Chapter 12

The pairing theorem

In this chapter, we prove the pairing theorem. We prove Theorems 1.3 and 1.4 of the introduction, as well as some generalizations and refinements.

12.1 The statement

We now state the pairing theorem (i.e., connected sum formula). In this section, we state the pairing theorem for links in terms of a box-tensor product. In Section 12.2, we unpack the statement in terms of the link surgery complexes. In Section 12.3, we prove the statement.

Theorem 12.1. *Suppose that L_1 and L_2 are two links in S^3, which are equipped with systems of arcs $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively. Form their connected sum $L_1\#L_2$ along two distinguished components $K_1 \subseteq L_1$ and $K_2 \subseteq L_2$. Assume the arc for K_1 is alpha-parallel and that the arc for K_2 is beta-parallel. Let $\mathcal{A}_1\#\mathcal{A}_2$ denote the connected sum, in the sense of Section 9.6 (see, in particular, (#-3)). Then*

$$\begin{aligned}&\mathcal{X}_{\Lambda_1+\Lambda_2}(L_1\#L_2,\mathcal{A}_1\#\mathcal{A}_2)^{\mathcal{L}_{\ell_1+\ell_2-1}}\\ &\qquad\simeq (\mathcal{X}_{\Lambda_1}(L_1,\mathcal{A}_1)^{\mathcal{L}_{\ell_1}}\otimes_{\mathbb{F}} \mathcal{X}_{\Lambda_2}(L_2,\mathcal{A}_2)^{\mathcal{L}_{\ell_2}})\mathbin{\hat{\boxtimes}}{}_{\mathcal{K}|\mathcal{K}}M^{\mathcal{K}}.\end{aligned}$$

Remark 12.2. In Theorem 12.1, the arc of $\mathcal{A}_1\#\mathcal{A}_2$ for $K_1\#K_2$ is the co-core of the band used to form the connected sum. See Figure 9.4.

12.2 Unpacking the statement

We now give a concrete reformulation of the pairing theorem from the previous section, in terms of the knot and link surgery complexes.

We begin by unpacking the claim for the Ozsváth–Szabó mapping cone complex $\mathbb{X}_\lambda(Y_1\#Y_2, K_1\#K_2)$. For $i\in\{1,2\}$, the complex $\mathbb{X}_{\lambda_i}(Y_i,K_i)$ is the mapping cone complex

$$\mathbb{X}_{\lambda_i}(Y_i,K_i)=\mathrm{Cone}\big(\mathbb{A}(K_i)\xrightarrow{v^{(Y_i,K_i)}+h_{\lambda_i}^{(Y_i,K_i)}}\mathbb{B}(K_i)\big).$$

We recall from Lemma 6.5 that we may identify $\mathbb{A}(K_i)$ with a completion of $\mathcal{CFK}(Y_i,K_i)$, and that we may likewise identify $\mathbb{B}(K_i)$ with a completion of

$\mathcal{V}^{-1}\mathcal{CFK}(Y_i, K_i)$. From the connected sum formula for knot Floer homology, we know that

$$\mathcal{CFK}(Y_1\#Y_2, K_1\#K_2) \cong \mathcal{CFK}(Y_1, K_1) \otimes_{\mathbb{F}[\mathcal{U},\mathcal{V}]} \mathcal{CFK}(Y_2, K_2)$$

and similarly the localized module $\mathcal{V}^{-1}\mathcal{CFK}(Y_1\#Y_2, K_1\#K_2)$ decomposes as a tensor product over the ring $\mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$.

