2; \mathbb{Z})$.

The cobordism $W_1 \cup W_2$ admits the following description in terms of surgery. Consider the link L^+ from Figure 35, and add an zero-framed, unknotted component $L_{\ell+3}$, which forms a Hopf link (in particular, has linking number one) with $L_{\ell+1}$, and is unlinked with the other components of L^+. Call the resulting framed link (L^{++}, Λ^{++}). Then, in the notation of Section 11.1,

$$W_1 \cup W_2 = W_{\Lambda^{++}}(L_1 \cup \dots \cup L_{\ell+1}, L^{++}).$$

In matrix form,

$$\Lambda^{++} = \left(\begin{array}{ccc|ccc} & & & 0 & 1 & 0 \\ & \Lambda & & 0 & 0 & 0 \\ & & & \vdots & \vdots & \vdots \\ & & & 0 & 0 & 0 \\ \hline 0\ 0 & \cdots & 0 & m_1 & 1 & 1 \\ 1\ 0 & \cdots & 0 & 1 & 0 & 0 \\ 0\ 0 & \cdots & 0 & 1 & 0 & 0 \end{array}\right)$$

Using Lemma 11.1(a), we see that

$$\mathrm{PD}(u) = [(v_1, \dots, v_\ell, j, 0, -m_1 j)] \in H_2(W_1 \cup W_2)$$

for some $j \in \mathbb{Z}$ and $\boldsymbol{v} = (v_1, \dots, v_\ell) \in H(L, \Lambda)^\perp$. We get

$$\langle u \smile u, [W_1 \cup W_2] \rangle = -m_1 j^2,$$

which is divisible by $\mathfrak{d}(\mathfrak{u})$ (because we assumed m_1 was so). Also, if $\mathfrak{u} = [\boldsymbol{s}]$ with $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \mathbb{H}(L)$, then we must have

$$\boldsymbol{t} = \left[\left(s_1 + \tfrac{1}{2}, s_2, \dots, s_\ell, m_1 + 1, s_{\ell+2}, s_{\ell+3}\right)\right] \in \mathbb{H}(L^{++})/H(L^{++}, \Lambda^{++}),$$

which gives

$$\langle c_1(\boldsymbol{t}) \smile u, [W_1 \cup W_2] \rangle = 2\boldsymbol{s} \cdot \boldsymbol{v} + 2m_1 j(1 - s_{\ell+3}).$$

This is divisible by $\mathfrak{d}(\mathfrak{u})$ because of (154) and by our choice of m_1.

Therefore, (155) holds, and $f^{\delta}_{1,\mathfrak{u}}$ can be made grading-preserving. The homotopy $H^{\mathfrak{d}}_{1,\mathfrak{u}}$ is also automatically grading-preserving. □

Corollary 11.27 *Suppose Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$, and fix $\delta > 0$. Then, for every $m_1 \gg 0$ suitably chosen and for every small Spin^c structure $\mathfrak{u}$ on $\mathrm{Spin}^c(Y_\Lambda(L))$, the mapping cone $\mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})$ admits a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading, and the quasi-isomorphism $(f^{\delta}_{1,\mathfrak{u}}, H^{\delta}_{1,\mathfrak{u}})$ is grading-preserving.*

Proof Since there are only finitely many small Spin^c structures $\mathfrak{u}$, we can choose m_1 to be a multiple of $\mathfrak{d}(\mathfrak{u})$ for all $\mathfrak{u}$ small. We then apply Proposition 11.24. □

Proposition 11.24 and Corollary 11.27 give a (partial) grading-preserving decomposition of the quasi-isomorphism $(f^{\delta}_1, H^{\delta}_1)$ from (133). Similar arguments can be applied to the quasi-isomorphism $(H^{\delta}_2, f^{\delta}_3)$ from (134). Indeed, one can check that an analogue of Lemma 11.22 holds for the cobordism W_3; ie the restriction map

$$\mathrm{Spin}^c(W_3) \to \mathrm{Spin}^c(\partial W_3) \cong \mathrm{Spin}^c(Y_{\Lambda'}(L')) \times \mathrm{Spin}^c(Y_\Lambda(L) \# L(m_1, 1))$$

is injective. Further, the maps f^{δ}_3 and H^{δ}_2 split into direct sums of maps $f^{\delta}_{3,\mathfrak{u}}$ and $H^{\delta}_{2,\mathfrak{u}}$, according to special Spin^c structures $\mathfrak{u}$ on $Y_\Lambda(L)$. A study of grading differences similar to that in the proof of Proposition 11.24 yields the following:

Proposition 11.28 *Fix $\delta > 0$ and suppose that Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$. Then, for every $m_1 \gg 0$ and for every special* Spin^c *structure $\mathfrak{u}$ on $\mathrm{Spin}^c(Y_\Lambda(L))$ such that m_1 is divisible by $\mathfrak{d}(\mathfrak{u})$, the quasi-isomorphism*

$$(H^\delta_{2,\mathfrak{u}}, f^\delta_{3,\mathfrak{u}})\colon \mathrm{Cone}(f^\delta_{2,\mathfrak{u}}) \xrightarrow{\sim} \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u})$$

respects the relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-gradings on the two sides.

To conclude this subsection we state a twisted coefficients generalization of Propositions 11.24 and 11.28. We work in the following setting. Recall that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w})$, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta}, \boldsymbol{w})$ are the Heegaard diagrams for $Y_\Lambda(L)$, $Y_{\overline{\Lambda}}(L)$, and $Y_{\Lambda'}(L')$, respectively. Also, for $i < g+k-1$, the curves β_i, γ_i and δ_i approximate one another. For $i = 1, \dots, g+k-1$, let us place points p_i and q_i on each side of the curve β_i so that they can be joined by an arc that intersects β_i, γ_i and δ_i once each, and does not intersect any of the alpha curves. (These are the analogues of w_1 and z_1.) Let $n_1, \dots, n_{g+k-2}$ be nonnegative integers. Consider the ring

$$\mathscr{S} = \mathbb{F}[S_1, \dots, S_{g+k-2}]/(S_1^{n_1} - 1, \dots, S_{g+k-1}^{n_{g+k-2}} - 1). \tag{156}$$

We can then construct Floer complexes with twisted coefficients

$$\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L); \mathscr{S}), \quad \boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L); \mathscr{S}), \quad \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'); \mathscr{S} \otimes \mathscr{T}_1)$$

and cobordism maps $f^\delta_{1;\mathscr{S}}$, $f^\delta_{2;\mathscr{S}}$ and $f^\delta_{3;\mathscr{S}}$ relating them, by counting all pseudoholomorphic strips and triangles using the coefficients

$$S_1^{n_{p_1}(\phi) - n_{q_1}(\phi)} \cdots S_{g+k-2}^{n_{p_{g+k-2}}(\phi) - n_{q_{g+k-2}}(\phi)},$$

where ϕ denotes the respective relative homology class. (This is all in addition to the power of T_1 coming from the twisting by $\mathscr{T}_1$.)

A straightforward generalization of Proposition 11.6 shows that $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L); \mathscr{S})$ is quasi-isomorphic to the mapping cone of $f^\delta_{2;\mathscr{S}}$. Note that twisting by $\mathscr{S}$ can have a nontrivial effect on the Floer homology of $Y_\Lambda(L)$ only when $b_1(Y_\Lambda(L)) > 0$. Supposing further that Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$ (as we did in this section), we have twisted coefficients analogues $f^\delta_{i,\mathfrak{u};\mathscr{S}}$ and $H^\delta_{i,\mathfrak{u};\mathscr{S}}$ of the maps $f^\delta_{i,\mathfrak{u}}$ and the homotopies $H^\delta_{i,\mathfrak{u}}$, respectively.

Proposition 11.29 *Fix $\delta > 0$ and suppose that Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$. Then, for every $m_1 \gg 0$ and for every special* Spin^c *structure $\mathfrak{u}$ on $\mathrm{Spin}^c(Y_\Lambda(L))$ such that m_1 is divisible by $\mathfrak{d}(\mathfrak{u}, \mathscr{S})$, the mapping cone* $\mathrm{Cone}(f^\delta_{2,\mathfrak{u};\mathscr{S}})$ *admits a relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u}, \mathscr{S})\mathbb{Z}$-grading, and the quasi-isomorphisms*

$$(f^\delta_{1,\mathfrak{u};\mathscr{S}}, H^\delta_{1,\mathfrak{u};\mathscr{S}})\colon \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u}; \mathscr{S}) \xrightarrow{\sim} \mathrm{Cone}(f^\delta_{2,\mathfrak{u};\mathscr{S}})$$

and

$$(H^\delta_{2,\mathfrak{u};\mathscr{S}}, f^\delta_{3,\mathfrak{u};\mathscr{S}})\colon \mathrm{Cone}(f^\delta_{2,\mathfrak{u};\mathscr{S}}) \xrightarrow{\sim} \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L), \mathfrak{u}; \mathscr{S})$$

respect the relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u}, \mathscr{S})\mathbb{Z}$-gradings on the two sides.

The proof of Proposition 11.29 is similar to those of Propositions 11.24 and 11.28.

11.7 Cobordism maps

Propositions 11.11, 11.15, 11.17 and 11.24 all describe quasi-isomorphisms between chain complexes of the form $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u})$ and mapping cones $\mathrm{Cone}(f_{2,\mathfrak{u}}^{\delta})$. Our goal in this section is to describe commutative diagrams which relate the inclusion of the target of $f_{2,\mathfrak{u}}$ into $\mathrm{Cone}(f_{2,\mathfrak{u}}^{\delta})$, with a cobordism map from the Floer complex (possibly with twisted coefficients) of $Y_{\Lambda'}(L')$ to the Floer complex of $Y_\Lambda(L)$. These diagrams will prove useful in our description of cobordism maps in terms of surgery, in Section 14.1.

We state several results, on a case-by-case basis, corresponding to the case analysis described above. We will always denote by $f_{3;\boldsymbol{t}}^{\delta}$ the map (with twisted coefficients) induced by the cobordism W_3 with a Spin^c structure $\boldsymbol{t}$. We start with Case I. In Case I(a), we have:

Proposition 11.30 *Fix $\delta>0$ and $i\in\mathbb{Z}$. Suppose Λ is nondegenerate, $a_1>0$, and m_1 is a sufficiently large multiple of a_1h. Suppose $\boldsymbol{s}_0\in\mathbb{H}(L)$ is such that $[\boldsymbol{s}_0]=\boldsymbol{t}_0\in\mathrm{Spin}^c(W_2)$ is the base Spin^c structure used to define the expression $T_1^{\boldsymbol{t}}$ in (137). If $\boldsymbol{s}\in\mathbb{H}(L)$ is such that $\boldsymbol{a}\cdot(\boldsymbol{s}-\boldsymbol{s}_0)=i$, let $\mathfrak{u}=[\boldsymbol{s}]\in\mathrm{Spin}^c(Y_\Lambda(L))\cong\mathbb{H}(L)/H(L,\Lambda)$, $\mathfrak{u}'=[\psi^{L_1}(\boldsymbol{s})]\in\mathrm{Spin}^c(Y_{\Lambda'}(L'))\cong\mathbb{H}(L')/H(L',\Lambda')$ and*

$$\boldsymbol{t}_i=[\boldsymbol{s}]\in\mathrm{Spin}^c(W_3)\cong\mathbb{H}(L)/\mathrm{Span}(\Lambda_2,\dots,\Lambda_\ell).$$

Then there is a diagram

$$\begin{tikzcd}
\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'),\mathfrak{u}') \arrow[r, "f^{\delta}_{3;\boldsymbol{t}_i}"] \arrow[d, "\cong"'] & \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u}) \arrow[d, "(f^{\delta}_{1,\mathfrak{u}},H^{\delta}_{1,\mathfrak{u}})"] \\
T_1^i\cdot\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'),\mathfrak{u}') \arrow[r] & \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})
\end{tikzcd}$$

commuting up to chain homotopy. Here, the bottom horizontal arrow is the inclusion into the mapping cone.

Proof The existence of such a diagram is a consequence of the proof of the fact that $(f_{1,\mathfrak{u}}^{\delta},H_{1,\mathfrak{u}}^{\delta})$ is a quasi-isomorphism; see the homological algebra lemma [Ozsváth and Szabó 2005, Lemma 4.4]; compare also [Ozsváth and Szabó 2008b, proof of Theorem 4.2]. A priori, the top horizontal arrow is a sum of all maps $f_{3;\boldsymbol{t}}^{\delta}$, over all $\boldsymbol{t}\in\mathrm{Spin}^c(W_3)$ such that $\boldsymbol{t}$ and $\boldsymbol{t}_i$ have the same restrictions to ∂W_3, and $\boldsymbol{a}\cdot(\boldsymbol{t}-[\boldsymbol{s}_0])\equiv i\ (\mathrm{mod}\ m_1)$. However, any such $\boldsymbol{t}$ differs from $\boldsymbol{t}_i$ by a multiple of $d\cdot\mathrm{PD}(\Sigma_3)$, where $\Sigma_3=\boldsymbol{s}$ is the generator of $H_2(W_3)$. Since i and δ are fixed, if m_1 (and hence $d=m_1/h$) is chosen sufficiently large, all $\boldsymbol{t}_i+dj\cdot\mathrm{PD}(\Sigma_3)$ produce trivial cobordism maps for $j\in\mathbb{Z}$ with $j\neq 0$. □

We have a similar result in Case I(b):

Proposition 11.31 *Fix $\delta>0$ and $i\in\mathbb{Z}$. Suppose Λ is nondegenerate, $a_1<0$ and m_1 is a sufficiently large multiple of a_1h. Suppose $\boldsymbol{s}_0,\boldsymbol{s}\in\mathbb{H}(L)$ (with $\boldsymbol{a}\cdot(\boldsymbol{s}-\boldsymbol{s}_0)=i$) and let $\mathfrak{u}$, $\mathfrak{u}'$ and $\boldsymbol{t}_i$ be as in*

Proposition 11.30. Then there is a commutative diagram

$$\begin{tikzcd} T_1^i \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r] \arrow[d, "\cong"'] & \operatorname{Cone}(f^{\delta}_{2,\mathfrak{u}}) \arrow[d, "(H^{\delta}_{2,\mathfrak{u}}, f^{\delta}_{3,\mathfrak{u}})"] \\ \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r, "f^{\delta}_{3;\boldsymbol{t}_i}"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}) \end{tikzcd}$$

where the top horizontal arrow is the inclusion into the mapping cone.

Proof This is simpler than Proposition 11.11 (and the diagram commutes on the nose, rather than only up to chain homotopy), because the relevant quasi-isomorphism (134) already involves $f_{3,\mathfrak{u}}$. The fact that the only contribution to $f_{3,\mathfrak{u}}$ comes from $\boldsymbol{t}_i$ (for δ, i fixed and m_1 large) is a consequence of the proof of Proposition 11.15. □

In Case I(c), we have:

Proposition 11.32 *Fix $\delta > 0$. Suppose Λ is nondegenerate, $a_1 = 0$ and m_1 is a sufficiently large multiple of h. Suppose*

$$\boldsymbol{t} \in \operatorname{Spin}^c(W_3) \cong \mathbb{H}(L)/\operatorname{Span}(\Lambda_2, \dots, \Lambda_\ell)$$

has restrictions $\mathfrak{u}'$ to $Y_{\Lambda'}(L')$ and $\mathfrak{u}$ to $Y_\Lambda(L)$. Then there is a diagram

$$\begin{tikzcd} \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{V}_1) \arrow[r, "f^{\delta}_{3;\boldsymbol{t}}"] \arrow[d, "\cong"'] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}) \arrow[d, "(f^{\delta}_{1,\mathfrak{u}}, H^{\delta}_{1,\mathfrak{u}})"] \\ T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{V}_1) \arrow[r] & \operatorname{Cone}(f^{\delta}_{2,\mathfrak{u}}) \end{tikzcd}$$

commuting up to chain homotopy, with the bottom horizontal arrow being inclusion into the mapping cone.

Proof The existence of the commutative diagram follows from the same reasoning as in Proposition 11.30, with the top arrow being the sum of cobordism maps over all possible Spin^c structures on W_3 with the given restrictions to the boundary. However, we claim that $\boldsymbol{t}$ is the unique such Spin^c structure. For this, it suffices to prove that the restriction

$$H^2(W_3) \to H^2(\partial W_3) \cong H^2(Y_\Lambda(L)) \oplus H^2(Y_{\Lambda'}(L')) \tag{157}$$

is injective. Suppose $[\boldsymbol{v}] \in H^2(W_3) = \mathbb{Z}^\ell/\operatorname{Span}(\Lambda_2, \dots, \Lambda_\ell)$ is in the kernel. Because $[\boldsymbol{v}]$ becomes trivial when projected to its last $\ell - 1$ components, without loss of generality we can assume $\boldsymbol{v} = (j, 0, \dots, 0)$ for some $j \in \mathbb{Z}$. We also know that $\boldsymbol{v} \in H(L, \Lambda)$, so j must be a multiple of h. Since $(h, 0, \dots, 0)$ is in the span of $\Lambda_2, \dots, \Lambda_\ell$ (because $a_1 = 0$), we must have $[\boldsymbol{v}] = 0$. □

Proposition 11.32 involves the map $f^{\delta}_{3;\boldsymbol{t}}$, which uses twisted coefficients. From here one can also similarly identify the usual (untwisted) map corresponding to the cobordism $W_3 = W_\Lambda(L', L)$. Indeed, the untwisted complex $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}')$ can be viewed as a subcomplex of $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{V}_1)$ via the inclusion

$$\iota\colon \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \hookrightarrow \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{V}_1), \quad \boldsymbol{x} \mapsto \boldsymbol{x}(1 + V_1 + \cdots + V_1^{d-1}). \tag{158}$$

The composition $f^{\delta}_{3;\boldsymbol{t}} \circ \iota$ is the untwisted chain map representing the cobordism W_3. By precomposing with ι in the diagram from Proposition 11.32, we get the following corollary:

Proposition 11.33 *Under the hypotheses of Proposition 11.32, we have a diagram*

$$\begin{tikzcd}
\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r, "f^{\delta}_{3;\boldsymbol{t}} \circ \iota"] \arrow[d, "\cong"'] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}) \arrow[d, "(f^{\delta}_{1,\mathfrak{u}}, H^{\delta}_{1,\mathfrak{u}})"] \\
T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r] & \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})
\end{tikzcd}$$

commuting up to chain homotopy, with the bottom horizontal map being the composition of ι and the inclusion into the mapping cone.

Finally, we have the following result in Case II:

Proposition 11.34 *Fix $\delta > 0$. Suppose Λ_1 is in the $\mathbb{Q}$-span of $\Lambda_2, \dots, \Lambda_\ell$ and m_1 is suitably chosen sufficiently large. Suppose*

$$\boldsymbol{t} \in \mathrm{Spin}^c(W_3) \cong \mathbb{H}(L)/\mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)$$

has restrictions $\mathfrak{u}'$ to $Y_{\Lambda'}(L')$ and $\mathfrak{u}$ to $Y_{\Lambda}(L)$ such that $\mathfrak{u}$ is special. Then:

(a) *There is a diagram*

$$\begin{tikzcd}
\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r, "f^{\delta}_{3;\boldsymbol{t}}"] \arrow[d, "\cong"'] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}) \arrow[d, "(f^{\delta}_{1,\mathfrak{u}}, H^{\delta}_{1,\mathfrak{u}})"] \\
T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r] & \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})
\end{tikzcd}$$

commuting up to chain homotopy, with the bottom horizontal arrow being inclusion into the mapping cone.

(b) *There is a diagram*

$$\begin{tikzcd}
T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r] \arrow[d, "\cong"'] & \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}}) \arrow[d, "(H^{\delta}_{2,\mathfrak{u}}, f^{\delta}_{3,\mathfrak{u}})"] \\
\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}') \arrow[r, "f^{\delta}_{3;\boldsymbol{t}}"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u})
\end{tikzcd}$$

commuting up to chain homotopy, with the top horizontal arrow being inclusion into the mapping cone.

Proof For both (a) and (b), just as in the proof of Proposition 11.32, it suffices to prove that the restriction map (157) is injective. As before, if $[\boldsymbol{v}]$ is in the kernel, we can assume that $\boldsymbol{v} = (j, 0, \dots, 0)$ for some $j \in \mathbb{Z}$. Lemma 11.19 then shows that $\boldsymbol{v} \in H(L, \Lambda)$ only if $j = 0$. $\square$

Let us also state the twisted coefficients generalization of Proposition 11.34, which has a similar proof:

Proposition 11.35 *Let δ, Λ, $\boldsymbol{t}$, $\mathfrak{u}$ and $\mathfrak{u}'$ be as in Proposition 11.34, and $\mathscr{S}$ a twisted coefficients ring as in* (156). *Then*:

(a) *There is a diagram*

$$\begin{tikzcd} \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{S}) \arrow[r, "f^{\delta}_{3;\boldsymbol{t},\mathscr{S}}"] \arrow[d, "\cong"'] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}; \mathscr{S}) \arrow[d, "(f^{\delta}_{1,\mathfrak{u};\mathscr{S}}, H^{\delta}_{1,\mathfrak{u};\mathscr{S}})"] \\ T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{S}) \arrow[r] & \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u};\mathscr{S}}) \end{tikzcd}$$

commuting up to chain homotopy, with the bottom horizontal arrow being inclusion into the mapping cone.

(b) *There is a diagram*

$$\begin{tikzcd} T_1^{\mathfrak{u}} \cdot \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{S}) \arrow[r] \arrow[d, "\cong"'] & \mathrm{Cone}(f^{\delta}_{2,\mathfrak{u};\mathscr{S}}) \arrow[d, "(H^{\delta}_{2,\mathfrak{u};\mathscr{S}}, f^{\delta}_{3,\mathfrak{u};\mathscr{S}})"] \\ \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'), \mathfrak{u}'; \mathscr{S}) \arrow[r, "f^{\delta}_{3;\boldsymbol{t};\mathscr{S}}"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda}(L), \mathfrak{u}; \mathscr{S}) \end{tikzcd}$$

commuting up to chain homotopy, with the top horizontal arrow being inclusion into the mapping cone.

12 Proof of the surgery theorem for link-minimal complete systems

The goal of this section is to prove Theorem 9.6. We first do so for the case of the basic systems from Definition 8.32. The proof in that case will be modeled on the proof of the formula for the Heegaard Floer homology of integral surgeries on knots; see [Ozsváth and Szabó 2008b]. We will need to combine the arguments in [loc. cit.] with the homological algebra from [Ozsváth and Szabó 2005]. At the end we will explain how the statement of Theorem 9.6 for basic systems implies the statement for all link-minimal complete systems of hyperboxes.

12.1 Large surgeries on links

Let $\vec{L} \subset Y$ be a link in an integral homology three-sphere as in Section 9. We let $\widetilde{\Lambda} \gg 0$ be a sufficiently large framing on L, meaning that the framing coefficients $\tilde{\lambda}_i$ on each component are sufficiently large, as for the framing denoted by $\widetilde{\Lambda}$ in Sections 10.3. We let $\widetilde{\Lambda}_i \in H_1(Y - L)$ be the induced framings on each component L_i, as usual. Recall that $H(L, \widetilde{\Lambda})$ denotes the lattice in $H_1(Y - L) \cong \mathbb{Z}^{\ell}$ generated by all $\widetilde{\Lambda}_i$.

We use the notation from Section 11.1. In particular, surgery on the framed link L produces a cobordism $W_{\widetilde{\Lambda}}(L)$ between Y and the surgered manifold $Y_{\widetilde{\Lambda}}(L)$. Since $\widetilde{\Lambda}$ is chosen sufficiently large, the manifold $Y_{\widetilde{\Lambda}}(L)$ is a rational homology three-sphere. Let $W'_{\widetilde{\Lambda}}(L)$ be the cobordism between $Y_{\widetilde{\Lambda}}(L)$ and Y, obtained by turning around the cobordism $-W_{\widetilde{\Lambda}}(L)$.

As in the proof of Lemma 11.1, choose a Seifert surface $F_i \subset Y$ for each link component L_i, and let $\widehat{F}_i$ be the surface obtained by capping off F_i in $W_{\widetilde{\Lambda}}(L)$. By a slight abuse of notation, we also denote by $\widehat{F}_i$ the corresponding surface in $W'_{\widetilde{\Lambda}}(L)$. The homology classes $[\widehat{F}_i]$ for $i = 1, \dots, \ell$ form a basis of $H_2(W'_{\widetilde{\Lambda}}(L))$. As in (122), we identify $H^2(W'_{\widetilde{\Lambda}}(L))$ with $\mathbb{Z}^\ell$ by sending a cohomology class c to $(\langle c, [\widehat{F}_1]\rangle, \dots, \langle c, [\widehat{F}_\ell]\rangle)$.

Given a Spin^c structure $\mathfrak{u}$ over $Y_{\widetilde{\Lambda}}(L)$, we can extend it to a Spin^c structure $\mathfrak{t}$ over $W'_{\widetilde{\Lambda}}(L)$. We can then find $\boldsymbol{s} \in \mathbb{H}(L)$ such that

$$c_1(\mathfrak{t}) \equiv 2\boldsymbol{s} - (\widetilde{\Lambda}_1 + \cdots + \widetilde{\Lambda}_\ell) \pmod{2H(L, \widetilde{\Lambda})}.$$

The correspondence $\mathfrak{u} \mapsto \boldsymbol{s}$ determined by the above formula induces an isomorphism

$$\mathrm{Spin}^c(Y_{\widetilde{\Lambda}}(L)) \to \mathbb{H}(L)/H(L, \widetilde{\Lambda}).$$

Let $P(\widetilde{\Lambda})$ be the intersection of the lattice $\mathbb{H}(L)$ with the hyperparallelepiped with vertices

$$\zeta + \tfrac{1}{2}(\pm\widetilde{\Lambda}_1 \pm \widetilde{\Lambda}_2 \pm \cdots \pm \widetilde{\Lambda}_\ell),$$

as in Section 10.3. This is a fundamental domain for $\mathbb{H}(L)/H(L, \widetilde{\Lambda})$; see (111). Hence, there is a bijection

$$\mathrm{Spin}^c(Y_{\widetilde{\Lambda}}(L)) \cong P(\widetilde{\Lambda});$$

see Section 9.2. From now on we will denote a Spin^c structure on $Y_{\widetilde{\Lambda}}(L)$ by the corresponding value $\boldsymbol{s} \in P(\widetilde{\Lambda}) \subset \mathbb{H}(L)$. For $\boldsymbol{s} \in P(\widetilde{\Lambda})$, we denote by $\mathfrak{x}_{\boldsymbol{s}}$ the Spin^c structure over $W'_{\widetilde{\Lambda}}(L)$ satisfying

$$c_1(\mathfrak{x}_{\boldsymbol{s}}) = 2\boldsymbol{s} - (\widetilde{\Lambda}_1 + \cdots + \widetilde{\Lambda}_\ell); \tag{159}$$

compare (124).

Let us choose a basic system $\mathscr{H}$ for $\vec{L} \subset Y$, as in Section 8.7. Recall that the initial Heegaard diagram $\mathscr{H}^L = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}, \boldsymbol{z})$ in the system contains ℓ beta curves $\beta_1, \dots, \beta_\ell$ such that the basepoints w_i and z_i lie one on each side of β_i. Let $\mathcal{R} = \mathbb{F}[\![U_1, \dots, U_\ell]\!]$.

Theorem 12.1 *For $\widetilde{\Lambda}$ sufficiently large, there exist quasi-isomorphisms of relatively $\mathbb{Z}$-graded complexes of $\mathcal{R}$-modules*

$$\Psi^-_{\widetilde{\Lambda},\boldsymbol{s}} \colon \boldsymbol{CF}^-(Y_{\widetilde{\Lambda}}(L), \boldsymbol{s}) \to \mathfrak{A}^-(\mathscr{H}^L, \boldsymbol{s})$$

for all $\boldsymbol{s} \in P(\widetilde{\Lambda})$.

Proof For each $i = 1, \dots, \ell$, we construct a curve δ_i by twisting the longitude of L_i $\tilde{\lambda}_i$ times along β_i, in a symmetric way as in [Ozsváth and Szabó 2004a, Figure 5] (but with the twisting done in the opposite direction). Thus, δ_i specifies the framing $\tilde{\Lambda}_i$ of the component L_i for $i = 1, \dots, \ell$. We complete this to a full set of attaching circles δ by taking curves $\delta_{\ell+1}, \dots, \delta_{\ell+g-1}$ that approximate (ie are small Hamiltonian translates of) $\beta_{\ell+1}, \dots, \beta_{\ell+g-1}$. The result is a triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta}, \boldsymbol{\beta}, \boldsymbol{w})$ for the cobordism $W'_{\tilde{\Lambda}}(L)$ such that we have the three-manifolds $Y_{\alpha,\beta} \cong Y$, $Y_{\alpha,\delta} \cong Y_{\tilde{\Lambda}}(L)$ and $Y_{\delta,\beta} \cong \#^{g-1}(S^1 \times S^2)$.

For $\boldsymbol{s} \in \mathbb{H}(L)$, we can define a chain map $\Psi^-_{\tilde{\Lambda},\boldsymbol{s}} \colon \boldsymbol{CF}^-(Y_{\tilde{\Lambda}}(L), \boldsymbol{s}) \to \mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ by

$$\Psi^-_{\tilde{\Lambda},\boldsymbol{s}}(\boldsymbol{x}) = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \Theta, \boldsymbol{y}) \\ \mu(\phi)=0 \\ n_{w_i}(\phi) - n_{z_i}(\phi) = A_i(\boldsymbol{y}) - s_i \text{ for all } i}} \#\mathcal{M}(\phi) \cdot \prod_{i=1}^{\ell} U_i^{\min(n_{z_i}(\phi), n_{w_i}(\phi))} \cdot \boldsymbol{y}, \tag{160}$$

where $\Theta \in \boldsymbol{CF}^-(Y_{\delta,\beta})$ is the top-degree generator in homology.

The proof that $\Psi_{\tilde{\Lambda},\boldsymbol{s}}$ is a quasi-isomorphism for $\tilde{\Lambda} \gg 0$ and $\boldsymbol{s} \in P(\tilde{\Lambda})$ then proceeds along the same lines as [ibid., proof of Theorem 4.4]; see also [Ozsváth and Szabó 2008b, Theorem 2.3; 2011, Theorem 4.1]. Roughly, the argument is as follows. There are ℓ winding regions on the surface Σ, that is, neighborhoods of the curves β_i for $i = 1, \dots, \ell$ in which the twisting of the corresponding curves δ_i takes place. A generator $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\delta$ is said to be *supported in the winding regions* if it contains points in all the ℓ winding regions. If this is the case, the i^{th} component of the Spin^c structure $\boldsymbol{s} \in P(\tilde{\Lambda})$ of $\boldsymbol{x}$ equals the depth of the respective point of $\boldsymbol{x}$ inside the i^{th} winding region, up to the addition of a constant; see [Ozsváth and Szabó 2004a, equation (14)]. A Spin^c structure $\boldsymbol{s}$ is said to be supported in the winding regions if all the generators $\boldsymbol{x} \in \boldsymbol{s}$ are supported in the winding regions. For such $\boldsymbol{s}$, the map $\Psi_{\tilde{\Lambda},\boldsymbol{s}}$ is actually an isomorphism of chain complexes, because it is approximated (with respect to area filtrations) by a "nearest-point" map which is a bijection.

Note that we have some freedom in choosing the winding region. Indeed, by replacing each δ_i for $i = 1, \dots, \ell$ with an isotopic curve δ'_i so that the number of twists to the left of β_i is changed (ie translating the twists to the right or left of the curve β_i), we obtain another strongly equivalent triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\delta}', \boldsymbol{w}, \boldsymbol{z})$. If δ'_i differs from δ_i by k_i twists, the set of Spin^c structures supported in the new winding regions is a translate of the old set by $(k_1, \dots, k_\ell)$. We have

$$(\tilde{\lambda}_1 - C_1) \cdot (\tilde{\lambda}_2 - C_2) \cdots (\tilde{\lambda}_\ell - C_\ell)$$

possibilities for the position of the attaching set $\boldsymbol{\delta}$, where $C_1, \dots, C_\ell$ are constants (independent of the framing coefficients). For each of these $\boldsymbol{\delta}$, the number of Spin^c structures not supported in the respective winding regions is of the order of

$$\tilde{\lambda}_1 \cdots \tilde{\lambda}_\ell \cdot \sum_{i=1}^{\ell} \frac{C'_i}{\tilde{\lambda}_i}$$

for some constants C'_i.

Let us choose 2^ℓ different sets of curves $\boldsymbol{\delta}^\varepsilon$ for $\varepsilon \in \{0,1\}^\ell$ such that $\boldsymbol{\delta}_i^\varepsilon$ and $\boldsymbol{\delta}_i^{\varepsilon'}$ differ by $\frac{1}{2}\tilde{\lambda}_i$ twists whenever $\varepsilon_i \neq \varepsilon_i'$ (and are the same curve if $\varepsilon_i = \varepsilon_i'$). If $\tilde{\lambda}_1, \dots, \tilde{\lambda}_i$ are sufficiently large, each Spin^c structure $\boldsymbol{s}$ on $Y_{\tilde{\Lambda}}(L)$ is supported in one of the winding regions for some $\boldsymbol{\delta}^\varepsilon$. The fact that the new map $\Psi^{\varepsilon,-}_{\tilde{\Lambda},\boldsymbol{s}}$ (using $\boldsymbol{\delta}^\varepsilon$ instead of $\boldsymbol{\delta}$) is an isomorphism implies that the original map $\Psi^-_{\tilde{\Lambda},\boldsymbol{s}}$ (which differs from $\Psi^{\varepsilon,-}_{\tilde{\Lambda},\boldsymbol{s}}$ by composition with chain homotopy equivalences) is a quasi-isomorphism. □

Let $\vec{L}^{\mathfrak{o}}$ be L with some arbitrary orientation $\mathfrak{o}$, and let $\mathfrak{x}^{\mathfrak{o}}_{\boldsymbol{s}}$ be the Spin^c structure on $W'_{\tilde{\Lambda}}(L)$ satisfying

$$c_1(\mathfrak{x}^{\mathfrak{o}}_{\boldsymbol{s}}) = c_1(\mathfrak{x}_{\boldsymbol{s}}) + 2\tilde{\Lambda}_{\vec{L},\vec{L}^{\mathfrak{o}}}, \tag{161}$$

where $\tilde{\Lambda}_{\vec{L},\vec{L}^{\mathfrak{o}}} = \sum_{i \in I_-(\vec{L},\vec{L}^{\mathfrak{o}})} \tilde{\Lambda}_i$ and $I_-(\vec{L},\vec{L}^{\mathfrak{o}})$ is the set of indices describing components of $\vec{L}$ oriented differently in $\vec{L}^{\mathfrak{o}}$; see Section 3.8.

We denote by $F^-_{W,\mathfrak{x}}$ the map on Heegaard Floer complexes induced by a particular cobordism W and Spin^c structure $\mathfrak{x}$; see [Ozsváth and Szabó 2006].

Theorem 12.2 *Fix $\vec{L}^{\mathfrak{o}} \subset Y$ as above. For any $\tilde{\Lambda} \gg 0$ and $\boldsymbol{s} \in P(\tilde{\Lambda})$, there is a commutative diagram*

$$\begin{tikzcd} \boldsymbol{CF}^-(Y_{\tilde{\Lambda}}(L), \boldsymbol{s}) \arrow[r, "F^-_{W'_{\tilde{\Lambda}}(L),\mathfrak{x}^{\mathfrak{o}}_{\boldsymbol{s}}}"] \arrow[d, "\Psi^-_{\tilde{\Lambda},\boldsymbol{s}}"'] & \boldsymbol{CF}^-(Y) \arrow[d, "\cong"] \\ \mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \arrow[r, "\mathcal{I}^{\vec{L}^{\mathfrak{o}}}_{\boldsymbol{s}}"] & \mathfrak{A}^-(\mathcal{H}^L, p^{\vec{L}^{\mathfrak{o}}}(\boldsymbol{s})) \end{tikzcd} \tag{162}$$

for all $\boldsymbol{s} \in P(\tilde{\Lambda})$.

The proof of Theorem 12.2 is similar to that of [Ozsváth and Szabó 2008b, Theorem 2.3]. In the diagram (162), we implicitly identified $\mathfrak{A}^-(\mathcal{H}^L, p^{\vec{L}^{\mathfrak{o}}}(\boldsymbol{s}))$ with $\mathfrak{A}^-(r_{\vec{L}^{\mathfrak{o}}}(\mathcal{H}^L), \psi^{\vec{L}^{\mathfrak{o}}}(\boldsymbol{s}))$; see (22). Also, it is worth mentioning that in the proof of Theorem 12.2, the map $F^-_{W'_{\tilde{\Lambda}}}(L)$ is defined using the triple Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\delta}, \boldsymbol{\beta}, \boldsymbol{w}^{\mathfrak{o}})$, where

$$\boldsymbol{w}^{\mathfrak{o}} = \{w_i \mid i \in I_+(\vec{L},\vec{L}^{\mathfrak{o}})\} \cup \{z_i \mid i \in I_-(\vec{L},\vec{L}^{\mathfrak{o}})\}.$$

(This is the set of basepoints that would be denoted by $\boldsymbol{w}^{\vec{L},\vec{L}^{\mathfrak{o}}}$ in the notation of Section 8.7.)

There is a more refined version of Theorem 12.2, as follows. Note that the cobordism $W'_{\tilde{\Lambda}}(L)$ consists of ℓ two-handle additions, which can be composed in any order. Different ways of composing are related by chain homotopies, forming a hypercube of chain complexes. In the refined version that we state below (Theorem 12.3) we replace the top arrow in (162) with this hypercube; this is isomorphic to another hypercube, replacing the bottom arrow in (162).

Let $\boldsymbol{\delta} = (\delta_1, \dots, \delta_{g+\ell-1})$ be a set of attaching curves as in the proof of Theorem 12.1. Given $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell) \in \mathbb{E}_\ell = \{0,1\}^\ell$, we define a new $(g+\ell-1)$-tuple of attaching circles $\boldsymbol{\eta}^\varepsilon$ by

$$\eta_i^\varepsilon \approx \begin{cases} \delta_i & \text{if } \varepsilon_i = 0, \\ \beta_i & \text{if } \varepsilon_i = 1. \end{cases}$$

For each $\varepsilon \in \mathbb{E}_\ell$, we denote by $\vec{L}^{\mathfrak{o},\varepsilon} \subseteq \vec{L}^{\mathfrak{o}}$ the oriented sublink consisting of those components L_i such that $\varepsilon_i = 1$, all taken with the orientation induced from $\mathfrak{o}$. We let L^ε be the underlying sublink. Note that the Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\eta}^\varepsilon, \boldsymbol{w}^{\mathfrak{o}})$ represents the three-manifold $Y_{\widetilde{\Lambda}|_{L-L^\varepsilon}}(L - L^\varepsilon)$.

Further, for every $\varepsilon < \varepsilon'$, the Heegaard diagram $(\Sigma, \boldsymbol{\eta}^\varepsilon, \boldsymbol{\eta}^{\varepsilon'}, \boldsymbol{w}^{\mathfrak{o}})$ represents a connected sum of some copies of $S^1 \times S^2$. We can arrange so that the Floer homology $HF(\mathbb{T}_{\eta^\varepsilon}, \mathbb{T}_{\eta^{\varepsilon'}})$, in the maximal degree with nontrivial homology (and in the torsion Spin^c structure), is represented by a unique intersection point, which we denote by $\Theta^{\mathrm{can}}_{\varepsilon,\varepsilon'}$. Set

$$\Theta_{\varepsilon,\varepsilon'} = \begin{cases} \Theta^{\mathrm{can}}_{\varepsilon,\varepsilon'} & \text{if } \|\varepsilon' - \varepsilon\| = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{163}$$

For $\varepsilon \leq \varepsilon'$, let

$$W'_{\widetilde{\Lambda}}(L - L^\varepsilon, L - L^{\varepsilon'}) \subseteq W'_{\widetilde{\Lambda}}(L)$$

be the cobordism from $Y_{\widetilde{\Lambda}|_{L-L^\varepsilon}}(L - L^\varepsilon)$ to $Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon'}}}(L - L^{\varepsilon'})$ obtained by reversing the surgery on $L^{\varepsilon'-\varepsilon}$. When $\mathfrak{x}$ is a Spin^c structure on $W'_{\widetilde{\Lambda}}(L)$, we keep the same notation $\mathfrak{x}$ for its restriction to $W'_{\widetilde{\Lambda}}(L - L^\varepsilon, L - L^{\varepsilon'})$.

Consider the polygon map

$$F(\varepsilon, \varepsilon', \boldsymbol{w}^{\mathfrak{o}})\colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}^{\mathfrak{o}}) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^{\varepsilon'}}, \boldsymbol{w}^{\mathfrak{o}}),$$
$$F(\varepsilon, \varepsilon', \boldsymbol{w}^{\mathfrak{o}})(\boldsymbol{x}) = \sum_{\varepsilon = \varepsilon^0 < \cdots < \varepsilon^p = \varepsilon'} f(\boldsymbol{x} \otimes \Theta_{\varepsilon^0,\varepsilon^1} \otimes \cdots \otimes \Theta_{\varepsilon^{p-1},\varepsilon^p}),$$

in the notation of Section 3.5, used here for polygon maps between ordinary Floer chain complexes, as in [Ozsváth and Szabó 2005, Section 4.2]. When $\varepsilon = \varepsilon'$, this is simply the differential ∂. When $\|\varepsilon' - \varepsilon\| = 1$, the map $F(\varepsilon, \varepsilon', \boldsymbol{w}^{\mathfrak{o}})$ is a triangle map representing the cobordism $W'_{\widetilde{\Lambda}}(L - L^\varepsilon, L - L^{\varepsilon'})$ and, as such, it decomposes as a sum of maps according to the Spin^c structures on that cobordism. For general $\varepsilon < \varepsilon'$, the map $F(\varepsilon, \varepsilon', \boldsymbol{w}^{\mathfrak{o}})$ is a higher order chain homotopy relating the different ways of splitting $W'_{\widetilde{\Lambda}}(L - L^\varepsilon, L - L^{\varepsilon'})$ into two-handle additions. It still decomposes as a sum of maps

$$F(\varepsilon, \varepsilon', \boldsymbol{w}^{\mathfrak{o}}, \mathfrak{x})$$

according to the Spin^c structures $\mathfrak{x}$ on the cobordism $W'_{\widetilde{\Lambda}}(L - L^\varepsilon, L - L^{\varepsilon'})$.

Theorem 12.3 *Fix $\widetilde{\Lambda} \gg 0$, $\boldsymbol{s} \in P(\widetilde{\Lambda})$ and an orientation $\mathfrak{o}$ on L as above. Then the hypercube with chain groups*

$$C^\varepsilon = \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}^{\mathfrak{o}}, \psi^{\vec{L}^{\mathfrak{o},\varepsilon}}(\boldsymbol{s})) = \boldsymbol{CF}^-(Y_{\widetilde{\Lambda}|_{L-L^\varepsilon}}(L - L^\varepsilon), \psi^{\vec{L}^{\mathfrak{o},\varepsilon}}(\boldsymbol{s}))$$

and maps

$$D^{\varepsilon'-\varepsilon}_\varepsilon = F(\varepsilon, \varepsilon', \boldsymbol{w}^{\mathfrak{o}}, \mathfrak{x}^{\mathfrak{o}}_{\boldsymbol{s}})$$

is quasi-isomorphic to the hypercube with chain groups

$$C^\varepsilon = \mathfrak{A}^-(\mathcal{H}^L, p^{\vec{L}^{\mathfrak{o},\varepsilon}}(\boldsymbol{s})) = \mathfrak{A}^-(r_{\vec{L}^{\mathfrak{o},\varepsilon}}(\mathcal{H}^L), \psi^{\vec{L}^{\mathfrak{o},\varepsilon}}(\boldsymbol{s}))$$

and maps

$$D_{\varepsilon}^{\varepsilon'-\varepsilon}=\begin{cases}\partial & \text{if } \varepsilon=\varepsilon',\\ \mathcal{I}^{\vec{L}_i^{\circ}}_{p^{\vec{L}^{\circ,\varepsilon}}(s)} & \text{if } \|\varepsilon'-\varepsilon\|=1,\ \vec{L}^{\circ,\varepsilon'}=\vec{L}^{\circ,\varepsilon}\amalg\vec{L}_i^{\circ},\\ 0 & \text{otherwise.}\end{cases}$$

Proof The maps

$$\Psi^{-}_{\widetilde{\Lambda}|_{L-L^{\varepsilon}},\psi^{\vec{L}^{\circ,\varepsilon}}(s)}\colon \boldsymbol{CF}^{-}(Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon}}}(L-L^{\varepsilon}),\psi^{\vec{L}^{\circ,\varepsilon}}(s))\to\mathfrak{A}^{-}(r_{\vec{L}^{\circ,\varepsilon}}(\mathcal{H}^{L}),\psi^{\vec{L}^{\circ,\varepsilon}}(s))$$

given by (160) can be generalized to give maps (increasing ε) which count higher polygons, and where the U powers are counted just as in (160). These form a chain map between the two hypercubes.

By definition, a quasi-isomorphism of hypercubes means that the corresponding ε-preserving maps are quasi-isomorphisms for all ε. Indeed, the maps $\Psi^{-}_{\widetilde{\Lambda}|_{L-L^{\varepsilon}},\psi^{\vec{L}^{\circ,\varepsilon}}(s)}$ are quasi-isomorphisms by a simple extension of Theorem 12.1. □

We can change the two hypercubes in Theorem 12.3 by chain homotopy equivalences, and arrive at the following:

Proposition 12.4 *Fix $\widetilde{\Lambda}\gg 0$ and $s\in P(\widetilde{\Lambda})$ and an orientation $\mathfrak{o}$ as above. Then the hypercube $\mathscr{H}_s^{\mathfrak{o}}$ with chain groups*

$$C^{\varepsilon}=\boldsymbol{CF}^{-}(\mathbb{T}_{\alpha},\mathbb{T}_{\eta^{\varepsilon}},\boldsymbol{w},\psi^{\vec{L}^{\circ,\varepsilon}}(s))=\boldsymbol{CF}^{-}(Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon}}}(L-L^{\varepsilon}),\psi^{\vec{L}^{\circ,\varepsilon}}(s))$$

and maps

$$D_{\varepsilon}^{\varepsilon'-\varepsilon}=F(\varepsilon,\varepsilon',\boldsymbol{w},\mathfrak{x}_s^{\mathfrak{o}})$$

is quasi-isomorphic to the hypercube $\mathcal{H}_s^{\mathfrak{o}}$ having chain groups

$$C^{\varepsilon}=\mathfrak{A}^{-}(\mathcal{H}^{L-L^{\varepsilon}},\psi^{\vec{L}^{\circ,\varepsilon}}(s))$$

and maps

$$D_{\varepsilon}^{\varepsilon'-\varepsilon}=\Phi^{\vec{L}^{\circ,\varepsilon'-\varepsilon}}_{\psi^{\vec{L}^{\circ,\varepsilon}}(s)}=D^{\vec{L}^{\circ,\varepsilon'-\varepsilon}}_{(p^{\vec{L}^{\circ,\varepsilon'-\varepsilon}}\circ\psi^{\vec{L}^{\circ,\varepsilon}})(s)}\circ\mathcal{I}^{\vec{L}^{\circ,\varepsilon'-\varepsilon}}_{\psi^{\vec{L}^{\circ,\varepsilon}}(s)},$$

in the notation of Section 9.1.

Proof Let us compare the first hypercube in Theorem 12.3 with the first hypercube in Proposition 12.4. Note that the Heegaard diagrams $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\eta}^{\varepsilon},\boldsymbol{w}^{\mathfrak{o}})$ and $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\eta}^{\varepsilon},\boldsymbol{w})$ both represent the same three-manifold $Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon}}}(L-L^{\varepsilon})$. Hence, the respective chain complexes are chain homotopy equivalent. In fact, we can describe the chain homotopy equivalence along the lines of Section 8.7. First, note that $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\eta}^{\varepsilon},\boldsymbol{w}^{\mathfrak{o}})$ is isotopic to $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\eta}''^{\varepsilon},\boldsymbol{w})$, where the collection $\boldsymbol{\eta}''^{\varepsilon}$ differs from $\boldsymbol{\eta}^{\varepsilon}$ by replacing every curve that approximates β_i with one that approximates β_i''. Here β_i'' is as in Section 8.7, and similarly we recall that we also have an intermediate curve β_i'. One can relate $\boldsymbol{CF}^{-}(\mathbb{T}_{\alpha},\mathbb{T}_{\eta''^{\varepsilon}},\boldsymbol{w})$ to $\boldsymbol{CF}^{-}(\mathbb{T}_{\alpha},\mathbb{T}_{\eta'^{\varepsilon}},\boldsymbol{w})$ and then to $\boldsymbol{CF}^{-}(\mathbb{T}_{\alpha},\mathbb{T}_{\eta^{\varepsilon}},\boldsymbol{w})$ via chain homotopy equivalences given by triangle maps with one vertex in the respective canonical generator. It is straightforward to lift these to chain homotopy equivalences between the respective hypercubes. Moreover, we can restrict everything to a Spin^c structure $\psi^{\vec{L}^{\circ,\varepsilon}}(s)$, which is left unchanged throughout.

Now let us compare the second hypercube in Theorem 12.3 with the second hypercube in Proposition 12.4. The former is reminiscent of a canonical hypercube as in Definition 5.22, but it has inclusions rather than identity maps along its edges. Nevertheless, we can construct a chain map very similar to the canonical inclusion from Section 5.9, as follows.

For $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell) \in \mathbb{E}_\ell$ and $i \in \{0, 1, \dots, \ell\}$, we let $\varepsilon[\le i]$ (resp. $\varepsilon[> i]$) be the multi-indices obtained from ε by changing all entries indexed by $j > i$ (resp. $j \le i$) into zeros; compare Section 5.9. We define an intermediate hypercube $H[i]$ to have chain groups

$$C[i]^\varepsilon = \mathfrak{A}^-(\mathcal{H}^{L-L^{\varepsilon[\le i]}}, (\psi^{\vec{L}^{\circ,\varepsilon[\le i]}} \circ p^{\vec{L}^{\circ,\varepsilon}})(s))$$

and maps

$$D[i]_\varepsilon^{\varepsilon'-\varepsilon} \colon \mathfrak{A}^-(\mathcal{H}^{L-L^{\varepsilon[\le i]}}, (\psi^{\vec{L}^{\circ,\varepsilon[\le i]}} \circ p^{\vec{L}^{\circ,\varepsilon}})(s)) \to \mathfrak{A}^-(\mathcal{H}^{L-L^{\varepsilon'[\le i]}}, (\psi^{\vec{L}^{\circ,\varepsilon'[\le i]}} \circ p^{\vec{L}^{\circ,\varepsilon'}})(s))$$

given by

$$D[i]_\varepsilon^{\varepsilon'-\varepsilon} = \begin{cases} D^{\vec{L}^{\circ,(\varepsilon'-\varepsilon)}} \circ \mathcal{I}^{\vec{L}^{\circ,(\varepsilon'-\varepsilon)}} & \text{if } \varepsilon[> i] = \varepsilon'[> i], \\ \mathcal{I}^{\vec{L}^{\circ,(\varepsilon'-\varepsilon)}} & \text{if } \varepsilon[\le i] = \varepsilon'[\le i] \text{ and } \|\varepsilon'[> i] - \varepsilon[> i]\| = 1, \\ 0 & \text{otherwise.} \end{cases}$$

We omitted here the subscripts in the maps D and $\mathcal{I}$, as they are uniquely determined by the domains of those maps.

Note that $H[0]$ is the second hypercube in Theorem 12.3, while $H[\ell]$ is the second hypercube in Proposition 12.4.

For $i = 1, \dots, \ell$, we define chain maps

$$F[i] \colon H[i-1] \to H[i]$$

to consist of

$$F[i]_\varepsilon^{\varepsilon'-\varepsilon} \colon \mathfrak{A}^-(\mathcal{H}^{L-L^{\varepsilon[\le(i-1)]}}, (\psi^{\vec{L}^{\circ,\varepsilon[\le(i-1)]}} \circ p^{\vec{L}^{\circ,\varepsilon}})(s)) \to \mathfrak{A}^-(\mathcal{H}^{L-L^{\varepsilon'[\le i]}}, (\psi^{\vec{L}^{\circ,\varepsilon'[\le i]}} \circ p^{\vec{L}^{\circ,\varepsilon'}})(s)),$$

$$F[i]_\varepsilon^{\varepsilon'-\varepsilon} = \begin{cases} D^{\vec{L}^{\circ,\varepsilon'[\le i]-\varepsilon[\le(i-1)]}} \circ \mathcal{I}^{\vec{L}^{\circ,(\varepsilon'-\varepsilon)[\le(i-1)]}} & \text{if } \varepsilon_i = 1,\ \varepsilon[> i] = \varepsilon'[> i], \\ \mathrm{Id} & \text{if } \varepsilon = \varepsilon' \text{ and } \varepsilon_i = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Note that when $\varepsilon = \varepsilon'$ the map $F[i]_\varepsilon^{\varepsilon'-\varepsilon}$ is either the identity (when $\varepsilon_i = 0$) or an edge map of the form $D^{\vec{L}_i^\circ}$ (when $\varepsilon_i = 1$); in either case, it is a chain homotopy equivalence. One can lift the respective chain homotopies to the level of the hypercubes. This shows that each $F[i]$ is a chain homotopy equivalence of hypercubes. The composition $F[\ell] \circ \dots \circ F[1]$ then represents a chain homotopy equivalence between $H[0]$ and $H[\ell]$.

The claim now follows from Theorem 12.3. □

12.2 Iterating the exact triangle

Our goal in this subsection is to present a generalization of the exact triangle (118) in the form of a description of $\boldsymbol{HF}^-(Y_\Lambda(L))$ for arbitrary surgery on a link L in an integral homology sphere Y. This will be based on iterating the more general exact triangle from Proposition 11.5.

Let $\widetilde{\Lambda} \gg 0$ be a new framing for L, as in Section 12.1. We denote by $\widetilde{\lambda}_i$ and λ_i the framing coefficients on the component L_i, coming from $\widetilde{\Lambda}$ and Λ, respectively. For each $i = 1, \dots, \ell$, set

$$m_i = \widetilde{\lambda}_i - \lambda_i \gg 0.$$

We keep all the notation from Section 12.1. In particular, we have chosen a basic system $\mathscr{H}$ for $\vec{L} \subset Y$, and we have collections of curves $\boldsymbol{\eta}^\varepsilon$ for each $\varepsilon \in \mathbb{E}_\ell = \{0,1\}^\ell$.

