

MEMOIRS OF THE EUROPEAN MATHEMATICAL SOCIETY


Ian Zemke

**Bordered Manifolds
with Torus Boundary
and the Link Surgery Formula**

MEMS Vol. 29 / 2026

EM
S 
PRESS

EM
S ■
PRESS

Memoirs of the European Mathematical Society

Edited by

Anton Alekseev (Université de Genève)

Hélène Esnault (Freie Universität Berlin)

Gerard van der Geer (Universiteit van Amsterdam)

Ari Laptev (Imperial College London)

Laure Saint-Raymond (Institut des Hautes Études Scientifiques)

Susanna Terracini (Università degli Studi di Torino)

The *Memoirs of the European Mathematical Society* publish outstanding research contributions in individual volumes, in all areas of mathematics and with a particular focus on works that are longer and more comprehensive than usual research articles.

The collection's editorial board consists of the editors-in-chief of the *Journal of the European Mathematical Society* and the *EMS Surveys in Mathematical Sciences*, along with editors of book series of the publishing house of the EMS as well as other distinguished mathematicians.

All submitted works go through a highly selective peer-review process.

Latest titles in this series:

M. Duerinckx, A. Gloria, *On Einstein's Effective Viscosity Formula*

R. Willett, G. Yu, *The Universal Coefficient Theorem for C^* -Algebras with Finite Complexity*

B. Janssens, K.-H. Neeb, *Positive Energy Representations of Gauge Groups I. Localization*

S. Dipierro, G. Giacomin, E. Valdinoci, *The Lévy Flight Foraging Hypothesis in Bounded Regions*

A. Naor, *Extension, Separation and Isomorphic Reverse Isoperimetry*

N. Lerner, *Integrating the Wigner Distribution on Subsets of the Phase Space, a Survey*

B. Adcock, S. Brugiapaglia, N. Dexter, S. Moraga, *On Efficient Algorithms for Computing Near-Best*

Polynomial Approximations to High-Dimensional, Hilbert-Valued Functions from Limited Samples

J. J. Carmona, K. Fedorovskiy, *Carathéodory Sets in the Plane*

D. Abramovich, Q. Chen, M. Gross, B. Siebert, *Punctured Logarithmic Maps*

T. Oh, M. Okamoto, L. Tolomeo, *Stochastic Quantization of the Φ_3^3 -Model*

H. Sasahira, M. Stoffregen, *Seiberg–Witten Floer Spectra for $b_1 > 0$*

J. KAAD, D. KYED, *The Quantum Metric Structure of Quantum $SU(2)$*

B. ZAVYALOV, *Almost Coherent Modules and Almost Coherent Sheaves*

J-P. LABESSE, B. LEMAIRE, *La formule des traces tordue pour un corps global de caractéristique $p > 0$*

S. MA, *Vector-Valued Orthogonal Modular Forms*

Y. ALMOG, B. HELFFER, *On the Stability of Symmetric Flows in a Two-Dimensional Channel*

A. BORS, D. PANARIO, Q. WANG, *Functional Graphs of Generalized Cyclotomic Mappings of Finite Fields*

P. SCHNEIDER, O. VENJAKOB, *Reciprocity Laws for (φ, Γ_1) -Modules over Lubin–Tate Extensions*

A.-L. DALIBARD, F. MARBACH, J. RAX, *Linear and Nonlinear Parabolic Forward-Backward Problems*

L. ERTZBISCHOFF, D. HAN-KWAN, *On Well-Posedness for Thick Spray Equations*

J. BETTINELLI, G. MIERMONT, *Compact Brownian Surfaces II. Orientable Surfaces*

P. BIRAN, O. CORNEA, J. ZHANG, *Triangulation, Persistence, and Fukaya Categories*



Ian Zemke

Bordered Manifolds with Torus Boundary and the Link Surgery Formula



EM
S ■
PRESS

Author

Ian Zemke
Department of Mathematics
University of Oregon
1021 East 13th Avenue
Eugene, OR 97403, USA
Email: izemke@uoregon.edu

Each volume of the *Memoirs of the European Mathematical Society* is available individually or as part of an annual subscription. It may be ordered from your bookseller, subscription agency, or directly from the publisher via subscriptions@ems.press.

ISSN 2747-9080, eISSN 2747-9099

ISBN 978-3-98547-110-2, eISBN 978-3-98547-610-7, DOI 10.4171/MEMS/29

© © The content of this volume is licensed under the CC BY 4.0 license, with the exception of the logo and branding of the European Mathematical Society and EMS Press, and where otherwise noted.

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>.

Published by EMS Press, an imprint of the

European Mathematical Society – EMS – Publishing House GmbH
Institut für Mathematik
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin, Germany

Email: info@ems.press

<https://ems.press>

© 2026 European Mathematical Society

Typesetting: Kira von Bülow, Hannover, Germany

Printed in Germany

© Printed on acid free paper

Abstract

In this memoir, we develop a theory of bordered HF^- using the link surgery formula of Manolescu and Ozsváth. We interpret their link surgery complexes as type-D modules over an associative algebra \mathcal{K} , which we introduce. We prove a connected sum formula, which we interpret as an A_∞ -tensor product over our algebra \mathcal{K} . Topologically, this connected sum formula may be viewed as a formula for gluing along torus boundary components.

We discuss several important examples. As a first example, if K_1 and K_2 are knots in S^3 , and Y is obtained by gluing the complements of K_1 and K_2 together using an orientation reversing diffeomorphism of their boundaries, then our theory may be used to compute $\text{CF}^-(Y)$ from $\text{CFK}^\infty(K_1)$ and $\text{CFK}^\infty(K_2)$. By computing the type-D modules for rationally framed solid tori, our theory gives a version of the link surgery formula for rationally framed links. As a final example, we use our theory to derive the Heegaard Floer homology of all 3-manifolds which bound the plumbing of a tree of disk bundles over 2-spheres.

Keywords: Heegaard Floer homology, Dehn surgery, link surgery, link Floer homology, link surgery complex, bordered Heegaard Floer homology

Mathematics Subject Classification (2020): 57K30 (primary); 57K31 (secondary)

Acknowledgments. The author would like to thank A. Alfieri, R. Lipshitz, C. Manolescu, and P. Ozsváth for helpful conversations. The author would like to thank K. Hendricks, J. Hom and M. Stoffregen for helpful conversations and the collaboration [14], where some of the results on the minimal model of the Hopf link complex were discovered. The author would also like to thank collaborators M. Borodzik and B. Liu for interesting discussions and helpful insights on related work. Finally, the author would like to thank the anonymous referee for their careful and thoughtful comments.

Funding. IZ was partially supported by NSF grants DMS-1703685 and 2204375 and a Sloan Fellowship.

Contents

1	Introduction	1
1.1	The knot surgery algebra	1
1.2	Surgery formulas as \mathcal{K} -modules	2
1.3	Pairing theorems	4
1.4	Splicing knot complements	5
1.5	Dual knots	6
1.6	Examples	7
1.7	Arc systems	7
1.8	General bordered manifolds with torus boundary	10
2	Hypercubes and Heegaard Floer homology	11
2.1	Background on Heegaard Floer homology	11
2.2	Hypercubes and hyperboxes of chain complexes	13
2.3	Compressing hyperboxes	14
2.4	A_∞ -categories and twisted complexes	17
2.5	Hypercubes of attaching curves	18
2.6	More hypercubes	20
3	A_∞-modules	23
3.1	Type-A and type-D modules	23
3.2	Bimodules	23
3.3	The Lipshitz–Ozsváth–Thurston box tensor product	24
3.4	Morphisms of modules and bimodules	25
3.5	External tensor products	28
3.6	Hypercubes as type-D modules	36
4	Homological perturbation theory	41
4.1	A_∞ -modules	41
4.2	DA bimodules	43
4.3	Hypercubes	43
4.4	Hypercubes of DA bimodules	45
5	Linear topological spaces	47
5.1	Linear maps	48
5.2	Direct products	48
5.3	Linear compactness	49

5.4	Tensor products	50
5.5	Associativity and commutativity relations	51
5.6	The tree topology	53
6	The link surgery formula	63
6.1	Statement of the link surgery formula	63
6.2	Gradings and algebraic reduction maps	65
6.3	On the algebra of the link surgery formula	68
7	The knot surgery algebra	73
7.1	The algebras \mathcal{K} and \mathcal{L}	73
7.2	Topologies on \mathcal{K} and \mathcal{L}	74
7.3	Alexander modules	77
7.4	Split Alexander modules	77
7.5	External tensor products	79
7.6	Box tensor products	83
7.7	Finitely generated \mathcal{K} -modules	84
8	Link surgery modules over \mathcal{K} and \mathcal{L}	87
8.1	The type-D module for a solid torus	87
8.2	The type-A module for a solid torus	87
8.3	The merge bimodule	89
8.4	The type-A identity bimodule	93
8.5	The knot surgery formula over \mathcal{K}	93
8.6	The link surgery formula over \mathcal{L}	95
9	σ-basic systems of Heegaard diagrams	99
9.1	Systems of arcs	99
9.2	σ -basic systems of Heegaard diagrams	100
9.3	Meridional σ -basic systems	102
9.4	Basic systems for general systems of arcs	103
9.5	Tensor products of hypercubes of attaching curves	106
9.6	Basic systems and connected sums	106
10	Hypercubes and disjoint unions	111
10.1	Small translate theorems	112
10.2	Disconnected Heegaard surfaces	114
11	Hypercubes and connected sums	121
11.1	Cylindrical boundary degenerations	124
11.2	The 0-dimensional case	125

11.3	Connected sums and hypercubes	133
12	The pairing theorem	143
12.1	The statement	143
12.2	Unpacking the statement	143
12.3	Proof of the pairing theorem	145
13	Arc systems and the link surgery formula	149
13.1	Basepoint moving maps on hypercubes	149
13.2	Basepoint actions on hypercubes	153
13.3	Homology actions on hypercubes	155
13.4	Relative homology actions and hypercubes	160
13.5	Free stabilization	165
13.6	Algebraic relations and the diffeomorphism action	167
13.7	Relating homology actions and base point actions	177
13.8	Properties of the U -action on link surgery complexes	180
14	The alpha-beta transformer	185
14.1	The bimodule	185
14.2	Basepoint actions as A_∞ -morphisms	186
14.3	The transformer as a mapping cone bimodule	188
14.4	Properties of the transformer bimodule	188
14.5	Proof of Theorem 14.1	190
15	The pair-of-pants bimodules	193
15.1	The $W_{\alpha\beta,\alpha}$ and $W_{\alpha\beta,\beta}$ bimodules	193
15.2	Proof of Theorem 15.1	194
15.3	Relation with the AA-identity bimodule	196
15.4	Relation with the bordered invariant of $S^1 \times P$	200
16	The link surgery complex of the Hopf link	203
16.1	Computation of the surgery complex	203
16.2	The Hopf link complex as a type-D module	211
16.3	The Hopf link with arc system $\mathcal{A}_{\alpha\beta}$	213
17	Minimal models for the Hopf link surgery complex	217
17.1	DA-perspectives on the Hopf link complex	217
17.2	Minimal models for the Hopf link Floer complex	218
17.3	Minimal models over \mathcal{K}	223
17.4	Comparison with the Eftekhary–Hedden–Levine model	228

18 Examples and basic properties	233
18.1 Meridional Dehn twists	233
18.2 The type-D invariants of solid tori	238
18.3 The surgery exact triangle	243
18.4 Finite generation	244
A Splicing operations	251
A.1 Splices and connected sums	251
A.2 The Hopf link and changes of parametrization	252
References	253

Chapter 1

Introduction

In this memoir, we define several algebraic objects associated to compact, oriented 3-manifolds with parametrized torus boundary, and prove a gluing formula. For the purposes of this memoir, we use the following definition (compare [27, Section 1.1]).

Definition 1.1. A *bordered manifold with torus boundary* consists of a compact and oriented manifold M such that $\partial M \cong \mathbb{T}^2$ and M is equipped with a choice of oriented basis (μ, λ) of $H_1(\partial M)$.

If M_1 and M_2 are two bordered manifolds with torus boundaries, a natural gluing operation is to glue μ_1 to μ_2 and glue λ_1 to $-\lambda_2$. This is motivated by the following fact. Suppose that M_1 and M_2 are the complements of two oriented knots K_1 and K_2 in S^3 , which are given integral framings λ_1 and λ_2 . Let $\phi: \partial M_1 \rightarrow \partial M_2$ denote the orientation reversing diffeomorphism which sends μ_1 to μ_2 and λ_1 to $-\lambda_2$. Then there is a diffeomorphism

$$M_1 \cup_{\phi} M_2 \cong S_{\lambda_1 + \lambda_2}^3(K_1 \# K_2).$$

This basic topological fact is reviewed in Appendix A.

The algebraic objects we define are reformulations of Manolescu and Ozsváth's link surgery formula [32], recast using the framework of type-D and type-A modules of Lipshitz, Ozsváth and Thurston [27].

Using the above topological observation about surgeries on connected sums, our gluing formula amounts to a formula for the behavior of the link surgery formula under connected sums.

1.1 The knot surgery algebra

We begin by describing our associative algebra \mathcal{K} , which is an algebra over the idempotent ring $\mathbf{I} = \mathbf{I}_0 \oplus \mathbf{I}_1$. Here each \mathbf{I}_i is a copy of $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$. We define

$$\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0 = \mathbb{F}[\mathcal{U}, \mathcal{V}] \quad \text{and} \quad \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1 = \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}],$$

where $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ denotes a polynomial ring in two variables and $\mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$ is its localization at \mathcal{V} . We set $\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_1 = \{0\}$. There are two special algebra elements $\sigma, \tau \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$, and we define

$$\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0 = \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}] \otimes \langle \sigma \rangle \oplus \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}] \otimes \langle \tau \rangle,$$

where $\langle \sigma \rangle$ and $\langle \tau \rangle$ denote copies of \mathbb{F} . The elements σ and τ satisfy the relations

$$\begin{aligned}\sigma\mathcal{U} &= \mathcal{U}\sigma & \text{and} & & \sigma\mathcal{V} &= \mathcal{V}\sigma, \\ \tau\mathcal{U} &= \mathcal{V}^{-1}\tau & \text{and} & & \tau\mathcal{V} &= \mathcal{U}\mathcal{V}^2\tau.\end{aligned}$$

We typically write U for the product $\mathcal{U}\mathcal{V}$.

The algebra \mathcal{K} is an associative algebra, i.e., we can think of it as an A_∞ -algebra which has vanishing actions μ_j for $j \neq 2$. Although seemingly asymmetric between \mathcal{U} and \mathcal{V} , as well as between σ and τ , the above algebra admits a more symmetrical description. See Remark 7.1. We work with an asymmetric presentation of the algebra \mathcal{K} because it mirrors an asymmetric construction of the Heegaard Floer surgery formulas which frequently simplifies computations.

1.2 Surgery formulas as \mathcal{K} -modules

Ozsváth and Szabó [43] proved that if $K \subseteq S^3$ is a knot, then the Heegaard Floer homology groups $\mathbf{HF}^-(S_\lambda^3(K))$ are computable from the knot Floer complex of K [37,47]. If λ is an integral framing on K , their construction produced a mapping cone complex $\mathbb{X}_\lambda(K)$, which satisfied

$$H_*(\mathbb{X}_\lambda(K)) \cong \mathbf{HF}^-(S_\lambda^3(K)).$$

(The bold font denotes coefficients in the power series ring $\mathbb{F}[[U]]$.)

Manolescu and Ozsváth [32] extended this formula to links in S^3 . They proved that if $L \subseteq S^3$ is a link with integral framing Λ , then there is a complex $\mathcal{C}_\Lambda(L)$, defined using the link Floer homology of L [41], such that

$$H_*(\mathcal{C}_\Lambda(L)) \cong \mathbf{HF}^-(S_\Lambda^3(L)).$$

In Section 8.5, we describe how to naturally reinterpret the data of Ozsváth and Szabó's mapping cone complex $\mathbb{X}_\lambda(K)$ both in terms of a right type-D module $\mathcal{X}_\lambda(K)^{\mathcal{K}}$ and a left type-A module ${}_{\mathcal{K}}\mathcal{X}_\lambda(K)$.

In Section 8.6, we extend this interpretation to the link surgery formula. If ℓ is a positive integer, we define the ℓ -component *link algebra* to be

$$\mathcal{L}_\ell = \bigotimes_{\mathbb{F}}^{\ell} \mathcal{K} = \mathcal{K} \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathcal{K}.$$

To an ℓ -component link $L \subseteq S^3$, we define in Section 8.6 a type-D module

$$\mathcal{X}_\Lambda(L)^{\mathcal{L}_\ell}.$$

Remark 1.2. We think of the map v in Ozsváth and Szabó's mapping cone formula as being encoded by the algebra element $\sigma \in \mathcal{K}$. The map h is encoded by the algebra element $\tau \in \mathcal{K}$.

The description of the link surgery formula given by Manolescu and Ozsváth is an infinite direct product of free modules over $\mathbb{F}[[U_1, \dots, U_\ell]]$. In particular, it is very large and sometimes challenging to do direct computations with. On the other hand, the type-D module $\mathcal{X}_\Lambda(L)^{\mathcal{L}_\ell}$ that we describe contains all of the algebraic information of the link surgery formula, but is always finite-dimensional over \mathbb{F} . For example, the complement of the 0-framed trefoil has the type-D module shown in Figure 1.1.




Figure 1.1. The type-D module for the 0-framed trefoil complement. The superscript indicates the idempotent.

In our memoir, the most fundamental object is the type-D module $\mathcal{X}_\Lambda(L)^{\mathcal{L}_\ell}$. To change a type-D algebra factor of \mathcal{K} to a type-A algebra factor, we tensor with an algebraically defined “identity bimodule”

$$\mathcal{K}|\mathcal{K}[\mathbb{I}^{\ni}].$$

Ignoring completions, $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\ni}]$ is a type-A module over $\mathcal{K} \otimes_{\mathbb{F}} \mathcal{K}$. (We write $\mathcal{K}|\mathcal{K}$ instead of $\mathcal{K} \otimes_{\mathbb{F}} \mathcal{K}$ to indicate a specific behavior with respect to completions, which we call the *split Alexander condition*.)

The bimodule $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\ni}]$ extends to a type-AA bimodule $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\ni}]_{\mathbb{F}[U]}$, which is also useful to consider. Unlike the type-D modules $\mathcal{X}_\Lambda(L)^{\mathcal{L}_\ell}$, which are all finitely generated over \mathbb{F} , the type-A and DA bimodules in our theory are usually not finitely generated.

Manolescu and Ozsváth’s link surgery formula $\mathcal{C}_\Lambda(L)$ is recovered by tensoring our type-D module $\mathcal{X}_\Lambda(L)^{\mathcal{L}_\ell}$ with ℓ -copies of the 0-framed solid torus module, which we denote by $\mathcal{K}[\mathcal{D}_0]_{\mathbb{F}[U]}$. Here, a “0-framed solid torus” refers to the complement of a 0-framed unknot in S^3 .

Algebraic completions (e.g., direct products and power series coefficients) interact with the algebra \mathcal{K} in subtle and interesting ways. The algebra \mathcal{K} has a natural filtration by right ideals

$$J_0 \supseteq J_1 \supseteq \dots$$

This filtration equips \mathcal{K} with a topology where the subspaces J_n form a basis of opens centered at 0. This equips \mathcal{K} with the structure of a *linear topological chiral algebra*, which is important for our theory. In Chapter 7, we define algebraic categories of *Alexander modules*, which consist of linear topological vector spaces with appropriately continuous actions of \mathcal{K} . These are the categories where our link surgery modules live.

1.3 Pairing theorems

One of the central results of our memoir is a connected sum formula for the Manolescu–Ozsváth integer surgery formula. Suppose L_1 and L_2 are two links in S^3 , which have framings Λ_1 and Λ_2 , as well as distinguished components $K_1 \subseteq L_1$ and $K_2 \subseteq L_2$. In Section 8.6, we describe how to compute $\mathcal{C}_{\Lambda_1+\Lambda_2}(L_1\#L_2)$ in terms of $\mathcal{C}_{\Lambda_1}(L_1)$ and $\mathcal{C}_{\Lambda_2}(L_2)$. Here, $L_1\#L_2$ denotes the connected sum of L_1 and L_2 along K_1 and K_2 , and $\Lambda_1 + \Lambda_2$ denotes the framing where we sum the framings on K_1 and K_2 and leave the others unchanged. We will interpret this connected sum formula as an A_∞ tensor product over the algebra \mathcal{K} . We use the algebraic formalism of Lipshitz, Ozsváth and Thurston [25, 27, 28], though our proof of the pairing theorem is independent of their proof in the standard bordered setting [27].

Our formula is simplest to state when we are working with knots.

Theorem 1.3. *Suppose that (Y_1, K_1) and (Y_2, K_2) are two knots in integer homology 3-spheres with integral framings λ_1 and λ_2 . Then*

$$\mathbb{X}_{\lambda_1+\lambda_2}(Y_1\#Y_2, K_1\#K_2) \simeq \mathcal{X}_{\lambda_1}(Y_1, K_1)^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{K}} \mathcal{X}_{\lambda_2}(Y_2, K_2).$$

Here, $\hat{\boxtimes}$ denotes a completed version of Lipshitz, Ozsváth and Thurston’s box tensor product [27]. There is also a version of the above theorem for links.

Theorem 1.4. *Suppose that L_1 and L_2 are two links in S^3 with integral framings Λ_1 and Λ_2 , and that we have two distinguished components $K_1 \subseteq L_1$ and $K_2 \subseteq L_2$. Then*

$$\mathcal{C}_{\Lambda_1+\Lambda_2}(L_1\#L_2) \simeq \mathcal{X}_{\Lambda_1}(L_1)^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{K}} \mathcal{X}_{\Lambda_2}(L_2).$$

In Section 8.6, we describe several more general versions of the pairing theorem. One version allows us to compute

$$\mathcal{X}_{\Lambda_1+\Lambda_2}(L_1\#L_2)^{\mathcal{L}_{\ell_1+\ell_2-1}} \hat{\boxtimes}_{\mathcal{K}} \mathcal{K}[\mathcal{D}_0]_{\mathbb{F}[U]}. \quad (1.1)$$

In the above, \mathcal{D}_0 denotes the module for the 0-framed solid torus. In the tensor product, we are tensoring the \mathcal{K} -action on \mathcal{D}_0 with the \mathcal{K} -factor of $\mathcal{L}_{\ell_1+\ell_2-1}$ which corresponds to $K_1\#K_2$. The reader should think of tensoring \mathcal{D}_0 as corresponding to gluing in a solid torus $K_1\#K_2$. See Remark 3.6 for more details about this tensor product.

We will show in Section 15.3 that the module in equation (1.1) is homotopy equivalent to an appropriate tensor product of the three modules

$$\mathcal{X}_{\Lambda_1}(L_1)^{\mathcal{L}\ell_1}, \quad \mathcal{X}_{\Lambda_2}(L_2)^{\mathcal{L}\ell_2}, \quad \text{and} \quad \mathcal{X}_{|\mathcal{K}}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}.$$

We think of tensoring with $\mathcal{X}_{|\mathcal{K}}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}$ as corresponding to gluing torus boundary components together.

In some sense, the tensor product formulas above use up the two actions of \mathcal{K} , and hence are not convenient for taking connected sums of more than two links along a single component. To take the connected sum of more than two components, we define *pair-of-pants* bimodules

$$\mathcal{X}_{|\mathcal{K}}W_{\alpha\beta,\alpha}^{\mathcal{K}} \quad \text{and} \quad \mathcal{X}_{|\mathcal{K}}W_{\alpha\beta,\beta}^{\mathcal{K}}.$$

In Chapter 15, we study these bimodules, and we prove that they can be used to compute the type-D module over \mathcal{K} after taking the connected sum of link components. There are additional modules for other combinations of α and β subscripts. There are multiple modules because of different choices of *arc systems* in the construction of the surgery formula. (See Section 1.7 for more details.)

In Section 15.4, we relate the pair-of-pants bimodules to a surgery description of the 3-manifold $P \times S^1$, where P is a pair-of-pants (i.e., a twice punctured disk).

Remark 1.5. The above pairing theorems give formulas for gluing two bordered manifolds along torus boundary components. Another important operation is to glue two boundary components of a single bordered 3-manifold together (i.e., self-gluing). Our pairing theorems do not cover this. We will investigate this in a future work.

1.4 Splicing knot complements

A basic topological operation is to glue the complements of two knots together using an orientation reversing diffeomorphism of their boundaries. 3-manifolds obtained in this manner are referred to as *splices*. Dehn surgery is a special case of splicing.

Splicing in the context of Heegaard Floer homology has been studied by many authors. See [5, 9, 10, 18] for some examples. Using bordered Heegaard Floer homology [27], it is possible to compute $\widehat{\text{HF}}$ of spliced manifolds. Despite the interest in splicing, there is to date no general description of HF^- of splices.

In this memoir, we give a description of HF^- of splices in terms of knot Floer homology. If $\phi: \mathbb{T}^2 \rightarrow \mathbb{T}^2$ is an orientation preserving diffeomorphism, then we will describe a bimodule $\mathcal{X}[\phi]^{\mathcal{K}}$. We think of this bimodule as the bordered invariant for $[0, 1] \times \mathbb{T}^2$ with basis $(\mu, -\lambda)$ on $\{0\} \times \mathbb{T}^2$ (where we write μ and λ for a standard basis of $H_1(\mathbb{T}^2)$), and $(\phi(\mu), \phi(\lambda))$ on $\{1\} \times \mathbb{T}^2$.

Theorem 1.6. *Let K_1 and K_2 be knots in S^3 and let M_1 and M_2 be their complements. Let $\phi: \partial M_1 \rightarrow -\partial M_2$ be an orientation preserving diffeomorphism. Then*

$$\mathbf{CF}^-(M_1 \cup_\phi M_2)_{\mathbb{F}[[U]]} \simeq \mathcal{X}_{\lambda_1}(K_1)^{\mathcal{K}} \widehat{\boxtimes} \mathcal{X}[\phi]^{\mathcal{K}} \widehat{\boxtimes} \mathcal{X}_{\lambda_2}(K_2)_{\mathbb{F}[[U]]}.$$

The bimodule $\mathcal{X}[\phi]^{\mathcal{K}}$ is constructed by factoring the diffeomorphism ϕ into the following generators of $\text{MCG}(\mathbb{T}^2)$:

- (1) $\phi(\mu) = \mu$ and $\phi(\lambda) = \lambda \pm \mu$.
- (2) $\phi(\mu) = \pm\lambda$ and $\phi(\lambda) = \mp\mu$.

In both of the above cases, the bimodule $\mathcal{X}[\phi]^{\mathcal{K}}$ is obtained from the surgery complex for the Hopf link. We remark that the Hopf link has complement $[0, 1] \times \mathbb{T}^2$, so taking a connected sum with the Hopf link and then surgering on the original knot has the effect of changing the boundary parametrization.

In general, we do not yet have a closed formula for the bimodule $\mathcal{X}[\phi]^{\mathcal{K}}$ for an arbitrary diffeomorphism ϕ , nor do our techniques imply independence of the bimodule $\mathcal{X}[\phi]^{\mathcal{K}}$ from the above presentation. We plan to address both of these questions in a future work.

Remark 1.7. When K_1 and K_2 are knots in S^3 , the modules $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{K}}$ and $\mathcal{X}_{\lambda_2}(K_2)_{\mathbb{F}[[U]]}$ are determined by the full knot Floer complexes $\text{CFK}^\infty(K_1)$ and $\text{CFK}^\infty(K_2)$, up to homotopy equivalence.

1.5 Dual knots

As a special case, our techniques recover the dual knot formula of Eftekhary [4] and Hedden–Levine [11] by taking a connected sum with the Hopf link. We recall that they computed the knot Floer complex of the dual knot $\mu \subseteq S_\lambda^3(K)$ in terms of the knot Floer complex $\text{CFK}^\infty(K)$. Note that the dual knot $\mu \subseteq S_\lambda^3(K)$ may be obtained by taking the connected sum of K with a Hopf link, and then performing λ -framed surgery on K .

In our present memoir case, taking the connected sum of K with a Hopf link and performing λ -framed surgery corresponds to taking the tensor product of $\mathcal{X}_\lambda(K)^{\mathcal{K}}$ with the Hopf link surgery complex, viewed as a DA bimodule $\mathcal{X}\mathcal{H}_\Lambda^{\mathcal{K}}$. Using homological perturbation theory, we describe a smaller model, which we denote by $\mathcal{X}\mathcal{Z}_\Lambda^{\mathcal{K}}$. Taking the box tensor product with this minimal model and then restricting to idempotent \mathbf{I}_0 recovers the model of Eftekhary and Hedden–Levine (modulo differing notational conventions). We sketch the equivalence of our model of the dual knot complex with the Eftekhary and Hedden–Levine model in Section 17.4.

Our analysis of the Hopf link complex builds on joint work by the author with Hendricks, Hom and Stoffregen in [14], where an extension of the dual knot formula is proven for Hendricks and Manolescu’s involutive knot Floer homology [15].

1.6 Examples

We compute many basic examples in this memoir. The most fundamental is the type-D module for the Hopf link. We compute this in Chapter 16, and compute a minimal model of the corresponding DA bimodule in Chapter 17. This example is important for applications and computations. Combined with the tensor product formulas, it also gives a computation of the Heegaard Floer homology of all 3-manifolds obtained as the boundary of a plumbing of disk bundles over S^2 , plumbed along a tree. This is explored further in [56], wherein the equivalence of Heegaard Floer homology and the lattice homology of Némethi [34, 35] is proven.

In Section 18.2, we describe the type-D modules of rationally framed solid tori $\mathcal{D}_{p/q}^{\mathcal{K}}$ for $p/q \in \mathbb{Q} \cup \{\infty\}$. The type-D modules $\mathcal{D}_{p/q}^{\mathcal{K}}$ can be viewed as naturally recovering the rational surgeries complex $\mathbb{X}_{p/q}(K)$ of Ozsváth and Szabó [44] in the sense that

$$\mathbb{X}_{p/q}(K) \cong \mathcal{D}_{p/q}^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{K}} \mathcal{X}_0(K)$$

for any knot K in S^3 .

Remark 1.8. If $L \subseteq S^3$, then by tensoring type-D modules of the form $\mathcal{D}_{p/q}^{\mathcal{K}}$ with the type-A module $\mathcal{X}_{|\dots|}^{\mathcal{K}}(L)$, we obtain a version of the link surgery formula which allows for rational surgeries on L .

As another example, we show that there is a homotopy equivalence of type-D modules

$$\mathcal{D}_{\infty}^{\mathcal{K}} \simeq \text{Cone}(f^1: \mathcal{D}_n^{\mathcal{K}} \rightarrow \mathcal{D}_{n+1}^{\mathcal{K}}),$$

for some type-D morphism f^1 . Tensoring with this homotopy equivalence recovers the surgery exact triangle in Heegaard Floer homology. This example is reminiscent of the original setting of bordered Heegaard Floer homology [27, Section 11.2].

1.7 Arc systems

We also prove several important technical results about the link surgery formula of Manolescu and Ozsváth [32]. We recall that the link surgery formula requires, for each knot component $K_i \subseteq L$, a choice of embedded arc $\mathcal{A}_i \subseteq S^3$ which connects two base points on K_i . We furthermore require the arcs to be pairwise disjoint and to be disjoint from L except at their endpoints. We call any such collection \mathcal{A} a *system of arcs*. To a system of arcs, Manolescu and Ozsváth's construction produces a chain complex $\mathcal{C}_{\Lambda}(L, \mathcal{A})$.

Given a knot $K \subseteq Y$, there are two natural choices of arc systems on K . These are obtained by connecting the two base points $w, z \in K$ with an arc that runs parallel to K . We say the arc system is *beta-parallel* if the arc runs from z to w and is oriented

in the same direction as K . We say that the arc system is *alpha-parallel* if it runs from z to w and is oriented oppositely to K . Given a Heegaard knot splitting Σ for K , the arc system is *alpha-parallel* if the arc is parallel to the subarc of K which lies in the alpha handlebody, and similarly for beta-parallel arcs. To a link $L \subseteq Y$, there are $2^{|L|}$ arc systems which can be constructed via this procedure, where we choose each arc to be either alpha-parallel or beta-parallel. Our theory also allows arc systems where some of the arcs are neither alpha-parallel nor beta-parallel. Although these seem less natural, they appear naturally in our proof of the connected sum formula.

There is an analogous asymmetry in Lipshitz, Ozsváth and Thurston's theory [27]. In their theory, for a bordered 3-manifold Y they consider Heegaard diagrams with boundary where some set of the attaching curves are properly embedded arcs. For each boundary component of Y , they must choose whether the arcs which abut that boundary component are alpha arcs or beta arcs. See Figure 1.2 for an example of a bordered Heegaard diagram with alpha arcs.




Figure 1.2. Left: A bordered Heegaard diagram for a solid torus in the Lipshitz–Ozsváth–Thurston theory. Right: A Heegaard diagram of an unknot with a beta-parallel arc system \mathcal{A} (passing through the alpha curve). In our theory, the corresponding bordered manifold would be represented by the module for an unknot (whose complement is a solid torus). The choice of alpha arcs on the boundary of the bordered Heegaard diagram on the left is analogous to the choice of a beta-parallel arc system on the right.

For a system of arcs \mathcal{A} we obtain a surgery complex $\mathcal{C}_\Lambda(L, \mathcal{A})$ and a type-D module $\mathcal{X}_\Lambda(L, \mathcal{A})^\mathcal{L}$. Manolescu and Ozsváth's result may be interpreted as showing if each arc of \mathcal{A} is alpha-parallel or beta-parallel, then

$$\mathbf{HF}^-(S_\Lambda^3(L)) \cong H_*(\mathcal{C}_\Lambda(L, \mathcal{A}))$$

as modules over $\mathbb{F}[[U]]$. They only consider arc systems which are beta-parallel. In our work, we study the dependence on the arcs system, which turns out to be rather subtle. As a first step, we prove the following theorem.

Theorem 1.9. *If \mathcal{A} is any system of arcs for L , then there is relatively graded isomorphism of vector spaces*

$$\mathbf{HF}^-(S_\Lambda^3(L)) \cong H_*(\mathcal{C}_\Lambda(L, \mathcal{A})).$$

If the arc $a_i \in \mathcal{A}$ is either alpha- or beta-parallel, then the isomorphism is of $\mathbb{F}[[U]]$ -modules, where U acts by U_i on $H_(\mathcal{C}_\Lambda(L, \mathcal{A}))$.*

The isomorphism between the link surgery complexes for different systems of arcs is somewhat subtle. For example, we will construct two systems of arcs \mathcal{A}_1 and \mathcal{A}_2 on the Hopf link H for which

$$\mathcal{X}_\Lambda(H, \mathcal{A}_1)^{\mathcal{L}_2} \not\cong \mathcal{X}_\Lambda(H, \mathcal{A}_2)^{\mathcal{L}_2} \quad (1.2)$$

as type-D modules. See Section 16.3. On the other hand, for type-D modules associated to knots, alpha-parallel and beta-parallel arc systems give the same modules. This follows from [56, Remark 5.7].

In Chapter 13, we will describe a general formula for the effect of changing the arcs in the link surgery formula. One important special case is when we want to change an arc from alpha-parallel to beta-parallel, or vice versa. This may be achieved by tensoring with a rank 1 DA bimodule $\mathcal{K}\mathcal{T}^{\mathcal{K}}$. In Chapter 14 we study this bimodule. It is algebraically related to the base point moving map investigated by Sarkar [49] and later by the author [53].

The arc systems also play an important role in the connected sum theorems and our construction of the pair-of-pants modules in Chapter 15. Note that in the Lipshitz, Ozsváth, Thurston bordered theory, the gluing formula requires one to tensor an “alpha bordered” module with an “alpha bordered” module. From a Heegaard diagrammatic perspective, this is natural because it is most natural to glue an “alpha bordered” Heegaard diagram to an “alpha bordered” Heegaard diagram to get a Heegaard diagram of the glued manifold (since we cannot glue alpha arcs to beta arcs).

It is helpful to extend the Lipshitz, Ozsváth and Thurston convention of gluing alpha-to-alpha and beta-to-beta into our theory. In our theory, gluing is accomplished by way of taking the tensor product with algebraically defined bimodules

$$\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}], \quad \mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}}, \quad \text{and} \quad \mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}.$$

(There are further modules for all combinations of α and β subscripts.) For connected sums involving $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]$, the most basic theorem we prove requires one module to be alpha bordered, and one to be beta bordered. Therefore, it is most natural to think of the module $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]$ as alpha bordered on one component, and beta bordered on the other. The labeling of the modules $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}}$, $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}$ indicates the borderedness. For example, tensoring two beta bordered type-D modules for knots into $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}$ gives the alpha bordered module for their connected sum.

1.8 General bordered manifolds with torus boundary

The connected sum formula of Theorem 1.4 and the topological interpretation of surgeries on connected sums in terms of gluing torus boundary components motivate a definition of a potential invariant of a bordered 3-manifold with torus boundary components.

We suppose that M is a bordered 3-manifold with n toroidal boundary components. We may describe M as surgery on a link L , in the complement of an unlink \mathbb{U}_n in S^3 , in a way which is compatible with the parametrization of ∂M .

We may define a type-D module $\mathcal{X}(M)^{\mathcal{L}_n}$ as follows. We begin with the link surgery formula for $L \cup \mathbb{U}_n$, interpreted as a type-D module over our algebra $\mathcal{L}_{n+\ell}$, where $\ell = |L|$. We write $\mathcal{X}_\Lambda(L \cup \mathbb{U}_n)^{\mathcal{L}_{n+\ell}}$ for this module. We define $\mathcal{X}(M)^{\mathcal{L}_n}$ by tensoring the module $\mathcal{X}_\Lambda(L \cup \mathbb{U}_n)^{\mathcal{L}_{n+\ell}}$ with ℓ type-A modules for 0-framed solid tori. This corresponds topologically to performing Dehn surgery on the components L , while identifying the components of $\partial\nu(\mathbb{U}_n)$ with ∂M . Theorem 1.4 may be interpreted as a pairing theorem for gluing two bordered 3-manifolds with torus boundary components together. We prove in Section 18.4 that the modules $\mathcal{X}(M)^{\mathcal{L}_n}$, as defined above, are always homotopy equivalent to finitely generated type-D modules, i.e., the underlying vector space $\mathcal{X}(M)$ may be taken to be finite-dimensional over \mathbb{F} .

In a future work, we hope to understand in what sense $\mathcal{X}(M)^{\mathcal{L}_n}$, defined as above, is an invariant of the bordered manifold M . In light of equation (1.2), the modules will in general depend on the choice of system of arcs \mathcal{A} in S^3 . We make the following conjecture.

Conjecture 1.10. *If M is a bordered manifold with toroidal boundary components which is equipped with a system of arcs $\mathcal{A} \subseteq M$, then the type-D module $\mathcal{X}(M, \mathcal{A})^{\mathcal{L}}$ is an invariant of (M, \mathcal{A}) up to homotopy equivalence. (In a subsequent work [57], we prove the above conjecture for arc systems where each arc is either alpha-parallel or beta-parallel.)*

Remark 1.11. Our bordered modules $\mathcal{X}(M)^{\mathcal{K}}$ give useful surgery formulas for homologically essential knots in 3-manifolds with $b_1 > 0$. This is a setting where the techniques of Ozsváth and Szabó [43] do not seem to have a simple adaptation. As a concrete example, the ∞ -framed solid torus may be viewed as the knot complement of a fiber in $S^1 \times \{\text{pt}\} \subseteq S^1 \times S^2$. We compute in Section 18.2 the associated type-D module, which is concentrated in idempotent 1 and has two generators, and has δ^1 encoded by the diagram

$$\mathbf{x}^1 \text{ --- } 1+\nu \text{ --- } \mathbf{y}^1.$$

Chapter 2

Hypercubes and Heegaard Floer homology

In this chapter, we describe some background on Heegaard Floer homology and homological algebra.

2.1 Background on Heegaard Floer homology

We will assume that the reader is familiar with the basics of Heegaard Floer homology [39], as well as its refinements for knots [37, 47] and links [41]. We provide a brief review to establish notation.

Definition 2.1 ([41, Definition 3.1]). We consider the following notions of Heegaard diagram in this memoir:

- (1) A *multi-pointed Heegaard diagram* consists of a tuple

$$(\Sigma, \alpha, \beta, \mathbf{w})$$

as follows. Here, Σ is a surface of genus g and $\mathbf{w} = \{w_1, \dots, w_n\}$ is a non-empty collection of base points. Also, $\alpha = \{\alpha_1, \dots, \alpha_{g+n-1}\}$, is a collection of pairwise disjoint embedded curves on Σ which are linearly independent in $H_1(\Sigma \setminus \mathbf{w})$. We assume $\beta = \{\beta_1, \dots, \beta_{g+n-1}\}$ satisfies the same condition. In particular, $\Sigma \setminus \alpha$ consists of n components, each of which contains a single base point $w_i \in \mathbf{w}$, and similarly for $\Sigma \setminus \beta$.

- (2) A *Heegaard link diagram with free base points* consists of a tuple

$$(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$$

satisfying the following. We assume that $\mathbf{w} = \{w_1, \dots, w_\ell\}$, $\mathbf{z} = \{z_1, \dots, z_m\}$ and $\mathbf{p} = \{p_1, \dots, p_m\}$ are collections of base points. Let $n = \ell + m$. We assume that $\alpha = \{\alpha_1, \dots, \alpha_{g+n-1}\}$ and $\beta = \{\beta_1, \dots, \beta_{g+n-1}\}$ are collections of pairwise disjoint, embedded curves which each span a rank g summand of $H_1(\Sigma)$. We assume that each component of $\Sigma \setminus \alpha$ (resp. $\Sigma \setminus \beta$) contains either a single base point of \mathbf{p} , or contains one base point from \mathbf{w} and one base point from \mathbf{z} .

Remark 2.2. Occasionally, it is helpful to allow the collections \mathbf{w} and \mathbf{z} to contain free base points. In this case, we assume that if $p \in \mathbf{w} \cup \mathbf{z}$ is the sole base point in a component of $\Sigma \setminus \alpha$, then it is the sole base point in a component of $\Sigma \setminus \beta$.

The curves $\alpha, \beta \subseteq \Sigma$ determine two handlebodies, U_α and U_β , with boundary Σ . A Heegaard link diagram with free base points determines a 3-manifold Y and an oriented link $L \subseteq Y$. In U_α , the link L is a push-off of a collection of smoothly embedded, pairwise disjoint arcs in $\Sigma \setminus \alpha$ which connect the link base points of \mathbf{w} to the link base points of \mathbf{z} . In U_β , we assume L is a push-off of a collection of smoothly embedded, pairwise disjoint arcs which connect the link base points of \mathbf{z} to the link base points of \mathbf{w} . We follow the orientation convention that L intersects Σ positively at \mathbf{z} and negatively at \mathbf{w} .

Remark 2.3. Typically we assume that if w_i and z_j are in a component of $\Sigma \setminus \alpha$, then they are also in a component of $\Sigma \setminus \beta$. This ensures that $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$ represents a link with ℓ components and each component has 2 base points.

If $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is a multi-pointed Heegaard diagram for Y and $\mathfrak{s} \in \text{Spin}^c(Y)$, we define $\text{CF}^-(\mathcal{H}, \mathfrak{s})$ to be the free $\mathbb{F}[U_1, \dots, U_n]$ -module generated by intersection points \mathbf{x} of the Lagrangian tori

$$\mathbb{T}_\alpha = \alpha_1 \times \cdots \times \alpha_{g+n-1} \quad \text{and} \quad \mathbb{T}_\beta = \beta_1 \times \cdots \times \beta_{g+n-1}$$

in the symmetric product $\text{Sym}^{g+n-1}(\Sigma)$, such that $\mathfrak{s}_\mathbf{w}(\mathbf{x}) = \mathfrak{s}$. See [39, Section 2.6] for the definition of the map $\mathfrak{s}_\mathbf{w}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \text{Spin}^c(Y)$.

The differential counts pseudo-holomorphic representatives of Maslov index 1 flowlines in $\text{Sym}^{g+n-1}(\Sigma)$ via the formula

$$\partial(\mathbf{x}) = \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi)=1}} \#(\mathcal{M}(\phi)/\mathbb{R}) U_1^{n_{w_1}(\phi)} \cdots U_n^{n_{w_n}(\phi)} \cdot \mathbf{y}, \quad (2.1)$$

extended equivariantly over $\mathbb{F}[U_1, \dots, U_n]$.

We will frequently make use of Lipshitz's *cylindrical reformulation* of Heegaard Floer homology [23]. In this formulation, the differential counts holomorphic curves of potentially higher genus in $\Sigma \times [0, 1] \times \mathbb{R}$. We refer the reader to [23] for additional background.

Definition 2.4 (Cf. [39, Section 4.2]). We make the following definitions regarding admissibility of Heegaard diagrams:

- (1) A *periodic domain* on a Heegaard multi-diagram $\mathcal{H} = (\Sigma, \boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_n, \mathbf{w})$ is a 2-chain P on Σ such that ∂P is an integral linear combination of the curves from $\boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_n$ and $n_\mathbf{w}(P) = 0$.
- (2) A Heegaard diagram for a pointed 3-manifold $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w})$ is *weakly admissible* if each non-trivial periodic domain has both positive and negative multiplicities.
- (3) If $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$ is a Heegaard link diagram with free base points, then we say that \mathcal{H} is weakly admissible if $(\Sigma, \alpha, \beta, \mathbf{q})$ is weakly admissible

for each collection $\mathbf{q} \subseteq \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$ such that \mathbf{q} contains all of the free base points and exactly one link base point from each link component.

- (4) We say that a Heegaard multi-diagram $(\Sigma, \boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_n, \mathbf{q})$ is *weakly admissible* if each periodic domain has positive and negative multiplicities. A Heegaard multi-diagram with link base points $(\Sigma, \boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_n, \mathbf{w}, \mathbf{z}, \mathbf{p})$ is called weakly admissible if it is weakly admissible for each collection $\mathbf{q} \subseteq \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$ as above.

Given a weakly admissible Heegaard diagram \mathcal{H} , we can define a completed version of the Heegaard Floer homology $\mathbf{CF}^-(\mathcal{H})$, which is freely generated by $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ (with no restriction on Spin^c structures) over the ring $\mathbb{F}[[U_1, \dots, U_n]]$.

We now discuss link Floer homology. In this memoir, we use a slightly different version of link Floer homology than appeared in [41]. If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard link diagram with no free base points and with $|\mathbf{w}| = |\mathbf{z}| = \ell$, we use a version of the link Floer complex which is a free, finitely generated chain complex over the ring $\mathbb{F}[[U_1, \dots, U_\ell, V_1, \dots, V_\ell]]$. This version was considered in [53] and [55]. We define $\mathcal{CF}\mathcal{L}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s})$ to be generated over $\mathbb{F}[[U_1, \dots, U_\ell, V_1, \dots, V_\ell]]$ by intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ such that $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$. The differential is similar to (2.1), except that we weight a holomorphic curve by

$$\mathcal{U}_1^{n_{w_1}(\phi)} \mathcal{V}_1^{n_{z_1}(\phi)} \dots \mathcal{U}_\ell^{n_{w_\ell}(\phi)} \mathcal{V}_\ell^{n_{z_\ell}(\phi)}.$$

In the case of a doubly pointed Heegaard knot diagram, $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$, we write $\mathcal{CF}\mathcal{K}(\mathcal{H}, \mathfrak{s})$ for the complex which is freely generated over $\mathbb{F}[[U, V]]$ by intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$.

2.2 Hypercubes and hyperboxes of chain complexes

We now describe Manolescu and Ozsváth's algebraic formalism of *hypercubes* and *hyperboxes* of chain complexes [32, Section 5]. Define

$$\mathbb{E}_n := \{0, 1\}^n \subseteq \mathbb{R}^n,$$

with the convention that $\mathbb{E}_0 = \{0\}$. If $\mathbf{d} = (d_1, \dots, d_n) \in \mathbb{Z}^{\geq 0}$, we define

$$\mathbb{E}(\mathbf{d}) := \{(i_1, \dots, i_n) \in \mathbb{Z}^n : 0 \leq i_j \leq d_j, \forall j\}.$$

If $\varepsilon, \varepsilon' \in \mathbb{E}(\mathbf{d})$, we write $\varepsilon \leq \varepsilon'$ if the inequality holds for each coordinate of ε and ε' . We write $\varepsilon < \varepsilon'$ if $\varepsilon \leq \varepsilon'$ and strict inequality holds at one or more coordinates.

Definition 2.5. An n -dimensional hypercube of chain complexes $(C_\varepsilon, D_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ consists of the following:

- (1) A group C_ε for each $\varepsilon \in \mathbb{E}_n$.
- (2) For each pair of indices $\varepsilon, \varepsilon' \in \mathbb{E}_n$ such that $\varepsilon \leq \varepsilon'$, a linear map

$$D_{\varepsilon, \varepsilon'}: C_\varepsilon \rightarrow C_{\varepsilon'}.$$

Furthermore, we assume the following compatibility condition holds whenever $\varepsilon \leq \varepsilon''$:

$$\sum_{\substack{\varepsilon' \in \mathbb{E}_n, \\ \varepsilon \leq \varepsilon' \leq \varepsilon''}} D_{\varepsilon', \varepsilon''} \circ D_{\varepsilon, \varepsilon'} = 0. \quad (2.2)$$

Definition 2.6. If $\mathbf{d} \in (\mathbb{Z}^{>0})^n$, a *hyperbox of chain complexes of size \mathbf{d}* consists of a collection of groups $(C_\varepsilon)_{\varepsilon \in \mathbb{E}(\mathbf{d})}$, together with a linear map $D_{\varepsilon, \varepsilon'}$ whenever $\varepsilon \leq \varepsilon'$ and $|\varepsilon' - \varepsilon|_{L^\infty} \leq 1$. Furthermore, we assume that the analog of (2.2) holds whenever $\varepsilon \leq \varepsilon''$ and $|\varepsilon'' - \varepsilon|_{L^\infty} \leq 1$.

Hypercubes and hyperboxes of the same size form dg-categories. If $\mathcal{C} = (C_\varepsilon, D_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ and $\mathcal{C}' = (C'_{\varepsilon}, D'_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ are two hypercubes of the same dimension, then a *morphism of hypercubes* F consists of a collection of linear maps

$$F_{\varepsilon, \varepsilon'}: C_\varepsilon \rightarrow C'_{\varepsilon'},$$

ranging over all $\varepsilon, \varepsilon' \in \mathbb{E}_n$ such that $\varepsilon \leq \varepsilon'$. We write $\text{Hom}(\mathcal{C}, \mathcal{C}')$ for the space of morphisms. There is a natural morphism differential ∂_{Mor} on $\text{Hom}(\mathcal{C}, \mathcal{C}')$, given by

$$\partial_{\text{Mor}}(F)_{\varepsilon, \varepsilon''} = \sum_{\substack{\varepsilon' \in \mathbb{E}_n, \\ \varepsilon \leq \varepsilon' \leq \varepsilon''}} D'_{\varepsilon', \varepsilon''} \circ F_{\varepsilon, \varepsilon'} + F_{\varepsilon', \varepsilon''} \circ D_{\varepsilon, \varepsilon'}. \quad (2.3)$$

Remark 2.7. In Section 3.6, we will give an alternate perspective on hypercubes, and show that the category of hypercubes of chain complexes is equivalent to the category of type-D modules over a certain algebra.

2.3 Compressing hyperboxes

Manolescu and Ozsváth [32, Section 5] describe an operation called *compression* which takes a hyperbox $\mathcal{C} = (C_\varepsilon, D_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}(\mathbf{d})}$ of size $\mathbf{d} = (d_1, \dots, d_n)$ and returns an n -dimensional hypercube $\widehat{\mathcal{C}} = (\widehat{C}_\varepsilon, \widehat{D}_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$. The underlying complexes are given by the formula

$$\widehat{C}_{(\varepsilon_1, \dots, \varepsilon_n)} = C_{(d_1 \varepsilon_1, \dots, d_n \varepsilon_n)}.$$

The hypercube structure maps are more complicated. We first illustrate the construction in the 1-dimensional case. Suppose \mathcal{C} is the 1-dimensional hyperbox of chain complexes

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

The compression of the above hyperbox is the hypercube

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

In Section 3.6, we give an alternate description of compression in the 1-dimensional case in terms of box-tensor products of DA bimodules over a certain algebra.

The same description also works for hyperboxes of size $(1, \dots, 1, d)$. That is, we view \mathcal{C} as a 1-dimensional hyperbox where the complex at each point of $\mathbb{E}(d)$ is a hypercube of dimension $n - 1$. Compression in this case is given by composition of hypercube morphisms.

For a hyperbox \mathcal{C} of size (d_1, \dots, d_n) , the compression may be defined by iterating the above construction. Doing so requires a choice of ordering of the axis directions. For concreteness, we pick the standard ordering of $1, \dots, n$. Given a hyperbox \mathcal{C} we define the partial compression of the i -th axis direction, denoted $c_i(\mathcal{C})$, to be the following hyperbox of size $(d_1, \dots, d_{i-1}, 1, d_{i+1}, \dots, d_n)$. We view the \mathcal{C} as a 1-dimensional hyperbox (of hyperboxes) of size (d_i) , where the complex at each point in $\mathbb{E}(d_i)$ is a hyperbox of size $(d_1, \dots, d_{i-1}, d_{i+1}, \dots, d_n)$. The hyperbox at point $j \in \{0, \dots, d_i\}$ has an underlying group consisting of the sum of all complexes at points $\varepsilon \in \mathbb{E}(\mathbf{d})$ with $\varepsilon_i = j$. We apply the composition perspective of compression for 1-dimensional hypercubes to obtain $c_i(\mathcal{C})$. We then define the compression

$$\widehat{\mathcal{C}} = (c_n \circ \cdots \circ c_1)(\mathcal{C}).$$

Compare [30, Section 4.1.2].

Of course, we can compress the hyperbox \mathcal{C} using any ordering of the axis directions. The following is well known, and we do not claim originality for the following, though we include it for completeness.

Lemma 2.8. *If \mathcal{C} is a hypercube of chain complexes and $\widehat{\mathcal{C}}$ is its compression, then $\widehat{\mathcal{C}}$ is independent up to homotopy equivalence of hypercubes from the choice of ordering of the axis directions.*

Proof. Given the inductive nature of the construction, it suffices to prove the case when $n = 2$. Let \mathcal{C} be a hyperbox of size (d_1, d_2) . Write $c_2(c_1(\mathcal{C}))$ and $c_1(c_2(\mathcal{C}))$ for the two compressions of \mathcal{C} . Note that the length 1 maps of $c_2(c_1(\mathcal{C}))$ and $c_1(c_2(\mathcal{C}))$ coincide, so it suffices to consider the diagonal maps. Since we are considering hypercubes of dimension 2, it suffices to show that the two diagonal maps are themselves chain homotopic since such a chain homotopy can be used to build a morphism of hypercubes $F: c_2(c_1(\mathcal{C})) \rightarrow c_1(c_2(\mathcal{C}))$ which is a homotopy equivalence. The components of F which preserve cube points act by the identity. We define F to have no components which increase only one cube point. We will define F to have a component from cube point $(0, 0)$ to cube point $(1, 1)$, which acts by the above chain homotopy between the length 2 maps of $c_2(c_1(\mathcal{C}))$ and $c_1(c_2(\mathcal{C}))$. Assuming we can

construct such a chain homotopy, the map F is clearly a homotopy equivalence. We now describe how to construct this chain homotopy.

Write $C_*(P_2)$ for the following chain complex over $\mathbb{F} = \mathbb{Z}/2$. There are two generators in degree 0, denoted e_1e_2 and e_2e_1 . We define $\partial(e_ie_j) = 0$. There is a single generator in degree 1, denoted $h_{1,2}$, which satisfies

$$\partial h = e_1e_2 + e_2e_1.$$

We view e_1e_2 and e_2e_1 as the paths along the boundary of a 2-dimensional cube from $(0, 0)$ to $(1, 1)$, and we think of $h_{1,2}$ as the diagonal. We think of $C_*(P_2)$ as the cell complex for the 2-dimensional permutohedron.

We now define another chain complex $C_*(P_{d_1, d_2})$. The generators consist of concatenations of e_1 , e_2 and $h_{1,2}$ such that the total number of 1 subscripts appearing is d_1 , and the total number of 2 subscripts appearing is d_2 . The differential sums over all ways of breaking a single $h_{1,2}$ into $(e_1e_2 + e_2e_1)$. We think of the generators of $C_*(P_{d_1, d_2})$ as corresponding to paths in $[0, d_1] \times [0, d_2]$ from $(0, 0)$ to (d_1, d_2) which consist of concatenations of three types of segments (horizontal unit length segments, vertical unit length segments, or diagonal segments of L^1 -length 2).

Note that the hyperbox structure maps of \mathcal{C} give a chain map

$$C_*(P_{d_1, d_2}) \rightarrow \text{Hom}(C_{0,0}, C_{d_1, d_2}).$$

The two compression procedures give chain maps

$$c_{1,2}, c_{2,1}: C_*(P_2) \rightarrow C_*(P_{d_1, d_2}).$$

Note that $c_{i,j}(e_1e_2) = e_1^{d_1}e_2^{d_2}$ and $c_{i,j}(e_2e_1) = e_2^{d_2}e_1^{d_1}$ for both choices of i, j . In order to show that the diagonals of $\widehat{\mathcal{C}}_{1,2}$ and $\widehat{\mathcal{C}}_{2,1}$ are chain homotopic, it suffices to show that the maps $c_{1,2}$ and $c_{2,1}$ are themselves chain homotopic. To do this, it suffices to show that the homology of P_{d_1, d_2} satisfies

$$H_n(P_{d_1, d_2}) \cong \begin{cases} \mathbb{F} & \text{if } n = 0, \\ 0 & \text{if } n > 0. \end{cases}$$

To establish this, we describe a filtration $\mathcal{F}_{\leq n}$ on $C_*(P_{d_1, d_2})$. Given an arrow sequence $\omega \in P_{d_1, d_2}$, view ω as giving an embedded piecewise linear curve in $[0, d_1] \times [0, d_2] \subseteq \mathbb{R}^2$. Consider the region R between ω and the arrow sequence $e_1^{d_1}e_2^{d_2}$. We define the filtration level of ω to be the number of unit squares in $[0, d_1] \times [0, d_2]$ with integral corners and whose interiors intersect R non-trivially. Note that the subspace $\mathcal{F}_{\leq n}$ is preserved by ∂ .

We now define a homotopy $H: C_*(P_{d_1, d_2}) \rightarrow C_{*+1}(P_{d_1, d_2})$. The map H sums over all ways of replacing one subword e_2e_1 with $h_{1,2}$. We observe that $\mathbb{I} + [\partial, H]$ acts by the identity on $e_1^{d_1}e_2^{d_2}$. We leave it to the reader to verify that if ω is a word

with filtration level $n > 0$, then $(\mathbb{I} + [\partial, H])(\omega)$ is a sum of arrow words of filtration level at most $n - 1$. Therefore, $(\mathbb{I} + [\partial, H])^N$ gives a homotopy equivalence between $C_*(P_{d_1, d_2})$ and \mathbb{F} (concentrated in degree 0) for $N \gg 0$. This completes the proof. ■

2.4 A_∞ -categories and twisted complexes

We now recall the notion of a *twisted complex* in an A_∞ -category. We refer to [50, Section 3.1] for more background on this construction.

We recall that an A_∞ -category \mathcal{C} consists of a set of objects $\mathcal{O}(\mathcal{C})$, a collection of morphisms $\text{Hom}(L, L')$ for $L, L' \in \mathcal{O}(\mathcal{C})$, together with a collection of higher compositions $(\mu_n)_{n \geq 1}$, as follows. We assume that if $L_0, L_1 \in \mathcal{O}(\mathcal{C})$, then $\text{Hom}(L_0, L_1)$ is a vector space. We assume that for each sequence $L_0, \dots, L_n \in \mathcal{O}(\mathcal{C})$, the map μ_n takes the form

$$\mu_n: \text{Hom}(L_0, L_1) \otimes \cdots \otimes \text{Hom}(L_{n-1}, L_n) \rightarrow \text{Hom}(L_0, L_1).$$

Furthermore, we assume that for all $\mathbf{x}_1, \dots, \mathbf{x}_n$ (where $\mathbf{x}_i \in \text{Hom}(L_{i-1}, L_i)$) the following associativity condition holds:

$$\sum_{i=1}^n \sum_{j=i}^n \mu_{n-j+i}(\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mu_{j-i+1}(\mathbf{x}_i, \dots, \mathbf{x}_j), \mathbf{x}_{j+1}, \dots, \mathbf{x}_n) = 0.$$

If \mathcal{C} is an A_∞ -category, the *additive enlargement* $\Sigma\mathcal{C}$ of \mathcal{C} is defined as follows. The objects of $\Sigma\mathcal{C}$ consist of formal sums of the form $\bigoplus_{i \in I} X_i \otimes V_i$. Here, X_i denote objects of \mathcal{C} , V_i denote graded vector spaces over $\mathbb{F} = \mathbb{Z}/2$, and I denotes a finite index set. The symbols \otimes and \bigoplus are just notation. (For our purposes, we can take $V_i = \mathbb{F}$ for all i .)

A morphism in $\Sigma\mathcal{C}$ from $\bigoplus_{i \in I} X_i \otimes V_i$ to $\bigoplus_{j \in J} Y_j \otimes W_j$ consists of a collection $(\theta_{i,j} \otimes f_{i,j})_{i \in I, j \in J}$ where $\theta_{i,j}: X_i \rightarrow X_j$ is a morphism in \mathcal{C} and $f_{i,j}: V_i \rightarrow W_j$ is a linear map of vector spaces. We will typically write such a morphism as $\sum_{(i,j) \in I \times J} \theta_{i,j} \otimes f_{i,j}$.

If \mathcal{C} is an A_∞ -category, then $\Sigma\mathcal{C}$ is naturally also an A_∞ -category, with compositions

$$\begin{aligned} \mu_n^{\Sigma\mathcal{C}} & \left(\sum_{i_0, i_1} \theta_{i_0, i_1} \otimes f_{i_0, i_1}, \dots, \sum_{i_{n-1}, i_n} \theta_{i_{n-1}, i_n} \otimes f_{i_{n-1}, i_n} \right) \\ & := \sum_{i_0, i_1, \dots, i_n} \mu_n^{\mathcal{C}}(\theta_{i_0, i_1}, \dots, \theta_{i_{n-1}, i_n}) \otimes (f_{i_{n-1}, i_n} \circ \cdots \circ f_{i_0, i_1}). \end{aligned}$$

It is straightforward to verify that $\Sigma\mathcal{C}$ is an A_∞ -category.

Finally, we define the category of twisted complexes $\text{Tw } \mathcal{C}$. A *twisted complex* (X, δ) in \mathcal{C} consists of an object $X = (X_i, V_i)_{i \in I}$ in $\Sigma\mathcal{C}$ where I is a finite partially

ordered set, together with an endomorphism $\delta \in \text{Hom}_{\Sigma\mathcal{C}}(X, X)$. If $\delta = \sum_{i,j} \delta_{i,j} \otimes f_{i,j}$, we assume that $\delta_{i,j} = 0$ unless $i < j$. Furthermore, we assume the following version of the Maurer–Cartan equation is satisfied:

$$\sum_{n \geq 1} \mu_n^{\Sigma\mathcal{C}}(\underbrace{\delta, \dots, \delta}_n) = 0.$$

A morphism in $\text{Tw}\mathcal{C}$ is the same as a morphism in $\Sigma\mathcal{C}$. Finally, we note that $\text{Tw}\mathcal{C}$ is itself an A_∞ -category, with compositions

$$\mu_n^{\text{Tw}}(\theta_1, \dots, \theta_n) := \sum_{i_0, \dots, i_n \geq 0} \mu_{n+i_0+\dots+i_n}^{\Sigma\mathcal{C}}(\underbrace{\delta, \dots, \delta}_{i_0}, \theta_1, \underbrace{\delta, \dots, \delta}_{i_1}, \theta_2, \dots, \theta_n, \underbrace{\delta, \dots, \delta}_{i_n}).$$

2.5 Hypercubes of attaching curves

The notion of a hypercube of chain complexes adapts to other categories. We now describe the appropriate notion of a hypercube in the Fukaya category. See [32, Section 8] and [24, Section 3], where such objects are referred to as *hyperboxes of Heegaard diagrams*, and *chain complexes of attaching curves*, respectively. These may naturally be viewed as twisted complexes in the Fukaya category, using the framework described in Section 2.4.

Definition 2.9. *A hypercube of handleslide equivalent attaching curves*

$$\mathcal{L}_\beta = (\beta_\varepsilon, \Theta_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$$

on (Σ, \mathbf{w}) consists of a collection of beta attaching curves β_ε , which are each pairwise related by a sequence of handleslides and isotopies, together with distinguished morphisms $\Theta_{\varepsilon, \varepsilon'} \in \mathbf{CF}^-(\Sigma, \beta_\varepsilon, \beta_{\varepsilon'}, \mathbf{w})$ whenever $\varepsilon < \varepsilon'$. We assume, furthermore, that the Heegaard multi-diagram containing all 2^n beta curves is weakly admissible (see Definition 2.4). Finally, we assume that for each pair $\varepsilon < \varepsilon' \in \mathbb{E}_n$, the following relation is satisfied:

$$\sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_n = \varepsilon'} f_{\beta_{\varepsilon_1}, \dots, \beta_{\varepsilon_n}}(\Theta_{\varepsilon_1, \varepsilon_2}, \dots, \Theta_{\varepsilon_{n-1}, \varepsilon_n}) = 0. \quad (2.4)$$

In the above, $f_{\beta_{\varepsilon_1}, \dots, \beta_{\varepsilon_n}}$ denotes the holomorphic polygon counting map.

Sometimes in Heegaard Floer theory it is helpful to have a notion of a hypercube of *beta* attaching curves, or a hypercube of *alpha* attaching curves. We define a hypercube of beta attaching curves to be the object described in Definition 2.9, however, we define a *hypercube of alpha attaching curves* as the following modification: for each $\varepsilon < \varepsilon'$ we have a chain $\Theta_{\varepsilon', \varepsilon} \in \mathbf{CF}^-(\Sigma, \alpha_{\varepsilon'}, \alpha_\varepsilon, \mathbf{w})$, and the compatibility condition in

equation (2.4) holds, as long as we write the indices in the opposite order. Clearly any hypercube of attaching curves can be viewed as either a hypercube of alpha attaching curves or beta attaching curves (by possibly reversing the order of indices), but the terminology is occasionally helpful.

If \mathcal{L}_α and \mathcal{L}_β are hypercubes of handleslide equivalent attaching curves on (Σ, \mathbf{w}) , of dimension m and n respectively, then there is an $(n + m)$ -dimensional hypercube of chain complexes

$$(C_{(\varepsilon, \nu)}, D_{(\varepsilon, \nu), (\varepsilon', \nu')})_{(\varepsilon, \nu) \in \mathbb{E}_n \times \mathbb{E}_m} = \mathbf{CF}^-(\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w}),$$

whenever the diagram containing all attaching curves is weakly admissible. For each $(\varepsilon, \nu) \in \mathbb{E}_{n+m}$, we set

$$C_{(\varepsilon, \nu)} = \mathbf{CF}^-(\Sigma, \alpha_\varepsilon, \beta_\nu, \mathbf{w}).$$

The structure maps are given by the following formula:

$$\begin{aligned} & D_{(\varepsilon, \nu), (\varepsilon', \nu')}(\mathbf{x}) \\ &= \sum_{\substack{\varepsilon = \varepsilon_1 < \dots < \varepsilon_i = \varepsilon', \\ \nu = \nu_1 < \dots < \nu_j = \nu'}} f_{\alpha_{\varepsilon_i}, \dots, \alpha_{\varepsilon_1}, \beta_{\nu_1}, \dots, \beta_{\nu_j}} (\Theta_{\varepsilon_i, \varepsilon_{i-1}}, \dots, \Theta_{\varepsilon_2, \varepsilon_1}, \mathbf{x}, \Theta_{\nu_1, \nu_2}, \dots, \Theta_{\nu_{j-1}, \nu_j}). \end{aligned}$$

It is straightforward to see that $\mathbf{CF}^-(\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w})$ is a hypercube of chain complexes.

The case when there are link base points \mathbf{w}, \mathbf{z} base points is a straightforward modification of the above construction, using variables \mathcal{U}_i and \mathcal{V}_i in the definition of the Floer complexes.

We observe that a hypercube of attaching curves $\mathcal{L} = (\gamma_\varepsilon, \theta_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ may naturally be interpreted as a twisted complex of Lagrangians in the Fukaya category. The index set I in the definition of a twisted complex is the cube \mathbb{E}_n , and the partial order is the natural ordering of cube points.

Correspondingly, there is a natural notion of a morphism between two hypercubes of attaching curves, and the resulting category is an A_∞ -category. If $\mathcal{L} = (\gamma_\varepsilon, \mathbf{x}_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ and $\mathcal{L}' = (\gamma'_{\varepsilon'}, \mathbf{y}_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ are n -dimensional hypercubes of attaching curves, then a morphism Φ from \mathcal{L} to \mathcal{L}' consists of a collection of chains $\Phi_{\varepsilon, \varepsilon'} \in \mathbf{CF}^-(\gamma_\varepsilon, \gamma'_{\varepsilon'}, \mathbf{w})$ ranging over all $\varepsilon, \varepsilon' \in \mathbb{E}_n$ with $\varepsilon \leq \varepsilon'$. Of course, the set of morphisms is none other than the Floer complex $\mathbf{CF}^-(\Sigma, \mathcal{L}, \mathcal{L}', \mathbf{w})$.

Suppose $\mathcal{L}_1, \dots, \mathcal{L}_n$ are n -dimensional hypercubes of attaching curves, where $\mathcal{L}_i = (\gamma_{i, \varepsilon}, (\Theta_i)_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$. The composition operator

$$\mu_n^{\text{Tw}}: \mathbf{CF}^-(\mathcal{L}_1, \mathcal{L}_2) \otimes \dots \otimes \mathbf{CF}^-(\mathcal{L}_{n-1}, \mathcal{L}_n) \rightarrow \mathbf{CF}^-(\mathcal{L}_1, \mathcal{L}_n)$$

is given as follows. If $\Phi_{i,i+1} \in \mathbf{CF}^-(\mathcal{L}_i, \mathcal{L}_{i+1})$, then

$$\begin{aligned} & \mu_n^{\text{Tw}}(\Phi_{1,2}, \dots, \Phi_{n-1,n}) \\ &= \sum_{i_1, \dots, i_n \geq 0} \mu_{i_1 + \dots + i_n + n} \left(\underbrace{\Theta_1, \dots, \Theta_1}_{i_1}, \Phi_{1,2}, \underbrace{\Theta_2, \dots, \Theta_2}_{i_2}, \right. \\ & \quad \left. \Phi_{2,3}, \dots, \Phi_{n-1,n}, \underbrace{\Theta_n, \dots, \Theta_n}_{i_n} \right). \end{aligned}$$

In the above, $\Theta_i = \sum_{\varepsilon < \varepsilon'} (\Theta_i)_{\varepsilon, \varepsilon'}$, viewed as an element in the group $\bigoplus_{\varepsilon < \varepsilon'} \mathbf{CF}^-(\gamma_{i,\varepsilon}, \gamma_{i,\varepsilon'})$, and similarly for $\Phi_{i,i+1}$. Additionally, $\mu_{i_1 + \dots + i_n + n}$ is declared to vanish on non-composable Floer chains, e.g., $\mu_2(\mathbf{x}, \mathbf{y}) = 0$ if $\mathbf{x} \in \mathbf{CF}^-(\gamma_0, \gamma_1)$ and $\mathbf{y} \in \mathbf{CF}^-(\gamma'_1, \gamma_2)$ where $\gamma_1 \neq \gamma'_1$.

2.6 More hypercubes

We now collect several additional constructions and basic properties. Firstly, if $\mathcal{C} = (C_\varepsilon, D_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ and $\mathcal{C}' = (C'_v, D'_{v, v'})_{v \in \mathbb{E}_m}$ are two hypercubes of chain complexes over a ring R , we may define their *external tensor product*

$$\mathcal{E} := \mathcal{C} \otimes_R \mathcal{C}',$$

as follows. The hypercube $\mathcal{E} = (E_{(\varepsilon, v)}, \delta_{(\varepsilon, v), (\varepsilon', v')})_{(\varepsilon, v) \in \mathbb{E}_n \times \mathbb{E}_m}$ is the tensor product in the ordinary sense: As groups, we set

$$E_{(\varepsilon, v)} = C_\varepsilon \otimes_R C'_v$$

and the structure maps are given by

$$\delta_{(\varepsilon, v), (\varepsilon', v')} = \begin{cases} D_{\varepsilon, \varepsilon'} \otimes \mathbb{I} & \text{if } v = v', \\ \mathbb{I} \otimes D'_{v, v'} & \text{if } \varepsilon = \varepsilon', \\ 0 & \text{if } \varepsilon < \varepsilon' \text{ and } v < v'. \end{cases}$$

Additionally, if \mathcal{C} and \mathcal{C}' are two hyperboxes of size $\mathbf{d}_0 \times \{d\}$ and $\mathbf{d}_0 \times \{d'\}$, such that the restrictions $\mathcal{C}|_{\mathbb{E}(\mathbf{d}_0) \times \{d\}}$ and $\mathcal{C}'|_{\mathbb{E}(\mathbf{d}_0) \times \{0\}}$ coincide (in both underlying groups and structure morphisms), then we may stack \mathcal{C} and \mathcal{C}' to obtain a hyperbox of size $\mathbf{d}_0 \times \{d + d'\}$, which we denote $\text{St}(\mathcal{C}, \mathcal{C}')$. The following lemma records some interactions between stacking and the external tensor product.

Lemma 2.10. *The operations of stacking and compressing satisfy the following:*

- (1) *Stacking hyperboxes of attaching curves commutes with pairing: If \mathcal{L}_α , \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are hyperboxes of attaching curves, and \mathcal{L}_β and $\mathcal{L}_{\beta'}$ may be stacked along a codimension 1 face, then*

$$\text{St}(\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta), \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_{\beta'})) \cong \mathbf{CF}^-(\mathcal{L}_\alpha, \text{St}(\mathcal{L}_\beta, \mathcal{L}_{\beta'})).$$

- (2) *Stacking hyperboxes of chain complexes commutes with tensoring hyperboxes: If \mathcal{C} and \mathcal{C}' are stackable hyperboxes of chain complexes, and \mathcal{D} is another hyperbox of chain complexes, then*

$$\mathrm{St}(\mathcal{C}, \mathcal{C}') \otimes \mathcal{D} \cong \mathrm{St}(\mathcal{C}, \mathcal{D}) \otimes \mathrm{St}(\mathcal{C}', \mathcal{D}).$$

- (3) *Compressing hyperboxes of chain complexes commutes with tensor products of hyperboxes. If \mathcal{C} and \mathcal{D} are hyperboxes of chain complexes, and $\widehat{\mathcal{C}}$ denotes compression, then there is a choice of ordering of the axis directions used to compute the compression so that*

$$\widehat{\mathcal{C}} \otimes \widehat{\mathcal{D}} \cong \widehat{\mathcal{C} \otimes \mathcal{D}}.$$

Proof. Claim (1) is essentially the definition of the pairing between hyperboxes. Similarly, (2) is clear. For claim (3), we use the inductive perspective on compression from Section 2.3. Suppose that \mathcal{C} is a hyperbox of size (d_1, \dots, d_n) and \mathcal{D} is a hyperbox of size (d'_1, \dots, d'_m) . We use the partial compression operations c described in Section 2.3. To show the main claim, it suffices to show that

$$c_i(\mathcal{C} \otimes \mathcal{D}) \cong c_i(\mathcal{C}) \otimes \mathcal{D} \quad \text{and} \quad c_{j+n}(\mathcal{C} \otimes \mathcal{D}) = \mathcal{C} \otimes c_j(\mathcal{D}) \quad (2.5)$$

if $1 \leq i \leq n$ and $1 \leq j \leq m$. We focus on the equation $c_i(\mathcal{C} \otimes \mathcal{D}) \cong c_i(\mathcal{C}) \otimes \mathcal{D}$ when $1 \leq i \leq n$. Recall that the differential on $\mathcal{C} \otimes \mathcal{D}$ is the tensor product $D_{\mathcal{C}} \otimes \mathbb{I} + \mathbb{I} \otimes D_{\mathcal{D}}$. Therefore, if we view $\mathcal{C} \otimes \mathcal{D}$ as a hyperbox of size (d_i) , then only the components of $D_{\mathcal{C}} \otimes \mathbb{I}$ will increase the $\mathbb{E}(d_i)$ index, while the components of $\mathbb{I} \otimes D_{\mathcal{D}}$ will preserve the $\mathbb{E}(d_i)$ index. From this observation, equation (2.5) is straightforward, and the main claim follows. ■

A more sophisticated extension of Lemma 2.10 concerns commuting a tensor product operation on attaching curves and the pairing operation on hypercubes of attaching curves. See Propositions 10.1 and 11.8.

Chapter 3

A_∞ -modules

In this chapter, we recall some background about A_∞ -modules. We follow the notation of [25, 27]. See [19, 20] for further background and historical context.

3.1 Type-A and type-D modules

In this section and throughout the memoir, we work over characteristic 2.

Definition 3.1. Suppose \mathcal{A} is an associative algebra. A *left A_∞ -module* ${}_{\mathcal{A}}M$ over \mathcal{A} is a left \mathbf{k} -module M equipped with a \mathbf{k} -module map

$$m_{j+1}: \mathcal{A}^{\otimes j} \otimes_{\mathbf{k}} M \rightarrow M$$

for each $j \geq 0$, such that if $a_n, \dots, a_1 \in \mathcal{A}$ and $\mathbf{x} \in M$, then

$$\begin{aligned} & \sum_{j=0}^n m_{n-j+1}(a_n, \dots, a_{j+1}, m_{j+1}(a_j, \dots, a_1, \mathbf{x})) \\ & + \sum_{k=1}^{n-1} m_n(a_n, \dots, a_{k+1}a_k, \dots, a_1, \mathbf{x}) = 0. \end{aligned}$$

An A_∞ -module ${}_{\mathcal{A}}M$ is *strictly unital* if $m_2(1, \mathbf{x}) = \mathbf{x}$ for all $\mathbf{x} \in M$, and such that for each $j > 1$, $m_{j+1}(a_j, \dots, a_1, \mathbf{x}) = 0$ if $a_i \in \mathbf{k}$ for some i .

Definition 3.2. If \mathcal{A} is an associative algebra over \mathbf{k} , a *right type-D module* $N^{\mathcal{A}}$ is right \mathbf{k} -module N , equipped with a \mathbf{k} -linear structure map

$$\delta^1: N \rightarrow N \otimes_{\mathbf{k}} \mathcal{A},$$

which satisfies

$$(\text{id}_N \otimes \mu_2) \circ (\delta^1 \otimes \text{id}_{\mathcal{A}}) \circ \delta^1 = 0.$$

Note that in the above, we are only considering associative algebras, instead of dg-algebras or A_∞ -algebras. See [25, Definition 2.2.23] for the general definition.

3.2 Bimodules

We also need to consider A_∞ -bimodules. Our exposition and notation follows [25], though we restrict to the case of associative algebras (as opposed to dg-algebras or A_∞ -algebras).

Definition 3.3. Suppose that \mathcal{A} and \mathcal{B} are algebras over \mathbf{j} and \mathbf{k} , respectively. A type-DA bimodule, denoted ${}_{\mathcal{A}}M^{\mathcal{B}}$, consists of a (\mathbf{j}, \mathbf{k}) -module M , equipped with (\mathbf{j}, \mathbf{k}) -linear structure morphisms

$$\delta_{j+1}^1: \mathcal{A}^{\otimes j} \otimes_{\mathbf{j}} M \rightarrow M \otimes_{\mathbf{k}} \mathcal{B},$$

for $j \geq 0$. These are required to satisfy the following structure relation:

$$\begin{aligned} & \sum_{j=0}^n ((\text{id}_M \otimes \mu_2) \circ (\delta_{n-j+1}^1 \otimes \text{id}_{\mathcal{B}}))(a_n, \dots, a_{j+1}, \delta_{j+1}^1(a_j, \dots, a_1, \mathbf{x})) \\ & + \sum_{k=1}^{n-1} \delta_n^1(a_n, \dots, a_{k+1}a_k, \dots, a_1, \mathbf{x}) = 0. \end{aligned}$$

3.3 The Lipshitz–Ozsváth–Thurston box tensor product

In this section, we recall the *box tensor product* of Lipshitz, Ozsváth and Thurston [27, Section 2.4].

Suppose that \mathcal{A} is an associative algebra over \mathbf{k} . If $N^{\mathcal{A}}$ is a type-D module and ${}_{\mathcal{A}}M$ is a type-A module, then the *box tensor product* $N^{\mathcal{A}} \boxtimes {}_{\mathcal{A}}M$ is the chain complex $N \otimes_{\mathbf{k}} M$, with differential ∂^{\boxtimes} as follows. Following [27], it is convenient to make the inductive notation that

$$\delta^j: N \rightarrow N \otimes \mathcal{A}^{\otimes j}$$

is the map given inductively by $\delta^j = (\delta^1 \otimes \text{id}_{\mathcal{A}^{\otimes j-1}}) \circ \delta^{j-1}$ and $\delta^0 = \text{id}$. Then the differential has the formula


$$\partial^{\boxtimes}(\mathbf{x} \otimes \mathbf{y}) = \sum_{j=0}^{\infty} (\text{id}_N \otimes m_{j+1})(\delta^j(\mathbf{x}), \mathbf{y}).$$

The differential is usually depicted via the following diagram:

$$\begin{array}{ccc} & \mathbf{x} & \mathbf{y} \\ & \downarrow & \downarrow \\ & \delta^1 & \\ & \downarrow & \downarrow \\ \partial^{\boxtimes}(\mathbf{x} \otimes \mathbf{y}) = \sum & \vdots & \\ & \downarrow & \downarrow \\ & \delta^1 & \\ & \downarrow & \downarrow \\ & m_* & \\ & \downarrow & \downarrow \end{array} \tag{3.1}$$

According to [27, Lemma 2.30], the map ∂^\boxtimes is a differential whenever one of M or N satisfies a *boundedness* assumption; see [25, Definitions 2.2.18 and 2.2.29].

There are also tensor products between various types of bimodules. The following diagram indicates schematically the box tensor products of the form $AD \boxtimes AD$, $AD \boxtimes A$ and $D \boxtimes AD$:



3.4 Morphisms of modules and bimodules

We now recall the definitions of morphisms of type-D, type-A and type-DA modules and bimodules from [25]. We assume that all algebras are associative algebras with vanishing differential. We refer the reader to [25] for the case of dg or A_∞ -algebras.

Suppose \mathcal{A} is an algebra over \mathbf{k} and $N^{\mathcal{A}}$ and $M^{\mathcal{A}}$ are two type-D modules. A morphism $f^1: N^{\mathcal{A}} \rightarrow M^{\mathcal{A}}$ is a \mathbf{k} -linear map

$$f^1: N \rightarrow M \otimes_{\mathbf{k}} \mathcal{A}.$$

There is a morphism differential, given by

$$\partial_{\text{Mor}}(f^1) = (\text{id}_M \otimes \mu_2) \circ (f^1 \otimes \text{id}_{\mathcal{A}}) \circ \delta^1 + (\text{id}_M \otimes \mu_2) \circ (\delta^1 \otimes \text{id}_{\mathcal{A}}) \circ f^1. \quad (3.2)$$

The map ∂_{Mor} squares to zero. A type-D *homomorphism* (or *cycle*) is a morphism which satisfies $\partial_{\text{Mor}}(f^1) = 0$.

A morphism f_* between two type-A modules ${}_{\mathcal{A}}M$ and ${}_{\mathcal{A}}N$ is a collection of \mathbf{k} -linear maps

$$f_{j+1}: \mathcal{A}^{\otimes j} \otimes_{\mathbf{k}} M \rightarrow N.$$

There is a morphism differential $\partial_{\text{Mor}}(f_*)$ given by

$$\begin{aligned} \partial_{\text{Mor}}(f_*)_{n+1}(a_n, \dots, a_1, \mathbf{x}) &= \sum_{j=0}^n m_{n-j+1}(a_n, \dots, a_{j+1}, f_{j+1}(a_j, \dots, a_1, \mathbf{x})) \\ &\quad + \sum_{j=0}^n f_{n-j+1}(a_n, \dots, a_{j+1}, m_{j+1}(a_j, \dots, a_1, \mathbf{x})) \\ &\quad + \sum_{i=1}^{n-1} f_n(a_n, \dots, a_{i+1}a_i, \dots, a_1, \mathbf{x}). \end{aligned}$$

Finally, if ${}_{\mathcal{A}}M^{\mathcal{B}}$ and ${}_{\mathcal{A}}N^{\mathcal{B}}$ are DA bimodules, then a morphism f_*^1 from ${}_{\mathcal{A}}M^{\mathcal{B}}$ to ${}_{\mathcal{A}}N^{\mathcal{B}}$ is a collection of maps

$$f_{j+1}^1: \mathcal{A}^{\otimes j} \otimes M \rightarrow N \otimes \mathcal{B}$$

for $j \geq 0$. The morphism differential is given by

$$\begin{aligned} \partial_{\text{Mor}}(f_*^1)_{n+1}(a_n, \dots, a_1, \mathbf{x}) &= \sum_{j=0}^n (\text{id}_N \otimes \mu_2)(\delta_{n-j+1}^1 \otimes \text{id}_{\mathcal{B}})(a_n, \dots, a_{j+1}, f_{j+1}^1(a_j, \dots, a_1, \mathbf{x})) \\ &\quad + \sum_{j=0}^n (\text{id}_N \otimes \mu_2)(f_{n-j+1}^1 \otimes \text{id}_{\mathcal{B}})(a_n, \dots, a_{j+1}, \delta_{j+1}^1(a_j, \dots, a_1, \mathbf{x})) \\ &\quad + \sum_{i=1}^{n-1} f_n^1(a_n, \dots, a_{i+1}a_i, \dots, a_1, \mathbf{x}). \end{aligned}$$

Of course, the morphism differential ∂_{Mor} also squares to zero for morphisms of type-A and DA morphisms. See [25, Definition 2.2.43].

If $f_*^1: {}_{\mathcal{A}}M^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}N^{\mathcal{B}}$ and $g_*^1: {}_{\mathcal{A}}N^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}Q^{\mathcal{B}}$ are DA bimodule morphisms, Lipshitz–Ozsváth–Thurston define the composition $f_*^1 \circ g_*^1$ via the following diagram:

$$g_*^1 \circ f_*^1 = \begin{array}{ccc} \mathbf{a} & & \mathbf{x} \\ \Downarrow & & \downarrow \\ \Delta & \searrow & f_*^1 \\ & \searrow & \downarrow \\ & & g_*^1 \\ & & \downarrow \\ & & \mu_2 \end{array} \tag{3.3}$$

In equation (3.3), Δ denotes the co-multiplication

$$\Delta: \mathcal{A}^{\otimes n} \rightarrow \bigoplus_{i=0}^n \mathcal{A}^{\otimes i} \otimes \mathcal{A}^{\otimes (n-i)}$$

given by


$$\Delta(a_1 | \dots | a_n) = \sum_{i=1}^n (a_1 | \dots | a_i) \otimes (a_{i+1} | \dots | a_n).$$

In the above, $|$ denotes tensor product.

We now discuss pairing morphisms of modules and bimodules. We focus on the case of DA bimodule morphisms. The other types of pairings which are relevant to our memoir are defined by forgetting about the input algebra or output algebra in the diagrams below. The natural algebraic operation is to pair a bimodule morphism with the identity (as opposed to pairing two morphisms together). Suppose that ${}_{\mathcal{A}}N^{\mathcal{B}}$ and ${}_{\mathcal{A}}M^{\mathcal{B}}$ and ${}_{\mathcal{B}}P^{\mathcal{C}}$ are DA bimodules, and

$$f_*^1: {}_{\mathcal{A}}N^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}M^{\mathcal{B}}$$

is a morphism of DA bimodules. There is a morphism of DA bimodules $f_*^1 \boxtimes \mathbb{I}: N \boxtimes P \rightarrow M \boxtimes P$, defined by the following diagram:




In the above, the map δ denotes the sum over $n \geq 0$ of the maps

$$\delta_{j+1}^n: \underbrace{\mathcal{A} \otimes \dots \otimes \mathcal{A}}_j \otimes M \rightarrow M \otimes \underbrace{\mathcal{B} \otimes \dots \otimes \mathcal{B}}_n$$

obtained by composing δ_*^1 repeatedly, but not multiplying the \mathcal{B} -outputs. (Note that by convention $\delta_*^0 = \mathbb{I}$.)

If instead $g_*^1: {}_{\mathcal{B}}P^{\mathcal{C}} \rightarrow {}_{\mathcal{B}}Q^{\mathcal{C}}$ is a morphism of type-DA modules, then $\mathbb{I} \boxtimes g_*^1$ is given by the following diagram:



3.5 External tensor products

In this section, we recall a basic algebraic operation called the *external tensor product*. We consider these for type-D and type-DA modules over associative algebras, and we prove some basic properties.

We begin with the setting of type-D modules. Let \mathcal{A} and \mathcal{B} algebras over rings \mathbf{i} and \mathbf{j} of characteristic 2, and let $N^{\mathcal{A}} = (N, \delta^1)$ and $M^{\mathcal{B}} = (M, d^1)$ be type-D modules. Then the *external tensor product*, denoted $(N \otimes_{\mathbb{F}} M)^{\mathcal{A} \otimes_{\mathbb{F}} \mathcal{B}}$, is the group $M \otimes_{\mathbb{F}} N$ equipped with the structure map $\text{id}_N \otimes 1_{\mathcal{A}} \otimes d^1 + \delta^1 \otimes \text{id}_M \otimes 1_{\mathcal{B}}$ (with tensor factors reordered).

The external tensor product of two type-DA bimodules is more complicated. In general, one must choose a cellular diagonal of the associahedron. We follow the presentation of Lipshitz, Ozsváth and Thurston [27, 28]. The main goal of this section is to give a concrete formula (not involving diagonals of the associahedron) of the external tensor product of two DA bimodules, which we use throughout the memoir. This formula appears in the statement of Proposition 3.7.

Our description is in terms of a more basic operation, which we call *extension of scalars* (and is also inspired by [25]). Suppose that \mathcal{A} , \mathcal{B} and \mathcal{C} are algebras over rings \mathbf{i} , \mathbf{j} and \mathbf{k} , respectively. If ${}_{\mathcal{A}}N^{\mathcal{B}}$ is a bimodule, then we will describe a bimodule ${}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(N; \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}}$ which we say is formed by extension of scalars. We will relate this operation with an external tensor product with the identity bimodule ${}_{\mathcal{C}}[\mathbb{I}]^{\mathcal{C}}$ in Proposition 3.18.

As an $(\mathbf{i} \otimes \mathbf{k}, \mathbf{j} \otimes \mathbf{k})$ -module, we declare $\mathcal{E}(N; \mathcal{C}) := N \otimes_{\mathbb{F}} \mathbf{k}$. If $k \in \mathbf{k}$, then we define

$$\delta_{j+1}^1(a_j | c_j, \dots, a_1 | c_1, \mathbf{x} | k) = \delta_{j+1}^1(a_j, \dots, a_1, \mathbf{x}) \otimes k \otimes (c_j \cdots c_1). \quad (3.4)$$

Note that $\delta_1^1(\mathbf{x} | k) = \delta_1^1(\mathbf{x}) \otimes k \otimes 1_{\mathcal{C}}$. (Note we are implicitly reordering the k factor and the \mathcal{B} factor above.)

Given a morphism of DA bimodules $f_*^1: {}_{\mathcal{A}}N^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}M^{\mathcal{B}}$, there is a morphism

$$\mathcal{E}(f)_*^1: {}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(N; \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}} \rightarrow {}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(M; \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}}$$

given by replacing δ_{j+1}^1 with f_{j+1}^1 in equation (3.4).

Remark 3.4. If ${}_{\mathcal{A}}M^{\mathcal{B}}$ is operationally bounded (i.e., $\delta_j^1 = 0$ for large j), then the same is true for $\mathcal{E}(M, \mathcal{C})$. A similar statement holds for morphisms.

Lemma 3.5. *Suppose that \mathcal{A} , \mathcal{B} , and \mathcal{C} are algebras and suppose that ${}_{\mathcal{A}}N^{\mathcal{B}}$ is a DA bimodule.*

- (1) ${}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(N; \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}}$ satisfies the DA bimodule structure relations.
- (2) If f_*^1 is a morphism of DA bimodules, then $\partial_{\text{Mor}}(\mathcal{E}(f_*^1)) = \mathcal{E}(\partial_{\text{Mor}}(f_*^1))$.
- (3) If ${}_{\mathcal{A}}M^{\mathcal{B}}$ and ${}_{\mathcal{B}}N^{\mathcal{D}}$ are bimodules and ${}_{\mathcal{B}}N^{\mathcal{D}}$ is operationally bounded, then there is a canonical isomorphism of DA bimodules

$${}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}({}_{\mathcal{A}}M^{\mathcal{B}} \boxtimes {}_{\mathcal{B}}N^{\mathcal{D}}; \mathcal{C})^{\mathcal{D} \otimes \mathcal{C}} \cong {}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(M; \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}} \boxtimes {}_{\mathcal{B} \otimes \mathcal{C}}\mathcal{E}(N; \mathcal{C})^{\mathcal{D} \otimes \mathcal{C}}.$$

- (4) With respect to the above isomorphism, we also have

$$\mathcal{E}(f_*^1 \boxtimes \mathbb{I}) = \mathcal{E}(f_*^1) \boxtimes \mathbb{I} \quad \text{and} \quad \mathcal{E}(\mathbb{I} \boxtimes g_*^1) = \mathbb{I} \boxtimes \mathcal{E}(g_*^1),$$

for any DA bimodule morphisms f_*^1 and g_*^1 .

Proof. The first claim follows from the bimodule relations for N , together with the fact that multiplication on \mathcal{C} is associative.

For the second claim, we note that each term of $\partial_{\text{Mor}}(\mathcal{E}(f_*^1))$, when evaluated on $(b_n|c_n), \dots, (b_1|c_1), \mathbf{x}|1_{\mathbf{k}}$, is equal to a corresponding term of $\partial_{\text{Mor}}(f_*^1)$. It is easy to see that all of the terms cancel.

For the third claim, the map is $f_1^1(\mathbf{x} \otimes \mathbf{y}) = \mathbf{x} \otimes \mathbf{y} \otimes 1_{\mathcal{A} \otimes \mathcal{C}}$ and $f_j^1 = 0$ if $j > 1$. It is easy to see that this is an isomorphism of DA bimodules.

The final claim is verified in a similar manner. ■

Remark 3.6. The external tensor product gives a way to tensor a type-D module $\mathcal{X}^{\mathcal{A} \otimes \mathcal{B}}$ with a type-A module ${}_{\mathcal{A}}\mathcal{Y}$. Namely, we take the ordinary box tensor product of $\mathcal{X}^{\mathcal{A} \otimes \mathcal{B}}$ with ${}_{\mathcal{A} \otimes \mathcal{B}}\mathcal{E}(\mathcal{Y}, \mathcal{B})^{\mathcal{B}}$. The reader may recognize that this procedure is identical to Lipshitz, Ozsváth and Thurston's box tensor product between a type-DD module and a type-A module [25, Figure 4].

We now consider an additional commutativity property of extension of scalars.

Proposition 3.7. *Suppose that ${}_{\mathcal{A}}M^{\mathcal{B}}$ and ${}_{\mathcal{C}}N^{\mathcal{D}}$ are DA bimodules over associative algebras. Then*

$${}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(M, \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}} \boxtimes {}_{\mathcal{B} \otimes \mathcal{C}}\mathcal{E}(N, \mathcal{B})^{\mathcal{B} \otimes \mathcal{D}}$$

and

$${}_{\mathcal{A} \otimes \mathcal{C}}\mathcal{E}(N, \mathcal{A})^{\mathcal{A} \otimes \mathcal{D}} \boxtimes {}_{\mathcal{A} \otimes \mathcal{D}}\mathcal{E}(M, \mathcal{D})^{\mathcal{B} \otimes \mathcal{D}}$$

are homotopy equivalent as DA bimodules. In fact, both are homotopy equivalent to the external tensor product of ${}_{\mathcal{A}}M^{\mathcal{B}}$ and ${}_{\mathcal{C}}N^{\mathcal{D}}$ along a cellular diagonal of the associahedron (defined below).

Remark 3.8. If one of M and N has the property that $\delta_j^1 = 0$ for $j > 2$, then the above proposition may be easily proven directly since the structure maps on the two tensor products are identical.

As alluded to in its statement, Proposition 3.7 is most naturally proven by considering a more general tensor product operation on bimodules, called the *external tensor product*. The construction requires a choice of a cellular diagonal of Stasheff’s associahedron [51, 52]. We recall some background about diagonals presently. Our exposition will follow the work of Lipshitz, Ozsváth and Thurston [28], though we refer the reader to earlier work of Saneblidze–Umble [48] on diagonals of the associahedron. See also [33] and [31].

We recall that the $(n - 2)$ -dimensional associahedron, denoted K_n , has a cell decomposition which admits a simple combinatorial description, as follows. We consider a chain complex $C_*^{\text{cell}}(K_n)$ over \mathbb{F} whose generators are identified with planar trees T , which have $n + 1$ valence 1 vertices, which are partitioned into n “input” vertices and one “output” vertex. Furthermore, trees T are assumed to have no vertices of valence 2. The dimension of a generator $[T]$ is given by

$$\text{deg}(T) = \sum_{v \in V_{\text{int}}(T)} (\text{val}(v) - 3),$$

where $V_{\text{int}}(T)$ denotes the interior vertices of T . The differential on $C_*^{\text{cell}}(K_n)$ is defined by setting $\partial(T)$ to be the sum of all trees obtained by splitting an interior vertex into two interior vertices by adding a single edge.