These tensor products of knot Floer complexes may themselves be described as box tensor products over the algebras $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ and $\mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$. Indeed, if $\mathbf{x}_1, \dots, \mathbf{x}_n$ is a free-basis of $\mathcal{CFK}(Y_1, K_1)$, then we may form a type-D module $\mathcal{CFK}(Y_1, K_1)^{\mathbb{F}[\mathcal{U},\mathcal{V}]}$, spanned over $\mathbb{F}$ by $\mathbf{x}_1, \dots, \mathbf{x}_n$, with structure map δ^1 encoding the differential on $\mathcal{CFK}(Y_1, K_1)$. Similarly, we may view $\mathcal{CFK}(Y_2, K_2)$ as a type-A module ${}_{\mathbb{F}[\mathcal{U},\mathcal{V}]}\mathcal{CFK}(Y_2, K_2)$, freely generated over $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ by intersection points, and with only m_1 and m_2 non-vanishing. Clearly,

$$\begin{aligned}&\mathcal{CFK}(Y_1, K_1)^{\mathbb{F}[\mathcal{U},\mathcal{V}]} \boxtimes {}_{\mathbb{F}[\mathcal{U},\mathcal{V}]}\mathcal{CFK}(Y_2, K_2)\\ &\qquad\cong \mathcal{CFK}(Y_1, K_1) \otimes_{\mathbb{F}[\mathcal{U},\mathcal{V}]} \mathcal{CFK}(Y_2, K_2).\end{aligned}$$

The maps $v^{(Y_i,K_i)}$ from the knot surgery formulas of K_1 and K_2 are the canonical inclusion maps for localization at $\mathcal{V}_i$. Hence, the map v on the tensor product is just the tensor product $v^{(Y_1,K_1)} \otimes v^{(Y_2,K_2)}$.

On the other hand, the maps $h_{\lambda_1}^{(Y_1,K_1)}$ and $h_{\lambda_2}^{(Y_2,K_2)}$ both satisfy $h_{\lambda_i}^{(Y_i,K_i)}(a\mathbf{x}) = \phi^\tau(a)h_{\lambda_i}^{(Y_i,K_i)}$ by Lemma 6.7. It is straightforward to see that this implies that their tensor product $h_{\lambda_1}^{(Y_1,K_1)} \otimes h_{\lambda_2}^{(Y_2,K_2)}$ is well defined.

In the box tensor product $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{K}} \mathbin{\hat{\boxtimes}} {}_{\mathcal{K}}\mathcal{X}_{\lambda_2}(K_2)$, there is a summand of the differential which corresponds to τ being output by $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{K}}$, and then input into ${}_{\mathcal{K}}\mathcal{X}_{\lambda_2}(K_2)$. This summand is identified with $h_{\lambda_1}^{(Y_1,K_2)} \otimes h_{\lambda_2}^{(Y_2,K_2)}$. Similarly, there is a summand of the differential that corresponds to a σ being output by $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{K}}$, and then input into ${}_{\mathcal{K}}\mathcal{X}_{\lambda_2}(K_2)$. This contributes $v_1 \otimes v_2$ to the differential.

We summarize the above observations with the following lemma.

Lemma 12.3. *If $K_1 \subseteq Y_1$ and $K_2 \subseteq Y_2$ are two knots in integer homology 3-spheres, then Theorem 12.1 is equivalent to the statement that the map $h_{\lambda_1+\lambda_2}^{(Y_1\#Y_2,K_1\#K_2)}$ in Ozsváth and Szabó's surgery formula $\mathbb{X}_{\lambda_1+\lambda_2}(Y_1\#Y_2, K_1\#K_2)$ satisfies*

$$h_{\lambda_1+\lambda_2}^{(Y_1\#Y_2,K_1\#K_2)} \simeq h_{\lambda_1}^{(Y_1,K_1)} \otimes h_{\lambda_2}^{(Y_2,K_2)}.$$

The pairing theorem for the link surgery formula has a similar restatement. Let $\mathcal{C}_{\Lambda_1}(L_1)$ and $\mathcal{C}_{\Lambda_2}(L_2)$ be the link surgery hypercubes of Manolescu and Ozsváth. Write

$$\mathcal{C}_{\Lambda_i}(L_i) \cong \operatorname{Cone}\big(\mathcal{C}_0(L_i) \xrightarrow{F^{-K_i}+F^{K_i}} \mathcal{C}_1(L_i)\big).$$

Here, $\mathcal{C}_\nu(L_i)$ consists of the complexes of all points of the cube $\mathbb{E}_{\ell_i}$ such that the coordinate for K_i is $\nu \in \{0,1\}$. Also, we are writing F^{K_i} (resp. F^{-K_i}) for the sum of the hypercube maps for all sublinks $\vec{N} \subseteq L_i$ which contain K_i (resp. $-K_i$).