We now define collections $\boldsymbol{\eta}^\varepsilon$ for all $\varepsilon \in \{0,1,\infty\}^\ell$ such that, when $\varepsilon \in \{0,1\}^\ell$, the respective collection coincides with the one already defined. Let γ_i be a simple closed curve in Σ disjoint from the basepoints and the beta curves, which specifies the framing Λ_i of the component L_i for $i = 1, \dots, \ell$. (In other words, this is the analogue of δ_i when we use Λ instead of $\widetilde{\Lambda}$.) We complete this to a full set of attaching circles $\boldsymbol{\gamma}$ by taking curves $\gamma_{\ell+1}, \dots, \gamma_{\ell+g-1}$ that approximate (ie are small Hamiltonian translates of) $\beta_{\ell+1}, \dots, \beta_{\ell_g-1}$. The Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{w})$ then represents the three-manifold $Y_\Lambda(L)$.

Given $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell) \in \{0,1,\infty\}^\ell$, we define the $(g+\ell-1)$-tuple of attaching circles $\boldsymbol{\eta}^\varepsilon$ by

$$\eta_i^\varepsilon \approx \begin{cases} \delta_i & \text{if } \varepsilon_i = 0, \\ \beta_i & \text{if } \varepsilon_i = 1, \\ \gamma_i & \text{if } \varepsilon_i = \infty. \end{cases}$$

For every $\varepsilon < \varepsilon'$, the Heegaard diagram $(\Sigma, \boldsymbol{\eta}^\varepsilon, \boldsymbol{\eta}^{\varepsilon'}, \boldsymbol{w})$ represents a connected sum of some copies of $S^1 \times S^2$ and lens spaces. As such, there is a canonical torsion Spin^c structure on this manifold; see [Ozsváth and Szabó 2008b, Definition 3.2] and Section 11.3. We arrange so that the Floer homology $\boldsymbol{HF}^-(\mathbb{T}_{\eta^\varepsilon}, \mathbb{T}_{\eta^{\varepsilon'}}, \boldsymbol{w})$ in that Spin^c structure, in the maximal degree with nonzero homology, is represented by a unique intersection point. We denote that point by $\Theta^{\text{can}}_{\varepsilon,\varepsilon'}$. We then define $\Theta_{\varepsilon,\varepsilon'}$ just as in (163).

Consider the ring

$$\mathscr{T} = \mathbb{F}[T_1, \dots, T_\ell]/(T_1^{m_1} - 1, \dots, T_\ell^{m_\ell} - 1).$$

Construct the chain complex with twisted coefficients $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}; \mathscr{T})$, which as a module is $\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}) \otimes_{\mathbb{F}} \mathscr{T}$, and comes equipped with the differential

$$\partial^- \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x},\boldsymbol{y}) \\ \mu(\phi)=1}} \#\mathcal{M}(\phi) \cdot \left(\prod_{i=1}^{\ell} T_i^{n_{w_i}(\phi) - n_{z_i}(\phi)} U_i^{n_{w_i}(\phi)} \right) \cdot \boldsymbol{y}. \tag{164}$$

Since Y is an integral homology sphere, all the periodic domains on the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w})$ are multiples of Σ. As a consequence, there exists an isomorphism of chain complexes

$$\boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}; \mathscr{T}) \cong \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}) \otimes_{\mathbb{F}} \mathscr{T} = \bigoplus^{m_1 m_2 \cdots m_\ell} \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w});$$

compare [ibid., equation (7)].

More generally, for $\varepsilon \in \{0, 1, \infty\}^\ell$, we let $\mathscr{T}^\varepsilon$ be the polynomial ring in variables T_i for those i such that $\varepsilon_i = 1$, and with relations $T_i^{m_i} = 1$. We consider the chain complex with twisted coefficients

$$\mathscr{C}^\varepsilon = \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}; \mathscr{T}^\varepsilon),$$

constructed as above, but in which we only keep track of the multiplicities $n_{w_i} - n_{z_i}$ (using T_i) for those i with $\varepsilon_i = 1$. Let

$$m^\varepsilon = \prod_{\varepsilon_i = 1} m_i.$$

We then have an identification of chain complexes

$$\mathscr{C}^\varepsilon = \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}; \mathscr{T}^\varepsilon) \cong \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}) \otimes_{\mathbb{F}} \mathscr{T}^\varepsilon = \bigoplus^{m^\varepsilon} \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}).$$

For multi-indices $\varepsilon \le \varepsilon'$ define linear maps

$$\mathscr{D}_\varepsilon^{\varepsilon'-\varepsilon} \colon \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}; \mathscr{T}^\varepsilon) \to \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^{\varepsilon'}}, \boldsymbol{w}; \mathscr{T}^{\varepsilon'}),$$
$$\mathscr{D}_\varepsilon^{\varepsilon'-\varepsilon}(\boldsymbol{x}) = \sum_p \sum_{\{\varepsilon = \varepsilon^0 < \cdots < \varepsilon^p = \varepsilon'\}} f(\boldsymbol{x} \otimes \Theta_{\varepsilon^0, \varepsilon^1} \otimes \cdots \otimes \Theta_{\varepsilon^{p-1}, \varepsilon^p}),$$

where f is the polygon map as in Section 3.5, which keeps track of the difference in multiplicities at w_i and z_i according to T_i, just as in (164) above. (Compare Section 8.3 and [Ozsváth and Szabó 2008b, Section 3; 2005, equation (9)].)

The direct sum $\bigoplus_{\varepsilon \in \{0,1,\infty\}^\ell} \mathscr{C}^\varepsilon$ forms a chain complex with differential $\mathscr{D} = \sum \mathscr{D}_\varepsilon^{\varepsilon'-\varepsilon}$. The proof that $\mathscr{D}^2 = 0$ is similar to that of [Ozsváth and Szabó 2005, Proposition 4.4], with the difference that here we use twisted coefficients as in the proof that $f_2^+ \circ f_1^+ \simeq 0$ in [Ozsváth and Szabó 2008b, Section 3]. Note that $\mathscr{D}$ is composed of ordinary twisted coefficients maps (and higher homotopies) between Floer complexes. In Section 11.3 (and, in particular, in Proposition 11.5) we needed a more complicated map denoted by f_3^-. In our current context, f_3^- would go from $\mathscr{C}^\varepsilon$ to $\mathscr{C}^{\varepsilon'}$ where in ε' we changed an ∞ coordinate of ε to 0. Maps of this type do not appear in $\mathscr{D}$.

Let us consider the quotient complex corresponding to $\varepsilon \in \mathbb{E}_\ell = \{0, 1\}^\ell \subset \{0, 1, \infty\}^\ell$. The projection of the differential (which we still denote by $\mathscr{D}$) turns this quotient complex into a hypercube of chain complexes, which we denote by $\mathscr{H}$.

Proposition 12.5 *The complex $\boldsymbol{CF}^-(Y_\Lambda(L)) = \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{w}) = \mathscr{C}^{(\infty,\ldots,\infty)}$ is quasi-isomorphic to the total complex of the hypercube $\mathscr{H} = (\mathscr{C}^\varepsilon, \mathscr{D}^\varepsilon)_{\varepsilon \in \mathbb{E}_\ell}$.*

Proof Iterate the quasi-isomorphism from Proposition 11.5 along the same lines as in the proof of [Ozsváth and Szabó 2005, Theorem 4.1]. Observe that in order to prove exactness in Proposition 11.5 we had to use the more complicated map f_3^-. However, having established in Proposition 11.5 that there are quasi-isomorphisms along the edges of our hypercube, we automatically obtain quasi-isomorphisms under iteration. Thus, there is no need to define higher homotopy versions of the map f_3^-. □

Remark 12.6 For now, we just regard Proposition 12.5 as a statement about ungraded complexes. The same goes for all the quasi-isomorphisms discussed below, until we deal with the question of relative gradings in Section 12.5.

We can give an alternative description of the hypercube $\mathscr{H}$ as follows. For $\varepsilon \in \mathbb{E}_\ell$, note that

$$\mathscr{C}^\varepsilon \cong \boldsymbol{CF}^-(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}) \otimes_{\mathbb{F}} \mathscr{T}^\varepsilon = \boldsymbol{CF}^-(Y_{\widetilde{\Lambda}|_{L-L^\varepsilon}}(L-L^\varepsilon)) \otimes_{\mathbb{F}} \mathscr{T}^\varepsilon,$$

where $L^\varepsilon \subset L$ is the sublink consisting of those components L_i such that $\varepsilon_i = 1$. Since $\widetilde{\Lambda} \gg 0$, the manifold $Y_{\widetilde{\Lambda}|_{L-L^\varepsilon}}(L-L^\varepsilon)$ is a rational homology three-sphere.

Consider now $\varepsilon, \varepsilon' \in \mathbb{E}_\ell$ with $\varepsilon \leq \varepsilon'$. Suppose $L^{\varepsilon'} = L^\varepsilon \cup L_{i_1} \cup \cdots \cup L_{i_p}$ for $p = \|\varepsilon' - \varepsilon\|$. If we set all the variables T_i equal to 1, the map $\mathscr{D}_\varepsilon^{\varepsilon'-\varepsilon}$ would simply be (several copies of) the polygon map $F(\varepsilon, \varepsilon', \boldsymbol{w})$ from Section 12.1. When we keep the variables T_i as they are, they keep track of the Spin^c structures on the cobordism $W'_{\widetilde{\Lambda}}(L-L^\varepsilon, L-L^{\varepsilon'})$, which is a two-handle attachment relating the manifolds $Y_{\widetilde{\Lambda}|_{L-L^\varepsilon}}(L-L^\varepsilon)$ and $Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon'}}}(L-L^{\varepsilon'})$. Indeed, by the argument used for [Ozsváth and Szabó 2008b, equation (8)], there is an identification

$$\operatorname{Spin}^c(W'_{\widetilde{\Lambda}}(L-L^\varepsilon, L-L^{\varepsilon'})) \xrightarrow{\cong} \mathbb{Z}^p, \quad \boldsymbol{k} \mapsto (k_1, \ldots, k_p), \tag{165}$$

such that

$$\mathscr{D}_\varepsilon^{\varepsilon'-\varepsilon} = \sum_{\boldsymbol{k} \in \mathbb{Z}^p} T^{\boldsymbol{k}} \cdot F(\varepsilon, \varepsilon', \boldsymbol{w}, \boldsymbol{k}), \tag{166}$$

where we formally wrote

$$T^{\boldsymbol{k}} = \prod_{j=1}^{p} T_{i_j}^{k_j}. \tag{167}$$

12.3 Modifying the hypercube

Fix an integer $\delta > 0$. As in [Ozsváth and Szabó 2008b], it is helpful to replace the Floer complexes $\boldsymbol{CF}^-$ by the corresponding complexes $\boldsymbol{CF}^{-,\delta}$; see Section 10.2. Proposition 12.5 has the following immediate consequence (which could also be obtained by iterating Proposition 11.6):

Proposition 12.7 *The complex $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L)) = \boldsymbol{CF}^{-,\delta}(\mathbb{T}_\alpha, \mathbb{T}_\gamma, \boldsymbol{w})$ is quasi-isomorphic to the total complex of the hypercube*

$$\mathscr{H}^\delta = (\mathscr{C}^{\varepsilon,\delta} = \boldsymbol{CF}^{-,\delta}(\mathbb{T}_\alpha, \mathbb{T}_{\eta^\varepsilon}, \boldsymbol{w}; \mathscr{T}^\varepsilon), \mathscr{D}^{\varepsilon,\delta})_{\varepsilon \in \mathbb{E}_\ell}.$$

Here is the vertical truncation of (166):

$$\mathscr{D}_\varepsilon^{\varepsilon'-\varepsilon,\delta} = \sum_{\boldsymbol{k} \in \mathbb{Z}^p} T^{\boldsymbol{k}} \cdot F(\varepsilon, \varepsilon', \boldsymbol{w}, \boldsymbol{k})^\delta. \tag{168}$$

Let $L^{\varepsilon'} = L^{\varepsilon} \amalg M$, with $M = L^{\varepsilon'-\varepsilon} = L_{i_1} \cup \cdots \cup L_{i_p}$. Note that among the Spin^c structures on the cobordism $W'_{\widetilde{\Lambda}}(L - L^{\varepsilon}, L - L^{\varepsilon'})$ there are some special ones, namely those of the form

$$\mathfrak{x}_s^{\mathfrak{o}} = \mathfrak{x}_s^{\mathfrak{o}}|_{W'_{\widetilde{\Lambda}}(L-L^{\varepsilon}, L-L^{\varepsilon'})}$$

(see Section 12.1) for all possible orientations $\mathfrak{o}$ of L and Spin^c structures s on $Y_{\widetilde{\Lambda}}(L)$. Note that $\mathfrak{x}_s^{\mathfrak{o}}$, viewed as a Spin^c structure on $W'_{\widetilde{\Lambda}}(L - L^{\varepsilon}, L - L^{\varepsilon'})$, only depends on the restriction of $\mathfrak{o}$ to M. Indeed, from (161), we deduce

$$c_1(\mathfrak{x}_s^{\mathfrak{o}}|_{W'_{\widetilde{\Lambda}}(L-L^{\varepsilon}, L-L^{\varepsilon'})}) = c_1(\mathfrak{x}_s|_{W'_{\widetilde{\Lambda}}(L-L^{\varepsilon}, L-L^{\varepsilon'})}) + 2(\widetilde{\Lambda}|_M)_{\vec{L}, \vec{M}^{\mathfrak{o}}}.$$

To simplify notation, for $\vec{M} \in \Omega(M)$ we denote by $\mathfrak{x}_s^{\vec{M}}$ the restriction to $W'_{\widetilde{\Lambda}}(L - L^{\varepsilon}, L - L^{\varepsilon'})$ of any Spin^c structure $\mathfrak{x}_s^{\mathfrak{o}}$ such that the restriction $\vec{M}^{\mathfrak{o}}$ of $\mathfrak{o}$ to M is $\vec{M}$. Then, with respect to the identification $W'_{\widetilde{\Lambda}}(L - L^{\varepsilon}, L - L^{\varepsilon'}) \cong \mathbb{Z}^p$ from (165),

$$\mathfrak{x}_s^{\vec{M}} = \mathfrak{x}_s + (\widetilde{\Lambda}|_M)_{\vec{L}, \vec{M}}. \tag{169}$$

The advantage of using truncated maps comes from the following:

Lemma 12.8 *Fix $\delta > 0$. Then, for sufficiently large $\widetilde{\Lambda} \gg 0$ (compared to δ) and for any $\varepsilon, \varepsilon' \in \mathbb{E}_\ell$ with $\varepsilon < \varepsilon'$,*

$$F(\varepsilon, \varepsilon', \boldsymbol{w}, \boldsymbol{k})^{\delta} = 0$$

whenever the Spin^c structure $\boldsymbol{k}$ on $W'_{\widetilde{\Lambda}}(L - L^{\varepsilon}, L - L^{\varepsilon'})$ is not of the form $\mathfrak{x}_s^{\vec{M}}$ for any $s \in \mathrm{Spin}^c(Y_{\widetilde{\Lambda}}(L)) = P(\widetilde{\Lambda})$ and $\vec{M} \in \Omega(L^{\varepsilon'-\varepsilon})$. Moreover, there is a constant $b^{\delta} > 0$ such that

$$F(\varepsilon, \varepsilon', \boldsymbol{w}, \mathfrak{x}_s^{\mathfrak{o}})^{\delta} = 0$$

whenever $s = (s_1, \dots, s_\ell) \in P(\widetilde{\Lambda}) \subset \mathbb{R}^\ell$ admits some $i \in \{1, \dots, \ell\}$ with the property that either

- *$s_i > b^{\delta}$ and $i \in I_-(\vec{L}, \vec{M})$, or*
- *$s_i < -b^{\delta}$ and $i \in I_+(\vec{L}, \vec{M})$.*

Proof Since both $Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon}}}(L - L^{\varepsilon})$ and $Y_{\widetilde{\Lambda}|_{L-L^{\varepsilon'}}}(L - L^{\varepsilon'})$ are rational homology three-spheres, the respective Floer complexes $\mathscr{C}^{\varepsilon,\delta}$ and $\mathscr{C}^{\varepsilon',\delta}$ admit absolute $\mathbb{Q}$-gradings; see [Ozsváth and Szabó 2006]. Because these complexes are vertically truncated, their absolute gradings lie in a finite range. Each map $F(\varepsilon, \varepsilon', \boldsymbol{w}, \boldsymbol{k})^{\delta}$ shifts the grading by a definite amount. The vanishing of the claimed maps happens because the respective shifts take the range for the initial complex to outside the grading range for the final complex. We leave the verification of the details to the interested reader. In the case $p = 1$, this was done in [Ozsváth and Szabó 2008b, Lemma 4.4 and Section 4.3]. □

Observe that the restriction of the Spin^c structure $\mathfrak{x}_s^{\vec{M}}$ to the cobordism $W'_{\widetilde{\Lambda}}(L - L^{\varepsilon}, L - L^{\varepsilon'})$ only depends on the value $\psi^{L^{\varepsilon}}(s) \in P(\widetilde{\Lambda}|_{L-L^{\varepsilon}})$. (As we recall from Remark 10.11, the hyperparallelepipeds $P(\widetilde{\Lambda})$ behave well with respect to restriction to sublinks.) Thus, we can write $\mathfrak{x}_{\bar{s}}^{\vec{M}}$ for $\mathfrak{x}_s^{\vec{M}}$ whenever $\bar{s} = \psi^{L^{\varepsilon}}(s) \in P(\widetilde{\Lambda}|_{L-L^{\varepsilon}})$.

By (168) and Lemma 12.8, for $\widetilde{\Lambda} \gg 0$,

$$\mathscr{D}_{\varepsilon}^{\varepsilon'-\varepsilon,\delta} = \sum_{\vec{M} \in \Omega(L^{\varepsilon'-\varepsilon})} \sum_{s \in P(\widetilde{\Lambda}|_{L-L^{\varepsilon}})} T^{\mathfrak{x}_s^{\vec{M}}} \cdot F(\varepsilon, \varepsilon', \boldsymbol{w}, \mathfrak{x}_{\boldsymbol{s}}^{\vec{M}})^{\delta}.$$

By adjusting the identifications (165) if necessary, and taking (169) into account, we can simply write

$$\mathscr{D}_{\varepsilon}^{\varepsilon'-\varepsilon,\delta} = \sum_{\vec{M} \in \Omega(L^{\varepsilon'-\varepsilon})} \sum_{s \in P(\widetilde{\Lambda}|_{L-L^{\varepsilon}})} T^{(\widetilde{\Lambda}|_{L^{\varepsilon'-\varepsilon}})_{\vec{L},\vec{M}}} \cdot F(\varepsilon, \varepsilon', \boldsymbol{w}, \mathfrak{x}_{\boldsymbol{s}}^{\vec{M}})^{\delta}.$$

We have $T^{\widetilde{\Lambda}_i - \Lambda_i} = 1$ for all i, by the definition of the T_i variables. Hence,

$$T^{(\widetilde{\Lambda}|_M)_{\vec{L},\vec{M}}} = T^{(\Lambda|_M)_{\vec{L},\vec{M}}}.$$

To summarize, the total complex $\mathscr{C}^{\delta}$ of the hypercube $\mathscr{H}^{\delta}$ is

$$\mathscr{C}^{\delta} = \bigoplus_{\varepsilon \in \mathbb{E}_{\ell}} \bigoplus_{s \in P(\widetilde{\Lambda}|_{L^{\varepsilon}})} \boldsymbol{CF}^{-,\delta}(\mathbb{T}_{\alpha}, \mathbb{T}_{\eta^{\varepsilon}}, \boldsymbol{w}, \boldsymbol{s}) \otimes \mathscr{T}^{\varepsilon} \tag{170}$$

with the differential on each summand being

$$\mathscr{D}^{\delta} = \sum_{\varepsilon' \geq \varepsilon} \sum_{\vec{M} \in \Omega(L^{\varepsilon'-\varepsilon})} T^{(\Lambda|_{L^{\varepsilon'-\varepsilon}})_{\vec{L},\vec{M}}} \cdot F(\varepsilon, \varepsilon', \boldsymbol{w}, \mathfrak{x}_{\boldsymbol{s}}^{\vec{M}})^{\delta}. \tag{171}$$

On the other hand, in Section 10.4 we constructed a folded truncated complex $\mathcal{C}^{-,\delta}(\mathcal{H}, \Lambda)\{\widetilde{\Lambda}\}$ quasi-isomorphic to $\mathcal{C}^{-,\delta}(\mathcal{H}, \Lambda)$. By rephrasing its description from (116), (117), we have

$$\mathcal{C}^{-,\delta}(\mathcal{H}, \Lambda)\{\widetilde{\Lambda}\} = \bigoplus_{\varepsilon \in \mathbb{E}_{\ell}} \bigoplus_{s \in P(\widetilde{\Lambda}|_{L-L^{\varepsilon}})} \mathfrak{A}^{-,\delta}(\mathcal{H}^{L-L^{\varepsilon}}, \boldsymbol{s}) \otimes \mathscr{T}^{\varepsilon} \tag{172}$$

with the differential

$$\mathcal{D}^{-,\delta} = \sum_{\varepsilon' \geq \varepsilon} \sum_{\vec{M} \in \Omega(L^{\varepsilon'-\varepsilon})} T^{(\Lambda|_{L^{\varepsilon'-\varepsilon}})_{\vec{L},\vec{M}}} \cdot \Phi_{\boldsymbol{s}}^{L^{\varepsilon'-\varepsilon},\delta}. \tag{173}$$

Looking at (171), we observe that the maps $F(\varepsilon, \varepsilon', \boldsymbol{w}, \mathfrak{x}_{\boldsymbol{s}}^{\vec{M}})^{\delta}$ also appear in the vertical truncations (by δ) of the hypercubes $\mathscr{H}_{\boldsymbol{s}}^{o}$ considered in the statement of Proposition 12.4. Similarly, looking at (173), we observe that the maps $\Phi_{\boldsymbol{s}}^{L^{\varepsilon'-\varepsilon},\delta}$ appear in the vertical truncation of the other hypercube $\mathcal{H}^{o}$ considered in the statement of Proposition 12.4. We seek to apply the result of Proposition 12.4 to obtain a quasi-isomorphism between the hypercubes $\mathscr{H}^{\delta}$ and $\mathcal{H}^{\delta}$. Before doing so, however, we need a basic result from homological algebra:

Lemma 12.9 *Let A, B, A' and B' be hypercubes of chain complexes (of the same dimension d, with differentials denoted by ∂), and*

$$F_1, F_2 \colon A \to B, \quad F_1', F_2' \colon A' \to B'$$

be chain maps. Suppose that, for $i = 1, 2$, the corresponding $(d+1)$-dimensional hypercubes for the maps F_i and F_i' (that is, the mapping cones $A \xrightarrow{F_i} B$ and $A' \xrightarrow{F_i} B'$) can be related by a quasi-isomorphism consisting of maps (Φ_i, Ψ_i, h_i) as in the diagram

$$\begin{tikzcd} A \arrow[r, "\Phi_i"] \arrow[d, "F_i"'] \arrow[dr, "h_i"] & A' \arrow[d, "F_i'"] \\ B \arrow[r, "\Psi_i"'] & B' \end{tikzcd}$$

Assume that Φ_1 and Ψ_1 are quasi-isomorphisms. Further, suppose that the maps Φ_1 and Φ_2 are chain homotopic, and so are Ψ_1 and Ψ_2. Then the mapping cones $A \xrightarrow{F_1+F_2} B$ and $A' \xrightarrow{F_1'+F_2'} B'$, viewed as $(d+1)$-dimensional hypercubes, are quasi-isomorphic as well.

Proof By hypothesis,

$$F_1'\Phi_1 - \Psi_1 F_1 = \partial h_1 + h_1\partial, \quad F_2'\Phi_2 - \Psi_2 F_2 = \partial h_2 + h_2\partial,$$

and there are homotopies $\phi\colon A \to A'$, $\psi\colon B \to B'$ such that

$$\Phi_1 - \Phi_2 = \partial\phi + \phi\partial, \quad \Psi_1 - \Psi_2 = \partial\psi + \psi\partial.$$

It follows that

$$(F_1' + F_2')\Phi_1 - \Psi_1(F_1 + F_2) = \partial g + g\partial,$$

where

$$g = h_1 + h_2 + F_2'\phi + \psi F_2.$$

Thus, (Φ_1, Ψ_1, g) form a chain map between the mapping cones $A \xrightarrow{F_1+F_2} B$ and $A' \xrightarrow{F_1'+F_2'} B'$. This map is a quasi-isomorphism because Φ_1 and Ψ_1 are so; compare Definition 5.20. $\square$

Proposition 12.10 *Given a basic system $\mathscr{H}$ for a link $\vec{L}$ in an integral homology sphere Y, there is an isomorphism*

$$H_*(\mathscr{C}^{-,\delta}(\mathscr{H}, \Lambda), \mathscr{D}^{-,\delta}) \cong \boldsymbol{HF}_*^{-,\delta}(Y_\Lambda(L)).$$

Proof Let us denote by $\mathscr{H}_s^{\mathfrak{o},\delta}$ and $\mathscr{H}_s^{\mathfrak{o},\delta}$ the vertical truncations by δ of the hypercubes $\mathscr{H}_s^{\mathfrak{o}}$ and $\mathscr{H}_s^{\mathfrak{o}}$ considered in Proposition 12.4. The result of that proposition implies (after vertical truncation) that, for any orientation $\mathfrak{o}$, the hypercubes $\mathscr{H}_s^{\mathfrak{o},\delta}$ and $\mathscr{H}_s^{\mathfrak{o},\delta}$ are related by a quasi-isomorphism. Further, if we change the orientation $\mathfrak{o}$, it is easy to see that the respective quasi-isomorphisms are chain homotopic.

The hypercube $\mathscr{H}^\delta = (\mathscr{C}^{-,\delta}(\mathscr{H}, \Lambda)\{\widetilde{\Lambda}\}, \mathscr{D}^{-,\delta})$ described in (172)–(173), is basically obtained from the hypercubes $\prod_s \mathscr{H}_s^{\mathfrak{o},\delta}$ for all possible choices of orientations $\mathfrak{o} \in \Omega(L)$, by gluing those hypercubes along their common parts (corresponding to sublinks). More precisely, this gluing process is an iteration of the one that made an appearance in Lemma 12.9: given two mapping cone hypercubes $A \xrightarrow{F_1} B$ and $A \xrightarrow{F_2} B$, we replace them by $A \xrightarrow{F_1+F_2} B$. (See Figure 38 for an illustration of the gluing procedure in

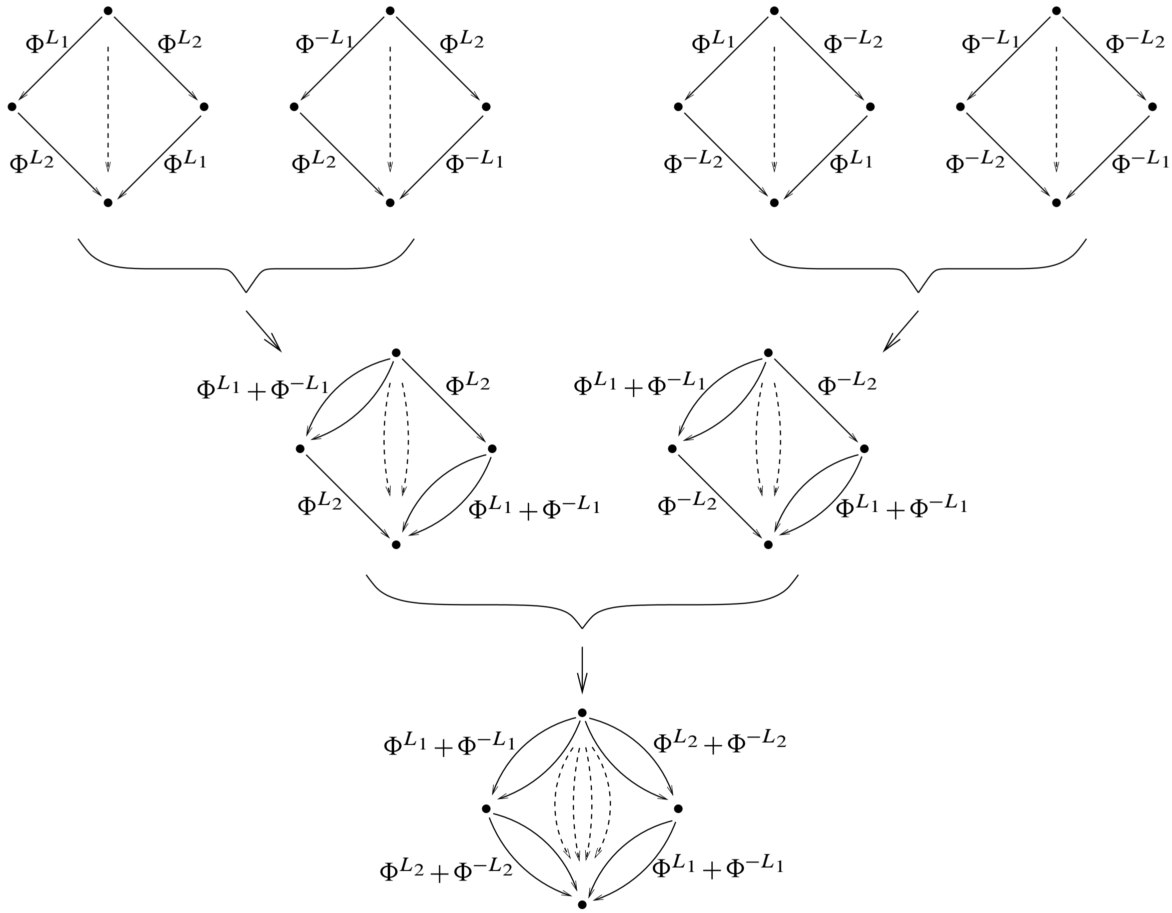

Figure 38: Gluing hypercubes in the proof of Proposition 12.10. Given a two-component link $\vec{L}=L_1\cup L_2$, each of the four squares in the top row represent one of the hypercubes $\prod_{\boldsymbol{s}}\mathcal{H}^{\mathfrak{o}}_{\boldsymbol{s}}$ (with the factors $\mathcal{H}^{\mathfrak{o}}_{\boldsymbol{s}}$ as defined in Proposition 12.4) for the four possible orientations $\mathfrak{o}$ of L. Combining these hypercubes (by adding up the respective differentials, as shown) yields the hypercube $\mathcal{H}$ at the bottom, which is the one appearing in the statement of the surgery theorem, Theorem 9.6. The dashed lines represent chain homotopies such as $\Phi^{L_1\cup L_2}$, $\Phi^{(-L_1)\cup L_2}$, etc.

the case $\ell=2$.) Observe also that the hypercube $\mathscr{H}^{\delta}$ described in (170)–(171) is obtained by from the hypercubes $\prod_{\boldsymbol{s}}\mathscr{H}^{\mathfrak{o},\delta}_{\boldsymbol{s}}$ by gluing them using the same process.

By applying Lemma 12.9 repeatedly, we can construct a quasi-isomorphism between $(\mathscr{C}^{\delta},\mathscr{D}^{\delta})$ and $(\mathcal{C}^{-,\delta}(\mathcal{H},\Lambda)\{\widetilde{\Lambda}\},\mathcal{D}^{-,\delta})$. The former complex is quasi-isomorphic to $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L))$ by Proposition 12.7, and the latter to $\mathcal{C}^{-,\delta}(\mathcal{H},\Lambda)$ by Proposition 10.15. $\square$

12.4 Spin^c structures

Recall from Section 9.2 that the complexes $\mathcal{C}^-(\mathcal{H},\Lambda)$ and $\boldsymbol{CF}^-(Y_\Lambda(L))$ both break into direct sums of complexes $\mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u})$ and $\boldsymbol{CF}^-(Y_\Lambda(L),\mathfrak{u})$ according to the set $\mathrm{Spin}^c(Y_\Lambda(L))\cong\mathbb{H}(L)/H(L,\Lambda)$. Of course, the same is true for their vertical truncations by δ. We would like to prove that the isomorphism in Proposition 12.10 preserves these decompositions.

As a warm-up exercise toward this goal, let us investigate to what extent we can make the total complex $\mathscr{C}^\delta$ of the hypercube $\mathscr{H}^\delta$ break into a direct sum according to Spin^c structures $\mathfrak{u} \in \mathbb{H}(L)/H(L,\Lambda)$.

We claim that this splitting can be realized when the lattice $H(L,\Lambda)$ is nondegenerate, ie the vectors Λ_i are linearly independent over $\mathbb{Q}$ or, equivalently, $Y_\Lambda(L)$ is a rational homology three-sphere. Indeed, from (166) we see that the complex $\mathscr{C}^\delta$ breaks into a direct sum according to equivalence classes of Spin^c structures on $Y_{\widetilde{\Lambda}}(L)$, where two structures in $\mathrm{Spin}^c(Y_{\widetilde{\Lambda}}(L)) \cong \mathbb{H}(L)/H(L,\widetilde{\Lambda})$ are equivalent if they differ by an element in the $\mathbb{Z}$-span of the vectors $m_i\tau_i = \widetilde{\Lambda}_i - \Lambda_i$ for $i = 1,\dots,\ell$. In other words, the decomposition is according to $\mathbb{H}(L)/H(L,\Lambda,\widetilde{\Lambda})$, where $H(L,\Lambda,\widetilde{\Lambda}) \subseteq \mathbb{Z}^\ell$ is the lattice generated by all Λ_i and $\widetilde{\Lambda}_i$. Since Λ is nondegenerate, we can arrange so that $H(L,\Lambda,\widetilde{\Lambda}) = H(L,\Lambda)$ by simply choosing $m_i \gg 0$ with $m_i\tau_i \in H(L,\Lambda)$, for all i. With this choice of $\widetilde{\Lambda}$, our claim about the splitting of $\mathscr{C}^\delta$ holds true.

When $H(L,\Lambda)$ is degenerate, we need to refine this approach, and settle for a splitting of a complex quasi-isomorphic to $\mathscr{C}^\delta$, rather than one of $\mathscr{C}^\delta$ itself. Indeed, since the vectors Λ_i do not span $\mathbb{Q}^\ell$, we cannot always find $m_i \in \mathbb{Z}$ such that $m_i\tau_i \in H(L,\Lambda)$. Nevertheless, let us first choose some arbitrary $m_i' \gg 0$ such that the resulting framing $\widetilde{\Lambda}'$ satisfies the conditions in Lemma 12.8. Next, suppose the vectors Λ_i span a subspace $\mathscr{V} \subset \mathbb{Q}^\ell$ of dimension $\ell - p$, with $p > 0$. Choose p coordinate vectors τ_i such that their span is complementary to $\mathscr{V}$. Without loss of generality, assume those coordinate vectors are τ_i for $i = 1,\dots,p$. Choose $m_i \gg m_i'$ arbitrarily for $i = 1,\dots,p$, and let $H(L,\Lambda)^+ \subset \mathbb{Z}^\ell$ be the lattice spanned by Λ_i's together with $m_1\tau_1,\dots,m_p\tau_p$. Now choose $m_i \gg m_i'$ for $i = p+1,\dots,\ell$ such that $m_i\tau_i \in H(L,\Lambda)^+$. The result is a framing $\widetilde{\Lambda} \gg \widetilde{\Lambda}'$.

Lemma 12.8 says that the complex $\mathscr{C}^\delta$ (constructed with respect to the framing $\widetilde{\Lambda}$) is similar in structure to the complex $\mathscr{C}^\delta = \mathscr{C}^{-,\delta}(\mathscr{H},\Lambda)$, compare Lemma 10.10. Hence, we can apply the combined truncation procedure from Section 10.4 to $\mathscr{C}^\delta$, using the smaller framing $\widetilde{\Lambda}'$. The result is a complex $\mathscr{C}^\delta\langle\!\langle\widetilde{\Lambda}'\rangle\!\rangle$ quasi-isomorphic to $\mathscr{C}^\delta$.

The advantage of the complex $\mathscr{C}^\delta\langle\!\langle\widetilde{\Lambda}'\rangle\!\rangle$ is that it splits as a direct sum according to $\mathbb{H}(L)/H(L,\Lambda)$. Indeed, we know it splits according to $\mathbb{H}(L)/H(L,\Lambda,\widetilde{\Lambda})$ just like $\mathscr{C}^\delta$. Moreover, it is supported roughly on the hyperparallelepiped $P(\widetilde{\Lambda}')$, which is very small compared to $P(\widetilde{\Lambda})$. The key observation is that, with our choice of the values m_i, if two elements in $P(\widetilde{\Lambda}')$ differ by an element in $H(L,\Lambda,\widetilde{\Lambda})$, they differ by an element in $H(L,\Lambda)$. This is true because $H(L,\Lambda,\widetilde{\Lambda})$ is contained in the union of parallel subspaces

$$\bigcup_{t_1,\dots,t_p\in\mathbb{Z}} (\mathscr{V} + t_1m_1\tau_1 + \dots + t_pm_p\tau_p). \tag{174}$$

Set $\Delta(P(\widetilde{\Lambda}')) = \{\mathbf{s} - \mathbf{s}' \mid \mathbf{s},\mathbf{s}' \in P(\widetilde{\Lambda}')\}$. If $\Delta(P(\widetilde{\Lambda}'))$ is sufficiently small compared to the values m_i, the only one of the parallel subspaces in the union (174) that intersects it nontrivially is $\mathscr{V}$ itself. This implies our claim about the decomposition of $\mathscr{C}^\delta\langle\!\langle\widetilde{\Lambda}'\rangle\!\rangle$ according to $\mathbb{H}(L)/H(L,\Lambda)$.

Observe that similar remarks are applicable to the folded truncated complex $\mathcal{C}^{-,\delta}(\mathcal{H},\Lambda)\{\widetilde{\Lambda}\}$. When $H(L,\Lambda)$ is nondegenerate, it splits according to Spin^c structures on $Y_{\widetilde{\Lambda}}(L)$, provided the values m_i are chosen suitably. When $H(L,\Lambda)$ is degenerate, we can replace it by a quasi-isomorphic complex

$$\mathcal{C}^{-,\delta}(\mathcal{H},\Lambda)\{\widetilde{\Lambda}\}\langle\!\langle\widetilde{\Lambda}'\rangle\!\rangle = \mathcal{C}^{-,\delta}(\mathcal{H},\Lambda)\langle\!\langle\widetilde{\Lambda}'\rangle\!\rangle,$$

which again has the desired splitting.

12.5 Relative gradings

As explained in Sections 9.2 and 9.3, for every $\mathfrak{u}\in\mathrm{Spin}^c(Y_\Lambda(L))$, both complexes

$$\mathcal{C}^{-,\delta}(\mathcal{H},\Lambda,\mathfrak{u}) \text{ and } \boldsymbol{CF}^{-,\delta}(Y_\Lambda(L),\mathfrak{u})$$

admit relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-gradings, where $\mathfrak{d}(\mathfrak{u})\in\mathbb{Z}$ depends on $\mathfrak{u}$. Our goal is to prove the following strengthening of Proposition 12.10:

Proposition 12.11 *Let $\mathcal{H}$ be a basic complete system for an oriented link $\vec{L}$ inside an integral homology sphere Y. We then have $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-grading-preserving isomorphisms*

$$H_*(\mathcal{C}^{-,\delta}(\mathcal{H},\Lambda,\mathfrak{u}),\mathcal{D}^{-,\delta}) \cong \boldsymbol{HF}_*^{-,\delta}(Y_\Lambda(L),\mathfrak{u}) \tag{175}$$

and

$$H_*^{\delta\leftarrow\delta'}(\mathcal{C}^{-}(\mathcal{H},\Lambda,\mathfrak{u})) \cong \boldsymbol{HF}_*^{-,\delta\leftarrow\delta'}(Y_\Lambda(L),\mathfrak{u}) \tag{176}$$

for all $\delta'\geq\delta\geq 0$ and $\mathfrak{u}\in\mathrm{Spin}^c(Y_{\widetilde{\Lambda}}(L))$.

Proof Recall that the quasi-isomorphism in Proposition 12.7 can be obtained by iterating Proposition 11.6. Indeed, let $\overline{\Lambda}$ be the framing on L obtained from Λ by adding m_1 to the coefficient of the first component, as in Section 11.3. Also, we let $L' = L - L_1$ and denote by Λ' the restriction of the framing Λ to L', Consider the ring $\mathcal{T}_1 = \mathbb{F}[T_1]/(T_1^{m_1}-1)$. The iteration process in the proof of Proposition 12.7 starts by applying Proposition 11.6 to get that $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L)) = \mathcal{C}^{(\infty,\infty,\dots,\infty),\delta}$ (in the notation of Section 12.2) is quasi-isomorphic to the mapping cone complex

$$\boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L)) \xrightarrow{f_2^\delta} \boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L');\mathcal{T}_1), \tag{177}$$

where the left-hand side is the Floer complex $\mathcal{C}^{(0,\infty,\dots,\infty),\delta}$, the right-hand side is the Floer complex $\mathcal{C}^{(1,\infty,\dots,\infty),\delta}$ with twisted coefficients, and the map f_2^δ is the triangle-counting map $\mathcal{D}^{(1,0,\dots,0),\delta}_{(0,\infty,\dots,\infty)}$ from (168).

The next step in the iteration process will be to show that each of the two sides in (177) is itself quasi-isomorphic to a mapping cone (for Floer complexes corresponding to multi-indices in which another one of the ∞ components is replaced by 0 and 1). These quasi-isomorphisms extend to give a quasi-isomorphism between $\boldsymbol{CF}^-(Y_\Lambda(L))$ and a two-dimensional hypercube of complexes. We continue this until we get a quasi-isomorphism between $\boldsymbol{CF}^-(Y_\Lambda(L))$ and the total complex of the hypercube $\mathscr{H}^\delta$.

Note that we have some freedom in this iteration: we could change the ordering of the components and start with L_2 instead of L_1, for example. We will choose the ordering as follows. For every $i = 1, \dots, \ell$, denote by $\Lambda^{(i)}$ the restriction of Λ to

$$L_i \cup L_{i+1} \cup \cdots \cup L_\ell,$$

and by $\Lambda_i^{(i)}, \dots, \Lambda_\ell^{(i)}$ the respective framing vectors. We require that, for every $i = 1, \dots, \ell$, either $\Lambda^{(i)}$ is nondegenerate, or else $\Lambda_i^{(i)}$ is in the $\mathbb{Q}$-span of $\Lambda_{i+1}^{(i)}, \dots, \Lambda_\ell^{(i)}$. Note that this can easily be arranged, by choosing which component we call L_1 first, then which component we call L_2, and so on.

Another degree of freedom in the iteration comes from the direction of the quasi-isomorphisms. In order to apply the results of Section 11, at each step we have to choose a quasi-isomorphism as in (133) or as in (134), depending on the framing Λ. At the first step, when we relate $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L))$ to $\mathrm{Cone}(f_2^\delta)$, we choose (133) if Λ is nondegenerate and the restriction of Λ to $\mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)^\perp$ is zero or positive definite. (This corresponds to Cases I(a) and I(c) discussed in Section 11.5.) We choose (133) if Λ is nondegenerate and the restriction of Λ to $\mathrm{Span}(\Lambda_2, \dots, \Lambda_\ell)^\perp$ is negative definite. When Λ is degenerate (so, by our choice of ordering, Λ_1 is in the span of the other framing vectors), we are free to choose either (133) or (134); compare Case II in Section 11.6.

At the second step, we need to combine a quasi-isomorphism relating $\boldsymbol{CF}^{-,\delta}(Y_{\overline{\Lambda}}(L))$ to a mapping cone, and one relating $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda'}(L'); \mathscr{T}_1)$ to a mapping cone. Of course, these quasi-isomorphisms should go in the same direction. We choose the direction according to the same recipe as at the first step. Precisely, if $\overline{\Lambda}$ is nondegenerate, we choose the direction based on the sign of the restriction of $\overline{\Lambda}$ to $\mathrm{Span}(\overline{\Lambda}_1, \Lambda_3, \dots, \Lambda_\ell)^\perp$. In particular, if Λ was nondegenerate to start with, then $\overline{\Lambda}$ is automatically nondegenerate and, for m_1 sufficiently large, the relevant sign is negative if and only if the sign of the restriction of $\overline{\Lambda}$ to $\mathrm{Span}(\tau_1, \Lambda_3, \dots, \Lambda_\ell)^\perp$ is negative; or, equivalently, if the sign of the restriction of Λ' to $\mathrm{Span}(\Lambda_3, \dots, \Lambda_\ell)^\perp$ is negative. If $\overline{\Lambda}$ is degenerate (for $m_1 \gg 0$), then Λ' is degenerate also, and we choose the direction of the quasi-isomorphisms arbitrarily.

We continue to choose directions this way at the following steps. At step i, we look at the framing matrix

$$\overline{\Lambda}^i = (\Lambda_1 + m_1\tau_1, \dots, \Lambda_{i-1} + m_{i-1}\tau_{i-1}, \Lambda_i, \dots, \Lambda_\ell)$$

for $m_1, \dots, m_{i-1} \gg 0$. If it is degenerate, we choose the direction arbitrarily. If it is nondegenerate, we choose it according to the sign of the restriction of $\overline{\Lambda}^i$ to the subspace

$$\mathrm{Span}(\Lambda_1 + m_1\tau_1, \dots, \Lambda_{i-1} + m_{i-1}\tau_{i-1}, \Lambda_{i+1}, \dots, \Lambda_\ell)^\perp.$$

With these choices of ordering and quasi-isomorphism directions, at every step in the iteration process we can apply one of the refinements of Proposition 11.6 discussed in Sections 11.5–11.6, namely

Proposition 11.11, 11.15, 11.17, 11.24 or 11.28. We first choose $m_1 \gg 0$ in such a way that the respective proposition applies, then we choose $m_2 \gg 0$, and so on. At the second step and later we may need to use the twisted coefficients variant, Proposition 11.29. We claim that in the end we obtain the desired isomorphisms (175), one for each $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$.

Indeed, at least in the case when Λ and all $\Lambda^{(i)}$'s are nondegenerate, all the complexes appearing in the proof of Proposition 12.10 decompose according to $\mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$. It is straightforward to check that the decompositions correspond to each other under the respective quasi-isomorphisms, and these quasi-isomorphisms preserve the relative $\mathbb{Z}$-gradings. Note that when iterating the exact sequences which give quasi-isomorphisms between mapping cones, at later steps in addition to the maps involved in Propositions 11.11 and 11.15 we also have certain higher homotopies. However, these decompose into Spin^c structures and preserve the relative $\mathbb{Z}$-gradings by the same arguments as those used in the discussion of the homotopy $H^{\delta}_{1,\mathfrak{u}}$ in Proposition 11.11, for example.

When $\Lambda = \Lambda^{(1)}$ or one of the other $\Lambda^{(i)}$'s is degenerate, there are two additional complications. First, in order to get a good decomposition into Spin^c structures we have to replace the complex $\mathscr{C}^{\delta}$ by a horizontally truncated, quasi-isomorphic one $\mathscr{C}^{\delta}\langle\langle\widetilde{\Lambda}'\rangle\rangle$, as discussed in Section 12.4. (When applying Propositions 11.24 or 11.28, this corresponds to focusing on a subset of all $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L))$, and choosing the respective value m_1 so that $\mathfrak{d}(\mathfrak{u})$ divides m_1 for all such $\mathfrak{u}$; compare Corollary 11.27. The unused Spin^c structures $\mathfrak{u}$ give rise to trivial complexes, so we can ignore them.) We then get a decomposition of $\mathscr{C}^{\delta}\langle\langle\widetilde{\Lambda}'\rangle\rangle$ according to Spin^c structures $\mathfrak{u}$. We have similar decompositions of $\mathfrak{C}^{\delta}\langle\langle\widetilde{\Lambda}'\rangle\rangle$, as well as of all the other complexes in the proof of Proposition 12.10, provided we truncate them with respect to $\widetilde{\Lambda}'$. The respective truncations are quasi-isomorphic to the original complexes. Putting everything together, we obtain the desired isomorphisms (175), for any $\mathfrak{u} \in \mathrm{Spin}^c(Y_{\widetilde{\Lambda}}(L))$.

The second complication has to do with the grading-preserving properties of the isomorphisms (175). If $\mathfrak{d}(\mathfrak{u}) = 0$, then, when we apply Proposition 11.17 at a step in the iteration process, the respective quasi-isomorphism preserves only the relative $\mathbb{Z}/2d\mathbb{Z}$-reduction of the relative $\mathbb{Z}$-grading. (Indeed, its target is only $\mathbb{Z}/2d\mathbb{Z}$-graded.) Thus, the resulting isomorphism (175) only preserves this $\mathbb{Z}/2d\mathbb{Z}$-grading. However, we can get such an isomorphism for any d in a sequence $\{d_n\}$ with $d_n \to \infty$. Both sides of (175) are finite-dimensional, relatively $\mathbb{Z}$-graded vector spaces, so if they are related by a $\mathbb{Z}/2d_n\mathbb{Z}$-grading-preserving isomorphism for all d_n, they must in fact be isomorphic as relatively $\mathbb{Z}$-graded vector spaces. This completes the proof of the claim about the existence of a grading-preserving isomorphism (175).

Finally, as δ varies, the isomorphisms (175) commute with the natural maps between the respective truncations, and we get the isomorphisms (176). □

Proof of Theorem 9.6 for basic systems Apply (175) and Lemma 10.4(b) for $\mathfrak{u}$ torsion, and (176) and Lemma 10.8(b) for $\mathfrak{u}$ nontorsion. □

12.6 Link-minimal complete systems

Now that we have established the truth of Theorem 9.6 for basic systems, we are left to do so for all link-minimal complete systems.

Proof of Theorem 9.6 Let $\mathscr{H}$ be a link-minimal complete system of hyperboxes for a link L, and let $\mathscr{H}_b$ be a basic system for L. According to Proposition 8.35(b), the system $\mathscr{H}$ can be obtained from $\mathscr{H}_b$ by a sequence of system moved that do not include any index zero/three link stabilizations; ie by a sequence of three-manifold isotopies, index one/two stabilizations and destabilizations, free index zero/three stabilizations, global shifts, elementary enlargements and contractions, and trajectory α-slides.

According to Proposition 8.19, except for the trajectory α-slides, all the other moves induce quasi-isomorphisms (in fact, chain homotopy equivalences) between the compressions of the respective hyperboxes of generalized Floer complexes $\mathfrak{A}^-(\mathscr{H}^{L',M}, s)$ and $\mathfrak{A}^-(\mathscr{H}_b^{L',M}, s)$. These compressions are the building blocks of the surgery complexes $\mathscr{C}^-(\mathscr{H}, \Lambda, \mathfrak{u})$ and $\mathscr{C}^-(\mathscr{H}_b, \Lambda, \mathfrak{u})$. The quasi-isomorphisms commute with the restriction maps relating these building blocks, so by putting them together, we obtain a quasi-isomorphism between the two surgery complexes.

The effect of a trajectory α-slide on the surgery complex of a complete system was studied in [Zemke 2021, Corollary 13.5]. In general, changing a trajectory c_w by concatenating it with a curve γ changes the surgery complex by a formula involving the action of the homology class $[\gamma] \in H_1(\Sigma; \mathbb{Z})$. Examining the formula for the action of γ in [ibid., Section 13.3], we see that this action vanishes identically on the span of the alpha curves. In our case, γ is a concatenation of the form $\tau * \alpha_j * \tau^{-1}$, where τ and α_j are as in Figure 25top left. In particular, γ is homologous to the curve α_j which bounds a disk in the alpha handlebody, so the action of γ on the surgery complex is trivial. Applying Zemke's result, it follows that the two surgery complexes are quasi-isomorphic.

Since Theorem 9.6 holds for the basic system $\mathscr{H}_b$, we deduce that it must also hold for $\mathscr{H}$. □

13 The general link surgery formula

This section contains the precise statement, as well as the proof, of the link surgery formula (Theorem 1.1). Up to now we have established the formula in the case of link-minimal complete systems. The case of general complete systems requires several new ingredients: explicit resolutions $\mathfrak{A}^-(\mathscr{H}, s)$ for the generalized Floer complexes $A^-(\mathscr{H}, s)$ (compare Section 3.6), as well as transition maps that allow one to define maps $\mathscr{I}_s^{\vec{M}}$ similar to the ones in Section 3.8; see Remark 3.25. Moreover, the proof of the general surgery formula is based on an analysis of the effect of an index zero/three link stabilization on a complete system, and uses the material from Section 7.

The new constructions are rather involved so, to help the reader, we build up the general statement by first considering some simpler cases. In Section 13.1 we describe the link surgery formula for a knot represented by a diagram with two w and two z basepoints. In Section 13.2 we study the simplest version of an index zero/three link stabilization, relating a knot diagram with one basepoint of each type to a diagram with two basepoints of each type. In Section 13.3 we explain how to define polygon maps between the resolutions $\mathfrak{A}^-(\mathcal{H}, \boldsymbol{s})$, again in the simplest nontrivial example, that of a knot with two w and two z basepoints. In Section 13.4 we describe the link surgery formula in more generality, still for a knot, but now with an arbitrary number of basepoints. In Section 13.5 we do yet another warm-up case, that of a link with two components, and with two basepoints of each type per component. Finally, in Section 13.6 we give the statement of the link surgery formula in full generality, and in Section 13.7 we prove it.

13.1 A knot with four basepoints

We consider the case of a knot $K \subset Y$, represented by a Heegaard diagram $\mathcal{H}^K$ with four linked basepoints w_1, z_1, w_2 and z_2. We suppose $\mathcal{H}^K$ is part of a complete system $\mathcal{H}$, that also contains a diagram $\mathcal{H}^\varnothing$ for Y, and two hyperboxes $\mathcal{H}^{K,K}$ and $\mathcal{H}^{K,-K}$; see Example 8.24.

