UCLA
UCLA Electronic Theses and Dissertations

Title
TQFT structures in Heegaard Floer homology

Permalink
https://escholarship.org/uc/item/46c1h5j3

Author
Zemke, Ian Michael

Publication Date
2017

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA
Los Angeles

TQFT constructions in Heegaard Floer homology

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Ian Michael Zemke

2017

© Copyright by
Ian Michael Zemke
2017

ABSTRACT OF THE DISSERTATION

TQFT constructions in Heegaard Floer homology

by

Ian Michael Zemke
Doctor of Philosophy in Mathematics
University of California, Los Angeles, 2017
Professor Ciprian Manolescu, Chair

In the early 2000s, Ozsváth and Szabó introduced a collection of invariants for 3–manifolds and 4–manifolds called Heegaard Floer homology. To a 3–manifold they constructed a group, and to a 4–manifold which cobounds two 3–manifolds, they constructed a homomorphism between the manifolds appearing on the ends. Their invariants satisfy many of the axioms of a TQFT as described by Atiyah, however their construction has some additional restrictions which prevent it from fitting into Atiyah's framework. There is a refinement of Heegaard Floer homology for 3–manifolds containing a knot, due to Ozsváth and Szabó, and independently Rasmussen, and a further refinement for 3–manifolds containing links, due to Ozsváth and Szabó. It's a natural question as to whether one can define functorial maps associated to link cobordisms.

In this thesis, we describe a package of cobordism maps for Heegaard Floer homology and link Floer homology. The cobordism maps satisfy an appropriate analogy of the axiomatic description of a TQFT formulated by Atiyah. To a ribbon graph cobordism between two based 3–manifolds, we associate a map between the Heegaard Floer homologies of the ends. To a decorated link cobordism, we obtain maps on the link Floer homologies of the ends. The maps associated to decorated link cobordisms reduce to the maps for ribbon graphs, in a natural way.

As applications, we describe several formulas for mapping class group actions on the Hee-

gaard Floer and knot Floer groups. We prove a new bound on a concordance invariant $\Upsilon_K(t)$ from knot Floer homology, and also see how the link cobordism maps give straightforward proofs of other bounds on concordance invariants from knot Floer homology. We also explore the interaction of the maps with conjugation actions on Heegaard Floer homology and link Floer homology, giving connected sum formulas for involutive Heegaard Floer homology and involutive knot Floer homology.

The dissertation of Ian Michael Zemke is approved.

Michael Gutperle

Ko Honda

Ciprian Manolescu, Committee Chair

University of California, Los Angeles

2017

To the wonderful landscape of Joshua Tree National Park...

Table of Contents

1 Introduction **1**

1.1 What is a TQFT? 2

1.2 A ribbon graph TQFT for multi-based manifolds 5

1.3 A TQFT for link Floer homology 18

2 Graph cobordisms and Heegaard Floer homology **36**

2.1 Introduction 36

2.2 Organization 43

2.3 Preliminaries on Heegaard Floer homology 45

2.4 Homological actions on the Heegaard Floer groups 63

2.5 Free stabilization maps 78

2.6 Graph action map 107

2.7 1– and 3–handles 123

2.8 2–handle maps 137

2.9 Invariance of path cobordism maps 148

2.10 Construction and invariance of the graph cobordism maps when $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ 156

2.11 The composition law for $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ 168

2.12 Cobordisms with $\pi_0(Y_i) \not\twoheadrightarrow \pi_0(W)$ 181

2.13 Moving basepoints, the π_1–action, and the relation between the graph and path cobordism maps 188

3 Quasi-stabilization and basepoint moving maps in link Floer homology . **232**

3.1 Introduction 232

3.2 Background, the complexes CFL^∞_{UV}, and $\mathfrak{P}$–filtrations 237
3.3 The maps Φ_w and Ψ_z 245
3.4 Cylindrical boundary degenerations 248
3.5 Preliminaries on the quasi-stabilization operation 249
3.6 Quasi-stabilization and triangle maps 264
3.7 Invariance of the quasi-stabilization maps 276
3.8 Commutation of quasi-stabilization maps 281
3.9 Further relations between the maps Ψ_z, Φ_w and $S^{\pm}_{w,z}$ 286
3.10 Basepoint moving maps 290

4 Link cobordisms and functoriality in link Floer homology 301

4.1 Introduction 301
4.2 Decorated link cobordisms, the complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ and the coherent chain homotopy type invariant $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ 314
4.3 Background on $\mathcal{C}FL^\circ$, quasi-stabilization, basepoint maps, and the disk stabilization maps 321
4.4 Maps for handles attached away from L 351
4.5 Quasi-stabilization and triangle maps 359
4.6 Definition of the band maps 367
4.7 Properties and relations involving the band maps 375
4.8 Maps for framed 0–spheres along L 402
4.9 Parametrized Kirby decompositions of link cobordisms 408
4.10 Morse theory and moves between parametrized Kirby decompositions 413
4.11 Construction and invariance of the cobordism maps 445
4.12 Maps for cobordisms where $\pi_0(\mathbb{L}_i) \to \pi_0(\Sigma)$ is not surjective 463

4.13 The composition law 467

4.14 The graph TQFT as the reduction of the link Floer TQFT 471

5 Link cobordisms and absolute gradings on link Floer homology 492

5.1 Introduction 492

5.2 Background and definitions 504

5.3 Heegaard triples and Spin^c structures 509

5.4 Kirby calculus for manifolds with boundary 518

5.5 Relative gradings 526

5.6 Coherent gradings on coherent chain homotopy types 532

5.7 Absolute Alexander multi-gradings 541

5.8 Absolute Maslov gradings 558

5.9 Link cobordisms and grading change formulas 561

5.10 Conjugation symmetry, equivalence with other constructions, and collapsing gradings 572

5.11 A new proof of a bound of Ozsváth and Szabó on τ 576

5.12 t-modified knot Floer homology and a bound on the $\Upsilon_K(t)$ invariant 578

5.13 The adjunction relation 586

5.14 The maps on $\mathcal{H}FL^\infty$ for surfaces in negative definite 4–manifolds 595

6 Connected sums and involutive knot Floer homology 599

6.1 Introduction 599

6.2 The group $\mathfrak{I}_K$ 609

6.3 Background on link Floer homology and the link Floer TQFT 622

6.4 Further properties of the link Floer TQFT 632

6.5 Link cobordisms and connected sums 643

6.6 The map ι_K on connected sums 658

6.7 The homomorphism $\mathcal{C} \to \mathfrak{I}_K$ 661

6.8 Computations and examples 663

6.9 Alternate proofs that G_i and E_i are homotopy inverses 669

List of Figures

1.1 Cobordisms in $(1+1)$–dimension closed string and open string TQFTs . . . 4

1.2 The "broken path" graph cobordism induces the map Φ_w. 12

1.3 The graph cobordism for a connected sum. 14

1.4 A trivalent graph. 17

1.5 A ribbon surface and its ribbon graph core. 26

1.6 Decorated link cobordisms, and ribbon graph cores. 26

1.7 Some relations for the graph TQFT 28

1.8 Computing the reduction of the link Floer TQFT on a genus 1 surface . . . 28

1.9 Decorated surfaces with divides for the maps Φ_w and Ψ_z. 33

1.10 The bypass relation. 33

2.1 A flow graph $\Gamma : V_0 \to V_1$, embedded in a 3–manifold Y^3. The dashed ovals indicate the vertex sets V_0 and V_1. 40

2.2 Decomposing the cobordism $X_{\alpha\beta\gamma\delta}$ into two triangular cobordisms. 58

2.3 Move (4) in Lemma 2.4.7 72

2.4 Multiplicities around the vertex in Move(4). 72

2.5 The standard diagram for (S^3, w, w'), with multiplicities labeled. The point p' is chosen in the same domain as w'. 79

2.6 The Heegaard triple $\overline{\mathcal{T}}$. 96

2.7 The Heegaard triple $(S^2, \alpha_0, \beta_0, \gamma_0, w, w')$ with $x^{\pm}, y^{\pm}, z^{\pm}$ and multiplicities m_1 and m_2 labelled. 97

2.8 Examples of the three types of elementary flow graphs. We view these graphs as flowing from bottom to top. We divide type (3) into two subtypes: type $(3i)$ and $(3ii)$. 109

2.9 Six moves which can be used to move between any two Cerf decompositions of a strong ribbon flow graph. 114

2.10 Vertices and edges in Lemma 2.6.9. 116

2.11 Vertices and edges used in the proof of invariance under move (2). 117

2.12 Vertices and edges used in the proof of invariance under move (3). Here $e_i = e * e' * e''$. 118

2.13 Edges and vertices in the proof of Lemma 2.6.11. The dashed lines indicate the incoming and outgoing ends of the flow graph. 120

2.14 The Heegaard Triple $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ used to perform $(1,2)$–stabilization of triangle maps. 140

2.15 The Heegaard triple $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$. The circles c_1, c_2 and c_3 are circles on which we need to stretch to compute the 1–handle map or to apply the stabilization lemma for triangle maps. 153

2.16 The various topological constructions used in the construction of the graph cobordism map. 160

2.17 A computation of a portion of the maps in $F^{\circ}_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}$ in Lemma 2.10.7. We see explicitly that the composition is $S^+_v \circ F^{\circ}_{\mathfrak{H},\mathfrak{s}|_{Y_1}}$. 167

2.18 The cobordisms (W_0, Γ_0) and (W_0', Γ_0') in Lemma 2.12.1 182

2.19 A sequence of diagrams which can be used to demonstrate that $F^{\circ}_{0,\mathfrak{s}} \circ F_{W_0',\boldsymbol{\gamma},\mathfrak{s}|_{W_0'}} = S^-_v$. 183

2.20 The cobordisms (W_i, Γ_i) and $(\widetilde{W}_{i0}, \Gamma_{i0})$, as cobordisms from bottom to top. . 187

2.21 We compute the above transition map, for appropriate choices of almost complex structures, in Theorem 2.13.1. 189

2.22 The intermediate diagram $\mathcal{H}_{1.5}$. 190

2.23 The almost complex structures J_α and J_β, shown in $\mathcal{H}_{1.5}$. 191

2.24 On the left are multiplicities on the diagram $(S^2, \alpha_1, \beta_2, w, w')$ used to analyze disks in $\mathcal{H}_{1.5}$, as well as the circles c and c_α. On the right, we see how the almost complex structure degenerates as we stretch J_α along c_α. 192

2.25 Multiplicities on $\mathcal{H}_1$. 195

2.26 The triple $\overline{\mathcal{T}}$ in Proposition 2.13.8, a stabilization of $\mathcal{T}$. We have included some multiplicities. The shaded regions are each examples of small triangles, which might be counted. This is used to compute the transition map for moving β_1 to β_2. 198

2.27 The Heegaard triple used to go from $\mathcal{H}_{1.5}$ to $\mathcal{H}_2$ in Proposition 2.13.9. The shaded regions are each examples of small triangles which might be counted. 202

2.28 The path λ in the diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$. 205

2.29 Breaking a path γ into a sequence of arcs $\lambda_n * \cdots * \lambda_1$, each of which crosses at most one $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curve. We will compute each move individually, by degenerating the almost complex structure carefully. 214

2.30 The almost complex structures J_α and J_β on a diagram freely-stabilized at w. On the bottom we've written in multiplicities of a homology disk. 216

2.31 A Heegaard triple which has been free-stabilized with the above triple. We stretch the almost complex structure along c and $c_{\beta\beta'}$. The shaded regions are each examples of small triangles which might be counted. 222

2.32 The map induced by $(\phi_i)_*$ along the top row can be computed by performing a handleslide of a $\boldsymbol{\beta}$ curve over the disk B, as and then pushing forward under Γ_i. We also need a change of almost complex structure map at the end, which is not shown. 226

3.1 The path λ and the surfaces Σ and $\overline{\Sigma}_p$. 252

3.2 Different choices of $\overline{\alpha}_s$ curve on D_1 interpolating $\alpha_s \setminus D_2$. 252

3.3 The diagram $\mathcal{H}_0$ used for quasi-stabilization 253

3.4 An example of the possible strip breaking which can occur for the homology class A. 262

3.5 A version of quasi-stabilization. 264

3.6 The diagram $\mathcal{H}_0 = (S^2, \alpha_0, \beta_0, w, w')$ in Lemma 3.6.3. 266

3.7 Two diagrams illustrating the strong positivity condition 269

3.8 Quasi-stabilizing a Heegaard triple 270

3.9 Multiplicities for a triangle on the diagram $(S^2, \alpha_0, \beta_0, \gamma_0, w, z)$. 271

3.10 Multiplicities on a doubly quasi-stabilized Heegaard diagram 283

3.11 An example of a Maslov index zero disk which might be counted by the map $\Phi_{\mathcal{J}(T_1,T_1')\to\mathcal{J}(T_2,T_2')}$. 285

3.12 A doubly quasi-stabilized Heegaard diagram 292

3.13 The two domains contributing to $\Psi_{z'}$ in Lemma 3.10.1 for the almost complex structure J_s stretched first on c', and then stretched on c. 292

3.14 A configuration of many basepoints and some arcs on a link. 297

3.15 Another configuration of basepoints and arcs on a link 299

4.1 An example of a colored link cobordism (W, F). The black dots are $\mathbf{w}$–basepoints and the white dots are the $\mathbf{z}$–basepoints. The shaded region is of type–$\mathbf{w}$, and the unshaded region is of type–$\mathbf{z}$. 305

4.2 Several simple dividing sets on cylinders and the maps they correspond to. These are surfaces of the form $\Sigma = L \times [0,1]$ inside of cobordisms $W = Y \times [0,1]$. 307

4.3 Building a diagram for an unlink where one component has four basepoints . 325

4.4 A quasi-stabilized diagram. 327

4.5 An example of quasi-stabilization maps not commuting 336

4.6 The diagram we consider in Lemma 4.3.23. 341

4.7 The diagram we consider in Lemma 4.3.25. 345

4.8 The diagram we insert into $\overline{D} \cap \Sigma$ for the maps $S^{\pm}_{\mathbb{U},D,\overline{D}}$. 350

4.9 Stretching along three necks 355

4.10 A version of quasi-stabilization 360

4.11 A double quasi-stabilization of a Heegaard triple 361

4.12 A free-stabilization inside a bigon. 364

4.13 Homology classes appearing which can make the a quasi-stabilized triangle map non-diagonal 367

4.14 Three configurations of quasi-stabilizations 376

4.15 A subcase in showing when quasi-stabilization and band maps commute. . . 378

4.16 The second case in showing when quasi-stabilization and band maps commute 379

4.17 A stabilized Heegaard triple used to show that disk stabilization followed by a band map is a quasi-stabilization 382

4.18 A modified free-stabilization to show that a band map, followed by disk destabilization is a quasi-destabilization 384

4.19 Two bands B^{α} and B^{β} related by an isotopy across $\mathbf{z}$–basepoints 387

4.20 Splitting a band into two bands and a disk 388

4.21 A configuration of the ends of the bands where a $\mathbf{w}$–band map commutes with a $\mathbf{z}$–band map 393

4.22 A diagram for an unlink in S^3 396

4.23 The remaining model computations needed to prove Proposition 4.7.12 without using any basepoint tricks. 397

4.24 A picture of why band maps don't commute in general 398

4.25 A picture of quasi-stabilization maps not commuting with band maps 402

4.26 Pushing the band out of a compound 1–handle/band 405

4.27 A local extrema of the dividing set crossing a saddle 444

4.28 The four configurations of local extrema of dividing arcs 447
4.29 The arcs C_- and C_+ on $\Sigma(\mathbb{S})$ 452
4.30 A picture of Move 10 460
4.31 Moving a saddle across a divide 462
4.32 Puncturing the link cobordism by removing a 4–ball along the dividing set . 465
4.33 Showing invariance under puncturing 466
4.34 Three examples of different choices of ribbon graph cores for ribbon surfaces. 474
4.35 Examples of elementary ribbon flow graphs. 479
4.36 Replacing a high valence vertex with a lower valence vertex and a trivalent vertex. 483
4.37 The H shaped graph. 484
4.38 **The curve λ is the concatenation of λ_1 and λ_2.** 485
4.39 Manipulating and decomposing a link cobordism and ribbon graph core which satisfies (4) into a sequence of link cobordisms with ribbon graph cores which satisfy (5). 489

5.1 An example of the dividing sets 501
5.2 The surface $\Sigma_{\alpha\beta\gamma}$ inside of $X_{\alpha\beta\gamma}$. 511
5.3 The 2–plane fields $\xi_{\mathbf{w}}(\psi)$ and $\xi_{\mathbf{z}}(\psi)$. 515
5.4 The move $\mathcal{O}_3$, which takes place in a solid torus in M. 521
5.5 The move $\mathcal{O}_3'$ between two parametrized Kirby diagrams 522
5.6 The diffeomorphisms Φ_H and $\Phi_{K',\mathcal{O}_3} \circ \Phi_{K,\mathcal{O}_3}^{-1}$ 524
5.7 Some Kirby calculus manipulations 525
5.8 Two diagrams for $S^1 \times S^2$, with a knot or link. 530
5.9 Blowing up a surgery triple away from the link. 547

5.10 Blowing up a surgery triple along the link I. 548

5.11 A movie presentation of the sphere E inside of $W(S^3, \mathbb{S}'_1)$. 549

5.12 A movie for $\Sigma'_j \subseteq W(S^3, \mathbb{S}'_1)$ 550

5.13 Blowing up a surgery triple II. 552

5.14 The two diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ for $(S^1 \times S^2, U, w, z)$ in a non-torsion grading, 557

5.15 The decorated link cobordisms associated to the quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$. 562

5.16 The triples used to compute the band maps $F^{\mathbf{o}}_{B_0}$ 567

5.17 The two knot cobordisms (W_i, Σ_i) between (S^3, K_-) and (S^3, K_+). 584

5.18 A decomposed link cobordism 589

5.19 Some homology classes, after cutting the link cobordism into pieces 589

6.1 Decorated link cobordisms for Φ and Ψ. 602

6.2 The bypass relation 605

6.3 Decorated link cobordisms for the quasi-stabilization maps. 636

6.4 The effect of conjugation on a quasi-stabilized diagram. 637

6.5 Computing the link cobordism map for Φ 641

6.6 A bypass triple 642

6.7 Obtaining the formula $\rho_* \simeq \mathrm{id} + \Psi \circ \Phi$ from the bypass relation. 643

6.8 The link cobordisms used to define the maps E_1, G_1, E_2 and G_2. 646

6.9 The link cobordisms $(W_1 \sqcup W_2, F_1 \sqcup F_2)$ and $(W_1 \# W_2, F_1 \# F_2)$ 649

6.10 The decorated surfaces M_1 and M_2. 651

6.11 A genus one diagram for $(S^1 \times S^2, \mathbb{L}_0)$. 653

6.12 The link cobordisms inducing the maps f, A and B, and a bypass relation between all three. 654

6.13 ink cobordisms representing compositions of maps 656

6.14 The decoration on the surface used in the cobordism defining the map G_1. . 659

6.15 The decoration on a link cobordism whose cobordism map is $\tau_{K_1\#K_2} \circ F_{W,F_1}$. 659

6.16 A bypass relation. 660

6.17 Isotoping the dividing sets in the bypass relation, to derive the connected sum formula. 660

6.18 The complexes for T_r, and the maps Φ and Ψ. 664

6.19 The complex $\mathcal{C}FL^\infty(T_r\#T_r)$, as well the conjugation map $\iota_K \approx \iota_1|\iota_2+\iota_1\Phi_1|\iota_2\Psi_2$.666

6.20 The complex $A_0^-(T_r\#T_r)$ on the left and the involution ι_0, on the right. . . . 666

6.21 Some knots and their involutive invariants of large surgeries. 668

6.22 Model complexes for $T_{3,4}, T_{4,5}, T_{5,6}$ and $T_{6,7}$. 669

6.23 A decomposition of $G_1 \circ E_1$ and a decomposition of the identity link cobordism.671

6.24 The connected sum region of the diagrams $\mathcal{H}_{K_1\#K_2}$ and $\mathcal{H}_{K_1\sqcup K_2}$ for $(Y_1\#Y_2, K_1\#K_2)$ and $(Y_1\#Y_2, K_1 \sqcup K_2)$. 672

6.25 The connected sum region of two triples which can compute the band maps $F^{\mathbf{z}}_B$ and $F^{\mathbf{z}}_{B'}$. 675

6.26 A decomposition of the link cobordism $(S^1 \times D^3, M_1)$. 677

6.27 Computing some terms of the maps $F^{\mathbf{z}}_B$ and Ψ_{z_1} and Φ_{w_1}. 678

6.28 Computing a 2–handle map. 680

6.29 The classes counted by $\Phi_{w_1}(\theta_1^-\theta_2^+)$ ad $\Phi_{w_1}(\theta_1^+\theta_2^-)$. 681

6.30 Showing $C_3 = C_3''$ by counting the ends of two 1–dimensional moduli spaces. 683

Acknowledgments

I would like to extend my gratitude to many people. First and foremost, I'm indebted to my advisor, Ciprian Manolescu, who patiently taught me about Heegaard Floer homology and gauge theory, and gave me many suggestions for problems to explore. I also am indebted to many members of the mathematical community for helpful conversations, mathematical suggestions, collaborative endeavors, and career advice. These people include Akram Alishahi, Maciej Borodzik, Skip Garibaldi, Kristen Hendricks, Jen Hom, Ko Honda, András Juhász, David Krcatovich, Tye Lidman, Jianfeng Lin, Robert Lipshitz, Yajing Liu, Andy Manion, Marco Marengon, Gordana Matić, Mike Miller, Yi Ni, Peter Ozsváth, Sucharit Sarkar, Matt Stoffregen, Dylan Thurston, and Faramarz Vafaee.

Finally I'd like to give thanks to my family, Anna, Cooper, Katy and Steve, for their support while I was in graduate school.

Vita

2012 B.S. (Mathematics), University of Washington

2012–2017 Research and teaching assistant, Mathematics Department, UCLA

Publications

Connected sums and involutive knot Floer homology (2017) (arXiv:1705.01117 [math.GT])

Link cobordisms and absolute gradings on link Floer homology (2017) (arXiv:1701.03454 [math.GT])

Link cobordisms and functoriality in link Floer homology (2016) (arXiv:1610.05207 [math.GT])

A connected sum formula for involutive Heegaard Floer homology (2016) (with Kristen Hendricks and Ciprian Manolescu). Selecta Mathematica, *to appear*.

Naturality and mapping class groups in Heegaard Floer homology (2016) (András Juhász and Dylan Thurston. My contribution: Section 9.3 on simple handleswap invariance, only) (arXiv: 1210.4996 [math.GT])

Quasi-stabilization and basepoint moving maps in Heegaard Floer homology. Algebraic and Geometric Topology, *to appear*.

Graph cobordisms and Heegaard Floer homology (2015) (arXiv:1512.01184 [math.GT])

A graph TQFT for hat Heegaard Floer homology (2015) (arXiv:1503.05846 [math.GT])

CHAPTER 1

Introduction

In the early 2000s, Ozsváth and Szabó introduced a collection of invariants for 3–manifolds and 4–manifolds called Heegaard Floer homology. To a based 3–manifold (Y, w) with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$ there are $\mathbb{F}_2[U]$–modules[1]

$$HF^-(Y, w, \mathfrak{s}), \qquad HF^\infty(Y, w, \mathfrak{s}), \qquad HF^+(Y, w, \mathfrak{s}), \qquad \widehat{HF}(Y, w, \mathfrak{s}).$$

The action of U is trivial on $\widehat{HF}(Y, w, \mathfrak{s})$. To a 4–manifold W with boundary $-Y_1 \sqcup Y_2$, they define a map

$$F^\circ_{W,\mathfrak{s}} : HF^\circ(Y_1, w_1, \mathfrak{s}_1) \to HF^\circ(Y_2, w_2, \mathfrak{s}_2).$$

The maps from Ozsváth and Szabó fit into a slightly restricted TQFT framework. However there are several natural properties of a TQFT which their maps cannot hope to satisfy, due to some limitations in the construction. Because of this, the phrase **secondary TQFT** has sometimes been used to describe Heegaard Floer homology. The purpose of this paper is to describe a construction which turns Heegaard Floer homology into a more robust TQFT package.

In this introductory chapter, we describe some formal properties of the TQFT, and highlight some of the main results. In the subsequent chapters, we will describe the constructions in more detail and prove the stated results.

[1] Ozsváth and Szabó described $\mathbb{Z}[U]$–modules, however more work would need to be done to upgrade the results of this thesis to $\mathbb{Z}$ coefficients

1.1 What is a TQFT?

A TQFT is approximately a functor from a cobordism category to an algebraic category which satisfies some natural axioms. Suppose that $\mathcal{C}$ is a category. So as not be excessively vague, we are interested in the case that $\mathcal{C}$ is a category of free chain complexes over a ring $\mathcal{R}$, perhaps with some extra structure, such as a filtration, or grading. A common assumption is that $\mathcal{C}$ be a **symmetric monoidal category**, though we will not really have use for the terminology. We also tacitly assume in this introduction that the objects of $\mathcal{C}$ are appropriately nice, e.g. free and finitely generated over $\mathcal{R}$. There are two operations which we need on $\mathcal{C}$, namely a notion of tensor product $\otimes$ and a notion of duality for which we will write $\vee$. Further, there must be some object $\mathcal{R} \in \mathcal{C}$ which satisfies $\mathcal{R} \otimes C \cong C$ for all $C \in \mathcal{C}$.

The historical origin of TQFTs go back to the late 80s and early 90s, and several highly influential papers due to Witten [Wit88]. Axiomatized versions of conformal field theories and topological quantum field theories were proposed by Segal [Seg88] and Atiyah [Ati88]. For our purposes, a convenient formulation of the axioms for a TQFT can be found in [BT06].

An $(n+1)$-dimensional TQFT $\mathcal{F}$ over the field k assigns a vector space $\mathcal{F}(Y)$ to each closed, oriented n manifold Y. To an $(n+1)$–manifold W with boundary $\partial W = -Y_1 \sqcup Y_2$ (the decomposition being part of the input data), it assigns a map

$$\mathcal{F}_W : \mathcal{F}(Y_1) \to \mathcal{F}(Y_2),$$

such that the following axioms are satisfied:

(TQFT.1) (Naturality) An orientation preserving diffeomorphism $f : Y \to Y'$ of an n–manifold induces a map $f_* : \mathcal{F}(Y) \to \mathcal{F}(Y')$, which is functorial under composition. An orientation preserving diffeomorphism between two cobordisms induces a commutative diagram, satisfying the obvious commutative diagram.

(TQFT.2) (Functoriality) If $W = W_2 \circ W_1$ is a composition of two cobordisms, then $\mathcal{F}_W = \mathcal{F}_{W_2} \circ \mathcal{F}_{W_1}$.

(TQFT.3) (Normalization) If Y is any n–manifold, then $\mathcal{F}_{Y\times[0,1]} : \mathcal{F}(Y) \to \mathcal{F}(Y)$ is the identity.

(TQFT.4) (Multiplicativity) There are functorial isomorphisms

$$\mathcal{F}(Y_1 \sqcup Y_2) \cong \mathcal{F}(Y_1) \otimes_{\mathcal{R}} \mathcal{F}(Y_2), \qquad \text{and} \qquad \mathcal{F}(\varnothing) \cong k.$$

The diffeomorphism $(Y_1 \sqcup Y_2) \sqcup Y_3 \cong Y_1 \sqcup (Y_2 \sqcup Y_3)$ induces the map corresponding to associativity on the level of tensor products. The diffeomorphism $Y \sqcup \varnothing \cong Y$ induces the natural map $\mathcal{F}(Y) \otimes \mathcal{R} \cong \mathcal{F}(Y)$. Similarly a disjoint union of cobordisms induces the tensor product of the two cobordism maps.

(TQFT.5) (Symmetry) The diffeomorphism $Y_1 \sqcup Y_2 \cong Y_2 \sqcup Y_1$ induces the map corresponding to switching the factors of a tensor product.

Note that the cobordism $Y \times [0,1]$ can be viewed as a cobordism from $Y \sqcup -Y$ to $\varnothing$ and hence induces a map $\mathcal{F}(Y) \otimes_{\mathcal{R}} \mathcal{F}(-Y) \to \mathcal{R}$. It gives a pairing, which can be seen to be nondegenerate by using Axiom (TQFT.3) and looking at an S-shaped cobordism which is diffeomorphic to $Y \times [0,1]$.

We should think of the above axioms as a guide for what a TQFT should look like, instead of an absolute definition of a TQFT. The TQFTs we will describe will require extra data, such as a collection of basepoints or links in the 3–manifolds, and a graph or surface inside of the cobordisms.

1.1.1 $(1+1)$–dimensional TQFTs, Frobenius algebras, closed strings and open strings

We now briefly discuss $(1+1)$–dimensional TQFTs. We do so for two reasons. Firstly, they are the simplest non-trivial example of TQFTs. Secondly, and more importantly, the ribbon graph cobordism category and the decorated link cobordism category can both be described in terms of 4–dimensional cobordisms which contain 2–dimensional cobordisms.

When thinking about $(1+1)$–dimensional TQFTs in this paper, it will sometimes be convenient to call a disjoint union of circles a **closed string** and a disjoint union of closed intervals an **open string**. I am not sure if this terminology corresponds to precisely to notions from physics in any reasonable way, however I have found it in mathematical literature. The definition of TQFT from the previous section corresponds to a closed string TQFT. Trying to not be overly-precise, morphisms in a $(1+1)$–dimensional open string cobordism category consist of 2–dimensional manifolds with boundary (and corners) such that the boundary is written as the union of a vertical boundary subset and a horizontal boundary subset, which meet along the corners of the surface. See Figure 1.1. One can also allow a TQFT with both closed strings and open strings [LP08].

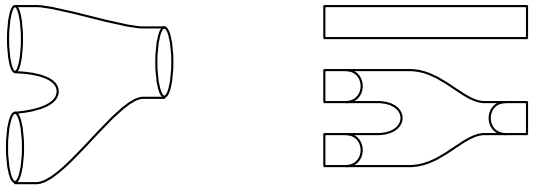

Figure 1.1: **A closed string cobordism (left) and two open string cobordisms (right).**

A classical result in the study of TQFTs is a classification of $(1+1)$–dimensional closed string TQFTs in terms of Frobenius algebras [Abr96] [Saw95].

For a $(1+1)$–dimensinal TQFT, the vector space associated to an arbitrary disjoint union of circles is determined by the vector space V assigned to a single circle. To the disjoint union of n circles, by the multiplicativity axiom, one assigns the vector space $\bigotimes^n V$.

A **Frobenius algebra** over k is a vector space V over k together with maps

$$\mu : V \otimes_k V \to V$$

$$\Delta : V \to V \otimes_k V$$

$$\eta : k \to V$$

$$\epsilon : V \to k,$$

called multiplication, comultiplication, unit and counit, respectively, such that certain compatibility relations are satisfied. We think of multiplication and comultiplication as being the maps associated to 2–dimensional pairs of pants, and the unit and counit as being the maps associated to caps and cups. The compatibility relations (which we have not listed) can be interpreted in terms of simple topological manipulations of surfaces.

1.2 A ribbon graph TQFT for multi-based manifolds

We now note that the axioms (TQFT.1)–(TQFT.5) do not quite hold for the cobordism maps described by Ozsváth and Szabó in [OS06]. There are several issues:

- Heegaard Floer homology is a functor of based manifolds. If $\phi : (Y, w) \to (Y, w)$ is a diffeomorphism preserving w, then there is an induced map

$$\phi_* : CF^\circ(Y, w, \mathfrak{s}) \to CF^\circ(Y, w, \mathfrak{s}).$$

  If γ is a loop in $\pi_1(Y, w)$, there is an induced diffeomorphism γ_* associated to isotoping the basepoint w along the loop γ. *A-priori* this map can be nontrivial. We will see that it is in fact trivial on homology for HF^-, HF^∞ and HF^+ (in the presence of a single basepoint), though it can be nontrivial on the chain level for CF^-, and on the level of homology for $\widehat{HF}$.

- The cobordisms considered by Ozsváth and Szabó must be connected, and have connected ends. One reason is that the 1–handle and 3–handle maps they define cannot connect or disconnect different components. Another reason is that they only consider cobordisms between manifolds with a single basepoint.

- With the maps they defined, there is no way to add, remove, or combine basepoints. Thus the duality $\widehat{CF}(-Y) \cong CF(Y)^\vee$ cannot be stated on the level of cobordisms, as their construction yields no map for the $Y \times [0, 1]$ viewed as a cobordism from $Y \sqcup -Y$ to $\varnothing$, since the incoming end of the cobordism has two basepoints, while the outgoing

end has none. Ozsváth and Szabó do state a version of duality by picking a basis of both objects which was induced by the Heegaard diagram.

In this section, we will describe the graph TQFT, which deals with the above issues.

1.2.1 Axioms for a ribbon graph TQFT

In this section, we state the analogs of Atiyah's axioms for the graph TQFT.

First, we state an important definition:

Definition 1.2.1. A **multi-based 3–manifold** $(Y, \mathbf{w})$ is a 3–manifold with a finite collection of basepoints $\mathbf{w} \subseteq Y$ such that each component of Y has at least one basepoint. A **graph cobordism** $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a 4–manifold manifold with $\partial W = -Y_1 \sqcup Y_2$ and Γ is an embedded graph such that $\Gamma \cap W$ consists of valence one vertices of Γ and $\Gamma \cap \partial W = \mathbf{w}_1 \cup \mathbf{w}_2$). A **ribbon graph** is a graph such that the edges adjacent to each vertex are assigned a cyclic ordering.

The following are the axioms for the ribbon graph TQFT. We state them for $\widehat{CF}$. The axioms also hold for CF^-, though one extra piece of data is needed for CF^- (a coloring of the components of the graph), which we will discuss after we state the axioms, and will suppress from the statements of the axioms.

(Γ-TQFT.1) (Naturality) If $\phi : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is an orientation preserving diffeomorphism of multi-based manifolds such that $\phi_*(\mathfrak{s}_1) = \mathfrak{s}_2$, then there is an induced map $\phi_* : \widehat{CF}(Y_1, \mathbf{w}_1, \mathfrak{s}_1) \to \widehat{CF}(Y_2, \mathbf{w}_2, \mathfrak{s}_2)$ which is well defined up to chain homotopy, and functorial under composition.

(Γ-TQFT.2) (Functoriality) If (W, Γ) and (W', Γ') are two composable ribbon graph cobordisms, then

$$F_{W',\Gamma',\mathfrak{s}'} \circ F_{W,\Gamma,\mathfrak{s}} \simeq \sum_{\substack{\mathfrak{t} \in \mathrm{Spin}^c(W \cup W') \\ \mathfrak{t}|_W = \mathfrak{s}, \mathfrak{t}|_{W'} = \mathfrak{s}'}} F_{W' \cup W, \Gamma' \cup \Gamma, \mathfrak{t}}.$$

(Γ-TQFT.3) (Normalization) If $(Y, \mathbf{w})$ is a multi-based 3–manifold, then

$$F_{Y\times[0,1],\mathbf{w}\times[0,1],\mathfrak{s}} \simeq \mathrm{id}\,|_{\widehat{CF}(Y,\mathbf{w},\mathfrak{s})}.$$

(Γ-TQFT.4) (Multiplicativity) There are functorial isomorphisms

$$\widehat{CF}(Y_1 \sqcup Y_2, \mathbf{w}_1 \sqcup \mathbf{w}_2, \mathfrak{s}_1 \sqcup \mathfrak{s}_2) \cong \widehat{CF}(Y_1, \mathbf{w}_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2} \widehat{CF}(Y_2, \mathbf{w}_2, \mathfrak{s}_2),$$

and

$$\widehat{CF}(\varnothing) \cong \mathbb{F}_2.$$

The diffeomorphism $(Y_1 \sqcup Y_2) \sqcup Y_3 \cong Y_1 \sqcup (Y_2 \sqcup Y_3)$ induces the map corresponding to associativity on the level of tensor products. Similarly the diffeomorphism $Y \sqcup \varnothing \cong Y$ induces the natural map $\widehat{CF}(Y, \mathbf{w}, \mathfrak{s}) \otimes \mathbb{F}_2 \cong \widehat{CF}(Y, \mathbf{w}, \mathfrak{s})$.

(Γ-TQFT.5) (Symmetry) The diffeomorphism $(Y_1, \mathbf{w}_1) \sqcup (Y_2, \mathbf{w}_2) \cong (Y_2, \mathbf{w}_2) \sqcup (Y_1, \mathbf{w}_1)$ induces the map corresponding to switching the factors of a tensor product.

(Γ-TQFT.6) (Weak duality) There are chain homotopy equivalences

$$\widehat{CF}(-Y, \mathbf{w}, \bar{\mathfrak{s}}) \simeq (\widehat{CF}(Y, \mathbf{w}, \mathfrak{s}))^\vee := \mathrm{Hom}_{\mathbb{F}_2}(\widehat{CF}(Y, \mathbf{w}, \mathfrak{s}), \mathbb{F}_2)$$

which are natural under change of diagrams maps. Furthermore if $W : Y_1 \to Y_2$ is a cobordism and $\mathfrak{s} \in \mathrm{Spin}^c(W)$, then

$$F_{W^\vee,\Gamma,\mathfrak{s}} = F^\vee_{W,\Gamma,\mathfrak{s}},$$

under the above pairing between $\widehat{CF}(Y, \mathbf{w}, \mathfrak{s})$ and $\widehat{CF}(-Y, \mathbf{w}, \bar{\mathfrak{s}})$.

(Γ-TQFT.7) (Strong duality) The cobordism maps for $(Y \times [0,1], \mathbf{w} \times [0,1])$, viewed as a cobordism from $(Y, \mathbf{w}) \sqcup (-Y, \mathbf{w})$ to $\varnothing$, or vice-versa, give a pairing between $\widehat{CF}(Y, \mathbf{w}, \mathfrak{s})$ and $\widehat{CF}(-Y, \mathbf{w}, \bar{\mathfrak{s}})$ which is chain homotopic to the pairing from the weak duality axiom.

Note that an easy algebraic computation shows that the axioms (Γ-TQFT.1)–(Γ-TQFT.5) and (Γ-TQFT.7) imply that a turned around cobordism induces the dual map, with respect to the pairing induced by the trace and cotrace cobordisms (one simply composes with the trace and cotrace cobordisms to turn around a cobordism).

In [Zem15] the author shows that that the above axioms (Γ-TQFT.1)–(Γ-TQFT.6) hold for the graph TQFT on $\widehat{HF}$. We will additional provide a construction in Chapter 2, which works for $\widehat{HF}$ as well as the other flavors.

Strong duality, axiom (Γ-TQFT.7), does in fact also hold, but will be proven in a later paper. Weak duality, axiom (Γ-TQFT.6), was proven by Ozsváth and Szabó in [OS06]. A cleaner way to axiomatize the TQFT would be to combine these two axioms, however at the time of submitting this thesis, although having proven the stronger duality result, I have not finished writing the paper [Zem].

A natural question is what can be said for CF^-, CF^∞ and CF^+, the more topologically sensitive flavors. To work more generally over the other flavors, we need the following definition:

Definition 1.2.2. A **coloring** of a collection of basepoints $\mathbf{w}$ is a pair $(\sigma, \mathfrak{P})$ where $\mathfrak{P}$ is a finite set $\sigma : \mathbf{w} \to \mathfrak{P}$ is a map. If Γ is a graph, a **coloring** of Γ is a pair $(\sigma, \mathfrak{P})$, where $\mathfrak{P}$ is a finite set and $\sigma : C(\Gamma) \to \mathfrak{P}$ is a map from the set, $C(\Gamma)$, of connected components of Γ to the set $\mathfrak{P}$.

If $\mathfrak{P} = \{p_1, \dots, p_n\}$ we write

$$\mathbb{F}_2[U_{\mathfrak{P}}] := \mathbb{F}_2[U_{p_1}, \dots, U_{p_n}].$$

The character $\mathfrak{P}$ is a Fraktur P, and stands for "palatte" (i.e. a collection of colors).

If $(Y, \mathbf{w})$ is a multi-pointed 3–manifold, the complex $CF^-(Y, \mathbf{w}, \mathfrak{s})$ is a module over $\mathbb{F}_2[U_{\mathbf{w}}] := \mathbb{F}_2[U_{w_1}, \dots, U_{w_n}]$. Given a coloring $(\sigma, \mathfrak{P})$, we can color the complex by taking the tensor product

$$CF^-(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}) := CF^-(Y, \mathbf{w}, \mathfrak{s}) \otimes_{\mathbb{F}_2[U_{\mathbf{w}}]} (\mathbb{F}_2[U_{\mathbf{w}}, U_{\mathfrak{P}}]/I_{\sigma,\mathfrak{P}}),$$

where $I_{\sigma,\mathfrak{p}}$ is the ideal generated by monomials of the form $w - \sigma(w)$, ranging over $w \in \mathbf{w}$. We will usually drop the coloring from the notation, and just write $CF^-(Y, \mathbf{w}, \mathfrak{s})$ even if we are implicitly using a coloring.

The axioms (Γ-TQFT.1)–(Γ-TQFT.7) be easily restated for colored ribbon graph cobordisms, using the ring $\mathbb{F}_2[U_{\mathfrak{p}}]$ instead of $\mathbb{F}_2$. We will prove in Chapter 2 that axioms (Γ-TQFT.1)–(Γ-TQFT.6) hold for CF^-. Again, axiom (Γ-TQFT.7), the strong duality axiom, does in fact hold, but will be proven in a later paper [Zem].

Note that given a ribbon graph Γ inside of W, there is a surface with boundary (and corners along ∂W) $R(\Gamma)$ inside of W formed by taking disks at the vertices of Γ, and gluing strips according to the cyclic ordering, in such a way that yields an oriented surface. It is easy to see that the isotopy class of the surface $R(\Gamma)$, relative $\partial W = -Y_1 \sqcup Y_2$ is well defined.

Definition 1.2.3. We say that two ribbon graphs Γ and Γ' in W are **ribbon equivalent** if $R(\Gamma)$ and $R(\Gamma')$ are isotopic in W.

(Γ-TQFT.8) (Ribbon graph equivalence) Two ribbon graph cobordisms which have ribbon equivalent graphs induce the same map.

Note that the fact that two ribbon equivalent graph cobordisms induce the same graph is not immediately obvious from the construction of the graph cobordism maps, since we associate a map to each edge and each vertex of the graph. One approach would be to compile a finite set of moves connecting any two ribbon graphs which are ribbon equivalent. We will not do this. Instead, the above axiom follows from Theorem C in Chapter 4, relating the graph cobordism maps to the link Floer cobordism maps.

Finally, we note that the maps $\widehat{F}_{W,\Gamma}$ on $\widehat{CF}$ depend only on a regular neighborhood of the graph, and do not depend on the choice of ribbon structure of the graph, by the work in [Zem15]. We state this as an additional axiom:

(Γ-TQFT.9) (Reduction to $\widehat{CF}$) The maps $\widehat{F}_{W,\Gamma}$, when restricted to $\widehat{CF}$ depend only on a regular neighborhood of Γ, and in particular are independent of the ribbon

structure on Γ.

1.2.2 Some properties and applications of the graph TQFT and its construction

The dependence on the basepoint in the original construction of cobordism maps from [OS06] was not made entirely explicit. The construction of Ozsváth and Szabó implicitly involved picking a path γ inside of W from w_1 to w_2, so we really should write $F^{\circ}_{W,\gamma,\mathfrak{s}}$ for their map. It was not clear from their construction how the maps depended on the choice of path. Nonetheless, for most applications, it's not important to know that the maps are independent of the choice of path.

The dependence on the basepoint was made precise in [Juh16], in the sense that the cobordism maps on $\widehat{HF}$ were shown to be invariants of the pair (W, γ). More generally, using the work of Juhász, one could also consider cobordisms where there are many basepoints on each end (the same number on each end), and a path from each incoming basepoint to an outgoing basepoint. It was not clear from the construction in [Juh16] whether the maps actually depend on the choice of path γ. More generally, Juhász did not provide a way of functorially adding and removing basepoints.

The graph TQFT that we construct in this paper resolves these issues. We note that there is a special construction that one can use for $\widehat{HF}$ ([Zem15]) using Juhász's sutured TQFT:

Theorem 1.2.4 ([Zem15])**.** *To a graph cobordism* $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ *there are cobordism maps* $F_{W,\Gamma} : \widehat{HF}(Y_1, \mathbf{w}_1) \to \widehat{HF}(Y_2, \mathbf{w}_2)$*, which are invariants of the pair* (W, Γ) *and are functorial under composition. There are refinements over* Spin^c *structures which satisfy (*Γ*-TQFT.1)–(*Γ*-TQFT.6).*

The above theorem is proven in [Zem15], which is not included in this thesis, since the results are also obtained in Chapter 2, using a more direct construction.

Juhasz's sutured TQFT uses the contact gluing map from [HKM08] to define maps. For a graph cobordism where the graph is a collection of paths from the incoming boundary to the outgoing boundary, the contact gluing map is not used in a nontrivial way. For a more

complicated graph, the construction uses the contact gluing map in a non-trivial way.

Also, by construction, the maps for $\widehat{HF}$ depend only on a regular neighborhood of the graph Γ, and do not require a ribbon structure to define.

By analyzing the graph TQFT for $\widehat{HF}$ defined using sutured Floer homology, we can compute the $\pi_1(Y, w)$ action on $\widehat{HF}$:

Theorem 1.2.5 ([Zem15]). *If $\gamma \in \pi_1(Y, w)$, the diffeomorphism map γ_* on $\widehat{HF}(Y, \mathbf{w})$ is equal to*

$$\gamma_* = \mathrm{id} + [\gamma] \circ (\partial_1)_*,$$

where $(\partial_1)_*$ is the map induced by counting holomorphic disks going over w exactly one, and $[\gamma]$ denotes the $H_1(Y; \mathbb{Z})/\operatorname{Tors}$ action described by Ozsváth and Szabó.

It is not hard to compute examples of 3–manifolds where this map does not vanish, and hence in general we see that the maps of Ozsváth and Szabó depend on a choice of path, for the hat flavor.

Juhász's construction does not extend beyond the hat flavor of Heegaard Floer homology, since it uses the Honda-Kazez-Matić contact gluing map from [HKM08]. We will be able to construct a graph TQFT for the other flavors directly though, by defining some new maps. However the structure of the TQFT is more subtle than the one we had for the hat flavor, since it requires an extra decoration on the graphs: a ribbon structure. Since the graph cobordism maps for $\widehat{HF}$ depend only on a regular neighborhood of the graph, and not the actual graph itself, this suggests that the contact gluing map probably does not admit a straightforward extension to HF^-.

We prove the following in Chapter 2:

Theorem 1.2.6. *Suppose that $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a colored, ribbon graph cobordism. If $\mathfrak{s} \in \mathrm{Spin}^c(W)$, there is a map*

$$F_{W,\Gamma,\mathfrak{s}} : CF^-(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to CF^-(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2})$$

which is an invariant, up $\mathbb{F}_2[U_{\mathfrak{p}}]$*–equivariant chain homotopies. The maps satisfy (*Γ*-TQFT.1)–(*Γ*-TQFT.6).*

The maps appearing in the previous theorem are between the colored modules, and are homomorphisms of $\mathbb{F}_2[U_{\mathfrak{p}}]$–modules. We are suppressing the coloring from the notation.

Analogously to the analysis we did for $\widehat{HF}$, we can compute the $\pi_1(Y,w)$ action on $CF^-(Y,\mathbf{w},\mathfrak{s})$ whenever $w \in \mathbf{w}$:

Theorem 1.2.7. *If* $\gamma \in \pi_1(Y,w)$ *then the induced map* γ_* *has* U_w*–equivariant chain homotopy type*

$$\gamma_* \simeq \mathrm{id} + \Phi_w \circ [\gamma],$$

for some map Φ_w*.*

The map Φ_w can be thought of as the formal derivative of the differential ∂ with respect to U_w. Another interesting property of the map Φ_w is its identification as the graph cobordism map of a peculiar graph cobordism: the "broken path" graph cobordism, shown in Figure 1.2.

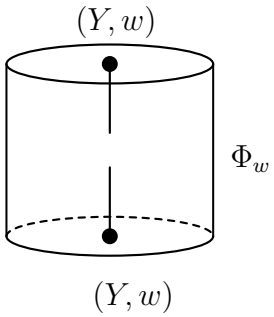

Figure 1.2: **The "broken path" graph cobordism induces the map** Φ_w**.**

An interesting property of the map Φ_w is that there is a chain homotopy

$$\Phi_w = H \circ \partial + \partial \circ H,$$

where H is the non $\mathbb{F}_2[U_{\mathbf{w}}]$–equivariant map

$$H = \frac{d}{dU_w}.$$

Hence the map Φ_w will vanish $HF^-(Y,w,\mathfrak{s})$, but not in general on $\widehat{HF}(Y,w,\mathfrak{s})$ since $\widehat{CF}(Y,w,\mathfrak{s})$ is the tensor product $CF^-(Y,w,\mathfrak{s}) \otimes_{\mathbb{F}_2[U_w]} (\mathbb{F}_2[U_w]/(U_w))$. The map H, being non-equivariant, doesn't necessarily induce a map on the tensor product.

The map Φ_w can also be described in terms of a path cobordism map, as follows. Since the graph cobordisms are invariant under puncturing and adding trivial strands, we can puncture the broken path cobordism twice, adding a new edge from each new boundary copy of S^3 to one of the two components of the broken path. However, the composition of the 4–handle and 0–handle maps is the identity map, so we can just connect up the punctures without changing the map. Reordering the handles involved, we see that the map Φ_w is also equal to the path cobordism map obtained by attaching a 3–handle along a 2–sphere in Y containing bounding a ball containing w, then adding a dual 1–handle, connecting the new copy of S^3, which contains w, back to Y.

1.2.3 Connected sum cobordisms and involutive Heegaard Floer homology

Heegaard Floer homology satisfies a well known connected sum formula:

$$CF^-(Y_1\#Y_2, w, \mathfrak{s}_1\#\mathfrak{s}_2) \cong CF^-(Y_1, w_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2[U]} CF^-(Y_2, w_2, \mathfrak{s}_2).$$

This was proven by Ozsváth and Szabó [OS04d] by constructing a map from the tensor product to the complex on the left, by producing a quasi-isomorphism as a composition of 1–handle maps and a triangle map. A 4–dimensional interpretation of their map is somewhat mysterious (at least to me).

There is a natural graph cobordism between $(Y_1\#Y_2, w)$ and $(Y_1 \sqcup Y_2, w_1 \sqcup w_2)$. The 4–manifold is obtained by attaching a 1–handle (to join Y_1 and Y_2) or a 3–handle (to separate $Y_1\#Y_2$), depending on which direction we want the cobordism to go. One puts a trivalent graph inside the cobordism. This is shown in Figure 1.3. Note that there are two cyclic orderings on the trivalent vertex, resulting in potentially two different maps.

A natural question is whether the maps induced by such cobordisms give chain homotopy equivalences. This was answered affirmatively in [HMZ16]:

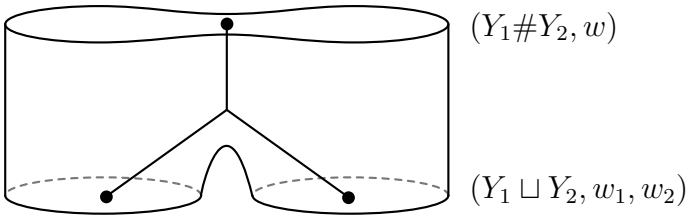

Figure 1.3: **The graph cobordism for a connected sum.** It is a 1–handle (or 3–handle) cobordism. Note that there are two cyclic orderings we can put on.

Theorem 1.2.8. *The graph cobordism maps for the connected sum graph cobordisms induce chain homotopy equivalences. Indeed, for appropriate choices of cyclic orderings, they yield two pairs of maps which are chain homotopy inverses of each other.*

The connected sum maps form [OS04d] have a similar ambiguity, since the role of Y_1 and Y_2 in their maps is not symmetric. Hence one can think of Ozsváth and Szabó's construction of maps associated to connected sums as defining two maps, with the same domain and target. Indeed in [Zem], we will identify each the maps of Ozsváth and Szabó with one of the two graph cobordism maps.

The interaction of ι, the conjugation action, with the Ozsváth-Szabó connected sum map is hard to decipher. The interaction with ι with the graph cobordism maps is more tractible. We have the following:

Theorem 1.2.9 ([HMZ16])**.** *Under the chain homotopy equivalence*

$$CF^-(Y_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2[U]} CF^-(Y_2, \mathfrak{s}_2) \cong CF^-(Y_1 \# Y_2, \mathfrak{s}_1 \# \mathfrak{s}_2)$$

given by the graph cobordism maps, the conjugation involution ι on the connected sum is chain homotopy equivalent to $\iota_1 \otimes \iota_2 + U(\Phi_1 \otimes \Phi_2)(\iota_1 \otimes \iota_2)$.

One should compare this formula to the one we will later prove for involutive knot Floer homology in Chapter 6. In [HMZ16] we also showed that, for algebraic reasons, the term $U(\Phi_1 \otimes \Phi_2)(\iota_1 \otimes \iota_2)$ vanishes on the chain level when $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are self conjugate.

1.2.4 Overview of the construction of the graph TQFT

The graph TQFT is constructed in a similar manner as the cobordism maps from [OS06]. One breaks the cobordism into pieces, and describes a map for each piece. The 4–dimensional handles induce maps which are roughly the same as the maps Ozsváth and Szabó described (though slightly more general, since now we allow things like a 1–handle which connects two different components). However there are new maps, called the *free-stabilization* and *relative homology maps*.

There are two free-stabilization maps, one for adding a basepoint (denoted S^+_w) and one for removing a basepoint (denotes S^-_w). The map S^+_w is a map

$$S^+_w : CF^-(Y, \mathbf{w}, \mathfrak{s}) \to CF^-(Y, \mathbf{w} \cup \{w\}, \mathfrak{s}),$$

and the map S^-_w is defined in the reverse direction.

Given a path λ from w_1 to w_2, both basepoints in $\mathbf{w}$, we can describe an endomorphism of $CF^-(Y, \mathbf{w}, \mathfrak{s})$, denoted A_λ. It is -1 graded and satisfies

$$A_\lambda \circ \partial + \partial \circ A_\lambda = U_{w_1} + U_{w_2},$$

i.e. it's a chain homotopy between the actions of U_{w_1} and U_{w_2}. Note that if we color so that $U_{w_1} = U_{w_2}$, then A_λ is a chain map.

If λ is a path from w_1 to w_2, the diffeomorphism ϕ_λ, corresponding to moving w_1 to w_2, induces a map

$$(\phi_\lambda)_* : CF^-(Y, \mathfrak{s}, w_1) \to CF^-(Y, \mathfrak{s}, w_2).$$

An essential observation of the graph TQFT is that the chain homotopy type of $(\phi_\lambda)_*$ is given by

$$(\phi_\lambda)_* \simeq S^-_{w_1} A_\lambda S^+_{w_2},$$

(assuming we color so that $U_{w_1} = U_{w_2}$). This relation is the key to establishing that the graph cobordism maps generalize the path cobordism maps. The above relation can be thought of as the normalization axiom (Γ-TQFT.3) in the definition we presented earlier.

In fact, it's not necessary to know that this axiom holds to construct the graph cobordism maps, prove invariance, and prove the composition law. However to see that the identity graph cobordism map induces the identity map, one needs the above relation.

Another observation is that if $w \notin \mathbf{w}$, and λ is any path from w to a point in $\mathbf{w}$ then

$$S^-_w A_\lambda S^+_w \simeq \mathrm{id},$$

as endomorphisms of $CF^-(Y, \mathbf{w}, \mathfrak{s})$. This should be thought of as invariance under adding an additional trivial strand to the graph.

As an illustration, we will describe the graph cobordism maps for graph cobordisms (W, Γ) which have underlying 4–manifold $W = Y \times [0, 1]$. For such graph cobordisms, the graph cobordism map is relatively simple, and is determined entirely by a composition of the maps $S^\pm_w$ and A_λ. We call this the **graph action map**. The nomenclature comes about because we think of it as encoding the $H_1(Y; \mathbb{Z})/\operatorname{Tors}$ action, among other things. Let us first consider the case, that Γ, when projected into Y, consists of just a single path λ, from w_1 to w_2. In this case, the graph action map is defined to be just

$$S^-_{w_1} A_\lambda S^+_{w_2}.$$

As we remarked, this is chain homotopic to the diffeomorphism map associated to moving w_1 to w_2 along the path λ.

The next interesting piece of the story is a trivalent graph. Suppose that Γ has three edges, λ_1, λ_2 and λ_3, one incoming vertex v^-, one trivalent vertex v_0 and two outgoing vertices v^+_1, v^+_2. Write V_{in} for $\{v^-\}$ and V_{out} for $\{v^+_1, v^+_2\}$. This configuration is shown in Figure 1.4.

The map

$$A_\Gamma : CF^-(Y, V_{\mathrm{in}}, \mathfrak{s}) \to CF^-(Y, V_{\mathrm{out}}, \mathfrak{s})$$

is defined by the formula

$$S^-_{v^-} S^-_{v_0} A_{\lambda_3} A_{\lambda_2} A_{\lambda_1} S^+_{v_0} S^+_{v^+_1} S^+_{v^+_2}.$$

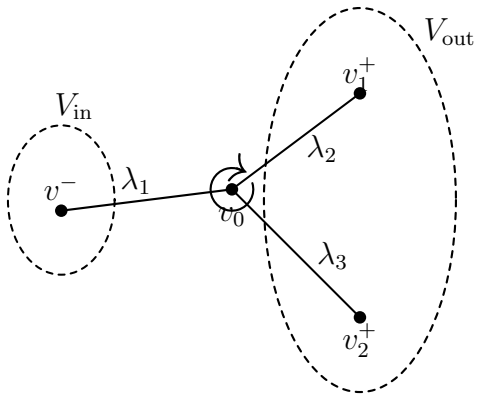

Figure 1.4: **A trivalent graph with cyclic ordering.**

We now show that this is invariant under cyclic permutation of the λ_i, but not under arbitrary permutation. Transposing λ_1 and λ_2 has the following effect:

$$\begin{aligned} S^-_{v^-,v_0} A_{\lambda_3} A_{\lambda_2} A_{\lambda_1} S^+_{v_0,v_1^+,v_2^+} &\simeq S^-_{v^-,v_0} A_{\lambda_3} A_{\lambda_1} A_{\lambda_2} S^+_{v_0,v_1^+,v_2^+} + U S^-_{v^-,v_0} A_{\lambda_3} S^+_{v_0,v_1^+,v_2^+} \\ &\simeq S^-_{v^-,v_0} A_{\lambda_3} A_{\lambda_1} A_{\lambda_2} S^+_{v_0,v_1^+,v_2^+} + U S^-_{v^-} S^+_{v_1^+,v_2^+}. \end{aligned}$$

The term $US^-_{v^-} S^+_{v_1,v_2}$ appears when we perform a transposition. However if we perform two transpositions (i.e. a cyclic permutation), we get an additional summand of $US^-_{v^-} S^+_{v_1,v_2}$, which cancels the one from the first transposition. Hence the map for a trivalent graph depend on the ordering of λ_1, λ_2 and λ_3, but only up to cyclic permutation.

Finally, to prove the formula for the π_1–action, the key is to write a curve γ as a concatenation of two arcs, λ_1 and λ_2, viewed as paths from w to a new basepoint w', and vice-versa. We can then write

$$\gamma_* \simeq (S^-_{w'} A_{\lambda_2} S^+_w)(S^-_w A_{\lambda_1} S^+_{w'}).$$

One notes that $S^+_w S^-_w \simeq \Phi_w$. From here, it is easy to use the algebra of the various maps involved to show that $\gamma_* \simeq \mathrm{id} + \Phi_w A_\lambda$ (see Chapter 2 for more details).

1.3 A TQFT for link Floer homology

In this section, we describe a construction of cobordism maps for (all flavors of) link Floer homology. Conjecturally the maps generalize Juhász's construction for $\widehat{HFL}$ [Juh16].

We first describe an axiomatic description of the link Floer TQFT. Then we provide a more detailed introduction to link Floer homology and describe the link Floer TQFT and some of its properties. We also explain the connection with the ribbon graph TQFT described in the previous section.

1.3.1 Axioms for a decorated link TQFT

Link Floer homology assigns a group to a link, decorated as follows:

Definition 1.3.1. A **oriented, multi-based link** $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ in Y is an embedded link with basepoints $\mathbf{w} \cup \mathbf{z}$ such that each component of L has at least two basepoints, and as one traverses L, the basepoints alternate between $\mathbf{w}$ and $\mathbf{z}$.

Naturally, to define cobordism maps, we must consider link cobordisms with some sort of decoration which respects the basepoints. The following is the definition we use in Chapter 4, and is essentially equivalent to the definition from [Juh16, Definition 4.5]:

Definition 1.3.2. We say a pair (W, F) is a **decorated link cobordism** between two 3–manifolds with multi-based links, and write $(W, F) : (Y_1, L_1, \mathbf{w}_1, \mathbf{z}_1) \to (Y_2, L_2, \mathbf{w}_2, \mathbf{z}_2)$, if

1. W is a 4–dimensional cobordism from Y_1 to Y_2;
2. $F = (\Sigma, \mathcal{A})$, where Σ is an oriented surface, properly embedded in W, and $\mathcal{A} \subseteq \Sigma$ is a properly embedded 1–manifold dividing Σ into two disjoint subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, which meet along $\mathcal{A}$;
3. $\partial\Sigma = -L_1 \sqcup L_2$
4. Each component of $L_i \setminus \mathcal{A}$ contains exactly one basepoint;

5. $\mathbf{w}_1 \cup \mathbf{w}_2 \subseteq \Sigma_{\mathbf{w}}$ and $\mathbf{z}_1 \cup \mathbf{z}_2 \subseteq \Sigma_{\mathbf{z}}$.

Juhász describes a set of cobordism maps on the hat flavor of link Floer homology [Juh16]. His construction uses the contact gluing map from [HKM08], and as such, his construction doesn't obviously extend to the other flavors of link Floer homology.

We now describe a axioms for a decorated link TQFT. We phrase them for the hat flavor, for simplicity, though they apply equally well to the $-$ and ∞ flavors (as described in the next section), upon changing $\mathbb{F}_2$ to the appropriate ring.

(Link-TQFT.1) (Naturality) If $\phi : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is an orientation preserving diffeomorphism of multi-based manifolds such that $\phi_*(\mathfrak{s}_1) = \mathfrak{s}_2$, then there is an induced map $\phi_* : \widehat{CFL}(Y_1, \mathbb{L}_1, \mathfrak{s}_1) \to \widehat{CFL}(Y_2, \mathbb{L}_2, \mathfrak{s}_2)$ which is well defined up to chain homotopy, and functorial under composition.

(Link-TQFT.2) (Functoriality) If (W, F) and (W', F') are two composable, decorated link cobordisms, then

$$F_{W',F',\mathfrak{s}'} \circ F_{W,F,\mathfrak{s}} \simeq \sum_{\substack{\mathfrak{t} \in \mathrm{Spin}^c(W \cup W') \\ \mathfrak{t}|_W = \mathfrak{s}, \mathfrak{t}|_{W'} = \mathfrak{s}'}} F_{W' \cup W, F' \cup F, \mathfrak{t}}.$$

(Link-TQFT.3) (Normalization) If $(Y, \mathbb{L})$ is a 3–manifold with multi-based link, then

$$F_{Y \times [0,1], F_{\mathrm{id}}, \mathfrak{s}} \simeq \mathrm{id}_{\widehat{CFL}(Y, \mathbb{L}, \mathfrak{s})},$$

where $F_{\mathrm{id}} = (L \times [0,1], \mathcal{A})$, and $\mathcal{A}$ are divides of the form $\{d\} \times [0,1]$ for a collection of points d between each basepoint of $\mathbb{L}$.

(Link-TQFT.4) (Multiplicativity) There are functorial isomorphisms

$$\widehat{CFL}(Y_1 \sqcup Y_2, \mathbb{L}_1 \sqcup \mathbb{L}_2, \mathfrak{s}_1 \sqcup \mathfrak{s}_2) \cong \widehat{CFL}(Y_1, \mathbb{L}_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2} \widehat{CFL}(Y_2, \mathbb{L}_2, \mathfrak{s}_2),$$

and

$$\widehat{CFL}(\varnothing) \cong \mathbb{F}_2.$$

The diffeomorphism $(Y_1 \sqcup Y_2) \sqcup Y_3 \cong Y_1 \sqcup (Y_2 \sqcup Y_3)$ induces the map corresponding to associativity on the level of tensor products. Similarly $Y \sqcup \varnothing \cong Y$ induces the natural map $\widehat{CFL}(Y, \mathbb{L}, \mathfrak{s}) \otimes \mathbb{F}_2 \cong \widehat{CFL}(Y, \mathbb{L}, \mathfrak{s})$.

(Link-TQFT.5) (Symmetry) The diffeomorphism $(Y_1, \mathbb{L}_1) \sqcup (Y_2, \mathbb{L}_2) \cong (Y_2, \mathbb{L}_2) \sqcup (Y_1, \mathbb{L}_1)$ induces the map corresponding to switching the factors of a tensor product.

(Link-TQFT.6) (Weak duality) There are equivalences

$$\widehat{CFL}(-Y, -\mathbb{L}, \bar{\mathfrak{s}}) \simeq (\widehat{CFL}(Y, \mathbb{L}, \mathfrak{s}))^\vee = \mathrm{Hom}_{\mathbb{F}_2}(\widehat{CFL}(Y, \mathbb{L}, \mathfrak{s}), \mathbb{F}_2)$$

which are natural under change of diagrams maps. Furthermore If $(W, F) : Y_1 \to Y_2$ is a cobordism and $\mathfrak{s} \in \mathrm{Spin}^c(W)$, then

$$F_{W^\vee, F^\vee, \mathfrak{s}} = F^\vee_{W, F, \mathfrak{s}},$$

under the above identification of dual spaces.

(Link-TQFT.7) (Strong duality) The cobordism maps for $(Y \times [0,1], L \times [0,1])$, viewed as a cobordism from $(Y, \mathbb{L}) \sqcup (-Y, -\mathbb{L})$ to $\varnothing$, or vice-versa, are equivalent to the algebraic trace and cotrace maps induced by the duality identifications from the weak duality axiom.

Additionally, we will consider more general link Floer modules over the ring $\mathbb{F}_2[U, V]$. We will denote these by $\mathcal{C}FL^-(Y, \mathbb{L}, \mathfrak{s})$. One can consider an additional axiom for this TQFT:

(Link-TQFT.8) (Reduction to a graph TQFT) The $\mathbb{F}_2[U]$–module

$$\mathcal{C}FL^-(Y, \mathbb{L}, \mathfrak{s})/(V = 1)$$

depends only on $(Y, \mathbf{w})$, $\mathfrak{s}$ and the homology class $[L] \in H_1(Y; \mathbb{Z})$, up to chain homotopy equivalence. Similarly the $\mathbb{F}_2[V]$–module

$$\mathcal{C}FL^-(Y, \mathbb{L}, \mathfrak{s}) \otimes_{\mathbb{F}_2[U,V]} /(U = 1)$$

depends only on $(Y, \mathbf{z})$, $\mathfrak{s}$ and the homology class $[L] \in H_1(Y;\mathbb{Z})$. After setting $U = 1$, the map $F_{W,F,\mathfrak{s}}$ depends only on the subsurface $\Sigma_{\mathbf{w}}$, $\mathfrak{s}$, and the homology class $[\Sigma] \in H_2(W, \partial W;\mathbb{Z})$. Similarly after setting $V = 1$, the map $F_{W,F,\mathfrak{s}}$ depends only on $\Sigma_{\mathbf{z}}$, $\mathfrak{s}$, and the homology class $[\Sigma] \in H_2(W, \partial W;\mathbb{Z})$.

The above result is Theorem C of Chapter 4.

1.3.2 The link Floer TQFT

The sutured Floer TQFT defined by Juhász provides cobordism maps for sutured Floer homology, usually denotes SFH. Sutured Floer homology generalizes $\widehat{HF}$ and $\widehat{HFL}$. By applying the sutured TQFT, I constructed the graph TQFT for $\widehat{HF}$ we described earlier. Juhász uses it to describe a TQFT for $\widehat{HFL}$. A sutured cobordism can be obtained from a graph cobordism by removing a regular neighborhood of the graph, and from a link cobordism by removing a regular neighborhood of the the surface inside the link cobordism. Doing this results in a cobordism with a "horizontal" boundary component and a "vertical" boundary component. The horizontal boundaries correspond to the ends of the original cobordism. To define the cobordism map, sutured Floer homology requires a contact structure on the vertical boundary. For the graph cobordisms, the boundary is equal to a connected sum of $S^1 \times S^2$s, with some balls removed (and standard sutures on the boundary spheres), and there is a unique tight contact structure on such a 3–manifold. For a link cobordism, the vertical boundary is an S^1 bundle M over Σ. According to [Hon00] or [Lut77], up to isotopy, there is a unique contact structure on M such that the induced dividing set induced by the contact structure is equal to our dividing set. In such a way, Juhász constructs a TQFT for $\widehat{HFL}$.

The graph TQFT for $\widehat{HF}$ does not generalize without change to the other flavors, however the link Floer TQFT essentially does. Using Juhász's link Floer TQFT as inspiration, I constructed the TQFT for the other flavors.

We consider the a version of the full link Floer complex, denoted by $\mathcal{C}FL^{\infty}(Y, \mathbb{L}, \mathfrak{s})$. This

is the module generated by monomials of the form $\mathbf{x} \cdot U_{\mathbf{w}}^I V_{\mathbf{z}}^J$ and $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$. There is an endomorphism ∂ (obtained by counting pseudo-holomorphic curves), which satisfies

$$\partial^2(\mathbf{x}) = k \cdot \mathbf{x},$$

for a constant $k \in \mathbb{F}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$. As such, the object $\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$ is naturally a **curved chain complex**. However in most cases we are interested, the curvature will be zero.

Analogously to the graph TQFT, for functorial maps, we need to work with colorings of the basepoints. Also if $F = (\Sigma, \mathcal{A})$ is a surface with divides, to define maps we need to color the components of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$.

Theorem 1.3.3. *To a decorated link cobordism $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ between two multi-based links, as well as a coloring $(\sigma, \mathfrak{P})$ of the components of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, there is a filtered, equivariant map*

$$F_{W,F,\mathfrak{s}} : \mathcal{C}FL^-(Y_1, \mathbb{L}_1, \mathfrak{s}|_{Y_1}) \to \mathcal{C}FL^-(Y_2, \mathbb{L}_2, \mathfrak{s}|_{Y_2})$$

which is invariant up to filtered, equivariant chain homotopy.

The maps above are morphisms of $\mathbb{F}_2[U_{\mathfrak{P}}]$ modules, and are filtered with respect to the filtration of $\mathcal{C}FL^-$ given by powers of the $U_{\mathfrak{P}}$–variables.

1.3.3 A first reduction to CF^-: adjunction inequalities and knot cobordisms in negative definition 4–manifolds

It's easy to see that $\mathcal{C}FL^-(Y, L, \mathbf{w}, \mathbf{z}, \mathfrak{s})$ is naturally identified with $CF^-(Y, \mathbf{w}, \mathfrak{s})$ upon setting the $V_{\mathbf{z}}$ variables to 1. Similarly $\mathcal{C}FL^-(Y, L, \mathbf{w}, \mathbf{z}, \mathfrak{s})$ is identified with $CF^-(Y, \mathbf{z}, \mathfrak{s} - PD[L])$ after setting the $U_{\mathbf{w}}$ variables equal to 1. A natural question is what the reductions of the link Floer cobordism maps to CF^- are, upon doing this. When we do this reduction, we get some expressions which are familiar from Seiberg-Witten.

For simplicity, let (W, F) be a knot cobordism from (Y_1, K_1) to (Y_2, K_2) such that $F = (\Sigma, \mathcal{A})$, $\mathcal{A}$ consists of two arcs, and $\Sigma \setminus \mathcal{A}$ consists of two connected components, $\Sigma_{\mathbf{w}}$

and $\Sigma_{\mathbf{z}}$. Let $A_1, B_1, \ldots, A_{g_1}, B_{g_1}$ be a geometric symplectic basis of $H_1(\Sigma_{\mathbf{w}}; \mathbb{Z})$, and let $A'_1, B'_1, \ldots, A'_{g_2}, B'_{g_2}$ be a geometric symplectic basis of $H_1(\Sigma_{\mathbf{z}}; \mathbb{Z})$. The terminology "geometric symplectic basis" is nonstandard; we mean a basis of $H_1(\Sigma_{\mathbf{o}}; \mathbb{Z})$ such that each A_i and B_i is represented by a simple closed curve on Σ and such that A_i intersects B_i exactly once, and there are no other intersections. We define the element $\xi(\Sigma_{\mathbf{w}}) \in \Lambda(H_1(\Sigma_{\mathbf{w}}; \mathbb{Z})) \otimes_{\mathbb{F}_2} \mathbb{F}_2[U]$ by the formula

$$\xi(\Sigma_{\mathbf{w}}) = \prod_{i=1}^{g_1} (U + [A_i] \cdot [B_i]),$$

and define an element $\xi(\Sigma_{\mathbf{z}})$ similarly.

Theorem 1.3.4. *Suppose that Σ is a surface in W (with boundary equal to possibly non-trivial knots in Y_1 and Y_2), and decorate Σ with a dividing set so that $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ are both connected, and each homeomorphic to a connected sum of D^2 with some number of 2–tori. The $V_{\mathbf{z}} = 1$ reduction of $F_{W,F,\mathfrak{s}}$ is*

$$F_{W,\mathfrak{s}}(\xi(\Sigma_{\mathbf{w}}) \otimes \cdot)$$

and the $U_{\mathbf{w}} = 1$ reduction is

$$F_{W,\mathfrak{s}-PD[\Sigma]}(\xi(\Sigma_{\mathbf{z}}) \otimes \cdot).$$

A particular case of interest is when we consider the link cobordisms induced by a closed surface Σ inside of W. We push Σ so that it intersects each of Y_1 and Y_2 in a disk, with boundary equal to two unknots, U_1 and U_2. We decorate Σ so that $\Sigma_{\mathbf{w}}$ is a genus zero strip from K_1 to K_2, and put two basepoints on each of U_1 and U_2. We can pick diagrams for (Y_1, U_1) and (Y_2, U_2) which have the $\mathbf{w}$ and $\mathbf{z}$ basepoints adjacent to each other. For such a diagram, the Alexander grading of $\mathbf{x} \cdot U^i V^j$ is just $j - i$. Hence, if we additionally assume that

$$\frac{\langle c_1(\mathfrak{s}), [\Sigma] \rangle - [\Sigma] \cdot [\Sigma]}{2} + g(\Sigma) = 0,$$

(this corresponds to the Alexander grading change of $F_{W,F,\mathfrak{s}}$ being zero) then the $U = 1$ and $V = 1$ reductions of $F_{W,F,\mathfrak{s}}$ coincide. Hence

$$F_{W,\mathfrak{s}}(\cdot) = F_{W,\mathfrak{s}-PD[\Sigma]}(\xi(\Sigma) \otimes \cdot).$$

This is normally called the adjunction relation. A marginally weaker version (where the element $\xi(\Sigma)$ is not explicitly identified) is proven in [OS04a, Theorem 3.1].

If $\mathbb{K}$ is a doubly based, null-homologous knot in Y, then

$$\mathcal{H}FL^\infty(Y,\mathbb{K},\mathfrak{s}) \cong \bigoplus_{k\in\mathbb{Z}} HF^\infty(Y,\mathfrak{s})\{k\},$$

where we think of HF^∞ as being concentrated in Alexander grading zero, and $\{k\}$ denotes a shift in the Alexander grading by k. By considering the reductions of the link cobordism maps $F_{W,F,\mathfrak{s}}$ as above, we get the following general principle, which follows from the analogous property of the maps defined by Ozsváth and Szabó:

Proposition 1.3.5. *Suppose that $(W,F) : (Y_1,\mathbb{K}_1) \to (Y_2,\mathbb{K}_2)$ is a knot cobordism, $\mathbb{K}_i$ are doubly based, null-homologous knots, $F = (\Sigma,\mathcal{A})$ is a surface with divides such that $\mathcal{A}$ consists of two arcs from $\mathbb{K}_1$ to $\mathbb{K}_2$. If $b_1(W) = b_2^+(W) = 0$, then*

$$F^\infty_{W,F,\mathfrak{s}} : \mathcal{H}FL^\infty(Y_1,\mathbb{K}_1,\mathfrak{s}_1) \to \mathcal{H}FL^\infty(Y_2,\mathbb{K}_2,\mathfrak{s}_2)$$

is an isomorphism for any $\mathfrak{s} \in \mathrm{Spin}^c(W)$.

The above result is very useful. Combined with the grading change formula for the Alexander grading, it can be used to prove our bound on $\Upsilon_K(t)$. It also lets us compute the cobordism maps for closed surfaces in S^4, and recover already known bounds on τ and V_s.

1.3.4 A second reduction to CF^-: ribbon surfaces and graph TQFTs

In this section, we consider the reduction from $\mathcal{C}FL^\infty$ to CF^∞ further, describing how the reduction of the link cobordism maps gives the graph cobordism maps we described earlier.

We can think of a multi-based knot or link as a union of two collections of open strings (i.e. closed intervals), partitioned into two collections, say W and Z, which meet along their boundaries. We require that as one traverses a link component, the strings alternate between W and Z. A string in W corresponds to a subarc of the link containing a $\mathbf{w}$ basepoint and similarly for the strings in Z. To a multi-based link, we obtained the link

Floer complexes $\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$. If we identify $U_{\mathbf{w}} = 1$ or $V_{\mathbf{z}} = 1$, (corresponding to ignoring one set of the strings), then we recover $CF^\infty(Y, \mathbf{w}, \mathfrak{s})$ or $CF^\infty(Y, \mathbf{z}, \mathfrak{s} - PD[L])$. In itself, this interpretation is not particularly helpful, however it has a natural interpretation in terms of decorated cobordisms between strings and links.

As a variation to Definition 1.2.1 of a ribbon graph, we make the following definition:

Definition 1.3.6. A **ribbon surface cobordism** between two 3–manifolds with collections of open strings $(W, R) : (Y_1, S_1) \to (Y_2, S_2)$ consists of a cobordism W from Y_1 to Y_2 with embedded, oriented surface R with corners such that

1. $R \cap Y_i = S_i$;
2. The corners of R occur exactly at ∂S_i;
3. R has no closed components.

Definition 1.3.7. If R is a ribbon surface, viewed as a ribbon graph cobordism between two collections of open strings S_1 and S_2, we say that an embedded ribbon graph $\Gamma \subseteq R$ is a **ribbon graph core** of R if the following are satisfied:

1. Γ has a single valence 1 vertex on each string in S_i;
2. At any vertex of Γ of valence 3 or more, the cyclic ordering of Γ agrees with the orientation of R;
3. The subsurface of R formed by adding disks at the vertices of Γ (half disks at the boundary vertices) and strips along the edges, is isotopic inside of R to the entire surface R.

It is easy to see that if R is a ribbon surface, then there is always a ribbon graph core for R. However given a ribbon surface, there are many non-isotopic ribbon graph cores.

If (W, F) is a decorated link cobordisms, then each of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ determines a ribbon surface cobordism. Let us write $\Gamma(\Sigma_{\mathbf{w}})$ and $\Gamma(\Sigma_{\mathbf{z}})$ for the (non-unique) ribbon graph cores

of $\Sigma_{\mathbf{w}}$. An example of a ribbon surface and a ribbon graph core is shown in Figure 1.5. An example of a decorated link cobordism, the corresponding subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, as well as ribbon graph cores of these surfaces are shown in Figure 1.6.

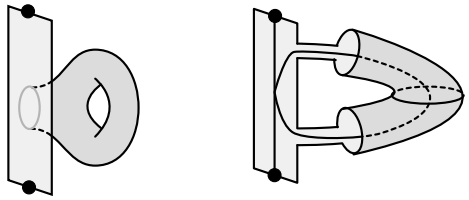

Figure 1.5: **A ribbon surface (left) and a ribbon graph core (right).** The ribbon surface on the right is redrawn slightly to make it obvious that the surface can be built by adding bands along the edges of the graph. The cyclic orders our counterclockwise with respect to the page.

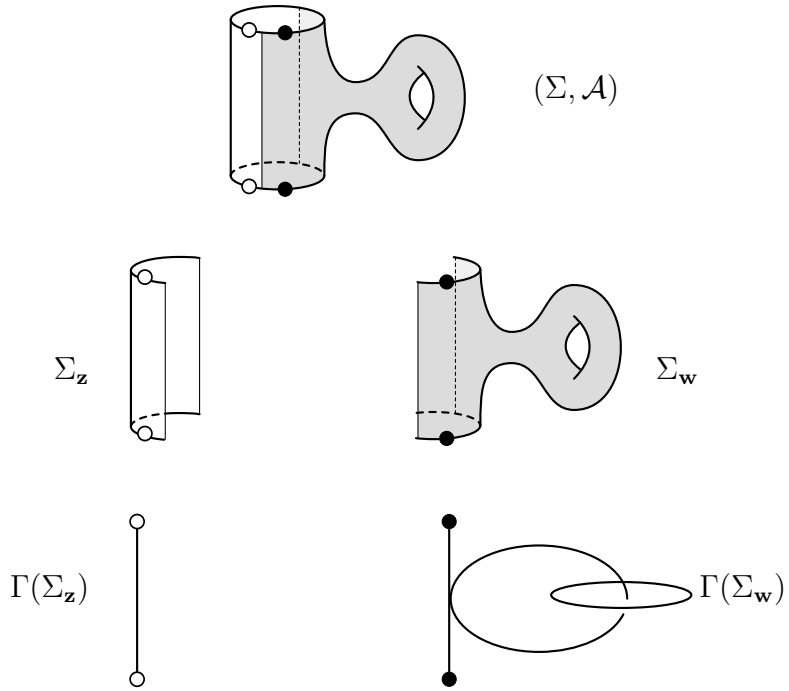

Figure 1.6: **A decorated link cobordism, the subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, as well as their ribbon graph cores. The cyclic orderings are counterclockwise with respect to the page.**

In Chapter 4, we will prove the following (Theorem C):

Theorem 1.3.8. *If (W, F) is any decorated link cobordism and $\Gamma(\Sigma_{\mathbf{w}})$ and $\Gamma(\Sigma_{\mathbf{z}})$ are ribbon graph cores for $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, respectively, then*

$$F_{W,F,\mathfrak{s}}|_{V_{\mathbf{z}}=1} = F^B_{W,\Gamma(\Sigma_{\mathbf{w}}),\mathfrak{s}}$$

and

$$F_{W,F,\mathfrak{s}}|_{U_{\mathbf{w}}=1} = F^A_{W,\Gamma(\Sigma_{\mathbf{z}}),\mathfrak{s}-PD[\Sigma]}.$$

See Chapter 4 for a description of the maps $F^A_{W,\Gamma,\mathfrak{s}}$ and $F^B_{W,\Gamma,\mathfrak{s}}$, which are minor variations of the constructions from Chapter 2.

As a corollary to Theorem, it follows that the link Floer TQFT satisfies axiom (Link-TQFT.8) and the ribbon graph TQFT satisfies Axiom (Γ-TQFT.8):

Corollary 1.3.9.

1. *Suppose that (W, F) is a decorated link cobordism. Then the $V_{\mathbf{z}} = 1$ reduction of $F_{W,F,\mathfrak{s}}$ depends only on the ribbon surface $\Sigma_{\mathbf{w}}$ and the* Spin^c *structure $\mathfrak{s}$. The $U_{\mathbf{w}} = 1$ reduction of $F_{W,F,\mathfrak{s}}$ depends only on the ribbon surface $\Sigma_{\mathbf{z}}$ and the* Spin^c *structure $\mathfrak{s} - PD[\Sigma]$.*

2. *Suppose that (W, Γ) is a ribbon graph cobordism. The graph cobordism map $F_{W,\Gamma,\mathfrak{s}}$ depends only up to isotopy on the isotopy class of the ribbon surface constructed from Γ, and not on the actual ribbon graph Γ.*

Remark 1.3.10. The dependence of the $U_{\mathbf{w}} = 1$ reduction on the homology class $[\Sigma]$ is only through the choice of Spin^c structure.

We will now show how Theorem 1.3.4 follows from Theorem 1.3.8 (in the case that $g = 1$). Let us compute the $V = 1$ reduction of the link cobordism map for a genus one surface as in Figure 1.6, formed by taking the connected sum of a closed surface with a cylinder. Let γ_1 and γ_2 be a symplectic basis (which can be used to construct a ribbon core of the surface, as in Figure 1.6). Some useful relations for the graph TQFT are shown in Figure 1.7.

Using the relations in Figure 1.7 (which we leave as an algebraic exercise for the reader to verify, using the definition of the graph action map and the relations from Chapter 4),

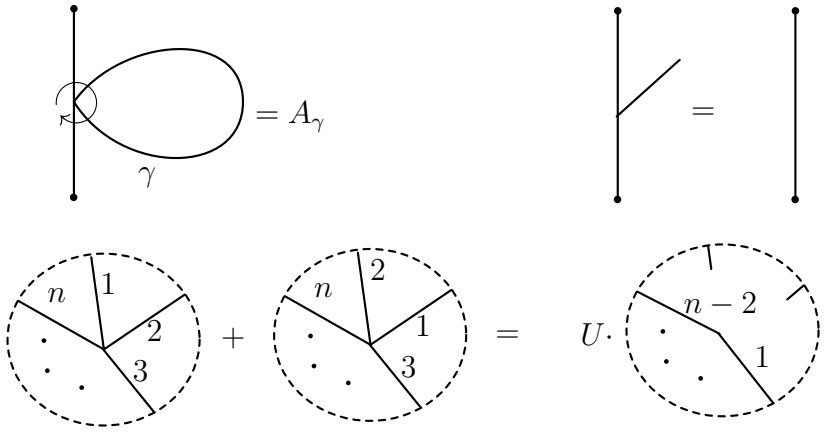

Figure 1.7: **Some relations for the graph TQFT.** (Top left) Splicing in a loop γ, using the cyclic ordering shown (so that the two strands of γ are adjacent) has the effect of the homology action A_γ. (Top right) splicing in a trivial strand has no effect on the maps. (Bottom row) Switching the relative ordering of two strands has the effect of breaking at the vertex and multiplying by U (as shown). Note that if we switch the order of two edges adjacent to a valence two vertex, the third term is interpreted to be zero.

we manipulate the ribbon graph core of the ribbon surface from Figure 1.6 to see that the induced map is $U + A_{\gamma_2} A_{\gamma_1}$.

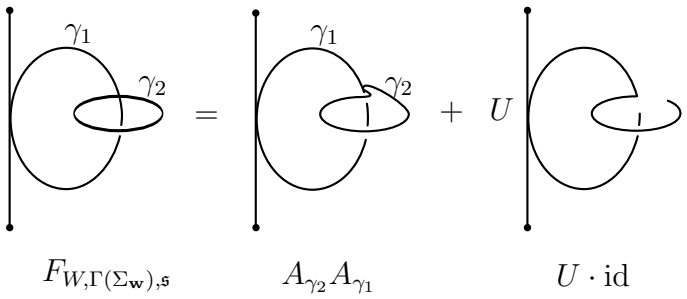

Figure 1.8: **Computing the reduction of the link Floer cobordism maps for a genus 1 surface by using relations of the graph TQFT.** We take the cyclic orderings at each vertex to be clockwise, with respect to the page they are drawn on.

To conclude this section, we note how the reduction from $\mathcal{C}FL^-$ to CF^- also yields the formula for the π_1 action on $CF^-(Y, w, \mathfrak{s})$ which we proved using the graph TQFT.

If $\gamma \in \pi_1(Y, w)$ is an embedded loop in Y, we treat γ as a knot in Y, add a z basepoint

to γ, and consider $\mathcal{C}FL^-(Y,\gamma,w,z,\mathfrak{s})$. We consider the diffeomorphism of Y corresponding to twisting in a full loop along γ. This induces a diffeomorphism ρ. Using Sarkar's formula (which we prove in Chapter 3), we have

$$\rho_* \simeq \mathrm{id} + \Phi_w \Psi_z.$$

The reduction of ρ_* to $CF^-(Y,w,\mathfrak{s})$ is the map γ_*. The map Φ_w on $\mathcal{C}FL^-$ reduces to Φ_w on CF^- once we set $V_z = 1$. However the map Ψ_z does not have as immediate a counterpart on CF^-. The reduction of the map Ψ_z counts Maslov index 1 disks with an extra factor of their multiplicity over the point z. The basepoint z is not distinguished on the multi-based 3–manifold $(Y,\mathbf{w})$. Instead, if we consider the knot shadow, it's not hard to argue that the reduction of Ψ_z is in fact

$$\Psi_z|_{V_z=1} \simeq A_\gamma + U_w \Phi_w.$$

Using the formula for ρ_*, we arrive at

$$(\gamma_*) \simeq (\mathrm{id} + \Phi_w \Psi_z)|_{V_z=1} \simeq \mathrm{id} + \Phi_w(A_\lambda + U_w \Phi_w) = \mathrm{id} + \Phi_w A_\lambda + U_w \Phi_w^2 \simeq \mathrm{id} + \Phi_w A_\lambda,$$

which is the formula that we will set in Chapter 2 by using the graph TQFT.

1.3.5 Grading changes and bounds on concordance invariants

An important property of the cobordism maps defined by Ozsváth and Szabó are the grading shifts. This is a very powerful result, which allows us to show that the d–invariant is an invariant of homology cobordism, and also derive inequalities on the d–invariant for negative definite cobordisms. Perhaps most spectacularly, Ozsváth and Szabo give a proof of Donaldson's diagonalizability theorem [Don83] in [OS03b] by considering the grading change formula for the cobordism maps and the fact that for negative definite 4–manfiolds with $b_1 = 0$, the maps $F_{W,\mathfrak{s}}$ are isomorphisms on HF^∞. Together with a theorem by Elkies on intersection forms over $\mathbb{Z}$ [Elk95], it follows that the intersection form of a negative definite smooth manifold can be diagonalized.

If $W : Y_1 \to Y_2$ is a cobordism, then the cobordism maps defined by Ozsváth and Szabó are graded, and satisfy the grading formula

$$\mathrm{gr}(F_{W,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}(\mathbf{x}) = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4},$$

for a homogeneously graded $\mathbf{x}$ (assuming also that $\mathfrak{s}$ is torsion on Y_1 and Y_2).

For the link cobordisms, there are similar grading change formulas. On link Floer homology, there are actually three gradings that are of interest. There are two homological gradings, which we denote by $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$, as well as an Alexander grading A. More generally, there is an Alexander multi-grading for links, though this can be collapsed into a single Alexander grading. In general, if $\mathfrak{s}$ is non-torsion, the Alexander grading will also require a choice of Seifert surface.

Theorem 1.3.11. *if $(W, F) : (Y_1, K_1) \to (Y_2, K_2)$ is a decorated link cobordism and $\mathbf{x}$ is a homogeneously graded element, then if S_1 and S_2 are Seifert surfaces for K_1 and K_2, then*

$$A_{S_2}(F_{W,F,\mathfrak{s}}(\mathbf{x})) - A_{S_1}(\mathbf{x}) = \frac{\langle c_1(\mathfrak{s}), [\widehat{\Sigma}]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}]}{2} + \frac{\chi(\Sigma_{\mathbf{w}}) - \chi(\Sigma_{\mathbf{z}})}{2}, \tag{1.1}$$

$$\mathrm{gr}_{\mathbf{w}}(F_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{w}}) \tag{1.2}$$

and

$$\mathrm{gr}_{\mathbf{z}}(F_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) = \frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{z}}). \tag{1.3}$$

Here

$$\widehat{\Sigma} = (-S_1) \cup \Sigma \cup S_2,$$

$\chi(\Sigma_{\mathbf{w}})$ denotes the Euler characteristic, and $\widetilde{\chi}(\Sigma_{\mathbf{w}})$ denotes the reduced Euler characteristic:

$$\widetilde{\chi}(\Sigma_{\mathbf{w}}) = \chi(\Sigma_{\mathbf{w}}) - \tfrac{1}{2}(|\mathbf{w}_1| + |\mathbf{w}_2|).$$

Using surgery theory, Ozsváth and Szabó prove in [OS03a] that if $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is a smooth, oriented knot cobordism with $b_1(W) = b_2^+(W) = 0$ then

$$\tau(K_2) - \tau(K_1) \leq -\frac{|[\Sigma]| - [\Sigma]\cdot[\Sigma]}{2} + g(\Sigma) = \min_{C \in \mathrm{Char}\, W} \frac{\langle C, [\Sigma]\rangle - [\Sigma]\cdot[\Sigma]}{2} + g(\Sigma).$$

As one would hope, the link cobordism maps can be used to give an alternate proof of this fact. By putting an appropriate choice of dividing set on Σ, and using the properties of the maps on negative definite knot cobordisms, the bound follows from a straightforward line of reasoning.

The link cobordism map argument for Ozsváth and Szabó's bound on τ also generalizes to give a bound on $\Upsilon_K(t)$. It is unclear (to me) how to do this with surgery theory.

Theorem 1.3.12. *If $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is a knot cobordism with $b_1(W) = b_2^+(W) = 0$, then*

$$\Upsilon_{K_2}(t) \geq \Upsilon_{K_1}(t) + M_t([\Sigma]) + g(\Sigma) \cdot (|t-1| - 1)$$

where

$$M_t([\Sigma]) = \max_{C \in \mathrm{Char}(W)} \frac{C^2 + b_2(W) - 2t\langle C, [\Sigma]\rangle + 2t([\Sigma] \cdot [\Sigma])}{4}.$$

For small t, one has $\Upsilon_K(t) = -\tau(K) \cdot t$. Near $t = 0$, the above bound on $\Upsilon_K(t)$ is the linearization of Ozsváth and Szabó's bound on τ.

Indeed it is also possible to prove Rasmussen's bound on the local h_i invariants by using a link cobordism argument. These are commonly denoted $V_s(K)$, in the literature. Rasmussen's proof went by way of surgery theory. The invariant $V_s(K)$ is defined by looking at the subcomplex $A_s^-(K)$ of $CFK^\infty(S^3, K)$ generated by monomials of the form $\mathbf{x} \cdot U^i V^j$ with $A(\mathbf{x}) + j - i = 0$ and $i \geq 0$ and $j \geq -s$. This has a Maslov grading from CFK^∞, and we can take the maximal grading of a non-torsion element in $H_*(A_s^-(K))$. If we multiply by $-\frac{1}{2}$, we get $V_s(K)$. Rasmussen proves [Ras04, Theorem 2.3] that if $0 \leq s \leq g_*(K)$, then

$$V_s(K) \leq \left\lceil \frac{g_*(K) - s}{2} \right\rceil.$$

For $s \geq g_*(K)$ the quantity $V_s(K)$ vanishes. By playing around with dividing sets which divide Σ into exactly two components, and using Theorems 1.3.4 and 1.3.11, one can recover an alternate proof of this bound.

1.3.6 Some interesting properties of the link TQFT

Like the broken path cobordism appearing in the graph TQFT, which induces the map Φ_w on CF^-, there are two interesting analogs on link Floer homology. These are the maps Φ_w and Ψ_z. They are defined by the formulas

$$\Phi_w(\mathbf{x}) = U_w^{-1} \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} n_w(\phi)\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)}\cdot\mathbf{y},$$

and

$$\Psi_z(\mathbf{x}) = V_z^{-1} \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} n_z(\phi)\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)}\cdot\mathbf{y}.$$

These maps were considered earlier in [Sar15] (on the associated graded complex, obtained by setting the $V_{\mathbf{z}}$ variables equal to 0) in the context of basepoint moving maps. There is a diffeomorphism ρ_K obtained by twisting a link component K of $\mathbb{L}$ in one full twist. Sarkar showed that on the associated graded complex of a link in S^3, the induced map took the form

$$(\rho_K)_* \simeq \mathrm{id} + \Phi_K\Psi_K,$$

where

$$\Phi_K = \sum_{w\in\mathbf{w}\cap K}\Phi_w, \qquad \text{and} \qquad \Psi_K = \sum_{z\in\mathbf{z}\cap K}\Psi_z.$$

The proof used grid diagrams in an essential way. In Chapter 3 we consider this question further, and show that Sarkar's formula holds for links in arbitrary 3–manifolds, and on the full link Floer complex:

Theorem 1.3.13. *The diffeomorphism map $\rho_K : \mathcal{C}FL^\infty(Y,\mathbb{L},\mathfrak{s}) \to \mathcal{C}FL^\infty(Y,\mathbb{L},\mathfrak{s})$ is filtered chain homotopic to the map*

$$(\rho_K)_* \simeq \mathrm{id} + \Phi_K\Psi_K.$$

I also gave the maps Φ_w and Ψ_z an interpretation in terms of decorated link cobordisms:

Proposition 1.3.14. *The maps Φ_w and Ψ_z are equal to the link cobordism maps for the surfaces with divides shown in Figure 1.9.*

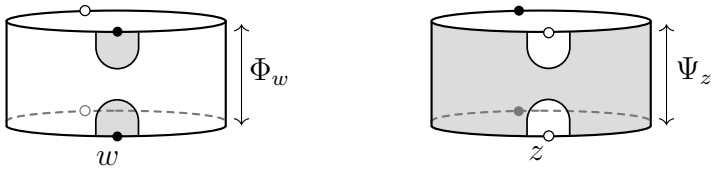

Figure 1.9: **Decorated surfaces with divides for the maps Φ_w and Ψ_z.**

Another interesting property of the link cobordism maps is the following bypass relation. It is inspired by some relations from the contact invariant in Heegaard Floer homology.

Theorem 1.3.15. *If $(W, F), (W, F')$ and (W, F'') fit into a bypass triangle, as in Figure 1.10, then*

$$F_{W,F,\mathfrak{s}} + F_{W,F',\mathfrak{s}} + F_{W,F'',\mathfrak{s}} \simeq 0.$$

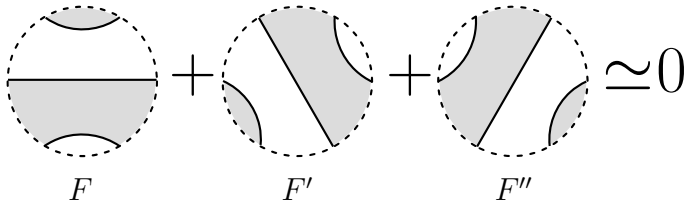

Figure 1.10: **The bypass relation.**

Using the bypass relation and the interpretation of maps from Proposition 1.3.14, we can prove Sarkar's formula for the basepoint moving map in Theorem 1.3.13. This is shown in Figure 6.7. Note that Sarkar's formula holds for links with many basepoints, though the picture proof we give only works for a link component with one basepoint. For links with many basepoints one could perform a picture argument, though it would take many bypasses.

1.3.7 Connected sums and involutive knot Floer homology

There is a natural conjugation action on knot Floer homology

$$\eta : \mathcal{C}FL^{\infty}(S^3, K, w, z) \to \mathcal{C}FL^{\infty}(S^3, K, z, w).$$

If we compose with a half twist diffeomorphism τ_K (which switches w and z), we get an endomorphism of $\mathcal{C}FL^\infty(S^3, K, w, z)$

$$\iota_K := \tau_K \circ \eta.$$

The map ι_K satisfies $A(\mathbf{x}) = -A(\iota_K(\mathbf{x}))$, so it preserves the portion in zero Alexander grading, which is naturally identified with the standard $CFK^\infty(S^3, K)$ complex. A natural question is how the map ι_K behaves with respect to the Kúnneth isomorphism on knot Floer homology, proven by Ozsváth and Szabó. I don't know of a straightforward way to analyze the interaction of ι_K with their maps[2], which like the Kúnneth isomorphisms for CF^-, involve many 1–handle maps and a triangle map.

Instead of using their maps, a natural approach is to see whether the maps from the link Floer TQFT yield chain homotopy equivalences, giving an alternate proof of the Künneth theorem for knot Floer homology. Indeed we will prove the following in Chapter 6:

Theorem 1.3.16. *There are four naturally occurring link cobordisms between* $(Y_1 \sqcup Y_2, K_1 \sqcup K_2, \{w_1, w_2\}, \{z_1, z_2\})$ *and* $(Y_1 \# Y_2, K_1 \# K_2, w, z)$ *(two in each direction) such that the induced link cobordisms are filtered, equivariant chain homotopy equivalences.*

The fact that their are two distinct link cobordism maps in each direction is an interesting feature. This is perhaps not surprising; the original maps by Ozsváth and Szabó in [OS04c] were not symmetric in the relative ordering between (Y_1, K_1) and (Y_2, K_2). Essentially their maps involved taking the connected sum at w_1 and z_2, but we could reverse the roles of Y_1 and Y_2 and take the connected sum at z_1 and w_2.

On the level of link cobordisms, the ambiguity can be seen as follows. All of the connected sum link cobordisms are obtained by attaching a 4–dimensional 1–handle and a band along the 1–handle. Inside of the surface, there are three dividing arcs: two from $K_1 \sqcup K_2$ to $K_1 \sqcup K_2$, and one from $K_1 \sqcup K_2$ to itself. Up to diffeomorphism, the surface with dividing arc

[2]In [Zem], we will identify their maps with the ones I describe below, so technically it can be done, but I don't know of a straightforward way of computing the interaction of the involution with the Ozsváth-Szabó connected sum maps.

is uniquely determined. However there is a $\mathbb{Z}/2\mathbb{Z}$ ambiguity in which region we identify with $\Sigma_{\mathbf{w}}$. The two choices are not diffeomorphic, since the ordering of Y_1, Y_2 and $Y_1 \# Y_2$ along the boundary of $\Sigma_{\mathbf{w}}$ is determined by the orientation of the surface.

By paying attention to this ambiguity, using the previous theorem, and some properties of the link cobordism maps (such as the bypass relation), we prove the following theorem in Chapter 6:

Theorem 1.3.17. *Suppose that $(Y_1, \mathbb{K}_1)$ and $(Y_2, \mathbb{K}_2)$ are two pairs of 3–manifolds with embedded, doubly based, null-homologous knots, and $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are self-conjugate* Spin^c *structures on Y_1 and Y_2. Writing $\mathbb{K}_1 \# \mathbb{K}_2$ for the connected sum, with exactly two basepoints, there are filtered, $\mathbb{F}_2[U, U^{-1}, V, V^{-1}]$–equivariant chain homotopy equivalences between*

$$\mathcal{CFL}^\infty(Y_1, \mathbb{K}_1, \mathfrak{s}_1) \otimes \mathcal{CFL}^\infty(Y_2, \mathbb{K}_2, \mathfrak{s}_2) \qquad \text{and} \qquad \mathcal{CFL}^\infty(Y_1 \# Y_2, \mathbb{K}_1 \# \mathbb{K}_2, \mathfrak{s}_1 \# \mathfrak{s}_2)$$

which intertwine $\iota_{K_1 \# K_2}$, on the latter chain complex, and

$$(\mathrm{id} \otimes \mathrm{id} + \Phi_{w_1} \otimes \Psi_{z_2})(\iota_{K_1} \otimes \iota_{K_2})$$

on the tensor product complex. There are also (different) chain homotopy equivalences of the two above complexes which intertwine $\iota_{K_1 \# K_2}$ with

$$(\mathrm{id} \otimes \mathrm{id} + \Psi_{w_1} \otimes \Phi_{z_2})(\iota_{K_1} \otimes \iota_{K_2}).$$

As usual, we can recover information about CF^- by using the analogous result on knot Floer homology, under the reduction $V_z = 1$. We can recover Theorem 1.2.9 by considering diagrams for unknots where we put a w basepoint immediately adjacent to a z basepoint.

CHAPTER 2

Graph cobordisms and Heegaard Floer homology

In this paper we construct graph cobordism maps for all flavors of Heegaard Floer homology and prove invariance of such maps. In the process we also finish the proof of invariance of the 4–manifold invariants constructed by Ozsváth and Szabó. The graph cobordism maps are defined for a very general class of cobordisms, including cobordisms which are disconnected, or have disconnected or empty ends. We also define a sort of TQFT for graphs inside a fixed 3–manifold. As an application we compute the chain homotopy type of the π_1–action on CF°, showing that it is chain homotopic to the identity map for the $+, -$ and ∞ flavors, but the chain homotopy is not U–equivariant. For the cobordisms originally considered by Ozsváth and Szabó (connected cobordisms with connected ends and a single path), this implies that, after taking homology, the cobordism maps don't depend on the path for the $+, -$ and ∞ flavors.

2.1 Introduction

Heegaard Floer homology, introduced by Ozsváth and Szabó associates functorial invariants to 3– and 4–manifolds. To an oriented, closed 3–manifold with a collection of basepoints, $(Y, \mathbf{w})$, and a Spin^c structure $\mathfrak{s}$, one associates homology groups

$$HF^\circ(Y, \mathbf{w}, \mathfrak{s})$$

for $\circ = +, -, \infty, \wedge$. The objects HF° are modules over the ring $\mathbb{Z}_2[U_{w_1}, \ldots, U_{w_n}]$ where here we let $\mathbf{w} = \{w_1, \ldots, w_n\}$. Originally defined only for connected 3–manifolds Y^3, in this paper we consider arbitrary closed and oriented Y which have at least one basepoint

in each component. Ozsváth and Szabó showed that the 3–manifold invariants fit into the framework of a (3+1)–dimensional TQFT:

Theorem 2.1.1 ([OS06])**.** *Given a connected 4–manifold W with connected ends, (Y_1, w_1) and (Y_2, w_2), and a 4–dimensional* Spin^c *structure $\mathfrak{s}$ on W, there are maps*

$$F^\circ_{W,\mathfrak{s},A} : HF^\circ(Y_1, w_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, w_2, \mathfrak{s}|_{Y_2})$$

for some choice extra data A.

The dependence on the extra data A was originally not understood, but it has been widely believed that the minimal amount of extra data needed is a path in W from w_1 to w_2. There have been many steps towards a proof of invariance for such maps, notably [OS06], [JT12], and [Juh16]. In this paper we provide the final steps towards proving invariance:

Theorem A. *Given a cobordism $(W, \boldsymbol{\gamma}) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where $\boldsymbol{\gamma}$ is a collection of paths from the basepoints in $\mathbf{w}_1$ to the basepoints in $\mathbf{w}_2$, there is a map*

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2}),$$

which is an invariant. Here W, Y_1 and Y_2 are allowed to be disconnected. The maps are $\mathbb{Z}_2[U_{\mathbf{w}}]$–module homomorphisms where we identify two variables U_{w_1} and U_{w_2} if there is a path in $\boldsymbol{\gamma}$ from $w_1 \in \mathbf{w}_1$ to $w_2 \in \mathbf{w}_2$.

In [OS06], for a closed, connected 4–manifold X with Spin^c structure $\mathfrak{s}$, Ozsváth and Szabó construct a map

$$\Phi_{X,\mathfrak{s}} : \mathbb{Z}_2[U] \otimes \Lambda^*(H_1(X)/\operatorname{Tors}) \to \mathbb{Z}_2.$$

To define such a map, one must choose two balls B_1, B_2 in X, an admissible cut $N \subseteq X$, as well as paths from a point $n \in N$ to ∂B_1 and ∂B_2. Given an admissible cut $N \subseteq X$, any two configurations of the above form are diffeomorphic, and hence Theorem A immediately shows that the untwisted version of this map, i.e. the map $\Phi_{X,\mathfrak{s}} : \mathbb{Z}_2[U] \to \mathbb{Z}_2$, depends only on N. In [OS06], it is shown to be independent of the admissible cut N.

In [Zem15] the author defines a "graph TQFT" for the hat Heegaard Floer homology using the sutured Floer homology TQFT developed by András Juhász in [Juh06] and [Juh16]. To a 4–manifold cobordism with embedded graph $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ such that $\Gamma \cap Y_i = \mathbf{w}_i$, the $\widehat{HF}$ graph TQFT associates a map

$$\widehat{F}_{W,\Gamma} : \widehat{HF}(Y_1, \mathbf{w}_1) \to \widehat{HF}(Y_2, \mathbf{w}_2),$$

which decomposes over Spin^c structures and is functorial under composition of graph cobordisms. The $\widehat{HF}$ graph TQFT defined in this manner allows for more cobordisms than in the original construction by Ozsváth and Szabó. It allows for disconnected 4–manifolds W and for disconnected or empty 3–manifolds Y_i as ends of 4–manifolds. The extra flexibility comes at the expense of keeping track of an embedded graph $\Gamma \subseteq W$. It is known that the $\pi_1(Y, p)$ action does not vanish on $\widehat{HF}(Y, p)$, and hence such a TQFT cannot be constructed without the additional data of a graph or collection of paths. There is a version of sutured Floer homology that generalizes HF^- for closed manifolds, developed by Alishahi and Eftekhary in [AE15], though we will use other techniques to construct a graph TQFT for the $+, -$ and ∞ flavors of Heegaard Floer homology.

In [Zem15], we showed that the graph cobordism maps in cylindrical cobordisms $W = Y \times I$ determine an action of $\Lambda^*(H_1(Y, \mathbf{w}; \mathbb{Z})/\operatorname{Tors})$ on $\widehat{HF}(Y, \mathbf{w})$, as well as maps corresponding to adding or removing a basepoint. Those maps, together with the path cobordism maps, uniquely determine the graph cobordism maps. In this paper we construct maps corresponding to adding or removing basepoints (which we call the free stabilization maps), and maps corresponding to a new relative homology action, and use these together with maps corresponding to handle attachment to define graph cobordism maps. To a path λ between two basepoints w_1 and w_2 which is embedded in the underlying Heegaard surface of a diagram $\mathcal{H}$ of $(Y, \mathbf{w})$, we define a map

$$A_\lambda : CF^\circ(\mathcal{H}, \mathfrak{s}) \to CF^\circ(\mathcal{H}, \mathfrak{s})$$

which satisfies

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

The situation is similar, but slightly different, for other Floer theories. For example in embedded contact homology (cf. [HT09, Sec. 2.5]), a complex $C_*(Y, \lambda, \Gamma)$ is defined independently of a basepoint and the U map is defined by first picking a basepoint w and counting certain holomorphic curves which depend on the basepoint w to get a map U_w. A path from w to w' yields a chain homotopy between the U_w and $U_{w'}$ maps on the complex $C_*(Y, \lambda, \Gamma)$. The situation is similar for the U–map in the monopole Floer homology theory developed by Kronheimer and Mrowka (cf. [KM07]).

Several novel features appear for our new graph TQFT for HF° which are not present in the hat flavor of the graph TQFT constructed in [Zem15]. Firstly we must now decorate the underlying graph in a graph cobordism with a ribbon structure, i.e. a cyclic ordering of the edges adjacent to each vertex. In the graph TQFT for $\widehat{HF}$ constructed previously, the maps depended only on the isotopy class of a regular neighborhood of the graph. In addition, Heegaard Floer homology has a module structure over a polynomial ring, which acts trivially over the hat flavor. To define functorial maps, we must work with colored sets of basepoints and colored graphs. If $\mathbf{w}$ is a set of basepoints, a coloring $(\sigma, \mathfrak{P})$ of $\mathbf{w}$ is a set $\mathfrak{P}$ indexing a collection of formal variables, as well as a map $\sigma : \mathbf{w} \to \mathfrak{P}$. If Γ is a graph, a coloring $(\sigma, \mathfrak{P})$ of Γ consists of a set $\mathfrak{P}$ and a map $\sigma : C(\Gamma) \to \mathfrak{P}$, where $C(\Gamma)$ denotes the set of connected components of Γ. We let $\mathbb{Z}_2[U_{\mathfrak{P}}]$ denote the free abelian algebra over $\mathbb{Z}_2$ and the variables U_p for $p \in \mathfrak{P}$. The colored Heegaard Floer homology groups

$$HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s})$$

are then defined by identifying the variable U_w with $U_{\sigma(w)}$ before taking homology. They are $\mathbb{Z}_2[U_{\mathfrak{P}}]$ modules. Variables corresponding to colors in $\mathfrak{P}$ which are not represented by any basepoints act essentially by extension of scalars. We will usually omit writing $\mathfrak{P}$ and write $HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{s})$ (or just $HF^\circ(Y, \mathbf{w}, \mathfrak{s})$) unless we need to emphasize the set $\mathfrak{P}$ or the coloring σ.

As a first step towards the full graph cobordism TQFT, we define a sort TQFT for graphs inside of a fixed, closed 3–manifold Y^3. We say that a graph Γ is a "flow graph" from vertices

V_0 to V_1 if each vertex in V_i has valence one in Γ and $V_0 \cap V_1 = \varnothing$. We use the functorial notation $\Gamma : V_0 \to V_1$. Notice that if $\Gamma : V_0 \to V_1$ is a flow graph, and $(\sigma, \mathfrak{P})$ is a coloring of Γ, then $(\sigma, \mathfrak{P})$ restricts to a coloring of V_0 and V_1. An example of a flow graph in a 3–manifold Y is shown in Figure 2.1. We prove the following:

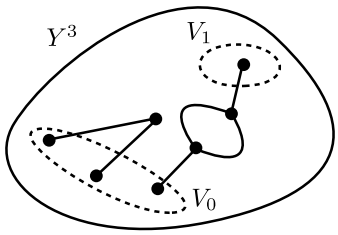

Figure 2.1: A flow graph $\Gamma : V_0 \to V_1$, embedded in a 3–manifold Y^3. The dashed ovals indicate the vertex sets V_0 and V_1.

Theorem B. *Given a ribbon flow graph $\Gamma : V_0 \to V_1$ embedded in Y^3 with coloring $(\sigma, \mathfrak{P})$, and* Spin^c *structure $\mathfrak{s}$ on Y, there is a morphism of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–modules*

$$A_{\Gamma,\mathfrak{s}} : HF^\circ(Y, V_0, \sigma|_{V_0}, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, V_1, \sigma|_{V_1}, \mathfrak{P}, \mathfrak{s})$$

which is functorial under concatenating graph components together.

We call the map $A_{\Gamma,\mathfrak{s}}$ the **graph action map**, since it encodes the $\Lambda^*(H_1(Y,\mathbb{Z})/\operatorname{Tors})$ action on HF°. In the case that $\Gamma : V_0 \to V_1$ consists of a collection of arcs from a point in V_0 to a point in V_1, the map $A_{\Gamma,\mathfrak{s}}$ is the map induced by the diffeomorphism corresponding to finger moves along the arcs in Γ.

We use the maps $A_{\Gamma,\mathfrak{s}}$ to define a full graph TQFT for cobordisms with embedded, colored ribbon graphs:

Theorem C. *Given a graph cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where Γ is a ribbon graph with coloring $(\sigma, \mathfrak{P})$, there is a homomorphism of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–modules*

$$F^\circ_{W,\Gamma,\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \sigma|_{\mathbf{w}_1}, \mathfrak{P}, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \sigma|_{\mathbf{w}_2}, \mathfrak{P}, \mathfrak{s}|_{Y_2}).$$

These are invariants of the cobordism with embedded colored ribbon graph.

In the above theorem, the cobordism W can be disconnected, and the ends Y_i can be disconnected or even empty There are also no restrictions on the embedded ribbon graph Γ other than that the vertices of the graph lying in Y_i must have valence 1. The new graph cobordism maps coincide with the path cobordism maps when the graph is a path:

Theorem D. *If $(W,\Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism and Γ is a collection of paths $\boldsymbol{\gamma}$ from the basepoints in Y_1 to the basepoints in Y_2, then*

$$F^{\circ}_{W,\Gamma,\mathfrak{s}} = F^{\circ}_{W,\boldsymbol{\gamma},\mathfrak{s}},$$

where the map on the left is the graph cobordism map, and the map on the right is the path cobordism map.

The graph cobordism maps essentially satisfy the standard composition law. The only caveat is that one must be careful about the coloring. In the following we will write $F^{\circ}_{W,\Gamma,\sigma,\mathfrak{P},\mathfrak{s}}$ to emphasize the coloring:

Theorem E. *Suppose that (W,Γ) is ribbon graph cobordism with coloring $(\sigma,\mathfrak{P})$. Suppose that (W,Γ) decomposes as $(W,\Gamma) = (W_2,\Gamma_2)\cup(W_1,\Gamma_1)$. If $\mathfrak{s}_i$ is a* Spin^c *structure on W_i, then we have the following:*

$$F^{\circ}_{W_2,\Gamma_2,\sigma|_{\Gamma_2},\mathfrak{P},\mathfrak{s}_2} \circ F^{\circ}_{W_1,\Gamma_1,\sigma|_{\Gamma_1},\mathfrak{P},\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2\\ \mathfrak{s}|_{W_1}=\mathfrak{s}_1}} F^{\circ}_{W,\Gamma,\sigma,\mathfrak{P},\mathfrak{s}}.$$

Notice that all of the maps in the above theorem are maps of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–modules.

It is interesting to compare the cobordism maps in Heegaard Floer homology to the cobordism maps in monopole Floer homology constructed in [KM07]. Due to the work of Kutluhan, Lee and Taubes in [KLT10a],[KLT10b],[KLT10c],[KLT11], and [KLT12], and due to Colin, Ghiggini and Honda in [CGH10], [CGH12a],[CGH12b] and [CGH12c], we now know that certain versions of Heegaard Floer homology, monopole Floer homology, and embedded contact homology are isomorphic. To the authors knowledge, a proof that these isomorphisms extend to the level of 4–manifold cobordism maps in Heegaard Floer homology and

monopole Floer homology has not appeared in the literature, though the two theories are strongly suspected to be equivalent. In light of this, it is natural to ask how the path cobordism maps in Heegaard Floer homology depend on the choices of paths, since one does not need to decorate cobordisms with paths in monopole Floer homology.

To this end, we first remark that we can consider a stronger invariant than the Heegaard Floer homology groups. Instead we can consider the Heegaard Floer chain homotopy types. The invariant $CF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s})$ is the set of all $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant chain complexes of the form $CF^-(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s})$, for diagrams $\mathcal{H}$ of Y, together with the change of diagrams maps $\Phi_{\mathcal{H}\to\mathcal{H}'}$ which are defined up to chain homotopy. A morphism between $CF^-(Y_1, \mathbf{w}_1, \mathfrak{s}_1)$ and $CF^-(Y_2, \mathbf{w}_2, \mathfrak{s}_2)$ is a collection of chain maps $F_{\mathcal{H}_1,\mathcal{H}_2}$, where $\mathcal{H}_i$ is a diagram for Y_i, which are defined up equivariant chain homotopy and commute with the change of diagrams maps up to equivariant chain homotopy.

Remark 2.1.2. Though we will often write things in terms of homology, all of the cobordism maps defined in this paper are defined on the stronger chain homotopy type invariant, up to equivariant chain homotopies.

In this paper, we show in this paper that the 3–dimensional graph action maps yield convenient expressions for the effect of moving basepoints along paths, and we are able to perform enough model computations to compute explicitly the $\mathbb{Z}_2[U_w]$–chain homotopy type of the $\pi_1(Y, w)$–action on all flavors of Heegaard Floer homology:

Theorem F. *If $(Y, \mathbf{w})$ is a 3–manifold and $w \in \mathbf{w}$, then the $\pi_1(Y, w)$–action has the $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariant chain homotopy type*

$$\gamma_* \simeq \mathrm{id} + \Phi_w \circ [\gamma]$$

on $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$ for $\circ \in \{+, -, \infty\}$, where Φ_w is the map

$$\Phi_w(\mathbf{x}) = U_w^{-1} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} n_w(\phi)\#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}$$

and $[\gamma]$ denotes the $H_1(Y, \mathbb{Z})/\operatorname{Tors}$ action of the curve γ. On the uncolored complexes, the map Φ_w is chain homotopic to 0, though the chain homotopy is not $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariant.

Hence γ_* is chain homotopic to id*, though not $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariantly so.*

Remark 2.1.3. In [Zem15] the author showed that the π_1–action on $\widehat{HF}$ took the form $\mathrm{id} + (\partial_1)_* \circ [\gamma]$, where ∂_1 is the map which counts disks going over the basepoint exactly once. The proof used the sutured Floer TQFT from [Juh16], but the formula coincides with the above theorem, once we set $U = 0$.

Algebraically, we can think of Φ_w as the formal derivative of the uncolored differential with respect to U_w. The non-equivariant chain homotopy from Φ_w to 0 follows from the Leibniz rule, appropriately interpreted. As a consequence, we have the following:

Theorem G. *If $(W, \boldsymbol{\gamma})$ is path cobordism where all the paths have distinct colors (for example if there is only one path), then the map*

$$F^{\circ}_{W,\boldsymbol{\gamma},\mathfrak{s}} : HF^{\circ}(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to HF^{\circ}(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2})$$

is independent of the choice of paths $\boldsymbol{\gamma}$ for $\circ = +, -, \infty$.

The stronger invariant, the map on chain complexes up to equivariant chain homotopy, does depend on the choice of paths—but after taking homology the dependence vanishes.

2.2 Organization

The organization of the paper is as follows:

- After discussing background material on Heegaard Floer homology in Section 2.3, we first define relative homology maps on the Heegaard Floer homology groups in Section 2.4. These are maps associated to closed loops or paths between two basepoints in Y, generalizing the action of $\Lambda^*(H_1(Y;\mathbb{Z})/\operatorname{Tors})$ on $HF^{\circ}(Y, w, \mathfrak{s})$ in [OS04e]. To a closed loop γ, we get a chain map A_γ, and to a path λ from w_1 to w_2, we get a chain homotopy A_λ between U_{w_1} and U_{w_2}.

- In Section 2.5 we define maps corresponding to adding or removing a basepoint, which we call free stabilization maps. We denote the free stabilization maps by S^+_w and S^-_w, depending on whether we are adding or removing the basepoint w.

- In Section 2.6 we then construct the graph action map associated to a ribbon flow graph embedded in a 3–manifold using the relative homology maps, and the free stabilization maps. In this section we prove Theorem B.

- In Sections 2.7–2.9, we define 1–,2–, and 3–handle maps, show that they are invariants, and show that they fit into the framework developed in [Juh16] and hence determine cobordism maps for cobordisms with embedded paths between the basepoints, proving Theorem A. Our contribution to this invariance result is mainly the gluing lemmas necessary to show that the 1–,2–, and 3–handle maps are invariants.

- In Section 2.10 we construct the graph cobordism maps for graph cobordisms $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where $\pi_0(Y_i) \hookrightarrow \pi_0(W)$ is a surjection for $i = 1, 2$. We show that they are invariant, proving a weak version of Theorem C. In Section 2.11 we show that the composition law holds for the cobordisms with $\pi_0(Y_i) \hookrightarrow \pi_0(W)$ a surjection, proving a weaker version of Theorem E. Finally in Section 2.12 we construct maps for cobordisms where $\pi_0(Y_i) \hookrightarrow \pi_0(W)$ is not a surjection. We then show that they are invariants, and that they satisfy the composition law, proving Theorems C and E in full generality.

- Finally in Section 2.13 we consider the effect of moving basepoints. Through a challenging model computation, we show that the effect of moving basepoints is captured by the graph action map. As a consequence, we arrive at a proof of Theorem D, showing that the path cobordism maps agree with the graph cobordism maps when the graph is a collection of paths. We also are able to explicitly compute the π_1 action on CF° and thus prove Theorem F. By exhibiting a chain homotopy which is not $U_{\mathbf{w}}$–equivariant, we arrive at Theorem G.

2.3 Preliminaries on Heegaard Floer homology

We begin with the construction and definition of the 3–manifold invariants. A multibased 3–manifold $(Y, \mathbf{w})$ is a 3–manifold Y with a finite collection of basepoints $\mathbf{w} \subseteq Y$ such that $\pi_0(\mathbf{w}) \hookrightarrow \pi_0(Y)$ is a surjection.

Definition 2.3.1. A **multipointed Heegaard diagram** $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ for $(Y, \mathbf{w})$ is a collection satisfying the following:

1. $\Sigma \subseteq Y$ is a closed, oriented surface of genus g embedded in Y such that $\mathbf{w} \subseteq \Sigma$;
2. $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_{g+\ell-1})$ is a collection of disjoint, simple closed curves on Σ bounding compressing disks in Y, where $|\mathbf{w}| = \ell$;
3. $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{g+\ell-1})$ is another collection of disjoint, simple closed curves in Σ bounding compressing disks in Y;
4. the collections $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ each span a g–dimensional sublattice of $H_1(\Sigma, \mathbb{Z})$;
5. each α_i and β_j intersect transversely.

We will also require that our diagrams satisfy admissibility requirements, which we will discuss in Section 2.3.5. Before we define the complexes precisely, we need to discuss the cylindrical formulation of Heegaard Floer homology.

2.3.1 Holomorphic disks in the cylindrical setup

In the original formulation of Heegaard Floer homology in [OS04e] and [OS06], one defines the flavors of Heegaard Floer homology as modified versions of Lagrangian intersection Floer homology of the tori

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

inside of $\mathrm{Sym}^{g+\ell-1}(\Sigma)$. An equivalent construction, the cylindrical formulation, was given by Robert Lipshitz in [Lip06]. Throughout the paper we will use this formulation, mainly because gluing lemmas are considerably easier to prove.

Following [Lip06], we define $W_{\mathbb{D}} = \Sigma \times I \times \mathbb{R}$ (where $I = [0,1]$) and define $\pi_{\mathbb{D}} : W_{\mathbb{D}} \to I \times \mathbb{R}$, $\pi_{\mathbb{R}} : W_{\mathbb{D}} \to \mathbb{R}$ and $\pi_{\Sigma} : W_{\mathbb{D}} \to \Sigma$ to be the projection maps. We define $C_\alpha = \boldsymbol{\alpha} \times \{1\} \times \mathbb{R}$ and $C_\beta = \boldsymbol{\beta} \times \{0\} \times \mathbb{R}$. We consider certain almost complex structures J which are perturbations of a split almost complex structure $\mathfrak{j}_\Sigma \times \mathfrak{j}_{\mathbb{D}}$ (see [Lip06] for the precise requirements).

We consider Riemann surfaces S with boundary and $(g + \ell - 1)$ positive punctures $\{p_1, \ldots, p_{g+\ell-1}\}$ and $(g + \ell - 1)$ negative punctures $\{q_1, \ldots, q_{g+\ell-1}\}$ on ∂S, such that S is compact away from the punctures. We will consider J–holomorphic curves $u : S \to W_{\mathbb{D}}$ such that

1. The map u is a smooth embedding;
2. $u(\partial S) \subseteq C_\alpha \cup C_\beta$;
3. u is nonconstant on each component of S;
4. u has finite energy;
5. for each i, $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ and $u^{-1}(\beta_i \times \{0\} \times \mathbb{R})$ consists of exactly one component of ∂S;
6. $\lim_{w \to p_i} \pi_{\mathbb{R}} \circ u(w) = -\infty$ and $\lim_{w \to q_i} \pi_{\mathbb{R}} \circ u(w) = +\infty$.

Let $\overline{W}_{\mathbb{D}} = \Sigma \times \mathbb{D}$ denote the compactification of $W_{\mathbb{D}}$, and let $e_\alpha = \partial\mathbb{D} \cap \{\mathrm{Re}(z) \geq 0\}$ and $e_\beta = \partial\mathbb{D} \cap \{\mathrm{Re}(z) \leq 0\}$. We let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the homology classes of maps

$$\phi : (S, \partial S) \to (W_{\mathbb{D}}, C_\alpha \cup C_\beta)$$

which converge to $\mathbf{x}$ and $\mathbf{y}$ at the punctures, where two maps are considered homologous if they define the same element in

$$H_2(\overline{W}_{\mathbb{D}}, (\boldsymbol{\alpha} \times e_\alpha) \cup (\boldsymbol{\beta} \times e_\beta); \mathbb{Z}).$$

Given a $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we now define $\mathcal{M}(\phi)$ to be the moduli space of J–holomorphic curves which represent $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. The space $\mathcal{M}(\phi)$ has a free $\mathbb{R}$–action, so we denote by $\widehat{\mathcal{M}}(\phi)$ the space $\mathcal{M}(\phi)/\mathbb{R}$. We will often have to consider the compactification of spaces $\widehat{\mathcal{M}}(\phi)$, which we denote by $\overline{\widehat{\mathcal{M}}(\phi)}$.

There is a Maslov index $\mu : \pi_2(\mathbf{x}, \mathbf{y}) \to \mathbb{Z}$ which computes the expected dimension of the moduli space $\mathcal{M}(\phi)$. The Maslov formula is preserved by splicing disks together, i.e. if $\phi_1 \in \pi_2(\mathbf{x}, \mathbf{y})$ and $\phi_2 \in \pi_2(\mathbf{y}, \mathbf{z})$, then

$$\mu(\phi_1 * \phi_2) = \mu(\phi_1) + \mu(\phi_2).$$

See Section 4 of [Lip06] for technicalities on the index. For our purposes, it will be more useful to use combinatorial formulas for the index, which we will discuss in Section 2.3.4.

2.3.2 The colored Heegaard Floer complexes

Given a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we define the local multiplicity at a point $p \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ to be the algebraic intersection number

$$n_p(\phi) = \#(\phi^{-1}(\{p\} \times I \times \mathbb{R})).$$

Given an ℓ–tuple, $\mathbf{w}$, we define $n_{\mathbf{w}} : \pi_w(\mathbf{x}, \mathbf{y}) \to \mathbb{Z}^\ell$ by

$$n_{\mathbf{w}}(\phi) = (n_{w_1}(\phi), \dots, n_{w_\ell}(\phi)).$$

An essential feature of Heegaard Floer homology is that it is graded over Spin^c structures. Indeed the admissibility requirements are an unfortunate necessity as a single diagram can often not be chosen to compute all flavors of HF° for all Spin^c structures. There is a map from intersection points to Spin^c structures:

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

(cf. eg. [OS08, Section 3.3])

To each basepoint $w \in \mathbf{w}$, we associate the formal variable U_w. To achieve functoriality of the graph cobordism maps, we must consider colorings of the basepoints, i.e. Heegaard

Floer complexes where we have formally set some of the variables equal to each other before taking homology. We also need to allow formal variables which don't correspond to any basepoints:

Definition 2.3.2. A **coloring**, $(\sigma, \mathfrak{P})$ of a collection of basepoints $\mathbf{w}$ is a finite index set $\mathfrak{P}$ and a function $\sigma : \mathbf{w} \to \mathfrak{P}$. A **coloring** of a graph Γ is a map $\sigma : C(\Gamma) \to \mathfrak{P}$, where $C(\Gamma)$ denotes the set of connected components of Γ.

We remark that a coloring $\sigma : \mathbf{w} \to \mathfrak{P}$ need not be surjective. When there is no chance of confusion, we will omit writing $\mathfrak{P}$ and just write σ for a coloring.

We will use the notation

$$\mathbb{Z}_2[U_{\mathbf{w}}] \stackrel{\text{def}}{=} \mathbb{Z}_2[U_{w_1}, \dots, U_{w_\ell}]$$

and

$$\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}] \stackrel{\text{def}}{=} \mathbb{Z}_2[U_{w_1}, \dots, U_{w_\ell}, U_{w_1}^{-1}, \dots, U_{w_\ell}^{-1}].$$

Similarly if $\mathfrak{P} = \{p_1, \dots, p_n\}$, we define

$$\mathbb{Z}_2[U_{\mathfrak{P}}] = \mathbb{Z}_2[U_{p_1}, \dots, U_{p_n}]$$

and $\mathbb{Z}_2[U_{\mathfrak{P}}, U_{\mathfrak{P}}^{-1}]$ similarly.

We make the following definition:

Definition 2.3.3. Given a coloring $(\sigma, \mathfrak{P})$ of basepoints $\mathbf{w}$, we let $\mathcal{C}_{\sigma,\mathfrak{P}}$ denote the $\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathfrak{P}}]$–module defined by

$$C_{\sigma,\mathfrak{P}} = \mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathfrak{P}}]/I_{\sigma,\mathfrak{P}}$$

where $I_{\sigma,\mathfrak{P}}$ is the ideal generated by elements of the form $U_w - U_{\sigma(w)}$.

We can now define the Heegaard Floer complexes. Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ be a strongly $\mathfrak{s}$–admissible diagram for $(Y, \mathbf{w}, \mathfrak{s})$ (see Section 2.3.5 for a discussion of admissibility). We let $\widehat{CF}(\mathcal{H}, \mathfrak{s})$ denote the free $\mathbb{Z}_2$–module generated by $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. The coloring has no effect on the hat flavor.

For the other flavors, we first define the uncolored complexes. Define $CF^\infty(\mathcal{H},\mathfrak{s})$ to be the $\mathbb{Z}_2[U_{\mathbf{w}},U_{\mathbf{w}}^{-1}]$–module generated by intersection points $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, i.e. formal sums of elements of the form $\mathbf{x}\cdot U_{w_1}^{i_1}\cdots U_{w_\ell}^{i_\ell}$ where $i_1,\dots,i_\ell$ are arbitrary integers, and $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ is an intersection point with $s_{\mathbf{w}}(\mathbf{x})=\mathfrak{s}$. Define $CF^-(\mathcal{H},\mathfrak{s})$ to be the free $\mathbb{Z}[U_{\mathbf{w}}]$–submodule generated by intersection points $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$, i.e. $CF^-(\mathcal{H},\mathfrak{s})$ consists of sums of elements of the form $\mathbf{x}\cdot U_{w_1}^{i_1}\cdots U_{w_\ell}^{i_\ell}$ where $i_1,\dots,i_\ell\geq 0$ and $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})=\mathfrak{s}$. We define $CF^+(\mathcal{H},\mathfrak{s})=CF^\infty(\mathcal{H},\mathfrak{s})/CF^-(\mathcal{H},\mathfrak{s})$.

We finally define the colored complexes

$$CF^-(\mathcal{H},\sigma,\mathfrak{P},\mathfrak{s})=\mathcal{C}_{\mathfrak{P},\sigma}\otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]}CF^-(\mathcal{H},\mathfrak{s}),$$

and similarly for the $+$, and ∞ flavors.

Let J_s be an almost complex structure on $\Sigma\times I\times\mathbb{R}$ which is appropriately generic. For $\widehat{CF}$, define

$$\widehat{\partial}_{\mathcal{H},\mathfrak{s},J_s}(\mathbf{x})=\sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1\\ n_{\mathbf{w}}(\phi)=\mathbf{0}}}\#\widehat{\mathcal{M}}(\phi)\cdot\mathbf{y}.$$

For the other flavors, define

$$\partial^\infty_{\mathcal{H},\mathfrak{s},J_s}(\mathbf{x})=\sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta}\sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}}\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{y},$$

which induces differentials on CF^- and CF^+.

A Gromov compactness argument, (cf. [OS08, Lemma 4.3]) shows that $\partial^\infty_{\mathcal{H},\mathfrak{s},J_s}$ square to zero.

2.3.3 Holomorphic triangle maps

Given three sets of curves $\boldsymbol{\alpha},\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ inside of a surface Σ, one can consider maps which count holomorphic triangles inside of $\mathrm{Sym}^{g+\ell-1}(\Sigma)$ with boundary on $\mathbb{T}_\alpha,\mathbb{T}_\beta$ and $\mathbb{T}_\gamma$. Following [Lip06], one can also compute these maps in the cylindrical formulation, which we briefly describe in this section.

A **multipointed Heegaard triple**, $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$, is a collection where each $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are collections of $(g + \ell - 1)$ pairwise disjoint, simple closed curves which each span a g–dimensional sublattice of $H_1(\Sigma; \mathbb{Z})$. We also require each σ_i and τ_j to be transverse for $\sigma, \tau \in \{\alpha, \beta, \gamma\}$. We also will need to assume an admissibility assumption, which we discuss in Section 2.3.5.

The triple $\mathcal{T}$ specifies three Heegaard diagrams, which we denote by $\mathcal{T}_{\alpha\beta}$, $\mathcal{T}_{\alpha\gamma}$ and $\mathcal{T}_{\beta\gamma}$. By attaching handles along these curves and then additional 0– and 3–handles, the triple specifies 3 closed manifolds $Y_{\alpha\beta}$, $Y_{\alpha\gamma}$, and $Y_{\beta\gamma}$.

The triple also specifies a compact 4–manifold $X_{\alpha\beta\gamma}$ with boundary

$$Y_{\alpha\beta} \sqcup Y_{\beta\gamma} \sqcup -Y_{\alpha\gamma}$$

obtained as follows: let U_α, U_β, U_γ denote the handlebodies specified by attaching handles along the α, β or γ curves respectively. Let $\overline{\Delta}$ be the closed triangle with edges e_α, e_β and e_γ. By gluing $U_\alpha \times [0,1], U_\beta \times [0,1]$ and $U_\gamma \times [0,1]$ to $\Sigma \times \overline{\Delta}$, we obtain $X_{\alpha\beta\gamma}$.

We let $v_{\alpha\beta}, v_{\beta\gamma}$ and $v_{\alpha\gamma}$ denote the vertices of $\overline{\Delta}$ and we let W_Δ denote $\Sigma \times \Delta$. Given intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta, \mathbf{y} \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, and $\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, we let $\pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ denote homology classes of maps $\psi : (S, \partial S) \to (\Sigma \times \Delta, \boldsymbol{\alpha} \times e_\alpha \cup \boldsymbol{\beta} \times e_\beta \cup \boldsymbol{\gamma} \cup e_\gamma)$ where S is a Riemann surface with $3(g + \ell - 1)$ punctures on the boundary, labeled as $p_i^{\alpha\beta}, p_i^{\beta\gamma}$, and $p_i^{\alpha\gamma}$. We further require that ψ extends continuously to the compactification $\overline{\psi} : (\overline{S}, \overline{\partial S}) \to \Sigma \times \overline{\Delta}$ such that $\pi_{\overline{\Delta}}\overline{\psi}(p_i^{\alpha\beta}) = v_{\alpha\beta}$ and similarly for the $p_i^{\alpha\gamma}$ and $p_i^{\beta\gamma}$ points.

To compute the triangle maps in the cylindrical setting, we wish to count holomorphic representatives of a homology class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$. We let $\mathcal{M}(\psi)$ denote the set of holomorphic curves $u : (S, \partial S) \to W_\Delta$ for an appropriately generic almost complex structure J on W_Δ satisfying the following:

1. u is homologous to ψ after forgetting the almost complex structure;
2. u is an embedding;
3. u has finite energy;

4. for each $\sigma \in \{\alpha, \beta, \gamma\}$ and $i = 1, \dots, (g+\ell-1)$, the set $u^{-1}(\sigma_i \times e_\sigma)$ consists of exactly one component of ∂S;

5. u is nonconstant on each component of S.

In general, the triangle maps in Heegaard Floer homology must be defined with reference to a particular 4–dimensional Spin^c structure. Following [OS06, Sec. 2.2], or [Lip06, Sec. 10.1.2], there is a map

$$s_{\mathbf{w}} : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathrm{Spin}^c(X_{\alpha\beta\gamma}),$$

such that

$$s_{\mathbf{w}}(\psi)|_{Y_{\alpha\beta}} = s_{\mathbf{w}}(\mathbf{x})$$

and similarly for $\mathbf{y}$ and $\mathbf{z}$.

Using this, we can thus define triangle maps

$$\begin{aligned}
\widehat{f}_{\alpha\beta\gamma,\mathfrak{s}} &: \widehat{CF}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}|_{\alpha\beta}) \otimes_{\mathbb{Z}_2} \widehat{CF}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to \widehat{CF}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
f^{\infty}_{\alpha\beta\gamma,\mathfrak{s}} &: CF^{\infty}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}|_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}]} CF^{\infty}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{\infty}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
f^{-}_{\alpha\beta\gamma,\mathfrak{s}} &: CF^{-}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]} CF^{-}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{-}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
f^{+-}_{\alpha\beta\gamma,\mathfrak{s}} &: CF^{+}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]} CF^{-}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{+}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),\\
f^{-+}_{\alpha\beta\gamma,\mathfrak{s}} &: CF^{-}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}]} CF^{+}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to CF^{+}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma}),
\end{aligned}$$

which we define using the formulas

$$\widehat{f}_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x} \otimes \mathbf{y}) = \sum_{\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\substack{\psi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \\ \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s} \\ \mu(\psi) = 0 \\ n_{\mathbf{w}}(\psi) = \mathbf{0}}} \#\mathcal{M}(\psi) \cdot \mathbf{z}$$

and

$$f^{\circ}_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x} \otimes \mathbf{y}) = \sum_{\mathbf{z} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma} \sum_{\substack{\psi \in \pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \\ \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s} \\ \mu(\psi) = 0}} \#\mathcal{M}(\psi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)} \cdot \mathbf{z}$$

for $\circ = +, -, \infty$.

Taking homology, we get a map

$$\widehat{F}_{\alpha\beta\gamma,\mathfrak{s}} : \widehat{HF}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}_{\alpha\beta}) \otimes_{\mathbb{Z}_2} \widehat{HF}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_{\beta\gamma}) \to \widehat{HF}(\mathcal{T}_{\alpha\gamma}, \mathfrak{s}|_{\alpha\gamma})$$

as well as maps $F^{\infty}_{\alpha\beta\gamma,\mathfrak{s}}, F^{-}_{\alpha\beta\gamma,\mathfrak{s}}, F^{+-}_{\alpha\beta\gamma,\mathfrak{s}}$ and $F^{-+}_{\alpha\beta\gamma,\mathfrak{s}}$ defined analogously.

2.3.4 Acceptable domains and combinatorial Maslov indices

One of the key features of Heegaard Floer homology is that homotopy classes of disks and triangles, as well as their Maslov indices can be computed completely combinatorially.

2.3.4.1 Domains

Let $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ have connected components $r_1, \ldots, r_n$ and let R_i be the closure of r_i. We call R_i **regions**. A **domain** is a formal $\mathbb{Z}$–linear combination of the R_i. To a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we can associate a domain

$$\mathcal{D}(\phi) = \sum_{i=1}^{n} n_{p_i}(\phi) R_i,$$

where p_i is any point in the region R_i. The domain of a triangle $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is defined analogously.

Given an intersection point $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ and an intersection point $x \in \alpha_i \cap \beta_j$, we define

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

Suppose $\mathcal{D}$ is a domain, and $x \in \alpha_i \cap \beta_j$ is an intersection point. There are four regions regions with a corner at x. While we traverse counterclockwise a small circle centered at x, in two of these regions we start at α_i and end at β_j. Let $m^{\alpha\beta}(\mathcal{D}, x)$ be the sum of the multiplicities in these regions. Similarly in the other two regions, we start at β_j and end at α_i. We let $m^{\beta\alpha}(\mathcal{D}, x)$ denote the sum of the multiplicities in these regions.

We say that a domain $\mathcal{D}$ is **acceptable for** $(\mathbf{x}, \mathbf{y})$ if

$$m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = \delta(\mathbf{x}, x) - \delta(\mathbf{y}, x) \tag{2.1}$$

for all i, and j and all $x \in \alpha_i \cap \beta_j$.

Lemma 2.3.4. *There is a bijection between domains $\mathcal{D}$ which are acceptable for $(\mathbf{x}, \mathbf{y})$ and elements of $\pi_2(\mathbf{x}, \mathbf{y})$.*

There is a similar relationship between acceptable domains and triangles.

Lemma 2.3.5. *There is a bijection between classes in $\pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ and $\mathcal{D}$ which satisfy*

$$m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = \delta(\mathbf{x}, x),$$
$$m^{\beta\gamma}(\mathcal{D}, y) - m^{\gamma\beta}(\mathcal{D}, y) = \delta(\mathbf{y}, y),$$
$$m^{\gamma\alpha}(\mathcal{D}, z) - m^{\alpha\gamma}(\mathcal{D}, z) = \delta(\mathbf{z}, z)$$

for all intersection points $x, y,$ and z.

Definition 2.3.6. We say that a homology disk $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ is a **periodic class** if $n_{\mathbf{w}}(\phi) = \mathbf{0}$. We say a domain $\mathcal{D}$ is a **periodic domain** if $m^{\alpha\beta}(\mathcal{D}, x) - m^{\beta\alpha}(\mathcal{D}, x) = 0$ for all $x \in \alpha_i \cap \beta_j$ and $n_{\mathbf{w}}(\phi) = \mathbf{0}$.

2.3.4.2 Combinatorial Maslov indices

The Maslov index $\mu : \pi_2(\mathbf{x}, \mathbf{y}) \to \mathbb{Z}$ and $\mu : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathbb{Z}$ compute the expected dimension of the moduli spaces of holomorphic triangles. The definition of μ, (cf. eg. [OS04e], [Lip06]) involves taking the index of the pseudodifferential $\overline{\partial}$ operator, which normally is not computationally tractable, but for our purposes it is sufficient to know that there is a combinatorial computation of the Maslov index, due in part to Lipshitz in [Lip06] and Sarkar in [Sar11b].

One first defines the Euler measure of a region R by

$$e(R) = \chi(R) - ac(R)/4 + ob(R)/4$$

where $ac(R)$ denotes the number of acute angles of R, and $ob(R)$ denotes the number of obtuse corners of R. We then extend e linearly to all domains.

A first formula for the Maslov index of a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ is due to [Lip06, Cor. 4.3]:

$$\mu(\phi) = e(\mathcal{D}) + \sum_{x \in \mathbf{x}} n_x(\mathcal{D}) + \sum_{y \in \mathbf{y}} n_y(\mathcal{D}), \tag{2.2}$$

where $\mathcal{D} = \mathcal{D}(\phi)$ and $n_p(\mathcal{D})$ denotes the average of the multiplicities in the four regions with vertex $p \in \alpha_i \cap \beta_j$.

Another useful formula which we will need holds for certain classes of homology classes. Let $\pi_2^{\alpha}(\mathbf{x})$ denote homology classes of disks with boundary in $\boldsymbol{\alpha}$ and define $\pi_2^{\beta}(\mathbf{x})$ similarly. For $\phi \in \pi_2^{\beta}(\mathbf{x})$ or $\phi \in \pi_2^{\alpha}(\mathbf{x})$, following [OS08, Lem. 5.4], we have

$$\mu(\phi) = 2 \sum_{i=1}^{\ell} n_{w_i}(\phi). \tag{2.3}$$

2.3.5 Admissibility

In order to compute HF°, one must use diagrams with certain admissibility properties. To compute HF^{+} and $\widehat{HF}$ one needs a weaker form of admissibility, usually referred to as *weak admissibility* or just *admissibility*. To compute HF^{-} and HF^{∞}, one needs to use a stronger form of admissibility, usually referred to as *strong* $\mathfrak{s}$*–admissibility*, which is specified to a particular Spin^c structure $\mathfrak{s}$. To compute cobordism maps in our framework, it will be sufficient to use weak admissibility for $\widehat{HF}$, but we will need strong admissibility for HF^{+}, HF^{-} and HF^{∞}. The reason we can't use weak admissibility for HF^{+} is essentially due to technical neck stretching arguments used to define the 1– and 3–handle maps.

2.3.5.1 Admissibility for Heegaard diagrams

Before we state the admissibility condition, we must first recall the Chern class formula for the Maslov index proven by Ozsváth and Szabó.

In their analysis of topological disks, Ozsváth and Szabó produce the following exact sequence [OS08, Eq. 4]:

$$0 \to \mathbb{Z} \to \pi_2^{\alpha}(\mathbf{x}) \oplus \pi_2^{\beta}(\mathbf{x}) \to \pi_2(\mathbf{x}, \mathbf{x}) \to H_2(Y; \mathbb{Z}) \to 0.$$

Given a $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$, we let $H(\phi)$ denote its image in $H_2(Y; \mathbb{Z})$ under the above map. The following Lemma is essentially [OS04d, Prop. 7.5], except for multipointed diagrams:

Lemma 2.3.7. *Suppose that $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$. Then*

$$\mu(\phi) = \langle c_1(\mathfrak{s}), H(\phi) \rangle + 2 \sum_{i=1}^{\ell} n_{w_i}(\phi).$$

Proof. This follows analogously to the proof of the case of a single basepoint ([OS04d, Prop. 7.5]). More precisely, instead of equation (5) in [OS04d], we replace $n_z(\mathcal{P} + m[\Sigma])$ with $\sum_{i=1}^{\ell} n_{w_i}(\mathcal{P} + m[\Sigma])$, since we must remove balls centered around the trajectories from index 0 to index 3 critical points passing through each basepoint w_i.

Performing the analogous computation as in [OS04d, Prop. 7.5] to produce their equation (5), we thus have if $\mathcal{D} = \mathcal{D}(\phi)$ is the domain of a $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ (not necessarily periodic), then

$$\langle c_1(\mathfrak{s}), H(\phi) \rangle = e(\mathcal{D} + m[\Sigma]) + 2 \left(\sum_{x \in \mathbf{x}} n_x(\mathcal{D} + m[\Sigma]) \right) - 2 \left(\sum_{i=1}^{\ell} n_{w_i}(\mathcal{D} + m[\Sigma]) \right)$$

for sufficiently large m. By expanding this and using Equations (2.2) and (2.3), we recover the desired formula. □

The strong notion of admissibility we will use for Heegaard diagrams is the following:

Definition 2.3.8. A diagram $\mathcal{H}$ is **strongly $\mathfrak{s}$–admissible** for a Spin^c structure $\mathfrak{s}$ if for every nontrivial periodic domain $\mathcal{P}$ with

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

then $\mathcal{P}$ has some local multiplicity strictly greater than N.

The following is only a small modification of [OS04e, Lem. 4.14]:

Lemma 2.3.9. *If $\mathcal{H}$ is strongly admissible for $\mathfrak{s}$, then for a fixed j there are only finitely many disks $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\phi) = j$ with $\mathcal{D}(\phi) \geq 0$.*

Proof. Fix any $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$. Any other disk $\psi \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\mu(\psi) = j$ satisfies

$$\mathcal{D}(\phi) = \mathcal{D}(\psi) + \mathcal{P} + \mathcal{A}$$

for a periodic domain $\mathcal{P}$ and a domain $\mathcal{A}$ which is a $\mathbb{Z}$–linear combination of the components of $\Sigma \setminus \boldsymbol{\alpha}$. If we assume $\mathcal{D}(\phi) \geq 0$, then we have

$$-\mathcal{D}(\psi) \leq \mathcal{P} + \mathcal{A}.$$

Suppose we have an infinite sequence of distinct ϕ_n. This yields an infinite sequence of distinct pairs $(\mathcal{P}_i, \mathcal{A}_i)$ satisfying the above inequality.

An infinite sequence of distinct such domains $\mathcal{P}_i + \mathcal{A}_i$ would then need to satisfy $\|\mathcal{P}_i + \mathcal{A}_i\| \to \infty$ where $\|\cdot\|$ denotes the Euclidean 2–norm, since they are elements of a finitely generated integer lattice. By renormalizing, note that both $\mathcal{P}_i + \mathcal{A}_i$ and $\mathcal{A}_i$ are bounded below, so we can extract a convergent subsequences so that $\mathcal{P}_i/\|\mathcal{P}_i + \mathcal{A}_i\|$ and $\mathcal{A}_i/\|\mathcal{P}_i + \mathcal{A}_i\|$ both converge to real domains $\mathcal{P}_\infty$ and $\mathcal{A}_\infty$ respectively. Since

$$0 = \mu(\mathcal{P}_i + \mathcal{A}_i) = \mu(\mathcal{A}_i) + \langle c_1(\mathfrak{s}), H(\mathcal{P}_i)\rangle = \left(\sum_{j=1}^{\ell} 2n_{w_j}(\mathcal{A}_i)\right) + \langle c_1(\mathfrak{s}), H(\mathcal{P}_i)\rangle$$

and the above equation is preserved by normalization, we know that

$$\left(\sum_{j=1}^{\ell} 2n_{w_j}(\mathcal{A}_\infty)\right) + \langle c_1(\mathfrak{s}), H(\mathcal{P}_\infty)\rangle = 0.$$

Since $\mathcal{A}_i$ were bounded below, we know that $\mathcal{A}_\infty \geq 0$. It is easy to see that given a real valued such pair $(\mathcal{A}_\infty, \mathcal{P}_\infty)$, there is a rational one $(\mathcal{A}, \mathcal{P})$ also with $\mathcal{A} \geq 0$ and $\mathcal{A} + \mathcal{P} \geq 0$. By clearing denominators we can assume that $\mathcal{A}$ and $\mathcal{P}$ are integral. We have, however, that

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

where $N = \sum_{j=1}^{\ell} n_{w_j}(\mathcal{A}) \geq 0$. Hence $-\mathcal{P}$ has a multiplicity which is greater than N by strong admissibility. Since $\mathcal{A} \geq 0$, this contradicts the fact that $\mathcal{A} + \mathcal{P} \geq 0$. □

2.3.5.2 Admissibility for higher multidiagrams

Ozsváth and Szabó produced similar admissibility conditions for Heegaard triples ([OS04e, Definition 8.8]), and Heegaard quadruples ([OS04e, Section 8.4.2.]). Instead of stating these directly, we will state a slightly more general description of higher Heegaard multidiagrams.

Definition 2.3.10. An abstract **Heegaard multi-diagram** $(\Sigma, \boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_n, \mathbf{w})$ is a surface Σ with sets of curves $\boldsymbol{\eta}_i = \{\eta_i^1, \dots, \eta_i^{g+\ell-1}\}$ such that

1. Σ is a closed, oriented, genus g surface with $\mathbf{w} \subseteq \Sigma$;
2. $\boldsymbol{\eta}_i = (\eta_i^1, \dots, \eta_i^{g+\ell-1})$ is a collection of disjoint, simple closed curves on Σ;
3. no more than two curves from any collection intersect at a single point;
4. all curves are transverse to each other.

Given a Heegaard multi-diagram $\mathcal{M} = (\Sigma, \boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_n, \mathbf{w})$, we can construct a compact 4–manifold $X_{\boldsymbol{\eta}_0 \dots \boldsymbol{\eta}_n}$ by gluing copies of $U_{\boldsymbol{\eta}_i} \times e_{\boldsymbol{\eta}_i}$ to $G_n \times \Sigma$, where G_n is the regular $(n+1)$–gon, with sides labelled $e_{\boldsymbol{\eta}_i}$, counterclockwise. The manifold $X_{\boldsymbol{\eta}_0 \dots \boldsymbol{\eta}_n}$ has boundary equal to

$$Y_{\boldsymbol{\eta}_0 \boldsymbol{\eta}_1} \sqcup \dots \sqcup Y_{\boldsymbol{\eta}_{n-1} \boldsymbol{\eta}_n} \sqcup Y_{\boldsymbol{\eta}_n \boldsymbol{\eta}_0}.$$

Similar to the map from intersection points on a Heegaard diagram to three dimensional Spin^c structures, there is a map

$$\mathfrak{s}_{\mathbf{w}} : \pi_2(\mathbf{x}_0, \dots, \mathbf{x}_n) \to \mathrm{Spin}^c(X_{\boldsymbol{\eta}_0 \dots \boldsymbol{\eta}_n})$$

such that

$$\mathfrak{s}_{\mathbf{w}}(\phi)|_{Y_{\boldsymbol{\eta}_j \boldsymbol{\eta}_{j+1}}} = \mathfrak{s}_{\mathbf{w}}(\mathbf{x}_j).$$

Such a map is defined for triangles and rectangles in [OS04e] (and those are the only cases that we will need), and for higher polygons in [?GW]. The following is analogous to [?GW, Proposition 3.9].

Lemma 2.3.11. *Let $\psi \in \pi_2(\mathbf{x}_0, \dots, \mathbf{x}_n)$ and $\psi' \in \pi_2(\mathbf{x}'_0, \dots, \mathbf{x}'_n)$ be two topological $(n+1)$–gons. Then $\mathfrak{s}(\psi) = \mathfrak{s}(\psi')$ iff*

$$\psi' = \psi + \phi_0 + \cdots \phi_n$$

for classes $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{x}'_i)$.

Proof. This follows as in [OS04e, Prop. 8.5] for triangles, or [?GW, Prop. 3.9] for higher n–gons. □

For 3–gons and 4–gons, Ozsváth and Szabó define notions of strong admissibility. In more generality, following [?GW] we can generalize these notions to higher n–gons. To be able to use counts of holomorphic n–gons to define maps for $n \geq 4$, we need a strong admissibility condition not just for a single Spin^c structure, but instead for certain families of Spin^c structures.

Given an n–gon G, we can consider a collection of disjoint arcs A on G such that each arc intersects exactly two edges of G, no two arcs intersect each other, and such that each component of $G \setminus A$ intersects exactly three sides of ∂G. We will call such a collection a **complete set of dividing arcs on** G. Given a complete set of dividing arcs, we get a decomposition of $X_{\boldsymbol{\eta}_0 \cdots \boldsymbol{\eta}_n}$ into triangular cobordisms of the form $X_{\boldsymbol{\eta}_i \boldsymbol{\eta}_j \boldsymbol{\eta}_k}$. We call such a decomposition of $X_{\boldsymbol{\eta}_0 \cdots \boldsymbol{\eta}_n}$ a **complete decomposition into triangular cobordisms**. An example of a decomposition for the cobordism $X_{\alpha\beta\gamma\delta}$ is shown in Figure 2.2.

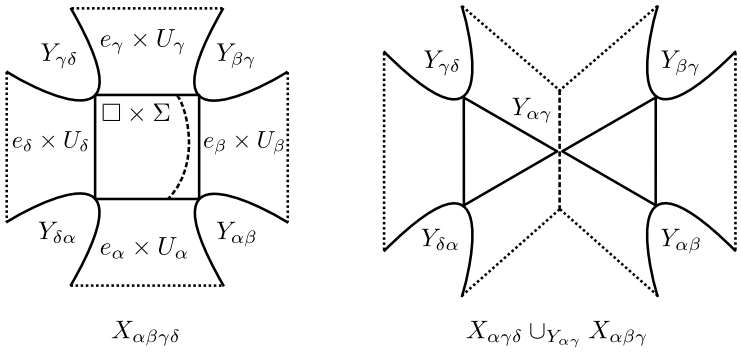

Figure 2.2: Decomposing the cobordism $X_{\alpha\beta\gamma\delta}$ into two triangular cobordisms.

Definition 2.3.12. We define two Spin^c structures $\mathfrak{s}, \mathfrak{s}' \in \text{Spin}^c(X_{\boldsymbol{\eta}_0 \cdots \boldsymbol{\eta}_n})$ to be **equivalent** if for every decomposition of $X_{\boldsymbol{\eta}_0 \cdots \boldsymbol{\eta}_n}$ into triangular cobordisms, we have

$$\mathfrak{s}|_{X_\Delta} = \mathfrak{s}'|_{X_\Delta}$$

for each of the triangular cobordisms X_Δ. We use the letter $\mathfrak{S}$ to denote such an equivalence class of Spin^c structures.

Remark 2.3.13. Given a rectangular cobordism $X_{\alpha\beta\gamma\delta}$, there are exactly two complete decompositions of $X_{\alpha\beta\gamma\delta}$, one of which is given by the arc from e_γ to e_α, and the other is given by the arc from e_δ to e_β. Using the Meyer-Vietoris sequence, we thus see that equivalence classes $\mathfrak{S}$ in Spin$^c(X_{\alpha\beta\gamma\delta})$ correspond exactly to $\delta H^1(Y_{\alpha\gamma}) + \delta H^1(Y_{\beta\delta}))$ orbits within Spin$^c(X_{\alpha\beta\gamma\delta})$, so our definition coincides for rectangles with the definition given in [OS04e]. Similarly note that there is only one triangular decomposition of a triangle cobordism, and hence two Spin^c structures are equivalent on a triangular cobordism iff they are equal.

Definition 2.3.14. We say that a multi-diagram $\mathcal{M} = (\Sigma, \boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_n, \mathbf{w})$ is strongly $\mathfrak{S}$–admissible for an equivalence class $\mathfrak{S}$ of Spin^c structures $\mathfrak{s} \in \text{Spin}^c(X_{\boldsymbol{\eta}_0 \cdots \boldsymbol{\eta}_n})$ if for each $\mathfrak{s} \in \mathfrak{S}$, if $\mathcal{P}$ is a domain which can be written as a sum of doubly-periodic domains

$$\mathcal{P} = \sum_{\{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}\} \subseteq \{\boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_n\}} \mathcal{P}_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}},$$

with

$$\sum_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}} \langle c_1(\mathfrak{s}|_{Y_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}}}), H(\mathcal{P}_{\boldsymbol{\eta}_{i_1}\boldsymbol{\eta}_{i_2}}) \rangle = 2N \geq 0,$$

then $\mathcal{P}$ has some local multiplicity strictly greater than N.

As remarked in [OS04e], for a Heegaard quadruple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{w})$, a $\delta H^1(Y_{\alpha\gamma}) + \delta H^1(Y_{\beta\delta})$ orbit may consist of infinitely many Spin^c structures, so we cannot necessarily perform the standard winding trick to achieve strong $\mathfrak{S}$–admissibility. For most of the Heegaard quadruples we will consider, one can achieve strong admissibility by winding. The following is explained in [OS04e, Sec. 8.4.2]:

Remark 2.3.15. If we assume that $\delta H^1(Y_{\beta\delta})|_{Y_{\alpha\gamma}} = 0$ and $\delta H^1(Y_{\alpha\gamma})|_{Y_{\beta\delta}} = 0$, then by winding sufficiently, we can achieve strong $\mathfrak{S}$–admissibility of a Heegaard quadruple.

The following is analogous to [OS04e, Lem. 8.9]:

Lemma 2.3.16. *Suppose that $\mathcal{M} = (\Sigma, \boldsymbol{\eta}_0, \dots, \boldsymbol{\eta}_n, \mathbf{w})$ is a strongly $\mathfrak{s}$–admissible Heegaard multidiagram and suppose that j is fixed. Then there are finitely many homotopy classes of $(n+1)$–gons $\psi \in \pi_2(\mathbf{x}_0, \dots, \mathbf{x}_n)$ with $\mu(\psi) = j, \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$ and $\mathcal{D}(\psi) \geq 0$.*

Proof. The proof is analogous to [OS04e, Lem. 8.9] in the case of triangles. For higher $(n+1)$–gons, the proof is analogous. Fix a Spin^c structure on $X_{\boldsymbol{\eta}_0,\dots,\boldsymbol{\eta}_n}$, an integer j and a reference $(n+1)$–gon ψ_0 with $\mathfrak{s}_{\mathbf{w}}(\psi_0) = \mathfrak{s}$. By Lemma 2.3.11 if ψ is any other $(n+1)$–gon with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, then we can write

$$\psi = \psi_0 + \phi_0 + \cdots + \phi_n$$

with $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{x}_i)$. Note that since each $\phi_i \in \pi_2(\mathbf{x}_i, \mathbf{x}_i)$, as domains we have that each $\phi_i \in \pi_2(\mathbf{y}, \mathbf{y})$ for any $\mathbf{y}$, because the vertex relations in Equation 2.1 are satisfied. Hence in particular we can write

$$\psi = \psi_0 + \mathcal{A} + \mathcal{P}$$

where $\mathcal{A}$ is an α–boundary degeneration and $\mathcal{P}$ is a periodic domain. The proof then proceeds identically to the proof of Lemma 2.3.9.

□

2.3.6 Heegaard moves and naturality of the 3–manifold invariants

Given a multi-based 3–manifold $(Y, \mathbf{w})$, there are many different Heegaard diagrams for $(Y, \mathbf{w})$. The following is standard:

Proposition 2.3.17. *Any two embedded Heegaard diagrams for $(Y, \mathbf{w})$ $(\Sigma_1, \boldsymbol{\alpha}_1, \boldsymbol{\beta}_1, \mathbf{w})$ and $(\Sigma_2, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \mathbf{w})$ can be connected by a sequence of the following moves:*

1. *isotopies of $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves not passing over the $\mathbf{w}$ basepoints;*

2. *handleslides of the $\boldsymbol{\alpha}$ curves across each other or handleslides of the $\boldsymbol{\beta}$ curves across each other;*

3. *(1,2)–stabilizations or destabilizations;*

4. *isotopies of $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ within Y which are fixed on $\mathbf{w}$.*

We use the phrase "$(1,2)$–stabilization" to distinguish from free-stabilization, which we will sometimes call $(0,3)$–stabilization, which occurs when we add a basepoint.

In it's original formulation in [OS04e], to an appropriately generic almost complex structure on $\mathrm{Sym}^k(\Sigma)$, Ozsváth and Szabó defined maps corresponding to moves (1)–(3) as well maps associated to changing the almost complex structure. Given two Heegaard diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$, as well as suitably generic almost complex structures J_i on $\mathrm{Sym}^{g_i}(\Sigma_i)$, Ozsváth and Szabó showed that there exists a chain homotopy equivalence

$$\Phi_{\mathcal{H}_1 \to \mathcal{H}_2} : CF^\circ_{J_1}(\mathcal{H}_1, \mathfrak{s}) \to CF^\circ_{J_2}(\mathcal{H}_2, \mathfrak{s}),$$

which is a composition of maps corresponding to moves of the above form, though they didn't show that the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ was independent of the choice of intermediate diagrams. This was resolved in [JT12] by Juhász and Thurston:

Theorem 2.3.18 ([JT12, Thm. 2.39])**.** *The map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ is independent up to chain homotopy of the choice of intermediate diagrams and almost complex structures.*

We will now describe the maps in more detail. Given sets of attaching curves $\boldsymbol{\nu}$ on Σ, we write $\boldsymbol{\nu} \sim \boldsymbol{\nu}'$ if $\boldsymbol{\nu}$ and $\boldsymbol{\nu}'$ are related by a sequence of handleslides or isotopies which don't cross $\mathbf{w}$. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple with $\boldsymbol{\beta} \sim \boldsymbol{\gamma}$. Note that $X_{\alpha\beta\gamma}$ has a subset of the boundary diffeomorphic to $Y_{\beta\gamma} \simeq \sqcup_i (S^1 \times S^2)^{\#n}$, which can be "filled in" with some number of 3–handles and 4–handles. There is a bijection from Spin^c structures on $Y_{\alpha\beta}$ and Spin^c structures on $X_{\alpha\beta\gamma}$ which extend over the additional 3–handles and 4–handles. Given $\mathfrak{s} \in \mathrm{Spin}^c(Y_{\alpha\beta})$ we let $\overline{\mathfrak{s}}$ denote the unique such extension on $X_{\alpha\beta\gamma}$. A similar comment holds if $\boldsymbol{\alpha} \sim \boldsymbol{\beta}$, and we use the same notation for the extension of a Spin^c structure $\mathfrak{s}$ on $Y_{\beta\gamma}$ to $X_{\alpha\beta\gamma}$.

The following lemma is proven by Juhász and Thurston [JT12, Lem. 9.2.]:

Lemma 2.3.19. *Let $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z})$ be a strongly (resp. weakly) $\mathfrak{s}$–admissible Heegaard triple. If $\boldsymbol{\beta} \sim \boldsymbol{\gamma}$ then $HF^-(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_0)$ and $\widehat{HF}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}_0)$ admit "top" degree generators $\Theta_{\beta\gamma}$, and for $\circ = \widehat{\ }, \infty, -$ we write $\Psi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\gamma}}$ for the map*

$$\Psi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\gamma}} = F^{\circ}_{\alpha\beta\gamma,\bar{\mathfrak{s}}}(\cdot \otimes \Theta_{\beta\gamma}) : HF^{\circ}(\mathcal{T}_{\alpha\beta}, \mathfrak{s}) \to HF^{\circ}(\mathcal{T}_{\alpha\gamma}, \bar{\mathfrak{s}}|_{\alpha\gamma}).$$

Similarly if $\boldsymbol{\alpha} \sim \boldsymbol{\beta}$, we define

$$\Psi^{\boldsymbol{\alpha}\to\boldsymbol{\beta}}_{\boldsymbol{\gamma}} = F^{\circ}_{\alpha\beta\gamma,\bar{\mathfrak{s}}}(\Theta_{\alpha\beta} \otimes \cdot) : HF^{\circ}(\mathcal{T}_{\beta\gamma}, \mathfrak{s}) \to HF^{\circ}(\mathcal{T}_{\alpha\gamma}, \bar{\mathfrak{s}}_{\alpha\gamma}).$$

Juhász and Thurston also reduced the number of types of Heegaard moves which one needs to perform to go between different Heegaard diagrams. Originally Ozsváth and Szabó defined continuation maps $\Gamma^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\gamma}}$ in [OS04e] to be the transition maps for move (1), above. These involved counting pseudo holomorphic disks with dynamic boundary conditions. According to [JT12, Lem. 9.5], if the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$ achieves admissibility, then

$$\Gamma^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'} = \Psi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'}.$$

If we don't assume admissibility, the map $\Gamma^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'}$ is instead equal to a composition of triangle maps over Heegaard triples which do achieve admissibility (cf. [JT12, Prop. 9.9]).

More generally, given admissible diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ and $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w})$ with $\boldsymbol{\alpha} \sim \boldsymbol{\alpha}'$ and $\boldsymbol{\beta} \sim \boldsymbol{\beta}'$, Juhász and Thurston define a map

$$\Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'} : HF^{\circ}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}) \to HF^{\circ}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathfrak{s})$$

which is a composition of triangle maps, generalizing the maps defined for individual moves. In particular this allows us to avoid using the continuation maps.

We now provide notation for the remaining maps. If $\mathcal{H}_2$ is a stabilization of $\mathcal{H}_1$, then we let $\sigma : HF^{\circ}(\mathcal{H}_1, \mathfrak{s}) \to HF^{\circ}(\mathcal{H}_2, \mathfrak{s})$ denote the stabilization isomorphism. If ϕ_t is an isotopy of Y fixing $\mathbf{w}$, and $\phi = \phi_1$, we define $\phi_* : HF^{\circ}(\mathcal{H}, \mathfrak{s}) \to HF^{\circ}(\phi_*(\mathcal{H}), \mathfrak{s})$ by the tautological formula

$$\phi_*(x_1, \ldots, x_{g+\ell-1}) = (\phi_* x_1, \ldots, \phi_* x_{g+\ell-1})$$

for an intersection point $(x_1, \dots, x_{g+\ell-1})$.

We summarize these comments into the following lemma:

Lemma 2.3.20. *Suppose that $\mathcal{H}_1$ and $\mathcal{H}_2$ are two embedded Heegaard diagrams for $(Y, \mathbf{w})$ and J_1 and J_2 are two suitably generic almost complex structures on $W_{\mathbb{D}}$. Then the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2} : CF^\circ_{J_1}(\mathcal{H}_1, \mathfrak{s}) \to CF^\circ_{J_2}(\mathcal{H}_2, \mathfrak{s})$ can be written as a composition of the following maps:*

1. *change of almost complex structure maps $\Phi_{J_s \to J'_s}$;*
2. *triangle maps of the form $F_{\alpha\beta\gamma,\bar{\mathfrak{s}}}(\Theta_{\alpha\beta} \otimes \cdot)$ if $\boldsymbol{\alpha} \sim \boldsymbol{\beta}$ or $F_{\alpha\beta\gamma,\bar{\mathfrak{s}}}(\cdot \otimes \Theta_{\beta\gamma})$ if $\boldsymbol{\beta} \sim \boldsymbol{\gamma}$;*
3. *stabilization maps σ or their inverses σ^{-1};*
4. *diffeomorphism maps ϕ_* for diffeomorphisms $\phi = \phi_1$ where ϕ_t is an isotopy of $(Y, \mathbf{w})$ fixing $\mathbf{w}$ for all t.*

2.4 Homological actions on the Heegaard Floer groups

We now introduce maps on $HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{s})$ which we call the relative homology maps. These roughly correspond to an action of $H_1(Y/\sigma; \mathbb{Z})$ on $HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{s})$ where σ is a coloring of the basepoints $\mathbf{w}$ and by Y/σ we mean the topological space given by identifying basepoints of the same color in Y. Given a path λ in Y between w_1 and w_2 which is immersed in the underlying surface of a Heegaard diagram $\mathcal{H}$ for Y, we define a map $A_\lambda : CF^\circ(\mathcal{H}, \mathfrak{s}) \to CF^\circ(\mathcal{H}, \mathfrak{s})$ which satisfies

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

For complexes colored so that $U_{w_1} = U_{w_2}$, we get that A_λ is a chain map. For a closed loop γ inside of the underlying surface of $\mathcal{H}$, we get the familiar map A_γ, which is a chain map satisfying $A_\gamma \circ A_\gamma \simeq 0$.

These maps generalize the more familiar action of $\Lambda^*(H_1(Y; \mathbb{Z})/\operatorname{Tors})$ on singly pointed Heegaard Floer homology described in, e.g. [OS04e], though there are some key differences. For one thing, the maps A_λ in this section don't always square to zero, and they also don't

always commute, even on homology. Nonetheless they are essential for the construction of graph cobordism maps. The work here is inspired by work in [Lip06] and [Ni14], and is technically very similar to the latter, where a very similar homological action is constructed on sutured Floer homology.

2.4.1 Construction of the relative homology actions

Fix a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ for $(Y, \mathbf{w})$. Suppose that $(K, \partial K) \subseteq (\Sigma \times I \times \mathbb{R}, \mathbf{w} \times \{0,1\} \times \mathbb{R})$ is an embedded surface which is $\mathbb{R}$–invariant. The case where $\partial K = \varnothing$ is allowed. Suppose further that the image $\pi_\Sigma(K)$ doesn't contain any intersection points between $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves and that $K \cap (\boldsymbol{\alpha} \cup \boldsymbol{\beta}) \times \partial I \times \mathbb{R}$ is trivial. One could try to define the map

$$A_K : CF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}) \to CF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s})$$

by

$$A_K(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} (\#K \cap \operatorname{im} \phi) \cdot \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y},$$

where $\#K \cap \operatorname{im} \phi$ is defined as the algebraic intersection number of K and $\operatorname{im} \phi$ as subsurfaces in $W_{\mathbb{D}} = \Sigma \times \mathbb{D}$ with boundaries fixed in

$$((\boldsymbol{\alpha} \cup \boldsymbol{\beta}) \times (e_\alpha \cup e_\beta)) \cup (\{\mathbf{w}\} \times \{0,1\} \times \mathbb{R}).$$

In the following motivating example, let $\rho : I \to [0, 1/2)$ be a smooth function which $\rho(0) = \rho(1) = 0$ and $\rho(t) > 0$ for $t \notin \{0,1\}$.

Example 2.4.1. If $\lambda : I \to \Sigma$ is a smooth immersed path from one basepoint w to another basepoint w', we define the following surface:

$$K_\lambda^{\alpha\alpha} = \{(\lambda(t), 1 - \rho(t), s) : t \in I, s \in \mathbb{R}\}$$

The map $A_{K_\lambda^{\alpha\alpha}}$ is one way to define the relative homology action.

The map $A_{K_\lambda^{\alpha\alpha}}$ defines a map for the path λ, but in proving invariance and relations involving $A_{K_\lambda^{\alpha\alpha}}$ we run into transversality issues involving holomorphic curves with marked

points. Instead we will define a map A_λ which should agree with $A_{K_\lambda^{\alpha\alpha}}$, but is defined more combinatorially. Our approach to defining the relative homology maps is essentially the same as the one taken in [Ni14] in defining an action of relative homology on sutured Floer homology. Note that since $A_{K_\lambda^{\alpha\alpha}}$ counts the algebraic intersection number, the choice of ρ does not affect the value of A_K. As a result, by letting ρ approach the function which is constantly zero, we obtain the following alternate characterization.

Suppose that λ intersects the curves $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ generically, i.e. it intersects each transversely and doesn't intersect any point in any $\alpha_i \cap \beta_j$. Let $p_1, \dots, p_m$ denote the points where λ intersects an $\boldsymbol{\alpha}$ curve. Given a homology disk ϕ, let $d_i(\lambda, \phi)$ denote the difference between the multiplicities of ϕ on the two sides of p_i. Define

$$a(\lambda, \phi) = \sum_{i=1}^{m} d_i(\lambda, \phi).$$

Then define

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

Example 2.4.2. If $\gamma : I \to \Sigma$ is a closed curve, then let $\widehat{\gamma} : I \to \Sigma \times (0,1)$ be an embedded closed curve which is homotopic to $\gamma \times \{1/2\} : I \to \Sigma \times I$. We define the surface

$$K_\gamma = \{(\widehat{\gamma}(t), s), t \in I, s \in \mathbb{R}\}.$$

The map A_{K_γ} was used in the definition of the $H_1(Y; \mathbb{Z})$ action on $HF^\circ(Y, w, \mathfrak{s})$ for singly pointed diagrams (cf. [Lip06, Prop. 8.6.]).

As we did with the relative homology maps, for a closed curve γ we can define maps A_γ more combinatorially by first defining $a(\gamma, \phi)$ as the sum of changes across the $\boldsymbol{\alpha}$ curves as before, and then defining

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

In principle $A_\gamma = A_{K_\gamma}$ and $A_\lambda = A_{K_\lambda}$, though it is easier to work with the maps A_γ than A_{K_γ}, so we will use A_γ and A_λ as the definitions of the (relative) homology maps.

Maps of the form A_K are often not chain maps. For the relative homology maps defined above we have the following:

Lemma 2.4.3. *Suppose that $\lambda : I \to \Sigma$ is an immersed path from w to w' which intersects the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves transversely and has trivial intersection with any $\alpha_i \cap \beta_j$. Then we have the following relation:*

$$A_\lambda \partial + \partial A_\lambda = U_w + U_{w'}.$$

If γ is a closed curve in Σ, then

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

Proof. We will show the first claim about A_λ and leave the necessary modifications for A_γ to the reader. This is analogous to [OS04e, Lem. 4.18]. Suppose that $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ is a homology disk with $\mu(\phi) = 2$. We will sum over the ends of $\widehat{\mathcal{M}}(\phi)$. There are an even number of ends, and hence

$$a(\lambda, \phi) \cdot \#(\partial \widehat{\mathcal{M}}(\phi)) = 0.$$

If $\mathbf{x} \neq \mathbf{z}$ then boundary degenerations cannot occur, so the ends of $\widehat{\mathcal{M}}(\phi)$ have the following form:

$$\partial(\widehat{\mathcal{M}}(\phi)) = \coprod_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi_1 \in \pi_2(\mathbf{x},\mathbf{y}), \phi_2 \in \pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1 \\ \phi_1 * \phi_2 = \phi}} \widehat{\mathcal{M}}(\phi_1) \times \widehat{\mathcal{M}}(\phi_2).$$

Note that if ϕ_1 and ϕ_2 are homology disks, then we have

$$a(\lambda, \phi_1 + \phi_2) = a(\lambda, \phi_1) + a(\lambda, \phi_2).$$

Combining these observations we have

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

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

In the case that $\mathbf{x} = \mathbf{z}$, we have extra contributions from $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ which can have holomorphic representatives degenerate into boundary degenerations. We have the additional counts of boundary degenerations, which add in a count of

$$\sum_{\substack{\phi \in \pi_2^\alpha(\mathbf{x}) \\ \mu(\phi)=2}} a(\lambda, \phi) \# \widehat{\mathcal{N}}^\alpha(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{x} + \sum_{\substack{\phi \in \pi_2^\beta(\mathbf{x}) \\ \mu(\phi)=2}} a(\lambda, \phi) \# \widehat{\mathcal{N}}^\beta(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{x}.$$

If $\phi \in \pi_2^\alpha(\mathbf{x})$ then ϕ can be represented as a sum of domains A_i where A_i is a connected component of $\Sigma \setminus \boldsymbol{\alpha}$. Similarly if $\phi \in \pi_2^\beta(\mathbf{x})$ then the domain of ϕ can be represented as a sum of the domains B_i, where B_i is a connected component of $\Sigma \setminus \boldsymbol{\beta}$. In light of the Maslov formula in Equation 2.3, we know that a nonnegative ϕ in $\pi_2^\alpha(\mathbf{x})$ or $\pi_2^\beta(\mathbf{x})$ with Maslov index 2 is of the form A_i or B_i for some i. On the other hand, it is clear from the definition that $a(\lambda, \phi) = 0$ unless $\phi = A_w$ or $\phi = A_{w'}$, where A_w and $A_{w'}$ are the components of $\Sigma \setminus \boldsymbol{\alpha}$ containing w or w' respectively. We can apply [OS08, Thm. 5.5], which shows that since we have multiple basepoints, we have $\#\widehat{\mathcal{N}}^\alpha(\phi) = 1 \in \mathbb{Z}_2$ if ϕ is a Maslov index 2 element of $\pi_2^\alpha(\mathbf{x})$.

In particular, summing over all $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$ with Maslov index 2, we have that the $\mathbf{x}$ component of

$$(A_\lambda \partial + \partial A_\lambda + U_w + U_{w'})(\mathbf{x})$$

is zero. Combining this with the computation in the case that $\mathbf{x} \neq \mathbf{z}$, we get the desired formula. □

2.4.2 Naturality of the relative homology maps

In this section we prove that the relative homology maps A_λ are independent of the choice of diagram and the representative of the homotopy class of curve λ, up to equivariant chain homotopy.

Lemma 2.4.4. *Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$ is an admissible Heegaard triple. On all of the subdiagrams of $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$, let A_λ denote the map obtained by counting the differences across the $\boldsymbol{\alpha}$–curves. Then*

$$F^\circ_{\alpha\beta\beta',\mathfrak{s}} \circ (A_\lambda \otimes 1) \simeq A_\lambda \circ F^\circ_{\alpha\beta\beta',\mathfrak{s}}.$$

Proof. For a homology triangle ψ, define $a(\lambda, \psi)$ to be the sum of the differences of multiplicities of ψ across any $\boldsymbol{\alpha}$ curve.

First note that $a(\lambda, \phi) = 0$ for any disk ϕ on $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$, as there are no differences over any $\boldsymbol{\alpha}$–curves for such a disk. Also, if ψ is a triangle and ϕ is a disk, then $a(\lambda, \psi + \phi) = a(\lambda, \psi) + a(\lambda, \phi)$. Now given a homology class $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ with $\mu(\psi) = 1$, we sum over the ends of $\mathcal{M}(\psi)$, which has a compactification as a compact 1–manifold. Since any nonconstant boundary degeneration or closed curve has Maslov index at least two, we know that the only degenerations possible correspond to a Maslov index 1 strip breaking off, so the ends are modeled on

$$\coprod_{\substack{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi\in\pi_2(\mathbf{x},\mathbf{w}),\psi'\in\pi_2(\mathbf{w},\mathbf{y},\mathbf{z}) \\ \mu(\phi)=1,\mu(\psi')=0 \\ \phi+\psi'=\psi}} \widehat{\mathcal{M}}(\phi)\times\mathcal{M}(\psi') \sqcup \coprod_{\substack{\mathbf{w}\in\mathbb{T}_\beta\cap\mathbb{T}_{\beta'} \\ \phi\in\pi_2(\mathbf{y},\mathbf{w}),\psi'\in\pi_2(\mathbf{x},\mathbf{w},\mathbf{z}) \\ \mu(\phi)=1,\mu(\psi')=0 \\ \phi+\psi'=\psi}} \widehat{\mathcal{M}}(\phi)\times\mathcal{M}(\psi') \tag{2.4}$$

$$\sqcup \coprod_{\substack{\mathbf{w}\in\mathbb{T}_\alpha\cap\mathbb{T}_{\beta'} \\ \phi\in\pi_2(\mathbf{w},\mathbf{z}),\psi'\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \\ \mu(\phi)=1,\mu(\psi')=0 \\ \phi+\psi'=\psi}} \mathcal{M}(\psi')\times\widehat{\mathcal{M}}(\phi).$$

Define the map $H^\lambda_{\alpha\beta\beta',\mathfrak{s}}$ by

$$H^\lambda_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) = \sum_{\substack{\mathbf{z}\in\mathbb{T}_\alpha\cap\mathbb{T}_{\beta'} \\ \psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \\ \mu(\psi)=0}} a(\lambda,\psi)\#\mathcal{M}(\psi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)}\cdot\mathbf{z}.$$

Since the total count of ends in Equation (2.4) is zero, multiplying the total count by $a(\lambda, \psi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)}$, and then summing over all $\psi \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ with Maslov index one, we get that the $\mathbf{z}$ coefficient of

$$F^\circ_{\alpha\beta\beta',\mathfrak{s}}(A_\lambda(\mathbf{x})\otimes\mathbf{y}) + A_\lambda(F^\circ_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y})) + \partial_{\alpha\beta'}H^\lambda_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) + H^\lambda_{\alpha\beta\beta',\mathfrak{s}}((\partial_{\alpha\beta}\mathbf{x})\otimes\mathbf{y} + \mathbf{x}\otimes(\partial_{\beta\beta'}\mathbf{y}))$$

is equal to zero. Note that there is an additional term, $F^{\circ}_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}\otimes A_\lambda(\mathbf{y}))$, which one would expect, though this term vanishes since $a(\lambda,\phi)=0$ for a homology class on $(\Sigma,\boldsymbol{\beta},\boldsymbol{\beta}',\mathbf{w})$, as there are no changes across any $\boldsymbol{\alpha}$–curves, for such a class. The desired relation follows. $\square$

The above lemma shows that the maps A_λ are independent of $\boldsymbol{\beta}$–isotopies and handleslides, which can be computed with a triangle map. More generally, it can also be used to show that certain 2–handle maps commute with the relative homology maps.

Due to the asymmetry between the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves in the definition of the map A_λ, the proof that A_λ commutes with $\boldsymbol{\alpha}$–isotopies and handleslides is naturally somewhat different than showing that it commutes with $\boldsymbol{\beta}$–isotopies and handleslides:

Lemma 2.4.5. *Suppose that $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ is an admissible Heegaard triple such that $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\mathbf{w})$ represents $(S^1\times S^2)^{\#k}$. Let A_λ denote the endomorphism of $CF^\circ(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ defined by counting differences over $\boldsymbol{\alpha}$–curves, and let A'_λ denote the endomorphism of $CF^\circ(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\beta})$ defied by counting differences across $\boldsymbol{\alpha}'$–curves. If $\Theta^+_{\alpha'\alpha}\in HF^\circ(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\mathbf{w})$ denotes the top degree generator for $\circ\in\{-,\wedge\}$, then*

$$F^{\circ}_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^+_{\alpha'\alpha}\otimes A_\lambda(\cdot))\simeq A'_\lambda\circ F^{\circ}_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^+_{\alpha'\alpha}\otimes\cdot).$$

Proof. For a disk or triangle ϕ on $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\boldsymbol{\beta})$ let $a(\lambda,\phi)$ denote sum of the changes of ϕ across the $\boldsymbol{\alpha}$–curves, and let $a'(\lambda,\phi)$ denote the sum of the changes over the $\boldsymbol{\alpha}'$–curves. For each Maslov index 1 homology triangle $\psi\in\pi_2(\mathbf{x},\mathbf{y},\mathbf{z})$, we count the ends of $\mathcal{M}(\psi)$, multiplying by $a(\lambda,\psi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)}\cdot\mathbf{z}$. As in the previous lemma, (noting that $a(\lambda,\phi)$ vanishes for any disk on $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\beta},\mathbf{w})$), we get exactly that

$$F^{\circ}_{\alpha'\alpha\beta}(A_\lambda\otimes 1)+F^{\circ}_{\alpha'\alpha\beta}(1\otimes A_\lambda)=\partial_{\alpha'\beta}H^\lambda_{A,\alpha'\alpha\beta}(1\otimes 1)+H^\lambda_{A,\alpha'\alpha\beta}(\partial_{\alpha\beta}\otimes 1+\mathbf{1}\otimes\partial_{\beta\beta'}),$$

as maps on the tensor product complex. Here $H^\lambda_{A,\alpha'\alpha\beta}$ counts triangles with multiplicities $a(\lambda,\psi)$ and 1 denotes the identity map. Doing the same, but instead with the coefficients $a'(\lambda,\phi)$, which count differences across the $\boldsymbol{\alpha}'$–curves, we see that

$$A'_\lambda F^{\circ}_{\alpha'\alpha\beta}(1\otimes 1)+F^{\circ}_{\alpha'\alpha\beta}(A'_\lambda\otimes 1)=\partial_{\alpha'\beta}H^\lambda_{A',\alpha'\alpha\beta}(1\otimes 1)+H^\lambda_{A',\alpha'\alpha\beta}(\partial_{\alpha\beta}\otimes 1+1\otimes\partial_{\beta\beta'}).$$

Summing the previous two equations we get

$$A'_\lambda F^\circ_{\alpha'\alpha\beta}(1\otimes 1)+F^\circ_{\alpha'\alpha\beta}(1\otimes A_\lambda)+F_{\alpha'\alpha\beta}((A_\lambda+A'_\lambda)\otimes 1)=\partial_{\alpha'\beta}H+H(\partial_{\alpha'\alpha}\otimes 1+1\otimes\partial_{\alpha\beta}).$$

We now put in $\Theta^+_{\alpha'\alpha}\in CF^\circ(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha})$ (for $\circ\in\{\wedge,-\}$) the top degree generator, and that

$$A'_\lambda F^\circ_{\alpha'\alpha\beta}(\Theta^+_{\alpha'\alpha}\otimes 1)+F^\circ_{\alpha'\alpha\beta}(\Theta^+_{\alpha'\alpha}\otimes A_\lambda)\simeq F_{\alpha'\alpha\beta}((A_\lambda+A'_\lambda)(\Theta^+_{\alpha'\alpha})\otimes 1).$$

As maps from $CF^\circ(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ to $CF^\circ(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\beta})$. The claim will be complete as long as we can show that $(A_\lambda+A'_\lambda)(\Theta^+_{\alpha'\alpha})$ is a boundary in $CF^-(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha})$, since inserting a boundary into one component of the triangle map yields a map of the other component which is chain homotopic to zero (by using associativity relations from the count of the ends of Maslov index 1 triangles). To this end, we note that the map $A_\lambda+A'_\lambda$ counts disks with a factor of $a(\lambda,\phi)+a'(\lambda,\phi)$. On the other hand, for a disk ϕ, the quantity $a(\lambda,\phi)+a'(\lambda,\phi)$ computes the sum of the total changes across either an $\boldsymbol{\alpha}$–curve or a $\boldsymbol{\alpha}'$–curve. As those are all the changes in multiplicities across λ for a disk ϕ on $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha})$, the sum telescopes and we see

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

where w_1 and w_2 are the endpoints of λ, or rather

$$A'_\lambda+A_\lambda=U_{w_1}\Phi_{w_1}+U_{w_2}\Phi_{w_2},$$

where Φ_{w_i} is the map

$$\Phi_{w_i}(\mathbf{x})=U_{w_i}^{-1}\sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} n_{w_i}(\phi)\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{y}.$$

It is easily verified that Φ_{w_i} is a chain map. On the other hand Φ_{w_i} is also $+1$ graded. As a consequence $\Phi_{w_i}(\Theta^+_{\alpha'\alpha})$ is zero in HF^- or $\widehat{HF}$, since $\Theta^+_{\alpha'\alpha}$ is the top graded element. More generally, using the chain homotopy $\Phi_{w_i}=\partial\circ\frac{d}{dU_{w_i}}+\frac{d}{dU_{w_i}}\circ\partial$ on the uncolored complex (see Equation (2.10)), it is not hard to see that $\Phi_{w_i}(x)$ is a boundary if x is cycle on the uncolored complex. □

The map Φ_w will make an appearance again, when we compute the π_1–action in Section 2.13. As a consequence of the previous Lemma, we see that the maps A_λ commute up to equivariant chain homotopy with triangle maps induced by handleslides or isotopies, as well as the 2–handle maps.

Lemma 2.4.6. *If J_s and J_s' are two almost complex structure on $W_{\mathbb{D}}$ which satisfy axioms $(J1)$—$(J6)$ in* [Lip06]*, then*

$$A_\lambda \Phi_{J_s \to J_s'} \simeq \Phi_{J_s \to J_s'} A_\lambda,$$

where $\Phi_{J_s \to J_s'}$ is the change of almost complex structures map.

Proof. Pick an almost complex structure $\widehat{J}$ on $W_{\mathbb{D}}$ which agrees with J_s on $\Sigma \times I \times [1, \infty)$ and with J_s' on $\Sigma \times I \times (-\infty, -1]$, and which satisfies the required axioms in [Lip06] to compute the map $\Phi_{J_s \to J_s'}$. The map $\Phi_{J_s \to J_s'}$ is computed by counting $\widehat{J}$–holomorphic disks of Maslov index zero.

We define a map $H_{\widehat{J},\lambda}$ by the formula

$$H_{\widehat{J},\lambda}(\mathbf{x}) = \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi \in \pi_2(\mathbf{x},\mathbf{y}) \\ \mu(\phi)=0}} a(\lambda, \phi) \# \mathcal{M}_{\widehat{J}}(\phi).$$

Counting the ends of a moduli space of the form $\mathcal{M}_{\widehat{J}}(\phi)$ where $\mu(\phi) = 1$, and then summing over all $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ yields that the $\mathbf{y}$ coefficient of

$$(A_\lambda \Phi_{J_s \to J_s'} + \Phi_{J_s \to J_s'} A_\lambda + \partial H_{\widehat{J},\lambda} + H_{\widehat{J},\lambda} \partial)(\mathbf{x})$$

is zero. □

Lemma 2.4.7. *Suppose that λ and λ' are two paths in Σ starting and ending at the basepoints w and w' which are homotopic relative their endpoints. Then*

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

Proof. The homotopy between λ and λ' can be replaced by a sequence of homotopies of the following form

1. homotopies which don't affect the intersection of λ which any $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves;

2. homotopies which push a piece of λ through a point on $\boldsymbol{\beta} \setminus \boldsymbol{\alpha}$;

3. homotopies which push a piece of λ through a point on $\boldsymbol{\alpha} \setminus \boldsymbol{\beta}$;

4. homotopies which push a piece of λ across an intersection point of $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$.

Homotopies of the first two types trivially don't affect $a(\lambda, \phi)$. It is an easy computation to see that move (3) also doesn't affect $a(\lambda, \phi)$. Move (4) is shown in Figure 2.3.

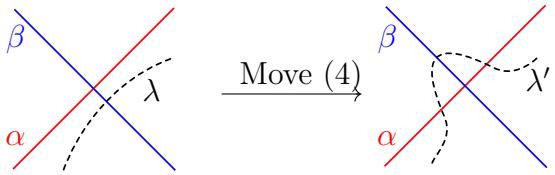

Figure 2.3: Move (4) in Lemma 2.4.7

Let p be the intersection point across which we are pushing λ. Let a, b, c, d be the multiplicities, labeled clockwise, as in Figure 2.4.

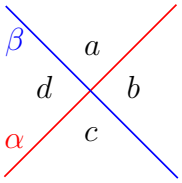

Figure 2.4: Multiplicities around the vertex in Move(4).

We compute that

$$a(\lambda', \phi) - a(\lambda, \phi) = (b - a) + (d - c).$$

Combining this with the vertex relations in Equation (2.1), we see

$$a(\lambda', \phi) - a(\lambda, \phi) = \delta(\mathbf{y}, p) - \delta(\mathbf{x}, p).$$

Thus we have that

$$A_{\lambda'}(\mathbf{x}) - A_{\lambda}(\mathbf{x}) = \sum_{\substack{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta \\ \phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 0}} (\delta(\mathbf{y}, p) - \delta(\mathbf{x}, p)) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}$$

$$= H_p(\partial \mathbf{x}) - \partial(H_p(\mathbf{x})),$$

where $H_p(\mathbf{z}) = \delta(\mathbf{z}, p) \cdot \mathbf{z}$, yielding the desired chain homotopy.

□

Lemma 2.4.8. *Suppose that λ is a path from w_1 to w_2 and λ' is a path from w_2 to w_3. Then*

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

Proof. This follows directly from the formula since $a(\lambda' * \lambda, \phi) = a(\lambda', \phi) + a(\lambda, \phi)$. □

Similarly, we also have the following:

Lemma 2.4.9. *If λ is a path from w to w', and γ and γ' are two closed loops which start and end at w, then we have*

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

and also

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

Proof. This follows as before since $a(\lambda' * \lambda, \phi) = a(\lambda', \phi) + a(\lambda, \phi)$. □

In contrast to the hat case, the relative homology maps do not commute. Instead we have the following:

Lemma 2.4.10. *Suppose λ_1 and λ_2 are two paths between basepoints. If $E(\lambda_i)$ denotes the endpoints of λ_i, then we have the following:*

$$A_{\lambda_1} A_{\lambda_2} + A_{\lambda_2} A_{\lambda_1} = \sum_{w \in E(\lambda_1) \cap E(\lambda_2)} U_w.$$

Proof. We proceed similarly to Lemma 2.4.3. We wish to count the ends of $\widehat{\mathcal{M}}(\phi)$ where $\phi \in \pi_2(\mathbf{x}, \mathbf{z})$ is a Maslov index 2 disk. There are two cases to consider: namely $\mathbf{x} \neq \mathbf{z}$ and $\mathbf{x} = \mathbf{z}$.

In the case that $\mathbf{x} \neq \mathbf{z}$, the ends of $\widehat{\mathcal{M}}(\phi)$ for a Maslov index two ϕ correspond only to strip breaking into two Maslov index one strips. Summing over all such ϕ, we have

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

Computing that

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

we can rewrite the above expression to yield

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

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

The above expression is exactly the $\mathbf{z}$ component of

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

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

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

In the case that $\mathbf{x}=\mathbf{z}$, we have the additional possibility of boundary degenerations. These make an additional contribution of

$$\sum_{\substack{\phi\in\pi_2^\alpha(\mathbf{x}) \\ \mu(\phi)=2}} a(\lambda_1,\phi)a(\lambda_2,\phi)\#\widehat{\mathcal{N}}^\alpha(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x} + \sum_{\substack{\phi\in\pi_2^\beta(\mathbf{x}) \\ \mu(\phi)=2}} a(\lambda_1,\phi)a(\lambda_2,\phi)\#\widehat{\mathcal{N}}^\beta(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x}.$$

We now note that $a(\lambda, \phi) = 0$ whenever $\phi \in \pi_2^{\boldsymbol{\beta}}(\mathbf{x})$ and also that $a(\lambda_1, \phi)a(\lambda_2, \phi)$ is zero unless $\mathcal{D}(\phi) = A_w$ for some $w \in E(\lambda_1) \cap E(\lambda_2)$, in which case it is 1. The desired formula thus follows. □

Lemma 2.4.11. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for Y and that γ is an immersed, closed curve in Σ which intersects no points of $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$ and also intersects the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves transversely. If γ is homotopic to an $\boldsymbol{\alpha}$ curve, or to a $\boldsymbol{\beta}$ curve, then $A_\gamma \simeq 0$.*

Proof. By Lemma 2.4.7 a homotopy of γ results in a chain homotopic A_γ. If γ is homotopic to an $\boldsymbol{\alpha}$ curve, then it is also homotopic to a small push off of a $\boldsymbol{\alpha}$ curve since Σ is orientable. Such a small push off does not intersect any $\boldsymbol{\alpha}$ curves, and hence the quantity $a(\gamma, \phi) = 0$ for all ϕ since $a(\gamma, \phi)$ is an empty sum. Hence $A_\gamma = 0$.

If γ is homotopic to a $\boldsymbol{\beta}$ curve, then by pushing off of the $\boldsymbol{\beta}$ curve, we can assume that γ doesn't intersect any $\boldsymbol{\beta}$ curve. In that case, $a(\gamma, \phi) = 0$ since $a(\gamma, \phi)$ is a sum of differences which telescope since γ doesn't cross any $\boldsymbol{\beta}$ curves. In particular $A_\gamma = 0$. □

In the context of a single basepoint in Heegaard Floer homology, the homology action squares to zero (cf. [OS04e, Prop. 4.17], [Lip06, Prop. 8.6]). If we were to proceed to prove the following lemma using those techniques, we would need to prove a transversality result for curves with two marked points, though it's not clear if such a transversality result even holds in our context. Our proof of the following lemma requires no such transversality result:

Lemma 2.4.12. *Suppose that λ is an immersed path from w to w' in Σ and γ is an immersed, closed loop in Σ. Then*

$$A_\lambda \circ A_\lambda \simeq U_w \simeq U_{w'}$$

and

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

Proof. Note that $U_w \simeq U_{w'}$ by Lemma 2.4.3. We will prove that $A_\lambda \circ A_\lambda \simeq U_w$. First pick an orientation of λ, so that we can assign an integer to each $a(\lambda, \phi)$. Let A_w and $A_{w'}$ be

the components of $\Sigma \setminus \boldsymbol{\alpha}$ containing w and w', respectively. We view A_w and $A_{w'}$ as Maslov index 2 classes of $\boldsymbol{\alpha}$ boundary degenerations. Note that $a(\lambda, A_w)$ and $a(\lambda, A_{w'})$ are both ± 1. In fact, one is $+1$ and one is -1, depending on the orientation of λ. Orient λ so that $a(\lambda, A_w) = +1$ and $a(\lambda, A_{w'}) = -1$.

Our choice of the orientation of λ yields a homotopy to U_w. The other choice yields a chain homotopy to $U_{w'}$. The proof of $A_\lambda \circ A_\lambda \simeq 0$ follows with only minor modification. Consider the chain homotopy

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

Here of course the expression $(a(\lambda,\phi_2)(a(\lambda,\phi_2)+1)/2$ depends on the orientation of λ and is first computed in $\mathbb{Z}$, then projected to $\mathbb{Z}_2$. We claim that

$$A_\lambda^2(\mathbf{x}) = (\partial H_\lambda + H_\lambda\partial + U_w)(\mathbf{x}).$$

In fact we will see that $H_{-\lambda}$ gives the chain homotopy between $A_\lambda^2(\mathbf{x})$ and $U_{w'}$.