The external tensor product operation has as input a pair of DA bimodules ${}_{\mathcal{A}}M^{\mathcal{B}}$ and ${}_{\mathcal{C}}N^{\mathcal{D}}$, and has as output a DA bimodule ${}_{\mathcal{A} \otimes \mathcal{C}}(M \otimes_{\Gamma} N)^{\mathcal{B} \otimes \mathcal{D}}$. The underlying vector space of $M \otimes_{\Gamma} N$ is the ordinary tensor product $M \otimes_{\mathbb{F}} N$. The DA bimodule structure map δ_{j+1}^1 depends on an extra piece of data, called a *cellular diagonal of the associahedron* Γ .

Definition 3.9 ([28, Definition 2.13]). A *cellular diagonal of the associahedron* Γ consists of a collection of degree preserving chain maps

$$\Gamma_n: C_*^{\text{cell}}(K_n) \rightarrow C_*^{\text{cell}}(K_n) \otimes_{\mathbb{F}} C_*^{\text{cell}}(K_n),$$

for $n \geq 2$, which are compatible with concatenation of trees (in the sense described below), and such that Γ_2 is the canonical isomorphism between $C_*^{\text{cell}}(K_2) \cong \mathbb{F}$ and $C_*^{\text{cell}}(K_2) \otimes_{\mathbb{F}} C_*^{\text{cell}}(K_2) \cong \mathbb{F}$.

The compatibility condition is as follows. Firstly, there is a natural splicing operation

$$\circ_i: C_*^{\text{cell}}(K_n) \otimes C_*^{\text{cell}}(K_m) \rightarrow C_*^{\text{cell}}(K_{n+m-1}),$$

where $T \circ_i S$ joins the output of S into the i -th input of T . We define \circ_i also for tensors of trees by the formula

$$(T_1 \otimes T_2) \circ_i (S_1 \otimes S_2) = (T_1 \otimes_i S_1) \otimes (T_2 \otimes_i S_2).$$

The compatibility condition of Definition 3.9 is that if T and S have n and m inputs, respectively, then

$$\Gamma_{n+m-1}(T \circ_i S) = \Gamma_n(T) \circ_i \Gamma_m(S).$$

We write $\Psi_n \in C_*^{\text{cell}}(K_n)$ for the tree with n inputs and exactly one interior vertex. The generator Ψ_n is called the *corolla* with n inputs and has homological degree $n - 2$. This is the maximal non-trivial grading of $C_*^{\text{cell}}(K_n)$.

If ${}_{\mathcal{A}}M^{\mathcal{B}} = (M, \delta_*^1)$ is a DA bimodule, then to each planar tree T with $(n + 1)$ -inputs, we may form an evaluation map $\delta_T^1: \mathcal{A}^{\otimes n} \otimes M \rightarrow M \otimes \mathcal{B}$ as follows. We first compose the maps δ_j^1 according to the tree T . This naturally gives a map from $\mathcal{A}^{\otimes n} \otimes M$ to $M \otimes \mathcal{B}^{\otimes k}$, where k is the number of vertices in the right-most path of T . We then apply $\mu_2^{\mathcal{B}}$ repeatedly to obtain a map from $\mathcal{A}^{\otimes n} \otimes M$ to $M \otimes \mathcal{B}$.

This gives a map

$$\text{ev}: C_*^{\text{cell}}(K_{n+1}) \rightarrow \text{Hom}_{(\mathbf{k}, \mathbf{j})}(\mathcal{A}^{\otimes n} \otimes M, M \otimes \mathcal{B}).$$

We equip $\text{Hom}_{(\mathbf{k}, \mathbf{j})}(\mathcal{A}^{\otimes n} \otimes M, M \otimes \mathcal{B})$ with a differential ∂ given by

$$\partial(f_{j+1}^1) = \delta_1^1 \circ f_{j+1}^1 + f_{j+1}^1 \circ \delta_1^1.$$

The statement that the structure maps δ_*^1 of ${}_{\mathcal{A}}M^{\mathcal{B}}$ satisfy the DA bimodule relations is equivalent to the statement that $M^{\mathcal{B}} \boxtimes_{\mathcal{B}} \mathcal{B}$ is a chain complex, and that ev is a chain map.

We define δ_1^1 on $M \otimes_{\mathbb{F}} N$ via the Leibniz rule, and for $n > 0$ we define the structure map δ_{n+1}^1 as $(\text{ev} \otimes \text{ev})(\Gamma_{n+1}(\Psi_{n+1}))$. The fact that the DA bimodule structure relations hold is immediate.

Since we are focused only on associative algebras, it is helpful to instead focus our attention on a slightly more coarse definition of a diagonal.

Definition 3.10. For $n \geq 2$, we define a chain complex $M_{n,*}^{\text{red}}$ as the following modification of $C_*^{\text{cell}}(K_n)$. We call $M_{n,*}^{\text{red}}$ the complex of *left-reduced module trees*. The vector space $M_{n,*}^{\text{red}}$ is generated by planar trees T with n -inputs, such that all vertices of valence 4 or more occur along the right-most path in T from an input to the output. Additionally, we quotient by the relation that $T \circ_i S = T \circ_i S'$ whenever $i < n$ and S and S' have degree 0 (i.e., only valence 3 internal vertices) and the same number of inputs. See Figure 3.1 for examples.

We define the differential ∂ on $M_{n,*}^{\text{red}}$ the same as in $C_*^{\text{cell}}(K_n)$. It is not hard to see that $\partial^2 = 0$.




Figure 3.1. Equalities of trees in $M_{n,*}^{\text{red}}$.

For the purposes of defining the diagonal tensor product of two DA bimodules over associative algebras, we will use the following notion.

Definition 3.11. A *left-reduced module diagonal* is a collection of grading preserving chain maps

$$\Gamma_n : M_{n,*}^{\text{red}} \rightarrow M_{n,*}^{\text{red}} \otimes M_{n,*}^{\text{red}}$$

which are compatible with splicing of trees.

Lemma 3.12. *The complex $M_{n,*}^{\text{red}}$ has homology \mathbb{F} in degree 0, and 0 in other degrees.*

Proof. We claim that the homology is generated by a tree

$$T_{\text{gen}} := \Psi_2 \circ_1 T_0,$$

where T_0 is a degree 0 tree with $n - 1$ inputs (all such trees are equal in $M_{n,*}^{\text{red}}$). To see this, we argue as follows. We define a map

$$H : M_{n,*}^{\text{red}} \rightarrow M_{n,*+1}^{\text{red}}$$

as follows. If T is a tree, we say the *right-most path* of T is the path from the right-most input of T to its root. We write $\text{RP}(T)$ for this path. Given a tree T , if T has more than one interior vertex along $\text{RP}(T)$, then we contract the top two interior vertices along this path to define a tree T' . We set $H(T) = T'$. Otherwise we define $H(T) = 0$.

We claim that if $N \gg 0$, the map $(\mathbb{I} + [\partial, H])^N$ is a projection onto the span of T_{gen} . Note that this clearly proves the claim. To see this, we argue as follows. Given a tree, we define the *right-path weight* to be

$$w_R(T) := \sum_{v \in \text{RP}(T)} (\text{val}(v) - 2).$$

We define the subspace $\mathcal{F}_j \subseteq M_{n,*}^{\text{red}}$ to be the span of trees T with $w_R(T) \leq j$.

It is straightforward to check that if $j > 1$, then $\mathbb{I} + [\partial, H]$ maps \mathcal{F}_j into \mathcal{F}_{j-1} . Furthermore, $\mathbb{I} + [\partial, H]$ is the identity on \mathcal{F}_1 . Therefore, for N large enough,

$(\mathbb{I} + [\partial, H])^N$ gives a homotopy equivalence onto \mathcal{F}_1 (viewing \mathcal{F}_1 as a subcomplex with vanishing differential). However, \mathcal{F}_1 is exactly the span of T_{gen} , so the proof is complete. \blacksquare

For our theory, we need to show that if ${}_{\mathcal{A}}M$ and ${}_{\mathcal{B}}N$ are A_∞ -modules over associative algebras \mathcal{A} and \mathcal{B} , then ${}_{\mathcal{A} \otimes \mathcal{B}}(M \otimes_\Gamma N)$ is independent of the choice of left-reduced module diagonal Γ . The argument is identical to the work in [28, Section 3], though we include it for completeness. Firstly, we introduce the following notion.

Definition 3.13. A *left-reduced module morphism tree* consists of a planar input tree T which has exactly one distinguished vertex along its right-most path, and whose vertices of valence 4 or more are all along the right-most path. We call the distinguished vertex the *morphism vertex*. This vertex has valence 2 or greater, whereas the other interior vertices have valence 3 or greater. We quotient this set of trees by the relation that $T \circ_i S = T \circ_i S'$ whenever T has n inputs, $i < n$, and S and S' have degree 0 and the same number of inputs.

The degree of a left-reduced module tree T is

$$\deg(T) = 1 + \sum_{v \in V_0(T)} (\text{val}(v) - 3).$$

We write $X_{n,*}^{\text{red}}$ for the vector space of left-reduced module morphism trees. We can naturally equip $X_{n,*}^{\text{red}}$ with a differential which sums over all ways of splitting one interior vertex into two interior vertices while leaving a valid left-reduced module morphism tree.

Lemma 3.14. *The chain complex $X_{n,*}^{\text{red}}$ has homology equal to \mathbb{F} in degree 0, and 0 in other degrees.*

The proof of Lemma 3.14 is very similar to the proof of Lemma 3.12, and we leave the details to the reader.

Definition 3.15. If Γ and Γ' are two left-reduced module diagonals, then a left-reduced module morphism diagonal from Γ to Γ' consists of a collection of chain maps

$$\hat{\Gamma}_n: X_{n,*}^{\text{red}} \rightarrow X_{n,*}^{\text{red}} \otimes X_{n,*}^{\text{red}},$$

such that $\hat{\Gamma}_1$ is an isomorphism, and $\hat{\Gamma}_n$ is compatible with splicing in the following sense. If $S \in X_{n,*}^{\text{red}}$ and $T \in M_{m,*}^{\text{red}}$, then

$$\hat{\Gamma}_{n+m-1}(S \circ_n T) = \hat{\Gamma}_n(S) \circ_n \Gamma_m(T) \quad \text{and} \quad \hat{\Gamma}_{n+m-1}(T \circ_m S) = \Gamma'_m(T) \circ_m \hat{\Gamma}_n(S).$$

Finally, if $T_0 \in M_{m,0}^{\text{red}}$ is a degree 0 tree and $S \in X_{n,*}^{\text{red}}$, then

$$\hat{\Gamma}_{n+m-1}(S \circ_i T_0) = \hat{\Gamma}_n(S) \circ_i (T_0 \otimes T_0).$$

Using Lemma 3.14, left-reduced module morphism diagonals may be constructed by a straightforward inductive argument.

Remark 3.16. If A and B are associative algebras and $f_*: {}_A M_1 \rightarrow {}_A M_2$ and $g_*: {}_B N_1 \rightarrow {}_B N_2$ are two A_∞ -module morphisms and $\hat{\Gamma}$ is a left-reduced module morphism diagonal, then we can define an A_∞ -module morphism

$$(f_* \otimes_{\hat{\Gamma}} g_*): {}_{A \otimes B}(M_1 \otimes_{\Gamma} N_1) \rightarrow {}_{A \otimes B}(M_2 \otimes_{\Gamma'} N_2).$$

Applying the above construction to the case that $M_1 = M_2 = M$ and $N_1 = N_2 = N$, we obtain a chain map I_* from $(M \otimes_{\Gamma} N)$ to $(M \otimes_{\Gamma'} N)$ which is a homotopy equivalence because $I_1 = \mathbb{I}_{M \otimes N}$ is an isomorphism. We may similarly use $\hat{\Gamma}$ to tensor DA bimodule morphisms together when the underlying algebras are associative algebras.

The motivation for left-reduced module diagonals is that it is straightforward to describe explicitly a diagonal over them. We now define two left-reduced module diagonals

$$L\Gamma_n, R\Gamma_n: M_n^{\text{red}} \rightarrow M_n^{\text{red}} \otimes M_n^{\text{red}}.$$

We only focus on $L\Gamma_n$, since $R\Gamma_n$ is obtained by switching the coordinates of the codomain.

To describe $L\Gamma_n$, we define the *right-join* of two input trees T_1 and T_2 , denoted $\text{RiJ}(T_1, T_2)$, to be the tree obtained by stacking the output of T_2 into the right-most input vertex of T_1 . (This is the mirror of Lipshitz, Ozsváth and Thurston’s *left-join*, as adapted to our handedness conventions on modules.) If T_1, \dots, T_n is a collection of trees, we write

$$\text{RiJ}(T_1, \dots, T_n) = \text{RiJ}(T_1, \dots, \text{RiJ}(T_{n-1}, T_n)),$$

If T_1, \dots, T_n are n trees, the *root-join* of T_1, \dots, T_n , denoted $\text{RoJ}(T_1, \dots, T_n)$, is obtained by adjoining the outputs of T_1, \dots, T_n into a single corolla Ψ_n . See [28, Definition 2.28].

For $n > 1$, we write $S_n \in H_0^{\text{cell}}(K_n)$ for the generator (thought of as a tree with n inputs and only valence 3 interior vertices).

The tree pairs appearing as summands of $L\Gamma_n(\Psi_n)$ consist of the sum of all pairs $T \otimes S$ where

$$T = \text{RiJ}(\Psi_{n_k+1}, \dots, \Psi_{n_1+1})$$

for some numbers $n_1, \dots, n_k > 0$ (where $k > 0$) satisfying

$$n_1 + \dots + n_k + 1 = n,$$

and

$$S = \text{RoJ}(S_{n_1}, \dots, S_{n_k}, \downarrow).$$

This root join is taken in such a way that \downarrow corresponds to the module input of the tree. See Figure 3.2 for an example.




Figure 3.2. An example of a tree contributing to $L\Gamma_9(\Psi_9)$.

Lemma 3.17. *The maps $R\Gamma_n$ and $L\Gamma_n$, defined above, are left-reduced module diagonals.*

The proof is straightforward and left to the reader.

Proof of Proposition 3.7. We observe that, immediately from the definitions, we have isomorphisms

$$\begin{aligned} \mathcal{A} \otimes_{\mathcal{C}} \mathcal{E}(M, \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}} \boxtimes_{\mathcal{B} \otimes \mathcal{C}} \mathcal{E}(N, \mathcal{B})^{\mathcal{B} \otimes \mathcal{D}} &\cong \mathcal{A} \otimes_{\mathcal{C}} (N \otimes_{R\Gamma} M)^{\mathcal{B} \otimes \mathcal{D}}, \\ \mathcal{A} \otimes_{\mathcal{C}} \mathcal{E}(N, \mathcal{A})^{\mathcal{A} \otimes \mathcal{D}} \boxtimes_{\mathcal{A} \otimes \mathcal{D}} \mathcal{E}(M, \mathcal{D})^{\mathcal{B} \otimes \mathcal{D}} &\cong \mathcal{A} \otimes_{\mathcal{C}} (N \otimes_{L\Gamma} M)^{\mathcal{B} \otimes \mathcal{D}}, \end{aligned}$$

where $R\Gamma$ and $L\Gamma$ denote the left-reduced module diagonals from Lemma 3.17. This implies the main claim, because any two choices of module diagonals give homotopy equivalent bimodules by the argument described in Remark 3.16. ■

We now prove a result which relates our extension of scalars operation with the external tensor product.

Proposition 3.18. *Let \mathcal{A} , \mathcal{B} and \mathcal{C} be associative algebras. If ${}_{\mathcal{A}}N^{\mathcal{B}}$ is a DA bimodule, then ${}_{\mathcal{A} \otimes_{\mathcal{C}}} \mathcal{E}(N, \mathcal{C})^{\mathcal{B} \otimes \mathcal{C}}$ is isomorphic to the external tensor product of ${}_{\mathcal{A}}N^{\mathcal{B}}$ and $\mathcal{C}[\mathbb{I}]^{\mathcal{C}}$, for any choice of cellular diagonal of the associahedron (or more generally, module diagonal or left-reduced module diagonal).*

The key to the above result is the following lemma.

Lemma 3.19. *Suppose that Γ is a cellular diagonal of the associahedron (or a module diagonal, or a left-reduced module diagonal). For each n , write $\Gamma_n(\Psi_n) = \Psi_n \otimes X + Y$, where X has degree 0 and $Y \in C_*^{\text{cell}}(K_n) \otimes C_{>0}^{\text{cell}}(K_n)$. Write $X = T_1 + \cdots + T_j$, where each T_i is a tree of degree 0. Then $j \equiv 1 \pmod{2}$.*

Proof. The argument is the same, regardless of what kind of diagonal Γ is, so we assume that Γ is an ordinary cellular diagonal of the associahedron. There is a quasi-isomorphism $\Pi: C_*^{\text{cell}}(K_n) \rightarrow \mathbb{F}$, which is 0 on $C_{>0}^{\text{cell}}(K_n)$. The map Π sends a tree in grading 0 to 1 in \mathbb{F} , and is 0 on all other trees. It is straightforward to check that Π is a chain map. By tensoring, we obtain a quasi-isomorphism $\mathbb{I} \otimes \Pi: C_*^{\text{cell}}(K_n) \otimes C_*^{\text{cell}}(K_n) \rightarrow C_*^{\text{cell}}(K_n)$.

The claim in the statement is equivalent to the claim that $(\mathbb{I} \otimes \Pi)(\Gamma_n(\Psi_n)) = \Psi_n$. This claim is equivalent to the claim that $(\mathbb{I} \otimes \Pi)(\Gamma_n(\Psi_n))$ is non-zero, since Ψ_n is the only tree in its grading. We argue by induction. The case that $n = 2$ is clear.

By induction and the fact that Γ respects splicing, we observe that

$$(\mathbb{I} \otimes \Pi)(\Gamma_n(\partial\Psi_n)) = \partial\Psi_n.$$

On the other hand,

$$(\mathbb{I} \otimes \Pi)(\Gamma_n(\partial\Psi_n)) = \partial((\mathbb{I} \otimes \Pi)(\Gamma_n(\Psi_n))).$$

For this to be $\partial\Psi_n$, we must have that $(\mathbb{I} \otimes \Pi)(\Gamma_n(\Psi_n))$ is non-zero, so the proof is complete. ■

Proof of Proposition 3.18. The proof is the same, regardless of what kind of diagonal Γ is, so we assume that Γ is an ordinary cellular diagonal of the associahedron. The map δ_{n+1}^1 on the external tensor product is determined by the structure map on N as well as the chain $\Gamma_{n+1}(\Psi_{n+1})$. If $T_i \otimes S_j$ is any summand of $\Gamma_{n+1}(\Psi_{n+1})$, where T_i and S_j are trees with $n + 1$ inputs of degree i and j , respectively, then the corresponding summand of δ_{n+1}^1 on $\mathcal{E}(N, \mathcal{C})$ will vanish unless $j = 0$ and $i = n - 1$ since \mathcal{C} is an associative algebra. In particular, the only contributions to δ_{n+1}^1 will be from tree pairs of the form $\Psi_{n+1} \otimes T_0$, where T_0 is a degree 0 tree. By Lemma 3.19 there are an odd number of such summands of $\Gamma_{n+1}(\Psi_{n+1})$. It follows that the external tensor product is exactly the extension of scalars construction we gave earlier. ■

3.6 Hypercubes as type-D modules

In this section, we present a helpful perspective of hypercubes and hyperboxes in terms of type-D modules over certain algebras which we call the *cube* and *box* algebras. Many natural hyperbox operations, such as compression, have natural interpretations as A_∞ -module operations with respect to these algebras.

We will not use the cube and box algebras explicitly in this memoir, but they are related to our reinterpretation of the Manolescu and Ozsváth’s link surgery formula in terms of type-D modules over the algebra \mathcal{L} .

We begin with the 1-dimensional cube algebra \mathcal{C}_1 . The algebra \mathcal{C}_1 is an algebra over the idempotent ring $\mathbf{I}_0 \oplus \mathbf{I}_1$, where $\mathbf{I}_\varepsilon \cong \mathbb{F}$. If $\varepsilon \in \{0, 1\}$, we write i_ε for the generator of \mathbf{I}_ε . We set $\mathbf{I}_\varepsilon \cdot \mathcal{C}_1 \cdot \mathbf{I}_\varepsilon = \mathbf{I}_\varepsilon$, and we set

$$\mathbf{I}_1 \cdot \mathcal{C}_1 \cdot \mathbf{I}_0 = \langle \theta_{0,1} \rangle,$$

i.e., we set $\mathbf{I}_1 \cdot \mathcal{C}_1 \cdot \mathbf{I}_0 \cong \mathbb{F}$, generated by an element $\theta_{0,1}$.

Next, we introduce the n -dimensional cube algebra \mathcal{C}_n . The definition is

$$\mathcal{C}_n = \mathcal{C}_1 \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathcal{C}_1 = \bigotimes_{\mathbb{F}}^n \mathcal{C}_1.$$

Note that this is an algebra over the idempotent ring $\mathbf{E}_n = \bigotimes_{\mathbb{F}}^n \mathbf{I}$. If $\varepsilon \in \{0, 1\}^n$, we will write \mathbf{E}_ε for $\mathbf{I}_{\varepsilon_1} \otimes \cdots \otimes \mathbf{I}_{\varepsilon_n}$. We will write i_ε for the generator of $\mathbf{E}_\varepsilon \cong \mathbb{F}$.

The importance of the cube algebra is the following.

Lemma 3.20. *The category of n -dimensional hypercubes of chain complexes is equivalent to the category of type-D modules over \mathcal{C}_n .*

Proof. We describe how to obtain a type-D module over \mathcal{C}_n from a hypercube of chain complexes. Suppose that $X = (X_\varepsilon, D_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ is a hypercube of chain complexes. We define a type-D module $\mathcal{X}^{\mathcal{C}_n}$ as follows. The underlying group of \mathcal{X} is the total space of X , i.e., $\mathcal{X} = \bigoplus_{\varepsilon \in \mathbb{E}_n} X_\varepsilon$. This is naturally a module over the idempotent ring \mathbf{E}_n .

We define $\delta^1: \mathcal{X} \rightarrow \mathcal{X} \otimes_{\mathbf{E}_n} \mathcal{C}_n$ as follows. If $\varepsilon \leq \varepsilon'$, there is an algebra element $\theta_{\varepsilon, \varepsilon'}$ which is the tensor product of $\theta_{0,1}$ over all i such that $\varepsilon_i < \varepsilon'_i$. If $\varepsilon = \varepsilon'$, we set $\theta_{\varepsilon, \varepsilon} = i_\varepsilon$. With this notation, we set

$$\delta^1(\mathbf{x}) = \sum_{\varepsilon \leq \varepsilon'} D_{\varepsilon, \varepsilon'}(\mathbf{x}) \otimes \theta_{\varepsilon, \varepsilon'}.$$

The type-D structure relations for (\mathcal{X}, δ^1) are straightforward to verify.

If $X = (X_\varepsilon, D_{\varepsilon, \varepsilon'})$ and $Y = (Y_\varepsilon, D'_{\varepsilon, \varepsilon'})$ are two n -dimensional hypercubes of chain complexes, it is straightforward to extend the above construction to give an isomorphism between morphism sets

$$F: \text{Mor}_{\mathbb{E}_n}(X, Y) \rightarrow \text{Mor}^{\mathcal{C}_n}(\mathcal{X}^{\mathcal{C}_n}, \mathcal{Y}^{\mathcal{C}_n}),$$

where $\text{Mor}_{\mathbb{E}_n}$ denotes the group of hypercube morphisms and $\text{Mor}^{\mathcal{C}_n}$ denotes the group of type-D morphisms. Clearly, F is an isomorphism and $F(f \circ g) = F(f) \circ F(g)$. Hence, we have constructed a functor from the category of n -dimensional hypercubes to the category of type-D modules over \mathcal{C}_n . An inverse functor is straightforward to describe. \blacksquare

The above construction extends naturally to the setting of hyperboxes, as we now describe. We begin with the 1-dimensional box algebra $\mathcal{B}_{(d)}$ of size $d > 0$. This is an algebra over the idempotent ring

$$\mathbf{I}_{(d)} := \mathbf{I}_0 \oplus \mathbf{I}_1 \oplus \cdots \oplus \mathbf{I}_d,$$

where each $\mathbf{I}_i \cong \mathbb{F}$. We set

$$\mathbf{I}_i \cdot \mathcal{B}_{(d)} \cdot \mathbf{I}_i = \mathbb{F}.$$

We set

$$\mathbf{I}_j \cdot \mathcal{B}_{(d)} \cdot \mathbf{I}_i = \begin{cases} \langle \theta_{i,i+1} \rangle & \text{if } j = i + 1, \\ \{0\} & \text{otherwise.} \end{cases}$$

Multiplication is given by setting $\theta_{j,j+1} \cdot \theta_{i,i+1} = 0$ for all i and j .

If $\mathbf{d} = (d_1, \dots, d_n)$ is a tuple of positive integers, we define the *hyperbox algebra* of size \mathbf{d} , denoted $\mathcal{B}_{\mathbf{d}}$ to be the tensor product

$$\mathcal{B}_{\mathbf{d}} := \mathcal{B}_{(d_1)} \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathcal{B}_{(d_n)}.$$

We have the following straightforward extension of Lemma 3.20.

Lemma 3.21. *The category of hyperboxes of size \mathbf{d} is equivalent to the category of type-D modules over the algebra $\mathcal{B}_{\mathbf{d}}$.*

Manolescu and Ozsváth’s compression operation also has a conceptually simple description in terms of the cube and box algebras. With the present perspective, compression takes the form of a DA bimodule

$$\mathcal{B}_{(d)}[\mathbb{C}_{(d)}]^{\mathcal{C}_1},$$

defined as follows.

If $d = 1$, we define $\mathbb{C}_{(d)}$ to be the identity bimodule. Henceforth, assume $d > 1$. As a group, $\mathbb{C}_{(d)}$ is a copy of $\mathbf{I} = \mathbf{I}_0 \oplus \mathbf{I}_1$. The right action of \mathbf{I} is the obvious one. The left action of the box idempotent ring $\mathbf{I}_{(d)}$ is given by having $\mathbf{I}_{(0)}$ and $\mathbf{I}_{(d)}$ act by \mathbf{I}_0 and \mathbf{I}_1 , respectively, and having \mathbf{I}_j act trivially for $j = 1, \dots, d - 1$.

The bimodule $[\mathbb{C}_{(d)}]$ has $\delta_j^1 = 0$ for all j except δ_2^1 and δ_{d+1}^1 . We set δ_2^1 to be the unique map satisfying $\delta_2^1(i_\varepsilon, i_\nu) = (i_\varepsilon \cdot i_\nu) \otimes 1$, where $\varepsilon \in \{0, \dots, d\}$ and $\nu \in \{0, 1\}$, and such that $\delta_2^1(\theta_{i,i+1}, i_\nu) = 0$ for all i . We define

$$\delta_{d+1}^1(\theta_{d-1,d}, \dots, \theta_{0,1}, i_0) = i_1 \otimes \theta.$$

Example 3.22. Consider a hyperbox of dimension 1,

$$X = (C_0 \xrightarrow{f_{1,0}} C_1 \xrightarrow{f_{2,1}} \cdots \xrightarrow{f_{d-1,d}} C_d).$$

The associated type-D module over $\mathcal{B}_{(d)}$ takes the following form:

$$\mathcal{X}^{\mathcal{B}_{(d)}} = (C_0 \xrightarrow{f_{0,1} \otimes \theta_{0,1}} C_1 \xrightarrow{f_{1,2} \otimes \theta_{1,2}} \cdots \xrightarrow{f_{d-1,d} \otimes \theta_{d-1,d}} C_d).$$

The underlying group of the box tensor product $\mathcal{X}^{\mathcal{B}_{(d)}} \boxtimes_{\mathcal{B}_{(d)}} [\mathbb{C}_{(d)}]^{\mathcal{C}_1}$ is $C_0 \oplus C_d$. One easily computes that the differential on this box tensor product is

$$\mathcal{X}^{\mathcal{B}_{(d)}} \boxtimes_{\mathcal{B}_{(d)}} [\mathbb{C}_{(d)}]^{\mathcal{C}_1} = (C_0 \xrightarrow{(f_{d-1,d} \circ \cdots \circ f_{0,1}) \otimes \theta} C_d).$$

The right-hand side is the compression of X , viewed as a type-D module over \mathcal{C}_1 , because it coincides with the function composition description of compression that is discussed in Section 2.3.

More generally, if $\mathbf{d} = (d_1, \dots, d_n)$ is an n -tuple of positive integers, we can define a bimodule $\mathcal{B}_{\mathbf{d}}[\mathcal{C}_{\mathbf{d}}]^{\mathcal{C}_n}$ as the external tensor product of the DA bimodules $\mathcal{B}_{(d_i)}[\mathcal{C}_{(d_i)}]^{\mathcal{C}_1}$ ranging over $i = 1, \dots, n$. Note that defining the external tensor product requires a choice of ordering of the axis directions, cf. Section 3.5.

Lemma 3.23. *If X is a hyperbox of size $\mathbf{d} = (d_1, \dots, d_n)$, and $\mathcal{X}^{\mathcal{B}_{\mathbf{d}}}$ is the corresponding type-D module, then the compression of X coincides with the hypercube corresponding to $\mathcal{X}^{\mathcal{B}_{\mathbf{d}}} \boxtimes \mathcal{B}_{\mathbf{d}}[\mathcal{C}_{\mathbf{d}}]^{\mathcal{C}_n}$.*

Proof. The bimodule $\mathcal{B}_{\mathbf{d}}[\mathcal{C}_{\mathbf{d}}]^{\mathcal{C}_n}$ requires a choice of module diagonal to construct. Alternatively, by the proof of Proposition 3.7, we can construct $\mathcal{B}_{\mathbf{d}}[\mathcal{C}_{\mathbf{d}}]^{\mathcal{C}_n}$ using the extension of scalars construction. Write $\mathbf{d}_i = (1, \dots, 1, d_i, \dots, d_n)$. We apply the extension of scalars construction to $\mathcal{B}_{d_i}[\mathcal{C}_{(d_i)}]^{\mathcal{C}_1}$ to obtain a bimodule

$$\mathcal{B}_{d_i}[\mathcal{C}_{\mathbf{d}_i, i}]^{\mathcal{B}_{d_i+1}}.$$

Using Proposition 3.7, we see that

$$\mathcal{B}_{\mathbf{d}}[\mathcal{C}_{\mathbf{d}}]^{\mathcal{C}_n} \simeq \mathcal{B}_{\mathbf{d}}[\mathcal{C}_{\mathbf{d}_1, 1}]^{\mathcal{B}_{d_2}} \boxtimes \dots \boxtimes \mathcal{B}_{\mathbf{d}_n}[\mathcal{C}_{\mathbf{d}_n, n}]^{\mathcal{C}_n}. \quad (3.5)$$

Tensoring with $\mathcal{B}_{d_i}[\mathcal{C}_{\mathbf{d}_i, i}]^{\mathcal{B}_{d_i+1}}$ has the same effect as applying the partial compression operation c_i from Section 2.3, so the main claim follows from equation (3.5). ■

Chapter 4

Homological perturbation theory

In this chapter, we describe several versions of the homological perturbation lemma. In this chapter and throughout the rest of the memoir, we assume that every ring and module has characteristic 2.

4.1 A_∞ -modules

Kadeishvili proved that A_∞ -algebra structures could be transferred along homotopy equivalences of the underlying chain complex, when coefficients were in a field [17]. A formula for the perturbed A_∞ -module structure in terms of trees may be found in work of Kontsevich and Soibelman [21, Section 6.4]. See [19, Section 3.3] or [50, Remark 1.15] for more historical context and development of ideas. The following is a statement of the homological perturbation lemma for A_∞ -modules.

Lemma 4.1. *Suppose that \mathcal{A} is an associative algebra over a ring \mathbf{k} , ${}_{\mathcal{A}}M$ is an A_∞ -module, (Z, m_1^Z) is a chain complex over \mathbf{k} , and that we have three maps of left \mathbf{k} -modules*

$$i: Z \rightarrow M, \quad \pi: M \rightarrow Z, \quad \text{and} \quad h: M \rightarrow M$$

satisfying the following:

- (1) i and π are chain maps.
- (2) $\pi \circ i = \text{id}_Z$.
- (3) $i \circ \pi = \text{id}_M + \partial_{\text{Mor}}(h)$.
- (4) $h \circ i = 0$.
- (5) $\pi \circ h = 0$.
- (6) $h \circ h = 0$.

In the above, $\partial_{\text{Mor}}(h)$ denotes the morphism differential for chain complexes, i.e., $\partial_{\text{Mor}}(h) = h \circ m_1^M + m_1^M \circ h$. Then there is an A_∞ -module structure on Z , extending m_1^Z , such that the maps π , i and h extend to A_∞ -module morphisms which satisfy all of the above relations as A_∞ -module maps.

One extremely useful aspect of the homological perturbation lemma is that all of the maps have a concrete description. The structure maps on Z are given by the diagrams shown in Figure 4.1. Therein, $m_{>1}$ denotes the A_∞ structure maps of M , with m_1 excluded.




Figure 4.1. The maps appearing in the homological perturbation lemma for A_∞ -modules.

4.2 DA bimodules

There is a homological perturbation lemma for DA bimodules, which we will use extensively. The statement and proof are similar to the statement for A_∞ -modules. If $({}_{\mathcal{A}}M^{\mathcal{B}}, \delta_j^1)$ is a DA bimodule, we will write $M^{\mathcal{B}}$ for the associated type-D module which has the same generators as ${}_{\mathcal{A}}M^{\mathcal{B}}$, and whose structure map is given by the map δ_1^1 from ${}_{\mathcal{A}}M^{\mathcal{B}}$.

Lemma 4.2. *Suppose that \mathcal{A} and \mathcal{B} are associative algebras, and that ${}_{\mathcal{A}}M^{\mathcal{B}}$ is a type-DA bimodule, $Z^{\mathcal{B}}$ is a type-D module over \mathcal{B} , and that we have three morphisms of type-D modules*

$$i^1: Z^{\mathcal{B}} \rightarrow M^{\mathcal{B}}, \quad \pi^1: M^{\mathcal{B}} \rightarrow Z^{\mathcal{B}}, \quad \text{and} \quad h^1: M^{\mathcal{B}} \rightarrow M^{\mathcal{B}}$$

satisfying the following:

- (1) i^1 and π^1 are homomorphisms of type-D modules (i.e., $\partial_{\text{Mor}}(i^1)$ and $\partial_{\text{Mor}}(\pi^1)$ vanish).
- (2) $\pi^1 \circ i^1 = \text{id}_Z$.
- (3) $i^1 \circ \pi^1 = \text{id}_M + \partial_{\text{Mor}}(h^1)$.
- (4) $h^1 \circ i^1 = 0$.
- (5) $\pi^1 \circ h^1 = 0$.
- (6) $h^1 \circ h^1 = 0$.

In the above, $\partial_{\text{Mor}}(h^1)$ denotes the morphism differential for type-D modules over \mathcal{B} ; see equation (3.2). Then there is an induced DA bimodule structure on ${}_{\mathcal{A}}Z^{\mathcal{B}}$, extending $Z^{\mathcal{B}}$, such that the maps π^1 , i^1 and h^1 extend to DA bimodule morphisms which satisfy all of the above relations as DA bimodule morphisms.

We refer the reader to [42, Lemma 2.12] for a proof. Note that the maps are similar to those appearing in Figure 4.1. The structure map δ_{j+1}^1 is indicated in Figure 4.2. The other maps i_j^1 , π_j^1 and h_j^1 are constructed by a similar modification of the other maps in Figure 4.1. Here, $\delta_{>1}^1$ denotes δ_j^1 for some $j > 1$. The final Π map means to multiply all of the inputs using several applications of μ_2 .

4.3 Hypercubes

In this section, we review a version of the homological perturbation lemma for hypercubes. See [14, Section 2.7] for more details. Lemma 4.3 is similar to other formulations of the homological perturbation lemma in the context of filtered chain complexes. See, e.g., [16]. A similar, though slightly less explicit, construction for transferring hypercube structures along homotopy equivalences is described by Liu [30, Section 5.6].




Figure 4.2. The structure relation δ_{j+1}^1 from the homological perturbation lemma for DA bimodules.

Lemma 4.3. *Suppose that $\mathcal{C} = (C_\varepsilon, D_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ is a hypercube of chain complexes, and $(Z_\varepsilon)_{\varepsilon \in \mathbb{E}_n}$ is a collection of chain complexes, indexed by $\varepsilon \in \mathbb{E}_n$. Furthermore, suppose that for each $\varepsilon \in \mathbb{E}_n$, we have chosen maps*

$$\pi_\varepsilon: C_\varepsilon \rightarrow Z_\varepsilon, \quad i_\varepsilon: Z_\varepsilon \rightarrow C_\varepsilon, \quad h_\varepsilon: C_\varepsilon \rightarrow C_\varepsilon,$$

satisfying

$$\begin{aligned} \pi_\varepsilon \circ i_\varepsilon &= \text{id}, & i_\varepsilon \circ \pi_\varepsilon &= \text{id} + \partial_{\text{Mor}}(h_\varepsilon), & h_\varepsilon \circ h_\varepsilon &= 0, \\ \pi_\varepsilon \circ h_\varepsilon &= 0, & \text{and } h_\varepsilon \circ i_\varepsilon &= 0, \end{aligned}$$

and such that π_ε and i_ε are chain maps. With the above data chosen, there are canonical hypercube structure maps $\delta_{\varepsilon, \varepsilon'}: Z_\varepsilon \rightarrow Z_{\varepsilon'}$ so that $\mathcal{Z} = (Z_\varepsilon, \delta_{\varepsilon, \varepsilon'})$ is a hypercube

of chain complexes, and also there are morphisms of hypercubes

$$\Pi: \mathcal{C} \rightarrow \mathcal{Z}, \quad I: \mathcal{Z} \rightarrow \mathcal{C}, \quad \text{and} \quad H: \mathcal{C} \rightarrow \mathcal{C}$$

such that I and Π are chain maps, and satisfy

$$\Pi \circ I = \text{id}, \quad I \circ \Pi = \text{id} + \partial_{\text{Mor}}(H), \quad H \circ H = 0, \quad \Pi \circ H = 0, \quad \text{and} \quad H \circ I = 0.$$

It is important to understand that the structure maps $\delta_{\varepsilon, \varepsilon'}$ and the morphisms Π and I are determined completely by an explicit formula. We now describe $\delta_{\varepsilon, \varepsilon'}$. Suppose that $\varepsilon < \varepsilon'$ are points in \mathbb{E}_n . The hypercube structure maps $\delta_{\varepsilon, \varepsilon'}$ are given by the following sum:

$$\delta_{\varepsilon, \varepsilon'} := \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} \pi_{\varepsilon'} \circ D_{\varepsilon_{j-1}, \varepsilon_j} \circ h_{\varepsilon_{j-1}} \circ D_{\varepsilon_{j-2}, \varepsilon_{j-1}} \circ \dots \circ h_{\varepsilon_2} \circ D_{\varepsilon_1, \varepsilon_2} \circ i_{\varepsilon}.$$

The component of the map I sending coordinate ε to coordinate ε' is given via the formula

$$I_{\varepsilon, \varepsilon'} := \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} h_{\varepsilon_j} \circ D_{\varepsilon_{j-1}, \varepsilon_j} \circ \dots \circ h_{\varepsilon_2} \circ D_{\varepsilon_1, \varepsilon_2} \circ i_{\varepsilon}.$$

Similarly, Π is given by the formula

$$\Pi_{\varepsilon, \varepsilon'} := \sum_{\varepsilon = \varepsilon_1 < \dots < \varepsilon_j = \varepsilon'} \pi_{\varepsilon'} \circ D_{\varepsilon_{j-1}, \varepsilon_j} \circ h_{\varepsilon_{j-1}} \circ D_{\varepsilon_{j-2}, \varepsilon_{j-1}} \circ \dots \circ D_{\varepsilon_1, \varepsilon_2} \circ h_{\varepsilon_1}.$$

4.4 Hypercubes of DA bimodules

There is a version of the homological perturbation lemma which naturally generalizes both the homological perturbation lemmas for DA bimodules and hypercubes of chain complexes.

Definition 4.4. A hypercube of DA bimodules consists of a DA bimodule (C, δ_{\ast}^1) , such that C decomposes as $C = \bigoplus_{\varepsilon \in \mathbb{E}_n} C_{\varepsilon}$, where each C_{ε} is itself a (\mathbf{j}, \mathbf{k}) -module. Furthermore, the structure map δ_{j+1}^1 decomposes as a sum of maps

$$\delta_{j+1, \varepsilon, \varepsilon'}^1: \mathcal{A}^{\otimes j} \otimes C_{\varepsilon} \rightarrow C_{\varepsilon'} \otimes \mathcal{B},$$

where $j \geq 0$, and $\varepsilon \leq \varepsilon'$.

The above definition is a special case of Lipshitz, Ozsváth and Thurston's notion of a filtered DA bimodule [24, Section 2.2].

Lemma 4.5. *Suppose that \mathcal{A} and \mathcal{B} are dg-algebras. Suppose also that ${}_{\mathcal{A}}\mathcal{C}^{\mathcal{B}} = ({}_{\mathcal{A}}C_{\varepsilon}^{\mathcal{B}}, D_{*,\varepsilon,\varepsilon'}^1)_{\varepsilon \in \mathbb{E}_n}$ is a hypercube of DA bimodules, and suppose that for each $\varepsilon \in \mathbb{E}_n$ we have chosen type-D modules $(Z_{\varepsilon}^{\mathcal{B}})_{\varepsilon \in \mathbb{E}_n}$, as well as morphisms of type-D modules*

$$\pi_{\varepsilon}^1: C_{\varepsilon}^{\mathcal{B}} \rightarrow Z_{\varepsilon}^{\mathcal{B}}, \quad i_{\varepsilon}^1: Z_{\varepsilon}^{\mathcal{B}} \rightarrow C_{\varepsilon}^{\mathcal{B}}, \quad \text{and} \quad h_{\varepsilon}^1: C_{\varepsilon}^{\mathcal{B}} \rightarrow C_{\varepsilon}^{\mathcal{B}}$$

satisfying

$$\begin{aligned} \pi_{\varepsilon}^1 \circ i_{\varepsilon}^1 &= \text{id}_{Z_{\varepsilon}^{\mathcal{B}}}, & i_{\varepsilon}^1 \circ \pi_{\varepsilon}^1 + \text{id}_{C_{\varepsilon}^{\mathcal{B}}} &= \partial_{\text{Mor}}(h_{\varepsilon}^1), & h_{\varepsilon}^1 \circ h_{\varepsilon}^1 &= 0, \\ h_{\varepsilon}^1 \circ i_{\varepsilon}^1 &= 0, & \text{and} & & \pi_{\varepsilon}^1 \circ h_{\varepsilon}^1 &= 0, \end{aligned}$$

and such that π_{ε}^1 and i_{ε}^1 are homomorphisms of type-D modules (i.e., cycles). Such data determines structure maps $\delta_{j+1,\varepsilon,\varepsilon'}^1: \mathcal{A}^{\otimes j} \otimes Z_{\varepsilon} \rightarrow Z_{\varepsilon'} \otimes \mathcal{B}$, which make ${}_{\mathcal{A}}\mathcal{Z}^{\mathcal{B}} = ({}_{\mathcal{A}}Z_{\varepsilon}^{\mathcal{B}}, \delta_{*,\varepsilon,\varepsilon'}^1)$ a hypercube of DA bimodules. Furthermore, there are morphisms of hypercubes of DA bimodules

$$\Pi_*^1: {}_{\mathcal{A}}\mathcal{C}^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}\mathcal{Z}^{\mathcal{B}}, \quad I_*^1: {}_{\mathcal{A}}\mathcal{Z}^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}\mathcal{C}^{\mathcal{B}}, \quad \text{and} \quad H_*^1: {}_{\mathcal{A}}\mathcal{C}^{\mathcal{B}} \rightarrow {}_{\mathcal{A}}\mathcal{C}^{\mathcal{B}},$$

which satisfy analogous relations to i_{ε}^1 , π_{ε}^1 and h_{ε}^1 .

We now describe the maps δ_*^1 , Π_*^1 , I_*^1 and H_*^1 . We focus on the structure map $\delta_{j+1,\varepsilon,\varepsilon'}^1$. This map is defined by modifying Figure 4.2 as follows. In place of each $\delta_{>1}^1$ labeled therein, we are allowed one of two maps:

- (1) A hypercube map $D_{j+1,\varepsilon,\varepsilon'}^1$ for \mathcal{C} , such that $\varepsilon < \varepsilon'$. Furthermore, we assume only that $j \geq 0$ (so that we allow the case that there are no algebra inputs here).
- (2) An internal structure map $D_{>1,\varepsilon,\varepsilon}^1$ of some ${}_{\mathcal{A}}C_{\varepsilon}^{\mathcal{B}}$ (so in this case we do not allow instances of the internal differential $D_{1,\varepsilon,\varepsilon}^1$).

The maps Π_*^1 , I_*^1 and H_*^1 are defined by small modifications, similar to Figure 4.1. We leave the proof of the above lemma to the reader, as it is a straightforward variation of Lemma 4.2.

Chapter 5

Linear topological spaces

In this chapter, we recall basic background on algebraic completions. See [1, Section 10] and [22, Chapter 2] for related background. This material is used to define our module categories in Chapter 7.

Definition 5.1. A *linear topological vector space* consists of a vector space \mathcal{X} equipped with a topology which admits a basis of open sets centered at 0 consisting of vector subspaces, such that furthermore the addition map $\mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$ is continuous.

If R is a ring, an *R -module with linear topology* is defined similarly, except that there is a basis of 0 consisting of R -submodules.

Vector spaces with linear topologies are usually specified by picking a decreasing filtration consisting of subspaces $X_i \subseteq \mathcal{X}$ indexed by $i \in I$, where I is some directed and partially ordered set. Such a filtration determines a topology, where we declare a basis to be the sets of the form $x + X_i$, where $i \in I$ and $x \in \mathcal{X}$.

Given a filtration $(X_i)_{i \in I}$ on a vector space \mathcal{X} , the *completion* is defined as the inverse limit

$$\mathfrak{X} := \varprojlim_{i \in I} \mathcal{X}/X_i.$$

Rather than working with inverse limits, it is sometimes easier to work with Cauchy sequences and nets. Since some of the spaces we are working with are not first countable, we work with Cauchy *nets* instead of the perhaps more familiar Cauchy *sequences*.

We recall that a *net* in a topological space \mathcal{X} is a map $\mathbf{x}: I \rightarrow \mathcal{X}$, where I is a directed set. We usually write $\mathbf{x} = (x_i)_{i \in I}$ for a net. When \mathcal{X} is a linear topological vector space, we say that a net $(x_i)_{i \in I}$ is *Cauchy* if for each open subspace $E \subseteq \mathcal{X}$, there is an $i_0 \in I$ such that if $i, j \geq i_0$, then $x_i - x_j \in E$.

If $(x_i)_{i \in I}$ and $(y_j)_{j \in J}$ are two nets in \mathcal{X} , their sum is the net $(x_i + y_j)_{(i,j) \in I \times J}$. Two Cauchy nets are *equivalent* if their difference converges to $0 \in \mathcal{X}$ as a net. We note that the collection of nets in \mathcal{X} forms a proper class, though the set of Cauchy equivalence classes is a set. The following is well known and is elementary to prove.

Lemma 5.2. *If \mathcal{X} is a linear topological space, then the completion \mathfrak{X} is canonically isomorphic to the vector space of equivalence classes of Cauchy nets in \mathcal{X} .*

5.1 Linear maps

If \mathcal{X} and \mathcal{Y} are linear topological vector spaces, then a continuous map $f: \mathcal{X} \rightarrow \mathcal{Y}$ naturally induces a map on completions. We recall the following basic principle, whose proof is elementary.

Lemma 5.3. *Suppose that \mathcal{X} and \mathcal{Y} are linear topological spaces. A linear map $f: \mathcal{X} \rightarrow \mathcal{Y}$ is continuous if and only if it is continuous at 0.*

Definition 5.4. We define the category of *linear topological vector spaces* over \mathbb{F} , denoted $\text{LTS}_{\mathbb{F}}$, as follows. The objects are linear topological vector spaces over \mathbb{F} . We define the set of morphisms from \mathcal{X} to \mathcal{Y} in this category to be the set of continuous linear maps from \mathcal{X} to \mathcal{Y} (i.e., continuous linear maps between the completed spaces).

Remark 5.5. By definition, \mathcal{X} and \mathfrak{X} are isomorphic in the category $\text{LTS}_{\mathbb{F}}$ for any linear topological \mathbb{F} -vector space \mathcal{X} .

5.2 Direct products

We now describe some helpful perspectives on the direct sum and product in the context of linear topological spaces. These are important for our work because of the role that direct products play in the surgery formula from [32].

Definition 5.6. If $(X_i)_{i \in I}$ is a family of linear topological spaces, the *direct product* of topological vector spaces $\mathcal{X} = \prod_{i \in I} X_i$ coincides with the direct product in the category of topological spaces. More explicitly, a subspace $E \subseteq \prod_{i \in I} X_i$ is open if and only if there is some finite set $S \subseteq I$ and a set of open subspaces $U_s \subseteq X_s$ ranging over $s \in S$, such that

$$\left(\prod_{i \in I \setminus S} X_i \right) \oplus \left(\prod_{s \in S} U_s \right) \subseteq E.$$

We declare the product topology on $\bigoplus_{i \in I} X_i$ to be the subspace topology induced by the inclusion $\bigoplus_{i \in I} X_i \subseteq \prod_{i \in I} X_i$.

Note that if $\mathcal{X} = \bigoplus_{i \in I} X_i$ is equipped with the product topology, then $\mathfrak{X} \cong \prod_{i \in I} \mathfrak{X}_i$.

In our work, we consider frequently the case where each X_i is equipped with the discrete topology. In this case, if $S \subseteq I$ is a finite set, we write

$$\mathcal{X}_{\text{co}(S)} = \bigoplus_{i \notin S} X_i$$

for the fundamental basis of opens.

Example 5.7. In this example, we illustrate the case that $I = \mathbb{Z}$ and $X_i = \mathbb{F}$. We consider $\mathcal{X} = \bigoplus_{i \in I} X_i$. Write x_i for the generator of X_i ; the basis of open subspaces as $\mathcal{X}_{\text{co}\{-n, \dots, n\}} = \text{Span}(x_i : |i| > n)$. We can identify $\mathcal{X}/\mathcal{X}_{\text{co}\{-n, \dots, n\}}$ with $\text{Span}(x_{-n}, \dots, x_n)$. The inverse limit of $\mathcal{X}/\mathcal{X}_{\text{co}\{-n, \dots, n\}}$ consists of all tuples $(a_i)_{i \in \mathbb{Z}}$ where $a_i \in X_i$. This is the direct product $\prod_{i \in \mathbb{Z}} X_i$.

The following example illustrates the meaning of a continuous map between direct products equipped with the product topologies.

Example 5.8. Suppose that $F: \prod_{i \in I} X_i \rightarrow \prod_{j \in J} Y_j$ is a linear map and each X_i and Y_j is equipped with the discrete topology. Write $F_{j,i}$ for $\Pi_j \circ F \circ I_i$, where Π_j and I_i are projections and inclusions. Then F is continuous with respect to the product topologies if and only if for each j , there are only finitely many i such that $F_{j,i} \neq 0$, and furthermore $F = \sum_{j,i} F_{j,i}$.

5.3 Linear compactness

We now recall the notion of linear compactness. This definition is important for the module categories in Chapter 7 since some of the operations we define (such as various versions of the external tensor product) require the underlying spaces of our modules to be linearly compact. All of the link surgery modules we construct are linearly compact.

Many basic statements about compact sets from point-set topology have analogs for linearly compact vector spaces, though linear compactness and point-set compactness are not equivalent in general. There are several equivalent definitions of linear compactness, though the following is the most useful for our purposes.

Definition 5.9. A linear topological space \mathcal{X} is *linearly compact* if each open subspace $U \subseteq \mathcal{X}$ has finite codimension, i.e., $\dim(\mathcal{X}/U) < \infty$.

Remark 5.10. (1) The notion of linear compactness may be found in the classical work of Lefschetz [22, Chapter 2]. Note that Lefschetz formulates linear compactness in terms of the finite intersection property for closed affine subspaces, though this turns out to be equivalent. Since we will not revisit the work of Lefschetz, we will not describe the equivalence.

(2) The direct product of linearly compact spaces is linearly compact. In particular, $\prod_{\mathbb{N}} \mathbb{F}$ is linearly compact.

5.4 Tensor products

Given linear topological vector spaces \mathcal{X} and \mathcal{Y} , there are many different ways to topologize the tensor product $\mathcal{X} \otimes \mathcal{Y}$. This phenomenon goes back to work of Grothendieck [8] in the setting of functional analysis and Banach spaces. Despite fundamental differences between the Banach spaces and linear topological vector spaces, similar phenomena occur in both settings.

In this memoir, we focus on three natural topologies on tensor products:

- (1) The *standard topology* $\mathcal{X} \otimes^! \mathcal{Y}$.
- (2) The **-topology* $\mathcal{X} \otimes^* \mathcal{Y}$.
- (3) The *chiral-topology* $\mathcal{X} \otimes^{\vec{\otimes}} \mathcal{Y}$.

The following inclusions are continuous:

$$\mathcal{X} \otimes^* \mathcal{Y} \subseteq \mathcal{X} \otimes^{\vec{\otimes}} \mathcal{Y} \subseteq \mathcal{X} \otimes^! \mathcal{Y}.$$

The topologies \otimes^* and $\otimes^{\vec{\otimes}}$ are due to Beilinson [2]. Additional exposition and analysis of these constructions is given by Positelski [46, Sections 12 and 13], wherein many basic properties, subtleties and counterexamples to natural conjectures are described. These topologies are defined as follows.

Definition 5.11. Suppose that \mathcal{X} and \mathcal{Y} are linear topological spaces.

- (1) A subspace $E \subseteq \mathcal{X} \otimes^! \mathcal{Y}$ is open if and only if the following holds: There are open subspaces $U \subseteq \mathcal{X}$ and $V \subseteq \mathcal{Y}$ such that $U \otimes \mathcal{Y} \subseteq E$ and $\mathcal{X} \otimes V \subseteq E$.
- (2) A subspace $E \subseteq \mathcal{X} \otimes^* \mathcal{Y}$ is open if and only if the following holds: There are open subspaces $U \subseteq \mathcal{X}$ and $V \subseteq \mathcal{Y}$ such that $U \otimes V \subseteq E$; and for each $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, there are open subspaces $U_y \subseteq \mathcal{X}$ and $V_x \subseteq \mathcal{Y}$ such that $x \otimes V_x \subseteq E$ and $U_y \otimes y \subseteq E$.
- (3) A subspace $E \subseteq \mathcal{X} \otimes^{\vec{\otimes}} \mathcal{Y}$ is open if and only if the following holds: There is an open subspace $U \subseteq \mathcal{X}$ so that $U \otimes \mathcal{Y} \subseteq E$; and for each $x \in \mathcal{X}$, there is an open $V_x \subseteq \mathcal{Y}$ such that $x \otimes V_x \subseteq E$.

The topologies \otimes^* and $\otimes^!$ are symmetric between \mathcal{X} and \mathcal{Y} , though $\otimes^{\vec{\otimes}}$ is asymmetric. The motivation for $\otimes^{\vec{\otimes}}$ is that it is a natural topology for an algebra whose topology is given by one-sided ideals. This perspective is of particular importance because our algebra \mathcal{K} is naturally topologized using a family of right ideals. See Lemma 7.7.

Remark 5.12. If \mathcal{X} and \mathcal{Y} are equipped with filtrations $(X_i)_{i \in \mathbb{N}}$ and $(Y_i)_{i \in \mathbb{N}}$, such that $X_0 = \mathcal{X}$ and $Y_0 = \mathcal{Y}$, then the topology on $\mathcal{X} \otimes^! \mathcal{Y}$ is equivalent to the one induced by the filtration $(\mathcal{X} \otimes \mathcal{Y})_n = \sum_{i+j=n} X_i \otimes Y_j$. This is because

$$X_{2n} \otimes \mathcal{Y} \oplus \mathcal{X} \otimes Y_{2n} \subseteq \sum_{i+j=2n} X_i \otimes Y_j \subseteq X_n \otimes \mathcal{Y} \oplus \mathcal{X} \otimes Y_n.$$

This is not in general true if we work with filtrations over \mathbb{Z} instead of \mathbb{N} , or if we relax the assumption that $X_0 = \mathcal{X}$ and $Y_0 = \mathcal{Y}$.

5.5 Associativity and commutativity relations

In this section, we recall some basic properties about the tensor products $\otimes^!$, $\vec{\otimes}$ and \otimes^* . We begin with associativity.

Lemma 5.13. *If \circ is one of $\{*, !, \rightarrow\}$, then the canonical map*

$$\mathcal{X} \otimes^\circ (\mathcal{Y} \otimes^\circ \mathcal{Z}) \cong (\mathcal{X} \otimes^\circ \mathcal{Y}) \otimes^\circ \mathcal{Z}$$

is a homeomorphism.

The reader may consult [46, Proposition 13.4] for details. In contrast, the different tensor product operations are not mutually associative. Nonetheless, there are some relations between different parenthesizations of the tensor products. One important and well-known example is the following (we give a proof for the benefit of the reader).

Lemma 5.14. *Suppose that $\mathcal{X}_1, \mathcal{X}_2, \mathcal{Y}_1$, and \mathcal{Y}_2 are linear topological vector spaces. The natural map*

$$(\mathcal{X}_1 \otimes^! \mathcal{Y}_1) \vec{\otimes} (\mathcal{X}_2 \otimes^! \mathcal{Y}_2) \rightarrow (\mathcal{X}_1 \vec{\otimes} \mathcal{X}_2) \otimes^! (\mathcal{Y}_1 \vec{\otimes} \mathcal{Y}_2)$$

is continuous.

Proof. A subspace $E \subseteq (\mathcal{X}_1 \vec{\otimes} \mathcal{X}_2) \otimes^! (\mathcal{Y}_1 \vec{\otimes} \mathcal{Y}_2)$ is open if

- (s) There are open $U_1 \subseteq \mathcal{X}_1$ and $V_1 \subseteq \mathcal{Y}_1$, and for each $x_1 \in \mathcal{X}_1$ and $y_1 \in \mathcal{Y}_1$ there are open $U_2^{x_1} \subseteq \mathcal{X}_2$ and $V_2^{y_1} \subseteq \mathcal{Y}_2$ so that $\mathcal{X}_1 \otimes \mathcal{X}_2 \otimes V_1 \otimes \mathcal{Y}_2$, $\mathcal{X}_1 \otimes \mathcal{X}_2 \otimes y_1 \otimes V_2^{y_1}$, $U_1 \otimes \mathcal{X}_2 \otimes \mathcal{Y}_1 \otimes \mathcal{Y}_2$, and $x_1 \otimes U_2^{x_1} \otimes \mathcal{Y}_1 \otimes \mathcal{Y}_2$ are contained in E .

A subspace $E \subseteq (\mathcal{X}_1 \otimes^! \mathcal{Y}_1) \vec{\otimes} (\mathcal{X}_2 \otimes^! \mathcal{Y}_2)$ is open if

- (t-1) There are open $U_1 \subseteq \mathcal{X}_1$ and $V_1 \otimes \mathcal{Y}_1$ so that $U_1 \otimes \mathcal{Y}_1 \otimes \mathcal{X}_2 \otimes \mathcal{Y}_2$, and $\mathcal{X}_1 \otimes V_1 \otimes \mathcal{X}_2 \otimes \mathcal{Y}_2$ are contained in E .
- (t-2) For each $x_1 \in \mathcal{X}_1$ and $y_1 \in \mathcal{Y}_1$, there is an open $U_2^{x_1, y_1} \subseteq \mathcal{X}_2$ so that $x_1 \otimes y_1 \otimes U_2^{x_1, y_1} \otimes \mathcal{Y}_2 \subseteq E$.
- (t-3) For each $x_1 \in \mathcal{X}_1$, $y_1 \in \mathcal{Y}_1$, $x_2 \in \mathcal{X}_2$, there is an open $V_2^{x_1, y_1, x_2} \subseteq \mathcal{Y}_2$ so that $x_1 \otimes y_1 \otimes x_2 \otimes V_2^{x_1, y_1, x_2} \subseteq E$.

Any subspace E satisfying (s) satisfies (t-1), (t-2) and (t-3), so the stated map is continuous. ■

Applying Lemma 5.14 repeatedly gives the following corollary.

Corollary 5.15. *Suppose that $\mathcal{X}_1, \dots, \mathcal{X}_n$ and $\mathcal{Y}_1, \dots, \mathcal{Y}_n$ are linear topological vector spaces.*

(1) *The map*

$$(\mathcal{X}_1 \otimes^! \mathcal{Y}_1) \vec{\otimes} \cdots \vec{\otimes} (\mathcal{X}_n \otimes^! \mathcal{Y}_n) \rightarrow (\mathcal{X}_1 \vec{\otimes} \cdots \vec{\otimes} \mathcal{X}_n) \otimes^! (\mathcal{Y}_1 \vec{\otimes} \cdots \vec{\otimes} \mathcal{Y}_n)$$

is continuous.

(2) *The map*

$$(\mathcal{X}_1 \otimes^! \cdots \otimes^! \mathcal{X}_n) \vec{\otimes} (\mathcal{Y}_1 \otimes^! \cdots \otimes^! \mathcal{Y}_n) \rightarrow (\mathcal{X}_1 \vec{\otimes} \mathcal{Y}_1) \otimes^! \cdots \otimes^! (\mathcal{X}_n \vec{\otimes} \mathcal{Y}_n)$$

is continuous.

The following is well known.

Lemma 5.16. *Suppose that \mathcal{X} and \mathcal{Y} are filtered spaces. If either*

(1) *\mathcal{X} is linearly compact, or*

(2) *\mathcal{Y} is discrete,*

then $\mathcal{X} \vec{\otimes} \mathcal{Y} \cong \mathcal{X} \otimes^! \mathcal{Y}$.

Proof. Consider the case that \mathcal{X} is linearly compact and suppose that $E \subseteq \mathcal{X} \vec{\otimes} \mathcal{Y}$ is an open subspace. By assumption, there is some open subspace $U \subseteq \mathcal{X}$ so that $U \otimes \mathcal{Y} \subseteq E$. Since \mathcal{X} is linearly compact, we may pick $x_1, \dots, x_n \in \mathcal{X}$ spanning a complementary vector space to U . By assumption, there are open subspaces $V_1, \dots, V_n \subseteq \mathcal{Y}$ so that $x_i \otimes V_i \subseteq E$. Therefore

$$U \otimes \mathcal{Y} + \mathcal{X} \otimes (V_1 \cap \cdots \cap V_n) \subseteq U \otimes \mathcal{Y} + x_1 \otimes V_1 + \cdots + x_n \otimes V_n \subseteq E$$

so the map $\mathcal{X} \otimes^! \mathcal{Y} \rightarrow \mathcal{X} \vec{\otimes} \mathcal{Y}$ is continuous. The map in the opposite direction is obviously continuous, so they are isomorphic.

The claim when \mathcal{Y} is discrete is similar. ■

Remark 5.17. If both \mathcal{X} and \mathcal{Y} are linearly compact, or both are discrete, then all three tensor products $\mathcal{X} \otimes^! \mathcal{Y}$, $\mathcal{X} \vec{\otimes} \mathcal{Y}$ and $\mathcal{X} \otimes^* \mathcal{Y}$ coincide. Similarly, if \mathcal{X} is linearly compact and discrete (i.e., finite-dimensional), then all three tensor products coincide.

Remark 5.18. Even if \mathcal{X} and \mathcal{Y} are first countable, it may not be the case that $\mathcal{X} \vec{\otimes} \mathcal{Y}$ is first countable. For example, if $\mathcal{X} = \bigoplus_{i \in \mathbb{Z}} \mathbb{F}$ is given the discrete topology and $\mathcal{Y} = \mathbb{F}[[U]]$ (equipped with the U -adic topology) then it is not hard to see that $\mathcal{X} \vec{\otimes} \mathcal{Y}$ is not first countable. See [46, Example 13.1 and Lemma 7.4].

Morphisms may also be tensored. We state the following well-known result.

Lemma 5.19. *All three tensor product operations $\otimes^!$, \otimes^* and $\bar{\otimes}$ are functorial with respect to morphisms, i.e., if f and g are continuous morphisms, then $f \otimes^\circ g$ is also continuous for $\circ \in \{!, *, \rightarrow\}$.*

See [46, Lemma 12.3] for a detailed exposition.

5.6 The tree topology

In this section, we introduce a topology on a tensor product of linear topological spaces indexed by vertices of certain directed trees. We use this notion when considering split Alexander modules in Section 7.4.

We say that a tree Γ is *strongly directed* if each edge is given an orientation, and each vertex has at most one outgoing edge. We write $v > v'$ if there is an oriented sequence of edges from v to v' . A strongly directed tree has a unique minimal vertex, which we refer to as the *root*.

If $(\mathcal{X}_v)_{v \in V(\Gamma)}$ is a collection of linear topological spaces, we will describe a linear topology on the tensor product

$$\mathcal{X}_\Gamma := \bigotimes_{v \in V(\Gamma)} \mathcal{X}_v.$$

The construction works more generally when the spaces \mathcal{X}_v are equipped with commuting actions of rings R_e for each edge e adjacent to v . (For our purposes, R_e will be an idempotent ring of a link algebra.)

We describe the topology after introducing some notation. We consider a set of symbols $\{p, \mathcal{O}, X\}$. We call a function

$$s: V(\Gamma) \rightarrow \{p, \mathcal{O}, X\}$$

an *admissible labeling* if it satisfies the following:

- (s-1) s is not constantly p .
- (s-2) If $s(v) = X$, then $s(v') = X$ for all $v' < v$.
- (s-3) If $s(v) = p$, then $s(v') = p$ for all $v' > v$.
- (s-4) If $s(v) = \mathcal{O}$, then $s(v') = p$ for all $v' > v$, and $s(v') = X$ for all $v' < v$.
- (s-5) If $s(v) = X$, then there is a vertex v' with an edge pointing to v such that $s(v') \in \{\mathcal{O}, X\}$.

See Figure 5.1 for an example of an admissible labeling. Note that the last axiom prohibits any maximal leaf from being labeled by X . Furthermore, an admissible labeling must label at least one vertex with \mathcal{O} .




Figure 5.1. An admissible labeling of a strongly directed tree.

Suppose that s is an admissible labeling of Γ . We define a *point-enhancement* \mathbf{p} of s to be a choice of an element $\mathbf{p}(v) \in \mathcal{X}_v$ for each $v \in s^{-1}(p)$. We say an *open-enhancement* \mathcal{U} of s is a choice of open subspaces $\mathcal{U}(v) \subseteq \mathcal{X}_v$ for each $v \in s^{-1}(\mathcal{O})$. Given point- and open-enhancements, \mathbf{p} and \mathcal{U} , of s , there is a subspace $\mathcal{W}(\mathcal{U}, \mathbf{p}) \subseteq \mathcal{X}_\Gamma$, obtained by tensoring the spans of $\mathbf{p}(v)$ from \mathbf{p} , with the open subspaces $\mathcal{U}(v)$, as well as \mathcal{X}_v ranging over $v \in s^{-1}(X)$.

We now define the *tree topology* on \mathcal{X}_Γ , when Γ is a strongly directed tree.


Definition 5.20. Suppose that Γ is a strongly directed tree. We say a subspace $E \subseteq \mathcal{X}_\Gamma$ is open if and only if for each admissible labeling s of Γ and each point-enhancement \mathbf{p} of s , there is an open-enhancement $\mathcal{U}_\mathbf{p}$ of s so that

$$\mathcal{W}(\mathcal{U}_\mathbf{p}, \mathbf{p}) \subseteq E.$$

Example 5.21. If Γ is a linearly ordered graph with vertices $v_n > v_{n-1} > \dots > v_1$, and \mathcal{X}_{v_i} are spaces, then

$$\mathcal{X}_\Gamma = \mathcal{X}_{v_n} \vec{\otimes} \dots \vec{\otimes} \mathcal{X}_{v_1}.$$

Remark 5.22. We warn the reader that there are some related (but subtly non-isomorphic) topologies that one can define on these tensor products. As an example, consider a graph Γ with a central root vertex v_0 , and two rays pointing into v_0 , labeled by vertices w_1, \dots, w_n and u_1, \dots, u_m . We designate v_0 is minimal, and w_n and u_m are maximal. See the following:



Note that the topology


$$\mathcal{Y} := ((\mathcal{X}_{w_1} \vec{\otimes} \dots \vec{\otimes} \mathcal{X}_{w_n}) \otimes^* (\mathcal{X}_{u_1} \vec{\otimes} \dots \vec{\otimes} \mathcal{X}_{u_m})) \vec{\otimes} \mathcal{X}_{v_0}$$

is similar to the tree topology \mathcal{X}_Γ , but is not isomorphic in general. The open subspaces of \mathcal{Y} may also be described in terms of admissible labelings of Γ , however, the description has some subtle differences. For example, given an admissible labeling s with point-enhancements \mathbf{p} , the open subspaces in an open-enhancement $\mathcal{U}_{\mathbf{p}}$ for \mathcal{X}_Γ may vary as we change \mathbf{p} . For \mathcal{Y} , the open set at a vertex labeled \mathcal{O} can only depend on the values of \mathbf{p} for vertices above that vertex. On the other hand, the map

$$\mathcal{X}_\Gamma \rightarrow \mathcal{Y}$$

is continuous. Compare Proposition 5.24 and Remark 5.26, below.

Remark 5.23. Suppose that Γ is a star with a minimal root v_0 , and n vertices w_1, \dots, w_n , each connected to v_0 (and no additional vertices):



If $\mathcal{X}_{v_0}, \mathcal{X}_{w_1}, \dots, \mathcal{X}_{w_n}$ are spaces indexed by the vertices, then

$$\mathcal{X}_\Gamma = (\mathcal{X}_{w_1} \otimes^* \dots \otimes^* \mathcal{X}_{w_n}) \bar{\otimes} \mathcal{X}_{v_0}.$$

Proposition 5.24. Suppose that Γ is a strongly directed tree and $\Gamma_0 \subseteq \Gamma$ is a connected subtree. Let $(\mathcal{X}_v)_{v \in V(\Gamma)}$ be a collection of spaces, indexed by vertices of Γ and let \mathcal{Y} be another linear topological space. Suppose that

$$f: \mathcal{X}_{\Gamma_0} \rightarrow \mathcal{Y}$$

is a continuous map. Let $\Gamma' = \Gamma/\Gamma_0$ be the strongly directed tree obtained by contracting all vertices of Γ_0 to a single vertex, denoted $[\Gamma_0]$. For $v \in V(\Gamma')$, write \mathcal{Y}_v for \mathcal{X}_v if $v \in V(\Gamma) \setminus V(\Gamma_0)$, and \mathcal{Y} if $v = [\Gamma_0]$. Then the map

$$f \otimes \text{id}: \mathcal{X}_\Gamma \rightarrow \mathcal{Y}_{\Gamma'}$$

is continuous.

See Figure 5.2 for a schematic.

Proof. We let $\mathcal{V}' \subseteq \mathcal{Y}_{\Gamma'}$ be open. Let s be an admissible labeling of Γ and \mathbf{p} a point-enhancement of s . Our goal is to construct an open-enhancement $\mathcal{U}_{\mathbf{p}}$ of s so that

$$(f \otimes \text{id})(\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p})) \subseteq \mathcal{V}'.$$

Our construction of $\mathcal{U}_{\mathbf{p}}$ varies depending on s .

We partition the admissible labelings s of Γ as follows:

- (1) (p -type) $s|_{\Gamma_0} \equiv p$.




Figure 5.2. The operation considered in Proposition 5.24. On the left is Γ and on the right Γ' . The regions enclosed by dashes indicate Γ_0 and $[\Gamma_0]$.

- (2) (\mathcal{O} -type) $s|_{\Gamma_0}$ is not constantly p , and all edges pointing into Γ_0 take value p .
- (3) (X -type) There is at least one edge pointing into Γ_0 which takes value \mathcal{O} or X .

If s is an admissible labeling, then we define a labeling s' of $V(\Gamma')$ as follows. For $v \neq [\Gamma_0]$, we can view v as a vertex of both Γ and Γ' . We set $s'(v) = s(v)$. If $v = [\Gamma_0]$, we define $s'(v)$ to be p if s is p -type, \mathcal{O} if s is \mathcal{O} -type, and X if s is X -type.

We make the following claims:

- (1) If s is an admissible labeling of Γ , then s' is an admissible labeling of Γ' .
- (2) If s is \mathcal{O} -type, then $s|_{\Gamma_0}$ is an admissible labeling.

We leave it to the reader to verify both of the above two claims, as they are straightforward.

We now describe $\mathcal{U}_{\mathbf{p}}$. The construction depends on whether s is p -, \mathcal{O} - or X -type. Suppose first that s is p -type. We let s' be the admissible labeling of Γ' constructed above, and we define a point-enhancement \mathbf{p}' of s' by using the points from \mathbf{p} for vertices in $V(\Gamma) \setminus V(\Gamma_0)$, and using $f(\mathbf{p}(v_1) \otimes \cdots \otimes \mathbf{p}(v_n))$ for the vertex $[\Gamma_0]$, where $V(\Gamma_0) = \{v_1, \dots, v_n\}$. Since \mathcal{V}' is open, there is an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{V}'. \quad (5.1)$$

We can use the same open subspaces from $\mathcal{U}_{\mathbf{p}'}$ to construct an open-enhancement $\mathcal{U}_{\mathbf{p}}$ of s . Equation (5.1) is equivalent to the statement that

$$(f \otimes \text{id})(\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p})) \subseteq \mathcal{V}'.$$

The construction of $\mathcal{U}_{\mathbf{p}}$ when s is X -type is similar. In this case, given a point-enhancement \mathbf{p} of s , we define a point-enhancement \mathbf{p}' of s' by using the points from \mathbf{p} for the vertices in $V(\Gamma') \setminus \{[\Gamma_0]\} = V(\Gamma) \setminus V(\Gamma_0)$. Since \mathcal{V}' is open, there is an

open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{V}'.$$

We construct an open-enhancement of s by using the subspaces from $\mathcal{U}_{\mathbf{p}'}$ for $v \notin V(\Gamma_0)$, and using the entire subspaces \mathcal{X}_v for $v \in V(\Gamma_0)$ with $s(v) = \mathcal{O}$. Since $s'([\Gamma_0]) = X$, we have that

$$(f \otimes \text{id})(\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p})) \subseteq \mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}')$$

which is contained in \mathcal{V}' by construction.

Finally, we consider the case that s is \mathcal{O} -type (or equivalently $s'([\Gamma_0]) = \mathcal{O}$). We let \mathbf{p} be a point-enhancement of s . We let \mathbf{p}' be the point-enhancement of s' obtained by restricting \mathbf{p} to the vertices of $V(\Gamma') \setminus \{[\Gamma_0]\}$. Since \mathcal{V}' is open, there is an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{V}'.$$

Let $U'_0 \subseteq \mathcal{Y}$ denote the open subspace $\mathcal{U}_{\mathbf{p}'}([\Gamma_0])$. Since f is continuous, there is an open subspace $U_0 \subseteq \mathcal{X}_{\Gamma_0}$ such that $f(U_0) \subseteq U'_0$.

Write s_0 for $s|_{V(\Gamma_0)}$. As described earlier, s_0 is an admissible labeling. Let \mathbf{p}_0 denote the restriction of \mathbf{p} to Γ_0 . Since U_0 is open, there is an open-enhancement \mathcal{U}_0 of s_0 so that

$$\mathcal{W}(\mathcal{U}_0, \mathbf{p}_0) \subseteq U_0.$$

We now define the open-enhancement $\mathcal{U}_{\mathbf{p}}$ of s as follows. For vertices $w \in V(\Gamma) \setminus V(\Gamma_0)$ with $s(w) = \mathcal{O}$, we set $\mathcal{U}_{\mathbf{p}}(w) = \mathcal{U}_{\mathbf{p}'}(w)$. For the vertices $w \in V(\Gamma_0)$ with $s(w) = \mathcal{O}$, we set $\mathcal{U}_{\mathbf{p}}(w) = \mathcal{U}_0(w)$. By construction

$$(f \otimes \text{id})(\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p})) \subseteq \mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{V}'.$$

It follows that $(f \otimes \text{id})^{-1}(\mathcal{V}')$ is an open subspace, so the proof is complete. ■

Remark 5.25. Note that we can extend the above result in several ways.

(1) If $(\mathcal{X}_v)_{v \in V(\Gamma)}$ is a family of linear topological spaces with $\Gamma_0 \subseteq \Gamma$, and if $f: \mathcal{X}_{\Gamma_0} \rightarrow \mathcal{Y}$ is a continuous map between completions, then there is a well-defined, continuous map

$$f \otimes \text{id}: \mathcal{X}_{\Gamma} \rightarrow \mathcal{Y}_{\Gamma'}.$$

To see this, note that it suffices to define a map from \mathcal{X}_{Γ} to $\mathcal{Y}_{\Gamma'}$. In this case, observe that if $x \in \mathcal{X}_{\Gamma_0}$, then by definition $f(x) \in \mathcal{Y}$ is the limit of a Cauchy net $\{y_{\beta}\}_{\beta \in B}$ where $y_{\beta} \in \mathcal{Y}$. If \mathbf{x} is in the tensor product of the \mathcal{X}_v for $v \notin V(\Gamma_0)$, then we claim that $\{y_{\beta} \otimes \mathbf{x}\}_{\beta \in B}$ is a Cauchy net in $\mathcal{Y}_{\Gamma'}$. It suffices to prove this when \mathbf{x} is an elementary tensor. Let s' be the admissible labeling on Γ/Γ_0 which assigns \mathcal{O} to $[\Gamma_0]$, assigns

X to each vertex below $[\Gamma_0]$, and assigns p to all other vertices. Let \mathbf{p}' be the point-enhancement of s' which assigns to each vertex w with $s(w') = p$ the corresponding factor from \mathbf{x} . If $\mathcal{V}' \subseteq \mathcal{Y}_{\Gamma'}$ is an open subspace, then by definition there is an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{V}'.$$

The open-enhancement $\mathcal{U}_{\mathbf{p}'}$ consists of a single open subspace $U_0 \subseteq \mathcal{Y}$, since $[\Gamma_0]$ is the only \mathcal{O} -labeled vertex. For all sufficiently large β , we have $y_{\beta'} - y_{\beta''} \in U_0$ whenever $\beta', \beta'' \geq \beta$. It follows that $y_{\beta'} \otimes \mathbf{x} - y_{\beta''} \otimes \mathbf{x} \in \mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}')$ so $\{y_{\beta} \otimes \mathbf{x}\}_{\beta \in B}$ is a Cauchy sequence in \mathcal{Y} . We therefore define $(f \otimes \text{id})(x \otimes \mathbf{x})$ to be the limit of the Cauchy net $\{y_{\beta} \otimes \mathbf{x}\}_{\beta \in B}$. A similar argument to the previous paragraph shows that this definition of $(f \otimes \text{id})$ is well defined (independent of the choice of net $\{y_{\beta}\}$). A straightforward extension of Proposition 5.24 shows that $(f \otimes \text{id}): \mathcal{X}_{\Gamma} \rightarrow \mathcal{Y}_{\Gamma'}$ is continuous.

(2) It follows from the previous argument that the completion \mathcal{X}_{Γ} is unchanged if we replace one \mathcal{X}_v in our family of spaces by its completion \mathcal{X}_v . To see this, let \mathcal{X}'_{Γ} denote the tree complex where one \mathcal{X}_v is replaced by \mathcal{X}_v . There is a natural map $\mathcal{X}_{\Gamma} \rightarrow \mathcal{X}'_{\Gamma}$. Applying the previous remark, there is a natural map in the opposite direction, because the identity map gives a continuous map from the completion of \mathcal{X}_v to the completion of \mathcal{X}_v . These maps are easily seen to be inverses.

We highlight an important application of Proposition 5.24.

Remark 5.26. If Γ is a strongly directed graph and $(\mathcal{X}_v)_{v \in V(\Gamma)}$ is a family of spaces, we can form another collection of spaces by “adding parentheses”, in the following sense. Let $\Gamma_0 \subseteq \Gamma$ be a tree, and consider the tree Γ/Γ_0 obtained by collapsing to vertices and edges of Γ_0 to a single vertex, denoted $[\Gamma_0]$. We can define a collection of spaces indexed by Γ/Γ_0 , by setting $\mathcal{Y}_v = \mathcal{X}_v$ if $v \notin V(\Gamma_0)$ and $\mathcal{Y}_{[\Gamma_0]} = \mathcal{X}_{\Gamma_0}$. Proposition 5.24 shows that the natural map

$$\mathbb{I}: \mathcal{X}_{\Gamma} \rightarrow \mathcal{Y}_{\Gamma/\Gamma_0}$$

is continuous.

In some cases, there is an isomorphism $\mathcal{X}_{\Gamma} \cong \mathcal{Y}_{\Gamma/\Gamma_0}$. We focus on the following simple case.

Lemma 5.27. *Let Γ be a strongly directed tree and $(\mathcal{X}_v)_{v \in V(\Gamma)}$ a family of spaces and let v_0 and v_1 be edges of Γ which are connected by an edge. Assume that $v_1 > v_0$ and that v_0 has no other incoming edges. See Figure 5.3. Let Γ_0 consist of a subgraph with vertices v_0 and v_1 and the edge connecting them. Let $\mathcal{Y}_{\Gamma/\Gamma_0}$ be the space above. If \mathcal{X}_{v_1} is linearly compact, then there is an isomorphism*

$$\mathcal{X}_{\Gamma} \cong \mathcal{Y}_{\Gamma/\Gamma_0}.$$




Figure 5.3. A tree from Lemma 5.27. Here, Γ_0 denotes the parenthesized subtree.

Proof. Remark 5.26 shows that the natural map

$$\mathbb{I}: \mathcal{X}_\Gamma \rightarrow \mathcal{Y}_{\Gamma/\Gamma_0}$$

is continuous. Therefore, it suffices to show that the natural map

$$\mathbb{I}: \mathcal{Y}_{\Gamma/\Gamma_0} \rightarrow \mathcal{X}_\Gamma$$

is continuous. Equivalently, it suffices to show that each open subspace of \mathcal{X}_Γ contains an open subspace of $\mathcal{Y}_{\Gamma/\Gamma_0}$. Let $\mathcal{V} \subseteq \mathcal{X}_\Gamma$ be an open subspace.

We consider admissible labelings s on Γ/Γ_0 . We wish to show that for each admissible labeling s' of Γ/Γ_0 and each point-enhancement \mathbf{p}' of s' , there is an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{V}.$$

We break the proof into three cases, depending on whether $s([\Gamma_0])$ is p , \emptyset or X .

If $s([\Gamma_0]) = X$, then we define a labeling s on Γ by setting $s(v_0) = s(v_1) = X$ and setting $s(w) = s'(w)$ for other vertices. If \mathbf{p}' is a point-enhancement of s' , then we can view \mathbf{p}' also as a point-enhancement of s , for which we write \mathbf{p} . Since \mathcal{V} is open, there is an open-enhancement $\mathcal{U}_{\mathbf{p}}$ so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p}) \subseteq \mathcal{V}.$$

Since v_0 and v_1 are labeled X by s' , we use the same open subspaces as in $\mathcal{U}_{\mathbf{p}}$ to form an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' . We observe that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') = \mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p}) \subseteq \mathcal{V}.$$

Next, we consider the case that $s'([\Gamma_0]) = p$. In this case, we form an admissible labeling s on Γ by setting $s(v_0) = s(v_1) = p$, and defining s and s' to coincide at other vertices. We let \mathbf{p}' be a point-enhancement of s' . There is a point assigned to $[\Gamma_0]$,

which consists of an element of the tensor product $\mathcal{X}_{v_0} \otimes \mathcal{X}_{v_1}$. We can write this as a finite sum of elementary tensors $\sum_{i=1}^n p_i \otimes q_i \in \mathcal{X}_{v_0} \otimes \mathcal{X}_{v_1}$. For each elementary tensor $p_i \otimes q_i$, we obtain a point-enhancement \mathbf{p}_i of Γ by using the elements of \mathbf{p}' for p -labeled vertices other than v_0 and v_1 , and by using p_i for v_0 and q_i for v_1 . Since \mathcal{V} is open, there is an open-enhancement $\mathcal{U}_{\mathbf{p}_i}$ of s so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}_i}, \mathbf{p}_i) \subseteq \mathcal{V}.$$

We construct an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' as follows. For a vertex $v \in V(\Gamma) \setminus \{v_0, v_1\}$ we set $\mathcal{U}_{\mathbf{p}'}(v) = \mathcal{U}_{\mathbf{p}_1}(v) \cap \cdots \cap \mathcal{U}_{\mathbf{p}_n}(v)$. We observe that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{W}(\mathcal{U}_{\mathbf{p}_1}, \mathbf{p}_1) + \cdots + \mathcal{W}(\mathcal{U}_{\mathbf{p}_n}, \mathbf{p}_n) \subseteq \mathcal{V}.$$

Finally, we consider the case that $s'([\Gamma_0]) = \mathcal{O}$. If $s'([\Gamma_0]) = \mathcal{O}$, then there are two labelings to consider on Γ . We denote these by $s_{p,\mathcal{O}}$ and $s_{\mathcal{O},X}$. The labeling $s_{p,\mathcal{O}}$ has $s_{p,\mathcal{O}}(v_1) = p$ and $s_{p,\mathcal{O}}(v_0) = \mathcal{O}$. The labeling $s_{\mathcal{O},X}$ has $s_{\mathcal{O},X}(v_1) = \mathcal{O}$ and $s_{\mathcal{O},X}(v_0) = X$. Both $s_{p,\mathcal{O}}$ and $s_{\mathcal{O},X}$ agree with s' for other vertices. It is straightforward to see that $s_{p,\mathcal{O}}$ and $s_{\mathcal{O},X}$ are admissible.

Since the s' and $s_{\mathcal{O},X}$ have the same p -labeled vertices, the collection \mathbf{p}' induces a point-enhancement \mathbf{p} (using the same points) of $s_{\mathcal{O},X}$. Since \mathcal{V} is open, there is an open-enhancement $\mathcal{U}_{\mathbf{p}}$ of $s_{\mathcal{O},X}$ so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p}) \subseteq \mathcal{V}.$$

Let U_{v_1} be the open set assigned to v_1 .

Since \mathcal{X}_{v_1} is linearly compact, by definition there are elements $p_1, \dots, p_n \in \mathcal{X}_{v_1}$ so that

$$\text{Span}_{\mathbb{F}}(p_1, \dots, p_n) + U_{v_1} = \mathcal{X}_{v_1}. \quad (5.2)$$

We now construct point-enhancements $\mathbf{p}_1, \dots, \mathbf{p}_n$ of $s_{p,\mathcal{O}}$. We use the points from \mathbf{p}' for the p -labeled vertices other than v_0 , and we use p_i for v_0 . By definition, there are open-enhancements $\mathcal{U}_{\mathbf{p}_1}, \dots, \mathcal{U}_{\mathbf{p}_n}$ so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}_i}, \mathbf{p}_i) \subseteq \mathcal{V}.$$

Write $U_{v_0,i} = \mathcal{U}_{\mathbf{p}_i}(v_0)$.

We now define an open-enhancement $\mathcal{U}_{\mathbf{p}'}$ of s' . For a vertex $v \neq [\Gamma_0]$ assigned \mathcal{O} , we use the open set

$$\mathcal{U}_{\mathbf{p}}(v) \cap \mathcal{U}_{\mathbf{p}_1}(v) \cap \cdots \cap \mathcal{U}_{\mathbf{p}_n}(v).$$

For the vertex $[\Gamma_0]$, we use the open set

$$U_{v_1} \otimes \mathcal{X}_{v_2} + \mathcal{X}_{v_1} \otimes (U_{v_0,1} \cap \cdots \cap U_{v_0,n}),$$

which is open in $\mathcal{X}_{\Gamma_0} \cong \mathcal{X}_{v_1} \overset{\vec{\otimes}}{\otimes} \mathcal{X}_{v_0}$ by Lemma 5.16 because \mathcal{X}_{v_1} is linearly compact.

We observe that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}'}, \mathbf{p}') \subseteq \mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p}) + \mathcal{W}(\mathcal{U}_{\mathbf{p}_1}, \mathbf{p}_1) + \cdots + \mathcal{W}(\mathcal{U}_{\mathbf{p}_n}, \mathbf{p}_n) \subseteq \mathcal{V}. \quad (5.3)$$

The above equation is proven as follows. For vertices $V(\Gamma/\Gamma_0)$ which are labeled p by s' , we use the same points in all of \mathbf{p}' , \mathbf{p} and \mathbf{p}_i . For the vertices of $V(\Gamma/\Gamma_0)$ which are labeled \mathcal{O} , other than $[\Gamma_0]$, we use the intersection of the opens from $\mathcal{U}_{\mathbf{p}}$, $\mathcal{U}_{\mathbf{p}_1}, \dots, \mathcal{U}_{\mathbf{p}_n}$, so these factors are contained in the corresponding factors of each summand in the middle of equation (5.3). Finally, for v_0 and v_1 , we observe that

$$\begin{aligned} U_{v_1} \otimes \mathcal{X}_{v_2} + \mathcal{X}_{v_1} \otimes (U_{v_{0,1}} \cap \cdots \cap U_{v_{0,n}}) \\ \subseteq U_{v_1} \otimes \mathcal{X}_{v_2} + p_1 \otimes U_{v_{0,1}} + \cdots + p_n \otimes U_{v_{0,n}} \end{aligned}$$

by equation (5.2). Together, the above implies equation (5.3). We conclude that \mathcal{V} is open in $\mathcal{X}_{\Gamma/\Gamma_0}$, completing the proof. \blacksquare

Chapter 6

The link surgery formula

In this chapter, we give some background on the link surgery formula of Manolescu and Ozsváth [32], and prove some important preliminary results.

6.1 Statement of the link surgery formula

In this section, we state Manolescu and Ozsváth's link surgery formula. We will describe more detail about the construction in Chapter 9.

If $L \subseteq S^3$ is a link, write K_1, \dots, K_ℓ for the components of L . Manolescu and Ozsváth define the affine lattice over \mathbb{Z}^ℓ

$$\mathbb{H}(L) := \prod_{i=1}^{\ell} \left(\frac{\text{lk}(K_i, L \setminus K_i)}{2} + \mathbb{Z} \right).$$

Whenever $M \subseteq L$, Manolescu and Ozsváth define a reduction map

$$\psi^M: \mathbb{H}(L) \rightarrow \mathbb{H}(L \setminus M)$$

as follows. If $K_{i_1} \cup \dots \cup K_{i_j} = L \setminus M$, then they set

$$\psi^M(\mathbf{s}) = (\psi_{i_1}^M(s_{i_1}), \dots, \psi_{i_j}^M(s_{i_j})),$$

where

$$\psi_i^M(s_i) = s_i - \frac{\text{lk}(K_i, M)}{2}.$$

Manolescu and Ozsváth also define a version of the map ψ^M for oriented sublinks of L . See [32, Section 3.7]. The above map ψ^M corresponds to the case that all components of M are oriented consistently with L .

Definition 6.1. Suppose $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathbf{p})$ is a Heegaard link diagram for (Y, L) with free base points \mathbf{p} , and link base points $\mathbf{w} \cup \mathbf{z}$. We say that \mathcal{H} is *link minimal* if each link component of L has exactly one base point from \mathbf{w} , and one base point from \mathbf{z} .

Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ be a link minimal Heegaard diagram for an oriented link $L \subseteq S^3$ with no free base points. If $M \subseteq L$, we write $\mathcal{H}^{L \setminus M}$ for the diagram obtained by deleting $z_i \in \mathbf{z}$ for each component $K_i \subseteq L$. Note that Manolescu also defines a reduction $\mathcal{H}^{L \setminus \vec{M}}$ when \vec{M} is equipped with an orientation. In this case, we delete z_i for each positively oriented component of M , and we delete w_i for each negatively

oriented component of M . The diagram $\mathcal{H}^{L \setminus M}$ is a link minimal diagram with $|M|$ free base points.

Recall that if $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is a link minimal Heegaard link diagram for L , then there is an ℓ -component Alexander grading

$$A = (A_1, \dots, A_\ell): \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \mathbb{H}(L).$$

Suppose \mathcal{H} is a link minimal Heegaard diagram for $L \subseteq S^3$ and \mathcal{H} has no free base points. If $\mathbf{s} \in \mathbb{H}(L)$, then Manolescu and Ozsváth define

$$\mathfrak{X}^-(\mathcal{H}^{L \setminus M}, \mathbf{s})$$

to be the free $\mathbb{F}[[U_1, \dots, U_\ell]]$ -module generated by monomials

$$U_1^{i_1} \dots U_n^{i_\ell} \cdot \mathbf{x},$$

where $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, which satisfy $i_j \geq 0$ for all j , and $A_j(\mathbf{x}) \leq s_j + i_j$ whenever $K_j \subseteq L \setminus M$. Manolescu and Ozsváth equip $\mathfrak{X}^-(\mathcal{H}^{L \setminus M}, \mathbf{s})$ with the differential which counts Maslov index 1 holomorphic disks, weighted by the product of $U_i^{nw_i(\phi)}$.

We will also occasionally write

$$A^-(\mathcal{H}^{L \setminus M}, \mathbf{s})$$

for the free $\mathbb{F}[U_1, \dots, U_n]$ -module which has the same generators as $\mathfrak{X}^-(\mathcal{H}^{L \setminus M}, \mathbf{s})$.

Let Λ denote an integral framing on L . The underlying group of the link surgery complex is defined to be

$$\mathcal{C}_\Lambda(L) = \bigoplus_{M \subseteq L} \prod_{\mathbf{s} \in \mathbb{H}(L)} \mathfrak{X}^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s})).$$

The differential $\mathcal{D}: \mathcal{C}_\Lambda(L) \rightarrow \mathcal{C}_\Lambda(L)$ decomposes as a sum

$$\mathcal{D} = \sum_{N \subseteq L} \sum_{\vec{N} \in \Omega(N)} \Phi^{\vec{N}},$$

where $\Omega(N)$ denotes the set of orientations on N .

Given an oriented sublink $\vec{N} \subseteq L$ (with orientation potentially different from L), Manolescu and Ozsváth define

$$\Lambda_{L, \vec{N}} \in \mathbb{Z}^\ell \cong H_1(S^3 \setminus L)$$

to be the sum of the longitudes of the negatively oriented components of \vec{N} .

Suppose $\vec{N} \subseteq L$ is an oriented sublink, and that $L_0 \subseteq L$ is a sublink which contains N . Write $L_1 = L_0 \setminus N$. In this case, the summand $\Phi^{\vec{N}}$ of \mathcal{D} maps $\mathfrak{X}^-(\mathcal{H}^{L_0}, \psi^{L \setminus L_0}(\mathbf{s}))$ to $\mathfrak{X}^-(\mathcal{H}^{L_1}, \psi^{L \setminus L_1}(\mathbf{s} + \Lambda_{L, \vec{N}}))$.

We write $\Phi_{L_0, L_1}^{\vec{N}}$ for the component of $\Phi^{\vec{N}}$ as follows:

$$\Phi_{L_0, L_1}^{\vec{N}}: \prod_{\mathbf{s} \in \mathbb{H}(L)} \mathfrak{X}^-(\mathcal{H}^{L_0}, \psi^{L \setminus L_0}(\mathbf{s})) \rightarrow \prod_{\mathbf{s} \in \mathbb{H}(L)} \mathfrak{X}^-(\mathcal{H}^{L_1}, \psi^{L \setminus L_1}(\mathbf{s} + \Lambda_{L, \vec{N}})).$$

Finally, we note that it is often convenient to view $\mathcal{C}_\Lambda(L, \mathcal{A})$ as a 1-dimensional mapping cone, so we establish some notation for this purpose. If $J \subseteq L$ is a component, and $\nu \in \{0, 1\}$, then we write $\mathcal{C}_\Lambda^\nu(L)$ for the subcube consisting of $\mathcal{C}_\varepsilon \subseteq \mathcal{C}_\Lambda(L, \mathcal{A})$ ranging over ε with J component equal to ν . We may decompose

$$\mathcal{C}_\Lambda(L) \cong \text{Cone}(\mathcal{C}_\Lambda^0(L) \xrightarrow{F^J + F^{-J}} \mathcal{C}_\Lambda^1(L)).$$

In the above, F^J (resp. F^{-J}) denotes the sum of $\Phi^{\vec{N}}$ for every oriented sublink \vec{N} of L such that $J \subseteq \vec{N}$ (resp. $-J \subseteq \vec{N}$).

We will review the construction of the hypercube maps $\Phi^{\vec{N}}$ in more detail in Chapter 9.

6.2 Gradings and algebraic reduction maps

In this section, we state several grading formulas which are useful when working with the link surgery formula.

Definition 6.2. Suppose that $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$ is a link minimal Heegaard diagram, where $\mathbf{w} \cup \mathbf{z}$ are link base points and \mathbf{p} are free base points. We call a subset $\mathcal{W} \subseteq \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$ a *complete set of base points* on \mathcal{H} if $\mathbf{p} \subseteq \mathcal{W}$ and \mathcal{W} contains exactly one base point from each link component.