In Section 3.2 we introduced the link Floer complex $\boldsymbol{CFL}^-(\mathcal{H}^K)$, which in our case is freely generated over $\mathcal{R} = \mathbb{F}[\![U_1, U_2]\!]$ by intersection points $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Let us consider the infinity version of this complex, obtained by inverting the U variables:

$$\boldsymbol{CFL}^\infty(\mathcal{H}^K) = (U_1, U_2)^{-1}\boldsymbol{CFL}^-(\mathcal{H}^K).$$

As a direct product, the complex $\boldsymbol{CFL}^\infty(\mathcal{H}^K)$ has generators (over $\mathbb{F}$) $U_1^{n_1}U_2^{n_2}\boldsymbol{x}$, with $n_i \in \mathbb{Z}$ and $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The differential takes into account intersection with w's, so we write $\boldsymbol{CFL}^\infty(\mathcal{H}^K) = C\{w\}$. On the other hand, $\boldsymbol{CFL}^-(\mathcal{H}^K)$ only has the generators that satisfy $n_1, n_2 \geq 0$, so we write

$$\boldsymbol{CFL}^-(\mathcal{H}^K) = C\{w; n_1, n_2 \geq 0\}.$$

For each $s \in \mathbb{Z}$, we also have a subcomplex

$$\boldsymbol{A}^-(\mathcal{H}^K, s) = C\{w; n_1, n_2 \geq 0,\ A(\boldsymbol{x}) - n_1 - n_2 \leq s\}.$$

For the surgery formula, we have an inclusion $\mathcal{I}_s^K$ of $\boldsymbol{A}^-(\mathcal{H}^K, s)$ into $\boldsymbol{CFL}^-(\mathcal{H}^K) = C\{w; n_1, n_2 \geq 0\}$; see Section 3.8. We can then use the Heegaard moves specified by the hyperbox $\mathcal{H}^{K,K}$ to relate $\boldsymbol{CFL}^-(\mathcal{H}^K)$ to $\boldsymbol{CF}^-(\mathcal{H}^\varnothing)$. The composition could play the role of the map Φ_s^K (just as in the case of link-minimal systems).

However, we also need a map Φ_s^{-K}, which would involve a chain map $\mathcal{I}_s^{-K}$ from $\boldsymbol{A}^-(\mathcal{H}^K, s)$ to $C\{z; n_1, n_2 \geq 0\}$, a complex with the differential counting z's. It does not seem possible to construct a suitable such map. To deal with this problem, we will replace $\boldsymbol{A}^-(\mathcal{H}^K, s)$ with an explicit resolution, $\mathfrak{A}^-(\mathcal{H}^K, s)$.

In the case of a knot with two basepoints, the map $\mathcal{I}_s^{-K}$ involves the inclusion of the complex $\boldsymbol{A}^-(\mathcal{H}^K, s) = C\{w; n \geq 0,\ A(\boldsymbol{x}) - n \leq s\}$ into $C\{w; A(\boldsymbol{x}) - n \leq s\}$. The latter complex is then identified with $C\{z; n \geq 0\}$ by using the isomorphism $\boldsymbol{x} \to U^{s-A(\boldsymbol{x})}\boldsymbol{x}$.

Similarly, for four basepoints, we have an inclusion of $\boldsymbol{A}^-(\mathcal{H}^K, s)$ into the intermediate complex

$$C_{\text{int}} = C\{w; n_2 \geq 0,\, A(\boldsymbol{x}) - n_1 - n_2 \leq s\}.$$

We will start by constructing a resolution of C_{int}, denoted by $\mathfrak{C}_{\text{int}}$, and then identify a subcomplex of $\mathfrak{C}_{\text{int}}$ as the resolution $\mathfrak{A}^-(\mathcal{H}^K, s)$ of $\boldsymbol{A}^-(\mathcal{H}^K, s)$. We will also construct a chain map (called a transition map) from $\mathfrak{C}_{\text{int}}$ to $C\{z; n_1, n_2 \geq 0\}$.

Consider the elements of C_{int}

$$\boldsymbol{x}_0 := U_1^{A(\boldsymbol{x})-s}\boldsymbol{x}, \quad \boldsymbol{x}_1 := U_1^{A(\boldsymbol{x})-s-1}U_2\boldsymbol{x}, \quad \boldsymbol{x}_2 := U_1^{A(\boldsymbol{x})-s-2}U_2^2\boldsymbol{x}, \quad \dots$$

over all possible $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Every element of C_{int} can be written as an infinite sum

$$\sum_{\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{i=1}^{\infty} U_1^{a_i} U_2^{b_i} \boldsymbol{x}_{n_i} = \sum_{\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{i=1}^{\infty} U_1^{A(\boldsymbol{x})-s-n_i+a_i} U_2^{b_i+n_i} \boldsymbol{x}, \tag{178}$$

where (a_i, b_i, n_i) is a sequence (depending on $\boldsymbol{x}$) consisting of distinct triples of nonnegative integers. Furthermore, in order to avoid infinitely many negative powers of U_1 in the sum, we need to impose the condition

$$(\exists K > 0)(\forall i) \quad n_i - a_i < K. \tag{179}$$

Observe also, that, if (179) is satisfied, then (178) is a well-defined sum. Indeed, because $a_i, n_i, b_i \geq 0$, we cannot have that the pair $(A(\boldsymbol{x}) - s - n_i + a_i, b_i + n_i)$ is some fixed (a, b) for infinitely many values of i.

Remark 13.1 Informally, we can think of the elements $\boldsymbol{x}_n$ as generators for C_{int}. If we had worked with polynomials in U_1 and U_2 instead of power series, then every element of the resulting C_{int} would have been a finite linear combination of the $\boldsymbol{x}_n$. With power series, however, we need to allow infinite sums satisfying (179).

The elements $\boldsymbol{x}_n$ are not free, as they satisfy relations of the form

$$U_2\boldsymbol{x}_n = U_1\boldsymbol{x}_{n+1}.$$

To get a resolution of C_{int}, we turn $\boldsymbol{x}_n$ into linearly independent elements, by introducing a new variable V_2, so that $\boldsymbol{x}_n$ corresponds to $V_2^n\boldsymbol{x}$. We also need to make sure $U_2V_2^n\boldsymbol{x} - U_1V_2^{n+1}\boldsymbol{x} = (U_2 + U_1V_2)V_2^n\boldsymbol{x}$ is in the image of the differential, so we introduce a further variable Y_2 so that

$$Y_2^2 = 0, \quad \partial Y_2 = U_2 + U_1V_2.$$

Furthermore, in order to determine the differential of some $\boldsymbol{x} \in \mathfrak{C}_{\text{int}}$, note that the differential of the corresponding element $U_1^{A(\boldsymbol{x})-s}\boldsymbol{x} \in C_{\text{int}}$ consists of terms of the form

$$U_1^{A(\boldsymbol{x})-s} \cdot U_1^{n_{w_1}(\phi)} U_2^{n_{w_2}(\phi)} \boldsymbol{y} = U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)} (U_1^{A(\boldsymbol{y})-s-n_{w_2}(\phi)} U_2^{n_{w_2}(\phi)} \boldsymbol{y}).$$

The last expression in parentheses corresponds to $V_2^{n_{w_2}(\phi)}\boldsymbol{y}$. Thus, $\partial\boldsymbol{x}$ in $\mathfrak{C}_{\mathrm{int}}$ should consist of terms of the form

$$U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)}V_2^{n_{w_2}(\phi)}\boldsymbol{y}.$$

Here is the precise definition of $\mathfrak{C}_{\mathrm{int}}$. We first consider the subring

$$\mathfrak{R}=\left\{\sum_{i=1}^{\infty}U_1^{a_i}U_2^{b_i}V_2^{n_i}\;\middle|\;n_i-a_i \text{ is bounded above}\right\}\subset\mathbb{F}[\![U_1,U_2,V_2]\!].$$

This is a module over $\mathbb{F}[\![U_1,U_2]\!]$, and contains the polynomial module $\mathbb{F}[\![U_1,U_2]\!][V_2]$. We will write

$$\mathfrak{R}=\mathbb{F}[\![U_1,U_2]\!][V_2]',$$

where the prime indicates that we allow infinite sums satisfying the boundedness condition (179).

We define $\mathfrak{C}^0_{\mathrm{int}}$ to be the free complex over $\mathfrak{R}$ with generators $\boldsymbol{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, and differential

$$\partial\boldsymbol{x}=\sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\;\sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}}\#(\mathcal{M}(\phi)/\mathbb{R})\cdot U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)}V_2^{n_{w_2}(\phi)}\boldsymbol{y}. \tag{180}$$

Lemma 13.2 *The map ∂ on $\mathfrak{C}^0_{\mathrm{int}}$ satisfies $\partial^2=0$.*

Proof The usual arguments in Floer theory apply, with ∂^2 counting the ends of a one-dimensional moduli space of pseudoholomorphic disks. One caveat is the presence of disk bubbles. According to the analysis in [Ozsváth and Szabó 2008a, Theorem 5.5], there are four such bubbles, one going over w_1 and z_1, one over z_1 and w_2, one over w_2 and z_2, and one over z_2 and w_1. Their contribution to $\partial^2\boldsymbol{x}$ is

$$U_1\boldsymbol{x}+U_1V_2\boldsymbol{x}+U_1V_2\boldsymbol{x}+U_1\boldsymbol{x}=0.$$

This shows that $\partial^2\boldsymbol{x}=0$. □

We now let $\mathfrak{C}_{\mathrm{int}}$ be the mapping cone

$$Y_2\cdot\mathfrak{C}^0_{\mathrm{int}}\xrightarrow{U_2+U_1V_2}\mathfrak{C}^0_{\mathrm{int}}. \tag{181}$$

Alternatively, we can simply define $\mathfrak{C}_{\mathrm{int}}$ starting from the free dg (differential graded) module over the dg algebra

$$\mathfrak{R}^Y=\mathbb{F}[\![U_1,U_2]\!][V_2,Y_2]'/(Y_2^2=0,\partial Y_2=U_2+U_1V_2),$$

with generators $\boldsymbol{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, by introducing the differential on generators as in (180). The homological grading on $\mathfrak{R}^Y$ is given by $\mathrm{gr}(U_i)=-2$, $\mathrm{gr}(V_2)=0$, $\mathrm{gr}(Y_2)=-1$. The prime again indicates the condition (179) in infinite sums.

However, most of the time we will view $\mathfrak{C}_{\mathrm{int}}$ as a complex over the ring $\mathcal{R}=\mathbb{F}[\![U_1,U_2]\!]$, rather than over $\mathfrak{R}$ or $\mathfrak{R}^Y$.

Define a chain map $P : \mathfrak{C}_{\mathrm{int}} \to C_{\mathrm{int}}$ by setting

$$P(V_2^n \boldsymbol{x}) = U_1^{A(\boldsymbol{x})-s-n} U_2^n \cdot \boldsymbol{x}, \quad P(V_2^n Y_2 \boldsymbol{x}) = 0, \tag{182}$$

and extending it linearly (including to the allowed infinite sums) over $\mathcal{R} = \mathbb{F}[\![U_1, U_2]\!]$.

Lemma 13.3 *P is a chain map.*

Proof In the expression $(\partial P + P\partial)(V_2^n \boldsymbol{x})$, for each holomorphic disk in a class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$, we obtain a term involving $\boldsymbol{y}$ with coefficient

$$U_1^{A(\boldsymbol{x})-s-n} U_2^n \cdot U_1^{n_{w_1}(\phi)} U_2^{n_{w_2}(\phi)} + U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)} \cdot U_1^{A(\boldsymbol{y})-s-n-n_{w_2}(\phi)} U_2^{n_{w_2(\phi)}+n}.$$

Since $A(\boldsymbol{x}) - A(\boldsymbol{y}) = n_{z_1}(\phi) + n_{z_2}(\phi) - n_{w_1}(\phi) - n_{w_2}(\phi)$, these terms cancel.

We also have

$$\begin{aligned} (\partial P + P\partial)(V_2^n Y_2 \boldsymbol{x}) &= P(\partial(V_2^n Y_2 \boldsymbol{x})) \\ &= P(U_2 V_2^n \boldsymbol{x} + U_1 V_2^{n+1} \boldsymbol{x} + V_2^n Y_2(\partial \boldsymbol{x})) \\ &= U_2 \cdot U_1^{A(\boldsymbol{x})-s-n} U_2^n \cdot \boldsymbol{x} + U_1 \cdot U_1^{A(\boldsymbol{x})-s-n-1} U_2^{n+1} \cdot \boldsymbol{x} + 0 \\ &= 0. \end{aligned}$$

This shows that P is a chain map. □

Lemma 13.4 *The map $P : \mathfrak{C}_{\mathrm{int}} \to C_{\mathrm{int}}$ exhibits $\mathfrak{C}_{\mathrm{int}}$ as a resolution of C_{int} over $\mathcal{R} = \mathbb{F}[\![U_1, U_2]\!]$; that is, P is a surjective quasi-isomorphism.*

Proof Clearly P is surjective; we need to show that it is also a quasi-isomorphism. Let G be the kernel of P, so that we have a short exact sequence

$$0 \to G \to \mathfrak{C}_{\mathrm{int}} \to C_{\mathrm{int}} \to 0.$$

By considering the long exact sequence in homology, we see that it suffices to prove that G is acyclic.

Note that G is freely generated over $\mathfrak{R} = \mathbb{F}[\![U_1, U_2]\!][V_2]'$ by the elements of the form $(U_2 + U_1 V_2)\boldsymbol{x}$ and $Y_2 \boldsymbol{x}$. The map $H : G \to G$ given by

$$H((U_2 + U_1 V_2)\boldsymbol{x}) = Y_2 \boldsymbol{x}, \quad H(Y_2 \boldsymbol{x}) = 0$$

is a chain homotopy between the identity and zero. Hence G is acyclic. □

Next, we define a filtration $\mathcal{F}$ on C_{int} by the negative of the exponent of U_1, ie by setting

$$\mathcal{F}(\boldsymbol{x}) = 0, \quad \mathcal{F}(U_1) = -1, \quad \mathcal{F}(U_2) = 0.$$

Note that $\boldsymbol{A}^-(\mathcal{H}^K, s)$ is the subcomplex of C_{int} given by $\mathcal{F} \leq 0$.

We also define a filtration $\mathfrak{F}$ on $\mathfrak{C}_{\mathrm{int}}$ by setting

$$\mathfrak{F}(\boldsymbol{x}) = -A(\boldsymbol{x}), \quad \mathfrak{F}(U_1) = -1, \quad \mathfrak{F}(U_2) = 0, \quad \mathfrak{F}(V_2) = 1, \quad \mathfrak{F}(Y_2) = 0. \tag{183}$$

Using (180) and (181), it is easy to check that the differential of $\mathfrak{C}_{\mathrm{int}}$ respects the filtration.

One can also verify that the map P from (182) is filtered of degree $-s$ with respect to $\mathfrak{F}$ and $\mathcal{F}$. Indeed,

$$\mathfrak{F}(V_2^n \boldsymbol{x}) = n - A(\boldsymbol{x}) = \mathcal{F}(U_1^{A(\boldsymbol{x})-s-n} U_2^n \boldsymbol{x}) - s.$$

We define the resolution $\mathfrak{A}^-(\mathcal{H}^K, s)$ of $\boldsymbol{A}^-(\mathcal{H}^K, s)$ to be the subcomplex of $\mathfrak{C}_{\mathrm{int}}$ corresponding to filtration level $\mathfrak{F} \le s$. Concretely, $\mathfrak{A}^-(\mathcal{H}^K, s)$ consists of infinite sums (with coefficients in $\mathcal{R}$, and a suitable boundedness condition) of the elements

$$\tilde{\boldsymbol{x}}_n^a := U_1^{\max(s+n-A(\boldsymbol{x}),0)} V_2^n Y_2^a \boldsymbol{x}, \qquad \boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta, \quad n \in \mathbb{Z}_{\ge 0}, \quad a \in \{0,1\}. \tag{184}$$

Let

$$\widetilde{P} : \mathfrak{A}^-(\mathcal{H}^K, s) \to \boldsymbol{A}^-(\mathcal{H}^K, s)$$

be the restriction of P from (182). In terms of the elements $\tilde{\boldsymbol{x}}_n^0$ and $\tilde{\boldsymbol{x}}_n^1$,

$$\widetilde{P}(\tilde{\boldsymbol{x}}_n^0) = U_1^{\max(A(\boldsymbol{x})-s-n,0)} U_2^n \cdot \boldsymbol{x}, \quad \widetilde{P}(\tilde{\boldsymbol{x}}_n^1) = 0. \tag{185}$$

Lemma 13.5 *The map $\widetilde{P} : \mathfrak{A}^-(\mathcal{H}^K, s) \to \boldsymbol{A}^-(\mathcal{H}^K, s)$ exhibits $\mathfrak{A}^-(\mathcal{H}^K, s)$ as a resolution of $\boldsymbol{A}^-(\mathcal{H}^K, s)$.*

Proof This is a filtered version of Lemma 13.4. Note that the kernel $\widetilde{G} = \mathrm{Ker}(\widetilde{P})$ is the subgroup of $G = \mathrm{Ker}(P)$ in filtration level $\mathfrak{F} \le s$. Moreover, the null-homotopy $H : G \to G$ from the proof of Lemma 13.4 is a filtered map; hence, by restricting it to $\widetilde{G}$, we get that $\widetilde{G}$ is acyclic. □

We now construct a *transition map*

$$\Xi : \mathfrak{C}_{\mathrm{int}} \to C\{z; n_1, n_2 \ge 0\}$$

by setting

$$\Xi(V_2^n Y_2 \boldsymbol{x}) = \boldsymbol{x} \tag{186}$$

and

$$\Xi(V_2^n \boldsymbol{x}) = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x},\boldsymbol{y}) \\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{n_{z_1}(\phi)} \frac{U_1^{n_{z_2}(\phi)} - U_2^{n_{z_2}(\phi)}}{U_1 - U_2} \boldsymbol{y}. \tag{187}$$

This definition is inspired from [Sarkar 2015]. Of course, by $(U_1^{n_{z_2}(\phi)} - U_2^{n_{z_2}(\phi)})/(U_1 - U_2)$ we actually mean the polynomial

$$U_1^{n_{z_2}(\phi)-1} + U_1^{n_{z_2}(\phi)-2} U_2 + \cdots + U_1 U_2^{n_{z_2}(\phi)-2} + U_2^{n_{z_2}(\phi)-1}.$$

We extend Ξ to be a module map over $\mathcal{R} = \mathbb{F}[\![U_1, U_2]\!]$.

Lemma 13.6 Ξ *is a chain map.*

Proof Note that Ξ is a module map over $\mathfrak{R} = \mathbb{F}[\![U_1, U_2]\!][V_2]'$, where V_2 acts on the target by the identity. Thus, it suffices to check that $\Xi\partial + \partial\Xi = 0$ on generators of the form $\boldsymbol{x}$ and $Y_2\boldsymbol{x}$.

Let us first check that
$$\Xi\partial(\boldsymbol{x}) + \partial\Xi(\boldsymbol{x}) = 0.$$
In the expression $\Xi\partial(\boldsymbol{x})$, a pair of holomorphic disks in classes $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{a})$ and $\psi \in \pi_2(\boldsymbol{a}, \boldsymbol{y})$ contributes
$$U_1^{n_{z_1}(\psi)} \frac{U_1^{n_{z_2}(\psi)} - U_2^{n_{z_2}(\psi)}}{U_1 - U_2} \cdot U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)} \boldsymbol{y}.$$
In $\partial\Xi(\boldsymbol{x})$, the same pair contributes
$$U_1^{n_{z_1}(\psi)} U_2^{n_{z_2}(\psi)} \cdot U_1^{n_{z_1}(\phi)} \frac{U_1^{n_{z_2}(\phi)} - U_2^{n_{z_2}(\phi)}}{U_1 - U_2} \boldsymbol{y}.$$
Adding these together, we get a total contribution of
$$U_1^{n_{z_1}(\phi*\psi)} \frac{U_1^{n_{z_2}(\phi*\psi)} - U_2^{n_{z_2}(\phi*\psi)}}{U_1 - U_2} \boldsymbol{y}. \tag{188}$$
This only depends on the combined domain $\phi * \psi$. The usual proof that $\partial^2 = 0$ applies here, and shows that adding up all these contributions gives 0. Note that when $\boldsymbol{x} = \boldsymbol{y}$, there are four periodic domains to take care of: one contains w_1 and z_1, one w_1 and z_2, one w_2 and z_1, and the last w_2 and z_2. Their contributions add up to zero, because the w's do not play a role in the expression (188).

Now let us check that
$$\Xi\partial(Y_2\boldsymbol{x}) + \partial\Xi(Y_2\boldsymbol{x}) = 0.$$
We have
$$\begin{aligned}
\Xi\partial(Y_2\boldsymbol{x}) &= \Xi\Bigg(U_2\boldsymbol{x} + U_1V_2\boldsymbol{x} + \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)} V_2^{n_{w_2}(\phi)} Y_2\boldsymbol{y}\Bigg) \\
&= \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \Bigg((U_2+U_1) \cdot U_1^{n_{z_1}(\phi)} \frac{U_1^{n_{z_2}(\phi)} - U_2^{n_{z_2}(\phi)}}{U_1 - U_2} + U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)}\Bigg) \boldsymbol{y} \\
&= \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) U_1^{n_{z_1}(\phi)} U_2^{n_{z_2}(\phi)} \boldsymbol{y}.
\end{aligned}$$
On the other hand, we also have
$$\partial\Xi(Y_2\boldsymbol{x}) = \partial\boldsymbol{x} = \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) U_1^{n_{z_1}(\phi)} U_2^{n_{z_2}(\phi)} \boldsymbol{y}. \qquad \square$$

We now have all the ingredients to write down the surgery formula for K (with four basepoints). This takes the form of a mapping cone, where the initial complex is the direct product of $\mathfrak{A}^-(\mathcal{H}^K, s)$ over $s \in \mathbb{Z}$, and the final complex is the direct product of copies of some $\boldsymbol{CF}^-(\mathcal{H}^\varnothing) = \boldsymbol{CF}^-(S^3)$.

The map Φ_s^K is defined to be the composition

$$\mathfrak{A}^-(\mathcal{H}^K, s) \xrightarrow{\widetilde{P}} \boldsymbol{A}^-(\mathcal{H}^K, s) \hookrightarrow C\{w; n_1, n_2 \geq 0\} \xrightarrow{\sim} \boldsymbol{CF}^-(\mathcal{H}^\varnothing), \tag{189}$$

with the last map being an equivalence obtained from the Heegaard moves in $\mathcal{H}^{K,K}$. In fact, we denote the composition of the first two maps in (189) by

$$\mathcal{I}_s^K : \mathfrak{A}^-(\mathcal{H}^K, s) \to C\{w; n_1, n_2 \geq 0\} = \boldsymbol{CFL}^-(\mathcal{H}^K) = \boldsymbol{CF}^-(p^K(\mathcal{H}^K)).$$

This plays the role of the inclusion map from Section 3.8. Since in our case it also includes the projection $\widetilde{P}$, we call $\mathcal{I}_s^K$ a *projection-inclusion map*.

The final equivalence in (189), induced by Heegaard moves, is a descent map

$$D^K : \boldsymbol{CF}^-(p^K(\mathcal{H}^K)) \to \boldsymbol{CF}^-(\mathcal{H}^\varnothing),$$

similar to the ones considered in Section 9.1. Thus, we can write

$$\Phi_s^K = D^K \circ \mathcal{I}_s^K,$$

just as before.

On the other hand, the map Φ_s^{-K} is defined to be the composition

$$\mathfrak{A}^-(\mathcal{H}^K, s) \hookrightarrow \mathfrak{C}_{\mathrm{int}} \xrightarrow{\Xi} C\{z; n_1, n_2 \geq 0\} \xrightarrow{\sim} \boldsymbol{CF}^-(\mathcal{H}^\varnothing), \tag{190}$$

with the last map again obtained from Heegaard moves, this time the ones from the hyperbox $\mathcal{H}^{K,-K}$. In this context, we let $\mathcal{I}_s^{-K}$ be the first inclusion above (into $\mathfrak{C}_{\mathrm{int}}$), and we define the descent map D^{-K} to be the composition of Ξ and the equivalences induced by Heegaard moves. We can still write $\Phi_s^{-K} = D^{-K} \circ \mathcal{I}_s^{-K}$.

Just as in the link-minimal case, by summing up the Φ_s^K and Φ_s^{-K} maps (with suitable shifts in s), we obtain the surgery complexes $\mathcal{C}^-(\mathcal{H}, n)$ for $n \in \mathbb{Z}$.

13.2 Invariance under an index zero/three link stabilization

As a warm-up to the general case of invariance, in this section we will study the effect of an index zero/three link stabilization on a minimally pointed complete system of hyperboxes for a knot $K \subset Y$.

Let $\overline{\mathcal{H}}$ denote such a complete system. Recall from Example 8.24 that $\overline{\mathcal{H}}$ consists of

- a Heegaard diagram $\overline{\mathcal{H}}^K$ for K with two basepoints w_1 and z_1;
- a Heegaard diagram $\overline{\mathcal{H}}^\varnothing$ for Y with the single basepoint w_1;
- a hyperbox (a sequence of Heegaard moves with Θ elements) $\overline{\mathcal{H}}^{K,K}$ from $p^K(\overline{\mathcal{H}}^K)$ to $\overline{\mathcal{H}}^\varnothing$;
- another hyperbox, $\overline{\mathcal{H}}^{K,-K}$, relating $p^{-K}(\overline{\mathcal{H}}^K)$ to $\overline{\mathcal{H}}^\varnothing$.

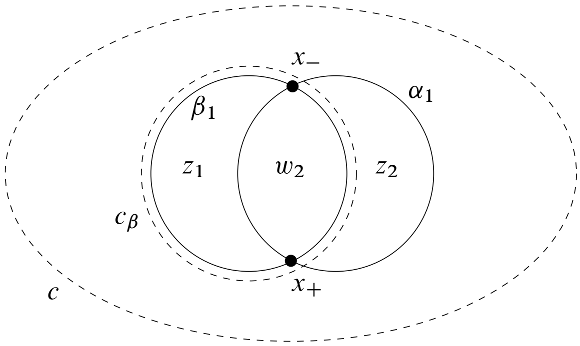

Figure 39: An index zero/three link stabilization.

From $\overline{\mathcal{H}}$ we can construct a surgery complex $\mathcal{C}^-(\overline{\mathcal{H}}, n)$ for each $n \in \mathbb{Z}$. This is a free complex over $\mathbb{F}[\![U_1]\!]$.

We now do a stabilization as in Figure 39, by introducing two new curves α_1 and β_1, and two new basepoints w_2 and z_2 in the neighborhood of the existing z_1. We let $\mathcal{H}$ be the complete system obtained from $\overline{\mathcal{H}}$ by this stabilization; see Section 8.9.

Thus, in the complete system $\mathcal{H}$, we have

- the diagram $\mathcal{H}^K$ obtained from $\overline{\mathcal{H}}^K$ by the link stabilization;
- the diagram $\mathcal{H}^\varnothing$ obtained from $\overline{\mathcal{H}}^\varnothing$ by a free index zero/three stabilization introducing α_1, β_1 and w_2;
- a hyperbox $\mathcal{H}^{K,K}$ from $p^K(\mathcal{H}^K)$ to $\mathcal{H}^\varnothing$ obtained from $\overline{\mathcal{H}}^{K,K}$ by doing free index zero/three stabilizations at each step;
- a hyperbox $\mathcal{H}^{K,-K}$ from $p^{-K}(\mathcal{H}^K)$ to $\mathcal{H}^\varnothing$ obtained $\overline{\mathcal{H}}^{K,-K}$ by doing free index zero/three stabilizations at each step, and then the move shown in Figure 18.

From $\mathcal{H}$ we can form a surgery complex $\mathcal{C}^-(\mathcal{H}, n)$ as explained in Section 13.1. This is a free complex over $\mathbb{F}[\![U_1, U_2]\!]$.

In Section 7.2 we described the behavior of holomorphic disks under an index zero/three link stabilization; see Proposition 7.2. For our calculations, we will choose almost complex structures as in that proposition, that is, with sufficiently stretched necks along the curves c and c_β from Figure 18.

Proposition 13.7 *For every $n \in \mathbb{Z}$, the surgery complexes $\mathcal{C}^-(\overline{\mathcal{H}}, n)$ and $\mathcal{C}^-(\mathcal{H}, n)$ are quasi-isomorphic over $\mathbb{F}[\![U_1]\!]$.*

Proof Let $\overline{C}\{w; n \geq 0\} = \boldsymbol{CF}^-(p^K(\overline{\mathcal{H}}^K))$ and $\overline{C}\{z; n \geq 0\} = \boldsymbol{CF}^-(p^{-K}(\overline{\mathcal{H}}^K))$ be the complexes obtained from the diagram $\overline{\mathcal{H}}^K$ by using the w_1 and the z_1 basepoint, respectively. Then, as in the proof of Theorem 3.16 or in the discussion from Section 7.2, we can identify $C\{w; n_1, n_2 \geq 0\} = \boldsymbol{CF}^-(p^K(\mathcal{H}^K))$ with the mapping cone

$$\overline{C}\{w; n \geq 0\}_-[\![U_2]\!] \xrightarrow{U_1 - U_2} \overline{C}\{w; n \geq 0\}_+[\![U_2]\!], \tag{191}$$

where the subscripts $-$ and $+$ indicate the point in $\alpha_1 \cap \beta_1$ that is part of the respective generator (x_- or x_+). Hence, by Corollary 7.3, we have a chain homotopy equivalence (and hence a quasi-isomorphism)

$$\rho\colon C\{w; n_1, n_2 \geq 0\} \to \overline{C}\{w; n \geq 0\},$$

acting on the domain of (191) by zero, and on the target by $U_1^{n_1} U_2^{n_2} \boldsymbol{x} \to U_1^{n_1+n_2} \boldsymbol{x}$.

By a slight abuse of notation, we will denote by ρ all quasi-isomorphisms of a similar form. For example, there are such equivalences ρ relating the complexes $\boldsymbol{CF}^-$ for all the diagrams in the hyperbox $\mathscr{H}^{K,K}$ to the corresponding ones in $\overline{\mathscr{H}}^{K,K}$. Furthermore, according to Proposition 6.21, these equivalences commute with the descent maps induced by Heegaard moves. Therefore, we have a commutative diagram

$$\begin{tikzcd} C\{w; n_1, n_2 \geq 0\} = \boldsymbol{CF}^-(p^K(\mathscr{H}^K)) \arrow[r, "D^K"] \arrow[d, "\rho"'] & \boldsymbol{CF}^-(\mathscr{H}^{\varnothing}) \arrow[d, "\rho"] \\ \overline{C}\{w; n \geq 0\} = \boldsymbol{CF}^-(p^K(\overline{\mathscr{H}}^K)) \arrow[r, "D^K"] & \boldsymbol{CF}^-(\overline{\mathscr{H}}^{\varnothing}) \end{tikzcd} \tag{192}$$

Theorem 3.16(b) gives yet another equivalence

$$\boldsymbol{A}^-(\mathscr{H}^K, s) \xrightarrow{\sim} \boldsymbol{A}^-(\overline{\mathscr{H}}^K, s).$$

We precompose this with the map $\widetilde{P}$ from (185), which is a quasi-isomorphism according to Lemma 13.5. The result is a quasi-isomorphism

$$\rho\colon \mathfrak{A}^-(\mathscr{H}^K, s) \to \boldsymbol{A}^-(\overline{\mathscr{H}}^K, s) = \mathfrak{A}^-(\overline{\mathscr{H}}^K, s). \tag{193}$$

By construction, we get a commutative diagram

$$\begin{tikzcd} \mathfrak{A}^-(\mathscr{H}^K, s) \arrow[r, "\mathscr{I}_s^K"] \arrow[d, "\rho"'] & \boldsymbol{CF}^-(p^K(\mathscr{H}^K)) \arrow[d, "\rho"] \\ \mathfrak{A}^-(\overline{\mathscr{H}}^K, s) \arrow[r, "\mathscr{I}_s^K"] & \boldsymbol{CF}^-(p^K(\overline{\mathscr{H}}^K)) \end{tikzcd}$$

Combining this with (192), we obtain a commutative diagram

$$\begin{tikzcd} \mathfrak{A}^-(\mathscr{H}^K, s) \arrow[r, "\Phi^K"] \arrow[d, "\rho"'] & \boldsymbol{CF}^-(\mathscr{H}^{\varnothing}) \arrow[d, "\rho"] \\ \mathfrak{A}^-(\overline{\mathscr{H}}^K, s) \arrow[r, "\Phi^K"] & \boldsymbol{CF}^-(\overline{\mathscr{H}}^{\varnothing}) \end{tikzcd} \tag{194}$$

We aim to construct a similar diagram involving Φ^{-K} instead of Φ^K:

$$\begin{tikzcd} \mathfrak{A}^-(\mathscr{H}^K, s) \arrow[r, "\Phi^{-K}"] \arrow[d, "\rho"'] \arrow[dr] & \boldsymbol{CF}^-(\mathscr{H}^{\varnothing}) \arrow[d, "\rho"] \\ \mathfrak{A}^-(\overline{\mathscr{H}}^K, s) \arrow[r, "\Phi^{-K}"] & \boldsymbol{CF}^-(\overline{\mathscr{H}}^{\varnothing}) \end{tikzcd} \tag{195}$$

This time, the maps ρ will not commute with Φ^{-K} on the nose, but only up to chain homotopy. Still, by introducing the chain homotopy as a diagonal map in the diagram above, we can form a chain map from the cone of the first row to the cone of the second row. This map is a quasi-isomorphism because the vertical maps are.

Recall that the Φ^K and Φ^{-K} maps are the building blocks in the surgery complex. Thus, once we have (194) and (195), we can form a quasi-isomorphism between the surgery complexes $\mathcal{C}^-(\mathcal{H}, n)$ and $\mathcal{C}^-(\overline{\mathcal{H}}, n)$.

We are left to construct (195). By Proposition 7.2, we can identify $C\{z; n_1, n_2 \geq 0\}$ with

$$\overline{C}\{z; n \geq 0\}_+^{U_1 \to U_2}[\![U_1]\!] \xrightarrow{U_1 - U_2} \overline{C}\{z; n \geq 0\}_-^{U_1 \to U_2}[\![U_1]\!], \tag{196}$$

and hence, by Corollary 7.4, we have an equivalence

$$\rho\colon C\{z; n_1, n_2 \geq 0\} \to \overline{C}\{z; n \geq 0\}.$$

We can do the same with the complexes appearing in the hyperboxes $\mathcal{H}^{K,-K}$ versus the ones in $\overline{\mathcal{H}}^{K,-K}$. We obtain a diagram, which commutes up to some diagonal chain homotopy,

$$\begin{tikzcd} C\{z; n_1, n_2 \geq 0\} = \boldsymbol{CF}^-(p^{-K}(\mathcal{H}^K)) \arrow[r, "\sim"] \arrow[d, "\rho"'] \arrow[dr] & \boldsymbol{CF}^-(\mathcal{H}^\varnothing) \arrow[d, "\rho"] \\ \overline{C}\{z; n \geq 0\} = \boldsymbol{CF}^-(p^{-K}(\overline{\mathcal{H}}^K)) \arrow[r, "\sim"] & \boldsymbol{CF}^-(\overline{\mathcal{H}}^\varnothing) \end{tikzcd} \tag{197}$$

where the horizontal maps are equivalences induced by the Heegaard moves in the respective hyperboxes. The fact that these maps commute with the projections ρ, up to chain homotopy, follows from Propositions 7.14 and 7.21. (The latter proposition applies for the final map in the stabilized hyperbox, the one shown in Figures 18 and 24.)

We claim that there is also a diagram

$$\begin{tikzcd} \mathfrak{A}^-(\mathcal{H}^K, s) \arrow[r, "\mathcal{I}_s^{-K}"] \arrow[d, "\rho"'] & \mathfrak{C}_{\text{int}} \arrow[r, "\Xi"] \arrow[d, "\rho"] \arrow[dr, "Z"] & C\{z; n_1, n_2 \geq 0\} \arrow[d, "\rho"] \\ \mathfrak{A}^-(\overline{\mathcal{H}}^K, s) \arrow[r, "\mathcal{I}_s^{-K}"] & \overline{C}\{z; n \geq 0\} \arrow[r, "\text{id}"] & \overline{C}\{z; n \geq 0\} \end{tikzcd} \tag{198}$$

where the left square in the diagram commutes, and the right square commutes up to the chain homotopy defined by Z; that is,

$$\rho + \rho \circ \Xi = \partial \circ Z + Z \circ \partial. \tag{199}$$

In the diagram (198), the two maps that we still have to define are the middle vertical map

$$\rho\colon \mathfrak{C}_{\text{int}} \to \overline{C}\{z; n \geq 0\} \tag{200}$$

and the chain homotopy Z.

We first note that, in view of (180) and the description of the holomorphic disks in Proposition 7.2, we can identify $\mathfrak{C}^0_{\mathrm{int}}$ with the cone

$$\overline{C}\{z; n \geq 0\}_-^{U_1 \to U_2}[\![U_1]\!][V_2]' \xrightarrow{V_2 - 1} \overline{C}\{z; n \geq 0\}_+^{U_1 \to U_2}[\![U_1]\!][V_2]'. \tag{201}$$

The same argument as in Lemma 3.14 shows that $\mathfrak{C}^0_{\mathrm{int}}$ is chain homotopy equivalent to

$$C\{z; n \geq 0\}^{U_1 \to U_2}[\![U_1]\!]$$

over $\mathbb{F}[\![U_1, U_2]\!]$, via the projection that takes

$$\boldsymbol{x} \times x_+ \to \boldsymbol{x}, \quad \boldsymbol{x} \times x_- \to 0, \quad V_2 \to 1 \tag{202}$$

for $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, where $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ are the totally real tori coming from the diagram $\overline{\mathcal{H}}^K$.

Furthermore, $\mathfrak{C}_{\mathrm{int}}$ is obtained from $\mathfrak{C}^0_{\mathrm{int}}$ by taking the cone of $U_2 + U_1 V_2$. Consequently, using (202), $\mathfrak{C}_{\mathrm{int}}$ is chain homotopy equivalent to the cone of $U_2 + U_1$ on $C\{z; n \geq 0\}^{U_1 \to U_2}[\![U_1]\!]$. In turn, this is chain homotopy equivalent to $C\{z; n \geq 0\}$ over $\mathbb{F}[\![U_1]\!]$, via the usual map taking the domain of the cone to 0, and projecting the target with $U_2 \mapsto U_1$; see Lemma 7.1.

The combined homotopy equivalence from $\mathfrak{C}_{\mathrm{int}}$ to $C\{z; n \geq 0\}$ is the desired map (200). We can write it as follows. The generators of $\mathfrak{C}_{\mathrm{int}}$ over $\mathbb{F}[\![U_1]\!]$ are

$$U_2^m V_2^n Y_2^a (\boldsymbol{x} \times r), \qquad m, n \geq 0, \quad a \in \{0, 1\}, \quad \boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta, \quad r \in \{x_+, x_-\}.$$

Then

$$\rho(U_2^m V_2^n Y_2^a (\boldsymbol{x} \times r)) = \begin{cases} U_1^m \boldsymbol{x} & \text{if } a = 0,\ r = x_+, \\ 0 & \text{otherwise.} \end{cases} \tag{203}$$

As for the homotopy Z from (198), we let it be the $\mathbb{F}[\![U_1]\!]$-module map defined on generators by

$$Z(U_2^m V_2^n Y_2^a (\boldsymbol{x} \times r)) = \begin{cases} m U_1^{m-1} \boldsymbol{x} & \text{if } a = 0,\ r = x_-, \\ 0 & \text{otherwise.} \end{cases} \tag{204}$$

As an aside, observe that the Alexander grading of $\boldsymbol{x} \times r$ is

$$A(\boldsymbol{x} \times r) = \begin{cases} A(\boldsymbol{x}) & \text{if } r = x_+, \\ A(\boldsymbol{x}) - 1 & \text{if } r = x_-. \end{cases} \tag{205}$$

Let us now check that the left square in (198) commutes. In preparation for that, note that the generators of $\mathfrak{A}^-(\mathcal{H}^K, s)$ over $\mathbb{F}[\![U_1]\!]$ can be written

$$U_2^m (\widetilde{\boldsymbol{x} \times r})_n^a, \qquad \boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta, \quad r \in \{x_+, x_-\}, \quad n \in \mathbb{Z}_{\geq 0}, \quad a \in \{0, 1\},$$

using the tilde notation from (184).

The left vertical map $\rho\colon \mathfrak{A}^-(\mathcal{H}^K, s) \to \mathfrak{A}^-(\overline{\mathcal{H}}^K, s)$ in (198) is constructed in (193), by composing $\widetilde{P}$ from (185) with a projection taking $U_2^m(\boldsymbol{x} \times r)$ to $U_1^m \boldsymbol{x}$ if $r = x_+$ and to zero if $r = x_-$. Therefore,

$$\rho(U_2^m (\widetilde{\boldsymbol{x} \times r})_n^a) = \begin{cases} U_1^{\max(A(\boldsymbol{x}) - s - n, 0) + m + n} \boldsymbol{x} & \text{if } a = 0,\ r = x_+, \\ 0 & \text{otherwise.} \end{cases}$$

The first map in the bottom row in (198) is the inclusion $\mathcal{I}_s^{-K}$ defined in (24) by $\boldsymbol{x} \mapsto U_1^{\max(s-A(\boldsymbol{x}),0)}\boldsymbol{x}$. However, that formula was based on the alternative description of $\mathfrak{A}^-(\overline{\mathcal{H}}^K)$ from Section 3.6, whereas here we use its description as a subcomplex of $\overline{C}\{w; n \geq 0\}$. The translation between the two is given by the inverse of (17), namely $\boldsymbol{x} \mapsto U_1^{-\max(A(\boldsymbol{x})-s,0)}\boldsymbol{x}$. Composing these two formulas we obtain

$$\mathcal{I}_s^{-K}(\boldsymbol{x}) = U_1^{s-A(\boldsymbol{x})}\boldsymbol{x}.$$

Therefore,

$$(\mathcal{I}_s^{-K} \circ \rho)(U_2^m (\widetilde{\boldsymbol{x} \times r})_n^a) = \begin{cases} U_1^{\max(s-A(\boldsymbol{x})+n,0)+m}\boldsymbol{x} & \text{if } a = 0,\ r = x_+, \\ 0 & \text{otherwise.} \end{cases} \tag{206}$$

With regard to the first map in the top row of (198), by (184), we have

$$\mathcal{I}_s^{-K}(U_2^m (\widetilde{\boldsymbol{x} \times r})_n^a) = U_1^{\max(s+n-A(\boldsymbol{x}\times r),0)} U_2^m V_2^n Y_2^a (\boldsymbol{x} \times r). \tag{207}$$

Using (203) and (205), $\mathcal{I}_s^{-K} \circ \rho$ is given by the same formula as (206), so the left square in (198) commutes. In fact, the middle vertical map $\rho\colon \mathfrak{C}_{\text{int}} \to C\{z; n \geq 0\}$, given by(203), is filtered (with respect to the filtration $\mathfrak{F}$ from (183) on the domain and the Alexander filtration on the target). Its restriction to filtration levels $\leq s$ is exactly the left vertical ρ map in (198).

Next, we check that the right square in (198) commutes up to the homotopy Z, ie that (199) is satisfied.

To compute the top horizontal map Ξ in this square, we use (186) and (187), as well as the description of holomorphic disks in $\mathcal{H}^K$ from Proposition 7.2. With our choice of almost complex structure, the holomorphic disks in $\mathcal{H}^K$ in classes ϕ with $n_{z_2}(\phi) \neq 0$ are

- the disks from $\boldsymbol{x} \times x_+$ to $\boldsymbol{x} \times x_-$, which contain z_2 once and no other basepoints, and
- disks from $\boldsymbol{x} \times x_\pm$ to $\boldsymbol{y} \times x_\pm$, which corresponds to disks in the destabilized diagram from $\boldsymbol{x}$ to $\boldsymbol{y}$, and go over z_2 as many times as they went over z_1 in the destabilized diagram.

It is helpful to introduce the following notation. If the differential on $\overline{C}\{z; n \geq 0\}$ is given by

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y}) U_1^{n(\boldsymbol{x},\boldsymbol{y})} \boldsymbol{y}, \tag{208}$$

we set

$$\partial^\Xi \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y}) \frac{U_1^{n(\boldsymbol{x},\boldsymbol{y})} - U_2^{n(\boldsymbol{x},\boldsymbol{y})}}{U_1 - U_2} \boldsymbol{y},$$

and

$$\partial' \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y}) n(\boldsymbol{x}, \boldsymbol{y}) U_1^{n(\boldsymbol{x},\boldsymbol{y})-1} \boldsymbol{y}.$$

Then

$$\Xi(U_2^m V_2^n Y_2^a (\boldsymbol{x} \times r)) = \begin{cases} U_2^m (\boldsymbol{x} \times r) & \text{if } a = 1, \\ U_2^m (\boldsymbol{x} \times x_- + (\partial^\Xi \boldsymbol{x}) \times x_+) & \text{if } a = 0,\ r = x_+, \\ (\partial^\Xi \boldsymbol{x}) \times x_- & \text{if } a = 0,\ r = x_-. \end{cases}$$

Moreover, in view of the identification (196), the rightmost vertical map in (198) is given by

$$\rho(U_2^m(\boldsymbol{x} \times r)) = \begin{cases} U_1^m \boldsymbol{x} & \text{if } r = x_-, \\ 0 & \text{if } r = x_+. \end{cases}$$

Therefore,

$$(\rho \circ \Xi)(U_2^m V_2^n Y_2^a(\boldsymbol{x} \times r)) = \begin{cases} U_1^m \boldsymbol{x} & \text{if } a = 1,\ r = x_- \text{ or } a = 0,\ r = x_+ \\ U_1^m(\partial' \boldsymbol{x}) & \text{if } a = 0,\ r = x_-, \\ 0 & \text{if } a = 1,\ r = x_+. \end{cases} \tag{209}$$

Combining this with (203), we get a formula for the left-hand side of (199):

$$U_2^m V_2^n Y_2^a(\boldsymbol{x} \times r) \mapsto \begin{cases} U_1^m \boldsymbol{x} & \text{if } a = 1,\ r = x_-, \\ U_1^m(\partial' \boldsymbol{x}) & \text{if } a = 0,\ r = x_-, \\ 0 & \text{if } r = x_+. \end{cases} \tag{210}$$

Let us compare this with the right-hand side of (199), which is the commutator $\partial \circ Z + Z \circ \partial$. Recall that Z is given by (204). Note that, since $\mathfrak{C}_{\text{int}}$ is the mapping cone of $U_2 + U_1 V_2$ on $\mathfrak{C}^0_{\text{int}}$, the differential ∂ on $\mathfrak{C}_{\text{int}}$ consists of two parts: one coming from $\mathfrak{C}^0_{\text{int}}$, and one taking Y_2 to $U_2 + U_1 V_2$. Further, in view of the identification (201), the first part consist of a term coming from the differential on $\overline{C}\{z; n \geq 0\}^{U_1 \to U_2}$, and a term taking $\boldsymbol{x} \times x_-$ to $(V_2 - 1)\boldsymbol{x} \times x_+$. With this in mind, we can compute the effect of $\partial \circ Z + Z \circ \partial$ on a generator $U_2^m V_2^n Y_2^a(\boldsymbol{x} \times r)$.

When $r = x_+$, note that ∂ cannot take a generator containing x_+ into one containing x_-. Since Z kills off all generators of the form $\boldsymbol{y} \times x_+$,

$$(\partial \circ Z + Z \circ \partial)(U_2^m V_2^n Y_2^a(\boldsymbol{x} \times x_+)) = 0,$$

in agreement with (210).

When $r = x_-$ and $a = 0$, if we write $\partial \boldsymbol{x} \in \overline{C}\{z; n \geq 0\}$ as in (208), then by using (201), we obtain

$$\partial(U_2^m V_2^n(\boldsymbol{x} \times x_-)) = (V_2 - 1) U_2^m V_2^n(\boldsymbol{x} \times x_+) + \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y}) U_2^{m + n(\boldsymbol{x}, \boldsymbol{y})} V_2^n(\boldsymbol{y} \times x_-)$$

and hence

$$(Z \circ \partial)(U_2^m V_2^n(\boldsymbol{x} \times x_-)) = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y})(m + n(\boldsymbol{x}, \boldsymbol{y})) U_1^{m + n(\boldsymbol{x}, \boldsymbol{y}) - 1} \boldsymbol{y}.$$

On the other hand,

$$(\partial \circ Z)(U_2^m V_2^n(\boldsymbol{x} \times x_-)) = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y}) m U_1^{m + n(\boldsymbol{x}, \boldsymbol{y}) - 1} \boldsymbol{y}.$$

From here, we get

$$(\partial \circ Z + Z \circ \partial)(U_2^m V_2^n(\boldsymbol{x} \times x_+)) = U_1^m(\partial' \boldsymbol{x}),$$

as we expected from (210).

Finally, when $r = x_-$ and $a = 1$,

$$\partial(U_2^m V_2^n Y_2(\boldsymbol{x} \times x_-))$$
$$= (U_2 + U_1 V_2) U_2^m V_2^n (\boldsymbol{x} \times x_-) + (V_2 - 1) U_2^m V_2^n Y_2(\boldsymbol{x} \times x_+) + \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} c(\boldsymbol{x}, \boldsymbol{y}) U_2^{m+n(\boldsymbol{x},\boldsymbol{y})} V_2^n Y_2(\boldsymbol{y} \times x_-),$$

so

$$(Z \circ \partial)(U_2^m V_2^n Y_2(\boldsymbol{x} \times x_-)) = ((m+1) + m) U_1^m \boldsymbol{x} + 0 + 0 = U_1^m \boldsymbol{x},$$

since we work modulo two. We also have

$$(\partial \circ Z)(U_2^m V_2^n Y_2(\boldsymbol{x} \times x_-)) = 0.$$

Therefore,

$$(\partial \circ Z + Z \circ \partial)(U_2^m V_2^n Y_2(\boldsymbol{x} \times x_+)) = U_1^m \boldsymbol{x},$$

again in agreement with (210).

This proves (199), ie that the right square in (198) commutes up to the given chain homotopy.

Putting the diagrams (197) and (198) side by side, we obtain (195), and the proof is complete. □

13.3 Maps between resolutions

Recall that one essential ingredient in the link surgery formula are the hyperboxes of strongly equivalent Heegaard diagrams. These are defined in Section 8.2, for arbitrary (not necessarily link-minimal) diagrams. The construction uses Θ-chain elements in generalized Floer complexes of the form $\mathfrak{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \mathbf{0}) \cong \boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \mathbf{0})$ or $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\alpha'}, \mathbf{0}) \cong \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\alpha'}, \mathbf{0})$, associated to link-minimal diagrams of an unlink in a connected sum of $(S^1 \times S^2)$'s. Furthermore, in Section 8.3, we showed that a hyperbox of strongly equivalent Heegaard diagrams produces one of generalized Floer complexes of the form $\boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s})$; this used the polygon maps between $\boldsymbol{A}^-$ complexes constructed in Section 3.5.

For the link surgery formula with diagrams that are not link-minimal, we will also need to see how the same Θ-chain elements produce hyperboxes made of the *resolutions* $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \boldsymbol{s})$ of generalized Floer complexes. To do so, it suffices to construct polygon maps between these resolutions. We will explain this here in the simplest nontrivial case, for two strongly equivalent diagrams representing a knot $K \subset Y$, with two w and two z basepoints:

$$\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w_1, w_2, z_1, z_2), \quad \mathcal{H}' = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}', w_1, w_2, z_1, z_2).$$

We assume that $\boldsymbol{\beta}$ and $\boldsymbol{\beta}'$ are strongly equivalent collections of curves, so together they form a (link-minimal, admissible) diagram representing a two-component unlink in a connected sum of $(S^1 \times S^2)$'s. Pick a cycle

$$\Theta \in \boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \mathbf{0}).$$

We can write Θ as a linear combination of intersection points,

$$\Theta = \sum_{\boldsymbol{a} \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}} c(\boldsymbol{a}) U_1^{u_1(\boldsymbol{a})} U_2^{u_2(\boldsymbol{a})} \boldsymbol{a}$$

for $c(\boldsymbol{a}) \in \{0,1\}$ and $u_1(\boldsymbol{a}), u_2(\boldsymbol{a}) \geq 0$.

There are two Alexander filtrations on $\mathbb{T}_\beta \cap \mathbb{T}_{\beta'}$, call them A_1 and A_2, and we are interested in their sum, $A = A_1 + A_2$. The fact that Θ is in bifiltration level $\leq \mathbf{0}$ means that

$$A_1(\boldsymbol{a}) - u_1(\boldsymbol{a}) \leq 0, A_2(\boldsymbol{a}) - u_2(\boldsymbol{a}) \leq 0 \quad \text{for all } \boldsymbol{a} \text{ with } c(\boldsymbol{a}) = 1.$$

and hence

$$A(\boldsymbol{a}) - u_1(\boldsymbol{a}) - u_2(\boldsymbol{a}) \leq 0 \quad \text{for all } \boldsymbol{a} \text{ with } c(\boldsymbol{a}) = 1. \tag{211}$$

The Θ element induces a triangle map (see Section 3.5)

$$\begin{aligned} &\boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, s) \to \boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta'}, s), \\ &\qquad \boldsymbol{x} \mapsto \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}} \sum_{\boldsymbol{a} \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{a}, \boldsymbol{y}) \\ \mu(\phi) = 0}} \#\mathcal{M}(\phi) \cdot c(\boldsymbol{a}) U_1^{n_{w_1}(\phi) + u_1(\boldsymbol{a})} U_2^{n_{w_2}(\phi) + u_2(\boldsymbol{a})} \boldsymbol{y}. \end{aligned}$$

Consider the resolutions $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, s)$ of $\boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, s)$ and $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta'}, s)$ of $\boldsymbol{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta'}, s)$, constructed as in Section 13.1. We claim that Θ also induces a chain map

$$\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, s) \to \mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_{\beta'}, s).$$

To construct it, let us first define an analogue of the complex $\mathfrak{C}^{\mathbf{0}}_{\text{int}}$ from Section 13.1 for the β–β' diagram. This complex, denoted by $\mathfrak{C}^{\mathbf{0}}_{\text{int}}(\mathbb{T}_\beta, \mathbb{T}_{\beta'})$, is free over $\mathfrak{R} = \mathbb{F}[\![U_1, U_2]\!][V_2]'$ with generators $\boldsymbol{a} \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}$, and differential similar to (180), counting holomorphic disks with coefficients

$$U_1^{n_{z_1}(\phi) + n_{z_2}(\phi)} V_2^{n_{w_2}(\phi)}.$$

Recall that in (182) we defined a map $P : \mathfrak{C}_{\text{int}} \to C_{\text{int}}$, whose restriction to $\mathfrak{C}^{\mathbf{0}}_{\text{int}}$ takes

$$\boldsymbol{x} \mapsto U_1^{A(\boldsymbol{x}) - s} \boldsymbol{x}, \quad V_2 \to U_1^{-1} U_2.$$

The idea is to "pull back" the element Θ under a map of this form (with $s = 0$), that is, to consider

$$\widetilde{\Theta} = \sum_{\boldsymbol{a} \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}} c(\boldsymbol{a}) U_1^{u_1(\boldsymbol{a}) + u_2(\boldsymbol{a}) - A(\boldsymbol{a})} V_2^{u_2(\boldsymbol{a})} \boldsymbol{a} \in \mathfrak{C}^{\mathbf{0}}_{\text{int}}(\mathbb{T}_\beta, \mathbb{T}_{\beta'}).$$

Note that the condition (211) implies that U_1 has nonnegative exponents. Further, it is easy to check that since Θ is a cycle, so is $\widetilde{\Theta}$.