As usual we will count the ends of moduli spaces. Let $\phi \in \pi_2(\mathbf{x},\mathbf{z})$ be a homotopy class of disks of Maslov index 2 and first suppose that $\mathbf{x} \neq \mathbf{z}$. Since $\mathbf{x} \neq \mathbf{z}$, the boundary of the 1–dimensional moduli spaces of Maslov index 2 disks corresponds exactly strip breaking. We note that the $\mathbf{z}$–component of $(\partial H_\lambda + H_\lambda\partial)(\mathbf{x})$ is

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \left(\frac{a(\lambda,\phi_1)(a(\lambda,\phi_1)+1)}{2}+\frac{a(\lambda,\phi_2)(a(\lambda,\phi_2)+1)}{2}\right)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1+\phi_2)}.$$

Rearranging, this becomes

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \left(\frac{a(\lambda,\phi_1+\phi_2)(a(\lambda,\phi_1+\phi_2)+1)}{2}\right)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1+\phi_2)}$$

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

The second summand is the $\mathbf{z}$ coefficient of $A_\lambda^2(\mathbf{x})$. Note however

$$\sum_{\substack{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta \\ \phi_1\in\pi_2(\mathbf{x},\mathbf{y}),\phi_2\in\pi_2(\mathbf{y},\mathbf{z}) \\ \mu(\phi_1)=\mu(\phi_2)=1}} \left(\frac{a(\lambda,\phi_1+\phi_2)(a(\lambda,\phi_1+\phi_2)+1)}{2}\right)\#\widehat{\mathcal{M}}(\phi_1)\#\widehat{\mathcal{M}}(\phi_2)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi_1+\phi_2)}=0,$$

is it is as it is the sum over all Maslov index 2 homology disks $\phi\in\pi_2(\mathbf{x},\mathbf{z})$ of

$$\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\#\partial\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)},$$

which is zero. Hence we conclude that the $\mathbf{z}$–component of $(H_\lambda\partial+dH_\lambda+A_\lambda^2)(\mathbf{x})$ is zero.

When we consider $\mathbf{z}=\mathbf{x}$, the computation is of course slightly different. The expression for the $\mathbf{x}$ component of $(\partial H_\lambda+H_\lambda\partial+A_\lambda^2)(\mathbf{x})$ is no longer the sum over all Maslov index 2 disks $\phi\in\pi_2(\mathbf{x},\mathbf{x})$ of

$$\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}\#\partial\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)},$$

because now there are nontrivial Maslov index 2 boundary degenerations which appear in the boundary of that moduli space. Hence, we instead have that the $\mathbf{x}$ component of $(\partial H_\lambda+H_\lambda\partial+A_\lambda^2)(\mathbf{x})$ is equal to

$$\sum_{\substack{\phi\in\pi_2^\alpha(\mathbf{x}) \\ \mu(\phi)=2}}\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}(\#\widehat{\mathcal{N}}^\alpha(\phi))U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x}+\sum_{\substack{\phi\in\pi_2^\beta(\mathbf{x}) \\ \mu(\phi)=2}}\frac{a(\lambda,\phi)(a(\lambda,\phi)+1)}{2}(\#\widehat{\mathcal{N}}^\beta(\phi))U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{x}.$$

We note that for any $\phi\in\pi_2^\beta(\mathbf{x})$, the quantity $a(\lambda,\phi)=0$, as there are no changes of ϕ over any $\boldsymbol{\alpha}$–curves. On the other hand it's easily seen that the only two Maslov index $\phi\in\pi_2^\alpha(\mathbf{x})$ which have nonzero $a(\lambda,\phi)$ are the two with domains A_w and $A_{w'}$. On the other hand, we picked the orientation on λ so that $a(\lambda,A_{w'})(a(\lambda,A_{w'})+1)=0\in\mathbb{Z}$ and $a(\lambda,A_w)(a(\lambda,A_w)+1)=2\in\mathbb{Z}$. It follows that

$$(\partial H_\lambda+H_\lambda\partial+A_\lambda^2)(\mathbf{x})=U_w\cdot\mathbf{x}.$$

The proof easily adapts to compute $A_\gamma\circ A_\gamma\simeq 0$. In that case, $a(\gamma,\phi)=0$ for any homotopy class $\phi\in\pi_2^\alpha(\mathbf{x})$ or $\phi\in\pi_2^\beta(\mathbf{x})$, so boundary degenerations have no contribution. □

We state the following fact, though we will not need it in this paper:

Lemma 2.4.13. *If γ is a closed loop in Σ representing an element of* $\mathrm{Tors} \subseteq H_1(Y)$*, then* $A_\gamma \simeq 0$.

Proof. This can be proven using the same argument in [Ni14, Lem. 2.4], using only a change of notation. □

2.5 Free stabilization maps

In this section we introduce maps associated to adding or removing a basepoint. These correspond to the *creation/termination* maps defined by the author in [Zem15] for the hat flavor. Suppose that $w \notin \mathbf{w}$ is a basepoint, and $(\sigma', \mathfrak{P})$ is a coloring on $\mathbf{w} \cup \{w\}$, which extends the coloring $(\sigma, \mathfrak{P})$ on $\mathbf{w}$. We will define homomorphisms of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–modules

$$S_w^+ : HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S_w^- : HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

which we will call the **free stabilization maps**.

We now summarize these maps. Given an embedded diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ with is strongly $\mathfrak{s}$–admissible, we say that $\mathcal{H}$ is **subordinate to** the point w if $w \in \Sigma$ and w does not intersect any of the $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. Now given a collection of data $\mathcal{J}$ (which we will describe more precisely below), we can form a diagram $\mathcal{H}_{\mathcal{J}}$

$$\mathcal{H}_{\mathcal{J}} = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w\})$$

where α_0 and β_0 are two small isotopic curves each bounding a small ball centered around the point w, which don't intersect any of the other $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves and such that $|\alpha_0 \cap \beta_0| = 2$. The two intersection points of $\alpha_0 \cap \beta_0$ are distinguished by the Maslov index. We let θ^+ and θ^- denote the higher and lower graded intersection points respectively.

For appropriate almost complex structures, we will define the maps $S_w^\pm$ by the formulas

$$S_w^+(\mathbf{x}) = \mathbf{x} \times \theta^+,$$

and

$$S_w^-(\mathbf{x} \times \theta^-) = \mathbf{x}, \qquad \text{and} \qquad S^-(\mathbf{x} \times \theta^+) = 0.$$

The rest of the section is devoted to proving Theorem 2.5.22, which shows that the free stabilization maps are natural.

2.5.1 Gluing data, neck stretching, and gluing moduli spaces

To define the free stabilization maps, as well as the 1– and 3–handle maps, we will need to perform neck stretching and gluing arguments repeatedly. We summarize the procedure in this subsection, and in the following subsection we will embark upon applying it to analyze the free stabilization maps. Consider the diagram $(S^2, \alpha_0, \beta_0, w, w')$ for (S^3, w, w') shown in Figure 2.5.

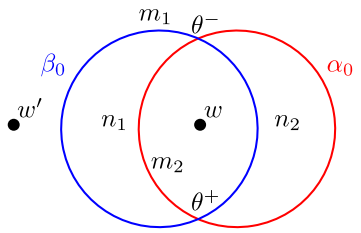

Figure 2.5: The standard diagram for (S^3, w, w'), with multiplicities labeled. The point p' is chosen in the same domain as w'.

We make the following definition:

Definition 2.5.1. Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(Y, \mathbf{w})$ which is subordinate to a point $w \in Y \setminus \mathbf{w}$. We call a collection $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ **gluing data**, if

1. B is a closed ball containing w on Σ which doesn't intersect any $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves;
2. B' is a closed ball on S^2 which doesn't intersect α_0 or β_0 and which is contained in the diagram $(S^2, \alpha_0, \beta_0, w, w')$ shown in Figure 2.5 such that B' contains the point $p' \in S^2$;
3. J_s is an almost complex structure on $\Sigma \times \mathbb{D}$ which is split on B;

4. J'_s is an almost complex structure on $S^2 \times \mathbb{D}$ which is split on B';

5. r and r' are real numbers such that $0 < r, r' < 1$;

6. using the unique (up to rotation) conformal identifications of (B, w) and (B', p') as $(\mathbb{D}, 0)$, where $\mathbb{D}$ denotes the unit complex disk, ι is an embedding of $\overline{S^2 \setminus r' \cdot B'}$ into $r \cdot B \subseteq \Sigma$ such that

$$\overline{r \cdot B} \setminus \iota(S^2 \setminus r' \cdot B')$$

is a closed annulus.

7. letting $\widetilde{A}, A$ and A' denote the closures of the annuli $B \setminus \iota(S^2 \setminus B')$, $B \setminus r \cdot B$ and $B' \setminus r' \cdot B'$ respectively, then

$$\phi : \widetilde{A} \to S^1 \times [-a, 1 + b]$$

is a diffeomorphism which sends the annulus A to $[-a, 0]$ and $\iota(A')$ to $[1, 1 + b]$ and is conformal on A and A'.

Given a diagram $\mathcal{H}$ with gluing data $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ and a $T > 1$, we can form the diagram

$$\mathcal{H}_{\mathcal{J}} = (\Sigma, \boldsymbol{\alpha} \cup \{\iota\alpha_0\}, \boldsymbol{\beta} \cup \{\iota\beta_0\}, \mathbf{w} \cup \{w_0\}),$$

and form an almost complex structure $\mathcal{J}(T)$ on $\mathcal{H}_{\mathcal{J}}$. The almost complex structure $\mathcal{J}(T)$ is defined to be J_s on $\Sigma \setminus (r \cdot B)$, and $\iota_* J'_s$ on $\iota(S^2 \setminus (r' \cdot B'))$. The map ϕ can be modified to a map $\phi_T : \widetilde{A} \to S^1 \times [-a, T + b]$ which is conformal on A and A' by stretching in the region $S^1 \times [0, 1]$. This can be done in an isotopically unique way, and the dependence is explored in Lemma 2.5.7. The complex structure $\mathcal{J}(T)$ is defined to be $(\phi_T)^* J_{S^1 \times [-a, T+b]}$.

The free stabilization maps will then be maps

$$S^+_w : CF^\circ_{J_s}(\mathcal{H}) \to CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}})$$

and

$$S^-_w : CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}}) \to CF^\circ_{J_s}(\mathcal{H}),$$

for sufficiently large T. The phrase "sufficiently large" is given a precise meaning in Definition 2.5.9.

As we let T approach $+\infty$, the almost complex structure $\mathcal{J}(T)$ approaches $J_s \vee J'_s$ on $(\Sigma \vee_{w=p'} S^2) \times I \times \mathbb{R}$. If u_{T_n} is a sequence of $\mathcal{J}(T_n)$–holomorphic disks on $\mathcal{H}_{\mathcal{J}}$, representing a homology class $\phi = \phi_1 \# \phi_2$, we can extract a Gromov limit of u_{T_n} to a broken holomorphic disk U_1 on $\Sigma \times I \times \mathbb{R}$ representing ϕ_1 and a broken holomorphic disk U_2 on $\Sigma_2 \times I \times \mathbb{R}$ representing ϕ_2. By a broken holomorphic disk U_i, we mean a collection of holomorphic curves (genuine holomorphic disks, boundary degenerations, closed curves) whose total homology class represents ϕ_i. In general such a collection may be quite complicated, but in the case that both U_1 and U_2 are each genuine holomorphic disks, u_1 and u_2, there is a matching condition which must be satisfied:

$$\rho^{p_1}(u_1) = \rho^{p_2}(u_2),$$

where $\rho^{p_i}(u_i)$ is the divisor in $\mathrm{Sym}^{n_{p_i}(u_i)}(\mathbb{D})$ defined by

$$\rho^{p_i}(u_i) = (\pi_{\mathbb{D}} \circ u_i) \circ (\pi_{\Sigma_i} \circ u_i)^{-1}(p_i).$$

Conversely, if u_1 and u_2 are holomorphic disks on $\Sigma \times I \times \mathbb{R}$ and $S^2 \times I \times \mathbb{R}$ respectively which satisfy $\rho^{p_1}(u_1) = \rho^{p_2}(u_2)$, we say that u_1 and u_2 are **preglued**. In general, given a preglued pair u_1 and u_2, we can glue u_1 and u_2 to get a $\mathcal{J}(T)$–holomorphic curve $u_1 \# u_2$ for all sufficiently large T which achieve transversality (more precisely, under appropriate assumptions, there is an isomorphism of the moduli space $\mathcal{M}(\phi)$ with the fibered product $\mathcal{M}(\phi_1) \times_\rho \mathcal{M}(\phi_2)$). Technical details can be found in, for example, [Lip06, App. A], [OS08, Thm. 5.1], [OS04e], or more generally [BEH$^+$03].

2.5.2 Counting holomorphic strips after free stabilization

In this subsection we begin our analysis of the free stabilization maps. We prove that the free stabilization maps are chain maps for sufficiently large T, and that they commute with the relative homology maps under appropriate circumstances.

Consider the diagram for $(S^2, \alpha_0, \beta_0, w, w')$ pictured in Figure 2.5. Let p' be a point in the region containing the basepoint w'. Given a homology disk ϕ and a divisor $\mathbf{d} \in \mathrm{Sym}^k(\mathbb{D})$, we define

$$\mathcal{M}(\phi, \mathbf{d}) = \{u \in \mathcal{M}(\phi) : \rho^{p_2}(u) = \mathbf{d}\}.$$

Ozsváth and Szabó prove the following using a Gromov compactness argument:

Lemma 2.5.2 ([OS08, Lem. 6.4])**.** *Let $(S^2, \alpha_0, \beta_0, w, w')$ be the doubly pointed diagram for (S^3, w, w') shown in Figure 2.5. If $\mathbf{d}$ is a generic divisor in $\mathrm{Sym}^k(\mathbb{D})$ for some integer $k > 0$, then*

$$\sum_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi)=2k \\ n_{w'}(\phi)=0}} \#\mathcal{M}(\phi, \mathbf{d}) = 1 \in \mathbb{Z}_2.$$

Lemma 2.5.3. *Suppose ϕ is a disk in the diagram $(S^2, \alpha_0, \beta_0, w, w')$ shown in Figure 2.5. If $\phi \in \pi_2(\theta^+, \theta^+)$ or $\phi \in \pi_2(\theta^-, \theta^-)$ then*

$$\mu(\phi) = 2m_1 + 2m_2.$$

If $\phi \in \pi_2(\theta^+, \theta^-)$ then

$$\mu(\phi) = 2m_1 + 2m_2 + 1.$$

If $\phi \in \pi_2(\theta^-, \theta^+)$, then

$$\mu(\phi) = 2m_1 + 2m_2 - 1.$$

Proof. Suppose that $\phi \in \pi_2(\theta^+, \theta^+)$ or $\phi \in \pi_2(\theta^-, \theta^-)$. Note that the stated formula holds if ϕ is a constant disk. Any other homotopy class is obtained by splicing in an α boundary degeneration or a β boundary degeneration, and the formula respects this.

Similarly for the case of $\phi \in \pi_2(\theta^+, \theta^-)$, we note that the formula holds for the disk with $n_1 = 1$ and all other multiplicities equal to zero. Any other disk can be obtained by splicing in boundary degenerations, and the formula respects that operation. The formula for $\phi \in \pi_2(\theta^-, \theta^+)$ follows similarly. □

Before we embark on the necessary gluing arguments needed to prove naturality, we first prove an admissibility result for stabilizing strongly $\mathfrak{s}$–admissible diagrams:

Lemma 2.5.4. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for $(Y, \mathbf{w})$ which is subordinate to $w \in Y \setminus \mathbf{w}$. Let $\mathcal{H}_{\mathcal{J}}$ be the diagram described above for $(Y, \mathbf{w} \cup \{w\})$ for some choice of gluing data $\mathcal{J}$. Suppose $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Then $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible iff $\mathcal{H}_{\mathcal{J}}$ is strongly $\mathfrak{s}$–admissible.*

Proof. Suppose $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible. Suppose $\mathcal{D} = \mathcal{D}_1 \# \mathcal{D}_2$ is a periodic domain on $\mathcal{H}_{\mathcal{J}}$ with $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2N$, where $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ and $\mathcal{D}_2$ is a

$$\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = \langle c_1(\mathfrak{s}), H(\mathcal{D}_1)\rangle,$$

so $\mathcal{D}_1$ has a multiplicity greater than N, and hence so does $\mathcal{D}$. It follows that $\mathcal{H}_{\mathcal{J}}$ is strongly $\mathfrak{s}$–admissible.

Conversely if $\mathcal{H}_{\mathcal{J}}$ is strongly $\mathfrak{s}$–admissible and $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ with

$$\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2N.$$

Let A_w and $A_{w'}$ be the components of $S^2 \setminus \alpha_0$ containing w and w' respectively. Let $\mathcal{D}_2$ be the domain on $(S^2, \alpha_0, \beta_0, w, w')$ defined by

$$\mathcal{D}_2 = n_w(\mathcal{D}_1) \cdot A_{w'}.$$

Since $\mathcal{H}_{\mathcal{J}}$ is strongly $\mathfrak{s}$–admissible, we know that $\mathcal{D}_1 \# \mathcal{D}_2$ has coefficients strictly bigger than N, and hence so does $\mathcal{D}_1$, so $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible. □

We now show that the free stabilization maps are chain maps:

Lemma 2.5.5. *Fix a Heegaard diagram $\mathcal{H}$ for $(Y, \mathbf{w})$ which is subordinate to $w \in Y \setminus \mathbf{w}$, as well as gluing data $\mathcal{J}$. If ∂_{J_s} denotes the differential of $CF^\circ_{J_s}(\mathcal{H}, \mathfrak{s})$ and $\partial_{\mathcal{J}(T)}$ denotes the differential of $CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}}, \mathfrak{s})$, then for all sufficiently large T which obtain transversality, we have an identification:*

$$\partial_{\mathcal{J}(T)} = \begin{pmatrix} \partial_{J_s} & * \\ 0 & \partial_{J_s} \end{pmatrix}.$$

Proof. This is essentially [OS08, Prop. 6.5], though we will prove it here as well, since later arguments will be very analogous. Also note that in [OS08, Prop. 6.5], Ozsváth and Szabó show that if we let p' approach a point on $\beta_0 \setminus \alpha_0$, then we can take $* = U_w - U_{w'}$, though we will not have use for that at the moment.

Write

$$\partial_{\mathcal{J}(T)} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We need to compute A, D and C. Let us first show that $A = D = \partial_{J_s}$ for sufficiently large T. Suppose that $\phi = \phi_1 \# \phi_2$ is a homology class of disks. Using the combinatorial formula for the Maslov index (cf. Eq. (2.2) or [Sar11b]), we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2.$$

We first show that any disk with $\mathcal{J}(T_i)$–holomorphic representative for arbitrary large i corresponds to gluing a preglued flowline. Given a sequence of disks u_T representing $\phi = \phi_1 \# \phi_2$ for the almost complex structure $\mathcal{J}(T)$, we know that as $T \to \infty$, the holomorphic strips u_T have a Gromov limit to a broken disk U_1 on Σ and a broken disk U_2 on S^2. Since $\mu(\phi) = 1$, and $m_2 \geq 0$, using the Maslov index computation above, we conclude that $\mu(\phi_1) = 1$ and $m_2 = 0$, so U_1 consists of a single holomorphic strip u_1. There must be a component u_2 of U_2 which satisfies the matching condition $\rho^{p_1}(u_1) = \rho^{p_2}(u_2)$. Say that u_2 represents the homology class ϕ_2'. We wish to show that u_2 is the only nonconstant flowline in U_2. We know that $\phi_1 \# \phi_2'$ is nonconstant and has holomorphic representative, and hence we must have $\mu(\phi_1 \# \phi_2') \geq 1$. On the other hand, we know that

$$\mu(\phi_1 \# \phi_2) = \mu(\phi_1 \# \phi_2') + \mu(0 \# (\phi_2 - \phi_2')) = \mu(\phi_1 \# \phi_2') + \mu(\phi_2 - \phi_2'). \tag{2.5}$$

Since $\phi_2 - \phi_2'$ has holomorphic representatives as a broken disk (the remaining components of U_2), we conclude that

$$\mu(\phi_2 - \phi_2') \geq 0,$$

with equality iff $\phi_2 = \phi_2'$. On the other hand, in light of the fact that the quantity in

Equation (2.5) is equal to 1, we conclude that $\phi_2 = \phi_2'$, and hence that u_2 represents all of the nonconstant components of ϕ_2.

In particular we conclude that any such sequence of u_T representing ϕ has a Gromov limit as a preglued flowline (u_1, u_2) with $m_2(u_2) = 0$ and $\mu(u_1) = 1$.

On the other hand, given a flowline u_1 on Σ representing ϕ_1, applying Lemma 2.5.2 we see that the total number of holomorphic strips u_2 representing a class $\phi_2 \in \pi_2(\theta^+, \theta^+)$ with Maslov index $2n_{p_1}(\phi_1)$ and $m_2(\phi_2) = 0$, with (u_1, u_2) preglued is $1 \in \mathbb{Z}_2$. The same statement holds for disks in $\pi_2(\theta^+, \theta^+)$.

Since we are working with strongly $\mathfrak{s}$–admissible diagrams, by Lemma 2.3.9, there are only finitely many homotopy classes of disks with nonnegative domain and Maslov index 1, so we can pick T large enough for the previous argument to apply to all of them simultaneously. Hence we conclude that $A = D = \partial_{J_s}$.

We finally need to show that $C = 0$. Given a disk $\phi = \phi_1 \# \phi_2$, we use the combinatorial formula for the Maslov index to see that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2 + 1.$$

If ϕ has holomorphic representatives u_T for arbitrarily large T, then from a weak limit argument we conclude that $\mu(\phi_1) \geq 0$. Since $m_2 \geq 0$, we conclude that $\mu(\phi_1) = 0$. By stretching the neck sufficiently, we see that ϕ_1 has a representative as a broken holomorphic strip, and since $\mu(\phi_1) = 0$, we conclude that ϕ_1 must be the constant homology class. Hence we see that

$$n_w(\phi_1) = n_p(\phi_2) = 0.$$

Since $n_{p'}(\phi_2) = 0$ as well, and $\phi_2 \in \pi_2(\theta^+, \theta^-)$, we know that ϕ_2 is constrained to one of two choices. Each admits a holomorphic representative, but modulo 2 these cancel, so we see that $C = 0$. □

A minor modification of the previous argument yields the following:

Lemma 2.5.6. *Suppose that* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ *is a diagram with* $w, w' \in \mathbf{w}$ *but* $w'' \notin \mathbf{w}$*. Let* $(\Sigma, \boldsymbol{\alpha} \cup \alpha_0, \boldsymbol{\beta} \cup \beta_0, \mathbf{w} \cup \{w''\})$ *be the diagram which is freely stabilized at* w''*. Suppose* λ *is a path from* w *to* w'*. If* $A_{\overline{\lambda}}$ *denotes the relative homology action on the diagram which is freely stabilized at* w''*, and* A_λ *is the map on the unstabilized diagram, then we have the following decomposition for sufficiently stretched almost complex structure:*

$$A_{\overline{\lambda}} = \begin{pmatrix} A_\lambda & * \\ 0 & A_\lambda \end{pmatrix}.$$

The analogous statement holds if we replace the relative homology action with the absolute homology action A_γ *for a closed curve* γ *which doesn't intersect* w''*. In particular we have that*

$$S^+_{w''} A_\lambda = A_{\overline{\lambda}} S^+_{w''} \qquad \text{and} \qquad S^-_{w''} A_{\overline{\lambda}} = A_\lambda S^-_{w''}.$$

Similarly

$$S^+_{w''} A_\gamma = A_{\overline{\gamma}} S^+_{w''} \qquad \text{and} \qquad S^-_{w''} A_{\overline{\gamma}} = A_\gamma S^-_{w''}.$$

Proof. The proof follows as in Lemma 2.5.5, but multiplying the counts of disks by the appropriate $a(\lambda, \phi)$. □

2.5.3 Dependence on gluing data

This subsection is devoted to proving that maps S^+_w and S^-_w do not depend on the gluing data. Recall that in constructed the almost complex structure $\mathcal{J}(T)$ for given gluing data $\mathcal{J}$, we had to stretch ϕ to get a map ϕ_T. There is not a unique way to do this, though we have the following:

Lemma 2.5.7. *Suppose that* $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ *is gluing data and that* ϕ_T *is as above. Let* $\mathcal{J}(T)$ *denote the almost complex structure defined above. Let* $\mathcal{J}'(T)$ *denote the almost complex structure constructed with a* ϕ'_T *instead of* ϕ_T*. Then we have that*

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

for any $T > 1$*.*

Proof. Given a fixed ϕ, the maps ϕ_T and ϕ'_T are determined up to isotopy which is constant on $A \cup A'$. As a result, there is an isotopy ψ_t of $\Sigma \# S^2$ which is constant outside of $\widetilde{A} \setminus (A \cup A')$ such that $\mathcal{J}'(T) = (\psi_1)_* \mathcal{J}(T)$. By naturality, the map $\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)}$ is chain homotopic to the diffeomorphism map $(\psi_1)_*$, which is the identity map on the obvious identification of the complexes $CF^\circ_{\mathcal{J}(T)}(\mathcal{H}_\mathcal{J})$ and $CF^\circ_{\mathcal{J}'(T)}(\mathcal{H}_\mathcal{J})$ (as groups), i.e.

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}'(T)} = (\psi_1)_* = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

□

We now consider the effect of increasing the neck length T.

Lemma 2.5.8. *Suppose that $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ is a choice of gluing data. There is an N such that if $T, T' > N$ and $\mathcal{J}(T)$ and $\mathcal{J}(T')$ both achieve transversality, then*

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}(T')} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. Consider a sequence of T_n and T'_n both approaching $+\infty$. For each n, pick a 1–parameter family of maps

$$\widehat{\phi}_{t,n} : \widetilde{A} \to S^1 \times \mathbb{R}$$

ranging over $t \in \mathbb{R}$, which interpolate ϕ_{T_n} and $\phi_{T'_n}$, i.e. for each fixed t, the map $\widehat{\phi}_{t,n}$ is conformal on A_0 and A_1 and for fixed $t \geq 1$ the map $\widehat{\phi}_{t,n}$ is a map from $\widetilde{A}$ to $S^1 \times [-a, T'_n + b]$, and for each $t \leq 1$ the map $\widehat{\phi}_{t,n}$ is a map from $\widetilde{A}$ to $S^1 \times [-a, T_n + b]$. Further assume that each of the complex structures $(\widehat{\phi}_{t,n})^*(S^1 \times \mathbb{R})$ on $\widetilde{A}$ has conformal length at least $a + b + \min(T_n, T'_n)$.

Let $\widehat{\mathcal{J}}_n$ denote the almost complex structure on $\Sigma \times I \times \mathbb{R}$ which is J_s on $(\Sigma \setminus r \cdot B)$, and is $\iota_* J'_s$ on $\iota(S^2 \setminus r' \cdot B')$. On $\widetilde{A}$, we define $\widehat{\mathcal{J}}_n$ to be $\mathfrak{j}_t \times \mathfrak{j}_\mathbb{D}$ where $\mathfrak{j}_t = \widehat{\phi}^*_{t,n}(J_{S^1 \times \mathbb{R}})$.

As $n \to \infty$, the almost complex structures $\widehat{\mathcal{J}}_n$ approach the structure $J_s \vee J'_s$ on $(\Sigma \vee S^2) \times \mathbb{R} \times I$.

The map $\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')}$ can be computed by counting Maslov index 0 curves which are $\widehat{\mathcal{J}}_n$–holomorphic. Write

$$\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We wish to show that $A = D = 1$ and $C = 0$ for sufficiently large n. First consider A and D. If $\phi = \phi_1 \# \phi_2$ is a Maslov index 0 homology disk in $\pi_2(\mathbf{x}\otimes\theta^+, \mathbf{y}\otimes\theta^+)$ or $\pi_2(\mathbf{x}\otimes\theta^-, \mathbf{y}\otimes\theta^-)$ in the diagram $\mathcal{H}_{\mathcal{J}}$, then using the combinatorial formula for the index, we have

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1.$$

By Lemma 2.5.3, we know that

$$\mu(\phi) = \mu(\phi_1) + 2m_2.$$

If ϕ has $\widehat{\mathcal{J}}_n$-holomorphic representatives for arbitrarily large n, then by passing to a Gromov limit of a subsequence, we know that ϕ_1 and ϕ_2 are represented by broken holomorphic disks. But then we know that $\mu(\phi_1) \geq 0$ and $\mu(\phi_2) \geq 0$ by transversality. Since $\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')}$ counts holomorphic disks with Maslov index zero, we conclude that $\mu(\phi_1) = m_2 = 0$ which implies that $\mu(\phi_2) = 0$ as well. Hence ϕ_1 has Maslov index zero, and since ϕ_1 has a J_s–holomorphic representative as a broken flowline, it must be a constant disk. By explicit examination of disks on the sphere, the only ϕ_2 in $\pi_2(\theta^+, \theta^+)$ or $\pi_2(\theta^-, \theta^-)$ with Maslov index zero are the constant disks. Hence we conclude that ϕ is a constant disk. Conversely the constant disks are always counted by $\Phi_{\mathcal{J}(T_n)\to\mathcal{J}(T_n')}$ so we conclude that $A = D = 1$.

We now consider the component C. Suppose that $\phi = \phi_1 \# \phi_2$ is homology disk which has holomorphic representatives for arbitrarily large n. Again we have

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1.$$

Applying Lemma 2.5.3 to compute $\mu(\phi_2)$ shows that

$$\mu(\phi) = \mu(\phi_1) + 2m_2 + 1.$$

Since $m_2 \geq 0$, we conclude that $\mu(\phi_1) \leq -1$. If ϕ has holomorphic representatives for arbitrarily large n, then we know that ϕ_1 and ϕ_2 admit broken J_s–holomorphic representatives. In particular $\mu(\phi_1) \geq 0$, a contradiction. Hence we conclude that no disks are counted and $C = 0$. □

Definition 2.5.9. We say that N is **sufficiently large** for gluing data $\mathcal{J}$, if $\Phi_{\mathcal{J}(T)\to\mathcal{J}(T')}$ is of the form in the previous lemma for all $T, T' \geq N$. We say T is **large enough to compute** $S_w^{\pm}$ for the gluing data $\mathcal{J}$ if $T > N$ for some N which is sufficiently large.

The maps $S_w^{\pm}$ are now defined to be between the complexes $CF_{J_s}^{\circ}(\mathcal{H})$ and $CF_{\mathcal{J}(T)}^{\circ}(\mathcal{H}_{\mathcal{J}})$ for any T which is large enough to compute $S_w^{\pm}$ for $\mathcal{J}$. We now show that the maps $S_w^{\pm}$ are invariant under changing the map ϕ.

Lemma 2.5.10. *Suppose that choices of gluing data $\mathcal{J}$ and $\widetilde{\mathcal{J}}$ differ only in choice of the map ϕ. There is an N such that if $T > N$ and $\mathcal{J}(T)$ and $\widetilde{\mathcal{J}}(T)$ both achieve transversality, then we have*

$$\Phi_{\mathcal{J}(T)\to\widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. Suppose that $\phi : \widetilde{A} \to S^1 \times [-a, 1+b]$ and $\widetilde{\phi} : \widetilde{A} \to S^1 \times [-a, 1+b]$ are two choices of diffeomorphisms which are both conformal on A and A'. Suppose $T_n \to \infty$ and let ϕ_{T_n} and $\widetilde{\phi}_{T_n}$ be as above. We now construct a 1–parameter family of diffeomorphisms

$$\phi_{t,n} : \widetilde{A} \to [-a, T_n + b]$$

such that $\phi_{t,n} = \widetilde{\phi}_{T_n}$ for $t \geq 1$ and $\phi_{t,n} = \phi_{T_n}$ for $t \leq -1$ and such that $\phi_{t,n}$ is conformal on A and A' for all t. The existence of such a $\phi_{t,n}$ is seen as follows. The space $\mathrm{Diff}(S^1 \times I)$ is generated by Dehn twists around $S^1 \times \{\frac{1}{2}\}$, which are themselves isotopic to the identity map (since we are not required to fix the boundary). Hence there is an isotopy between ϕ_{T_n} and $\widetilde{\phi}_{T_n}$. Such an isotopy can clearly be taken to be only a rotation on $[-a, 0]$ and $[T, T+b]$. Since rotations are conformal maps on $S^1 \times \mathbb{R}$, we know that the isotopy between ϕ_{T_n} and $\widetilde{\phi}_{T_n}$ can be taken to be a conformal on A and A'. Hence such a $\phi_{t,n}$ exists. Here t denotes the $\mathbb{R}$ component of $\mathbb{D} = I \times \mathbb{R}$.

We now consider the almost complex structure $\widehat{\mathcal{J}}_n$ on $(\Sigma \# S^2) \times I \times \mathbb{R}$ defined to be J_s outside of $r \cdot B$ and J'_s outside of $r' \cdot B'$ and to be $\mathfrak{j}_t \times \mathfrak{j}_{\mathbb{D}}$ on $\widetilde{A} \times I \times \mathbb{R}$, where $\mathfrak{j}_t = \phi^*_{t,n}(J_{S^1 \times \mathbb{R}})$.

As we let $n \to \infty$, we get that $\widehat{\mathcal{J}}_n$ approaches the structure $J_s \vee J'_s$ on $(\Sigma \vee S^2) \times I \times \mathbb{R}$.

The map $\Phi_{\mathcal{J}(T_n) \to \widetilde{\mathcal{J}}(T_n)}$ can be computed by counting Maslov index zero $\widehat{\mathcal{J}}_n$–holomorphic disks on $\mathcal{H}_{\mathcal{J}}$. Write

$$\Phi_{\mathcal{J}(T_n) \to \widetilde{\mathcal{J}}(T_n)} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

An argument identical to the one in Lemma 2.5.8 shows that for all sufficiently large n, we have that $A = D = 1$ and $C = 0$, completing the claim. □

We now show that the maps $S^{\pm}_w$ are independent of the constants r and r' chosen in the gluing data.

Lemma 2.5.11. *Suppose that $\mathcal{J}$ and $\widetilde{\mathcal{J}}$ are two choices of gluing data which vary only in the choice of r and r'. Suppose also that ι and $\widetilde{\iota}$ are also related by restriction. Then there is an N such that if $T > N$ we have*

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. Notice that if $r < \widetilde{r}$, and $r' = \widetilde{r}'$, then the map ϕ_T can be translated in the $\mathbb{R}$ direction to make a map $\widetilde{\phi}_{T+c}$ for some $c > 0$. In particular we know that almost complex structure $\widetilde{\mathcal{J}}(T)$ is actually just equal to $\mathcal{J}(T + c)$ for some $c > 0$. Applying Lemma 2.5.8, then yields the desired form of $\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)}$. The same argument works if $\widetilde{r} < r$, and to change r'. □

We now show that the maps $S^{\pm}_w$ are independent of the almost complex structure J_s and the map ϕ used in the gluing data:

Lemma 2.5.12. *Suppose that $\mathcal{J} = (J_s, J'_s, B, B', r, r', p', \iota, \phi)$ is a choice of gluing data and $\widetilde{J}_s$ is another choice of almost complex structure on $\Sigma \times I \times \mathbb{R}$ which is split in a neighborhood containing B and such that ι is still an embedding of $\overline{S^2 \setminus r' \cdot B'} \hookrightarrow r \cdot B$. Then there is a*

$\widetilde{\phi}$ such that if $\widetilde{\mathcal{J}}$ denotes the gluing data $\widetilde{\mathcal{J}} = (J_s, \widetilde{J}'_s, B, B', r, r', p, \iota, \widetilde{\phi})$, then there is an N such that if $T > N$, we have an identification

$$\Phi_{\mathcal{J}(T)\to\widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} \Phi_{J_s\to\widetilde{J}_s} & * \\ 0 & \Phi_{J_s\to\widetilde{J}_s} \end{pmatrix}.$$

Proof. This proceeds roughly as before. We first construct an almost complex structure $\widehat{\mathcal{J}}(T)$ which can be used to count the transition map $\Phi_{\mathcal{J}(T)\to\mathcal{J}'(T)}$, and then we degenerate the almost complex structure by letting T approach $+\infty$, which will yield the total counts of disks via a Maslov index computation and a gluing argument.

First pick an interpolating conformal structure $\widehat{J}$ on $\Sigma \times I \times \mathbb{R}$ which agrees with J_s for $t \leq -1$ and with J'_s for $t \geq 1$. We can assume that $\widehat{J}$ is of the form $\mathfrak{j}_t \times \mathfrak{j}_{\mathbb{D}}$ for $t \in \mathbb{R}$ on a set containing $B \times I \times \mathbb{R}$. Pick a family $\psi_t : (B, w) \to (\mathbb{D}, 0)$ of conformal maps for the structure $\mathfrak{j}_t$. The map ψ_t is uniquely determined up to multiplication by a map $\omega : \mathbb{R} \to S^1$. A family of annuli $A_t \subseteq B$ for $t \in \mathbb{R}$ are defined by $B \setminus (r \cdot B)$ (notice that the set $r \cdot B$ depends on the conformal structure and hence varies with t). Using Lemma 2.5.11, we can assume that r is close enough to 1 so that $\overline{\iota(S^2 \setminus r' \cdot B')} \cap A_t = \varnothing$ for all t.

Let T_n be a sequence of numbers approaching $+\infty$. Let $\phi_{T_n} : \widetilde{A} \to S^1 \times [-a, T_n + b]$ be a choice of the stretched map ϕ.

Pick a 1–parameter family of smooth embeddings

$$\phi_{t,n} : \widetilde{A} \to S^1 \times \mathbb{R}$$

ranging over $t \in \mathbb{R}$, such that and if $t \leq -1$ then $\phi_{t,n} = \phi_{T_n}$ and such that $\phi_{t,n}$ maps A_t conformally to $[-a, 0]$ and maps A' conformally to $[T, T + b]$, and such that $\phi_{t,n}$ is constant in t for $t \geq 1$. Let $\widetilde{\phi}_{T_n}$ be $\phi_{t,n}$ for any $t \geq 1$. A map $\widetilde{\phi}$ can be constructed from $\widetilde{\phi}_{T_n}$ by scaling in the $\mathbb{R}$–component in the region $S^1 \times [0, T]$. We now define an almost complex structure $\widehat{\mathcal{J}}_n$ on $\Sigma \# S^2 \times I \times \mathbb{R}$ by defining it to be J_s outside on $\Sigma \setminus r \cdot B$, and and defining it to be J'_s on $S^2 \setminus r' \cdot B'$. On $\widetilde{A}$ it is $\mathfrak{j}_t \times \mathfrak{j}_{\mathbb{D}}$ where $\mathfrak{j}_t = \phi^*_{n,t}(J_{S^1\times\mathbb{R}})$.

As $n \to \infty$, the almost complex structure $\widehat{\mathcal{J}}_n$ approaches the almost complex structure $\widehat{J} \vee J'_s$ on $\Sigma \vee S^2 \times I \times \mathbb{R}$.

Write

$$\Phi_{\mathcal{J}(T_n)\to\widetilde{\mathcal{J}}(T_n)} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

First consider A and D. Suppose that $\phi = \phi_1\#\phi_2$ has Maslov index zero and has $\widehat{\mathcal{J}}_n$–holomorphic representatives for arbitrarily large n and that ϕ is in $\pi_2(\mathbf{x}\otimes\theta^+, \mathbf{y}\otimes\theta^+)$ or $\pi_2(\mathbf{x}\otimes\theta^-, \mathbf{y}\otimes\theta^-)$. Passing to a Gromov limit yields broken holomorphic strips U_1 and U_2 representing ϕ_1 and ϕ_2. By Lemma 2.5.3 we know that

$$0 = \mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 = \mu(\phi_1) + 2m_2.$$

Since ϕ_1 has holomorphic representatives we know $\mu(\phi_1) \geq 0$. By positivity of domains $m_2 \geq 0$. Hence $\mu(\phi_1) = 0$ and $m_2 = 0$. In principle the broken holomorphic strip U_1 could consist of many Maslov index zero strips, but only one strip can be $\widehat{J}$–holomorphic and the others must be J_s– or $\widetilde{J}_s$–holomorphic. Since generically there are no nonconstant J_s– or $\widetilde{J}_s$–holomorphic strips with Maslov index zero, we conclude that U_1 contains only a single nonconstant component u_1. There must be a matching component in U_2, namely a component u_2 with

$$\rho^{p_1}(u_1) = \rho^{p_2}(u_2).$$

We now claim that there are no other nonconstant components. Let ϕ_2' represent the homology class of u_2. Note that $m_1(u_1) = m_1(u_2)$ and so by positivity we know that $m_1(\phi_2 - \phi_2') = 0$. Since $\phi_2 - \phi_2'$ has holomorphic representatives as broken a broken holomorphic curve on $(S^2, \alpha_0, \beta_0, w, w')$, we conclude that $\mu(\phi_2 - \phi_2') \geq 0$ with equality iff $\mathcal{D}(\phi_2') = 0$. But then by an easy Maslov computation (the same one done in Lemma 2.5.5) we see that

$$0 = \mu(\phi) = \mu(\phi_1\#\phi_2') + \mu(\phi_2 - \phi_2').$$

Since $\phi_1\#\phi_2'$ has the holomorphic representative $u_1\#u_2$, we conclude that $\mu(\phi_1\#\phi_2') \geq 0$, and hence $\mu(\phi_2 - \phi_2') = 0$. Hence $\mathcal{D}(\phi_2 - \phi_2') = 0$ and the remaining components of U_2 are constant.

Conversely, given a $\widehat{J}$–holomorphic disk of Maslov index zero counted by $\Phi_{J_s\to\widetilde{J}_s}$, we can glue, as in Lemma 2.5.5 using Lemma 2.5.2. Hence we conclude that $A = D = \Phi_{J_s\to\widetilde{J}_s}$.

We finally consider component C. This is zero by an identical argument as the one in Lemma 2.5.8. $\square$

We now consider changing the almost complex structure J_s' on S^2:

Lemma 2.5.13. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is a choice of gluing data and $\widetilde{J}_s'$ is another choice of almost complex structure on $S^2 \times I \times \mathbb{R}$ which is split in a neighborhood of B such that ι is still an embedding of $\overline{S^2 \setminus r' \cdot B'} \hookrightarrow r \cdot b$. Then there is a $\widetilde{\phi}$ such that if $\widetilde{\mathcal{J}}$ denotes the gluing data*

$$\widetilde{\mathcal{J}} = (J_s, J_s', B, B', r, r', p', \iota, \phi),$$

then there is an N such that if $T > N$, we have an identification o

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. The argument is nearly identical to the proof of Lemma 2.5.12, so we omit it. $\square$

We now show that the maps $S_w^{\pm}$ are invariant under the choice of embedding $\iota : \overline{S^2 \setminus r' \cdot B'} \hookrightarrow r \cdot B$.

Lemma 2.5.14. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is gluing data and suppose that $\widetilde{\iota}$ is another choice of embedding of $(\overline{S^2 \setminus r' \cdot B'}, w) \hookrightarrow (r \cdot B, w)$. Then there is a map $\widetilde{\phi}$ such that*

$$\widetilde{\mathcal{J}} = (J_s, J_s', B, B', r, r', p', \widetilde{\iota}, \widetilde{\phi})$$

is gluing data and

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

for any $T > 1$.

Proof. The maps ι and $\widetilde{\iota}$ are isotopic relative $\mathbf{w} \cup \{w\}$. Let ψ_t be an isotopy of Σ fixing $\mathbf{w} \cup \{w\}$ such that $\psi_1 \circ \iota = \widetilde{\iota}$ and such that ψ_t is fixed outside of $r \cdot B$. We simply define

$\widetilde{\phi} = \phi \circ \psi_1^{-1}$. We tautologically have $\widetilde{\mathcal{J}}(T) = (\psi_1)_* \mathcal{J}(T)$. By naturality of the Heegaard Floer complexes, we thus know that

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq (\psi_1)_* = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

□

We now consider the dependence on the choice of ball $B' \subseteq S^2$ and connected sum point $p' \in B'$:

Lemma 2.5.15. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is gluing data and that $\widetilde{B}'$ and $\widetilde{p}'$ are different choices of balls and connected sum point p'. Then there are choices $\widetilde{J}_s'$ of almost complex structure on S^2, embedding $\widetilde{\iota}$ and map $\widetilde{\phi}$ such that $\widetilde{\mathcal{J}} = (J_s, \widetilde{J}_s', B, \widetilde{B}', r, r', \widetilde{p}', \widetilde{\iota}, \widetilde{\phi})$ is gluing data and*

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

for all $T > 1$.

Proof. Let ψ_t be an isotopy of S^2 fixing $w \in S^2$, such that ψ_1 maps B' to $\widetilde{B}'$ and p' to $\widetilde{p}'$. We now set $\widetilde{J}_s' = (\psi_1)_* J_s$ and $\widetilde{\iota} = \iota \circ \psi_1^{-1}$. We can actually just set $\widetilde{\phi}$ equal to ϕ as functions. Tautologically have that $\mathcal{H}_{\mathcal{J}} = \mathcal{H}_{\widetilde{\mathcal{J}}}$ and $\mathcal{J}(T) = \widetilde{\mathcal{J}}(T)$, and hence $\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)}$ is tautologically the identity map. □

We now explore dependence on the final piece of a data, the ball $B \subseteq \Sigma$.

Lemma 2.5.16. *Suppose that $\mathcal{J} = (J_s, J_s', B, B', r, r', p', \iota, \phi)$ is gluing data and suppose that $\widetilde{B}$ is another ball in Σ containing w which doesn't intersect any of the $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. Then there is gluing data*

$$\widetilde{\mathcal{J}} = (\widetilde{J}_s, J_s', \widetilde{B}, B', r, r', p', \widetilde{\iota}, \widetilde{\phi})$$

with

$$\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Proof. We proceed similarly to the last two lemmas. Let ψ_t be an isotopy of Σ which fixes $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \{w\})$ and such that ψ_1 maps B to $\widetilde{B}$. We simply define $\widetilde{J}_s = (\psi_1)_* J_s$, $\widetilde{\iota} = \psi_1 \circ \iota$, and $\widetilde{\phi} = \phi \circ \psi_1^{-1}$. We have that $\widetilde{\mathcal{J}}(T) = (\psi_1)_* \mathcal{J}(T)$. By naturality we have that $\Phi_{\mathcal{J}(T) \to \widetilde{\mathcal{J}}(T)} \simeq (\psi_1)_*$, which clearly has the desired form. □

2.5.4 Free stabilization of the triangle maps

In this section we prove a theorem about free stabilization and the triangle maps. The main consequence is that the free stabilization maps commute with maps which can be computed as a sequence of triangle maps, for instance the handleslide maps or isotopy maps. The result has some similarities with the quasi-stabilization argument in [MO10], though our argument is somewhat different.

Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a strongly $\mathfrak{s}$–admissible Heegaard triple for $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma}})$ and suppose that $w \in \Sigma \setminus (\mathbf{w} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma})$ is a new basepoint. Fix curves α_0, β_0 and γ_0 which are pairwise isotopic and each bound a small disk containing w and which don't intersect any other $\boldsymbol{\alpha}, \boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ curves. Further assume that $|\sigma_0 \cap \tau_0| = 2$ for distinct $\sigma, \tau \in \{\alpha, \beta, \gamma\}$. Define

$$\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\alpha_0\}$$

and define $\overline{\boldsymbol{\beta}}$ and $\overline{\boldsymbol{\gamma}}$ similarly. Define

$$\overline{\mathcal{T}} = (\Sigma, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}}, \mathbf{w} \cup \{w\}).$$

In light of Lemma 2.3.11, a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma}})$ uniquely determines a Spin^c structure $\overline{\mathfrak{s}} \in \mathrm{Spin}^c(X_{\overline{\boldsymbol{\alpha}}\overline{\boldsymbol{\beta}}\overline{\boldsymbol{\gamma}}})$, An easy argument shows that $\mathcal{T}$ is strongly $\mathfrak{s}$–admissible iff $\overline{\mathcal{T}}$ is strongly $\overline{\mathfrak{s}}$ admissible. The diagram $\overline{\mathcal{T}}$ is shown in Figure 2.6.

The main result of this section is the following:

Theorem 2.5.17. *If $\mathcal{J}$ is gluing data, as in the previous section, then for all sufficiently*

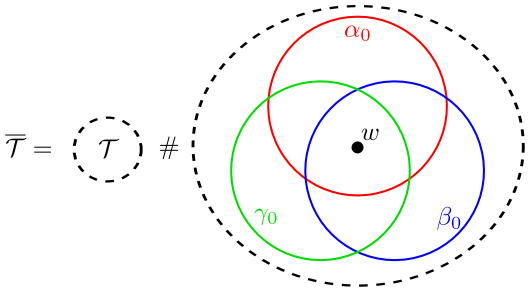

Figure 2.6: The Heegaard triple $\overline{\mathcal{T}}$.

large T such that $\mathcal{J}(T)$ achieves transversality, we have that

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix},$$

where the matrix on the right is an expansion of the map on the left in the variable $\cdot$, which can be either x^+ or x^-.

Similarly, for sufficiently stretched almost complex structure, we have

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}((\mathbf{x}\times x^+)\otimes(\mathbf{y}\times\cdot)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix}.$$

As with the previous neck stretching arguments, we need a lemma about counts of disks in $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$ satisfying a certain matching condition. This is similar to [MO10, Lem. 5.16], which is itself very similar to [OS08, Lem. 6.4]. The first step is a Maslov index computation. Consider the Heegaard triple $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$ shown in Figure 2.7. We first perform a Maslov index computation for triangles on $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$:

Lemma 2.5.18. *Suppose $\psi\in\pi_2(a,b,c)$ is a triangle on $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$. If $\psi\in\pi_2(x^+,y^+,z^+),\pi_2(x^-,y^+,z^-)$, or $\pi_2(x^+,y^-,z^-)$. Then*

$$\mu(\psi) = 2m_1 + 2m_2.$$

If $\psi\in\pi_2(x^+,y^+,z^-)$ then

$$\mu(\psi) = 2m_1 + 2m_2 + 1.$$

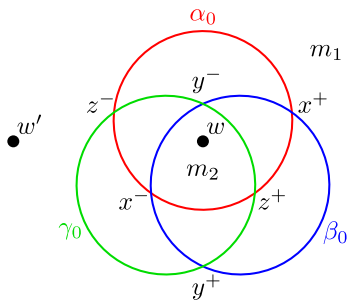

Figure 2.7: The Heegaard triple $(S^2, \alpha_0, \beta_0, \gamma_0, w, w')$ with $x^{\pm}, y^{\pm}, z^{\pm}$ and multiplicities m_1 and m_2 labelled.

If $\psi \in \pi_2(x^-, y^+, z^+)$ *or* $\pi_2(x^+, y^-, z^+)$*, then*

$$\mu(\psi) = 2m_1 + 2m_2 - 1.$$

Proof. In all cases, the formula can be easily verified for a particular "simple" homology class of triangles. The formula in fully generality then follow since both respect splicing in homology classes of boundary degenerations, and any two homology triangles with the same vertices differ by a sum of domains corresponding to boundary degenerations. □

Let $p \in S^2$ be a point in the same region as w'. Let $x \in \alpha_0 \cap \beta_0$, $y \in \beta_0 \cap \gamma_0$ and $z \in \alpha_0 \cap \gamma_0$. Given a divisor $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$, and a homology triangle $\psi \in \pi_2(x, y, z)$ we define

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

and

$$M_{(x,y,z)}(\mathbf{d}) = \sum_{\substack{\psi \in \pi_2(x,y,z) \\ n_w(\psi)=0 \\ \mu(\psi)=2k}} \#\mathcal{M}(\psi, \mathbf{d}).$$

Lemma 2.5.19. *If* (x, y, z) *is one of the tuples* $(x^+, y^+, z^+), (x^-, y^+, z^-)$ *or* (x^+, y^-, z^-)*, then the number* $M_{(x,y,z)}(\mathbf{d})$ *is independent of* $\mathbf{d}$ *for generic* $\mathbf{d}$ *and in fact*

$$M_{(x,y,z)}(\mathbf{d}) \equiv 1 (\mathrm{mod}\, 2).$$

Proof. The proof is nearly identical to the proof of [OS08, Lem. 6.4]. We first show that $M_{(x,y,z)}(\mathbf{d})$ is independent of $\mathbf{d}$ using a Gromov compactness argument. Suppose that $\mathbf{d}_t$ is a path in $\mathrm{Sym}^k(\Delta)$ between $\mathbf{d}_0$ and $\mathbf{d}_1$. We then consider the 1–manifold

$$\mathcal{M}_{(x,y,z)}(\mathbf{d}_t) = \bigcup_{\substack{t\in I\\ \psi\in\pi_2(x,y,z)\\ n_w(\psi)=0\\ \mu(\psi)=2k}} \mathcal{M}(\psi,\mathbf{d}_t).$$

We count the ends of $\mathcal{M}_{(x,y,z)}(\mathbf{d}_t)$. There are several types of ends. The first type corresponds to setting $t=0$ or $t=1$, namely

$$\bigcup_{\substack{\psi\in\pi_2(x,y,z)\\ n_w(\psi)=0\\ \mu(\psi)=2k}} \mathcal{M}(\psi,\mathbf{d}_i).$$

The other types of ends correspond to homology classes of triangles ψ degenerating into broken holomorphic triangles. Let x', y' and z' denote the "other" generators of $\alpha_0\cap\beta_0$, $\beta_0\cap\gamma_0$ and $\alpha_0\cap\gamma_0$ respectively, i.e. the generators in the opposite grading as x, y and z. In principle a sequence of holomorphic strips of index $2k$ could break into a collection of many holomorphic strips, boundary degenerations, and closed surfaces, whose total homology represented the original strip. In the ends of $\mathcal{M}_{(x,y,z)}(\mathbf{d}_t)$, one of the components must satisfy the matching condition with a divisor $\mathbf{d}_t$, and hence using the Maslov index computation in Lemma 2.5.18, the possible degenerations are limited to the following types:

$$\bigcup_{\substack{t\in I\\ \phi\in\pi_2(x,x')\\ \psi\in\pi_2(x',y,z)\\ n_w(\phi)=n_w(\psi)=n_{w'}(\phi)=0\\ \mu(\phi)=1,\mu(\psi)=2k-1}} \mathcal{M}(\psi,\mathbf{d}_t)\times\widehat{\mathcal{M}}(\phi), \tag{a}$$

$$\bigcup_{\substack{t\in I\\ \phi\in\pi_2(y,y')\\ \psi\in\pi_2(x,y',z)\\ n_w(\phi)=n_w(\psi)=n_{w'}(\phi)=0\\ \mu(\phi)=1,\mu(\psi)=2k-1\\ n_{w'}(\phi)=0}} \mathcal{M}(\psi,\mathbf{d}_t)\times\widehat{\mathcal{M}}(\phi), \tag{b}$$

and

$$\bigcup_{\substack{t\in I\\ \phi\in\pi_2(z',z)\\ \psi\in\pi_2(x,y,z')\\ n_w(\phi)=n_w(\psi)=n_{w'}(\phi)=0\\ \mu(\phi)=1,\mu(\psi)=2k-1}} \mathcal{M}(\psi,\mathbf{d}_t)\times\widehat{\mathcal{M}}(\phi). \tag{c}$$

The total number of elements in any of the ends of type (a),(b) or (c) is even, since the differential of the hat complexes $\widehat{CF}(S^2,\sigma,\tau,w,w')$ vanishes for any two distinct $\sigma,\tau\in\{\alpha_0,\beta_0,\gamma_0\}$.

Hence the total number of ends, modulo 2, is just

$$\#\mathcal{M}(\mathbf{d}_1)-\#\mathcal{M}(\mathbf{d}_0),$$

and hence that quantity is zero.

We now wish to show that $M(\mathbf{d})\equiv 1$, modulo 2. Let $\mathbf{d}_T$ consist of k points spaced at least T apart in Δ such that as $T\to\infty$ the points approach e_β. As $T\to\infty$, the set of curves splits into k Maslov index 2 β–boundary degenerations. There is a unique Maslov index 2 homotopy class of β boundary degeneration ϕ with $n_w(\phi)=0$, namely the domain which has multiplicity one in the complement of the region in $S^2\setminus\beta_0$ which doesn't contain w, and zero elsewhere. By [OS08, Thm. 5.5], we know that a boundary degeneration of Maslov index 2 has a unique holomorphic representative modulo the action of $PSL_2(\mathbb{R})$. Hence we conclude that $\lim_{T\to\infty}\#\mathcal{M}(\psi,\mathbf{d}_T)=1(\mathrm{mod}\,2)$ by gluing, from which the claim follows.

□

We are now in position to prove the main result of this subsection:

Proof of Theorem 2.5.17. We will focus on computing the map $F_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+))$ in terms of $F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y})$, computing the second map in the theorem statement follows an identical argument.

Write

$$F_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^+))=\begin{pmatrix}A(\mathbf{x},\mathbf{y}) & B(\mathbf{x},\mathbf{y})\\ C(\mathbf{x},\mathbf{y}) & D(\mathbf{x},\mathbf{y})\end{pmatrix}.$$

We will start by computing A and D. Suppose that $\psi = \psi_1 \# \psi_2$ is a homology triangle in $\pi_2(\mathbf{x} \times x^+, \mathbf{y} \times y^+, \mathbf{z} \times z^+)$ or $\pi_2(\mathbf{x} \times x^-, \mathbf{y} \times y^+, \mathbf{z} \times z^-)$ of Maslov index zero. Using the combinatorial formula for the Maslov index, we have that

$$0 = \mu(\psi_1 \# \psi_2) = \mu(\psi_1) + \mu(\psi_2) - 2m_1.$$

Using Lemma 2.5.18, we have

$$\mu(\psi) = \mu(\psi_1) + 2m_2.$$

If ψ has $\mathcal{J}(T)$–holomorphic representatives u_T for arbitrarily large T, then a Gromov compactness argument shows that ψ_1 and ψ_2 have representatives as broken holomorphic triangles, U_1 and U_2 respectively. In particular $\mu(\psi_1) \geq 0$. In light of the equality $0 = \mu(\psi_1) + 2m_2$, from positivity of domains we conclude that $\mu(\psi_1) = 2m_2 = 0$. The collection U_1 is a broken holomorphic triangle, which means that it is a collection of a single holomorphic triangle, a collection of holomorphic strips, boundary degenerations, and closed surfaces. Since the total Maslov index is zero, we conclude that the only nonconstant component of U_1 is a single holomorphic triangle, u_1. There must be a component of U_2, say u_2, with the matching condition

$$\rho^{p_1}(u_1) = \rho^{p_2}(u_2).$$

Say that u_2 is represented by the homology triangle ψ_2'. Then $\psi_2 - \psi_2'$ admits a representative as a broken holomorphic curve, and hence $\mu(\psi_2') \geq 0$, with equality iff $\mathcal{D}(\psi_2') = 0$, since ψ_2' is represented by holomorphic strips, boundary degenerations, and closed surfaces.

Since $\psi_1 \# \psi_2'$ has a holomorphic representative, $u_1 \# u_2$, we conclude that $\mu(\psi_1 \# \psi_2') \geq 0$. Because of the matching condition,we have $m_1(\psi_2') = m_1(\psi_1) = m_1(\psi_2)$ and hence

$$m_1(\psi_2 - \psi_2') = 0.$$

Hence we have that

$$0 = \mu(\psi_1 \# \psi_2) = \mu(\psi_1 \# \psi_2') + \mu(\psi_2 - \psi_2').$$

Combining this with our previous observations, we conclude that $\mu(\psi_1 \# \psi_2') = 0$ and $\mu(\psi_2 - \psi_2') = 0$ and hence $\psi_2' = \psi_2$. Hence the only nonconstant component of U_2 is u_2. We thus conclude that the pair (u_1, u_2) were a preglued pair.

Conversely, given a preglued pair, we can glue to form a genuine holomorphic triangle for sufficiently stretched almost complex structure. By Lemma 2.5.19, given a holomorphic triangle u_1 representing a $\psi_1 \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$, among holomorphic triangles representing classes in $\pi_2(x^+, y^+, z^+)$, there is one (modulo 2), holomorphic triangle satisfying the matching condition with u_1. Hence for sufficiently stretched almost complex structures, we have an identification of

$$A(\mathbf{x}, \mathbf{y}) = F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}).$$

The same reasoning shows that we have an identification

$$D(\mathbf{x}, \mathbf{y}) = F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}).$$

We finally wish to show that $B(\mathbf{x}, \mathbf{y}) = 0$. Suppose that $\psi = \psi_1 \# \psi_2 \in \pi_2(\mathbf{x} \times x^+, \mathbf{y} \times y^+, \mathbf{z} \times z^-)$ is a triangle with Maslov index zero. By 2.5.18 we have that

$$0 = \mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1) + 2m_2 + 1.$$

If ψ has $\mathcal{J}(T)$–holomorphic representatives for arbitrarily large T, then we conclude that ψ_1 must have a representative as a broken holomorphic triangle, and hence $\mu(\psi_1) \geq 0$. Since $m_2 \geq 0$ by positivity, we conclude that

$$\mu(\psi_1) + 2m_2 + 1 \geq 1,$$

and hence no triangles are counted for sufficiently stretched almost complex structure, so $B(\mathbf{x}, \mathbf{y}) = 0$. □

2.5.5 Topological preliminaries about based embeddings of Heegaard surfaces

In this section we prove a lemma about isotopies of Heegaard surfaces which will be useful for proving naturality of the free stabilization maps, as well as the 1– and 3–handle maps. We begin with a two parameter isotopy extension theorem. We leave the proof to the reader, since it is elementary:

Lemma 2.5.20. *Let M be a closed manifold and let $\Phi_t : M \to M$ be an isotopy. Let N be an embedded submanifold and let $\phi_t = \Phi_t|_N$. Suppose that ψ_t is another isotopy of N such that $\psi_0 = \phi_0$ and $\psi_1 = \phi_1 = \mathrm{id}$ and suppose there is a homotopy $(h_t)_s$ through isotopies between ϕ_t and ψ_t such that $(h_0)_s$ and $(h_1)_s$ are constant in s. Then there is another isotopy $\Psi_t : M \to M$ such that $\Psi_0 = \Phi_0 = \mathrm{id}_M$ and $\Psi_1 = \Phi_1$ and $\Psi_t|_N = \psi_t$. Furthermore the isotopy Ψ_t can be taken to be equal to Φ_t outside of a given open neighborhood of $h(N \times I \times I) \subseteq M$.*

Lemma 2.5.21. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ and $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha'}, \boldsymbol{\beta'}, \mathbf{w})$ are Heegaard diagrams for $(Y, \mathbf{w})$ and that $\mathbf{p} \subseteq Y$ is a finite collection of points such that $\mathbf{p} \subseteq \Sigma \cap \Sigma'$. Suppose further that $\mathcal{H}$ and $\mathcal{H}'$ are related by an isotopy of Y fixing $\mathbf{w}$. Then $\mathcal{H}$ and $\mathcal{H}'$ are related by an isotopy F_t fixing $\mathbf{w}$ such that $\mathbf{p} \subseteq F_t(\Sigma)$ for all t.*

Proof. Let $G_t : Y \to Y$ be an isotopy fixing $\mathbf{w}$ which maps $\mathcal{H}$ to $\mathcal{H}'$. Note that $\mathbf{p} \subseteq G_t(\Sigma)$ for all t iff $G_t^{-1}(\mathbf{p}) \subseteq \Sigma$ for all t. Let $\Phi_t = G_t^{-1}$ and let $\phi_t = \Phi_t|_{\mathbf{p}}$. Note that ϕ_t is a collection of paths starting at $\mathbf{p}$ and ending at a collection of points $\mathbf{q} \subseteq \Sigma$.

Since $\pi_1(\Sigma)$ surjects onto $\pi_1(Y)$, we can homotope ϕ_t with a homotopy $(h_t)_s$ so that $\mathrm{im}\,\phi_t \subseteq \Sigma$. We can assume that the image of $(h_t)_s$ avoids the basepoints $\mathbf{w}$. Applying Lemma 2.5.20, there is an isotopy Ψ_t of Y such that $\Psi_t(\mathbf{p}) \subseteq \Sigma$ for all t, and such that $\Psi_0 = \Phi_0$ and $\Psi_1 = \Phi_1$. Furthermore, since the image of $(h_t)_s$ avoids $\mathbf{w}$, we can assume that Ψ_t fixes $\mathbf{w}$, since Φ_t also fixes $\mathbf{w}$. Set $F_t = \Psi_t^{-1}$. Since $\Psi_1 = \Phi_1$, we know that F_t is an isotopy between $\mathcal{H}$ and $\mathcal{H}'$. By construction, $F_t^{-1}(\mathbf{p}) \subseteq \Sigma$. □

2.5.6 Invariance of free stabilization maps

In this subsection we combine the work of the previous subsections to show that the free stabilization maps yield well defined maps on Heegaard Floer homology.

Theorem 2.5.22. *The free stabilization maps $S_w^{\pm}$ described above yield well defined maps*

$$S_w^+ : HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S_w^- : HF^\circ(Y, \mathbf{w} \cup \{w\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

on homology.

Proof. Suppose that $\mathcal{H}_1 = (\Sigma_1, \boldsymbol{\alpha}_1, \boldsymbol{\beta}_1, \mathbf{w})$ and $\mathcal{H}_2 = (\Sigma_2, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_2, \mathbf{w})$ are two Heegaard diagrams for $(Y, \mathbf{w})$ which are subordinate to $w \in Y \setminus \mathbf{w}$, and that $\mathcal{J}_i = (J_{s,i}, J'_{s,i}, B_i, B'_i, r_i, r'_i, p', \iota_i, \phi_i)$ are choices of almost complex structures for $\mathcal{H}_i$. We wish to show that that the following diagram commutes up to chain homotopy:

$$\begin{array}{ccc} CF^\circ_{J_{s,1}}(\mathcal{H}_1) & \xrightarrow{S_w^+} & CF^\circ_{\mathcal{J}_1(T_1)}((\mathcal{H}_1)_{\mathcal{J}_1}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{(\mathcal{H}_1)_{\mathcal{J}_1} \to (\mathcal{H}_2)_{\mathcal{J}_2}}} \\ CF^\circ_{J_{s,2}}(\mathcal{H}_2) & \xrightarrow{S_w^+} & CF^\circ_{\mathcal{J}_2(T_2)}((\mathcal{H}_2)_{\mathcal{J}_2}). \end{array}$$

for any T_i which are large enough to compute $S_w^\pm$ for the gluing data $\mathcal{J}_i$. According to Lemma 2.3.20, the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ can be written as a composition of moves corresponding to four moves: change of almost complex structure, triangle maps (isotopies and handleslides), (1,2)–stabilization and destabilization, and isotopies of Σ inside of Y fixing $\mathbf{w}$.

We first address changing the gluing data $\mathcal{J}$ for a fixed Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$. Combining Lemmas 2.5.7, 2.5.8, 2.5.10, 2.5.11, 2.5.12, 2.5.13, 2.5.14, 2.5.15, and 2.5.16, we see that the change of almost complex structure map

$$\Phi_{\mathcal{J}_1(T_1) \to \mathcal{J}_2(T_2)} \simeq \begin{pmatrix} \Phi_{J_{s,1} \to J_{s,2}} & * \\ 0 & \Phi_{J_{s,1} \to J_{s,2}} \end{pmatrix},$$

for T_i which are large enough to compute $S_w^\pm$ for gluing data $\mathcal{J}_i$, which immediately implies that S_w^+ and S_w^- commutes with the change of almost complex structure map. We leave it to the reader to combine these lemmas to produce the above matrix since it is quite easy, though we do make a comment about the definition of T being "large enough to compute $S_w^\pm$". In going between between the two almost complex structures $\mathcal{J}_1(T_1)$ and $\mathcal{J}_2(T_2)$, we need to go through intermediate almost complex structures $\mathcal{J}_i(T)$ for T which is large enough to compute $S_w^\pm$ for both $\mathcal{J}_1$ and $\mathcal{J}_2$ (and potentially even larger, so that we know the forms

of the intermediate transition maps). To do this we need Lemma 2.5.8 to show that as we increase T, eventually the change of almost complex structure maps become predictable.

We now consider maps which can be computed with a triangle map (handleslide maps and isotopy maps). By Theorem 2.5.17, the free stabilization maps commute with these triangles maps. Upon direct examination, the (1,2)–stabilization maps commute with the free stabilization maps.

We finally consider maps induced by isotopies of Y relative the points $\mathbf{w}$. Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram with $w \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ and suppose $\phi_t : Y \to Y$ is an isotopy which fixes $\mathbf{w}$ and such that $w \in \phi_1(\Sigma)$. Letting $\phi = \phi_1$, we need to show that the following diagram commutes up to chain homotopy for some choices of gluing data $\mathcal{J}_1$ on $\mathcal{H}$ and $\mathcal{J}_2$ on $\phi_*\mathcal{H}$:

$$\begin{array}{ccc} CF_{J_{s,1}}(\mathcal{H}) & \xrightarrow{S^+_w} & CF^\circ_{\mathcal{J}_1(T_1)}(\mathcal{H}_{\mathcal{J}_1}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}\to\phi_*\mathcal{H}}} & & \downarrow{\scriptstyle \Phi_{\mathcal{H}_{\mathcal{J}_1}\to(\phi_*\mathcal{H})_{\mathcal{J}_2}}} \\ CF_{J_{s,2}}(\phi_*\mathcal{H}) & \xrightarrow{S^+_w} & CF^\circ_{\mathcal{J}_2(T_2)}((\phi_*\mathcal{H})_{\mathcal{J}_2}) \end{array} \tag{2.6}$$

To this end, by Lemma 2.5.21, we can assume that $w \in \phi_t(\Sigma)$ for all t. Let $w_t = \phi_t^{-1}(w)$. The map w_t can be viewed as a smooth isotopy of the point w in Σ. Extend the isotopy w_t to an isotopy of $\Sigma \times I \to \Sigma$ which fixes $\mathbf{w}$, and then extend this to an isotopy $W_t : Y \times I \to Y$ which fixes $\mathbf{w}$ and maps Σ to Σ. Consider the two isotopies

$$\rho_t = W_t$$

and

$$\psi_t = \phi_t \circ W_t^{-1}.$$

Note that

$$\psi_1 \circ \rho_1 = \phi_1.$$

Also note that ψ_t fixes $\mathbf{w} \cup \{w\}$, and ρ_t fixes $\mathbf{w}$ and maps Σ diffeomorphically to Σ. Hence to show that the diagram in (2.6) commutes for an arbitrary isotopy ϕ_t fixing $\mathbf{w}$, it is sufficient to show that the diagram in (2.6) commutes if ϕ_t either

1. fixes $\mathbf{w} \cup \{w\}$;

2. fixes $\mathbf{w}$ and maps Σ into Σ.

Note that if ϕ_t fixes $\mathbf{w} \cup \{w\}$, then the commutativity of the diagram in (2.6) is tautological because we can take $\mathcal{J}_2 = \phi_* \mathcal{J}_1$ and $\Phi_{\mathcal{H}_{\mathcal{J}_1} \to (\phi_* \mathcal{H})_{\mathcal{J}_2}} = (\phi_1)_*$.

Now suppose that ϕ_t is an isotopy which fixes $\mathbf{w}$ but which maps Σ to Σ. The map $\Phi_{\mathcal{H} \to (\phi_* \mathcal{H})}$ can be computed simply as the tautological map $(\phi_1)_*$, but by naturality (cf. [JT12, Thm. 2.39]), it can also be computed as a composition of isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves, and change of almost complex structure maps. For our present purpose, we can thus assume that the underlying Heegaard surfaces for $\mathcal{H}$ and $\phi_* \mathcal{H}$ are the same, which allows us to pick the gluing data $\mathcal{J}_1$ and $\mathcal{J}_2$ to be equal. The maps corresponding to isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves can be computed as compositions of triangle maps. To such a move on $\mathcal{H}$, there is a corresponding move on $\mathcal{H}_{\mathcal{J}}$, which can also be computed as a triangle map. By Theorem 2.5.17 we know the free stabilization maps commute with such moves with the triangle maps. Hence we know that the diagram in (2.6) commutes up to chain homotopy for some choice of gluing data, completing the proof. □

2.5.7 Commutation of free-stabilization maps

In this subsection, we prove that free-stabilization and de-stabilizaiton maps always commute. Upon inspection of the formulas, this appears obvious, since we can realize both free-stabilizations simultaneously on the same diagram. There is a subtlety that the almost complex structures used to compute one free-stabilization are not *a-priori* the same almost complex structures used to compute another. Nonetheless, we have the following:

Proposition 2.5.23. *Suppose* $(Y, \mathbf{w})$ *is a multi-based 3–manifold, and* $w_1, w_2 \notin \mathbf{w}$ *and* $\circ_1, \circ_2 \in \{+, -\}$. *Then we have equality*

$$S^{\circ_1}_{w_1} S^{\circ_2}_{w_2} = S^{\circ_2}_{w_2} S^{\circ}_{w_1}$$

as maps on $HF^\circ(Y, \mathbf{w}, \mathfrak{s}), HF^\circ(Y, \mathbf{w} \cup \{w_i\}, \mathfrak{s})$ *and* $HF^\circ(Y, \mathbf{w} \cup \{w_1, w_2\}, \mathfrak{s})$ *(the invariants*

from naturality).

Proof. This follows from a change of almost complex structure computation analogous to the one in Lemma 2.5.8. Let us consider first the claim involving the maps $S^+_{w_1}$ and $S^+_{w_2}$. We take gluing data $\mathcal{J}$ for doing both free-stabilizations and stretching on two necks, c_1 and c_2, for w_1 and w_2, respectively. Note that the map $S^+_{w_1}S^+_{w_2}$ is computed by stretching on c_2 first, and then stretching (possibly a lot more) on c_1. The map $S^+_{w_2}S^+_{w_1}$ is defined by stretching on c_1 first, and then on c_2. The key point is that the final almost complex structure for $S^+_{w_2}S^+_{w_1}$ may be stretched on c_2 way more than on c_1, while the resulting almost complex structure for $S^+_{w_1}S^+_{w_2}$ may be stretched way more on c_1 than c_2. Thus we need to compute the change of almost complex structure map $\Phi_{\mathcal{J}(T_1,T_2)\to\mathcal{J}(T_1',T_2')}$. We make the following claim: if T_1, T_2, T_1', and T_2' are all sufficiently large, then

$$\Phi_{\mathcal{J}(T_1,T_2)\to\mathcal{J}(T_1',T_2')}(\mathbf{x}\times\theta_1^+\times\theta_2^+) = \mathbf{x}\times\theta_1^+\times\theta_2^+.$$

To prove this, we take a homology disk

$$\phi = \phi_\Sigma \# \phi_1 \# \phi_2 \in \pi_2(\mathbf{x}\times\theta_1^+\times\theta_2^+, \mathbf{y}\times y_1\times y_2),$$

and note that by modifying the Maslov index computation of Lemma 2.5.3, we see that

$$\mu(\phi) = \mu(\phi_\Sigma) + \Delta(\theta_1^+, y_1) + \Delta(\theta_2^+, y_2) + 2n_{w_1}(\phi_1) + 2n_{w_2}(\phi_2),$$

where $\Delta(\theta_i^+, y_i)$ denotes the drop in relative grading from θ_i^+ to y_i. For sufficiently large T_i and T_i' we can extract a weak limit, of ϕ_Σ, ϕ_1 and ϕ_2. Since all the terms in the above equation are nonnegative, we conclude that for such degenerated almost complex structures, the change of almost complex structure map, which counts Maslov index zero disks, counts only disks where all of the above terms are zero. Of course the limiting almost complex structure on $\Sigma\times[0,1]\times\mathbb{R}$ (on the unstabilized diagram) is cylindrical, so we see that $\mu(\phi_\Sigma)=0$, and admits a weak representative and hence is a constant disk. Hence $\Phi_{\mathcal{J}(T_1,T_2)\to\mathcal{J}(T_1',T_2')}(\mathbf{x}\times\theta_1^+\times\theta_2^+) = \mathbf{x}\times\theta_1^+\times\theta_2^+$.

To show that the other maps commute, we simply adapt the above argument and observe that the only disks counted by the change of almost complex structure map which are relevant

to verifying $S^{\circ_1}_{w_1} \circ S^{\circ_2}_{w_2} = S^{\circ_1}_{w_2} \circ S^{\circ_2}_{w_1}$ have both Δ terms nonnegative, so the above argument applies. $\square$

2.6 Graph action map

In this section we describe a sort of 3–dimensional graph TQFT as an intermediate step to defining the graph cobordisms. Roughly speaking, to a graph $\Gamma \subseteq Y$ with two distinguished disjoint subsets of vertices, V_0 and V_1, we will construct a map from $HF^\circ(Y_1, V_0, \mathfrak{s})$ to $HF^\circ(Y_2, V_1, \mathfrak{s})$, called the graph action map. Due to some technical requirements, we will have to construct it in a slightly larger generality, allowing additional vertices $\mathbf{w}_0$ which don't interact with the graph Γ. We now precisely define the graphs involved:

Definition 2.6.1. A **ribbon graph** Γ is a graph equipped with a cyclic ordering of the edges incident to each vertex. A **strong ribbon graph** is a graph Γ equipped with an absolute ordering on the vertices incident to each vertex.

Definition 2.6.2. Suppose that Y^3 is a closed 3–manifold. We call a tuple $\mathcal{G} = (\Gamma, V_0, V_1, (\sigma, \mathfrak{P}), \mathbf{w}_0)$ a strong ribbon (resp. ribbon) **flow graph** if

1. Γ is an embedded, strong ribbon (resp. ribbon) graph;
2. $(\sigma, \mathfrak{P})$ is a coloring of Γ and the basepoints $\mathbf{w}_0$;
3. $\mathbf{w}_0$ is a collection of basepoints in Y with $\mathbf{w}_0 \cap \Gamma = \varnothing$;
4. $V_0, V_1 \subseteq V(\Gamma)$ are nonempty subsets of the vertices $V(\Gamma)$ with $V_0 \cap V_1 = \varnothing$;
5. the inclusions $\pi_0(V_0 \cup \mathbf{w}_0) \to \pi_0(Y)$ and $\pi_0(V_1 \cup \mathbf{w}_0) \to \pi_0(Y)$ are both surjective.
6. each vertex in V_0 or V_1 has valence 1;
7. there are no vertices of $V(\Gamma)$ which are not incident to an edge.

The main theorem of this section is the following:

Theorem B. *Suppose that* $(\mathcal{G}, \Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ *is a strong ribbon flow graph in* Y *and* $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. *If* $V \subseteq V(\Gamma)$ *is a collection of vertices with* $V_0 \subseteq V$, *then then there is a well defined map*

$$A^V_{\mathcal{G},\mathfrak{s}} : HF^\circ(Y, V \cup \mathbf{w}_0, \sigma|_{V \cup \mathbf{w}_0}, \mathfrak{P}, \mathfrak{s}) \to HF^\circ(Y, V_1 \cup \mathbf{w}_0, \sigma|_{V_1} \cup \mathbf{w}_0, \mathfrak{P}, \mathfrak{s}).$$

The map $A^V_{\mathcal{G},\mathfrak{s}}$ *depends only on the ordering of the edges incident to a vertex* $v \notin V$ *up to cyclic permutation. The map* $A^V_{\mathcal{G},\mathfrak{s}}$ *is functorial in the following sense. Suppose that* $\mathcal{G} = (\Gamma, V_0, V_2, \sigma, \mathbf{w}_0)$ *is a flow graph which decomposes as the composition of two flow graphs* $\mathcal{G} = \mathcal{G}_2 \circ \mathcal{G}_1$ *where* $\mathcal{G}_i = (\Gamma_i, V_{i-1}, V_i, \sigma_i)$. *Then*

$$A^{V_1}_{\mathcal{G}_2,\mathfrak{s}} \circ A^{V_0}_{\mathcal{G}_1,\mathfrak{s}} = A^{V_0}_{\mathcal{G},\mathfrak{s}}.$$

Of course the zero maps satisfy the above theorem, but our maps will be much more interesting than the zero maps. To define the graph action maps, we will break a graph into small pieces, define the graph action map for each piece, and then compose all of the maps together.

Definition 2.6.3. We say that a flow graph $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is **elementary** if it has one of the following forms:

1. $|V_1| = |V_0|$ and each edge of Γ connects a vertex of V_0 to a vertex of V_1;
2. There is a single vertex v of Γ which is not in $V_0 \cup V_1$ and all edges of Γ connect either V_0 to V_1 or to a point of $V_0 \cup V_1$ to v;
3. $|V_1| = |V_0| \pm 2$, all edges except for one connect a vertex in V_0 to a vertex in V_1 and there is a single edge which connects two vertices in V_0 to each other, or two vertices in V_1 to each other. We divide elementary flow graphs of type (3) into two subtypes: type $(3i)$ corresponds to the edge connecting two vertices in V_0, and type $(3ii)$ corresponds to the edge connecting two vertices in V_1.

Examples of the three types of elementary flow graphs are depicted in Figure 2.8 .

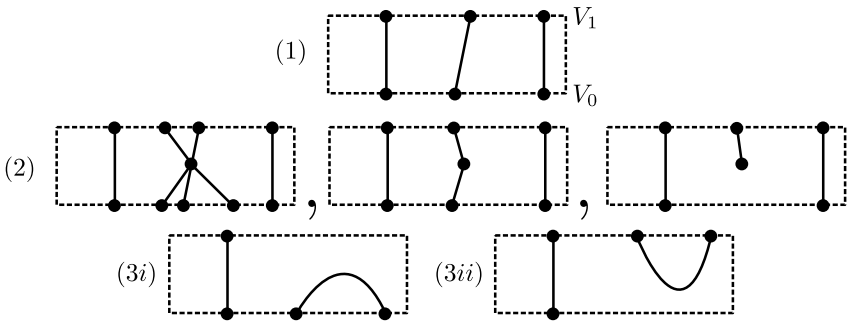

Figure 2.8: Examples of the three types of elementary flow graphs. We view these graphs as flowing from bottom to top. We divide type (3) into two subtypes: type (3i) and (3ii).

Definition 2.6.4. Given a flow graph $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ with a distinguished set of vertices $V \subseteq V(\Gamma)$ with $V_0 \subseteq V$, we define a **Cerf decomposition** of a flow graph $(\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ to be a collection of flow graphs $\mathcal{G}_i = (\Gamma_i, V_{0,i}, V_{1,i}, \sigma_i, \mathbf{w}_0 \cup (V \setminus V(\Gamma_i)))$ for $i = 0, \dots n$ such that

1. $V_0 = V_{0,0}$ and $V_1 = V_{1,n}$;
2. each $\mathcal{G}_i$ is elementary;
3. $V_{0,i} = V_{1,i-1}$ for $i > 0$;
4. $\cup_i \Gamma_i$ is a subdivision of Γ;
5. the $V_{0,i}$ and $V_{1,i}$ do not contain any vertices of Γ (the original, nonsubdivided graph) except for $V_{0,0}$ and $V_{1,n}$;
6. if $\mathcal{G}_i$ is an elementary flow graph of type (2), then the middle vertex, v, is a vertex in the nonsubdivided graph Γ.
7. $\sigma_i = \sigma|_{\Gamma_i \cup \mathbf{w}_0}$.

We emphasize the distinction between the original vertices of Γ, and the vertices of the Cerf decomposition. We refer to the graph Γ with the original collection of vertices as the nonsubdivided graph, and we refer to the vertices of the Cerf decomposition as the subdivided

graph. A vertex of Γ which is not in V_0 or V_1 is never an incoming or outgoing vertex for any of the elementary flow graphs in a Cerf decomposition. Instead a vertex of Γ which is not in V_0 or V_1 is always the middle vertex of an elementary flow graph of type 2.

A Cerf decomposition can be obtained by picking a Morse function $f : \Gamma \to [0, 1]$ with $f^{-1}(i) = V_i$ for $i = 0, 1$ which has connected level sets. Each edge e can be viewed as the image of an immersion of $e : I \to Y$, which is embedded except possibly at the endpoints, where $e(0)$ and $e(1)$ might coincide. By a Morse function $f : \Gamma \to [0, 1]$, we mean a function f such that $f \circ e$ is Morse for each edge $e : I \to Y$. A Cerf decomposition then results by picking level sets of a generic Morse function f such that between any two consecutive level sets, there is at most 1 vertex of Γ or critical point of f. Conversely any Cerf decomposition arises from a Morse function on Γ in this way, since given a Cerf decomposition of a flow graph, a Morse function can easily be built on each elementary flow graph, and glued together to form a Morse function on the entire flow graph which yields the original Cerf decomposition.

Moving between two Cerf decompositions corresponds to doing Cerf theory on compact 1–manifolds with boundary. With the above in mind, we have the following:

Lemma 2.6.5. *Any two Cerf decompositions of a flow graph can be connected by moves corresponding to the following set of moves and their inverses on level of Morse functions:*

1. *inserting a level set, splitting an elementary flow graph of type* (n) *(for* $n \in \{1, 2, 3i, 3ii\}$*), with two elementary flow graphs, one of type* (n) *and one of type (1);*

2. *performing a birth–death move for a pair of critical points along the interior of an edge, replacing two elementary flow graphs of type* $(3i)$ *and* $(3ii)$ *respectively with one of type (1);*

3. *eliminating a critical point along an edge by sending it to a vertex;*

4. *exchanging the relative values of two critical points on different edges of the subdivided graph which don't share a common vertex;*

5. *exchanging the relative values of a vertex v of Γ (the unsubdivided graph) and a critical point which is on an edge such that there is not an edge (of the subdivided graph) which is incident to both v and contains a vertex incident to the edge which has the critical point;*

6. *exchanging the relative values of two vertices v and v' such that there is no pair of edges e and e' in the subdivided graph $\cup_i \Gamma_i$ such that e is incident to v, e' is incident to v' and e and e' share a common vertex;*

Examples of these are shown in Figure 2.9.

Proof. This follows essentially immediately from elementary Cerf theory for 1–manifolds with boundary, though there are several comments we need to make. Given two Morse functions f_0, f_1 on Γ (the unsubdivided graph) we consider a path of smooth functions f_t between f_0 and f_1. Generically such a function f_t will be Morse except at finitely many t. Similarly generically f_t will also only have finitely many t where the vertices of Γ and the critical points of f_t do not have distinct values.

Suppose that $\mathcal{C}_0$ and $\mathcal{C}_1$ are Cerf decompositions obtained by taking a collection of values

$$0 = s_0^i < \cdots < s_{n_i}^i = 1$$

such that $f_i^{-1}([s_j^i, s_{j+1}^i])$ contains at most one critical point of f_i or one vertex of Γ. By applying moves of type (1), we can assume that there is exactly one critical point of f_i or vertex of Γ in each $f_i^{-1}([s_j^i, s_{j+1}^i])$.

Now if there are no points $t \in [a, b]$ where f_t has not Morse and all of the critical values and vertex values of f are distinct, then let $c_1^t, \ldots, c_n^t$ denote the critical values and vertex values of f, ordered so that $c_1^t < \cdots c_n^t$. Pick values $s_j^t \in [0, 1]$ extending s_j^i for $i \in \{0, 1\}$ such that

$$0 = s_0^t < c_1^t < s_1^t < \cdots < c_n^t < s_n^t = 1.$$

The sets $\Gamma_{i,t} = f_t^{-1}([s_i^t, s_{i+1}^t)$ yield the underlying graph for a Cerf decomposition of $\mathcal{G}$. If $\mathcal{C}_t$

denotes the Cerf decomposition of $\mathcal{G}$ arising from the subsets $\Gamma_{i,t}$, then it is easy to see that the $\mathcal{C}_a$ and $\mathcal{C}_b$ are related by a sequence of moves of type (1).

Suppose a birth-death singularity of f_t occurs at a point $t = c \in (0, 1)$. Then if there are no other singularities of f_t in $[c - \epsilon, c + \epsilon]$, and $\mathcal{C}_{c-\epsilon}, \mathcal{C}_{c+\epsilon}$ are Cerf decompositions as above for $f_{c-\epsilon}$ and $f_{c+\epsilon}$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ are related by a composition of moves of type (1) and moves of type (2).

Similarly if a critical point is destroyed or created along on edge at a vertex incident to that edge at a time $c \in (0, 1)$, and there are no other singularities in $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ are related by a move of type (3) and moves of type (1).

Suppose two critical points exchange relative value at a time $c \in (0, 1)$, and there are no other singularities in $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ are related by a move of type (4) and moves of type (1) as long as the edges containing the critical two points do not contain a common vertex. Observe that if the edges with a critical point contain a common vertex (of the subdivided graph), then at $t = c$ the derivative at all of the points between the two critical points must be zero, which is a non-generic singularity. Hence generically the two critical points switching relative values corresponds to a composition of moves (4) and (1).

If a critical point and a vertex exchange relative value at a point $c \in (0, 1)$ and there are no other singularities in $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}$ and $\mathcal{C}_{c+\epsilon}$ relative by a move of type (5) and moves of type (1) so long as there is no edge which is both incident to both the vertex v and one of the two vertices on the edge with a critical point. Note that once again, if there were such an edge, then at the derivative of f_c along this edge must vanish, which is a nongeneric singularity. Hence generically we can assume that there is no such edge (of the subdivided graph).

Finally if two vertices v and v' (of the nonsubdivided graph Γ) exchange relative values and there are no other singularities in the interval $[c - \epsilon, c + \epsilon]$, then $\mathcal{C}_{c-\epsilon}, \mathcal{C}_{c+\epsilon}$ are related by moves (1) and (6), so long as there are not edges e and e' in the subdivided graph such that e is incident to v, e' is incident to v', and e and e' contain a common vertex. As in the

previous cases, if there were two such edges e and e', then f_c would have vanishing derivative on all of $e \cup e'$, which is a nongeneric singularity.

We note that these are the only singularities that occur generically for a 1–parameter family of functions f_t defined on a graph Γ.

We finally comment that for a Morse function to yield a Cerf decomposition as we've described, we must have $\mathbf{w}_0 \cup f_t^{-1}(s)$ nonempty for any s. If $\mathbf{w}_0 \neq \varnothing$, this is trivially satisfied. In the case that $\mathbf{w}_0 = \varnothing$, given two Morse functions on a flow graph which both have nonempty level sets, it is possible to connect them via a generic path of functions which have nonempty level sets. This can be achieved by pulling a small portion of a terminal edge down through the level sets and fixing the value on this portion through the homotopy of Morse functions, and then moving it back into its final position.

□

We now construct the graph action map. Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is an embedded strong ribbon flow graph in Y^3 and that $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Suppose that $V \subseteq V(\Gamma)$ is a distinguished set of vertices with $V_0 \subseteq V$.

We first define the graph action map for an elementary strong ribbon flow graph. Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is an elementary strong ribbon flow graph of type $(1), (3i)$, or $(3ii)$. Then we define

$$A^V_{\mathcal{G},\mathfrak{s}} = \prod_{v \in V(\Gamma)\setminus V_1} S^-_{v,\mathfrak{s}} \circ \prod_{e \in E(\Gamma)} A_{e,\mathfrak{s}} \circ \prod_{v \in V(\Gamma)\setminus V} S^+_{v,\mathfrak{s}}.$$

Notice that this formula is well defined because no two edges share a common vertex, so by Lemma 2.4.10 we know that the order of the various A_e is not important. Similarly the S^+_v and S^-_v can be commuted amongst each other.

Now suppose that $\mathcal{G}$ is an elementary strong ribbon flow graph of type 2. Let v denote the unique vertex which is not in V_0 or V_1. Let $e_1, \ldots, e_n$ be the edges adjacent to v with

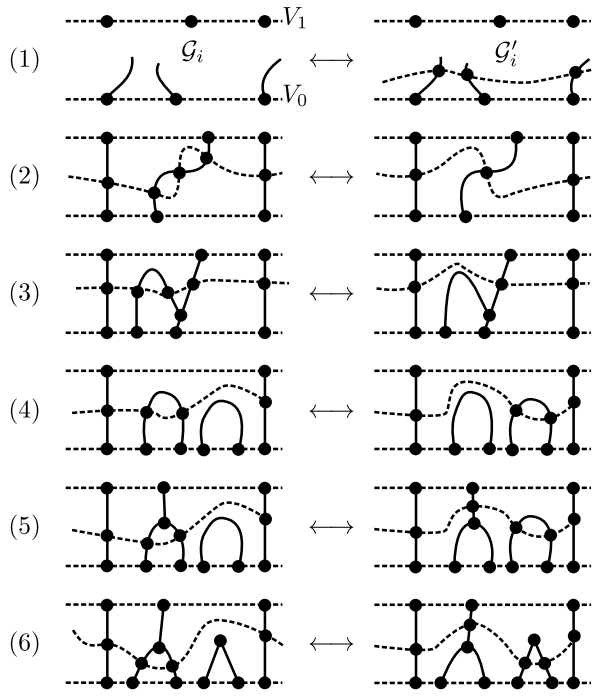

Figure 2.9: Six moves which can be used to move between any two Cerf decompositions of a strong ribbon flow graph. All of the dashed lines indicate level sets of Morse functions, and the space between them indicates an elementary strong ribbon flow graph. The dashed lines may be somewhat misleading, because strictly speaking the Morse function is only defined on the graph.

the ordering from Γ. We define the graph action map by the formula

$$A^V_{\mathcal{G},\mathfrak{s}} = \left(\prod_{v \in V(\Gamma) \setminus V_1} S^-_{v,\mathfrak{s}} \right) \circ \left(\prod_{e \in E(\Gamma) \setminus \{e_1,\dots,e_n\}} A_{e,\mathfrak{s}} \right) \circ (A_{e_n,\mathfrak{s}} \cdots A_{e_1,\mathfrak{s}}) \circ \left(\prod_{v \notin V} S^+_{v,\mathfrak{s}} \right).$$

Now given an arbitrary strong ribbon flow graph $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ on Y^3, and a collection $V \subseteq V(\Gamma)$ such that $V_0 \subseteq V$, we pick a Cerf decomposition $(\mathcal{G}_1, \dots, \mathcal{G}_n)$ of $\mathcal{G}$, and define

$$A^V_{\mathcal{G},\mathfrak{s}} : HF^\circ(Y, V \cup \mathbf{w}_0, \mathfrak{s}, \sigma|_{V \cup \mathbf{w}_0}) \to HF^\circ(Y, V_1 \cup \mathbf{w}_0, \mathfrak{s}, \sigma|_{V_1 \cup \mathbf{w}_0})$$

as the composition

$$A^V_{\mathcal{G},\mathfrak{s}} = \prod_{i=1}^{n} A^{V_{\geq i}}_{\mathcal{G}_i,\mathfrak{s}},$$

where

$$V_{\geq i} = V \cap \bigcup_{j \geq i} V(\mathcal{G}_j)$$

Theorem 2.6.6. *If $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph with V, then the map $A_{\mathcal{G}}$ is independent of the Cerf decomposition $(\mathcal{G}_1, \ldots, \mathcal{G}_n)$.*

We need to show that the composition $A_{\mathcal{G}}$ is unchanged by moves (1) through (6) in Lemma 2.6.5. Before we prove the above theorem, we will first state a model computation that we will use repeatedly, but will prove later:

Lemma 2.6.7. *Suppose that e is a path from v to v'. Then there exists a diagram and an almost complex structure such that, viewing CF° as begin freely-stabilized at v, we have that*

$$A_e = \begin{pmatrix} 0 & U_v \\ 1 & 0 \end{pmatrix},$$

and furthermore this diagram and almost complex structure can be used to compute S_v^+ and S_v^-. In particular

$$A_e^2 = U_v$$

and

$$S_v^- A_e S_v^+ = 1.$$

The proof will be given later, as Lemma 2.13.12. The latter two relations follow easily from the first claim. Note that $A_e^2 \simeq U_v$ was previously proven in Lemma 2.4.12, by counting the ends of moduli spaces of Maslov index 2–disks.

The following lemma states essentially that the graph action map is invariant under subdivision of an edge.

Lemma 2.6.8. *Suppose e_1 is a path from w to w' and e_2 is a path from w' to w''. Let $A_{e_2 * e_1}$ denote the relative homology map on a diagram which does not have w' as a basepoint. Then we have*

$$A_{e_2 * e_1} = S_{w'}^- A_{e_2} A_{e_1} S_{w'}^+.$$

Proof. By Lemma 2.6.7 we have $S^-_{w'} A_{e_1} S^+_{w'} = 1$. Hence we have that

$$\begin{aligned} A_{e_2 * e_1} &= A_{e_2 * e_1}(S^-_{w'} A_{e_1} S^+_{w'}) \\ &= S^-_{w'} A_{e_2 * e_1} A_{e_1} S^+_{w'} \\ &= S^-_{w'}(A_{e_2} + A_{e_1}) A_{e_1} S^+_{w'} \\ &= S^-_{w'} A_{e_2} A_{e_1} S^+_{w'} + S^-_{w'} U_{w'} S^+_{w'} \\ &= S^-_{w'} A_{e_2} A_{e_1} S^+_{w'}, \end{aligned}$$

as we wanted. □

We need another lemma, about subdividing an edge into thirds. From here on, let $S^+_{w_1 w_2 \cdots w_n}$ abbreviate $S^+_{w_1} \circ \cdots \circ S^+_{w_n}$ and similarly for the destabilization maps.

Lemma 2.6.9. *Suppose e_1, e_2 and e_3 are edges from w_0 to w_1, from w_1 to w_2, and from w_2 to w_3, respectively, as is depicted in Figure 2.10. Suppose also that $\sigma : \{1,2,3\} \to \{1,2,3\}$ is a bijection, then*

$$S^-_{w_1 w_2} A_{e_{\sigma(3)}} A_{e_{\sigma(2)}} A_{e_{\sigma(1)}} S^+_{w_1 w_2} = S^-_{w_1 w_2} A_{e_3} A_{e_2} A_{e_1} S^+_{w_1 w_2}.$$

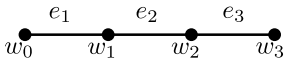

Figure 2.10: Vertices and edges in Lemma 2.6.9.

Proof. We will show that if e_i, e_j and e_k are the edges, then

$$S^-_{w_1 w_2} A_{e_i} A_{e_j} A_{e_k} S^+_{w_1 w_2} = S^-_{w_1 w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1 w_2}.$$

The same argument will apply to show that the expression is invariant under switching the order of the left two relative homology maps as well.

There are two cases: if e_j and e_k share a vertex, or if they do not. If e_i and e_j do not share a vertex, then $A_{e_j} A_{e_k} = A_{e_k} A_{e_j}$ so the above equality trivially holds.

Now suppose that e_j and e_k share the vertex $w_n \in \{w_1, w_2\}$. Let w_m denote the other vertex in $\{w_1, w_2\}$.

We simply compute that

$$\begin{aligned} S^-_{w_1w_2} A_{e_i} A_{e_j} A_{e_k} S^+_{w_1w_2} &= S^-_{w_1w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1w_2} + S^-_{w_1w_2} A_{e_i} U_{w_n} S^+_{w_1w_2} \\ &= S^-_{w_1w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1w_2} + S^-_{w_m} U_{w_n} S^-_{w_n} S^+_{w_n} A_{e_i} S^+_{w_m} \\ &= S^-_{w_1w_2} A_{e_i} A_{e_k} A_{e_j} S^+_{w_1w_2}. \end{aligned}$$

□

We now have the necessary tools to prove the invariance of the graph action map:

Proof of Theorem 2.6.6. **Move (1):** In the case that the elementary flow graph $\mathcal{G}$ is of type (1), or (2). This follows from Lemma 2.6.8. In the case that the graph is of type $(3i)$ or $(3ii)$, invariance follows from Lemmas 2.6.8 and 2.6.9.

Move (2): Suppose that v_1, v_2, v_3, v_4, and v_5 are vertices, and e_1, e_2, e_3 and e_4 are edges, as in Figure 2.11.

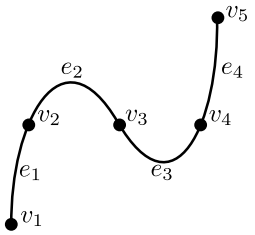

Figure 2.11: Vertices and edges used in the proof of invariance under move (2).

Write e for the concatenation of e_1, e_2, e_3 and e_4. To prove invariance under move (2), it is sufficient to show that

$$S^-_{v_1} A_e = (S^-_{v_2v_3v_4} A_{e_4} A_{e_2})(S^-_{v_1} A_{e_3} A_{e_1} S^+_{v_2v_3v_4}).$$

To this end, we compute as follows:

$$
\begin{aligned}
S^-_{v_1}A_e &= S^-_{v_1v_2v_3v_4}A_{e_4}A_{e_3}A_{e_2}A_{e_1}S^+_{v_2v_3v_4} && \text{(Lemma 2.6.8)}\\
&= S^-_{v_1v_2v_3v_4}A_{e_4}A_{e_2}A_{e_3}A_{e_1}S^+_{v_2v_3v_4} && \text{(Lemma 2.6.9)}\\
&= S^-_{v_2v_3v_4}A_{e_4}A_{e_2}S^-_{v_1}A_{e_3}A_{e_1}S^+_{v_2v_3v_4} && \text{(Lemma 2.5.6)},
\end{aligned}
$$

as we wanted.

Move (3): We consider move (3) shown in Figure 2.9. There is another move which we would need to check obtained by flipping this picture upside down, but we leave checking that move to the reader since the proof is identical to what we will present. Let v, v', v'' and $e_2, \ldots, e_n, e, e'$, and e'' be as in Figure 2.12. We let e_n be the concatenation of e, e' and e''.

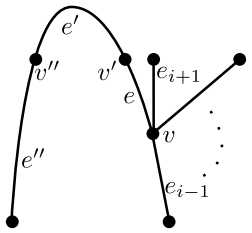

Figure 2.12: Vertices and edges used in the proof of invariance under move (3). Here $e_i = e * e' * e''$.

To show invariance under move (3), it is sufficient to show that

$$
A_{e_n}\cdots A_{e_1}S^+_v = (S^-_{v'v''}A_{e'})(A_{e''}A_{e_n}\cdots A_{e_{i+1}}A_eA_{e_{i-1}}\cdots A_{e_1}S^+_{v'v''}).
$$

To this end, we compute:

$$
\begin{aligned}
A_{e_n}\cdots A_{e_i}\cdots A_{e_1} &= A_{e_n}\cdots A_{e_{i+1}}(S^-_{v'}A_{e'*e''}A_eS^+_{v'})A_{e_{i-1}}\cdots A_{e_1} && \text{(Lemma 2.6.8)}\\
&= A_{e_n}\cdots A_{e_{i+1}}(S^-_{v'}(S^-_{v''}A_{e'}A_{e''}S^+_{v''})A_eS^+_{v'})A_{e_{i-1}}\cdots A_{e_1} && \text{(Lemma 2.6.8)}\\
&= S^-_{v''v'}A_{e_n}\cdots A_{e_{i+1}}A_{e''}A_{e'}A_eA_{e_{i-1}}\cdots A_{e_1}S^+_{v'v''} && \text{(Lemma 2.5.6)}\\
&= (S^-_{v''v'}A_{e'})(A_{e''}A_{e_n}\cdots A_{e_{i+1}}A_eA_{e_{i-1}}\cdots A_{e_1}S^+_{v'v''}) && \text{(Lemma 2.4.10)}
\end{aligned}
$$

as we wanted.

Moves (4)–(6): Invariance under moves (4)–(6) follows similarly from Lemmas 2.6.8, 2.6.9 and 2.4.10. □

Having established the independence of the graph action map from the choice of Cerf decomposition, we now prove two lemmas about the graph action map which will be useful later when we construct and analyze the graph cobordism maps. Firstly, the graph action map is invariant under subdivision, in the following sense:

Lemma 2.6.10. *Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph in Y^3 and $\mathcal{G}'$ is a subdivision of $\mathcal{G}$ formed by adding a single vertex v to the interior of an edge in Γ. Then if V is a collection of vertices containing V_0 with $v \notin V$, then*

$$A^V_{\mathcal{G}} = A^{V\cup\{v\}}_{\mathcal{G}'} \circ S^+_v,$$

for either absolute ordering of the two edges adjacent to v.

Proof. Pick a Cerf decomposition of $\mathcal{G}$ such that the vertex v appears in an elementary flow graph $\mathcal{G}_v$ of type (1) (a trivial flow graph). Let e_v denote the edge of $\mathcal{G}_v$ containing v. Let e_1 and e_2 denote the two components of $e_v \setminus \{v\}$, with either ordering. By Corollary 2.9.3 we know that

$$A_{e_v} = S^-_v A_{e_1} A_{e_2} S^+_v.$$

The left hand side A_{e_v} is a factor in $A^V_{\mathcal{G}}$. If we insert the right hand expression into the formula for $A^V_{\mathcal{G}}$ and then commute S^+_v with everything to the right, we get exactly $A^{V\cup\{v\}}_{\mathcal{G}'} \circ S^+_v$. □

We now consider the effect of adding a "trivial strand" to a strong ribbon flow graph:

Lemma 2.6.11. *Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \mathbf{w}_0)$ is a strong ribbon flow graph in Y and let Γ_0 denote a graph obtained by adding a new vertex $v \in Y \setminus \Gamma$ to Γ and also adding an edge connecting v to a point on the interior of an edge of Γ. Let $\mathcal{G}_0 = (\Gamma_0, V_0, V_1, \sigma, \mathbf{w}_0)$, $\mathcal{G}' = (\Gamma_0, V_0 \cup \{v\}, V_1, \sigma, \mathbf{w}_0)$ and $\mathcal{G}'' = (\Gamma_0, V_0, V_1 \cup \{v\}, \sigma, \mathbf{w}_0)$. Then*

$$A_{\mathcal{G}_0} = A_{\mathcal{G}},$$

$$S_v^- \circ A_{\mathcal{G}'} = A_{\mathcal{G}},$$

and

$$A_{\mathcal{G}''} \circ S_v^+ = A_{\mathcal{G}}.$$

This holds for any ordering of the vertices adjacent to v.

Proof. We will first show that $A_{\mathcal{G}_0} = A_{\mathcal{G}}$ with the absolute ordering of e_1, e_2, λ. By taking a Cerf decomposition of $\mathcal{G}$ which we have subdivided enough, one easily computes that it is sufficient to show that

$$S_v^- S_{v_2}^- A_\lambda A_{e_2} A_{e_1} S_{v_2}^+ S_v^+ = S_{v_2}^- A_{e_2} A_{e_1} S_{v_2}^+,$$

where $v, v_1, v_2, \lambda, e_1$ and e_2 are as in Figure 2.13.

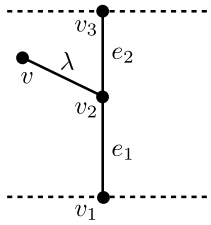

Figure 2.13: Edges and vertices in the proof of Lemma 2.6.11. The dashed lines indicate the incoming and outgoing ends of the flow graph.

By Lemma 2.6.7, we have that $S_v^- A_\lambda S_v^+ = 1$, and hence we have

$$\begin{aligned} S_{v_2}^- A_{e_2} A_{e_1} S_{v_2}^+ &= S_{v_2}^- (S_v^- A_\lambda S_v^+) A_{e_2} A_{e_1} S_{v_2}^+ \\ &= S_v^- S_{v_2}^- A_\lambda A_{e_2} A_{e_1} S_{v_2}^+ S_v^+, \end{aligned}$$

as we wanted. We conclude that $A_{\mathcal{G}} = A_{\mathcal{G}_0}$. To see that the statement holds for any other absolute ordering of the edges e_1, e_2, and λ, we observe that in the above computation

$$S_{v_2}^- A_{e_2} A_{e_1} S_{v_2}^+ = S_{v_2}^- A_{e_1} A_{e_2} S_{v_2}^+,$$

allowing us to reorder e_1 and e_2. Also the expression $S_v^- A_\lambda S_v^+$ can be inserted anywhere in $S_{v_2} A_{e_1} A_{e_2} S_{v_2}^+$, allowing us to reorder λ and either of the edges e_1 and e_2.

We now consider $A_{\mathcal{G}'}$ and $A_{\mathcal{G}''}$. We note that we've established that $A_{\mathcal{G}} = A_{\mathcal{G}_0}$ for any Cerf decomposition of $\mathcal{G}_0$. In particular, a Cerf decomposition of $\mathcal{G}'$ naturally induces a Cerf decomposition of $\mathcal{G}_0$. The only difference in the two formulas defining $A_{\mathcal{G}'}$ and $A_{\mathcal{G}_0}$ is a factor of S_v^- which is present in $A_{\mathcal{G}_0}$ but not in $A_{\mathcal{G}'}$ and hence

$$S_v^- \circ A_{\mathcal{G}'} = A_{\mathcal{G}_0},$$

establishing the desired equality.

Similarly to show that $A_{\mathcal{G}''} \circ S_v^+ = A_{\mathcal{G}_0}$, we pick a Cerf decomposition of $\mathcal{G}''$, which naturally induces one for $\mathcal{G}_0$. The only difference in the two formulas defining these operations is a factor of S_v^+ to the left of all other terms in $A_{\mathcal{G}''}$ but not in $A_{\mathcal{G}_0}$, establishing the desired equality.

□

We finally show that the graph action map is invariant of cyclic reordering of the edges incident to a vertex $v \in V(\Gamma)$ as long as $v \notin V$, the initial set of vertices of the graph action map.

Lemma 2.6.12. *Suppose that $e_1, \ldots, e_n$ are edges which each share a vertex v (and v has valence 1 in each edge e_i). Then*

$$S_v^- A_{e_n} \cdots A_{e_2} A_{e_1} S_v^+ = S_v^- A_{e_{n-1}} \cdots A_{e_1} A_{e_n} S_v^+.$$

Proof. We first prove the cases $n = 1, 2$ and $n = 3$ directly. The cases $n > 3$ will follow by an induction argument. First note that $n = 1$ is vacuous. The $n = 2$ is relatively easy. We just compute directly that

$$S_v^- A_{e_2} A_{e_1} S_v^+ = S_v^- U_v S_v^+ + S_v^- A_{e_1} A_{e_2} S_v^+ = S_v^- A_{e_1} A_{e_2} S_v^+,$$

by Lemma 2.4.10. The $n = 3$ is slightly more involved, though still relatively easy. We

simply compute that

$$\begin{aligned} S_v^- A_{e_2} A_{e_1} A_{e_0} S_v^+ &= S_v^- A_{e_1} A_{e_2} A_{e_0} S_v^+ + U_v S_v^- A_{e_0} S_v^+ \\ &= S_v^- A_{e_1} A_{e_0} A_{e_2} S_v^+ + U_v S_v^- (A_{e_0} + A_{e_2}) S_v^+ \\ &= S_v^- A_{e_1} A_{e_0} A_{e_2} S_v^+ + U_v A_{e_1 * e_0} S_v^- S_v^+ \\ &= S_v^- A_{e_1} A_{e_0} A_{e_2} S_v^+. \end{aligned}$$

We now prove the case of $n > 3$. This is more involved. First pick an extra vertex w which is not an endpoint of any of the e_i and a path e from v to w. We first claim that

$$A_{e_2} A_{e_1} = S_w^- A_{e_2 * e} A_{e_1 * e} A_e S_w^+.$$

To see this, note that $A_{e_2 * e} = A_{e_2} + A_e$. Hence we just compute that

$$\begin{aligned} S_w^- A_{e_2 * e} A_{e_1 * e} A_e S_w^+ &= S_w^- (A_{e_2} + A_e)(A_{e_1} + A_e) A_e S_w^+ \\ &= S_w^- A_{e_2} A_{e_1} A_e S_w^+ + S_w^- (A_{e_2} A_e^2 + A_e A_{e_1} A_e + A_e^3) S_w^+ \\ &= A_{e_2} A_{e_1} S_w^- A_e S_w^+ + S_w^- (A_{e_2} U_w + A_{e_1} U_w + U_w A_e + U_w A_e) S_w^+ \\ &= A_{e_2} A_{e_1} + U_w S_w^- A_{e_2 * e_1} S_w^+ \\ &= A_{e_1} A_{e_1} + U_w A_{e_2 * e_1} S_w^- S_w^+ \\ &= A_{e_1} A_{e_1}, \end{aligned}$$

as we wanted.

Let $\widetilde{e}_1 = e_1 * e$ and let $\widetilde{e}_2 = e_2 * e$. We compute as follows:

$$\begin{aligned} S_v^- A_{e_n} \cdots A_{e_2} A_{e_1} S_v^+ &= S_v^- A_{e_n} \cdots A_{e_3} (S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_e S_w^+) S_v^+ \\ &= S_v^- S_w^- (A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_{e_n} \cdots A_{e_3} A_e) S_w^+ S_v^+ \\ &= S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_n} \cdots A_{e_3} A_e S_v^+) S_w^+. \end{aligned} \tag{2.7}$$

By induction, Equation (2.7) is equal to

$$\begin{aligned} S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_n} \cdots A_{e_3} A_e S_v^+) S_w^+ &= S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_e A_{e_n} \cdots A_{e_3} S_v^+) S_w^+ \\ &= S_v^- S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_e S_w^+ A_{e_n} \cdots A_{e_3} S_v^+ \\ &= S_v^- A_{e_2} A_{e_1} A_{e_n} \cdots A_{e_3} S_v^+. \end{aligned}$$

This implies that the expression $S_v^- A_{e_n} \cdots A_{e_1} S_v^+$ is invariant under permuting the e_i cyclically by two elements at a time.

Similarly, by induction we know that Equation (2.7) is also equal to

$$\begin{aligned} S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_n} \cdots A_{e_3} A_e S_v^+) S_w^+ &= S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} (S_v^- A_{e_3} A_e A_{e_n} \cdots A_{e_4} S_v^+) S_w^+ \\ &= S_v^- A_{e_3} (S_w^- A_{\widetilde{e}_2} A_{\widetilde{e}_1} A_e S_w^+) A_{e_3} A_{e_n} \cdots A_{e_4} S_v^+ \\ &= S_v^- A_{e_3} A_{e_2} A_{e_1} A_{e_n} \cdots A_{e_4} S_v^+, \end{aligned}$$

which shows that the expression $S_v^- A_{e_n} \cdots A_{e_1} S_v^+$ is invariant under permuting the e_i cyclically three elements at a time. Since it is also invariant under permuting the e_i cyclically two elements at a time, we know that it is invariant under permuting them cyclically one at a time, so the main claim now follows. □

Theorem B now follows from Theorem 2.6.6 Lemma 2.6.12.

2.7 1– and 3–handles

In this section we construct and prove invariance of 1– and 3–handle maps. Our construction of the maps is very similar to the ones in [OS06] and [Juh16], though it has some notable differences. Unlike the construction in [OS06], our construction works for 1–handles with feet in different components, allowing us to consider cobordisms between manifolds with disconnected ends. Our construction is somewhat more flexible than the one in [Juh16], since we don't require the feet of our 1–handles to be in a neighborhood of the basepoints, a trick that eliminates most of the need for gluing lemmas for the hat flavor, but which isn't useful for the $-, \infty$ and $+$ flavors.

2.7.1 Preliminaries on 1– and 3–handles

Suppose that $p_1, p_2 \in Y$ are two points. By choosing an orientation on $\{p_1, p_2\}$, we can view $\{p_1, p_2\}$ as a framed zero sphere $\mathbb{S}_0 \subseteq Y$. Surgery on $\mathbb{S}_0$ corresponds to removing two disjoint balls containing p_1 and p_2, then gluing a copy of $S^2 \times I$ connecting the two resulting boundary components. We call the resulting 3–manifold $Y(\mathbb{S}_0)$, though it depends on the choice of balls containing the points p_1 and p_2 (though of course any two choices of balls yield manifolds which are diffeomorphic with diffeomorphism which is unique up to isotopy). In terms of cobordisms, we can construct a cobordism $W(\mathbb{S}_0)$ with boundary $-Y \sqcup Y(\mathbb{S}_0)$ by attaching a 1–handle to $Y \times I$ along $\mathbb{S}_0 \times \{1\}$.

Dually if Y has an embedded 2–sphere $\mathbb{S}_2$ with trivial normal bundle, then after removing a regular neighborhood of $\mathbb{S}_2$, we get a three manifold with two boundary components, both of which are 2–spheres. By filling in the boundary components with balls, we get a closed three manifold $Y(\mathbb{S}_2)$. In terms of cobordisms, we can form a cobordism $W(\mathbb{S}_2)$ by attaching a 3–handle to $Y \times I$ along $\mathbb{S}_2 \times \{1\}$.

Definition 2.7.1. Given an embedded 0–sphere $\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y \setminus \mathbf{w}$ and a choice of open balls $\mathcal{B}_1, \mathcal{B}_2 \subseteq Y \setminus \mathbf{w}$ such that $p_i \in \mathcal{B}_i$ and $\mathcal{B}_1 \cap \mathcal{B}_2 = \varnothing$, we say that a Heegaard splitting $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is **subordinate** to $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ if $p_1, p_2 \in \Sigma$, and $\mathcal{B}_i \cap \Sigma$ is a disk which is disjoint from $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w}$.

Given a Spin^c structure $\mathfrak{s}$ on Y, and a framed sphere $\mathbb{S}_0$, there is a unique Spin^c structure $\overline{\mathfrak{s}}$ on $W(\mathbb{S}_0)$ extending $\mathfrak{s}$. Let $\mathfrak{s}'$ denote the restriction of $\overline{\mathfrak{s}}$ to $Y(\mathbb{S}_0)$. The map $\mathfrak{s} \mapsto \mathfrak{s}'$ is an injection. Dually, given a framed 2–sphere $\mathbb{S}_2 \subseteq Y$ and a Spin^c structure $\mathfrak{s}' \in \text{Spin}^c(Y)$, if $\mathfrak{s}'$ extends over $W(\mathbb{S}_2)$, then it does so uniquely. We let $\overline{\mathfrak{s}}$ denote the extension over $W(\mathbb{S}_2)$ and we let $\mathfrak{s}$ denote the restriction of $\overline{\mathfrak{s}}$ to $Y(\mathbb{S}_2)$.

We will define maps

$$F^\circ_{Y,\mathbb{S}_0,\overline{\mathfrak{s}}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}) \to HF^\circ(Y(\mathbb{S}_0), \mathbf{w}, \mathfrak{s}')$$

and

$$F^{\circ}_{Y,\mathbb{S}_2,\bar{\mathfrak{s}}} : HF^{\circ}(Y, \mathbf{w}, \mathfrak{s}') \to HF^{\circ}(Y(\mathbb{S}_2), \mathbf{w}, \mathfrak{s}).$$

If we let $\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y(\mathbb{S}_2)$ where each p_i is in one of the balls that we used to fill in the boundary components of $Y \setminus N(\mathbb{S}_2)$, we clearly have $Y(\mathbb{S}_2)(\mathbb{S}_0) = Y$, so instead of defining maps from $HF^{\circ}(Y, \mathbf{w}, \mathfrak{s})$ to $HF^{\circ}(Y(\mathbb{S}_2), \mathbf{w}, \mathfrak{s}')$ for embedded 2–spheres, it is equivalent to define maps $HF^{\circ}(Y(\mathbb{S}_0), \mathbf{w}, \mathfrak{s}') \to HF^{\circ}(Y, \mathbf{w}, \mathfrak{s})$ for embedded 0–spheres $\mathbb{S}_0$, which is the strategy we will take.

Now given a Heegaard diagram $\mathcal{H}$ which is subordinate to $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$, and some additional data $\mathcal{J}$ (gluing data) which we will define in the next subsection, we will define maps

$$F^{\circ}_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}} : CF^{\circ}_{J_s}(\mathcal{H}, \mathfrak{s}) \to CF^{\circ}_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}}, \mathfrak{s}')$$

and

$$G^{\circ}_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}} : CF_{\mathcal{J}(T)}(\mathcal{H}_{\mathcal{J}}, \mathfrak{s}') \to CF^{\circ}_{J_s}(\mathcal{H}, \mathfrak{s}).$$

If $\mathcal{H}$ is the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$, then $\mathcal{H}_{\mathcal{J}} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \mathbf{w})$ is a diagram formed by removing disks from Σ centered at p_1 and p_2, then connecting the corresponding boundary components in $Y(\mathbb{S}_0)$ by an annulus. We define $\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\alpha_0\}$ and $\overline{\boldsymbol{\beta}} = \boldsymbol{\beta} \cup \{\beta_0\}$, where α_0 and β_0 are homotopically nontrivial, simple closed curves in the annulus which are isotopies of one another and satisfy $|\alpha_0 \cap \beta_0| = 2$.

Note that on the level of complexes, we have the following:

$$CF^{\circ}(\mathcal{H}_{\mathcal{J}}, \mathfrak{s}') = CF^{\circ}(\mathcal{H}, \mathfrak{s}) \otimes_{\mathbb{Z}_2} \mathbb{Z}_2^2 = CF^{\circ}(\mathcal{H}, \mathfrak{s}) \oplus CF^{\circ}(\mathcal{H}, \mathfrak{s}).$$

We will define the 1–handle cobordism maps by

$$F^{\circ}_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

where θ^+ is the top degree intersection point of $\alpha_0 \cap \beta_0$. We will define the 3–handle map by

$$G^{\circ}_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}}(\mathbf{x} \times \theta^+) = 0, \qquad \text{and} \qquad G^{\circ}_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J},\bar{\mathfrak{s}}}(\mathbf{x} \times \theta^-) = \mathbf{x},$$

for appropriate choice of almost complex structure.

Lemma 2.7.2. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for $(Y, \mathbf{w})$ which is subordinate to the zero sphere $\mathbb{S}_0 = \{p, p'\} \subseteq \Sigma$. Let $\overline{\mathcal{H}}$ be the diagram described above for $Y(\mathbb{S}_0)$. Suppose $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ and that $\mathfrak{s}' \in \mathrm{Spin}^c(Y(\mathbb{S}_0))$ is the Spin^c structure corresponding to $\mathfrak{s}$ under the correspondence in the previous paragraphs. Then $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible iff $\overline{\mathcal{H}}$ is strongly $\overline{\mathfrak{s}}$–admissible.*

Proof. Suppose $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible. Suppose $\mathcal{D} = \mathcal{D}_1 \# \mathcal{D}_2$ is a periodic domain on $\overline{\mathcal{H}}$ with $\langle c_1(\overline{\mathfrak{s}}), H(\mathcal{D})\rangle = 2N$, where $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ and $\mathcal{D}_2$ is a periodic domain on our standard diagram for $(S^2, \alpha_0, \beta_0, w, w')$. Using Lemmas 2.3.7 and 2.5.3, an easy computation shows that

$$\langle c_1(\overline{\mathfrak{s}}), H(\mathcal{D})\rangle = \langle c_1(\mathfrak{s}), H(\mathcal{D}_1)\rangle,$$

so $\mathcal{D}_1$ has a multiplicity greater than N, and hence so does $\mathcal{D}$. It follows that $\overline{\mathcal{H}}$ is strongly $\overline{\mathfrak{s}}$–admissible.

Conversely if $\overline{\mathcal{H}}$ is strongly $\overline{\mathfrak{s}}$–admissible and $\mathcal{D}_1$ is a periodic domain on $\mathcal{H}$ with $\langle c_1(\mathfrak{s}), H(\mathcal{D})\rangle = 2N$. Let A_w and $A_{w'}$ be the components of $S^2 \setminus \alpha_0$ containing w and w' respectively. Let $\mathcal{D}_2$ be the domain on $(S^2, \alpha_0, \beta_0, w, w')$ defined by

$$\mathcal{D}_2 = n_p(\mathcal{D}_1)A_w + n_{p'}(\mathcal{D}_1)A_{w'}.$$

Since $\overline{\mathcal{H}}$ is strongly $\overline{\mathfrak{s}}$–admissible, we know that $\mathcal{D}_1 \# \mathcal{D}_2$ has coefficients strictly bigger than N, and hence so does $\mathcal{D}_1$, so $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible. □

2.7.2 Gluing data for 1– and 3–handle maps

Very similarly to the free stabilization maps, we now define "gluing data" for the 1– and 3–handle maps. The only real difference is that we are gluing at two points instead of one. In the following definition, let $(S^2, \alpha_0, \beta_0, w, w')$ be the diagram for (S^3, w, w') shown in Figure 2.5.

Definition 2.7.3. Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(Y, \mathbf{w})$ which is subordinate to a framed sphere $\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y \setminus \{w\}$. We call a collection

$$\mathcal{J} = (\mathcal{A}, J_s, J_s', B_1, B_1', B_2, B_2', r_1, r_1', r_2, r_2', p_1', p_2', \iota, \phi_1, \phi_2)$$

gluing data for $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ where

1. $\mathcal{A}$ is a choice of annulus in $S^2 \times I \subseteq Y(\mathbb{S}_0)$ extending the circles $\Sigma \cap \partial\mathcal{B}_i$ such that $\mathcal{A} \cap (S^2 \times \{t\})$ is a circle for all t;

2. $\mathbb{S}_0 = \{p_1, p_2\} \subseteq Y$ is a 0–sphere and $\mathcal{B}_i$ is an open ball in Y containing p, such that $Y(\mathbb{S}_0)$ is formed by gluing a copy of $S^2 \times I$ to $Y \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$.

3. B_i is a closed ball on Σ which doesn't intersect any $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w}$ and which contains the point $p_i \in \Sigma$ and which contains $\mathcal{B}_i \cap \Sigma$, and such that $B_1 \cap B_2 = \varnothing$;

4. the points p_i' are on S^2 and p_1' is in the same region as w' and p_2' is in the same region as w.

5. B_i' is a closed ball on S^2 containing p_i' which doesn't intersect $\alpha_0 \cup \beta_0 \cup \{w, w'\}$, and $B_1' \cap B_2' = \varnothing$;

6. J_s is an almost complex structure on $\Sigma \times I \times \mathbb{R}$ which is split on B_1 and B_2;

7. J_s' is an almost complex structure on $S^2 \times I \times \mathbb{R}$ which is split on B_1' and B_2';

8. r_1, r_2 are real numbers such that $0 < r_1, r_2 < 1$ and such that $\mathcal{B}_i \cap \Sigma$ is contained in $\mathrm{int}(r_i \cdot B_i)$;

9. r_1', r_2' are two real numbers such that $0 < r_1', r_2' < 1$;

10. $\iota : \overline{S^2 \setminus (r_1' \cdot B_1' \cup r_2' \cdot B_2')} \hookrightarrow \mathcal{A}$ is an embedding;

11. if $\widetilde{A}_i$ denotes the annulus inside of $(\Sigma \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)) \cup \widetilde{A}$ with boundary equal to $\partial B_i \cup \partial B_i'$, then

$$\phi_i : \widetilde{A}_i \to S^1 \times [-a_i, 1 + b_i]$$

is a diffeomorphism which restricts to a conformal map of $B_i \setminus (r_i \cdot B_i)$ to $S^1 \times [-a_i, 0]$ and a conformal map of $B_i' \setminus (r_i' \cdot B_i')$ to $S^1 \times [1, 1 + b_i]$.

As in the case of the free stabilization maps, gluing data $\mathcal{J}$ determines a diagram $\mathcal{H}_{\mathcal{J}}$ as well as an almost complex structure $\mathcal{J}(T_1, T_2)$ for $T_i > 1$ which have necks length T_1 and T_2.

2.7.3 Counting holomorphic disks after surgery on $\mathbb{S}_0$

In this section, we prove most of the holomorphic curve components of the proof of naturality of the 1– and 3–handle maps. Formally the 1– and 3–handle maps are nearly identical to the free stabilization maps. This is essentially because the free stabilization maps are secretly a 0– or 4–handle map composed with a 1– or 3–handle map, a fact which we will uncover later when we consider subdivision invariance of the graph cobordism maps.

The strategy will be very similar to our analysis of the free stabilization map, the only real difference that we are counting disks which satisfy a matching condition at two points, instead of one. Let $(S^2, \alpha_0, \beta_0, w, w')$ be the diagram in Figure 2.5. Let p_1 be a point in the same region as w and let p_2 be a point in the same region as w'. Given a homotopy class of disks $\phi \in \pi_2(a, b)$ for $a, b \in \{\theta^+, \theta^-\}$ and a pair of divisors $\mathbf{d}_1 \in \mathrm{Sym}^{k_1}(\mathbb{D})$ and $\mathbf{d}_2 \in \mathrm{Sym}^{k_2}(\mathbb{D})$ with $\mu(\phi) = 2k_1 + 2k_2$, we define

$$\mathcal{M}(\phi, \mathbf{d}_1, \mathbf{d}_2) = \{u \in \mathcal{M}(\phi) : \rho^{p_1}(u) = \mathbf{d}_1, \rho^{p_2}(u) = \mathbf{d}_2\}.$$

The following lemma is very similar to [OS08, Lem. 6.4.]:

Lemma 2.7.4. *Let $(S^2, \alpha_0, \beta_0, w, w')$ and p_1 and p_2 be points in the same regions as w and w' respectively. Let $\mathbf{d}_1 \in \mathrm{Sym}^{k_1}(\mathbb{D})$ and $\mathbf{d}_2 \in \mathrm{Sym}^{k_2}(\mathbb{D})$ be two divisors. Write $k = k_1 + k_2$. Then*

$$\sum_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi) = 2k}} \#\mathcal{M}(\phi, \mathbf{d}_1, \mathbf{d}_2) \equiv 1 (\mathrm{mod}\, 2).$$

Proof. We will first show that the quantity

$$M(a, \mathbf{d}_1, \mathbf{d}_2) \stackrel{\text{def}}{=} \sum_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi)=2k}} \#\mathcal{M}(\phi, \mathbf{d}_1, \mathbf{d}_2)$$

is independent modulo 2 of $\mathbf{d}_1$ and $\mathbf{d}_2$. This follows from the following Gromov compactness argument. Suppose that $(\mathbf{d}_1^0, \mathbf{d}_2^0)$ and $(\mathbf{d}_1^1, \mathbf{d}_2^1)$ are two pairs of divisors and that $\mathbf{d}_i^t$ is a path from $\mathbf{d}_i^0$ to $\mathbf{d}_i^1$. Consider the 1–dimensional moduli space

$$\mathcal{M}(a, \mathbf{d}_1, \mathbf{d}_2) = \bigcup_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi)=2k}} \mathcal{M}(\phi, \mathbf{d}_1^t, \mathbf{d}_2^t).$$

There are four kinds of ends of this space. Since $\mathcal{M}(a, \mathbf{d}_1, \mathbf{d}_2)$ can be compactified to a compact 1–manifold, the total number of ends is zero. The first two types correspond to setting $t = 0$ or $t = 1$, i.e. we have one end of the form

$$\bigcup_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi)=2k}} \mathcal{M}(\phi, \mathbf{d}_1^0, \mathbf{d}_2^0)$$

and another of the form

$$\bigcup_{\substack{\phi \in \pi_2(a,a) \\ \mu(\phi)=2k}} \mathcal{M}(\phi, \mathbf{d}_1^1, \mathbf{d}_2^1).$$

The other two ends correspond to a degenerations of a collections of homologous holomorphic strips into a broken holomorphic strip, i.e. collections of holomorphic strips, boundary degenerations, and closed surfaces. Since one of the components must satisfy a matching condition with a pair of divisors $(\mathbf{d}_1^t, \mathbf{d}_2^t)$ for some t, using the Maslov index calculation in Lemma 2.5.3, we can constrain the types of degenerations which can occur, and we are left with the following two types of ends:

$$\bigcup_{\substack{\phi_1 \in \pi_2(a,b) \\ \phi_2 \in \pi_2(b,a) \\ t \in I \\ \mu(\phi_1)=1 \\ n_p(\phi_1)=n_{p'}(\phi_1)=0 \\ \mu(\phi_2)=2k-1}} \widehat{\mathcal{M}}(\phi_1) \times \mathcal{M}(\phi_2, \mathbf{d}_1^t, \mathbf{d}_2^t)$$

and

$$\bigcup_{\substack{\phi_1\in\pi_2(a,b)\\ \phi_2\in\pi_2(b,a)\\ t\in I\\ \mu(\phi_1)=2k-1\\ n_{p_1}(\phi_2)=n_{p_2}(\phi_2)=0\\ \mu(\phi_2)=1}} \mathcal{M}(\phi_1,\mathbf{d}_1^t,\mathbf{d}_2^t)\times\widehat{\mathcal{M}}(\phi_2).$$

Each of the above two spaces have an even number of ends because the differential on $\widehat{CF}(S^2,\alpha_0,\beta_0,w,w')$ vanishes. Hence modulo 2, the total number of ends is

$$M(a,\mathbf{d}_1^0,\mathbf{d}_2^0)-M(a,\mathbf{d}_1^1,\mathbf{d}_2^2)$$

and hence the number $M(a,\mathbf{d}_1,\mathbf{d}_2)$ is independent of the divisors $\mathbf{d}_i$.

To show that $M(a\mathbf{d}_1,\mathbf{d}_2)\equiv 1(\mathrm{mod}\,2)$, pick divisors $\mathbf{d}_1^T$ and $\mathbf{d}_2^T$ such that $\mathbf{d}_1^T\cup\mathbf{d}_2^T$ is a collection with vertical spacing at least T between each pair of points and as $T\to\infty$ the points in $\mathbf{d}_1^T\cup\mathbf{d}_2^T$ all approach $\{0\}\times\mathbb{R}\subseteq\mathbb{D}$. As we let $T\to\infty$, one argues as in [OS08, Lem. 6.4] that for sufficiently large T, any such strip is equal to k_1+k_2 boundary degenerations glued to the constant strip, and hence $M(a,\mathbf{d}_1^T,\mathbf{d}_2^T)$ is 1 for sufficiently large T. □

Lemma 2.7.5. *Suppose that $\mathbb{S}_0=\{p_1,p_2\}\subseteq Y$ is an embedded 0–sphere and that $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ is a strongly $\mathfrak{s}$–admissible diagram for $(Y,\mathbf{w})$ which is subordinate to $\mathbb{S}_0$. Suppose that $\mathcal{J}$ is gluing data with almost complex structure J_s on $\mathcal{H}$. For sufficiently large T and T' we have an identification of differentials*

$$\partial_{\mathcal{J}(T,T')}=\begin{pmatrix}\partial_{J_s} & *\\ 0 & \partial_{J_s}\end{pmatrix}.$$

Proof. The proof is very similar to the proof of Lemma 2.5.5. Write

$$\partial_{\mathcal{J}(T,T')}=\begin{pmatrix}A & B\\ C & D\end{pmatrix}.$$

We first compute A and D. Suppose that $\phi=\phi_1\#\phi_2$ is a Maslov index one homology disk and $\phi\in\pi_2(\mathbf{x}\times\theta^+,\mathbf{y}\times\theta^+)$ or $\phi\in\pi_2(\mathbf{x}\times\theta^-,\mathbf{y}\times\theta^-)$, which admits holomorphic representatives $u_{T,T'}$ for T and T' which are both arbitrarily large. Then by extracting a

convergent subsequence, we know that ϕ_1 and ϕ_2 each have representatives U_1 and U_2 as broken holomorphic strips. Now using the Maslov index computation in Lemma 2.5.3 as well as the combinatorial formula for the Maslov index, we see that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 - 2m_2 = \mu(\phi_1).$$

We conclude that $\mu(\phi_1) = 1$. In particular, the broken holomorphic strip U_1 can have at most one nonconstant component, u_1. There must be a component u_2 of U_2 satisfying the matching condition

$$\rho^{p_1}(u_1) = \rho^{p_1'}(u_2) \qquad \text{and} \qquad \rho^{p_2}(u_1) = \rho^{p_2'}(u_2).$$

We now claim that the remaining components of U_2 must be constant. To see this, let u_2 represent the homology class ϕ_2'. Since $\phi_1 \# \phi_2'$ has a holomorphic representative $u_1 \# u_2$ given by gluing, we know that $\mu(\phi_1 \# \phi_2') \geq 1$. From the matching condition, we know that $n_p(\phi_2 - \phi_2') = n_{p'}(\phi_2 - \phi_2') = 0$. Hence an easy Maslov index computation shows that

$$1 = \mu(\phi) = \mu(\phi_1 \# \phi_2') + \mu(0 \# (\phi_2 - \phi_2')) = \mu(\phi_1 \# \phi_2') + \mu(\phi_2 - \phi_2'),$$

from which it follows that $\mu(\phi_2 - \phi_2') = 0$, and hence $\phi_2 - \phi_2'$ has no nonconstant components since it admits a representative as a broken holomorphic curve. Hence we conclude that if both T and T' are sufficiently large, then any holomorphic strip u representing $\phi_1 \# \phi_2$ is of the form $u_1 \# u_2$ where u_1 and u_2 are preglued holomorphic curves.

Conversely, given a Maslov index 1 homology class ϕ_1 on $\mathcal{H}$ which admits a holomorphic representative u_1, by Lemma 2.7.4, we know that among homology classes $\phi_2 \in \pi_2(\theta^+, \theta^+)$, there is a one (modulo 2) holomorphic representative u_2 such that

$$\rho^{p_1}(u_1) = \rho^{p_1'}(u_2) \qquad \text{and} \qquad \rho^{p_2}(u_1) = \rho^{p_2'}(u_2).$$

The same statement holds for $\phi_2 \in \pi_2(\theta^-, \theta^-)$. Combining this with the previous paragraph yields an identification $A = D = \partial_{J_s}$.

We now wish to show that $B = 0$. Suppose that $\phi = \phi_1 \# \phi_2 \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$ is a Maslov index one homology disk. Note that by Lemma 2.5.3, we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) - 2m_1 - 2m_2 = \mu(\phi_1) + 1.$$

Hence $\mu(\phi_1) = 0$. If ϕ has representatives for T and T' both arbitrarily large, then ϕ_1 and ϕ_2 have broken holomorphic representatives U_1 and U_2. Since $\mu(\phi_1) = 0$, we conclude that U_1 consists only of constant components. Hence

$$n_{p_1}(\phi_1) = n_{p_2}(\phi_1) = n_{p_1'}(\phi_2) = n_{p_2'}(\phi_2) = 0.$$

By direct examination, ϕ_2 is limited to one of two homology disks. Each has a unique holomorphic representative, which cancel modulo 2. Hence $C = 0$. □

Similar to the free stabilization maps, we also need to show that the relative homology maps commute with the 1– and 3–handle maps.

Lemma 2.7.6. *Suppose that λ is a path in Y and $\mathbb{S}$ is a framed, embedded 0– or 2–sphere. Suppose that $\overline{\mathfrak{s}}$ is a* Spin^c *structure on $W(Y, \mathbb{S})$. Assume that λ is an embedded path in Y which doesn't intersect $\mathbb{S}$. Let $\overline{\lambda}$ be the induced path in $Y(\mathbb{S})$. The following relation holds:*

$$A_{\overline{\lambda}} \circ F^{\circ}_{Y,\mathbb{S},\overline{\mathfrak{s}}} = F^{\circ}_{Y,\mathbb{S},\overline{\mathfrak{s}}} \circ A_{\lambda}.$$

Proof. Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ be a diagram for $(Y, \mathbf{w})$ which is subordinate to the 1–handle data $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ and suppose that $\mathcal{J}$ is gluing data for $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$. Let $\mathcal{H}_{\mathcal{J}}$ denote the resulting diagram for $Y(\mathbb{S}_0)$. Suppose that λ is a curve between two basepoints in $\mathbf{w}$ which doesn't intersect the points $\mathbb{S}_0$, and let $\overline{\lambda}$ denote the induced curve in $Y(\mathbb{S}_0)$. It is sufficient to show that

$$A_{\overline{\lambda}} = \begin{pmatrix} A_{\lambda} & * \\ 0 & A_{\lambda} \end{pmatrix},$$

from which the main claim follows. This follows similar to the proof of Lemma 2.7.5, but with an additional factor of $a(\lambda, \phi)$ on the counts of holomorphic disks. □

Lemma 2.7.7. *Suppose $\mathbb{S} \subseteq Y$ is a framed 0– or 2–sphere and $\mathbf{w}$ is a collection of basepoints for Y with $\mathbf{w} \cap \mathbb{S} = \varnothing$. If $w \in Y \setminus (\mathbf{w} \cup \mathbb{S})$, then*

$$S^{\pm}_w \circ F^{\circ}_{Y,\mathbb{S},\mathfrak{s}} = F^{\circ}_{Y,\mathbb{S},\mathfrak{s}} \circ S^{\pm}_w$$

for a $\mathfrak{s} \in$ Spin$^c(W(Y, \mathbb{S}))$.

Proof. This is follows immediately from the formulas for $S^{\pm}_w$ and $F^{\circ}_{Y,\mathbb{S},\mathfrak{s}}$, so long as an almost complex structure can be picked which puts the maps in their standard forms. It is not hard to argue that this can be done (cf. the subclaim at the end of Theorem 2.9.2). □

2.7.4 Counting holomorphic triangles after surgery on $\mathbb{S}_0$

As we did in Section 2.5.4 for the free stabilization maps, we would like to prove an analogous stabilization result for triangle maps after adding a 1–handle. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple which is strongly admissible for a Spin^c structure $\mathfrak{s}$ and which is subordinate to the 1–handle data $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ with $\mathbb{S}_0 = \{p_1, p_2\}$. Let $\mathcal{J}$ be gluing data for the triple $\mathcal{T}$. We form the 1–handle stabilized Heegaard triple $\overline{\mathcal{T}} = \{\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}}, \mathbf{w})$ where $\overline{\Sigma}$ is the surface obtained by removing $\Sigma \cap (\mathcal{B}_1 \cup \mathcal{B}_2)$, and connecting the resulting boundary components with an annulus as specified by $\mathcal{J}$. Also we define

$$\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\alpha_0\}$$

and similarly for $\overline{\boldsymbol{\beta}}$ and $\overline{\boldsymbol{\gamma}}$, where α_0, β_0 and γ_0 are homotopically nontrivial simple closed curves in the annulus which are all pairwise isotopic to one another, and such that

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

analogous to Figure 2.7.

In light of Lemma 2.3.11, we have

$$\mathfrak{s}_{\mathbf{w}}(\psi_1 \# \psi_2) = \mathfrak{s}_{\mathbf{w}}(\psi_1' \# \psi_2') \qquad \text{iff} \qquad \mathfrak{s}_{\mathbf{w}}(\psi_1) = \mathfrak{s}_{\mathbf{w}}(\psi_1'),$$

if ψ_1 is a triangle on $\mathbb{T}$ and ψ_2 is a triangle on $(S^2, \alpha_0, \beta_0, \gamma_0, w, w')$. In particular a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(X_{\alpha\beta\gamma})$ uniquely determines a $\overline{\mathfrak{s}} \in \text{Spin}^c(X_{\overline{\alpha}\overline{\beta}\overline{\gamma}})$ such that $\mathfrak{s}_{\mathbf{w}}(\psi_1 \# \psi_2) = \overline{\mathfrak{s}}$ iff $\mathfrak{s}_{\mathbf{w}}(\psi_1) = \mathfrak{s}$.

We have the following stabilization theorem, analogous to Theorem 2.5.17.

Theorem 2.7.8. *For sufficiently large T, the stabilized triangle maps $\overline{\mathcal{T}}$ take the form*

$$F^{-}_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x}\times\cdot)\otimes(\mathbf{y}\times y^{+})) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix},$$

where the matrix on the right is an expansion of the map on the left in the variable $\cdot$, which can be either x^{+} or x^{-}. The maps $F^{\circ}_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}$ are computed with respect to the almost complex structure $\mathcal{J}(T)$, and $F^{\circ}_{\mathcal{T},\mathfrak{s}}$ is computed with the almost complex structure J_s on Σ.

Similarly, for sufficiently large T, we have

$$F^{-}_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x}\times x^{+})\otimes(\mathbf{y}\times\cdot)) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}\otimes\mathbf{y}) \end{pmatrix}.$$

As with the analysis of the free stabilization maps, to prove Theorem 2.7.8, we need to perform a count of holomorphic triangles which satisfy a certain matching condition. Suppose that $\psi\in\pi_2(x,y,z)$ is a homology triangle on $(S^2,\alpha_0,\beta_0,\gamma_0,w,w')$, with $x\in\alpha_0\cap\beta_0, y\in\beta_0\cap\gamma_0$ and $z\in\alpha_0\cap\gamma_0$. Suppose $\mathbf{d}_1\in\mathrm{Sym}^{k_1}(\mathbb{D})$ and $\mathbf{d}_2\in\mathrm{Sym}^{k_2}(\mathbb{D})$. We then define

$$\mathcal{M}(\psi,\mathbf{d}_1,\mathbf{d}_2)=\{u\in\mathcal{M}(\psi):\rho^{p_1'}(u)=\mathbf{d}_1,\rho^{p_2'}(u)=\mathbf{d}_2\},$$

and

$$M_{(x,y,z)}(\mathbf{d}_1,\mathbf{d}_2)=\sum_{\substack{\psi\in\pi_2(x,y,z)\\ \mu(\psi)=2k_1+2k_2}}\#\mathcal{M}(\psi,\mathbf{d}_1,\mathbf{d}_2).$$

Lemma 2.7.9. *If (x,y,z) is one of the triples $(x^{+},y^{+},z^{+}),(x^{-},y^{+},z^{-})$ or (x^{+},y^{-},z^{-}), then for generic $\mathbf{d}_1$ and $\mathbf{d}_2$ we have*

$$M_{(x,y,z)}(\mathbf{d}_1,\mathbf{d}_2)\equiv 1(\mathrm{mod}\,2).$$

Proof. The proof is identical to the proof of Lemma 2.5.19. □

Proof of Theorem 2.7.8. With Lemma 2.7.9, the proof is now follows essentially identically to the proof of 2.5.17. □

2.7.5 Dependence on gluing data

Similar to the free stabilization maps, we have the following:

Lemma 2.7.10. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a strongly $\mathfrak{s}$–admissible diagram for $(Y, \mathbf{w})$ and let $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ be 1–handle data in Y where $\mathbb{S}_0 = \{p_1, p_2\}$ is a framed 0–sphere in Y. Suppose that $\mathcal{J}$ is gluing data for $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$. Write J_s for the almost complex structure on $\Sigma \times I \times \mathbb{R}$ specified by $\mathcal{J}$. Then there is an N such that if $T, T', T'', T''' > N$, and all almost complex structures achieve transversality, then the map*

$$\Phi_{\mathcal{J}(T,T')\to\mathcal{J}(T'',T''')} : CF^\circ_{\mathcal{J}(T,T')}(\mathcal{H}_\mathcal{J}, \mathfrak{s}') \to CF^\circ_{\mathcal{J}(T'',T''')}(\mathcal{H}_\mathcal{J}, \mathfrak{s}')$$

is homotopy equivalent to

$$\Phi_{\mathcal{J}(T,T')\to\mathcal{J}(T'',T''')} \simeq \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Proof. The proof is identical to the proof of Lemma 2.5.8, though we use Lemma 2.7.4 to find preglued disks. □

Definition 2.7.11. We say that N is **sufficiently large** for gluing data $\mathcal{J}$, if $\Phi_{\mathcal{J}(T,T')\to\mathcal{J}(T'',T''')}$ is of the form in the previous lemma for all $T, T', T'', T''' \geq N$. We say T is **large enough to compute** the 1– or 3–handle maps if $T > N$ for some N which is sufficiently large.

Lemma 2.7.12. *Suppose that $\mathcal{J}_1$ and $\mathcal{J}_2$ are two choices of gluing data for $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$ and that $\mathcal{H}$ is strongly $\mathfrak{s}$–admissible and subordinate to $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$. Then if T_i and T_i' are large enough to the compute the 1– and 3–handle maps for $\mathcal{J}_i$ using the above definition, then*

$$\Phi_{\mathcal{J}_1(T_1,T_1')\to\mathcal{J}_2(T_2,T_2')} \simeq \begin{pmatrix} \Phi_{J_{s,1}\to J_{s,2}} & * \\ 0 & \Phi_{J_{s,1}\to J_{s,2}} \end{pmatrix}$$

where $J_{s,i}$ is the almost complex structure on $\Sigma \times I \times \mathbb{R}$ associated to $\mathcal{J}_i$.

Proof. The effect of changing all of the pieces of gluing data except for the annulus $\mathcal{A}$ filling in $\Sigma \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$ follows nearly identically to the analogous result for the free stabilization

maps. One simply verifies that the proofs of Lemmas 2.5.10, 2.5.11, 2.5.12, 2.5.13, 2.5.14, 2.5.15, and 2.5.16 adapt essentially without change from the free stabilization setting to the 1– and 3–handle setting. Then one verifies that the above moves are sufficient to move between $\mathcal{J}_1$ to $\mathcal{J}_2$ as long as we use the same annulus, $\mathcal{A}$, for both $\mathcal{J}_1$ and $\mathcal{J}_2$ is the same.

The annulus $\mathcal{A}$ is specified up to isotopy by the following argument: inside of $Y(\mathbb{S}_0)$ there is a subset identified with $S^2 \times I$. By assumption $\mathcal{A} \cap (S^2 \times \{t\})$ is a single circle. Such an embedding of $S^1 \times I$ into $S^2 \times I$ always extends to an embedding of $\mathbb{D}^2 \times I$ into $S^2 \times I$, and any two embeddings are isotopic to each other via an isotopy fixing $S^2 \times \partial I$ since S^2 is simply connected.

Hence given two annuli $\mathcal{A}$ and $\widetilde{\mathcal{A}}$ which fill in $\Sigma \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$ in $Y(\mathbb{S}_0)$, there is an isotopy ψ_t of $Y(\mathbb{S}_0)$ which is fixed outside of $Y \setminus (\mathcal{B}_1 \cup \mathcal{B}_2)$ which maps $\mathcal{A}$ to $\widetilde{\mathcal{A}}$. If $\mathcal{J}$ is gluing data, then we can move to gluing data $\widetilde{\mathcal{J}} = (\psi_1)_* \mathcal{J}$, which has annulus $\widetilde{\mathcal{A}}$ and the change of almost complex structures map $\Phi_{\mathcal{J}(T,T') \to \widetilde{\mathcal{J}}(T,T')}$ is tautologically the isotopy map $(\psi_1)_*$, which is the identity matrix once we identity the two complexes. □

2.7.6 Invariance of 1– and 3–handle maps

In this section we prove invariance of the 1– and 3–handle maps.

Theorem 2.7.13. *Suppose that $(W, \mathbf{w})$ is a based 3–manifold with 1–handle data $(\mathbb{S}_0, \mathcal{B}_1, \mathcal{B}_2)$. Suppose that $\overline{\mathfrak{s}} \in \mathrm{Spin}^c(W(\mathbb{S}_0))$, $\mathfrak{s} = \overline{\mathfrak{s}}|_Y$, and suppose that $\mathcal{H}_1$ and $\mathcal{H}_2$ are two strongly $\mathfrak{s}$–admissible diagrams for Y. Suppose that $\mathcal{J}_1$ and $\mathcal{J}_2$ are choices of gluing data for $\mathcal{H}_1$ and $\mathcal{H}_2$ respectively. Then the following diagram and its analog for the 3–handle maps commutes up to chain homotopy:*

$$\begin{array}{ccc} CF^\circ_{J_{s,1}}(\mathcal{H}_1) & \xrightarrow{F_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J}_1,\overline{\mathfrak{s}}}} & CF^\circ_{\mathcal{J}_1(T_1)}((\mathcal{H}_1)_{\mathcal{J}_1}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{(\mathcal{H}_1)_{\mathcal{J}_1} \to (\mathcal{H}_2)_{\mathcal{J}_2}}} \\ CF^\circ_{J_{s,2}}(\mathcal{H}_2) & \xrightarrow{F_{(\mathbb{S}_0,\mathcal{B}_1,\mathcal{B}_2),\mathcal{J}_2,\overline{\mathfrak{s}}}} & CF^\circ_{\mathcal{J}_2(T_2)}((\mathcal{H}_2)_{\mathcal{J}_2}). \end{array}$$

Proof. The proof of the analogous statement for the free stabilization maps (Theorem 2.5.22)

holds in this context essentially without change. Invariance from gluing data follows from Lemma 2.7.12. Invariance from $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–handleslides and isotopies follows from Theorem 2.7.8. Invariance from isotopies of Y which fix the basepoints $\mathbf{w}$ follows identically to Theorem 2.5.22. □

2.8 2–handle maps

In this section we construct and analyze the 2–handle maps associated to an embedded framed link $\mathbb{L} \subseteq Y$. These maps were essentially defined in their current form by Ozsváth and Szabó in [OS06], though not in the cylindrical setting, and not in the case of multiple basepoints. Our contribution to these maps is essentially limited to gluing lemmas which we use to show that the 2–handle maps cancel the 1– and 3–handle maps in the appropriate sense, as well as showing that the 2–handle maps commute with the free stabilization, relative homology, and 1– and 3–handle maps under appropriate circumstances. In addition to the original construction in [OS06, Sec. 4], we will rely heavily on work from [Juh16, Sec. 8], which builds the theory for 3–manifolds with multiple basepoints.

2.8.1 Maps associated to a framed link

Definition 2.8.1. A **framed link** $\mathbb{L}$ in a 3–manifold Y is a collection of n pairwise disjoint embedded circles $K_1, \dots, K_n \subseteq Y$ together with a **framing**, i.e. a choice of homology classes $\ell_i \in H_1(\partial N(K_i))$ satisfying $m_i \cdot \ell_i = 1$ where m_i is the meridian of K_i.

Definition 2.8.2. A **bouquet** $B(\mathbb{L})$ for a link $\mathbb{L}$ is an embedded 1–complex in Y which is the union of $\mathbb{L}$ with a collection of arcs $a_1, \dots, a_n$ where each a_i is an arc from a component K_i of $\mathbb{L}$ to a basepoint in $\mathbf{w}$. The arcs a_i do not all need to connect to the same basepoint. We define F_i to be the punctured torus $\partial N(K_i) \setminus N(a_i)$.

Definition 2.8.3. We say a Heegaard triple

$$\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}) = (\Sigma, \{\alpha_1, \dots, \alpha_{g+\ell-1}\}, \{\beta_1, \dots, \beta_{g+\ell-1}\}, \{\gamma_1, \dots, \gamma_{g+\ell-1}\}, \mathbf{w}),$$

is **subordinate to the bouquet** $B(\mathbb{L})$ if

1. $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is an embedded Heegaard diagram for $(Y, \mathbf{w})$;

2. $(\Sigma, \{\alpha_1, \dots, \alpha_{g+\ell-1}\}, \{\beta_{n+1}, \dots, \beta_{g+\ell-1}\}, \mathbf{w})$ is a diagram for $(Y \setminus N(B(\mathbb{L}))$. By this we mean the following: write U_α and U_β for the closures of the components of $Y \setminus \Sigma$ and let $\mathbf{w}^\alpha$ denote small push-offs of the basepoints $\mathbf{w}$ into U_α. Then the $\alpha_1, \dots, \alpha_{g+\ell+1}$ curves bound compressing disks and after compressing Σ along these disks in U_α we get a surface which is isotopic to the spheres $\partial N(\mathbf{w}^\alpha))$. Let $(B(\mathbb{L}) \cup \mathbf{w})^\beta$ denote a pushoff of $B(\mathbb{L}) \cup \mathbf{w}$ into U_β. The curves $\beta_{n+1}, \dots, \beta_{g+\ell-1}$ bound compressing disks in $U_\beta \setminus N(B(\mathbb{L}))$, and after compressing Σ along these disks, we get a surface which is isotopic to $\partial N((B(\mathbb{L}) \cup \mathbf{w})^\beta)$;

3. the curves $\gamma_{n+1}, \dots, \gamma_{g+\ell-1}$ are small isotopies of the curves $\beta_{n+1}, \dots, \beta_{g+\ell-1}$;

4. after compressing Σ along the curves $\beta_{n+1}, \dots, \beta_{g+\ell-1}$, the induced curves β_i and γ_i (for $i = 1, \dots, n$) lie in the punctured torus F_i;

5. for $i = 1, \dots, n$ the curve β_i represents a meridian of K_i and $\beta_i \cap \gamma_j = \varnothing$ unless $i = j$, in which case β_i and γ_i meet at a single transverse intersection point;

6. for $i = 1, \dots, n$ the homology class of γ_i corresponds to the framing ℓ_i under an identification of
$$H_1(\partial N(\cup_i K_i)) \cong H_1(\partial N(B(\mathbb{L}))).$$

Definition 2.8.4. If $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple subordinate to $B(\mathbb{L}) \subseteq Y$, and $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$ is another Heegaard triple subordinate to $B(\mathbb{L})$, we say that $\mathcal{T}'$ is the **stabilization** of $\mathcal{T}$ if Σ' is the internal connected sum of Σ and a genus one surface $E \subseteq Y \setminus N(B(\mathbb{L}))$ with $E \cap \Sigma = \varnothing$ and $\boldsymbol{\tau}' = \boldsymbol{\tau} \cup \{\tau_0\}$ for $\tau \in \{\alpha, \beta, \gamma\}$ where α_0, β_0 and γ_0 are curves on E such that $|\alpha_0 \cap \beta_0| = 1$ and β_0 and γ_0 are small isotopies of each other with $|\beta_0 \cap \gamma_0| = 2$.

The following is an analog of [OS06, Prop. 4.3] for the case of multiple basepoints:

Lemma 2.8.5. *Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple subordinate to the bouquet $B(\mathbb{L})$ in Y. Then $X_{\alpha\beta\gamma}$ is a cobordism from Y to*

$$Y(\mathbb{L}) \sqcup (\sqcup_i \#^{k_i}(S^1 \times S^2))$$

where k_i is the number of $\boldsymbol{\alpha}$ curves in the ith component of Σ. The monodiagram $(\Sigma, \boldsymbol{\beta})$ defines a cobordism X_β from $U_\beta \cup_\Sigma (-U_\beta)$ to $\varnothing$. Let $W_\beta(\mathbb{L})$ denote the cobordism obtained from W_β by attaching 4–dimensional 2–handles to $\mathbb{L}$ along U_β. The manifold $W_\beta(\mathbb{L})$ is a cobordism from $Y_{\beta\gamma}$ to $\varnothing$. We can glue $X_{\alpha\beta\gamma}$ and $W_\beta(\mathbb{L})$ along $Y_{\beta\gamma}$ to obtain a cobordism equivalent to $W(\mathbb{L})$, and this equivalence is well-defined up to isotopy fixing $\mathbf{w} \times I \subseteq W(\mathbb{L})$.

Proof. This follows from [Juh16, Prop. 6.6]. □

Lemma 2.8.6. *Suppose $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is subordinate to $B(\mathbb{L})$, and $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$. Under the inclusion $Y_{\beta\gamma} \subseteq X_{\alpha\beta\gamma} \subseteq W(\mathbb{L})$, the Spin^c structure $\mathfrak{s}|_{Y_{\beta\gamma}}$ is the torsion Spin^c structure.*

Proof. This follows as in [OS06, Prop. 4.3]. The cobordism $W(\mathbb{L})$ is diffeomorphic to $X_{\boldsymbol{\alpha\beta\gamma}}$ with 3– and 4–handles attached along $Y_{\beta\gamma}$. The only Spin^c structure extending over these 3– and 4–handles is the torsion Spin^c structure. □

We can now define the two handle maps. Given a framed link $\mathbb{L} \subseteq Y$, a bouquet $B(\mathbb{L})$, an $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$, and a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ which is subordinate to $B(\mathbb{L})$ and which is strongly $\mathfrak{s}|_{X_{\alpha\beta\gamma}}$–admissible, we can define

$$F^\circ_{Y,\mathbb{L},\mathcal{T},\mathfrak{s}} = HF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}|_Y) \to HF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{L})})$$

by the formula

$$F^\circ_{Y,\mathbb{L},\mathcal{T},\mathfrak{s}}(\mathbf{x}) = F^\circ_{\mathcal{T},\mathfrak{s}|_{X_{\alpha\beta\gamma}}}(\mathbf{x} \otimes \Theta^+),$$

where $\Theta^+ \in CF^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathfrak{s}_0)$ is the top degree generator. We now show invariance of this map from the Heegaard triple $\mathcal{T}$, and then from the bouquet $B(\mathbb{L})$. We will address independence of this map from the bouquet and the intermediate Heegaard diagram, but first we must prove several versions of $(1,2)$–stabilization invariance of the triangle maps.

2.8.2 (1,2)–stabilization and triangle maps

We now prove $(1,2)$–stabilization invariance of the triangle maps. In principle this is "standard" since it is needed to prove many foundational results, for example the invariance of the 3–manifold invariants, but to the authors knowledge a proof of this fact in the cylindrical context has not appeared in the literature. The following argument proceeds along the same lines as Theorem 2.5.17, first performing a count of triangles on a simple diagram with a certain matching condition, and then performing a stretching argument to show that all triangles on the stabilized diagram limit to preglued triangles as we let the neck length approach ∞. Let $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ be the diagram shown in Figure 2.14. There is a new feature of this argument, which is not present in the neck stretching arguments we performed for free stabilization and for the 1– and 3–handle maps. To avoid closed holomorphic surfaces from appearing, which do not necessarily achieve transversality, we need to consider somewhat more harshly perturbed almost complex structures, similar to those satisfying $(J5')$ in [Lip06].

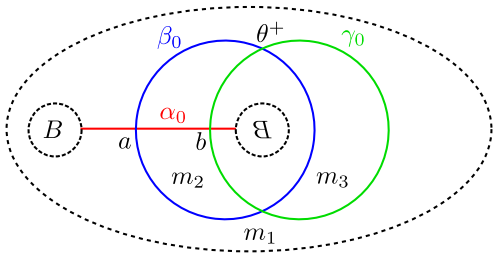

Figure 2.14: The Heegaard Triple $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ used to perform $(1,2)$–stabilization of triangle maps.

Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is strongly admissible and let $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ be the triple shown in Figure 2.14. Given an $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma}})$, then in light of Lemma 2.3.11, $\mathfrak{s}$ determines a unique $\overline{\mathfrak{s}}$ in $\mathrm{Spin}^c(X_{\overline{\boldsymbol{\alpha}}\overline{\boldsymbol{\beta}}\overline{\boldsymbol{\gamma}}})$ for the $(1,2)$–stabilized triple $\overline{\mathcal{T}} = (\Sigma \# \mathbb{T}^2, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\gamma}}, \mathbf{w})$, where $\overline{\boldsymbol{\tau}} = \boldsymbol{\tau} \cup \{\tau_0\}$ for $\tau \in \{\alpha, \beta, \gamma\}$. Using this correspondence between Spin^c structures, we have the following:

Theorem 2.8.7. *Suppose $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a strongly $\mathfrak{s}$–admissible with a distinguished*

point $q \in \Sigma \setminus (\mathbf{w} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma})$ *and* $\overline{\mathcal{T}}$ *is the Heegaard triple*

$$\overline{\mathcal{T}} = (\Sigma \# \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w}),$$

where the connected sum is taken at $p \in \mathbb{T}^2$ *and* $q \in \Sigma$. *Then* $\overline{\mathcal{T}}$ *is strongly* $\bar{\mathfrak{s}}$*–admissible and if* θ^+ *is the top generator of* $\beta_0 \cap \gamma_0$, *then for sufficiently stretched almost complex structure (which is suitably generic in a sense described below) there is an identification of the* $\mathbf{z} \times b$ *component of*

$$F^{\circ}_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}((\mathbf{x} \times a) \otimes (\mathbf{y} \times \theta^+))$$

and the $\mathbf{z}$ *component of*

$$F^{\circ}_{\mathcal{T},\mathfrak{s}}(\mathbf{x} \otimes \mathbf{y})$$

where $\{a\} = \alpha_0 \cap \beta_0$ *and* $\{b\} = \alpha_0 \cap \gamma_0$.

Suppose $\overline{\mathcal{T}}'$ *instead denotes the triple defined similarly to* $\overline{\mathcal{T}}$, *but instead with* α_0 *and* β_0 *isotopic and with* $|\alpha_0 \cap \beta_0| = 2$ *and* $|\alpha_0 \cap \gamma_0| = |\beta_0 \cap \gamma_0| = 1$. *Letting* θ^+ *denote the top degree generator of* $\alpha_0 \cap \beta_0$ *and* $\{a\} = \beta_0 \cap \gamma_0$ *and* $\{b\} = \alpha_0 \cap \gamma_0$, *then for sufficiently stretched and generic almost complex structure there is an identification of the* $\mathbf{z} \times b$ *component of*

$$F^{\circ}_{\overline{\mathcal{T}}',\mathfrak{s}}((\mathbf{x} \times \theta^+) \otimes (y \times a))$$

and the $\mathbf{z}$ *component of*

$$F^{\circ}_{\mathcal{T},\mathfrak{s}}(\mathbf{x} \times \mathbf{y}).$$

We will only prove the above claim for $\overline{\mathcal{T}}$, since the proof for $\overline{\mathcal{T}}'$ is nearly identical. We also remark that admissibility is easy to check, so we leave it to the reader (cf. Lemma 2.7.2). There is a somewhat subtle issue in that to make the correct statements, we must work with almost complex structures for which closed surfaces mapped into the interior of Δ are non generic. In [Lip06], Lipshitz considers almost complex structures on $\Sigma \times I \times \mathbb{R}$ satisfying five axioms, labeled $(J1), \ldots, (J5)$. For proving stabilization invariance of the chain complexes, he considers an alternate fifth axiom $(J5')$ which is slightly more generic, and for which closed holomorphic curves mapped into the interior of Δ are non generic. In constructing

the holomorphic triangle maps, he considers almost complex structures on $\Sigma \times \Delta$ satisfying axioms labeled $(J'1) - -(J'4)$. Following Lipshitz's approach to stabilization invariance of the chain complexes, we will consider almost complex structures satisfying a slightly more generic condition $(J'4')$, which is simply the following:

Definition 2.8.8. We say an almost complex structure J on $\Sigma \times \Delta$ satisfies condition $(J'4')$ if there is a 2–plane distribution ξ on $\Delta \times \Sigma$ such that the restriction of ω to ξ is nondegenerate, J preserves ξ, and the restriction of J to ξ is compatible with ω. Also ξ is tangent to Σ near $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma}) \times \Delta$ and near $\Sigma \times \partial\Delta$. We also assume that the cylindrical almost complex structures on the ends of $\Sigma \times \Delta$ satisfy $(J5')$.

We now proceed with a Lemma to count holomorphic triangles in the diagram $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ satisfying a matching condition. Suppose that $\psi \in \pi_2(a, x, b)$ is a homology triangle and $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$ is a divisor. We define

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

and similarly

$$\mathcal{M}_{(a,x,b)}(\mathbf{d}) = \coprod_{\substack{\psi \in \pi_2(a,x,b) \\ \mu(\psi) = 2k}} \mathcal{M}(\psi, \mathbf{d}).$$

Finally we define

$$M(\mathbf{d}) = \#\mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}).$$

Lemma 2.8.9. *If $\psi \in \pi_2(a, \theta^+, b)$ then with the multiplicities m_1, m_2, m_3 shown in Figure 2.14, we have*

$$\mu(\psi) = 2m_1 = m_2 + m_3.$$

If $\psi \in \pi_2(a, \theta^-, b)$ then

$$\mu(\psi) = 2m_1 - 1 = m_2 + m_3.$$

Proof. As usual we verify that the formulas hold for the "trivial" Maslov index zero triangle $\psi_0 \in \pi_2(a, \theta^+, b)$ and then show that the above formula respects splicing in elements of $\pi_2(s, t)$ for various s and t. □

Lemma 2.8.10. *For generic* $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$ *and for generic almost complex structures on* $\Sigma \times \Delta$ *satisfying* $(J'4')$*, we have that*

$$M(\mathbf{d}) \equiv 1 (\mathrm{mod}\, 2).$$

Proof. This proceeds very similarly to the proof of Lemma 2.5.19. We first show that $M(\mathbf{d})$ is independent of $\mathbf{d}$. To do this we pick a path $\mathbf{d}_t$ between $\mathbf{d}_0$ and $\mathbf{d}_1$. We now consider the space

$$\coprod_t \mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}_t).$$

First note that the almost complex structure $(J'4')$ prohibits closed surfaces mapped into $(\mathrm{int}\, \Delta)$ from appearing. Hence the only kinds of degenerations which can appear as ends correspond to strip breaking and boundary degenerations, as well as setting $t = 0$ or 1. The first two types of ends correspond to

$$\mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}_i)$$

for $i = 0, 1$. The other two correspond to degenerations of a family of holomorphic triangles into a broken holomorphic triangle. Using the Maslov index in Lemma 2.8.9 and the fact that the triangular component must satisfy a matching condition with a $\mathbf{d}_t$, we easily see that the only type of degenerations that can occur corresponds to the following space:

$$\coprod_{\substack{t \in I \\ \psi \in \pi_2(a,\theta^-,b) \\ \phi \in \pi_2(\theta^+,\theta^-) \\ \mu(\psi)=2k-1, \mu(\phi)=1}} \widehat{\mathcal{M}}(\phi) \times \mathcal{M}(\psi, \mathbf{d}).$$

Since the complex $\widehat{CF}(\mathbb{T}^2, \beta_0, \gamma_0, w)$ has vanishing differential (modulo 2), the above space has an even number of ends, and hence contributes nothing to the total count of ends. We thus conclude that $M(\mathbf{d}_0) - M(\mathbf{d}_1)$ is even, as we wanted.

We now wish to show that $M(\mathbf{d}) \equiv 1$ for appropriate choice of $\mathbf{d}$. To prove this, let $\mathbf{d}_T$ be a divisor with k points spaced at least T apart in Δ, such that as $T \to \infty$, the points in $\mathbf{d}_T$ each approach $v_{\alpha\beta} \subseteq \Delta$. In the cylindrical end of Δ limiting to $v_{\alpha\beta}$, we assume that

the I–component of each point in $\mathbf{d}_T$ approaches a value in $(0,1)$. A sequence of curves in $\mathcal{M}_{(a,\theta^+,b)}(\mathbf{d}_T)$ limits to a collection of k Maslov index 2 curves u_i representing classes in $\pi_2(a,a)$ each of which satisfies a matching condition $\rho^p(u_i) = d_i$ for some $d_i \in I \times \mathbb{R}$, as well as a triangular component τ with $[\tau] \in \pi_2(a,\theta^+,b)$ with $\mu([\tau]) = 0$. The triangular component τ is easily seen to be restricted to the "trivial" triangular domain of Maslov index zero, which has a unique representative. Using the proof of [Lip06, Prop. A.3], for almost complex structures on $\Sigma \times \mathbb{D}$ satisfying $(J5')$, we know that the space of holomorphic curves in $\pi_2(a,a)$ with Maslov index 2 satisfying the matching condition with d_i has a unique element. Thus by gluing, we conclude that for sufficiently large T, $M(\mathbf{d}_T) \equiv 1 (\mathrm{mod}\, 2)$, from which the result follows. □

We are now in position to prove Theorem 2.8.7.

Proof of Theorem 2.8.7. We first consider almost complex structures satisfying $(J'4')$. Suppose that $\psi = \psi_1 \# \psi_2 \in \pi_2(\mathbf{x} \times a, \mathbf{y} \times \theta^+, \mathbf{z} \times b)$ and suppose that ψ is a Maslov index zero homology triangle which has representatives v_T for arbitrarily large neck length T. Using the combinatorial Maslov index as well as Lemma 2.8.9, we have that

$$\mu(\psi) = \mu(\psi_1) + \mu(\psi_2) - 2m_1 = \mu(\psi_1).$$

Thus we conclude that $\mu(\psi_1) = 0$. Letting $T \to \infty$ we can extract a limit of v_T as a pair of broken triangles V_1 on $\mathcal{T}$ and V_2 on $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$. Since $\mu(\psi_1) = 0$, we conclude that V_1 cannot contain any flowlines, boundary degenerations, or closed surfaces, and hence must be represented by a single triangle v_1. Now in V_2 there must be a matching component, i.e. a triangular component v_2 with

$$\rho^q(v_1) = \rho^p(v_2).$$

We now claim that there are no other components of v_2. The remaining components are represented by flowlines, boundary degenerations, and closed surfaces, they must each have positive Maslov index if they are nonconstant. But since the remaining components have total Maslov index zero, they must be constant. Hence we know that for all sufficiently large T, v_T is obtained by gluing a preglued triangle (v_1, v_2).

Conversely given a holomorphic triangle v_1 on $\mathcal{T}$, according to Lemma 2.8.10 there is one (modulo 2) holomorphic triangle v_2 on $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ satisfying

$$\rho^q(v) = \rho^p(v_2).$$

Hence we know from gluing that for sufficiently large T, if $\psi_1 \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$, the total count of $\mathcal{M}(\psi_1)$ is equal to the total count of

$$\coprod_{\substack{\psi_2 \in \pi_2(a, \theta^+, b) \\ n_p(\psi_2) = n_q(\psi_1)}} \mathcal{M}(\psi_1 \# \psi_2).$$

□

We need one additional gluing lemma. Theorem 2.8.7 shows that the triangle maps are invariant under $(1,2)$–stabilization. It also will be used to show that 2–handle maps cancel 1–handle maps. We need the following lemma to show that 3–handle maps cancel 2–handle maps.

Theorem 2.8.11. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is strongly $\mathfrak{s}$–admissible and let $(\mathbb{T}^2, \alpha_0, \beta_0, \gamma_0, w)$ be the Heegaard triple similar to Figure 2.14 with α_0 and γ_0 isotopic curves with $|\alpha_0 \cap \beta_0| = |\beta_0 \cap \gamma_0| = 1$ and $|\alpha_0 \cap \gamma_0| = 2$. Denote $\overline{\mathcal{T}} = (\Sigma \# \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w})$. Suppose that $\mathfrak{s} \in \mathrm{Spin}^c(X_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma}})$ and that $\overline{\mathfrak{s}}$ denotes the corresponding Spin^c structure in $\mathrm{Spin}^c(X_{\overline{\boldsymbol{\alpha}}\overline{\boldsymbol{\beta}}\overline{\boldsymbol{\gamma}}})$. Then for sufficiently stretched and generic almost complex structure, there is an identification of the $\mathbf{z} \times \theta^-$ component of*

$$F^{\circ}_{\overline{\mathcal{T}}, \overline{\mathfrak{s}}}((\mathbf{x} \times a) \otimes (\mathbf{y} \times b))$$

and the $\mathbf{z}$ component of

$$F^{\circ}_{\mathcal{T}, \mathfrak{s}}(\mathbf{x} \otimes \mathbf{y}).$$

The proof is identical to the proof of Theorem 2.8.7, so we omit it.

2.8.3 Invariance of maps associated to a framed link

We are now in position to prove invariance of the maps associated to a framed link. We first show invariance of the map for a fixed bouquet, and then show invariance from the bouquet.

Lemma 2.8.12. *Suppose we are given a framed link $\mathbb{L} \subseteq Y$ and a bouquet $B(\mathbb{L})$, and two Heegaard triples $\mathcal{T}_i = (\Sigma_i, \boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \mathbf{w})$ for $i = 1, 2$, which are both strongly $\mathfrak{s}$–admissible for an $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$ and subordinate to $B(\mathbb{L})$. Let $\mathcal{H}_i$ denote the diagram $(\Sigma_i, \boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \mathbf{w})$ for Y and let $\overline{\mathcal{H}}_i$ denote the diagram $(\Sigma_i, \boldsymbol{\alpha}_i, \boldsymbol{\gamma}_i, \mathbf{w})$ for $Y(\mathbb{L})$. Then the following diagram commutes:*

$$\begin{array}{ccc} HF^\circ(\mathcal{H}_1, \mathfrak{s}|_Y) & \xrightarrow{F_{Y,\mathbb{L},\mathcal{T}_1,\mathfrak{s}}} & HF^\circ(\overline{\mathcal{H}}_1, \mathfrak{s}|_{Y(\mathbb{L})}) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{\overline{\mathcal{H}}_1 \to \overline{\mathcal{H}}_2}} \\ HF^\circ(\mathcal{H}_2, \mathfrak{s}|_Y) & \xrightarrow{F_{Y,\mathbb{L},\mathcal{T}_2,\mathfrak{s}}} & HF^\circ(\overline{\mathcal{H}}_2, \mathfrak{s}|_{Y(\mathbb{L})}) \end{array}$$

Before we proceed with the proof of Lemma 2.8.12, we state the following lemma:

Lemma 2.8.13. *Suppose that Y is a 3–manifold with a framed link $\mathbb{L}$ and a bouquet $B(\mathbb{L})$. Then there is a Heegaard triple subordinate to $B(\mathbb{L})$ and any two such diagrams can be connected by a sequence of the following moves:*

1. *isotopies and handleslides amongst the $\boldsymbol{\alpha}$ curves;*
2. *isotopies and handleslides amongst the $\beta_{n+1}, \dots, \beta_{g+\ell-1}$, while performing the the same move to the corresponding $\gamma_{n+1}, \dots, \gamma_{g+\ell-1}$ as well;*
3. *stabilizations and destabilizations;*
4. *for $i \in \{1, \dots, n\}$ an isotopy of β_i, or a handleslide of β_i across a β_j with $j \in \{n+1, \dots, g+\ell-1\}$;*
5. *for $i \in \{1, \dots, n\}$ an isotopy of γ_i or a handleslide of γ_i across a γ_j with $j \in \{n+1, \dots, g+\ell-1\}$*
6. *a diffeomorphism isotopic to the identity in Y relative $N(B(\mathbb{L})) \cup \mathbf{w}$.*

Proof. This follows immediately from [Juh16, Lem. 6.5], where the analogous statement is proven in the sutured setting. □

We are now in position to prove that the maps $F^\circ_{\mathbb{L},\mathcal{T},\mathfrak{s}}$, are invariant of the Heegaard diagram subordinate to B:

Proof of Lemma 2.8.12. It is sufficient to show that the maps $F^\circ_{\mathbb{L},\mathcal{T},\mathfrak{s}}$ commute with the moves listed in Lemma 2.8.13.

First note that invariance under move (6) is immediate since an isotopy which fixes $N(B(\mathbb{L}))\cup\mathbf{w}$ induces an isotopy of the Heegaard surface Σ in both Y and $Y(\mathbb{L})$ making the obvious diagram commute.

Invariance under move (3) follows immediately from Theorem 2.8.7.

Invariance under moves (1),(2),(4) and (5) follow identically to [Juh16, Thm. 6.9], the proof in the sutured context. This essentially follows from associativity of the triangle maps, as well as invariance of the top degree generator of $HF^-((S^1\times S^2)^{\#n},\mathfrak{s}_0)$. □

Given a framed link $\mathbb{L}$ and a bouquet $B(\mathbb{L})$, we let $F^\circ_{Y,\mathbb{L},\mathfrak{s},B}$ denote $F^\circ_{Y,\mathbb{L},\mathfrak{s},\mathcal{T}}$ for any triple $\mathcal{T}$ subordinate to B.

Lemma 2.8.14. *If B and B' are two different bouquets for a framed link $\mathbb{L}\subseteq Y$, then the maps $F^\circ_{\mathbb{L},\mathfrak{s},B}$ and $F^\circ_{\mathbb{L},\mathfrak{s},B'}$ agree.*

Proof. The proof is identical to the proof of the statement in the sutured context (the end of the proof of [Juh16, Thm. 6.9]). We can move from one bouquet to another by a sequence of small isotopies of edges of the bouquet, and replacing one edge at a time. If we want to replace an edge e with another edge e' connected to the same link component, if e and e' are disjoint and B denotes the bouquet with edge e and B' denotes the bouquet with edge e', then we can find two Heegaard triples $\mathcal{T}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{w})$ and $\mathcal{T}'=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta}',\boldsymbol{\gamma}',\mathbf{w})$ subordinate to B and B' respectively, such that $\boldsymbol{\beta}'$ is obtained from $\boldsymbol{\beta}$ by a sequence of handleslides and isotopies, and $\boldsymbol{\gamma}'$ is obtained from $\boldsymbol{\gamma}$ by a sequence of handleslides and isotopies. The

commutativity of the appropriate diagram then follows as in proof of invariance of the maps $F^\circ_{\mathbb{L},\mathfrak{s},\mathcal{T}}$ under moves (1),(2), (4) and (5). □

We write simply $F^\circ_{Y,\mathbb{L},\mathfrak{s}}$ for the map $F^\circ_{Y,\mathbb{L},\mathfrak{s},B}$ for any bouquet B.

Lemma 2.8.15. *If $\mathbb{L}_1$ and $\mathbb{L}_2$ are two disjoint, framed links in Y, then*

$$F^\circ_{Y(\mathbb{L}_1),\mathbb{L}_2,\mathfrak{s}_2} \circ F^\circ_{Y,\mathbb{L}_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W(\mathbb{L}_1)(\mathbb{L}_2)) \\ \mathfrak{s}|_{W(\mathbb{L}_1)}=\mathfrak{s}_1, \mathfrak{s}|_{W(\mathbb{L}_2)}=\mathfrak{s}_2}} F^\circ_{Y,\mathbb{L}_1\cup\mathbb{L}_2,\mathfrak{s}}.$$

Proof. This follows as in [OS06, Prop. 4.9] from associativity of the triangle maps [OS04e, Thm. 8.16]. □

2.9 Invariance of path cobordism maps

In this section we prove invariance of cobordism maps associated to path cobordisms. This will follow from the framework established in [Juh16] and [?JTQFT], as well as the gluing lemmas we've proven already. For our purposes, the main result we need is [Juh16, Thm. 8.9]. Juhász phrases these results in terms of cobordisms of sutured manifolds. Given a multibased 3–manifold $(Y, \mathbf{w})$, we can remove regular neighborhoods of the basepoints and add a single suture to each resulting boundary component to get a sutured manifold $Y(\mathbf{w})$. By removing a regular neighborhood of the paths in our cobordism, we get a "special cobordism" and hence can apply Juhász' result directly. A special cobordism is a sutured cobordism from (M, γ) to (M', γ'), where the Z part of the boundary is of the form $\partial M \times I$. On such a cobordism, one can put the unique tight contact structure induced by the sutures on ∂M, so the contact geometry plays no interesting role. It is important to note that if $(W, \boldsymbol{\gamma})$ is a path cobordism, and f is a Morse function on W which has $\boldsymbol{\gamma}$ as flowlines, then f naturally induces a Morse function on the sutured cobordism obtained by removing a regular neighborhood of $\boldsymbol{\gamma}$ from W.

According to [Juh16, Thm. 8.9], in order to define maps associated to special cobordisms,

it is sufficient to define handle attachment maps

$$F^\circ_{Y,\mathbb{S},\mathfrak{s}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) \to HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})})$$

for $\mathbb{S}$ a framed 0–sphere, link, or 2–sphere and $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}))$, such that 7 axioms are satisfied. We list the axioms for the convenience of the reader. We have rephrased the axioms slightly since Juhász considers sutured manifolds in full generality, but we need his result only for sutured manifolds of the form $Y(\mathbf{w})$ for a closed $(Y^3, \mathbf{w})$.

Theorem 2.9.1 ([Juh16, Thm. 8.9])**.** *To define a functorial cobordism maps*

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2})$$

satisfying the standard Spin^c *composition law for path cobordisms* $(W, \boldsymbol{\gamma}, \mathfrak{s}) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$*, it is sufficient to define for each framed* 0*–sphere, link, or* 2*–sphere* $\mathbb{S} \subseteq Y \setminus \mathbf{w}$ *and* Spin^c *structure* $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}))$ *a map*

$$F^\circ_{Y,\mathbb{S},\mathfrak{s}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) \to HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})})$$

which satisfy the following axioms:

1. *If* $d \in \mathrm{Diff}_0(Y, \mathbf{w})$*, then* $d_* : HF^\circ(Y, \mathbf{w}, \mathfrak{s}) \to HF^\circ(Y, \mathbf{w}, \mathfrak{s})$ *is the identity.*

2. *Consider a diffeomorphism* $d : (Y, \mathbf{w}) \to (Y', \mathbf{w}')$ *between closed 3–manifolds, and let* $\mathbb{S} \subseteq Y \setminus \mathbf{w}$ *be a framed 0–sphere, link, or 2–sphere and let* $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}))$*. Let* $\mathbb{S}' = d(\mathbb{S})$*. Let*

   $$d^{\mathbb{S}} : (Y(\mathbb{S}), \mathbf{w}) \to (Y'(\mathbb{S}'), \mathbf{w}') \text{ and } D^{\mathbb{S}} : W(\mathbb{S}) \to W(\mathbb{S}')$$

   be the induced diffeomorphisms, and write $\mathfrak{s}' = D^{\mathbb{S}}(\mathfrak{s}) \in \mathrm{Spin}^c(W(\mathbb{S}'))$*. Then the following diagram is commutative:*

$$\begin{array}{ccc} HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) & \xrightarrow{F^\circ_{Y,\mathbb{S},\mathfrak{s}}} & HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})}) \\ \downarrow{\scriptstyle d_*} & & \downarrow{\scriptstyle (d^{\mathbb{S}})_*} \\ HF^\circ(Y', \mathbf{w}', \mathfrak{s}'|_{Y'}) & \xrightarrow{F^\circ_{Y',\mathbb{S}',\mathfrak{s}'}} & HF^\circ(Y'(\mathbb{S}'), \mathbf{w}', \mathfrak{s}'|_{Y'(\mathbb{S}')}) \end{array}.$$

3. *If $(Y, \mathbf{w})$ is a 3–manifold and $\mathbb{S}$ and $\mathbb{S}'$ are disjoint framed 0–sphere, link, or 2–sphere in Y then $Y(\mathbb{S})(\mathbb{S}') = Y(\mathbb{S}')(\mathbb{S})$ and we denote this manifold by $Y(\mathbb{S}, \mathbb{S}')$. There is an isotopically unique equivalence*

$$\mathcal{W} \stackrel{\text{def}}{=} W(Y(\mathbb{S}), \mathbb{S}') \circ W(Y, \mathbb{S}) \cong W(Y(\mathbb{S}'), \mathbb{S}) \circ W(Y, \mathbb{S}').$$

Given $\mathfrak{s} \in \mathrm{Spin}^c(\mathcal{W})$, this equivalence allows us to define $\mathfrak{s}_{\mathbb{S}} = \mathfrak{s}|_{W(Y,\mathbb{S})}$, $\mathfrak{s}_{\mathbb{S}'} = \mathfrak{s}|_{W(\mathbb{S}')}$, $\mathfrak{s}_{\mathbb{S},\mathbb{S}'} = \mathfrak{s}|_{W(Y(\mathbb{S}),\mathbb{S}')}$ and $\mathfrak{s}_{\mathbb{S}',\mathbb{S}} = \mathfrak{s}|_{W(Y(\mathbb{S}'),\mathbb{S})}$. If at least one of $\mathbb{S}$ and $\mathbb{S}'$ is not 1–dimensional, then the following diagram is commutative:

$$\begin{array}{ccc} HF^\circ(Y, \mathbf{w}, \mathfrak{s}|_Y) & \xrightarrow{F^\circ_{Y,\mathbb{S},\mathfrak{s}|_{\mathbb{S}}}} & HF^\circ(Y(\mathbb{S}), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S})}) \\ \downarrow{\scriptstyle F^\circ_{Y,\mathbb{S}',\mathfrak{s}|_{\mathbb{S}'}}} & & \downarrow{\scriptstyle F^\circ_{Y(\mathbb{S}),\mathbb{S}',\mathfrak{s}|_{\mathbb{S},\mathbb{S}'}}} \\ HF^\circ(Y(\mathbb{S}'), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S}')}) & \xrightarrow{F^\circ_{Y(\mathbb{S}'),\mathbb{S},\mathfrak{s}|_{\mathbb{S}',\mathbb{S}}}} & HF^\circ(Y(\mathbb{S}, \mathbb{S}'), \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S},\mathbb{S}')}). \end{array}$$

4. *Let $\mathbb{L} = \mathbb{L}_1 \cup \mathbb{L}_2$ be a partition of a framed link in Y and write $\mathcal{W}_1 = W(Y, \mathbb{L}_1), \mathcal{W}_2 = W(Y(\mathbb{L}_1), \mathbb{L}_1)$ and $\mathcal{W} = \mathcal{W}_2 \circ \mathcal{W}_1 \cong W(\mathbb{L})$. If $\mathfrak{s}_i \in \mathrm{Spin}^c(\mathcal{W}_i)$ for $i = 1, 2$ then*

$$F^\circ_{Y(\mathbb{L}_1),\mathbb{L}_2,\mathfrak{s}_2} \circ F^\circ_{Y,\mathbb{L}_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(\mathcal{W}) \\ \mathfrak{s}|_{\mathcal{W}_2} = \mathfrak{s}_2 \\ \mathfrak{s}|_{\mathcal{W}_1} = \mathfrak{s}_1}} F^\circ_{Y,\mathbb{L},\mathfrak{s}}.$$

5. *If $\mathbb{S}' \subseteq Y(\mathbb{S})$ intersects the belt sphere of the handle attached along $\mathbb{S}$ once transversely, then there is a diffeomorphism $\phi : Y \to Y(\mathbb{S})(\mathbb{S}')$ (which is the identity on $Y \cap Y(\mathbb{S})(\mathbb{S}')$ and is unique up to isotopy), such that for every $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}') \circ W(\mathbb{S}))$, if we write $\mathfrak{s}|_{\mathbb{S}} = \mathfrak{s}|_{W(\mathbb{S})}$ and $\mathfrak{s}|_{\mathbb{S}'} = \mathfrak{s}|_{W(\mathbb{S}')}$, then $\mathfrak{s}|_{Y(\mathbb{S})(\mathbb{S}')} = \phi_*(\mathfrak{s}|_Y)$ and*

$$F^\circ_{Y(\mathbb{S}),\mathbb{S}',\mathfrak{s}_{\mathbb{S}'}} \circ F^\circ_{Y,\mathbb{S},\mathfrak{s}_{\mathbb{S}}} = \phi_*.$$

6. *Suppose that $\mathbb{L}$ is a framed link in $Y \setminus \mathbf{w}$ and $\mathbb{L}'$ is obtained by sliding a component L_i over L_j along a framed arc a. Let $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{L}))$ and let $\mathfrak{s}' = (D_1^{L_1})_*(\mathfrak{s})$, $\mathfrak{t} = \mathfrak{s}|_Y, \mathfrak{t}_{\mathbb{L}} = \mathfrak{s}|_{Y(\mathbb{L})}, \mathfrak{t}_{\mathbb{L}'} = \mathfrak{s}'|_{Y(\mathbb{L}')}$, where $D_1^{L_1} : W(\mathbb{L}) \to W(\mathbb{L}')$ is the diffeomorphism described just prior to Theorem 8.9 in* [Juh16]. *If $d_1^{L_i} : Y(\mathbb{L}) \to Y(\mathbb{L}')$ is the restriction of $D_1^{L_i}$, then*

the following diagram is commutative:

$$\begin{array}{ccc} HF^\circ(Y,\mathbf{w},\mathfrak{t}) & \xrightarrow{\mathrm{Id}} & HF^\circ(Y,\mathbf{w},\mathfrak{t}) \\ \downarrow{\scriptstyle F^\circ_{Y,\mathbb{L},\mathfrak{s}}} & & \downarrow{\scriptstyle F^\circ_{Y,\mathbb{L}',\mathfrak{s}'}} \\ HF^\circ(Y(\mathbb{L}),\mathbf{w},\mathfrak{t}_{\mathbb{L}}) & \xrightarrow{(d_1^{L_i})_*} & HF^\circ(Y(\mathbb{L}'),\mathbf{w},\mathfrak{t}_{\mathbb{L}'}) \end{array} .$$

7. *For every framed 0–sphere, link, or 2–sphere* $\mathbb{S}$ *and* $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S})) = \mathrm{Spin}^c(W(\mathbb{S}))$*, we have* $F^\circ_{Y,\mathbb{S},\mathfrak{s}} = F^\circ_{Y,\overline{\mathbb{S}},\mathfrak{s}}$ *where* $\overline{\mathbb{S}}$ *denotes the sphere* $\mathbb{S}$ *with the opposite orientation.*

The axioms are verified in [Juh16] for sutured Floer homology, and hence as a special case, for $\widehat{HF}$. We will verify them in this section for HF^--, HF^+ and HF^∞. Much of the work in verifying the axioms for these flavors follows analogously to the case of $\widehat{HF}$. For sutured Floer homology, since disks and triangles aren't allowed to go near the boundary of the sutured Heegaard surface, one can require 1–handles and stabilizations to occur only in a neighborhood of the boundary, eliminating the need for some of the gluing lemmas we have proven. Our main contributions are the various gluing arguments needed to show relations between maps defined by counting holomorphic curves when curves are allowed to go over the basepoints.

Theorem 2.9.2. *The maps* $F^\circ_{Y,\mathbb{S},\mathfrak{s}}$ *defined previously for framed 1– and 3– handles, as well as the maps* $F^\circ_{Y,\mathbb{L},\mathfrak{s}}$ *for framed links satisfy axioms (1)–(7) of* [Juh16, Thm. 8.9]*. In particular, Theorem A follows.*

Proof. Axioms (1) and (2) follow from the invariance of the Heegaard Floer homology groups, as well as the invariance of the 1–,2– and 3–handle maps. Axiom (3) is immediate if neither framed sphere is 1–dimensional. In the case that one of the $\mathbb{S}$ is a framed link, the diagram in axiom (3) commutes by Theorem 2.7.8. Axiom (4) follows from Lemma 2.8.15. Axiom (6), invariance under handleslides in a link $\mathbb{L}$, follows without change as the proof in the sutured setting in [Juh16, Thm. 8.11]. Reversing the orientation of a framed link or sphere has no effect on our maps, so axiom (7) is satisfied.

The remaining axiom, axiom (5), requires more work in our setting. We will consider only the case of canceling 1– and 2–handles in detail, since the case of 2– and 3–handles

is formally the same. The work in [Juh16, Thm. 8.11] is not quite sufficient since in that context disks and triangles cannot go over the basepoints, so weaker gluing arguments can be used in that context.

Suppose that $\mathbb{S}_0 \subseteq Y \setminus \mathbf{w}$ is a framed 0–sphere and let $\mathbb{S}_2 \subseteq Y(\mathbb{S}_0)$ denote the canonical 2–sphere in $Y(\mathbb{S}_0)$. Suppose $\mathbb{K} \subseteq Y(\mathbb{S}_0) \setminus \mathbf{w}$ is a framed knot transverse to $\mathbb{S}_2$ with $|\mathbb{K} \cap \mathbb{S}_2| = 1$. According to [Juh16, Thm. 8.11.] there is a diffeomorphism $\phi : Y \to Y(\mathbb{S}_0)(\mathbb{K})$ which is unique up to isotopy fixing $\mathbf{w}$. Suppose that $\mathfrak{s} \in \mathrm{Spin}^c(W(\mathbb{S}_0)(\mathbb{K}))$. Axiom (5) states that the following relation should hold:

$$F^\circ_{Y(\mathbb{S}_0),\mathbb{K},\mathfrak{s}|_{W(\mathbb{K})}} \circ F^\circ_{Y,\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}} = \phi_*.$$

Following [Juh16, Lem. 8.3], we can pick a bouquet $B(\mathbb{K})$, a diagram

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

and a strongly admissible triple

$$\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w}) = ((\Sigma \setminus B_1 \cup B_2) \cup A, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w})$$

such that

1. $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a strongly admissible Heegaard diagram of $(Y, \mathbf{w})$ subordinate to $\mathbb{S}_0$ (i.e. $\mathbb{S}_0 \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w})$);
2. B_1 and B_2 are two small balls in Σ centered around the points of $\mathbb{S}_0 \subseteq \Sigma$ which don't intersect $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w}$;
3. A is an annulus in $Y(\mathbb{S}_0)$ such that $\partial A = \partial B_1 \cup \partial B_2$;
4. the curve α_0 is a homotopically nontrivial curve in A and β_0 is a small Hamiltonian isotopy of α_0 with $|\alpha_0 \cap \beta_0| = 2$;
5. the curves $\boldsymbol{\gamma}$ are small Hamiltonian isotopies of the $\boldsymbol{\beta}$ curves which intersect $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ transversely;

6. γ_0 intersects α_0 and β_0 each transversally at a single point and γ_0 intersects no other $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves;

7. $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$ is a triple subordinate to the bouquet $B(\mathbb{K})$ and β_0 is a meridian of K, and γ_0 is the framing of K;

Now $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\gamma}', \mathbf{w})$ is a diagram for $(Y(\mathbb{S}_0)(\mathbb{K}), \mathbf{w})$ which is the $(1,2)$–stabilization of $\phi_*\mathcal{H}$. Diagrammatically the situation is depicted in Figure 2.15.

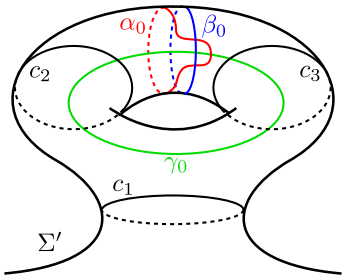

Figure 2.15: The Heegaard triple $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \mathbf{w})$. The circles c_1, c_2 and c_3 are circles on which we need to stretch to compute the 1–handle map or to apply the stabilization lemma for triangle maps.

We need to show that

$$F^\circ_{\mathbb{K},\mathfrak{s}|_{W(\mathbb{K})}} \circ F^\circ_{\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}} = \phi_*.$$

Let $a = \alpha_0 \cap \gamma_0$ and $b = \beta_0 \cap \gamma_0$. In terms of diagrams, the desired equality becomes

$$F^\circ_{\mathcal{T}',\mathfrak{s}|_{W(\mathbb{K})}}(F^\circ_{\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}}(\mathbf{x}) \otimes (\Theta^+ \times b)) = \Phi_{(\phi_*\mathcal{H})\to\mathcal{H}'} \circ \phi_*.$$

By stretching along c_1 sufficiently, the map $\Phi_{\mathcal{H}\to\mathcal{H}'}$ is the stabilization isomorphism $\mathbf{z} \mapsto \mathbf{z} \times a$. For almost complex structures which are sufficiently stretched along c_2 and c_3, the 1–handle map $F_{\mathbb{S}_0,\mathfrak{s}|_{W(\mathbb{S}_0)}}$ takes the form $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$ where θ^+ is the higher degree point in $\alpha_0 \cap \beta_0$. Also for almost complex structure sufficiently stretched along c_1, by Lemma 2.8.7, we have an identification of

$$F^\circ_{\mathcal{T}',\mathfrak{s}|_{W(\mathbb{K})}}((\mathbf{x} \times \theta^+) \otimes (\Theta^+ \times b)) = \sigma \circ \phi_*,$$

where σ is the stabilization isomorphism. Thus in particular, the desired equality will be satisfied if there is some almost complex structure such that the 1–handle map takes the

form $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$, and such the stabilization map takes the form $\mathbf{y} \mapsto \mathbf{y} \times b$ and the 2–handle maps take form in Theorem 2.8.7.

To this end, we will need the following subclaim: if J_s is an almost complex structure on Σ' which is split in neighborhoods of c_1, c_2 and c_3, then if we let $J_s(T_1, T_2, T_3)$ denote the almost complex structure obtained by inserting neck lengths of T_1, T_2 and T_3, then there is an N such that if $T_i, T_i' > N$ for $i \in \{1, 2, 3\}$, then the maps

$$\Phi_{J_s(T_1,T_2,T_3) \to J_s(T_1',T_2',T_3')} : CF^\circ(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S}_0)}) \to CF^\circ(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathfrak{s}|_{Y(\mathbb{S}_0)})$$

are $\mathbb{Z}_2[U_{\mathbf{w}}]$–chain homotopic to the map

$$\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}$$

.

Before we prove the subclaim, let's see how it proves the desired result. If we pick such an N, and if $T_2, T_3 > N$, then the almost complex structure $J_s(T, T_2, T_3)$ is "long enough to compute the 1–handle map" for any $T > N$. Hence for any T, the 1–handle map takes the form $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$ with respect to the almost complex structure $J_s(T, T_2, T_3)$. By Theorem 2.8.7 to the triple $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w})$, we know that for sufficiently large T, we have

$$F^\circ_{\mathcal{T}', \mathfrak{s}|_{W(\mathbb{K})}}((\mathbf{x} \times \theta^+) \otimes (\Theta^+ \times b)) = \sigma \circ \phi_*,$$

implying by our previous observations that the following equality is satisfied for a particular choice of diagrams

$$F^\circ_{\mathbb{K}, \mathfrak{s}|_{W(\mathbb{K})}} \circ F^\circ_{\mathbb{S}_0, \mathfrak{s}|_{W(\mathbb{S}_0)}} = \phi_*.$$

We now prove the subclaim. To see this pick a sequence of 6–tuples of the form $(T_{1,n}, T_{2,n}, T_{3,n}, T'_{1,n}, T'_{2,n}, T'_{3,n})$ where $T_{i,n}$ and $T'_{i,n}$ all approach $+\infty$. Pick an interpolating almost complex structure $\widehat{J}_n$ which is constant outside of a neighborhood of the circles c_i. Suppose that u_n is a sequence of $\widehat{J}_n$–holomorphic strips which represent a Maslov index zero homology class. The almost complex structure $\widehat{J}_n$ converge to an almost complex structure

$\widehat{J}$ on $\Sigma_\infty = \Sigma \cup_{c_1} S^2 \cup_{c_2,c_3} S^2$ which is topologically obtained by collapsing the circles c_i to points. Suppose that all u_n represent the homology class $\phi = \phi_1 \# \phi_2 \# \phi_3$. Taking a Gromov limit of u_n, we get holomorphic representatives of each ϕ_i by broken holomorphic strips U_i. If $\phi \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^+)$ or $\pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^-)$, then by Lemma 2.5.3 we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) + \mu(\phi_3) - 6m = \mu(\phi_1),$$

where m is the multiplicity in the region containing the circles c_1, c_2 and c_3. Hence if $\mu(\phi) = 0$, then $\mu(\phi_1) = 0$. Since U_1 is a broken holomorphic strip on an $\mathbb{R}$-invariant almost complex structure on Σ with Maslov index zero, we conclude that U_1 can only have constant components. Hence the diagonal entries of the transition matrix are equal to the identity. Similarly if $\phi \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$ has Maslov index zero, then by Lemma 2.5.3, we have that

$$\mu(\phi) = \mu(\phi_1) + \mu(\phi_2) + \mu(\phi_3) - 6m = \mu(\phi_1) + 1.$$

If ϕ has $\widehat{J}_n$–holomorphic representatives for arbitrarily large n, then taking a weak limit yields a broken representative U_1 of ϕ_1, implying that $\mu(\phi_1) \geq 0$. Hence $\mu(\phi) = \mu(\phi_1) + 1 \geq 1$, implying that no such u_n are counted in the transition matrix. Hence the lower left entry of the transition map is zero, completing the proof of the subclaim.