Note that similar to the case of ordinary link Floer complexes, the connected sum formula [37, Section 7] provides an identification of vector spaces

$$\mathcal{C}_0(L_1\#L_2) \cong \mathcal{C}_0(L_1) \,\hat{\otimes}_{\mathbb{F}[\mathscr{U},\mathscr{V}]}\, \mathcal{C}_0(L_2)$$

and

$$\mathcal{C}_1(L_1\#L_2) \cong \mathcal{C}_1(L_1) \,\hat{\otimes}_{\mathbb{F}[\mathscr{U},\mathscr{V},\mathscr{V}^{-1}]}\, \mathcal{C}_1(L_2).$$

Here, $\hat{\otimes}$ denotes the completed tensor product.

The maps F^{K_i} are $\mathbb{F}[\mathscr{U},\mathscr{V}]$-equivariant, while F^{-K_i} are $[T]$-equivariant. In particular, both tensor products $F^{K_1} \otimes F^{K_2}$ and $F^{-K_1} \otimes F^{-K_2}$ are well defined. The same logic as with the case of connected sums of knots gives the following restatement of the pairing theorem for the surgery hypercubes.

Lemma 12.4. *Theorem 12.1 is equivalent to the statement that the surgery hypercube* $\mathcal{C}_{\Lambda_1+\Lambda_2}(L_1\#L_2, \mathcal{A}_1\#\mathcal{A}_2)$ *is homotopy equivalent to the* $(\ell_1+\ell_2-1)$*-dimensional hypercube*

$$\operatorname{Cone}\Big(\mathcal{C}_0(L_1,\mathcal{A}_1) \,\hat{\otimes}\, \mathcal{C}_0(L_2,\mathcal{A}_2) \xrightarrow{F^{K_1}\otimes F^{K_2} + F^{-K_1}\otimes F^{-K_2}} \mathcal{C}_1(L_1,\mathcal{A}_1) \,\hat{\otimes}\, \mathcal{C}_1(L_2,\mathcal{A}_2)\Big).$$

For each $\nu \in \{0,1\}$*, the above complexes* $\mathcal{C}_\nu(L_1) \,\hat{\otimes}\, \mathcal{C}_\nu(L_2)$ *are equipped with the tensor product differential* $D_1^\nu \otimes \mathrm{id} + \mathrm{id} \otimes D_2^\nu$*, where* D_j^ν *is the total differential of the hypercube* $\mathcal{C}_\nu(L_j)$ *(i.e., the sum of the internal differentials as well as the hypercube maps).*

12.3 Proof of the pairing theorem

We now prove the pairing theorem, Theorem 12.1.

Proof of Theorem 12.1. We will use the restatement in terms of mapping cones from Lemma 12.4. Our strategy is to pick a σ-basic system for which we can apply the hypercube tensor product formulas from Chapter 11.

We consider first the map $F^{-(K_1\#K_2)}$ and the corresponding Heegaard diagrams. These maps will in general have summands of length greater than 1, but since we are using a σ-basic system, the only non-trivial summands correspond to sublinks $\vec{M} \subseteq -(L_1\#L_2)$, all of whose components are oriented oppositely to $L_1\#L_2$. (See the second paragraph of [32, Section 8.7].)

We focus on the largest hyperbox in a σ-basic system, i.e., the one for $\vec{M} = -(L_1\#L_2)$, since all of the smaller hypercubes are determined by this one. We suppose that $\mathcal{H}_1$ and $\mathcal{H}_2$ are σ-basic systems of Heegaard diagrams for L_1 and L_2, which use the systems of arcs $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively. We consider the σ-basic system $\mathcal{H}_\#$ constructed in Section 9.6.