There is a filtration $\mathfrak{F}$ on $\mathfrak{C}^{\mathbf{0}}_{\text{int}}(\mathbb{T}_\beta, \mathbb{T}_{\beta'})$ similar to the one from Section 13.1, ie

$$\mathfrak{F}(\boldsymbol{x}) = -A(\boldsymbol{x}), \quad \mathfrak{F}(U_1) = -1, \quad \mathfrak{F}(U_2) = 0, \quad \mathfrak{F}(V_2) = 1.$$

Note that all the terms in the expression for $\widetilde{\Theta}$ are in filtration degree $\mathfrak{F} \leq 0$.

We have a triangle map

$$f\colon \mathfrak{C}_{\mathrm{int}}(\mathbb{T}_\alpha,\mathbb{T}_\beta)\otimes\mathfrak{C}^0_{\mathrm{int}}(\mathbb{T}_\beta,\mathbb{T}_{\beta'})\to\mathfrak{C}_{\mathrm{int}}(\mathbb{T}_\alpha,\mathbb{T}_{\beta'})$$

given by

$$f(\boldsymbol{x}\otimes\boldsymbol{a})=\sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_{\beta'}}\sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{a},\boldsymbol{y})\\ \mu(\phi)=0}}\#\mathcal{M}(\phi)\cdot U_1^{n_{z_1}(\phi)+n_{z_2}(\phi)}V_2^{n_{w_2}(\phi)}\boldsymbol{y}, \tag{212}$$

and equivariant with respect to the actions of U_1, U_2, V_2 and Y_2 (with Y_2 acting by 1 on $\mathfrak{C}^0_{\mathrm{int}}(\mathbb{T}_\beta,\mathbb{T}_{\beta'})$). This map is filtered with respect to the filtrations $\mathfrak{F}$.

By plugging in the $\widetilde{\Theta}$ element and restricting to $\mathfrak{F}\le 0$, we obtain the desired map $\mathfrak{A}^-(\mathbb{T}_\alpha,\mathbb{T}_\beta,s)\to\mathfrak{A}^-(\mathbb{T}_\alpha,\mathbb{T}_{\beta'},s)$, given by

$$\boldsymbol{x}\to\sum_{\boldsymbol{a}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'}}c(\boldsymbol{a})U_1^{u_1(\boldsymbol{a})+u_2(\boldsymbol{a})-A(\boldsymbol{a})}V_2^{u_2(\boldsymbol{a})}f(\boldsymbol{x}\otimes\boldsymbol{a}). \tag{213}$$

Higher polygon maps between the resolutions are constructed in a similar manner.

13.4 A knot with many basepoints

We now consider a knot K represented by a Heegaard diagram $\mathcal{H}^K$ with $2p$ basepoints, denoted (in order) by $w_1,z_1,w_2,z_2,\dots,w_p,z_p$.

Just as we had to construct a resolution of $C_{\mathrm{int}}=C\{w;n_2\ge 0,\ A(\boldsymbol{x})-n_1-n_2\le s\}$ for a knot with four basepoints, we now start by constructing a resolution $\mathfrak{C}_{\mathrm{int}}$ of the intermediate complex

$$C_{\mathrm{int}}=C\{w;n_2,\dots,n_p\ge 0,\ A(\boldsymbol{x})-n_1-\dots-n_p\le s\},$$

with elements of the form

$$\boldsymbol{x}_{\boldsymbol{n}}:=U_1^{A(\boldsymbol{x})-s-|\boldsymbol{n}|}U_2^{n_2}\cdots U_p^{n_p}\boldsymbol{x},\qquad \boldsymbol{x}\in\mathbb{T}_\alpha\cap\mathbb{T},\quad \boldsymbol{n}=(n_2,\dots,n_p)\in(\mathbb{Z}_{\ge 0})^{p-1}, \tag{214}$$

where we wrote $|\boldsymbol{n}|=n_2+\dots+n_p$. Every element of C_{int} can be written as an infinite sum

$$\sum_{\boldsymbol{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\sum_{i=1}^{\infty}U_1^{a_i}U_2^{b_i^2}\cdots U_n^{b_i^n}\boldsymbol{x}_{\boldsymbol{n}_i}$$

for some sequence of distinct tuples $(a_i,b_i^2,\dots,b_i^n,\boldsymbol{n}_i=(n_{i,2},\dots,n_{i,p}))_{i\ge 1}$ satisfying a boundedness condition similar to (179), namely

$$(\exists K>0)(\forall i)\quad |\boldsymbol{n}_i|-a_i<K.$$

In the construction of $\mathfrak{C}_{\mathrm{int}}$, we turn $\boldsymbol{x}_{\boldsymbol{n}}$ into linearly independent elements over $\mathbb{F}[\![U_1,\dots,U_p]\!]$, by introducing new variables $V_2,\dots,V_p$, so that $\boldsymbol{x}_{\boldsymbol{n}}$ corresponds to $V_2^{n_2}\cdots V_p^{n_p}\boldsymbol{x}$. More precisely, we proceed as follows. Given a monomial $\boldsymbol{m}$ in several variables, we will write $e_{\boldsymbol{m}}(X)$ for the exponent of a variable X in $\boldsymbol{m}$. We define the subring $\mathfrak{R} \subset \mathbb{F}[\![U_1, U_2, \dots, U_p, V_2, \dots, V_p]\!]$ to consist of infinite sums of monomials $\boldsymbol{m}$ such that

$$e_{\boldsymbol{m}}(V_2) + \cdots + e_{\boldsymbol{m}}(V_p) - e_{\boldsymbol{m}}(U_1)$$

is bounded above (over all $\boldsymbol{m}$ in a sum). Note that $\mathfrak{R}$ is a module over $\mathbb{F}[\![U_1, U_2]\!]$, and contains the ring $\mathbb{F}[\![U_1, U_2, \dots, U_p]\!][V_2, \dots, V_p]$. We will write

$$\mathfrak{R} = \mathbb{F}[\![U_1, U_2, \dots, U_p]\!][V_2, \dots, V_p]'.$$

Let $\mathfrak{C}^0_{\text{int}}$ be the complex freely generated over $\mathfrak{R}$ by the intersection points $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, with differential

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi) = 1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{n_{z_1}(\phi) + \cdots + n_{z_p}(\phi)} V_2^{n_{w_2}(\phi)} \cdots V_p^{n_{w_p}(\phi)} \boldsymbol{y}. \tag{215}$$

The fact that it is a complex ($\partial^2 = 0$) follows from [Ozsváth and Szabó 2008a, Theorem 5.5]: There are $2p$ boundary degenerations to account for, each containing a pair of basepoints,

$$(w_1, z_1), (z_1, w_2), (w_2, z_2), \dots, (w_p, z_p), (z_p, w_1).$$

Since we keep track of all z basepoints through the same variable U_1, the contributions from boundary degenerations cancel in pairs.

Next, we construct $\mathfrak{C}_{\text{int}}$ from $\mathfrak{C}^0_{\text{int}}$ by taking iterated mapping cones for maps of the form $U_i + U_1 V_i$, that is, we introduce new variables $Y_2, \dots, Y_p$ satisfying

$$Y_i^2 = 0, \quad \partial Y_i = U_i + U_1 V_i.$$

We then take the tensor product

$$\mathfrak{C}_{\text{int}} = \mathfrak{C}^0_{\text{int}} \otimes_{\mathfrak{R}} \bigotimes_{i=2}^{p} \mathfrak{Y}_i,$$

where $\mathfrak{Y}_i$ is the free $\mathfrak{R}$-complex with two generators 1 and Y_i, and with differential $\partial Y_i = U_i + U_1 V_i$.

In other words, $\mathfrak{C}_{\text{int}}$ is the dg module over the dga

$$\mathfrak{R}^Y = \mathbb{F}[\![U_1, \dots, U_p]\!][V_2, \dots, V_p, Y_2, \dots, Y_p]'/(Y_i^2 = 0, \partial Y_i = U_i + U_1 V_i)$$

generated by $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, with differential (215).

We define a projection map $P : \mathfrak{C}_{\text{int}} \to C_{\text{int}}$ by

$$P(V_2^{n_2} \cdots V_p^{n_p} Y_2^{a_2} \cdots Y_p^{a_p} \boldsymbol{x}) = \begin{cases} U_1^{A(\boldsymbol{x}) - s - |\boldsymbol{n}|} U_2^{n_2} \cdots U_p^{n_p} \boldsymbol{x} & \text{for } a_2 = \cdots = a_p = 0, \\ 0 & \text{otherwise,} \end{cases} \tag{216}$$

and extending it linearly (to the allowed infinite sums) over $\mathcal{R} = \mathbb{F}[\![U_1, U_2, \dots, U_p]\!]$.

Lemma 13.8 *The map P is a chain map, and exhibits $\mathfrak{C}_{\text{int}}$ as a resolution of C_{int}.*

Proof The fact that P commutes with the differential follows as in the proof of Lemma 13.3.

To show that P is a quasi-isomorphism, observe that P is surjective, and denote its kernel by G. We seek to prove that G is acyclic.

Note that G is generated over $\mathfrak{R}$ by the elements that contain nontrivial powers of some Y_i's, as well as by the sums $(U_i + U_1 V_i)\boldsymbol{x}$, corresponding to relations between the generators $\boldsymbol{x_n}$ of C_{int}; see (214).

Consider an area filtration $\mathscr{A}$ on the generators of G, as follows. Let Π be the space of periodic domains ϕ on the Heegaard diagram $\mathscr{H}^K$ that satisfy

$$n_{z_1}(\phi) + \cdots + n_{z_p}(\phi) = 0, \quad n_{w_2}(\phi) = \cdots = n_{w_p}(\phi) = 0.$$

Equip the Heegaard surface with an area form such that all domains in Π have zero area. Pick a base generator $\boldsymbol{x}_b \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and set $\mathscr{A}(\boldsymbol{x}_b) = 0$. Then, for any other generator of the form $U_1^m V_2^{n_2} \cdots V_p^{n_p} \boldsymbol{x}$, pick a class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{x}_b)$ with

$$n_{z_1}(\phi) + \cdots + n_{z_p}(\phi) = m, \quad n_{w_2}(\phi) = n_2, \quad \ldots, \quad n_{w_p}(\phi) = n_p$$

and let $\mathscr{A}(U_1^m V_2^{n_2} \cdots V_p^{n_p} \boldsymbol{x})$ be the area of ϕ. Finally, let $\mathscr{A}$ be unchanged when we multiply by variables Y_i or U_i for $i \geq 2$.

In view of (215), $\mathscr{A}$ defines a filtration on the complex $\mathfrak{C}_{\mathrm{int}}$, and hence on its subcomplex G. The associated graded of $\mathscr{A}$ on $\mathfrak{C}_{\mathrm{int}}$ is a direct sum, over $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, of Koszul complexes for the regular sequence

$$(U_2 + U_1 V_2, U_3 + U_1 V_3, \ldots, U_p + U_1 V_p)$$

over the ring $\mathfrak{R}$. The homology of such a Koszul complex is concentrated in degree zero, with respect to the grading given by the sum of the exponents of the Y_i variables. It follows that the same is true for the associated graded for $\mathscr{A}$ on G. Moreover, in degree zero, all the generators $(U_i + U_1 V_i)\boldsymbol{x}$ of G are in the image of the differential ∂ on the associated graded (because $\partial Y_i = U_i + U_1 V_i$). We conclude that the associated graded of G is acyclic, and from here that G is acyclic. $\square$

Next, as in the case of four basepoints, we define a filtration $\mathscr{F}$ on C_{int} by the negative of the exponent of U_1. Then $\boldsymbol{A}^-(\mathscr{H}^K, s)$ is the subcomplex of C_{int} given by $\mathscr{F} \leq 0$.

We also define a filtration $\mathfrak{F}$ on $\mathfrak{C}_{\mathrm{int}}$ by setting

$$\mathfrak{F}(\boldsymbol{x}) = -A(\boldsymbol{x}), \quad \mathfrak{F}(U_1) = -1, \quad \mathfrak{F}(U_i) = 0, \quad \mathfrak{F}(V_i) = 1, \quad \mathfrak{F}(Y_i) = 0, \quad i \geq 2.$$

One can check that the map P from (216) is filtered of degree $-s$ with respect to $\mathfrak{F}$ and $\mathscr{F}$.

We define a resolution $\mathfrak{A}^-(\mathscr{H}^K, s)$ of $\boldsymbol{A}^-(\mathscr{H}^K, s)$ to be the subcomplex of $\mathfrak{C}_{\mathrm{int}}$ corresponding to $\mathfrak{F} \leq s$. Thus, $\mathfrak{A}^-(\mathscr{H}^K, s)$ consists of infinite sums (with $\mathscr{R}$ coefficients, and a suitable boundedness condition) of the elements

$$\tilde{\boldsymbol{x}}_{\boldsymbol{n}}^{\boldsymbol{a}} := U_1^{\max(s+|\boldsymbol{n}|-A(\boldsymbol{x}),0)} V_2^{n_2} \cdots V_p^{n_p} Y_2^{a_2} \cdots Y_p^{a_p} x,$$

where $\boldsymbol{n} = (n_2, \dots, n_p)$, $\boldsymbol{a} = (a_2, \dots, a_p)$.

We let

$$\widetilde{P} : \mathfrak{A}^-(\mathcal{H}^K, s) \to \boldsymbol{A}^-(\mathcal{H}^K, s)$$

to be the restriction of P from (216). In terms of the generators $\tilde{\boldsymbol{x}}_{\boldsymbol{n}}^{\boldsymbol{a}}$,

$$\widetilde{P}(\tilde{\boldsymbol{x}}_{\boldsymbol{n}}^{\boldsymbol{a}}) = \begin{cases} U_1^{\max(A(\boldsymbol{x})-s-|\boldsymbol{n}|,0)} U_2^{n_2} \cdots U_p^{n_p} \boldsymbol{x} & \text{for } a_2 = \cdots = a_p = 0, \\ 0 & \text{otherwise.} \end{cases}$$

By the same arguments as in Lemma 13.8, we can show that $\widetilde{P}$ exhibits $\mathfrak{A}^-(\mathcal{H}^K, s)$ as a resolution of $\boldsymbol{A}^-(\mathcal{H}^K, s)$.

As in Section 13.1, we define the map Φ_s^K in the surgery formula as the composition

$$\mathfrak{A}^-(\mathcal{H}^K, s) \xrightarrow{\widetilde{P}} \boldsymbol{A}^-(\mathcal{H}^K, s) \hookrightarrow C\{w; n_1, \dots, n_p \geq 0\} \xrightarrow{\sim} \boldsymbol{CF}^-(\mathcal{H}^{\varnothing}),$$

Recall also that in Section 13.1 we defined the map Φ_s^{-K} as the composition

$$\mathfrak{A}^-(\mathcal{H}^K, s) \hookrightarrow \mathfrak{C}_{\mathrm{int}} \xrightarrow{\Xi} C\{z; n_1, n_2 \geq 0\} \xrightarrow{\sim} \boldsymbol{CF}^-(\mathcal{H}^{\varnothing}),$$

with the last map again obtained from Heegaard moves.

In the general case, we have a similar sequence, but instead of the transition map Ξ we have a composition of $p-1$ maps

$$\mathfrak{C}_{\mathrm{int}} = \mathfrak{C}_0 \xrightarrow{\Xi_0} \mathfrak{C}_1 \xrightarrow{\Xi_1} \cdots \xrightarrow{\Xi_{p-2}} \mathfrak{C}_{p-1} = C\{z; n_1, \dots, n_p \geq 0\}. \tag{217}$$

Here, for $0 \leq j \leq p-1$, we consider the dga

$$\mathfrak{R}_j^Y = \mathbb{F}[\![U_1, \dots, U_p]\!][V_{j+2}, \dots, V_p, Y_{j+2}, \dots, Y_p]' / (Y_i^2 = 0, \partial Y_i = U_i + U_1 V_i),$$

where the prime refers to allowing infinite sums of monomials $\boldsymbol{m}$ in the given variables, provided that

$$e_{\boldsymbol{m}}(V_{j+2}) + \cdots + e_{\boldsymbol{m}}(V_p) - e_{\boldsymbol{m}}(U_1)$$

is bounded above. The complex $\mathfrak{C}_j$ is defined as the dg module over $\mathfrak{R}_j^Y$ generated by $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, with differential

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi) = 1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{n_{z_1}(\phi)} \cdots U_j^{n_{z_j}(\phi)} U_{j+1}^{n_{z_{j+1}}(\phi) + \cdots + n_{z_p}(\phi)} V_{j+2}^{n_{w_{j+2}}(\phi)} \cdots V_p^{n_{w_p}(\phi)}.$$

The transition maps $\Xi_j : \mathfrak{C}_j \to \mathfrak{C}_{j+1}$ are constructed as follows. Recall that for $p = 2$ we had equations (186) and (187). In general, we set

$$\Xi_j(V_{j+2}^n Y_{j+2} \boldsymbol{x}) = \boldsymbol{x} \tag{218}$$

and

$$\Xi_j(V_{j+2}^n \boldsymbol{x}) = \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_1^{n_{z_1}(\phi)} \cdots U_{j+1}^{n_{z_{j+1}}(\phi)} \cdot \frac{U_{j+1}^{n_{z_{j+2}}(\phi)+\cdots+n_{z_p}(\phi)} - U_{j+2}^{n_{z_{j+2}}(\phi)+\cdots+n_{z_p}(\phi)}}{U_{j+1}-U_{j+2}} \cdot V_{j+3}^{n_{w_{j+3}}(\phi)} \cdots V_p^{n_{w_p}(\phi)} \boldsymbol{y}, \tag{219}$$

then extend these maps equivariantly with respect to the action of the variables $U_1, \dots, U_p, V_{j+3}, \dots, V_p, Y_{j+3}, \dots, Y_p$.

Lemma 13.9 Ξ_j *is a chain map.*

Proof This is entirely similar to the proof of Lemma 13.6, with the variables U_{j+1} and U_{j+2} playing the roles of U_1 and U_2. Furthermore, all holomorphic disks get counted with an additional factor of

$$U_1^{n_{z_1}(\phi)} \cdots U_j^{n_{z_j}(\phi)} V_{j+3}^{n_{w_{j+3}}(\phi)} \cdots V_p^{n_{w_p}(\phi)}. \qquad \square$$

13.5 A link with two components and eight basepoints

As the next warm-up for the general case of setting up the hyperboxes, let us consider a link $L = L_1 \cup L_2$ with two components. Suppose this is represented by a complete system of hyperboxes $\mathscr{H}$, where the initial diagram $\mathscr{H}^L$ has four basepoints on each component:

$$w_{1,1}, z_{1,1}, w_{1,2}, z_{1,2} \quad \text{on } L_1,$$

and

$$w_{2,1}, z_{2,1}, w_{2,2}, z_{2,2} \quad \text{on } L_2.$$

The system $\mathscr{H}$ consists of several diagrams and hyperboxes, as listed in Example 8.25. We seek to construct resolutions $\mathfrak{A}^-(\mathscr{H}^M, \boldsymbol{s})$ of $\boldsymbol{A}^-(\mathscr{H}^M, \boldsymbol{s})$, for each sublink $M \subseteq L$, and also maps

$$\Phi_{\boldsymbol{s}}^{\vec{M}} : \mathfrak{A}^-(\mathscr{H}^{L'}, \boldsymbol{s}) \to \mathfrak{A}^-(\mathscr{H}^{L'-M}, \psi^{\vec{M}}(\boldsymbol{s})) \quad \text{for } M \subseteq L' \subseteq L,$$

associated to the hyperboxes $\mathscr{H}^{\vec{L}', \vec{M}}$.

The complexes $\mathfrak{A}^-(\mathscr{H}^{L_1}, \boldsymbol{s})$, $\mathfrak{A}^-(\mathscr{H}^{\varnothing}) = \boldsymbol{CF}^-(\mathscr{H}^{\varnothing})$ and the chain map associated to $\mathscr{H}^{\vec{L}_1, \pm\vec{L}_1}$ are constructed just as in Section 13.1, with L_1 playing the role of K, except that when counting holomorphic disks and triangles, we also keep track of the basepoints on L_2, with coefficients

$$U_{2,1}^{n_{w_{2,1}}(\phi)} U_{2,2}^{n_{w_{2,2}}(\phi)}.$$

Similar remarks apply to $\mathfrak{A}^-(\mathscr{H}^{L_2}, \boldsymbol{s})$ and the map associated to $\mathscr{H}^{\vec{L}_2, \pm\vec{L}_2}$.

Let us construct $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$. By analogy with what we did in Section 13.1, we first define $\mathfrak{R} = \mathbb{F}[\![U_{1,1}, U_{1,2}, U_{2,1}, U_{2,2}]\!][V_{1,2}, V_{2,2}]'$ to consist of power series made of monomials $\boldsymbol{m}$ such that

$$e_{\boldsymbol{m}}(V_{1,2}) - e_{\boldsymbol{m}}(U_{1,1}) \quad \text{and} \quad e_{\boldsymbol{m}}(V_{2,2}) - e_{\boldsymbol{m}}(U_{2,1})$$

are bounded above. We let $\mathfrak{C}^0_{\text{int}}(\mathcal{H}^L)$ to be freely generated by $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ over $\mathfrak{R}$, with differential counting holomorphic disks with coefficients

$$U_{1,1}^{n_{z_{1,1}}(\phi)+n_{z_{1,2}}(\phi)} V_{1,2}^{n_{w_{1,2}}(\phi)} U_{2,1}^{n_{z_{2,1}}(\phi)+n_{z_{2,2}}(\phi)} V_{2,2}^{n_{w_{2,2}}(\phi)}.$$

The complex $\mathfrak{C}_{\text{int}}(\mathcal{H}^L)$ is obtained from $\mathfrak{C}^0_{\text{int}}(\mathcal{H}^L)$ by adjoining variables $Y_{1,2}, Y_{2,2}$ with relations

$$Y_{1,2}^2 = 0, \quad \partial Y_{1,2} = U_{1,2} + U_{1,1}V_{1,2}, \tag{220}$$

$$Y_{2,2}^2 = 0, \quad \partial Y_{2,2} = U_{2,2} + U_{2,1}V_{2,2}. \tag{221}$$

There are now two filtrations, $\mathfrak{F}_1$ and $\mathfrak{F}_2$, on $\mathfrak{C}_{\text{int}}(\mathcal{H}^L)$, given by

$$\mathfrak{F}_1(\boldsymbol{x}) = -A_1(\boldsymbol{x}), \quad \mathfrak{F}_1(U_{1,1}) = -1, \quad \mathfrak{F}(U_{1,2}) = 0, \quad \mathfrak{F}(V_{1,2}) = 1, \quad \mathfrak{F}(Y_{1,2}) = 0,$$
$$\mathfrak{F}_1(U_{2,1}) = \mathfrak{F}(U_{2,2}) = \mathfrak{F}(V_{2,2}) = \mathfrak{F}(Y_{2,2}) = 0$$

and similarly for $\mathfrak{F}_2$. For $\boldsymbol{s} = (s_1, s_2)$, we let $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ be the subcomplex of $\mathfrak{C}_{\text{int}}(\mathcal{H}^L)$ in bifiltration degrees $\mathcal{F}_1 \leq s_1, \mathcal{F}_2 \leq s_2$.

Next, consider the one-dimensional hyperboxes $\mathcal{H}^{L, \pm L_1}$. The corresponding maps

$$\Phi_{\boldsymbol{s}}^{\pm L_1} : \mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \to \mathfrak{A}^-(\mathcal{H}^{L_2}, \psi^{\pm L_1}(\boldsymbol{s}))$$

are constructed as in Section 13.1, except that at each step, we keep track of the basepoints on L_2 through coefficients

$$U_{2,1}^{n_{z_{2,1}}(\phi)+n_{z_{2,2}}(\phi)} V_{2,2}^{n_{w_{2,2}}(\phi)}$$

and we also keep the variable $Y_{2,2}$ with relations (221). Note that, as part of the construction, the Θ elements that are part of the data in $\mathcal{H}^{L, \pm L_1}$ are used to define an equivalence from $\mathfrak{A}^-(r_{\pm L_1}(\mathcal{H}^L), \psi^{\pm L_1}(\boldsymbol{s}))$ to $\mathfrak{A}^-(\mathcal{H}^{L_2}, \psi^{\pm L_1}(\boldsymbol{s}))$. This equivalence is the composition of several triangle maps, induced by the Θ elements as explained in Section 13.3.

The maps associated to $\mathcal{H}^{L, \pm L_2}$ are defined in the same way.

The last ingredient in the surgery formula are chain homotopies associated to the four two-dimensional hyperboxes in $\mathcal{H}$. We leave the constructions corresponding to $\mathcal{H}^{L, L_1 \cup L_2}$, $\mathcal{H}^{L, L_1 \cup -L_2}$ and $\mathcal{H}^{L, (-L_1) \cup L_2}$ as an exercise for the reader, and focus on the most complicated case, that of the hyperbox $\mathcal{H}^{L, -L_1 \cup -L_2}$. (These constructions are all particular cases of the general definition, which will be given in Section 13.6.)

We seek to construct the map $\Phi_{\boldsymbol{s}}^{-L}$, where $\boldsymbol{s} = (s_1, s_2)$. This will be a composition

$$\Phi_{\boldsymbol{s}}^{-L} = D^{-L}_{p^{-L}(\boldsymbol{s})} \circ \mathcal{I}_{\boldsymbol{s}}^{-L}.$$

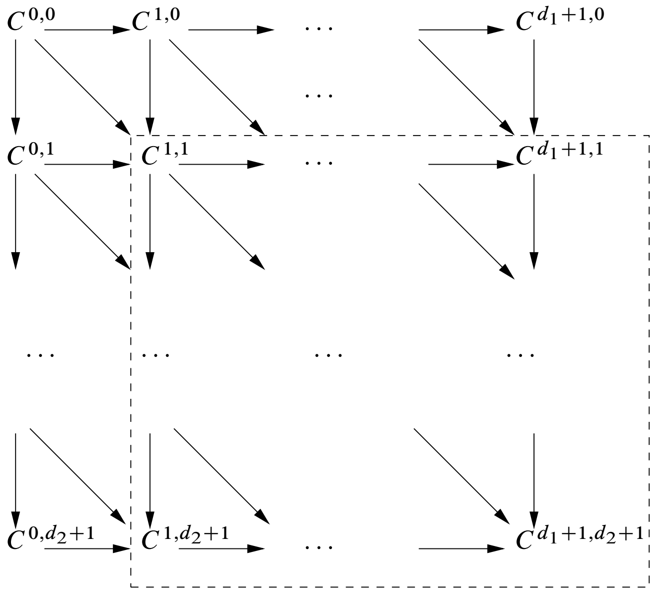

Figure 40: A hyperbox of chain complexes of size (d_1+1, d_2+1). This is the hyperbox $\mathfrak{C}(\mathcal{H}^{L,-L})$. The initial vertex is $\mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$, and the rectangle bounded by dashed lines is the hyperbox obtained by applying $\boldsymbol{CF}^-$ to the diagrams in $\mathcal{H}^{L,-L}$.

Here, $\mathcal{I}_{\boldsymbol{s}}^{-L}$ is just the inclusion of $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ into $\mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$. The descent map $D_{\boldsymbol{s}}^{-L}$ will be obtained by compressing a certain hyperbox, which we proceed to describe. Note that, in this case, $p^{-L}(\boldsymbol{s}) = (-\infty, -\infty)$ does not depend on $\boldsymbol{s}$, so we can simply write D^{-L} or $D_{\boldsymbol{s}}^{-L}$.

As part of the data we have the two-dimensional hyperbox $\mathcal{H}^{L,-L}$ of strongly equivalent Heegaard diagrams (representing $\varnothing \subset S^3$). This hyperbox starts at $r_{-L}(\mathcal{H}^L)$, the diagram obtained from $\mathcal{H}^L$ by deleting the w basepoints, and ends at $\mathcal{H}^{\varnothing}$. Suppose that this hyperbox (rectangle) is of size $\boldsymbol{d} = (d_1, d_2)$, and let us denote by $\mathcal{H}_{\varepsilon}^{L,-L}$ the diagram at position $\varepsilon = (\varepsilon_1, \varepsilon_2)$, with $\varepsilon_i \in \{0, \dots, d_i\}$.

We also have the one-dimensional hyperbox $\mathcal{H}^{L,-L_1}$ of size d_1, consisting of strongly equivalent Heegaard diagrams that represent L_2. This hyperbox starts at $r_{-L_1}(\mathcal{H}^L)$, the diagram obtained from $\mathcal{H}^L$ by deleting the w basepoints on L_1 only. Further, by the compatibility relation between hyperboxes (see Definition 8.22), when we apply r_{-L_2} to the hyperbox $\mathcal{H}^{L,-L_1}$, we should obtain the first row in the two-dimensional hyperbox $\mathcal{H}^{L,-L}$. Thus, if we denote by $\mathcal{H}_{\varepsilon_1}^{L,-L_1}$ the diagram at position ε_1, we have $r_{-L_2}(\mathcal{H}_{\varepsilon_1}^{L,-L_1}) = \mathcal{H}_{\varepsilon_1,0}^{L,-L}$.

There is also a similar hyperbox $\mathcal{H}^{L,-L_2}$, of size d_2, in the other direction. (Compare Example 8.25 and Figure 21.) Its reduction with respect to L_1 is the first column in $\mathcal{H}^{L,-L}$.

Note that from $\mathcal{H}^{L,-L}$ we could get a hyperbox of chain complexes of size (d_1, d_2), by applying $\boldsymbol{CF}^-$ to each diagram $\mathcal{H}_{\varepsilon}^{L,-L}$, as in Section 8.3. However, we want to include transition maps into this hyperbox, so we add an additional row and column in the beginning. The resulting hyperbox, denoted by $\mathfrak{C}(\mathcal{H}^{L,-L})$

$$\begin{tikzcd}[column sep=huge, row sep=huge]
C^{0,0} \arrow[loop left, "U_{1,1}^{m_1}U_{2,1}^{m_2}"] \arrow[r, "\frac{U_{1,1}^{m_1}-U_{1,2}^{m_1}}{U_{1,1}-U_{1,2}}U_{2,1}^{m_2}"] \arrow[d, "U_{1,1}^{m_1}\frac{U_{2,1}^{m_2}-U_{2,2}^{m_2}}{U_{2,1}-U_{2,2}}"'] \arrow[dr, bend right, "\frac{U_{1,1}^{m_1}-U_{1,2}^{m_1}}{U_{1,1}-U_{1,2}}\frac{U_{2,1}^{m_2}-U_{2,2}^{m_2}}{U_{2,1}-U_{2,2}}"] & C^{1,0} \arrow[loop right, "U_{1,2}^{m_1}U_{2,1}^{m_2}"] \arrow[d, "U_{1,2}^{m_1}\frac{U_{2,1}^{m_2}-U_{2,2}^{m_2}}{U_{2,1}-U_{2,2}}"] \\
C^{0,1} \arrow[loop left, "U_{1,1}^{m_1}U_{2,2}^{m_2}"] \arrow[r, "\frac{U_{1,1}^{m_1}-U_{1,2}^{m_1}}{U_{1,1}-U_{1,2}}U_{2,2}^{m_2}"'] & C^{1,1} \arrow[loop right, "U_{1,2}^{m_1}U_{2,2}^{m_2}"]
\end{tikzcd}$$

Figure 41: The first square in the hyperbox $\mathfrak{C}(\mathcal{H}^{L,-L})$. Each arrow is a map counting holomorphic disks. We indicated the coefficients that keep track of the basepoints $z_{1,2}$ and $z_{2,2}$.

and pictured in Figure 40, is of size (d_1+1, d_2+1). At its vertices we have complexes

$$C^{\varepsilon_1,\varepsilon_2} = \begin{cases} \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L) & \text{if } \varepsilon_1=\varepsilon_2=0, \\ \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^{L,-L_1}_{\varepsilon_1-1}) & \text{if } \varepsilon_1\neq 0,\, \varepsilon_2=0, \\ \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^{L,-L_2}_{\varepsilon_2-1}) & \text{if } \varepsilon_1=0,\, \varepsilon_2\neq 0, \\ \boldsymbol{CF}^-(\mathcal{H}^{L,-L}_{\varepsilon_1-1,\varepsilon_2-1}) & \text{if } \varepsilon_1,\varepsilon_2\neq 0. \end{cases}$$

We need to specify the chain maps along the horizontal and vertical segments in the hyperbox, and the chain homotopies along diagonals. In the bottom-right rectangle of size (d_1, d_2), they are the usual maps in $\boldsymbol{CF}^-(\mathcal{H}^{L,-L})$, given by counting holomorphic triangles or quadrilaterals, with some vertices fixed at Θ elements that are part of the data in $\mathcal{H}^{L,-L}$. Note that these holomorphic triangles and quadrilaterals, in classes ϕ, are counted with coefficients

$$U_{1,1}^{n_{z_{1,1}}(\phi)}U_{1,2}^{n_{z_{1,2}}(\phi)}U_{2,1}^{n_{z_{2,1}}(\phi)}U_{2,2}^{n_{z_{2,2}}(\phi)}.$$

We are left to specify the chain maps and chain homotopies in the first row and column of the hyperbox $\mathfrak{C}(\mathcal{H}^{L,-L})$.

Let us start by looking at the top-left square, pictured in Figure 41. Observe that, at its vertices, we have chain complexes associated to the same diagram $\mathcal{H}^L$, in which we keep track of the basepoints in various ways. With regard to the basepoints $z_{1,1}$ and $z_{2,1}$, at all four vertices, the coefficients in the differential pick up factors of

$$U_{1,1}^{n_{z_{1,1}}(\phi)}U_{2,1}^{n_{z_{2,1}}(\phi)}.$$

On the other hand, with regard to $z_{1,2}$ and $z_{2,2}$, these are counted differently, depending on the vertex. If we set

$$m_1 := n_{z_{1,2}}(\phi), \quad m_2 := n_{z_{2,2}}(\phi),$$

then m_1 appears as an exponent of either $U_{1,1}$ or $U_{1,2}$, and m_2 appears as an exponent of either $U_{2,1}$ or $U_{2,2}$ as indicated in Figure 41. We also have the w basepoints, which appear as exponents of V variables. In $C^{0,0} = \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$, we have

$$V_{1,2}^{n_{w_{1,2}}(\phi)} V_{2,2}^{n_{w_{2,2}}(\phi)},$$

then in $C^{1,0}$ we set $V_{1,2} = 1$ (and keep $V_{2,2}$ as before). In $C^{0,1}$ we set $V_{2,2} = 1$, and in $C^{1,1}$ both V variables are set to 1. A similar pattern holds for the Y variables: in $C^{0,0}$ we have both $Y_{1,2}$ and $Y_{2,2}$, with relations (220) and (221); in $C^{1,0}$ we drop $Y_{1,2}$, in $C^{0,1}$ we drop $Y_{2,2}$, and in $C^{1,1}$ we drop both.

With regard to the chain maps along the edges of the square in Figure 41, these are transition maps similar to the map Ξ from Section 13.1. For example, the top horizontal map

$$D_{0,0}^{1,0} \colon C^{0,0} \to C^{1,0}$$

is defined as follows. This map is equivariant with respect to the action of the U and V variables, as well to the action of $Y_{2,2}$ (but not $Y_{1,2}$, which is a variable that does not appear in the target complex $C^{1,0}$). The map $D_{0,0}^{1,0}$ takes

$$Y_{1,2}\boldsymbol{x} \mapsto \boldsymbol{x}, \quad \boldsymbol{x} \mapsto \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x},\boldsymbol{y}) \\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_{1,1}^{n_{z_{1,1}}(\phi)} \frac{U_{1,1}^{n_{z_{1,2}}(\phi)} - U_{1,2}^{n_{z_{1,2}}(\phi)}}{U_{1,1} - U_{1,2}} U_{2,1}^{n_{z_{2,1}}(\phi) + n_{z_{2,2}}(\phi)} \boldsymbol{y}.$$

The bottom horizontal map, $D_{0,1}^{1,0}$, is defined in the same way, with the only difference that when counting holomorphic disks, the quantity $n_{z_{2,2}}(\phi)$ appears as an exponent of $U_{2,2}$ instead of $U_{2,1}$. The two vertical chain maps have similar definitions, with the roles of L_1 and L_2 (and the corresponding variables) switched.

Finally, the chain homotopy along the diagonal of the square,

$$D_{0,0}^{1,1} \colon C^{0,0} \to C^{1,1}$$

is taken to be equivariant with respect to the U and V variables, to take the $Y_{1,2}$ and $Y_{2,2}$ variables to zero, and to satisfy

$$\boldsymbol{x} \mapsto \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x},\boldsymbol{y}) \\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_{1,1}^{n_{z_{1,1}}(\phi)} \frac{U_{1,1}^{n_{z_{1,2}}(\phi)} - U_{1,2}^{n_{z_{1,2}}(\phi)}}{U_{1,1} - U_{1,2}} U_{2,1}^{n_{z_{2,1}}(\phi)} \frac{U_{2,1}^{n_{z_{2,1}}(\phi)} - U_{2,2}^{n_{z_{2,2}}(\phi)}}{U_{2,1} - U_{2,2}} \boldsymbol{y}.$$

Lemma 13.10 *With these definitions, the total differential D on the first square in the hyperbox $\mathfrak{C}(\mathcal{H}^{L,-L})$ satisfies $D^2 = 0$.*

Proof The one interesting case to check is when we apply D^2 to a generator $\mathbf{x} \in C^{0,0}$ (without any Y variables). Then $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ (viewed as an element in $C^{1,1}$) appears in $D^2(\mathbf{x})$ with a sum of coefficients that come from pairs of holomorphic disks, in classes $\phi \in \pi_2(\mathbf{x}, \mathbf{a})$, $\psi \in \pi_2(\mathbf{a}, \mathbf{y})$, over all $\mathbf{a} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. For each pair (ϕ, ψ), there are four such contributions, from each of the four terms in

$$D^2 = D^{0,1}_{1,0} \circ D^{1,0}_{0,0} + D^{1,0}_{0,1} \circ D^{0,1}_{0,0} + D^{1,1}_{0,0} \circ D^{0,0}_{0,0} + D^{0,0}_{1,1} \circ D^{1,1}_{0,0}. \tag{222}$$

In other words, we go from the initial to the final vertex in Figure 41 either by following the edges, or along the diagonal combined with a differential (self-map) on either the initial or the final vertex.

We claim that these terms cancel in pairs. To see this, we need to track the powers of the U variables in each of them. In all the terms we have factors of the form

$$U_{1,1}^{n_{z_{1,1}}(\phi*\psi)} U_{2,1}^{n_{z_{2,1}}(\phi*\psi)}.$$

We will ignore these, and focus on the appearances of $n_{z_{1,2}}$ and $n_{z_{2,2}}$. These are indicated in Figure 41, for each map. Since we now have two different classes ϕ and ψ, let us define

$$m_1 := n_{z_{1,2}}(\phi), \quad m_2 := n_{z_{2,2}}(\phi), \quad m_1' := n_{z_{1,2}}(\psi), \quad m_2' := n_{z_{2,2}}(\psi).$$

For each pair of disks in classes (ϕ, ψ), we get the following contributions to the coefficient of $\mathbf{y}$ in $\partial^2 \mathbf{x}$, corresponding to the four terms in (222):

$$\begin{aligned} &\frac{U_{1,1}^{m_1} - U_{1,2}^{m_1}}{U_{1,1} - U_{1,2}} U_{2,1}^{m_2} U_{1,2}^{m_1'} \frac{U_{2,1}^{m_2'} - U_{2,2}^{m_2'}}{U_{2,1} - U_{2,2}} + U_{1,1}^{m_1} \frac{U_{2,1}^{m_2} - U_{2,2}^{m_2}}{U_{2,1} - U_{2,2}} \frac{U_{1,1}^{m_1'} - U_{1,2}^{m_1'}}{U_{1,1} - U_{1,2}} U_{2,2}^{m_2'} \\ &\qquad + U_{1,1}^{m_1} U_{2,1}^{m_2} \frac{U_{1,1}^{m_1'} - U_{1,2}^{m_1'}}{U_{1,1} - U_{1,2}} \frac{U_{2,1}^{m_2'} - U_{2,2}^{m_2'}}{U_{2,1} - U_{2,2}} + \frac{U_{1,1}^{m_1} - U_{1,2}^{m_1}}{U_{1,1} - U_{1,2}} \frac{U_{2,1}^{m_2} - U_{2,2}^{m_2}}{U_{2,1} - U_{2,2}} U_{1,2}^{m_1'} U_{2,2}^{m_2'}. \end{aligned}$$

We obtain a fraction whose denominator is $(U_{1,1} - U_{1,2})(U_{2,1} - U_{2,2})$, and whose numerator is a sum of sixteen terms. Twelve of those terms cancel one another in pairs, and the remaining four terms can be combined into the product

$$(U_{1,1}^{m_1+m_1'} - U_{1,2}^{m_1+m_1'})(U_{2,1}^{m_2+m_2'} - U_{2,2}^{m_2+m_2'}).$$

Observe that $m_1 + m_1' = n_{z_{1,2}}(\phi * \psi)$ and $m_2 + m_2' = n_{z_{2,2}}(\phi * \psi)$. Thus, we are left with a count of pairs of holomorphic disks, where the exponents depend only on the combined class $\phi * \psi$. By the usual arguments in Floer theory, these pairs of holomorphic disks form the boundary of a one-dimensional moduli space of holomorphic disks, in the class $\phi * \psi$. Since the coefficients only depend on $\phi * \psi$, it follows that the relevant count of pairs of disks is zero. The contributions from the boundary degenerations cancel out in pairs, because the w's do not appear in our expressions. We conclude that $D^2 = 0$. □

This completes the description of the top-left square in the hyperbox $\mathfrak{C}(\mathcal{H}^{L,-L})$. We now proceed to describe a typical square in the first row of that hyperbox. (The squares in the first column will be described similarly.)

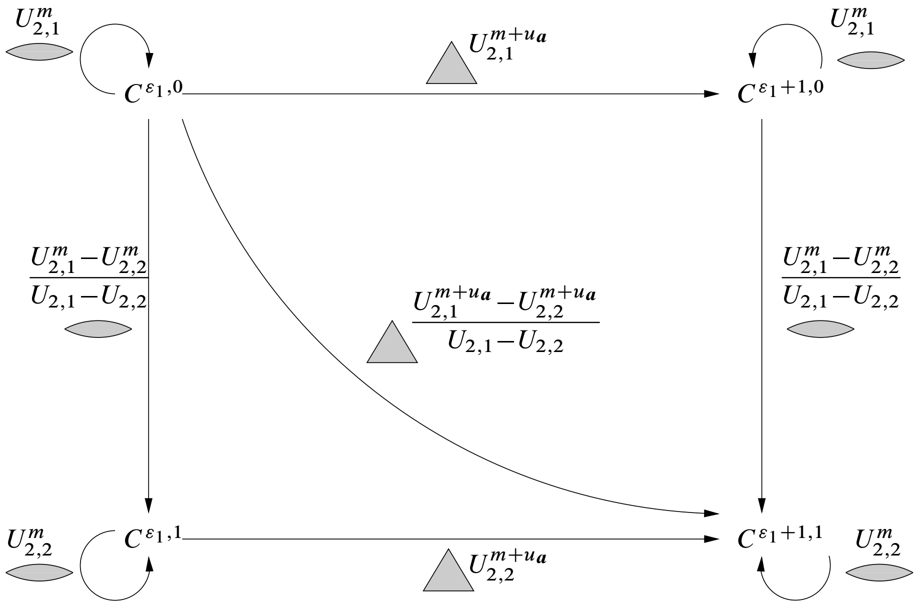

Figure 42: A further square in the first row of the hyperbox $\mathfrak{C}(\mathscr{H}^{L,-L})$. Each arrow is a map counting holomorphic disks or triangles, as indicated by the nearby picture. We also wrote the coefficients that involve the quantities $m = n_{z_{2,2}}(\phi)$ and $u_{\boldsymbol{a}} = u_{2,2}(\boldsymbol{a}) - A_{2,2}(\boldsymbol{a})$.

Consider the square in the first row that starts at $C^{\varepsilon_1,0}$. This square is pictured in Figure 42. It consists of data coming from two strongly equivalent Heegaard diagrams for L_2, namely $\mathscr{H}^{L,-L_1}_{\varepsilon_1}$ and $\mathscr{H}^{L,-L_1}_{\varepsilon_1+1}$. Without loss of generality, let us suppose these differ by a strong equivalence among the beta curves; we denote the sets of curves on $\mathscr{H}^{L,-L_1}_{\varepsilon_1}$ by $(\boldsymbol{\alpha},\boldsymbol{\beta})$, and those on $\mathscr{H}^{L,-L_1}_{\varepsilon_1+1}$ by $(\boldsymbol{\alpha},\boldsymbol{\beta}')$. Moreover, note that both of these diagrams have the same six basepoints: the free basepoints $z_{1,1}$ and $z_{1,2}$, and the four basepoints on L_2.

With this notation, the two complexes on the top edge of Figure 42 are

$$C^{\varepsilon_1,0} = \mathfrak{C}_{\mathrm{int}}(\mathbb{T}_\alpha,\mathbb{T}_\beta), \quad C^{\varepsilon_1+1,0} = \mathfrak{C}_{\mathrm{int}}(\mathbb{T}_\alpha,\mathbb{T}_{\beta'}).$$

The respective differentials count holomorphic disks with coefficients

$$U_{1,1}^{n_{z_{1,1}}(\phi)} U_{1,2}^{n_{z_{1,2}}(\phi)} U_{2,1}^{n_{z_{2,1}}(\phi)+n_{z_{2,2}}(\phi)} V_{2,2}^{n_{w_{2,2}}(\phi)}, \tag{223}$$

and also have terms from tensoring with a mapping cone complex $\mathscr{Y}_{2,2}$ (with a variable $Y_{2,2}$).

By contrast, the complexes on the bottom edge of Figure 42 are

$$C^{\varepsilon_1,1} = \boldsymbol{CF}^-(\mathbb{T}_\alpha,\mathbb{T}_\beta), \quad C^{\varepsilon_1+1,1} = \boldsymbol{CF}^-(\mathbb{T}_\alpha,\mathbb{T}_{\beta'}).$$

They don't involve the basepoints $w_{1,2}$ and $w_{2,2}$, and have no variables $V_{2,2}$ or $Y_{2,2}$. Their differentials count holomorphic disks with coefficients

$$U_{1,1}^{n_{z_{1,1}}(\phi)} U_{1,2}^{n_{z_{1,2}}(\phi)} U_{2,1}^{n_{z_{2,1}}(\phi)} U_{2,2}^{n_{z_{2,2}}(\phi)}. \tag{224}$$

The vertical maps in Figure 42 are transition maps defined by the formulas

$$Y_{2,2}\boldsymbol{x} \mapsto \boldsymbol{x} \tag{225}$$

and

$$\boldsymbol{x} \mapsto \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot U_{1,1}^{n_{z_{1,1}}(\phi)} U_{1,2}^{n_{z_{1,2}}(\phi)} U_{2,1}^{n_{z_{2,1}}(\phi)} \frac{U_{2,1}^{n_{z_{2,2}}(\phi)} - U_{2,2}^{n_{z_{2,2}}(\phi)}}{U_{2,1}-U_{2,2}} \boldsymbol{y}. \tag{226}$$

The horizontal maps in Figure 42 are triangle maps induced by the strong equivalence from $(\boldsymbol{\alpha},\boldsymbol{\beta})$ to $(\boldsymbol{\alpha},\boldsymbol{\beta}')$. They are specified by the choice of a cycle

$$\Theta \in \boldsymbol{A}^-(\mathbb{T}_\beta,\mathbb{T}_{\beta'},\boldsymbol{0})$$

that represents the maximal-degree element in the homology of $\boldsymbol{A}^-(\mathbb{T}_\beta,\mathbb{T}_{\beta'},\boldsymbol{0})$. Note that the diagram formed by the β and β' curves represents an unlink in a connected sum of $(S^1\times S^2)$'s, and this diagram is link-minimal. The diagram contains the basepoints $z_{1,1}$ and $z_{1,2}$ from L_1, and all the four basepoints coming from L_2. The unlink consists of two components, one specified by $w_{2,1}$ and $z_{2,1}$, the other by $w_{2,2}$ and $z_{2,2}$. As such, the intersection points $\boldsymbol{a}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'}$ admit two Alexander gradings $A_{2,1}$ and $A_{2,2}$, whose sum is denoted by A_2.

We can write Θ as a linear combination of intersection points

$$\Theta = \sum_{\boldsymbol{a}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'}} c(\boldsymbol{a}) \prod_{i_1,i_2\in\{1,2\}} U_{i_1,i_2}^{u_{i_1,i_2}(\boldsymbol{a})}\boldsymbol{a}$$

for $c(\boldsymbol{a})\in\{0,1\}$ and some nonnegative integers $u_{i_1,i_2}(\boldsymbol{a})$.

By analogy with the construction of the triangle maps in Section 13.3 — see (212) and (213) — we define the map along the top edge of Figure 42 by

$$\begin{aligned}\boldsymbol{x} \mapsto \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_{\beta'}} \sum_{\boldsymbol{a}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'}} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{a},\boldsymbol{y})\\ \mu(\phi)=0}} &\#\mathcal{M}(\phi)\cdot c(\boldsymbol{a}) U_{1,1}^{n_{z_{1,1}}(\phi)+u_{1,1}(\boldsymbol{a})} U_{1,2}^{n_{z_{1,2}}(\phi)+u_{1,2}(\boldsymbol{a})}\\ &\cdot U_{2,1}^{n_{z_{2,1}}(\phi)+n_{z_{2,2}}(\phi)+u_{2,1}(\boldsymbol{a})+u_{2,2}(\boldsymbol{a})-A_2(\boldsymbol{a})} V_{2,2}^{n_{w_{2,2}}(\phi)+u_{2,2}(\boldsymbol{a})}\boldsymbol{y}.\end{aligned} \tag{227}$$

On the other hand, for the map along the bottom edge, we use the image of Θ under the inclusion into the complex that uses the z basepoints only; see Remark 8.8. Using (25), this element can be written

$$\sum_{\boldsymbol{a}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'}} c(\boldsymbol{a}) \prod_{i\in\{1,2\}} U_{1,i}^{u_{1,i}(\boldsymbol{a})} \prod_{i\in\{1,2\}} U_{2,i}^{u_{2,i}(\boldsymbol{a})-A_{2,i}(\boldsymbol{a})}\boldsymbol{a}.$$

By analogy with (224), we let

$$\begin{aligned}\boldsymbol{x} \mapsto \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_{\beta'}} \sum_{\boldsymbol{a}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'}} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{a},\boldsymbol{y})\\ \mu(\phi)=0}} &\#\mathcal{M}(\phi)\cdot c(\boldsymbol{a}) U_{1,1}^{n_{z_{1,1}}(\phi)+u_{1,1}(\boldsymbol{a})} U_{1,2}^{n_{z_{1,2}}(\phi)+u_{1,2}(\boldsymbol{a})}\\ &\cdot U_{2,1}^{n_{z_{2,1}}(\phi)+u_{2,1}(\boldsymbol{a})-A_{2,1}(\boldsymbol{a})} U_{2,2}^{n_{z_{2,2}}(\phi)+u_{2,2}(\boldsymbol{a})-A_{2,2}(\boldsymbol{a})}\boldsymbol{y}.\end{aligned} \tag{228}$$

Finally, for the chain homotopy along the diagonal of Figure 42, by analogy with (225) and (226), we let

$$Y_{2,2}\boldsymbol{x} \mapsto 0 \tag{229}$$

and

$$\boldsymbol{x} \mapsto \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R})\cdot c(\boldsymbol{a}) U_{1,1}^{n_{z_{1,1}}(\phi)} U_{1,2}^{n_{z_{1,2}}(\phi)} U_{2,1}^{n_{z_{2,1}}(\phi)+u_{2,1}(\boldsymbol{a})-A_{2,1}(\boldsymbol{a})} \cdot \frac{U_{2,1}^{n_{z_{2,2}}(\phi)+u_{2,2}(\boldsymbol{a})-A_{2,2}(\boldsymbol{a})} - U_{2,2}^{n_{z_{2,2}}(\phi)+u_{2,2}(\boldsymbol{a})-A_{2,2}(\boldsymbol{a})}}{U_{2,1}-U_{2,2}}\boldsymbol{y}. \tag{230}$$

Lemma 13.11 *With these definitions, the total differential D on the square pictured in Figure 42 satisfies $D^2=0$.*

Proof Just as in Lemma 13.10, the interesting case to check is when we apply D^2 to a generator $\boldsymbol{x}$ in the initial complex $C^{\varepsilon_1,0}$ (without the $Y_{2,2}$ variable). The result is a sum of counts of pairs consisting of a holomorphic bigon and a holomorphic triangle. There are two cases, according to whether the triangle or the bigon comes first in the composition.

Let us consider the case when we first have a triangle (in a class ϕ), and then a bigon (in a class ψ). Such contributions come from going along the top edge, and then the right edge in Figure 42, and also from going along the diagonal and then using the differential on the final complex $C^{\varepsilon_1+1,1}$. In the coefficients of some $\boldsymbol{y}$ in $D^2\boldsymbol{x}$, we always obtain factors of the form

$$c(\boldsymbol{a}) U_{1,1}^{n_{z_{1,1}}(\sigma)+u_{1,1}(\boldsymbol{a})} U_{1,2}^{n_{z_{1,2}}(\sigma)+u_{1,2}(\boldsymbol{a})} U_{2,1}^{n_{z_{2,1}}(\sigma)+u_{2,1}(\boldsymbol{a})-A_{2,1}(\boldsymbol{a})},$$

where σ is the combined class $\phi*\psi$. These get multiplied by factors having to do with the basepoint $z_{2,2}$, as shown in Figure 42. If we write $m=n_{z_{2,2}}(\phi)$, $m'=n_{z_{2,2}}(\psi)$ and $u_{\boldsymbol{a}}=u_{2,2}(\boldsymbol{a})-A_{2,2}(\boldsymbol{a})$, the sum of the latter factors is

$$U_{2,1}^{m+u_{\boldsymbol{a}}}\frac{U_{2,1}^{m'}-U_{2,2}^{m'}}{U_{2,1}-U_{2,2}} + \frac{U_{2,1}^{m+u_{\boldsymbol{a}}}-U_{2,2}^{m+u_{\boldsymbol{a}}}}{U_{2,1}-U_{2,2}}U_{2,2}^{m'} = \frac{U_{2,1}^{m+m'+u_{\boldsymbol{a}}}-U_{2,2}^{m+m'+u_{\boldsymbol{a}}}}{U_{2,1}-U_{2,2}}.$$

Note that $m+m'=n_{z_{2,2}}(\phi*\psi)$, so the result only depends on the combined class $\sigma=\phi*\psi$.