The analogous statement holds for 3–handles which cancel a 2–handle. The only difference is that we must use the gluing result in Theorem 2.8.11 for the model computation. □

The following is a corollary of the above theorem and the definition of a functor which is split over Spin^c structures ([Juh16, Def. 8.7])

Corollary 2.9.3. *If $(W, \boldsymbol{\gamma})$ is a path cobordism which decomposes as*

$$(W, \boldsymbol{\gamma}) = (W_2, \boldsymbol{\gamma}_2) \circ (W_1, \boldsymbol{\gamma}_1)$$

then the following composition law holds:

$$F^\circ_{W_2,\boldsymbol{\gamma}_2,\mathfrak{s}_2} \circ F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}}.$$

2.10 Construction and invariance of the graph cobordism maps when $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$

In this section, we construct the graph cobordism maps using the path cobordism maps, and the graph action map, though in this section we focus only on graph cobordisms $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where the inclusions $\pi_0(Y_i) \to \pi_0(W)$ are surjective for $i = 1, 2$. The construction in this section actually works if we assume only that $\pi_0(Y_2) \to \pi_0(W)$ is surjective, but it provides no benefit to work in this intermediate level of generality. Thus we assume in this section that $\pi_0(Y_i) \to \pi_0(W)$ is surjective, and consider the fully general case in Section 2.12, once we have established the composition law for cobordisms with $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$.

In this section, we prove the following weaker version of Theorem C:

Theorem 2.10.1. *Given a graph cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ where Γ is a ribbon graph with coloring $(\sigma, \mathfrak{P})$, and $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ is surjective, there is a homomorphism of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–modules*

$$F_{W,\Gamma,\mathfrak{s}} : HF^\circ(Y_1, \mathbf{w}_1, \sigma|_{\mathbf{w}_1}, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_2, \mathbf{w}_2, \sigma|_{\mathbf{w}_2}, \mathfrak{s}|_{Y_2}).$$

These maps are invariants of the cobordism with embedded colored ribbon graph cobordism.

2.10.1 0– and 4–handle maps

In this subsection we define the 0– and 4–handle maps. Unlike the 1–,2–, or 3–handle maps, the 0– and 4–handle maps are easy to define.

Let $(\mathbb{T}^2, \alpha_0, \beta_0, w_0)$ be the standard diagram for (S^3, w_0) where $\alpha_0 \cap \beta_0 = \{c\}$. Given a three manifold $(Y, \mathbf{w})$ and a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ of $(Y, \mathbf{w})$, the zero handle map

$$F^\circ_{0,\mathfrak{s}} : HF^\circ(Y, \mathbf{w}, \mathfrak{s}) \to HF^\circ(Y \sqcup S^3, \mathbf{w} \cup \{w_0\}, \mathfrak{s} \sqcup \mathfrak{s}_0)$$

is defined on complexes as

$$F^\circ_{0,\mathfrak{s}} : CF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s}) \to CF^\circ(\Sigma \sqcup \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w_0\}, \mathfrak{s} \sqcup \mathfrak{s}_0)$$

by the formula

$$F^\circ_{0,\mathfrak{s}}(\mathbf{x}) = \mathbf{x} \times c.$$

The map is natural in the sense that it commutes with change of diagrams maps on Y.

In the opposite direction, given a based 3–manifold $(Y, \mathbf{w})$, there is a 4–handle map

$$F^\circ_{4,\mathfrak{s}} : HF^\circ(Y \sqcup S^3, \mathbf{w} \cup \{w_0\}, \mathfrak{s} \sqcup \mathfrak{s}_0) \to HF^\circ(Y, \mathbf{w}, \mathfrak{s}),$$

which is defined on complexes as a map

$$F^\circ_{4,\mathfrak{s}} : CF^\circ(\Sigma \sqcup \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w_0\}, \mathfrak{s} \sqcup \mathfrak{s}_0) \to CF^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathfrak{s})$$

by the formula

$$F^\circ_{4,\mathfrak{s}}(\mathbf{x} \times c) = \mathbf{x}.$$

When there is no ambiguity in doing so, we will often omit the Spin^c structure $\mathfrak{s}$ and write only F°_0 or F°_4.

2.10.2 Main construction

We now proceed with the main construction of the graph cobordism maps. For the construction, we need to first fix some data about the graph in a graph cobordism (W, Γ):

- fix a subdivision of Γ such that there are no edges from a vertex to itself (in subsection 2.10.5 we will show that the graph cobordism maps are invariant under subdivision);
- fix an absolute ordering on the edges incident to each vertex (in subsection 2.10.6 we will show that the graph cobordism maps are invariant under cyclic reordering of the vertices);

Suppose that W^4 is a compact, oriented 4–manifold with boundary and Γ is an embedded graph. Let $N(\Gamma)$ denote a regular neighborhood of Γ. The graph Γ yields a handle decomposition, $\mathfrak{H}$, of $N(\Gamma)$ in the following sense: to each vertex in $V(\Gamma)$ which is not in ∂W, we associate a 0–handle, i.e. a subset diffeomorphic to a ball B^4 centered at the vertex

v. To each edge of Γ, we associate a 1–handle, i.e. a subset diffeomorphic to $I \times B^3$. To a vertex $v \in V(\Gamma) \cap \operatorname{int} W^4$, the 1–handles we associated to edges incident to v are glued to the 0–handle associated to v.

We fix some basic notation about embeddings and isotopies of submanifolds and graphs:

Definition 2.10.2. If W and Y are compact manifolds with boundary and $Y \hookrightarrow W$ is an embedding, we say that $Y \subseteq W$ is a **neatly embedded submanifold** if $\partial Y \subseteq \partial W$ and Y is transverse to ∂W. If $\Phi_t : W \times I \to W$ is a 1–parameter family of diffeomorphisms, we say that Φ_t is **neat isotopy** if Φ_t maps ∂W into ∂W. If $\phi_t : Y \to W$ is a 1–parameter family of embeddings of Y into W, we say ϕ_t is a **neat isotopy** if $\phi_t(Y)$ is a neat submanifold for all t. If $\Gamma \subseteq W$ is a graph with a distinguished set of vertices denoted $\partial\Gamma$, we say that Γ is a **neatly embedded graph** if $\partial\Gamma$ consists only of 1–valent vertices in Γ, $\Gamma \cap \partial W = \partial\Gamma$ and Γ intersects ∂W transversely. If Γ_t is a 1–parameter family of graphs in W (i.e. a continuous family of maps $\phi_t : \Gamma \to W$ for a fixed graph Γ, with a fixed set of distinguished boundary vertices, $\partial\Gamma$), we say that Γ_t is a **neat graph isotopy** if the $\phi_t(\Gamma)$ is neatly embedded for all t.

The graph Γ specifies a graph $\widetilde{\Gamma}$ in $\partial N(\Gamma)$, isotopic to Γ inside of $N(\Gamma)$, which we now describe. The handle decomposition of $N(\Gamma)$ induced by Γ induces a decomposition of $\partial N(\Gamma)$ into pieces as well. More precisely, for each vertex $v \in V(\Gamma) \cap \operatorname{int} W$, there is a subset $M(v)$ of $\partial N(\Gamma)$ defined by

$$M(v) = H(v) \cap \partial N(\Gamma)$$

where $H(v)$ denotes the 0–handle corresponding to the vertex v. Note that if v has valence n, then $\partial H(v) \cap \partial N(\Gamma)$ is diffeomorphic to $S^3 \setminus (n \text{ balls})$.

For an edge e, there is a subset $M(e)$ of $\partial N(\Gamma)$ defined by

$$M(e) = H(e) \cap \partial N(\Gamma),$$

where $H(e)$ is the 1–handle associated to the edge e. Note that $H(e)$ is diffeomorphic to $S^2 \times I$.

The graph $\widetilde{\Gamma}$ in $N(\Gamma)$ can now be described as follows. First arbitrarily pick a point on each of the embedded 2–spheres $S \subseteq \partial N(\Gamma)$ corresponding to the boundaries of the subsets $M(e)$ and $M(v)$. In each subset of the form $M(v)$ for $v \in V(\Gamma) \cap \operatorname{int} W$, we place a vertex $\widetilde{v}$ arbitrarily in the interior of $M(v)$, and connect it with n pairwise disjoint paths, each from $\widetilde{v}$ to the marked points on the boundary of $M(v)$. in each subset of the form $M(e)$ for e an edge, there are two boundary spheres, and we pick a path arbitrarily between the distinguished points on the boundary spheres. Patching these all together, we get a graph $\widetilde{\Gamma} \subseteq \partial N(\Gamma)$. There is, of course, ambiguity in choosing the graph $\widetilde{\Gamma}$, but we have the following:

Lemma 2.10.3. *Any graph $\widetilde{\Gamma}$ constructed using the above procedure is isotopic to Γ by a neat isotopy $\phi_t : W \to W$ which maps $\widetilde{\Gamma}$ into $N(\Gamma)$ for all t. Also any two $\widetilde{\Gamma}$ and $\widetilde{\Gamma}'$ constructed in this manner are neatly isotopic in $\partial N(\Gamma)$.*

Proof. To see that $\widetilde{\Gamma}$ and Γ are isotopic, we simply first isotope $\widetilde{\Gamma}$ so the vertices are in the same locations as the vertices of Γ. After performing this isotopy to $\widetilde{\Gamma}$, the edges of $\widetilde{\Gamma}$ and Γ are clearly homotopic in $N(\Gamma)$, and hence isotopic since $N(\Gamma)$ is a 4–manifold.

We now show that $\widetilde{\Gamma}$ and $\widetilde{\Gamma}'$ are neatly isotopic in $\partial N(\Gamma)$. This is a generalization of the "lightbulb theorem", i.e. the fact that any knot in $S^2 \times S^1$ which intersects $S^2 \times \{*\}$ transversely at a single point is isotopic to the knot $\{p\} \times S^1$. If $v \in V(\Gamma) \cap \operatorname{int} W$ then the portion of $\widetilde{\Gamma}$ inside of $M(v)$ consists of a vertex $\widetilde{v}$ connected to each boundary component of $M(v)$ by a path. We need to show that any two such graphs are isotopic relative $\partial M(v)$. Clearly any two graphs are homotopic. The homotopy can be taken to be an isotopy except at finitely many points of time, where a crossing change occurs. We can change any crossing via the following trick: given a crossing of two arcs a and a', traveling along a' in one direction leads to a boundary component of $M(v)$. We simply isotope a along a' in this direction, then around the boundary component and then back to its original position, but with the crossing changed. □

We now continue to describe the construction of the graph cobordism maps. Before we define them, we need to define more auxiliary data and one additional intermediate

map, corresponding to attaching 0– and 1–handles according to $\mathfrak{H}$. Suppose that $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism and Γ is a cyclically ordered graph. Subdivide Γ further so that the edges connecting to Y_2 have as boundary a 1–valent vertex (in Y_2) and a 2–valent vertex (in $\operatorname{int} W$). Let $\boldsymbol{\gamma}_f$ denote the collection of edges connected to Y_2. Let Γ_0 denote the graph $\Gamma \setminus \boldsymbol{\gamma}_f$. Pick a regular neighborhood $N(\Gamma_0)$ of Γ_0. Following Lemma 2.10.3, let Φ_t be a neat isotopy of W which is fixed outside of a small neighborhood of $\overline{N(\Gamma_0)}$, which maps Γ_0 to a graph $\widetilde{\Gamma}_0$ in $\partial N(\Gamma_0)$ and such that $\Phi_t|_{Y_1}$ is an isotopy of Y_1 which is fixed outside a small neighborhood of $\overline{N(\Gamma) \cap Y_1} = \overline{N(\mathbf{w}_1)}$. Let $\widetilde{\boldsymbol{\gamma}}_f$ denote the image of $\boldsymbol{\gamma}_f$ under the isotopy Φ_t. Let $\phi = \Phi_1|_{Y_1}$. Most of this data is illustrated in Figure 2.16.

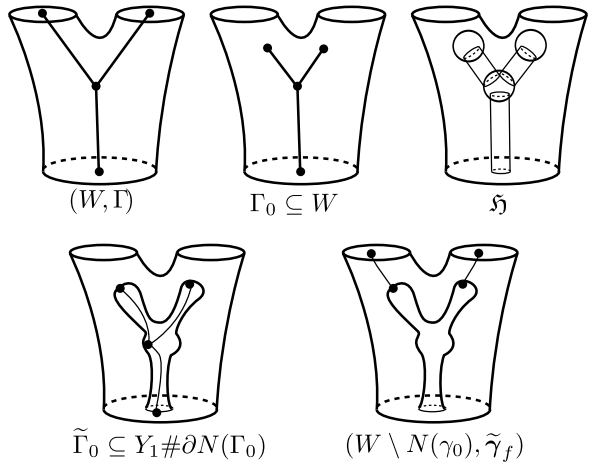

Figure 2.16: The various topological constructions used in the construction of the graph cobordism map.

Suppose that $\mathfrak{t} \in \mathrm{Spin}^c(Y_1)$. We now define a map

$$F^\circ_{\mathfrak{H},\mathfrak{t}} : HF^\circ(Y_1, \phi_* \mathbf{w}, \mathfrak{t}) \to HF^\circ(Y_1 \# \partial N(\Gamma_0), \phi_* \mathbf{w} \cup \mathbf{v}, \mathfrak{t}'),$$

where $\mathbf{v} = \{\widetilde{v} \in V(\widetilde{\Gamma}_0) \cap \operatorname{int} W\}$, which is a composition of 0– and 1–handle maps according to $\mathfrak{H}$. Here $\mathfrak{t}'$ is a Spin^c on $Y_1 \# \partial N(\Gamma_0)$ which is uniquely determined by $\mathfrak{t}$. If $\mathfrak{t} = \mathfrak{s}|_{Y_1}$ for some $\mathfrak{s} \in \mathrm{Spin}^c(W)$, then $\mathfrak{t}' = \mathfrak{s}|_{Y_1 \# \partial N(\Gamma_0)}$. Write $H_1, \ldots, H_k$ for the 0–handles of $\mathfrak{H}$ and $H_{k+1}, \ldots, H_n$ for the 1–handles. Let $G_i \subseteq \Gamma_0$ denote the union of vertices and edges corresponding to handles $H_1, \ldots, H_i$ in the above ordering. If H_i is a zero handle, there

is a diffeomorphism $\phi_i : (Y_1\#\partial N(G_{i-1})) \sqcup S^3 \to Y_1\#\partial N(G_i)$ which is the identity on the component which is not S^3. Since $MCG(S^3, w_0) = 1$, this diffeomorphism is well defined up to isotopy fixing w_0. If G_i is obtained from G_{i-1} by adding a vertex v_i, we define a map

$$f^\circ_{i,\mathfrak{t}} : HF^\circ(Y_1\#\partial N(G_{i-1}), \phi_*\mathbf{w} \cup \mathbf{v}^{\leq i-1}, \mathfrak{t}) \to HF^\circ(Y_1\#\partial N(G_i), \phi_*\mathbf{w} \cup \mathbf{v}^{\leq i}, \mathfrak{t} \sqcup \mathfrak{t}_0)$$

by

$$f^\circ_{i,\mathfrak{t}} = (\phi_i)_* \circ F^\circ_{0,\mathfrak{t}}$$

where $F^\circ_{0,\mathfrak{t}}$ is the 0–handle map. Here $\mathbf{v}^{\leq i} = \mathbf{v}^{\leq i-1} \cup \{\widetilde{v}_i\}$ with $\mathbf{v}^{\leq 0} = \varnothing$, where $\widetilde{v}_i$ denotes the vertex of $\widetilde{\Gamma}_0$ in H_i. The diffeomorphism ϕ_i is defined up to isotopy, so the map $(\phi_i)_*$ is well defined up to chain homotopy.

If G_i is an edge, we define a map

$$f^\circ_{i,\mathfrak{t}} : HF^\circ(Y_1\#\partial N(G_{i-1}), \phi_*\mathbf{w} \cup \mathbf{v}, \mathfrak{t}) \to HF^\circ(Y_1\#\partial N(G_i), \phi_*\mathbf{w} \cup \mathbf{v}, \mathfrak{t}')$$

simply by

$$f^\circ_{i,\mathfrak{t}} = F^\circ_{(Y_1\#\partial N(G_{i-1})\times I)\cup H_i, (\phi_*\mathbf{w}\cup\mathbf{v})\times I, \bar{\mathfrak{t}}},$$

the path cobordism map associated to path cobordism $((Y_1\#\partial N(G_{i-1}) \times I) \cup H_i, (\phi_*\mathbf{w}) \times I)$. Here $\bar{\mathfrak{t}}$ is the unique Spin^c structure on $(Y_1\#\partial N(G_{i-1})) \times I \cup H_i$ which restricts to $\mathfrak{t}$ on $Y_1\#\partial N(G_{i-1})$.

We now define the map $F^\circ_{\mathfrak{H},\mathfrak{t}}$ as the composition

$$F^\circ_{\mathfrak{H},\mathfrak{t}} = \prod_{i=k+1}^{n} f^\circ_{i,\mathfrak{t}_i},$$

where $\mathfrak{t}_1 = \mathfrak{t}$, and $\mathfrak{t}_i$ is either $\mathfrak{t}_{i-1} \sqcup \mathfrak{t}_0$ or $\mathfrak{t}'_{i-1}$, using the above notation, depending on whether G_i is obtained from G_{i-1} by adding a vertex or an edge.

We define the graph cobordism map $F^\circ_{\Gamma,\mathfrak{s}}$ as follows. Write W' for the 4–manifold $W \setminus N(\Gamma_0)$. Write $\mathcal{W}'$ for the path cobordism $\mathcal{W}' = (W', \widetilde{\gamma}_f)$. Notice that W' has boundary equal to $-(Y_1\#\partial N(\Gamma_0)) \sqcup Y_2$. Define the graph cobordism map as

$$F^\circ_{W,\Gamma,\mathfrak{s}} \stackrel{\text{def}}{=} F^\circ_{\mathcal{W}',\mathfrak{s}|_{W'}} \circ A^V_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} \circ \phi_*. \tag{2.8}$$

The set V of vertices consists of all vertices in $V(\widetilde{\Gamma}_0)$. Notice that since V consists of all vertices of $V(\widetilde{\Gamma}_0)$, *a-priori* $F^\circ_{W,\Gamma,\mathfrak{s}}$ could depend on the absolute ordering of the edges adjacent to each vertex, since Lemma 2.6.12 only applies to vertices which are not in the set V. We show in subsection 2.10.6 that it is invariant under cyclically reordering the edges adjacent to each vertex.

As we proceed to analyze the graph cobordism maps, if we are given a handlebody decomposition $\mathfrak{H}$ of the regular neighborhood of a graph Γ, we will often write

$$F^\circ_{\mathfrak{H}}$$

instead of $F^\circ_{\mathfrak{H},\mathfrak{s}}$ since there is no ambiguity.

Having now defined the graph cobordism maps in terms of some additional choices, we now proceed to show that it is invariant under the additional data and depends only on the Spin^c decorated ribbon-graph cobordism $(W, \Gamma, \mathfrak{s})$.

2.10.3 Invariance of the graph cobordism maps from $\widetilde{\Gamma}_0$

In the construction of the graph cobordism maps, we needed to isotope Γ into $\partial N(\Gamma)$ to get a graph $\widetilde{\Gamma}$. We first need to show that the graph cobordism map is independent of the choice of isotopic graph $\widetilde{\Gamma}_0$ or the choice of isotopy $\Phi_t : W \times I \to I$ sending Γ_0 to $\widetilde{\Gamma}_0$.

Lemma 2.10.4. *Suppose that $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ are two choices of graphs in $\partial N(\Gamma)$ as constructed in the previous section, and Φ_t and Φ'_t are two isotopies of W mapping Γ_0 to $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ respectively, which are each the identity outside of a neighborhood of $\overline{N(\Gamma_0)}$. Then the graph cobordism maps computed with respect to the pairs $(\widetilde{\Gamma}_0, \Phi_t)$ and $(\widetilde{\Gamma}'_0, \Phi'_t)$ agree.*

Proof. By Lemma 2.10.3 we can find a neat isotopy $(\widetilde{\Gamma}_0)_t$ between $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ inside of $\partial N(\Gamma_0)$. Such an isotopy of graphs induces an isotopy between $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}'_0$ inside of $Y_1 \# \partial N(\Gamma_0)$. First extend the isotopy of graphs to an isotopy of $Y_1 \# \partial N(\Gamma_0)$ which is fixed outside of a small neighborhood of $\partial N(\Gamma_0)$. Since this isotopy can be chosen to map $\partial(\partial N(\Gamma_0))$, a collection of 2–spheres, to itself, we can then extend this to a neat isotopy $\Psi_t : W \times I \to W$ which

is also fixed outside of a small neighborhood of $\partial N(\Gamma_0)$. Define $\psi : Y_1 \to Y_1$ by $\psi = \Psi_1|_{Y_1}$. Define $\phi = \Phi_1|_{Y_1}$ and $\phi' = \Phi_1'|_{Y_1}$.

To show that the graph cobordism maps computed with respect to the pairs $(\widetilde{\Gamma}_0, \Phi_t)$ and $(\widetilde{\Gamma}_0', \Phi_t')$ agree, it is sufficient to show that each square in the following diagram commutes (the diagram has been broken into two lines):

$$\begin{array}{ccccccc}
HF^\circ(Y_1, \mathbf{w}_1) & \xrightarrow{\phi_*} & HF^\circ(Y_1, \widetilde{\mathbf{w}}_1) & \xrightarrow{F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}} & HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_1) & \xrightarrow{A_{\widetilde{\Gamma}_0}} & \cdots \\
\downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \psi_*} & & \downarrow{\scriptstyle (\Psi_1)_*} & & \\
HF^\circ(Y_1, \mathbf{w}_1) & \xrightarrow{\phi_*'} & HF^\circ(Y_1, \widetilde{\mathbf{w}}_1') & \xrightarrow{F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}} & HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_1') & \xrightarrow{A_{\widetilde{\Gamma}_0'}} & \cdots
\end{array}$$

$$\begin{array}{ccc}
\cdots HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_f) & \xrightarrow{F^\circ_{(W \setminus N(\Gamma_0), \widetilde{\gamma}_f), \mathfrak{s}|_{W \setminus N(\Gamma_0)}}} & HF^\circ(Y_2, \mathbf{w}_2) \\
\downarrow{\scriptstyle (\Psi_1)_*} & & \downarrow{\scriptstyle (\Psi_1)_* = \mathrm{id}} \\
\cdots HF^\circ(Y_1 \# \partial N(\Gamma_0), \widetilde{\mathbf{w}}_f') & \xrightarrow{F^\circ_{(W \setminus N(\Gamma_0), \widetilde{\gamma}_f'), \mathfrak{s}|_{W \setminus N(\Gamma_0)}}} & HF^\circ(Y_2, \mathbf{w}_2)
\end{array}$$

In the above expression, $\widetilde{\mathbf{w}}_1$ denotes $\widetilde{\Gamma_0} \cap Y_1$ and $\widetilde{\mathbf{w}}_f$ denotes the endpoints of the $\widetilde{\boldsymbol{\gamma}}_f$ arcs which lie in $\operatorname{int} W$.

The composition of the top row is the graph cobordism map computed with respect to $(\widetilde{\Gamma}_0, \Phi_t)$ and the second is the graph cobordism map computed with respect to $(\widetilde{\Gamma}_0', \Phi_t')$. All of the squares except for the first one commute simply because of diffeomorphism invariance of the maps involved. Some amount of care is needed in the first square since ϕ' is not necessarily equal to $\psi \circ \phi$. On the other hand, we claim that $\psi \circ \phi$ and ϕ' are isotopic relative $\mathbf{w}_1$. To see this note that ϕ, ψ and ϕ' are all equal to the identity of a neighborhood $N(\mathbf{w}_1) \subseteq Y_1$ of the basepoints. The map $(\phi')^{-1} \circ \psi \circ \phi$ is isotopic to the map resulting from finger moves of the basepoints $\mathbf{w}_1$ around loops which are contained in $N(\mathbf{w}_1)$. Hence $(\phi')^{-1} \circ \psi \circ \phi \simeq \mathrm{id}$ relative $\mathbf{w}$, and hence $\phi' \simeq \psi \circ \phi$ relative $\mathbf{w}$. Hence the diagram commutes, and the graph cobordism map doesn't depend on the choice of $\widetilde{\Gamma}_0$ and Φ_t. $\square$

2.10.4 Invariance of the cobordism maps from $N(\Gamma_0)$

The next step towards proving invariance is show that the graph cobordism maps don't depend on the choice of regular neighborhood $N(\Gamma_0)$. To this end, we note that in the construction of the graph cobordism maps, we created a sort of handle body decomposition of $N(\Gamma_0)$, the combinatorics of which is described by Γ_0. We prove the following lemma about invariance up to isotopy of these handlebody decompositions:

Lemma 2.10.5. *Suppose that $\Gamma \subseteq W$ is a neatly embedded graph and that $N(\Gamma)$ and $N'(\Gamma)$ are two regular neighborhoods and $\mathfrak{H}$ and $\mathfrak{H}'$ are two handlebody decompositions of the type described in the previous section for $N(\Gamma)$ and $N'(\Gamma)$. Then there is a neat isotopy of W which can be taken to be fixed outside of a given neighborhood of $\overline{N(\Gamma) \cup N'(\Gamma)}$ which maps each i–handle of $\mathfrak{H}$ diffeomorphically onto the corresponding i–handle of $\mathfrak{H}'$.*

Proof. We use invariance up to isotopy of the tubular neighborhood of a neatly embedded submanifold $Y \subseteq W$, [Hir94, Sec. 6]. We first apply isotopy invariance to each 0–handle, which are tubular neighborhoods of vertices in $V(\Gamma) \cap \operatorname{int} W$. This yields an isotopy Φ_t fixing the vertices which maps each 0–handle of $\mathfrak{H}$ diffeomorphically onto the corresponding 0–handle of $\mathfrak{H}'$. It is easy to see that such an isotopy can be taken to be fixed outside of a neighborhood of the union of the 4–balls associated to each 0–handle. Similarly, it is also easy to arrange the isotopy Φ_t so $\Phi_1 H(e)$ is a regular neighborhood of $e \setminus \cup_{v \in V} \Phi_1 H(v)$. Hence we can assume that the handlebody decompositions $\mathfrak{H}$ and $\mathfrak{H}'$ have 0–handles which are equal. Let $B_1^0, \ldots, B_k^0$ denote the 0–handles of the handlebody decompositions $\mathfrak{H}$ and $\mathfrak{H}'$. Let $e_1, \ldots, e_n$ denote the edges of Γ. We proceed by induction. Suppose that an intermediate isotopy Φ_t^i has been constructed so that all of the 1–handles corresponding to edges e_j with $j \leq i$ have corresponding 1–handles in $\Phi_1^i(\mathfrak{H})$ and $\mathfrak{H}'$ which are equal and such that Φ_1^i preserves Γ. Let U_i denote the interior of the union of all of the 0–handles and all of the 1–handles corresponding to edges e_j with $j \leq i$ (for either handlebody decomposition $\Phi_1^i(\mathfrak{H})$ or $\mathfrak{H}'$, since these handles agree by induction). Then the 1–handles corresponding to e_{j+1} in the two handlebody decompositions are both regular neighborhoods of a 1–dimensional neatly

embedded submanifold of $W \setminus U_i$, and hence there is a (neat) isotopy of $W \setminus U_i$ mapping the 1–handle in $\Phi_1^i(\mathfrak{H})$ to the 1–handle of $\mathfrak{H}'$. It is easy to see that this isotopy can be taken to be fixed outside of a given neighborhood of the union of the two 1–handles corresponding to e_{j+1} in $\Phi_1^i(\mathfrak{H})$ and $\mathfrak{H}'$. The isotopy Φ_t^{j+1} can then be extended to U_i. Proceeding by induction yields the total isotopy as a composition. □

Corollary 2.10.6. *The graph cobordism maps are invariant under the choice of regular neighborhood $N(\Gamma_0)$ and the choice of a handlebody decomposition of $N(\Gamma_0)$ into 0–handles and 1–handles (the combinatorics of which is determined by the graph Γ_0, as in the construction of the graph cobordism map).*

Proof. One simply uses Lemma 2.10.5 to create an isotopy which maps one regular neighborhood to the other, then a commutative diagram nearly identical to the one in Lemma 2.10.4 proves invariance of the graph cobordism map. □

2.10.5 Invariance of the graph cobordism maps from subdivision

One important property of the graph cobordism maps is that they are invariant under subdivision of an edge, with respect to either absolute ordering of the edges adjacent to the new vertex. It follows that the ordering of the edges at a 2–valent vertex does not affect the graph cobordism maps, though more generally we will show in Section 2.10.6 that the graph cobordism maps are invariant under the cyclic reordering of the edges incident to a vertex.

Lemma 2.10.7. *Suppose $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism with Γ a strong ribbon graph and $\mathfrak{s} \in \mathrm{Spin}^c(W)$, and that v is a point on the interior of Γ. Suppose that Γ' is the graph resulting from adding v to the interior of an edge. Then $F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\Gamma',\mathfrak{s}}$. The result holds for either ordering of the two edges adjacent to the new vertex.*

Proof. First note that subdividing an edge does not change the regular neighborhood, so we can take $N(\Gamma_0) = N(\Gamma_0')$. The handlebody decomposition of $N(\Gamma_0)$ does change slightly however. There is an extra 0–handle corresponding to the vertex v, and the edge e which

contained v is broken into two edges e_1 and e_2. Let $\mathfrak{H}$ denote the handlebody decomposition of $N(\Gamma_0)$ corresponding to Γ, and let $\mathfrak{H}'$ denote the handlebody decomposition of $N(\Gamma_0') = N(\Gamma_0)$ according to Γ'.

Pick graphs $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}_0'$ inside of $\partial N(\Gamma_0)$ which are neatly isotopic to Γ_0 in $N(\Gamma_0)$ such that $\widetilde{\Gamma}_0'$ is a subdivision of $\widetilde{\Gamma}_0$, as well as a neat isotopy $\Phi_t : W \times I \to W$ mapping Γ_0 to $\widetilde{\Gamma}_0$ and $\widetilde{\Gamma}_0'$. Let $\phi = \Phi_1|_{Y_1}$.

Using the formula for the graph cobordism map in Equation (2.8), we want to show that

$$F^\circ_{\mathcal{W}',\mathfrak{s}|_{W'}} \circ A^V_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} \circ \phi_* = F^\circ_{\mathcal{W}',\mathfrak{s}|_{W'}} \circ A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'} \circ F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}} \circ \phi_*.$$

To this end, it is sufficient to show that

$$A^V_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} = A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'} \circ F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}.$$

Applying Lemma 2.6.10, we see that for either ordering of the edges e_1 and e_2, the desired equality becomes

$$A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'} \circ S^+_v \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} = A^{V\cup\{v\}}_{\widetilde{\Gamma}_0'} \circ F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}.$$

To show this, it is sufficient to show that

$$S^+_v \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} = F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}.$$

This is a relatively easy model computation. In Figure 2.17, we perform a model computation of $F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}$ and see explicitly that this is equal to $S^+_v \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}$, completing the proof.

□

2.10.6 Dependence on cyclic ordering

In the construction of the graph cobordism maps, we picked an absolute ordering of the edges incident to a vertex representing a given cyclic ordering. In this section, we show that the graph cobordism maps are invariant under cyclic reordering of the edges adjacent to a given vertex. The graph cobordism map is defined as a composition of several maps, including a

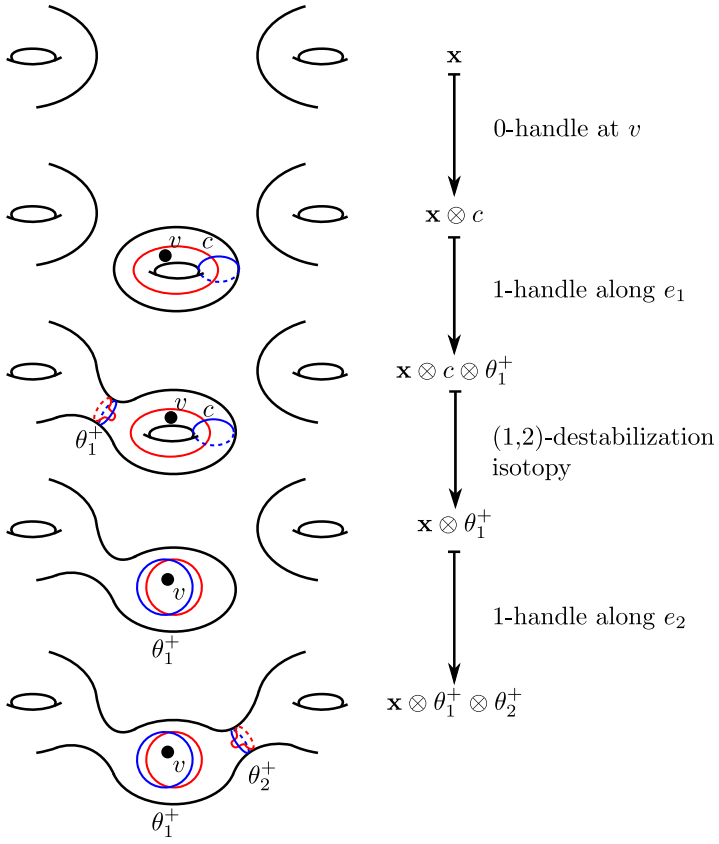

Figure 2.17: A computation of a portion of the maps in $F^\circ_{\mathfrak{H}',\mathfrak{s}|_{Y_1}}$ in Lemma 2.10.7. We see explicitly that the composition is $S^+_v \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}}$.

graph action map. Although the graph action map does depend on the absolute ordering since we have included all of the vertices in the incoming set V, the composition does not:

Lemma 2.10.8. *The graph cobordism maps depend only on the cyclic ordering of the edges incident to a vertex.*

Proof. Let $\Gamma_0, \widetilde{\Gamma}_0$ and $\mathfrak{H}$ be as in the construction of the graph cobordism map. We will show that $A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H}}$ depends only on the cyclic ordering at a vertex. To do this, we will factor $A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H}}$ in an interesting way, and then apply Lemma 2.6.12.

Let $\mathcal{G}$ denote the flow graph that we associate to $\widetilde{\Gamma}_0$ in the graph cobordism maps and let $\mathcal{G}_m \circ \cdots \circ \mathcal{G}_1$ denote a Cerf decomposition of $\mathcal{G}$. Let $\mathcal{G}_v$ denote the elementary flow graph associated to the vertex v. Pick two distinct edges e_i and e_j which are incident to v,

arbitrarily. By Lemma 2.10.7, subdividing the graph Γ does not effect the graph cobordism map, so we can assume that e_i and e_j are distinct edges which both connect v to vertices which are not v. Let $\mathfrak{H}'$ denote the handlebody decomposition of $N(\Gamma_0)$ obtained by replacing the union of the 0–handle associated to v and the 1–handles associated to e_i and e_j, with a single 1–handle. We now claim that

$$F^\circ_{\mathfrak{H}} = S^+_v \circ F_{\mathfrak{H}'}.$$

This is essentially the same model computation that we used to prove subdivision invariance in Lemma 2.10.7, and is shown in Figure 2.17. Having established the above equality, letting $\widetilde{e}_i$ denote the portion of the edge e_i in the elementary flow graph $\mathcal{G}_v$, we now just compute that

$$\begin{aligned} A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H}} &= A_{\mathcal{G}_m} \cdots A_{\mathcal{G}_v} \cdots A_{\mathcal{G}_1} \circ F^\circ_{\mathfrak{H}} \\ &= A_{\mathcal{G}_m} \cdots (S^-_v A_{\widetilde{e}_n} \cdots A_{\widetilde{e}_1}) \cdots A_{\mathcal{G}_1} \circ S^+_v \circ F^\circ_{\mathfrak{H}'} \\ &= A_{\mathcal{G}_m} \cdots (S^-_v A_{\widetilde{e}_n} \cdots A_{\widetilde{e}_1} S^+_v) \cdots A_{\mathcal{G}_1} \circ F^\circ_{\mathfrak{H}'}. \end{aligned}$$

The only term in the above expression that depends on the ordering of the edges $e_1, \ldots, e_n$ is $(S^-_v A_{\widetilde{e}_n} \cdots A_{\widetilde{e}_1} S^+_v)$, so applying Lemma 2.6.12, shows that it only depends on the cyclic ordering of the edges. □

2.11 The composition law for $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$

The main goal of this section is to prove Lemma 2.11.7, i.e. the composition law for graph cobordisms $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ with $\pi_0(Y_i) \to \pi_0(W)$ a surjection for $i = 1, 2$. In working towards a proof of this, much of this section is devoted to showing the relationship between the graph cobordism maps, the graph action maps and the path cobordism maps. The simplest such result is Lemma 2.11.3, which states that the graph cobordism maps for a cylindrical cobordism $(Y \times I, \Gamma)$ are the same as a the graph action map for the graph projected into Y. We begin with some technical lemmas:

Lemma 2.11.1. *Suppose that $(W, \boldsymbol{\gamma}) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a path cobordism and E is a neatly embedded path in W between two points in Y_1. Let $N(E)$ denote a tubular neighborhood of E and let $\mathfrak{E}$ denote the handlebody decomposition of $N(E)$ as a single 1–handle as in the construction of the graph cobordism maps. Let $\mathfrak{s} \in \mathrm{Spin}^c(W)$. Letting $F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}$ denote the handlebody map in the construction of the graph cobordism maps, then*

$$F^\circ_{W\setminus N(E),\boldsymbol{\gamma},\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}}.$$

Proof. By construction

$$F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} = F_{(Y_1\times I)\cup N(E),\mathbf{w}_1\times I,\mathfrak{s}'},$$

where $\mathfrak{s}'$ is the unique Spin^c structure on $Y_1 \times I \cup N(E)$ extending $\mathfrak{s}|_{Y_1}$. By the composition law for path cobordism maps (Corollary 2.9.3) we have

$$F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} = F^\circ_{W\cup(Y_1\times I),\boldsymbol{\gamma}\cup(\mathbf{w}_1\times I),\bar{\mathfrak{s}}},$$

where $\bar{\mathfrak{s}}$ is the unique Spin^c structure on $W \cup (Y_1 \times I)$ extending $\mathfrak{s}$. Using this and the composition law for path cobordisms, we have that

$$\begin{aligned} F_{W,\boldsymbol{\gamma},\mathfrak{s}} &= F^\circ_{W\cup(Y_1\times I),\boldsymbol{\gamma}\cup(\mathbf{w}_1\times I),\bar{\mathfrak{s}}} \\ &= F^\circ_{W\setminus N(E),\boldsymbol{\gamma},\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{(Y_1\times I)\cup N(E),(\mathbf{w}_1\times I),\bar{\mathfrak{s}}} \\ &= F^\circ_{W\setminus N(E),\boldsymbol{\gamma},\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}, \end{aligned}$$

as we wanted. □

Lemma 2.11.2. *Suppose that W is a cobordism from Y_1 to Y_2 with a Morse function f which has critical points which are only index 1,2 or 3 and such that any index i critical point has value less than any index $i+1$ critical point. Suppose that $\mathcal{G} = (\Gamma, V_0, V_1, \sigma, \varnothing)$ is a flow graph in Y_1, which has trivial intersection with the descending manifolds of the critical points of f. Let $\boldsymbol{\gamma}_0$ denote the paths determined by the flow under f of the points V_0 and similarly let $\boldsymbol{\gamma}_1$ denote the paths determined by the flow of the points V_1. Then there is a*

flow graph $\mathcal{G}' = (\Gamma', V_0', V_1', \sigma, \varnothing)$ *in* Y_2 *such that*

$$F^{\circ}_{W,\boldsymbol{\gamma}_1,\mathfrak{s}} \circ A_{\mathcal{G}} = A_{\mathcal{G}'} \circ F^{\circ}_{W,\boldsymbol{\gamma}_0,\mathfrak{s}}.$$

After perturbing to be neatly embedded in W, $\Gamma' * \boldsymbol{\gamma}_0$ *and* $\boldsymbol{\gamma}_1 * \Gamma$ *are isotopic.*

Proof. Using the composition law for path cobordisms we can assume that f either has a single critical point of index 1 or 3, or has only critical points of index 2. The Morse function f induces a diffeomorphism $\psi_f : Y_1(\mathbb{L}) \to Y_2$, where $\mathbb{L}$ is the collection of framed spheres (all of the same dimension) consisting of the the intersection of the descending manifolds of critical points of f and Y_1. Assuming that Γ intersects these framed spheres trivially, the map ψ_f then naturally induces a flow graph $\mathcal{G}' = (\psi_f)_*\mathcal{G}$, which is the image of $\mathcal{G}$ under ψ_f. The paths $\boldsymbol{\gamma}_i$ are simply the flow under the gradient flow of f of the basepoints V_i. The basepoints V_i' are the images under ψ_f of V_i.

To see that we have the relation

$$F^{\circ}_{W,\boldsymbol{\gamma}_1,\mathfrak{s}} \circ A_{\mathcal{G}} = A_{\mathcal{G}'} \circ F^{\circ}_{W,\boldsymbol{\gamma}_0,\mathfrak{s}},$$

we pick a Heegaard surface $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, V_0)$ for (Y_1, V_0). We can isotope Σ so that V_1 sits inside of Σ. Now by stabilizing the Heegaard surface Σ and isotoping the edges and additional vertices of Γ relative the points V_0 and V_1, we can assume they sit inside of Σ. We can assume that the final image of the graph does intersect the embedded spheres corresponding to the descending manifold of the critical points of f, since we assumed that the graph Γ initially did not intersect those spheres. The Morse function f also induces a diffeomorphism of $\Phi : W \to W(Y_1, \mathbb{L})$. Given a nonempty collection of basepoints $\mathbf{w} \subseteq \Sigma$, if $\boldsymbol{\gamma}_{\mathbf{w}}$ is the collection of paths in W given by flowing $\mathbf{w}$ along the gradient flow of f, then by definition

$$F^{\circ}_{W,\boldsymbol{\gamma}_{\mathbf{w}},\mathfrak{s}} = \psi_f \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}}.$$

In particular, since the maps associated to framed 0–spheres, framed links, and framed 2–spheres all commute with the free stabilization and destabilization maps (Theorem 2.5.17,

Lemma 2.7.7) , as well as the relative homology maps (Lemmas 2.4.4 and 2.7.6), we compute that

$$\begin{aligned} F^{\circ}_{W,\boldsymbol{\gamma}_1,\mathfrak{s}} \circ A_{\mathcal{G}} &= \psi_f \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}} \circ A_{\mathcal{G}} \\ &= \psi_f \circ A_{\mathcal{G}} \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}} \\ &= A_{\mathcal{G}'} \circ \psi_f \circ F^{\circ}_{\mathbb{L},\Phi_*\mathfrak{s}} \\ &= A_{\mathcal{G}'} \circ F^{\circ}_{W,\boldsymbol{\gamma}_0,\mathfrak{s}}, \end{aligned}$$

as we wanted to show.

As graphs which are perturbed to be neatly embedded in W, $\Gamma' * \boldsymbol{\gamma}_0$ and $\boldsymbol{\gamma}_1 * \Gamma$ can be seen to be isotopic in W because Γ' is simply the flow of Γ under the gradient flow of f. $\square$

Lemma 2.11.3. *Suppose that $(W,\Gamma) : (Y,\mathbf{w}_1) \to (Y,\mathbf{w}_2)$ is a graph cobordism with $W = Y \times I$. Assume that Γ is perturbed so that Γ projects to an embedded graph Γ_Y in Y under the projection $\pi : Y \times I \to Y$. Then*

$$F^{\circ}_{W,\Gamma,\mathfrak{s}} = A_{\mathcal{G},\mathfrak{s}}$$

where $\mathcal{G} = (\Gamma_Y, \mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$.

Proof. Let $\Gamma_0, N(\Gamma_0), \widetilde{\Gamma}_0, \phi_*, \widetilde{\boldsymbol{\gamma}}_f$ and $\mathfrak{H}$ be as in the construction of the graph cobordism maps. Let

$$V_W = V(\Gamma_0) \cap \operatorname{int} W$$

denote the collection of vertices of Γ_0 which are in the interior of W. Let V_Y denote the projection of these vertices onto $Y \times \{0\}$. Let $\widetilde{V}_W$ denote the vertices of $\widetilde{\Gamma}_0$ corresponding to vertices of V_W. By assuming Γ is in generic position, we can pick $\widetilde{v} \in \widetilde{V}_W$ so that $\pi_Y \widetilde{v} = \pi_Y v$. Pick a set of paths E in W from points in $Y \times \{0\} \setminus N(\mathbf{w})$ to the vertices in V_W such that the paths $e \in E$ are constant in the Y component. By picking an appropriate regular neighborhood $N(\Gamma_0)$, we can assume that the paths E naturally induce a set of paths E' in $W \setminus N(\Gamma_0)$ from points in $Y \times \{0\} \setminus N(\mathbf{w}_1)$ to points in $\partial N(v)$ for each $v \in V_W$. Let $\mathfrak{E}$ denote the handle decomposition of $N(E')$ induced by these paths.

Pick $\widetilde{\Gamma}_0$ and $\widetilde{\boldsymbol{\gamma}}_f$ so that $\widetilde{\boldsymbol{\gamma}}_f$ are also constant in the Y component. Let $\mathbf{w}_f$ denote the basepoints in $\widetilde{\boldsymbol{\gamma}}_f \cap \widetilde{\Gamma}_0$. Pick paths $\boldsymbol{\gamma}'$ from the vertices $\widetilde{v} \in \widetilde{V}_W \setminus \mathbf{w}_f$ to points in $Y \times \{1\}$ and pick the paths in $\boldsymbol{\gamma}'$ to all be constant in Y as well.

By definition we have that

$$F^{\circ}_{W,\Gamma,\mathfrak{s}} = F_{W\setminus N(\Gamma_0),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|_{W\setminus N(\Gamma_0)}} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}} \circ \phi_*.$$

Write $\mathfrak{H} = \mathfrak{H}_0 \cup \mathfrak{H}_1$ where $\mathfrak{H}_0$ are the 0–handles corresponding to vertices in V_W and $\mathfrak{H}_1$ are the 1–handles corresponding to edges of Γ_0. Using the above observation, and Lemma 2.11.1, we perform the following manipulation (in the following, we will abbreviate any restriction of $\mathfrak{s}$ by $\mathfrak{s}|$):

$$\begin{aligned} F_{W\setminus N(\Gamma_0),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}} \circ \phi_* &= F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ F^{\circ}_{\mathfrak{E}} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ F^{\circ}_{\mathfrak{H}_0} \circ \phi_* \\ &= F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ F^{\circ}_{\mathfrak{E}} \circ F^{\circ}_{\mathfrak{H}_0} \circ \phi_*. \end{aligned}$$

The second manipulation results from commuting $F^{\circ}_{\mathfrak{E}}$ with the graph action map and the 1–handle maps in $F^{\circ}_{\mathfrak{H}_1}$. Letting $\widetilde{\mathbf{w}}_1 = \phi(\mathbf{w}_1)$, there is a diffeomorphism $\Phi : (Y, \widetilde{\mathbf{w}}_1 \cup V_Y) \to (Y\#\partial(N(E) \cup N(V_W)), \widetilde{\mathbf{w}}_1 \cup \widetilde{V}_W)$ mapping a vertex $\widetilde{v} \in V_Y$ to the corresponding $v \in \widetilde{V}_W$. The map Φ can be taken to be the identity outside a small neighborhood of $(Y\#\partial(N(E) \cup N(V_W))) \setminus Y$. A model computation similar to the one in Lemma 2.10.7 shows that

$$F^{\circ}_{\mathfrak{E}} \circ F^{\circ}_{\mathfrak{H}_0} = \Phi_* \circ S^{+}_{V_Y}.$$

Continuing the manipulations above, we thus have that

$$F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ F^{\circ}_{\mathfrak{E}} \circ F^{\circ}_{\mathfrak{H}_0} \circ \phi_*$$

$$= F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ S^{+}_{V_Y} \circ \phi_*.$$

We now claim that $W \setminus (N(\Gamma_0) \cup N(E))$ has Morse function which has only index 2 critical points, and which has the paths in $\widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}'$ as flow lines. To see this, first observe that there

is a diffeomorphism of $W \setminus (N(E) \cup N(V_W))$ with $W = Y \times I$ which extends Φ and maps the paths $\widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}'$ to paths of the form $(V_Y \cup \widetilde{\mathbf{w}}_1) \times I$. On $W = Y \times I$, we first take the Morse function $f_0(y,t) = t$. We then modify f_0 to create a Morse function f by introducing an index 1–2 critical point pair for each edge of Γ_0. We then compose f_0 with an isotopy of W which moves the descending manifold of each index 1 critical point to the core of the corresponding 1–handle in $\mathfrak{H}_1$. This can be done without affecting the flowlines corresponding to the points $V_Y \cup \widetilde{\mathbf{w}}_1$. Such a Morse function f_0 induces a decomposition $W = W_2 \cup_{Y'} W_1$ where Y' is a level set of f_0 and $f_0|_{W_1}$ has only index 1 critical points and $f_0|_{W_2}$ has only index 2 critical points. Furthermore, there is an equivalence

$$\Psi : (W \setminus (N(\Gamma_0) \cup N(E)), \widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}') \to (W_2, (\widetilde{\boldsymbol{\gamma}}_f \cup \boldsymbol{\gamma}') \cap W_2).$$

We then define a Morse function f on $W \setminus (N(\Gamma_0) \cup N(E))$ by $f = f_0 \circ \Psi$. Let ψ_f denote the diffeomorphism

$$\psi_f : (Y \# \partial(N(E) \cup N(\Gamma_0))(\mathbb{L}), \widetilde{\mathbf{w}}_1 \cup \widetilde{V}_W) \to (Y, \widetilde{\mathbf{w}}_1 \cup V_Y)$$

induced by the Morse function $f|_{W_2}$, where $\mathbb{L} \subseteq Y \# \partial(N(E) \cup N(\Gamma_0))$ is a framed link.

We now observe that if $\mathbb{L}_0 \subseteq Y_0$ is a framed link in a 3–manifold Y_0, and $\mathcal{G} = (\Gamma, V_0, V_1, V, \sigma)$ is a flow graph in Y_0, then

$$F^\circ_{\mathbb{L}_0,\mathfrak{s}} \circ A_{\mathcal{G}} = A_{\mathcal{G}'} \circ F^\circ_{\mathbb{L}_0,\mathfrak{s}}$$

as maps from $HF^\circ(Y_0, V, \mathfrak{s}|_{Y_1}) \to HF^\circ(Y_0(\mathbb{L}), V_1, \mathfrak{s}|_{Y_0(\mathbb{L})})$, where $\mathcal{G}'$ is the induced map on $Y_0(\mathbb{L}_0)$. This follows since the free stabilization maps and the relative homology maps commute with the triangle maps (Theorem 2.5.17 and Lemma 2.4.4).

We thus have that

$$\begin{aligned}
&F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= \psi_f \circ F^{\circ}_{\mathbb{L},\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= \psi_f \circ A_{\widetilde{\Gamma}_0} \circ F^{\circ}_{\mathbb{L},\mathfrak{s}|} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ \psi_f \circ F^{\circ}_{\mathbb{L},\mathfrak{s}|} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ S^+_{V_Y} \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ \psi_f \circ F^{\circ}_{\mathbb{L},\mathfrak{s}|} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E),((\Phi\circ\phi)_*\mathbf{w}_1)\times I,\mathfrak{s}|} \\
&\qquad \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ \phi_*.
\end{aligned}$$

By Lemma 2.11.1, we have

$$A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^{\circ}_{W\setminus(N(\Gamma_0)\cup N(E),(\Phi\circ\phi)_*\mathbf{w}_1,\mathfrak{s}|} \circ F^{\circ}_{\mathfrak{H}_1} \circ \Phi_* \circ \phi_* = A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^{\circ}_{W\setminus N(E),(\Phi\circ\phi)_*\mathbf{w}_1,\mathfrak{s}|} \circ \Phi_* \circ \phi_*$$

Since $\Phi : Y \to Y\#\partial N(E)$ extends to an equivalence of $W \setminus N(E)$ with $W = Y \times I$ which is the identity on $Y \times \{1\}$, we have

$$\begin{aligned}
A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^{\circ}_{W\setminus N(E),(\Phi\circ\phi)_*\mathbf{w}_1\times I,\mathfrak{s}|} \circ \Phi_* \circ \phi_* &= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ F^{\circ}_{W,\phi_*\mathbf{w}_1\times I,\mathfrak{s}|} \circ \phi_* \\
&= A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ \phi_*.
\end{aligned}$$

At this point we are almost done. Note that in the above expression $A_{(\psi_f)_*\widetilde{\Gamma}_0}$ is the graph action map corresponding to a flow graph $\mathcal{G}'_f = ((\psi_f)_*\widetilde{\Gamma}_0, V_Y \cup \phi_*\mathbf{w}_1, \phi_*\mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$. Letting $\mathcal{G}_f = ((\psi_f)_*\widetilde{\Gamma}_0, \phi_*\mathbf{w}_1, \phi_*\mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$, by Lemma 2.6.10, we know that

$$\begin{aligned}
A_{(\psi_f)_*\widetilde{\Gamma}_0} \circ S^+_{\psi_f V_Y} \circ \phi_* &= A_{\mathcal{G}'_f} \circ S^+_{\psi_f V_Y} \circ \phi_* \\
&= A_{\mathcal{G}_f} \circ \phi_*.
\end{aligned}$$

It remains to see what the graph $(\psi_f)_*\widetilde{\Gamma}_0$ is. Note that the diffeomorphism ψ_f is induced by the flow of f. The flow of f preserves the basepoints $V_Y \cup \phi_*\mathbf{w}_1$ and the images of the edges of $\widetilde{\Gamma}_0$ are induced by the flow of f. Hence if e is an edge of Γ and e_Y is its projection, then e_Y, e and $\psi_f(e)$ are all isotopic via an isotopy which maps the endpoints of e into

$(\phi_*\mathbf{w}_1 \cup V_Y) \times I$. In particular the relative homology map in Y along the edge e_Y and $\psi_f(e)$ agree. If $\mathcal{G} = (\Gamma, \mathbf{w}_1, \mathbf{w}_2, \sigma, \varnothing)$, we thus have that

$$A_{\mathcal{G}_f} = A_{\phi_*\mathcal{G}},$$

and hence combining all of the computations we've done so far, we thus have that

$$\begin{aligned} F^\circ_{W,\Gamma,\mathfrak{s}} &= A_{\phi_*\mathcal{G}} \circ \phi_* \\ &= \phi_* \circ A_{\mathcal{G}} \\ &= A_{\mathcal{G}}. \end{aligned}$$

The last equality follows since ϕ is isotopic to the identity relative the basepoints $\mathbf{w}_2$. □

Lemma 2.11.4. *Suppose that* $(W, \Gamma') = (W_2, \boldsymbol{\gamma}) \circ (W_1, \Gamma)$ *where* $(W_2, \boldsymbol{\gamma}_2)$ *is a path cobordism. Then*

$$F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^\circ_{W,\Gamma',\mathfrak{s}},$$

where $F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2}$ *is the path cobordism map, and the remaining maps in the above expression are graph cobordism maps.*

Proof. This follows essentially from the construction of the graph cobordism maps and the composition law for path cobordisms (Corollary 2.9.3). Pick Γ_0, $N(\Gamma_0), \widetilde{\Gamma}_0, \phi_*, \widetilde{\boldsymbol{\gamma}}_f$ and $\mathfrak{H}$ as in the construction of the graph cobordism maps. Then by construction

$$F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma,\mathfrak{s}_1} = F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F_{W_1 \setminus N(\Gamma_0), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}_1|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}_1|_{Y_1}} \circ \phi_*.$$

Writing $(W', \boldsymbol{\gamma}')$ for $(W_2, \boldsymbol{\gamma}) \circ (W_1 \setminus N(\Gamma_0), \widetilde{\boldsymbol{\gamma}}_f)$ we see that

$$F^\circ_{W_2,\boldsymbol{\gamma},\mathfrak{s}_2} \circ F_{W_1 \setminus N(\Gamma_0), \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{s}_1|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}_1|_{Y_1}} \circ \phi_*$$

$$\left(\sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(W') \\ \mathfrak{s}|_{W_2} = \mathfrak{s}_2 \\ \mathfrak{s}|_{W_1 \setminus N(\Gamma)_0} = \mathfrak{s}_1|_{W_1 \setminus N(\Gamma_0)}}} F_{W',\boldsymbol{\gamma}',\mathfrak{s}} \right) \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}_1|_{Y_1}} \circ \phi_*,$$

but any Spin^c structure on W' extends uniquely over $N(\Gamma_0)$ to a Spin^c structure on W, since W is obtained by adding 3– and 4–handles to W'. Hence the above expression is equal to

$$\left(\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{W_1}=\mathfrak{s}_1}} F^\circ_{W',\boldsymbol{\gamma}',\mathfrak{s}|_{W'}} \right) \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|_{Y_1}} \circ \phi_*,$$

which by definition is equal to

$$\sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_2}=\mathfrak{s}_2 \\ \mathfrak{s}|_{W_1}=\mathfrak{s}_1}} F^\circ_{W,\Gamma',\mathfrak{s}},$$

as claimed. □

Lemma 2.11.5. *Suppose that $(W,\Gamma) : (Y_1,\mathbf{w}_1) \to (Y_2,\mathbf{w}_2)$ is a graph cobordism and suppose that Γ is isotopic (necessarily not neatly) to a graph of the form $\Gamma' \cup \boldsymbol{\gamma}$ where $\Gamma' \subseteq Y_1$ is a flow graph and $\boldsymbol{\gamma}$ is a collection of paths from the terminal basepoints of Γ' in Y_1 to the basepoints $\mathbf{w}_2 \subseteq Y_2$. Then*

$$F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} \circ A_{\Gamma'}.$$

Proof. Pick a collar neighborhood $N(Y_1) = Y_1 \times I$ inside of W. By assumption we can isotope Γ so that W is equivalent to the composition $(W,\boldsymbol{\gamma}) \circ (N(Y_1),\Gamma'')$ where Γ'' is a graph in $N(Y_1) \times I$ which projects to Γ'. By Lemmas 2.11.4 and 2.11.3 we know that

$$F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} \circ F^\circ_{N(Y_1),\Gamma'',\mathfrak{s}|_{N(Y_1)}} = F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}} \circ A_{\Gamma'},$$

as we wanted. □

Lemma 2.11.6. *Suppose $(W,\Gamma) : (Y_1,\mathbf{w}_1) \to (Y_2,\mathbf{w}_2)$ is a graph cobordism with $\pi_0(Y_i) \to \pi_0(W)$ surjective. Then there is a decomposition of W into*

$$(W,\Gamma) = (W_3,\boldsymbol{\gamma}_3) \circ (W_2,\boldsymbol{\gamma}_2) \circ (N,\Gamma') \circ (W_1,\boldsymbol{\gamma}_2),$$

where W_i has a Morse function with only index–i critical points and $N \cong Y' \times I$ is a cylindrical cobordism. Furthermore

$$F^\circ_{W,\Gamma,\mathfrak{s}} = F^\circ_{W_3,\boldsymbol{\gamma}_3,\mathfrak{s}|_{W_3}} \circ F^\circ_{W_2,\boldsymbol{\gamma}_2,\mathfrak{s}|_{W_2}} \circ A_{\Gamma'_{Y'}} \circ F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}.$$

Proof. First pick a Morse function f on W with critical points of increasing index and no index 0 or index 4 critical points, which is possible since $\pi_0(Y_i) \to \pi_0(W)$ is surjective for $i = 1, 2$. By picking appropriate level sets of f, we get a decomposition of W into cobordisms $W = W_3 \cup W_2 \cup N \cup W_1$ where $f|_{W_i}$ is a Morse function with only index i critical points and $f|_N$ has no critical points. Writing Y' for the common boundary between W_1 and N, we first note that $\pi_1(Y', q) \to \pi_1(W, q)$ is surjective for any choice of q. As a consequence, we can isotope Γ within W and perturb f slightly so that $\Gamma \cap W_i$ is a collection of paths which are flowlines of f.

Let W' denote $W_3 \cup W_2 \cup N$ and let (W', Γ') denote the graph cobordism $(W', \Gamma \cap W')$. Let the basepoints of $Y' = N \cap W_1$ be denoted by $\mathbf{w}'$. By Lemmas 2.11.4, 2.11.5 we see that

$$F^\circ_{W',\Gamma',\mathfrak{s}'} = F^\circ_{W_3,\boldsymbol{\gamma}_3,\mathfrak{s}|_{W_3}} \circ F^\circ_{W_2,\boldsymbol{\gamma}_2,\mathfrak{s}|_{W_2}} \circ A_{\Gamma'}.$$

Hence it is sufficient to show that

$$F^\circ_{W',\Gamma',\mathfrak{s}|_{W'}} \circ F_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}} = F^\circ_{W,\Gamma,\mathfrak{s}}.$$

For this, let E consist of all of the descending manifolds of the index 1 critical points of f inside of W. Let $\mathfrak{E}$ denote the induced handle decomposition of $N(E)$ with one 1–handle for each arc of E. The flow of f yields a diffeomorphism of $\Phi : (Y_1 \# \partial N(E), \mathbf{w}_1) \to (Y', \mathbf{w}')$. As in Lemma 2.11.1 the 1–handle map $F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}}$ satisfies

$$F^\circ_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}} = \Phi_* \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}}.$$

On the other hand, there is an equivalence

$$\widehat{\Phi} : (W \setminus N(E), \Gamma) \to (W', \Gamma')$$

extending Φ. Hence we have

$$\begin{aligned} F^\circ_{W',\Gamma',\mathfrak{s}'} \circ F_{W_1,\boldsymbol{\gamma}_1,\mathfrak{s}|_{W_1}} &= F^\circ_{W',\Gamma',\mathfrak{s}'} \circ \Phi_* \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} \\ &= F^\circ_{W \setminus N(E),\Gamma,\mathfrak{s}|_{W \setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} \end{aligned}$$

Let $\Gamma_0, \widetilde{\Gamma}_0, N(\Gamma_0), \phi, \widetilde{\boldsymbol{\gamma}}_f, \mathfrak{H}$ be chosen as in the construction of the graph cobordism maps for the graph cobordism $(W \setminus N(E), \Gamma)$. By construction

$$F^\circ_{W\setminus N(E),\Gamma,\mathfrak{s}|_{W\setminus N(E)}} = F^\circ_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|} \circ \phi_*.$$

Hence we have that

$$\begin{aligned} F^\circ_{W\setminus N(E),\Gamma,\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} &= F^\circ_{W\setminus(N(\Gamma_0)\cup N(E)),\widetilde{\boldsymbol{\gamma}}_f,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|} \circ \phi_* \circ F^\circ_{\mathfrak{E},\mathfrak{s}|_{Y_1}} \\ &= F^\circ_{W\setminus(N(E)\cup N(\Gamma_0)),\Gamma,\mathfrak{s}|_{W\setminus N(E)}} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|} \circ \phi_*. \end{aligned}$$

The last equality follows since the map $F^\circ_{\mathfrak{E},\mathfrak{s}|}$ commutes with 0–handle, 1–handle maps, the free stabilization maps, and the relative homology maps. It also commutes with ϕ_* (or more precisely a choice of two different ϕ maps, one on Y_1 and one on $Y_1 \# \partial N(E)$) because ϕ_* is fixed outside of a neighborhood of the basepoints, and the 1–handles in $\mathfrak{E}$ are attached away from the basepoints. Finally, by applying Lemma 2.11.1, we see that

$$\begin{aligned} F^\circ_{W\setminus(N(E)\cup N(\Gamma_0)),\Gamma,\mathfrak{s}|} \circ F^\circ_{\mathfrak{E},\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|} \circ \phi_* &= F^\circ_{W\setminus N(\Gamma_0)),\Gamma,\mathfrak{s}|} \circ A_{\widetilde{\Gamma}_0} \circ F^\circ_{\mathfrak{H},\mathfrak{s}|} \circ \phi_* \\ &= F^\circ_{W,\Gamma,\mathfrak{s}}, \end{aligned}$$

completing the proof. □

Lemma 2.11.7. *Suppose that $(W, \Gamma) = (W_2, \Gamma_2) \circ (W_1, \Gamma_1)$ is a decomposition of a graph cobordism. Writing $(W_i, \Gamma_i) : (Y_i, \mathbf{w}_i) \to (Y_i', \mathbf{w}_i)$, suppose that $\pi_0(Y_i) \to \pi_0(W_i)$ and $\pi_0(Y_i') \to \pi_0(W_i)$ are all surjective. Then*

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Proof. Using Lemma 2.11.6 to decompose each cobordism W_i into a composition

$$(W_{3i}, \boldsymbol{\gamma}_{3i}) \circ (W_{2i}, \boldsymbol{\gamma}_{2i}) \circ (N_i, \Gamma_i') \circ (W_{1i}, \boldsymbol{\gamma}_{1i}),$$

by the same Lemma we have that

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1}$$

$$= F^{\circ}_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}} \circ F^{\circ}_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}} \circ A_{\Gamma'_{2,\pi}} \circ F^{\circ}_{W_{12},\boldsymbol{\gamma}_{12},\mathfrak{s}_{12}} \circ F^{\circ}_{W_{31},\boldsymbol{\gamma}_{31},\mathfrak{s}_{31}} \circ F^{\circ}_{W_{21},\boldsymbol{\gamma}_{21},\mathfrak{s}_{21}|} \circ A_{\Gamma'_{1,\pi}} \circ F^{\circ}_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}},$$

where $\mathfrak{s}_{ij} = \mathfrak{s}_j|_{W_{ij}}$ and $A_{\Gamma'_{i,\pi}}$ denotes the projection of Γ'_i into the the incoming end of N_i given by the flow of the Morse function f.

We can assume that such a decomposition resulted from the level sets of a Morse function f_i on W_i, and assume further that f_i are the restrictions of a Morse function on $W = W_2 \cup W_1$. Also we can assume that $\mathfrak{s}_{12}|_{Y_2} = \mathfrak{s}_{31}|_{Y'_1}$ since otherwise the composition is zero, which satisfies the composition law.

The Morse function $f|_{W_{12}\cup W_{31}}$ has index 1 and index 3 critical points. By using a standard rearrangement theorem we can replace $f|_{W_{21}\cup W_{31}}$ with a Morse function f' which has only index 1 and index 3 critical points, but for which the index 1 critical points have lower value than the index 3 critical points. We can assume that f' still has $\boldsymbol{\gamma}_{12} * \boldsymbol{\gamma}_{31}$ as flowlines since the rearrangement can be assumed to modify the Morse function $f|_{W_{1,2}\cup W_{3,1}}$ only in a neighborhood of the ascending and descending manifolds of the critical points. This allows us to rewrite $(W_{12} \cup W_{31}, \boldsymbol{\gamma}_{12} * \boldsymbol{\gamma}_{31})$ as $(W'_{31} \cup W'_{12}, \boldsymbol{\gamma}'_{31} * \boldsymbol{\gamma}'_{12})$ where W'_{31} has a Morse function with only index 3 critical points and W'_{12} has a Morse function with only index 1 critical points.

Using the composition rule for path cobordisms (Corollary 2.9.3) for the cobordism $W_{12} \cup W_{31} = W'_{31} \cup W'_{12}$, we get that

$$F^{\circ}_{W_2,\Gamma_2,\mathfrak{s}_2} \circ F^{\circ}_{W_1,\Gamma_1,\mathfrak{s}_1}$$

$$= F^{\circ}_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}} \circ F^{\circ}_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}} \circ A_{\Gamma'_{2,\pi}} \circ F^{\circ}_{W'_{31},\boldsymbol{\gamma}'_{31},\mathfrak{s}'_{31}} \circ F^{\circ}_{W'_{12},\boldsymbol{\gamma}'_{12},\mathfrak{s}'_{12}} \circ F^{\circ}_{W_{21},\boldsymbol{\gamma}_{21},\mathfrak{s}_{21}} \circ A_{\Gamma'_{1,\pi}} \circ F^{\circ}_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}.$$

Similarly commuting the 1–handle cobordism $(W_{12}, \boldsymbol{\gamma}_{12})$ across W_{21} we get that the above expression is equal to

$$F^{\circ}_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}} \circ F^{\circ}_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}} \circ A_{\Gamma'_{2,\pi}} \circ F^{\circ}_{W'_{31},\boldsymbol{\gamma}'_{31},\mathfrak{s}'_{31}} \circ F^{\circ}_{W'_{21},\boldsymbol{\gamma}'_{21},\mathfrak{s}'_{21}} \circ F^{\circ}_{W''_{12},\boldsymbol{\gamma}''_{12},\mathfrak{s}''_{12}} \circ A_{\Gamma'_{1,\pi}} \circ F^{\circ}_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}$$

for cobordisms $(W''_{12}, \boldsymbol{\gamma}''_{12})$ and $(W'_{21}, \boldsymbol{\gamma}'_{21})$. Applying Lemma 2.11.2, we commute the graph action maps across the cobordism maps to see that the above expression is equal to

$$F^{\circ}_{W_{32},\boldsymbol{\gamma}_{32},\mathfrak{s}_{32}} \circ F^{\circ}_{W_{22},\boldsymbol{\gamma}_{22},\mathfrak{s}_{22}} \circ F^{\circ}_{W'_{31},\boldsymbol{\gamma}''_{31},\mathfrak{s}'_{31}} \circ F^{\circ}_{W'_{21},\boldsymbol{\gamma}''_{21},\mathfrak{s}'_{21}} \circ A_{\Gamma''_{2,\pi}} \circ A_{\Gamma''_{1,\pi}} \circ F^{\circ}_{W''_{12},\boldsymbol{\gamma}''_{12},\mathfrak{s}''_{12}} \circ F^{\circ}_{W_{11},\boldsymbol{\gamma}_{11},\mathfrak{s}_{11}}$$

for paths $\boldsymbol{\gamma}''_{31}, \boldsymbol{\gamma}''_{21}$ and $\boldsymbol{\gamma}''_{12}$ and graphs $\Gamma''_{2,\pi}$ and $\Gamma''_{1,\pi}$. By perturbing the Morse functions, graphs and paths slightly we may assume that $\Gamma''_{1,\pi} * \Gamma''_{2,\pi}$ is an embedded flow graph, where $*$ denotes concatenation at initial and terminal vertices. By construction the graph action map is functorial under concatenation, and hence $A_{\Gamma''_{2,\pi}} \circ A_{\Gamma''_{1,\pi}} = A_{\Gamma''_{2,\pi} * \Gamma''_{1,\pi}}$. We can also perform a critical point switch between $(W_{22}, \boldsymbol{\gamma}_{22})$ and $(W'_{31}, \boldsymbol{\gamma}''_{31})$ to decompose $(W_{22}, \boldsymbol{\gamma}_{22}) \circ (W'_{31}, \boldsymbol{\gamma}''_{31})$ as $(W''_{31}, \boldsymbol{\gamma}'''_{31}) \circ (W''_{22}, \boldsymbol{\gamma}_{22})$. Hence the above expression becomes

$$F^{\circ}_{W_{32}, \boldsymbol{\gamma}_{32}, \mathfrak{s}_{32}} \circ F^{\circ}_{W''_{31}, \boldsymbol{\gamma}'''_{31}, \mathfrak{s}''_{31}} \circ F^{\circ}_{W'_{22}, \boldsymbol{\gamma}'_{22}, \mathfrak{s}'_{22}} \circ F^{\circ}_{W'_{21}, \boldsymbol{\gamma}''_{21}, \mathfrak{s}'_{21}} \circ A_{\Gamma''_{2,\pi} * \Gamma''_{1,\pi}} \circ F^{\circ}_{W''_{12}, \boldsymbol{\gamma}''_{12}, \mathfrak{s}''_{12}} \circ F^{\circ}_{W_{11}, \boldsymbol{\gamma}_{11}, \mathfrak{s}_{11}}.$$

We now use the composition law for path cobordisms (Corollary 2.9.3) on the 2–handles, and see that this becomes

$$\sum_{\substack{\mathfrak{s}' \in \mathrm{Spin}^c(W'_{22} \circ W'_{21}) \\ \mathfrak{s}'|_{W_{2i}} = \mathfrak{s}'_{2i}}} F^{\circ}_{W_{32}, \boldsymbol{\gamma}_{32}, \mathfrak{s}_{32}} \circ F^{\circ}_{W''_{31}, \boldsymbol{\gamma}'''_{31}, \mathfrak{s}''_{31}} \circ F^{\circ}_{W'_{22} \cup W'_{21}, \boldsymbol{\gamma}'_{22} * \boldsymbol{\gamma}''_{21}, \mathfrak{s}'} \circ A_{\Gamma''_{2,\pi} * \Gamma''_{1,\pi}} \circ F^{\circ}_{W''_{12}, \boldsymbol{\gamma}''_{12}, \mathfrak{s}''_{12}} \circ F^{\circ}_{W_{11}, \boldsymbol{\gamma}_{11}, \mathfrak{s}_{11}}.$$

A Spin^c structure $\mathfrak{s}'$ on $W'_{22} \circ W'_{21}$ uniquely extends over the 1–handles and 3–handles to a Spin structure $\mathfrak{s}$ on $W_2 \circ W_1$. It is also easy to see that $\mathfrak{s}|_{W'_{2i}} = \mathfrak{s}'_{2i}$ for $i = 1, 2$ iff $\mathfrak{s}|_{W_{2i}} = \mathfrak{s}_i|_{W_{2i}}$ for $i = 1, 2$, which itself occurs iff $\mathfrak{s}|_{W_i} = \mathfrak{s}_i$.

Applying the composition rule for path cobordisms to the 1–handle maps and the 3–handle maps, we thus see that the above expression becomes

$$\sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^{\circ}_{W_{32} \cup W''_{31}, \boldsymbol{\gamma}_{32} * \boldsymbol{\gamma}''_{31}, \mathfrak{s}|} \circ F^{\circ}_{W'_{22} \cup W'_{21}, \boldsymbol{\gamma}'_{22} * \boldsymbol{\gamma}''_{21}, \mathfrak{s}|} \circ A_{\Gamma''_{2,\pi} * \Gamma''_{1,\pi}} \circ F^{\circ}_{W''_{12} \cup W_{11}, \boldsymbol{\gamma}''_{12} * \boldsymbol{\gamma}_{11}, \mathfrak{s}|},$$

which is exactly

$$\sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^{\circ}_{W, \Gamma, \mathfrak{s}}$$

by Lemma 2.11.6. □

2.12 Cobordisms with $\pi_0(Y_i) \not\twoheadrightarrow \pi_0(W)$

In this section we discuss graph cobordisms which have components with empty ends. The general idea is to remove balls and add paths arbitrarily connecting the new boundary components to the graph whenever a component of W has an empty incoming or outgoing end. The results of this section follow from a technical lemma which we prove first. We then proceed to construct the graph cobordism maps for arbitrary graph cobordisms, and then complete the proof of the graph composition law for arbitrary graph cobordisms. In this section we will finish the proofs of Theorems C and E.

2.12.1 Main technical lemma

The following is the main technical lemma of this section:

Lemma 2.12.1. *Suppose $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ is a graph cobordism with $\pi_0(Y_i) \to \pi_0(W)$ surjective, and suppose that $(W', \Gamma') : (Y_1, \mathbf{w}_1) \to (Y_2 \sqcup S^3, \mathbf{w}_2 \cup \{w_0\})$ denotes the cobordism obtained by removing a ball from W, and connecting the resulting S^3 arbitrarily with a path to the interior of an edge of Γ. Then*

$$F^\circ_{4,\mathfrak{s}|_{Y_2}} \circ F^\circ_{W',\Gamma',\mathfrak{s}|_{W'}} = F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Similarly if $(W'', \Gamma'') : (Y_1 \sqcup S^3, \mathbf{w}_1 \cup \{w_0\}) \to (Y_2, \mathbf{w}_2)$ is the graph cobordism resulting from removing a ball inside of W and connecting the basepoint in S^3 with a path arbitrarily to the interior of an edge of Γ, then

$$F^\circ_{W'',\Gamma'',\mathfrak{s}|_{W''}} \circ F^\circ_{0,\mathfrak{s}|_{Y_1 \sqcup S^3}} = F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Proof. We will only prove the statement about the cobordism (W', Γ') since the statement about (W'', Γ'') is proven nearly identically. Using Lemma 2.11.7 (the composition law) it is sufficient to show that the statement holds for a cobordism of the form $(W_0, \Gamma_0) : (Y, \mathbf{w} \cup \{w_0, w_1\}) \to (Y, \mathbf{w})$ where $W_0 = Y \times I$ and $\Gamma_0 = (\mathbf{w} \times I) \cup \widetilde{\lambda}$ where $\widetilde{\lambda}$ is a neatly embedded perturbation of a path λ from w_0 to w_1 in $Y \times \{0\}$. To see that this is sufficient

to show the claim for an arbitrary graph cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ with $\pi_0(Y_i) \to \pi_0(W)$ surjective, note that given any ball $B \subseteq W$ and any choice of path from ∂B to the interior of an edge in Γ, we can isotope that edge, as well as B and the path to the edge, so that they lie in a collar neighborhood of Y_2. Using the composition for such graph cobordisms (Lemma 2.11.7), we see that showing the claim for (W_0, Γ_0) implies the result for a general graph cobordism (W, Γ).

Let (W_0', Γ_0') be the cobordism obtained by removing a ball from W_0 and connecting a basepoint in the new boundary S^3 to a point on $\widetilde{\lambda}$. The cobordisms are depicted in Figure 2.18.

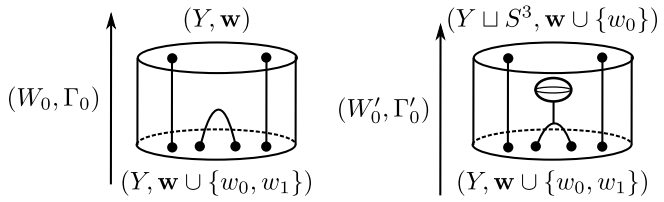

Figure 2.18: The cobordisms (W_0, Γ_0) and (W_0', Γ_0') in Lemma 2.12.1

We can isotope portions of the graphs Γ_0 and Γ_0' into the incoming ends of their respective cobordisms. Isotope Γ_0 into $Y \times \{0\}$ so that it consists of an arc λ in $Y \times \{0\}$. Isotope Γ_0' in W_0' so that it becomes λ with an additional edge, e added between a point on the interior of λ, to a vertex $(v, 0) \in Y \times \{0\}$, and then a path γ in W_0 from $(v, 0)$ to a point in S^3. We can assume that γ is a subset of $\{v\} \times I$. Let $(W_0', \boldsymbol{\gamma})$ be the path cobordism $(W_0', (\mathbf{w} \times I) \cup \gamma)$. Let $\mathcal{G} = (\lambda, \{w_0, w_1\}, \varnothing, \sigma, \mathbf{w})$ be the flow graph in Y consisting of the path λ. Let $\mathcal{G}'$ be the flow graph $(\lambda \cup e, \{w_0, w_1\}, \{v\}, \sigma, \mathbf{w})$ in $Y \times \{0\}$. By Lemma 2.11.5, we know that

$$F^\circ_{W_0, \Gamma_0, \mathfrak{s}} = F^\circ_{W_0, \mathbf{w} \times I} \circ A_{\mathcal{G}} = A_{\mathcal{G}}$$

and

$$F^\circ_{W_0', \Gamma_0', \mathfrak{s}|_{W_0'}} = F^\circ_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}} \circ A_{\mathcal{G}'}.$$

Hence the desired equality is equivalent to

$$A_{\mathcal{G}} = F^\circ_{4, \mathfrak{s}} \circ F^\circ_{W_0', \boldsymbol{\gamma}, \mathfrak{s}|_{W_0'}} \circ A_{\mathcal{G}'}.$$

By Lemma 2.6.11 we know that

$$A_{\mathcal{G}} = S_v^- \circ A_{\mathcal{G}'},$$

and hence it is sufficient to show that

$$F^\circ_{4,\mathfrak{s}} \circ F^\circ_{W_0',\boldsymbol{\gamma},\mathfrak{s}|_{W_0'}} = S_v^-.$$

To see this, put the Morse function $f(y,t) = t$ on $W_0 = Y \times I$. By introducing an index 3–4 critical point pair, and removing a ball B centered at the index 4 critical point, we get a Morse function f for $W_0' = W_0 \setminus B$ which has a single index 3 critical point. A sequence of diagrams encoding the composition is shown in Figure 2.19.

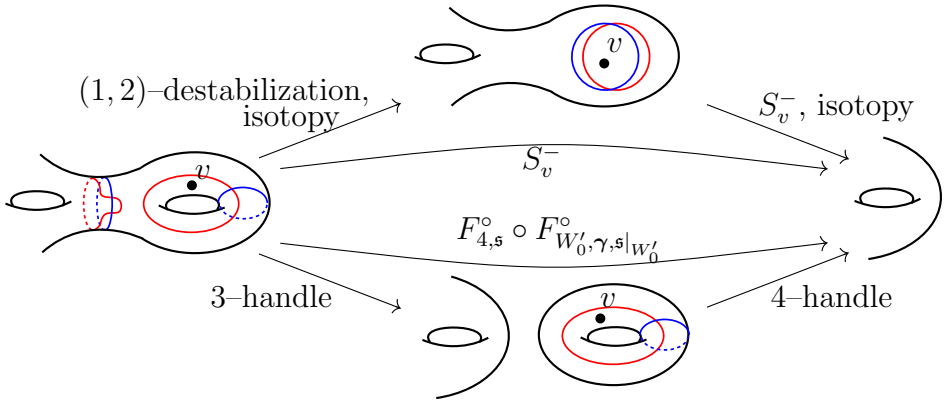

Figure 2.19: A sequence of diagrams which can be used to demonstrate that $F^\circ_{0,\mathfrak{s}} \circ F_{W_0',\boldsymbol{\gamma},\mathfrak{s}|_{W_0'}} = S_v^-$.

The path on top is S_v^-, and some change of diagrams maps. The path on the bottom is $F^\circ_{4,\mathfrak{s}} \circ F^\circ_{W_0',\boldsymbol{\gamma},\mathfrak{s}|_{W_0'}}$. Using the definitions of these maps, one sees immediately that the compositions are equal. □

2.12.2 Construction of graph cobordism maps for arbitrary cobordisms

We first define the complexes $CF^\circ(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P})$ by

$$\begin{aligned}
\widehat{CF}(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= \mathbb{Z}_2, \\
CF^-(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= \mathbb{Z}_2[U_\mathfrak{P}], \\
CF^\infty(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= \mathbb{Z}_2[U_\mathfrak{P}, U_\mathfrak{P}^{-1}] \\
CF^+(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}) &= CF^\infty(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P})/CF^-(\varnothing, \mathfrak{s}_\varnothing, \mathfrak{P}),
\end{aligned}$$

all with trivial differentials.

Given an arbitrary cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$, remove a balls $B_1, \dots, B_k$ from each component of W which doesn't contain a nonempty component of Y_1, and remove additional balls $B_1', \dots, B_n'$ from each component of W from each component of W which doesn't contain a component of Y_2. Let S_i denote the sphere ∂B_i, and let S_i' denote the sphere $\partial B_i'$. Let

$$(W_0, \Gamma_0, \mathfrak{s}_0) : (Y_1 \sqcup \coprod_i S_i, \mathbf{w}_1 \cup \{w_1, \dots, w_k\}) \to (Y_2 \sqcup \coprod_i S_i', \mathbf{w}_2 \cup \{w_1', \dots, w_n'\})$$

be the graph cobordism where

$$W_0 = W \setminus \left((\bigcup_i B_i) \cup (\bigcup_i B_i') \right).$$

and Γ_0 is the graph obtained by adjoining a path between the basepoint w_i in each S_i and points on Γ on the interior of an edge, and similarly adjoining a path between each of the points $w_i' \in S_i'$ and points on the interior of edges of Γ. The Spin^c structure $\mathfrak{s}_0$ is $\mathfrak{s}|_{W_0}$. We then define the graph cobordism map

$$F^\circ_{W,\Gamma,\mathfrak{s}} \stackrel{\text{def}}{=} \prod_{i=1}^n F^\circ_{4,S_i',\mathfrak{s}|} \circ F_{W_0,\Gamma_0,\mathfrak{s}_0} \circ \prod_{i=1}^k F^\circ_{0,S_i,\mathfrak{s}|},$$

where $\mathfrak{s}|$ denotes the appropriate restriction of $\mathfrak{s}$.

Lemma 2.12.2. *The map $F^\circ_{W,\Gamma,\mathfrak{s}}$ defined above is independent of the the choice of balls B_i and B_i', as well as the paths from points on the S_i and S_i' to Γ.*

Proof. Suppose we wish to replace one of the balls B_i and the path λ_i from $w_i \in \partial B_i$ to a point on Γ, with another ball $\widetilde{B_i}$ and path $\widetilde{\lambda}_i$. By diffeomorphism invariance of the cobordism map $F_{W_0,\Gamma_0,\mathfrak{s}_0}$, performing an isotopy of B_i or λ_i which fixes Γ leaves the cobordism map unchanged. Hence we can assume that $(B_i \cup \lambda_i)$ and $(\widetilde{B}_i \cup \widetilde{\lambda}_i)$ are disjoint. Let $(W_0, \Gamma_0, \mathfrak{s}_0)$ be the cobordism obtained by removing B_i and adding the path λ_i, and let $(\widetilde{W_0}, \widetilde{\Gamma}_0, \widetilde{\mathfrak{s}_0})$ be the graph cobordism obtained by removing $\widetilde{B}_i$ instead of B_i ad adding the path $\widetilde{\lambda}_i$. Let $(\overline{W}_0, \overline{\Gamma}_0, \overline{\mathfrak{s}}_0)$ denote the cobordism obtained by removing both B_i and $\widetilde{B_i}$ and adding both paths λ_i and $\widetilde{\lambda}_i$. Applying Lemma 2.12.1, we see that

$$\begin{aligned} F^\circ_{W_0,\Gamma_0,\mathfrak{s}_0} \circ F^\circ_{0,S_i,\mathfrak{s}|} &= F^\circ_{\overline{W}_0,\overline{\Gamma}_0,\bar{\mathfrak{s}}_0} \circ F^\circ_{0,\widetilde{S}_i,\mathfrak{s}|} \circ F_{0,S_i,\mathfrak{s}|} \\ &= F^\circ_{\overline{W}_0,\overline{\Gamma}_0,\bar{\mathfrak{s}}_0} \circ F_{0,S_i,\mathfrak{s}|} \circ F^\circ_{0,\widetilde{S}_i,\mathfrak{s}|} \\ &= F^\circ_{\widetilde{W}_0,\widetilde{\Gamma}_0,\widetilde{\mathfrak{s}}_0} \circ F^\circ_{0,\widetilde{S}_i,\mathfrak{s}|}, \end{aligned}$$

showing the the replacing B_i and λ_i with $\widetilde{B}_i$ with $\widetilde{\lambda}_i$ doesn't affect the graph cobordism map. The same argument also shows invariance from the B_i' balls and choice of paths. □

We can now prove Theorem C:

Proof of Theorem C. Invariance of the graph cobordism maps for arbitrary graph cobordisms now follows from Lemma 2.12.2 as well as invariance of the graph cobordism maps for graph cobordisms with $\pi_0(Y_i) \twoheadrightarrow \pi_0(W)$ surjective (Theorem 2.10.1). □

2.12.3 Composition law for arbitrary cobordisms

In this subsection we prove Theorem E, the composition law for arbitrary cobordisms:

Theorem E. *Suppose that* $(W_1, \Gamma_1) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ *and* $(W_2, \Gamma_2) : (Y_2, \mathbf{w}_2) \to (Y_3, \mathbf{w}_3)$ *are two graph cobordism maps and* $\mathfrak{s}_i \in \mathrm{Spin}^c(W_i)$ *are* Spin^c *structures. Writing* (W, Γ) *for*

the composition $(W_2, \Gamma_2) \circ (W_1, \Gamma_1)$*, we have the following composition law:*

$$F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2} \circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1} = \sum_{\substack{\mathfrak{s} \in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F^\circ_{W,\Gamma,\mathfrak{s}}.$$

Proof. Pick balls B_{1j} and B_{2j} to remove from W_1 and W_2 respectively so that each component of W_i contains a nonempty component of the incoming boundary. Similarly pick balls B'_{1j} and B'_{2j} in W_1 and W_2 respectively so that each component of W_i contains a nonempty component of the outgoing boundary. Let $S_{ij} = \partial B_{ij}$ and $S'_{ij} = \partial B'_{ij}$. Let W_{i0} be the manifold obtained by removing the balls B_{ij} and B'_{ij}. Pick paths from points on the new boundary spheres in W_{i0} to Γ_i. Let $\widetilde{\mathbf{w}}_i$ denote the set of basepoints in the spheres S_{ij} and let $\widetilde{\mathbf{w}}'_i$ denote the set of basepoints in the spheres S'_{ij}. Let (W_{i0}, Γ_{i0}) be the resulting graph cobordisms. Define a graph cobordism

$$(\widetilde{W}_{10}, \widetilde{\Gamma}_{10}) : (Y_1 \sqcup_i S_{1i} \sqcup_i S_{2i}, \mathbf{w}_1 \cup \widetilde{\mathbf{w}}_1 \cup \widetilde{\mathbf{w}}_2) \to (Y_2 \sqcup_i S'_{1i} \sqcup_i S_{2i}, \mathbf{w}_2 \cup \widetilde{\mathbf{w}}'_1 \cup \widetilde{\mathbf{w}}_2)$$

by

$$\widetilde{W}_{10} = W_{10} \sqcup_i (S_{2i} \times I)$$

and

$$\widetilde{\Gamma}_{10} = \Gamma_{10} \sqcup (\widetilde{\mathbf{w}}_2 \times I).$$

Similarly define a graph cobordism

$$(\widetilde{W}_{20}, \widetilde{\Gamma}_{20}) : (Y_2 \sqcup_i S_{2i} \sqcup_i S'_{1i}, \mathbf{w}_2 \cup \widetilde{\mathbf{w}}_2 \cup \widetilde{\mathbf{w}}'_1) \to (Y_3 \sqcup_i S'_{2i} \sqcup_i S'_{1i}, \mathbf{w}_3 \cup \widetilde{\mathbf{w}}'_2 \cup \widetilde{\mathbf{w}}'_1)$$

by

$$\widetilde{W}_{20} = W_{20} \sqcup_i (S'_{1i} \times I)$$

and

$$\widetilde{\Gamma}_{20} = \Gamma_{20} \sqcup (\widetilde{\mathbf{w}}'_1 \times I).$$

The cobordisms $(\widetilde{W}_{i0}, \Gamma_{i0})$ are shown in Figure 2.20.

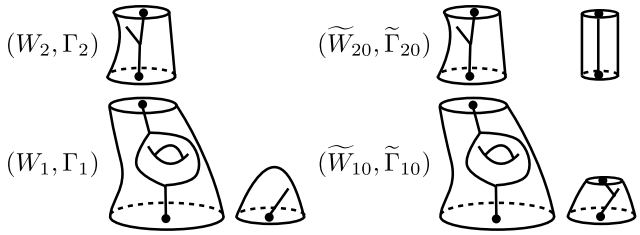

Figure 2.20: The cobordisms (W_i,Γ_i) and $(\widetilde{W}_{i0},\Gamma_{i0})$, as cobordisms from bottom to top.

We now observe that

$$\begin{aligned}
&F^\circ_{W_2,\Gamma_2,\mathfrak{s}_2}\circ F^\circ_{W_1,\Gamma_1,\mathfrak{s}_1}\\
&=\left(\prod_i F^\circ_{4,S'_{2i}}\right)F^\circ_{W_{20},\Gamma_{20},\mathfrak{s}_{20}}\left(\prod_i F^\circ_{0,S_{2i}}\right)\left(\prod_i F^\circ_{4,S'_{1i}}\right)F^\circ_{W_{10},\Gamma_{10},\mathfrak{s}_{10}}\left(\prod_i F^\circ_{0,S_{1i}}\right)\\
&=\left(\prod_i F^\circ_{4,S'_{2i}}\right)\left(\prod_i F^\circ_{4,S'_{1i}}\right)\left(F^\circ_{\widetilde{W}_{20},\widetilde{\Gamma}_{20},\widetilde{\mathfrak{s}}_{20}}\circ F^\circ_{\widetilde{W}_{10},\widetilde{\Gamma}_{10},\widetilde{\mathfrak{s}}_{10}}\right)\left(\prod_i F^\circ_{4,S_{2i}}\right)\left(\prod_i F^\circ_{4,S_{1i}}\right),
\end{aligned}$$

where $\mathfrak{s}_{i0}=\mathfrak{s}_i|_{W_{i0}}$ and $\widetilde{\mathfrak{s}}_{i0}$ is the unique Spin^c structure on $\widetilde{W}_{i0}$ which restricts to $\mathfrak{s}_{i0}$ on W_{i0}. The last equality follows since the 0– and 4–handle maps trivially commute with 1–,2– and 3–handle maps, free (de)stabilization maps, and the relative homology maps. Let $(\widetilde{W}_0,\widetilde{\Gamma}_0)$ denote the composition $(\widetilde{W}_{20},\widetilde{\Gamma}_{20})\circ(\widetilde{W}_{10},\widetilde{\Gamma}_{10})$. Using Lemma 2.11.7, we have that

$$\begin{aligned}
&\left(\prod_i F^\circ_{4,S'_{2i}}\right)\left(\prod_i F^\circ_{4,S'_{1i}}\right)\left(F^\circ_{\widetilde{W}_{20},\widetilde{\Gamma}_{20},\widetilde{\mathfrak{s}}_{20}}\circ F^\circ_{\widetilde{W}_{10},\widetilde{\Gamma}_{10},\widetilde{\mathfrak{s}}_{10}}\right)\left(\prod_i F^\circ_{4,S_{2i}}\right)\left(\prod_i F^\circ_{4,S_{1i}}\right)\\
&=\left(\prod_i F^\circ_{4,S'_{2i}}\right)\left(\prod_i F^\circ_{4,S'_{1i}}\right)\left(\sum_{\substack{\mathfrak{s}_0\in\mathrm{Spin}^c(\widetilde{W}_0)\\ \mathfrak{s}_0|_{\widetilde{W}_{i0}}=\widetilde{\mathfrak{s}}_{i0}}}F^\circ_{\widetilde{W}_0,\widetilde{\Gamma}_0,\mathfrak{s}_0}\right)\left(\prod_i F^\circ_{4,S_{2i}}\right)\left(\prod_i F^\circ_{4,S_{1i}}\right),
\end{aligned}$$

which by definition is just

$$\sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(\widetilde{W})\\ \mathfrak{s}|_{\widetilde{W}_i}=\mathfrak{s}_i}}F^\circ_{W,\Gamma,\mathfrak{s}},$$

since $\widetilde{W}_0$ is obtained by removing balls from W and a Spin^c structure on $\widetilde{W}_0$ uniquely extends over the 4–balls to a Spin^c structure on W. □

2.13 Moving basepoints, the π_1–action, and the relation between the graph and path cobordism maps

In this section we will prove Theorem D, showing that the graph cobordism maps agree with the path cobordism maps when a graph is a collection of paths. We also prove Theorem F by providing an explicit description of the chain homotopy type of the π_1–action. Using this we prove Theorem G by showing that for path cobordisms where each path has a distinct color, the cobordism maps on homology do not depend on the choice of paths.

The main result is a computation of the equivariant chain homotopy type of the diffeomorphism map obtained by moving a basepoint $w \in \mathbf{w}$ along a path λ to another point $w' \notin \mathbf{w}$. We call this diffeomorphism the **basepoint moving map**, and we denote the induced map by λ_*. In Theorem 2.13.13, we show that the map λ_* is equivariantly chain homotopic to $S^-_w A_\lambda S^+_{w'}$. For this, we will need a transition map computation similar to a computation for $\widehat{HF}$ found in [Zem15, Lemma 8.1]. For a closed loop γ, by writing the map γ_* as a composition of two basepoint moving maps along arcs, an easy algebraic computation will yield the formula for γ_*.

Finally, in Subsection 2.13.3 we refine some of these techniques and give an alternate proof of the formula for the π_1–action, by showing that for multi-pointed diagrams we can explicitly describe all of the holomorphic curves counted in a sequence of change of diagrams maps associated to the π_1–action, and see how the formula appears.

2.13.1 A transition map computation

For a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{z\})$ we consider two diagrams,

$$\mathcal{H}_1 = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_1\}, \boldsymbol{\beta} \cup \{\beta_1\}, \mathbf{w}_0 \cup \{w, w'\}) \quad \text{and} \quad \mathcal{H}_2 = (\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_2\}, \boldsymbol{\beta}' \cup \{\beta_2\}, \mathbf{w}_0 \cup \{w, w'\}),$$

obtained by removing the basepoint z, and inserting the diagrams shown in Figure 2.21 into the region which contained z. The curves $\boldsymbol{\alpha}'$ and $\boldsymbol{\beta}'$ are small Hamiltonian isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. The case that $\mathbf{w}_0 = \varnothing$ is not excluded. We will need to compute the transition map

between the two diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$.

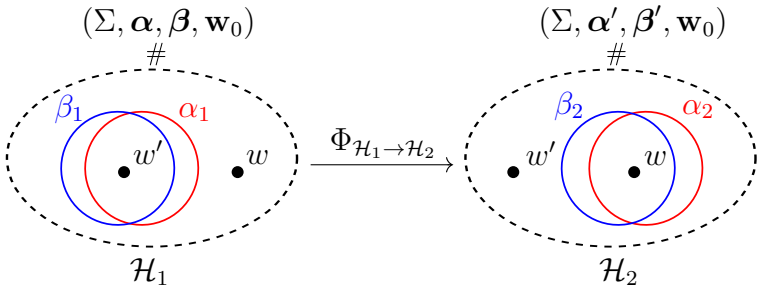

Figure 2.21: We compute the above transition map, for appropriate choices of almost complex structures, in Theorem 2.13.1.

Theorem 2.13.1. *For the diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ in Figure 2.21, with $\boldsymbol{\alpha}'$ and $\boldsymbol{\beta}'$ small Hamiltonian isotopies of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, there are choices of almost complex structure J_i on $\mathcal{H}_i$ which can be used to compute $S^{\pm}_{w'}$ on $\mathcal{H}_1$ and $S^{\pm}_{w}$ on $\mathcal{H}_2$ (respectively), such that the transition map $\Phi_{(\mathcal{H}_1,J_1)\to(\mathcal{H}_2,J_2)}$ is chain homotopic to a map of the form*

$$\Phi_{(\mathcal{H}_1,J_1)\to(\mathcal{H}_2,J_2)} = \begin{pmatrix} (\Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'})_{w'} \circ (\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}})_{w} & 0 \\ * & (\Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'})_{w'} \circ (\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}})_{w} \end{pmatrix}.$$

Here $(\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}})_{w}$ denotes the triangle map computing the $\boldsymbol{\beta}$–isotopy on the unstabilized diagram, with the basepoint w placed in the region of the unstabilized diagram where curves α_i and β_i are placed. The map $(\Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'})_{w'}$ is defined similarly, though with the basepoint w' placed inside.

Remark 2.13.2. In fact, we can explicitly identify the component marked with a $*$ in the above proposition. We can take it to be

$$(\Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'})_{w'} \circ \left(\sum_{i,j\geq 0} U_w^i U_{w'}^j \partial_{i+j+1} \right) \circ (\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}})_{w},$$

though this component isn't necessary for our computation of the diffeomorphism map λ_*. Here the term ∂_{i+j+1} is obtained from the differential on the unstabilized diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{z\})$ (where z is a basepoint added in the region where we are adding w, w', α_i

and β_i) by writing

$$\partial = \sum_{i=0}^{\infty} \partial_i U_z^i.$$

Note that if we color so that $U_w = U_{w'}$, then the map appearing in the middle becomes $\sum_{n=1}^{\infty} n\partial_n U^n$, where $U = U_w = U_{w'}$, which makes an appearance later as the map Φ_z.

The proof involves counting holomorphic disks and triangles on diagrams where the almost complex structure has been degenerated. The degenerations we use are similar to the "quasi-stabilization" degeneration considered in [MO10] and [Zem16]. We introduce an intermediate diagram, $\mathcal{H}_{1.5}$, for the computation, shown in Figure 2.22. The computation of Theorem 2.13.1 will consist of four main computations:

1. An explicit description of the holomorphic disks appearing on $\mathcal{H}_1$, $\mathcal{H}_{1.5}$ and $\mathcal{H}_2$;
2. A computation of a triangle map for the transition map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_{1.5}}$;
3. A computation of a change of almost complex structure map on $\mathcal{H}_{1.5}$;
4. A computation of a triangle map for the transition map $\Phi_{\mathcal{H}_{1.5} \to \mathcal{H}_2}$.

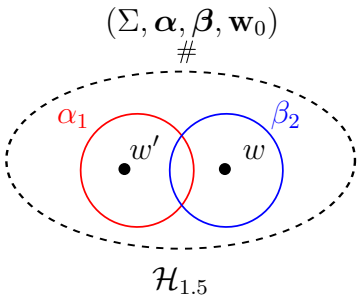

Figure 2.22: The intermediate diagram $\mathcal{H}_{1.5}$.

These holomorphic curve computations are similar to the differential computation of Lemma 2.5.5, and the triangle map computations of Theorem 2.5.17, where we degenerated the almost complex structure by stretching along a circle in the surface. In this section, we will need to do a more refined degeneration.

On the diagrams $\mathcal{H}_1$, $\mathcal{H}_{1.5}$, and $\mathcal{H}_2$ we will write J_α for an almost complex structure which is stretched along the circles c and c_α in Figure 2.23. On those diagrams, we will write J_β for an almost complex structure which is stretched along the circles c and c_β in Figure 2.23.

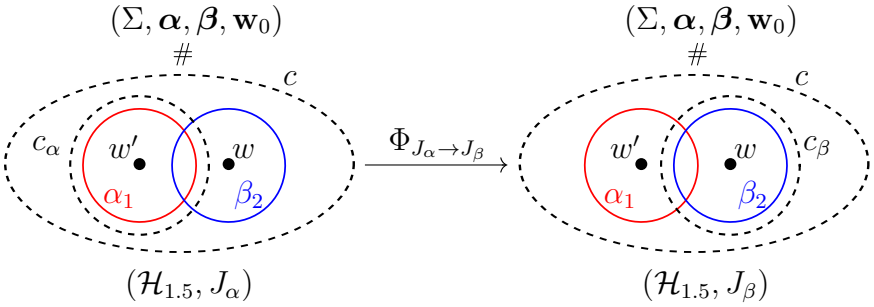

Figure 2.23: The almost complex structures J_α and J_β, shown in $\mathcal{H}_{1.5}$.

We begin with a refinement of the differential computation of Lemma 2.5.5. Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{z\})$ be the unstabilized diagram, as before. Let ∂_w be the differential on the unstabilized diagram if we replace z with w, and let $\partial_{w'}$ be the differential if we replace z with w'. The holomorphic curve counts going into the two differentials ∂_w and $\partial_{w'}$ are identical, but $\partial_{w'}$ has a $U_{w'}$ everywhere that ∂_w has a U_w.

Lemma 2.13.3. *Let ∂_w and $\partial_{w'}$ be the differentials on the unstabilized diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w'\})$ respectively. For the almost complex structure J_α, stretched appropriately along both c and c_α, the differential on the diagram $\mathcal{H}_{1.5}$ with almost complex structure J_α takes the form*

$$\partial_{\mathcal{H}_{1.5}, J_\alpha} = \begin{pmatrix} \partial_w & U_w + U_{w'} \\ 0 & \partial_w \end{pmatrix}$$

and the differential on $\mathcal{H}_{1.5}$ with J_β, stretched appropriately on both c and c_β, takes the form

$$\partial_{\mathcal{H}_{1.5}, J_\beta} = \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix}.$$

Proof. We note that the differential on $\mathcal{H}_{1.5}$ counts essentially the same disks as on $\mathcal{H}_1$ and $\mathcal{H}_2$, but on $\mathcal{H}_{1.5}$ the basepoint configuration is different, resulting in possibly a different

differential than on $\mathcal{H}_1$ or $\mathcal{H}_2$, as the powers of U_w and $U_{w'}$ may now be different. Let m_1, m_2, n_1 and n_2 denote the multiplicities of the regions of $(S^2, \alpha_2, \beta_1, w, w')$ as shown in Figure 2.24.

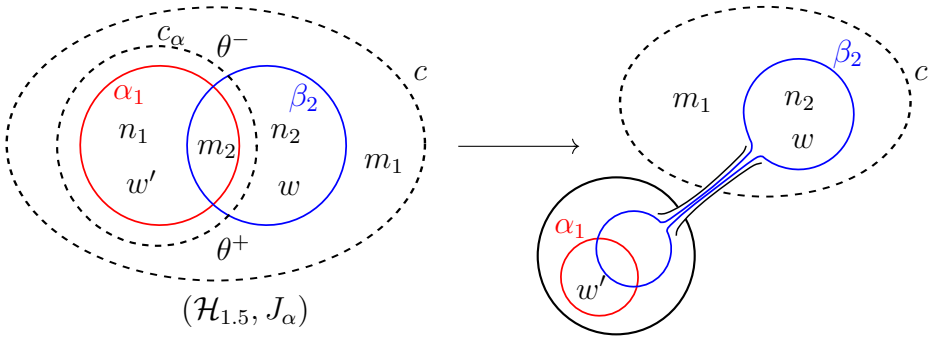

Figure 2.24: On the left are multiplicities on the diagram $(S^2, \alpha_1, \beta_2, w, w')$ used to analyze disks in $\mathcal{H}_{1.5}$, as well as the circles c and c_α. On the right, we see how the almost complex structure degenerates as we stretch J_α along c_α.

Let us consider $\partial_{\mathcal{H}_{1.5}, J_\alpha}$, as the computation for J_β is a straightforward modification. Suppose $\phi = \phi_\Sigma \# \phi_0$ is a homology disk on $\mathcal{H}_{1.5}$, with $\phi_\Sigma \in \pi_2(\mathbf{x}, \mathbf{y})$ a homology disk on the unstabilized diagram, and $\phi_0 \in \pi_2(x, y)$ a homology disk on (S^2, α_1, β_2). We note that the Maslov index computation of Lemma 2.5.3 is easily adapted to show that

$$\mu(\phi) = \mu(\phi_\Sigma) + 2m_2 - \Delta(x, y),$$

where $\Delta(x, y)$ is the drop in relative Maslov grading from x to y. Write

$$\partial_{\mathcal{H}_{1.5}, J_\alpha} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

Let us first consider the component B. These correspond to curves with $\Delta(x, y) = -1$. By stretching sufficiently, we can assume that $\mu(\phi_\Sigma) \geq 0$. The above Maslov index computation then implies that $\mu(\phi_\Sigma) = m_2 = 0$, so that ϕ_Σ is the constant disk. The only remaining possibility is that $\phi_\Sigma \# \phi_0$ consists of one of the two bigons going over w or w', each of which has a unique holomorphic representative. Hence the upper right component is $U_w + U_{w'}$.

We now consider the components A and D. These correspond to $\Delta(x,y) = 0$. In this case, the Maslov index computation yields that $\mu(\phi_\Sigma) = 1$. By the argument in the proof of Lemma 2.5.5, if we fix a neck length along c_α, then for almost complex structure sufficiently stretched along c, the moduli space $\mathcal{M}(\phi_\Sigma \# \phi_0)$ can be identified with the fibered product

$$\{(u_\Sigma, u_0) \in \mathcal{M}(\phi_\Sigma) \times \mathcal{M}(\phi_0) : \rho^p(u_\Sigma) = \rho^{p_0}(u_0)\},$$

where p and p_0 are the connected sum points on Σ and S^2, respectively, arising when we collapse the circle c by stretching the almost complex structure. On the other hand, since $\widehat{\mathcal{M}}(\phi_\Sigma)$ is a discrete space, it is sufficient to count $\mathcal{M}(\phi_0, \mathbf{d})$, the set of holomorphic representatives of ϕ_0 which match a generic $\mathbf{d} \in \mathrm{Sym}^{m_1}(\mathbb{D})$. By the proof of Lemma 2.5.5, we know that

$$\sum_{\substack{\phi_0 \in \pi_2(x,x) \\ m_2(\phi_0)=0}} \#\mathcal{M}(\phi_0, \mathbf{d}) \equiv 1 \pmod 2.$$

Note however, that for such a ϕ_0, we have $n_1 + n_2 = m_1$, so the total multiplicity over w and w' of $\phi_\Sigma \# \phi_0$ is equal to $m_1(\phi_\Sigma)$, though we don't immediately know how multiplicity is spread out between n_1 and n_2.

To this end we make the following claim: if $\mathbf{d}$ is a fixed divisor in $\mathrm{Sym}^{m_1}(\mathbb{D})$, then by stretching the almost complex structure on S^2 along c_α, we can ensure that $n_2(\phi_0) \geq m_1(\phi_0)$ for any homology disk $\phi_0 \in \pi_2(x,x)$ with $m_2(\phi_0) = 0$ and $\mathcal{M}(\phi_0, \mathbf{d}) \neq \varnothing$. Assuming this for the moment, since $m_1(\phi_0) = n_1(\phi_0) + n_2(\phi_0)$ for these disks (as $m_2(\phi_0) = 0$), it would follow that $n_2(\phi_0) = m_1(\phi_0)$ and $n_1(\phi_0) = 0$. As a consequence, it would follow that any holomorphic disk counted by A or D would go only over w, and have zero multiplicity over w'. The stated identification of the diagonal entries of the differential would then follow.

To show the stated claim in the previous paragraph, we let $J^0_\alpha(T)$ denote an almost complex structure on $S^2 \times [0,1] \times \mathbb{R}$ obtained by inserting neck length T along the curve c_α. We note that as we increase T, we can extract a weak limit of curves in the space $\mathcal{M}_{J^0_\alpha(T)}(\phi_0, \mathbf{d})$ to broken curves on the two diagrams (S^2, α_1, β_2) and (S^2, β_2), as in Figure 2.24. As the curves in $\mathcal{M}_{J^0_\alpha(T)}(\phi_0, \mathbf{d})$ all match $\mathbf{d}$, we must have a component of the broken curve on (S^2, β_2) which also matches $\mathbf{d}$. Since (S^2, β_2) has no $\boldsymbol{\alpha}$ curves, however, the maximum principle

applied to $\pi_{\mathbb{D}} \circ u$ implies that the only possible curve which could match $\mathbf{d}$ would be a collection of closed surfaces mapped into subsets $S^2 \times \{d\}$ for the various $d \in \mathbf{d}$. Such a closed curve contributes to both m_1 and n_2 equally, and since there must be at least m_1 of them, we conclude that $n_2 \geq m_1$, completing the proof of the subclaim, and hence the computation of the diagonal entries of the differential.

We now consider the component C of the differential. As $\Delta(x, y) = -1$ in this case, we see that m_2 is constrained to be either 0 or 1. If $m_2 = 1$, then $\mu(\phi_\Sigma) = 0$ and ϕ is just the bigon $m_2 = 1$ with domain supported on (S^2, α_1, β_2) (which has a unique holomorphic representative). If $m_2 = 0$, then the only remaining possibility is that $\mu(\phi_\Sigma) = 2$ and $\Delta(x, y) = -1$. We claim that for sufficiently stretched almost complex structure along c_α, the only possibility is that ϕ_Σ has the homology class of a Maslov index 2 β-boundary degeneration or $[\Sigma]$ (if the latter has Maslov index 2). This can be seen by observing that as we stretch along the two curves, a curve in the limit representing ϕ_Σ must satisfy a matching condition with a broken curve on (S^2, β_2). If there is a legitimate flow line in the limit of the broken curve representing ϕ_Σ, then the matching component on (S^2, β_2) consists of a collection of closed surfaces, each contributing 1 to both m_1 and n_2. On the other hand, the vertex multiplicities imply that

$$m_1 = n_1 + n_2 + 1 > n_2,$$

so in particular the flow line in the broken limit of ϕ_Σ cannot represent the entire homology class. Hence there must be either a boundary degeneration or a closed surface in the weak limit representing ϕ_Σ. But any such curve has Maslov index at least 2, which prohibits the existence of any other curves appearing in the limit, since such an additional curve would raise the Maslov index of ϕ_Σ strictly above 2.

Hence $\pi_2(\mathbf{x} \times \theta^+, \mathbf{y} \times \theta^-)$ has no Maslov index 1 classes with holomorphic representatives unless $\mathbf{x} = \mathbf{y}$. If $\mathbf{x} = \mathbf{y}$, there are exactly two classes which might have representatives. One is the bigon with $m_2 = 1$ and all other multiplicities equal to zero. The other has $m_1 = 1$ and has ϕ_Σ equal to a Maslov index 2 $\boldsymbol{\beta}$–degeneration or $\phi_\Sigma = [\Sigma]$. The bigon with $m_2 = 1$

clearly has a unique holomorphic representative. We now show how to compute that count of holomorphic representatives of the homology class with $m_2 = 0$ is $1 \in \mathbb{Z}_2$. To see this, let $\phi \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{x} \times \theta^-)$ denote this homology class with $m_1 = 1$ and $m_2 = 0$. Let ϕ' denote the Maslov index 1 bigon going over w'. Splicing the two homology classes together, we get the Maslov index 2 homology class $B = \phi' * \phi \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{x} \times \theta^+)$. The moduli space $\widehat{\mathcal{M}}(B)$ is 1–dimensional, so we count it's ends. One end corresponds to $\widehat{\mathcal{N}}^\beta(B)$, which has (modulo 2) a unique representative by [OS08, Theorem 5.5]. The other ends correspond to strip breaking. Due to our degeneration of the almost complex structure, our previous argument implies that a Maslov index one holomorphic disk with $n_1 = 1$ has domain equal to the bigon ϕ'. This forces the other holomorphic disk appearing in end of $\widehat{\mathcal{M}}(A)$ to have domain $A - \mathcal{D}(\phi') = \mathcal{D}(\phi)$. In particular, the only end of $\widehat{\mathcal{M}}(A)$ resulting from strip breaking is the end corresponding to A breaking into ϕ and ϕ'. Hence, summing over the ends of $\widehat{\mathcal{M}}(B)$ we get exactly that

$$\#\widehat{\mathcal{N}}^\beta(B) + \#\widehat{\mathcal{M}}(\phi) \cdot \#\widehat{\mathcal{M}}(\phi') = 0.$$

Since $\#\widehat{\mathcal{N}}^\beta(B) = \#\widehat{\mathcal{M}}(\phi') = 1$, we conclude that $\#\widehat{\mathcal{M}}(\phi) = 1$. □

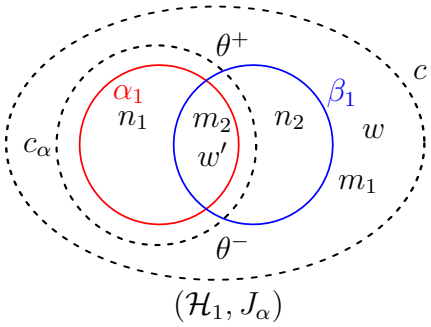

Figure 2.25: Multiplicities on $\mathcal{H}_1$.

Similarly, we can compute all of the curves computed by the differentials on $\mathcal{H}_1$ and $\mathcal{H}_2$:

Lemma 2.13.4. *On $\mathcal{H}_1$, for both J_α and J_β, the free-stabilized differential takes the form*

$$\partial_{\mathcal{H}_1} = \begin{pmatrix} \partial_w & U_w + U_{w'} \\ 0 & \partial_w \end{pmatrix}.$$

Furthermore, using the multiplicities from Figure 2.25 we have the following more precise description of the Maslov index 1 holomorphic disks which appear for the almost complex structure J_α:

1. *All of the Maslov index 1 holomorphic disks representing classes in $\pi_2(\mathbf{x}\times\theta^+,\mathbf{y}\times\theta^+)$ and $\pi_2(\mathbf{x}\times\theta^-,\mathbf{y}\times\theta^-)$ have $n_1=m_2=0$ and $n_2=m_1$;*

2. *There are exactly two domains which have holomorphic representatives in any $\pi_2(\mathbf{x}\times\theta^+,\mathbf{y}\times\theta^-)$. For such disks, we have $\mathbf{x}=\mathbf{y}$. The domains are the bigons with exactly one of n_1 or n_2 equal to 1, and all other multiplicities equal to zero.*

3. *There are exactly two domains which have holomorphic representatives in any $\pi_2(\mathbf{x}\times\theta^-,\mathbf{y}\times\theta^+)$. For such disks, we have $\mathbf{x}=\mathbf{y}$. One is the bigon with $m_2=1$. The other is the bigon on (S^2,α_1,β_1) with $m_1=1$ glued to a class of Maslov index 2 $\boldsymbol{\beta}$-degeneration (or the class $[\Sigma]$) with $m_1=1$. For any class $\phi\in\pi_2(\mathbf{x}\times\theta^-,\mathbf{x}\times\theta^+)$ with one of those domains, there is one representative (modulo 2).*

A similar statement holds for the diagram $\mathcal{H}_2$, or for the almost complex structure J_β.

Proof. The proof is a straightforward modification of the proof of Lemma 2.13.3. In fact, the same holomorphic curves are counted for $\partial_{\mathcal{H}_1,J_\alpha}$ and $\partial_{\mathcal{H}_{1.5},J_\alpha}$. One needs only to figure out which components of the differential they contribute to, and what multiplicity they have over the various basepoints. □

We can also need to compute several components of the change of almost complex structure map $\Phi_{J_\alpha\to J_\beta}$ on $\mathcal{H}_{1.5}$.

Lemma 2.13.5. *For sufficiently stretched almost complex structures along c, c_α and c_β, the change of almost complex structure map $\Phi_{J_\alpha\to J_\beta}$ on $\mathcal{H}_{1.5}$ takes the form*

$$\Phi_{J_\alpha\to J_\beta}=\begin{pmatrix}\mathrm{id} & 0\\ * & \mathrm{id}\end{pmatrix}.$$

Proof. We pick a dynamic almost complex structure $\widetilde{J}$ between J_α and J_β (with all three depending on a parameter T of neck length along c). As we stretch T, all three almost complex structures degenerate to ones on $(\Sigma \vee S^2) \times [0,1] \times \mathbb{R}$, though we can assume that the limiting almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$ is a cylindrical almost complex structure achieving transversality. If we take a homology disk $\phi_\Sigma \# \phi_0$, with multiplicities in the left side of Figure 2.24, then we can write

$$\mu(\phi) = \mu(\phi_\Sigma) + 2m_2(\phi) - \Delta(x,y).$$

For sufficiently stretched almost complex structure, ϕ_Σ has a broken representative, so $\mu(\phi_\Sigma)$ is nonnegative, so in particular $\mu(\phi_\Sigma)$ is forced to be 0 or 1. In the lemma statement, we are considering only disks with $\Delta(x,y) \leq 0$. If $\Delta(x,y) = -1$, then we must have $\mu(\phi_\Sigma) = -1$, so by transversality no disks would be counted, so the upper right component of the transition matrix is zero. If $\Delta(x,y) = 0$ (corresponding to the diagonal entries), then $\mu(\phi_\Sigma) = 0$, but this forces ϕ_Σ to be the constant disk, as the almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$ is cylindrical and achieves transversality. □

Remark 2.13.6. We can actually compute that the entry $(*)$ is equal to $\sum_{i,j \geq 0} U_w^i U_{w'}^j \partial_{i+j+1}$, where ∂_{i+j+1} is obtained by taking the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{z\})$, (where z is in the region where we add α_i, β_i, w and w') and writing the differential as

$$\partial = \sum_{k=0}^{\infty} \partial_k U_z^k.$$

To see this, write C for the component marked $*$ in the change of almost complex structures map. Using Lemma 2.13.3, and the fact that the change of almost complex structure map $\Phi_{J_\alpha \to J_\beta}$ is a chain map with respect to the differentials $\partial_{\mathcal{H}_{1.5}, J_\alpha}$ and $\partial_{\mathcal{H}_{1.5}, J_\beta}$, we see that

$$\partial_w + (U_w + U_{w'})C = \partial_{w'}.$$

It's easy to see that this algebraically forces C to be the expression we wrote above. In Proposition 2.13.22 of the next subsection, we will refine this argument, and see a way to count the curves in a more direct manner.

Lemma 2.13.7. *Suppose that* $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w})$ *is a diagram such that the curves in* $\boldsymbol{\alpha}'$ *are small isotopies of the curves in* $\boldsymbol{\alpha}$*, satisfying* $|\alpha_i' \cap \alpha_j| = 2\delta_{ij}$*. If* $\phi \in \pi_2(\Theta^+_{\alpha'\alpha}, \mathbf{y})$ *is a nonnegative homology disk, then* $\mu(\phi) - n_{w_0}(\phi) \geq 0$ *for any* $w_0 \in \mathbf{w}$*, with equality iff* ϕ *is the constant disk. Here* $\Theta^+_{\alpha'\alpha}$ *denotes the top graded intersection point of* $\mathbb{T}_{\alpha'} \cap \mathbb{T}_{\alpha}$*.*

Proof. By the formula for the relative Maslov grading in [OS04e], we have

$$\mathrm{gr}(\Theta^+_{\alpha'\alpha}, \mathbf{y}) = \mu(\phi) - 2\sum_{w \in \mathbf{w}} n_w(\phi) \geq 0.$$

Hence

$$\mu(\phi) - n_{w_0}(\phi) \geq n_{w_0}(\phi) + 2\sum_{w \neq w_0} n_w(\phi).$$

Hence $\mu(\phi) - n_{w_0}(\phi) \geq 0$, and we see that equality occurs iff $n_w(\phi) = 0$ for all $w \in \mathbf{w}$ and $\mu(\phi) = 0$. However it's easy to see that the only nonnegative disk in any $\pi_2(\Theta^+_{\alpha'\alpha}, \mathbf{y})$ with $\mu(\phi) = 0$ and $n_w(\phi) = 0$ for all $w \in \mathbf{w}$ is the constant disk. □

We now perform the first of our triangle map computations, for a triple which has been stabilized as in Figure 2.26.

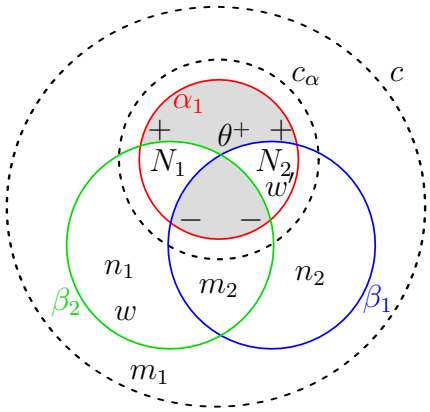

Figure 2.26: The triple $\overline{\mathcal{T}}$ in Proposition 2.13.8, a stabilization of $\mathcal{T}$. We have included some multiplicities. The shaded regions are each examples of small triangles, which might be counted. This is used to compute the transition map for moving β_1 to β_2.

Proposition 2.13.8. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}_0 \cup \{w\})$ is a Heegaard triple with $\boldsymbol{\beta}'$ small isotopies of the $\boldsymbol{\beta}$ curves. Let $\overline{\mathcal{T}}$ be the Heegaard triple $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_1\}, \boldsymbol{\beta} \cup \{\beta_1\}, \boldsymbol{\beta}' \cup \{\beta_2\}, \mathbf{w} \cup \{w, w'\})$ shown in Figure 2.26. The triangle map associated to the triple $\overline{\mathcal{T}}$ satisfy*

$$F_{\overline{\mathcal{T}}}(\mathbf{x} \times \cdot, \Theta^+_{\beta\beta'} \times \theta^+) = \begin{pmatrix} F_{\mathcal{T}}(\mathbf{x}, \Theta^+_{\beta\beta'}) & 0 \\ 0 & F_{\mathcal{T}}(\mathbf{x}, \Theta^+_{\beta\beta'}) \end{pmatrix}$$

with respect to the Maslov grading on the intersection points of $\alpha_1 \cap \beta_1$ and $\alpha_1 \cap \beta_2$, for almost complex structure J_α which has been stretched appropriately along the two dashed curves c and c_α shown in Figure 2.26. Here $\Theta^+_{\beta\beta'} \in \mathbb{T}_\beta \cap \mathbb{T}_{\beta'}$ is the top graded intersection point. In particular no holomorphic triangles are counted which go over w', and the multiplicity of holomorphic triangles over w is the same in the stabilized and unstabilized Heegaard triples.

Proof. Let $\alpha_1 \cap \beta_1 = \{x^-, x^+\}$ and let $\alpha_1 \cap \beta_2 = \{z^-, z^+\}$. Suppose that $\psi = \psi_\Sigma \# \psi_0$ is a homology triangle in $\pi_2(\mathbf{x} \times x, \Theta^+_{\beta\beta'} \times \theta^+, \mathbf{z} \times z)$, where $\Theta^+_{\beta\beta'} \times \theta^+$ is the top degree generator. This connected sum decomposition is the one obtained by cutting the homology class along c. We first note that

$$\mu(\psi_0) = m_1 + m_2 + N_1 + N_2.$$

This can be seen by verifying that the formula holds for one of the Maslov index zero triangles supported in a single region, and that every bigon in the diagram $(S^2, \alpha_1, \beta_1, \beta_2)$ contains exactly one of the regions containing N_1, N_2, m_1 or m_2, so that the formula respects splicing in any of the bigons in the Heegaard triple. Hence we can write

$$\mu(\psi) = \mu(\psi_\Sigma) + \mu(\psi_0) - 2m_1(\psi_0) = \mu(\psi_\Sigma) + (m_2 - m_1 + N_1 + N_2)(\psi_0). \tag{2.9}$$

Let c and c_α be the two circles which we are stretching on, as labeled in Figure 2.26. For a fixed neck length on c, when we stretch along c_α, a sequence of holomorphic curves representing a homology class ψ splits into a collection of curves on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta} \cup \{\beta_1\}, \boldsymbol{\beta}' \cup \{\beta_2\})$ and a collection on $(S^2, \alpha_1, \beta_1, \beta_2)$.

Using the same argument as in [MO10, Proposition 5.2] (cf. also [Zem16, Proposition 5.3]), no curve appearing in the weak limit on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta} \cup \{\beta_1\}, \boldsymbol{\beta}' \cup \{\beta_2\})$ with a boundary

component mapping to an $\boldsymbol{\alpha}$–curve can also have a boundary component which maps to β_1 or β_2.

Hence, a weak limit can be extracted and arranged into the following collections of curves:

1. a broken holomorphic triangle u^0_Σ on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ representing a class $\psi^0_\Sigma \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$;
2. a broken holomorphic disk $u_{\beta\beta'}$ on $(\Sigma, \boldsymbol{\beta} \cup \{\beta_1\}, \boldsymbol{\beta}' \cup \{\beta_2\})$ representing a class $\phi_{\beta\beta'}$ in $\pi_2(\Theta^+_{\beta\beta'} \times \theta^+, \mathbf{y} \times \theta^+)$;
3. a broken holomorphic triangle u_0 on $(S^2, \alpha_1, \beta_1, \beta_2)$ representing $\psi_0 \in \pi_2(x, \theta^+, z)$.

Let $\phi^0_{\beta\beta'}$ denote the homology class in $\pi_2(\Theta^+_{\beta\beta'}, \mathbf{y})$ induced by $\phi_{\beta\beta'}$ on $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}')$ when we remove a ball containing the β_1 and β_2 curves.

We note that $\psi_\Sigma = \psi^0_\Sigma + \phi^0_{\beta\beta'}$, and hence

$$\mu(\psi_\Sigma) = \mu(\psi^0_\Sigma) + \mu(\phi^0_{\beta\beta'}).$$

We note that

$$(m_2 - m_1)(\psi) = (m_2 - m_1)(\psi^0_\Sigma + \phi_{\beta\beta'}) = (m_2 - m_1)(\phi_{\beta\beta'}),$$

as $m_2(\psi^0_\Sigma) = m_1(\psi^0_\Sigma)$. Using Equation (2.9) we thus have

$$\mu(\psi) = \mu(\psi^0_\Sigma) + (\mu(\phi^0_{\beta\beta'}) - m_1(\phi^0_{\beta\beta'})) + m_2(\phi_{\beta\beta'}) + (N_1 + N_2)(\psi_0).$$

By Lemma 2.13.7, we note that $\mu(\phi^0_{\beta\beta'}) - m_1(\phi^0_{\beta\beta'}) \geq 0$. If we count triangles ψ on the stabilized diagram with Maslov index zero, we conclude that for almost complex structures sufficiently stretched on c_α, one has

$$\mu(\psi^0_\Sigma) = \mu(\phi^0_{\beta\beta'}) - m_1(\phi^0_{\beta\beta'}) = m_2(\phi_{\beta\beta'}) = N_1(\psi_0) = N_2(\psi_0) = 0.$$

Another application of Lemma 2.13.7 shows that $\phi^0_{\beta\beta}$ is the constant homology class on $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}')$. As $(n_1 + n_2)(\phi_{\beta\beta'}) = (m_1 + m_2)(\phi_{\beta\beta'}) = 0$ by the local vertex multiplicities at θ^-, we conclude that $\phi_{\beta\beta'}$ is also a constant homology class. Hence the weak limit contains

just a Maslov index zero holomorphic triangle on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$, and a broken holomorphic triangle on $(S^2, \alpha_1, \beta_1, \beta_2)$, representing a homology class ψ_0.

Since we also have $N_1(\psi) = N_2(\psi) = 0$, it is easy to see that $\psi_0 \in \pi_2(x^+, \theta^+, z^+)$ or $\psi_0 \in \pi_2(x^-, \theta^+, z^-)$ (i.e. the off diagonal entries of the matrix in the Proposition statement are zero).

Finally, to count the curves in the homology class $\psi_\Sigma \# \psi_0$, we can fix a neck length on c, and stretch c_α. As the broken curve appearing in the limit on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta} \cup \{\beta_1\}, \boldsymbol{\beta}' \cup \{\beta_2\})$ is a single Maslov index zero holomorphic triangle on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$, by stretching sufficiently along c_α, we can use standard gluing results to reduce the problem of counting the curves in the moduli space $\mathcal{M}(\psi_\Sigma \# \psi_0)$ to the problem of counting the holomorphic triangles in $\{u \in \mathcal{M}(\psi_0) : \rho^{\theta^-}(u) = \mathbf{d}\}$ (a set of curves on $(S^2, \alpha_1, \beta_1, \beta_2)$) for a generic divisor $\mathbf{d} \in \mathrm{Sym}^{m_1}(\Delta)$. In the latter moduli space we are considering holomorphic $u : S \to S^2$ which only map interior points of the source S to $\theta^- \in S^2$. The argument of Lemma 2.5.19 readily adapts to this situation. For a generic path $\mathbf{d}_t$ between two divisors $\mathbf{d}_0$ and $\mathbf{d}_1$, we can form a 1–dimensional moduli space associated to curves matching divisors along $\mathbf{d}_t$, and it is easily observed that there are in fact no ends of the 1–dimensional moduli space except those corresponding to $\mathbf{d}_0$ and $\mathbf{d}_1$. If $\mathbf{d}$ is a generic divisor, we consider the path starting at $\mathbf{d}$ and approaching the $\boldsymbol{\alpha}$–edge of Δ, but spaced farther and farther apart, as in Lemma 2.5.19. As in that lemma, the broken curve appearing in the end consists of $|\mathbf{d}|$ Maslov index 2 $\boldsymbol{\alpha}$–boundary degenerations, each satisfying a matching condition with a single point $p \in (-\infty, 0] \times \mathbb{R}$, as well as a Maslov index zero holomorphic triangle, whose domain is just an empty triangle on $(S^2, \alpha_1, \beta_1, \beta_2)$, and hence admits a unique holomorphic representative. By gluing, we thus obtain that the count for a general $\mathbf{d} \in \mathrm{Sym}^{m_1}(\Delta)$ is $1 \in \mathbb{Z}_2$. This implies that the diagonal entries of the stabilized triangle matrix are identified with the unstabilized triangle counts, completing the proof. □

Proposition 2.13.9. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w'\})$ is a Heegaard triple with $\boldsymbol{\alpha}'$ small isotopies of the $\boldsymbol{\alpha}$ curves. Let $\overline{\mathcal{T}}$ be the Heegaard triple $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_2\}, \boldsymbol{\alpha} \cup \{\alpha_1\}, \boldsymbol{\beta} \cup$*

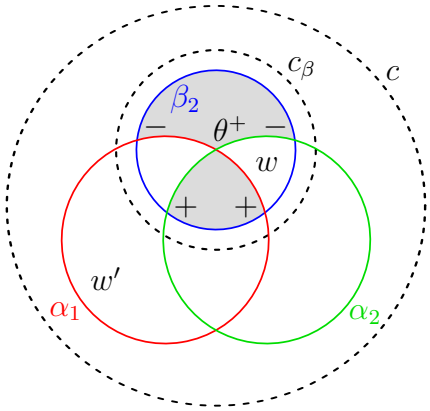

Figure 2.27: The Heegaard triple used to go from $\mathcal{H}_{1.5}$ to $\mathcal{H}_2$ in Proposition 2.13.9. The shaded regions are each examples of small triangles which might be counted.

$\{\beta_2\}, \mathbf{w} \cup \{w, w'\}$) *shown in Figure 2.27. The triangle counts associated to the triple* $\overline{\mathcal{T}}$ *satisfy*

$$F_{\overline{\mathcal{T}}}(\Theta^+_{\alpha'\alpha} \times \theta^+, \mathbf{x} \times \cdot) = \begin{pmatrix} F_{\mathcal{T}}(\Theta^+_{\alpha'\alpha}, \mathbf{x}) & 0 \\ 0 & F_{\mathcal{T}}(\Theta^+_{\alpha'\alpha}, \mathbf{x}) \end{pmatrix}$$

with respect to the Maslov grading on the intersection points of $\alpha_1 \cap \beta_2$ *and* $\alpha_2 \cap \beta_2$*, for almost complex structure* J_β *which has been stretched appropriately along the two dashed curves* c *and* c_β *shown in Figure 2.26. Here* $\Theta^+_{\alpha'\alpha} \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_\alpha$ *is the top graded intersection point. Note in particular that no holomorphic triangles are counted in the above expression which go over* w*, and the multiplicity of these holomorphic triangles over* w' *is the same in the stabilized and unstabilized Heegaard triples.*

Proof. The proof is identical to the proof of Proposition 2.13.8. □

Lemma 2.13.10. *We can pick almost complex structures* J_α *and* J_β *so that Lemma 2.13.3 applies to compute the differential on* $\mathcal{H}_{1.5}$ *and Propositions 2.13.8 and 2.13.9 can be used to compute the triangle maps. Also the almost complex structure* J_α *can be used on* $\mathcal{H}_1$ *to compute* $S^{\pm}_{w'}$ *and the almost complex structure* J_β *can be used to compute* $S^{\pm}_{w}$ *on the diagram* $\mathcal{H}_2$*.*

Proof. This is a modification of Lemma 2.13.5. The issue is that for Lemma 2.13.3 we stretched along c, while for Propositions 2.13.8 and 2.13.9 we stretched along c_α and c_β. The

key point of this argument is that if we stretch J_α enough on both c and c_α on $\mathcal{H}_1$, then the change of almost complex structure maps for stretching further on either one of them eventually only count the constant disks. The argument is easily adapted to to consider the diagrams $\mathcal{H}_{1.5}, \mathcal{H}_2$ and the almost complex structure J_β.

Let m_1, m_2, n_1 and n_2 denote the multiplicities as in Figure 2.25. If $\phi = \phi_\Sigma \# \phi_0 \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y)$, then we have, as before

$$\mu(\phi) = \mu(\phi_\Sigma) + 2m_2(\phi) + \Delta(x, y).$$

Arguing as in Lemma 2.13.5 shows that the only holomorphic disks with $\Delta(x, y) \geq 0$, counted for a sufficiently stretched dynamic almost complex structure, are the constant disks. The disks that are not considered by that argument are those with $\Delta(x, y) = -1$. In this case, we could conceivably count disks with $\mu(\phi_\Sigma) = 1$ and $m_2(\phi_0) = 0$. Note as we stretch along c and c_α, dynamic almost complex structures on $\Sigma \# S^2 \# S^2 \times [0, 1] \times \mathbb{R}$ degenerate to yield a cylindrical almost complex structure on $\Sigma \vee S^2 \vee S^2 \times [0, 1] \times \mathbb{R}$, as in Figure 2.24. Noting that $\mu(\phi_\Sigma) = 1$, we conclude that ϕ_Σ has a representative as a genuine holomorphic disk u_Σ. On the other hand, the middle copy of S^2 has only the β_1 curve on it, and no $\boldsymbol{\alpha}$–curve. We can extract a weak limit to a curve on (S^2, β_1). It must satisfy the a matching condition with u_Σ at the connected sum point, but u_Σ is a genuine disk, so this is only possible if the matching curves on (S^2, β_1) are all closed surfaces mapped into sets of the form $S^2 \times \{pt\}$. As such, they contribute to both $m_1(\phi)$ and $n_2(\phi)$ equally, so we conclude that

$$n_2(\phi) \geq m_1(\phi).$$

On the other hand, since $\Delta(x, y) = -1$, we know that $n_1(\phi) + n_2(\phi) + 1 = m_1(\phi) + m_2(\phi)$, or rather

$$m_2(\phi) = n_1(\phi) + (n_2(\phi) - m_1(\phi)) + 1,$$

from which it follows that $m_2(\phi) \geq 1$, contradicting our earlier computation that $m_2(\phi) = 0$. Hence we conclude that no holomorphic disks are counted with $\Delta(x, y) = -1$.

Finally, to see that J_α and J_β can be used to compute $S^{\pm}_{w'}$ and $S^{\pm}_{w}$, one can modify the above

argument to see that changing J_α and J_β to ones in the definition of the free-stabilization maps preserves the maps $S^{\pm}_{w'}$ and $S^{\pm}_{w}$. See also Lemma 2.5.13. □

We can now compute $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ for appropriate choices of almost complex structures and prove Theorem 2.13.1:

Proof of Theorem 2.13.1. We pick the almost complex structures J_1 and J_2 on $\mathcal{H}_1$ and $\mathcal{H}_2$ to be J_α on $\mathcal{H}_1$ and J_β on $\mathcal{H}_2$. We can decompose $\Phi_{(\mathcal{H}_1, J_\alpha) \to (\mathcal{H}_2, J_\beta)}$ as the composition of three maps:

$$\Phi_{(\mathcal{H}_1, J_\alpha) \to (\mathcal{H}_2, J_\beta)} = \Phi_{(\mathcal{H}_{1.5}, J_\beta) \to (\mathcal{H}_2, J_\beta)} \circ \Phi_{(\mathcal{H}_{1.5}, J_\alpha) \to (\mathcal{H}_{1.5}, J_\beta)} \circ \Phi_{(\mathcal{H}_1, J_\alpha) \to (\mathcal{H}_{1.5}, J_\alpha)}.$$

Using Lemma 2.13.5 and Propositions 2.13.8 and 2.13.9, this becomes simply

$$\begin{pmatrix} (\Phi^{\boldsymbol{\beta'}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha'}})_{w'} & 0 \\ 0 & (\Phi^{\boldsymbol{\beta'}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha'}})_{w'} \end{pmatrix} \begin{pmatrix} \mathrm{id} & 0 \\ * & \mathrm{id} \end{pmatrix} \begin{pmatrix} (\Phi^{\boldsymbol{\beta} \to \boldsymbol{\beta'}}_{\boldsymbol{\alpha}})_{w} & 0 \\ 0 & (\Phi^{\boldsymbol{\beta} \to \boldsymbol{\beta'}}_{\boldsymbol{\alpha}})_{w} \end{pmatrix}$$

$$= \begin{pmatrix} (\Phi^{\boldsymbol{\beta'}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha'}})_{w'} \circ (\Phi^{\boldsymbol{\beta} \to \boldsymbol{\beta'}}_{\boldsymbol{\alpha}})_{w} & 0 \\ * & (\Phi^{\boldsymbol{\beta'}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha'}})_{w'} \circ (\Phi^{\boldsymbol{\beta} \to \boldsymbol{\beta'}}_{\boldsymbol{\alpha}})_{w} \end{pmatrix},$$

as we wanted. Furthermore, by Lemma 2.13.10, the almost complex structure J_α can be used to compute $S^{\pm}_{w'}$ on $\mathcal{H}_1$, and an analogous argument shows that J_β can be chosen so that it can be used to compute $S^{\pm}_{w}$ on $\mathcal{H}_2$. □

2.13.2 Basepoint moving maps and the π_1–action

In this section, we compute the diffeomorphism map induced by moving a basepoint along an path, and use this to compute the π_1–action for moving a basepoint in a closed loop.

We first need a computation of the relative homology map for the diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ considered in the previous section. In the diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$, let λ be the path shown in Figure 2.28.

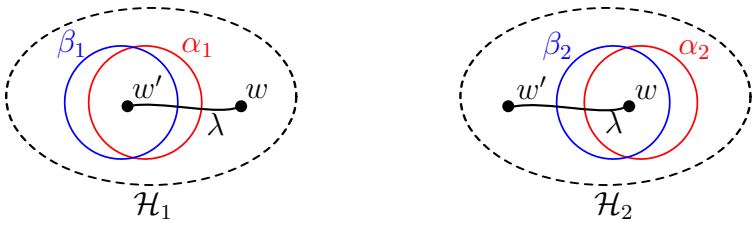

Figure 2.28: The path λ in the diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$.

Lemma 2.13.11. *In the diagram $\mathcal{H}_1$, with almost complex structure J_α, the map A_λ takes the form*

$$(A_\lambda)_{\mathcal{H}_1,J_\alpha} = \begin{pmatrix} U_w\Phi_w & U_w \\ 1 & U_w\Phi_w \end{pmatrix}.$$

On the diagram $\mathcal{H}_2$, with almost complex structure J_β, the map A_λ takes the form

$$(A_\lambda)_{\mathcal{H}_2,J_\beta} = \begin{pmatrix} 0 & U_w \\ 1 & 0 \end{pmatrix}.$$

Here Φ_w is the map on the unstabilized complex $CF^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\})$ defined by the formula

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

Proof. This follows from Lemma 2.13.4, which uses the holomorphic disk counts from Lemma 2.13.3, applied to the diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$. It is straightforward to determine the contribution of disks to the off diagonal entries. For the diagonal entries, the key point is to observe that using the almost complex structure J_α, for the a holomorphic disk counted along the diagonal of the matrix for A_λ, there is no change in multiplicity across the β_2 curve, and a drop of $n_w(\phi)$ across α_1. Hence when we count curves with the coefficient $a(\lambda,\phi) = n_w(\phi)$, we get exactly $U_w\Phi_w$, appearing along the diagonal. Similarly if we use the almost complex structure J_β on $\mathcal{H}_2$, for the disks contributing to the diagonal entries of the matrix, there is no change across the curve α_2, so the diagonal entries of A_λ are zero. □

We now pay an earlier debt, and prove Lemma 2.6.7:

Lemma 2.13.12. *If λ is a path with w and w' as endpoints, then*

$$S_w^- A_\lambda S_w^+ \simeq \mathrm{id}\,.$$

The chain homotopy is $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant for any coloring which gives w and w' the same color. Furthermore, there is a diagram and almost complex structure such that

$$A_\lambda = \begin{pmatrix} 0 & U_w \\ 1 & 0 \end{pmatrix},$$

with respect to our usual matrix notation.

Proof. We use the diagram and almost complex structure $(\mathcal{H}_2, J_\beta)$. The form of the matrix in $(A_\lambda)_{\mathcal{H}_2,J_\beta}$ is computed in Lemma 2.13.11. To prove the relation $S_w^- A_\lambda S_w^+ \simeq \mathrm{id}$, one can just multiply matrices:

$$S_w^- A_\lambda S_w^+ \simeq \begin{pmatrix} 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & U_w \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = 1.$$

We note that for this relation, we only need the form of the bottom left component of A_λ. Adapting the differential computation from Lemma 2.5.5, we see that for any almost complex structure which can be used to compute the maps $S_w^\pm$, the bottom right component is id. □

Theorem 2.13.13. *If λ is a path from w to w', then we have an equivalence of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant chain homotopy types*

$$\lambda_* \simeq S_w^- A_\lambda S_{w'}^+,$$

for any coloring which gives w and w' the same color.

Proof. First pick a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\})$ such that λ is embedded in Σ and doesn't intersect any $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. Let ϕ_t be the isotopy of Σ obtained by moving w to w' along λ, so that $\lambda_* = (\phi_1)_*$. We can assume that ϕ_t is fixed outside a small neighborhood of λ, which doesn't intersect any of the $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. If J_s is an almost complex structure on Σ, then pushing forward under ϕ_1 yields a map

$$\Gamma : CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\}) \to CF^-_{(\phi_1)_* J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w'\}).$$

The map

$$\lambda_* : CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\}) \to CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w'\})$$

is then the composition

$$\lambda_* = \Phi_{(\phi_1)_* J_s \to J_s} \circ \Gamma.$$

We first claim that the map λ_* above is chain homotopic to the tautological map between the two complexes. To see this, note that Γ is the tautological map between the two complexes, though the almost complex structures associated to the two complexes are different. The map $\Phi_{(\phi_1)_* J_s \to J_s}$ is obtained by counting Maslov index zero holomorphic disks for a dynamic almost complex structure which interpolates $(\phi_1)_* J_s$ and J_s. On the other hand, the isotopy ϕ_t gives a diffeomorphism Φ of $\Sigma \times [0,1] \times \mathbb{R}$ with itself, given by

$$\Phi(x, s, t) = (\phi_t(x), s, t).$$

Given J_s, we can simply push J_s forward along Φ to get a dynamic almost complex structure, interpolating J_s and $(\phi_1)_* J_s$. However $\Phi_*(J_s)$–holomorphic disks on $\Sigma \times [0,1] \times \mathbb{R}$ can be pulled back under Φ to get J_s–holomorphic disks on $\Sigma \times [0,1] \times \mathbb{R}$. Note that the assumption that ϕ_t is fixed near all of the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves implies that the pushforward of a J_s–holomorphic disk will have acceptable boundary values. By transversality of J_s, a Maslov index zero $\Phi_*(J_s)$–holomorphic disk must thus be a constant disk. Thus the change of almost complex structures map $\Phi_{(\phi_1)_* J_s \to J_s}$ is just the tautological map.

Hence, as a map from $CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w} \cup \{w\})$ to $CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w} \cup \{w'\})$, the map $\lambda_* = \Phi_{(\phi_1)_* J_s \to J_s} \circ \Gamma$ is the tautological map.

Using this, and our computations of the various transition maps and relative homology maps, the result is now a computation. The map $S^-_w A_\lambda S^+_{w'}$ viewed as a map from $CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w\})$ to $CF^-_{J_s}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0 \cup \{w'\})$, can be computed as the composition

$$S^-_w A_\lambda S^+_{w'} = \Phi^{\boldsymbol{\beta}' \to \boldsymbol{\beta}}_{\boldsymbol{\alpha}' \to \boldsymbol{\alpha}} \circ S^-_w \circ (A_\lambda)_{\mathcal{H}_2, J_\beta} \circ \Phi_{(\mathcal{H}_1, J_\alpha) \to (\mathcal{H}_2 \to J_\beta)} \circ S^+_{w'}.$$

Using the matrix forms from Theorem 2.13.1, Lemma 2.13.5, and Lemma 2.13.11 this is

equivalent to

$$\Phi^{\boldsymbol{\beta}'\to\boldsymbol{\beta}}_{\boldsymbol{\alpha}'\to\boldsymbol{\alpha}}\begin{pmatrix}0 & 1\end{pmatrix}\begin{pmatrix}0 & U_w\\ 1 & 0\end{pmatrix}\begin{pmatrix}\Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}} & 0\\ * & \Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}}\end{pmatrix}\begin{pmatrix}1\\ 0\end{pmatrix},$$

which multiplies out to $\Phi^{\boldsymbol{\beta}'\to\boldsymbol{\beta}}_{\boldsymbol{\alpha}'\to\boldsymbol{\alpha}}\Phi^{\boldsymbol{\beta}'}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}}$, which is chain homotopic to the tautological map by naturality. We already showed that λ_* was the tautological map, so we get the equivalence stated in the theorem. □

Remark 2.13.14. Note that in the previous lemma we used the expression for $(A_\lambda)_{\mathcal{H}_2,J_\beta}$ from Lemma 2.13.11, but on examination of Lemma 2.13.11, the map $(A_\lambda)_{\mathcal{H}_1,J_\alpha}$ looks more complicated. Nevertheless, for the theory to make sense, we must get the same answer if we use $(A_\lambda)_{\mathcal{H}_1,J_\alpha}$ instead. If we were to do the above computation instead using $(A_\lambda)_{\mathcal{H}_1,J_\alpha}$ at the appropriate spot, upon performing the matrix multiplication we get

$$\Phi^{\boldsymbol{\beta}'\to\boldsymbol{\beta}}_{\boldsymbol{\alpha}'\to\boldsymbol{\alpha}}(1+U_w\cdot C\Phi_w)\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta}'}_{\boldsymbol{\alpha}},$$

for the map $C=*$ appearing in the lower left of the change of almost complex structure map $\Phi_{J_\alpha\to J_\beta}$ on the diagram $\mathcal{H}_{1.5}$. As described in Remarks 2.13.2 and 2.13.6, the term appearing the $*$ component of the matrix is Φ_w (once we identify U_w and $U_{w'}$), so doing the computation this way yields an extra summand of $U_w\cdot\Phi_w^2$. Fortunately, we'll see in Lemma 2.13.19, that $\Phi_w^2\simeq 0$, equivariantly, so we get the same answer in the end.

As a corollary, we see that the graph cobordism maps and the path cobordism maps agree if the graph is just a collection of paths:

Corollary 2.13.15. *If $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ is a graph cobordism with Γ a collection of paths from the incoming boundary to the outgoing boundary, then the path cobordism maps agree with the graph cobordism maps.*

Proof. We first pick a collar neighborhood N of Y_1 and an identification of N as $Y_1\times I$. By choosing N appropriately we can assume that $N\cap\Gamma$ is a collection of paths from points in $Y_1\times\{0\}$ to points in $Y_1\times\{1\}$. Using Lemma 2.11.4 and the composition law for path

cobordisms (Lemma 2.9.3), it is sufficient to show that the path and graph cobordisms agree for the cobordism $(N, \Gamma \cap N)$. Letting Γ_{Y_1} denote the projection of $\Gamma \cap N$ into Y_1, we note that the path cobordism map of $(N, \Gamma \cap N)$ is equal to the basepoint moving map associated to the paths in Γ_{Y_1}. Using Lemma 2.11.3, we know the graph cobordism map for $(N, \Gamma \cap N)$ is equal to the graph action map for Γ_{Y_1}. By Theorem 2.13.13, the graph action map in this case is equal to the basepoint moving map, and hence equal to the path cobordism map, completing the proof. □

As an additional consequence, we can now compute the $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariant chain homotopy type of the π_1–action on $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$. We first introduce some notation. Suppose that $w \in \mathbf{w}$ and $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for $(Y, \mathbf{w})$. We define an endomorphism Φ_w by the formula

$$\Phi_w(\mathbf{x}) = U_w^{-1} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_w(\phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}.$$

We can interpret Φ_w as the formal derivative of the differential, with respect to U_w. An analog of the map Φ_w was discovered by Sarkar on link Floer homology in [Sar15] when studying certain automorphisms of link Floer homology.

Lemma 2.13.16. *The map Φ_w commutes with change of diagrams maps up to $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariant chain homotopy. If $w \in \mathbf{w}$ and there are enough basepoints for $S_w^+ S_w^-$ to be defined, then*

$$\Phi_w \simeq S_w^+ S_w^-.$$

Proof. We first show that Φ_w commutes with change of diagrams maps up to equivariant chain homotopies. To see this, note that the transition maps commute with ∂, so upon applying d/dU_w to

$$\partial \Phi_{\mathcal{H}_1 \to \mathcal{H}_2} + \Phi_{\mathcal{H}_1 \to \mathcal{H}_2} \partial = 0$$

and applying the Leibniz rule, we see that Φ_w commutes with the transition maps up to chain homotopy.

To see that $S^+_w S^+_w \simeq \Phi_w$, one picks a diagram for $(Y, \mathbf{w})$ which is free-stabilized at w, near another basepoint w', and picks an almost complex structure so that

$$\partial = \begin{pmatrix} \partial_{w'} & U_w + U_{w'} \\ 0 & \partial_{w'} \end{pmatrix}$$

which can also be used to compute $S^{\pm}_w$ (for example the pair $(\mathcal{H}_2, J_\alpha)$ from Lemma 2.13.4). Differentiating with respect to U_w yields

$$\Phi_w = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$

which is $S^+_w S^-_w$. □

We proceed with a lemma.

Lemma 2.13.17. *If λ is a path from w to w', we have that*

$$A_\lambda \Phi_w + \Phi_w A_\lambda \simeq 1.$$

If w and w' are distinct basepoints of Y, we have

$$\Phi_w S^{\pm}_{w'} + S^{\pm}_{w'} \Phi_w \simeq 0.$$

Proof. For the first equivalence, take the expression

$$A_\lambda \partial + \partial A_\lambda = U_w + U_{w'},$$

and differentiate with respect to U_w. We view the latter expression as an expression of matrices, over the points of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Using the Leibniz rule, we get that

$$A_\lambda \partial' + \partial' A_\lambda + \partial A'_\lambda + A'_\lambda \partial = 1.$$

Using the interpretation of Φ_w as ∂', we get exactly that

$$A_\lambda \Phi_w + \Phi_w A_\lambda \simeq 1.$$

For the second equivalence, we note that

$$\partial S^{\pm}_{w'} + S^{\pm}_{w'}\partial = 0,$$

and upon differentiating with respect to U_w, we get the desired claim. □

We can now prove Theorem F:

Theorem F. *If γ is a closed loop in Y starting and ending at w, the action of $\gamma \in \pi_1(Y, w)$ has the $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariant chain homotopy type*

$$\gamma_* \simeq 1 + \Phi_w \circ A_\gamma.$$

Proof. Break γ into the concatenation of two paths, and write $\gamma = \lambda_2 * \lambda_1$, where λ_1 is a path from w to w' and λ_2 is a path from w' to w. Using Theorem 2.13.13, we then have that

$$(\gamma)_* \simeq (\lambda_2)_* \circ (\lambda_1)_* \simeq (S^-_{w'} A_{\lambda_2} S^+_w)(S^-_w A_{\lambda_1} S^+_{w'}).$$

We can thus compute

$$\begin{aligned}
(\gamma)_* &\simeq S^-_{w'} A_{\lambda_2} \Phi_w A_{\lambda_1} S^+_{w'} && \text{(Lemma 2.13.16)} \\
&\simeq S^-_{w'} A_{\lambda_2} (A_{\lambda_1} \Phi_w + 1) S^+_{w'} && \text{(Lemma 2.13.17)} \\
&\simeq S^-_{w'} A_{\lambda_2} A_{\lambda_1} \Phi_w S^+_{w'} + S^-_{w'} A_{\lambda_2} S^+_{w'} && \\
&\simeq S^-_{w'} A_{\lambda_2} A_{\lambda_1} S^+_{w'} \Phi_w + S^-_{w'} A_{\lambda_2} S^+_{w'} && \text{(Lemma 2.13.16)} \\
&\simeq A_\gamma \Phi_w + S^-_{w'} A_{\lambda_2} S^+_{w'} && \text{(Lemma 2.6.8)} \\
&\simeq A_\gamma \Phi_w + 1 && \text{(Lemma 2.6.7).}
\end{aligned}$$

We finally note that $A_\gamma \Phi_w + A_\gamma \Phi_w \simeq 0$, which can be seen by differentiating $A_\gamma \partial + \partial A_\gamma = 0$. □

We now remark that if we apply $\frac{d}{dU_w}$ to the expression $\partial(x)$, we must use the Leibniz rule, which yields

$$\frac{d}{dU_w}(\partial(x)) = \left(\frac{d}{dU_w}\partial\right)(x) + \partial\left(\frac{d}{dU_w}x\right)$$

or rather

$$\Phi_w = \frac{d}{dU_w} \circ \partial + \partial \circ \frac{d}{dU_w}. \tag{2.10}$$

The map Φ_w is a $\mathbb{Z}_2[U_{\mathbf{w}}]$–module homomorphism and hence induces a $\mathbb{Z}_2[U_{\mathfrak{P}}]$–module homomorphism on the colored complexes. Note however that the formal derivative $\frac{d}{dU_w}$ is not $\mathbb{Z}_2[U_{\mathbf{w}}]$–module homomorphism and hence doesn't necessarily induce a chain homotopy on the tensor product of the uncolored CF° with the coloring module.

Nonetheless, the chain homotopy persists on any coloring where w is not given the same coloring as any other basepoint. We summarize with the following corollary:

Corollary 2.13.18. *The $\pi_1(Y,w)$–action vanishes on $HF^\circ(Y,\mathbf{w},\mathfrak{s},\sigma,\mathfrak{P})$ for $\circ \in \{+,-,\infty\}$ as long as w has a coloring distinct from all the other basepoints. In particular, it vanishes on the uncolored modules $HF^\circ(Y,\mathbf{w},\mathfrak{s})$ for $\circ \in \{+,-,\infty\}$. The chain homotopy is not $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant, and hence the $\pi_1(Y,w)$–action does not necessarily vanish on $CF^\circ(Y,\mathbf{w},\mathfrak{s})$, the $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariant chain homotopy type invariant.*

Note that the map $\frac{d}{dU_w}$ is not well defined on the hat versions of the complexes since they are obtained from CF^- by tensoring over $\mathbb{Z}_2[U_{\mathbf{w}}]$ with the module $\mathbb{Z}_2[U_{\mathbf{w}}]/(U_{w_1} = 0,\dots,U_{w_n} = 0)$. One can find examples where the π_1–action does not vanish on $\widehat{HF}(Y,w)$ for singly pointed 3–manifolds. For multi-pointed 3–manifolds, the action will be nontrivial whenever the homology action $[\gamma]$ is nontrivial.

We can finally prove Theorem G:

Theorem G. *If $(W,\boldsymbol{\gamma}) : (Y_1,\mathbf{w}_1) \to (Y_2,\mathbf{w}_2)$ is a path cobordism and each path is given a distinct coloring (for example if there is only one path) then the map $F^\circ_{W,\boldsymbol{\gamma},\mathfrak{s}}$ is independent of the curves $\boldsymbol{\gamma}$ on homology for $\circ = +,-,\infty$.*

Proof. This follows from the composition law. Suppose $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ are two collections of paths. Decompose W into $W_2 \circ W_1$, where W_1 is obtained by adding 1-handles and W_2 is obtained by adding 2- and 3-handles. Write $\partial W_1 = -Y_1 \sqcup Y'$. Observe that $\pi_1(Y') \to \pi_1(W)$ is surjective, so we can homotope $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma'}$ so that they only differ in a collar neighborhood of

Y'. Since W is a 4–manifold, such a homotopy can be taken to be an isotopy. Furthermore we can also assume that $\boldsymbol{\gamma} \cap W_i$ and $\boldsymbol{\gamma}' \cap W_i$ are paths for $i = 1, 2$. Using the composition law and the fact the π_1–action is trivial on homology for manifolds where all basepoints have distinct colors, the theorem follows. □

As we've remarked, the map Φ_w is not necessarily $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariantly chain homotopic to zero, though we do have the following (compare [Sar15, Lemma 4.6], and also Lemma 2.4.12):

Lemma 2.13.19. *The map $(\Phi_w)^2$ is $\mathbb{Z}_2[U_{\mathbf{w}}]$–equivariantly chain homotopic to zero.*

Proof. Write $\partial = \sum_{n=0}^{\infty} \partial_n U_w^n$ (on the uncolored complexes), where ∂_n doesn't involve any powers of U_w. We take

$$H = \sum_{n=2}^{\infty} \frac{n(n-1)}{2} \partial_n U_w^{n-2},$$

(the quantity $n(n-1)/2$ is computed over $\mathbb{Z}$, then projected to $\mathbb{Z}_2$). We note that the fact that $\partial^2 = 0$ implies that

$$\sum_{i+j=k} \partial_i \partial_j = 0,$$

for each $k \geq 0$. We now simply compute that

$$\begin{aligned}
\partial H + H\partial &= \sum_{i,j \geq 0} \left(\frac{i(i-1)}{2} + \frac{j(j-1)}{2} \right) \partial_i \partial_j U_w^{i+j-2} \\
&= \sum_{k=2}^{\infty} \left(\frac{k(k-1)}{2} \sum_{i+j=k} \partial_i \partial_j \right) U_w^{k-2} + \sum_{i,j \geq 0} ij \partial_i \partial_j U_w^{i+j-2} \\
&= 0 + \Phi_w^2,
\end{aligned}$$

completing the proof. □

2.13.3 An alternate, explicit proof of the π_1–action

In this section we give an alternate proof of the formula for the π_1–action. The key point is that on a multi–pointed diagram, one can explicitly perform the computation using counts of

holomorphic triangles and holomorphic disks. This section is unnecessary from the perspective of just proving the formula for the π_1–action, but the explicit verification offers some insights into the type of algebra that these maps induce. Our strategy will be to break a closed loop γ, starting and ending at $w \in \mathbf{w}$ into a sequence of arcs, each of which crosses a single $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curve. We explicitly compute the map associated to moving a free-stabilization at the incoming end of the arc to the outgoing end of the arc, but keeping the rest of the diagram fixed. This is illustrated in Figure 2.29.

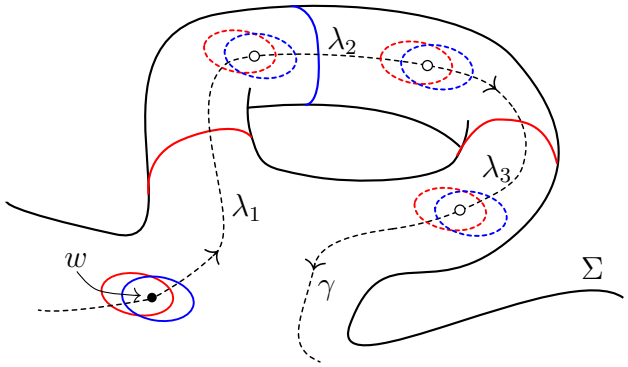

Figure 2.29: Breaking a path γ into a sequence of arcs $\lambda_n * \cdots * \lambda_1$, each of which crosses at most one $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curve. We will compute each move individually, by degenerating the almost complex structure carefully.

We first prove that knowing the formula on multi-pointed diagrams is sufficient for knowing it on singly pointed diagrams:

Lemma 2.13.20. *If the equivalence $\gamma_* \simeq \mathrm{id} + \Phi_w A_\gamma$ holds for diagrams with at least two basepoints, then the formula holds for singly pointed diagrams as well.*

Proof. The idea is that we can add a new basepoint w' using Lemma 2.13.12. Let w' be a

new basepoint, colored the same as w, and let λ be a path between w and w'. We have

$$\begin{aligned}
\gamma_* &\simeq (S^-_{w'} A_\lambda S^+_{w'})(\gamma_*) && \text{(Lemma 2.13.12)} \\
&\simeq S^-_{w'} A_\lambda(\gamma_*) S^+_{w'} && \text{(Theorem 2.5.22)} \\
&\simeq S^-_{w'} A_\lambda(\mathrm{id} + \Phi_w A_\gamma) S^+_{w'} && \text{(Hypothesis)} \\
&\simeq S^-_{w'} A_\lambda S^+_{w'}(\mathrm{id} + \Phi_w A_\gamma) && \text{(Lemmas 2.5.6, 2.13.17)} \\
&\simeq \mathrm{id} + \Phi_w A_\gamma && \text{(Lemma 2.13.12),}
\end{aligned}$$

completing the proof. □

Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0)$ be a Heegaard diagram, and let $\overline{\mathcal{H}} = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w}_0 \cup \{w\})$ be its free-stabilization at a point $w \in \Sigma$. Suppose that γ is a closed loop in Σ, starting and ending at w. We now wish to provide an alternate proof of Theorem F, by showing explicitly that the diffeomorphism map γ_* has the chain homotopy type

$$\gamma_* \simeq 1 + A_\gamma \Phi_w.$$

Let us consider the almost complex structure J_β, stretched along the circle c_β parallel to curve β_0. Analogously, we will write J_α for an almost complex structure which is stretched around on c_α. These are shown in Figure 2.30. Using the differential computation of Lemma 2.13.3, we see that, for either J_α or J_β,

$$A_\gamma = \begin{pmatrix} A_\gamma & U_w + U_{w'} \\ 0 & A_\gamma \end{pmatrix} \qquad \text{and} \qquad \Phi_w = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

Multiplying this out, we get that in this matrix notation, we have

$$1 + A_\gamma \Phi_w = \begin{pmatrix} 1 & A_\gamma \\ 0 & 1 \end{pmatrix}. \tag{2.11}$$

For convenience, we state a differential computation for a free-stabilization which does not occur immediately next to another basepoint of the diagram:

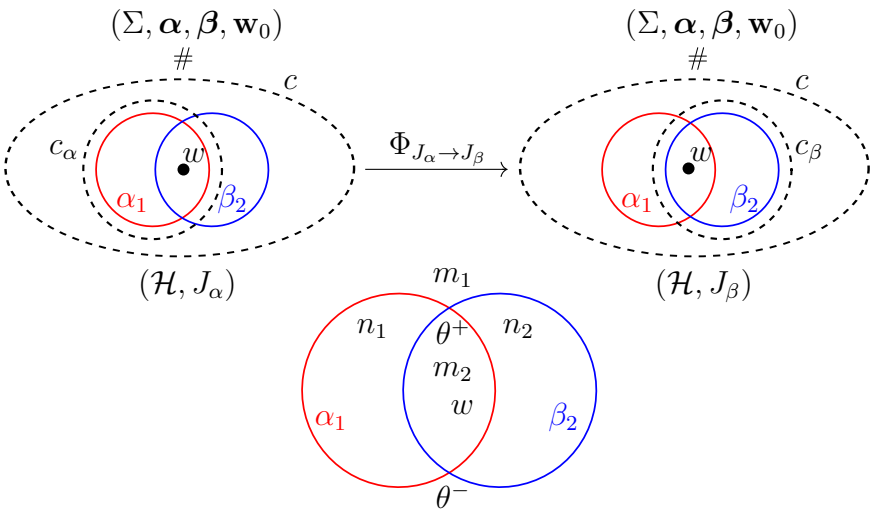

Figure 2.30: The almost complex structures J_α and J_β on a diagram freely-stabilized at w. On the bottom we've written in multiplicities of a homology disk.

Lemma 2.13.21. *If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}_0)$ is a diagram and $\overline{\mathcal{H}}$ is its free stabilization at a point $w \notin \mathbf{w}_0 \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta}$, then for the almost complex structures J_α and J_β discussed before, we have*

$$\partial_{\overline{\mathcal{H}}, J_\alpha} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w_\beta} \\ 0 & \partial_{\mathcal{H}} \end{pmatrix}, \qquad \text{and} \qquad \partial_{\overline{\mathcal{H}}, J_\beta} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w_\alpha} \\ 0 & \partial_{\mathcal{H}} \end{pmatrix},$$

where w_β denotes the basepoint in the component of $\Sigma \setminus \boldsymbol{\beta}$ containing w, and w_α denotes the basepoint in the component of $\Sigma \setminus \boldsymbol{\alpha}$ which contains w.

Proof. The proof is identical to the proofs of Lemmas 2.13.3 and 2.13.4. □

We need to compute the change of almost complex structure map

$$\Phi_{J_\beta \to J_\alpha}.$$

In Lemma 2.13.5, we computed several components of a similar transition map. Using our current basepoint configuration, the analogous computation would yield, for sufficiently stretched almost complex structure (on the outer circle c), that

$$\Phi_{J_\beta \to J_\alpha} = \begin{pmatrix} \mathrm{id} & * \\ 0 & \mathrm{id} \end{pmatrix}.$$

Since the entry of the map $\Phi_{J_\beta \to J_\alpha}$ marked a $*$ has moved from the lower left to the upper right since we now have a different basepoint configuration, we cannot just use the trick of knowing that $\Phi_{J_\beta \to J_\alpha}$ is a chain map to compute the $*$ component. We can however, perform this computation on the level of individual homology classes of disks:

Proposition 2.13.22. *For sufficiently stretched almost complex structures, the change of almost complex structures map $\Phi_{J_\beta \to J_\alpha}$ takes the form*

$$\Phi_{J_\beta \to J_\alpha} = \begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_p \\ 0 & \mathrm{id} \end{pmatrix},$$

where $\widetilde{\Psi}_p$ is the map on the unstabilized complex defined by

$$\widetilde{\Psi}_p(\mathbf{x}) = \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_p(\phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}.$$

Remark 2.13.23. We note that the map $\widetilde{\Psi}_p$ appearing in the previous proposition is not necessarily a chain map. Instead it is a chain homotopy between U_{w_α} and U_{w_β}, where w_α is the basepoint in the same component of $\Sigma \setminus \boldsymbol{\alpha}$ containing p, and w_β is the basepoint in the same component of $\Sigma \setminus \boldsymbol{\beta}$ as p.

Proof of Proposition 2.13.22. We do the same degeneration of the almost complex structure as in Lemma 2.13.5, stretching a dynamic almost complex structure $\widetilde{J}$, which interpolates almost complex structures J_α and J_β, along an outer circle c.

Adapting the argument of Lemma 2.13.5, the transition matrix takes the form

$$\Phi_{J_\beta \to J_\alpha} = \begin{pmatrix} \mathrm{id} & * \\ 0 & \mathrm{id} \end{pmatrix}.$$

Note that the $*$ has moved from the lower left, as it was in Lemma 2.13.5 to the upper right since the basepoint configuration is now different. We now need to compute the disks appearing in the $*$ component of the transition map, i.e. the number of representatives of Maslov index zero disks of the form $\phi_\Sigma \# \phi_0 \in \pi_2(\mathbf{x} \times \theta^-, \mathbf{y} \times \theta^+)$. Note that for such disks

the Maslov index now reads

$$\mu(\phi) = \mu(\phi_\Sigma) + 2m_2 + \Delta(\theta^-, \theta^+) = \mu(\phi_\Sigma) + 2m_2 - 1.$$

For such a disk $\phi = \phi_\Sigma \# \phi_0$ to have a representative for sufficiently stretched almost complex structures, we conclude that $\mu(\phi_\Sigma) = 1$ and $2m_2 = 0$. If $\phi_\Sigma \in \pi_2(\mathbf{x}, \mathbf{y})$ has $m_1(\phi_\Sigma) = k$, we note that there are k nonnegative homology disks $\phi_0 \in \pi_2(\theta^-, \theta^+)$ with $m_2(\phi_0) = 0$ and $m_1(\phi_0) = k$. Similarly, there are $k+1$ nonnegative homology disks $\phi_0 \in \pi_2(\theta^+, \theta^+)$ with $m_2(\phi_0) = 0$ and $m_1(\phi_0) = k$. Among these classes, if $k_1 + k_2 = k - 1$. we will write ϕ_{k_1,k_2} for the homology disk in $\pi_2(\theta^-, \theta^+)$ on (S^2, α_0, β_0) with

$$m_2(\phi_{k_1,k_2}) = 0, \qquad m_1(\phi_{k_1,k_2}) = k, \qquad n_1(\phi_{k_1,k_2}) = k_1, \qquad \text{and} \qquad n_2(\phi_{k_1,k_2}) = k_2.$$

If $K_1 + K_2 = k$, we will write $\phi^+_{K_1,K_2}$ for the homology disk in $\pi_2(\theta^+, \theta^+)$ defined similarly. To compute the unknown component of the almost complex structure map, it is thus sufficient to count the moduli spaces

$$\#\widehat{\mathcal{M}}_{\tilde{J}}(\phi_\Sigma \# \phi_{k_1,k_2})$$

ranging over Maslov index 1 homology classes $\phi_\Sigma \in \pi_2(\mathbf{x}, \mathbf{y})$ for nonnegative k_1 and k_2 with $k_1 + k_2 = m_1(\phi_\Sigma) - 1 = k - 1$. To this end, we will show that

$$\#\mathcal{M}_{\tilde{J}}(\phi_\Sigma \# \phi_{k_1,k_2}) \equiv \#\widehat{\mathcal{M}}(\phi_\Sigma) \pmod 2$$

for each of the k pairs of nonnegative (k_1, k_2) with $k_1 + k_2 = k - 1$. From this count, the proposition follows immediately, since $n_p(\phi_\Sigma) = m_1(\phi_\Sigma)$. Here $\widehat{\mathcal{M}}(\phi_\Sigma)$ is the moduli space of holomorphic representatives for the cylindrical almost complex structure on the unstabilized diagram.

To establish these counts, we use the philosophy that the change of almost complex structures are chain maps. More precisely, we will count the moduli spaces of Maslov index 0 $\widetilde{J}$–holomorphic disks by observing that the appear in the ends of Maslov index 1 $\widetilde{J}$–holomorphic disks, but also observing that our degeneration of the almost complex structure constrains which other curves can appear in the ends.

We consider the ends of $\mathcal{M}(\phi_\Sigma \# \phi^+_{K_1,K_2})$ for $K_1 + K_2 = k$. The ends (with no assumptions on the almost complex structure) correspond to strip breaking, i.e. the ends correspond to a sequence of Maslov index 1 $\widetilde{J}$–holomorphic disks breaking into a single Maslov index 0 $\widetilde{J}$–holomorphic disk, and a Maslov index 1 J_α– or J_β–holomorphic disk. We first note that due to our degeneration, the only Maslov index 0 $\widetilde{J}$–holomorphic disks which can appear are either the constant disks, or ones of the form $\phi_\Sigma \# \phi_{k_1,k_2}$ for a disk ϕ_Σ of Maslov index 1 with $m_1(\phi_\Sigma) > 0$, and (k_1, k_2) a pair of nonnegative integers with $k_1 + k_2 = k - 1$.

If $\widetilde{J}$–holomorphic representatives of $\phi_\Sigma \# \phi^+_{K_1,K_2}$ split into two holomorphic curves representing classes $\phi^1_\Sigma \# \phi^1_0$ and $\phi^2_\Sigma \# \phi^2_0$, we note that $\phi_\Sigma = \phi^2_\Sigma * \phi^1_\Sigma$, but by our neck stretching argument, both ϕ^i_Σ have holomorphic representatives for a cylindrical almost complex structure. As

$$1 = \mu(\phi_\Sigma) = \mu(\phi^1_\Sigma) + \mu(\phi^2_\Sigma),$$

we conclude that if a pair of homology classes $\phi^1_\Sigma \# \phi^1_0$ and $\phi^2_\Sigma \# \phi^2_0$ are represented as a broken holomorphic disk in the ends of $\mathcal{M}_{\tilde{J}}(\phi_\Sigma \# \phi^+_{K_1,K_2})$, then one of ϕ^1_Σ and ϕ^2_Σ, has Maslov index zero, and is hence a constant disk.

We note that there are clearly no nonconstant, nonnegative, Maslov index zero homology classes $\phi_\Sigma \# \phi_0$ with ϕ_Σ a constant disk and $m_2 = 0$. On the other hand, there are exactly three Maslov index 1 homology classes of the form $\phi_\Sigma \# \phi_0$ with ϕ_Σ a constant disk: these are the three bigons which are supported on (S^2, α_0, β_0). Let $\phi^0_{1,0}$ denote the homology class of the bigon with $n_1 = 1$, and let $\phi^0_{0,1}$ denote the homology class of the bigon with $n_2 = 1$. The third bigon will not appear in our argument, as $m_2 \neq 0$.

If $\mathbf{z}$ is an intersection point on the stabilized diagram, let us write $e_{\mathbf{z}}$ for the homology class of a constant disk in $\pi_2(\mathbf{z}, \mathbf{z})$. By our previous reduction of possible strip breaking of holomorphic disks in $\mathcal{M}_{\tilde{J}}(\phi_\Sigma \# \phi^+_{K_1,K_2})$ for a disk $\phi_\Sigma \in \pi_2(\mathbf{x}, \mathbf{y})$ and an arbitrary pair of nonnegative K_1 and K_2 with $K_1 + K_2 = k$, we have that

$$\partial\mathcal{M}_{\tilde{J}}(\phi_\Sigma\#\phi^+_{K_1,K_2}) = \widehat{\mathcal{M}}_{J_\alpha}(\phi_\Sigma\#\phi^+_{K_1,K_2})\times\mathcal{M}_{\tilde{J}}(e_{\mathbf{x}\times\theta^+}) \quad \sqcup \quad \mathcal{M}_{\tilde{J}}(e_{\mathbf{y}\times\theta^+})\times\widehat{\mathcal{M}}_{J_\beta}(\phi_\Sigma\#\phi^+_{K_1,K_2})$$
$$\sqcup \quad \mathcal{M}_{\tilde{J}}(\phi_\Sigma\#\phi_{K_1-1,K_2})\times\widehat{\mathcal{M}}_{J_\beta}(\phi^0_{1,0}) \quad \sqcup \quad \widehat{\mathcal{M}}_{J_\alpha}(\phi^0_{0,1})\times\mathcal{M}_{\tilde{J}}(\phi_\Sigma\#\phi_{K_1,K_2-1}).$$

Noting that the bigons and constant disks have unique holomorphic representative, we get the relation

$$\#\widehat{\mathcal{M}}_{J_\alpha}(\phi_\Sigma\#\phi^+_{K_1,K_2})+\#\widehat{\mathcal{M}}_{J_\beta}(\phi_\Sigma\#\phi^+_{K_1,K_2})+\#\mathcal{M}_{\tilde{J}}(\phi_\Sigma\#\phi_{K_1-1,K_2})+\#\mathcal{M}_{\tilde{J}}(\phi_\Sigma\#\phi_{K_1,K_2-1})=0.$$

Finally we note that by the differential computation from Lemma 2.13.3, we know that

$$\#\widehat{\mathcal{M}}_{J_\beta}(\phi_\Sigma\#\phi^+_{K_1,K_2}) = \begin{cases} \#\widehat{\mathcal{M}}(\phi_\Sigma) & \text{if } (K_1,K_2)=(k,0) \\ 0 & \text{otherwise,} \end{cases}$$

while

$$\#\widehat{\mathcal{M}}_{J_\alpha}(\phi_\Sigma\#\phi^+_{K_1,K_2}) = \begin{cases} \#\widehat{\mathcal{M}}(\phi_\Sigma) & \text{if } (K_1,K_2)=(0,k) \\ 0 & \text{otherwise.} \end{cases}$$

By starting at $(K_1,K_2)=(k,0)$ and using the above recursion formula, to go through all pairs (K_1,K_2) with $K_1+K_2=k$ we see immediately that this implies that

$$\#\widehat{\mathcal{M}}_{\tilde{J}}(\phi_\Sigma\#\phi_{k_1,k_2}) \equiv \#\widehat{\mathcal{M}}(\phi_\Sigma) \pmod 2$$

for any nonnegative pair (k_1,k_2) with $k_1+k_2=k-1$. □

We remark that the map $\widetilde{\Psi}_p$ appearing in the previous proposition is not a natural map, in that it doesn't commute with arbitrary change of diagrams maps. However if p is in the same region as a basepoint w on the diagram, then we have $\widetilde{\Psi}_w = U_w\Phi_w$. Nonetheless, it is natural for appropriately small moves of the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves, in the following sense:

Lemma 2.13.24. *Suppose that $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w})$ is a Heegaard triple such that $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha})$ is a diagram for $(S^1\times S^2)^{\#k}$ and $p\in\Sigma\setminus(\boldsymbol{\alpha}'\cup\boldsymbol{\alpha}\cup\boldsymbol{\beta})$ is a marked point. If p is in a component of $\Sigma\setminus(\boldsymbol{\alpha}'\cup\boldsymbol{\alpha})$ which contains a basepoint of $\mathbf{w}$, then*

$$F\circ\widetilde{\Psi}_p \simeq \widetilde{\Psi}_p\circ F,$$

where $F : CF^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}) \to CF^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}')$ *is the map*

$$F(\mathbf{y}) = F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^+_{\alpha'\alpha} \otimes \mathbf{y}).$$

Proof. Compare Lemma 2.4.5. Using our usual strategy, we count the ends of moduli spaces of Maslov index 1 holomorphic triangles. For a general $\mathbf{x} \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_\alpha$ and $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we get

$$F_{\alpha'\alpha\beta}(\widetilde{\Psi}_p(\mathbf{x})\otimes\mathbf{y})+F_{\alpha'\alpha\beta}(\mathbf{x}\otimes\widetilde{\Psi}_p(\mathbf{y}))+\widetilde{\Psi}_p(F_{\alpha'\alpha\beta}(\mathbf{x}\otimes\mathbf{y})) = H_{\alpha'\alpha\beta,p}(\partial_{\alpha'\alpha}(\mathbf{x})\otimes\mathbf{y}+\mathbf{x}\otimes\partial_{\alpha\beta}(\mathbf{y}))+\partial_{\alpha'\beta}H_{\alpha'\alpha\beta,p}(\mathbf{x}\otimes\mathbf{y}).$$

We insert $\Theta^+_{\alpha'\alpha}$ in for $\mathbf{x}$, and we see that

$$(F \circ \widetilde{\Psi}_p)(\mathbf{y}) + F_{\alpha'\alpha\beta}(\widetilde{\Psi}_p(\Theta^+_{\alpha'\alpha}) \otimes \mathbf{y}) + (\widetilde{\Psi}_p \circ F)(\mathbf{y}) = (H \circ \partial_{\alpha\beta})(\mathbf{y}) + (\partial_{\alpha'\beta} \circ H)(\mathbf{y}).$$

Hence we will be done if we can show that $\widetilde{\Psi}_p(\Theta^+_{\alpha'\alpha}) = \partial\eta$ for some $\eta \in CF^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$, as inserting such an element into the triangle map $F_{\alpha'\alpha\beta}$ yields a map which is chain homotopic to zero in the other component. For that, we note that since p is in the same domain of $\Sigma \setminus (\boldsymbol{\alpha}' \cup \boldsymbol{\alpha})$ as a basepoint w, the map $\widetilde{\Psi}_p$ on $CF^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ is in fact just $U_w\Phi_w$. Now Φ_w is $+1$ graded, and a chain map, so $\Phi_w(\Theta^+_{\alpha'\alpha}) \in CF^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ is a boundary, as $\Theta^+_{\alpha'\alpha}$ is by assumption the top degree element of $HF^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$. The claim now follows. □

We need a final triangle map computation, for a free-stabilization with a degenerated almost complex structure as in Figure 2.31. Compare this to Proposition 2.13.8, where we assumed that $\boldsymbol{\beta}'$ were small Hamiltonian isotopies of the $\boldsymbol{\beta}$ curves. In the following, we do not need that assumption.

We also note that the following follows immediately from [MO10, Proposition 5.2]. Indeed the argument we provide is a trivial modification of theirs.

Lemma 2.13.25. *If* $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$ *is any Heegaard triple, then for a fixed neck length on* c*, and sufficiently stretched almost complex structure on* $c_{\beta\beta'}$*, the triangle maps on the freely-stabilized triple in Figure 2.31 satisfy*

$$F_{\overline{\mathcal{T}}}(\mathbf{x} \times \cdot, \mathbf{y} \times \theta^+) = \begin{pmatrix} F_{\mathcal{T}}(\mathbf{x}, \mathbf{y}) & 0 \\ 0 & F_{\mathcal{T}}(\mathbf{x}, \mathbf{y}) \end{pmatrix},$$

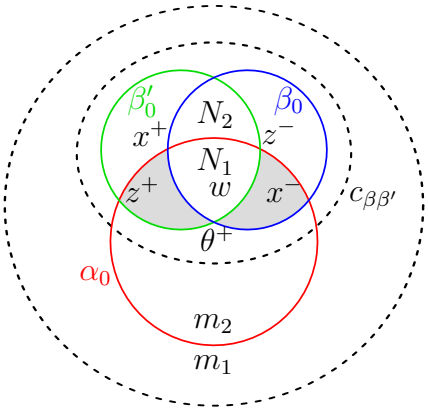

Figure 2.31: A Heegaard triple which has been free-stabilized with the above triple. We stretch the almost complex structure along c and $c_{\beta\beta'}$. The shaded regions are each examples of small triangles which might be counted.

as a matrix decomposition over $\alpha_0 \cap \beta_0$ *and* $\alpha_0 \cap \beta_0'$

Proof. Suppose that $\psi \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times \theta^+, \mathbf{z} \times z)$ is a homology class of triangles with Maslov index zero. When one stretches the along $c_{\beta\beta'}$, from a sequence of holomorphic representatives ofψ we can extract a weak limit to a collection of curves on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ and $(S^2, \alpha_0, \beta_0, \beta_0')$. However by the argument of [MO10, Proposition 5.2], there are no holomorphic curves on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ which have boundary on α_0 as well as the curves in $\boldsymbol{\beta}$ or $\boldsymbol{\beta}'$. Hence we can organize the homology classes into the following collections:

1. A broken holomorphic triangle v_Σ on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ representing a class $\psi_\Sigma^0 \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$;
2. A broken holomorphic triangle v_0 on $(S^2, \alpha_0, \beta_0, \beta_0')$ representing the class $\psi_0 \in \pi_2(x, \theta^+, z)$;
3. A holomorphic curve v_α on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\})$, representing a homology class ϕ_α of $\boldsymbol{\alpha} \cup \{\alpha_0\}$–boundary degenerations.

Write ϕ_α^0 for the homology class $\boldsymbol{\alpha}$–boundary degeneration on $(\Sigma, \boldsymbol{\alpha})$ obtained by chopping off a disk containing α_0 and the region marked with multiplicity m_2. By modifying the Maslov

index computation from Proposition 2.13.8, we see

$$\mu(\psi) = \mu(\psi^0_\Sigma) + \mu(\phi^0_\alpha) + m_1(\phi_\alpha) + m_2(\phi_\alpha) + N_1 + N_2.$$

However all of these terms are nonnegative for sufficiently stretched almost complex structure. Hence if $\mu(\psi) = 0$, we conclude all are zero. From here, the proof proceeds essentially the same as Proposition 2.13.8. One notes that since $N_1(\psi_0) = N_2(\psi_0) = 0$ and $m_1(\psi_0) = m_2(\psi_0)$, it's easy to see that the homology class $\psi_0 \in \pi_2(x, \theta^+, y)$ appearing is constrained to have $(x, y) = (x^+, y^+)$ or (x^-, y^-) (i.e. both x and y have the same grading). Using an adaptation of the argument from Proposition 2.13.8 to count curves on $(S^2, \alpha_0, \beta_0, \gamma_0)$ which match a given divisor, the desired curve count thus follows from a gluing argument. □

Using the above two holomorphic curve computations, we can now give an explicit proof of the π_1–action by computing explicitly a sequence of transition maps on a multi-pointed diagram.

Theorem 2.13.26. *Suppose that $(Y, \mathbf{w})$ is a multi-based 3–manifold and $w \in \mathbf{w}$, and γ is a path on Y which starts and ends at w. Then*

$$\gamma_* \simeq \mathrm{id} + A_\gamma \Phi_w.$$

Proof. We can assume that Y is connected, and by Lemma 2.13.20 we can assume that $\mathbf{w}$ has at least two basepoints. Pick a diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w} \setminus \{w\})$ and consider it's free-stabilization at w. Let us assume that γ is an immersed curve on Σ, starting and ending at w, and intersecting no other basepoints. Let $B \subseteq \Sigma$ be a small disk containing the free-stabilization, but not intersecting $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup (\mathbf{w} \setminus \{w\})$. We wish to consider the effect of pushing B along the path γ, throughout the diagram. Subdivide γ into a concatenation of arcs $\lambda_1, \dots, \lambda_n$ on Σ, such that each λ_i intersects at most one $\boldsymbol{\alpha}$ curve or $\boldsymbol{\beta}$ curve, as in Figure 2.29. Let ϕ_i denote the diffeomorphism corresponding to the basepoint moving map along λ_i.

We can compute the induced map γ_* as the composition

$$\gamma_* = (\lambda_n)_* \circ \cdots \circ (\lambda_1)_* = (\phi_n)_* \circ \cdots \circ (\phi_1)_*.$$

Write w_{i-1} for the incoming basepoint of λ_i and w_i for the outgoing basepoint (so λ_1 is a path from $w = w_0$ to w_1, and λ_n is a path from w_{n-1} to $w_n = w$). We will write

$$p_i = w_i,$$

when we wish to refer to p_i as a point on Σ, in the unstabilized diagram.

Write α_0 and β_0 for the curves of the free-stabilization at w. Define

$$\alpha_i = \phi_i(\alpha_{i-1}), \qquad \beta_i = \phi_i(\beta_{i-1}).$$

We assume that $\alpha_n = \alpha_0$ and $\beta_n = \beta_0$.

Finally, for each $p_i \in \Sigma$, pick choices of gluing data $\mathcal{J}_{p_i}$ for free-stabilizing at $p_i = w_i$, such that the underlying complex structure on the unstabilized diagram is a fixed J_s for all p_i.

Our goal will be to compute each induced map $(\phi_i)_*$, viewed as a map between the complexes

$$\begin{aligned}(\phi_i)_* : CF^-_{\mathcal{J}_{p_{i-1}}(T)}(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_{i-1}\}, \boldsymbol{\beta} \cup \{\beta_{i-1}\}, \mathbf{w} \cup \{w_{i-1}\}) \\ \to CF^-_{\mathcal{J}_{p_i}(T)}(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_i\}, \boldsymbol{\beta} \cup \{\beta_i\}, \mathbf{w} \cup \{w_i\}). \end{aligned} \tag{2.12}$$

If we can compute the explicit chain homotopy of this map in terms of the natural matrix decomposition, then by simply multiplying matrices we can compute the homotopy type of γ_*.

We define a "tautological diffeomorphism map" Γ_i as

$$\Gamma_i : CF^-_{\mathcal{J}_{p_{i-1}}(T)}(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_{i-1}\}, \boldsymbol{\beta} \cup \{\beta_{i-1}\}, \mathbf{w} \cup \{w_{i-1}\})$$

$$\to CF^-_{(\phi_i)_*(\mathcal{J}_{p_{i-1}}(T))}(\Sigma, \phi_i(\boldsymbol{\alpha} \cup \{\alpha_{i-1}\}), \phi_i(\boldsymbol{\beta} \cup \{\beta_{i-1}\}), \mathbf{w} \cup \{w_i\}),$$

simply by pushing forward an intersection point on a diagram to its image under ϕ_i.[1] As an abuse of notation, we will use Γ_i to denote such the map on both the stabilized complexes, and the unstabilized complexes, when convenient.

As a warm up, let us first consider the case that λ_i intersects no $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves. We can assume that the diffeomorphism ϕ_i also fixes the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves.

We can write the $(\phi_i)_*$ map from Equation (2.12) as

$$(\phi_i)_* = \Phi_{(\phi_i)_*\mathcal{J}_{p_{i-1}} \to \mathcal{J}_{p_i}} \circ \Gamma_i.$$

Using Lemma 2.5.12, with the degeneration from Lemma 2.13.10, we can compute that the change of almost complex structure map $\Phi_{\phi_*(\mathcal{J}_{p_{i-1}}(T) \to \mathcal{J}_{p_i}(T)}$ is the diagonal matrix with entries $\Phi_{(\phi_i)_* J_s \to J_s}$ as long as the almost complex structure is stretched along either c_β or c_β. Similarly, the map Γ_i is tautologically the matrix with Γ_i (now on the unstabilized complexes) along the diagonal. Hence

$$(\phi_i)_* = \begin{pmatrix} \Phi_{(\phi_i)_* J_s \to J_s} \Gamma_i & 0 \\ 0 & \Phi_{(\phi_i)_* J_s \to J_s} \Gamma_i \end{pmatrix} \simeq \begin{pmatrix} \mathrm{id} & 0 \\ 0 & \mathrm{id} \end{pmatrix}.$$

The chain homotopy appearing in the above equation can be constructed by observing that $\Phi_{(\phi_i)_* J_s \to J_s} \Gamma_i \simeq \mathrm{id}$ on the unstabilized complexes by naturality, and the differential on the domain and target complexes of $(\phi_i)_*$ are equal, by our computation of the free-stabilized differential in Lemma 2.13.21. To construct a chain homotopy on the stabilized complexes we simply put the chain homotopy from the unstabilized complexes along the diagonal.

In general, we need to consider the case that λ_i intersects a single $\boldsymbol{\alpha}$ curve or a $\boldsymbol{\beta}$ curve. We can compute the diffeomorphism map $(\phi_i)_*$ as a composition by handlesliding the curve that intersects λ_i over B, performing Γ_i, then changing the almost complex structure.

[1] We write Γ_i for this, since it coincides with a continuation map, which counts Maslov index zero disks with time dependent boundary and time dependent almost complex structure. The time dependent almost complex structure and boundary conditions are given by taking an isotopy ϕ_t between ϕ_i and id which is fixed for $t \notin [-1, 1]$ and using it to build a diffeomorphism of $\Sigma \times [0,1] \times \mathbb{R}$ with itself. The time dependent boundary conditions and almost complex structure are given by pushing forward $\boldsymbol{\alpha} \times \{0,1\} \times \mathbb{R}$ and $\boldsymbol{\beta} \times \{0,1\} \times \mathbb{R}$, and the almost complex structure J_s under this diffeomorphism.

First suppose that λ_i intersects a $\boldsymbol{\beta}$ curve. Let $\boldsymbol{\beta}'$ be the curves resulting from handlesliding the $\boldsymbol{\beta}$ curve over the region B containing the free-stabilization. If we destabilize the diagram by removing the curves in B, then the curves in $\boldsymbol{\beta}'$ can be taken to just be Hamiltonian isotopies of the curves in $\boldsymbol{\beta}$. We can, in fact, assume that the diffeomorphism ϕ_i sends $\boldsymbol{\beta}'$ to $\boldsymbol{\beta}$, and maps $\boldsymbol{\alpha}$ to $\boldsymbol{\alpha}$. Hence we can compute the diffeomorphism map $(\phi_i)_*$ from Equation (2.12) as the composition

$$(\phi_i)_* \simeq \Phi_{(\phi_i)_*\mathcal{J}_{p_{i-1}}(T)\to\mathcal{J}_{p_i}(T)} \circ \Gamma_i \circ \Phi^{\boldsymbol{\beta}\cup\{\beta_{i-1}\}\to\boldsymbol{\beta}'\cup\{\beta'_{i-1}\}}_{\boldsymbol{\alpha}\cup\{\alpha_{i-1}\}}.$$

See Figure 2.32, for a schematic of this composition.

By our previous argument, if we stretch sufficiently on c_β, the change of almost complex structure map becomes

$$\Phi_{(\phi_i)_*\mathcal{J}_{p_{i-1}}(T)\to\mathcal{J}_{p_i}(T)} = \begin{pmatrix} \Phi_{(\phi_i)_*J_s\to J_s} & 0 \\ 0 & \Phi_{(\phi_i)_*J_s\to J_s} \end{pmatrix}.$$

Similarly Γ_i is by definition the diagonal matrix with Γ_i along the diagonal.

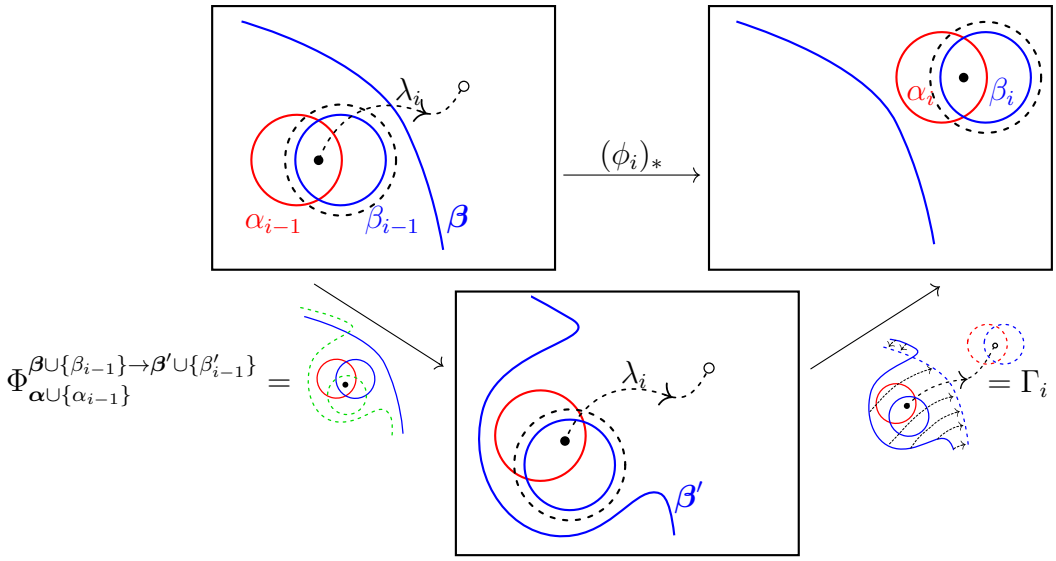

Figure 2.32: The map induced by $(\phi_i)_*$ along the top row can be computed by performing a handleslide of a $\boldsymbol{\beta}$ curve over the disk B, as and then pushing forward under Γ_i. We also need a change of almost complex structure map at the end, which is not shown.

We can compute the triangle map $\Phi^{\boldsymbol{\beta}\cup\{\beta_{i-1}\}\to\boldsymbol{\beta'}\cup\{\beta'_{i-1}\}}_{\boldsymbol{\alpha}\cup\{\alpha_{i-1}\}}$ associated to a handleslide of a $\boldsymbol{\beta}$ curve across β_0 in terms of the map $\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}}$ on the unstabilized complexes by using Lemma 2.13.25, as long as we are using an almost complex structure J_β sufficiently stretched on a circle c_β around the curve β_{i-1}. For such almost complex structures, the triangle counts take the form

$$\Phi^{\boldsymbol{\beta}\cup\{\beta_{i-1}\}\to\boldsymbol{\beta'}\cup\{\beta'_{i-1}\}}_{\boldsymbol{\alpha}\cup\{\alpha_{i-1}\}} = \begin{pmatrix} \Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} & 0 \\ 0 & \Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} \end{pmatrix}.$$

Hence the diffeomorphism map from Equation (2.12) takes the form

$$(\phi_i)_* = \begin{pmatrix} \Phi_{(\phi_i)_*J_s\to J_s} & 0 \\ 0 & \Phi_{(\phi_i)_*J_s\to J_s} \end{pmatrix} \begin{pmatrix} \Gamma_i & 0 \\ 0 & \Gamma_i \end{pmatrix} \begin{pmatrix} \Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} & 0 \\ 0 & \Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} \end{pmatrix}$$

$$= \begin{pmatrix} \Phi_{(\phi_i)_*J_s\to J_s}\circ\Gamma_i\circ\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} & 0 \\ 0 & \Phi_{(\phi_i)_*J_s\to J_s}\circ\Gamma_i\circ\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} \end{pmatrix} \simeq \begin{pmatrix} \mathrm{id} & 0 \\ 0 & \mathrm{id} \end{pmatrix}.$$

The last chain homotopy is given as follows. We note that as $\Phi_{(\phi_i)_*J_s\to J_s}\circ\Gamma_i\circ\Phi^{\boldsymbol{\beta}\to\boldsymbol{\beta'}}_{\boldsymbol{\alpha}} \simeq \mathrm{id}$ on the unstabilized complex by naturality. The differentials of the domain and target of the above matrices are computed in Lemma 2.13.21. As we are using J_β, and we move the free-stabilization point across a $\boldsymbol{\beta}$ curve, the upper right entry of the differential does not change, and a chain homotopy on the unstabilized complex induces a chain homotopy on the stabilized complex.