We obtain two important hyperboxes and hypercubes:

(1) A hyperbox $\mathbf{CF}^-(\mathcal{H}_\#)$ over

$$R = \mathbb{F}[U_1, \dots, U_{\ell_1 - 1}, U_{\ell_1}, U_{\ell_1+1}, \dots, U_{\ell_1+\ell_2-1}],$$

where $U_1, \dots, U_{\ell_1}$ are the variables for L_1, and $U_{\ell_1}, \dots, U_{\ell_1+\ell_2-1}$ are the variables for L_2. Here, U_{ℓ_1} is viewed as the variable for the special component $K_1\#K_2$. This hypercube is generated over R by intersection points on each constituent Heegaard diagram. If $\mathbf{d}$ is the size of the hypercube, then at each point $\varepsilon \in \mathbb{E}(\mathbf{d})$, the diagrams of $\mathcal{H}_\#$ are equipped with a complete collection of base points $\mathcal{W}_\varepsilon$, consisting of exactly one base point for each link component. When $\varepsilon = 0$, these coincide with the $\mathbf{z}$-base points. As one increases the coordinates of the cube, the $\mathbf{z}$-base points are moved into the position of the $\mathbf{w}$-base points. A holomorphic curve representing a class ϕ is weighted by the product of $U_i^{n_{p_{i,\varepsilon}}(\phi)}$, where $p_{i,\varepsilon} \in \mathcal{W}_\varepsilon$ is the base point corresponding to the link component in index i.

(2) A hypercube $\mathcal{C}_{\Lambda_1+\Lambda_2}^{-(L_1\#L_2)}$, which has the same underlying groups and internal differential as the link surgery hypercube $\mathcal{C}_{\Lambda_1+\Lambda_2}(L_1\#L_2)$ (computed using $\mathcal{H}_\#$), but has only the differentials for negatively oriented sublinks of $L_1\#L_2$.

Write $c\mathbf{CF}^-(\mathcal{H}_\#)$ for the compression of $\mathbf{CF}^-(\mathcal{H}_\#)$. If $\varepsilon \in \mathbb{E}_{\ell_1+\ell_2-1}$, write $\mathcal{C}_\varepsilon \subseteq \mathcal{C}_{\Lambda_1+\Lambda_2}(L)$ and $\boldsymbol{C}_\varepsilon \subseteq c\mathbf{CF}^-(\mathcal{H}_\#)$ for the underlying chain complexes at the point $\varepsilon \in \mathbb{E}_{\ell_1+\ell_2-1}$.

The hypercube structure maps of $\mathcal{C}_{\Lambda_1+\Lambda_2}^{-(L_1\#L_2)}$ are determined by the hypercube structure maps for $\mathbf{CF}^-(\mathcal{H}_\#)$, as we now describe. If $\varepsilon \in \mathbb{E}_{\ell_1+\ell_2-1}$, write

$$\mathcal{S}_\varepsilon \subseteq \mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_{\ell_1+\ell_2-1}, \mathcal{V}_{\ell_1+\ell_2-1}]$$

for the multiplicatively closed subset generated by $\mathcal{U}_i$ for each i such that $\varepsilon_i = 0$ and $\mathcal{V}_i$ for each i such that $\varepsilon_i = 1$.