The case when we first have a bigon and then a triangle is similar, and we get the same coefficients in front of $\boldsymbol{y}$, with dependence only on σ.

Thus, it is natural to consider the one-dimensional moduli space of triangles with vertices at $\boldsymbol{x}$, $\boldsymbol{a}$ and $\boldsymbol{y}$, in the class σ, and weigh it by the same factor

$$c(\boldsymbol{a}) U_{1,1}^{n_{z_{1,1}}(\sigma)+u_{1,1}(\boldsymbol{a})} U_{1,2}^{n_{z_{1,2}}(\sigma)+u_{1,2}(\boldsymbol{a})} U_{2,1}^{n_{z_{2,1}}(\sigma)+u_{2,1}(\boldsymbol{a})-A_{2,1}(\boldsymbol{a})} \cdot \frac{U_{2,1}^{n_{z_{2,2}}(\sigma)+u_{2,2}(\boldsymbol{a})-A_{2,2}(\boldsymbol{a})} - U_{2,2}^{n_{z_{2,2}}(\sigma)+u_{2,2}(\boldsymbol{a})-A_{2,2}(\boldsymbol{a})}}{U_{2,1}-U_{2,2}}. \tag{231}$$

The boundary of this moduli space represents pairs of bigons and triangles, including the two cases described above. The other possibility involves a triangle with vertices at $\boldsymbol{x}, \boldsymbol{b}$ and $\boldsymbol{y}$ (for some $\boldsymbol{b} \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}$), combined with a bigon from $\boldsymbol{a}$ to $\boldsymbol{b}$. However, the fact that Θ is a cycle in $\boldsymbol{A}^-(\mathbb{T}_\beta, \mathbb{T}_{\beta'}, \boldsymbol{0})$ implies that those contributions sum up to zero. It follows that the coefficient of any $\boldsymbol{y}$ in $D^2\boldsymbol{x}$ vanishes. □

This completes the description of the hyperbox $\mathfrak{C}(\mathcal{H}^{L,-L})$. By compressing this hyperbox, we obtain the descent map

$$D^{-L}\colon \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L) \to CF(\mathcal{H}^{\varnothing}).$$

After precomposing this with the inclusion $\mathcal{I}_{\boldsymbol{s}}^{-L}\colon \mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \to \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$, we arrive at the map

$$\Phi_{\boldsymbol{s}}^{-L}\colon \mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \to CF(\mathcal{H}^{\varnothing}),$$

which can be used as a chain homotopy in the surgery complex.

13.6 The general statement

By building up on the particular cases presented in Sections 13.1, 13.4 and 13.5, we are now ready to give the general statement of the link surgery formula, Theorem 1.1. In other words, we plan to define the complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$, where $\mathcal{H}$ is an arbitrary complete system of hyperboxes for the link $\vec{L} \subset Y$, and Λ is a framing. Let $L_1, \dots, L_\ell$ be the components of $\vec{L}$.

The surgery complex $\mathcal{C}^-(\mathcal{H}, \Lambda)$ has roughly the same form as in the link-minimal case presented in Section 9.2:

$$\mathcal{C}^-(\mathcal{H}, \Lambda) = \bigoplus_{M \subseteq L} \prod_{\boldsymbol{s} \in \mathbb{H}(L)} \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s})),$$

with differential

$$\begin{aligned} \mathcal{D}^-(\boldsymbol{s}, \boldsymbol{x}) &= \sum_{N \subseteq L-M} \sum_{\vec{N} \in \Omega(N)} (\boldsymbol{s} + \Lambda_{\vec{L}, \vec{N}}, \Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}(\boldsymbol{x})) \\ &\in \bigoplus_{N \subseteq L-M} \bigoplus_{\vec{N} \in \Omega(N)} \mathfrak{A}^-(\mathcal{H}^{L-M-N}, \psi^{M \cup \vec{N}}(\boldsymbol{s})) \subseteq \mathcal{C}^-(\mathcal{H}, \Lambda). \end{aligned}$$

Our job is to define the complexes $\mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$ and the maps $\Phi^{\vec{N}}_{\psi^M(\boldsymbol{s})}$ in the general setting. It suffices to describe $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ and $\Phi^{\vec{M}}_{\boldsymbol{s}}$; then the same definitions will apply to the complexes for the sublinks of L, and to the maps having these as domains.

To define $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$, we denote the basepoints on the component L_i, in order, by

$$w_{i,1}, z_{i,1}, w_{i,2}, z_{i,2}, \dots, w_{i,p_i}, z_{i,p_i}.$$

We also allow some free basepoints, denoted by

$$w_{\circ,1}, w_{\circ,2}, \dots, w_{\circ,p}.$$

Consider first the dga

$$\mathfrak{R}^Y = \mathbb{F}[\![(U_{i,j})_{\substack{1\le i\le \ell,\\ 1\le j\le p_i}}, (U_{\circ,j})_{1\le j\le p}]\!][(V_{i,j}, Y_{i,j})_{\substack{1\le i\le \ell,\\ 2\le j\le p_i}}]'/(Y_{i,j}^2 = 0, \partial Y_{i,j} = U_{i,j} + U_{i,1}V_{i,j}),$$

where the prime refers to allowing infinite sums of monomials $\boldsymbol{m}$ such that, for each $i = 1, \dots, \ell$,

$$e_{\boldsymbol{m}}(V_{i,2}) + \dots + e_{\boldsymbol{m}}(V_{i,p_i}) - e_{\boldsymbol{m}}(U_{i,1})$$

is bounded above.

We define the complex $\mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$ as the dg module over $\mathfrak{R}^Y$ generated by $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, with differential

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot \prod_{i=1}^{\ell} U_{i,1}^{n_{z_{i,1}}(\phi)+\dots+n_{z_{i,p_i}}(\phi)} V_{i,2}^{n_{w_{i,2}}(\phi)} \cdots V_{i,p_i}^{n_{w_{i,p_i}}(\phi)} \cdot \prod_{j=1}^{p} U_{\circ,j}^{n_{w_{\circ,j}}(\phi)} \boldsymbol{y}. \tag{232}$$

This complex admits filtrations $\mathfrak{F}_i$ for $1 \le i \le \ell$ such that

$$\mathfrak{F}_i(\boldsymbol{x}) = -A_i(\boldsymbol{x}), \quad \mathfrak{F}_i(U_{i,1}) = -1, \quad \mathfrak{F}(V_{i,j}) = 1, \quad j \ge 2,$$

and $\mathfrak{F}_i$ takes all the other U, V and Y variables to zero (ie the action of those variables respects $\mathfrak{F}_i$).

For $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \mathbb{H}(L)$, we define $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ to be the filtered part of $\mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$ given by

$$\mathfrak{F}_i \le s_i, \quad i = 1, \dots, \ell.$$

Let $\mathcal{R}$ be the power series ring freely generated over $\mathbb{F}$ by the U variables (without V's or Y's). The complex $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ admits an $\mathcal{R}$-linear projection $\widetilde{P}$ to $\boldsymbol{A}^-(\mathcal{H}^L, \boldsymbol{s})$, given by

$$\widetilde{P}\biggl(\prod_{i,j} V_{i,j}^{n_{i,j}} Y_{i,j}^{a_{i,j}} \boldsymbol{x}\biggr) = \begin{cases} \prod_{i=1}^{\ell} U_{i,1}^{A_i(\boldsymbol{x})-s_i-\sum_{j=2}^{p_i} n_{i,j}} U_{i,2}^{n_{i,2}} \cdots U_{i,p_i}^{n_{i,p_i}} \boldsymbol{x} & \text{if all } a_{i,j} = 0, \\ 0 & \text{otherwise.} \end{cases}$$

The proofs of Lemmas 13.5 and 13.8 extend to our setting, showing that $\widetilde{P}$ exhibits $\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s})$ as a resolution of $\boldsymbol{A}^-(\mathcal{H}^L, \boldsymbol{s})$, over the ring $\mathcal{R}$.

Next, let $M \subseteq L$ a sublink, with an orientation $\vec{M}$. As usual, $I_+(\vec{L}, \vec{M})$ (resp. $I_-(\vec{L}, \vec{M})$) denotes the set of indices i such that the component L_i is in M and its orientation induced from $\vec{L}$ is the same as (resp. opposite to) the one induced from $\vec{M}$. Let

$$I(\vec{L}, \vec{M}) = I_+(\vec{L}, \vec{M}) \cup I_-(\vec{L}, \vec{M}).$$

Also, as in Section 8.6, we set

$$M_\pm = \bigcup_{i \in I_\pm(\vec{L}, \vec{M})} L_i,$$

so that $M = M_+ \amalg M_-$.

We seek to define maps

$$\Phi_{\boldsymbol{s}}^{\vec{M}} : \mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \to \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^{\vec{M}}(\boldsymbol{s})).$$

Similarly to what we did in Section 9, these maps will be constructed as compositions

$$\Phi_{s}^{\vec{M}} = D_{p^{\vec{M}}(s)}^{\vec{M}} \circ \mathcal{I}_{s}^{\vec{M}},$$

where $\mathcal{I}_{s}^{\vec{M}}$ is a projection-inclusion and $D_{p^{\vec{M}}(s)}^{\vec{M}}$ is a descent map; compare (189) and (190).

Let us first define the map

$$\mathcal{I}_{s}^{\vec{M}} : \mathfrak{A}^-(\mathcal{H}^L, s) \to \mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(s)).$$

Here, the target $\mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(s))$ is a complex that looks like a resolution $\mathfrak{A}^-$ with respect to $L - M$, like $\boldsymbol{CF}^- = C\{w\}$ with respect to M_+, and like $\mathfrak{C}_{\text{int}}$ with respect to M_-. Precisely, we first consider $\mathfrak{C}_{\text{int}}(r_{M_+}(\mathcal{H}^L))$, a complex constructed similarly to $\mathfrak{C}_{\text{int}}(\mathcal{H}^L)$, except that for $i \in I_+(\vec{L}, \vec{M})$ we do not introduce variables $V_{i,j}$ and $Y_{i,j}$; instead, for such i, we use $n_{w_{i,j}}(\phi)$ as exponents for $U_{i,j}$ in the formula for the differential:

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi) = 1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot \prod_{i \notin I_+(\vec{L}, \vec{M})} U_{i,1}^{n_{z_{i,1}}(\phi) + \cdots + n_{z_{i,p_i}}(\phi)} V_{i,2}^{n_{w_{i,2}}(\phi)} \cdots V_{i,p_i}^{n_{w_{i,p_i}}(\phi)} \cdot \prod_{i \in I_+(\vec{L}, \vec{M})} U_{i,1}^{n_{w_{i,1}}(\phi)} \cdots U_{i,p_i}^{n_{w_{i,p_i}}(\phi)} \prod_{j=1}^{p} U_{\circ,j}^{n_{w_{\circ,j}}(\phi)} \boldsymbol{y}. \tag{233}$$

The complex $\mathfrak{C}_{\text{int}}(r_{M_+}(\mathcal{H}^L))$ has filtrations $\mathfrak{F}_i$ for $i \notin I_+(\vec{L}, \vec{M})$, where we use the convention that the Alexander gradings A_i of $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ are the ones from the diagram $\mathcal{H}^L$ (rather than those from $r_{M_+}(\mathcal{H}^L)$, where s becomes $\psi^{M_+}(s)$).

We let $\mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(s))$ be the subcomplex of $\mathfrak{C}_{\text{int}}(r_{M_+}(\mathcal{H}^L))$ in filtration degrees

$$\mathfrak{F}_i \le s_i \quad \text{for } L_i \not\subseteq M.$$

The map $\mathcal{I}_{s}^{\vec{M}}$ is the composition of a projection similar to $\widetilde{P}$, but taken only with respect to the indices $i \in I_+(\vec{L}, \vec{M})$, and the natural inclusion into $\mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(s))$. Specifically, we set

$$\mathcal{I}_{s}^{\vec{M}}\Bigl(\prod_{i \in I_+(\vec{L}, \vec{M})} \prod_j V_{i,j}^{n_{i,j}} Y_{i,j}^{a_{i,j}} \boldsymbol{x}\Bigr) = \begin{cases} \prod_{i \in I_+(\vec{L}, \vec{M})} U_{i,1}^{A_i(\boldsymbol{x}) - s_i - \sum_{j=2}^{p_i} n_{i,j}} U_{i,2}^{n_{i,2}} \cdots U_{i,p_i}^{n_{i,p_i}} \boldsymbol{x} & \text{if all } a_{i,j} = 0 \text{ for } i \in I_+(\vec{L}, \vec{M}), \\ 0 & \text{otherwise.} \end{cases} \tag{234}$$

We then extend $\mathcal{I}_{s}^{\vec{M}}$ to be equivariant with respect to the action of the variables $V_{i,j}$ and $Y_{i,j}$ for $i \notin I_+(\vec{L}, \vec{M})$, as well as to that of all the U variables.

We are left to describe the descent map

$$D_{p^{\vec{M}}(s)}^{\vec{M}} : \mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(s)) \to \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^{\vec{M}}(s)). \tag{235}$$

This is associated to the hyperbox $\mathscr{H}^{\vec{L},\vec{M}}$ in the system $\mathscr{H}$. Recall that the hyperbox $\mathscr{H}^{\vec{L},\vec{M}}$ consists of diagrams for $L-M$ such that the initial diagram is $r_{\vec{M}}(\mathscr{H}^L)$ and the final one is $\mathscr{H}^{\vec{L},\vec{M}}$, which is surface isotopic to $\mathscr{H}^{L-M}$. The dimension of $\mathscr{H}^{\vec{L},\vec{M}}$ is the number of components of M, denoted by m; we label its coordinate directions by the indices $i \in \{1,\ldots,\ell\}$ such that $L_i \subseteq M$, and order them accordingly.

Let d_i be the side length of the hyperbox $\mathscr{H}^{\vec{L},\vec{M}}$ in the direction i. To construct the descent map $D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$, we will compress a hyperbox of chain complexes associated to $\mathscr{H}^{\vec{L},\vec{M}}$, of the same dimension m but possibly of larger size. The models are

- the precomposition with the $p-1$ transition maps Ξ_j in Section 13.4, to construct the descent map for a knot with $2p$ basepoints of each type;
- the construction of the hyperbox $\mathfrak{C}(\mathscr{H}^{L,-L})$ in Section 13.5, where we added one extra row and column at the beginning; see Figure 40.

In the general setting, we construct an m-dimensional hyperbox of chain complexes

$$C = \mathfrak{C}(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$$

whose sides are labeled by $i \in I(\vec{L},\vec{M})$, and such that the side length in direction i is

$$d_i' = \begin{cases} d_i & \text{if } i \in I_+(\vec{L},\vec{M}), \\ d_i + p_i - 1 & \text{if } i \in I_-(\vec{L},\vec{M}). \end{cases} \tag{236}$$

(Recall that p_i denotes the number of basepoints of each type on the component L_i.) We let $\boldsymbol{d}'$ be the vector with d_i' as entries.

At a vertex $\varepsilon = (\varepsilon_i)_{i \in I(\vec{L},\vec{M})}$ in the hyperbox, we place a complex C^{ε} defined as follows. Let

$$I(\varepsilon) = \{ i \in I_-(\vec{L},\vec{M}) \mid \varepsilon_i < p_i - 1\}$$

and

$$M_\varepsilon = \bigcup_{i \in I(\varepsilon)} L_i \subseteq M_-.$$

(For example, in Section 13.5, the part of the hyperbox $\mathfrak{C}(\mathscr{H}^{L,-L})$ corresponding to ε such that $I(\varepsilon) = \varnothing$ is the dashed rectangle in Figure 40.)

We split the information in ε into two parts, $\varepsilon^<$ and $\varepsilon^>$, corresponding to $i \in I(\varepsilon)$ and $i \notin I(\varepsilon)$, respectively. Specifically, $\varepsilon^<$ is a vector, whose entries are indexed by $i \in I(\varepsilon)$, such that

$$\varepsilon_i^< = \varepsilon_i.$$

The second vector, $\varepsilon^>$, has entries indexed by $i \in I(\vec{L},\vec{M}) - I(\varepsilon)$ and is such that

$$\varepsilon_i^> = \begin{cases} \varepsilon_i & \text{if } i \in I_+(\vec{L},\vec{M}), \\ \varepsilon_i - (p_i - 1) & \text{if } i \in I_-(\vec{L},\vec{M}) - I(\varepsilon). \end{cases}$$

In the second case, we subtract $p_i - 1$ so that we are in agreement with the indexing of the entries in the hyperbox $\mathscr{H}^{\vec{L},\vec{M}}$. Indeed, note that the values of $\varepsilon_i^>$ vary between 0 and d_i.

Consider the hyperbox $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}$, whose reduction at M_ε is identified with a subhyperbox of $\mathscr{H}^{\vec{L},\vec{M}}$; see Definitions 8.22 and 8.23. The complex C^ε will be associated to the diagram at position $\varepsilon^>$ in the hyperbox $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}$. We write

$$C^\varepsilon = \mathfrak{C}_{\varepsilon^<}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{\varepsilon^>}, M_\varepsilon, \psi^{\vec{M}}(\boldsymbol{s})). \tag{237}$$

This is a Floer complex whose construction is a mix of those in Sections 13.4 and 13.5. Roughly speaking, C^ε looks like

- the resolution $\mathfrak{A}^-(\,\cdot\,, \psi^{\vec{M}}(\boldsymbol{s}))$ in the directions i such that $L_i \subseteq L - M$, ie $i \notin I(\vec{L}, \vec{M})$;
- the complexes $\mathfrak{C}_j$ from (217), with $j = \varepsilon_i$, in the directions $i \in I(\varepsilon)$;
- the complex $\boldsymbol{CF}^- = C\{z\}$ in the directions $i \in I_-(\vec{L}, \vec{M}) - I(\varepsilon)$;
- the complex $\boldsymbol{CF}^- = C\{w\}$ in the directions $i \in I_+(\vec{L}, \vec{M})$.

We will construct C^ε as a subcomplex of a complex denoted by

$$\mathfrak{C}_{\varepsilon^<}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{\varepsilon^>}, M_\varepsilon),$$

which is similar, except it looks like $\mathfrak{C}_{\mathrm{int}}$ in the directions $i \notin I(\vec{L}, \vec{M})$.

Let us make these definitions fully rigorous. First, recall that we have a base ring

$$\mathscr{R} = \mathbb{F}[\![(U_{i,j})_{1\le i\le \ell, 1\le j\le p_i}(U_{\circ,j})_{1\le j\le p}]\!].$$

We introduce the dga

$$\mathfrak{R}^Y_{\varepsilon^<}(\vec{L}, \vec{M}, M_\varepsilon)$$

generated over $\mathscr{R}$ (in the sense of infinite sums with boundedness conditions) by new variables $V_{i,j}$ and $Y_{i,j}$, for indices i and j such that

$$\big(i \notin I(\vec{L}, \vec{M}),\ 2 \le j \le p_i\big) \quad \text{or} \quad \big(i \in I(\varepsilon),\ \varepsilon_i + 2 \le j \le p_i\big) \tag{238}$$

with relations

$$Y_{i,j}^2 = 0, \quad \partial Y_{i,j} = U_{i,j} + U_{i,1}V_{i,j}.$$

We let $\mathfrak{C}_{\varepsilon^<}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{\varepsilon^>}, M_\varepsilon)$ be the complex freely generated over $\mathfrak{R}^Y_{\varepsilon^<}(\vec{L}, \vec{M}, M_\varepsilon)$ by $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, with differential

$$\begin{aligned}\partial \boldsymbol{x} = \sum_{\boldsymbol{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\boldsymbol{x},\boldsymbol{y})\\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) \cdot \prod_{i\notin I(\vec{L},\vec{M})} U_{i,1}^{n_{z_{i,1}}(\phi)+\cdots+n_{z_{i,p_i}}(\phi)} V_{i,2}^{n_{w_{i,2}}(\phi)} \cdots V_{i,p_i}^{n_{w_{i,p_i}}(\phi)} \\ \cdot \prod_{i\in I(\varepsilon)} U_{i,1}^{n_{z_{i,1}}(\phi)} \cdots U_{i,\varepsilon_i}^{n_{z_{i,\varepsilon_i}}(\phi)} U_{i,\varepsilon_i+1}^{n_{z_{i,\varepsilon_i+1}}(\phi)+\cdots+n_{z_{i,p_i}}(\phi)} V_{i,\varepsilon_i+2}^{n_{w_{i,\varepsilon_i+2}}(\phi)} \cdots V_{i,p_i}^{n_{w_{i,p_i}}(\phi)} \\ \cdot \prod_{i\in I_-(\vec{L},\vec{M})-I(\varepsilon)} U_{i,1}^{n_{z_{i,1}}(\phi)} \cdots U_{i,p_i}^{n_{z_{i,p_i}}(\phi)} \prod_{i\in I_+(\vec{L},\vec{M})} U_{i,1}^{n_{w_{i,1}}(\phi)} \cdots U_{i,p_i}^{n_{w_{i,p_i}}(\phi)} \prod_{j=1}^{p} U_{\circ,j}^{n_{w_{\circ,j}}(\phi)} \boldsymbol{y}.\end{aligned} \tag{239}$$

This complex admits filtrations $\mathfrak{F}_i$ for $i \notin I(\vec{L}, \vec{M})$ such that

$$\mathfrak{F}_i(\boldsymbol{x}) = -A_i(\boldsymbol{x}), \quad \mathfrak{F}_i(U_{i,1}) = -1, \quad \mathfrak{F}_i(V_{i,j}) = 1, \quad j \geq 2,$$

and $\mathfrak{F}_i$ takes all the other U, V and Y variables to zero. (Once again, the Alexander gradings A_i are taken to be the ones induced from the diagram $\mathscr{H}^L$.)

We then define $C^{\varepsilon} = \mathfrak{C}_{\varepsilon^<}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{\varepsilon^>}, M_\varepsilon, \psi^{\vec{M}}(\boldsymbol{s}))$ to be the subcomplex of $\mathfrak{C}_{\varepsilon^<}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{\varepsilon^>}, M_\varepsilon)$ given by

$$\mathfrak{F}_i \leq s_i, \quad i \notin I(\vec{L}, \vec{M}).$$

In particular, when $\varepsilon = \boldsymbol{0}$, observe that $M_{\boldsymbol{0}} = M_-$, $I(\boldsymbol{0}) = I_-(\vec{L}, \vec{M})$ and (by the compatibility conditions between hyperboxes) $\mathscr{H}^{\vec{L},\vec{M}-M_{\boldsymbol{0}}}_{\boldsymbol{0}} = r_{M+}(\mathscr{H}^L)$. Therefore, we have identifications

$$\mathfrak{C}_{\boldsymbol{0}}(\mathscr{H}^{\vec{L},\vec{M}-M_{\boldsymbol{0}}}_{\boldsymbol{0}}, M_{\boldsymbol{0}}) \cong \mathfrak{C}_{\mathrm{int}}(r_{M_+}(\mathscr{H}^L))$$

and

$$C^{\boldsymbol{0}} = \mathfrak{C}_{\boldsymbol{0}}(\mathscr{H}^{\vec{L},\vec{M}-M_{\boldsymbol{0}}}_{\boldsymbol{0}}, M_{\boldsymbol{0}}, \psi^{\vec{M}}(\boldsymbol{s})) \cong \mathfrak{A}^-(\mathscr{H}^L, \vec{M}, p^{\vec{M}}(\boldsymbol{s})). \tag{240}$$

At the other extreme, for ε such that $I(\varepsilon) = \varnothing$, we have $M_\varepsilon = \varnothing$ and

$$C^{\varepsilon} \cong \mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}_{\varepsilon^>}, \psi^{\vec{M}}(\boldsymbol{s})).$$

Thus, the final complex in the hyperbox $\mathfrak{C}(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$, with $\varepsilon = \boldsymbol{d}'$, is

$$C^{\boldsymbol{d}'} \cong \mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}(M), \psi^{\vec{M}}(\boldsymbol{s})). \tag{241}$$

For general ε and ε' such that $\varepsilon \leq \varepsilon'$ and $\varepsilon, \varepsilon'$ are neighbors, we now proceed to define the maps

$$D^{\varepsilon'-\varepsilon}_{\varepsilon} \colon C^{\varepsilon} \to C^{\varepsilon'}$$

that are part of the hyperbox $\mathfrak{C}(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$.

Since ε and ε' are neighbors, for any i we must have $\varepsilon'_i = \varepsilon_i$ or $\varepsilon'_i = \varepsilon_i + 1$. Observe that $I(\varepsilon') \subseteq I(\varepsilon)$. When an index i belongs to $I(\varepsilon) - I(\varepsilon')$, it means that $\varepsilon_i = p_i - 2$ and $\varepsilon'_i = p_i - 1$, so $(\varepsilon')^>$ has 0 in position i. Thus, while $(\varepsilon')^>$ may have more entries that $\varepsilon^>$, all these additional entries are zero; let us denote by $(\varepsilon^>)'$ the vector obtained from $(\varepsilon')^>$ by deleting these entries. In the same manner, $(\varepsilon')^<$ may have fewer entries than $\varepsilon^<$, in which case the additional entries of $\varepsilon^<$ are $p_i - 2$; we let $(\varepsilon^<)'$ be the vector obtained from $(\varepsilon^<)'$ by introducing new entries of $p_i - 1$ in those positions. Thus, $(\varepsilon^>)'$ has the same length as $\varepsilon^>$, and $(\varepsilon^<)'$ has the same length as $\varepsilon^<$.

Note that (by the compatibility between hyperboxes) $\mathscr{H}^{\vec{L},\vec{M}-M_{\varepsilon'}}_{(\varepsilon')^>}$, the diagram used to define $C^{\varepsilon'}$, is a reduction of a diagram in the hyperbox $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}$:

$$\mathscr{H}^{\vec{L},\vec{M}-M_{\varepsilon'}}_{(\varepsilon')^>} = r_{M_\varepsilon - M_{\varepsilon'}}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{(\varepsilon^>)'}).$$

Thus, the complex $C^{\varepsilon'}$ can be interpreted as being associated to the diagram $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{(\varepsilon^>)'}$. Indeed, from the definitions,

$$\begin{aligned} C^{\varepsilon'} &= \mathfrak{C}_{(\varepsilon')^<}(\mathscr{H}^{\vec{L},\vec{M}-M_{\varepsilon'}}_{(\varepsilon')^>}, M_{\varepsilon'}, \psi^{\vec{M}}(\boldsymbol{s})) \\ &= \mathfrak{C}_{(\varepsilon')^<}(r_{M_\varepsilon - M_{\varepsilon'}}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{(\varepsilon^>)'}), M_{\varepsilon'}, \psi^{\vec{M}}(\boldsymbol{s})) \\ &= \mathfrak{C}_{(\varepsilon^<)'}(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_{(\varepsilon^>)'}, M_\varepsilon, \psi^{\vec{M}}(\boldsymbol{s})). \end{aligned}$$

Similar remarks apply to the intermediate indices between ε and ε'. Specifically, we can write the corresponding complexes as

$$\mathfrak{C}_\nu(\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}_\eta, M_\varepsilon, \psi^{\vec{M}}(\boldsymbol{s}))$$

for all ν and η with

$$\varepsilon^< \le \nu \le (\varepsilon^<)', \quad \varepsilon^> \le \eta \le (\varepsilon^>)'.$$

These complexes have almost the same definition as the one for index ε, except that changing $\varepsilon^>$ into η results in a change in the underlying Heegaard diagram (a different vertex in the hyperbox $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}$), whereas changing $\varepsilon^<$ into ν results in a change in the exponents of $U_{i,j}$ for some $i \in I(\varepsilon)$ in (239); namely, ε_i gets replaced by ν_i, which may equal $\varepsilon_i + 1$.

In light of this discussion, to define $D^{\varepsilon'-\varepsilon}_\varepsilon$ we will focus on the hyperbox

$$\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon},$$

and specifically on the subhypercube in that hyperbox corresponding to indices η such that

$$\varepsilon^> \le \eta \le (\varepsilon^>)'.$$

We denote this hypercube by

$$\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}(\varepsilon^>, (\varepsilon^>)').$$

This is a υ-dimensional hypercube of strongly equivalent Heegaard diagrams, with

$$\upsilon = \|(\varepsilon^>)' - \varepsilon^>\|.$$

Recall that, in the link-minimal case, the maps $D^{\varepsilon'-\varepsilon}_\varepsilon$ were defined in Section 8.3 as counts of holomorphic k-gons, where $k \le (\|\varepsilon' - \varepsilon\| + 2)$. In our setting, they will be given by counts of holomorphic k-gons for $k \le \upsilon + 2$, coming from the above hypercube.

The formula in the link-minimal case was

$$D^{\varepsilon'-\varepsilon}_\varepsilon(\boldsymbol{x}) = \sum_{l,q} \sum_{\{\varepsilon'^\alpha = \gamma^0 > \cdots > \gamma^l = \varepsilon^\alpha\}} \sum_{\{\varepsilon^\beta = \zeta^0 < \cdots < \zeta^q = \varepsilon'^\beta\}} f(\Theta^\alpha_{\gamma^0,\gamma^1} \otimes \cdots \otimes \Theta^\alpha_{\gamma^{l-1},\gamma^l} \otimes \boldsymbol{x} \otimes \Theta^\beta_{\zeta^0,\zeta^1} \otimes \cdots \otimes \Theta^\beta_{\zeta^{q-1},\zeta^q}).$$

In general, our hypercube still has bipartition maps (see Section 8.2), so we can talk about indices

$$(\varepsilon^>)^\alpha \le (\varepsilon^>)'^\alpha, \quad (\varepsilon^>)^\beta \le (\varepsilon^>)'^\beta.$$

For any intermediate γ between (and including) $(\varepsilon^>)^\alpha$ and $(\varepsilon^>)'^\alpha$, as part of the hyperbox we have a collection of curves $\boldsymbol{\alpha}^\gamma$ such that all these collections are strongly equivalent. In particular, we write $\boldsymbol{\alpha} = \boldsymbol{\alpha}^{(\varepsilon^>)^\alpha}$ and $\boldsymbol{\alpha}' = \boldsymbol{\alpha}^{(\varepsilon^>)'^\alpha}$. Furthermore, for any two intermediate indices γ and γ' with

$$(\varepsilon^>)^\alpha \le \gamma \le \gamma' \le (\varepsilon^>)'^\alpha,$$

we have a chain element

$$\Theta^\alpha_{\gamma',\gamma} \in \boldsymbol{A}^-(\mathbb{T}_{\alpha^{\gamma'}}, \mathbb{T}_{\alpha^\gamma}, \mathbf{0}).$$

We can write this element explicitly in terms of intersection points as

$$\Theta^\alpha_{\gamma',\gamma} = \sum_{\boldsymbol{a}^\alpha_{\gamma'\gamma} \in \mathbb{T}_{\alpha^{\gamma'}} \cap \mathbb{T}_{\alpha^\gamma}} c(\boldsymbol{a}^\alpha_{\gamma'\gamma}) \prod_{i,j} U_{i,j}^{u_{i,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma})} \prod_j U_{\circ,j}^{u_{\circ,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma})} \boldsymbol{a}^\alpha_{\gamma'\gamma},$$

with coefficients $c(\boldsymbol{a}^\alpha_{\gamma'\gamma}) \in \{0,1\}$ and exponents $u_{i,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma}), u_{\circ,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma}) \ge 0$.

Note that $\Theta^\alpha_{\gamma',\gamma}$ are part of the data in the hyperbox $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}$, where the $U_{i,j}$ basepoints keep track of the basepoints $z_{i,j}$ for $i \in I_-(\vec{L},\vec{M}) - I(\varepsilon)$, and of the basepoints $w_{i,j}$ for other i.

Further, the $\alpha^{\gamma'}$–α^γ diagram is link-minimal, and therefore the intersection points $\boldsymbol{a}^\alpha_{\gamma'\gamma}$ admit Alexander gradings $A_{i,j}$ for $L_i \not\subseteq M - M_\varepsilon$, that is, for $i \notin I(\vec{L},\vec{M})$ as well as for $i \in I(\varepsilon)$, and for all corresponding j. We will need to pull back the Θ elements as in Section 13.3, so that the corresponding exponents are in terms of the z basepoints rather than the w basepoints. It helps to introduce the notation

$$\bar{u}_{i,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma}) := u_{i,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma}) - A_{i,j}(\boldsymbol{a}^\alpha_{\gamma'\gamma}).$$

The hypercube $\mathscr{H}^{\vec{L},\vec{M}-M_\varepsilon}(\varepsilon^>, (\varepsilon^>)')$ also contains collections of beta curves and corresponding Θ elements. For ζ between $(\varepsilon^>)^\beta$ and $(\varepsilon^>)'^\beta$, we have a collection of curves $\boldsymbol{\beta}^\zeta$, and for

$$(\varepsilon^>)^\beta \le \zeta \le \zeta' \le (\varepsilon^>)'^\beta$$

we have a chain element

$$\Theta^\beta_{\zeta,\zeta'} \in \boldsymbol{A}^-(\mathbb{T}_{\beta^\zeta}, \mathbb{T}_{\beta^{\zeta'}}, \mathbf{0})$$

of the form

$$\Theta^\beta_{\zeta,\zeta'} = \sum_{\boldsymbol{a}^\beta_{\zeta\zeta'} \in \mathbb{T}_{\beta^\zeta} \cap \mathbb{T}_{\beta^{\zeta'}}} c(\boldsymbol{a}^\beta_{\zeta\zeta'}) \prod_{i,j} U_{i,j}^{u_{i,j}(\boldsymbol{a}^\beta_{\zeta\zeta'})} \prod_j U_{\circ,j}^{u_{\circ,j}(\boldsymbol{a}^\beta_{\zeta\zeta'})} \boldsymbol{a}^\beta_{\zeta\zeta'}.$$

Similar remarks to those for the alpha curves apply here. In particular, we set

$$\bar{u}_{i,j}(\boldsymbol{a}^\beta_{\gamma'\gamma}) := u_{i,j}(\boldsymbol{a}^\beta_{\zeta\zeta'}) - A_{i,j}(\boldsymbol{a}^\beta_{\zeta\zeta'}).$$

The map $D^{\varepsilon'-\varepsilon}_\varepsilon \colon C^\varepsilon \to C^{\varepsilon'}$ are defined as follows. For an intersection point $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we set

$$D^{\varepsilon'-\varepsilon}_\varepsilon(\boldsymbol{x}) = \sum_E \#\mathcal{M}(\phi) \cdot \boldsymbol{c}(E) \boldsymbol{U}^{u(E)} \boldsymbol{y}, \tag{242}$$

where the summation is over data E consisting of

- $l, q \ge 0$,
- $(\varepsilon^>)'^\alpha = \gamma^0 > \cdots > \gamma^l = (\varepsilon^>)^\alpha$,

- $(\varepsilon^>)^\beta = \zeta^0 < \cdots < \zeta^q = (\varepsilon^>)'^\beta$,
- $\boldsymbol{a}^\alpha_{\gamma^k\gamma^{k+1}} \in \mathbb{T}_{\alpha_{\gamma^k}} \cap \mathbb{T}_{\alpha_{\gamma^{k+1}}}$ for $k = 0, \ldots, l-1$,
- $\boldsymbol{a}^\beta_{\zeta^k\zeta^{k+1}} \in \mathbb{T}_{\beta_{\zeta^k}} \cap \mathbb{T}_{\beta_{\zeta^{k+1}}}$ for $k = 0, \ldots, q-1$,
- $\boldsymbol{y} \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_{\beta'}$,
- $\phi \in \pi_2(\boldsymbol{a}^\alpha_{\gamma^0\gamma^1}, \ldots, \boldsymbol{a}^\alpha_{\gamma^{l-1}\gamma^l}, \boldsymbol{x}, \boldsymbol{a}^\beta_{\zeta^0\zeta^1}, \ldots, \boldsymbol{a}^\beta_{\zeta^{q-1}\zeta^q}, \boldsymbol{y})$ with $\mu(\phi) = 1 - l - q$.

In (242), by $\#\mathcal{M}(\phi)$ we mean the count of holomorphic polygons in the class ϕ (where we divide by the action of $\mathbb{R}$ if we count bigons). Moreover, the coefficient $\boldsymbol{c}(E) \in \{0, 1\}$ is

$$\boldsymbol{c}(E) := \prod_{i,j} c(\boldsymbol{a}^\alpha_{\gamma^0\gamma^1}) \cdots c(\boldsymbol{a}^\alpha_{\gamma^{l-1}\gamma^l}) \cdot c(\boldsymbol{a}^\beta_{\zeta^0\zeta^1}) \cdots c(\boldsymbol{a}^\beta_{\zeta^{q-1}\zeta^q})$$

whereas the factor $\boldsymbol{U}^{u(E)}$ is a product of powers of $U_{i,j}$ that will be described soon. Let us first introduce the shorthand notation

$$\begin{aligned}
n^{\boldsymbol{a}}_{w_{i,j}}(\phi) &:= n_{w_{i,j}}(\phi) + \sum_{k=0}^{l-1} u_{i,j}(\boldsymbol{a}^\alpha_{\gamma^k\gamma^{k+1}}) + \sum_{k=0}^{q-1} u_{i,j}(\boldsymbol{a}^\beta_{\zeta^k\zeta^{k+1}}),\\
n^{\boldsymbol{a}}_{z_{i,j}}(\phi) &:= n_{z_{i,j}}(\phi) + \sum_{k=0}^{l-1} \bar{u}_{i,j}(\boldsymbol{a}^\alpha_{\gamma^k\gamma^{k+1}}) + \sum_{k=0}^{q-1} \bar{u}_{i,j}(\boldsymbol{a}^\beta_{\zeta^k\zeta^{k+1}}),\\
n^{\boldsymbol{a}}_{w_{\circ,j}}(\phi) &:= n_{w_{\circ,j}}(\phi) + \sum_{k=0}^{l-1} u_{\circ,j}(\boldsymbol{a}^\alpha_{\gamma^k\gamma^{k+1}}) + \sum_{k=0}^{q-1} u_{\circ,j}(\boldsymbol{a}^\beta_{\zeta^k\zeta^{k+1}}).
\end{aligned}$$

Then, by analogy with the U-powers in the differentials (219), (230) and (239), we set

$$\begin{aligned}
(243)\quad \boldsymbol{U}^{u(E)} := &\prod_{i \notin I(\vec{L},\vec{M})} U_{i,1}^{n^{\boldsymbol{a}}_{z_{i,1}}(\phi)+\cdots+n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)} V_{i,2}^{n^{\boldsymbol{a}}_{w_{i,2}}(\phi)} \cdots V_{i,p_i}^{n_{w_{i,p_i}}(\phi)}\\
&\cdot \prod_{\substack{i \in I(\varepsilon)\\ \nu_i = \varepsilon_i}} U_{i,1}^{n^{\boldsymbol{a}}_{z_{i,1}}(\phi)} \cdots U_{i,\varepsilon_i}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i}}(\phi)} U_{i,\varepsilon_i+1}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+1}}(\phi)+\cdots+n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)} V_{i,\varepsilon_i+2}^{n^{\boldsymbol{a}}_{w_{i,\varepsilon_i+2}}(\phi)} \cdots V_{i,p_i}^{n^{\boldsymbol{a}}_{w_{i,p_i}}(\phi)}\\
&\cdot \prod_{\substack{i \in I(\varepsilon)\\ \nu_i = \varepsilon_i+1}} U_{i,1}^{n^{\boldsymbol{a}}_{z_{i,1}}(\phi)} \cdots U_{i,\varepsilon_i+1}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+1}}(\phi)} \frac{U_{i,\varepsilon_i+1}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+2}}(\phi)+\cdots+n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)} - U_{i,\varepsilon_i+2}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+2}}(\phi)+\cdots+n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)}}{U_{i,\varepsilon_i+1} - U_{i,\varepsilon_i+2}}\\
&\qquad\qquad\cdot V_{i,\varepsilon_i+3}^{n^{\boldsymbol{a}}_{w_{i,\varepsilon_i+3}}(\phi)} \cdots V_{i,p_i}^{n^{\boldsymbol{a}}_{w_{i,p_i}}(\phi)}\\
&\cdot \prod_{i \in I_-(\vec{L},\vec{M}) - I(\varepsilon)} U_{i,1}^{n^{\boldsymbol{a}}_{z_{i,1}}(\phi)} \cdots U_{i,p_i}^{n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)} \prod_{i \in I_+(\vec{L},\vec{M})} U_{i,1}^{n^{\boldsymbol{a}}_{w_{i,1}}(\phi)} \cdots U_{i,p_i}^{n^{\boldsymbol{a}}_{w_{i,p_i}}(\phi)} \prod_{j=1}^{p} U_{\circ,j}^{n^{\boldsymbol{a}}_{w_{\circ,j}}(\phi)}.
\end{aligned}$$

We have described the values of $D_\varepsilon^{\varepsilon'-\varepsilon}$ at intersection points $\boldsymbol{x}$. To define the map in full, we extend it to be equivariant with respect to the action of all the $U_{i,j}$, $V_{i,j}$, $Y_{i,j}$ and $U_{\circ,j}$ variables that appear in the construction of the target $C^{\varepsilon'}$. In view of the conditions (238), note that the domain C^{ε} may have a few more variables, namely

$$V_{i,\varepsilon_i+2}, Y_{i,\varepsilon_i+2} \quad \text{for } i \in I(\varepsilon) \text{ such that } \varepsilon'_i = \varepsilon_i + 1.$$

We let $D^{\varepsilon'-\varepsilon}_{\varepsilon}$ be equivariant with respect to V_{i,ε_i+2}, where we let this variable act on the target by 1. With regard to Y_{i,ε_i+2}, this behaves somewhat like in an identity hyperbox, with identity maps on the edges and zeros on the diagonals. Precisely, we set

$$D^{\varepsilon'-\varepsilon}_{\varepsilon}(Y_{i,\varepsilon_i+2}\boldsymbol{x}) = 0 \quad \text{when } \|\varepsilon' - \varepsilon\| > 1,$$

whereas when $\|\varepsilon' - \varepsilon\| = 1$ there is a unique i with $\varepsilon'_i = \varepsilon_i + 1$, and (for the extra variable Y to exist) we must have $i \in I(\varepsilon)$; that is,

$$(\varepsilon^<)'_i = \varepsilon^<_i + 1, \quad (\varepsilon^>)'_i = \varepsilon^>_i.$$

In this particular case and for that particular i, by analogy with (218) and (219), the formula for $D^{\varepsilon'-\varepsilon}_{\varepsilon}(Y_{i,\varepsilon_i+2}\boldsymbol{x})$ differs from the one for $D^{\varepsilon'-\varepsilon}_{\varepsilon}(\boldsymbol{x})$ by replacing

$$U_{i,1}^{n^{\boldsymbol{a}}_{z_{i,1}}(\phi)} \cdots U_{i,\varepsilon_i+1}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+1}}(\phi)} \frac{U_{i,\varepsilon_i+1}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+2}}(\phi)+\cdots+n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)} - U_{i,\varepsilon_i+2}^{n^{\boldsymbol{a}}_{z_{i,\varepsilon_i+2}}(\phi)+\cdots+n^{\boldsymbol{a}}_{z_{i,p_i}}(\phi)}}{U_{i,\varepsilon_i+1} - U_{i,\varepsilon_i+2}} V_{i,\varepsilon_i+3}^{n^{\boldsymbol{a}}_{w_{i,\varepsilon_i+3}}(\phi)} \cdots V_{i,p_i}^{n^{\boldsymbol{a}}_{w_{i,p_i}}(\phi)}$$

on the third line of (243) with 1.

We have now completed the description of the hyperbox $\mathfrak{C}(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$. This plays the role of the hyperbox $\mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$ from the link-minimal case presented in Section 9.1. Just as there, we let

$$\hat{D}^{\vec{M}}_{\boldsymbol{s}} : \mathfrak{A}^-(\mathscr{H}^L, \vec{M}, \boldsymbol{s}) \to \mathfrak{A}^-(\mathscr{H}^{\vec{L},\vec{M}}(M), \psi^{\vec{M}}(\boldsymbol{s}))$$

be the longest diagonal map in the hypercube obtained from $\mathfrak{C}(\mathscr{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$ by compression; compare (240) and (241). After identifying $\mathscr{H}^{\vec{L},\vec{M}}(M)$ with $\mathscr{H}^{L-M}$, we get the desired descent map $D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$ as in (235). By precomposing it with the corresponding projection-inclusion map, we get

$$\Phi^{\vec{M}}_{\boldsymbol{s}} = D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})} \circ \mathscr{I}^{\vec{M}}_{\boldsymbol{s}}.$$

These are the maps used to define the surgery complex $\mathscr{C}^-(\mathscr{H}, \Lambda)$. The fact that the differential on this complex squares to zero is an extension of Proposition 9.4, with a similar proof.

13.7 Proof of the theorem

In this section we prove Theorem 1.1 for arbitrary complete systems $\mathscr{H}$. In Section 12, we have already given the proof in the case where $\mathscr{H}$ is link-minimal. Recall that any complete system $\mathscr{H}$ can be obtained from a basic one by a sequence of the system moves listed in Section 8.9; see Proposition 8.35. Thus, to prove Theorem 1.1 it suffices to check that the quasi-isomorphism type of the surgery complex $\mathscr{C}^-(\mathscr{H}, \Lambda)$ is unchanged by these moves. Invariance under most of these moves (three-manifold isotopies, index one/two stabilizations and destabilizations, free index zero/three stabilizations, global shifts, and

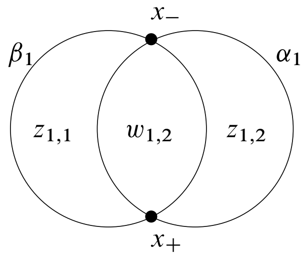

Figure 43: An index zero/three link stabilization.

elementary enlargements and contractions) follows by the same arguments as in the link-minimal case; see Section 12.6.

It remains to establish the following:

Proposition 13.12 *Let $\overline{\mathscr{H}}$ be a complete system of hyperboxes for a link $\vec{L} \subset Y$, and let $\mathscr{H}$ be the system obtained from $\overline{\mathscr{H}}$ by an index zero/three link stabilization. Then, for any framing Λ, the surgery complexes $\mathscr{C}^-(\overline{\mathscr{H}}, \Lambda)$ and $\mathscr{C}^-(\mathscr{H}, \Lambda)$ are quasi-isomorphic over $\mathbb{F}[[U]]$.*

Proof We keep the notation from Section 13.6 for everything in the construction of $\mathscr{C}^-(\mathscr{H}, \Lambda)$. Furthermore, we assume that the stabilization from $\overline{\mathscr{H}}$ to $\mathscr{H}$ is done on the link component L_1, and introduces the points labeled $w_{1,2}$ and $z_{1,2}$ in the neighborhood of an existing $z_{1,1}$; see Figure 43.[5] Thus, the basepoints on L_1 in $\overline{\mathscr{H}}$ are $w_{1,1}, z_{1,1}, w_{1,3}, z_{1,3}, \dots, w_{1,p_1}, z_{1,p_1}$.

Note that $\mathscr{C}^-(\mathscr{H}, \Lambda)$ is a complex over the ring

$$\mathscr{R} = \mathbb{F}[[(U_{i,j})_{1\le i\le \ell,\, 1\le j\le p_i}, (U_{\circ,j})_{1\le j\le p}]]$$

and $\mathscr{C}^-(\overline{\mathscr{H}}, \Lambda)$ is a complex over the smaller ring $\overline{\mathscr{R}}$, in which we do not have the variable $U_{1,2}$. We will in fact prove that the complexes $\mathscr{C}^-(\overline{\mathscr{H}}, \Lambda)$ and $\mathscr{C}^-(\mathscr{H}, \Lambda)$ are quasi-isomorphic over $\overline{\mathscr{R}}$. This will imply that they are quasi-isomorphic over $\mathbb{F}[[U]]$, where U is any of the variables in $\overline{\mathscr{R}}$.

The proof will be very similar to the one in Section 13.2.

Let us start by comparing the complexes $\mathfrak{A}^-(\mathscr{H}^L, s)$ and $\mathfrak{A}^-(\overline{\mathscr{H}}^L, s)$. These are subcomplexes of $\mathfrak{C}_{\text{int}}(\mathscr{H}^L)$ and $\mathfrak{C}_{\text{int}}(\overline{\mathscr{H}}^L)$, respectively. Recall that $\mathfrak{C}_{\text{int}}(\mathscr{H}^L)$ was described in Section 13.6 as a dg module over the dga

$$\mathfrak{R}^Y = \mathbb{F}[[(U_{i,j})_{1\le i\le \ell,1\le j\le p_i} (U_{\circ,j})_{1\le j\le p}]][(V_{i,j}, Y_{i,j})_{1\le i\le \ell,2\le j\le p_i}]'/(Y_{i,j}^2 = 0, \partial Y_{i,j} = U_{i,j} + U_{i,1}V_{i,j}).$$

[5] Strictly speaking, we lose some generality by choosing the stabilization on L_1, because the compression procedure from Section 5 involves an ordering of the components. However, this issue will not be relevant here, and the arguments below apply equally well to all other components, with only cosmetic modifications.

Similarly, $\mathfrak{C}_{\mathrm{int}}(\overline{\mathcal{H}}^L)$ is a dg module over a dga $\overline{\mathfrak{R}}^Y$; compared to $\mathfrak{R}^Y$, in $\overline{\mathfrak{R}}^Y$ we do not have the variables $U_{1,2}$, $V_{1,2}$ and $Y_{1,2}$. Further, by looking at the formula (232) for the differential on $\mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$, and its analogue for $\mathfrak{C}_{\mathrm{int}}(\overline{\mathcal{H}}^L)$, we deduce that $\mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)$ is obtained from $\mathfrak{C}_{\mathrm{int}}(\overline{\mathcal{H}}^L)$ by first considering the cone

$$\mathfrak{C}_{\mathrm{int}}(\overline{\mathcal{H}}^L)_-^{U_{1,1}\to U_{1,2}}[\![U_{1,1}]\!][V_{1,2}]' \xrightarrow{V_{1,2}-1} \mathfrak{C}_{\mathrm{int}}(\overline{\mathcal{H}}^L)_+^{U_{1,1}\to U_{1,2}}[\![U_{1,11}]\!][V_{1,2}]' \tag{244}$$

and then taking the cone of $U_{1,2}+U_{1,1}V_{1,2}$, ie introducing the variable $Y_{1,2}$ with $\partial Y_{1,2}=U_{1,2}+U_{1,1}V_{1,2}$. We are using here the description of the holomorphic disks in the stabilized diagram; see Proposition 7.2, and the subscripts $-$ and $+$ refer to containing the points x_- and x_+ from Figure 43. The situation is entirely similar to the one in Section 13.2, where $\mathfrak{C}_{\mathrm{int}}$ was obtained from $\overline{C}\{z;n\geq 0\}$ in this way. As explained there, this implies that the complexes are homotopy equivalent over $\overline{\mathcal{R}}$. In our case, the equivalence is given by the projection

$$\rho\colon \mathfrak{C}_{\mathrm{int}}(\mathcal{H}^L)\to \mathfrak{C}_{\mathrm{int}}(\overline{\mathcal{H}}^L), \qquad \rho(U_{1,2}^m V_{1,2}^n Y_{1,2}^a(\boldsymbol{x}\times r)) = \begin{cases} U_{1,1}^m \boldsymbol{x} & \text{if } a=0,\ r=x_+, \\ 0 & \text{otherwise.}\end{cases} \tag{245}$$

Compare (203). The map ρ is filtered with respect to the filtrations $\mathfrak{F}_i$, and as such it induces an equivalence between the subcomplexes in filtration levels $\mathfrak{F}_i\leq s_i$:

$$\rho\colon \mathfrak{A}^-(\mathcal{H}^L,\boldsymbol{s})\to \mathfrak{A}^-(\overline{\mathcal{H}}^L,\boldsymbol{s}).$$

Similar remarks apply when we compare the complexes $\mathfrak{A}^-(\mathcal{H}^{L-M},\boldsymbol{s})$ and $\mathfrak{A}^-(\overline{\mathcal{H}}^{L-M},\boldsymbol{s})$ for sublinks M such that $L_1\not\subseteq M$. We get that these complexes are equivalent via a projection ρ, defined by the same formula (245).

When $L_1\subseteq M$, the situation is simpler, as the basepoints $z_{1,1}$ and $z_{1,2}$ disappear. The diagram $\mathcal{H}^{L-M}$ is obtained from $\overline{\mathcal{H}}^{L-M}$ by a free index zero/three stabilization. Hence, we can identify $\mathfrak{A}^-(\mathcal{H}^{L-M},\boldsymbol{s})$ with the cone

$$\mathfrak{A}^-(\overline{\mathcal{H}}^{L-M},\boldsymbol{s})_+[\![U_{1,2}]\!] \xrightarrow{U_{1,2}-U_{1,1}} \mathfrak{A}^-(\overline{\mathcal{H}}^{L-M},\boldsymbol{s})_-[\![U_{1,2}]\!].$$

Thus, we can still construct an equivalence

$$\rho\colon \mathfrak{A}^-(\mathcal{H}^{L-M},\boldsymbol{s})\to \mathfrak{A}^-(\overline{\mathcal{H}}^{L-M},\boldsymbol{s}),$$

this time given by

$$\rho(U_{1,2}^m(\boldsymbol{x}\times r)) = \begin{cases} U_{1,1}^m\boldsymbol{x} & \text{if } r=x_+, \\ 0 & \text{otherwise.}\end{cases} \tag{246}$$

We can view the surgery complexes $\mathcal{C}^-(\mathcal{H},\Lambda)$ and $\mathcal{C}^-(\overline{\mathcal{H}},\Lambda)$ as ℓ-dimensional hypercubes (as in Section 4.1). We aim to construct a chain map between these hypercubes, in the sense of Definition 5.15, by starting from the equivalences

$$\rho\colon \mathfrak{A}^-(\mathcal{H}^{L-M},\boldsymbol{s})\to \mathfrak{A}^-(\overline{\mathcal{H}}^{L-M},\boldsymbol{s}),$$

and then adding diagonal maps as in the diagram

$$
\begin{tikzcd}
\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \arrow[r, "\Phi_{\boldsymbol{s}}^{\vec{M}}"] \arrow[d, "\rho"'] \arrow[dr] & \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^{\vec{M}}(\boldsymbol{s})) \arrow[d, "\rho"] \\
\mathfrak{A}^-(\overline{\mathcal{H}}^L, \boldsymbol{s}) \arrow[r, "\Phi_{\boldsymbol{s}}^{\vec{M}}"] & \mathfrak{A}^-(\overline{\mathcal{H}}^{L-M}, \psi^{\vec{M}}(\boldsymbol{s}))
\end{tikzcd}
\tag{247}
$$

For these to form a chain map between hypercubes, we need for example that when M is a knot (so that $\Phi_{\boldsymbol{s}}^{\vec{M}}$ are chain maps), the projections ρ commute with $\Phi_{\boldsymbol{s}}^{\vec{M}}$ up to the chain homotopy given by the diagonal. In general, we need the diagonal maps to be higher chain homotopies, ie to satisfy the relations (60).