We now consider the case that λ_i intersects an $\boldsymbol{\alpha}$ curve. Analogously to the previous case when λ_i intersected a $\boldsymbol{\beta}$ curve, we can write the diffeomorphism $(\phi_i)_*$ from Equation (2.12) as the composition

$$(\phi_i)_* = \Phi_{(\phi_i)_*\mathcal{J}_{p_{i-1}}(T)\to\mathcal{J}_{p_i}(T)}\circ\Gamma_i\circ\Phi^{\boldsymbol{\beta}\cup\{\beta_{i-1}\}}_{\boldsymbol{\alpha}\cup\{\alpha_{i-1}\}\to\boldsymbol{\alpha'}\cup\{\alpha'_{i-1}\}}.$$

We can modify Lemma 2.13.25 to allow us to compute $\boldsymbol{\alpha}$ handleslides across α_{i-1}, but we would need to use the almost complex structure J_α which we considered before, which was stretched around the boundary of a ball containing α_{i-1} and its Hamiltonian isotopy α'_{i-1}, in the free-stabilized Heegaard triple. If we want to do an $\boldsymbol{\alpha}$ handleslide using this triangle

map computation, we need to first switch from J_β to J_α, perform the triangle map using the modification of Lemma 2.13.25, then switch back from J_α to J_β. We can compute the change of almost complex structure maps using Proposition 2.13.22. Using matrix notation, we now get that the triangle map, computed with J_β, is chain homotopic to

$$\Phi^{\boldsymbol{\beta}\cup\{\beta_{i-1}\}}_{\boldsymbol{\alpha}\cup\{\alpha_{i-1}\}\to\boldsymbol{\alpha}'\cup\{\alpha'_{i-1}\}} \simeq \begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_{p_i} \\ 0 & \mathrm{id} \end{pmatrix} \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} & 0 \\ 0 & \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} \end{pmatrix} \begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_{p_{i-1}} \\ 0 & \mathrm{id} \end{pmatrix},$$

where $\widetilde{\Psi}_p$ is the map described in Proposition 2.13.22.

We now claim that

$$\begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_{p_i} \\ 0 & \mathrm{id} \end{pmatrix} \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} & 0 \\ 0 & \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} \end{pmatrix} \simeq \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} & 0 \\ 0 & \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} \end{pmatrix} \begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_{p_i} \\ 0 & \mathrm{id} \end{pmatrix}.$$

Multiplying this out, it's sufficient to show that

$$\begin{pmatrix} 0 & \widetilde{\Psi}_{p_i}\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} + \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\widetilde{\Psi}_{p_i} \\ 0 & 0 \end{pmatrix} \simeq 0,$$

as maps on the stabilized complexes. By Lemma 2.13.24, we have that $\widetilde{\Psi}_{p_i}\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} + \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\widetilde{\Psi}_{p_i} \simeq 0$ on the unstabilized complexes. By using the differential computation of Lemma 2.13.4, we can construct a chain homotopy $\overline{H}$ on the stabilized complexes by picking a chain homotopy H on the unstabilized complexes satisfying $\widetilde{\Psi}_{p_i}\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} + \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}\widetilde{\Psi}_{p_i} = H\partial_{\alpha\beta} + \partial_{\alpha'\beta}H$ and putting H into the upper right component of $\overline{H}$, while defining the other components of $\overline{H}$ to be zero.

Hence, multiplying out the previous contributions, we get that

$$(\phi_i)_* \simeq \Phi_{(\phi_i)_*\mathcal{J}_{p_{i-1}}(T)\to\mathcal{J}_{p_i}(T)} \circ \Gamma_i \circ \Phi^{\boldsymbol{\beta}\cup\{\beta_{i-1}\}}_{\boldsymbol{\alpha}\cup\{\alpha_{i-1}\}\to\boldsymbol{\alpha}'\cup\{\alpha'_{i-1}\}}$$

$$\simeq \begin{pmatrix} \Phi_{(\phi_i)_*J_s\to J_s}\Gamma_i\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} & 0 \\ 0 & \Phi_{(\phi_i)_*J_s\to J_s}\Gamma_i\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} \end{pmatrix} \begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_{p_i} + \widetilde{\Psi}_{p_{i-1}} \\ 0 & \mathrm{id} \end{pmatrix}.$$

We can view the left matrix in the above product as an endomorphism of the single free-stabilized complex $CF^-_{\mathcal{J}_{p_i}(T)}(\Sigma, \boldsymbol{\alpha}\cup\{\alpha_i\}, \boldsymbol{\beta}\cup\{\beta_i\}, \mathbf{w}\cup\{w_i\})$. On the unstabilized complexes,

we have $\Phi_{(\phi_i)_* J_s \to J_s} \Gamma_i \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} \simeq \mathrm{id}$ by naturality. In light of the explicit description of the differential for a free-stabilized complex with differential J_β from Lemma 2.13.21, we note that it's easy to use such a chain homotopy on the unstabilized complex to build a chain homotopy on the free-stabilized complex (just put the chain homotopy along the diagonal).

Hence the $(\phi_i)_*$ map from Equation (2.12) satisfies the equivalence

$$(\phi_i)_* \simeq \begin{pmatrix} \mathrm{id} & \widetilde{\Psi}_{p_{i-1}} + \widetilde{\Psi}_{p_i} \\ 0 & \mathrm{id} \end{pmatrix}.$$

Note that $\widetilde{\Psi}_{p_{i-1}} + \widetilde{\Psi}_{p_i}$ counts Maslov index one disks with an extra factor of the difference in multiplicities across the $\boldsymbol{\alpha}$ curve that we handleslid.

Thus we can compute the composition γ_* then as the product of the matrices involved in the composition $(\phi_n)_* \circ \cdots \circ (\phi_1)_*$. Using our above computation of the homotopy type of each $(\phi_i)_*$, the composition of all of them is chain homotopic to

$$\begin{pmatrix} \mathrm{id} & A \\ 0 & \mathrm{id} \end{pmatrix}$$

where A is the sum of the expressions

$$\widetilde{\Psi}_{p_{i-1}} + \widetilde{\Psi}_{p_i}$$

over the i such that λ_i crossed a single $\boldsymbol{\alpha}$ curve (and p_{i-1} and p_i are the points on either side of the $\boldsymbol{\alpha}$ curve). By definition, this sum is equal to A_γ on the unstabilized complex. As this coincides with the matrix for $\mathrm{id} + \Phi_w A_\gamma$ from Equation (2.11), we conclude that $\gamma_* \simeq \mathrm{id} + A_\gamma \Phi_w$. □

2.13.4 Basepoint swapping map

As a final application of the computation of the basepoint moving map, we consider a final diffeomorphism. If w_1 and w_2 are two basepoints and λ is a path between the two, we can consider the effect of swapping w_1 and w_2 along the path λ. In terms of cobordisms, this

corresponds to the path cobordism with underlying 4–manifold $Y \times [0,1]$ with a path from $\{w_1\} \times \{0\}$ to $\{w_2\} \times \{1\}$ and a path from $\{w_2\} \times \{0\}$ to $\{w_1\} \times \{1\}$ such that the projection of each path into Y is λ.

If $\phi : (Y, \mathbf{w}) \to (Y, \mathbf{w})$ denotes this basepoint swapping diffeomorphism, we note that ϕ only induces an equivariant map on $CF^-(Y, \mathbf{w}, \mathfrak{s})$ if we give w_1 and w_2 the same color, i.e. if we identity the two variables. Of course if they aren't given the same color, then we can't color the strands of the path cobordism to get a graph cobordism with the proper decoration.

Proposition 2.13.27. *If $\phi : (Y, \mathbf{w}) \to (Y, \mathbf{w})$ is the diffeomorphism swapping w_1 and w_2 along λ, described above, then the induced map ϕ_*, an endomorphism of the complex where U_{w_1} and U_{w_2} have been identified, has the equivariant chain homotopy type*

$$\phi_* \simeq \Phi_{w_1} A_\lambda + A_\lambda \Phi_{w_2} \simeq A_\lambda \Phi_{w_1} + \Phi_{w_2} A_\lambda.$$

Proof. First note that the second equivalence follows immediately from Lemma 2.13.17.

For the first equivalence, one strategy would be to use our computations to explicitly compute the map on the level of diagrams. The diffeomorphism ϕ induces the tautological map between $CF^-(\mathcal{H}_1)$ and $CF^-(\mathcal{H}_2)$. To get an endomorphism of $CF^-(\mathcal{H}_1)$, we could compose with the change of diagrams map $\Phi_{\mathcal{H}_2 \to \mathcal{H}_1}$. We computed this map in Proposition 2.13.1, but we remark that we'd also have to compose with a (nontrivial) change of almost complex structure map, since the transition map in that Proposition went between $(\mathcal{H}_1, J_\alpha)$ and $(\mathcal{H}_2, J_\beta)$. Although quite doable, this is somewhat unenlightening, so instead we will use our computation of the basepoint moving map in terms of other natural maps from Theorem 2.13.13, and the fact that the maps appearing in the basepoint moving map are natural.

Let $w' \notin \mathbf{w}$ be a new basepoint, given the same color as w_1 and w_2, which we will add to help out. Pick paths λ_1 and λ_2 from w_1 to w' and from w' to w_2 (respectively) such that the concatenation $\lambda_1 * \lambda_2$ is homotopic to λ. Using Theorem 2.13.13 to decompose ϕ_* into a composition of three basepoint moving maps, we have

$$\phi_* \simeq (S^-_{w'} A_{\lambda_2} S^+_{w_2})(S^-_{w_2} A_\lambda S^+_{w_1})(S^-_{w_1} A_{\lambda_1} S^+_{w'}),$$

i.e. we first move w_1 to w' along λ_1, then move w_2 to w_1 along λ, and then finally move w' to w_2 along λ_2. Using the relations

$$S_w^+ S_w^- \simeq \Phi_w, \qquad \Phi_w A_e + A_e \Phi_w \simeq \mathrm{id}, \qquad \text{and} \qquad S_{w'} \Phi_w + \Phi_w S_{w'} \simeq 0$$

(for a path e with exactly one endpoint equal to w) from Lemma 2.13.17, repeatedly, the above equivalence yields

$$\begin{aligned}
\phi_* &\simeq S_{w'} A_{\lambda_2} \Phi_{w_2} A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ \\
&\simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ + S_{w'}^- A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ \\
&\simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda A_{\lambda_1} S_{w'}^+ \Phi_{w_1} + \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda S_{w'}^+ + S_{w'}^- A_\lambda \Phi_{w_1} A_{\lambda_1} S_{w'}^+ \\
&\simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda A_{\lambda_1} S_{w'}^+ \Phi_{w_1} + \Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda S_{w'}^+ + \Phi_{w_1} S_{w'}^- A_\lambda A_{\lambda_1} S_{w'}^+ + S_{w'}^+ A_{\lambda_1} S_{w'}^+.
\end{aligned}$$

We note that since $A_\lambda = A_{\lambda_2} + A_{\lambda_1}$ when all are defined, we have

$$A_{\lambda_2} A_\lambda A_{\lambda_1} \simeq A_{\lambda_2}(A_{\lambda_1} + A_{\lambda_2}) A_{\lambda_1} \simeq A_{\lambda_2} A_{\lambda_1}^2 + A_{\lambda_2}^2 A_{\lambda_1} \simeq U(A_{\lambda_1} + A_{\lambda_2}) U A_\lambda.$$

Here we are writing U for any of U_{w_1}, U_{w_2} or $U_{w'}$ (since we've formally identified all three). Using this, we see

$$\Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda A_{\lambda_1} S_{w'}^+ \Phi_{w_1} \simeq U S_{w'}^- A_\lambda S_{w'}^+ \Phi_{w_1} + \simeq U A_\lambda S_{w'}^- S_{w'}^+ \Phi_{w_1} \simeq 0.$$

Also, using Lemma 2.5.6 to move $S_{w'}^\pm$ past A_λ, we see

$$\Phi_{w_2} S_{w'}^- A_{\lambda_2} A_\lambda S_{w'}^+ \simeq \Phi_{w_2} S_{w'}^- A_{\lambda_2} S_{w'}^+ A_\lambda \simeq \Phi_{w_2} A_\lambda,$$

$$\Phi_{w_1} S_{w'}^- A_\lambda A_{\lambda_1} S_{w'}^+ \simeq \Phi_{w_1} A_\lambda S_{w'}^- A_{\lambda_1} S_{w'}^+ \simeq \Phi_{w_1} A_\lambda$$

and

$$S_{w'}^+ A_{\lambda_1} S_{w'}^+ \simeq \mathrm{id}\,.$$

Adding up these contributions, we see that

$$\phi_* \simeq \Phi_{w_2} A_\lambda + \Phi_{w_1} A_\lambda + \mathrm{id},$$

which is homotopy equivalent to the expression in the proposition statement since $\Phi_{w_1} A_\lambda + A_\lambda \Phi_{w_1} \simeq \mathrm{id}$ by Lemma 2.13.17. □

CHAPTER 3

Quasi-stabilization and basepoint moving maps in link Floer homology

We analyze the effect of adding, removing, and moving basepoints on link Floer homology. We prove that adding or removing basepoints via a procedure called quasi-stabilization is a natural operation on a certain version of link Floer homology, which we call CFL^{∞}_{UV}. We consider the effect on the full link Floer complex of moving basepoints, and develop a simple calculus for moving basepoints on the link Floer complexes. We apply it to compute the effect of several diffeomorphisms corresponding to moving basepoints. Using these techniques we prove a conjecture of Sarkar about the map on the full link Floer complex induced by a finger move along a link component.

3.1 Introduction

Introduced by Ozsváth and Szabó, Heegaard Floer homology associates algebraic invariants to closed three manifolds. To a three manifold Y with embedded nullhomologous knot K, there is a refinement of Heegaard Floer homology called knot Floer homology, introduced by Ozsváth and Szabó in [OS04c] and independently by Rasmussen in [Ras03]. A similar invariant was defined by Ozsváth and Szabó for links in [OS08].

To a nullhomologous knot $K \subseteq Y$ with two basepoints z and w and a relative Spin^c structure $\mathfrak{t} \in \text{Spin}^c(Y, K)$, in [OS04c], Ozsváth and Szabó define a $\mathbb{Z} \oplus \mathbb{Z}$–filtered chain complex $CFK^{\infty}(Y, K, w, z, \mathfrak{t})$. The $\mathbb{Z} \oplus \mathbb{Z}$–filtered chain homotopy type of $CFK^{\infty}(Y, K, w, z, \mathfrak{t})$ is an invariant of the data $(Y, K, w, z, \mathfrak{t})$.

One of the nuances of Heegaard Floer homology is the dependence on basepoints. In the case of closed three manifolds, if $\mathbf{w} \subseteq Y$ is a collection of basepoints, $w \in \mathbf{w}$ and γ is a curve in $\pi_1(Y, w)$, then one can consider the diffeomorphism ϕ_γ resulting from a finger move along γ. According to [JT12], the based mapping class group $MCG(Y, w)$ acts on $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$ and hence there is an induced map $(\phi_\gamma)_*$ on the closed three manifold invariant $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$, which is a $\mathbb{Z}_2[U]$–equivariant chain homotopy type. In Chapter 2, the author computes the equivariant chain homotopy type of $(\phi_\gamma)_*$ to be

$$(\phi_\gamma)_* \simeq \mathrm{id} + [\gamma] \circ \Phi_w,$$

where $[\gamma]$ is the $\Lambda^*(H_1(Y; \mathbb{Z})/\operatorname{Tors})$ action and Φ_w is an analog of a map appearing also on link Floer homology, which we describe below.

In this paper, we consider the analogous question about basepoint dependence for link Floer homology. In link Floer homology, the basepoints are constrained to be on the link component, so the analogous operation is to consider the map on link Floer homology induced by a finger move ς around a link component, in the positive direction according to the link's orientation. Using grid diagrams, Sarkar computes in [`?Smovingbasepoints`] the map associated to the diffeomorphism ς on a certain version of link Floer homology (the associated graded complex) for links in S^3. For links in arbitrary three manifolds, and for the induced map on the full link Floer complex, he conjectures the formula. We prove his formula in full generality (Theorem B), but before we state that theorem we will provide a brief description of the complexes and maps which appear.

We will work with a slightly different version of CFL^∞ than the one which most often appears in the literature. To a multibased link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ inside of Y, as well as a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, we construct a chain complex

$$CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s}),$$

which is a module over the polynomial ring $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$, generated by variables U_w with $w \in \mathbf{w}$ and V_z with $z \in \mathbf{z}$. The module $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$ has generators of the form

$$\mathbf{x} \cdot U_{\mathbf{w}}^I V_{\mathbf{z}}^J = \mathbf{x} \cdot U_{w_1}^{i_1} \cdots U_{w_n}^{i_n} V_{z_1}^{j_1} \cdots V_{z_n}^{j_n},$$

for multi-indices $I = (i_1, \dots, i_n)$ and $J = (j_1, \dots, j_n)$, though we identify two variables V_z and $V_{z'}$ if z and z' are on the same link component. As such, $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$ has a filtration by $\mathbb{Z}^{|\mathbf{w}|} \oplus \mathbb{Z}^{|L|}$ given by filtering over powers of the variables, where $|L|$ denotes the set of components of L.

As with the free-stabilization maps from Chapter 2, to define functorial maps corresponding to adding or removing basepoints in link Floer homology, we must work with colored complexes. A coloring $(\sigma, \mathfrak{P})$ of a link with basepoints, $(L, \mathbf{w}, \mathbf{z})$, is a set $\mathfrak{P}$ indexing a collection of formal variables, together with a map $\sigma\colon\ \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ which maps all $\mathbf{z}$–basepoints on a component of L to the same color. Given a coloring $(\sigma, \mathfrak{P})$ of a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, we create a $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain complex

$$CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}).$$

The powers of the $U_{\mathfrak{P}}$ variables yield a filtration by $\mathbb{Z}^{\mathfrak{P}}$, which we call the $\mathfrak{P}$–filtration.

In the context link Floer homology, analogous to adding or removing a free basepoint in a closed 3–manifold, one can add or remove a pair of adjacent basepoints, w and z, on a link component. Some authors refer to this operation as a "special stabilization". Manolescu and Ozsváth consider certain questions about the operation in [MO10], calling it "quasi-stabilization", which is the phrase we will use. A full description and proof of naturality of the operation has not been completed, so we do that in this paper:

Theorem A. *Suppose z and w are new basepoints on a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, ordered so that w comes after z, which aren't separated by any basepoints in $\mathbf{w}$ or $\mathbf{z}$. If $\sigma\colon\ \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ is a coloring which is extended by $\sigma'\colon\ \mathbf{w} \cup \mathbf{z} \cup \{w, z\} \to \mathfrak{P}$, then there are $\mathfrak{P}$–filtered $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain maps*

$$S^+_{w,z}\colon\ CFL^\infty_{UV}(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}) \to CFL^\infty_{UV}(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S^-_{w,z}\colon\ CFL^\infty_{UV}(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to CFL^\infty_{UV}(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

which are well defined invariants, up to $\mathfrak{P}$–filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant chain homotopy. If z comes after w, there are maps $S^+_{z,w}$ and $S^-_{z,w}$ defined analogously.

Following [Sar15], we consider endomorphisms Φ_w and Ψ_z of $CFL^\infty_{UV}(Y,\mathbb{L},\mathfrak{s})$ (denoted $\Phi_{i,j}$ and $\Psi_{i,j}$ in his notation). We can think of the maps Φ_w and Ψ_z as formal derivatives of the differential ∂ with respect to the variables U_w and V_z respectively. The maps Φ_w and Ψ_z are invariants of $CFL^\infty_{UV}(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$ up to $\mathfrak{P}$–filtered chain homotopy.

The maps Ψ_z can be thought of as an analog for the relative homology maps A_λ defined in Chapter 2 for the closed three manifold invariants, since they play the role in the basepoint moving maps for link Floer homology that the relative homology maps introduced in Chapter 2 played in the basepoint moving maps for the closed three-manifolds invariants. Indeed the objects $CFL^\infty_{UV}(Y,\mathbb{L},\mathfrak{s})$ and the maps Ψ_z and $S^\pm_{w,z}$ fit into the framework of a "graph TQFT" for surfaces embedded in four-manifolds with some extra decoration, similar to the TQFT for $\widehat{HFL}$ constructed using sutured Floer homology by Juhász in [Juh16] and considered further in [JM16b] for concordances. Such a TQFT construction for CFL^∞_{UV} will appear in a future paper.

We finally state Sarkar's conjecture, cast into the framework of $CFL^\infty_{UV}(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$:

Theorem B. *Suppose that $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ is a multibased link in an arbitrary 3–manifold Y and K is a component of L. Suppose that the basepoints on K are $z_1,w_1,\dots,z_n,w_n$. Letting ς denote the diffeomorphism resulting from a finger move around a link component K, the induced map ς_* on $CFL^\infty_{UV}(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$ has the $\mathfrak{P}$–filtered equivariant chain homotopy type*

$$\varsigma_* \simeq \mathrm{id} + \Phi_K\Psi_K$$

where

$$\Phi_K=\sum_{j=1}^n \Phi_{w_j} \qquad \text{and} \qquad \Psi_K=\sum_{j=1}^n \Psi_{z_j}.$$

Sarkar's conjecture for the effect on the filtered link Floer complex, which we will denote by $CFL^\infty(Y,\mathbb{L},\mathfrak{t})$ for $\mathfrak{t}$ a relative $\mathrm{Spin}^c(Y,L)$ structure, follows by setting $(\sigma,\mathfrak{P})$ to be the trivial coloring (i.e. $\mathfrak{P}=\mathbf{w}\cup|L|$ and $\sigma\colon\ (\mathbf{w}\cup\mathbf{z})\to(\mathbf{w}\cup|L|)$ the natural map) since the complex

$CFL^\infty(Y,\mathbb{L},\mathfrak{t})$ becomes a $\mathbb{Z}_2$–subcomplex of $CFL^\infty_{UV}(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$ which is preserved by ς_*, where $\mathfrak{s}$ is the underlying Spin^c structure associated to the relative Spin^c structure $\mathfrak{t}$.

There are several other formulations of the above conjecture for different versions of link Floer homology. For example, the conjectured formula for ς_* on $CFK^\infty(S^3,K)$ for $K\subseteq S^3$ is useful for computations in the involutive Heegaard Floer homology theory developed by Hendricks and Manolescu (see [HM15, Sec. 6]). In their notation, for a choice of diagrams, the complex $CFK^\infty(S^3,K)$ for a knot $K\subseteq S^3$ is generated by elements of the form $[\mathbf{x},i,j]$ where i and j satisfy $A(\mathbf{x})=i-j$, and A denotes the Alexander grading. In their notation, the U map takes the form $U\cdot[\mathbf{x},i,j]=[\mathbf{x},i-1,j-1]$. Again the complex $CFK^\infty(S^3,K)$ is a $\mathbb{Z}_2$–subcomplex

$$CFK^\infty(S^3,K)\subseteq CFL^\infty_{UV}(S^3,K,w,z,\mathfrak{s}_0)$$

which is preserved by ς_*. Recasting Theorem B into this notation and recalling that we are using coefficients in $\mathbb{Z}_2$, we arrive at the following:

Corollary C. *For a knot $K\subseteq S^3$, the involution ς_* on $CFK^\infty(S^3,K)$ takes the form*

$$\varsigma_*\simeq 1+U^{-1}\Bigg(\sum_{\substack{i,j\geq 0\\ i\ odd}}\partial_{ij}\Bigg)\circ\Bigg(\sum_{\substack{i,j\geq 0\\ j\ odd}}\partial_{ij}\Bigg),$$

if we write the differential $\partial=\sum_{i,j\geq 0}\partial_{ij}$. Here ∂_{ij} decreases the first filtration by i, and the second filtration by j.

For other flavors, such as $\widehat{CFL}$ or CFL^-, the formula conjectured by Sarkar also follows, since those cases correspond to setting various variables equal to zero in the formula for ς_*.

In addition, we consider the effect of another diffeomorphism which naturally appears in the involutive setting (cf. [HM15]). Suppose that K is a component of a link $\mathbb{L}$ and suppose that the basepoints of K are $z_1,w_1,\cdots,z_n$ and w_n, appearing in that order. We can consider the diffeomorphism $\tau\colon (Y,\mathbb{L})\to(Y,\mathbb{L})$ which twists $(\frac{1}{n})^{\text{th}}$ of the way around K. The diffeomorphism τ maps z_i and w_i to z_{i+1} and w_{i+1} respectively, with indices taken modulo n. If $(\sigma,\mathfrak{P})$ is a coloring of $\mathbb{L}$ which sends all of the $\mathbf{w}$–basepoints on K to the same

color, then τ naturally induces an automorphism of

$$CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}).$$

Using the techniques of this paper, we can compute the following:

Theorem D. *Suppose that $\mathbb{L}$ is an embedded link in Y, and K is a component of $\mathbb{L}$ with basepoints $z_1, w_1, \dots, z_n$, and w_n, appearing in that order. Assume that $n > 1$. If τ denotes the diffeomorphism induced by twisting $(\frac{1}{n})^{th}$ of the way around K, then for a coloring where all $\mathbf{w}$–basepoints on K have the same color, we have*

$$\tau_* \simeq (\Psi_{z_1}\Phi_{w_1}\Psi_{z_2}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{z_n}\Phi_{w_n}) + (\Phi_{w_1}\Psi_{z_2}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{z_n}).$$

3.1.1 Organization

In Section 3.2 we define the complexes which will appear in this paper, as well as their algebraic structures as $\mathfrak{P}$–filtered chain complexes over certain modules. In Section 3.3 we define the maps Φ_w and Ψ_z which feature prominently in this paper. In Sections 3.5–3.7 we define quasi-stabilization maps $S^{\pm}_{w,z}$ and show that they are independent of the choice of diagrams and auxiliary data, proving Theorem A. In Sections 3.8 and 3.9 we prove useful relations amongst the maps Ψ_z, Φ_w and $S^{\pm}_{w,z}$. In Section 3.10 we compute several maps associated with moving basepoints, proving Theorems B and D.

3.2 Background, the complexes CFL^∞_{UV}, and $\mathfrak{P}$–filtrations

In this section, we provide some background, and describe the complexes $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$ which will appear.

3.2.1 Spin^c structures and relative Spin^c structures

In [OS04e], Ozsváth and Szabó define a Spin^c structure on Y to be a homology class of nonvanishing vector fields on Y. To a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ they define a

map

$$\mathfrak{s}_{\mathbf{w}}\colon\ \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \mathrm{Spin}^c(Y).$$

The vector field $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})$ is obtained by taking an upward gradient-like vector field associated to a Morse function yielding $\mathcal{H}$, and modifying it in a neighborhood of the flowlines passing through the points in $\mathbf{w}$ and $\mathbf{x}$ to obtain a nonvanishing vector field. In [OS08], Ozsváth and Szabó provide a notion of relative Spin^c structures for a link in a 3–manifold. These are homology classes of vector fields on $Y \setminus N(L)$ which are tangent to the torus $\partial N(L)$. They define a map

$$\mathfrak{s}_{\mathbf{w},\mathbf{z}}\colon\ \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \mathrm{Spin}^c(Y,L).$$

We note that in general there are two natural ways to obtain an absolute Spin^c structure from a relative Spin^c structure. We take the convention that the filling map covers the map $\mathfrak{s}_{\mathbf{w}}$ (compare [OS08, Section 3.7]). In more generality, one has $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) - \mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = PD[L]$, so if we restrict to links whose total homology class vanishes, then there is no distinction. Since we include the versions of link Floer homology which use relative Spin^c structures only for comparisons sake, whenever we consider relative Spin^c structures, we will assume that Y is a integer homology sphere. We will primarily be interested in working with the version CFL^∞_{UV}, which uses absolute Spin^c structures.

3.2.2 The complex $CFL^\infty_{UV}(Y,\mathbb{L},\mathfrak{s})$

Here we describe the uncolored complex $CFL^\infty_{UV}(Y,\mathbb{L},\mathfrak{s})$. We first describe an intermediate object, $CFL^\infty_{UV,0}(Y,\mathbb{L},\mathfrak{s})$.

Let $\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}, V_{\mathbf{z}}, V_{\mathbf{z}}^{-1}]$ denote the ring generated by variables U_w, V_z and their inverses U_w^{-1}, and V_z^{-1} for $w \in \mathbf{w}$ and $z \in \mathbf{z}$. Given a diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for $(Y, L, \mathbf{w}, \mathbf{z})$, we define $CFL^\infty_{UV,0}(\mathcal{H},\mathfrak{s})$ to be the free $\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}, V_{\mathbf{z}}, V_{\mathbf{z}}^{-1}]$–module generated by $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$. We refer the reader to, e.g., [OS08], for the definition of a Heegaard diagram for a link, though we emphasize that in light of the results of [JT12], we must assume that

$$\mathbf{w} \cup \mathbf{z} \subseteq \Sigma \subseteq Y$$

and that the embedding of Σ in Y is part of the data of a Heegaard splitting.

We now define a map

$$\partial\colon CFL^{\infty}_{UV,0}(\mathcal{H},\mathfrak{s}) \to CFL^{\infty}_{UV,0}(\mathcal{H},\mathfrak{s})$$

by

$$\partial(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

The map ∂ does not square to zero, but we do have the following:

Lemma 3.2.1. *The map $\partial\colon CFL^{\infty}_{UV,0}(\mathcal{H},\mathfrak{s}) \to CFL^{\infty}_{UV,0}(\mathcal{H},\mathfrak{s})$ satisfies*

$$\partial^2 = \sum_{K\in|L|} (U_{w_{K,1}} V_{z_{K,1}} + V_{z_{K,1}} U_{w_{K,2}} + U_{w_{K,2}} V_{z_{K,2}} + \cdots V_{z_{K,n_K}} U_{w_{K,1}}),$$

where $w_{K,1}, z_{K,1}, \ldots, w_{K,n_K}, z_{K,n_K}$ are the basepoints on the link component K, in the order that they appear on K.

Proof. This follows from the usual proof that the differential squares to zero, now just counting boundary degenerations carefully. If there are exactly two basepoints, there are no boundary degenerations by [OS08, Theorem 5.5], and the above formula is satisfied. If there are more than two, then each $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–degeneration has a unique holomorphic representative by [OS08, Theorem 5.5] and each crosses over a $\mathbf{w}$–basepoint and a $\mathbf{z}$–basepoint. The formula follows. □

To get a chain complex, we must color $CFL^{\infty}_{UV,0}(Y,\mathbb{L},\mathfrak{s})$ by setting certain variables equal. Let $C_{\mathbb{L}}$ denote the ideal generated by elements of the form $V_{z_i} - V_{z_j}$ where z_i and z_j are in the same link component. We let $\mathcal{L} = \mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}, V_{\mathbf{z}}, V_{\mathbf{z}}^{-1}]/C_{\mathbb{L}}$.

We now define

$$CFL^{\infty}_{UV}(\mathcal{H},\mathfrak{s}) = CFL^{\infty}_{UV,0}(\mathcal{H},\mathfrak{s}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}},U_{\mathbf{w}}^{-1},V_{\mathbf{z}},V_{\mathbf{z}}^{-1}]} \mathcal{L}.$$

We have the following:

Lemma 3.2.2. *The map ∂ defined above is a differential on $CFL^\infty_{UV}(\mathcal{H},\mathfrak{s})$, i.e. $\partial^2=0$.*

Proof. This follows from the formula in Lemma 3.2.1 since the module $\mathcal{L}$ simply identifies all V_z variables for z which lie in the same link component. □

Remark 3.2.3. There are other modules that we could tensor with to make the differential square to zero. The module $\mathcal{L}$ is actually a quite natural choice. As we will see in the proof of Proposition 3.5.3, terms of the form $V_z+V_{z'}$ appear in the differential after quasi-stabilization, and these terms must be zero for $S^{\pm}_{w,z}$ to be chain maps.

The $\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]$–module $CFL^\infty_{UV}(\mathcal{H},\mathfrak{s})$ has a natural $\mathbb{Z}^{|\mathbf{w}|}\oplus\mathbb{Z}^{|L|}$ filtration given by filtering over powers of the variables $U_{\mathbf{w}}$ and $V_{\mathbf{z}}$.

There are of course many different Heegaard diagrams $\mathcal{H}$ for a given multi-based link $(L,\mathbf{w},\mathbf{z})$. As in the case of closed three manifolds, using [JT12], given two diagrams $\mathcal{H}$ and $\mathcal{H}'$, there is a $\mathbb{Z}^{|\mathbf{w}|}\oplus\mathbb{Z}^{|L|}$–filtered map

$$\Phi_{\mathcal{H}\to\mathcal{H}'}\colon CFL^\infty_{UV}(\mathcal{H},\mathfrak{s})\to CFL^\infty_{UV}(\mathcal{H}',\mathfrak{s})$$

which is a filtered chain homotopy equivalence, and is an invariant up to $\mathbb{Z}^{|\mathbf{w}|}\oplus\mathbb{Z}^{|L|}$–filtered chain homotopy. The maps $\Phi_{\mathcal{H}\to\mathcal{H}'}$ are functorial in the sense that if $\mathcal{H},\mathcal{H}'$ and $\mathcal{H}''$ are three diagrams, then

$$\Phi_{\mathcal{H}'\to\mathcal{H}''}\circ\Phi_{\mathcal{H}\to\mathcal{H}'}\simeq\Phi_{\mathcal{H}\to\mathcal{H}''}.$$

The strongest invariant, which we will occasionally refer to as the **coherent filtered chain homotopy type**, is the collection of all of the complexes $CFL^\infty_{UV}(\mathcal{H},\mathfrak{s})$ for all admissible diagrams, $\mathcal{H}$, for $(Y,L,\mathbf{w},\mathbf{z})$, as well as the maps $\Phi_{\mathcal{H}\to\mathcal{H}'}$. We let

$$CFL^\infty_{UV}(Y,\mathbb{L},\mathfrak{s})$$

denote this invariant. Note that since we are working with embedded Heegaard surfaces, the set of Heegaard diagrams for a link is a set, and not a proper class.

Remark 3.2.4. As we remarked earlier, since $CFL^\infty_{UV}(Y,\mathbb{L},\mathfrak{s})$ is generated by intersection points with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})=\mathfrak{s}$, there is some asymmetry between the $\mathbf{w}$ and $\mathbf{z}$ basepoints in the

construction of CFL^∞_{UV}. We note that $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})-\mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = PD[L]$, so if L is null-homologous, this doesn't affect the chain complexes. As a toy example, one can consider $(S^1 \times S^2, S^1 \times \{pt\})$ to see how the modules change over different choices of $\mathfrak{s}$.

3.2.3 Other versions of the link Floer complex

Supposing for simplicity that Y is an integer homology sphere, we briefly describe a complex $CFL^\infty(Y, \mathbb{L}, \mathfrak{t})$, for a relative Spin^c structure $\mathfrak{t} \in \text{Spin}^c(Y, L)$. It will not feature in any of the sections after this, but we describe it as a comparison with CFL^∞_{UV}. Let μ_i be a positive meridian of the ith link component. The complex $CFL^\infty(Y, \mathbb{L}, \mathfrak{t})$ is defined as the subcomplex of $CFL^\infty_{UV}(\mathcal{H}, \mathfrak{s})$ generated over $\mathbb{Z}_2$ by elements of the form $\mathbf{x} \cdot U^I_{\mathbf{w}} V^J_{\mathbf{z}}$, where

$$J \cdot PD[M] = (\mathfrak{t} - \mathfrak{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x})) + I \cdot PD[M], \tag{3.1}$$

where PD denotes Poincaré duality. Here if $J = (j_1, \dots, j_\ell)$ then $J \cdot PD[M]$ is defined to be $j_1 \cdot PD[\mu_1] + \cdots + j_\ell \cdot PD[\mu_\ell]$ and $I \cdot PD[M]$ is defined similarly.

In the case that $\mathbb{L} = (K, w, z)$ is a knot with exactly two basepoints, we see that $CFL^\infty(Y, K, w, z, \mathfrak{t})$ is generated by elements of the form $\mathbf{x} \cdot U^i_w V^j_z$ with

$$j \cdot PD[\mu] = (\mathfrak{t} - \mathfrak{s}_{\mathbf{w},\mathbf{z}}(\mathbf{x})) + i \cdot PD[\mu],$$

which is exactly the complex $CFK^\infty(K, \mathfrak{t})$ found in [OS04c]. More often one writes $[\mathbf{x}, i, j]$ for what we write $\mathbf{x} \cdot U^{-i}_w V^{-j}_z$. Most authors also write U for the action defined by $U \cdot [\mathbf{x}, i, j] = [\mathbf{x}, i-1, j-1]$, though in our notation this action corresponds to multiplication by $U_w V_z$. It's also common to consider an object $CFK^\infty(Y, K)$, generated by monomials $U^i_w V^j_z \cdot \mathbf{x}$ satisfying $A(\mathbf{x}) = i - j$, where A is the symmetrized Alexander grading.

Given a relative Spin^c structure $\mathfrak{t} \in \text{Spin}^c(Y, L)$, we write

$$\iota_{\mathfrak{t}} \colon CFL^\infty(Y, L, \mathfrak{t}) \hookrightarrow CFL^\infty_{UV}(Y, L, \mathfrak{s})$$

for inclusion.

As a direct sum of $\mathbb{Z}_2$–modules, we have

$$CFL^\infty_{UV}(\mathcal{H},\mathfrak{s}) = \bigoplus_{\mathfrak{t}\in \mathrm{Spin}^c(Y,L)} CFL^\infty(\mathcal{H},\mathfrak{t}).$$

Write $\pi_{\mathfrak{t}}\colon CFL^\infty_{UV}(\mathcal{H},\mathfrak{s}) \to CFL^\infty(\mathcal{H},\mathfrak{t})$ for the projection onto $CFL^\infty(\mathcal{H},\mathfrak{t})$.

Finally, we note that in $CFL^\infty(Y,\mathbb{L},\mathfrak{t})$ the multi-index J in a monomial $U^I_{\mathbf{w}}V^J_{\mathbf{z}}\cdot\mathbf{x} \in CFL^\infty(Y,\mathbb{L},\mathfrak{t})$ is determined by multi-index I, as well as the choice of $\mathfrak{t}$. As such the full link Floer complex in [OS08] is described instead as the module generated by monomials $U^I_{\mathbf{w}}\cdot\mathbf{x}$, but with a filtration by $\mathrm{Spin}^c(Y,L)$. Given a $\mathfrak{t}\in\mathrm{Spin}^c(Y,L)$, it is straightforward to write down an isomorphism of filtered chain complexes between these two objects.

3.2.4 Colorings and $\mathfrak{P}$–filtrations

As was the case in Chapter 2, to define functorial maps it is important to work in a category of chain complexes over a fixed ring. As the link Floer complexes are modules over a ring which depends on the link, we need to formally "color" the complexes to make them modules over a fixed ring. Different choices of base rings will be useful for different applications, but for a single computation, a single ring must be fixed.

If $\mathfrak{P}$ is a finite set, we let $\mathbb{Z}_2[U_{\mathfrak{P}},U_{\mathfrak{P}}^{-1}]$ denote the ring generated by the formal variables U_p and their inverses U_p^{-1} for $p\in\mathfrak{P}$.

Definition 3.2.5. If $\mathfrak{P}$ is a finite set, a **$\mathfrak{P}$–filtered chain complex** is a chain complex with a filtration of $\mathbb{Z}^{\mathfrak{P}}$, i.e. if C is a chain complex, then a $\mathfrak{P}$–filtration is a collection of subcomplexes $\mathcal{F}_I\subseteq C$ ranging over $I\in\mathbb{Z}^{\mathfrak{P}}$ such that if $I\le I'$, then $\mathcal{F}_{I'}\subseteq\mathcal{F}_I$. A $\mathfrak{P}$–filtered homomorphism is a homomorphism $\phi\colon C\to C'$ where C and C' are $\mathfrak{P}$–filtered with filtrations $\mathcal{F}_I$ and $\mathcal{F}'_I$ such that

$$\phi(\mathcal{F}_I)\subseteq\mathcal{F}'_I.$$

Definition 3.2.6. A **coloring of a multi-based link** $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ in Y is a pair $(\sigma,\mathfrak{P})$ where $\mathfrak{P}$ is a finite set and $\sigma\colon \mathbf{w}\cup\mathbf{z}\to\mathfrak{P}$ is a map which sends all of the $\mathbf{z}$ basepoints on a

given link component to the same color (this condition ensures that the differential squares to zero).

If $(\sigma, \mathfrak{P})$ is a coloring of a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, let $\mathcal{C}_{\sigma,\mathfrak{P}}$ denote the module $\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}, V_{\mathbf{z}}, V_{\mathbf{z}}^{-1}, U_{\mathfrak{P}}, U_{\mathfrak{P}}^{-1}]/I_{\sigma,\mathfrak{P}}$ where $I_{\sigma,\mathfrak{P}}$ is the submodule generated by elements of the form $U_w - U_{\sigma(w)}$ and $V_z - U_{\sigma(z)}$. The colored complex is then defined as

$$CFL^{\infty}_{UV}(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s}) = CFL^{\infty}_{UV,0}(\mathcal{H}, \sigma) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}}, U_{\mathbf{w}}^{-1}, V_{\mathbf{z}}, V_{\mathbf{z}}^{-1}]} \mathcal{C}_{\sigma,\mathfrak{P}}.$$

Given a coloring $(\sigma, \mathfrak{P})$ of $\mathbf{w} \cup \mathbf{z}$ of a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ in Y, the colored complexes $CFL^{\infty}_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ naturally obtain a $\mathfrak{P}$–filtration by powers of the variables U_p. An element of $CFL^{\infty}_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ is uniquely written as a sum of elements of the form $\mathbf{x} \cdot U_{\mathfrak{P}}^I$, and given an $I \in \mathbb{Z}^{\mathfrak{P}}$ we define $\mathcal{F}_I$ to be the $\mathbb{Z}_2[U_{\mathfrak{P}}]$–submodule generated by $\mathbf{x} \cdot U_{\mathfrak{P}}^J$ with $J \geq I$.

Remark 3.2.7. Asking that a $\mathbb{Z}_2[U_{\mathfrak{P}}$–equivariant map $F\colon CFL^{\infty}_{UV}(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s}) \to CFL^{\infty}_{UV}(\mathcal{H}', \sigma', \mathfrak{P}, \mathfrak{s}')$ be $\mathfrak{P}$–filtered is just asking that F can be written as

$$F(\mathbf{x}) = \sum_{I \geq 0} U_{\mathfrak{P}}^I \cdot H_I(\mathbf{x})$$

where the maps H_I do not involve the $U_{\mathfrak{P}}$ variables. Most maps which appear in Heegaard Floer homology are $\mathfrak{P}$–filtered. The differential, the triangle maps, the quadrilateral maps, as well as the maps Φ_w, Φ_z and $S^{\pm}_{w,z}$ are all $\mathfrak{P}$–filtered.

Given an arbitrary coloring $(\sigma, \mathfrak{P})$ of basepoints $\mathbf{w} \cup \mathbf{z}$, we may not always be able to define submodules corresponding to relative Spin^c structures $\mathfrak{t}$. However if no two basepoints from distinct link components are given the same color, then one can use a modification of Equation (3.1) to define a $\mathfrak{P}$–filtered $\mathbb{Z}_2$–submodule $CFL^{\infty}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{t})$. For our purposes, we just observe that in the case that $(\sigma, \mathfrak{P})$ is the trivial coloring (i.e. $\mathfrak{P} = \mathbf{w} \cup |L|$ and σ is the map sending $w \in \mathbf{w}$ to w and $z \in \mathbf{z}$ to the link component containing it), then $CFL^{\infty}_{UV}(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s})$ is equal to just $CFL^{\infty}_{UV}(\mathcal{H}, \mathfrak{s})$ and the maps $\pi_{\mathfrak{t}}$ and $\iota_{\mathfrak{t}}$ are still defined. The following lemma is essentially trivial, though it is useful for relating endomorphisms of $CFL^{\infty}_{UV}(Y, \mathbb{L}, \mathfrak{s})$ to endomorphisms of the subcomplexes $CFL^{\infty}(Y, \mathbb{L}, \mathfrak{t})$:

Lemma 3.2.8. *Suppose that $(L, \mathbf{w}, \mathbf{z})$ is a link in an integer homology sphere Y and that $(\sigma, \mathfrak{P})$ is the trivial coloring. Suppose f and g are $\mathfrak{P}$–filtered $\mathbb{Z}_2[U_{\mathfrak{P}}]$–module endomorphisms of $CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$, such that f and g are chain homotopic via a chain homotopy which is $\mathfrak{P}$–filtered on $CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$. Then $\pi_{\mathfrak{t}} \circ f \circ \iota_{\mathfrak{t}}$ and $\pi_{\mathfrak{t}} \circ g \circ \iota_{\mathfrak{t}}$ are $\mathbb{Z}^{|\mathbf{w}|} \oplus \mathbb{Z}^{|L|}$–filtered $\mathbb{Z}_2[\overline{U}_{\mathbf{w}}]$–chain homotopic.*

Proof. First note that the filtration on $CFL^\infty(\mathcal{H}, \mathfrak{t})$ is just the pullback under $\iota_{\mathfrak{t}}$ of the $\mathfrak{P}$–filtration on $CFL^\infty_{UV}(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s})$. If f and g are $\mathfrak{P}$–filtered chain homotopic, we have that

$$f - g = \partial H + H\partial$$

for a $\mathfrak{P}$ filtered map H. Pre– and post–composing with the $\mathfrak{P}$–filtered maps $\iota_{\mathfrak{t}}$ and $\pi_{\mathfrak{t}}$ yields a $\mathbb{Z}^{|\mathbf{w}|} \oplus \mathbb{Z}^{|L|}$–filtered chain homotopy between $\pi_{\mathfrak{t}} \circ f \circ \iota_{\mathfrak{t}}$ and $\pi_{\mathfrak{t}} \circ g \circ \iota_{\mathfrak{t}}$ because the maps $\pi_{\mathfrak{t}}$ and $\iota_{\mathfrak{t}}$ are $\mathfrak{P}$–filtered chain maps. The chain homotopy is $\mathbb{Z}_2[\overline{U}_{\mathbf{w}}]$–equivariant since the maps $\iota_{\mathfrak{t}}$ and $\pi_{\mathfrak{t}}$ are $\mathbb{Z}_2[\overline{U}_{\mathbf{w}}]$–equivariant, as we mentioned above (recall that $\overline{U}_w = U_w V_z$, where z is any base point on the link component containing z). □

3.2.5 Why we use the larger $CFL^\infty_{UV}(Y, L, \mathfrak{s})$ instead of other versions

We briefly explain why we use the object $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$ to prove formulas for basepoint moving maps, instead of other versions of link Floer homology. In the next sections, we will define maps Φ_w and Ψ_z, which are endomorphisms of $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$. However, due to the extra factors of U_w^{-1} or V_z^{-1} in the definitions, these do not preserve $CFL^\infty(Y, \mathbb{L}, \mathfrak{t})$ for a relative Spin^c structure. Instead they change the relative Spin^c structure by $\pm PD[\mu]$, where μ is the meridian of the component containing w and z (note, however that the composition $\Phi_w \Psi_z$ does actually preserve relative Spin^c structure). Although this is not insurmountable, what's worse is that the maps $S^+_{w,z}$ and $S^-_{w,z}$ are not even endomorphisms of the same complex, and since $S^+_{w,z} S^-_{w,z} \simeq \Phi_w$, we know that they can't preserve relative Spin^c structures. Similarly, one could try to use the version of link Floer homology described in [OS08] as a Spin$^c(Y, L)$–filtration on $CF^\infty(Y)$, but we have the same problem since the map Φ_w is not Spin$^c(Y, L)$–filtered.

The solution is clearly to work with the larger complexes $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$.

There are also other algebraic advantages to working with $CFL^\infty_{UV}(Y, \mathbb{L}, \mathfrak{s})$. For instance we can think of Φ_w and Ψ_z as formal derivatives of the differential. Using our expression for ∂^2, we can quickly derive many relations between various Φ_w and Ψ_z maps.

3.3 The maps Φ_w and Ψ_z

Here we now define maps Φ_w and Ψ_z, which are endomorphisms of $CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$. These are denoted by $\Phi_{i,j}$ and $\Psi_{i,j}$ in [Sar15]. We define $\Phi_w : CFL^\infty_{UV}(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s}) \to CFL^\infty_{UV}(\mathcal{H}, \sigma, \mathfrak{P}, \mathfrak{s})$ by the formula

$$\Phi_w(\mathbf{x}) = U_w^{-1} \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_w(\phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y},$$

which we can alternatively think of as $(\frac{d}{dU_w}\partial)$. Similarly we define

$$\Psi_z(\mathbf{x}) = V_z^{-1} \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_z(\phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y},$$

which we can alternatively write as $(\frac{d}{dV_z}\partial)$, where the derivative is taken on $CFL^\infty_{UV,0}$ (i.e. before tensoring with the module $\mathcal{L}$ and thus setting all of the V_z on a link component equal to each other).

We have the following (compare [Sar15, Lem. 4.1]):

Lemma 3.3.1. *On $CFL^\infty_{UV}(\mathcal{H}, \mathfrak{s})$, we have $\Phi_w \partial + \partial \Phi_w = 0$. Also $\Psi_z \partial + \partial \Psi_z = U_w + U_{w'}$, where w and w' are the $\mathbf{w}$ basepoints adjacent to z.*

Proof. One takes the derivative of $\partial \circ \partial$ with respect to either V_z or U_w, before one tensors $CFL^\infty_{UV,0}$ with $\mathcal{L}$. The map ∂^2 is computed in Lemma 3.2.1. After tensoring with $\mathcal{L}$, one immediately arrives at the equalities described above. □

In addition, we have the following (compare [Sar15, Theorem 4.2.]):

Lemma 3.3.2. *The maps Φ_w and Ψ_z commute with change of diagram maps $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}$ up to $\mathfrak{P}$–filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain homotopy.*

Proof. Consider the complex $CFL^{\infty}_{UV,0}$ (i.e. the complex before we set all of the V_z variables on each link component equal to each other). The differential doesn't square to zero, though we can still consider the maps $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}$. These can be written as a composition of change of almost complex structure maps, triangle maps (corresponding to $\boldsymbol{\alpha}$– or $\boldsymbol{\beta}$– isotopies or handleslides), (1,2)–stabilization maps, as well as maps corresponding to isotoping the Heegaard surface inside of Y via an isotopy which fixes $\mathbb{L}$. We claim that the maps $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}$ satisfy

$$\Phi_{\mathcal{H}_1\to\mathcal{H}_2}\partial + \partial\Phi_{\mathcal{H}_1\to\mathcal{H}_2} = 0,$$

even before tensoring with $\mathcal{L}$. The maps $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}$ are defined as a composition of maps which count holomorphic triangles (handleslides or isotopy maps), holomorphic disks with dynamic almost complex structure (change of almost complex structure maps) or maps which are defined via simple, explicit formulas (stabilization and diffeomorphism). The maps which associated to $(1,2)$–stabilizations and diffeomorphisms obviously satisfy $\partial\phi + \phi\partial = 0$ before tensoring with $\mathcal{L}$. The maps induced by counting disks with dynamic almost complex structure also satisfy $\partial\phi + \phi\partial = 0$ before tensoring with ∂, since that follows from a Gromov compactness argument. Maps induced by handleslides or isotopies of the $\boldsymbol{\alpha}$–curves take the form

$$\mathbf{x} \mapsto F_{\boldsymbol{\alpha}'\boldsymbol{\alpha}\boldsymbol{\beta}}(\Theta \otimes \mathbf{x})$$

where Θ is the top degree generator of a complex $CFL^{-}_{UV}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ and $F_{\boldsymbol{\alpha}'\boldsymbol{\alpha}\boldsymbol{\beta}}$ is the map which counts holomorphic triangles. For this to be a chain map before tensoring with $\mathcal{L}$, it is sufficient that $\partial\Theta = 0$ before tensoring with $\mathcal{L}$, since the triangle map

$$F_{\boldsymbol{\alpha}'\boldsymbol{\alpha}\boldsymbol{\beta}}\colon CFL^{-}_{UV,0}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z}) \otimes CFL^{\infty}_{UV,0}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}) \to CFL^{\infty}_{UV,0}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$$

is a chain map by a Gromov compactness argument. We note now that the diagram $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ represents an unlink embedded in $(S^1 \times S^2)^{\#n}$ for some n, and this unlink

has exactly two basepoints per link component. By the differential computation in Lemma 3.2.1, the complex

$$CFL^-_{UV,0}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$$

is a chain complex before tensoring with anything, and in particular the homology group $HFL^-_{UV,0}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$, is well defined even before tensoring with anything. An easy computation shows that if $HFL^-_{UV,0,\max}$ denotes the subset of maximal homological grading (here the homological grading is obtained by ignoring the $\mathbf{z}$–basepoints, and assigning U variables degree -2, and V variables degree 0), then we have an isomorphism

$$HFL^-_{UV,0,\max}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z}) \cong \mathbb{Z}_2[V_{\mathbf{z}}],$$

and in particular $HFL^-_{UV,0}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ admits a "generator" Θ which is distinguished by the property of generating the maximally graded subset as a $\mathbb{Z}_2[V_{\mathbf{z}}]$–module. In particular $\partial\Theta = 0$ even before tensoring with $\mathcal{L}$, as we needed.

Hence

$$\Phi_{\mathcal{H}_1\to\mathcal{H}_2}\partial + \partial\Phi_{\mathcal{H}_1\to\mathcal{H}_2} = 0,$$

even before tensoring with $\mathcal{L}$. Differentiating with respect to U_w yields that

$$\Phi'_{\mathcal{H}_1\to\mathcal{H}_2}\partial + \Phi_{\mathcal{H}_1\to\mathcal{H}_2}\Phi_w + \Phi_w\Phi_{\mathcal{H}_1\to\mathcal{H}_2} + \partial\Phi'_{\mathcal{H}_1\to\mathcal{H}_2} = 0,$$

immediately implying that $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}\Phi_w + \Phi_w\Phi_{\mathcal{H}_1\to\mathcal{H}_2} \simeq 0$. The only point to check is that the chain homotopy $H = \Phi'_{\mathcal{H}_1\to\mathcal{H}_2}$ is $\mathfrak{P}$–filtered and $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant. The equivariance condition is trivial. The filtration condition is also easy to check, since whenever F has a decomposition with only nonnegative powers of U_w and V_z, the map $(\frac{d}{dU_w}F)$ also has a decomposition with nonnegative powers of U_w and V_z. □

Remark 3.3.3. Using the Leibniz rule we have that

$$\Phi_w = \frac{d}{dU_w} \circ \partial + \partial \circ \frac{d}{dU_w},$$

as long as U_w doesn't share the same color as any other basepoint. Similarly $\Psi_z \simeq 0$ if z doesn't share the same color with any other basepoint, though in both cases the chain homotopy $H = \frac{d}{dU_w}$ or $H = \frac{d}{dV_z}$ is neither $\mathfrak{P}$–filtered nor $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant.

3.4 Cylindrical boundary degenerations

We consider holomorphic curves whose boundary is mapped to only the $\boldsymbol{\alpha}$–curves, or only the $\boldsymbol{\beta}$ curves. These will be called **cylindrical $\boldsymbol{\alpha}$–boundary degenerations** or **cylindrical $\boldsymbol{\beta}$–boundary degenerations**.

We now define cylindrical $\boldsymbol{\alpha}$–boundary degenerations. Suppose that S is a Riemann surface with d punctures $\{p_1, \dots, p_d\}$ on its boundary. We consider holomorphic maps

$$u\colon\ S \to \Sigma \times (-\infty, 1] \times \mathbb{R}$$

such that the following hold:

1. u is smooth;
2. $u(\partial S) \subseteq (\boldsymbol{\alpha} \times \{1\} \times \mathbb{R})$;
3. $\pi_{\mathbb{D}} \circ u$ is non-constant on each component of S;
4. for each i, $u^{-1}(\alpha_i \times \{1\} \times \mathbb{R})$ consists of exactly one component of $\partial S \setminus \{p_1, \dots, p_d\}$;
5. the energy of u is finite;
6. u is an embedding;
7. if $z_i \in S$ is a sequence of points approaching a puncture p_j, then $(\pi_{\mathbb{D}} \circ u)(z_i)$ approaches -1 in the compactification of $(-\infty, 1] \times \mathbb{R}$ as the unit complex disk (with the point at ∞ identified with -1).

We organize such curves into moduli spaces $\mathcal{N}(\phi)$ for $\phi \in \pi_2^{\alpha}(\mathbf{x})$, modding out by automorphisms of the source, as usual. There is an action of $\mathrm{PSL}_2(\mathbb{R})$ on $\mathcal{N}(\phi)$, which is just the action on the $(-\infty, 1] \times \mathbb{R}$ coordinate of a disk u, and we denote the quotient space $\widehat{\mathcal{N}}(\phi)$. One defines cylindrical $\boldsymbol{\beta}$–boundary degenerations analogously. We now discuss transversality. In the original set up (singly pointed diagrams and disks mapped into $\mathrm{Sym}^g(\Sigma)$), a generic almost complex structure J on $\mathrm{Sym}^g(\Sigma)$ in a neighborhood of $\mathrm{Sym}^g(\mathbf{j})$ achieves

transversality for Maslov index 2 holomorphic $\boldsymbol{\alpha}$–degenerate disks ([OS04e, Prop. 3.14]). In the cylindrical setup, if a sequence of holomorphic strips for the almost complex structures considered in [?LCylindrical] has a cylindrical boundary degeneration in its limit, the boundary degeneration will be $\mathfrak{j}_\Sigma \times \mathfrak{j}_{\mathbb{D}}$–holomorphic. Thus we need transversality for cylindrical boundary degenerations for split almost complex structures. For the standard proof that $\partial^2 = 0$ and for purposes of this paper, we only need transversality for Maslov index 2 boundary degenerations. Each of these domains has multiplicity 1 in one component of $\Sigma \setminus \boldsymbol{\alpha}$, and zero everywhere else. If $u\colon S \to \Sigma \times (-\infty, 1] \times \mathbb{R}$ is a component of a holomorphic curve representing an element of $\mathcal{N}(\phi)$ for a $\phi \in \pi_2^{\alpha}(\mathbf{x})$ such that $\pi_\Sigma \circ u$ is nonconstant, then by easy complex analysis $u|_C$ is injective, where $C \subseteq \partial S$ is the part of S mapping to $\partial\mathcal{D}(\phi)$. Adapting the strategy of perturbing boundary conditions instead of almost complex structures, as in [?LCylindrical, Prop. 3.9], [OS04e, Prop. 3.9] or [Oh96], for generic choice of $\boldsymbol{\alpha}$–curves, we can thus achieve transversality for Maslov index 2 cylindrical boundary degenerations.

An important result for our purposes is a count of Maslov index 2 boundary degenerations produced by Ozsváth and Szabó:

Theorem 3.4.1 ([OS08, Theorem 5.5])**.** *Consider a surface Σ of genus g equipped with a set of attaching circles $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_{g+\ell-1}\}$ which span a g–dimensional lattice in $H_1(\Sigma; \mathbb{Z})$. If $\mathcal{D}(\phi) \geq 0$ and $\mu(\phi) = 2$, then $\mathcal{D}(\phi) = A_i$ for some i, and indeed*

$$\#\widehat{\mathcal{N}}(\phi) = \begin{cases} 0 \pmod 2 & \text{if } \ell = 1 \\ 1 \pmod 2 & \text{if } \ell > 1 \end{cases}.$$

Here A_i denotes a component of $\Sigma \setminus \boldsymbol{\alpha}$.

3.5 Preliminaries on the quasi-stabilization operation

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is an oriented link in Y and that w and z are two points, both in a single component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ such that $(L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$ has basepoints which alternate

between $\mathbf{w}$ and $\mathbf{z}$ as one traverses the link. We assume that according to the orientation of L the point w comes after z. In Section 3.7 we prove invariance for quasi-stabilization maps

$$S^+_{w,z}\colon CFL^\infty_{UV}(Y,L,\mathbf{w},\mathbf{z},\sigma,\mathfrak{P},\mathfrak{s}) \to CFL^\infty_{UV}(Y,L,\mathbf{w}\cup\{w\},\mathbf{z}\cup\{z\},\sigma',\mathfrak{P},\mathfrak{s})$$

and

$$S^-_{w,z}\colon CFL^\infty_{UV}(Y,L,\mathbf{w}\cup\{w\},\mathbf{z}\cup\{z\},\sigma',\mathfrak{P},\mathfrak{s}) \to CFL^\infty_{UV}(Y,L,\mathbf{w},\mathbf{z},\sigma,\mathfrak{P},\mathfrak{s}),$$

which are defined up to $\mathfrak{P}$–filtered $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain homotopy. Here σ' is a coloring which extends σ.

Though it will take several sections to construct the maps and prove they are well defined, we now summarize that the maps will be defined by the formulas

$$S^+_{w,z}(\mathbf{x}) = \mathbf{x}\times\theta^+$$

and

$$S^-_{w,z}(\mathbf{x}\times\theta^+) = 0, \qquad \text{and} \qquad S^-_{w,z}(\mathbf{x}\times\theta^-) = \mathbf{x},$$

for suitable choices of Heegaard diagrams and almost complex structures.

To define the quasi-stabilization map, we use the special connected sum operation from [MO10]. In [MO10], Manolescu and Ozsváth describe a way of adding new w and z basepoints to Heegaard multi-diagrams. They prove that for multi-diagrams with at least three sets of attaching curves (e.g. Heegaard triples or quadruples), there is an identification of certain moduli spaces of holomorphic curves on the unstabilized diagram and certain moduli spaces of holomorphic curves on the stabilized diagram. They conjecture an analogous result for the holomorphic curves on a Heegaard diagram with two sets of attaching curves (i.e. for the differentials of quasi-stabilized diagrams), but only prove the result for grid diagrams using somewhat ad-hoc techniques, since in general they run into transversality issues. We will soon prove Proposition 3.5.3 below, computing the differential on quasi-stabilized diagrams for appropriate almost complex structures, showing how to avoid any transversality issues and using no more gluing technology than is used in showing that $\partial^2 = 0$ on multi-pointed diagrams.

3.5.1 Topological preliminaries on quasi-stabilization

Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, L, \mathbf{w}, \mathbf{z})$. Given new basepoints w, z in the same component of $L \setminus (\mathbf{w} \cup \mathbf{z})$, such that w occurs after z, we now describe a new diagram $\overline{\mathcal{H}}_{p,\alpha_s}$, which depends on a choice of point $p \in \Sigma$ and curve $\alpha_s \subseteq \Sigma \setminus \boldsymbol{\alpha}$ which passes through the point p. For fixed α_s and p, the diagram $\overline{\mathcal{H}}_{p,\alpha_s}$ will be defined up to an isotopy of Y relative $\mathbf{w} \cup \mathbf{z} \cup \{w, z\}$ which maps L to L.

Given a diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ as above, let A denote the component of $\Sigma \setminus \boldsymbol{\alpha}$ which contains the basepoints adjacent to w and z on L. Let $p \in A \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$ be a point. If U_α denotes the handlebody component of $Y \setminus \Sigma$ such that the $\boldsymbol{\alpha}$–curves bound compressing disks in U_α, then there is a path λ in U_α from p to a point on L between w and z. Such a curve λ is specified up to an isotopy fixing $\mathbf{w} \cup \mathbf{z} \cup \{w, z\}$ which maps L to L by requiring that λ be isotopic in $\overline{U}_\alpha$ to a segment of L concatenated with an embedded arc on $\Sigma \setminus \boldsymbol{\alpha}$.

Let $N(\lambda)$ denote a regular neighborhood of λ inside of U_α such that $\partial N(\lambda)$, the boundary of $N(\lambda)$ inside of U_α, satisfies

$$\partial N(\lambda) \cap L = \{w, z\}.$$

Topologically $\partial N(\lambda)$ is just a disk, which we denote by D_1. Also let D_2 denote $N(\lambda) \cap \Sigma$, which we note is also a disk. We can assume that $D_2 \cap (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z}) = \varnothing$. Define

$$\overline{\Sigma}_p = (\Sigma \setminus D_2) \cup D_1.$$

The surface $\overline{\Sigma}_p$ is specified up to an isotopy which fixes $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$. Figure 3.1 shows the situation schematically.

We wish to extend the arc $\alpha_s \setminus D_2$ over all of D_1 to get a curve $\overline{\alpha}_s$ on $\overline{\Sigma}_p$. As is demonstrated in Figure 3.2, there is not an isotopically unique way to do this relative the new basepoints w and z.

The set of such curves is easily seen to be those generated by the images of the curve on the left in Figure 3.2 under finger moves of w around z. Fortunately the arc $\alpha_s \setminus D_2$ can be extended over D_1 uniquely (up to isotopy) by requiring the resulting curve $\overline{\alpha}_s \subseteq \overline{\Sigma}_p$ to bound

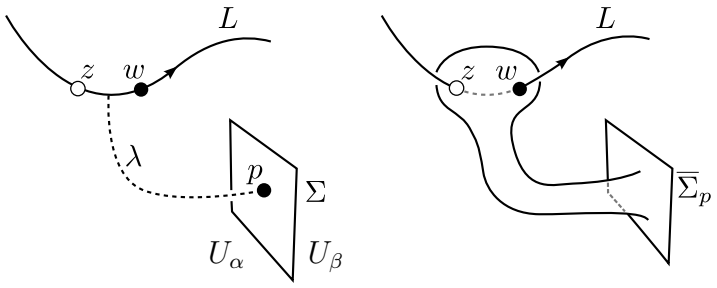

Figure 3.1: **The path λ and the surfaces Σ and $\overline{\Sigma}_p$.**

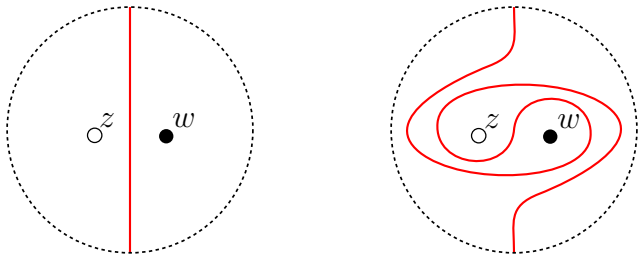

Figure 3.2: **Different choices of $\overline{\alpha}_s$ curve on D_1 interpolating $\alpha_s \setminus D_2$.** There is a unique isotopy class of such curves such that the resulting $\overline{\alpha}_s$ curve on $\overline{\Sigma}_p$ bounds a compressing disk which doesn't intersect L.

a compressing disk in $U_{\overline{\alpha}}$ which doesn't intersect L, where here $U_{\overline{\alpha}}$ denotes the component of $Y \setminus \overline{\Sigma}_p$ in which the $\boldsymbol{\alpha}$–curves bound compressing disks.

Suppose $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram, and $p \in \Sigma \setminus \boldsymbol{\alpha}$ is a chosen point, and let $N(p_0)$ denote a neighborhood of the connected sum point p_0 on S^2, which intersects α_0 in an arc and doesn't intersect β_0. Given a diffeomorphism $\psi\colon \Sigma \to \overline{\Sigma}_p$ which is the identity outside of D_2, as well as an embedding $\iota\colon S^2 \setminus N(p_0) \to D_1$ which sends $\alpha_0 \setminus N(p_0)$ to $\overline{\alpha}_s$, we can form a diagram

$$\overline{\mathcal{H}}_{p,\alpha_s} = (\overline{\Sigma}_p, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}),$$

where $\overline{\boldsymbol{\alpha}} = \boldsymbol{\alpha} \cup \{\overline{\alpha}_s\}$ and $\overline{\boldsymbol{\beta}} = \boldsymbol{\beta} \cup \{\iota(\beta_0)\}$. Such a choice of ψ and ι will be part of a larger collection of data $\mathcal{J}$ which we will consider in the next section and will call "gluing data". If we need to emphasize the distinction, we will write $\overline{\boldsymbol{\alpha}}$ for $\boldsymbol{\alpha} \cup \{\overline{\alpha}_s\}$, the curves on $\overline{\Sigma}_p$, and we will write $\boldsymbol{\alpha}^+$ for $\boldsymbol{\alpha} \cup \{\alpha_s\}$, the curves on Σ. By abuse of notation, we will often write α_s to denote both $\alpha_s \subseteq \Sigma$ and $\overline{\alpha}_s \subseteq \overline{\Sigma}_p$. Similarly, when no confusion will arise, we will write

β_0 for both the curve β_0 on S^2 as well as the curve $\iota(\beta_0)$ on $\overline{\Sigma}_p$.

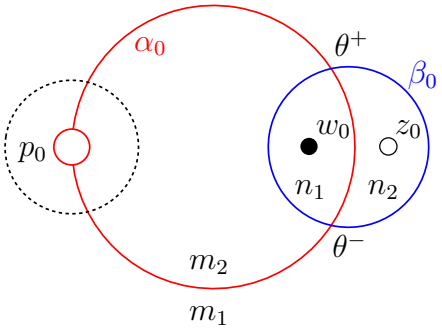

Figure 3.3: **The diagram $\mathcal{H}_0$ used for quasi-stabilization,** with multiplicities labeled. The dashed circle denotes where we will perform the neck stretching in the special connected sum.

3.5.2 Gluing data and almost complex structures

In Chapter 2 the author describes a systematic way of constructing and proving invariance of maps corresponding to adding or removing a basepoint from a closed 3–manifold. A key ingredient was a choice of auxiliary data which we call "gluing data" for piecing patching two almost complex structures together in a systematic way. Here we introduce the analogous idea for quasi-stabilization.

Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ and w, and z are two new consecutive basepoints on L with z following w. Suppose $p \in A \setminus (\mathbf{w} \cup \mathbf{z} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta})$ is a distinguished point, where here A denotes the component of $\Sigma \setminus \boldsymbol{\alpha}$ containing the basepoints on L adjacent to w and z. Let $\overline{\Sigma}_p$ denote the Heegaard surface described in the previous subsection. We have the following definition:

Definition 3.5.1. We define **gluing data** to be a collection $\mathcal{J} = (\psi, J_s, J_{s,0}, B, B_0, r, r_0, p_0, \iota, \phi)$ where

1. $\psi\colon \Sigma \to \overline{\Sigma}_p$ is a diffeomorphism which is fixed outside of D_2 and maps α_s to $\overline{\alpha}_s$;
2. $B \subseteq \Sigma$ is a closed ball containing p which doesn't intersect $(\mathbf{w} \cup \mathbf{z} \cup \boldsymbol{\alpha} \cup \boldsymbol{\beta})$ and such

that $\alpha_s \cap B$ is a closed arc;

3. the point $p_0 \in \alpha_0 \setminus \beta_0$ is the connected sum point;

4. $B_0 \subseteq S^2$ is a closed ball containing p_0 which doesn't intersect β_0 and such that $B_0 \cap \alpha_0$ is a closed arc;

5. J_s is an almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$ which is split on B;

6. $J_{s,0}$ is an almost complex structure on $S^2 \times [0,1] \times \mathbb{R}$ which is split on B_0;

7. r and r_0 are real numbers such that $0 < r, r_0 < 1$;

8. using the unique (up to rotation) conformal identifications of (B, p) and (B_0, p_0) as $(\mathbb{D}, 0)$, where $\mathbb{D}$ denotes the unit complex disk, ι is an embedding of $\overline{S^2 \setminus r_0 \cdot B_0}$ into $r \cdot B \subseteq \Sigma$ such that

$$\iota(\alpha_0) \subseteq \alpha_s,$$

$$(\psi \circ \iota)(z_0) = z, \qquad (\psi \circ \iota)(w_0) = w,$$

and

$$(r \cdot B) \setminus \iota(S^2 \setminus r_0 \cdot B_0)$$

is a closed annulus;

9. letting $\widetilde{A}, A$ and A_0 denote the closures of the annuli $B \setminus \iota(S^2 \setminus B_0)$, $B \setminus r \cdot B$ and $B_0 \setminus r_0 \cdot B_0$ respectively, then

$$\phi\colon\ \widetilde{A} \to S^1 \times [-a, 1+b]$$

is a diffeomorphism which sends the annulus A to $[-a, 0]$ and $\iota(A_0)$ to $[1, 1+b]$ and is conformal on A and A_0.

The space of embeddings ι is connected since if a denotes the arc on the left side of Figure 3.2, the space of diffeomorphisms $f\colon\ B \to B$ mapping $\partial B \cup a$ to itself and fixing $\{z, w\}$ is connected. That the space of diffeomorphisms $\psi\colon\ \Sigma \to \overline{\Sigma}_p$ in the definition is also connected follows for similar reasons.

Gluing data $\mathcal{J}$ and a choice of neck length T determines an almost complex structure $\mathcal{J}(T)$ on $\overline{\Sigma}_p \times [0,1] \times \mathbb{R}$.

3.5.3 Computing the quasi-stabilized differential

We now wish to compute the differential after performing quasi-stabilization. We have the following Maslov index computation:

Lemma 3.5.2. *If ϕ_0 is a homology class of disks on $\mathcal{H}_0$, then using the multiplicities in Figure 3.3, we have*

$$\mu(\phi_0) = (n_1 + n_2 + m_1 + m_2)(\phi_0).$$

Proof. The formula is easily verified for the constant disk and respects splicing in any of the Maslov index 1 strips. □

Let $\mathcal{J}$ denote gluing data as in the previous section, and let $\mathcal{J}(T)$ denote the almost complex structure on $\overline{\mathcal{H}}_{p,\alpha_s}$ determined by $\mathcal{J}$ for a choice of neck length T. We have the following:

Proposition 3.5.3. *Suppose that $\mathcal{H}$ is a strongly $\mathfrak{s}$–admissible diagram and that $\mathcal{J}$ is gluing data with almost complex structure J_s on $\Sigma \times [0,1] \times \mathbb{R}$. Then $\overline{\mathcal{H}}_{p,\alpha_s}$ is also strongly $\mathfrak{s}$–admissible and for sufficiently large T there is an identification of uncolored differentials (i.e. before we tensor with $\mathcal{L}$)*

$$\partial_{\overline{\mathcal{H}}_{p,\alpha_s},\mathcal{J}(T)} = \begin{pmatrix} \partial_{\mathcal{H},J_s} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H},J_s} \end{pmatrix},$$

where the basepoints w and z are placed between w' and z' on $\mathbb{L}$.

Proof. Suppose that u_i is a sequence of Maslov index 1 holomorphic curves on $\overline{\mathcal{H}}_{p,\alpha_s}$ in a fixed homology class for the almost complex structure $\mathcal{J}(T_i)$ for a sequence of T_i with $T_i \to \infty$. From the sequence u_i we can extract a weak limit of curves on the diagrams $(\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ and $\mathcal{H}_0$. Let U_Σ, and U_0 denote these collections of curves. The curves in U_Σ consist of

flowlines on $(\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ as well as $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–boundary degenerations on $(\Sigma, \boldsymbol{\alpha}^+)$ and $(\Sigma, \boldsymbol{\beta})$, and closed surfaces mapped into Σ. The holomorphic curves are now allowed to have a puncture along the α–boundary which is mapped asymptotically to p.

We first note that any flowline in the limit (i.e. a map $u\colon S \to \Sigma \times [0,1] \times \mathbb{R}$ which maps ∂S to $(\boldsymbol{\beta} \times \{0\} \cup \boldsymbol{\alpha}^+ \times \{1\} \times \mathbb{R}$ such that each component of S has both $\boldsymbol{\alpha}^+$ and $\boldsymbol{\beta}$ components) on the diagram $\mathcal{H}^+ = (\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ must actually be a legitimate flow line on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$. This is because if u is any holomorphic curve which is part of a weak limit of the curves u_{T_i}, then u cannot have a puncture asymptotic to an intersection point $\alpha_s \cap \beta_j$ for $\beta_j \in \boldsymbol{\beta}$. Hence if u is part of the weak limit of u_i, and if S denotes the source of u, then if ∂S has any points mapped to α_s, then S must have boundary component with a single puncture which is mapped to α_s. Projecting to $[0,1] \times \mathbb{R}$, we note that either $u|_{\partial S}$ obtains a local extrema, or u is asymptotic to both $+\infty$ and $-\infty$ as one approaches the puncture. If $u|_{\partial S}$ obtains a local extrema, then one can use the doubling trick to create an analytic function mapping $\mathbb{D}$ into $\mathbb{D}$ (where here $\mathbb{D} = \{z : |z| < 1\}$) which maps $\mathbb{D} \cap \{\mathrm{im}(z) \geq 0\}$ to $\mathbb{D} \cap \{\mathrm{im}(z) \geq 0\}$ but which satisfies $f'(z) = 0$ for some $z \in \mathbb{R}$, which is impossible by writing down a local model. The case that u is asymptotic to both $+\infty$ and $-\infty$ at the puncture is impossible since u must extend to a continuous function over the punctures. Hence any such u must be constant in the $[0,1] \times \mathbb{R}$ component, which implies that u cannot have any portion of ∂S mapped onto a $\boldsymbol{\beta}$–curve. Hence the curves in the weak limit can be taken to be holomorphic disks on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$, $\boldsymbol{\alpha}^+$– or $\boldsymbol{\beta}$–degenerations, or closed surfaces.

Though not essential for our argument, to avoid "annoying" curves (i.e. maps into $\Sigma \times [0,1] \times \mathbb{R}$ which are constant in the $[0,1] \times \mathbb{R}$–component) among the $\boldsymbol{\beta}$– or $\boldsymbol{\alpha}^+$–degenerations, we observe that by rescaling the $[0,1] \times \mathbb{R}$ component, we could instead get curves that map into $\Sigma \times (-\infty, 1] \times \mathbb{R}$ or $\Sigma \times S^2$ such that the $(-\infty, 1] \times \mathbb{R}$ or S^2 components are nonconstant. Maps into $\Sigma \times (-\infty, 1] \times \mathbb{R}$ are cylindrical $\boldsymbol{\alpha}^+$–boundary degenerations.

We now wish to compute exactly which of the above degenerations can occur in a weak limit of the sequence u_i of Maslov index one $\mathcal{J}(T_i)$–holomorphic curves. Assume without loss of generality that all of the u_i are in the same homology class ϕ.

Suppose that U_Σ consists of a collection U'_Σ of curves on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ (flowlines, boundary degenerations, closed surfaces) and a collection of curves $\mathcal{A}$ in $(\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta})$ which have a boundary component which maps to α_s. As we've already remarked, the collection $\mathcal{A}$ consists exactly of cylindrical $\boldsymbol{\alpha}^+$–boundary degenerations. Letting ϕ'_Σ denote the underlying homology class of U'_Σ, we define a combinatorial Maslov index for U_Σ by

$$\mu(U_\Sigma) = \mu(\phi'_\Sigma) + m_1(\mathcal{A}) + m_2(\mathcal{A}) + 2 \sum_{\substack{\mathcal{D} \in C(\Sigma \setminus \boldsymbol{\alpha}), \\ \alpha_s \cap \mathcal{D} = \varnothing}} n_{\mathcal{D}}(\mathcal{A}),$$

where $C(\Sigma \setminus \boldsymbol{\alpha})$ denotes the connected components of $\Sigma \setminus \boldsymbol{\alpha}$. By Lemma 3.5.2 the Maslov index of U_0 satisfies

$$\mu(U_0) = m_1(\phi) + m_2(\phi) + n_1(\phi) + n_2(\phi).$$

The formula for $\mu(U_\Sigma)$ does not necessarily count the expected dimension of anything since we've only defined it combinatorially, though we can compute the Maslov index $\mu(\phi)$ using these formulas, as follows:

$$\mu(\phi) = \mu(U_\Sigma) + \mu(U_0) - m_1(\phi) - m_2(\phi)$$

$$= \mu(\phi'_\Sigma) + n_1(\phi) + n_2(\phi) + m_1(\mathcal{A}) + m_2(\mathcal{A}) + 2 \sum_{\substack{\mathcal{D} \in C(\Sigma \setminus \boldsymbol{\alpha}) \\ \alpha_s \cap \mathcal{D} = \varnothing}} n_{\mathcal{D}}(\mathcal{A}). \tag{3.2}$$

In the above sum, each term is nonnegative, and by assumption the total sum is equal to one. If $\mu(\phi) = 1$, we immediately have that $n_{\mathcal{D}}(\mathcal{A}) = 0$ for $\mathcal{D} \in C(\Sigma \setminus \boldsymbol{\alpha})$, with $\alpha_s \cap \mathcal{D} = \varnothing$. Now $\mu(\phi'_\Sigma)$ does actually count the expected dimension of the moduli space of ϕ'_Σ, and in particular if ϕ'_Σ has a representative as a broken curve, we must have $\mu(\phi'_\Sigma) \geq 0$ with equality iff ϕ'_Σ is the constant disk.

As a consequence we see that if $\mu(\phi) = 1$, we have that exactly one of $\mu(\phi'_\Sigma), n_1(\phi), n_2(\phi), m_1(\mathcal{A})$, or $m_2(\mathcal{A})$ is equal to 1, and the rest are zero. The cases of $n_1(\phi) = 1$ or $n_2(\phi) = 1$ respectively are easy to analyze, and those possibilities contribute summands of

$$\begin{pmatrix} 0 & U_w \\ 0 & 0 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} 0 & 0 \\ V_z & 0 \end{pmatrix},$$

respectively, to $\partial_{\overline{\mathcal{H}}_{p,\alpha_s}}$.

We now consider broken curves in the limit with $\mu(\phi'_\Sigma) = 1$ and the remaining terms zero. In this case case we have that $m_1(\mathcal{A}) = m_2(\mathcal{A}) = 0$ (so $\mathcal{A} = 0$) and $n_1(\phi) = n_2(\phi) = 0$. In this case, we observe that ϕ'_Σ is represented by U_Σ and $\mu(\phi'_\Sigma) = 1$, so the limit cannot contain any boundary degenerations or closed surfaces. Furthermore, by Maslov index considerations we have that U_Σ consists of a single Maslov index one flowline, which we denote by u_Σ.

We now consider the curves in U_0. There must be a component of U_0 which satisfies a matching condition with u_Σ. Note also that since U_Σ consists only of a single disk on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, there cannot be any curves in U_0 which have a point on the boundary mapped to p_0 or be asymptotic to p_0.

Let u_0 denote the component of U_0 which satisfies the matching condition

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

In particular, this forces $m_1(u_0) = m_2(u_0)$. We also have $n_1(u_0) = n_2(u_0) = 0$. Here, if $u\colon S \to \Sigma \times [0,1] \times \mathbb{R}$ is a holomorphic disk, $\rho^q(u)$ is the divisor

$$(\pi_{\mathbb{D}} \circ u)(\pi_\Sigma \circ u)^{-1}(q) \in \mathrm{Sym}^{n_q(u)}(\mathbb{D}).$$

Any additional components u'_0 of of U_0 must also satisfy $n_1(u'_0) = n_2(u'_0) = 0$ and also can't have an interior point or boundary point mapped to p_0, and hence must be constant. Hence U_0 consists exactly of a holomorphic strip u_0 with $m_2(u_0) = m_1(u_0)$ which satisfies a matching condition with u_Σ, i.e. (u_Σ, u_0) are prematched strips.

We thus have shown that if $\mu(\phi'_\Sigma) = 1$, then the weak limit of the curves u_i is a prematched strip, so following standard gluing arguments (cf. [`?LCylindrical`, App. A]), the count of $\widehat{\mathcal{M}}_{\mathcal{J}(T)}(\phi)$ is equal to the count of prematched strips with total homology class ϕ, for sufficiently large T. If ϕ_0 is a homology class of disks in $\pi_2(\theta^+, \theta^+)$ or $\pi_2(\theta^-, \theta^-)$, let $\mathcal{M}(\phi_0, \mathbf{d})$ denote the set of holomorphic strips u representing ϕ_0 with $\rho^{p_0}(u) = \mathbf{d}$. Note that there is a unique homology class of disks $\phi_0 \in \pi_2(\theta^+, \theta^+)$ with $m_1(\phi_0) = m_2(\phi_0) = |\mathbf{d}|$ and $n_1(\phi_0) = n_2(\phi_0) = 0$.

We claim that

$$\mathcal{M}(\phi_0, \mathbf{d}) \equiv 1 \pmod 2$$

if $m_1(\phi_0) = m_2(\phi_0) = |\mathbf{d}|$ and $n_1(\phi_0) = n_2(\phi_0) = 0$. We consider a path $\mathbf{d}_t$ between two divisors $\mathbf{d}_0$ and $\mathbf{d}_1$ and consider the 1–dimensional space

$$\mathcal{M} = \coprod_{t \in [0,1]} \mathcal{M}(\phi_0, \mathbf{d}_t).$$

We count the ends of $\mathcal{M}$. There are ends corresponding to $\mathcal{M}(\phi_0, \mathbf{d}_0)$ and $\mathcal{M}(\phi_0, \mathbf{d}_1)$. On the other hand, there are ends corresponding to strip breaking or other types of degenerations. Any curve in the degeneration cannot have p_0 in it's boundary, which constrains any degeneration to be into disks of the form $\pi_2(\theta^+, \theta^+)$ or $\pi_2(\theta^-, \theta^-)$. But if any nontrivial strip breaking occurs, the Maslov index of the matching component drops, contradicting the formula for the Maslov index. Hence the only ends of $\mathcal{M}$ correspond to $\mathcal{M}(\phi_0, \mathbf{d}_0)$ and $\mathcal{M}(\phi_0, \mathbf{d}_1)$, implying that

$$\#\mathcal{M}(\phi_0, \mathbf{d}_0) \equiv \#\mathcal{M}(\phi_0, \mathbf{d}_1) \pmod 2.$$

We now consider a path of divisors $\mathbf{d}_T$ consisting of k points in $[0,1] \times \mathbb{R}$ spaced at least T apart which approach the line $\{0\} \times \mathbb{R}$ as $T \to \infty$. Letting $T \to \infty$, since p_0 is not on β_0, we know that the Gromov limit of the curves in $\mathcal{M}(\phi_0, \mathbf{d}_T)$ consists of k cylindrical β_0–degenerations, and a single constant holomorphic strip. Applying Theorem 3.4.1, we get that the total count of the boundary of

$$\overline{\sqcup_T \mathcal{M}(\phi_0, \mathbf{d}_T)}$$

is $\#\mathcal{M}(\phi_0, \mathbf{d}_1) + 1$, implying the claim.

Disks which consist of preglued flowline (u_Σ, u_0) glued together thus provide a total contribution of

$$\begin{pmatrix} \partial_{\mathcal{H}, J_s} & 0 \\ 0 & \partial_{\mathcal{H}, J_s} \end{pmatrix},$$

to the differential.

We now consider the last contribution to the differential. These correspond to having $\mathcal{D}(\phi_1') = 0$ and $n_1 = n_2 = 0$, and $m_1(\mathcal{A}) + m_2(\mathcal{A}) = 1$. In this case $\mathcal{A}$ is an $\boldsymbol{\alpha}^+$–boundary degeneration on $(\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta})$. Since

$$\sum_{\substack{\mathcal{D} \in C(\Sigma \setminus \boldsymbol{\alpha}) \\ \alpha_s \cap \mathcal{D} = \varnothing}} n_{\mathcal{D}}(\mathcal{A}) = 0,$$

we have that $\mathcal{D}(\mathcal{A})$ is constrained to one of two domains. For each of the two possible domains of $\mathcal{D}(\mathcal{A})$, there is exactly one corresponding choice of domain for ϕ_0, the homology class of U_0, which is just the domain with exactly one of m_1 and m_2 equal to one, and the other equal to zero, as well as n_1 and n_2 also zero.

On $\overline{\mathcal{H}}_{p,\alpha_s}$ these correspond to exactly two homology classes of disks. We now describe two strategies to count such disks. The first would be to perform a gluing argument to glue holomorphic representatives of the bigon on $\mathcal{H}_0$ to Maslov index 2 $\boldsymbol{\alpha}$–boundary degeneration on $(\Sigma, \boldsymbol{\alpha}^+)$ at punctures along their boundaries. As we remarked, once could rescale the curves so that they were genuine cylindrical $\boldsymbol{\alpha}^+$, and by perturbing the $\boldsymbol{\alpha}^+$ curves we could achieve transversality since the domains of such curves are α–injective. By a gluing argument, one could prove that the count on $\overline{\mathcal{H}}_{p,\alpha_s}$ was equal to the product of the counts for the two pieces, for sufficiently stretched neck. Although the author isn't aware of any obstruction to do this, we will describe another approach which uses more established gluing results and a nice trick.

Let $\mathbf{x}$ be an intersection point on the unstabilized diagram. By our work up to now, there are two homology classes we have left to count: a disk $\phi_{z'} \in \pi_2(\mathbf{x} \times \theta^+, \mathbf{x} \times \theta^-)$ which goes over z' once, and a disk $\phi_{w'} \in \pi_2(\mathbf{x} \times \theta^-, \mathbf{x} \times \theta^+)$, which goes over w' once. To count the number of representatives of $\phi_{w'}$ and $\phi_{z'}$, we consider the ends of the moduli spaces associated to certain Maslov index 2 homology classes in $\pi_2(\mathbf{x} \times \theta^+, \mathbf{x} \times \theta^+)$. On $\overline{\mathcal{H}}_{p,\alpha_s}$, we consider the two components of $\overline{\Sigma}_p \setminus \overline{\boldsymbol{\alpha}}$ which have boundary along α_s. For an intersection point $\mathbf{x}$ on the unstabilized diagram, each of these two domains yields a homotopy class $\pi_2(\mathbf{x} \times \theta^+, \mathbf{x} \times \theta^+)$. Let us call these homotopy classes $A_{w'}$ and $A_{z'}$, depending on whether they go over w' or z'. Let us consider the ends of the 1–dimensional space of holomorphic

disks $\widehat{\mathcal{M}}(A_{w'})$. The ends correspond to boundary degenerations and strip breaking. By our work so far, for sufficiently stretched almost complex structure, there is a single domain which can appear as the domain of a Maslov index 1 homology class with $n_1 \neq 0$ and which admits a holomorphic representative, namely the the bigon going over w once. Let us call this bigon b_w. Hence if a 1-parameter family of holomorphic disks in $\widehat{\mathcal{M}}(A_{z'})$ breaks into a pair of holomorphic disks, one of them must be have domain equal to b_w. This forces the other have domain $A_{z'} - b_w$, i.e. the homology class must be $\phi_{z'}$. As no other homotopy classes can appear in strip breaking, due to our degenerated almost complex structure, we conclude that

$$\#\widehat{\mathcal{N}}(A_{z'}) + \#\widehat{\mathcal{M}}(\phi_{z'})\#\widehat{\mathcal{M}}(b_w) = 0,$$

as the latter count is the number of ends of $\widehat{\mathcal{M}}(A_{z'})$. Since $\#\widehat{\mathcal{N}}(A_{z'}) = 1$ by Theorem 3.4.1, and the bigon certainly has a unique holomorphic representative, we conclude that $\widehat{\mathcal{M}}(\phi_{z'}) = 1 \in \mathbb{Z}_2$. By an analogous argument, we conclude that $\widehat{\mathcal{M}}(\phi_{w'}) = 1$, as well.

With the above count we see that such curves make contributions of

$$\begin{pmatrix} 0 & U_{w'} \\ 0 & 0 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} 0 & 0 \\ V_{z'} & 0 \end{pmatrix}.$$

Summing together all of the contributions, we see that the differential takes the form

$$\partial_{\overline{\mathcal{H}}_{p,\alpha_s},\mathcal{J}(T)} = \begin{pmatrix} \partial_{\mathcal{H},J_s} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H},J_s} \end{pmatrix}.$$

□

Note that we tensor $CFL^{\infty}_{UV,0}(\mathcal{H},\mathfrak{s})$ with $\mathcal{L}$ so that the differential squares to zero. When we do this, we set $V_z = V_{z'}$, and the bottom left entry of the differential vanishes.

Example 3.5.4. We now briefly give an example which helps to illustrate the technique we used to count some of the disks appearing in the off diagonal entries of the differential. We consider a nested quasi-stabilization, shown in Figure 3.4, where we stretch along the dashed curve on the inside quasi-stabilization. We haven't drawn any basepoints in the figure. We've

illustrated a homology class $A \in \pi_2(\mathbf{x}\times\theta, \mathbf{x}\times\theta)$ whose domain is just a component of $\Sigma \setminus (\overline{\alpha})$ and the ends of $\widehat{\mathcal{M}}(A)$. We've illustrated how the homology class can split into either pairs of Maslov index 1 disks, or a boundary degeneration. When we stretch the neck sufficiently, the weak limits argument from the previous proposition prohibits the middle pair of homology disks from both having a representative. Hence the boundary of $\widehat{\mathcal{M}}(\phi)$, consists of exactly $\widehat{\mathcal{N}}(A)$ (which has a unique representative), as well as $\widehat{\mathcal{M}}(b) \times \widehat{\mathcal{M}}(\phi)$, where b is a bigon, and $\phi \in \pi_2(\mathbf{x}\times\theta, \mathbf{x}\times\theta')$ is one of the disks we were trying to count at the end of the previous proposition.

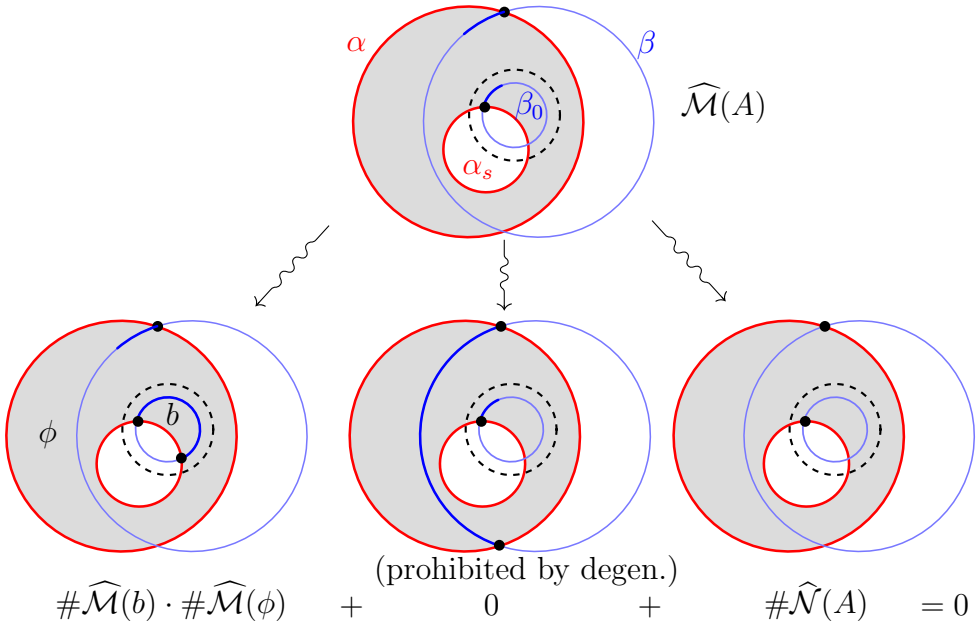

Figure 3.4: **An example of the possible strip breaking which can occur for the homology class A.** By degenerating the almost complex structure, our weak limits argument in the previous proposition rules out the middle degeneration from occurring. This allows us to count the representatives for the Maslov index 1 homology class ϕ, appearing on the left, which is otherwise hard to count.

3.5.4 Dependence of quasi-stabilization on gluing data

In this subsection we prove some initial results about quasi-stabilization and change of almost complex structure maps. The reader should compare this to Section 2.5, where the analogous arguments are presented for free stabilization.

Lemma 3.5.5. *Suppose that $\mathcal{J}$ is gluing data. Then there is an N such that if $T, T' > N$ and $\mathcal{J}(T)$, and $\mathcal{J}(T')$ achieve transversality, then*

$$\Phi_{\mathcal{J}(T)\to\mathcal{J}(T')} \simeq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The proof is analogous to the proof of Lemma 2.5.8, using the techniques of Proposition 3.5.3. Note that in Chapter 2, the upper right entry appeared as a $*$. Due the Maslov index computation in Equation (3.2), both off diagonal entries are forced to be zero in our case. Philosophically this is because quasi-stabilization is a stronger degeneration than free-stabilization. As in Chapter 2, we make the following definition:

Definition 3.5.6. We say that N is **sufficiently large** for gluing data $\mathcal{J}$, if $\Phi_{\mathcal{J}(T)\to\mathcal{J}(T')}$ is of the form in the previous lemma for all $T, T' \geq N$. We say T is **large enough to compute** $S^{\pm}_{w,z}$ for the gluing data $\mathcal{J}$ if $T > N$ for some N which is sufficiently large.

Adapting the proofs of Lemmas 2.5.10–2.5.16, we have the following:

Lemma 3.5.7. *If $\mathcal{J}$ and $\overline{\mathcal{J}}$ are two choices of gluing data with almost complex structures J_s and $\overline{J}_s$ on $\Sigma \times [0,1] \times \mathbb{R}$, respectively, then there is an N such that if $T > N$, and $\mathcal{J}(T)$ and $\overline{\mathcal{J}}(T)$ achieve transversality, then*

$$\Phi_{\mathcal{J}(T)\to\overline{\mathcal{J}}(T)} \simeq \begin{pmatrix} \Phi_{J_s\to\overline{J}_s} & 0 \\ 0 & \Phi_{J_s\to\overline{J}_s} \end{pmatrix}.$$

The previous lemma will be used to show that the quasi-stabilization maps are independent of the choice of gluing data.

3.6 Quasi-stabilization and triangle maps

In this section we prove several results about quasi-stabilizing Heegaard triples, which we will use to prove invariance of the quasi-stabilization maps. The results for quasi-stabilization of Heegaard triples along a single α_s curve are established in [MO10], so we focus on quasi-stabilizing a Heegaard triple along two curves, α_s and β_s. To compute the quasi-stabilization maps $S^{\pm}_{w,z}$, we pick a curve α_s in the surface Σ, but there are many choices of such an α_s curve, so in order to address invariance of the quasi-stabilization maps, we need to show that the maps $S^{\pm}_{w,z}$ commute with the change of diagram map corresponding to moving α_s to α'_s, which can be computed using a Heegaard triple which has been quasi-stabilized along two curves. Our main result is Theorem 3.6.5, which is an analog to our computation of the differential after quasi-stabilizing in Proposition 3.5.3, but for certain Heegaard triples which we have quasi-stabilized along two curves which are allowed to travel throughout the diagram..

3.6.1 Quasi-stabilizing Heegaard triples along a single curve

We now consider Heegaard triples which are quasi-stabilized along a single α_s curve which is allowed to run through the diagram. This was first considered in [MO10]. We state [MO10, Prop. 5.2], which considers the quasi-stabilized configuration shown in Figure 3.5. The result will be useful in showing that the quasi-stabilization maps are invariant under $\boldsymbol{\beta}$–handleslides and $\boldsymbol{\beta}$–isotopies.

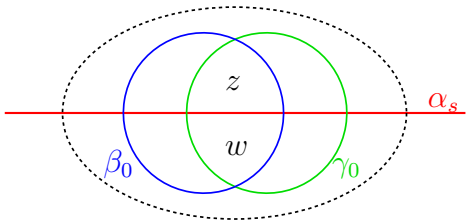

Figure 3.5: **The version of quasi-stabilization discussed in Lemma 3.6.1,** which is considered in [MO10, Prop. 5.2]

Lemma 3.6.1 ([MO10, Prop. 5.2]). *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a strongly $\mathfrak{s}$–admissible triple and suppose that α_s is a new $\boldsymbol{\alpha}$–curve, passing through the point $p \in \Sigma$. Let $\overline{\mathcal{T}}_{\alpha_s,p}$ denote the Heegaard triple resulting from quasi-stabilizing along α_s at p, as in Figure 3.5. If $\mathcal{J}$ is gluing data, then for sufficiently large T, with the almost complex structure $\mathcal{J}(T)$ we have the following identifications:*

$$F_{\overline{\mathcal{T}}_{\alpha_s,p},\overline{\mathfrak{s}}}(\mathbf{x} \times \cdot, \mathbf{y} \times y^+) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) & 0 \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) \end{pmatrix}$$

where $\cdot \in \{x^+, x^-\} = \alpha_s \cap \beta_0$ and the matrix on the right denotes the expansion into the upper and lower generator components of $\alpha_s \cap \beta_0$ and $\alpha_s \cap \gamma_0$.

3.6.2 Strong positivity condition for diagrams of $(S^1 \times S^2)^{\#k}$

In this subsection we describe a class of simple diagrams for $(S^1 \times S^2)^{\#k}$ which we will use in a technical condition in Theorem 3.6.5 for quasi-stabilizing Heegaard triples along two curves, α_s and β_s, passing through a Heegaard triple.

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(S^1 \times S^2)^{\#k}$ such that

$$|\alpha_i \cap \beta_j| = \begin{cases} 1 \text{ or } 2 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

and that if $\alpha_i \cap \beta_i = \{p_i^-, p_i^+\}$, then p_i^- and p_i^+ differ by Maslov grading one. We do not assume that the α_i curves are small isotopies of the β_j curves. Let $\theta^+ = p_1^+ \times \cdots \times p_{n-1}^+$ denote the top graded (partial) intersection point. Assume that $|\alpha_n \cap \beta_n| = 2$ and write p_n^- and p_n^+ for the two points of $\alpha_n \cap \beta_n$.

Definition 3.6.2. Under the assumptions of the previous paragraph, we say that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is **strongly positive** with respect to p_n^+ if for every nonnegative disk $\phi \in \pi_2(\theta^+ \times p_n^+, \mathbf{y} \times p_n^+)$ we have that

$$(\mu - (m_1 + m_2))(\phi) = (\mu - (n_1 + n_2))(\phi) \geq 0,$$

with equality to zero iff ϕ is the constant disk. Here m_1, m_2, n_1, n_2 denote the multiplicities adjacent to the point p_n^+, appearing in the following counterclockwise order: n_1, m_1, n_2, then m_2.

Note that $m_1 + m_2 = n_1 + n_2$ for any disk $\phi \in \pi_2(\mathbf{x} \times p_n^+, \mathbf{y} \times p_n^+)$, by the vertex relations.

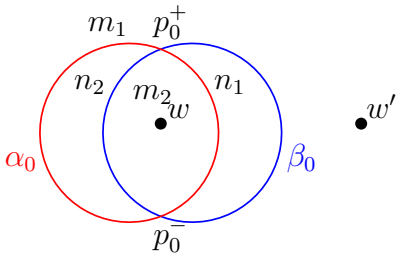

Figure 3.6: **The diagram $\mathcal{H}_0 = (S^2, \alpha_0, \beta_0, w, w')$ in Lemma 3.6.3** which is strongly positive at p_0^+, as well as the multiplicities m_1, n_1, m_2 and n_2.

We now describe a class of diagrams which are strongly positive at an intersection point, which will be sufficient for our purposes:

Lemma 3.6.3. *The diagram $\mathcal{H}_0 = (S^2, \alpha_0, \beta_0, w, w')$ in Figure 3.6 is strongly positive with respect to p_0^+, the intersection point of α_0 and β_0 with higher relative grading. If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram with a distinguished intersection point $p_n^+ \in \alpha_n \cap \beta_n$ where $|\alpha_n \cap \beta_n| = 2$, and $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}')$ is the result of any of the following moves, then $\mathcal{H}'$ is strongly positive with respect to p_n^+ iff $\mathcal{H}$ is strongly positive with respect to p_n^+:*

1. *$(1,2)$–stabilization[1];*
2. *taking the disjoint union of $\mathcal{H}$ with the standard diagram $(\mathbb{T}^2, \alpha_0, \beta_0, w)$ for (S^3, w).*
3. *performing surgery on an embedded 0–sphere $\{q_1, q_2\} \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ by removing small disks from Σ, and connecting the resulting boundary components with an annulus with new α_0 and β_0 curves with $|\alpha_0 \cap \beta_0| = 2$, and which are isotopic to each other, and homotopically nontrivial in the annulus;*

[1] *If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a Heegaard diagram for Y, a $(1,2)$–stabilization of $\mathcal{H}$ is obtained by taking an embedded torus $\mathbb{T}^2$ inside of a 3–ball in $Y \setminus \Sigma$, together with curves α_0 and β_0 with $|\alpha_0 \cap \beta_0| = 1$, which bound compressing disks with boundary on Σ, and letting $\mathcal{H}'$ be the diagram $(\Sigma \# \mathbb{T}^2, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w})$, where $\Sigma \# \mathbb{T}^2$ is the internal connected sum.*

Proof. We first note that $\mathcal{H}_0$ is strongly positive with respect to p_0^+, because the Maslov index of any disk is given by

$$\mu(\phi) = (m_1 + m_2 + n_1 + n_2)(\phi),$$

by Lemma 3.5.2, where m_1, m_2, n_1, n_2 are multiplicities appearing in the counterclockwise order m_1, n_1, m_2, n_2 around p_0^+, as in Figure 3.6. Hence for any disk ϕ, we have

$$(\mu - (n_1 + n_2))(\phi) = (m_1 + m_2)(\phi),$$

which is certainly nonnegative. For a disk $\phi \in \pi_2(p_0^+, p_0^+)$ we also have $(m_1 + m_2)(\phi) = (n_1 + n_2)(\phi)$, so the above quantity is positive iff ϕ has positive multiplicities.

We now address moves (1)–(3).

We first consider move (1). If $\mathcal{H}$ is a diagram, and $\mathcal{H}'$ is the result of $(1,2)$–stabilization, we note that there is an isomorphism

$$\sigma_*\colon\ \pi_2^{\mathcal{H}}(\mathbf{x} \times p_n^+, \mathbf{y} \times p_n^+) \to \pi_2^{\mathcal{H}'}(\mathbf{x} \times c \times p_n^+, \mathbf{y} \times c \times p_n^+),$$

where c is the intersection of the new $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves. Furthermore we note that

$$(\mu - (n_1 + n_2))(\phi) = (\mu - (n_1 + n_2))(\sigma_*\phi),$$

from which the claim follows easily.

We now consider move (2). Suppose $\mathcal{H}'$ is formed from $\mathcal{H}$ by taking the disjoint union of $\mathcal{H}$ with a diagram $(\mathbb{T}^2, \alpha_0, \beta_0, w)$. Homology classes on $\mathcal{H}'$ are of the form $\phi \sqcup k \cdot [\mathbb{T}^2]$, where ϕ is a homology disk on $\mathcal{H}$. One has

$$\mu(\phi \sqcup k \cdot [\mathbb{T}^2]) = \mu(\phi) + 2k,$$

from which the claim follows easily.

Finally we consider move (3), corresponding to surgering on an embedded 0–sphere $\{q_1, q_2\} \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w})$. Write α_0 and β_0 for the new curves on the annulus, and

$\{\theta_0^+, \theta_0^-\} = \alpha_0 \cap \beta_0$. Suppose that $y \in \alpha_0 \cap \beta_0$ is a choice of intersection point. We can define (noncanonically) an injection

$$\iota_y \colon \pi_2^{\mathcal{H}}(\theta^+ \times p_n^+, \mathbf{y} \times p_n^+) \to \pi_2^{\mathcal{H}'}(\theta^+ \times \theta_0^+ \times p_n^+, \mathbf{y} \times y \times p_n^+).$$

For $y = \theta_0^+$, we define $\iota_{\theta_0^+}(\phi)$ to be the disk on the surgered diagram which has no change across the curve β_0, but which agrees with the disk ϕ away from the α_0 and β_0 curves. For $y = \theta_0^-$, a map $\iota_{\theta_0^-}$ can be defined by defining it to be the the map $\iota_{\theta_0^+}$, defined above, composed with the map on disks obtained by splicing in a choice of one of the bigons from θ_0^+ to θ_0^-. An easy computation shows that

$$(\mu - (m_1 + m_2))(\iota_{\theta_0^+}(\phi)) = (\mu - (m_1 + m_2))(\phi)$$

while

$$(\mu - (m_1 + m_2))(\iota_{\theta_0^-}(\phi)) = (\mu - (m_1 + m_2))(\phi) + 1.$$

Any disk in $\pi_2(\theta^+ \times \theta_0^+ \times p_n^+, \mathbf{y} \times y \times p_n^+)$ is equal to one which is in the image of ι_y, with $n \cdot \mathcal{P}$ spliced in, where $\mathcal{P}$ is the periodic domain which is $+1$ in one of the small strips between α_0 and β_0, and -1 in the other. We note that

$$\mu(\mathcal{P}) = m_1(\mathcal{P}) = m_2(\mathcal{P}) = 0.$$

From these observations it follows easily that $\mathcal{H}$ is strongly positive with respect to p_n^+ iff $\mathcal{H}'$ is. □

Remark 3.6.4. Suppose $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w})$ is obtained by taking attaching curves $\boldsymbol{\alpha}$ and letting $\boldsymbol{\alpha}'$ be small Hamiltonian isotopies of the curves in $\boldsymbol{\alpha}$. Let α_s be a new curve in $\Sigma \setminus \boldsymbol{\alpha}$ which doesn't intersect any $\boldsymbol{\alpha}'$–curves. If α_s' is the result of handlesliding α_s across a curve in $\boldsymbol{\alpha}$, then $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_s'\}, \boldsymbol{\alpha} \cup \{\alpha_s\}, \mathbf{w} \cup \{w\})$ is strongly positive at p^+, the intersection point of $\alpha_s' \cap \alpha_s$ with higher grading. Here w is a new basepoint in one of the regions adjacent to p^+.

Similarly if $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w})$ is the result of handlesliding a curve in $\boldsymbol{\alpha}$ across another curve in $\boldsymbol{\alpha}$, and α_s and α_s' are two new curves which are Hamiltonian isotopies of each other,

then $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_s'\}, \boldsymbol{\alpha} \cup \{\alpha_s\}, \mathbf{w} \cup \{w\})$ is strongly positive with respect to the intersection point of $\alpha_s' \cap \alpha_s$ of higher relative grading.

Note that a diagram $(\Sigma_g, \boldsymbol{\alpha}', \boldsymbol{\alpha}, w)$ where the curves in $\boldsymbol{\alpha}'$ are small isotopies of the curves in $\boldsymbol{\alpha}$ with $g(\Sigma_g) = |\boldsymbol{\alpha}'| = |\boldsymbol{\alpha}| = g$, is not a strongly positive diagram at any point, since $[\Sigma_g]$ represents a positive homology class in $\pi_2(\theta^+, \theta^+)$ with $\mu([\Sigma_g]) - m_1 - m_2 = 2 - 1 - 1 = 0$. Strongly positive diagrams always have multiple basepoints. The prototypical example is the one resulting from handlesliding a quasi-stabilization curve α_s across an $\boldsymbol{\alpha}$–curve, as in Figure 3.7.

A more interesting example of diagram which doesn't satisfy the strong positivity condition would be the pair $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_s'\}, \boldsymbol{\alpha} \cup \{\alpha_s\}, \mathbf{w})$ arising in a triple $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_s'\}, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w})$ for handlesliding an $\boldsymbol{\alpha}$ curve across α_s. Fortunately such a move is not required in the proof of invariance of the quasi-stabilization maps.

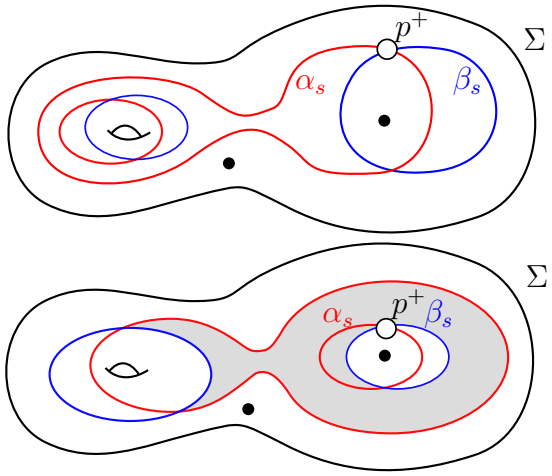

Figure 3.7: **Two diagrams illustrating the strong positivity condition.** The diagram on the top is strongly positive with respect to the point p^+. The curves α_s and β_s are curves on which one could perform the quasi-stabilization operation of triangles in Theorem 3.6.5. The diagram on the bottom is not, and the nonzero, nonnegative domain of a disk ϕ with $\mu(\phi) - (n_1 + n_2)(\phi) = 0$ is shown.

3.6.3 Quasi-stabilizing Heegaard triples along two curves

We now consider the effect on the triangle maps of quasi-stabilizing along two curves. Our analysis follows a similar spirit to the proof of [MO10, Prop. 5.2]. Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with a distinguished point $p^+ \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w} \cup \mathbf{z})$. Suppose also that α_s and β_s are two choices of curves in $\Sigma \setminus (\boldsymbol{\alpha} \cup \mathbf{w} \cup \mathbf{z})$ and $(\Sigma \setminus (\boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$ respectively, which intersect only at p^+, and another point $p^- \in \Sigma$. We can form the diagram $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$, obtained by quasi-stabilizing along both α_s and β_s, simultaneously, at the point p^+. This corresponds to removing a small disk containing p^+, and inserting the diagram shown in Figure 3.8, with a disk centered around p_0^- removed.

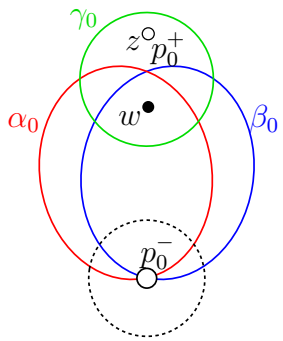

Figure 3.8: **The diagram we insert into a Heegaard triple diagram $\mathcal{T}$ along the curves α_s and β_s to form the diagram $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$.** We cut out the solid circle marked with p_0^- and stretch the almost complex structure along the dashed circle.

Theorem 3.6.5. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with curves α_s and β_s, intersecting at two points p^+ and p^- and let $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$ be the Heegaard triple result from quasi-stabilization, as described above. If $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ is a strongly positive diagram for $(S^1 \times S^2)^{\#k}$ with respect to the point p^+ (Definition 3.6.2), and $\mathcal{J}$ is gluing data for stretching along the dashed circle, then for sufficiently large T, with respect to the almost complex structure $\mathcal{J}(T)$, there are identifications*

$$F_{\mathcal{T}^+_{\alpha_s, \beta_s, p^+}, \bar{\mathfrak{s}}, \mathcal{J}(T)}(\Theta^+_{\alpha\beta} \times p_0^+, \mathbf{x} \times \cdot) = \begin{pmatrix} F_{\mathcal{T}, \mathfrak{s}, J}(\Theta^+_{\alpha\beta}, \mathbf{x}) & 0 \\ 0 & F_{\mathcal{T}, \mathfrak{s}, J}(\Theta^+_{\alpha\beta}, \mathbf{x}) \end{pmatrix},$$

where $\cdot$ denotes $x^{\pm} \in \beta_0 \cap \gamma_0$ and the matrix on the right denotes the matrix decomposition of the map based on the decompositions given $x^{\pm}$ and $y^{\pm}$.

As usual, the argument proceeds by a Maslov index calculation, which we use to put constraints on the homology classes of holomorphic curves which can appear in a weak limit as we let the parameter T approach $+\infty$. Once we determine which homology classes of triangles can appear, we can use standard gluing results to explicitly count holomorphic curves.

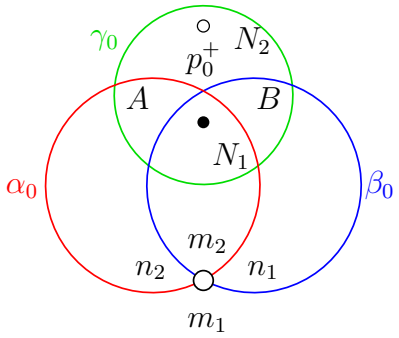

Figure 3.9: **Multiplicities for a triangle on the diagram $(S^2, \alpha_0, \beta_0, \gamma_0, w, z)$.**

Lemma 3.6.6. *Suppose $\psi \in \pi_2(x, y, z)$ is a homology disk on $(S^2, \alpha_0, \beta_0, \gamma_0, w, z)$, shown in Figure 3.9. Then*

$$\mu(\psi) = n_1 + n_2 + N_1 + N_2.$$

Proof. The formula is easily checked for any of the Maslov index zero small triangles, and respects splicing in any Maslov index 1 strip. Since any two triangles on this diagram differ by splicing in some number of the Maslov index 1 strips, the formula follows in full generality. □

We can now prove Theorem 3.6.5.

Proof of Theorem 3.6.5. Suppose that u_i is a sequence of holomorphic triangles of Maslov index zero representing a class $\psi \in \pi_2(\Theta^+_{\alpha\beta} \times p_0^+, \mathbf{x} \times x, \mathbf{y} \times y)$, for almost complex structure $\mathcal{J}(T_i)$, where T_i is a sequence of neck lengths approaching $+\infty$. Adapting the proof of

Proposition 3.5.3, the limiting curves which appear can be arranged into three classes of broken holomorphic curves:

1. a broken holomorphic triangle u_Σ representing a homology class ψ_Σ on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ which has no boundary components on α_s or β_s;

2. a broken holomorphic disk $u_{\alpha\beta}$ on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\})$ representing a class $\phi_{\alpha\beta} \in \pi_2(\Theta^+_{\alpha\beta} \times p^+, \mathbf{y} \times p^+)$, for some $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$;

3. a broken holomorphic triangle u_0 on $(S^2, \alpha_0, \beta_0, \gamma_0)$ which represents a homology triangle $\psi_0 \in \pi_2(p_0^+, x, y)$.

We now wish to write down the Maslov index of ψ in terms of the Maslov indices and multiplicities of ψ_Σ, ψ_0 and $\phi_{\alpha\beta}$. Let $m_1(\cdot), m_2(\cdot), n_1(\cdot)$, and $n_2(\cdot)$ denote the the multiplicities of a homology curve in the regions surrounding p^+ or p_0^-, as in Figure 3.9. In [Sar11b], Sarkar derives a formula for the Maslov index of a homology triangle $\rho \in \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z})$ which can be computed entirely from the domain $\mathcal{D}(\rho)$. Writing $\mathcal{D} = \mathcal{D}(\rho)$, the formula reads

$$\mu(\rho) = e(\mathcal{D}) + n_{\mathbf{x}}(\mathcal{D}) + n_{\mathbf{y}}(\mathcal{D}) + a(\mathcal{D}).c(\mathcal{D}) - \frac{d}{2},$$

where $d = |\boldsymbol{\alpha}| = |\boldsymbol{\beta}| = |\boldsymbol{\gamma}|$. Here $a(\mathcal{D})$ is defined intersection $\partial\mathcal{D} \cap \boldsymbol{\alpha}$ (viewed as a 1–chain), and $c(\mathcal{D})$ is defined similarly, using the $\boldsymbol{\gamma}$–curves. The quantity $a(\mathcal{D}).c(\mathcal{D})$ is defined as the average of the four algebraic intersection numbers of $a'(\mathcal{D})$ and $c(\mathcal{D})$, where $a'(\mathcal{D})$ is a translate of $a(\mathcal{D})$ in any of the four "diagonal directions". If $s \in \alpha_i \cap \beta_j$, then $n_s(\mathcal{D})$ is the average of the multiplicities in the regions surrounding s, and if $\mathbf{s}$ is a set of such intersection points, then $n_{\mathbf{s}}(\mathcal{D})$ is the sum of the $n_s(\mathcal{D})$ ranging over $s \in \mathbf{s}$.

For a homology triangle $\psi \in \pi_2(\Theta^+_{\alpha\beta} \times p_0^+, \mathbf{x} \times x, \mathbf{y} \times y)$ which can be decomposed into homology classes $\psi_\Sigma, \phi_{\alpha\beta}$, and ψ_0 as above (as any homology class admitting holomorphic representatives for arbitrarily large neck length can) we observe that $\mu(\psi)$ can be computed by adding up $\mu(\psi_\Sigma), \mu(\phi_{\alpha\beta})$ and $\mu(\psi_0)$, then subtracting the quantities which are over-counted. This corresponds to subtracting $\frac{1}{2}(m_1 + m_2 + n_1 + n_2)(\psi)$, which is the excess of

Euler measure resulting from removing balls centered at p^+ and at p_0^- (note that the Euler measure of a quarter disk is $\frac{1}{4}$), and subtracting

$$2n_{p^+}(\phi_{\alpha\beta}) = \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta}),$$

which is the quantity in the expression

$$\mu(\phi_{\alpha\beta}) = e(\mathcal{D}(\phi_{\alpha\beta})) + n_{\Theta^+_{\alpha\beta}\times p^+}(\phi_{\alpha\beta}) + n_{\mathbf{y}\times p^+}(\phi_{\alpha\beta})$$

$$= e(\mathcal{D}(\phi_{\alpha\beta})) + n_{\Theta^+_{\alpha\beta}}(\phi_{\alpha\beta}) + n_{\mathbf{y}}(\phi_{\alpha\beta}) + \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta})$$

which does not contribute to $\mu(\psi)$. Adding these contributions, we get

$$\mu(\psi) = \mu(\psi_\Sigma)+\mu(\psi_0)+\mu(\phi_{\alpha\beta})-\tfrac{1}{2}(n_1+n_2+m_1+m_2)(\psi)-\tfrac{1}{2}(n_1+n_2+m_1+m_2)(\phi_{\alpha\beta}). \quad (3.3)$$

Writing $\psi_0 \in \pi_2(p_0^+, x, y)$, using the vertex multiplicity relations around p_0^-, it is an easy computation that

$$(m_1 + m_2)(\psi) = (n_1 + n_2)(\psi).$$

Note also that $m_i(\psi_0) = m_i(\psi)$ and similarly for the multiplicities n_i, since we are grouping all holomorphic curves on $(S^2, \alpha_0, \beta_0, \gamma_0)$ appearing in the weak limit into the homology class ψ_0. Using the Maslov index formula from Lemma 3.6.6 for ψ_0, we get from Equation (3.3) that

$$\mu(\psi) = \mu(\psi_\Sigma) + \mu(\phi_{\alpha\beta}) + (m_1 + m_2 + N_1 + N_2)(\psi_0)$$

$$-\tfrac{1}{2}(m_1 + m_2 + n_1 + n_2)(\psi_0) - \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta})$$

which reduces to

$$\mu(\psi) = \mu(\psi_\Sigma) + (N_1 + N_2)(\psi_0) + \mu(\phi_{\alpha\beta}) - \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta}), \quad (3.4)$$

since ψ_0 does not have p^+ as a vertex, so the vertex relations at p^+ yield

$$\tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\psi_0)$$

$$= (n_1 + n_2)(\psi_0) = (m_1 + m_2)(\psi_0).$$

We now use the assumption that $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w})$ is strongly positive with respect to p^+, and hence

$$\mu(\phi_{\alpha\beta}) - \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta}) \geq 0,$$

with equality iff $\phi_{\alpha\beta}$ is a constant disk. We note that $\mu(\psi_\Sigma) \geq 0$, since ψ_Σ admits a broken holomorphic representative. Hence the formula for $\mu(\psi)$ in Equation (3.4) can be written as a sum of nonnegative expressions, and hence each must be zero if $\mu(\psi) = 0$. Thus

$$0 = \mu(\psi) = \mu(\psi_\Sigma) = \mu(\phi_{\alpha\beta}) = N_1 = N_2.$$

Note first that this implies that $\phi_{\alpha\beta}$ is a constant disk, since generically there are no nonconstant, Maslov index zero disks which admit broken holomorphic representatives. From here the argument proceeds in a familiar manner. Note that ψ_Σ is a Maslov index 0 broken holomorphic triangle, and hence must be a genuine holomorphic triangle, as nonconstant holomorphic disks, boundary degenerations, and closed surfaces all have strictly positive Maslov index. As such, since ψ_Σ is represented by a holomorphic triangle with no boundary components mapped to α_s or β_s, we must have $m_1(\psi) = m_2(\psi) = n_1(\psi) = n_2(\psi)$.

We now claim that this implies that the off diagonal entries of the matrix representing the triangle map in the statement are zero. Writing $\psi_0 \in \pi_2(p_0^+, x, y)$, and using the multiplicities in Figure 3.9, the vertex relations at p_0^+ read

$$A + B = 1,$$

and hence exactly one of A and B is 1, and the other is 0, since $A, B \geq 0$. Since $n_1 = n_2 = m_1 = m_2$, we know that by subtracting some number of copies of the γ_0–boundary degeneration of Maslov index 2 with $N_1 = N_2 = 0$, we get a homology triangle in $\pi_2(p_0^+, x, y)$ with N_1, N_2, m_1, m_2, n_1, and n_2 all zero. There are two only homology triangles satisfying that condition. One is in $\pi_2(p_0^+, x^+, y^+)$ and the other is in $\pi_2(p_0^+, x^-, y^-)$, implying that ψ_0 itself must be in one of those sets. Hence the off diagonal entries of the matrix are zero.

Since u_Σ is a genuine Maslov index zero holomorphic triangle, there must be a curve in u_0'

in the broken holomorphic triangle u_0 which matches, i.e. which satisfies

$$\rho^{p^+}(u_\Sigma) = \rho^{p_0}(u_0').$$

Recall that if $u\colon\ S \to \Sigma \times \Delta$ is a holomorphic map and $q \in \Sigma$ is a point, we define

$$\rho^q(u) = (\pi_\Delta \circ u)(\pi_\Sigma \circ u)^{-1}(q) \in \mathrm{Sym}^{n_q(u)}(\Delta).$$

Since this in particular forces $n_{p^+}(u_\Sigma) = n_{p_0^-}(u_0')$, it is easy to see that there can be no other curves in the broken curve u_0 since there are no multiplicities on $(S^2, \alpha_0, \beta_0, \gamma_0)$ which could be increased without increasing n_1, m_1, n_2, m_2, N_1, and N_2 while still preserving the vertex relations. By standard gluing arguments (cf. [`?LCylindrical`, App. A]) the count of prematched triangles[2] is equal to the count holomorphic triangles in $\#\mathcal{M}_{\mathcal{J}(T)}(\psi_\Sigma \# \psi_0)$ for sufficiently large T.

For x and y of the same relative grading (both the top intersection points or both the lower intersection points), for each k there is a unique homology class ψ_k on $(S^2, \alpha_0, \beta_0, \gamma_0)$ in $\pi_2(p_0^+, x, y)$ with $n_1 = n_2 = m_1 = m_2 = k$. Thus, adapting the proof of Proposition 3.5.3, it is sufficient to count holomorphic triangles on $(S^2, \alpha_0, \beta_0, \gamma_0)$ of homology class ψ_k which match a fixed divisor $\mathbf{d} \in \mathrm{Sym}^k(\Delta)$. The count of such holomorphic triangles matching $\mathbf{d}$ is generically 1 (mod 2), as can be seen from a Gromov compactness argument nearly identical to the one done at the end of Proposition 3.5.3. Hence the diagonal entries of the triangle map matrix are as claimed, completing the proof. □

Remark 3.6.7. Without the "strongly positive" condition, the counts of the previous theorem are false. The bottom of Figure 3.7 shows a diagram which is not strongly positive, as well as a Maslov index 1 disk which could appear on $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ when we degenerate. Fortunately such pairs $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ don't appear proving invariance of the quasi-stabilization maps, as we don't need to handleslide other $\boldsymbol{\alpha}$ curves across α_s.

[2]Recall that a prematched triangle is a pair (u_Σ, u_0) where u_Σ and u_0 are holomorphic triangles representing ψ_Σ and ψ_0 on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ and $(S^2, \alpha_0, \beta_0, \gamma_0)$ respectively and $\rho^{p^+}(u_\Sigma) = \rho^{p_0}(u_0')$.

3.7 Invariance of the quasi-stabilization maps

In this section, we combine the results of the previous section to prove invariance of the quasi-stabilization maps:

Theorem A. *Assume that $(\sigma, \mathfrak{P})$ is a coloring of the basepoints $\mathbf{w} \cup \mathbf{z}$ which is extended by the coloring $(\sigma', \mathfrak{P})$ of the basepoints $\mathbf{w} \cup \mathbf{z} \cup \{w, z\}$. The quasi-stabilization operation induces well defined maps*

$$S^+_{w,z} \colon CFL^\infty_{UV}(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}) \to CFL^\infty_{UV}(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S^-_{w,z} \colon CFL^\infty_{UV}(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to CFL^\infty_{UV}(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

which are well defined invariants up to $\mathfrak{P}$–filtered $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain homotopy.

The proof is to construct the maps for choices of Heegaard diagram and auxiliary data, and show that the maps we describe are independent from that auxiliary data and the choice of diagram.

If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, recall from Subsection 3.5.1 that if A denotes the component of $\Sigma \setminus \boldsymbol{\alpha}$ containing the basepoints of $(L, \mathbf{w}, \mathbf{z})$ adjacent to w and z, then we pick a point $p \in A \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$ and a simple closed curve $\alpha_s \subseteq A \setminus \boldsymbol{\alpha}$ to form a diagram $\overline{\mathcal{H}}_{p,\alpha_s}$. Let $\mathcal{J}$ denote gluing data (cf. Subsection 3.5.2) for performing the special connected sum operation at $p \in \Sigma$ and $p_0 \in S^2$ and gluing almost complex structures on $\mathcal{H}$ and $\mathcal{H}_0$ together.

We now define maps

$$S^+_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J},T} \colon CFL^\infty_{UV,J_s}(\mathcal{H}, \mathfrak{s}) \to CFL^\infty_{UV,\mathcal{J}(T)}(\overline{\mathcal{H}}_{p,\alpha_s}, \mathfrak{s})$$

and

$$S^-_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J},T} \colon CFL^\infty_{UV,\mathcal{J}(T)}(\overline{\mathcal{H}}_{p,\alpha_s}, \mathfrak{s}) \to CFL^\infty_{UV,J_s}(\mathcal{H}, \mathfrak{s})$$

by the formulas

$$S^+_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

and

$$S^-_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}(\mathbf{x} \times \theta^+) = 0 \qquad \text{and} \qquad S^+_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}(\mathbf{x} \times \theta^-) = \mathbf{x},$$

where J_s denotes the almost complex structure on Σ associated to the gluing data $\mathcal{J}$ and T is sufficiently large. Here θ^+ denotes the top degree intersection point with respect to the Maslov grading (the grading obtained by using the $\mathbf{w}$–basepoints and ignoring the $\mathbf{z}$–basepoints).

The maps $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J},T}$ can be extended to the entire $\mathfrak{P}$–filtered chain homotopy-type invariant by pre- and post-composing with change of diagram maps and change of almost complex structure maps. By functoriality of the change of diagrams maps, we get well defined maps $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J},T}$ between the coherent chain homotopy type invariants $CFL^{\infty}_{UV}(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s})$ and $CFL^{\infty}_{UV}(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s})$, though of course we need to show independence from the choices of $\mathcal{H}, p, \alpha_s, \mathcal{J}$, and T.

We note that any diagram $\overline{\mathcal{H}}$ for $(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$ can be connected to one of the form $\overline{\mathcal{H}}_{p,\alpha_s}$ for a diagram $\mathcal{H}$ and a choice of p and α_s by a sequence of handleslides and isotopies of the attaching curves, $(1,2)$–stabilizations, and isotopies of Y relative the basepoints and preserving the link, since we can always find a diagram for the unstabilized link and quasi-stabilize it, and any two diagrams for the same multi-based link can be connected by a sequence of Heegaard moves by [JT12, Proposition 2.37], for example. This is somewhat unsatisfying since it would be nice to have an actual algorithm for reducing an arbitrary Heegaard diagram for the quasi-stabilized link to a quasi-stabilized diagram, but for our purposes it is sufficient to know that such a path exists.

We now begin our proof of invariance of the maps $S^{\pm}_{w,z}$. We first address independence from $\mathcal{J}$ and the parameter T. Recalling Lemma 3.5.5, there is an N such that if $T, T' > N$ we have

$$\Phi_{\mathcal{J}(T) \to \mathcal{J}(T')} \simeq \begin{pmatrix} \mathrm{id} & 0 \\ 0 & \mathrm{id} \end{pmatrix}.$$

Hence, as with the free stabilization maps defined in Chapter 2, we define the maps $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}$ to be between the complexes $CFL^{\infty}_{J_s,UV}(\mathcal{H},\mathfrak{s})$ and $CFL^{\infty}_{\mathcal{J}(T),UV}(\overline{\mathcal{H}}_{p,\alpha_s},\mathfrak{s})$ for any T greater than any such N.

Lemma 3.7.1. *The maps $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}$ are independent of $\mathcal{J}$ and the parameter T.*

Proof. Lemma 3.5.5 implies that $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J},T}$ is independent of T for any T which is sufficiently large. Let $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}$ denote the common map. Lemma 3.5.7 implies that the maps $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s,\mathcal{J}}$ are independent of $\mathcal{J}$. We denote the common map from now on by $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s}$. □

Lemma 3.7.2. *For a fixed diagram $\mathcal{H}$ with fixed $p \in \Sigma$, the maps $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s}$ are independent of α_s.*

Proof. This follows from the triangle map computation in Theorem 3.6.5, which allows one to change the α_s curve through a sequence of isotopies of the α_s curve, and handleslides of the α_s curve over other $\boldsymbol{\alpha}$–curves, each of which can be realized by quasi-stabilizing a Heegaard triple $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z})$ along two curves α'_s and α_s with $\alpha'_s \cap \alpha_s = \{p^+,p^-\}$, such that $(\Sigma,\boldsymbol{\alpha}' \cup \{\alpha'_s\},\boldsymbol{\alpha} \cup \{\alpha_s\},\mathbf{w})$ is strongly positive at p^+ (cf. Remark 3.6.4). For each of these moves, by Theorem 3.6.5 the change of diagrams map can be written as

$$\Phi^{\boldsymbol{\alpha}\cup\{\alpha_s\}\to\boldsymbol{\alpha}'\cup\{\alpha'_s\}}_{\boldsymbol{\beta}\cup\{\beta_s\}} = \begin{pmatrix} \Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}} & 0 \\ 0 & \Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}} \end{pmatrix},$$

where the curves in $\boldsymbol{\alpha}'$ are small Hamiltonian isotopies of the curves in $\boldsymbol{\alpha}$. Similarly by Theorem 3.6.5 we also have

$$\Phi^{\boldsymbol{\alpha}\cup\{\alpha_s\}\to\boldsymbol{\alpha}'\cup\{\alpha_s\}}_{\boldsymbol{\beta}\cup\{\beta_s\}} = \begin{pmatrix} \Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}} & 0 \\ 0 & \Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}} \end{pmatrix},$$

completing the proof. □

We now let $S^{\pm}_{w,z,\mathcal{H},p}$ denote the map $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s}$ for any choice of α_s. We now prove independence from the choice of point $p \in \Sigma$.

Lemma 3.7.3. *Given a fixed diagram $\mathcal{H}$, the maps $S^{\pm}_{w,z,\mathcal{H},p}$ are independent of the choice of point $p \in \Sigma$.*

Proof. Let A denote the component of $\Sigma \setminus \boldsymbol{\alpha}$ containing the basepoints on $\mathbb{L}$ which are adjacent to w and z. Let p and p' be two choices of points in $A \setminus (\boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$. Let ϕ_t be an isotopy $\phi_t \colon \Sigma \to \Sigma$ which fixes $\Sigma \setminus A$ and maps p to p'. Recall that the surfaces $\overline{\Sigma}_p$ were well defined up to an isotopy fixing $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z}$ and mapping L to L. Extend $\phi = \phi_1$ to an isotopy of Y which fixes $(\Sigma \setminus A) \cup \mathbf{w} \cup \mathbf{z} \cup \{w, z\}$ and maps L to L. By definition

$$(\phi)_*(\overline{\Sigma}_p) = \overline{\Sigma}_{p'},$$

as embedded surfaces. The diffeomorphism ϕ fixes all the curves in $\boldsymbol{\alpha}$, but moves some of the $\boldsymbol{\beta}$–curves which pass through the region A.

Observe the following factorizations:

$$\Phi_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},J_s)\to(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\phi_* J_s)} \simeq \Phi^{\boldsymbol{\alpha}}_{\phi_*\boldsymbol{\beta}\to\boldsymbol{\beta}} \circ \phi_* \tag{3.5}$$

and similarly

$$\Phi_{(\overline{\Sigma}_p,\boldsymbol{\alpha}\cup\{\alpha_s\},\boldsymbol{\beta}\cup\{\beta_0\},\mathcal{J}(T))\to(\overline{\Sigma}_{p'},\boldsymbol{\alpha}\cup\{\phi_*\alpha_s\},\boldsymbol{\beta}\cup\{\phi_*\beta_0\},\phi_*\mathcal{J}(T))} \simeq \Phi^{\boldsymbol{\alpha}\cup\{\phi_*\alpha_s\}}_{(\phi_*\boldsymbol{\beta})\cup\{\phi_*\beta_0\}\to\boldsymbol{\beta}\cup\{\phi_*\beta_0\}} \circ \phi_*.$$

Using Theorem 3.6.5, for sufficiently stretched almost complex structure we can write

$$\Phi^{\boldsymbol{\alpha}\cup\{\phi_*\alpha_s\}}_{(\phi_*\boldsymbol{\beta})\cup\{\phi_*\beta_0\}\to\boldsymbol{\beta}\cup\{\phi_*\beta_0\}} \simeq \begin{pmatrix} \Phi^{\boldsymbol{\alpha}}_{\phi_*\boldsymbol{\beta}\to\boldsymbol{\beta}} & 0 \\ 0 & \Phi^{\boldsymbol{\alpha}}_{\phi_*\boldsymbol{\beta}\to\boldsymbol{\beta}} \end{pmatrix},$$

and hence

$$\Phi^{\boldsymbol{\alpha}\cup\{\phi_*\alpha_s\}}_{(\phi_*\boldsymbol{\beta})\cup\{\phi_*\beta_0\}\to\boldsymbol{\beta}\cup\{\phi_*\beta_0\}} \circ \phi_* \simeq \begin{pmatrix} \Phi^{\boldsymbol{\alpha}}_{\phi_*\boldsymbol{\beta}\to\boldsymbol{\beta}} \circ \phi_* & 0 \\ 0 & \Phi^{\boldsymbol{\alpha}}_{\phi_*\boldsymbol{\beta}\to\boldsymbol{\beta}} \circ \phi_* \end{pmatrix}.$$

Combining this with Equation (3.5), we see that for sufficiently large T, we have

$$\begin{aligned} &\Phi_{(\overline{\Sigma}_p,\boldsymbol{\alpha}\cup\{\alpha_s\},\boldsymbol{\beta}\cup\{\beta_0\},\mathcal{J}(T))\to(\overline{\Sigma}_{p'},\boldsymbol{\alpha}\cup\{\phi_*\alpha_s\},\boldsymbol{\beta}\cup\{\phi_*\beta_0\},\phi_*\mathcal{J}(T))} \\ &\simeq \begin{pmatrix} \Phi_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},J_s)\to(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\phi_* J_s)} & 0 \\ 0 & \Phi_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},J_s)\to(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\phi_* J_s)} \end{pmatrix}. \end{aligned}$$

Since the above map is upper triangular with diagonal entries equal to the change of diagrams maps, the change of diagrams maps commute with the maps $S^{\pm}_{w,z,\mathcal{H},p}$ and $S^{\pm}_{w,z,\mathcal{H},p'}$ in the appropriate sense. Since any two choices of points p and p' at which we can perform quasi-stabilization are in the region A, we know that the maps on the filtered chain homotopy invariant CFL^{∞}_{UV} induced by $S^{\pm}_{w,z,\mathcal{H},p,\alpha_s\mathcal{J}}$ and $S^{\pm}_{w,z,\mathcal{H},p',\phi_*\alpha_s,\phi_*\mathcal{J}}$ are equal. Since we proved invariance from the gluing data $\mathcal{J}$ in Lemma 3.5.7, and we proved invariance from the curve α_s in Lemma 3.7.2, the proof is thus complete. □

We let $S^{\pm}_{w,z,\mathcal{H}}$ denote the map $S^{\pm}_{w,z,\mathcal{H},p}$ for any choice of p in the component of $\Sigma \setminus \boldsymbol{\alpha}$ containing the basepoints of L adjacent to w and z.

Lemma 3.7.4. *If $\mathcal{H}$ and $\mathcal{H}'$ are two diagrams for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, then the maps $S^{\pm}_{w,z,\mathcal{H}}$ and $S^{\pm}_{w,z,\mathcal{H}'}$ are filtered chain homotopic.*

Proof. Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ and $\mathcal{H}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ are two diagrams for $(L, \mathbf{w}, \mathbf{z})$. The diagrams $\mathcal{H}$ and $\mathcal{H}'$ are related by a sequence of the following moves:

1. $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–handleslides and isotopies;
2. (1,2)–stabilizations away from L;
3. isotopies of Σ inside of Y which fix $\mathbf{w} \cup \mathbf{z}$, and map L to L.

The maps corresponding to $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–handleslides on the unstabilized diagram $\mathcal{H}$ can be computed using triangle maps. For moves of the $\boldsymbol{\beta}$–curves, we simply apply Lemma 3.6.1 to see that the maps $S^{\pm}_{w,z,\mathcal{H}}$ are invariant under $\boldsymbol{\beta}$–isotopies and handleslides. Theorem 3.6.5 implies independence under $\boldsymbol{\alpha}$–moves of $\mathcal{H}$ such that there are curves α_s and α'_s in Σ, with top graded intersection point $p \in \alpha'_s \cap \alpha_s$, such that $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha'_s\}, \boldsymbol{\alpha} \cup \{\alpha_s\}, \mathbf{w})$ is strongly positive with respect to p. An arbitrary $\boldsymbol{\alpha}$–move can be realized as a sequence of such moves, as well as moves of the point p inside of the region of $\Sigma \setminus \boldsymbol{\alpha}$. Since we've already shown invariance under each of these smaller moves, the maps $S^{\pm}_{w,z,\mathcal{H}}$ are unchanged by handleslides and isotopies of the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves.

The maps $S^{\pm}_{w,z,\mathcal{H}}$ obviously commute with the $(1,2)$–stabilization maps.

We now consider isotopies $\phi_t\colon Y \to Y$ which fix $\mathbf{w} \cup \mathbf{z}$ and map L to L. We note that tautologically we have that

$$\phi_* \circ S^{\pm}_{w,z,\mathcal{H},p,\mathcal{J}} = S^{\pm}_{w,z,\phi_*\mathcal{H},\phi_* p,\phi_*\mathcal{J}} \circ \phi_*.$$

Since we already know that $S^{\pm}_{w,z,\mathcal{H},p,\mathcal{J}}$ is invariant from p and $\mathcal{J}$, we thus conclude that the maps $S^{\pm}_{w,z,\mathcal{H}}$ and $S^{\pm}_{w,z,\phi_*\mathcal{H}}$ agree. □

We can now write $S^{\pm}_{w,z}$ for the quasi-stabilization maps, completing the proof of Theorem A.

Remark 3.7.5. Given that the triangle map computations in Lemma 3.6.1 and Theorem 3.6.5 showed that change of diagrams maps were not only upper triangular, but diagonal, one may ask why it is natural to define $S^{+}_{w,z}$ by $\mathbf{x} \mapsto \mathbf{x} \times \theta^+$ and not $\mathbf{x} \mapsto \mathbf{x} \times \theta^-$. We remark that $\mathbf{x} \mapsto \mathbf{x} \times \theta^-$ is only a chain map when w is given the same color as the other $\mathbf{w}$–basepoint adjacent to z, and indeed $\mathbf{x} \mapsto \mathbf{x} \times \theta^-$ is equal to $\Psi_z S^{+}_{w,z}$. The map Ψ_z is only a chain map if the $\mathbf{w}$–basepoints adjacent to z have the same color.

Remark 3.7.6. We have defined quasi-stabilization maps $S^{\pm}_{w,z}$ in the case that w comes after z and showed that such maps were invariants, i.e. that they yielded well defined maps $S^{\pm}_{w,z}$ on the coherent filtered chain homotopy type invariant. These maps were only constructed if w came after z on the link component. In the case that z comes after w, we can define quasi-stabilization maps $S^{\pm}_{z,w}$ analogously, picking a choice of β_s. We could call such an operation β quasi-stabilization. There is no ambiguity between α quasi-stabilizations or β quasi-stabilizations because $S^{\pm}_{w,z}$ is always an α quasi-stabilization and $S^{\pm}_{z,w}$ is always a β quasi-stabilization.

3.8 Commutation of quasi-stabilization maps

In this section we show that if $\{w,z\} \cap \{w',z'\} = \varnothing$, then the maps $S^{\pm}_{w,z}$ and $S^{\pm}_{w',z'}$ all commute. In Chapter 2 we showed that the free stabilization maps commute, though com-

mutation was easier to show in that setting, since we could just pick a diagram and an almost complex structure where both free-stabilization maps could be computed, and by simply looking at the formulas, one could observe that the maps commuted. In the case of quasi-stabilization, we cannot always pick an almost complex structure which can be used to compute both maps. Nevertheless, we can compute enough components of the change of almost complex structure map to show that quasi-stabilization maps commute:

Theorem 3.8.1. *Suppose that $(L, \mathbf{w}, \mathbf{z})$ is a multibased link in Y^3 and that w, z, w' and z' are new basepoints, such that (w, z) and (w', z') are each pairs of adjacent basepoints on $(L, \mathbf{w} \cup \{w, w'\}, \mathbf{z} \cup \{z, z'\})$. Then*

$$S^{\circ_1}_{w,z} \circ S^{\circ_2}_{w',z'} \simeq S^{\circ_2}_{w',z'} \circ S^{\circ_1}_{w,z}$$

for any $\circ_1, \circ_2 \in \{+, -\}$.

Pick a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for $(L, \mathbf{w}, \mathbf{z})$, and let α_s and α'_s be curves in $\Sigma \setminus \boldsymbol{\alpha}$ along which we can perform quasi-stabilization for (w, z) and (w', z') respectively. Let β_0 and β'_0 denote the new $\boldsymbol{\beta}$–curves. Let $\mathcal{J}$ denote gluing data for stretching along circles bounding β_0 and β'_0. There are two distinct cases to consider, corresponding to whether the pairs (w, z) and (w', z') are adjacent or not: either α_s and α'_s lie in the same component of $\Sigma \setminus \boldsymbol{\alpha}$ (this case corresponds to having the pair (w, z) be adjacent to the pair (w', z')), or α_s and α'_s lie in different components of $\Sigma \setminus \boldsymbol{\alpha}$ (this case corresponds to the pair (w, z) not being adjacent to (w', z')).

The first case is the easier to consider. In this case, we now show that we can pick an almost complex structure which computes both quasi-stabilization maps. To this end, we have the following lemma:

Lemma 3.8.2. *Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram as in the previous paragraph with new curves α_s and α'_s for quasi-stabilizing at (w, z) and (w', z') respectively. If α_s and α'_s are not in the same component of $\Sigma \setminus \boldsymbol{\alpha}$, then for all sufficiently large T_1, T'_1, T_2, T'_2, we have*

$$\Phi_{\mathcal{J}(T_1,T'_1)\to\mathcal{J}(T_2,T'_2)} \simeq \mathrm{id},$$

with respect to the obvious identification between the two complexes.

Proof. To show this, we will perform a computation similar to the one in Proposition 3.5.3, but for Maslov index zero disks. Let A be the component of $\Sigma \setminus \boldsymbol{\alpha}$ which contains α_s and let A' denote the component of $\Sigma \setminus \boldsymbol{\alpha}$ which contains α'_s. By assumption $A \neq A'$. Let A_1 and A_2 denote the two components of $A \setminus \alpha_s$. Let A'_1 and A'_2 denote the components of $A' \setminus \alpha'_s$. Suppose that ϕ is a Maslov index 0 homology disk on the doubly quasi-stabilized diagram $(\overline{\Sigma}, \boldsymbol{\alpha} \cup \{\alpha_s, \alpha'_s\}, \boldsymbol{\beta} \cup \{\beta_0, \beta'_0\}, \mathbf{w} \cup \{w, w'\}, \mathbf{z} \cup \{z, z'\})$. Write $\phi = \phi_\Sigma \# \phi_0 \# \phi'_0$ where ϕ_Σ is a homology class on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s, \alpha'_s\}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$, ϕ_0 is a homology class on $(S^2, \alpha_0, \beta_0, w, z)$ and $(S^2, \alpha'_0, \beta'_0, w, z)$. Suppose that $T_{1,n}, T'_{1,n}, T_{2,n}$, and $T'_{2,n}$ are sequences of neck lengths all approaching $+\infty$ and let $\widehat{\mathcal{J}}_n$ denote an interpolating almost complex structures between $\mathcal{J}(T_{1,n}, T'_{1,n})$ and $\mathcal{J}(T_{2,n}, T'_{2,n})$. Pick $\widehat{\mathcal{J}}_n$ so that as $n \to \infty$ the almost complex structures $\widehat{\mathcal{J}}_n$ split into $J_s \vee J_{S^2} \vee J_{S^2}$ on $(\Sigma \vee S^2 \vee S^2) \times [0,1] \times \mathbb{R}$. If u_n is a sequence of Maslov index 0, $\widehat{\mathcal{J}}_n$–holomorphic curves representing ϕ, we can extract a weak limit to broken curves U_Σ, U_0 and U'_0 on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s, \alpha'_s\}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$, (S^2, α_0, β_0) and $(S^2, \alpha'_0, \beta'_0)$ representing ϕ_Σ, ϕ_0 and ϕ'_0 respectively. As in Proposition 3.5.3, the curves in U_Σ consist of a broken holomorphic strip U'_Σ on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, as well as a collection $\mathcal{A}$ of cylindrical $(\boldsymbol{\alpha} \cup \{\alpha_s\} \cup \{\alpha'_s\})$–boundary degenerations. Let ϕ'_Σ denote the underlying homology class of U'_Σ. Let $m_1, m_2, n_1, n_2, m'_1, m'_2, n'_1$, and n'_2 be multiplicities as in Figure 3.10.

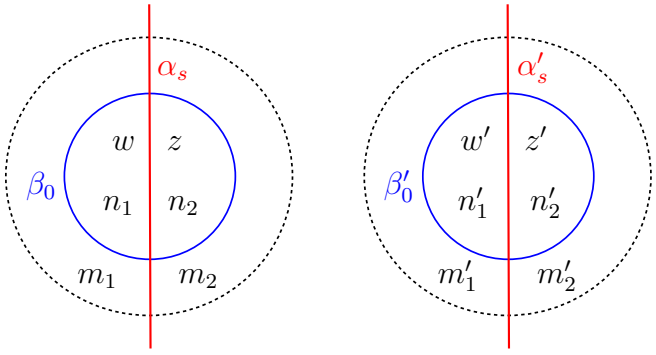

Figure 3.10: **Multiplicities of a disk ϕ near new basepoints w, z, w' and z' on a diagram which has been quasi-stabilized twice.**

Adapting the Maslov index computation from Proposition 3.5.3, we see that

$$\mu(\phi) = \mu(\phi'_\Sigma) + n_1(\phi) + n_2(\phi) + n'_1(\phi) + n'_2(\phi)$$

$$+m_1(\mathcal{A}) + m_2(\mathcal{A}) + m'_1(\mathcal{A}) + m'_2(\mathcal{A}) + 2 \sum_{\mathcal{D} \in C(\Sigma \setminus \boldsymbol{\alpha}), \mathcal{D} \neq A} n_{\mathcal{D}}(\mathcal{A})$$

where $C(\Sigma \setminus \boldsymbol{\alpha})$ denotes the connected components of $\Sigma \setminus \boldsymbol{\alpha}$.

Since U'_Σ is a broken holomorphic curve for an $\mathbb{R}$–invariant almost complex structure, we conclude that U'_Σ consists only of constant flowlines. Since all of the other summands are zero, it's easy to see that this forces all multiplicities to be zero. Hence only constant disks are counted by the change of almost complex structures map, concluding the proof of the lemma. □

In the case that α_s and α'_s are in the same component, the change of almost complex structure maps will often be nontrivial. Nevertheless we have the following:

Lemma 3.8.3. *Suppose that the pairs (w', z') and (w, z) are adjacent on $(L, \mathbf{w} \cup \{w, w'\}, \mathbf{z} \cup \{z, z'\})$ and that (w', z') comes after (w, z). Let $\theta^\pm$ and $(\theta')^\pm$ denote the intersection points corresponding to quasi-stabilization. If T_1, T'_1, T_2, T'_2 are all sufficiently large, then writing $F = \Phi_{\mathcal{J}(T_1, T'_1) \to \mathcal{J}(T_2, T'_2)}$, we have*

$$\begin{aligned}
F(\mathbf{x} \times \theta^+ \times (\theta')^+) &= \mathbf{x} \times \theta^+ \times (\theta')^+ \\
F(\mathbf{x} \times \theta^+ \times (\theta')^-) &= \mathbf{x} \times \theta^+ \times (\theta')^- + C \cdot \mathbf{x} \times \theta^- \times (\theta')^+ \\
F(\mathbf{x} \times \theta^- \times (\theta')^+) &= \mathbf{x} \times \theta^- \times (\theta')^+ \\
F(\mathbf{x} \times \theta^- \times (\theta')^-) &= \mathbf{x} \times \theta^- \times (\theta')^-,
\end{aligned}$$

for some C (which is not independent of T_i and T'_i).

Proof. We proceed similar to the previous lemma. Now a single component, which we denote by A, of $\Sigma \setminus \boldsymbol{\alpha}$ contains both α_s and α'_s. Write A_1, A_2 and A_3 for the three different components of $A \setminus (\alpha_s \cup \alpha'_s)$. Two of the A_i share boundary with exactly one of the other

A_j, and one of the A_i shares boundary with both of the other A_i. Without loss of generality assume that A_1 shares boundary with A_2, and that A_2 also shares boundary with A_3.

As before, as we simultaneously stretch the necks, a sequence of Maslov index zero disks u_i has a weak limit as before. Now, however, the Maslov index computation is different. Let $a_i(\mathcal{A})$ denote the multiplicity of the $(\boldsymbol{\alpha} \cup \{\alpha_s\} \cup \{\alpha_s'\})$–degeneration $\mathcal{A}$ in the region A_i. One computes that the Maslov index now satisfies

$$\mu(\phi) = \mu(\phi_\Sigma') + n_1(\phi) + n_2(\phi) + n_1'(\phi) + n_2'(\phi) + a_1(\mathcal{A}) + a_3(\mathcal{A}) + 2 \sum_{\mathcal{D} \in C(\Sigma \setminus \boldsymbol{\alpha}), \mathcal{D} \neq A} n_{\mathcal{D}}(\mathcal{A}).$$

As usual this implies that all of the above terms are zero. Hence ϕ_Σ', which has a broken representative for a cylindrical almost complex structure, must be the constant disk by transversality. The only multiplicities which may be nonzero are $a_2(\mathcal{A})$, and the multiplicities $m_i(\phi)$ and $m_i'(\phi)$, none of which appear in the above sum. As is easily observed, this constrains the disk ϕ to be in $\pi_2(\mathbf{x} \times \theta^+ \times (\theta')^-, \mathbf{x} \times \theta^- \times (\theta')^+)$, completing the proof. An example of a disk which might appear in the change of almost complex structure map is shown in Figure 3.11. □

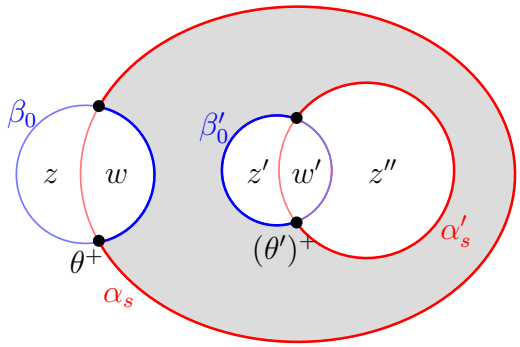

Figure 3.11: **An example of a Maslov index zero disk which might be counted by** $\Phi_{\mathcal{J}(T_1,T_1')\to\mathcal{J}(T_2,T_2')}$ in Lemma 3.8.3 for arbitrarily large T_i and T_i'.

Proof of Theorem 3.8.1. The proof is easy algebra in all cases using Lemmas 3.8.2 and 3.8.3. □

3.9 Further relations between the maps Ψ_z, Φ_w and $S^{\pm}_{w,z}$

In this section we prove several relations between the maps $S^{\pm}_{w,z}$, Ψ_z and Φ_w. We highlight the convenience of viewing Φ_w and Ψ_z as formal derivatives of the differential, since basically all of the relations in this section are derived by either formally differentiating the expression for $\partial\circ\partial$ from Lemma 3.2.1, or by differentiating our expression of the quasi-stabilized differential in Proposition 3.5.3.

Lemma 3.9.1. *If w and z are not adjacent, or if w and z are the only basepoints on a link component, then*

$$\Phi_w\Psi_z + \Psi_z\Phi_w \simeq 0.$$

If w and z are adjacent and there are other basepoints on the link component, then

$$\Phi_w\Psi_z + \Psi_z\Phi_w \simeq \mathrm{id}\,.$$

Proof. Take the expression for ∂^2 from Lemma 3.2.1 and differentiate it with respect to U_w. We obtain

$$\partial\Phi_w + \Phi_w\partial = V_{z'} + V_{z''}$$

where z' and z'' are the variables adjacent to w on its link component. Suppose first that $z' \neq z''$, i.e. that w and z are not the only basepoints on their link component.

Differentiating the above expression with respect to V_z, we see that

$$\Psi_z\Phi_w + \Phi_w\Psi_z \simeq \begin{cases} \mathrm{id} & \text{if } w \text{ is adjacent to } z \\ 0 & \text{if } w \text{ is not adjacent to } z. \end{cases},$$

from which the claim follows as long as w and z are not the only basepoints on their link component.

If w and z are the only basepoints on their link component, then $z = z' = z''$ and the above argument is easily modified to give the stated result. □

Lemma 3.9.2. *We have*

$$\Psi_z \Psi_{z'} + \Psi_{z'} \Psi_z \simeq 0$$

and

$$\Phi_w \Phi_{w'} + \Phi_{w'} \Phi_w \simeq 0$$

for any choice of z, z', w, and w'.

Proof. This is proven identically to the previous lemma. □

As with the free stabilization maps in Chapter 2, we have the following:

Lemma 3.9.3. *The following relation holds:*

$$S^+_{w,z} S^-_{w,z} = \Phi_w.$$

Proof. The differential on the uncolored quasi-stabilized diagram takes the form

$$\partial_{\overline{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H}} \end{pmatrix}$$

where $\partial_{\mathcal{H}}$ is the differential on the unstabilized diagram. After taking the U_w derivative we get

$$\Phi_w = \begin{pmatrix} 0 & \mathrm{id} \\ 0 & 0 \end{pmatrix},$$

which is exactly $S^+_{w,z} S^-_{w,z}$. □

We now consider commutators of the quasi-stabilization maps and the maps Φ_w and Ψ_z.

Lemma 3.9.4. *Suppose that w and z are two new basepoints for a link. If w is not adjacent to z', then*

$$S^\pm_{w,z} \Psi_{z'} + \Psi_{z'} S^\pm_{w,z} \simeq 0.$$

With no assumptions on adjacency, we have

$$S^\pm_{w,z} \Phi_{w'} + \Phi_{w'} S^\pm_{w,z} \simeq 0.$$

Proof. Suppose that w and z are inserted between basepoints z'' and w'' on the link $\mathbb{L}$. The quasi-stabilized differential takes the form

$$\partial_{\overline{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w''} \\ V_z + V_{z''} & \partial_{\mathcal{H}} \end{pmatrix},$$

by Proposition 3.5.3. By assumption $z' \neq z''$. Differentiating with respect to $V_{z'}$ thus yields

$$\widetilde{\Psi}_{z'} = \begin{pmatrix} \Psi_{z'} & 0 \\ 0 & \Psi_{z'} \end{pmatrix},$$

where $\widetilde{\Psi}_{z'}$ denotes the map on the stabilized diagram, and $\Psi_{z'}$ denotes the map on the unstabilized diagram. In matrix notation the maps $S^{\pm}_{w,z}$ take the form

$$S^{+}_{w,z} = \begin{pmatrix} \mathrm{id} \\ 0 \end{pmatrix} \qquad \text{and} \qquad S^{-}_{w,z} = \begin{pmatrix} 0 & \mathrm{id} \end{pmatrix}. \tag{3.6}$$

The stated equality involving $\Psi_{z'}$ now follows from matrix multiplication.

The equality involving $\Phi_{w'}$ follows similarly. □

We also have the following:

Lemma 3.9.5. *Suppose that z' is adjacent to w and that $z' \neq z$. Then we have*

$$S^{+}_{w,z}\Psi_{z'} \simeq (\Psi_{z'} + \Psi_z)S^{+}_{w,z}$$

and

$$\Psi_{z'}S^{-}_{w,z} \simeq S^{-}_{w,z}(\Psi_{z'} + \Psi_z).$$

Proof. Once again we consider the quasi-stabilized differential, which is

$$\partial_{\overline{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H}} \end{pmatrix}.$$

Differentiating with respect to z' yields

$$\widetilde{\Psi}_{z'} = \begin{pmatrix} \Psi_{z'} & 0 \\ \mathrm{id} & \Psi_{z'} \end{pmatrix}$$

and

$$\widetilde{\Psi}_z = \begin{pmatrix} 0 & 0 \\ \mathrm{id} & 0 \end{pmatrix}.$$

Here $\widetilde{\Psi}_z$ denotes the map on the complex after quasi-stabilization, and Ψ_z denotes the map on the complex before quasi-stabilization. Using the matrix notation from Equation (3.6), the desired relations follow from matrix multiplication. □

The reader can compare the following lemma to Lemma 2.6.7, the analogous result for the closed three manifold invariants.

Lemma 3.9.6. *Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link in Y^3, such that w and z are two new, consecutive basepoints on $\mathbb{L}$ such that w follows z. If z' is one of the two $\mathbf{z}$–basepoints adjacent to w, then we have*

$$S^-_{w,z} \Psi_{z'} S^+_{w,z} \simeq \mathrm{id}\,.$$

Proof. This follows from our usual strategy. Pick a diagram $\mathcal{H}$ for $(L, \mathbf{w}, \mathbf{z})$ and let $\overline{\mathcal{H}}$ denote a diagram which has been quasi-stabilized at w and z. Let z'' and w'' denote the basepoints adjacent to w and z on $\mathbb{L}$. Using Proposition 3.5.3 we have that

$$\partial_{\overline{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w''} \\ V_z + V_{z''} & \partial_{\mathcal{H}} \end{pmatrix}.$$

By assumption, either $z' = z$ or $z' = z''$ (but not both). In both cases, we have that

$$\Psi_{z'} = \left(\frac{d}{dV_{z'}} \partial_{\overline{\mathcal{H}}} \right) = \begin{pmatrix} * & * \\ \mathrm{id} & * \end{pmatrix},$$

where the $*$ terms are unimportant. Using the matrix notation from Equation (3.6), we get the desired equality immediately from matrix multiplication. □

The reader should compare the following to [Sar15, Lem. 4.4.].

Lemma 3.9.7. *We have $\Psi_z^2 \simeq 0$ and $\Phi_w^2 \simeq 0$, as $\mathfrak{P}$–filtered maps of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–modules.*

Proof. The proof follows identically to the proof of Lemma 2.13.19. □

3.10 Basepoint moving maps

In this section, we compute several basepoint moving maps. The procedure for computing maps induced by moving basepoints is in a similar spirit to the author's computation of the π_1–action on the Heegaard Floer homology of a closed three manifold in Chapter 2. We first compute the effect of moving basepoints along a small arc on a link component via a model computation. We then use this to prove Theorem B, the effect of moving all of the basepoints on a link component in one full loop. The final computation is Theorem D, where we compute the effect on certain colored complexes of moving each $\mathbf{w}$–basepoint to the next $\mathbf{w}$–basepoint, and moving each $\mathbf{z}$–basepoint to the next $\mathbf{z}$–basepoint.

3.10.1 Moving basepoints along an arc

Suppose we Y^3 is a three manifold with embedded multipointed link $\mathbb{L}_0 = (L, \mathbf{w}_0, \mathbf{z}_0)$, though we allow the case that one of the components of $\mathbb{L}_0$ has no basepoints. Suppose that z, w, z', w' are all points on a single component of $L \setminus (\mathbf{w}_0 \cup \mathbf{z}_0)$, appearing in that order according to the orientation of $\mathbb{L}$. Let

$$\mathbf{w} = \mathbf{w}_0 \cup \{w\}, \qquad \mathbf{z} = \mathbf{z}_0 \cup \{z\}, \qquad \mathbf{w}' = \mathbf{w}_0 \cup \{w'\}, \qquad \text{and} \qquad \mathbf{z}' = \mathbf{z}_0 \cup \{z'\}.$$

Finally assume that $(L, \mathbf{w}, \mathbf{z})$ has basepoints in each component of L. There is an isotopically unique diffeomorphism of Y which maps L to itself and fixes $\mathbf{w}_0 \cup \mathbf{z}_0$ and maps w to w' and z to z', which is isotopic to the identity relative $\mathbf{w}_0 \cup \mathbf{z}_0$ through isotopies which map L to itself. Let ς_0 denote this diffeomorphism. The diffeomorphism ς_0 induces a map

$$(\varsigma_0)_*\colon CFL^\infty_{UV}(Y, L, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to CFL^\infty_{UV}(Y, L, \mathbf{w}', \mathbf{z}', \mathfrak{s}).$$

The map $(\varsigma_0)_*$ is defined as a tautology. That is, if $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, L, \mathbf{w}, \mathbf{z})$ with almost complex structure J_s, we just apply the map ς_0 to the diagram $\mathcal{H}$ to get a new diagram

$$\varsigma_0(\mathcal{H}) = (\varsigma_0(\Sigma), \varsigma_0(\boldsymbol{\alpha}), \varsigma_0(\boldsymbol{\beta}), \varsigma_0(\mathbf{w}), \varsigma_0(\mathbf{z})),$$

with almost complex structure $(\varsigma_0)_* J_s$. The diffeomorphism ς_0 tautologically determines a chain map

$$(\varsigma_0)_* \colon CFL^\infty_{UV,J_s}(\mathcal{H},\mathfrak{s}) \to CFL^\infty_{UV,\varsigma_0(J_s)}(\varsigma_0(\mathcal{H}),\mathfrak{s})$$

defined by

$$(\varsigma_0)_*(\mathbf{x}) = \varsigma_0(\mathbf{x}).$$

By the naturality results of [JT12], this yields a well defined morphism on the coherent chain homotopy type invariants (i.e. it commutes with change of diagrams maps, in the appropriate sense).

Note that $(\varsigma_0)_*$ "appears" like the identity map since it just maps an intersection point $\mathbf{x}$ to it's image under the diffeomorphism ς_0. With this in mind, we prove the following:

Lemma 3.10.1. *The induced map $(\varsigma_0)_*$ is filtered chain homotopic to*

$$(\varsigma_0)_* \simeq S^-_{w,z}\Psi_{z'}S^+_{w',z'}.$$

Proof. We first prove the result in the case that the link component containing w and z has at least one extra pair of basepoints. Let z'' denote the basepoint occurring immediately after w'. In this case, we can pick a diagram like the one shown in Figure 3.12, where the dashed lines show two circles along which the almost complex structure will be stretched. In this diagram, we assume that α_s and α'_s each bound disks on Σ and that $\alpha_s, \alpha'_s, \beta_0$ and β'_0 do not intersect any other $\boldsymbol{\alpha}$- or $\boldsymbol{\beta}$–curves. With this diagram, we can compute all of the maps $\Psi_{z'}, S^-_{w,z}$, and $S^+_{w',z'}$ explicitly. We must be careful though, since we cannot use the same almost complex structure for all of the maps. Instead we will need to use the change of almost complex structure computation from Lemma 3.8.3. Let J_s be an almost complex structure which is sufficiently stretched along c to compute compute $S^\pm_{w,z}$, and let J'_s be an almost complex structure which is sufficiently stretched along c' to compute $S^\pm_{w',z'}$, and assume that both are stretched sufficiently so that the change of almost complex structure map $\Phi_{J'_s\to J_s}$ takes the form described in Lemma 3.8.3. We wish to compute $S^-_{w,z}\circ\Psi_{z'}\circ\Phi_{J'_s\to J_s}\circ S^+_{w',z'}$.

Let $\theta^\pm$ denote the intersection points of $\alpha_s\cap\beta_0$ and let $(\theta')^\pm$ denote the intersection points of α'_s and β'_0.

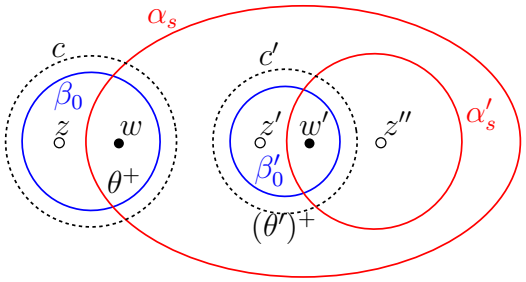

Figure 3.12: **A diagram for Lemma 3.10.1 when we have another basepoint z'' on the link component containing w, z, w' and z'.** The curves α_s and α_s' each bound disks, and $\alpha_s, \alpha_s', \beta_0, \beta_0'$ do not intersect any other $\boldsymbol{\alpha}$– or $\boldsymbol{\beta}$–curves.

Using the analysis in Proposition 3.5.3, we see that for J_s there are exactly two domains which are the domain of Maslov index one disks ϕ which support holomorphic representatives with $n_{z'}(\phi) > 0$. These domains are shown in Figure 3.13. Also every homology disk ϕ which has one of these domains has $\#\widehat{\mathcal{M}}_{J_s}(\phi) = 1$.

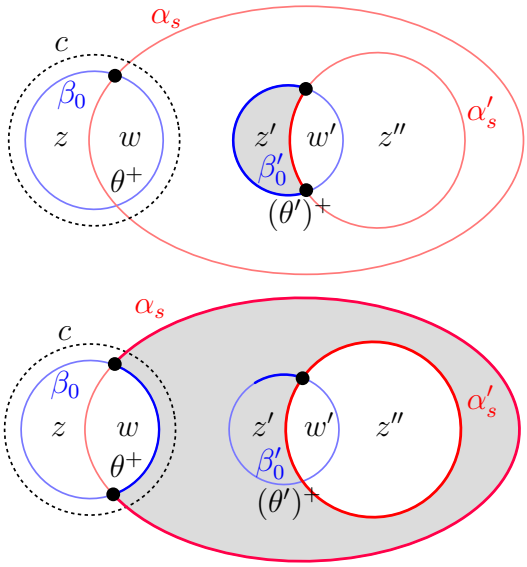

Figure 3.13: **The two domains contributing to $\Psi_{z'}$ in Lemma 3.10.1 for the almost complex structure J_s,** which is stretched first on c', and then on c (possibly much more than on c'). Also drawn in are two examples of holomorphic disks with those domains.

We wish to show that $S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J_s' \to J_s} \circ S^+_{w',z'} = (\varsigma_0)_*$, where ς_0 is the diffeomorphism induced by simply pushing w and z to w' and z' respectively. To this end, it is sufficient to

show that

$$(S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s} \circ S^+_{w',z'})(\mathbf{x} \times \theta^{\pm}) = \mathbf{x} \times (\theta')^{\pm},$$

since $(\varsigma_0)_*(\mathbf{x} \times \theta^{\pm}) = \mathbf{x} \times (\theta')^{\pm}$.

Homology disks with the left domain in Figure 3.13 yield a contribution of

$$\mathbf{x} \times \theta^{\pm} \times (\theta')^{+} \longrightarrow \mathbf{x} \times \theta^{\pm} \times (\theta')^{-}$$

to $\Psi_{z'}$.

Homology disks with the right domain in Figure 3.13 yield a contribution of

$$\mathbf{x} \times \theta^{+} \times (\theta')^{\pm} \longrightarrow \mathbf{x} \times \theta^{-} \times (\theta')^{\pm}$$

to $\Psi_{z'}$.

We first compute $(S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s} \circ S^+_{w',z'})(\mathbf{x} \times \theta^+)$. Using the above computation of $\Psi_{z'}$ and the computation of $\Phi_{J'_s \to J_s}$ from Lemma 3.8.3, we have that

$$\begin{aligned}
(S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s} \circ S^+_{w',z'})(\mathbf{x} \times \theta^+) &= (S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s})(\mathbf{x} \times \theta^+ \times (\theta')^+) \\
&= (S^-_{w,z} \circ \Psi_{z'})(\mathbf{x} \times \theta^+ \times (\theta')^+) \\
&= S^-_{w,z}(\mathbf{x} \times \theta^+ \times (\theta')^- + \mathbf{x} \times \theta^- \times (\theta')^+) \\
&= \mathbf{x} \times (\theta')^+
\end{aligned}$$

We now compute $(S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s} \circ S^+_{w',z'})(\mathbf{x} \times \theta^-)$. Once again using our previous computation of $\Psi_{z'}$ and Lemma 3.8.3, we have that

$$\begin{aligned}
(S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s} \circ S^+_{w',z'})(\mathbf{x} \times \theta^-) &= (S^-_{w,z} \circ \Psi_{z'} \circ \Phi_{J'_s \to J_s})(\mathbf{x} \times \theta^- \times (\theta')^+) \\
&= (S^-_{w,z} \circ \Psi_{z'})(\mathbf{x} \times \theta^- \times (\theta')^+) \\
&= S^-_{w,z}(\mathbf{x} \times \theta^- \times (\theta')^-) \\
&= \mathbf{x} \times (\theta')^-,
\end{aligned}$$

completing the proof of the claim if $\mathbf{w}$ and $\mathbf{z}$ each have at least two basepoints on L.

We now consider the case that $\mathbb{L}$ doesn't have any basepoints other than w and z. In this case we just introduce two new basepoints w'', z'' which are on the component of $L \setminus \{w, w', z, z'\}$ which is going from w' to z. Note that ς_0 is isotopic relative $\{w, z\}$ to a diffeomorphism which fixes w'' and z''. Hence $(\varsigma_0)_* S^-_{w'',z''} = S^-_{w'',z''}(\varsigma_0)_*$. We just compute that

$$\begin{aligned}
(\varsigma_0)_* &\simeq (\varsigma_0)_*(S^-_{w'',z''}\Psi_{z''}S^+_{w'',z''}) && \text{(Lemma 3.9.6)}\\
&\simeq S^-_{w'',z''}(\varsigma_0)_*\Psi_{z''}S^+_{w'',z''} && \text{(observation above)}\\
&\simeq S^-_{w'',z''}(S^-_{w,z}\Psi_{z'}S^+_{w',z'})\Psi_{z''}S^+_{w'',z''} && \text{(previous case)}\\
&\simeq (S^-_{w,z}\Psi_{z'}S^+_{w',z'})(S^-_{w'',z''}\Psi_{z''}S^+_{w'',z''}) && \text{(Theorem 3.8.1, Lemma 3.9.4)}\\
&\simeq S^-_{w,z}\Psi_{z'}S^+_{w',z'}, && \text{(Lemma 3.9.6)}
\end{aligned}$$

as we wanted.

□

3.10.2 Sarkar's formula for moving basepoints in a full twist around a link component

In this section, we prove Theorem B, which is Sarkar's conjectured formula for the effect of moving basepoints on a link component in a full twist around the link component for the full link Floer complex. The main technical tool is Lemma 3.10.1, which computes the effect of moving basepoints on a small arc on a link component. By writing the diffeomorphism of a full twist as a composition of many smaller moves of the previous form, we will obtain Sarkar's formula.

Theorem B. *Suppose ς is the diffeomorphism corresponding to a positive Dehn twist around a link component K of $\mathbb{L}$. Suppose that the basepoints on K are $w_1, z_1, \dots, w_n, z_n$. The induced map ς_* on $CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ has the $\mathfrak{P}$–filtered $\mathbb{Z}_2[U_\mathfrak{P}]$ chain homotopy type*

$$\varsigma_* \simeq \mathrm{id} + \Phi_K\Psi_K$$

where

$$\Phi_K = \sum_{j=1}^{n} \Phi_{w_j} \qquad \textit{and} \qquad \Psi_K = \sum_{j=1}^{n} \Psi_{z_j}.$$

To the reader who is not interested in colorings, we note that one can just take $\mathfrak{P} = \mathbf{w} \cup |L|$, where $|L|$ denotes the components of L.

In this section, we also introduce some new formalism to make the computation easier. The maps $\Psi_{z'}$ and $S^{\pm}_{w,z}$ interact strangely (e.g. Lemma 3.9.5), which leads to challenging and messy algebra if we are not careful. Suppose that A is an arc on L between two $\mathbf{w}$–basepoints which share the same color. We define the map

$$\Psi_A = \sum_{z \in A \cap \mathbf{z}} \Psi_z.$$

The maps Ψ_A can be thought of as defining an action of

$$\Lambda^* H_1(L/(\mathbf{w}, \sigma); \mathbb{Z})$$

on $CFL^{\infty}_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$, where $L/(\mathbf{w}, \sigma)$ denotes the space obtained by identifying two $\mathbf{w}$–basepoints if they share the same color. This formalism is intriguing, but we will only have use for maps Ψ_A for arcs A between $\mathbf{w}$–basepoints of the same color.

Given an arc A between two $\mathbf{w}$–basepoints, we define an **endpoint** of A to be a basepoint w such that the sets $\overline{K \setminus A}$ and $\overline{A}$ both contain w (so if $A = K$, A has no endpoints).

We now proceed to prove some basic properties of the maps Ψ_A, all of which are recastings of previous lemmas proven about the maps Ψ_z.

Lemma 3.10.2. *We have*

$$S^{\pm}_{w,z} \Psi_A + \Psi_A S^{\pm}_{w,z} \simeq 0,$$

as long as w is not an endpoint of A.

Proof. This follows immediately from Lemmas 3.9.4 and 3.9.5. □

Lemma 3.10.3. *If A and A' are two arcs between $\mathbf{w}$–basepoints, then*

$$\Psi_A \Psi_{A'} + \Psi_{A'} \Psi_A \simeq 0.$$

Proof. This follows from Lemma 3.9.2. □

Lemma 3.10.4. *If A is an arc on L, then we have*

$$\Psi_A^2 \simeq 0,$$

as filtered equivariant maps.

Proof. Simply write $\Psi_A = \sum_{z \in A \cap \mathbf{z}} \Psi_z$, multiply out Ψ_A^2, then apply Lemmas 3.9.2 and 3.9.7. □

Lemma 3.10.5. *Suppose $A \subseteq K$ is an arc between $\mathbf{w}$–basepoints and let $c(A)$ denote the arc $\overline{K \setminus A}$. Then*

$$\Psi_K \Psi_A = \Psi_{c(A)} \Psi_A.$$

Proof. Write $\Psi_K = \Psi_A + \Psi_{c(A)}$ and then use the previous lemma to compute that

$$\Psi_K \Psi_A = (\Psi_A + \Psi_{c(A)})\Psi_A = \Psi_A^2 + \Psi_{c(A)}\Psi_A = \Psi_{c(A)}\Psi_A.$$

□

Lemma 3.10.6. *If w is an endpoint of A then we have*

$$\Psi_A \Phi_w + \Phi_w \Psi_A \simeq \mathrm{id}\,.$$

If w is not an endpoint of A, then we have

$$\Psi_A \Phi_w + \Phi_w \Psi_A \simeq 0.$$

Proof. The first claim follows from Lemma 3.9.1. The second claim follows from Lemma 3.10.2 since we can always write $\Phi_w = S^+_{w,z} S^-_{w,z}$ for z the basepoint immediately preceding w on $\mathbb{L}$. □

We can now prove Theorem B:

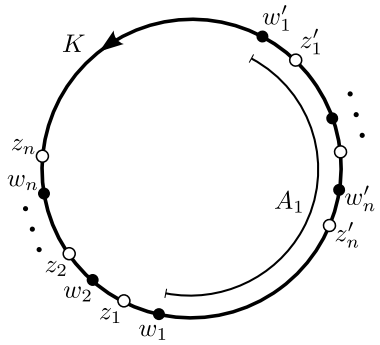

Figure 3.14: **The basepoints $z_1, w_1, \ldots, z_n, w_n$ and $z_1', w_1', \ldots z_n', w_n'$, as well as the arcs A_i.**

Proof of Theorem B. Let $w_1, z_1, \ldots w_n, z_n$ be the basepoints on K, in the reverse order that they appear on K according to the orientation of K. Let $w_1' z_1', \ldots w_n', z_n'$ be new basepoints on K in the interval between z_n and w_1. Let A_j be the arc on K from w_j to w_j', as in Figure 3.14.

Write

$$\mathbf{w} = \{w_1, \ldots, w_n\}, \qquad \mathbf{z} = \{z_1, \ldots, z_n\}, \qquad \mathbf{w}' = \{w_1', \ldots, w_n'\}, \quad \text{and} \quad \mathbf{z}' = \{z_1', \ldots, z_n'\}.$$

As usual, we write ς as a composition of two diffeomorphisms $\varsigma = \varsigma_2 \circ \varsigma_1$, where ς_1 moves the basepoints $\mathbf{w}$ and $\mathbf{z}$ to $\mathbf{w}'$ and $\mathbf{z}'$ respectively and ς_2 moves the basepoints $\mathbf{w}'$ and $\mathbf{z}'$ to $\mathbf{w}$ and $\mathbf{z}$ respectively. Let $c(A_i) = \overline{K \setminus A_i}$. Note that by Lemma 3.9.6 we have

$$\prod_{j=1}^{n} (S^-_{w_j', z_j'} \Psi_{A_j} S^+_{w_j', z_j'}) \simeq \mathrm{id}\,. \tag{3.7}$$

Write $S^{\pm}_{\mathbf{w},\mathbf{z}}$ for $\prod_{j=1}^{n} S^{\pm}_{w_j,z_j}$ and similarly for $S^{\pm}_{\mathbf{w'},\mathbf{z'}}$. We compute as follows:

$$\begin{aligned}
\varsigma_* &= (\varsigma_2)_* \circ (\varsigma_1)_* \\
&= \left(\prod_{j=1}^{n} S^{-}_{w'_j,z'_j} \Psi_{c(A_j)} S^{+}_{w_j,z_j}\right)\left(\prod_{j=1}^{n} S^{-}_{w_j,z_j} \Psi_{A_j} S^{+}_{w'_j,z'_j}\right) && \text{(Lemma 3.10.1)} \\
&= S^{-}_{\mathbf{w'},\mathbf{z'}}\left(\prod_{j=1}^{n} \Psi_{c(A_j)}\right) S^{+}_{\mathbf{w},\mathbf{z}} S^{-}_{\mathbf{w},\mathbf{z}}\left(\prod_{j=1}^{n} \Psi_{A_j}\right) S^{+}_{\mathbf{w'},\mathbf{z'}} && \text{(Lemmas 3.10.2, 3.8.1)} \\
&= S^{-}_{\mathbf{w'},\mathbf{z'}}\left(\prod_{j=1}^{n} \Psi_{c(A_j)}\right)\left(\prod_{j=1}^{n} \Phi_{w_j}\right)\left(\prod_{j=1}^{n} \Psi_{A_j}\right) S^{+}_{\mathbf{w'},\mathbf{z'}} && \text{(Lemmas 3.9.3, 3.8.1)} \\
&= S^{-}_{\mathbf{w'},\mathbf{z'}} \sum_{s\in\{0,1\}^n}\left(\prod_{j=1}^{n} \Phi^{s_j}_{w_j}\right)\left(\prod_{j=1}^{n} \Psi^{s_j}_{c(A_j)}\right)\left(\prod_{j=1}^{n} \Psi_{A_j}\right) S^{+}_{\mathbf{w'},\mathbf{z'}} && \text{(Lemma 3.10.6)} \\
&= S^{-}_{\mathbf{w'},\mathbf{z'}} \sum_{s\in\{0,1\}^n}\left(\prod_{j=1}^{n} \Phi^{s_j}_{w_j}\right)\left(\prod_{j=1}^{n} \Psi^{s_j}_{K}\right)\left(\prod_{j=1}^{n} \Psi_{A_j}\right) S^{+}_{\mathbf{w'},\mathbf{z'}} && \text{(Lemmas 3.10.3, 3.10.5)} \\
&= \sum_{s\in\{0,1\}^n}\left(\prod_{j=1}^{n} \Phi^{s_j}_{w_j}\right)\left(\prod_{j=1}^{n} \Psi^{s_j}_{K}\right) S^{-}_{\mathbf{w'},\mathbf{z'}}\left(\prod_{j=1}^{n} \Psi_{A_j}\right) S^{+}_{\mathbf{w'},\mathbf{z'}} && \text{(Lemmas 3.10.2, 3.9.4)} \\
&= \sum_{s\in\{0,1\}^n}\left(\prod_{j=1}^{n} \Phi^{s_j}_{w_j}\right)\left(\prod_{j=1}^{n} \Psi^{s_j}_{K}\right) && \text{(Equation (3.7)).}
\end{aligned}$$

We now note that by Lemma 3.10.4, if $s \in \{0,1\}^n$ then

$$\left(\prod_{j=1}^{n} \Psi^{s_j}_{K}\right) \simeq 0$$

if s_j is nonzero for more than one j. Hence the above sum reduces to

$$\varsigma_* \simeq \mathrm{id} + \sum_{j=1}^{n} \Phi_{w_i}\Psi_K = \mathrm{id} + \Phi_K\Psi_K,$$

completing the proof. □

3.10.3 The map associated to a partial twist around a link component

In this section, we perform an additional basepoint moving map computation and prove Theorem D. Suppose that $\mathbb{L}$ is a multibased link and K is a component with basepoints $z_1, w_1, z_2, w_2, \ldots, z_n$, and w_n, appearing in that order. Let τ be the diffeomorphism induced

by twisting $(\frac{1}{n})^{\text{th}}$ of the way around K, sending z_i to z_{i+1} and w_i to w_{i+1} (with indices taken modulo n). In the case that we pick a coloring $(\sigma, \mathfrak{P})$ where all of the $\mathbf{w}$–basepoints have the same color, the map τ induces a map on the complex

$$CFL^{\infty}_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}).$$

We have the following:

Theorem D. *Suppose that $\mathbb{L}$ is an embedded multibased link in Y, and K is a component of $\mathbb{L}$ with basepoints $z_1, w_1, \dots, z_n$, and w_n, appearing in that order. Assume that $n > 1$. If τ denotes the $(\frac{1}{n})^{th}$–twist map, then for a coloring where all $\mathbf{w}$–basepoints on K have the same color, we have*

$$\tau_* \simeq (\Psi_{z_1}\Phi_{w_1}\Psi_{z_2}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{z_n}\Phi_{w_n}) + (\Phi_{w_1}\Psi_{z_2}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{z_n}).$$

Proof. Let A_i be the arc from w_i to w_{i+1}, respecting the orientation of K. Let w' and z' be new basepoints in the region between z_n and w_1. Let A' denote the arc from w_n to w' and let A'' denote the arc from w' to w_1. This is illustrated in Figure 3.15.

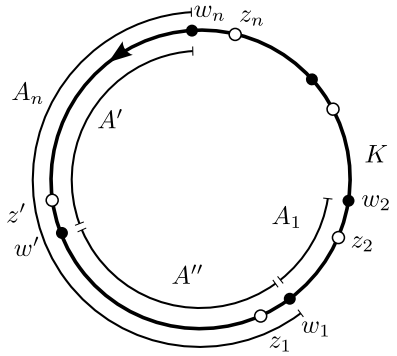

Figure 3.15: **The basepoints $z_1, w_1, \dots, z_n, w_n, z', w'$ as well as the arcs A_i, A' and A'' from Theorem D.**

Using Lemma 3.10.1 repeatedly, we have that

$$\tau_* \simeq (S^-_{w'}\Psi_{A''}S^+_{w_1})(S^-_{w_1}\Psi_{A_1}S^+_{w_2})\cdots(S^-_{w_{n-1}}\Psi_{A_{n-1}}S^+_{w_n})(S^-_{w_n}\Psi_{A'}S^+_{w'}).$$

Using this, we perform the following computation:

$$
\begin{aligned}
\tau_* &\simeq S^-_{w'}\Psi_{A''}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\Phi_{w_n}\Psi_{A'}S^+_{w'} && \text{(Lemma 3.9.3)}\\
&\simeq S^-_{w'}\Psi_{A''}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}(\Psi_{A'}\Phi_{w_n}+1)S^+_{w'} && \text{(Lemma 3.10.6)}\\
&\simeq S^-_{w'}\Psi_{A''}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\Psi_{A'}\Phi_{w_n}S^+_{w'}\\
&\quad + S^-_{w'}\Psi_{A''}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}S^+_{w'}\\
&\simeq S^-_{w'}(\Psi_{A''}\Psi_{A'})\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\Phi_{w_n}S^+_{w'} && \text{(Lemmas 3.10.6, 3.10.3)}\\
&\quad + S^-_{w'}\Psi_{A''}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}S^+_{w'}\\
&\simeq S^-_{w'}(\Psi_{A''}\Psi_{A_n})\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\Phi_{w_n}S^+_{w'} && \text{(Lemma 3.10.4)}\\
&\quad + S^-_{w'}\Psi_{A''}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}S^+_{w'}\\
&\simeq (S^-_{w'}\Psi_{A''}S^+_{w'})\Psi_{A_n}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\Phi_{w_n} && \text{(Lemma 3.10.2)}\\
&\quad + (S^-_{w'}\Psi_{A''}S^+_{w'})\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\\
&\simeq \Psi_{A_n}\Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}\Phi_{w_n} && \text{(Lemma 3.9.6)}\\
&\quad + \Phi_{w_1}\Psi_{A_1}\Phi_{w_2}\cdots\Phi_{w_{n-1}}\Psi_{A_{n-1}}
\end{aligned}
$$

completing the proof since by definition $\Psi_{A_i} = \Psi_{z_{i+1}}$ on the complex $CFL^\infty_{UV}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$.

□

CHAPTER 4

Link cobordisms and functoriality in link Floer homology

We construct cobordism maps associated to decorated link cobordisms for link Floer homology. The maps are associated to oriented surfaces with divides embedded in 4-dimensional cobordisms between 3-manifolds with decorated links. The maps are defined on an equivariant, filtered, curved chain homotopy type invariant, denoted $\mathcal{C}FL$. We describe the construction, and prove invariance. We also make a comparison with the graph TQFT defined by the author for Heegaard Floer homology, showing that the maps associated to cobordisms with embedded, cyclically ordered graphs can be recovered from the link Floer TQFT, giving an alternate construction and proof of invariance of those maps.

4.1 Introduction

In [OS04e] and [OS06], Ozsváth and Szabó construct algebraic invariants of 3–manifolds, called Heegaard Floer homology, such that cobordisms between 3–manifolds induce maps between the Heegaard Floer homologies of the ends. Defined by Ozsváth and Szabó in [OS04c] and independently by Rasmussen in [`?RasmussenCFr`], there is a refinement of Heegaard Floer homology for knots embedded in 3–manifolds, called knot Floer homology. In [OS08], Ozsváth and Szabó define a generalization called link Floer homology, for links embedded in 3–manifolds.

To a four manifold W with properly embedded surface $\Sigma \subseteq W$, the ends $\partial\Sigma$ are links in ∂W. In light of the cobordism maps in [OS06] associated to 4–manifolds, it is natural to

expect functorial maps between the link Floer homologies of the ends of the link cobordism (W, Σ). In this paper we provide such a construction of a $(3+1)$–dimensional TQFT for a version of the full link Floer complex, and prove invariance of the maps.

We call the version of link Floer homology we will use **curved link Floer homology**. A multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link embedded in a 3–manifold with basepoints $\mathbf{w}$ and $\mathbf{z}$, such that each component has at least two basepoints, and the basepoints alternate between $\mathbf{w}$ and $\mathbf{z}$ as one travels along the link. To an oriented multi-based link $\mathbb{L}$ in a 3–manifold Y with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, we will define algebraic objects

$$\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}).$$

In this paper we will focus on the versions $\circ \in \{-, \infty\}$. If $S = \{s_1, \dots, s_n\}$ is a set, we define $\mathbb{Z}_2[U_S]$ to be the polynomial ring generated by the formal variables $U_{s_1}, \dots, U_{s_n}$. To a diagram $\mathcal{H}$ of the link $\mathbb{L}$, we will define a $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$–module $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ with a filtration of $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$. The module $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ also has an endomorphism ∂ which satisfies

$$\partial^2 = k \cdot \text{id},$$

for some $k \in \mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$, depending on the link $\mathbb{L}$. The precise constant k can be found in Lemma 4.2.9, and is determined by the configuration of basepoints on the link. The fact that k can be nonzero results from the bubbling of Maslov index two holomorphic disks when computing ∂^2. Indeed the fact that such bubbling occurs should be thought of as an essential feature of the link Floer package, as it leads to interesting relations between the maps which appear.

Modules with an endomorphism squaring to the action of a scalar in the ground ring are often referred to as "matrix factorizations", and appear in the context of Khovanov-Rozansky homology. See [KR08]. We will call a module with such an endomorphism a **curved chain complex**.

Thus for each diagram $\mathcal{H}$, we get a curved, filtered chain complex $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ over $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$. The filtered, $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$–equivariant, curved chain homotopy type is an invariant

of $(Y, \mathbb{L}, \mathfrak{s})$. To define maps associated to link cobordisms, we need to consider a more subtle invariant, involving canonical homotopy equivalences between the complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ for various diagrams $\mathcal{H}$. To this end, we make the following definition:

Definition 4.1.1. If $\mathcal{C}$ is an category and A is a set, a **transitive system over $\mathcal{C}$ indexed by** A is a map $C : A \to \mathcal{C}$, together with distinguished morphisms

$$\Phi_{a\to a'} : C(a) \to C(a')$$

for $a, a' \in A$ such that

- $\Phi_{a\to a} = \mathrm{id}$;
- $\Phi_{a'\to a''} \circ \Phi_{a\to a'} = \Phi_{a\to a''}$.

If C_1 and C_2 are two transitive systems over $\mathcal{C}$, indexed by A_1 and A_2 respectively, then a **morphism of transitive systems** over $\mathcal{C}$ is a collection of morphisms F_{a_1,a_2}, indexed by pairs $(a_1, a_2) \in A_1 \times A_2$, such that

$$F_{a_1,a_2} : C_1(a_1) \to C_2(a_2)$$

and the following diagram commutes for any two pairs $(a_1, a_2), (a_1', a_2') \in A_1 \times A_2$:

$$\begin{array}{ccc} C_1(a_1) & \xrightarrow{F_{a_1,a_2}} & C_2(a_2) \\ \downarrow{\scriptstyle \Phi_{a_1\to a_2}} & & \downarrow{\scriptstyle \Phi_{a_2\to a_2'}} \\ C_1(a_1') & \xrightarrow{F_{a_1',a_2'}} & C_2(a_2') \end{array}.$$

The complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ fit into the structure of a transitive system over $\mathcal{C}$ indexed by A, where A is the set of all admissible diagrams for $(Y, \mathbb{L})$, and $\mathcal{C}$ is the category of chain $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$–equivariant chain complexes with a filtration of $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$. The morphisms of this category of $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$–equivariant homomorphisms modulo filtered equivariant chain homotopies. To a pair of diagrams, $\mathcal{H}$ and $\mathcal{H}'$ for a link, we can define a map

$$\Phi_{\mathcal{H}\to\mathcal{H}'} : \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H}', \mathfrak{s}),$$

which is specified up to filtered, $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$–equivariant chain homotopy using the results of [JT12]. We define $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ to be this transitive system over $\mathcal{C}$ and call it the **coherent, filtered, equivariant chain homotopy type**.

If $\mathcal{C}$ is a category of chain complexes, to a transitive system over $\mathcal{C}$, we can associate a single group by first taking the homology of each group, then taking the inverse limit using the isomorphisms $\Phi_{a\to a'}$. The homology groups of link Floer homology do not contain as much information as the chain complexes (and furthermore we can't take homology if the differential doesn't square to zero), so this isn't as useful for our purposes. As the chain homotopies $\Phi_{\mathcal{H}\to\mathcal{H}'}$ are chain homotopy equivalences and not chain isomorphisms, and are only defined up to filtered chain homotopy equivalence, it is unclear how to define a single curved chain complex using an inverse limit construction. As such, we will work with the notion of transitive systems indexed by a set. We note that a morphism between transitive systems is determined by a single one of the maps F_{a_1,a_2}.

To define cobordism maps, we need to work over a fixed category $\mathcal{C}$, while the objects $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ are transitive systems over a category $\mathcal{C}$ which depends on Y and $\mathbb{L}$. To deal with this, we must work with **colored complexes**. A coloring of a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a pair $(\sigma, \mathfrak{P})$ where $\mathfrak{P}$ is a collection of formal variables and $\sigma : \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ is a map assigning each basepoint a formal variable.

We form the colored complexes

$$\mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$$

by tensoring with a certain module to identify variables of the same color. The objects $\mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ are now transitive systems over $\mathcal{C}_{\mathfrak{P}}$, where $\mathcal{C}_{\mathfrak{P}}$ is the category of $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant, $\mathbb{Z}^{\mathfrak{P}}$–filtered chain complexes, and morphisms are chain maps up to to $\mathbb{Z}_2[U^{\mathfrak{P}}]$–equivariant, $\mathbb{Z}^{\mathfrak{P}}$–filtered chain homotopy.

A **decorated link cobordism**

$$(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$$

is a pair where W is a 4–manifold with $\partial W = -Y_1 \sqcup Y_2$, and F is a tuple $F = (\Sigma, \mathcal{A}, \sigma)$ such that

- $\Sigma \subseteq W$ is a properly embedded, oriented surface with $\partial\Sigma = -L_1 \sqcup L_2$;[1]
- $\mathcal{A}$ is a properly embedded 1–manifold in Σ such that the components of $\Sigma \setminus \mathcal{A}$ are partitioned into two subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, such that $\Sigma_{\mathbf{w}}$ contains all of the $\mathbf{w}$–basepoints on the ends, and $\Sigma_{\mathbf{z}}$ contains all of the $\mathbf{z}$–basepoints on the ends. Further, all curves in $\mathcal{A}$ are part of the boundary of both $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$;
- $\sigma : C(\Sigma \setminus \mathcal{A}) \to \mathfrak{P}$ is a coloring of the components of $\Sigma \setminus \mathcal{A}$.

We call $F = (\Sigma, \mathcal{A}, \sigma)$ a **colored surface with divides**. We include a picture of the kind of cobordism we are considering in Figure 4.1.

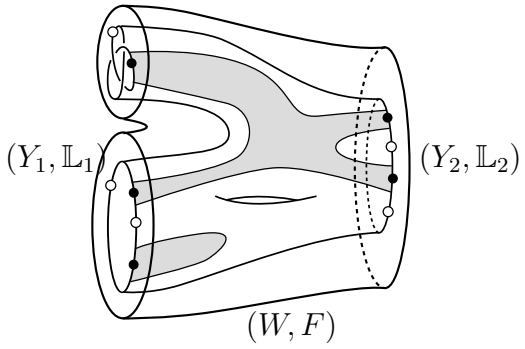

Figure 4.1: An example of a colored link cobordism (W, F). The black dots are $\mathbf{w}$–basepoints and the white dots are the $\mathbf{z}$–basepoints. The shaded region is of type–$\mathbf{w}$, and the unshaded region is of type–$\mathbf{z}$.

4.1.1 Main results

Our paper is centered on proving the following:

[1]Orientation conventions will be described in Section 4.2.

Theorem A. *To an oriented, decorated link cobordism $(W,F):(Y_1,\mathbb{L}_1)\to(Y_2,\mathbb{L}_2)$, where F is a surface with divides with coloring $(\sigma,\mathfrak{P})$, for $\mathfrak{s}\in \mathrm{Spin}^c(W)$, there are maps*

$$F_{W,F,\mathfrak{s}}:\mathcal{C}FL^\circ(Y_1,\mathbb{L}_1,\sigma|_{\mathbb{L}_1},\mathfrak{P},\mathfrak{s}|_{Y_1})\to\mathcal{C}FL^\circ(Y_2,\mathbb{L}_2,\sigma|_{\mathbb{L}_2},\mathfrak{P},\mathfrak{s}|_{Y_2}).$$

They are invariants, up to $\mathbb{Z}^{\mathfrak{P}}$–filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$–equivariant chain homotopy.

In [OS04e] and [OS06], Heegaard Floer homology is defined only for connected 3–manifolds, and maps are only defined for connected 4–manifolds. The essential limitation in that context was the definition of 1– and 3–handles. For the hat flavor, this is resolved in [Juh16], but for HF^-, HF^+ and HF^∞, this is resolved in Chapter 2, and $CF^\circ(Y,\mathbf{w},\mathfrak{s})$ is defined for possibly disconnected 3–manifolds (as long as there is a basepoint in each component), and maps are defined for possibly disconnected 4–manifolds (at the expense of keeping track of a graph or collection of paths which connects the basepoints in the ends). In our present context, we make the following remark:

Remark 4.1.2. The objects $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ can be defined as long as $\pi_0(\mathbb{L})\to\pi_0(Y)$ is a surjection. Cobordism maps for a link cobordisms (W,F) are well defined as long as $\pi_0(\partial\Sigma)\to\pi_0(\partial W)$, $\pi_0(\Sigma)\to\pi_0(W)$, and $\pi_0(\mathcal{A})\to\pi_0(\Sigma)$ are all surjections (note that we don't even require that each component of W have an incoming end or an outgoing end).

As one would expect, the standard Spin^c composition law is also satisfied:

Theorem B. *Suppose that $(W,F):(Y_1,\mathbb{L}_1)\to(Y_2,\mathbb{L}_2)$ is an arbitrary link cobordism and*

$$(W,F)=(W_2,F_2)\circ(W_1,F_1)$$

where each (W_i,F_i) is a link cobordism. If $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are Spin^c structures on W_1 and W_2 respectively, then we have

$$F_{W_2,F_2,\mathfrak{s}_2}\circ F_{W_1,F_1\mathfrak{s}_1}\simeq\sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}}F_{W,F,\mathfrak{s}}.$$

The construction of the maps $F_{W,F,\mathfrak{s}}$ goes by way of decomposing the pair (W, F) into elementary pieces, and then defining the maps for each piece. We will define maps corresponding to 1–,2– and 3–handle attachment away from the link, as well as maps for cylindrical cobordisms with simple dividing sets, as well as maps for attaching a band to a link in a 3–manifold. In Figure 4.2, we summarize the maps associated to several simple dividing sets which can appear on cylindrical surfaces $L \times [0,1]$ inside of a copy of $Y \times [0,1]$.