At each point $\varepsilon \in \mathbb{E}_{\ell_1+\ell_2-1}$, we write I_ε for the inclusion map

$$I_\varepsilon \colon \mathcal{C}_\varepsilon \to \mathcal{S}_\varepsilon^{-1} \cdot \mathcal{C}_\varepsilon.$$

For each ε, there is also an isomorphism

$$\theta_\varepsilon \colon \mathcal{S}_\varepsilon^{-1}\mathcal{C}_\varepsilon \to \boldsymbol{C}_\varepsilon \otimes \mathbb{F}[\mathbb{H}(L)].$$

This map is gotten by sending a generator $a \cdot \mathbf{x}$ (where $\mathbf{x}$ is an intersection point and a is an algebra element) to $a' \cdot \mathbf{x}$, where a' is the unique element in $\mathbb{F}[U_1, \dots, U_{\ell_1+\ell_2-1}] \otimes \mathbb{F}[\mathbb{H}(L)]$ satisfying the following:

(1) $A(a') = A(\mathbf{x}) + A(a)$ (where $A(a')$ is defined by setting $A(U_i) = 0$ and $A(T^{\mathbf{s}}) = \mathbf{s}$).

(2) If $\varepsilon_i = 0$, then the power of U_i in a' is equal to the power of $\mathcal{V}_i$ in a.

(3) If $\varepsilon_i = 1$, then the power of U_i in a' is equal to the power of $\mathcal{U}_i$ in a.

Write $\Phi_{\varepsilon,\varepsilon'}$ for the hypercube structure map for $\mathcal{C}^{-(L_1\#L_2)}_{\Lambda_1+\Lambda_2}$ from $\mathcal{C}_\varepsilon$ to $\mathcal{C}_{\varepsilon'}$, and write $f_{\varepsilon,\varepsilon'}$ for the hypercube structure map for $c\mathbf{CF}^-(\mathcal{H}_\#)$ from $\boldsymbol{C}_\varepsilon$ to $\boldsymbol{C}_{\varepsilon'}$. The map $\Phi_{\varepsilon,\varepsilon'}$ is equal to

$$\Phi_{\varepsilon,\varepsilon'} = \theta_{\varepsilon'}^{-1} \circ (f_{\varepsilon,\varepsilon'} \otimes T^{\Lambda_{\varepsilon,\varepsilon'}}) \circ \theta_\varepsilon \circ I_\varepsilon. \tag{12.1}$$

In the above, $\Lambda_{\varepsilon,\varepsilon'} \in \mathbb{Z}^{\ell_1+\ell_2-1}$ denotes the sum of the columns of the framing matrix of $L_1\#L_2$ corresponding to components K_i of $L_1\#L_2$ for which $\varepsilon'_i > \varepsilon_i$.

Hence, to show the main claim it suffices to show that $c\mathbf{CF}^-(\mathcal{H}_\#)$ admits a tensor product decomposition analogous to the claimed one for $\mathcal{C}(\mathcal{H}_\#)$, and also to show that each summand of $F^{-K_1} \otimes F^{-K_2}$ has Alexander grading consistent with the surgery formula. The claim about Alexander gradings is obvious, so it suffices to address the claim about $c\mathbf{CF}^-(\mathcal{H}_\#)$.

We now consider in more detail the σ-basic system of Heegaard diagrams $\mathcal{H}_\#$ constructed in Section 9.6. The hypercubes of attaching curves which appear in this σ-basic system have a simple description in terms of tensor products of hypercubes of attaching curves. See equation (9.1). Recall that this hyperbox was constructed from σ-basic systems of Heegaard diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ for L_1 and L_2, respectively. We may assume that these hypercubes are algebraically rigid by Lemma 9.8 (see also Lemma 9.7). Recall that in Section 9.6 we defined

$$\mathcal{H}_\# = \mathrm{St}(\mathcal{H}_1\#\mathcal{H}_2^{(0)}, \mathcal{H}_1^{(1)}\#\mathcal{H}_2).$$