Observe that, if we can construct such a chain map, then it would be filtered with respect to the filtration given by the value of $-\|\varepsilon\|$ in the hypercube. The maps on the associated graded would be just the projections ρ, which are quasi-isomorphisms over $\overline{\mathcal{R}}$. This would imply that the whole map from $\mathcal{C}^-(\mathcal{H}, \Lambda)$ and $\mathcal{C}^-(\overline{\mathcal{H}}, \Lambda)$ is a quasi-isomorphism, and complete the proof of the proposition.

We are left to find diagonal maps as in (247). Recall that each $\Phi_{\boldsymbol{s}}^{\vec{M}}$ is the composition of a projection-inclusion $\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}$ and a descent map $D_{p^{\vec{M}}(\boldsymbol{s})}^{\vec{M}}$. We will first show that the projections ρ commute with the projection-inclusions (on the nose), and then that they commute with the descent maps up to chain homotopies that fit into suitable hypercubes.

Thus, with regard to the projection-inclusions, we seek commutative diagrams of the form

$$
\begin{tikzcd}
\mathfrak{A}^-(\mathcal{H}^L, \boldsymbol{s}) \arrow[r, "\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}"] \arrow[d, "\rho"'] & \mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(\boldsymbol{s})) \arrow[d, "\rho"] \\
\mathfrak{A}^-(\overline{\mathcal{H}}^L, \boldsymbol{s}) \arrow[r, "\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}"] & \mathfrak{A}^-(\overline{\mathcal{H}}^L, \vec{M}, p^{\vec{M}}(\boldsymbol{s}))
\end{tikzcd}
\tag{248}
$$

Here, the map ρ on the right-hand side is defined just as before, through formulas of the type (246) if $L_1 \subseteq M_+$, or (245) otherwise.

To check that (248) commutes, we use the formula (234) for the projection-inclusion maps. If $L_1 \not\subseteq M_+$, the verification is straightforward, since both vertical maps are given by the same formula (245). If $L_1 \subseteq M_+$, ie $1 \in I_+(\vec{L}, \vec{M})$, then we use (245) for the left vertical map and (246) for the right vertical map. We calculate

$$
\begin{aligned}
&\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}\biggl(\rho\biggl(U_{1,2}^m \cdot \prod_{i \in I_+(\vec{L},\vec{M})} \prod_j V_{i,j}^{n_{i,j}} Y_{i,j}^{a_{i,j}}(\boldsymbol{x} \times r)\biggr)\biggr) \\
&\qquad = \begin{cases} \mathcal{I}_{\boldsymbol{s}}^{\vec{M}}\bigl(U_{1,1}^m \cdot \prod_{j=3}^{p_1} V_{i,j}^{n_{i,j}} Y_{i,j}^{a_{i,j}} \cdot \prod_{i \in I_+(\vec{L},\vec{M}), i \neq 1} \prod_j V_{i,j}^{n_{i,j}} Y_{i,j}^{a_{i,j}} \boldsymbol{x}\bigr) & \text{if } a_{1,2} = 0,\ r = x_+, \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}
$$

which gives

$$
U_{1,1}^{m+A_1(\boldsymbol{x})-s_1-\sum_{j=3}^{p_1} n_{1,j}} U_{1,3}^{n_{1,3}} \cdots U_{1,p_1}^{n_{1,p_1}} \cdot \prod_{\substack{i \in I_+(\vec{L},\vec{M}) \\ i \neq 1}} U_{i,1}^{A_i(\boldsymbol{x})-s_i-\sum_{j=2}^{p_i} n_{i,j}} U_{i,2}^{n_{i,2}} \cdots U_{i,p_i}^{n_{i,p_i}} \boldsymbol{x}
$$

if $r = x_+$ and all $a_{i,j} = 0$ for $i \in I_+(\vec{L}, \vec{M})$, and gives zero otherwise. This is the same answer as

$$\rho\biggl(\mathcal{I}_{\boldsymbol{s}}^{\vec{M}}\biggl(U_{1,2}^{m}\cdot \prod_{i\in I_+(\vec{L},\vec{M})}\prod_{j} V_{i,j}^{n_{i,j}} Y_{i,j}^{a_{i,j}}(\boldsymbol{x}\times r)\biggr)\biggr)$$

$$= \begin{cases} \rho\bigl(U_{1,2}^{m}\cdot\prod_{i\in I_+(\vec{L},\vec{M})} U_{i,1}^{A_i(\boldsymbol{x})-s_i-\sum_{j=2}^{p_i} n_{i,j}} U_{i,2}^{n_{i,2}}\cdots U_{i,p_i}^{n_{i,p_i}}(\boldsymbol{x}\times r)\bigr) & \text{if all } a_{i,j}=0 \text{ for } i\in I_+(\vec{L},\vec{M}),\\ 0 & \text{otherwise.}\end{cases}$$

Hence, the diagram (248) commutes.

Next, we look at descent maps of the form (235). When M is a knot, we seek a diagram

$$\begin{tikzcd} \mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(\boldsymbol{s})) \arrow[r, "D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}"] \arrow[d, "\rho"'] \arrow[dr] & \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^{\vec{M}}(\boldsymbol{s})) \arrow[d, "\rho"] \\ \mathfrak{A}^-(\overline{\mathcal{H}}^L, \vec{M}, p^{\vec{M}}(\boldsymbol{s})) \arrow[r, "D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}"] & \mathfrak{A}^-(\overline{\mathcal{H}}^{L-M}, \vec{M}, \psi^{\vec{M}}(\boldsymbol{s})) \end{tikzcd} \tag{249}$$

that commutes up to the homotopy given by the diagonal map. For general M, we need higher chain homotopies that fit into a chain map between hypercubes.

Recall from Section 13.6 that the map

$$D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}\colon \mathfrak{A}^-(\mathcal{H}^L, \vec{M}, p^{\vec{M}}(\boldsymbol{s})) \to \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^{\vec{M}}(\boldsymbol{s}))$$

is the largest diagonal in the compression of a hyperbox

$$C = \mathfrak{C}(\mathcal{H}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s})).$$

Similarly, the corresponding map $D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$ coming from $\overline{\mathcal{H}}$ is obtained by compressing a hyperbox

$$\overline{C} = \mathfrak{C}(\overline{\mathcal{H}}^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s})).$$

The relation between the hyperboxes of chain complexes $\mathcal{H}^{\vec{L},\vec{M}}$ and $\overline{\mathcal{H}}^{\vec{L},\vec{M}}$ was described in Section 8.9. When $L_1 \not\subseteq M_-$, these two hyperboxes have the same size, and the diagrams in $\mathcal{H}^{\vec{L},\vec{M}}$ are obtained from the corresponding ones in $\overline{\mathcal{H}}^{\vec{L},\vec{M}}$ by an index zero/three link stabilization. As such, there are equivalences ρ that relate the corresponding complexes in the hyperboxes $\mathcal{H}^{\vec{L},\vec{M}}$ and $\overline{\mathcal{H}}^{\vec{L},\vec{M}}$. These equivalences commute (up to chain homotopy) with the polygon maps in the hyperboxes, by Propositions 6.20 and 7.16. Thus, the maps ρ commute with $D^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$ up to chain homotopy, and we obtain commuting diagrams of the form (249).

The more interesting case is when $L_1 \subseteq M_-$. Then, as explained in Section 8.9, the hyperbox $\mathcal{H}^{\vec{L},\vec{M}}$ is obtained from $\overline{\mathcal{H}}^{\vec{L},\vec{M}}$ by stabilizing all diagrams as in Figure 14, and then also increasing the length of the hyperbox by one in the direction $i = 1$, by adding the move shown in Figure 18. Furthermore, recall from (236) that the length of the hyperbox C in the direction $i = 1$ is $d_1 + p_1 - 1$, where d_1 is the length of $\mathcal{H}^{\vec{L},\vec{M}}$ in that direction. A similar formula holds for $\overline{C}$ and $\overline{\mathcal{H}}^{\vec{L},\vec{M}}$. Since there is one extra pair of basepoints in $\mathcal{H}^{\vec{L},\vec{M}}$ than in $\overline{\mathcal{H}}^{\vec{L},\vec{M}}$, we deduce that the first side length of C is two more than that of $\overline{C}$.

To be able to compare the hyperboxes, we replace $\overline{C}$ by a new hyperbox $\widetilde{C}$, which is obtained from $\overline{C}$ by two elementary enlargements, as in Section 5.8 both in the direction $i = 1$, with one done at the beginning and one at the end. The beginning one has to do with the additional basepoint (and thus additional transition maps) present in C, and the ending one with the move from Figure 18.

Now, the hyperboxes C and $\widetilde{C}$ have the same size. We claim that we can construct a chain map

$$Z\colon C \to \widetilde{C}$$

between these hyperboxes, in the sense of Definition 5.15, such that the maps $Z^{\mathbf{0}}_{\varepsilon} = \rho\colon C^{\varepsilon} \to \widetilde{C}^{\varepsilon}$ are chain homotopy equivalences over $\overline{\mathscr{R}}$. (Recall that a chain map between hyperboxes consists of maps $Z^{\mathbf{0}}_{\varepsilon}$ that preserve the index ε, as well as higher diagonal maps $Z^{\varepsilon'-\varepsilon}_{\varepsilon}\colon C^{\varepsilon} \to \widetilde{C}^{\varepsilon'}$ when $\varepsilon' > \varepsilon$ are neighbors.) After compression, our map Z would produce a chain map between the resulting hypercubes. Furthermore, by Lemma 5.21, elementary enlargements leave the compressed hypercubes unchanged. We would thus get the desired chain map between the compressions of C and $\overline{C}$, providing diagonal maps in (249).

We start by constructing the quasi-isomorphisms $\rho\colon C^{\varepsilon} \to \widetilde{C}^{\varepsilon}$. Recall from Section 13.6 that C^{ε} looks like

- the resolution $\mathfrak{A}^-(\,\cdot\,, \psi^{\vec{M}}(\mathbf{s}))$ in the directions $i \notin I(\vec{L}, \vec{M})$,
- the complexes $\mathfrak{C}_j$ from (217), with $j = \varepsilon_i$, in the directions $i \in I(\varepsilon)$,
- the complex $\boldsymbol{CF}^- = C\{z\}$ in the directions $i \in I_-(\vec{L}, \vec{M}) - I(\varepsilon)$,
- the complex $\boldsymbol{CF}^- = C\{w\}$ in the directions $i \in I_+(\vec{L}, \vec{M})$.

A similar description applies to $\widetilde{C}^{\varepsilon}$. In our case, since $L_1 \subseteq M_-$, either $1 \in I(\varepsilon)$ or $1 \in I_-(\vec{L}, \vec{M}) - I(\varepsilon)$, according to whether $\varepsilon_1 < p_1 - 1$ or not.

Observe that ε_1 can take values from 0 to $d_1' = d_1 + p_1 - 1$. If $0 < \varepsilon_1 < d_1'$, then (regardless of whether ε_1 is less or greater than $p_1 - 1$), the basepoint $w_{2,1}$ does not play a role in the complex C^{ε}, and we are in the situation of Figure 14. Thus, as in (76), we can identify C^{ε} with the cone

$$(\widetilde{C}^{\varepsilon})_+^{U_{1,1}\to U_{1,2}}[\![U_{1,1}]\!] \xrightarrow{U_{1,1}-U_{1,2}} (\widetilde{C}^{\varepsilon})_-^{U_{1,1}\to U_{1,2}}[\![U_{1,1}]\!].$$

We get an equivalence $\rho\colon C^{\varepsilon} \to \widetilde{C}^{\varepsilon}$ given by

$$\rho(U_{1,2}^m(\boldsymbol{x} \times r)) = \begin{cases} U_{1,1}^m \boldsymbol{x} & \text{if } r = x_-, \\ 0 & \text{if } r = x_+. \end{cases}$$

On the other hand, when $\varepsilon_1 = 0$, then the complex C^{ε} is obtained from $\widetilde{C}^{\varepsilon}$ by taking cones on $V_{1,2} - 1$ and $U_{1,2} + U_{1,1}V_{1,2}$, as in (244), and therefore we have an equivalence ρ given by (245):

$$\rho(U_{1,2}^m V_{1,2}^n Y_{1,2}^a(\boldsymbol{x} \times r)) = \begin{cases} U_{1,1}^m \boldsymbol{x} & \text{if } a = 0,\ r = x_+, \\ 0 & \text{otherwise.} \end{cases}$$

Lastly, when $\varepsilon_1 = d_1'$, then C^{ε} is constructed from the Heegaard diagram at the end of the move in Figure 18, ie using the α_1 and γ_1 curves. This is just a free index zero/three stabilization of the diagram that gives $\widetilde{C}^{\varepsilon}$, and therefore we can identify C^{ε} with the cone

$$\widetilde{C}_+^{\varepsilon}[\![U_{1,2}]\!] \xrightarrow{U_{1,1}-U_{1,2}} \widetilde{C}_-^{\varepsilon}[\![U_{1,2}]\!].$$

We obtain an equivalence $\rho\colon C^{\varepsilon} \to \widetilde{C}^{\varepsilon}$ given by

$$\rho(U_{1,2}^m(\boldsymbol{x} \times r)) = \begin{cases} U_{1,1}^m \boldsymbol{x} & \text{if } r = y_-, \\ 0 & \text{if } r = y_+. \end{cases}$$

We now construct the other components of the chain map Z, ie $Z_{\varepsilon}^{\varepsilon'-\varepsilon}$ for $\varepsilon < \varepsilon'$.

When $\varepsilon_1 > 0$, we simply set $Z_{\varepsilon}^{\varepsilon'-\varepsilon} = 0$. If we denote by $D_{\varepsilon}^{\varepsilon'-\varepsilon}$ and $\widetilde{D}_{\varepsilon}^{\varepsilon'-\varepsilon}$ the maps in the hyperboxes C and $\widetilde{C}$, we claim that the diagrams

$$\begin{tikzcd} C^{\varepsilon} \arrow[r, "D_{\varepsilon}^{\varepsilon'-\varepsilon}"] \arrow[d, "\rho"'] & C^{\varepsilon'} \arrow[d, "\rho"] \\ \widetilde{C}^{\varepsilon} \arrow[r, "\widetilde{D}_{\varepsilon}^{\varepsilon'-\varepsilon}"] & \widetilde{C}^{\varepsilon'} \end{tikzcd}$$

commute on the nose. Indeed, when $1 < \varepsilon_1 < d_1'-1$, this follows from the study of holomorphic polygons in the stabilized diagram; see Proposition 7.14. For $\varepsilon_1 = d_1'-2$ and $\varepsilon_1' = d_1'-1$, we can use the description of holomorphic polygons under the move in Figure 18, that is, Proposition 7.21. When $\varepsilon_1 = \varepsilon_1' = d_1'-1$, we use the similar description of a free index zero/three stabilization; see Proposition 6.21.

When $\varepsilon_1 = 0$, we again set $Z_{\varepsilon}^{\varepsilon'-\varepsilon} = 0$ (for $\varepsilon < \varepsilon'$), with one exception: if ε' differs from ε only in position $i = 1$, where $\varepsilon_1' = 1$ (in other words if $\varepsilon' = \varepsilon + \tau_1$ in the notation of Section 5.1). When $\varepsilon_1 = 0$ and $\varepsilon' = \varepsilon + \tau_1$, we set

$$Z_{\varepsilon}^{\varepsilon'-\varepsilon}(U_{1,2}^m V_{1,2}^n Y_{1,2}^a(\boldsymbol{x} \times r)) = \begin{cases} mU_{1,1}^{m-1}\boldsymbol{x} & \text{if } a = 0,\ r = x_-, \\ 0 & \text{otherwise.} \end{cases}$$

This is exactly analogous to the situation described in Section 13.2, where we had the formula (204) for the diagonal map Z. Checking that $Z_{\varepsilon}^{\varepsilon'-\varepsilon}$ fits into a chain map between the hypercubes in C and $\widetilde{C}$ is entirely similar to the proof of (199) in Section 13.2.

This completes the construction of the chain map between C and $\widetilde{C}$. From here we obtain the desired quasi-isomorphism between the surgery complexes $\mathcal{C}^-(\mathcal{H}, \Lambda)$ and $\mathcal{C}^-(\overline{\mathcal{H}}, \Lambda)$. □

14 Beyond the surgery theorem

We discuss here several extensions of Theorem 1.1.

14.1 Maps induced by surgery

We work in the setting of Section 13.7, with $\vec{L} \subset Y$ being a link in an integral homology three-sphere, and $\mathcal{H}$ a complete system of hyperboxes for $\vec{L}$. Let $L' \subseteq L$ be a sublink, with the orientation induced from $\vec{L}$. The hyperboxes $\mathcal{H}^{\vec{L}'', \vec{M}}$ with $M \subseteq L'' \subseteq L'$ form a complete system of hyperboxes for $\vec{L}'$, which we denote by $\mathcal{H}|_{L'}$.

Following the notation from Section 11.1, we let $W_\Lambda(L', L)$ be the cobordism from $Y_{\Lambda|_{L'}}(L')$ to $Y_\Lambda(L)$ given by surgery on $L - L'$ (framed with the restriction of Λ). Recall that in Lemma 11.2 we established an identification

$$\mathrm{Spin}^c(W_\Lambda(L', L)) \cong \mathbb{H}(L)/H(L, \Lambda|_{L'})$$

such that the natural projection

$$\pi^{L,L'}\colon \bigl(\mathbb{H}(L)/H(L, \Lambda|_{L'})\bigr) \to (\mathbb{H}(L)/H(L, \Lambda))$$

corresponds to restricting the Spin^c structures to $Y_\Lambda(L)$, and the map

$$\psi^{L-L'}\colon \bigl(\mathbb{H}(L)/H(L, \Lambda|_{L'})\bigr) \to \bigl(\mathbb{H}(L')/H(L', \Lambda|_{L'})\bigr)$$

corresponds to restricting them to $Y_{\Lambda|_{L'}}(L')$.

Lemma 14.1 *Fix an equivalence class $\boldsymbol{t} \in \mathbb{H}(L)/H(L, \Lambda|_{L'})$. Then the map*

(250) $$\psi^{L-L'}\colon \{\boldsymbol{s} \in \mathbb{H}(L) \mid [\boldsymbol{s}] = \boldsymbol{t}\} \to \{\boldsymbol{s}' \in \mathbb{H}(L') \mid [\boldsymbol{s}'] = \psi^{L-L'}(\boldsymbol{t})\}$$

is always surjective, and it is injective if and only if every component $L_j \subseteq L - L'$ is rationally null-homologous inside the surgered manifold $Y_{\Lambda|_{L'}}(L')$.

Proof Using affine identifications between spaces of Spin^c structures and second cohomology, surjectivity and injectivity of (250) are equivalent to the same conditions for the restriction map:

(251) $$\mathrm{Ker}\bigl(H^2(W_\Lambda(L)) \to H^2(W_\Lambda(L', L))\bigr) \to \mathrm{Ker}\bigl(H^2(W_{\Lambda|_{L'}}(L')) \to H^2(Y_{\Lambda|_{L'}}(L'))\bigr).$$

Using long exact sequences for pairs, we can rewrite (251) as

$$\mathrm{Im}\bigl(H^2(W_\Lambda(L), W_\Lambda(L', L)) \to H^2(W_\Lambda(L))\bigr) \to \mathrm{Im}\bigl(H^2(W_{\Lambda|_{L'}}(L'), Y_{\Lambda|_{L'}}(L')) \to H^2(W_{\Lambda|_{L'}}(L'))\bigr).$$

Surjectivity of the last map follows directly from the existence of a commutative diagram

$$\begin{tikzcd}
H^2(W_\Lambda(L), W_\Lambda(L', L)) \arrow[r] \arrow[d, "\cong"'] & H^2(W_\Lambda(L)) \arrow[d] \\
H^2(W_{\Lambda|_{L'}}(L'), Y_{\Lambda|_{L'}}(L')) \arrow[r] & H^2(W_{\Lambda|_{L'}}(L'))
\end{tikzcd}$$

Injectivity of (251) is equivalent to injectivity of the combined restriction map

(252) $$H^2(W_\Lambda(L)) \to H^2(W_\Lambda(L', L)) \oplus H^2(W_{\Lambda|_{L'}}(L')).$$

In turn, using the Mayer–Vietoris sequence in cohomology, injectivity of (252) is equivalent to surjectivity of the map

(253) $$H^1(W_\Lambda(L', L)) \oplus H^1(W_{\Lambda|_{L'}}(L')) \to H^1(Y_{\Lambda|_{L'}}(L')).$$

Since $H^1(W_{\Lambda|_{L'}}(L')) = 0$ and H^1 groups have no torsion, we can rephrase surjectivity of (253) as injectivity of the dual map with rational coefficients:

$$H_1(Y_{\Lambda|_{L'}}(L'); \mathbb{Q}) \to H_1(W_\Lambda(L', L); \mathbb{Q}).$$

In turn, this is equivalent to the vanishing of the boundary map

$$H_2(W_\Lambda(L', L), Y_{\Lambda|_{L'}}(L'); \mathbb{Q}) \to H_1(Y_{\Lambda|_{L'}}(L'); \mathbb{Q}). \tag{254}$$

Since the domain of (254) is generated by the cores of the 2-handles attached along link components $L_j \subseteq L - L'$, the conclusion follows. □

Observe that, for every equivalence class $\boldsymbol{t} \in \mathbb{H}(L)/H(L, \Lambda|_{L'})$,

$$\widetilde{\mathcal{C}}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}} := \bigoplus_{L-L'\subseteq M\subseteq L} \prod_{\substack{\boldsymbol{s}\in\mathbb{H}(L)\\ [\boldsymbol{s}]=\boldsymbol{t}}} \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$$

is a subcomplex of $\mathcal{C}^-(\mathcal{H}, \Lambda, \pi^{L,L'}(\boldsymbol{t})) \subseteq \mathcal{C}^-(\mathcal{H}, \Lambda)$.

If the injectivity condition in Lemma 14.1 holds, so that the map (250) is a bijection, then the complex $\widetilde{\mathcal{C}}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$ is isomorphic to

$$\mathcal{C}^-(\mathcal{H}|_{L'}, \Lambda|_{L'}, \psi^{L-L'}(\boldsymbol{t})) = \bigoplus_{M'\subseteq L'} \prod_{\substack{\boldsymbol{s}'\in\mathbb{H}(L')\\ [\boldsymbol{s}']=\psi^{L-L'}(\boldsymbol{t})}} \mathfrak{A}^-(\mathcal{H}^{L'-M'}, \psi^{M'}(\boldsymbol{s}')).$$

Indeed, the isomorphism is induced by taking M to $M' = M - (L - L')$ and $\boldsymbol{s}$ to $\boldsymbol{s}' = \psi^{L-L'}(\boldsymbol{s})$.

In the general case, even if the map (250) is not bijective, we still have an injective chain map

$$\iota\colon \mathcal{C}^-(\mathcal{H}|_{L'}, \Lambda|_{L'}, \psi^{L-L'}(\boldsymbol{t})) \to \widetilde{\mathcal{C}}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$$

given by summing up all the natural identifications

$$\mathfrak{A}^-(\mathcal{H}^{L'-M'}, \psi^{M'}(\boldsymbol{s}')) \xrightarrow{\cong} \mathfrak{A}^-(\mathcal{H}^{L-M}, \psi^M(\boldsymbol{s}))$$

such that $M' = M - (L - L')$, $[\boldsymbol{s}] = \boldsymbol{t}$ and $\boldsymbol{s}' = \psi^{L-L'}(\boldsymbol{s})$.

We denote by $\mathcal{C}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$ the image of the inclusion ι. This is a subcomplex of $\mathcal{C}^-(\mathcal{H}, \Lambda, \pi^{L,L'}(\boldsymbol{t}))$ that is isomorphic to $\mathcal{C}^-(\mathcal{H}|_{L'}, \Lambda|_{L'}, \psi^{L-L'}(\boldsymbol{t}))$. When the injectivity condition in Lemma 14.1 holds, $\mathcal{C}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$ coincides with the previously defined $\widetilde{\mathcal{C}}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$.

Example 14.2 Suppose $L = L_1 \cup L_2$ is a two-component link, with $\mathrm{lk}(L_1, L_2) = 1$. Consider $(0, 0)$ surgery on L, so that $\Lambda = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$. Let $L' = L_1$. Pick $s' \in \mathbb{Z}$ and set

$$\boldsymbol{t} = \left[\left(s' + \tfrac{1}{2}, \tfrac{1}{2}\right)\right] \in \mathbb{H}(L)/H(L, \Lambda|_{L'}) \cong \left(\mathbb{Z} + \tfrac{1}{2}\right)^2/(0, 1).$$

Then $\mathcal{C}^-(\mathcal{H}|_{L'}, \Lambda|_{L'}, \psi^{L-L'}(\boldsymbol{t}))$ is the complex associated to 0 surgery on L_1, and consists of two generalized Floer complexes

$$\mathfrak{A} := \mathfrak{A}^-(\mathcal{H}^{L_1}, s'), \quad \mathfrak{B} := \mathfrak{A}^-(\mathcal{H}^{\varnothing}),$$

Figure 44: The inclusion of $\mathcal{C}^-(\mathcal{H}|_{L'}, \Lambda|_{L'}, \psi^{L-L'}(\boldsymbol{t}))$ into $\widetilde{\mathcal{C}}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$. The blue and red arrows are copies of the maps $\Phi^{+L_1}_{s'}$ and $\Phi^{-L_1}_{s'}$, respectively. The inclusion ι is given by sending each generator $\boldsymbol{x} \in \mathfrak{A}$ to $(\dots, \boldsymbol{x}, \boldsymbol{x}, \boldsymbol{x}, \dots)$ in the infinite direct product on the right-hand side; and similarly for generators $\boldsymbol{y} \in \mathfrak{B}$. The image of ι is denoted by $\mathcal{C}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$.

related by the sum of the maps $\Phi^{+L_1}_{s'}$ and $\Phi^{-L_1}_{s'}$. On the other hand, $\widetilde{\mathcal{C}}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$ is the direct product of infinitely many copies of $\mathfrak{A}$ and $\mathfrak{B}$, related by maps $\Phi^{+L_1}_{s'}$ and $\Phi^{-L_1}_{s'}$ arranged in a zigzag, as in Figure 44. The inclusion ι is similar to the one defined in (158) from Section 11.7, except now the target has infinitely many factors.

Theorem 1.1 implies that the homology of $\mathcal{C}^-(\mathcal{H}|_{L'}, \Lambda|_{L'}, \psi^{L-L'}(\boldsymbol{t}))$, and hence also the homology of $\mathcal{C}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}$, are isomorphic to

$$\boldsymbol{HF}^-_*(Y_{\Lambda|_{L'}}(L'), \boldsymbol{t}|_{Y_{\Lambda|_{L'}}(L')}).$$

Ozsváth and Szabó [2006] associated a map $F^-_{W,\boldsymbol{t}}$ to any cobordism W between connected three-manifolds, and Spin^c structure $\boldsymbol{t}$ on that cobordism. In the case when the cobordism W consists only of two-handles (ie is given by integral surgery on a link), the following theorem gives a way of looking at the map $F^-_{W,\boldsymbol{t}}$ in terms of complete systems of hyperboxes:

Theorem 14.3 *Let $\vec{L} \subset Y$ be a link in an integral homology three-sphere, $L' \subseteq L$ a sublink, $\mathcal{H}$ a complete system of hypercubes for $\vec{L}$, and Λ a framing of L. Then, for any $\boldsymbol{t} \in \mathrm{Spin}^c(W_\Lambda(L', L)) \cong \mathbb{H}(L)/H(L, \Lambda|_{L'})$, the following diagram commutes*:

$$\begin{tikzcd}
H_*(\mathcal{C}^-(\mathcal{H}, \Lambda)^{L',\boldsymbol{t}}) \arrow[r] \arrow[d, "\cong"'] & H_*\big(\mathcal{C}^-(\mathcal{H}, \Lambda, \pi^{L,L'}(\boldsymbol{t}))\big) \arrow[d, "\cong"] \\
\boldsymbol{HF}^-_*(Y_{\Lambda|_{L'}}(L'), \boldsymbol{t}|_{Y_{\Lambda|_{L'}}(L')}) \arrow[r, "F^-_{W_\Lambda(L',L),\boldsymbol{t}}"] & \boldsymbol{HF}^-_*(Y_\Lambda(L), \boldsymbol{t}|_{Y_{\Lambda|_L}(L)})
\end{tikzcd}$$

Here, the top horizontal map is induced from the inclusion of chain complexes, while the two vertical isomorphisms are the ones from Theorem 1.1.

Proof We first discuss the proof in the case when $\mathcal{H}$ is a basic system. The argument is similar to the one in [Ozsváth and Szabó 2008b, Theorem 4.2]; basically, one has to keep track of the surgery maps all throughout the arguments in Section 12. The key point is to find commutative diagrams that relate the inclusion maps of subhypercubes of $\mathscr{H}^\delta$ to cobordism maps between the respective Floer complexes. This is done by applying the results of Section 11.7 (precisely, Propositions 11.30, 11.31, 11.33 and 11.35) repeatedly, as we follow the iteration process in the proof of Proposition 12.11.

We emphasize that, even though in the proof of Proposition 12.11 we have chosen a particular ordering of the components of L (such that, if Λ is degenerate, Λ_1 is in the span of the other framing vectors, etc), regardless of the ordering we can apply the results of Section 11.7 to obtain an identification of the cobordism maps associated to $W_\Lambda(L', L)$ for *any* sublink $L' \subset L$. For the sake of concreteness, we explain how this works in the case of a link of two components $L = L_1 \cup L_2$.

Suppose that the framing matrix Λ is degenerate, so we had to choose the ordering of the components such that Λ_1 is a multiple of Λ_2. (This is a constraint when Λ_1 is the zero vector.) To pick a particular situation, let us assume that Λ_1 is indeed zero (so, in particular, L_1 and L_2 have zero linking number), whereas the surgery coefficient of L_2 is positive. In the iteration process that leads to Proposition 12.7, we have a diagram of maps

$$\begin{tikzcd}
\boldsymbol{CF}^{-,\delta}(Y_{\Lambda_1}(L_1); \mathscr{T}_2) \arrow[r, "g_1^\delta"] \arrow[d, "j_3^\delta"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda_1+m_1\tau_1}(L_1); \mathscr{T}_2) \arrow[r, "g_2^\delta"] \arrow[d, "k_3^\delta"] & \boldsymbol{CF}^{-,\delta}(Y; \mathscr{T}) \arrow[d, "l_3^\delta"] \\
\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L)) \arrow[r, "f_1^\delta"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda+m_1\tau_1}(L)) \arrow[r, "f_2^\delta"] \arrow[d, "k_1^\delta"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda_2}(L_2); \mathscr{T}_1) \arrow[d, "l_1^\delta"] \\
& \boldsymbol{CF}^{-,\delta}(Y_{\widetilde{\Lambda}}(L)) \arrow[r, "h_2^\delta"] \arrow[d, "k_2^\delta"] & \boldsymbol{CF}^{-,\delta}(Y_{\Lambda_2+m_2\tau_2}(L_2); \mathscr{T}_1) \arrow[d, "l_2^\delta"] \\
& \boldsymbol{CF}^{-,\delta}(Y_{\Lambda_1+m_1\tau_1}(L_1); \mathscr{T}_2) \arrow[r, "g_2^\delta"] & \boldsymbol{CF}^{-,\delta}(Y; \mathscr{T})
\end{tikzcd} \tag{255}$$

Here $\mathscr{T}_1 = \mathbb{F}[T_1]/(T_1^{m_1} - 1)$, $\mathscr{T}_2 = \mathbb{F}[T_2]/(T_2^{m_2} - 1)$, and we have dropped the chain homotopies (corresponding to diagonals in the squares above) from notation for simplicity. Note that the bottom two rows in (255) form the complex $\mathscr{C}^\delta$.

In the proof of Proposition 12.11, we first look at a quasi-isomorphism (f_1^δ, H_1^δ) from $\boldsymbol{CF}^{-,\delta}(Y_\Lambda(L))$ to the mapping cone $\operatorname{Cone}(f_2^\delta)$, and then at a quasi-isomorphism from the latter to the mapping square $\mathscr{C}^\delta$. Let $\mathfrak{u}$ be a Spin^c structure on $Y_\Lambda(L)$. An application of Proposition 11.24 shows that the first quasi-isomorphism above decomposes into a direct sum of several quasi-isomorphisms; one of them relates $\boldsymbol{CF}^{-,\delta}(Y_\Lambda, \mathfrak{u})$ to a mapping cone $\operatorname{Cone}(f_{2,\mathfrak{u}}^\delta)$, in a grading-preserving fashion. Then a double application of Proposition 11.11 gives a grading-preserving quasi-isomorphism from $\operatorname{Cone}(f_{2,\mathfrak{u}}^\delta)$ to a direct summand $\mathscr{C}_\mathfrak{u}^\delta \subset \mathscr{C}^\delta$.

Our new claim is that under the composition of these quasi-isomorphisms, the inclusion of the subcomplex $\operatorname{Cone}(g_{2,\mathfrak{u}}^\delta)$ (resp. $\operatorname{Cone}(l_{2,\mathfrak{u}}^\delta)$) into $\mathscr{C}_\mathfrak{u}^\delta$ corresponds to a cobordism map (coming from a unique Spin^c structure) from $Y_{\Lambda_1}(L_1)$ to $Y_\Lambda(L)$ (resp. from $Y_{\Lambda_2}(L_2)$ to $Y_\Lambda(L)$).

Indeed, in the case of $\operatorname{Cone}(g_{2,\mathfrak{u}}^\delta)$, a double application of Proposition 11.30 gives a commutative diagram between its inclusion into $\mathscr{C}_\mathfrak{u}^\delta$ and a map (consisting of $k_{3,\mathfrak{u}}^\delta, l_{3,\mathfrak{u}}^\delta$ and a diagonal chain homotopy) from

$\mathrm{Cone}(g^{\delta}_{2,\mathfrak{u}})$ to $\mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})$. Then another commutative diagram relates the latter map to a cobordism map from $Y_{\Lambda_1}(L_1)$ to $Y_{\Lambda}(L)$, using a double application of Proposition 11.29.

In the case of $\mathrm{Cone}(l^{\delta}_{2,\mathfrak{u}})$, a double application of Proposition 11.11 gives a commutative diagram between its inclusion into $\mathscr{C}^{\delta}_{\mathfrak{u}}$ and the inclusion of a summand of $\boldsymbol{CF}^{-,\delta}(Y_{\Lambda_2}(L_2);\mathscr{T}_1)$ into $\mathrm{Cone}(f^{\delta}_{2,\mathfrak{u}})$. Then another commutative diagram relates the latter inclusion to a cobordism map from $Y_{\Lambda_1}(L_2)$ to $Y_{\Lambda}(L)$, by applying Proposition 11.34.

It is straightforward to extend this argument to links of several components (and arbitrary sublinks). This leads to a proof of Theorem 14.3 in the case when the complete system $\mathscr{H}$ is basic.

For general complete systems, note that the quasi-isomorphisms used in the proof of Theorem 1.1 in Sections 12.6 and 13.7 respect the inclusion maps. We obtain the desired commutative diagram, except that a priori, the bottom row is a more general cobordism map than the one considered in [Ozsváth and Szabó 2006]. More precisely, it counts holomorphic triangles between multipointed Heegaard diagrams (for the respective three-manifolds) that may have more than one basepoint; the original cobordism maps $F^-_{W_{\Lambda}(L',L),\boldsymbol{t}}$ as defined in [loc. cit.], were going between singly pointed diagrams. Nevertheless, after some handleslides, isotopies and index one/two stabilizations and destabilizations, we can arrange so that the multipointed Heegaard triple diagrams involved are all obtained from basic ones by a sequence of index zero/three (free and link) stabilizations. The fact that the bottom row can be identified with $F^-_{W_{\Lambda}(L',L),\boldsymbol{t}}$ then follows from Propositions 6.21 and 7.16. □

14.2 Other versions

The chain complex $\mathscr{C}^-(\mathscr{H},\Lambda,\mathfrak{u})$ from Section 9.2 was constructed so that the version of Heegaard Floer homology appearing in Theorem 15.9 is $\boldsymbol{HF}^-$. We now explain how one can construct similar chain complexes $\widehat{\mathscr{C}}(\mathscr{H},\Lambda,\mathfrak{u})$, $\mathscr{C}^+(\mathscr{H},\Lambda,\mathfrak{u})$ and $\mathscr{C}^\infty(\mathscr{H},\Lambda,\mathfrak{u})$, corresponding to the theories $\widehat{HF}$, HF^+ and $\boldsymbol{HF}^\infty$.

The chain complex $\widehat{\mathscr{C}}(\mathscr{H},\Lambda,\mathfrak{u})$ is simply obtained from $\mathscr{C}^-(\mathscr{H},\Lambda,\mathfrak{u})$ by setting one of the variables U_i equal to zero. Its homology computes $\widehat{HF}(Y_\Lambda(L),\mathfrak{u})$.

The chain complex $\mathscr{C}^\infty(\mathscr{H},\Lambda,\mathfrak{u})$ is obtained from $\mathscr{C}^-(\mathscr{H},\Lambda,\mathfrak{u})$ by inverting all the U_i variables. It is a module over the ring of Laurent semi-infinite polynomials

$$\mathscr{R}^\infty=\mathbb{F}[\![U_1,\dots,U_k;U_1^{-1},\dots,U_k^{-1}]=(U_1,\dots,U_k)^{-1}\mathscr{R}.$$

In other words, $\mathscr{R}^\infty$ consists of those power series in U_i's that are sums of monomials with degrees bounded from below.

Note that $C^-(\mathscr{H},\Lambda,\mathfrak{u})$ is a subcomplex of $\mathscr{C}^\infty(\mathscr{H},\Lambda,\mathfrak{u})$. We denote the respective quotient complex by $\mathscr{C}^+(\mathscr{H},\Lambda,\mathfrak{u})$. Theorems 1.1 and 14.3 admit the following extension:

Theorem 14.4 *Fix a complete system of hyperboxes $\mathscr{H}$ for an oriented, ℓ-component link $\vec{L}$ in an integral homology three-sphere Y, and fix a framing Λ of L. Pick $\mathfrak{u} \in \mathrm{Spin}^c(Y_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$. Then there are vertical isomorphisms and horizontal long exact sequences making the following diagram commute*:

$$\begin{tikzcd}
\cdots \arrow[r] & H_*(\mathscr{C}^-(\mathscr{H},\Lambda,\mathfrak{u})) \arrow[r] \arrow[d, "\cong"] & H_*(\mathscr{C}^\infty(\mathscr{H},\Lambda,\mathfrak{u})) \arrow[r] \arrow[d, "\cong"] & H_*(\mathscr{C}^+(\mathscr{H},\Lambda,\mathfrak{u})) \arrow[r] \arrow[d, "\cong"] & \cdots \\
\cdots \arrow[r] & \boldsymbol{HF}^-_*(Y_\Lambda(L),\mathfrak{u}) \arrow[r] & \boldsymbol{HF}^\infty_*(Y_\Lambda(L),\mathfrak{u}) \arrow[r] & HF^+_*(Y_\Lambda(L),\mathfrak{u}) \arrow[r] & \cdots
\end{tikzcd}$$

Furthermore, the maps in these diagrams behave naturally with respect to cobordisms, in the sense that there are commutative diagrams analogous to those in Theorem 14.3, involving the cobordism maps $F^-_{W_\Lambda(L',L),\mathfrak{t}}$, $F^\infty_{W_\Lambda(L',L),\mathfrak{t}}$ and $F^+_{W_\Lambda(L',L),\mathfrak{t}}$.

Proof Inverting the U_i variables is an exact operation on modules; see for example [Atiyah and Macdonald 1969, Proposition 3.3]. Hence the quasi-isomorphisms relating $\mathscr{C}^-(\mathscr{H},\Lambda,\mathfrak{u})$ and $\boldsymbol{CF}^-(Y_\Lambda(L),\mathfrak{u})$ induce similar ones between the respective infinity versions. The five lemma then implies that the resulting maps between the plus versions are quasi-isomorphisms as well. Naturality with respect to the cobordism maps is clear from the construction. □

Remark 14.5 For the plus version, at least when $c_1(\mathfrak{u})$ is torsion, we can replace direct products with direct sums in the link surgery formula, and thus obtain a formula closer in spirit to [Ozsváth and Szabó 2008b]. Indeed, when $c_1(\mathfrak{u})$ is torsion, we have a relative $\mathbb{Z}$-grading on the surgery complex $\mathscr{C}^+(\mathscr{H},\Lambda,\mathfrak{u})$, and only finitely many terms in the direct product are nonzero in each grading level. Therefore, if we are interested in a fixed grading level, direct sums and direct products give the same answer. This implies that the plus surgery complexes constructed with direct sums are the same as those with direct products, overall. (The U actions also coincide, because they coincide in each grading level.)

By contrast, recall that, for $\boldsymbol{CF}^-$, we do need to use direct products, as explained in Section 4.3.

14.3 Mixed invariants of closed four-manifolds

Let us recall the definition of the closed four-manifold invariant from [Ozsváth and Szabó 2006]. Let X be a closed, oriented four-manifold with $b_2^+(X) \geq 2$. By deleting two four-balls from X we obtain a cobordism W from S^3 to S^3. We can cut W along a three-manifold N so as to obtain two cobordisms W_1, W_2 with $b_2^+(W_i) > 0$; further, the manifold N can be chosen such that $\delta H^1(N;\mathbb{Z}) \subset H^2(W;\mathbb{Z})$ is trivial. (If this is the case, N is called an *admissible cut*.) Let $\mathfrak{t}$ be a Spin^c structure on X and $\mathfrak{t}_1, \mathfrak{t}_2$ its restrictions to W_1, W_2. In this situation, the cobordism maps

$$F^-_{W_1,\mathfrak{t}_1} : \boldsymbol{HF}^-(S^3) \to \boldsymbol{HF}^-(N, \mathfrak{t}|_N) \quad \text{and} \quad F^+_{W_2,\mathfrak{t}_2} : HF^+(N,\mathfrak{t}|_N) \to HF^+(S^3)$$

factor through $HF_{\mathrm{red}}(N,\mathfrak{t}|_N)$, where

$$HF_{\mathrm{red}} = \mathrm{Coker}(\boldsymbol{HF}^\infty \to HF^+) \cong \mathrm{Ker}(\boldsymbol{HF}^- \to \boldsymbol{HF}^\infty).$$

By composing them we obtain the mixed map

$$F^{\mathrm{mix}}_{W,\mathfrak{t}} \colon \boldsymbol{HF}^-(S^3) \to HF^+(S^3),$$

which changes degree by the quantity $d(\mathfrak{t})+2$, where

$$d(\mathfrak{t}) = \tfrac{1}{4}(c_1(\mathfrak{t})^2 - 2\chi(X) - 3\sigma(X)).$$

Let Θ_- be the maximal-degree generator in $\boldsymbol{HF}^-(S^3)$, lying in degree -2. Clearly the map $F^{\mathrm{mix}}_{W,\mathfrak{t}}$ can be nonzero only when $d(\mathfrak{t})$ is even and nonnegative. If this is the case, the value

$$\Phi_{X,\mathfrak{t}} = U^{d(\mathfrak{t})/2} \cdot F^{\mathrm{mix}}_{W,\mathfrak{t}}(\Theta_-) \in HF^+_0(S^3) \cong \mathbb{F} \tag{256}$$

is an invariant of the four-manifold X and the Spin^c structure $\mathfrak{t}$. It is conjecturally the same as the Seiberg–Witten invariant.

Remark 14.6 In [Ozsváth and Szabó 2006, Section 9], the mixed invariant was defined as a map

$$\mathbb{F}[U] \otimes \Lambda^*(H_1(X)/\mathrm{Tors}) \to \mathbb{F}.$$

We only discuss here the value of this map at 1, which is exactly $\Phi_{X,\mathfrak{t}}$ as defined in (256).

The following definition was sketched in the introduction:

Definition 14.7 Let X be a closed, oriented four-manifold with $b_2^+(X) \geq 2$. A *cut link presentation for* X consists of a link $L \subset S^3$, a decomposition of L as a disjoint union

$$L = L_1 \amalg L_2 \amalg L_3,$$

and a framing Λ for L (with restrictions Λ_i to L_i for $i = 1, \dots, 3$) with the following properties:

- $S^3_{\Lambda_1}(L_1)$ is a connected sum of m copies of $S^1 \times S^2$ for some $m \geq 0$. We denote by W_1 the cobordism from S^3 to $\#^m(S^1 \times S^2)$ given by m one-handle attachments.
- $S^3_{\Lambda_1 \cup \Lambda_2 \cup \Lambda_3}(L_1 \cup L_2 \cup L_3)$ is a connected sum of m' copies of $S^1 \times S^2$ for some $m' \geq 0$. We denote by W_4 the cobordism from $\#^{m'}(S^1 \times S^2)$ to S^3 given by m' three-handle attachments.
- If we denote by W_2 (resp. W_3) the cobordisms from $S^3_{\Lambda_1}(L_1)$ to $S^3_{\Lambda_1 \cup \Lambda_2}(L_1 \cup L_2)$ (resp. from $S^3_{\Lambda_1 \cup \Lambda_2}(L_1 \cup L_2)$ to $S^3_{\Lambda_1 \cup \Lambda_2 \cup \Lambda_3}(L_1 \cup L_2 \cup L_3)$) given by surgery on L_2 (resp. L_3), ie consisting of two-handle additions, then

  $$W = W_1 \cup W_2 \cup W_3 \cup W_4$$

  is the cobordism from S^3 to S^3 obtained from X by deleting two copies of B^4.
- The manifold $N = S^3_{\Lambda_1 \cup \Lambda_2}(L_1 \cup L_2)$ is an admissible cut for W, ie $b_2^+(W_1 \cup W_2) > 0$, $b_2^+(W_3 \cup W_4) > 0$ and $\delta H^1(N) = 0$ in $H^2(W)$.

Lemma 14.8 *Any closed, oriented four-manifold X with $b_2^+(X) \geq 2$ admits a cut link presentation.*

Proof Start with a decomposition $W = W' \cup_N W''$ along an admissible cut. Split W' into three cobordisms

$$W' = W_1' \cup W_2' \cup W_3'$$

such that W_i' consists of i-handle additions only. It is easy to check that the decomposition

$$W = (W_1' \cup W_2') \cup (W_3' \cup W'')$$

is still along an admissible cut. Next, split the cobordism $W_3' \cup W''$ into

$$W_1'' \cup W_2'' \cup W_3''$$

such that W_i' consists of i-handle additions only. Finally, adjoin the one-handles from W_1'' to $W_1' \cup W_2'$ and rearrange the handles to obtain a decomposition

$$W_1' \cup W_2' \cup W_1'' = W_1 \cup W_2,$$

where W_i for $i = 1, 2$ consists of i-handle additions only. If we set $W_3 = W_2''$ and $W_4 = W_3''$, we obtain a decomposition along admissible cut of the form

$$W = (W_1 \cup W_2) \cup (W_3 \cup W_4),$$

We can then find a framed link $L = L_1 \cup L_2 \cup L_3$ such that surgery on L_1 produces the same three-manifold as at the end of the cobordism W_1 (made of one-handles), whereas surgery on L_2 and L_3 is represented by the cobordisms W_2 and W_3 (made of two-handles), respectively. □

Definition 14.9 Let X be a closed, oriented four-manifold with $b_2^+(X) \geq 2$. A *hyperbox presentation* Γ for X consists of a cut link presentation $(L = L_1 \cup L_2 \cup L_3, \Lambda)$ for X, together with a complete system of hyperboxes for L.

The four-manifold invariant $\Phi_{X,\mathfrak{t}}$ can be expressed in terms of a hyperbox presentation Γ for X as follows. Using Theorem 14.4, we can express the maps $F^-_{W_2,\mathfrak{t}|_{W_2}}$ and $F^+_{W_3,\mathfrak{t}|_{W_3}}$ in terms of counts of holomorphic polygons on a symmetric product of the surface. We can combine these maps using their factorization through HF_{red}, and obtain a mixed map

$$F^{\mathrm{mix}}_{W_2 \cup W_3, \mathfrak{t}|_{W_2 \cup W_3}} : \boldsymbol{HF}^-\bigl(\#^m (S^1 \times S^2)\bigr) \to HF^+\bigl(\#^{m'} (S^1 \times S^2)\bigr).$$

On the other hand, by composing the maps induced on homology by natural inclusions of chain complexes (of the kind used in Theorem 14.3), via factoring through a reduced group

$$\operatorname{Ker}\bigl(H_*(\mathscr{C}^-(\mathscr{H}, \Lambda)^{L_1 \cup L_2, \boldsymbol{t}|_{W_3}}) \to H_*(\mathscr{C}^\infty(\mathscr{H}, \Lambda)^{L_1 \cup L_2, \boldsymbol{t}|_{W_3}})\bigr)$$

we can construct a map

$$F^{\mathrm{mix}}_{\Gamma,\mathfrak{t}} : H_*(\mathscr{C}^-(\mathscr{H}, \Lambda)^{L_1, \mathfrak{t}|_{W_2 \cup W_3}}) \to H_*(\mathscr{C}^+(\mathscr{H}, \Lambda)^{L_1 \cup L_2 \cup L_3, \mathfrak{t}|_{\#^{m'}(S^1 \times S^2)}}).$$

Theorem 14.4 implies that $F^{\mathrm{mix}}_{\Gamma,\mathfrak{t}}$ is the same as $F^{\mathrm{mix}}_{W_2\cup W_3,\mathfrak{t}|_{W_2\cup W_3}}$, up to compositions with isomorphisms on both the domain and the target. Note, however, that at this point we do not know how to identify elements in the domains (or targets) of the two maps in a canonical way. For example, we know that there is an isomorphism

$$H_*(\mathcal{C}^-(\mathcal{H},\Lambda)^{L_1,\mathfrak{t}|_{W_2\cup W_3}}) \cong \boldsymbol{HF}^-\big(\#^m(S^1\times S^2)\big), \tag{257}$$

but it may be difficult to pinpoint what the isomorphism is, in terms of $\mathcal{H}$. Nevertheless, the good news is that $\boldsymbol{HF}^-\big(\#^m(S^1\times S^2)\big)$ has a unique maximal-degree element $\Theta^m_{\max}$. We can identify what $\Theta^m_{\max}$ corresponds to on the left-hand side of (257) by simply computing degrees. Let us denote the respective element by

$$\Theta^{\Gamma}_{\max} \in H_*(\mathcal{C}^-(\mathcal{H},\Lambda)^{L_1,\mathfrak{t}|_{W_2\cup W_3}}).$$

The following proposition says that one can decide whether $\Phi_{X,\mathfrak{t}}\in\mathbb{F}$ is zero or one from information in the hyperbox presentation Γ:

Theorem 14.10 *Let X be a closed, oriented four-manifold X with $b_2^+(X)\geq 2$, with a Spin^c structure $\mathfrak{t}$ with $d(\mathfrak{t})\geq 0$ even. Let Γ be a hyperbox presentation for X. Then $\Phi_{X,\mathfrak{t}}=1$ if and only if $U^{d(\mathfrak{t})/2}\cdot F^{\mathrm{mix}}_{\Gamma,\mathfrak{t}}(\Theta^{\Gamma}_{\max})$ is nonzero.*

Proof We have

$$\Phi_{X,\mathfrak{t}} = U^{d(\mathfrak{t})/2}\cdot F^{\mathrm{mix}}_{W,\mathfrak{t}}(\Theta_-) = F^+_{W_4,\mathfrak{t}|_{W_4}}\big(U^{d(\mathfrak{t})/2}\cdot F^{\mathrm{mix}}_{W_2\cup W_3,\mathfrak{t}|_{W_2\cup W_3}}\big(F^-_{W_1,\mathfrak{t}|_{W_1}}(\Theta_-)\big)\big).$$

By the definition of the one-handle addition maps from [Ozsváth and Szabó 2006, Section 4.3],

$$F^-_{W_1,\mathfrak{t}|_{W_1}}(\Theta_-) = \Theta^m_{\max}.$$

Note that $U^{d(\mathfrak{t})/2}\cdot F^{\mathrm{mix}}_{W_2\cup W_3,\mathfrak{t}|_{W_2\cup W_3}}(\Theta^m_{\max})$ lies in the minimal degree k for which

$$HF^+_k(S^3_{\Lambda_1\cup\Lambda_2\cup\Lambda_3}(L_1\cup L_2\cup L_3)) = HF^+_k\big(\#^{m'}(S^1\times S^2)\big)$$

is nonzero, namely $k=-\frac{1}{2}m'$. There is a unique nonzero element in the Floer homology in that degree, which is taken to 1 by the three-handle addition map $F^+_{W_4,\mathfrak{t}|_{W_4}}$; see [loc. cit.].

We deduce from here that $\Phi_{X,\mathfrak{t}}=1$ if and only if $U^{d(\mathfrak{t})/2}\cdot F^{\mathrm{mix}}_{W_2\cup W_3,\mathfrak{t}|_{W_2\cup W_3}}(\Theta^m_{\max})$ is nonzero. The claim then follows from the fact that the maps $F^{\mathrm{mix}}_{\Gamma,\mathfrak{t}}$ and $F^{\mathrm{mix}}_{W_2\cup W_3,\mathfrak{t}|_{W_2\cup W_3}}$ are the same up to pre- and post-composition with isomorphisms. □

14.4 The link surgeries spectral sequence

Our goal here will be to explain how the link surgeries spectral sequence from [Ozsváth and Szabó 2005, Section 4] can be understood in terms of complete systems of hyperboxes for links in S^3.