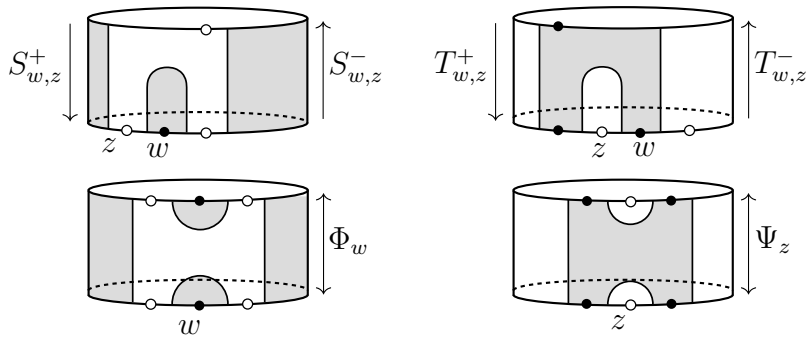

Figure 4.2: Several simple dividing sets on cylinders and the maps they correspond to. These are surfaces of the form $\Sigma = L \times [0,1]$ inside of cobordisms $W = Y \times [0,1]$.

The maps on $\mathcal{C}FL^\circ$ for the surfaces appearing in Figure 4.2 appear (using slightly different notation) in Chapter 3, where they are used to prove a formula for the map induced by moving all of the basepoints in a loop around a link component. We now briefly discuss the interplay between the algebraic properties of the maps and the colorings of the surface. The map Ψ_z is an endomorphism on $\mathcal{C}FL^\circ$ which satisfies

$$\partial\Psi_z + \Psi_z\partial = U_w + U_{w'},$$

where w and w' are basepoints adjacent to z. If we set $U_w = U_{w'}$ by coloring, then Ψ_z becomes a chain map. Likewise if we look at the picture in Figure 4.2 for the cylinder corresponding Ψ_z, we see that in order to color the surface corresponding to this diagram, the basepoints w and w' must have the same color. The interplay between the algebraic requirements for maps to be well defined on $\mathcal{C}FL^\circ$ and the coloring of surfaces will be a common thread through this paper.

Additionally, we define maps associated to oriented bands attached to links in a 3–manifold. A generic Morse function on Σ will have critical points which do not intersect the divides $\mathcal{A}$. As such, a critical point can occur in either a type–$\mathbf{w}$ region or a type–$\mathbf{z}$ region. Similarly, it turns out that there are two natural ways to define band maps on the modules $\mathcal{C}FL^\circ$. As such, we define two band maps, denoted $F_B^{\mathbf{w}}$ and $F_B^{\mathbf{z}}$, depending on whether the band is to take the designation of type–$\mathbf{w}$ or type–$\mathbf{z}$ as a part of the surface with divides.

By decomposing a link cobordism into pieces, we can define link cobordism maps as a composition involving the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ for adding and removing basepoints, the two band maps, and handle attachment maps for attaching four dimensional handles away from the link L. We then develop the Morse theory necessary to prove that the maps associated to a link cobordism do not depend on the decomposition.

4.1.2 Algebraic reduction to the graph TQFT

We will also compare the maps constructed in this paper for $\mathcal{C}FL^\circ$ to the cobordism maps constructed in Chapter 2 for cobordisms (W,Γ) with embedded ribbon graphs between closed, multi-pointed 3–manifold. For the purposes of this section, let us restrict to the case where all $\mathbf{w}$ regions and basepoints are assigned the variable U, and all $\mathbf{z}$ regions and basepoints are assigned the variable V.

Definition 4.1.3. A **ribbon graph cobordism** $(W,\Gamma) : (Y_1,\mathbf{w}_1) \to (Y_2,\mathbf{w}_2)$ between two multi-based 3–manifolds is a pair (W,Γ) such that

1. W is a cobordism from Y_1 to Y_2;
2. Γ is an embedded graph in W such that $\Gamma \cap \partial W = \mathbf{w}_1 \cup \mathbf{w}_2$ and each basepoint of $\mathbf{w}_1$ and $\mathbf{w}_2$ has valence 1 in Γ;
3. Γ is decorated with a ribbon structure, i.e., a choice of cyclic ordering of the edges adjacent to each vertex.

A **coloring** of a graph is an assignment of a formal variable to each component of a graph.

In Chapter 2, the author proves that there are Spin^c–functorial maps $F_{W,\Gamma,\mathfrak{s}}$ associated to colored ribbon graph cobordisms (W,Γ) with a 4–dimensional Spin^c structure.

The requirement in Chapter 2 that a graph cobordism be equipped with a cyclic orderings of the vertices seems very suggestive of a relationship between the graph TQFT and a TQFT involving surfaces embedded in a 4–manifold. We make the following definition:

Definition 4.1.4. A **ribbon surface cobordism** (W,R) is between two multi-based 3–manifolds $(Y_1,\mathbf{w}_1)$ and $(Y_2,\mathbf{w}_2)$ is a pair such that W is a cobordism from Y_1 to Y_2, and $R\subseteq W$ is an oriented surface with boundary and corners such that the following holds:

1. $R\cap Y_i$ consists of a finite collection of closed intervals, such that each closed interval of $R\cap Y_i$ contains exactly one of the basepoints of $\mathbf{w}_i$;
2. The corners of R correspond exactly to the boundaries of the closed intervals forming $R\cap Y_i$;
3. R contains no closed components.

From a ribbon graph cobordism (W,Γ), one can construct a surface with boundary and corners in W, for which we write R_Γ, which is defined up to appropriate isotopy. The surface R_Γ intersects ∂W in a collection of closed intervals, such that each basepoint is contained in exactly one such closed intervals. We say that two graph cobordisms (W,Γ_1) and (W,Γ_2) are **ribbon equivalent** if R_{Γ_1} and R_{Γ_2} are isotopic through ribbon surfaces. See Section 4.14 for a more precise account of ribbon surfaces, and ribbon equivalence.

Given a decorated link cobordism $(W,F):(Y_1,L_1,\mathbf{w}_1,\mathbf{z}_1)\to(Y_2,L_2,\mathbf{w}_2,\mathbf{z}_2)$, by forgetting about half of the link cobordism, we get two ribbon surface cobordisms

$$(W,\Sigma_{\mathbf{w}}):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2),\qquad (W,\Sigma_{\mathbf{z}}):(Y_1,\mathbf{z}_1)\to(Y_2,\mathbf{z}_2).$$

Given a ribbon surface cobordism (W,R), we can consider ribbon graphs embedded in R, such that R can be built by attaching disks along at the vertices of Γ, and strips along

the edges of Γ (see Section 4.14 for a precise definition). The surface in a ribbon surface cobordism always possesses a ribbon graph core. For a link cobordism with subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, let us write $\Gamma(\Sigma_{\mathbf{w}})$ and $\Gamma(\Sigma_{\mathbf{z}})$ for any choice of ribbon graph cores for $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ respectively.

If we tensor the $\mathbb{Z}_2[U,V]$–module $\mathcal{C}FL^-(Y_1,\mathbb{L}_1,\mathfrak{s}_1)$ with $\mathbb{Z}_2[U,V]/(V-1)$, where then we recover

$$CF^-(Y_1,\mathbf{w}_1,\mathfrak{s}_1).$$

If we instead tensor $\mathcal{C}FL^-(Y_1,\mathbb{L}_1,\mathfrak{s}_1)$ with $\mathbb{Z}_2[U,V]/(U-1)$ we instead recover

$$CF^-(Y_1,\mathbf{z}_1,\mathfrak{s}_1-PD[L]),$$

where $[\Sigma]$ is the underlying homology class of the surface in the link cobordism (W,F), viewed as an element of $H_2(W,\partial W;\mathbb{Z})$. The map $F_{W,F,\mathfrak{s}}$ associated to the link cobordism (W,F) thus yield maps on CF^-, for which we write

$$F_{W,F,\mathfrak{s}}|_{U=1}, \qquad \text{and} \qquad F_{W,F,\mathfrak{s}}|_{V=1}.$$

Theorem C. *Given a ribbon graph cobordism $(W,\Gamma):(Y_1,\mathbf{w}_1)\to(Y_2,\mathbf{w}_2)$ with subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ with choices of ribbon graph cores $\Gamma(\Sigma_{\mathbf{w}})$ and $\Gamma(\Sigma_{\mathbf{z}})$, then the reductions satisfy*

$$F_{W,F,\mathfrak{s}}|_{V=1}\simeq F^B_{W,\Gamma(\Sigma_{\mathbf{w}}),\mathfrak{s}}, \qquad \textit{and} \qquad F_{W,F,\mathfrak{s}}|_{U=1}\simeq F^A_{W,\Gamma(\Sigma_{\mathbf{z}}),\mathfrak{s}-PD[\Sigma]}.$$

In the above theorem, the maps $F^A_{W,\Gamma,\mathfrak{s}}$, are the graph cobordism maps, as defined in Chapter 2. The maps $F^B_{W,\Gamma,\mathfrak{s}}$ are minor variations, which we will describe in Section 4.14.

As corollaries of the above theorem, we have the following:

Corollary D. *If (W,Γ) and (W,Γ') are ribbon equivalent ribbon graph cobordisms, then*

$$F^A_{W,\Gamma,\mathfrak{s}}\simeq F^A_{W,\Gamma',\mathfrak{s}},$$

and similarly for the maps $F^B_{W,\Gamma,\mathfrak{s}}$.

Corollary E. *The $V = 1$ reduction of $F_{W,F,\mathfrak{s}}$ induces a map*

$$F_{W,F,\mathfrak{s}}|_{V=1} : CF^-(Y_1, \mathbf{w}_1, \mathfrak{s}|_{Y_1}) \to CF^-(Y_2, \mathbf{w}_2, \mathfrak{s}|_{Y_2})$$

which only depends on $\mathfrak{s}$ and the ribbon surface $\Sigma_{\mathbf{w}}$. The $U = 1$ reduction of $F_{W,F,\mathfrak{s}}$ determines a map

$$F_{W,F,\mathfrak{s}}|_{U=1} : CF^-(Y_1, \mathbf{z}_1, \mathfrak{s} - PD[L_1]) \to CF^-(Y_2, \mathbf{z}_2, \mathfrak{s} - PD[L_2])$$

which depends only on the ribbon surface $\Sigma_{\mathbf{z}}$, and the Spin^c *structure $\mathfrak{s} - PD[\Sigma]$.*

4.1.3 Comparison of curved link Floer homology to other versions of link Floer homology

In [OS04c] and [OS08], Ozsváth and Szabó define knot and link Floer homologies for knots and links in rational homology spheres. Curved link Floer homology recovers these theories. For example, if K is a knot in an integer homology sphere, then the object $CFK^\infty(Y, K)$, defined by Ozsváth and Szabó, is a $\mathbb{Z}_2$ subcomplex of $\mathcal{C}FL^\infty(Y, K, \mathfrak{s}_0)$, where $\mathfrak{s}_0$ is the unique element of Spin$^c(Y)$. In fact, in Chapter 5, we will define an absolute Alexander grading on $\mathcal{C}FL^\infty(Y, K, \mathfrak{s}_0)$ and $CFK^\infty(Y, K)$ can naturally be identified with $\mathcal{C}FL^\infty(Y, K, \mathfrak{s}_0)_0$, the subcomplex in zero Alexander grading. The version of the full link Floer complex from [OS08] can be recovered similarly, by first tensoring $\mathcal{C}FL^\circ$ with a module to identify all of the $V_{\mathbf{z}}$ variables on each link component, and then taking a subcomplex in a single Alexander grading.

One of the advantages of using curved link Floer homology is that it allows us to define maps without paying attention to Alexander gradings or relative Spin^c structures. For example, the maps Φ_w and Ψ_z are not endomorphisms of $CFK^\infty(Y, K)$, as they change the relative Spin^c structure (though the product $\Phi_w\Psi_z$ is an endomorphism). If we were very careful about gradings, we could work with the uncurved complexes, though this would be extremely tedious. Additionally, we will need to work with links which may have components which are not nullhomologous, even rationally. To deal with this, one would need to adapt the more standard $CFL^\infty(Y, L)$ complexes to deal with this.

4.1.4 Comparison to other link cobordism map constructions

There are several examples of constructions like ours already in the literature. The most similar is the one in [Juh16], where Juhász constructs a TQFT for $\widehat{HFL}$ associated to cobordisms with embedded surfaces, decorated with a dividing set. The resulting maps arising in the cases of knot concordances are studied in [JM16b]. It would be interesting to compare our maps, to those defined by Juhász. Indeed our maps do yield maps on $\widehat{HFL}$ by simply setting all of the variables to zero. In such a case, the colorings become irrelevant, and our definition of a decorated link cobordism becomes essentially equivalent to the one in [Juh16]. To this end, we make the following conjecture:

Conjecture 4.1.5. *The maps* $\widehat{F}_{W,F} = \sum_{\mathfrak{s}\in \mathrm{Spin}^c(W)} \widehat{F}_{W,F,\mathfrak{s}}$ *defined in this paper for the hat flavor coincide with the maps* $F_{W,F}$ *defined in* [Juh16] *using sutured Floer homology.*

In [JM16a], Juhász and Marengon provide a full description of all of the link cobordism maps arising from the sutured Floer TQFT, as maps which defined in terms of generators of the complexes, as well as maps defined by counting holomorphic curves on Heegaard triple diagrams. In principle, verification of the above conjecture should be tractable using their description, but we will not pursue it in this paper.

There are additional partial constructions of cobordism maps. For example, in [Sar11a], Sarkar defines maps associated to surfaces in $S^3 \times [0,1]$ in order to provide an alternate proof of the Milnor conjecture using grid diagrams. In [OSS15], Ozsváth, Stipsicz and Szabó use grid diagrams to analyze the unoriented four-ball genus of knots and links.

Additionally, in [AE], Alishahi and Eftekhary provide a notion of cobordism between tangles and give a construction of cobordism maps on a minus version of sutured Floer homology. The construction has some similarities to the construction in this paper, such as a construction of maps associated to bands using holomorphic triangle counts. As a specialization, their construction does define maps for certain link cobordisms, though the maps in this paper are defined for a more broad class of link cobordisms, and the algebraic properties of the maps and the induced TQFT are exploited differently in this paper.

Despite these constructions, there is still plenty of room for a new theory, since most constructions either are defined using grid diagrams, admit no straightforward generalization to the full knot Floer complex, or are limited in terms of what 4–manifolds and surfaces they can be applied to. The grid diagram approaches feature the advantage of being defined in a relatively combinatorial fashion, though proving invariance is harder in the setting of grid diagrams, and there is no obvious generalization to links in arbitrary 3–manifolds.

4.1.5 Organization

In Section 4.2 we review the definition of the chain complexes $\mathcal{C}FL^\circ$ and the category of decorated link cobordisms. In Section 4.3 we describe the maps Φ_w, Ψ_z, the quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$, and the maps $S^{\pm}_{\mathbb{U},D}$, which appear in our construction of the link Floer TQFT. We then prove several important properties and relations of these maps. In Section 4.4 we describe handle attachment maps for handles which are attached along framed spheres which don't intersect L. In Section 4.5 we perform some technical triangle map computations. In Section 4.6 we define maps associated to bands, and prove some important relations. In Section 4.8 we define cobordism maps associated to a 4–dimensional 1–handle with a 2–dimensional band inside of it. In Section 4.9, we define the notions of a parametrized cobordism and a parametrized decomposition of a link cobordism. In Section 4.10 we prove that we can move between any two parametrized decompositions via a set of moves. In Section 4.11 we give a precise definition of the cobordism maps involved, in terms of maps on Heegaard Floer homology and parametrized decompositions. In Section 4.12, we provide the construction when the cobordism has a component with an empty end. In Section 4.13, we prove the Spin^c composition law. In Section 4.14 we show how the construction here recovers the graph TQFT defined in Chapter 2, and we explain the connection.

4.2 Decorated link cobordisms, the complexes $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$ and the coherent chain homotopy type invariant $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$

In this section, we describe background about the category of decorated link cobordisms, and the definitions of the objects $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ which will appear in this paper.

4.2.1 The category of decorated link cobordisms

In this section we provide the definitions of multi-based links in 3–manifolds, and link cobordisms of colored surfaces with divides between them.

Definition 4.2.1. A **three manifold with multi-based link** is a pair $(Y,\mathbb{L})$ where Y is an oriented 3–manifold and $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ is a tuple where L is an oriented link in Y, and $\mathbf{w}$ and $\mathbf{z}$ are collections of basepoints on L, such that each component of L has at least two basepoints, and the basepoints alternate between those in $\mathbf{w}$ and those in $\mathbf{z}$ as one traverses the link.

Definition 4.2.2. A **surface with divides** is a pair $(\Sigma,\mathcal{A})$ where Σ is a compact, oriented[2] surface together with a collection $\mathcal{A}$ of properly embedded 1–manifolds such that $\Sigma\setminus\mathcal{A}$ is partitioned into subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, such that $\Sigma_{\mathbf{w}}$ is on one side of $\mathcal{A}$ and $\Sigma_{\mathbf{z}}$ is on the other.

To achieve well defined, functorial maps in our link TQFT, we need to consider colorings of the basepoints and the surfaces with divides.

Definition 4.2.3. A coloring of a multibased link $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ is a pair $(\sigma,\mathfrak{P})$ where $\mathfrak{P}$ is a finite set of colors and $\sigma:(\mathbf{w}\cup\mathbf{z})\to\mathfrak{P}$ is an assignment of a color to each basepoint. A coloring of a surface with divides $(\Sigma,\mathcal{A})$ is a pair $(\sigma,\mathfrak{P})$ where $\sigma:C(\Sigma\setminus\mathcal{A})\to\mathfrak{P}$ is an assignment of a color to each component of $\Sigma\setminus\mathcal{A}$.

[2]Our orientation convention that if N is an outward pointing normal vector and $v\in T(\partial\Sigma)$, then v is positively oriented if the pair (v,N) is an oriented basis for $T\Sigma$. This is merely a convention, though it affects the statement of Theorem C, as well as an Alexander grading formula, to appear in Chapter 5. This is the natural orientation to choose, as it gives $L\times[0,1]$, the product orientation, when viewed as a cobordism from $L\times\{0\}$ to $L\times\{1\}$.

We now define a notion of morphism between 3–manifolds with multi-based links.

Definition 4.2.4. A pair (W, F) is called a **decorated link cobordism** from $(Y_1, (L_1, \mathbf{w}_1, \mathbf{z}_1))$ to $(Y_2, (L_2, \mathbf{w}_2, \mathbf{z}_2)$ if W is a compact 4–manifold with $\partial W = -Y_1 \sqcup Y_2$, and $F = (\Sigma, \mathcal{A}, \sigma)$ is a colored surface with divides, $\Sigma \subseteq W$ is a properly embedded, compact, oriented surface with $\Sigma \cap W = \partial L_i$. Furthermore, each component of $L_i \setminus \mathcal{A}$ contains exactly one basepoint of $\mathbf{w}_i \cup \mathbf{z}_i$, which is of the same type as the component of $\Sigma \setminus \mathcal{A}$ which it is in.

To properly define a link cobordism category, we need to work with cobordisms which have a parametrization of the ends. More precisely we define a parametrized link cobordism $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ to be a link cobordism $(W, F) : (Y_1', \mathbb{L}_1') \to (Y_2', \mathbb{L}_2')$ in the previous sense, together with a choice of diffeomorphism $\phi_i : (Y_i', \mathbb{L}_i') \to (Y_i, \mathbb{L}_i)$. Such a definition is necessary to define a proper category, since otherwise the cobordism category lacks identity morphisms, for example. We should always remember that this is the category we are working in, though it's usually cumbersome to make this explicit.

Remark 4.2.5. One could define a category where link cobordisms with nonorientable surfaces were allowed. It may be possible to define cobordism maps for nonorientable surfaces on a version of $\mathcal{C}FL^\infty$ obtained by identifying the $U_{\mathbf{w}}$ and $V_{\mathbf{z}}$ variables, as in [OSS15], but this is work for a future paper. In this paper we restrict to oriented surfaces.

4.2.2 Curved chain complexes

For the purposes of this section, assume that $\mathcal{R}$ is a ring of characteristic two. In this section we describe the notion of a curved chain complex, which is basically a chain complex where the differential squares to a scalar in $\mathcal{R}$. We could avoid using these if we were willing to either only have two variables (a U–variable and a V–variable), or we are willing to lose some amount of symmetry between the $\mathbf{w}$–basepoints and the $\mathbf{z}$–basepoints, as we did in Chapter 3, where we set all of the $V_{\mathbf{z}}$ variables on a link component equal to each other. The setting of curved chain complexes seems like the most natural, so we will use it.

We now precisely describe some terminology (which is elementary in the case of actual chain complexes):

Definition 4.2.6. If $\mathcal{R}$ is a ring, a **curved $\mathcal{R}$–chain complex** is an $\mathcal{R}$–module C, together with an endomorphism $\partial : C \to C$ such that

$$\partial^2 = k \cdot \mathrm{id},$$

for some $k \in \mathcal{R}$. We say an $\mathcal{R}$–homomorphism $F : (C, \partial) \to (C', \partial')$ of curved $\mathcal{R}$–chain complexes is a **chain map** if $F\partial = \partial F$. An $\mathcal{R}$–equivariant **chain homotopy** between morphisms $F, G : (C, \partial) \to (C', \partial')$ is an $\mathcal{R}$–homomorphism $H : C \to C'$ such that $F + G = \partial' H + H\partial$.

Note that being chain homotopic is an equivalence relation. We note that the class of curved $\mathcal{R}$–chain complexes is a category, with morphisms being ($\mathcal{R}$–equivariant) chain maps up to ($\mathcal{R}$–equivariant) chain homotopy.

Definition 4.2.7. A $\mathbb{Z}^n$ filtration on a curved $\mathcal{R}$–chain complex (C, ∂) is a collection of subsets $\mathcal{F}_I \subseteq C$ ranging over $I \in \mathbb{Z}^n$ such that $\partial \mathcal{F}_I \subseteq \mathcal{F}_I$ and such that if $I \leq J$ (i.e. each component satisfies that inequality) then $\mathcal{F}_I \subseteq \mathcal{F}_J$.

4.2.3 Heegaard diagrams for multi-based links

If $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multi-based link in Y^3, we define a **Heegaard diagram** $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ to be an embedded Heegaard diagram of Y such that

1. $\Sigma \cap L = \mathbf{w} \cup \mathbf{z}$, and Σ divides Y into two solid handlebodies, U_α and U_β,
2. Σ is oriented with a normal vector pointing into U_β;
3. L intersects Σ positively at the $\mathbf{z}$–basepoints and negatively at the $\mathbf{w}$–basepoints;
4. $L \cap U_\alpha$ is isotopic in U_α, in the complement of compressing disks for the $\boldsymbol{\alpha}$ curves, relative to its endpoints to a collection of arcs in Σ;

5. $L \cap U_\beta$ is isotopic in U_β, in the complement of compressing disks for the $\boldsymbol{\beta}$ curves, relative to its endpoints to a collection of arcs in Σ.

Alternatively, we can connect choose paths on Σ from the $\mathbf{z}$ basepoints to the $\mathbf{w}$ basepoints in the complement of the $\boldsymbol{\beta}$ curves, and choose paths from the $\mathbf{w}$ basepoints to the $\mathbf{z}$ basepoints in the complement of the $\boldsymbol{\alpha}$ curves. Pushing off these collections of curves into U_β and U_α respectively yields a link which is isotopic to L.

To achieve functorial maps, we must work with embedded Heegaard diagrams, i.e. Heegaard diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ equipped with an embedding

$$(\Sigma, \mathbf{w}, \mathbf{z}) \hookrightarrow (Y, \mathbf{w}, \mathbf{z}).$$

An "abstract" Heegaard diagram (a diagram without a choice of embedding) is enough to describe the isomorphism class of the various Heegaard Floer homology invariants $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$, though to describe functorial maps between the complexes $\mathcal{C}FL^\circ$ for various links and 3–manifolds, one needs canonical change of diagrams maps

$$\Phi_{\mathcal{H}\to\mathcal{H}'} : \mathcal{C}FL(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL(\mathcal{H}', \mathfrak{s}),$$

which are defined up to equivariant, filtered chain homotopy. According to [JT12], to describe such maps it is necessary to work with embedded Heegaard diagrams.

4.2.4 Spin^c structures

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is an ℓ–pointed Heegaard diagram for Y. We can define a map

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

If $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is an intersection point, then we define $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})$, by picking a Morse function defining $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and removing regular neighborhoods of the flowlines passing through the points in $\mathbf{x}$, as well as the basepoints $\mathbf{w}$. This yields a nonvanishing vector field on the complement of this neighborhood of these flowlines, which will extend over these balls, since

the flowlines connected critical points with indices of opposite parity. The map $\mathfrak{s}_{\mathbf{w}}$ depends on the choice of basepoints, $\mathbf{w}$, and indeed picking a different basepoints in the same Heegaard diagram changes the map $\mathfrak{s}_{\mathbf{w}}$. We have the following:

Lemma 4.2.8. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, L, \mathbf{w}, \mathbf{z})$ and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we have*

$$\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) - \mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = PD[L].$$

Proof. This follows from reasoning identical to the proof of [OS04e, Lem. 2.19]. Alternatively, $\mathrm{Spin}^c(Y)$ is described as the set of homology classes of nonvanishing vector fields on Y. Given a closed path $\gamma \in H_1(Y;\mathbb{Z})$, the action of $\gamma \in H_1(Y;\mathbb{Z})$, can be described by a procedure sometimes referred to as Reeb surgery (cf. [Nic03, Section 3.2]). Given the explicit description from [OS04e] of the vector fields $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$, it is easy to see that they differ by Reeb surgery on the link L. □

Hence if the sum of all the components of a link is trivial in $H_1(Y;\mathbb{Z})$, then it doesn't matter whether we take intersection points with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$ or $\mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = \mathfrak{s}$. For a link which is not nullhomologous in this sense, the resulting modules will in general vary. An instructive example is $Y = (S^1 \times S^2)$ with $K = S^1 \times \{pt\}$.

4.2.5 The chain complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$

Given a strongly–$\mathfrak{s}$ admissible Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for a link $\mathbb{L}$ in Y^3, we define

$$\mathcal{C}FL^-(\mathcal{H}, \mathfrak{s})$$

to be the free $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$–module with generators of the form

$$U_{\mathbf{w}}^I V_{\mathbf{z}}^J \cdot \mathbf{x},$$

where I and J are nonnegative multi-indices and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is an intersection point with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$. The module

$$\mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s})$$

is defined as the module with the same generators, but allowing any multi-indices I and J. The module $\mathcal{C}FL^+(\mathcal{H},\mathfrak{s})$ can be defined as the quotient

$$\mathcal{C}FL^+(\mathcal{H},\mathfrak{s}) = \mathcal{C}FL^\infty(\mathcal{H},\mathfrak{s})/\mathcal{C}FL^-(\mathcal{H},\mathfrak{s}).$$

Now given a multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, we will need to consider diagrams which satisfy an extra admissibility assumption. The admissibility assumption is called strong $\mathfrak{s}$–admissibility, which is described in [OS04e, Def. 4.10], and ensures necessary finiteness results on counts of flowlines.

Given an almost complex structure J_s on $\Sigma \times [0,1] \times \mathbb{R}$ which is suitably generic in the sense of [Lip06], we define an endomorphism

$$\partial : \mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$$

by the formula

$$\partial(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

The map ∂ does not square to zero, in general. Instead we have the following:

Lemma 4.2.9 (Lemma 3.2.1)**.** *The map $\partial : \mathcal{C}FL^\infty(\mathcal{H},\mathfrak{s}) \to \mathcal{C}FL^\infty(\mathcal{H},\mathfrak{s})$ satisfies*

$$\partial^2 = \sum_{K\in|L|} (U_{w_{K,1}} V_{z_{K,1}} + V_{z_{K,1}} U_{w_{K,2}} + U_{w_{K,2}} V_{z_{K,2}} + \cdots V_{z_{K,n_K}} U_{w_{K,1}}),$$

where $w_{K,1}, z_{K,1}, \ldots, w_{K,n_K}, z_{K,n_K}$ are the basepoints on the link component K, in the order that they appear on K.

Thus $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$ is a curved chain complex over the ring $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$.

Given two diagrams, $\mathcal{H}$ and $\mathcal{H}'$, for $(Y,\mathbb{L})$, by composing change of diagrams maps corresponding to elementary Heegaard moves (i.e $(1,2)$–stabilization, isotopy relative L, handleslides and isotopies of the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves) we can define maps

$$\Phi_{\mathcal{H}\to\mathcal{H}'} : \mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H}',\mathfrak{s})$$

as a composition of maps corresponding to elementary moves. Since the proof of simple handleswap invariance from [JT12] follows from a holomorphic triangle argument, invariance from simple handleswaps holds for $\mathcal{C}FL^\circ$. It follows from [JT12] that up to equivariant, filtered chain homotopy, these maps are independent of the choice of elementary Heegaard moves connecting $\mathcal{H}$ and $\mathcal{H}'$.

Thus we can define $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ as the transitive system over $\mathcal{C}$ indexed by A (cf. Definition 4.1.1), where A is the set of all strongly $\mathfrak{s}$–admissible diagrams for $(Y,\mathbb{L})$ and $\mathcal{C}$ is the category of $\mathbb{Z}^{\mathbf{w}}\oplus\mathbb{Z}^{\mathbf{z}}$–filtered, $\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]$–equivariant curved chain complexes. We call $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ the **coherent, equivariant, filtered, curved chain homotopy type**, which we will usually abbreviate by "coherent chain homotopy type".

4.2.6 Colorings of complexes, and the category of curved $\mathbb{Z}^{\mathfrak{P}}$–filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain complexes

An important algebraic step in constructing our link Floer TQFT is coloring complexes. Recall that a coloring $(\sigma,\mathfrak{P})$ of a link $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ is a finite set $\mathfrak{P}$ indexing a collection of formal variables, as well as a map

$$\sigma:\mathbf{w}\cup\mathbf{z}\to\mathfrak{P}.$$

Since we are working with curved complexes instead of standard chain complexes, we no longer need to make any assumptions on $(\sigma,\mathfrak{P})$ to get a well defined curved chain complex (in Chapter 3 we assumed that two $\mathbf{z}$–basepoints on a link component were given the same color if they were on the same link component). Given a link $\mathbb{L}$ and a coloring $(\sigma,\mathfrak{P})$, we define $I_{\sigma,\mathfrak{P}}$ as the ideal of $\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}},U_{\mathfrak{P}}]$ generated by elements of the form $U_w-U_{\sigma(w)}$ and $V_z-U_{\sigma(z)}$. Then we define the coloring module as

$$C_{\sigma,\mathfrak{P}}=\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}},U_{\mathfrak{P}}]/I_{\sigma,\mathfrak{P}}.$$

Finally we define the colored complexes

$$\mathcal{C}FL^-(\mathcal{H},\sigma,\mathfrak{P},\mathfrak{s})=\mathcal{C}FL^-(\mathcal{H},\mathfrak{s})\otimes_{\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]}C_{\sigma,\mathfrak{P}},$$

and similarly for the ∞ flavor.

The objects $\mathcal{C}FL^\circ(\mathcal{H},\sigma,\mathfrak{P},\mathfrak{s})$ become transitive systems over $\mathcal{C}$, indexed by A, where $\mathcal{C}$ is the category of $\mathbb{Z}^{\mathfrak{P}}$–filtered (the filtration being induced simply by the powers of the variables), curved $\mathbb{Z}_2[U_{\mathfrak{P}}]$–chain complexes, and A is the set of strongly $\mathfrak{s}$–admissible Heegaard diagrams for $(Y,\mathbb{L})$.

We also note that including colorings is essential for functoriality. Without colorings, we do not have a fixed category to work in. Also, most of the maps we define for decorations on the surface are not chain maps until we color the complexes.

Although we will primarily work with curved chain complexes, we will occasionally have the need to actually take homology. In the case that $(\sigma,\mathfrak{P})$ is a coloring of $\mathbb{L}$ such that $\partial^2=0$ on $\mathcal{C}FL^\circ(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$, we will write

$$\mathcal{H}FL^\circ(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$$

for the associated homology group.

4.3 Background on $\mathcal{C}FL^\circ$, quasi-stabilization, basepoint maps, and the disk stabilization maps

In this section, we include several important computations, as well some basic results about the maps $S^{\pm}_{w,z}, T^{\pm}_{w,z}$ and $S^{\pm}_{\mathbb{U},D}$, which appear in the construction of the link Floer TQFT.

4.3.1 The modules $\mathcal{H}FL^-$ for unlinks in $(S^1\times S^2)^{\#g}$

In this section, we analyze $\mathcal{H}FL^-$ for unlinks in connected sums of $(S^1\times S^2)$ with various configurations of basepoints. The analogous statements hold for disjoint unions of such manifolds, though we phrase them only for connected manifolds, for notational simplicity.

We note that when $c_1(\mathfrak{s})$ is torsion, the objects $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$ have two relative homological gradings, each analogous to the grading on closed 3–manifold invariants defined by Ozsváth

and Szabó. We denote one grading by $\mathrm{gr}_{\mathbf{w}}$ and the other by $\mathrm{gr}_{\mathbf{z}}$. The the $U_{\mathbf{w}}$ variables are -2 graded with respect to $\mathrm{gr}_{\mathbf{w}}$ and the $V_{\mathbf{z}}$ variables are 0–graded with respect to $\mathrm{gr}_{\mathbf{w}}$. Similarly the $V_{\mathbf{z}}$ variables are -2 graded with respect to $\mathrm{gr}_{\mathbf{z}}$ and the $U_{\mathbf{w}}$ variables are 0–graded.

For a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$, we define

$$\mathrm{gr}_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = \mu(\phi) - 2\sum_{w \in \mathbf{w}} n_w(\phi),$$

and

$$\mathrm{gr}_{\mathbf{z}}(\mathbf{x}, \mathbf{y}) = \mu(\phi) - 2\sum_{z \in \mathbf{z}} n_z(\phi).$$

If $c_1(\mathfrak{s})$ is torsion, then $\mathrm{gr}_{\mathbf{w}}$ is independent of the disk ϕ, using the Chern class formula from [OS04d]. Analogously, using Lemma 4.2.8, if $c_1(\mathfrak{s} - PD[L])$ is torsion, then $\mathrm{gr}_{\mathbf{z}}$ will be independent of the disk ϕ.

For $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$ we will define $\mathcal{H}FL^-_{\mathrm{max},\mathbf{o}}(Y, \mathbb{L}, \mathfrak{s})$ to denote the subset of $\mathcal{H}FL^-(Y, \mathbb{L}, \mathfrak{s}_0)$ of maximal homological grading with respect to $\mathrm{gr}_{\mathbf{o}}$.

Lemma 4.3.1. *Suppose $\mathbb{U}$ is an unlink in $(S^1 \times S^2)^{\#g}$ with ℓ components, and each component has exactly two basepoints. Then there is an isomorphism*

$$\mathcal{H}FL^-_{\mathrm{max},\mathbf{w}}((S^1 \times S^2)^{\#g}, \mathbb{U}, \mathfrak{s}_0) \cong \mathbb{Z}_2[V_{\mathbf{z}}].$$

In particular, there is an element which generates the maximally $\mathrm{gr}_{\mathbf{w}}$–graded subset with respect to the $\mathbb{Z}_2[V_{\mathbf{z}}]$–module action. Similarly

$$\mathcal{H}FL^-_{\mathrm{max},\mathbf{z}}((S^1 \times S^2)^{\#g}, \mathbb{U}, \mathfrak{s}_0) \cong \mathbb{Z}_2[U_{\mathbf{w}}],$$

and hence the latter object has an algebraic generator, as well.

The two generators are equal in $\mathcal{H}FL^-$.

Proof. We will prove the claim about $\mathcal{H}FL^-_{\mathrm{max},\mathbf{w}}$ since the claim about $\mathcal{H}FL^-_{\mathrm{max},\mathbf{z}}$ follows identically. We can pick g homologically independent curves $\alpha_1, \dots, \alpha_g$ on Σ_g, add a pair of basepoints w_1, and z_1 in the (single) region of $\Sigma_g \setminus (\alpha_1 \cup \dots \cup \alpha_g)$.

Now pick $\ell - 1$ other points in $\Sigma \setminus (\alpha_1 \cup \cdots \cup \alpha_g)$, and then add curves $\alpha_{g+1}, \ldots, \alpha_{g+\ell-1}$ which each bound a disk containing exactly one of the chosen $\ell - 1$ points. Call the totality of these curves $\boldsymbol{\alpha}$ and pick $\boldsymbol{\beta}$ to be small Hamiltonian isotopies of the $\boldsymbol{\alpha}$ curves. Now pick basepoints w_i, z_i for $2 \leq i \leq \ell$ in each of the disks bounded by the curves $\alpha_{g+1}, \ldots, \alpha_{g+\ell-1}$ and their corresponding $\boldsymbol{\beta}$–curves.

In this diagram, there is a canonical top degree intersection point Θ. The top degree subset of $\mathcal{C}FL^-(\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$ is easily seen to be generated by sums of elements of the form $V_{\mathbf{z}}^J \cdot \Theta$. Now we claim that any sum of elements of the form $V_{\mathbf{z}}^J \cdot \Theta$ is a cycle, and no nontrivial sum of such elements is a boundary.

To see that any sum of elements of the form $V_{\mathbf{z}}^J \cdot \Theta$ is a cycle, one simply analyzes the proof of [OS04e, Lem. 9.1]. Using the relative Maslov grading, one sees that if $\phi \in \pi_2(\Theta, \mathbf{y})$ is a Maslov index 1 disk, then

$$2 \sum_{w \in \mathbf{w}} n_w(\phi) \leq 0,$$

hence any such disk admitting holomorphic representatives cannot cross any of the $\mathbf{w}$–basepoints. Since the diagram is so simple, it is easy to see that the disks which don't cross over any $\mathbf{w}$–basepoints come in pairs, and hence cancel when counted by the differential. Thus $\partial\Theta = 0$.

To see that no nontrivial sum of terms of the form $V_{\mathbf{z}}^J \cdot \Theta$ is a boundary, we note that any disk $\phi \in \pi_2(\mathbf{y}, \Theta)$ satisfies

$$\mu(\phi) - 2 \sum_{w \in \mathbf{w}} n_w(\phi) \leq 0,$$

by the relative homological grading formula, and hence

$$2 \sum_{w \in \mathbf{w}} n_w(\phi) \geq 1,$$

implying that if ϕ has a holomorphic representative which is counted by the differential, it must have nonzero multiplicities on the $\mathbf{w}$–basepoints. In particular ϕ must increase the $U_{\mathbf{w}}$–filtration, and hence all cycles live in the subset of $\mathcal{C}FL^-$ of $U_{\mathbf{w}}$–filtration at least 1, and

in particular nothing which is a sum of elements of the form $V_{\mathbf{z}}^J \cdot \Theta$ is a cycle. Hence

$$\mathcal{H}FL^-_{\max,\mathbf{w}}((S^1 \times S^2)^{\#g}, \mathbb{U}, \mathfrak{s}_0) \cong \mathbb{Z}_2[V_{\mathbf{z}}].$$

That the generators for the two maximal subsets agree is obvious, from the diagram we've constructed. □

We now consider the case that all components except for one of the unlink $\mathbb{U} \subseteq (S^1 \times S^2)^{\#g}$ have exactly two basepoints, and that the remaining component of $\mathbb{U}$ has exactly four basepoints. These will appear when we define maps associated to bands.

Lemma 4.3.2. *Suppose that $\mathbb{U}$ is an unlink in $(S^1 \times S^2)^{\#g}$ such that all components of $\mathbb{U}$ except for one have exactly two basepoints, and one component of $\mathbb{U}$ has exactly four basepoints. Let z and z' denote the two $\mathbf{z}$–basepoints which are on the same link component. Suppose $\mathfrak{P}_0 = \mathbf{w} \cup (\mathbf{z}/(z = z'))$ where $\mathbf{z}/(z = z')$ is the set of $\mathbf{z}$–basepoints, modulo the relation that $z = z'$, and σ_0 is the natural map from $\mathbf{w} \cup \mathbf{z}$ to $\mathfrak{P}_0$. Then we have an isomorphism*

$$\mathcal{H}FL^-_{\max,\mathbf{w}}((S^1 \times S^2)^{\#g}, \mathbb{U}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0) \cong \mathbb{Z}_2[V_{\mathbf{z}}]/(V_z - V_{z'}).$$

Similarly if w and w' are the two $\mathbf{w}$ basepoints on this link component, and $(\sigma_0', \mathfrak{P}_0')$ is the coloring with $\mathfrak{P}_0' = (\mathbf{w}/(w = w')) \cup \mathbf{z}$ and σ_0' is the obvious map, we have

$$\mathcal{H}FL^-_{\max,\mathbf{z}}((S^1 \times S^2)^{\#g}, \mathbb{U}, \sigma_0', \mathfrak{P}_0', \mathfrak{s}_0) \cong \mathbb{Z}_2[U_{\mathbf{w}}]/(U_w - U_{w'}).$$

Proof. Again we will only prove the claim about $\mathcal{H}FL^-_{\max,\mathbf{w}}$, since the other one follows identically, by simply switching the w's and z's. This is a modification of the proof of Lemma 4.3.1. Let $\mathbb{U}_0$ denote the unlink which has all of the components of $\mathbb{U}$ which have exactly two basepoints. Let U denote the component of $\mathbb{U}$ which has four basepoints. Pick a diagram $(\Sigma_g, \boldsymbol{\alpha}_0, \boldsymbol{\beta}_0, \mathbf{w}_0, \mathbf{z}_0)$ for $\mathbb{U}_0$ as in the proof of Lemma 4.3.1 (in the case that $\mathbb{U}_0 = \varnothing$, simply pick a diagram for $(S^1 \times S^2)^{\#g}$ which has no basepoints). Now pick a point in the region of large area, and insert the diagram shown in Figure 4.3 (in the case that $\mathbb{U}_0 = \varnothing$, we insert the diagram as in Figure 4.3, but without the outer α– and β–curves). Let $\mathcal{H} = (\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ denote the resulting diagram.

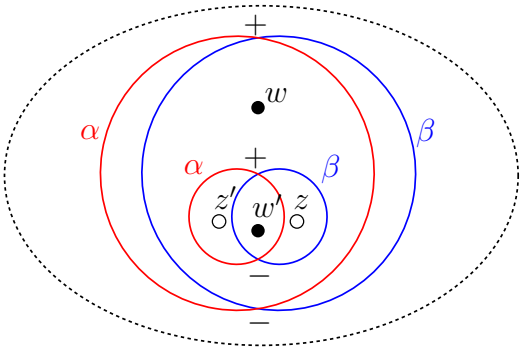

Figure 4.3: **Building a diagram for an unlink where one component has four basepoints** in the case that $\mathbb{U}_0 \neq \varnothing$. We insert this subdiagram into a diagram $(\Sigma_g, \boldsymbol{\alpha}_0, \boldsymbol{\beta}_0, \mathbf{w}_0, \mathbf{z}_0)$ for $((S^1 \times S^2)^{\#g}, \mathbb{U}_0)$, to get a diagram for $\mathbb{U} = \mathbb{U}_0 \cup U$. The component U has basepoints z, w, z' and w'. If $\mathbb{U}_0 = \varnothing$, we insert the above diagram, but with the outer α– and β–curves omitted.

The same argument that we used in the proof of Lemma 4.3.1, shows that sums of elements of the form $\Theta \cdot V_{\mathbf{z}}^J$ generate the subset of $\mathcal{C}FL^-(\mathcal{H}, \mathfrak{s})$ of top $\mathrm{gr}_{\mathbf{w}}$–homological degree. Also as before, we note that none of these elements are boundaries.

However, it is important to note that on the uncolored module $\mathcal{C}FL^-$, the element Θ is not a cycle. Instead $\partial\Theta = (V_z + V_{z'})\tau$, for some element τ, as is easily seen in Figure 4.3. On the other hand, the coloring $(\sigma_0, \mathfrak{P}_0)$ identifies V_z and $V_{z'}$, so Θ becomes a cycle in the colored chain complex. Hence

$$\mathcal{H}FL^-_{\max,\mathbf{w}}(\mathcal{H}, \sigma_0, \mathfrak{P}_0, \mathfrak{s})$$

is exactly

$$\mathcal{C}FL^-_{\max,\mathbf{w}}(\mathcal{H}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}) = \mathcal{C}FL^-_{\max,\mathbf{w}}(\mathcal{H}, \mathfrak{s})/(V_z - V_{z'}),$$

which is generated by elements of the form $V_{\mathbf{z}}^J \cdot \Theta$, and has the single relation that $V_z - V_{z'} = 0$. □

Remark 4.3.3. We will often make use of the existence of an element Θ distinguished by the property of generating a maximally graded subset. For example if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ represents an unlink in $(S^1 \times S^2)^{\#g}$, then the map F

defined by

$$F(\mathbf{x}) = F_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma},\mathfrak{s}}(\Theta \otimes \mathbf{x})$$

where $F_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma},\mathfrak{s}}$ counts holomorphic triangles, is a well defined, $\mathfrak{P}$–filtered map up to a choice of Θ. It's easy to see that if $\eta \in \mathcal{C}FL^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, then adding $\partial\eta$ to Θ induces an equivariant, filtered chain homotopy, and hence the map F is well defined up to equivariant, filtered chain homotopy as long as the element Θ is defined on homology in $\mathcal{H}FL^-$. We will use such triangle maps to define maps corresponding to surgery on a link, as well as adding bands between components of a link.

4.3.2 Quasi-stabilization

In Chapter 3, the author described maps associated to adding or removing adjacent basepoints, w and z along a link component. The operation had already been described in [MO10], but a proof that the operation yielded natural chain maps hadn't been completed. In Chapter 3, we colored the the curved complexes $\mathcal{C}FL^\circ$ to get actual chain complexes, which we denoted by CFL°_{UV}.

If w follows z, and $\sigma : \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ is a coloring which is extended by $\sigma' : \mathbf{w} \cup \mathbf{z} \cup \{w, z\} \to \mathfrak{P}$, and such that z is given the same color as the other $\mathbf{z}$–basepoint adjacent to w, then there are filtered chain maps

$$S^+_{w,z} : \mathcal{C}FL^\circ(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S^-_{w,z} : \mathcal{C}FL^\circ(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

which are well defined on the coherent chain homotopy type invariant. If z comes after w, there are maps $S^+_{z,w}$ and $S^-_{z,w}$ defined analogously. Note that the conditions on the coloring were automatic in Chapter 3, since in that setting all $V_{\mathbf{z}}$ variables on each link component were identified.

The above statement follows from Theorem A, though in Chapter 3, the result was stated for the version of link Floer homology CFL°_{UV}.

We will shortly describe the necessary modifications to the proof in Chapter 3. Additionally, we will describe maps

$$T^+_{w,z} : \mathcal{C}FL^\circ(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$T^-_{w,z} : \mathcal{C}FL^\circ(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

defined when w has the same color as the other $\mathbf{w}$ basepoint adjacent to z.

We first outline the quasi-stabilization operation.

Suppose we are given a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$. Given a choice of new, adjacent basepoints w and z, such that w comes after z, we let A denote the component of $\Sigma \setminus \boldsymbol{\alpha}$ containing the basepoints adjacent to w and z. We pick a point $p \in A \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$ and a choice of new α_s curve in A passing through p. We then form the "special connected sum" of $\mathcal{H}$ with a diagram $(S^2, \alpha_0, \beta_0, w, z)$ at a point $p_0 \in \alpha_0$ chosen, as in Figure 4.4.

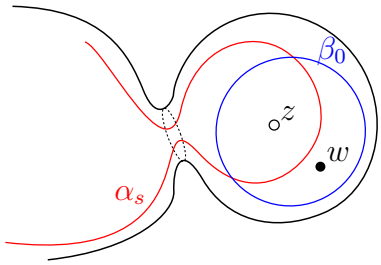

Figure 4.4: **A quasi-stabilized diagram.** We stretch the almost complex structure along the dashed curve.

Note that there are two new intersection points, which are distinguished by the relative Maslov grading. Of course we now have two Maslov gradings, namely $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. Let $\theta^+_{\mathbf{w}}$ denote the higher generator with respect to $\mathrm{gr}_{\mathbf{w}}$ and let $\theta^+_{\mathbf{z}}$ denote the higher generator with respect to $\mathrm{gr}_{\mathbf{z}}$. By inspection, we have that

$$\theta^{\pm}_{\mathbf{w}} = \theta^{\mp}_{\mathbf{z}}.$$

In Lemma 3.2.1, the quasi-stabilized differential is computed:

Proposition 4.3.4. *If $\mathcal{H}$ is a diagram for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and $\overline{\mathcal{H}}$ denotes the quasi-stabilized diagram, then for appropriately stretched almost complex structures, there is an identification of differentials*

$$\partial_{\overline{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H}} \end{pmatrix},$$

where w' and z' are basepoints on $\mathbb{L}$ adjacent to w and z. Here the matrix is in terms of the generators $\theta^{\pm}_{\mathbf{w}}$ (i.e. $\theta^{+}_{\mathbf{w}}$ is corresponds to the first row and column and $\theta^{-}_{\mathbf{w}}$ corresponds to the second row and column).

In particular, if $(\sigma, \mathfrak{P})$ is a coloring of the basepoints of $\mathbb{L}$ which sends z and z' to the same color, then we can define maps by the formulas

$$S^{+}_{w,z}(\mathbf{x}) = \mathbf{x} \otimes \theta^{+}_{\mathbf{w}}$$

and

$$S^{-}_{w,z}(\mathbf{x} \otimes \theta^{+}_{\mathbf{w}}) = 0, \qquad \text{and} \qquad S^{-}_{w,z}(\mathbf{x} \otimes \theta^{-}_{\mathbf{w}}) = \mathbf{x},$$

and these maps are chain maps.

We similarly, if $(\sigma, \mathfrak{P})$ is a coloring which gives w and w' the same color, then we can define chain maps

$$T^{+}_{w,z}(\mathbf{x}) = \mathbf{x} \otimes \theta^{+}_{\mathbf{z}}$$

and

$$T^{-}_{w,z}(\mathbf{x} \otimes \theta^{+}_{\mathbf{z}}) = 0, \qquad \text{and} \qquad T^{-}_{w,z}(\mathbf{x} \otimes \theta^{-}_{\mathbf{z}}) = \mathbf{x}.$$

Under these coloring assumptions, the proof of Theorem A carries over without modification to show that the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ each commute up to filtered, equivariant chain homotopy with the change of diagrams maps $\Phi_{\mathcal{H}\to\mathcal{H}'}$. The key computation is a challenging triangle map computation, which is stated in this paper as Theorem 4.5.3.

In the setting of Chapter 3, the maps $T^{\pm}_{w,z}$ correspond to $\Psi_z S^{+}_{w,z}$ and $S^{-}_{w,z}\Psi_z$. In our setting, we have the following:

Lemma 4.3.5. *Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link and w and z are new basepoints, in a component of $L \setminus (\mathbf{w} \cup \mathbf{z})$. Say that w and z are between w' and z', two basepoint of $\mathbb{L}$. If w and w' are given the same color, and z and z' are given the same color, then*

$$T^+_{w,z} \simeq \Psi_z S^+_{w,z}, \qquad \text{and} \qquad T^-_{w,z} \simeq S^-_{w,z} \Psi_z.$$

Similarly

$$S^+_{w,z} \simeq \Phi_w T^+_{w,z}, \qquad \text{and} \qquad S^-_{w,z} \simeq T^-_{w,z} \Phi_w.$$

Proof. This is an easy computation following from the definitions of the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$, as well as Φ_w and Ψ_z, which can be easily computed from Proposition 4.3.4. Note that the requirement on the coloring guarantees that all of the above maps are chain maps. □

4.3.3 The maps Φ_w and Ψ_z and relations amongst them

Two additional maps which appear in the link cobordism maps are the maps Φ_w and Ψ_z. These correspond to "formally differentiating the differential with respect to the variables U_w or V_z". They are defined by the formulas

$$\Phi_w(\mathbf{x}) = U_w^{-1} \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_w(\phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y},$$

and

$$\Psi_z(\mathbf{x}) = V_z^{-1} \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_z(\phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

According to Lemma 3.3.2, these maps commute with the change of diagrams maps, up to equivariant, filtered chain homotopy. We now summarize several important properties of the maps Φ_w and Ψ_z in the following lemmas.

By differentiating the expression in Lemma 4.2.9 once, with respect to a single variable, we have the following:

Lemma 4.3.6 (Lemma 3.3.1)**.** *If z is a basepoint, which is adjacent to basepoints w and w', then we have*

$$\partial\Psi_z + \Psi_z\partial = U_w + U_{w'}.$$

If w is a basepoint which is adjacent to z and z', then we have that

$$\partial\Phi_w + \Phi_w\partial = V_z + V_{z'}.$$

Note that Φ_w and Ψ_z are not chain maps for the trivial coloring. In fact, they induce chain maps iff the variables U_w and $U_{w'}$ are identified, or the variables V_z and $V_{z'}$ are identified, respectively. Notice that this corresponds exactly to the condition that the coloring of the basepoints agree with a coloring of the appropriate decorated surface in Figure 4.2.

By differentiating the expression in Lemma 4.2.9 twice, the commutator of various Φ_w and Ψ_z is computed in Lemmas 3.9.1 and 3.9.2. Those two lemmas can be summarized as follows:

Lemma 4.3.7. *If w, w', z and z' are basepoints on $\mathbb{L}$, we have*

$$\Phi_w\Phi_{w'} + \Phi_{w'}\Phi_w \simeq 0,$$

and

$$\Psi_z\Psi_{z'} + \Psi_{z'}\Psi_z \simeq 0.$$

We also have

$$\Phi_w\Psi_z + \Psi_z\Phi_w \simeq N(w, z),$$

where $N(w, z)$ is the number of components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which have both w and z as boundary.

We make the following remark (compare Corollary 2.13.18):

Remark 4.3.8. If $\mathcal{H}$ is a diagram for $(Y, \mathbb{L})$, then on the uncolored curved complex $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$, the maps Φ_w are nonequivariantly chain homotopic to zero. The chain homotopy is given by

$$\Phi_w = \frac{d}{dU_w} \circ \partial + \partial \circ \frac{d}{dU_w}.$$

The chain homotopy may not persist if we color and set some of the $U_{\mathbf{w}}$–variables equal to each other. Notice also that we work in the category of filtered, equivariant homomorphisms up to *filtered, equivariant* chain homotopies, and hence Φ_w is not in general the zero morphism, despite the above chain homotopy.

Similarly on the uncolored complex, we have

$$\Psi_z = \frac{d}{dV_z} \circ \partial + \partial \circ \frac{d}{dV_z}$$

on $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$. Again this chain homotopy may not persist after we color the complexes.

The maps Φ_w and Ψ_z are homotopy differentials (cf. Lemma 3.9.7):

Lemma 4.3.9. *The endomorphisms Φ_w and Ψ_z satisfy*

$$\Phi_w^2 \simeq 0, \qquad \textit{and} \qquad \Psi_z^2 \simeq 0,$$

through filtered, equivariant chain homotopies.

Proof. This is a purely algebraic. We will show that $\Phi_w^2 \simeq 0$; the proof for Ψ_z will be identical. Write $\partial = \sum_{k=0}^{\infty} \partial_k U_w^k$, where ∂_i does not involve U_w. Using this notation, we have

$$\Phi_w = \sum_{k=0}^{\infty} k \partial_k U^{k-1}, \qquad \text{and} \qquad \partial^2 = \sum_{k \geq 0} \sum_{n+m=k} \partial_n \partial_m U^{n+m}.$$

Applying Lemma 4.2.9 for any $k \geq 2$, we have that

$$\sum_{m+n=k} \partial_n \partial_m = 0. \tag{4.1}$$

Define

$$H = \sum_{k=0}^{\infty} \frac{k(k-1)}{2} \partial_k U_w^{k-2}.$$

One now simply computes that

$$\Phi_w^2 + (\partial H + H\partial) = \sum_{m,n \geq 0} \frac{(m+n)(m+n-1)}{2} \partial_n \partial_m U_w^{n+m-2} = \sum_{k=0}^{\infty} \frac{k(k-1)}{2} U_w^{k-2} \sum_{n+m=k} \partial_n \partial_m,$$

which we see vanishes, because the $k = 0$ and $k = 1$ terms vanish trivially and the terms for $k \geq 2$ vanish by Equation (4.1). □

Remark 4.3.10. The fact that $\Phi_w^2 \simeq 0$ and $\Psi_z^2 \simeq 0$ is a purely algebraic fact arising from our computation of ∂^2 in Lemma 4.2.9, and presents no real obstruction to a complex coming from Heegaard Floer homology. Suppose $\mathcal{R}$ is a polynomial ring (over $\mathbb{Z}$ or $\mathbb{Z}_2$) and C is a free module over $\mathcal{R}$ with a chosen free basis. Suppose E is an $\mathcal{R}$–equivariant endomorphism and U is a variable of the polynomial ring. If $E^2 = F_0 + F_1 U$, for endomorphisms F_0 and F_1 with no U terms (e.g. if $E^2 = 0$), then

$$(E')^2 + EH + HE = 0,$$

where E' is the formal derivative of E with respect to U, and H is a map defined like the chain homotopy in the previous lemma.

4.3.4 Relations between the quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ and the maps Φ_w and Ψ_z

In this section, we summarize some of the important relations satisfied by the quasi-stabilization maps and the maps Φ_w and Ψ_z. Similar statements can be found in Chapter 3.

By considering the quasi-stabilized differential in Proposition 4.3.4, one can derive the following:

Lemma 4.3.11 (Lemma 3.9.3)**.** *If w is not the only $\mathbf{w}$–basepoint on the link component containing w (so that $S^-_{w,z,\mathbf{w}}$ can be defined), we have that, as endomorphisms of $\mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ for a fixed coloring $(\sigma, \mathfrak{P})$,*

$$\Phi_w \simeq S^+_{w,z} S^-_{w,z}, \qquad \text{or} \qquad \Phi_w \simeq S^+_{z,w} S^-_{z,w}$$

for z an adjacent $\mathbf{z}$–basepoint.

Proof. This can be proven by explicit examination of the quasi-stabilized differential in Proposition 4.3.4, as well as the formulas for $S^{\pm}_{w,z}$. □

Note that Φ_w is defined as an endomorphism of a particular complex. Nonetheless the maps $S^+_{w,z}S^-_{w,z}$ can also be defined as maps between complexes when w is given a different coloring on the first complex than the second. In the latter case, a map Φ_w can no longer be defined using the formula we used above.

Symmetrically, we have the following:

Lemma 4.3.12. *Suppose z is not the only $\mathbf{z}$–basepoint on the link component containing z (so that $T^-_{w,z}$ can be defined), then as endomorphisms of $\mathcal{C}FL^\circ(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$, for a fixed coloring $(\sigma,\mathfrak{P})$, we have that*

$$\Psi_z = T^+_{w,z}T^-_{w,z}$$

if w comes immediately after z. If w comes immediately before z, we have

$$\Psi_z = T^+_{z,w}T^-_{z,w}.$$

Proof. The proof is identical to the proof of Lemma 4.3.11. □

Additionally we have the following:

Lemma 4.3.13. *The quasi-stabilization maps satisfy*

$$S^-_{w,z}S^+_{w,z} \simeq 0, \qquad T^-_{w,z}T^+_{w,z} \simeq 0, \qquad T^-_{w,z}S^+_{w,z} \simeq \mathrm{id}, \qquad \text{and} \qquad S^-_{w,z}T^+_{w,z} \simeq \mathrm{id},$$

as well as

$$S^+_{w,z}T^-_{w,z} + T^+_{w,z}S^-_{w,z} \simeq \mathrm{id}\,.$$

Proof. These are all just easy applications of the formulas involved. □

Remark 4.3.14. The latter two formulas should be thought of as an analog to Lemma 3.9.6. In that context, we had $S^-_{w,z}\Psi_z S^+_{w,z} \simeq \mathrm{id}$. In our context, $T^-_{w,z} \simeq S^-_{w,z}\Psi_z$. Note that if add basepoints (w,z) in between w' and z', and w' is colored the same as w, and z' is colored the same as z, then using Lemma 4.3.5, we see that indeed we still have $S^-_{w,z}\Psi_z S^+_{w,z} \simeq \mathrm{id}$.

In Chapter 3, we show the following result about quasi-stabilization maps commuting:

Proposition 4.3.15 (Theorem 3.8.1)**.** *The quasi-stabilization maps $S^{\pm}_{w,z}$ commute. That is if (w,z) and (w',z') are disjoint pairs of basepoints, then*

$$S^{\circ}_{w,z}S^{\circ'}_{w',z'} \simeq S^{\circ'}_{w',z'}S^{\circ}_{w,z},$$

for any choice of $\circ, \circ' \in \{+,-\}$. Commutation still holds if both of the pairs of basepoints (w,z) and (w',z') appear with the opposite ordering (i.e. if z appears after w, and we use the maps $S^{+}_{z,w}$, instead).

The proof can be found in Section 3.8. The proof in Chapter 3 of the previous theorem also carries over without change if we replace both S–quasi-stabilizations by T–quasi-stabilizations:

Proposition 4.3.16. *The quasi-stabilization maps $T^{\pm}_{w,z}$ commute. That is if (w,z) and (w',z') are disjoint pairs of basepoints, then*

$$T^{\circ}_{w,z}T^{\circ'}_{w',z'} \simeq T^{\circ'}_{w',z'}T^{\circ}_{w,z},$$

for any choice of $\circ, \circ' \in \{+,-\}$. Commutation still holds if either or both of the pairs of basepoints (w,z) and (w',z') appears with the opposite ordering (i.e. if z appears after w, and we use the maps $T^{+}_{z,w}$, instead).

The previous two results don't immediately tell us whether $S^{\pm}_{w,z}$ and $T^{\pm}_{w',z'}$ commute. In general they will not always commute (see Figure 4.5 and Lemma 4.3.18), but we have the following:

Proposition 4.3.17. *Suppose (w,z) are (w',z') are disjoint pairs of adjacent basepoints, with w following z and w' following z'. If w' is not adjacent to z, then*

$$T^{\circ}_{w,z}S^{\circ'}_{w',z'} \simeq S^{\circ'}_{w',z'}T^{\circ}_{w,z},$$

for any choice of $\circ, \circ' \in \{+,-\}$. The same result holds if z follows w and z' follows w', with $T^{\circ}_{w,z}$ changed to $T^{\circ}_{z,w}$ and $S^{\circ'}_{w',z'}$ changed to $S^{\circ'}_{z',w'}$, as long as z and w' are not adjacent.

Proof. We will consider only the the first statement, since the second follows by a straightforward modification. We refer the reader to Section 3.8, since our argument here will be a straightforward modification. We take a diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for the unstabilized link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and let $p, p' \in \Sigma \setminus \boldsymbol{\alpha}$ be the two points at which we perform the quasi-stabilization. We form the quasi-stabilized diagram $\overline{\mathcal{H}}$, as in Figure 3.10. Note that the formulas for the two quasi-stabilization maps obviously commute, but it's not clear whether we can use the same almost complex structure to compute both.

There are two cases to consider. The first case is that p and p' are in separate components of $\Sigma \setminus \boldsymbol{\alpha}$, and the second is that p and p' occur in the same component of $\Sigma \setminus \boldsymbol{\alpha}$. The case that p and p' are in separate components of $\Sigma \setminus \boldsymbol{\alpha}$ corresponds to (w, z) and (w', z') being in separate components of $L \setminus (\mathbf{w} \cup \mathbf{z})$. There are two necks to stretch, and we let $\mathcal{J}$ denote gluing data to stretch on both necks. According to Lemma 3.8.2, for all sufficiently large T_1, T_1', T_2, and T_2', the change of almost complex structure map

$$\Phi_{\mathcal{J}(T_1,T_1')\to\mathcal{J}(T_2,T_2')} = \mathrm{id},$$

under the obvious identification of the two complexes. Hence in this case, the maps $T^{\pm}_{w,z} S^{\pm}_{w',z'} \simeq S^{\pm}_{w',z'} T^{\pm}_{w,z}$ obviously commute.

We now consider the second case, where p and p' are in the same component of $\Sigma \setminus \boldsymbol{\alpha}$. This case corresponds to (w, z) and (w', z') being adjacent on $L \setminus (\mathbf{w} \cup \mathbf{z})$. Since w' is not adjacent to z, we note that this implies that the basepoints appear in the following order z, w, z' then w'. The change of almost complex structure maps $\Phi_{\mathcal{J}(T_1,T_1')\to\mathcal{J}(T_2,T_2')}$ are computed in Lemma 3.8.3. By the proof of that Lemma, if $F = \Phi_{\mathcal{J}(T_1,T_1')\to\mathcal{J}(T_2,T_2')}$ then

$$\begin{aligned}
F(\mathbf{x} \times \theta^-_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^+) &= \mathbf{x} \times \theta^-_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^+ \\
F(\mathbf{x} \times \theta^-_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^-) &= \mathbf{x} \times \theta^-_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^- + C \cdot \mathbf{x} \times \theta^+_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^+ \\
F(\mathbf{x} \times \theta^+_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^+) &= \mathbf{x} \times \theta^+_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^+ \\
F(\mathbf{x} \times \theta^+_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^-) &= \mathbf{x} \times \theta^+_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^-,
\end{aligned}$$

for some C (which is not independent of T_i and T_i').

It is now easy (though somewhat tedious) algebra to check that the maps commute. □

The strategy used in the previous proof fails to show, for example, that $S^-_{w,z}T^+_{w',z'}$ and $T^+_{w',z'}S^-_{w,z}$ are chain homotopic if z' and w are adjacent and we have the following ordering of basepoints: z, w, z' then w'. Using the interpretation of the quasi-stabilization maps in terms of the dividing sets shown in Figure 4.2, we note that we should definitely not expect these maps to commute, as the dividing sets for the two compositions are clearly different. In Figure 4.5 we show the two dividing sets of these maps on the left. Instead, we prove the following:

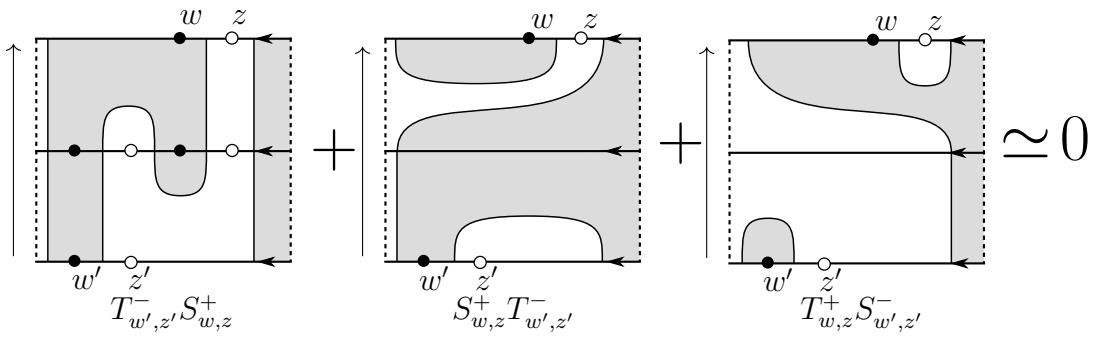

Figure 4.5: **An example of quasi-stabilization maps not commuting.** Shown is the relation proven in Lemma 4.3.18 when commutation of quasi-stabilization fails. The two dividing sets for the compositions $S^-_{w,z}T^+_{w',z'}$ and $T^+_{w',z'}S^-_{w,z}$ are shown on the left. Proposition 4.3.17 doesn't imply composition of these maps. The formulas for the two maps look like they should commute, but when we change between almost complex structures which can compute the two maps, we cannot prevent Maslov index zero holomorphic disks from appearing which prevent the two quasi-stabilization maps from commuting.

Lemma 4.3.18. *If w', z', w and z are consecutive basepoints on a link component that also contains other basepoints, and the basepoints appear in the order z, w, z' then w', then*

$$T^-_{w',z'}S^+_{w,z} + S^+_{w,z}T^-_{w',z'} + T^+_{w,z}S^-_{w',z'} \simeq 0.$$

The interpretation in terms of surfaces with divides is shown in Figure 4.5.

Proof. We start with the relation

$$T^+_{w',z'}S^-_{w',z'} + S^+_{w',z'}T^-_{w',z'} + \mathrm{id} \simeq 0,$$

from Lemma 4.3.13 and then multiply on the left by $T^-_{w',z'}S^+_{w,z}$ to get

$$T^-_{w',z'}S^+_{w,z}T^+_{w',z'}S^-_{w',z'} + T^-_{w',z'}S^+_{w,z}S^+_{w',z'}T^-_{w',z'} + T^-_{w',z'}S^+_{w,z} \simeq 0.$$

We now can manipulate the terms to see that

$$\begin{aligned}
0 &\simeq T^-_{w',z'}S^+_{w,z}T^+_{w',z'}S^-_{w',z'} + T^-_{w',z'}S^+_{w,z}S^+_{w',z'}T^-_{w',z'} + T^-_{w',z'}S^+_{w,z} \\
&\simeq T^-_{w',z'}S^+_{w,z}T^+_{w',z'}S^-_{w',z'} + T^-_{w',z'}S^+_{w',z'}S^-_{w,z}T^+_{w',z'} + T^-_{w',z'}S^+_{w,z} && \text{(Proposition 4.3.15)} \\
&\simeq T^-_{w',z'}S^+_{w,z}T^+_{w',z'}S^-_{w',z'} + S^-_{w,z}T^+_{w',z'} + T^-_{w',z'}S^+_{w,z} && \text{(Lemma 4.3.13)} \\
&\simeq T^-_{w',z'}\Phi_w T^+_{w,z}T^+_{w',z'}S^-_{w',z'} + S^-_{w,z}T^+_{w',z'} + T^-_{w',z'}S^+_{w,z} && \text{(Lemma 4.3.5)} \\
&\simeq T^-_{w',z'}\Phi_w T^+_{w',z'}T^+_{w,z}S^-_{w',z'} + S^-_{w,z}T^+_{w',z'} + T^-_{w',z'}S^+_{w,z} && \text{(Proposition 4.3.16)} \\
&\simeq T^+_{w,z}S^-_{w',z'} + S^-_{w,z}T^+_{w',z'} + T^-_{w',z'}S^+_{w,z} && \text{(Lemmas 4.3.5, 4.3.13),}
\end{aligned}$$

completing the proof. □

We now give simple repackaging of Lemma 3.9.4 and Lemma 3.9.5.

Suppose that A is a subarc of L, between two $\mathbf{w}$–basepoints. We do not exclude the case that A consists of a single point (which we treat as the empty arc), or an entire link component of L. We define

$$\Psi_A = \sum_{z \in A} \Psi_z.$$

Note that Ψ_z is a chain map iff the endpoints of A share the same color.

Lemma 4.3.19. *If w and z are new, adjacent basepoints, we have*

$$\Psi_A S^{\pm}_{w,z} + S^{\pm}_{w,z}\Psi_A \simeq 0.$$

Similarly

$$\Psi_A T^{\pm}_{w,z} + T^{\pm}_{w,z}\Psi_A \simeq 0.$$

Proof. Let $\mathfrak{P} = \mathbf{w} \cup \mathbf{z} \cup \{w \cup z\}$, and let $\sigma : (\mathbf{w} \cup \mathbf{z}) \to \mathfrak{P}$ denote the natural inclusion, and let $\sigma' : \mathbf{w} \cup \mathbf{z} \cup \{w \cup z\} \to \mathfrak{P}$ denote the identity. Let $\mathcal{H}^-$ denote a diagram for $(L, \mathbf{w}, \mathbf{z})$, and let $\mathcal{H}^+$ denote it's quasi-stabilization resulting from adding the basepoints w and z. Consider the colored complexes $\mathcal{C}FL^\circ(\mathcal{H}^-, \sigma, \mathfrak{P}, \mathfrak{s})$ and $\mathcal{C}FL^\circ(\mathcal{H}^+, \sigma', \mathfrak{P}, \mathfrak{s})$. We note that using Proposition 4.3.4, we can write

$$\partial_{\mathcal{H}^\pm} S^\pm_{w,z} + S^\pm_{w,z} \partial_{\mathcal{H}^\mp} + (V_z + V_{z'})F = 0, \tag{4.2}$$

where F is a map which has does not involve the variables in $U_{\mathfrak{P}}$. Here z' is the $\mathbf{z}$–basepoint adjacent to z. We view these as matrices with coefficients in $\mathbb{Z}_2[U_{\mathfrak{P}}]$. Let

$$D_A = \sum_{z \in A} \frac{d}{dV_z}.$$

Noting that $\frac{d}{dV_z}\partial_{\mathcal{H}^-} = 0$, we have that $D_A(\partial_{\mathcal{H}^\circ}) = \Psi_A$ for both $\circ = +$ and $\circ = -$. We also note that on the quasi-stabilized link, either z and z' are in A, or neither are in A, as w is not an endpoint of A. Hence

$$D_A((V_z + V_{z'})F) = 0,$$

as is easily computed using the Leibniz rule, separately, in the case that $z, z' \in A$ and the case that $z, z' \notin A$, using also the fact that $D_A F = 0$, as the matrix expression for F involves none of the variables in A, upon direct inspection. Thus applying D_A to Equation (4.2) yields

$$\Psi_A S^\pm_{w,z} + S^\pm_{w,z} \Psi_A \simeq 0.$$

The statement involving $T^\pm_{w,z}$ follows from similar considerations. □

We can prove the analogous statement for the maps Φ_w. If A is an arc on $(L, \mathbf{w}, \mathbf{z})$ between two $\mathbf{z}$–basepoints, define

$$\Phi_A = \sum_{w \in A} \Phi_w.$$

The map Φ_A is a chain map iff the $\mathbf{z}$–basepoints on the ends share the same color. Analogously to the previous lemma, we have the following:

Lemma 4.3.20. *If A is an arc on L between two $\mathbf{z}$–basepoints of $(L, \mathbf{w}, \mathbf{z})$, and w and z are two new, adjacent basepoints on L, then*

$$S^{\pm}_{w,z}\Phi_A + \Phi_A S^{\pm}_{w,z} \simeq 0,$$

and

$$T^{\pm}_{w,z}\Phi_A + \Phi_A T^{\pm}_{w,z} \simeq 0.$$

The proof is identical to the proof of Lemma 4.3.19.

We now highlight several useful special cases of the previous lemmas:

Lemma 4.3.21 (Lemma 3.9.4)**.** *If w and w' are distinct basepoints, w follows z, and the $\mathbf{z}$–basepoints adjacent to w' have the same color, then*

$$S^{\pm}_{w,z}\Phi_{w'} + \Phi_{w'} S^{\pm}_{w,z} \simeq 0.$$

A similar statement holds for $S^{\pm}_{z,w}$ if z comes after w.

If z' is not adjacent to a basepoint w, and z comes immediately before w, then

$$S^{\pm}_{w,z}\Psi_{z'} + \Psi_{z'} S^{\pm}_{w,z} \simeq 0.$$

Proof. The first relation is obtained by picking an arc A between the two $\mathbf{z}$–basepoints adjacent to w', before quasi-stabilizing. Applying Lemma 4.3.19 yields the result immediately in the case that A does not contain the basepoints w and z. In the case that A contains the basepoints w and z, we note that the result follows from the additional observation that $S^{-}_{w,z}\Phi_w \simeq 0$ and $\Phi_w S^{+}_{w,z} \simeq 0$.

The second claim follows from the same strategy. □

Lemma 4.3.22 (Lemma 3.9.5)**.** *If z and z' are the basepoints immediately adjacent to w, and w comes after z (and $z \neq z'$), then we have*

$$S^{+}_{w,z}\Psi_{z'} \simeq (\Psi_z + \Psi_{z'}) S^{+}_{w,z}$$

and

$$S^-_{w,z}(\Psi_z + \Psi_{z'}) \simeq \Psi_{z'} S^-_{w,z}.$$

A similar statement holds of the maps $S_{z,w}$, in the case that z comes after w.

Proof. This follows immediately from Lemma 4.3.19, picking A to be the arc on between the $\mathbf{w}$–basepoints on the unstabilized link adjacent to z'. □

4.3.5 Quasi-stabilization and basepoint moving maps I

In this section we consider the relation between the quasi-stabilization maps, $S^{\pm}_{w,z}, S^{\pm}_{z,w}, T^{\pm}_{w,z}$ and $T^{\pm}_{z,w}$, and certain basepoint moving maps. We will later define four classes of simple dividing sets, which we will denote by I^+, I^-, II^+ and II^-, as shown in Figure 4.28. For a given dividing set as in Figure 4.28, there are two maximal identifications of a subset of the incoming boundary with a subset of the outgoing boundary, respecting the divides. We now prove several results which will imply that cobordism maps we will define in Subsection 4.11.2 do not depend on the choice of these two identifications.

4.3.5.1 Relations for the maps $S^{\pm}_{w,z}$

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link, and $z' \in \mathbf{z}$.

Now suppose that w and z are two adjacent basepoints on L which are not in $\mathbf{w} \cup \mathbf{z}$, such that, as one traverses the link according to its orientation, these basepoints appear in the order z', w then z. Let ρ be the diffeomorphism

$$\rho_* : (Y, L, \mathbf{w}, \mathbf{z}) \to (Y, L, \mathbf{w}, (\mathbf{z} \setminus \{z\}) \cup \{z'\})$$

such that $\rho(z') = z$, and is fixed outside of a neighborhood of the arc between z' and z (note that $w, z \notin \mathbf{w} \cup \mathbf{z}$).

Lemma 4.3.23. *We have a chain homotopy equivalence*

$$S^+_{w,z} \simeq S^+_{z',w} \rho_*.$$

Proof. Consider the diagram shown in Figure 4.6.

$(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z} \setminus \{z'\}) \quad \#$

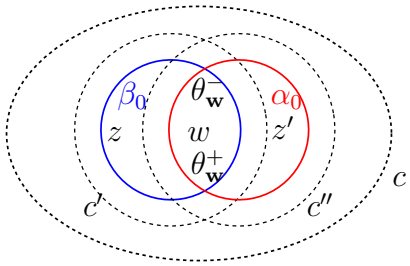

Figure 4.6: **The diagram we consider in Lemma 4.3.23.** We must consider the change of almost complex structure map as we change from an α–quasi-stabilization (stretched along c'') and a β–quasi-stabilization (stretched along c').

If we stretch along the circle c', we get a complex structure which can be used to compute $S^+_{w,z}$. If we change the almost complex structure so that it is stretched along c'', we get a diagram which we can use to $S^+_{z',w}$. On the diagram occurring before quasi-stabilization (with only the z' basepoint), the map ρ_* has the effect of just moving z' to z along the surface. Thus it is sufficient to compute the effect of changing the almost complex structure by switching from a β–quasi-stabilization to an α–quasi-stabilization.

Let $\mathcal{H}_0 = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \setminus\{z'\})$ denote the partial Heegaard diagram. The partial diagram $\mathcal{H}_0$ has a distinguished point p, the connected sum point. Pick a generic almost complex structure J_s on $\mathcal{H}_0$, i.e., an almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$, which is split in a neighborhood of p. In Definition 2.5.1 we define "gluing data" $\mathcal{J}$ to be a collection of data, including a choice of almost complex structures J_s on $\Sigma \times [0,1] \times \mathbb{R}$ and J'_s on $S^2 \times [0,1] \times \mathbb{R}$, as well as some auxiliary data which allow one to construct an almost complex structure $\mathcal{J}(T)$, for a choice of necklength T (see Subsection 2.5.1 for more details). As $T \to \infty$, the almost complex manifolds

$$((\Sigma \# S^2) \times [0,1] \times \mathbb{R}, \mathcal{J}(T))$$

approach

$$((\Sigma \vee S^2) \times [0,1] \times \mathbb{R}, J_s \vee J'_s).$$

Performing the obvious adaptation, we can consider a collection of gluing data $\mathcal{J}_{c'}$ for

stretching along c and c'. We let $\mathcal{J}_{c,c'}(T,T')$ denote the resulting almost complex structure, with neck lengths T and T' inserted along c and c'' respectively. We define a collection of gluing data $\mathcal{J}_{c,c''}$ for stretching along c and c'' analogously.

We make the following claim: *if T is sufficiently large, then for any $T', T'' > 0$ we can pick a dynamic almost complex structure $\widehat{J}$ such that the change of almost complex structure map, computed with the dynamic almost complex structure $\widehat{J}$, takes the form*

$$\Phi_{\mathcal{J}_{c,c'}(T,T')\to\mathcal{J}_{c,c''}(T,T'')} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}. \tag{4.3}$$

Here the first component corresponds to the $\theta^+_{\mathbf{w}}$ generator and the second corresponds to the $\theta^-_{\mathbf{w}}$ generator. We claim the choice of T can be made independent of T' and T''.

Note that, assuming this, the lemma statement follows immediately, since with respect to $\mathcal{J}_{c,c'}(T,T')$ (for sufficiently large T') the map $S^+_{w,z}$ takes the form

$$\mathbf{x} \mapsto \mathbf{x} \times \theta^+_{\mathbf{w}},$$

and with respect to $\mathcal{J}_{c,c''}(T,T'')$ the map $S^+_{z',w}$ takes the same form. Assuming the change of almost complex structure map $\Phi_{\mathcal{J}_{c,c'}(T,T')\to\mathcal{J}_{c,c''}(T,T'')}$ takes the form claimed in Equation (4.3), we see that $S^+_{w,z}$ and $S^+_{z',w}\rho_*$ are equal.

Hence it is sufficient to show that $\Phi_{\mathcal{J}_{c,c'}(T,T')\to\mathcal{J}_{c,c''}(T,T'')}$ takes the form in Equation (4.3) for sufficiently large T, independent of T' and T''.

We argue by the following weak limits argument. First let us emphasize that J_s is a cylindrical almost complex structure, and hence admits no nonconstant Maslov index zero disks, by transversality. Suppose that there were representatives u_{T_i}, for dynamic almost complex structures $\widehat{\mathcal{J}}_i$, which are cylindrical on a subset $(\Sigma \setminus B) \cup (S^1 \times [-T_i, T_i]) \subseteq \Sigma$, for arbitrarily large choices of T_i. Without loss of generality suppose that u_{T_i} all represent the same homology class ϕ. Write

$$\phi = \phi_\Sigma \# \phi_{S^2} \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y),$$

where ϕ_Σ is a homology class on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and ϕ_{S_2} is a homology class on (S^2, α_0, β_0). One computes easily that

$$\mu(\phi) = \mu(\phi_1) + \Delta_{\mathbf{w}}(x, y) + 2n_w(\phi_{S^2})$$

where $\Delta_{\mathbf{w}}(x, y)$ denotes the drop in grading from x to y, with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. By our assumptions on $\widehat{\mathcal{J}}_i$, we can extracting a weak limit to a broken curve representing ϕ_1, and hence we conclude that $\mu(\phi_1) \geq 0$. Hence for $\Delta(x, y) \geq 0$ (i.e. the diagonal and lower diagonal entries) we conclude that for a representative to exist, we must have

$$\mu(\phi) = \Delta_{\mathbf{w}}(x, y) = n_w(\phi_{S^2}) = 0.$$

Since ϕ has a broken representative for J_s, a cylindrical almost complex structure, we conclude by transversality that ϕ must be a constant disk. On the other hand, constant disks are always counted. Since the lower left entry of the transition matrix corresponds to $\Delta_{\mathbf{w}}(x, y) = 1$, no disks are counted by that component, since the existence of a broken representative of ϕ_1 would force

$$\mu(\phi) = \mu(\phi_1) + \Delta(x, y) + 2n_w(\phi_{S^2}) \geq 1.$$

Hence the entries of the transition matrix are as claimed. □

We now consider the relation that will be useful for type I^- cylindrical cobordisms. Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with a basepoint $z_0 \in \mathbf{z}$. Let z, w, w', z' be new basepoints in the two components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to z_0. Assume that these basepoints appear in the order z, w, z_0, w' then z', according to the orientation of the link. Note that there is a well defined diffeomorphism

$$\rho' : (Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}) \to (Y, L, \mathbf{w} \cup \{w'\}, \mathbf{z} \cup \{z'\})$$

which is the identity outside of a neighborhood of the arc from z' to z, and satisfies

$$\rho'(z) = z_0, \qquad \rho'(w) = w', \qquad \text{and} \qquad \rho'(z_0) = z'.$$

Analogously to Lemma 4.3.23, we have the following:

Lemma 4.3.24. *We have that*

$$S^-_{w,z} \simeq S^-_{z',w'}\rho'_*.$$

Proof. The proof follows easily from the same computation of the change of almost complex structure maps performed in Lemma 4.3.23 □

4.3.5.2 Relations for the maps $T^{\pm}_{w,z}$

We now wish to prove the analogs of Lemmas 4.3.23 and 4.3.24, but for the maps $T^{\pm}_{w,z}$.

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with a basepoint $w' \in \mathbf{w}$. Suppose that z and w are two basepoints in $L \setminus (\mathbf{w} \cup \mathbf{z})$ which appear in the order w', z then w, with no other basepoints in between these basepoints. There is a well defined diffeomorphism

$$\rho : (Y, L, \mathbf{w}, \mathbf{z}) \to (Y, L, (\mathbf{w} \setminus \{w'\}) \cup \{w\}, \mathbf{z})$$

which is the identity outside of a neighborhood of the arc from w to w', and satisfies

$$\rho(w') = w.$$

We have the following:

Lemma 4.3.25. *We have that*

$$T^+_{w,z} \simeq T^+_{z,w'}\rho_*.$$

Proof. As with the proof of Lemma 4.3.23, the key is a change of almost complex structures map. The analogous diagram is shown in Figure 4.7.

Again we just need to compute the matrix form of each map and show they agree. Both $T^+_{w,z}$ and $T^+_{z,w'}$ take the matrix form $\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$, where now we put the generator of higher $\mathrm{gr}_{\mathbf{z}}$ grading as the first generator. Note that if $\phi = \phi_\Sigma \# \phi_{S^2} \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y)$, is a homology disk then

$$\mu(\phi) = \mu(\phi_\Sigma) + 2n_{z'}(\phi_{S_2}) + \Delta_{\mathbf{z}}(x, y).$$

Thus to show the equality of the lemma statement, we need to compute the change of almost complex structures map and show that it preserves $\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$. The weak limits argument from

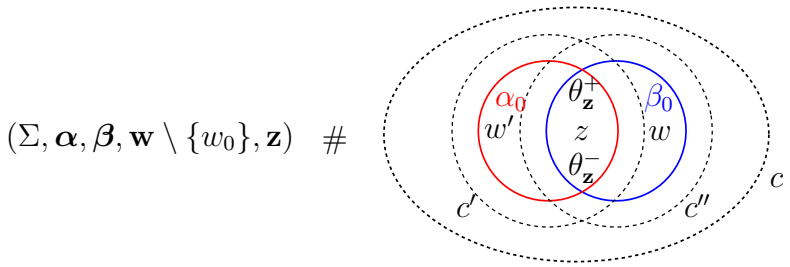

Figure 4.7: **The diagram we consider in Lemma 4.3.25.**

Lemma 4.3.23 now applies to show that the relevant change of almost complex structure map is now

$$\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix},$$

which clearly preserves the map $\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right)$.[3] □

We now consider the analogous statement for model cylindrical cobordisms of type II^-. Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with a basepoint $w_0 \in \mathbf{w}$. Let w', z', z, w be new basepoints in the two components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to z_0. Assume that these basepoints appear in the order w', z', w_0, z then w, according to the orientation of the link. Note that there is a well defined diffeomorphism

$$\rho' : (Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}) \to (Y, L, \mathbf{w} \cup \{w'\}, \mathbf{z} \cup \{z'\})$$

which is the identity outside of a neighborhood of the arc from w' to w, and satisfies

$$\rho'(w) = w_0, \qquad \rho'(z) = z', \qquad \text{and} \qquad \rho'(w_0) = w'.$$

Analogously to Lemma 4.3.24, we have the following:

Lemma 4.3.26. *We have that*

$$T^-_{w,z} \simeq T^-_{z',w'} \rho'_*.$$

[3]Note that this the matrix decomposition for the change of almost complex structure map is now in the $\mathrm{gr}_{\mathbf{z}}$–grading, whereas in Lemma 4.3.23 it was in the $\mathrm{gr}_{\mathbf{w}}$–grading. The basepoint configuration has also changed, however.

Proof. The proof is a straightforward modification of Lemma 4.3.25. □

4.3.6 Quasi-stabilization and basepoint moving maps II

In Chapter 3, we proved that if (w, z) and (w', z') are adjacent pairs of basepoints, then the map induced by the diffeomorphism ϱ_0 resulting from moving the pair (w, z) to (w', z') is chain homotopic to

$$(\varrho_0)_* \simeq S^-_{w,z}\Psi_{z'}S^+_{w',z'}. \tag{4.4}$$

A proof can be found in Lemma 3.10.1. In this section, we will give a new proof of the above formula, and prove some related results.

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link and $z \in \mathbf{z}$. Let w and z' be new basepoints, in one of the components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to z. There is a well defined diffeomorphism $\phi_{z\to z'}$ obtained from just moving z to z' and fixing all the other basepoints. Suppose that the basepoints appear in the order z, w, then z'.

Lemma 4.3.27. *The map induced by $\phi_{z\to z'}$ takes the form*

$$(\phi_{z\to z'})_* \simeq T^-_{w,z}S^+_{z',w} \simeq S^-_{w,z}T^+_{z',w}.$$

Proof. Let us first consider $T^-_{z,w}S^+_{w,z'}$. By Lemma 4.3.23, we have that

$$S^+_{w,z}\phi_{z\to z'} \simeq S^+_{z',w}.$$

We now just compose with $T^-_{w,z}$ to get

$$T^-_{w,z}S^+_{w,z}\phi_{z\to z'} \simeq T^-_{w,z}S^+_{z',w}.$$

Noting that $T^-_{w,z}S^+_{w,z} \simeq \mathrm{id}$ by Lemma 4.3.13, the first equivalence now follows. The second equality follows similarly, using Lemma 4.3.25 instead. □

The above lemma also holds in the case that the basepoints z, w and z' appear in the opposite order, as long as we replace the maps $T^+_{w,z}$ with $T^+_{z,w}$ and $S^-_{z',w}$ with $S^-_{w,z'}$. Analogously, can consider moving a $\mathbf{w}$ basepoint. Suppose $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link with $w \in \mathbf{w}$,

and let w' and z' be new basepoints in a component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to w, such that the order of the basepoints is w, z' then w'. Let $\phi_{w \to w'}$ denote the diffeomorphism induced by moving w to w' and fixing all of the other basepoints. We have the following:

Lemma 4.3.28. *The map induced by $\phi_{w \to w'}$ takes the form*

$$(\phi_{w \to w'})_* \simeq S^-_{z',w} T^+_{w',z'} \simeq T^-_{z',w} S^+_{w',z'}.$$

Proof. Applying Lemma 4.3.25 (note that our labeling of the basepoints is slightly different here), shows that

$$T^+_{w',z'} \simeq T^+_{z',w} (\phi_{w \to w'})_*.$$

Postcomposing with $S^-_{z',w}$ yields

$$S^-_{z',w} T^+_{w',z'} \simeq (\phi_{w \to w'})_*,$$

yielding the first equivalence. The second equivalence follows from a similar argument, using Lemma 4.3.23. □

We now give a new proof of Equation (4.4), for curved link Floer homology.

Lemma 4.3.29. *Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with w and z adjacent basepoints in $\mathbf{w}$ and $\mathbf{z}$, such that w follows z. Suppose that w' and z' are new adjacent basepoints which are in the components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ immediately following w. Let ϱ_0 denote the diffeomorphism obtained by moving the pair (w, z) to (w', z'), while fixing all the other basepoints. If z has the same color as the other $\mathbf{z}$ basepoint adjacent to w, then*

$$(\varrho_0)_* \simeq S^-_{w,z} \Psi_{z'} S_{w',z'}.$$

If w has the same color as the other $\mathbf{w}$ basepoint adjacent to z, then

$$(\varrho_0)_* \simeq T^-_{w,z} \Phi_w T^+_{w,z}.$$

For any coloring, we have

$$(\varrho_0)_* \simeq S^-_{w,z} T^+_{w',z'} \simeq T^-_{w,z} S^+_{w',z'}.$$

Proof. We first consider the case of arbitrary colorings. We compute, using our previous computations of the basepoint moving maps in this section

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq S^-_{w,z}T^+_{z',w}(\phi_{w\to w'})_* && \text{(Lemma 4.3.27)} \\ &\simeq S^-_{w,z}T^+_{w',z'} && \text{(Lemma 4.3.25).}\end{aligned}$$

Analogously, we can also write

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq T^-_{w,z}S^+_{z',w}(\phi_{w\to w'})_* && \text{(Lemma 4.3.27)} \\ &\simeq T^-_{w,z}S^+_{w',z'} && \text{(Lemma 4.3.23).}\end{aligned}$$

We now consider the case where we have an extra assumption on the colorings of the basepoints. The extra assumptions on the colorings of the basepoints are necessary so that the maps Ψ_z or Φ_w can be defined. With the appropriate assumptions on the colorings, we have by Lemmas 4.3.27 and 4.3.28 that

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq S^-_{w,z}T^+_{z',w}T^-_{z',w}S^+_{w',z'} && \text{(Lemmas 4.3.27 and 4.3.28)} \\ &\simeq S^-_{w,z}\Psi_{z'}S^+_{w',z'} && \text{(Lemma 4.3.12)}\end{aligned}$$

and similarly

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq T^-_{w,z}S^+_{z',w}S^-_{z',w}T^+_{w',z'} && \text{(Lemmas 4.3.27 and 4.3.28)} \\ &\simeq T^-_{w,z}\Phi_w T^+_{w',z'} && \text{(Lemma 4.3.11),}\end{aligned}$$

completing the proof. □

4.3.7 Disk-stabilization maps $S^{\pm}_{\mathbb{U},D}$

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link in Y^3 and that $\mathbb{U} = (U, w, z)$ is an unknot which is unknotted from all of the components L. Suppose in addition that D is a choice of embedded disk in $Y \setminus L$ which has boundary U. Write $\mathbb{L} \cup \mathbb{U}$ for the link $(L \cup U, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$. Given a coloring σ of $\mathbb{L}$ which is extend by a coloring σ' of $\mathbb{L} \cup \mathbb{U}$, we will now define maps

$$S^{+}_{\mathbb{U},D} : \mathcal{C}FL^{\circ}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^{\circ}(Y, \mathbb{L} \cup \mathbb{U}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S^{-}_{\mathbb{U},D} : \mathcal{C}FL^{\circ}(Y, \mathbb{L} \cup \mathbb{U}, \sigma', \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^{\circ}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$$

which are well defined up to filtered equivariant chain homotopy. Formally these are constructed very similarly to the free stabilization maps defined in Chapter 2. The Floer homological component of the proof of invariance of these maps is identical to the proof of invariance of the free-stabilization maps in Section 2.5, so we will focus on the topology in this section.

We pick a regular neighborhood $\overline{D}$ of D which doesn't intersect L and define maps $S^{\pm}_{\mathbb{U},D,\overline{D}}$ depending on the choice of regular neighborhood $\overline{D}$. We consider only diagrams $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ such that

1. $\overline{D} \cap \Sigma$ is a disk;
2. $\overline{D} \cap (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z}) = \varnothing$.

We call such diagrams **subordinate** to $\overline{D}$. As in Chapter 2 and 3 we also need to choose a collection of data $\mathcal{J}$, which we call gluing data. The collection $\mathcal{J}$ consists of a choice of almost complex structures on $\mathcal{H}$, a choice of embedding of the diagram shown in Figure 4.8 into Σ, as well as some additional data.

Given such a subordinate diagram $\mathcal{H}$ and a choice of gluing data $\mathcal{J}$, we define a diagram $\overline{\mathcal{H}}_{\overline{D},\mathcal{J}}$ by simply replacing $\Sigma \cap \overline{D}$ with a disk shown in Figure 4.8, according to an embedding

specified by $\mathcal{J}$. For a given choice of neck length T, a choice of gluing data $\mathcal{J}$ determines an almost complex structure $\mathcal{J}(T)$ on the diagram $\overline{\mathcal{H}}_{\overline{D},\mathcal{J}}$.

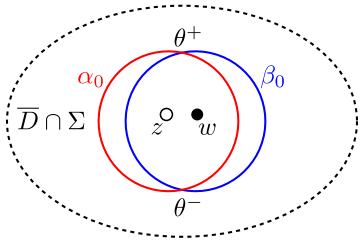

Figure 4.8: **The diagram we insert into $\overline{D} \cap \Sigma$ for the maps $S^{\pm}_{\mathbb{U},D,\overline{D}}$.**

The two intersection points of α_0 and β_0 are distinguished by the relative grading, and the designation of top and bottom generators is the same for $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. With respect to these diagrams and almost complex structures, we have the following differential computation:

Lemma 4.3.30. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard diagram for $\mathbb{L}$ and that $\overline{D} \cap \Sigma$ is contained in the component A of $\Sigma \setminus \boldsymbol{\alpha}$ containing, and the component B of $\Sigma \setminus \boldsymbol{\beta}$. Let w_A and z_A (resp. w_B and z_B) denote the basepoints in A (resp. B). Then for sufficiently large T, there is an identification of*

$$\partial_{\mathcal{H},\mathcal{J}(T)} = \begin{pmatrix} \partial_{\mathcal{H}} & * \\ 0 & \partial_{\mathcal{H}} \end{pmatrix}.$$

For a choice of almost complex structure with connected sum point sufficiently close to the α_0 curve, there is an identification of $*$ as $U_w V_z + U_{w_A} V_{z_A}$. For connected sum point sufficient close to the β_0 curve, there is an identification of $*$ as $U_w V_z + U_{w_B} V_{z_B}$.

Proof. The proof is identical to the proof of [OS08, Proposition 6.5]. Alternatively, all components except the top right are computed in Lemma 2.5.5, though the proof in the latter is essentially identical to the former argument. □

We define a map

$$S^+_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T} : \mathcal{C}FL^\circ_{J_s}(\mathcal{H}, \sigma_1, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ_{\mathcal{J}(T)}(\overline{\mathcal{H}}_{\overline{D}}, \sigma_2, \mathfrak{P}, \mathfrak{s}),$$

and a map $S^-_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$, defined in the opposite direction, by the familiar formulas of

$$S^+_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

and

$$S^-_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}(\mathbf{x} \times \theta^-) = \mathbf{x} \qquad \text{and} \qquad S^-_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}(\mathbf{x} \times \theta^+) = 0.$$

For large enough T, the maps $S^\pm_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$ are chain maps by Lemma 4.3.30, and also for all sufficiently large T, they commute with increasing the parameter T. As with the free-stabilization and quasi-stabilization maps, we can show that $S^\pm_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$ are independent of T and $\mathcal{J}$ by a computation of the change of almost complex structure map. Independence from $\mathcal{H}$ follows identically to the free-stabilization maps. Independence from the choice of regular neighborhood $\overline{D}$ follows since any two choices of $\overline{D}$ are isotopic via an isotopy fixing D and L, and the relevant diagrams tautologically commute because the change of diagram maps are induced by isotopies. We have the following:

Proposition 4.3.31. *For fixed D, the maps $S^\pm_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$ are independent of the choices of $\overline{D}, \mathcal{H}, \mathcal{J}$ and T.*

Proof. The proof follows identically to Theorem 2.5.22. A key component of the proof is a triangle count in Theorem 2.5.17 for freely stabilized triangles. □

We denote the common map by $S^\pm_{\mathbb{U},D}$. We make no claim about independence from the disk D, and indeed we expect the map to depend on the choice of embedded disk D.

4.4 Maps for handles attached away from L

In this section we define maps associated to handle attachments away from the link L. We define maps associated to surgery on a framed 0-sphere, a framed 1-dimensional link, or a framed 2-sphere. The 1-handle and 3-handle maps are formally identical to the maps defined in Chapter 2 (which are, of course, modifications of the original maps defined in [OS06]). The 2-handle maps are defined nearly identically to the maps from [OS06], though there is a

minor difference as in our setting one must pick either the **w**-basepoints or the **z**-basepoints to compute Spin^c structures. We use the **w**-basepoints for this.

Adapting the terminology of [Juh15], we make the following definition:

Definition 4.4.1. We say $\mathbb{S}$ is a **framed k-sphere** in Y, away from L, if $\mathbb{S}$ is an embedding of $S^k \times D^{3-k}$ into Y.

We usually write $\iota : S^k \times D^{3-k} \to Y$ for the actual map.

4.4.1 Model link cobordisms for handles away from L

We now describe model cobordisms, corresponding to handle attachments away from L. Suppose that L is a multi-based link in Y, and $\mathbb{S}$ is collection of pairwise disjoint, framed spheres in Y, all away from L. If $\mathbb{S}$ has only a single component, we define

$$Y(\mathbb{S}) = (Y \setminus \iota(S^k \times \operatorname{int} D^{3-k})) \cup_\iota (D^{k+1} \times S^{3-k-1}).$$

If $\mathbb{S}$ has multiple components, we iterate this procedure for each component. We allow the possibility that $\mathbb{S} = \varnothing$.

We define the model cobordism $W(Y, L, \mathbb{S})$ for attaching handles along $\mathbb{S}$ by $W(Y, L, \mathbb{S}) = (W, \Sigma)$ where

$$W = Y \times [0,1] \cup_{\mathbb{S}} H, \qquad \text{and} \qquad \Sigma = L \times [0,1].$$

Here H is a disjoint union consisting of one copy of $D^{k+1} \times D^{3-k}$ for each framed k-sphere of $\mathbb{S}$.

4.4.2 Maps for framed 0- or 2-spheres

For $(Y, \mathbb{L})$ a closed 3-manifold with embedded multibased link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, we now define maps associated to framed 0-spheres $\mathbb{S}_0 \subseteq Y \setminus L$ and 2-spheres $\mathbb{S}_2 \subseteq Y \setminus L$. Given a 0-sphere $\mathbb{S}_0 = \{p_1, p_2\} \times D^3$, and a Spin^c structure $\mathfrak{s}$ on Y, there is a unique Spin^c structure $\bar{\mathfrak{s}}$ on the cobordism $W(Y, \mathbb{S}_0)$ which extends $\mathfrak{s}$. We let $\mathfrak{s}'$ denote the restriction to $Y(\mathbb{S}_0)$.

The 1-handle maps will take the form

$$F_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}}} : \mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s}) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_0),\mathbb{L},\mathfrak{s}').$$

Given a framed 2-sphere $\mathbb{S}_2 \subseteq Y$, there is a canonical "dual" framed zero sphere $\mathbb{S}_0 \subseteq Y(\mathbb{S}_2)$. If $\iota : S^2 \times D^1 \to Y$ is the map associated to $\mathbb{S}_2$, then by definition

$$Y(\mathbb{S}_2) = (Y \setminus \iota(S^2 \times D^1)) \cup_\iota (D^3 \times S^0).$$

The subset $D^3 \times S^0$ is the canonical dual framed zero sphere, $\mathbb{S}_0$.

Since $\mathbb{S}_0$ is a framed sphere, i.e. an embedding of $S^0 \times D^3$ into Y, the manifold $Y(\mathbb{S}_2)(\mathbb{S}_0)$ is canonically diffeomorphic to Y. Thus, to define maps for framed 2-spheres, it is sufficient to instead define maps

$$G_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}}} : \mathcal{C}FL^\circ(Y(\mathbb{S}_0),\mathbb{L},\mathfrak{s}') \to \mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s}),$$

for $\bar{\mathfrak{s}}$ the unique Spin^c structure on $W(Y(\mathbb{S}_0),\mathbb{S}_2)$ extending $\mathfrak{s}'$ on $Y(\mathbb{S}_0)$.

We will only consider diagrams $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for $(Y,\mathbb{L})$ which are **subordinate** to the framed 0-sphere $\mathbb{S}_0$, i.e. $\mathbb{S}_0 \cap \Sigma$ consists of two disks which don't intersect $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$. For a choice of data $\mathcal{J}$, which as usual we call gluing data, we define a diagram $\mathcal{H}_\mathcal{J} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \mathbf{w}, \mathbf{z})$, which is formed by removing small disks around the points p_1 and p_2 and connecting the two new boundary components with an annulus. On the annulus we add two homotopically nontrivial curves, α_0 and β_0, which intersect each other transversely at two points. The two intersection points are distinguished by the relative grading (note that both $\mathrm{gr}_{\mathbf{w}}$ or $\mathrm{gr}_{\mathbf{z}}$ give the same decomposition). The 1- and 3- handle maps take the familiar form:

$$F_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}},\mathcal{H},\mathcal{J}}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

and

$$G_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}},\mathcal{H},\mathcal{J}}(\mathbf{x} \times \theta^+) = 0, \qquad \text{and} \qquad G_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}},\mathcal{H},\mathcal{J}}(\mathbf{x} \times \theta^-) = \mathbf{x}.$$

This gives a 1-handle map from $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$ to $\mathcal{C}FL^\circ(\overline{\mathcal{H}},\mathfrak{s}')$ and a 3-handle map from $\mathcal{C}FL^\circ(\overline{\mathcal{H}},\mathfrak{s}')$ to $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$. For a choice of $\mathcal{H}$ and $\mathcal{J}$, by pre- and post-composing with

change of diagrams maps, we get a well defined map

$$F_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}},\mathcal{H},\mathcal{J}} : \mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s}) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_0),\mathbb{L},\mathfrak{s}')$$

and a map $G_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}}}$ in the opposite direction, both of which are maps of coherent filtered chain homotopy types. The above maps *a-priori* depend on the choice of $\mathcal{H}$ and $\mathcal{J}$, though we have the following:

Lemma 4.4.2. *The maps $F_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}},\mathcal{H},\mathcal{J}}$ and $G_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}},\mathcal{H},\mathcal{J}}$ are independent of $\mathcal{H}$ and $\mathcal{J}$.*

The proof in Section 2.7 carries over without change.

Lemma 4.4.3. *The maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ commute with the 1- and 3-handle maps $F_{Y,\mathbb{S}_0,\bar{\mathfrak{s}}}$ and $F_{Y,\mathbb{S}_2,\bar{\mathfrak{s}}}$.*

Proof. The formulas for the maps obviously commute, but it is not clear whether a single almost complex structure can be used for both maps. However both maps are defined for particular choices of almost complex structures, and it is not clear whether the same structure can be used for both. Instead we will show that the change of almost complex structure map to go between almost complex structures which can be used to compute the respective maps preserves the formulas for the maps.

Let $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z})$ be a diagram for $\mathbb{L}$, the link in the 3-manifold Y before surgery on $\mathbb{S}_0$ or quasi-stabilization. Suppose that $\mathbb{S}_0 = \{p_1,p_2\}\times D^3$ is a framed 0-sphere, $p \in \Sigma \setminus (\boldsymbol{\alpha}\cup\boldsymbol{\beta}\cup\mathbf{w}\cup\mathbf{z})$ is a point at which we will perform quasi-stabilization. Let α_0 and β_0 denote the new curves from the 1-handle operation. Let α_0' and β_0' denote the two new curves from quasi-stabilization. These are shown in Figure 4.9.

To show that the 1- and 3-handle maps, as well as the quasi-stabilization maps can all be computed simultaneously we will show that any nonconstant disks counted by the change of almost complex structures map

$$\Phi_{\mathcal{J}(T_1,T_2,T_3)\to\mathcal{J}(T_1',T_2',T_3')},$$

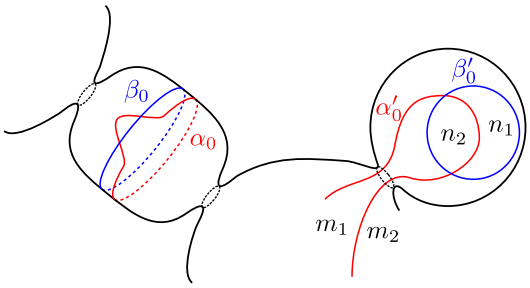

Figure 4.9:

The three necks we stretch in Lemma 4.4.3. Multiplicities m_1, m_2, n_1, and n_2 are also labeled.

occur for disks

$$\phi \in \pi_2(\mathbf{x} \times x \times x', \mathbf{y} \times y \times y')$$

with $\Delta(x,y) < 0$, where $\Delta(x,y) = \Delta_{\mathbf{w}}(x,y) = \Delta_{\mathbf{z}}(x,y)$ denotes the drop in grading from x to y. Here $x, y \in \alpha_0 \cap \beta_0$ correspond to the intersection points resulting from the 1-handle addition, and $x', y' \in \alpha_0' \cap \beta_0'$ are the intersection points from quasi-stabilization. It is easy to check that this is algebraically sufficient to imply that the quasi-stabilization maps commute with the 1- and 3-handle maps.

Suppose that u_i is a sequence of Maslov index zero holomorphic disks with dynamic almost complex structure representing a homology class ϕ. As we let $i \to \infty$, we can extract a weak limit to a collection of curves U_Σ on $(\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ and U_S on $(S^2, \alpha_0', \beta_0')$. As in Proposition 3.5.3, the curves in U_Σ consist exactly of a broken holomorphic curve U_Σ' on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ and boundary degenerations $\mathcal{A}$ on $(\Sigma, \boldsymbol{\alpha}^+)$.

Using the Maslov index computations from Lemma 3.5.3 and Lemma 2.5.3, we see that

$$\mu(\phi) = \mu(U_\Sigma) + \Delta(x,y) + n_1(\phi) + n_2(\phi) + m_1(\mathcal{A}) + m_2(\mathcal{A}) + 2\sum_{\substack{w_0 \in \mathbf{w} \\ w_0 \neq w'}} n_{w_0}(\phi),$$

where n_1, n_2, m_1, and m_2 are the multiplicities shown in Figure 4.9, and w' is the $\mathbf{w}$-basepoint adjacent to z (the new $\mathbf{z}$-basepoint from quasi-stabilization). If $\Delta(x,y) \geq 0$, then all of the above terms must be zero for $\mu(\phi) = 0$. In particular we conclude that $\mu(U_\Sigma) = 0$, so U_Σ must be a constant disk and $\mathcal{A} = \varnothing$. Note that if $\Delta(x,y) = 1$, no disks are counted.

As we mentioned before, with this, it's relatively easy to check all of the possible combinations of $S^{\pm}_{w,z}$ and $F_{Y,\mathbb{S}_0,\bar{\mathfrak{s}}}$ and $G_{Y(\mathbb{S}_0),\mathbb{S}_0,\bar{\mathfrak{s}}}$ to see that they commute. □

Lemma 4.4.4. *Suppose that* $\mathbb{S}$ *and* $\mathbb{S}'$ *are two framed spheres, which are either 0-dimensional or 2-dimensional. If* $\mathfrak{s} \in \mathrm{Spin}^c(W(Y,\mathbb{S},\mathbb{S}'))$*, then*

$$F_{Y(\mathbb{S}),\mathbb{L},\mathbb{S}',\mathfrak{s}|_{W(Y(\mathbb{S}),\mathbb{S}')}} F_{Y,\mathbb{L},\mathbb{S},\mathfrak{s}|_{W(Y,\mathbb{S})}} \simeq F_{Y(\mathbb{S}'),\mathbb{L},\mathbb{S},\mathfrak{s}|_{W(Y(\mathbb{S}'),\mathbb{S})}} F_{Y,\mathbb{L},\mathbb{S}',\mathfrak{s}|_{W(Y,\mathbb{S}')}}.$$

Proof. The formulas obviously commute, though one needs to pay attention to the almost complex structures. The appropriate almost complex structure map can be computed using a modification of the argument in Lemma 4.4.3. In general, the change of almost complex structure map may have nontrivial entries, but using the analogous computation as in Lemma 4.4.3, any component which counts nonconstant and nonzero disks must have at least one $\Delta(x_i, y_i)$ negative. It is easy to verify algebraically that such disks do not affect the commutator of the two 1- or 3-handle attachment maps. □

4.4.3 Maps for framed 1-dimensional links

Given a multibased link $\mathbb{L}$ in Y with $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and a framed 1-dimensional link $\mathbb{S}_1 \subseteq Y \setminus L$, we now describe maps $F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ associated to a framed link $\mathbb{S}_1$ and the Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(W(Y,\mathbb{S}_1))$.

Let $\mathbb{L}_\alpha$ denote the arcs of L which go from $\mathbf{w}$-basepoints to $\mathbf{z}$-basepoints. Let $\mathbb{L}_\beta$ denote the arcs of L which go from $\mathbf{z}$-basepoints to $\mathbf{w}$-basepoints.

Definition 4.4.5. An **α-bouquet**, $\mathcal{B}^\alpha$, for a framed link $\mathbb{S}_1$ in $Y \setminus L$ is a collection of arcs which connect links in $\mathbb{S}_1$ to the arcs in $\mathbb{L}_\alpha$. We assume that there is exactly one arc per component of $\mathbb{S}_1$, and each arc has one endpoint on $\mathbb{S}_1$ and one endpoint in $\mathbb{L}_\alpha$. A **β-bouquet**, $\mathcal{B}^\beta$, for $\mathbb{S}_1$ is defined analogously, but with endpoints on $\mathbb{L}_\beta$ instead.

Given a bouquet $\mathcal{B}^\alpha$ or $\mathcal{B}^\beta$ for a framed link $\mathbb{S}_1 \subseteq Y \setminus L$, we can consider the sutured manifolds $Y(\mathcal{B}^\alpha)$ and $Y(\mathcal{B}^\beta)$ obtained by removing regular neighborhoods of $L \cup \mathcal{B}^\alpha$ and

$L \cup \mathcal{B}^\beta$, and adding sutures corresponding to the basepoints $\mathbf{w}$ and $\mathbf{z}$. We have the following definition:

Definition 4.4.6. We say that that a triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to the β-bouquet $\mathcal{B}^\beta$ for a framed 1-dimensional link $\mathbb{S}_1 \subseteq Y \setminus L$ if

1. $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_{k+1}, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ yields a sutured diagram for $Y(\mathcal{B}^\beta)$;
2. The curves $\beta_1, \dots, \beta_k$ are each meridians of a different component of $\mathbb{S}_1$ and hence $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;
3. The curves $\gamma_{k+1}, \dots, \gamma_n$ are each small Hamiltonian isotopies of the curves $\beta_{k+1}, \dots, \beta_n$ respectively, with $|\beta_i \cap \gamma_j| = 2\delta_{ij}$;
4. for $i = 1, \dots, k$, the curve γ_i is induced by the framing of the corresponding link component that β_i is a meridian of.

If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is subordinate to a β-bouquet, then $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ represents an unlink in $(S^1 \times S^2)^{\#(g-k)}$ with exactly two basepoints per component. As such, regardless of the coloring, $\mathcal{H}FL^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$ admits a generator $\Theta_{\beta\gamma}$ of the top degree subset by Lemma 4.3.1.

The 4-manifold $X_{\alpha\beta\gamma}$ becomes diffeomorphic to $W(Y, \mathbb{S}_1)$ after filling in $Y_{\beta\gamma}$ with 3- and 4-handles (see [OS06, Proposition 4.3]). As such, given a $\mathfrak{s} \in \mathrm{Spin}^c(W(Y, \mathbb{S}_1))$, we naturally can restrict an $\mathfrak{s} \in \mathrm{Spin}^c(W(Y, \mathbb{S}_1))$ to $\mathrm{Spin}^c(X_{\alpha\beta\gamma})$. Hence given an $\mathfrak{s} \in \mathrm{Spin}^c(W(Y, \mathbb{S}_1))$, we can define the β-subordinate 2-handle map

$$F^\beta_{Y, \mathbb{L}, \mathbb{S}_1, \mathfrak{s}, \mathcal{B}^\beta} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}|_Y) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_1), \mathbb{L}, \mathfrak{s}|_{Y(\mathbb{S}_1)})$$

by the formula

$$\mathbf{x} \mapsto F_{\alpha\beta\gamma, \mathfrak{s}}(\mathbf{x} \otimes \Theta_{\beta\gamma}) = \sum_{\substack{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta\gamma}, \mathbf{y}) \\ \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}}} \#\mathcal{M}(\psi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

The proofs in [OS06], [Juh16], or Chapter 2 show that these don't depend on the choice of bouquet or Heegaard triple subordinate to them.

One similarly defines what it means for a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ to be subordinate to an α-bouquet. The complex $\mathcal{C}FL^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$ will have a generator $\Theta_{\alpha\beta}$ of the top degree, and we define the α-subordinate 2-handle map simply by

$$F^{\alpha}_{Y,\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\alpha}(\mathbf{x}) = F_{\alpha\beta\gamma,\mathfrak{s}}(\Theta_{\alpha\beta} \otimes \mathbf{x}).$$

As before the original argument from [OS06] goes through essentially without change, and we see that $F^{\alpha}_{Y,\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\alpha}$ is independent of the choice of α-bouquet or triple subordinate to it.

Remark 4.4.7. In [OS06], Ozsváth and Szabó construct a map $\mathfrak{s}_{\mathbf{w}} : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathrm{Spin}^c(X_{\alpha\beta\gamma})$. Analogously to the maps $\mathfrak{s}_{\mathbf{w}} : \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \mathrm{Spin}^c(Y_{\alpha\beta})$, the map on triangles depends on the choice of basepoints (cf. Lemma 4.2.8). Indeed the maps $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ are not equal, though this asymmetry should be viewed as an essential feature of the theory, as it leads to an Alexander grading change formula, to appear in Chapter 5

We now wish to show that $F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ and $F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ are equal. The proof given by Ozsváth and Szabó in [OS06, Lem. 5.2.] for their invariants essentially carries over without change. We state it as a lemma here and outline their argument briefly.

Lemma 4.4.8. *The maps $F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ and $F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ induce filtered chain homotopy equivalent maps on the coherent filtered chain homotopy invariants.*

Proof sketch. Let $\mathbb{S}_1'$ be an isotopic copy of $\mathbb{S}_1$, in the direction of the framing of $\mathbb{S}_1$. One finds a Heegaard quadruple $(\Sigma, \boldsymbol{\delta}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ such that $(\Sigma, \boldsymbol{\delta}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is triple subordinate to an α-bouquet of $\mathbb{S}_1'$, and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a triple subordinate to a β-bouquet of $\mathbb{S}_1$. After surgering on $\mathbb{S}_1$, the collection $\mathbb{S}_1'$ becomes a collection of 0-framed unknots. Similarly after surgering on $\mathbb{S}_1'$, the link $\mathbb{S}_1$ becomes a collection of 0-framed unknots. The standard associativity relations shows that,

$$F^{\alpha}_{Y(\mathbb{S}_1),\mathbb{L},\mathbb{S}_1',\mathfrak{s}_0} \circ F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F^{\beta}_{Y(\mathbb{S}_1'),\mathbb{L},\mathbb{S}_1,\mathfrak{s}_0} \circ F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1',\mathfrak{s}}$$

for appropriate choices of Spin^c structures on the two triangle maps. The second map in each of the compositions can be canceled by a composition of 3-handle maps, implying that $F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1}$ and $F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1'}$ are related by post-composition by a diffeomorphism. □

Since the triangle maps are graded over 4-dimensional Spin^c structures, and there isn't a canonical way to compose Spin^c structures, they instead satisfy a Spin^c composition law. We have the following:

Proposition 4.4.9. *If $\mathbb{S}^1_1$ and $\mathbb{S}^2_1$ are two disjoint framed 1-dimensional links in Y, then*

$$F_{Y(\mathbb{S}^1_1),\mathbb{L},\mathbb{S}^2_1,\mathfrak{s}_2} \circ F_{Y,\mathbb{L},\mathbb{S}^1_1,\mathfrak{s}_1} \simeq \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W(Y,\mathbb{S}^1_1\cup\mathbb{S}^2_1)) \\ \mathfrak{s}|_{W(Y,\mathbb{S}^1_1)}=\mathfrak{s}_1,\mathfrak{s}|_{W(Y(\mathbb{S}^1_1),\mathbb{S}^2_1)}=\mathfrak{s}_2}} F_{Y,\mathbb{L},\mathbb{S}^1_1\cup\mathbb{S}^2_1,\mathfrak{s}}.$$

Proof. This follows from associativity of the the triangle maps [OS04e, Theorem 8.16]. □

4.5 Quasi-stabilization and triangle maps

In Chapter 3, the quasi-stabilization operation was described in order to define the maps $S^{\pm}_{w,z}$ and $S^{\pm}_{z,w}$. There were three main holomorphic curve counts which were necessary to show naturality of the maps $S^{\pm}_{w,z}$, namely a count of disks necessary to show that the maps $S^{\pm}_{w,z}$ were chain maps, and two counts of holomorphic triangles necessary to show that the maps commuted with the change of diagrams maps induced by moves of the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves. In this section, we review the counts of holomorphic triangles in quasi-stabilized diagrams, and we also modify the holomorphic triangle counts from Theorem 4.5.3, to get a count of holomorphic triangles which will allow us to compute the relations between the band maps and the quasi-stabilization maps. Later, in Remark 4.7.14, we sketch a strategy which one could use to construct the quasi-stabilization maps without these computations if they were willing to define the quasi-stabilization maps instead as a composition of a disk stabilization map, a band map, and a diffeomorphism map.

All of statements about higher and lower graded generators in this section will be with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. Symmetrically, all statements also hold if one replaces $\mathrm{gr}_{\mathbf{w}}$ with $\mathrm{gr}_{\mathbf{z}}$.

4.5.1 Strongly positive diagrams for $(S^1 \times S^2)^{\#k}$ and quasi-stabilization of Heegaard triples

In this subsection we review the results from Chapter 3 about quasi-stabilization and triangle maps. We first review a result from [MO10], which will be useful in showing that quasi-stabilization commutes with the change of diagram maps induced by $\boldsymbol{\beta}$–moves.

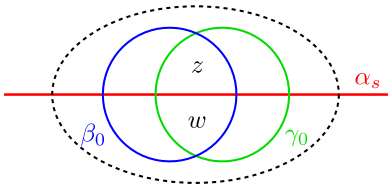

Figure 4.10: **The version of quasi-stabilization discussed in Lemma 4.5.1.** It was considered in [MO10, Prop. 5.2]

Given a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$, we quasi-stabilize it along a single α_s curve passing through a point $p \in \Sigma$, as in Figure 4.10. We have the following:

Lemma 4.5.1 ([MO10, Prop. 5.2])**.** *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a strongly $\mathfrak{s}$–admissible triple and suppose that α_s is a new $\boldsymbol{\alpha}$–curve, passing through the point $p \in \Sigma$. Let $\overline{\mathcal{T}}_{\alpha_s,p}$ denote the Heegaard triple resulting from quasi-stabilizing along α_s at p, as in Figure 4.10. If $\mathcal{J}$ is gluing data, then for sufficiently large T, with the almost complex structure $\mathcal{J}(T)$ we have the following identifications:*

$$F_{\overline{\mathcal{T}}_{\alpha_s,p},\bar{\mathfrak{s}}}(\mathbf{x} \times \cdot, \mathbf{y} \times y^+) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x},\mathbf{y}) & 0 \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x},\mathbf{y}) \end{pmatrix}$$

where $\cdot \in \{x^+, x^-\} = \alpha_s \cap \beta_0$ and the matrix on the right denotes the expansion into the upper and lower generator components of $\alpha_s \cap \beta_0$ and $\alpha_s \cap \gamma_0$.

We also need an count which shows that the map $S^{\pm}_{w,z}$ is independent of the choice of curve α_s. This requires a more subtle result about how curves can degenerate. Indeed taking an arbitrary Heegaard triple and performing quasi-stabilization along two curves α_s and β_s which intersect at two points in general will not result in curve counts as in the above

theorem. Nevertheless, if we restrict which diagrams can appear as $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, we can achieve such a curve count. We make the following definition:

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(S^1 \times S^2)^{\#k}$ such that

$$|\alpha_i \cap \beta_j| = \begin{cases} 1 \text{ or } 2 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

and that if $\alpha_i \cap \beta_i = \{p_i^-, p_i^+\}$, then p_i^- and p_i^+ differ by Maslov grading one. We do not assume that the α_i curves are small isotopies of the β_j curves. Let $\theta^+ = p_1^+ \times \cdots \times p_{n-1}^+$ denote the top graded (partial) intersection point. Assume that $|\alpha_n \cap \beta_n| = 2$ and write p_n^- and p_n^+ for the two points of $\alpha_n \cap \beta_n$.

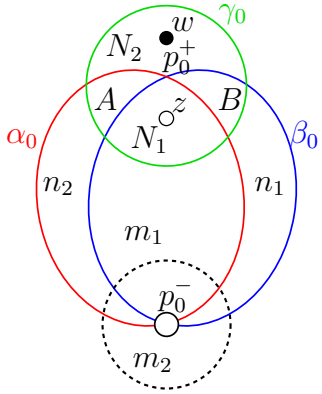

Figure 4.11: **The diagram we insert into a Heegaard triple diagram $\mathcal{T}$ along the curves α_s and β_s to form the diagram $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$.** We cut out the solid circle marked with p_0^- and stretch the almost complex structure along the dashed circle.

Definition 4.5.2. Under the assumptions of the previous paragraph, we say that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is **strongly positive** with respect to p_n^+ if for every nonnegative disk $\phi \in \pi_2(\theta^+ \times p_n^+, \mathbf{y} \times p_n^+)$ we have that

$$(\mu - (m_1 + m_2))(\phi) = (\mu - (n_1 + n_2))(\phi) \geq 0,$$

with equality to zero iff ϕ is the constant disk. Here m_1, m_2, n_1, n_2 denote the multiplicities adjacent to the point p_n^+, appearing in the following counterclockwise order: n_1, m_1, n_2, then m_2.

Indeed Heegaard triples where $(\Sigma, \boldsymbol{\alpha}\cup\{\alpha_s\}, \boldsymbol{\beta}\cup\{\beta_s\}, \mathbf{w}\cup\{w\}, \mathbf{z}\cup\{z\})$ is strongly positive with respect to the quasi-stabilization point are ones where we have a reasonable count of curves, as the following theorem demonstrates:

Theorem 4.5.3 (Theorem 3.6.5)**.** *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with curves α_s and β_s, intersecting at two points p^+ and p^- and let $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$ be the Heegaard triple result from quasi-stabilization, as described above. If $(\Sigma, \boldsymbol{\alpha}\cup\{\alpha_s\}, \boldsymbol{\beta}\cup\{\beta_s\}, \mathbf{w}\cup\{w\})$ is a strongly positive diagram for $(S^1\times S^2)^{\#k}$ with respect to the point p^+ (Definition 4.5.2), and $\mathcal{J}$ is gluing data for stretching along the dashed circle, then for sufficiently large T, with respect to the almost complex structure $\mathcal{J}(T)$, there are identifications*

$$F_{\mathcal{T}^+_{\alpha_s,\beta_s,p^+},\bar{\mathfrak{s}},\mathcal{J}(T)}(\Theta^+_{\alpha\beta}\times p_0^+, \mathbf{x}\times\cdot) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s},J_s}(\Theta^+_{\alpha\beta},\mathbf{x}) & 0 \\ 0 & F_{\mathcal{T},\mathfrak{s},J_s}(\Theta^+_{\alpha\beta},\mathbf{x}) \end{pmatrix},$$

where $\cdot$ denotes $x^\pm\in\beta_0\cap\gamma_0$ and the matrix on the right denotes the matrix decomposition of the map based on the decompositions given $x^\pm$ and $y^\pm$. The decompositions are with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

4.5.2 Weakly positive diagrams and further holomorphic triangle computations

Unfortunately Theorem 4.5.3 only covers triangle diagrams which achieve the strong positivity condition. Indeed it does not hold for general triangle diagrams, even for the ones which will occur in the band maps.

Instead, we need a weaker condition. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple such that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(S^1\times S^2)^{\#k}$ such that each $\alpha_i\cap\beta_i$ consists of two points differing by grading 1, or just one point. Further assume that $\alpha_i\cap\beta_j = \varnothing$ if $i\neq j$. Suppose $\alpha_n\cap\beta_n = \{p_n^-, p_n^+\}$.

Definition 4.5.4. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a triple as in the previous paragraph, and α_s and β_s are new $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves, intersecting each other twice, with top intersection point p^+, and let $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$ denote the quasi-stabilized triple, as in Figure 4.11. We say that $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$ is **weakly positive** at p^+ if for every disk $\phi\in\pi_2(\theta^+\times p^+, \mathbf{y}\times p^+)$ on $(\Sigma, \boldsymbol{\alpha}\cup$

$\{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ which is nonnegative, we have that

$$(\mu - (m_1 + m_2))(\phi) = (\mu - (n_1 + n_2))(\phi) \geq 0,$$

and if we have equality to zero, then $m_2(\phi) = n_w(\phi) = 0$, where m_2 is the multiplicity in Figure 4.11.

Clearly any diagram which is strongly positive at p_n^+ is also weakly positive at p_n^+.

Lemma 4.5.5. *Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(S^1 \times S^2)^{\#k}$ where the curves in $\boldsymbol{\alpha}$ are all small Hamiltonian isotopies of the curves in $\boldsymbol{\beta}$ with $\alpha_i \cap \beta_j = 2\delta_{ij}$. If $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+} = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w} \cup \{w\})$ is a free-stabilization[4] at any point in $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w})$, then $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$ is weakly positive at p^+.*

Proof. The proof breaks into cases, depending on the region of $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ where the free-stabilization takes place. If it takes place in a region already containing a basepoint, the diagram $\mathcal{H}$ is actually strongly positive at p^+ by Lemma 3.6.3, and hence weakly positive as a consequence.

The other alternative is that it occurs in one of the small bigons corresponding to a Maslov index one disk from p_k^+ to p_k^-. Suppose ϕ is a disk in $\pi_2(\theta^+ \times p^+, \mathbf{y} \times p^+)$. Consider the multiplicities $m_1, m_2, n_1, n_2, m_2', n_1'$, and n_2' shown in Figure 4.12. Let $\phi_0 \in \pi_2(\theta^+, \mathbf{y})$ denote the disk induced by ϕ, on the unstabilized diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$.

By assumption $p_k^+ \in \theta^+$. Either p_k^+ or p_k^- is in $\mathbf{y}$. Let $\delta = \delta(p_k^+, \mathbf{y}) \in \{0, 1\}$ be 1 if $p_k^+ \in \mathbf{y}$ and 0, otherwise.

By assumption on the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$, the regions with multiplicities n_1' and n_2' must each have a basepoint (though they could be the same region). Hence $n_{\mathbf{w}}(\phi_0) \geq \frac{1}{2}(n_1' + n_2')$. Since $\phi_0 \in \pi_2(\theta^+, \mathbf{y})$, and θ^+ is the top degree intersection point, using the grading formula, we have

$$\mathrm{gr}(\theta^+) - \mathrm{gr}(\mathbf{y}) = \mu(\phi_0) - 2n_{\mathbf{w}}(\phi_0) \geq \delta,$$

[4] *As in Chapter 2, a free-stabilization is where α_s, β_s and γ_0 are all Hamiltonian isotopic, each pair intersecting in two points, such α_s, β_s, and γ_0 are contained in a ball B on Σ which has trivial intersection with $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w}$. The basepoint w is added inside the region bounded by all three curves.*

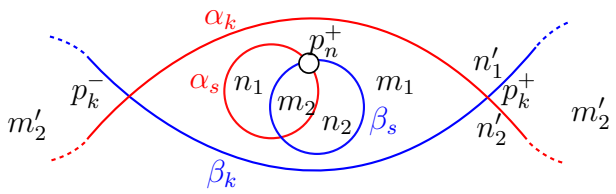

Figure 4.12: **A free-stabilization inside a bigon.** Multiplicities for a disk ϕ are labeled.

which yields

$$\mu(\phi_0) \geq 2n_{\mathbf{w}}(\phi_0) + \delta.$$

Using our bound for $n_{\mathbf{w}}(\phi_0)$, above, we have

$$\mu(\phi_0) \geq n_1' + n_2' + \delta.$$

Using the formula for the Maslov index of a disk after performing free-stabilization (cf. Lemma 2.5.3), we have that

$$\mu(\phi) \geq n_1' + n_2' + 2m_2 + \delta.$$

Since $n_1' + n_2' + \delta = m_1 + m_2'$ we have

$$\mu(\phi) - m_1 - m_2 \geq m_1 + m_2' + 2m_2 - (m_1 + m_2) = m_2' + m_2,$$

which is certainly a nonnegative quantity. On the other hand, if $\mu(\phi) - m_1 - m_2 = 0$, then $m_2' + m_2 = 0$ and hence in particular $m_2 = 0$. □

We now show that weakly positive double quasi-stabilizations have reasonable triangle map counts.

Proposition 4.5.6. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with curves α_s and β_s, intersecting at two points p^+ and p^- and let $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$ be the Heegaard triple result from quasi-stabilization, which satisfies the weak positivity condition, with respect to the top generator, p^+, of $\alpha_s \cap \beta_s$. Suppose $\mathcal{J}$ is gluing data for stretching the almost complex structure around γ_0. then for sufficiently large T, with respect to the almost complex structure $\mathcal{J}(T)$,*

there are identifications

$$F_{\mathcal{T}^+_{\alpha_s,\beta_s,p^+},\bar{\mathfrak{s}},\mathcal{J}(T)}(\Theta^+_{\alpha\beta}\times p_0^+,\mathbf{x}\times\cdot)=\begin{pmatrix}F_{\mathcal{T},\mathfrak{s},J_s}(\Theta^+_{\alpha\beta},\mathbf{x}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s},J_s}(\Theta^+_{\alpha\beta},\mathbf{x})\end{pmatrix},$$

where $\cdot$ denotes $x^\pm\in\beta_0\cap\gamma_0$ and the matrix on the right denotes the matrix decomposition of the map based on the decompositions given $x^\pm$ and $y^\pm$. The decomposition is taken with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

Proof. We perform the same degeneration as in Theorem 3.6.5. Indeed if ψ is a Maslov index zero triangle in $\pi_2(\theta^+\times p_n^+,\mathbf{x}\times x,\mathbf{y}\times y)$ with holomorphic representatives for arbitrarily large neck lengths, then we extract a weak limit as a collection of curves, as follows:

1. a broken holomorphic triangle u_Σ representing a homology class ψ_Σ on $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma})$ which has no boundary components on α_s or β_s;
2. a broken holomorphic disk $u_{\alpha\beta}$ on $(\Sigma,\boldsymbol{\alpha}\cup\{\alpha_s\},\boldsymbol{\beta}\cup\{\beta_s\})$ representing a class $\phi_{\alpha\beta}\in\pi_2(\Theta^+_{\alpha\beta}\times p^+,\mathbf{y}\times p^+)$, for some $\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\gamma$;
3. a broken holomorphic triangle u_0 on $(S^2,\alpha_0,\beta_0,\gamma_0)$ which represents a homology triangle $\psi_0\in\pi_2(p_0^+,x,y)$.

As in Theorem 3.6.5, the Maslov index can be computed to be

$$\mu(\psi)=\mu(\psi_\Sigma)+(N_1+N_2)(\psi_0)+\mu(\phi_{\alpha\beta})-\tfrac{1}{2}(n_1+n_2+m_1+m_2)(\phi_{\alpha\beta}).$$

By weak positivity, we know that the sum above can be expressed as a sum of nonnegative integers, and hence each must be zero. That is,

$$\mu(\psi_\Sigma)=N_1(\psi_0)+N_2(\psi_0)=\mu(\phi_{\alpha\beta})-\tfrac{1}{2}(n_1+n_2+m_1+m_2)(\phi_{\alpha\beta})=0.$$

By weak positivity, we also know that $m_2(\phi_{\alpha\beta})=0$.

Now since $\mu(\psi_\Sigma)=0$, we know that u_Σ, the broken triangle representing ψ_Σ, must actually be a genuine holomorphic triangle. As such, there is a genuine holomorphic triangle, u_0^m, in the broken curve u_0, which matches u_Σ in the sense that $\rho^{p^+}(u_\Sigma)\subseteq\rho^{p_0^-}(u_0^m)$.

First we note that ψ_Σ and $\phi_{\alpha\beta}$ are the only two classes appearing on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \boldsymbol{\gamma})$ and since $m_2(\phi_{\alpha\beta}) = 0$, we must have $m_2(\psi_0) = m_2(\psi) = m_2(\psi_\Sigma)$. Since u_Σ doesn't have any boundary components on α_s or β_s, we know that it must contribute to m_1, m_2, n_1 and n_2 equally. Hence we conclude that

$$n_2(\psi_0), m_1(\psi_0), n_1(\psi_0) \geq m_2(\psi_0).$$

Since $N_1(\psi) = N_2(\psi) = 0$, we must have $A(\psi) + B(\psi) = 1$, as p_0^+ is a vertex of ψ. Since p_0^+ is also a vertex of ψ_0, we similarly conclude that one of $A(\psi_0)$ and $B(\psi_0)$ is one, and the other is zero. Let us first consider the case that $A(\psi_0) = 1$. As a homology class, we can thus write the domain of ψ_0 as a sum of the small triangle with $A = 1$, $m_2(\psi)$ of the Maslov index 2 γ_0 degenerations with $A = B = N_1 = N_2 = 0$, as well as some sum of nonnegative domains supported in the regions marked with n_2, m_1 and n_1. Let n_2', m_1' and n_1' denote the multiplicities of $\phi_{\alpha\beta}$ in the regions corresponding to those marked n_2, m_1 and n_1 on S^2. We first have

$$n_2' + n_1' = m_1',$$

by the vertex relations. Secondly, it is easy to see that the vertex relations imply that n_1' and n_2' both take values in $\{0, 1\}$. Now since $B = 0$, we note that the vertex relations imply $n_1' = 0$. Hence the only possibility is that $n_2' = m_1' = 1$. In this case, we note that $\psi \in \pi_2(\Theta \times p_0^+, \mathbf{x} \times x^-, \mathbf{y} \times y^+)$, a component which contributes to the top right of the matrix.

Similarly, the remaining case is that $B = 1$ and $A = 0$. In this case, we have $n_2' = 0$ while it is possible that $m_1' = n_1' = 1$. In this case we also have that $\psi \in \pi_2(\Theta_{\alpha\beta}^+ \times p_0^+, \mathbf{x} \times x^-, \mathbf{y} \times y^+)$, so such triangles do not contribute to the counts we are interested in.

An example of how the broken triangle u_0 representing ψ_0 might split for triangles in $\pi_2(\Theta_{\alpha\beta}^+ \times p_0^+, \mathbf{y} \times y^-, \mathbf{z} \times z^+)$ (the ones we are not interested in) is shown in Figure 4.13.

For triangles not contributing to the upper right corner of the matrix, we must have $(m_1 + m_2 + n_1 + n_2)(\phi_{\alpha\beta}) = 0$. As $\mu(\phi_{\alpha\beta}) - \frac{1}{2}(m_1 + m_2 + n_1 + n_2)(\phi_{\alpha\beta}) = 0$, we conclude that $\mu(\phi_{\alpha\beta}) = 0$, and since $\phi_{\alpha\beta}$ admits a broken holomorphic representative, we conclude that $\phi_{\alpha\beta}$ is the constant disk. Hence ψ is represented by exactly a prematched triangle, and

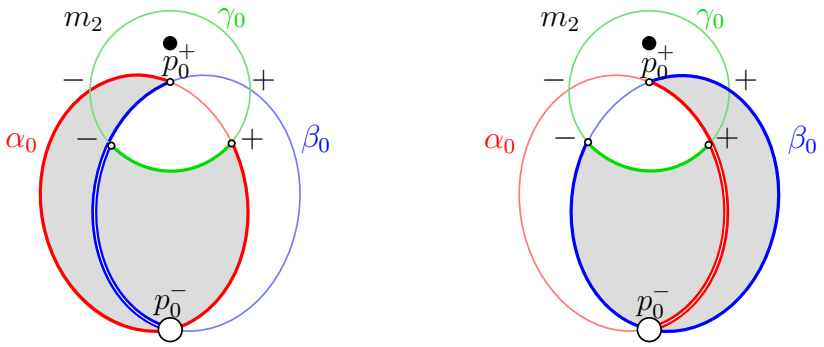

Figure 4.13: **Two broken triangles which could appear in 4.5.6 for homology triangles contributing to the top right component of the transition matrix**. The component was marked with a $*$.

ψ is a class in $\pi_2(\Theta^+_{\alpha\beta} \times p_0^+, \mathbf{x} \times x, \mathbf{y} \times y)$ for x and y of the same relative grading. From here the proof follows the proof of Theorem 3.6.5 identically. □

4.6 Definition of the band maps

In this section, we construct maps associated to surgering a link $\mathbb{L} \subseteq Y$ along a band B embedded in Y. A band B on a link L is an embedded copy of $[-1,1] \times [-1,1]$ in Y such that $B \cap L$ consists of exactly $\{-1,1\} \times [-1,1]$. In terms of link cobordisms, a band corresponds to a saddle critical point of the surface in the link cobordism. If L is an oriented link, an **oriented** band B is a band such that the orientation of B is opposite the orientation of L on $B \cap L$. If L is an oriented link and B is an oriented band for L, we can form the surgered link, $L(B)$, which is also an oriented link. We will assume all bands are oriented, and also omit stating that bands are oriented.

In terms of decorated cobordisms, we can assume that if $(W, (\Sigma, \mathcal{A}))$ is a decorated link cobordism, and f is a Morse function on W which restricts to a Morse function on Σ, then generically the index 1 critical points of $f|_\Sigma$ occur away from the divides $\mathcal{A}$. They can occur either inside a type–$\mathbf{w}$ region or inside of a type–$\mathbf{z}$ region. We will define two maps, which we will denote by $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ corresponding to bands consisting of a type–$\mathbf{z}$ region or a

type–$\mathbf{w}$ region, respectively.

Maps associated to band moves between links have been described (implicitly or explicitly) by many authors. For example, such maps have been defined explicitly by Sarkar in [Sar11a], Ozsváth, Stipsicz and Szabó in [OSS15]. Usually they are defined by a triangle map, or using grid diagrams. Alternatively, they can be defined in the hat flavor of link Floer homology using the construction in [Juh16], though this approach is somewhat different than the others in that it requires the gluing map in [HKM08] defined using contact geometry. In this section we construct maps associated to bands on $\mathcal{C}FL^\circ$, and prove that the maps are invariant from the choices in the construction.

Additionally, it's worth mentioning that in [JM16a], Juhász and Marengon prove a skein exact sequence for the hat version of link Floer homology, using link cobordism maps. In light of this, it is reasonable to suspect that many of the maps fitting into such exact sequences, which are often computed by counting holomorphic triangles, are secretly link cobordism maps. Examples of such constructions can be found in, for example, [OS04c] or [Man07].

4.6.1 Topological preliminaries on bands

We wish to consider bands on oriented, multi-based links, and we now define the type of bands which will be useful for our theory.

Definition 4.6.1. We say that a band B^α of an oriented multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is an α–**band** if the the ends of B^α occur in components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which go from $\mathbf{w}$–basepoints to $\mathbf{z}$–basepoints. We say that a band B^β is a β–**band** if the ends of B occur in components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ going from $\mathbf{z}$–basepoints to $\mathbf{w}$–basepoints. For a band B of either type, we always assume that the ends of B lie in distinct components of $L \setminus (\mathbf{w} \cup \mathbf{z})$.

Notice that a band can never be both an α–band and a β–band. If $\mathbb{L}$ is a multibased link and B is a band (of either type) for a link $\mathbb{L}$ in Y, we can form the surgered link $L(B)$ in the obvious way. Since the ends of B are assumed to lie distinct components of $L \setminus (\mathbf{w} \cup \mathbf{z})$,

the link with basepoints

$$\mathbb{L}(B) = (L(B), \mathbf{w}, \mathbf{z}),$$

is a multi-based link in the sense of Definition 4.2.1.

Definition 4.6.2. We say that that a Heegaard triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to an α–band B^α for $\mathbb{L}$, if

1. the diagram $(\Sigma, \beta_1, \dots, \beta_{n-1}, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ induces a sutured diagram for the sutured manifold $(Y \setminus N(L \cup B), \partial N(\mathbf{w} \cup \mathbf{z}))$, where $\partial N(\mathbf{w} \cup \mathbf{z}) \subseteq \Sigma$ are given orientations compatible to the orientation of L, as in [Juh16, Definition 4.4];

2. $(\Sigma, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;

3. $\alpha_1, \dots, \alpha_{n-1}$ are small Hamiltonian isotopies of the curves $\beta_1, \dots, \beta_{n-1}$ with $|\alpha_i \cap \beta_j| = \delta_{ij}$, for $i, j \in \{1, \dots, n-1\}$;

4. the curve α_n is induced by the band B^α and $(\Sigma, \alpha_1, \dots, \alpha_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha))$.

Using [JT12, Proposition 2.37], we know that any two sutured Heegaard diagrams for $Y \setminus N(L \cup B)$ can be connected by a set of Heegaard moves for embedded Heegaard surfaces, and hence we have the following:

Lemma 4.6.3. *Given an α–band B^α, any two triples subordinate to B^α can be connected by a sequence of the following moves*

1. *isotopies of Σ inside of Y which fix $L \cup B$.*

2. *isotopies or handleslides amongst the $\boldsymbol{\gamma}$–curves;*

3. *isotopies or handleslides amongst the $\alpha_1, \dots, \alpha_{n-1}$ curves followed by the corresponding handleslide or isotopy of the corresponding curve in $\beta_1, \dots, \beta_{n-1}$;*

4. *isotopies or handleslides of the curve β_n across other $\boldsymbol{\beta}$–curves;*

5. *isotopies or handleslides of the curve γ_n across other $\boldsymbol{\gamma}$–curves;*

6. $(1,2)$*–stabilizations of the Heegaard triple, i.e., taking the internal connected sum of Σ in Y with a torus $\mathbb{T}^2 \subseteq Y$ such that there is a 3–ball $B^3 \subseteq Y$ containing $\mathbb{T}^2$ with $B^3 \cap (\Sigma \cup L \cup B) = \varnothing$. On $\mathbb{T}^2$ we put three curves, α_0,β_0 and γ_0, with $|\beta_0 \cap \gamma_0| = 1$ and α_0 a small Hamiltonian isotopy of β_0 with $|\alpha_0 \cap \beta_0| = |\alpha_0 \cap \gamma_0| = 2$.*

There is an analogous lemma and definition for triples subordinate to β–bands B:

Definition 4.6.4. We say a Heegaard triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to a β–band B^β if

1. $(\Sigma \setminus N(\mathbf{w} \cup \mathbf{z}), \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_{n-1}, \partial N(\mathbf{w} \cup \mathbf{z}))$ induces a sutured diagram for the sutured manifold $(Y \setminus N(L \cup B^\beta), \partial N(\mathbf{w} \cup \mathbf{z}))$, where the circles $\partial N(\mathbf{w} \cup \mathbf{z})$ are given orientations compatible to the orientation of L, as in [Juh16, Definition 4.4];

2. $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;

3. $\gamma_1, \dots, \gamma_{n-1}$ are small Hamiltonian isotopies of the curves $\beta_1, \dots, \beta_{n-1}$, with $|\beta_i \cap \gamma_j| = \delta_{ij}$, for $i, j \in \{1, \dots, n-1\}$;

4. the curve γ_n is induced by the band B^β and $(\Sigma, \alpha_1, \dots, \alpha_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\beta))$.

As before we have the following:

Lemma 4.6.5. *Given a β–band B^β, any two triples subordinate to B^β can be connected by a sequence of the following moves*

1. *isotopies of Σ inside of Y which fix $L \cup B^\beta$.*

2. *isotopies or handleslides amongst the $\boldsymbol{\alpha}$–curves;*

3. *isotopies or handleslides amongst the $\beta_1, \dots, \beta_{n-1}$ curves followed by the corresponding handleslide or isotopy of the corresponding curve in $\gamma_1, \dots, \gamma_{n-1}$;*

4. *isotopies or handleslides of the curve β_n across other $\boldsymbol{\beta}$–curves;*

5. *isotopies or handleslides of the curve γ_n across other $\boldsymbol{\gamma}$–curves;*

6. *$(1,2)$–stabilizations of the Heegaard triple, by taking the internal connected sum of Σ in Y with a torus $\mathbb{T}^2 \subseteq Y$ such that there is a 3–ball $B^3 \subseteq Y$ containing $\mathbb{T}^2$ with $B^3 \cap (\Sigma \cup L \cup B^3) = \varnothing$. On $\mathbb{T}^2$ we put three curves, α_0,β_0 and γ_0, with $|\alpha_0 \cap \beta_0| = 1$ and γ_0 a small Hamiltonian isotopy of β_0 with $|\alpha_0 \cap \gamma_0| = |\beta_0 \cap \gamma_0| = 2$.*

We now define maps associated to bands. We first consider α–bands.

Lemma 4.6.6. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a diagram subordiate to an α–band B, then $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for an unlink in (a disjoint union of copies of) $(S^1 \times S^2)^{\#k}$, and every component of the unlink has exactly two basepoints, except one component, which has four.*

Proof. Let $(Y_{\alpha\beta}, \mathbb{L}_{\alpha\beta})$ denote the 3–manifold with embedded link obtained by attaching 2–handles to $\Sigma \times [0,1]$ and filling in the resulting boundary spheres in the unique way. The link $\mathbb{L}_{\alpha\beta}$ is obtained by concatenating a simple closed curve in each of the boundary spheres. Note that each pair (α_i, β_i) for $i = 1, \dots, n-1$ determines an embedded 2–sphere S_i in $Y_{\alpha\beta}$ which doesn't intersect $\mathbb{L}_{\alpha\beta}$. By surgering out S_i, we can reduce to the case that we are left with $|L|$ genus zero components of Σ. There are $|L|-1$ of these components which will have no $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves, but will have a single pair of $\mathbf{w}$ and $\mathbf{z}$ basepoints (thus determining an unknotted component of $\mathbb{L}_{\alpha\beta}$ in $Y_{\alpha\beta}$ with exactly two basepoints). There will be one component of the resulting Heegaard diagram which contains the curves α_n and β_n. In such a manner we reduce the claim to the case that the band B^α is a band between two unknots in S^3, each with two basepoints. A model diagram shows that for this link, α_n and β_n can be taken intersect twice on a genus zero Heegaard surface such that each of the resulting regions on Σ contains a single basepoint, completing the proof. □

4.6.2 Maps for z–bands

Given a $\mathbf{z}$–band B, for a multi-based link $\mathbb{L}$ with coloring $\sigma : (\mathbf{w} \cup \mathbf{z}) \to \mathfrak{P}$, we now define a map

$$F_B^{\mathbf{z}} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, \mathbb{L}(B), \sigma, \mathfrak{P}, \mathfrak{s}).$$

These are defined only when the $\mathbf{z}$–basepoints adjacent to the ends of the band have the same color.

As we showed in Lemma 4.3.2, there is a distinguished generator as a $\mathbb{Z}_2[V_{\mathbf{z}}]$–module

$$\Theta^{\mathbf{w}}_{\alpha\beta} \in \mathcal{H}FL^-_{\max,\mathbf{w}}(\mathcal{H}_{\alpha\beta}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0) \subseteq \mathcal{H}FL^-(\mathcal{H}_{\alpha\beta}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0)$$

where $\mathcal{H}FL^-_{\max,\mathbf{w}}$ denotes the subgroup of top homological degree with respect to the grading $\mathrm{gr}_{\mathbf{w}}$, and $(\sigma_0, \mathfrak{P}_0)$ is the coloring that sets V_z equal to $V_{z'}$, as in Lemma 4.3.2. Note that for any other coloring which gives z and z' the same color, there is a chain map

$$\mathcal{H}FL^-(\mathcal{H}_{\alpha\beta}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0) \to \mathcal{H}FL^-(\mathcal{H}_{\alpha\beta}, \sigma, \mathfrak{P}, \mathfrak{s}_0),$$

so the element $\Theta^{\mathbf{w}}_{\alpha\beta}$ is a well defined element of $\mathcal{H}FL^-$ for any coloring that assigns z and z' the same color.

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple subordinate to an α–band. The triangle maps require a choice of Spin^c structure on $X_{\alpha\beta\gamma}$, but we note that $X_{\alpha\beta\gamma}$ becomes diffeomorphic to $Y \times [0,1]$ after filling in $Y_{\alpha\beta}$ with 3– and 4–handles. Letting $\mathrm{Spin}^c(X_{\alpha\beta\gamma})_0 \subseteq \mathrm{Spin}^c(X_{\alpha\beta\gamma})$ denote the set of Spin^c structures which restrict to the torsion Spin^c structure on $Y_{\alpha\beta}$, we have

$$\mathrm{Spin}^c(X_{\alpha\beta\gamma})_0 \cong \mathrm{Spin}^c(Y).$$

Hence given an α–band B^α and a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, we can define a band map

$$F^{\mathbf{z}}_{B^\alpha}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\Theta^{\mathbf{w}}_{\alpha\beta} \otimes \cdot) : \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}).$$

Note that by $F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}$ we mean the map which counts holomorphic triangles representing classes ψ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$. For links where the sum of all of the components are nullhomologous,

this presents no asymmetry between the $\mathbf{w}$ and $\mathbf{z}$ basepoints as $\mathfrak{s}_{\mathbf{w}} = \mathfrak{s}_{\mathbf{z}}$ by Lemma 4.2.8 (or rather, it's analog for Heegaard triples, see Chapter 5). For links which aren't nullhomologous in this sense, this introduces an asymmetry between the $\mathbf{w}$ and $\mathbf{z}$ basepoints.

We now claim that these induce well defined maps on the coherent chain homotopy type:

Lemma 4.6.7. *The map $F^{\mathbf{z}}_{B^\alpha}$ is independent (up to equivariant, filtered chain homotopy) of the choice of Heegaard triple subordinate to a framed arc for the band B^α, and hence descends to a well defined map on the coherent filtered chain homotopy type invariants.*

Proof. We just need to check independence from each of the moves appearing Lemma 4.6.3. Each of these moves is considered in the maps associated to surgery on framed 1–dimensional links embedded in 3–manifolds, and can be addressed in the same way (cf. Lemma 2.8.12 and [Juh16, Theorem 6.9]). □

Analogously, if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a triangle subordinate to a β–band B^β, there is an element

$$\Theta^{\mathbf{w}}_{\beta\gamma} \in \mathcal{H}FL^-_{\max,\mathbf{w}}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$$

which generates $\mathcal{H}FL^-_{\max,\mathbf{w}}$ as a $\mathbb{Z}_2[V_{\mathbf{z}}]$–module. Thus as long as $\mathbf{z}$–basepoints adjacent to the ends of the band B^β are given the same color, we can define a band map

$$F^{\mathbf{z}}_{B^\beta}(\cdot) = F_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma},\mathfrak{s}}(\cdot \otimes \Theta^{\mathbf{w}}_{\beta\gamma}) : \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}),$$

which descends to the coherent $\mathfrak{P}$–filtered chain homotopy type invariant.

Remark 4.6.8. In [OS06], Ozsváth and Szabó define cobordism maps for performing surgery on framed links by taking Heegaard triples which are subordinate to the bouquets of the links. In that setting, one can take Heegaard triples which are either left or right subordinate, corresponding to putting the top degree generator on the right or the left side in the Heegaard triple. In the case of band maps, there is no ambiguity of whether a band is an α–band or a β–band, and hence there is no ambiguity in whether a triple is right or left subordinate to a band, as there was in triples subordinate to a link. As an analog of switching between

right and left subordinate diagrams for a framed link, we will prove a relation between the α–band maps and the β–band maps in Proposition 4.7.8.

We also make the following remark:

Remark 4.6.9. One can work with Heegaard triples that were obtained from triples subordinate to a band by winding some of the curves to achieve admissibility. We note that using the basic properties of the band maps, as well as the fact that the top degree generators $\Theta^{\mathbf{w}}_{\alpha\beta}$ or $\Theta^{\mathbf{w}}_{\beta\gamma}$ are defined as elements of homology (which doesn't depend on the choice of diagrams), one can use such triples for which we have wound some of the $\boldsymbol{\alpha}, \boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ curves in order to achieve admissibility to compute band maps.

4.6.3 Maps for w–bands

Suppose that B^α is an α–band for a link $\mathbb{L}$ in Y. Let w and w' be the $\mathbf{w}$–basepoints adjacent to the ends of B^α. Again by Lemma 4.3.2, there is a distinguished generator as a $\mathbb{Z}_2[U_{\mathbf{w}}]$–module

$$\Theta^{\mathbf{z}}_{\alpha\beta} \in \mathcal{H}FL^-_{\max,\mathbf{z}}(\mathcal{H}_{\alpha\beta}, \sigma, \mathfrak{P}, \mathfrak{s}_0) \subseteq \mathcal{H}FL^-(\mathcal{H}_{\alpha\beta}, \sigma, \mathfrak{P}, \mathfrak{s}_0)$$

where $\mathcal{H}FL^-_{\max,\mathbf{z}}$ denotes the subgroup of top homological degree with respect to the grading $\mathrm{gr}_{\mathbf{z}}$, as long as $(\sigma, \mathfrak{P})$ assigns the same color to w and w'.

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple subordinate to a band. For $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ we can define

$$F^{\mathbf{w}}_{B^\alpha}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\Theta^{\mathbf{z}}_{\alpha\beta} \otimes \cdot) : \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}).$$

Analogously for a β–band B^β, we can pick a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ subordinate to B^β and define

$$F^{\mathbf{w}}_{B^\beta}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\cdot \otimes \Theta^{\mathbf{z}}_{\beta\gamma}) : \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}),$$

the map which counts holomorphic triangles with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s} \in \mathrm{Spin}^c(X_{\alpha\beta\gamma})_0 \cong \mathrm{Spin}^c(Y)$. Again $\mathrm{Spin}^c(X_{\alpha\beta\gamma})_0$ denotes the Spin^c structures which restrict to the torsion one on $Y_{\beta\gamma}$.

The same argument used for the $\mathbf{z}$ band maps can be used to show that the $\mathbf{w}$ yield well defined maps up to equivariant, filtered chain homotopy on $\mathcal{C}FL^\circ$.

4.7 Properties and relations involving the band maps

In this section, we prove some important properties of the band maps.

4.7.1 Band maps and quasi-stabilization

We now consider trying to commute the band maps and the quasi-stabilization maps. One can often do this, but there are several cases where the bands don't commute, as one would expect, given the interpretation of band maps in terms of dividing sets on surfaces.

Proposition 4.7.1. *Suppose that B is a band for the link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and w and z are two new basepoints which are adjacent in $L \setminus (\mathbf{w} \cup \mathbf{z} \cup B)$. Then, regardless of the configuration, one has*

$$F^{\mathbf{z}}_B S^{\pm}_{w,z} \simeq S^{\pm}_{w,z} F^{\mathbf{z}}_B, \qquad \text{if } w \text{ follows } z, \text{ and}$$

$$F^{\mathbf{z}}_B S^{\pm}_{z,w} \simeq S^{\pm}_{z,w} F^{\mathbf{z}}_B, \qquad \text{if } z \text{ follows } w.$$

If the pair (w, z) does not separate the ends of the band B from either of $\mathbf{z}$–basepoints adjacent to B, then

$$F^{\mathbf{z}}_B T^{\pm}_{w,z} \simeq T^{\pm}_{w,z} F^{\mathbf{z}}_B, \qquad \text{if } w \text{ follows } z, \text{ and}$$

$$F^{\mathbf{z}}_B T^{\pm}_{z,w} \simeq T^{\pm}_{z,w} F^{\mathbf{z}}_B, \qquad \text{if } z \text{ follows } w.$$

Proof. This will follow from the triangle map computations of performed in Theorem 4.5.3, Proposition 4.5.6, and Lemma 4.5.1 and (this last one is [MO10, Proposition 5.2]). Let us first consider the case that $B = B^\alpha$ is an α–band. Pick a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ which is subordinate to B^α.

Let us first consider the case that z follows w. In this case, the adjacency assumption required when considering the $T^{\pm}_{z,w}$ maps is automatically satisfied since the basepoints w

and z cannot be in the component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which contain the ends of B^α (as those components are going from $\mathbf{w}$–basepoints to $\mathbf{z}$–basepoints). The maps $S^{\pm}_{z,w}$ and $T^{\pm}_{z,w}$ are computed using a β–quasi-stabilization, i.e. we add a small α_0 curve, and a β_s curve which is allowed to travel around the diagram. Since the pair (z, w) is not adjacent to the ends of B^α, there is a basepoint $w_0 \in \mathbf{w}$, adjacent to z on $L \setminus B^\alpha$. We can now quasi-stabilize the triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ by adding a small α_0 and β_0 encircling z and w, and a γ_s which encircles z and w_0, as in the top left diagram of Figure 4.14. Now applying Lemma 4.5.1, we know that for appropriately stretched almost complex structure, with respect to the decomposition given by the $\mathrm{gr}_{\mathbf{w}}$–grading of the new intersection points on the quasi-stabilized triple, the quasi-stabilized band map takes the form

$$F^{\mathbf{z}}_{B^\alpha} = \begin{pmatrix} F^{\mathbf{z}}_{B^\alpha} & 0 \\ 0 & F^{\mathbf{z}}_{B^\alpha} \end{pmatrix}.$$

Hence $F^{\mathbf{z}}_{B^\alpha}$ commutes with both $S^{\pm}_{z,w}$ and $T^{\pm}_{z,w}$.

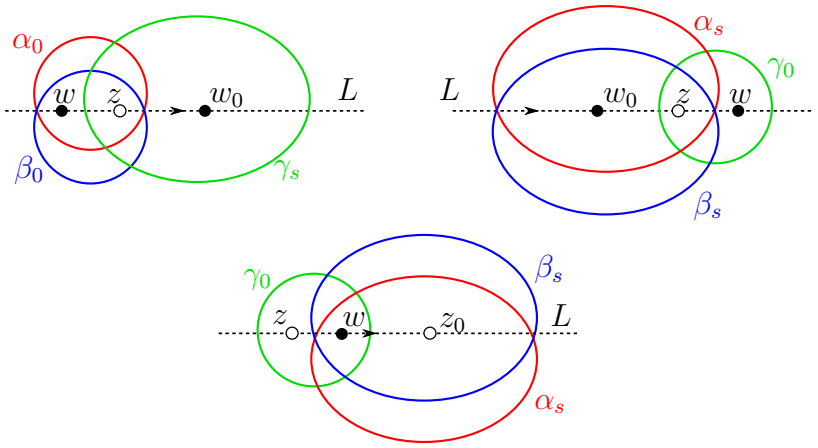

Figure 4.14: **Three configurations of quasi-stabilizations considered in Proposition 4.7.1.** The dashed lines represent shadows of the link L.

We now consider the case that w follows z. Write $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_n\}$ and $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_n\}$ where $\alpha_1, \dots, \alpha_{n-1}$ are small Hamiltonian isotopies of $\beta_1, \dots, \beta_{n-1}$, and α_n and β_n are as in the construction of the band maps.

Suppose first that the pair (w, z) is not between the ends of the band B^α and one of z_1

and z_2. In this case, there is basepoint $w_0 \in \mathbf{w}$ such that w_0 is adjacent to z on $L \setminus B^\alpha$. In this case, we can form the quasi-stabilized diagram for the band map as in the top right diagram of Figure 4.14. The quasi-stabilized triple shown there is now a triple subordinate to the band B^α, added to the quasi-stabilized link. We claim that $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w})$ is strongly positive with respect to the point p_s^+, the higher $\mathrm{gr}_{\mathbf{w}}$ graded generator of $\alpha_s \cap \beta_s$. To see this, we apply Lemma 3.6.3, which shows that if $\mathcal{H}$ is a diagram with a chosen point $p^+ \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $\overline{\mathcal{H}}$ is the diagram obtained by performing surgery on a 0–sphere, then $\mathcal{H}$ is strongly positive at p^+ iff $\overline{\mathcal{H}}$ is strongly positive at p^+. For $i \in \{1, \dots, n-1\}$, the pair of curves α_i and β_i determine an embedded 2–sphere. Upon surgering all of these out and forgetting about the $\mathbf{z}$–basepoints, there are two possibilities for the resulting diagram, depending on the configuration of the basepoints.

The first subcase is that w_0 is adjacent on $(L, \mathbf{w}, \mathbf{z})$ to an end of B^α. Given the ordering of w and z, this is equivalent to the pair (w, z) being adjacent to an end of B^α. The second subcase is that the basepoint w_0 is not adjacent to the ends of B^α, which is equivalent to the pair (w, z) not being adjacent to an end of the band.

In the first subcase, where w_0 is adjacent to an end of B^α, the resulting diagram is

$$(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w_0, w_0', w)$$

shown in Figure 4.15, where w_0' is the other $\mathbf{w}$–basepoint adjacent to the end of B^α. The curves α_n and β_n are isotopic relative the $\mathbf{w}$–basepoints in the diagram, and hence determine a framed 2–sphere. Upon surgering this out, we are left with the diagram $(S^2, \alpha_s, \beta_s, w, w_0)$, which is clearly strongly positive at p_s^+.

The other subcase we need to consider is if w_0 is not adjacent to an end of B^α. Given the ordering of the basepoints w and z, this is equivalent to (w, z) not being adjacent to the ends of the band B^α. then the resulting diagram is

$$(S^2, \alpha_n, \beta_n, w_1, w_2) \sqcup (S^2, \alpha_s, \beta_s, w_0, w)$$

where w_1 and w_2 are now the two basepoints adjacent to the ends of B^α. The curves α_n and β_n each intersect each other twice, and similarly α_s and β_s both intersect each other twice.

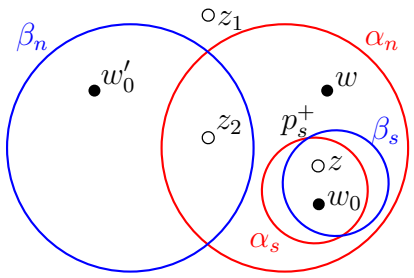

Figure 4.15: **The first subcase, where (w, z) is between an end of B^α and the basepoint w_0.** Once we surger out the 2–spheres determined by the pairs α_i and β_i for $1 \leq i \leq n-1$, we are left with the diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$ shown. The basepoints z_1 and z_2 are the $\mathbf{z}$–basepoints adjacent to the ends of B^α.

Consider now the case shown in Figure 4.15, where (w, z) is between the basepoint $w_0 \in \mathbf{w}$ and an end of the band B^α. In the second case, where w_0 is not adjacent to an end of B^α, the resulting diagram is the disjoint union of two diagrams on genus zero Heegaard surfaces, each with a pair of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves intersecting exactly twice. This diagram is obviously strongly positive at p_s^+. Hence the original diagram was as well, since it can be obtained by surgering on a collection of 0–spheres. In both cases, the triangle computation of Theorem 4.5.3 implies that the quasi-stabilized triangle map for $F^{\mathbf{z}}_{B^\alpha}$ will take the form

$$F^{\mathbf{z}}_{B^\alpha} = \begin{pmatrix} F^{\mathbf{z}}_{B^\alpha} & 0 \\ 0 & F^{\mathbf{z}}_{B^\alpha} \end{pmatrix},$$

and hence will commute with both the $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ maps.

We now have to consider the final case, which is when the pair (w, z) is between an end of the band B^α and one of the $\mathbf{z}$–basepoints adjacent to the ends of B^α. Call this basepoint z_0. In this case, we can quasi-stabilize near z_0, as in the bottom diagram of in Figure 4.14. If p_s^+ is the higher $\mathrm{gr}_{\mathbf{w}}$–graded intersection point of $\alpha_s \cap \beta_s$, we note that the diagram $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ is no longer strongly positive at p_s^+. Indeed we proceed as before. The proof of Lemma 3.6.3 (giving a procedure for constructing strongly positive diagrams) applies equally well to weak positivity, instead of just strong positivity, and hence we can surger out all of the 2–spheres determined by the pairs α_i and β_i for $1 \leq$

$i \leq n-1$, leaving the diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$, for basepoints $w_1, w_2 \in \mathbf{w}$. The diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$ is shown in Figure 4.16, and once we forget about the $\mathbf{z}$–basepoints, we can view it as a free-stabilization at w with the curves α_s and β_s.

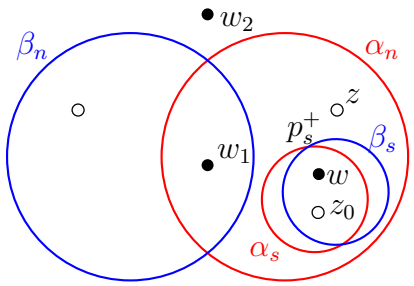

Figure 4.16: **Another case in showing when quasi-stabilization and band maps commute** In the case that (w, z) is between an end of B^α and a $\mathbf{z}$ basepoint, once we surger out the 2–spheres determined by the pairs α_i and β_i for $1 \leq i \leq n-1$, we are left with the diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$.

By Lemma 4.5.5, the pair $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$ is thus weakly positive at p_s^+. As we remarked, surgering out the 2–spheres induced by the pairs α_i and β_i doesn't affect weak positivity, by the same argument as Lemma 3.6.3, and hence we conclude that the original diagram $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ is weakly positive at p_s^+. By Proposition 4.5.6, we can thus write

$$F_{B^\alpha}^{\mathbf{z}} = \begin{pmatrix} F_{B^\alpha}^{\mathbf{z}} & * \\ 0 & F_{B^\alpha}^{\mathbf{z}} \end{pmatrix},$$

with respect to the decomposition given by higher and lower generators, with respect to the grading $\mathrm{gr}_{\mathbf{w}}$. This now commutes with $S_{w,z}^{\pm}$ since the bottom left entry is zero, but not necessarily $T_{w,z}^{\pm}$ since the top right entry is not zero.

If B is instead a β–band, the above proof adapts using the obvious adaptations of the triangle lemmas involved obtained by switching the roles of the $\boldsymbol{\alpha}$–, $\boldsymbol{\beta}$– and $\boldsymbol{\gamma}$–curves appropriately. □

Remark 4.7.2. Notice that the case which was not covered by the previous proposition is when the pair (w,z) separates an end of B from one of the $\mathbf{z}$–basepoints adjacent to B. If (w,z) separates B from an adjacent $\mathbf{z}$–basepoint, then $F_B^{\mathbf{z}}T^+_{w,z} + T^+_{w,z}F_B^{\mathbf{z}} \simeq S^+_{w,z}F_B^{\mathbf{w}}$ (in particular $T^+_{w,z}$ and $F_B^{\mathbf{z}}$ do not commute). We will prove this in Lemma 4.7.17, after we prove a few more useful relations. See Figure 4.25, which shows that the interpretation of these two maps as dividing sets is different.

By symmetry (or just rewriting the above proof), we have the following:

Proposition 4.7.3. *Suppose that B is a band for the link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and w and z are two new basepoints which are adjacent in $L \setminus (\mathbf{w} \cup \mathbf{z} \cup B)$. Then, regardless of the configuration, one has*

$$F_B^{\mathbf{w}}T^{\pm}_{w,z} \simeq T^{\pm}_{w,z}F_B^{\mathbf{w}}, \qquad \text{if } w \text{ follows } z, \text{ and}$$

$$F_B^{\mathbf{w}}T^{\pm}_{z,w} \simeq T^{\pm}_{z,w}F_B^{\mathbf{w}}, \qquad \text{if } z \text{ follows } w.$$

If the pair (w,z) does not separate the ends of the band B from either of $\mathbf{w}$–basepoints adjacent to B, then

$$F_B^{\mathbf{w}}S^{\pm}_{w,z} \simeq S^{\pm}_{w,z}F_B^{\mathbf{w}}, \qquad \text{if } w \text{ follows } z, \text{ and}$$

$$F_B^{\mathbf{w}}S^{\pm}_{z,w} \simeq S^{\pm}_{z,w}F_B^{\mathbf{w}}, \qquad \text{if } z \text{ follows } w.$$

4.7.2 Band maps and disk stabilization

The band maps correspond to index 1 critical points on the surface, and the maps $S^{\pm}_{\mathbb{U},D}$ correspond to index 0 and index 2 critical points. In this section we prove that these maps can cancel each other, in an appropriate sense.

Suppose that $\mathbb{U} = (U, w, z)$ is an unknot with two basepoints, spanned by the disk $D \subseteq Y$, which is disjoint from L. Suppose that we attach a band B to $L \cup U$. The surface $B \cup D$, yields an isotopically unique diffeomorphism

$$\phi : (Y, L) \to (Y, (L \cup U)(B)),$$

which is the identity outside of a neighborhood of $B \cup D$. Let $\overline{B}$ denote the dual band, attached to $(Y, L \cup U)(B)$, yielding the link $L \cup U$. We prove the following:

Proposition 4.7.4. *Let B and $\overline{B}$ be the bands described in the previous paragraph. If B is an α–band then we have*

$$F^{\mathbf{z}}_B S^+_{(U,w,z),D} \simeq S^+_{w,z}\phi_* \qquad \text{and} \qquad F^{\mathbf{w}}_B S^+_{(U,w,z),D} \simeq T^+_{w,z}\phi_*.$$

Similarly,

$$S^-_{(U,w,z),D} F^{\mathbf{z}}_{\overline{B}} \simeq S^-_{w,z}\phi_* \qquad \text{and} \qquad S^-_{(U,w,z),D} F^{\mathbf{z}}_{\overline{B}} \simeq T^-_{w,z}\phi_*.$$

If B is a β–band, we have

$$F^{\mathbf{z}}_B S^+_{(U,w,z),D} \simeq S^+_{z,w}\phi_* \qquad \text{and} \qquad F^{\mathbf{w}}_B S^+_{(U,w,z),D} \simeq T^+_{z,w}\phi_*,$$

and

$$S^-_{(U,w,z),D} F^{\mathbf{w}}_{\overline{B}} \simeq T^-_{z,w}\phi_* \qquad \text{and} \qquad S^-_{(U,w,z),D} F^{\mathbf{w}}_{\overline{B}} \simeq T^-_{z,w}\phi_*.$$

The core of the argument are two triangle map computations. Consider the Heegaard triple shown in Figure 4.17, which we view as a modified free-stabilization of a Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ at a point $z' \in \mathbf{z}$. In the following two triangle map lemmas, we don't require that $\boldsymbol{\alpha}'$ be related to $\boldsymbol{\alpha}$ in any way (other than assuming the triple assumes the appropriate admissibility assumption). Since we will only apply the lemma in the case that $\boldsymbol{\alpha}'$ is a result of a small Hamiltonian isotopy of $\boldsymbol{\alpha}$, we label the first set of attaching curves as $\boldsymbol{\alpha}'$ for expository purposes. Additionally, recall that for the $\mathbf{z}$–band maps to be defined, we need the $\mathbf{z}$–basepoints adjacent to the end of the band to have the same color. As such, in the following two Lemmas, we will identify both V_z and $V_{z'}$ with a formal variable V.

Lemma 4.7.5. *If $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with $z' \in \mathbf{z}$, and $\overline{\mathcal{T}}$ denotes the stabilization shown in Figure 4.17, near the point z', then if $\mathfrak{s} \in \mathrm{Spin}^c(X_{\alpha'\alpha\beta})$ we have identifications*

$$F_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}(\mathbf{x} \times x^+, \mathbf{y} \times \cdot) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) & * \\ 0 & V \cdot F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) \end{pmatrix},$$

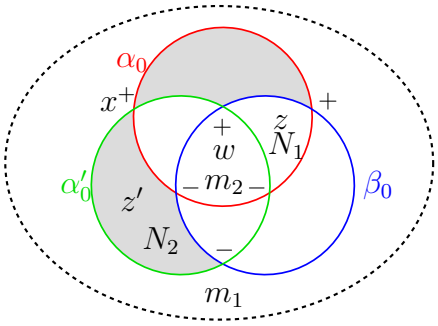

Figure 4.17: **A Heegaard triple to show that disk-stabilization, followed by a band map is a quasi-stabilization.** Shown is a modified free-stabilization of a Heegaard triple $(\Sigma, , \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ at a point $z' \in \mathbf{z}$. Multiplicities and integers m_1, m_2, N_1 and N_2 are shown. This is considered in Lemma 4.7.5. Triangles like this can be used to compute the α–band map for a band which merges an unknot with two basepoints onto a link component. The two shaded regions are each examples of small triangles, which could be counted.

for almost complex structure sufficiently stretched along the dashed circle, and we have identified $V = V_z = V_{z'}$. Here $x^+ \in \alpha_0' \cap \alpha_0$, where here the matrix is an expansion in the intersection point $\cdot \in \alpha_0 \cap \beta_0$, with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

Proof. The proof is similar to the proof of Theorem 2.5.17. For a homology triangle $\psi_1 \# \psi_2 \in \pi_2(\mathbf{x} \times x^+, \mathbf{y} \times y, \mathbf{z} \times z)$, we have that

$$\mu(\psi_1 \# \psi_2) = \mu(\psi_1) + 2m_2 + \Delta_{\mathbf{w}}(y, z),$$

where $\Delta_{\mathbf{w}}(y, z)$ denotes the drop in $\mathrm{gr}_{\mathbf{w}}$ grading from y to z. By stretching the neck sufficiently we can ensure that $\mu(\psi_1)$ and $2m_2$ are nonnegative. In particular if $\Delta_{\mathbf{w}}(y, z) = 1$, there are no holomorphic triangles representing a $\psi_1 \# \psi_2$ with Maslov index 0, so the lower left entry of the matrix is zero. We now consider the diagonal entries, where $\Delta_{\mathbf{w}}(y, z) = 0$. We are left with $\mu(\psi_1) = m_2 = 0$. Using the identical argument as in the proof of Theorem 2.5.17, we see that for fixed y and z in the same grading, the number of points in $\mathcal{M}(\psi_1)$ is equal to the total number of points in moduli spaces of the form $\mathcal{M}(\psi_1 \# \psi_2)$ where ψ_2 ranges over triangles in $\pi_2(x^+, y, z)$ with $m_2(\psi_2) = 0$. We must now just pay attention to

the multiplicities of these triangles at the $\mathbf{z}$–basepoints. Using the multiplicities in Figure 4.17, for a homology triangle $\psi_2 \in \pi_2(x^+, y^+, z^+)$, we claim that

$$m_1 + m_2 = N_1 + N_2.$$

To see this, note that this equation holds for small triangle in $\pi_2(x^+, y^+, z^+)$ shown in Figure 4.17 (the top small triangle). Any other homology element in $\pi_2(x^+, y^+, z^+)$ in this diagram can be obtained by splicing in α_0'–, α_0– or β_0–boundary degenerations of Maslov index $+2$ or -2 (the obvious disks with boundary on α_0', α_0 or β_0, with multiplicity $+1$ or -1, respectively), which changes both $m_1 + m_2$ and $N_1 + N_2$ both by 1 or both by -1, respectively. From this observation the equality $m_1 + m_2 = N_1 + N_2$ follows for any triangle in $\pi_2(x^+, y^+, z^+)$. The triangle ψ_1 has $V_{z'}$ multiplicity m_1, and we thus conclude that any $\psi_1 \# \psi_2$ which is counted has total V_z and $V_{z'}$ multiplicity $N_1 + N_2 = m_1 + m_2 = m_1$, from which we conclude the top left entry of the matrix is as stated.

Analogously, to see that the bottom right entry is as stated, we observe that if $\psi_2 \in \pi_2(x^+, y^-, z^-)$, then

$$N_1 + N_2 = m_1 + m_2 + 1,$$

since it holds for the formula holds for the stated small triangle, and respects splicing in boundary degeneration homology classes. Using the same reasoning as before, we conclude that the bottom right entry has an extra factor of $V = V_z = V_{z'}$. □

Analogously, we will also have use for the following lemma:

Lemma 4.7.6. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with $z' \in \mathbf{z}$, and $\overline{\mathcal{T}}$ denotes the stabilization shown in Figure 4.18, near the point z'. Then if $\mathfrak{s} \in \mathrm{Spin}^c(X_{\alpha'\alpha\beta})$, we have an identification*

$$F_{\overline{\mathcal{T}},\bar{\mathfrak{s}}}(\mathbf{x} \times x^+, \mathbf{y} \times \cdot) = \begin{pmatrix} V \cdot F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) \end{pmatrix},$$

for almost complex structure sufficiently stretched along the dashed circle, assuming we have identified $V = V_z = V_{z'}$. Here $x^+ \in \alpha_0' \cap \alpha_0$ and the matrix expansion is in the intersection point $\cdot \in \alpha_0 \cap \beta_0$, with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

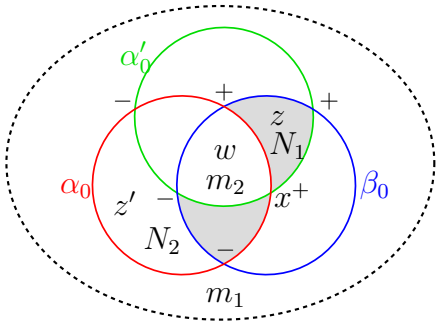

Figure 4.18: **A Heegaard triple to show that a band map followed by a disk-destabilization is a quasi-destabilization.** Shown is a modified free-stabilization of a Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ at a point $z' \in \mathbf{z}$. Multiplicities m_1, m_2, N_1 and N_2 are shown. This is considered in Lemma 4.7.6, and can be used to compute an α–band map for an α–band which splits off an unknot with two basepoints from a link component. The two shaded regions are each small triangles which might be counted.

Proof. The proof is nearly identical to the proof of Lemma 4.7.5. When we stretch the almost complex structure, the same degenerations occur as in that Lemma, but the only change is the multiplicity of these triangles over the $\mathbf{z}$–basepoints. If $\psi_2 \in \pi_2(x^+, y^+, z^+)$, then we have

$$N_1 + N_2 = m_1 + m_2 + 1,$$

by the same argument as before (one checks that the formula holds for a small triangle, then check that it respects splicing in Maslov index ± 2 boundary degenerations). Using the same reasoning as in Lemma 4.7.5, we conclude that the entry in the top left corner has an extra factor of V_z. Similarly, if $\psi_2 \in \pi_2(x^+, y^-, z^-)$, then we have

$$N_1 + N_2 = m_1 + m_2,$$

and so the same reasoning as before shows that the each holomorphic triangle counted in the bottom right entry will have the same combined V_z and $V_{z'}$ multiplicity as the corresponding triangle had in $F_{\mathcal{T},\mathfrak{s}}$. □

Proof of Proposition 4.7.4. We first prove that $F^{\mathbf{z}}_B S^+_{(U,w,z),D} \simeq S^+_{w,z}\phi_*$ in the case that B is an α–band. This follows from the model triangle map computation in Lemma 4.7.5, as follows.

Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ be a Heegaard diagram for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$. Let $\boldsymbol{\alpha}'$ be small Hamiltonian isotopies of the $\boldsymbol{\alpha}$ curves. We can perform disk stabilization to get a diagram $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$. We can assume that the band B is embedded in Σ, and define α_0' to be as in Figure 4.17. The triple $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_0'\}, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\})$ is a triple subordinate to the band B. By Lemma 4.7.5, we can write

$$F_B^{\mathbf{z}} = \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} & * \\ 0 & V \cdot \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} \end{pmatrix},$$

where V is identified with both V_z and $V_{z'}$ (recall we need the complexes colored so that z and z' have the same color in order for the band maps to be defined).

Noting that to compute $S^+_{w,z}$, we need to stretch the almost complex structure around an α_0 curve so that the diagram is quasi-stabilized at w and z. Using the change of almost complex structure computation from Lemma 4.3.23, we know that we can ensure that this map takes the form

$$\Phi_{J \to J'} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Hence we have that

$$\Phi_{J \to J'} F_B^{\mathbf{z}} S^+_{(U,w,z),D} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} & * \\ 0 & V \cdot \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} \\ 0 \end{pmatrix},$$

which is tautologically just $S^+_{w,z} \phi_*$.

To show $S^-_{(U,w,z),D} F^{\mathbf{z}}_B \simeq S^-_{w,z} \phi_*$, one follows similar reasoning, though one needs to use the triangle map computation from Lemma 4.7.6. The statement about $\mathbf{w}$–bands follows by symmetry, and the statements about β–bands follow from an easy modification. □

4.7.3 Commuting band maps I

In this subsection, we prove that band maps can always be commuted, if one is an α–band, and the other is a β–band. We will later consider commuting two band maps when both bands are α–bands or both are β–bands.

Proposition 4.7.7. *Suppose that B_1^α is an α–band and B_2^β is a β–band such that $\mathbb{L}(B_1^\alpha), \mathbb{L}(B_2^\beta)$ and $\mathbb{L}(B_1^\alpha, B_2^\beta)$ all have at least two basepoints on each component. Then, for any choice of $\mathbf{o}_1, \mathbf{o}_2 \in \{\mathbf{w}, \mathbf{z}\}$, we have*

$$F^{\mathbf{o}_1}_{B_1^\alpha} F^{\mathbf{o}_2}_{B_2^\beta} \simeq F^{\mathbf{o}_2}_{B_2^\beta} F^{\mathbf{o}_1}_{B_1^\alpha}.$$

Proof. This follows from the associativity relations for Heegaard Floer homology. Pick curves as follows. Let $(\Sigma, \alpha_1, \dots, \alpha_{n-1}, \beta_1, \dots, \beta_{n-1}, \mathbf{w}, \mathbf{z})$ represent $Y \setminus N(L \cup B_1^\alpha \cup B_2^\beta)$. Let $\alpha_1', \dots, \alpha_{n-1}'$ be small Hamiltonian isotopies of the $\alpha_1, \dots, \alpha_{n-1}$ curves, and let $\beta_1', \dots, \beta_{n-1}'$ be small Hamiltonian isotopies of the $\beta_1, \dots, \beta_{n-1}$ curves. Pick remaining curves $\alpha_n, \beta_n, \alpha_n'$ and β_n' such that

- $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L})$
- $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1', \dots, \beta_n', \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B_2^\beta))$
- $(\Sigma, \alpha_1', \dots, \alpha_n', \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B_1^\alpha))$,
- and $(\Sigma, \alpha_1', \dots, \alpha_n', \beta_1', \dots, \beta_n', \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B_1^\alpha, B_2^\beta))$.

Let us write $\boldsymbol{\alpha}' = (\alpha_1', \dots, \alpha_n')$ and $\boldsymbol{\beta}' = (\beta_1', \dots, \beta_n')$. Let $\Theta^{\mathbf{o}_1}_{\alpha'\alpha}$ denote the algebraic generator of $\mathcal{H}FL^-_{\max,\mathbf{o}_1}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ and let $\Theta^{\mathbf{o}_2}_{\beta\beta'}$ be the algebraic generator of $\mathcal{H}FL^-_{\max,\mathbf{o}_2}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}')$.

Noting that

$$\mathrm{Spin}^c(X_{\alpha'\beta\beta'})_0 \cong \mathrm{Spin}^c(X_{\alpha'\alpha\beta})_0 \cong \mathrm{Spin}^c(X_{\alpha'\alpha\beta'})_0 \cong \mathrm{Spin}^c(X_{\alpha\beta\beta'})_0 \cong \mathrm{Spin}^c(Y),$$

where $\mathrm{Spin}^c(X)_0$ denotes the Spin^c structures which restrict to the torsion one on the appropriate boundary component corresponding to $(S^1 \times S^2)^{\#k}$, we have by the associativity relations for the Heegaard Floer triangle maps that

$$F_{\alpha'\beta\beta',\mathfrak{s}}(F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha'\alpha}, \mathbf{x}), \Theta^{\mathbf{o}_2}_{\beta\beta'}) \simeq F_{\alpha'\alpha\beta',\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha'\alpha}, F_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}, \Theta^{\mathbf{o}_2}_{\beta\beta'}))$$

for a choice of $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, thus showing that

$$F^{\mathbf{o}_1}_{B_1^\alpha} F^{\mathbf{o}_2}_{B_2^\beta} \simeq F^{\mathbf{o}_2}_{B_2^\beta} F^{\mathbf{o}_1}_{B_1^\alpha}.$$

□

4.7.4 Band maps and basepoint moving maps

We now consider the relation between α–bands and β–bands. Whether a band B is an α–band or a β–band is determined by the configuration of the ends of B with respect to the basepoints, and hence a band B can never be both an α–band and a β–band. On the other hand, given an α–band B^α, we can isotope the ends of B^α across the adjacent $\mathbf{z}$–basepoints to get a β–band B^β. Let z_1 and z_2 denote the adjacent $\mathbf{z}$–basepoints. There is a well defined diffeomorphism $\phi : (Y, \mathbb{L}(B^\alpha)) \to (Y, \mathbb{L}(B^\beta))$ which fixes all basepoints in $\mathbf{w} \cup \mathbf{z}$ except z_1 and z_2. The diffeomorphism ϕ switches z_1 and z_2. This is illustrated in Figure 4.19.

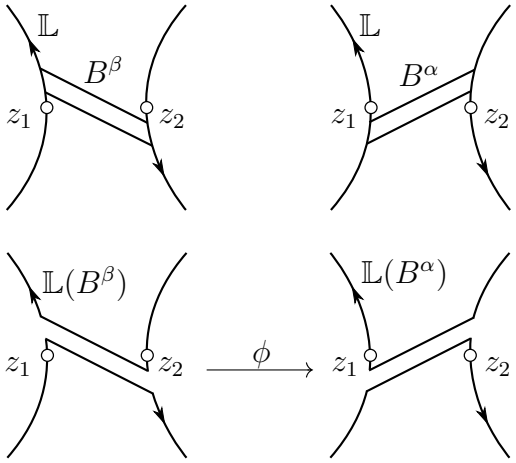

Figure 4.19: **Two bands B^α and B^β related by an isotopy across z–basepoints,** and the diffeomorphism ϕ between $(Y, \mathbb{L}(B^\alpha))$ and $(Y, \mathbb{L}(B^\beta))$, which maps z_1 to z_2 and maps z_2 to z_1.

Proposition 4.7.8. *If B^α, B^β and ϕ are as in the previous paragraph, then*

$$\phi_* F^{\mathbf{z}}_{B^\alpha} \simeq F^{\mathbf{z}}_{B^\beta}.$$

As a step towards proving the above proposition, consider the following situation. Suppose that D is a disk embedded in $Y \setminus L$, w and z are basepoints on $\partial D = U$, and B^α_0, B^β_0 are bands connecting D to $\mathbb{L}$, as in Figure 4.20. Define diffeomorphisms ϕ_1 and ϕ_2, as in Figure 4.20. Also let ϕ_0 be the basepoint moving map, shown in Figure 4.20.

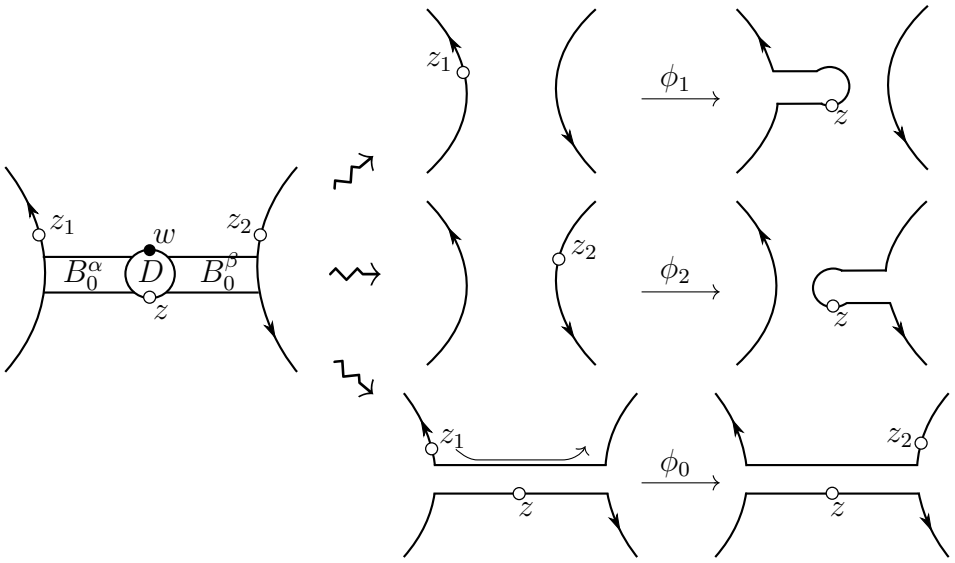

Figure 4.20: **The disk D, the bands B_0^α and B_0^β, and the diffeomorphisms ϕ_0, ϕ_1 and ϕ_2.**

Note that to prove Proposition 4.7.8, using diffeomorphism invariance of the band maps, it is sufficient to show

$$(\phi_0)_* F^{\mathbf{z}}_{B_0^\alpha}(\phi_2)_* \simeq F^{\mathbf{z}}_{B_0^\beta}(\phi_1)_*. \tag{4.5}$$

To this end, we first prove the following:

Lemma 4.7.9. *The diffeomorphisms ϕ_0, ϕ_1 and ϕ_2 induce maps*

$$(\phi_0)_* \simeq T^-_{z_1,w} S^+_{w,z_2}, \qquad (\phi_1)_* \simeq T^-_{z_1,w} F^{\mathbf{z}}_{B_0^\alpha} S^+_{(U,w,z),D}, \qquad \text{and} \qquad (\phi_2)_* \simeq T^-_{w,z_2} F^{\mathbf{z}}_{B_0^\beta} S^+_{(U,w,z),D}.$$

Proof. The formula for $(\phi_0)_*$ follows immediately Lemma 4.3.27.

We now consider the formula for $(\phi_1)_*$. Let $\phi^\alpha : (Y, L) \to (Y, (L \cup U)(B_0^\alpha))$ be a diffeomorphism which is fixed outside of a neighborhood of $D \cup B_0^\alpha$. By Proposition 4.7.4 we have that

$$F^{\mathbf{z}}_{B_0^\alpha} S^+_{(U,w,z),D} \simeq S^+_{w,z} \phi^\alpha_*.$$

Post-composing with $T^-_{z_1,w}$ we get

$$T^-_{z_1,w} F^{\mathbf{z}}_{B_0^\alpha} S^+_{(U,w,z),D} \simeq T^-_{z_1,w} S^+_{w,z} \phi^\alpha_*.$$

By Lemma 4.3.27, the composition $T^-_{z_1,w}S^+_{w,z}$ is equivalent to the basepoint moving map $(\phi_{z\to z_1})_*$. Since $\phi_{z\to z_1}\circ\phi^\alpha$ is clearly the diffeomorphism ϕ_1, we thus have

$$T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D}\simeq(\phi_{z\to z_1})_*\phi^\alpha_*\simeq(\phi_1)_*.$$

The claim about $(\phi_2)_*$ follows from identical reasoning to $(\phi_1)_*$. □

We can now prove Proposition 4.7.8:

Proof of Proposition 4.7.8. As we remarked earlier, it is sufficient to show Equation (4.5). Using Lemma 4.7.9, we thus need to show

$$T^-_{z_1,w}S^+_{w,z_2}F^{\mathbf{z}}_{B_0^\alpha}T^-_{w,z_2}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D}\simeq F^{\mathbf{z}}_{B_0^\beta}T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D}.$$

We note that

$$\begin{aligned}T^-_{z_1,w}S^+_{w,z_2}F^{\mathbf{z}}_{B_0^\alpha}T^-_{w,z_2}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D}&\simeq T^-_{z_1,w}S^+_{w,z_2}T^-_{w,z_2}F^{\mathbf{z}}_{B_0^\alpha}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D}&&\text{(Lemma 4.7.1)}\\&\simeq T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D}&&\text{(Lemma 4.7.7).}\end{aligned}$$

On the other hand we have

$$\begin{aligned}F^{\mathbf{z}}_{B_0^\beta}T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D}&\simeq T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\beta}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D}&&\text{(Lemma 4.7.1)}\\&\simeq T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D}&&\text{(Proposition 4.7.7).}\end{aligned}$$

These two expressions are clearly equal, completing the proof. □

Remark 4.7.10. We note that the proof of Proposition 4.7.8 can be interpreted similarly to [OS06, Lemma 5.2], which says that one can can compute maps induced surgery on framed links using either left subordinate or right subordinate Heegaard triples. The argument uses the fact that the 3–handle maps cancel both the left subordinate, and right subordinate maps, for 0–surgery on an unlink. Analogously, our proof of Proposition 4.7.8 follows from a similar observation that the addition of an unknot via the maps $S^+_{\mathbb{U},D}$ can be canceled by either the α–band maps, or the β–band maps.

The analogous statement also holds for the $\mathbf{w}$ band maps. Suppose that B^α is an α–band and B^β is a β–band resulting by sliding the ends of B^α across the two adjacent $\mathbf{w}$–basepoints (call them w_1 and w_2). Then there is a well defined diffeomorphism $\phi : (Y, \mathbb{L}(B^\alpha)) \to (Y, \mathbb{L}(B^\beta))$ which moves w_1 to w_2 along the link (as in Figure 4.19, but with basepoints w_1 and w_2 replacing the z_1 and z_2 which appear in that picture.) We have the following:

Proposition 4.7.11. *If B^α, B^β and ϕ are as in the previous paragraph, then*

$$F^{\mathbf{w}}_{B^\alpha} = \phi_* F^{\mathbf{w}}_{B^\beta}.$$

The proof is, of course, to just rewrite the proof of Proposition 4.7.8, switching Ss with Ts and $\mathbf{w}$s with $\mathbf{z}$s.

4.7.5 Commuting band maps II

We now revisit the question of commuting two bands maps in the case that both are α–bands, or both are β–bands. We phrase the result for α–bands, but of course the proof carries over without change to β–bands.

Proposition 4.7.12. *Suppose that B_1 and B_2 are two α–bands on the link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ such that $\mathbb{L}(B_1), \mathbb{L}(B_2)$ and $\mathbb{L}(B_1, B_2)$ all have at least two basepoints on each component. Then we have*

$$F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} \simeq F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1},$$

and

$$F^{\mathbf{w}}_{B_1} F^{\mathbf{w}}_{B_2} \simeq F^{\mathbf{w}}_{B_2} F^{\mathbf{w}}_{B_1}.$$

Suppose further that there exist points $\mathbf{d} \subseteq L \setminus (\mathbf{w} \cup \mathbf{z})$ such that each adjacent pair of points in $\mathbf{w} \cup \mathbf{z}$ is separated by exactly one point of $\mathbf{d}$. If B_2 has its ends in regions of $L \setminus \mathbf{d}$ containing $\mathbf{w}$–basepoints, and B_1 has its ends in regions of $L \setminus \mathbf{z}$ containing $\mathbf{z}$–basepoints, then

$$F^{\mathbf{z}}_{B_1} F^{\mathbf{w}}_{B_2} \simeq F^{\mathbf{w}}_{B_2} F^{\mathbf{z}}_{B_1}.$$

We will give two proofs of this fact. The first is very fast, and mostly unenlightening. The second is longer, but more enlightening.

Proof 1. Suppose that B_1 and B_2 are two bands, and suppose that we assign type–$\mathbf{z}$ to B_1 and type–$\mathbf{o}_2$ to B_2. The only case of the proposition which isn't covered is the case that we were to assign a designation of type–$\mathbf{w}$ to both B_1 and B_2, though this case is an easy modification of the case where B_1 and B_2 are both assigned type–$\mathbf{z}$.

It is easily verified the assumptions of the Proposition statement imply that there are two $\mathbf{z}$–basepoints which are adjacent to the ends of one of the bands on both $\mathbb{L}$ and $\mathbb{L}(B_2)$. In the case that both B_1 and B_2 are given type–$\mathbf{z}$, we let B_1 denote this band. In the case that only B_1 is of type–$\mathbf{z}$, the configuration of the points in $\mathbf{d}$ forces this band to be B_1. Let B_1' be the β–band obtained by sliding B_1 across the $\mathbf{z}$–basepoints adjacent to B_1. Let $\phi : L(B_1) \to L(B_1')$ be the obvious modification of the diffeomorphism pictured in Figure 4.19. Note that ϕ fixes the band B_2, so we can define a diffeomorphism $\phi^{B_2} : L(B_1)(B_2) \to L(B_1')(B_2)$, which is equal to ϕ as a map on Y. We now apply Propositions 4.7.7 and 4.7.8. We have that

$$\begin{aligned} F^{\mathbf{z}}_{B_1} F^{\mathbf{o}_2}_{B_2} &\simeq \phi^{B_2}_* F^{\mathbf{z}}_{B_1'} F^{\mathbf{o}_2}_{B_2} \\ &\simeq \phi^{B_2}_* F^{\mathbf{o}_2}_{B_2} F^{\mathbf{z}}_{B_1'} \\ &\simeq F^{\mathbf{o}_2}_{B_2} \phi_* F^{\mathbf{z}}_{B_1'} \\ &\simeq F^{\mathbf{o}_2}_{B_2} F^{\mathbf{z}}_{B_1}, \end{aligned}$$

completing the proof. Note that $\phi^{B_2}_* F^{\mathbf{o}_2}_{B_2} \simeq F^{\mathbf{o}_2}_{B_2} \phi_*$ since ϕ fixes the band B_2. □

We offer an alternate proof, which is somewhat more enlightening and more in the spirit with other standard proofs of this type from Heegaard Floer homology. The reader can safely ignore this if they are satisfied with the first proof, though we include it to give an alternate perspective.