Suppose \mathcal{H} is a diagram for a link L in Y , each of whose components is null-homologous, and $\mathfrak{s} \in \text{Spin}^c(Y)$ is a torsion. Given a complete set of base points \mathcal{W} on \mathcal{H} , there is a well-defined Maslov grading $\text{gr}_{\mathcal{W}}$ on $\mathcal{CF}\mathcal{L}(\mathcal{H}, \mathfrak{s})$. The grading $\text{gr}_{\mathcal{W}}$ is induced from the absolute grading of $\text{CF}^-(Y, \mathcal{W}, \mathfrak{s})$ by declaring the natural quotient map $\mathcal{CF}\mathcal{L}(\mathcal{H}, \mathfrak{s}) \rightarrow \text{CF}^-(Y, \mathcal{W}, \mathfrak{s})$ to be grading preserving. There is also an ℓ -component Alexander grading $A^{\mathcal{W}} = (A_1^{\mathcal{W}}, \dots, A_\ell^{\mathcal{W}})$. Note that a complete set of base points also determines an orientation on L , via the convention that L intersects the Heegaard surface Σ negatively at the link base points in \mathcal{W} . In particular, we may equivalently specify the Alexander grading A^L by picking an orientation on L .

The variable \mathcal{U}_i has $\text{gr}_{\mathcal{W}}$ -grading -2 if $w_i \in \mathcal{W}$, and 0 if $w_i \notin \mathcal{W}$. Similarly, \mathcal{V}_i has $\text{gr}_{\mathcal{W}}$ -grading -2 if $z_i \in \mathcal{W}$ and 0 otherwise. The variable U_i for a free base point $p_i \in \mathbf{p}$ has $\text{gr}_{\mathcal{W}}$ -grading -2 .

Remark 6.3. There are two natural conventions for defining the absolute grading for multi-pointed 3-manifolds. We recall that adding a base point has the effect on $\widehat{\text{HF}}$

of tensoring with a 2-dimensional vector space $V = \mathbb{F} \oplus \mathbb{F}$. See [41, Section 6]. In this memoir, we use the convention that the top degree element of V has absolute grading 0. This is different from the TQFT convention of [54], where V is supported in gradings $1/2$ and $-1/2$.

The following is implicit in the work of Manolescu and Ozsváth [32] (see, in particular, [32, Section 12]). In particular, the results are not new to our memoir, however the presentation is helpful.

Lemma 6.4. *Let $L = K_1 \cup \cdots \cup K_\ell$ denote an ℓ -component link in a 3-manifold Y , such that each component is null-homologous, and suppose Y is equipped with a torsion Spin^c structure \mathfrak{s} . Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathbf{p})$ be a link minimal Heegaard diagram, where \mathbf{w} and \mathbf{z} are link base points and \mathbf{p} are free base points. Let \mathcal{C} denote the link Floer complex $\mathcal{CFL}(\mathcal{H}, \mathfrak{s})$ over $\mathbb{F}[\mathcal{U}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell, U_1, \dots, U_n]$, where $\ell = |\mathbf{w}| = |\mathbf{z}|$ and $n = |\mathbf{p}|$.*

- (1) *If \mathcal{W} is a complete collection of base points on \mathcal{H} , and $z_i \in \mathbf{z}$ but $z_i \notin \mathcal{W}$, then the quotient map*

$$\mathcal{Q}_{\mathcal{V}_i}: \mathcal{C} \rightarrow \mathcal{C}/(\mathcal{V}_i - 1)$$

is grading preserving with respect to $\text{gr}_{\mathcal{W}}$. Similarly, the quotient map

$$\mathcal{Q}_{\mathcal{U}_i}: \mathcal{C} \rightarrow \mathcal{C}/(\mathcal{U}_i - 1),$$

is grading preserving if $w_i \in \mathbf{w}$ but $w_i \notin \mathcal{W}$.

- (2) *Suppose that K_j is a component of L , and let $\mathcal{W}_0 \cup \{w_j\}$ be a complete collection of base points, and write $\{w_j, z_j\} = K_j \cap (\mathbf{w} \cup \mathbf{z})$. Let L and K_j have orientations induced by $\mathcal{W}_0 \cup \{w_j\}$. Then*

$$\text{gr}_{\mathcal{W}_0 \cup \{w_j\}} - \text{gr}_{\mathcal{W}_0 \cup \{z_j\}} = 2A_j^{\mathcal{W}_0 \cup \{w_j\}} - \text{lk}(L \setminus K_j, K_j).$$

- (3) *Suppose that \mathcal{W} and \mathcal{W}' are two complete collections of base points which coincide at index j . Then*

$$A_j^{\mathcal{W}} = A_j^{\mathcal{W}'}$$

- (4) *Let $K_i \subseteq L$, and orient both K_i and L to intersect Σ negatively at \mathbf{w} , and suppose $i \neq j$. Give \mathcal{C} and $\mathcal{C}/(\mathcal{V}_i - 1)$ the Alexander grading A_j induced by the base points \mathbf{w} and $\mathbf{w} \setminus \{w_i\}$. For $j \neq i$,*

$$A_j(\mathcal{Q}_{\mathcal{V}_i}) = -\frac{1}{2}\text{lk}(K_i, K_j).$$

- (5) *Let $K_i \subseteq L$, and orient K_i and L to intersect Σ negatively at \mathbf{w} , and suppose $j \neq i$. Give \mathcal{C} and $\mathcal{C}/(\mathcal{U}_i - 1)$ the Alexander grading A_j induced by the base points \mathbf{w} and $\mathbf{w} \setminus \{w_i\}$, respectively. Then,*

$$A_j(\mathcal{Q}_{\mathcal{U}_i}) = \frac{1}{2}\text{lk}(K_i, K_j).$$

Proof. Following [54], the Alexander and Maslov gradings may be defined by representing $Y \setminus N(L)$ as Dehn surgery on a framed link J in the complement of an ℓ -component unlink in S^3 . We pick a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}_0, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ which represents the 2-handle cobordism $X(J)$. We pick a triangle class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ on \mathcal{T} . In [54, Section 5.5], the following formulas for the Maslov and Alexander gradings are proven:

$$\begin{aligned} \text{gr}_{\mathcal{W}}(\mathbf{z}) &= \text{gr}_{\mathcal{W}}(\mathbf{x}) + \text{gr}_{\mathcal{W}}(\mathbf{y}) - \frac{1}{2}(g(\Sigma) + |\mathcal{W}| - 1) - \mu(\psi) + 2n_{\mathcal{W}}(\psi) \\ &\quad + \frac{c_1(\mathfrak{s}_{\mathcal{W}}(\psi))^2 - 2\chi(X(J)) - 3\sigma(X(J))}{4}, \\ A_j^{\mathcal{W}}(\mathbf{z}) &= A_j^{\mathcal{W}}(\mathbf{x}) + A_j^{\mathcal{W}}(\mathbf{y}) + (n_{\mathcal{W}} - n_{(\mathbf{w} \cup \mathbf{z} \cup \mathbf{p}) \setminus \mathcal{W}})(\psi) \\ &\quad + \frac{\langle c_1(\mathfrak{s}_{\mathcal{W}}(\psi)), \widehat{S}_j \rangle - [\widehat{S}] \cdot [\widehat{S}_j]}{2}. \end{aligned} \tag{6.1}$$

In the above, if W is obtained by attaching 2-handles to $S^3 \times [0, 1]$ in the complement of an unlink with components $u_1 \cup \cdots \cup u_\ell$, then S_j is the surface $u_j \times [0, 1]$ and $S = S_1 + \cdots + S_\ell$. The classes \widehat{S} and \widehat{S}_j are obtained by capping these surfaces with Seifert surfaces in either end.

With the above in place, claim (1) follows because the formula for $\text{gr}_{\mathcal{W}}$ does not change if we delete a base point which is not in \mathcal{W} .

Consider now claim (2). An identical argument to [54, Lemma 3.8] implies that if ψ is a homology class of triangles on the surgery diagram for the above cobordism, then

$$\mathfrak{s}_{\mathbf{w}_0 \cup \{w_j\}}(\psi) - \mathfrak{s}_{\mathbf{w}_0 \cup \{z_j\}}(\psi) = \text{PD}[S_j]. \tag{6.2}$$

We may use the formula in equation (6.1) to evaluate $\text{gr}_{\mathbf{w}_0 \cup \{w_j\}} - \text{gr}_{\mathbf{w}_0 \cup \{z_j\}}$. Firstly, we recall that by straightforward computation $\widehat{\text{HF}}\text{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}_0, \mathbf{w}, \mathbf{z}, \mathbf{p})$ is isomorphic to $(\mathbb{F} \oplus \mathbb{F})^{\otimes (|\mathbf{w}| + |\mathbf{p}| - 1)}$. This computation can be obtained by using topological invariance of the Heegaard Floer group, and picking a convenient genus 0 Heegaard diagram, so that each w_i is immediately adjacent to z_i , and so that the attaching curves come in isotopic pairs. Similarly, $\widehat{\text{HF}}\text{L}(\boldsymbol{\beta}_0, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathbf{p})$ is isomorphic to $(\mathbb{F} \oplus \mathbb{F})^{\otimes (g(\Sigma) - |J| + |\mathbf{w}| + |\mathbf{p}| - 1)}$. Furthermore, for both groups, the non-zero elements of homology decompose into homogeneously graded summands which have the property that

$$\text{gr}_{\mathcal{W}}(\mathbf{x}) = \text{gr}_{\mathcal{W}'}(\mathbf{x}) \quad \text{and} \quad A^{\mathcal{W}}(\mathbf{x}) = A^{\mathcal{W}'}(\mathbf{x}) = 0 \tag{6.3}$$

for all complete collections $\mathcal{W}, \mathcal{W}' \subseteq \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$. Therefore, in our grading formulas, we may assume that the intersection points \mathbf{x} and \mathbf{y} satisfy equation (6.3). Using equation (6.2), we therefore obtain

$$\begin{aligned} &\text{gr}_{\mathbf{w}_0 \cup \{w_j\}}(\mathbf{z}) - \text{gr}_{\mathbf{w}_0 \cup \{z_j\}}(\mathbf{z}) \\ &= \frac{c_1^2(\mathfrak{s}_{\mathbf{w}_0 \cup \{w_j\}}(\psi)) - (c_1(\mathfrak{s}_{\mathbf{w}_0 \cup \{w_j\}}(\psi)) - 2\text{PD}[S_j])^2}{4} - 2(n_{z_j} - n_{w_j})(\psi). \end{aligned}$$

We may rearrange the above to see that

$$\begin{aligned}
 & \text{gr}_{\mathcal{W}_0 \cup \{w_j\}}(\mathbf{z}) - \text{gr}_{\mathcal{W}_0 \cup \{z_j\}}(\mathbf{z}) \\
 &= \langle c_1(\mathfrak{s}_{\mathcal{W}_0 \cup \{w_j\}}(\psi)), \widehat{S}_j \rangle - [\widehat{S}_j] \cdot [\widehat{S}_j] - 2(n_{z_j} - n_{w_j})(\psi) \\
 &= \langle c_1(\mathfrak{s}_{\mathcal{W}_0 \cup \{w_j\}}(\psi)), \widehat{S}_j \rangle - [\widehat{S}] \cdot [\widehat{S}_j] - 2(n_{z_j} - n_{w_j})(\psi) + [\widehat{S} \setminus \widehat{S}_j] \cdot [\widehat{S}] \\
 &= 2A_j^{\mathcal{W}_0 \cup \{w_j\}}(\mathbf{z}) + [\widehat{S} \setminus \widehat{S}_j] \cdot [\widehat{S}_j] \\
 &= 2A_j^{\mathcal{W}_0 \cup \{w_j\}}(\mathbf{z}) - \text{lk}(L \setminus K_j, K_j).
 \end{aligned}$$

In the last line, we used the equality $[\widehat{S} \setminus \widehat{S}_j] \cdot [\widehat{S}_j] = -\text{lk}(L \setminus K_j, K_j)$, which follows since we capped with negative Seifert surfaces of L (and since we are using the outward normal first convention for boundary orientations).

We now consider claim (3). We consider the case when \mathcal{W} and \mathcal{W}' differ only at a single index $i \neq j$. Assume that \mathcal{W} contains w_i and that \mathcal{W}' contains z_i . Write L' for L with the orientation of K_i reversed. Computing directly from the definition, shows that

$$2(A_j^{\mathcal{W}} - A_j^{\mathcal{W}'}) = \langle c_1(\mathfrak{s}_{\mathcal{W}}(\psi)) - c_1(\mathfrak{s}_{\mathcal{W}'}(\psi)), \widehat{S}_j \rangle - \widehat{S} \cdot \widehat{S}_j + \widehat{S}' \cdot \widehat{S}_j.$$

Here, $\widehat{S}' = \widehat{S} - 2\widehat{S}_i$. As before, $\mathfrak{s}_{\mathcal{W}} - \mathfrak{s}_{\mathcal{W}'} = \text{PD}[S_i]$. Hence, the above equation gives

$$A_j^{\mathcal{W}} - A_j^{\mathcal{W}'} = 0.$$

Consider claim (4) now. For this claim, we apply claim (2), with $\mathcal{W}_0 = \mathcal{W} \setminus \{w_j\}$. Since $z_i \notin \mathcal{W}_0 \cup \{w_j\}$ and $z_i \notin \mathcal{W}_0 \cup \{z_j\}$, by part (1) we know that $\mathcal{Q}_{\mathcal{V}_i}$ is homogeneously graded with respect to $\text{gr}_{\mathcal{W}_0 \cup \{w_j\}}$ and $\text{gr}_{\mathcal{W}_0 \cup \{z_j\}}$. Using claim (2), we obtain

$$2\Delta(A_j) = \text{lk}(L \setminus (K_i \cup K_j), K_j) - \text{lk}(L \setminus K_j, K_j) = -\text{lk}(K_i, K_j).$$

We now consider claim (5). By claim (3), it is sufficient to compute the grading change of $A_j^{\mathcal{W}'}$ where $\mathcal{W}' = (\mathcal{W} \setminus \{w_i\}) \cup \mathbf{p} \cup \{z_i\}$. We use claim (4), but note that with respect to \mathcal{W}' , the orientation of K_i is reversed, while the orientation of K_j is unchanged. ■

6.3 On the algebra of the link surgery formula

In this section, we reformulate the algebra of the link surgery formula slightly. We begin with a convenient description of the underlying group.

Lemma 6.5. *Suppose that \mathcal{H} is a link minimal Heegaard diagram for a link L in S^3 , which has no free base points. Let M be a sublink of L . Write $S_M \subseteq \mathbb{F}[\mathcal{U}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$ for the multiplicatively closed subset generated by \mathcal{V}_i for i such that $K_i \subseteq M$. Then there is an $\mathbb{F}[U_1, \dots, U_\ell]$ -equivariant chain isomorphism*

$$\bigoplus_{\mathbf{s} \in \mathbb{H}(L)} A^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s})) \cong S_M^{-1} \cdot \mathcal{CF}\mathcal{L}(\mathcal{H}),$$

where we view U_i as acting by $\mathcal{U}_i \mathcal{V}_i$ on the right-hand side. Furthermore, if $\mathbf{s} \in \mathbb{H}(L)$, this isomorphism intertwines the summand $A^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s}))$ with the subspace of $S_M^{-1} \cdot \mathcal{CF}\mathcal{L}(\mathcal{H})$ in an Alexander multi-grading \mathbf{s} .

Proof. The module $S_M^{-1} \cdot \mathcal{CF}\mathcal{L}(\mathcal{H})$ decomposes over Alexander gradings $\mathbf{s} \in \mathbb{H}(L)$. Hence, it is sufficient to identify the subspace in an Alexander grading \mathbf{s} with

$$A^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s}))$$

in a way which commutes with the action of $\mathbb{F}[U_1, \dots, U_\ell]$. We recall that $A^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s}))$ is generated by monomials $U_1^{i_1} \cdots U_\ell^{i_\ell} \cdot \mathbf{x}$ such that $i_j \geq 0$ for all j , and

$$A_j^{L \setminus M}(\mathbf{x}) - i_j \leq \psi_j^M(s_j)$$

for each j such that $K_j \not\subseteq M$. Here, we are writing $\mathbf{s} = (s_1, \dots, s_\ell)$. Note that by definition of $\psi_j^M(s_j)$ (see [32, Section 3.7]), we have

$$A_j^L(\mathbf{x}) - A_j^{L \setminus M}(\mathbf{x}) = s_j - \psi_j^M(s_j).$$

Compare Lemma 6.4. Hence we may think of $A^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s}))$ as being generated by monomials $U_1^{i_1} \cdots U_\ell^{i_\ell} \cdot \mathbf{x}$ such that $i_j \geq 0$ for all j and such that

$$A_j^L(\mathbf{x}) - i_j \leq s_j$$

whenever $K_j \not\subseteq M$. We may define a map

$$A^-(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s})) \rightarrow S_M^{-1} \cdot \mathcal{CF}\mathcal{L}(\mathcal{H})$$

via the formula

$$U_1^{i_1} \cdots U_\ell^{i_\ell} \cdot \mathbf{x} \mapsto \mathcal{U}_1^{i_1} \cdots \mathcal{U}_\ell^{i_\ell} \mathcal{V}_1^{s_1 - A_1^L(\mathbf{x}) + i_1} \cdots \mathcal{V}_\ell^{s_\ell - A_\ell^L(\mathbf{x}) + i_\ell} \cdot \mathbf{x}. \quad \blacksquare$$

Remark 6.6. In their construction of the link surgery formula, Manolescu and Ozsváth also define inclusion maps $\mathcal{I}_s^{\vec{N}}$ between the complexes $\mathfrak{X}^-(\mathcal{H}, \mathbf{s})$ and some generalizations of these complexes. See [32, Section 3.8]. We note that the map $\mathcal{I}_s^{\vec{N}}$ can be identified with the map for localizing at the variables \mathcal{U}_i , for i such that $+K_i \subseteq \vec{N}$, and localizing at variables \mathcal{V}_i , for i such that $-K_i \subseteq \vec{N}$.

We now discuss some algebraic properties of the maps in the link surgery hypercube. Manolescu and Ozsváth prove that the hypercube maps in their surgery formula $\mathcal{C}_\Lambda(L)$ are $\mathbb{F}[U_1, \dots, U_\ell]$ -equivariant. By the previous lemma, there is a more refined action of $\mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$ on the complex $\mathcal{C}_\Lambda(L)$, though the hypercube differential does not commute with this action. We now describe how the hypercube maps interact with the action of this larger ring.

Firstly, we define a ring homomorphism

$$\phi^\tau: \mathbb{F}[U, \mathcal{V}] \rightarrow \mathbb{F}[U, \mathcal{V}, \mathcal{V}^{-1}]$$

via the formulas

$$\phi^\tau(U) = \mathcal{V}^{-1} \quad \text{and} \quad \phi^\tau(\mathcal{V}) = U\mathcal{V}^2.$$

We write $\phi^\sigma: \mathbb{F}[U, \mathcal{V}] \rightarrow \mathbb{F}[U, \mathcal{V}, \mathcal{V}^{-1}]$ for the canonical inclusion.

We view $S_M^{-1} \cdot \mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$ as the tensor product over \mathbb{F} of ℓ different rings, which are each isomorphic to $\mathbb{F}[U, \mathcal{V}]$ or $\mathbb{F}[U, \mathcal{V}, \mathcal{V}^{-1}]$.

If $N \subseteq L \setminus M$, then we define the ring homomorphism

$$\phi^{\vec{N}}: S_M^{-1} \cdot \mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell] \rightarrow S_{M \cup N}^{-1} \cdot \mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$$

by tensoring the map ϕ^τ for each component $K_i \subseteq \vec{N}$ which is oriented oppositely to L , tensoring ϕ^σ for each component $K_i \subseteq \vec{N}$ which is oriented consistently with L , and tensoring the identity map for the remaining components of L .

Lemma 6.7. *Suppose that $M \subseteq L$ and $N \subseteq L \setminus M$. Let \vec{N} denote N with a choice of orientation. If $a \in S_M^{-1} \cdot \mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$, and $\mathbf{x} \in S_M^{-1} \cdot \mathcal{CFL}(\mathcal{H})$, then*

$$\Phi^{\vec{N}}(a \cdot \mathbf{x}) = \phi^{\vec{N}}(a) \cdot \Phi^{\vec{N}}(\mathbf{x}).$$

Proof. We focus on the length 1 maps in the surgery hypercube. The maps of higher length are analyzed using the same argument. In this case, write $N = K_i$, where K_i is a component of L . To simplify the notation further, we focus on the case that $M = \emptyset$. Let K_i , (resp. $-K_i$), denote K_i equipped with the orientation which coincides with (resp. is opposite to) the orientation from L . The map $\Phi_{L, L \setminus K_i}^{K_i}$ is the canonical inclusion map

$$\mathcal{CFL}(\mathcal{H}) \hookrightarrow \mathcal{V}_i^{-1} \cdot \mathcal{CFL}(\mathcal{H}),$$

which is clearly $\mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$ -equivariant. Note that ϕ^{σ_i} is just the canonical inclusion of $\mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$ into $\mathcal{V}_i^{-1} \cdot \mathbb{F}[U_1, \dots, U_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$.

The map $\Phi_{L, L \setminus K_i}^{-K_i}$ is more complicated, as we now describe. Similarly to Lemma 6.5, there are canonical isomorphisms

$$\begin{aligned} \theta_i: U_i^{-1} \cdot \mathcal{CFL}(\mathcal{H}) &\rightarrow \mathcal{CFL}(\mathcal{H}) / (U_i - 1) \otimes \mathbb{F}[\mathbb{H}_i(L)], \\ \rho_i: \mathcal{V}_i^{-1} \cdot \mathcal{CFL}(\mathcal{H}) &\rightarrow \mathcal{CFL}(\mathcal{H}) / (\mathcal{V}_i - 1) \otimes \mathbb{F}[\mathbb{H}_i(L)]. \end{aligned}$$

Here, $\mathbb{F}[\mathbb{H}_i(L)]$ denotes the $\mathbb{F}[T_i, T_i^{-1}]$ -module generated by T_i^α where $\alpha \in \mathbb{H}_i(L)$. Here, T_i is a formal variable. We also view $\mathcal{CFL}(\mathcal{H})/(\mathcal{U}_i - 1)$ and $\mathcal{CFL}(\mathcal{H})/(\mathcal{V}_i - 1)$ as $\mathbb{F}[U_i, T_i, T_i^{-1}]$ -modules, where U_i acts by \mathcal{V}_i on $\mathcal{CFL}(\mathcal{H})/(\mathcal{U}_i - 1)$, and by \mathcal{U}_i on $\mathcal{CFL}(\mathcal{H})/(\mathcal{V}_i - 1)$.

The map θ_i is given as follows. Write

$$\mathcal{Q}_{\mathcal{U}_i}: \mathcal{U}_i^{-1} \mathcal{CFL}(\mathcal{H}) \rightarrow \mathcal{CFL}(\mathcal{H})/(\mathcal{U}_i - 1)$$

for the quotient map, which sends \mathcal{U}_i to 1 and which sends \mathcal{V}_i to U_i . Then θ_i is given by the formula

$$\theta_i(\mathbf{x}) = \mathcal{Q}_{\mathcal{U}_i}(\mathbf{x}) \otimes T_i^{A_i^L(\mathbf{x})}.$$

The map ρ_i is defined similarly.

We write \mathcal{I}_i for the inclusion of $\mathcal{CFL}(\mathcal{H})$ into $\mathcal{U}_i^{-1} \cdot \mathcal{CFL}(\mathcal{H})$.

Moving the base point z_i to w_i along a subarc of K_i determines a homotopy equivalence

$$\Psi_{z_i \rightarrow w_i}: \mathcal{CFL}(\mathcal{H})/(\mathcal{U}_i - 1) \rightarrow \mathcal{CFL}(\mathcal{H})/(\mathcal{V}_i - 1)$$

which preserves $A_j^{L \setminus K_i}$ for $j \neq i$. The map $\Psi_{z_i \rightarrow w_i}$ is $\mathbb{F}[U_i, T_i, T_i^{-1}]$ -equivariant, and is equivariant with respect to the variables for other link components.

The map $\Phi_{L, L \setminus K_i}^{-K_i}$ is defined as the composition

$$\Phi_{L, L \setminus K_i}^{-K_i} := \rho_i^{-1} \circ (\Psi_{z_i \rightarrow w_i} \otimes T_i^{\lambda_i}) \circ \theta_i \circ \mathcal{I}_i.$$

We now consider the interaction of $\Phi_{L, L \setminus K_i}^{-K_i}$ with \mathcal{U}_i and \mathcal{V}_i . The map \mathcal{I}_i commutes with \mathcal{U}_i and \mathcal{V}_i . The map θ_i has the property that

$$\theta_i(\mathcal{U}_i^n \cdot \mathbf{x}) = T_i^{-n} \cdot \theta_i(\mathbf{x}) \quad \text{and} \quad \theta_i(\mathcal{V}_i^n \cdot \mathbf{x}) = U_i^n T_i^n \cdot \theta_i(\mathbf{x}).$$

Similarly,

$$\rho_i(\mathcal{U}_i^n \cdot \mathbf{x}) = U_i^n T_i^{-n} \cdot \rho_i(\mathbf{x}) \quad \text{and} \quad \rho_i(\mathcal{V}_i^n \cdot \mathbf{x}) = T_i^n \cdot \rho_i(\mathbf{x}),$$

which implies that

$$\rho_i^{-1}(T_i^n \cdot \mathbf{x}) = \mathcal{V}_i^n \cdot \rho_i^{-1}(\mathbf{x}) \quad \text{and} \quad \rho_i^{-1}(U_i^n \cdot \mathbf{x}) = \mathcal{U}_i^n \mathcal{V}_i^n \cdot \rho_i^{-1}(\mathbf{x}).$$

From these relations, it follows that

$$\Phi_{L, L \setminus K_i}^{-K_i}(\mathcal{U}_i \cdot \mathbf{x}) = \mathcal{V}_i^{-1} \cdot \Phi_{L, L \setminus K_i}^{-K_i}(\mathbf{x}) \quad \text{and} \quad \Phi_{L, L \setminus K_i}^{-K_i}(\mathcal{V}_i \cdot \mathbf{x}) = \mathcal{U}_i \mathcal{V}_i^2 \cdot \Phi_{L, L \setminus K_i}^{-K_i}(\mathbf{x}),$$

completing the proof. \blacksquare

Example 6.8. We now describe the surgery complex for the unknot \mathbb{O} in our present notation, and compare it to the traditional notation of [43]. The complex $\mathcal{CFK}(\mathbb{O})$ has one generator, 1, over $\mathbb{F}[\mathcal{U}, \mathcal{V}]$. In our notation, v sends $\mathcal{U}^i \mathcal{V}^j \in \mathcal{CFK}(\mathbb{O})$ to $\mathcal{U}^i \mathcal{V}^j \in \mathcal{V}^{-1} \mathcal{CFK}(\mathbb{O})$. Similarly, h_λ sends $\mathcal{U}^i \mathcal{V}^j$ to $\mathcal{V}^\lambda \phi^\tau(\mathcal{U}^i \mathcal{V}^j) = \mathcal{U}^j \mathcal{V}^{2j-i+\lambda}$.

Traditionally, one works with the mapping cone formula over $\mathbb{F}[U]$. We identify both $A(\mathbb{O}, s) \subseteq \mathcal{CFK}(\mathbb{O})$ and $B(\mathbb{O}, s) \subseteq \mathcal{V}^{-1} \mathcal{CFK}(\mathbb{O})$ with $\mathbb{F}[U]$ (where U acts by $\mathcal{U}\mathcal{V}$). We identify $\mathcal{U}^i \mathcal{V}^j \in \mathcal{CFK}(\mathbb{O})$ with $U^{\min(i,j)} \in A(\mathbb{O}, s)$, and we identify $\mathcal{U}^i \mathcal{V}^j \in \mathcal{V}^{-1} \mathcal{CFK}(\mathbb{O})$ with $U^i \in B(\mathbb{O}, s)$. We can view the previous identification as giving two isomorphisms

$$\mathcal{CFK}(\mathbb{O}) \cong \bigoplus_{s \in \mathbb{Z}} \mathbb{F}[U] \quad \text{and} \quad \mathcal{V}^{-1} \mathcal{CFK}(\mathbb{O}) \cong \bigoplus_{s \in \mathbb{Z}} \mathbb{F}[U].$$

In the original notation of Ozsváth and Szabó,

$$v(U^i) = \begin{cases} U^{i-s} & \text{if } s \leq 0, \\ U^i & \text{if } s \geq 0, \end{cases} \quad h_\lambda(U^i) = \begin{cases} U^i & \text{if } s \leq 0, \\ U^{i+s} & \text{if } s \geq 0. \end{cases}$$

The map h_λ sends $A(\mathbb{O}, s)$ to $B(\mathbb{O}, s + \lambda)$. It is straightforward to see that the above maps form a commutative diagram

$$\begin{array}{ccc} \mathcal{CFK}(\mathbb{O}) & \xrightarrow{f} & \mathcal{V}^{-1} \mathcal{CFK}(\mathbb{O}) \\ \downarrow \cong & & \downarrow \cong \\ \bigoplus_{s \in \mathbb{Z}} A(\mathbb{O}, s) & \xrightarrow{f} & \bigoplus_{s \in \mathbb{Z}} B(\mathbb{O}, s) \end{array}$$

whenever f is one of v or h_λ .

Chapter 7

The knot surgery algebra

In this chapter, we introduce our algebras \mathcal{K} and \mathcal{L} . We also describe natural filtrations on \mathcal{K} and \mathcal{L} . Using these filtrations, we describe our module categories.

7.1 The algebras \mathcal{K} and \mathcal{L}

We recall the *knot surgery algebra* \mathcal{K} from the introduction. We define \mathcal{K} to be an algebra over the idempotent ring

$$\mathbf{I} = \mathbf{I}_0 \oplus \mathbf{I}_1,$$

where each of \mathbf{I}_i is rank 1 over $\mathbb{F} = \mathbb{Z}/2$. We write i_0 and i_1 for the generators of \mathbf{I}_0 and \mathbf{I}_1 , respectively. We set

$$\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0 = \mathbb{F}[\mathcal{U}, \mathcal{V}] \quad \text{and} \quad \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1 \cong \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}].$$

Also,

$$\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_1 = 0.$$

We define

$$\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0 = \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}] \otimes \langle \sigma, \tau \rangle.$$

That is, elements of $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$ can be written as sums of monomials of the form $\mathcal{U}^i \mathcal{V}^j \sigma$ and $\mathcal{U}^i \mathcal{V}^j \tau$ where $i \geq 0$ and $j \in \mathbb{Z}$. These algebra elements satisfy the relations

$$\sigma \cdot a = \phi^\sigma(a) \cdot \sigma \quad \text{and} \quad \tau \cdot a = \phi^\tau(a) \cdot \tau,$$

for $a \in \mathbb{F}[\mathcal{U}, \mathcal{V}] = \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$. Here, $\phi^\tau: \mathbb{F}[\mathcal{U}, \mathcal{V}] \rightarrow \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$ is the algebra homomorphism satisfying $\phi^\tau(\mathcal{U}) = \mathcal{V}^{-1}$ and $\phi^\tau(\mathcal{V}) = \mathcal{U}\mathcal{V}^2$. The map $\phi^\sigma: \mathbb{F}[\mathcal{U}, \mathcal{V}] \rightarrow \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$ is the canonical inclusion.

In particular, $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$ is generated by two special algebra elements σ and τ , together with the left action of $\mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$, which satisfy the relations

$$\begin{aligned} \sigma \mathcal{U} &= \mathcal{U} \sigma & \text{and} & & \sigma \mathcal{V} &= \mathcal{V} \sigma, \\ \tau \mathcal{U} &= \mathcal{V}^{-1} \tau & \text{and} & & \tau \mathcal{V} &= \mathcal{U} \mathcal{V}^2 \tau. \end{aligned}$$

The link algebra \mathcal{L}_ℓ is defined as

$$\mathcal{L}_\ell := \mathcal{K} \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathcal{K}.$$

We often write just \mathcal{L} , when ℓ is determined by context. We view \mathcal{L}_ℓ as being an algebra over the idempotent ring

$$\mathbf{E}_\ell := \mathbf{I} \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathbf{I}.$$

Remark 7.1. The algebra \mathcal{K} admits a more symmetric basis, as follows. We may view $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$ as being $\mathbb{F}[U, T, T^{-1}]$, where $U = \mathcal{U}\mathcal{V}$ and $T = \mathcal{V}$. Then the above relations become

$$\begin{aligned} \sigma\mathcal{U} &= UT^{-1}\sigma & \text{and} & & \sigma\mathcal{V} &= T\sigma, \\ \tau\mathcal{U} &= T^{-1}\tau & \text{and} & & \tau\mathcal{V} &= UT\tau. \end{aligned}$$

Remark 7.2. The above symmetry may be packaged into an algebra homomorphism

$$\mathcal{E}: \mathcal{K} \rightarrow \mathcal{K}$$

as follows. On $\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$, we set

$$\mathcal{E}(\mathcal{U}) = \mathcal{V} \quad \text{and} \quad \mathcal{E}(\mathcal{V}) = \mathcal{U}.$$

On $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$, we set

$$\mathcal{E}(U) = U \quad \text{and} \quad \mathcal{E}(\mathcal{V}^i) = \mathcal{V}^{-i}.$$

On $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$, we set

$$\mathcal{E}(\sigma) = \tau \quad \text{and} \quad \mathcal{E}(\tau) = \sigma.$$

7.2 Topologies on \mathcal{K} and \mathcal{L}

In this section, we describe the topologies on \mathcal{K} and \mathcal{L} which we use throughout the memoir.

Definition 7.3. Suppose that $n \in \mathbb{N}$ is fixed. We define $J_n \subseteq \mathcal{K}$ to be the \mathbb{F} span of the following set of generators:

- (1) In $\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$, the generators $\mathcal{U}^i\mathcal{V}^j$, for $i \geq n$ or $j \geq n$ (i.e., $\max(i, j) \geq n$).
- (2) In $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$, the generators $\mathcal{U}^i\mathcal{V}^j\sigma$ for $i \geq n$ or $j \geq n$.
- (3) In $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$, the generators $\mathcal{U}^i\mathcal{V}^j\tau$ for $j \leq 2i - n$ or $i \geq n$.
- (4) In $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$, the generators $\mathcal{U}^i\mathcal{V}^j$ where $i \geq n$.

The subspace J_n is illustrated in Figure 7.1.

We observe that

$$J_0 = \mathcal{K}, \quad J_{n+1} \subseteq J_n, \quad \text{and} \quad \bigcap_{n \in \mathbb{N}} J_n = \{0\}.$$




Figure 7.1. The subspace $J_n \subseteq \mathcal{K}$, indicated by the shaded regions. Here, J_n^σ and J_n^τ denote the subspaces of J_n which are in the \mathcal{K} span of σ and τ . The dots indicate monomials in the algebra. The horizontal direction indicates the power of \mathcal{U} , and the vertical direction indicates the power of \mathcal{V} .

We will use the subspaces J_n as a fundamental basis of 0 in our topology on \mathcal{K} . The condition that $\bigcap_{n \in \mathbb{N}} J_n = \{0\}$ is equivalent to the inclusion $\mathcal{K} \rightarrow \mathcal{K}$ being injective.

Remark 7.4. The subspaces J_n^σ and J_n^τ can also be described as follows:

$$\begin{aligned} J_n^\sigma &= (U^n) \cdot \sigma + \sigma \cdot (U^n, \mathcal{V}^n), \\ J_n^\tau &= (U^n) \cdot \tau + \tau \cdot (U^n, \mathcal{V}^n). \end{aligned}$$

Here, $(U^n) \subseteq \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$ denotes the ideal generated by $U = \mathcal{U}\mathcal{V}$, and $(U^n, \mathcal{V}^n) \subseteq \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$ denotes the ideal spanned by the elements U^n and \mathcal{V}^n .

We now consider the link algebra. For our purposes, the most natural topology on the link algebra is the !-topology,

$$\mathcal{L}_\ell = \mathcal{K} \otimes_{\mathbb{F}}^! \cdots \otimes_{\mathbb{F}}^! \mathcal{K}.$$

This topology is first countable, with a basis of opens given by

$$\mathcal{J}_{\ell,n} := \sum_{i=0}^{\ell-1} \mathcal{K}^{\otimes i} \otimes_{\mathbb{F}} J_n \otimes_{\mathbb{F}} \mathcal{K}^{\otimes \ell-i-1} \subseteq \mathcal{L}_\ell.$$

A fundamental result to our memoir is the following.

Proposition 7.5. *The map $\mu_2: \mathcal{K} \overset{\rightarrow}{\otimes}_I \mathcal{K} \rightarrow \mathcal{K}$ is continuous.*

Remark 7.6. Multiplication is not continuous on $\mathcal{K} \otimes_1^! \mathcal{K}$. As an example, $\mathcal{V}^{-i} \otimes \mathcal{V}^i \sigma \rightarrow 0$ in $\mathcal{K} \otimes_I \mathcal{K}$ as $i \rightarrow \infty$, whereas $\mu_2(\mathcal{V}^{-i} \otimes \mathcal{V}^i \sigma) = \sigma \not\rightarrow 0$.

Proposition 7.5 follows immediately from the following lemma and the definition of $\overset{\rightarrow}{\otimes}$ in Section 5.4.

Lemma 7.7. *Let $J_n \subseteq \mathcal{K}$ be the subspace defined in Definition 7.3.*

- (1) *For all n , $\mu_2(J_n \otimes \mathcal{K}) \subseteq J_n$ (i.e., J_n is a right ideal of \mathcal{K}).*
- (2) *Suppose $x \in \mathcal{K}$ and $n \in \mathbb{N}$ are fixed. Then for sufficiently large m ,*

$$\mu_2(x \otimes J_m) \subseteq J_n.$$

Proof. We consider the first claim. We wish to show that $\mu_2(x \otimes a) \in J_n$ whenever $x \in J_n$ and $a \in \mathcal{K}$. We break the argument into cases, depending on the idempotents of x and a . If $x, a \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$, the claim follows because $w_U(x \cdot a) \geq w_U(x) + w_U(a)$, where w_U is the U -adic weight (i.e., $w_U(a) = \max\{i : \exists a', a = U^i \cdot a'\}$). The same argument applies if $x \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$ and $a \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$. The situation that $x, a \in \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$ is also clear. We now consider the case that $x \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$ and $a \in \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$. Suppose first that $x = \mathcal{U}^i \mathcal{V}^j \sigma$ and $a = \mathcal{U}^s \mathcal{V}^t$ (for $s, t \geq 0$). Then $x \cdot a = \mathcal{U}^{i+s} \mathcal{V}^{j+t} \sigma$, which is clearly also in J_n . Similarly, if $x = \mathcal{U}^i \mathcal{V}^j \tau$ and $a = \mathcal{U}^s \mathcal{V}^t$, then

$$x \cdot a = \mathcal{U}^{i+t} \mathcal{V}^{j-s+2t} \tau.$$

By assumption, either $i \geq n$ or $j \leq 2i - n$. If $i \geq n$, then $i + t \geq n$ so $x \cdot a \in J_n$. The second case is equivalent to $2i - j - n \geq 0$. We note that

$$2(i + t) - (j - s + 2t) - n = s + 2i - j - n \geq 0$$

so $x \cdot a \in J_n$. This completes the first part of the claim.

We leave the second claim to the reader, as it may be handled similarly. ■

Combining Corollary 5.15 and Proposition 7.5, we also obtain the following.

Corollary 7.8. *Suppose \mathcal{L} is topologized as $\mathcal{K} \otimes^1 \cdots \otimes^1 \mathcal{K}$. Then the map*

$$\mu_2: \mathcal{L} \overset{\rightarrow}{\otimes}_E \mathcal{L} \rightarrow \mathcal{L}$$

is continuous.

7.3 Alexander modules

In this section, we describe our categories of type-D, A and DA modules. We will refer to these categories as the categories of *Alexander modules*.

Definition 7.9. A *type-D Alexander module* consists of a pair (\mathcal{X}, δ^1) where \mathcal{X} is a linear topological right \mathbf{E} -module and

$$\delta^1: \mathcal{X} \rightarrow \mathcal{X} \vec{\otimes}_{\mathbf{E}} \mathcal{L}$$

is a linear topological morphism which satisfies

$$(\text{id} \otimes \mu_2) \circ (\delta^1 \otimes \text{id}_{\mathcal{L}}) \circ \delta^1 = 0.$$

Recall that a linear topological morphism is the same as a continuous linear map on completions. If $\mathcal{X}^{\mathcal{L}}$ and $\mathcal{Y}^{\mathcal{L}}$, we define $\text{Mor}(\mathcal{X}^{\mathcal{L}}, \mathcal{Y}^{\mathcal{L}})$ to be the space of linear topological morphisms

$$f^1: \mathcal{X} \rightarrow \mathcal{Y} \vec{\otimes}_{\mathbf{E}} \mathcal{L}.$$

We write $\text{Mod}_{\alpha}^{\mathcal{K}}$ for the category of Alexander type-D modules.

Type- A Alexander modules are similar, as we see in the following definition.

Definition 7.10. A *type-A Alexander module* (\mathcal{X}, m_j) over \mathcal{L} consists of a linear topological left \mathbf{E} -module \mathcal{X} and a collection of linear topological morphisms

$$m_{j+1}: \mathcal{L} \vec{\otimes}_{\mathbf{E}} \cdots \vec{\otimes}_{\mathbf{E}} \mathcal{L} \vec{\otimes}_{\mathbf{E}} \mathcal{X} \rightarrow \mathcal{X},$$

satisfying the A_{∞} -relations.

Type- DA Alexander modules are defined by the obvious amalgamation of the above notions.

7.4 Split Alexander modules

We now define the notion of a *split* Alexander type- A or DA module. Consider the link algebra \mathcal{L}_{ℓ} and suppose that \mathbb{P} is a partition of $\{1, \dots, \ell\}$. Suppose that \mathcal{X} is a linear topological left \mathbf{E}_{ℓ} -module. A \mathbb{P} -*split continuous* linear topological morphism

$$f_{j+1}: \mathcal{L}_{\ell} \otimes_{\mathbf{E}} \cdots \otimes_{\mathbf{E}} \mathcal{L}_{\ell} \otimes_{\mathbf{E}} \mathcal{X} \rightarrow \mathcal{X}$$

is a continuous map on completions, where we give $cL_{\ell}^{\otimes j} \otimes \mathcal{X}$ the tree topology from Section 5.6, where \mathcal{X} is the central vertex and we have $|\mathbb{P}|$ -rays of length j extending from this vertex. If $p_1, \dots, p_{|\mathbb{P}|}$ are the sizes of the elements of \mathbb{P} , then each non-root vertex of the i -th ray is labeled with \mathcal{L}_{p_i} (topologized using the !-topology).

Note that \mathbb{P} -split continuity is preserved by A_∞ -composition of such maps by Proposition 5.24.

Definition 7.11. Suppose that \mathbb{P} is a partition of $\{1, \dots, \ell\}$. A \mathbb{P} -split type-A Alexander module consists of a linear topological $\mathbf{E} = \mathbf{E}_\ell$ -module \mathcal{X} , equipped with a \mathbb{P} -split linear topological morphism

$$m_{j+1}: \mathcal{L}_\ell \otimes_{\mathbf{E}} \cdots \otimes_{\mathbf{E}} \mathcal{L}_\ell \otimes_{\mathbf{E}} \mathcal{X} \rightarrow \mathcal{X},$$

which satisfies the A_∞ -module relations.

If \mathbb{P} is a partition of $\{1, \dots, \ell\}$ and i_1, \dots, i_n are the sizes of the partition elements (so $i_1 + \cdots + i_n = \ell$), we will write

$$\mathcal{L}_{i_1 | \cdots | i_n} \mathcal{X}$$

for a \mathbb{P} -split Alexander module.

Once we introduce box tensor products of split Alexander modules, we will show in Lemma 7.16 that the split-Alexander condition is weaker than the Alexander condition. In particular, any type-A Alexander module \mathcal{X} is also a split type-A Alexander module for any partition.

\mathbb{P} -split Alexander type-DA modules can be defined similarly to the above, though we typically need to assume that the underlying space \mathcal{X} of the module is linearly compact for the type-DA structure relations to be meaningful, and for the A_∞ -composition of morphisms to be defined. (Note that all of the link surgery modules we consider in this memoir are linearly compact.) We explain this in the case of a $\{\{1\}, \{2\}\}$ -split Alexander module $\mathcal{K}|\mathcal{K} \mathcal{X}^{\mathcal{K}}$. Consider the spaces involved in the composition of $\delta_{i+1}^1, \delta_{j+1}^1$ and $\mathbb{I}_{\mathcal{X}} \otimes \mu_2$, applied to

$$\underbrace{(\mathcal{K}|\mathcal{K}) \otimes_{\mathbf{E}_2} \cdots \otimes_{\mathbf{E}_2} (\mathcal{K}|\mathcal{K}) \otimes_{\mathbf{E}_2} \mathcal{X}}_{i+j}.$$

We topologize the above space using the tree topology for the tree where there are 2 length $i + j$ rays attached to a lowest vertex. The lowest vertex is assigned the space \mathcal{X} . Proposition 5.24 shows that $\delta_{i+1}^1 \otimes \text{id}$ gives a map from the above space to

$$\underbrace{(\mathcal{K}|\mathcal{K}) \otimes_{\mathbf{E}_2} \cdots \otimes_{\mathbf{E}_2} (\mathcal{K}|\mathcal{K}) \otimes_{\mathbf{E}_2} \mathcal{X}}_j \bar{\otimes} \mathcal{K}.$$

The above space is topologized using a tree with a single lowest vertex, and 2 length i rays extending from this ray. The lowest vertex is assigned the space $\mathcal{X} \bar{\otimes} \mathcal{K}$. Lemma 5.27 implies that if \mathcal{X} is linearly compact, then the above topological vector space is isomorphic to the same tensor product, equipped with the tree topology where the tree Γ has 2 length i rays connected to a vertex v_1 , which is above a single vertex v_0 . We assign v_1 the vector space \mathcal{X} , and v_0 the vector space \mathcal{K} . See Figure 7.2 for a schematic when $i = j = 1$.




Figure 7.2. Maps appearing in the DA-structure relation for a split Alexander DA bimodule. Here, the arrows labeled \cong correspond to the isomorphism from Lemma 5.27, which requires \mathcal{X} to be linearly compact.

7.5 External tensor products

We now discuss external tensor products, focusing on subtleties involving completions.

We first discuss type-D modules. Suppose that $\mathcal{X}^{\mathcal{L}_m}$ and $\mathcal{Y}^{\mathcal{L}_n}$ are type-D Alexander modules. To form the external tensor products, we assume that \mathcal{X} and \mathcal{Y} are linearly compact, so that $\mathcal{X} \bar{\otimes} \mathcal{L}_m \cong \mathcal{X} \otimes^! \mathcal{L}_m$ by Lemma 5.16, and similarly for $\mathcal{Y} \bar{\otimes} \mathcal{L}_n$. Under this assumption, we may define the structure map on the tensor product as $\delta_{\mathcal{X}}^1 \otimes \text{id}_{\mathcal{Y}} \otimes 1_{\mathcal{L}_n} + \text{id}_{\mathcal{X}} \otimes 1_{\mathcal{L}_m} \otimes \delta_{\mathcal{Y}}^1$, with tensor factors reordered, as indicated below:

$$\begin{aligned} \mathcal{X} \otimes^! \mathcal{Y} &\rightarrow (\mathcal{X} \bar{\otimes} \mathcal{L}_m) \otimes^! (\mathcal{Y} \bar{\otimes} \mathcal{L}_n) \cong \mathcal{X} \otimes^! \mathcal{L}_m \otimes^! \mathcal{Y} \otimes^! \mathcal{L}_n \\ &\cong (\mathcal{X} \otimes^! \mathcal{Y}) \bar{\otimes} (\mathcal{L}_m \otimes^! \mathcal{L}_n). \end{aligned}$$

Next, we consider tensoring (split) type-A modules. Suppose that \mathbb{P} is a partition of $\{1, \dots, s\}$ and \mathbb{P}' is a partition of $\{1, \dots, r\}$. Let us write i_1, \dots, i_m for the sizes of the elements of \mathbb{P} , and j_1, \dots, j_n for those of \mathbb{P}' . Suppose that we have \mathbb{P} and \mathbb{P}' split modules

$$\mathcal{L}_{i_1|\dots|i_m} \mathcal{X} \quad \text{and} \quad \mathcal{L}_{j_1|\dots|j_n} \mathcal{Y}.$$

We will describe two versions of the external tensor product, namely

$$\mathcal{L}_{i_1|\dots|i_m+j_1|\dots|j_n} (\mathcal{X} \otimes^! \mathcal{Y}) \quad \text{and} \quad \mathcal{L}_{i_1|\dots|i_m|\mathcal{L}_{j_1|\dots|j_n}} (\mathcal{X} \otimes^! \mathcal{Y}).$$

We topologize $\mathcal{L}_{r+s}^{\otimes j} \otimes \mathcal{X} \otimes \mathcal{Y}$ using the tree topology for a star with $m + n - 1$ (resp. $m + n$) rays). We label the root $\mathcal{X} \otimes^! \mathcal{Y}$, and we label each other vertex with a $\otimes^!$ -tensor product of copies of \mathcal{K} .

To construct the external tensor product, one must also pick a cellular diagonal of the associahedron. This construction is recalled in Section 3.5. The fact that the external tensor product behaves well with completions is implied by the following result.

Proposition 7.12. *Suppose that \mathbb{P} and \mathbb{P}' are partitions of $\{1, \dots, r\}$ and $\{1, \dots, s\}$, respectively, and that*

$$f_{j+1}: \mathcal{L}_r^{\otimes j} \otimes \mathcal{X} \rightarrow \mathcal{X} \quad \text{and} \quad g_{j+1}: \mathcal{L}_s^{\otimes j} \otimes \mathcal{Y} \rightarrow \mathcal{Y}$$

are \mathbb{P} and \mathbb{P}' split continuous, then their tensor product $f_{j+1} \otimes g_{j+1}$ is both $\mathbb{P} \sqcup \mathbb{P}'$ and $(\mathbb{P} \sqcup \mathbb{P}')_{p \sim p'}$ split continuous, for any $p \in \mathbb{P}$ and $p' \in \mathbb{P}'$. (Here, we write $(\mathbb{P} \sqcup \mathbb{P}')_{p \sim p'}$ for the partition obtained by merging p and p' in $\mathbb{P} \sqcup \mathbb{P}'$.)

We will prove Proposition 7.12 for the module $\mathcal{L}_{i_1 | \dots | \mathcal{L}_{i_m+j_1} | \dots | \mathcal{L}_{j_n}}(\mathcal{X} \otimes^! \mathcal{Y})$. We will later see in Lemma 7.16 that the module $\mathcal{L}_{i_1 | \dots | \mathcal{L}_{i_m} | \mathcal{L}_{j_1} | \dots | \mathcal{L}_{j_n}}(\mathcal{X} \otimes^! \mathcal{Y})$ is a box tensor product of this module with a split Alexander bimodule $\mathcal{L}_{i_m | \mathcal{L}_{j_1}}[\mathbb{I}]^{\mathcal{L}_{i_m+j_1}}$, and hence is also a split Alexander module.

The proof Proposition 7.12 is an immediate consequence of the following general result about tree topologies.

Lemma 7.13. *Let Γ_n^k denote the star which has n rays of length k , which all point towards a root vertex v_0 . Let $(\mathcal{X}_v)_{v \in V(\Gamma_n^k)}$ and $(\mathcal{Y}_v)_{v \in V(\Gamma_n^k)}$ denote two families of spaces. Pick a distinguished ray from each of Γ_n^k and Γ_m^k . Form a new collection of spaces $(\mathcal{Z}_v)_{v \in \Gamma_{m+n-1}^k}$ as follows. View Γ_{m+n-1}^k as being obtained by merging vertices of the same height along the distinguished rays. We define \mathcal{Z}_v to be either \mathcal{X}_v or \mathcal{Y}_v for v not along the merged ray. For v along the merged ray, we define \mathcal{Z}_v to be $\mathcal{X}_v \otimes^! \mathcal{Y}_v$. Then the natural map*

$$\mathcal{Z}_{\Gamma_{m+n-1}^k} \rightarrow \mathcal{X}_{\Gamma_n^k} \otimes^! \mathcal{Y}_{\Gamma_m^k}$$

is continuous. See Figure 7.3 for a schematic where $m = n = 2$.




Figure 7.3. A schematic of Lemma 7.13. Solid dots denote \mathcal{X}_v and open dots denote \mathcal{Y}_v .

Proof. We abbreviate $\mathcal{X}_{\Gamma_n^k}$, $\mathcal{Y}_{\Gamma_m^k}$ and $\mathcal{Z}_{\Gamma_{m+n-1}^k}$ by \mathcal{X} , \mathcal{Y} and \mathcal{Z} , respectively.

Let

$$U_{\mathcal{X}} \otimes \mathcal{Y} + \mathcal{X} \otimes V_{\mathcal{Y}} \subseteq \mathcal{X} \otimes^! \mathcal{Y}$$

be an open subspace, where $U_{\mathcal{X}} \subseteq \mathcal{X}$ and $V_{\mathcal{Y}} \subseteq \mathcal{Y}$ are open subspaces. Let s be an admissible labeling of Γ_{n+m-1}^k . Our goal is to show that if \mathbf{p} is a point-enhancement of s , then there is an open-enhancement $\mathcal{U}_{\mathbf{p}}$ so that

$$\mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p}) \subseteq U_{\mathcal{X}} \otimes \mathcal{Y} + \mathcal{X} \otimes V_{\mathcal{Y}}.$$

There are two induced labelings $s_{\mathcal{X}}$ and $s_{\mathcal{Y}}$ of Γ_n^k and Γ_m^k , obtained by restricting s . We make the following claims:

- (1) At least one of $s_{\mathcal{X}}$ and $s_{\mathcal{Y}}$ is admissible.
- (2) If there is an \mathcal{O} -labeled vertex along the merged ray, then both $s_{\mathcal{X}}$ and $s_{\mathcal{Y}}$ are admissible.

The above claims are straightforward to prove, so we leave the argument to the reader.

Consider first the case that the merged ray has no vertices labeled \mathcal{O} . In this case, the root is labeled X , but all other vertices along the merged ray are labeled p . Also, in this case, at least one of $s_{\mathcal{X}}$ and $s_{\mathcal{Y}}$ is admissible. Suppose, without loss of generality, that $s_{\mathcal{X}}$ is admissible. Write v_1, \dots, v_{k-1} for the p -labeled vertices along the merged ray. If v is a p -labeled vertex along the merged ray, write $\mathbf{p}(v)$ as a sum of finitely many elementary tensors. For each $v \in \{v_1, \dots, v_{k-1}\}$, write $x_{v,1}, \dots, x_{v,i_v}$ for these elementary tensors. We consider the $N := i_{v_1} \cdots i_{v_{k-1}}$ point-enhancements $\mathbf{p}_{\mathcal{X},j}$, where $\mathbf{p}_{\mathcal{X},j}(w) = \mathbf{p}(w)$ if w is not along the merged ray, and such that $\mathbf{p}_{\mathcal{X},j}$ enumerate all combinations of the $x_{v_i,r}$ for $i \in \{1, \dots, k-1\}$ and $r \in \{1, \dots, i_{v_i}\}$.

Since $U_{\mathcal{X}}$ is open, for each $\mathbf{p}_{\mathcal{X},i}$, there is an open-enhancement $\mathcal{U}_{\mathcal{X},i} \subseteq \mathcal{X}$ so that

$$\mathcal{W}(\mathcal{U}_{\mathcal{X},i}, \mathbf{p}_{\mathcal{X},i}) \subseteq U_{\mathcal{X}}.$$

We define an open-enhancement of s as follows. If v is in Γ_n^k (the tree for \mathcal{X}), we define

$$\mathcal{U}_{\mathbf{p}}(v) = \bigcap_{i=1}^N \mathcal{U}_{\mathcal{X},i}(v).$$

If v is in Γ_m^k (the tree for \mathcal{Y}), set $\mathcal{U}_{\mathbf{p}}(v) = \mathcal{Y}_v$. Then

$$\begin{aligned} \mathcal{W}(\mathcal{U}_{\mathbf{p}}, \mathbf{p}) &\subseteq \mathcal{W}(\mathcal{U}_{\mathcal{X},1}, \mathbf{p}_{\mathcal{X},1}) \otimes \mathcal{Y} + \cdots + \mathcal{W}(\mathcal{U}_{\mathcal{X},N}, \mathbf{p}_{\mathcal{X},N}) \otimes \mathcal{Y} \\ &\subseteq U_{\mathcal{X}} \otimes \mathcal{Y} + \mathcal{X} \otimes V_{\mathcal{Y}}, \end{aligned}$$

as we wanted to show.

Next, we consider the case that s labels a vertex v_0 along the merged ray by \mathcal{O} . In this case, both $s_{\mathcal{X}}$ and $s_{\mathcal{Y}}$ are admissible. Consider a point-enhancement \mathbf{p} of s . Write v_1, \dots, v_l for the vertices above v_0 (which are labeled p by s). Similar to before, by separating the elements of \mathbf{p} into elementary tensors, we obtain point-enhancements

$\mathbf{p}_{x,i}$ and $\mathbf{p}_{y,j}$ for $i = 1, \dots, N$ and $j = 1, \dots, M$. For each i and j , there are open-enhancements $\mathcal{U}_{x,i}$ and $\mathcal{U}_{y,j}$ of s_x and s_y (respectively) so that

$$\mathcal{W}(\mathcal{U}_{x,i}, \mathbf{p}_{x,i}) \subseteq U_x \quad \text{and} \quad \mathcal{W}(\mathcal{U}_{y,j}, \mathbf{p}_{x,j}) \subseteq V_y.$$

We then define an open-enhancement \mathcal{U}_p of s as follows. Suppose $s(v) = \emptyset$. If v is in Γ_n^k not in the merged ray, we set

$$\mathcal{U}_p(v) = \bigcap_{i=1}^N \mathcal{U}_{x,i}(v).$$

If v is in Γ_m^k but not in the merged ray, we set

$$\mathcal{U}_p(v) = \bigcap_{j=1}^M \mathcal{U}_{y,j}(v).$$

If v is in the merged ray, we set

$$\mathcal{U}_p(v) = \left(\bigcap_{i=1}^N \mathcal{U}_{x,i}(v) \right) \otimes \mathcal{Y} + \mathcal{X} \otimes \left(\bigcap_{j=1}^M \mathcal{U}_{y,j}(v) \right).$$

We observe that

$$\begin{aligned} \mathcal{W}(\mathcal{U}_p, \mathbf{p}) &\subseteq (\mathcal{W}(\mathcal{U}_{x,1}, \mathbf{p}_{x,1}) + \dots + \mathcal{W}(\mathcal{U}_{x,N}, \mathbf{p}_{x,N})) \otimes \mathcal{Y} \\ &\quad + \mathcal{X} \otimes (\mathcal{W}(\mathcal{U}_{y,1}, \mathbf{p}_{y,1}) + \dots + \mathcal{W}(\mathcal{U}_{y,M}, \mathbf{p}_{y,M})) \\ &\subseteq U_x \otimes \mathcal{Y} + \mathcal{X} \otimes V_y. \end{aligned}$$

This completes the proof. ■

Taking the external tensor product of two type-DA Alexander modules, when the underlying spaces of the modules are linearly compact, is a straightforward extension.

Remark 7.14. We can see more directly that if \mathcal{X}_i and \mathcal{Y}_j are type-A Alexander modules, then $\mathcal{X}_{i+j}(\mathcal{X} \otimes^! \mathcal{Y})$ is also an Alexander module. Consider two continuous maps

$$f_{j+1}: \tilde{\mathcal{T}}^j \mathcal{L}_n \bar{\otimes} \mathcal{X} \rightarrow \mathcal{X} \quad \text{and} \quad g_{j+1}: \tilde{\mathcal{T}}^j \mathcal{L}_m \bar{\otimes} \mathcal{X} \rightarrow \mathcal{X},$$

where

$$\tilde{\mathcal{T}}^j \mathcal{L}_n = \underbrace{\mathcal{L}_n \bar{\otimes} \dots \bar{\otimes} \mathcal{L}_n}_j.$$

The external tensor product of f_{j+1} and g_{j+1} is continuous since it may be written as the following sequence of continuous maps:

$$\begin{aligned} &\tilde{\mathcal{T}}^j(\mathcal{L}_n \otimes^! \mathcal{L}_m) \bar{\otimes}(\mathcal{X} \otimes^! \mathcal{Y}) \\ &\rightarrow (\tilde{\mathcal{T}}^j \mathcal{L}_n \bar{\otimes} \mathcal{X}) \otimes^! (\tilde{\mathcal{T}}^j \mathcal{L}_m \bar{\otimes} \mathcal{Y}) \rightarrow \mathcal{X} \otimes^! \mathcal{Y}. \end{aligned}$$

The first map is from Corollary 5.15 and the second map is $f_{j+1} \otimes g_{j+1}$.

7.6 Box tensor products

In this section, we describe the interaction of Lipshitz, Ozsváth and Thurston's box tensor product operation with Alexander modules. The simplest case is boxing an Alexander type-D module $\mathcal{X}^{\mathcal{L}}$ with an Alexander type-A module \mathcal{Y} . In this case,

$$\delta^1: \mathcal{X} \rightarrow \mathcal{X} \overset{\circlearrowleft}{\otimes}_{\mathbb{E}} \mathcal{L} \quad \text{and} \quad m_{j+1}: \mathcal{L} \overset{\circlearrowleft}{\otimes}_{\mathbb{E}} \cdots \overset{\circlearrowleft}{\otimes}_{\mathbb{E}} \mathcal{L} \overset{\circlearrowleft}{\otimes}_{\mathbb{E}} \mathcal{Y} \rightarrow \mathcal{Y}$$

are continuous linear maps on the completions. Hence, assuming an appropriate boundedness condition so that only finitely many terms contribute to the box tensor product (e.g., $m_j = 0$ for $j \gg 0$), the differential ∂^{\boxtimes} will be a sum of finitely many continuous maps. Type-DA modules present no additional complications.

Box tensor products involving split Alexander type-A and DA modules are more subtle. If $\mathcal{L}_{i_1|\dots|\mathcal{L}_{i_n}} \mathcal{Y}$ is a split Alexander type-A module, it seems only possible to tensor with an Alexander type-D module which is itself an external tensor product of n modules

$$\mathcal{X}_1^{\mathcal{L}_{i_1}}, \dots, \mathcal{X}_n^{\mathcal{L}_{i_n}}.$$

In this case, we write

$$(\mathcal{X}_1^{\mathcal{L}_{i_1}}, \dots, \mathcal{X}_n^{\mathcal{L}_{i_n}}) \hat{\boxtimes}_{\mathcal{L}_{i_1}|\dots|\mathcal{L}_{i_n}} \mathcal{Y}$$

for the resulting tensor product. Ignoring completions, the tensor product is formed by taking the external tensor product of the type-D modules $\mathcal{X}_j^{\mathcal{L}_{i_j}}$ to obtain a type-D module over $\mathcal{L}_{i_1}|\dots|\mathcal{L}_{i_n}$. This type-D module is then tensored with the type-A module using the standard definition (see equation (3.1)). We topologize the underlying space of the tensor product

$$(\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n) \otimes \mathcal{Y}$$

using the graph topology from Section 5.6 for a tree with one root and n rays. Here, \mathcal{Y} is labeled as the root, and $\mathcal{X}_1, \dots, \mathcal{X}_n$ each form a ray, with an edge pointing to \mathcal{Y} . Proposition 5.24 shows that the structure operations are continuous.

We may also form the tensor product when \mathcal{Y} or some of the \mathcal{X}_i are split DA Alexander bimodules. The output will be a split Alexander DA bimodule, whose tensor splitting of the incoming algebras is the concatenation of the tensor-splittings for each of the \mathcal{X}_i . To illustrate, if $\mathcal{L}_{i_1|\mathcal{L}_{i_2}} \mathcal{Y}^{\mathcal{L}}$, $\mathcal{L}_{j_1|\dots|\mathcal{L}_{j_m}} \mathcal{X}_1^{\mathcal{L}_{i_1}}$ and $\mathcal{L}_{k_1|\dots|\mathcal{L}_{k_t}} \mathcal{X}_2^{\mathcal{L}_{i_2}}$ are split Alexander DA bimodules, then their tensor product, N , will be a split module as indicated in the following equation:

$$(\mathcal{L}_{j_1|\dots|\mathcal{L}_{j_m}} \mathcal{X}_1^{\mathcal{L}_{i_1}}, \mathcal{L}_{k_1|\dots|\mathcal{L}_{k_t}} \mathcal{X}_2^{\mathcal{L}_{i_2}}) \hat{\boxtimes}_{\mathcal{L}_{i_1}|\mathcal{L}_{i_2}} \mathcal{Y}^{\mathcal{L}} = \mathcal{L}_{j_1|\dots|\mathcal{L}_{j_m}|\mathcal{L}_{k_1}|\dots|\mathcal{L}_{k_t}} N^{\mathcal{L}}.$$

Remark 7.15. The box tensor product operation of split Alexander modules seems definable using our techniques only when the underlying vector spaces of our modules

are linearly compact. In this case, the underlying vector space of the tensor product $(\mathcal{X}^{\mathcal{K}}, \mathcal{Y}^{\mathcal{K}}) \hat{\boxtimes}_{\mathcal{K}|\mathcal{K}} \mathcal{Z}$ is the standard tensor product $\mathcal{X} \otimes^! \mathcal{Y} \otimes^! \mathcal{Z}$ by Lemma 5.16. Furthermore, assuming linear compactness, the completed box tensor product operation satisfies the following associativity property:

$$\begin{aligned} (\mathcal{X}^{\mathcal{K}}, \mathcal{Y}^{\mathcal{K}}) \hat{\boxtimes}_{\mathcal{K}|\mathcal{K}} \mathcal{Z} &\cong \mathcal{X}^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{K}} ((\mathcal{K}^{\mathbb{I}^{\mathcal{K}}}, \mathcal{Y}^{\mathcal{K}}) \hat{\boxtimes}_{\mathcal{K}|\mathcal{K}} \mathcal{Z}) \\ &\cong \mathcal{Y}^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{K}} ((\mathcal{X}^{\mathcal{K}}, \mathcal{K}^{\mathbb{I}^{\mathcal{K}}}) \hat{\boxtimes}_{\mathcal{K}|\mathcal{K}} \mathcal{Z}). \end{aligned}$$

Note that all of the bordered link surgery modules considered in this memoir are countable direct products of \mathbb{F} , so are linearly compact.

We now describe a split type-DA Alexander bimodule ${}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}]^{\mathcal{L}_2}$ which transforms any Alexander module ${}_{\mathcal{L}_2} \mathcal{X}$ into a split module ${}_{\mathcal{K}|\mathcal{K}} \mathcal{X}$. The structure maps are the same as for the identity module.

Lemma 7.16. *The bimodule ${}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}]^{\mathcal{L}_2}$ is a split Alexander bimodule. More generally, if \mathbb{P} is a partition of $\{1, \dots, \ell\}$, and ℓ_1, \dots, ℓ_j are the sizes of the partition elements, then ${}_{\mathcal{L}_{\ell_1|\dots|\ell_j}}[\mathbb{I}]^{\mathcal{L}_\ell}$ is a split Alexander bimodule.*

Proof. Consider first ${}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}]^{\mathcal{L}_2}$. In this case, the claim follows from Remark 5.23 and the fact that the natural map

$$\mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K} \rightarrow \mathcal{K} \otimes_{\mathbb{F}}^! \mathcal{K}$$

is continuous. The case of general ℓ and \mathbb{P} is proven by the same argument. ■

Remark 7.17. Lemma 7.16 implies that the ordinary Alexander module condition is stronger than the split Alexander module condition. In particular, given a type-A Alexander module ${}_{\mathcal{L}_\ell} \mathcal{X}$, we may always view \mathcal{X} as a split Alexander module for any partition of $\{1, \dots, \ell\}$.

7.7 Finitely generated \mathcal{K} -modules

In this section, we discuss the category of finitely generated type-D modules over \mathcal{L} .

We denote by \mathcal{K} the completion of the algebra \mathcal{K} with respect to the topology described above. As a vector space, we have the following isomorphisms:

$$\begin{aligned} \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0 &\cong \mathbb{F}[\mathcal{U}, \mathcal{V}], \\ \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0 &\cong \mathbb{F}[\mathcal{V}, \mathcal{V}^{-1}][\mathcal{U}] \langle \tau \rangle \oplus \mathbb{F}[\mathcal{V}, \mathcal{V}^{-1}][\mathcal{U}] \langle \sigma \rangle, \\ \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1 &\cong \mathbb{F}[\mathcal{V}, \mathcal{V}^{-1}][\mathcal{U}]. \end{aligned}$$

We may consider the completion of the link algebra \mathcal{L} as well.

Since multiplication is continuous by Proposition 7.5 and Corollary 7.8, we obtain a well-defined map on the completed algebras

$$\mu_2: \mathcal{L} \overset{\sim}{\otimes}_E \mathcal{L} \rightarrow \mathcal{L}.$$

Since \mathcal{L} is an algebra, we may also consider the ordinary category of type-D modules over this algebra (i.e., ordinary modules and algebras, with no topologies). Of particular interest to us is the category of finitely generated type-D modules. We write $\text{Mod}_{\text{fg}}^{\mathcal{L}}$ for this category.

There is a related category, $\text{Mod}_{\text{fg},\alpha}^{\mathcal{L}}$, consisting of Alexander type-D modules over \mathcal{L} which are finitely generated over \mathbb{F} .

Proposition 7.18. *The categories $\text{Mod}_{\text{fg},\alpha}^{\mathcal{L}}$ and $\text{Mod}_{\text{fg}}^{\mathcal{L}}$ are equivalent.*

Proof. We define functors in both directions. The functor

$$F: \text{Mod}_{\text{fg},\alpha}^{\mathcal{L}} \rightarrow \text{Mod}_{\text{fg}}^{\mathcal{L}}$$

is obtained by taking completions, noting that the completion of $\mathcal{X} \overset{\sim}{\otimes} \mathcal{L}$ coincides with the ordinary tensor product $\mathcal{X} \otimes \mathcal{L}$ when \mathcal{X} is finitely generated. (Note that when \mathcal{X} is finite-dimensional, $\mathcal{X} \cong \mathcal{X}$ if and only if \mathcal{X} is Hausdorff.)

The functor

$$G: \text{Mod}_{\text{fg}}^{\mathcal{L}} \rightarrow \text{Mod}_{\text{fg},\alpha}^{\mathcal{L}}$$

sends a type-D module (\mathcal{X}, δ^1) to itself, where \mathcal{X} is equipped with the discrete topology. This is well defined due to the fact that if \mathcal{X} and \mathcal{Y} are equipped with the discrete topology, then a map

$$f^1: \mathcal{X} \rightarrow \mathcal{Y} \overset{\sim}{\otimes} \mathcal{L}$$

is automatically continuous, since $\{0\} \subseteq \mathcal{X}$ is open. We observe that $F \circ G$ is the identity functor.

On the other hand, $G \circ F$ maps (\mathcal{X}, δ^1) to (\mathcal{X}, δ^1) . Note that in $\text{mod}_{\text{fg},\alpha}^{\mathcal{L}}$, the objects (\mathcal{X}, δ^1) to (\mathcal{X}, δ^1) are canonically isomorphic, since morphisms are defined to be continuous linear maps on completions. See Definition 5.4 and Remark 5.5. In particular, there is a natural transformation from $G \circ F$ to the identity. The proof is complete. \blacksquare

Chapter 8

Link surgery modules over \mathcal{K} and \mathcal{L}

In this chapter, we describe some basic link surgery modules over the algebras \mathcal{K} and \mathcal{L} . We begin by defining the type-D and type-A modules for a solid torus and the algebraic merge module. Subsequently, we interpret the knot and link surgery formulas as type-D modules over \mathcal{K} and \mathcal{L} .

8.1 The type-D module for a solid torus

If λ is an integer, we refer to the complement of a λ -framed unknot in S^3 as the λ -framed solid torus. We will define a type-D module $\mathcal{D}_\lambda^{\mathcal{K}}$ for the λ -framed solid torus as follows. We set

$$\mathcal{D}_\lambda^{\mathcal{K}} \cdot \mathbf{I}_0 = \langle \mathbf{x}^0 \rangle \quad \text{and} \quad \mathcal{D}_\lambda^{\mathcal{K}} \cdot \mathbf{I}_1 = \langle \mathbf{x}^1 \rangle,$$

where $\langle \mathbf{x}^\varepsilon \rangle = \mathbb{F}$, spanned by a generator \mathbf{x}^ε . Here, $\varepsilon \in \{0, 1\}$ denotes the idempotent. We define the structure map via the formula

$$\delta^1(\mathbf{x}^0) = \mathbf{x}^1 \otimes (\sigma + \mathcal{V}^n \tau).$$

8.2 The type-A module for a solid torus

Suppose that λ is an integer. We now define the type-A module $\mathcal{K}\mathcal{D}_\lambda$ for the solid torus as follows. We set

$$\mathbf{I}_0 \cdot \mathcal{D}_\lambda = \mathbb{F}[\mathcal{U}, \mathcal{V}] \quad \text{and} \quad \mathbf{I}_1 \cdot \mathcal{D}_\lambda = \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}].$$

We define the type-A structure map m_j on \mathcal{D}_λ to be 0 unless $j = 2$. We define m_2 on \mathcal{D}_λ as follows. If $f \in \mathbf{I}_i \cdot \mathcal{K} \cdot \mathbf{I}_i$ and $x \in \mathbf{I}_j \cdot \mathcal{D}_j$, then we define $m_2(f, x)$ to be $f \cdot x$ (ordinary multiplication of polynomials) if $i = j$ and to be 0 otherwise. If $x \in \mathbf{I}_0 \cdot \mathcal{D}_\lambda$, we define

$$m_2(\sigma, x) = \phi^\sigma(x) \in \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}] = \mathbf{I}_1 \cdot \mathcal{D}_\lambda,$$

where ϕ^σ is the canonical inclusion of localization. Similarly, we define

$$m_2(\tau, x) = \mathcal{V}^\lambda \cdot \phi^\tau(x) \in \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}],$$

where \cdot denotes ordinary multiplication of polynomials.

We view \mathcal{D}_λ as being the direct sum of the 1-dimensional spans of monomials, and give \mathcal{D}_λ the product topology with respect to this decomposition. See Definition 5.6. With respect to this topology, the completion is given by

$$\mathbf{I}_0 \cdot \mathcal{D}_\lambda = \mathbb{F}[[\mathcal{U}, \mathcal{V}]], \quad \text{and} \quad \mathbf{I}_1 \cdot \mathcal{D}_\lambda = \mathbb{F}[[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]].$$

(Note that the latter object is defined only as a vector space, and not as a ring, since we are taking completions with respect to \mathcal{V} and \mathcal{V}^{-1} .)

Lemma 8.1. *The map*

$$m_2: \mathcal{K} \overset{\sim}{\otimes}_1 \mathcal{D}_\lambda \rightarrow \mathcal{D}_\lambda$$

is continuous.

Proof. Write \mathcal{D} for \mathcal{D}_λ . Continuity amounts to two claims. The first is that if $E \subseteq \mathcal{D}$ is an open subspace, then there is some n so that $m_2(J_n \otimes \mathcal{D}) \subseteq E$. Additionally, we need to show that if E is as above and $a \in \mathcal{K}$, then there is some open $V_a \subseteq \mathcal{D}$ so that $m_2(a \otimes V_a) \subseteq E$. By definition, we may assume that E is $\mathcal{D}_{\text{co}(S)}$ for some finite set of basis elements in \mathcal{D} . (Recall that $\mathcal{D}_{\text{co}(S)}$ is the span of the basis elements not in S .)

For the first claim, it suffices to show that if n is sufficiently large, then $J_n \otimes \mathcal{D}$ is mapped into $\mathcal{D}_{\text{co}(S)}$. This may be proven by considering each idempotent of \mathcal{K} separately. To ensure $(\mathbf{I}_0 \cdot J_n \cdot \mathbf{I}_0) \otimes \mathcal{D}$ is sent into $\mathcal{D}_{\text{co}(S)}$, it is sufficient to pick n larger than $\max(p, q)$ ranging over all monomials $\mathcal{U}^p \mathcal{V}^q$ appearing in $\mathbf{I}_0 \cdot S$. For $(\mathbf{I}_1 \cdot J_n \cdot \mathbf{I}_1) \otimes \mathcal{D}$ to be sent into $\mathcal{D}_{\text{co}(S)}$, we pick n larger than p for each monomial $\mathcal{U}^p \mathcal{V}^q$ appearing in $\mathbf{I}_1 \cdot S$. Finally, to ensure that $(\mathbf{I}_1 \cdot J_n \cdot \mathbf{I}_0) \otimes \mathcal{D}$ is sent into $\mathcal{D}_{\text{co}(S)}$, it is sufficient to pick n larger than $\max(p, q, 2p - q)$, ranging over monomials $\mathcal{U}^p \mathcal{V}^q$ appearing in a summand of $\mathbf{I}_1 \cdot S$. To see that this is sufficient, note that if $\mathcal{U}^i \mathcal{V}^j \tau \otimes \mathcal{U}^s \mathcal{V}^t \in J_n \otimes \mathcal{D}$, then $\max(i, 2i - j) \geq n$. We observe that $\mathcal{U}^i \mathcal{V}^j \tau \mathcal{U}^s \mathcal{V}^t = \mathcal{U}^{i+t} \mathcal{V}^{j-s+2t} \tau$ lies in $\mathcal{D}_{\text{co}(S)}$, since

$$\max(i + t, 2(i + t) - (j - s + 2t)) = \max(i + t, 2i - j + s) \geq \max(i, 2i - j) \geq n.$$

The fact that $\mathcal{U}^i \mathcal{V}^j \sigma \otimes \mathcal{U}^s \mathcal{V}^t \in J_n \otimes \mathcal{D}$ is sent into $\mathcal{D}_{\text{co}(S)}$ is similar.

Next, we need to show that if $a \in \mathcal{K}$ is fixed, then there is a finite subset T of generators of \mathcal{D} so that $a \otimes \mathcal{D}_{\text{co}(T)}$ is mapped into $\mathcal{D}_{\text{co}(S)}$. Without loss of generality, we assume that a is a monomial. We observe the following basic fact: if $a \in \mathbf{I}_\varepsilon \cdot \mathcal{K} \cdot \mathbf{I}_\varepsilon$ and $y \in \mathbf{I}_{\varepsilon'} \cdot \mathcal{D}$ are non-zero monomials, then there is at most 1 (in particular, finitely many) $x \in \mathbf{I}_\varepsilon \cdot \mathcal{D}$ such that $m_2(a, x) = y$. In particular, there exists a finite set of generators T such that $a \otimes \mathcal{D}_{\text{co}(T)}$ is sent into $\mathcal{D}_{\text{co}(S)}$, completing the proof. ■

Remark 8.2. Note that m_2 is *not* continuous as a map from $\mathcal{K} \otimes^! \mathcal{D}$ to \mathcal{D} . Compare Remark 7.6. In fact, m_2 does not induce a map from the completion of $\mathcal{K} \otimes^! \mathcal{D}_\lambda$

to \mathcal{D}_λ . For example, in the completion of $\mathcal{K} \otimes^! \mathcal{D}_\lambda$ the sum $\sum_{i=0}^\infty \mathcal{V}^i \otimes \mathcal{V}^{-i} i_1$ converges since $\mathcal{V}^{-i} i_1 \rightarrow 0$ in the topology of \mathcal{D}_λ , while m_2 cannot be defined on such infinite sums.

The module for a solid torus extends to a type-AA bimodule

$$\mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}.$$

Here, we have U act on the right via polynomial multiplication with $\mathcal{U}\mathcal{V}$ (summed over both idempotents). Since $[\mathcal{U}\mathcal{V}, \tau] = 0$ and $[\mathcal{U}\mathcal{V}, \sigma] = 0$, the construction above defines an AA-bimodule structure on \mathcal{D}_λ . Note that only $m_{1,1,0}$ and $m_{0,1,1}$ are non-trivial.

Additionally, we can also think of the module for a solid torus as corresponding instead to a DA bimodule $\mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}$, in such a way that

$$\mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]} \hat{\boxtimes}_{\mathbb{F}[U]} \mathbb{F}[U]_{\mathbb{F}[U]} \cong \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}.$$

For this description, the generators of $\mathbf{I}_0 \cdot \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}$ are \mathcal{U}^i and \mathcal{V}^j for $i, j \geq 0$. The generators of $\mathbf{I}_1 \cdot \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}$ are \mathcal{V}^i for $i \in \mathbb{Z}$. This module has $\delta_j^1 = 0$ unless $j = 2$. We illustrate the structure map δ_2^1 of $\mathcal{K}[\mathcal{D}_0]_{\mathbb{F}[U]}$ in Figure 8.1 and leave the structure relations to the reader for other framings.




Figure 8.1. The DA bimodule $\mathcal{K}[\mathcal{D}_0]_{\mathbb{F}[U]}$. An arrow decorated with $a|b$ from \mathbf{x} to \mathbf{y} means that $\delta_2^1(a, \mathbf{x}) = \mathbf{y} \otimes b$.

8.3 The merge bimodule

In this section, we define a bimodule $\mathcal{K}_{|\mathcal{K}} M^{\mathcal{K}}$, which we call the *merge bimodule*.

Whenever R is a commutative algebra over \mathbb{F} , there is a morphism of algebras $\phi: R \otimes_{\mathbb{F}} R \rightarrow R$, given by $\phi(a \otimes b) = ab$. A ring homomorphism $\psi: A \rightarrow B$ always determines a bimodule ${}_A[\psi]^B$ whose underlying vector space is the ground ring, and has structure map $\delta_1^1 = 0$ and $\delta_2^1(a \otimes 1) = 1 \otimes \psi(a)$. In particular, there is a natural bimodule ${}_{R \otimes R}[\phi]^R$ for any commutative ring R .

The algebra \mathcal{K} is not commutative, though the bimodule ${}_{\mathcal{K}|\mathcal{K}}M^{\mathcal{K}}$ is somewhat analogous.

We now describe the bimodule M . To make the notation clear, in this section we will write $|$ for $\otimes_{\mathbb{F}}$ and \otimes for $\otimes_{\mathbb{I}}$.

As a vector space, M is $\mathbf{I} = \mathbf{I}_0 \oplus \mathbf{I}_1$. There is a left action of $\mathbf{I}|\mathbf{I}$, given by

$$(i_1|i_2) \cdot i = i_1 i_2 \cdot i.$$

The right \mathbf{I} -action is the standard action.

On M , only δ_2^1 and δ_3^1 are non-trivial. The map δ_2^1 is defined as follows. Suppose that $a_1|a_2$ is an elementary tensor in either

$$(\mathbf{I}_0|\mathbf{I}_0) \cdot (\mathcal{K}|\mathcal{K}) \cdot (\mathbf{I}_0|\mathbf{I}_0) \quad \text{or} \quad (\mathbf{I}_1|\mathbf{I}_1) \cdot (\mathcal{K}|\mathcal{K}) \cdot (\mathbf{I}_1|\mathbf{I}_1).$$

In this case, we set

$$\delta_2^1(a_1|a_2 \otimes i) = i \otimes a_1 a_2.$$

On any other elementary tensor, we set δ_2^1 to vanish.

We define δ_3^1 as follows. We set

$$\delta_3^1(1|\sigma, \sigma|1 \otimes i_0) = i_1 \otimes \sigma \quad \text{and} \quad \delta_3^1(1|\tau, \tau|1 \otimes i_0) = i_1 \otimes \tau.$$

More generally, if a, b, c, d are monomials concentrated in single idempotents, then we set

$$\delta_3^1(a|b\sigma, c\sigma|d, i_0) = i_0 \otimes abc\sigma d$$

and similarly for the τ terms. We set

$$\delta_3^1(\sigma|1, 1|\sigma \otimes i_0) = 0,$$

and similarly if τ replaces σ . The DA bimodule relations are straightforward to verify.

Lemma 8.3. *The merge module ${}_{\mathcal{K}|\mathcal{K}}M^{\mathcal{K}}$ is a split Alexander module (using the discrete partition on the two incoming algebra factors).*

Remark 8.4. The structure maps of the merge module are *not* continuous when we view the bimodule as a non-split module over $\mathcal{K} \otimes^! \mathcal{K}$ or $\mathcal{K} \otimes^* \mathcal{K}$, i.e., as ${}_{\mathcal{K} \otimes^! \mathcal{K}}M^{\mathcal{K}}$ or ${}_{\mathcal{K} \otimes^* \mathcal{K}}M^{\mathcal{K}}$. To see this, observe that discontinuity with respect to the topology on $\mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K}$ implies discontinuity with respect to $\mathcal{K} \otimes_{\mathbb{F}}^! \mathcal{K}$ since there is a continuous map $\mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K} \rightarrow \mathcal{K} \otimes_{\mathbb{F}}^! \mathcal{K}$. Note that a sequence of tensors

$$\mathbf{x}_n = (x_n|y_n) \otimes (z_n|w_n) \in (\mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K}) \otimes_{\mathbb{E}} (\mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K})$$

will converge to 0 if x_n is constant in n and $y_n \rightarrow 0$. In particular, we need not make any assumptions about z_n or w_n . We note that $\mathcal{V}^n \sigma \rightarrow 0$ in \mathcal{K} , as each of the ideals

J_m (which form a basis of open sets centered at 0 for our topology on \mathcal{K}) contains $\mathcal{V}^n \sigma$ for all sufficiently large n . Therefore,

$$(1|\mathcal{V}^n \sigma) \otimes (\mathcal{V}^{-n} \sigma|1) \rightarrow 0.$$

On the other hand,

$$\delta_3^1(1|\mathcal{V}^n \sigma, \mathcal{V}^{-n} \sigma|1, i_0) = i_1 \otimes \sigma \not\rightarrow 0,$$

so δ_3^1 is not continuous.

Proof of Lemma 8.3. We consider first δ_2^1 . In this case, the only non-trivial actions on i_0 or i_1 are from pairs $a|b$ where a and b are both in $\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$ or both in $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$. In this case, the claim is clear.

We now consider δ_3^1 . Recall that verifying the split Alexander condition amounts to showing that δ_3^1 is continuous if we use the tree topology from Section 5.6 using the following tree:

$$\begin{array}{ccc} \mathcal{K} & \longrightarrow & \mathcal{K} & \searrow & & \mathbf{I} \\ & & & & \nearrow & \\ \mathcal{K} & \longrightarrow & \mathcal{K} & \nearrow & & \end{array}$$

We define a function

$$\tilde{\mu}: \mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K} \rightarrow \mathcal{K}$$

given on elementary tensors of monomials by

$$\tilde{\mu}(a\sigma, b\sigma) = ab\sigma \quad \text{and} \quad \tilde{\mu}(a\tau, b\tau) = ab\tau,$$

with $\tilde{\mu}$ vanishing on other elementary tensors of monomials. Additionally, we consider a map

$$\Pi: (\mathcal{K} \otimes_{\mathbf{I}} \mathcal{K}) \otimes_{\mathbb{F}} (\mathcal{K} \otimes_{\mathbf{I}} \mathcal{K}) \rightarrow (\mathcal{K} \otimes_{\mathbf{I}} \mathcal{K}) \otimes_{\mathbb{F}} (\mathcal{K} \otimes_{\mathbf{I}} \mathcal{K})$$

such that Π is the identity on the elementary tensors of monomials

$$(a\sigma \otimes a')|(b \otimes b'\sigma) \quad \text{and} \quad (a\tau \otimes a')|(b \otimes b'\tau),$$

and such that Π vanishes on other elementary tensors of monomials.

We observe that

$$\delta_3^1(\mathbf{a}, 1) = 1 \otimes (\tilde{\mu} \circ (\mu_2|\mu_2) \circ \Pi)(\mathbf{a}).$$

The map Π is clearly continuous with respect to the tree topology, and furthermore, $\mu_2|\mu_2$ is continuous by Propositions 7.5 and 5.24. Note that if we topologize $(\mathcal{K}|\mathcal{K}) \otimes_{\mathbf{I}|\mathbf{I}} \mathbf{I}$ with the tree topology, it is isomorphic to $\mathcal{K} \otimes_{\mathbb{F}}^* \mathcal{K}$. (See Remark 5.23.) Hence, it suffices to show that $\tilde{\mu}$ is continuous. This amounts to the following:

(m -1) For each $n \in \mathbb{N}$, there are $k, m \in \mathbb{N}$ so that

$$\tilde{\mu}(J_k \otimes J_m) \subseteq J_n.$$

(m-2) For each $n \in \mathbb{N}$ and $x \in \mathcal{K}$, there is an $m \in \mathbb{N}$ so that

$$\tilde{\mu}(x \otimes J_m) \subseteq J_n \quad \text{and} \quad \tilde{\mu}(J_m \otimes x) \subseteq J_n.$$

Both of the arguments are elementary, and we illustrate the idea. Recall that by our construction in Definition 7.3, $\mathcal{U}^i \mathcal{V}^j \sigma \in J_n$ if and only if $i \geq 0$ and $\max(i, j) \geq n$. If $\mathcal{U}^i \mathcal{V}^j \sigma$ and $\mathcal{U}^s \mathcal{V}^t \sigma$ are two monomials in J_n , then at least one of two situations occurs: both $j, t \geq n$; or one of i and s is at least n . If $j, t \geq n$, then $\mathcal{U}^{i+s} \mathcal{V}^{j+t} \sigma \in J_n$ since $j + t \geq 2n$. If one of i and s is at least n , then $i + s \geq n$. A symmetric analysis holds for the multiples of τ , and we see that

$$\tilde{\mu}(J_n \otimes J_n) \subseteq J_n,$$

so (m-1) holds. A similar and elementary argument implies (m-2). ■

There is an asymmetry in the merge module. Recall that we set

$$\delta_3^1(1|\sigma, \sigma|1, i_0) = i_1 \otimes \sigma \quad \text{and} \quad \delta_3^1(\sigma|1, 1|\sigma, i_0) = 0.$$

We define a different module ${}_{\mathcal{K}|\mathcal{K}}N^{\mathcal{K}}$ by instead setting

$$\delta_3^1(1|\sigma, \sigma|1, i_0) = 0 \quad \text{and} \quad \delta_3^1(\sigma|1, 1|\sigma, i_0) = i_1 \otimes \sigma.$$

We make a similar change to the map δ_3^1 on τ inputs.

This construction turns out to produce homotopy equivalent bimodules.


Lemma 8.5. *There is a homotopy equivalence of DA bimodules*

$${}_{\mathcal{K}|\mathcal{K}}M^{\mathcal{K}} \simeq {}_{\mathcal{K}|\mathcal{K}}N^{\mathcal{K}}.$$

Proof. We construct a morphism f_{j+1}^1 between the two bimodules. We set $f_1^1(\mathbf{x}) = \mathbf{x} \otimes 1$. The only other non-vanishing term is f_2^1 . This map is determined by the relations

$$f_2^1(\sigma|\sigma, i_0) = i_1 \otimes \sigma \quad \text{and} \quad f_2^1(\tau|\tau, i_0) = i_1 \otimes \tau.$$

The DA bimodule morphism structure relations for two algebra inputs which both change idempotent are given by the schematic



which is easily verified. The structure relations for $n \neq 2$ algebra inputs are also easily verified. ■

8.4 The type-A identity bimodule

It is helpful to introduce another bimodule, which we denote by

$$\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus] := \mathcal{K}|\mathcal{K}M^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{D}_0.$$

We call $\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]$ the *type-A identity module*. We will use $\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]$ to transform type-D actions into type-A actions. The type-A module $\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]$ appears in the pairing theorem, stated in Theorem 1.4.

There is a further refinement which takes into account the right $\mathbb{F}[U]$ -action on \mathcal{D}_0 . We define

$$\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]_{\mathbb{F}[U]} := \mathcal{K}|\mathcal{K}M^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}[\mathcal{D}_0]_{\mathbb{F}[U]}.$$

Verifying the pairing theorem with the bimodule $\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]_{\mathbb{F}[U]}$ turns out to be more subtle than for $\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]$. Nonetheless, we prove the pairing theorem with $\mathcal{K}|\mathcal{K}[\mathbb{I}^\ominus]_{\mathbb{F}[U]}$ in Section 15.3 once we introduce the pair-of-pants bimodules.

8.5 The knot surgery formula over \mathcal{K}

We now describe how to view the mapping cone formula of Ozsváth and Szabó in terms of the algebra \mathcal{K} .

Recall that Ozsváth and Szabó's mapping cone formula [43, Theorem 1.1] states that if λ is an integral framing on $K \subseteq S^3$, then

$$\mathbf{CF}^-(S_\lambda^3(K)) \cong \mathbb{X}_\lambda(K) = \text{Cone}(\mathbb{A}(K) \xrightarrow{v+h_\lambda} \mathbb{B}(K)).$$

Recall, as in Lemma 6.5, that $\mathbb{A}(K)$ may be identified with a completion of $\mathcal{CFK}(K)$ and $\mathbb{B}(K)$ may be identified with a completion of $\mathcal{V}^{-1}\mathcal{CFK}(K)$.

We will describe the following algebraic perspectives of the mapping cone formula:

- (1) A chain complex $\mathbb{X}_\lambda(K)$ over $\mathbb{F}[U]$.
- (2) A right type-D module $\mathcal{X}_\lambda(K)^{\mathcal{K}}$.
- (3) A left type-A module ${}_{\mathcal{K}}\mathcal{X}_\lambda(K)$.

These will satisfy the following relations:

$$\mathbb{X}_\lambda(K) \cong \mathcal{X}_\lambda(K)^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}[\mathcal{D}_0] \simeq \mathcal{D}_0^{\mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{X}_\lambda(K).$$

We begin with the description as a type-D module. Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be a free basis of $\mathcal{CFK}(K)$ over $\mathbb{F}[U, V]$. We declare the underlying \mathbf{I} -module of $\mathcal{X}_\lambda(K)^{\mathcal{K}}$ to be

$$\mathcal{X}_\lambda(K)^{\mathcal{K}} = \text{Span}_{\mathbb{F}}(\mathbf{x}_1, \dots, \mathbf{x}_n) \otimes_{\mathbb{F}} \mathbf{I}.$$

In particular, over \mathbb{F} each \mathbf{x}_i contributes one generator in each idempotent. We denote these generators by \mathbf{x}_i^0 and \mathbf{x}_i^1 .

The map δ^1 on $\mathcal{X}_\lambda(K)^{\mathcal{K}}$ is as follows. Firstly, there are internal δ^1 summands from the differential of $\mathcal{CFK}(K)$. If $\partial(\mathbf{x})$ contains a summand of $\mathbf{y} \cdot \mathcal{U}^n \mathcal{V}^m$, then $\delta^1(\mathbf{x}^\varepsilon)$ contains a summand of $\mathbf{y}^\varepsilon \otimes \mathcal{U}^n \mathcal{V}^m$, for $\varepsilon \in \{0, 1\}$.

In Ozsváth and Szabó's mapping cone formula, v is the canonical inclusion of $\mathcal{CFK}(K)$ into $\mathcal{V}^{-1}\mathcal{CFK}(K)$. Correspondingly, $\delta^1(\mathbf{x}^0)$ contains a summand of the form $\mathbf{x}^1 \otimes \sigma$.

If $\mathbf{y} \cdot \mathcal{U}^i \mathcal{V}^j$ is a summand of $h_\lambda(\mathbf{x})$ then we define $\delta^1(\mathbf{x}^0)$ to have a summand of the form

$$\mathbf{y}^1 \otimes \mathcal{U}^i \mathcal{V}^j \tau.$$

Lemma 8.6. $\mathcal{X}_\lambda(K)^{\mathcal{K}}$ is a type-D module.

Proof. The proof is a formal consequence of the fact that $\partial^2 = 0$ on $\mathcal{CFK}(K)$, that v and h_λ are chain maps, and that v and h_λ satisfy the equivariance properties of Lemma 6.7.

In more detail, suppose that $\mathbf{x}_1, \dots, \mathbf{x}_n$ is a free basis of $\mathcal{CFK}(K)$, and consider

$$((\text{id}_{\mathcal{X}} \otimes \mu_2) \circ (\delta^1 \otimes \text{id}_{\mathcal{K}}) \circ \delta^1)(\mathbf{x}_i^\varepsilon) \quad (8.1)$$

For concreteness, consider the case when $\varepsilon = 0$. There are three sources of terms in the above expression: those arising from two applications of ∂ , those arising from one v and one ∂ , and those arising from one h_λ and one ∂ . Consider the terms with two ∂ terms. Write $\partial(\mathbf{x}_i) = \sum_{j=1}^n \mathbf{x}_j \cdot f_{j,i}$. The terms with two ∂ terms are simply

$$\sum_{k=1}^n \mathbf{x}_k^0 \otimes \left(\sum_{j=1}^n f_{k,j} f_{j,i} \right),$$

which is 0 since $\partial^2 = 0$ on $\mathcal{CFK}(K)$. Similarly, the terms corresponding to one ∂ and one v are given by the formula

$$\sum_{j=1}^n \mathbf{x}_j^1 \otimes (\sigma f_{j,i} + f_{j,i} \sigma),$$

which vanishes because $\sigma f_{j,i} = f_{j,i} \sigma$ in \mathcal{K} .