We recall the main result from [loc. cit.]. Let $M = M_1 \cup \cdots \cup M_\ell$ be a framed ℓ-component link in a three-manifold Y. For each $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell) \in \mathbb{E}_\ell = \{0,1\}^\ell$, we let $Y(\varepsilon)$ be the three-manifold obtained from Y by doing ε_i-framed surgery on M_i for $i = 1, \dots, \ell$.

When ε' is an immediate successor to ε (that is, when $\varepsilon < \varepsilon'$ and $\|\varepsilon' - \varepsilon\| = 1$), the two-handle addition from $Y(\varepsilon)$ to $Y(\varepsilon')$ induces a map on Heegaard Floer homology

$$F^-_{\varepsilon<\varepsilon'} \colon \boldsymbol{HF}^-(Y(\varepsilon)) \to \boldsymbol{HF}^-(Y(\varepsilon')).$$

The following is the link surgery spectral sequence [ibid., Theorem 4.1], but phrased here in terms of $\boldsymbol{HF}^-$ rather than $\widehat{HF}$ or HF^+:

Theorem 14.11 (Ozsváth and Szabó) *There is a spectral sequence whose E^1 term is $\bigoplus_{\varepsilon\in\mathbb{E}_\ell} \boldsymbol{HF}^-(Y(\varepsilon))$, whose d_1 differential is obtained by adding the maps $F^-_{\varepsilon<\varepsilon'}$ (for ε' an immediate successor to ε), and which converges to $E^\infty \cong \boldsymbol{HF}^-(Y)$.*

To relate this to the constructions in this paper, we represent $Y(0,\dots,0)$ itself as surgery on a framed link (L', Λ') inside S^3. Let $L'_1, \dots, L'_{\ell'}$ be the components of L'. There is another framed link $(L = L_1 \cup \cdots \cup L_\ell, \Lambda)$ in S^3, disjoint from L', such that surgery on each component L_i (with the given framing) corresponds exactly to the 2-handle addition from $Y(0,\dots,0)$ to $Y(0,\dots,0,1,0,\dots,0)$, where the 1 is in position i. For $\varepsilon \in \mathbb{E}_\ell$, we denote by L^ε the sublink of L consisting of those components L_i such that $\varepsilon_i = 1$.

Let $\mathscr{H}$ be a complete system of hyperboxes for the link $L' \cup L \subset S^3$. As mentioned in Section 14.1, inside the surgery complex $\mathscr{C}^-(\mathscr{H}, \Lambda' \cup \Lambda)$ (which is an $(\ell'+\ell)$-dimensional hypercube of chain complexes) we have various subcomplexes corresponding to surgery on the sublinks on $L' \cup L$. We will restrict our attention to those sublinks that contain L', and use the respective subcomplexes to construct a new, ℓ-dimensional hypercube of chain complexes $\mathscr{C}^-(\mathscr{H}, \Lambda' \cup \Lambda /\!/ L)$ as follows.

At a vertex $\varepsilon \in \mathbb{E}_\ell$ we put the complex

$$\mathscr{C}^-(\mathscr{H}, \Lambda' \cup \Lambda /\!/ L)^\varepsilon = \mathscr{C}^-(\mathscr{H}|_{L'\cup L^\varepsilon}, \Lambda' \cup \Lambda|_{L^\varepsilon}).$$

Consider now an edge from ε to $\varepsilon' = \varepsilon + \tau_i$ in the hypercube $\mathbb{E}_\ell$. The corresponding complex $\mathscr{C}^-(\mathscr{H}|_{L'\cup L^\varepsilon}, \Lambda' \cup \Lambda|_{L^\varepsilon})$ decomposes as a direct product over all Spin^c structures $\boldsymbol{s}$ on

$$Y(\varepsilon) = S^3(L' \cup L^\varepsilon, \Lambda' \cup \Lambda|_{L^\varepsilon}).$$

As we explained in Section 14.1, each factor $\mathscr{C}^-(\mathscr{H}|_{L'\cup L^\varepsilon}, \Lambda' \cup \Lambda|_{L^\varepsilon}, \boldsymbol{s})$ admits an inclusion into $\mathscr{C}^-(\mathscr{H}|_{L'\cup L^{\varepsilon'}}, \Lambda' \cup \Lambda|_{L^{\varepsilon'}})$ as a subcomplex. In fact, there are several such inclusion maps, one for each Spin^c structure $\boldsymbol{t}$ on the 2-handle cobordism from $Y(\varepsilon)$ to $Y(\varepsilon')$ such that $\boldsymbol{t}$ restricts to $\boldsymbol{s}$ on $Y(\varepsilon)$. Adding up all the inclusion maps on each factor, one obtains a combined map

$$G^-_{\varepsilon<\varepsilon'} \colon \mathscr{C}^-(\mathscr{H}|_{L'\cup L^\varepsilon}, \Lambda' \cup \Lambda|_{L^\varepsilon}) \to \mathscr{C}^-(\mathscr{H}|_{L'\cup L^{\varepsilon'}}, \Lambda' \cup \Lambda|_{L^{\varepsilon'}}).$$

We take $G^-_{\varepsilon<\varepsilon'}$ to be the edge map in the hypercube of chain complexes $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$. Since the edge maps are just sums of inclusions of subcomplexes, they commute on the nose along each face of the hypercube. Therefore, in the hypercube $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$ we can take the diagonal maps to be zero, along all faces of dimension at least two.

This completes the construction of $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$. As an ℓ-dimensional hypercube of chain complexes, its total complex admits a filtration by $-\|\varepsilon\|$, which induces a spectral sequence; we will refer to the filtration by $-\|\varepsilon\|$ as the *depth filtration* on $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$.

Theorem 14.12 *Fix a complete system of hyperboxes $\mathcal{H}$ for an oriented link $\vec{L}' \cup \vec{L}$ in S^3, and fix framings Λ for L and Λ' for L'. Suppose $\mathcal{H}$ has k basepoints of type w and p colors, and that L has ℓ components $L_1, \dots, L_\ell$. Let $Y(0, \dots, 0) = S^3_{\Lambda'}(L')$, and let $Y(\varepsilon)$ be obtained from $Y(0, \dots, 0)$ by surgery on the components $L_i \subseteq L$ with $\varepsilon_i = 1$ (for any $\varepsilon \in \mathbb{E}_\ell$). Then there is an isomorphism between the link surgeries spectral sequence from Theorem 14.11 and the spectral sequence associated to the depth filtration on $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$.*

Proof Theorem 1.1 gives a quasi-isomorphism between $\boldsymbol{CF}^-(Y(1, \dots, 1)) = \boldsymbol{CF}^-(S^3_{\Lambda' \cup \Lambda}(L' \cup L))$ and the surgery complex $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda)$. Let us summarize the main steps in the construction of this quasi-isomorphism. It suffices to construct the quasi-isomorphism at the level of the vertical truncations $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}, \Lambda' \cup \Lambda)$ and $\boldsymbol{CF}^{-,\delta\leftarrow\delta'}(Y(1, \dots, 1))$. (In the proof, we used vertical truncations by δ for torsion Spin^c structures and by $\delta\leftarrow\delta'$ for nontorsion Spin^c structures. However, it is clear that always truncating by $\delta\leftarrow\delta'$ works, too.) We consider a basic system $\mathcal{H}_b$ for $L' \cup L$. By relating $\mathcal{H}$ to $\mathcal{H}_b$ via some moves on complete systems, we find a quasi-isomorphism between $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}, \Lambda' \cup \Lambda)$ and $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}_b, \Lambda' \cup \Lambda)$. By a version of the large surgeries theorem, the hypercube $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}_b, \Lambda' \cup \Lambda)$ is shown to be quasi-isomorphic to a hypercube $\mathscr{H}^{\delta\leftarrow\delta'}$, in which at each vertex we have the truncated Floer complex of some surgery on a sublink of $L' \cup L$, possibly with twisted coefficients. Finally, iterating a variant of the surgery exact triangle, we obtain a quasi-isomorphism between $\boldsymbol{CF}^{-,\delta\leftarrow\delta'}(Y(1, \dots, 1))$ and the total complex of $\mathscr{H}^{\delta\leftarrow\delta'}$.

Recall that from the $(\ell+\ell')$-dimensional hypercube $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda)$ we construct an ℓ-dimensional hypercube $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$. In the new hypercube, at each vertex $\varepsilon \in \mathbb{E}_\ell$ we have an $\|\varepsilon\|$-dimensional subhypercube of the original $\mathcal{C}^-(\mathcal{H}, \Lambda' \cup \Lambda)$; along the edges we have corresponding inclusion maps, and along higher-dimensional faces the diagonal maps are trivial. We can apply an analogous procedure to the truncated hypercubes $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}, \Lambda' \cup \Lambda)$, $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}_b, \Lambda' \cup \Lambda)$ and $\mathscr{H}^{\delta\leftarrow\delta'}$, and obtain ℓ-dimensional hypercubes from them; we denote these by $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$, $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}_b, \Lambda' \cup \Lambda /\!\!/ L)$ and $\mathscr{H}^{\delta\leftarrow\delta'} /\!\!/ L$. The quasi-isomorphisms between $(\ell+\ell')$-dimensional hypercubes constructed in the proof of Theorem 1.1 all preserve the corresponding depth filtrations, and thus induce quasi-isomorphisms between respective subhypercubes. As a consequence, we can construct filtered quasi-isomorphisms (with respect to the depth filtration) between $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}, \Lambda' \cup \Lambda /\!\!/ L)$, $\mathcal{C}^{-,\delta\leftarrow\delta'}(\mathcal{H}_b, \Lambda' \cup \Lambda /\!\!/ L)$ and

$(\mathscr{H}^{\delta\leftarrow\delta'} /\!\!/ L)$. This implies that the corresponding spectral sequences (induced by the depth filtrations on ℓ-dimensional hypercubes) are isomorphic.

It remains to find an isomorphism between the spectral sequence induced by the depth filtration on the hypercube $\mathscr{H}^{\delta\leftarrow\delta'} /\!\!/ L$, and the $\delta\leftarrow\delta'$ truncation of the link surgeries spectral sequence from Theorem 14.11. Let us first explain how this is done in the simplest case, when $\ell = 1$ so that the link L has a single component $L_1 = K$. We further assume that K has linking number zero with each component of L'. We will drop the truncation symbol $\delta\leftarrow\delta'$ from notation for simplicity (in fact, in this step of the argument the quasi-isomorphism exists also at the level of untruncated complexes). The spectral sequence from Theorem 14.11 is simply associated to the depth filtration on a one-dimensional hypercube, which is the mapping cone

$$\boldsymbol{CF}^-(S^3_{\Lambda'}(L')) \to \boldsymbol{CF}^-(S^3_{\Lambda'\cup\Lambda}(L'\cup K)). \tag{258}$$

On the other hand, the hypercube $\mathscr{H} /\!\!/ L$ is a mapping cone

$$(\mathscr{H} /\!\!/ L)^0 \to (\mathscr{H} /\!\!/ L)^1, \tag{259}$$

where $(\mathscr{H} /\!\!/ L)^0$ is $\boldsymbol{CF}^-(S^3_{\Lambda'}(L'))$ and $(\mathscr{H} /\!\!/ L)^1$ is itself a mapping cone

$$\boldsymbol{CF}^-(S^3_{\Lambda'\cup(\Lambda+m)}(L'\cup K)) \to \bigoplus^m \boldsymbol{CF}^-(S^3_{\Lambda'}(L')) \tag{260}$$

for some $m \gg 0$.

The surgery exact triangle (Proposition 11.5) says that the second term in (258) is quasi-isomorphic to the second term in (259). In fact, we can take a quasi-isomorphism given by a triangle-counting map from $\boldsymbol{CF}^-(S^3_{\Lambda'\cup\Lambda}(L'\cup K))$ to the first term in (260), and a quadrilateral-counting map to the second term in (258). We can extend this quasi-isomorphism to one between the mapping cone (258) and the mapping cone (259), by taking the identity between their first terms $\boldsymbol{CF}^-(S^3_{\Lambda'}(L'))$, and also adding a diagonal map from $\boldsymbol{CF}^-(S^3_{\Lambda'}(L'))$ to the mapping cone (259). This diagonal map consists of a quadrilateral-counting map from $\boldsymbol{CF}^-(S^3_{\Lambda'}(L'))$ to the first term in (260), and a pentagon-counting map from $\boldsymbol{CF}^-(S^3_{\Lambda'}(L'))$ to the second term in (260). This produces a chain map between (258) and (259), which is a quasi-isomorphism because it is so on the level of the associated graded of the depth filtrations.

Let us now discuss how this construction can be generalized to the case when the link L has an arbitrary number ℓ of connected components. The link surgeries spectral sequence from Theorem 14.11 is associated to the depth filtration on an ℓ-dimensional hypercube, where at each vertex we have a Heegaard Floer complex $\boldsymbol{CF}^-(Y(\varepsilon))$, along the edges we have triangle-counting maps (producing the cobordism maps $F^-_{\varepsilon<\varepsilon'}$ on the E^1 page), and along the higher-dimensional faces we have higher polygon-counting maps. We can construct a filtered quasi-isomorphism between this hypercube and $\mathscr{H} /\!\!/ L$ as follows: between corresponding vertices we use the quasi-isomorphisms given iterating the surgery exact triangle (Proposition 11.5), and then we complete this to a chain map between the hypercubes, by adding further polygon-counting maps along the diagonals. A filtered quasi-isomorphism between ℓ-dimensional hypercubes produces an isomorphism between the respective spectral sequences, and this completes the proof. □

15 The surgery theorem applied to grid diagrams with free markings

In this section we state a variant of the surgery theorem, in terms of counts of polygons on grid diagrams with free markings. The proof involves applying Theorem 1.1 to a special kind of complete system, associated to the grid.

15.1 Grid diagrams with free markings

Toroidal grid diagrams are a particular kind of Heegaard diagrams for a link in S^3. In [Manolescu et al. 2009a; 2007], they have been used to give combinatorial descriptions to link Floer complexes.

We introduce here a slightly more general concept, that of a *toroidal grid diagram with free markings*. (An example is shown in Figure 45.) Such a diagram G consists of a torus $\mathcal{T}$, viewed as a square in the plane with the opposite sides identified, and split into n annuli (called rows) by n horizontal circles $\alpha_1, \dots, \alpha_n$, and into n other annuli (called columns) by n vertical circles $\beta_1, \dots, \beta_n$. Further, we are given several markings on the torus, of two types X and O, such that

- each row and each column contains exactly one O marking;
- each row and each column contains at most one X marking;
- if the row of an O marking contains no X markings, then the column of that O marking contains no X markings either. An O marking of this kind is called a *free marking*.

It follows that G contains exactly n O markings and $n-q$ X markings, where q is the number of free markings. A marking that is not free is called *linked*. The number n is called the *grid number* of G.

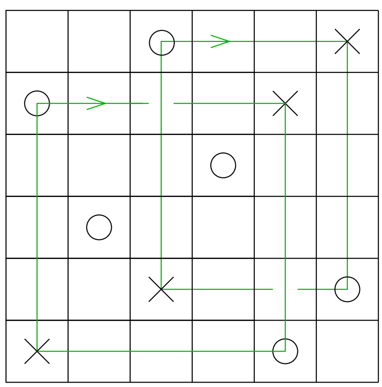

Figure 45: A grid diagram for the Hopf link, with two free markings. The link is drawn in green.

Given G as above, we draw horizontal arcs between the (linked) markings in the same row, and vertical arcs between the markings in the same column. Letting the vertical arcs be overpasses whenever they intersect the horizontal arcs, we then obtain a planar diagram for a link $\vec{L} \subset S^3$, which we orient so that all horizontal arcs go from an O to an X. We denote by ℓ the number of components of L.

A toroidal grid is a type of multipointed Heegaard diagram for $\vec{L}$, with the horizontal circles being the alpha curves $\alpha_1, \dots, \alpha_n$, and the vertical curves being the beta curves $\beta_1, \dots, \beta_n$. We let $\boldsymbol{S} = \boldsymbol{S}(G) = \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, where $\mathbb{T}_\alpha = \alpha_1 \times \dots \times \alpha_n$ and $\mathbb{T}_\beta = \beta_1 \times \dots \times \beta_n$ are tori in the symmetric product $\mathrm{Sym}^n(\mathcal{T})$.

Following the notation of [Manolescu et al. 2007], we use X and O markings to play the role of the basepoints z_i and w_i, respectively. In particular, we change the notation from Section 3.2 and write X_i and O_i for n_{z_i} and n_{w_i}, respectively. Further, we let $\mathbb{X}_i$ and $\mathbb{O}_i$ be the sets of X and O markings on a link component L_i. We write $\mathbb{X} = \bigcup_i \mathbb{X}_i$ for the set of all X markings, and $\mathbb{O}$ for the set of all O markings (including the linked ones). We also fix an ordering of the basepoints in each set $\mathbb{X}_i$ and $\mathbb{O}_i$. Note that the free markings are exactly the free basepoints; we arrange so that they are denoted by $O_{n-q+1}, \dots, O_n$.

For $\boldsymbol{x}, \boldsymbol{y} \in \boldsymbol{S}$, we let $\mathrm{Rect}(\boldsymbol{x}, \boldsymbol{y})$ be the set of rectangles from $\boldsymbol{x}$ to $\boldsymbol{y}$, as in [loc. cit.]. A rectangle from $\boldsymbol{x}$ to $\boldsymbol{y}$ can exist only when $\boldsymbol{x}$ and $\boldsymbol{y}$ differ in exactly two rows (and therefore, also in exactly two columns). In that case, those two rows and two columns split the torus $\mathcal{T}$ into four rectangles, and two of them are said to go from $\boldsymbol{x}$ to $\boldsymbol{y}$: those such that, when their horizontal boundary has the orientation induced from $\mathcal{T}$, it consists of arcs oriented from a point in $\boldsymbol{x}$ to a point in $\boldsymbol{y}$. Furthermore, a rectangle is called *empty* if it contains no points of $\boldsymbol{x}$ (or, equivalently, $\boldsymbol{y}$) in its interior; it may still contain some O and X markings. We let $\mathrm{Rect}^\circ(\boldsymbol{x}, \boldsymbol{y})$ be the set of empty rectangles from $\boldsymbol{x}$ to $\boldsymbol{y}$.

The arguments in [Manolescu et al. 2009a] show that isolated holomorphic disks in $\mathrm{Sym}^n(\mathcal{T})$ with boundaries on $\mathbb{T}_\alpha, \mathbb{T}_\beta$ are in a natural one-to-one correspondence to empty rectangles on the grid G. Thus, the Floer chain complex $\boldsymbol{CFL}^-(G) = \boldsymbol{CFL}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta)$ is the free module over $\mathcal{R} = \mathbb{F}[\![U_1, \dots, U_n]\!]$ generated by $\boldsymbol{S}$, and endowed with the differential

$$\partial \boldsymbol{x} = \sum_{\boldsymbol{y} \in \boldsymbol{S}} \sum_{r \in \mathrm{Rect}^\circ(\boldsymbol{x}, \boldsymbol{y})} U_1^{O_1(r)} \cdots U_n^{O_n(r)} \boldsymbol{y}. \tag{261}$$

We have Alexander gradings A_i on $\boldsymbol{S}$ for $i = 1, \dots, \ell$, which produce filtrations on $\boldsymbol{CFL}^-(G)$. For $\boldsymbol{s} = (s_1, \dots, s_\ell) \in \overline{\mathbb{H}}(L)$, we denote by $\boldsymbol{A}^-(G, \boldsymbol{s}) \subset \boldsymbol{CFL}^-(G)$ the subcomplex given by $A_i \le s_i$ for $i = 1, \dots, \ell$.

We have a resolution $\mathfrak{A}^-(G, \boldsymbol{s})$ of $\boldsymbol{A}^-(G, \boldsymbol{s})$ constructed as in Section 13.6. Specifically, let us relabel the basepoints on each component $L_i \subseteq L$ as

$$O_{i,1}, X_{i,1}, O_{i,2}, X_{i,2}, \dots, O_{i,p_i}, X_{i,p_i}.$$

We also give the alternative name $U_{i,j}$ to the U variable associated to $O_{i,j}$.

Consider the complex $\mathfrak{C}_{\mathrm{int}}(G)$, which is the dg module over the dga

$\mathfrak{R}^Y =$

$\mathbb{F}[\![(U_{i,j})_{1\le i\le \ell, 1\le j\le p_i}(U_j)_{n-q+1\le j\le p}]\!][(V_{i,j},Y_{i,j})_{1\le i\le\ell,2\le j\le p_i}]'/(Y_{i,j}^2=0, \partial Y_{i,j}=U_{i,j}+U_{i,1}V_{i,j})$,

generated by $\boldsymbol{x}\in \boldsymbol{S}$, with differential

$$\partial \boldsymbol{x}=\sum_{\boldsymbol{y}\in\boldsymbol{S}}\sum_{r\in\mathrm{Rect}^\circ(\boldsymbol{x},\boldsymbol{y})}\prod_{i=1}^{\ell}U_{i,1}^{X_{i,1}(r)+\cdots+X_{i,p_i}(r)}V_{i,2}^{O_{i,2}(r)}\cdots V_{i,p}^{O_{i,p}(r)}\cdot\prod_{j=n-q+1}^{n}U_j^{O_j(r)}\boldsymbol{y}.$$

This complex admits filtrations $\mathfrak{F}_i$ for $1\le i\le \ell$, as in Section 13.6. For $\boldsymbol{s}=(s_1,\dots,s_\ell)\in\mathbb{H}(L)$, we define $\mathfrak{A}^-(\mathscr{H}^L,\boldsymbol{s})$ to be the filtered part of $\mathfrak{C}_{\mathrm{int}}(\mathscr{H}^L)$ given by $\mathfrak{F}_i\le s_i$ for $i=1,\dots,\ell$.

If $M\subseteq L$ is a sublink with an orientation $\vec{M}$, we have a complex $\mathfrak{A}^-(G,\vec{M},p^{\vec{M}}(\boldsymbol{s}))$ and a projection-inclusion map

$$\mathscr{I}_{\boldsymbol{s}}^{\vec{M}}:\mathfrak{A}^-(G,\boldsymbol{s})\to\mathfrak{A}^-(G,\vec{M},p^{\vec{M}}(\boldsymbol{s})), \tag{262}$$

defined as in (234).

We also introduce the following notation: $\mathbb{X}^{\vec{M}}\subseteq\mathbb{X}$ will be the subset consisting of the X basepoints on $L-M$, and

$$\mathbb{O}^{\vec{M}}:=\Bigl(\mathbb{O}-\bigcup_{i\in I_-(\vec{L},\vec{M})}\mathbb{O}_i\Bigr)\cup\bigcup_{i\in I_-(\vec{L},\vec{M})}\mathbb{X}_i.$$

Thus, $(\mathscr{T},\boldsymbol{\alpha},\boldsymbol{\beta},\mathbb{O}^{\vec{M}},\mathbb{X}^{\vec{M}})$ is the reduction $r_{\vec{M}}(G)$, in the sense of Definition 3.22.

15.2 Handleslides over a set of markings

From now on we will specialize to the situation in which G is a toroidal grid diagram *with at least one free marking.*

Consider a subset $\mathfrak{Z}=\{Z_1,\dots,Z_k\}\subseteq\mathbb{X}\cup\mathbb{O}$ consisting only of linked markings. We say that $\mathfrak{Z}$ is *consistent* if, for any i, at most one of the sets $\mathfrak{Z}\cap\mathbb{O}_i$ and $\mathfrak{Z}\cap\mathbb{X}_i$ is nonempty. From now on we shall assume that $\mathfrak{Z}$ is consistent.

We let $L(\mathfrak{Z})\subseteq L$ be the sublink consisting of those components L_i such that at least one of the markings on L_i is in $\mathfrak{Z}$. We orient $L(\mathfrak{Z})$, as $\vec{L}(\mathfrak{Z})$, so that a component L_i is given the orientation coming from $\vec{L}$ when $\mathfrak{Z}\cap\mathbb{O}_i\neq\varnothing$, and is given the opposite orientation when $\mathfrak{Z}\cap\mathbb{X}_i\neq\varnothing$.

Let us define a new set of curves $\boldsymbol{\beta}^{\mathfrak{Z}}=\{\beta_j^{\mathfrak{Z}}\mid j=1,\dots,n\}$ on the torus $\mathscr{T}$. Let j_i be the index corresponding to the vertical circle β_{j_i} just to the left of a marking $Z_i\in\mathfrak{Z}$. We let $\beta_{j_i}^{\mathfrak{Z}}$ be a circle encircling Z_i and intersecting β_{j_i}, as well as the alpha curve just below Z_i, in two points each; in other words, $\beta_{j_i}^{\mathfrak{Z}}$ is obtained from β_j by handlesliding it over the vertical curve just to the right of Z_i. For those j that are not j_i for any $Z_i\in\mathfrak{Z}$, we let $\beta_j^{\mathfrak{Z}}$ be a curve isotopic to β_j and intersecting it in two points.

Remark 15.1 Our assumption on the existence of a free marking is crucial here, because it ensures that $\boldsymbol{\beta}^{\mathscr{X}}$ is a good set of attaching curves. Indeed, since $\mathscr{X}$ has strictly fewer than n markings, at least one beta curve in $\boldsymbol{\beta}^{\mathscr{X}}$ remains homologically nontrivial in $H_1(\mathscr{T})$.

We denote

$$\mathbb{T}_\beta^{\mathscr{X}} = \beta_1^{\mathscr{X}} \times \cdots \times \beta_n^{\mathscr{X}} \subset \mathrm{Sym}^n(\mathscr{T}).$$

Observe that

$$(\mathscr{T}, \boldsymbol{\alpha}, \boldsymbol{\beta}^{\mathscr{X}}, \mathbb{O}^{\vec{L}(\mathscr{X})}, \mathbb{X}^{\vec{L}(\mathscr{X})})$$

is a multipointed Heegaard diagram representing the link $\vec{L} - L(\mathscr{X})$. For each $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, we can associate to this diagram a resolved Floer complex

$$\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathscr{X}}, \psi^{\vec{L}(\mathscr{X})}(\boldsymbol{s})),$$

as in Section 13.6. For simplicity, we omit the basepoints from the notation for this chain complex. The map $\psi^{\vec{L}(\mathscr{X})} : \overline{\mathbb{H}}(L) \to \overline{\mathbb{H}}(L - L(\mathscr{X}))$ is as in Section 3.7.

More generally, let $\vec{M}$ be any sublink of L containing $\vec{L}(\mathscr{X})$ such that the restriction to $L(\mathscr{X})$ of the orientation on $\vec{M}$ coincides with $\vec{L}(\mathscr{X})$. We can then consider a diagram

$$(\mathscr{T}, \boldsymbol{\alpha}, \boldsymbol{\beta}^{\mathscr{X}}, \mathbb{O}^{\vec{M}}, \mathbb{X}^{\vec{M}})$$

and resolved Floer complexes

$$\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathscr{X}}, \psi^{\vec{M}}(\boldsymbol{s})).$$

By writing $\psi^{\vec{M}}(\boldsymbol{s})$ we implicitly mean that the basepoints are taken to be those in $\mathbb{O}^{\vec{M}}$ and $\mathbb{X}^{\vec{M}}$.

When we have two collections of markings $\mathscr{X}$ and $\mathscr{X}'$ such that $\mathscr{X} \cup \mathscr{X}'$ is consistent, we will require that $\beta_i^{\mathscr{X}}$ and $\beta_i^{\mathscr{X}'}$ intersect in exactly two points. Hence, there is always a unique maximal-degree intersection point $\Theta^{\mathrm{can}}_{\mathscr{X},\mathscr{X}'} \in \mathbb{T}_\beta^{\mathscr{X}} \cap \mathbb{T}_\beta^{\mathscr{X}'}$. See Figure 46.

Of course, for the above Floer complexes to be well defined, we need to make sure that the underlying Heegaard diagrams are admissible. This follows from the stronger admissibility result in Lemma 15.2 below.

15.3 Handleslides for a sublink

Let $M \subseteq \vec{L}$ be a sublink, endowed with an arbitrary orientation $\vec{M}$. We seek to define a descent map for the sublink $\vec{M}$, along the lines of Section 9.1.

Set

$$\mathscr{X}(\vec{M}) = \bigcup_{i \in I_+(\vec{L}, \vec{M})} \mathbb{O}_i \cup \bigcup_{i \in I_-(\vec{L}, \vec{M})} \mathbb{X}_i.$$

We proceed to construct a hyperbox $\mathscr{H}_G^{\vec{L},\vec{M}}$ for the pair $(\vec{L}, \vec{M})$ (in the sense of Definition 8.20), as follows. Order the components of M according to their ordering as components of L:

$$M = L_{i_1} \cup \cdots \cup L_{i_m}, \quad i_1 < \cdots < i_m.$$

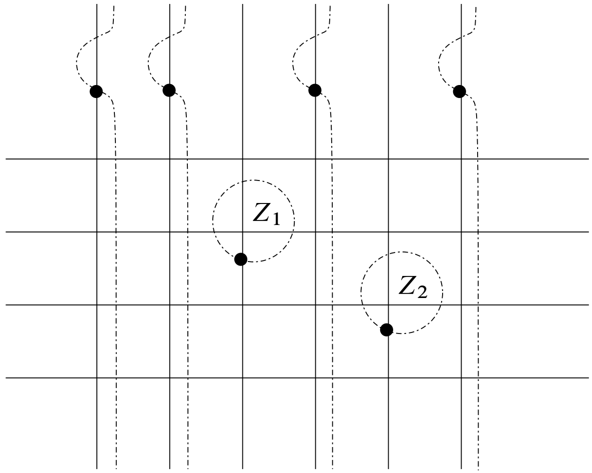

Figure 46: A new collection of curves. We show here a part of a grid diagram, with the horizontal segments lying on curves in $\boldsymbol{\alpha}$ and the straight vertical segments lying on curves in $\boldsymbol{\beta}$. The interrupted curves (including the two circles) represent curves in $\boldsymbol{\beta}^{\mathcal{Z}}$, where $\mathcal{Z}$ consists of the two markings Z_1 and Z_2. The maximal-degree intersection point $\Theta^{\mathrm{can}}_{\varnothing,\mathcal{Z}}$ is represented by the black dots.

For $j = 1, \dots, m$, let us denote $M_j = L_{i_j}$ for simplicity, and equip M_j with the orientation $\vec{M}_j$ induced from $\vec{M}$. Then $\mathcal{Z}(\vec{M}_j)$ is either $\mathbb{O}_{i_j}$ or $\mathbb{X}_{i_j}$. In either case, we have an ordering of its elements, so we can write

$$\mathcal{Z}(\vec{M}_j) = \{Z_1^{\vec{M}_j}, \dots, Z_{d_{i_j}}^{\vec{M}_j}\},$$

where d_{i_j} is the cardinality of $\mathcal{Z}(\vec{M}_j)$.

The hyperbox $\mathcal{H}_G^{\vec{L},\vec{M}}$ will be m-dimensional, of size

$$\boldsymbol{d}^M = (d_{i_1}, \dots, d_{i_m}).$$

It will be a pure β-hyperbox, ie the alpha curves remain fixed: they are the horizontal circles in the grid diagram G. For each multi-index $\varepsilon = (\varepsilon_1, \dots, \varepsilon_m) \in \mathbb{E}(\boldsymbol{d}^M)$, we let $\mathcal{Z}(\vec{M})^{\varepsilon} \subseteq \mathcal{Z}(\vec{M})$ be the collection of markings

$$\mathcal{Z}(\vec{M})^{\varepsilon} = \bigcup_{j=1}^{m} \{Z_1^{\vec{M}_j}, \dots, Z_{\varepsilon_j}^{\vec{M}_j}\}.$$

We then let

$$\boldsymbol{\beta}^{\varepsilon} = \boldsymbol{\beta}^{\mathcal{Z}(\vec{M})^{\varepsilon}}$$

be the collection of beta curves handleslid at the points of $\mathcal{Z}(\vec{M})^{\varepsilon}$. For each ε, consider the Heegaard diagram

$$\mathcal{H}_{G,\varepsilon}^{\vec{L},\vec{M}} = (\mathcal{T}, \boldsymbol{\alpha}, \boldsymbol{\beta}^{\varepsilon}, \mathbb{O}^{\vec{M}}, \mathbb{X}^{\vec{M}}).$$

This diagram represents the link $\vec{L} - M$.

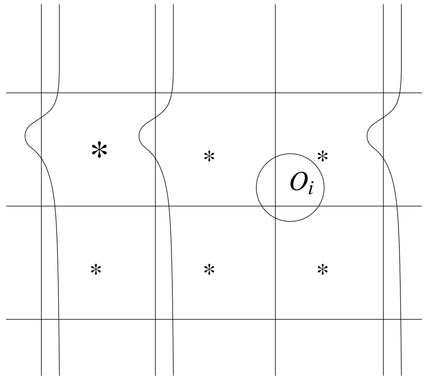

Figure 47: Asterisks in the Heegaard multidiagram.

According to Definitions 8.7 and 8.20, to have a well-defined hyperbox for a pair, we need admissibility for the Heegaard multidiagrams corresponding to each hypercube in the hyperbox. In fact, it is convenient to consider the bigger multidiagram that includes the curve collections $\boldsymbol{\beta}^{\mathfrak{L}}$ for all subsets $\mathfrak{L} \subseteq \mathfrak{L}(\vec{M})$. Once we prove that diagram is admissible, it will follow that all its subdiagrams (such as the ones relevant for our hypercubes) are admissible as well.

Lemma 15.2 *Consider the Heegaard multidiagram consisting of the curve collections $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}^{\mathfrak{L}}$ on $\mathcal{T}$ for all possible $\mathfrak{L} \subseteq \mathfrak{L}(\vec{M})$, with basepoints given by $\mathbb{O}^{\vec{M}}$ and $\mathbb{X}^{\vec{M}}$. Then this multidiagram is admissible, in the sense of Definition 3.5.*

Proof The markings in $\mathbb{X}^{\vec{M}}$ are irrelevant for admissibility, so we can ignore them. There is one marking in $\mathbb{O}^{\vec{M}}$ in each row and in each column of the grid G. Without loss of generality, let us assume that $\vec{M}$ is oriented as in $\vec{L}$, so that $\mathbb{O}^{\vec{M}} = \mathbb{O} = \{O_1, \dots, O_n\}$.

Recall from Definition 3.5 that a multiperiodic domain is a $\mathbb{Z}$-linear combination of regions such that its boundary is a $\mathbb{Z}$-linear combination of the curves in the multidiagram, and such that its multiplicity at any O_i is zero. Admissibility means there are no nontrivial multiperiodic domains that have only nonnegative local multiplicities. By way of contradiction, suppose we had such a domain D.

Let us mark an asterisk in each square of the original grid diagram G. When we construct the new beta curves $\beta_j^{\mathfrak{L}}$, we make sure that each beta curve encircling some O_i does not go around an asterisk, and also that, whenever we isotope a beta curve to obtain a new beta curve intersecting it at two points, these isotopies do not cross the asterisks. See Figure 47.

Let $a_j \in \mathbb{Z}$ be the coefficient with which the curve α_j appears in ∂D. This means that, whenever we have two regions in the multidiagram separated by a segment in α_j, the difference between the multiplicities of D on these two regions equals a_j. Observe that there are two such regions such that one on top of α_j contains an O marking. The multiplicity of D is supposed to be zero at that marking and nonnegative at the region below. It follows that $a_j \leq 0$, and therefore the multiplicity increases (or stays constant)

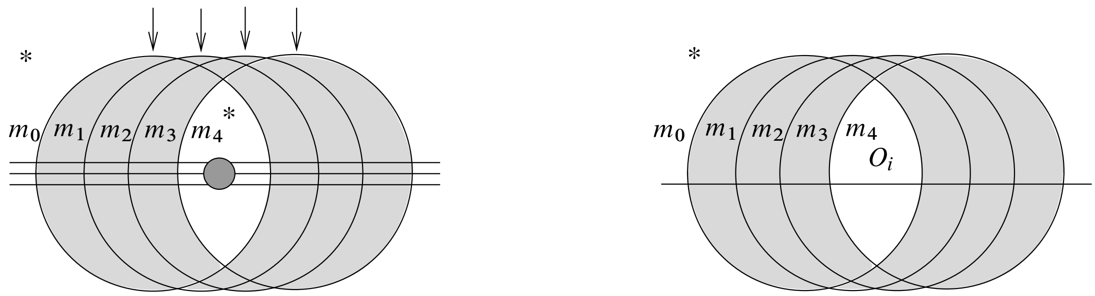

Figure 48: The Hamiltonian translates of a beta curve. On the left we have the neighborhood of a vertical beta curve on the grid, shown here as a circle. The other circles are Hamiltonian translates of the same curve. The horizontal lines are alpha curves, and the disk in the middle is a handle that is connected to the exterior of the picture to produce the torus $\mathcal{T}$. The shaded area is the potential support of a periodic domain, in the proof of Lemma 15.2. On the right we show a similar neighborhood of a beta curve encircling an O marking.

whenever we cross an alpha curve downward. In particular, if we look at the multiplicities of D at the regions containing asterisks in a given column, they should be nondecreasing in cyclic order, which means that they are constant.

We can do a similar argument with the beta curves, by letting b_j be the sum of the coefficients in ∂D of β_j and all of its Hamiltonian translates. In this case b_j is the difference in multiplicities as we cross all the translates of β_j. By looking at the region containing the O marking in each column, we find that b_j are all of the same sign. This shows that the multiplicities of D at the regions containing asterisks in a given row are constant.

Combining these two results, we deduce that D has the same multiplicity at all the regions containing asterisks. However, we know there is at least one free O marking in G, so that no handleslides are done over that marking. Therefore, the region containing that marking has an asterisk, and D has multiplicity zero there. It follows that D has multiplicity zero at all asterisks.

This means that D is a sum of disjoint periodic domains, each supported in the area bounded by the Hamiltonian translates of the same curve: either a vertical beta curve, or a curve encircling an O marking; see Figure 48. We can study these domains separately.

Let us consider the neighborhood of a vertical beta curve, as on the left-hand side of Figure 48. (The other case is similar.) Let $c_1, \dots, c_p$ be the coefficients in ∂D of each of the beta curves in the picture. As we cross a beta curve (say, in the directions shown by arrows in Figure 48), the multiplicity of D should changed by the corresponding c_i. Note that the outside region at the top of the picture has an asterisk and hence multiplicity zero. Therefore, $c_i \geq 0$ for all i. Let $m_0 \leq m_1 \leq \cdots \leq m_p$ be the multiplicities at the regions shown in the figure, as we go from left to right. We must have $m_0 = m_p = 0$, because those regions contain asterisks. Also, $m_i = m_{i-1} + c_i$, which implies that

$$0 = m_0 \leq m_1 \leq \cdots \leq m_{p-1} \leq m_p = 0.$$

This shows that $m_i = 0$ and $c_i = 0$ for all i, which implies that the domain is trivial. We arrived at a contradiction, and therefore we have admissibility. $\square$

To finish the construction of the hyperbox $\mathcal{H}_G^{\vec{L},\vec{M}}$, whenever $\varepsilon, \varepsilon' \in \mathbb{E}(\boldsymbol{d}^M)$ with $\varepsilon < \varepsilon'$ are neighbors, we need to pick chain elements $\Theta_{\varepsilon,\varepsilon'} \in \boldsymbol{A}^-(\mathbb{T}_{\beta^\varepsilon}, \mathbb{T}_{\beta^{\varepsilon'}}, \boldsymbol{0})$. We choose

$$\Theta_{\varepsilon,\varepsilon'} = \begin{cases} \Theta^{\mathrm{can}}_{\mathfrak{X}(\vec{M})^\varepsilon, \mathfrak{X}(\vec{M})^{\varepsilon'}} & \text{if } \|\varepsilon' - \varepsilon\| = 1, \\ 0 & \text{otherwise.} \end{cases}$$

The proof of (90) for these chains is similar to those of [Ozsváth and Szabó 2004b, Lemma 9.7; 2005, Lemma 4.3].

As in Section 13.6, for each $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, we have an associated hyperbox of generalized Floer complexes

$$\mathfrak{C}(\mathcal{H}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s})),$$

of size $\boldsymbol{d}'^M := (d'_{i_1}, \dots, d'_{i_m})$, where

$$d'_{i_j} = \begin{cases} d_{i_j} & \text{if } i_j \in I_+(\vec{L}, \vec{M}), \\ 2d_{i_j} - 1 & \text{if } i_j \in I_-(\vec{L}, \vec{M}). \end{cases}$$

(Note that in our setting, the side length d_{i_j} of the hyperbox $\mathcal{H}_G^{\vec{L},\vec{M}}$ is the same as the number p_{i_j} of w basepoints on the respective link component.)

At each vertex $\varepsilon \in \mathbb{E}(\boldsymbol{d}'^M)$, we have the Floer complex

$$C^\varepsilon = \mathfrak{C}_{\varepsilon^<}(\mathcal{H}_{G,\varepsilon^>}^{\vec{L},\vec{M}-M_\varepsilon}, M_\varepsilon, \psi^{\vec{M}}(\boldsymbol{s})),$$

and along the faces we have linear maps $D_\varepsilon^{\varepsilon'-\varepsilon} \colon C^\varepsilon \to C^{\varepsilon'}$.

We compress the hyperbox $\mathfrak{C}(\mathcal{H}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$ as in Section 5.6, and define

$$\widehat{D}^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})} \colon \mathfrak{A}^-(G, \vec{M}, p^{\vec{M}}(\boldsymbol{s})) \cong \mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta, \vec{M}, \psi^{\vec{M}}(\boldsymbol{s})) \to \mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}, \psi^{\vec{M}}(\boldsymbol{s})) \tag{263}$$

to be the longest diagonal map in the resulting hypercube.

For example, when $M = L_i$ is a single component with orientation $\vec{M} = \vec{L}$, the map $\widehat{D}^{\vec{M}}_{p^{\vec{M}}(\boldsymbol{s})}$ is a composition of the triangle maps corresponding to handleslides over the basepoints in $\mathfrak{X}(\vec{M})$, in the given order. When M has several components, it is a sum of compositions of more complicated polygon maps, corresponding to chain homotopies (of higher order) between compositions of the handleslide maps.

Remark 15.3 The formulas used to define $D_\varepsilon^{\varepsilon'-\varepsilon}$ in Section 13.6 simplify somewhat in the setting of grids. Indeed, all the hypercubes of strongly Heegaard diagrams are of pure β type, meaning that the curve collection $\boldsymbol{\alpha}$ stays constant; that is, the bipartition maps take only the value β. Further, the chain elements $\Theta^{\beta}_{\zeta,\zeta'}$ consist of a single intersection point (in Alexander gradings zero) when $\|\zeta'-\zeta\|=1$, and are zero otherwise. Thus,

$$c(\boldsymbol{a}^{\beta}_{\gamma'\gamma}) = \begin{cases} 1 & \text{if } \boldsymbol{a} = \Theta^{\beta}_{\zeta,\zeta'}, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\bar{u}_{i,j}(\boldsymbol{a}^{\beta}_{\gamma'\gamma}) := u_{i,j}(\boldsymbol{a}^{\beta}_{\zeta\zeta'}) = 0.$$

From here, we find that, for $\boldsymbol{a} = \Theta^{\beta}_{\zeta,\zeta'}$ (which is the only intersection point with a nonzero contribution), the quantities $n^{\boldsymbol{a}}$ that appear in (243) are the same as the corresponding basepoint multiplicities n (with the same subscripts).

15.4 The handleslid complex

For $O_j \in \mathbb{O}_i$, we let $U_{j'}$ be the variable corresponding to the row exactly under the row through O_j. Given a sublink $M \subseteq L$, we define a chain complex

$$\mathcal{K}(M) := \bigotimes_{M_i \subseteq M} \bigotimes_{j \in \mathbb{O}_i} (\mathcal{R} \xrightarrow{U_j - U_{j'}} \mathcal{R}). \tag{264}$$

Given $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$ and an orientation $\vec{M}$ of M, we seek to describe the complex $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}, \psi^{\vec{M}}(\boldsymbol{s}))$ explicitly. Before stating the result, let us introduce some terminology. We call the Heegaard diagram

$$\mathcal{H}^{\vec{L},\vec{M}}_{G,\boldsymbol{d}^M} = (\mathcal{T}, \boldsymbol{\alpha}, \boldsymbol{\beta}^{\mathfrak{X}(\vec{M})}, \mathbb{O}^{\vec{M}}, \mathbb{X}^{\vec{M}})$$

the *handleslid diagram*. It contains two kinds of beta curves: those that approximate curves in the original collection $\boldsymbol{\beta}$, and those that encircle one of the points in $\mathfrak{X}(\vec{M})$. We call the former curves *old*, and the latter *recent*. Each recent curve bounds a disk in $\mathcal{T}$, which we call a *marked disk*. The curves in $\boldsymbol{\alpha}$ are also of two types: those that intersect a recent $\beta^{\mathfrak{X}(\vec{M})}$ curve, and those that do not. We call the first alpha curves *special*, and the latter *normal*. The connected components of $\mathcal{T} \setminus \boldsymbol{\alpha}$ are called rows, and a row is called *special* if at least one of its boundaries is special. If both of its alpha boundaries are special, the row is called *very special*.

Observe that each recent beta curve intersects a unique (special) alpha curve, and it does so in two points. Hence, each $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}$ must contain exactly one of these two points.

Note that there is a *quasidestabilized grid diagram* G^{L-M} obtained from G by eliminating all rows and columns on which M is supported; compare Section 6. The diagram G^{L-M} represents the link $L-M$, with the orientation induced from $\vec{L}$.

Proposition 15.4 *For any $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, for a suitable choice of almost complex structure on* $\mathrm{Sym}^n(\mathcal{T})$, *there is an isomorphism of chain complexes of $\mathcal{R}$-modules*

$$\Psi^{\vec{M}}_{\psi^{\vec{M}}(\boldsymbol{s})} : \mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}, \psi^{\vec{M}}(\boldsymbol{s})) \to \mathfrak{A}^-(G^{L-M}, \psi^{\vec{M}}(\boldsymbol{s}))[\![\{U_i\}_{L_i \subseteq M}]\!] \otimes_{\mathcal{R}} \mathcal{K}(M). \tag{265}$$

Here, the square brackets mean adjoining the variables U_i, in order to make the respective group into an $\mathcal{R}$-module.

Observe that, on the right-hand side of (265), we have the complex associated to the quasidestabilized grid diagram G^{L-M}, for which holomorphic disks in the symmetric product exactly correspond to empty rectangles on the grid; see [Manolescu et al. 2009a]. In proving Proposition 15.4, the challenge will be to show that the complex on the left-hand side of (265) has the same description. The proof will occupy the rest of this subsection.

The curves $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}^{\mathfrak{X}(\vec{M})}$ split the torus $\mathcal{T}$ into a number of regions. These can be rectangles, bigons, hexagons, and octagons. The bigons and hexagons appear near the markings in $\mathfrak{X}(\vec{M})$, while the octagons only appear when two markings in $\mathfrak{X}(M)$ lie in adjacent rows and adjacent columns of G. The hexagons and octagons are always supported in special rows.

Let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ be a homology class of Whitney disks in $\mathrm{Sym}^n(\mathcal{T})$ with boundaries on $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}$, for $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}$. The class ϕ has an associated domain; see Section 6.3. Recall that the Maslov index $\mu(\phi)$ can be calculated in terms of the domain $\mathcal{D} = \mathcal{D}(\phi)$ using Lipshitz's formula (62):

$$\mu(\phi) = \sum_{x \in \boldsymbol{x}} n_x(\mathcal{D}) + \sum_{y \in \boldsymbol{y}} n_y(\mathcal{D}) + e(\mathcal{D}). \tag{266}$$

Suppose $\mathcal{D}$ is written as a linear combination of regions

$$\mathcal{D} = \sum a_i R_i, \quad a_i \in \mathbb{Z}.$$

For any region R, we set $n(R, \boldsymbol{x}, \boldsymbol{y}) = n_{\boldsymbol{x}}(R) + n_{\boldsymbol{y}}(R)$ and $\mu(R, \boldsymbol{x}, \boldsymbol{y}) = n(R, \boldsymbol{x}, \boldsymbol{y}) + e(R)$. Since the Euler measure and vertex multiplicities are additive, equation (266) gives

$$\mu(\phi) = \sum a_i \mu(R_i, \boldsymbol{x}, \boldsymbol{y}). \tag{267}$$

The differential on $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathfrak{X}(\vec{M})}, \psi^{\vec{M}}(\boldsymbol{s}))$ involves counts of holomorphic disks in homology classes ϕ with $\mu(\phi) = 1$. The following are two necessary conditions for a domain $\mathcal{D} = \mathcal{D}(\phi)$ of index one to admit a holomorphic representative:

(i) All local multiplicities a_i of $\mathcal{D}$ must be nonnegative.

(ii) The support of the domain $\mathcal{D}$ (that is, the closure of the union of the regions R_i such that $a_i \neq 0$) must be connected.

The first condition is a consequence of the principle of positivity of intersection for holomorphic objects; see [Ozsváth and Szabó 2004b, Lemma 3.2]. The second condition needs to be satisfied because if the support were disconnected, the class ϕ would be the sum of two domains ϕ_1 and ϕ_2, with each ϕ_i required to admit pseudoholomorphic representatives generically. Hence $\mu(\phi_1), \mu(\phi_2) \geq 1$, which would contradict the fact that $\mu(\phi_1) + \mu(\phi_2) = \mu(\phi) = 1$.

A domain $\mathcal{D}$ that satisfies the two conditions above is called *positive* and *connected*.

For our handleslid Heegaard diagram, we can completely characterize the index one domains that are positive and connected. Indeed, given a class $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ that has such a domain, let us understand the quantities $\mu(R_i, \boldsymbol{x}, \boldsymbol{y})$ that appear with nonzero multiplicity $a_i > 0$ in (267):

- If R_i is a rectangle, then $e(R_i) = 0$ and $\mu(R_i, \boldsymbol{x}, \boldsymbol{y}) \in [0, 1]$, with the exact value depending on the number of vertices in $\boldsymbol{x}, \boldsymbol{y}$ among the corners of R_i.
- If R_i is a bigon, then $e(R_i) = \frac{1}{2}$ and $\mu(R_i, \boldsymbol{x}, \boldsymbol{y}) = 1$. Indeed, note that $\boldsymbol{x}$ contains exactly one of the two points of intersection between the respective new beta curve and a special alpha curve. Hence $\boldsymbol{x}$ contributes $\frac{1}{4}$ to the quantity $n(R_i, \boldsymbol{x}, \boldsymbol{y})$. The same goes for $\boldsymbol{y}$, while the Euler measure of the bigon is $\frac{1}{2}$.
- If R_i is a hexagon, then $e(R_i) = -\frac{1}{2}$ and $\mu(R_i, \boldsymbol{x}, \boldsymbol{y}) \in \left[0, \frac{1}{2}\right]$. Indeed, there is a contribution of $\frac{1}{2}$ to the vertex multiplicity coming from points on the recent beta curve, and there may be an additional contribution from intersections between old beta curves and normal alpha curves. We denote this additional contribution by $n'(R_i, \boldsymbol{x}, \boldsymbol{y})$. It is at most $\frac{1}{2}$.
- If R_i is an octagon, then $e(R_i) = -1$ and $\mu(R_i, \boldsymbol{x}, \boldsymbol{y}) = 0$, because the vertex multiplicity contribution is always one, coming from points on the two recent beta curves on the boundary of the octagon.

Hence,

$$\mu(\phi) = \sum_{\text{bigons}\, R_i} a_i + \sum_{\text{rectangles}\, R_i} a_i n(R_i, \boldsymbol{x}, \boldsymbol{y}) + \sum_{\text{hexagons}\, R_i} a_i n'(R_i, \boldsymbol{x}, \boldsymbol{y}) \geq 0. \tag{268}$$

If $\mu(\phi) = 1$, it follows that we can have at most one bigon in the support of $\mathscr{D}$, and that the bigon (if it exists) must appear with multiplicity one. Further, no points of $\boldsymbol{x} \cap \boldsymbol{y}$ can be contained in the interior of D.

We distinguish several cases for $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ positive, connected and of index one. These cases are shown in Figures 49–52.

(I) There exists a normal alpha curve α_j such that $\alpha_j \cap \boldsymbol{x} \neq \alpha_j \cap \boldsymbol{y}$. Then the two points $\alpha_j \cap \boldsymbol{x}$ and $\alpha_j \cap \boldsymbol{y}$ are on different (old) beta curves, so there must exist another normal alpha curve $\alpha_k \neq \alpha_j$ with $\alpha_k \cap \boldsymbol{x} \neq \alpha_k \cap \boldsymbol{y}$. We get a contribution of at least $\frac{1}{4}$ to $n(\mathscr{D})$ from each of the four intersection points on α_j and α_k. Since $\mu(\phi) = 1$, there can be no other contribution. It follows that the support of $\mathscr{D}$ has no bigons. One possibility is that $\mathscr{D}$ is a big rectangle with vertices the four intersection points $\alpha_j \cap \boldsymbol{x}$, $\alpha_j \cap \boldsymbol{y}$, $\alpha_k \cap \boldsymbol{x}$ and $\alpha_k \cap \boldsymbol{y}$, with all the marked disks in the interior of the rectangle removed. This is Case I(a), shown in Figure 49. Note that there is a one-to-one correspondence between such rectangles and empty rectangles in the quasidestabilized grid diagram G^{L-M}.