Proof 2. We now offer a second proof, which is somewhat more explicit. There are three cases to consider:

1. B_1 and B_2 have ends on four separate components of $L \setminus (\mathbf{w} \cup \mathbf{z})$;

2. B_1 and B_2 each have exactly one end on the same component of $L \setminus (\mathbf{w} \cup \mathbf{z})$, but the other ends are on other components;

3. B_1 and B_2 each have ends in a total of two components of $L \cup (\mathbf{w} \cup \mathbf{z})$.

We first show that the case (2) can be reduced to case (1), and case (3) can be reduced to case (2). Consider case (2). On the component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which contains an end of B_1 and an end of B_2, let $d \in \mathbf{d}$ denote a point separating the bands. We let z and w denote new basepoints near a, ordered so that w comes after z. Let $\mathbf{o}_1, \mathbf{o}_2 \in \{\mathbf{w}, \mathbf{z}\}$ denote the type designations of B_1 and B_2 respectively. If $\mathbf{o}_1 = \mathbf{o}_1 = \mathbf{z}$, then we note that $S^-_{w,z}$ commutes with $F^{\mathbf{z}}_{B_i}$ by Proposition 4.7.1, and hence we simply compute, using the result for case (1), that

$$\begin{aligned}
F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} &\simeq F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} S^-_{w,z} T^+_{w,z} && \text{(Lemma 4.3.13)} \\
&\simeq S^-_{w,z} F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} T^+_{w,z} && \text{(Proposition 4.7.1)} \\
&\simeq S^-_{w,z} F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1} T^+_{w,z} && \text{(Case (1))} \\
&\simeq F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1} S^-_{w,z} T^+_{w,z} && \text{(Proposition 4.7.1)} \\
&\simeq F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1} && \text{(Lemma 4.3.13).}
\end{aligned}$$

The same strategy can be used to reduce Case (3) to Case (2). Using $T^-_{w,z}$ instead of $S^-_{w,z}$ establishes the result in the case that both $\mathbf{o}_1 = \mathbf{o}_2 = \mathbf{w}$. In the case that $\mathbf{o}_1 = \mathbf{w}$ and $\mathbf{o}_2 = \mathbf{z}$, we let w and z be new basepoints as before. We note that the pair (w, z) does not separate B_1 from the $\mathbf{w}$ basepoints adjacent to the ends of B_1, as can be seen in Figure 4.21. Applying Propositions 4.7.1 and 4.7.3 shows that $S^-_{w,z}$ commutes with both $F^{\mathbf{w}}_{B_1}$ and $F^{\mathbf{z}}_{B_2}$ so we can reduce from Case (2) to Case (1), and from Case (3) to Case (2), as before.

Hence we need only to show the statement in the case that the ends of B_1 and B_2 are in separate components of $L \setminus (\mathbf{w} \cup \mathbf{z})$. Note that in this case, the condition of the Proposition

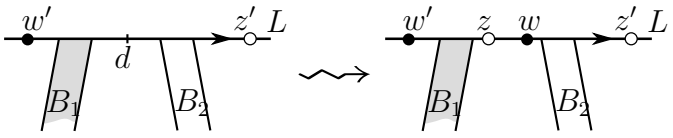

Figure 4.21: **A configuration of basepoints and ends of bands such that the band maps commute.** Shown are the ends of the bands B_1 and B_2, and the adjacent $\mathbf{w}$–basepoint w'. Also shown are the new basepoints w and z, and the point d between the ends of B_1 and B_2.

statement about a set of points $\mathbf{d} \subseteq L \setminus (\mathbf{w} \cup \mathbf{z} \cup B_1 \cup B_2)$ is trivially satisfied (it was only necessary to ensure that the quasi-stabilization maps commuted appropriately, so we could reduce to this case). The remainder of the proof now follows from the associativity relations, and a model computation, as we now describe.

We pick a Heegaard diagram $(\Sigma, \alpha_1, \dots, \alpha_{n-2}, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ which induces a sutured diagram for the sutured manifold $Y \setminus N(L \cup B_1^\alpha \cup B_2^\alpha)$, with sutures induced by the basepoints $\mathbf{w}$ and $\mathbf{z}$. For $i \in \{1, \dots, n-2\}$, let α_i^1, α_i^2 and α_i^{12} be small isotopies of α_i, in such a way that each pair of curves with the same index i intersects in exactly two points, but curves with differing indices do not intersect (e.g. $|\alpha_i \cap \alpha_j^1| = 2\delta_{ij}$).

Let α_{n-1} and α_n be curves such that $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $\mathbb{L}$. Now pick curves α_{n-1}^τ and α_n^τ for $\tau \in \{1, 2, 12\}$ such that

- $(\Sigma, \alpha_1^1, \dots, \alpha_n^1, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B_1^\alpha))$;
- $(\Sigma, \alpha_1^2, \dots, \alpha_n^2, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B_2^\alpha))$;
- $(\Sigma, \alpha_1^{12}, \dots, \alpha_n^{12}, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B_1^\alpha, B_2^\alpha)$;
- α_n^2 is a small isotopy of α_n;
- α_{n-1}^1 is a small isotopy of α_{n-1};
- α_{n-1}^{12} is a small isotopy of α_{n-1}^2;
- α_n^{12} is a small isotopy of α_n^1.

Define $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$, $\boldsymbol{\alpha}^1 = (\alpha_1^1, \dots, \alpha_n^1)$ and $\boldsymbol{\alpha}^2$ and $\boldsymbol{\alpha}^{12}$ similarly. Note that

- $(\Sigma, \boldsymbol{\alpha}^1, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate the band B_1^α, attached to the link $\mathbb{L}$;
- $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^1, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate to the band B_2^α, attached to $\mathbb{L}(B_1^\alpha)$;
- $(\Sigma, \boldsymbol{\alpha}^2, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate to band B_2^α, attached to $\mathbb{L}$;
- $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^2, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate to the band B_1^α, attached to $\mathbb{L}(B_2^\alpha)$.

We note that by construction $(\Sigma, \boldsymbol{\alpha}^1, \boldsymbol{\alpha}, \boldsymbol{\beta})$, $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^1, \boldsymbol{\beta})$, $(\Sigma, \boldsymbol{\alpha}^2, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^2, \boldsymbol{\beta})$ all represent unlinks in (disjoint unions of) $(S^1 \times S^2)^{\#k}$ where each component of the unlink has exactly two basepoint, except for one component, which has four. As such, the associated $\mathcal{H}FL^-_{\max,\mathbf{o}_i}$ groups have a generator of top degree. We write

$$\begin{aligned}
\Theta^{\mathbf{o}_1}_{\alpha_1\alpha} &\in \mathcal{H}FL^-_{\max,\mathbf{o}_1}(\Sigma, \boldsymbol{\alpha}_1, \boldsymbol{\alpha}),\\
\Theta^{\mathbf{o}_2}_{\alpha_2\alpha} &\in \mathcal{H}FL^-_{\max,\mathbf{o}_2}(\Sigma, \boldsymbol{\alpha}_2, \boldsymbol{\alpha}),\\
\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1} &\in \mathcal{H}FL^-_{\max,\mathbf{o}_2}(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_1),\\
\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2} &\in \mathcal{H}FL^-_{\max,\mathbf{o}_1}(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_2),
\end{aligned}$$

for these algebraic generators.

For $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, the associativity relations for the quadruple $(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_1, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ read

$$F_{\alpha_{12}\alpha_1\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1}, F_{\alpha_1\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha_1\alpha}, \mathbf{x})) \simeq F_{\alpha_{12}\alpha\beta,\mathfrak{s}}(F_{\alpha_{12}\alpha_1\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1}, \Theta^{\mathbf{o}_1}_{\alpha_1\alpha}), \mathbf{x}).$$

Similarly the associativity relations for the quadruple $(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_2, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ read

$$F_{\alpha_{12}\alpha_2\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2}, F_{\alpha_2\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_2}_{\alpha_2\alpha}, \mathbf{x})) \simeq F_{\alpha_{12}\alpha\beta,\mathfrak{s}}(F_{\alpha_{12}\alpha_2\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2}, \Theta^{\mathbf{o}_2}_{\alpha_2\alpha}), \mathbf{x}).$$

Hence we have that $F^{\mathbf{o}_1}_{B_1} F^{\mathbf{o}_2}_{B_2} \simeq F^{\mathbf{o}_2}_{B_2} F^{\mathbf{o}_1}_{B_1}$ as long as

$$F_{\alpha_{12}\alpha_1\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1}, \Theta^{\mathbf{o}_1}_{\alpha_1\alpha}) = F_{\alpha_{12}\alpha_2\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2}, \Theta^{\mathbf{o}_2}_{\alpha_2\alpha}),$$

as elements in the homology group $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$.

Let $\boldsymbol{\alpha}'$ be small isotopies of the $\boldsymbol{\alpha}$ curves. Let $\widetilde{B}_1$ and $\widetilde{B}_2$ denote the bands B_1 and B_2, now viewed as bands for unlinks inside of $(S^1 \times S^2)^{\#g}$. Let Θ be the top degree generator of $\mathcal{H}FL^-_{\max}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$. Note that $\Theta^{\mathbf{o}_1}_{\alpha_1\alpha}$ is equal, on homology to

$$F^{\mathbf{o}_1}_{\widetilde{B}_1}(\Theta),$$

because tautologically the triangle map which computes that band map can be interpreted, when applied to Θ, as a change of diagrams map applied to the element $\Theta^{\mathbf{o}_1}_{\alpha_1\alpha}$. Analogously, $\Theta^{\mathbf{o}_2}_{\alpha_2\alpha}$ is equal on homology to $F^{\mathbf{o}_2}_{\widetilde{B}_2}(\Theta)$. Upon these considerations, it is thus sufficient to show that

$$F^{\mathbf{o}_1}_{\widetilde{B}_1} F^{\mathbf{o}_2}_{\widetilde{B}_2}(\Theta) = F^{\mathbf{o}_2}_{\widetilde{B}_2} F^{\mathbf{o}_1}_{\widetilde{B}_1}(\Theta),$$

as elements of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha})$. This can now be done via a model computation. Note that we just need to verify this on homology, for a single choice of diagrams. By using the triangle map computation of Theorem 2.7.8, since we are applying the band maps to Θ, the top degree generator, by surgering out the 2–spheres corresponding to pairs α'_i and α_i for $i \neq n-1, n$, we can assume that Σ is genus zero, and that there are only four components of our link, each with two basepoints, and that the bands connect two pairs of these. Considering the diagram in Figure 4.22, we can eliminate the middle pair of $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ curves by applying Theorem 2.7.8 once again. This reduces the problem to the case that the two pairs of unknots are each in different components of $S^3 \sqcup S^3$, and now commutation is trivial.

□

Remark 4.7.13. We still used some trickery in reducing the complexity of the model computations, in the previous alternate proof, since we used the trick of adding then removing basepoints. If one didn't want to do this, they would need to do model computations for cases (2) and (3) separately. Case (2) would be where there are three unknots, and we want add 2 bands, each with two ends on one of the unknots. One would then need to check that the band maps commuted, when applied to the top degree generator. Case (2) can actually be performed, without too much trouble. One can check the desired commutations on a

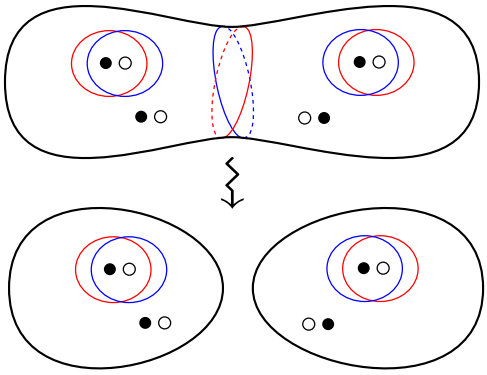

Figure 4.22: **A diagram $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ for unknots in S^3 which we can reduce to for our model band computation.** By eliminating the middle $\boldsymbol{\alpha}'$ and $\boldsymbol{\alpha}$ curves, we reduce to the case that the bands $\widetilde{B}_1$ and $\widetilde{B}_2$ are in different components of $S^3 \sqcup S^3$.

genus zero diagram, with $|\boldsymbol{\alpha}| = |\boldsymbol{\beta}| = 2$ and $|\mathbf{w}| = |\mathbf{z}| = 3$. Using the triangle counts from Theorem 2.7.8 cleverly, one can see that for appropriate almost complex structure, the band maps commute for the configuration described in the above lemma. When the configuration of bands is not as in the previous lemma, the above triangle count actually doesn't apply (we end up counting triangles which correspond to the $*$ component of the triangle map). Case (3) would be to add two bands to two unknots. The resulting link is slightly more complicated. It ends up being a two component link in $S^1 \times S^2$, with each component being isotopic to $S^1 \times \{*\}$, but going in opposite directions. Similar holomorphic triangle counts can be used. We've included a picture showing the starting and ending links in both cases (2) and (3), if one were to use this strategy.

Remark 4.7.14. We now also describe an alternate route to defining the quasi-stabilization maps. The quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ are fairly challenging to define, as they involve challenging holomorphic triangle arguments. An alternate approach to defining them is as follows. One can define the disk stabilization maps $S^{\pm}_{(U,w,z),D}$ and the band maps $F^{\mathfrak{o}}_B$. In light of Proposition 4.7.4, the quasi-stabilization maps are then equal to a composition of disk stabilization, band maps and diffeomorphism maps. Disk stabilization maps commute with each by the same argument that 1–handle maps or free-stabilization maps commute with each other. Commutation (in the appropriate cases) of the quasi-stabilization maps,

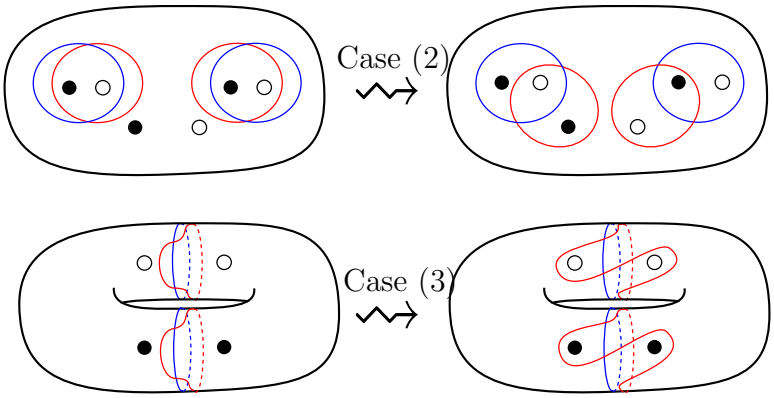

Figure 4.23: **The remaining model computations needed to prove Proposition 4.7.12 without using any basepoint tricks.** Each of the squiggly arrows corresponds to a composition of two band maps. Each band map moves one of the two $\boldsymbol{\alpha}$ curves. In both cases, for appropriately stretched almost complex structure, the maps can be computed by using Theorem 2.7.8, as long as the configuration of the types of the bands is as in the Proposition statement. In the first diagram, we can try to view the diagram as a free stabilization of a diagram with one $\boldsymbol{\alpha}$– and one $\boldsymbol{\beta}$–curve, and in the second diagram, we can try to view the diagram as a 1–handle stabilization of a diagram with with one $\boldsymbol{\alpha}$– and one $\boldsymbol{\beta}$–curve. In each case, there are two ways to view the entire diagram as a stabilization of a smaller diagram. The configuration of basepoints with respect to the ends ensures that one of the configurations can legitimately be viewed as a free stabilization or 1–handle stabilization with respect to one of the gradings $\mathrm{gr}_{\mathbf{w}}$ or $\mathrm{gr}_{\mathbf{z}}$, with two intersection points being the top degree generator. In such a case, we can use the triangle map computation of Theorem 2.7.8 to arrive at the expected triangle count. In the case that the bands are both of the "wrong" type for the configuration of ends, when we try to view the diagrams as a free stabilization or 1–handle stabilization, we will be counting triangles with vertices on the bottom generators. One can explicitly find examples where "unexpected" triangles are counted, resulting in the band maps not commuting.

such as Proposition 4.3.15, would then follow by using the second proof of Proposition 4.7.12 (without any of the basepoint tricks, as in Remark 4.7.13). One would of course need

to perform several triangle map computations to show the relation between this alternate definition of the quasi-stabilization maps and the basepoint moving maps. We will not pursue this further, but it's instructive to know that there are other routes to the theory.

As a final remark, we consider the case that was not covered by Proposition 4.7.12:

Remark 4.7.15. If B_1 and B_2 are two bands added to a link, and the ends of B_2 separate the an end of B_1 from a $\mathbf{z}$–basepoint, then in general we do not expect $F^{\mathbf{z}}_{B_1}$ and $F^{\mathbf{w}}_{B_2}$ to commute. We interpret the two compositions as surfaces with divides in Figure 4.24, and the dividing sets are clearly not diffeomorphic.

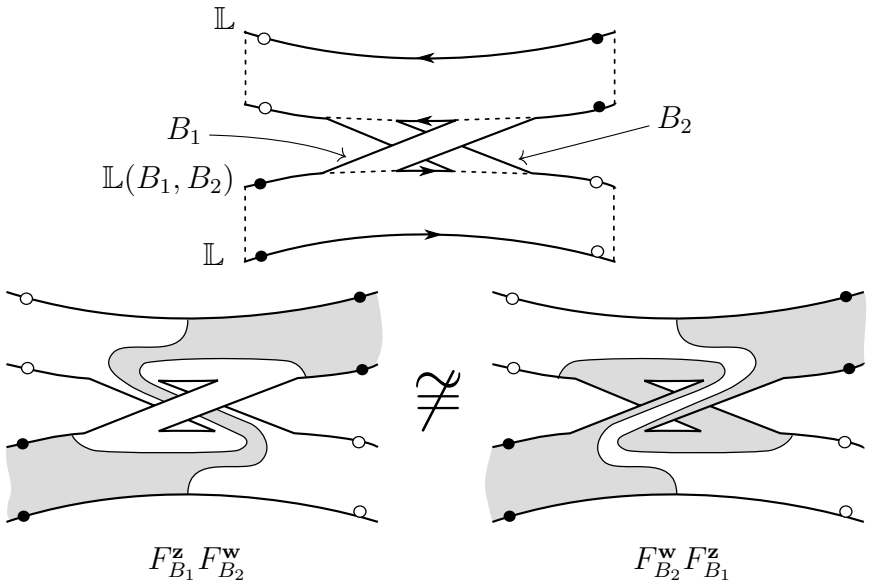

Figure 4.24: On top is the surface formed by adding two α–bands to $L \times [0,1]$, viewed as having incoming boundary $\mathbb{L}$ and outgoing boundary $\mathbb{L}(B_1, B_2)$. Below are the non-diffeomorphic dividing sets for the two compositions $F^{\mathbf{z}}_{B_1} F^{\mathbf{w}}_{B_2}$ and $F^{\mathbf{w}}_{B_2} F^{\mathbf{z}}_{B_1}$ for a configuration of bands not covered in Lemma 4.7.12. This was one of the cases where the argument from Proposition 4.7.12 failed to show that the band maps commuted.

4.7.6 A final relation between the $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ maps.

There is one final relation that the maps $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ satisfy, assuming a restriction on the coloring.

Proposition 4.7.16. *Suppose that B is a band (either α– or β–). Suppose that the ends of the band B are adjacent to the basepoints z_1, z_2, w_1 and w_2. If both w_1 and w_2 have the same color, and both z_1 and z_2 have the same color, then*

$$F_B^{\mathbf{w}} \simeq \Psi_{z_1} F_B^{\mathbf{z}} + F_B^{\mathbf{z}} \Psi_{z_1} \simeq \Psi_{z_2} F_B^{\mathbf{z}} + F_B^{\mathbf{z}} \Psi_{z_2}$$

and

$$F_B^{\mathbf{z}} \simeq \Phi_{w_1} F_B^{\mathbf{w}} + F_B^{\mathbf{w}} \Phi_{w_1} \simeq \Phi_{w_2} F_B^{\mathbf{w}} + F_B^{\mathbf{w}} \Phi_{w_2}.$$

Proof. Let us consider the case that B is an α–band; the case that B is a β–band is analogous. We will show the first equivalence. The second follows from symmetry. Let $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ be a triple subordinate to the band B. The diagram $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ is a diagram of an unlink in $(S^1 \times S^2)^{\#g}$, where all but one component of the unlink has exactly two basepoints, and a single component has four. Let $\Theta^{\mathbf{o}}_{\alpha'\alpha} \in \mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ denote the generator of $\mathcal{H}FL^-_{\max,\mathbf{o}}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$. We claim that $\Psi_{z_i} \Theta^{\mathbf{w}}_{\alpha'\alpha} = \Theta^{\mathbf{z}}_{\alpha'\alpha}$ as elements of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$, for $i = 1, 2$. As $\Theta^{\mathbf{o}}_{\alpha'\alpha}$ is a well defined element of homology, and the maps Ψ_{z_i} are well defined on the coherent chain homotopy types, it is sufficient to show this relation for a particular choice of diagram and almost complex structure. We prove the statement by induction on the number of attaching curves. Note first that when $\Sigma = S^2$ and $|\boldsymbol{\alpha}'| = |\boldsymbol{\alpha}| = 1$, the claim is obvious, since we simply pick $\boldsymbol{\alpha}'$ and $\boldsymbol{\alpha}$ to intersect at exactly two points, dividing $S^2 \setminus (\boldsymbol{\alpha}' \cup \boldsymbol{\alpha})$ into four bigons, each with a single basepoint. There is a single Maslov index one holomorphic disk which goes over a given z_i, and it has multiplicity one and goes from $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ to $\Theta^{\mathbf{z}}_{\alpha'\alpha}$. More generally, we can proceed by induction to show the claim using the differential computation of Lemma 2.7.5. This implies that if α_n and α'_n are two curves on Σ, small isotopies of each other, and intersecting exactly twice, then the differential can be written as

$$\partial = \begin{pmatrix} \partial_0 & * \\ 0 & \partial_0 \end{pmatrix},$$

for appropriate almost complex structure, where ∂_0 is the differential on the diagram obtained by surgering out α'_n and α_n. The matrix decomposition is in terms of the relative grading

of the two intersection points of α'_n and α_n (note that the designation of higher and lower intersection point of $\alpha'_n \cap \alpha_n$ agrees for $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ in this case). Hence

$$\Psi_{z_i} = \begin{pmatrix} \Psi_{z_i} & * \\ 0 & \Psi_{z_i} \end{pmatrix}.$$

Noting that in terms of this decomposition, on the complex with α'_n and α_n, the element $\Theta^{\mathbf{o}}_{\alpha'\alpha}$ is equal to $\begin{pmatrix} \Theta^{\mathbf{o}}_{\alpha'\alpha} \\ 0 \end{pmatrix}$, where the entry $\Theta^{\mathbf{o}}_{\alpha'\alpha}$ inside the vector refers to the top $\mathrm{gr}_{\mathbf{o}}$ generator on the complex without the α'_n and α_n curves. Since $\Psi_{z_i}\Theta^{\mathbf{w}}_{\alpha'\alpha} = \Theta^{\mathbf{z}}_{\alpha'\alpha}$ on the complex with α'_n and α_n surgered out, we see that

$$\begin{pmatrix} \Psi_{z_i} & * \\ 0 & \Psi_{z_i} \end{pmatrix} \begin{pmatrix} \Theta^{\mathbf{w}}_{\alpha'\alpha} \\ 0 \end{pmatrix} = \begin{pmatrix} \Theta^{\mathbf{z}}_{\alpha'\alpha} \\ 0 \end{pmatrix},$$

showing that the same relation holds on the complex with α'_n and α_n present.

Having established that $\Psi_{z_i}\Theta^{\mathbf{w}}_{\alpha'\alpha} = \Theta^{\mathbf{z}}_{\alpha'\alpha}$, we proceed with a computation.

Gromov compactness yields that

$$F_{\alpha'\alpha\beta} \circ (\partial_{\alpha'\alpha} \otimes 1 + 1 \otimes \partial_{\alpha\beta}) + \partial_{\alpha'\beta} \circ F_{\alpha'\alpha\beta} = 0.$$

Differentiating with respect to the variable V_{z_i}, for either $i = 1$ or $i = 2$, we see that the quantity

$$F_{\alpha'\alpha\beta} \circ (\Psi_{z_i} \otimes 1 + 1 \otimes \Psi_{z_i}) + \Psi_{z_i} \circ F_{\alpha'\alpha\beta} + \left(\frac{d}{dV_{z_i}} F_{\alpha'\alpha\beta}\right) \circ (\partial_{\alpha'\alpha} \otimes 1 + 1 \otimes \partial_{\alpha\beta}) + \partial_{\alpha'\beta} \circ \left(\frac{d}{dV_{z_i}} F_{\alpha'\alpha\beta}\right).$$

vanishes Here $\frac{d}{dV_{z_i}} F_{\alpha'\alpha\beta}$ can be thought of as the map which counts holomorphic triangles, but with the additional factor of $n_{z_i}(\psi)$ for each triangle, and an additional factor of $V_{z_i}^{-1}$, on the entire map. Now color so that $V_{z_1} = V_{z_2}$ and $U_{w_1} = U_{w_2}$ (so that $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ and $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ are both well defined elements of homology), and apply the previous map to elements $\Theta^{\mathbf{w}}_{\alpha'\alpha} \otimes \mathbf{x}$. Noting that $\partial_{\alpha'\alpha}\Theta^{\mathbf{w}}_{\alpha'\alpha} = 0$, after coloring, we see that

$$F^{\mathbf{w}}_B = \left(\mathbf{x} \mapsto F_{\alpha'\alpha\beta}(\Theta^{\mathbf{z}}_{\alpha'\alpha}, \mathbf{x})\right) = \left(\mathbf{x} \mapsto F_{\alpha'\alpha\beta}(\Psi_{z_i}\Theta^{\mathbf{w}}_{\alpha'\alpha}, \mathbf{x})\right) \simeq \Psi_{z_i} F^{\mathbf{z}}_B + F^{\mathbf{z}}_B \Psi_{z_i}.$$

The equivalence $F^{\mathbf{z}}_B \simeq \Phi_{w_i} F^{\mathbf{w}}_B + F^{\mathbf{w}}_B \Phi_{w_i}$ follows by symmetry. □

Using the previous relation, we can now compute the commutator of $F_B^{\mathbf{z}}$ and $T^+_{w,z}$ when the pair (w,z) is between the end of B and a $\mathbf{z}$–basepoint. This was the case excluded by Proposition 4.7.1 (compare Lemma 4.3.18).

Lemma 4.7.17. *If (w,z) is an adjacent pair of basepoints between an end of the band B and a $\mathbf{z}$–basepoint, then*

$$F_B^{\mathbf{z}}T^+_{w,z} + T^+_{w,z}F_B^{\mathbf{z}} + S^+_{w,z}F_B^{\mathbf{w}} \simeq 0.$$

Proof. Let w' be one of the $\mathbf{w}$–basepoints adjacent to the ends of B, as in Figure 4.25, so that w' is not adjacent to z before adding the band B. We then compute:

$$\begin{aligned} F_B^{\mathbf{z}}T^+_{w,z} + T^+_{w,z}F_B^{\mathbf{z}} &\simeq (\Phi_{w'}F_B^{\mathbf{w}} + F_B^{\mathbf{w}}\Phi_{w'})T^+_{w,z} + T^+_{w,z}(\Phi_{w'}F_B^{\mathbf{w}} + F_B^{\mathbf{w}}\Phi_{w'}) \quad \text{(Proposition 4.7.16)} \\ &\simeq F_B^{\mathbf{w}}(\Phi_{w'}T^+_{w,z} + T^+_{w,z}\Phi_{w'}) + (\Phi_{w'}T^+_{w,z} + T^+_{w,z}\Phi_{w'})F_B^{\mathbf{w}} \quad \text{(Lemma 4.7.3).} \end{aligned}$$

By assumption, we know that w' is not adjacent to z before we add B, and hence

$$F_B^{\mathbf{z}}(\Phi_{w'}T^+_{w,z} + T^+_{w,z}\Phi_{w'})+ \simeq 0$$

by Lemma 4.3.20. Once we add the band B, however, we have that z and w' are adjacent, so that

$$\Phi_{w'}T^+_{w,z} + T^+_{w,z}\Phi_{w'} \simeq \Phi_w T^+_{w,z}.$$

Using the relations $\Phi_w \simeq S^+_{w,z}S^-_{w,z}$ we thus have that

$$F_B^{\mathbf{z}}T^+_{w,z} + T^+_{w,z}F_B^{\mathbf{z}} \simeq \Phi_w T^+_{w,z}F_B^{\mathbf{w}} \simeq S^+_{w,z}S^-_{w,z}T^+_{w,z}F_B^{\mathbf{w}} \simeq S^+_{w,z}F_B^{\mathbf{w}},$$

since also $S^-_{w,z}T^+_{w,z} \simeq \mathrm{id}$ by Lemma 4.3.13. □

Remark 4.7.18. Commutation of band maps and quasi-stabilization maps was considered in Propositions 4.7.1 and 4.7.3, but the configuration considered in Lemma 4.7.17 was excluded, since our holomorphic curve counts did not apply in that case. Note that using our interpretation of band maps and quasi-stabilization maps in terms of surfaces with divides,

one should not expect $T^+_{w,z}$ and $F^{\mathbf{z}}_B$ to commute in the situation considered in Lemma 4.7.17, since the induced dividing sets are different. Instead, one has the relation from Lemma 4.7.17, which is shown in Figure 4.25.

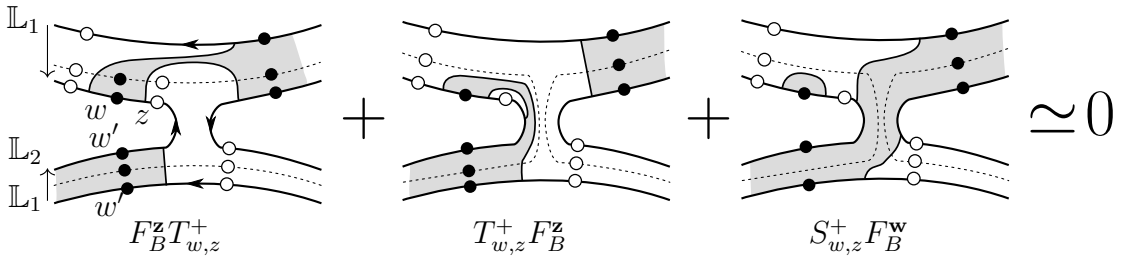

Figure 4.25: **An example of quasi-stabilization maps not commuting with band maps**The dividing sets involved in the relation $F^{\mathbf{z}}_B T^+_{w,z} + T^+_{w,z} F^{\mathbf{z}}_B + S^+_{w,z} F^{\mathbf{w}}_B \simeq 0$. These are shown on a surface which is formed by adding a band to $L_1 \times [0,1]$ (i.e. a saddle cobordism, viewed as a cobordism from the top and bottom to the middle). The dashed lines indicate level sets of a Morse function on the surface. This is viewed as a link cobordism from $\mathbb{L}$ to $\mathbb{L}(B)$.

4.8 Maps for framed 0–spheres along L

In this section, we define maps for performing surgery on a framed 0–sphere $\mathbb{S}_0$ which occurs along the link L. The maps are defined as a composition involving 1–handle maps away from L, band maps, and maps induced by diffeomorphisms.

The reader may notice that this construction may force us to consider links which are not nullhomologous (in the sense that each component may not be nullhomologous). We give a brief discussion, later in this section.

We begin with the definition:

Definition 4.8.1. Suppose L is an oriented link in Y. We say that a framed 0–sphere $\mathbb{S}_0$ in Y, with embedding $\iota : S^0 \times D^3 \to Y$, is a **framed 0–sphere along** L if the following holds:

1. Viewing D^3 as $\{(y,w,z) : y^2 + w^2 + z^2 \leq 1\}$, we have that $\iota^{-1}(L) = S^0 \times \{(y,0,0) :$

$-1 \leq y \leq 1\}$;

2. Each of $\iota(S^0 \times \{(1,0,0\})$ and $\iota(S^0 \times \{(-1,0,0)\})$ contains an incoming point of L, and an outgoing point of L, according to the orientation of L.

4.8.1 Model compound 1–handle/band cobordisms

We introduce a model cobordism for a link cobordism obtained by adding a 4–dimensional 1–handle to $Y \times [0,1]$ with a 2–dimensional 1–handle added to $L \times [0,1]$ inside of the 4–dimensional handle.

Suppose we are given a framed 0–sphere $\mathbb{S}_0$ along $\mathbb{L}$ (with embedding $\iota : S^0 \times D^3 \to Y$). We form the surgered 3–manifold

$$Y(\mathbb{S}_0) = (Y \setminus \operatorname{int} \iota(S^0 \times D^3)) \sqcup_\iota ([-1,1] \times S^2),$$

using the same construction as if $\mathbb{S}_0$ were a framed 0–sphere embedded away from L.

Inside of $Y(\mathbb{S}_0)$, we can form the link $L(\mathbb{S}_0)$, which is

$$L(\mathbb{S}_0) = (L \setminus \iota(\mathbb{S}_0)) \cup_\iota ([-1,1] \times \{\pm 1, 0, 0\}).$$

Condition (1) implies that the link $L(\mathbb{S}_0)$ is well defined, while condition (2) implies that $L(\mathbb{S}_0)$ can be given an orientation compatible with L.

Let $W(Y, \mathbb{S}_0)$ denote the 1–handle cobordism defined as

$$W(Y, \mathbb{S}_0) = (Y \times [0,1]) \cup (D^1 \times D^3) \cup (Y(\mathbb{S}) \times [1,2])$$

We define a surface, $\Sigma(\mathbb{S}_0)$, inside of $W(Y, \mathbb{S}_0)$, by

$$\Sigma(\mathbb{S}_0) = (L \times [0,1]) \cup B \cup (L(\mathbb{S}_0) \times [1,2])$$

where $B \subseteq [-1,1] \times D^3$ is the subset $B = [-1,1] \times \{(y,w,z) : -1 \leq y \leq 1, z = w = 0\}$.

Note that B is canonically identified with $[-1,1] \times [-1,1]$, and there is a distinguished point $(0,0)$ in B.

4.8.2 Construction of the compound 1–handle/band maps

Suppose that $\mathbb{S}_0$ is a framed 0–sphere along L, which doesn't intersect any of the basepoints in $\mathbf{w} \cup \mathbf{z}$. If $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ and $\overline{\mathfrak{s}}$ denotes the unique extension to the 4–manifold $W(Y, \mathbb{S}_0)$ and $\mathfrak{s}'$ denotes its restriction to $Y(\mathbb{S}_0)$, we will define a type–$\mathbf{z}$ compound 1–handle/band map

$$F^{\mathbf{z}}_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}} : \mathcal{C}FL^{\circ}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^{\circ}(Y(\mathbb{S}_0), \mathbb{L}(\mathbb{S}_0), \sigma', \mathfrak{P}, \mathfrak{s}')$$

as well as a type–$\mathbf{w}$ compound handle/band map

$$F^{\mathbf{w}}_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}} : \mathcal{C}FL^{\circ}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^{\circ}(Y(\mathbb{S}_0), \mathbb{L}(\mathbb{S}_0), \sigma', \mathfrak{P}, \mathfrak{s}').$$

Our definition of the compound 1–handle/band map is motivated by the following observation. If we pick an element $v \in S^2 \setminus \{(\pm 1, 0, 0)\}$, we can push the surface $\Sigma(\mathbb{S}_0) \subseteq W(Y, \mathbb{S}_0)$ in the direction v towards $[-1, 1] \times S^2 \subseteq [-1, 1] \times D^3$, while keeping the band fixed along $\{(\pm 1, 0, 0)\} \times S^2$, where $[-1, 1] \times D^3$ is the 1–handle of the cobordism, so that the band B becomes a subset of the 3–manifold $Y(\mathbb{S}_0) \times \{1\}$. Thinking of the link cobordism in terms of a sequence of moves of the 4–manifold and surface, this corresponds to first pushing L off of the framed spheres $\mathbb{S}_0$ in the direction of v, then attaching a 4–dimensional 1–handle, and then attaching a band to L (or rather L pushed away from the framed sphere) inside of $Y(\mathbb{S}_0)$, which now goes across the subset $[-1, 1] \times S^2 \subseteq Y(\mathbb{S}_0)$. Such a presentation of the link cobordism is determined by the choice of $v \in S^2 \setminus \{\pm 1, 0, 0\}$. Similarly, a path between any two choices of such v yields an isotopy of the resulting presentation of the link cobordism through such presentations. As $S^2 \setminus \{(\pm 1, 0, 0)\}$ is connected, we conclude that any two presentations are isotopic in this sense. With this observation in mind, we now proceed to formally define the maps.

For the type $\mathbf{o}$ maps (where $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$), we require that the two $\mathbf{o}$–basepoints adjacent to the points in $\mathbb{S}_0$ have the same color.

Given a framed 0–sphere $\mathbb{S}_0$ along L, which doesn't intersect any of the basepoints, we first pick a vector $\theta \in S^1 = \{(0, w, z) : w^2 + z^2 = 1\}$. Such a vector determines a half disk H_i in each $\{p_i\} \times D^3$ (the set of points with spherical coordinates which have angle θ), and

an arc a on S^2 (the boundary of the half disk). There is naturally a square $S = [-1,1] \times a$ determined in $[-1,1] \times S^2$.

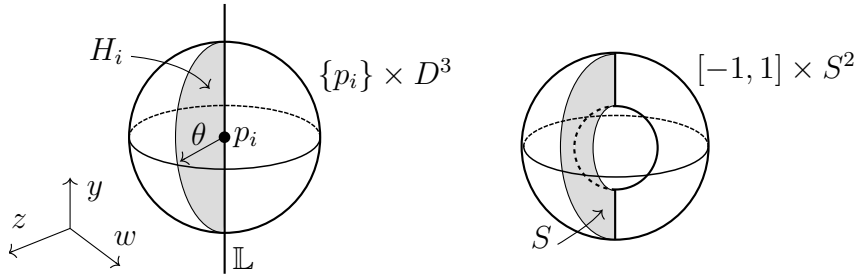

Figure 4.26: **Pushing the band out of a compound 1–handle/band**On the left is the vector $\theta \in S^1$, as well as a half disk H_i, the link $\mathbb{L}$ and the balls $\{p_i\} \times D^3$. On the right is $[-1,1] \times S^2$, and the square S with boundary on $\{-1,1\} \times S^2 \cup [-1,1] \times \{(\pm 1,0,0)\}$.

Extend the half disks to slightly larger half disks $\overline{H}_1$ and $\overline{H}_2$ which extend just past the balls $\{p_i\} \times D^3$. The half disks $\overline{H}_1$ and $\overline{H}_2$ determine a diffeomorphism, ψ^θ of Y which pushes $\mathbb{L}$ across the half disks $\overline{H}_i$ to a link $\mathbb{L}_\theta$ just outside of the balls $\{p_i\} \times D^3$ which doesn't intersect $\mathbb{S}_0$, and is the identity outside a neighborhood of the $\{p_i\} \times D^3$. The extensions $\overline{H}_i$ of the disks H_i depend on a contractible set of choices.

Inside of $Y(\mathbb{S}_0)$, there is a natural band

$$B_\theta = \overline{H}_1 \cup S \cup \overline{H}_2$$

from $\mathbb{L}_\theta$ to $\mathbb{L}(\mathbb{S}_0)$.

For $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$. We now define the type–$\mathbf{o}$ compound 1–handle/band to be the composition

$$F^{\mathbf{o}}_{Y,\mathbb{S}_0,\mathbb{L},\bar{\mathfrak{s}}} = F^{\mathbf{o}}_{B_\theta} \circ F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}} \circ \psi^\theta_*.$$

Lemma 4.8.2. *The maps $F^{\mathbf{o}}_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}}}$ (for $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$) are independent of the choice of $\theta \in S^1$, as well as the extensions $\overline{H}_i$.*

Proof. Fix a choice of $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, for the argument. The extended half disks $\overline{H}_i$ depend on a contractible set of choices (a vector field pointing out of the boundary of $\mathbb{S}_0$, and a

Riemannian metric), and hence it is easy to see that the maps don't depend on the extensions $\overline{H}_i$.

We now consider dependence on $\theta \in S^1$. Let θ and θ' be two elements in S^1, and let ψ^θ and $\psi^{\theta'}$ be the two associated diffeomorphisms, and let B_θ and $B_{\theta'}$ be the two bands. We pick diffeomorphism Ψ of Y which rotates $\mathbb{L}_\theta$ to $\mathbb{L}_{\theta'}$, which is supported in a neighborhood of $\mathbb{S}_0$, but is constant on $\mathbb{S}_0$. Note that since Ψ is constant on $\mathbb{S}_0$, we can consider the diffeomorphism $\Psi^{\mathbb{S}_0}$ on $Y(\mathbb{S}_0)$.

Note that $\Psi^{\mathbb{S}_0}$ sends the band B_θ to a band which is isotopic relative its boundary to $B_{\theta'}$. As a consequence, we conclude that

$$F^{\mathbf{o}}_{B_{\theta'}} \circ \Psi^{\mathbb{S}_0}_* \simeq \Psi^{\mathbb{S}_0}_* \circ F^{\mathbf{z}}_{B_\theta}\,, \tag{4.6}$$

by diffeomorphism invariance of the band maps.

Furthermore, since Ψ fixes $\mathbb{S}_0$, we have

$$\Psi^{\mathbb{S}_0}_* \circ F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}} \simeq F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}} \circ \Psi_*. \tag{4.7}$$

Note too that

$$\Psi \circ \psi^\theta \sim \psi^{\theta'}, \tag{4.8}$$

where $\sim$ denotes isotopy relative $\mathbb{L}$.

Finally, we note that $\Psi^{\mathbb{S}_0}$ is isotopic to the identity relative $\mathbb{L}(\mathbb{S}_0)$, so

$$\Psi^{\mathbb{S}_0}_* \simeq 1 \tag{4.9}$$

on $\mathcal{C}FL^\circ(Y(\mathbb{S}_0), \mathbb{L}(\mathbb{S}_0), \mathfrak{s}')$. Hence

$$\begin{aligned}
F^{\mathbf{o}}_{B_\theta} F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}} \psi^\theta_* &\simeq (\Psi^{\mathbb{S}_0}_*)^{-1} F^{\mathbf{o}}_{B_{\theta'}} \Psi^{\mathbb{S}_0}_* F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}} \psi^\theta_* && \text{(Equation (4.6))} \\
&\simeq (\Psi^{\mathbb{S}_0}_*)^{-1} F^{\mathbf{o}}_{B_{\theta'}} F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}} \Psi_* \psi^\theta_* && \text{(Equation (4.7))} \\
&\simeq (\Psi^{\mathbb{S}_0}_*)^{-1} F^{\mathbf{o}}_{B_{\theta'}} F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}} \psi^{\theta'}_* && \text{(Equation (4.8))} \\
&\simeq F^{\mathbf{o}}_{B_{\theta'}} F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}} \psi^{\theta'}_*\,, && \text{(Equation (4.9))}
\end{aligned}$$

completing the proof. □

4.8.3 Alexander and Maslov gradings

One can use a modification of the techniques from [OS06] to define absolute Maslov and Alexander gradings on link Floer homology, using surgery presentations of the link complement. This is explored in more detail in Chapter 5, so we just summarize the overall strategy here.

Using the absolute grading from $CF^-(Y,\mathbf{w},\mathfrak{s})$, one can define a grading $\mathrm{gr}_{\mathbf{w}}$ on $\mathcal{C}FL^-(Y,\mathbb{L},\mathfrak{s})$ as long as $c_1(\mathfrak{s})$ is torsion. The $U_{\mathbf{w}}$ variables are -2 graded, while the $V_{\mathbf{z}}$ variables are 0 graded.

If $\mathfrak{s}-PD[L]$ is torsion, then one can use the grading on $CF^-(Y,\mathbf{z},\mathfrak{s}-PD[L])$ to get a second Maslov grading $\mathrm{gr}_{\mathbf{z}}$ on $\mathcal{C}FL^-(Y,\mathbb{L},\mathfrak{s})$. For this grading, one has that the $U_{\mathbf{w}}$ variables are 0 graded, while the $V_{\mathbf{z}}$ variables are -2 graded.

As in [OS08], one can also define an Alexander multi-grading on $\mathcal{C}FL^-(Y,\mathbb{L},\mathfrak{s})$. To an ℓ–component link in S^3, there will be an ℓ-component Alexander grading on $\mathcal{C}FL^-(Y,\mathbb{L},\mathfrak{s})$. For links in other 3–manifolds, the situation is somewhat more complicated. If $\pi : C(L)\to J$ is a map from the set of components of L to an index set J, then we say that L is J–**null-homologous** if $\pi^{-1}(j)=0\in H_1(Y;\mathbb{Z})$ for each $j\in J$. To a multi-based link $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ which is J–null-homologous, the chain complexes $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ have a natural Alexander multi-grading over $\mathbb{Q}^J$.

The relative grading between two intersection points $\mathbf{x}$ and $\mathbf{y}$ is defined by

$$A(\mathbf{x},\mathbf{y})_j=\sum_{z\in\mathbf{z}\cap\pi^{-1}(j)}n_z(\phi)-\sum_{w\in\mathbf{w}\cap\pi^{-1}(j)}n_w(\phi),$$

for a homology disk $\phi\in\pi_2(\mathbf{x},\mathbf{y})$.

For a link L with equipped with a map $\pi : C(L)\to J$, it is important to note that one does not need each component of L to be null-homologous for this to be well defined. For the above formula to be well defined, one must show that it is independent of splicing in periodic domains to ϕ (i.e. two chains on Σ with boundary equal to a sum of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves). If P is a periodic domain, there is a natural homology class $H(P)\in H_2(Y;\mathbb{Z})$, obtained by capping

the the 2–chain determined by P with the compressing disks along the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves. The above expression for $A(\mathbf{x},\mathbf{y})_j$ changes by exactly $\#H(P)\cap\pi^{-1}(j)$ when we splice in P, and hence is well defined exactly when L is (rationally) J–null-homologous.

We remark that attaching a band may produce link components which are not null-homologous. Nonetheless, if an oriented link is given a grading assignment $\pi : C(L) \to J$, and B is an oriented band attached to L, then L is J–null-homologous iff $L(B)$ is. Thus by working with link cobordisms (W,Σ) where the components of Σ are themselves given an assignment $\pi : C(\Sigma) \to J$, the link cobordism maps will naturally be graded assuming some favorable assumptions on the colorings of the link and surface. This is explored in more detail in Chapter 5.

4.9 Parametrized Kirby decompositions of link cobordisms

In this section, we define notions of parametrized cobordisms and elementary parametrized cobordism. We then define a way of dividing a link cobordism into elementary parametrized cobordisms, which we call a parametrized Kirby decomposition, borrowing the terminology of [Juh16].

4.9.1 Parametrized cobordisms and elementary parametrized cobordisms

We now define a parametrized link cobordism:

Definition 4.9.1. If $(W,\Sigma) : (Y,L) \to (Y',L')$ is an undecorated link cobordism, a **parametrization** of (W,Σ), is a diffeomorphism of link cobordisms

$$\Phi : W(Y,L,\mathbb{S}) \to (W,\Sigma),$$

where $\mathbb{S}$ is a pairwise disjoint collection of framed spheres in Y (along L or away from L), and $W(Y,L,\mathbb{S})$ is a model link cobordism obtained by attaching 4-dimensional handles, and 2-dimensional handles, as described in Sections 4.4.1 and 4.8.1.

The following fact will be useful:

Remark 4.9.2. If $\mathbb{S}$ and $\mathbb{S}'$ are disjoint collections of framed spheres in Y, the link cobordisms $W(Y, L, \mathbb{S} \cup \mathbb{S}')$, $W(Y(\mathbb{S}), L, \mathbb{S}') \circ W(Y, L, \mathbb{S})$, and $W(Y(\mathbb{S}'), L, \mathbb{S}) \circ W(Y, L, \mathbb{S}')$ are all diffeomorphic. The diffeomorphism is canonically specified up to isotopy fixing $Y \times \{0\}$, and preserving Σ. To see this, it is sufficient to specify the diffeomorphism from the first cobordism to the second. Notice that we can think of the second as being equal to the first with an extra copy of $Y(\mathbb{S}) \times [0, 1]$ inserted. A choice of bicollar neighborhood of $Y(\mathbb{S}) \times \{0\}$ allows us to specify such a diffeomorphism, up to isotopy. Since any two bicollar neighborhoods are isotopic, a diffeomorphism between the two cobordisms is specified uniquely up to isotopy fixing $Y \times \{0\}$ and preserving Σ.

Now suppose $(W, \Sigma) : (Y_1, L_1) \to (Y_2, L_2)$ and $(W', \Sigma') : (Y_2, L_2) \to (Y_3, L_3)$ are two link cobordisms with parametrizations

$$\Phi : W(Y_1, L_1, \mathbb{S}_1) \to (W, \Sigma), \qquad \text{and} \qquad \Phi' : W(Y_2, L_2, \mathbb{S}_2) \to (W', \Sigma').$$

For notational simplicity, assume that no components of $\mathbb{S}_1$ or $\mathbb{S}_2$ are framed 0-spheres along L_1 or L_2. Let $\mathbb{S}_2' \subseteq Y_1(\mathbb{S}_1)$ denote $\Phi|^{-1}_{Y_1(\mathbb{S}_1)}(\mathbb{S}_2)$. We can form the **concatenation** of the two parametrizations

$$\Phi' * \Phi : W(Y_1(\mathbb{S}_1), L_1, \mathbb{S}_2') \circ W(Y_1, L_1, \mathbb{S}_1) \to (W', \Sigma') \circ (W, \Sigma),$$

which is defined as Φ on $W(Y_1, L_1, \mathbb{S}_1)$ and $\Phi' \circ \Phi|_{Y_1(\mathbb{S}_1)}$ on $W(Y_1(\mathbb{S}_1), L_1, \mathbb{S}_2')$. If $\mathbb{S}_2' \subseteq Y_1 \setminus \mathbb{S}_1$, then in light of Remark 4.9.2, the concatenated diffeomorphisms induce a parametrization

$$\Phi' * \Phi : W(Y_1, L_1, \mathbb{S}_1 \cup \mathbb{S}_2') \to (W', \Sigma') \circ (W, \Sigma)$$

which is well defined up to isotopies (of link cobordisms) fixing $Y_1 \times \{0\}$. The concatenation of the two parametrizations can also be defined in the case that $\mathbb{S}_1$ or $\mathbb{S}_2$ is a framed link along L_1 or L_2, with only the obvious notational changes to the domain of $\Phi' * \Phi$.

Analogous to [Juh15], we need a notion of elementary parametrized cobordisms, corresponding to link cobordisms which have a Morse function with at most one critical value of $f|_{W \setminus \Sigma}$, $f|_\Sigma$ or $f|_\mathcal{A}$:

Definition 4.9.3. A tuple $\mathcal{W} = (W, F, \Phi, \mathbb{S})$ is called an **elementary, parametrized link cobordism** if $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a decorated link cobordism with $\partial W = -Y_1 \sqcup Y_2$, $F = (\Sigma, \mathcal{A})$ is a surface with divides, and $\mathbb{S}$ is a (possibly empty) collection of framed spheres (0-,1-, or 2-spheres away from L, or 0-spheres along L) in Y_1 which are pairwise disjoint and such that exactly one of the following holds:

1. $\mathbb{S} = \varnothing$ and $\Phi : W(Y_1, L_1, \varnothing) \to (W, \Sigma)$ is a parametrization which is the identity on $Y_1 \times \{0\}$ is defined up to isotopies of link cobordisms which fix $Y_1 \times \{0\}$. Furthermore, each arc of the divides $\mathcal{A}$ goes from the incoming boundary to the outgoing boundary.

2. $\mathbb{S} \neq \varnothing$ and $\mathbb{S}$ is a framed 0-sphere, 1-dimensional link or 2-sphere away from L_1 and Φ is a parametrization
$$\Phi : W(Y_1, L_1, \mathbb{S}) \to (W, \Sigma)$$
which is the identity on $Y_1 \times \{0\}$, and which is defined up to isotopies of link cobordisms which fix $Y_1 \times \{0\}$. Furthermore, each arc of the divides $\mathcal{A}$ goes from the incoming boundary to the outgoing boundary.

3. $\mathbb{S} = \varnothing$ and Φ is a parametrization
$$\Phi : W(Y_1, L_1, \varnothing) \to (W, \Sigma),$$
which is the identity on $Y_1 \times \{0\}$, and is defined up to isotopies fixing $Y_1 \times \{0\}$. Furthermore, all divides of $\mathcal{A}$, except for exactly one arc A, go from the incoming boundary to the outgoing boundary. There are four subtypes:

   I^+ The arc A has two ends on the outgoing boundary, and above the arc is a type-**w** region.

   I^- The arc A has two ends on the incoming boundary, and below the arc is a type-**w** region.

   II^+ The arc A has two ends on the outgoing boundary, and above the arc is a type-**z** region.

II^- The arc A has two ends on the incoming boundary, and below the arc is a type-$\mathbf{z}$ region.

These dividing sets can be found in Figure 4.28.

4. $\mathbb{S}$ is a framed 0-sphere, embedded along L_1, and Φ is a parametrization

$$\Phi : W(Y_1, L_1, \mathbb{S}) \to (W, \Sigma),$$

which is the identity on $Y_1 \times \{0\}$, and such that the divides in $\Phi^{-1}(\mathcal{A}) \subseteq \Sigma(\mathbb{S})$ do not intersect $B \subseteq \Sigma(\mathbb{S})$. The parametrization is well defined up to isotopies which fix $Y_1 \times \{0\}$ and the point $(0,0) \in B$ Furthermore, all of the arcs in $\mathcal{A}$ are arcs from the incoming boundary to the outgoing boundary.

The type ($\mathbf{w}$ or $\mathbf{z}$) of each region of $\Sigma(\mathbb{S}) \setminus \Phi^{-1}(\mathcal{A})$ is well defined. Hence there are two subtypes, (type $\mathbf{w}$ and type $\mathbf{z}$) corresponding to the type of the region with the band under the parametrization.

Remark 4.9.4. In the definition of an elementary parametrized link cobordism, the diffeomorphism Φ in the definition is one of undecorated link cobordisms, and doesn't necessarily respect the divides beyond what is stated for type (4) link cobordisms. Furthermore, the parametrization Φ is only defined up to isotopies which fix $Y_1 \times \{0\}$, and in the case of (4), isotopies fixing $(0,0)$.

We make the following definition, borrowing the terminology from [Juh16] (see also [Juh15, Definition 2.8]):

Definition 4.9.5. For a link cobordism $(W, F) : (Y, \mathbb{L}) \to (Y', \mathbb{L}')$, we define a **parametrized Kirby decomposition**, $\mathcal{K}$, of the decorated link cobordism (W, F) to be the following collection of data:

1. A decomposition of (W, F) as

$$(W, F) = (W_0, F_0) \cup_{Y_1} \cdots \cup_{Y_n} (W_n, F_n)$$

where (W_i, F_i) is a link cobordism with $\partial W_i = -Y_i \sqcup Y_{i+1}$ and $\partial F_i = -L_i \sqcup L_{i+1}$, for 3-manifolds with links (Y_i, L_i) and surface with divides F_i (all embedded in W), such that $Y_0 = Y$ and $Y_{n+1} = Y'$ and $\mathcal{A}$ intersects each component of each L_i nontrivially and transversally;

2. A choice of basepoints $\mathbf{w}_i$ and $\mathbf{z}_i$ on each L_i, with one basepoint per component of $L_i \setminus \mathcal{A}$, such that the type of each basepoint is compatible with the type induced by the divides $\mathcal{A}$.

3. An embedded sphere or link, $\mathbb{S}_i \subseteq Y_i$ and a parameterizing diffeomorphism Φ_i, making each $(W_i, F_i, \Phi_i, \mathbb{S}_i)$ an elementary, parametrized link cobordism.

Additionally, we assume that all terms of the decomposition of type (2) occur after all of the terms of type (3) or (4).

Finally we assume that the terms of the decomposition of type (2) appear with nondecreasing index (i.e. all 0-spheres occur before 1-dimensional links, which themselves occur before 2-spheres), and there is at most one term of the decomposition which has a 1-dimensional link.

We make the following definition (compare [Juh15, Section 1.3]):

Definition 4.9.6. Suppose $\mathbb{S} \subseteq Y$ is a framed k sphere, and $\mathbb{S}' \subseteq Y(\mathbb{S})$ is a framed $(k+1)$ sphere which intersects the belt sphere of $\mathbb{S}$, transversely at a single point. We say a diffeomorphism

$$\Psi : W(Y(\mathbb{S}), L_i, \mathbb{S}') \circ W(Y, L, \mathbb{S}) \to W(Y, L, \varnothing),$$

is **formed by a handle cancellation** if it is constructed by the following procedure: take two Morse functions and gradient like vector fields on $W(Y(\mathbb{S}), L, \mathbb{S}')$ and $W(Y, L, \mathbb{S})$ which each have exactly one critical point, and induce the parametrizations the cobordisms which are the identity on the boundary. Such Morse functions exist by Lemma 4.10.25. The stable manifold of the higher critical point is transverse to the unstable manifold of the lower critical point and there is a single flowline between them. Using the argument from [Juh15, Section

1.3], there is a homotopically unique path of functions f_t and gradient like vector fields v_t, which are constant outside a neighborhood of the flowline between the two critical points, which realizes a birth-death singularity. The pair (f_1, v_1) on $W(Y(\mathbb{S}), L, \mathbb{S}') \circ W(Y, L, \mathbb{S})$ determines a diffeomorphism $\Psi : W(Y, L, \varnothing) \to W(Y(\mathbb{S}), L, \mathbb{S}') \circ W(Y, L, \mathbb{S})$, which is the identity on the incoming boundary, $Y \times \{0\}$, and is well defined up to isotopy relative $Y \times \{0\}$, since (f_1, v_1) is determined up to homotopy (through paths of Morse functions and gradient like vector fields, with no intermediate singularities).

4.10 Morse theory and moves between parametrized Kirby decompositions

In this section we describe a set of moves that one can be used to connect two parametrized Kirby decompositions. Our strategy is motivated by the observation that a Morse function f_0 on a submanifold can be used to define a Morse function on a regular neighborhood by the formula $f = f_0 + r^2$ (where r is the distance function induced by a Riemannian metric on the normal bundle). The function f has an index k critical point for each index k critical point of f_0. In this manner, we can construct a Morse function on all of W, by first constructing one on a regular neighborhood of Σ in this manner, and then extending it to all of W. Performing an appropriately adapted Cerf theory will then allow us to move between two presentations given by performing this procedure.

In Subsection 4.10.1, we list a set of moves to move between two parametrized Kirby decompositions. In Subsection 4.10.2, we prove several technical results about spaces of coordinate functions, which we will need in order to do the Morse theory. In Subsection 4.10.3, we show how to nicely connect any two Morse functions which are equal on Σ. In Subsection 4.10.4, we show how to change the Morse function on Σ. In Subsection 4.10.5, we show how a (suitably nice) Morse function gives rise to a parametrized Kirby decomposition. Finally in Subsection 4.10.6, we put everything together, and show that the set of moves from Subsection 4.10.1 are sufficient to connect any two parametrized Kirby decompositions.

4.10.1 Weak equivalences of parametrized Kirby decompositions

In this section we describe a set of moves to go between two Kirby decompositions.

We first describe a notion of isotopy between two parametrized Kirby decompositions, which will be one of the moves. Suppose that $\mathcal{K}$ is a parametrized Kirby decomposition of (W, F) and that $\phi : W \to W$ is a diffeomorphism which is the identity on ∂W. Then ϕ induces a decomposition $\phi_*\mathcal{K}$ of (W, F) into parametrized cobordisms, though the decomposition may not be a parametrized Kirby decomposition in the sense of Definition 4.9.5, as we now describe. If

$$(W, F) = (W_0, \Sigma_0) \cup_{Y_1} \cdots \cup_{Y_n} (W_n, \Sigma_n)$$

is the decomposition of (W, Σ) associated to $\mathcal{K}$, then the decomposition associated to $\phi_*\mathcal{K}$ is

$$(\phi(W_0), \phi(\Sigma_0)) \cup_{\phi(Y_1)} \cdots \cup_{\phi(Y_n)} (\phi(W_n), \phi(\Sigma_n)). \tag{4.10}$$

The parametrizations, Φ_i, of $\mathcal{K}$ can also be pushed forward along ϕ to give diffeomorphisms

$$\phi_*(\Phi_i) : W(\phi(Y_i), \phi(L_i), \phi(\mathbb{S}_i)) \to (\phi(W_i), \phi(\Sigma_i)).$$

The parametrized pieces may fail to be elementary parametrized cobordisms in the sense of Definition 4.9.3, since the divides $\mathcal{A} \cap \phi_*\Sigma_i$ may not make each $(\phi_*W_i, \phi_*\Sigma_i, \phi_*\Sigma_i \cap \mathcal{A})$ into an elementary parametrized cobordism. With this in mind, we make the following definition:

Definition 4.10.1. Suppose that (W, F) is a link cobordism with $F = (\Sigma, \mathcal{A})$. If $\mathcal{K}$ and $\mathcal{K}'$ are two parametrized Kirby decompositions of (W, F), we say that $\mathcal{K}$ and $\mathcal{K}'$ are **strongly isotopic** if there is a 1-parameter family of diffeomorphisms $\phi_t : (W, \Sigma) \to (W, \Sigma)$ such that following holds:

1. $\phi_t|_{\partial W} = \mathrm{id}$ for all t;
2. $(\phi_t)_*\mathcal{K}$ is a valid parametrized Kirby decomposition for all t;
3. $(\phi_1)_*\mathcal{K} = \mathcal{K}'$.

In the previous definition, note that if ϕ_t is such an isotopy of W and Φ_i is a parameterizing diffeomorphism from a term of $\mathcal{K}$ corresponding to a framed 0-sphere along L_i, then $(\phi_t)_*\Phi_i$ maps the point $(0,0)$ in the model band to a path on Σ. For ϕ_t to induce a strong isotopy, this path cannot intersect an arc of $\mathcal{A}$, since at the point of time when a point crosses the band, the map $(\phi_t)_*\Phi_i$ would fail to be a valid parametrization for an elementary parametrized cobordism, as it would map the center of the band onto an arc of $\mathcal{A}$.

We now define a more general definition of equivalence between two parametrized Kirby decompositions.

Definition 4.10.2. We say two parametrized Kirby decompositions are **weakly equivalent** if they are related by the following set of moves (or their inverses):

1. (strong isotopies): See Definition 4.10.1.

2. (isotopy of an attaching sphere, $\mathbb{S}_i$, away from L_i): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathbb{S}_i^t$ is a 1-parameter family of framed k-spheres in $Y_i \setminus L_i$ starting at $\mathbb{S}_i$ and ending at a framed sphere $\mathbb{S}_i'$. Let d_i^t denote an extension of this isotopy to Y, we define

$$D : W(Y_i, L_i, \mathbb{S}_i) \to W(Y_i, L_i, \mathbb{S}_i)$$

by

$$D(x,t) = (d_i^t(x), t)$$

for $(x,t) \in Y \times [0,1]$. We define D to map one handle to the other, under their identifications as $D^{k+1} \times D^{3-k}$. There is an induced diffeomorphism $d_i^{\mathbb{S}_i}$ from $Y_i(\mathbb{S}_i)$ to $Y_i(\mathbb{S}_i')$, which D restricts to. The elementary parametrized cobordism $\mathcal{W}_i$ is replaced with an elementary cobordism $\mathcal{W}_i' = (W_i, F_i, \Phi_i', \mathbb{S}_i')$ where $\Phi_i' \circ D = \Phi_i$.

3. (isotopy of a framed 0-sphere sphere $\mathbb{S}_i$ along L_i): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathbb{S}_i^t$ is a 1-parameter family of framed 0-spheres along L_i from $\mathbb{S}_i$ to $\mathbb{S}_i'$. The diffeomorphism sending $\mathbb{S}_i$ to $\mathbb{S}_i^t$ extends to an isotopy d_i^t of Y_i which preserves L_i, and sends $\mathbb{S}_i$ to $\mathbb{S}_i^t$. The isotopy d_i^t defines a diffeomorphism

$$D : W(Y_i, L_i, \mathbb{S}_i) \to W(Y_i, L_i, \mathbb{S}_i')$$

as follows. It is defined by

$$D(x,t) = (d_i^t(x), t),$$

for $(x,t) \in Y \times [0,1]$,

$$D(x,t) = (d_i^1(x), t),$$

for $(x,t) \in Y(\mathbb{S}) \times [1,2]$, and D is the identity on the 1-handles of the cobordisms, under the identification of both as $D^1 \times D^3$. Note that since d_i^t preserves L_i, the map D sends $\Sigma(\mathbb{S})$ to $\Sigma(\mathbb{S}')$, diffeomorphically, and maps the bands to each other as the identity, under the identification of both as $[-1,1] \times [-1,1]$.

The parametrized cobordism $(W_i, F_i, \Phi_i, \mathbb{S}_i)$ is replaced by $\mathcal{W}_i' = (W_i, F_i, \Phi_i', \mathbb{S}_i')$, where $\Phi_i' \circ D = \Phi_i$.

4. (adding a trivial cylinder): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are adjacent terms of $\mathcal{K}$, with at least one of $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ being of type (1), then the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ is replaced with a single $\mathcal{W}_i' = (W_{i+1} \cup W_i, F_{i+1} \cup F_i, \Phi_i', \mathbb{S}_i')$. Under the canonical identification of

$$W(Y_{i+1}, L_{i+1}, \mathbb{S}_{i+1}) \cup_{\Phi_i} W(Y_i, L_i, \mathbb{S}_i) \cong W(Y_i, L_i, \mathbb{S}_i \cup \Phi_i^{-1}(\mathbb{S}_{i+1})),$$

the maps Φ_i' and the concatenation $\Phi_{i+1} * \Phi_i$ are isotopic through diffeomorphisms fixing $Y_i \times \{0\}$ (or $(Y_i \times \{0\}) \cup \{(0,0)\}$).

5. (changing the orientation of a framed sphere): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ is an elementary parametrized cobordism with $\mathbb{S}_i \neq \varnothing$, a framed 0-sphere, 2-sphere, or 1-dimensional link, away from L_i. We view $\mathbb{S}_i$ as a map from a $S^k \times D^{3-k}$ to Y_i, and give $S^k \times D^{3-k} \subseteq D^{k+1} \times D^{3-k} \subseteq \mathbb{R}^4$ the coordinates (x,y,w,z) from $\mathbb{R}^4$. Then we can replace $\mathbb{S}_i$ with $\mathbb{S}_i'$ formed by precomposing with the map

$$\sigma(x,y,w,z) = (-x,y,w,-z).$$

If $\mathbb{S}_i$ is a framed 0-sphere along L_i, we can replace $\mathbb{S}_i$ with the map $\mathbb{S}_i'$ induced by precomposing $\mathbb{S}_i$ with the map

$$\sigma(x,y,w,z) = (-x,-y,w,z).$$

The parameterizing diffeomorphism Φ is also changed by precomposition with σ on $D^{k+1} \times D^{3-k}$.

6. (An index 1/2 birth-death singularity away from Σ): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are adjacent terms of the decomposition, both of type (2), such that $\mathcal{W}_i$ has framed 0-sphere $\mathbb{S}_i$ away from L_i and parametrization

$$\Phi_i : W(Y_i, L_i, \mathbb{S}_i) \to (W_i, \Sigma_i),$$

and $\mathcal{W}_{i+1}$ has framed 1-dimensional link $\mathbb{S}_{i+1}$, and parametrization

$$\Phi_{i+1} : W(Y_{i+1}, L_{i+1}, \mathbb{S}_{i+1}) \to (W_{i+1}, \Sigma_{i+1}).$$

Suppose the belt sphere of $\mathbb{S}_i$ in Y_{i+1} (under the parametrization) intersects a single component, $\mathbb{K}$, of $\mathbb{S}_{i+1}$ exactly once.

The move is to replace $\mathcal{W}_{i+1}$ and $\mathcal{W}_i$ with a pair $\mathcal{W}'_{i+1}$ and $\mathcal{W}'_i$, such that the following holds: There is decomposition of $\mathcal{W}_{i+1}$ as $\mathcal{W}'_{i+1} \circ \mathcal{W}''_{i+1}$ as parametrized cobordisms, where

$$\mathcal{W}'_{i+1} = (W'_{i+1}, F'_{i+1}, \Phi'_{i+1}, \mathbb{S}'_{i+1}), \qquad \text{and} \qquad \mathcal{W}''_{i+1} = (W''_{i+1}, F''_{i+1}, \Phi''_{i+1}, \mathbb{K})$$

and $(\Phi''_{i+1})^{-1}(\mathbb{S}'_{i+1}) = \mathbb{S}_{i+1} \setminus \mathbb{K}$. The pair $\mathcal{W}''_{i+1} \circ \mathcal{W}_i$ is replaced by an elementary parametrized cobordism $\mathcal{W}'_i = (W'_i, F'_i, \Phi'_i, \varnothing)$ where

$$\Phi'_i \circ \Psi \simeq \Phi''_{i+1} * \Phi_i,$$

where

$$\Psi : W(Y_i(\mathbb{S}_i), L_i, \Phi_i^{-1}(\mathbb{K})) \circ W(Y_i, L_i, \mathbb{S}_i) \to W(Y_i, L_i, \varnothing)$$

is a diffeomorphism obtained by handle cancellation, as in Definition 4.9.6.

7. (An index 2/3 birth-death singularity away from Σ): This move is defined analogously to Move (6), but instead we start with a pair of adjacent terms $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ where $\mathbb{S}_i$ is a framed 1-dimensional link, and $\mathbb{S}_{i+1}$

is a framed 2-sphere, which cancels a component of $\mathbb{S}_i$. The pair is replaced with two elementary parametrized cobordisms, $\mathcal{W}'_{i+1}$ and $\mathcal{W}'_i$, with $\mathcal{W}'_{i+1}$ being a trivial cobordism and $\mathcal{W}'_i$ corresponding to surgery on the remaining components of $\mathbb{S}_i$.

8. (A critical value switch between framed 0- or 2-spheres, away from Σ): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are both of type (2), for 0- or 2-spheres $\mathbb{S}_i \subseteq Y_i$ and $\mathbb{S}_{i+1} \subseteq Y_{i+1}$ such that under the images of the framings of $\mathbb{S}_i$ and $\Phi_i^{-1}(\mathbb{S}_{i+1})$ in Y_i are disjoint, then we replace the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ with a pair

$$\mathcal{W}'_i = (W'_i, F'_i, \Phi'_i, \mathbb{S}'_i) \qquad \text{and} \qquad \mathcal{W}'_{i+1} = (W'_{i+1}, F'_{i+1}, \Phi'_{i+1}, \mathbb{S}'_{i+1})$$

both of type (2), such that $W_i \cup W_{i+1} = W'_i \cup W'_{i+1}$ and $F_i \cup F_{i+1} = F'_i \cup F'_{i+1}$, and also

$$\Phi_i(\mathbb{S}'_i) = \mathbb{S}_{i+1} \qquad \text{and} \qquad \Phi'_i(\mathbb{S}_i) = \mathbb{S}'_{i+1}.$$

Furthermore, under the canonical identifications

$$W(Y_i(\mathbb{S}_i), L_i, \mathbb{S}'_i) \circ W(Y_i, L_i, \mathbb{S}_i) \cong W(Y_i(\mathbb{S}'_i), L_i, \mathbb{S}_i) \circ W(Y_i, L_i, \mathbb{S}'_i),$$

we have that $\Phi_{i+1} * \Phi_i$ and $\Phi'_{i+1} * \Phi'_i$ are isotopic through diffeomorphisms which fix $Y_i \times \{0\}$.

9. (A handleslide amongst the framed 1-dimensional spheres): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ is of type (2), with $\mathbb{S}_i$ a framed 1-dimensional link, then handlesliding a component of $\mathbb{S}_i$ across another component yields a framed link $\mathbb{S}'_i$, and a diffeomorphism $D : W(Y_i, L_i, \mathbb{S}'_i) \to W(Y_i, L_i, \mathbb{S}_i)$ fixing $Y_1 \times \{0\}$, well defined up to isotopy[5]. The parametrized cobordism $\mathcal{W}_i$ is replaced by

$$\mathcal{W}'_i = (W_i, F_i, \Phi_i \circ D, \mathbb{S}'_i).$$

10. (A birth-death singularity along $\mathcal{A}$): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \varnothing)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \varnothing)$ are elementary parametrized link cobordisms, both of type (3),

[5]An isotopy of an attaching map yields a diffeomorphism of the spaces obtained by attaching a handle, well defined up to isotopy. Handlesliding a handle across another corresponds to an isotopy of the attaching sphere of a handle on the boundary of the 4-manifold obtained by attaching another handle. As such, the diffeomorphism D resulting from handlesliding is well defined up to isotopy.

such that the two critical arcs $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ intersect at exactly one point on the common boundary $W_i \cap W_{i+1}$, then we replace the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ with

$$\mathcal{W}_i' = (W_i \cup W_{i+1}, F_i \cup F_{i+1}, \Phi_{i+1} * \Phi_i, \varnothing),$$

an elementary, parametrized link cobordism of type (1).

11. (A critical point switch along $\mathcal{A}$): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \varnothing)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \varnothing)$ are both of type (3), then they are replaced by a pair $\mathcal{W}_i' = (W_i', F_i', \Phi_i', \varnothing)$ and $\mathcal{W}_{i+1}' = (W_{i+1}', F_{i+1}', \Phi_{i+1}', \varnothing)$, such that if a_i and a_{i+1} are the arcs of $\mathcal{A}|_{F_i \cup F_{i+1}}$ which contain the critical arcs of F_i and F_{i+1}, respectively, then $a_i \neq a_{i+1}$ and a_i contains the critical arc of F_{i+1}' and a_{i+1} contains the critical arc of F_i'. We also have that $\Phi_{i+1} * \Phi_i$ and $\Phi_{i+1}' * \Phi_i'$ are isotopic through diffeomorphisms which fix $Y_i \times \{0\}$.

12. (A critical point switch along Σ): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are two elementary parametrized cobordisms of type (4), then the pair is replaced by a pair

$$\mathcal{W}_i' = (W_i', F_i', \Phi_i', \mathbb{S}_i') \qquad \text{and} \qquad \mathcal{W}_{i+1}' = (W_{i+1}', F_{i+1}', \Phi_{i+1}', \mathbb{S}_{i+1}'),$$

both of type (4), with $W_i \cup W_{i+1} = W_i' \cup W_{i+1}'$ and $F_i \cup F_{i+1} = F_i' \cup F_{i+1}'$, and

$$\Phi_i|_{(Y_i \setminus \mathbb{S}_i) \times \{2\}}(\mathbb{S}_i') = \mathbb{S}_{i+1} \qquad \text{and} \qquad \Phi_i'|_{(Y_i \setminus \mathbb{S}_i') \times \{2\}}(\mathbb{S}_i) = \mathbb{S}_{i+1}'.$$

Furthermore, under the canonical identifications

$$W(Y_i(\mathbb{S}_i), L_i(\mathbb{S}_i), \mathbb{S}_i') \circ W(Y_i, L_i, \mathbb{S}_i) \cong W(Y_i(\mathbb{S}_i'), L_i(\mathbb{S}_i'), \mathbb{S}_i) \circ W(Y_i, L_i, \mathbb{S}_i'),$$

we have that $\Phi_{i+1} * \Phi_i$ and $\Phi_{i+1}' * \Phi_i'$ are isotopic through diffeomorphisms which fix $Y_i \times \{0\}$, as well as the center points of the two bands.

13. (A critical point switch between critical points of $\mathcal{A}$ and Σ): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \varnothing)$ is of type (3) and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ is of type (4), then we replace the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ with a pair

$$\mathcal{W}_i' = (W_i', F_i', \Phi_i', \mathbb{S}_i'), \qquad \text{and} \qquad \mathcal{W}_{i+1}' = (W_{i+1}', F_{i+1}', \Phi_{i+1}', \varnothing),$$

such that $\mathcal{W}_i'$ is of type (4) and $\mathcal{W}_{i+1}'$ is of (3). Furthermore, $W_{i+1} \cup W_i = W_{i+1}' \cup W_i'$ and $F_{i+1} \cup F_i = F_{i+1}' \cup F_i'$, and $\Phi_i(\mathbb{S}_i') = \mathbb{S}_{i+1}$. Furthermore, under the canonical identification (well defined up to isotopy)

$$W(Y_i, L_i, \mathbb{S}_i') \circ W(Y_i, L_i, \varnothing) \cong W(Y_i(\mathbb{S}_i'), L_i, \varnothing) \circ W(Y_i, L_i, \mathbb{S}_i'),$$

we have that $\Phi_{i+1} * \Phi_i$ and $\Phi_{i+1}' * \Phi_i'$ are isotopic through diffeomorphisms which fix $(Y_i \times \{0\}) \cup \{(0,0)\}$.

14. (A critical point of Σ crossing $\mathcal{A}$): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \varnothing)$ is of type $(3.I^+)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ is of type $(4.\mathbf{w})$, then $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ are replaced by a pair

$$\mathcal{W}_i' = (W_i', F_i', \Phi_i', \varnothing) \qquad \text{and} \qquad \mathcal{W}_{i+1}' = (W_{i+1}', F_{i+1}', \Phi_{i+1}', \mathbb{S}_{i+1}'),$$

where $\mathcal{W}_i'$ is of type $(3.II^+)$ and $\mathcal{W}_{i+1}'$ is of type $(4.\mathbf{z})$. Furthermore $W_{i+1} \cup W_i = W_{i+1}' \cup W_i'$ and $F_{i+1} \cup F_i = F_{i+1}' \cup F_i'$. Finally, we have that $\Phi_i^{-1}(\mathbb{S}_{i+1})$ and $\Phi_i^{-1}(\mathbb{S}_{i+1}')$ are equal. The maps $\Phi_{i+1} * \Phi_i$ and $\Phi_{i+1}' * \Phi_i'$ are isotopic diffeomorphisms, relative $Y_i \times \{0\}$. The isotopy between $\Phi_{i+1} * \Phi_i$ and $\Phi_{i+1}' * \Phi_i'$ does not fix the point $(0,0)$ of the band. Instead the image of $(0,0)$ on the surface $\Sigma_i \cup \Sigma_{i+1}$ crosses a single arc of $\mathcal{A}$, transversely, at a single point (cf. Figure 4.27).

To prove invariance of the link Floer cobordism maps, we will prove the following:

Proposition 4.10.3. *If (W, F) is a link cobordism with $F = (\Sigma, \mathcal{A})$, and each component of W and each component of Σ have a nonempty incoming and outgoing end, then any two parametrized Kirby decompositions of (W, F) are weakly equivalent.*

The remainder of this section will be devoted to the proof of Proposition 4.10.3.

4.10.2 Spaces of adapted coordinates

We now prove several useful lemmas about spaces of coordinate functions, which will be necessary for various Morse theoretic arguments to come.

If p is a critical point of a Morse function f of index k, we let $\mathcal{CO}_p(f)$ denote the space of local coordinates in which f takes the form

$$f = f(p) - (x_1^2 + \cdots + x_k^2) + (x_{k+1}^2 + \cdots + x_n^2)$$

(we give $\mathcal{CO}_p(f)$ the C^∞ topology). Note that by picking a base set of coordinates, the set $\mathcal{CO}_p(f)$ is homeomorphic to the space of germs of self diffeomorphisms of $\mathbb{R}^n$ which map $\mathbf{0}$ to $\mathbf{0}$ and preserve the function $f = f(p) - (x_1^2 + \cdots + x_k^2) + (x_{k+1}^2 + \cdots + x_n^2)$. We let $\mathcal{D}(f)$ denote this space of germs of diffeomorphisms. Let $L\mathcal{D}(f) \subseteq \mathcal{D}(f)$ denote the subset of linear maps in $\mathcal{D}(f)$. The following fact is well known (cf. [Cer70, pg 168]). We will use this trick repeatedly.

Lemma 4.10.4. *The space $L\mathcal{D}(f)$ is a deformation retraction of $\mathcal{D}(f)$. Furthermore, $L\mathcal{D}(f)$ is isomorphic to $SO(k, n-k)$.*

Proof. The homotopy is given by linearizing. We define

$$H_t(\phi)(\mathbf{x}) = \begin{cases} \frac{1}{t}\phi(t\mathbf{x}), & \text{if } t \neq 0 \\ d_{\mathbf{0}}(\phi)(\mathbf{x}), & \text{if } t = 0 \end{cases}.$$

First note that that if $\phi \in \mathcal{D}(f)$, then $H_t(\phi) \in \mathcal{D}(f)$ for all t by the following easy computation:

$$f(H_t(\phi)(\mathbf{x})) = f(\tfrac{1}{t}\phi(t\mathbf{x})) = \tfrac{1}{t^2}(f(\phi(t\mathbf{x})) = \tfrac{1}{t^2}f(t\mathbf{x}) = f(\mathbf{x}).$$

To see that the map H_t is a continuous map $\mathcal{D}(f) \times [0,1] \to \mathcal{D}(f)$, we note that if I is a multi-index, then

$$\left(\frac{\partial^I}{\partial \mathbf{x}^I} H_t(\phi)\right)(\mathbf{x}) = t^{|I|-1}\left(\frac{\partial^I}{\partial \mathbf{x}^I}\phi\right)(t\mathbf{x}),$$

from which continuity in the C^∞ topology is straightforward. The space $L\mathcal{D}(f)$ is exactly the set of linear maps which preserve the quadratic form f, which is $SO(k, n-k)$ by definition.

□

We prove a lemma about the space of coordinates adapted to certain Morse functions on 4-manifolds with embedded surfaces:

Lemma 4.10.5. *Let (x, y, w, z) be the standard coordinates on $\mathbb{R}^4$, let Σ denote the subset $\{(x, y, 0, 0)\} \subseteq \mathbb{R}^4$, and let $f = -x^2 + y^2 + w^2 + z^2$. Let $\mathcal{D}_\Sigma(f)$ denote the space of germs of self diffeomorphisms of $\mathbb{R}^4$ mapping $\mathbf{0}$ to $\mathbf{0}$, mapping Σ to itself, and preserving the Morse function f. Then $\mathcal{D}_\Sigma(f)$ deformation retracts onto the subset of linear diffeomorphisms $L\mathcal{D}_\Sigma(f) \subseteq \mathcal{D}_\Sigma(f)$, which is itself canonically isomorphic to $S(O(1,1) \times O(2))$, which has four components. The subset of these which preserve a given orientation of Σ has two components. The set of vector fields $\mathcal{V}$ arising as $\phi_*(v_0)$ for*

$$v_0 = 2\left(-x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y} + w\frac{\partial}{\partial w} + z\frac{\partial}{\partial z}\right)$$

and $\phi \in \mathcal{D}_\Sigma(f)$ is connected.

Proof. The deformation retraction of Lemma 4.10.4 is easily seen to be through maps which preserve Σ and f, so $L\mathcal{D}_\Sigma(f)$ is a deformation retraction of $\mathcal{D}_\Sigma(f)$. Write an element $M \in L\mathcal{D}_\Sigma(f)$ as

$$M = \begin{pmatrix} A & B \\ 0 & D \end{pmatrix},$$

for 2×2 matrices A, B, and D (here the first component corresponds to the space spanned by $\partial/\partial x$, and $\partial/\partial y$ and the second corresponds to the space spanned by $\partial/\partial w$ and $\partial/\partial z$). Asking that M preserves f is easily seen to imply

$$A \in O(1,1), \qquad B = 0, \qquad \text{and} \qquad D \in O(2).$$

Since M must have positive determinant, we see that A and D must both be positively oriented, or both negatively oriented. Hence $L\mathcal{D}_\Sigma(f)$ has four components. Two of these components will preserve the orientation of Σ, and two will reverse it.

To see that $\mathcal{V}$ is connected, it is sufficient to note that the vector field v_0 itself is preserved by the diffeomorphisms obtained by replacing any number of the coordinate functions τ with $-\tau$. We can go between any components of $\mathcal{D}_\Sigma(f)$ in this manner. □

Lemma 4.10.6. *Let (x, y, w, z) be the standard coordinates on $\mathbb{R}^4$ and let $\Sigma = \{(x, y, 0, 0)\}$. Let $\mathcal{D}_\Sigma(f_\Sigma)$ denote the space of germs of diffeomorphisms which fix $\mathbf{0}$, preserve Σ, and which preserve $f_\Sigma(x, y) = -x^2 + y^2$ on Σ. Letting $L\mathcal{D}_\Sigma(f_\Sigma)$ denote the subset consisting of linear diffeomorphisms, then we have that $\mathcal{D}_\Sigma(f_\Sigma)$ deformation retracts on $L\mathcal{D}_\Sigma(f_\Sigma)$, and that $L\mathcal{D}_\Sigma(f_\Sigma)$ has four components, though the set $\mathcal{F}$ of Morse functions of the form $f \circ \phi$, where $\phi \in \mathcal{D}_\Sigma(f_\Sigma)$, is connected.*

Proof. We note that if ϕ fixes $\mathbf{0}$, preserves Σ, and preserves f_Σ on Σ, then the linearizations $H_t(\phi)$ will do the same for all t. Hence as before we get a deformation retraction of $\mathcal{D}_\Sigma(f_\Sigma)$ onto $L\mathcal{D}_\Sigma(f_\Sigma)$. Write an element $M \in L\mathcal{D}_\Sigma(f_\Sigma)$ as

$$M = \begin{pmatrix} A & B \\ 0 & D \end{pmatrix},$$

we see that see that B can be anything, and $A \in O(1, 1)$. Deformation retracting onto the subset with $B = 0$ we see that $\mathcal{D}_\Sigma(f_\Sigma)$ is homotopy equivalent to $(O(1, 1) \times GL(\mathbb{R}^2))^+$, which has four components.

The images of the function $f = -x^2 + y^2 + w^2 + z^2$ under all of these components are easily seen to coincide, since replacing any of the coordinate functions τ with $-\tau$ preserves f. □

We also need to consider spaces of gradient like vector fields which take a standard form near critical points on Σ. We make the following definition:

Definition 4.10.7. Suppose that (W, F) is a link cobordism and $F = (\Sigma, \mathcal{A})$ and that f_Σ is a fixed Morse function on Σ, with only index 1 critical points. We define $\mathcal{M}(f_\Sigma)$ to be the space of Morse functions on W which restrict to f_Σ on Σ, and such that

1. $\mathrm{Crit}(f_\Sigma) = \mathrm{Crit}(f) \cap \Sigma$;
2. There are coordinates (x, y, w, z) (oriented, for W) near each point $p \in \mathrm{Crit}(f_\Sigma)$ such that $f = f(p) - x^2 + y^2 + w^2 + z^2$, and Σ is locally given by $\{(x, y, 0, 0)\}$,

3. All critical points of $f|_{W\setminus\Sigma}$ have distinct critical values.

If f_Σ is a fixed Morse function on Σ and $f \in \mathcal{M}(f_\Sigma)$, we call a vector field v a **gradient like vector field** for f if the following hold:

1. $v(f) > 0$ away from $\mathrm{Crit}(f)$;

2. v is tangent to Σ along Σ;

3. near any critical point $p \in \mathrm{Crit}(f) \setminus \Sigma$ of index k, there are coordinates (x_1, x_2, x_3, x_4) such that
$$f = -x_1^2 - \cdots - x_k^2 + x_{k+1}^2 + \cdots + x_4^2,$$
$$\text{and} \quad v = 2\left(-x_1\tfrac{\partial}{\partial x_1} - \cdots - x_k\tfrac{\partial}{\partial x_k} + x_{k+1}\tfrac{\partial}{\partial x_k} + \cdots + x_4\tfrac{\partial}{\partial x_4}\right);$$

4. Near each $p \in \mathrm{Crit}(f_\Sigma)$, there are coordinates (x, y, w, z), (oriented, for W and Σ) such that Σ is locally given by $\{(x, y, 0, 0)\}$,
$$f = f(p) - x^2 + y^2 + w^2 + z^2, \qquad \text{and} \qquad v = 2\left(-x\tfrac{\partial}{\partial x} + y\tfrac{\partial}{\partial y} + w\tfrac{\partial}{\partial w} + z\tfrac{\partial}{\partial z}\right).$$

We write $\mathcal{V}(f)$ for the set of gradient like vector fields, in the above sense, for f.

Remark 4.10.8. If $f \in \mathcal{M}(f_\Sigma)$, the existence of coordinates in a neighborhood of each $p \in \mathrm{Crit}(f_\Sigma)$ as in the above definition imply that p is an index 1 critical point for both f and f_Σ.

Lemma 4.10.9. *If f_Σ is a fixed Morse function on Σ and $f \in \mathcal{M}(f_\Sigma)$, then the set $\mathcal{V}(f)$ of gradient like vector fields for f is nonempty, and connected.*

Proof. We first show $\mathcal{V}(f)$ is nonempty. First note that at the critical points of f_Σ, we can locally find a gradient like vector field, simply as the gradient under a choice of local coordinates that puts the critical point into the local form in the definition of $\mathcal{M}(f_\Sigma)$. Similarly near a critical point of $f|_{W\setminus\Sigma}$, we can find coordinates which give a gradient like vector field of the correct form. We can now pick a finite set $\mathcal{U}$ of open sets U covering W such that the following hold:

1. U is one of the coordinate charts for a critical point (either on Σ or away from Σ), constructed above. Such open sets can be chosen to be pairwise disjoint;

2. U does not intersect Σ or any of the critical points of f;

3. U has coordinates (x, y, w, z) centered at a point of Σ, such that locally $(x, y, 0, 0)$ corresponds to Σ, and on Σ we have $f|_\Sigma = x$.

On the coordinate charts of the first type, we let associate a vector field v_U, by taking the gradient of f using the local coordinates. On sets of the second type, we let v_U be ∇f, for any Riemannian metric f. On sets of the third type, we let $v_U = \partial/\partial x$. Noting that $v_U(f|_\Sigma) > 0$ for vector fields of this third type, we can shrink the sets of the third type and add sets of the second so that $v_U(f) > 0$ on the entirety of a set U of the third type. We now pick a partition of unity for $\mathcal{U}$, and let v be the associated sum.

We now show that $\mathcal{V}(f)$ is connected. If v and v' are equal on neighborhoods of the critical points of f, then any convex combination will also be a gradient like vector field with a valid local model near each critical point. By Lemma 4.10.5, we can find a path ϕ_t of local diffeomorphisms of W defined near p, which fix p, preserve f and Σ, and such that $\phi_0 = \mathrm{id}$ and $(\phi_1)_*(v) = v'$. Let ρ be a smooth function which is constantly 1 near p, and has support inside an open set U containing p which is in the domain and image of all of the ϕ_t. We then define

$$v_t|_x = \rho(x) \cdot (\phi_t)_*(v|_{\phi_t^{-1}(x)}) + (1 - \rho(x)) \cdot v|_x. \tag{4.11}$$

Outside of U, we have $v_t = v$, and hence each v_t is gradient like away from U. Inside of U, but away from p, v_t is pointwise a convex combination of two vector fields V_i^t which each have $V_i^t(f) > 0$, and hence the pointwise convex combination also satisfies that inequality. In a neighborhood of p, $v_t = (\phi_t)_*(v)$, which by assumption takes has the desired local form for the coordinates induced by ϕ_t. Hence v_t is a path of vector fields in $\mathcal{V}(f)$. In this way, we reduced to the case that the two gradient like vector fields are equal in a neighborhood of the critical points, completing the proof by our previous argument. □

4.10.3 Changing the Morse function away from Σ

In this subsection, we construct appropriately nice paths between Morse functions which agree on Σ.

Lemma 4.10.10. *If f and f' are an arbitrary pair of Morse functions in $\mathcal{M}(f_\Sigma)$, then they can be connected by a path of functions f_t, such that $f_t \in \mathcal{M}(f_\Sigma)$ for all but finitely many t, and at the finite points of time when f_t fails to be in $\mathcal{M}(f_\Sigma)$, one of the following occurs:*

1. *A critical point birth-death occurs away from Σ;*
2. *Two critical points exchange relative value.*

Furthermore, for all t, we have $\mathrm{Crit}(f_\Sigma) = \mathrm{Crit}(f_t) \cap \Sigma$. In addition, for each t, there is a neighborhood of each $p \in \mathrm{Crit}(f_\Sigma)$, with coordinates (x, y, w, z) with $f_t = f_\Sigma(p) - x^2 + y^2 + w^2 + z^2$ where Σ is locally given by $\{(x, y, 0, 0)\}$.

Proof. Let us first prove the claim in the case that f and f' agree in a regular neighborhood of each critical point $p \in \mathrm{Crit}(f_\Sigma)$. In this case, we can simply consider $f_t = tf + (1-t)f'$. Notice that f_t has no critical points along Σ except for those in $\mathrm{Crit}(f_\Sigma)$ since it is constant in t on Σ, and in a neighborhood of $\mathrm{Crit}(f_\Sigma)$. A generic perturbation of f_t near the set of critical points on $W \setminus \Sigma$ will have only birth death singularities, and critical point value switches. Since f_t is already generic near each $p \in \mathrm{Crit}(f_\Sigma)$ (as it's constant near each p), one does not need to perturb near those critical points. On $\Sigma \setminus \mathrm{Crit}(f_\Sigma)$, the functions f_t are constant in t, and also have no critical points since they are equal to f_Σ. Hence no perturbation is needed near Σ.

Using the previous argument, we now just need to show that if f and f' are arbitrary Morse functions in $\mathcal{M}(f_\Sigma)$, then there is a path f_t in $\mathcal{M}(f_\Sigma)$ such that $f_0 = f$ and f_1 agrees with f' in a neighborhood of each point of $\mathrm{Crit}(f_\Sigma)$. Using Lemma 4.10.6 the space of oriented local diffeomorphisms preserving Σ and f_Σ has four components, but the spaces of Morse functions $-x^2 + y^2 + w^2 + z^2$ induced by these four components are the same. In

particular, we can find a path ϕ_t of local diffeomorphisms near a given critical point p, which preserve p, Σ and f_Σ, such that $f' = f \circ \phi_1$, near p. Pick a closed ball B centered at p such that B is in the domain and image of all of the ϕ_t, and let ρ be a bump function which is 1 in a neighborhood of p, and zero outside of B. Consider the time dependent vector field V_t defined by

$$V_t(x) = \rho(x) \cdot \left.\frac{d}{ds}\right|_{s=0} \phi_{t+s}(\phi_t^{-1}(x)).$$

We integrate $V_t(x)$, to define a 1-parameter family Φ_t of diffeomorphisms. We note that Φ_t preserves Σ, f_Σ and p. We note that in a neighborhood of p, we have that Φ_t and ϕ_t agree. Defining $f_t = f \circ \Phi_t$, we have thus connected f by a path in $\mathcal{M}(f_\Sigma)$ to a Morse function which agrees with f' in a neighborhood of p. Doing this for each $p \in \mathrm{Crit}(f_\Sigma)$ reduces the lemma to the previous case, finishing the proof. □

We now define a first class of Morse functions, which we call "nice Morse functions". Since there is no reference to the set of divides, they cannot quite be used to define the link cobordism maps. We will later define a class of Morse functions called "very nice Morse functions", which will also respect the divides, and which can be used to compute the cobordism maps.

Definition 4.10.11. If (W, F) is a decorated link cobordism with $F = (\Sigma, \mathcal{A})$, and f_Σ is a Morse function on Σ with only index 1 critical points we say that $f \in \mathcal{M}(f_\Sigma)$ is **nice** if

1. All critical points of $f|_\Sigma$, occur before any critical points of $f|_{W\setminus\Sigma}$;
2. $f|_{W\setminus\Sigma}$ has only index 1,2 or 3 critical points;
3. All index 1 critical points of $f|_{W\setminus\Sigma}$ occur before the index 2 and 3 critical points, and the index 2 critical points all occur before the index 3 critical points.
4. All critical points of $f|_{W\setminus\Sigma}$ have distinct critical values.

We let $\mathcal{M}^{\mathrm{nice}}(f_\Sigma)$ denote the set of nice Morse functions restricting to f_Σ on Σ.

Definition 4.10.12. We say that a smooth path of functions f_t is a **nice path** if f_t is in $\mathcal{M}^{\text{nice}}(f_\Sigma)$ for all but finitely many t, and at those points of time, one of the following singularities occurs:

1. two critical points of $f|_{W\setminus\Sigma}$ exchange relative value;
2. a birth-death singularity occurs, with higher critical value than any of the critical values of of $f|_\Sigma$.

Remark 4.10.13. Notice that if f_t is a nice path, then necessarily all birth deaths occur away from Σ, since f_t is constant on Σ. Similarly, if a birth-death singularity occurs, then it necessarily must be an index 1/2 or 2/3 birth-death, since the Morse functions before and after the singularity are nice, by definition.

Lemma 4.10.14. *Given any two Morse functions $f_0, f_1 \in \mathcal{M}^{\text{nice}}(f_\Sigma)$, there is a nice path f_t from f_0 to f_1. If $\mathcal{A}$ is a set of divides on Σ, and the functions f_0 and f_1 map W into $[0,2]$, all critical values of $f_i|_{W\setminus\Sigma}$ are greater than 1, and all critical values of $f_i|_\Sigma$ and $f_i|_\mathcal{A}$ are less than 1, then we can further assume that all critical values of $f_t|_{W\setminus\Sigma}$ are also greater than 1 for all t.*

Proof. Using Lemma 4.10.10, we can find a generic path f_t from f_0 to f_1, which is in $\mathcal{M}(f_\Sigma)$ for all but finitely many t, where a critical point value switch occurs, or a birth death singularity occurs. We pick a path of gradient like vector fields v_t for f_t. We note that by assumption v_t is a gradient like vector field such that the descending manifolds of $\text{Crit}(f_\Sigma)$ do not intersect the ascending manifolds of any critical points of $f_t|_{W\setminus\Sigma}$, since all of the descending flowlines from points in $\text{Crit}(f_\Sigma)$ stay in Σ, by assumption. Hence we can modify the path f_t so that all critical points of any $f_t|_{W\setminus\Sigma}$ occur above all of critical points in $\text{Crit}(f_\Sigma)$, using the "independent trajectories principle" ([Kir78, pg. 40]). Indeed we can also assume that they occur above any critical point of $f|_\mathcal{A}$. This can be done while preserving $f_t|_\Sigma$.

Now one can use the techniques of Cerf graphics, developed in [Kir78] (cf. [Cer70] and [HW73]) to manipulate the Cerf graphic associated to the path f_t so that all index 1 critical

points occur before index 2 critical points, which occur before index 3 critical points, and that there are also no index 0 or 4 critical points. There are finitely many points of time when a critical point birth death may occur, or when a critical point value swap may occur. If the the critical values of each $f_t|_{W\setminus\Sigma}$ are greater than 1, then the resulting path (after performing Cerf theory) can also be constructed to satisfy that. □

4.10.4 Changing the Morse function f_Σ on Σ

We now consider different Morse functions f_Σ on Σ. We first define the a set of Morse functions on Σ which respect the divides $\mathcal{A}$ well enough that we can use them to compute the link cobordism maps.

Definition 4.10.15. Given a surface with divides $(\Sigma, \mathcal{A})$ with boundary partitioned into subsets $-L_1$ and L_2, we say a Morse function f on Σ is **$\mathcal{A}$-admissible** if the following hold:

1. The level sets of f are nonempty;
2. Each component of $f^{-1}(t)$ intersects $\mathcal{A}$ nontrivially for each t;
3. There are no index 0 or index 2 critical points;
4. $f|_{\mathcal{A}}$ is Morse;
5. All critical points of f and $f|_{\mathcal{A}}$ have distinct values;

Remark 4.10.16. If f is an $\mathcal{A}$-admissible Morse function on Σ, then the critical points of f cannot intersect $\mathcal{A}$, since such a critical point would be a critical point of f and of $f|_{\mathcal{A}}$, contradicting condition (5) of the previous definition.

Lemma 4.10.17. *If $(\Sigma, \mathcal{A})$ is a surface with divides with boundary partitioned into incoming and outgoing components, such that each component of Σ has nonempty incoming and outgoing boundary components, then there exists an $\mathcal{A}$-admissible Morse function on Σ, and any two $\mathcal{A}$-admissible Morse functions, f_0 and f_1, on Σ can be connected by a path of Morse func-*

tions f_t such that there are only finitely many points of time t where f_t is not $\mathcal{A}$-admissible. At those points of time, one of the following singularities occurs:

1. *A birth-death singularity of $f_t|_{\mathcal{A}}$ occurs.*
2. *Two critical points of $f_t|_{\mathcal{A}}$ switch relative ordering.*
3. *Two critical points of f_t exchange relative ordering.*
4. *A critical point of $f_t|_{\mathcal{A}}$ and a critical point of f exchange relative ordering.*
5. *A critical point of f_t crosses $\mathcal{A}$, transversely.*

Each of the above events can be assumed to occur at distinct times t.

Proof. Suppose that f_0 and f_1 are two $\mathcal{A}$-admissible Morse functions. Using standard Morse theory, one can find a path of functions f_t which are Morse on Σ except at finitely many t, where a birth/death singularity occurs, and such that there are no index 0 or index 2 critical points, throughout (cf. [Juh15, Theorem 2.21]). Since there are no index 0 or index 2 critical points of f_t for any t, there can be no birth/death singularities at all, so f_t is Morse for all t. To ensure that each component of $(f_t)^{-1}(s)$ intersects $\mathcal{A}$ nontrivially, one can then change the path f_t of Morse functions by first taking f_0 and considering the path of Morse functions $f_0 \circ \phi_t$, where ϕ_t is an isotopy of Σ which "scrambles" the surface suitably, concatenating with the path of Morse functions $f_t \circ \phi_1$, then concatenating with the path of Morse functions $f_1 \circ \phi_{1-t}$. Such a path has the property that $(f_t)^{-1}(s)$ intersects $\mathcal{A}$ nontrivially for each t and s. After performing a generic perturbation of f_t, for all but finitely many t, f_t will be Morse when restricted to $\mathcal{A}$, all critical values of f_t and $f_t|_{\mathcal{A}}$ will be distinct and occur at distinct points on Σ. At finitely many t, one of the singularities listed can occur. □

Definition 4.10.18. If $(W, \Sigma, \mathcal{A})$ is a decorated link cobordism, we will say that a Morse function $f : W \to [0, 2]$ is **very nice**, and write $f \in \mathcal{M}^{\text{nice}}(\mathcal{A})$, if the following holds:

1. f is a nice Morse function;

2. $f|_\Sigma$ is $\mathcal{A}$-admissible;

3. All critical values of $f|_\Sigma$ or $f|_\mathcal{A}$ are less than 1, and all critical values of $f|_{W\setminus\Sigma}$ are greater than 1;

4. All critical values of f and $f|_\mathcal{A}$ are distinct.

We say that a path of smooth functions $f_t : W \to [0,2]$ is a **very nice path** if each component of $(f_t)^{-1}(s)$ intersects $\mathcal{A}$ nontrivially for each s and t, f_t is a nice path, and each f_t is a very nice Morse function for all but finitely many t, where f_t fails to be a nice Morse function, or one of the singularities listed in Lemma 4.10.17 occurs. Furthermore, all such singularities occur at distinct times t.

We now prove the following:

Lemma 4.10.19. *Suppose (W,F) is a link cobordism and $(f_\Sigma)_0$ and $(f_\Sigma)_1$ are two $\mathcal{A}$-admissible Morse functions from Σ to $[0,2]$ such that all critical points of $(f_\Sigma)_0, (f_\Sigma)_1, (f_\Sigma)_0|_\mathcal{A}$ and $(f_\Sigma)_1|_\mathcal{A}$ have value less than 1. Then there is a very nice path f_t of smooth functions from f_0 to f_1 such that $(f_0)|_\Sigma = (f_\Sigma)_0$ and $(f_1)|_\Sigma = (f_\Sigma)_1$.*

Our proof is by explicit construction. We first describe the construction, and then prove that it can be used to construct the very nice path in the previous lemma.

Equip a neighborhood of $Y_1 \cup \Sigma$ with a Riemannian metric. This yields an identification of $N_{r_0}(\Sigma)$, the disk bundle of radius r_0 inside of the normal bundle, with a regular neighborhood of Σ in W. Let

$$\pi : N_{r_0}(\Sigma) \to \Sigma,$$

denote projection of the normal bundle onto Σ. Suppose $f_\Sigma : \Sigma \to [0,2]$ is a Morse function which has only index 1 critical points, and all critical values of f_Σ and $f_\Sigma|_\mathcal{A}$ are strictly less than 1. Writing Σ_1 for $f_\Sigma^{-1}([0,1])$ we assume that $N_2(\Sigma_1)$ (the radius 2 closed disk bundle) intersects Y_2 (the outgoing boundary of W) trivially. Pick a Morse function $h : N(Y_1) \to [0,2]$ on a collar neighborhood of Y_1 such that $h = k \cdot (f_\Sigma \circ \pi)$ for a constant $k > 1$, on the

intersection of $N_{r_0}(\Sigma)$ and $N(Y_1)$. We assume that $N(Y_1) \cap Y_2 = \varnothing$ as well. We also assume that h and $N(Y_1)$ have been selected so that f_Σ and $f_\Sigma|_{\mathcal{A}}$ have no critical points on $N(Y_1)$. Let

$$r : N_{r_0}(\Sigma) \to [0, r_0]$$

denote the radial distance function. For notational simplicity, normalize the metric so that $r_0 = 2$. We will occasionally write

$$N_s(Y_1) \subseteq N(Y_1)$$

for the set of points x where $h(x) \leq s$.

Let

$$\rho : [0, 2] \to [0, 1]$$

be a smooth function such that the following hold:

1. $0 \leq \rho \leq 1$;
2. ρ is nondecreasing, and $\rho'(t) > 0$ at any t where $\rho(t) \neq 1$;
3. $\rho(t) = t$ for t in a neighborhood of 0, and ρ is constantly 1 in a neighborhood of $t = 2$.

Let $F : N(Y_1) \cup N_2(\Sigma) \to \mathbb{R}$ be defined by $f_\Sigma \circ \pi$, or h/k, whichever is defined (by assumption they are equal when they are both defined).

We now consider the Morse function

$$f = F + (\rho \circ r)^2(\rho \circ h), \tag{4.12}$$

viewed as a function on $N(Y_1) \cup N_2(\Sigma)$. We define

$$W_0^f = f^{-1}([0, 1]).$$

Lemma 4.10.20. *The function f is Morse on $W_0^f = f^{-1}([0, 1])$, a neighborhood of $Y_1 \cup \Sigma_1$. The critical points of f coincide exactly with the critical points of f_Σ, and each critical point of f_Σ (which is of index 1 by assumption) corresponds to an index 1 critical point of f. There*

are local coordinates (x, y, w, z) *near any critical point* p *on* Σ *such that* Σ *is locally given by* $\{(x, y, 0, 0)\}$ *and*

$$f = f_\Sigma(p) - x^2 + y^2 + w^2 + z^2.$$

Furthermore, W_0^f *is a cobordism from* Y_1 *to the closed, properly embedded submanifold,* $f^{-1}(1)$*, of* W*.*

Proof. Note first that if $\rho(r(x)) = \rho(h(x)) = 1$, then $f(x) > 1$, so $x \notin W_0^f$. Hence if $x \in W_0^f$, at least one of $\rho(r)$ and $\rho(h)$ must be strictly less than 1. If $0 < \rho(r(x)) < 1$, then $x \in N_2(\Sigma)$, and we can compute

$$\frac{d}{dr} f = 2\rho(r)\rho'(r)\rho(h),$$

which is strictly positive when $0 < \rho(r) < 1$ and $h > 0$. Hence there are no critical points with $0 < \rho(r) < 1$. If $\rho(r) = 1$ at a point $x \in W_0^f$, then $\rho(h) < 1$ at x, by what we have said, so $x \in N_2(Y_1)$. Hence $f = h/k + \rho(h)$, which is easily seen to not have any critical points on $N_2(Y_1)$.

Hence the only remaining case where critical points can occur is $\rho(r) = 0$, which corresponds exactly to being on Σ. Since $f|_\Sigma = f_\Sigma$, critical points of f can only occur at critical points of f_Σ. Noting f_Σ has no critical points when $\rho(h) < 1$, we see that at any critical point of f_Σ, we have $\rho(h) = 1$. In particular, near any critical point for f, we have that locally,

$$f = (f_\Sigma \circ \pi) + r^2.$$

Picking coordinates (x, y) on Σ near a critical point p of f_Σ, such that $f_\Sigma = -x^2 + y^2$, and extending these to coordinates (x, y, w, z) near p on the $SO(2)$ disk bundle $N_2(\Sigma)$ such that $r^2 = w^2 + z^2$, we have

$$f = f_\Sigma(p) - x^2 + y^2 + w^2 + z^2.$$

We now need to show that $Y' = f^{-1}(1)$ is a closed, properly embedded submanifold of W, and that W_0^f is a cobordism from Y_1 to Y'. The claim that it is properly embedded is equivalent to showing that it doesn't intersect ∂W. Note that $Y_1 = f^{-1}(0)$, so $Y_1 \cap Y' = \varnothing$.

We note $W_0^f \subseteq N_2(Y_1) \cup N_2(\Sigma_1)$ (where $\Sigma_1 = f_\Sigma^{-1}([0,1])$), and the latter two sets intersect Y_2 trivially by assumption. Hence Y' is a compact, embedded submanifold in $\operatorname{int} W$ and $W_0^f \subseteq W$ is a cobordism from Y_1 to Y'. □

As an intermediate technical results, we now prove a very slight modification of [GWW12, Lemma 3.1].

Lemma 4.10.21. *Suppose that $f_t : W \to [0,1]$ is a path of Morse functions on (W, Σ) such that $f_t \in \mathcal{M}(f_t|_\Sigma)$ for all t (in particular all critical values are distinct, and there are no birth death singularities of f_t, but we are assuming nothing about a dividing set). Then there is a smooth isotopy $\psi_t : W \to W$ such that $\psi_t|_{\partial W} = \mathrm{id}$ for all t, and an isotopy $\phi_t : \mathbb{R} \to \mathbb{R}$ such that*

$$f_t = \phi_t \circ f_0 \circ \psi_t^{-1}.$$

Proof. If we don't require ψ_t to map Σ to Σ, then [GWW12, Lemma 3.1] gives us the result immediately. We briefly sketch the modification, to our setting. First pick a $\phi_t : \mathbb{R} \to \mathbb{R}$ such that $\phi_t \circ f_t$ and f_0 have the same critical values. Such a ϕ_t exists because by assumption all critical values of f_t are distinct, and there are no critical point birth death singularities. Write g_t for $\phi_t \circ f_t$. Pick a path of gradient like vector fields V_t for g_t such that $V_t|_\Sigma$ is tangent to Σ. Away from Σ and the critical points of g_t, we define a vector field

$$\overline{V}_t = -\frac{\partial_t g_t}{V_t(g_t)} \cdot V_t.$$

Note that

$$dg_t(\overline{V}_t) + \partial_t g_t = 0.$$

Extend $\overline{V}_t$ over any critical points p of g_t by picking local coordinates $\theta_t : B^4 \to W$ such that, if $p \in \Sigma$, we have

$$(g_t \circ \theta_t)(x, y, w, z) = c - x^2 + y^2 + w^2 + z^2$$

(with Σ locally given by $\{(x, y, 0, 0)\}$), and in a neighborhood of the critical point p defining

$$\overline{V}_t(y) = \left.\frac{d}{ds}\right|_{s=0} \theta_{s+t}(\theta_t^{-1}(y)).$$

We use a similar formula near critical points which are not on Σ. We patch together the various definition of $\overline{V}_t$ using partitions of unity, and note that they satisfy

$$dg_t(\overline{V}_t) + \partial_t g_t = 0.$$

We define ψ_t to be the flow of the time dependent vector field $\overline{V}_t$, and note that the latter differential equation implies $g_t \circ \psi_t$ is constant, implying the claim. $\square$

We can now prove Lemma 4.10.19:

Proof of Lemma 4.10.19. Suppose $(f_\Sigma)_0 : \Sigma \to [0,2]$ and $(f_\Sigma)_1 : \Sigma \to [0,2]$ are two $\mathcal{A}$-admissible Morse functions, such that all critical values of $(f_\Sigma)_0, (f_\Sigma)_1, (f_\Sigma)_0|_{\mathcal{A}}$ and $(f_\Sigma)_1|_{\mathcal{A}}$ are less than 1. Using Lemma 4.10.17, we can find a path $(f_\Sigma)_t$ of Morse functions on Σ from $(f_\Sigma)_0$ to $(f_\Sigma)_1$, which are $\mathcal{A}$-admissible for all but finitely many t, where one of the listed singularities occurs. We arrange so that all critical values of $(f_\Sigma)_t$, and $(f_\Sigma)_t|_{\mathcal{A}}$ occur inside of the interval $(0,1)$, for all t. We then use the construction in Equation (4.12) to extend each $(f_\Sigma)_t$ smoothly to a smooth function $f_{t,0}$ which is defined on neighborhoods of $Y_1 \cup \Sigma$. We now extend $f_{0,0}$ to a Morse function f_0 on all of W with $f_0|_\Sigma = (f_\Sigma)_0$. Recalling the notation $W_0^{f_{0,0}} = f_{0,0}^{-1}([0,1])$, since each component of $W \setminus W_0^{f_{0,0}}$ has a nonempty incoming and outgoing end, and $\Sigma \cap (W \setminus W_0^{f_{0,0}})$ is diffeomorphic to a collection of cylinders with trivial divides, we can pick the extension f_0 on W such that f_0 has only critical points of index 1,2 and 3, all critical values are distinct, and all index 1 critical points occur before index 2 and 3 critical points, and index 2 critical points occur before index 3 critical points. Furthermore, we can assume that on $W \setminus W_0^{f_{0,0}}$, the functions $f_0|_\Sigma$ and $f_0|_{\mathcal{A}}$ have no critical points.

By adapting the proof of Lemma 4.10.21, we can find a 1-parameter family of diffeomorphisms

$$\psi_t : W \setminus W_0^{f_{0,0}} \to W \setminus W_0^{f_{t,0}},$$

mapping $\Sigma \setminus W_0^{f_{0,0}}$ to $\Sigma \setminus W_0^{f_{t,0}}$ such that $(f_\Sigma)_t \circ \psi_t = (f_\Sigma)_0$. Using this family ϕ_t we can simply push forward the Morse function f_0 on $W \setminus W_0^{f_{0,0}}$ to smoothly extend the family $f_{t,0}$

to a family of smooth functions f_t on all of W, which restrict to $f_{t,0}$ on $W_0^{f_{t,0}}$. This causes no new birth death singularities and no critical point value switches away from Σ, since $f_t|_{W \setminus W_0^{f_{t,0}}}$ is simply the push forward of $f_{t,0}$ to $W \setminus W_0^{f_{t,0}}$. Also all critical values of $f_t|_\Sigma$ and $f_t|_{\mathcal{A}}$ are less than 1, and all critical values of $f_t|_{W \setminus \Sigma}$ are greater than 1. At the finitely many points where $(f_\Sigma)_t$ is not $\mathcal{A}$-admissible, the functions f_t are nice but fail to be very nice, and at each such t one of the listed moves occurs. □

Combining Lemmas 4.10.14 and 4.10.19, we get the following:

Lemma 4.10.22. *Suppose $(W, \Sigma, \mathcal{A})$ is a decorated link cobordism, and that f_0 and f_1 are two very nice Morse functions from W to $[0, 2]$. Then there is smooth path of functions $f_t : W \to [0, 2]$ which is a very nice path.*

Proof. Using Lemma 4.10.19, we can find a very nice path of Morse functions f'_t such that $f'_0|_\Sigma = f_0|_\Sigma$ and $f'_1|_\Sigma = f_1|_\Sigma$. Now by Lemma 4.10.14, we can find nice paths of functions from f'_0 to f_0 and a nice path of functions from f'_1 to f_1. By definition these are paths of functions which are constant along Σ (and hence their restrictions to Σ are $\mathcal{A}$-admissible), and are nice Morse functions, except at finitely many t, when a critical value swap away from Σ, or a birth-death singularity away from Σ, may occur. Furthermore, by the second part of Lemma 4.10.14, we can also assume that the critical values of the two paths restricted to $W \setminus \Sigma$ are all in the interval $(1, 2)$, and that the critical values of the functions restricted to Σ and $\mathcal{A}$ all lie in $(0, 1)$.

Concatenating these three paths, we get a very nice path of functions from f_0 to f_1. □

4.10.5 Morse data and parametrized Kirby decompositions

We now describe how a collection of Morse data on a decorated link cobordism gives a parametrized Kirby decomposition, well defined up to weak equivalence.

Definition 4.10.23. Suppose that $(W, \Sigma, \mathcal{A})$ is a decorated link cobordism and that $f : W \to [0, 2]$ is a Morse function and $\mathbf{b} \subseteq [0, 2]$ is a collection of regular values of f, with

$\mathbf{b} = (b_1, \ldots, b_n)$ with $0 = b_1 < \cdots < b_n = 2$. Define $W_i = f^{-1}([b_i, b_{i+1}])$ and $\Sigma_i = \Sigma \cap W_i$. We say that $(f, \mathbf{b})$ is **very nice Morse data** if f is a very nice Morse function and on each (W_i, Σ_i), exactly one of the following holds

1. f and $f|_{\mathcal{A}}$ have no critical points in W_i (i.e. $(W_i, \Sigma_i, \mathcal{A}_i)$ is a trivial cobordism);
2. $f|_{W \setminus \Sigma}$ has a single critical point in W_i, which is of index 1 or 3;
3. $f|_{W \setminus \Sigma}$ has (arbitrarily many) index 2 critical points in W_i, and no other W_i contains any index 2 critical points;
4. $f|_{\Sigma}$ has a single index 1 critical point in W_i (which by assumption is also an index 1 critical point of f);
5. $f|_{\mathcal{A}}$ has a single critical point in W_i.

To a collection of very nice Morse data, $(f, \mathbf{b})$, on a decorated link cobordism $(W, \Sigma, \mathcal{A})$ we now describe a parametrized Kirby decomposition

$$\mathcal{K}(f, \mathbf{b}).$$

In Lemma 4.10.24 we show that $\mathcal{K}(f, \mathbf{b})$ is determined by f and $\mathbf{b}$ up to weak equivalence.

Write $W_i = f^{-1}([b_i, b_{i+1}])$, $\Sigma_i = \Sigma \cap W_i$ and $\mathcal{A}_i = \Sigma_i \cap \mathcal{A}$. Also define $F_i = (\Sigma_i, \mathcal{A}_i)$. Each (W_i, Σ_i) is an (unparameterized) elementary cobordism in the sense of Definition 4.9.3, with a Morse function $f_i = f|_{W_i}$, which restricts to be Morse on Σ_i and $\mathcal{A}_i$. To construct a parametrized Kirby decomposition, we need to construct a parametrization on each (W_i, Σ_i). Using Lemma 4.10.9, we can find a gradient like vector field $v \in \mathcal{V}(f)$, which restricts to Σ to be a gradient like vector field of $f_i|_{\Sigma_i}$.

Suppose first that the Morse function f_i on (W_i, Σ_i) has no critical points. We set $\mathbb{S}_i = \varnothing$, and the flow of $v/v(f)$ immediately gives a parameterizing diffeomorphism Φ_i. Also note that by Lemma 4.10.9, the space of gradient like vector fields for f_i is connected, so any two parametrizations Φ_i made in this way are isotopic relative Y_i. The parametrized cobordism

$(W_i, F_i, \Phi_i, \varnothing)$ is then an elementary parametrized cobordism of type (1) or (3), depending on whether $f_i|_{\mathcal{A}}$ has a critical point or not.

Now suppose that $f_i|_{W_i \setminus \Sigma_i}$ has a single index 1 or index 3 critical point, or arbitrarily many index 2 critical points. We will construct and framed sphere or framed link and a parameterizing diffeomorphism for (W_i, F_i) turning it into an elementary parametrized cobordism of type (2). Suppose first that f_i has a single critical point of index 1, 2 or 3. The vector field v and a choice of local coordinates in a neighborhood of the critical point give a well defined parametrization. One description of this procedure can be found in [Juh15, Section 2.3.1], though we will describe an approach that is slightly different, and more obviously gives a smooth parametrization. This approach is described in detail in [Pal63, Section 11]. A similar account, using "spherical modification" instead of handle attachment, can be found in [Kos93, Section VII.2].

Let p be an index k critical point of $f|_{W_i}$ for $k \in \{1, 2, 3\}$, away from Σ, and suppose that p is the only critical point of f inside of W_i. Let $Y_i = f^{-1}(b_i)$ and $Y_{i+1} = f^{-1}(b_{i+1})$, and define L_i and L_{i+1} as the intersections of Y_i and Y_{i+1} with Σ. For notational convenience, we will describe the handle attachment procedure for arbitrary ambient dimension n (though we will always be assuming that $n = 2$ or $n = 4$). Pick coordinates $(x_1, \dots, x_n)$ near p such that $f = f(p) - (x_1^2 + \cdots + x_k)^2 + (x_{k+1}^2 + \cdots + x_n^2)$. Pick a gradient like vector field v for f on W_i as well, which is equal to

$$v = 2\left(-x_1\frac{\partial}{\partial x_1} - \cdots - x_k\frac{\partial}{\partial x_k} + x_{k+1}\frac{\partial}{\partial x_k} + \cdots + x_n\frac{\partial}{\partial x_n}\right)$$

near p. The descending manifold is $\{(x_1, \dots x_k, 0, \dots, 0)\}$ and the ascending manifold is $\{(0, \dots, 0, x_{k+1}, \dots, x_n)\}$. Fix a small $\epsilon > 0$, such that the ball of radius 2ϵ is contained in the domain of the coordinates centered at p.

We pick a smooth function $\lambda : \mathbb{R} \to \mathbb{R}$ which has $\lambda'(t) \leq 0$ and satisfies $\lambda(x) = 1$ for $x \leq \frac{1}{2}$, $\lambda(x) > 0$ if $x < 1$ and $\lambda(x) = 0$ if $x \geq 1$. Notice that the space of such functions λ is contractible, since convex combinations of such functions satisfy the those properties as

well. We consider the function

$$g(x) = f(x) - \frac{3\epsilon}{2}\lambda\left(\frac{x_{k+1}^2 + \cdots + x_n^2}{\epsilon}\right).$$

One notes that

$$f^{-1}([b_i, c-\epsilon]) \subseteq g^{-1}([b_i, c-\epsilon]).$$

Write

$$H = g^{-1}([b_i, c-\epsilon]) \setminus f^{-1}([b_i, c-\epsilon]).$$

According to [Pal63, pg. 316], H is homeomorphic to $D^k \times D^{n-k}$, via a homeomorphism which can be canonically specified in terms of λ, ϵ and the choice of coordinates. Furthermore this homeomorphism restricts to a diffeomorphism on $(\text{int } D^k) \times D^{n-k}$, and also restricts to a smooth embedding of $S^{k-1} \times D^{n-k}$ into $f^{-1}(c-\epsilon)$, which is a framed k-sphere in $f^{-1}(c-\epsilon)$. The normalized gradient like vector field, $v/v(f)$, induces a well defined diffeomorphism $Y_i \times [0,1]$ with $f^{-1}([b_i, c-\epsilon])$. The framed k-sphere in $f^{-1}([b_i, c-\epsilon])$ thus naturally induces a framed k-sphere $\mathbb{S}_i$ in $Y_i \times \{0\}$, by using this identification. The homeomorphism of H with $D^k \times D^{n-k}$ described above thus yields a diffeomorphism of $(Y_i \times [0,1]) \cup_{\mathbb{S}_i} (D^k \times D^{n-k})$ with $g^{-1}([b_i, c-\epsilon])$, where the smooth structure on the former space is obtained by attaching the handle and then smoothing corners. By the Proposition from [Pal63, page 319], the gradient like vector field v is transverse to the level sets of g, except at the single critical point p of g (which is also the critical point of f). Hence by flowing along v, the set $f^{-1}([b_i, b_{i+1}]) \setminus g^{-1}([b_i, c-\epsilon])$ is diffeomorphic to $g^{-1}(c-\epsilon) \times [0,1]$, by a diffeomorphism which is unique up to isotopy, relative $g^{-1}(c-\epsilon)$. By picking a collar neighborhood (which is determined up to isotopy) of $g^{-1}(c-\epsilon)$. We thus get a diffeomorphism from $(Y_i \times [0,1]) \cup_{\mathbb{S}_i} (D^k \times D^{n-k})$ to W which is determined up to isotopies fixing $Y_i \times \{0\}$.

Note that even for fixed gradient like vector field and coordinates near p, the construction still relies on a choice of ϵ and λ, though the space of such λ is contractible, as we've already remarked, so any two choices yield framed spheres which are isotopic away from L_i, resulting in Move (2).

In the case that (W_i, F_i) is an elementary cobordism with only index 2 critical points, one can still use the above strategy to arrive at a parametrization $\Phi_i : W(Y_i, L_i, \mathbb{S}_i) \to (W_i, F_i)$. The space of gradient like vector fields for $f|_{W_i}$ is still connected. A generic gradient like vector field will have ascending and descending manifolds of the critical points which are transverse, and hence generically the intersection of the ascending and descending manifolds of two index 2 critical points will be empty, allowing us to perform the above construction whenever the choice of gradient like vector field is generic (paths of gradient like vector fields may of course result in handleslides, as we discuss in the following lemma).

We lastly consider cobordisms of type (4), where the critical point p is of index 1, and is along Σ. In this case, we pick a gradient like vector field v such that in oriented coordinates (x, y, w, z) near the critical point

$$f = f(p) - x^2 + y^2 + w^2 + z^2, \qquad \Sigma = \{(x, y, 0, 0)\},$$

$$\text{and} \qquad v = 2\left(-x\tfrac{\partial}{\partial x} + y\tfrac{\partial}{\partial y} + w\tfrac{\partial}{\partial w} + z\tfrac{\partial}{\partial z}\right).$$

The space of such coordinates where f takes this form has four connected components by Lemma 4.10.5. The space of coordinates which restrict to oriented coordinates on Σ has two components. Using one of these sets of coordinates, we can now adapt the above procedure from [Pal63], as we did for critical points which were away from Σ. In this case, using function g constructed above, the procedure gives a parametrization of (W, Σ) as $(Y_i \times [0,1], L_i \times [0,1]) \cup H \cup (g^{-1}(c+\epsilon) \times [1,2])$. Furthermore H is canonically identified as $D^1 \times D^3$. Our choice of local coordinates identified Σ with the subset $\{(x, y, 0, 0)\} \subseteq \mathbb{R}^4$. It is straightforward to check that using this identification, the set $B \subseteq D^1 \times D^3$ defined by $B = \{(x, y, 0, 0) : -1 \leq x \leq 1, -1 \leq y \leq 1\}$ is mapped to Σ, and the identification also sends $(0, 0) \in B$ to the critical point of f.

The construction of the parametrized Kirby decomposition $\mathcal{K}(f, \mathbf{b})$ involved choices of local coordinates, and gradient like vector fields. Nevertheless, we have the following:

Lemma 4.10.24. *Given a Morse function f and a collection of regular values $\mathbf{b}$ as in Definition 4.10.23, any two choices of gradient like vector fields and local coordinates in a*

neighborhood of the critical points of f yield weakly equivalent parametrized Kirby decompositions $\mathcal{K}(f, \mathbf{b})$.

Proof. We first consider elementary cobordisms of type (1) or (2). We note that the spaces of local coordinates in a neighborhood of each critical point which put f into the correct local model are not connected, though the components are related by changing the sign of two of the coordinate functions appropriately, by Lemma 4.10.4. The resulting vector fields are unaffected, and hence the associated spaces of vector fields are connected. By performing Move (5) to change the orientation of the framed sphere $\mathbb{S}_i$, we can assume that the two choices of local coordinates are in the same component in the appropriate space of local coordinates.

On an elementary parametrized cobordism with a Morse function with no critical points, or exactly one critical point, for two choices of local coordinates and associated gradient like vector fields with the correct local model, if the coordinates are in the same component in the appropriate space of local coordinates, then we can simultaneously change the coordinates and the gradient like vector field (cf. Equation (4.11)). Such a move results in an isotopy of the framed sphere away from L_i (Move (2)).

We now consider the case that f_i has multiple index 2 critical points, away from L_i. In this case, the descending and ascending manifolds of a generic gradient like vector field of each critical point will be transverse, and hence disjoint. For a generic path of gradient like vector fields, there will be finitely many points of time when there is a single flow line between two critical points. At such a point of time, a handleslide of one of the framed 1-spheres occurs, which results in Move (9). Over the course of time when the path of gradient like vector fields have no flow lines amongst each other, the associated parametrized Kirby decompositions are related to each other again by a sequence of Moves (2) and (3).

We lastly consider elementary cobordisms of type (4) (framed 0-spheres along L). By Lemmas 4.10.5 and 4.10.9, different choices of local coordinates or gradient like vector field yield framed spheres and parametrizations which are related by an isotopy of the framed

sphere along L_i (Move (3)) or a switch of the orientation of the parametrization (Move (5)). □

Lemma 4.10.25. *Given a parametrized Kirby decomposition $\mathcal{K}$ of (W, F), one can find very nice Morse data $(f, \mathbf{b})$ for (W, F) such that*

$$\mathcal{K} = \mathcal{K}(f, \mathbf{b}).$$

Proof. One can modify the construction from [Juh15, Lemma 2.14] to get a Morse function on each elementary cobordism which induce the correct parametrization for a choice of gradient like vector field. By modifying them near the boundary, one can ensure that they glue together to form a smooth function on all of W. □

4.10.6 Proof of Proposition 4.10.3

We now have the tools to prove that all parametrized Kirby decompositions of a given (W, F) are weakly equivalent:

Proof of Proposition 4.10.3. Given two parametrized Kirby decompositions, $\mathcal{K}_1$ and $\mathcal{K}_2$, one uses Lemma 4.10.25 to find Morse functions f_1 and f_2, together with collections of regular values, $\mathbf{b}_1$ and $\mathbf{b}_2$, and choices of local coordinates and gradient like vector fields so that

$$\mathcal{K}_1 = \mathcal{K}(f_1, \mathbf{b}_1), \qquad \text{and} \qquad \mathcal{K}_2 = \mathcal{K}(f_2, \mathbf{b}_2).$$

Using Lemma 4.10.22, we can find a very nice path f_t between f_1 and f_2.

At all but finitely many t, the functions f_t are very nice Morse functions, and hence for some choice of regular values $\mathbf{b}_t$ yield very nice Morse data $(f_t, \mathbf{b}_t)$, which determine a parametrized Kirby decomposition $\mathcal{K}(f_t, \mathbf{b}_t)$ for a choice of gradient like vector field v_t and local coordinates near each critical point. By Lemma 4.10.24, changing the choices of local coordinates or gradient like vector fields v_t results in a weak equivalence. Changing the set of regular values $\mathbf{b}$ can be achieved by adding and removing points, and hence by applications of Move (4).

At finitely many points of time, a critical point value switch or a critical point birth-death occurs for the functions $f_t|_{W\setminus\Sigma}$. The birth death singularities are limited to index 1/2 and index 2/3 birth death singularities, since nice Morse functions cannot have index 0 or index 4 critical points. At such points of time, by perturbing the path of functions slightly, we can assume that for a small subinterval of time around the singularity, the path of functions f_t is supported in a small ball containing the critical points. The existence of such normal forms is established in [Cer70] (cf. [Juh15, pg. 17]). As such the parametrized Kirby decompositions immediately preceding and following the singularity may be related by Moves (6), (7) and (8).

If $[a, b]$ is an interval of time which contains no critical point birth-deaths or critical value swaps of f_t, then by Lemma 4.10.21, we can find an isotopy $\phi_t : \mathbb{R} \to \mathbb{R}$ (ranging over $t \in [a, b]$, with $\psi_a = \mathrm{id}$) and an isotopy $\psi_t : (W, \Sigma) \to (W, \Sigma)$, which is the identity on ∂W for all t, such that $f_t = \phi_t \circ f_a \circ \psi_t$. In particular, given a choice of $\mathbf{b}_a$, we can take the parametrized Kirby decomposition $\mathcal{K}(f_a, \mathbf{b}_a)$, and push it forward along the diffeomorphism ψ_t to get a parametrized decomposition (of undecorated link cobordisms) $(\psi_t)_*\mathcal{K}(f_a, \mathbf{b}_a)$. A description of the pushforward decomposition is described immediately before Definition 4.10.1.

For all but finitely many $t \in [a, b]$, the critical values of $f_t|_{\mathcal{A}}$ will be disjoint from the critical values of f_t. Additionally, at all but finitely many t, the parametrization $(\psi_t)_*\Phi_i$ for an embedded 0-sphere along a link will map the point $(0, 0) \in B$ into $\Sigma \setminus \mathcal{A}$. If there are no such t where $f_t|_{\mathcal{A}}$ and f_t share a critical value, and there are no t where the point $(0, 0)$ in a band is mapped to $\mathcal{A}$, then $\mathcal{K}(f_a, \mathbf{b}_a)$ and $\mathcal{K}(f_b, \mathbf{b}_b)$ are related by the strong isotopy move (Move (1)) by Lemma 4.10.21, where $\mathbf{b}_b = \phi^{-1}(\mathbf{b}_a)$. More generally, at finitely many t in $[a, b]$ one of the singularities listed in Lemma 4.10.19 may occur. By simply adding and removing points to $\mathbf{b}$ at intermediate steps of the isotopy, we can find a $\mathbf{b}_b$ such that $\mathcal{K}(f_a, \mathbf{b}_a)$ and $\mathcal{K}(f_b, \mathbf{b}_b)$ are related by a sequence of strong isotopies (Move (1)), critical point value switch between a critical point of $f_t|_{\Sigma}$ and a critical point of $f_t|_{\mathcal{A}}$ (Move (13)), critical point value switches between critical points of $f|_{\mathcal{A}}$ (Move (11)), critical point birth-deaths of $f_t|_{\mathcal{A}}$

(Move (10)), or isotopies of a parametrization Φ_i for a framed 0-sphere along a link, which send the center point of the band transversely across an arc in $\mathcal{A}$ (Move (14)). We finally remark that throughout the process, we may need to change the gradient like vector fields used, but for a fixed Morse function f and choice of regular values $\mathbf{b}$, according to Lemma 4.10.24, any two gradient like vector fields, v and v' yield parametrized Kirby decompositions which are weakly equivalent.

We make an additional remark about Move (14), where a critical point of Σ crosses $\mathcal{A}$. By a sequence of Moves (10) and (11), we can isotope the Morse function in a neighborhood of the singularity to arrange for it to take the form shown in Figure 4.27. For such a path of Morse functions f_t on Σ yielding the singularity shown in Figure 4.27, the surfaces involved in the new elementary parametrized cobordisms $\mathcal{W}_i'$ and $\mathcal{W}_{i+1}'$, appearing in the statement of Move (14), are the ones labeled in that Figure. See also Figure 4.31.

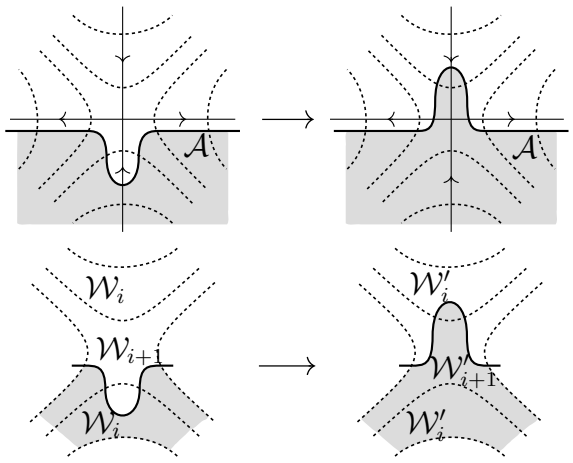

Figure 4.27: **A local model for Move** (14)**.** By possibly composing the path of Morse functions with an isotopy of the surface, we can assume that a singularity of type (5) takes the local model shown on top. The thick solid lines indicate $\mathcal{A}$, while the dashed lines indicate the level sets of f_t, and the thin solid lines with arrows are upward gradient flowlines of f_t. On the bottom, we indicate the change in associated parametrized Kirby decomposition described in the definition of the move. The dotted lines correspond to level sets of the Morse function.

□

4.11 Construction and invariance of the cobordism maps

In this section, we provide a precise construction of the maps $F_{W,F,\mathfrak{s}}$, for a link cobordism (W, F), and prove invariance from choices involved in the construction.

To define maps associated to a link cobordism (W, F), we take a parametrized Kirby decomposition, $\mathcal{K}$, of (W, F). If we can define maps for each elementary parametrized link cobordism, then we can simply define $F_{W,F,\mathfrak{s}}$ to be the composition of all of them. Note that since only one term of a parametrized Kirby decomposition is allowed to have 2–handles, this agrees with the composition law. In the following subsections, we will combine the maps defined in previous sections for quasi-stabilizations, framed spheres and bands to define maps associated to elementary parametrized Kirby decompositions. Then in Subsection 4.11.4, we prove invariance from the choice of parametrized Kirby decomposition.

4.11.1 Maps induced by elementary parametrized link cobordisms of type (1) and (2)

We now consider elementary parametrized link cobordisms $(W, F, \Phi, \mathbb{S}) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ where $\mathbb{S}$ is a framed 0–sphere, 2–sphere or 1–dimensional link which doesn't intersect L (corresponding to type (2) cobordisms), or the empty set (corresponding to type (1) cobordisms). In both cases, the set of divides is a collection of arcs from L_1 to L_2 and the underlying surface of F is a disjoint union of annuli.

Let $\mathcal{A}_0$ denote the pull back of the divides $\mathcal{A}$ to $L_1 \times [0, 1]$, under the diffeomorphism Φ. Now $\mathcal{A}_0$ is clearly isotopic relative $L_1 \times \{0, 1\}$ to a set of divides, $\mathcal{A}'_0$, which are transverse to each $L_1 \times \{t\}$. The divides $\mathcal{A}'_0$ uniquely specify an isotopy of the basepoints $\mathbf{w}_1$ and $\mathbf{z}_1$, along L_1. We then let $\psi_t : (Y_1, L_1) \to (Y_1, L_1)$ denote an extension of this isotopy to all of

Y_1. Setting $\psi = \psi_1$, we define the map

$$F_{W,F,\Phi,\mathbb{S},\mathfrak{s}} = (\Phi|_{Y_1(\mathbb{S})})_* F_{Y_1,\mathbb{S},\Phi^{-1}(\mathfrak{s})}\psi_*.$$

We need to show this is invariant under the choice of isotopic $\mathcal{A}_0'$, as well as invariant under isotopies of Φ, relative Y_1. Although straightforward in our current setting, we prove the following, since we will have to prove a less obvious analog for surfaces with bands later:

Definition 4.11.1. If $\mathcal{A}_t$ is a 1–parameter family of dividing sets on $L \times [0,1]$, which are fixed on $L \times \{0,1\}$, we say that $\mathcal{A}_t$ is an **admissible isotopy** if every arc going from $L \times \{0\}$ to $L \times \{1\}$ is transverse to the sets $L \times \{s\}$ for all s.

Lemma 4.11.2. *Suppose $(L \times [0,1], \mathcal{A})$ is a cylinder with dividing set, which is a trivial cylinder (i.e. all arcs go from $L \times \{0\}$ to $L \times \{1\}$) or is of type I^+, I^-, II^+ or II^- (there is a single arc which doesn't go from $L \times \{0\}$ to $L \times \{1\}$). Suppose $\mathcal{A}_1$ and $\mathcal{A}_2$ are isotopic to $\mathcal{A}$ through isotopies fixing $L \times \{0,1\}$, and suppose that all of the arcs of $\mathcal{A}_1$ and $\mathcal{A}_2$ which go from $L \times \{0\}$ to $L \times \{1\}$ are transverse to the sets $L \times \{s\}$ for all s. Then $\mathcal{A}_1$ and $\mathcal{A}_2$ are isotopic relative $L \times \{0,1\}$ through a path of dividing sets $\mathcal{A}_t$ such that for all t, all arcs of $\mathcal{A}_t$ which go from $L \times \{0\}$ to $L \times \{1\}$ are transverse to the level sets $L \times \{s\}$.*

Proof. For notational simplicity, assume that L has only one component; the argument easily adapts to the case of multiple components. Suppose $A \in \mathcal{A}$ is any of the arcs going from $L \times \{0\}$ to $L \times \{1\}$ (the set of such arcs is nonempty, since by assumption the links on each ends have basepoints on each end). Given two dividing sets $\mathcal{A}_1$ and $\mathcal{A}_2$ on $L \times [0,1]$, which are both of type I^+, I^-, II^+, II^-, and have the same endpoints, and such that two endpoints of arcs are connected by an arc in $\mathcal{A}_1$ iff they are connected by an arc in $\mathcal{A}_2$, then it is easy to that they are admissibly isotopic iff they share the same algebraic intersection number $\#(A \cap (\{p\} \times [0,1]))$ for any $p \in L$ which is not the endpoint of an arc on either end. The algebraic intersection number is unchanged by arbitrary isotopies which fix $L \times \{0,1\}$, and hence two dividing sets are isotopic relative $L \times \{0,1\}$ iff they are admissibly isotopic relative $L \times \{0,1\}$. □

Lemma 4.11.3. *Isotopies of the parameterizing diffeomorphism Φ through diffeomorphisms fixing $Y_1 \times \{0\}$ do not affect $F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}$ for type* (1) *and* (2) *elementary cobordisms.*

Proof. Isotopies of Φ which fix Σ have no effect. Similarly isotopies of Φ which are supported only in a neighborhood of Σ, but which fix $L_1 \times \{1\}$ only affect the composition by changing the pullback of the dividing set, and hence possibly changing the map ψ. By Lemma 4.11.2, such an isotopy of the dividing set can be taken to be an admissible isotopy, which clearly has no effect (since the resulting ψ_* is unchanged). It remains to just check the effect of isotopies supported in a neighborhood of $L_1 \times \{1\}$. If τ denotes a diffeomorphism which twists a neighborhood of $L_1 \times \{1\}$, then precomposing Φ with τ changes $(\Phi|_{Y_1\times\{1\}})_*$ to $(\Phi|_{Y_1\times\{1\}})_*(\tau_*)$ but also changes ψ_* to $(\tau^{-1})_*(\psi_*)$ (since the pullback under Φ of the dividing set changes). Since $(\tau^{-1})_*$ commutes with $F_{Y_1,\mathbb{S},\Phi^{-1}(\mathfrak{s})}$, as τ doesn't move the framed sphere $\mathbb{S}$, the factors of $(\tau)_*$ and $(\tau)_*^{-1}$ cancel and the composition is unchanged. □

4.11.2 Maps induced by elementary parametrized link cobordisms of type (3)

Suppose $(W, F, \Phi, \mathbb{S}) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a parametrized link cobordism of type (3). Define $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$, a collection of divides on $L_1 \times [0, 1]$.

Recall that there are four subtypes of elementary cobordisms of type (3), which we denote by I^+, I^-, II^+, and II^-. These are shown in Figure 4.28.

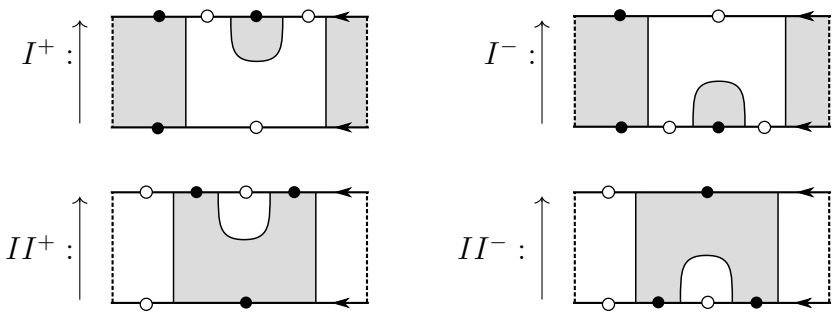

Figure 4.28: **The four subtypes of elementary cobordisms of type** (3)**.**

4.11.2.1 Type I^+ elementary cobordisms

Suppose $(W, F, \Phi, \varnothing) : (Y_1, (L_1, \mathbf{w}_1, \mathbf{z}_1)) \to (Y_2, (L_2, \mathbf{w}_2, \mathbf{z}_2))$ is a parametrized cylindrical cobordism of type $(3I^+)$.

There is a $\mathbf{z}$–region which has three basepoints on its boundary. Let $z_0 \in L_1$ be the basepoint on the incoming end, and let z and z' be the two basepoints on the outgoing end. Let w be the basepoint between z and z', and suppose the basepoints appear with in the following order: z, w then z'. One then isotopes the divides $\mathcal{A}_0$, relative $L_1 \times \{0,1\}$ to get isotopic divides $\mathcal{A}_0'$ such that all arcs from $L_1 \times \{0\}$ to $L_1 \times \{1\}$ are transverse to the sets $L_1 \times \{t\}$. The divides $\mathcal{A}_0'$ determine two diffeomorphisms, ψ and ψ', of Y, which are induced by twisting the link L_1 according to the divides $\mathcal{A}_0'$. The diffeomorphisms ψ and ψ' satisfy

$$\psi(z_0) = z', \qquad \text{and} \qquad \psi'(z_0) = z.$$

By Lemma 4.11.2, the diffeomorphisms ψ and ψ' are independent (up to isotopies fixing the basepoints) of the choice of $\mathcal{A}_0'$ which is isotopic to $\mathcal{A}_0$.

We now define two maps. The first is

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* S^+_{w,z} \psi_*$$

and the second is

$$F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* S^+_{z',w} \psi'_*.$$

Note that there really is no canonical reason to choose one map over the other, since we have to choose which of z and z' we want to identify z_0 with. However an application of Lemma 4.3.23 shows that the two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

4.11.2.2 Type I^- elementary cobordisms

Suppose $(W, F, \Phi, \varnothing)$ is of type I^-. As with elementary cylindrical cobordisms of type I^-, the divides $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$ determine two well defined diffeomorphisms ψ and ψ' of Y, each

coming from one of the two maximal identification of a subset of the basepoints in $\mathbf{w}_1 \cup \mathbf{z}_1$ with $\Phi^{-1}(\mathbf{w}_2 \cup \mathbf{z}_2)$ in such a way which respects the divides. Suppose that z and z' are the basepoints on $(L_1, \mathbf{w}_1, \mathbf{z}_1)$ which are in the same region of $\Sigma \setminus \mathcal{A}$, and let w be the basepoint between them. There is one basepoint, $z_0 \in \Phi^{-1}(\mathbf{w}_2 \cup \mathbf{z}_2)$, in the same region on as z and z' on $(L_1 \times [0,1]) \setminus \mathcal{A}_0$, and assume that the ordering is z, w then z'. Assume that

$$\psi(z') = z_0 \qquad \text{and} \qquad \psi'(z) = z_0.$$

We then define two maps:

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* S^-_{\psi(w),\psi(z)} \psi_*, \qquad \text{and} \qquad F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* S^-_{\psi'(z'),\psi'(w)} \psi'_*.$$

By Lemma 4.3.24, these two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

4.11.2.3 Type II^+ elementary cobordisms

Suppose $(W, F, \Phi, \varnothing)$ is of type II^+. Let $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$. In the region of $(L_1 \times [0,1]) \setminus \mathcal{A}_0$ which has three $\mathbf{w}$–basepoints, let w_0 be the basepoint on $L_1 \times \{0\}$, and let w and w' be the basepoints on $L_1 \times \{1\}$, in this region. Let z denote the basepoint between w and w'. Assume that on $L_1 \times \{0\}$, we have the following ordering of basepoints: w', z then w. Using Lemma 4.11.2, the divides $\mathcal{A}_0$ determine two diffeomorphisms, ψ and ψ', of Y corresponding to twisting a link component as prescribed by the divides. Assume that

$$\psi(w_0) = w' \qquad \text{and} \qquad \psi'(w_0) = w.$$

We define two maps:

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* T^+_{w,z} \psi_*, \qquad \text{and} \qquad F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* T^+_{z,w'} \psi'_*.$$

By Lemma 4.3.25, these two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

4.11.2.4 Type II^- elementary cobordisms

Finally, suppose $(W, F, \Phi, \varnothing)$ is of type II^-. Let $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$. In the region of $(L_1 \times [0,1]) \setminus \mathcal{A}_0$ which has three $\mathbf{w}$–basepoints, let w_0 be the basepoint on $L_1 \times \{1\}$, and let w and w' be the basepoints on $L_1 \times \{0\}$, in this region. Let z denote the basepoint between w and w', and suppose that the basepoints appear in the order w', z then w. Using Lemma 4.11.2, the divides $\mathcal{A}_0$ determine two diffeomorphisms, ψ and ψ', of Y corresponding to twisting a link component as prescribed by the divides. Assume that

$$\psi(w') = w_0 \qquad \text{and} \qquad \psi'(w) = w_0.$$

We define two maps:

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* T^-_{\psi(w),\psi(z)} \psi_* \qquad \text{and} \qquad F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* T^-_{\psi'(z),\psi'(w')} \psi'_*.$$

By Lemma 4.3.26, these two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

Since the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ commute with maps induced by diffeomorphisms, adapting Lemma 4.11.3 immediately yields the following:

Lemma 4.11.4. *Isotopies of the parameterizing diffeomorphism Φ through diffeomorphisms fixing $Y_1 \times \{0\}$, do not affect the maps $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$ for elementary parametrized cobordisms of type* (3).

Proof. This follows identically to Lemma 4.11.3, using Lemma 4.11.2. □

4.11.3 Maps induced by elementary parametrized link cobordisms of type (4)

Suppose that $(W, F, \Phi, \mathbb{S})$ is an elementary cobordism of type (4) from $(Y_1, \mathbb{L}_1)$ to $(Y_2, \mathbb{L}_2)$. Here

$$\Phi : W(Y_1, L_1, \mathbb{S}) \to (W, F)$$

is the parameterizing diffeomorphism, such that $\Phi^{-1}(\mathcal{A})$ doesn't intersect B, and (it could however, wind around many times on the subsets of $\Sigma(\mathbb{S})$ identified with $L_1 \times [0,1]$ or $L_1(\mathbb{S}) \times [1,2]$) and the set of divides $\mathcal{A}$ consists entirely of arcs going from the incoming boundary to the outgoing boundary.

Recall the model surface $\Sigma(\mathbb{S}) = (L_1 \times [0,1]) \cup B \cup (L_1(\mathbb{S}) \times [1,2])$ inside of $W(Y_1, \mathbb{S})$ is a diffeomorphic to a disjoint union of a pair of pants together with some cylinders. There is a subset, B, identified with $[-1,1] \times [-1,1]$. Define the set

$$\mathcal{L} = (L \times \{1\}) \cap (L(\mathbb{S}) \times \{1\}),$$

which can be thought of as the level set, outside of the band, of a Morse function with a single index one critical point on the band.

Definition 4.11.5. We will call an isotopy $\mathcal{A}_t$ of divides on $\Sigma(\mathbb{S})$, which is fixed on $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$, an **admissible isotopy** if each arc of $\mathcal{A}_t$ intersects $\mathcal{L}$ at a exactly one point, transversely, and no arcs ever intersect B.

Lemma 4.11.6. *Suppose that $\mathcal{A}_1$ and $\mathcal{A}_2$ are collections of divides on $\Sigma(\mathbb{S}) = (L \times [0,1]) \cup B \cup (L(\mathbb{S}) \times [1,2])$ and that each arc of both divides goes from $L \times \{0\}$ to $L(\mathbb{S}) \times \{2\}$. Suppose further that each arc of $\mathcal{A}_1$ and $\mathcal{A}_2$ intersects $\mathcal{L}$ exactly once, transversely, and that no arcs of either $\mathcal{A}_i$ intersect B. If $\mathcal{A}_1$ and $\mathcal{A}_2$ are isotopic through divides $\mathcal{A}_t$ which do not intersect the point $(0,0) \in B$ and are fixed on $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$, then $\mathcal{A}_1$ and $\mathcal{A}_2$ are admissibly isotopic.*

Proof. On any component of $\Sigma(\mathbb{S})$ which doesn't intersect B, the claim is obvious, so suppose without loss of generality that L is a one (resp. two) component link which is separated (resp. joined) by the band B. Suppose for definiteness, that L is a two component link which is joined by B. The case that L is a one component link which is separated by the band follows from nearly identical reasoning.

Let $\{p_-, p_+\} \subseteq L$ denote the underlying 0–sphere of $\mathbb{S}$. Let C_- and C_+ be curves on $(L \times [0,1]) \cup B$, each consisting of a path on $L \times [0,1]$ from a point on $L \times \{0\}$ which is

not an endpoint of $\mathcal{A}_i$, to $(p_-, 1)$ or $(p_+, 1)$, respectively, concatenated with one of the arcs $[-1, 0] \times \{0\}$ or $[0, 1] \times \{0\}$, respectively, on B. These are shown in Figure 4.29.

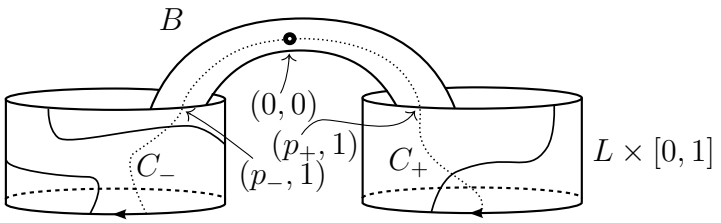

Figure 4.29: **The arcs C_- and C_+ on** $\Sigma(\mathbb{S})$, as well as the band B and the point $(0, 0) \in B$.

Similarly let C'_+ and C'_- be two similarly defined curves on $B \cup (L(\mathbb{S}) \times [1, 2])$ from $(0, 0)$ to two points on $L(\mathbb{S}) \times \{2\}$ which are not endpoints of any of the arcs of $\mathcal{A}_i$.

The admissible isotopy class of the divides $\mathcal{A}_1$ and $\mathcal{A}_2$ relative $L \times \{0\} \cup L(\mathbb{S}) \times \{2\}$ is determined by the ordering that the arcs appear on $\mathcal{L}$, and the intersection numbers of arcs with C_-, C_+, C'_-, and C'_+.

The configuration of endpoints of arcs on $L \times \{1\}$ is determined by the region of $\Sigma \setminus \mathcal{A}$ that the band is in and the order that the arcs appear on $L \times \{0\}$. This is unchanged by isotopies of $\Sigma(\mathbb{S})$ which may not be admissible, but which don't map any arcs to $(0, 0)$. Similarly the algebraic intersection numbers mentioned above are not changed by isotopies of the arcs which fix the ends of the arcs on $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$, and don't map any arcs to $(0, 0)$. Hence if two collections of divides, $\mathcal{A}_1$ and $\mathcal{A}_2$, each consisting of arcs going from $L \times \{0\}$ to $L(\mathbb{S}) \times \{2\}$ and sharing the same end points, are isotopic relative $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$ through isotopies not sending an arc to $(0, 0)$, then they are admissibly isotopic.

□

We now arbitrarily pick basepoints $\mathbf{w}_{1.5} \cup \mathbf{z}_{1.5}$ on the $L_1 \times \{1\}$ which make $(L_1 \times \{1\}, \mathbf{w}_{1.5}, \mathbf{z}_{1.5})$ an acceptably based link, and which make the band B either an α–band or a β–band. Note that the type of the band as either type–$\mathbf{w}$ or type–$\mathbf{z}$ is well defined by the configuration of the divides. For a choice of $\mathbf{w}_{1.5}$ and $\mathbf{z}_{1.5}$ there is no ambiguity

in whether B is an α–band or a β–band, but up to isotopy which fixes B, there are two choices of basepoints $\mathbf{w}_{1.5}$ and $\mathbf{z}_{1.5}$ which respect the divides (these two choices are isotopic through isotopies which don't fix B). Let $\mathbf{w}^\tau_{1.5}$ and $\mathbf{z}^\tau_{1.5}$ denote these basepoints, for a choice of $\tau \in \{\alpha, \beta\}$.

Adapting Lemma 4.11.2 in the obvious way, the divides $\Phi^{-1}(\mathcal{A})$ on $L_1 \times [0,1]$ induce an isotopy ψ^τ_1 of (Y, L) which sends $(\mathbf{w}_1, \mathbf{z}_1)$ to $(\mathbf{w}^\tau_{1.5}, \mathbf{z}^\tau_{1.5})$. Similarly the divides $\Phi^{-1}(\mathcal{A})$ induce an isotopy ψ^τ_2 which sends $(\mathbf{w}^\tau_{1.5}, \mathbf{z}^\tau_{1.5})$ to $(\Phi^{-1}(\mathbf{w}_2), \Phi^{-1}(\mathbf{z}_2))$. We now define, with respect to these choices

$$F^\tau_{W,F,\Phi,\mathbb{S},\mathfrak{s}} = \Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi^\tau_2)_* F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\Phi^{-1}(\mathfrak{s})}(\psi^\tau_1)_*.$$

There is no ambiguity in the choice of $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, though τ can be either α or β, as determined by the choice of basepoints $\mathbf{w}^\tau_{1.5}$ and $\mathbf{z}^\tau_{1.5}$. By Propositions 4.7.8 and 4.7.11 we know that the α– and β–band maps are related by the diffeomorphism resulting from moving basepoints, which is tautologically encoded into ψ^τ_1 and ψ^τ_2. From these considerations we conclude that

$$F^\alpha_{W,F,\Phi,\mathbb{S},\mathfrak{s}} \simeq F^\beta_{W,F,\Phi,\mathbb{S},\mathfrak{s}},$$

and we denote the common map

$$F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}.$$

We have the following:

Lemma 4.11.7. *For an elementary parametrized cobordism $(W, F, \Phi, \mathbb{S})$ of type (4), isotopies of the parametrization Φ which fix $Y_1 \times \{0\}$ and don't map the point $(0,0) \in B$ to $\mathcal{A}$ do not affect the map $F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}$.*

Proof. Isotopies which are fixed on $\Sigma(\mathbb{S})$ clearly have no effect, so we can just consider isotopies of $\Phi|_{\Sigma(\mathbb{S})}$, extended to isotopies of all of W. By the same argument as in Lemma 4.11.3, the maps are independent of isotopies of Φ which are supported in a neighborhood of $L_1(\mathbb{S}) \times \{2\}$. Hence it remains to see that the maps are invariant under isotopies which are fixed on $Y_1 \times \{0\}$ and $Y_1(\mathbb{S}) \times \{2\}$. Since Φ_t is fixed on $Y(\mathbb{S})$, the map $F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}$ only

depends on Φ through the dividing set $\Phi^{-1}(\mathcal{A})$. If Φ_t is such an isotopy, we consider the dividing sets $\mathcal{A}_t = \Phi_t^{-1}(\mathcal{A})$ on $\Sigma(\mathbb{S})$. By Lemma 4.11.6, we can assume that $\mathcal{A}_0$ and $\mathcal{A}_1$ are admissibly isotopic, i.e. the dividing sets $\mathcal{A}_t$ never pass through the band B, and the curves $\mathcal{A}_t$ remain transverse to $\mathcal{L}$ throughout.

Such an isotopy of $\mathcal{A}_t$ induces an isotopy of both (Y_1, L_1) and $(Y_1(\mathbb{S}), L_1(\mathbb{S}))$ which fixes $\mathbb{S} \subseteq Y$ and $D^1 \times S^2 \subseteq Y(\mathbb{S})$, respectively, resulting from moving the basepoints according to the intersection of $\mathcal{A}_t$ with $\mathcal{L}$. We let τ denote this isotopy on Y_1, and let $\tau^{\mathbb{S}_1}$ denote this isotopy on $Y_1(\mathbb{S})$. The effect of the isotopy of $\mathcal{A}$ is to replace ψ_1 with $\psi_1' = \tau \circ \psi_1$ and to replace ψ_2 with $\psi_2' = \psi_2 \circ (\tau^{\mathbb{S}_i})^{-1}$. In general, the band maps $F_B^{\mathbf{o}}$ are not invariant under diffeomorphisms which don't fix the band. Instead we canonically have $\phi_* F_B = F_{\phi_* B} \phi_*$ for a diffeomorphism ϕ. Since τ fixes $\mathbb{S} \subseteq Y$, and $\tau^{\mathbb{S}}$ fixes $D^1 \times S^2 \subseteq Y(\mathbb{S})$, we have in particular that $\tau^{\mathbb{S}}$ fixes the band $B \subseteq Y_1(\mathbb{S})$. Hence we have

$$\Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2')_* F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\mathfrak{s}}(\psi_1')_* \simeq \Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2)_*(\tau^{\mathbb{S}_i})_*^{-1} F^{\mathbf{o}}_{Y,\mathbb{L},\mathbb{S},\mathfrak{s}} \tau_*(\psi_1)_*$$

$$\simeq \Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2)_* F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\mathfrak{s}}(\tau)_*^{-1} \tau_*(\psi_1)_* \simeq \Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2)_* F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\mathfrak{s}}(\psi_1)_*,$$

showing that the maps are unchanged. □

4.11.4 Definition and invariance of the cobordism maps

If (W, F) is a link cobordism such that each component of W has a nonempty incoming and outgoing end, and each component of Σ has a nonempty incoming and outgoing end, we take a parametrized Kirby decomposition, $\mathcal{K}$, with elementary parametrized cobordisms $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$, and define the maps

$$F_{W,F,\mathfrak{s},\mathcal{K}} = \prod_i F_{W_i,F_i,\Phi_i,\mathbb{S}_i,\mathfrak{s}|_{W_i}},$$

where the maps $F_{W_i,F_i,\Phi_i,\mathbb{S}_i,\mathfrak{s}|_{W_i}}$ are the maps defined in the previous subsections.

We now prove a slightly weaker version of Theorem A, for cobordisms which have nonempty incoming and outgoing ends of each component of W and Σ. We will extend this to the most general version of Theorem A in Section 4.12.

Theorem A′. *The maps $F_{W,F,\mathfrak{s},\mathcal{K}}$ are independent of the parametrized Kirby decomposition, $\mathcal{K}$, up to filtered, equivariant chain homotopy, and hence define invariants of the decorated link cobordisms with* Spin^c *structure,* $(W, F, \mathfrak{s})$.

Proof. By Proposition 4.10.3, it is sufficient to show that the maps $F_{W,F,\mathfrak{s},\mathcal{K}}$ are equal for weakly equivalent parametrized Kirby decompositions. Hence it is sufficient to check invariance from the moves of Definition 4.10.2.

Move 1: We consider strong isotopies (Definition 4.10.1). First note that if

$$\phi : (W, \Sigma) \to (W, \Sigma)$$

is any diffeomorphism with $\phi|_{\partial W} = \mathrm{id}$, and $\mathcal{K}$ is a parametrized Kirby decomposition for (W, F), then using Equation (4.10) we get a decomposition $\phi_*\mathcal{K}$ of (W, F) into parametrized cobordisms, though in general, as we remarked after Definition 4.10.1, the decomposition may fail to be a parametrized Kirby decomposition for (W, F). Instead, tautologically, $\phi_*\mathcal{K}$ yields a parametrized Kirby decomposition for (W, ϕ_*F), where $\phi_*F = (\Sigma, \phi_*\mathcal{A})$. Indeed, we tautologically have that

$$F_{W,F,\mathfrak{s},\mathcal{K}} \simeq F_{W,\phi_*F,\phi_*\mathfrak{s},\phi_*\mathcal{K}},$$

or more explicitly

$$F_{W,(\Sigma,\mathcal{A}),\mathfrak{s},\mathcal{K}} \simeq F_{W,(\Sigma,\phi_*\mathcal{A}),\phi_*\mathfrak{s},\phi_*\mathcal{K}}.$$

In the case where $\phi = \phi_1$ for a strong isotopy ϕ_t, we note that our desired equality now becomes

$$F_{W,(\Sigma,(\phi_1)_*\mathcal{A}),(\phi_1)_*\mathfrak{s},(\phi_1)_*\mathcal{K}} \simeq F_{W,(\Sigma,\mathcal{A},)(\phi_1)_*\mathfrak{s},(\phi_1)_*\mathcal{K}}.$$

Hence, to show invariance under strong isotopies, it is sufficient to show that isotopies of the arcs in $\mathcal{A}$, which are fixed on ∂W, which additionally preserve the property that each term in the decomposition is an elementary parametrized cobordism, do not affect the maps. Such an isotopy can be arranged into a sequence of isotopies which are supported in a neighborhood of an individual Y_i, and isotopies which are fixed on all of the Y_i.

For isotopies of $\mathcal{A}$ which are fixed on all of the Y_i, invariance follows from Lemmas 4.11.3, 4.11.4, and 4.11.7. Hence it remains only to show invariance under isotopies of $\mathcal{A}$ which are supported in a neighborhood of a fixed Y_i which is in the interior of W. For such isotopies, let $\mathcal{W}_{i-1} = (W_{i-1}, F_{i-1}, \Phi_{i-1}, \mathbb{S}_{i-1})$ and $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ denote the elementary parametrized cobordisms adjacent to Y_i for $\mathcal{K}$. Let $\tau : W \to W$ be the result of taking an isotopy τ_t of L_i, and extending it to an isotopy of Y_i, then extending that to an isotopy of all of W which preserves Σ, and setting $\tau = \tau_1$. It is straightforward to check that for each type of elementary parametrized cobordism that

$$F_{W_{i-1},(\Sigma_{i-1},\tau(\mathcal{A})\cap\Sigma_{i-1}),\Phi_{i-1},\mathfrak{s}|_{W_{i-1}}} \simeq \tau_* \circ F_{W_{i-1},(\Sigma_{i-1},\mathcal{A}\cap\Sigma_{i-1}),\Phi_{i-1},\mathfrak{s}|_{W_{i-1}}}$$

and

$$F_{W_i,(\Sigma_i,\tau(\mathcal{A})\cap\Sigma_i),\Phi_i,\mathfrak{s}_i} \simeq F_{W_i,(\Sigma_i,\mathcal{A}_i),\Phi_i,\mathfrak{s}_i} \circ (\tau^{-1})_*.$$

When the two maps are composed, the effects of τ cancel and the composition is unchanged.

Move 2: We now consider isotopies of a framed sphere, away from L_i. We will omit the subscript i, and the Spin^c structures, for notational simplicity.

Suppose that $\mathbb{S}^t$ is an isotopy of framed spheres away from L, starting at $\mathbb{S}$ and ending at $\mathbb{S}'$. Let d be the (isotopically unique) diffeomorphism of Y extending the canonical diffeomorphism sending $\mathbb{S}$ to $\mathbb{S}'$ along this isotopy. Let $d^{\mathbb{S}} : Y(\mathbb{S}) \to Y(\mathbb{S}')$ denote the induced map. Let ψ denote the diffeomorphism of (Y, L) obtained by pulling back the divides under Φ to the model cobordism $W(Y, L, \mathbb{S})$, and defining ψ to be the map induced by twisting along L, according to the divides. Notice that Φ and Φ' can be taken to be equal along $L \times [0, 1] \subseteq W(Y, L, \mathbb{S})$, since the isotopy of framed spheres occurs away from L. By definition we have that

$$F_{W,\mathbb{L},\mathbb{S},\Phi} = (\Phi|_{Y(\mathbb{S})})_* F_{Y,\mathbb{L},\mathbb{S}} \psi_*$$

and

$$F_{W,\mathbb{L},\mathbb{S}',\Phi'} = (\Phi'|_{Y(\mathbb{S}')})_* F_{Y,\mathbb{L},\mathbb{S}'} \psi_*.$$

By definition of the move, we have that

$$(\Phi'|_{Y(\mathbb{S}')})_* d^{\mathbb{S}}_* = (\Phi|_{Y(\mathbb{S})})_*.$$

Similarly, by tautological we have

$$F_{Y,\mathbb{L},\mathbb{S}'}d_* = d_*^{\mathbb{S}}F_{Y,\mathbb{L},\mathbb{S}}.$$

We note that ψ_* and d_* commute, since they are induced by isotopies supported on disjoint subsets of Y. As a consequence we have that

$$\begin{aligned}
F_{W,\mathbb{L},\mathbb{S}',\Phi'}d_* &= (\Phi'|_{Y(\mathbb{S}')})_*F_{Y,\mathbb{L},\mathbb{S}'}\psi_*d_* \\
&= (\Phi'|_{Y(\mathbb{S}')})_*F_{Y,\mathbb{L},\mathbb{S}'}d_*\psi_* \\
&= (\Phi'|_{Y(\mathbb{S}')})_*d_*^{\mathbb{S}}F_{Y,\mathbb{L},\mathbb{S}}\psi_* \\
&= (\Phi|_{Y(\mathbb{S})})_*F_{Y,\mathbb{L},\mathbb{S}}\psi_* \\
&= F_{W,\mathbb{L},\mathbb{S},\Phi}.
\end{aligned}$$

Now d is a diffeomorphism of Y which is isotopic to the identity relative L, and hence we conclude that $d_* \simeq 1$. Hence

$$F_{W,\mathbb{L},\mathbb{S}',\Phi'} \simeq F_{W,\mathbb{L},\mathbb{S},\Phi},$$

showing the maps to be invariant under this move.

Move 3: We now prove invariance of isotopies of a framed 0–sphere along a link L_i. This follows similarly to Move (2), and is essentially a tautology. Suppose that $(W, F, \mathbb{S}, \Phi)$ is an elementary parametrized link cobordism of type (4). The type, $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, of the elementary parametrized cobordism is determined uniquely by the type of the region of F containing the image of $(0,0) \in B$, but by construction the image of the point $(0,0)$ in the model band does not change. For notational simplicity we omit the i subscript. Let d_t be an isotopy of $\mathbb{S}$ along L, define $d = d_1$ and let $\mathbb{S}'$ be the image of $\mathbb{S}$ and let Φ' be the parametrization described in Move (3). Let ψ_1 and ψ_2 be the diffeomorphisms in the definition of the maps (cf. Subsection 4.11.3) associated to Φ, and let ψ_1' and ψ_2' be the diffeomorphisms associated to the parametrization Φ'. By definition

$$F_{W,F,\mathbb{S},\Phi} = (\Phi|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2)_*F^{\mathbf{o}}_{Y,\psi_1(\mathbb{L}),\mathbb{S}}(\psi_1)_*$$

and

$$F_{W,F,\mathbb{S}',\Phi'} = (\Phi'|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2')_*F^{\mathbf{o}}_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'}(\psi_1')_*.$$

It is easy to see that by construction we have that

$$(\psi_1')_* = d_*(\psi_1)_*.$$

Tautologically we have

$$(d^{\mathbb{S}})_*^{-1} F_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'} d_* = F_{Y,\psi_1(\mathbb{L}),\mathbb{S}}.$$

Now also by construction, we have

$$\psi_2' = (d^{\mathbb{S}})_*(\psi_2)_*(d^{\mathbb{S}})_*^{-1}.$$

Finally we note that, by construction

$$(\Phi|_{Y(\mathbb{S})\times\{2\}})_* = (\Phi'|_{Y(\mathbb{S}')\times\{2\}})_* d_*^{\mathbb{S}}.$$

As a consequence, we see that

$$\begin{aligned} F_{W,F,\mathbb{S}',\Phi'} &= (\Phi'|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2')_* F^{\mathbf{o}}_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'}(\psi_1')_* \\ &= (\Phi'|_{Y(\mathbb{S})\times\{2\}})_*(d^{\mathbb{S}})_*(\psi_2)_*(d^{\mathbb{S}})_*^{-1} F^{\mathbf{o}}_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'} d_*(\psi_1)_* \\ &= (\Phi|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2)_* F^{\mathbf{o}}_{Y,\psi_1(\mathbb{L}),\mathbb{S}}(\psi_1)_* \\ &= F_{W,F,\mathbb{S},\Phi}, \end{aligned}$$

completing the proof of invariance under Move (3).

Move 4: Invariance under adding or removing a trivial cylinder is essentially a tautology. The map induced by a trivial cylinder is a diffeomorphism map, and all of the other maps involve pre- or post-composition with an analogous diffeomorphism. It is easy to check that combining a trivial cylinder with another elementary cobordism (from the left or the right) does not affect the maps.

Move 5: We now consider invariance under changing the orientation of the framed sphere $\mathbb{S}$. We argue in the case of a framed sphere $\mathbb{S}$, attached away from L. The case that $\mathbb{S}$ is attached along L follows from similar considerations. Let $\sigma : \mathbb{R}^4 \to \mathbb{R}^4$ denote the map $\sigma(x, y, w, z) = (-x, y, w, -z)$. If $\iota : S^k \times D^{3-k} \to Y$ is the embedding of $\mathbb{S}$, let $\overline{\mathbb{S}}$ denote the

framed sphere with embedding $\iota \circ \sigma$. Note that σ induces a map $\sigma : Y(\mathbb{S}) \to Y(\overline{\mathbb{S}})$, as well as $\sigma : W(Y, L, \mathbb{S}) \to W(Y, L, \overline{\mathbb{S}})$. Indeed one easily sees that

$$\sigma \circ F_{Y,\mathbb{S},\mathfrak{s}} = F_{Y,\overline{\mathbb{S}},\sigma_*\mathfrak{s}},$$

which follows from the fact that the maps don't depend on a choice of orientation for the framed spheres $\mathbb{S}$. Since the parametrizations Φ and Φ' also differ by composition with σ, the composed maps are unchanged.

Move 6: Invariance under index 1/2 critical point birth-deaths of $f|_{W\setminus\Sigma}$ follows from the same triangle map computation used to prove the analogous result about the closed 3-manifold invariants (Theorem 2.8.7). We remark that by construction, the handles are attached by framed spheres which are in the complement of the link, allowing the same construction to be used. We note that the above count of holomorphic triangles is an adaptation of the original argument by Ozsváth and Szabó in [OS06, Lemma 4.16] to handle our slightly different definition of the 1–handle maps than in the original papers.

Move 7: Invariance under index 2/3 critical point birth-deaths of $f|_{W\setminus\Sigma}$ follows from Theorem 2.8.11. Note that although generically a framed 2-sphere in Y could intersect the link at a collection of points, by our construction, any framed 2-sphere appearing in a parametrized Kirby decomposition does not intersect the link L, which allows us to use the same model computation as in the case of Theorem 2.8.11. Again note that an analogous count of holomorphic triangles was proven by Ozsváth and Szabó in [OS06, Lemma 4.17].

Move 8: Invariance under critical value swaps between two index 1 or two index 3 critical points of $f|_{W\setminus\Sigma}$ follows from the commutation of 1– and 3–handle maps amongst each other, proven in Lemma 4.4.4.

Move 9: Invariance from handleslides amongst the components of a framed 1–dimensional link follows identically as it did in [OS06, Lemma 4.14]. In summary, a handleslide of components of a framed link is reflected on a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ subordinate to an α–bouquet for the framed link by handlesliding the corresponding $\boldsymbol{\alpha}$–curves across each other. Invariance from such handleslide follows using the associativity relations, and a model

computation (performed in [OS06]).

Move 10: Invariance from a birth-death singularity of $f|_{\mathcal{A}}$ follows from the identities $S^-_{w,z}T^+_{w,z} \simeq T^-_{w,z}S^+_{w,z} \simeq \mathrm{id}$ proven in Lemma 4.3.13. This is demonstrated in Figure 4.30. If instead we switched the orientation of the links with respect to this picture, independence would follow by replacing $S^{\pm}_{w,z}$ with $S^{\pm}_{z,w}$ and $T^{\pm}_{w,z}$ with $T^{\pm}_{z,w}$.

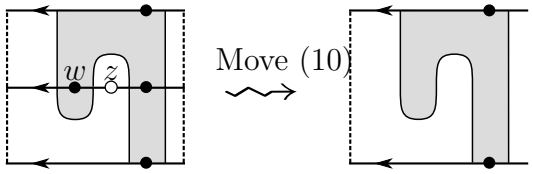

Figure 4.30: **A picture of Move 10.** The map induced on the left is $T^-_{w,z}S^+_{w,z}$ if the cobordism is going from bottom to top, and $S^-_{w,z}T^+_{w,z}$, if the cobordism is going from top to bottom. The map on the right is the identity.

Move 11: For a critical value switch of critical points of $f|_{\mathcal{A}}$, there are 10 possible configurations, depending on whether the two type (3) elementary parametrized cobordisms are of subtype I^+, I^-, II^+, or II^-. The cases where both are of type I^+ or I^- are handled by Proposition 4.3.15. The cases where both are of type II^+ or II^- are covered by Proposition 4.3.16, and the cases that one is of type $I^{\pm}$ and the other is of type $II^{\pm}$ is covered by Proposition 4.3.17.

It is important to note that in this last case, one does not expect $T^{\pm}_{w,z}$ and $S^{\pm}_{w',z'}$ to always commute. Indeed in Proposition 4.3.17, there were requirements on the configuration of basepoints which can appear. It is easy to check that the requirements on the configuration of basepoints is always satisfied when the basepoints arise in critical value swap of critical points of $f|_{\mathcal{A}}$ where one induces a type $II^{\pm}$ map, and the other induces a type $I^{\pm}$ map (see Figure 4.28 for a picture of type $I^{\pm}$ and $II^{\pm}$ decorations for elementary parametrized cobordisms of type (3)).

Move 12: We now consider the effect of switching the order of two index 1 critical points on Σ. The maps associated to critical points along Σ are a composition of the 1–handle maps, as well as the **z**–band or **w**–band maps.

We note that by the triangle map computation of Theorem 2.7.8, the 1–handle maps commute with the band maps. By Lemma 4.4.4 the 1–handle maps commute with each other. Thus it remains just to show that the band maps can be commuted amongst themselves. This now follows from Proposition 4.7.12. Note that if both band maps are $\mathbf{z}$–band maps or both are $\mathbf{w}$–band maps, then the maps always commute, though if one band map is type–$\mathbf{w}$ and the other is type–$\mathbf{z}$, there are assumptions from Proposition 4.7.12 about the configuration of ends of the band. These are automatically satisfied in the case that the configuration of bands came from a surface with divides with a Morse function f such that f has two index one critical points, and $f|_{\mathcal{A}}$ has no critical points.

Move 13: We now consider a critical value switch between critical points of $f|_{\Sigma}$ and $f|_{\mathcal{A}}$. This follows from a similar argument as the last move. Again, the 1–handle maps commute with quasi-stabilization maps by Lemma 4.4.3 and the band maps by the triangle map computation of Theorem 2.7.8. Thus it is sufficient to show that the band maps commute with the maps induced by elementary cobordism maps of type (3).

Commutation between the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ and the band maps $F^{\mathbf{w}}_B$ and $F^{\mathbf{z}}_B$ now follows from Propositions 4.7.1 and 4.7.3. We note that the hypotheses of those two propositions regarding the configuration of basepoints involved are always satisfied when the configuration of basepoints and bands are induced by a surface with divides with a Morse function f which has a single index one critical point, and such that $f|_{\mathcal{A}}$ also has a single index 1 critical point.

Move 14: We now consider move corresponding to a critical point of Σ crossing $\mathcal{A}$. Notice that the 1–handle maps commute with everything, so we can bring them below the quasi-stabilization maps and the band maps, so we can just show the equality of maps related to the quasi-stabilization maps and the band maps. In this move, we have arranged for the crossing to locally be as in Figure 4.27. With basepoints, we can redraw this as in Figure 4.31.

The maps corresponding to the bands involve postcomposition with a diffeomorphism map which fixes the link and moves the basepoints according to the divides. Hence it is sufficient to show that the indicated maps from Figure 4.31 become equal after moving (w, z) to

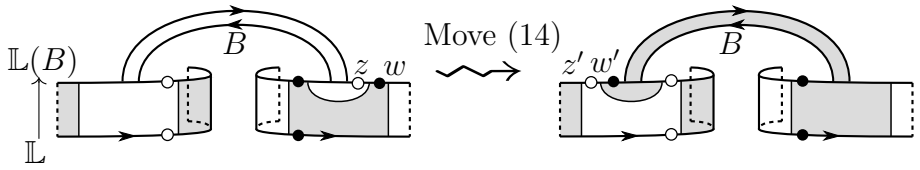

Figure 4.31: **The two configurations we need to show induce the same map, after moving basepoints, for Move 14.**

(w', z'). The map on the right induces

$$F_B^{\mathbf{w}} S^+_{w',z'},$$

while the map on the left induces

$$F_B^{\mathbf{z}} T^+_{w,z}.$$

The diffeomorphism moving (z', w') to (z, w) takes the form $S^-_{w',z'} T^+_{w,z}$ by Lemma 4.3.29. Thus we wish to show

$$S^-_{w',z'} T^+_{w,z} F_B^{\mathbf{w}} S^+_{w',z'} \simeq F_B^{\mathbf{z}} T^+_{w,z}.$$

This will be an algebraic computation. We note first that

$$\Psi_z T^+_{w,z} \simeq 0, \tag{4.13}$$

which follows immediately by direct computation using the quasi-stabilized differential in Proposition 4.3.4 and the definition of $T^+_{w,z}$. We thus compute:

$$\begin{aligned}
S^-_{w',z'} T^+_{w,z} F_B^{\mathbf{w}} S^+_{w',z'} &\simeq S^-_{w',z'} F_B^{\mathbf{w}} T^+_{w,z} S^+_{w',z'} && \text{(Proposition 4.7.3)} \\
&\simeq S^-_{w',z'} (F_B^{\mathbf{z}} \Psi_z + \Psi_z F_B^{\mathbf{z}}) T^+_{w,z} S^+_{w',z'} && \text{(Proposition 4.7.16)} \\
&\simeq S^-_{w',z'} \Psi_z F_B^{\mathbf{z}} T^+_{w,z} S^+_{w',z'} && \text{(Equation (4.13))} \\
&\simeq S^-_{w',z'} \Psi_z F_B^{\mathbf{z}} S^+_{w',z'} T^+_{w,z} && \text{(Proposition 4.3.17)} \\
&\simeq S^-_{w',z'} \Psi_z S^+_{w',z'} F_B^{\mathbf{z}} T^+_{w,z} && \text{(Proposition 4.7.1)} \\
&\simeq F_B^{\mathbf{z}} T^+_{w,z}, && \text{(Remark 4.3.14)}
\end{aligned}$$

completing the proof. □

4.12 Maps for cobordisms where $\pi_0(\mathbb{L}_i) \to \pi_0(\Sigma)$ is not surjective

In this section we consider cobordisms where $\pi_0(\mathbb{L}_0) \to \pi_0(\Sigma)$ is not surjective. As in Chapter 2, to construct these maps, we will need to pre- or post-compose with 0- and 4-handle maps, which we now define. In Chapter 2, we introduced them in the case that we were adding a copy of S^3 with a single new basepoint. We now wish to define them in the case that we are adding a copy of S^3 containing an unknot with two basepoints.

Definition 4.12.1. If $B^4 \subseteq \operatorname{int} W$ is a 4-ball, intersecting Σ along an arc of $\mathcal{A}$, we say B is **standard** if $B \cap \mathcal{A}$ is connected and there are coordinates (x, y, w, z) on a neighborhood of $B \subseteq \operatorname{int} W$ such that $B = \{(x, y, w, z) : x^2 + y^2 + w^2 + z^2 < 1\}$ and in these coordinates $\Sigma \cap B = \{(x, y, 0, 0)\} \cap B$.

In this section, we will implicitly assume all 4-balls B along $\mathcal{A}$ are standard. Note that as we did in Subsection 4.8.2, when we constructed the compound 1-handle/band maps, if we pick a vector in the unit sphere bundle $U((T_{(0,0,0,0)}\Sigma)^{\perp}) \cong S^1$, we can push the surface Σ off of B and into ∂B. As S^1 is connected, any two ways of doing this yield isotopic disks in $\partial B = S^3$, relative the unknot $U = \partial B \cap \Sigma$. From another perspective, if $U \subseteq S^3$ is an unknot then any two disks with boundary on U are isotopic relative U, which can be seen from other considerations. With this in mind, we proceed to define the maps for general cobordisms which might not have enough ends.

Let $\mathbb{U} \subseteq S^3$ denote an unknot with exactly two basepoints. We now define the 0-handle and 4-handle maps for link Floer homology:

$$F_{Y,0,\mathbb{U}} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y \sqcup S^3, \mathbb{L} \sqcup \mathbb{U}, \mathfrak{s} \sqcup \mathfrak{s}_0)$$

and

$$F_{Y,4,\mathbb{U}} : \mathcal{C}FL^\circ(Y \sqcup S^3, \mathbb{L} \sqcup \mathbb{U}, \mathfrak{s} \sqcup \mathfrak{s}_0) \to \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}).$$

These take the obvious form, namely

$$F_{Y,0,\mathbb{U}}(\mathbf{x}) = \mathbf{x} \times c \qquad \text{and} \qquad F_{Y,4,\mathbb{U}}(\mathbf{x} \times c) = \mathbf{x}$$

for c the generator of $\mathcal{C}FL^-(S^3, \mathbb{U}, \mathfrak{s}_0)$.

Remark 4.12.2. Earlier, we defined maps for a "compound 1-handle/band", corresponding to attaching a 4-dimensional 1-handle for W, which contained a 2-dimensional 1-handle for Σ. Analogously, the 0-handle and 4-handle maps correspond to attaching a 4-dimensional 0- or 4-handle for W, which contains a 2-dimensional 0- or 2-handle for Σ.

Suppose one attaches a 0-handle to $Y \times [0,1]$, and then attaches a 1-handle with framed 0-sphere $\mathbb{S}_0 = \{p_1, p_2\} \times D^3$, with one foot at a point p_1 in Y and one foot at a point p_2 in the new copy of S^3. There is a diffeomorphism $\phi : (Y \sqcup S^3)(\mathbb{S}_0) \to Y$, which is the identity outside of a neighborhood of $\{p_1\} \times D^3$. In the case of the closed three manifold invariants, we had that if $w \in S^3$ was a basepoint, then

$$\phi_* \circ F_{Y \sqcup S^3, \mathbb{S}_0, \bar{\mathfrak{s}}} \circ F_{Y,0} = S^+_{\phi(w)}$$

and a similar relation held for the 4-handle maps and S^-_w. As an analog for links, we prove the following:

Lemma 4.12.3. *Suppose that $\mathbb{U}$ is an unknot in S^3 and D is a disk bounding $\mathbb{U}$. If ϕ is the diffeomorphism described in the previous paragraph, and D is a disk in S^3 for $\mathbb{U}$, we have the following:*

$$F_{Y \sqcup S^3, \mathbb{S}_0, \bar{\mathfrak{s}}} \circ F_{Y,0,\mathbb{U}} \simeq S^+_{\phi(\mathbb{U}), \phi(D)}.$$

If $\mathbb{S}_2$ denotes the canonical framed 2-sphere in $(Y \sqcup S^3)(\mathbb{S}_0)$, then

$$F_{Y \sqcup S^3, 4, \mathbb{U}} \circ F_{(Y \sqcup S^3)(\mathbb{S}_0), \mathbb{S}_2, \bar{\mathfrak{s}}} \simeq S^-_{\phi(\mathbb{U}), \phi(D)}.$$

Proof. The proof is an easy model computation, nearly identical to the analogous statement for the 3- and 4-manifold invariants (cf. Lemma Lemma 2.12.1). □

To introduce maps for link cobordisms where a component of Σ doesn't have an end in one of the $\mathbb{L}_i$, we pick standard 4-balls along arcs of $\mathcal{A}$. The configuration is shown in Figure 4.32. After removing enough 4-balls, we can use Theorem A′ to construct cobordism maps for the punctured link cobordism, which are invariants (for a given choice of 4-balls).

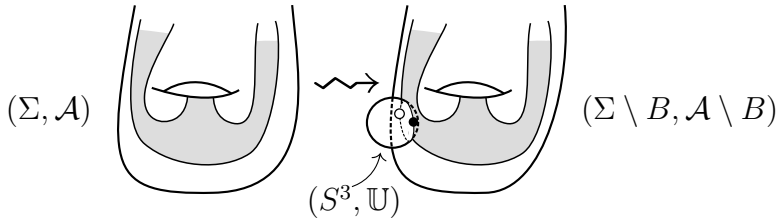

Figure 4.32: **The process of removing a standard 4-ball from $(W, \Sigma, \mathcal{A})$ which intersect Σ along an arc of $\mathcal{A}$.**

We need to show that the induced map is independent of the choice of standard 4-balls. It is not the case that any two configurations of 4-balls induce diffeomorphic cobordisms, so we cannot just apply the invariance result of Theorem A′. Instead we show that in general the maps are invariant under removing additional 4-balls along arcs of $\mathcal{A}$. By a sequence of removing additional 4-balls or filling in 4-balls, we can go from one collection to another. To this end, we prove the following lemma:

Lemma 4.12.4. *Suppose $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a decorated link cobordism with $F = (\Sigma, \mathcal{A})$, and B is a 4-ball intersecting Σ in a disk along an arc of $\mathcal{A}$. Let*

$$(W_0, F_0) : (Y_1 \sqcup S^3, \mathbb{L}_1 \sqcup \mathbb{U}) \to (Y_2, \mathbb{L}_2)$$

be the link cobordism obtained by removing B from W and Σ, then

$$F_{W,F,\mathfrak{s}} \simeq F_{W_0,F_0,\mathfrak{s}_0} \circ F_{Y_1,0,\mathbb{U}}.$$

Proof. This is a model computation. To a Morse function inducing a parametrized Kirby decomposition of (W, F), we introduce a canceling index 0/1 pair such that the index 0 critical point occurs inside of B, along an arc of $\mathcal{A}$, and the index 1 critical point occurs on Σ, in a $\mathbf{z}$-region. By Lemma 4.12.3, the composition of the 0-handle and the 1-handle map is $S^+_{(U,w,z),D)}$, for an unknot U bounding a disk D in the complement of the rest of the link.

Using the schematic in Figure 4.33, it thus sufficient to show that

$$\phi_* T^-_{w,z} F^{\mathbf{z}}_B S^+_{(U,w,z),D} \simeq \mathrm{id}$$

for the diffeomorphism $\phi : (Y, (L \cup U)(B)) \to (Y, L)$ which is the identity outside of a neighborhood of $B \cup D$. This relation follows from Proposition 4.7.4 and Lemma 4.3.13.

□

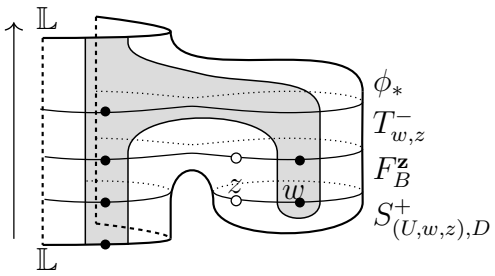

Figure 4.33: **Showing invariance under puncturing is equivalent to showing this sequence of maps is the identity.** Shown is the surface interpretation of the composition of maps in Lemma 4.12.4. The dashed lines indicate level sets of a Morse function on the surface.

Of course the dual statement is also true, that if we remove a ball from W along an arc of $\mathcal{A}$ in Σ, and then add the new boundary sphere to the outgoing end $(Y_2, \mathbb{L}_2)$, the map is unaffected:

Lemma 4.12.5. *If $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a link cobordism with $F = (\Sigma, \mathcal{A}, \sigma)$, and B is a 4-ball intersecting Σ in a disk along an arc of $\mathcal{A}$, and*

$$(W_0, F_0) : (Y_1, \mathbb{L}_1) \to (Y_2 \sqcup S^3, \mathbb{L}_2 \sqcup \mathbb{U})$$

is the link cobordism obtained by removing B from W and Σ, then

$$F_{W,F,\mathfrak{s}} \simeq F_{Y_2 \sqcup S^3, 4, \mathbb{U}} \circ F_{W_0, F_0, \mathfrak{s}_0}.$$

The proof of the above lemma is essentially identical to the proof of Lemma 4.12.4.

Indeed following this strategy, we can construct maps for arbitrary cobordisms, even where the 4-manifold has empty ends, just as in Section 2.12.2. Summarizing, we can conclude the proof of Theorem A:

Proof of Theorem A. Using Theorem A′, to a link cobordism where not every component of Σ has an incoming or outgoing end, we remove some number of standard 4-balls which intersect Σ along an arc of $\mathcal{A}$, so that each component of Σ has an incoming and outgoing end. To see that the resulting maps are independent of the choice of 4-balls, we use Lemmas 4.12.4 and 4.12.5 to see that we can remove as many standard 4-balls as we'd like, along the arcs $\mathcal{A}$, without changing the maps. Through a sequence of removing standard 4-balls along $\mathcal{A}$, or filling them back in, we can go between any two configurations of standard 4-balls. □

4.13 The composition law

In this Section we prove Theorem B, the Spin^c composition law. Given a link cobordism which is the composition of two link cobordisms, we will take two parametrized Kirby decompositions (one for each link cobordism), and then commute pieces past each other to get a decomposition into elementary parametrized link cobordisms of the composed cobordism, which is almost a parametrized Kirby decomposition, except there are two elementary cobordisms which have 1-dimensional framed links. Most of the Floer homology has already been done, but we have a few lemmas which we need to prove.

Lemma 4.13.1. *Suppose $\mathbb{L} \subseteq Y$ is a link, and let $\mathbb{L}'$ be the link obtained by adding adjacent basepoints (w, z). If $\mathbb{S} \subseteq Y \setminus L$ is a framed 1-dimensional link, and $\mathfrak{s} \in W(Y, \mathbb{S})$ is a* Spin^c *structure, then*

$$S^{\pm}_{w,z} F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F_{Y,\mathbb{L}',\mathbb{S}_1,\mathfrak{s}} S^{\pm}_{w,z}, \qquad \text{and} \qquad T^{\pm}_{w,z} F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F_{Y,\mathbb{L}',\mathbb{S}_1,\mathfrak{s}} T^{\pm}_{w,z}$$

Proof. This follows by writing $F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ as $F^{\beta}_{Y,\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\beta}$ for a β-bouquet $\mathcal{B}^\beta$ of the link $\mathbb{S}_1$, and then using [MO10, Prop. 5.2], which is restated using our matrix notation in Lemma 4.5.1. Of course if one wants to use an α-bouquet $\mathcal{B}^\alpha$ for $\mathbb{S}_1$, then one can use the triangle map computation in Theorem 4.5.3 instead. □

Lemma 4.13.2. *Suppose that B is a band for the link $\mathbb{L}$ in Y, and $\mathbb{S}_1$ is a framed 1-dimensional link in Y, which is disjoint from $L \cup B$ and $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}))$. If $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$,*

then

$$F_{Y,\mathbb{L}(B),\mathbb{S}_1,\mathfrak{s}} F^{\mathbf{o}}_B \simeq F^{\mathbf{o}}_B F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}.$$

Proof. Without loss of generality, suppose that B is an α-band. We now pick a β-bouquet, $\mathcal{B}$, for $\mathbb{S}_1$ in the complement of $B \cup L$. We also pick a Heegaard quadruple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ so that

1. $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is Heegaard diagram for $(Y, \mathbb{L})$;
2. $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a triple subordinate to the band $B \subseteq Y$;
3. $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a triple subordinate to the bouquet $\mathcal{B}$ in $(Y, \mathbb{L})$;
4. $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a triple subordinate to the band B in $Y(\mathbb{S}_1)$;
5. $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a triple subordinate to the bouquet $\mathcal{B}$ in $(Y, \mathbb{L}(B))$.

Let $\Theta^{\mathbf{o}}_{\alpha'\alpha} \in \mathcal{H}FL^-_{\max,\mathbf{o}}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ denote the generator. Similarly let $\Theta_{\beta\beta'}$ be the generator of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}')$ (see Lemma 4.3.1). The associativity relations applied to the quadruple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ yield the equivalence

$$F_{\alpha'\alpha\beta',\mathfrak{s}|_{Y(\mathbb{S})}}(\Theta^{\mathbf{o}}_{\alpha'\alpha}, F_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}, \Theta_{\beta\beta'})) \simeq F_{\alpha'\beta\beta,\mathfrak{s}}(F_{\alpha'\alpha\beta,\mathfrak{s}|_Y}(\Theta^{\mathbf{o}}_{\alpha'\alpha}, \mathbf{x}), \Theta_{\beta\beta'}),$$

showing that

$$F^{\beta}_{Y,\mathbb{L}(B),\mathbb{S}_1,\mathfrak{s}} F^{\mathbf{o}}_B \simeq F^{\mathbf{o}}_B F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}.$$

Noting that $F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$, by Lemma 4.4.8, the proof is complete. □

We now prove a slightly weaker version of Theorem B:

Lemma 4.13.3. *Suppose that* $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ *is an arbitrary link cobordism and*

$$(W, F) = (W_2, F_2) \circ (W_1, F_1)$$

where each (W_i, F_i) is a link cobordisms such that each component of W has a nonempty incoming and outgoing end, and each component of Σ contains a nonempty incoming and outgoing end. If $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are Spin^c *structures on W_1 and W_2 respectively, then we have*

$$F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1,\mathfrak{s}_1} \simeq \sum_{\substack{\mathfrak{s}\in \mathrm{Spin}^c(W) \\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}} F_{W,F,\mathfrak{s}}.$$

Proof. Take parametrized Kirby decompositions $\mathcal{K}_1 = \mathcal{K}(f_1, \mathbf{b}_1)$ of (W_1, F_1) and $\mathcal{K}_2 = \mathcal{K}(f_2, \mathbf{b}_2)$ of (W_2, F_2), induced by very nice Morse functions f_1 and f_2 with collections of regular values $\mathbf{b}_1$ and $\mathbf{b}_2$. Also pick gradient like vector fields v_1 and v_2, inducing the decompositions. We can assume that f_1 and f_2 patch together, and v_1 and v_2 patch together to give Morse functions and gradient like flowlines on (W, F). The collections of regular values $\mathbf{b}_1$ and $\mathbf{b}_2$ induce a collection $\mathbf{b}$ for f. This yields a decomposition, $\mathcal{K}(f, \mathbf{b})$, of (W, F) into elementary parametrized cobordisms, which is the one obtained by just stacking $\mathcal{K}_1$ and $\mathcal{K}_2$. We note that $\mathcal{K}(f, \mathbf{b})$ may not be a parametrized Kirby decomposition since the terms are in the wrong order, and there may be two 2-handle cobordisms. Nonetheless, we can still take the composition of the cobordism maps for each elementary parametrized cobordism in the decomposition, which results in just $F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1\mathfrak{s}_1}$.

Notice that the index 1, 2 and 3 dimensional critical points of f_1 which are away from Σ_1 have unstable manifolds which do not intersect Σ, and similarly the critical points of f_2 away from Σ have stable manifolds which don't intersect Σ. Hence, assuming that (f, v) is generic, we can modify the Morse functions to pull all of the critical points of $f_1|_{W_1\setminus\Sigma_1}$ above any critical point of $f_2|_{\Sigma_2}$ along Σ, and then pull the index 3 critical points of f_1 above the index 1 and 2 critical points of f_2, and pull the index 2 critical points of f_1 above the index 1 critical points of f_1.

The composition of the map for each piece of the stacked decomposition is invariant under pulling the 3-handles of f_1 above the 1-handles of f_2 by Lemma 4.4.4. The composition is also invariant of pulling the 3-handles of f_1 past the 2-handles of f_2 by the triangle map computation of Theorem 2.7.8.

Hence it remains to show that the composition is unchanged when we pull the 1-,2- and 3-handles of $f_1|_{W_1\setminus\Sigma_1}$ past the elementary cobordisms associated to critical points of $f_2|_\Sigma$ and $f_2|_{\mathcal{A}}$. The critical points of $f_2|_{\mathcal{A}}$ induce maps which are a composition of the basepoint moving maps along the link components, and the quasi-stabilization maps. The maps for surgery on framed spheres away from the link commute with the basepoint moving maps by diffeomorphism invariance of the maps. The 1- and 3-handles of f_1 can be pulled past the critical points of $f_2|_{\mathcal{A}}$ without changing the composition of the maps by Lemma 4.4.3 (showing that the 1-handle and 3-handle maps commute with quasi-stabilization). The 1- and 3-handle maps can be commuted past critical points of $f_2|_\Sigma$ using Lemma 4.4.4 (to commute a 1- or 3-handle past the 1-handle map from the critical point along Σ_2) as well as the triangle map computation from Theorem 2.7.8 (to commute a 1- or 3-handle past the band map for the critical point along Σ_2). The 2-handles of f_1 can be commuted past the critical points of $f_2|_{\mathcal{A}}$ by Lemma 4.13.1. Similarly the 2-handles of f_1 can be commuted past the critical points of $f_2|_\Sigma$ by Lemma 4.13.2 (to commute the 2-handle maps past the band maps), and Theorem 2.7.8 (to commute the 2-handle maps past the 1-handle map).

Thus we get that $F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1,\mathfrak{s}_1}$ is equal to the composition of maps induced by a decomposition of (W,F) into parametrized elementary cobordisms, which has elementary cobordisms satisfying the requirements of a parametrized Kirby decomposition, except that there are exactly two adjacent terms corresponding to surgery on framed 1-dimensional links. As W is obtained from the union of the two 2-handle pieces by adding 1-handles and 3-handles, a Spin^c structure defined on the union of the two 2-handle cobordisms extends uniquely to a Spin^c structure on all of W. Hence Lemma 4.4.9 (the composition law for the 2-handle maps) implies the Lemma. □

Using the previous Lemma, we can now prove Theorem B:

Proof Theorem B. This now follows from Lemma 4.13.3, and the construction of the general cobordism maps for cobordisms with empty ends defined in Section 4.12 using the same strategy as the final component of the proof of Theorem E found in Subsection 2.12.3 of that

paper. See also [Juh16, Theorem 11.3]. □

4.14 The graph TQFT as the reduction of the link Floer TQFT

Throughout this section, we will restrict to considering only surfaces with divides which are colored using exactly two colors, one for the $\mathbf{w}$ basepoints and regions, and one for the $\mathbf{z}$ basepoints and regions. Let U denote the variable corresponding to the $\mathbf{w}$ basepoints and regions, and let V denote the variable for the $\mathbf{z}$ regions. It is straightforward to generalize the results of this section to more general colorings, as long as no $\mathbf{w}$ basepoint or region shares the same color as a $\mathbf{z}$ basepoint or region.