Finally, we consider the terms corresponding to one ∂ and one h_λ . Write $h_\lambda(\mathbf{x}_i) = \sum_{j=1}^n \mathbf{x}_j \cdot g_{j,i}$, where $g_{j,i} \in \mathbb{F}[\mathcal{U}, \mathcal{V}, \mathcal{V}^{-1}]$. The corresponding terms of equation (8.1) are

$$\sum_{k=1}^n \mathbf{x}_k^1 \otimes \left(\sum_{j=1}^n g_{k,j} \tau f_{j,i} + f_{k,j} g_{j,i} \tau \right) = \sum_{k=1}^n \mathbf{x}_k^1 \otimes \left(\sum_{j=1}^n g_{k,j} \phi^\tau(f_{j,i}) + f_{k,j} g_{j,i} \right) \tau. \quad (8.2)$$

However, $\sum_{j=1}^n g_{k,j} \phi^\tau(f_{j,i}) + f_{k,j} g_{j,i}$ is the \mathbf{x}_k coefficient of $[\partial, h_\lambda](\mathbf{x}_i)$, since h_λ satisfies $h_\lambda(\mathbf{x} \cdot a) = h_\lambda(\mathbf{x}) \cdot \phi^\tau(a)$, by Lemma 6.7. In particular, equation (8.2) vanishes. The case when $\varepsilon = 1$ is similar. \blacksquare

Having defined the type-D module $\mathcal{X}_\lambda(K)^{\mathcal{K}}$, we may now define the type-A module

$${}_{\mathcal{K}} \mathcal{X}_\lambda(K) := \mathcal{X}_\lambda(K)^{\mathcal{K}} \widehat{\boxtimes}_{\mathcal{K}|\mathcal{K}} [\mathbb{I}^\boxplus].$$

8.6 The link surgery formula over \mathcal{L}

In this section, we describe how to view Manolescu and Ozsváth’s link surgery formula as a type-D module.

Recall that we define

$$\mathcal{L}_\ell := \mathcal{K}_1 \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathcal{K}_\ell,$$

where each \mathcal{K}_i denotes a copy of \mathcal{K} .

Let L be an ℓ -component link in S^3 with integral framing Λ . The link surgery complexes require a choice of auxiliary data, which we call a *system of arcs* \mathcal{A} . We will discuss these in detail in Section 9.1 below. If \mathcal{A} is chosen, Manolescu and Ozsváth’s construction produces a chain complex $\mathcal{C}_\Lambda(L, \mathcal{A})$ over $\mathbb{F}[U_1, \dots, U_\ell]$. In this section, we will describe how to construct a type-D module

$$\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}_\ell},$$

based on their construction of $\mathcal{C}_\Lambda(L, \mathcal{A})$.

We view the link algebra \mathcal{L}_ℓ as being an algebra over the idempotent ring

$$\mathbf{E}_\ell := \underbrace{\mathbf{I} \otimes_{\mathbb{F}} \cdots \otimes_{\mathbb{F}} \mathbf{I}}_\ell.$$

We can view \mathbf{E}_ℓ as being a ring with 2^ℓ elementary idempotents, each identified with a point in the cube \mathbb{E}_ℓ . If $\varepsilon \in \mathbb{E}_\ell$, we write

$$\mathbf{E}_\ell := \mathbf{I}_{\varepsilon_1} \otimes \cdots \otimes \mathbf{I}_{\varepsilon_\ell} \cong \mathbb{F}$$

for the corresponding idempotent.

If $\varepsilon' \geq \varepsilon$, we can understand $\mathbf{E}_{\varepsilon'} \cdot \mathcal{L}_\ell \cdot \mathbf{E}_\varepsilon$ as follows. Write $I_{\varepsilon', \varepsilon} = \{i_1, \dots, i_n\} = (\varepsilon' - \varepsilon)^{-1}(1)$, i.e., the set of indices i where $\varepsilon'_i > \varepsilon$. Then $\mathbf{E}_{\varepsilon'} \cdot \mathcal{L}_\ell \cdot \mathbf{E}_\varepsilon$ is generated by elements of the form

$$\alpha_1 \cdots \alpha_\ell \phi_{i_1} \cdots \phi_{i_n},$$

where

- (1) a_i is in $\mathbb{F}[\mathcal{U}_i, \mathcal{V}_i]$ if $\varepsilon'_i = 0$,
- (2) a_i is in $\mathbb{F}[\mathcal{U}_i, \mathcal{V}_i, \mathcal{V}_i^{-1}]$ if $\varepsilon'_i = 1$,
- (3) each ϕ_{i_j} is either σ_{i_j} or τ_{i_j} .

If $\mathbf{x}_1, \dots, \mathbf{x}_n$ is a free basis of $\mathcal{CF}\mathcal{L}(L)$ over $\mathbb{F}[\mathcal{U}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell]$, then we define the generators of $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}\ell}$ to be

$$\text{Span}_{\mathbb{F}}(\mathbf{x}_1, \dots, \mathbf{x}_n) \otimes_{\mathbb{F}} \mathbf{E}_\ell.$$

In particular, if \mathbf{x} is a basis element of $\mathcal{CF}\mathcal{L}(L)$, we have a generator \mathbf{x}^ε for each $\varepsilon \in \mathbb{E}_\ell$.

If $M \subseteq L$, write $\varepsilon(M) \in \mathbb{E}_\ell$ for the coordinate such that $\varepsilon(M)_i = 0$ if $L_i \notin M$ and $\varepsilon(M)_i = 1$ if $L_i \in M$.

Suppose that \mathbf{x} is a basis element of $\mathcal{CF}\mathcal{L}(L)$ and $M \subseteq L$. By Lemma 6.5, we may view $\mathbf{x}^{\varepsilon(M)}$ as an element of the group $\mathcal{C}_\Lambda(L, \mathcal{A}) = \prod_{\mathbf{s} \in \mathbb{H}(L)} \mathfrak{X}(\mathcal{H}^{L \setminus M}, \psi^M(\mathbf{s}))$. Suppose also that \vec{N} is an oriented sublink of $L \setminus M$, and that $\Phi^{\vec{N}}(\mathbf{x}^{\varepsilon(M)})$ has a summand of $\mathbf{y}^{\varepsilon(M \cup N)} \cdot f$, where we view f as an element of the 2ℓ -variable polynomial ring, localized at the variables \mathcal{V}_i such that $L_i \subseteq N$. We may naturally view f as being an element of

$$\mathbf{E}_{\varepsilon(M \cup N)} \cdot \mathcal{L}_\ell \cdot \mathbf{E}_{\varepsilon(M \cup N)}.$$

There is an algebra element

$$t_{\varepsilon(M), \varepsilon(M \cup N)}^{\vec{N}} \in \mathbf{E}_{\varepsilon(M \cup N)} \cdot \mathcal{L}_\ell \cdot \mathbf{E}_{\varepsilon(M)}$$

which is the tensor of σ_i for i such that $L_i \subseteq \vec{N}$ and L_i is oriented the same as L , τ_i for i such that $L_i \subseteq \vec{N}$ and L_i is oriented oppositely from L , and 1 for i such that $L_i \notin N$. With this notation, we declare $\delta^1(\mathbf{x}^{\varepsilon(M)})$ to have the summand

$$\mathbf{y}^{\varepsilon(M \cup N)} \otimes f \cdot t_{\varepsilon(M), \varepsilon(M \cup N)}^{\vec{N}}.$$

Lemma 8.7. $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}\ell}$ is a type-D module.

Proof. The proof is similar to Lemma 8.6. We can decompose the structure map δ^1 as

$$\delta^1 = \sum_{\vec{M} \subseteq L} \delta_{\vec{M}}^1,$$

where $\delta_{\vec{M}}^1$ consists of all terms which are weighted by idempotent-preserving multiples of the algebra element $t^{\vec{M}}$.

The map $\delta_{\vec{M}}^1$ encodes the map $\Phi^{\vec{M}}$ on the surgery formula, in the sense that if $\Phi^{\vec{M}}(\mathbf{x})$ has a summand $a \cdot \mathbf{y}$ for some monomial a in the \mathcal{U}_i and $\mathcal{V}_i^{\pm 1}$, then $\delta_{\vec{M}}^1$ has a summand $\mathbf{y} \otimes at^{\vec{M}}$.

Let $\vec{N} \subseteq L$ be an oriented link, and consider the components of

$$((\text{id}_{\mathcal{X}} \otimes \mu_2) \circ (\delta^1 \otimes \text{id}_{\mathcal{X}_\ell}) \circ \delta^1)(\mathbf{x}^\varepsilon)$$

whose algebra elements are weighted by idempotent-preserving multiples of $t^{\vec{N}}$. These can be identified with

$$\sum_{\substack{\vec{M}_1 \cup \vec{M}_2 = \vec{N}, \\ \vec{M}_1 \cap \vec{M}_2 = \emptyset}} (\mathbb{I}_{\mathcal{X}} \otimes \mu_2) \circ (\delta_{\vec{M}_2}^1 \otimes \mathbb{I}_{\mathcal{X}}) \circ \delta_{\vec{M}_1}^1. \quad (8.3)$$

We consider a component of the above map corresponding to sublinks \vec{M}_1 and \vec{M}_2 , applied to some generator \mathbf{x} . Suppose that $\delta_{\vec{M}_1}^1(\mathbf{x})$ contains a summand $\mathbf{y} \otimes at^{\vec{M}_1}$. Such a summand corresponds to a summand of $\mathbf{y} \cdot a$ in $\Phi^{\vec{M}_1}(\mathbf{x})$. Next, consider a summand $\mathbf{z} \otimes bt^{\vec{M}_2}$ in $\delta_{\vec{M}_2}^1(\mathbf{y})$. In the composition appearing in equation (8.3), there is a corresponding term of

$$\mathbf{z} \otimes bt^{\vec{M}_2} at^{\vec{M}_1} = \mathbf{z} \otimes b\phi^{\vec{M}_2}(a)t^{\vec{N}}.$$

(The second equality follows from the definition of the algebra \mathcal{L}_ℓ .) Using Lemma 6.7, which states that $\Phi^{\vec{M}_2}(a\mathbf{y}) = \phi^{\vec{M}_2}(a) \cdot \Phi^{\vec{M}_2}(\mathbf{y})$, we see that there is a corresponding summand in $(\Phi^{\vec{M}_2} \circ \Phi^{\vec{M}_1})(\mathbf{x})$ of $\mathbf{z} \cdot b\phi^{\vec{M}_2}(a)$. Therefore, the fact equation (8.3) vanishes follows from the fact that

$$\sum_{\substack{\vec{M}_1 \cup \vec{M}_2 = \vec{N}, \\ \vec{M}_1 \cap \vec{M}_2 = \emptyset}} \Phi^{\vec{M}_2} \circ \Phi^{\vec{M}_1} = 0,$$

which is proven in [32, Proposition 9.4]. ■

Given a collection of integers $\Lambda = (\lambda_1, \dots, \lambda_\ell)$, there is a type-A module

$$\mathcal{L}_\ell \mathcal{D}_\Lambda$$

defined by taking the external tensor product of the modules $\mathcal{K}[\mathcal{D}_{\lambda_i}]$ from Section 8.2. There is also a right action of $\mathbb{F}[U_1, \dots, U_\ell]$, and \mathcal{D}_Λ may be viewed as an AA-bimodule

$$\mathcal{L}_\ell [\mathcal{D}_\Lambda]_{\mathbb{F}[U_1, \dots, U_\ell]}.$$

The following is immediate from the definition of $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}_\ell}$.

Proposition 8.8. *There is a chain isomorphism*

$$\mathcal{C}_\Lambda(L)_{\mathbb{F}[U_1, \dots, U_\ell]} \cong \mathcal{X}_\Lambda(L)^{\mathcal{L}_\ell} \hat{\boxtimes} \mathcal{L}_\ell [\mathcal{D}_{(0, \dots, 0)}]_{\mathbb{F}[U_1, \dots, U_\ell]}.$$

We will also need to consider type-A and type-DA versions of the link surgery formula. An important special case is when we have a distinguished component $K_i \subseteq L$. We define a bimodule ${}_{\mathcal{K}}\mathcal{X}_{\Delta}(L)^{\mathcal{L}^{\ell-1}}$, where the \mathcal{K} -input corresponds to K_i by setting

$${}_{\mathcal{K}}\mathcal{X}_{\Delta}(L)^{\mathcal{L}^{\ell-1}} := \mathcal{X}_{\Delta}(L)^{\mathcal{L}^{\ell}} \hat{\boxtimes} {}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}^{\mathfrak{D}}],$$

where the tensor product is taken along the algebra component corresponding to K_i . In the above equation, ${}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}^{\mathfrak{D}}]$ is the bimodule from Section 8.4. More generally, we may turn more algebra components from type-D to type-A by tensoring additional copies of ${}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}^{\mathfrak{D}}]$.

Chapter 9

σ -basic systems of Heegaard diagrams

In this chapter, we define the notion of a σ -basic system of Heegaard diagrams. The notion is a very small generalization of Manolescu and Ozsváth's *basic system*. These are collections of Heegaard diagrams which can be used to compute the link surgery formula. We use the small distinction in notation to disambiguate the notions, and to highlight the algebraic significance of these systems. For these systems, the maps $\Phi^{\vec{N}}$ vanish unless $\vec{N} = +K$ (i.e., a single component, oriented consistently with L), or \vec{N} contains only negatively oriented components relative to L . We also describe how to construct a σ -basic system for the connected sum of two links.

9.1 Systems of arcs

In this section, we define the notion of a *system of arcs*, which is a necessary piece of auxiliary data used to define the link surgery formula.

Definition 9.1. Suppose that $(S^3, L, \mathbf{w}, \mathbf{z})$ is an oriented, multi-pointed link, which is link minimal (i.e., each component of L contains exactly one base point of \mathbf{w} , and one base point of \mathbf{z}).

- (1) A *system of arcs* for L consists of a set $\mathcal{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_\ell\}$ of pairwise disjoint arcs satisfying that $\partial\mathcal{A}_i = \mathcal{A}_i \cap L = \{w_i, z_i\}$.
- (2) We say that an arc \mathcal{A}_i is *beta-parallel* if \mathcal{A}_i is isotopic to a push-off of the subarc of L_i which goes from z_i to w_i .
- (3) We say that \mathcal{A}_i is *alpha-parallel* if \mathcal{A}_i is isotopic to a push-off of the subarc of L_i which goes from w_i to z_i .

Note that an arc \mathcal{A}_i is beta-parallel if for any Heegaard diagram of $(Y, L, \mathbf{w}, \mathbf{z})$, \mathcal{A}_i is isotopic to the subarc of $K_i \subseteq L$ which lies in the beta handlebody U_β . A similar remark holds for alpha-parallel arcs.

Given a link minimal $(L, \mathbf{w}, \mathbf{z})$ there are two distinguished systems of arcs, \mathcal{A}_α and \mathcal{A}_β , which are the ones which have only alpha-parallel arcs, or only beta-parallel arcs, respectively.

Remark 9.2. Manolescu and Ozsváth's notion of *good sets of trajectories* [32, Definition 8.26] corresponds to our notion of a system of alpha-parallel arcs.

9.2 σ -basic systems of Heegaard diagrams

In this section, we define the notion of a σ -basic system of Heegaard diagrams. Our notion is a small adaptation of Manolescu and Ozsváth's construction.

Definition 9.3. Suppose that $(L, \mathbf{w}, \mathbf{z})$ is a minimally pointed, n -component link in a 3-manifold Y , and suppose that $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_\ell)$ is a system of arcs for $(Y, L, \mathbf{w}, \mathbf{z})$. A σ -basic system of Heegaard diagrams for (Y, L, \mathcal{A}) consists of the following data:

- (1) A tuple of positive integers $\mathbf{d} = (d_1, \dots, d_\ell)$ where $n = |\mathbf{d}|$.
- (2) A collection of Heegaard diagrams $\mathcal{H}(M) = (\mathcal{H}_\varepsilon)_{\varepsilon \in \mathbb{E}(\mathbf{d})}$. We assume all Heegaard diagrams have the same underlying Heegaard surface Σ , which we also assume contains the arcs $\mathcal{A}_1, \dots, \mathcal{A}_\ell$. We write $\mathcal{H} = (\Sigma, \alpha_\varepsilon, \beta_\varepsilon, \mathbf{w}_\varepsilon, \mathbf{z}_\varepsilon, \mathbf{p}_\varepsilon)$, where \mathbf{w}_ε and \mathbf{z}_ε are link base points and \mathbf{p}_ε are free base points.
- (3) For each component $i \in \{1, \dots, n\}$, we have a formal labeling of each of the intervals $[0, 1], \dots, [d_i - 1, d_i]$ as being either an *alpha incrementing interval*, a *beta incrementing interval* or a *surface isotopy interval*.
- (4) A decomposition of each arc \mathcal{A}_j into a concatenation of oriented subarcs $l_{i,1}, \dots, l_{i,j_i}$, as well as a choice of surface isotopies $\phi_{i,1}, \dots, \phi_{i,j_i}$. We assume $\phi_{i,j}$ is supported in a small neighborhood of $l_{i,j}$ and moves the initial endpoint of $l_{i,j}$ to the terminal endpoint.

We assume that the following conditions are satisfied:

- (1) $\mathbf{w}_\varepsilon = \{w_i : \varepsilon_i = 0\}$ and $\mathbf{z}_\varepsilon = \{z_i : \varepsilon_i = 0\}$.
- (2) Let $\mathcal{H}_\varepsilon = (\Sigma, \alpha_\varepsilon, \beta_\varepsilon, \mathbf{w}_\varepsilon, \mathbf{z}_\varepsilon, \mathbf{p}_\varepsilon)$ and let $e_i = (0, \dots, 1, \dots, 0) \in \{0, 1\}^n$ denote a length 1 vector (i.e., a direction). As described above, the interval $I = [\varepsilon_i, \varepsilon_i + e_i]$ is labeled as either an alpha incrementing interval, a beta incrementing interval, or a surface isotopy interval.
 - (a) If I is an alpha incrementing interval, then $\beta_\varepsilon = \beta_{\varepsilon+e_i}$ while $\alpha_{\varepsilon+e_i}$ and α_ε differ. Furthermore, $\mathbf{p}_{\varepsilon+e_i} = \mathbf{p}_\varepsilon$ unless $\varepsilon_i = 0$, in which case $\mathbf{p}_{\varepsilon+e_i} = \mathbf{p}_\varepsilon \cup \{z_i\}$.
 - (b) If I is a beta incrementing interval, then the same holds as above except with the alpha and beta curves interchanged.
 - (c) If I is a surface isotopy interval, then $\mathcal{H}_{\varepsilon+e_i}$ is the image of \mathcal{H}_ε under one of the surface isotopies $\phi_{i,j}$ described above. Similarly, $\mathbf{w}_{\varepsilon+e_i} = \mathbf{w}_\varepsilon$ and $\mathbf{z}_{\varepsilon+e_i} = \mathbf{z}_\varepsilon$. If $\varepsilon_i > 0$, then $\mathbf{p}_{\varepsilon+e_i} = \phi_{i,j}(\mathbf{p}_\varepsilon)$. If $\varepsilon_i = 0$, then $\mathbf{p}_{\varepsilon+e_i} = \phi_{i,j}(\mathbf{p}_\varepsilon \cup \{z_i\})$.
- (3) The faces on opposite sides of the box have the same Heegaard diagrams. In more detail, if $\varepsilon \in \mathbb{E}(\mathbf{d})$ and $\varepsilon_i = 0$, then

$$\alpha_{\varepsilon+d_i e_i} = \alpha_\varepsilon, \quad \beta_{\varepsilon+d_i e_i} = \beta_\varepsilon, \quad \mathbf{p}_{\varepsilon+d_i e_i} = \mathbf{p}_\varepsilon \cup \{w_i\}.$$

We illustrate an example of a σ -basic system in Figure 9.1.




Figure 9.1. An example of a σ -basic system of Heegaard diagrams for the Hopf link. At the top, we show the underlying Heegaard link diagram of the Hopf link, with an arc system consisting of two beta-parallel arcs \mathcal{A}_1 and \mathcal{A}_2 (shown as dashed arcs). For this example, the subdivisions of \mathcal{A}_1 and \mathcal{A}_2 are trivial (i.e., consist of a single arc). The labels of the arrows indicate whether the corresponding subintervals are “alpha incrementing intervals”, “beta incrementing intervals” or “surface isotopies intervals”. We label the \mathbf{w} base points as solid circles, the \mathbf{z} base points as open circles, and the free base points as squares.

Remark 9.4. The terminology σ -basic system refers to the fact that when we use such systems of Heegaard diagrams to construct the link surgery formula, the resulting complexes have Φ^{+K_i} equal to inclusions for localization, and $\Phi^{\vec{M}} = 0$ if \vec{M} has more than one component and contains a positively oriented component. In particular, using these diagrams we obtain models of the type-D modules $\mathcal{X}_\Delta(L)^{\mathcal{L}^\ell}$ where the components weighted by σ_i are very simple. Our definition is slightly more general than Manolescu and Ozsváth’s notion of a basic system of Heegaard diagrams, which has the same algebraic properties with respect to the differential.

The following definition plays an important role in our proof of the connected sum formula for surgery hypercubes.

Definition 9.5. Suppose that $\mathcal{L}_\beta = (\beta_\varepsilon, \theta_{\varepsilon, \varepsilon'})_{\varepsilon \in \mathbb{E}_n}$ is a weakly admissible hypercube of handleslide equivalent beta attaching curves on a pointed Heegaard surface $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ (where \mathbf{w}, \mathbf{z} are link base points and \mathbf{p} are free base points). We say \mathcal{L}_β is *algebraically rigid* if each $\mathbb{T}_{\beta_\varepsilon} \cap \mathbb{T}_{\beta_{\varepsilon'}}$ has $2^{g(\Sigma) + |\mathbf{w}| + |\mathbf{p}| - 1}$ intersection points (the minimal possible number). We make a parallel definition for alpha hyperboxes. We say a hyperbox of attaching curves is algebraically rigid if each sub-hypercube is algebraically rigid.

Note that a hypercube of attaching curves being algebraically rigid is a property of the attaching curves, and has nothing to do with the chains $\theta_{\varepsilon, \varepsilon'}$. Therefore, abusing notation slightly, we will say a σ -basic system of Heegaard diagrams \mathcal{H} is algebraically rigid if it has the property that $\mathbb{T}_{\beta_\varepsilon} \cap \mathbb{T}_{\beta_{\varepsilon'}}$ has $2^{g(\Sigma) + |\mathbf{w}_\varepsilon| + |\mathbf{p}_\varepsilon|}$ intersection points whenever $|\varepsilon' - \varepsilon|_{L^\infty} \leq 1$ and none of the intervals $[\varepsilon, \varepsilon + e_i]$ is a surface isotopy interval for i such that $\varepsilon_i < \varepsilon'_i$. We make the same assumption for the alpha curves.

9.3 Meridional σ -basic systems

In this section, we describe a particular family of σ -basic systems which are useful in our proof of the pairing system. Similar diagrams are ubiquitous in Heegaard Floer theory. For example, they are the diagrams which Ozsváth and Szabó use to prove the mapping cone formula [43].

Definition 9.6. We say that a σ -basic system of Heegaard diagrams \mathcal{H} for (Y, L) , with a system of arcs \mathcal{A} , is a *meridional σ -basic system* if the following hold:

- (1) Each arc of \mathcal{A} is either beta-parallel or alpha-parallel.
- (2) The diagram $\mathcal{H}(\emptyset) = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ has the following property. Suppose $K_i \in L$ is a component with corresponding arc $\mathcal{A}_i \in \mathcal{A}$ and base points w_i, z_i . Then \mathcal{A}_i is either alpha-parallel or beta-parallel. If \mathcal{A}_i is alpha-parallel, then \mathcal{A}_i is disjoint from the alpha curves, and intersects a single beta curve β_i^s .

A similar assumption is made if \mathcal{A}_i is beta-parallel. In particular, if $|L| = \ell$, there are n special alpha and beta circles, denoted α_i^s and β_j^s .

- (3) In each Heegaard diagram of \mathcal{H} , there are similarly ℓ special alpha and beta curves, and the rest are designated as non-special. The non-special alpha curves of each subcube of \mathcal{H} consist of small Hamiltonian translates of the non-special curves of \mathcal{H} .
- (4) If $K_i \subseteq L$ is a component whose arc \mathcal{A}_i is beta-parallel, then the component of $\Sigma \setminus \beta$ which contains the base points of K_i is a punctured disk which contains a special beta circle β_i^s as two of its boundary components. Gluing the two β_i^s boundary components together, we obtain a punctured torus. We assume that K_i -axis direction of \mathcal{H} realizes the isotopy of β_i^s in a loop around this punctured torus (handlesliding β_i^s over the boundary punctures). We make an analogous assumption for components with beta-parallel decoration.

Lemma 9.7. *Suppose L is an ℓ -component link in S^3 and \mathcal{A} is a system of arcs for L such that each arc is either alpha-parallel or beta-parallel. Then there exists a link minimal, weakly admissible, algebraically rigid σ -basic system of Heegaard diagrams for (S^3, L, \mathcal{A}) .*

Proof. Except for the claim about admissibility, existence of such a σ -basic system is clear since the construction is given in Definition 9.6.

Admissibility is verified as follows. We first verify admissibility for each hypercube of alpha and beta Lagrangians appearing in the construction. For concreteness, suppose that \mathcal{L}_α is a hypercube of alpha attaching curves appearing in the construction. For appropriately chosen Hamiltonian translates, admissibility of \mathcal{L}_α is equivalent to admissibility of the diagram where we delete the tuples of curves which appear only as small translates of each other, and surger Σ along these curves as well. The resulting Heegaard diagram is a union of tori, and admissibility is straightforward to arrange. See Figure 9.2.

Admissibility of each pair $(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ may be arranged by winding one of the cubes (say \mathcal{L}_α) relative to \mathcal{L}_β , following the procedure described in [39, Section 4.2.2] and [41, Section 3.4]. ■

9.4 Basic systems for general systems of arcs

We now consider general systems of arcs. Generalizing Lemma 9.7, we prove the following.

Lemma 9.8. *Suppose that \mathcal{A} is a system of arcs for $L \subseteq S^3$. Then there is an algebraically rigid σ -basic system of Heegaard diagrams for (L, \mathcal{A}) .*




Figure 9.2. Some of the beta curves and a base point translation from Lemma 9.7. The first arrow is a base point translation, and the remainder are holomorphic polygon counting maps. The dashed arc on the top is c_i .

Proof. The proof is similar to the proof of Lemma 9.7. We begin with a minimally pointed Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ for (Y, L) such that each arc \mathcal{A}_i is embedded in Σ and such that the arcs \mathcal{A}_i are pairwise disjoint. By performing an isotopy of α and β supported in a neighborhood of the arcs \mathcal{A}_i , we may assume that each \mathcal{A}_i is the concatenation of two arcs A_i^α and A_i^β such that A_i^α is disjoint from β and A_i^β is disjoint from α . We assume that $\partial A_i^\beta = \{z_i, x_i\}$ and $\partial A_i^\alpha = \{w_i, x_i\}$, for some point $x_i \in \Sigma$.

We now stabilize the Heegaard diagram $2|L|$ times, as shown in Figure 9.3. For each \mathcal{A}_i , we introduce two stabilizations. One of the stabilizations corresponds to attaching a 1-handle with feet near z_i and x_i . We introduce a new alpha and a new beta curve. The alpha curve runs parallel to A_i^β . The new beta curve, denoted β_i^s , is the belt-sphere of the 1-handle. Similarly, the other stabilization corresponds to attaching a 1-handle with feet near x_i and w_i . Here, the new beta curve runs parallel to A_i^α . The new alpha curve, α_i^s , is the belt sphere of the 1-handle.




Figure 9.3. Top: the Heegaard surface $(\Sigma, \alpha, \beta, w, z)$ and the arc \mathcal{A}_i from Lemma 9.8. Bottom: the stabilized Heegaard surface $(\Sigma', \alpha', \beta', w, z)$ and the arc \mathcal{A}'_i . Also shown are the special curves α_i^s and β_i^s .

We write $(\Sigma', \alpha', \beta', w, z)$ for the stabilized Heegaard diagram. We handleslide the arcs \mathcal{A}_i into the new stabilization tubes to obtain arcs $\mathcal{A}'_i \subseteq \Sigma'$. The arcs $\mathcal{A}'_1, \dots, \mathcal{A}'_\ell$ are still isotopic to the arcs in \mathcal{A} .

By definition of a Heegaard link diagram, the base points $w_i, z_i \in \Sigma$ are contained in a single component $A_i \subseteq \Sigma \setminus \alpha$ and also a single component $B_i \subseteq \Sigma \setminus \beta$. Note that the point $x_i \in \Sigma$ is contained in A_i since there is a path A_i^β from x_i to z_i which is disjoint from α . Similarly, $x_i \in B_i$. Write A'_i and B'_i for the analogous components of $\Sigma' \setminus \alpha'$ and $\Sigma' \setminus \beta'$. Note that A'_i is obtained by removing four disks from A_i , and similarly for B'_i .

We define our σ -basic system of Heegaard diagrams as follows. In the i -th axis direction, we first move β_i^s around B_i by a sequence of handleslides across the curves of β . This is similar to the case of meridional σ -basic systems in Lemma 9.7. We perform a sequence of isotopies and handleslides to move β_i^s to the other side of z_i , so that z_i may be moved to x_i without crossing any attaching curves. Write z'_i for x_i , thought of as a base point. The next step in the i -th axis direction is to move α_i^s around A_i so that it moves to the other side of z'_i , allowing z'_i to be moved to w_i .

Note that since β_i^s moves only around B'_i , and α_i^s moves only around A'_i , the β_i^s may be moved around simultaneously (for different i) and similarly different α_i^s may be moved simultaneously. Additionally, the original curves α and β are also unchanged in this process. In particular, we may build the different axis directions of the hypercube simultaneously to build the σ -basic system of Heegaard diagrams. By a similar argument to Lemma 9.7, we may choose the constituent hypercubes of attaching circles to be algebraically rigid and weakly admissible. ■

9.5 Tensor products of hypercubes of attaching curves

If \mathcal{L}_β is a hypercube of attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z})$, and $\mathcal{L}_{\beta'}$ is a hypercube of attaching curves on $(\Sigma', \mathbf{w}', \mathbf{z}')$, of dimensions n and m , respectively, we now define an $(n + m)$ -dimensional hypercube $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ on $(\Sigma \sqcup \Sigma', \mathbf{w} \cup \mathbf{w}', \mathbf{z} \cup \mathbf{z}')$. We note that the construction of $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is essentially identical to Lipshitz, Ozsváth and Thurston’s construction of a *connected sum of chain complexes of attaching circles* [26, Definition 3.40].

We begin with the underlying attaching curves of $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$, which we denote by $(\delta_{(\varepsilon, \nu)})_{(\varepsilon, \nu) \in \mathbb{E}_n \times \mathbb{E}_m}$, where $n = \dim(\mathcal{L}_\beta)$ and $m = \dim(\mathcal{L}_{\beta'})$. We set $\delta_{(\varepsilon, \nu)}$ to be a small translation of $\beta_\varepsilon \cup \beta'_\nu$. We assume the small translations are chosen so that the copies of β_ε from $\delta_{(\varepsilon, \nu)}$ and $\delta_{(\varepsilon, \nu')}$ intersect in $2^{g(\Sigma) + |\mathbf{w}|-1}$ points, and similarly for the translated curves on Σ' .

Before defining the morphisms of $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$, we begin with some helpful terminology.

Definition 9.9. If $\varepsilon, \varepsilon' \in \mathbb{E}_n$, we say that there is an *arrow* from ε to ε' if $\varepsilon < \varepsilon'$.

- (1) We say that an arrow from $(\varepsilon, \nu) \in \mathbb{E}_n \times \mathbb{E}_m$ to (ε', ν') is *mixed* if $\varepsilon < \varepsilon'$ and $\nu < \nu'$.
- (2) We say that an arrow from (ε, ν) to (ε', ν') is *non-mixed* if $\varepsilon = \varepsilon'$ or $\nu = \nu'$.

We now describe the morphisms of $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. We define the chains in $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ for each mixed arrow to be zero. For the arrow from (ε, ν) to (ε, ν') , where $\nu < \nu'$, we use the chain

$$\Theta_{\varepsilon, \varepsilon}^+ \otimes \Theta_{\nu, \nu'},$$

where $\Theta_{\varepsilon, \varepsilon}^+$ denotes the top degree generator (recall that we have perturbed one copy of β_ε slightly in the definition), and $\Theta_{\nu, \nu'}$ denotes the chain from $\mathcal{L}_{\beta'}$. Similarly, for the arrow from (ε, ν) to (ε', ν) , where $\varepsilon < \varepsilon'$, we use the chain

$$\Theta_{\varepsilon, \varepsilon'} \otimes \Theta_{\nu, \nu}^+.$$

Remark 9.10. We will only use the above construction in the case that \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are algebraically rigid. In this case, $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is also algebraically rigid, and the hypercube relations are automatic.

9.6 Basic systems and connected sums

In this section, we define a connected sum operation on σ -basic systems of Heegaard diagrams.

We begin by defining a connected sum operation for systems of arcs. Let L_1 and L_2 be links in S^3 , with distinguished components K_1 and K_2 , respectively. If \mathcal{A}_1

and \mathcal{A}_2 are two systems of arcs for L_1 and L_2 , we can form their connected sum $\mathcal{A}_1\#\mathcal{A}_2$ by taking the connected sum of L_1 and L_2 at $w_1 \in L_1$ and $z_2 \in L_2$. We then concatenate the arcs for the components K_1 and K_2 to obtain the arc for $K_1\#K_2$. *A priori*, this operation is asymmetric between L_1 and L_2 .

The most important special cases of this construction are the following:

- (#-1) $\alpha\alpha$ *connected sums*: Suppose \mathcal{A}_1 and \mathcal{A}_2 are systems of arcs for L_1 and L_2 , such that the arcs for K_1 and K_2 are both alpha-parallel. We form a system of arcs for $L_1\#L_2$ by using an alpha-parallel arc for $K_1\#K_2$, and using the remaining arcs without change.
- (#-2) $\beta\beta$ *connected sums*: These are analogous to $\alpha\alpha$.
- (#-3) $\alpha\beta$ *connected sums*: Suppose that \mathcal{A}_1 and \mathcal{A}_2 are systems of arcs for L_1 and L_2 , such that the arc for K_1 is alpha-parallel while the arc for K_2 is beta-parallel. We form the system $\mathcal{A}_1\#\mathcal{A}_2$ by using the co-core of the connected sum band as the arc for $K_1\#K_2$.

These special cases are illustrated in Figure 9.4.

We now describe how to realize the above connected sum operations on the level of σ -basic systems of Heegaard diagrams. Suppose that we have σ -basic systems \mathcal{H}_1 and \mathcal{H}_2 for L_1 and L_2 . We focus on the largest hyperbox of the system, since the smaller hyperboxes are determined by the compatibility condition with respect to inclusion.

We take the connected sum of the Heegaard surfaces Σ_1 and Σ_2 at $w_1 \in K_1$ and $z_2 \in K_2$, deleting those two base points, and leaving z_1 and w_2 . See Figure 9.5.

We write $\mathcal{H}_1^{(0)}, \mathcal{H}_1^{(1)} \subseteq \mathcal{H}_1$ for the codimension 1 subboxes where the K_1 -component is 0 (resp. where the K_1 -component is maximal). Write $\mathcal{H}_2^{(0)}$ and $\mathcal{H}_2^{(1)}$ for the analogous subboxes of \mathcal{H}_2 .




Figure 9.4. The connected sums of two alpha-parallel arcs (top row), and the connected sum of one alpha-parallel arc and one beta-parallel arc (bottom row).




Figure 9.5. Forming the connected sum of two σ -basic systems of Heegaard diagrams. The dashed arcs denote the arcs in our σ -basic system. On the left, we begin with an alpha-parallel arc for K_1 and a beta-parallel arc for K_2 . On the right, we have the arc for the connected sum.

We first construct a hyperbox of Heegaard diagrams on $\Sigma_1 \# \Sigma_2$, which we denote $\mathcal{H}_1 \# \mathcal{H}_2^{(0)}$. We build this hyperbox similarly to the tensor product procedure from Section 9.5. We build an analogous hyperbox of Heegaard diagrams $\mathcal{H}_1^{(1)} \otimes \mathcal{H}_2$ of dimension $(\ell_1 + \ell_2 - 1)$. Note that since these hypercubes take place on the connected sum $\Sigma_1 \# \Sigma_2$, as opposed to the disjoint union, we may not be able to build the chains in the constituent hypercubes of attaching curves using the tensor product operation. Nonetheless, we may always use the curves so-constructed, and fill in the morphisms of these hypercubes using the standard filling construction of Manolescu–Ozsváth [32, Lemma 8.6].

Both $\mathcal{H}_1 \otimes \mathcal{H}_2^{(0)}$ and $\mathcal{H}_1^{(1)} \otimes \mathcal{H}_2$ share $\mathcal{H}_1^{(1)} \otimes \mathcal{H}_2^{(0)}$ as a codimension 1 subspace. In particular, we may stack $\mathcal{H}_1 \otimes \mathcal{H}_2^{(0)}$ and $\mathcal{H}_1^{(1)} \otimes \mathcal{H}_2$. Hence, for the connected sum $L_1 \# L_2$, we use

$$\mathcal{H}_1 \# \mathcal{H}_2 := \text{St}(\mathcal{H}_1 \otimes \mathcal{H}_2^{(0)}, \mathcal{H}_1^{(1)} \otimes \mathcal{H}_2), \tag{9.1}$$

where St denotes stacking hyperboxes. This construction yields a σ -basic system of Heegaard diagrams for $L_1 \# L_2$, whose largest hyperbox is given by equation (9.1).

Remark 9.11. If \mathcal{H}_1 and \mathcal{H}_2 are algebraically rigid, then $\mathcal{H}_1 \# \mathcal{H}_2$ is also algebraically rigid.

Remark 9.12. Note that $\text{St}(\mathcal{H}_1^{(0)} \otimes \mathcal{H}_2, \mathcal{H}_1^{(1)} \otimes \mathcal{H}_2)$ does not naturally define a σ -basic system of Heegaard diagrams. This is because the Heegaard diagrams must

encode an isotopy of the base point z_1 to the base point w_2 on the connected sum of the Heegaard surfaces. The stacking in equation (9.1) naturally encodes an isotopy (concatenating the isotopy from z_1 to w_1 , the isotopy of the base point across the connected sum tube from w_1 to z_2 , and the isotopy of z_2 to w_2). The alternate stacking does not naturally encode an isotopy of this sort.

Chapter 10

Hypercubes and disjoint unions

In this chapter, we consider hypercubes of attaching curves on disconnected Heegaard surfaces. Recall that in Section 9.5, we defined a tensor product operation on hypercubes of attaching curves. The goal of this chapter is to prove the following tensor product formula for pairing such hypercubes of attaching curves.

Proposition 10.1. *Suppose that \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are admissible hypercubes of handleslide equivalent beta attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z})$ and $(\Sigma', \mathbf{w}', \mathbf{z}')$, respectively. Suppose additionally that \mathcal{L}_α and $\mathcal{L}_{\alpha'}$ are hypercubes of alpha attaching curves, satisfying the same assumptions. Suppose also that if α and β are curves in \mathcal{L}_α and \mathcal{L}_β , then $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is the diagram of a link in a rational homology 3-sphere, and similarly for $\mathcal{L}_{\alpha'}$ and $\mathcal{L}_{\beta'}$.*

- (1) *If the small translates in the construction of $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ are chosen suitably small, then $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is a hypercube of attaching curves. The same holds for $\mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}$.*
- (2) *If the translations used in the constructions of $\mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}$ and $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ are suitably small, then the canonical map of vector spaces*

$$\begin{aligned} \mathbf{CF}_{\mathcal{J} \sqcup \mathcal{J}'}^-(\Sigma \sqcup \Sigma', \mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}, \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}, \mathbf{w} \cup \mathbf{w}', \mathbf{z} \cup \mathbf{z}') \\ \cong \mathbf{CF}_{\mathcal{J}}^-(\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w}, \mathbf{z}) \otimes_{\mathbb{F}} \mathbf{CF}_{\mathcal{J}'}^-(\Sigma', \mathcal{L}_{\alpha'}, \mathcal{L}_{\beta'}, \mathbf{w}', \mathbf{z}'), \end{aligned}$$

is a chain isomorphism.

The same statements hold if Σ and Σ' have free base points in addition to link base points.

Remark 10.2. In our proof, the assumption that $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is a diagram for a link in a rational homology 3-sphere is only used to simplify several statements about admissibility. For general three-manifolds, the statement above holds as long as we restrict to a finite set of Spin^c structures on the hypercubes and assume the diagrams satisfy a stronger version of admissibility (see [39, Definition 4.10]).

Our proof of Proposition 10.1 follows from the same line of reasoning as [26, Section 3.5], though we present it for the benefit of the reader.

10.1 Small translate theorems

We now review a preliminary technical result concerning Heegaard diagrams with repeated attaching curves. These are based on the results of [26, Section 3] and [12, Section 11].

Suppose $\mathcal{D} = (\Sigma, \alpha_j, \dots, \alpha_1, \beta_1, \dots, \beta_k, \mathbf{w}, \mathbf{z})$ is a multi-pointed Heegaard multi-diagram, the curves β_1, \dots, β_k are pairwise handleslide equivalent, and the curves $\alpha_1, \dots, \alpha_j$ are pairwise handleslide equivalent. Let $\mathbb{I} = (i_1, \dots, i_j)$ and $\mathbb{I}' = (i'_1, \dots, i'_k)$ be tuples of positive integers. For $s \in \{1, \dots, j\}$, pick curves $\alpha_s^1, \dots, \alpha_s^{i_s}$ which are small Hamiltonian translates of α_s . Similarly, for $t \in \{1, \dots, k\}$, pick attaching curves $\beta_t^1, \dots, \beta_t^{i'_t}$ which are small translates of β_t . We define the diagram

$$\mathbb{I}\mathcal{D}_{\mathbb{I}'} := (\Sigma, \alpha_j^{i_j}, \dots, \alpha_j^1, \dots, \alpha_1^{i_1}, \dots, \alpha_1^1, \beta_1^{i'_1}, \dots, \beta_1^1, \dots, \beta_k^{i'_k}, \dots, \beta_k^1, \mathbf{w}, \mathbf{z}).$$

We assume that each $\mathbb{T}_{\alpha_s^n} \cap \mathbb{T}_{\alpha_s^m}$ contains exactly $2^{g(\Sigma) + |\mathbf{w}| - 1}$ points, for each s, n and m , and similarly for the beta-translates.

Lemma 10.3. *Suppose $\mathcal{D} = (\Sigma, \alpha_j, \dots, \alpha_1, \beta_1, \dots, \beta_k, \mathbf{w}, \mathbf{z})$ is a multi-diagram such that the α_i are all pairwise handleslide equivalent, and the β_i are also all pairwise handleslide equivalent. Assume that $(\Sigma, \alpha_1, \beta_1, \mathbf{w}, \mathbf{z})$ represents a link in a rational homology 3-sphere. Assume that \mathcal{D} is admissible for each complete collection of base points $\mathcal{W} \subseteq \mathbf{w} \cup \mathbf{z}$. Let $\mathbb{I}\mathcal{D}_{\mathbb{I}'}$ be the diagram constructed above. Assume the translations are chosen so that $\mathbb{I}\mathcal{D}_{\mathbb{I}'}$ is also weakly admissible. Suppose that $(J_y)_{y \in K_{\ell-1}}$ is a generically chosen family of almost complex structures for counting holomorphic ℓ -gons, where ℓ is the total number of attaching curves on $\mathbb{I}\mathcal{D}_{\mathbb{I}'}$. Here, $K_{\ell-1}$ is Stasheff's associahedron on $\ell - 1$ inputs. Equivalently, $K_{\ell-1}$ is the moduli space of complex disks with ℓ boundary marked points, one of which is distinguished as the “output”. If the translations in the construction are chosen suitably small, then the following holds: If ψ is a class of ℓ -gons representing \approx such that the input of each $\mathbf{CF}^-(\alpha_s^{i_s+1}, \alpha_s^{i_s})$ and $\mathbf{CF}^-(\beta_t^{i'_t+1}, \beta_t^{i'_t})$ is the top degree generator and ψ has a J_y -holomorphic representative for some $y \in K_{\ell-1}$, then*

$$\mu(\psi) \geq \min(0, 3 - j - k).$$

Equivalently, if $S := \sum_{s=1}^j (i_s - 1) + \sum_{t=1}^k (i'_t - 1)$ denotes the number of special inputs from small translates, then

$$\mu(\psi) \geq \min(0, 3 - \ell + S).$$

Remark 10.4. The condition that $(\Sigma, \alpha_1, \beta_1)$ represents a rational homology 3-sphere is to obtain finiteness in the number of classes that the polygon maps count. In this case, the existence of the Maslov gradings $\text{gr}_{\mathcal{W}}$ for each complete collection

$\mathcal{W} \subseteq \mathbf{w} \cup \mathbf{z}$ will ensure that for each N , there are only finitely many nonnegative classes in a given

$$\pi_2(\Theta_{n,n-1}^\alpha, \dots, \Theta_{2,1}^\alpha, \mathbf{x}, \Theta_{1,2}^\beta, \dots, \Theta_{m-1,m}^\beta, \mathbf{y})$$

with Maslov index N . Similar arguments can be made for diagrams of general 3-manifolds, though one would need to require a stronger version of admissible (such as strong \mathfrak{s} -admissibility). See [39, Definition 4.10].

We now briefly sketch the proof of Lemma 10.3. We assume for simplicity of notation that $j = 1$ and $i_1 = 1$, so that there are only beta-translates. Consider a class of ℓ -gons ψ on $\mathbb{I}\mathcal{D}\mathbb{I}'$, such that for each i, t , the input from $\mathbb{T}_{\beta_t^i} \cap \mathbb{T}_{\beta_{t+1}^i}$ is the top degree generator. There is a well-defined approximating class ψ^{app} on \mathcal{D} , and $\mu(\psi) = \mu(\psi^{\text{app}})$ (at this step, one needs the special inputs to be the top degree generator; see [12, Lemma 11.3], which generalizes easily from triangles to ℓ -gons). For each i and j , we pick a sequence $\{\beta_{t,n}^i\}_{n \in \mathbb{N}}$ such that as $n \rightarrow \infty$, the curves $\beta_{t,n}^i$ approach β_t^i (in the sense that there are maps $\beta_{t,n}^i$ is the time 1 translation by a Hamiltonian vector field $X_{H_{i,t,n}}$ for a sequence of functions $H_{i,t,n}$ on Σ which converge to 0 in the C^∞ topology). Write \mathcal{D}'_n for the diagram obtained by replacing each β_t^i with $\beta_{t,n}^i$. Given a sequence u_n of holomorphic representatives of ψ on \mathcal{D}'_n , we may extract a subsequence which converges to a representative of ψ^{app} , which is a $(k + 1)$ -gon. In particular, by applying transversality for $(k + 1)$ -gons, one obtains that

$$\mu(\psi) \geq \min(0, 3 - (k + 1)).$$

As noted in Remark 10.4, we use the assumption that $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is a diagram for a rational homology sphere so that there are only finitely many classes of ℓ -gons of index $3 - \ell$ representing a given Spin^c structure, and we apply the above argument to each class.

We refer the reader to [26, Section 3] and [12, Section 11] for additional details.

There is an additional refinement of the above lemma for the case when $\mathcal{D} = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is an ordinary Heegaard diagram, and we have only one translate β' of the curves β . In this case, if β' are chosen to be suitably small Hamiltonian translates of β , then there is a canonical nearest point map

$$\Phi_{\text{np}}: \mathbb{T}_\alpha \cap \mathbb{T}_\beta \rightarrow \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}.$$

Extending this map linear over the variables, we obtain a map

$$\Phi_{\text{np}}: \mathbf{CF}^-(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z}) \rightarrow \mathbf{CF}^-(\Sigma, \alpha, \beta', \mathbf{w}, \mathbf{z}).$$

Lemma 10.5. *Suppose that $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ is a diagram for a link in a rational homology 3-sphere which is weakly admissible for each complete collection $\mathcal{W} \subseteq \mathbf{w} \cup \mathbf{z}$,*

and suppose that β' are suitably small translates of the curves β . Then we have an equality of maps

$$\Phi_{\text{np}}(-) = f_{\alpha, \beta, \beta'}(-, \Theta_{\beta, \beta'}^+).$$

The proof in the case of $\widehat{\text{CF}}$ was given by Lipshitz, Ozsváth and Thurston in [26, Lemma 3.38]. This was extended to CF^- (mostly by adapting the notation of [26]) in [12, Proposition 11.1].

10.2 Disconnected Heegaard surfaces

Suppose that \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z})$ and $(\Sigma', \mathbf{w}', \mathbf{z}')$, of dimension n and m , respectively. In this section, we prove part (1) of Proposition 10.1, i.e., for suitable choices of translates, the diagram $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is a hypercube of attaching curves (compare [26, Proposition 3.52]).

We begin with a technical result, from which we will derive the hypercube relations. In the following, we will write $\kappa = (\varepsilon, \nu)$ for points in $\mathbb{E}_{n+m} \cong \mathbb{E}_n \times \mathbb{E}_m$. Also, we write $(\delta_\kappa)_{\kappa \in \mathbb{E}_{n+m}}$ for the attaching curves of \mathbb{E}_{n+m} . If $\kappa_1 < \dots < \kappa_\ell$ is an increasing sequence of indices, we write $f_{\kappa_1, \dots, \kappa_\ell}$ for the holomorphic ℓ -gon map $f_{\delta_{\kappa_1}, \dots, \delta_{\kappa_\ell}}$. We similarly write $f_{\varepsilon_1, \dots, \varepsilon_\ell}$ and $f_{\nu_1, \dots, \nu_\ell}$ for holomorphic polygon maps on Σ and Σ' . We write $\Theta_{\kappa_{i-1}, \kappa_i}$ for the chains on $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. We recall that these are given by

$$\Theta_{\kappa_i, \kappa_{i+1}} = \begin{cases} \Theta_{\varepsilon_i, \varepsilon_{i+1}} \otimes \Theta_{\nu_i}^+ & \text{if } \nu_i = \nu_{i+1}, \\ \Theta_{\varepsilon_i}^+ \otimes \Theta_{\nu_i, \nu_{i+1}} & \text{if } \varepsilon_i = \varepsilon_{i+1}, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 10.6. *Let \mathcal{L}_β and $\mathcal{L}_{\beta'}$ be hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z})$ and $(\Sigma', \mathbf{w}', \mathbf{z}')$, respectively, and let $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ denote the hypercube on $(\Sigma \sqcup \Sigma', \mathbf{w} \cup \mathbf{w}', \mathbf{z} \cup \mathbf{z}')$ described in Section 9.5. Suppose that $\kappa_1 < \dots < \kappa_\ell$ is an increasing sequence in \mathbb{E}_{n+m} , and let $\Theta_{\kappa_{i-1}, \kappa_i}$ denote the chains in $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. If $\ell \neq 3$, then*

$$\begin{aligned} & f_{\kappa_1, \dots, \kappa_\ell}(\Theta_{\kappa_1, \kappa_2}, \dots, \Theta_{\kappa_{\ell-1}, \kappa_\ell}) \\ &= \begin{cases} f_{\varepsilon_1, \dots, \varepsilon_\ell}(\Theta_{\varepsilon_1, \varepsilon_2}, \dots, \Theta_{\varepsilon_{\ell-1}, \varepsilon_\ell}) \otimes \Theta_{\nu_1, \nu_1}^+ & \text{if } \nu_1 = \nu_\ell, \\ \Theta_{\varepsilon_1, \varepsilon_1}^+ \otimes f_{\nu_1, \dots, \nu_\ell}(\Theta_{\nu_1, \nu_2}, \dots, \Theta_{\nu_{\ell-1}, \nu_\ell}) & \text{if } \varepsilon_1 = \varepsilon_\ell, \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{10.1}$$

When $\ell = 2$, we interpret f_{κ_1, κ_2} as the ordinary Floer differential. When $\ell = 3$, equation (10.1) holds when $\nu_1 = \nu_3$ or $\varepsilon_1 = \varepsilon_3$. If $\ell = 3$ and instead $\varepsilon_1 < \varepsilon_2 = \varepsilon_3$ and $\nu_1 = \nu_2 < \nu_3$, then

$$f_{\kappa_1, \kappa_2, \kappa_3}(\Theta_{\kappa_1, \kappa_2}, \Theta_{\kappa_2, \kappa_3}) = \Theta_{\varepsilon_1, \varepsilon_2} \otimes \Theta_{\nu_2, \nu_3},$$

where we identify complexes whose attaching curves are small approximations of each other. Similarly, if $\varepsilon_1 = \varepsilon_2 < \varepsilon_3$ and $\nu_1 < \nu_2 = \nu_3$ we have

$$f_{\kappa_1, \kappa_2, \kappa_3}(\Theta_{\kappa_1, \kappa_2}, \Theta_{\kappa_2, \kappa_3}) = \Theta_{\varepsilon_2, \varepsilon_3} \otimes \Theta_{\nu_1, \nu_2}.$$

Proof. Consider first the claim when $\ell = 2$ (i.e., for holomorphic disks). Suppose that $(\varepsilon_1, \nu_1) < (\varepsilon_2, \nu_2)$ and consider $\Theta_{\kappa_1, \kappa_2}$. By assumption, $\Theta_{\kappa_1, \kappa_2}$ is only non-trivial if $\varepsilon_1 = \varepsilon_2$ or $\nu_1 = \nu_2$, so assume for concreteness that $\varepsilon_1 < \varepsilon_2$ and $\nu_1 = \nu_2$. By definition, $\Theta_{\kappa_1, \kappa_2} = \Theta_{\varepsilon_1, \varepsilon_2} \otimes \Theta_{\nu_1, \nu_1}^+$. The differential on the disjoint union of two diagrams is clearly tensorial, so

$$\partial(\Theta_{\varepsilon_1, \varepsilon_2} \otimes \Theta_{\nu_1, \nu_1}^+) = \partial\Theta_{\varepsilon_1, \varepsilon_2} \otimes \Theta_{\nu_1, \nu_1}^+ + \Theta_{\varepsilon_1, \varepsilon_2} \otimes \partial\Theta_{\nu_1, \nu_1}^+ = \partial\Theta_{\varepsilon_1, \varepsilon_2} \otimes \Theta_{\nu_1, \nu_1}^+,$$

proving the claim in this case. The case that $\ell = 3$ is similarly straightforward to verify.

We assume now that $\ell > 3$. The map $f_{\kappa_1, \dots, \kappa_\ell}$ counts holomorphic ℓ -gons of Maslov index $3 - \ell$. The holomorphic curves map into $(\Sigma \sqcup \Sigma') \times D_\ell$. To define the polygon counting map, we pick a family of almost complex structures $(\mathcal{J}_y)_{y \in K_{\ell-1}}$ on $(\Sigma \sqcup \Sigma') \times D_\ell$. Such a family of almost complex structures induces two families, $(J_y)_{y \in K_{\ell-1}}$ and $(J'_y)_{y \in K_{\ell-1}}$, on $\Sigma \times D_\ell$ and $\Sigma' \times D_\ell$, respectively.

The holomorphic polygon map $f_{\kappa_1, \dots, \kappa_\ell}$ on the disjoint union may be equivalently described by counting pairs (u, u') representing pairs of classes of ℓ -gons (ψ, ψ') satisfying

$$\mu(\psi) + \mu(\psi') = 3 - \ell, \tag{10.2}$$

which have the same almost complex structure parameter $y \in K_{\ell-1}$. In particular, we may naturally view the moduli space for a class (ψ, ψ') (as would be counted by the polygon map $f_{\kappa_1, \dots, \kappa_\ell}$) as the fibered product

$$\begin{aligned} & \bigcup_{y \in K_{\ell-1}} \mathcal{M}_{\mathcal{J}_y}(\psi, \psi') \times \{y\} \\ &= \left(\bigcup_{y \in K_{\ell-1}} \mathcal{M}_{J_y}(\psi) \times \{y\} \right) \times_{\text{ev}} \left(\bigcup_{y \in K_{\ell-1}} \mathcal{M}_{J'_y}(\psi') \times \{y\} \right), \end{aligned} \tag{10.3}$$

where ev is the evaluation map to $K_{\ell-1}$, which sends a pair (u, y) to the parameter y .

Consider the last line of equation (10.1). This equation concerns sequences $\kappa_1 < \dots < \kappa_\ell$ which increment both the \mathbb{E}_n -coordinates and the \mathbb{E}_m -coordinates. Consider a class of ℓ -gons (ψ, ψ') on $(\Sigma \sqcup \Sigma') \times D_\ell$, potentially counted by the ℓ -gon map $f_{\kappa_1, \dots, \kappa_\ell}$. Let k be the number of inputs of ψ which do not increment the \mathbb{E}_n -coordinate. Let k' be the number of inputs of ψ' which do not increment the \mathbb{E}_m -coordinate.

By Lemma 10.3, if (ψ, ψ') has a holomorphic representative, we must have

$$\mu(\psi) \geq \min(3 - \ell + k, 0) \quad \text{and} \quad \mu(\psi') \geq \min(3 - \ell + k', 0), \tag{10.4}$$

because each input of ψ which does not increase the \mathbb{E}_n -coordinate contributes a top degree generator in a pair of translated curves as input. Since $\mu(\psi) + \mu(\psi') = 3 - \ell$, by assumption, for a holomorphic representative to exist we must have

$$\min(3 - \ell + k, 0) + \min(3 - \ell + k', 0) \leq 3 - \ell.$$

We add $2(\ell - 3)$ to the above to obtain

$$\min(k, \ell - 3) + \min(k', \ell - 3) \leq \ell - 3. \quad (10.5)$$

On the other hand, there are $\ell - 1$ total inputs to the map $f_{\kappa_1, \dots, \kappa_\ell}$, and each input increments either \mathbb{E}_n or \mathbb{E}_m , but not both (by definition of $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$). Hence

$$k + k' = \ell - 1. \quad (10.6)$$

Combining the above equation with equation (10.5), we deduce that it is not possible that both $k \leq \ell - 3$ and $k' \leq \ell - 3$, since then equation (10.5) would imply $k + k' \leq \ell - 3$, which contracts equation (10.6). Hence, we assume without loss of generality that $k > \ell - 3$. Equation (10.5) implies that

$$\min(k', \ell - 3) \leq 0,$$

which implies that either $k' = 0$ or $\ell \in \{2, 3\}$. Since we assumed that $\ell > 3$, we conclude that $k' = 0$. Equation (10.6) implies that $k = \ell - 1$. Symmetrically, if instead we assumed $k' > \ell - 3$, we would have $(k, k') = (0, \ell - 1)$. This proves the last line of equation (10.1).

We now consider the first two lines of equation (10.1). These correspond to the cases that (k, k') is $(\ell - 1, 0)$ or $(0, \ell - 1)$. Consider the case that $(k, k') = (0, \ell - 1)$. In this case, the holomorphic curve count occurs on the diagram

$$(\Sigma, \beta_{\varepsilon_1}, \dots, \beta_{\varepsilon_\ell}, \mathbf{w}, \mathbf{z}) \sqcup (\Sigma', \beta'_v, \dots, \beta'_v, \mathbf{w}', \mathbf{z}'),$$

where the ℓ attaching curves on the right-hand diagram are small translates and we write v for v_1 . Equation (10.4) implies in this case that $\mu(\psi) \geq 3 - \ell$ and $\mu(\psi') \geq 0$. Together with equation (10.2) we see that

$$\mu(\psi) = 3 - \ell \quad \text{and} \quad \mu(\psi') = 0. \quad (10.7)$$

If $\psi' \in \pi_2(\Theta_{v,v}^+, \dots, \Theta_{v,v}^+, \mathbf{z})$, the Maslov index formula for the grading implies that

$$\mu(\psi') = n_{\mathbf{w}'}(\psi') + \text{gr}_{\mathbf{w}'}(\Theta_{v,v}^+, \mathbf{z}).$$

Since $\mu(\psi') = 0$, we conclude that $n_{\mathbf{w}'}(\psi') = 0$ and $\mathbf{z} = \Theta_{v,v}^+$.

For the main claim, it suffices to show that

$$\# \bigcup_{y \in K_{\ell-1}} \mathcal{M}_{J_y}(\psi) \equiv \sum_{\substack{\psi' \in \pi_2(\Theta_{v,v}^+, \dots, \Theta_{v,v}^+), \\ \mu(\psi')=0, \\ y \in K_{\ell-1}}} \mathcal{M}_{\mathcal{J}_y}(\psi, \psi') \pmod{2}.$$

Equation (10.7) implies that

$$\dim \bigcup_{y \in K_{\ell-1}} \mathcal{M}_{J_y}(\psi) \times \{y\} = 0$$

and

$$\dim \bigcup_{y \in K_{\ell-1}} \mathcal{M}_{J'_y}(\psi') \times \{y\} = \dim(K_{\ell-1}) = \ell - 3.$$

By the fibered product equation in equation (10.3) and the above dimension counts, it therefore suffices to show that the map

$$\text{ev}: \bigcup_{\substack{\psi' \in \pi_2(\Theta_{v,v}^+, \dots, \Theta_{v,v}^+), \\ \mu(\psi')=0, \\ y \in K_{\ell-1}}} \mathcal{M}_{J_y}(\psi') \rightarrow K_{\ell-1}$$

has odd degree. (Compare [26, Lemma 3.50].) This is proven by considering the preimage of a path $\gamma: [0, 1] \rightarrow K_{\ell-1}$, such that $\gamma(0) = y$, $\gamma(t) \in \text{int } K_{\ell-1}$ for $t \in [0, 1)$ and $\gamma(1)$ is a point in $\partial K_{\ell-1}$ which realizes a degeneration of a holomorphic ℓ -gon into $\ell - 2$ triangles (i.e., a point in the lowest-dimensional boundary strata of $K_{\ell-1}$). The codimension 1 degenerations along the image of $(0, 1)$ under γ consist of index 1 holomorphic disks breaking off. The holomorphic disks which break off at the input generator cancel in pairs, since each $\Theta_{v,v}^+$ is a cycle. There are no disks which bubble off in the output, since they would leave an index -1 ℓ -gon in $\pi_2(\Theta_{v,v}^+, \dots, \Theta_{v,v}^+, \mathbf{y})$, for some \mathbf{y} , as well as an index 1 disk $\phi \in \pi_2(\mathbf{y}, \Theta_{v,v}^+)$ with $n_{\mathbf{w}}(\phi) = 0$. The existence of such a degeneration would imply that $\text{gr}(\mathbf{y}, \Theta_{v,v}^+) = 1$, which is impossible since there are no generators with grading higher than $\Theta_{v,v}^+$.

The cardinality of the limit at $t = 1$ corresponds to the component of $\Theta_{v,v}^+$ in the $\ell - 2$ fold composition of maps of the form $f_{\beta'_v, \beta'_v, \beta'_v}(\Theta_{v,v}^+, -)$, applied to the element $\Theta_{v,v}^+$ (e.g., by the grading preserving invariance of the Heegaard Floer complex of connected sums of $S^1 \times S^2$). Clearly this is 1, modulo 2. This establishes the first line of (10.1). The second line follows from the same reasoning, establishing equation (10.1) and completing the proof. ■

We now use the previous lemma to finish our proof that $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is a hypercube of attaching curves.

Proof of part (1) of Proposition 10.1. Suppose that $\kappa < \kappa'$ are points in \mathbb{E}_{n+m} . We wish to show the hypercube relations for $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. Write $\kappa = (\varepsilon, \nu)$ and $\kappa' = (\varepsilon', \nu')$. There are two cases to consider:

- (1) $\varepsilon = \varepsilon'$ or $\nu = \nu'$. I.e., the arrow (κ, κ') is *non-mixed*.
- (2) $\varepsilon < \varepsilon'$ and $\nu < \nu'$. I.e., the arrow (κ, κ') is *mixed*.

In the case of a non-mixed arrow, Lemma 10.6 shows that the hypercube relations for $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ follow immediately from the hypercube relations on \mathcal{L}_β and $\mathcal{L}_{\beta'}$.

For a mixed arrow (κ, κ') , Lemma 10.6 implies that there are exactly two terms which contribute to the hypercube relation. These correspond to the two broken arrow sequences $(\kappa, \kappa_0, \kappa')$ where $\kappa_0 \in \{(\varepsilon, \nu'), (\varepsilon', \nu)\}$. By Lemma 10.6 both of these sequences contribute to $\Theta_{\varepsilon, \varepsilon'} \otimes \Theta_{\nu, \nu'}$, which cancel, so the hypercube relations are satisfied. \blacksquare

We are now able to prove the remainder of Proposition 10.1.

Proof of part (2) of Proposition 10.1. The proof is in the same spirit as part (1) of the proposition. To simplify the notation, we will assume that \mathcal{L}_α and $\mathcal{L}_{\alpha'}$ are both 0-dimensional, and consist of ordinary sets of attaching circles α and α' . If $(\varepsilon, \nu) < (\varepsilon', \nu')$ are points in \mathbb{E}_{n+m} , then the hypercube map from (ε, ν) to (ε', ν') in

$$\mathbf{CF}^-(\Sigma \sqcup \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}, \mathbf{w} \cup \mathbf{w}') \quad (10.8)$$

is obtained by summing over all increasing sequences

$$(\varepsilon, \nu) = (\varepsilon_1, \nu_1) < \cdots < (\varepsilon_\ell, \nu_\ell) = (\varepsilon', \nu'),$$

the holomorphic $(\ell + 1)$ -gon map which has special inputs from $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. We say such a sequence is *mixed* if $\varepsilon < \varepsilon'$ and $\nu < \nu'$. We claim that in the pairing of the two hypercubes of attaching curves, mixed sequences make trivial contribution. This is an approximation argument similar to part (1) of the proposition, as we now describe.

We recall that $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ has no mixed arrows which are assigned a non-zero chain. Hence we may consider only broken arrow paths in \mathbb{E}_{n+m} where each individual arrow increases exactly one of the \mathbb{E}_n or \mathbb{E}_m coordinates. For such a sequence of length $\ell \geq 1$, let k be the number of arrows which do not increase the \mathbb{E}_n coordinate, and let k' be the number of arrows which do not increase the \mathbb{E}_m coordinate. In our present situation, $k + k' = \ell - 1$. The contribution of this arrow path is a count of holomorphic $(\ell + 1)$ -gons, with $\ell - 1$ special inputs from the hypercube $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$, and one input from $\mathbf{CF}^-(\alpha \cup \alpha', \beta_\varepsilon \cup \beta'_\nu)$. Suppose (ψ, ψ') is a homology class which could potentially contribute. By construction, $\mu(\psi) + \mu(\psi') = 2 - \ell$. Equation (10.4) adapts to show that

$$\mu(\psi) \geq \min(2 - \ell + k, 0) \quad \text{and} \quad \mu(\psi') \geq \min(2 - \ell + k', 0),$$

so

$$\min(2 - \ell + k, 0) + \min(2 - \ell + k', 0) \leq 2 - \ell,$$

and hence

$$\min(k, \ell - 2) + \min(k', \ell - 2) \leq \ell - 2.$$

Similar to the argument from part (1), the only allowable configurations are $(k, k') \in \{(\ell - 1, 0), (0, \ell - 1)\}$. This implies that there are no mixed arrows in the hypercube of equation (10.8).

It remains to verify that the non-mixed arrows of equation (10.8) coincide with the tensor product differential. Arguing similarly to the proof of part (1), it suffices to show that if $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta_\varepsilon}$, then the map

$$\text{ev: } \bigcup_{\substack{\psi \in \pi_2(\mathbf{x}, \Theta_{\varepsilon, \varepsilon}^+, \dots, \Theta_{\varepsilon, \varepsilon}^+, \mathbf{y}), \\ \mu(\psi) = 0, \\ \mathbf{y} \in K_\ell}} \mathcal{M}_{J, \mathbf{y}}(\psi) \rightarrow K_\ell$$

is odd degree if \mathbf{y} is the canonical nearest point to \mathbf{x} , $\mathbf{y} = \mathbf{x}_{\text{np}}$, and is even degree otherwise. Similarly to part (1), we consider the preimage of a path $\gamma: [0, 1] \rightarrow K_\ell$, which connects a generic $y \in \text{int } K_\ell$ to a point in ∂K_ℓ of maximal codimension. The only possible generic degenerations on the interior consist of an index 1 disk breaking off. Disks breaking off at the $\Theta_{\varepsilon, \varepsilon}^+$ inputs cancel in pairs. Index 1 disks breaking off at the \mathbf{x} input or the \mathbf{y} output are impossible, since they leave an index -1 $(\ell + 1)$ -gon which has $\ell - 1$ inputs equal to the top degree generator, whose existence would violate Lemma 10.3. Hence, we identify the preimage over y with the preimage of $\gamma(1)$. By the nearest point map argument of Lemma 10.5, the count is odd if and only if $\mathbf{y} = \mathbf{x}_{\text{np}}$, completing the proof. ■

Chapter 11

Hypercubes and connected sums

In this chapter, we prove a connected sum formula for hypercubes of attaching curves. We begin with some preliminary definitions before stating our result in Proposition 11.8.

Definition 11.1. A Heegaard surface with matched link base points $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ consists of a surface with a finite collection of points $\mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$, which is equipped with a matching function

$$m: \mathbf{w} \rightarrow \mathbf{z},$$

which is a bijection. We call $\mathbf{w} \cup \mathbf{z}$ the *link base points*, and \mathbf{p} the *free base points*.

Definition 11.2. Suppose that $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and $(\Sigma', \mathbf{w}', \mathbf{z}', \mathbf{p}')$ are two Heegaard surfaces with matched link base points, as in Definition 11.1. We say that $(\Sigma \# \Sigma', \mathbf{w}'', \mathbf{z}'', \mathbf{p}'')$ is formed by an *admissible connected sum* if it is formed by one of the two following procedures (possibly with the roles of Σ and Σ' reversed).

- (1) (Connected sum at a free base point.) The connected sum is taken at a free base point $p \in \mathbf{p}$, and at some point $x' \in \Sigma' \setminus (\mathbf{w}' \cup \mathbf{z}')$. The base points are

$$\mathbf{w}'' = \mathbf{w} \cup \mathbf{w}', \quad \mathbf{z}'' = \mathbf{z} \cup \mathbf{z}', \quad \text{and} \quad \mathbf{p}'' = (\mathbf{p} \setminus \{p\}) \cup \mathbf{p}'.$$

- (2) (Connected sum at link base points.) The connected sum is taken at two link base points $w \in \mathbf{w}$ and $z' \in \mathbf{z}'$. The base points are

$$\mathbf{w}'' = (\mathbf{w} \setminus \{w\}) \cup \mathbf{w}', \quad \mathbf{z}'' = \mathbf{z} \cup (\mathbf{z}' \setminus \{z'\}), \quad \text{and} \quad \mathbf{p}'' = \mathbf{p} \cup \mathbf{p}'.$$

Suppose that we form an admissible connected sum $(\Sigma \# \Sigma', \mathbf{w}'', \mathbf{z}'', \mathbf{p}'')$ from the pointed surfaces $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and $(\Sigma', \mathbf{w}', \mathbf{z}', \mathbf{p}')$. Suppose \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are hypercubes of handleslide equivalent attaching curves on Σ and Σ' , respectively. We may construct a hypercube-shaped diagram $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ on $(\Sigma \# \Sigma', \mathbf{w}'', \mathbf{z}'', \mathbf{p}'')$ via the same procedure as on the disjoint union of the two diagrams in Section 9.5. We do not claim that the hypercube relations are satisfied, however we have the following remark.

Remark 11.3. If \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are algebraically rigid, then the top degree chains coincide on disjoint unions of diagrams and on admissible connected sums, justifying the use of the notation $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ for both situations. Also, if \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are algebraically rigid, then $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is automatically a hypercube of attaching curves.

To state our connected sum theorem for hypercubes, we need an additional condition on our hypercubes of attaching curves.

Definition 11.4. Suppose that \mathcal{L}_α is an algebraically rigid hypercube of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$, and x is a point in the complement of $\Sigma \setminus \bigcup_{\alpha \in \mathcal{L}_\alpha} \alpha$. We say that x is *base point-esque* for \mathcal{L}_α if for every composable sequence $\Theta_{v_j, v_{j-1}}, \dots, \Theta_{v_2, v_1}$ of chains in \mathcal{L}_α , and every nonnegative class $\psi \in \pi_2(\Theta_{v_j, v_{j-1}}, \dots, \Theta_{v_2, v_1}, \mathbf{x})$, we have

$$\mu(\psi) \geq 2n_x(\psi).$$

(Note \mathbf{x} is the outgoing intersection point in the class ψ .) We make a similar definition for hypercubes of algebraically rigid beta hypercubes.

Lemma 11.5. *If \mathcal{L}_α is an algebraically rigid hypercube of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$, and $x \in \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$, then x is base point-esque for \mathcal{L}_α .*

Proof. Assume for concreteness that $x \in \mathbf{w}$. If $\psi \in \pi_2(\Theta_{v_j, v_{j-1}}, \dots, \Theta_{v_2, v_1}, \mathbf{x})$ is a class, then the relation between the Maslov index and the absolute grading on the Heegaard Floer homology (see [40, Section 7]) implies

$$\mu(\psi) = 2n_{\mathbf{w} \cup \mathbf{p}}(\psi) + \text{gr}_{\mathbf{w} \cup \mathbf{p}}(\Theta_{v_j, v_{j-1}}) + \dots + \text{gr}_{\mathbf{w} \cup \mathbf{p}}(\Theta_{v_2, v_1}) - \text{gr}_{\mathbf{w} \cup \mathbf{p}}(\mathbf{x}),$$

where $\text{gr}_{\mathbf{w} \cup \mathbf{p}}$ is the absolute grading, normalized so that the top degree cycle of

$$\mathbf{HF}^-(\Sigma, \alpha_{v_i}, \alpha_{v_{i-1}}, \mathbf{w}, \mathbf{z}, \mathbf{p})$$

has grading 0. Since \mathcal{L}_α is algebraically rigid, we also know that $\text{gr}_{\mathbf{w} \cup \mathbf{p}}(\mathbf{x}) \leq 0$ and $\text{gr}_{\mathbf{w} \cup \mathbf{p}}(\Theta_{v_{i+1}, v_i}) = 0$ for all i and hence

$$\mu(\psi) \geq 2n_{\mathbf{w} \cup \mathbf{p}}(\psi).$$

If ψ is a nonnegative class, we have $n_{\mathbf{w} \cup \mathbf{p}}(\psi) \geq n_x(\psi)$, so we see that x is base point-esque for \mathcal{L}_α . ■

Remark 11.6. Suppose

$$\mathcal{L}_\alpha = \left(\alpha \xrightarrow{\Theta_{\alpha', \alpha}^+} \alpha' \right),$$

where α' is obtained by performing a Hamiltonian isotopy to α which crosses x , then x is not base point-esque for \mathcal{L}_α because there is a bigon of index 1 which covers x once. See Figure 11.1.

Suppose that \mathcal{L}_α is a hypercube of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and $x \in \Sigma$ is in the complement of each curve in \mathcal{L}_α . If α is in \mathcal{L}_α , write E_x^α for the product of the variables of the base points in the component of $\Sigma \setminus \alpha$ containing x .

Lemma 11.7. *Let \mathcal{L}_α be a hypercube of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and suppose that x is base point-esque for \mathcal{L}_α . Then E_x^α is independent of $\alpha \in \mathcal{L}_\alpha$.*




Figure 11.1. An example of a point x which is not base point-esque from Remark 11.6.

Proof. Let $\alpha_1, \alpha_2 \in \mathcal{L}_\alpha$. Consider the endomorphism A_x of $\mathbf{CF}^-(\alpha_2, \alpha_1)$ which counts holomorphic disks of index 1, which are given a multiplicative factor of $n_x(\phi)$. It is straightforward to see that $[\partial, A_x] = (E_x^{\alpha_1} + E_x^{\alpha_2}) \cdot \text{id}$. On the other hand, if $\Theta_{\alpha_2, \alpha_1}$ is a cycle generating the top degree of homology, then $\partial(\Theta_{\alpha_2, \alpha_1}) = 0$ since $\Theta_{\alpha_2, \alpha_1}$ is a cycle and $A_x(\Theta_{\alpha_2, \alpha_1}) = 0$ since x is base point-esque for \mathcal{L}_α so there are no Maslov index 1 and nonnegative homology classes of disks in any $\pi_2(\Theta_{\alpha_2, \alpha_1}, \mathbf{x})$ with multiplicity 1 on x . Since $\mathbf{CF}^-(\alpha_2, \alpha_1)$ is free and $0 = [\partial, A_x](\Theta_{\alpha_2, \alpha_1}) = \Theta_{\alpha_2, \alpha_1} \cdot (E_x^{\alpha_1} + E_x^{\alpha_2})$, it follows that $E_x^{\alpha_1} + E_x^{\alpha_2} = 0$, completing the proof. ■

The main result of this chapter is the following.

Proposition 11.8. *Suppose that \mathcal{L}_α and \mathcal{L}_β are algebraically rigid hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and $\mathcal{L}_{\alpha'}$ and $\mathcal{L}_{\beta'}$ are algebraically rigid hypercubes of handleslide equivalent attaching curves on $(\Sigma', \mathbf{w}', \mathbf{z}', \mathbf{p}')$. Suppose that we form $(\Sigma \# \Sigma', \mathbf{w}'', \mathbf{z}'', \mathbf{p}'')$ by an admissible connected sum of $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and $(\Sigma', \mathbf{w}', \mathbf{z}', \mathbf{p}')$ at points $x \in \Sigma$ and $x' \in \Sigma'$. Form the hypercubes $\mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}$ and $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ as described above. If x is base point-esque for \mathcal{L}_β and x' is base point-esque for $\mathcal{L}_{\alpha'}$, then there is a natural homotopy equivalence of hypercubes*

$$\begin{aligned} & \mathbf{CF}^-(\Sigma \# \Sigma', \mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}, \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}, \mathbf{w}'', \mathbf{z}'', \mathbf{p}'') \\ & \simeq \mathbf{CF}^-(\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w}, \mathbf{z}, \mathbf{p}) \otimes_R \mathbf{CF}^-(\Sigma', \mathcal{L}_{\alpha'}, \mathcal{L}_{\beta'}, \mathbf{w}', \mathbf{z}', \mathbf{p}'). \end{aligned}$$

Also, R is as follows:

- (1) *If the connected sum is taken along a link component, then $R = \mathbb{F}[\mathcal{U}, \mathcal{V}]$ is the ring for the link components along which the connected sum is taken.*
- (2) *If the connected sum is taken at a free base point $x = p_i \in \mathbf{p}$, then the tensor product is taken over $\mathbb{F}[U]$, where we have U act by U_i on the left factor, and $E_x^{\alpha'}$ on the right factor. (Cf. Lemma 11.7.)*
- (3) *If the connected sum is taken at a free base point $x' = p'_i \in \mathbf{p}'$ in $\mathbf{w}' \cup \mathbf{z}'$, the tensor product is taken over $\mathbb{F}[U]$, where we have U act by U_i on the right factor and E_x^β on the left factor.*

Unlike in Chapter 10, the homotopy equivalence of Proposition 11.8 is not given by the identity map on the level of groups. Instead, it involves counting holomorphic curves.

Remark 11.9. If instead x is base point-esque for \mathcal{L}_α and x' is base point-esque for $\mathcal{L}_{\beta'}$, then there is also a homotopy equivalence of hypercubes realizing a tensor product formula. Note that in the case that if x and x' are base point-esque for all of the hypercubes of Lagrangians, then there are two distinct maps of hypercubes realizing the homotopy equivalence. These are referred to as $\Phi^{\bullet < \circ'}$ and $\Phi^{\circ' < \bullet}$ below.

11.1 Cylindrical boundary degenerations

In this section, we recall some important facts about boundary degenerations.

Definition 11.10. Suppose β is a set of attaching curves on (Σ, \mathbf{w}) , and $\mathbf{x} \in \mathbb{T}_\beta$. A *cylindrical beta boundary degeneration* at \mathbf{x} consists of a tuple (u, S, j) such that (S, j) is a Riemann surface and

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

is a map satisfying the following:

- (1) u is (j, J) -holomorphic.
- (2) u is proper.
- (3) For each $t \in \mathbb{R}$ and $i \in \{1, \dots, g(\Sigma) + |\mathbf{w}| - 1\}$, the set $u^{-1}(\beta_i \times \{0\} \times \{t\})$ consists of a single point.
- (4) u has finite energy.
- (5) $\pi_{\mathbb{H}} \circ u$ is non-constant on each component of S , where $\mathbb{H} = [0, \infty) \times \mathbb{R}$.
- (6) S has a collection of $n = g(\Sigma) + |\mathbf{w}| - 1$ boundary punctures p_1, \dots, p_n . If $\mathbf{x} = (x_1, \dots, x_n)$, where $x_i \in \beta_i$, then

$$\lim_{z \rightarrow p_i} (\pi_{\mathbb{H}} \circ u)(z) = \infty \quad \text{and} \quad \lim_{z \rightarrow p_i} (\pi_{\Sigma} \circ u)(z) = x_i.$$

Cylindrical alpha boundary degenerations are defined by analogy. Of fundamental importance is the mod 2 count of boundary degenerations.

Proposition 11.11. *Suppose β is a set of attaching curves on (Σ, \mathbf{w}) , and B is a Maslov index 2 class of boundary degenerations. For an appropriate choice of almost complex structures on $\Sigma \times [0, \infty) \times \mathbb{R}$, the moduli space of boundary degenerations $\mathcal{N}(B, \mathbf{x})$ is transversely cut out. Furthermore, the parametrized moduli space*

$$\bigcup_{\mathbf{x} \in \mathbb{T}_\beta} \mathcal{N}(B, \mathbf{x}) \times \{\mathbf{x}\}$$

is also transversely cut out. Furthermore,

$$\#\mathcal{N}(B, \mathbf{x}) / \text{Aut}(\mathbb{H}) \equiv 1.$$

The case when $|\mathbf{w}| > 1$ is proven by Ozsváth and Szabó [41, Theorem 5.5]. Ozsváth and Szabó also proved that if $|\mathbf{w}| = 1$ there are generically no boundary degenerations for split almost complex structures on $\Sigma \times [0, \infty) \times \mathbb{R}$. However, for their boundary degenerations, the parametrized moduli space $\bigcup_{\mathbf{x} \in \mathbb{T}_\beta} \mathcal{N}(B, \mathbf{x}) \times \{\mathbf{x}\}$ might in general be non-empty, and there are always broken boundary degenerations at which transversality is not achieved (e.g., the union of a closed curve representing Σ and a constant disk at \mathbf{x}). This is sufficient for showing that $\partial^2 = 0$, but is not sufficient for our purposes. The $|\mathbf{w}| = 1$ case was revisited in [12, Section 7.6], where it was shown that for appropriately generic choices of almost complex structures, the count of Proposition 11.11 holds.

11.2 The 0-dimensional case

In this section, we present the proof of Proposition 11.8 in the case that \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are both 0-dimensional (i.e., we recover the standard connected sum formula of Ozsváth and Szabó [38]). The analytic details we present here will be the basis of the higher-dimensional cases of Proposition 11.8, which we consider in the subsequent section. Our argument is inspired by work of Ozsváth and Szabó in the bordered setting [45]. See also [12, Section 19.4].

Suppose $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ and $\mathcal{H}' = (\Sigma', \alpha', \beta', \mathbf{w}', \mathbf{z}')$ are multi-pointed diagrams, and suppose that J and J' are almost complex structures on $\Sigma \times [0, 1] \times \mathbb{R}$ and $\Sigma' \times [0, 1] \times \mathbb{R}$, respectively. The notation $J \wedge J'$ means the data of the pair J and J' together with distinguished connected sum points, $x \in \Sigma$ and $x' \in \Sigma'$. We assume that the connected sum is admissible, in the sense of Definition 11.2.

Remark 11.12. In Proposition 11.8, we additionally had the assumption that x and x' were base point-esque for some of the hypercubes of attaching curves. In the present case, $\mathcal{L}_\alpha, \mathcal{L}_{\alpha'}, \mathcal{L}_\beta$ and $\mathcal{L}_{\beta'}$ are all 0-dimensional, and this condition is automatic.

For notational simplicity, we focus on the case where we have just one \mathcal{U} variable, and one \mathcal{V} variable, no free base points, and we are doing a connected sum of link components. The first step is to do a neck-stretching degeneration along the connected sum tube. The result of this degeneration is a chain complex, freely generated over the ground ring $R = \mathbb{F}[\mathcal{U}, \mathcal{V}]$ by pairs $\mathbf{x} \times \mathbf{x}'$, where $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $\mathbf{x}' \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_{\beta'}$. The differential is given by the formula

$$\partial_{J \wedge J'}(\mathbf{x} \times \mathbf{x}') = \sum_{\substack{\mathbf{y} \times \mathbf{y}' \in (\mathbb{T}_\alpha \cap \mathbb{T}_\beta) \times (\mathbb{T}_{\alpha'} \cap \mathbb{T}_{\beta'}), \\ \phi \in \pi_2(\mathbf{x}, \mathbf{y}), \\ \phi' \in \pi_2(\mathbf{x}', \mathbf{y}'), \\ n_x(\phi) = n_{x'}(\phi'), \\ \mu(\phi) + \mu(\phi') - 2n_x(\phi) = 1}} \#(\mathcal{MM}(\phi, \phi')/\mathbb{R}) \cdot \mathcal{U}^{n_{\mathbf{w}'}(\phi + \phi')} \mathcal{V}^{n_{\mathbf{z}'}(\phi + \phi')} \cdot \mathbf{y} \times \mathbf{y}', \tag{11.1}$$

where $\mathcal{M}\mathcal{M}(\phi, \phi')$ is the *perfectly matched moduli space*, consisting of pairs (u, u') satisfying the following matching condition. In the following, we write $(S^{\mathbf{q}}, j)$ and $(T^{\mathbf{q}'}, j')$ for the sources curves of u and u' , where (S, j) and (T, j') are Riemann surfaces, and \mathbf{q} and \mathbf{q}' are ordered collections of $n_x(\phi) = n_{x'}(\phi)$ marked points, for which we write $\mathbf{q} = \{q_1, \dots, q_n\}$ and $\mathbf{q}' = \{q'_1, \dots, q'_n\}$. We define

$$\begin{aligned} \mathcal{M}\mathcal{M}_{J \wedge J'}(\phi, \phi') := \{ & (u, u') : u \text{ is } J\text{-holomorphic,} \\ & u' \text{ is } J'\text{-holomorphic,} \\ & (\pi_{\Sigma} \circ u)(q_i) = x, (\pi_{\Sigma'} \circ u')(q'_i) = x', \\ & (\pi_{\mathbb{D}} \circ u)(q_i) = (\pi_{\mathbb{D}'} \circ u')(q'_i)\}. \end{aligned} \quad (11.2)$$

It is not hard to see that this gives a chain complex, which we denote by $\mathbf{CF}^-_{J \wedge J'}(\mathcal{H}, \mathcal{H}')$. (Compare [12, Section 19.4].) The construction is inspired by the matched moduli spaces which appear in [27, Section 9.1].

If I is a (non-singular) almost complex structure on $\Sigma \# \Sigma' \times [0, 1] \times \mathbb{R}$, then one may define a chain homotopy equivalence

$$\mathbf{CF}^-_I(\mathcal{H} \# \mathcal{H}') \simeq \mathbf{CF}^-_{J \wedge J'}(\mathcal{H}, \mathcal{H}')$$

by counting index 0 curves for a non-cylindrical almost complex structure on $\Sigma \# \Sigma' \times [0, 1] \times \mathbb{R}$ which interpolates an ordinary cylindrical almost complex structure on $\Sigma \# \Sigma' \times [0, 1] \times (-\infty, t_1]$, and a degenerate almost complex structure (i.e., one with infinite neck length) on $\Sigma \wedge \Sigma' \times [0, 1] \times [t_2, \infty)$, for some $t_1 \ll 0$ and $t_2 \gg 0$. See [12, Section 19.4] for more details.

We now describe a chain homotopy equivalence

$$\Phi^{\bullet < \sigma'} : \mathbf{CF}^-_{J \wedge J'}(\mathcal{H}, \mathcal{H}') \rightarrow \mathbf{CF}^-_J(\mathcal{H}) \otimes_{\mathbb{R}} \mathbf{CF}^-_{J'}(\mathcal{H}').$$

If $t_0 \in \mathbb{R}$, write

$$\lambda_{t_0} := [0, 1] \times \{t_0\} \subseteq [0, 1] \times \mathbb{R}.$$

Let us write

$$s : \Sigma \times [0, 1] \times \mathbb{R} \rightarrow [0, 1] \quad \text{and} \quad t : \Sigma \times [0, 1] \times \mathbb{R} \rightarrow \mathbb{R}$$

for the projection maps.

Definition 11.13. Suppose $t_0 \in \mathbb{R}$. A $(\Phi, \bullet < \sigma', t_0)$ -*matched $J \wedge J'$ -holomorphic curve pair* consists of a pair of marked $J \wedge J'$ holomorphic disks (u, u') , equipped with the following data:

- (1) A partition of the marked points of u (resp. u') into three sets, \mathbf{S} , \mathbf{C} and \mathbf{N} (resp. \mathbf{S}' , \mathbf{C}' and \mathbf{N}').
- (2) A bijection $\phi : \mathbf{S} \rightarrow \mathbf{S}'$.

We assume the following are satisfied:

- (1) If q_i is a marked point of u , then $(\pi_\Sigma \circ u)(q_i) = x$, and similarly for all marked points of u' .
- (2) If $q \in \mathbf{S}$ and $q' \in \mathbf{S}'$ and $\phi(q) = q'$, then

$$(t \circ u)(q) = (t \circ u')(q') < 0 \quad \text{and} \quad (s \circ u)(q) = (s \circ u')(q').$$

- (3) If $q \in \mathbf{N}$ then $(t \circ u)(q) > t_0$. The analogous statement holds for $q' \in \mathbf{N}'$.
- (4) $|\mathbf{C}| = |\mathbf{C}'|$. Furthermore, if $q \in \mathbf{C}$ then $(t \circ u)(q) = t_0$, and similarly for the marked points of \mathbf{C}' . The marked points of \mathbf{C} and \mathbf{C}' alternate between those of \mathbf{C} and those of \mathbf{C}' along λ_{t_0} . Finally, the left-most marked point along this line is contained in \mathbf{C} .

See Figure 11.2 for a schematic of a $(\Phi, \bullet < \circ', t_0)$ -matched curve pair.

If (ϕ, ϕ') are a pair of homology classes, M is a pair of marked source curves for ϕ and ϕ' , and $t_0 \in \mathbb{R}$, we write $\mathcal{M}\mathcal{M}^{\Phi, \bullet < \circ'}(\phi, \phi', M, t_0)$ for the moduli space of $(\Phi, \bullet < \circ', t_0)$ -matched disks representing ϕ and ϕ' . We consider the parametrized moduli space

$$\mathcal{M}\mathcal{M}^{\Phi, \bullet < \circ'}(\phi, \phi', M) = \bigcup_{t_0 \in \mathbb{R}} \mathcal{M}\mathcal{M}^{\Phi, \bullet < \circ'}(\phi, \phi', M, t_0) \times \{t_0\}.$$

If t_0 is fixed, the expected dimension of $\mathcal{M}\mathcal{M}^{\Phi, \bullet < \circ'}(\phi, \phi', M, t_0)$ is given by the formula

$$\text{ind}(\phi, \phi', M) = \mu(\phi) + \mu(\phi') - 2(|\mathbf{S}| + |\mathbf{C}|).$$

The map $\Phi^{\bullet < \circ'}$ counts pairs of curves (u, u') in $\mathcal{M}\mathcal{M}^{\Phi, \bullet < \circ'}(\phi, \phi', M)/\mathbb{R}$, ranging over tuples with $\text{ind}(\phi, \phi', M) = 0$. Equivalently, we can think of $\Phi^{\bullet < \circ'}$ as counting elements of the non-parametrized moduli spaces $\mathcal{M}\mathcal{M}^{\Phi, \bullet < \circ'}(\phi, \phi', M, t_0)$ for some fixed t_0 . (The parametrized perspective becomes helpful later.)




Figure 11.2. The projection to $[0, 1] \times \mathbb{R}$ of marked points of a $(\Phi, \bullet < \circ', t_0)$ -matched holomorphic curve pair. Solid dots indicate the marked points of u , while open dots are the marked points of u' . The map $\Psi^{\bullet < \circ'}$ counts similar configurations.

Next, we explain the \mathcal{U} and \mathcal{V} weights of a curve. The base points not involved in the connected sum contribute powers of variables as normal. The marked points in $\mathbf{S} \cup \mathbf{S}' \cup \mathbf{C} \cup \mathbf{C}'$ contribute algebra elements which are the same as would be counted on the connected sum of the two diagrams, where we have removed any base points used for the connected sum. For the marked points in \mathbf{N} and \mathbf{N}' , the algebraic contribution is the same as on the disjoint union of the two diagrams, where we treat the punctures as base points.

Concretely, if the connected sum is formed along a link component, and $x = w \in \mathbf{w}$ and $x' = z' \in \mathbf{z}'$, then the algebra element contributed by the marked points would be

$$\mathcal{U}^{|\mathbf{N}|\mathcal{V}^{|\mathbf{N}'|}}. \tag{11.3}$$

(There would be an additional contribution from the other base points not involved in the connected sum.) If instead the connected sum is formed at a free base point $x = w \in \Sigma$ and some non-base point $x' \in \Sigma'$, then the algebra contribution from the marked points would be $\mathcal{U}^{|\mathbf{N}|}$.

We define a similar map

$$\Psi^{\bullet < \circ'} : \mathbf{CF}_{\mathcal{J}}^-(\mathcal{H}) \otimes_R \mathbf{CF}_{\mathcal{J}' }^-(\mathcal{H}') \rightarrow \mathbf{CF}_{\mathcal{J} \wedge \mathcal{J}' }^-(\mathcal{H}, \mathcal{H}')$$

by counting holomorphic curves with similar conditions, but with the roles of the \mathbf{S} and \mathbf{N} labeled marked points switched.

Remark 11.14. Constant disks are counted by both $\Phi^{\bullet < \circ'}$ and $\Psi^{\bullet < \circ'}$.

Lemma 11.15. *The maps $\Phi^{\bullet < \circ'}$ and $\Psi^{\bullet < \circ'}$ are chain maps.*

Proof. We focus on $\Phi^{\bullet < \circ'}$ since the argument for $\Psi^{\bullet < \circ'}$ is not substantially different. Our proof is modeled on work of Ozsváth and Szabó in the setting of bordered knot Floer homology [45]. The proof is to count the ends of $(\Phi, \bullet < \circ', t_0)$ -matched moduli spaces for triples (ϕ, ϕ', M) with $\text{ind}(\phi, \phi', M) = 1$. Modulo the \mathbb{R} -action on the parametrized moduli space, such moduli spaces are 1-dimensional. In the present situation, it is sufficient to fix $t_0 = 0$, and consider only $(\Phi, \bullet < \circ', 0)$ -matched holomorphic curves. The ends are constrained generically to the following configurations:

- (Φ -1) Two paired marked points in \mathbf{S} and \mathbf{S}' may collide with the line λ_0 , away from the marked points in \mathbf{C} and \mathbf{C}' .
- (Φ -2) A pair of punctures in \mathbf{C} and \mathbf{C}' may collide along λ_0 .
- (Φ -3) A puncture of \mathbf{N} or \mathbf{N}' may collide with λ_0 (in the complement of \mathbf{C} and \mathbf{C}'). There are two subcases:
 - (a) After the degeneration, the marked points along λ_0 do not alternate between those of u and u' .
 - (b) After the degeneration, the marked points along λ_0 do alternate between those of u and u' .

- (Φ-4) A puncture of \mathbf{C} may degenerate into an index 2 beta boundary degeneration at height $t = 0$.
- (Φ-5) A puncture of \mathbf{C}' may degenerate into an index 2 alpha boundary degeneration at height $t = 0$.
- (Φ-6) Strip breaking may occur, leaving a $(\Phi, \bullet < o', 0)$ -matched disk of index 0 as well as a holomorphic disk of index 1 in either positive or negative direction. There are two subcases:
 - (a) The index 1 holomorphic disk degenerates towards $-\infty$, and is perfectly-matched.
 - (b) The index 1 holomorphic disk degenerates towards $+\infty$, and has trivial matching (i.e., the projections of all marked points to $[0, 1] \times \mathbb{R}$ are distinct).

Most of these ends appear in canceling pairs, and the rest correspond to the relation $[\partial, \Phi^{\bullet < o'}] = 0$, as we describe presently.

The ends (Φ-1) cancel with the ends (Φ-2).

In an end of type (Φ-3a), there are two adjacent marked points along λ_0 which are both from u or both from u' . Such a curve appears twice in the boundary of the moduli spaces. See Figure 11.3.

The ends of type (Φ-3b) cancel with the end of type (Φ-4) and (Φ-5). Here we are using the mod 2 count of boundary degenerations from Proposition 11.11. Note also the algebra contributions coincide for the two canceling degenerations of type (Φ-3b) and type (Φ-4) or (Φ-5) (cf. equation (11.3)). If the connected sum is taken along a link component, each boundary degeneration which forms along λ_{t_0} will contain one puncture from $\mathbf{C} \cup \mathbf{C}'$ and also one link base point. The degeneration which cancels this boundary degeneration formation consists of a \mathbf{N} or \mathbf{N}' puncture colliding with λ_{t_0} . The algebra weight from the \mathbf{N} or \mathbf{N}' puncture coincides with the weight of the base point which is lost in the boundary degeneration. The case of a connected sum at a free base point is similar.

The remaining ends are (Φ-6), which correspond exactly to the commutator $[\partial, \Phi^{\bullet < o'}]$. Summing all ends, we conclude that

$$\partial \circ \Phi^{\bullet < o'} + \Phi^{\bullet < o'} \circ \partial = 0,$$

completing the proof. ■

Lemma 11.16. *The maps $\Phi^{\bullet < o'}$ and $\Psi^{\bullet < o'}$ are homotopy inverses.*

Proof. We define two maps H_\wedge and H_\sqcup , which we prove satisfy

$$\text{id} + \Psi^{\bullet < o'} \circ \Phi^{\bullet < o'} = [\partial, H_\wedge] \quad \text{and} \quad \text{id} + \Phi^{\bullet < o'} \circ \Psi^{\bullet < o'} = [\partial, H_\sqcup]. \quad (11.4)$$

We focus on the map H_\wedge , since H_\sqcup is constructed by a straightforward modification.