The other possibility is that in addition to a big rectangle, $\mathscr{D}$ also contains some additional rectangle, hexagon, or octagon regions. These have to be supported in some very special rows. Further, all nonbigon regions in those rows are contained in the domain $\mathscr{D}$, because the multiplicity of regions in a very special

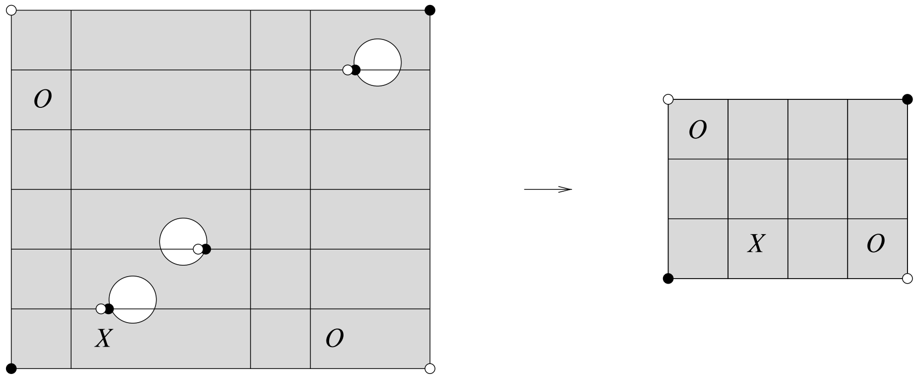

I(a)

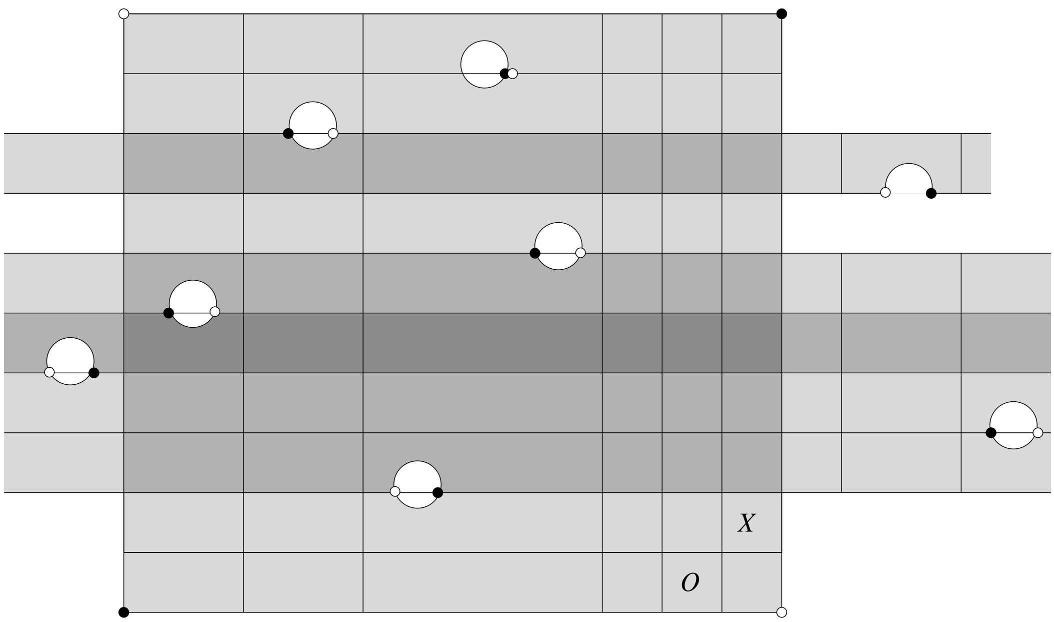

I(b)

Figure 49: Positive connected domains of index one: Case I. In each subcase, the domain is shaded, with darker shading corresponding to higher local multiplicity. In subcase (a) we only have multiplicities zero and one, while in subcase (b) we can arbitrarily large multiplicities; in the figure, they go up to multiplicity three. The black dots represent components of the initial generator $\boldsymbol{x}$, and the white dots components of $\boldsymbol{y}$. There can be various markings (X and O) in the domains, though not in the rows containing markings in $\mathbb{X}(\vec{M})$. In subcase (a), apart from the domain in the handleslid diagram $(\mathcal{T}, \boldsymbol{\alpha}, \boldsymbol{\beta}^{\mathbb{X}(M)})$ on the left-hand side, we also show the corresponding domain in the quasidestabilized grid diagram G^{L-M}, on the right-hand side.

row can only change when we pass the vertical curves through $\alpha_j \cap \boldsymbol{x}$ and $\alpha_j \cap \boldsymbol{y}$. Thus, the domain consists of a big rectangle as in Case I(a), plus one or more very special rows intersecting the big rectangle, minus the bigons in those rows. The multiplicities in the special rows cannot change by more than one as

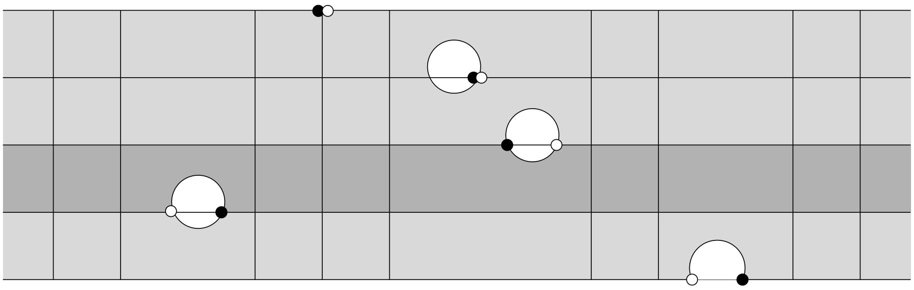

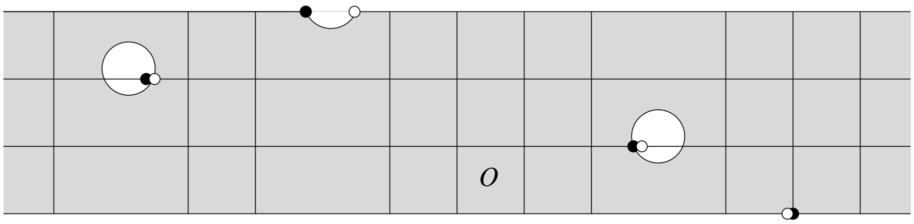

Figure 50: Positive connected domains of index one: Case II. We use the same conventions as in Case I. In subcase (a) we show a domain of width four and having local multiplicity two in the darkly shaded regions. In (b) we show a domain of width three with only zero and one local multiplicities. Of course, in each subcase the widths and multiplicities can be arbitrarily large, just as we saw in Case I(b).

we pass a special alpha curve, because the support of $\mathscr{D}$ contains no bigons. A domain of this type is said to be in Case I(b). An example is shown in Figure 49.

(II) We have $\alpha_j \cap \boldsymbol{x} = \alpha_j \cap \boldsymbol{y}$ for all normal alpha curves α_j, and the support of $\mathscr{D}$ contains no bigons. Then $\mathscr{D}$ is supported in some adjacent special rows. Further, there are no old beta curves such that when we pass them the multiplicity of $\mathscr{D}$ gets changed. Hence, if the support of $\mathscr{D}$ contains a nonbigon region in a special row, it must contain all the nonbigon regions in that row, with the same multiplicity. Further, one of the boundaries of the support must be a normal alpha curve, which produces all nonzero contributions to $\mu(\phi)$ in (268). There are two subcases, II(a) and II(b), according to whether the normal alpha curve is the top or the bottom boundary of the support. In both cases, the multiplicity in the row bounded by the normal alpha curve must be one. The multiplicities in the other rows can be higher, but they are constant in each row and cannot change by more than one as we pass a special alpha curve. See Figure 50.

(III) We have $\alpha_j \cap \boldsymbol{x} = \alpha_j \cap \boldsymbol{y}$ for all normal alpha curves α_j, and the support of $\mathscr{D}$ includes a bigon which contains a marking in $\mathscr{X}(\vec{M})$. Then the domain $\mathscr{D}$ is supported in some adjacent special rows, and the bigon gives the only nonzero contributions to $\mu(\phi)$ in (268). Further, if $\mathscr{D}$ contains a rectangle, hexagon or octagon in a special row, it must contains all the rectangles, hexagons, and octagons in that row, with the same multiplicity. There are three subcases III(a)–(c), according to whether

(a) $\mathscr{D}$ is the bigon itself;

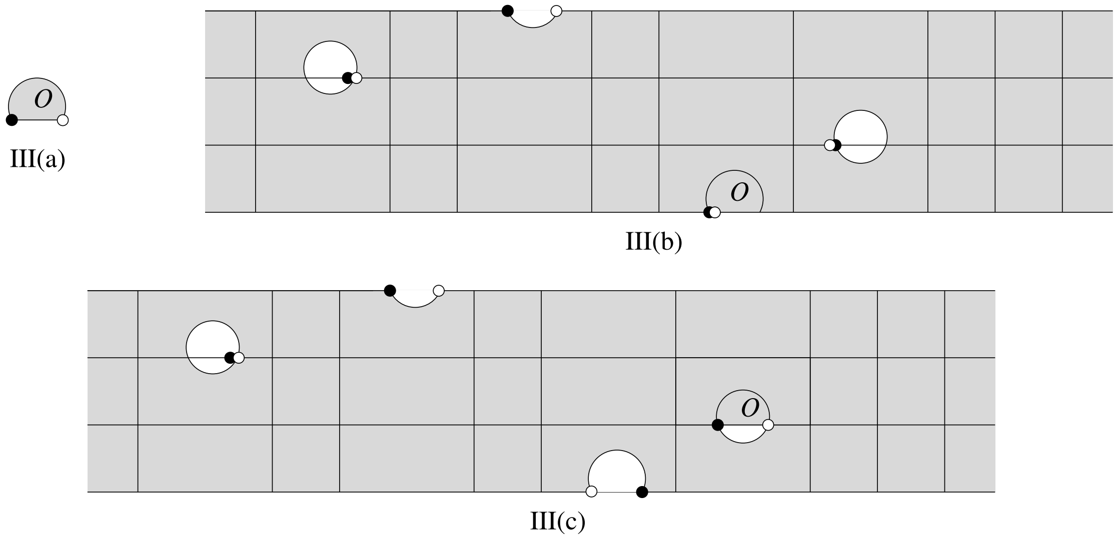

Figure 51: Positive connected domains of index one: Case III. Same conventions as before. In subcases (b) and (c) the widths and multiplicities can be arbitrarily large, just as in Case II.

(b) $\mathscr{D}$ contains some special rows, but only above the bigon; or

(c) $\mathscr{D}$ contains special rows both below and above the bigon.

See Figure 51.

(IV) We have $\alpha_j \cap \boldsymbol{x} = \alpha_j \cap \boldsymbol{y}$ for all normal alpha curves α_j, and the support of $\mathscr{D}$ includes a bigon which does not contain a marking in $\mathscr{X}(\vec{M})$. This is similar to Case III, in that we have the same observations about the support consisting of some adjacent special rows, and three subcases:

(a) $\mathscr{D}$ is the bigon itself;

(b) $\mathscr{D}$ contains some special rows, but only below the bigon; or

(c) $\mathscr{D}$ contains special rows both below and above the bigon.

See Figure 52.

When a domain $\mathscr{D}$ is in one of the cases above, we say that it is of the corresponding *type*. For example, a domain in Case II(b) is called of type II(b).

Definition 15.5 If $\mathscr{D}$ is a domain on the handleslid grid, the number of rows that have nontrivial intersection with its support is called the *width* of $\mathscr{D}$.

We would like to count the number of holomorphic representatives for each class ϕ with a positive, connected domain $\mathscr{D}$ of index one. One should note that many of the domains in Cases (I)–(IV) above are decomposable, meaning that we can write them as a sum $\mathscr{D}_1 + \mathscr{D}_0$, with $\mathscr{D}_i$ of index $i \in \{0, 1\}$ being

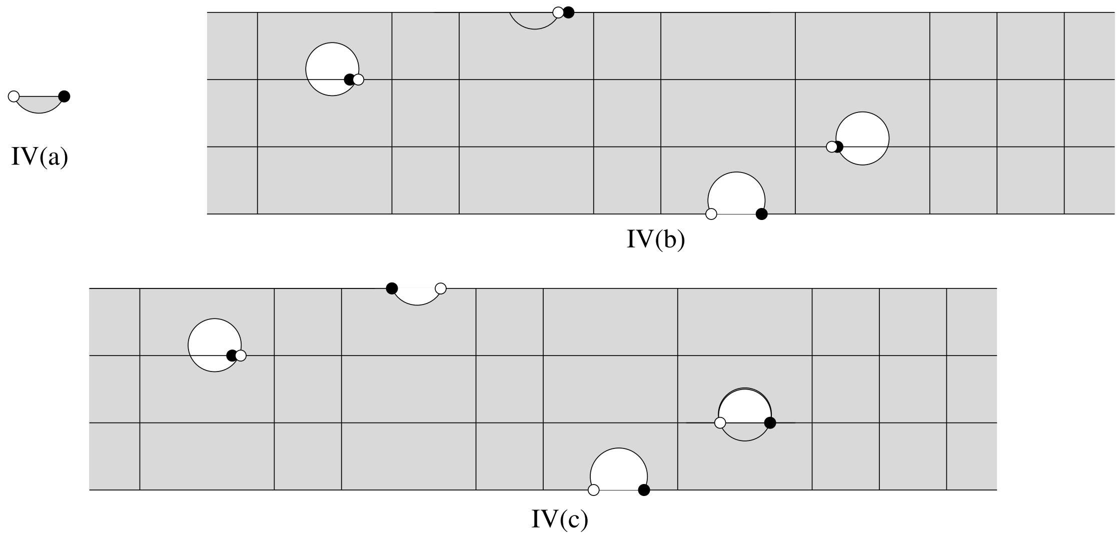

Figure 52: Positive connected domains of index one: Case IV. Same conventions as before. In subcases (b) and (c) the widths and multiplicities can be arbitrarily large, just as in Cases II and III.

positive and connected. This means that the number of holomorphic representatives for $\mathscr{D}$ may depend on the choice of almost complex structure on the symmetric product. We will explain how to make a suitable choice soon.

When counting holomorphic representatives, it is convenient to use Lipshitz's cylindrical formulation of Heegaard Floer homology [Lipshitz 2006], as in Section 6.2. In his setting, we need to choose an almost complex structure on $W = \mathscr{T} \times [0,1] \times \mathbb{R}$, rather than on the symmetric product.

Suppose

$$\mathscr{Z}(\vec{M}) = \{Z_1, \dots, Z_k\}.$$

The labeling of the markings is as follows. Choose a normal alpha curve α_1 on the handleslid diagram. Go vertically down around the torus, starting at α_1 and ending at α_1, and number the markings in $\mathscr{Z}(\vec{M})$ as $Z_1, \dots, Z_k$, in the order in which they are encountered.

For each $Z_i \in \mathscr{Z}(M)$, there is a recent beta curve $\beta_{j_i}^{\mathscr{Z}(\vec{M})}$, which is the boundary of the marked disk containing Z_i. Choose a slightly bigger curve around this beta curve, and introduce a long cylindrical neck of length T_i there, for $T_i \gg 0$. We can then choose an almost complex structure $J(T_1, \dots, T_k)$ on W which only depends on T_i on the neck around Z_i, and is split on these necks. This is the analogue of the "stretching the neck" process from Section 6.4, only now we consider flow lines rather than triangles. In fact, we can view the handleslid diagram as a special connected sum of the quasidestabilized grid diagram G^{L-M} and k genus zero diagrams $\mathscr{S}_1, \dots, \mathscr{S}_k$. Each $\mathscr{S}_i$ is obtained from the diagram in Figure 11 by deleting the gamma curve. See Figure 53.

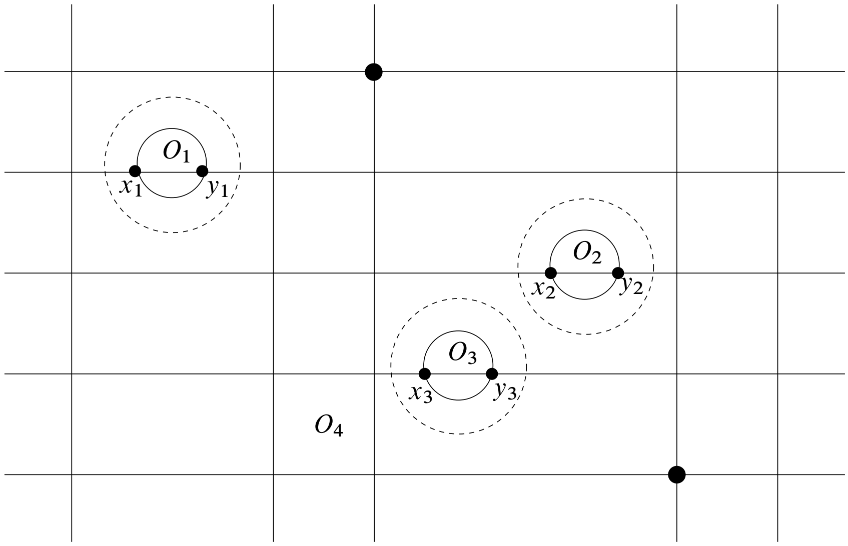

Figure 53: Degenerations. We choose the complex structure by stretching the neck along the dashed curves. The bigger black dots are components of the generator $\bar{\boldsymbol{x}}$ in Lemma 15.7.

Given a marking Z_i, we denote by M_i the marked disk containing it, and by x_i and y_i the two points of intersection of $\beta_{j_i}^{\mathbb{X}(M)}$ with the corresponding alpha curve. The point x_i is on the left of y_i inside the marked disk.

Lemma 15.6 *Let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ be a homology class of Whitney disks in the handleslid diagram. Suppose $\mu(\phi) = 1$ and the domain $\mathcal{D} = \mathcal{D}(\phi)$ is positive and connected. Suppose α_j is a special alpha curve that intersects a recent beta curve β_j. Let M_i be the marked disk whose boundary is β_j, and let Z_i be the marking inside. Suppose $\alpha_j \cap \boldsymbol{x} = x_i$, $\alpha_j \cap \boldsymbol{y} = y_i$, and that $\mathcal{D}$ has multiplicity zero everywhere inside the marked disk M_i. Denote by r the row whose upper boundary is α_j, and denote by b the bigon $r \cap M_i$. If the row r contains a marking $Z_{i+1} \in \mathbb{X}(M)$, fix the corresponding neck-length T_{i+1}. If the class ϕ admits holomorphic representatives for a sequence of almost complex structures $J(T_1^{(n)}, \dots, T_k^{(n)})$ such that $T_i^{(n)} \to \infty$ (where $T_i^{(n)}$ corresponds to Z_i) and $T_{i+1}^{(n)} = T_{i+1}$ is fixed, then the support of the domain $\mathcal{D}$ is exactly the union of all the regions in the row r, except the bigon b.*

Proof For simplicity, let us first assume that all neck-lengths $T_s^{(n)} = T_s$ are fixed for $s \neq i$, and $T_i^{(n)} \to \infty$.

Let m_1 and m_2 the multiplicities of $\mathcal{D}$ on each side of α_j, near the connected sum curve. By an analogue of Proposition 6.14, in the limit $T_i \to \infty$ the presumed holomorphic representatives of ϕ degenerate into a broken flow line v on $\mathcal{S}_i$, and the union of a broken flow line v' on the other side (ie on the diagram G' obtained by deleting β_j) with an annoying α-degeneration with domain $\mathcal{P}$. Let ψ and ψ' be the homology classes of v and v', respectively. Lemmas 6.10 and 6.16 can be applied to homology classes of flow lines just as well as to triangles. Therefore,

$$\mu(\phi) = \mu(\psi) + \mu(\psi') + \mu(\mathcal{P}) - m_1 - m_2 \quad \text{and} \quad \mu(\psi) = m_1 + m_2.$$

In principle, the latter equality should also involve the multiplicities of $\mathcal{D}$ inside the two bigons that form M_i, but those multiplicities are zero by hypothesis. We get

$$1 = \mu(\phi) = \mu(\psi') + \mu(\mathcal{P}).$$

Further, we must have $m_2 = m_1 + 1$ because of the acceptability condition (61). This means that $\partial\mathcal{P}$ contains α_j with multiplicity -1. In particular, $\mathcal{P} \neq 0$, so $\mu(\mathcal{P}) \geq 1$; see Lemma 6.13 for $\mu(\mathcal{P})$. We deduce that $\mu(\psi') = 0$. Since ψ' admits holomorphic representatives (for an almost complex structure which can be chosen to be generic away from the neck), it must be the case that ψ' is trivial. As for the remaining domain $\mathcal{P}$, it is positive, of index one, and satisfies $m_2 = m_1 + 1$. The only possibility is that it is the row r, viewed in the quasidestabilized diagram G'; compare (71). The claim follows, provided that the neck-lengths T_i stay fixed for $s \neq i$.

Now allow the other neck-lengths to vary as well, except T_{i+1}. In the limit some of them go to infinity, so the holomorphic representatives of ϕ degenerate into broken flow lines on a number of surfaces $\mathcal{S}_s$, plus the union of a broken flow line and α-boundary degenerations on a diagram G' obtained by deleting several recent beta curves. Nevertheless, most of the arguments above still apply, with some notational modifications. The broken flow line on G' must be trivial, so all we are left with is some α-boundary degenerations on G'. The main difference is that now, instead of (71), we have

$$\mu(\mathcal{P}) = \sum_{\alpha} m(\mathcal{P}, \alpha),$$

where the sum is over all curves α that did not intersect any of the deleted beta curves, and the quantity $m(\mathcal{P}, \alpha)$ is the sum of the multiplicities of $\mathcal{P}$ in the two rows that have α as part of their boundary.

It follows that the inequality $\mu(\mathcal{P}) \geq 0$ in Lemma 6.13 still holds true, but equality can happen for some nonzero domains on G', namely those composed of very special rows such that the normal beta curves intersecting their boundaries have been deleted. As such, a priori there are more possibilities for the positive domain $\mathcal{P}$ of index one that represents the α-boundary degenerations. However, the fact that the recent beta curve just below β_j (if it exists) has not been deleted, together with the fact that $\mathcal{P}$ is connected and satisfies $m_2 = m_1 + 1$, suffice to determine $\mathcal{P}$ uniquely as being the row below α_j. □

Lemma 15.6 suggests the following way of choosing the neck-lengths $T_1, \dots, T_k$ in the construction of the almost complex structure $J(T_1, \dots, T_k)$. We first choose $T_k \gg 0$ such that the conclusion of the Lemma holds true for all possible classes ϕ satisfying the hypothesis, with $i = k$. This is possible because there are a finite number of such classes, and the row r below Z_k does not contain any marking in $\mathcal{X}(\vec{M})$. Next, choose $T_{k-1} \gg T_k$ such that the conclusion of the Lemma holds for all classes with $i = k - 1$. Iterate this procedure until we get to T_1, so that

$$T_1 \gg T_2 \gg \cdots \gg T_k \gg 0.$$

Lemma 15.6 implies that, for this choice of almost complex structure, the domains of type I(b), III(c) and IV(c) have holomorphic disk counts equal to zero (mod 2), so they do not contribute to the differential. The same goes for domains of type II(b) or III(b) that have width larger than one.

On the other hand, the holomorphic disk counts are one (mod 2) for

- simple bigons, that is, domains of types III(a) and IV(a);
- domains of type I(a); see [Ozsváth and Szabó 2009, Lemma 3.11];
- annular domains of type II(b) and width one; see [Ozsváth and Szabó 2004c, proof of Lemma 3.4];
- annular domains of type III(b) and width one; see [Ozsváth and Szabó 2004b, Lemma 9.4] — indeed, for these the conformal angle of the beta part of the upper boundary is very small, because of our choice of neck-lengths, which means that there exists a choice of cut length on the lower boundary that makes the annulus holomorphic.

This leaves us with the domains of types II(a) and IV(b). It is more difficult to establish the holomorphic disk counts for these. Instead, we will settle for computing their net effect on the differential ∂ on the handleslid complex.

More precisely, let us split the markings in $\mathscr{Z}(\vec{M})$ into equivalence classes, with the equivalence relation being generated by the relations $Z \sim Z'$ if Z and Z' lie in adjacent rows. Then an equivalence class consists of all the markings in $\mathscr{Z}(\vec{M})$ that lie between two normal alpha curves.

Let us consider such an equivalence class. Without loss of generality, say it is composed of the markings $Z_1 = O_1, \dots, Z_p = O_p$, for some $p \le k$. (See Figure 53 for the case $p = 3$.) Let $\overline{G}$ be the grid diagram obtained by deleting the recent beta curves $\beta_1, \dots, \beta_p$ that encircle $O_1, \dots, O_p$, as well as the special alpha curves $\alpha_1, \dots, \alpha_p$ that intersect $\beta_1, \dots, \beta_p$. Given a generator $\boldsymbol{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta^{\mathscr{Z}(\vec{M})}$, let $\bar{\boldsymbol{x}}$ be the corresponding generator in $\overline{G}$ obtained from $\boldsymbol{x}$ by deleting its components on α_i, for $i = 1, \dots, p$. Define $I(\boldsymbol{x})$ to be the Maslov index of $\bar{\boldsymbol{x}}$ in $\overline{G}$. Then I defines a filtration on the handleslid complex. Indeed, if $\boldsymbol{y}$ contributes a nonzero term to the differential $\partial\boldsymbol{x}$, we must have $I(\boldsymbol{x}) \ge I(\boldsymbol{y})$, with equality if and only if $\bar{\boldsymbol{x}} = \bar{\boldsymbol{y}}$.

Thus, the associated graded of the handleslid complex $\mathfrak{A}^-(\mathbb{T}_\alpha, \mathbb{T}_\beta^{\mathscr{Z}(\vec{M})}, \psi^{\vec{M}}(\boldsymbol{s}))$ splits as a direct sum of terms $C(\bar{\boldsymbol{x}})$, where $C(\bar{\boldsymbol{x}})$ is generated by all possible $\boldsymbol{x}$ with the fixed reduction $\bar{\boldsymbol{x}}$ on $\overline{G}$. Pick such a direct summand $C(\bar{\boldsymbol{x}})$.

Lemma 15.7 *For the almost complex structure $J(T_1, \dots, T_k)$ chosen as before (with $T_1 \gg T_2 \gg \cdots \gg T_k \gg 0$), the complex $C(\bar{\boldsymbol{x}})$ is isomorphic to the tensor product of the complexes*

$$\mathscr{R} \xrightarrow{U_i - U_{i+1}} \mathscr{R},$$

for $i = 1, \dots, p$.

Proof Recall that the intersection points between α_i and β_i are denoted by x_i and y_i; compare Figure 53. The generators of $C(\bar{\boldsymbol{x}})$ can be written as $\boldsymbol{x}^A$, where $A \subseteq \{1, \dots, p\}$ is the set of indices i such that $x_i \in \boldsymbol{x}^A$.

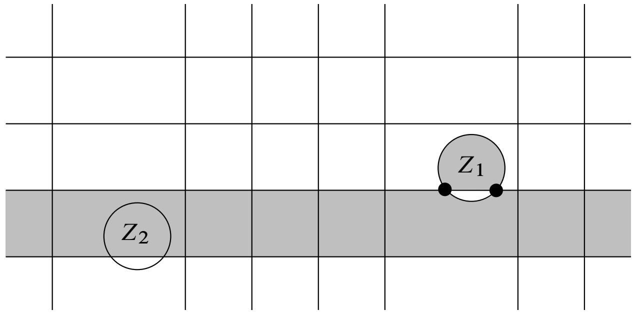

Figure 54: The origin of the complexes $\mathcal{H}(M)$. The figure shows part of a grid diagram with some arcs on the α and $\beta^{\{Z_1,Z_2,\dots,Z_k\}}$ curves drawn. There are two intersection points (marked as bullets) between the alpha curve below the marking Z_1, and the corresponding beta curve. There are two differentials going from the left to the right generator: a bigon containing Z_1 and an annulus containing Z_2, both drawn shaded in the diagram. This produces a factor of the form $\mathcal{R} \xrightarrow{U_j - U_{j'}} \mathcal{R}$ in the definition of the complexes $\mathcal{H}(M)$.

The contributions to the differential on $C(\bar{x})$ come from domains completely supported in the rows containing $O_1, \dots, O_{p+1}$. By our previous observations, if their contribution is nonzero, the domains must be of types II(b) of width one, III(b) of width one, III(a), IV(a), II(a), or IV(b). We already know that the first four cases in this series produce nonzero contributions mod 2. More precisely, the first three cases give a term $(U_i - U_{i+1})\boldsymbol{x}^{A\setminus\{i\}}$ in $\partial \boldsymbol{x}^A$, whenever $i \in A$.

The challenge is to figure out the contributions from domains of type II(a) and IV(b). Lemma 15.6 implies that nonzero contributions can come only from domains that have no local multiplicities bigger than one. Hence, the differential on the complex $C(\bar{x})$ takes the form

$$\partial \boldsymbol{x}^A = \sum_{i \in A} (U_i - U_{i+1}) \cdot \boldsymbol{x}^{A\setminus\{i\}} + \sum_{i \notin A} n_i^A \cdot \boldsymbol{x}^{A\cup\{i\}}$$

for some values $n_i^A \in \{0, 1\}$. These values are constrained by the requirement that $\partial^2 = 0$. In fact, it suffices to look at the coefficient of $\boldsymbol{x}^A$ in $\partial^2(\boldsymbol{x}^A)$:

$$\sum_{i \in A} n_i^{A\setminus\{i\}} \cdot (U_i - U_{i+1}) + \sum_{i \notin A} n_i^A \cdot (U_i - U_{i+1}) = 0.$$

Since the values $U_i - U_{i+1}$ for $i = 1, \dots, p$, are linearly independent, we deduce that $n_i^A = 0$ for all possible A and $i \notin A$. This completely determines the complex $C(\bar{x})$. $\square$

Proof of Proposition 15.4 For our choice of almost complex structure, all the domains that contribute to the differential on the handleslid complex are either of type I(a), or appear in the differential on a summand $C(\bar{x})$ of an associated graded.

The holomorphic count for all domains of type I(a) is always ± 1. As previously noted, such domains are in one-to-one correspondence with empty rectangles on the quasidestabilized grid G^{L-M}. The result now follows by combining this observation with Lemma 15.7. See Figure 54. $\square$

15.5 The surgery theorem

By composing the maps (262), (263) and (265), we construct a map

$$\Phi_{s}^{\vec{M}} = \Psi^{\vec{M}}_{\psi^{\vec{M}}(s)} \circ \hat{D}^{\vec{M}}_{p^{\vec{M}}(s)} \circ \mathcal{I}^{\vec{M}}_{s} : \mathfrak{A}^-(G,s) \to \mathfrak{A}^-(G^{L-M}, \psi^{\vec{M}}(s))[[\{U_{i,j}\}_{L_i \subseteq M}]] \otimes_{\mathcal{R}} \mathcal{H}(M), \tag{269}$$

defined for any $s \in \mathbb{H}(L)$.

This is the analogue of the map $\Phi_s^{\vec{M}}$ from (100). There is also an analogue of Proposition 9.4:

Proposition 15.8 *For any $\vec{M}$ and $s \in \overline{\mathbb{H}}(L)$,*

$$\sum_{\vec{M}_1 \amalg \vec{M}_2 = \vec{M}} \Phi^{\vec{M}_2}_{\psi^{\vec{M}_1}(s)} \circ \Phi^{\vec{M}_1}_{s} = 0, \tag{270}$$

where $\vec{M}_1$ and $\vec{M}_2$ are only considered with the orientations induced from $\vec{M}$.

Proof This is similar to the proof of Proposition 9.4, except we have the additional factors $\Psi^{\vec{M}}_{\psi^{\vec{M}}(s)}$ in (269). These factors commute with the descent maps $\hat{D}$ corresponding to disjoint sublinks. Indeed, if $\vec{M}_1 \cap \vec{M}_2 = \varnothing$, then Proposition 6.18 implies that

$$\Psi^{\vec{M}_1}_{s'} \circ \hat{D}^{\vec{M}_2}_{s'} = \hat{D}^{\vec{M}_2}_{s'} \circ \Psi^{\vec{M}_1}_{s'},$$

where $s' = p^{\vec{M}_2}(\psi^{\vec{M}_1}(s))$.

The maps $\Psi^{\vec{M}_1}$ also commute with the inclusion maps $\mathcal{I}^{\vec{M}_2}$ corresponding to disjoint sublinks. With these observations in mind, the proof of Proposition 9.4 can be easily adapted to the present context. □

Using the new maps $\Phi_s^{\vec{M}}$, we can exactly mimic Section 9.2, and state a version of the surgery theorem for grid diagrams instead of complete systems of hyperboxes.

We transfer most of the notation from Section 9.2. In particular, we have a framing Λ for the link $\vec{L}$. For a sublink $N \subseteq L$, recall that $\Omega(N)$ is the set of all possible orientations on N. For $\vec{N} \in \Omega(N)$, the quantity $\Lambda_{\vec{L},\vec{N}} \in H_1(S^3 - L)$ is the sum of the framings of those components of N oriented differently in $\vec{N}$ as in $\vec{L}$.

We consider the $\mathcal{R}$-module

$$\mathcal{C}^-(G,\Lambda) = \bigoplus_{M \subseteq L} \prod_{s \in \mathbb{H}(L)} \bigl(\mathfrak{A}^-(G^{L-M}, \psi^M(s))[[\{U_{i,j}\}_{L_i \subseteq M}]]\bigr) \otimes_{\mathcal{R}} \mathcal{H}(M),$$

where ψ^M simply means $\psi^{\vec{M}}$ with $\vec{M}$ being the orientation induced from the one on $\vec{L}$.

We equip $\mathcal{C}^-(G,\Lambda)$ with a boundary operator $\mathcal{D}^-$ as follows.

For $\boldsymbol{s} \in \mathbb{H}(L)$ and $\boldsymbol{x} \in \big(\mathfrak{A}^-(G^{L-M}, \psi^M(\boldsymbol{s}))[\![\{U_{i,j}\}_{L_i \subseteq M}]\!]\big) \otimes_{\mathcal{R}} \mathcal{K}(M)$, we set

$$\begin{aligned}\mathcal{D}^-(\boldsymbol{s},\boldsymbol{x}) &= \sum_{N \subseteq L-M} \sum_{\vec{N} \in \Omega(N)} (\boldsymbol{s} + \Lambda_{\vec{L},\vec{N}}, \Phi_{\boldsymbol{s}}^{\vec{N}}(\boldsymbol{x})) \\ &\in \bigoplus_{N \subseteq L-M} \bigoplus_{\vec{N} \in \Omega(N)} \big(\mathfrak{A}^-(G^{L-M-N}, \psi^{M \cup \vec{N}}(\boldsymbol{s}))[\![\{U_{i,j}\}_{L_i \subseteq M \cup N}]\!]\big) \otimes_{\mathcal{R}} \mathcal{K}(M \cup N) \subset \mathcal{C}^-(G, \Lambda).\end{aligned}$$

According to (270), $\mathcal{C}^-(G,\Lambda)$ is a chain complex. Just like the complex $\mathcal{C}^-(\mathcal{H},\Lambda)$ from Section 9.2, it splits into a direct product of complexes $\mathcal{C}^-(G,\Lambda,\mathfrak{u})$, according to $\mathfrak{u} \in \mathrm{Spin}^c(S^3_\Lambda(L)) \cong \mathbb{H}(L)/H(L,\Lambda)$. The complexes $\mathcal{C}^-(G,\Lambda,\mathfrak{u})$ admit relative $\mathbb{Z}/\mathfrak{d}(\mathfrak{u})\mathbb{Z}$-gradings, constructed just as for $\mathcal{C}^-(\mathcal{H},\Lambda,\mathfrak{u})$.

Theorem 15.9 *Fix a grid diagram G (with at least one free marking) for an oriented link $\vec{L}$ in S^3, and fix a framing Λ of L. Then, for every $\mathfrak{u} \in \mathrm{Spin}^c(S^3_\Lambda(L))$, we have an isomorphism of relatively graded $\mathbb{F}[\![U]\!]$-modules*

$$H_*(\mathcal{C}^-(G,\Lambda,\mathfrak{u}), \mathcal{D}^-) \cong \boldsymbol{HF}^-_*(S^3_\Lambda(L),\mathfrak{u}). \tag{271}$$

15.6 A complete system associated to the grid

Theorem 15.9 is a consequence of Theorem 1.1, albeit not an immediate one. A first guess would be to construct a complete system of hyperboxes $\mathcal{H}_G$ out of a grid diagram G (with at least one free marking) as follows: as noted in Section 15.3, the grid G gives rise to hyperboxes $\mathcal{H}_G^{\vec{L},\vec{M}}$ for the pairs $(\vec{L},\vec{M})$. One is tempted to try to include these into a complete system. Unfortunately, this is not possible: in a complete system, the diagrams $\mathcal{H}^{\vec{L},\vec{M}}(M)$ and $\mathcal{H}^{\vec{L},-\vec{M}}(M)$ are required to be surface isotopic, while for a grid G, typically the diagrams $\mathcal{H}^{\vec{L},\vec{M}}_{G,\boldsymbol{d}^M}$ and $\mathcal{H}^{\vec{L},-\vec{M}}_{G,\boldsymbol{d}^M}$ are not surface isotopic. For example, if $\vec{M}$ has the orientation induced from $\vec{L}$, then in the diagram $\mathcal{H}_G^{\vec{L},\vec{M}}$ the new beta curves are obtained by handlesliding over the O markings on M, whereas in $\mathcal{H}_G^{\vec{L},-\vec{M}}$ they are obtained by handlesliding over the X markings on M.

Nevertheless, we can still associate to a grid G a complete system of hyperboxes $\overline{\mathcal{H}}_G$ such that the (compressed) maps on Floer complexes coming from $\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}$ are essentially the same as those coming from $\mathcal{H}_G^{\vec{L},\vec{M}}$.

In fact, if $\vec{M}$ is the orientation induced from $\vec{L}$, we define $\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}$ to be exactly the same as $\mathcal{H}_G^{\vec{L},\vec{M}}$. On the other hand, if $\vec{M}$ has some components with a different orientation, we need to change $\mathcal{H}_G^{\vec{L},\vec{M}}$ using some more beta handleslides.

Before constructing $\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}$ in general, let us make an observation. Let O_j, X_j be two markings on the same row in the grid G. Then the two collections of beta curves $\boldsymbol{\beta}^{O_j}$ and $\boldsymbol{\beta}^{X_j}$ are strongly equivalent, being obtained from each other through handleslides of the beta curves between O_j and X_j, as in Figure 55.

For simplicity, let us first consider the case when $\vec{M} = -\vec{L}_i$ is a single component of L, oriented oppositely from its orientation in $\vec{L}$. Then $\mathcal{H}_G^{\vec{L},-L_i}$ is a one-dimensional hyperbox, consisting of a sequence of moves from the reduction $r_{-L_i}(G)$ to the diagram $\mathcal{H}^{\vec{L},-\vec{L}_i}_{G,\boldsymbol{d}^{L_i}}$ obtained from it by handleslides over all $X_j \in \mathbb{X}_i$. In

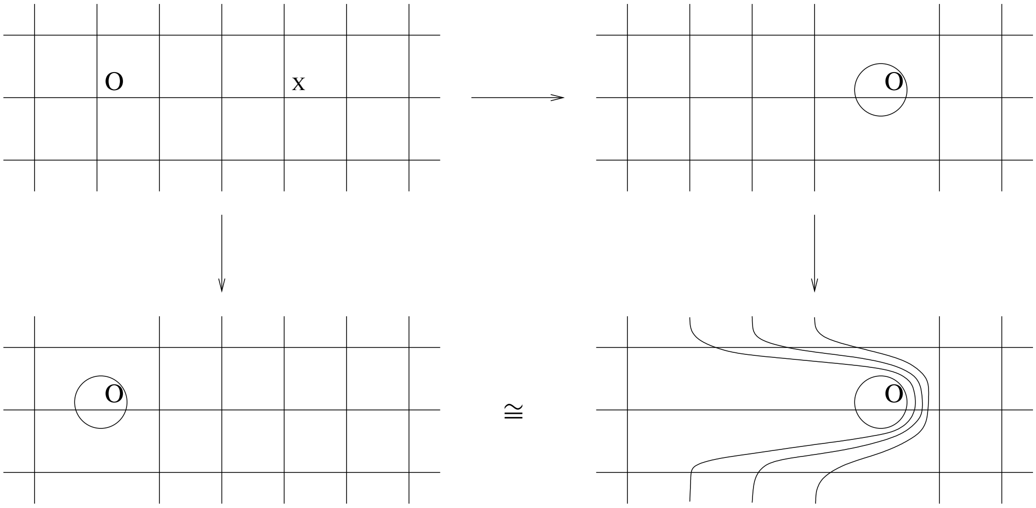

Figure 55: Relating the handleslides over O and over X. The top left picture shows a part of a grid diagram, with an O and X markings in the same row. The vertical arrow on the left describes a handleslide over the O. The top arrow describes a handleslide over the X, as well as relabeling the X as an O. The vertical arrow on the right shows a sequence of handleslides, the end result being a diagram surface isotopic with the one on the bottom left.

the new one-dimensional hyperbox $\overline{\mathcal{H}}_G^{\vec{L},-\vec{L}_i}$, we add at the end of this sequence a series of handleslides as in Figure 55, so that we end at the diagram $\mathcal{H}^{\vec{L},\vec{L}_i}_{G,\boldsymbol{d}^{L_i}}$, which is $r_{\vec{L}_i}(G)$ handleslid over all $O_{i,j} \in \mathbb{O}_i$. See Figure 56 for an example. There are unique choices for the intersection points that play the role of Θ-chain elements.

More generally, for $L_i \subset L'$, we define the one-dimensional hyperboxes $\overline{\mathcal{H}}_G^{\vec{L}',\vec{L}_i}$ and $\overline{\mathcal{H}}_G^{\vec{L}',\vec{L}_i}$ to consist of similar moves as in the case $L' = L$, except starting at the diagram handleslid over all the O markings that are on the components of $L - L'$.

Next, to define the higher-dimensional hyperboxes $\overline{\mathcal{H}}_G^{\vec{L}',\vec{M}}$ for all L' and $\vec{M}$, note that we already know the sequence of moves on their edges. The moves in one direction of the hyperbox are independent of those in a different direction, so it is straightforward to fill in the hyperbox by combining these moves. The Θ-chain elements are taken to be zero on higher-dimensional faces, just as we did for the hyperboxes $\mathcal{H}_G^{\vec{L},\vec{M}}$ from Section 15.3. The result is a complete system of hyperboxes $\overline{\mathcal{H}}_G$ for $\vec{L}$. Indeed, for the associated good set of trajectories, we can take horizontal paths on the grid joining each O marking to the X marking in the same row.

Proof of Theorem 15.9 In light of Theorem 1.1, it suffices to show that the chain complexes $\mathcal{C}^-(G,\Lambda,\mathfrak{u})$ and $\mathcal{C}^-(\overline{\mathcal{H}}_G,\Lambda,\mathfrak{u})$ are chain homotopy equivalent. Indeed, for $\boldsymbol{s} \in \overline{\mathbb{H}}(L)$, let us look at the hyperbox of Floer complexes $\mathfrak{C}(\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$ associated to a hyperbox $\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}$ in $\overline{\mathcal{H}}_G$, as in Section 13.6. This contains as a subhyperbox $\mathfrak{C}(\mathcal{H}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(\boldsymbol{s}))$.

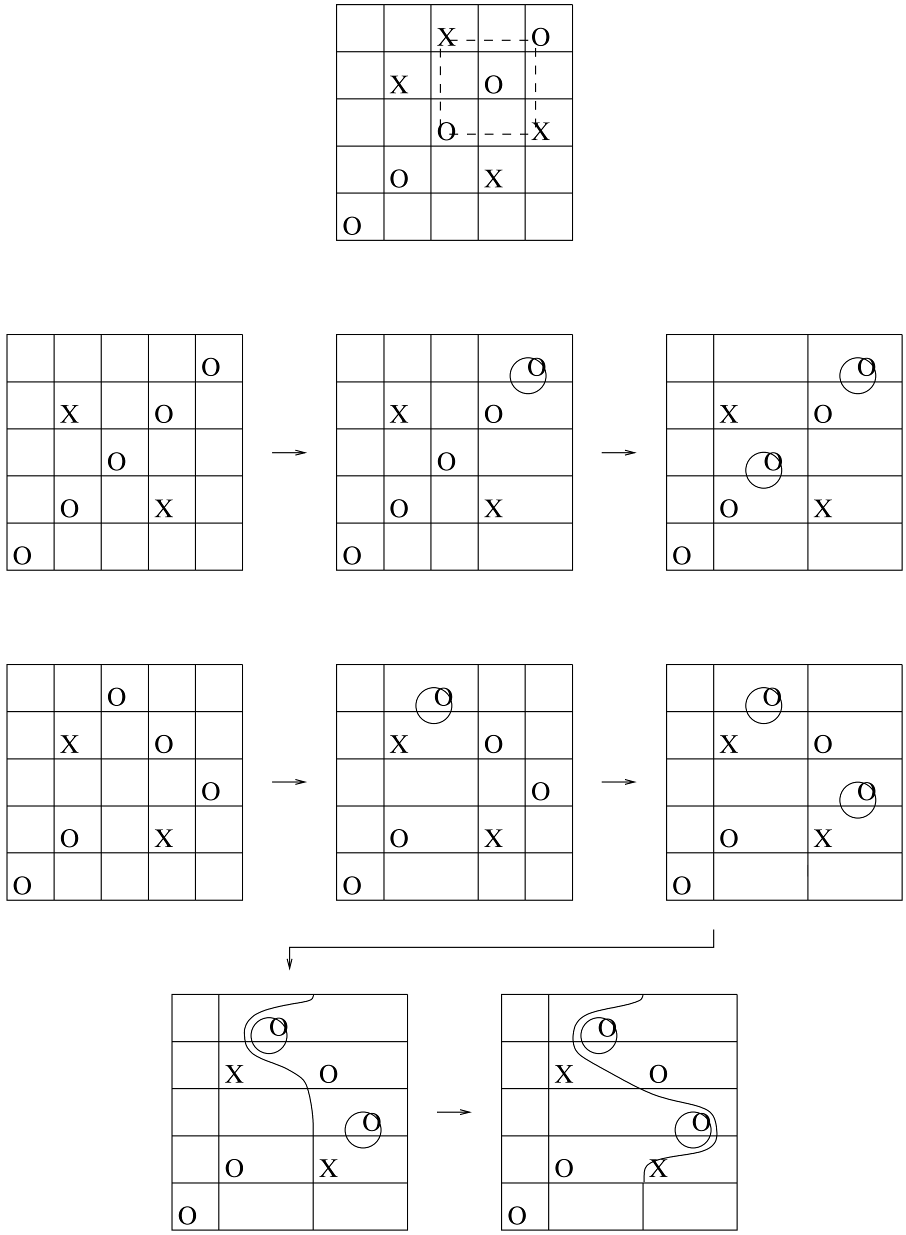

Figure 56: Some hyperboxes in a complete system for the Hopf link. We show a grid diagram G for the Hopf link $\vec{L}$ (with one free marking) and two one-dimensional hyperboxes in the complete system $\overline{\mathscr{H}}_G$ associated to the grid. At the top we have the grid G, with one component L_1 shown by the dashed lines. The second row is the hyperbox $\overline{\mathscr{H}}_G^{\vec{L},\vec{L}_1}$ of size (2), where $\vec{L}_1$ has the orientation induced from $\vec{L}$. The two bottom rows show the hyperbox $\overline{\mathscr{H}}_G^{\vec{L},-\vec{L}_1}$, of size (4). We describe each hyperbox as a sequence of Heegaard diagrams, from the initial to the final vertex. The two initial diagrams are obtained from the grid G by reduction at $\vec{L}_1$ and $-\vec{L}_1$. The two final diagrams are surface isotopic.

In fact, $\mathfrak{C}(\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(s))$ is obtained from $\mathfrak{C}(\mathcal{H}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(s))$ by adding polygon maps corresponding to handleslides over marked disks as on the right of Figure 55. According to Proposition 6.18, if we choose suitable almost complex structures, these extra higher polygon maps are the same as the corresponding polygon maps in the quasidestabilized diagrams. In the quasidestabilized diagrams, the handleslides from the right-hand side of Figure 55 are nothing more than curve isotopies. On the level of hyperboxes of Floer complexes, isotopies have (up to chain homotopy) the effect of identity shifts; compare Lemma 8.15 and the discussion of elementary enlargements in Section 8.4. It follows that $\mathfrak{C}(\overline{\mathcal{H}}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(s))$ is obtained from $\mathfrak{C}(\mathcal{H}_G^{\vec{L},\vec{M}}, \psi^{\vec{M}}(s))$ by elementary enlargements, in the sense of Section 5.8. Elementary enlargements leave unchanged the corresponding compressed hypercubes; see Lemma 5.21. Hence, the maps involved in the complexes $\mathcal{C}^-(G, \Lambda, \mathfrak{u})$ and $\mathcal{C}^-(\overline{\mathcal{H}}_G, \Lambda, \mathfrak{u})$ are the same, up to compatible chain homotopy equivalences. □

Remark 15.10 Theorems 14.3, 14.4 and 14.10 also admit straightforward adaptations to the case of grids, with the respective complexes all being expressed in terms of hyperboxes $\mathcal{H}_G^{\vec{L},\vec{M}}$. The exact statements are given in [Manolescu et al. 2009b].

References

[Atiyah and Macdonald 1969] **M F Atiyah**, **I G Macdonald**, *Introduction to commutative algebra*, Addison-Wesley, Reading, MA (1969) MR Zbl

[Baldwin 2011] **J A Baldwin**, *On the spectral sequence from Khovanov homology to Heegaard Floer homology*, Int. Math. Res. Not. 2011 (2011) 3426–3470 MR Zbl

[Bourgeois et al. 2003] **F Bourgeois**, **Y Eliashberg**, **H Hofer**, **K Wysocki**, **E Zehnder**, *Compactness results in symplectic field theory*, Geom. Topol. 7 (2003) 799–888 MR Zbl

[Fink 2009] **E Fink**, *On the twisted Floer homology of mapping tori of periodic diffeomorphisms*, preprint (2009) arXiv 0910.2010

[Juhász 2006] **A Juhász**, *Holomorphic discs and sutured manifolds*, Algebr. Geom. Topol. 6 (2006) 1429–1457 MR Zbl

[Khovanov 2000] **M Khovanov**, *A categorification of the Jones polynomial*, Duke Math. J. 101 (2000) 359–426 MR Zbl

[Khovanov and Rozansky 2008a] **M Khovanov**, **L Rozansky**, *Matrix factorizations and link homology*, Fund. Math. 199 (2008) 1–91 MR Zbl

[Khovanov and Rozansky 2008b] **M Khovanov**, **L Rozansky**, *Matrix factorizations and link homology, II*, Geom. Topol. 12 (2008) 1387–1425 MR Zbl

[Kronheimer and Mrowka 2007] **P Kronheimer**, **T Mrowka**, *Monopoles and three-manifolds*, New Math. Monogr. 10, Cambridge Univ. Press (2007) MR Zbl

[Lipshitz 2006] **R Lipshitz**, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. 10 (2006) 955–1096 MR Zbl

[Manolescu et al. 2007] **C Manolescu**, **P Ozsváth**, **Z Szabó**, **D Thurston**, *On combinatorial link Floer homology*, Geom. Topol. 11 (2007) 2339–2412 MR Zbl

[Manolescu et al. 2009a] **C Manolescu**, **P Ozsváth**, **S Sarkar**, *A combinatorial description of knot Floer homology*, Ann. of Math. 169 (2009) 633–660 MR Zbl

[Manolescu et al. 2009b] **C Manolescu**, **P Ozsváth**, **D Thurston**, *Grid diagrams and Heegaard Floer invariants*, preprint (2009) arXiv 0910.0078

[Ozsváth and Stipsicz 2010] **P Ozsváth**, **A I Stipsicz**, *Contact surgeries and the transverse invariant in knot Floer homology*, J. Inst. Math. Jussieu 9 (2010) 601–632 MR Zbl

[Ozsváth and Szabó 2003] **P Ozsváth**, **Z Szabó**, *Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary*, Adv. Math. 173 (2003) 179–261 MR Zbl

[Ozsváth and Szabó 2004a] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and knot invariants*, Adv. Math. 186 (2004) 58–116 MR Zbl

[Ozsváth and Szabó 2004b] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. 159 (2004) 1027–1158 MR Zbl

[Ozsváth and Szabó 2004c] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. 159 (2004) 1159–1245 MR Zbl

[Ozsváth and Szabó 2004d] **P Ozsváth**, **Z Szabó**, *Holomorphic triangle invariants and the topology of symplectic four-manifolds*, Duke Math. J. 121 (2004) 1–34 MR Zbl

[Ozsváth and Szabó 2005] **P Ozsváth**, **Z Szabó**, *On the Heegaard Floer homology of branched double-covers*, Adv. Math. 194 (2005) 1–33 MR Zbl

[Ozsváth and Szabó 2006] **P Ozsváth**, **Z Szabó**, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math. 202 (2006) 326–400 MR Zbl

[Ozsváth and Szabó 2008a] **P Ozsváth**, **Z Szabó**, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, Algebr. Geom. Topol. 8 (2008) 615–692 MR Zbl

[Ozsváth and Szabó 2008b] **P S Ozsváth**, **Z Szabó**, *Knot Floer homology and integer surgeries*, Algebr. Geom. Topol. 8 (2008) 101–153 MR Zbl

[Ozsváth and Szabó 2009] **P Ozsváth**, **Z Szabó**, *A cube of resolutions for knot Floer homology*, J. Topol. 2 (2009) 865–910 MR Zbl

[Ozsváth and Szabó 2011] **P S Ozsváth**, **Z Szabó**, *Knot Floer homology and rational surgeries*, Algebr. Geom. Topol. 11 (2011) 1–68 MR Zbl

[Rasmussen 2003] **J A Rasmussen**, *Floer homology and knot complements*, PhD thesis, Harvard University (2003) Available at `https://www.proquest.com/docview/305332635`

[Sarkar 2011] **S Sarkar**, *Maslov index formulas for Whitney n–gons*, J. Symplectic Geom. 9 (2011) 251–270 MR Zbl

[Sarkar 2015] **S Sarkar**, *Moving basepoints and the induced automorphisms of link Floer homology*, Algebr. Geom. Topol. 15 (2015) 2479–2515 MR Zbl

[Sarkar and Wang 2010] **S Sarkar**, **J Wang**, *An algorithm for computing some Heegaard Floer homologies*, Ann. of Math. 171 (2010) 1213–1236 MR Zbl

[Weibel 1994] **C A Weibel**, *An introduction to homological algebra*, Cambridge Stud. Adv. Math. 38, Cambridge Univ. Press (1994) MR Zbl

[Witten 1994] **E Witten**, *Monopoles and four-manifolds*, Math. Res. Lett. 1 (1994) 769–796 MR Zbl

[Zemke 2015] **I Zemke**, *Graph cobordisms and Heegaard Floer homology*, preprint (2015) arXiv 1512.01184

[Zemke 2021] **I Zemke**, *Bordered manifolds with torus boundary and the link surgery formula*, preprint (2021) arXiv 2109.11520

Department of Mathematics, Stanford University
Stanford, CA, United States

Department of Mathematics, Princeton University
Princeton, NJ, United States

cm5@stanford.edu, petero@math.princeton.edu

Proposed: Tomasz S Mrowka
Seconded: András I Stipsicz, Cameron Gordon

Received: 6 April 2017
Revised: 18 November 2022

GEOMETRY & TOPOLOGY

msp.org/gt

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2025 is US $865/year for the electronic version, and $1210/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Geometry & Topology

Volume 29 Issue 6 (pages 2783–3343) 2025