Here, $\mathcal{H}_1^{(0)}$ (resp. $\mathcal{H}_1^{(1)}$) is the subbox of $\mathcal{H}_1$ where the coordinate for K_1 is 0 (resp. maximal). We define $\mathcal{H}_2^{(0)}$ and $\mathcal{H}_2^{(1)}$ similarly. Write $x_1 \in \Sigma_1$ and $x_2 \in \Sigma_2$ for the connected sum points. Since $\mathcal{H}_1$ is alpha-parallel at K_1 (i.e., we use a special beta-curve as a meridian of K_1) it is straightforward to see that x_1 is base point-esque for all of the constituent alpha-hypercubes of $\mathcal{H}_1$, in the sense of Definition 11.4. (Note that x_1 is typically not base point-esque for the beta-hyperboxes in the construction, cf. Remark 11.6.) Similarly, since $\mathcal{H}_2$ is beta-parallel at K_2, x_2 will be base point-esque for all of the constituent beta-hypercubes of $\mathcal{H}_2$. See Figure 9.5. Note that x_1 and x_2 will be base point-esque for both the alpha and beta hypercubes of $\mathcal{H}_i^{(j)}$.

By Lemma 2.10, the operations of stacking and pairing hyperboxes of attaching curves commute. Using our results about connected sums of hypercubes in Propositions 10.1 and 11.8, we obtain a homotopy equivalence between $\mathbf{CF}^-(\mathcal{H}_\#)$ and the hyperbox obtained by stacking

$$\mathbf{CF}^-(\mathcal{H}_1) \otimes_{\mathbb{F}[U_\ell]} \mathbf{CF}^-(\mathcal{H}_2^{(0)}) \quad \text{and} \quad \mathbf{CF}^-(\mathcal{H}_1^{(1)}) \otimes_{\mathbb{F}[U_\ell]} \mathbf{CF}^-(\mathcal{H}_2).$$

To compress, we use the inductive construction described in Section 2.3. In doing so, we may choose to compress the axis corresponding to the special link components K_1 and K_2 last. We obtain that the compression of $\mathbf{CF}^-(\mathcal{H}_\#)$ is homotopy equivalent to the compression of a hyperbox

$$\begin{aligned} \boldsymbol{C}_0(L_1) \otimes_{\mathbb{F}[U_\ell]} \boldsymbol{C}_0(L_2) &\xrightarrow{f_1 \otimes \mathrm{id}} \boldsymbol{C}_1(L_1) \otimes_{\mathbb{F}[U_\ell]} \boldsymbol{C}_0(L_2) \\ &\xrightarrow{\mathrm{id} \otimes f_2} \boldsymbol{C}_1(L_1) \otimes_{\mathbb{F}[U_\ell]} \boldsymbol{C}_1(L_2). \end{aligned}$$

In the above, $\boldsymbol{C}_\nu(L_i)$ denotes the compression of $\mathbf{CF}^-(\mathcal{H}_i^{(\nu)})$, for $\nu \in \{0,1\}$, and f_i is the map so that the compression of $\mathbf{CF}^-(\mathcal{H}_i)$ is isomorphic to $\mathrm{Cone}(f_i)$. This is the analog for $c\mathbf{CF}^-(\mathcal{H}_\#)$ of the formula in the main statement for $\mathcal{C}_{\Lambda_1+\Lambda_2}^{-(L_1\#L_2)}$. Since the other maps in equation (12.1) are tensorial, we obtain $F^{-(K_1\#K_2)} = F^{-K_1} \otimes F^{-K_2}$, as in the statement.

We now consider the map $F^{K_1\#K_2}$. Since we are using a σ-basic system, the maps F^{K_1}, F^{K_2} and $F^{K_1\#K_2}$ are all the canonical inclusions for localizing at the $\mathcal{V}$ variable for K_1, K_2 and $K_1\#K_2$. We claim that the homotopy equivalence described above intertwines $F^{K_1\#K_2}$ with $F^{K_1} \otimes F^{K_2}$. Concretely, this amounts to the claim that the homotopy equivalence, defined above, counts the same curves on $\boldsymbol{C}_0(L_1) \otimes_{\mathbb{F}[U_\ell]} \boldsymbol{C}_0(L_2)$ as it does on $\boldsymbol{C}_1(L_1) \otimes_{\mathbb{F}[U_\ell]} \boldsymbol{C}_1(L_2)$. This is immediate from the definition. The main claim follows. ∎