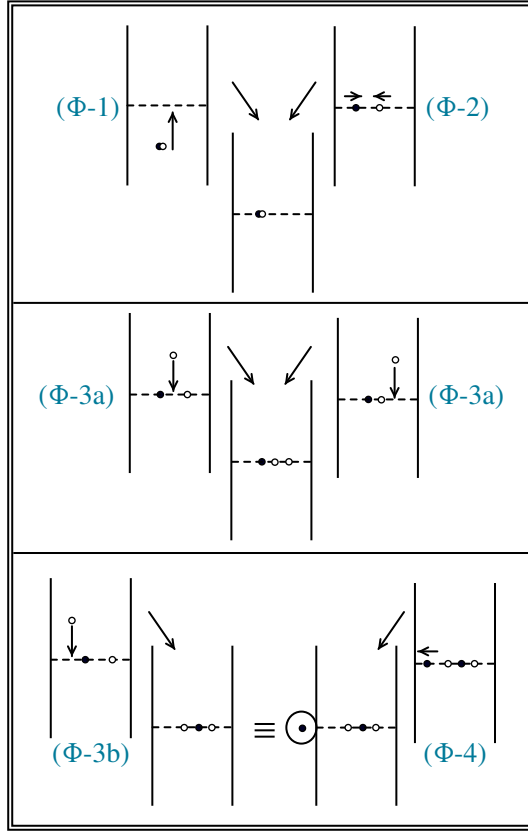


Figure 11.3. Some cancellations in the proof that $\Phi^{\bullet < o'}$ and $\Psi^{\bullet < o'}$ are chain maps.

The map H_\wedge counts certain curve pairs (u, u') where u has five collections of marked points, \mathbf{NN} , \mathbf{NC} , \mathbf{M} , \mathbf{SC} and \mathbf{SS} , and u' has five collections of marked points \mathbf{NN}' , \mathbf{NC}' , \mathbf{M}' , \mathbf{SC}' and \mathbf{SS}' .

Definition 11.17. Suppose that $t_0 < t_1$ are real numbers. We say a pair of marked holomorphic strips (u, u') is (H_\wedge, t_0, t_1) -*matched* if the following are satisfied:

- (1) The marked points of \mathbf{NN} and \mathbf{NN}' are perfectly matched. Furthermore, their projection to \mathbb{R} lies above t_1 .
- (2) The marked points in \mathbf{NC} and \mathbf{NC}' both project to λ_{t_1} . Furthermore, $|\mathbf{NC}| = |\mathbf{NC}'|$. As one travels along λ_{t_1} , the marked points alternate between $|\mathbf{NC}|$ and $|\mathbf{NC}'|$, and the left-most marked point is from \mathbf{NC} , while the right-most is from \mathbf{NC}' .
- (3) The marked points of \mathbf{M} and \mathbf{M}' have no matching condition. Their projection to $[0, 1] \times \mathbb{R}$ lies between λ_{t_0} and λ_{t_1} .

- (4) The marked points **SC** and **SC'** satisfy the same conditions as **NC** and **NC'**, except that they project to λ_{t_0} .
- (5) The marked points in **SS** and **SS'** are perfectly matched, and their projection to \mathbb{R} lies below t_0 .

See Figure 11.4 for a schematic.

If (ϕ, ϕ') are two homology classes of disks, equipped with decorated sources and matching data M , and $t_0 < t_1$ are fixed, then the expected dimension of the moduli space of curves which are (H_\wedge, t_0, t_1) -matched and represent (ϕ, ϕ', M) is given by

$$\text{ind}(\phi, \phi', M) = \mu(\phi) + \mu(\phi') - 2|\text{SS}| - 2|\text{NN}| - 2|\text{SC}| - 2|\text{NC}|.$$

Write $\mathcal{MM}(\phi, \phi', M)$ for this moduli space. The parametrized moduli space

$$\mathcal{MM}(\phi, \phi', M) := \bigcup_{t_0 < t_1} \mathcal{MM}(\phi, \phi', M, t_0, t_1) \times \{(t_0, t_1)\}$$

has a free \mathbb{R} -action, corresponding to overall translation (acting diagonally on t_0 and t_1).

Further, the map H_\wedge counts $\mathcal{MM}(\phi, \phi', M)/\mathbb{R}$ for triples (ϕ, ϕ', M) satisfying $\text{ind}(\phi, \phi', M) = -1$. In this case, the expected dimension of $\mathcal{MM}(\phi, \phi', M)$ is 1.

To verify equation (11.4), we count the ends of the parametrized moduli spaces $\mathcal{MM}(\phi, \phi', M)/\mathbb{R}$ where $\text{ind}(\phi, \phi', M) = 0$ (so that $\dim \mathcal{MM}(\phi, \phi', M)/\mathbb{R} = 1$).

There are ends which occur at $0 < t_1 - t_0 < \infty$. These are the obvious analogs of the ends in the case of one special line $(\Phi-1), \dots, (\Phi-5)$. They cancel by an identical proof to the case of one special line.

It remains to analyze the curves which appear as $t_1 - t_0 \rightarrow 0$ or $t_1 - t_0 \rightarrow \infty$. The ends which appear as $t_1 - t_0 \rightarrow \infty$ are easy to analyze: they correspond exactly to the composition $\Psi^{\bullet < \circ'} \circ \Phi^{\bullet < \circ'}$. We now focus on the ends which appear as $t_1 - t_0 \rightarrow 0$. These correspond to the lines λ_{t_0} and λ_{t_1} colliding to form a single line, for which




Figure 11.4. A (H_\wedge, t_0, t_1) -matched curve pair. On the left side, we indicate the regions where the marked points **NN**, **NC**, **M**, **SC**, and **SS** are sent.

we write $\lambda_{t'}$. The limiting curves live in a moduli space of pairs (u, u') where the points of \mathbf{SS} are perfectly matched with \mathbf{SS}' and \mathbf{NN} are perfectly matched to \mathbf{NN}' , and the rest of the marked points are matched to $\lambda_{t'}$. For fixed t' , this moduli space has expected dimension

$$\mu(\phi) + \mu(\phi') - 2|\mathbf{SS}| - 2|\mathbf{NN}| - 2|\mathbf{SC}| - 2|\mathbf{NC}| - |\mathbf{M}| - |\mathbf{M}'|.$$

By assumption, the above quantity is $-|\mathbf{M}| - |\mathbf{M}'|$. The moduli space which is parametrized over all t' has dimension $-|\mathbf{M}| - |\mathbf{M}'| + 1$, but since it also has a free \mathbb{R} action, we conclude that $-|\mathbf{M}| - |\mathbf{M}'| \geq 0$, so $|\mathbf{M}| = |\mathbf{M}'| = 0$.

There are two remaining cases for the curves appearing as $t_1 - t_0 \rightarrow 0$:

$(H_{\wedge-1})$ $|\mathbf{SC}| = |\mathbf{NC}| = |\mathbf{SC}'| = |\mathbf{NC}'| = 0.$

$(H_{\wedge-2})$ At least one of $|\mathbf{SC}|, |\mathbf{NC}|, |\mathbf{SC}'|, |\mathbf{NC}'|$ is non-empty.

For the ends $(H_{\wedge-1})$, there are no marked points along the special line and the limiting curve pair (u, u') is perfectly matched, similar to the curves counted by $\partial_{J \wedge J'}$ in equation (11.1), except (u, u') has index 0 (i.e., the expected dimension, ignoring the special line, is 0). Using transversality and expected dimension counts, we see that u and u' must both represent the constant class. These ends contribute $\text{id}_{\text{CF}^-_{J \wedge J'}(\mathcal{H}, \mathcal{H}')}$ to the left equation of (11.4).

We now consider the ends labeled $(H_{\wedge-2})$. We claim that these appear in canceling pairs. Indeed, each such end appears with even multiplicity, corresponding to switching the roles of \mathbf{NC} and \mathbf{SC} , and switching \mathbf{NC}' and \mathbf{SC}' (i.e., having the special lines pass through each other). See Figure 11.5.

The construction of the homotopy H_{\cup} is defined via the obvious modification, and a similar analysis goes through to establish the right equation of (11.4). ■




Figure 11.5. The cancellation of the ends labeled $(H_{\wedge-2})$.

11.3 Connected sums and hypercubes

We now extend the work of the previous section to handle the more general case of hypercubes of dimension greater than 0. We focus on the case that \mathcal{L}_α and $\mathcal{L}_{\alpha'}$ are 0-dimensional, i.e., $\mathcal{L}_\alpha = \alpha$ and $\mathcal{L}_{\alpha'} = \alpha'$, while \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are of arbitrary dimension.

Similar to the setting of ordinary Floer complexes in Section 11.2, it is straightforward to construct a neck-stretching homotopy equivalence

$$\mathbf{CF}_I^-(\Sigma \# \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) \simeq \mathbf{CF}_{J \wedge J'}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}).$$

The construction is essentially the same as the construction of a homotopy equivalence of hypercubes for changing the almost complex structure. Compare [12, Section 14.2].

The goal of this section is to describe a homotopy equivalence of hypercubes

$$\Phi^{\bullet < \circ'} : \mathbf{CF}_{J \wedge J'}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) \rightarrow \mathbf{CF}_{J \sqcup J'}^-(\Sigma \sqcup \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'})$$

under certain restrictions on \mathcal{L}_β and $\mathcal{L}_{\beta'}$. In the domain of $\Phi^{\bullet < \circ'}$, the curves counted are perfectly-matched. In codomain, the curves have trivial matching (i.e., they are counted in the normal way for disconnected Heegaard surfaces, as in Section 10.2). The map $\Phi^{\bullet < \circ'}$ is an extension of the map defined in Section 11.2, which we describe momentarily. Note that by Proposition 10.1, the codomain of $\Phi^{\bullet < \circ'}$ is homotopy equivalent to the tensor product of $\mathbf{CF}_J^-(\Sigma, \alpha, \mathcal{L}_\beta)$ and $\mathbf{CF}_{J'}^-(\Sigma', \alpha', \mathcal{L}_{\beta'})$, if the hypercube $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ is constructed with sufficiently small Hamiltonian translations.

Similar to the case of connected sums, we will define the map $\Phi^{\bullet < \circ'}$ to count holomorphic polygons with a special line.

To simplify the exposition and notation, it is helpful to establish a few conventions. Firstly, we recall that the holomorphic ℓ -gon maps require a choice of a family of almost complex structures $(J_y)_{y \in K_{\ell-1}}$. The space $K_{\ell-1}$ may be identified with the moduli space of disks with ℓ marked points along its boundary. For the sake of notation, it is more convenient to consider the moduli space of $\ell - 2$ marked points along the line

$$\{0\} \times \mathbb{R} \subseteq [0, 1] \times \mathbb{R}.$$

Overall translation by the \mathbb{R} -action gives equivalent elements in the moduli space. For our purposes, it is actually more convenient to not initially quotient by this action, and consider a family of almost complex structures $(J_{y,\tau})_{(y,\tau) \in K_{\ell-1} \times \mathbb{R}}$. Let T_r denote translation of the \mathbb{R} component of $\Sigma \times [0, 1] \times \mathbb{R}$ by r units. We require that

$$J_{y,\tau} = (T_r^*)(J_{y,\tau+r}),$$

for each $y \in K_{\ell-1}$, and $\tau, r \in \mathbb{R}$. Note that this does *not* imply that each $J_{y,\tau}$ is translation invariant under the \mathbb{R} -action on $[0, 1] \times \mathbb{R}$.

With this convention, the normal ℓ -gon counting maps can be viewed as counting the 1-dimensional parametrized moduli spaces

$$\bigcup_{(y,\tau) \in K_{\ell-1} \times \mathbb{R}} \mathcal{M}_{J_y}(\psi) \times \{(y, \tau)\},$$

modulo their free \mathbb{R} -action.

If $t_0 \in \mathbb{R}$, the notion of a $(\Phi, \bullet < \circ', t_0)$ -matched pair of $(\ell + 1)$ -gons is easily adapted from Definition 11.13. (Note that to define the map $\Phi^{\bullet < \circ'}$, the indexing is somewhat more natural if we count $(\ell + 1)$ -gons than ℓ -gons.) Here, t_0 denotes the height of the special line, where the marked points \mathbf{C} and \mathbf{C}' are matched. If (ψ, ψ', M) is a triple consisting of two homology classes of $(\ell + 1)$ -gons, and M is a pair of decorated source curves and matching data, and $(y, \tau) \in K_\ell \times \mathbb{R}$ and $t_0 \in \mathbb{R}$ is fixed, then the expected dimension of the moduli space of $(\Phi, \bullet < \circ', t_0)$ -matched $(\ell + 1)$ -gons, denoted $\mathcal{M}\mathcal{M}_{J_{y,\tau}}^{\Phi, \bullet < \circ'}(\psi, \psi', M, t_0)$, is given by

$$\text{ind}(\psi, \psi', M) = \mu(\psi) + \mu(\psi') - 2|\mathbf{S}| - 2|\mathbf{C}|.$$

This is because the moduli space with $J_{y,\tau}$ fixed but with no puncture constraints has expected dimension $\mu(\psi) + \mu(\psi')$, and the perfect matching below t_0 and the matching along the special line give a codimension $2(|\mathbf{S}| + |\mathbf{C}|)$ constraint.

If $J = (J_{(y,\tau)})_{(y,\tau) \in K_\ell \times \mathbb{R}}$ is a family, we define the parametrized moduli space

$$\mathcal{M}\mathcal{M}_J^{\Phi, \bullet < \circ'}(\psi, \psi', M) = \bigcup_{\substack{(y,\tau) \in K_\ell \times \mathbb{R}, \\ t_0 \in \mathbb{R}}} \mathcal{M}\mathcal{M}_{J_{(y,\tau)}}^{\Phi, \bullet < \circ'}(\psi, \psi', M, t_0) \times \{(y, \tau, t_0)\},$$

which has expected dimension

$$e \dim \mathcal{M}\mathcal{M}_J^{\Phi, \bullet < \circ'}(\psi, \psi', M) = \text{ind}(\psi, \psi', M) + \ell, \quad (11.5)$$

and has a free \mathbb{R} -action.

We now define our hypercube morphism $\Phi^{\bullet < \circ'}$. If $\kappa \leq \kappa'$ and $\kappa, \kappa' \in \mathbb{E}_{n+m}$, we define the component $\Phi_{\kappa \rightarrow \kappa'}^{\bullet < \circ'}$ of $\Phi^{\bullet < \circ'}$ as follows. If $\kappa = \kappa' = (\varepsilon, \nu)$, then the map

$$\Phi_{\kappa \rightarrow \kappa'}^{\bullet < \circ'}: \mathbf{CF}_{J \wedge J'}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \beta_\varepsilon \cup \beta'_\nu) \rightarrow \mathbf{CF}_{J \sqcup J'}^-(\Sigma \sqcup \Sigma', \alpha \cup \alpha', \beta_\varepsilon \cup \beta'_\nu)$$

is the tensor product map defined in Section 11.2.

If $\kappa_1 < \dots < \kappa_\ell$ is an increasing sequence of points in \mathbb{E}_{n+m} , then we write $\kappa_j = (\varepsilon_j, \nu_j)$ and we define a map

$$\begin{aligned} \varphi_{\kappa_1 < \dots < \kappa_\ell}^{\bullet < \circ'}: \mathbf{CF}_{J \wedge J'}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \beta_{\varepsilon_1} \cup \beta'_{\nu_1}) \\ \rightarrow \mathbf{CF}_{J \sqcup J'}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \beta_{\varepsilon_\ell} \cup \beta'_{\nu_\ell}) \end{aligned}$$

by counting $(\Phi, \bullet < \circ', t_0)$ -matched holomorphic $(\ell + 1)$ -gons which have $\text{ind}(\psi, \psi', M) = 1 - \ell$, and all of whose special inputs come from $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. By equation (11.5)

the parametrized moduli spaces of such families have expected dimension 1, and also have a free \mathbb{R} -action.

If $\kappa < \kappa'$, then we define

$$\Phi_{\kappa \rightarrow \kappa'}^{\bullet < \circ'} := \sum_{\kappa = \kappa_1 < \dots < \kappa_\ell = \kappa'} \varphi_{\kappa_1 < \dots < \kappa_\ell}$$

We may also define a map

$$\Psi^{\bullet < \circ'} : \mathbf{CF}_{J \sqcup J'}^-(\Sigma \sqcup \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) \rightarrow \mathbf{CF}_{J \wedge J'}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'})$$

in the opposite direction via the obvious modification.

Lemma 11.18. *Suppose that \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are hypercubes of beta attaching curves on Σ and Σ' , respectively, and $x \in \Sigma$ and $x' \in \Sigma'$ are choices of connected sum points. Suppose the following are satisfied:*

- (1) \mathcal{L}_β and $\mathcal{L}_{\beta'}$ are algebraically rigid.
- (2) x is base point-esque for \mathcal{L}_β .

The following are the generic codimension 1 degenerations in the parametrized moduli spaces of $(\Phi, \bullet < \circ', t_0)$ -matched $(\ell + 1)$ -gon curve pairs on $(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'})$ representing classes with matching data satisfying $\text{ind}(\psi, \psi', M) = 2 - \ell$ (i.e., of expected dimension 1 after quotienting by the free \mathbb{R} -action), where the special beta inputs are from the hypercube $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$:

- $(\Phi^\ell-1)$ *Two paired marked points in \mathbf{S} and \mathbf{S}' may collide with the special line (away from the marked points \mathbf{C} and \mathbf{C}'). None of the punctures in $\mathbf{P} \cup \mathbf{P}'$ have the same height as the special line.*
- $(\Phi^\ell-2)$ *A pair of punctures of \mathbf{C} and \mathbf{C}' may collide along λ_{t_0} (and the \mathbf{P} and \mathbf{P}' punctures have different heights than λ_{t_0}).*
- $(\Phi^\ell-3)$ *A puncture of \mathbf{N} or \mathbf{N}' may collide with λ_{t_0} (in the complement of \mathbf{C} and \mathbf{C}' , and with different height than \mathbf{P} and \mathbf{P}').*
- $(\Phi^\ell-4)$ *A puncture of \mathbf{C} may degenerate into an index 2 beta boundary degeneration or alpha boundary degeneration. The height λ_{t_0} is different from any puncture in \mathbf{P} or \mathbf{P}' .*
- $(\Phi^\ell-5)$ *The holomorphic pair (u, u') may break into two \mathbb{R} -levels, exactly one of which contains the special line.*
- $(\Phi^\ell-6)$ *Degenerations may occur which involve the punctures \mathbf{P} and \mathbf{P}' . These are constrained to the following:*
 - (a) *An index 1 disk pair may bubble at two punctures of $\mathbf{P} \cup \mathbf{P}'$. If it degenerates above the special line, then the matching is trivial. If it degenerates below the special line, the matching is perfect.*

- (b) *Four punctures of $\mathbf{P} \cup \mathbf{P}'$ collide and bubble off a pair of index 0 holomorphic triangles. If they form below the special line, the matching is perfect. If they occur above the special line, the matching is trivial.*

See Figure 11.6 for a schematic of the above degenerations.

Proof. The claim that these degenerations have codimension 1 is clear. We claim that there are no further degenerations. There are two points which require explanation:

- (P-1) There are generically no degenerations involving the punctures $\mathbf{P} \cup \mathbf{P}'$ which leave a puncture along $\{0\} \times \mathbb{R}$ of the same height as the special line λ_{t_0} .
- (P-2) The generic degenerations involving the punctures \mathbf{P} or \mathbf{P}' along $\{0\} \times \mathbb{R}$ (which necessarily occur at heights other than that of the special line by (P-1)) are constrained to those listed in $(\Phi^\ell-6)$. In particular, they are constrained only to index 1 disk bubbling and index 0 triangle degenerations. Such curves are perfectly matched if they occur below the special line, and trivially matched if they occur above the special line. Degenerations of pairs of j -gons at these punctures are prohibited for $j > 3$.

See Figure 11.7 for a schematic of a degeneration which we claim is non-generic.




Figure 11.6. The generic degenerations of $(\Phi, \bullet < \circ', t_0)$ -matched holomorphic ℓ -gon pairs. The large boundary dots denote punctures of \mathbf{P} and \mathbf{P}' .




Figure 11.7. A putative example of a degeneration of type (P-1). Big dots are boundary punctures. Small dots are marked points. The left disk is identified with \mathbb{H} .

We now argue that degenerations (P-1) are non-generic. Consider a pair (u, u') of broken curves obtained as the limit of a 1-parameter family of $(\Phi, \bullet < \circ', t_0)$ -matched $(\ell + 1)$ -gons, representing a pair of classes (ψ, ψ') , all of whose special inputs are the top degree generators. Assume that each of u and u' consist of a j -gon and a k -gon for

$$j + k = \ell + 3. \tag{11.6}$$

Such a degeneration corresponds to the codimension 1 strata of K_ℓ . Our argument extends easily to further degenerations into a tree of curves, corresponding to the higher codimension strata of K_ℓ , though we focus on the codimension 1 strata to simplify the notation.

Suppose that u degenerates into a pair u_r and u_l , where u_r is a map into $\Sigma \times [0, 1] \times \mathbb{R}$, and we view u_l as a map into $\Sigma \times \mathbb{H}$ where $\mathbb{H} = [0, \infty) \times \mathbb{R}$. Here, u_r and u_l have additional boundary punctures, and we think of u_r as a holomorphic k -gon, and we think of u_l as a holomorphic j -gon. We assume that u' similarly consists of a pair u'_l and u'_r , which are holomorphic j -gons and k -gons.

We assume for the sake of the argument that $j, k \geq 2$. Note that we automatically have $k \geq 2$, since there are the punctures at $\pm\infty$. The case that $j = 1$ corresponds to a boundary degeneration. In this case, we may fill in the boundary puncture for the curves u_r and u'_r and obtain an $(\ell + 1)$ -gon. We leave it to the reader to analyze the case of boundary degenerations by adapting the index arguments we present.

In the limit, write $\mathbf{N}_l, \mathbf{N}_r$, (resp. $\mathbf{C}_l, \mathbf{C}_r$; resp. $\mathbf{S}_l, \mathbf{S}_r$) for the marked points of u_l and u_r that lie above (resp. on; resp. below) the special line. Write \mathbf{N}'_l (and so forth) for the analogous marked points of u'_l and u'_r . Here, we use the same designation of \mathbf{N}, \mathbf{C} and \mathbf{S} as from the 1-parameter family (so in principle, it is possible for an \mathbf{N} -marked point to lie on the special line, though it will follow from our argument that these degenerations occur in too high of codimension to be generic).

By construction of the 1-parameter family relevant to the degeneration, we have

$$\mu(\psi) + \mu(\psi') + \ell - 2 - 2(|\mathbf{C}| + |\mathbf{S}|) = 0. \tag{11.7}$$

For these families, the parametrized moduli space has expected dimension 2 by equation (11.5), and also has a free \mathbb{R} -action. We note that $|\mathbf{N}| = |\mathbf{N}_l| + |\mathbf{N}_r|$, and similarly for the other collections of marked points.

We may view the pair (u_l, u'_l) as determining a map to $(\Sigma \sqcup \Sigma') \times [0, \infty) \times \mathbb{R}$. Excluding the puncture at ∞ , there are $j - 1$ punctures along $\{0\} \times \mathbb{R}$. We normalize the limit so that the special line coincides with $[0, \infty) \times \{0\}$. There is a 1-parameter family of conformal automorphisms of the half plane which preserve this line. They are multiplication by positive real numbers.

We claim that

$$\mu(\psi_l) + \mu(\psi'_l) + j - 2 - 2|\mathbf{S}_l| - \max(|\mathbf{C}_l| + |\mathbf{C}'_l|, 1) \geq 0. \tag{11.8}$$

The reasoning for the above equation is as follows. We may view the limiting curves as being elements of a moduli space of holomorphic curves which is parametrized by the moduli space of $j - 1$ marked points along $\{0\} \times \mathbb{R}$. This is a $(j - 1)$ -dimensional moduli space, however, there is a free action by $\mathbb{R} \times (0, \infty)$. The \mathbb{R} -factor translates in the i -direction, and the $(0, \infty)$ -factor scales by a real positive number. If $|\mathbf{C}_l| + |\mathbf{C}'_l| = 0$, then both of these actions preserve the moduli space of curves which are parametrized over the moduli space of $j - 1$ points along $\{0\} \times \mathbb{R}$, and normalized so that the special line is $\{0\} \times [0, \infty)$. This parametrized moduli space has expected dimension $\mu(\psi_l) + \mu(\psi'_l) + j - 1 - 2|\mathbf{S}_l|$. Hence, $\mu(\psi_l) + \mu(\psi'_l) + j - 3 - 2|\mathbf{S}_l| \geq 0$ if this moduli space is non-empty, as claimed. If $|\mathbf{C}_l| + |\mathbf{C}'_l| \geq 1$, then the \mathbb{R} -action which translates in the i -direction no longer preserves the matched moduli space, but the $(0, \infty)$ -action does preserve the matched moduli space, so instead we only obtain

$$\mu(\psi_l) + \mu(\psi'_l) + j - 2 - 2|\mathbf{S}_l| - |\mathbf{C}_l| - |\mathbf{C}'_l| \geq 0,$$

as claimed.

Similarly, (u_r, u'_r) can be viewed as living in a moduli space of k -gons where the marked points \mathbf{C}_r and \mathbf{C}'_r lie on the line λ_{t_0} , the marked points \mathbf{S}_r and \mathbf{S}'_r are perfectly matched, and also that there is a boundary puncture which has the same height as the special line. Since the expected dimension of this moduli space must be nonnegative if it admits a representative, we conclude that

$$\mu(\psi_r) + \mu(\psi'_r) + k - 3 - 2|\mathbf{S}_r| - |\mathbf{C}_r| - |\mathbf{C}'_r| \geq 0. \tag{11.9}$$

We may rearrange equations (11.7) and (11.6) to obtain

$$\begin{aligned} 0 &= \mu(\psi) + \mu(\psi') + \ell - 2 - 2(|\mathbf{C}| + |\mathbf{S}|) \\ &= \mu(\psi_l) + \mu(\psi'_l) + j - 2 - |\mathbf{C}_l| - |\mathbf{C}'_l| - 2|\mathbf{S}_l| \\ &\quad + \mu(\psi_r) + \mu(\psi'_r) + k - 3 - |\mathbf{C}_r| - |\mathbf{C}'_r| - 2|\mathbf{S}_r|. \end{aligned} \tag{11.10}$$

Next, since \mathcal{L}_β is algebraically rigid and x is base point-esque for \mathcal{L}_β , we obtain

$$\mu(\psi_l) \geq 2n_p(\psi_l) = 2|\mathbf{C}_l| + 2|\mathbf{S}_l| + 2|\mathbf{N}_l|. \tag{11.11}$$

See Definition 11.4. In particular, we obtain that

$$\begin{aligned} \mu(\psi_l) + \mu(\psi'_l) + j - 2 - |\mathbf{C}_l| - |\mathbf{C}'_l| - 2|\mathbf{S}_l| \\ \geq \mu(\psi'_l) + j - 2 + |\mathbf{C}_l| - |\mathbf{C}'_l| + 2|\mathbf{N}_l|. \end{aligned} \tag{11.12}$$

Since we are performing the $\bullet < \circ'$ degeneration, we know also that

$$|\mathbf{C}_l| \geq |\mathbf{C}'_l| \geq |\mathbf{C}_l| - 1, \tag{11.13}$$

since the left-most puncture along the special line will be in \mathbf{C}_l . By transversality for ordinary j -gons, we know that

$$\mu(\psi'_l) + j - 2 \geq 1. \tag{11.14}$$

In particular, from equations (11.12) and (11.14) we obtain that

$$\mu(\psi_l) + \mu(\psi'_l) + j - 2 - |\mathbf{C}_l| - |\mathbf{C}'_l| - 2|\mathbf{S}_l| \geq 1. \tag{11.15}$$

Equations (11.8) and (11.15) imply that equation (11.10) is the sum of two non-negative integers, one of which is at least 1. Hence the equation is never satisfiable. This implies claim (P-1).

It remains to consider degenerations into polygons where the special line does not occur at the same height as one of the punctures in $\mathbf{P} \cup \mathbf{P}'$. In such cases, $|\mathbf{C}_l| = |\mathbf{C}'_l| = 0$. Furthermore, equations (11.8) and (11.9) adapt to show that

$$\begin{aligned} \mu(\psi_l) + \mu(\psi'_l) + j - 3 - 2|\mathbf{S}_l| &\geq 0, \\ \mu(\psi_r) + \mu(\psi'_r) + k - 2 - 2|\mathbf{S}_r| - |\mathbf{C}_r| - |\mathbf{C}'_r| &\geq 0. \end{aligned} \tag{11.16}$$

Arguing similarly to equation (11.10), we see that both of the two inequalities above are equalities.

We consider first the case that (u_l, u'_l) from below the special line. In this case, we observe that $\psi_l \# \psi'_l$ may be viewed as a class on $\Sigma \# \Sigma'$ since they have the same multiplicity at the connected sum point. The excision principle for the index (i.e., deleting disks at the connected sum points) implies that

$$\mu(\psi_l \# \psi'_l) = \mu(\psi_l) + \mu(\psi'_l) - 2|\mathbf{S}_l|.$$

By the absolute grading formula, if \mathbf{z} is the outgoing intersection point of $\psi_l \# \psi'_l$, we have

$$\mu(\psi_l \# \psi'_l) = n_{\mathbf{w}''}(\psi_l \# \psi'_l) + \text{gr}(\Theta^+, \mathbf{z}),$$

since the inputs of $\psi_l \# \psi'_l$ consist of only the top degree intersection points. We conclude that $\mu(\psi_l \# \psi'_l) \geq 0$. The first line of equation (11.16) (now known to be an equality) now implies that $j - 3 \leq 0$. The only possibilities are $j = 2, 3$, as claimed. Degenerations of j -gons above the special line are handled via an essentially identical argument. ■

Lemma 11.19. *Under the assumptions in Lemma 11.18, the maps $\Phi^{\bullet < \circ'}$ and $\Psi^{\bullet < \circ'}$ are chain maps of hypercubes.*

Proof. We focus on the map $\Phi^{\bullet < \circ'}$, since the analysis for $\Psi^{\bullet < \circ'}$ is nearly identical. Suppose $\kappa \leq \kappa' \in \mathbb{E}_{n+m}$, and consider the hypercube relations for $\text{Cone}(\Phi^{\bullet < \circ'})$. If $\kappa = \kappa'$, then the hypercube relations follow from Lemma 11.15, so assume $\kappa < \kappa'$.

Suppose that $\kappa = \kappa_1 < \dots < \kappa_\ell = \kappa'$. We count the ends of 1-dimensional families of $(\Phi, \bullet < \circ', t_0)$ -matched $\ell + 1$ -gons, which have inputs from $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$. The codimension 1 degenerations are analyzed in Lemma 11.18. Many of these ends have a similar cancellation pattern as in Lemma 11.15. For example, the ends $(\Phi^\ell-1)$ and $(\Phi^\ell-2)$ cancel modulo 0. Similarly, each curve satisfying $(\Phi^\ell-3)$ appears twice in the boundary of the modulo space, unless the limiting marked points alternate between Σ and Σ' , in which case the curves cancel the degenerations in $(\Phi^\ell-4)$. The curves appearing in $(\Phi^\ell-6)$ cancel modulo 2 because of the hypercube relations for $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ (on both $\Sigma \sqcup \Sigma'$ and $\Sigma \wedge \Sigma'$).

The remaining curves are those which appear in the degenerations labeled $(\Phi^\ell-5)$. These ends correspond exactly to the hypercube relations for the mapping cone of $\Phi^{\bullet < \circ'}$. This verifies that $\Phi^{\bullet < \circ'}$ is a chain map. ■

Remark 11.20. In the above, we assumed that $\mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}$ was 0-dimensional. In the case that both $\mathcal{L}_\alpha \otimes \mathcal{L}_{\alpha'}$ and $\mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}$ have positive dimension, the argument is essentially the same. The only additional subtlety concerns for which hypercubes x and x' are base point-esque for. In the above, we used that x was base point-esque for \mathcal{L}_β to obtain equations (11.11) and (11.13) for degenerations along $\{0\} \times \mathbb{R}$. To obtain the analogous bounds for degenerations along $\{1\} \times \mathbb{R}$ we instead would need p' to be base point-esque for $\mathcal{L}_{\alpha'}$, as in the statement of Proposition 11.8.

Lemma 11.21. *Under the assumptions in Lemma 11.18, the maps $\Phi^{\bullet < \circ'}$ and $\Psi^{\bullet < \circ'}$ are homotopy inverses of each other.*

Proof. The proof is similar to the proof of Lemma 11.16. Indeed, one can define (H_\wedge, t_0) -matched holomorphic $(\ell + 1)$ -gons by analogy to the case of disks. Counting such curves gives a morphism of hypercubes

$$H_\wedge: \mathbf{CF}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) \rightarrow \mathbf{CF}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}).$$

By a straightforward adaptation of Lemma 11.18 to the case where there are two special lines, and by adapting the argument from Lemma 11.16 to handle the limiting

curves appearing as the special lines collide, we obtain exactly that the following diagram is an $(n + m + 2)$ -dimensional hypercube of chain complexes:

$$\begin{array}{ccc}
 \mathbf{CF}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) & \xrightarrow{\Phi^{\bullet < \alpha'}} & \mathbf{CF}^-(\Sigma \sqcup \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) \\
 \downarrow \text{id} & \dashrightarrow^{H_\wedge} & \downarrow \Psi^{\bullet < \alpha'} \\
 \mathbf{CF}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'}) & \xrightarrow{\text{id}} & \mathbf{CF}^-(\Sigma \wedge \Sigma', \alpha \cup \alpha', \mathcal{L}_\beta \otimes \mathcal{L}_{\beta'})
 \end{array}$$

The above diagram realizes the relation

$$\Psi^{\bullet < \alpha'} \circ \Phi^{\bullet < \alpha'} + \text{id} = [\partial, H_\wedge],$$

completing the proof. A homotopy equivalence in the other direction is defined by the obvious adaptation. ■

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*

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

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

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

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,

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

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 $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{X}} \boxtimes_{\mathcal{X}} \mathcal{X}_{\lambda_2}(K_2)$, there is a summand of the differential which corresponds to τ being output by $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{X}}$, and then input into $\mathcal{X}_{\lambda_2}(K_2)$. This summand is identified with $h_{\lambda_1}^{(Y_1, K_1)} \otimes h_{\lambda_2}^{(Y_2, K_2)}$. Similarly, there is a summand of the differential that corresponds to a σ being output by $\mathcal{X}_{\lambda_1}(K_1)^{\mathcal{X}}$, and then input into $\mathcal{X}_{\lambda_2}(K_2)$. 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

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

Here, $\mathcal{C}_\nu(L_i)$ consists of the complexes of all points of the cube \mathbb{E}_{ℓ_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

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

and

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

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

The maps F^{K_i} are $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -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*

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

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 $D_1^\nu \otimes \text{id} + \text{id} \otimes D_2^\nu$, 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_{ℓ_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}_ε , 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 $\mathbf{cCF}^-(\mathcal{H}_\#)$ 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 $\mathbf{C}_\varepsilon \subseteq \mathbf{cCF}^-(\mathcal{H}_\#)$ 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

$$I_\varepsilon: \mathcal{C}_\varepsilon \rightarrow \mathcal{S}_\varepsilon^{-1} \cdot \mathcal{C}_\varepsilon.$$

For each ε , there is also an isomorphism

$$\theta_\varepsilon: \mathcal{S}_\varepsilon^{-1} \mathcal{C}_\varepsilon \rightarrow \mathbf{C}_\varepsilon \otimes \mathbb{F}[\mathbb{H}(L)].$$

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 $A(T^s) = s$).
- (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 $\mathcal{C}_{\Lambda_1+\Lambda_2}^{-(L_1\#L_2)}$ from \mathcal{C}_ε to $\mathcal{C}_{\varepsilon'}$, and write $f_{\varepsilon, \varepsilon'}$ for the hypercube structure map for $\mathbf{cCF}^-(\mathcal{H}_\#)$ from \mathcal{C}_ε to $\mathcal{C}_{\varepsilon'}$. 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 $\mathbf{cCF}^-(\mathcal{H}_\#)$ 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 $\mathbf{cCF}^-(\mathcal{H}_\#)$.

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

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

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

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

In the above, $\mathbf{C}_\nu(L_i)$ 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 $\text{Cone}(f_i)$. This is the analog for $\mathbf{cCF}^-(\mathcal{H}_\#)$ 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 $\mathbf{C}_0(L_1) \otimes_{\mathbb{F}[U_\ell]} \mathbf{C}_0(L_2)$ as it does on $\mathbf{C}_1(L_1) \otimes_{\mathbb{F}[U_\ell]} \mathbf{C}_1(L_2)$. This is immediate from the definition. The main claim follows. ■

Chapter 13

Arc systems and the link surgery formula

For any system of arcs \mathcal{A} for a link L , the construction of Manolescu and Ozsváth gives a hypercube of a chain complex $\mathcal{C}_\Lambda(L, \mathcal{A})$, which is a module over $\mathbb{F}[[U_1, \dots, U_\ell]]$. When \mathcal{A} consists only of beta-parallel arcs, their link surgery formula states that

$$H_*(\mathcal{C}_\Lambda(L, \mathcal{A})) \cong \mathbf{HF}^-(Y_\Lambda(L)).$$

When \mathcal{A} consists only of arcs which are alpha-parallel or beta-parallel, their proof also adapts without substantial change to show the above isomorphism. However, for a general set of arcs, their work does not prove the isomorphism. In this chapter, we prove the following.

Theorem 13.1. *Suppose that \mathcal{A} and \mathcal{A}' are two systems of arcs for a framed link $L \subseteq S^3$. Suppose that \mathcal{A} differs from \mathcal{A}' by changing only the arc for the component $K_1 \subseteq L$. Then there is a homotopy equivalence of hypercubes*

$$\mathcal{C}_\Lambda(L, \mathcal{A}) \simeq \mathcal{C}_\Lambda(L, \mathcal{A}')$$

which is equivariant over $\mathbb{F}[[U_2, \dots, U_\ell]]$.

We note that our argument does not imply that the type-D modules $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}}$ are homotopy equivalent for all choices of \mathcal{A} . Instead, our proof gives the following statement about the bordered link modules.

Proposition 13.2. *Let \mathcal{A} and \mathcal{A}' be two systems of arcs for $L \subseteq S^3$, and suppose that \mathcal{A} differs from \mathcal{A}' by changing only the arc for $K_1 \subseteq L$. Then there is a homotopy equivalence of type-D modules over $\mathcal{L}_{\ell-1}$:*

$$\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}_\ell} \hat{\boxtimes} \mathcal{K}_1 \mathcal{D}_0 \simeq \mathcal{X}_\Lambda(L, \mathcal{A}')^{\mathcal{L}_\ell} \hat{\boxtimes} \mathcal{K}_1 \mathcal{D}_0.$$

In the above, \mathcal{K}_1 denotes the algebra factor corresponding to the component $K_1 \subseteq L$.

13.1 Basepoint moving maps on hypercubes

Suppose that \mathcal{L}_α and \mathcal{L}_β are two hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$. We suppose that $\gamma \subseteq \Sigma$ is an immersed loop, starting and ending at a free base point p . We now describe an endomorphism γ_* of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$.

We begin by picking a diffeomorphism of Σ , which corresponds to an isotopy which moves p in a loop along γ , and is the identity outside of a small neighborhood of γ .

There is a canonical map

$$\phi_*: \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\phi\mathcal{L}_\alpha, \phi\mathcal{L}_\beta)$$

obtained by pushing forward an intersection point tautologically under the diffeomorphism ϕ , and extending equivariantly over \mathcal{U}_i , \mathcal{V}_i and U_j . We now build a hyperbox of handleslide equivalent attaching curves

$$\phi\mathcal{L}_\alpha \longrightarrow \mathcal{L}'_\alpha \longrightarrow \mathcal{L}_\alpha. \quad (13.1)$$

We pick \mathcal{L}'_α by winding the curves of \mathcal{L}_α sufficiently so that $(\phi\mathcal{L}_\alpha, \mathcal{L}'_\alpha)$ and $(\mathcal{L}'_\alpha, \mathcal{L}_\alpha)$ are both weakly admissible. We use the standard filling procedure of Manolescu and Ozsváth [32, Lemma 8.6] to fill in the remaining chains of the above hyperbox.

By pairing the hyperbox in equation (13.1) with $\phi_*\mathcal{L}_\beta$ and compressing, we obtain a map

$$\Psi_{\phi\mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha}^{\phi\mathcal{L}_\beta}: \mathbf{CF}^-(\phi\mathcal{L}_\alpha, \phi\mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \phi\mathcal{L}_\beta).$$

A map $\Psi_{\mathcal{L}_\alpha}^{\phi\mathcal{L}_\beta \rightarrow \mathcal{L}_\beta}: \mathbf{CF}^-(\mathcal{L}_\alpha, \phi\mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ is defined similarly. We define the map γ_* as a composition

$$\gamma_* := \Psi_{\mathcal{L}_\alpha}^{\phi\mathcal{L}_\beta \rightarrow \mathcal{L}_\beta} \circ \Psi_{\phi\mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha}^{\phi\mathcal{L}_\beta} \circ \phi_*.$$

It is not hard to see that the map γ_* is well defined up to chain homotopies of hypercube morphisms by repeating the same procedure to build hypercubes of higher dimension.

Remark 13.3. The above construction may be packaged into a form of naturality for hypercubes. If $\mathcal{H} = (\mathcal{L}_\alpha, \mathcal{L}_\beta)$ and $\mathcal{H}' = (\mathcal{L}_{\alpha'}, \mathcal{L}_{\beta'})$ are two hypercubes of handleslide equivalent attaching curves, such that each set of curves in \mathcal{L}_α is handleslide equivalent to each set of curves in $\mathcal{L}_{\alpha'}$ and $\dim(\mathcal{L}_\alpha) = \dim(\mathcal{L}_{\alpha'})$, and similarly for \mathcal{L}_β and $\mathcal{L}_{\beta'}$, then the above construction produces a map $\Psi_{\mathcal{H} \rightarrow \mathcal{H}'}: \mathbf{CF}^-(\mathcal{H}) \rightarrow \mathbf{CF}^-(\mathcal{H}')$. Furthermore,

$$\Psi_{\mathcal{H}' \rightarrow \mathcal{H}} \circ \Psi_{\mathcal{H} \rightarrow \mathcal{H}'} \simeq \text{id}_{\mathbf{CF}^-(\mathcal{H})}. \quad (13.2)$$

Equation (13.2) is proven by using associativity to reduce to the case that \mathcal{H}' consist of small translates of \mathcal{L}_α and \mathcal{L}_β , and then by either adapting Lipshitz's proof in the case of triangles [23] or by using the small translate theorems for holomorphic polygons in Lemma 10.3.

We now state our formula for the map γ_* . The statement is formally similar to the case of the ordinary Floer complexes [58, Theorem D]. Our formula involves two endomorphisms Φ_w and \mathcal{A}_γ which are defined in Sections 13.3 and 13.2.

Theorem 13.4. *Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$. Let $p \in \mathbf{p}$ be a free base point and let γ be a closed path on Σ , which is based at p . If γ_* is the associated base point moving hypercube morphism, then*

$$\gamma_* \simeq \text{id} + \Phi_p \circ \mathcal{A}_\gamma$$

as morphisms of hypercubes.

The proof of the above result will be given in Section 13.6. As a first application, we can compute the effect of changing one arc in a system of arcs for a link $L \subseteq S^3$.

Corollary 13.5. *Suppose that $L \subseteq S^3$ is a framed link and that \mathcal{A} is a system of arcs for L . Let \mathcal{H} be a σ -basic system of Heegaard diagrams for (L, \mathcal{A}) . Let γ be an embedded closed loop in S^3 which is based at some $z_i \in K_i \subseteq L$ but is otherwise disjoint from L and \mathcal{A} . Assume γ is contained in the Heegaard surface for Σ . Let \mathcal{A}' be the system of arcs obtained by concatenating the arc for K_i with γ . Let us view $\mathcal{C}_\Lambda(\mathcal{H})$ as*

$$\mathcal{C}_\Lambda(\mathcal{H}) \cong \text{Cone}(\mathcal{C}_0(\mathcal{H}) \xrightarrow{F^{K_i} + F^{-K_i}} \mathcal{C}_1(\mathcal{H})), \quad (13.3)$$

where \mathcal{C}_i is a $|L| - 1$ -dimensional hypercube and F^{-K_i} denotes the sum of all hypercube maps for oriented sublinks of L which contain $-K_i$, and similarly for F^{K_i} . Then there is a σ -basic system of Heegaard diagrams \mathcal{H}' for (L, \mathcal{A}') such that

$$\mathcal{C}_\Lambda(\mathcal{H}') \simeq \text{Cone}(\mathcal{C}_0(\mathcal{H}) \xrightarrow{F^{K_i} + (\text{id} + \mathcal{V}_i^{-1} \mathcal{A}_\gamma \circ \Phi_{w_i}) F^{-K_i}} \mathcal{C}_1(\mathcal{H})), \quad (13.4)$$

where \mathcal{A}_γ and Φ_{w_i} are viewed as endomorphisms of $\mathcal{C}_1(\mathcal{H})$.

We now sketch the proof of Theorem 13.1 using Corollary 13.5 and a few other basic results from later in this chapter.

Proof of Theorem 13.1. The map Φ_{w_i} will be introduced below in more detail in Section 13.3. According to Remark 13.11, below, the map Φ_{w_i} is null-homotopic as a hypercube endomorphism via the homotopy

$$\Phi_{w_i} = \partial_{\text{Mor}}(\partial_{\mathcal{U}_i}) = \partial \circ \partial_{\mathcal{U}_i} + \partial_{\mathcal{U}_i} \circ \partial.$$

Here, ∂ denotes the hypercube differential on $\mathcal{C}_1(\mathcal{H})$, and $\partial_{\mathcal{U}_i}$ is the derivative with respect to \mathcal{U}_i (i.e., $\partial_{\mathcal{U}_i}(\mathcal{U}_i^n \mathbf{x}) = n\mathcal{U}_i^{n-1} \mathbf{x}$), when \mathbf{x} is a generator which has no \mathcal{U}_i powers. See Remark 13.11 for more details on this null-homotopy.

The null-homotopy $\partial_{\mathcal{U}_i}$ of Φ_{w_i} gives a chain homotopy equivalence between the hypercubes in equations (13.3) and (13.4). The map $\partial_{\mathcal{U}_i}$ does not commute with \mathcal{U}_i or $U_i = \mathcal{U}_i \mathcal{V}_i$, however, it will commute with \mathcal{U}_j and \mathcal{V}_j for $j \neq i$. In particular, the null-homotopy $\partial_{\mathcal{U}_i}$ cannot be used to induce a morphism between $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{X}_\ell}$ and $\mathcal{X}_\Lambda(L, \mathcal{A}')^{\mathcal{X}_\ell}$. However, after tensoring with \mathcal{D}_0 , the endomorphism $\mathbb{I} \otimes \partial_{\mathcal{U}_i}$ can be

defined, so we may package the null-homotopy of Φ_{w_i} as a homotopy equivalence of type-D modules

$$\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}_\ell} \hat{\boxtimes}_{\mathcal{X}} \mathcal{D}_0 \simeq \mathcal{X}_\Lambda(L, \mathcal{A}')^{\mathcal{L}_\ell} \hat{\boxtimes}_{\mathcal{X}} \mathcal{D}_0$$

as in the statement of Proposition 13.2. ■

We now describe the proof of Corollary 13.5.

Proof of Corollary 13.5. One first builds the surgery hyperbox for sublinks of $-L$. We compress the direction for K_i . The resulting hyperbox takes the form

$$\text{Cone}(\tilde{\mathcal{E}}_0(\mathcal{H}) \xrightarrow{F^{-K_i}} \tilde{\mathcal{E}}_1(\mathcal{H})),$$

where $\tilde{\mathcal{E}}_j(\mathcal{H})$ compresses to $\mathcal{E}_j(\mathcal{H})$. To build the hyperbox for \mathcal{H}' , we stack the above hyperbox with the base point moving map hyperbox for γ_* to form the following hyperbox:

$$(\tilde{\mathcal{E}}_0(\mathcal{H}) \xrightarrow{F^{-K_i}} \tilde{\mathcal{E}}_1(\mathcal{H}) \xrightarrow{\gamma_*} \tilde{\mathcal{E}}_1(\mathcal{H})).$$

If $\tilde{\mathcal{E}}_1(\mathcal{H})$ has dimension $n - 1$ and size \mathbf{d} , the hyperbox for γ_* has dimension n and size $\mathbf{d} \times \{1\}$.

Theorem 13.4 implies that the above hyperbox has homotopy equivalent compression to the compression of

$$(\tilde{\mathcal{E}}_0(\mathcal{H}) \xrightarrow{F^{-K_i}} \tilde{\mathcal{E}}_1(\mathcal{H}) \xrightarrow{\mathcal{V}_i^{-1} \Phi_{w_i}} \tilde{\mathcal{E}}_1(\mathcal{H}) \xrightarrow{\mathcal{A}_\gamma} \tilde{\mathcal{E}}_1(\mathcal{H})).$$

The factor of \mathcal{V}_i^{-1} in front of Φ_{w_i} is to make the map for sublinks of L which contain $-K_i$ have the correct Alexander grading, since \mathcal{A}_γ has A_i -Alexander grading 0 while Φ_{w_i} has A_i -Alexander grading 1. (Note that \mathcal{A}_γ and Φ_{w_i} will not necessarily be homogeneous in the other Alexander gradings, since they are hypercube morphisms, and hence may increment the cube directions for other link components.) We recall that for these maps, the powers of \mathcal{V}_i are determined by the overall Alexander grading change of the map. Compressing the above hyperbox gives the statement. ■

Remark 13.6. We observe that in the above corollary, there is some ambiguity in the definition of Φ_{w_i} . In practice, we would like to compute Φ_{w_i} on the link surgery hypercube, as opposed computing Φ_{w_i} using a large hyperbox whose compression is the link surgery hypercube. It is an easy consequence of the Leibniz rule that Φ_{w_i} commutes with the compression operation, so either choice gives the same answer.

The following remark indicates some subtleties of the hypercube homology actions \mathcal{A}_γ and the base point moving map formulas.

Remark 13.7. The statement of Corollary 13.5 is false if γ is not disjoint from the curves \mathcal{A} on the Heegaard diagram. An example is the genus 0 diagram for the Hopf link in Figure 16.1. In this case, composing with the base point moving map for moving one w_i around one of the components will switch between two models of the Hopf link complex (for different arc systems) which we show are non-isomorphic type-D modules in Chapter 16. On the other hand, any loop on this Heegaard diagram bounds a disk since the surface is S^3 , so Lemma 13.13 can be used to show that the homology action is null-homotopic. In this case, the hypotheses are not satisfied, since the curve $\gamma = K_2$ is not disjoint from the arc \mathcal{A}_1 for K_1 .

13.2 Basepoint actions on hypercubes

In this section, we define the base point action on hypercubes. The base point actions are analogs of natural endomorphisms which appear on the ordinary Heegaard Floer complexes, especially in the context of base point moving maps and diffeomorphism maps. See [49, 53, 58], for example.

Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$. Here and throughout, we assume \mathbf{w} and \mathbf{z} are link base points and \mathbf{p} are free base points. In this section, we will consider the Floer complex $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ which is a free module over $\mathbb{F}[\mathcal{U}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_1, \dots, \mathcal{V}_\ell, U_1, \dots, U_t]$, where \mathcal{U}_j is the variable for $w_j \in \mathbf{w}$, \mathcal{V}_j is the variable for $z_j \in \mathbf{z}$, and U_j is the variable for $p_j \in \mathbf{p}$.

If $w_i \in \mathbf{w}$, we now describe an endomorphism

$$\Phi_{w_i} : \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta).$$

If $z_i \in \mathbf{z}$, or $p_i \in \mathbf{p}$, there will be analogous endomorphisms, denoted Ψ_{z_i} and Φ_{p_i} , all defined by essentially the same construction.

We define the map Φ_{w_i} by formally differentiating the hypercube differential ∂ on $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ with respect to \mathcal{U}_i . More precisely, if \mathbf{x} and \mathbf{y} are intersection points $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ and $\partial(\mathbf{x})$ has a summand of $\mathbf{a} \cdot \mathbf{y}$ for some $\mathbf{a} \in \mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell, U_1, \dots, U_t]$, then we define $\Phi_{w_i}(\mathbf{x})$ to have a summand of $\partial_{\mathcal{U}_i}(\mathbf{a}) \cdot \mathbf{y}$. We extend Φ_{w_i} to all of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ by declaring it to be $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell, U_1, \dots, U_t]$ -equivariant.

Lemma 13.8. *The map Φ_{w_i} satisfies $\partial_{\text{Mor}}(\Phi_{w_i}) = 0$, where ∂_{Mor} is the morphism differential for hypercube morphisms (see equation (2.3)).*

Proof. The proof is the same as in the setting of the ordinary Floer chain complexes. See, e.g., [53, Lemma 3.1]. This is obtained by taking the equation $\partial^2 = 0$, and differentiating with respect to \mathcal{U}_i . The Leibniz rule yields $\Phi_{w_i} \circ \partial + \partial \circ \Phi_{w_i} = 0$. ■

Remark 13.9. In Lemma 13.8, we are implicitly using the fact that the differential commutes with the action of \mathcal{U}_i . We note that in the link surgery formula $\mathcal{C}_\Lambda(L)$, although the group is naturally a completion of a free-module over $\mathbb{F}[\mathcal{U}_i]$, the differential does not commute with \mathcal{U}_i (cf. Lemma 6.7). However, if $w \in K$ and we restrict to a subcube of $\mathcal{C}_\Lambda(L)$ where the coordinate for K is constant, then the differential does commute with \mathcal{U}_i and the map Φ_{w_i} can be defined. This is the only case that we consider in the present memoir.

When \mathcal{L}_α and \mathcal{L}_β are algebraically rigid, there is an alternate definition of Φ_{w_i} , for which we write Φ_w^0 . The map Φ_w^0 is defined as follows. Suppose that \mathcal{L}_α is m -dimensional and \mathcal{L}_β is n -dimensional. Given $\varepsilon_1 < \dots < \varepsilon_j$ in \mathbb{E}_m and $\nu_1 < \dots < \nu_k$ in \mathbb{E}_n , we define the w_i -weighted polygon counting map, $f_{\alpha_{\varepsilon_j}, \dots, \alpha_{\varepsilon_1}, \beta_{\nu_1}, \dots, \beta_{\nu_k}}^{w_i}$, to count holomorphic polygons of index $3 - j - k$ with a multiplicative weight of

$$n_{w_i}(\psi) \mathcal{U}_1^{n_{w_1}(\psi)} \mathcal{V}_1^{n_{z_1}(\psi)} \dots \mathcal{U}_\ell^{n_{w_\ell}(\psi)} \mathcal{V}_\ell^{n_{w_\ell}(\psi)} U_1^{n_{p_1}(\psi)} \dots U_t^{n_{p_t}(\psi)}.$$

We define

$$(\phi_{w_i})_{\varepsilon_1 < \dots < \varepsilon_j}^{\nu_1 < \dots < \nu_k} : \mathbf{CF}^-(\alpha_{\varepsilon_1}, \beta_{\nu_1}) \rightarrow \mathbf{CF}^-(\alpha_{\varepsilon_j}, \beta_{\nu_k}) \quad (13.5)$$

via the formula

$$(\phi_{w_i})_{\varepsilon_1 < \dots < \varepsilon_j}^{\nu_1 < \dots < \nu_k}(\mathbf{x}) := \mathcal{U}_i^{-1} f_{\alpha_{\varepsilon_j}, \dots, \alpha_{\varepsilon_1}, \beta_{\nu_1}, \dots, \beta_{\nu_k}}^{w_i} (\Theta_{\alpha_{\varepsilon_j}, \alpha_{\varepsilon_{j-1}}, \dots, \Theta_{\alpha_{\varepsilon_2}, \alpha_{\varepsilon_1}}, \mathbf{x}, \Theta_{\beta_{\nu_1}, \beta_{\nu_2}}, \dots, \Theta_{\beta_{\nu_{k-1}}, \beta_{\nu_k}}).$$

We define $\Phi_{w_i}^0 : \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ via the formula

$$\Phi_{w_i}^0 := \sum_{\substack{\nu_1 < \dots < \nu_k, \\ \varepsilon_1 < \dots < \varepsilon_j}} (\phi_{w_i})_{\varepsilon_1 < \dots < \varepsilon_j}^{\nu_1 < \dots < \nu_k}.$$

Lemma 13.10. *If \mathcal{L}_α and \mathcal{L}_β are algebraically rigid hypercubes of handleslide equivalent attaching curves, then $\Phi_{w_i} = \Phi_{w_i}^0$.*

Proof. For algebraically rigid hypercubes of attaching curves, the only special inputs for the hypercube differential are top degree intersection points, which are uniquely determined and have no \mathcal{U}_i -powers. In particular, the \mathcal{U}_i -power of an arrow in the differential of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ is determined entirely by the n_{w_i} multiplicity of the corresponding holomorphic ℓ -gon, which is exactly what the map $\Phi_{w_i}^0$ records. ■

Remark 13.11. Similar to the setting of the 3-manifold invariants, the map Φ_{w_i} is a null-homotopic hypercube endomorphism, though the natural homotopy does not commute with the action of \mathcal{U}_i . A similar comment holds for the endomorphisms Ψ_{z_i} and Φ_{p_i} , when p_i is a free base point. As in [58, equation (14.34)], we have

$$\Phi_{w_i} = [\partial, \partial_{\mathcal{U}_i}] = \partial \circ \partial_{\mathcal{U}_i} + \partial_{\mathcal{U}_i} \circ \partial, \quad (13.6)$$

where ∂ denotes the hypercube differential and ∂_{u_i} is the map

$$\partial_{u_i}(u_i^k \mathbf{x}) = k u_i^{k-1} \mathbf{x}$$

extended equivariantly over the actions of the other variables. Conceptually, we can think of equation (13.6) as an application of the Leibniz rule, in the following sense. If we think of Φ_{w_i} and ∂ as matrices with coefficients in $\mathbb{F}[u_i]$ and a generator \mathbf{x} as a column vector with entries in $\mathbb{F}[u_i]$, then $\Phi_{w_i}(\mathbf{x})$ and $\partial(\mathbf{x})$ can be thought of as matrix products. The Leibniz rule for derivatives implies that

$$\partial_{u_i}(\partial \mathbf{x}) = (\partial_{u_i}(\partial))(\mathbf{x}) + \partial(\partial_{u_i}(\mathbf{x})).$$

By definition $\partial_{u_i}(\partial) = \Phi_{w_i}$, so rearranging the above equation yields the null-homotopy $\Phi_{w_i} = [\partial, \partial_{u_i}]$.

13.3 Homology actions on hypercubes

We now recall the construction of various homological actions on hypercubes from [13, Section 6.2]. We will describe an action for both closed curves on Σ and also arcs which have boundary on two base points $\{w_1, w_2\} \subseteq \mathbf{w}$.

In the case of closed curves, the construction is an adaptation of the original case [39, Section 4.2.5]. For arcs with boundary, the construction extends [36] and [58, Section 5].

We consider first the case of a closed curve $\gamma \subseteq \Sigma$. Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z})$. We now describe an endomorphism

$$\mathcal{A}_\gamma: \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$$

as follows. Suppose

$$\psi \in \pi_2(\mathbf{x}_{\alpha_{\varepsilon_j}, \alpha_{\varepsilon_{j-1}}}, \dots, \mathbf{x}_{\alpha_{\varepsilon_2}, \alpha_{\varepsilon_1}}, \mathbf{x}, \mathbf{x}_{\beta_{v_1}, \beta_{v_2}}, \dots, \mathbf{x}_{\beta_{v_{k-1}}, \beta_{v_k}}, \mathbf{y}),$$

where each α_{ε_i} and β_{ε_i} are curves in \mathcal{L}_α or \mathcal{L}_β , respectively, and each intersection point labeled $\mathbf{x}_{\gamma, \gamma'}$ is in $\mathbb{T}_\gamma \cap \mathbb{T}_{\gamma'}$. We define the 1-chain $\partial_\alpha(\psi) \subseteq \Sigma$ to be the boundary of ψ which lies in $\alpha_{\varepsilon_1} \cup \dots \cup \alpha_{\varepsilon_j}$. We define a quantity $a(\psi, \gamma) \in \mathbb{F}$ via the formula

$$a(\psi, \gamma) = \#(\partial_\alpha(\psi) \cap \gamma).$$

We now define an endomorphism

$$\mathcal{A}_\gamma: \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta).$$

The map counts holomorphic polygons as would be counted in the hypercube differential for $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$, except with an extra multiplicative weight of $a(\psi, \gamma)$.

Lemma 13.12. *Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and γ is a closed curve on Σ . Then $\partial_{\text{Mor}}(\mathcal{A}_\gamma) = 0$, where ∂_{Mor} is the morphism differential for hypercube morphisms.*

Proof. The proof that $\partial_{\text{Mor}}(\mathcal{A}_\gamma) = 0$ follows from a Gromov compactness argument, similar to the proof of the hypercube relations for $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$.

Suppose that $\mathbf{x} \in \mathbb{T}_{\alpha_\varepsilon} \cap \mathbb{T}_{\beta_\nu}$ and $\mathbf{y} \in \mathbb{T}_{\alpha_{\varepsilon'}} \cap \mathbb{T}_{\beta_{\nu'}}$, for $\varepsilon \leq \varepsilon'$ and $\nu \leq \nu'$. We consider a class of polygons

$$\psi \in \pi_2(\mathbf{x}_{\alpha_{\varepsilon_j}, \alpha_{\varepsilon_{j-1}}}, \dots, \mathbf{x}_{\alpha_{\varepsilon_2}, \alpha_{\varepsilon_1}}, \mathbf{x}, \mathbf{x}_{\beta_{\nu_1}, \beta_{\nu_2}}, \dots, \mathbf{x}_{\beta_{\nu_{k-1}}, \beta_{\nu_k}}, \mathbf{y}),$$

where each $\mathbf{x}_{\alpha_{\varepsilon_k}, \alpha_{\varepsilon_{k-1}}}$ is an intersection point appearing as a summand of a chain from \mathcal{L}_α , and similarly for the beta labeled intersection points. We suppose that

$$\mu(\psi) = 4 - j - k$$

so that $\mathcal{M}(\psi)$ is 1-dimensional. We consider the ends of the 1-dimensional moduli space $\mathcal{M}(\psi)$.

If $j + k > 2$, the ends of $\mathcal{M}(\psi)$ generically consist of ψ breaking into a pair of holomorphic polygons (with at least 2 sides, each) representing homology classes ψ_1, ψ_2 satisfying $\psi = \psi_1 + \psi_2$. It is helpful to organize possible degenerations into three subtypes (up to reordering ψ_1 and ψ_2):

- (e-1) ψ_1, ψ_2 both have boundary on both the alpha and beta curves.
- (e-2) ψ_1 has boundary on both alpha and beta curves, while ψ_2 has boundary only on the alpha curves.
- (e-3) ψ_1 has boundary on both alpha and beta curves, while ψ_2 has boundary only on the beta curves.

See Figure 13.1 for a schematic.

We consider the sum over all sequences of intersections points $\mathbf{x}_{\alpha_{\varepsilon_j}, \alpha_{\varepsilon_{j-1}}}, \dots, \mathbf{x}_{\alpha_{\varepsilon_2}, \alpha_{\varepsilon_1}}$ and $\mathbf{x}_{\beta_{\nu_1}, \beta_{\nu_2}}, \dots, \mathbf{x}_{\beta_{\nu_{k-1}}, \beta_{\nu_k}}$ from the corresponding chains of \mathcal{L}_α and \mathcal{L}_β , as well as classes ψ , as above, of

$$a(\psi, \gamma) \# \partial \mathcal{M}(\psi) \mathcal{U}_1^{n_{w_1}(\psi)} \mathcal{V}_1^{n_{z_1}(\psi)} \dots \mathcal{U}_\ell^{n_{w_\ell}(\psi)} \mathcal{V}_\ell^{n_{z_\ell}(\psi)} U_1^{n_{p_1}(\psi)} \dots U_t^{n_{p_t}(\psi)}.$$

(We also multiply by any variables \mathcal{U}_i or \mathcal{V}_i appearing in the corresponding chains of \mathcal{L}_α or \mathcal{L}_β , but we omit this from the above equation.) We observe that if $\psi = \psi_1 + \psi_2$, as above, then

$$a(\psi, \gamma) = a(\psi_1, \gamma) + a(\psi_2, \gamma).$$

Therefore, when summing over the ends of $\# \partial \mathcal{M}(\psi)$, instead of summing each end with weight $a(\psi, \gamma)$, we can instead sum each end twice, once with weight $a(\psi_1, \gamma)$ and once with weight $a(\psi_2, \gamma)$.




Figure 13.1. The degenerations (e-1), (e-2), and (e-3) from Lemma 13.12.

Ends (e-1), summed as above, contribute the y coefficient of

$$(\partial \circ \mathcal{A}_\gamma + \mathcal{A}_\gamma \circ \partial)(x).$$

When weighted by $a(\psi_1, \gamma)$, we claim that the ends (e-2) and (e-3) cancel modulo 2 by the hypercube structure relation for \mathcal{L}_α and \mathcal{L}_β .

We will show that when weighted by $a(\psi_2, \gamma)$, the ends (e-2) and (e-3) cancel modulo 2. To establish this, we claim more generally that if ψ_2 is a class of polygons on a subdiagram of $(\Sigma, \mathcal{L}_\alpha)$ or $(\Sigma, \mathcal{L}_\beta)$, then

$$a(\psi_2, \gamma) = 0, \tag{13.7}$$

which clearly will imply the claim. If ψ_2 is a class of polygons on $(\Sigma, \mathcal{L}_\beta)$ this claim is trivial since the alpha boundary of ψ_2 is trivial. If ψ_2 is instead a class on a subdiagram of $(\Sigma, \mathcal{L}_\alpha)$, equation (13.7) follows from the fact that $\partial_\alpha(\psi_2) = \partial(D(\psi_2))$, where we are writing $D(\psi_2)$ for the domain of the class ψ_2 . We recall that $D(\psi_2)$ is the formal integral combination of the components of $\Sigma \setminus \bigcup_{\alpha \in \mathcal{L}_\alpha} \alpha$ whose multiplicity at a point $p \in \Sigma \setminus \bigcup_{\alpha \in \mathcal{L}_\alpha} \alpha$ coincides with $n_p(\psi_2)$. Therefore,

$$a(\psi_2, \gamma) = \#\partial_\alpha(\psi_2) \cap \gamma = \#\partial(D(\psi_2)) \cap \gamma \equiv \#D(\psi_2) \cap \partial\gamma = 0, \tag{13.8}$$

since $\partial\gamma = \emptyset$. This establishes equation (13.7).

Finally, we consider the ends of the moduli spaces ψ , as above, when $k = j = 1$ (i.e., the ψ is a class of disks). In this case, ends of the form (e-2) and (e-3) do not appear, though ends of the form (e-1) contribute the \mathbf{y} coefficient of $(\mathcal{A}_\gamma \circ \partial + \partial \circ \mathcal{A}_\gamma)(\mathbf{x})$ as before. In this case, there is an additional type of end which appears when $\mathbf{x} = \mathbf{y}$:

- (e-4) A constant holomorphic strip at \mathbf{x} , together with a Maslov index 2 boundary degeneration.

See [41, Section 5] for more on boundary degenerations in Heegaard Floer theory. We claim that boundary degenerations make trivial algebraic contribution to the weighted sum of the ends of the moduli spaces under consideration. This follows from the fact that if ψ_2 is a class of boundary degenerations, then the same argument as equation (13.8) establishes that $a(\psi_2, \gamma) = 0$.

Summing all ends as above, we conclude that $\partial_{\text{Mor}}(\mathcal{A}_\gamma) = 0$. ■

There is a more streamlined presentation of the proof of the above lemma. It is helpful to view γ itself as a type of Floer morphism from a set of attaching curves α to itself. We refer to γ as a *formal endomorphism* of α . We can define holomorphic polygon counts with γ as an input as follows. If $\delta_1, \dots, \delta_n$ are attaching curves, we define

$$f_{\delta_1, \dots, \delta_j, \delta_j, \dots, \delta_n}(\Theta_{1,2}, \dots, \Theta_{j-1,j}, \gamma, \Theta_{j,j+1}, \dots, \Theta_{n-1,n})$$

to count holomorphic n -gons of Maslov index $3 - n$ which are weighted by a factor or $\#\partial_{\delta_j}(\psi) \cap \gamma$. If we formally set $\partial\gamma = 0$, then it is straightforward to adapt the argument from the proof of Lemma 13.12 to see that standard associativity relations hold if we allow for a formal endomorphism as an argument.

We can extend the above formalism by defining a formal endomorphism of hypercubes of attaching curves

$$F_\gamma: \mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha.$$

The morphism L_γ has only length 1 components, all of which are γ . There are no higher length components. We will think of F_γ as a morphism of twisted complexes (see Section 2.4).

Equation (13.7) immediately implies

$$\mu_1^{\text{Tw}}(F_\gamma) = 0.$$

Furthermore, by definition

$$\mathcal{A}_\gamma = \mu_2^{\text{Tw}}(F_\gamma, -).$$

The A_∞ -associativity conditions for the category of twisted complexes implies that

$$\partial_{\text{Mor}}(\mathcal{A}_\gamma) := \mu_1^{\text{Tw}}(\mu_2^{\text{Tw}}(F_\gamma, -)) + \mu_2^{\text{Tw}}(F_\gamma, \mu_1^{\text{Tw}}(-)) = 0, \tag{13.9}$$

which is the statement of Lemma 13.12.

Another important property of the homology action is the following.

Lemma 13.13. *Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and that γ is a closed 1-chain on Σ . Suppose that $C \subseteq \Sigma$ is an integral 2-chain such that $\partial C = \gamma + S_\alpha + S_\beta$ where S_α are closed 1-chains which are disjoint from all curves in \mathcal{L}_α and S_β are closed 1-chains which are disjoint from all curves in \mathcal{L}_β . Then*

$$\mathcal{A}_\gamma \simeq 0$$

as morphisms of hypercubes.

Proof. We construct the following diagram to realize the chain homotopy:

$$\begin{array}{ccc} \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) & \xrightarrow{\mathcal{A}_\gamma} & \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \\ \downarrow \text{id} & \dashrightarrow^{H_C} & \downarrow \text{id} \\ \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) & \xrightarrow{0} & \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \end{array}$$

We define the map H_C to have only length 2 chains in the above diagram (i.e., to increment the γ - and id-directions, but no other directions), and to send $\mathbf{x} \in \mathbf{CF}^-(\alpha_\varepsilon, \beta_\nu)$ to $n_{\mathbf{x}}(C) \cdot \mathbf{x} \in \mathbf{CF}^-(\alpha_\varepsilon, \beta_\nu)$, and to be equivariant in actions of the variables.

We claim that the hypercube relations are satisfied. To see this, it is sufficient to show that if

$$\varepsilon_1 < \cdots < \varepsilon_n \quad \text{and} \quad \nu_1 < \cdots < \nu_m$$

are increasing sequences in the cubes for \mathcal{L}_α and \mathcal{L}_β , respectively, then

$$H_C \circ D_{\varepsilon_1 < \cdots < \varepsilon_n}^{\nu_1 < \cdots < \nu_m} + D_{\varepsilon_1 < \cdots < \varepsilon_n}^{\nu_1 < \cdots < \nu_m} \circ H_C = (\mathcal{A}_\gamma)_{\varepsilon_1 < \cdots < \varepsilon_n}^{\nu_1 < \cdots < \nu_m}. \quad (13.10)$$

In the above, $D_{\varepsilon_1 < \cdots < \varepsilon_n}^{\nu_1 < \cdots < \nu_m}$ denotes the summands of the hypercube differential of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$ which counts curves on $(\Sigma, \alpha_{\varepsilon_n}, \dots, \alpha_{\varepsilon_1}, \beta_{\nu_1}, \dots, \beta_{\nu_m})$.

In order to prove equation (13.10), we consider a class

$$\psi \in \pi_2(\Theta_{\varepsilon_n, \varepsilon_{n-1}, \dots, \varepsilon_2, \varepsilon_1}, \mathbf{x}, \Theta_{\nu_1, \nu_2, \dots, \nu_{m-1}, \nu_m}, \mathbf{y})$$

of $(n+m)$ -gons, and consider $\partial_\alpha(\psi) \cap C \subseteq \Sigma$. Since $\partial_\alpha(\psi)$ is 1-dimensional and C is 2-dimensional, the intersection $\partial_\alpha(\psi) \cap C$ is 1-dimensional. We compute $\partial(\partial_\alpha(\psi) \cap C)$ using the Leibniz rule for intersections, which yields

$$\begin{aligned} \partial(\partial_\alpha(\psi) \cap C) &\equiv \#(\partial \partial_\alpha(\psi)) \cap C + \# \partial_\alpha(\psi) \cap \partial C \pmod{2} \\ &\equiv n_{\mathbf{x}}(C) + n_{\mathbf{y}}(C) + \# \partial_\alpha(\psi) \cap (\gamma + S_\alpha + S_\beta) \pmod{2}. \end{aligned}$$

We first note that $\partial_\alpha(\psi) \cap S_\alpha = \partial_\beta(\psi) \cap S_\beta = \emptyset$. Furthermore, $\partial_\alpha(\psi)$ and $\partial_\beta(\psi)$ are homologous via the 2-chain $D(\psi)$, so we also have $\# \partial_\alpha(\psi) \cap S_\beta \equiv 0$. Hence

$$n_{\mathbf{x}}(C) + n_{\mathbf{y}}(C) \equiv \# \partial_\alpha(\psi) \cap \gamma \pmod{2}.$$

This implies equation (13.10), completing the proof. \blacksquare

13.4 Relative homology actions and hypercubes

We now define an endomorphism \mathcal{A}_λ for arcs λ on Σ which have boundary on two base points. We focus on the case that $\partial\lambda$ consists of two free base points, p_i and p_j , to simplify the exposition.

In the case of the 3-manifold invariants, if (Y, \mathbf{p}) is a multi-pointed 3-manifold and λ is an arc connecting two base points $p_i, p_j \in \mathbf{p}$, then there is an endomorphism

$$A_\lambda: \mathbf{CF}^-(Y, \mathbf{p}) \rightarrow \mathbf{CF}^-(Y, \mathbf{p})$$

which satisfies

$$\partial_{\text{Mor}}(A_\lambda) = U_i + U_j.$$

See [58, Section 5].

We now describe how to extend the construction of this map into the setting of hypercubes. When \mathcal{L}_α and \mathcal{L}_β are hypercubes of strongly equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and λ connects two free base points p_i, p_j , we will construct a morphism of hypercubes

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

which satisfies

$$\partial_{\text{Mor}}(\mathcal{A}_\lambda) = U_i + U_j.$$

Similar to the case of closed curves γ , we will define

$$\mathcal{A}_\lambda = \mu_2^{\text{Tw}}(F_\lambda, -)$$

for some formal endomorphism

$$F_\lambda: \mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha.$$

The map F_λ will have length 1 components equal to the formal endomorphisms $\lambda: \alpha_\varepsilon \rightarrow \alpha_\varepsilon$. However, unlike for closed curves, we will need to add higher length arrows to the map F_λ . The main issue is that if ψ is a class of ℓ -gons on a subdiagram of $(\Sigma, \mathcal{L}_\alpha, \mathbf{w}, \mathbf{z}, \mathbf{p})$, then in contrast to equation (13.7) we have

$$a(\psi, \lambda) \equiv n_{p_i}(\psi) - n_{p_j}(\psi) \pmod{2}, \quad (13.11)$$

which may be non-zero. Therefore the proof of Lemma 13.12 does not carry over without modification.

To construct F_λ and understand its formal properties, we have to first expand our notation of a formal endomorphism slightly. If α is a set of attaching curves, we will also consider formal morphisms of the form $a \cdot \mathbb{I}: \alpha \rightarrow \alpha$, where a is a polynomial in the $\mathcal{U}_i, \mathcal{V}_i$ and U_j variables. We define holomorphic polygon maps with $a \cdot \mathbb{I}$ as an

input by declaring them to be \mathcal{U}_i , \mathcal{V}_i and U_i -equivariant, and also strictly unital. That is, we declare

$$f_{\alpha_1, \dots, \alpha_j, \alpha_j, \dots, \alpha_n}(\mathbf{x}_1, \dots, a \cdot \mathbb{I}, \dots, \mathbf{x}_{n-1})$$

to vanish unless $n = 2$. We define

$$f_{\alpha_1, \alpha_2, \alpha_2}(\mathbf{x}, a \cdot \mathbb{I}) = a \cdot \mathbf{x} \quad \text{and} \quad f_{\alpha_1, \alpha_1, \alpha_2}(a \cdot \mathbb{I}, \mathbf{x}) = a \cdot \mathbf{x}.$$

We also declare $\partial(\mathbb{I}) = 0$.

We can extend the above construction to define a formal endomorphism $\mathbb{I}: \mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha$ for any hypercube of attaching curves \mathcal{L}_α . The morphism \mathbb{I} has only length 1 components, all of which are of the form $\mathbb{I}_{\alpha_\varepsilon}$.

We will construct F_λ so that

$$\mu_1^{\text{Tw}}(F_\lambda) = (U_i + U_j) \cdot \mathbb{I}.$$

Once we construct F_λ , we will define

$$A_\lambda(-) := \mu_2^{\text{Tw}}(F_\lambda, -).$$

In the following lemma, we prove that it is always possible to construct the higher length chains of F_λ when \mathcal{L}_α is a hypercube of handleslide equivalent curves.

Lemma 13.14. *Suppose that \mathcal{L}_α is a hypercube of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$, and let λ be an immersed arc on Σ .*

- (1) *If $\partial\lambda = \{p_i, p_j\}$, where p_i, p_j are free base points, then L_λ may be constructed so that*

$$\mu_1^{\text{Tw}}(F_\lambda) = U_i + U_j.$$

- (2) *More generally, if $\partial\lambda = \{x_1, x_2\} \subseteq \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$, then F_λ may be constructed so that*

$$\mu_1^{\text{Tw}}(F_\lambda) = E_{x_1}^\alpha + E_{x_2}^\alpha,$$

where $E_{x_i}^\alpha$ denotes the product of all variables for base points in the component of $\Sigma \setminus \alpha_\varepsilon$ which contains x_i . (Note that this is independent of ε , since \mathcal{L}_α consists of handleslide equivalent attaching curves.)

Remark 13.15. (1) It follows from the associativity relations for μ_i^{Tw} (as in equation (13.9)) that if \mathcal{L}_β is another hypercube of attaching curves and $\partial\lambda = \{p_i, p_j\}$ consists of free base points, then

$$\partial_{\text{Mor}}(\mathcal{A}_\lambda) = (U_i + U_j) \cdot \mathbb{I}$$

as an endomorphism of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$. The second part of the lemma has a similar consequence.

(2) As an example of the second part of the lemma, if $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ represents a link $L \subseteq Y$ for some $\alpha \in \mathcal{L}_\alpha$ and $\beta \in \mathcal{L}_\beta$, and $x_1 = w_i$ and $x_2 = w_j$ are two link base points on components K_i and K_j of L , then

$$\partial_{\text{Mor}}(\mathcal{A}_\lambda) = (\mathcal{U}_i \mathcal{V}_i + \mathcal{U}_j \mathcal{V}_j) \cdot \text{id}.$$

Proof of Lemma 13.14. We focus on the first part of the lemma, as the second part is not substantially different. The proof is described in [13, Section 6.2] when we set $U_i = U_j$, though we review it in detail since we need to prove additional properties of the maps \mathcal{A}_λ .

We begin by considering the length 2 arrows which increment the λ -direction. If α_1 is an immediate successor of α_0 in \mathcal{L}_α , then we wish to construct a chain $\eta_{\alpha_1, \alpha_0}$ so that the diagram below is a hypercube (except for the length 1 relations along λ):

$$\begin{array}{ccc} \alpha_0 & \xrightarrow{\lambda} & \alpha_0 \\ \downarrow \Theta_{\alpha_1, \alpha_0} & \dashrightarrow \eta_{\alpha_1, \alpha_0} & \downarrow \Theta_{\alpha_1, \alpha_0} \\ \alpha_1 & \xrightarrow{\lambda} & \alpha_1 \end{array}$$

Concretely, this amounts to finding $\eta_{\alpha_1, \alpha_0} \in \mathbf{CF}^-(\Sigma, \alpha_1, \alpha_0)$ such that

$$\partial \eta_{\alpha_1, \alpha_0} = A_\lambda^{\alpha_0}(\Theta_{\alpha_1, \alpha_0}) + A_\lambda^{\alpha_1}(\Theta_{\alpha_1, \alpha_0})$$

Above, the map $A_\lambda^{\alpha_0}$ weights curves by $\#(\partial_{\alpha_0}(\phi) \cap \lambda)$ while $A_\lambda^{\alpha_1}$ weights curves by $\#(\partial_{\alpha_1}(\phi) \cap \lambda)$.

We recall that for the ordinary Floer complexes,

$$A_\lambda^{\alpha_0}(\Theta_{\alpha_1, \alpha_0}) + A_\lambda^{\alpha_1}(\Theta_{\alpha_1, \alpha_0}) = U_{p_i} \Phi_{p_i}(\Theta_{\alpha_1, \alpha_0}) + U_{p_j} \Phi_{p_j}(\Theta_{\alpha_1, \alpha_0}).$$

See [58, Lemma 5.7]. However, $\Phi_{p_i}(\Theta_{\alpha_1, \alpha_0})$ and $\Phi_{p_j}(\Theta_{\alpha_1, \alpha_0})$ are cycles of a Maslov $(\text{gr}_{\mathbf{w} \cup \mathbf{p}}, \text{gr}_{\mathbf{z} \cup \mathbf{p}})$ -bigrading $(1, 1)$ larger than the canonical generator of $\mathbf{HF}^-(\Sigma, \alpha_1, \alpha_0)$, and hence are each boundaries. We note that $\eta_{\alpha_1, \alpha_0}$ may be chosen to be in the submodule of $\mathbf{CF}^-(\Sigma, \alpha_1, \alpha_0)$ generated by the ideal (U_{p_i}, U_{p_j}) .

We now define the length 3 morphisms of F_λ which increase the λ -direction. Suppose that α_{00} and α_{11} differ by two coordinates of \mathcal{L}_α . We consider the 3-dimensional cube spanned by $\alpha_{00}, \alpha_{10}, \alpha_{01}, \alpha_{11}$, extended into the λ -direction. All length 2 hypercube relations are satisfied. We sum the terms which would appear in the length 3 relation of this hypercube except for contribution of the length 3 arrow. Write

$$C_{\alpha_{11}, \alpha_{00}} \in \mathbf{CF}^-(\alpha_{11}, \alpha_{00})$$

for this element. We observe three facts:

- (C-1) $C_{\alpha_{11}, \alpha_{00}}$ is a cycle.

(C-2) $C_{\alpha_{11}, \alpha_{00}}$ has Maslov bigrading equal to that of the canonical top degree generator of $\mathbf{HF}^-(\alpha_{11}, \alpha_{00})$.

(C-3) $C_{\alpha_{11}, \alpha_{00}}$ is in the submodule generated by the ideal (U_{p_i}, U_{p_j}) .

The first two equations are easily verified. To see the third, we observe that the length 3 relation is obtained by summing over all ways of taking a consecutive sequence of morphisms in the cube, and then inserting λ between 2 terms or at the start or beginning and then applying the holomorphic polygon maps interpreted as above. As an example, the expression $C_{\alpha_{11}, \alpha_{00}}$ contains the sum

$$\begin{aligned} & f_{\alpha_{11}, \alpha_{01}, \alpha_{00}, \alpha_{00}}(\Theta_{\alpha_{11}, \alpha_{01}}, \Theta_{\alpha_{01}, \alpha_{00}}, \lambda) \\ & + f_{\alpha_{11}, \alpha_{01}, \alpha_{01}, \alpha_{00}}(\Theta_{\alpha_{11}, \alpha_{01}}, \lambda, \Theta_{\alpha_{01}, \alpha_{00}}) \\ & + f_{\alpha_{11}, \alpha_{11}, \alpha_{01}, \alpha_{00}}(\lambda, \Theta_{\alpha_{11}, \alpha_{01}}, \Theta_{\alpha_{01}, \alpha_{00}}). \end{aligned}$$

The above expression may be collapsed into a count of holomorphic triangles which are weighted by $n_{p_i}(\psi) + n_{p_j}(\psi)$. In particular, the output lies in the submodule generated by U_{p_i} and U_{p_j} . The same reasoning holds for the other terms in $C_{\alpha_{11}, \alpha_{00}}$.

Further, since α_{11} and α_{00} are handleslide equivalent, the canonical generator of $\mathbf{HF}^-(\alpha_{11}, \alpha_{00})$ has both the maximal $\text{gr}_{\mathbf{w}\cup\mathbf{p}}$ - and maximal $\text{gr}_{\mathbf{z}\cup\mathbf{p}}$ -grading of any non-zero element, and furthermore, the canonical element is the unique non-zero element of homology in this grading. Furthermore, it is straightforward to verify (e.g., by a model computation on a simple diagram) that the quotient map

$$\Pi: \mathbf{CF}^-(\alpha_{11}, \alpha_{00}) \rightarrow \mathbf{CF}^-(\alpha_{11}, \alpha_{00}) / (U_i, U_j) \tag{13.12}$$

is non-vanishing on the canonical generator of top degree. This map Π vanishes on $C_{\alpha_{11}, \alpha_{00}}$ because $C_{\alpha_{11}, \alpha_{00}}$ is in the ideal (U_i, U_j) . We note that $C_{\alpha_{11}, \alpha_{00}}$ has the same grading as the top degree element of $\mathbf{HF}^-(\alpha_{11}, \alpha_{00})$, and hence $C_{\alpha_{11}, \alpha_{00}}$, being zero on homology, must be a boundary in $\mathbf{CF}^-(\alpha_{11}, \alpha_{00})$. We pick any homogeneously graded chain $\omega_{\alpha_{11}, \alpha_{00}}$ such that $\partial\omega_{\alpha_{11}, \alpha_{00}} = C_{\alpha_{11}, \alpha_{00}}$ to be the length 3 map in our cube.

Filling in the higher-dimensional cubes proceeds in a similar, though slightly simpler manner. In this case, when we want to fill over an n -dimensional subcube for $n > 3$, all of whose morphisms except for the longest morphism are defined, the corresponding element C is a cycle of bidegree $(n - 3, n - 3)$ greater than the canonical generator, and hence will automatically be a cycle. ■

Remark 13.16. If one restricts to the case that \mathcal{L}_α is algebraically rigid, then we can take F_λ to have only length 1 components, which consist of the formal endomorphisms λ . This is because if $\psi \in \pi_2(\Theta_{\alpha_{\varepsilon_n}, \alpha_{\varepsilon_{n-1}}}, \dots, \Theta_{\alpha_{\varepsilon_2}, \alpha_{\varepsilon_1}}, \mathbf{y})$ is a class of polygons with expected dimension 0 on $(\Sigma, \alpha_{\varepsilon_n}, \dots, \alpha_{\varepsilon_1}, \mathbf{w}, \mathbf{z}, \mathbf{p})$, then the Maslov grading formula in Heegaard Floer theory implies that

$$\mu(\psi) = \text{gr}_{\mathbf{w}\cup\mathbf{p}}(\Theta_{\alpha_{\varepsilon_n}, \alpha_{\varepsilon_1}}, \mathbf{y}) + 2n_{\mathbf{w}\cup\mathbf{p}}(\psi).$$

Since \mathcal{L}_α is algebraically rigid, $\text{gr}_{\mathbf{w} \cup \mathbf{p}}(\Theta_{\alpha_{\varepsilon n}, \alpha_{\varepsilon 1}}, \mathbf{y}) \geq 0$. It follows that, if $\mu(\psi) = 3 - n$ (so that $\mathcal{M}(\psi)$ is 0-dimensional) then there are no holomorphic representatives of ψ if $n \geq 3$ and if $n = 2$ or $n = 3$, then we have $n_{\mathbf{w} \cup \mathbf{p}}(\psi) = 0$. In particular, $n_{p_i}(\psi) = n_{p_j}(\psi) = 0$ and hence

$$\#\partial_\alpha(\psi) \cap \lambda = 0$$

by equation (13.11). In particular, we can argue as in Lemma 13.12 to see that $\partial(F_\lambda) = (U_i + U_j) \cdot \mathbb{I}$, when F_λ is defined with only length 1 morphisms. Therefore, when \mathcal{L}_α is algebraically rigid, we can define $\mathcal{A}_\lambda = \mu_2^{\text{Tw}}(F_\lambda, -)$ (with F_λ having no higher length arrows) just as we defined \mathcal{A}_γ in Lemma 13.12, whereas in general we need to add higher length arrows to F_λ .

The following is a technical result which is useful when we prove properties about \mathcal{A}_λ (cf. [13, Lemma 6.2]).

Lemma 13.17. *If $\partial\lambda = \{x_1, x_2\}$, where $x_i \in \mathbf{w} \cup \mathbf{z} \cup \mathbf{p}$, then all of the higher length morphisms of F_λ may be chosen to be in the submodule spanned by (E_{x_1}, E_{x_2}) .*

Proof. We focus on the case that $\partial\lambda = \{p_i, p_j\}$ are free base points, so $E_{x_1} = U_i$ and $E_{x_2} = U_j$, since the general case is no different. In the proof of Lemma 13.14, we observed that the length 2 components of F_λ could be chosen to be in the submodule generated by the ideal (U_i, U_j) . Consider the length 3 arrows of F_λ (the argument for higher length chains is essentially the same). Following the proof of Lemma 13.14, let $C_{\alpha_{11}, \alpha_{00}}$ be the sum of the terms in the length 3 relation, excluding the term $\partial\omega_{\alpha_{11}, \alpha_{00}}$, which has yet to be defined.

We observed in Lemma 13.14 that $C_{\alpha_{11}, \alpha_{00}}$ was a boundary, and we picked some $\omega_{\alpha_{11}, \alpha_{00}}$ in bigrading $(1, 1)$ higher than the canonical generator such that $\partial\omega_{\alpha_{11}, \alpha_{00}} = C_{\alpha_{11}, \alpha_{00}}$. Write

$$\omega_{\alpha_{11}, \alpha_{00}} = U_i\omega_i + U_j\omega_j + \omega_0,$$

where ω_0 has no powers of U_i or U_j . Note that the image $\Pi(\omega)$ of ω_0 in $\mathbf{CF}^-(\alpha_{11}, \alpha_{00})/(U_i, U_j)$ is a cycle because in claim (C-3) in the proof of Lemma 13.14 we showed that $C_{\alpha_{11}, \alpha_{00}}$ is in the submodule spanned by (U_i, U_j) . Further, ω has bigrading $(1, 1)$ higher than the top degree. Hence $\Pi(\omega_0)$ is a boundary, i.e., $\Pi(\omega_0) = \partial(\mathbf{x}_0)$. We lift \mathbf{x}_0 to $\mathbf{CF}^-(\alpha_{11}, \alpha_{00})$ to get a chain \mathbf{x} of homogeneous bigrading $(1, 1)$ higher than ω_0 , so that

$$\partial\mathbf{x} = \omega_0 + U_i\mathbf{y}_i + U_j\mathbf{y}_j,$$

for some $\mathbf{y}_i, \mathbf{y}_j$. Note that

$$\partial(\omega_{\alpha_{11}, \alpha_{00}} + \partial\mathbf{x}) = C_{\alpha_{11}, \alpha_{00}}$$

and $\omega_{\alpha_{11}, \alpha_{00}} + \partial \mathbf{x}$ is in the submodule spanned by (U_i, U_j) , so we use $\omega_{\alpha_{11}, \alpha_{00}} + \partial \mathbf{x}$ in place of $\omega_{\alpha_{11}, \alpha_{00}}$. The proof for chains of length $n > 3$ follows from essentially the same logic. ■

Remark 13.18. Applying essentially the same construction as above shows that F_λ is well defined up to chain homotopy. If F_λ and F'_λ are two constructions of this morphism, then we can apply the filling construction to build a diagonal morphism h in the following diagram:

$$\begin{array}{ccc}
 \mathcal{L}_\alpha & \xrightarrow{F_\lambda} & \mathcal{L}_\alpha \\
 \downarrow \parallel & \searrow h & \downarrow \parallel \\
 \mathcal{L}_\alpha & \xrightarrow{F'_\lambda} & \mathcal{L}_\alpha
 \end{array}$$

which gives a chain homotopy of generalized endomorphisms

$$\partial_{\text{Mor}}(h) = F_\lambda + F'_\lambda.$$

13.5 Free stabilization

In [58, Section 6], the author describes maps for adding and removing base points to the Heegaard Floer complexes. These are referred to as *free-stabilization* maps. One can also define a free-stabilization map on the level of hypercubes. These are described in [13, Section 7.1], and we recall the construction presently.

Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of attaching curves on a surface $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$, and that p is a point on $\Sigma \setminus (\mathbf{w} \cup \mathbf{z} \cup \mathbf{p})$, which is not contained in any curves of \mathcal{L}_α or \mathcal{L}_β . We define stabilized hypercubes \mathcal{L}_α^+ and \mathcal{L}_β^+ . We construct the attaching curves of \mathcal{L}_α^+ and \mathcal{L}_β^+ as follows. For each α in \mathcal{L}_α , we adjoin a new component which consists of a small circle bounding a disk centered at p . We define \mathcal{L}_β^+ similarly. We translate these curves slightly to achieve admissibility. If $\varepsilon < \varepsilon'$ are points in the cube for \mathcal{L}_α , and $\Theta_{\alpha_{\varepsilon'}, \alpha_\varepsilon}$ is the corresponding chain of \mathcal{L}_α , then we define the corresponding chain in \mathcal{L}_α^+ to be $\Theta_{\alpha_{\varepsilon'}, \alpha_\varepsilon} \otimes \theta^+$, where θ^+ is the top graded generator of the new attaching curves centered at p . We make a similar definition for the chains of \mathcal{L}_β^+ .

We define a map

$$\mathcal{S}_p^+ : \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha^+, \mathcal{L}_\beta^+)$$

via the formula $\mathbf{x} \mapsto \mathbf{x} \otimes \theta^+$, extended equivariantly over the variables $\mathcal{U}_i, \mathcal{V}_i$ and U_j .

There is also a map \mathcal{S}_p^- in the opposite direction, defined dually by the formula $\mathbf{x} \otimes \theta^+ \mapsto 0$ and $\mathbf{x} \otimes \theta^- \mapsto \mathbf{x}$, extended equivariantly over the variables $\mathcal{U}_i, \mathcal{V}_i$ and U_j .

Lemma 13.19. *The hypercube free-stabilization maps, \mathcal{S}_p^+ and \mathcal{S}_p^- , are chain maps.*

The proof follows from a standard index computation and gluing argument. See [13, Proposition 5.5] for a proof in the general case, and [58, Propositions 6.5 and Theorem 6.7] for the case of holomorphic disks and triangles.

The fact that \mathcal{S}_p^+ and \mathcal{S}_p^- are chain maps of hypercubes may be interpreted as implying a form of naturality for the free-stabilization maps, similar to Remark 13.3. Suppose $\mathcal{H} = (\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta)$ and $\mathcal{H}' = (\Sigma, \mathcal{L}_{\alpha'}, \mathcal{L}_{\beta'})$ are diagrams of hypercubes of attaching curves and that the curves of \mathcal{L}_α and $\mathcal{L}_{\alpha'}$ are pairwise handleslide equivalent, and similarly for \mathcal{L}_β and $\mathcal{L}_{\beta'}$. Write \mathcal{H}_p^+ and \mathcal{H}'_p^+ for their stabilizations. Using the notation of Remark 13.3, we see that Lemma 13.19 implies the following (strict) commutation of the diagram:

$$\begin{array}{ccc} \mathbf{CF}^-(\mathcal{H}) & \xrightarrow{\Psi_{\mathcal{H} \rightarrow \mathcal{H}'}} & \mathbf{CF}^-(\mathcal{H}') \\ \downarrow \mathcal{S}_p^+ & & \downarrow \mathcal{S}_p^+ \\ \mathbf{CF}^-(\mathcal{H}_p^+) & \xrightarrow{\Psi_{\mathcal{H}_p^+ \rightarrow \mathcal{H}'_p^+}} & \mathbf{CF}^-(\mathcal{H}'_p^+) \end{array}$$

Similar to the setting of the graph TQFT, we have the following relation between the free-stabilization maps and the base point actions from Section 13.2.

Lemma 13.20. *For appropriately degenerated almost complex structures, we have $\mathcal{S}_p^+ \mathcal{S}_p^- = \Phi_p$.*

Proof. The argument is essentially the same as [13, Proposition 5.5], though we repeat the argument here for completeness. We consider a class of polygons

$$\begin{aligned} \psi \# \psi' \in \pi_2(\Theta_{\alpha_n, \alpha_{n-1}} \times \theta^+, \dots, \Theta_{\alpha_2, \alpha_1} \times \theta^+, \mathbf{x} \times \theta, \Theta_{\beta_1, \beta_2} \times \theta^+, \dots, \\ \Theta_{\beta_{m-1}, \beta_m} \times \theta^+, \mathbf{y} \times \theta') \end{aligned}$$

which potentially contributes to the differential on the hypercube $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$. Here, $\theta, \theta' \in \{\theta^+, \theta^-\}$. Here, ψ' has is a class in the stabilization region, and ψ is a class on the unstabilized diagram. By a standard computation

$$\mu(\psi \# \psi') = \mu(\psi) + 2n_p(\psi') + \text{gr}(\theta, \theta'). \quad (13.13)$$

See [58, equation 6.3]. The hypercube differential counts solutions with $\mu(\psi \# \psi') = 3 - m - n$.