If $(Y, \mathbb{L})$ is a 3–manifold with multi-based link, we note that

$$\mathcal{C}FL^-(Y, \mathbb{L}, \mathfrak{s}) \otimes_{\mathbb{Z}_2[U,V]} \mathbb{Z}_2[U,V]/(V-1) \cong CF^-(Y, \mathbf{w}, \mathfrak{s})$$

and similarly

$$\mathcal{C}FL^-(Y, \mathbb{L}, \mathfrak{s}) \otimes_{\mathbb{Z}_2[U,V]} \mathbb{Z}_2[U,V]/(U-1) \cong CF^-(Y, \mathbf{z}, \mathfrak{s} - PD[L])$$

The change in Spin^c structure of the $U = 1$ reduction is a consequence of Lemma 4.2.8.

The link cobordism maps $F_{W,F,\mathfrak{s}}$ thus naturally induce maps on CF^-, which we call the *reductions* of $F_{W,F,\mathfrak{s}}$. We will commonly write $F_{W,F,\mathfrak{s}}|_{U=1}$ or $F_{W,F,\mathfrak{s}}|_{V=1}$, for these reductions.

In Chapter 2, we constructed a graph TQFT for Heegaard Floer homology. We recall from the introduction the following definition:

Definition 4.14.1. A **ribbon graph cobordism** $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$ between two multi-based 3–manifolds is a pair (W, Γ) such that

1. W is a cobordism from Y_1 to Y_2;
2. Γ is an embedded graph in W such that $\Gamma \cap \partial W = \mathbf{w}_1 \cup \mathbf{w}_2$ and each basepoint in $\mathbf{w}_1 \cup \mathbf{w}_2$ has valence 1 in Γ;

3. Γ is decorated with a ribbon structure, i.e., a choice of cyclic ordering of the edges adjacent to each vertex;

The above definition naturally inspires the following definition:

Definition 4.14.2. A **ribbon surface cobordism** (W, R) between two multi-based 3–manifolds $(Y_1, \mathbf{w}_1)$ and $(Y_2, \mathbf{w}_2)$ is a pair such that W is a cobordism from Y_1 to Y_2, and $R \subseteq W$ is an oriented surface with boundary and corners such that the following holds:

1. $R \cap Y_i$ consists of a finite collection of closed intervals, each containing exactly one of the basepoints of $\mathbf{w}_1 \cup \mathbf{w}_2$. Furthermore, each basepoint in $\mathbf{w}_1 \cup \mathbf{w}_2$ is contained in a component of $R \cap Y_i$;
2. The corners of R correspond exactly to the boundaries of the closed intervals forming $R \cap Y_i$;
3. R contains no closed components.

Lemma 4.14.3. *Given a ribbon graph cobordism $(W, \Gamma) : (Y_1, \mathbf{w}_1) \to (Y_2, \mathbf{w}_2)$, there is ribbon surface cobordism (W, R_Γ), which is well defined up to isotopies of R_Γ through ribbon surfaces $(R_\Gamma)_t$ in W, such that $(W, (R_\Gamma)_t)$ is a ribbon graph cobordism from $(Y_1, \mathbf{w}_1)$ to $(Y_2, \mathbf{w}_2)$ for all t.*

Proof. We first show how to construct the surface R_Γ. Pick an orientation of each edge of Γ. Isotope Γ inside of W near each vertex of $\Gamma \cap \operatorname{int} W$, so that the edges meet in a 2–plane, and the cyclic ordering that they appear in the two plane (with respect to a choice of orientation of the 2–plane) agrees with the ordering from the ribbon structure. Call this graph $\widetilde{\Gamma}$. To each vertex in $\widetilde{\Gamma} \cap \operatorname{int} W$, we attach an oriented disk centered at the vertex, such that the edges of Γ are tangent to the tangent 2–plane of the disk, and the ordering of the edges agrees with the ordering from the ribbon structure. Along each edge, we pick a vector field w_e, tangent to e, and positive, with respect to the orientation of the edge e. We also pick a non-vanishing vector v_e along e in the orthogonal complement of Te. The restriction of v_e

to the disks associated to the vertices can be specified by requiring the pair (v_e, w_e) be an oriented basis for the tangent space space of the disks appearing at the ends of e. At the vertices of $\Gamma \cap \partial W$, we have no requirement on w_e. The restriction of w_e to ∂e determines the vector field w_e on all of e, up to homotopy, as the space of such vector fields is homotopy equivalent to the space of maps from $(e, \partial e)$ to $(S^2, \{pt\})$. Note that the resulting surface is orientable, and is independent of the choice of orientation along each edge.

To show uniqueness up to isotopy of the surface R_Γ, we note that the construction involved several steps: choosing an orientation of each edge, isotoping Γ to a graph $\widetilde{\Gamma}$ where the vertices meet in a plane, and then picking the vector field w_e. As we described above, the surface is independent of the choice of orientation of the edges, as well as the choice of vector field. Thus we just need to show independence of the choice of isotopic $\widetilde{\Gamma}$. By construction, the graph $\widetilde{\Gamma}$ is formed by isotoping the edges inside a neighborhood of the vertices. Clearly any two choices of isotopic $\widetilde{\Gamma}$ are homotopic. Since W is a 4–manifold, a homotopy of graphs inside of W can be taken to be an isotopy. □

Remark 4.14.4. We will implicitly round the edges of the surface R_Γ near where the strips have been attached to the disks.

Definition 4.14.5. We say that two graph cobordisms (W, Γ_1) and (W, Γ_2), both from $(Y_1, \mathbf{w}_1)$ to $(Y_2, \mathbf{w}_2)$ are **ribbon equivalent** if the two ribbon surfaces R_{Γ_1} and R_{Γ_2} are isotopic in W through ribbon surfaces.

Given a ribbon surface cobordism (W, R), we can go in the other direction, and construct a ribbon graph cobordism (W, Γ_R) such that $\Gamma_R \subseteq R$ is essentially a deformation retraction of the graph, in the following sense:

Definition 4.14.6. If (W, R) is a ribbon surface cobordism from $(Y_1, \mathbf{w}_1)$ to $(Y_2, \mathbf{w}_2)$, we say that a ribbon graph $\Gamma \subseteq R$ is a **ribbon graph core** of R if the cyclic orders of Γ correspond to the orientation of R, and if R_Γ is isotopic (viewed as a map $R_\Gamma \hookrightarrow R$) to a diffeomorphism of $R_\Gamma \cong R$.

Note that there are many non-isotopic ribbon graph cores for a fixed ribbon surface (in fact there are many which are not even isomorphic as graphs). We will not endeavor to create a set of moves between two ribbon graph cores of ribbon surfaces, though a set of moves could probably be made by a Morse theory argument. In Figure 4.34 we show a few examples of ribbon graph cores of surfaces.

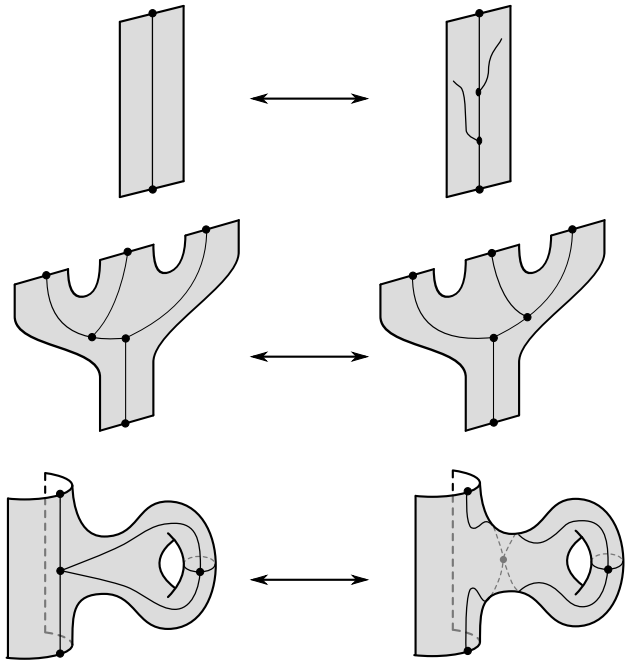

Figure 4.34: **Three examples of different choices of ribbon graph cores for ribbon surfaces.** The ribbon graph cores are, by definition, ribbon equivalent. The pairs of ribbon graphs are not isomorphic as ribbon graphs. The top pair of ribbon graph cores are related by adding trivial strands. The middle pair are related by sliding an edge across a vertex, in a way which is compatible with the cyclic orderings. The bottom pair can be related by a sequence of such moves. Note that the bottom graphs are isomorphic as graphs, but not as graphs with cyclic ordering.

Theorem C. *Suppose that* $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ *is a decorated link cobordism with* $F = (\Sigma, \mathcal{A})$. *If* $\Gamma(\Sigma_{\mathbf{w}})$ *and* $\Gamma(\Sigma_{\mathbf{z}})$ *are choices of ribbon graph cores of* $\Sigma_{\mathbf{w}}$ *and* $\Sigma_{\mathbf{z}}$ *(resp.), then*

$$F_{W,F,\mathfrak{s}}|_{V=1} = F_{W,\Gamma(\Sigma_{\mathbf{w}}),\mathfrak{s}}$$

and

$$F_{W,F,\mathfrak{s}}|_{U=1} = F_{W,\Gamma(\Sigma_{\mathbf{z}}),\mathfrak{s}-PD[\Sigma]}.$$

The following two corollaries follow from the above theorem:

Corollary D. *If (W,Γ) and (W,Γ') are ribbon equivalent ribbon graph cobordisms, then*

$$F^A_{W,\Gamma,\mathfrak{s}} \simeq F^A_{W,\Gamma',\mathfrak{s}},$$

and similarly for the maps $F^B_{W,\Gamma,\mathfrak{s}}$.

Corollary E. *The $V=1$ reduction of $F_{W,F,\mathfrak{s}}$ induces a map*

$$F_{W,F,\mathfrak{s}}|_{V=1} : CF^-(Y_1,\mathbf{w}_1,\mathfrak{s}|_{Y_1}) \to CF^-(Y_2,\mathbf{w}_2,\mathfrak{s}|_{Y_2})$$

which only depends on $\mathfrak{s}$ and the ribbon surface $\Sigma_{\mathbf{w}}$. The $U=1$ reduction of $F_{W,F,\mathfrak{s}}$ induces a map

$$F_{W,F,\mathfrak{s}}|_{U=1} : CF^-(Y_1,\mathbf{z}_1,\mathfrak{s}-PD[L_1]) \to CF^-(Y_2,\mathbf{z}_2,\mathfrak{s}-PD[L_2])$$

which depends only on the ribbon surface $\Sigma_{\mathbf{z}}$, $\mathfrak{s}$ and the homology class $[\Sigma]\in H_2(W,\partial W;\mathbb{Z})$.

4.14.1 Background on the graph TQFT

In this section, we recall the construction of the graph cobordism maps from Chapter 2, and prove some basic results, which will be helpful for our purposes.

There are several components of the graph cobordism maps. Firstly, there are 1–handle, 2–handle and 3–handle maps, analogously to the maps defined by Ozsváth and Szabó [OS06], and similar to the ones we defined for the link cobordism maps. There are also 0–handle and 4–handle maps, corresponding to adding or removing a copy of (S^3,w). These have the obvious effect on the chain complexes. A key component of the construction corresponded to defining maps for graphs embedded in a 3–manifold. To this end, we make the following definition.

Definition 4.14.7. A **ribbon flow graph** $\mathcal{G}=(\Gamma,V_1,V_2)$ in Y^3 is a tuple consisting of a ribbon flow graph $\Gamma\subseteq Y$, with collections of vertices V_1 and V_2 such that the following hold:

1. V_1 and V_2 both intersect each component of Y non-trivially;

2. Each vertex in V_1 and V_2 has valence 1 in Γ;

3. Each vertex in Γ has valence at least 1.

Composition of flow graphs is given by concatenation of graphs.

In Chapter 2, a *graph action map* was defined for a ribbon flow graph. If $\mathcal{G} = (\Gamma, V_1, V_2)$ is a colored ribbon flow graph in a fixed 3–manifold Y, we will construct a chain map

$$A_{\mathcal{G},\mathfrak{s}} : CF^-(Y, V_1, \mathfrak{s}) \to CF^-(Y, V_2, \mathfrak{s}).$$

These are described in more detail in Chapter 2. There is a natural variation, $B_{\mathcal{G},\mathfrak{s}}$, with the same domain and range, which will discuss later.

In the next section, we recall the maps which are used in the construction.

4.14.2 Ingredients of the graph TQFT

There are two key ingredients of the graph action map: the relative homology maps, and the free-stabilization maps. Given a path λ between basepoints w_1 and w_2 in a collection of basepoints $\mathbf{w}$ in Y, there is a -1 graded endomorphism

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

which satisfies

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

On the level of Heegaard diagrams, the map A_λ is defined by homotoping the path λ in Y so that it lies in the Heegaard surface, then defining

$$A_\lambda(\mathbf{x}) = \sum_{\substack{\phi \in \pi_2(\mathbf{x},\mathbf{y}) \\ \mu(\phi)=1}} a(\lambda, \phi) \# \widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}.$$

The quantity $a(\lambda, \phi)$ is equal to the sum of changes of the class ϕ over the $\boldsymbol{\alpha}$ curves, as one traverses λ. We could of course define the map B_λ analogously, and count the sum of

changes across the $\boldsymbol{\beta}$ curves instead. The sum $A_\lambda + B_\lambda$ then counts disks with an extra factor equal to the sum of changes across all curves, i.e. which is just the difference between the multiplicities at w_1 and w_2. Hence

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

The following relations are obtained in Chapter 2, by counting the ends of index 2 moduli spaces of holomorphic curves:

$$A_\lambda^2 \simeq U_{w_1} \simeq U_{w_2} \tag{4.14}$$

$$A_{\lambda_1}A_{\lambda_2} + A_{\lambda_2}A_{\lambda_1} \simeq \sum_{w\in(\partial\lambda_1)\cap(\partial\lambda_2)} U_w. \tag{4.15}$$

The second ingredient of the graph action map are the *free-stabilization* maps, associated to adding or removing basepoints. If $(Y, \mathbf{w})$ is a multi-based 3–manifold and $w \notin \mathbf{w}$, we can define a map

$$S_w^+ : CF^-(Y, \mathbf{w}, \mathfrak{s}) \to CF^-(Y, \mathbf{w}\cup\{w\}, \mathfrak{s})$$

and a map S_w^- in the opposite direction. The free stabilization map is essentially defined by adding in a copy of (S^3, w_0) via a 0–handle, and then attaching the new component via a 1–handle with a foot near w and a foot near w_0. As such the formula for S_w^+ is the same as the 1–handle map defined in Section 4.4.2. In Chapter 2, it is shown that

$$S_{w_1}^{\circ}S_{w_2}^{\circ'} \simeq S_{w_2}^{\circ'}S_{w_1}^{\circ} \tag{4.16}$$

for any distinct w_1 and w_2 and $\circ, \circ' \in \{+,-\}$ (see also Lemma 4.4.4). In light of the above relation, if $\{w_1, \dots, w_n\}$ is a collection of basepoints, we will write

$$S_{w_n,\dots,w_1}^+ := S_{w_n}^+ \circ \cdots \circ S_{w_1}^+,$$

noting that the composition is independent of the ordering of $w_1, \dots, w_n$.

Also,

$$S_w^{\pm}A_\lambda \simeq A_\lambda S_w^{\pm} \tag{4.17}$$

as long as $w \notin \partial\lambda$. Additionally

$$S_w^+ S_w^- \simeq \Phi_w \tag{4.18}$$

and

$$S_w^- S_w^+ \simeq 0. \tag{4.19}$$

If λ is a path from w_1 to w_2, and $w_1 \in \mathbf{w}$ but $w_2 \notin \mathbf{w}$, then

$$S_{w_2}^- A_\lambda S_{w_2}^+ \simeq \mathrm{id}\,. \tag{4.20}$$

Sometimes we refer to this as the *trivial strand relation*, since it corresponds to invariance from the manipulation of graphs shown on the top of Figure 4.34.

Finally, there is the *basepoint moving relation*. If λ is a path from w_1 to w_2 and $w_1, w_2 \notin \mathbf{w}$, then the diffeomorphism which moves w_1 to w_2 along λ induces a chain homotopy equivalence

$$\phi_\lambda : CF^-(Y, \mathbf{w} \cup \{w_1\}, \mathfrak{s}) \to CF^-(Y, \mathbf{w} \cup \{w_2\}, \mathfrak{s}).$$

The basepoint moving relation is that

$$(\phi_\lambda)_* \simeq S_{w_1}^- A_\lambda S_{w_2}^+. \tag{4.21}$$

4.14.3 The graph action map

To define the graph action map $A_{\mathcal{G},\mathfrak{s}}$, one subdivides the graph Γ by adding extra vertices, and writes $\mathcal{G}$ (now subdivided) as a composition of *elementary flow graphs*. Such a decomposition is termed a *Cerf decomposition* of the graph. By definition, there are three types of elementary flow graphs $\mathcal{G} = (\Gamma, V_1, V_2)$:

(EFG 1) $V(\Gamma) = V_1 \sqcup V_2$ and all edges go from V_1 to V_2;

(EFG 2) $V(\Gamma) = V_1 \sqcup V_2 \sqcup \{v_0\}$ and all edges either connect a vertex in V_1 to V_2 or connect a vertex in $V_1 \cup V_2$ to $\{v_0\}$;

(EFG 3) $V(\Gamma) = V_1 \sqcup V_2$ and there is exactly one edge which connects two vertices in V_1 or two vertices in V_2.

The three subtypes of elementary ribbon flow graphs are shown in Figure 4.35. Note than any flow graph can be subdivided so that it is a concatenation of elementary flow graphs. In Chapter 2, it is shown that the graph action map is invariant under subdivision, and furthermore, independent of the choice of decomposition into elementary flow graphs.

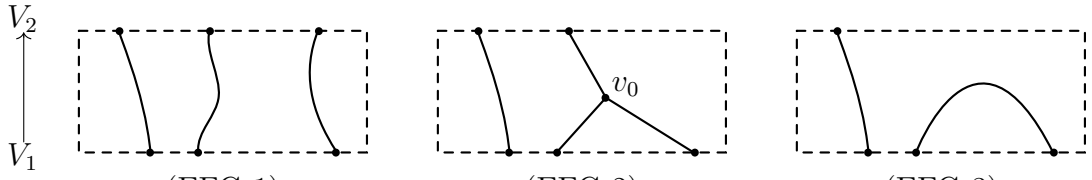

Figure 4.35: **Examples of elementary ribbon flow graphs.**

To an elementary flow graph $\mathcal{G} = (\Gamma, V_2, V_1)$ of Type (1), the graph action map is, by definition,

$$A_{\mathcal{G},\mathfrak{s}} := \Big(\prod_{v\in V_2} S_v^-\Big)\Big(\prod_{e\in E(\Gamma)} A_e\Big)\Big(\prod_{v\in V_1} S_v^+\Big). \tag{EFG 1}$$

The map doesn't depend on the ordering of the terms in each factor, by Equations (4.16) and (4.15). To an elementary flow graph $\mathcal{G}$ of Type (2), the graph action map is defined by

$$A_{\mathcal{G},\mathfrak{s}} := \Big(\prod_{v\in V_1\cup\{v_0\}} S_v^-\Big)\Big(\prod_{\substack{e\in E(\Gamma)\\ v_0\notin\partial e}} A_e\Big)(A_{e_n}\cdots A_{e_1})\Big(\prod_{v\in V_2\cup\{v_0\}} S_v^+\Big), \tag{EFG 2}$$

where $e_1, \ldots, e_n$ are the edges adjacent to v_0, in the order that they appear according to the cyclic ordering. In Chapter 2 it is shown that this expression is invariant under cyclic permutation of the edges $e_1, \ldots, e_n$.

Finally for a type (3) elementary flow graph, the graph action is defined by

$$A_{\mathcal{G},\mathfrak{s}} := \Big(\prod_{v\in V_2} S_v^-\Big)\Big(\prod_{e\in E(\Gamma)} A_e\Big)\Big(\prod_{v\in V_1} S_v^+\Big). \tag{EFG 3}$$

Maps $B_{\mathcal{G},\mathfrak{s}}$ can be defined by the same formulas, except using the maps B_λ instead of A_λ. A natural question is how the maps $A_\mathcal{G}$ and $B_\mathcal{G}$ are related. To this end we have the following:

Proposition 4.14.8. *If $\mathcal{G} = (\Gamma, V_1, V_2)$ is a ribbon flow graph, and $\overline{\mathcal{G}} = (\overline{\Gamma}, V_1, V_2)$ denotes the ribbon flow graph with all cyclic orders reversed, then*

$$B_{\overline{\mathcal{G}},\mathfrak{s}} = A_{\mathcal{G},\mathfrak{s}}.$$

We will not have a need for the above relation. A proof for trivalent graphs can be found in [HMZ16, Lemma 5.9]. For general graphs, one can use Lemma 4.14.10, below, to reduce to the case of trivalent graphs.

4.14.4 Relating the graph action map to link Floer homology

In this section, we prove several useful computations of graph action maps, and compute the reductions of certain maps appearing in link Floer homology. In the following section, we will use these results to prove Theorem **??**.

A first relation, which will be useful for us, is related to subdivision invariance of the graph action map:

Lemma 4.14.9. *Suppose that e_1 and e_2 are two paths which share a single endpoint w. Then, writing $e_1 * e_2$ for the concatenation, one has*

$$A_{e_2 * e_1} \simeq S_w^- A_{e_2} A_{e_1} S_w^+.$$

Proof. On a diagram which has w as a basepoint, we note that $A_{e_2 * e_1} \simeq A_{e_2} + A_{e_1}$. Hence $A_{e_2} = A_{e_2 * e_1} + A_{e_1}$. Hence

$$S_w^- A_{e_2} A_{e_1} S_w^+ \simeq S_w^-(A_{e_2 * e_1} + A_{e_1}) A_{e_1} S_w^+ \simeq A_{e_2 * e_1} S_w^- A_{e_1} S_w^+ + U S_w^- S_w^+ \simeq A_{e_2 * e_1},$$

as we wanted. □

Clearly any ribbon flow graph in Y is ribbon equivalent to a trivalent ribbon graph. At any vertex of valence greater than three, one can split off a consecutive pair of edges, reducing the number of edges adjacent to the vertex, while adding a new trivalent vertex to the graph. We show that the graph action map is invariant under this move:

Lemma 4.14.10. *If $\mathcal{G} = (\Gamma, V_1, V_2)$ is a ribbon flow graph in Y, and $\mathcal{G}' = (\Gamma', V_1, V_2)$ is related to $\mathcal{G}$ by splitting off a consecutive pair of edges at a given vertex of Γ, and forming a new trivalent vertex, then*

$$A_{\mathcal{G},\mathfrak{s}} \simeq A_{\mathcal{G}',\mathfrak{s}}.$$

Proof. Since the graph action map is constructed by breaking a subdivision of the graph into a composition of elementary flow graphs, it is sufficient to prove the claim for elementary flow graphs satisfying (EFG 2), the only type of elementary ribbon flow graphs which can have a vertex of valence greater than three. Suppose that v_0 is the middle vertex and suppose that $e_1, \ldots, e_n$ are the edges adjacent to v_0, indexed according to the cyclic ordering assigned to v_0. Let e_0', e_1', e_2', denote three new edges adjacent to a new vertex v_0', as in Figure 4.36, which form a ribbon equivalent graph in Y. By definition

$$A_{\mathcal{G},\mathfrak{s}} = \Big(\prod_{v \in V_1 \cup \{v_0\}} S_v^- \Big) \Big(\prod_{\substack{e \in E(\Gamma) \\ v_0 \notin \partial e}} A_e \Big) (A_{e_n} \cdots A_{e_1}) \Big(\prod_{v \in V_2 \cup \{v_0\}} S_v^+ \Big).$$

By commuting the various S_v^+ amongst themselves, we can isolate a factor of $A_{e_2} A_{e_1} S_{v_0}^+$. We now claim that

$$S_{v_0'}^- A_{e_0'} A_{e_2'} A_{e_1'} S_{v_0', v_0}^+ \simeq A_{e_2} A_{e_1} S_{v_0}^+. \tag{4.22}$$

To this end, the easiest strategy is to identify both expressions with a third expression, namely

$$S_{v_0'}^- A_{e_0'} (A_{e_2'} + A_{e_0'})(A_{e_1'} + A_{e_0'}) S_{v_0', v_0}^+. \tag{4.23}$$

To show the left side of Equation (4.22) is equal to the Equation (4.23), we compute by

multiplying everything out

$$\begin{aligned}
S^-_{v_0'}A_{e_0'}(A_{e_2'}+A_{e_0'})(A_{e_1'}+A_{e_0'})S^+_{v_0',v_0} &\simeq S^-_{v_0'}(A_{e_0'}A_{e_2'}A_{e_1'}+A^2_{e_0'}A_{e_1'}+A_{e_0'}A_{e_2'}A_{e_0'}+A^2_{e_0'}A_{e_1'}+A^3_{e_0'})S^+_{v_0',v_0}\\
&\simeq S^-_{v_0'}(A_{e_0'}A_{e_2'}A_{e_1'}+A^2_{e_0'}A_{e_1'}+A^2_{e_0'}A_{e_2'}+UA_{e_0'}+A^2_{e_0'}A_{e_1'}+A^3_{e_0'})S^+_{v_0',v_0}\\
&\simeq S^-_{v_0'}(A_{e_0'}A_{e_2'}A_{e_1'}+UA_{e_1'}+UA_{e_2'}+UA_{e_0'}+UA_{e_1'}+UA_{e_0'})S^+_{v_0',v_0}\\
&\simeq S^-_{v_0'}A_{e_0'}A_{e_2'}A_{e_1'}S^+_{v_0',v_0}+US^-_{v_0'}(A_{e_1'}+A_{e_2'}+A_{e_0'}+A_{e_1'}+A_{e_0'})S^+_{v_0',v_0}\\
&\simeq S^-_{v_0'}A_{e_0'}A_{e_2'}A_{e_1'}S^+_{v_0',v_0}\\
&\qquad +U(S^-_{v_0'}A_{e_1'}S^+_{v_0'}+S^-_{v_0'}A_{e_2'}S^+_{v_0'}+S^-_{v_0'}A_{e_1'}S^+_{v_0'}+S^-_{v_0'}A_{e_0'}S^+_{v_0'})S^+_{v_0}\\
&\simeq S^-_{v_0'}A_{e_0'}A_{e_2'}A_{e_1'}S^+_{v_0',v_0}+4\cdot US^+_{v_0}\\
&\simeq S^-_{v_0'}A_{e_0'}A_{e_2'}A_{e_1'}S^+_{v_0',v_0}.
\end{aligned}$$

To see that the right side of Equation (4.22) is equal to the expression from Equation (4.23), we compute that

$$\begin{aligned}
S^-_{v_0'}A_{e_0'}(A_{e_2'}+A_{e_0'})(A_{e_1'}+A_{e_0'})S^+_{v_0',v_0} &\simeq S^-_{v_0'}A_{e_0'}(A_{e_2'*e_0'})(A_{e_1'*e_0'})S^+_{v_0',v_0}\\
&\simeq S^-_{v_0'}A_{e_0'}S^+_{v_0'}(A_{e_2'*e_0'})(A_{e_1'*e_0'})S^+_{v_0}\\
&\simeq A_{e_2}A_{e_1}S^+_{v_0}.
\end{aligned}$$

Hence we conclude that

$$A_{\mathcal{G},\mathfrak{s}}\simeq\Big(\prod_{v\in V_1\cup\{v_0\}}S^-_v\Big)\Big(\prod_{\substack{e\in E(\Gamma)\\ v_0\notin\partial e}}A_e\Big)(A_{e_n}\cdots A_{e_3})(S^-_{v_0'}A_{e_0'}A_{e_2'}A_{e_1'}S^+_{v_0',v_0})\Big(\prod_{v\in V_2}S^+_v\Big).$$

This expression, however, is not the expression for a Cerf decomposition of $\mathcal{G}'$. To obtain a Cerf decomposition of $\mathcal{G}'$, we use Lemma 4.14.9 to subdivide the edge e_0', and also all the edges which are not adjacent to v_0. We leave this last step to the reader, however the process is shown in Figure 4.36. □

Remark 4.14.11. The above lemma implies invariance of the graph action map under ribbon moves like the second row of Figure 4.34 by applying the above lemma twice. To move from the left ribbon graph to the right ribbon graph, one first merges the two trivalent vertices

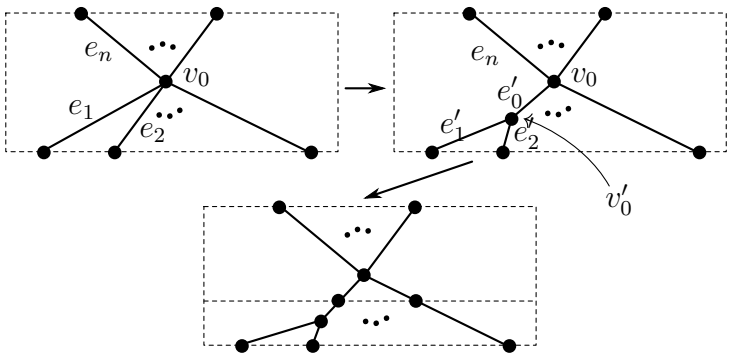

Figure 4.36: **Replacing a high valence vertex with a lower valence vertex and a trivalent vertex.** On the top, we show the manipulation performed from Lemma 4.14.10. In order to turn this into a Cerf decomposition of the flow graph, one must additionally also subdivide several edges of the graph. Cyclic orders are counterclockwise with respect to the page.

to get a ribbon graph with a 4–valent vertex, then one applies the previous lemma to split them again, to get the graph on the right.

For the purposes of comparing the graph TQFT to the reductions of the link Floer TQFT, we will need to compute the graph action map for the graph shown in Figure 4.36.

Lemma 4.14.12. *Suppose λ is a path in Y from w_1 to w_2. Let Γ_H denote the graph $(\{w_1, w_2\} \times [0,1]) \cup \lambda \times \{\frac{1}{2}\}$, with cyclic orders as shown in Figure 4.37. Then*

$$F^A_{Y\times[0,1],\Gamma_H,\mathfrak{s}} \simeq A_\lambda + U\Phi_{w_1}, \qquad \text{and} \qquad F^B_{Y\times[0,1],\Gamma_H,\mathfrak{s}} \simeq B_\lambda + U\Phi_{w_1}.$$

The same formula holds in the case that Y has additional basepoints $\mathbf{w}$, and we add $\mathbf{w}\times[0,1]$ to Γ_H.

Proof. In Figure 4.36, we show a Cerf decomposition for $\mathcal{H}$ into two elementary ribbon flow graphs, both satisfying (EFG 2). The ribbon graph $\mathcal{H} = (H, \{w_1, w_2\}, \{w_1', w_2'\})$ is obtained by pushing Γ_H into $Y \times \{0\}$. The proof follows from some symbol pushing, using the relations of the maps involved. If e is a path, we will write ϕ_e for the diffeomorphism obtained by pushing a basepoint along e. By definition, using the Cerf decomposition shown, the graph

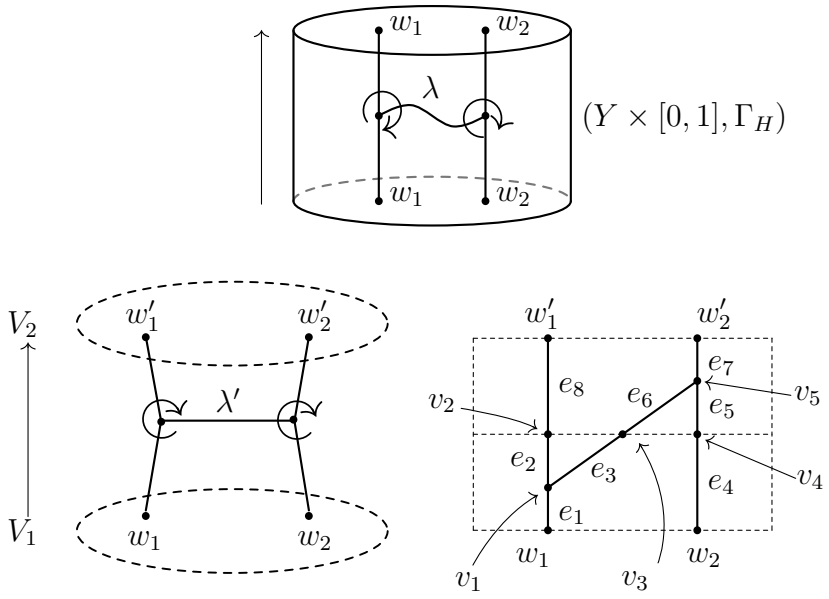

Figure 4.37: **The graph cobordism** $(Y \times [0,1], \Gamma_H)$ **and the ribbon flow graph** $\mathcal{H} = (H, \{w_1, w_2\}, \{w_1', w_2'\})$**, obtained by pushing** H **into** $Y \times \{0\}$**.** In Lemma 4.14.12, we compute cobordism map for the graph cobordism $(Y \times [0,1], \Gamma_H)$, by computing the graph action map for $\mathcal{H}$. A Cerf decomposition of the flow graph is shown on the bottom right. The graph cobordism map for $F_{Y\times[0,1],\Gamma_H,\mathfrak{s}}$ is equal to the graph action map $A_{\mathcal{H},\mathfrak{s}}$, composed with a basepoint moving map to move w_1' to w_1 and w_2' to w_2.

action map takes the form

$$
\begin{aligned}
&S^-_{w_1w_2v_1v_2v_3v_5} A_{e_8} A_{e_7} A_{e_6} A_{e_5} A_{e_4} A_{e_3} A_{e_2} A_{e_1} S^+_{v_1v_2v_3v_4v_5w_1'w_2'} \\
&\quad \simeq (S^-_{v_5} A_{e_7} S^+_{w_2'}) S^-_{v_1v_2v_4} A_{e_8} (S^-_{v_3} A_{e_6} A_{e_3} S^+_{v_3}) A_{e_2} A_{e_1} S^+_{v_2w_1'} (S^-_{w_1} A_{e_1} S^+_{v_1}) (S^-_{w_2} (S^-_{v_4} A_{e_5} A_{e_4} S^+_{v_4}) S^+_{v_5}) \\
&\quad \simeq (\phi_{e_7})_* S^-_{v_1v_2} A_{e_8} A_{e_6*e_3} A_{e_2} S^+_{v_2w_1'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad \simeq U(\phi_{e_7})_* S^-_{v_1v_2} A_{e_8} S^+_{v_2w_1'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* + (\phi_{e_7})_* S^-_{v_1v_2} A_{e_8} A_{e_2} A_{e_6*e_3} S^+_{v_2w_1'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad \simeq U(\phi_{e_7})_* S^-_{v_1} S^+_{w_1'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* + (\phi_{e_7})_* (S^-_{v_1}) S^-_{v_2} A_{e_8} A_{e_2} S^+_{v_2}) S^+_{w_1'} A_{e_6*e_3} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad \simeq U(\phi_{e_7})_* (S^-_{v_1} A_{e_2*e_8} S^+_{v_1}) S^-_{v_1} S^+_{w_1'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* + (\phi_{e_7})_* (\phi_{e_8*e_2})_* A_{e_6*e_3} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad \simeq U(\phi_{e_7})_* S^-_{v_1} A_{e_2*e_8} \Phi_{v_1} S^+_{w_1'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* + (\phi_{e_7})_* (\phi_{e_8*e_2})_* A_{e_6*e_3} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad \simeq U(\phi_{e_7})_* S^-_{v_1} A_{e_2*e_8} S^+_{w_1'} \Phi_{v_1} (\phi_{e_1})_* (\phi_{e_4*e_5})_* + (\phi_{e_7})_* (\phi_{e_8*e_2})_* A_{e_6*e_3} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad \simeq U(\phi_{e_7})_* (\phi_{e_8*e_2})_* \Phi_{v_1} (\phi_{e_1})_* (\phi_{e_4*e_5})_* + (\phi_{e_7})_* (\phi_{e_8*e_2})_* A_{\lambda'} (\phi_{e_1})_* (\phi_{e_4*e_5})_* \\
&\quad = \phi'_* (A_{\lambda'} + U\Phi_{v_1}) \phi_*,
\end{aligned}
$$

where ϕ' and ϕ denote basepoint moving maps for moving along the "vertical" edges of H. completing the proof for $\mathcal{A}_{\mathcal{H},\mathfrak{s}}$. The statement for $B_{\mathcal{H},\mathfrak{s}}$ follows by replacing the A's with B's in the above argument. □

The graph in the previous lemma resembles a ribbon graph core of the $\Sigma_{\mathbf{w}}$ subsurface of the decorated surface for Ψ_z. This is no accident:

Lemma 4.14.13. *Suppose that w_1, z, and w_2 are consecutive basepoints on a link $\mathbb{L}$ (in that order, though w_1 and w_2 need not be distinct), and λ is the subarc of L from w_1 to w_2. Then after we have*

$$\Psi_z|_{V=1} \simeq A_\lambda + U_{w_2}\Phi_{w_2} \simeq B_\lambda + U_{w_1}\Phi_{w_1}.$$

Proof. Write λ as the concatenation of two paths, λ_1 and λ_2, as in Figure 4.38.

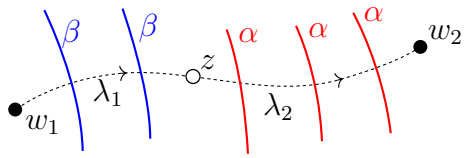

Figure 4.38: **The curve λ is the concatenation of λ_1 and λ_2.**

As λ is a subarc of L, we can assume that λ_1 does not intersect any of the $\boldsymbol{\alpha}$-curves, and that λ_2 does not intersect any of the $\boldsymbol{\beta}$-curves. Hence when adding up the changes across the $\boldsymbol{\alpha}$-curves to compute $a(\lambda, \phi)$, we will count only the differences of intersections of λ_2 with the $\boldsymbol{\alpha}$-curves. Since λ_2 also does not intersect any $\boldsymbol{\beta}$-curves, the quantity $a(\lambda, \phi)$ will just be the total change of the multiplicities of ϕ along λ_2, i.e.,

$$a(\lambda, \phi) = n_{w_2}(\phi) - n_z(\phi).$$

Multiplying $\#\widehat{\mathcal{M}}(\phi)U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}$ by this quantity and summing over all disks, we get that

$$A_\lambda(\mathbf{x}) = U_{w_2}\Phi_{w_2}(\mathbf{x}) + \Psi_z(\mathbf{x}),$$

for this diagram. Analogously, the map B_λ computes differences across the $\boldsymbol{\beta}$-curves, and the same line of reasoning yields

$$B_\lambda(\mathbf{x}) = U_{w_1}\Phi_{w_1}(\mathbf{x}) + \Psi_z(\mathbf{x}).$$

Note that these two formulas are consistent with the relation $A_\lambda + B_\lambda = U_{w_1}\Phi_{w_1} + U_{w_2}\Phi_{w_2}$. □

Analogously, we can compute the reduction of the maps Φ_w, after setting $U = 1$.

Lemma 4.14.14. *If $w \in \mathbf{w}$ is a basepoint of $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, let λ denote the component of $L \setminus \mathbf{z}$ which contains w. Suppose that λ, oriented as a subset of L, goes from z_1 to z_2. Then*

$$\Phi_w|_{U=1} \simeq A_\lambda + V_{z_1}\Psi_{z_1} \simeq B_\lambda + V_{z_2}\Psi_{z_2}.$$

Proof. The proof is identical to the proof of Lemma 4.14.13. □

4.14.5 Proof of Theorem C

We can now prove our stated theorem about the algebraic reductions of $F_{W,F,\mathfrak{s}}$.

Proof of Theorem C. Let (W, F) be a decorated link cobordism with surface with divides $F = (\Sigma, \mathcal{A})$. Let us first consider the $V = 1$ reduction of the link cobordism maps. First note that the composition law holds for both the graph cobordism maps, and the link cobordism maps. Noting also that the claim is obviously true for the 0–handle and 4–handle cobordism maps, we can assume that each component of Σ intersects the incoming boundary and the outgoing boundary nontrivially.

Given a ribbon core $\Gamma(\Sigma_{\mathbf{w}})$ of $\Sigma_{\mathbf{w}}$, we note that by Lemma 4.14.10, we can assume that $\Gamma(\Sigma_{\mathbf{w}})$ has no vertices of valence more than three. Now, assuming that each component of Σ has nonempty incoming and outgoing boundaries, we can decompose (W, F) into a composition $(W_n, F_n) \circ \cdots \circ (W_1, F_1)$ such that each (W_i, F_i) is one of the following:

1. (W_i, F_i) is diffeomorphic to the link cobordism induced by an isotopy of the link and 3–manifold, and $\Gamma(\Sigma_{\mathbf{w}}) \cap W_i$ consists of a collection of paths from the incoming boundary to the outgoing boundary;

2. (W_i, F_i) is a cobordism for attaching a 1–handle, 3–handle, or collection of 2–handles, away from the link. Furthermore F_i is diffeomorphic to a collection of annuli, with

trivial dividing sets. Furthermore, the graph $\Gamma(\Sigma_{\mathbf{w}}) \cap F_i$ consists of a collection of paths from the incoming boundary to the outgoing boundary;

3. (W_i, F_i) is diffeomorphic to a cobordism obtained by attaching a $\mathbf{z}$ band to the link. Furthermore $\Gamma(\Sigma_{\mathbf{w}})$ consists of a collection of paths from the incoming boundary to the outgoing boundary.

4. (W_i, F_i) is equivalent (as an undecorated link cobordism) to a cylindrical link cobordism. The graph $\Gamma(\Sigma_{\mathbf{w}}) \cap W_i$ is an elementary flow graph, of satisfying one of (EFG 1), (EFG 2) or (EFG 3). The subsurface $\Sigma_{\mathbf{w}} \cap W_i$ is a regular neighborhood of $\Sigma_{\mathbf{w}}$.

We now show the claim for a link cobordism satisfying condition (1). In this case, the induced link cobordism is the diffeomorphism map induced by the isotopy ϕ. It is clear that the reduction of the diffeomorphism map on link Floer homology is the reduction of the same diffeomorphism map on Heegaard Floer homology.

We now consider the claim for link cobordisms satisfying condition (2). In this case, the link cobordism maps are defined by using the handle attachment maps. The 1–handle and 3–handle maps on link Floer homology obviously reduce to the 1–handle and 3–handle maps in the graph cobordism maps. The 2–handle maps on link Floer homology count holomorphic triangles representing homology classes of triangles with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$. The graph cobordism maps for $\Gamma(\Sigma_{\mathbf{w}})$ count the same triangles. Hence the $V = 1$ reduction of the link cobordism map $F_{W_i, F_i, \mathfrak{s}|_{W_i}}$ coincides with $F_{W_i, \Gamma(\Sigma_{\mathbf{w}}) \cap W_i, \mathfrak{s}|_{W_i}}$, if (W_i, F_i) satisfies condition (2). Note that the $U = 1$ reduction of the two handle maps is not handled by the above argument. See below, when we consider the $U = 1$ reduction in more detail.

We now consider the claim for link cobordisms satisfying condition (3). In this case, the link cobordism map is a type $\mathbf{z}$ band map. The band map is defined by counting holomorphic triangles by the formula

$$F_{B^{\mathbf{z}}}(\mathbf{x}) = F_{\alpha'\alpha\beta}(\Theta^{\mathbf{w}}_{\alpha'\alpha}, \mathbf{x}),$$

where $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ denotes the generator of top $\mathrm{gr}_{\mathbf{w}}$ degree. After setting $V = 1$, this simply becomes the change of diagrams map associated to moving the $\boldsymbol{\alpha}$ curves, which is the identity map,

on coherent chain complexes. On the other hand, the graph cobordism $(Y \times [0,1], \Gamma(\Sigma_{\mathbf{w}}))$ is the identity graph cobordism, so the induced graph cobordism map is the identity map.

We now consider link cobordisms satisfying condition (4). Given a cobordism (W_i, F_i) satisfying condition (4), we can subdivide the cobordism and add trivial strands to the graph, and replace (W_i, F_i) with a sequence of cobordisms satisfying the following:

5. (W_i, F_i) is equivalent (as an undecorated link cobordism) to a cylindrical link cobordism. The subsurface $\Sigma_{\mathbf{w}} \cap W_i$ is a regular neighborhood of $\Gamma(\Sigma_{\mathbf{w}}) \cap W_i$. The graph $\Gamma(\Sigma_{\mathbf{w}}) \cap W_i$ satisfies one of the following:

   (a) $\Gamma(\Sigma_{\mathbf{w}})$ is an elementary flow graph, of type (2), and the vertex v_0 in the definition of such an elementary flow graph, has valence 1;

   (b) $\Gamma(\Sigma_{\mathbf{w}})$ is the H–shaped graph from Figure 4.37.

The process of manipulating and decomposing link cobordisms (W_i, F_i) satisfying condition (4) into a sequence of cobordisms satisfying condition (5) is shown in Figure 4.39.

Thus, to show that the $V = 1$ reduction of the link cobordism maps agree with the graph cobordism maps for a ribbon graph core of $\Sigma_{\mathbf{w}}$, it is sufficient instead to show the claim for link cobordisms with embedded ribbon graph cores satisfying one of conditions (5a) or (5b).

If (W_i, F_i) satisfies condition (5a), then the link cobordism map is equal to $S^+_{w,z}$ or $S^-_{w,z}$, where w is the basepoint on the single component of $(\Sigma_{\mathbf{w}}) \cap W_i$ which is a bigon. Similarly, the graph cobordism map is easily seen to be S^+_w or S^-_w. From the definition of the quasi-stabilization maps, it is clear that

$$S^{\pm}_{w,z}|_{V=1} = S^{\pm}_w,$$

proving the claim if (W_i, F_i) satisfies condition (5a).

If (W_i, F_i) satisfies condition (5b), then the graph $\Gamma(\Sigma_{\mathbf{w}})$ is the H graph from Figure 4.37 and $\Sigma_{\mathbf{w}}$ is a regular neighborhood of $\Gamma(\Sigma_{\mathbf{w}})$. We can write the link cobordism maps as a composition of two quasi-stabilizations:

$$F_{W,F,\mathfrak{s}} = T^+_{w,z} T^-_{w,z} = \Psi_z,$$

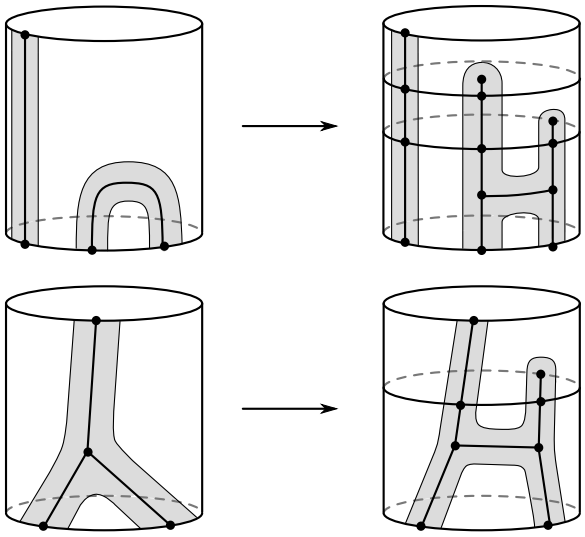

Figure 4.39: **Manipulating and decomposing a link cobordism and ribbon graph core which satisfies** (4) **into a sequence of link cobordisms with ribbon graph cores which satisfy** (5)**.** Shown are only the dividing sets on the surfaces $L \times [0,1]$, as well as the graph $\Gamma(\Sigma_{\mathbf{w}})$. These are inside of the link cobordism $(Y \times [0,1])$ (which is not shown).

where z is the basepoint in bounded by either of the two bigons of type $\Sigma_{\mathbf{z}}$. By Lemma 4.14.13, the $V = 1$ reduction of the link cobordism map is

$$\Psi_z|_{V=1} \simeq B_\lambda + U_{w_1}\Phi_{w_1}.$$

By Lemma 4.14.12, this is equal to the type-B graph cobordism map $F^B_{Y\times[0,1],\Gamma(\Sigma_{\mathbf{W}}),\mathfrak{s}}$, completing the proof of the claim if (W_i, F_i) is satisfies condition (5b).

Having proven the claim for link cobordisms satisfying one of conditions (1)–(4), above, it now follows from the composition law that $F_{W,F,\mathfrak{s}}|_{V=1} \simeq F^B_{W,\Gamma(\Sigma_{\mathbf{W}}),\mathfrak{s}}$ for a general decorated link cobordism (W, F).

We now consider the $U = 1$ reduction of the link cobordism maps. The above strategy is easily adapted to consider these maps, however there are several slight differences. Firstly, the 2–handle maps for $F_{W,F,\mathfrak{s}}$ count holomorphic triangles with underlying homology class ψ satisfying $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, while the 2–handle maps for $F_{W,\Gamma(\Sigma_{\mathbf{z}}),\mathfrak{s}}$ count holomorphic triangles

which satisfy $\mathfrak{s}_{\mathbf{z}}(\psi) = \mathfrak{s}$. However, by Lemma 5.3.3, one has

$$\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi) = PD[\Sigma],$$

where $[\Sigma]$ denotes the homology class of the surface in the link cobordism, viewed as an element of $H_2(W, \partial W; \mathbb{Z})$. Hence, for a 2–handle cobordism, one has

$$F_{W_i,F_i,\mathfrak{s}|_{W_i}}|_{U=1} = F_{W_i,\Gamma(\Sigma_{\mathbf{z}})\cap W_i,(\mathfrak{s}-PD[\Sigma])|_{W_i}}.$$

Finally for link cobordisms satisfying (5b) (but now with $\Sigma_{\mathbf{z}}$ having an H-shaped component, instead of $\Sigma_{\mathbf{w}}$), we note that the induced link cobordism map is Φ_w, for a w basepoint in the boundary of one of the bigon components of $\Sigma_{\mathbf{w}}$. By Lemma 4.14.14, the $U = 1$ reduction of Φ_w is $A_\lambda + V_{z_1}\Psi_{z_1}$, which is the type–$A$ graph cobordism map $F^A_{W_i,\Gamma(\Sigma_{\mathbf{z}})\cap W_i,(\mathfrak{s}-PD[\Sigma])|_{W_i}}$ (note that the change in Spin^c structures, in this case, is simply due to the change in Spin^c structures of the reductions of the link Floer complexes, from Lemma 4.2.8). Hence we conclude that, in general

$$F_{W,F,\mathfrak{s}}|_{U=1} \simeq F^A_{W,\Gamma(\Sigma_{\mathbf{z}}),\mathfrak{s}-PD[\Sigma]}.$$

□

It's clear from the the previous argument that the $F_{W,F,\mathfrak{s}}|_{V=1}$ reduction depends only on the graph $\Gamma(\Sigma_{\mathbf{w}})$ and $\mathfrak{s}$. Similarly it's clear that the reduction $F_{W,F,\mathfrak{s}}|_{U=1}$ depends only on $\Gamma(\Sigma_{\mathbf{z}})$ and $\mathfrak{s} - PD[\Sigma]$. Hence Corollary E follows.

4.14.6 Constructing the graph TQFT from the link Floer TQFT

In this section, we show that given a ribbon graph cobordism (W, Γ) one can always associate a decorated link cobordism $F(\Gamma)$, up to a small ambiguity. We will show that the ambiguity does not effect the maps. It follows that we get an alternate construction and proof of invariance of the graph cobordism maps from Chapter 2.

By Lemma 4.14.3, to a ribbon graph cobordism (W, Γ), we can associate a ribbon surface cobordism (W, R_Γ), which is well defined up to isotopy. The normal bundle to the surface

R_Γ is an oriented 2–plane bundle, over a space homotopy equivalent to a graph, so it is trivializable. We can pick a section v of the normal bundle, and use it to push off a copy R'_Γ of R_Γ (in such a way which keeps $\partial R_\Gamma \cap \operatorname{int} W$ fixed). We define a decorated link cobordism $F(\Gamma)$ to be the union of R_Γ and R'_Γ (giving R'_Γ the opposite orientation). By Theorem C, we have that

$$F_{W,F(\Gamma),\mathfrak{s}}|_{V=1} \simeq F^B_{W,\Gamma,\mathfrak{s}}.$$

The surface $F(\Gamma)$ is not quite well defined, since the vector v is not uniquely specified, even up to isotopy, since a choice of such a vector field v is equivalent to a choice of map from Γ to S^1. The ambiguity thus lies in twisting along the edges of the graph. However the induced link cobordism map for different choices of the vector field v can differ by the diffeomorphism map induced by twisting along doubly based unknot, but the map induced by such a twist is clearly the identity map (one simply considers a diagram where the two basepoints of the knot are right next to each other). Hence in particular the map $F_{W,F(\Gamma),\mathfrak{s}}$ is independent of the choice of vector field v, normal to R_Γ.

CHAPTER 5

Link cobordisms and absolute gradings on link Floer homology

We give an alternate construction of absolute Maslov and Alexander gradings on link Floer homology, using a surgery presentation of the link complement. We show that the link cobordism maps defined by the author are graded and satisfy a grading change formula. We also show that our construction agrees with the original definition of the symmetrized Alexander grading defined by Ozsváth and Szabó. As an application, we compute the link cobordism maps associated to a closed surface in S^4 with a simple dividing set. For simple dividing sets, the maps are determined by the genus of the surface. We show how the grading formula quickly recovers some known bounds on the concordance invariants $\tau(K)$ and $\Upsilon_K(t)$, and use it to prove a new bound for $\Upsilon_K(t)$ for knot cobordisms in negative definite 4–manifolds. As a final application, we recover an adjunction relation for the ordinary Heegaard Floer cobordism maps as a consequence of our grading formula and properties of the link Floer cobordism maps. This gives a link Floer interpretation of adjunction relations and inequalities appearing in Seiberg–Witten and Heegaard Floer theories.

5.1 Introduction

In [OS04e] and [OS06], Ozsváth and Szabó define a collection of invariants, known as Heegaard Floer homology, for closed, oriented 3–manifolds. To a 3–manifold Y and a Spin^c structure, they construct modules $HF^\circ(Y,\mathfrak{s})$ for $\circ \in \{-,+,\infty,\wedge\}$. If $W : Y_1 \to Y_2$ is a cobordism and $\mathfrak{s} \in \text{Spin}^c(W)$ is a Spin^c structure, Ozsváth and Szabó define cobordism

maps from $HF^\circ(Y_1, \mathfrak{s}|_{Y_1})$ to $HF^\circ(Y_2, \mathfrak{s}|_{Y_2})$. An important property of the cobordism maps is that they are graded with respect to an absolute grading, and induce grading change

$$\frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4}.$$

The cobordism maps have particularly nice properties with respect to negative definite 4–manifolds, and the previously mentioned grading change formula gives, for example, a proof of Donaldson's theorem on the intersection forms of negative definite 4–manifolds (see [OS03b]).

There is a refinement of Heegaard Floer homology for knots embedded in a 3–manifold, discovered by Ozsváth and Szabó in [OS04c] and independently by Rasmussen in [Ras03]. Ozsváth and Szabó also define a refinement for links in [OS08]. In Chapter 4, the author describes a collection of cobordism maps forming a "link Floer TQFT", which are defined on a version of the full link Floer complex. In this paper, we consider the grading changes associated to the maps constructed in Chapter 4, and consider also some consequences for several knot invariants defined using knot Floer homology.

5.1.1 Bounds on the concordance invariants $\Upsilon_K(t)$ and $\tau(K)$

In [OS03a], using knot Floer homology Ozsváth and Szabó construct a homomorphism τ from the smooth concordance group to $\mathbb{Z}$. It follows from the bound they prove on τ that if $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is an oriented knot cobordism with $b_1(W) = b_2^+(W) = 0$, then

$$\tau(K_2) \leq \tau(K_1) - \frac{\left|[\Sigma]\right| + [\Sigma] \cdot [\Sigma]}{2} + g(\Sigma) \tag{5.1}$$

where $\left|[\Sigma]\right|$ is the "L^1 norm" of $[\Sigma] \in H_2(W; \mathbb{Z})$. In Section 5.11, we will give a new proof of this result, and see how it follows quickly from a grading change formula for the link cobordism maps from Chapter 4 which we prove in this paper. We state the more general grading change formulas in the next section of the introduction.

Using the previously mentioned grading change formulas for the Alexander and Maslov gradings, we prove a similar bound on the concordance invariant $\Upsilon_K(t)$, defined in [OSS14].

The invariant $\Upsilon_K(t)$ is a piecewise linear function from $[0,2]$ to $\mathbb{R}$. It determines a homomorphism from the concordance group to the group of piecewise linear functions from $[0,2]$ to $\mathbb{R}$.

If s is an integer and $t \in [0,2]$ is a real number, we define the quantity

$$M_t(s) = \max_{a \in 2\mathbb{Z}+1} \frac{-a^2 + 1 + 2ast - 2s^2t}{4}.$$

If W is a negative definite 4–manifold with boundary equal to two copies of S^3, we pick an orthonormal basis $e_1, \dots, e_n$ of $H_2(W;\mathbb{Z})$. If $[\Sigma] \in H_2(W;\mathbb{Z})$ is a homology class, we write $[\Sigma] = s_1 \cdot e_1 + \cdots + s_n \cdot e_n$ for integers s_i and define

$$M_t([\Sigma]) = \sum_{i=1}^{n} M_t(s_i),$$

a piecewise linear function over $t \in [0,2]$. Note that $M_t([\Sigma])$ also has the more invariant description

$$M_t([\Sigma]) = \max_{C \in \mathrm{Char}(W)} \frac{C^2 + b_2(W) - 2t\langle C, [\Sigma]\rangle + 2t([\Sigma]\cdot[\Sigma])}{4},$$

where $\mathrm{Char}(W)$ is the set of characteristic vectors of $H^2(W;\mathbb{Z})$.

We prove the following bound on $\Upsilon_K(t)$:

Theorem 5.1.1. *If $(W,\Sigma) : (S^3, K_1) \to (S^3, K_2)$ is an oriented knot cobordism with $b_1(W) = b_2^+(W) = 0$, then*

$$\Upsilon_{K_2}(t) \geq \Upsilon_{K_1}(t) + M_t([\Sigma]) + g(\Sigma)\cdot(|t-1|-1).$$

For small t, one has $\Upsilon_K(t) = -\tau(K)\cdot t$, and correspondingly for small t our bound reads

$$\Upsilon_{K_2}(t) \geq \Upsilon_{K_1}(t) + t\cdot\left(\frac{\big|[\Sigma]\big| + [\Sigma]\cdot[\Sigma]}{2} - g(\Sigma)\right),$$

reflecting the Ozsváth–Szabó bound on τ in Equation (5.1).

Interestingly, using the computation of $\Upsilon_{T_{n,n+1}}(t)$ in [OSS14] for the $(n, n+1)$–torus knot, we note that in fact

$$M_t(n) = \Upsilon_{T_{n,n+1}}(t).$$

In more generality, the $\Upsilon_K(t)$ invariant is computed for torus knots by Feller and Krcatovich in [FK16]. Comparing our bound with their computation shows that our bound is sharp, when applied to naturally appearing knot cobordisms between positive torus knots.

In Section 5.12 we consider some further applications and examples of the above bound. We give several alternate proofs and extensions of results from [OSS14]. For example if $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is an oriented knot cobordism such that W is a rational homology cobordism, the above theorem immediately yields

$$|\Upsilon_{K_2}(t) - \Upsilon_{K_1}(t)| \leq t \cdot g(\Sigma).$$

As another example we give a new proof of the bound from [OSS14] for $\Upsilon_K(t)$ after crossing changes.

5.1.2 Absolute gradings on link Floer homology and general grading change formulas

The previous theorem follows from a much more general result about the grading changes associated to the link cobordism maps from Chapter 4. We now outline our results about grading changes on link Floer homology.

Definition 5.1.2. An **oriented multi-based link** $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ in Y^3 is an oriented link L with two disjoint collections of basepoints $\mathbf{w} = \{w_1, \dots, w_n\}$ and $\mathbf{z} = \{z_1, \dots, z_n\}$, such that as one traverses the link the basepoints alternate between $\mathbf{w}$ and $\mathbf{z}$. Also, each component of L has at least two basepoints, and each component of Y contains a component of L.

To a multi-based link $\mathbb{L}$ in Y^3 equipped with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, we can construct a "curved" chain complex

$$\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$$

over the ring $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$ which is filtered by $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$. Here $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$ denotes the polynomial ring generated by the variables variables U_w and V_z for $w \in \mathbf{w}$ and $z \in \mathbf{z}$. A curved chain

complex is a module with an endomorphism ∂ squaring to the action of a scalar in the ground ring (sometimes these are called matrix factorizations).

Suppose that J is an index set and L is a link. We call a map $\pi : C(L) \to J$, from the set of components of L to J, a **grading assignment**. We say that L is J–**null-homologous** if

$$[\pi^{-1}(j)] = 0 \in H_1(Y; \mathbb{Z})$$

for all $j \in J$. We define a J–**Seifert surface** of L to be a collection of immersed surfaces S_j (indexed over $j \in J$) such that each S_j has oriented boundary equal to the sublink $-\pi^{-1}(j) \subseteq L$. The components S_j may intersect for various j.

In [OS08], Ozsváth and Szabó, one can define a distinguished Alexander grading on $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ for a link $\mathbb{L}$ in an integer homology sphere Y. However it isn't immediately obvious how to compute the grading change induced by the link cobordism maps from Chapter 4 using their construction. Using surgery presentations of link complements, we provide a description which is more natural from the link cobordism perspective:

Theorem 5.1.3. *(a) If $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multi-based link in Y, with a grading assignment $\pi : C(L) \to J$ such that L is J–null-homologous, and S is a choice of J–Seifert surface, then the modules $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ possess a distinguished absolute Alexander multi-grading $A_{Y,\mathbb{L},S}$ over $\mathbb{Q}^J$.*

(b) If $\mathfrak{s}$ is torsion, the multi-grading $A_{Y,\mathbb{L},S}$ is independent of the J–Seifert surface S. More generally if S and S' are two choices of J–Seifert surfaces, then

$$A_{Y,\mathbb{L},S'}(\mathbf{x})_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = \frac{\langle c_1(\mathfrak{s}), [S'_j \cup -S_j] \rangle}{2}.$$

(c) Under only the assumption that $\mathfrak{s}$ is torsion, there is an absolute Maslov grading $\mathrm{gr}_{\mathbf{w}}$ on $\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$. Under only the assumption that $\mathfrak{s} - PD[L]$ is torsion, there is an absolute Maslov grading $\mathrm{gr}_{\mathbf{z}}$.

(d) If $[L] = 0 \in H_1(Y; \mathbb{Z})$ and $\mathfrak{s}$ is torsion, then all gradings are defined, and $A = \frac{1}{2}(\mathrm{gr}_{\mathbf{w}} - \mathrm{gr}_{\mathbf{z}})$, where A denotes the collapsed Alexander grading.

The $\mathrm{gr}_{\mathbf{w}}$–grading is obtained by paying attention to only the $\mathbf{w}$–basepoints, while the $\mathrm{gr}_{\mathbf{z}}$–grading is obtained by paying attention to only the $\mathbf{z}$–basepoints. For a J–null-homologous link L, the j–component, $(A_{Y,\mathbb{L},\mathfrak{s}})_j$, of the multi-grading takes values in $\mathbb{Z} + \frac{1}{2}\ell k(L \setminus L_j, L_j)$ (see Remark 5.7.1).

The following table shows the grading changes associated to the action of the variables for the various gradings:

Variable	A_j–grading change	$\mathrm{gr}_{\mathbf{w}}$–grading change	$\mathrm{gr}_{\mathbf{z}}$–grading change
U_w	$-\delta(\pi(w), j)$	-2	0
V_z	$+\delta(\pi(z), j)$	0	-2

Here $\delta(\pi(w), j)$ is 1 if w is on a link component assigned grading j by π, and zero otherwise.

We show in Section 5.10 that our construction of the absolute Alexander grading has a symmetry with respect to conjugation which also characterizes the absolute grading defined in [OS08] (for links in integer homology spheres with two basepoints per component). In particular, our construction of the absolute grading agrees with the one due to Ozsváth and Szabó for links in those cases.

As our grading agrees with the one defined by Ozsváth and Szabó in the most interesting cases, the above theorem is not of great interest in itself. Nonetheless, our description of the gradings has the advantage that we can compute the grading change associated to the maps in Chapter 4 induced by link cobordisms. The maps from Chapter 4 use the following notion of morphism between multi-based links in 3–manifolds, originally from [Juh16]:

Definition 5.1.4. We say a pair (W, F) is a **decorated link cobordism** between two 3–manifolds with multi-based links, and write

$$(W, F) : (Y_1, (L_1, \mathbf{w}_1, \mathbf{z}_1)) \to (Y_2, (L_2, \mathbf{w}_2, \mathbf{z}_2)),$$

if

1. W is a 4–dimensional cobordism from Y_1 to Y_2;

2. $F = (\Sigma, \mathcal{A})$ consists of an oriented surface Σ with a properly embedded 1–manifold $\mathcal{A}$ dividing Σ into two disjoint subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, meeting along $\mathcal{A}$;

3. $\partial\Sigma = -L_1 \sqcup L_2$

4. Each component of $L_i \setminus \mathcal{A}$ contains exactly one basepoint;

5. The $\mathbf{w}$–basepoints are all in $\Sigma_{\mathbf{w}}$ and the $\mathbf{z}$–basepoints are all in $\Sigma_{\mathbf{z}}$.

In [Juh16], Juhász gave a construction of link cobordism maps for $\widehat{HFL}$, the hat flavor of link Floer homology, for decorated link cobordisms in the above sense. The construction used the contact gluing map from [HKM08] as well as handle attachment maps similar to the ones from [OS06]. To such a decorated link cobordism, in Chapter 4 the author constructed maps

$$F^{\circ}_{W,F,\mathfrak{s}} : \mathcal{C}FL^{\circ}(Y_1, \mathbb{L}_1, \mathfrak{s}|_{Y_1}) \to \mathcal{C}FL^{\circ}(Y_2, \mathbb{L}_2, \mathfrak{s}|_{Y_2}),$$

which are invariants up to filtered, equivariant chain homotopy. These maps are expected to agree with the maps defined by Juhász when restricted to the hat flavor.

To define functorial maps, we actually need an extra piece of data: a "coloring". A coloring is an assignment of formal variables to the regions of $\Sigma \setminus \mathcal{A}$. Different colorings are useful for different applications. After coloring, the maps become homomorphisms over the ring generated by our collection of formal variables. In Section 5.2.3, we define a notion of a J–graded coloring, which is a coloring that is compatible with a grading assignment over J, and allows for the maps and chain complexes to be graded.

We prove the following grading change formula for the Alexander grading:

Theorem 5.1.5. *Suppose that $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a decorated link cobordism with $F = (\Sigma, \mathcal{A})$ and and $\pi : C(\Sigma) \to J$ is a grading assignment, such that the links on the ends are J–null-homologous with respect to the induced grading assignment. Suppose further that S_1 and S_2 are J–Seifert surfaces for $\mathbb{L}_1$ and $\mathbb{L}_2$, respectively. For a J–graded coloring of F, the map $F^{\circ}_{W,F,\mathfrak{s}}$ is graded with respect to the absolute Alexander multi-grading over J.*

For $j \in J$ the map satisfies

$$A_{Y_2,\mathbb{L}_2,S_2}(F^\circ_{W,F,\mathfrak{s}}(\mathbf{x}))_j - A_{Y_1,\mathbb{L}_1,S_1}(\mathbf{x})_j = \frac{\langle c_1(\mathfrak{s}), \widehat{\Sigma}_j\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]}{2} + \frac{\chi(\Sigma_{\mathbf{w},j}) - \chi(\Sigma_{\mathbf{z},j})}{2},$$

for a homogeneous $\mathbf{x}$. Here $\Sigma_j = \pi^{-1}(j)$, and $\widehat{\Sigma}_j$ is the surface obtained by capping Σ_j with $-(S_1)_j$ and $(S_2)_j$. Similarly $\Sigma_{j,\mathbf{w}}$ and $\Sigma_{j,\mathbf{z}}$ denote the type–$\mathbf{w}$ and –$\mathbf{z}$ regions of $\Sigma_j \setminus \mathcal{A}$.

We also prove grading change formulas for the $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ gradings:

Theorem 5.1.6. *Suppose that $(W,F) : (Y_1,\mathbb{L}_1) \to (Y_2,\mathbb{L}_2)$ is a decorated link cobordism with a J–graded coloring. If $\mathfrak{s}$ is torsion on the ends of W, then the maps $F^\circ_{W,F,\mathfrak{s}}$ are graded with respect to $\mathrm{gr}_{\mathbf{w}}$ and satisfy*

$$\mathrm{gr}_{\mathbf{w}}(F^\circ_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{w}}),$$

for homogeneously graded $\mathbf{x}$. Similarly, if $\mathfrak{s} - PD[\Sigma]$ is torsion on the ends of W, then the maps $F^\circ_{W,F,\mathfrak{s}}$ are graded with respect to $\mathrm{gr}_{\mathbf{z}}$ and satisfy

$$\mathrm{gr}_{\mathbf{z}}(F^\circ_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) = \frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{z}}),$$

for homogeneous $\mathbf{x}$.

In the above, the quantity $\widetilde{\chi}(\Sigma_{\mathbf{w}})$ is defined to be

$$\widetilde{\chi}(\Sigma_{\mathbf{w}}) = \chi(\Sigma_{\mathbf{w}}) - \tfrac{1}{2}(|\mathbf{w}_{\mathrm{in}}| + |\mathbf{w}_{\mathrm{out}}|).$$

The term $\widetilde{\chi}(\Sigma_{\mathbf{z}})$ is defined analogously.

5.1.3 Adjunction relations from link Floer homology

In a different direction, we consider the maps on the ordinary versions of Heegaard Floer homology obtained from the link cobordism maps by ignoring one set of the basepoints. When we do this, several formulas naturally appear which also appear in Seiberg–Witten theory. Using our grading formula, and some basic properties of the maps from Chapter 4, we prove the following adjunction relation for the ordinary cobordism maps in Heegaard Floer homology:

Theorem 5.1.7. *Suppose that Σ is a connected, closed, oriented, embedded surface in a cobordism $W : Y_1 \to Y_2$ with W, Y_1 and Y_2 connected, and that $\mathfrak{s} \in \mathrm{Spin}^c(W)$ is a Spin^c structure with*

$$\langle c_1(\mathfrak{s}), [\Sigma]\rangle - [\Sigma]\cdot[\Sigma] = -2g(\Sigma).$$

Then there is an element $\xi(\Sigma) \in \mathbb{Z}_2[U] \otimes \Lambda^*(H_1(\Sigma;\mathbb{Z}))$, such that

$$F^{\circ}_{W,\mathfrak{s}}(\cdot) = F^{\circ}_{W,\mathfrak{s}-PD[\Sigma]}(\iota_*(\xi(\Sigma)) \otimes \cdot),$$

as maps from $(\mathbb{Z}_2[U] \otimes \Lambda^*(H_1(W;\mathbb{Z})/\mathrm{Tors})) \otimes HF^{\circ}(Y_1, \mathfrak{s}|_{Y_1})$ to $HF^{\circ}(Y_2, \mathfrak{s}|_{Y_2})$ for $\circ \in \{-, +, \infty\}$ and $\iota_* : H_1(\Sigma;\mathbb{Z}) \to H_1(W;\mathbb{Z})$ is the map induced by inclusion. In fact, we can take $\xi(\Sigma) = \prod_{i=1}^{g(\Sigma)}(U + A_i \cdot B_i)$ for any symplectic basis $A_1, \dots, A_n, B_1, \dots B_n$ of $H_1(\Sigma;\mathbb{Z})$, arising from a collection of simple closed curves A_i and B_i with geometric intersection number $|A_i \cap B_j| = \delta_{ij}$.

Note that we don't require that $g(\Sigma) > 0$. The above should be compared to [OS04b, Theorem 3.1] (for $g(\Sigma) > 0$ and Heegaard Floer), [OS00b, Theorem 1.3] (for $g(\Sigma) > 0$ and Seiberg-Witten) and [FS95, Lemma 5.2] (for $g(\Sigma) = 0$ and Seiberg-Witten).

By applying the previous adjunction relation to the identity cobordism $Y \times [0, 1]$, we get an alternate, link Floer theoretic proof of the standard, 3–dimensional adjunction inequality for Heegaard Floer homology from [OS04d], i.e. if $\Sigma \subseteq Y$ is an embedded, oriented surface of genus $g(\Sigma) > 0$ and $HF^+(Y, \mathfrak{s}) \neq 0$, then

$$|\langle c_1(\mathfrak{s}), [\Sigma]\rangle| \leq 2g(\Sigma) - 2.$$

5.1.4 The maps on $\mathcal{H}FL^{\infty}$ for surfaces in negative definite 4–manifolds

Applying the techniques developed in the previous section, we compute the maps on homology associated to knot cobordisms where the underlying 4–manifold is a negative definite and the dividing set is relatively simple.

Theorem 5.1.8. *Suppose that $(W, F) : (S^3, \mathbb{K}_1) \to (S^3, \mathbb{K}_2)$ is a knot cobordism such that $b_1(W) = b_2^+(W) = 0$, $F = (\Sigma, \mathcal{A})$ is a surface with divides such that Σ is connected, and $\mathbb{K}_i$*

are two knots, each with two basepoints. Suppose further that $\mathcal{A}$ consists of two arcs going from $\mathbb{K}_1$ to $\mathbb{K}_2$ and $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ are both connected. Then the induced map on homology

$$F^{\infty}_{W,F,\mathfrak{s}} : \mathcal{H}FL^{\infty}(Y,\mathbb{K}_1) \to \mathcal{H}FL^{\infty}(S^3,\mathbb{K}_2)$$

is an isomorphism. In fact, under the identification $\mathcal{H}FL^{\infty}(S^3,\mathbb{K}_i) \cong \mathbb{Z}_2[U,V,U^{-1},V^{-1}]$, it is the map

$$1 \mapsto U^{-d_1/2}V^{-d_2/2},$$

where

$$d_1 = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} - 2g(\Sigma_{\mathbf{w}})$$

and

$$d_2 = \frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4} - 2g(\Sigma_{\mathbf{z}}).$$

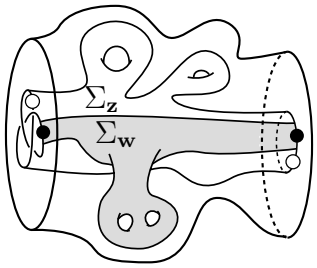

Figure 5.1: **An example of the dividing sets considered in Theorem 5.1.8 and Corollary 5.1.9.** Here $\mathcal{A}$ consists of two arcs, and $\Sigma \setminus \mathcal{A}$ consists of two connected components, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$.

More generally, we can compute explicitly the maps $F^{\infty}_{W,F,\mathfrak{s}}$ on the homology group $\mathcal{H}FL^{\infty}$ in terms of the maps $F^{\infty}_{W,\mathfrak{s}}$ on HF^{∞} when the dividing set is as in the previous theorem, even when we don't have $b_1(W) = b_2^+(W) = 0$, though they don't take quite as simple of a form and we can't guarantee they are isomorphisms (see Remark 5.14.1).

As a consequence of the above theorem, we can compute the maps associated to closed surfaces in S^4.

Corollary 5.1.9. *Suppose that $\Sigma \subseteq S^4$ is a closed, oriented surface, and $\mathcal{A}$ is a simple closed curve on Σ which divides Σ into two connected subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, then the link cobordism map*

$$F^-_{S^4,F,\mathfrak{s}_0} : \mathcal{C}FL^-(\varnothing,\varnothing) \to \mathcal{C}FL^-(\varnothing,\varnothing)$$

is the map

$$1 \mapsto U^{g(\Sigma_{\mathbf{w}})}V^{g(\Sigma_{\mathbf{z}})}.$$

We remark that $\mathcal{C}FL^-(\varnothing,\varnothing)$ is identified with $\mathbb{Z}_2[U,V]$. The cobordism map $F^-_{S^4,F,\mathfrak{s}_0}$ is obtained by puncturing S^4 at two points along the arc $\mathcal{A}$ (i.e. removing two balls and computing the map induced by the knot cobordism from $(S^3,\mathbb{U})$ to $(S^3,\mathbb{U})$).

5.1.5 Further remarks

In [JM16a] Juhász and Marengon compute the cobordism maps from [Juh16] on $\widehat{HFL}$ associated to link cobordisms in $S^3\times[0,1]$ and in doing so, compute the grading change for such link cobordism maps. For link cobordisms in $S^3\times[0,1]$, all terms in our grading formula except the various Euler characteristics of subsets of Σ vanish, and their formula for the change in the Alexander grading and Maslov gradings agrees with ours. See also [JM16b], where their maps are analyzed for concordances.

In the literature, it is more common to use a slightly different version of the full link Floer complex than the "curved link Floer complex" we are using. If $\mathbb{K}=(K,w,z)$ is a doubly based knot in an integer homology sphere Y, one normally considers a $\mathbb{Z}\oplus\mathbb{Z}$–filtered chain complex $CFK^\infty(Y,K)$ over $\mathbb{Z}_2[U,U^{-1}]$. Under our construction of the Alexander grading, we have that

$$CFK^\infty(Y,K) = \mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s}_0)_0,$$

where $\mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s}_0)_0$ denotes the homogeneous subset of zero Alexander grading. Indeed, decomposing over Alexander gradings, we can write

$$\mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s}_0) = \bigoplus_{k\in\mathbb{Z}} CFK^\infty(Y,K)[k].$$

The action of U on $CFK^\infty(Y,K)$ corresponds to the action of $U_w V_z$ on $\mathcal{CFL}^\infty(Y,\mathbb{K},\mathfrak{s}_0)$. Noting that from Theorem 5.1.3 that $A = \frac{1}{2}(\mathrm{gr}_{\mathbf{w}} - \mathrm{gr}_{\mathbf{z}})$, it follows that the two Maslov gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ coincide on $\mathcal{CFL}^\infty(Y,\mathbb{K},\mathfrak{s}_0)_0 = CFK^\infty(Y,K)$, reflecting the usual convention that CFK^∞ has a single Maslov grading.

We note that in general, if (W,F) is a decorated link cobordism, the cobordism map $F^\infty_{W,F,\mathfrak{s}}$ may not send $CFK^\infty(S^3,K_1)$ (the standard complex) to $CFK^\infty(S^3,K_2)$, since the Alexander grading change may be nonzero. Instead, the link cobordism maps will send $CFK^\infty(S^3,K_1)$ to $CFK^\infty(S^3,K_2)[k]$ for some shift $k \in \mathbb{Z}$ in the Alexander grading (i.e. the cobordism map sends monomials $\mathbf{x}\cdot U^i V^j$ with $A(\mathbf{x}) + (j-i) = 0$ to sums of monomials of the form $\mathbf{y}\cdot U^{i'}V^{j'}$ with $A(\mathbf{y}) + (j'-i') = k$). As such, it is often more convenient to work with the curved link Floer complexes.

Finally, we warn the reader that our normalization conventions for the gradings is somewhat different than is common in the literature, in the case that there are more than two basepoints per link component. We normalize so that if $\mathbb{U} = (U,\mathbf{w},\mathbf{z})$ is an unlink in S^3 (with possibly many basepoints), the ranks of $\widehat{HFL}(S^3,\mathbb{U})$ in each Maslov grading are symmetric about zero, and similarly the ranks of each Alexander grading are symmetric about zero. Sometimes authors take the convention that $\widehat{HFL}(S^3,\mathbb{U})$ has top degree generator in degree zero for any number of basepoints per link component, so that Heegaard Floer homology is invariant under adding basepoints to a link component. This is less natural from the perspective of link cobordisms. Our convention ensures for example that cobordisms for connected sums like the one from [HMZ16] are zero graded.

5.1.6 Organization

In Sections 5.2 and 5.3 we provide some background on link Floer homology, link cobordism maps, and also prove a useful Spin^c structure formula for doubly pointed Heegaard triples. In Section 5.4 we describe a set of Kirby moves for describing link complements as surgery on a standard unlink complement in S^3. In Sections 5.5 and 5.6 we describe some preliminaries

about the relative gradings. In Sections 5.7 and 5.8 we define the absolute gradings and prove invariance from the choices in the construction. In Section 5.9 and 5.10 we prove the grading change formula for link cobordisms, and some basic properties, such as conjugation invariance and the equivalence with other constructions. In Sections 5.11 and 5.12 we prove bounds on $\tau(K)$ and $\Upsilon_K(t)$ using our grading formulas. In Section 5.13 we see how the adjunction inequality follows from Alexander grading and link cobordism considerations. In Section 5.14 we compute the link cobordism maps on $\mathcal{H}FL^\infty$ for many negative definite knot cobordisms.

5.2 Background and definitions

In this section we describe background material about the complexes $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$, and some preliminary definitions which will be used throughout the paper.

5.2.1 Curved link Floer homology

Knot Floer homology was originally constructed by Ozsváth and Szabó in [OS04c], and independently by Rasmussen in [Ras03]. Link Floer homology is a generalization to links constructed by Ozsváth and Szabó in [OS08]. The link cobordism maps from Chapter 4 are defined on a variation of the construction from [OS08]. We recall some definitions and preliminary results from Chapter 3 and Chapter 4.

To an oriented multi-pointed link $\mathbb{L}$ in Y, we can construct a multi-pointed Heegaard diagram $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z})$. We consider the two tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ in $\mathrm{Sym}^n(\Sigma)$ defined by

$$\mathbb{T}_\alpha=\alpha_1\times\cdots\times\alpha_n,\qquad\text{and}\qquad\mathbb{T}_\beta=\beta_1\times\cdots\times\beta_n$$

where $n=|\boldsymbol{\alpha}|=|\boldsymbol{\beta}|$. Ozsváth and Szabó define a map $\mathfrak{s}_{\mathbf{w}}:\mathbb{T}_\alpha\cap\mathbb{T}_\beta\to\mathrm{Spin}^c(Y)$ in [OS04e]. For a choice of $\mathfrak{s}\in\mathrm{Spin}^c(Y)$ and under appropriate admissibility assumptions, we consider the $\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]$–module generated by intersection points $\mathbf{x}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta$ with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})=\mathfrak{s}$, which we denote by $\mathcal{C}FL^-(\mathcal{H},\mathfrak{s})$. If we choose an almost complex structure on $\Sigma\times[0,1]\times\mathbb{R}$, then

by counting pseudo-holomorphic curves we can construct an endomorphism ∂, depending on the almost complex structure. Because of nontrivial Maslov index 2 boundary degenerations bubbling off, the map ∂ doesn't square to zero for a general link. Instead we have that

$$\partial^2(\mathbf{x}) = \sum_{K \in C(L)} (U_{w_{K,1}} V_{z_{K,1}} + V_{z_{K,1}} U_{w_{K,2}} + U_{w_{K,2}} V_{z_{K,2}} + \cdots V_{z_{K,n_K}} U_{w_{K,1}}) \cdot \mathbf{x},$$

where $w_{K,1}, z_{K,1}, \ldots, w_{K,n_K}, z_{K,n_K}$ are the basepoints on the link component $K \in C(L)$, in the order that they appear on K (see Lemma 3.2.1). As such, we call $\mathcal{C}FL^\circ$ a "curved" chain complex. If each link component has exactly two basepoints, then $\partial^2 = 0$. Also, if we only have two colors (one for the U–variables and one for the V–variables), then $\partial^2 = 0$ as well. The modules $\mathcal{C}FL^-(\mathcal{H}, \mathfrak{s})$ also have a natural filtration of $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$, which is given by powers of the variables. The modules $\mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s})$ are defined by allowing all powers of the variables. The more standard complexes $\widehat{CFL}(\mathcal{H}, \mathfrak{s})$ are defined by setting all the variables equal to zero.

There are many diagrams $\mathcal{H}$ for a link. If $\mathcal{H}$ and $\mathcal{H}'$ differ by an elementary Heegaard move, then using the constructions of Ozsváth and Szabó, we can associate a filtered, equivariant chain map $\Phi_{\mathcal{H} \to \mathcal{H}'} : \mathcal{C}FL^\circ(\mathcal{H}) \to \mathcal{C}FL^\circ(\mathcal{H}')$. To a general pair of diagrams $\mathcal{H}$ and $\mathcal{H}'$, we can define a map $\Phi_{\mathcal{H} \to \mathcal{H}'}$ as a composition of maps associated to elementary Heegaard moves. The main result of [JT12] implies that the composition is a chain homotopy equivalence and is independent of the choice of intermediate diagrams, up to filtered, equivariant chain homotopy. We define $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ to be the collection of all of the complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$, together with the change of diagrams maps. We call the object $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ the **coherent, equivariant chain homotopy type**. In Chapter 4, we described the natural notion of a morphism between two coherent chain complexes (essentially as a collection of morphisms of chain complexes, which commute up to filtered equivariant chain homotopy with the change of diagrams maps).

5.2.2 Colorings of links and surfaces, and link cobordism maps

To define link cobordism maps as in Chapter 4, we must algebraically modify the complexes described in the previous section. Since the module $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ is a module over $\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]$, a ring which depends on $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$, it doesn't make sense to specify a module homomorphism between the complexes associated to two different links. This ensures that the link cobordism maps are homomorphisms of a fixed module. Also, at many points of the construction in Chapter 4, in order for the maps to be chain maps or to be well defined up to chain homotopy equivalence, one needs to identify certain variables on the complexes. The decoration of the surface with divides encoded exactly which variables needed to be identified to get well defined maps which are chain maps.

To handle these issues, we make the following two definitions:

Definition 5.2.1. A **coloring of a multi-based link** $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ is a pair $(\sigma,\mathfrak{P})$ where $\mathfrak{P}$ is a finite set, indexing a collection of formal variables, and $\sigma:\mathbf{w}\cup\mathbf{z}\to\mathfrak{P}$ is a map.

Analogously, we make the following definition:

Definition 5.2.2. A **coloring of a surface with divides** $(\Sigma,\mathcal{A})$ is a pair $(\sigma,\mathfrak{P})$ where $\mathfrak{P}$ is a finite set, indexing a collection of formal variables, and $\sigma:C(\Sigma_{\mathbf{w}})\cup C(\Sigma_{\mathbf{z}})\to\mathfrak{P}$ is a map from the set of components of $\Sigma\setminus\mathcal{A}$ to $\mathfrak{P}$.

If $(\sigma,\mathfrak{P})$ is a coloring of the multi-based link $\mathbb{L}=(L,\mathbf{w},\mathbf{z})$ in Y and $\circ\in\{-,\infty\}$, we can form the colored complex

$$\mathcal{C}FL^\circ(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$$

by tensoring $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ with a "coloring module" (see Section 3.2), to formally identify the variable $U_w\in\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]$ with $U_{\sigma(w)}\in\mathbb{Z}_2[U_{\mathfrak{P}}]$, and similarly for the $V_{\mathbf{z}}$ variables.

It is important to note that different choices of colorings will be useful for different applications. In the next section, we describe conditions on colorings which will allow us to define Alexander and Maslov gradings.

If (W, F) is a decorated link cobordism and F is given a coloring $(\sigma, \mathfrak{P})$, then the construction of Chapter 4 yields a chain map

$$F^{\circ}_{W,F,\mathfrak{s}} : \mathcal{C}FL^{\circ}(Y_1, \mathbb{L}_1, \sigma|_{Y_1}, \mathfrak{P}, \mathfrak{s}|_{Y_1}) \to \mathcal{C}FL^{\circ}(Y_2, \mathbb{L}_2, \sigma|_{Y_2}, \mathfrak{P}, \mathfrak{s}|_{Y_2}).$$

We usually suppress the coloring from the notation.

5.2.3 Grading assignments, J–null-homologous links, and J–graded colorings of links

In this subsection we provide some definitions necessary to define the gradings. We begin with a first definition:

Definition 5.2.3. If L is an embedded link in Y^3, we say that a pair (π, J) is a **grading assignment** of L if J is a finite set, and $\pi : C(L) \to J$ is a map from the set of components of L to J. Analogously, if Σ is an embedded surface in W, we say that a **grading assignment** of Σ is a map $\pi : C(\Sigma) \to J$.

For notational convenience, if s is a point on L we will sometimes write $\pi(s)$ for $\pi(K)$ where K is the component of L which contains s.

Definition 5.2.4. Suppose that L is a link in Y and $\pi : L \to J$ is a grading assignment. We say that L is J–**null-homologous** if

$$[\pi^{-1}(j)] = 0 \in H_1(Y; \mathbb{Z}),$$

for each $j \in J$.

Analogously, we could consider links L which are **rationally** J–**null-homologous**, i.e.

$$[\pi^{-1}(j)] = 0 \in H_1(Y; \mathbb{Q})$$

for each $j \in J$. For expository reasons we will focus on integrally J–null-homologous links.

Remark 5.2.5. It is important that our links be oriented. A link can be J–null-homologous with one orientation, but not with another. For example consider $L = S^1 \times \{p_1, p_2\} \subseteq S^1 \times S^2$. For the grading assignment sending both components to a single grading, L is J–null-homologous if the two components are given opposite orientations, but is not J–null-homologous if they are given the same orientation.

As we must color the link Floer complexes to get functorial cobordism maps, we also need to describe which colorings will allow for the colored modules $\mathcal{C}FL^\infty(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ to admit Alexander and Maslov gradings.

Definition 5.2.6. If $(L, \mathbf{w}, \mathbf{z})$ is a multi-based link in Y, we say a tuple $(\sigma, \mathfrak{P}, \widehat{\pi}, \pi, J)$ is a J–**graded coloring** if the following holds:

1. $(\sigma, \mathfrak{P})$ is a coloring of $\mathbb{L}$ and (π, J) is a grading assignment of $\mathbb{L}$;
2. the set of colors $\mathfrak{P}$ is partitioned into two subsets $\mathfrak{P} = \mathfrak{P}_{\mathbf{w}} \sqcup \mathfrak{P}_{\mathbf{z}}$;
3. $\sigma(\mathbf{w}) \subseteq \mathfrak{P}_{\mathbf{w}}$ and $\sigma(\mathbf{z}) \subseteq \mathfrak{P}_{\mathbf{z}}$;
4. $\widehat{\pi} : \mathfrak{P} \to J$ is a map such that
$$\pi \circ \iota = \widehat{\pi} \circ \sigma,$$
where $\iota : \mathbf{w} \cup \mathbf{z} \to C(L)$ is the map induced by inclusion.

Example 5.2.7. One natural coloring that is graded is obtained by collapsing all of $U_{\mathbf{w}}$ to a single variable U, and collapsing all of the $V_{\mathbf{z}}$ variables to a single V. The set J is $\{*\}$ and $\widehat{\pi}$ and π map every color and basepoint to $*$.

Another coloring of interest is the "tautological coloring", obtained by setting $\mathfrak{P} = \mathbf{w} \cup \mathbf{z}$, and defining σ to be the identity map. The tautological coloring determines a J–graded coloring with $J = C(L)$, the set of components of L.

Similarly we can define J–graded colorings of surfaces with divides, for which the link cobordism maps of Chapter 4 will be graded and satisfy several grading change formulas:

Definition 5.2.8. Suppose that $F = (\Sigma, \mathcal{A})$ is a surface with divides, and let $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ be the two subsurfaces of Σ, which meet along $\mathcal{A}$, with their designation as type $\mathbf{w}$ or $\mathbf{z}$. A **grading assignment** of F is a map $\pi : C(\Sigma) \to J$. A J–graded coloring of F is a tuple $(\sigma, \mathfrak{P}, \widehat{\pi}, \pi, J)$ such that

1. $(\sigma, \mathfrak{P})$ is a coloring of the components of $\Sigma \setminus \mathcal{A}$ and (π, J) is a grading assignment of F;
2. the set of colors $\mathfrak{P}$ is partitioned into two subsets $\mathfrak{P} = \mathfrak{P}_{\mathbf{w}} \sqcup \mathfrak{P}_{\mathbf{z}}$;
3. $\sigma(C(\Sigma_{\mathbf{w}})) \subseteq \mathfrak{P}_{\mathbf{w}}$ and $\sigma(C(\Sigma_{\mathbf{z}})) \subseteq \mathfrak{P}_{\mathbf{z}}$;
4. $\widehat{\pi} : \mathfrak{P} \to J$ is a map such that
$$\pi \circ \iota = \widehat{\pi} \circ \sigma,$$
where $\iota : C(\Sigma_{\mathbf{w}}) \cup C(\Sigma_{\mathbf{z}}) \to C(L)$ denotes inclusion.

We now define a notion of Seifert surface for J–null-homologous links.

Definition 5.2.9. A J–**Seifert surface** S for a J–null-homologous link L with grading assignment $\pi : L \to J$ is a collection of oriented, immersed surfaces $S_j \subseteq Y$, such that
$$\partial S_j = -\pi^{-1}(j),$$
where $\pi^{-1}(j) \subseteq L$ is given the orientation of L.

The various surfaces S_j of a J–Seifert surface will of course intersect, in general.

We remark that under our orientation conventions, a Seifert surface can be thought of as determining a link cobordism *from* $\pi^{-1}(j)$ *to* the empty link.

5.3 Heegaard triples and Spin^c structures

In this section we describe some constructions using Heegaard triples.

5.3.1 Constructing 4–manifolds and Spin^c structures from Heegaard triples

In [OS04e], Ozsváth and Szabó associate to a singly pointed Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, w)$, a 4–manifold with three boundary components, denoted $X_{\alpha\beta\gamma}$. The same construction works for Heegaard triples with multiple $\mathbf{w}$–basepoints. To a Heegaard triple that has two types of basepoints, $\mathbf{w}$ and $\mathbf{z}$, we describe an embedded surface $\Sigma_{\alpha\beta\gamma}$ inside of the 4–manifold $X_{\alpha\beta\gamma}$.

Definition 5.3.1. We say that a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a **doubly multi-pointed Heegaard triple** if each component of $\Sigma \setminus \boldsymbol{\alpha}$ has exactly one $\mathbf{w}$–basepoint, and exactly one $\mathbf{z}$–basepoint, and the same holds for the $\boldsymbol{\beta}$– and $\boldsymbol{\gamma}$–curves.

Using the construction of Ozsváth and Szabó, we can construct a 4–manifold $X_{\alpha\beta\gamma}$ as the union

$$X_{\alpha\beta\gamma} = \big((\Sigma \times \Delta) \cup (e_\alpha \times U_\alpha) \cup (e_\beta \times U_\beta) \cup (e_\gamma \times U_\gamma)\big) / \sim,$$

where Δ is the triangle with edges (in clockwise order) e_α, e_β and e_γ. Also U_α, U_β and U_γ are genus $g(\Sigma)$ handlebodies with boundary Σ, which have $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ as the boundaries of compressing curves (respectively). The 4–manifold $X_{\alpha\beta\gamma}$ naturally has boundary

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

If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a doubly multi-pointed Heegaard triple with two types of basepoints, then there is a naturally defined, properly embedded surface $\Sigma_{\alpha\beta\gamma}$ inside of $X_{\alpha\beta\gamma}$. It is formed as the union

$$\Sigma_{\alpha\beta\gamma} = ((-\mathbf{w} \cup \mathbf{z}) \times \Delta) \cup (e_\alpha \times f_\alpha) \cup (e_\beta \times f_\beta) \cup (e_\gamma \times f_\gamma).$$

Here $f_\alpha, f_\beta, f_\gamma$ are compact 1–manifolds, defined as follows. Inside of the handlebody U_α, one takes a Morse function which is 0 on Σ, has maximum value of 1, such that the intersection of Σ with the descending manifolds of the index 2 critical points of Σ are the $\boldsymbol{\alpha}$ curves. One then takes f_α to be the union of the (finitely many) flowlines which intersect a basepoint in $\mathbf{w} \cup \mathbf{z} \subseteq \Sigma$. The 1–manifolds f_β and f_γ are defined similarly.

Note that the ends of $X_{\alpha\beta\gamma}$ are $-Y_{\alpha\beta}, -Y_{\beta\gamma}$ and $Y_{\alpha\gamma}$. Inside of each of these closed three manifolds is a link, defined by $\mathbb{L}_{\tau\sigma} = f_\tau \cup f_\sigma$. The link $\mathbb{L}_{\tau\sigma}$ is oriented in $Y_{\tau\sigma}$ by assuming that the intersections with Σ are negative at the $\mathbf{w}$–basepoints, and positive at the $\mathbf{z}$–basepoints. We include a picture in Figure 5.2.

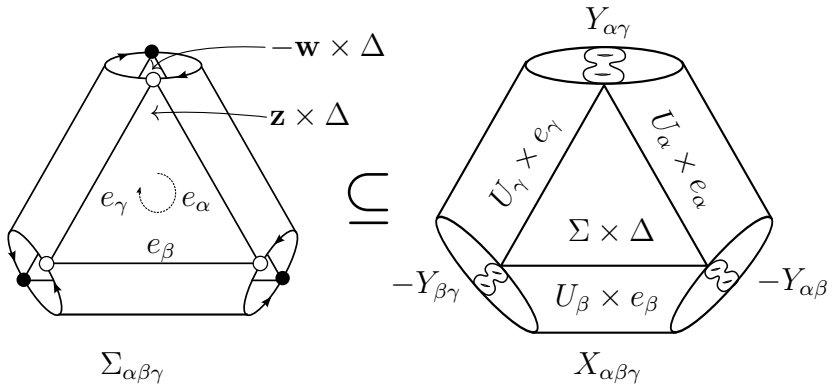

Figure 5.2: **The surface $\Sigma_{\alpha\beta\gamma}$ inside of $X_{\alpha\beta\gamma}$.** Orientations are shown. Note that the picture is somewhat simplistic, since there could be many basepoints, and the number of components on the ends need not be equal.

Using the orientation induced by an outward pointing normal vector N, (i.e. $v \in TL$ is positive if (v, N) is a positive basis of $T\Sigma_{\alpha\beta\gamma}$) we have that

$$\partial\Sigma_{\alpha\beta\gamma} = -L_{\alpha\beta} \sqcup -L_{\beta\gamma} \sqcup L_{\alpha\gamma}.$$

Now if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a triple which is subordinate to β–bouquet of a framed 1–dimension link $\mathbb{S}_1 \subseteq Y_{\alpha\beta} \setminus L_{\alpha\beta}$, then the cobordism $W(Y_{\alpha\beta}, \mathbb{S}_1)$ can be obtained from $X_{\alpha\beta\gamma}$ by filling in $Y_{\beta\gamma}$ with 3– and 4–handles (see Section 5.3.2 for the definition of a β–bouquet). Inside of $W(Y_{\alpha\beta}, \mathbb{S}_1)$, there is a surface $L_{\alpha\beta} \times [0,1]$, which is equal to $\Sigma_{\alpha\beta\gamma}$ after removing the intersection with the 3–handles and 4–handles (see [Juh16, Proposition 6.6])

We now make a remark on orientations. We note that our orientation conventions specify that $Y \times [0,1]$, as a cobordism from $Y \times \{0\}$ to $Y \times \{1\}$, gets the product orientation. We make the same definition for surfaces. The reader should beware of the following though. To define an orientation on the 4–manifold $X_{\alpha\beta\gamma}$, the triangle Δ, viewed as a triangle with edges

e_α, e_β, and e_γ (appearing in that order, counterclockwise), is given the clockwise orientation, as a manifold with boundary. This is somewhat unfortunate, as the triangle maps counts holomorphic maps into $\Sigma \times \Delta$, but Δ is given the counterclockwise orientation, in that context. Hopefully this isn't too confusing. It is important to make this distinction, since it affects formulas arising from intersection numbers.

Let $\mathcal{P}_{\alpha\beta\gamma}$ denote the set 2–chains on a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ which have boundary equal to a linear combination of $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ curves. We can define a map from $\mathcal{P}_{\alpha\beta\gamma}$ to $H_2(X_{\alpha\beta\gamma}; \mathbb{Z})$ as in [OS04e, Proposition 8.3] or [OS06, pg. 9]. This is defined by taking such a domain $\mathcal{D}$ in Σ, and including it into $\{pt\} \times \Sigma \subseteq X_{\alpha\beta\gamma}$. We then extend the 2–chain outward towards the boundary of Δ, and then cap off with disks in U_α, U_β and U_γ to get a closed 2–chain $H(\mathcal{D}) \in H_2(X_{\alpha\beta\gamma}; \mathbb{Z})$. For Heegaard triples with a single basepoint, such domains which also don't pass over the basepoints are usually called **triply periodic domains**. In this paper we will call any domain in $\mathcal{P}_{\alpha\beta\gamma}$ a triply periodic domain if it has boundary equal to a sum of the $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ curves. We note that it is easy to compute the intersection of $H(\mathcal{D})$ and $\Sigma_{\alpha\beta\gamma}$, since they intersect only at $(\mathbf{w} \cup \mathbf{z}) \times \{pt\}$, and the multiplicity of the intersection points is given by the multiplicity of the domain $\mathcal{D}$ at the basepoints. Noting the above remarks about orientations, we have that:

$$\langle PD[H(\mathcal{D})], [\Sigma_{\alpha\beta\gamma}] \rangle = \#(H(\mathcal{D}) \cap \Sigma_{\alpha\beta\gamma}) = n_{\mathbf{z}}(\mathcal{D}) - n_{\mathbf{w}}(\mathcal{D}), \tag{5.2}$$

and indeed, more generally, if Σ_j is a union of connected components of $\Sigma_{\alpha\beta\gamma}$, then the analogous equation holds, if we only sum over the basepoints in $\mathbf{w}$ and $\mathbf{z}$ corresponding to Σ_j.

In [OS04e], to a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ for Y, Ozsváth and Szabó describe a map

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

One starts with the upward gradient vector field for a Morse function inducing the Heegaard diagram. In a neighborhood of the flowlines passing through $\mathbf{w}$, and the intersection points of $\mathbf{x}$, one modifies the vector field so that it is nonvanishing (which can be done since these

chosen flowlines connect critical points of opposite parities). This of course depends on the choice of basepoints $\mathbf{w}$, though we have the following:

Lemma 5.3.2. *If $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, (L, \mathbf{w}, \mathbf{z}))$, and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, then*

$$\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) - \mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = PD[L].$$

Proof. This follows from [OS04e, Lemma 2.19]. □

Analogously, to a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$, in [OS04e, Section 8] Ozsváth and Szabó also describe a map

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

Again this depends on the choice of basepoints $\mathbf{w}$. As an analog of the above lemma, we have the following:

Lemma 5.3.3. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a doubly multi-pointed Heegaard triple, then*

$$\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi) = PD[\Sigma_{\alpha\beta\gamma}].$$

Proof. The proof is analogous to [OS04e, Lemma 2.19]. See [OS04e, Section 8] for a precise description of the map $\mathfrak{s}_{\mathbf{w}} : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathrm{Spin}^c(X_{\alpha\beta\gamma})$. Ozsváth and Szabó associate an oriented two plane field to a triangle ψ on $X_{\alpha\beta\gamma}$, on the complement of a contractible set. In Figure 5 of [OS04e], Ozsváth and Szabó describe a singular codimension 1 foliation on $X_{\alpha\beta\gamma}$. Let $B \subseteq \operatorname{int} \Delta$ be a ball containing the image of the singular leaf of the foliation under the projection map $\Sigma \times \Delta \to \Delta$. Upon inspection, the two 2–plane fields associated to $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ differ only on a neighborhood of $\Sigma_{\alpha\beta\gamma}$ (though they are not defined on all of a neighborhood of $\Sigma_{\alpha\beta\gamma}$). Instead, they are defined on a neighborhood of a surface

$$\Sigma^0_{\alpha\beta\gamma} = \Sigma_{\alpha\beta\gamma} \setminus (\Sigma \times B).$$

Define also

$$X^0_{\alpha\beta\gamma} = X_{\alpha\beta\gamma} \setminus (\Sigma \times B).$$

By construction, on each leaf of the foliation on $X_{\alpha\beta\gamma} \setminus (\Sigma \times B)$ shown in Figure 5 of [OS04e], the 2–plane field of $\mathfrak{s}_{\mathbf{w}}(\psi)$ is the orthogonal complement of $T\Sigma^0_{\alpha\beta\gamma}$ inside of the leaf of the foliation. In fact, we can identify each leaf of the foliation away from $\Sigma \times B$ either as one of the closed manifolds $Y_{\sigma\tau}$ or one of the manifolds with boundary U_τ, for various τ and σ. On the other hand, the 2–plane field of $\mathfrak{s}_{\mathbf{z}}(\psi)$ is given by the orthogonal complement of Reeb surgery along the intersection of $\Sigma_{\alpha\beta\gamma}$ with the leaves of the foliation on $X_{\alpha\beta\gamma}$. This is shown in Figure 5.3. Let $\xi_{\mathbf{w}}$ and $\xi_{\mathbf{z}}$ denote the oriented 2–plane fields for $\mathfrak{s}_{\mathbf{w}}(\psi)$ and $\mathfrak{s}_{\mathbf{z}}(\psi)$ used to define the Spin^c structures. The vector bundles $\xi^\perp_{\mathbf{w}}$ and $\xi^\perp_{\mathbf{z}}$, being oriented 2–plane fields, naturally obtain the structure of a complex line bundle, and thus each of the pairs $(\xi_{\mathbf{w}}, \xi^\perp_{\mathbf{w}})$ and $(\xi_{\mathbf{z}}, \xi^\perp_{\mathbf{z}})$ give $TX_{\alpha\beta\gamma}|_{X_{\alpha\beta\gamma}\setminus C}$ an almost complex structure (for a contractible subset $C \subseteq X_{\alpha\beta\gamma}$), which uniquely determines a Spin^c structure extending over all of $X_{\alpha\beta\gamma}$ by an obstruction theory argument. The first Chern class of $\mathfrak{s}_{\mathbf{w}}(\psi)$ or $\mathfrak{s}_{\mathbf{z}}(\psi)$ evaluated on submanifolds $F \subseteq X_{\alpha\beta\gamma}$, which the 2–plane field is defined on, can be computed as the first Chern class of complex bundle $TX_{\alpha\beta\gamma}|_F$, evaluated on the fundamental class of F.

Since the 2–plane fields used to define $\mathfrak{s}_{\mathbf{w}}(\psi)$ and $\mathfrak{s}_{\mathbf{z}}(\psi)$ agree outside of a neighborhood of $\Sigma_{\alpha\beta\gamma}$, we know that $\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi)$ is an element of $H^2(X_{\alpha\beta\gamma}; \mathbb{Z})$ which is supported in a neighborhood of $\Sigma_{\alpha\beta\gamma}$.

As described in [?LCylindrical, Lemma 10.5], the restriction map Spin$^c(X_{\alpha\beta\gamma}) \to$ Spin$^c(X^0_{\alpha\beta\gamma})$ is injective, and since restriction of Spin^c structures is natural with respect to the action of H^2, it is sufficient to show that $(\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi))|_{N(\Sigma^0_{\alpha\beta\gamma})}$ differ by the class $PD[\Sigma^0_{\alpha\beta\gamma}] \in H^2(N(\Sigma^0_{\alpha\beta\gamma}), \partial N(\Sigma^0_{\alpha\beta\gamma}); \mathbb{Z})$, where $N(\Sigma^0_{\alpha\beta\gamma})$ is closed unit disk bundle of the normal bundle, and $\partial N(\Sigma^0_{\alpha\beta\gamma})$ is the unit sphere bundle. The Thom isomorphism theorem yields an isomorphism $H^2(N(\Sigma^0_{\alpha\beta\gamma}), \partial N(\Sigma^0_{\alpha\beta\gamma}); \mathbb{Z}) \cong H^0(\Sigma^0_{\alpha\beta\gamma}; \mathbb{Z})$. Noting that $H^*(\Sigma^0_{\alpha\beta\gamma}; \mathbb{Z})$ is torsion free, if $[F]$ is any element of $H_2(N(\Sigma^0_{\alpha\beta\gamma}), \partial N(\Sigma^0_{\alpha\beta\gamma}); \mathbb{Z})$, represented by an embedded surface F, transverse to $\Sigma^0_{\alpha\beta\gamma}$, it is sufficient to show that

$$(\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi))|_{N(\Sigma^0_{\alpha\beta\gamma})}[F] = \#(\Sigma^0_{\alpha\beta\gamma} \cap F).$$

Noting as well that $c_1(\mathfrak{s}_{\mathbf{o}}(\psi) + H) = c_1(\mathfrak{s}_{\mathbf{o}}(\psi)) + 2H$ for $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, it is further sufficient to

show that

$$(c_1(\mathfrak{s}_{\mathbf{w}}(\psi)) - c_1(\mathfrak{s}_{\mathbf{z}}(\psi)))|_{(N(\Sigma^0_{\alpha\beta\gamma}),\partial N(\Sigma^0_{\alpha\beta\gamma}))}[F] = 2\#(\Sigma^0_{\alpha\beta\gamma} \cap F).$$

By what we've said, this will imply that $\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi) = PD[\Sigma_{\alpha\beta\gamma}]$. This can now be performed as a model computation.

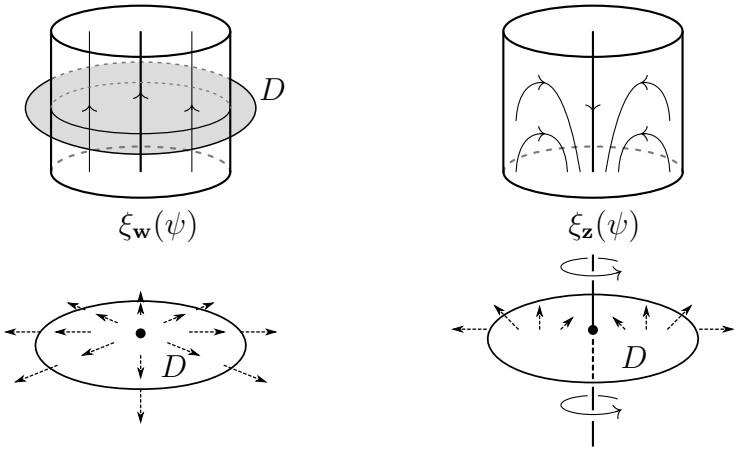

Figure 5.3: **The 2–plane fields $\xi_{\mathbf{w}}(\psi)$ and $\xi_{\mathbf{z}}(\psi)$** are the orthogonal complements of the vector fields shown on the top row, which are shown inside of a (3–dimensional) leaf of the foliation.The disk $D \subseteq F$ is a neighborhood of the intersection point in $F \cap \Sigma^0_{\alpha\beta\gamma}$. On the bottom row, we show a sections of $\xi_{\mathbf{w}}(\psi)$ and $\xi_{\mathbf{z}}(\psi)$ on D. The section for $\xi_{\mathbf{w}}(\psi)$ is shown only on a diameter of D. To obtain it on all of D, one just rotates about the axis shown.

Suppose that F intersects $\Sigma^0_{\alpha\beta\gamma}$ transversely, and near any intersection points is tangent (positively or negatively) to the 2–plane field $\xi_{\mathbf{w}}$. Let $D \subseteq F$ be a disk centered at one of these intersection points. We can pick sections $v_{\mathbf{w}}$ and $v_{\mathbf{z}}$ of the 2–plane bundles $\xi_{\mathbf{w}}(\psi)$ and $\xi_{\mathbf{z}}(\psi)$ (respectively) which agree on ∂D, as follows. For $\xi_{\mathbf{w}}(\psi)$ we pick simply the outward radial vector field on D. For $\xi_{\mathbf{z}}(\psi)$, we take the section shown on a radius of D in Figure 5.3, and rotate around D to form a section $v_{\mathbf{z}}$. To get a well define section at the origin, one must scale $v_{\mathbf{z}}$ near the origin so that it vanishes there. The two sections $v_{\mathbf{w}}(\psi)$ and $v_{\mathbf{z}}(\psi)$ agree on ∂D, but the image of $v_{\mathbf{w}}$ in the total space of $\xi_{\mathbf{w}}(\psi)$ intersects the zero section transversely once, with sign $+1$ if the tangent space of D agrees with $\xi_{\mathbf{w}}(\psi)$ and with sign -1 if the tangent space of D is opposite the orientation of $\xi_{\mathbf{w}}(\psi)$. On the other hand, $v_{\mathbf{z}}$

intersects the zero section transversely once, but with opposite sign. Also note that one can (carefully) check that D has the same orientation as the 2–plane field for $\xi_{\mathbf{w}}(\psi)$ iff the corresponding point $F \cap \Sigma^0_{\alpha\beta\gamma}$ is a positive intersection point).

Finally, we note that $\xi_{\mathbf{w}}(\psi)^{\perp}$ and $\xi_{\mathbf{z}}(\psi)^{\perp}$ both admit the same non-vanishing section on a neighborhood of $\Sigma^0_{\alpha\beta\gamma}$ (a vector in the e_α, e_β, or e_γ direction), and hence the contributions to the Chern classes from these 2–plane bundles are equal.

We conclude that $(c_1(\mathfrak{s}_{\mathbf{w}}(\psi)) - c_1(\mathfrak{s}_{\mathbf{z}}(\psi)))[F] = 2\#(F \cap \Sigma^0_{\alpha\beta\gamma})$ for any $F \in H_2(N(\Sigma^0_{\alpha\beta\gamma}), \partial N(\Sigma^0_{\alpha\beta\gamma}); \mathbb{Z})$, and we view $(c_1(\mathfrak{s}_{\mathbf{w}}(\psi)) - c_1(\mathfrak{s}_{\mathbf{z}}(\psi)))$ as a class in $H^2(N(\Sigma^0_{\alpha\beta\gamma}); \partial N(\Sigma^0_{\alpha\beta\gamma}); \mathbb{Z})$. We conclude

$$(\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi))|_{X^0_{\alpha\beta\gamma}} = PD[\Sigma^0_{\alpha\beta\gamma}].$$

Hence

$$\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi) = PD[\Sigma_{\alpha\beta\gamma}],$$

as elements of $\mathrm{Spin}^c(X_{\alpha\beta\gamma})$. □

5.3.2 Heegaard triples and bouquets of framed 1–dimensional links

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multi-based link in Y. Let $\mathbb{L}_\alpha$ denote the arcs of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which go from $\mathbf{w}$–basepoints to $\mathbf{z}$–basepoints. Let $\mathbb{L}_\beta$ denote the arcs of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which go from $\mathbf{z}$–basepoints to $\mathbf{w}$–basepoints.

Definition 5.3.4. A β–**bouquet** $\mathcal{B}^\beta$ for a framed link $\mathbb{S}_1$ in $Y \setminus L$ is a collection of arcs which connect links in $\mathbb{S}_1$ to the interior of $\mathbb{L}_\beta$. We assume that there is exactly one arc per component of $\mathbb{S}_1$, and each arc has one endpoint on $\mathbb{S}_1$ and one endpoint in $\mathbb{L}_\beta$.

An α–bouquet can be defined analogously. We will use β–bouquets to define the grading, but we will show in Section 5.7.3 that either α–bouquets or β–bouquets can be used.

Given a β–bouquet $\mathcal{B}_\beta$ for a framed link $\mathbb{S}_1 \subseteq Y \setminus L$, we can consider the sutured manifold $Y(\mathcal{B}_\beta)$ obtained by removing regular neighborhoods of $L \cup \mathcal{B}_\alpha$ and $L \cup \mathcal{B}_\beta$, and adding sutures corresponding to the basepoints $\mathbf{w}$ and $\mathbf{z}$. We have the following definition:

Definition 5.3.5. We say that that a triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to the β–bouquet $\mathcal{B}_\beta$ for a framed 1–dimensional link $\mathbb{S}_1 \subseteq Y \setminus L$ if

1. $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_{k+1}, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ yields a sutured diagram for $Y(\mathcal{B}_\beta)$;
2. The curves $\beta_1, \dots, \beta_k$ are each meridians of a different component of $\mathbb{S}_1$ and hence $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;
3. γ_{k+1} through γ_n are each small Hamiltonian isotopies of the curves $\beta_{k+1}, \dots, \beta_n$ respectively, with $|\beta_i \cap \gamma_j| = 2\delta_{ij}$;
4. for $i = 1, \dots, k$, the curve γ_i is induced by the framing of the corresponding link component that β_i is a meridian of.

For moving between Heegaard triples subordinate to a fixed bouquet, we have the following:

Lemma 5.3.6. *Any two β–bouquets for $\mathbb{S}_1$ can be connected by the following moves:*

1. *isotopies and handleslides amongst the $\boldsymbol{\alpha}$–curves;*
2. *isotopies and handleslides amongst the the curves $\{\beta_{k+1}, \dots, \beta_n\}$, while performing the analogous isotopy or handleslide of the curve $\{\gamma_{k+1}, \dots, \gamma_n\}$;*
3. *$(1,2)$–stabilizations or destabilizations;*
4. *for $1 \leq i \leq k$, an isotopy of β_i, or a handleslide of β_i across a β_j with $k+1 \leq j \leq n$;*
5. *for $1 \leq i \leq k$, an isotopy of γ_i, or a handleslide of γ_i across a γ_j with $k+1 \leq j \leq n$;*
6. *A self diffeomorphism of $(Y, L \cup \mathcal{B}^\beta, \mathbf{w} \cup \mathbf{z})$ isotopic to the identity, through such diffeomorphisms.*

Proof. This follows immediately from [Juh16, Lemma 6.5]. □

Remark 5.3.7. Note that in general the surface $\Sigma_{\alpha\beta\gamma}$ associated to a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ may be complicated and have high genus, though if $\mathcal{T}$ is subordinate to an α– or β–bouquet for a framed 1–dimensional link, then the surface $\Sigma_{\alpha\beta\gamma}$ is diffeomorphic to a collection of triply punctured spheres.

5.4 Kirby calculus for manifolds with boundary

Our strategy for constructing an absolute grading on $\mathcal{C}FL^\circ$ will parallel the construction of the absolute $\mathbb{Q}$–gradings on the groups $HF^\circ(Y, \mathfrak{s})$ in [OS06]. We will define a notion of a Kirby diagram for 3–manifolds with an embedded, multi-based link $(Y, \mathbb{L})$, essentially by presenting the link complements $Y \setminus N(\mathbb{L})$ as surgery on the standard unlink complement $S^3 \setminus N(\mathbb{U})$ for an unlink $\mathbb{U}$ in S^3, then consider a Kirby calculus argument for how two such presentations differ.

It will be useful for our purposes to first define a more general notion of surgery presentations, not specific to link complements:

Definition 5.4.1. If M and M' are oriented 3–manifolds with boundary and $\phi : \partial M \to \partial M'$ is a fixed, orientation preserving diffeomorphism, we say that **parametrized surgery data** for (M, M', ϕ) is a pair $(\mathbb{S}_1, f)$ where $\mathbb{S}_1 \subseteq \operatorname{int} M$ is a framed link and

$$f : M(\mathbb{S}_1) \to M'$$

is a diffeomorphism which extends ϕ.

It is not hard to see that if M and M' are connected, oriented 3–manifolds and $\phi : \partial M \to \partial M'$ is an orientation preserving diffeomorphism, then there exists parametrized surgery data for (M, M', ϕ). This can be seen by the following argument (cf. [Rob97]). Using the diffeomorphism ϕ, we form the closed, oriented three manifold $-M \cup (\partial M \times [0, 1]) \cup M'$. This bounds a compact oriented 4–dimensional manifold W. We can view such a manifold as a cobordism of manifolds with boundary from M to M', (a "special" cobordism in the language of [Juh16]). We think of M and M' as the "horizontal" parts of ∂W and $[0, 1] \times \partial M$

as the "vertical" part of the boundary. We can find a Morse function which is 0 on M, t on $\partial M \times [0,1]$ and 1 on M', which has only index 1,2, and 3 critical points. Using a standard trick we can replace index 1 and 3 critical points with index 2 critical points (by changing the 4–manifold), to get a cobordism from M to M' which has a Morse function with only index 2 critical points. If we take a gradient like vector field on W which is $\partial/\partial t$ on $\partial M \times [0,1]$, then the descending manifolds from the index 2 critical points yield a framed link $\mathbb{S}_1$ in M, and the Morse function and gradient like vector field determine a diffeomorphism $f : M(\mathbb{S}_1) \to M'$ which is well defined up to isotopy and extends ϕ.

In [Kir78], Kirby creates a set of moves between any two parametrized surgery decompositions of the form $(S^3, Y, \varnothing)$ for a closed, oriented 3–manifold Y. The moves between any two decompositions were blow-ups, blow-downs, isotopies of f or $\mathbb{S}_1$, and handleslides. In [FR79], Fenn and Rourke extended the calculus to arbitrary oriented 3–manifolds, instead of just S^3, though an extra move is required which is supported in a solid torus. In [Rob97], this is extended to oriented 3–manifolds with boundary, as long as a diffeomorphism of the boundaries is fixed.

We note that in our definition of parametrized surgery data, there is a choice of diffeomorphism $f : M(\mathbb{S}_1) \to M'$. The diffeomorphism is important for our purposes. We note that a Kirby move between two framed links $\mathbb{S}_1$ and $\mathbb{S}_1'$ in M canonically yields a diffeomorphism $M(\mathbb{S}_1) \to M(\mathbb{S}_1')$ as we describe in the following paragraph. In particular, if $(\mathbb{S}_1, f)$ is parametrized surgery data, and $\mathbb{S}_1'$ is the result of one of the above moves on $\mathbb{S}_1$, then there is a diffeomorphism $f' : M(\mathbb{S}_1') \to M'$ which is canonically specified. In such a way, blow-ups, blow-downs and handleslides determine moves not only between framed links, but instead between collections of parametrized surgery data.

We now describe the canonical diffeomorphism from $M(\mathbb{S}_1)$ to $M(\mathbb{S}_1')$ resulting from a Kirby move. A good account can be found in [GS99, pg. 160]. Suppose for the sake of demonstration that we are in the case of standard Kirby calculus, so $M = S^3$ and $\partial M = \partial M' = \varnothing$. Suppose that $\mathbb{S}_1$ and $\mathbb{S}_1'$ are two framed links, such that $\mathbb{S}_1' = \mathbb{S}_1 \cup \{U\}$, where U is a ± 1 framed unknot which is contained in a ball $B \subseteq S^3 \setminus \mathbb{S}_1$. In this case, the manifolds

$S^3(\mathbb{S}_1) \setminus B = S^3(\mathbb{S}'_1) \setminus B(U)$ are canonically diffeomorphic, via the identity map. Noting that B and $B(U)$ are both 3–balls with an identification of their boundaries (via the identity map), our diffeomorphism of $S^3(\mathbb{S}_1) \setminus B$ and $S^3(\mathbb{S}'_1) \setminus B(U)$ can be extended over B, uniquely up to isotopy, since any diffeomorphism of the 2–sphere extends over the 3–ball, uniquely up to isotopy. Hence we get a canonical diffeomorphism from $S^3(\mathbb{S}_1)$ to $S^3(\mathbb{S}'_1)$. Similarly if $\mathbb{S}'_1$ is the result of handlesliding a link component in $\mathbb{S}_1$ across another link component, then a diffeomorphism from $S^3(\mathbb{S}'_1)$ to $S^3(\mathbb{S}_1)$ is canonically specified as follows. Note that $S^3(\mathbb{S}'_1) \setminus H = S^3(\mathbb{S}_1) \setminus H'$, where H and H' are genus 2 handlebodies containing the link components involved in the handleslide. Now ∂H and $\partial H'$ are canonically identified, and H and H' are diffeomorphic by a diffeomorphism which extends this identification. But since H and H' are handlebodies, a diffeomorphism from H to H' is canonically specified up to isotopy by its restriction to the boundary since $\mathrm{MCG}(H_g, \partial H_g) = \{*\}$, for H_g a 3–dimensional, genus g handlebody. In such a manner, a sequence of Kirby moves between two framed links $\mathbb{S}_1$ and $\mathbb{S}'_1$ canonically determines a diffeomorphism $M(\mathbb{S}_1)$ and $M(\mathbb{S}'_1)$. As a specific example, a diffeomorphism $\psi : M \to M$ which is the identity on ∂M may be presented as a sequence of Kirby moves on framed links in M, starting and ending at the empty link in $\operatorname{int} M$.

We state the following version of the result from [Rob97]:

Theorem 5.4.2. *Any two pairs of parametrized surgery data for (M, M', ϕ) can be connected by the following moves:*

1. *(Move $\mathcal{O}_0$): isotopies of f or $\mathbb{S}_1$ which fix ∂M;*
2. *(Move $\mathcal{O}_1$): handleslides of link components amongst each other;*
3. *(Move $\mathcal{O}_2$): blow ups or blow downs along a ± 1 framed unknot in $\operatorname{int} M$;*
4. *(Move $\mathcal{O}_3$): adding or removing two link components, one with zero framing, and one with arbitrary framing, inside of a solid torus, as shown in Figure 5.4.*

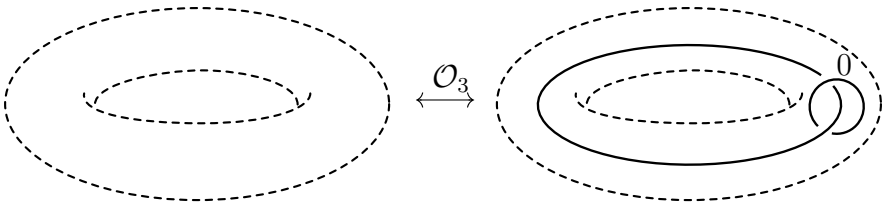

Figure 5.4: **The move $\mathcal{O}_3$, which takes place in a solid torus in M.** The framing on the link component which is homotopically nontrivial in the solid torus can be arbitrary.

We note that in [Rob97], the result is stated without referencing the diffeomorphism f, though for our purposes, it is important to keep track of the diffeomorphism f. We now consider the implications of the previous theorem when applied to link complements.

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multi-based link in a 3–manifold Y and let $\mathbb{U} = (U, \mathbf{w}_0, \mathbf{z}_0)$ denote a multi-based unlink in S^3. Define

$$Y_L = Y \setminus N(L), \qquad \text{and} \qquad S^3_U = S^3 \setminus N(U).$$

Let $\phi_0 : \mathbb{U} \to \mathbb{L}$ be a fixed homeomorphism of based links. We wish to describe diffeomorphisms from S^3_U to Y_L. We add meridianal sutures (closed 1–manifolds) $\gamma_{\mathbf{w}}$, $\gamma_{\mathbf{z}}$, $\gamma_{\mathbf{w}_0}$, and $\gamma_{\mathbf{z}_0}$ in the boundaries of Y_L and S^3_U, corresponding to the basepoints of $\mathbf{w}, \mathbf{z}, \mathbf{w}_0$ and $\mathbf{z}_0$.

A choice of framing λ of L specifies a diffeomorphism

$$\phi_\lambda : (\partial(S^3_U), \gamma_{\mathbf{w}_0}, \gamma_{\mathbf{z}_0}) \to (\partial(Y_L), \gamma_{\mathbf{w}}, \gamma_{\mathbf{z}}).$$

The diffeomorphism ϕ_λ is determined by ϕ_0 and λ up to isotopy preserving the circles $\gamma_{\mathbf{z}}$ and $\gamma_{\mathbf{w}}$. We make the following definition:

Definition 5.4.3. We call a tuple $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ a **parametrized Kirby diagram** for a multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ in Y, if $\phi_0 : \mathbb{U} \to \mathbb{L}$ is a homeomorphism of multi-based 1–manifolds, from an unlink $\mathbb{U} = (U, \mathbf{w}, \mathbf{z})$ in S^3 to $\mathbb{L}$, λ is a framing of L, $\mathbb{S}_1$ is a framed link in S^3_U, and

$$f : (S^3_U)(\mathbb{S}_1) \to Y_L$$

is a diffeomorphism which extends the diffeomorphism ϕ_λ described in the previous paragraph.

We will describe gradings

$$A_{Y,\mathbb{L},S,\mathbb{P}}$$

where S is a choice of J–Seifert surface, $\mathbb{P}$ is a parametrized Kirby decomposition for $(Y,\mathbb{L})$. We will show that the above grading depends only on $(Y,\mathbb{L})$ and S.

We now consider a new move, which we will call $\mathcal{O}_3'$. Given a framed link (L,λ) in Y, consider the effect of performing ± 1 surgery on a knot which encircles a component of L, as in Figure 5.5. Suppose $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a choice of parametrized Kirby diagram for $(Y,\mathbb{L})$, with $f : S^3_U(\mathbb{S}_1) \to Y_L$ a diffeomorphism extending ϕ_λ. We get a diffeomorphism

$$f_K : S^3_U(\mathbb{S}_1 \cup \{K\}) \to Y_L(f(K)).$$

There is a canonical diffeomorphism of $Y_L(f(K))$ with Y_L which is the identity outside of a solid torus K whose boundary intersects ∂Y_L in an annulus. Hence we get a diffeomorphism

$$f_K : S^3_U(\mathbb{S}_1 \cup \{K\}) \to Y_L.$$

Of course on the boundary f_K no longer restricts to ϕ_λ, but instead $\phi_{\lambda'}$, where λ' is a new framing which differs by ± 1 on the component that K encircled.

We call this move $\mathcal{O}_3'$. After performing move $\mathcal{O}_3'$, we get a new parametrized Kirby diagram $\mathbb{P}_K = (\phi_0, \lambda', \mathbb{S}_1 \cup K, f_K)$ for $(Y,\mathbb{L})$.

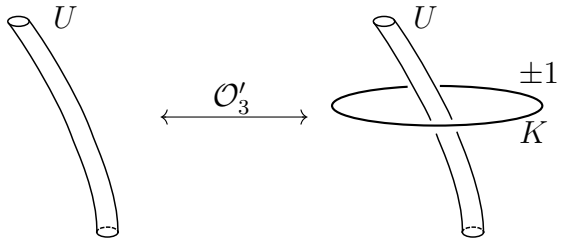

Figure 5.5: **The move $\mathcal{O}_3'$ between two parametrized Kirby diagrams**, and the configuration of the new component K of $\mathbb{S}_1$ with respect to the unlink $U \subseteq S^3$.

We now reformulate Theorem 5.4.2 to describe a sufficient set of moves between any two parametrized Kirby diagrams of a link:

Lemma 5.4.4. *Any two parametrized Kirby diagrams of a 3–manifold with a multi-based link can be connected by a sequence of the following moves:*

1. *(Move $\mathcal{O}_0$): Isotopies of f or $\mathbb{S}_1$ which preserve ∂Y_L, and also preserve $\gamma_{\mathbf{w}}$ and $\gamma_{\mathbf{z}}$;*
2. *(Move $\mathcal{O}_1$): handleslides amongst the components of $\mathbb{S}_1$;*
3. *(Move $\mathcal{O}_2$): blowing up or down along a ± 1 framed unknot which is null-isotopic in $S^3 \setminus (U \cup \mathbb{S}_1)$;*
4. *(Move $\mathcal{O}_3'$): blowing up along a ± 1 framed unknot which encircles a single component of U, exactly once, and no other components of U or $\mathbb{S}_1$;*
5. *(Move $\mathcal{O}_4$): Replacing $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ with $\mathbb{P}' = (\phi_0 \circ \psi, \lambda, \psi(\mathbb{S}_1), f \circ (\psi^{\mathbb{S}_1})^{-1})$, for $\psi : (S^3, \mathbb{U}) \to (S^3, \mathbb{U})$ which is orientation preserving for both S^3 and $\mathbb{U}$.*

Proof. For a fixed homeomorphism ϕ_0 and framing λ, Theorem 5.4.2 implies that Moves $\mathcal{O}_0$, $\mathcal{O}_1$, $\mathcal{O}_2$ and $\mathcal{O}_3$ suffice.

We first claim that for fixed ϕ_0 and λ, we can use only $\mathcal{O}_3$ moves where the solid torus encircles a single component of U, but is unlinked from all other components of $\mathbb{S}_1$ and U. Let us write $\mathcal{O}_3^0$ this move. Note that we need to pay attention to both the framed links, as well as the parameterizing diffeomorphisms resulting from the move. Suppose $\mathbb{S}_1$ is a framed link in S^3_U and $f : S^3_U(\mathbb{S}_1) \to Y_L$ is a diffeomorphism. Let K denote a new link component (with any framing) and let us perform Move $\mathcal{O}_3$ on K, adding K (with some framing) and a zero framed unknot U_K encircling K to $\mathbb{S}_1$. Let $\Phi_{\mathcal{O}_3,K}$ denote the diffeomorphism from $S^3_U(\mathbb{S}_1)$ to $S^3_U(\mathbb{S}_1 \cup \{K \cup U_K\})$ which is the identity outside of a solid torus containing K and U_K. Now let K' differ from K by a crossing change of K with itself, or a crossing change of K with another component of $\mathbb{S}_1$. The link $\mathbb{S}_1 \cup \{K' \cup U_{K'}\}$ can be obtained by a handleslide of a link component across U_K (if the crossing change is of K with itself, then K is handleslid across U_K; if the crossing change is of K with another component of $\mathbb{S}_1$, then the other component is handleslide across U_K) followed by an isotopy of the framed links.

Let $\Phi_H : S^3_U(\mathbb{S}_1 \cup \{K \cup U_K\}) \to S^3_U(\mathbb{S}_1 \cup \{K' \cup U_{K'}\})$ denote the diffeomorphism resulting from the composition of this handleslide and isotopy. We need to show that

$$\Phi_H \circ \Phi_{K,\mathcal{O}_3} \simeq \Phi_{K',\mathcal{O}_3}. \tag{5.3}$$

Suppose that we are handlesliding another component K_0 of $\mathbb{S}_1$ across U_K. Let $N(K_0)$ and $N(K)$ be regular neighborhoods of K_0 and K, inside of S^3_U, such that $N(K)$ contains U_K, as well. Let a be the arc over which we handleslide K_0 across U_K. Noting that $(N(K) \cup N(K_0) \cup N(a))(K_0, K, U_K)$ is a genus two handlebody, and that Φ_H and $\Phi_{K',\mathcal{O}_3} \circ \Phi^{-1}_{K,\mathcal{O}_3}$ both restrict to diffeomorphisms

$$(N(K) \cup N(K_0) \cup N(a))(K_0, K, U_K) \to (N(K) \cup N(K_0) \cup N(a))(K_0, K', U_{K'}),$$

which are the identity on the boundary, we conclude that they are isotopic, as $MCG(H_g, \partial H_g) = \{*\}$, where H_g is a genus g handlebody. Hence Equation (5.3) holds. This is shown in Figure 5.6. An analogous argument holds for changing a crossing of K with itself. In that case, we let a be an arc from U_K to K, and consider a neighborhood N of $K \cup a$, which contains U_K. Now N is a genus two handlebody, and again, Φ_H and $\Phi_{K',\mathcal{O}_3} \circ \Phi^{-1}_{K,\mathcal{O}_3}$, differ only inside of N, and hence must be isotopic.

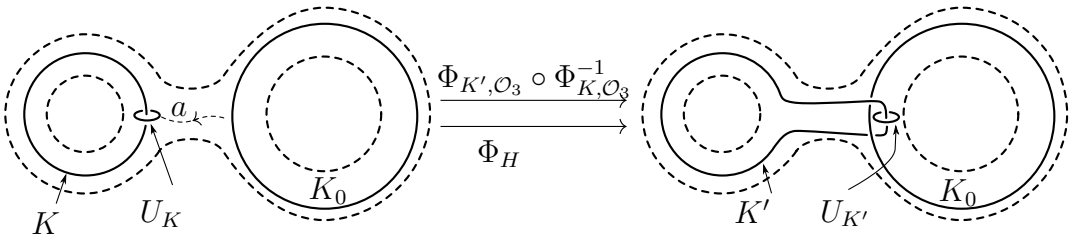

Figure 5.6: **The diffeomorphisms Φ_H and $\Phi_{K',\mathcal{O}_3} \circ \Phi^{-1}_{K,\mathcal{O}_3}$ are equal outside of a genus two handlebody**, and hence they are isotopic.

Hence in particular, instead of performing Move $\mathcal{O}_3$ on K, we can perform Move $\mathcal{O}_3$ on K' (the knot obtained by changing a crossing of K with either itself or with another component of $\mathbb{S}_1$), and then performing a handleslide, and the resulting framed link $\mathbb{S}_1$ and parameterizing diffeomorphism f will be the same.

Next, $K \subseteq S^3_U$ is a knot, and K' is the result of changing a crossing of K with U, then we need to show that the Move $\mathcal{O}_3$ performed on K can be written as a composition of Move $\mathcal{O}_3$, performed on K', as well as Moves $\mathcal{O}_0, \mathcal{O}_1$, and $\mathcal{O}_2$ and $\mathcal{O}^0_3$. The procedure for doing this is shown in Figure 5.7. We perform Move $\mathcal{O}^0_3$ around U, then perform a sequence of handleslides, and then perform the inverse of Move $\mathcal{O}^0_3$. The resulting framed links in S^3_U that we perform surgery on are clearly equal. As before, the parameterizing diffeomorphism resulting from these two moves can be seen to be equal, since they agree outside of a genus two handlebody.

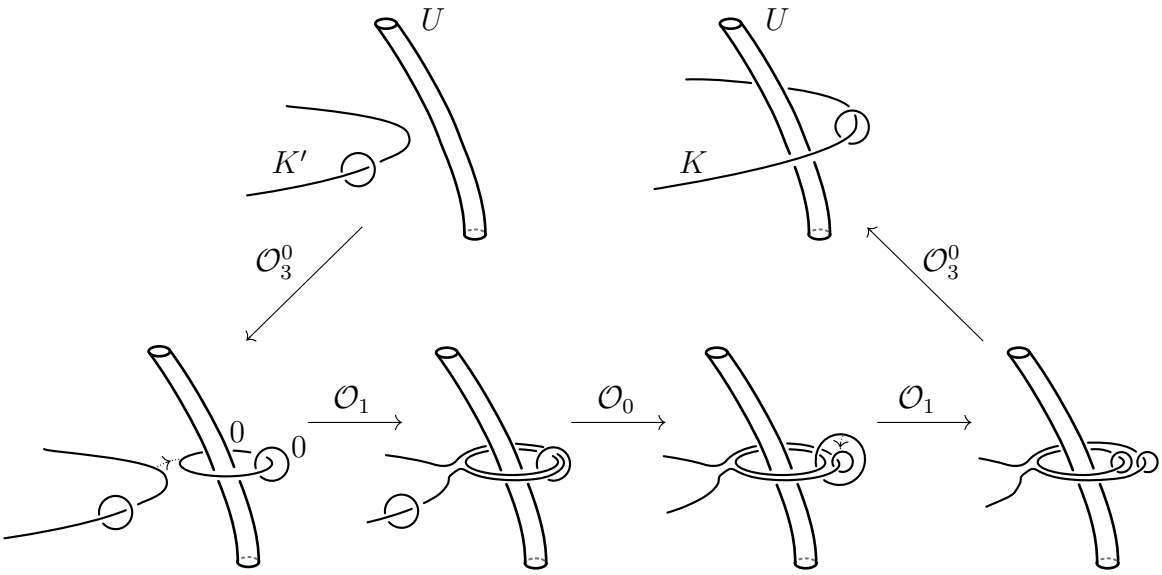

Figure 5.7: **Move $\mathcal{O}_3$ performed on K, is equal to Move $\mathcal{O}_3$ performed on K', composed with a sequence of Moves $\mathcal{O}_0, \mathcal{O}_1, \mathcal{O}_2$ and $\mathcal{O}^0_3$.**

In such manner, by reducing K to an unknot encircling just one (or no) components of U, as in Move $\mathcal{O}^0_3$, we can write an arbitrary $\mathcal{O}_3$ move on a knot K, we can write an arbitrary $\mathcal{O}_3$ move as a composition of the moves $\mathcal{O}_0, \mathcal{O}_1, \mathcal{O}_2$ and $\mathcal{O}^0_3$.

Finally we note that Move $\mathcal{O}^0_3$ can be written as a composition of two $\mathcal{O}'_3$ moves (and possibly some of the other moves, depending on the framings).

Hence for a fixed ϕ_0, Moves $\mathcal{O}_0$, $\mathcal{O}_1, \mathcal{O}_2$ and $\mathcal{O}'_3$ suffice. We finally note that any two ϕ_0 differ by pre-composition with an orientation preserving diffeomorphism from $\mathbb{U}$ to itself,

and any such diffeomorphism extends to an orientation preserving diffeomorphism of $(S^3, \mathbb{U})$ with itself. Hence Move $\mathcal{O}_4$ is sufficient to move between any two ϕ_0 maps. □

5.5 Relative gradings

In this section, we describe the three relative gradings, $A, \mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$, on $\mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s})$, for a diagram $\mathcal{H}$ of $(Y, \mathbb{L})$.

5.5.1 The relative Maslov gradings

We now describe the two Maslov gradings, $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. The grading $\mathrm{gr}_{\mathbf{w}}$ is defined on generators by

$$\mathrm{gr}_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) = \mu(\phi) - 2\sum_{w \in \mathbf{w}} n_w(\phi), \tag{5.4}$$

for a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. The Chern class formula from [OS04d] shows that this is independent of the choice of disk ϕ, as long as $c_1(\mathfrak{s}_{\mathbf{w}}(\mathbf{x})) = c_1(\mathfrak{s})$ is torsion. We extend $\mathrm{gr}_{\mathbf{w}}$ to the entire complex, be declaring all $U_{\mathbf{w}}$ variables to have grading change -2, and all $V_{\mathbf{z}}$ variables to have grading change 0.

Analogously, we can define a relative grading $\mathrm{gr}_{\mathbf{z}}$. On generators it is defined as

$$\mathrm{gr}_{\mathbf{z}}(\mathbf{x}, \mathbf{y}) = \mu(\phi) - 2\sum_{z \in \mathbf{z}} n_z(\phi), \tag{5.5}$$

for a disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. We extend $\mathrm{gr}_{\mathbf{z}}$ to the entire complex, be declaring all $V_{\mathbf{z}}$ variables to have grading change -2, and all $U_{\mathbf{w}}$ variables to have grading change 0. By the Chern class formula of Ozsváth and Szabó, this is independent of the disk in the case that $c_1(\mathfrak{s}_{\mathbf{z}}(\mathbf{x})) = c_1(\mathfrak{s} - PD[L])$ is torsion.

5.5.2 The relative Alexander multi-grading

Suppose that $\pi : L \to J$ is a grading assignment. We now define a relative $\mathbb{Z}^J$ grading A on $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$, in the case that L is J–null-homologous. Given a disk ϕ, and an element

$j \in J$, define

$$n_{\mathbf{z}}(\phi)_j = \sum_{\substack{z \in \mathbf{z} \\ \pi(z)=j}} n_z(\phi), \qquad \text{and} \qquad n_{\mathbf{w}}(\phi)_j = \sum_{\substack{w \in \mathbf{w} \\ \pi(w)=j}} n_w(\phi).$$

As usual, the grading is defined on generators as

$$A(\mathbf{x}, \mathbf{y})_j = (n_{\mathbf{z}} - n_{\mathbf{w}})(\phi)_j, \tag{5.6}$$

for a homology disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ and an index $j \in J$. This is extended to the entire complex $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ by declaring V_z to be $+1$ graded in index $\pi(z)$, and declaring U_w to be -1 graded in index $\pi(w)$.

Of course the above formula depends *a-priori* on the choice of homology disk $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$. We have the following (cf. eg. [OS04e, Lemma 2.18]):

Lemma 5.5.1. *If L is a link with a grading assignment $\pi : C(L) \to J$, then Equation* (5.6) *is independent of the disk ϕ iff $\mathbb{L}$ is (rationally) J–null-homologous.*

Proof. Two disks in $\pi_2(\mathbf{x}, \mathbf{y})$ differ by splicing in elements $\pi_2(\mathbf{x}, \mathbf{x})$, so we just need to show that the expression $(n_{\mathbf{z}} - n_{\mathbf{w}})(\phi)_j$ vanishes for any element $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$.

There is a surjective map

$$H : \pi_2(\mathbf{x}, \mathbf{x}) \to H_2(Y; \mathbb{Z})$$

given by sending the domain of a disk, a 2–chain on Σ with boundary equal to a linear combination of the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves, to the 2–cycle in Y formed by taking $\mathcal{D}(\phi)$ and adding multiples of the compression disks in the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$– handlebodies, along the boundary curves of $\mathcal{D}(\phi)$. Given a class $\phi \in \pi_2(\mathbf{x}, \mathbf{x})$, the integer $(n_{\mathbf{z}} - n_{\mathbf{w}})(\phi)_j$ satisfies

$$(n_{\mathbf{z}} - n_{\mathbf{w}})(\phi)_j = \#(H(\phi) \cap \pi^{-1}(j)), \tag{5.7}$$

where $\#(H(\phi) \cap \pi^{-1}(j))$ is the oriented intersection number. The claim now follows. □

Using the relative Alexander grading on the uncolored complexes, we can define an Alexander grading on the colored complexes, for J–graded colorings.

Lemma 5.5.2. *Suppose that $\mathbb{L}$ is a multi-based link in Y and $(\sigma, \mathfrak{P}, \widehat{\pi}, \pi, J)$ is a J–graded coloring of $\mathbb{L}$. If $\mathbb{L}$ is J–null-homologous, then there is a well defined relative Alexander multi-grading on $\mathcal{C}FL^{\circ}(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$, taking values in $\mathbb{Z}^J$. Recalling that $\mathfrak{P} = \mathfrak{P}_{\mathbf{w}} \sqcup \mathfrak{P}_{\mathbf{z}}$ by definition of a J–graded coloring, if $t \in \mathfrak{P}_{\mathbf{z}}$ then V_t is $+1$ graded if $t \in \mathfrak{P}_{\mathbf{z}}$ then U_t is -1 graded. Similarly there are well defined relative Maslov gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ on the colored complexes.*

Proof. By definition of a J–graded coloring, we have $\mathfrak{P} = \mathfrak{P}_{\mathbf{w}} \sqcup \mathfrak{P}_{\mathbf{z}}$ and $\sigma(\mathbf{w}) \subseteq \mathfrak{P}_{\mathbf{w}}$ and $\sigma(\mathbf{z}) \subseteq \mathfrak{P}_{\mathbf{z}}$. If $\mathbf{x}$ and $\mathbf{y}$ are two intersection points, the quantity $A(\mathbf{x}, \mathbf{y})_j$ is equal to the same formula as before. Since the colors in $\mathfrak{P}$ are partitioned into $\mathfrak{P}_{\mathbf{w}}$ and $\mathfrak{P}_{\mathbf{z}}$, and since the grading assignment $\pi : C(L) \to J$ satisfies

$$\pi \circ \iota = \widehat{\pi} \circ \sigma$$

for a map $\widehat{\pi} : \mathfrak{P} \to J$ (where $\iota : \mathbf{w} \cup \mathbf{z} \to C(L)$ is the natural inclusion), the relative grading has a well defined extension to monomials with nonzero powers of the variables. Similar considerations apply for the Maslov gradings. □

Lemma 5.5.3. *Given a J–null-homologous link, the subset of $\mathcal{C}FL^{\circ}(Y, \mathbb{L}, \mathfrak{s})$ generated by monomials of a particular relative Alexander grading are subcomplexes.*

Proof. This is an easy computation. □

Finally, we note that the relative Alexander grading is preserved by the change of diagrams maps:

Lemma 5.5.4. *Suppose $\mathbb{L}$ is a link in Y, with a grading assignment $\pi : C(L) \to J$, and that $\mathbb{L}$ is J–null-homologous. If $(\mathcal{H}, J_s)$ and $(\mathcal{H}', J_s')$ are two choices of diagrams and almost complex structures for $(Y, \mathbb{L})$, and $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ is a torsion Spin^c structure, then the change of diagrams map*

$$\Phi_{(\mathcal{H}, J_s) \to (\mathcal{H}', J_s')} : \mathcal{C}FL^{\circ}_{J_s}(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^{\circ}_{J_s'}(\mathcal{H}', \mathfrak{s})$$

preserves the relative Alexander multi-grading over $\mathbb{Z}^J$, and the two relative Maslov gradings.

Proof. To prove this, we check the statement for the maps associated to isotopies and handleslides of the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–curves, $(1,2)$–stabilizations and destabilizations, diffeomorphisms of the 3–manifold, and changes of the almost complex structure. Let us consider a handleslide or isotopy of the $\boldsymbol{\alpha}$–curves. We will leave the other moves to the reader. Suppose that $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a triple for an $\boldsymbol{\alpha}$–equivalence, and suppose that $\mathbf{x}$ and $\mathbf{x}'$ are two intersection points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Suppose that $\psi \in \pi_2(\theta, \mathbf{x}, \mathbf{y})$ and $\pi_2(\theta', \mathbf{x}', \mathbf{y}')$ are two triangles with Maslov index zero, with θ and θ' in the same Alexander and Maslov gradings as the top generator of $\mathcal{H}FL^\circ(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ (for $\circ \in \{\wedge, -\}$). Now since ψ and ψ' represent the same Spin^c structure, there are disks

$$\phi_{\alpha'\alpha} \in \pi_2(\theta', \theta), \qquad \phi_{\alpha\beta} \in \pi_2(\mathbf{x}', \mathbf{x}), \qquad \text{and} \qquad \phi_{\alpha'\beta} \in \pi_2(\mathbf{y}, \mathbf{y}')$$

such that

$$\psi' = \psi + \phi_{\alpha'\alpha} + \phi_{\alpha\beta} + \phi_{\alpha'\beta}.$$

Hence

$$\begin{aligned} A(\mathbf{y} \cdot U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\psi)}, \mathbf{y}' \cdot U_{\mathbf{w}}^{n_{\mathbf{w}}(\psi')} V_{\mathbf{z}}^{n_{\mathbf{z}}(\psi')})_j &= (n_{\mathbf{z}} - n_{\mathbf{w}})(\phi_{\alpha'\beta} + \psi - \psi')_j \\ &= (n_{\mathbf{z}} - n_{\mathbf{w}})(-\phi_{\alpha\beta})_j - (n_{\mathbf{z}} - n_{\mathbf{w}})(\phi_{\alpha'\alpha})_j \\ &= A(\mathbf{x}, \mathbf{x}')_j, \end{aligned}$$

since $(n_{\mathbf{z}} - n_{\mathbf{w}})(\phi_{\alpha'\alpha})_j = 0$ as θ and θ' are in the same Alexander grading. Hence $\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'}$ preserves the relative Alexander multi-grading. Easy adaptations of the above argument show that $\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'}$ preserves the other relative gradings, and also that the maps associated to $\boldsymbol{\beta}$–equivalences, changes of almost complex structure, $(1,2)$-stabilization and diffeomorphisms of Y all preserve the relative gradings. □

5.5.3 Two simple examples

We briefly consider a few examples, focusing on cases where one grading may be defined but not another.

Example 5.5.5. Consider $Y = S^1 \times S^2$, with $K = S^1 \times \{pt\}$. A diagram is shown in Figure 5.8. For $\mathfrak{s}$ the torsion Spin^c structure, we see that

$$\mathcal{C}FL^-(Y, \mathbb{L}, \mathfrak{s}) \cong \mathbb{Z}_2[U, V] \xrightarrow{1+V} \mathbb{Z}_2[U, V].$$

The homology is $\mathbb{Z}[U, V]/(V - 1)$. The grading $\mathrm{gr}_{\mathbf{w}}$ is defined, but $\mathrm{gr}_{\mathbf{z}}$ cannot be defined (for example V acts by the identity on homology). Note that $\mathfrak{s} - PD[K]$ is not torsion, so we wouldn't expect $\mathrm{gr}_{\mathbf{z}}$ to be well defined.

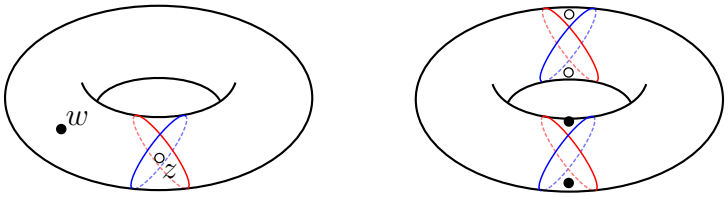

Figure 5.8: **Two diagrams for $S^1 \times S^2$, with a knot or link.** The left is for the knot $K = S^1 \times \{p\}$ and the right is for $K = S^1 \times \{p_1, p_2\}$, with the components going in opposite directions. All intersection points are mapped to the torsion Spin^c structure by $\mathfrak{s}_{\mathbf{w}}$.

Example 5.5.6. Another instructive example is $S^1 \times S^2$ with $L = S^1 \times \{p_1, p_2\}$, giving each component two basepoints, and the two components opposite orientations. A diagram is shown on the right side of Figure 5.8. The intersection points represent the torsion Spin^c structure with respect to $\mathfrak{s}_{\mathbf{w}}$. It's easy to see that the collapsed Alexander grading can be defined, but a two component Alexander grading cannot be defined. One can get a diagram for the link where both components are going in the same direction by switching the roles of the two $\mathbf{w}$ and $\mathbf{z}$ basepoints on a component. The intersection points then represent a Spin^c structure $\mathfrak{s}_1$ where $\langle c_1(\mathfrak{s}_1), \{pt\} \times S^2 \rangle = \pm 2$.

5.5.4 Absolute gradings on $\mathcal{C}FL^\circ((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$ for unlinks $\mathbb{U}$

As a first step to describing the absolute gradings in general, we first declare the absolute gradings for unlinks in $(S^1 \times S^2)^{\#k}$, with various configurations of basepoints. It is sufficient to define the gradings on the generators of $\widehat{\mathcal{C}FL}$, since we can just extend over all elements

of the full chain complex by declaring the grading of the variables. For simplicity assume that Σ is connected. For disconnected manifolds we define the grading of a generator as the sum of the gradings of the constituent generators on each component (note our convention is that $HF^-(S^3)$ has top degree generator in degree zero).

Lemma 5.5.7. *The $\mathbb{Z}_2$–module $\widehat{HFL}((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$ has rank $2^{|\mathbf{w}|+k-1}$. Furthermore, the group $\widehat{HFL}((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$ has a top degree generator with respect to each of the gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$, for which we write $\Theta^{\mathbf{w}}$ or $\Theta^{\mathbf{z}}$.*

Proof. By Lemma 5.5.4, the relatively graded isomorphism type of the group $\widehat{\mathcal{HFL}}((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$ is an invariant, so we need only check the claim for a particularly simple diagram. One can prove this by induction. A new unknot component with exactly two basepoints can be added by adding it in a sphere disjoint from $(S^1 \times S^2)^{\#k}$, and then attaching it with a 1-handle. It's easy to verify that the new complex has two top generators $\Theta^{\mathbf{w}}$ and $\Theta^{\mathbf{z}}$, since the old complex did. Using, for example, the holomorphic disk computation from Proposition 3.5.3 of a quasi-stabilized differential, one can see that the same is true when we add two basepoints to a link component. □

For an arbitrary diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ of $((S^1 \times S^2)^{\#k}, \mathbb{U})$, we can thus define absolute Maslov and Alexander gradings by declaring the values of the gradings on $\Theta^{\mathbf{w}}$. Note that it is sufficient to define the Alexander $\mathbb{Z}^J$ grading when $J = C(\mathbb{U})$ and the grading assignment $\pi : C(\mathbb{U}) \to J$ is the identity map. For any other choice of π and J, we define the absolute grading over J to be obtained by collapsing and extending gradings, i.e. if $\pi : C(\mathbb{U}) \to J$, then we define

$$A_{(S^1\times S^2)^{\#k},\mathbb{U},\pi}(\mathbf{x})_j = \sum_{K\in\pi^{-1}(j)} A_{(S^1\times S^2)^{\#k},\mathbb{U}}(\mathbf{x})_K. \tag{5.8}$$

We declare

$$\mathrm{gr}_{\mathbf{w}}(\Theta^{\mathbf{w}}) = \mathrm{gr}_{\mathbf{z}}(\Theta^{\mathbf{z}}) = \tfrac{1}{2}(k + |\mathbf{w}| - 1) = \tfrac{1}{2}(k + |\mathbf{z}| - 1), \tag{5.9}$$

where $|\mathbf{w}| = |\mathbf{z}|$ is the total number of $\mathbf{w}$–basepoints. Similarly we declare that

$$A_{(S^1\times S^2)^{\#k},\mathbb{U},S}(\Theta^{\mathbf{w}})_K = \tfrac{1}{2}(n_K - 1), \tag{5.10}$$

where $K \in J = C(\mathbb{U})$ is a component of $\mathbb{U}$ and n_K is half the total number of basepoints on K (i.e. the number of $\mathbf{w}$–basepoints on K). From now on, we will write only A_K for $(A_{(S^1\times S^2)^{\#k},\mathbb{U},S})_K$. Note that these declarations determine $\mathrm{gr}_{\mathbf{w}}, \mathrm{gr}_{\mathbf{z}}$ and A_K on all of $\mathcal{C}FL^\circ((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$.

If $\mathbb{U}$ has exactly two basepoints per component, then it's easy to see that $\Theta^{\mathbf{w}} = \Theta^{\mathbf{z}}$, and hence we have $\mathrm{gr}_{\mathbf{w}} = \mathrm{gr}_{\mathbf{z}}$ and $A_K = 0$ on $\widehat{HFL}((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$.

Computing for diagrams which are obtained by performing a sequence of quasi-stabilizations, we can compute the following equations for a general configuration of basepoints on $\mathbb{U} \subseteq (S^1 \times S^2)^{\#k}$:

$$\begin{aligned}
\mathrm{gr}_{\mathbf{w}}(\Theta^{\mathbf{z}}) &= \mathrm{gr}_{\mathbf{w}}(\Theta^{\mathbf{w}}) - \sum_{K\in C(\mathbb{U})} (n_K - 1) \\
\mathrm{gr}_{\mathbf{z}}(\Theta^{\mathbf{w}}) &= \mathrm{gr}_{\mathbf{z}}(\Theta^{\mathbf{z}}) - \sum_{K\in C(\mathbb{U})} (n_K - 1) \\
A(\Theta^{\mathbf{w}})_K &= \tfrac{1}{2}(n_K - 1) \\
A(\Theta^{\mathbf{z}})_K &= -\tfrac{1}{2}(n_K - 1).
\end{aligned}$$

Finally we note that

$$A = \tfrac{1}{2}(\mathrm{gr}_{\mathbf{w}} - \mathrm{gr}_{\mathbf{z}}), \tag{5.11}$$

where A denotes the collapsed Alexander grading, since the equation is obviously true for the relative gradings, and is true for the absolute gradings evaluated at $\Theta^{\mathbf{w}}$.

5.6 Coherent gradings on coherent chain homotopy types

The objects $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ are not actual chain complexes. Instead they are "coherent chain homotopy types", indexed by the set of admissible diagrams for $(Y, \mathbb{L})$. Naturally we would like to make precise what it means to specify a grading on $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$, since the proof of invariance of the gradings does not factor through a single diagram. Additionally, as

the change of diagrams are not isomorphisms, instead just chain homotopy equivalences on curved chain complexes, we need to be somewhat careful about pushing forward a grading along the change of diagrams maps.

For an admissible diagram $\mathcal{H}$ of $(Y, \mathbb{L})$, with grading assignment $\pi : C(L) \to J$, we define

$$\mathbb{A}(\mathcal{H}, \pi, \mathfrak{s}),$$

to be the set of absolute $\mathbb{Q}$–lifts of the relative Alexander grading defined in Section 5.5. Analogously, we define

$$\mathbb{G}_{\mathbf{w}}(\mathcal{H}, \mathfrak{s}), \qquad \text{and} \qquad \mathbb{G}_{\mathbf{z}}(\mathcal{H}, \mathfrak{s})$$

to be the set of absolute $\mathbb{Q}$–lifts of the relative Maslov gradings define in Section 5.5. In this section, we define a notion of a **coherent grading**, which is essentially one which respects the change of diagrams maps. For the set of coherent gradings for a grading assignment π, we will later introduce the notation $\overrightarrow{\mathbb{A}}(Y, \mathbb{L}, \pi, \mathfrak{s})$. We will define it as the direct limit over the sets $\mathbb{A}(\mathcal{H}, \pi, \mathfrak{s})$, using the change of diagrams maps which we define below. This is somewhat pedantic, but results in cleaner arguments later on.

5.6.1 Change of diagrams maps for Alexander gradings

Suppose $\mathcal{H}$ and $\mathcal{H}'$ are two admissible diagrams for $(Y, \mathbb{L})$. Analogously to the change of diagrams maps for the chain complexes, we now define maps

$$\Phi_{\mathcal{H} \to \mathcal{H}'} : \mathbb{A}(\mathcal{H}, \pi, \mathfrak{s}) \to \mathbb{A}(\mathcal{H}', \pi, \mathfrak{s})$$

on gradings. These will be defined as a composition of maps associated to individual moves. If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram, and $\boldsymbol{\alpha}'$ is another choice of $\boldsymbol{\alpha}$ curves such that $\boldsymbol{\alpha}'$ is the result of a sequence of handleslides or isotopies of the $\boldsymbol{\alpha}$ curves and $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ satisfies the appropriate admissibility requirements, then we can define a map

$$\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} : \mathbb{A}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \pi, \mathfrak{s}) \to \mathbb{A}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}, \pi, \mathfrak{s}),$$

as follows. We note that $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ is a diagram is for an unlink in $(S^1 \times S^2)^{\#k}$, with two basepoints on each component. If $A \in \mathbb{A}(\mathcal{H}, \pi, \mathfrak{s})$, we define $\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'}(A)$ by the formula

$$\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'}(A)(\mathbf{y})_j = A(\theta)_j + A(\mathbf{x})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi)_j,$$

for a homology triangle $\psi \in \pi_2(\theta, \mathbf{x}, \mathbf{y})$ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}|_{X_{\alpha'\alpha\beta}}$ under the identification of $X_{\alpha'\alpha\beta}$ as a subset of $Y \times [0,1]$. The grading $\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'}(A)$ is independent of the choice of the intersection points $\mathbf{x}$ and θ, as well as the homology triangle ψ, since any other homology triangle representing $\mathfrak{s}$ with the same endpoints can be obtained by splicing in homology disks on $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$, $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, and $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta})$ which doesn't affect the formula.

Analogously, if $\boldsymbol{\beta}'$ differs from $\boldsymbol{\beta}$ by a sequence of handleslides or isotopies, and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ achieves the appropriate admissibility hypothesis, we can define a map $\Phi_{\boldsymbol{\beta} \to \boldsymbol{\beta}'}^{\boldsymbol{\alpha}}$.

We have the following:

Lemma 5.6.1. *Suppose that $(\Sigma, \boldsymbol{\alpha}'', \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ is an admissible quadruple and $\boldsymbol{\alpha}, \boldsymbol{\alpha}'$ and $\boldsymbol{\alpha}''$ are all related to each other by a sequence of handleslides and isotopies. If $A \in \mathbb{A}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \pi, \mathfrak{s})$, then*

$$\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}''}(A) = (\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha}' \to \boldsymbol{\alpha}''} \circ \Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'})(A).$$

Proof. Suppose that $\mathbf{x} \in \mathbb{T}_{\alpha''} \cap \mathbb{T}_{\beta}$ and pick intersection points $\theta_{\alpha''\alpha'} \in \mathbb{T}_{\alpha''} \cap \mathbb{T}_{\alpha'}$, $\theta_{\alpha'\alpha} \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_{\alpha}$, $\theta_{\alpha''\alpha} \in \mathbb{T}_{\alpha''} \cap \mathbb{T}_{\alpha}$, $\mathbf{x}_{\alpha\beta} \in \mathbb{T}_{\alpha} \cap \mathbb{T}_{\beta}$ and $\mathbf{x}_{\alpha'\beta} \in \mathbb{T}_{\alpha'} \cap \mathbb{T}_{\beta}$, such that $\theta_{\alpha''\alpha}, \theta_{\alpha'\alpha}$ and $\theta_{\alpha''\alpha'}$ represent the torsion Spin^c structure. Pick triangles $\psi_{\alpha''\alpha'\alpha}, \psi_{\alpha''\alpha\beta}, \psi_{\alpha'\alpha\beta}$, and $\psi_{\alpha''\alpha'\beta}$ with the previous intersection points as vertices, such that

$$\psi_{\alpha''\alpha\beta} + \psi_{\alpha''\alpha'\alpha} = \psi_{\alpha'\alpha\beta} + \psi_{\alpha''\alpha'\beta}. \tag{5.12}$$

By definition, we have

$$\Phi_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}''}^{\boldsymbol{\beta}}(A)(\mathbf{x})_j = A(\mathbf{x}_{\alpha\beta})_j + A(\theta_{\alpha''\alpha})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha''\alpha\beta})_j$$

and

$$(\Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha}' \to \boldsymbol{\alpha}''} \circ \Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'})(A)(\mathbf{x})_j = A(\mathbf{x}_{\alpha\beta})_j + A(\theta_{\alpha'\alpha})_j + A(\theta_{\alpha''\alpha'})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha''\alpha'\beta} + \psi_{\alpha'\alpha\beta})_j.$$

The second minus the first is

$$-A(\theta_{\alpha''\alpha})_j + A(\theta_{\alpha'\alpha})_j + A(\theta_{\alpha''\alpha'})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha''\alpha'\alpha})_j,$$

which we need to show is zero. We note that the previous value is independent of the endpoints $\theta_{\alpha''\alpha}, \theta_{\alpha'\alpha}$ or $\theta_{\alpha''\alpha'}$, as well as the homology triangle $\psi_{\alpha''\alpha'\alpha}$, as any other homology triangle representing the same Spin^c structure on $X_{\alpha''\alpha'\alpha}$ can be obtained from $\psi_{\alpha''\alpha'\alpha}$ by splicing in homology disks on the three diagrams representing the ends of $X_{\alpha''\alpha'\alpha}$, and the above quantity is clearly independent of splicing in such disks.

To see that the above quantity is zero, we simply use Lemma 5.5.4, which shows that the relative Alexander grading is preserved by the change of diagrams maps. As such the top degree generator $\Theta^{\mathbf{w}} = \Theta^{\mathbf{z}}$ is preserved by the change of diagrams map (note that $\Theta^{\mathbf{w}} = \Theta^{\mathbf{z}}$ since there are exactly two basepoints per component). Since we defined the absolute Alexander grading of $\Theta^{\mathbf{w}}$ to be equal in two diagrams $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ and $(\Sigma, \boldsymbol{\alpha}'', \boldsymbol{\alpha})$, it follows that the above quantity is zero. The lemma statement now follows. □

The analog of the previous lemma for two moves of the $\boldsymbol{\beta}$–curves follows from the same argument. Similarly we can commute the maps induced by $\boldsymbol{\alpha}$–moves with the maps induced by $\boldsymbol{\beta}$–moves.

Lemma 5.6.2. *Suppose that $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a Heegaard quadruple achieving weak admissibility, such that $\boldsymbol{\alpha}'$ results from $\boldsymbol{\alpha}$ by a sequence of handleslides and isotopies, and $\boldsymbol{\beta}'$ is obtained from $\boldsymbol{\beta}$ by a sequence of handleslides or isotopies. If $A \in \mathbb{A}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, then*

$$(\Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}'} \circ \Phi^{\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'})(A) = (\Phi^{\boldsymbol{\alpha}'}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'} \circ \Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}})(A).$$

Proof. This follows from the obvious adaptation of the proof of Lemma 5.6.1, using associativity on the level of homology classes. □

Using the previous two lemmas, we can adapt the strategy of [JT12, Section 9] to show that if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}')$ are each admissible Heegaard diagrams for $(Y, \mathbb{L})$ with the same underlying Heegaard surface, such that $\boldsymbol{\alpha}'$ is the result of a sequence of handleslides and

isotopies amongst the $\boldsymbol{\alpha}$ curves, and $\boldsymbol{\beta}'$ is the result of handleslides and isotopies amongst the $\boldsymbol{\beta}$ curves, then there are well defined maps

$$\Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}\to\boldsymbol{\beta}'} : \mathbb{A}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}) \to \mathbb{A}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}'),$$

which are functorial under composition, and also satisfy $\Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}}_{\boldsymbol{\beta}\to\boldsymbol{\beta}} = \mathrm{id}$.

If the diagram $\mathcal{H}_\sigma$ is the result of a $(1,2)$–stabilization of the diagram $\mathcal{H}$, and $A \in \mathbb{A}(\mathcal{H}, \pi, \mathfrak{s})$, then we can define a grading $\Phi_\sigma(A) \in \mathbb{A}(\mathcal{H}_\sigma, \pi, \mathfrak{s})$ by the formula

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

where c is the new intersection point on $\mathcal{H}_\sigma$.

Lemma 5.6.3. *If $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ is an admissible triple associated to a sequence of $\boldsymbol{\alpha}$ handleslides or isotopies, and $(\Sigma, \overline{\boldsymbol{\alpha}}', \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}})$ is the triple resulting from performing $(1,2)$–stabilization to $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$, then we have*

$$(\Phi_\sigma \circ \Phi^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}_{\boldsymbol{\beta}})(A) = (\Phi^{\overline{\boldsymbol{\alpha}}\to\overline{\boldsymbol{\alpha}}'}_{\overline{\boldsymbol{\beta}}} \circ \Phi_\sigma)(A).$$

Proof. To a homology triangle on $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$, we can pick a stabilized triangle on $(\Sigma, \overline{\boldsymbol{\alpha}}', \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}})$, which has vertices at the new, higher degree intersection point of $\overline{\boldsymbol{\alpha}}' \cap \overline{\boldsymbol{\alpha}}$, as well as the two new intersection points of $\overline{\boldsymbol{\alpha}}' \cap \overline{\boldsymbol{\beta}}$ and $\overline{\boldsymbol{\alpha}} \cap \overline{\boldsymbol{\beta}}$. An easy computation using this homology class shows that the induced gradings satisfy the stated relation. □

A similar lemma shows that the maps Φ_σ commute in the appropriate sense with handleslides and isotopies of the $\boldsymbol{\beta}$ curves.

Similarly we have the following:

Lemma 5.6.4. *If σ and σ' are two $(1,2)$–stabilizations, then*

$$(\Phi_\sigma \circ \Phi_{\sigma'})(A) = (\Phi_{\sigma'} \circ \Phi_\sigma)(A).$$

Proof. This is obvious from the formulas. □

Additionally, to a diffeomorphism $\phi : (Y, \mathbb{L}) \to (Y, \mathbb{L})$ which maps L to L and is the identity on $\mathbf{w} \cup \mathbf{z}$, there is a tautological map

$$\phi_* : \mathbb{A}(\mathcal{H}, \pi, \mathfrak{s}) \to \mathbb{A}(\phi_*\mathcal{H}, \phi_*\pi, \phi_*\mathfrak{s}).$$

We have the following:

Lemma 5.6.5. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$ and $\phi : (\Sigma, \mathbf{w} \cup \mathbf{z}) \to (\Sigma, \mathbf{w} \cup \mathbf{z})$ is a diffeomorphism which is isotopic to $\mathrm{id}|_\Sigma$, relative $\mathbf{w} \cup \mathbf{z}$, then*

$$\Phi^{\boldsymbol{\beta} \to \phi_*\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \phi_*\boldsymbol{\alpha}} = \phi_*.$$

Proof. This follows from an adaptation of the proof of [JT12, Proposition 9.24] (except that we don't need to pay attention to whether triangles have holomorphic representatives). One decomposes ϕ into a composition of diffeomorphisms, such that for each, the diffeomorphism map can be defined as a composition of two small triangle maps (one for the $\boldsymbol{\alpha}$–equivalence, and one for a $\boldsymbol{\beta}$–equivalence) (see the proof in [JT12] for more details). The small triangle map induces a map on gradings, which is equal by definition to the maps induced by $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ equivalences. □

From [JT12], there is an important class of loops in the space of Heegaard diagrams, called simple handleswaps, and we need to check that there is no monodromy around such loops. The simple handleswap loop can be thought of as generating π_1 of the space of Heegaard diagrams. See [JT12, Definition 2.32] for a precise description. The move consists of one $\boldsymbol{\alpha}$–equivalence, one $\boldsymbol{\beta}$–equivalence, and one diffeomorphism, all supported in a ball. We have the following, easier, version of [JT12, Proposition 9.25]:

Lemma 5.6.6. *If $H_1 \xrightarrow{e} H_2 \xrightarrow{f} H_3 \xrightarrow{g} H_1$ is a simple handleswap, then*

$$\Phi_g \circ \Phi_f \circ \Phi_e = \mathrm{id},$$

as a map on $\mathbb{A}(\mathcal{H}, \pi, \mathfrak{s})$.

Proof. The maps Φ_e and Φ_f are maps induced by $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–handleslides. The map Φ_g is induced by a diffeomorphism. Since they are computed (in our context) by picking any homology triangle, we can simply pick the small triangle representatives. For such a choice of triangles, it is easy to verify the above identity. □

Proposition 5.6.7. *For two admissible diagrams $\mathcal{H}$ and $\mathcal{H}'$ for $(Y, \mathbb{L})$, there are maps $\Phi_{\mathcal{H}\to\mathcal{H}'} : \mathbb{A}(\mathcal{H}, \pi, \mathfrak{s}) \to \mathbb{A}(\mathcal{H}', \pi, \mathfrak{s})$, defined as a composition of the maps described above, for a choice of Heegaard moves from $\mathcal{H}$ to $\mathcal{H}'$. The maps $\Phi_{\mathcal{H}\to\mathcal{H}'}$ are independent of the choice of intermediate Heegaard diagrams.*

Proof. We will show that the spaces $\mathbb{A}(\mathcal{H}, \pi, \mathfrak{s})$ define a "strong Heegaard invariant" from the set of diffeomorphism types of sutured manifolds which are link complements ([JT12, Definition 2.33]) to the category of affine sets over $\mathbb{Q}^J$. Axiom (1) is that we can separately define functorial maps associated to the subcategories of sutured manifolds corresponding to $\boldsymbol{\alpha}$–equivalences, $\boldsymbol{\beta}$–equivalences, and diffeomorphisms. Axiom (2) states that the five distinguished rectangles commute. The first four are proven in Lemmas 5.6.1, 5.6.2, 5.6.3, and 5.6.4. The fifth states that the (1,2)–stabilization maps commute with diffeomorphism maps, which is obvious. Similarly, Axiom (3) states that the map induced by an isotopy of Σ is equal to map induced by $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$–equivalence, which is verified in Lemma 5.6.5. Finally Axiom (4), handleswap invariance, is verified in Lemma 5.6.6. The result now follows from [JT12, Theorem 2.39]. □

Lemma 5.6.8. *If $A \in \mathbb{A}(\mathcal{H}, \pi, \mathfrak{s})$, then the change of diagrams maps $\Phi_{\mathcal{H}\to\mathcal{H}'} : \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H}', \mathfrak{s})$ are 0–graded maps with respect to the two gradings A and $\Phi_{\mathcal{H}\to\mathcal{H}'}(A)$.*

Proof. This is essentially a tautology from the definitions. One simply checks this for equivalences of the $\boldsymbol{\alpha}$– and $\boldsymbol{\beta}$– curves, as well as $(1, 2)$–stabilizations, and diffeomorphisms $\phi : (Y, \mathbb{L}) \to (Y, \mathbb{L})$. □

5.6.2 Change of diagrams maps for Maslov gradings

We now sketch how to adapt the ideas of the previous section to define change of diagrams maps on the sets of Maslov gradings of complexes.

From our declaration of the absolute gradings for $\mathcal{C}FL^-(\mathcal{H},\mathfrak{s})$ for a diagram $\mathcal{H}=(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ for $((S^1\times S^2)^{\#k},\mathbb{U})$, when $\mathbb{U}$ is an unlink with two basepoints per component, the absolute Maslov gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ are determined by declaring the top degree $\mathrm{gr}_{\mathbf{o}}$ generator of $\widehat{HFL}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ to be in $\mathrm{gr}_{\mathbf{o}}$ grading $\frac{1}{2}(k+|\mathbf{w}|-1)$. By picking a simple diagram, it is easily computed that $\Theta_{\mathbf{w}}=\Theta_{\mathbf{z}}$ and $\mathrm{gr}_{\mathbf{w}}=\mathrm{gr}_{\mathbf{z}}$, on homology.

Recalling that $\mathbb{G}_{\mathbf{w}}(\mathcal{H},\mathfrak{s})$ and $\mathbb{G}_{\mathbf{z}}(\mathcal{H},\mathfrak{s})$ are defined as the sets of lifts of the relative $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ gradings on $\mathcal{C}FL^\infty(\mathcal{H},\mathfrak{s})$, we note that if $(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\alpha},\boldsymbol{\beta})$ is an admissible triple, we can define a map

$$\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}:\mathbb{G}_{\mathbf{o}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})\to\mathbb{G}_{\mathbf{o}}(\Sigma,\boldsymbol{\alpha}',\boldsymbol{\beta}),$$

for $\mathbf{o}\in\{\mathbf{w},\mathbf{z}\}$. If $G\in\mathbb{G}_{\mathbf{o}}(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta})$ we define $\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}(G)$ by the formula

$$\Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}(G)(\mathbf{y})=G(\mathbf{x})+\mathrm{gr}_{\mathbf{o}}(\boldsymbol{\theta})-\tfrac{1}{2}(k+|\mathbf{w}|-1)-\mu(\psi)+2n_{\mathbf{o}}(\psi),$$

for a homology triangle $\psi\in\pi_2(\theta,\mathbf{x},\mathbf{y})$ with $\mathfrak{s}_{\mathbf{w}}(\psi)|_{Y_{\alpha\beta}}=\mathfrak{s}$ and $c_1(\mathfrak{s}_{\mathbf{w}}|_{Y_{\alpha'\alpha}})$ torsion. Note that the formula is invariant under splicing in disks on any of the ends, and hence is independent of the choice of intersection points $\mathbf{x}$ and θ, as well as the triangle ψ. We can additionally define maps for moves of the $\boldsymbol{\beta}$–curves, as well as diffeomorphisms and (1,2)–(de)stabilizations. The argument in the previous section goes through with only trivial modification to yield well defined, functorial change of diagrams maps $\Phi_{\mathcal{H}\to\mathcal{H}'}:\mathbb{G}_{\mathbf{o}}(\mathcal{H},\mathfrak{s})\to\mathbb{G}_{\mathbf{o}}(\mathcal{H}',\mathfrak{s})$.

Lemma 5.6.9. *If $G\in\mathbb{G}_{\mathbf{o}}(\mathcal{H},\mathfrak{s})$, then the change of diagrams maps $\Phi_{\mathcal{H}\to\mathcal{H}'}:\mathcal{C}FL^\infty(\mathcal{H},\mathfrak{s})\to\mathcal{C}FL^\infty(\mathcal{H}',\mathfrak{s})$ are 0–graded maps with respect to the two gradings G and $\Phi_{\mathcal{H}\to\mathcal{H}'}(G)$.*

Proof. As in Lemma 5.6.8, this is essentially a tautology, and is checked easily for each elementary Heegaard move. □

5.6.3 Coherent gradings

Using the previous two sections, we can define the notion of a **coherent grading**. Of course this notion is already implicit in Heegaard Floer homology, but we state it for convenience. Let $D(Y,\mathbb{L},\mathfrak{s})$ denote the set of strongly $\mathfrak{s}$–admissible diagrams for $(Y,\mathbb{L})$. We let $\overrightarrow{\mathbb{A}}(Y,\mathbb{L},\pi,\mathfrak{s})$ denote the subset

$$\overrightarrow{\mathbb{A}}(Y,\mathbb{L},\pi,\mathfrak{s}) \subseteq \prod_{\mathcal{H}\in D(Y,\mathbb{L},\mathfrak{s})} \mathbb{A}(\mathcal{H},\pi,\mathfrak{s})$$

consisting of tuples $(A_{\mathcal{H}})_{\mathcal{H}\in D(Y,\mathbb{L},\mathfrak{s})}$ such that $\Phi_{\mathcal{H}\to\mathcal{H}'}(A_{\mathcal{H}}) = A_{\mathcal{H}'}$. We call such an element a **coherent Alexander grading**. We define the coherent Maslov gradings $\overrightarrow{\mathbb{G}}_{\mathbf{w}}(Y,\mathbb{L},\mathfrak{s})$ and $\overrightarrow{\mathbb{G}}_{\mathbf{z}}(Y,\mathbb{L},\mathfrak{s})$ analogously.

Note that there are canonical isomorphisms

$$\eta_{\mathcal{H},\mathbb{A}} : \mathbb{A}(\mathcal{H},\pi,\mathfrak{s}) \to \overrightarrow{\mathbb{A}}(Y,\mathbb{L},\pi,\mathfrak{s}),$$

as well as similar isomorphisms $\eta_{\mathcal{H},\mathbb{G}_{\mathbf{w}}}$ and $\eta_{\mathcal{H},\mathbb{G}_{\mathbf{z}}}$ for the Maslov gradings. In particular, a choice of absolute Alexander or Maslov grading on a single diagram determines a coherent grading, under the map $\eta_{\mathcal{H},\mathbb{A}}$.

Lemma 5.6.10. *The absolute Alexander and Maslov gradings we declared for unlinks in $(S^1\times S^2)^{\#k}$ are coherent gradings.*

Proof. Let $\mathcal{H}$ and $\mathcal{H}'$ be two diagrams, and let $G_{\mathcal{H}}$ and $G_{\mathcal{H}'}$ be two gradings we declared in Section 5.5.4, for the diagrams $\mathcal{H}$ and $\mathcal{H}'$ (assume both $G_{\mathcal{H}}$ and $G_{\mathcal{H}'}$ are both lifts of the Alexander grading, or both lifts of $\mathrm{gr}_{\mathbf{w}}$, or both lifts of $\mathrm{gr}_{\mathbf{z}}$). Let $\Theta^{\mathbf{w}}_{\mathcal{H}} \in \widehat{\mathcal{H}FL}(\mathcal{H},\mathfrak{s}_0)$ denote the top $\mathrm{gr}_{\mathbf{w}}$–graded element. It is sufficient to show that

$$\Big(\Phi_{\mathcal{H}\to\mathcal{H}'}(G_{\mathcal{H}})\Big)(\Theta^{\mathbf{w}}_{\mathcal{H}'}) = G_{\mathcal{H}'}(\Theta^{\mathbf{w}}_{\mathcal{H}'}).$$

By Lemma 5.5.4, the change of diagrams maps preserve the relative gradings, and as a consequence, $\Phi_{\mathcal{H}\to\mathcal{H}'}(\Theta^{\mathbf{w}}_{\mathcal{H}}) = \Theta^{\mathbf{w}}_{\mathcal{H}'}$. Combining this with Lemmas 5.6.8 and 5.6.9, we have

$$\Big(\Phi_{\mathcal{H}\to\mathcal{H}'}(G_{\mathcal{H}})\Big)(\Theta^{\mathbf{w}}_{\mathcal{H}'}) = \Big(\Phi_{\mathcal{H}\to\mathcal{H}'}(G_{\mathcal{H}})\Big)(\Phi_{\mathcal{H}\to\mathcal{H}'}(\Theta^{\mathbf{w}}_{\mathcal{H}})) = G_{\mathcal{H}}(\Theta^{\mathbf{w}}_{\mathcal{H}}),$$

which is equal by definition, to $G_{\mathcal{H}'}(\Theta^{\mathbf{w}}_{\mathcal{H}'})$. Hence $(\Phi_{\mathcal{H}\to\mathcal{H}'}(G_{\mathcal{H}}))(\Theta^{\mathbf{w}}_{\mathcal{H}'}) = G_{\mathcal{H}'}(\Theta^{\mathbf{w}}_{\mathcal{H}'})$, completing the proof. □

5.7 Absolute Alexander multi-gradings

In this section, we describe a distinguished, coherent Alexander multi-grading on $\mathcal{C}FL^{\circ}(Y,\mathbb{L},\mathfrak{s})$ when L is J–null-homologous with respect to a grading assignment $\pi : C(L) \to J$.

Our approach is modeled on the construction of the absolute grading in [OS06]. Given a based link $\mathbb{L}$ in a 3–manifold Y, we will pick a parametrized Kirby diagram $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ of $(Y,\mathbb{L})$ and take Heegaard triples representing surgery on $\mathbb{S}_1$ in $S^3 \setminus N(U)$ (where U is a collection of unknots) which represent $Y \setminus N(L)$. The Alexander grading is then defined using these Heegaard triples. Using the Kirby calculus argument from Section 5.4, we show that the grading is independent of the parametrized Kirby diagram and the Heegaard triple.

5.7.1 Construction of the absolute grading $A_{Y,\mathbb{L},S,\mathbb{P}}$

We now define the grading, in terms of a choice of parametrized Kirby diagram and some other auxiliary data. Suppose that $(Y,\mathbb{L})$ is a multi-based link in a three manifold, $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, and $\pi : C(L) \to J$ is a grading assignment, and that L is J–null-homologous. Suppose $S \subseteq Y$ is a J–Seifert surface, and λ is a choice of framing for L. Finally, suppose that $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a parametrized Kirby diagram of $(Y,\mathbb{L})$.

We now pick a triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ which is subordinate to a β–bouquet for $\mathbb{S}_1$. By assumption $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(S^3, \mathbb{U})$ where $\mathbb{U} = (U, \mathbf{w}, \mathbf{z})$ is an unlink, and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y,\mathbb{L})$. Furthermore, the homeomorphism ϕ_0 in $\mathbb{P}$ gives an identification between the components of $\mathbb{L}$ and $\mathbb{U}$, so we naturally get a grading assignment $\pi \circ \phi_0 : C(U) \to J$, by using the grading assignment for $\mathbb{L}$.

The diagram $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ represents an unlink, say $\mathbb{U}' = (U', \mathbf{w}, \mathbf{z})$, in $(S^1 \times S^2)^{\#k}$,

though the number of components of $\mathbb{U}'$ may be greater than $\mathbb{U}$, since each component of $\mathbb{U}'$ has exactly two basepoints. Note however, that the two basepoints on a link component of $\mathbb{U}'$ are on the same link component of the link $\mathbb{L}$ in Y. Hence there is a natural map $C(U') \to C(L)$, which we can compose with the grading assignment for $\mathbb{L}$ to get a grading assignment $\pi' : C(U') \to J$, hence allowing us to define an Alexander J–grading on $\mathcal{C}FL^\circ((S^1 \times S^2)^{\#k}, \mathbb{U}', \mathfrak{s}_0)$ by collapsing the standard grading defined in Section 5.5.4.

Given a homology triangle $\psi \in \pi_2(\mathbf{x}, \theta, \mathbf{y})$ and a grading $j \in J$, we define

$$A_{\mathcal{T}}(\mathbf{y})_j = A_{S^3,\mathbb{U}}(\mathbf{x})_j + A_{(S^1\times S^2)^{\#k},\mathbb{U}'}(\theta)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi)_j + \frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi)), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]}{2}, \quad (5.13)$$

where $\widehat{\Sigma}_j$ is the closed surface in $W(S^3, \mathbb{S}_1)$ obtained by taking $\mathbb{U}_j \times [0,1]$ in $W(S^3, L)$ and capping it off with any choice of Seifert surface for $\mathbb{U}_j$ in S^3, and $f^{-1}(S_j)$ in $S^3(\mathbb{S}_1)$. Here $\mathbb{U}_j$ denotes the components of $\mathbb{U}_j$ assigned grading j. Here θ is any intersection point, but it's worth noting that $\mathbb{U}'$ has exactly two basepoints per component, so $\widehat{\mathcal{H}FL}((S^1 \times S^2)^{\#k}, \mathbb{U}')$ is supported in a single Alexander grading. Since S^3 is an integer homology sphere, all choices of J–Seifert surfaces are homologous, so the above formula doesn't depend on the choice of Seifert surface in S^3.

We then define $A_{Y,\mathbb{L},S,\mathbb{P}} \in \overrightarrow{\mathbb{A}}(Y, \mathbb{L}, \pi, \mathfrak{s})$ to be the coherent grading

$$A_{Y,\mathbb{L},S,\mathbb{P}} = \eta_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma}),\mathbb{A}}(A_{\mathcal{T}}),$$

where $\eta_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma}),\mathbb{A}} : \mathbb{A}((\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}), \pi, \mathfrak{s}) \to \overrightarrow{\mathbb{A}}(Y, \mathbb{L}, \pi, \mathfrak{s})$ is the inclusion map defined in the previous section. In the next section, we show that this is well defined.

Remark 5.7.1. It is easy to see that $\langle c_1(\mathfrak{s}), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}_j]\cdot[\widehat{\Sigma}_j]$ is always an even integer. As such, the j–component of the grading takes values in

$$\mathbb{Z} + \tfrac{1}{2}[\widehat{\Sigma} \setminus \widehat{\Sigma}_j]\cdot[\widehat{\Sigma}_j].$$

On the other hand, we note that it is easily seen that $[\widehat{\Sigma} \setminus \widehat{\Sigma}_j]\cdot[\widehat{\Sigma}_j] = \pm\ell k(L \setminus L_j, L_j)$, since the linking numbers of two null-homologous links are defined by picking a Seifert surface for one, and computing the intersection number of the other link with the Seifert surface.

5.7.2 Invariance of the absolute Alexander multi-grading

In this subsection, we prove invariance of the absolute Alexander multi-gradings from the Heegaard triple $\mathcal{T}$, subordinate to a bouquet of $\mathbb{S}_1$, and the parametrized Kirby diagram $\mathbb{P}$. Invariance of the grading under Moves $\mathcal{O}_1$ and $\mathcal{O}_2$ follows from a similar argument to the one in [OS06]. Our argument deviates from [OS06] when we prove invariance from Move $\mathcal{O}_3'$ (and hence the choice of framing λ). Throughout this section, let $\pi : C(L) \to J$ be a fixed grading assignment and suppose that $\mathbb{L}$ is J–null-homologous.

Lemma 5.7.2. *The coherent grading $A_{\mathcal{T}}$ is independent of isotopies of f (relative $\mathbb{L}$) or isotopies of $\mathbb{S}_1$ (relative $\mathbb{U}$).*

Proof. This is essentially a tautology. Isotopies of f, relative $\mathbb{L}$, for a fixed link $\mathbb{S}_1 \subseteq S^3_U$, result in an isotopy of the pullback of S under f. The topological quantities appearing at the end of the formula are unchanged. Thus an isotopy of f results in a grading which is related to the original grading by the change of diagrams map (on sets of gradings) associated to the isotopy. Isotopies of $\mathbb{S}_1$ are handled analogously. □

Lemma 5.7.3. *For a fixed choice of β–bouquet for $\mathbb{S}_1$, and a fixed Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ subordinate to it, the grading $A_{\mathcal{T}}$ defined in Equation* (5.13) *defines an element of $\mathbb{A}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \pi, \mathfrak{s})$ which is independent of the choice of intersection points $\mathbf{x}$ and θ, as well as the homology triangle $\psi \in \pi_2(\mathbf{x}, \theta, \mathbf{y})$.*

Proof. To see that $A_{\mathcal{T}} \in \mathbb{A}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \pi, \mathfrak{s})$, we note that if $\mathbf{y}$ and $\mathbf{y}'$ are two intersection points representing $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, then there is a disk $\phi \in \pi_2(\mathbf{y}', \mathbf{y})$, and one sees from the definition that $A_{\mathcal{T}}$ satisfies

$$A_{\mathcal{T}}(\mathbf{y}')_j - A_{\mathcal{T}}(\mathbf{y})_j = (n_{\mathbf{z}} - n_{\mathbf{w}})(\phi)_j = A(\mathbf{y}', \mathbf{y})_j.$$

We now show that $A_{\mathcal{T}}$ is independent of the triangle ψ and the choice of $\mathbf{x}$ and θ. We first show that for fixed $\mathbf{x}$ and θ, it is independent of ψ. If $\psi, \psi' \in \pi_2(\mathbf{x}, \theta, \mathbf{y})$ are two homology triangles which share the same endpoints, we can write

$$\psi' = \psi + \mathcal{P},$$

for a triply periodic domain $\mathcal{P}$. We have by [OS04e, Proposition 8.5], that

$$\mathfrak{s}_{\mathbf{w}}(\psi') = \mathfrak{s}_{\mathbf{w}}(\psi) + q_* PD[H(\mathcal{P})],$$

where $q_* : H^2(X_{\alpha\beta\gamma}, \partial X_{\alpha\beta\gamma}; \mathbb{Z}) \to H^2(X_{\alpha\beta\gamma}; \mathbb{Z})$ is the map in the long exact sequence of cohomology.

Let A denote the grading defined with the triangle ψ, and let A' denote the grading defined with ψ'. We now simply compute that

$$\begin{aligned}
A'(\mathbf{y})_j &= A(\mathbf{x})_j + A(\theta)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi')_j + \frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi')), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]}{2} \\
&= A(\mathbf{x})_j + A(\theta)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\mathcal{P})_j \\
&\quad + \frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi)) + 2q_* PD[H(\mathcal{P})], [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]}{2} \\
&= A(\mathbf{x})_j + A(\theta)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\mathcal{P})_j + \langle q_* PD[H(\mathcal{P})], \widehat{\Sigma}_j\rangle \\
&\quad + \frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi)), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]}{2} \\
&= A(\mathbf{y})_j
\end{aligned}$$

since

$$\begin{aligned}
(n_{\mathbf{w}} - n_{\mathbf{z}})(\mathcal{P})_j + \langle q_* PD[H(\mathcal{P})], \widehat{\Sigma}_j\rangle &= (n_{\mathbf{w}} - n_{\mathbf{z}})(\mathcal{P})_j + \langle PD[H(\mathcal{P})], \Sigma_j\rangle \\
&= (n_{\mathbf{w}} - n_{\mathbf{z}})(\mathcal{P})_j + \langle PD[H(\mathcal{P})], (\Sigma_{\alpha\beta\gamma})_j\rangle \\
&= 0
\end{aligned}$$

by Equation (5.2). □

We now prove that the grading is invariant under the particular Heegaard triple subordinate to a given β–bouquet of $\mathbb{S}_1$.

Lemma 5.7.4. *For a fixed β–bouquet $\mathcal{B}^\beta$ of $\mathbb{S}_1$, the coherent grading $\eta_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma},\mathbb{A})}(A_{\mathcal{T}})$ on $\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$ is independent from the Heegaard triple $\mathcal{T}$ subordinate to $\mathcal{B}^\beta$.*

Proof. Any two triples subordinate to $\mathcal{B}^\beta$ can be connected by the six moves of Lemma 5.3.6. Hence we need to show that if $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ and $\mathcal{T}' = (\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}')$ differ by a move in

the above list, then

$$\Phi_{(\Sigma,\boldsymbol{\alpha},\boldsymbol{\gamma})\to(\Sigma',\boldsymbol{\alpha}',\boldsymbol{\gamma}')}(A_{\mathcal{T}}) = A_{\mathcal{T}'}.$$

Consider Move (1). Supposing that $\boldsymbol{\alpha}'$ differs from $\boldsymbol{\alpha}$ by a sequence of handleslides or isotopies, we need to show that

$$A_{\mathcal{T}'} = \Phi_{\boldsymbol{\gamma}}^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}(A_{\mathcal{T}}).$$

By functoriality of the change of diagrams maps on gradings, we can assume that the quadruple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ is admissible.

Suppose that $\psi_{\alpha\beta\gamma} \in \pi_2(\mathbf{x}_{\alpha\beta}, \theta_{\beta\gamma}, \mathbf{y}_{\alpha\gamma})$ and $\psi_{\alpha'\alpha\beta} \in \pi_2(\mathbf{y}_{\alpha'\alpha}, \theta_{\alpha\beta}, \mathbf{y})$ are homology triangles. Since $\mathrm{Spin}^c(X_{\alpha'\beta\gamma}) \cong \mathrm{Spin}^c(X_{\alpha\beta\gamma})$, we can write

$$\psi_{\alpha\beta\gamma} + \psi_{\alpha'\alpha\gamma} = \psi_{\alpha'\alpha\beta} + \psi_{\alpha'\beta\gamma}$$

for homology triangles $\psi_{\alpha'\alpha\beta} \in \pi_2(\theta_{\alpha'\alpha}, \mathbf{x}_{\alpha\beta}, \mathbf{x}_{\alpha'\beta})$ and $\psi_{\alpha'\beta\gamma} \in \pi_2(\mathbf{x}_{\alpha'\beta}, \theta_{\beta\gamma}, \mathbf{y})$ with $\mathfrak{s}_{\mathbf{w}}(\psi_{\alpha\beta\gamma}) = \mathfrak{s}_{\mathbf{w}}(\psi_{\alpha'\beta\gamma})$. By definition of the gradings, we have

$$A_{\mathcal{T}'}(\mathbf{y})_j - \Phi_{\boldsymbol{\gamma}}^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}(A_{\mathcal{T},\mathfrak{s}})(\mathbf{y})_j = A(\mathbf{x}_{\alpha'\beta})_j - A(\mathbf{x}_{\alpha\beta})_j - A(\theta_{\alpha'\alpha})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha'\beta\gamma} - \psi_{\alpha\beta\gamma} - \psi_{\alpha'\alpha\gamma})_j$$

$$= A(\mathbf{x}_{\alpha'\beta})_j - A(\mathbf{x}_{\alpha\beta})_j - A(\theta_{\alpha'\alpha})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha'\alpha\beta})_j = A(\mathbf{x}_{\alpha'\beta})_j - \Phi_{\boldsymbol{\beta}}^{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'}(\mathbf{x}_{\alpha\beta})_j,$$

which is zero since our declared absolute Alexander gradings on unlinks in $(S^1 \times S^2)^{\#k}$ are coherent by Lemma 5.6.10.

Independence from Moves (2), (4) and (5) follow in an identical fashion. Move (3) is an easy computation, analogous to Lemma 5.6.3, and Move (6) is a tautology. □

We now consider dependence on the choice of β–bouquet $\mathcal{B}^\beta$ for $\mathbb{S}_1$:

Lemma 5.7.5. *The absolute grading is independent from the choice of bouquet for a fixed, framed one dimensional link $\mathbb{S}_1 \subseteq Y \setminus L$.*

Proof. This follows from the proof of [Juh16, Thm. 6.9]. We change the bouquet one arc at a time. A change of a single arc in the bouquet to a different arc can be achieved by

performing a sequence of handleslides and isotopies amongst the $\boldsymbol{\beta}$–curves, and amongst the $\boldsymbol{\gamma}$–curves. Invariance under these moves is then achieved as in Lemma 5.7.4. □

Lemma 5.7.6. *The absolute grading is independent from handleslides amongst the components of the framed link $\mathbb{S}_1$. In particular, the absolute grading is invariant under move $\mathcal{O}_1$.*

Proof. This follows again as in [OS06]. Handlesliding components of $\mathbb{S}_1$ across each other can be realized as handleslides of the $\boldsymbol{\beta}$– and $\boldsymbol{\gamma}$–curves across each other, and we've already shown invariance under those moves. □

We now consider move $\mathcal{O}_2$, corresponding to blowing up or blowing down along a ± 1 framed unknot K which is unlinked from $\mathbb{S}_1$ and U:

Lemma 5.7.7. *Suppose that K is an unknot in S^3 which is contained in a ball which is disjoint from $\mathbb{S}_1$ and U. Suppose that $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a parametrized Kirby diagram for $(Y, \mathbb{L})$ and let $\mathbb{P}' = (\phi_0, \lambda, \mathbb{S}_1 \cup \{K\}, f_K)$ denote the parametrized Kirby diagram obtained by adding K to $\mathbb{S}_1$ with framing ± 1, and let f_K be the induced diffeomorphism. The gradings $A_{Y,\mathbb{L},S,\mathbb{P}}$ and $A_{Y,\mathbb{L},S,\mathbb{P}'}$ agree.*

Proof. If $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ represents surgery on $\mathbb{S}_1$, then we can form a triple representing surgery on $\mathbb{S}_1 \cup \{K\}$ by taking the connected sum of $\mathcal{T}$ with the diagram shown in Figure 5.9.

If ψ is a triangle on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$, we can pick a homology triangle ψ' with the same multiplicity in the connect sum region and form $\psi \# \psi'$. We note that

$$\mathfrak{s}_{\mathbf{w}}(\psi \# \psi') = \mathfrak{s}_{\mathbf{w}}(\psi) \# \mathfrak{s}' \tag{5.14}$$

for some $\mathfrak{s}'$ in $\mathrm{Spin}^c(\mathbb{CP}^2)$ or $\mathrm{Spin}^c(\overline{\mathbb{CP}}^2)$. A diagram is shown in Figure 5.9, with a homology triangle drawn in. Using Equation (5.14)

$$\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi \# \psi')), \widehat{\Sigma}_j \rangle = \langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi)), \widehat{\Sigma}_j \rangle$$

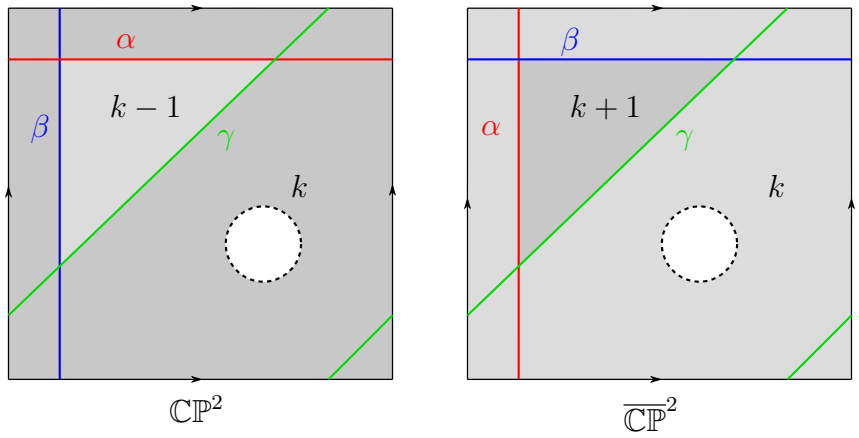

Figure 5.9: **Taking connected sum of a surgery triple with one of the two triples shown above, at the dashed circle shown, results in a surgery triple for blowing up away from Σ.** Multiplicities of a homology triangles are shown.

and that $[\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]$ is also unchanged. Since $(n_{\mathbf{w}} - n_{\mathbf{z}})(\psi)_j$ is also unchanged, the Alexander grading is unchanged and $A_{Y,\mathbb{L},S,\mathbb{P}}$ and $A_{Y,\mathbb{L},S,\mathbb{P}'}$ agree. □

We will now show that the absolute gradings are invariant under move $\mathcal{O}'_3$, for -1 framed unknots.