In equation (13.13), the term $\text{gr}(\theta, \theta')$ is in $\{-1, 0, 1\}$. For sufficiently stretched almost complex structures, we know that $\mu(\psi) \geq 3 - m - n$ if $m + n \geq 3$, by transversality. Therefore, it follows that when $m + n \geq 3$, we have

$$(\mu(\psi), \text{gr}(\theta, \theta'), n_p(\psi')) \in \{(3 - m - n, 0, 0), (4 - m - n, -1, 0)\}.$$

In particular, $n_p(\psi) = 0$ so such curves do not contribute to Φ_p because $n_p(\psi') = 0$.

When $m + n = 2$, it is possible for $\mu(\psi) = 2 - m - n = 0$ and for ψ to have a representative, though such curves must represent the constant class. The only possibility with $n_p(\psi') = 1$ is when $\text{gr}(\theta, \theta') = -1$, $\mu(\psi) = 0$ and $n_p(\psi') = 1$. These curves have domain equal to the bigon in the stabilization region which covers p . These are the only curves which are counted by the hypercube map Φ_p , and we see that

$$\Phi_p = \mathcal{S}_p^+ \circ \mathcal{S}_p^-. \quad \blacksquare$$

The same neck-stretching argument as in Lemma 13.19 can be used to prove the following (cf. [58, Corollary 6.6]).

Lemma 13.21. *Suppose that λ is an arc or a closed path which is disjoint from p . For an appropriately degenerated almost complex structure, we have*

$$[\mathcal{A}_\lambda, \mathcal{S}_p^\pm] = 0.$$

We leave the proof of the above result to the reader, as it is very similar to Lemmas 13.19 and 13.20.

Similarly, we have the following analog of [58, Proposition 6.14], whose proof we also leave to the reader.

Lemma 13.22. *If p and p' are distinct free base points, then $[\mathcal{S}_p^{\circ_1}, \mathcal{S}_{p'}^{\circ_2}] \simeq 0$, where $\circ_1, \circ_2 \in \{+, -\}$.*

13.6 Algebraic relations and the diffeomorphism action

In this section, we prove some important algebraic relations involving the maps \mathcal{A}_λ and \mathcal{S}_p^+ . The relations we prove are analogs of the relations from the graph TQFT [58].

In the next lemma, we collect several first relations about the maps \mathcal{A}_λ and \mathcal{S}_p^\pm .

Lemma 13.23. *Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$.*

- (1) *If $p \in \mathbf{p}$, then $\Phi_p^2 \simeq 0$. The same holds for base points in $\mathbf{w} \cup \mathbf{z}$.*
- (2) *Suppose that p, p' are distinct free base points, and that λ is an arc on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$ such that $\partial\lambda = \{p, p'\}$. Then*

$$[\mathcal{A}_\lambda, \Phi_p] \simeq [\mathcal{A}_\lambda, \Phi_{p'}] \simeq \text{id}.$$

- (3) *Suppose that $p \in \mathbf{p}$ is a free base point, and $p' \in \Sigma \setminus (\mathbf{w} \cup \mathbf{z} \cup \mathbf{p})$ is a point which is not contained in \mathcal{L}_α or \mathcal{L}_β . If λ is an arc with $\partial\lambda = \{p, p'\}$, then*

$$\mathcal{S}_{p'}^- \mathcal{A}_\lambda \mathcal{S}_{p'}^+ \simeq \text{id}$$

as endomorphisms of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$. Here, we view $\mathcal{S}_{p'}^+$ as having codomain $\mathbf{CF}^-(\mathcal{L}_\alpha^+, \mathcal{L}_\beta^+, \mathbf{w}, \mathbf{z}, \mathbf{p} \cup \{p'\}) / (U_p + U_{p'})$.

Proof. The claim that $\Phi_p^2 \simeq 0$ follows from the argument in [49, Lemma 4.4]. Cf. [58, Lemma 14.8].

The proof of the second claim is essentially the same as in the case of the ordinary 3-manifold invariants [58, Lemma 14.6]. We differentiate the relation $[\partial, \mathcal{A}_\lambda] = (U_p + U_{p'}) \cdot \text{id}$ from Lemma 13.14 with respect to the U_p to obtain the relation $[\Phi_p, \mathcal{A}_\lambda] \simeq \text{id}$. The same logic gives $[\Phi_{p'}, \mathcal{A}_\lambda] \simeq \text{id}$.

The final claim that $\mathcal{S}_{p'}^- \mathcal{A}_\lambda \mathcal{S}_{p'}^+ \simeq \text{id}$ is proven in [13, Proposition 7.2]. Compare [58, Lemma 7.10] for the case of the ordinary Heegaard Floer complexes. ■

The following lemma is a hypercube analog of [39, Proposition 4.17] and [58, Lemma 5.5].

Lemma 13.24. *Let \mathcal{L}_α and \mathcal{L}_β be hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$. If γ is a closed curve on Σ then*

$$\mathcal{A}_\gamma^2 \simeq 0.$$

If λ is an arc which connects two distinct free base points $p_1, p_2 \in \mathbf{p}$, then

$$\mathcal{A}_\lambda^2 \simeq U_1 \cdot \text{id} \simeq U_2 \cdot \text{id}.$$

Proof. We begin by considering the claim for \mathcal{A}_γ , when γ is a closed curve. Our strategy will be to interpret the expression $\mu_2^{\text{Tw}}(F_\gamma, F_\gamma)$ and show that

$$\mu_2^{\text{Tw}}(F_\gamma, F_\gamma) = 0.$$

Assuming such a formalism can be established, the associativity relations for twisted complexes would show that

$$\mathcal{A}_\gamma \circ \mathcal{A}_\gamma = \mu_2^{\text{Tw}}(F_\gamma, \mu_2^{\text{Tw}}(F_\gamma, -)) = \partial_{\text{Mor}}(\mu_3^{\text{Tw}}(F_\gamma, F_\gamma, -)). \quad (13.14)$$

We now describe how to interpret $\mu_2^{\text{Tw}}(F_\gamma, F_\gamma)$ and $\mu_3^{\text{Tw}}(F_\gamma, F_\gamma, -)$. To define these, it suffices to explain how to define the holomorphic polygon counting maps when there are two inputs which are γ . Let $\delta_1, \dots, \delta_n$ be attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$. We define

$$f_{\delta_1, \dots, \delta_i, \delta_i, \dots, \delta_j, \delta_j, \dots, \delta_n}(\Theta_{1,2}, \dots, \gamma, \Theta_{i,i+1}, \dots, \Theta_{j-1,j}, \gamma, \dots, \Theta_{n-1,n})$$

(where the two γ 's are non-adjacent inputs) to count holomorphic n -gons of index $3 - n$ with a multiplicative weight of

$$(\#\partial_{\delta_i}(\psi) \cap \gamma) \cdot (\#\partial_{\delta_j}(\psi) \cap \gamma).$$

We define

$$f_{\delta_1, \dots, \delta_i, \delta_i, \delta_i, \delta_i, \dots, \delta_n}(\Theta_{1,2}, \dots, \Theta_{i-1,i}, \gamma, \gamma, \Theta_{i,i+1}, \dots, \Theta_{n-1,n})$$

to count holomorphic triangles with weight $3 - n$ and multiplicative weight

$$\frac{(\#\partial_{\delta_i}(\psi) \cap \gamma + 1)(\#\partial_{\delta_i}(\psi) \cap \gamma)}{2}.$$

The A_∞ -associativity relations are not hard to verify if we interpret $\mu_1(\gamma) = 0$ and $\mu_2(\gamma, \gamma) = 0$. This allows us to interpret all of the expressions in equation (13.14). These numerical choices can be motivated by thinking of the homology action as counting curves with marked points along their alpha boundaries. See [29, Section 3].

To establish equation (13.14), using associativity for μ_n^{Tw} , it suffices to show that

$$\mu_2^{\text{Tw}}(F_\gamma, F_\gamma) = 0. \tag{13.15}$$

To establish equation (13.15), we argue as follows. To streamline the argument, we define the quantities

$$a_i^\psi := \#(\partial_{\alpha_i}(\psi) \cap \gamma) \quad \text{and} \quad a_{i,i}^\psi := \frac{a_i^\psi(a_i^\psi + 1)}{2}.$$

The length 1 components of $\mu_2^{\text{Tw}}(F_\gamma, F_\gamma)$ are equal to $\mu_2(\gamma, \gamma)$, which is 0 by definition. The higher length components count sums as in

$$\int a_{\varepsilon_1, \dots, \alpha_{\varepsilon_i}, \alpha_{\varepsilon_i}, \dots, \alpha_{\varepsilon_j}, \alpha_{\varepsilon_j}, \dots, \alpha_{\varepsilon_n}} (\Theta_{1,2}, \dots, \gamma, \dots, \gamma, \dots, \Theta_{n-1,n}).$$

Summing over all ways of inserting two γ terms into the above expression, we see that a given class ψ is counted with weight

$$\sum_{i < j} a_i^\psi a_j^\psi + \sum_{i=1}^n a_{ii}^\psi.$$

It is straightforward to see that the above quantity is equal, modulo 2, to

$$\frac{(\#\partial_\alpha(\psi) \cap \gamma)(\#\partial_\alpha(\psi) \cap \gamma + 1)}{2}.$$

By equation (13.7), $\#\partial_\alpha(\psi) \cap \gamma \equiv 0 \pmod{2}$, so we conclude that $\mu_2^{\text{Tw}}(F_\gamma, F_\gamma) = 0$.

We now move on to the case that $\partial\lambda = \{p_1, p_2\}$ and p_1, p_2 are two free base points. In this case, we will construct a formal endomorphism $h_{\lambda,\lambda}: \mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha$ such that

$$\mu_2^{\text{Tw}}(F_\lambda, F_\lambda) + \mu_1^{\text{Tw}}(h_{\lambda,\lambda}) = U_1 \cdot \mathbb{I}. \tag{13.16}$$

When defining holomorphic polygon counts with two λ terms as inputs, we use the same convention as above, except we use

$$\mu_2(\lambda, \lambda) = U_j \cdot \mathbb{I}$$

and

$$\mu_1(\lambda) = (U_i + U_j) \cdot \mathbb{I}.$$

We return to the general question of verifying the associativity conditions on a diagram $(\Sigma, \delta_1, \dots, \delta_n, \mathbf{w}, \mathbf{z}, \mathbf{p})$. We remark that both our definition of $\mu_2(\lambda, \lambda)$ and the definition of our polygon maps when λ appears twice are asymmetric between U_i and U_j . Note also that the quantity a_i^ψ requires a choice of orientation on λ . We orient λ so that $\#(\partial(\psi) \cap \lambda) = n_{p_1}(\psi) - n_{p_2}(\psi)$. This implies that if ψ_1 and ψ_2 are the nonnegative classes of index 2 δ_j -boundary degenerations which cover p_1 and p_2 , respectively, then

$$a_{j,j}^{\psi_1} \equiv 1 \pmod{2} \quad \text{and} \quad a_{j,j}^{\psi_2} \equiv 0 \pmod{2}. \quad (13.17)$$

Furthermore, if ψ is any other nonnegative class of index 2 boundary degeneration, then

$$a_{j,j}^\psi = 0. \quad (13.18)$$

With these conventions, we now claim that the A_∞ -associativity relations hold for polygons with two λ inputs. We consider the desired associativity relations applied to the tuple $(\Theta_{1,2}, \dots, \Theta_{i-1,i}, \lambda, \dots, \lambda, \Theta_{j,j+1}, \dots, \Theta_{n-1,n})$. There are two cases to consider: $i < j$ and $i = j$. Suppose that $\psi = \psi_1 * \psi_2$ is a class of n -gons. We consider first the case that $i < j$. Here, the associativity relations follow from the fact that

$$a_i^{\psi_1} a_j^{\psi_1} + a_i^{\psi_1} a_j^{\psi_2} + a_i^{\psi_2} a_j^{\psi_1} + a_i^{\psi_2} a_j^{\psi_2} \equiv a_i^{\psi_1 * \psi_2} a_j^{\psi_1 * \psi_2} \pmod{2}.$$

Since the right-hand side is independent of the decomposition of $\psi = \psi_1 * \psi_2$, we may obtain the associativity relations in this case summing $a_i^\psi a_j^\psi \# \partial \mathcal{M}(\psi)$ (multiplied by the appropriate powers of $\mathcal{U}_i, \mathcal{V}_i$ and U_j) over n -gon classes ψ of index $4 - n$. If $n = 2$, there are additional ends corresponding to boundary degenerations. Since we are considering the case that $i < j$, this corresponds to considering the associativity relations for $(\lambda, \Theta_{1,2}, \lambda)$. In this case, boundary degenerations make trivial contribution to the associativity relations since $a_1^\psi \cdot a_2^\psi = 0$ for any nonnegative, index 2 class of boundary degenerations ψ (since ψ will only have non-trivial boundary on one of the two sets of attaching curves).

We now consider the associativity relation in the case that $i = j$. We observe the relation

$$a_{i,i}^{\psi_1} + a_{i,i}^{\psi_2} + a_i^{\psi_1} a_i^{\psi_2} \equiv a_{i,i}^{\psi_1 * \psi_2} \pmod{2}.$$

When $n > 2$, the associativity relations follow from the above relation by summing over the ends of moduli spaces of n -gons with Maslov index $4 - n$, with a weight of $a_{i,i}^\psi$. When $n = 2$, there are additional ends corresponding to boundary degenerations. These cases are relevant to the tuples $(\lambda, \lambda, \Theta_{1,2})$ and $(\Theta_{1,2}, \lambda, \lambda)$. Focus on $(\lambda, \lambda, \Theta_{1,2})$, since the other configuration is handled by a similar argument. The values of $a_{i,i}^\psi$ on index 2 boundary degenerations are described in equations (13.18)


and (13.17). When weighted by $a_{1,1}^\psi$, the total weighted count of boundary degenerations is U_1 , and this is contributed by the δ_1 -degeneration covering p_1 . Similarly, when weighted by $a_{2,2}^\psi$, the total count is also U_1 , and this is contributed by the δ_2 -degeneration covering p_1 . Both of these are encoded by the relation $\mu_2(\lambda, \lambda) = U_1 \cdot \mathbb{I}$. This completes the proof of associativity.

We now construct $h_{\lambda,\lambda}$ by a filling procedure similar to our construction of the morphism F_λ . We assume that the morphisms F_λ have already been constructed, as in Lemma 13.14. We think of $h_{\lambda,\lambda}$ as being the diagonal morphism in a diagram of the following shape:

$$\begin{array}{ccc} \mathcal{L}_\alpha & \xrightarrow{F_\lambda} & \mathcal{L}_\alpha \\ & \searrow h_{\lambda,\lambda} & \downarrow F_\lambda \\ & & \mathcal{L}_\alpha \end{array}$$

As such, we say that the *length* of a component of $h_{\lambda,\lambda}$ from α_ε to $\alpha_{\varepsilon'}$ is $|\varepsilon' - \varepsilon|_{L^1} + 2$.

We now construct the components of $h_{\lambda,\lambda}$ by induction on their length. We define the length 2 components of $h_{\lambda,\lambda}$ to vanish. We consider the problem of filling arrows of length $n > 2$ of $h_{\lambda,\lambda}$. We illustrate the case that $n = 3$ below:



We let C_{α_2, α_1} denote the terms appearing in the length 3 relation, except for $\partial\omega_{\alpha_2, \alpha_1}$ (which has yet to be defined). We make two claims:

- (1) C_{α_2, α_1} is a cycle of bidegree $(1, 1)$ less than the top degree of homology.
- (2) If F_λ is constructed as in Lemma 13.17 so that its components of length 2 or more lie in the submodule generated by the ideal (U_1, U_2) , then C_{α_2, α_1} is also in the submodule generated by the ideal (U_1, U_2) .

The claim that C_{α_2, α_1} is a cycle is an easy consequence of the associativity relations, interpreted above, which we leave to the reader.

The fact that C_{α_2, α_1} may be assumed to be in the submodule spanned by (U_1, U_2) may be seen as follows. By assumption each summand of C_{α_2, α_1} which involves $\eta_{\alpha_2, \alpha_1}$ is in this submodule. The remaining terms involve counts of holomorphic disks

of index 1 which have $\Theta_{\alpha_2, \alpha_1}$ as an input. There are multiple ways for such a holomorphic disk to contribute to C_{α_2, α_1} , depending on the locations of the two λ -inputs. We observe that the total count of a class $\phi \in \pi_2(\Theta_{\alpha_2, \alpha_1}, \mathbf{y})$ to C_{α_2, α_1} is weighted by

$$a_1^\phi a_2^\phi + a_{1,1}^\phi + a_{2,2}^\phi.$$

It is straightforward to see that the above quantity is congruent modulo 2 to

$$\frac{(a_1^\phi + a_2^\phi)(a_1^\phi + a_2^\phi + 1)}{2}.$$

We observe that $a_1^\phi + a_2^\phi = \pm(n_{p_1}(\phi) - n_{p_2}(\phi))$, since $\partial D(\phi) = \partial_{\alpha_1}(\phi) + \partial_{\alpha_2}(\phi)$. In particular, if a disk ϕ has $n_{p_1}(\phi) = n_{p_2}(\phi) = 0$, then it is weighted by 0 in C_{α_2, α_1} . Hence C_{α_2, α_1} is in the submodule spanned by (U_1, U_2) .

We have established that C_{α_2, α_1} is a cycle in bigrading $(1, 1)$ less than the canonical top degree generator, and furthermore is in the span of U_1 and U_2 . We observe that the map Π in equation (13.12) for quotienting by U_1 and U_2 is a chain map which is injective on the subspace of homology which lies in the grading of the canonical generator and also in the subspace of grading $(1, 1)$ less than the top degree of $\mathbf{HF}^-(\alpha_2, \alpha_1)$. Hence, $C_{\alpha_2, \alpha_1} = \partial\omega_{\alpha_2, \alpha_1}$ for some $\omega_{\alpha_2, \alpha_1}$ in the top degree.

We now argue that $\omega_{\alpha_2, \alpha_1}$ may be taken to be in the submodule (U_1, U_2) . The proof is essentially the same as Lemma 13.17. The image $\Pi(\omega_{\alpha_2, \alpha_1})$ is a cycle, and by possibly adding a top degree generator of $\mathbf{HF}^-(\alpha_2, \alpha_1)$, we may assume that $[\Pi(\omega_{\alpha_2, \alpha_1})] = 0$. From here, the argument of Lemma 13.17 may be applied verbatim.

From here, we continue filling the higher length chains of $h_{\lambda, \lambda}$. The argument for the higher length chains is essentially identical to the argument for the length 3 chains, so we leave it to the reader. This verifies equation (13.16). The main claim follows from this and the associativity relations for morphisms of twisted complexes of attaching curves. ■

Remark 13.25. We now make an algebraic digression. We note that the map \mathcal{A}_λ is in general only defined up to chain homotopy, and is not a cycle unless we set $U_1 = U_2$. Some care is required when composing maps which are not cycles. For example, if f_1, f_2, g_1, g_2 are in $\text{End}(C)$, (but not necessarily chain maps) and $f_1 \simeq f_2$ and $g_1 \simeq g_2$, it is not necessarily the case that $f_1 \circ g_1 \simeq f_2 \circ g_2$. For the above lemma to be meaningful, it must be the case that changing \mathcal{A}_λ by a chain homotopy also changes \mathcal{A}_λ^2 by a chain homotopy. We observe the following: If $f_1 \simeq f_2$ and $\partial_{\text{Mor}}(f_1) = \partial_{\text{Mor}}(f_2)$ is central in $\text{End}(C)$, then each f_i^2 is a chain map and

$$f_1^2 \simeq f_2^2.$$

Similarly, if f_1, f_2, g_1, g_2 are maps so that $f_1 \simeq f_2, g_1 \simeq g_2$ and $\partial_{\text{Mor}}(f_i)$ and $\partial_{\text{Mor}}(g_i)$ are central in $\text{End}(C)$, then $[f_i, g_i]$ are chain maps ($i = 1, 2$) and

$$[f_1, g_1] \simeq [f_2, g_2].$$

To see that $f_1^2 \simeq f_2^2$, we let j be a chain homotopy between f_1 and f_2 , i.e., $\partial_{\text{Mor}}(j) = f_1 + f_2$. Then

$$\partial_{\text{Mor}}(f_1 \circ j + j \circ f_2) = [\partial_{\text{Mor}}(f_1), j] + f_1^2 + f_2^2 = f_1^2 + f_2^2.$$

(Note here that $\partial_{\text{Mor}}(f_1) = \partial_{\text{Mor}}(f_2)$ since $f_1 \simeq f_2$.) Similarly, if $\partial_{\text{Mor}}(j) = f_1 + f_2$ and $\partial_{\text{Mor}}(h) = g_1 + g_2$, then

$$\begin{aligned} \partial_{\text{Mor}}([j, g_1] + [f_2, h]) &= [f_1 + f_2, g_1] + [j, \partial_{\text{Mor}}(g_1)] + [f_2, g_1 + g_2] + [\partial_{\text{Mor}}(f_2), h] \\ &= [f_1, g_1] + [f_2, g_2]. \end{aligned}$$

Lemma 13.26. *Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$, where \mathbf{p} consists of free base points. Suppose that λ_1 and λ_2 are closed curves or arcs on Σ with boundary in \mathbf{p} . Then*

$$[\mathcal{A}_{\lambda_1}, \mathcal{A}_{\lambda_2}] \simeq \sum_{p_i \in \partial\lambda_1 \cap \partial\lambda_2} U_i.$$

Proof. The proof is similar to the proof of Lemma 13.24. We are going to construct a formal endomorphism

$$h_{[\lambda_1, \lambda_2]}: \mathcal{L}_\alpha \rightarrow \mathcal{L}_\alpha$$

which satisfies

$$\mu_1^{\text{Tw}}(h_{[\lambda_1, \lambda_2]}) = \mu_2^{\text{Tw}}(F_{\lambda_1}, F_{\lambda_2}) + \mu_2^{\text{Tw}}(F_{\lambda_2}, F_{\lambda_1}). \tag{13.19}$$

We think of constructing $h_{[\lambda_1, \lambda_2]}$ as corresponding to building a diagram of the form

$$\mathcal{L}_\alpha^{[\lambda_1, \lambda_2]} := \begin{array}{ccc} \mathcal{L}_\alpha & \xrightarrow{\lambda_1} & \mathcal{L}_\alpha \\ | & \searrow & | \\ \lambda_2 & h_{[\lambda_1, \lambda_2]} & \lambda_2 \\ \downarrow & & \downarrow \\ \mathcal{L}_\alpha & \xrightarrow{\lambda_1} & \mathcal{L}_\alpha \end{array}$$

which satisfies a suitable version of the hypercube relations.

Similar to our proof of Lemma 13.24, we need to describe how to interpret the holomorphic polygon counts where we formally have both λ_1 and λ_2 as an input.

We have already defined the holomorphic polygon maps when there is a single input of λ_1 or λ_2 . When there are two inputs, we make the following definitions. Suppose $\delta_1, \dots, \delta_n$ is a collection of attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$. If $j < k$, $s, t \in \{1, 2\}$, and $s \neq t$, we define

$$f_{\delta_1, \dots, \delta_j, \delta_j, \dots, \delta_k, \delta_k, \dots, \delta_n}(\Theta_{1,2}, \dots, \lambda_s, \dots, \lambda_t, \dots, \Theta_{n-1,n})$$

to count polygons weighted by $\#(\partial_{\delta_j}(\psi) \cap \lambda_s)\#(\partial_{\delta_k}(\psi) \cap \lambda_t)$. Next, we define

$$f_{\delta_1, \dots, \delta_j, \delta_j, \delta_j, \dots, \gamma_n}(\Theta_{1,2}, \dots, \lambda_2, \lambda_1, \dots, \Theta_{n-1,n})$$

to also count curves weighted by $\#(\partial_{\delta_j}(\psi) \cap \lambda_2)\#(\partial_{\delta_j}(\psi) \cap \lambda_1)$. Finally, we define

$$f_{\delta_1, \dots, \delta_j, \delta_j, \delta_j, \dots, \delta_n}(\Theta_{1,2}, \dots, \lambda_1, \lambda_2, \dots, \Theta_{n-1,n}) = 0.$$

We claim that these satisfy a slightly restricted version of the A_∞ -associativity relation. Instead of showing an A_∞ -associativity relation on all tuples, we will show that associativity is satisfied on only some generators. We will show that associativity holds on tuples of the form

$$(\Theta_{1,2}, \dots, \Theta_{j-1,j}, \lambda_t, \dots, \lambda_s, \Theta_{k,k+1}, \dots, \Theta_{n-1,n}) \quad (13.20)$$

whenever $j < k$ and $t \neq s$. Additionally, we will show that the associativity relations also hold on

$$\begin{aligned} &(\Theta_{1,2}, \dots, \Theta_{j-1,j}, \lambda_1, \lambda_2, \Theta_{j,j+1}, \dots, \Theta_{n-1,n}) \\ &+ (\Theta_{1,2}, \dots, \Theta_{j-1,j}, \lambda_2, \lambda_1, \Theta_{j,j+1}, \dots, \Theta_{n-1,n}) \end{aligned} \quad (13.21)$$

as long as we make the following conventions: We assume the polygon maps are strictly unital with the λ -inputs; $\partial\lambda_i = 0$; and $[\lambda_1, \lambda_2] = \sum_{p_i \in \partial\lambda_1 \cap \partial\lambda_2} U_i$. This weaker notion of associativity is sufficient for the purposes of constructing $h_{[\lambda_1, \lambda_2]}$ and proving equation (13.19). Note that instead of the above approach, we could attempt to prove associativity for general tuples, but this would require defining the polygon maps when there are inputs such as $\lambda_1 \cdot \lambda_2$, which is an unnecessary detour.

The proof of associativity, as described above, is broken into several cases. The first case is when $n > 2$ and $j < k$. In this case, the generic degenerations of a 1-dimensional family of holomorphic n -gons will consist of pairs of a p -gon and a q -gon for $p + q = n + 2$. Write ϕ and ψ for the classes of the p - and q -gons, respectively. If $s \in \{1, 2\}$, write a_j^ϕ for $\#(\partial_{\delta_j}(\phi) \cap \lambda_1)$ and b_j^ϕ for $\#(\partial_{\delta_j}(\phi) \cap \lambda_2)$. In this case, the associativity relation applied to tuples as in equation (13.20) follows from the additive relation

$$a_j^\phi b_k^\psi + a_j^\psi b_k^\phi + a_j^\phi b_k^\phi + a_j^\psi b_k^\psi = a_j^{\phi*\psi} b_k^{\phi*\psi}.$$

Associativity for tuples as in equation (13.21), when $j = k$ and $n > 2$, follows from the similar additive relation

$$a_j^\phi b_j^\psi + a_j^\psi b_j^\phi + a_j^\phi b_j^\phi + a_j^\psi b_j^\psi = a_j^{\phi*\psi} b_j^{\phi*\psi}.$$

Lastly, in the case that $n = 2$, there are additional ends corresponding to boundary degenerations. These ends contribute in the associativity relations for $(\Theta_{1,2}, \lambda_1, \lambda_2) +$

$(\Theta_{1,2}, \lambda_2, \lambda_1)$. The relation $[\lambda_1, \lambda_2] = \sum_{p \in \partial\lambda_1 \cap \partial\lambda_2} U_p$ corresponds to the fact that $a_2^\phi b_2^\phi \equiv 1$ if ϕ is a nonnegative, index 2 class of δ_2 -boundary degenerations which covers a point of $\partial\lambda_1 \cap \partial\lambda_2$. The associativity relations for $(\lambda_1, \lambda_2, \Theta_{1,2}) + (\lambda_2, \lambda_1, \Theta_{1,2})$ are similar. Finally, we consider the relations for $(\lambda_1, \Theta_{1,2}, \lambda_2)$. In this case, boundary degenerations do not make any contribution, since $a_1^\phi b_2^\phi \equiv 0$ for any boundary degeneration because a boundary degeneration has boundary on only one set of attaching curves.

Having described the necessary associativity relations, filling the higher length chains of $h_{[\lambda_1, \lambda_2]}$ and verifying equation (13.19) follows the same line of reasoning as in the proof of Lemma 13.24. We leave the remaining details of the construction to the reader. ■

We prove a final lemma.

Lemma 13.27. *Suppose that λ_1 and λ_2 are two 1-chains on $(\Sigma, \mathbf{w}, \mathbf{z}, \mathbf{p})$, such that $\partial\lambda_1 = \{x_1, x_2\}$ and $\partial\lambda_2 = \{x_2, x_3\}$. Then*

$$\mathcal{A}_{\lambda_1} + \mathcal{A}_{\lambda_2} \simeq \mathcal{A}_{\lambda_1 * \lambda_2}.$$

Proof. We observe first that for any class of curves ϕ which has boundary on α_i , we have

$$\#(\partial_{\alpha_i}(\phi) \cap (\lambda_1 + \lambda_2)) = \#(\partial_{\alpha_i}(\phi) \cap \lambda_1) + \#(\partial_{\alpha_i}(\phi) \cap \lambda_2).$$

In particular, we may take the formal endomorphism $F_{\lambda_2 * \lambda_1}$ to be $F_{\lambda_2} + F_{\lambda_1}$. With this definition, $\mathcal{A}_{\lambda_1 * \lambda_2} = \mathcal{A}_{\lambda_1} + \mathcal{A}_{\lambda_2}$. Since $F_{\lambda_1 * \lambda_2}$ is well defined up to chain homotopy by Remark 13.18, the claim in the statement follows for an arbitrary construction of $F_{\lambda_2 * \lambda_2}$. ■

We now prove a helpful formula for understanding base point moving maps on the link surgery hypercubes. See [58, Theorem 14.11] for an analogous result in the setting of the ordinary Heegaard Floer complexes.

Lemma 13.28. *Let $\mathcal{H} = (\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$, where \mathcal{L}_α and \mathcal{L}_β are hypercubes of handleslide equivalent attaching curves, $\mathbf{w} \cup \mathbf{z}$ are link base points and \mathbf{p} are free base points. Suppose that $p_1 \in \mathbf{p}$, and that $p_2 \in \Sigma \setminus \mathbf{p}$ is in the complement of the curves in \mathcal{L}_α and \mathcal{L}_β . Suppose also that $(\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \mathbf{w}, \mathbf{z}, \mathbf{p})$ and $(\Sigma, \mathcal{L}_\alpha, \mathcal{L}_\beta, \{p_2\} \cup \mathbf{p} \setminus \{p_1\})$ are both weakly admissible. Let $\mathcal{H}_{p_i}^+$ denote the diagram obtained by free-stabilizing \mathcal{H} at p_i . If λ is a path from p_1 to p_2 , then*

$$S_{p_1}^- \Psi_{\mathcal{H}_{p_2}^+ \rightarrow \mathcal{H}_{p_1}^+} \mathcal{A}_\lambda S_{p_2}^+$$

is chain homotopic to the diffeomorphism map which moves p_1 to p_2 along λ . In the above, we work on the complexes where we have identified U_{p_1} and U_{p_2} so that \mathcal{A}_λ is a chain map. Also $\Psi_{\mathcal{H}_{p_2}^+ \rightarrow \mathcal{H}_{p_1}^+}$ denotes the map defined in Section 13.1.

Proof. We give a slightly simpler proof than appeared in the setting of 3-manifold invariants in [58, Theorem 14.11], as follows. Let q be a point very near to p_2 . We will view λ as the concatenation of a path λ_1 from p_1 to q , and a path λ_2 from q to p_2 . In the following, we drop the transition maps $\Psi_{\mathfrak{I}c_{p_2}^+ \rightarrow \mathfrak{I}c_{p_1}^+}$ from the notation. We will write U for

$$U = U_{p_1} = U_{p_2}.$$

Let ϕ_* denote the point pushing a diffeomorphism map which moves p_2 to p_1 along λ , composed with the naturality map. It is sufficient to show that

$$\text{id} \simeq \phi_* \mathcal{S}_{p_1}^- \mathcal{A}_\lambda \mathcal{S}_{p_2}^+. \quad (13.22)$$

As a first step, we claim that the map $\mathcal{S}_{p_1}^- \mathcal{A}_\lambda \mathcal{S}_{p_2}^+$ is functorial under concatenation of arcs, i.e., that

$$\mathcal{S}_{p_1}^- \mathcal{A}_\lambda \mathcal{S}_{p_2}^+ \simeq \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \mathcal{S}_{p_2}^+ \mathcal{S}_{p_1}^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+. \quad (13.23)$$

To establish equation (13.23), first note that we may commute free-stabilization maps with other terms by Lemmas 13.21 and 13.22 to obtain

$$\begin{aligned} \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \mathcal{S}_{p_2}^+ \mathcal{S}_{p_1}^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ \simeq \mathcal{S}_{p_1}^- \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \mathcal{S}_{p_2}^+. \end{aligned}$$

Next, we use the relations $\mathcal{A}_{\lambda_i}^2 \simeq U$ and $\mathcal{A}_{\lambda_1 + \lambda_2} = \mathcal{A}_{\lambda_1} + \mathcal{A}_{\lambda_2}$ from Lemmas 13.24 and 13.27 to see that

$$\mathcal{S}_{p_1}^- \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \mathcal{S}_{p_2}^+ \simeq \mathcal{S}_{p_1}^- \mathcal{S}_q^- (\mathcal{A}_{\lambda_2 + \lambda_1} \mathcal{A}_{\lambda_1} + U) \mathcal{S}_q^+ \mathcal{S}_{p_2}^+.$$

We use the relation $\mathcal{S}_q^- \mathcal{S}_q^+ = 0$ (which is immediate from the formulas for \mathcal{S}_q^+ and \mathcal{S}_q^-) and the fact that $\lambda_2 + \lambda_1 = \lambda$ to see that the above is homotopic to

$$\mathcal{S}_{p_1}^- \mathcal{S}_q^- \mathcal{A}_\lambda \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \mathcal{S}_{p_2}^+.$$

Next, we use the fact that λ is disjoint from q to commute \mathcal{S}_q^- with \mathcal{A}_λ using Lemma 13.21, and then we use the relation $\mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \simeq \text{id}$ from Lemma 13.23 to see that the above equation is homotopic to

$$\mathcal{S}_{p_1}^- \mathcal{A}_\lambda \mathcal{S}_{p_2}^+.$$

This establishes equation (13.23).

To obtain the main statement in equation (13.22), we combine equation (13.23) with the naturality of all of the other maps with respect to diffeomorphisms, via the following manipulation:

$$\begin{aligned} \phi_* \mathcal{S}_{p_1}^- \mathcal{A}_\lambda \mathcal{S}_{p_2}^+ &\simeq \phi_* \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \mathcal{S}_{p_2}^+ \mathcal{S}_{p_1}^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \mathcal{S}_q^- \mathcal{A}_{\phi(\lambda_2)} \phi_* \mathcal{S}_{p_2}^+ \mathcal{S}_{p_1}^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \mathcal{S}_q^- \mathcal{A}_{\phi(\lambda_2)} \mathcal{S}_{p_1}^+ \mathcal{S}_{p_1}^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+. \end{aligned}$$

We may easily arrange that $\phi(\lambda_2)$ and λ_1 are isotopic on Σ (relative to its boundary). Hence $\mathcal{A}_{\phi(\lambda_2)} \simeq \mathcal{A}_{\lambda_1}$ by Lemma 13.13. The expression is easily shown to be homotopic to the identity using the graph TQFT relations, as we briefly recall. We obtain from Lemma 13.20 that

$$\mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{S}_{p_1}^+ \mathcal{S}_{p_1}^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \simeq \mathcal{S}_q^- \mathcal{A}_{\lambda_1} \Phi_{p_1} \mathcal{A}_{\lambda_1} \mathcal{S}_q^+.$$

The above expression is homotopic to $\mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{A}_{\lambda_1} \Phi_{p_1} \mathcal{S}_q^+ + \mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+$ by Lemma 13.23. By Lemma 13.24 we see that $\mathcal{S}_q^- \mathcal{A}_{\lambda_1}^2 \Phi_{p_1} \mathcal{S}_q^+ \simeq U \mathcal{S}_q^- \mathcal{S}_q^+ \Phi_{p_1} = 0$ and by Lemma 13.23 we have that $\mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \simeq \text{id}$. This completes the proof. ■

Remark 13.29. The above argument also applies to the setting of ordinary Heegaard Floer complexes to give a shorter proof of the normalization axiom of the graph TQFT [58, Theorem 14.11].

We now prove the main base point moving map result.

Proof of Theorem 13.4. The proof follows from the same manipulation as in the setting of the graph TQFT, see [58, Section 14.4]. We produce the argument for the reader. We decompose γ into the concatenation of two arcs λ_1, λ_2 , and view λ_1 as a path from p to a new base point q , and λ_2 as a path from q to p . In all of the Floer complexes we consider, we identify the variables for p and q with a single variable U . We compute

$$\begin{aligned} \gamma_* &\simeq \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \mathcal{S}_p^+ \mathcal{S}_p^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \mathcal{S}_q^- \mathcal{A}_{\lambda_2} \Phi_p \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ + \mathcal{S}_q^- \Phi_p \mathcal{A}_{\lambda_2} \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \text{id} + \Phi_p \mathcal{S}_q^- (\mathcal{A}_{\lambda_1} + \mathcal{A}_\gamma) \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \text{id} + \Phi_p \mathcal{S}_q^- U \mathcal{S}_q^+ + \Phi_p \mathcal{A}_\gamma \mathcal{S}_q^- \mathcal{A}_{\lambda_1} \mathcal{S}_q^+ \\ &\simeq \text{id} + \Phi_p \mathcal{A}_\gamma, \end{aligned}$$

completing the proof. ■

13.7 Relating homology actions and base point actions

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard link diagram for (Y, L) and that $K_i \subseteq L$ is a knot component which has base points $w_i \in \mathbf{w}$ and $z_i \in \mathbf{z}$. We may represent K_i on Σ as an immersed curve $[K_i]$ on Σ which is the concatenation of two arcs λ_α and λ_β which both have boundary $\{w_i, z_i\}$. The arc λ_α intersects only alpha curves while λ_β intersects only beta curves. See Figure 13.2. If ϕ is a class of disks, then

$$a(\phi, \gamma) = n_{w_i}(\phi) - n_{z_i}(\phi) \pmod{2}. \quad (13.24)$$




Figure 13.2. A knot shadow K_i (dashed) decomposed as a concatenation of two subarcs λ_α and λ_β .

The map $A_{[K_i]}$ weights a holomorphic disk ϕ by $a(\phi, \gamma)$. On the other hand, we can define a map which counts holomorphic disks weighted by $n_{w_i}(\phi)$. This map is clearly equal to $\mathcal{U}_i \Phi_{w_i}$. Similarly, the map which counts disks weighted by $n_{z_i}(\phi)$ is equal to $\mathcal{V}_i \Psi_{z_i}$. Therefore, equation (13.24) translates into the equality

$$A_{[K_i]} = \mathcal{U}_i \Phi_{w_i} + \mathcal{V}_i \Psi_{z_i},$$

as endomorphisms of $\mathcal{CF}\mathcal{L}(Y, L)$. Compare [55, Lemma 14.12]. In this section, we prove a hypercube version of the above statement.

We consider hypercubes of handleslide equivalent attaching curves \mathcal{L}_α and \mathcal{L}_β on $(\Sigma, \mathbf{w}, \mathbf{z})$. We suppose that there is a pair of arcs $\lambda_\alpha, \lambda_\beta \subseteq \Sigma$ such that

$$\partial\lambda_\alpha = \partial\lambda_\beta = \{w_i, z_i\}, \quad K_i = \lambda_\alpha * \lambda_\beta,$$

and such that λ_α intersects only curves from \mathcal{L}_α while λ_β intersects only curves from \mathcal{L}_β . Write K_i for the concatenation of λ_α and λ_β . We recall from Section 13.3 that there is an endomorphism

$$\mathcal{A}_{[K_i]}: \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta) \rightarrow \mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta).$$

Since K_i is a closed curve on Σ , the map $\mathcal{A}_{[K_i]}$ counts holomorphic polygons (as would be counted by the ordinary hypercube differential) with a weight of $a(\phi, K_i)$, where $a(\phi, K_i)$ is the sum over $\#(\partial_{\alpha_\varepsilon} \phi \cap K)$ ranging over all α_ε in \mathcal{L}_α .

In Section 13.3, we constructed base point actions Φ_{w_i} and Ψ_{z_i} on the hypercube $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$. The map Φ_{w_i} was defined as the formal differential of the hypercube differential with respect to \mathcal{U}_i . We note that in general powers of \mathcal{U}_i are contributed by both the holomorphic curves in the hypercube differential, and also the powers of \mathcal{U}_i present in the input chains of the polygon maps used to construct the hypercube. We defined a map $\Phi_{w_i}^0$ only taking into account the multiplicities of holomorphic curves at the base point w_i . We proved in Lemma 13.10 that if \mathcal{L}_α and \mathcal{L}_β are algebraically rigid, then $\Phi_{w_i}^0 = \Phi_{w_i}$ and $\Psi_{w_i}^0 = \Psi_{w_i}$. In general, this is not the case. Nonetheless, we do have the following.

Lemma 13.30. *Suppose that \mathcal{L}_α and \mathcal{L}_β are hypercubes of handleslide equivalent attaching curves on $(\Sigma, \mathbf{w}, \mathbf{z})$, and $K_i \subseteq \Sigma$ is a concatenation of two arcs λ_α and λ_β , as above, where λ_α intersects only the curves in \mathcal{L}_α , and \mathcal{L}_β intersects only the curves in \mathcal{L}_β . Then $\mathcal{U}_i \Phi_{w_i}^0 + \mathcal{V}_i \Phi_{z_i}^0$ is a chain map (regardless of whether \mathcal{L}_α and \mathcal{L}_β are algebraically rigid) and*

$$\mathcal{A}_{[K_i]} = \mathcal{U}_i \Phi_{w_i}^0 + \mathcal{V}_i \Psi_{z_i}^0 = \mathcal{U}_i \Phi_{w_i} + \mathcal{V}_i \Psi_{z_i},$$

as hypercube morphisms.

Proof. Since $a(\phi, K_i) \equiv n_{w_i}(\phi) - n_{z_i}(\phi) \pmod{2}$ for any holomorphic polygon counted by the hypercube differential of $\mathbf{CF}^-(\mathcal{L}_\alpha, \mathcal{L}_\beta)$, we conclude that

$$\mathcal{A}_{[K_i]} = \mathcal{U}_i \Phi_{w_i}^0 + \mathcal{V}_i \Psi_{z_i}^0,$$

since the two sides count the same holomorphic curves, with the same weights. Since the left-hand side is a chain map by Lemma 13.12, the right-hand side must also be a chain map.

It remains to show that

$$\mathcal{U}_i \Phi_{w_i}^0 + \mathcal{V}_i \Psi_{z_i}^0 = \mathcal{U}_i \Phi_{w_i} + \mathcal{V}_i \Psi_{z_i}.$$

To see this, we observe that $\mathcal{U}_i \Phi_{w_i} + \mathcal{V}_i \Psi_{z_i}$ may be described as summing over sequences $\varepsilon_1 < \dots < \varepsilon_n$ and $\nu_1 < \dots < \nu_m$ the maps

$$\mathcal{U}_i \cdot (\phi_{w_i})_{\varepsilon_1 < \dots < \varepsilon_n}^{\nu_1 < \dots < \nu_m} + \mathcal{V}_i \cdot (\psi_{z_i})_{\varepsilon_1 < \dots < \varepsilon_n}^{\nu_1 < \dots < \nu_m} \quad (13.25)$$

(which count holomorphic curves weighted by their multiplicities over w_i and z_i , respectively, as in equation (13.5)) and also summing expressions of the form

$$\mathbf{x} \mapsto f_{\alpha_{\varepsilon_n}, \dots, \alpha_{\nu_1}, \beta_{\nu_1}, \dots, \beta_{\nu_m}}(\Theta_{n, n-1}^\alpha, \dots, (\Theta_{j+1, j}^\alpha)', \dots, \Theta_{2, 1}^\alpha, \mathbf{x}, \Theta_{1, 2}^\beta, \dots, \Theta_{m-1, m}^\beta), \quad (13.26)$$

as well as terms where the primed term is a beta-chain. In the above, we write $\Theta_{j+1, j}^\alpha$ for $\Theta_{\alpha_{\varepsilon_{j+1}}, \alpha_{\varepsilon_j}}$ and $(\Theta_{j+1, j}^\alpha)'$ denotes $(\mathcal{U}_i \partial_{\mathcal{U}_i} + \mathcal{V}_i \partial_{\mathcal{V}_i})(\Theta_{j+1, j}^\alpha)$, and similarly for the beta terms. That is, we sum over the holomorphic polygon maps where we apply $(\mathcal{U}_i \partial_{\mathcal{U}_i} + \mathcal{V}_i \partial_{\mathcal{V}_i})$ to exactly one of the inputs from \mathcal{L}_α or \mathcal{L}_β . The sum of maps in equation (13.25) is exactly the map $\mathcal{U}_i \Phi_{w_i}^0 + \mathcal{V}_i \Psi_{z_i}^0$, so it suffices to show that the terms in equation (13.26) vanish. To see this, we will show that if Θ is any chain in \mathcal{L}_α or \mathcal{L}_β , then $(\mathcal{U}_i \partial_{\mathcal{U}_i} + \mathcal{V}_i \partial_{\mathcal{V}_i})(\Theta) = 0$. We argue as follows, focusing on \mathcal{L}_α . We observe that on the diagram $(\Sigma, \mathcal{L}_\alpha, \mathbf{w}, \mathbf{z})$, the base points w_i and z_i are immediately adjacent (by virtue of the existence of λ_β , which is disjoint from the alpha curves). We note that each subdiagram $(\Sigma, \alpha', \alpha, \mathbf{w}, \mathbf{z})$ represents an unlink \mathbb{U} in a connected sum of several copies of $S^1 \times S^2$. There is a $|\mathbb{U}|$ -component Alexander grading on $\mathcal{CFL}(\mathbb{U})$, and the top Maslov bidegree generators lie in the Alexander grading

$(0, \dots, 0)$. By construction, all of the chains of \mathcal{L}_α also lie in the Alexander grading $(0, \dots, 0)$. Write A_i for the component of the Alexander grading corresponding to K_i . Since λ_β connects w_i to z_i and is disjoint from all curves of \mathcal{L}_α , if $\mathbf{x} \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_\alpha$, we have $A_i(\mathbf{x}) = 0$. In particular, any chain in \mathcal{L}_α must have equal powers of \mathcal{U}_i and \mathcal{V}_i , i.e., it must be of the form $\mathcal{U}_i^N \mathcal{V}_i^N \cdot \mathbf{y}$, where \mathbf{y} has no powers of \mathcal{U}_i or \mathcal{V}_i . Clearly,

$$(\mathcal{U}_i \partial_{\mathcal{U}_i} + \mathcal{V}_i \partial_{\mathcal{V}_i})(\mathcal{U}_i^N \mathcal{V}_i^N \cdot \mathbf{y}) = 0,$$

as claimed. The same argument works for \mathcal{L}_β , so the proof is complete. ■

13.8 Properties of the U -action on link surgery complexes

We now explore some of the complexities of the surgery theorem when we allow arbitrary arcs for our σ -basic systems.

Lemma 13.31. *Suppose that \mathcal{A} is a system of arcs for $L \subseteq S^3$, and suppose that the arcs for knot components $K_1, K_2 \subseteq L$ are alpha-parallel or beta-parallel (the case that one is alpha-parallel and the other is beta-parallel is allowed). Then the actions of U_1 and U_2 are homotopic on $\mathcal{X}_\Lambda(L, \mathcal{A})^\mathcal{L}$ in the sense that there is a type-D module morphism $H^1: \mathcal{X}_\Lambda(L, \mathcal{A})^\mathcal{L} \rightarrow \mathcal{X}_\Lambda(L, \mathcal{A})^\mathcal{L}$ such that*

$$\partial_{\text{Mor}}(H^1) = \text{id} \otimes (U_1 + U_2).$$

In the above, U_i denotes $\mathcal{U}_i \mathcal{V}_i$ and $\text{id} \otimes (U_1 + U_2)$ means the map which sends \mathbf{x} to $\mathbf{x} \otimes (U_1 + U_2)$ in all idempotents.

Proof. Let \mathcal{H} be a σ -basic system of Heegaard diagrams for (L, \mathcal{A}) . We may assume that \mathcal{H} is meridional for K_1 and K_2 , in the sense of Section 9.5. We begin with the case that K_1 and K_2 are both alpha-parallel or both beta-parallel. For concreteness, assume that K_1 and K_2 are both alpha-parallel, so that the arcs of \mathcal{A} are both disjoint from the alpha curves. We let $\lambda: [0, 1] \rightarrow \Sigma$ denote a path from z_1 to z_2 which is disjoint from all of the arcs of \mathcal{A} .

Given hypercubes of alpha and beta Lagrangians \mathcal{L}_α and \mathcal{L}_β on a pointed surface $(\Sigma, \mathbf{w}, \mathbf{z})$, as well as an arc λ connecting two points $x_1, x_2 \in \Sigma$, we obtain a map

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

as in Section 13.3, satisfying

$$\partial_{\text{Mor}}(\mathcal{A}_\lambda) = E_1^\alpha + E_2^\alpha,$$

where E_i^α denotes the product of the variables for the base points in the unique alpha degeneration containing x_i . Compare Lemma 13.14. We view the link surgery hypercube as a 2-dimensional hypercube by internalizing all axis directions except those of

to the other length 1 maps $\mathcal{A}_{\lambda_\varepsilon}$ so we conclude that the length 1 maps satisfy

$$\mathcal{A}_{\lambda_\varepsilon} = \mathcal{A}_{\lambda_{0,0}} \tag{13.28}$$

for all $\varepsilon \in \mathbb{E}_2$.

We now consider the length 2 maps, we claim that

$$h^{-K_i} = j^{-K_i} \quad \text{for } i = \{1, 2\}. \tag{13.29}$$

The argument is similar to the argument for the length 1 maps. Namely, the maps h^{-K_2} and j^{-K_2} count the same holomorphic curves. Some of the curves contributing to h^{-K_2} are counted with a weight of $a(\psi, \lambda)$ whereas the same curves are counted by $a(\psi, \phi_1(\lambda))$ for j^{-K_2} . Since $\phi_1(\lambda)$ is obtained by concatenating λ with an arc which passes through only beta curves, we see that $a(\psi, \lambda) = a(\psi, \phi_1(\lambda))$, so these expressions are equal. We conclude that $h^{-K_i} = j^{-K_i}$.

We now define an endomorphism $\mathcal{A}_\lambda^{\mathbb{X}}$ of the link surgery hypercube $\mathcal{C}_\Lambda(L)$ to be the endomorphism induced by the components of the diagram in (13.27) which come out of the page (i.e., the $\mathcal{A}_{\lambda_\varepsilon}$ -direction). By construction, the maps labeled h^{-K_i} have the same equivariance properties as the link surgery maps F^{-K_i} , so the morphism $\mathcal{A}_\lambda^{\mathbb{X}}$ is induced by a morphism of type-D modules

$$\mathcal{A}_\lambda^{\mathbb{X},1}: \mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}} \rightarrow \mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}}.$$

(Note that outside of this lemma, we will write just \mathcal{A}_λ or \mathcal{A}_λ^1 for the maps labeled $\mathcal{A}_\lambda^{\mathbb{X}}$ and $\mathcal{A}_\lambda^{\mathbb{X},1}$ above.)

The components labeled $\mathcal{A}_{\lambda_\varepsilon}$ preserve the idempotents for K_1 and K_2 . The arrows labeled $-K_1$ are weighted in H^1 by multiples of the algebra element τ_1 . The arrows labeled $-K_2$ are weighted by multiples of τ_2 . The dotted arrow labeled ω^{-K_1, K_2} is weighted by a multiple of $\tau_1 \tau_2$. None of the components of $\mathcal{A}_\lambda^{\mathbb{X},1}$ are weighted by algebra elements containing a factor of σ_1 or σ_2 .

As a morphism from the back face of the diagram in (13.27) to the front face, the map $\mathcal{A}_\lambda^{\mathbb{X}}$ satisfies

$$\partial_{\text{Mor}}(\mathcal{A}_\lambda^{\mathbb{X}}) = (U_1 + U_2) \cdot \text{id}$$

by Remark 13.15. We claim that the same equation is satisfied if we view $\mathcal{A}_\lambda^{\mathbb{X}}$ as a morphism from $\mathcal{C}_\Lambda(L)$ to itself, and furthermore, that the type-D endomorphism $\mathcal{A}_\lambda^{\mathbb{X},1}$ satisfies $\partial_{\text{Mor}}(\mathcal{A}_\lambda^{\mathbb{X},1}) = \text{id} \otimes (U_1 + U_2)$ as an endomorphism of $\mathcal{X}_\Lambda(L)^{\mathcal{L}}$. Note that the hypercube differential in the diagram in (13.27) does not include the differentials F^{K_1} or F^{K_2} . Since we are using σ -basic systems, these maps are identified with the maps induced by localization at \mathcal{V}_1 and \mathcal{V}_2 . Once we include F^{K_1} and F^{K_2} into the differentials on the front and back faces, the fact that $\partial_{\text{Mor}}(\mathcal{A}_\lambda^{\mathbb{X}})$ is still $(U_1 + U_2) \cdot \text{id}$ follows immediately from equations (13.28) and (13.29). The same argument works for type-D morphisms. This completes the proof when both components are alpha-parallel.

We now consider the case that K_1 is alpha-parallel and K_2 is beta-parallel. In this case, we must modify the map $\mathcal{A}_\lambda^{\mathcal{X},1}$. Recall that $\mathcal{A}_\lambda^{\mathcal{X},1}$ was built from the map \mathcal{A}_λ . We instead replace this by a map $\mathcal{A}_\lambda^{\mathcal{X},1}$ built from $\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}$.

If we view the complex as a 1-dimensional cone in the K_2 -direction

$$\mathcal{C}_\Lambda(L, \mathcal{A}) = \text{Cone}(\mathcal{C}_0 \xrightarrow{F^{K_2} + F^{-K_2}} \mathcal{C}_1),$$

then we can define endomorphisms $\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}: \mathcal{C}_i \rightarrow \mathcal{C}_i$ similar to the case when K_1 and K_2 were both alpha-parallel. The map $\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}: \mathcal{C}_0 \rightarrow \mathcal{C}_0$ consists of the components labeled $\mathcal{A}_{\lambda_{0,0}}$, $\mathcal{A}_{\lambda_{1,0}}$ and h^{-K_1} . As an endomorphism of \mathcal{C}_1 , the map is similar. Since K_1 is alpha-parallel, we see that

$$\mathcal{A}_{\lambda_{0,0}} = \mathcal{A}_{\lambda_{1,0}} = \mathcal{A}_\lambda,$$

and so the same argument as before shows that

$$\partial_{\text{Mor}}(\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}) = (U_1 + U_2) \cdot \text{id}$$

as endomorphisms of \mathcal{C}_i .

On the other hand, the point pushing map for K_2 moves z_2 to w_2 through exactly the alpha circles which intersect the trace of K_2 on the knot diagram. Hence

$$\mathcal{A}_{\lambda_{0,1}} = \mathcal{A}_{\lambda_{1,1}} = \mathcal{A}_\lambda + \mathcal{A}_{[K_2]}.$$

Here, $\mathcal{A}_{[K_2]}$ denotes the hypercube homology action for the trace of K_2 on the Heegaard diagram. We conclude that

$$F^{-K_2}\mathcal{A}_\lambda \simeq (\mathcal{A}_\lambda + \mathcal{A}_{[K_2]})F^{-K_2}, \quad (13.30)$$

as maps from \mathcal{C}_0 to \mathcal{C}_1 . On the other hand, since F^{-K_2} moves z_2 to w_2 , we have

$$F^{-K_2}\mathcal{V}_2\Psi_{z_2} \simeq \mathcal{U}_2\Phi_{w_2}F^{-K_2}. \quad (13.31)$$

We recall from Lemma 13.30 that $\mathcal{A}_{[K_2]} \simeq \mathcal{U}_2\Phi_{w_2} + \mathcal{V}_2\Psi_{z_2}$, as endomorphisms of \mathcal{C}_0 . Combining the relations in equations (13.30), (13.31) with this equation for $\mathcal{A}_{[K_2]}$, we obtain

$$F^{-K_2}(\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}) \simeq (\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2})F^{-K_2},$$

as maps from \mathcal{C}_0 to \mathcal{C}_1 . Note that since F^{K_2} is the map for localizing at \mathcal{V}_2 , we also have

$$F^{-K_2}(\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}) = (\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2})F^{-K_2}.$$

Arranging the maps $\mathcal{A}_\lambda + \mathcal{V}_2\Psi_{z_2}$ (as maps from \mathcal{C}_i to \mathcal{C}_i) as well as the above homotopies (which map \mathcal{C}_0 to \mathcal{C}_1) into an endomorphism H of $\mathcal{C}_\Lambda(L)$, we see that $\partial_{\text{Mor}}(H) = (U_1 + U_2) \cdot \text{id}$. The same logic gives a type-D homotopy, so the proof is complete. \blacksquare

Chapter 14

The alpha-beta transformer

In this chapter, we describe a bimodule ${}_{\mathcal{K}}\mathcal{T}^{\mathcal{K}}$, which we call the *alpha-beta transformer bimodule*. We will prove the following.

Theorem 14.1. *Suppose that $L \subseteq S^3$ is a framed link with a distinguished component K . Let \mathcal{A} be a system of arcs for L such that the arc for K is alpha-parallel. Let \mathcal{A}' be a system of arcs where K is beta-parallel, but otherwise the arcs are the same as \mathcal{A} . Then*

$$\mathcal{X}_{\Lambda}(L, \mathcal{A})^{\mathcal{L}} \simeq \mathcal{X}_{\Lambda}(L, \mathcal{A}')^{\mathcal{L}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{T}^{\mathcal{K}}.$$

The same formula holds if we switch the roles of \mathcal{A} and \mathcal{A}' .

14.1 The bimodule

The bimodule \mathcal{T} has rank 1 over \mathbf{I} , and has $\delta_1^1 = 0$. The map δ_2^1 is the unique \mathbf{I} -linear map which satisfies

$$\delta_2^1(a, 1) = 1 \otimes a,$$

for $a \in \mathcal{K}$. (This coincides with the identity bimodule.)


Additionally, there is a δ_4^1 term, which is non-vanishing on tuples of the form $(a, b, f\tau, i_0)$ where f, a and b are concentrated in a single idempotent, and which vanishes on other elementary tensors. On such elements, we have the formula

$$\delta_4^1(a, b, f\tau, i_0) = i_1 \otimes (\partial_U(a) \cdot \partial_V(b) \cdot f\tau).$$

Lemma 14.2. *The structure maps on ${}_{\mathcal{K}}\mathcal{T}^{\mathcal{K}}$ satisfy the DA bimodule relations.*

Proof. The only interesting structure relations to be verified are the ones with 5 inputs. There are only two non-vacuous relations to be verified, which are those for sequences of the form (a, b, c, τ, i_0) and those of the form (a, b, τ, c, i_0) .

The structure relations for sequences of the form (a, b, τ, c, i_0) are straightforward, so we consider the tuple (a, b, c, τ, i_0) . The structure diagrams with non-trivial contribution are the following:



The vanishing of the above equation follows from the algebraic relation

$$\partial_U(ab)\partial_V(c) + \partial_U(a)\partial_V(bc) + \partial_U(a)\partial_V(b)c + a\partial_U(b)\partial_V(c) = 0,$$

which is straightforward to verify. ■

In Section 14.3, we describe another perspective which makes the DA bimodule relations more straightforward to verify.

14.2 Basepoint actions as A_∞ -morphisms

We now give a helpful reformulation of algebraic base point actions in terms of A_∞ morphisms. This framework helps us to understand the structure maps appearing in the alpha-to-beta transformer bimodule.

If C is a free $\mathbb{F}[U]$ -module, then there is natural base point action $\Phi_U: C \rightarrow C$, defined as follows. We pick a free $\mathbb{F}[U]$ basis B of C . If $\mathbf{x} \in B$, then we write $\partial(\mathbf{x}) = \sum_{\mathbf{y} \in B} c_{\mathbf{y},\mathbf{x}} U^{n_{\mathbf{y},\mathbf{x}}} \cdot \mathbf{y}$, where $c_{\mathbf{y},\mathbf{x}} \in \mathbb{F}$ and $n_{\mathbf{y},\mathbf{x}} \in \mathbb{N}$. The map Φ_U is defined via the formula

$$\Phi_U(\mathbf{x}) = \sum_{\mathbf{y} \in B} n_{\mathbf{y},\mathbf{x}} c_{\mathbf{y},\mathbf{x}} U^{n_{\mathbf{y},\mathbf{x}}-1} \cdot \mathbf{y}. \tag{14.1}$$

The induced map $(\Phi_U)_*$ on $H_*(C)$ is trivial (cf. Remark 13.11), though the map Φ_U is often not null-homotopic through an $\mathbb{F}[U]$ -equivariant homotopy.

Since C is free and the differential is $\mathbb{F}[U]$ -equivariant, we may view C as corresponding to a type-D module $C^{\mathbb{F}[U]}$. Of course,

$$C^{\mathbb{F}[U]} \cong C^{\mathbb{F}[U]} \boxtimes_{\mathbb{F}[U][\mathbb{I}]} \mathbb{F}[U].$$

There is a DA bimodule endomorphism ϕ_*^1 of $\mathbb{F}[U][\mathbb{I}]^{\mathbb{F}[U]}$ defined as follows. We set $\phi_1^1 = 0$, and we set

$$\phi_2^1(\alpha, 1) = 1 \otimes \partial_U(\alpha).$$

We set $\phi_j^1 = 0$ if $j > 2$. The DA bimodule relations are straightforward to verify, and easily translated into the Leibniz rule for derivatives.

Lemma 14.3. *The map Φ_U in equation (14.1), viewed as a type-D morphism*

$$\Phi_U: C^{\mathbb{F}[U]} \rightarrow C^{\mathbb{F}[U]},$$

satisfies

$$\Phi_U := \mathbb{I}_C \boxtimes \phi_*^1.$$

The proof is immediate.

We recall that the map Φ_U satisfies $\Phi_U^2 \simeq 0$. See Lemma 13.23. It is helpful to note that this relation may be translated into a fact about the DA bimodule endomorphism ϕ_*^1 .

Lemma 14.4. *The DA bimodule morphism ϕ_*^1 satisfies*

$$\phi_*^1 \circ \phi_*^1 \simeq 0.$$

Proof. We define a DA bimodule endomorphism

$$H_*^1: \mathbb{F}[U][\mathbb{I}]^{\mathbb{F}[U]} \rightarrow \mathbb{F}[U][\mathbb{I}]^{\mathbb{F}[U]}$$

via the formula $H_1^1 = 0$ and

$$H_2^1(U^n, 1) = 1 \otimes \frac{n(n-1)}{2} U^{n-2}.$$

The map $\partial_{\text{Mor}}(H_*^1)$ has three summands, shown as graphs below:

$$\partial_{\text{Mor}}(H_*^1) = \begin{array}{c} \swarrow \downarrow \\ \mu_2 \\ \downarrow \\ \swarrow \downarrow \\ H_2^1 \\ \downarrow \searrow \\ \mu_2 \end{array} + \begin{array}{c} \swarrow \downarrow \\ H_2^1 \\ \swarrow \downarrow \\ \delta_2^1 \\ \downarrow \searrow \\ \mu_2 \end{array} + \begin{array}{c} \swarrow \downarrow \\ \delta_2^1 \\ \swarrow \downarrow \\ H_2^1 \\ \downarrow \searrow \\ \mu_2 \end{array}$$

Hence,

$$\begin{aligned} \partial_{\text{Mor}}(H_*^1)(U^n, U^m, 1) &= 1 \otimes \frac{n(n-1) + m(m-1) + (n+m)(n+m-1)}{2} U^{n+m-2} \\ &\equiv 1 \otimes nm U^{n+m-2} \pmod{2} \\ &= (\phi_* \circ \phi_*)_2^1(U^n, U^m, 1), \end{aligned}$$

completing the proof. ■

We will also need to consider analogous base point actions on complexes over the 2-variable polynomial ring $\mathbb{F}[\mathcal{U}, \mathcal{V}]$. In this case, we may define two DA bimodule endomorphisms of $\mathbb{F}[\mathcal{U}, \mathcal{V}][\mathbb{I}]^{\mathbb{F}[\mathcal{U}, \mathcal{V}]}$, denoted ϕ_*^1 and ψ_*^1 , which differentiate with respect to \mathcal{U} and \mathcal{V} , respectively. Lemma 14.4 adapts to show that $\phi_*^1 \circ \phi_*^1 \simeq 0$ and $\psi_*^1 \circ \psi_*^1 \simeq 0$. Additionally, we have the following.

Lemma 14.5. *As endomorphisms of $\mathbb{F}[\mathcal{U}, \mathcal{V}][\mathbb{I}]^{\mathbb{F}[\mathcal{U}, \mathcal{V}]}$, we have*

$$\phi_*^1 \circ \psi_*^1 \simeq \psi_*^1 \circ \phi_*^1.$$

Proof. We define a DA bimodule endomorphism H_*^1 of $\mathbb{F}[\mathcal{U}, \mathcal{V}][\mathbb{I}]^{\mathbb{F}[\mathcal{U}, \mathcal{V}]}$ whose only non-vanishing action is given by the formula

$$H_2(\alpha, 1) = 1 \otimes \partial_{\mathcal{U}, \mathcal{V}}^2(\alpha),$$

where $\partial_{\mathcal{U}, \mathcal{V}}^2(a) = \partial_{\mathcal{U}}(\partial_{\mathcal{V}}(a))$. We observe the equality

$$\partial_{\text{Mor}}(H_*^1) = \phi_*^1 \circ \psi_*^1 + \psi_*^1 \circ \phi_*^1$$

which follows from the easily verified relation

$$\partial_{\mathcal{U}, \mathcal{V}}^2(a \cdot b) + a \cdot \partial_{\mathcal{U}, \mathcal{V}}^2(b) + \partial_{\mathcal{U}, \mathcal{V}}^2(a) \cdot b = \partial_{\mathcal{U}}(a) \cdot \partial_{\mathcal{V}}(b) + \partial_{\mathcal{V}}(a) \cdot \partial_{\mathcal{U}}(b). \quad \blacksquare$$

14.3 The transformer as a mapping cone bimodule

In order to efficiently prove statements about the transformer bimodule, it is helpful to discuss some generalities about constructing bimodules. If $f_*^1: \mathcal{A}M^{\mathcal{B}} \rightarrow \mathcal{A}N^{\mathcal{B}}$ is a morphism of DA bimodules such that $\partial_{\text{Mor}}(f_*^1) = 0$, then we may construct another bimodule $\text{Cone}(f_*^1)$ which is also a DA bimodule. If f_*^1 and g_*^1 are two homotopic bimodule morphisms, then $\text{Cone}(f_*^1)$ and $\text{Cone}(g_*^1)$ are homotopy equivalent.

It is helpful to understand that the identity and transformer bimodules are mapping cone bimodules. For $\varepsilon \in \{0, 1\}$, write $\mathcal{X}[\mathbf{i}_\varepsilon]^{\mathcal{K}}$, denote $\mathbf{I}_\varepsilon \cdot \mathcal{X}[\mathbb{I}]^{\mathcal{K}} \cdot \mathbf{I}_\varepsilon$, with the restriction of the structure map from the identity bimodule $\mathcal{X}[\mathbb{I}]^{\mathcal{K}}$. There are two morphisms

$$s_*^1, t_*^1: \mathcal{X}[\mathbf{i}_0]^{\mathcal{K}} \rightarrow \mathcal{X}[\mathbf{i}_1]^{\mathcal{K}}$$

as follows. The map $t_j^1 = 0$ for $j \neq 2$, and $t_2^1(f\tau, i_0) = i_1 \otimes f\tau$ and $t_2^1(f\sigma, i_0) = 0$. The map s_*^1 is similar, but with the roles of σ and τ reversed. It is not hard to see that the DA bimodule morphisms s_*^1 and t_*^1 are cycles. Furthermore,

$$\mathcal{X}[\mathbb{I}]^{\mathcal{K}} = \text{Cone}(s_*^1 + t_*^1: \mathcal{X}[\mathbf{i}_0]^{\mathcal{K}} \rightarrow \mathcal{X}[\mathbf{i}_1]^{\mathcal{K}}).$$

The A_∞ base point actions ϕ_*^1 and ψ_*^1 from Section 14.2 also induce endomorphisms of $\mathcal{X}[\mathbf{i}_\varepsilon]^{\mathcal{K}}$ for $\varepsilon \in \{0, 1\}$. In particular, we observe that the transformer bimodule may be described as

$$\mathcal{X}\mathcal{T}^{\mathcal{K}} = \text{Cone}(s_*^1 + (\text{id} + \phi_*^1 \circ \psi_*^1) \circ t_*^1: \mathcal{X}[\mathbf{i}_0]^{\mathcal{K}} \rightarrow \mathcal{X}[\mathbf{i}_1]^{\mathcal{K}}).$$

Remark 14.6. The above perspective gives a more streamlined proof that $\mathcal{X}\mathcal{T}^{\mathcal{K}}$ satisfies the DA bimodule structure relations. Indeed, it is sufficient to verify that ϕ_*^1 , ψ_*^1 , s_*^1 and t_*^1 are cycles, which is easy to verify.

14.4 Properties of the transformer bimodule

In this section, we prove some important properties of the transformer bimodule.

We begin by showing that in analogy to Sarkar’s map for a Dehn twist [49], the transformer bimodule is involutive.

Lemma 14.7. *The transformer bimodule satisfies*

$$\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{T}^{\mathcal{K}} \simeq \mathcal{K}[\mathbb{I}]^{\mathcal{K}}.$$

Proof. We compute that $\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{T}^{\mathcal{K}}$ is equal to the mapping cone bimodule

$$\text{Cone}(s_*^1 + (\text{id} + \phi_*^1 \circ \psi_*^1 \circ \phi_*^1 \circ \psi_*^1) \circ t_*^1: \mathcal{K}[\mathbf{i}_0]^{\mathcal{K}} \rightarrow \mathcal{K}[\mathbf{i}_1]^{\mathcal{K}}).$$

On the other hand, $\phi_*^1 \circ \psi_*^1 \circ \phi_*^1 \circ \psi_*^1 \simeq 0$ by Lemmas 14.4 and 14.5. Hence, $\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{T}^{\mathcal{K}}$ is homotopy equivalent to $\text{Cone}(s_*^1 + t_*^1)$, which is $\mathcal{K}[\mathbb{I}]^{\mathcal{K}}$. ■

Next, we consider the tensor product of the transformer with the modules for a solid torus.

Lemma 14.8. *The transformer bimodule satisfies*

$$\mathcal{D}_\lambda^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{T}^{\mathcal{K}} \simeq \mathcal{D}_\lambda^{\mathcal{K}} \quad \text{and} \quad \mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{D}_\lambda^{\mathbb{F}[U]} \simeq \mathcal{K}\mathcal{D}_\lambda^{\mathbb{F}[U]}.$$

Proof. The relation $\mathcal{D}_\lambda^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{T}^{\mathcal{K}} \simeq \mathcal{D}_\lambda^{\mathcal{K}}$ is trivial, so we focus on the relation $\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{D}_\lambda^{\mathbb{F}[U]} \simeq \mathcal{K}\mathcal{D}_\lambda^{\mathbb{F}[U]}$. To simplify the notation, it is easier to construct a homotopy equivalence

$$\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]} \simeq \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]},$$

though the maps we write down can be easily adapted to give the statement about DA bimodules, since they are strictly $\mathbb{F}[U]$ -equivariant.

We will view both bimodules as mapping cone bimodules, in the sense of Section 14.3. Define

$$\mathcal{K}[\mathbf{d}_\varepsilon]_{\mathbb{F}[U]} := \mathbf{I}_\varepsilon \cdot \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]},$$

for $\varepsilon \in \{0, 1\}$. We observe that both $\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}$ and $\mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]}$ are the mapping cones of type-AA module maps from \mathbf{d}_0 to \mathbf{d}_1 . Indeed, there are type-AA module maps

$$\Sigma_*, T_*: \mathcal{K}[\mathbf{d}_0]_{\mathbb{F}[U]} \rightarrow \mathcal{K}[\mathbf{d}_1]_{\mathbb{F}[U]},$$

where $\Sigma_{1,1,0}(f\sigma, \mathbf{x}) = m_2(f\sigma, \mathbf{x})$ (where m_2 is the action from \mathcal{D}_λ) and $T_{1,1,0}(f\tau, \mathbf{x}) = m_2(f\tau, \mathbf{x})$. Furthermore, $\Sigma_{1,1,0}(f\tau, \mathbf{x}) = 0$ and $T_{1,1,0}(f\sigma, \mathbf{x}) = 0$. Also Σ_* and T_* vanish on other configurations of inputs. It is straightforward to verify that Σ_* and T_* are cycles. Additionally,

$$\mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]} = \text{Cone}(\Sigma_* + T_*: \mathcal{K}[\mathbf{d}_0]_{\mathbb{F}[U]} \rightarrow \mathcal{K}[\mathbf{d}_1]_{\mathbb{F}[U]}) \quad (14.2)$$

and

$$\mathcal{K}\mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}[\mathcal{D}_\lambda]_{\mathbb{F}[U]} = \text{Cone}(\Sigma_* + (\text{id} + \phi_* \circ \psi_*) \circ T_*: \mathcal{K}[\mathbf{d}_0]_{\mathbb{F}[U]} \rightarrow \mathcal{K}[\mathbf{d}_1]_{\mathbb{F}[U]}). \quad (14.3)$$

Here, we are writing ϕ_* and ψ_* for $\phi_*^1 \boxtimes \mathbb{I}_{\mathcal{D}_\lambda}$ and $\psi_*^1 \boxtimes \mathbb{I}_{\mathcal{D}_\lambda}$, where ϕ_*^1 and ψ_*^1 are the endomorphisms of $\mathcal{X}[\mathbf{i}_1]^{\mathcal{K}}$ described above.

We observe that

$$\phi_* \circ \psi_* \simeq \mathcal{V}^{-1} \phi_* \circ (\mathcal{U}\phi_* + \mathcal{V}\psi_*), \tag{14.4}$$

as endomorphisms of $[\mathbf{d}_1]$, since $\phi_*^2 \simeq 0$ by Lemma 14.4.

We write $\partial_{\mathcal{W}} = \mathcal{U}\partial_{\mathcal{U}} + \mathcal{V}\partial_{\mathcal{V}}$. We define an endomorphism J_* of $\mathcal{X}[\mathbf{d}_1]_{\mathbb{F}[U]}$ via the formula

$$J_{0,1,0}(\mathbf{x}) = \partial_{\mathcal{W}}(\mathbf{x}).$$

We set $J_{i,1,j} = 0$ unless $i = j = 0$. Observe that $J_{0,1,0}$ commutes with the right action of U . We claim that

$$\partial_{\text{Mor}}(J_*) = \mathcal{U}\phi_* + \mathcal{V}\psi_*.$$

To see that, we note that there are exactly two structure trees which contribute to $\partial_{\text{Mor}}(J_*)$. These are shown below:

$$\partial_{\text{Mor}}(J_*)_{1,1,0}(a, \mathbf{x}) = \begin{array}{ccc} a & \searrow & \mathbf{x} \\ & & \downarrow \\ & & \delta_2^1 \\ & & \downarrow \\ & & J_{0,1,0} \\ & & \downarrow \end{array} + \begin{array}{ccc} a & \searrow & \mathbf{x} \\ & & \downarrow \\ & & J_{0,1,0} \\ & & \downarrow \\ & & \delta_2^1 \\ & & \downarrow \end{array}$$

Hence,

$$\partial_{\text{Mor}}(J_*)_{1,1,0}(a, \mathbf{x}) = \partial_{\mathcal{W}}(a \cdot \mathbf{x}) + a \cdot \partial_{\mathcal{W}}(\mathbf{x}).$$

By the Leibniz rule we have that the above is $\partial_{\mathcal{W}}(a) \cdot \mathbf{x}$. Therefore,

$$\partial_{\text{Mor}}(J_*) = \mathcal{U}\phi_* + \mathcal{V}\psi_*,$$

as endomorphisms of $[\mathbf{d}_1]$. We conclude that $\phi_* \circ \psi_* \simeq 0$ as endomorphisms of $[\mathbf{d}_1]$ by equation (14.4). Hence the two mapping cones in equations (14.2) and (14.3) are homotopy equivalent. The proof is complete. ■

14.5 Proof of Theorem 14.1

We now prove Theorem 14.1, which states that changing from an alpha-parallel arc to a beta-parallel arc may be achieved by tensoring with the transformer bimodule $\mathcal{X}\mathcal{T}^{\mathcal{K}}$.

Proof of Theorem 14.1. We use our base point moving map formula from Corollary 13.5.

Let \mathcal{A} be an arc system for $L \subseteq S^3$, and assume that $K_i \subseteq L$ is a component whose arc \mathcal{A}_i is beta-parallel. Let \mathcal{A}'_i denote an arc for K_i which is alpha-parallel, and let \mathcal{A}' denote the arc system which has \mathcal{A}'_i in place of \mathcal{A}_i .

We pick a Heegaard link diagram $\mathcal{H} = (\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$ for (S^3, L) so that all of the arcs of \mathcal{A} are embedded in Σ and pairwise disjoint. We assume the arc \mathcal{A}_i is disjoint from β . The arc \mathcal{A}'_i may also be viewed as being embedded on Σ . With respect to this embedding, \mathcal{A}'_i is disjoint from α . In principle, \mathcal{A}'_i might intersect some of the other arcs of \mathcal{A} . However, by stabilizing the Heegaard diagram $(\Sigma, \alpha, \beta, \mathbf{w}, \mathbf{z})$, we may assume that \mathcal{A}'_i is disjoint from all of the other arcs of \mathcal{A} .

We build a hyperbox of Heegaard diagrams using the diagram \mathcal{H} as the initial diagram. If $K_i \subseteq L$, then the K_i -direction of this hypercube involves the composition of a tautological diffeomorphism map to move z_i along \mathcal{A}_i , followed by a holomorphic polygon counting maps which change the alpha and beta curves back to their original position. We do this in each axis direction to build a σ -basic system used for the link surgery formula.

For this σ -basic system, we may change \mathcal{A}_i to \mathcal{A}'_i by composing the K_i -direction with a final base point moving hypercube which moves w_i in a loop along the curve $[K_i] = \mathcal{A}'_i * \mathcal{A}_i^{-1}$. Corollary 13.5 implies that the effect on the link surgery hypercube is to replace F^{-K_i} (the sum of the hypercube maps ranging over all sublinks of L containing $-K_i$) with the map $(\text{id} + \mathcal{V}_i^{-1} \mathcal{A}_{[K_i]} \circ \Phi_{w_i}) \circ F^{-K_i}$.

The map $(\text{id} + \mathcal{V}_i^{-1} \mathcal{A}_{[K_i]} \circ \Phi_{w_i})$ is viewed as an endomorphism of the $(\ell - 1)$ -dimensional subcube of $\mathcal{C}_\Lambda^1(L, \mathcal{A}) \subseteq \mathcal{C}_\Lambda(L, \mathcal{A})$ which has K_i -component 1. On this subcube, the arc \mathcal{A}'_i intersects only beta curves, and \mathcal{A}_i intersects only alpha curves. By Lemma 13.30, we have

$$\mathcal{A}_{[K_i]} \simeq \mathcal{U}_i \Phi_{w_i} + \mathcal{V}_i \Psi_{z_i},$$

as endomorphisms of $\mathcal{C}_\Lambda^1(L, \mathcal{A})$. Hence

$$\text{id} + \mathcal{V}_i^{-1} \Phi_{w_i} \mathcal{A}_{[K_i]} = \text{id} + \mathcal{V}_i^{-1} \Phi_{w_i} (\mathcal{U}_i \Phi_{w_i} + \mathcal{V}_i \Psi_{z_i}) \simeq \text{id} + \Phi_{w_i} \Psi_{z_i},$$

since $\Phi_{w_i}^2 \simeq 0$. In particular, it follows that if we write

$$\mathcal{C}_\Lambda(L, \mathcal{A}) \cong \text{Cone}(\mathcal{C}_\Lambda^0(L, \mathcal{A}) \xrightarrow{F^{K_i} + F^{-K_i}} \mathcal{C}_\Lambda^1(L, \mathcal{A})),$$

then $\mathcal{C}_\Lambda(L, \mathcal{A}')$ is homotopic to

$$\text{Cone}(\mathcal{C}_\Lambda^0(L, \mathcal{A}) \xrightarrow{F^{K_i} + (\text{id} + \Phi_{w_i} \Psi_{z_i}) F^{-K_i}} \mathcal{C}_\Lambda^1(L, \mathcal{A})).$$

On the level of type-D structures over \mathcal{L} , we claim that the above complex corresponds to the tensor product of $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}}$ with the transformer bimodule $\mathcal{X} \mathcal{T}^{\mathcal{X}}$. To

see this, observe that the components of the box tensor differential of $\mathcal{X}_\Lambda(L, \mathcal{A})^{\mathcal{L}} \hat{\boxtimes}$ $\mathcal{X}\mathcal{T}^{\mathcal{K}}$ which involve the map δ_2^1 of $\mathcal{X}\mathcal{T}^{\mathcal{K}}$ contribute the internal differentials of $\mathcal{C}_\Lambda^\varepsilon(L, \mathcal{A})$ (for $\varepsilon \in \{0, 1\}$), as well as the maps F^{K_i} and F^{-K_i} to the mapping cone differential. The terms of the box tensor product differential which involve δ_4^1 of $\mathcal{X}\mathcal{T}^{\mathcal{K}}$ contribute the term $\Phi_{w_i} \Psi_{z_i} F^{-K_i}$ to the above mapping cone differential (compare Lemma 14.3). This completes the proof. ■

Chapter 15

The pair-of-pants bimodules

In this chapter, we describe several bimodules $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}}$ and $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}$. The bimodules are related by the transformer bimodule

$$\mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}\mathcal{T}^{\mathcal{K}} \simeq \mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}.$$

We call these the *pair-of-pants* bimodules. Other choices of α and β in the subscripts can also be constructed by tensoring in copies $\mathcal{K}\mathcal{T}^{\mathcal{K}}$ to switch from α to β and vice-versa.

After defining these bimodules and proving several basic properties, we will prove the following.

Theorem 15.1. *Suppose that $L_1, L_2 \subseteq S^3$ are framed links with distinguished components $K_1 \subseteq L_1$ and $K_2 \subseteq L_2$. Write $n = |L_1|$ and $m = |L_2|$. Let $\mathcal{A}_{1,\alpha}$ and $\mathcal{A}_{2,\beta}$ be systems of arcs for L_1 and L_2 . Suppose that the arc for K_1 is alpha-parallel and the arc for K_2 is beta-parallel. (The other arcs need not be alpha- or beta-parallel.) Let \mathcal{A}_α and \mathcal{A}_β be the system of arcs for $L_1\#L_2$ which use an alpha-parallel (resp. beta-parallel) arc for K_1 (resp. K_2). Then*

$$\begin{aligned} & (\mathcal{X}_{\Lambda_1}(L_1, \mathcal{A}_{1,\alpha})^{\mathcal{K} \otimes \mathcal{L}_{n-1}}, \mathcal{X}_{\Lambda_1}(L_2, \mathcal{A}_{2,\beta})^{\mathcal{K} \otimes \mathcal{L}_{m-1}}) \hat{\boxtimes} \mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}} \\ & \simeq \mathcal{X}_{\Lambda_1 + \Lambda_2}(L_1\#L_2, \mathcal{A}_\alpha)^{\mathcal{L}_{n+m-1}} \end{aligned}$$

and

$$\begin{aligned} & (\mathcal{X}_{\Lambda_1}(L_1, \mathcal{A}_{1,\alpha})^{\mathcal{K} \otimes \mathcal{L}_{n-1}}, \mathcal{X}_{\Lambda_1}(L_2, \mathcal{A}_{2,\beta})^{\mathcal{K} \otimes \mathcal{L}_{m-1}}) \hat{\boxtimes} \mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}} \\ & \simeq \mathcal{X}_{\Lambda_1 + \Lambda_2}(L_1\#L_2, \mathcal{A}_\beta)^{\mathcal{L}_{n+m-1}}. \end{aligned}$$

15.1 The $W_{\alpha\beta,\alpha}$ and $W_{\alpha\beta,\beta}$ bimodules

We now define the $W_{\alpha\beta,\alpha}$ and $W_{\alpha\beta,\beta}$ bimodules. We consider $W_{\alpha\beta,\beta}$ first. This bimodule has a δ_3^1 which is identical to the merge module M . Additionally, there is a δ_5^1 , determined by the following formula:

$$\delta_5^1(a|b, a'|b', 1|\tau, \tau|1, i_0) = i_1 \otimes \mathcal{V}^{-1} \partial_{\mathcal{U}}(ab)a'(\mathcal{U}\partial_{\mathcal{U}} + \mathcal{V}\partial_{\mathcal{V}})(b')\tau. \quad (15.1)$$

In the above, $a, b, a', b' \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$.

The module $W_{\alpha\beta,\alpha}$ is similar, except it has

$$\delta_5^1(a|b, a'|b', 1|\tau, \tau|1, i_0) = i_1 \otimes \mathcal{V}^{-1} \partial_{\mathcal{U}}(ab)b'(\mathcal{U}\partial_{\mathcal{U}} + \mathcal{V}\partial_{\mathcal{V}})(a')\tau.$$

Lemma 15.2. *The bimodules $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}$ and $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}}$ satisfy the DA bimodule structure relations. Furthermore, both are split Alexander bimodules when we give the incoming algebra factors the discrete partition.*

Proof. We consider the structure relations first. We take the perspective of mapping cone bimodules from Section 14.3. We begin by considering the merge module $\mathcal{K}|\mathcal{K}M^{\mathcal{K}}$. We define bimodules $\mathcal{K}|\mathcal{K}[\mathbf{m}_\varepsilon]^{\mathcal{K}}$ for $\varepsilon \in \{0, 1\}$ by

$$(\mathbf{I}_\varepsilon \otimes \mathbf{I}_\varepsilon) \cdot \mathcal{K}|\mathcal{K}M^{\mathcal{K}} \cdot \mathbf{I}_\varepsilon.$$

The bimodule structure relations are easily seen to be satisfied. There are two DA bimodule morphisms

$$\mathcal{S}_*^1, \mathcal{T}_*^1: \mathcal{K}|\mathcal{K}[\mathbf{m}_0]^{\mathcal{K}} \rightarrow \mathcal{K}|\mathcal{K}[\mathbf{m}_1]^{\mathcal{K}}$$

as follows. We set $\mathcal{S}_j^1 = 0$ if $j \neq 3$. We set $\mathcal{S}_3^1(1|\sigma, \sigma|1, i_0) = i_1 \otimes \sigma$. The map \mathcal{T}_*^1 is similar, but with τ replacing σ . The DA bimodule endomorphisms \mathcal{S}_*^1 and \mathcal{T}_*^1 are easily seen to be cycles. Furthermore,

$$\mathcal{K}|\mathcal{K}M^{\mathcal{K}} = \text{Cone}(\mathcal{S}_*^1 + \mathcal{T}_*^1: \mathcal{K}|\mathcal{K}[\mathbf{m}_0]^{\mathcal{K}} \rightarrow \mathcal{K}|\mathcal{K}[\mathbf{m}_1]^{\mathcal{K}}).$$

Note that we may form four DA bimodule endomorphisms $\phi_*^l, \phi_*^r, \psi_*^l$ and ψ_*^r of $\mathcal{K}|\mathcal{K}[\mathbf{m}_1]^{\mathcal{K}}$ by adapting the construction from Section 14.2. Here, ϕ_*^l differentiates with respect to \mathcal{U} on the left factor of $\mathcal{K} \otimes \mathcal{K}$, while ϕ_*^r differentiates with respect to \mathcal{U} on the right factor. We may also view these maps as endomorphisms of $\mathcal{K}|\mathcal{K}[\mathbf{m}_1]^{\mathcal{K}}$ by boxing with the identity map on $[\mathbf{m}_1]$. With respect to the above notation, we observe that

$$\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}} = \text{Cone}([\mathbf{m}_0] \xrightarrow{\mathcal{S}_*^1 + (\text{id} + \mathcal{V}^{-1}(\phi_*^l + \phi_*^r)) \circ (\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l)} \circ \mathcal{T}_*^1} [\mathbf{m}_1]). \quad (15.2)$$

The map

$$(\text{id} + \mathcal{V}^{-1}(\phi_*^l + \phi_*^r)) \circ (\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^r) \circ \mathcal{T}_*^1$$

is a cycle since it is the composition of cycles. Similarly, \mathcal{S}_*^1 is a cycle by direct computation. Hence the DA bimodule relations are satisfied for $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}}$. The same argument works for $\mathcal{K}|\mathcal{K}W_{\alpha\beta,\beta}^{\mathcal{K}}$.

Finally, the claim about split continuity is proven by an easy modification of the proof for the merge module, from Lemma 8.3. \blacksquare

15.2 Proof of Theorem 15.1

We now give a proof of Theorem 15.1.

Proof of Theorem 15.1. The proof is based on our formula for changing the path in the link surgery formula of Corollary 13.5. It is easier to analyze the surgery hypercubes, instead of their interpretation as type-D modules, so we focus on the statement in terms of hypercube complexes first.

Letting $K_1 \subseteq L_1$ and $K_2 \subseteq L_2$ be the special components, and suppose that $\mathcal{A}_{1,\alpha}$ and $\mathcal{A}_{2,\beta}$ are systems of arcs such that K_1 has an alpha-parallel arc and K_2 has a beta-parallel arc. We use the meridional σ -basic systems considered in Section 9.3, which we assume are algebraically rigid. These were constructed in Lemma 9.8. By Theorem 12.1, we know that

$$\begin{aligned} & \mathcal{C}_{\Lambda_1 + \Lambda_2}(L_1 \# L_2, \mathcal{A}_{1,\alpha} \# \mathcal{A}_{2,\beta}) \\ & \cong \text{Cone}(\mathcal{C}^0(L_1) \otimes \mathcal{C}^0(L_2) \xrightarrow{F^{K_1} | F^{K_2} + F^{-K_1} | F^{-K_2}} \mathcal{C}^1(L_1) \otimes \mathcal{C}^1(L_2)). \end{aligned}$$

(Recall that we are writing $F^{\pm K_i}$ for the sum of all surgery maps for oriented sublinks of L_i which contain $\pm K_i$.)

The above surgery hypercube is for the system of arcs $\mathcal{A}_{1,\alpha} \# \mathcal{A}_{2,\beta}$. The arc in $\mathcal{A}_{1,\alpha} \# \mathcal{A}_{2,\beta}$ for $K_1 \# K_2$ is the co-core of the connected sum band. The present statement instead involves σ -basic systems of Heegaard diagrams for $L_1 \# L_2$ which use an alpha- or beta-parallel arc on $K_1 \# K_2$. If we want to change the arc to be alpha-parallel, we may stack the above hypercube for $\mathcal{C}_{\Lambda_1 + \Lambda_2}(L_1 \# L_2, \mathcal{A}_{1,\alpha} \# \mathcal{A}_{2,\beta})$ with the hypercube for moving the base point in a loop following K_2 in the subcube $\mathcal{C}^1(L_1) \otimes \mathcal{C}^1(L_2)$. We may compute this cube using Corollary 13.5. The effect on the resulting surgery hypercube is to replace $F^{-K_1} | F^{-K_2}$ with the expression

$$(\text{id} + \mathcal{V}^{-1} \Phi_w \circ \mathcal{A}_{[K_2]})(F^{-K_1} | F^{-K_2}): \mathcal{C}^1(L_1 \# L_2) \rightarrow \mathcal{C}^1(L_1 \# L_2).$$

We are using Proposition 11.8 to identify $\mathcal{C}^\varepsilon(L_1 \# L_2) \cong \mathcal{C}^\varepsilon(L_1) \otimes \mathcal{C}^\varepsilon(L_2)$, for $\varepsilon \in \{0, 1\}$. The differential on each $\mathcal{C}^\varepsilon(L_1 \# L_2)$ is the tensor product differential, i.e., $\partial_1 \otimes \text{id} + \text{id} \otimes \partial_2$. In particular, the map Φ_w on $\mathcal{C}^\varepsilon(L_1 \# L_2)$ is intertwined with the map

$$\Phi_{w_1} \otimes \text{id} + \text{id} \otimes \Phi_{w_2}$$

under the connected sum isomorphism $\mathcal{C}^1(L_1 \# L_2) \cong \mathcal{C}^1(L_1) \otimes \mathcal{C}^1(L_2)$. Similarly, using the connected sum formulas for hypercubes in Propositions 10.1 and 11.8, it is straightforward to see that with respect to the isomorphism $\mathcal{C}^1(L_1 \# L_2) \cong \mathcal{C}^1(L_1) \otimes \mathcal{C}^1(L_2)$ the map $\mathcal{A}_{[K_2]}$ is intertwined with $\text{id} \otimes \mathcal{A}_{[K_2]}$. By Lemma 13.30 we see that

$$\text{id} \otimes \mathcal{A}_{[K_2]} = \text{id} \otimes (\mathcal{U} \Phi_{w_2} + \mathcal{V} \Psi_{z_2}),$$

as endomorphisms of $\mathcal{C}^\varepsilon(L_1) \otimes \mathcal{C}^\varepsilon(L_2)$. We conclude that

$$\Phi_w \circ \mathcal{A}_{[K_2]} = (\Phi_{w_1} \otimes \text{id} + \text{id} \otimes \Phi_{w_2}) \circ (\text{id} \otimes (\mathcal{U} \Phi_{w_2} + \mathcal{V} \Psi_{z_2})).$$

The addition of this term is encoded exactly by the extra δ_5^1 term in equation (15.1) (cf. Lemma 14.3), so the proof is complete. \blacksquare

15.3 Relation with the AA-identity bimodule

We recall that the AA-identity bimodule $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]$ from Section 8.4 admits an extension to a bimodule

$$\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}.$$

It follows from Theorem 12.1 and Proposition 13.2 that we may use $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]$ to obtain a valid model of the link surgery formula when taking the connected sum of two links (i.e., the resulting hypercube has the correct module actions on the remaining knot components). It is not clear from this argument whether tensoring with $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}$ gives a valid model. In this section, we show that indeed tensoring with $\mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}$ produces a valid model, via the following algebraic result.

Proposition 15.3. *There is a homotopy equivalence*

$$\mathcal{K}|\mathcal{K}W_{\alpha\beta,\alpha}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K}[\mathcal{D}_0]_{\mathbb{F}[U]} \simeq \mathcal{K}|\mathcal{K}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}.$$

The same holds if we replace $W_{\alpha\beta,\alpha}$ with $W_{\alpha\beta,\beta}$.

The above proposition will be verified in several steps.

Lemma 15.4. *The operations ϕ_*^1 and ψ_*^1 satisfy the following relations:*

(1) *As morphisms from $\mathcal{K}[\mathbf{i}_0]^{\mathcal{K}}$ to $\mathcal{K}[\mathbf{i}_1]^{\mathcal{K}}$, we have*

$$\phi_*^1 \circ s_*^1 \simeq s_*^1 \circ \phi_*^1 \quad \text{and} \quad \psi_*^1 \circ s_*^1 \simeq s_*^1 \circ \psi_*^1.$$

Similarly,

$$t_*^1 \circ \phi_*^1 \simeq \mathcal{V}^2 \psi_*^1 \circ t_*^1 \quad \text{and} \quad t_*^1 \circ \psi_*^1 \simeq \mathcal{V}^{-2} \phi_*^1 \circ t_*^1.$$

(2) *As morphisms from $\mathcal{K}|\mathcal{K}[\mathbf{m}_0]^{\mathcal{K}}$ to $\mathcal{K}|\mathcal{K}[\mathbf{m}_1]^{\mathcal{K}}$ we have*

$$\phi_*^l \circ \mathcal{S}_*^1 \simeq \mathcal{S}_*^1 \circ \phi_*^l \quad \text{and} \quad \psi_*^l \circ \mathcal{S}_*^1 \simeq \mathcal{S}_*^1 \circ \psi_*^l$$

and

$$\mathcal{T}_*^1 \circ \phi_*^l \simeq \mathcal{V}^2 \psi_*^l \circ \mathcal{T}_*^1 \quad \text{and} \quad \mathcal{T}_*^1 \circ \psi_*^l \simeq \mathcal{V}^{-2} \phi_*^l \circ \mathcal{T}_*^1.$$

The same relations hold with ϕ_*^r and ψ_*^r in place of ϕ_*^l and ψ_*^l .

Proof. We begin with two algebraic relations on \mathcal{K} . If $f \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$ and $g \in \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$ are such that $\tau a = b \tau$ then

$$\tau \partial_{\mathcal{U}}(a) = \mathcal{V}^2 \partial_{\mathcal{V}}(b) \tau \quad \text{and} \quad \tau \partial_{\mathcal{V}}(a) = \mathcal{V}^{-2} \partial_{\mathcal{U}}(b) \tau.$$

We verify the above equations. If $a = \mathcal{U}^i \mathcal{V}^j$, then $b = \phi^\tau(a) = \mathcal{U}^j \mathcal{V}^{2j-i}$. Then

$$\tau \partial_{\mathcal{U}}(a) = \tau \partial_{\mathcal{U}}(\mathcal{U}^i \mathcal{V}^j) = \tau i \mathcal{U}^{i-1} \mathcal{V}^j = i \mathcal{U}^j \mathcal{V}^{2j-i+1} \tau = \mathcal{V}^2 \partial_{\mathcal{V}}(b) \tau$$

as claimed. Similarly,

$$\tau \partial_{\mathcal{V}}(a) = \tau j \mathcal{U}^i \mathcal{V}^{j-1} = j \mathcal{U}^{j-1} \mathcal{V}^{2j-2-i} \tau = \mathcal{V}^{-2} \partial_{\mathcal{U}}(b) \tau.$$

We now consider the equation $\phi_*^1 \circ s_*^1 \simeq s_*^1 \circ \phi_*^1$. We define a morphism $h_*^1: \mathcal{X}[\mathbf{i}_0]^{\mathcal{K}} \rightarrow \mathcal{X}[\mathbf{i}_1]^{\mathcal{K}}$ via the formula

$$h_*^1(f\sigma, 1) = 1 \otimes \partial_{\mathcal{U}}(f)\sigma.$$

One checks easily that $\partial_{\text{Mor}}(h_*^1) = \phi_*^1 \circ s_*^1 + s_*^1 \circ \phi_*^1$. The other relations with t_*^1 and ϕ_*^1 and ψ_*^1 are proven similarly.

We now consider the relations for $\mathbb{S}_*^1 \circ \phi_*^l \simeq \phi_*^l \circ \mathbb{S}_*^1$. Define $f_*^1 := \mathbb{S}_*^1 \circ \phi_*^l$ and note that f_*^1 has f_4^1 non-trivial, and this map satisfies

$$f_4^1(\sigma|1, 1|\sigma, a|a', i_0) = i_1 \otimes \sigma \partial_{\mathcal{U}}(a)a'.$$

Note that more generally, $f_4^1(a\sigma|a', b|b'\sigma, c|c', i_0) = i_1 \otimes aa'bb'\sigma \partial_{\mathcal{U}}(c)c'$, though we omit the extra possible coefficients on the σ terms to simplify the notation. Define $g_*^1 := \phi_*^l \circ \mathbb{S}_*^1$ and note that g_*^1 has only g_4^1 non-trivial, and

$$g_4^1(a|a', \sigma|1, 1|\sigma, i_0) = i_1 \otimes \partial_{\mathcal{U}}(a)a'\sigma.$$

There is a third map of interest, k_*^1 , which has

$$k_4^1(\sigma|1, a|a', 1|\sigma, i_0) = i_1 \otimes a'\sigma \partial_{\mathcal{U}}(a).$$

A homotopy h_*^1 between f_*^1 and k_*^1 is given by

$$h_3^1(a\sigma|a', b|b'\sigma, i_0) = i_1 \otimes aa'b'\sigma \partial_{\mathcal{U}}(b).$$

A homotopy between k_*^1 and g_*^1 is constructed similarly. The other relations in the lemma statement are proven by minor modifications of the above argument. \blacksquare

Lemma 15.5. *There is a homotopy equivalence*

$$\mathcal{X}|\mathcal{X} W_{\alpha\beta, \alpha}^{\mathcal{K}} \hat{\boxtimes} \mathcal{X}[\mathcal{D}_0]_{\mathbb{F}[U]} \simeq (\mathcal{X}\mathcal{T}^{\mathcal{K}}, \mathcal{X}\mathcal{T}^{\mathcal{K}}) \hat{\boxtimes} \mathcal{X}|\mathcal{X}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}.$$

The same holds if we replace $W_{\alpha\beta, \alpha}$ with $W_{\alpha\beta, \beta}$.

Remark 15.6. The above statement does not hold if we replace $W_{\alpha\beta, \alpha}$ with $W_{\alpha\alpha, \alpha}$, $W_{\alpha\alpha, \beta}$, $W_{\beta\beta, \alpha}$ or $W_{\beta\beta, \beta}$.

Proof of Lemma 15.5. The argument is similar to the proof of Lemma 14.8. The key is to observe that we may define a homotopy $\partial_{\mathcal{W}} = \mathcal{U}\partial_{\mathcal{U}} + \mathcal{V}\partial_{\mathcal{V}}$, and that $\partial_{\mathcal{W}}$ commutes with the action of $U = \mathcal{U}\mathcal{V}$.

In our present context, we want to view both modules in the statement as AA-bimodule mapping cones. For $\varepsilon \in \{0, 1\}$, write

$$\mathcal{X}|\mathcal{X}[\mathbf{md}_{\varepsilon}]_{\mathbb{F}[U]} := (\mathbf{I}_{\varepsilon} \otimes \mathbf{I}_{\varepsilon}) \cdot \mathcal{X}|\mathcal{X}[\mathbb{I}^{\mathfrak{D}}]_{\mathbb{F}[U]}.$$

Both of the bimodules in the statement may be viewed as a mapping cone of an AA-bimodule morphism from $[\mathbf{md}_0]$ to $[\mathbf{md}_1]$.

Similar to equation (15.2), we may view $W_{\alpha\beta,\alpha} \hat{\boxtimes} \mathcal{D}_0$ as the mapping cone of

$$\mathcal{S}_* + (\text{id} + \mathcal{V}^{-1}(\phi_*^l + \phi_*^r)(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l)) \circ \mathcal{T}_*. \quad (15.3)$$

We observe that the morphism $H_{0,1,0}: [\mathbf{md}_1] \rightarrow [\mathbf{md}_1]$ given by $H_{0,1,0}(\mathbf{x}) = \partial_{\mathcal{W}}(\mathbf{x})$ satisfies

$$\partial_{\text{Mor}}(H_{0,1,0}) = \mathcal{U}\phi_*^l + \mathcal{U}\phi_*^r + \mathcal{V}\psi_*^l + \mathcal{V}\psi_*^r.$$

In particular, we obtain that

$$\begin{aligned} & \text{id} + \mathcal{V}^{-1}(\phi_*^l + \phi_*^r)(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l) \\ &= \text{id} + \mathcal{V}^{-1}\phi_*^l(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l) + \mathcal{V}^{-1}\phi_*^r(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l) \\ &\simeq \text{id} + \mathcal{V}^{-1}\phi_*^l(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l) + \mathcal{V}^{-1}\phi_*^r(\mathcal{U}\phi_*^r + \mathcal{V}\psi_*^r) \\ &\simeq \text{id} + \phi_*^l\psi_*^l + \phi_*^r\psi_*^r \\ &\simeq (\text{id} + \phi_*^l\psi_*^l)(\text{id} + \phi_*^r\psi_*^r). \end{aligned}$$

In the last equivalence, we are using the fact that

$$\phi_*^l\psi_*^l\phi_*^r\psi_*^r \simeq 0$$

since $\phi_*^l\psi_*^l \simeq \mathcal{V}^{-1}\phi_*^l(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l)$ (and similarly for $\phi_*^r\psi_*^r$), so

$$\phi_*^l\psi_*^l\phi_*^r\psi_*^r \simeq \mathcal{V}^{-2}\phi_*^l\phi_*^r(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l)(\mathcal{U}\phi_*^r + \mathcal{V}\psi_*^r) \simeq \mathcal{V}^{-2}\phi_*^l\phi_*^r(\mathcal{U}\phi_*^l + \mathcal{V}\psi_*^l)^2 \simeq 0.$$

Hence, the map in equation (15.3) is homotopy equivalent to

$$\mathcal{S}_* + (\text{id} + \phi_*^l\psi_*^l)(\text{id} + \phi_*^r\psi_*^r)\mathcal{T}_*.$$

Arguing similarly to the proof of Lemma 15.4, one sees that the above map is homotopic to the map whose cone is

$$(\mathcal{X}\mathcal{T}^{\mathcal{K}}, \mathcal{X}\mathcal{T}^{\mathcal{K}}) \hat{\boxtimes}_{\mathcal{X}|\mathcal{K}} [\mathbb{I}^{\boxplus}]_{\mathbb{F}[U]},$$

completing the proof. ■

Lemma 15.7. *There is a homotopy equivalence*

$$(\mathcal{X}\mathcal{T}^{\mathcal{K}}, \mathcal{X}\mathcal{T}^{\mathcal{K}}) \hat{\boxtimes}_{\mathcal{X}|\mathcal{K}} M^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{X}|\mathcal{K}} \mathcal{X}\mathcal{T}^{\mathcal{K}} \simeq_{\mathcal{X}|\mathcal{K}} M^{\mathcal{K}}.$$

Proof. The proof is similar to the proof of Lemma 15.5. We observe first by computing algebraically using the techniques of Lemma 15.4 that $(\mathcal{X}\mathcal{T}^{\mathcal{K}}, \mathcal{X}\mathcal{T}^{\mathcal{K}}) \hat{\boxtimes}_{\mathcal{X}|\mathcal{K}} M^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{X}|\mathcal{K}} \mathcal{X}\mathcal{T}^{\mathcal{K}}$ is homotopy equivalent to the mapping cone

$$[\mathbf{m}_0] \xrightarrow{\mathcal{S}_*^1 + (\text{id} + (\phi_*^l + \phi_*^r)(\psi_*^l + \psi_*^r))(\text{id} + \phi_*^l\psi_*^l)(\text{id} + \phi_*^r\psi_*^r)\mathcal{T}_*^1} [\mathbf{m}_1]. \quad (15.4)$$

One observes by direct computation that

$$(\text{id} + (\phi_*^l + \phi_*^r)(\psi_*^l + \psi_*^r))(\text{id} + \phi_*^l \psi_*^l)(\text{id} + \phi_*^r \psi_*^r) \simeq (\text{id} + \phi_*^l \psi_*^r)(\text{id} + \psi_*^l \phi_*^r).$$

We define the mapping cone complexes

$$\mathcal{K}|\mathcal{K}X := [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \phi_*^l \psi_*^r)(\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1} [\mathbf{m}_1],$$

and

$$\mathcal{K}|\mathcal{K}Y := [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1(\text{id} + \phi_*^l \psi_*^r)} [\mathbf{m}_1].$$

We now claim that $\mathcal{K}|\mathcal{K}X$ and $\mathcal{K}|\mathcal{K}Y$ are homotopy equivalent. We construct a homotopy equivalence F by way of the following diagram:

$$\begin{array}{ccc} \mathcal{K}|\mathcal{K}X^{\mathcal{K}} & [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \phi_*^l \psi_*^r)(\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1} & [\mathbf{m}_1] \\ \downarrow F = & \text{id} + \phi_*^l \psi_*^r \downarrow & \text{id} + \phi_*^l \psi_*^r \downarrow \\ \mathcal{K}|\mathcal{K}Y^{\mathcal{K}} & [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1(\text{id} + \phi_*^l \psi_*^r)} & [\mathbf{m}_1] \end{array}$$

The length 2 map above is a morphism which realizes the chain homotopies

$$(\text{id} + \phi_*^l \psi_*^r)^2(\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1 \simeq (\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1(\text{id} + \phi_*^l \psi_*^r)^2,$$

and

$$[\text{id} + \phi_*^l \psi_*^r, S_*^1] \simeq 0.$$

Note that by a similar construction we can construct a map in the opposite direction:

$$\begin{array}{ccc} \mathcal{K}|\mathcal{K}Y^{\mathcal{K}} & [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1(\text{id} + \phi_*^l \psi_*^r)} & [\mathbf{m}_1] \\ \downarrow G = & \text{id} + \phi_*^l \psi_*^r \downarrow & \text{id} + \phi_*^l \psi_*^r \downarrow \\ \mathcal{K}|\mathcal{K}X^{\mathcal{K}} & [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \phi_*^l \psi_*^r)(\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1} & [\mathbf{m}_1] \end{array}$$

Note that since $(\text{id} + \phi_*^l \psi_*^r)^2 \simeq \text{id}$ as a map from $[\mathbf{m}_i]$ to $[\mathbf{m}_i]$, it follows that $G \circ F$ is homotopic to a map $J: \mathcal{K}|\mathcal{K}X^{\mathcal{K}} \rightarrow \mathcal{K}|\mathcal{K}X^{\mathcal{K}}$ which takes the following form:

$$\begin{array}{ccc} \mathcal{K}|\mathcal{K}X^{\mathcal{K}} & [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \phi_*^l \psi_*^r)(\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1} & [\mathbf{m}_1] \\ \downarrow J = & \text{id} \downarrow & \text{id} \downarrow \\ \mathcal{K}|\mathcal{K}X^{\mathcal{K}} & [\mathbf{m}_0] \xrightarrow{S_*^1 + (\text{id} + \phi_*^l \psi_*^r)(\text{id} + \psi_*^l \phi_*^r)\mathcal{T}_*^1} & [\mathbf{m}_1] \end{array}$$

Note that $J^2 = \text{id}$. Therefore, the map F has a left homotopy inverse $J \circ G$ since $J \circ G \circ F \simeq \text{id}$. One may modify the construction above to construct a right inverse

to F by similar reasoning. By basic algebra, the left and right inverses of F must be homotopic. Hence F is a homotopy equivalence.

Finally, we use Lemma 15.4 to see that

$$(\text{id} + \psi_*^l \phi_*^r) \mathcal{T}_*^1 (\text{id} + \phi_*^l \psi_*^r) \simeq (\text{id} + \psi_*^l \phi_*^r)^2 \mathcal{T}_*^1 \simeq \mathcal{T}_*^1.$$

It follows that the $(\mathcal{K} \mathcal{T}^{\mathcal{K}}, \mathcal{K} \mathcal{T}^{\mathcal{K}}) \hat{\boxtimes} \mathcal{K} |_{\mathcal{K}} M^{\mathcal{K}} \boxtimes \mathcal{K} \mathcal{T}^{\mathcal{K}}$ (shown in equation (15.4)) is homotopy equivalent to the mapping cone

$$\text{Cone}(\mathcal{S}_*^1 + \mathcal{T}_*^1: \mathcal{K} |_{\mathcal{K}} [\mathbf{m}_0]^{\mathcal{K}} \rightarrow \mathcal{K} |_{\mathcal{K}} [\mathbf{m}_1]^{\mathcal{K}})$$

which is the definition of $\mathcal{K} |_{\mathcal{K}} M^{\mathcal{K}}$. ■

Proposition 15.3 is now proven as follows. Lemma 15.5 implies that

$$\mathcal{K} |_{\mathcal{K}} W_{\alpha\beta,\alpha}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} [\mathcal{D}_0]_{\mathbb{F}[U]} \simeq (\mathcal{K} \mathcal{T}^{\mathcal{K}}, \mathcal{K} \mathcal{T}^{\mathcal{K}}) \hat{\boxtimes} \mathcal{K} |_{\mathcal{K}} M^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} [\mathcal{D}_0]_{\mathbb{F}[U]}.$$

Lemmas 14.7 and 15.7 imply that

$$(\mathcal{K} \mathcal{T}^{\mathcal{K}}, \mathcal{K} \mathcal{T}^{\mathcal{K}}) \hat{\boxtimes} \mathcal{K} |_{\mathcal{K}} M^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} [\mathcal{D}_0]_{\mathbb{F}[U]} \simeq \mathcal{K} |_{\mathcal{K}} M^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} \mathcal{T}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} [\mathcal{D}_0]_{\mathbb{F}[U]}.$$

Lemma 15.5 implies that the above is homotopy equivalent to

$$\mathcal{K} |_{\mathcal{K}} M^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} [\mathcal{D}_0]_{\mathbb{F}[U]},$$

which is the definition of $\mathcal{K} |_{\mathcal{K}} [\mathbb{I}^{\ni}]_{\mathbb{F}[U]}$.

15.4 Relation with the bordered invariant of $S^1 \times P$

In this section, we sketch why the bimodules $\mathcal{K} |_{\mathcal{K}} W_{\alpha\beta,\beta}^{\mathcal{K}}$ and $\mathcal{K} |_{\mathcal{K}} W_{\alpha\beta,\alpha}^{\mathcal{K}}$ can naturally be interpreted as the bordered invariants for $S^1 \times P$, where P denotes a 2-dimensional disk with 2 subdisks removed. In fact, this is a consequence of the connected sum formula and a computation from [3] of the invariant for the identity mapping cylinder over the 2-torus.

Proposition 15.8. *The bimodules $\mathcal{K} |_{\mathcal{K}} W_{\alpha\beta,\beta}^{\mathcal{K}}$ and $\mathcal{K} |_{\mathcal{K}} W_{\alpha\beta,\alpha}^{\mathcal{K}}$ are the bordered invariants for $S^1 \times P$, for various choices of arc systems on the boundary components.*

Remark 15.9. The arc systems on the boundary are compatible with the following gluing convention:

- (1) We glue alpha bordered boundary components to alpha bordered boundary components, and vice-versa.
- (2) $\mathcal{K} |_{\mathcal{K}} [\mathbb{I}^{\ni}]$ has one alpha bordered boundary and one beta bordered boundary component.

Since we switch from type-D to type-A by tensoring with $\mathcal{X}|\mathcal{X}\mathbb{I}^\mathcal{D}$, this convention means that a type-A component of \mathcal{K} is “alpha bordered” if the arc system is beta-parallel, whereas a type-D component of \mathcal{K} is “alpha bordered” if the corresponding arc system is alpha-parallel.

Proof of Proposition 15.8. A surgery description of $S^1 \times P$ is given in Figure 15.1. We can view the 5-component link appearing in this surgery description as being obtained as the connected sum of two copies of a 3-component link J , as shown in Figure 15.1. (The link J is a connected sum of two Hopf links.) Write $\mathcal{A}_{\alpha,\beta}$ for an arc system on J where the top component is alpha-parallel, and the bottom component is beta-parallel. The middle component can have either alpha- or beta-parallel arc. Write $\mathcal{X}\mathcal{X}_0(J, \mathcal{A}_{\alpha,\beta})^\mathcal{K}$ for the DA bimodule where the left action is the top-most component of J , and the right action is for the bottom-most component. The middle component is surgered out. (I.e., we take the type-D module for the 3-component link, tensor $\mathcal{X}\mathcal{D}_0$ to the middle component, and tensor $\mathcal{X}|\mathcal{X}\mathbb{I}^\mathcal{D}$ to the top component.) It is proven in [3, Theorem 1.7] that

$$\mathcal{X}\mathcal{X}_0(J, \mathcal{A}_{\alpha,\beta})^\mathcal{K} \simeq \mathcal{X}\mathcal{X}_0(J, \mathcal{A}_{\beta,\alpha})^\mathcal{K} \simeq \mathcal{X}[\mathbb{I}]^\mathcal{K}. \tag{15.5}$$

(This computation is obtained by tensoring two copies of the Hopf link complex which we compute in Chapter 16 of the present memoir.) It follows from our connected sum formula that if $\mathcal{X}|\mathcal{X}\mathcal{X}(L)^\mathcal{K}$ denotes the surgery module for the 5-component link describing $S^1 \times P$, then

$$\mathcal{X}|\mathcal{X}\mathcal{X}(L)^\mathcal{K} \simeq (\mathcal{X}\mathcal{X}_0(J, \mathcal{A}_{\alpha,\beta})^\mathcal{K}, \mathcal{X}\mathcal{X}_0(J, \mathcal{A}_{\beta,\alpha})^\mathcal{K}) \hat{\boxtimes} \mathcal{X}|\mathcal{X}W_{\beta\alpha,\alpha}^\mathcal{K}.$$

By equation (15.5), the above is homotopy equivalent to $\mathcal{X}|\mathcal{X}W_{\beta\alpha,\alpha}^\mathcal{K}$, completing the proof. ■




Figure 15.1. A Kirby diagram for $S^1 \times P$. The components marked 0 are surgered, while a tubular neighborhood of the unmarked components is removed. On the right, we describe the 5-component link as a connected sum of two copies of a 3-component link J .

Chapter 16

The link surgery complex of the Hopf link

In this chapter, we compute the link surgery complex of the Hopf link. Our main result is Proposition 16.1.

16.1 Computation of the surgery complex

In this section, we compute the surgery hypercube for the Hopf link. We use the following model of the Hopf link:

$$\mathcal{EFL}(H) \cong \begin{array}{ccc} \mathbf{a} \leftarrow u_2 - & \mathbf{b} & \\ \uparrow & & \downarrow \\ v_1 & & u_1 \\ | & & \downarrow \\ \mathbf{c} - v_2 \rightarrow & \mathbf{d} & \end{array} \quad (16.1)$$

This complex is realized by the diagram in Figure 16.1. This complex is for the *negative* Hopf link, i.e., the two components have linking number -1 . The complex of the positive Hopf link has a similar model. The techniques of this section may also be applied, essentially verbatim, to the positive Hopf link, though we focus on the negative one for concreteness.

The Alexander gradings $A = (A_1, A_2)$ and Maslov gradings $\text{gr} = (\text{gr}_{\mathbf{w}}, \text{gr}_{\mathbf{z}})$ of the generators are

$$\begin{aligned} A(\mathbf{a}) &= \left(-\frac{1}{2}, \frac{1}{2}\right), & \text{gr}(\mathbf{a}) &= (0, 0), \\ A(\mathbf{b}) &= \left(-\frac{1}{2}, -\frac{1}{2}\right), & \text{gr}(\mathbf{b}) &= (-1, 1), \\ A(\mathbf{c}) &= \left(\frac{1}{2}, \frac{1}{2}\right), & \text{gr}(\mathbf{c}) &= (1, -1), \\ A(\mathbf{d}) &= \left(\frac{1}{2}, -\frac{1}{2}\right), & \text{gr}(\mathbf{d}) &= (0, 0). \end{aligned}$$

Note that we are using the conventions where the top degree generator of $\text{HF}^-(S^3, \mathbf{w})$ is in Maslov grading 0. See Remark 6.3.

The main result of this section is the following result.

Proposition 16.1. *Suppose that $\Lambda = (\lambda_1, \lambda_2)$ is an integral framing on the negative Hopf link H . Let $\mathcal{A}_{\alpha\alpha}$ be a system of arcs for the Hopf link where both arcs are*

alpha-parallel. The link surgery hypercube $\mathcal{C}_\Lambda(H, \mathcal{A}_{\alpha\alpha})$ has the following maps (up to overall homotopy equivalence):

- (1) *The maps $\Phi_{H,L_2}^{L_1}$ and $\Phi_{L_1,\emptyset}^{L_1}$ are the canonical inclusions of localization. The maps $\Phi_{H,L_2}^{-L_1}$ and $\Phi_{L_1,\emptyset}^{-L_1}$ are both given by the following formula:*

$$\Phi^{-L_1} = \begin{array}{l} \mathbf{a} \longmapsto \mathcal{V}_1^{\lambda_1-1} \mathbf{d} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathcal{V}_1^{\lambda_1+1} \mathbf{b} + \mathcal{V}_1^{\lambda_1} \mathcal{U}_2 \mathbf{c} \\ \mathbf{d} \longmapsto \mathcal{V}_1^{\lambda_1} \mathcal{U}_2 \mathbf{d} \end{array}$$

- (2) *The maps $\Phi_{H,L_1}^{L_2}$ and $\Phi_{L_2,\emptyset}^{L_2}$ are the canonical inclusions of localization. The maps $\Phi_{H,L_1}^{-L_2}$ and $\Phi_{L_2,\emptyset}^{-L_2}$ are both given by the following formula:*

$$\Phi^{-L_2} = \begin{array}{l} \mathbf{a} \longmapsto \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \mathbf{a} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathcal{V}_2^{\lambda_2+1} \mathbf{b} + \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \mathbf{c} \\ \mathbf{d} \longmapsto \mathcal{V}_2^{\lambda_2-1} \mathbf{a} \end{array}$$

- (3) *The length 2 map $\Phi_{H,\emptyset}^{-H}$ is given by the following formula:*

$$\Phi_{H,\emptyset}^{-H} = \begin{array}{l} \mathbf{a} \longmapsto \mathcal{V}_1^{\lambda_1-2} \mathcal{V}_2^{\lambda_2-1} \mathbf{c} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2} \mathbf{d} \\ \mathbf{d} \longmapsto \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2-2} \mathbf{c} \end{array}$$

The length 2 maps for other orientations of the Hopf link vanish.

We now compute the link surgery hypercube maps for the Hopf link. For the purposes of computation, it is easier to compute the induced maps after quotienting variables, instead of after localizing. Since the link surgery hypercube maps are defined by computing holomorphic polygons on diagrams where we have deleted certain base points, this is easily seen to contain equivalent information as the ordinary surgery hypercube. Compare Lemma 6.5.

As an example, if $L_1 \subseteq H$, then the descent map from surgery hypercube takes the form

$$\Phi_{H,H \setminus L_1}^{-L_1} : \mathcal{CF}\mathcal{L}(H) \rightarrow \mathcal{V}_1^{-1} \cdot \mathcal{CF}\mathcal{L}(H).$$

The map $\Phi_{H,H\setminus L_1}^{-L_1}$ is completely determined by the choice of framing on H , and the induced homotopy equivalence on the level of quotients

$$\tilde{\Phi}_{H,H\setminus L_1}^{-L_1} : \mathcal{CFL}(H)/(\mathcal{U}_1 - 1) \rightarrow \mathcal{CFL}(H)/(\mathcal{V}_1 - 1).$$

In this section, we write $\tilde{\Phi}^{\vec{N}}$ for the descent map on the level of quotients of $\mathcal{CFL}(H)$.

Lemma 16.2. *Write L_1 and L_2 for the two components of the negative Hopf link. Up to overall equivalence, the maps for the Hopf link are as follows:*

(1) *The map*

$$\tilde{\Phi}_{H,L_2}^{-L_1} : \mathcal{CFL}(H)/(\mathcal{U}_1 - 1) \rightarrow \mathcal{CFL}(H)/(\mathcal{V}_1 - 1)$$

is given by the following formula:

$$\tilde{\Phi}_{H,L_2}^{-L_1} = \begin{array}{l} \mathbf{a} \longmapsto \mathbf{d} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathbf{b} + \mathcal{U}_2\mathbf{c} \\ \mathbf{d} \longmapsto \mathcal{U}_2\mathbf{d} \end{array}$$

extended $\mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]$ -equivariantly, and sending \mathcal{V}_1 to \mathcal{U}_1 .

(2) *The map*

$$\tilde{\Phi}_{L_1,\emptyset}^{-L_1} : \mathcal{CFL}(H)/(\mathcal{U}_1 - 1, \mathcal{V}_2 - 1) \rightarrow \mathcal{CFL}(H)/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1)$$

is given by the same formula as $\tilde{\Phi}_{H,L_2}^{-L_1}$, except with \mathcal{V}_2 set to 1.

(3) *The map*

$$\tilde{\Phi}_{H,L_1}^{-L_2} : \mathcal{CFL}(H)/(\mathcal{U}_2 - 1) \rightarrow \mathcal{CFL}(H)/(\mathcal{V}_2 - 1)$$

is given by the following formula:

$$\tilde{\Phi}_{H,L_1}^{-L_2} = \begin{array}{l} \mathbf{a} \longmapsto \mathcal{U}_1\mathbf{a} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathbf{b} + \mathcal{U}_1\mathbf{c} \\ \mathbf{d} \longmapsto \mathbf{a} \end{array}$$

extended equivariantly over $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1]$, and sending \mathcal{V}_2 to \mathcal{U}_2 .

(4) *The map*

$$\tilde{\Phi}_{L_2,\emptyset}^{-L_2} : \mathcal{CFL}(H)/(\mathcal{V}_1 - 1, \mathcal{U}_2 - 1) \rightarrow \mathcal{CFL}(H)/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1)$$

is given by the same formula as $\tilde{\Phi}_{H,L_1}^{-L_2}$, except with \mathcal{V}_1 set to 1.

(5) *The length 2 map*

$$\tilde{\Phi}_{H,\emptyset}^{-H} : \mathcal{CFL}(H)/(\mathcal{U}_1 - 1, \mathcal{U}_2 - 1) \rightarrow \mathcal{CFL}(H)/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1)$$

is given by the following formula:

$$\tilde{\Phi}_{H,\emptyset}^{-H} = \begin{array}{l} \mathbf{a} \longmapsto \mathbf{c} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathbf{d} \\ \mathbf{d} \longmapsto \mathbf{c} \end{array}$$

extended to send \mathcal{V}_1 to \mathcal{U}_1 and \mathcal{V}_2 to \mathcal{U}_2 .

(6) *The length 1 maps with positively oriented components, e.g., $\tilde{\Phi}_{H,L_2}^{L_1}$, are the canonical quotient maps. The length 2 maps vanish for any orientation other than $-H$.*

Remark 16.3. As we will see in the proof of the above lemma, all of the maps in the surgery hypercube for H are uniquely determined by the chain complex $\mathcal{CFL}(H)$, except for the length 2 map $\tilde{\Phi}_{H,L_2}^{L_1}$, and up to homotopy there are two possible choices. We will use additional algebraic restrictions on the link surgery hypercube to determine which choice of length 2 map is correct.

Proof of Lemma 16.2. We focus first on the length 1 maps. Consider the following complexes:

$$\mathcal{C}/(\mathcal{U}_1 - 1) = \begin{array}{ccc} \mathbf{a} \leftarrow \mathcal{U}_2 - \mathbf{b} & & \\ \uparrow & & | \\ \mathcal{V}_1 & & \mathbf{1} \\ | & & \downarrow \\ \mathbf{c} - \mathcal{V}_2 \rightarrow \mathbf{d} & & \end{array} \quad \text{and} \quad \mathcal{C}/(\mathcal{V}_1 - 1) = \begin{array}{ccc} \mathbf{a} \leftarrow \mathcal{U}_2 - \mathbf{b} & & \\ \uparrow & & | \\ \mathbf{1} & & \mathcal{U}_1 \\ | & & \downarrow \\ \mathbf{c} - \mathcal{V}_2 \rightarrow \mathbf{d} & & \end{array}$$

The map $\tilde{\Phi}_{H,L_2}^{-L_1}$ is a homotopy equivalence between these two complexes, which is $\mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]$ -equivariant, and sends \mathcal{V}_1 to \mathcal{U}_1 . We perform the following changes of basis:

$$\mathcal{C}/(\mathcal{U}_1 - 1) \cong \begin{array}{ccc} \mathbf{a} & \mathbf{b} & \\ \uparrow & | & \\ \mathcal{V}_1 + \mathcal{U}_2 \mathcal{V}_2 & \mathbf{1} & \\ | & \downarrow & \\ \mathbf{c} + \mathcal{V}_2 \mathbf{b} & \mathbf{d} + \mathcal{U}_2 \mathbf{a} & \end{array} \quad \text{and} \quad \mathcal{C}/(\mathcal{V}_1 - 1) \cong \begin{array}{ccc} \mathbf{a} + \mathcal{V}_2 \mathbf{d} & \mathbf{b} + \mathcal{U}_2 \mathbf{c} & \\ \uparrow & | & \\ \mathbf{1} & \mathcal{U}_1 + \mathcal{U}_2 \mathcal{V}_2 & \\ | & \downarrow & \\ \mathbf{c} & \mathbf{d} & \end{array} \tag{16.2}$$

We may view each of the above complexes as the knot Floer complex for an unknot, with an extra free base point. In the following, we write $\text{gr}_{\mathbf{w}}$ for the grading induced by the complete collection of base points consisting of the \mathbf{w} -type link base points (i.e., the ones where the unknot intersects the Heegaard surface negatively), and all the free base points. We write $\text{gr}_{\mathbf{z}}$ for the grading induced by the complete collection

consisting of the \mathbf{z} -type base points and the free base points. We may easily compute the $(\text{gr}_w, \text{gr}_z)$ -gradings of all of the quotiented complexes, viewed in this manner:

$$\begin{array}{ll}
 \mathcal{C}/(\mathcal{U}_1 - 1) : & \begin{array}{l} \mathbf{a} \xrightarrow{(\text{gr}_w, \text{gr}_z)} (0, 0) \\ \mathbf{b} \longrightarrow (-1, 1) \\ \mathbf{c} \longrightarrow (-1, -1) \\ \mathbf{d} \longrightarrow (-2, 0) \end{array} & \mathcal{C}/(\mathcal{V}_1 - 1) : & \begin{array}{l} \mathbf{a} \longrightarrow (0, -2) \\ \mathbf{b} \longrightarrow (-1, -1) \\ \mathbf{c} \longrightarrow (1, -1) \\ \mathbf{d} \longrightarrow (0, 0) \end{array} \\
 \\
 \mathcal{C}/(\mathcal{U}_2 - 1) : & \begin{array}{l} \mathbf{a} \longrightarrow (-2, 0) \\ \mathbf{b} \longrightarrow (-1, 1) \\ \mathbf{c} \longrightarrow (-1, -1) \\ \mathbf{d} \longrightarrow (0, 0) \end{array} & \mathcal{C}/(\mathcal{V}_2 - 1) : & \begin{array}{l} \mathbf{a} \longrightarrow (0, 0), \\ \mathbf{b} \longrightarrow (-1, -1) \\ \mathbf{c} \longrightarrow (1, -1) \\ \mathbf{d} \longrightarrow (0, -2) \end{array}
 \end{array}$$

The homotopy equivalence $\tilde{\Phi}_{H,L_2}^{-L_1}$ is grading preserving with respect to the above gradings, since it is the map for changing Heegaard diagrams. It is straightforward to see from equation (16.2) that there is a unique choice of a non-zero chain map from $\mathcal{C}/(\mathcal{U}_1 - 1)$ to $\mathcal{C}/(\mathcal{V}_1 - 1)$ which is $(\text{gr}_w, \text{gr}_z)$ -grading preserving, is $\mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]$ -equivariant, and sends \mathcal{V}_1 to \mathcal{U}_1 . (Note here we are using the $(\text{gr}_w, \text{gr}_z)$ -bigrading obtained by viewing $\mathcal{C}/(\mathcal{U}_1 - 1)$ and $\mathcal{C}/(\mathcal{V}_1 - 1)$ as the complex of an unknot, shown above; we are not using the gradings naively inherited from the Hopf link.) This map is the homotopy equivalence $\tilde{\Phi}_{H,L_2}^{-L_1}$ in the map in the statement. By a symmetric argument, one obtains the formula for $\tilde{\Phi}_{H,L_1}^{-L_2}$ in the statement.

Next, by the definition of a σ -basic system, the map $\tilde{\Phi}_{L_1, \emptyset}^{-L_1}$ is obtained by quotienting the domain and range of $\tilde{\Phi}_{H,L_2}^{-L_1}$ by $\mathcal{V}_2 - 1$. The maps $\tilde{\Phi}_{H,L_1}^{-L_2}$ and $\tilde{\Phi}_{L_2, \emptyset}^{-L_2}$ are related by a similar relation.

We now consider the diagonal map of the cube. Firstly, since we are working with a basis system, the length 2 map vanishes except for the orientation $-H$ (i.e., all components oriented negatively). In our present notation, this is the map

$$\tilde{\Phi}_{H, \emptyset}^{-H} : \mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{U}_2 - 1) \rightarrow \mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1).$$

We define the map

$$E := \tilde{\Phi}_{L_1, \emptyset}^{-L_1} \circ \tilde{\Phi}_{H,L_1}^{-L_2} + \tilde{\Phi}_{L_2, \emptyset}^{-L_2} \circ \tilde{\Phi}_{H,L_2}^{-L_1}. \tag{16.3}$$

Using our computations of the length 1 maps of the cube, we compute E to be

$$E = \begin{array}{l} \mathbf{a} \longrightarrow \mathbf{a} + \mathbf{d} \\ \mathbf{b} \longrightarrow 0 \\ \mathbf{c} \longrightarrow (\mathcal{U}_1 + \mathcal{U}_2)\mathbf{c} \\ \mathbf{d} \longrightarrow \mathbf{a} + \mathbf{d}, \end{array}$$

extended linearly by sending \mathcal{V}_1 to \mathcal{U}_1 and \mathcal{V}_2 to \mathcal{U}_2 .

Up to an overall chain homotopy intertwining $\mathbb{F}[\mathcal{V}_1, \mathcal{V}_2]$ and $\mathbb{F}[\mathcal{U}_1, \mathcal{U}_2]$, there are two choices for a graded null-homotopy of the map E . These are the maps

$$H_1 := \begin{array}{l} \mathbf{a} \longmapsto \mathbf{c} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto 0 \\ \mathbf{d} \longmapsto \mathbf{c}, \end{array} \quad \text{and} \quad H_2 := \begin{array}{l} \mathbf{a} \longmapsto \mathbf{c} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathbf{d} \\ \mathbf{d} \longmapsto \mathbf{c}. \end{array}$$

To prove this, we first observe that H_1 and H_2 are easily checked to be valid homotopies. Any other valid homotopy will differ from H_1 by a $+1$ graded chain map from $\mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{U}_2 - 1)$ to $\mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1)$. However, it is easily checked that up to chain homotopy, there is exactly one non-zero $+1$ graded map from $\mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{U}_2 - 1)$ to $\mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1)$ which sends \mathcal{V}_i to \mathcal{U}_i . This is the map which sends \mathbf{c} to \mathbf{d} , sends \mathcal{V}_i to \mathcal{U}_i , and vanishes on other generators. Adding this map to H_1 yields H_2 . Hence, H_1 and H_2 are the only possible choices, up to further homotopy. We prove in the subsequent Lemma 16.4 that H_1 is not a valid choice when we use the arc system $\mathcal{A}_{\alpha\alpha}$, and hence the length 2 map of the surgery hypercube is H_2 . ■

We let $(S^2, \alpha, \beta, \{w_1, w_2\}, \{z_1, z_2\})$ be the diagram shown in Figure 16.1, and we have $p \in \alpha$ for the point shown therein. We define a map A_p by counting disks weighted by the difference in multiplicities on the two sides of the point p . Note that the map A_p may be identified with the relative homology map considered in Section 13.4 if we pick a path λ which connects a base point on L_1 to a base point on L_2 and satisfies $\lambda \cap \alpha = \{p\}$.

The map A_p is defined on $\mathcal{CF}\mathcal{L}(H)$, and is easily computed by counting holomorphic bigons to be given by the formula

$$A_p = \begin{array}{ccc} \mathbf{a} & \xrightarrow{\mathcal{V}_2} & \mathbf{b} \\ \downarrow \mathcal{U}_1 & & \uparrow \mathcal{V}_1 \\ \mathbf{c} & \xleftrightarrow[\mathcal{U}_2]{} & \mathbf{d} \end{array} = \begin{array}{l} \mathbf{a} \longmapsto \mathcal{V}_2\mathbf{b} + \mathcal{U}_1\mathbf{c} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathcal{V}_2\mathbf{d} \\ \mathbf{d} \longmapsto \mathcal{U}_2\mathbf{c} + \mathcal{V}_1\mathbf{b}. \end{array}$$

The map A_p has an extension to hypercubes, similar to the construction described in Section 13.4. We can describe this by considering a formal endomorphism

$$F_p: \alpha \rightarrow \alpha,$$

and setting the hypercube action \mathcal{A}_p to be $\mu_2^{\text{Tw}}(F_p, -)$.

As in Lemma 13.31 we can extend the above construction to give an endomorphism \mathcal{A}_p of the link surgery formula when we are using a σ -basic system of Heegaard diagrams for H where the arc system consists of only alpha-parallel arcs. In




Figure 16.1. A diagram for the Hopf link with a special point $p \in \alpha$.

this case, the point-pushing diffeomorphisms never change α . We consider this hypercube on the 2-dimensional hypercube used to define all of the link surgery structure maps for negatively oriented sublinks of H . With these considerations in mind, we obtain a diagram of the following form:

$$\begin{array}{ccccc}
 \mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{U}_2 - 1) & \xrightarrow{\tilde{\Phi}^{-L_1}} & \mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{U}_2 - 1) & & \\
 \downarrow \tilde{\Phi}^{-L_2} & \searrow A_p & \downarrow h_p^{-L_1} & \swarrow A_p & \\
 & \mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{U}_2 - 1) & \xrightarrow{\tilde{\Phi}^{-L_1}} & \mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{U}_2 - 1) & \\
 \downarrow \tilde{\Phi}^{-L_2} & \searrow H & \downarrow \tilde{\Phi}^{-L_2} & \swarrow H & \downarrow \tilde{\Phi}^{-L_2} \\
 \mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{V}_2 - 1) & \xrightarrow{\tilde{\Phi}^{-L_1}} & \mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1) & & \\
 \downarrow A_p & \searrow h_p^{-L_2} & \downarrow \omega_p & \swarrow h_p^{-L_2} & \downarrow A_p \\
 \mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{V}_2 - 1) & \xrightarrow{\tilde{\Phi}^{-L_1}} & \mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1) & & \\
 & \downarrow h_p^{-L_1} & & & \\
 & \mathcal{C}/(\mathcal{U}_1 - 1, \mathcal{V}_2 - 1) & \xrightarrow{\tilde{\Phi}^{-L_1}} & \mathcal{C}/(\mathcal{V}_1 - 1, \mathcal{V}_2 - 1) &
 \end{array}
 \tag{16.4}$$

In the above diagram, both maps labeled $h_p^{-L_2}$ are induced by a single map from $\mathcal{C}/(\mathcal{U}_2 - 1)$ to $\mathcal{C}/(\mathcal{V}_2 - 1)$. The map labeled $h_p^{-L_1}$ is similar.

The diagram in equation (16.4) does not satisfy the hypercube relations. Instead, if we view it as $\text{Cone}(\mathcal{A}_p)$, where \mathcal{A}_p is a morphism from the back face to the front face, then

$$[\partial, \mathcal{A}_p] = (\mathcal{U}_1 \mathcal{V}_1 + \mathcal{U}_2 \mathcal{V}_2) \cdot \text{id}.$$

We now analyze the maps labeled $h_p^{-L_i}$. Consider first $h_p^{-L_2}$. Both instances of this map are induced by a single map from $\mathcal{C}/(\mathcal{U}_2 - 1)$ to $\mathcal{C}/(\mathcal{V}_2 - 1)$, for which we also write $h_p^{-L_2}$. This map preserves the Maslov bigradings, if we view $\mathcal{C}/(\mathcal{U}_2 - 1)$ and $\mathcal{C}/(\mathcal{V}_2 - 1)$ as the link Floer complexes for an unknot with an extra free base point. Extending the observation from the proof of Lemma 16.2 that we used to compute the length 1 maps, there were exactly two grading preserving chain maps from $\mathcal{C}/(\mathcal{U}_2 - 1)$ to $\mathcal{C}/(\mathcal{V}_2 - 1)$ which are $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1]$ -equivariant, and send \mathcal{V}_2 to \mathcal{U}_2 : the zero map, and the map labeled $\tilde{\Phi}^{-L_2}$. This observation constrains $h_p^{-L_2}$ to be one of two maps, which differ by the map $\tilde{\Phi}^{-L_2}$. The same comment applies for $h_p^{-L_1}$.

We compute that directly

$$[\tilde{\Phi}^{-L_1}, A_p] = 0.$$

Hence, by the above observation, we conclude that

$$h_p^{-L_1} = s \cdot \tilde{\Phi}^{-L_1} \tag{16.5}$$

for some $s \in \mathbb{F}$.

Similarly, we compute that

$$[\tilde{\Phi}^{-L_2}, A_p] = \begin{array}{l} \mathbf{a} \longmapsto 0 \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathcal{U}_2 \mathbf{a} + \mathcal{U}_1 \mathbf{d} \\ \mathbf{d} \longmapsto 0. \end{array}$$

Hence,

$$h_p^{-L_2} = (\mathbf{c} \mapsto \mathbf{b}) + t \cdot \tilde{\Phi}^{-L_2}, \tag{16.6}$$

for some $t \in \mathbb{F}$.

Using the above observations, we are now able to compute the length 2 map of the link surgery hypercube of the Hopf link.

Lemma 16.4. *The length 3 hypercube relation in equation (16.4) is satisfiable with H_2 , but not with H_1 . Hence, H_2 is the length 2 map of the Hopf link surgery hypercube.*

Proof. For $i \in \{1, 2\}$, let C_i denote the sum of compositions of maps in (16.4) which would appear in the length 3 relation, except for $[\partial, \omega_p]$. Here, C_i is computed with the map $H = H_i$. The hypercube relations are satisfiable with H_i if and only if $C_i \simeq 0$.

We have determined the maps $h_p^{-L_2}$ and $h_p^{-L_1}$ up to a minor ambiguity, as shown in equations (16.5) and (16.6). Note that changing the value of s in $h_p^{-L_1}$ changes C_i

by addition of the map E from (16.3), which we already know to be null-homotopic. Changing t by adding 1 has the same effect. Hence, it is sufficient to consider the case where $s = t = 0$.

One easily computes

$$C_1 = [h_p^{-L_2}, \tilde{\Phi}^{-L_1}] + [H_1, A_p] = \begin{array}{l} \mathbf{a} \longmapsto \mathbf{d} \\ \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathbf{b} + \mathcal{U}_2\mathbf{c} \\ \mathbf{d} \longmapsto \mathbf{d}, \end{array}$$

which is not null-homotopic. Hence, H_1 is not a valid choice. It is easily checked that in fact

$$[h_p^{-L_2}, \tilde{\Phi}^{-L_1}] + [H_2, A_p] = 0.$$

Hence, the diagonal map must be H_2 . ■

Remark 16.5. View the maps H_1 and H_2 as being $\mathbb{F}[U_1, U_2]$ -equivariant, where U_i acts by \mathcal{V}_i on the domain, and \mathcal{U}_i on the codomain. The maps H_1 and H_2 are homotopic over $\mathbb{F}[U_1]$ or $\mathbb{F}[U_2]$, but not over $\mathbb{F}[U_1, U_2]$. If we wanted to compute the Heegaard Floer homology of surgeries on the Hopf link, either H_1 or H_2 would give the same answer. Compare [30, Section 5.2]. We will see in Section 16.3 that using H_1 to build the link surgery hypercube gives $\mathcal{C}_\Lambda(H, \mathcal{A}_{\alpha\beta})$ where $\mathcal{A}_{\alpha\beta}$ is a system of arcs where one arc is alpha-parallel and the other is beta-parallel.

Since we will be taking tensor products, we need to know the length 2 map over the ring $\mathbb{F}[U_1, U_2]$, and there are algebraically two distinct models.

It remains to describe the proof of Proposition 16.1.

Proof of Proposition 16.1. The proof is obtained by filling in powers of \mathcal{V}_i in the maps from Lemma 16.2 according to Alexander grading changes. We leave this mostly to the reader, but we verify the claim for the map $\Phi_{H,L_2}^{-L_1}$. This map sends Alexander grading \mathbf{s} to $\mathbf{s} + (\lambda_1, -1)$. Furthermore, it is $\mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]$ -equivariant, and is T -equivariant with respect to $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1]$. As an example, the equivariance properties and the computation of the quotient map in Lemma 16.2 constrain $\Psi_{H,L_2}^{-L_1}(\mathbf{a}) = \mathcal{V}_1^N \mathbf{d}$ for some N . Since $A(\mathbf{a}) = (-\frac{1}{2}, \frac{1}{2})$ and $A(\mathbf{d}) = (\frac{1}{2}, -\frac{1}{2})$, we conclude that $N = \lambda_1 - 1$. Similar arguments hold for all of the other components. ■

16.2 The Hopf link complex as a type-D module

We now describe the type-D module over $\mathcal{L}_2 = \mathcal{K} \otimes \mathcal{K}$ for the Hopf link complex. We write $\mathcal{H}_\Lambda^{\mathcal{K} \otimes \mathcal{K}}$ for the type-D module for the negative Hopf link with an alpha-parallel system of arcs. We recall that by definition, for each $\varepsilon \in \mathbb{E}_2$, we define $\mathcal{H} \cdot \mathbf{I}_\varepsilon$

to have the same generators over \mathbb{F} as a free $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{U}_2, \mathcal{V}_2]$ -basis of $\mathcal{CFL}(H)$. Hence, we write

$$\mathcal{H}_\Lambda^{\mathcal{K} \otimes \mathcal{K}} \cdot \mathbf{I}_\varepsilon = \text{Span}_{\mathbb{F}}(\mathbf{a}_\varepsilon, \mathbf{b}_\varepsilon, \mathbf{c}_\varepsilon, \mathbf{d}_\varepsilon).$$

It is helpful to write \mathcal{H}_ε for $\mathcal{H}_\Lambda \cdot \mathbf{I}_\varepsilon$. With this notation, we may decompose \mathcal{H}_Λ as

$$\mathcal{H}_\Lambda = \begin{array}{cc} \mathcal{H}_{0,0} & \mathcal{H}_{1,0} \\ \mathcal{H}_{0,1} & \mathcal{H}_{1,1} \end{array}$$

We now describe the structure maps.

Lemma 16.6. *Let $\mathcal{H}_\Lambda^{\mathcal{K} \otimes \mathcal{K}}$ be the type-D module associated to the negative Hopf link with an alpha-parallel system of arcs. The type-D module $\mathcal{H}_\Lambda^{\mathcal{K} \otimes \mathcal{K}}$ is as follows:*

$$\begin{array}{ccccc} \partial^1 \curvearrowright & \mathcal{H}_{0,0} & \xrightarrow{L_1P^1 + ^{-}L_1P^1} & \mathcal{H}_{1,0} & \curvearrowleft \partial^1 \\ & \downarrow & \dashrightarrow & \downarrow & \\ & L_2F^1 + ^{-}L_2F^1 & -H \Omega^1 & L_2G^1 + ^{-}L_2G^1 & \\ & \downarrow & & \downarrow & \\ \partial^1 \curvearrowright & \mathcal{H}_{0,1} & \xrightarrow{L_1Q^1 + ^{-}L_1Q^1} & \mathcal{H}_{1,1} & \curvearrowleft \partial^1 \end{array}$$

where the maps are as follows:

- (1) The map ∂^1 is the ordinary differential of the Hopf link complex. See equation (16.1).
- (2) The maps denoted L_1P^1 , $^{-}L_1P^1$, L_1Q^1 and $^{-}L_1Q^1$ correspond to the length 1 surgery maps for $\pm L_1$. The maps L_1P^1 and L_1Q^1 share the same formula as each other, and the maps $^{-}L_1P^1$ and $^{-}L_1Q^1$ share the same formula as each other. They are given by the following formulas:

$$(a) \quad \begin{array}{l} \mathbf{a} \longmapsto \mathbf{a} \otimes \sigma_1 \\ L_1P^1 = \mathbf{b} \longmapsto \mathbf{b} \otimes \sigma_1 \\ \mathbf{c} \longmapsto \mathbf{c} \otimes \sigma_1 \\ \mathbf{d} \longmapsto \mathbf{d} \otimes \sigma_1. \end{array}$$

$$(b) \quad \begin{array}{l} \mathbf{a} \longmapsto \mathbf{d} \otimes \mathcal{V}_1^{\lambda_1 - 1} \tau_1 \\ ^{-}L_1P^1 = \mathbf{b} \longmapsto 0 \\ \mathbf{c} \longmapsto \mathbf{b} \otimes \mathcal{V}_1^{\lambda_1 + 1} \tau_1 + \mathbf{c} \otimes \mathcal{U}_2 \mathcal{V}_1^{\lambda_1} \tau_1 \\ \mathbf{d} \longmapsto \mathbf{d} \otimes \mathcal{U}_2 \mathcal{V}_1^{\lambda_1} \tau_1. \end{array}$$

(3) The maps L_2F^1 and $-L_2F^1$ are the surgery maps for the component L_2 , with L_1 unreduced. They are given by the following formulas:

$$\begin{aligned}
 \text{(a)} \quad & \mathbf{a} \longmapsto \mathbf{a} \otimes \sigma_2 \\
 L_2F^1 = & \mathbf{b} \longmapsto \mathbf{b} \otimes \sigma_2 \\
 & \mathbf{c} \longmapsto \mathbf{c} \otimes \sigma_2 \\
 & \mathbf{d} \longmapsto \mathbf{d} \otimes \sigma_2. \\
 \\
 \text{(b)} \quad & \mathbf{a} \longmapsto \mathbf{a} \otimes \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \tau_2 \\
 -L_2F^1 = & \mathbf{b} \longmapsto 0 \\
 & \mathbf{c} \longmapsto \mathbf{b} \otimes \mathcal{V}_2^{\lambda_2+1} \tau_2 + \mathbf{c} \otimes \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \tau_2 \\
 & \mathbf{d} \longmapsto \mathbf{a} \otimes \mathcal{V}_2^{\lambda_2-1} \tau_2.
 \end{aligned}$$

The maps L_2G^1 and $-L_2G^1$ are the surgery maps L_2 , with L_1 already reduced. They are given by the same formulas.

(4) The map $-H\Omega^1$ is the length 2 surgery map for $-H = -L_1 \cup -L_2$. This map takes the following form:

$$\begin{aligned}
 -H\Omega^1 = & \mathbf{a} \longmapsto \mathbf{c} \otimes \mathcal{V}_1^{\lambda_1-2} \mathcal{V}_2^{\lambda_2-1} \tau_1 \tau_2 \\
 & \mathbf{b} \longmapsto 0 \\
 & \mathbf{c} \longmapsto \mathbf{d} \otimes \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2} \tau_1 \tau_2 \\
 & \mathbf{d} \longmapsto \mathbf{c} \otimes \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2-2} \tau_1 \tau_2.
 \end{aligned}$$

We leave the verification of the above lemma to the reader as it is mostly a restatement of Proposition 16.1 using the notation from Section 8.6.

16.3 The Hopf link with arc system $\mathcal{A}_{\alpha\beta}$

In this section, we compute the type-D module of the Hopf link complex when we use an arc system such that L_1 is alpha-parallel, and L_2 is beta-parallel. We write $\overline{\mathcal{H}}_{\Lambda}^{\mathcal{K} \otimes \mathcal{K}}$ for the associated type-D module.

Proposition 16.7. *The type-D module $\overline{\mathcal{H}}_{\Lambda}^{\mathcal{K} \otimes \mathcal{K}}$ is identical to the module $\mathcal{H}_{\Lambda}^{\mathcal{K} \otimes \mathcal{K}}$, except that we delete the term $\mathbf{c} \mapsto \mathbf{d} \otimes \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2} \tau_1 \tau_2$ from δ^1 .*


Proof. We may use Theorem 14.1 and the transformer bimodule to change one of the arcs from alpha-parallel to beta-parallel. Accordingly, we have that

$$\overline{\mathcal{H}}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2} = \mathcal{H}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2} \boxtimes \mathcal{K}_2 \mathcal{T} \mathcal{K}_2.$$

We claim that the above tensor product description of $\overline{\mathcal{H}}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2}$ gives exactly the module structure described in the statement (i.e., with structure maps as in Lemma 16.6, except with the additional length 2 term $\mathbf{c} \mapsto \mathbf{d} \otimes \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2} \tau_1 \tau_2$). This is verified by a direct computation. We compute $\mathcal{H}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2} \widehat{\boxtimes} \mathcal{K}_2 \mathcal{T}^{\mathcal{K}_2}$. The module $\mathcal{K}_2 \mathcal{T}^{\mathcal{K}_2}$ has a structure map term of the form $\delta_2^1(a, 1) = 1 \otimes a$, for all a . These contribute all of the differentials from $\mathcal{H}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2}$ to $\overline{\mathcal{H}}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2}$. Additionally, $\mathcal{K}_2 \mathcal{T}^{\mathcal{K}_2}$ has a δ_4^1 term, which is given by

$$\delta_4^1(a, b, c\tau, i_0) = i_1 \otimes \partial_{\mathcal{U}}(a) \partial_{\mathcal{V}}(b) c\tau.$$

The structure map diagram which adds the extra length 2 differential is the following:



In the above, Φ^{-L_2} , ∂e and Φ^{-L_1} denote specific summands of the structure map δ^1 on $\mathcal{H}_{\Lambda}^{\mathcal{K}_1 \otimes \mathcal{K}_2}$. The specific sequence of generators contributing to the above diagram is the following:

$$\begin{aligned} \mathbf{c}_{00} &\xrightarrow{\Phi^{-L_2}} \mathbf{c}_{01} \otimes \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \tau_2 \xrightarrow{\partial e} \mathbf{d}_{01} \otimes \mathcal{V}_2 \otimes \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \tau_2 \\ &\xrightarrow{\Phi^{-L_1}} \mathbf{d}_{11} \otimes \mathcal{U}_2 \mathcal{V}_1^{\lambda_1} \tau_1 \otimes \mathcal{V}_2 \otimes \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \tau_2. \end{aligned}$$

(Here, we omit all terms which evaluate to zero under $\mathbb{I} \otimes \delta_4^1$.) Applying δ_4^1 , we obtain the extra term

$$\mathbf{c}_{00} \mapsto \mathbf{d}_{11} \otimes \mathcal{V}_1^{\lambda_1} \tau_1 \mathcal{U}_1 \mathcal{V}_2^{\lambda_2} \tau_2 = \mathbf{d}_{11} \otimes \mathcal{V}_1^{\lambda_1-1} \mathcal{V}_2^{\lambda_2} \tau_1 \tau_2,$$

as claimed in the statement. Note that when evaluating δ_2^1 and δ_4^1 in the tensor product, we are using the extension of scalars construction from Section 3.5 to view $\mathcal{K}_1 \mathcal{T}^{\mathcal{K}_1}$ as a DA bimodule over $(\mathcal{K}_1 \otimes \mathcal{K}_2, \mathcal{K}_1 \otimes \mathcal{K}_2)$. ■

Remark 16.8. In our proof of Proposition 16.1, we observed that all of the maps in the Hopf link complex were uniquely determined except for the length 2 map, and there were two choices, up to overall homotopy. We wrote H_1 and H_2 for these maps. We proved that H_2 was the correct choice when the system of arcs was alpha-parallel. We observe that the above model of $\overline{\mathcal{H}}_{\Lambda}^{\mathcal{L}_2}$ corresponds exactly to the other choice of homotopy.

Chapter 17

Minimal models for the Hopf link surgery complex

In this chapter, we consider the link surgery complex for the Hopf link as a DA bimodule. We do this via a tensor product

$${}_{\mathcal{K}}\mathcal{H}_{\Lambda}^{\mathcal{K}} := \mathcal{H}_{\Lambda}^{\mathcal{K} \otimes \mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}^{\otimes 2}].$$

In the above, just one algebra output of $\mathcal{H}_{\Lambda}^{\mathcal{K} \otimes \mathcal{K}}$ is input into ${}_{\mathcal{K}|\mathcal{K}}[\mathbb{I}^{\otimes 2}]$.

We note that the number of generators in the complex $\mathcal{H}_{\Lambda}^{\mathcal{K} \otimes \mathcal{K}}$ can be reduced by homotopy equivalence. To this end, we define the maximal ideal $\mathfrak{m} \subseteq \mathcal{K}$, to be the one which is generated by $\mathcal{U}, \mathcal{V} \in \mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$, $\sigma, \tau \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_0$ and $\mathcal{U} \in \mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$. We make the following definition (compare [6, page 476]).

Definition 17.1. We say that a type-DA bimodule ${}_{\mathcal{K}}\mathcal{X}^{\mathcal{K}}$ is *minimal* if $\delta_1^1: \mathcal{X} \rightarrow \mathcal{X} \otimes \mathcal{K}/\mathfrak{m}$ is zero.

In the Heegaard Floer literature, one frequently uses the term *reduced* instead of minimal. In this chapter, we will explore a minimal model of the Hopf link complex ${}_{\mathcal{K}}\mathcal{H}_{\Lambda}^{\mathcal{K}}$, for which we will write ${}_{\mathcal{K}}\mathcal{Z}^{\mathcal{K}}$. We will also compute a minimal model of the complex ${}_{\mathcal{K}}\overline{\mathcal{H}}_{\Lambda}^{\mathcal{K}}$ obtained by using both alpha and beta-parallel arcs on the Hopf link.

17.1 DA-perspectives on the Hopf link complex

In this section, we describe how to view the Hopf link complex as a type-DA module, and we compute a minimal model. This perspective was first explored in [14].

Write $H = L_1 \cup L_2$ for the Hopf link. Write $P_1 = \mathbb{F}[\mathcal{U}_1, \mathcal{V}_1]$ and $P_2 = \mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]$, where P_i contains the variables for L_i . Let us write \mathcal{C} for $\mathcal{CFL}(H)$.

We may view \mathcal{C} as a type-D module over P_2 . We do this by viewing \mathcal{C}^{P_2} as freely generated by $\mathbf{a}, \mathbf{b}, \mathbf{c}$ and \mathbf{d} over P_1 . The structure map δ^1 is shown below:

$$\mathcal{C}^{P_2} \cong \begin{array}{ccc} \mathbf{a}[\mathcal{U}_1, \mathcal{V}_1] \leftarrow 1|\mathcal{U}_2 & \text{--} & \mathbf{b}[\mathcal{U}_1, \mathcal{V}_1] \\ \uparrow \text{---} \mathcal{V}_1|1 & & \downarrow \text{---} \mathcal{U}_1|1 \\ \mathbf{c}[\mathcal{U}_1, \mathcal{V}_1] \text{---} 1|\mathcal{V}_2 & \rightarrow & \mathbf{d}[\mathcal{U}_1, \mathcal{V}_1] \end{array}$$

An arrow from \mathbf{x} to \mathbf{y} labeled by $a|b$ indicates that $\delta^1(\mathbf{x})$ has a summand of $(a \cdot \mathbf{y}) \otimes b$.

As a type-D module, \mathcal{C}^{P_2} is not minimal. We now describe a homotopy equivalent complex which is minimal. It is helpful to realize that the above description of \mathcal{C}^{P_2}

may be further decomposed as a box tensor product, as follows. Let \mathcal{E}_2 denote the exterior algebra on two generators,

$$\mathcal{E}_2 := \Lambda^*(\phi_2, \psi_2).$$

We consider the following type-DD bimodule ${}^{\mathcal{E}_2}\mathcal{K}^{P_2}$. As a vector space, $\mathcal{K} \cong \mathbb{F}$. The type-DD structure map is given by the formula

$$\delta^{1,1}(1) = \phi_2|1|\mathcal{U}_2 + \psi_2|1|\mathcal{V}_2.$$

We may define a type-A module $\mathcal{C}_{\mathcal{E}_2}$, generated freely over P_1 by **a**, **b**, **c** and **d**, with module structure given by the following diagram:

$$\mathcal{C}_{\mathcal{E}_2} = \begin{array}{ccc} \mathbf{a}[\mathcal{U}_1, \mathcal{V}_1] & \leftarrow \phi_2 & \mathbf{b}[\mathcal{U}_1, \mathcal{V}_1] \\ \uparrow \text{---} \nu_1 & & \text{---} \mathcal{U}_1 \downarrow \\ \mathbf{c}[\mathcal{U}_1, \mathcal{V}_1] & \xrightarrow{\psi_2} & \mathbf{d}[\mathcal{U}_1, \mathcal{V}_1] \end{array}$$

In the above, solid arrows denote the m_2 actions of \mathcal{E}_2 , while the dashed arrows denote the m_1 actions (i.e., internal differentials). To illustrate the notation, if $f \in \mathbb{F}[\mathcal{U}_1, \mathcal{V}_1]$, then

$$m_2(f \cdot \mathbf{b}, \phi_2) = f \cdot \mathbf{a} \quad \text{and} \quad m_1(f \cdot \mathbf{b}) = \mathcal{U}_1 \cdot f \cdot \mathbf{d}.$$

The action of ψ_2 is similar.

Clearly, there is an isomorphism of type-D modules

$$\mathcal{C}^{P_2} \cong \mathcal{C}_{\mathcal{E}_2} \boxtimes {}^{\mathcal{E}_2}\mathcal{K}^{P_2}.$$

In the above, we have forgotten about the natural action of P_1 , given by ordinary multiplication. In fact, we may incorporate additionally this action to obtain bimodules ${}_{P_1}\mathcal{C}^{P_2}$ and ${}_{P_1}\mathcal{C}_{\mathcal{E}_2}$, such that

$${}_{P_1}\mathcal{C}^{P_2} \cong {}_{P_1}\mathcal{C}_{\mathcal{E}_2} \boxtimes {}^{\mathcal{E}_2}\mathcal{K}^{P_2}.$$

The module ${}_{P_1}\mathcal{C}^{P_2}$ has only δ_1^1 and δ_2^1 non-trivial, and ${}_{P_1}\mathcal{C}_{\mathcal{E}_2}$ has only $m_{0,1,0}$, $m_{1,1,0}$ and $m_{0,1,1}$ non-trivial. In the next section, we explore minimal models.

17.2 Minimal models for the Hopf link Floer complex

We apply the homological perturbation lemma to obtain minimal models of the Hopf link bimodules from the previous section. We will define a DA bimodule ${}_{P_1}Z^{P_2}$ which is homotopy equivalent to ${}_{P_1}\mathcal{C}^{P_2}$ and which has $\delta_1^1 = 0$. The techniques of this section formalize the construction in [14].

First, define the chain complex \mathcal{C}_0 by forgetting about the \mathcal{E}_2 -action on $\mathcal{C}_{\mathcal{E}_2}$. Over \mathbb{F}_2 , \mathcal{C}_0 is homotopy equivalent to its homology, which we view as the vector space

$$Z := \mathbf{a}[\mathcal{U}_1] \oplus \mathbf{d}[\mathcal{V}_1].$$

Here, $\mathbf{a}[\mathcal{U}_1]$ denotes a copy of $\mathbb{F}[\mathcal{U}_1]$, generated by \mathbf{a} , and similarly for $\mathbf{d}[\mathcal{V}_1]$. There are natural maps

$$i: Z \rightarrow \mathcal{C}_0, \quad \pi: \mathcal{C}_0 \rightarrow Z, \quad \text{and} \quad h: \mathcal{C}_0 \rightarrow \mathcal{C}_0$$

such that i and π are chain maps, $i \circ \pi = \text{id} + [m_1, h]$, $\pi \circ i = \text{id}$, $h \circ i = 0$, $\pi \circ h = 0$ and $h \circ h = 0$. The map i is given by

$$i(\mathcal{U}_1^n \mathbf{a}) = \mathcal{U}_1^n \mathbf{a} \quad \text{and} \quad i(\mathcal{V}_1^n \mathbf{d}) = \mathcal{V}_1^n \mathbf{d}.$$

The map π is given by

$$\pi(\mathcal{U}_1^n \mathcal{V}_1^m \mathbf{a}) = \begin{cases} \mathcal{U}_1^n \mathbf{a} & \text{if } m = 0, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad \pi(\mathcal{U}_1^n \mathcal{V}_1^m \mathbf{d}) = \begin{cases} \mathcal{V}_1^m \mathbf{d} & \text{if } n = 0, \\ 0 & \text{otherwise.} \end{cases}$$

The map h is given by

$$h(\mathcal{U}_1^n \mathcal{V}_1^m \mathbf{a}) = \begin{cases} \mathcal{U}_1^n \mathcal{V}_1^{m-1} \mathbf{c} & \text{if } m \geq 1, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad h(\mathcal{U}_1^n \mathcal{V}_1^m \mathbf{d}) = \begin{cases} \mathcal{U}_1^{n-1} \mathcal{V}_1^m \mathbf{b} & \text{if } n \geq 1, \\ 0 & \text{otherwise.} \end{cases}$$

The homological perturbation lemma for A_∞ -modules, Lemma 4.1, endows Z with the structure of a right A_∞ -module over the exterior algebra \mathcal{E}_2 , for which we write $Z_{\mathcal{E}_2}$, which is A_∞ -homotopy equivalent to $\mathcal{C}_{\mathcal{E}_2}$. In fact, the homotopy equivalence is given explicitly by the homological perturbation lemma. The maps i , π and h extend to A_∞ -module morphisms i_* , π_* and h_* . We box these morphisms with the identity map on ${}^{\mathcal{E}_2}\mathcal{K}^{P_2}$ to obtain maps of type-D modules

$$\Pi^1: \mathcal{C}^{P_2} \rightarrow Z^{P_2}, \quad I^1: Z^{P_2} \rightarrow \mathcal{C}^{P_2}, \quad \text{and} \quad H^1: \mathcal{C}^{P_2} \rightarrow \mathcal{C}^{P_2}.$$

Lemma 17.2. *The morphisms I^1 and Π^1 are type-D homomorphisms (i.e., $\partial_{\text{Mor}}(I^1) = 0$ and $\partial_{\text{Mor}}(\Pi^1) = 0$). Furthermore,*

- (1) $\Pi^1 \circ I^1 = \text{id}$.
- (2) $H^1 \circ H^1 = 0$.
- (3) $H^1 \circ I^1 = 0$.
- (4) $\Pi^1 \circ H^1 = 0$.
- (5) $I^1 \circ \Pi^1 = \text{id} + \partial_{\text{Mor}}(H^1)$.

Proof. Lemma 4.1 (the homological perturbation lemma) implies that the stated formulas hold for i_* , h_* and π_* . Boxing with the identity map preserves these relations since the algebra is \mathcal{E}_2 (which is an associative algebra), so

$$(f_* \boxtimes \mathbb{I}_{\mathcal{K}}) \circ (g_* \boxtimes \mathbb{I}_{\mathcal{K}}) = ((f_* \circ g_*) \boxtimes \mathbb{I}_{\mathcal{K}}),$$

by [25, Remark 2.2.28]. Furthermore, boxing with the identity is a chain map

$$\partial_{\text{Mor}}(f_* \boxtimes \mathbb{I}_{\mathcal{K}}) = (\partial_{\text{Mor}}(f_*) \boxtimes \mathbb{I}_{\mathcal{K}}).$$

See [25, Lemma 2.3.3 (1)]. ■

It is enlightening to compute concrete formulas for the morphisms Π^1 , H^1 and I^1 .

Lemma 17.3. *The type-D module maps I^1 , Π^1 and H^1 are given by the following formulas:*

- (1) I^1 is given by $I^1(\mathbf{x}) = \mathbf{x} \otimes 1$, for $\mathbf{x} \in Z$.
- (2) Π^1 vanishes on $\text{Span}(\mathbf{b}, \mathbf{c})$, and satisfies

$$\begin{aligned} \Pi^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}) &= \begin{cases} \mathcal{U}_1^{i-j} \mathbf{a} \otimes \mathcal{U}_2^j \mathcal{V}_2^j & \text{if } i \geq j, \\ \mathcal{V}_1^{j-i-1} \mathbf{d} \otimes \mathcal{U}_2^i \mathcal{V}_2^{i+1} & \text{if } i < j, \end{cases} \\ \Pi^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{d}) &= \begin{cases} \mathcal{U}_1^{i-j-1} \mathbf{a} \otimes \mathcal{U}_2^{j+1} \mathcal{V}_2^j & \text{if } i > j, \\ \mathcal{V}_1^{j-i} \mathbf{d} \otimes \mathcal{U}_2^i \mathcal{V}_2^i & \text{if } i \leq j. \end{cases} \end{aligned}$$

- (3) H^1 vanishes on $\text{Span}(\mathbf{b}, \mathbf{c})$, and maps $\text{Span}(\mathbf{a}, \mathbf{d})$ to $\text{Span}(\mathbf{b}, \mathbf{c}) \otimes P_2$. It is given by the formula

$$H^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}) = \begin{cases} \mathcal{U}_1^i \mathcal{V}_1^{j-1} \mathbf{c} \otimes 1 + \mathcal{U}_1^{i-1} \mathcal{V}_1^{j-1} \mathbf{b} \otimes \mathcal{V}_2 + \cdots \\ \quad + \mathcal{U}_1^{i-j} \mathbf{b} \otimes \mathcal{U}_2^{j-1} \mathcal{V}_2^j & \text{if } j \leq i, \\ \mathcal{U}_1^i \mathcal{V}_1^{j-1} \mathbf{c} \otimes 1 + \mathcal{U}_1^{i-1} \mathcal{V}_1^{j-1} \mathbf{b} \otimes \mathcal{V}_2 + \cdots \\ \quad + \mathcal{V}_1^{j-i-1} \mathbf{c} \otimes \mathcal{U}_2^i \mathcal{V}_2^i & \text{if } j > i, \end{cases}$$

and

$$H^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{d}) = \begin{cases} \mathcal{U}_1^{i-1} \mathcal{V}_1^j \mathbf{b} \otimes 1 + \mathcal{U}_1^{i-1} \mathcal{V}_1^{j-1} \mathbf{c} \otimes \mathcal{U}_2 + \cdots \\ \quad + \mathcal{U}_1^{i-j-1} \mathbf{b} \otimes \mathcal{U}_2^j \mathcal{V}_2^j & \text{if } i > j \\ \mathcal{U}_1^{i-1} \mathcal{V}_1^j \mathbf{b} \otimes 1 + \mathcal{U}_1^{i-1} \mathcal{V}_1^{j-1} \mathbf{c} \otimes \mathcal{U}_2 + \cdots \\ \quad + \mathcal{V}_1^{j-i} \mathbf{c} \otimes \mathcal{U}_2^i \mathcal{V}_2^{i-1} & \text{if } i \leq j. \end{cases}$$

Proof. Lemma 4.1 gives a concrete formula for the A_∞ -action on $Z_{\mathcal{E}_2}$. We will prove the first statement, and also compute $\Pi^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a})$ when $i \geq j$ in order to illustrate the technique. We will leave the remaining cases to the reader.

Firstly, we begin with I^1 . We are boxing with $\mathbb{I}_{\mathcal{K}}$, so suppose that $a_1 | \cdots | a_n | 1 | b_n | \cdots | b_1$ is the output of repeated applications of $\delta^{1,1}$ on $\mathcal{E}_2 \mathcal{K}^{P_2}$. The homological perturbation lemma gives a recipe for $m_{n+1}^Z(\mathbf{x}, a_1, \dots, a_n)$. See Figure 4.1. The recipe is to include \mathbf{x} into \mathcal{C} via i . We then apply $m_2(-, a_1)$, then h , then $m_2(-, a_2)$, then h , and so forth, until one applies $m(-, a_n)$. Then we apply π . The algebra output is the product $b_n \cdots b_1$. The m_2 action of \mathcal{E}_2 on \mathcal{C} vanishes on the image of i , so there are no contributions unless $n = 0$. The formula follows.

We now compute $\Pi^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a})$ for $i \geq j$. We begin at $\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a} \in \mathcal{C}$. If $j = 0$, then $h(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}) = 0$, so our only option is to apply π , which gives $\mathcal{U}_1^i \mathbf{a}$ as claimed. If $j > 0$, the only option is to apply h to get $\mathcal{U}_1^i \mathcal{V}_1^{j-1} \mathbf{c}$. Then we apply $m_2(-, \psi_2)$ which gives $\mathcal{U}_1^i \mathcal{V}_1^{j-1} \mathbf{d} \otimes \mathcal{V}_2$. We then apply h and $m_2(-, \phi_2)$ and we get $\mathcal{U}_1^{i-1} \mathcal{V}_1^{j-1} \mathbf{a} \otimes \mathcal{U}_2 \mathcal{V}_2$. We repeat this procedure until we cannot apply h or m_2 anymore. In the case that $i \geq j$, the final term in this sequence will be $\mathcal{U}_1^{i-j} \mathbf{a} \otimes \mathcal{U}_2^j \mathcal{V}_2^j$. All of the remaining claims follow a similar analysis. ■

Lemma 17.2 allows us to apply the homological perturbation lemma of DA bimodules, Lemma 4.2, which equips Z with a DA bimodule structure ${}_{P_1} Z^{P_2}$, and supplies morphisms of DA bimodules

$$\Pi_*^1: {}_{P_1} \mathcal{C}^{P_2} \rightarrow {}_{P_1} Z^{P_2}, \quad I_*^1: {}_{P_1} Z^{P_2} \rightarrow {}_{P_1} \mathcal{C}^{P_2}, \quad \text{and} \quad H_*^1: {}_{P_1} \mathcal{C}^{P_2} \rightarrow {}_{P_1} \mathcal{C}^{P_2}$$

which satisfy relations identical to Lemma 17.2.

Lemma 17.4. *The bimodule ${}_{P_1} Z^{P_2}$ and the morphisms Π_*^1 , I_*^1 and H_*^1 satisfy the following relations:*

- (1) *The structure maps δ_j^1 on ${}_{P_1} Z^{P_2}$ vanish if $j \neq 2$.*
- (2) *The maps Π_j^1 and H_j^1 vanish unless $j = 1$.*
- (3) *The map I_j^1 vanishes if $j > 2$.*
- (4) *The maps δ_*^1 and I_*^1 are strictly unital, i.e., they vanish if 1 is an input, except for δ_2^1 , which satisfies $\delta_2^1(1 \otimes x) = x \otimes 1$. (Note that H_j^1 and Π_j^1 are trivially unital, since they are only non-trivial if $j = 1$.)*

Proof. The proofs of all statements are by explicit examination of the maps from the homological perturbation lemma.

We begin with the statements about δ_j^1 . Consider first a sequence of algebra elements (a_n, \dots, a_1) in P_1 . The rule for computing δ_j^1 is to input via I^1 , then apply m_2 , then we apply pairs of H^1 followed by m_2 until we exhaust (a_n, \dots, a_1) , and then finally we apply Π_1 . The algebra elements outputs are multiplied together by applying μ_2 repeatedly. However, applying m_2 does not move an elements position in \mathcal{C} (i.e., \mathbf{a} , \mathbf{b} , \mathbf{c} or \mathbf{d}). The map H^1 does change the generator, and it maps $\text{Span}(\mathbf{a}, \mathbf{d})$ to $\text{Span}(\mathbf{c}, \mathbf{b}) \otimes P_2$, and vanishes on $\text{Span}(\mathbf{c}, \mathbf{b})$. If we apply another m_2 , we remain in $\text{Span}(\mathbf{c}, \mathbf{b})$. We note that Π^1 and H^1 both vanish on $\text{Span}(\mathbf{c}, \mathbf{b})$. In particular, the

only terms making non-trivial contribution are those with one δ_2^1 , and no H^1 term. The same argument shows that Π_k^1 and H_k^1 vanish if $k > 1$.

To compute the map $I_{j+1}^1(a_j, \dots, a_1, \mathbf{x})$, the recipe is to first include \mathbf{x} into $\mathcal{C} \otimes P_2$ via I^1 . This includes \mathbf{x} into $\text{Span}(\mathbf{a}, \mathbf{d}) \otimes P_2$. We then apply $\delta_2^1(a_1, -)$. This preserves $\text{Span}(\mathbf{a}, \mathbf{d}) \otimes P_2$. Next we apply H^1 , which maps $\text{Span}(\mathbf{a}, \mathbf{d}) \otimes P_2$ to $\text{Span}(\mathbf{b}, \mathbf{c}) \otimes P_2$. Any further applications of $\delta_2^1(a_n, -)$ followed by H^1 would map to zero. Hence we may have I_2^1 , but not I_j^1 for $j > 2$.

The final statement about being strictly unital follows from the homological perturbation lemma. \blacksquare

It is helpful to explicitly compute δ_2^1 on $P_1 Z^{P_2}$. Since there is no δ_j^1 for $j > 2$ by Lemma 17.4, we compute only $\delta_2^1(\mathcal{U}_1, -)$ and $\delta_2^1(\mathcal{V}_1, -)$.

Lemma 17.5. *The map δ_2^1 on $P_1 Z^{P_2}$ satisfies the following:*

- (1) $\delta_2^1(\mathcal{U}_1, \mathcal{U}_1^n \mathbf{a}) = \mathcal{U}_1^{n+1} \mathbf{a} \otimes 1.$
- (2) $\delta_2^1(\mathcal{V}_1, \mathcal{U}_1^n \mathbf{a}) = \begin{cases} \mathcal{U}_1^{n-1} \mathbf{a} \otimes \mathcal{U}_2 \mathcal{V}_2 & \text{if } n > 0, \\ \mathbf{d} \otimes \mathcal{V}_2 & \text{if } n = 0. \end{cases}$
- (3) $\delta_2^1(\mathcal{V}_1, \mathcal{V}_1^m \mathbf{d}) = \mathcal{V}_1^{m+1} \mathbf{d} \otimes 1.$
- (4) $\delta_2^1(\mathcal{U}_1, \mathcal{V}_1^m \mathbf{d}) = \begin{cases} \mathcal{V}_1^{m-1} \mathbf{d} \otimes \mathcal{U}_2 \mathcal{V}_2 & \text{if } m > 0, \\ \mathbf{a} \otimes \mathcal{U}_2 & \text{if } m = 0. \end{cases}$

Finally, the complex $\mathcal{V}_1^{-1} \mathcal{C}$ will also be important to understand. Similarly to \mathcal{C} , the complex $\mathcal{V}_1^{-1} \mathcal{C}$ can be reduced in size via a homotopy equivalence. We note that $\mathcal{V}_1^{-1} \mathcal{C}$ is the following DA bimodule:

$$\mathcal{V}_1^{-1} \mathcal{C} \cong \begin{array}{ccc} \mathbf{a}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{V}_1^{-1}] \leftarrow {}_1|\mathcal{U}_2 & - & \mathbf{b}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{V}_1^{-1}] \\ \uparrow \mathcal{V}_1|1 & & \downarrow \mathcal{U}_1|1 \\ \mathbf{c}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{V}_1^{-1}] \leftarrow {}_1|\mathcal{V}_2 & \rightarrow & \mathbf{d}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{V}_1^{-1}] \end{array}$$

In the above, all arrows denote δ_1^1 . The actions of δ_2^1 are given by ordinary polynomial multiplication. Since \mathcal{V}_1 is invertible, the arrow labeled $\mathcal{V}_1|1$ and the two generators \mathbf{a} and \mathbf{c} may be completely canceled from the complex. Hence, the above is homotopy equivalent to a DA bimodule whose underlying type-D structure is shown below:

$$\mathbf{b}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{V}_1^{-1}] \xrightarrow{\mathcal{U}_1|1} \mathbf{d}[\mathcal{U}_1, \mathcal{V}_1, \mathcal{V}_1^{-1}].$$

This complex above may be further reduced as a type-D structure, giving a minimal model, for which we write $\mathcal{V}_1^{-1} P_1 W^{P_2}$. As a vector space, $W \cong \mathbf{d}[\mathcal{V}_1, \mathcal{V}_1^{-1}]$. The same

homological perturbation argument as before equips W with the bimodule structure $\mathcal{V}_1^{-1}P_1W^{P_2}$ which has $\delta_j^1 = 0$ unless $j = 2$. For completeness, we record the actions.

Lemma 17.6. *The bimodule $\mathcal{V}_1^{-1}P_1W^{P_2}$ has the following action:*

$$\delta_2^1(\mathcal{U}_1^i\mathcal{V}_1^j, \mathcal{V}_1^n\mathbf{d}) = \mathcal{V}_1^{j+n-i}\mathbf{d} \otimes \mathcal{U}_2^i\mathcal{V}_2^i.$$

We leave the proofs of Lemmas 17.5 and 17.6 to the reader.

17.3 Minimal models over \mathcal{K}

In this section, we describe a minimal model of the DA bimodule $\mathcal{K}_1\mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$. We will denote the minimal model by

$$\mathcal{K}_1\mathcal{Z}_{(\lambda_1,0)}^{\mathcal{K}_2}.$$

The existence of the minimal model follows from the homological perturbation lemma for hypercubes of DA bimodules. From the description of $\mathcal{K}_1\mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$ it is clear that there is a filtration by the cube \mathbb{E}_2 . In Section 17.2, we described minimal models of $\mathcal{CF}\mathcal{L}(H)^{\mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]}$ and $\mathcal{V}_1^{-1}\mathcal{CF}\mathcal{L}(H)^{\mathbb{F}[\mathcal{U}_2, \mathcal{V}_2]}$. Note that we may identify the underlying type-D module $\mathcal{H}_\varepsilon^{\mathcal{K}_2}$ with one of these type-D modules, for $\varepsilon \in \mathbb{E}_2$. Hence, the construction from Section 17.2 may be viewed as giving morphisms of type-D structures

$$\Pi_\varepsilon^1: \mathcal{H}_\varepsilon^{\mathcal{K}_2} \rightarrow \mathcal{Z}_\varepsilon^{\mathcal{K}_2}, \quad I_\varepsilon^1: \mathcal{Z}_\varepsilon^{\mathcal{K}_2} \rightarrow \mathcal{H}_\varepsilon^{\mathcal{K}_2}, \quad \text{and} \quad H_\varepsilon^1: \mathcal{H}_\varepsilon^{\mathcal{K}_2} \rightarrow \mathcal{H}_\varepsilon^{\mathcal{K}_2}$$

which induce a homotopy equivalence of type-D structures, such that furthermore the algebraic assumptions of the homological perturbation lemma are satisfied. We obtain a homotopy equivalent hypercube of DA bimodules $\mathcal{K}_1\mathcal{Z}_{(\lambda_1,0)}^{\mathcal{K}_2}$ by applying the homological perturbation lemma for hypercubes of DA bimodules, Lemma 4.5. We now schematically sketch the induced type-D structure map:

$$\begin{array}{ccc} \mathcal{Z}_{0,0} & \mathcal{Z}_{1,0} & m_2^1 \xrightarrow{\quad} \mathcal{Z}_{0,0} \xrightarrow{-L_1p_2^1 + -L_1p_2^1} \mathcal{Z}_{1,0} \xleftarrow{\quad} m_2^1 \\ \delta_1^1 = \begin{array}{c} \downarrow \scriptstyle L_2f_1^1 + -L_2f_1^1 \quad \downarrow \scriptstyle L_2g_1^1 + -L_2g_1^1 \\ \mathcal{Z}_{0,1} & \mathcal{Z}_{1,1} \end{array} & \delta_2^1 = \begin{array}{c} \downarrow \scriptstyle -L_1q_2^1 + -L_1q_2^1 \\ \mathcal{Z}_{0,1} & \mathcal{Z}_{1,1} \end{array} \xleftarrow{\quad} m_2^1 \end{array}$$

$$\delta_3^1 = \begin{array}{ccc} \mathcal{Z}_{0,0} & & \mathcal{Z}_{1,0} \\ & \searrow \scriptstyle -H\omega_3^1 & \\ & & \mathcal{Z}_{1,1} \\ \mathcal{Z}_{0,1} & & \end{array}$$

Proposition 17.7. *Give the negative Hopf link framing $(\lambda_1, 0)$. The structure maps of the minimal model $\mathcal{X}_1 \mathcal{Z}_{(\lambda_1, 0)}^{\mathcal{K}_2}$ are as follows:*

- (1) *The maps m_2^1 are the same as the δ_2^1 maps in Lemmas 17.5 and 17.6.*
- (2) *The maps $L_1 p_2^1$ and $L_1 q_2^1$ are given by the same formulas as each other, as are $-L_1 p_2^1$ and $-L_1 q_2^1$. They are determined by the following formulas:*
 - (a) $L_1 p_2^1(\sigma_1, \mathcal{U}_1^i \mathbf{a}) = \mathcal{V}_1^{-i-1} \mathbf{d} \otimes \mathcal{U}_2^i \mathcal{V}_2^{i+1}$ and $L_1 p_2^1(\sigma_1, \mathcal{V}_1^j \mathbf{d}) = \mathcal{V}_1^j \mathbf{d} \otimes 1$.
 - (b) $-L_1 p_2^1(\tau_1, \mathcal{U}_1^i \mathbf{a}) = \mathcal{V}_1^{-i-1+\lambda_1} \mathbf{d} \otimes 1$ and $-L_1 p_2^1(\tau_1, \mathcal{V}_1^j \mathbf{d}) = \mathcal{V}_1^{j+\lambda_1} \mathbf{d} \otimes \mathcal{U}_2^{j+1} \mathcal{V}_2^j$.
 - (c) $L_1 p_2^1(\tau_1, -) = 0$ and $-L_1 p_2^1(\sigma_1, -) = 0$.
- (3) *The maps for $\pm L_2$ are as follows:*
 - (a) $L_2 f_1^1(\mathcal{U}_1^i \mathbf{a}) = \mathcal{U}_1^i \mathbf{a} \otimes \sigma_2$ and $L_2 f_1^1(\mathcal{V}_1^j \mathbf{d}) = \mathcal{V}_1^j \mathbf{d} \otimes \sigma_2$.
 - (b) $-L_2 f_1^1(\mathcal{U}_1^i \mathbf{a}) = \mathcal{U}_1^{i+1} \mathbf{a} \otimes \tau_2$ and

$$-L_2 f_1^1(\mathcal{V}_1^j \mathbf{d}) = \begin{cases} \mathbf{a} \otimes \mathcal{V}_2^{-1} \tau_2 & \text{if } j = 0, \\ \mathcal{V}_1^{j-1} \mathbf{d} \otimes \tau_2 & \text{if } j > 0. \end{cases}$$

- (c) $L_2 g_1^1(\mathcal{V}_1^i \mathbf{d}) = \mathcal{V}_1^i \mathbf{d} \otimes \sigma_2$.
- (d) $-L_2 g_1^1(\mathcal{V}_1^i \mathbf{d}) = \mathcal{V}_1^{i-1} \mathbf{d} \otimes \tau_2$.
- (4) *The map $-^H \omega_3^1$ is determined by the relations*

$$\begin{aligned} -^H \omega_3^1(\tau_1, \mathcal{V}_1^m, \mathcal{U}_1^i \mathbf{a}) &= \min(i+1, m) \mathcal{V}_1^{\lambda_1+m-i-2} \mathbf{d} \otimes \mathcal{U}_2^{m-1} \mathcal{V}_2^{m-1} \tau_2, \\ -^H \omega_3^1(\tau_1, \mathcal{U}_1^i, \mathcal{U}_1^i \mathbf{a}) &= 0, \\ -^H \omega_3^1(\tau_1, \mathcal{U}_1^n, \mathcal{V}_1^j \mathbf{d}) &= \min(n, j) \mathcal{V}_1^{j-n+\lambda_1-1} \mathbf{d} \otimes \mathcal{U}_2^{j-1} \mathcal{V}_2^{j-2} \tau_2, \\ -^H \omega_3^1(\tau_1, \mathcal{V}_1^m, \mathcal{V}_1^j \mathbf{d}) &= 0, \end{aligned}$$

and that $-^H \omega_3^1$ vanishes if an algebra input is a multiple of σ_1 . The map $-^H \omega_3^1$ also vanishes on pairs of algebra elements with other configurations of idempotents.

Remark 17.8. We have not enumerated a complete list of the structure maps. Additional powers of \mathcal{U}_1 and \mathcal{V}_1 may be added to arguments in the above maps. There are no terms of δ_{j+1}^1 for $j > 2$, however. A more minimal list could also have been made by specifying the length 1 maps on only \mathbf{a} and \mathbf{d} . As an example of several relations which are forced by the DA bimodule relations, we have


$$-^H \omega_3^1(\tau_1, \mathcal{U}_1^i \mathcal{V}_1^j, \mathcal{U}_1^n \mathbf{a}) = -^H \omega_3^1(\tau, \mathcal{V}_1^j, \mathcal{U}_1^{i+n} \mathbf{a})$$

and

$$-^H \omega_3^1(\tau_1, \mathcal{U}_1^i \mathcal{V}_1^j, \mathcal{V}_1^n \mathbf{d}) = -^H \omega_3^1(\tau, \mathcal{U}_1^i, \mathcal{V}_1^{j+n} \mathbf{d}).$$

Proof. All of these computations are performed algorithmically using the homological perturbation lemma for hypercubes of DA bimodules, Lemma 4.5. For the first two sets of equations the computation is essentially straightforward. We first apply the map I_ε^1 , then we apply a length 1 map of the cube (either with no algebra input, as for the maps labeled f_1^1 and g_1^1 , or with an algebra input of σ_1 or τ_1 , as for the maps p_2^1 and q_2^1). We leave these computations to the reader, as they are straightforward.

The map ${}^{-H}\omega_3^1$ is more interesting. There is exactly one configuration of a morphism graph which gives a non-trivial evaluation. This occurs for elements $\mathbf{x}_{0,0}$ in the idempotent $\varepsilon = (0, 0)$. The structure graph is shown below:



In the above, the map ${}^{-H}\Omega_2^1$ denotes the component of δ_2^1 on $\mathcal{H}_\Lambda^{\mathcal{K}_1 \otimes \mathcal{K}_2} \boxtimes \mathcal{K}_1 | \mathcal{K}_1 [\mathbb{I}^\boxtimes]$ contributed by the map component ${}^{-H}\Omega^1$ of δ^1 of $\mathcal{H}_\Lambda^{\mathcal{K}_1 \otimes \mathcal{K}_2}$. Concretely, it is given by the formula

$$\begin{aligned} \mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}_{00} &\longmapsto b \mathcal{V}_1^{\lambda_1 - 2} \phi^\tau (\mathcal{U}_1^i \mathcal{V}_1^j) \mathbf{c}_{11} \otimes \mathcal{V}_2^{-1} \tau_2 \\ \Omega_2^1(b\tau_1, -) = \mathcal{U}_1^i \mathcal{V}_1^j \mathbf{b}_{00} &\longmapsto 0 \\ \mathcal{U}_1^i \mathcal{V}_1^j \mathbf{c}_{00} &\longmapsto b \mathcal{V}_1^{\lambda_1 - 1} \phi^\tau (\mathcal{U}_1^i \mathcal{V}_1^j) \mathbf{d}_{11} \otimes \tau_2 \\ \mathcal{U}_1^i \mathcal{V}_1^j \mathbf{d}_{00} &\longmapsto b \mathcal{V}_1^{\lambda_1 - 1} \phi^\tau (\mathcal{U}_1^i \mathcal{V}_1^j) \mathbf{c}_{11} \otimes \mathcal{V}_2^{-2} \tau_2. \end{aligned}$$

We focus on the case that $b = i_1$. Consider first the case that $\mathbf{x} = \mathcal{U}_1^i \mathbf{a}$. If $a = \mathcal{U}_1^n$, then $\delta_2^1(\mathcal{U}_1^n, \mathcal{U}_1^i \mathbf{a}) = \mathcal{U}_1^{i+n} \mathbf{a} \otimes 1$, however, H^1 vanishes on this element, so we conclude that

$${}^{-H}\omega_3^1(\tau_1, \mathcal{U}_1^n, \mathcal{U}_1^i \mathbf{a}) = 0.$$

The same argument implies that ${}^{-H}\omega_3^1(\tau_1, \mathcal{V}_1^m, \mathcal{V}_1^j \mathbf{d}) = 0$.

We now consider the case that $\mathbf{x} = \mathcal{U}_1^i \mathbf{a}$ and $a = \mathcal{V}_1^m$. In this case, $H^1(\mathcal{U}_1^i \mathcal{V}_1^m \mathbf{a})$ is a sum involving both \mathbf{c} and \mathbf{b} . The next term in equation (17.1) is an application of ${}^{-H}\Omega_2^1$, which vanishes on multiples of \mathbf{b} . In particular, we only need to consider the terms of $H^1(\mathcal{U}_1^i \mathcal{V}_1^m \mathbf{a})$ involving \mathbf{c} . There are $\min(i + 1, m)$ such terms. They are

$$\mathcal{U}_1^i \mathcal{V}_1^{m-1} \mathbf{c} \otimes 1 + \mathcal{U}_1^{i-1} \mathcal{V}_1^{m-2} \mathbf{c} \otimes \mathcal{U}_2 \mathcal{V}_2 + \mathcal{U}_1^{i-2} \mathcal{V}_1^{m-3} \mathbf{c} \otimes \mathcal{U}_2^2 \mathcal{V}_2^2 + \cdots .$$

(The sum is over all terms in the sequence above where \mathcal{U}_1 and \mathcal{V}_1 both have nonnegative exponent.) We then apply ${}^{-H}\Omega_2^1$, Π^1 , and then multiply the outgoing algebra elements by repeatedly applying μ_2 . It is easy check that the application of ${}^{-H}\Omega_2^1$, Π^1 , and μ_2 on each of the above summands coincide, so we will only consider their evaluation on $\mathcal{U}_1^i \mathcal{V}_1^{m-1} \mathbf{c} \otimes 1$. Applying ${}^{-H}\Omega_2^1$ gives

$$\mathcal{V}_1^{\lambda_1-1} \phi^\tau(\mathcal{U}_1^i \mathcal{V}_1^{m-1}) \mathbf{d} \otimes \tau_2 \otimes 1 = \mathcal{U}_1^{m-1} \mathcal{V}_1^{2m+\lambda_1-3-i} \mathbf{d} \otimes \tau_2 \otimes 1.$$

Applying Π^1 to the above generator, and then multiplying the algebra elements and the coefficient $\min(i + 1, m)$ gives

$$\omega_3^1(\tau_1, \mathcal{V}_1^m, \mathcal{U}_1^i \mathbf{a}) = \min(i + 1, m) \mathcal{V}_1^{m-i-2+\lambda_1} \mathbf{d} \otimes \mathcal{U}_2^{m-1} \mathcal{V}_2^{m-1} \tau_2,$$

which is the stated formula.

We now consider ${}^{-H}\omega_3^1(\tau_1, \mathcal{U}_1^n, \mathcal{V}_1^j \mathbf{d})$. The terms of $H^1(\mathcal{U}_1^n \mathcal{V}_1^j \mathbf{d})$ which involve \mathbf{c} are

$$\mathcal{U}_1^{n-1} \mathcal{V}_1^{j-1} \mathbf{c} \otimes \mathcal{U}_2 + \mathcal{U}_1^{n-2} \mathcal{V}_1^{j-2} \mathbf{c} \otimes \mathcal{U}_2^2 \mathcal{V}_2 + \cdots .$$

As before, the sum contains all such elements in the sequence which have nonnegative powers of both \mathcal{U}_1 and \mathcal{V}_1 . There are $\min(n, j)$ terms in this sum. As before, it is sufficient to evaluate ${}^{-H}\Omega_2^1$ and Π^1 only on the first term, and then multiply the result by $\min(n, j)$. Applying ${}^{-H}\Omega_2^1$, we obtain

$$\mathcal{V}_1^{\lambda_1-1} \phi^\tau(\mathcal{U}_1^{n-1} \mathcal{V}_1^{j-1}) \mathbf{d} \otimes \tau_2 \otimes \mathcal{U}_2 = \mathcal{U}_1^{j-1} \mathcal{V}_1^{2j-2+\lambda_1-n} \mathbf{d} \otimes \tau_2 \otimes \mathcal{U}_2.$$

Applying Π^1 and $\text{id} \otimes \mu_2$, and multiplying by the coefficient $\min(n, j)$ gives

$$\omega_3^1(\tau_1, \mathcal{V}_1^m, \mathcal{U}_1^i \mathbf{a}) = \min(n, j) \mathcal{V}_1^{j-1-n+\lambda_1} \mathbf{d} \otimes \mathcal{U}_2^{j-1} \mathcal{V}_2^{j-2} \tau_2,$$

which proves the statement. ■

Remark 17.9. Recall that we write $\mathcal{K}_1 \overline{\mathcal{H}}_\Lambda^{\mathcal{K}_2}$ for the Hopf link complex obtained by using one alpha-parallel arc and one beta-parallel arc. The associated type-D module over $\mathcal{K} \otimes \mathcal{K}$ is computed in Proposition 16.7. Write $\mathcal{K}_1 \overline{\mathcal{Z}}_\Lambda^{\mathcal{K}_2}$ for the minimal model of the associated type-DA bimodule. We leave it to the reader to verify that $\mathcal{K}_1 \overline{\mathcal{Z}}_{(\lambda_1, 0)}^{\mathcal{K}_2}$ has an identical description to $\mathcal{K}_1 \mathcal{Z}_{(\lambda_1, 0)}^{\mathcal{K}_2}$ except that we omit the δ_3^1 term.

The DA bimodule structure maps of $\mathcal{K}_1 \overline{\mathcal{Z}}_{(0, 0)}^{\mathcal{K}_2}$ are illustrated in Figure 17.1.




Figure 17.1. The DA bimodule of the negative Hopf link $\mathcal{K}\overline{\mathcal{Z}}_{(0,0)}^{\mathcal{K}}$. The gray arrows denote δ_1^1 . Subscripts on algebra elements indicating link components are omitted. This coincides with the bimodule $\mathcal{K}\mathcal{Z}_{(0,0)}^{\mathcal{K}}$ except for the lack of the ω_3^1 differential.

We now prove that $\mathcal{K}_1\mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$ and $\mathcal{K}_1\mathcal{Z}_{(\lambda_1,0)}^{\mathcal{K}_2}$ are both Alexander modules, and are furthermore homotopy equivalent.

Lemma 17.10. *The DA bimodule structure maps on $\mathcal{K}_1\mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$ and $\mathcal{K}_1\mathcal{Z}_{(\lambda_1,0)}^{\mathcal{K}_2}$ are continuous.*

Furthermore, all of the maps involved in the homotopy equivalences described in Proposition 17.7 are continuous.

Proof. The DA bimodule structure maps on $\mathcal{K}_1 \mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$ are continuous by virtue of the facts that the map δ^1 on $\mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_1 \otimes \mathcal{K}_2}$ being continuous, since $\mathcal{H}_{(\lambda_1,0)}$ is a finitely generated vector space, and that $\mathcal{K}_1 \mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$ is obtained by a tensor product of $\mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_1 \otimes \mathcal{K}_2}$ with $\mathcal{K}_1 | \mathcal{K}_1 [\mathbb{I}^{\otimes 3}]$.

The structure maps on $\mathcal{K}_1 \mathcal{Z}_{(\lambda_1,0)}^{\mathcal{K}_2}$ and also the equivalence with $\mathcal{K}_1 \mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$ are given by the homological perturbation lemma. These maps are finite compositions of the maps Π^1 , H^1 , I^1 , as well as the map δ_2^1 of $\mathcal{K}_1 \mathcal{H}_{(\lambda_1,0)}^{\mathcal{K}_2}$. Hence, it is sufficient to show that each of these maps is continuous with respect to the appropriate topology. The map I^1 is obviously continuous, since it is given by $\mathbf{x} \mapsto \mathbf{x} \otimes 1$.

Consider the map Π^1 applied to elements in $\mathbf{I}_0 \cdot \mathcal{H}_{(\lambda_1,0)} \cdot \mathbf{I}_0$. This map is computed in Lemma 17.3. As an example, consider $\Pi^1(x)$ when $x = \mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}$ for $i \geq j \geq 0$. In this case,

$$\Pi^1(\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}) = \mathcal{U}_1^{i-j} \mathbf{a} \otimes \mathcal{U}_2^j \mathcal{V}_2^j.$$

Given a finite set of $S \subseteq \mathbb{N}$ and some $n \in \mathbb{N}$, we wish to show that all but finitely many $\mathcal{U}_1^i \mathcal{V}_1^j \mathbf{a}$ are mapped into $\text{Span}(\mathcal{U}_1^s \mathbf{a} \otimes \mathcal{U}_2^j \mathcal{V}_2^j : s \in \mathbb{N} \setminus S \text{ or } j \geq n)$. This is the case, since of course there are only finitely many $i, j \geq 0$ such that $0 \leq j \leq n$ and $i - j \in S$. A similar computation holds for the rest of $\mathbf{I}_0 \cdot \mathcal{H}_{(\lambda_1,0)} \cdot \mathbf{I}_0$, so Π^1 is continuous in these idempotents. Essentially the same argument applies for the other idempotents of \mathcal{H} .

The map H^1 is verified to be continuous by a very similar argument. ■

17.4 Comparison with the Eftekhary–Hedden–Levine model

We now compare our Hopf link complex with the dual knot formulas of Hedden–Levine [11] and Eftekhary [4].

If $K \subseteq S^3$, we will write $\text{EHL}_n(K)^{\mathbb{F}[u, \nu]}$ for the complex described in the introduction of [11], which is a model for $\mathcal{CFZ}(\mathcal{S}_n^3(K), \mu)^{\mathbb{F}[u, \nu]}$ where μ is a dual of K inside of the Dehn surgery.

Proposition 17.11. *If K is a knot in S^3 , then there is a canonical isomorphism*

$$\text{EHL}_n(K)^{\mathbb{F}[u, \nu]} \cong \mathcal{X}_n(K)^{\mathcal{K}} \hat{\boxtimes}_{\mathcal{K}} \mathcal{Z}^{\mathbb{F}[u, \nu]}.$$

Proof sketch. Since we do not need this result for any later results, we will not spell out all details. Instead, we will sketch several important details from which the interested reader can easily work out the rest of the argument.

Both complexes are mapping cone complexes with similar structures. We will abbreviate our complex by $\mathbb{X}_n^\mu(K)$. We write

$$\text{EHL}_n(K)^{\mathbb{F}[u, \nu]} \cong \text{Cone}(v' + h'_n : \mathbb{A}_{\text{EHL}}^-(K) \rightarrow \mathbb{B}_{\text{EHL}}^-(K))$$

and

$$\mathbb{X}_n^\mu(K)^{\mathbb{F}[\mathcal{U}, \mathcal{V}]} \cong \text{Cone}(v^\mu + h_n^\mu: \mathbb{A}^\mu(K) \rightarrow \mathbb{B}^\mu(K)).$$

We will only consider the claim when $n = 1$. Furthermore, we will only show that

$$\mathbb{A}_{\text{EHL}}^-(K) \cong \mathbb{A}^\mu(K)$$

as (infinitely generated) type-D modules over $\mathbb{F}[\mathcal{U}, \mathcal{V}]$. Most of the remaining details are straightforward extensions of the ideas we present.

We follow the description of $\text{EHL}_{+1}(K)$ given by Hedden and Levine [11]. They focus on a version of the knot Floer complex which is denoted $\text{CFK}^\infty(K)$. This takes the form of a free chain complex over $\mathbb{F}[U, U^{-1}]$ which is filtered by $\mathbb{Z} \oplus \mathbb{Z}$. The generators are of the form $[\mathbf{x}, i, j]$, where $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and $A(\mathbf{x}) = j - i$. The variable U acts by $U \cdot [\mathbf{x}, i, j] = [\mathbf{x}, i - 1, j - 1]$. The components i and j are the two components of the $\mathbb{Z} \oplus \mathbb{Z}$ -filtration. There is additionally a Maslov grading $\text{gr}_w([\mathbf{x}, i, j]) = \text{gr}_w(\mathbf{x}) + 2i$. Given such a filtered chain complex, we can recover the chain complex $\mathcal{CFK}(K)$ by replacing each $\mathbb{F}[U, U^{-1}]$ basis element $[\mathbf{x}, i, j]$ with a single generator \mathbf{x} , viewed as having Alexander grading $j - i$ and gr_w -grading $\text{gr}_w([\mathbf{x}, i, j]) - 2i$.

We will write $\mathbb{A}_{\text{EHL}}^\infty(K)$ for the infinity version of the Hedden–Levine model. (We will later reformulate this to get the version which is a type-D module over $\mathbb{F}[\mathcal{U}, \mathcal{V}]$.) By definition,

$$\mathbb{A}_{\text{EHL}}^\infty(K) \cong \prod_{s \in \mathbb{Z}} A_s^\infty(K),$$

for some finitely generated $\mathbb{Z} \oplus \mathbb{Z}$ -filtered complexes $A_s^\infty(K)$, as follows. For each \mathbf{x} (an intersection point for a Heegaard diagram of K), the generators are of the form $[\mathbf{x}, i, j]$, where $A(\mathbf{x}) = j - i$, where $A(\mathbf{x})$ is Alexander grading from $\mathcal{CFK}(K)$. We will write $\mathbf{x}_{i,j}$ for $[\mathbf{x}, i, j]$. They describe two filtrations \mathcal{I} and \mathcal{J} on these generators given by the formulas

$$\mathcal{I}(\mathbf{x}_{i,j}) = \max(i, j - s), \quad \mathcal{J}(\mathbf{x}_{i,j}) = \max(i - 1, j - s) + s.$$

Then the $\mathbb{Z} \oplus \mathbb{Z}$ -filtration of $\mathbf{x}_{i,j}$ is $(\mathcal{I}(\mathbf{x}_{i,j}), \mathcal{J}(\mathbf{x}_{i,j}))$. They also define a Maslov grading

$$\text{gr}_w(\mathbf{x}_{i,j}) = \text{gr}_w(\mathbf{x}) + 2i + \frac{(2s - 1)^2 - 1}{4}.$$

See [11, equations (1.5)–(1.10)], noting that we are setting $s = s_l, k = 1$ and $d = 1$ in their formulas.

We now rewrite \mathcal{I} and \mathcal{J} as follows:

$$\mathcal{I}(\mathbf{x}_{i,j}) = i + \max(0, j - i - s) \quad \text{and} \quad \mathcal{J}(\mathbf{x}_{i,j}) = \max(i - j + s - 1, 0) + j. \tag{17.2}$$

It is helpful to consider the two cases $A(\mathbf{x}) \geq s$ and $A(\mathbf{x}) < s$ separately. Recalling that $A(\mathbf{x}) = j - i$, we compute, using equation (17.2), that

$$(\mathcal{I}, \mathcal{J})(\mathbf{x}_{i,j}) = \begin{cases} (j - s, j) & \text{if } A(\mathbf{x}) \geq s, \\ (i, i + s - 1) & \text{if } A(\mathbf{x}) < s. \end{cases}$$

We now modify the above algebraic generators to get a type-D module over $\mathbb{F}[\mathcal{U}, \mathcal{V}]$. We do this using the same procedure as we described to go from $\text{CFK}^\infty(K)$ to $\mathcal{CFK}^-(K)$. We write $\mathbb{A}_{\text{EHL}}^-(K) := \prod_{s \in \mathbb{Z}} A_s^-(K)$ for the type-D module constructed from $\mathbb{A}_{\text{EHL}}^\infty(K)$ in this manner. For each generator \mathbf{x} in $\mathcal{CFK}^-(K)$, we get a generator \mathbf{x}'_s of $A_s^-(K)$ with Alexander grading

$$A(\mathbf{x}'_s) = \begin{cases} s & \text{if } A(\mathbf{x}) \geq s, \\ s - 1 & \text{if } A(\mathbf{x}) < s. \end{cases} \quad (17.3)$$

We have, additionally, that

$$\text{gr}_{\mathbf{w}}(\mathbf{x}'_s) = \text{gr}_{\mathbf{w}}(\mathbf{x}_{i,j}) - 2\mathcal{I}(\mathbf{x}_{i,j}) = \text{gr}_{\mathbf{w}}(\mathbf{x}) + 2 \min(0, s - A(\mathbf{x})) + \frac{(2s - 1)^2 - 1}{2}. \quad (17.4)$$

We now consider the powers of \mathcal{U} and \mathcal{V} which appear in the differential. Recall that in general, if \mathbf{x} and \mathbf{y} are generators of a knot Floer complex $\mathcal{CFK}(K)$ (for a knot K in an integer homology 3-sphere) and there is a differential from \mathbf{x} to \mathbf{y} which is weighted by $\mathcal{U}^i \mathcal{V}^j$, then

$$j - i = A(\mathbf{x}) - A(\mathbf{y}) \quad \text{and} \quad \text{gr}_{\mathbf{w}}(\mathbf{y}) = \text{gr}_{\mathbf{w}}(\mathbf{x}) - 1 + 2i. \quad (17.5)$$

We now compare this with the type-D structure $\mathbb{A}^\mu(K)$. By construction, this type-D structure is the tensor product of $\mathcal{CFK}(K)^{\mathbb{F}[\mathcal{U}, \mathcal{V}]}$ with the DA bimodule shown below:

$$\dots \mathcal{U}^2 \mathbf{a} \begin{array}{c} \xrightarrow{\mathcal{V}|U} \\ \xleftarrow{\mathcal{U}|1} \end{array} \mathcal{U} \mathbf{a} \begin{array}{c} \xrightarrow{\mathcal{V}|U} \\ \xleftarrow{\mathcal{U}|1} \end{array} \mathbf{a} \begin{array}{c} \xrightarrow{\mathcal{V}|\mathcal{V}} \\ \xleftarrow{\mathcal{U}|\mathcal{U}} \end{array} \mathbf{d} \begin{array}{c} \xrightarrow{\mathcal{V}|1} \\ \xleftarrow{\mathcal{U}|U} \end{array} \mathcal{V} \mathbf{d} \begin{array}{c} \xrightarrow{\mathcal{V}|1} \\ \xleftarrow{\mathcal{U}|U} \end{array} \mathcal{V}^2 \mathbf{d} \dots$$

We can give $\mathbb{A}^\mu(K)$ a similar description to $\mathbb{A}^-(K)$. The generators of $\mathbb{A}^\mu(K)$ are of the form $\mathbf{x} \otimes \mathbf{y}$, where \mathbf{x} is a generator of $\mathcal{CFK}^-(K)$, and \mathbf{y} is of the form $\mathcal{U}^i \mathbf{a}$ or $\mathcal{V}^j \mathbf{d}$. We define an Alexander grading on generators via the formula

$$A'(\mathcal{U}^i \mathbf{a}) = -i \quad \text{and} \quad A'(\mathcal{V}^j \mathbf{d}) = j + 1.$$

We then define $A_s^\mu(K) \subseteq \mathbb{A}^\mu(K)$ to be the \mathbb{F} span of pairs $\mathbf{x} \otimes \mathbf{y}$ where

$$A(\mathbf{x}) + A'(\mathbf{y}) = s.$$

It is straightforward to see that $A_s^\mu(K)$ is a subcomplex of \mathbb{A}^μ . We claim that there is an isomorphism of type-D modules

$$A_s^\mu(K)^{\mathbb{F}[\mathcal{U}, \mathcal{V}]} \cong A_s^-(K)^{\mathbb{F}[\mathcal{U}, \mathcal{V}]}.$$

Note that there is a canonical bijective correspondence between $\mathbb{F}[\mathcal{U}, \mathcal{V}]$ -generators of these complexes, since the generators of both complexes are bijectively identified with generators of $\mathcal{CFK}(K)$. Therefore, we have a canonical isomorphism of vector spaces between these modules. We claim that this isomorphism intertwines the differential. To see this, note that the differentials on both complexes are also identified with the ordinary differential of $\mathcal{CFK}(K)$, except with powers of \mathcal{U} and \mathcal{V} changed. Therefore, it suffices to show that the powers of \mathcal{U} and \mathcal{V} appearing in the two differentials coincide.

We consider a differential in $\mathcal{CFK}(K)$ from \mathbf{x} to \mathbf{y} , weighted by $a \in \mathbb{F}[\mathcal{U}, \mathcal{V}]$. We consider its induced weight in $A_s^\mu(K)$ and $A'_s(K)$. Suppose that in $A_s^\mu(K)$, it is weighted by an algebra element a^μ , and suppose that in $A'_s(K)$, it is weighted by an algebra element a' . Our goal is to show that $a^\mu = a'$.

Using equations (17.3) and (17.5), we observe that

$$A(a') = \begin{cases} 0 & \text{if } A(\mathbf{x}), A(\mathbf{y}) \geq s, \\ 0 & \text{if } A(\mathbf{x}), A(\mathbf{y}) < s, \\ 1 & \text{if } A(\mathbf{x}) \geq s \text{ and } A(\mathbf{y}) < s, \\ -1 & \text{if } A(\mathbf{x}) < s \text{ and } A(\mathbf{y}) \geq s. \end{cases} \tag{17.6}$$

Next, we compute, using equations (17.4) and (17.5), that

$$\text{gr}_w(a') = \text{gr}_w(\mathbf{x}) - \text{gr}_w(\mathbf{y}) - 1 + 2 \min(0, s - A(\mathbf{x})) - 2 \min(0, s - A(\mathbf{y})). \tag{17.7}$$

Note that $\text{gr}_w(a')$ and $A(a')$ uniquely determine a' .

Next, we compute a^μ and show that it is equal to a' . It is helpful to break the argument into four cases, parallel to equation (17.6). We will consider two of the four cases, and leave the rest to the reader.

We consider first the case that $A(\mathbf{x}), A(\mathbf{y}) \geq 0$. The corresponding generators of $A_s^\mu(K)$ are of the form

$$\mathbf{x} \otimes \mathcal{U}^{A(\mathbf{x})-s} \mathbf{a} \quad \text{and} \quad \mathbf{y} \otimes \mathcal{U}^{A(\mathbf{y})-s} \mathbf{a}.$$

We can write

$$a = U^{-\text{gr}_w(a)/2} \mathcal{V}^{A(a)} = U^{-(\text{gr}_w(\mathbf{x})-\text{gr}_w(\mathbf{y})-1)/2} \mathcal{V}^{A(\mathbf{x})-A(\mathbf{y})}.$$

Note that

$$\delta_2^1(a, \mathcal{U}^{A(\mathbf{x})-s} \mathbf{a}) = \mathcal{U}^{A(\mathbf{y})-s} \mathbf{a} \otimes U^{-(\text{gr}_w(\mathbf{x})-\text{gr}_w(\mathbf{y})-1)/2 + A(\mathbf{x})-A(\mathbf{y})}.$$

(Some care must be taken to verify the above formula when $A(\mathbf{x}) - A(\mathbf{y})$ is negative, but the formula holds regardless of the sign of $A(\mathbf{x}) - A(\mathbf{y})$.) Hence,

$$a^\mu = U^{-(\text{gr}_w(\mathbf{x}) - \text{gr}_w(\mathbf{y}) - 1)/2 + A(\mathbf{x}) - A(\mathbf{y})}.$$

We observe that $A(a^\mu) = 0 = A(a')$. Also,

$$\text{gr}_w(a') = \text{gr}_w(\mathbf{x}) - \text{gr}_w(\mathbf{y}) - 1 + 2(A(\mathbf{y}) - A(\mathbf{x})),$$

which is the same as $\text{gr}_w(a^\mu)$.

We now consider the third case in equation (17.6), where $A(\mathbf{x}) \geq s$ and $A(\mathbf{y}) < s$. Since $A(\mathbf{x}) \geq s$, the corresponding generator of $A_s^\mu(K)$ is of the form $\mathbf{x} \otimes \mathcal{U}^{A(\mathbf{x})-s} \mathbf{a}$. Also, the generator of $A_s^\mu(K)$ corresponding to \mathbf{y} will be of the form $\mathbf{y} \otimes \mathcal{V}^{s-A(\mathbf{y})-1} \mathbf{d}$. We write

$$a = U^{-(\text{gr}_w(\mathbf{x}) - \text{gr}_w(\mathbf{y}) - 1)/2} \mathcal{V}^{A(\mathbf{x}) - A(\mathbf{y})}.$$

In this case, we have $A(\mathbf{x}) - A(\mathbf{y}) > 0$. Therefore we can read a^μ as follows. Firstly, we have a factor of $U^{-(\text{gr}_w(\mathbf{x}) - \text{gr}_w(\mathbf{y}) - 1)/2}$. Then, we have an additional factor which is obtained by composing all of the right moving arrows from $\mathcal{U}^{A(\mathbf{x})-s} \mathbf{a}$ to $\mathcal{V}^{s-A(\mathbf{y})-1} \mathbf{d}$. This will give us factors of $U^{A(\mathbf{x})-s} \mathcal{V}$. Therefore,

$$a^\mu = U^{-(\text{gr}_w(\mathbf{x}) - \text{gr}_w(\mathbf{y}) - 1)/2 + A(\mathbf{x}) - s} \mathcal{V}.$$

This must coincide with a' , since it has the same (gr_w, A) -bigrading by equations (17.6) and (17.7).

It remains to verify that $a' = a^\mu$ in the cases that $A(\mathbf{x}), A(\mathbf{y}) < s$ and when $A(\mathbf{x}) < s$ and $A(\mathbf{y}) \geq s$. The argument follows from the same line of reasoning as the two previously analyzed cases. We leave the details to the reader.

The analysis of $\mathbb{B}^\mu(K)$ and the maps v^μ and h_n^μ follows from similar, albeit somewhat tedious, reasoning, so we leave the details to the interested reader. ■

Chapter 18

Examples and basic properties

In this chapter, we perform several example computations. In Section 18.1, we compute the effect of adding a ± 1 -framed meridian to a link component. We show that it corresponds to a simple bimodule. In Section 18.2, we compute the type-D module for a p/q -framed solid torus. We show that the type-D module for a p/q -framed solid torus recovers the rational surgeries formula of Ozsváth and Szabó [44]. Additionally, we show that the type-D modules for solid tori recovers the surgery exact triangle by exhibiting a homotopy equivalence

$$\mathcal{D}_\infty^{\mathcal{K}} \simeq \text{Cone}(f^1: \mathcal{D}_n^{\mathcal{K}} \rightarrow \mathcal{D}_{n+1}^{\mathcal{K}})$$

for all $n \in \mathbb{Z}$.

18.1 Meridional Dehn twists

In this section, we consider the effect of adding a ± 1 -framed meridian to a link component. Adding a meridian may be encoded by taking the connected sum with a Hopf link, which has a predictable algebraic effect.

We first define our algebraic candidate, denoted ${}_{\mathcal{K}}\mathcal{B}_{\pm 1}^{\mathcal{K}}$. These are the bimodules for two very simple algebra endomorphisms of \mathcal{K} . As an (\mathbf{I}, \mathbf{I}) -module, $\mathcal{B}_{\pm 1} \cong \mathbf{I}$ with the natural \mathbf{I} -action. The structure map δ_1^1 vanishes. If $a \in \mathbf{I}_j \cdot \mathcal{K} \cdot \mathbf{I}_j$ for either $j = 0$ or $j = 1$, then we set $\delta_2^1(a, i) = i \otimes a$ for $i \in \mathbf{I}$. Additionally, we set

$$\delta_2^1(\sigma, i_0) = i_1 \otimes \sigma \quad \text{and} \quad \delta_2^1(\tau, i_0) = i_1 \otimes \mathcal{V}^{\pm 1} \tau.$$

We write ${}_{\mathcal{K}}\mathfrak{M}_{\pm 1}^{\mathcal{K}}$ for the type-D module

$${}_{\mathcal{K}}\mathfrak{M}_{\pm 1}^{\mathcal{K}} = \mathcal{H}_{\Lambda(\pm 1)}^{\mathcal{K}_0 \otimes \mathcal{K}_1} \hat{\boxtimes} {}_{\mathcal{K}_0} \mathcal{D}_0 \hat{\boxtimes} {}_{\mathcal{K}_1 | \mathcal{K}} W_{\alpha\beta, \alpha}^{\mathcal{K}}, \quad (18.1)$$

where $W_{\alpha\beta, \alpha}$ is the pair-of-pants module from Chapter 15 and

$$\Lambda(\pm 1) = (\pm 1, 0).$$

In the above, \mathcal{K}_0 and \mathcal{K}_1 denote the algebras associated to different components of the Hopf link. We write \mathcal{K}_0 for the algebra associated to the component of the Hopf link which becomes the new meridian (and is framed by ± 1).

By the pairing theorem and Theorem 15.1, adding a ± 1 -framed meridian to a component $K \subseteq L$ and tensoring ${}_{\mathcal{K}}\mathcal{D}_0$ into the factor of the meridian has the effect of tensoring with the bimodule ${}_{\mathcal{K}}\mathfrak{M}_{\pm 1}^{\mathcal{K}}$.

Proposition 18.1. *Let $\mathcal{K}\mathfrak{M}_{\pm 1}^{\mathcal{K}}$ denote the DA bimodule in equation (18.1) for adding a ± 1 -framed meridian. Then*

$$\mathcal{K}\mathfrak{M}_{\pm 1}^{\mathcal{K}} \simeq \mathcal{K}\mathcal{B}_{\mp 1}^{\mathcal{K}}.$$

The proof is a direct computation, which we do in several steps. Firstly, we observe the very basic isomorphism

$$\mathcal{K}_0\mathcal{D}_0 \cong \mathcal{D}_0^{\mathcal{K}_2} \hat{\boxtimes} \mathcal{K}_{0|\mathcal{K}_2}[\mathbb{I}^{\boxplus}].$$

(The subscripts on different copies of \mathcal{K} are meant to indicate only how the tensor product is formed.) Hence, we may manipulate several terms in equation (18.1) as follows:

$$\begin{aligned} \mathcal{H}_{\Lambda(\pm 1)}^{\mathcal{K}_0 \otimes \mathcal{K}_1} \hat{\boxtimes} \mathcal{K}_0\mathcal{D}_0 &\simeq \mathcal{D}_0^{\mathcal{K}_2} \hat{\boxtimes} (\mathcal{H}_{\Lambda(\pm 1)}^{\mathcal{K}_0 \otimes \mathcal{K}} \hat{\boxtimes} \mathcal{K}_{0|\mathcal{K}_2}[\mathbb{I}^{\boxplus}]) \\ &:= \mathcal{D}_0^{\mathcal{K}_2} \hat{\boxtimes} \mathcal{K}_2\mathcal{H}_{\Lambda(\pm 1)}^{\mathcal{K}_1} \\ &\simeq \mathcal{D}_0^{\mathcal{K}_2} \hat{\boxtimes} \mathcal{K}_2\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}_1}. \end{aligned} \quad (18.2)$$

Therefore, combining equations (18.1) and (18.2), we see that

$$\begin{aligned} \mathcal{K}\mathfrak{M}_{\pm 1}^{\mathcal{K}} &= \mathcal{H}_{\Lambda(\pm 1)}^{\mathcal{K}_0 \otimes \mathcal{K}_1} \hat{\boxtimes} \mathcal{K}_0\mathcal{D}_0 \hat{\boxtimes} \mathcal{K}_{1|\mathcal{K}}W_{\alpha\beta,\alpha}^{\mathcal{K}} \\ &\simeq \mathcal{D}_0^{\mathcal{K}_2} \hat{\boxtimes} \mathcal{K}_2\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}_1} \hat{\boxtimes} \mathcal{K}_{1|\mathcal{K}}W_{\alpha\beta,\alpha}^{\mathcal{K}}. \end{aligned} \quad (18.3)$$

On the other hand, we observe by direct computation that

$$\mathcal{K}\mathcal{B}_{\mp 1}^{\mathcal{K}} \simeq \mathcal{D}_{\mp 1}^{\mathcal{K}_0} \hat{\boxtimes} \mathcal{K}_{0|\mathcal{K}}W_{\alpha\beta,\alpha}^{\mathcal{K}}. \quad (18.4)$$

Comparing equations (18.3) and (18.4), we observe that to prove Proposition 18.1, it suffices to show

$$\mathcal{D}_0^{\mathcal{K}_1} \hat{\boxtimes} \mathcal{K}_1\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}} \simeq \mathcal{D}_{\mp 1}^{\mathcal{K}}. \quad (18.5)$$

Let us write

$$\mathcal{Z}_{\Lambda}^{\mathcal{K}} := \mathcal{D}_0^{\mathcal{K}_1} \hat{\boxtimes} \mathcal{K}_1\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}}.$$

We note that the type-D module $\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}}$ is not minimal, in the sense of Definition 17.1. Using the notation of Proposition 17.7, the complex $\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}}$ may be written as follows:

$$\mathcal{Z}_{\Lambda(\pm 1)}^{\mathcal{K}} = \begin{array}{ccc} \mathcal{Z}_{0,0}^{\mathcal{K}} & \xrightarrow{p_2^1(\sigma_1 + \tau_1, -)} & \mathcal{Z}_{1,0}^{\mathcal{K}} \\ \downarrow f_1^1 & & \downarrow g_1^1 \\ \mathcal{Z}_{0,1}^{\mathcal{K}} & \xrightarrow{q_2^1(\sigma_1 + \tau_1, -)} & \mathcal{Z}_{1,1}^{\mathcal{K}} \end{array} \quad (18.6)$$

In the above, we are writing f_1^1 for the map $L_2 f_1^1 + -L_2 f_1^1$, and similarly for all of the other maps. Note that $-^H \omega_3^1$ does not make any contribution.

Our argument will go by way of viewing the above diagram as a mapping cone from the top line to the bottom line, and then applying the homological perturbation lemma. Let us write

$$\mathcal{Q}_{\Lambda(\pm 1), \varepsilon}^{\mathcal{K}} = \text{Cone}(\mathcal{Z}_{0, \varepsilon}^{\mathcal{K}} \xrightarrow{p_2^1(\sigma_1 + \tau_1, -)} \mathcal{Z}_{1, \varepsilon}^{\mathcal{K}}).$$

We will write

$$J_{\Lambda(\pm 1)}^1 : \mathcal{Q}_{\Lambda(\pm 1); 0}^{\mathcal{K}} \rightarrow \mathcal{Q}_{\Lambda(\pm 1); 1}^{\mathcal{K}}$$

for the total morphism from the top line to the bottom line of equation (18.6) (i.e., $f_1^1 \oplus g_1^1$), so that $\mathcal{Z}_{(\pm 1), 0}^{\mathcal{K}}$ is $\text{Cone}(J_{\Lambda(\pm 1)}^1)$.

For $\varepsilon \in \{0, 1\}$, write $\mathbf{i}_\varepsilon^{\mathcal{K}}$ for the type-D module which is concentrated in idempotent \mathbf{I}_ε , has a single generator, and has vanishing δ^1 . Equation (18.5), and consequently Proposition 18.1, follows from the subsequent two lemmas.

Lemma 18.2. *For $\varepsilon \in \{0, 1\}$, there are maps*

$$I_\varepsilon^1 : \mathbf{i}_\varepsilon^{\mathcal{K}} \rightarrow \mathcal{Q}_{\Lambda(-1); \varepsilon}^{\mathcal{K}}, \quad \Pi_\varepsilon^1 : \mathcal{Q}_{\Lambda(-1); \varepsilon}^{\mathcal{K}} \rightarrow \mathbf{i}_\varepsilon^{\mathcal{K}}, \quad \text{and} \quad H_\varepsilon^1 : \mathcal{Q}_{\Lambda(-1); \varepsilon}^{\mathcal{K}} \rightarrow \mathcal{Q}_{\Lambda(-1); \varepsilon}^{\mathcal{K}}$$

such that I_ε^1 and Π_ε^1 are cycles, and $I_\varepsilon^1 \circ \Pi_\varepsilon^1 = \text{id} + \partial_{\text{Mor}}(H_\varepsilon^1)$ and $\Pi_\varepsilon^1 \circ I_\varepsilon^1 = \text{id}$. Furthermore,

$$\Pi_1^1 \circ J_{\Lambda(-1)}^1 \circ I_0^1 = (i_0 \mapsto i_1 \otimes (\sigma + \mathcal{V}\tau)). \tag{18.7}$$

Here, i_ε denotes the generator of $\mathbf{i}_\varepsilon^{\mathcal{K}}$. In particular, $\mathcal{Z}_{\Lambda(-1)}^{\mathcal{K}}$ is homotopy equivalent to the solid torus module $\mathcal{D}_{+1}^{\mathcal{K}}$ defined in Section 8.1.

Proof. As a first step, we record the mapping cone of $\mathcal{Q}_{\Lambda(-1); \varepsilon}^{\mathcal{K}}$. This is the following type-D module:

$$\begin{array}{ccccccccccc} \dots & u_1^3 \mathbf{a} & u_1^2 \mathbf{a} & u_1^1 \mathbf{a} & \mathbf{a} & \mathbf{d} & \mathcal{V}_1 \mathbf{d} & \mathcal{V}_1^2 \mathbf{d} & \mathcal{V}_1^3 \mathbf{d} & \dots \\ & \downarrow u_2^3 v_2^4 & \downarrow u_2^2 v_2^3 & \downarrow u_2 v_2^2 & \downarrow v_2 & \downarrow u_2 & \downarrow u_2 v_2 & \downarrow u_2^3 v_2^2 & \downarrow u_2^4 v_2^3 & \downarrow u_2^5 v_2^4 & \downarrow u_2^6 v_2^5 \\ \dots & \mathcal{V}_1^{-4} \mathbf{d} & \mathcal{V}_1^{-3} \mathbf{d} & \mathcal{V}_1^{-2} \mathbf{d} & \mathcal{V}_1^{-1} \mathbf{d} & \mathbf{d} & \mathcal{V}_1 \mathbf{d} & \mathcal{V}_1^2 \mathbf{d} & \mathcal{V}_1^3 \mathbf{d} & \dots \end{array}$$

We define the type-D module $N^\mathcal{E}$, where \mathcal{E} is the exterior algebra on one generator, θ . The module N has a single generator, and differential $\delta^1(1) = 1 \otimes \theta$. Then $\mathcal{Q}_{\Lambda(-1), \varepsilon}^{\mathcal{K}}$ is obtained by boxing $N^\mathcal{E}$ with the DA bimodule over $(\mathcal{E}, \mathcal{K})$ shown below:

$$\begin{array}{ccccccccccc} \dots & u_1^3 \mathbf{a} & u_1^2 \mathbf{a} & u_1^1 \mathbf{a} & \mathbf{a} & \mathbf{d} & \mathcal{V}_1 \mathbf{d} & \mathcal{V}_1^2 \mathbf{d} & \mathcal{V}_1^3 \mathbf{d} & \dots \\ & \downarrow \theta | u_2^3 v_2^4 & \downarrow \theta | u_2^2 v_2^3 & \downarrow \theta | u_2 v_2^2 & \downarrow \theta | v_2 & \downarrow \theta | u_2 & \downarrow \theta | u_2 v_2 & \downarrow \theta | u_2^3 v_2^2 & \downarrow \theta | u_2^4 v_2^3 & \downarrow \theta | u_2^5 v_2^4 & \downarrow \theta | u_2^6 v_2^5 \\ \dots & \mathcal{V}_1^{-4} \mathbf{d} & \mathcal{V}_1^{-3} \mathbf{d} & \mathcal{V}_1^{-2} \mathbf{d} & \mathcal{V}_1^{-1} \mathbf{d} & \mathbf{d} & \mathcal{V}_1 \mathbf{d} & \mathcal{V}_1^2 \mathbf{d} & \mathcal{V}_1^3 \mathbf{d} & \dots \end{array}$$

In the above diagram, the solid lines denote the δ^1 action, while the dashed lines denote $\delta_2^1(\theta, -)$. A solid arrow from \mathbf{x} to \mathbf{y} with a 1 means that $\delta^1(\mathbf{x})$ has a summand of $\mathbf{y} \otimes 1$. If we ignore the type-A action, then the type-D module is equivalent to the one spanned by $\mathcal{V}_1^{-1}\mathbf{d}$ with vanishing δ^1 . There is a canonical inclusion map i^1 and projection map π^1 . A homotopy h^1 is defined by moving backwards along the solid arrows. Clearly, i^1 and π^1 are type-D homomorphisms, and

$$i^1 \circ \pi^1 = \text{id} + \partial_{\text{Mor}}(h^1), \quad \pi^1 \circ i^1 = \text{id}.$$

Similarly, $\pi^1 \circ h^1 = 0$, $h^1 \circ h^1 = 0$, and $h^1 \circ i^1 = 0$. By the homological perturbation lemma, Lemma 4.2, i^1 , π^1 , and h^1 extend to a homotopy equivalence of DA bimodules, which we denote by i_*^1 , π_*^1 , and h_*^1 . We define the type-D module morphisms in the statement by

$$I_\varepsilon^1 = \mathbb{I}_N \boxtimes i_*^1, \quad \Pi_\varepsilon^1 = \mathbb{I}_N \boxtimes \pi_*^1, \quad \text{and} \quad H_\varepsilon^1 = \mathbb{I}_N \boxtimes h_*^1.$$

We leave it to the reader to check that the above expressions (which involve infinite sums on the completion) determine continuous morphisms.

All of the remaining claims in the statement are clear, except equation (18.7). We compute that I_0^1 maps i_0 to $\mathcal{V}_1^{-1}\mathbf{d} \otimes 1$. By the computation of Proposition 17.7, the map $J_{\Lambda(-1)}^1$ sends this to

$$\mathcal{V}_1^{-2}\mathbf{d} \otimes \tau_2 + \mathcal{V}_1^{-1}\mathbf{d} \otimes \sigma_2.$$

Then Π_1^1 maps this to

$$i_1 \otimes \mathcal{V}_2\tau_2 + i_1 \otimes \sigma_2.$$

Note that the application of Π_1^1 to $\mathcal{V}_1^{-2}\mathbf{d} \otimes \tau_2$ is slightly subtle. Indeed, the recipe from the homological perturbation lemma is to move backwards along the arrow from \mathbf{a} to $\mathcal{V}_1^{-2}\mathbf{d}$ and move forward along the arrow from \mathbf{a} to $\mathcal{V}_1^{-1}\mathbf{d}$, while picking up a factor of \mathcal{V}_2 . The proof is complete. ■

We now consider the claim for the framing $\Lambda(+1)$.

Lemma 18.3. *There are maps*


$$I_\varepsilon^1: \mathfrak{I}_\varepsilon^{\mathcal{K}} \rightarrow \mathcal{Q}_{\Lambda(+1); \varepsilon}^{\mathcal{K}}, \quad \Pi_\varepsilon^1: \mathcal{Q}_{\Lambda(+1); \varepsilon}^{\mathcal{K}} \rightarrow \mathfrak{I}_\varepsilon^{\mathcal{K}}, \quad \text{and} \quad H_\varepsilon^1: \mathcal{Q}_{\Lambda(+1); \varepsilon}^{\mathcal{K}} \rightarrow \mathcal{Q}_{\Lambda(+1); \varepsilon}^{\mathcal{K}}$$

such that I_ε^1 and Π_ε^1 are type-D homomorphisms, and $I_\varepsilon^1 \circ \Pi_\varepsilon^1 = \text{id} + \partial_{\text{Mor}}(H_\varepsilon^1)$ and $\Pi_\varepsilon^1 \circ I_\varepsilon^1 = \text{id}$. Furthermore,


$$\Pi_1^1 \circ J_{\Lambda(+1)}^1 \circ I_0^1 = (i_0 \mapsto i_1 \otimes (\sigma + \mathcal{V}^{-1}\tau)).$$

Here, i_ε denotes the generator of $\mathcal{Q}_{+1, \varepsilon}^{\mathcal{K}}$. In other words, $\mathcal{Z}_{+1}^{\mathcal{K}} \simeq \mathcal{D}_{-1}^{\mathcal{K}}$.

Proof. The proof is very much the same as the proof with framing -1 . We write down the complex $\mathcal{Q}_{\Lambda(+1); \varepsilon}^{\mathcal{K}}$ below:



Similar to the framing -1 case, we realize the above diagram as the box tensor product of N^ε with the following DA bimodule over $(\mathcal{E}, \mathcal{K})$:



We forget about the left \mathcal{E} -action and simplify the type-D structure. We build a homotopy equivalence with the module $\mathcal{W}_\varepsilon^{\mathcal{K}}$, by picking type-D module maps i^1 , π^1 , and h^1 . The map i^1 sends i_ε to $(\mathbf{a} + \mathbf{d}) \otimes 1$. The map π^1 requires a choice. It is either $(\mathbf{a} \mapsto i_\varepsilon \otimes 1, \mathbf{d} \mapsto 0)$, or $(\mathbf{a} \mapsto 0, \mathbf{d} \mapsto i_\varepsilon \otimes 1)$. The choice of h^1 depends on our choice of π^1 . If π^1 sends \mathbf{a} to i_ε , then h^1 maps \mathbf{d} (on the bottom row) to $\mathbf{d} \otimes 1$ (on the top row). If π^1 maps \mathbf{d} to $i_\varepsilon \otimes 1$, then h^1 maps \mathbf{d} (bottom row) to $\mathbf{a} \otimes 1$ (top row). Arbitrarily, pick π^1 to map \mathbf{d} to $i_\varepsilon \otimes 1$. Away from the central region, the map h^1 just maps backwards along the arrows marked with 1. One easily verifies that π^1 , i^1 , and h^1 satisfy the assumptions of the homological perturbation lemma.

The homological perturbation lemma now gives us extensions π_*^1 , i_*^1 and h_*^1 . Boxing these with \mathbb{I}_N gives homotopy equivalences between $\mathcal{Q}_{\Lambda(+1); \varepsilon}^{\mathcal{K}}$ and $\mathbf{i}_\varepsilon^{\mathcal{K}}$ described in the statement. Note that the maps I_ε^1 , Π_ε^1 , and H_ε^1 sometimes involve infinite sums, however, it is straightforward to verify that the induced maps are continuous.

It remains to perform the computation involving $J_{\Lambda(+1)}^1$. We easily compute

$$I_0^1(i_0) = \dots + \mathcal{U}_1 \mathbf{a} \otimes \mathcal{V}_2 + \mathbf{a} \otimes 1 + \mathbf{d} \otimes 1 + \mathcal{V}_1 \mathbf{d} \otimes \mathcal{U}_2 + \dots .$$

Using part (2) of Proposition 17.7, we compute that the map $J_{\Lambda(+1)}^1$ sends the above to

$$\begin{aligned}
 & \dots + \mathcal{U}_1 \mathbf{a} \otimes \sigma_2 \mathcal{V}_2 + \mathbf{a} \otimes \sigma_2 + \mathbf{d} \otimes \sigma_2 + \mathcal{V}_1 \mathbf{d} \otimes \sigma_2 \mathcal{U}_2 + \dots , \\
 & \dots + \mathcal{U}_1^2 \mathbf{a} \otimes \tau_2 \mathcal{V}_2 + \mathcal{U}_1 \mathbf{a} \otimes \tau_2 + \mathbf{a} \otimes \mathcal{V}_2^{-1} \tau_2 + \mathbf{d} \otimes \tau_2 \mathcal{U}_2 + \dots .
 \end{aligned}$$

The map Π_1^1 evaluates the above $i_1 \otimes (\sigma_2 + \mathcal{V}_2^{-1} \tau_2)$ (regardless of which choice we made in the construction of π^1). Note, we are using here that $\tau_2 \mathcal{U}_2 = \mathcal{V}_2^{-1} \tau_2$. The proof is complete. ■

18.2 The type-D invariants of solid tori

We now describe the type-D module for a solid torus with $p/q \in \mathbb{Q} \cup \{\infty\}$ framing. The most convenient description is to take a standard unknot complement (with standard meridian and 0-framed longitude) and perform $-q/p$ surgery to a meridian of U . See Figure 18.1. Additionally, we will consider the ∞ -framed solid torus, which we view as being obtained by performing 0 surgery to a meridian of U .




Figure 18.1. A slam dunk on a p/q -framed solid torus (top row), and an ∞ -framed solid torus (bottom). The component U with an arrow on it denotes the component which we associate to \mathcal{K} .

We now define our candidate type-D modules. We begin by recalling that by definition the type-D module for an n -framed solid torus $\mathcal{D}_n^{\mathcal{K}}$ is the same as the surgery complex for an n -framed unknot, and hence has two generators \mathbf{x}^0 and \mathbf{x}^1 , which live in idempotent 0 and 1, respectively. The structure map is given by

$$\delta^1(\mathbf{x}^0) = \mathbf{x}^1 \otimes (\sigma + \mathcal{V}^n \tau).$$

Generalizing this, we now define the candidate p/q -framed solid torus module $\mathcal{D}_{p/q}^{\mathcal{K}}$, where p, q are coprime and $q > 0$. We define $\mathcal{D}_{p/q}^{\mathcal{K}} \cdot \mathbf{I}_\varepsilon$ to be spanned by generators $\mathbf{x}_0^\varepsilon, \dots, \mathbf{x}_{q-1}^\varepsilon$, for $\varepsilon \in \{0, 1\}$. The structure map is given by the following formula:

$$\delta^1(\mathbf{x}_i^0) = \mathbf{x}_{(i+p) \bmod q}^1 \otimes \mathcal{V}^{\lfloor (i+p)/q \rfloor} \tau + \mathbf{x}_i^1 \otimes \sigma.$$

The examples $\mathcal{D}_{\pm 1/3}^{\mathcal{K}}$ are shown in Figure 18.2.

Remark 18.4. The bimodule $\mathcal{D}_{p/q}^{\mathcal{K}}$ recovers the rational surgery mapping cone complex $\mathbb{X}_{p/q}(K)$ of Ozsváth and Szabó [44] in the sense that if $K \subseteq S^3$ is a knot, then

$$\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{X}_0(K) \cong \mathbb{X}_{p/q}(K).$$




Figure 18.2. The modules $\mathcal{D}_{1/3}^{\mathcal{K}}$ (left) and $\mathcal{D}_{-1/3}^{\mathcal{K}}$ (right).

Finally, we define the type-D module $\mathcal{D}_{\infty}^{\mathcal{K}}$ for the ∞ -framed solid torus. This module $\mathcal{D}_{\infty}^{\mathcal{K}} \cdot \mathbf{I}_0 = \{0\}$ and $\mathcal{D}_{\infty}^{\mathcal{K}} \cdot \mathbf{I}_1 = \text{Span}_{\mathbb{F}}(\mathbf{x}, \mathbf{y})$. The structure map is given by

$$\delta^1(\mathbf{x}) = \mathbf{y} \otimes (1 + \mathcal{V}).$$

The link surgery formula naturally produces a type-D module $\mathcal{X}_{p/q}^{\mathcal{K}}$ for the p/q -framed solid torus, as follows. We may write p/q as a continued fraction expansion

$$p/q = [a_n, \dots, a_1]^- = a_n - \frac{1}{a_{n-1} - \frac{1}{a_{n-2} - \dots}}.$$

We may define $\mathcal{X}_{p/q}^{\mathcal{K}}$ as the link surgery complex of a linear plumbing with weights a_n, \dots, a_1 . This produces a type-D module over \mathcal{L}_n . We tensor 0-framed solid tori to the components labeled a_{n-1}, \dots, a_1 . The remaining component, weighted a_n , corresponds to the type-D algebra action of $\mathcal{X}_{p/q}^{\mathcal{K}}$.

Proposition 18.5. *The type-D module for a p/q -framed solid torus $\mathcal{X}_{p/q}^{\mathcal{K}}$ is homotopy equivalent to $\mathcal{D}_{p/q}^{\mathcal{K}}$.*

Proof. The proof is by induction. By considering the continued fraction decomposition above, we will assume that the formula holds for p/q , and use this to show that it also holds for

$$m - \frac{1}{p/q} = m - \frac{q}{p},$$

when $m \in \mathbb{Z}$.

Note that increasing the framing by m may be performed by tensoring with m meridional Dehn twist bimodules $\mathcal{K} \mathcal{B}_{\pm 1}^{\mathcal{K}}$ from Section 18.1. This tensor product is easy to understand, and indeed it is easy to check that

$$\mathcal{D}_{p/q}^{\mathcal{K}} \hat{\boxtimes} \mathcal{K} \mathcal{B}_{+1}^{\mathcal{K}} \cong \mathcal{D}_{p/q+1}^{\mathcal{K}}.$$

In particular, it suffices to show that if the claim is true for some p/q , then it is true for $-q/p$.

On the algebraic level, it is sufficient to show that

$$\mathcal{D}_{-q/p}^{\mathcal{K}} \simeq \mathcal{D}_{p/q}^{\mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{Z}_{(0,0)}^{\mathcal{K}},$$

where ${}_{\mathcal{K}}\mathcal{Z}_{(0,0)}^{\mathcal{K}}$ is the Hopf link bimodule.

In this case, we compute the type-D module in a similar manner to our computation of the meridional Dehn twist module from Section 18.1. We will only consider the case that $p < 0$ and $q > 0$. The case that $p > 0$ and $q > 0$ follows from a straightforward modification of this argument.

We compute $\mathcal{D}_{p/q}^{\mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_0$ first. The tensor product has generators

$$x_j^i := \mathbf{x}_j^0 u_1^i \mathbf{a} \quad \text{and} \quad y_j^i := \mathbf{x}_j^0 v_1^i \mathbf{d},$$

ranging over $i \in \mathbb{N}$ and $j \in \mathbb{Z}/q$, and generators

$$z_j^i := \mathbf{x}_j^1 v_1^i \mathbf{d}$$

ranging over $i \in \mathbb{Z}$ and $j \in \mathbb{Z}/q$. We display the complex in Figure 18.3.

Let us write $|p| = a \cdot q + r$, where $a \geq 0$ and $0 \leq r \leq q - 1$. The minimal model of $\mathcal{D}_{p/q}^{\mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_0$ is generated by the $|p|$ generators

$$\mathbf{w}_0^0 = z_{q-1}^{-1}, \mathbf{w}_1^0 = z_{q-2}^{-1}, \dots, \mathbf{w}_{|p|-1}^0 = z_{q-r}^{-a-1}.$$




Figure 18.3. The tensor product $\mathcal{D}_{-1/5}^{\mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_\varepsilon$, for either $\varepsilon \in \{0, 1\}$. The boxed generators are the generators of the minimal model of $\mathcal{D}_{p/q}^{\mathcal{K}} \hat{\boxtimes} {}_{\mathcal{K}}\mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_\varepsilon$. The dashed arrows are weighted by 1. The dotted arrow indicates a term of δ^1 which maps from idempotent 0 to idempotent 1. (Note that technically this arrow goes from the copy of this complex in idempotent 0 to the copy in idempotent 1.)

We think of the above generators as all of the generators between z_{q-r}^{-a-1} and z_{q-1}^{-1} if we arrange the generators in a row as in the bottom row of Figure 18.3. (Recall here we are assuming that $p < 0$ and $q > 0$.) The same argument works for $\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_1$, and we write $\mathbf{w}_0^1, \dots, \mathbf{w}_{|p|-1}^1$ for these generators.

This gives the stated identification of the generators of the module. It remains now to understand δ^1 . This is obtained via homological perturbation (cf. Lemma 4.5) by first applying the inclusion map from $\mathcal{D}_{-q/p}^{\mathcal{K}} \cdot \mathbf{I}_0$ into $\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_0$, applying δ^1 of $\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}}$ and then projecting $\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_1$ to $\mathcal{D}_{-q/p}^{\mathcal{K}} \cdot \mathbf{I}_1$.

There is clearly a summand of δ^1 which sends \mathbf{w}_i^0 to $\mathbf{w}_i^1 \otimes \sigma$, for each $i \in \mathbb{Z}/q$. There is a remaining term of δ^1 terms which contributes τ . We recall that via homological perturbation theory, these terms are obtained by first including, applying δ^1 , and then projecting. After including, the map δ^1 moves a generator of $\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}}$ to the left by q places. The outgoing algebra element is 1. We then apply the projection map. If the application of δ^1 sent \mathbf{w}_i^0 to $\mathbf{w}_j^1 \otimes \tau$ for $\mathbf{w}_j^1 \in \{\mathbf{w}_0^1, \dots, \mathbf{w}_{p-1}^1\}$, then the projection map sends \mathbf{w}_j^1 to $\mathbf{w}_j^1 \otimes 1$. In general, the projection map will involve traveling to the generators $\mathbf{w}_0^1, \dots, \mathbf{w}_{p-1}^1$ via a zigzag of arrows in the larger complex $\mathcal{D}_{p/q}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}}$ (more precisely, such a zigzag may be encoded by the homological perturbation lemma, similar to our argument from Section 18.1). We pick up one power of \mathcal{V} for each zigzag we have to do to get back to $\mathbf{w}_{(i+q) \bmod p}^1$. The number of zigzags that we have to do to map $\delta^1(\mathbf{w}_i^0)$ to $\mathbf{w}_{(i+q) \bmod p}^1$ is exactly $\lfloor (i+q)/(-p) \rfloor$. This coincides with $\mathcal{D}_{q/(-p)}^{\mathcal{K}}$, completing the proof. ■

Proposition 18.6. *The type-D module for an ∞ -framed solid torus is $\mathcal{D}_{\infty}^{\mathcal{K}}$.*

Proof. We compute the tensor product

$$\mathcal{D}_{\infty}^{\mathcal{K}} := \mathcal{D}_0^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}}.$$

We note for both $\varepsilon = 0, 1$, that $\mathcal{D}_0 \widehat{\boxtimes} \mathcal{K} \mathcal{Z}_{(0,0)}^{\mathcal{K}} \cdot \mathbf{I}_{\varepsilon}$ is isomorphic to the following complex:

$$\begin{array}{cccccc} \dots & \mathcal{U}_1 \mathbf{a} & \mathbf{a} & \mathbf{d} & \mathcal{V}_1 \mathbf{d} & \dots \\ & \downarrow 1+\mathcal{U}\mathcal{V}^2 & \downarrow 1+\mathcal{V} & \downarrow 1+\mathcal{U} & \downarrow 1+\mathcal{U}^2\mathcal{V} & \\ \dots & \mathcal{V}_1^{-2} \mathbf{d} & \mathcal{V}_1^{-1} \mathbf{d} & \mathbf{d} & \mathcal{V}_1 \mathbf{d} & \dots \end{array}$$

In idempotent \mathbf{I}_0 , this complex is acyclic. For example, $\text{Cone}(1 + \mathcal{U} : \mathbf{d} \rightarrow \mathbf{d})$ is acyclic because $1 + \mathcal{U}$ is a unit in the completion of $\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$. Its inverse is given by $\sum_{i=0}^{\infty} \mathcal{U}^i$. In idempotent \mathbf{I}_1 , each summand is acyclic except for $\text{Cone}(1 + \mathcal{V} : \mathbf{a} \rightarrow \mathcal{V}_1^{-1} \mathbf{d})$. Deleting each acyclic subcomplex yields $\mathcal{D}_{\infty}^{\mathcal{K}}$, completing the proof. ■

There is a related bimodule which is useful for applications, defined by

$$\mathcal{K} \mathcal{D}_{\infty}^{\mathbb{F}[U]} := \mathcal{D}_{\infty}^{\mathcal{K}} \widehat{\boxtimes} \mathcal{K}_{|\mathcal{K}}[\mathbb{I}^{\ni}]^{\mathbb{F}[U]}.$$

The following lemma is helpful.

Lemma 18.7. *The DA bimodule $\mathcal{X}\mathcal{D}_\infty^{\mathbb{F}[U]}$ is homotopy equivalent to the DA bimodule with a single generator \mathbf{z}^1 , concentrated in idempotent \mathbf{I}_1 , such that $\delta_2^1(\mathcal{U}^i \mathcal{V}^j, \mathbf{z}^1) = \mathbf{z}^1 \otimes U^i$.*

Proof. Write $\mathcal{X}\mathcal{Z}_\infty^{\mathbb{F}[U]}$ for the small model in the statement. We view the completion of $\mathcal{X}\mathcal{D}_\infty^{\mathbb{F}[U]}$ as being the vector space

$$\mathbb{F}[[\mathcal{V}, \mathcal{V}^{-1}]]\langle \mathbf{x}^1 \rangle \oplus \mathbb{F}[[\mathcal{V}, \mathcal{V}^{-1}]]\langle \mathbf{y}^1 \rangle.$$

We have

$$\delta_1^1(\mathcal{V}^t \mathbf{x}^1) = (1 + \mathcal{V})\mathcal{V}^t \mathbf{y}^1 \otimes 1.$$

Furthermore,

$$\delta_2^1(i_1 \mathcal{V}^n U^m i_1, \mathcal{V}^t \mathbf{x}^1) = \mathcal{V}^{n+t} \mathbf{x}^1 \otimes U^m,$$

and similarly for \mathbf{y}^1 . Apply the forgetful functor to the type-A actions on $\mathcal{X}\mathcal{D}_\infty^{\mathbb{F}[U]}$ and $\mathcal{X}\mathcal{Z}_\infty^{\mathbb{F}[U]}$. Define type-D morphisms

$$i^1: \mathcal{Z}_\infty^{\mathbb{F}[U]} \rightarrow \mathcal{D}_\infty^{\mathbb{F}[U]}, \quad \pi^1: \mathcal{D}_\infty^{\mathbb{F}[U]} \rightarrow \mathcal{Z}_\infty^{\mathbb{F}[U]}, \quad \text{and} \quad h^1: \mathcal{D}_\infty^{\mathbb{F}[U]} \rightarrow \mathcal{D}_\infty^{\mathbb{F}[U]}$$

via the following formulas. We set

$$i^1(\mathbf{z}^1) = \left(\sum_{i \in \mathbb{Z}} \mathcal{V}^i \right) \mathbf{x}^1 \otimes 1, \quad \pi^1(\mathcal{V}^i \mathbf{x}^1) = \begin{cases} \mathbf{z}^1 \otimes 1 & \text{if } i = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Additionally, we set

$$\pi^1(\mathcal{V}^i \mathbf{y}^1) = 0 \quad \text{and} \quad h^1(\mathcal{V}^i \mathbf{y}^1) = \begin{cases} \sum_{j < i} \mathcal{V}^j \mathbf{x}^1 \otimes 1 & \text{if } i \leq 0, \\ \sum_{j \geq i} \mathcal{V}^j \mathbf{x}^1 \otimes 1 & \text{if } i > 0. \end{cases}$$

We leave it to the reader to check that the above maps are continuous, and furthermore

$$\begin{aligned} \pi^1 \circ i^1 &= \text{id}, & i^1 \circ \pi^1 &= \partial_{\text{Mor}}(h^1), & h^1 \circ h^1 &= 0, \\ \pi^1 \circ h^1 &= 0, & \text{and} & & h^1 \circ i^1 &= 0. \end{aligned}$$

Applying the homological perturbation lemma from Lemma 4.2 induces a type-DA Alexander bimodule structure on $\mathcal{Z}_\infty^{\mathbb{F}[U]}$ which is homotopy equivalent to $\mathcal{X}\mathcal{D}_\infty^{\mathbb{F}[U]}$. This bimodule structure is easily checked to coincide with our definition of $\mathcal{X}\mathcal{Z}_\infty^{\mathbb{F}[U]}$, completing the proof. \blacksquare

18.3 The surgery exact triangle

We now show that our bordered invariants recover the surgery exact triangle, in the following sense.

Proposition 18.8. *There is a type-D morphism $f^1: \mathcal{D}_n^{\mathcal{K}} \rightarrow \mathcal{D}_{n+1}^{\mathcal{K}}$ such that*

$$\text{Cone}(f^1: \mathcal{D}_n^{\mathcal{K}} \rightarrow \mathcal{D}_{n+1}^{\mathcal{K}}) \simeq \mathcal{D}_\infty^{\mathcal{K}},$$

for any $n \in \mathbb{Z}$.

Remark 18.9. This is unsurprising since the knot and link surgery formulas are proven using surgery exact triangles, though it highlights the similarities with the original bordered theory, cf. [27, Section 11.2].

Remark 18.10. It is straightforward to adapt our proof to construct morphisms and homotopy equivalences

$$\mathcal{D}_{n+1}^{\mathcal{K}} \simeq \text{Cone}(g^1: \mathcal{D}_\infty^{\mathcal{K}} \rightarrow \mathcal{D}_n^{\mathcal{K}}) \quad \text{and} \quad \mathcal{D}_n^{\mathcal{K}} \simeq \text{Cone}(h^1: \mathcal{D}_{n+1}^{\mathcal{K}} \rightarrow \mathcal{D}_\infty^{\mathcal{K}}).$$

Proof. Write $\mathbf{x}_0, \mathbf{x}_1$ for the generators of $\mathcal{D}_n^{\mathcal{K}}$ and $\mathbf{y}_0, \mathbf{y}_1$ for the generators of $\mathcal{D}_{n+1}^{\mathcal{K}}$. We define f^1 via the formula

$$f^1(\mathbf{x}_\varepsilon) = \mathbf{y}_\varepsilon \otimes (1 + \alpha)$$

for $\varepsilon \in \{0, 1\}$, where

$$\alpha = \sum_{s \geq 1} (\mathcal{U}^s + \mathcal{V}^s) \mathcal{U}^{s(s-1)/2} = (\mathcal{U} + \mathcal{V}) + (\mathcal{U}^2 + \mathcal{V}^2)\mathcal{U} + (\mathcal{U}^3 + \mathcal{V}^3)\mathcal{U}^3 + \dots$$

(The element α is discovered by writing $\beta = 1 + \alpha$ as $\beta = \sum_{i \in \mathbb{Z}} \beta_i$ where β_i has Alexander grading i , and then recursively solving the equation $\mathcal{V}\phi^\tau(\beta_i) = \beta_{i+1}$ starting with $\beta_0 = 1$.)

It is straightforward to see that $\partial_{\text{Mor}}(f^1) = 0$. Write:

$$\text{Cone}(f^1) = \begin{array}{ccc} \mathbf{x}_0 & \xrightarrow{1+\alpha} & \mathbf{y}_0 \\ \downarrow \sigma + \mathcal{V}^n \tau & & \downarrow \sigma + \mathcal{V}^{n+1} \tau \\ \mathbf{x}_1 & \xrightarrow{1+\alpha} & \mathbf{y}_1 \end{array}$$

Since $1 + \alpha$ is a unit in the completion of $\mathbf{I}_0 \cdot \mathcal{K} \cdot \mathbf{I}_0$, the above complex is homotopy equivalent to the bottom row, $\text{Cone}(1 + \alpha : \mathbf{x}_1 \rightarrow \mathbf{y}_1)$. To construct a homotopy equivalence between $\text{Cone}(1 + \alpha : \mathbf{x}_1 \rightarrow \mathbf{y}_1)$ and $\mathcal{D}_\infty^{\mathcal{K}}$, it is sufficient to show that $1 + \alpha = (1 + \mathcal{V})u$ for some unit u in the completion of $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$. To see this, we

recall that $\mathcal{U} = U\mathcal{V}^{-1}$. We use this expression for \mathcal{U} in our formula for α and then regroup terms based on the powers of U , and we obtain the following expression:

$$1 + \alpha = \sum_{s \geq 1} (\mathcal{V}^s + \mathcal{V}^{-s+1}) U^{s(s-1)/2}.$$

We observe that

$$(\mathcal{V}^s + \mathcal{V}^{-s+1}) U^{s(s-1)/2} = (1 + \mathcal{V}^{2s-1}) \mathcal{V}^{-s+1} U^{s(s-1)/2},$$

which clearly has a factor of $(1 + \mathcal{V})$, for all $s \geq 0$. Hence, $1 + \alpha = (1 + \mathcal{V})(1 + Uf)$ for some series f , and $(1 + Uf)$ is a unit in the completion of $\mathbf{I}_1 \cdot \mathcal{K} \cdot \mathbf{I}_1$, completing the proof. ■

18.4 Finite generation

In this section, we prove a basic property of the surgery modules that is useful in applications.

Proposition 18.11. *Suppose that $L \subseteq S^3$ is a link with framing Λ , and $L = K_1 \cup \dots \cup K_\ell$. Suppose further that $0 < j < \ell$ and write $L_{1,\dots,j} = K_1 \cup \dots \cup K_j$ and $L_{j+1,\dots,\ell} = K_{j+1} \cup \dots \cup K_\ell$. Let $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}_{\ell-j}}$ denote the type-D module obtained by tensoring $\mathcal{X}\mathcal{D}_0$ into each algebra component for $L_{1,\dots,j}$. Then $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}_{\ell-j}}$ is homotopy equivalent to a finitely generated type-D module.*

Our argument is similar to Manolescu and Ozsváth’s truncation procedure [32, Section 10]. We will prove the claim in several steps. If $j < \ell$, we write $\Lambda_{(j)}$ for the induced framing matrix on $L_{1,\dots,j}$. We will first verify the claim when $\Lambda_{(j)}$ is positive definite. We will then use the exact triangle from Section 18.3 to verify the claim for arbitrary framings.

If $0 < j < \ell$ and we decompose L as $L = L_{1,\dots,j} \cup L_{j+1,\dots,\ell}$, then we may consider the subcube

$$\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}_{\ell-j}}$$

generated by points of the cube \mathbb{E}_ℓ which have $L_{j+1,\dots,\ell}$ components 0. We will view this as a type-D module over the algebra $\mathbb{F}[\mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$ which is continuous with respect to the topology induced by

$$\mathbf{E}_0 \cdot \mathcal{L}_{\ell-j} \cdot \mathbf{E}_0 \cong \mathbb{F}[\mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell].$$

Recall from Section 8.2 that the solid torus module $\mathcal{X}\mathcal{D}_0$ may be decomposed as a tensor product

$$\mathcal{X}[\mathcal{D}_0]^{\mathbb{F}[U]} \hat{\boxtimes}_{\mathbb{F}[U]} \mathbb{F}[U].$$

Hence, there is a type-D module $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathbb{F}[U_1,\dots,U_j] \otimes \mathcal{L}_{n-j}}$, from which we get $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}_{n-j}}$ by tensoring with $\mathbb{F}[U_1,\dots,U_j]^{\mathbb{F}[U_1, \dots, U_j]}$.

Lemma 18.12. *Suppose $L = K_1 \cup \dots \cup K_\ell$ is a link in S^3 , and $0 \leq j < \ell$. Suppose that $\Lambda_{(j)}$ is positive definite. Then*

$$\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathbb{F}[U_1,\dots,U_j, \mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]}$$

is homotopy equivalent to a finitely generated type-D module. Furthermore, the maps appearing in the homotopy equivalence may be taken to be continuous if we equip $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})$ with the product topology over a monomial basis, and we equip

$$\mathbb{F}[U_1, \dots, U_j, \mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$$

with the (U_1, \dots, U_ℓ) -adic topology.

Remark 18.13. The fact that the morphisms in the homotopy equivalence may be taken to be continuous with respect to the (U_1, \dots, U_ℓ) -adic topology, as opposed to being continuous with respect to the $(U_1, \dots, U_j, \mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell)$ -adic topology, allows us to use the homotopy equivalence constructed in the lemma for the complex at each subcube $\mathbb{E}_j \times \{\varepsilon\}$, for all $\varepsilon \in \mathbb{E}_{n-j}$.

Proof. Our proof will be by induction on j . The case that $j = 0$ is automatic since $\mathcal{X}_\Lambda^{(*,0)}(\emptyset, L_{1,\dots,\ell}) \cong \mathcal{CF}\mathcal{L}(L)$, which is a finitely generated type-D module over $\mathbb{F}[\mathcal{U}_1, \mathcal{V}_1, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$.

We define a shifted Alexander A_j -grading on $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})$, as follows. We set

$$\omega(\mathbf{s}) = -(\text{lk}(K_1, K_j), \dots, \text{lk}(K_{j-1}, K_j)) \Lambda_{(j-1)}^{-1} (s_1, \dots, s_{j-1})^T. \tag{18.8}$$

In the above, $\Lambda_{(j-1)}^{-1}$ is the inverse matrix (over \mathbb{Q}) and T denotes transpose. Also, if $\mathbf{s} \in \mathbb{H}(L)$ we are writing $\mathbf{s} = (s_1, \dots, s_\ell)$. On the link surgery formula, we define the ω -shifted Alexander grading

$$A_j^\omega(\mathbf{s}) := s_j + \omega(\mathbf{s}).$$

We observe that if $i \neq j$, then

$$\omega(\mathbf{s} + \Lambda_{L,-K_i}) - \omega(\mathbf{s}) = -\text{lk}(K_i, K_j), \tag{18.9}$$


where $\Lambda_{L,-K_i}$ (as defined in Section 6.1) consists of the vector whose j -th component is $\text{lk}(K_i, K_j)$ if $i \neq j$, and whose i -th component is λ_i , the framing of K_i . Equation (18.9) follows from the fact that the i -th column of $\Lambda_{(j-1)}$ is obtained by deleting the last component of $\Lambda_{L,-K_i}$. Equation (18.9) verifies that if $\vec{M} \subseteq L_{1,\dots,j}$ is a sublink which does not contain $-K_j$, then $\Phi^{\vec{M}}$ preserves A_j^ω .

On the other hand, we compute directly from equation (18.8) and the definition of $\Lambda_{L,-K_j}$ that

$$\begin{aligned} & \lambda_j + \omega(\mathbf{s} + \Lambda_{L,-K_j}) - \omega(\mathbf{s}) \\ &= \lambda_j - (\text{lk}(K_1, K_j), \dots, \text{lk}(K_{j-1}, K_j)) \Lambda_{(j-1)}^{-1} (\text{lk}(K_1, K_j), \dots, \text{lk}(K_{j-1}, K_j))^T. \end{aligned} \tag{18.10}$$

Note that the above quantity, which we denote by λ'_j , is the Schur complement of $\Lambda_{(j-1)}$ inside of $\Lambda_{(j)}$. In particular, $\Lambda_{(j)}$ is positive definite if and only if $\Lambda_{(j-1)}$ is positive definite and $\lambda'_j > 0$. It follows from equation (18.10) that if $-K_j \subseteq \bar{M} \subseteq L_{1,\dots,j}$, then $\Phi^{\bar{M}}$ shifts A_j^ω by λ'_j .

We may thus decompose the complex $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})$ as a sum of staircases:



The vertical direction is the direction of K_j . Here, we are writing F^{-K_j} for the hypercube maps for all oriented sublinks $\bar{M} \subseteq L_{1,\dots,j}$ which contain $-K_j$, and we define F^{K_j} similarly. In the above, we are writing $\mathcal{A}_s^\omega \subseteq \mathcal{C}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})$ for the subspace of $\mathbb{F}[U_1, \dots, U_j, \mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$ -module generators with K_j -component 0 and Alexander A_j^ω -grading $s \in \mathbb{Q}$. The subspace \mathcal{B}_s^ω is similar, but has K_j -component 1. Since $\mathbb{H}(L_{1,\dots,j})/\Lambda_{(j)}$ is finite, there are only finitely many such staircases. Here, the subscript s of \mathcal{A}_s^ω and \mathcal{B}_s^ω denotes their shifted Alexander grading.

Standard arguments (see [32, Lemma 10.1]) imply the following:

- (h-1) If s is sufficiently large, the map F^{K_j} is a homotopy equivalence between \mathcal{A}_s^ω and \mathcal{B}_s^ω .
- (h-2) If s is sufficiently negative, the map F^{-K_j} is a homotopy equivalence between \mathcal{A}_s^ω and $\mathcal{B}_{s+\lambda'_j}^\omega$.

Note that we may view $\mathcal{A}_s^\omega \otimes \mathbb{F}[U_j]$ as being a subspace of $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j-1}, L_{j+1,\dots,\ell}) \otimes \mathbb{F}[\mathcal{U}_j, \mathcal{V}_j]$ in A_j^ω -Alexander grading s . Similarly, we may view $\mathcal{B}_s^\omega \otimes \mathbb{F}[U_j]$ as being the subspace of $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j-1}, L_{j+1,\dots,\ell}) \otimes \mathbb{F}[\mathcal{U}_j, \mathcal{V}_j, \mathcal{V}_j^{-1}]$ in A_j^ω -Alexander grading s . By induction, $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j-1}, L_{j+1,\dots,\ell})$ is homotopy equivalent to a finitely generated type-D module over $\mathbb{F}[U_1, \dots, U_{j-1}, \mathcal{U}_j, \mathcal{V}_j, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$ with respect to the topologies in the statement of the lemma.

In particular, we may replace $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j-1}, L_{j+1,\dots,\ell})$ with a finitely generated model \mathcal{C}' , and also replace the subcomplexes \mathcal{A}_s^ω and \mathcal{B}_s^ω with the finite-dimensional

subspaces of $\mathcal{C}' \otimes \mathbb{F}[U_j, \mathcal{V}_j]$ and $\mathcal{C}' \otimes \mathbb{F}[U_j, \mathcal{V}_j, \mathcal{V}_j^{-1}]$ which form an $\mathbb{F}[U_j]$ basis in Alexander grading s . Properties (h-1) and (h-2) are preserved by homotopy equivalence, and hence still persist on the finitely generated models. Abusing notation, we write \mathcal{A}_s^ω and \mathcal{B}_s^ω also for the finitely generated models.

Using the same logic as for the ordinary knot surgery mapping cone formula [43], each such staircase above admits a finitely generated truncation. We need additionally to show that the maps appearing in a homotopy equivalence with the finite truncation are continuous with respect to the topologies in the main statement. These maps may be described easily using homological perturbation theory. Compare the proof of Lemma 18.2. Since \mathcal{C}' is finitely generated, for $s \gg 0$, each generator of \mathcal{A}_s^ω must be a multiple of $\mathcal{V}_j^{r(s)}$ for some function $r: \mathbb{Q} \rightarrow \mathbb{N}$ such that $r(s) \rightarrow \infty$ as $s \rightarrow \infty$. Similarly, for $s \ll 0$, each generator of \mathcal{A}_s^ω must be a multiple of $\mathcal{U}_j^{l(s)}$ for some $l: \mathbb{Q} \rightarrow \mathbb{N}$ such that $l(s) \rightarrow \infty$ as $s \rightarrow -\infty$. Using the equivariance of the maps F^{K_j} and F^{-K_j} from Lemma 6.7, we see that for $s \ll 0$, the map F^{K_j} sends \mathcal{A}_s^ω into the ideal $(U_j^{l(s)})$. Similarly, for $s \gg 0$, F^{-K_j} maps \mathcal{A}_s^ω into the ideal $(U_j^{r(s)})$. This is sufficient to show that the maps appearing in the homotopy equivalence are continuous as stated.

Since $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,j})$ decomposes as a sum of finitely many staircases, each of which is homotopy equivalent to a finitely generated truncation, as above, the direct sum is homotopy equivalent to a finitely generated complex, completing the proof. ■

We now prove a basic lemma concerning homological algebra.

Lemma 18.14. *Write $\mathcal{R} = \mathbb{F}[U_1, \dots, U_{j-1}, U_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$, where $0 < j < \ell$. Write $\mathcal{R}[U_j]$ for $\mathcal{R} \otimes \mathbb{F}[U_j]$. Suppose that $\mathcal{X}^{\mathcal{R}[U_j]}$ is a finitely generated type-D module which admits a relative $\mathbb{Q} \times \mathbb{Q}$ -valued $(\text{gr}_w, \text{gr}_z)$ -bigrading. Suppose that $U_j \simeq \mathcal{U}_\ell \mathcal{V}_\ell$ as type-D morphisms (i.e., there is a morphism J^1 satisfying $\partial_{\text{Mor}}(J^1) = \text{id} \otimes (U_j + \mathcal{U}_\ell \mathcal{V}_\ell)$). Then*

$$(\mathcal{X}^{\mathcal{R}[U_j]} \boxtimes_{\mathbb{F}[U_j]} \mathbb{F}[U_j])^{\mathcal{R}}$$

is homotopy equivalent to a finitely generated type-D module over \mathcal{R} .

Remark 18.15. The condition on the $(\text{gr}_w, \text{gr}_z)$ -bigrading is likely not absolutely necessary, though it simplifies the discussion on completions.

Proof. Our proof uses homological perturbation theory. We view $\mathcal{X}^{\mathcal{R}[U_j]}$ as a being obtained from a DA bimodule ${}_{\mathcal{E}}\mathcal{W}^{\mathcal{R}[U_j]}$ where \mathcal{E} denotes the exterior algebra on one generator. We define $\delta_{k+1}^1(\theta, \dots, \theta, \mathbf{x})$ on \mathcal{W} to consist of the components of δ^1 of \mathcal{X} which are weighted by an algebra element of total \mathcal{R} -degree equal to k . Here, the \mathcal{R} -degree of a monomial denotes the sum of powers of variables from \mathcal{R} appearing (but not of powers of U_j). It is straightforward to see that \mathcal{W} satisfies the DA bimodule structure relations.

Note that δ_1^1 on \mathcal{W} consists of exactly the terms of the differential of \mathcal{X} which are weighted by elements of $\mathbb{F}[U_j]$.

Let N^ε be the rank 1 type-D module with $\delta^1(1) = 1 \otimes \theta$, so that

$$N^\varepsilon \boxtimes_\varepsilon \mathcal{W}^{\mathcal{R}[U_j]} = \mathcal{X}^{\mathcal{R}[U_j]}.$$

We consider the type-D module $\mathcal{Y}^{\mathcal{R}[U_j]}$ obtained by applying the forgetful functor to

$${}_\varepsilon \mathcal{W}^{\mathcal{R}[U_j]}.$$

The type-D module $\mathcal{Y}^{\mathcal{R}[U_j]}$ has δ^1 with algebra elements in $\mathbb{F}[U_j]$, i.e., we may view it as a finitely generated free chain complex over $\mathbb{F}[U_j]$. Setting all variables except U_j equal to 0, the map J^1 induces a map h^1 on $\mathcal{Y}^{\mathcal{R}[U_j]}$ which satisfies $\partial_{\text{Mor}}(j^1) = \text{id} \otimes U_j$. The classification theorem for finitely generated chain complexes over a PID implies that we may find a basis of \mathcal{Y} so that $\mathcal{Y}^{\mathcal{R}[U_j]}$ decomposes as a direct sum of 1-step complexes (i.e., generators with vanishing δ^1), and 2-step complexes (i.e., summands generated by two generators with $\delta^1(\mathbf{y}) = 0$ and $\delta^1(\mathbf{x}) = \mathbf{y} \otimes \alpha$). Our assumptions about the existence of a relative bigrading imply that α may always be taken to be a power of U_j . The fact that $\partial_{\text{Mor}}(j^1) = \text{id} \otimes U_j$ implies that there are no 1-step summands, and each 2-step summand is of the form $\delta^1(\mathbf{x}) = \mathbf{y} \otimes U_j$.

We define the finite-dimensional subspace $\mathcal{Q} \subseteq \mathcal{W} \boxtimes \mathbb{F}[U_j]$ as follows. Enumerate the generators of the U_j^1 -weighted 2-step complexes as \mathbf{x}_t and \mathbf{y}_t , so that $\delta^1(\mathbf{x}_t) = \mathbf{y}_t \otimes U_j$, for t in some finite set. We define \mathcal{Q} to be the \mathbb{F} span of $\mathbf{y}_t \otimes 1$, ranging over all t . We view \mathcal{Q} as being a finitely generated type-D module over \mathcal{R} with vanishing differential. There are maps forming a homotopy equivalence of type-D modules

$$\pi^1: (\mathcal{W} \boxtimes \mathbb{F}[U_j])^{\mathcal{R}} \rightarrow \mathcal{Q}^{\mathcal{R}}, \quad i^1: \mathcal{Q}^{\mathcal{R}} \rightarrow (\mathcal{W} \boxtimes \mathbb{F}[U_j])^{\mathcal{R}},$$

and

$$h^1: (\mathcal{W} \boxtimes \mathbb{F}[U_j])^{\mathcal{R}} \rightarrow (\mathcal{W} \boxtimes \mathbb{F}[U_j])^{\mathcal{R}},$$

which satisfy the assumptions of the homological perturbation lemma stated in Lemma 4.2. The map π^1 sends $\mathbf{y}_t \otimes U_j^s$ to $\mathbf{y}_t \otimes 1$ if $s = 0$, and is zero on all generators. The map i^1 sends \mathbf{y}_t to $(\mathbf{y}_t \otimes 1) \otimes 1$. The map h^1 sends $\mathbf{y}_t \otimes U_j^s$ to $(\mathbf{y}_t \otimes U_j^{s-1}) \otimes 1$ if $s > 0$, and h^1 vanishes on all other generators. Applying the homological perturbation lemma endows $\mathcal{Q}^{\mathcal{R}}$ with a left action of ε ; and also yields a homotopy equivalence ${}_\varepsilon \mathcal{Q}^{\mathcal{R}} \simeq {}_\varepsilon \mathcal{W}^{\mathcal{R}[U_j]} \boxtimes_{\mathbb{F}[U_j]} \mathbb{F}[U_j]$. Note that the induced A_∞ -module structure is operationally bounded by grading considerations. Finally, boxing with the identity morphism of N^ε gives a homotopy equivalence as in the statement. ■

We are now able to prove finite generation in general.

Proof of Proposition 18.11. Suppose that $L = K_1 \cup \dots \cup K_\ell$ is a link in S^3 , and Λ is an integral framing on L . Let $0 \leq j < \ell$. It is sufficient to show that

$$\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathbb{F}[\mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]}$$

is homotopy equivalent to a finitely generated type-D module over $\mathbb{F}[\mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]$. This is sufficient since we may use this model at each point of the cube \mathbb{E}_{n-j} to obtain a finitely generated model of $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}^{n-j}}$.

By Proposition 13.2, changing the arc system on the components K_1, \dots, K_j does not change the homotopy type of $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}^{n-j}}$. Hence, we assume all of the arcs for K_1, \dots, K_j are alpha-parallel.

Lemma 18.12 verifies that if $\Lambda_{(j)}$ is positive definite, then

$$\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathbb{F}[U_1, \dots, U_j, \mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]}$$

is homotopy equivalent to a finitely generated type-D module. Furthermore, since $\Lambda_{(j)}$ is positive definite (in particular, non-singular), it is straightforward to see that this complex admits a relative $\mathbb{Q} \times \mathbb{Q}$ -valued $(\text{gr}_w, \text{gr}_z)$ -bigrading (compare [32, Section 9.3]). Lemma 13.31 implies that each U_1, \dots, U_j is chain homotopic to $\mathcal{U}_\ell \mathcal{V}_\ell$ as a type-D endomorphism on this subcube. (Note that on this subcube, the K_ℓ -direction is never incremented, so the arc for K_ℓ is irrelevant.) Lemma 18.14 now implies that $\mathcal{X}_\Lambda^{(*,0)}(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathbb{F}[\mathcal{U}_{j+1}, \mathcal{V}_{j+1}, \dots, \mathcal{U}_\ell, \mathcal{V}_\ell]}$ is homotopy equivalent to a finitely generated type-D module.

Since $\mathcal{X} \mathcal{D}_\infty^{\mathbb{F}[U]}$ is homotopy equivalent to a finitely generated type-DA module by Lemma 18.7, we observe that the claim still holds if we replace any number of framings in a positive definite $\Lambda_{(j)}$ with $+\infty$. The mapping cone of two finitely generated complexes is finitely generated, so the surgery exact triangle from Proposition 18.8 (cf. Remark 18.10), implies that $\mathcal{X}_\Lambda(L_{1,\dots,j}, L_{j+1,\dots,\ell})^{\mathcal{L}^{n-j}}$ is homotopy equivalent to a finitely generated type-D module for any framing, concluding the proof. ■

Appendix A

Splicing operations

In this chapter, we describe some topology related to splicing knot complements. Everything in this chapter is well known, and we only include it as a reference.

A.1 Splices and connected sums

Suppose that K_1 and K_2 are knots in S^3 . The complements of K_1 and K_2 are manifolds with torus boundary. The following is well known, though we give a proof for the convenience of the reader. Compare [7, Section 7].

Lemma A.1. *Suppose that λ_1 and λ_2 are integral framings on knots K_1 and K_2 in S^3 , and let $[\lambda_1]$ and $[\lambda_2]$ denote the induced elements of $H_1(\partial S^3 \setminus N(K_1))$ and $H_1(\partial S^3 \setminus N(K_2))$. Then $S^3_{\lambda_1 + \lambda_2}(K_1 \# K_2)$ is diffeomorphic to the 3-manifold obtained by gluing $S^3 \setminus N(K_1)$ to $S^3 \setminus N(K_2)$ using the diffeomorphism which sends*

$$[\lambda_1] \mapsto -[\lambda_2] \quad \text{and} \quad [\mu_1] \mapsto [\mu_2].$$

Proof. Consider the manifold Z obtained by gluing the complements of K_1 and K_2 as described above. We may describe Z alternatively by first attaching a 3-dimensional 1-handle to connect $S^3 \setminus \nu(K_1)$ and $S^3 \setminus \nu(K_2)$, and then attaching two 2-handles and one 3-handle. The first 2-handle is attached along $\mu_1 * \bar{\mu}_2$ (where $\bar{\mu}_2$ denotes μ_2 with orientation reversed), and the second 2-handle is attached along $\lambda_1 * \lambda_2$. In both cases, the curves are concatenated across the 1-handle.

We observe that joining $S^3 \setminus \nu(K_1)$ and $S^3 \setminus \nu(K_2)$ with the 1-handle and the 2-handle, tracing $\mu_1 * \bar{\mu}_2$ is the same as gluing $S^3 \setminus \nu(K_1)$ and $S^3 \setminus \nu(K_2)$ along meridians of K_1 and K_2 . This gives $S^3 \setminus \nu(K_1 \# K_2)$. Gluing the second 2-handle and then the 3-handle is the same as performing Dehn surgery on $K_1 \# K_2$ with framing $\lambda_1 + \lambda_2$. ■

Remark A.2. The argument above also holds for arbitrary knots in 3-manifolds, as long as the framings are Morse.

Remark A.3. The same argument works for links. Suppose that L_1 and L_2 are links in S^3 with framings Λ_1 and Λ_2 . Suppose that K_1 and K_2 are two chosen components of L_1 and L_2 , respectively. Then

$$S^3_{\Lambda_1 + \Lambda_2}(L_1 \# L_2) \cong M_{\Lambda_1}(L_1, K_1) \cup_{\phi} M_{\Lambda_2}(L_2, K_2),$$

where $M_{\Lambda_i}(L_i, K_i)$ is the manifold obtained by surgering along $L_i - K_i$, and removing a neighborhood of K_i . The diffeomorphism ϕ which identifies $\partial N(K_1)$ with $\partial N(K_2)$ sends λ_1 to $-\lambda_2$, and μ_1 to μ_2 . The framing $\Lambda_1 + \Lambda_2$ is obtained by summing the framings for K_1 and K_2 , and leaving the remaining framings unchanged.

A.2 The Hopf link and changes of parametrization

We may naturally view the Hopf link as a link which changes the boundary parametrization. There are two ways to do this:

- (1) (Blow up) If K is the special component, and we connect the sum with a Hopf link, and leave K the special component. If we give the other component of the Hopf link framing ± 1 , then we change the framing on $\partial N(K)$ by a Dehn twist along the meridian.
- (2) (Rotating layer) Noting that the complement of the Hopf link is $\mathbb{T}^2 \times [0, 1]$, we may use the description from the previous remark to see the effect of a connect summing the Hopf link to the special component, and using the new Hopf link component as the special component. The effect is to change the boundary parametrization by composing with one of the two maps:

$$\mu \mapsto \lambda \quad \text{and} \quad \lambda \mapsto -\mu,$$

or

$$\mu \mapsto -\lambda \quad \text{and} \quad \lambda \mapsto \mu.$$

These two cases are distinguished by the sign of the Hopf link which is used in the connected sum.

References

- [1] M. F. Atiyah and I. G. Macdonald, *Introduction to commutative algebra*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969
- [2] A. Beilinson, [Remarks on topological algebras](#). *Mosc. Math. J.* **8** (2008), no. 1, 1–20, 183
- [3] D. Chen, I. Zemke, and H. Zhou, L-space satellite operations and knot Floer homology. 2024, arXiv:[2412.05755](#)
- [4] E. Eftekhary, Heegaard Floer homology and Morse surgery. 2006, arXiv:[math/0603171v2](#)
- [5] E. Eftekhary, [Floer homology and splicing knot complements](#). *Algebr. Geom. Topol.* **15** (2015), no. 6, 3155–3213
- [6] D. Eisenbud, *Commutative algebra*. Grad. Texts in Math. 150, Springer, New York, 1995
- [7] C. M. Gordon, [Dehn surgery and satellite knots](#). *Trans. Amer. Math. Soc.* **275** (1983), no. 2, 687–708
- [8] A. Grothendieck, Produits tensoriels topologiques et espaces nucléaires. In *Séminaire Bourbaki, Vol. 2*, pp. 193–200, Soc. Math. France, Paris, 1954
- [9] J. Hanselman, [Splicing integer framed knot complements and bordered Heegaard Floer homology](#). *Quantum Topol.* **8** (2017), no. 4, 715–748
- [10] M. Hedden and A. S. Levine, [Splicing knot complements and bordered Floer homology](#). *J. Reine Angew. Math.* **720** (2016), 129–154
- [11] M. Hedden and A. S. Levine, [A surgery formula for knot Floer homology](#). *Quantum Topol.* **15** (2024), no. 2, 229–336
- [12] K. Hendricks, J. Hom, M. Stoffregen, and I. Zemke, Surgery exact triangles and involutive Heegaard Floer homology. [v1] 2020, [v2] 2025, arXiv:[2011.00113v2](#)
- [13] K. Hendricks, J. Hom, M. Stoffregen, and I. Zemke, Naturality and functoriality in involutive Heegaard Floer homology. [v1] 2021, [v3] 2025, arXiv:[2201.12906v3](#)
- [14] K. Hendricks, J. Hom, M. Stoffregen, and I. Zemke, A dual knot mapping cone formula for involutive Heegaard Floer homology. [v1] 2022, [v2] 2025 arXiv:[2205.12798v2](#)
- [15] K. Hendricks and C. Manolescu, [Involutive Heegaard Floer homology](#). *Duke Math. J.* **166** (2017), no. 7, 1211–1299
- [16] J. Huebschmann and T. Kadeishvili, [Small models for chain algebras](#). *Math. Z.* **207** (1991), no. 2, 245–280
- [17] T. V. Kadeishvili, The algebraic structure in the homology of an $A(\infty)$ -algebra (in Russian). *Soobshch. Akad. Nauk Gruzin. SSR* **108** (1982), no. 2, 249–252
- [18] Ç. Karakurt, T. Lidman, and E. Tweedy, [Heegaard Floer homology and splicing homology spheres](#). *Math. Res. Lett.* **28** (2021), no. 1, 93–106
- [19] B. Keller, [Introduction to \$A\$ -infinity algebras and modules](#). *Homology Homotopy Appl.* **3** (2001), no. 1, 1–35

- [20] B. Keller, Addendum to: “Introduction to A -infinity algebras and modules” [Homology Homotopy Appl. **3** (2001), no. 1, 1–35]. *Homology Homotopy Appl.* **4** (2002), no. 1, 25–28
- [21] M. Kontsevich and Y. Soibelman, Homological mirror symmetry and torus fibrations. In *Symplectic geometry and mirror symmetry (Seoul, 2000)*, pp. 203–263, World Sci. Publ., River Edge, NJ, 2001
- [22] S. Lefschetz, *Algebraic Topology*. Amer. Math. Soc. Colloq. Publ., Vol. 27, American Mathematical Society, New York, 1942
- [23] R. Lipshitz, A cylindrical reformulation of Heegaard Floer homology. *Geom. Topol.* **10** (2006), 955–1096
- [24] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, Bordered Floer homology and the spectral sequence of a branched double cover I. *J. Topol.* **7** (2014), no. 4, 1155–1199
- [25] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, Bimodules in bordered Heegaard Floer homology. *Geom. Topol.* **19** (2015), no. 2, 525–724
- [26] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, Bordered Floer homology and the spectral sequence of a branched double cover II: the spectral sequences agree. *J. Topol.* **9** (2016), no. 2, 607–686
- [27] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, Bordered Heegaard Floer homology. *Mem. Amer. Math. Soc.* **254** (2018), no. 1216
- [28] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston, Diagonals and A -infinity tensor products. *Studia Sci. Math. Hungar.* **62** (2025), no. 2–3, 97–304
- [29] R. Lipshitz and S. Sarkar, Khovanov homology detects split links. *Amer. J. Math.* **144** (2022), no. 6, 1745–1781
- [30] Y. Liu, Heegaard Floer homology of surgeries on two-bridge links. *Algebr. Geom. Topol.* **22** (2022), no. 2, 473–557
- [31] J.-L. Loday, The diagonal of the Stasheff polytope. In *Higher structures in geometry and physics*, pp. 269–292, Progr. Math. 287, Birkhäuser/Springer, New York, 2011
- [32] C. Manolescu and P. Ozsváth, Heegaard Floer homology and integer surgeries on links. *Geom. Topol.* **29** (2025), no. 6, 2783–3062
- [33] M. Markl and S. Shnider, Associahedra, cellular W -construction and products of A_∞ -algebras. *Trans. Amer. Math. Soc.* **358** (2006), no. 6, 2353–2372
- [34] A. Némethi, On the Ozsváth–Szabó invariant of negative definite plumbed 3-manifolds. *Geom. Topol.* **9** (2005), 991–1042
- [35] A. Némethi, Lattice cohomology of normal surface singularities. *Publ. Res. Inst. Math. Sci.* **44** (2008), no. 2, 507–543
- [36] Y. Ni, Homological actions on sutured Floer homology. *Math. Res. Lett.* **21** (2014), no. 5, 1177–1197
- [37] P. Ozsváth and Z. Szabó, Holomorphic disks and knot invariants. *Adv. Math.* **186** (2004), no. 1, 58–116
- [38] P. Ozsváth and Z. Szabó, Holomorphic disks and three-manifold invariants: properties and applications. *Ann. of Math. (2)* **159** (2004), no. 3, 1159–1245

- [39] P. Ozsváth and Z. Szabó, [Holomorphic disks and topological invariants for closed three-manifolds](#). *Ann. of Math. (2)* **159** (2004), no. 3, 1027–1158
- [40] P. Ozsváth and Z. Szabó, [Holomorphic triangles and invariants for smooth four-manifolds](#). *Adv. Math.* **202** (2006), no. 2, 326–400
- [41] P. Ozsváth and Z. Szabó, [Holomorphic disks, link invariants and the multi-variable Alexander polynomial](#). *Algebr. Geom. Topol.* **8** (2008), no. 2, 615–692
- [42] P. Ozsváth and Z. Szabó, [Kauffman states, bordered algebras, and a bigraded knot invariant](#). *Adv. Math.* **328** (2018), 1088–1198
- [43] P. S. Ozsváth and Z. Szabó, [Knot Floer homology and integer surgeries](#). *Algebr. Geom. Topol.* **8** (2008), no. 1, 101–153
- [44] P. S. Ozsváth and Z. Szabó, [Knot Floer homology and rational surgeries](#). *Algebr. Geom. Topol.* **11** (2011), no. 1, 1–68
- [45] P. Ozsváth and Z. Szabó, [Algebras with matchings and knot Floer homology](#). 2019, arXiv:1912.01657
- [46] L. Positelski, [Exact categories of topological vector spaces with linear topology](#). *Mosc. Math. J.* **24** (2024), no. 2, 219–286
- [47] J. A. Rasmussen, [Floer homology and knot complements](#). Ph.D. thesis, Harvard University, 2003
- [48] S. Sanedidze and R. Umble, [Diagonals on the permutahedra, multiplihedra and associahedra](#). *Homology Homotopy Appl.* **6** (2004), no. 1, 363–411
- [49] S. Sarkar, [Moving basepoints and the induced automorphisms of link Floer homology](#). *Algebr. Geom. Topol.* **15** (2015), no. 5, 2479–2515
- [50] P. Seidel, [Fukaya categories and Picard–Lefschetz theory](#). Zur. Lect. Adv. Math., European Mathematical Society (EMS), Zürich, 2008
- [51] J. D. Stasheff, [Homotopy associativity of H-spaces. I](#). *Trans. Am. Math. Soc.* **108** (1963), no. 2, 275–292
- [52] J. D. Stasheff, [Homotopy associativity of H-spaces. II](#). *Trans. Am. Math. Soc.* **108** (1963), no. 2, 293–312
- [53] I. Zemke, [Quasistabilization and basepoint moving maps in link Floer homology](#). *Algebr. Geom. Topol.* **17** (2017), no. 6, 3461–3518
- [54] I. Zemke, [Link cobordisms and absolute gradings on link Floer homology](#). *Quantum Topol.* **10** (2019), no. 2, 207–323
- [55] I. Zemke, [Link cobordisms and functoriality in link Floer homology](#). *J. Topol.* **12** (2019), no. 1, 94–220
- [56] I. Zemke, [The equivalence of lattice and Heegaard Floer homology](#). *Duke Math. J.* **174** (2025), no. 5, 857–910
- [57] I. Zemke, [A general Heegaard Floer surgery formula](#). *Adv. Math.* **480** (2025), Part B, article no. 110489
- [58] I. Zemke, [Graph cobordisms and Heegaard Floer homology](#). *Geom. Topol.* **30** (2026), no. 2, 389–528

Ian Zemke

Bordered Manifolds with Torus Boundary and the Link Surgery Formula

In this memoir, we develop a theory of bordered HF^- using the link surgery formula of Manolescu and Ozsváth. We interpret their link surgery complexes as type-D modules over an associative algebra \mathcal{K} , which we introduce. We prove a connected sum formula, which we interpret as an A_∞ -tensor product over our algebra \mathcal{K} . Topologically, this connected sum formula may be viewed as a formula for gluing along torus boundary components.

We discuss several important examples. As a basic example, if K_1 and K_2 are knots in S^3 , and Y is obtained by gluing the complements of K_1 and K_2 together using an orientation reversing diffeomorphism of their boundaries, then our theory may be used to compute $\text{CF}^-(Y)$ from $\text{CFK}^\infty(K_1)$ and $\text{CFK}^\infty(K_2)$. By computing the type-D modules for rationally framed solid tori, our theory gives a version of the link surgery formula for rationally framed links. As a final example, we use our theory to derive the Heegaard Floer homology of all 3-manifolds which bound the plumbing of a tree of disk bundles over 2-spheres.

<https://ems.press>

ISSN 2747-9080

ISBN 978-3-98547-110-2



EM
S 
PRESS