Lemma 5.7.8. *Suppose that $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a parametrized Kirby diagram for $(Y, \mathbb{L})$ and that K is an unknot which is unlinked from all components of $\mathbb{S}_1$ and U, except one component K of $\mathbb{S}_1$, and to K it is linked as in move $\mathcal{O}'_3$. Suppose K has framing -1. Let $\mathbb{P}' = (\phi_0, \lambda', \mathbb{S}_1 \cup \{K\}, f_K)$ where f_K is the induced diffeomorphism, and λ' is the new framing on L. The gradings $A_{Y,\mathbb{L},S,\mathbb{P}}$ and $A_{Y,\mathbb{L},S,\mathbb{P}'}$ agree.*

Proof. A Heegaard triple for surgery on $\mathbb{S}_1 \cup \{K\}$ can be obtained by taking the connected sum of a diagram for surgery on $\mathbb{S}_1$ and the diagram shown in Figure 5.10, near a $\mathbf{z}$–basepoint on the link component that K encircles (note the placement of the z–basepoint in the connected sum). Let $j \in J$ be the grading assigned to the basepoint z.

We note that a there is a diffeomorphism between $W(S^3, \mathbb{S}'_1)$ and $W(S^3, \mathbb{S}_1) \# \overline{\mathbb{CP}}^2$ which is the identity outside a neighborhood of $B \times \{1\} \subseteq W(S^3, \mathbb{S}_1)$ for a ball $B \subseteq S^3$ con-

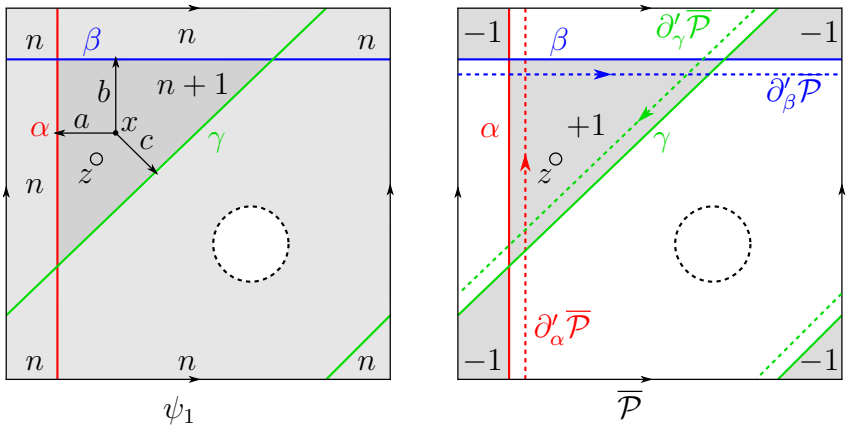

Figure 5.10:

A diagram for surgery on $\mathbb{S}_1 \cup \{K\}$ where K has framing -1. This is move $\mathcal{O}_3'$. On the right is the triply periodic domain $\overline{\mathcal{P}}$ with $H(\overline{\mathcal{P}}) \in H_2(\overline{\mathbb{CP}}^2;\mathbb{Z})$ a generator. On the left is the homology triangle ψ_1. On the left is also a dual spider with basepoint x and legs a, b and c. On the right are also the translates $\partial'_\alpha\overline{\mathcal{P}}, \partial'_\beta\overline{\mathcal{P}}$ and $\partial'_\gamma\overline{\mathcal{P}}$.

taining K. Under this connected sum decomposition, we can describe $H_2(W(S^3,\mathbb{S}_1');\mathbb{Z})$ as $H_2(W(S^3,\mathbb{S}_1);\mathbb{Z})\oplus\mathbb{Z}$ where the copy of $\mathbb{Z}$ is generated by an embedded sphere E in $W(S^3,\mathbb{S}_1')$ formed by taking the obvious Seifert disk for K in $S^3\times\{1\}$ and attaching the core of the 2–handle for K. Let Σ_j denote the surface $U\times[0,1]$ in $W(S^3,\mathbb{S}_1)$ and let Σ_j' denote the analogous surface in $W(S^3,\mathbb{S}_1')$. Under the identification of $W(S^3,\mathbb{S}_1')$ as $W(S^3,\mathbb{S}_1)\#\overline{\mathbb{CP}}^2$ we now claim that the surface Σ_j' is equal to the connected sum of Σ_j (pushed outside of the connected sum region) and the sphere E. To see this, we note that we can push E radially out of the 2–handle attached along K, and give E the movie description shown in Figure 5.11, i.e. first an addition of a disk D_1 to the surface, then a 2–handle attached along K, then a disk D_2 attached to the surface. Inside of the 4–manifold $W(S^3,\mathbb{S}_1')$, the sphere E is the union of D_1 and D_2.

We can manipulate the movie for $\widehat{\Sigma}_j'$ by creating a 2–dimensional index 0/1 birth death which together cancel a small isotopy of the link, followed by surgery along K, then an additional 2–dimensional index 1/2 birth death followed by an isotopy of the link across the surgered solid torus. This is shown in Figure 5.12. We can move the band B_2 below

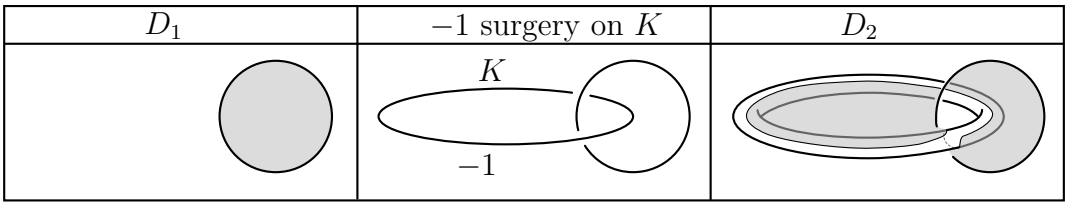

Figure 5.11: **A movie presentation of the sphere E inside of $W(S^3, \mathbb{S}'_1)$.** The disks D_1 and D_2 are shaded. On the right, we are showing only the complement of a neighborhood of K, inside of the surgered manifold $S^3(\mathbb{S}'_1)$. An annulus of D_2 is shown. The rest is D_2 compressing disk inside of the solid torus, which is not shown.

the 2–handle attachment since they don't intersect, and we note that B_1 and B_2 are "dual bands". This says that precisely Σ'_j is formed by taking Σ_j, pushed out of B (via the isotopies appearing as the first and last moves of the decomposition in Figure 5.12, and forming Σ'_j as the connected sum with the sphere E from Figure 5.11.

Note also that the homology element $H(\overline{\mathcal{P}})$ determined by the triply periodic domain $\overline{\mathcal{P}}$ can be seen to be $\pm E$, directly from its construction. The sign can be determined to be -1 since writing $\Sigma'_j = \Sigma_j \# aE$ and applying Equation (5.2) yields

$$1 = (n_{\mathbf{z}}(\overline{\mathcal{P}}) - n_{\mathbf{w}}(\overline{\mathcal{P}}))_j = \#(\Sigma'_j \cap H(\overline{\mathcal{P}})) = [\widehat{\Sigma}_j + aH(\overline{\mathcal{P}})] \cdot [H(\overline{\mathcal{P}})] = -a,$$

so $a = -1$.

Now letting ψ_1 be the triangle on the left of Figure 5.10, we compute $\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_1)), H(\overline{\mathcal{P}})\rangle$ using [OS06, Prop. 6.3]. According to [OS06, Prop. 6.3], we have

$$\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_1)), H(\overline{\mathcal{P}})\rangle = e(\overline{\mathcal{P}}) + \#(\partial\overline{\mathcal{P}}) - 2n_{\mathbf{w}}(\overline{\mathcal{P}}) + 2\sigma(\psi_1, \overline{\mathcal{P}}),$$

where e denotes the Euler measure, $\#\partial\overline{\mathcal{P}} = 3$ is the number of boundary components of $\overline{\mathcal{P}}$, and $\sigma(\psi_1, \overline{\mathcal{P}})$ is the "dual spider number", a quantity defined in [OS06], by the formula

$$\sigma(\phi_1, \overline{\mathcal{P}}) = n_{\phi_1(x)}(\overline{\mathcal{P}}) + \#(a \cap \partial'_a\overline{\mathcal{P}}) + \#(b \cap \partial'_\beta\overline{\mathcal{P}}) + \#(c \cap \partial'_\gamma\overline{\mathcal{P}}),$$

where a, b and c are arcs of the "dual spider", and $\partial'_\tau\overline{\mathcal{P}}$ denotes the translation of the boundary component $\partial_\tau\overline{\mathcal{P}}$ of $\overline{\mathcal{P}}$, in the direction of the inward normal vector field, according

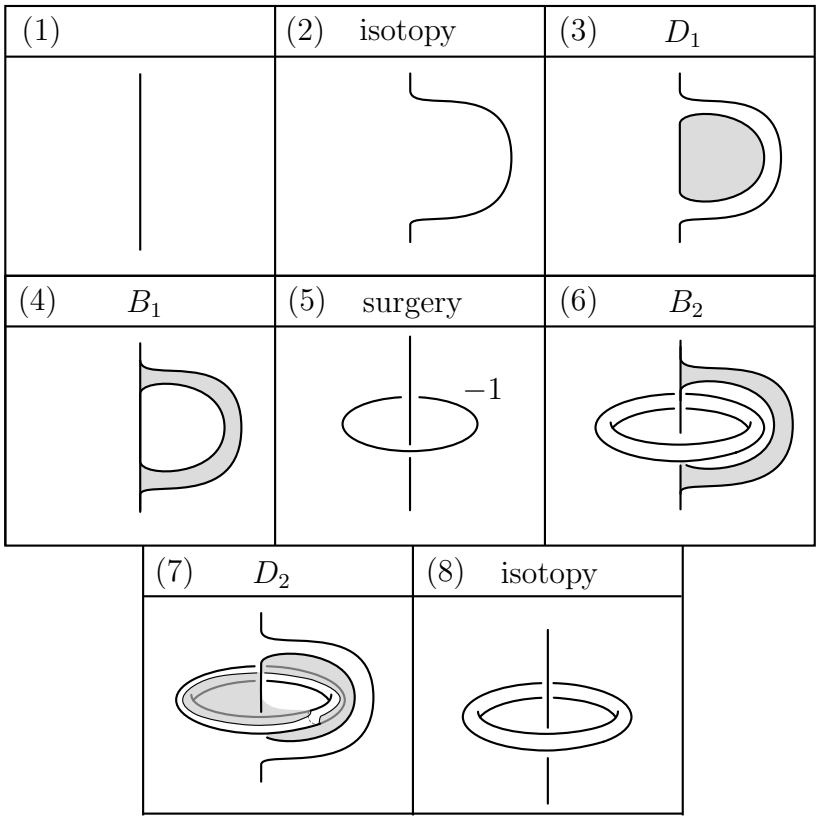

Figure 5.12: **A movie for $\Sigma'_j \subseteq W(S^3, \mathbb{S}'_1)$,** obtained by adding canceling 2–dimensional 0/1 and 1/2 critical points to Σ'_j. The shaded regions are the disks D_1 and D_2, as well as the bands B_1 and B_2. Moving the band B_2 below the 2–handle added along K, we get a description of Σ'_j inside of $W(S^3, \mathbb{S}_1)\#\overline{\mathbb{CP}}^2$ as $\Sigma_j \# E$, with Σ_j pushed out of the connected sum region. In the steps after -1 surgery on K, we are only showing $S^3 \setminus N(K) \subseteq S^3(K)$.

to the periodic domain $\overline{\mathcal{P}}$. We compute now that $\sigma(\psi_1, \overline{\mathcal{P}}) = 1 - 1 - 1 - 1 = -2$. Adding the above terms up we get

$$\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_1)), H(\overline{\mathcal{P}})\rangle = 3 - 4 = -1.$$

Now let ψ_k be the triangle defined by

$$\psi_k = \psi_1 + (k-1) \cdot \overline{\mathcal{P}}.$$

We have

$$\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_k)), H(\overline{\mathcal{P}})\rangle = \langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_1)), H(\overline{\mathcal{P}})\rangle + 2(k-1)\langle PD[H(\overline{\mathcal{P}})], H(\overline{\mathcal{P}})\rangle$$

$$= -1 - 2(k-1) = 1 - 2k.$$

Notice also that $(n_{\mathbf{w}})_j$ is unchanged, but $(n_{\mathbf{z}})_j$ increases by $n_z(\psi_k) - n = k$ for ψ_k. Hence the j component of the Alexander grading changes by

$$\frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_k)), -H(\overline{\mathcal{P}})\rangle - [-H(\overline{\mathcal{P}})] \cdot [-H(\overline{\mathcal{P}})]}{2} - (n_z(\psi_k) - n)$$

$$= \frac{2k - 1 + 1}{2} - k = 0.$$

In a similar way to Lemma 5.7.7, one sees that the other components of the Alexander grading are unchanged. □

We now prove invariance under move $\mathcal{O}_3'$ in the case that the unknot K has framing $+1$. This is similar to the last lemma.

Lemma 5.7.9. *Suppose that $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a parametrized Kirby diagram for $(Y, \mathbb{L})$ and suppose that K is an unknot which is unlinked from all components of $\mathbb{S}_1$ and U, except one component of U, and it is linked as in move $\mathcal{O}_3'$. Suppose K has framing $+1$. Let $\mathbb{P}' = (\phi_0, \lambda', \mathbb{S}_1 \cup \{K\}, f_K)$, where f_K is the induced diffeomorphism, and λ' is the new framing on L. The two gradings $A_{Y,\mathbb{L},S,\mathbb{P}}$ and $A_{Y,\mathbb{L},S,\mathbb{P}'}$ agree.*

Proof. The proof follows analogously to the proof of Lemma 5.7.8. Again let j denote the index of the grading that K is assigned to. Let ψ_{-1} and $\mathcal{P}$ be the homology triangle and triply periodic domain shown in Figure 5.13.

Let Σ_j' denote the surface $U \times [0,1]$ in $W(S^3, \mathbb{S}_1')$ and let Σ_j denote the surface $U \times [0,1]$ in $W(S^3, \mathbb{S}_1)$. As in Lemma 5.7.8, we can write $\Sigma_j' = \Sigma_j \# aH(\mathcal{P})$ if we identify $W(S^3, \mathbb{S}_1')$ as $W(S^3, \mathbb{S}_1) \# \mathbb{CP}^2$. Arguing as before, we have now

$$1 = (n_{\mathbf{z}}(\mathcal{P}) - n_{\mathbf{w}}(\mathcal{P}))_j = \#(\Sigma_j' \cap H(\mathcal{P})) = a[H(\mathcal{P})] \cdot [H(\mathcal{P})] = a,$$

implying that $\Sigma_j' = \Sigma_j \#(+H(\mathcal{P}))$. Computing as before we have that

$$\sigma(\psi_{-1}, \mathcal{P}) = 1 - 1 - 1 - 1 = -2$$

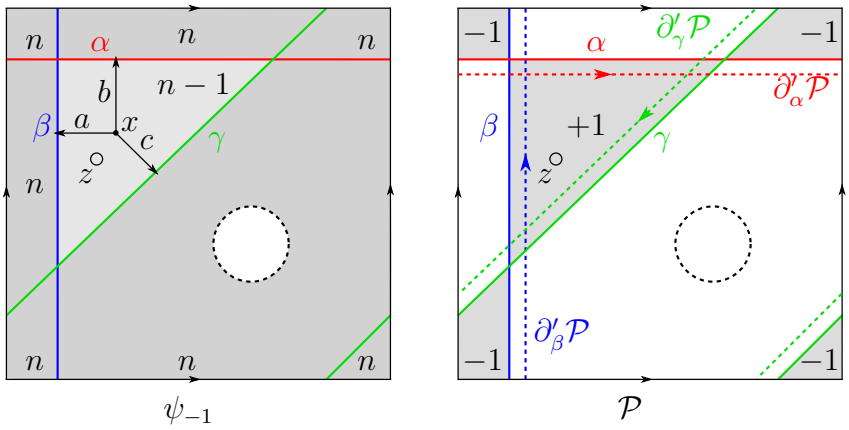

Figure 5.13: **A diagram for surgery on $\mathbb{S}_1 \cup \{K\}$ where K has framing $+1$.** This is move $\mathcal{O}_3'$. Here $\mathcal{P}$ is the periodic domain shown on the right which represents a generator $H(\mathcal{P}) \in H_2(\mathbb{CP}^2;\mathbb{Z})$. On the left is the homology triangle ψ_{-1}. We can assume that the diagram has a $\mathbf{w}$–basepoint in the same region as the dashed circle. On the left is also a dual spider with basepoint x and legs a, b and c. On the right are the translates $\partial'_\alpha\mathcal{P}, \partial'_\beta\mathcal{P}$ and $\partial'_\gamma\mathcal{P}$.

$$\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_{-1})), H(\mathcal{P})\rangle = 0 + 3 + 0 - 4 = -1.$$

Let ψ_k be the triangle defined by $\psi_k = \psi_{-1} + (k+1)\mathcal{P}$. We have

$$\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_k)), H(\mathcal{P})\rangle = \langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_{-1})), H(\mathcal{P})\rangle + 2(k+1)\langle PD[H(\mathcal{P})], H(\mathcal{P})\rangle$$

$$= 2k + 1.$$

Hence blowing up along K and using the Spin^c structure $\mathfrak{s}_{\mathbf{w}}(\psi_k)$ thus produces a net change in the j component Alexander grading of

$$\frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_k)), H(\mathcal{P})\rangle - [H(\mathcal{P})]\cdot[H(\mathcal{P})]}{2} - (n_z(\psi_k) - n) = \frac{(2k+1)-1}{2} - k = 0.$$

As in the previous two lemmas, the Alexander grading is unchanged in the components other than j. □

We now consider changing the identification ϕ_0 of $\mathbb{U}$ with $\mathbb{L}$. Suppose that $\psi : (S^3, \mathbb{U}) \to (S^3, \mathbb{U})$ is an orientation preserving diffeomorphism (of 3–manifolds with multi-based links),

and suppose that $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a parametrized Kirby diagram for $(Y, \mathbb{L})$. Note that ψ determines a diffeomorphism $\psi^{\mathbb{S}_1} : S^3_U(\mathbb{S}_1) \to S^3_U(\psi(\mathbb{S}_1))$. Hence we can form the parametrization $\mathbb{P}' = (\phi_0, \lambda, \psi(\mathbb{S}_1), f \circ (\psi^{\mathbb{S}_1})^{-1})$.

Lemma 5.7.10. *Suppose that $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ is a parametrized Kirby diagram of $(Y, \mathbb{L})$ and that $\psi : (S^3, \mathbb{U}) \to (S^3, \mathbb{U})$ is an orientation preserving diffeomorphism (of 3–manifolds with multi-based links). Then $A_{Y,\mathbb{L},S,\mathbb{P}}$ and $A_{Y,\mathbb{L},S,\mathbb{P}'}$ agree, where $\mathbb{P}' = (\phi_0 \circ \psi, \lambda, \psi(\mathbb{S}_1), f \circ (\psi^{\mathbb{S}_1})^{-1})$.*

Proof. We take a triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ subordinate to a β–bouquet of $\mathbb{S}_1$. The triple $\psi_*\mathcal{T} = (\psi_*\Sigma, \psi_*\boldsymbol{\alpha}, \psi_*\boldsymbol{\beta}, \psi_*\boldsymbol{\gamma})$ is subordinate to a β–bouquet for $\psi(\mathbb{S}_1)$. If $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ and $\psi_*\mathbf{y}$ denotes the corresponding intersection point on the surface $\psi_*\Sigma \subseteq S^3$, then

$$f_*(\mathbf{y}) = (f \circ (\psi^{\mathbb{S}_1})^{-1})_*(\psi_*\mathbf{y}).$$

Since the absolute gradings for unlinks in S^3 are determined by the multi-grading of the top $\mathrm{gr}_{\mathbf{w}}$–graded element, and ψ_* preserves the top graded element and also preserves the relative Alexander multi-grading, we know that $A(\psi_*(\theta))_j = A(\theta)_j$ for $\theta \in \mathbb{T}_\beta \cap \mathbb{T}_\gamma$, and similarly $A(\psi_*(\mathbf{x}))_j = A(\mathbf{x})_j$ for $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$. Similarly, the value of $\frac{1}{2}(\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi)), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j])$ is tautologically unchanged. As such, the formulas for $A_{Y,\mathbb{L},S,\mathbb{P}}$ and $A_{Y,\mathbb{L},S,\mathbb{P}'}$ agree, so the gradings are the same. □

Combining the results of this section, we can prove part (a) of Theorem 5.1.3:

Theorem 5.1.3(a). *For a fixed J–Seifert surface S of $\mathbb{L}$, the absolute gradings $A_{Y,\mathbb{L},S,\mathbb{P}}$ depend only on the tuple $(Y, \mathbb{L}, S)$.*

Proof. By Lemma 5.7.3, for fixed $\mathbb{P}$, the gradings $A_{Y,\mathbb{L},S,\mathbb{P}}$ depend at most on the Heegaard triple $\mathcal{T}$ and $\mathbb{P}$. By Lemma 5.7.4 the grading is independent of the choice of Heegaard triple $\mathcal{T}$ subordinate to a fixed bouquet for $\mathbb{S}_1$. By Lemma 5.7.5, the grading is independent of the choice of bouquet subordinate to $\mathbb{S}_1$. It follows that $A_{Y,\mathbb{L},S,\mathbb{P}}$ depends on only $Y, \mathbb{L}, S$ and $\mathbb{P}$. By Lemma 5.4.4, independence from $\mathbb{P}$ will follow from showing independence of

Moves $\mathcal{O}_0, \mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3'$ and $\mathcal{O}_4$. Invariance from $\mathcal{O}_0$ follows from Lemma 5.7.2. Invariance under Move $\mathcal{O}_1$ follows from Lemma 5.7.6. Invariance from Move $\mathcal{O}_2$ follows from Lemma 5.7.7. Invariance from Move $\mathcal{O}_3'$ follows from Lemmas 5.7.8 and 5.7.9. Finally, invariance from Move $\mathcal{O}_4$ follows from Lemma 5.7.10.

□

5.7.3 Gradings using α–bouquets

We defined the absolute grading using β–bouquets, but it will be useful to know that α–bouquets can be used as well. For definiteness, let us write $A^{\beta}_{Y,\mathbb{L},S}$ for the gradings defined in the previous section.

Repeating the above procedure, but using α–bouquets instead, yields an absolute grading $A^{\alpha}_{Y,\mathbb{L},S}$ on $\mathcal{C}FL^{\circ}(Y,\mathbb{L},\mathfrak{s})$, defined using α–bouquets.

Lemma 5.7.11. *The gradings $A^{\alpha}_{Y,\mathbb{L},S}$ and $A^{\beta}_{Y,\mathbb{L},S}$ are equal.*

Proof. We can think of this as a grading analog of [OS06, Lemma 5.2]. If $\mathbb{S}_1$ is a framed 1–dimensional link in $Y \setminus L$, and S' is a choice of J–Seifert surface in $Y(\mathbb{S}_1)$ for L, and $\mathfrak{s} \in \mathrm{Spin}^c(W(Y,\mathbb{S}_1))$, we can define two maps, $F^{\alpha}_{\mathbb{S}_1,S,S',\mathfrak{s}}$ and $F^{\beta}_{\mathbb{S}_1,S,S',\mathfrak{s}}$ from $\overrightarrow{\mathbb{A}}(Y,\mathbb{L},\pi,\mathfrak{s}|_Y)$ to $\overrightarrow{\mathbb{A}}(Y,\mathbb{L},\pi,\mathfrak{s}|_{Y(\mathbb{S}_1)})$. Each is defined by taking a surgery diagram for a bouquet (either an α–bouquet or a β–bouquet, as appropriate) and then defining the grading $F^{\tau}_{\mathbb{S}_1,S,S',\mathfrak{s}}(A)$ by the formula

$$F_{\mathbb{S}_1,S,S',\mathfrak{s}}(A)(\mathbf{y})_j = A(\mathbf{x})_j + A_{(S^1\times S^2)^{\#k},\mathbb{U}}(\theta)_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi)_j + \frac{\langle c_1(\mathfrak{s}), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]}{2}$$

where $\widehat{\Sigma}_j$ is formed by capping of the surface $L\times[0,1] \subseteq W(Y,\mathbb{S}_1)$ with the J–Seifert surfaces S and S', and ψ is a homology triangle with vertices on $\mathbf{x}, \theta$ and $\mathbf{y}$ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, and with θ representing the torsion Spin^c structure. As in the proof of invariance of the gradings, the maps $F^{\tau}_{\mathbb{S}_1,S,S',\mathfrak{s}}$ only depend on $\mathfrak{s}$ through it's restriction to the ends, by the same argument as in Lemma 5.7.3. Using the same associativity argument as in Lemma 5.6.1 (and noting that the homological quantities involving the capped J–Seifert surface are additive) we see

that

$$F^{\alpha}_{\mathbb{S}_1,S,S',\mathfrak{s}}(A^{\alpha}_{Y,\mathbb{L},S}) = A^{\alpha}_{Y(\mathbb{S}_1),\mathbb{L},S'}$$

and similarly for the β–gradings and β–maps. Essentially the same argument we used to show invariance of the gradings $A^{\tau}_{Y,\mathbb{L},S}$ now shows that the maps $F^{\tau}_{\mathbb{S}_1,S,S',\mathfrak{s}}$ yield well defined maps on coherent gradings.

Given a framed 2–sphere $\mathbb{S}_2 \subseteq Y$ and a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(Y(\mathbb{S}_2))$, we can define a map on gradings $F_{\mathbb{S}_2}$, using the obvious formula, for diagrams which can be used to compute the 2–handle map. The same argument which shows that the 2–handle maps are well defined on the coherent chain complexes now shows that $F_{\mathbb{S}_2}$ yields a well defined map on coherent gradings. Furthermore, if S is a Seifert surface for $\mathbb{L}$ which doesn't intersect $\mathbb{S}_2$, then $F_{\mathbb{S}_2}(A^{\tau}_{Y,\mathbb{L},S}) = A^{\tau}_{Y(\mathbb{S}_2),\mathbb{L},S}$.

Note that if U is a 0–framed unknot contained in a ball in $Y \setminus L$, and $\mathbb{S}_2$ is the canonical framed 2–sphere in the surgered $Y_0(U)$ which cancels it, then the Seifert surface can be isotoped to miss $\mathbb{S}_2$. Furthermore, the same computation as the one in [OS06, Lemma 14.6], except done only for homology triangles, shows that

$$F_{\mathbb{S}_2} \circ F^{\tau}_{U,S,S,\mathfrak{s}} = \mathrm{id},$$

where $\mathfrak{s}$ is the unique Spin^c structure which extends to the entire cobordism $W(Y,\mathbb{S}_1)(\mathbb{S}_2) \cong Y \times [0,1]$, and which restricts to $\mathfrak{t}$ on Y. This holds for both $\tau = \alpha$ and $\tau = \beta$. Using this, the argument now proceeds by straightforward modification of [OS06, Lemma 5.2]. □

5.7.4 Dependence on the Seifert surface S

We now prove part (b) of Theorem 5.1.3:

Theorem 5.1.3(b). *If S and S' are two J–Seifert surfaces for a colored link $\mathbb{L}$ in Y, then*

$$A_{Y,\mathbb{L},S'}(\mathbf{x})_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = \frac{\langle c_1(\mathfrak{s}), [S'_j \cup -S_j]\rangle}{2}.$$

In particular, if $c_1(\mathfrak{s})$ is torsion, then the absolute Alexander grading does not depend on the choice of J–Seifert surface.

Proof. Let S and S' be two choices of J–Seifert surfaces. Pick a parametrized Kirby diagram of $(Y,\mathbb{L})$ with framed 1–dimensional link $\mathbb{S}_1$, and pick a Heegaard triple subordinate to a β–bouquet of $\mathbb{S}_1$. Let $\widehat{\Sigma}$ denote the surface in the cobordism obtained by capping with S and let $\widehat{\Sigma}'$ be the surface obtained from capping with S'. Let $[F]=[S'\cup -S]$. As elements of $H_2(W(S^3,\mathbb{S}_1);\mathbb{Z})$, we have $[\widehat{\Sigma}']=[\widehat{\Sigma}]+[F]$. We similarly define $[F_j]$, $[\widehat{\Sigma}'_j]$ and $[\widehat{\Sigma}_j]$, and the analogous relation holds. Using the formula for the absolute grading, we have

$$\begin{aligned}A_{Y,\mathbb{L},S'}(\mathbf{x})_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j &= \frac{\langle c_1(\mathfrak{s}),[\widehat{\Sigma}'_j]\rangle - \langle c_1(\mathfrak{s}),[\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}']\cdot[\widehat{\Sigma}'_j]+[\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]}{2}\\ &= \frac{\langle c_1(\mathfrak{s}),[F_j]\rangle}{2}+\frac{-[F]\cdot[\widehat{\Sigma}_j]-[\widehat{\Sigma}]\cdot[F_j]-[F]\cdot[F_j]}{2}\\ &= \frac{\langle c_1(\mathfrak{s}),[F_j]\rangle}{2},\end{aligned}$$

since $[F]$ and $[F_j]$ are in the image of the inclusion map $H_2(\partial W(S^3,\mathbb{S}_1);\mathbb{Z}) \to H_2(W(S^3,\mathbb{S}_1);\mathbb{Z})$, and hence the intersection number of either with anything in $H_2(W(S^3,\mathbb{S}_1);\mathbb{Z})$ vanishes. □

If $\mathfrak{s}$ is non-torsion, then in general the absolute Alexander grading depends on the choice of Seifert surface. In general, we can pick a diagram for a doubly based unlink $\mathbb{U}=(U,w,z)$ in Y by taking a diagram for (Y,w) and simply placing the z basepoint in the same region as w. The absolute Alexander grading on $\mathcal{C}FL^\circ(Y,\mathbb{U},\mathfrak{s})$ of $U_w^iV_z^j\cdot\mathbf{x}$ is given by just $j-i$, for the Seifert surface obtained by taking a path in surface between w and z which doesn't cross any $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves, and then pushing it off of Σ into both handlebodies U_α and U_β to get a disk with boundary U. The resulting Seifert surface of course depended on our choice of diagram and path in the surface, so in general one must be somewhat careful doing this. We illustrate this with an instructive example:

Example 5.7.12. Consider $Y=S^1\times S^2$ with a doubly based unknot $\mathbb{U}=(U,w,z)$. In Figure 5.14, two diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ are shown. In each $\mathcal{H}_i$ we have shown a path λ_i from w to z. We can specify two Seifert disks D_1 and D_2, such that $D_i\cap\Sigma=\lambda_i$. Consider a Spin^c structure $\mathfrak{s}$ represented by the intersection points in the diagram $\mathcal{H}_1$. Note that

$$\langle c_1(\mathfrak{s}),H(P_i)\rangle = 2,$$

where P_1 and P_2 are the periodic domains on $\mathcal{H}_1$ and $\mathcal{H}_2$ (resp.) with multiplicity -1 in the bigon containing the two basepoints, and multiplicity $+1$ in the empty bigon. Note that $H(P_i) \in H_2(Y;\mathbb{Z})$ can be represented by isotopic spheres.

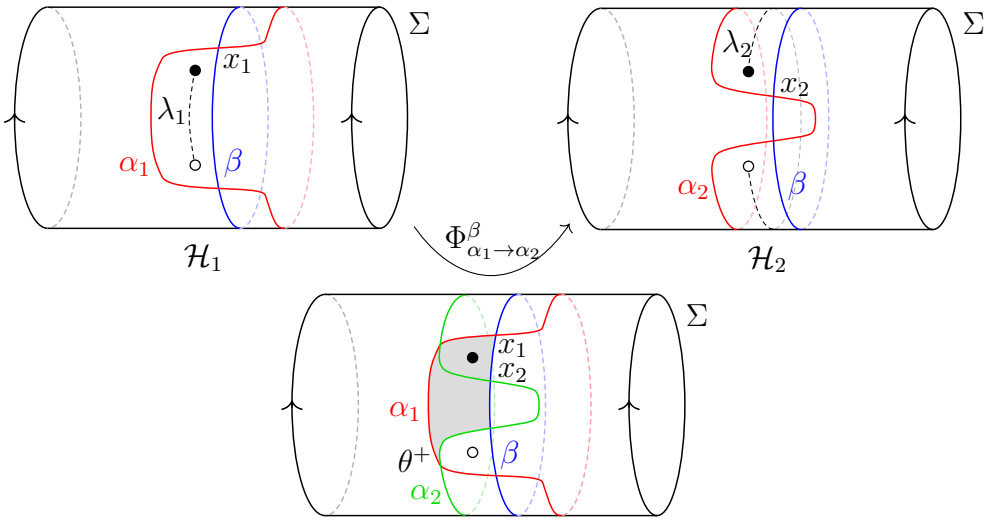

Figure 5.14: **The two diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ for $(S^1 \times S^2, U, w, z)$ in a non-torsion grading, and the paths λ_1 and λ_2.** On the bottom is a Heegaard triple which can compute $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}$. The disks D_1 and D_2 intersect Σ along λ_1 and λ_2. Shaded is a the domain of a homology triangle which can be used to compute $\Phi_{\mathcal{H}_1\to\mathcal{H}_2}(A_{D_1})$.

Let D_1 denote the disk constructed from λ_1, and let D_2 denote the disk constructed using λ_2. Using our orientation conventions, it is not hard to see that

$$D_2 \cup (-D_1) = -H(P_i).$$

To compute the signs, it's easiest to use Equation (5.7) to compute the intersection number of $H(P_i)$ with a knot of the form $S^1 \times \{pt\}$. By drawing a picture of the Seifert disks, and using our orientation conventions, it's easy to see that $D_2 \cup (-D_1)$ intersects the same knot with the opposite orientation. Hence

$$\langle c_1(\mathfrak{s}), D_2 \cup (-D_1)\rangle = -2.$$

On both diagrams $\mathcal{H}_i$, the absolute Alexander grading with Seifert surface D_i is given by $A_{D_i}(U^iV^j \cdot \mathbf{x}) = j - i$ for any intersection point x. We can move from $\mathcal{H}_1$ to $\mathcal{H}_2$ by a single

$\boldsymbol{\alpha}$ move, shown in Figure 5.14. We now explicitly compute that, as gradings on $\mathcal{H}_2$

$$\Phi_{\mathcal{H}_1\to\mathcal{H}_2}(A_{D_1})(x_2) = A_{D_1}(x_1) + (n_w - n_z)(\psi) = A_{D_1}(x_1) + 1.$$

Since $A_{D_i}(x_i) = 0$ for both $i = 1$ and $i = 2$, in terms of coherent gradings this yields

$$A_{D_2} = A_{D_1} - 1 = A_{D_1} + \frac{\langle c_1(\mathfrak{s}), D_2 \cup (-D_1)\rangle}{2},$$

agreeing with the previous theorem.

5.8 Absolute Maslov gradings

In this section, we prove parts (c) and (d) of Theorem 5.1.3, and describe the two absolute Maslov gradings, $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$, on $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$. The grading $\mathrm{gr}_{\mathbf{w}}$ is defined when $\mathfrak{s}$ is torsion, and $\mathrm{gr}_{\mathbf{z}}$ is defined when $\mathfrak{s} - PD[L]$ is torsion.

The construction of the absolute Maslov gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ mimics the construction from [OS06]. To simplify the construction from [OS06], we first need to state a formula due to Sarkar for the change of the Maslov index due to adding a triple periodic domain to a homology triangle.

Lemma 5.8.1. *Suppose that $\mathcal{P}$ is triply periodic domain and ψ is a homology triangle, on the Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$. Then*

$$\mu(\psi + \mathcal{P}) - \mu(\psi) = 2n_{\mathbf{w}}(\mathcal{P}) + \frac{c_1(\mathfrak{s}_{\mathbf{w}}(\psi + \mathcal{P}))^2 - c_1(\mathfrak{s}_{\mathbf{w}}(\psi))^2}{4}.$$

The proof can be found in [Sar11b, Section 5.1].

We now describe the absolute Maslov gradings. The $\mathrm{gr}_{\mathbf{w}}$–grading will be defined on intersection points $\mathbf{y}$ with $\mathfrak{s}_{\mathbf{w}}(\mathbf{y})$ torsion, and the $\mathrm{gr}_{\mathbf{z}}$–grading will be defined on intersection points with $\mathfrak{s}_{\mathbf{z}}(\mathbf{y})$ torsion. For a link where the sum of the link components is not rationally null-homologous, we may be able to define one of these, but not the other (see Example 5.5.5).

We pick a parametrized Kirby diagram $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ of $(Y, \mathbb{L})$, and a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ which is subordinate to a β–bouquet for $\mathbb{S}_1$. By definition $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ represents a collection of unlinks in $(S^1 \times S^2)^{\#k}$ for some k. In Section 5.5.4, we defined two absolute gradings, $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ on $\widehat{\mathcal{CFL}}((S^1 \times S^2)^{\#k}, \mathbb{U}, \mathfrak{s}_0)$, for an unlink $\mathbb{U}$ with various configurations of basepoints. If $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ and $\psi \in \pi_2(\mathbf{x}, \theta, \mathbf{y})$ is a homology triangle, we define

$$\mathrm{gr}_{\mathbf{w}}(\mathbf{y}) = \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) + \mathrm{gr}_{\mathbf{w}}(\theta) - \tfrac{1}{2}(k + |\mathbf{w}| - 1) - \mu(\psi) + 2n_{\mathbf{w}}(\psi) + \frac{c_1(\mathfrak{s}_{\mathbf{w}}(\psi))^2 - 2\chi(W) - 3\sigma(W)}{4},$$

where $W = W(S^3, \mathbb{S}_1)$. Similarly we define

$$\mathrm{gr}_{\mathbf{z}}(\mathbf{y}) = \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) + \mathrm{gr}_{\mathbf{z}}(\theta) - \tfrac{1}{2}(k + |\mathbf{w}| - 1) - \mu(\psi) + 2n_{\mathbf{z}}(\psi) + \frac{c_1(\mathfrak{s}_{\mathbf{z}}(\psi))^2 - 2\chi(W) - 3\sigma(W)}{4}.$$

Here θ is an intersection point which represents the torsion Spin^c structure on $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma})$.

We can now prove part (c) of Theorem 5.1.3.

Theorem 5.1.3(c). *The absolute gradings* $\mathrm{gr}_{\mathbf{w}}$ *and* $\mathrm{gr}_{\mathbf{z}}$*, listed above, are independent of the intersection points* $\mathbf{y}$ *and* $\mathbf{x}$*, the homology triangle* ψ*, as well as the parametrized Kirby diagram* $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$*.*

Proof. Most of the proof proceeds as in the proof of Theorem 5.1.3(a). The main difference is in the proof that the formulas are independent of the homology triangle ψ and intersection points. Independence of the absolute grading from $\mathbf{x}$ and θ can be proven by splicing in disks on those ends and seeing that the formula doesn't change. To see more generally that the formula is invariant under the choice of ψ, we note that any two homology triangles with the same endpoints differ by a triply periodic domain. Suppose that ψ and $\psi + \mathcal{P}$ are two such triangles differing by a triply periodic domain $\mathcal{P}$. Letting $\mathrm{gr}^{\psi}_{\mathbf{w}}(\mathbf{y})$ and $\mathrm{gr}^{\psi+\mathcal{P}}_{\mathbf{w}}(\mathbf{y})$ denote the formulas, computed with ψ or $\psi + \mathcal{P}$, respectively, we observe that the difference is

$$\mathrm{gr}^{\psi+\mathcal{P}}_{\mathbf{w}}(\mathbf{y}) - \mathrm{gr}^{\psi}_{\mathbf{w}}(\mathbf{y}) = -\mu(\psi + \mathcal{P}) + \mu(\psi) + 2n_{\mathbf{w}}(\mathcal{P}) + \frac{c_1(\mathfrak{s}_{\mathbf{w}}(\psi + \mathcal{P}))^2 - c_1(\mathfrak{s}_{\mathbf{w}}(\psi))^2}{4},$$

which is zero by Sarkar's formula from Lemma 5.8.1. Analogously, using the computations from the proof of invariance of the Alexander gradings in the previous subsection, invariance

from the moves of Lemma 5.4.4 follows similarly. The same argument works for showing invariance of $\mathrm{gr}_{\mathbf{z}}$. □

The following is part (d) of Theorem 5.1.3:

Theorem 5.1.3(d). *The absolute Maslov and collapsed Alexander grading satisfy*

$$A = \tfrac{1}{2}(\mathrm{gr}_{\mathbf{w}} - \mathrm{gr}_{\mathbf{z}}).$$

Proof. We pick a parametrized Kirby diagram and associated surgery triple for $(Y, \mathbb{L})$. We compute $\frac{1}{2}$ times the difference between the expressions defining $\mathrm{gr}_{\mathbf{w}}(\mathbf{x})$ and $\mathrm{gr}_{\mathbf{z}}(\mathbf{x})$. We note that by Equation (5.11), the formula $A = \frac{1}{2}(\mathrm{gr}_{\mathbf{w}} - \mathrm{gr}_{\mathbf{z}})$ holds for the gradings associated to unlinks in $(S^1 \times S^2)^{\#k}$. By Lemma 5.3.3, we have that $\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi) = PD[\Sigma_{\alpha\beta\gamma}]$. Under the inclusion $X_{\alpha\beta\gamma} \hookrightarrow W(Y, \mathbb{S}_1)$, we have that $PD[\Sigma_{\alpha\beta\gamma}] = PD[\Sigma]$. Noting that

$$\frac{c_1(\mathfrak{s}_{\mathbf{w}}(\psi))^2 - c_1(\mathfrak{s}_{\mathbf{z}}(\psi))^2}{4} = c_1(\mathfrak{s}_{\mathbf{w}}(\psi)) \cup PD[\Sigma] - PD[\Sigma] \cup PD[\Sigma] = \langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi)), [\widehat{\Sigma}] \rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}],$$

we see that the expression defining $\frac{1}{2}(\mathrm{gr}_{\mathbf{w}}(\mathbf{x}) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}))$ becomes exactly the expression defining $A(\mathbf{x})$. □

Remark 5.8.2. Note that if L is null-homologous and $\mathfrak{s}$ is torsion, if we tensor $\mathcal{C}FL^{\circ}(Y, \mathbb{L}, \mathfrak{s})$ with $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]/(V_{\mathbf{z}} - 1)$ or $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]/(U_{\mathbf{w}} - 1)$, the gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ both give the standard absolute gradings on $\widehat{CF}(Y, \mathbf{w}, \mathfrak{s})$ or $\widehat{CF}(Y, \mathbf{z}, \mathfrak{s})$ (respectively), since a parametrized Kirby diagram of a link naturally gives a parametrized Kirby diagram of the 3–manifold, by forgetting about half of the basepoints. An alternate way to compute the collapsed collapsed Alexander grading is then to just compute the two absolute gradings on $\widehat{CF}(Y, \mathbf{w}, \mathfrak{s})$ or $\widehat{CF}(Y, \mathbf{z}, \mathfrak{s})$, and then define $A = \frac{1}{2}(\mathrm{gr}_{\mathbf{w}} - \mathrm{gr}_{\mathbf{z}})$, on intersection points.

Remark 5.8.3. The absolute gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ we've written here are normalized so that $HF^-(S^3)$ has a top degree generator in grading 0. In [OS06] the gradings are normalized so that $HF^-(S^3)$ has top degree generator in grading -2. Similarly the gradings in [HMZ16] are normalized so that HF^- of disjoint unions of S^3 have top degree generator in grading -2. If one wants to go from one convention to the other, one simply subtracts 2 from our

Maslov gradings to get ones consistent with the other theories. Hopefully this won't cause confusion.

5.9 Link cobordisms and grading change formulas

In this section we discuss the grading change for link cobordisms. We first consider the maps associated to 4–dimensional handles added away from the link, then we consider the maps associated to the surface (band maps, quasi-stabilization, and disk-stabilization) separately, and then prove that the formulas are additive to obtain the stated grading change formulas in general.

5.9.1 Grading changes of maps associated to elementary link cobordisms

In this subsection, we compute the Alexander grading change for the constituent maps appearing in the link cobordism maps. We first consider the disk-stabilization maps:

Lemma 5.9.1. *The disk stabilization maps $S^{\pm}_{\mathbb{U},D}$ satisfy $A_{Y,\mathbb{L}\cup\mathbb{U},S\cup D}(S^{\pm}_{\mathbb{U},D}(\mathbf{x}))_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = 0$. Similarly the maps $S^{\pm}_{\mathbb{U},D}$ both induce $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ grading changes of $+\frac{1}{2}$.*

Proof. First consider the Alexander grading change. We take a parametrized Kirby diagram and a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ representing surgery on the associated link $\mathbb{S}_1$ in $S^3 \setminus U$, which results in $(Y, \mathbb{L})$, with the additional assumption that $w, z \in \Sigma$ and $D \cap \Sigma$ in an arc from w to z which is disjoint from $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w} \cup \mathbf{z})$.. We then form the triple $\overline{\mathcal{T}}$ by disk-stabilizing $\mathcal{T}$ by adding in isotopic α_0, β_0 and γ_0 curves which each bound a ball containing $D \cap \Sigma$. The triple $\overline{\mathcal{T}}$ represents surgery on the same framed link $\mathbb{S}_1$ in S^3, but now the link $\mathbb{L}$ has an extra unknot component U. By picking $\mathbb{S}_1$ so that it doesn't intersect D, we view D as being embedded in S^3, $(S^1 \times S^2)^{\#k}$ and Y. Pick a triangle $\psi_{\alpha\beta\gamma} \in \pi_2(\mathbf{x}_{\alpha\beta}, \theta_{\beta\gamma}, \mathbf{x}_{\alpha\gamma})$ and a stabilization $\overline{\psi} \in \pi_2(\mathbf{x}_{\alpha\beta} \times \theta^+, \theta_{\beta\gamma} \times \theta^+, \mathbf{x}_{\alpha\gamma} \times \theta^+)$. The surface $\widehat{\overline{\Sigma}}_j$ appearing in the link cobordism associated to $\mathcal{T}$ is obtained from $\widehat{\Sigma}_j$ by adding a null-homologous sphere, and hence the homological quantities in the grading formulas are

unchanged. An easy computation using triangles ψ and $\overline{\psi}$ as above shows that

$$A_{\overline{\mathcal{T}}}(\mathbf{x}_{\alpha\gamma} \times \theta^+)_j - A_{\mathcal{T}}(\mathbf{x}_{\alpha\gamma})_j = A(\mathbf{x}_{\alpha\beta} \times \theta^+)_j - A(\mathbf{x}_{\alpha\beta})_j + A(\theta_{\beta\gamma} \times \theta^+)_j - A(\theta_{\beta\gamma})_j,$$

and hence it is sufficient to show that the maps $S^{\pm}_{\mathbb{U},D}$ have grading change zero for unlinks in S^3, which is obvious from the definition of the Alexander grading on unlinks in $(S^1 \times S^2)^{\#k}$.

We now consider the Maslov grading changes. Adapting the above argument, by stabilizing the homology triangle by adding in a small triangle with Maslov index 0 which crosses over none of the basepoints, we need only consider the grading change associated to performing the disk stabilization maps to an unlink in S^3. It follows immediately from the formula that $S^+_{\mathbb{U},D}$ and $S^-_{\mathbb{U},D}$ are both $+\frac{1}{2}$ graded with respect to both $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. □

We now prove the analogous result for the quasi-stabilization maps. We recall that the quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ were defined in Chapter 4 and were associated to the decorated link cobordisms shown in Figure 5.15.

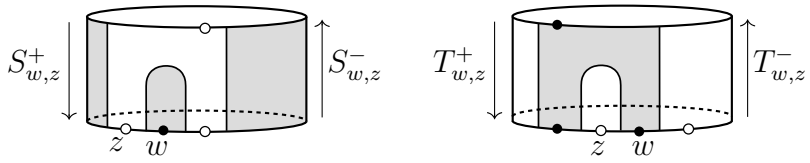

Figure 5.15: **The decorated link cobordisms associated to the quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$.** Shown are the decorated surfaces $L \times [0,1]$, which sit inside of $Y \times [0,1]$.

Lemma 5.9.2. *The quasi-stabilization maps $S^{\pm}_{w,z}$ induce Alexander grading change $+\frac{1}{2}\delta(\pi(K), j)$ in the j component of the Alexander multi-grading, where K is the component of L containing w and z. Similarly the maps $T^{\pm}_{w,z}$ induce an Alexander grading change of $-\frac{1}{2}\delta(\pi(K), j)$.*

The maps $S^{\pm}_{w,z}$ also induce a $+\frac{1}{2}$ grading change of $\mathrm{gr}_{\mathbf{w}}$ and a $-\frac{1}{2}$ grading change of $\mathrm{gr}_{\mathbf{z}}$. Analogously $T^{\pm}_{w,z}$ induces a $-\frac{1}{2}$ grading change of $\mathrm{gr}_{\mathbf{w}}$ and a $+\frac{1}{2}$ grading change of $\mathrm{gr}_{\mathbf{z}}$.

Proof. This follows analogously to the proof of Lemma 5.9.1. Consider first the maps $S^{\pm}_{w,z}$. We take a surgery triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ and its quasi-stabilization along an α_s curve (we

add w and z, the α_s curve is allowed to travel around the diagram, but we add two curves β_0 and γ_0 which both bound disks in Σ containing w and z and no other basepoints or curves). We note that the diagram $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma})$ becomes disk-stabilized, not quasi-stabilized. If $\mathbf{x}_{\alpha\beta}, \mathbf{x}_{\alpha\gamma}$ and $\theta_{\beta\gamma}$ are intersection points, as in the previous lemma, we have that

$$A_{\overline{\mathcal{T}}}(\mathbf{x}_{\alpha\gamma} \times \theta^+_{\mathbf{w}})_j - A_{\mathcal{T}}(\mathbf{x}_{\alpha\gamma})_j = A(\mathbf{x}_{\alpha\beta} \times \theta^+_{\mathbf{w}})_j - A(\mathbf{x}_{\alpha\beta})_j + A(\theta_{\beta\gamma} \times \theta^+)_j - A(\theta_{\beta\gamma})_j.$$

Now $A(\mathbf{x}_{\alpha\beta} \times \theta^+_{\mathbf{w}})_j - A(\mathbf{x}_{\alpha\beta})_j$ is equal to the grading induced by the quasi-stabilization map $S^+_{w,z}$ on unlinks in S^3, which is easily seen to be $+\frac{1}{2}$ by definition of the absolute grading on those links. Now $A(\theta_{\beta\gamma} \times \theta^+)_j - A(\theta_{\beta\gamma})_j$ is equal to the grading change of disk stabilization on unlinks in $(S^1 \times S^2)^{\#k}$, which is zero by the definition of the absolute grading on unlinks in $(S^1 \times S^2)^{\#k}$. The grading change for $T^{\pm}_{w,z}$, as well as the Maslov gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ are described analogously. □

If B is a band for a link L, with its ends on components of L of the same grading, with J–Seifert surface S, then we naturally get an J–Seifert surface $S \cup B$ for $L(B)$. Recall that a band B for L, intersects L only along the boundary of B. Note that we defined a J–Seifert surface to be an immersed surface, not necessarily an embedded surface, so it isn't important that B doesn't intersect S, since we only need a homology class defined by $S \cup B$ to compute the grading. In general, there is no hope of assuming that a band B occurs in the complement of the Seifert surfaces. Nonetheless, if B is a with ends on link components with grading assignment j and B intersects $S_j = \pi^{-1}(j) \subseteq S$ in an a collection of arcs and closed 1–manifolds, it is easy to surger $S_j \cup B$ to make it an embedded surface without changing the homology class. This of course is only allowable if the intersection of B is with S_j and B is also attached to link components with the grading assignment j.

Lemma 5.9.3. *Suppose $(Y, \mathbb{L})$ is a 3–manifold with multi-based link and J–Seifert surface S, and B is a band for $\mathbb{L}$, between link components with the same grading assignment, j_0, then*

$$A_{Y,\mathbb{L}(B),S\cup B}(F^{\mathbf{z}}_B(\mathbf{x}))_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = +\tfrac{1}{2}\delta(j, j_0),$$

and

$$A_{Y,\mathbb{L}(B),S\cup B}(F_B^{\mathbf{w}}(\mathbf{x}))_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = -\tfrac{1}{2}\delta(j,j_0).$$

Proof. We pick a parametrized Kirby diagram $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ which is nicely adapted to the band, as we now describe. Let $\mathbb{U} \subseteq S^3$ be a multi-based unlink and let $\phi_0 : \mathbb{U} \to \mathbb{L}$ be a diffeomorphism of multi-based, oriented 1–manifolds. We pick a band B_0 in S^3 between the two components of $\mathbb{U}$ which correspond under ϕ_0 to the components connected by B, and that the configuration of basepoints with respect to the bands is the same for $\mathbb{U}$ and B_0 and $\mathbb{L}$ and B. We can assume that B_0 is a "standard" band for $\mathbb{U}$, i.e. we can assume that all components of $\mathbb{U}$ can be put on an embedded sphere, and the band B_0 also lies on that sphere.

We can add sutures to $\partial N(\mathbb{U} \cup B_0)$ and $\partial N(\mathbb{L} \cup B)$ corresponding to the basepoints. A choice of framing of the link $\mathbb{L}$ determines a diffeomorphism ϕ_λ of $\partial N(\mathbb{U}\cup B_0)$ to $\partial N(\mathbb{L}\cup B)$, which is compatible with ϕ_0 in the obvious sense. As we described after Definition 5.4.1, we can find parametrized surgery data for the tuple $(S^3 \setminus N(\mathbb{U} \cup B_0), Y \setminus (\mathbb{L} \cup B), \phi_\lambda)$. This consists of a framed 1–dimensional link $\mathbb{S}_1 \subseteq S^3 \setminus N(\mathbb{U} \cup B_0)$ and a diffeomorphism $f : (S^3 \setminus N(\mathbb{U} \cup B_0)) \to (Y \setminus N(\mathbb{L} \cup B))$ which extends ϕ_λ. This naturally induces a parametrized Kirby diagram $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ for $\mathbb{L}$, and a parametrized Kirby diagram $\mathbb{P}(B) = (\phi_0(B), \lambda(B), \mathbb{S}_1, f)$ for $(Y, \mathbb{L}(B))$, where $\phi_0(B)$ and $\lambda(B)$ are the diffeomorphisms and framings induced by ϕ_0, λ and B.

Assume now, for the sake of demonstration that B is a α–band for $\mathbb{L}$. We now pick a Heegaard quadruple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ such that

- $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ is a triple subordinate to a β–bouquet of the framed link $\mathbb{S}_1 \subseteq S^3 \setminus (\mathbb{U}\cup B)$;
- $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a triple subordinate to the band B_0 for $\mathbb{U}$;
- (the pushforward under f of) $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\gamma})$ is a triple subordinate to the band B for $\mathbb{L}$;
- $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}, \boldsymbol{\gamma})$ is a triple subordinate to a β–bouquet of the link $\mathbb{S}_1 \subseteq S^3 \setminus (\mathbb{U} \cup B)$.

The band B_0 connects two different components of $\mathbb{U}$, and $\mathbb{U}(B_0)$ is an unlink as well.

Let $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$ be an intersection point. We wish to compute the grading of $F_{B^{\mathbf{o}}}(\mathbf{x})$. Pick triangles

$$\psi_{\alpha'\alpha\beta} \in \pi_2(\theta_{\alpha'\alpha}, \mathbf{x}_{\alpha\beta}, \mathbf{x}_{\alpha'\beta}), \qquad \psi_{\alpha'\beta\gamma} \in \pi_2(\mathbf{x}_{\alpha'\beta}, \theta_{\beta\gamma}, \mathbf{x}_{\alpha'\gamma}), \qquad \psi_{\alpha\beta\gamma} \in \pi_2(\mathbf{x}_{\alpha\beta}, \theta_{\beta\gamma}, \mathbf{x}),$$

$$\text{and} \qquad \psi_{\alpha'\alpha\gamma} \in \pi_2(\theta_{\alpha'\alpha}, \mathbf{x}_{\alpha\gamma}, \mathbf{x}_{\alpha'\gamma})$$

with

$$\psi_{\alpha'\alpha\gamma} + \psi_{\alpha\beta\gamma} = \psi_{\alpha'\beta\gamma} + \psi_{\alpha'\alpha\beta}.$$

By splicing in disks to the ends of the triangles if necessary, we can assume that $\theta_{\alpha'\alpha}$ is in the same Alexander grading as $\Theta^{\mathbf{o}}$ and also that $\theta_{\beta\gamma}$ is in the same Alexander grading as the generator of top degree of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma})$. In fact, since these triangles represent the same Spin^c structures as those counted by $F^{\mathbf{o}}_B$ and $F^{\mathbf{o}}_{B_0}$, by splicing in disks on the ends, we can obtain any homology triangles which could be counted by the band maps $F^{\mathbf{o}}_B$ or $B^{\mathbf{o}}_{B_0}$.

Let $\widehat{\Sigma}$ denote the surface in $W(S^3, \mathbb{S}_1)$ obtained by capping off $\mathbb{U} \times [0,1]$ with$f^{-1}(S)$ in $S^3(\mathbb{S}_1)$ and a J–Seifert surface in S^3. Let $\widehat{\Sigma}'$ denote the surface in $W(S^3, \mathbb{S}_1)$ obtained by capping off $\mathbb{U}(B) \times [0,1]$ with $f^{-1}(S \cup B)$ and a J–Seifert surface in S^3. Note that $\widehat{\Sigma}$ and $\widehat{\Sigma}'$ are homologous surfaces. We similarly define $\widehat{\Sigma}_j$ and $\widehat{\Sigma}'_j$ for each $j \in J$ and note that each $\widehat{\Sigma}_j$ is homologous to $\widehat{\Sigma}'_j$.

By definition of the absolute grading, we have that

$$A_{Y,\mathbb{L}(B),S\cup B}(\mathbf{x}_{\alpha'\gamma})_j = A(\mathbf{x}_{\alpha'\beta})_j + A(\theta_{\beta\gamma})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha'\beta\gamma})_j + \frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_{\alpha'\beta\gamma})), [\widehat{\Sigma}'_j]\rangle - [\widehat{\Sigma}'] \cdot [\widehat{\Sigma}'_j]}{2},$$

and

$$A_{Y,\mathbb{L},S}(\mathbf{x})_j = A(\mathbf{x}_{\alpha\beta})_j + A(\theta_{\beta\gamma})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha\beta\gamma})_j + \frac{\langle c_1(\mathfrak{s}_{\mathbf{w}}(\psi_{\alpha\beta\gamma})), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]}{2}.$$

As we've remarked, the homological terms appearing at the end are equal. Hence we have just that

$$A_{Y,\mathbb{L}(B),S\cup B}(\mathbf{x}_{\alpha'\gamma})_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = A(\mathbf{x}_{\alpha'\beta})_j - A(\mathbf{x}_{\alpha\beta})_j + (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha'\alpha\gamma} - \psi_{\alpha'\alpha\beta}),$$

which we can rewrite as

$$A_{Y,\mathbb{L}(B),S\cup B}(F_B^{\mathbf{o}}(\mathbf{x}))_j - A_{Y,\mathbb{L},S}(\mathbf{x})_j = A(\mathbf{x}_{\alpha'\beta})_j - A(\mathbf{x}_{\alpha\beta})_j - (n_{\mathbf{w}} - n_{\mathbf{z}})(\psi_{\alpha'\alpha\beta})_j.$$

The above equation can be interpreted as saying that the grading change for the band maps $F_B^{\mathbf{o}}$ for $B \subseteq Y$ are the same as the grading change for the model band maps for the band B_0, the band in S^3 between two components of an unlink, which is standard in the way we've described above.

This can now be performed as a model computation. As usual, we can pick any diagram to show this. The diagram we pick is obtained as follows. Suppose for moment that B_0 joins two components of $\mathbb{U}$. Pick a diagram for each component of $\mathbb{U}$ in S^3, then form a diagram for $\mathbb{U}$ by attaching the diagrams together using 1–handles. If the ends of B_0 are adjacent to the basepoints w_1, w_2, z_1 and z_2, then we can assume that we attach the 1–handles near z_1 and z_2. A triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ for the band B_0 is shown on the top of Figure 5.16, along with intersection points of $\Theta^{\mathbf{w}}$ and $\Theta^{\mathbf{z}}$. If B_0 splits a single component of $\mathbb{U}$, then we modify this procedure to get the Heegaard triple shown on the bottom row of Figure 5.16.

Outside of the region shown in Figure 5.16, curves $\boldsymbol{\alpha}'$ are just small Hamiltonian isotopies of the curves $\boldsymbol{\alpha}$. In the top left of Figure 5.16, we show a portion of a triangle in $\pi_2(\Theta^{\mathbf{w}}, \Theta^{\mathbf{w}}, \Theta^{\mathbf{w}})$, which would be counted by $F_{B_0}^{\mathbf{z}}$. Outside of the region shown, we simply add in the small triangles (as the $\boldsymbol{\alpha}'$–curves are small isotopies of the $\boldsymbol{\alpha}$–curves). As this triangle class goes over no basepoints, we see that $F_{B_0}^{\mathbf{z}}$ maps the Alexander grading containing $\Theta^{\mathbf{w}}$ in $\mathcal{C}FL^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ to the Alexander grading containing $\Theta^{\mathbf{w}}$ in $\mathcal{C}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta})$. By our declaration of the absolute grading of unlinks in S^3, this induces a grading change of $+\frac{1}{2}$ for the grading j_0, and 0 in the other gradings.

A similar consideration of the other three triples shown in Figure 5.16 yields the rest of the grading changes. □

Analogous to the previous lemma, we have the following:

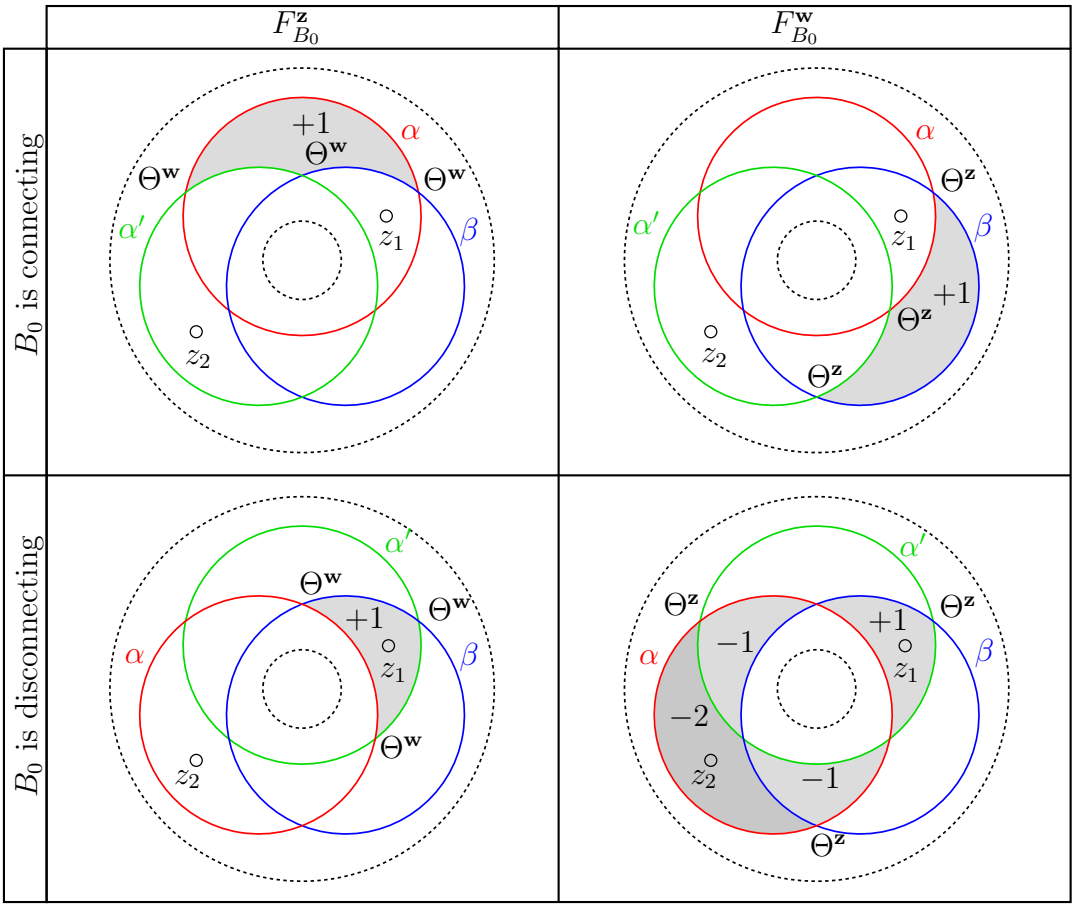

Figure 5.16: **The triples used to compute the band maps $F^{\mathbf{o}}_{B_0}$.** Also shaded are portions of homology triangles. Outside of the 1–handle regions shown, the rest of the homology triangle consists of small triangles. The dashed circles are where the connected sum is taking place.

Lemma 5.9.4. *Suppose $(Y, \mathbb{L})$ is a 3-manifold with multi-based link and J–Seifert surface S, and B is a band for $\mathbb{L}$, between link components with the same grading assignment, j_0, and suppose that $\mathfrak{s}$ is torsion. Then*

$$\mathrm{gr}_{\mathbf{w}}(F^{\mathbf{z}}_B(\mathbf{x})) - \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) = 0, \qquad \textit{and} \qquad \mathrm{gr}_{\mathbf{w}}(F^{\mathbf{w}}_B(\mathbf{x})) - \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) = -1.$$

Similarly, if $\mathfrak{s} - PD[L]$ is torsion, then

$$\mathrm{gr}_{\mathbf{z}}(F^{\mathbf{z}}_B(\mathbf{x})) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) = -1, \qquad \textit{and} \qquad \mathrm{gr}_{\mathbf{z}}(F^{\mathbf{w}}_B(\mathbf{x})) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) = 0.$$

Proof. Let us first consider $\mathrm{gr}_{\mathbf{w}}$. The same argument as in the previous lemma reduces the

model computation to computing the band map in S^3 between two unknots. The same model computation in Figure 5.16, together with our declaration of the absolute Maslov gradings for unlinks in S^3 immediately yields the stated grading change formulas. □

Lemma 5.9.5. *The 0–handle and 4–handle maps $F_{0,\mathbb{U}}$ and $F_{4,\mathbb{U}}$, which add or remove a copy of (S^3, U, w, z), are both 0–graded, for all gradings.*

Proof. One simply takes a parametrized Kirby diagram, and adds a copy of $(S^3, \mathbb{U})$ together with the empty framed link in S^3, to the parametrized Kirby diagram. The formula for the grading is clearly unchanged. □

To package the grading changes induced by quasi-stabilization maps, band maps, and disk-stabilization maps, we define several homological quantities associated to surfaces with divides.

Recall that if $(\Sigma, \mathcal{A})$ is a surface with divides, then $\Sigma \backslash \mathcal{A}$ is partitioned into two subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$. If $\pi : C(\Sigma) \to J$ is a grading assignment, let $\Sigma_{\mathbf{w},j}$ and $\Sigma_{\mathbf{z},j}$ denote the components of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ assigned the grading j. Given a link cobordism (W, F) from $(Y_1, \mathbb{L}_1)$ to $(Y_2, \mathbb{L}_2)$ with $\mathbb{L}_i = (L_i, \mathbf{w}_i, \mathbf{z}_i)$, we define the **reduced Euler characteristics**

$$\widetilde{\chi}(\Sigma_{\mathbf{w},j}) = \chi(\Sigma_{\mathbf{w},j}) - \tfrac{1}{2}(|\mathbf{w}_1| + |\mathbf{w}_2|), \qquad \text{and} \qquad \widetilde{\chi}(\Sigma_{\mathbf{z},j}) = \chi(\Sigma_{\mathbf{z},j}) - \tfrac{1}{2}(|\mathbf{z}_1| + |\mathbf{z}_2|).$$

Note that

$$\widetilde{\chi}(\Sigma_{\mathbf{w},j}) - \widetilde{\chi}(\Sigma_{\mathbf{z},j}) = \chi(\Sigma_{\mathbf{w},j}) - \chi(\Sigma_{\mathbf{z},j}).$$

We prove the following (cf. [HMZ16, Lemma 4.3]):

Lemma 5.9.6. *The reduced Euler characteristics $\widetilde{\chi}$ defined above are additive under composition of link cobordisms.*

Proof. This follows from the inclusion–exclusion principle on the Euler characteristic of a union. □

We can now prove our Alexander grading formula for a general cobordism:

Theorem 5.1.5. *If $(W,F):(Y_1,\mathbb{L}_1)\to(Y_2,\mathbb{L}_2)$ is a J–graded link cobordism which is given a J–graded coloring, and S_1 and S_2 are two J–Seifert surfaces for $\mathbb{L}_1$ and $\mathbb{L}_2$, respectively, then we have the following:*

$$A_{Y_2,\mathbb{L}_2,S_2}(F^{\circ}_{W,F,\mathfrak{s}}(\mathbf{x}))_j - A_{Y_1,\mathbb{L}_1,S_1}(\mathbf{x})_j = \frac{\langle c_1(\mathfrak{s}),\widehat{\Sigma}_j\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]}{2} + \frac{\chi(\Sigma_{\mathbf{w},j})-\chi(\Sigma_{\mathbf{z},j})}{2}.$$

Proof. It is sufficient to verify that the formula is additive, and that it holds for the elementary cobordisms we've described. The homological terms involving $\widehat{\Sigma}_j$ are obviously additive. The terms involving the Euler characteristics are also easily seen to be additive, since each link component has the same number of $\mathbf{w}$ basepoints as $\mathbf{z}$ basepoints.

It remains to verify the formula for elementary cobordisms. Consider first a cobordism which is of the form $(Y\times[0,1],L\times[0,1])$ with a dividing set which has a single local max or local min along one of the arcs. The dividing sets are shown in Figure 5.15. The cobordism map is just a quasi-stabilization map, either $S^+_{w,z}, S^-_{w,z}, T^+_{w,z}$ or $T^-_{w,z}$, depending on the configuration of the arc with respect to the incoming or outgoing boundaries. For example if the arc has two endpoints on the outgoing boundary, and the small region bounded is of type–$\mathbf{w}$, then term $\frac{1}{2}(\chi(\Sigma_{\mathbf{w},j})-\chi(\Sigma_{\mathbf{z},j}))$ is equal to $+\frac{1}{2}$. Similarly the grading change of the quasi-stabilization map corresponding to this dividing set, $S^+_{w,z}$, is also $+\frac{1}{2}$, according to Lemma 5.9.2. The other three configurations and quasi-stabilization maps can be analyzed similarly.

If the cobordism is diffeomorphic to $(Y\times[0,2],(L\times[0,1])\cup B)\cup(Y\times[1,2],L(B)\times[1,2])$ (after smoothing the corners of the surface), with a dividing set of the form $\mathbf{p}\times[0,1]$, and B is a type $\mathbf{o}$ band (for $\mathbf{o}\in\{\mathbf{w},\mathbf{z}\}$), then $\frac{1}{2}(\chi(\Sigma_{\mathbf{w},j})-\chi(\Sigma_{\mathbf{z},j}))$ is equal to $+\frac{1}{2}$ if $\mathbf{o}=\mathbf{z}$, and is equal to $-\frac{1}{2}$ if $\mathbf{o}=\mathbf{w}$. This agrees with the grading of the band maps, as computed in Lemma 5.9.3.

Now suppose the cobordism is obtained by adding a 4–dimensional 1–handle away from the link, with both feet in the same component. We use the same Seifert surfaces on both sides. We stabilize the Heegaard triple by attaching a torus representing zero surgery on an unknot in a ball not intersecting the framed surgery link or the unlink U. An easy

computation shows that the cobordism map induces Alexander grading change 0 if we use the same Seifert surface on both ends, agreeing with our formula for the grading. For a 1–handle with its feet attached in different components, we can just use the Heegaard triple formed by attaching the two components of the Heegaard surface with a 1-handle, and putting isotopic α_0, β_0 and γ_0 curves. A simple model computation shows that again the Alexander grading change is zero.

Similarly, if S is a Seifert surface which doesn't intersect a framed 2–sphere, then the 3–handle cobordism map is zero graded, with respect to the gradings induced by putting the same Seifert surface on both ends.

Finally if the cobordism is obtained by attaching 2–handles, then a homology associativity argument analogous to Lemma 5.6.1 shows that the change of grading in the j–component of the grading is

$$\frac{\langle c_1(\mathfrak{s}), \widehat{\Sigma}_j\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]}{2}.$$

We finally consider 4–handle and 0–handle maps. By Lemma 5.9.5, the cobordism maps are 0–graded. On the other hand the terms involving the homology class $[\widehat{\Sigma}_j]$ vanish, and $\chi(\Sigma_{\mathbf{w},j}) = \chi(\Sigma_{\mathbf{z},j})$, so formula for the grading change reads zero as well.

Noting that (after removing some balls and applying the 0–handle or 4–handle maps) a parametrized Kirby decomposition in the sense of Chapter 4 gives us a decomposition in terms of the cobordisms described in the previous paragraph. The formula follows. □

We can now prove our formula for the grading change of the Maslov gradings.

Theorem 5.1.6. *If $(W,F):(Y_1,\mathbb{L}_1)\to(Y_2,\mathbb{L}_2)$ is a link cobordism with grading assignment and a J–graded coloring, and $\mathfrak{s}\in \mathrm{Spin}^c(W)$ is a* Spin^c *structure which is torsion on Y_1 and Y_2, then the link cobordism maps satisfy*

$$\mathrm{gr}_{\mathbf{w}}(F^{\circ}_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{w}}),$$

for homogeneous $\mathbf{x}$ with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. Similarly if $\mathfrak{s} - PD[\Sigma]$ restricts to a torsion Spin^c *structure on the ends, then for $\mathbf{x}$ which is homogeneous with respect to $\mathrm{gr}_{\mathbf{z}}$, we*

have

$$\mathrm{gr}_{\mathbf{z}}(F^{\circ}_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) = \frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{z}}).$$

Proof. Let us first consider the $\mathrm{gr}_{\mathbf{w}}$ grading. An easy computation shows that the decorated surfaces for associated to the quasi-stabilization maps $S^{\pm}_{w,z}$ satisfy $\widetilde{\chi}(\Sigma_{\mathbf{w}}) = \frac{1}{2}$. Analogously, the maps $S^{+}_{w,z}$ and $S^{-}_{w,z}$ are $+\frac{1}{2}$ graded with respect to $\mathrm{gr}_{\mathbf{w}}$, by Lemma 5.9.2. Similarly the dividing sets associated to the quasi-stabilization maps $T^{+}_{w,z}$ and $T^{-}_{w,z}$ have $\widetilde{\chi}(\Sigma_{\mathbf{w}}) = -\frac{1}{2}$, while the maps $T^{+}_{w,z}$ and $T^{-}_{w,z}$ induce $\mathrm{gr}_{\mathbf{w}}$–grading change equal to $-\frac{1}{2}$, again by Lemma 5.9.2.

The band maps induce the expected grading change by Lemma 5.9.4. That 1– and 3–handles produce the expected grading change by adding a handle with three isotopic curves to the Heegaard surface involved in a surgery triple, as in [OS06]. The 2–handle maps count triangles with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, and hence induce the correct $\mathrm{gr}_{\mathbf{w}}$ grading change, by the same argument as in [OS06]. The 0–handle and 4–handle maps are zero graded, while $\chi(W) = 1$ and $\widetilde{\chi}(\Sigma_{\mathbf{w}}) = \frac{1}{2}$ for a 0–handle or 4–handle cobordism, so the proposed formula also reads zero. The formula for the grading change is additive under composition by Lemma 5.9.6, and hence the grading change formula for $\mathrm{gr}_{\mathbf{w}}$ follows in general.

The same argument goes through for the $\mathrm{gr}_{\mathbf{z}}$ grading, except in the case of 2–handles. The map associated to two handles counts triangles with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, whereas the absolute grading $\mathrm{gr}_{\mathbf{z}}$ is defined using the Spin^c structure $\mathfrak{s}_{\mathbf{z}}(\psi)$. If ψ is counted by the 2–handle map, then $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$, while the 2–handle map induces $\mathrm{gr}_{\mathbf{z}}$ grading change

$$\frac{c_1(\mathfrak{s}_{\mathbf{z}}(\psi))^2 - 2\chi(W) - 3\sigma(W)}{4}.$$

Using Lemma 5.3.3, we see that $\mathfrak{s}_{\mathbf{z}}(\psi) = \mathfrak{s}_{\mathbf{w}}(\psi) - PD[\Sigma] = \mathfrak{s} - PD[\Sigma]$, and hence the $\mathrm{gr}_{\mathbf{z}}$ grading will be changed by

$$\frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4}.$$

Using additivity under composition of cobordisms, the proof is complete. □

5.10 Conjugation symmetry, equivalence with other constructions, and collapsing gradings

In this section we describe several basic properties of the Alexander gradings we've defined, such as conjugation symmetry and the effect of collapsing the index set J.

5.10.1 Conjugation symmetry

We now prove that our construction of the absolute Alexander multi-grading agrees with the construction from definition in [OS08] by showing that it satisfies the conjugation symmetry property. Undoubtedly one could show this directly, but it's useful to have the conjugation symmetry result we state in this paper, so we state the equivalence of our grading and theirs as a consequence of this. In [OS08], the Alexander grading for an oriented link L in S^3 (with two basepoints per component) is presented over an affine $\mathbb{Z}^\ell$ lattice denoted $\mathbb{H}(L)$. Using the notation and Alexander gradings from [OS08, pg. 616], one has

$$\widehat{HFL}_d(L,h) = \widehat{HFL}_{d-2\delta(h)}(L,-h). \tag{5.15}$$

Here d denotes the homological grading, and h denotes the Alexander grading. The element h is an element of an affine, integral lattice $\mathbb{H}(L)$ and if $h = \sum_{i=1}^{n} a_i \cdot PD[\mu_i]$, (where μ_i is the meridian of the ith link component and $a_i \in \mathbb{Q}$) then $\delta(h) = a_1 + \cdots + a_n$ (see [OS08] for more about this notation).

If $\mathbb{L} = (L, \mathbf{p}, \mathbf{q})$ is a multi-based link, then we let $\overline{\mathbb{L}} = (L, \mathbf{q}, \mathbf{p})$ denote the same link, but with the roles of the $\mathbf{w}$ and $\mathbf{z}$ basepoints switched. For notational simplicity, we will write U_t for the variable t associated to $t \in \mathbf{p} \cup \mathbf{q}$ and $\mathbb{Z}_2[U_{\mathbf{p}\cup\mathbf{q}}]$ for the associated ring. There is a natural conjugation action on $\text{Spin}^c(Y)$. To a non-vanishing vector field v representing $\mathfrak{s}$, we define $\overline{\mathfrak{s}}$ to be the Spin^c structure determined by $-v$.

We now describe the conjugation action η on $\mathcal{C}FL^\circ$. Given a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{p}, \mathbf{q})$ for $\mathbb{L} = (L, \mathbf{p}, \mathbf{q})$, we consider the conjugate diagram $\overline{\mathcal{H}} = (-\Sigma, \boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{q}, \mathbf{p})$ for $(Y, \overline{\mathbb{L}})$. There is an obvious correspondence between the intersection points on the two

diagrams. We note however, that $\overline{\mathfrak{s}_{\mathcal{H},\mathbf{p}}(\mathbf{x})} = \mathfrak{s}_{\overline{\mathcal{H}},\mathbf{p}}(\eta(\mathbf{x}))$, by explicit examination of the vector fields constructed by Ozsváth and Szabó in [OS04e], but the types ($\mathbf{w}$ or $\mathbf{z}$) of the basepoints $\mathbf{p}$ and $\mathbf{q}$ has been switched. On the other hand,

$$\mathfrak{s}_{\overline{\mathcal{H}},\mathbf{q}}(\eta(\mathbf{x})) = \overline{\mathfrak{s}_{\mathcal{H},\mathbf{p}}(\mathbf{x})} + PD[L],$$

from Lemma 5.3.2. We extend the map η to arbitrary generators (not just intersection points) by declaring it to be linear in the U_t variables for $t \in \mathbf{p} \cup \mathbf{q}$. By the above formula of Spin^c structures, we see that η maps $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ to $\mathcal{C}FL^\circ(Y, \overline{\mathbb{L}}, \overline{\mathfrak{s}} + PD[L])$. As in [HM15], the map η commutes with change of diagrams maps, up to equivariant filtered chain homotopy.

In this section, we prove the following:

Proposition 5.10.1. *Suppose that $\mathbb{L} = (L, \mathbf{p}, \mathbf{q})$ is a link in Y with a grading assignment $\pi : C(L) \to J$, which is J–null-homologous. The map η induces a chain homotopy equivalence $\eta : \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, \overline{\mathbb{L}}, \overline{\mathfrak{s}})$ which is $\mathbb{Z}_2[U_{\mathbf{p}\cup\mathbf{q}}]$–equivariant and $\mathbb{Z}^{\mathbf{p}} \oplus \mathbb{Z}^{\mathbf{q}}$ filtered. The map η maps homogeneous graded elements (with respect to $A_{Y,\mathbb{L},S}$, $\mathrm{gr}_{\mathbf{w}}$ or $\mathrm{gr}_{\mathbf{z}}$) to homogeneous elements, and satisfies*

$$A_{Y,\overline{\mathbb{L}},S}(\eta(\mathbf{x})) = -A_{Y,\mathbb{L},S}(\mathbf{x}), \qquad \mathrm{gr}_{\mathbf{w}}(\eta(\mathbf{x})) = \mathrm{gr}_{\mathbf{z}}(\mathbf{x}), \qquad \text{and} \qquad \mathrm{gr}_{\mathbf{z}}(\eta(\mathbf{x})) = \mathrm{gr}_{\mathbf{w}}(\mathbf{x}),$$

for a J–Seifert surface S.

Proof. We first consider how η affects the relative Alexander and Maslov gradings. We note that since $\mathbb{L}$ and $\overline{\mathbb{L}}$ have the roles of the type–$\mathbf{w}$ basepoints and type–$\mathbf{z}$ basepoints switched, an easy computation shows that the relative gradings satisfy

$$A(\eta(x), \eta(y))_j = -A(x, y)_j, \qquad \mathrm{gr}_{\mathbf{w}}(\eta(x), \eta(y)) = \mathrm{gr}_{\mathbf{z}}(x, y),$$

$$\text{and} \qquad \mathrm{gr}_{\mathbf{z}}(\eta(x), \eta(y)) = \mathrm{gr}_{\mathbf{w}}(x, y),$$

whenever x and y are homogeneous with respect to the appropriate gradings. Thus it is sufficient to just consider how η affects the absolute gradings of intersection points.

We take a parametrized Kirby diagram $\mathbb{P} = (\phi_0, \lambda, \mathbb{S}_1, f)$ and a surgery triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{p}, \mathbf{q})$, which is subordinate to a β–bouquet for $\mathbb{S}_1$. If $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$, and

$\psi \in \pi_2(\mathbf{x}, \theta, \mathbf{y})$ is a homology triangle, with any other intersection points $\mathbf{x}$ and θ, then the absolute Alexander gradings for $\mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ are defined by the formula

$$A_{Y,\mathbb{L},S}(\mathbf{y})_j = A(\mathbf{x})_j + A(\theta)_j + (n_{\mathbf{p}}(\psi) - n_{\mathbf{q}}(\psi))_j + \frac{\langle c_1(\mathfrak{s}_{\mathbf{p}}(\psi)), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]}{2}.$$

We can form the conjugate triple $\overline{\mathcal{T}} = (-\Sigma, \overline{\boldsymbol{\gamma}}, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\alpha}}, \mathbf{q}, \mathbf{p})$ of $\mathcal{T}$. Here the $\overline{\boldsymbol{\alpha}}$ curves are the same on $-\Sigma$ as the $\boldsymbol{\alpha}$ curves on Σ, and similarly for the other sets of curves, but the roles of the various curves is now changed. The β–bouquet for $\mathbb{S}_1$ now becomes an α–bouquet for $\mathbb{S}_1$, and $\overline{\mathcal{T}}$ is now subordinate to this α–bouquet for the same framed link $\mathbb{S}_1$ in $S^3 \setminus U$.

Using Lemma 5.7.11, we can use $\overline{\mathcal{T}}$ to compute the grading of $\eta(\mathbf{y})$. Notice that there is a canonical identification

$$X_{\alpha\beta\gamma} \cong X_{\overline{\gamma}\overline{\beta}\overline{\alpha}},$$

as oriented manifolds, and similarly this identification respects the embedding of both in $W(S^3, \mathbb{S}_1)$. This diffeomorphism also identifies

$$\Sigma_{\alpha\beta\gamma} \cong \Sigma_{\overline{\gamma}\overline{\beta}\overline{\alpha}},$$

as oriented surfaces, and their images inside of $W(S^3, \mathbb{S}_1)$ coincide as subsurfaces of $\widehat{\Sigma}$. The homology triangle ψ induces a triangle, $\overline{\psi}$, on the conjugate Heegaard triple.

The Alexander grading of $\eta(\mathbf{y})$ can be computed using $\overline{\mathcal{T}}$ and the triangle $\overline{\psi}$, and indeed we see that

$$A(\eta(\mathbf{y}))_j = A(\eta(\mathbf{x}))_j + A(\eta(\theta))_j + (n_{\mathbf{q}}(\overline{\psi}) - n_{\mathbf{p}}(\overline{\psi}))_j + \frac{\langle c_1(\mathfrak{s}_{\mathbf{q}}(\overline{\psi})), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}] \cdot [\widehat{\Sigma}_j]}{2}.$$

We claim that $A(\eta(\mathbf{y}))_j = -A(\mathbf{y})_j$, from which the claim will follow. First observe that $A(\eta(\mathbf{x})) = -A(\mathbf{x})$ and similarly for θ, by an easy model computation for unlinks in S^3 or $(S^1 \times S^2)^{\#k}$ (the map η and the Alexander gradings are natural, so such a computation can be checked on any convenient diagram).

Note that obviously

$$(n_{\mathbf{q}}(\psi) - n_{\mathbf{p}}(\psi))_j = -(n_{\mathbf{p}}(\overline{\psi}) - n_{\mathbf{q}}(\overline{\psi}))_j.$$

We now consider the homological terms involving the homology class of the surface $\Sigma_{\alpha\beta\gamma}$ appearing in the formula for the grading. We note that $\mathfrak{s}_{\mathbf{p}}(\overline{\psi}) = \overline{\mathfrak{s}_{\mathbf{p}}(\psi)}$, the conjugate Spin^c structure. Similarly $\mathfrak{s}_{\mathbf{q}}(\overline{\psi}) = \mathfrak{s}_{\mathbf{p}}(\overline{\psi}) + PD[\Sigma_{\overline{\gamma}\overline{\beta}\overline{\alpha}}]$ by Lemma 5.3.3. Note that $\Sigma_{\alpha\beta\gamma}$ can be capped off with J–Seifert surfaces in the ends of $X_{\alpha\beta\gamma}$ to yield a surface which is homologous to $\widehat{\Sigma}$. Hence we just compute

$$\begin{aligned}
\langle c_1(\mathfrak{s}_{\mathbf{q}}(\overline{\psi})), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j] &= \langle c_1(\mathfrak{s}_{\mathbf{p}}(\overline{\psi}) + PD[\widehat{\Sigma}]), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j] \\
&= \langle c_1(\overline{\mathfrak{s}_{\mathbf{p}}(\psi)} + PD[\widehat{\Sigma}]), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j] \\
&= -\langle c_1(\mathfrak{s}_{\mathbf{p}}(\psi)), [\widehat{\Sigma}_j]\rangle + 2\langle PD[\widehat{\Sigma}], [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j] \\
&= -\big(\langle c_1(\mathfrak{s}_{\mathbf{p}}(\psi)), [\widehat{\Sigma}_j]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}_j]\big).
\end{aligned}$$

Thus each term in $A(\mathbf{y})_j$ is changed to its negative in $A(\eta(\mathbf{y}))_j$, from which we conclude that $A(\eta(\mathbf{y}))_j = -A(\mathbf{y})_j$.

The proof of the statement about the Maslov gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ follows similarly. □

The map η restricts to the $\mathbb{Z}_2$–modules $\widehat{CFL}$ to thus yield an isomorphism between $\widehat{HFL}(Y,\mathbb{L})_s$ and $\widehat{HFL}(Y,\mathbb{L})_{-s}$ where s denotes our Alexander grading. Since the gradings from [OS08] also satisfy this relation with respect to the conjugation action, the two gradings must coincide.

Indeed we also note that that the homological grading shift from Equation (5.15) can be read off in our setting as well, as $\mathrm{gr}_{\mathbf{w}}(\eta(x)) = \mathrm{gr}_{\mathbf{z}}(x) = \mathrm{gr}_{\mathbf{w}}(x) - 2A(x)$, using our notation.

5.10.2 Collapsing gradings

Finally we make a remark on collapsing grading assignments. If $\pi : L \to J$ is a grading assignment, and $f : J \to J'$ is a map, then the map $\pi' = f \circ \pi : L \to J'$ is also a grading assignment.

Lemma 5.10.2. *The multi-grading $A_{Y,\mathbb{L},S,\pi'}$ is obtained from $A_{Y,\mathbb{L},S,\pi}$ by collapsing gradings,*

in the sense that

$$(A_{Y,\mathbb{L},\mathcal{S},\pi'})_{j'} = \sum_{j \in f^{-1}(j')} (A_{Y,\mathbb{L},\mathcal{S},\pi})_j.$$

The empty sum is interpreted as zero.

Proof. This follows from the fact that the stated equation holds for unlinks in $(S^1 \times S^2)^{\#k}$ by Equation (5.8), and the fact that the expressions $\langle c_1(\mathfrak{s}), \widehat{\Sigma}_j \rangle$ and $[\Sigma] \cdot [\Sigma_j]$ each also satisfy the analogous equation associated to collapsing gradings. □

5.11 A new proof of a bound of Ozsváth and Szabó on τ

In [OS03a], Ozsváth and Szabó define a homomorphism $\tau : \mathcal{C} \to \mathbb{Z}$, and prove a bound on how τ behaves with respect to knot cobordisms in negative definite 4–manifolds. In this section, we give a new proof of this bound, as a consequence of our grading formula.

We briefly recall the definition of τ and some other notation. The differential on $\widehat{\mathcal{CFL}}(S^3, K, w, z)$ is defined by counting disks which don't go over any of the w or z basepoints. If we allow disks to go over z but not w, the resulting chain complex is $\widehat{CF}(S^3, w)$, but with a $\mathbb{Z}$–filtration. Given a diagram $\mathcal{H}$ for (S^3, K), we let $\mathcal{F}_s(\mathcal{H}, K)$, denote the subgroup of $\widehat{CF}(\mathcal{H})$ generated by $\mathbf{x}$ with $A(\mathbf{x}) \leq s$. The filtered chain homotopy type is a knot invariant, so we will omit $\mathcal{H}$ from the notation, and write $\mathcal{F}_s(K)$. The integer $\tau(K)$ is defined as the minimal s such that the map

$$H_*(\mathcal{F}_s(K)) \to \widehat{HF}(S^3) \cong \mathbb{Z}_2$$

induced by inclusion is nontrivial. It is easy to check that when we set $V = 1$, and $U = 0$ a map $F : \mathcal{CFL}^-(S^3, K_1) \to \mathcal{CFL}^-(S^3, K_2)$ which changes Alexander grading by $+\Delta$ induces a map $\mathcal{F}_s(K_1) \to \mathcal{F}_{s+\Delta}(K_2)$. If we can ensure the induced map $\widehat{HF}(S^3) \to \widehat{HF}(S^3)$ is an isomorphism, then we can quickly obtain bounds on the invariant τ.

Suppose as in [OS03a] that W is an oriented 4–manifold with $\partial W = S^3$ and $b_2^+(W) =$

$b_1(W) = 0$. If $e_1, \dots, e_b$ is an orthonormal basis of $H^2(W; \mathbb{Z})$, and

$$[\Sigma] = s_1 \cdot e_1 + \cdots + s_b \cdot e_b,$$

we define

$$\big|[\Sigma]\big| = |s_1| + \cdots + |s_b|,$$

a well defined integer (see [OS03a]).

The following theorem is essentially [OS03a, Theorem 1.1], for oriented surfaces in negative definite 4–manifolds. We give a new proof, using our grading formula:

Theorem 5.11.1. *Suppose that $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is an oriented knot cobordism with $b_2^+(W) = 0 = b_1(W)$. Then*

$$\tau(K_2) \leq \tau(K_1) - \frac{\big|[\Sigma]\big| + [\Sigma] \cdot [\Sigma]}{2} + g(\Sigma).$$

Proof. As above, we write $b = b_2(W) = b_2^-(W)$. Decorate Σ by picking a path from one boundary to the other, and letting $\Sigma_{\mathbf{w}}$ be a regular neighborhood of this path. Let F be the resulting decorated link cobordism. Using [OS03b, Section 9], if $\mathfrak{s}$ is a Spin^c structure on $W \setminus B$ with $c_1(\mathfrak{s})^2 + b = 0$, then the map $\widehat{F}_{W,\gamma,\mathfrak{s}} : \widehat{HF}(S^3) \to \widehat{HF}(S^3)$ is an isomorphism. Here γ is the path we chose on Σ (though the fact that the map is an isomorphism doesn't depend on the path). By the grading formula, the maps $F_{W,F,\mathfrak{s}}$ induce a map from $\mathcal{F}_{\tau(K_1)}(K_1)$ to $\mathcal{F}_{\tau(K_1)+\Delta}(K_2)$ which is equal to the map $\widehat{F}_{W,\mathfrak{s}}$ on homology, where the grading change Δ is given by the formula

$$\Delta = \frac{\langle c_1(\mathfrak{s}), [\Sigma]\rangle - [\Sigma] \cdot [\Sigma]}{2} + g(\Sigma).$$

There are 2^b characteristic vectors C with $C \cdot C = -b$ (these are sums of ± 1 times each e_i), and upon direct inspection, we see that

$$-\big|[\Sigma]\big| = \min_{C \in \mathrm{Char}(W) : C \cdot C = -b} \langle C, [\Sigma]\rangle.$$

Also every characteristic vector of the intersection form is the first Chern class of a Spin^c structure. Pick an $\mathfrak{s} \in \mathrm{Spin}^c(W)$ with $c_1(\mathfrak{s})^2 + b = 0$ such that $c_1(\mathfrak{s})$ realizes the above

equality. Pick an $a \in \mathcal{F}_{\tau(K_1)}$ such that $[a] \neq 0 \in \widehat{HF}(S^3)$. Then $F_{W,F,\mathfrak{s}}(a) \in \mathcal{F}_{\tau(K_1)+\Delta}(K_2)$ is an element which is nonzero in $\widehat{HF}(S^3)$. Hence

$$\tau(K_2) \leq \tau(K_1) + \Delta = \tau(K_1) + \frac{-\big|[\Sigma]\big| - [\Sigma]\cdot[\Sigma]}{2} + g(\Sigma).$$

□

For example, if $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is a knot cobordism with $g(\Sigma) = g$ and Σ connected, such that W is a rational homology cobordism, then $|\tau(K_1) - \tau(K_2)| \leq g$.

In light of the proof of the previous bound, we see that homomorphism $\tau : \mathcal{C} \to \mathbb{Z}$ naturally factors through the **integer homology concordance group** $\mathcal{C}^3_{\mathbb{Z}}$, which is generated by pairs (Y^3, K) of knots in integer homology spheres modulo the relation that two pairs are equivalent if there is an integer homology cobordism between the 3–manifolds, which contains an oriented genus zero surface between the two knots.

5.12 t-modified knot Floer homology and a bound on the $\Upsilon_K(t)$ invariant

In [OSS14], Ozsváth, Stipsicz and Szabó define a homomorphism from the smooth concordance group $\mathcal{C}$ to the group of piecewise linear, $\mathbb{R}$–valued functions over $[0, 2]$. In this section, we recall their construction, and show how our grading formula and link cobordism maps yields a bound for $\Upsilon_K(t)$.

We recall (a version of) their construction. Suppose that $K \subseteq S^3$ is an oriented knot, and $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$ is a diagram for (S^3, K, w, z). If $0 \leq t = \frac{m}{n} \leq 2$, we define the t–grading on intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ by

$$\mathrm{gr}_t(\mathbf{x}) = (1 - \tfrac{t}{2})\,\mathrm{gr}_{\mathbf{w}}(\mathbf{x}) + \tfrac{t}{2}\,\mathrm{gr}_{\mathbf{z}}(\mathbf{x}).$$

Note that this agrees with the formula from [OSS14]. We then consider the chain complex $tCFK^-(K)$, a chain complex over $\mathbb{Z}_2[v^{1/n}]$, generated by intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$

with differential

$$\partial(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} \#\widehat{\mathcal{M}}(\phi) v^{tn_z(\phi)+(2-t)n_w(\phi)} \cdot \mathbf{y}.$$

The grading gr_t descends to give a grading on $tCFK^-(K)$. The differential lowers degree by 1, and the action of v also lowers degree by 1. The number $\Upsilon_K(t) \in \mathbb{R}$ is defined as the maximal gr_t–grading of any homogeneous non-torsion element of $tHFK^-(K)$.

To obtain bounds on $\Upsilon_K(t)$, we first show that we can recover $tCFK^-(K)$ from $\mathcal{C}FL^-(S^3, K, w, z)$. Define the ring

$$R_t^- = \mathbb{Z}_2[U, V, v^{1/n}]/(U - v^{2-t}, V - v^t),$$

and define R_t^∞ similarly.

Lemma 5.12.1. *The chain complex $\mathcal{C}FL^-(S^3, K, w, z) \otimes_{\mathbb{Z}_2[U,V]} R_t^-$ is naturally isomorphic to $tCFK^-(K)$ as a $\mathbb{Z}_2[v^{1/n}]$–module.*

Proof. We first describe an isomorphism between the rings R_t^- and $\mathbb{Z}_2[v^{1/n}]$. Noting that $t = \frac{m}{n}$, the map from $\mathbb{Z}_2[U, V, v^{1/n}]/(U - v^{2-t}, V - v^t)$ to $\mathbb{Z}_2[v^{1/n}]$ is

$$U^i V^j v^s \mapsto v^{i(2-t)+jt+s}$$

and a map in the opposite direction is

$$v^s \mapsto v^s.$$

To define maps between the chain complexes, we use the above maps on rings, extended over linear combinations of intersection points. That these maps are chain maps is immediate. It's also clear that these two maps are inverses of each other. □

Phrased another way, R_t^- is in fact just $\mathbb{Z}_2[v^{1/n}]$ with a module action of $\mathbb{Z}_2[U, V]$ declared. In a similar manner the chain complex $tCFK^\infty(S^3, K)$ is isomorphic to $\mathcal{C}FL^\infty(S^3, K) \otimes_{\mathbb{Z}_2[U,V,U^{-1},V^{-1}]} R_t^\infty$ as a $\mathbb{Z}_2[v^{1/n}, v^{-1/n}]$–module.

As a consequence, if $(W, \Sigma) : (S^3, K_1) \to (S^3, K_2)$ is a link cobordism and $\mathfrak{s} \in \mathrm{Spin}^c(W)$, then the link cobordism maps $F^-_{W,F,\mathfrak{s}}$ induce a map from $tCFK^-(K_1)$ to $tCFK^-(K_2)$.

Lemma 5.12.2. *Suppose that (W, Σ) is a knot cobordism from (S^3, K_1) to (S_3, K_2) and $\mathfrak{s} \in \mathrm{Spin}^c(W)$. If the 3–manifold cobordism map $F^\infty_{W,\mathfrak{s}}$ is an isomorphism on $HF^\infty(S^3)$, and we decorate Σ with a dividing set consisting of two arcs going from K_1 to K_2, such that the induced region $\Sigma_{\mathbf{w}}$ is equal to a strip on Σ from K_1 to K_2, then the link cobordism map*

$$F^-_{W,F,\mathfrak{s}} : tHFK^-(K_1) \to tHFK^-(K_2)$$

sends non-torsion elements of $tHFK^-(K_1)$ to non-torsion elements of $tHFK^-(K_2)$.

Proof. We consider $tCFK^\infty(K_i)$, which are modules over $\mathbb{Z}_2[v^{-1/n}, v^{1/n}]$. There is a natural map from $tHFK^-(K_i)$ to $tHFK^\infty(K_i)$, given by inclusion of complexes. It's easily seen that an element of $tHFK^-(K)$ is non-torsion iff it maps to a nonzero element of $tHFK^\infty(K)$. On the other hand, there is an isomorphism

$$tHFK^\infty(K) \cong H_*\big(\mathcal{C}FL^\infty(S^3, K) \otimes_{\mathbb{Z}_2[U,V,U^{-1},V^{-1}]} R^\infty_t\big).$$

We claim that this is isomorphic to just $H_*(\mathcal{H}FL^\infty(S^3, K)) \otimes_{\mathbb{Z}_2[U,V,U^{-1},V^{-1}]} R^\infty_t$. To see this, it's easiest to note that $\mathcal{C}FL^\infty(S^3, K)$ decomposes over Alexander gradings, and the 0–graded part is the standard $CFK^\infty(S^3, K)$ complex, viewed as a module over $\mathbb{Z}_2[\overline{U}, \overline{U}^{-1}]$ where $\overline{U} = UV$. Using the classification of chain complexes over the PID $\mathbb{Z}_2[\overline{U}, \overline{U}^{-1}]$ (see [HMZ16, Lemma 6.1]) and extending over all of $\mathbb{Z}_2[U, V, U^{-1}, V^{-1}]$, we can write $\mathcal{C}FL^\infty(S^3, K)$ as a sum of two step complexes $(a_i \longrightarrow b_i)$ for $i = 1, \dots, n$, as well as a single summand (a_0) with vanishing differential in zero $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ gradings (of course this decomposition doesn't respect the $\mathbb{Z} \oplus \mathbb{Z}$ filtration). Tensoring with R^∞_t preserves this decomposition of the chain complex, and hence $tHFK^\infty(S^3, K)$ is the free $\mathbb{Z}_2[v^{1/n}, v^{-1/n}]$–module generated by $[a_0]$.

Phrased another way, the natural map

$$\mathcal{H}FL^\infty(S^3, k) \otimes_{\mathbb{Z}_2[U,V,U^{-1},V^{-1}]} R^\infty_t \to tHFK^\infty(K)$$

is an isomorphism

Additionally, since $\mathcal{H}FL^\infty(S^3, K_i) \cong \mathbb{Z}_2[U, V, U^{-1}, V^{-1}]$, on homology the link cobordism map is determined by its value on $1 \in \mathbb{Z}_2[U, V, U^{-1}, V^{-1}] \cong \mathcal{H}FL^\infty(S^3, K_1)$. Since the

link cobordism maps are graded, $F^\infty_{W,F,\mathfrak{s}}$ must be $c \cdot [x]$, where $[x]$ is a nonzero element of $\mathcal{H}FL^\infty(S^3, K_2)$ in the expected grading and $c \in \mathbb{Z}_2$. As $F^\infty_{W,F,\mathfrak{s}}$ covers $F^\infty_{W,\mathfrak{s}}$ when we set $V = 1$, we see that $c = 1$ since by assumption $F^\infty_{W,\mathfrak{s}}$ is an isomorphism. In particular, the map induced by $F^\infty_{W,F,\mathfrak{s}}$ on $tHFL^\infty$ is an isomorphism. From these considerations applied to the commutative diagram

$$\begin{array}{ccc} tHFK^-(K_1) & \xrightarrow{\iota} & tHFK^\infty(K_1) \\ \downarrow{\scriptstyle F^-_{W,F,\mathfrak{s}}} & & \cong\downarrow{\scriptstyle F^\infty_{W,F,\mathfrak{s}}} \\ tHFK^-(K_2) & \xrightarrow{\iota} & tHFK^\infty(K_2), \end{array}$$

the claim now follows. □

Finally, we note that if $(W, F, \mathfrak{s})$ is a link cobordism, the gr_t–grading change can be computed using the $\mathrm{gr}_{\mathbf{w}}$– and $\mathrm{gr}_{\mathbf{z}}$–grading changes we've computed in previous sections, as gr_t is a convex combination of $\mathrm{gr}_{\mathbf{x}}$ and $\mathrm{gr}_{\mathbf{z}}$. We compute that if $\mathbf{x}$ is a homogeneously graded element, then

$$\begin{aligned} &\mathrm{gr}_t(F^-_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_t(\mathbf{x}) \\ &= \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} + t \cdot \left(\frac{-\langle c_1(\mathfrak{s}), \widehat{\Sigma}\rangle + [\widehat{\Sigma}] \cdot [\widehat{\Sigma}]}{2} \right) \qquad (5.16) \\ &\quad + \left(1 - \frac{t}{2}\right) \cdot \widetilde{\chi}(\Sigma_{\mathbf{w}}) + \frac{t}{2} \cdot \widetilde{\chi}(\Sigma_{\mathbf{z}}), \qquad (5.17) \end{aligned}$$

using the grading change formulas for $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$.

We now proceed to prove Theorem 5.1.1. Recall from the introduction that if s is an integer and $t \in [0, 2]$, we define

$$M_t(s) = \max_{a \in 2\mathbb{Z}+1} \frac{-a^2 + 1 + 2ast - 2s^2 t}{4}.$$

If $e_1, \ldots, e_n$ is an orthonormal basis for $H_2(W; \mathbb{Z})$, where W is a negative definite 4-manifold, and $[\Sigma] = s_1 \cdot e_1 + \cdots + s_n \cdot e_n \in H_2(W; \mathbb{Z})$, we define

$$M_t([\Sigma]) = \sum_{i=1}^{n} M_t(s_i).$$

Theorem 5.1.1. *Suppose that $(W,\Sigma):(S^3,K_1)\to(S^3,K_2)$ is an oriented knot cobordism with $b_2^+(W)=0=b_1(W)$. Then*

$$\Upsilon_{K_2}(t)\geq\Upsilon_{K_1}(t)+M_t([\Sigma])+g(\Sigma)\cdot(|t-1|-1).$$

Proof. We construct a decorated link cobordism (W,F) with $F=(\Sigma,\mathcal{A})$, where the divides $\mathcal{A}$ are the boundary of a neighborhood of a path on the surface from K_1 to K_2, which we let designate as type-**w**. By the proof [OS03b, Theorem 9.1], the map $F^{\infty}_{W,\mathfrak{s}}:HF^{\infty}(S^3)\to HF^{\infty}(S^3)$ is an isomorphism for any Spin^c structure $\mathfrak{s}\in\text{Spin}^c(W)$. By Lemma 5.12.2, the map on $tHFL^-$ induced by the link cobordism map $F^-_{W,F,\mathfrak{s}}$ sends non-torsion elements to non-torsion elements, for each t. Using the gr_t–grading change from Equation (5.17), the map $F^-_{W,F,\mathfrak{s}}$ induces a gr_t–grading change of

$$\frac{c_1(\mathfrak{s})^2+b_2(W)-2t\cdot\langle c_1(\mathfrak{s}),[\Sigma]\rangle+2t\cdot[\Sigma]\cdot[\Sigma]}{4}-t\cdot g(\Sigma)$$

on $tHFL^-$. The set of Chern classes of Spin^c structures on W is equal to the set of characteristic vectors. We note that the set of characteristic vectors in $H^2(W;\mathbb{Z})$ is equal to (the Poincaré duals of) the elements $a_1\cdot e_1+\cdots a_b\cdot e_b$ where each a_i is an odd integer. Thus if $PD[c_1(\mathfrak{s})]=a_1\cdot e_1+\cdots a_b\cdot e_b$, we have that

$$\frac{c_1(\mathfrak{s})^2+b_2(W)-2t\cdot\langle c_1(\mathfrak{s}),[\Sigma]\rangle+2t\cdot[\Sigma]\cdot[\Sigma]}{4}=\sum_{i=1}^{b}\frac{-a_i^2+1+2ta_is_i-2ts_i^2}{4}.$$

Since the summands are independent from each other in i, we can maximize together or separately. Taking the maximum over $\mathfrak{s}\in\text{Spin}^c(W)$ for a fixed t, we get

$$\Upsilon_{K_2}(t)\geq\Upsilon_{K_1}(t)+M_t([\Sigma])-t\cdot g(\Sigma).$$

Letting $F_a(t)=\frac{1}{4}(-a^2+1+2ast-2s^2t)$, we note that

$$F_a(t)=F_{-a+2s}(2-t),$$

so that $M_t(n)=M_{2-t}(n)$. Using also the symmetry of the $\Upsilon_K(t)$ invariant, we see that

$$\Upsilon_{K_2}(t)\geq\Upsilon_{K_1}(t)+M_t([\Sigma])-(2-t)\cdot g(\Sigma).$$

Combining this bound with first yields the stated inequality. □

Remark 5.12.3. One could hope to refine the above bound even further, by considering different surfaces with divides and trying to optimize the expression

$$\left(1-\frac{t}{2}\right)\cdot\widetilde{\chi}(\Sigma_{\mathbf{w}})+\frac{t}{2}\cdot\widetilde{\chi}(\Sigma_{\mathbf{z}}).$$

We now remark that this doesn't yield any stronger bounds. If we pick a dividing set on Σ consisting of two arcs from K_1 to K_2 which cut it into two connected components $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, then $\widetilde{\chi}(\Sigma_{\mathbf{w}})=-2g(\Sigma_{\mathbf{w}})$ and similarly for $\Sigma_{\mathbf{z}}$. Adapting Theorem 5.1.8, which we prove in the next section, we can actually see that the map $F^{\infty}_{W,F,\mathfrak{s}}$ on $\mathcal{H}FL^{\infty}$ is an isomorphism. On the other hand, using $k=g(\Sigma_{\mathbf{w}})$, we see that we can replace the summand $-t\cdot g(\Sigma)$ appearing in the gr_t–grading change formula with

$$\max_{\substack{k\in\mathbb{Z}\\0\leq k\leq g(\Sigma)}} -(2-t)k-t(g(\Sigma)-k)=g(\Sigma)\cdot(|t-1|-1),$$

so the bound we get by varying the set of divides is no better than the one we proved using the divides where $\Sigma_{\mathbf{w}}$ is a strip on Σ, and then using the symmetry of $\Upsilon_K(t)$.

5.12.1 Additional examples of the bound

As a corollary, we see that the above theorem recovers the genus bound from [OSS14]:

Corollary 5.12.4. *If $(W,\Sigma):(S^3,K_1)\to(S^3,K_2)$ is an oriented knot cobordism, and W is a rational homology cobordism, then*

$$|\Upsilon_{K_2}(t)-\Upsilon_{K_1}(t)|\leq t\cdot g(\Sigma).$$

Proof. This follows immediately from Theorem 5.1.1. □

Also, as another example, we recover [OSS14, Proposition 1.10]:

Corollary 5.12.5. *If K_- and K_+ are knots in S^3, which differ by a crossing change, then*

$$\Upsilon_{K_+}(t)\leq\Upsilon_{K_-}(t)\leq\Upsilon_{K_+}(t)+1-|t-1|.$$

Proof. Suppose that K_- and K_+ are two knots which differ by a crossing change (and K_- has the negative crossing and K_+ has the positive crossing). We can construct a negative definite link cobordism (W_1, Σ_1) from (S^3, K_-) to (S^3, K_+) and also a negative definite link cobordism (W_2, Σ_2) from (S^3, K_+) to (S^3, K_-). Each is formed by adding a 2-handle with framing -1, around the crossing, as shown in Figure 5.17. The homology class of the surface can be read off from the intersection number of the knot on the incoming end with the Seifert disk for the -1 framed unknot.

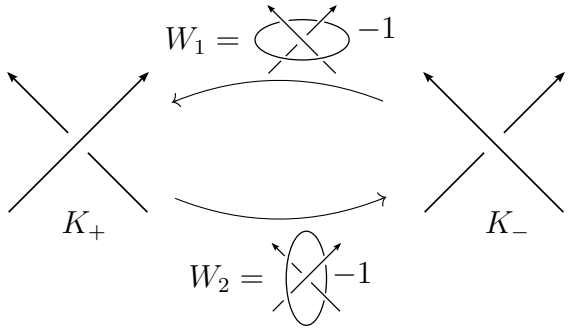

Figure 5.17: **The two knot cobordisms** (W_i, Σ_i) **between** (S^3, K_-) **and** (S^3, K_+)**.**

If the algebraic intersection number is n, then in $H_2(W; \mathbb{Z})$ the surface Σ induces homology class $[\Sigma] = n \cdot E$, where $E \subseteq \overline{\mathbb{CP}}^2$ is a sphere with self intersection -1. Hence $[\Sigma_1] = 2 \cdot E$ and $[\Sigma_2] = 0$. We note that

$$M_t(0) = 0, \qquad \text{and} \qquad M_t(2) = -(1 - |t - 1|).$$

Applying Theorem 5.1.1, we see

$$\Upsilon_{K_+}(t) \leq \Upsilon_{K_-}(t) \leq \Upsilon_{K_+}(t) + 1 - |t - 1|.$$

□

5.12.2 Positive torus knots

We show that that the bound from Theorem 5.1.1 applied repeatedly to torus knots, is sharp.

According to [OSS14, Proposition 5], if n is a positive integer and $T_{n,n+1}$ denotes the

$(n, n+1)$–torus knot, then for $t \in \left[\frac{2i}{n}, \frac{2i+2}{n}\right]$, we have

$$\Upsilon_{T_{n,n+1}}(t) = -i(i+1) - \frac{1}{2}n(n-1-2i)t.$$

Using this, Feller and Krcatovich show in [FK16, Proposition 6] that $\Upsilon_K(t)$ can be computed inductively for positive torus knots, using the recursion

$$\Upsilon_{T_{a,b}}(t) = \Upsilon_{T_{a,b-a}}(t) + \Upsilon_{T_{a,a+1}}(t).$$

On the other hand, note that there is a natural knot cobordism from $(S^3, T_{a,b-a})$ to $(S^3, T_{a,b})$, obtained by -1 surgery on an unknot encircling the torus which $T_{a,b-a}$ is embedded on. If D is the Seifert disk for this unknot, we can orient D so that $T_{a,b-a}$ intersects D at a points, positively. Applying Theorem 5.1.1 yields

$$\Upsilon_{T_{a,b}}(t) \geq \Upsilon_{T_{a,b-a}}(t) + M_t(a). \tag{5.18}$$

A special case is the knot cobordism from $T_{n,1}$ (the unknot) to $T_{n,n+1}$. In this case, our bound reads

$$\Upsilon_{T_{n,n+1}}(t) \geq 0 + M_t(n) = \max_{a \in 2\mathbb{Z}+1} \frac{-a^2 + 1 + 2atn - 2tn^2}{4}.$$

On the other hand, using the computation from [OSS14] of $\Upsilon_{T_{n,n+1}}(t)$, and the following computation, we see the bound from Equation (5.18) is sharp:

Lemma 5.12.6. *If $t \in \left[\frac{2i}{n}, \frac{2i+2}{n}\right]$ then*

$$M_t(n) = -i(i+1) - \frac{1}{2}n(n-1-2i)t.$$

Proof. We claim that for t in the stated interval, the maximum of $\frac{1}{4}(-a^2 + 1 + 2tan - 2tn^2)$ occurs at $a = 2i+1$. Write $F_a(t) = \frac{1}{4}(-a^2 + 1 + 2atn - 2tn^2)$. Note that plugging $a = 2i+1$ into the expression for $F_a(t)$ yields the expression in the lemma statement. For real a and a fixed t, the maximum of $F_a(t)$, a parabola in a, occurs at $a = tn \in [2i, 2i+2]$. Hence for a fixed t the maximum of $F_t(a)$ over odd integers a will occur at one of $2i-1, 2i+1$, or $2i+3$. It is now an easy matter to compute that the lines $F_{2i-1}(t)$ and $F_{2i+1}(t)$, intersect

at $t = 2i/n$ and the lines $F_{2i+1}(t)$ and $F_{2i+3}(t)$ intersect at $t = 2(i+1)/n$. As the slope of the line $F_a(t)$ is increasing in a, it follows that on the stated interval, $M_t(n) = F_{2i+1}(t)$, completing the proof. □

5.13 The adjunction relation

As another application of our grading formula, we prove an adjunction relation similar to [OS04b, Theorem 3.1]. The standard adjunction inequality for $HF^+(Y, \mathfrak{s})$ follows from this relation.

If Σ is a closed surface and $A_1, \dots, A_g, B_1, \dots, B_g$ is a symplectic basis of $H_1(\Sigma; \mathbb{Z})$ which is obtained by taking a collection of simple closed curves A_i and B_i on the surface with geometric intersection numbers $|A_i \cap B_j| = \delta_{ij}$, we let $\xi(\Sigma)$ denote the element

$$\xi(\Sigma) = \prod_{i=1}^{g(\Sigma)} (U + A_i \cdot B_i) \in \mathbb{Z}_2[U] \otimes \Lambda^*(H_1(\Sigma; \mathbb{Z})).$$

The element $\xi(\Sigma)$ is defined in terms of the chosen symplectic basis.

Theorem 5.1.7. *Suppose that Σ is a connected, closed, oriented, embedded surface (of any genus), inside of a cobordism $W : Y_1 \to Y_2$ with W, Y_1 and Y_2 connected, and that $\mathfrak{s} \in \mathrm{Spin}^c(W)$ is a Spin^c structure with*

$$\langle c_1(\mathfrak{s}), [\Sigma] \rangle - [\Sigma] \cdot [\Sigma] = -2g(\Sigma),$$

then

$$F^\circ_{W,\mathfrak{s}}(\cdot) = F^\circ_{W,\mathfrak{s}-PD[\Sigma]}(\iota_*(\xi(\Sigma)) \otimes \cdot),$$

for $\circ \in \{+, -, \infty\}$, where $\iota_* : H_1(\Sigma; \mathbb{Z}) \to H_1(W; \mathbb{Z})$ is the map induced by inclusion. Here $\xi(\Sigma)$ is the element described above, for a choice of symplectic basis of $H_1(\Sigma; \mathbb{Z})$.

Remark 5.13.1. The above theorem is phrased in terms of the $+, -$ and ∞ flavors, with a single basepoint, so we don't have to keep track of paths. For the hat flavor, an analogous statement holds, but one needs to keep track of a path.

Remark 5.13.2. When $g(\Sigma) > 0$, analogous adjunction relations can be found in [OS04b, Proposition 3.1] (for Heegaard Floer) and [OS00b, Theorem 1.3] (for Seiberg–Witten). Let us briefly compare our version to theirs, which is corrected by a factor of $\epsilon = \pm 1$ on $PD[\Sigma]$, equal to the sign of $\langle c_1(\mathfrak{s}), \Sigma\rangle$. If $\epsilon = -1$, then our formulas agree. If $\epsilon = +1$, then the surface $-\Sigma$ obtained by reversing orientation satisfies $\langle c_1(\mathfrak{s}), [-\Sigma]\rangle - [-\Sigma]\cdot[-\Sigma] \leq -2g(-\Sigma)$. Add n homologically trivial handles to the type–$\mathbf{z}$ region of $-\Sigma$ to turn the previous inequality into an equality, and let Σ' be the resulting surface. We note that $[\Sigma'] = -[\Sigma]$. Hence our adjunction relation now reads

$$F^{\circ}_{W,\mathfrak{s}}(\cdot) = F^{\circ}_{W,\mathfrak{s}+\epsilon PD[\Sigma]}(\iota_*(\xi(\Sigma')) \otimes \cdot),$$

and now we observe that from our explicit description of ξ, we have $\iota_*(\xi(\Sigma')) = U^n \cdot \iota_*(\xi(\Sigma))$, which fits nicely with the formula from [OS00b, Theorem 1.3].

Remark 5.13.3. If Σ has genus zero, then the situation is similar to [FS95, Lemma 5.2]. To see that our result is consistent, suppose that S is an embedded sphere of negative self intersection with $|\langle c_1(\mathfrak{s}), S\rangle| + S\cdot S \geq 0$ (this is the situation they consider). In this case, let $\epsilon \in \{\pm 1\}$ be the sign of $\langle c_1(\mathfrak{s}), S\rangle$. We have

$$\langle c_1(\mathfrak{s}), -\epsilon S\rangle - S\cdot S \leq -2g(S).$$

Adding n null-homologous handles to the surface to achieve equality in the previous equation, and then applying our result, we see that

$$F^{\circ}_{W,\mathfrak{s}} = U^n \cdot F^{\circ}_{W,\mathfrak{s}+\epsilon PD[S]}.$$

If S is an embedded sphere in a 4–manifold X which admits an admissible cut which is disjoint from S, we recover an analog of [FS95, Lemma 5.2], for the Ozsváth–Szabó mixed invariant.

To prove the above theorem, we will consider what happens to the maps $F^{-}_{W,F,\mathfrak{s}}$ when we set either $U = 1$ or $V = 1$. To find out what the maps are, we need two results from Chapter 4. The first is Proposition 4.7.16, which states that if the ends of B are adjacent to

basepoints w_1, w_2, z_1 and z_2, then

$$F_B^{\mathbf{z}} \simeq \Phi_{w_1} F_B^{\mathbf{w}} + F_B^{\mathbf{w}} \Phi_{w_1}. \tag{5.19}$$

(Note also that the equivalence is easily seen to also hold if we replace both instances of w_1 with w_2). The second relation we need is Lemma 4.14.13, which implies that after we set all $U_{\mathbf{w}}$ variables equal to 1, if λ is the arc of $\mathbb{L}$ going from z to z', containing only the basepoint w, then we have

$$\Phi_w|_{U_{\mathbf{w}}=1} \simeq A_\lambda + V_z \Psi_z. \tag{5.20}$$

We note here we are writing Ψ_z for the map on CF^- that we would normally write Φ_z for, using the notation from Chapter 2. In the case that the link component containing w has only two basepoints, the relative homology map A_λ is replaced with the map for the normal homology action for that link component.

We use the previous two relations to prove the following lemma:

Lemma 5.13.4. *Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link in Y and B_1 and B_2 are two bands attached to L. Further assume that both ends of B_1 are attached to the same component of $L \setminus (\mathbf{w} \cup \mathbf{z})$, and that the ends of B_2 are both attached to opposite sides of B_1 (see Figure 5.19). Write (W, F) for the decorated link cobordism where $W = Y \times [0,1]$ and F is formed by adding the bands B_1 and B_2, to $\mathbb{L}$, both as type–$\mathbf{z}$ bands. Let γ_1 and γ_2 two simple closed curves on Σ be defined by taking γ_1 to be the core of B_1 concatenated with an arc λ_0 of $K \setminus (\mathbf{w} \cup \mathbf{z})$ between the two ends of B_1, and define γ_2 to be the core of B_2, as in Figure 5.19. After tensoring to set $U = 1$, the map $F^-_{W,F,\mathfrak{s}}$ on $CF^-(Y, z, \mathfrak{s} - PD[L])$ is equal to*

$$(V + A_{\gamma_2} A_{\gamma_1}).$$

Proof. The extra components of L play no role in the proof, so let us assume for the sake of simplicity that $\mathbb{L} = \mathbb{K} = (K, \mathbf{w}, \mathbf{z})$, for a knot K. Also for notational simplicity, we will assume that $\mathbb{K}$ has exactly two basepoints, w_1 and z_1. For knots with more basepoints, the assumption that B_1 has ends in the same component of $K \setminus (\mathbf{w} \cup \mathbf{z})$ will be needed, though the proof will be essentially unchanged.

The proof is to decompose the surface nicely, and then use relations in the previous paragraphs from Chapter 2. We consider the decomposition shown in Figure 5.18.

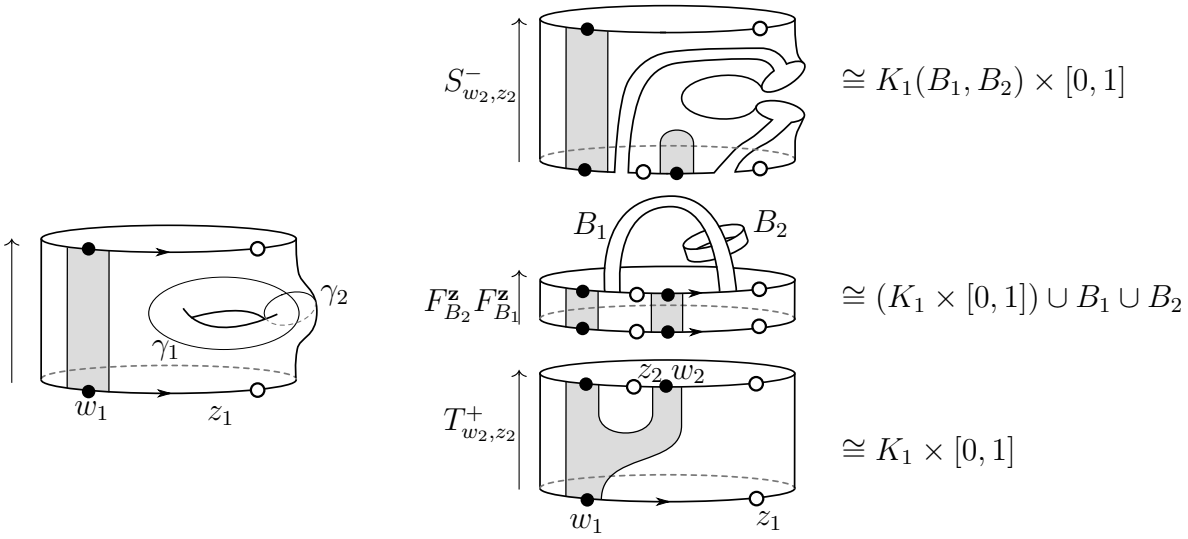

Figure 5.18: **A decomposition of the surface with divides F into three pieces which we can use to compute the map.** The first piece is corresponds to a $T^+_{w_2,z_2}$ quasi-stabilization. The second piece is two type–$\mathbf{z}$ bands, and the third piece is an $S^-_{w_2,z_2}$ quasi-destabilization.

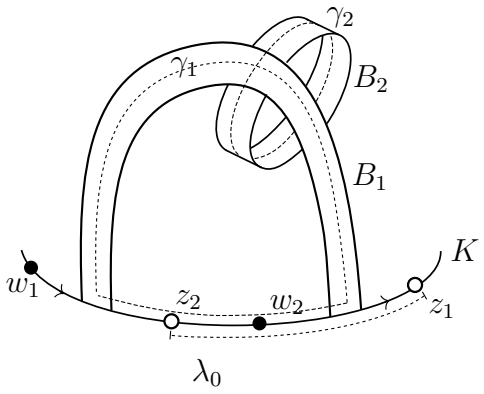

Figure 5.19: **The homology classes γ_1 and γ_2, as well as the relative homology class λ_0.**

Using this decomposition, the maps are

$$F^\circ_{W,F,\mathfrak{s}} \simeq S^-_{w_2,z_2}F^{\mathbf{z}}_{B_2}F^{\mathbf{z}}_{B_1}T^+_{w_2,z_2}.$$

Note that by direct computation of the quasi-stabilized differential in Chapter 2, we can compute that $S^-_{w_2,z_2} \simeq T^-_{w_2,z_2}\Phi_{w_2}$. Thus by using Equations (5.19) and (5.20), we see that $F^\circ_{W,F,\mathfrak{s}}$ is equivalent to

$$T^-_{w_2,z_2}\Phi_{w_2}(\Phi_{w_2}F^{\mathbf{w}}_{B_2} + F^{\mathbf{w}}_{B_2}\Phi_{w_2})(\Phi_{w_2}F^{\mathbf{w}}_{B_1} + F^{\mathbf{w}}_{B_1}\Phi_{w_2})T^+_{w_2,z_2}$$

which is equivalent to

$$T^-_{w_2,z_2}\Phi_{w_2}F^{\mathbf{w}}_{B_2}\Phi_{w_2}F^{\mathbf{w}}_{B_1}\Phi_{w_2}T^+_{w_2,z_2}$$

since $\Phi^2_{w_2} \simeq 0$.

Note that when we set $U = 1$, the map Ψ_z (on $\mathcal{C}FL^-$) becomes Ψ_z (on CF^-).[1] Similarly, the quasi-stabilization map $T^\pm_{w_2,z_2}$ becomes the free stabilization map $S^\pm_{z_2}$, once we set $U = 1$. Hence, after tensoring with $\mathbb{Z}_2[U,V]/(U-1)$, and using Equation (5.20), we get that this is now equivalent to

$$S^-_{z_2}(A_{\lambda_0} + A_{\gamma_2} + V\Psi_{z_2})(A_{\gamma_1} + V\Psi_{z_2})(A_{\lambda_0} + V\Psi_{z_2})S^+_{z_2},$$

as maps on CF^- (note that after tensoring with that module, $T^\pm_{w_2,z_2}$ becomes $S^\pm_{z_2}$). Note also that $\Psi_{z_2}S^+_{z_2} \simeq 0$ and $S^-_{z_2}\Psi_{z_2} \simeq 0$, as $\Psi_{z_2} \simeq S^+_{z_2}S^-_{z_2}$ and $S^-_{z_2}S^+_{z_2} \simeq 0$. Hence the above expression can be reduced to

$$S^-_{z_2}(A_{\lambda_0} + A_{\gamma_2})(A_{\gamma_1} + V\Psi_{z_2})(A_{\lambda_0})S^+_{z_2},$$

which becomes

$$S^-_{z_2}A_{\lambda_0}A_{\gamma_1}A_{\lambda_0}S^+_{z_2} + S^-_{z_2}A_{\gamma_2}V\Psi_{z_2}A_{\lambda_0}S^+_{z_2} + S^-_{z_2}A_{\lambda_0}V\Psi_{z_2}A_{\lambda_0}S^+_{z_2} + S^-_{z_2}A_{\gamma_2}A_{\gamma_1}A_{\lambda_0}S^+_{z_2}.$$

We now manipulate each term individually, using the relations of the homology actions, the relative homology maps, and the free-stabilization maps derived in Chapter 2. We compute as follows:

$$S^-_{z_2}A_{\lambda_0}A_{\gamma_1}A_{\lambda_0}S^+_{z_2} \simeq S^-_{z_2}A^2_{\lambda_0}S^-_{z_2}A_{\gamma_1} \simeq S^-_{z_2}VS^+_{z_2}A_{\gamma_1} \simeq 0;$$

[1] Using the notation from Chapter 2, we would normally write Φ_z for this map on CF^-, but we will persist in writing Ψ_z.

$$S^-_{z_2} A_{\gamma_2} V \Psi_{z_2} A_{\lambda_0} S^+_{z_2} \simeq A_{\gamma_2} V S^-_{z_2} \Psi_{z_2} A_{\lambda_0} S^+_{z_2} \simeq 0;$$

$$S^-_{z_2} A_{\lambda_0} V \Psi_{z_2} A_{\lambda_0} S^+_{z_2} \simeq V S^-_{z_2} \Psi_{z_2} A_{\lambda_0} A_{\lambda_0} S^+_{z_2} + V S^-_{z_2} A_{\lambda_0} S^+_{z_2} \simeq V;$$

and

$$S^-_{z_2} A_{\gamma_2} A_{\gamma_1} A_{\lambda_0} S^+_{z_2} \simeq A_{\gamma_2} A_{\gamma_1} S^-_{z_2} A_{\lambda_0} S^+_{z_2} \simeq A_{\gamma_2} A_{\gamma_1}.$$

Adding up the above contributions yields $V + A_{\gamma_2} A_{\gamma_1}$, completing the proof. □

Proof of Theorem 5.1.7. Push Σ so that a disk D_1 of Σ lies in Y_1 and a disk D_2 of Σ lies in Y_2. We then form the decorated link cobordism $(W, F) = (W, \Sigma \setminus (B_1 \cup B_2), \mathcal{A})$, where $\mathcal{A}$ are two parallel arcs on Σ running from $\mathbb{U}_1 = \partial B_1$ to $\mathbb{U}_2 = \partial B_2$. We set the region between the two parallel arcs to be $\Sigma_{\mathbf{w}}$, and we set the rest to be $\Sigma_{\mathbf{z}}$. We will consider the link cobordism maps $F^-_{W,F,\mathfrak{s}}$.

Using the disks D_1 and D_2 for Seifert surfaces of $\mathbb{U}_1$ and $\mathbb{U}_2$, the induced Alexander grading due to the map $F^-_{W,F,\mathfrak{s}}$ on link Floer homology is zero, since

$$\frac{\langle c_1(\mathfrak{s}), \Sigma \rangle - [\Sigma] \cdot [\Sigma]}{2} + g(\Sigma) = 0,$$

by assumption. We pick diagrams for Y_i so that the basepoints w and z are in the same region on the Heegaard surface, and the Seifert disks D_i intersect the Heegaard surface in an arc between the two basepoints, which is also contained in the same region as the basepoints. We note that $\mathcal{C}FL^-(Y_i, \mathbb{U}_i, \mathfrak{s}|_{Y_i})$ is generated by monomials $U^m V^n \cdot \mathbf{x}$ with $m, n \geq 0$. The subcomplex in zero Alexander grading is generated by monomials with $m = n$. The fact that the Alexander grading change is zero implies that the zero graded portion of $\mathcal{C}FL^-(Y_1, \mathbb{U}_1, \mathfrak{s}|_{Y_1})$ is mapped to the zero graded portion of $\mathcal{C}FL^-(Y_2, \mathbb{U}_2, \mathfrak{s}|_{Y_1})$. Note that there are two ways to identify $\mathcal{C}FL^-(Y_i, \mathbb{U}_i, \mathfrak{s}|_{Y_i})$ with $CF^-(Y_i, \mathfrak{s}|_{Y_i})$: we can set $U = 1$ or we can set $V = 1$. Let U_0 be a formal variable, which we will use for clarity. Since the Alexander grading change is zero, the map induced by setting $U = 1$ and $V = U_0$ is equal to the map induced by setting $U = U_0$ and $V = 1$. Hence our strategy is thus to show that the $V = 1$ map can be naturally identified with $F^-_{W,\mathfrak{s}}(\cdot)$, and the $U = 1$ map can be naturally identified with $F^-_{W,\mathfrak{s}-PD[\Sigma]}(\iota_*(\xi(\Sigma)) \otimes \cdot)$.

Take a parametrized Kirby decomposition of (W, F) (see Definition 4.9.5) obtained from a Morse function f with gradient like vector field v (in the sense of Definition 4.10.7) and a collection of regular values $\mathbf{b}$. Write $\mathcal{K} = \mathcal{K}_n \circ \cdots \circ \mathcal{K}_1$ where each $\mathcal{K}_i$ is an elementary parametrized link cobordism. That is, each $\mathcal{K}_i$ either corresponds to a single 1–handle or 3–handle away from the link, a collection of 2–handles away from the link, a single critical point of the divides, or a 4–dimensional index 1 critical point which occurs along Σ. We arrange so that the index 1 critical points along Σ and critical points of the divides occur first, then afterwards come the index 1 critical points away from Σ, then the index 2 critical points, and then finally the index 3 critical points. The maps for index 1 critical points along Σ are a composition of a 1–handle map, a band map, and some diffeomorphism maps.

The 1–handle maps commute with the band maps by the triangle map computation of Theorem 2.7.8, and hence we can pull all of the 1–handles down, to write the link cobordism map as a composition of 1–handle maps, then band maps, then a single 2–handle map (adding all the 2–handles at once), finally followed by some 3–handle maps. Write $W = W_3 \cup W_2 \cup W_S \cup W_1$, where W_i is formed by adding i–handles away from the link, and W_S is diffeomorphic to $Y \times [0,1]$ with a surface formed by adding bands inside of Y, and the dividing set on W_S contains all of the critical points of the divides. Write $F_{W,F,\mathfrak{s}}|_{U=1}$ for the map $F^-_{W,F,\mathfrak{s}}$ after we set $U = 1$ and $V = U_0$ (for a formal variable U_0), and set $F_{W,F,\mathfrak{s}}|_{V=1}$ to be the map obtained by setting $V = 1$ and $U = U_0$. Write (W_i, F_i) for $(W_i, F \cap W_i)$ and write $\mathfrak{s}_i$ for $\mathfrak{s}|_{W_i}$, and define $F_S \subseteq W_S$ and $\mathfrak{s}_S$ similarly. Write Σ_i for the underlying surface of F_i.

By the composition law, we have

$$F^-_{W,F,\mathfrak{s}} = F^-_{W_3,F_3,\mathfrak{s}_3} \circ F^-_{W_2,F_2,\mathfrak{s}_2} \circ F^-_{W_S,F_S,\mathfrak{s}_S} \circ F^-_{W_1,F_1,\mathfrak{s}_1}. \tag{5.21}$$

We first note that as any $\mathbf{z}$–band map turns into a change of diagrams map once we set $V = 1$, as we use the top $\mathrm{gr}_{\mathbf{w}}$–graded intersection point for the $\mathbf{z}$–band maps, which is the same intersection point we use for the change of diagrams map once we forget about the $\mathbf{z}$–basepoints. Hence, as in Section 4.14, since $\Sigma_{\mathbf{w}}$ is just a strip, the map $F^-_{W_S,F_S,\mathfrak{s}_S}|_{V=1}$

becomes the graph cobordism map for a graph which is equal to a path with some trivial strands added. Since we can always remove the trivial strands by basic relations of the graph cobordism maps, once we set $V = 1$ the map $F^-_{W_S,F_S,\mathfrak{s}_S}$ is just the identity map. Upon direct inspection, we have that

$$F^-_{W_i,F_i,\mathfrak{s}_i}|_{V=1} = F^-_{W_i,\mathfrak{s}_i}$$

for the handle attachment cobordisms (W_i, F_i). Using the composition law again, we see that $F^-_{W,F,\mathfrak{s}}|_{V=1} = F^-_{W,\mathfrak{s}}$.

We now consider the map $F^-_{W,F,\mathfrak{s}}|_{U=1}$. Note that the 1–handle maps and 3–handle maps in the composition for $F^-_{W,F,\mathfrak{s}}|_{U=1}$ agree with the maps for $F^-_{W,\mathfrak{s}-PD[\Sigma]}$. Similarly, noting that $(\mathfrak{s} - PD[\Sigma])|_{W_S} = \mathfrak{s}_S$, we have

$$F^-_{W_S,F_S,(\mathfrak{s}-PD[\Sigma])|_{W_S}}|_{U=1} = F^-_{W_S,F_S,\mathfrak{s}_S}|_{U=1} = \iota_* \prod_{i=1}^{g(\Sigma)} (V + A_i \cdot B_i),$$

by Lemma 5.13.4 and the composition law.

We finally note that the triangle maps used in the construction of the link cobordism maps in Chapter 4 are not entirely symmetric between the $\mathbf{w}$– and the $\mathbf{z}$–basepoints. The triangle maps involved in 2–handle maps count triangles which satisfy $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$. The band maps also count holomorphic triangles, but since all of the links resulting from taking $\mathbb{U}_1$ and adding bands are null-homologous (in the sense that the total class is null-homologous), the maps $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ appearing in the band maps coincide. For the 2–handle maps, the maps $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ may be different. The map $F_{W_2,\mathfrak{s}_2}$ counts triangles with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}_2$, but by Lemma 5.3.3, these are exactly the same triangles which satisfy $\mathfrak{s}_{\mathbf{z}}(\psi) = \mathfrak{s}_2 - PD[\Sigma_2]$. These are the triangles counted by $F_{W_2,\mathfrak{s}_2-PD[F_2]}$. Hence we conclude that $F^-_{W,F,\mathfrak{s}}|_{U=1}$ is equal to

$$F^-_{W,F,\mathfrak{s}}|_{U=1} = F^-_{W_3,\mathfrak{s}_3} \circ F^-_{W_2,\mathfrak{s}_2-PD[\Sigma_2]} \circ [\iota_*(\xi(\Sigma))] \circ F^-_{W_1,\mathfrak{s}_1}.$$

Hence, using the composition law, we see that $F^-_{W,F,\mathfrak{s}}|_{U=1}$ is the map

$$F^-_{W,\mathfrak{s}-PD[\Sigma]}(\iota_*(\xi(\Sigma)) \otimes \cdot),$$

completing the proof. □

More generally, one could put more exotic sets of divides on Σ, such as those considered in Theorem 5.1.8, and would presumably recover relations of the Heegaard Floer cobordism maps analogous to those from [OS00a] for the Seiberg-Witten invariant.

We now use the adjunction relation above to give a link Floer homology proof of the standard adjunction inequality, proven by Ozsváth and Szabó in [OS04d].

Corollary 5.13.5. *If Σ is a closed, oriented surface in Y with $g(\Sigma) > 0$ and $HF^+(Y, \mathfrak{s}) \neq 0$, then*

$$|\langle c_1(\mathfrak{s}), \Sigma\rangle| \leq 2g(\Sigma) - 2.$$

Proof. We apply the previous adjunction relation to the identity cobordism $W = Y \times [0, 1]$. If Σ is a closed, oriented surface in Y which violates the inequality, we reverse the orientation of Σ if necessary, and add null-homologous handles so that

$$\langle c_1(\mathfrak{s}), \Sigma\rangle = -2g(\Sigma) \tag{5.22}$$

(note that $\langle c_1(\mathfrak{s}), \Sigma\rangle$ is always even). Applying the previous theorem to the identity cobordism $Y \times [0, 1]$, we see that

$$\mathrm{id} = [\iota_*(\xi(\Sigma))] \cdot F^+_{Y\times[0,1],\mathfrak{s}-PD[\Sigma]}.$$

As $[\Sigma] = 0 \in H_2(W, \partial W; \mathbb{Z}) \cong H^2(W; \mathbb{Z})$, we have that $F^+_{W,\mathfrak{s}-PD[\Sigma]} = F^+_{W,\mathfrak{s}} = \mathrm{id}$. Hence

$$\mathrm{id} = [\iota_*(\xi(\Sigma))].$$

On the other hand, $\mathfrak{s}$ must be non-torsion for Equation (5.22) to be satisfied. As a consequence, we have that $U^n \cdot HF^+(Y, \mathfrak{s}) = 0$ for sufficiently large n, by [OS04b, Lemma 2.3]. Similarly $[\gamma] \cdot [\gamma] = 0$ for any $[\gamma] \in H_1(Y; \mathbb{Z})$. As such, if n is sufficiently large, the action of $[\iota_*(\xi(\Sigma))]^n$ will be zero on $HF^+(Y, \mathfrak{s})$. Hence $\mathrm{id} = \mathrm{id}^n = [\iota_*(\xi(\Sigma))]^n = 0$ as maps on $HF^+(Y, \mathfrak{s})$, so $HF^+(Y, \mathfrak{s})$ must vanish. □

As in [OS06, Theorem 1.5], the adjunction inequality for the Ozsváth–Szabó mixed invariant for surfaces with nonnegative self intersection inside of a 4–manifold follows from

the above adjunction inequality by using the blowup formula to reduce to the case that $[\Sigma] \cdot [\Sigma] = 0$, and factoring the cobordism map on HF^+ through a regular neighborhood of Σ.

5.14 The maps on $\mathcal{H}FL^\infty$ for surfaces in negative definite 4–manifolds

In this section we compute the maps associated to knot cobordisms on homology when the 4–manifold is negative definite and the dividing set is relatively simple.

Theorem 5.1.8. *Suppose that $(W, F) : (S^3, \mathbb{K}_1) \to (S^3, \mathbb{K}_2)$ is a knot cobordism such that $b_1(W) = b_2^+(W) = 0$, $F = (\Sigma, \mathcal{A})$ is a connected surface, and $\mathbb{K}_i$ are two null-homologous knots, each with two basepoints. Suppose further that $\mathcal{A}$ consists of two arcs going from K_1 to K_2 and $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ are both connected. Then the induced map on homology*

$$F^\infty_{W,F,\mathfrak{s}} : \mathcal{H}FL^\infty(S^3, \mathbb{K}_1) \to \mathcal{H}FL^\infty(S^3, \mathbb{K}_2)$$

is an isomorphism. In fact, under the identification $\mathcal{H}FL^\infty(S^3, \mathbb{K}_i) \cong \mathbb{Z}_2[U, V, U^{-1}, V^{-1}]$, it is the map

$$1 \mapsto U^{-d_1/2} V^{-d_2/2},$$

where

$$d_1 = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} - 2g(\Sigma_{\mathbf{w}})$$

and

$$d_2 = \frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4} - 2g(\Sigma_{\mathbf{z}}).$$

Proof. Our absolute gradings give a canonical identification of $\mathcal{H}FL^\infty(S^3, \mathbb{K}_i)$ with $\mathbb{Z}_2[U, V, U^{-1}, V^{-1}]$. Using our grading formulas, the map $F^\infty_{W,F,\mathfrak{s}} : \mathbb{Z}_2[U, V, U^{-1}, V^{-1}] \to \mathbb{Z}_2[U, V, U^{-1}, V^{-1}]$ is equal to $1 \mapsto c \cdot U^{-d_1/2} V^{-d_2/2}$ for some $c \in \mathbb{Z}_2$. We just need to argue that $c = 1$. For this, we note that when we set $V = 1$, the map goes from $CF^\infty(S^3)$ to $CF^\infty(S^3)$. We can apply Lemma 5.13.4 repeatedly, as well as the composition law, as

in the proof of Theorem 5.1.7 to see that the map induced by $F^{\infty}_{W,F,\mathfrak{s}}$ on $CF^{\infty}(S^3)$ is equal to $F^{\infty}_{W,\mathfrak{s}}(\iota_*(\xi(\Sigma_{\mathbf{w}}))\otimes\cdot)$, where $\xi(\Sigma_{\mathbf{w}})$ is defined by picking an appropriate symplectic basis $A_1, B_1, \dots, A_{g(\Sigma_{\mathbf{w}})}, B_{g(\Sigma_{\mathbf{w}})}$ of $H_1(\Sigma_{\mathbf{w}};\mathbb{Z})$ and defining $\xi(\Sigma_{\mathbf{w}}) = \prod_{i=1}^{g(\Sigma_{\mathbf{w}})}(U + A_i\cdot B_i)$. However we assumed that $b_1(W) = 0$, so the actions of A_i and B_i vanish, and we see that $F^{\infty}_{W,F,\mathfrak{s}}$ simply induces the map $U^{g(\Sigma_{\mathbf{w}})}\cdot F^{\infty}_{W,\mathfrak{s}}$. The map $F^{\infty}_{W,\mathfrak{s}}$ is an isomorphism once again by the argument in the proof of [OS03b, Theorem 9.1]. In particular, we conclude that $c = 1$ so $F^{\infty}_{W,F,\mathfrak{s}}$ is an isomorphism on homology. □

Remark 5.14.1. We can use the above strategy to compute the maps on CF^- induced by the knot cobordism maps for more general knot cobordisms. As we've seen, there are two ways to obtain a map on CF^- from the link cobordism maps $F^{\infty}_{W,F,\mathfrak{s}}$. The first is obtained by setting $V = 1$ and only paying attention to the $\mathbf{w}$–basepoints. The second is obtained by setting $U = 1$, and paying attention to only the $\mathbf{z}$–basepoints. If $(W, F) : (Y_1, \mathbb{K}_1) \to (Y_2, \mathbb{K}_2)$ is a knot cobordism with $\mathbb{K}_i$ two knots each with exactly two basepoints, with $F = (\Sigma, \mathcal{A})$ a connected surface with divides, such that $\mathcal{A}$ consists of exactly two arcs from $\mathbb{K}_1$ to $\mathbb{K}_2$ which divide Σ into two connected subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ meeting along $\mathcal{A}$, a straightforward adaptation of the proofs of Theorem 5.1.7 and Theorem 5.1.8 shows that the $V = 1$ map is exactly $F^{\infty}_{W,\mathfrak{s}}(\iota_*(\xi(\Sigma_{\mathbf{w}}))\otimes\cdot)$ and the $U = 1$ map is exactly $F^{\infty}_{W,\mathfrak{s}-PD[\Sigma]}(\iota_*(\xi(\Sigma_{\mathbf{z}}))\otimes\cdot)$.

Corollary 5.1.9. *Suppose that $\Sigma \subseteq S^4$ is a closed, oriented surface, and $\mathcal{A}$ is a simple closed curve on Σ which divides Σ into two connected subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, then the link cobordism map*

$$F^-_{S^4,F,\mathfrak{s}_0} : \mathcal{C}FL^-(\varnothing,\varnothing) \to \mathcal{C}FL^-(\varnothing,\varnothing)$$

is the map

$$1 \mapsto U^{g(\Sigma_{\mathbf{w}})}V^{g(\Sigma_{\mathbf{z}})}.$$

In particular the maps for 2–knots in S^4 (or any homotopy S^4) are the identity.

We can apply this to compute the the effect on the link cobordism maps of taking the internal connected sum of a surface Σ with a closed surface Σ_0 in a 4–ball not intersecting Σ.

Corollary 5.14.2. *Suppose that $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a link cobordism with $F = (\Sigma, \mathcal{A})$ and suppose there is a 4–ball $B \subseteq W$ which doesn't intersect Σ, and that $\Sigma_0 \subseteq B$ is a closed surface. Let $F' = (\Sigma \# \Sigma_0, \mathcal{A})$ denote the internal connected sum, taken at a point in $\Sigma_{\mathbf{z}} \subseteq \Sigma$, and the entire surface Σ_0 is given the designation of type $\mathbf{z}$. If t denotes the color of Σ_0, then*

$$F^-_{W,F',\mathfrak{s}} = V_t^{g(\Sigma_0)} \cdot F^-_{W,F,\mathfrak{s}}.$$

Proof. The key idea is that the cobordism maps from Chapter 4 allow us to puncture the 4–manifold and surface, creating a new end, which we can push the surface Σ_0 into.

Let B' be a 4–ball containing Σ_0 and intersecting Σ in a disk along an arc of $\mathcal{A}$, and assume that the internal connected sum is taken inside of the ball B'. Define $\mathcal{B}_1 = (B', F_1)$ to be the link cobordism from $\varnothing$ to $(S^3, \mathbb{U})$ (here $\mathbb{U}$ is the unknot in S^3 with exactly two basepoints), formed by setting $F_1 = (\Sigma_0 \# \Sigma) \cap B'$. Define $F_2 = \Sigma \cap B' \subseteq B'$, and let $\mathcal{B}_2 = (B', F_2)$. Let

$$(W_0, F_0) = (W \setminus B', F \setminus B'),$$

viewed as a cobordism from $(Y_1 \sqcup S^3, \mathbb{L}_1 \sqcup \mathbb{U})$ to $(Y_2, \mathbb{L}_2)$. Applying Corollary 5.1.9 and Lemma 4.12.4, we have

$$\begin{aligned} F^-_{W,F',\mathfrak{s}} &= F^-_{W_0,F_0,\mathfrak{s}_0} \circ F^-_{\mathcal{B}_1,\mathfrak{s}_0} \\ &= F^-_{W_0,F_0,\mathfrak{s}_0} \circ (V_t^{g(\Sigma_0)} \cdot F^-_{\mathcal{B}_2,\mathfrak{s}_0}) \\ &= V_t^{g(\Sigma_0)} \cdot F^-_{W_0,F_0,\mathfrak{s}_0} \circ F^-_{\mathcal{B}_2,\mathfrak{s}_0} \\ &= V_t^{g(\Sigma_0)} \cdot F^-_{W,F,\mathfrak{s}}. \end{aligned}$$

□

A natural question in the spirit of the above computations is the following:

Question 5.14.3. Can the link cobordism maps differentiate two slice disks of a given knot?

Using the previous corollary, we see that if D_1 and D_2 are two slice disks for a knot $\mathbb{K}$ such that there are 2–knots S_1 and S_2, such that S_i is contained in a 4–ball not intersecting D_i

such that $D_1 \# S_1$ and $D_2 \# S_2$ are isotopic, then the maps from $\mathbb{Z}_2[U,V] = \mathcal{C}FL^-(S^3, U, w, z)$ to $\mathcal{C}FL^-(S^3, \mathbb{K})$ associated to D_1 and D_2 agree.

CHAPTER 6

Connected sums and involutive knot Floer homology

We analyze the interaction between the conjugation action on link Floer homology and maps induced by link cobordisms. We derive a formula for the chain homotopy type of the conjugation action on the connected sum of two knots. We then construct a homomorphism from the concordance group to an abelian group consisting of chain complexes with homotopy automorphisms, modulo an equivalence relation. Using our connected sum formula, we perform some example computations of Hendricks and Manolescu's involutive invariants on large surgeries of connected sums of knots.

6.1 Introduction

Heegaard Floer homology is an invariant associated to 3–manifolds, introduced by Ozsváth and Szabó in [OS04e] and [OS04d]. By taking into account a conjugation involution, ι, Hendricks and Manolescu construct a refinement, called involutive Heegaard Floer homology [HM15]. They define a module $HFI^-(Y,\mathfrak{s})$ over $\mathbb{F}_2[U,Q]/(Q^2)$ as the homology of the mapping cone

$$CFI^-(Y,\mathfrak{s}) := \text{Cone}\left(CF^-(Y,\mathfrak{s}) \xrightarrow{Q\cdot(\text{id}+\iota)} Q\cdot CF^-(Y,\mathfrak{s})\right),$$

where Q is a formal variable. They define two correction terms, $\overline{d}$ and $\underline{d}$, as analogs of the Ozsváth-Szabó d invariant from Heegaard Floer homology.

An important property of Heegaard Floer homology is the Künneth theorem for connected sums. In [OS04d], if Y_1 and Y_2 are two connected 3–manifolds, Ozsváth and Szabó construct

a quasi-isomorphism

$$CF^-(Y_1\#Y_2,\mathfrak{s}_1\#\mathfrak{s}_2)\cong CF^-(Y_1,\mathfrak{s}_1)\otimes_{\mathbb{F}_2[U]}CF^-(Y_2,\mathfrak{s}_2).$$

In [HMZ16], the behavior of the conjugation involution on connected sums is described in terms of the above quasi-isomorphism, giving a version of the Künneth theorem for involutive Heegaard Floer homology.

Knot Floer homology is an invariant of knots in 3–manifolds introduced by Ozsváth and Szabó in [OS04c], and independently by Rasmussen in [Ras03]. Link Floer homology is a generalization for links in 3–manifolds, developed by Ozsváth and Szabó in [OS08]. There is a similar conjugation action on knot Floer homology, and also a version of the Künneth theorem for knot Floer homology.

In this paper we consider the conjugation action on knot Floer homology, its form on connected sums of knots, and more generally the interaction between the conjugation action and maps induced by decorated link cobordisms.

To a knot $\mathbb{K}=(K,w,z)$ with two basepoints in a 3–manifold Y equipped with a Spin^c structure $\mathfrak{s}\in\text{Spin}^c(Y)$, we consider a chain complex

$$\mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s}),$$

over the ring

$$\mathcal{R}=\mathbb{F}_2[U,V,U^{-1},V^{-1}],$$

the Laurent polynomial ring generated by two variables. This is a slight variation of the standard $CFK^\infty(Y,K)$ complex from [OS04c]. The chain complex $\mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s})$ has a filtration over $\mathbb{Z}\oplus\mathbb{Z}$. If K is null-homologous and $\mathfrak{s}$ is a self conjugate Spin^c structure, Hendricks and Manolescu consider a conjugation map

$$\iota_K:\mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s})\to\mathcal{C}FL^\infty(Y,\mathbb{K},\mathfrak{s}).$$

Unlike the conjugation action ι on $CF^-(Y,\mathfrak{s})$, the map ι_K is not a homotopy involution. Instead, the map ι_K satisfies

$$\iota_K^2\simeq\rho_*$$

where ρ_* is the map induced by the diffeomorphism $\rho : (Y, K, w, z) \to (Y, K, w, z)$ obtained by twisting the knot K in one full twist, in the direction of its orientation. The map ρ_* is considered in [Sar15] and Chapter 3. Over $\mathbb{F}_2$, however, one has $\iota_K^4 \simeq \mathrm{id}$.

6.1.1 The conjugation map on connected sums of knots

Our first result is the chain homotopy type of the map ι_K on connected sums of knots:

Theorem 6.1.1. *Suppose that $(Y_1, \mathbb{K}_1)$ and $(Y_2, \mathbb{K}_2)$ are two pairs of 3–manifolds with embedded, doubly based, null-homologous knots, and $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are self-conjugate* Spin^c *structures on Y_1 and Y_2. Writing $\mathbb{K}_1 \# \mathbb{K}_2$ for the connected sum, with exactly two basepoints, there are filtered, $\mathcal{R}$–equivariant chain homotopy equivalences between*

$$\mathcal{C}FL^\infty(Y_1, \mathbb{K}_1, \mathfrak{s}_1) \otimes_{\mathcal{R}} \mathcal{C}FL^\infty(Y_2, \mathbb{K}_2, \mathfrak{s}_2) \qquad \text{and} \qquad \mathcal{C}FL^\infty(Y_1 \# Y_2, \mathbb{K}_1 \# \mathbb{K}_2, \mathfrak{s}_1 \# \mathfrak{s}_2)$$

which intertwine $\iota_{K_1 \# K_2}$, on the latter chain complex, and

$$(\mathrm{id} \,|\, \mathrm{id} + \Phi_1 | \Psi_2)(\iota_{K_1} | \iota_{K_2})$$

on the tensor product complex, for endomorphisms Φ_i and Ψ_i of $\mathcal{C}FL^\infty(Y_i, \mathbb{K}_i, \mathfrak{s}_i)$. There are also (different) chain homotopy equivalences of the two above complexes which intertwine $\iota_{K_1 \# K_2}$ with

$$(\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2)(\iota_{K_1} | \iota_{K_2}).$$

In the above theorem, the vertical bar, |, denotes the tensor product of maps. The maps Φ_i and Ψ_i appeared in Chapters 3 and 4, however they appeared earlier in [Sar15], in the context of basepoint moving maps. Using our formulation, they can be thought of as the formal derivatives of the differential with respect to U or V, respectively. Notably, the maps Φ_i and Ψ_i are the link cobordism maps for the decorated link cobordisms shown in Figure 6.1 (see Definition 6.1.2). This interpretation of the maps Φ_i and Ψ_i in terms of decorated link cobordisms turns out to be useful in our proof of Theorem 6.1.1. See Section 6.4.2 for an explicit definition of the maps Φ and Ψ in terms of holomorphic disks on the link Floer complexes.

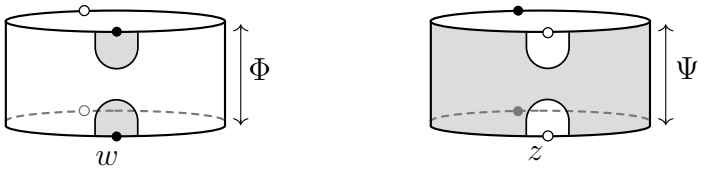

Figure 6.1: **Decorated link cobordisms for Φ and Ψ.** The underlying undecorated link cobordism is $K \times [0,1] \subseteq Y \times [0,1]$.

6.1.2 Conjugation invariance of the link Floer TQFT

To arrive at Theorem 6.1.1, we will study the interaction between the link Floer TQFT and the conjugation action more generally. If $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is an oriented link in Y with collections of basepoints $\mathbf{w}$ and $\mathbf{z}$ in L which alternate between $\mathbf{w}$ and $\mathbf{z}$ as one travels along L, we consider a version of the full link Floer complex, which we denote by

$$\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s}).$$

If $\mathbf{w} = \{w_1, \dots, w_n\}$ and $\mathbf{z} = \{z_1, \dots, z_n\}$, these are modules over the ring

$$\mathbb{F}_2[U_{\mathbf{w}}, V_{\mathbf{z}}] := \mathbb{F}_2[U_{w_1}, \dots, U_{w_n}, V_{z_1}, \dots, V_{z_n}].$$

For a general link $\mathbb{L}$, the module $\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$ has an endomorphism ∂, which satisfies

$$\partial^2 = k \cdot \mathrm{id},$$

for a constant $k \in \mathbb{F}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$. The constant k is determined by the configuration of basepoints on the link components (see Equation (6.8)). In many natural situations, the constant k can be taken to be zero. For example if we identify all $U_{\mathbf{w}}$ variables with a single variable U, and we identify all $V_{\mathbf{z}}$ variables with a single variable V, the constant k will vanish. This situation is sufficient for proving Theorem 6.1.1.

There is a tautological conjugation map η on link Floer homology. It is a map

$$\eta : \mathcal{C}FL^\infty(Y, L, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\infty(Y, L, \mathbf{z}, \mathbf{w}, \overline{\mathfrak{s}} + PD[L]).$$

It's important to note that on link Floer homology, the basepoints $\mathbf{w}$ and $\mathbf{z}$ play distinct roles. Hence the map η canonically squares to the identity, but is not an involution since it is a map between the link Floer homologies of two different based links.

When $\overline{\mathfrak{s}} = \mathfrak{s}$ and $\mathbb{K} = (K, w, z)$ is a doubly based knot which is null-homologous, the maps η and ι_K are related by the formula

$$\iota_K := \tau_K \circ \eta,$$

where τ_K is the diffeomorphism associated to a half twist of K, in the direction of the orientation of K.

The chain homotopy equivalences appearing in Theorem 6.1.1 are link cobordism maps for certain decorated link cobordisms, using the construction of link cobordism maps from Chapter 4. The maps from Chapter 4 are expected to generalize the link cobordism maps from [Juh16], which are defined on $\widehat{HFL}$. The first step towards proving Theorem 6.1.1, is to analyze the interaction of the link cobordism maps from Chapter 4 and the conjugation map η. The maps from Chapter 4 use the following notion of a decorated link cobordisms (adapted from [Juh16, Definition 4.5]):

Definition 6.1.2. We say a pair (W, F) is a **decorated link cobordism** between two 3–manifolds with multi-based links, and write $(W, F) : (Y_1, L_1, \mathbf{w}_1, \mathbf{z}_1) \to (Y_2, L_2, \mathbf{w}_2, \mathbf{z}_2)$, if

1. W is a 4–dimensional cobordism from Y_1 to Y_2;
2. $F = (\Sigma, \mathcal{A})$, where Σ is an oriented surface, properly embedded in W, and $\mathcal{A} \subseteq \Sigma$ is a properly embedded 1–manifold dividing Σ into two disjoint subsurfaces $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, which meet along $\mathcal{A}$;
3. $\partial\Sigma = -L_1 \sqcup L_2$
4. Each component of $L_i \setminus \mathcal{A}$ contains exactly one basepoint;
5. $\mathbf{w}_1 \cup \mathbf{w}_2 \subseteq \Sigma_{\mathbf{w}}$ and $\mathbf{z}_1 \cup \mathbf{z}_2 \subseteq \Sigma_{\mathbf{z}}$.

Additionally, to define functorial link cobordism maps, we need to work with a set of formal variables and assign a formal variable to each basepoint in $\mathbf{w} \cup \mathbf{z}$ and each region of $\Sigma \setminus \mathcal{A}$. Such an assignment of formal variables is called a **coloring**. See Section 6.3 for more

details about colorings. For the purposes of proving Theorem 6.1.1, we can pick the coloring to assign the variable U to all components of $\Sigma_{\mathbf{w}}$ and each $\mathbf{w}$ basepoint, and to assign the variable V to all components of $\Sigma_{\mathbf{z}}$ and all $\mathbf{z}$ basepoints.

We prove the following version of conjugation invariance for the link Floer cobordism maps:

Theorem 6.1.3. *Suppose that $(W,F) : (Y_1, L_1, \mathbf{w}_1, \mathbf{z}_1) \to (Y_2, L_2, \mathbf{w}_2, \mathbf{z}_2)$ is a decorated link cobordism with a coloring of the components of $\Sigma \setminus \mathcal{A}$. Let $(W, \overline{F}) : (Y_1, L_1, \mathbf{z}_1, \mathbf{w}_1) \to (Y_2, L_2, \mathbf{z}_2, \mathbf{w}_2)$ denote the cobordism obtained by switching the roles of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ (but using the same coloring). Then the following diagram commutes up to filtered, equivariant chain homotopy:*

$$\begin{array}{ccc} \mathcal{C}FL^\infty(Y_1, L_1, \mathbf{w}_1, \mathbf{z}_1, \mathfrak{s}_1) & \xrightarrow{\eta} & \mathcal{C}FL^\infty(Y_1, L_1, \mathbf{z}_1, \mathbf{w}_1, \overline{\mathfrak{s}}_1 + PD[L_1]) \\ \downarrow{\scriptstyle F_{W,F,\mathfrak{s}}} & & \downarrow{\scriptstyle F_{W,\overline{F},\overline{\mathfrak{s}}+PD[\Sigma]}} \\ \mathcal{C}FL^\infty(Y_2, L_2, \mathbf{w}_2, \mathbf{z}_2, \mathfrak{s}_2) & \xrightarrow{\eta} & \mathcal{C}FL^\infty(Y_2, L_2, \mathbf{z}_2, \mathbf{w}_2, \overline{\mathfrak{s}}_2 + PD[L_2]). \end{array}$$

6.1.3 A bypass relation for link Floer homology

Theorem 6.1.1 does not follow immediately from Theorem 6.1.3, as the maps ι_K also involve a diffeomorphism corresponding to a half twist of the knot K. However, $\iota_{K_1\#K_2}$ involves a half twist on $K_1\#K_2$, while $\iota_{K_1}|\iota_{K_2}$ involves half twists on each of K_1 and K_2. To relate these two, we will find a general relation satisfied by triples of dividing sets on a surface which fit into a "bypass triple".

The link cobordism maps constructed by Juhász in [Juh16] for the hat flavor of link Floer homology are defined by using the contact gluing map from [HKM08], for a contact structure on the boundary of a regular neighborhood of Σ which is determined by the dividing set on Σ. There is relation on the contact invariants in the sutured Floer homology of $S^1 \times \Sigma$ for a triple of contact structures fitting into a bypass triple (see [HKM08, Section 7]). Inspired by this, we show the following relation on the link cobordism maps from Chapter 4:

Lemma 6.1.4. *Let (W,F) (W,F') and (W,F'') be decorated link cobordisms from $(Y_1, \mathbb{L}_1)$*

to $(Y_2, \mathbb{L}_2)$*, which form a bypass triple, as in Figure 6.2. Then*

$$F_{W,F,\mathfrak{s}} + F_{W,F',\mathfrak{s}} + F_{W,F'',\mathfrak{s}} \simeq 0.$$

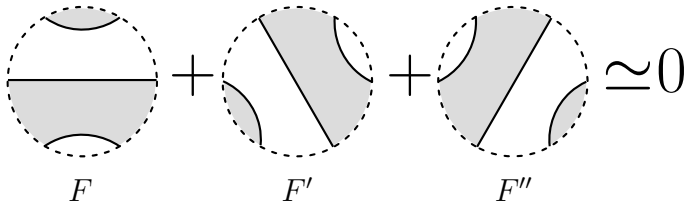

Figure 6.2: **The bypass relation** for decorated link cobordisms $(W, F), (W, F')$ (W, F'') which agree outside of the region shown.

Our proof of the above lemma involves no contact geometry, and instead follows from several simple relations between maps defined on the link Floer complexes, and their interpretation in terms of decorated surfaces.

Using the bypass relation, in Section 6.4.3 we give an entirely pictorial proof of Sarkar's formula that

$$\rho_* \simeq \mathrm{id} + \Phi_w \circ \Psi_z,$$

where ρ_* is the diffeomorphism map for twisting a link component in one full twist, in the case that the link component has exactly two basepoints, w and z. This was proven on the associated graded complex for links in S^3 by Sarkar in [Sar15], and on the full link Floer complex for links in arbitrary 3–manifolds by the author in Chapter 3.

6.1.4 A homomorphism from the smooth concordance group

Using Theorem 6.1.1 we construct a homomorphism from the smooth concordance group $\mathcal{C}$ to an algebraically defined, abelian group $\mathfrak{I}_K$. The construction is inspired by constructions of Hom in [Hom14] and Stoffregen in [Sto15]. See also [HMZ16] for a similar construction.

In Section 6.2, we will define a notion of ι_K–complexes, which is a tuple (C, ∂, B, ι), where (C, ∂) is a free chain complex over $\mathcal{R}$ with basis B, and ι is an $\mathcal{R}$–skew equivariant (i.e. it switches the actions of U and V), $\mathbb{Z} \oplus \mathbb{Z}$–skew filtered map, which squares to $\mathrm{id} + \Phi \circ \Psi$.

We define an abelian group $\mathfrak{I}_K$, consisting of ι_K complexes modulo a relation we call local equivalence (see Section 6.2.3 for a precise definition). To a smooth concordance between two knots (or more generally a homology concordance; see below) the link cobordism maps from Chapter 4 will induce local equivalence. We prove the following:

Theorem 6.1.5. *The natural map $\mathcal{C} \to \mathfrak{I}_K$ defined by*

$$K \mapsto [\mathcal{C}FL^\infty(S^3, K)] \in \mathfrak{I}_K.$$

is a well defined homomorphism.

More generally, we can extend the above theorem to knots in integer homology spheres. One can define a group $\mathcal{C}_3^{\mathbb{Z}}$ generated by pairs (Y, K) where Y is an integer homology sphere and $K \subseteq Y$ is an oriented knot. We identify two pairs if there is an integer homology cobordism between the 3–manifolds, which contains a smooth, oriented, genus zero knot cobordism between the two knots. Multiplication is given by taking the connected sum of both the 3–manifolds and the knots. It is not hard to see that $\mathcal{C}_3^{\mathbb{Z}}$ is an abelian group, where the inverse of (Y, K) is given by $(-Y, -K)$.

Remark 6.1.6. The homomorphism $\mathcal{C} \to \mathfrak{I}_K$ from Theorem 6.1.5 factors through a homomorphism $\mathcal{C}_3^{\mathbb{Z}} \to \mathfrak{I}_K$.

6.1.5 Example computations

In Section 6.8, we use the above formula to compute the involutive correction terms $\overline{V}_0$ and $\underline{V}_0$ from large surgeries from [HM15] for several examples of connected sums of knots. The following is a straightforward consequence of several example computations which we perform:

Proposition 6.1.7. *None of the knots $T_{4,5}\#T_{4,5}$, $T_{4,5}\#T_{4,5}\#T_{5,6}$, $T_{6,7}\#T_{6,7}$, $T_{4,5}\#T_{6,7}$, and $T_{3,4}^{-1}\#T_{4,5}^{-1}\#T_{5,6}$ are concordant to a thin knot, an L-space knot, or the mirror of an L-space knot.*

The proof of the above proposition uses the computation of the triple $(\overline{V}_0, V_0, \underline{V}_0)$ for thin knots, L-space knots, and mirrors of L-space knots from [HM15]. For such knots, the involutive correction terms have a simple pattern: either all three are nonnegative and $\underline{V}_0 - \overline{V}_0 \leq 1$, or $\overline{V}_0 \leq 0 = V_0 = \underline{V}_0$. Using our formula we will compute that none of the knots listed above have involutive invariants fitting into one of these two patterns. See [BH16] for some deeper applications of involutive large surgery invariants as well as some applications of the connected sum formula from this paper.

Note that looking at the triple $(\overline{V}_0, V_0, \underline{V}_0)$ is not the only way to detect various parts of the above result. For example $\Upsilon_K(t)$ is very simple for thin knots, so can be used as an obstruction to knots being concordant to thin knots. Similarly $V_0 = 0$ for any mirror of an L-space knot, and $V_0 \neq 0$ for the above knots. The function $\Upsilon_K(t)$ is convex for an L-space knot [OSS14, Theorem 6.2], which can provide an obstruction to knots being concordant to L-space knots, however except for $T_{3,4}^{-1}\#T_{4,5}^{-1}\#T_{5,6}$, the above knots have convex $\Upsilon_K(t)$ invariants.

6.1.6 Further Remarks

In [HMZ16, Theorem 1.1], a similar formula for the involution ι on CF^- of connected sums of closed 3–manifolds is proven, by considering the interaction between the conjugation action and the graph cobordism maps from Chapter 2. Writing ι_i for the involution on $CF^-(Y_i, \mathfrak{s}_i)$, we showed that the involution on $CF^-(Y_1\#Y_2, \mathfrak{s}_1\#\mathfrak{s}_2)$ is $\mathbb{F}_2[U]$–equivariantly chain homotopic to

$$(\mathrm{id} \,|\, \mathrm{id} + U(\Phi_1|\Phi_2))(\iota_1|\iota_2)$$

under a chain homotopy equivalence with $CF^-(Y_1, \mathfrak{s}_1) \otimes_{\mathbb{F}_2[U]} CF^-(Y_2, \mathfrak{s}_2)$. In the context of the 3–manifold invariants, the maps Φ_i are analogous to the maps Φ_i and Ψ_i appearing in Theorem 6.1.1. By specializing Theorem 6.1.1 of this paper to the case $\mathbb{K}_1$ and $\mathbb{K}_2$ are doubly based unknots, one recovers the formula from [HMZ16]. In the setting of 3–manifold invariants, the term $U(\Phi_1|\Phi_2)$ turns out to be $\mathbb{F}_2[U]$–equivariantly null-homotopic when each

$\mathfrak{s}_i$ is self conjugate, by an algebraic computation. As we will see in Section 6.8, the term $\Phi_1|\Psi_2$ is usually non-trivial in the context of link Floer homology.

More generally, using the interpretation of the graph cobordism maps in terms of decorated link cobordisms between unknots from Section 4.14, it is not hard to prove the relations from [HMZ16] about the interaction of the graph TQFT and the conjugation involution by using Theorem 6.1.3 form this paper, which is about link cobordisms.

For a doubly based knot $\mathbb{K} = (K, w, z)$ in S^3, one more often considers a $\mathbb{Z} \oplus \mathbb{Z}$ filtered chain complex $CFK^\infty(K)$ over $\mathbb{F}_2[U]$. The complexes $\mathcal{C}FL^\infty$ are graded by the Alexander grading and

$$CFK^\infty(K) = \mathcal{C}FL^\infty(S^3, \mathbb{K})_0,$$

where $\mathcal{C}FL^\infty(S^3, \mathbb{K})_0 \subseteq \mathcal{C}FL^\infty(S^3, \mathbb{K})$ is the subset in zero Alexander grading. The standard action of U on $CFK^\infty(K)$ corresponds to the action of the product UV on $\mathcal{C}FL^\infty(S^3, \mathbb{K})$, in the notation of this paper. As the map ι_K satisfies $A(\iota_K(\mathbf{x})) = -A(\mathbf{x})$, the map ι_K preserves the subset of zero Alexander grading, so Theorem 6.1.1 can be restated for the more standard CFK^∞ complexes. Nonetheless it's worth noting that Φ_i and Ψ_i are not individually endomorphisms of $CFK^\infty(K_i)$ (as they shift Alexander grading by $+1$ and -1 respectively), but the tensor product $\Phi_1|\Psi_2$ shifts Alexander grading by zero, and hence is an endomorphism of the tensor product of the two CFK^∞ complexes. For the purposes of this paper, it is easier to use the complexes $\mathcal{C}FL^\infty$ than the CFK^∞ complexes, since it's convenient to work with the maps Φ_i and Ψ_i individually.

In a different direction, due to the work of Kutluhan, Lee and Taubes in [KLT10a, KLT10b, KLT10c, KLT11, KLT12], as well as Colin, Ghiggini and Honda in [CGH10, CGH12a, CGH12b, CGH12c] there are isomorphisms between certain versions of Heegaard Floer homology and the monopole Floer homology groups developed by Kronheimer and Mrowka in [KM07]. Monopole Floer homology is an S^1 equivariant theory, however the S^1 action can be enlarged to a Pin(2) action, and constructions which take the full Pin(2) action into account can be found in [Man16] and [Lin14]. The modules $HFI^-(Y, \mathfrak{s})$ from [HM15] are expected to correspond to the $\mathbb{Z}_4$–equivariant theory for the

subgroup $\mathbb{Z}_4 = \langle j \rangle \subseteq \mathrm{Pin}(2)$. In [Lin16] Lin proves a connected sum formula for Pin(2)–equivariant monopole Floer homology, in terms of an $\mathcal{A}_\infty$ tensor product. Our connected sum formula for involutive link Floer homology is different, since the involutive complex is determined entirely by the chain homotopy type of the map ι_K. It seems likely that there is a description of our formula in terms of an $\mathcal{A}_\infty$ tensor product, but we will not pursue it in this paper.

6.1.7 Organization

In Section 6.2 we construct the group $\mathfrak{I}_K$ of ι_K–complexes modulo local equivalence. In Section 6.3 we provide background on the link Floer complexes and the link Floer TQFT. In Section 6.4, we prove some new results about the link Floer TQFT, such as Theorem 6.1.3 and Lemma 6.1.4, about conjugation invariance and the bypass relation, respectively. In Section 6.5, we prove that several naturally defined link cobordism maps induce filtered chain homotopy equivalences for connected sums of 3–manifolds and knots. In Section 6.6 we prove Theorem 6.1.1, the connected sum formula for the map ι_K. In Section 6.7, we combine our results to construct a homomorphism from $\mathcal{C}$ to $\mathfrak{I}_K$. In Section 6.8 we compute some examples, mostly using the computer algebra system *Macaulay2* [GS]. Finally in Appendix 6.9, we provide some alternate proofs of some of the results of the paper.

6.2 The group $\mathfrak{I}_K$

In this section we define $\mathfrak{I}_K$, the group generated by chain complexes with an involution, up to local equivalence. Along the way, we provide some algebraic preliminaries, such as some facts about the maps Φ and Ψ, which feature in our connected sum formula.

6.2.1 ι_K–complexes and the definition of $\mathfrak{I}_K$

In this section we define the group $\mathfrak{I}_K$. In the following subsections we will prove some basic facts, and also prove that $\mathfrak{I}_K$ is an abelian group. Throughout, we will let $\mathcal{R}$ denote the ring

$$\mathcal{R} := \mathbb{F}_2[U, V, U^{-1}, V^{-1}].$$

Definition 6.2.1. Suppose that C and C' are two chain complexes over $\mathcal{R}$ and $F : C \to C'$ is a homomorphism of abelian groups. We say F is $\mathcal{R}$ **skew-equivariant** if $F(U \cdot \mathbf{x}) = V \cdot F(\mathbf{x})$ and $F(V \cdot \mathbf{x}) = U \cdot F(\mathbf{x})$. If C and C' have a filtration by $\mathbb{Z} \oplus \mathbb{Z}$, denoted $\mathcal{F}$ and $\mathcal{F}'$ respectively, we say that F is **skew-filtered** if $F(\mathcal{F}_{(i,j)}) \subseteq \mathcal{F}'_{(j,i)}$. If C and C' are both bigraded with two gradings gr_U and gr_V, we will say that F is **skew-graded** if $\mathrm{gr}_U(F(\mathbf{x})) = \mathrm{gr}_V(\mathbf{x})$ and $\mathrm{gr}_V(F(\mathbf{x})) = \mathrm{gr}_U(\mathbf{x})$ for homogeneous elements $\mathbf{x}$.

We will write $\simeq$ for $\mathcal{R}$–equivariant, filtered chain homotopy equivalence, unless we specify otherwise. Similarly $\approx$ will mean $\mathcal{R}$ skew-equivariant, skew-filtered chain homotopy equivalence, unless specify otherwise.

Suppose (C, ∂) is a chain complex which is a free $\mathcal{R}$–module and such that ∂ is a homomorphism of $\mathcal{R}$–modules. If B is a basis over $\mathcal{R}$ for C, we can define a filtration $\mathcal{F}^B$ on C over $\mathbb{Z} \oplus \mathbb{Z}$, by defining $\mathcal{F}^B_{(i,j)}$ to be the subset generated by elements of the form $U^m V^n \cdot \mathbf{x}$, for $\mathbf{x} \in B$ and $m \geq i$ and $n \geq j$. If $\mathbf{x} \in B$, we can write

$$\partial(\mathbf{x}) = \sum_{\mathbf{y} \in B} P_{\mathbf{xy}} \cdot \mathbf{y},$$

for a unique $P_{\mathbf{xy}} \in \mathcal{R}$. Since ∂ is an $\mathcal{R}$–homomorphism, the polynomials $P_{\mathbf{xy}}$ determine ∂ on all of C.

If all of the $P_{\mathbf{xy}}$ are in $\mathbb{F}_2[U, V]$ (i.e. involve only nonnegative powers of U and V) then (C, ∂) is a $\mathbb{Z} \oplus \mathbb{Z}$ filtered chain complex with respect to the filtration $\mathcal{F}^B$. Given a basis B, we can formally differentiate each $P_{\mathbf{xy}}$, with respect to U or V. We define

$$\Phi_B(\mathbf{x}) = \sum_{\mathbf{y} \in B} \left(\tfrac{d}{dU} P_{\mathbf{xy}}\right) \cdot \mathbf{y} \qquad \text{and} \qquad \Psi_B(\mathbf{x}) = \sum_{\mathbf{y} \in B} \left(\tfrac{d}{dV} P_{\mathbf{xy}}\right) \cdot \mathbf{y}. \tag{6.1}$$

We call these the **formal derivatives** of the differential ∂. Note that ∂ is a filtered map with respect to $\mathcal{F}^B$ iff the matrix for ∂ in the basis B involves only nonnegative powers of U and V. Thus if ∂ is filtered, the maps Φ_B and Ψ_B are also filtered.

We now define the objects which will form the elements of the group $\mathfrak{I}_K$:

Definition 6.2.2. We say a tuple (C, ∂, B, ι) is an ι_K–complex if

1. (C, ∂) is a finitely generated free chain complex over $\mathcal{R}$ with basis B;
2. ∂ is a filtered map with respect to $\mathcal{F}^B$;
3. The elements of B are assigned two gradings, gr_U and gr_V, taking values in $\mathbb{Z}$. Extending gr_U and gr_V to the entire complex, by declaring U to have gr_U grading change -2, V to have gr_U grading change 0, and V to have gr_V grading change -2 and U to have gr_V grading change zero, the map ∂ is -1 graded with respect to both gr_U and gr_V;
4. There is a grading preserving isomorphism $H_*(C, \partial) \cong \mathcal{R}$ where $1 \in \mathcal{R}$ has gr_U and gr_V grading 0;
5. ι is a skew-filtered, skew-graded, $\mathcal{R}$ skew-equivariant endomorphism of C;
6. $\iota^2 \simeq \mathrm{id} + \Phi_B \circ \Psi_B$.

Remark 6.2.3. It will be convenient to combine the two gradings gr_U and gr_V into an **Alexander grading**, by defining $A = \frac{1}{2}(\mathrm{gr}_U - \mathrm{gr}_V)$. Notice that since ι switches gr_U and gr_V, we have that $A(\iota(\mathbf{x})) = -A(\mathbf{x})$ for homogeneous $\mathbf{x}$. Hence given an ι_K complex (C, ∂, B, ι), we can consider the subset C_0 of C concentrated in zero Alexander grading. On C_0, the two gradings gr_U and gr_V coincide, and furthermore ι restricts to a homotopy automorphism of C_0. The group C_0 is no longer an $\mathcal{R}$-module, but instead a $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$-module, where $\widehat{U} = UV$.

There is an obvious notion of homotopy equivalence between two ι_K complexes, but this notion is too strong to define a homomorphism from the concordance group. Instead we use

a weaker notion of equivalence from [Sto15] in the definition of $\mathfrak{I}_K$:

Definition 6.2.4. We will say two ι_K complexes, $\mathcal{C}_1 = (C_1, \partial_1, B_1, \iota_1)$ and $\mathcal{C}_2 = (C_2, \partial_2, B_2, \iota_2)$, are **locally equivalent** if there are filtered, grading preserving $\mathcal{R}$–equivariant chain maps

$$F : C_1 \to C_2, \qquad \text{and} \qquad G : C_2 \to C_1,$$

such that

$$\iota_2 F + F\iota_1 \approx 0, \qquad \text{and} \qquad \iota_1 G + G\iota_2 \approx 0$$

such that F and G are isomorphisms on homology. If in addition we have that $F \circ G \simeq \mathrm{id}$ and $G \circ F \simeq \mathrm{id}$, we will say that $\mathcal{C}_1$ and $\mathcal{C}_2$ are **homotopy equivalent**.

Note that homotopy equivalent ι_K–complexes are also locally equivalent.

Given two ι_K–complexes $\mathcal{C}_1 = (C_1, \partial_1, B_1, \iota_1)$ and $\mathcal{C}_2 = (C_2, \partial_2, B_2, \iota_2)$, there are two ι_K–complexes which can naturally be described as the product of $\mathcal{C}_1$ and $\mathcal{C}_2$. We define two products, $\times_1$ and $\times_2$, by the formulas

$$(\mathcal{C}_1 \times_i \mathcal{C}_2) = (C_1 \otimes_{\mathcal{R}} C_2, \partial_1 |\, \mathrm{id} + \mathrm{id}\, | \partial_2, B_1 \times B_2, \iota_1 \times_i \iota_2),$$

where

$$\iota_1 \times_1 \iota_2 = \iota_1 | \iota_2 + \Phi_1 \iota_1 | \Psi_2 \iota_2 \qquad \text{and} \qquad \iota_1 \times_2 \iota_2 = \iota_1 | \iota_2 + \Psi_1 \iota_1 | \Phi_2 \iota_2. \tag{6.2}$$

We can now define the group $\mathfrak{I}_K$:

Definition 6.2.5. We define $\mathfrak{I}_K$ to be the set generated local equivalence classes of ι_K–complexes. Group multiplication is given by either of $\times_1$ or $\times_2$. The inverse of $\mathcal{C} = (C, \partial, B, \iota)$ is given by $\mathcal{C}^\vee = (C^\vee, \partial^\vee, B^\vee, \iota^\vee)$, where $C^\vee = \mathrm{Hom}_{\mathcal{R}}(C, \mathcal{R})$ and the other terms are defined similarly. If $\mathbf{x} \in B$, then one defines $\mathrm{gr}(\mathbf{x}^\vee) = -\mathrm{gr}(\mathbf{x})$, for gr equal to either of gr_U or gr_V. The identity complex is $\mathcal{C}_e = (\mathcal{R}, 0, \{1\}, \iota_e)$, where $\iota_e : \mathcal{R} \to \mathcal{R}$ fixes $1 \in \mathcal{R}$ and switches U and V.

We will prove the following over the course of the following subsections:

Proposition 6.2.6. *The pair $(\mathfrak{I}_K, \times)$ is a well defined abelian group.*

6.2.2 Properties of Φ_B and Ψ_B

In this section, we describe some useful properties of the maps Φ_B and Ψ_B defined earlier. Recall that if B is a basis of (C, ∂) over $\mathcal{R}$, we defined Φ_B and Ψ_B to be the formal derivatives of ∂ with respect to U and V (respectively) in terms of the basis B.

As a first property, we note that the Leibniz rule can be interpreted as saying exactly that

$$\Phi_B = \partial \circ \tfrac{d}{dU}\big|_B + \tfrac{d}{dU}\big|_B \circ \partial \tag{6.3}$$

and

$$\Psi_B = \partial \circ \tfrac{d}{dV}\big|_B + \tfrac{d}{dV}\big|_B \circ \partial. \tag{6.4}$$

In particular the maps Φ_B and Ψ_B vanish on homology. Note that the derivative maps $(d/dU)|_B$ and $(d/dV)|_B$ are neither filtered nor $\mathcal{R}$–equivariant, so the maps Φ_B and Ψ_B may still be nonzero in the group of chain maps modulo $\mathcal{R}$–equivariant, filtered chain homotopy.

Definition 6.2.7. Given an ι_K complex $\mathcal{C} = (C, \partial, B, \iota)$ we define the **skew** of $\mathcal{C}$, denoted $\widetilde{\mathcal{C}} = (\widetilde{C}, \widetilde{\partial}, \widetilde{B}, \widetilde{\iota})$ as follows. As a $\mathbb{F}_2$–module with differential, we define $(\widetilde{C}, \widetilde{\partial}) = (C, \partial)$, but we give $\widetilde{C}$ a new $\mathcal{R}$–action. The element U now acts by V on $\widetilde{C}$ and the element V now acts by U, on $\widetilde{C}$. We set $\widetilde{B} = B$ and $\widetilde{\iota} = \iota$. For $\mathbf{x} \in B$ we define $\widetilde{\mathrm{gr}}_U(\mathbf{x}) = \mathrm{gr}_V(\mathbf{x})$ and $\widetilde{\mathrm{gr}}_V(\mathbf{x}) = \mathrm{gr}_U(\mathbf{x})$.

Note that by definition $\mathcal{F}^{\widetilde{B}}_{(i,j)} = \mathcal{F}^{B}_{(j,i)}$. Also, there is a natural map

$$s : C \to \widetilde{C},$$

which is the identity map on the $\mathbb{F}_2$–module C, but as a map between ι_K–complexes is skew-equivariant, skew-filtered and skew-graded. Tautologically, we have

$$s \circ \Phi_B = \Psi_{\widetilde{B}} \circ s \qquad \text{and} \qquad s \circ \Psi_B = \Phi_{\widetilde{B}} \circ s. \tag{6.5}$$

Lemma 6.2.8. *Suppose that (C_1, ∂_1, B_1) and (C_2, ∂_2, B_2) are two complexes over $\mathcal{R}$ with bases B_1 and B_2. If $F : C_1 \to C_2$ is a $\mathcal{R}$–equivariant then*

$$\Phi_2 \circ F + F \circ \Phi_1 \simeq 0 \qquad \textit{and} \qquad \Psi_2 \circ F + F \circ \Psi_1 \simeq 0,$$

through $\mathcal{R}$–equivariant chain homotopies. Similarly, if $G : C_1 \to C_2$ is a $\mathcal{R}$ skew-equivariant, then

$$\Psi_2 \circ G + G \circ \Phi_1 \eqsim 0 \qquad \textit{and} \qquad \Phi_2 \circ G + G \circ \Psi_1 \eqsim 0,$$

through skew-equivariant chain homotopies. If F (resp. G) is filtered (resp. skew-filtered) with respect to $\mathcal{F}^B$ then the chain homotopies can be taken to be filtered (resp. skew-filtered).

Proof. Suppose F is $\mathcal{R}$–equivariant. Writing ∂_1, ∂_2 and F in terms of matrices with respect to B_1 and B_2, we take the expression $\partial_2 \circ F + F \circ \partial_1 = 0$ and differentiate with respect to U, using the Leibniz rule, to see

$$\Phi_{B_2} \circ F + F \circ \Phi_{B_1} = F' \circ \partial_1 + \partial_2 \circ F',$$

where F' denotes the result of differentiating the matrix for F with respect to U. If F is a filtered map, then F' will be as well. The analogous argument works for Ψ_B.

Now suppose that $G : C_1 \to C_2$ is skew-equivariant. Note that $G = G \circ s \circ s$, and $G \circ s$ is an $\mathcal{R}$–equivariant map. Hence we just apply the previous result together with Equation (6.5) to see that

$$\Psi_2 \circ G + G \circ \Phi_1 \eqsim 0 \qquad \text{and} \qquad \Phi_2 \circ G + G \circ \Psi_1 \eqsim 0.$$

If G is skew-filtered, then $G \circ s$ is filtered, so the previous argument applies to show that the chain homotopies can be taken to be skew-filtered. □

We now prove that the maps Φ_B and Ψ_B depend only on the filtration $\mathcal{F}^B$ and not on the particular choice of basis, in the following sense:

Corollary 6.2.9. *If B and B' are two bases for the complex (C, ∂), then*

$$\Phi_B \simeq \Phi_{B'} \qquad \textit{and} \qquad \Psi_B \simeq \Psi_{B'},$$

through $\mathcal{R}$–equivariant chain homotopies. If in addition $\mathcal{F}^B = \mathcal{F}^{B'}$, then the chain homotopies can be taken to be filtered with respect to $\mathcal{F}^B = \mathcal{F}^{B'}$.

Proof. If B and B' are arbitrary bases, we apply the first part of Lemma 6.2.8 to the map $\mathrm{id}_C : C \to C$, which is $\mathcal{R}$–equivariant, to see that $\Phi_B \simeq \Phi_{B'}$ and $\Psi_B \simeq \Psi_{B'}$ through $\mathcal{R}$–equivariant chain homotopies. If $\mathcal{F}^B = \mathcal{F}^{B'}$, then id_C is also filtered, so applying the second part of Lemma 6.2.8 shows that the chain homotopies can be taken to be filtered. □

Note that if B and B' are two different bases, the matrix for the identity map between B and B' will not be the identity matrix. As such, the chain homotopy in the previous lemma will not vanish in general.

Lemma 6.2.10. *If (C, ∂) is a free chain complex over $\mathcal{R}$ with basis B, then Φ_B and Ψ_B are chain maps, and*

$$\Phi_B \circ \Psi_B \simeq \Psi_B \circ \Phi_B,$$

through $\mathcal{R}$–equivariant chain homotopies. If ∂ is filtered with respect to $\mathcal{F}^B$, then the chain homotopies can be taken to be filtered with respect to $\mathcal{F}^B$.

Proof. Taking the matrix expression for ∂ in the basis B, one differentiates the expression $\partial^2 = 0$ with respect to U to get $\Phi_B \circ \partial + \partial \circ \Phi_B = 0$. Taking $\Phi_B \circ \partial + \partial \circ \Phi_B = 0$ and differentiating it with respect to V, one sees that $\Phi_B \circ \Psi_B + \Psi_B \circ \Phi_B \simeq 0$. If ∂ is filtered, then so is Φ_B. If Φ_B is filtered, then so is the chain homotopy between $\Phi_B \circ \Psi_B + \Psi_B \circ \Phi_B$ and 0, since it is obtained by differentiating the matrix for Φ_B with respect to V. □

From Equations (6.3) and (6.4), we see that the maps Φ_B and Ψ_B are chain homotopic to zero, though not necessarily through filtered $\mathcal{R}$–equivariant chain homotopies. On the other hand, we have the following:

Lemma 6.2.11. *If (C, ∂) is a free chain complex over $\mathcal{R}$, with a basis B, then $(\Phi_B)^2 \simeq 0$ and $(\Psi_B)^2 \simeq 0$, through $\mathcal{R}$–equivariant chain homotopies. If ∂ is a filtered map with respect to $\mathcal{F}^B$, then the chain homotopies can also be taken to be filtered.*

Proof. The proof is the same as Lemma 4.3.9. For Φ_B, the chain homotopy is defined by first writing $\partial = \sum_{n=-\infty}^{\infty} P_n \cdot U^n$ where P_n is a matrix (using the basis B) with entries in

$\mathbb{F}_2[V, V^{-1}]$, then defining

$$H_B = \sum_{n=-\infty}^{\infty} \left(\frac{n(n-1)}{2}\right) P_n \cdot U^{n-2}.$$

It is easy to check that

$$\Phi_B^2 = \partial \circ H_B + H_B \circ \partial.$$

Finally, if ∂ is filtered with respect to $\mathcal{F}^B$, then P_n vanishes for $n < 0$ and P_n involves only nonnegative powers of V, so H_B is filtered as well. □

Remark 6.2.12. Since we are working over $\mathbb{F}_2$, the latter relations imply that if (C, ∂, B, ι) is an ι_K–complex, then $\iota^4 \simeq \mathrm{id}$, since

$$\iota^4 \simeq (\mathrm{id} + \Phi\Psi)^2 \simeq \mathrm{id} + 2\Phi\Psi + \Phi^2\Psi^2 \simeq \mathrm{id}\,.$$

6.2.3 $\mathfrak{I}_K$ is an abelian group

In this section we verify that our description determines an abelian group structure on $\mathfrak{I}_K$.

We first show that the product of two ι_K–complexes is an ι_K–complex:

Lemma 6.2.13. *If $\mathcal{C}_1 = (C_1, \partial_1, B_1, \iota_1)$ and $\mathcal{C}_2 = (C_2, \partial_2, B_2, \iota_2)$ are two ι_K–complexes, then $\mathcal{C}_1 \times_1 \mathcal{C}_2$ and $\mathcal{C}_1 \times_2 \mathcal{C}_2$ are both ι_K–complexes.*

Proof. Only conditions (4) and (6) of Definition 6.2.2 are not entirely straightforward to check. Condition (4) states $H_*(C_1 \otimes_{\mathcal{R}} C_2) \cong \mathcal{R}$. By our grading assumptions, we can write C_1 and C_2 each as a direct sum (over $\mathbb{F}_2$) over Alexander gradings (where $A = \frac{1}{2}(\mathrm{gr}_U - \mathrm{gr}_V)$). Write $(C_i)_j$ for the portion of C_i in Alexander grading j. We note that each $(C_i)_j$ is an $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$–module, where $\widehat{U} = UV$. Now the differential is zero graded with respect to the Alexander grading, so this gives C_i the structure of a direct sum of chain complexes over $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$. But $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$ is a PID, so using the classification of chain complexes over a PID (see e.g. [HMZ16, Lemma 6.1]), we can write $(C_i)_0$ as a direct sum of 1–step and 2–step complexes. A 1–step complex is a copy of $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$ with vanishing differential, and a 2–step complex is a complex over $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$, with two generators a and b, such that $\partial(a) = f(\widehat{U}) \cdot b$

for some $f(\widehat{U}) \in \mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$. Since $(C_i)_0$ has a $\mathbb{Z}$ grading, it is not hard to see that the 2–step complexes must be of the form $\partial(a) = \widehat{U}^n b$. Clearly This decomposition of $(C_i)_0$ as a sum of 1–step and 2–step complexes over $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$ can be extended linearly over $\mathcal{R}$ to a decomposition of C_i over $\mathcal{R}$ as a sum 1–step and 2–step complexes.

Over $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$, 1–step complexes of the form $\partial(a) = \widehat{U}^n b$ are $\mathcal{R}$–equivariantly chain homotopy equivalent to the zero complex, since $\widehat{U}^n$ is invertible (though of course not filtered chain homotopy equivalent to the zero complex). The homology $H_*(C_1 \otimes_{\mathcal{R}} C_2)$ only depends on the chain homotopy type of C_i over $\mathcal{R}$, so such 2–step complexes can be deleted from each C_i. Since by assumption the homology of C_i is isomorphic to $\mathcal{R}$, with $1 \in \mathcal{R}$ given zero gr_U and gr_V grading, we know that both C_1 and C_2 are $\mathcal{R}$–equivariantly chain homotopy equivalent (though not filtered chain homotopy equivalent) to a 1-step complex. Hence $H_*(C_1 \otimes_{\mathcal{R}} C_2) \cong \mathcal{R}$.

The other condition worth mentioning is condition (6), which requires that $(\iota_1 \times_i \iota_2)^2 \simeq \mathrm{id} + \Phi \circ \Psi$. To see this, note that on $C_1 \otimes_{\mathcal{R}} C_2$, we have

$$\mathrm{id} + \Phi \circ \Psi = \mathrm{id} \,|\, \mathrm{id} + (\Phi_1 |\, \mathrm{id} + \mathrm{id} \,|\Phi_2)(\Psi_1 |\, \mathrm{id} + \mathrm{id} \,|\Psi_2). \tag{6.6}$$

On the other hand

$$\begin{aligned}
(\iota_1 \times_1 \iota_2)^2 &= (\iota_1 | \iota_2 + \Phi_1 \iota_1 | \Psi_2 \iota_2)(\iota_1 | \iota_2 + \Phi_1 \iota_1 | \Psi_2 \iota_2) \\
&= \iota_1^2 | \iota_2^2 + \Phi_1 \iota_1^2 | \Psi_2 \iota_2^2 + \iota_1 \Phi_1 \iota_1 | \iota_2 \Psi_2 \iota_2 + \Phi_1 \iota_1 \Phi_1 \iota_1 | \Psi_2 \iota_2 \Psi_2 \iota_2 \\
&\simeq (\mathrm{id} + \Phi_1 \Psi_1) | (\mathrm{id} + \Phi_2 \Psi_2) + \Phi_1 (\mathrm{id} + \Phi_1 \Psi_1) | \Psi_2 (\mathrm{id} + \Phi_2 \Psi_2) \\
&\qquad + \Psi_1 (\mathrm{id} + \Phi_1 \Psi_1) | \Phi_2 (\mathrm{id} + \Phi_2 \Psi_2) + \Phi_1 \Psi_1 (\mathrm{id} + \Phi_1 \Psi_1) | \Psi_2 \Phi_2 (\mathrm{id} + \Phi_2 \Psi_2) \\
&\simeq \mathrm{id} \,|\, \mathrm{id} + \Phi_1 | \Psi_2 + \Psi_1 | \Phi_2 + \Phi_1 \Psi_1 |\, \mathrm{id} + \mathrm{id} \,| \Phi_2 \Psi_2
\end{aligned}$$

which is the same as Equation (6.6). A similar computation yields that $(\iota_1 \times_2 \iota_2)^2 \simeq \mathrm{id} + \Phi \circ \Psi$, as well. □

We now show that the two multiplications $\times_1$ and $\times_2$ coincide on homotopy classes of ι_K–complexes (and hence local equivalence classes):

Lemma 6.2.14. *If $\mathcal{C}_1 = (C_1, \partial_1, B_1, \iota_1)$ and $\mathcal{C}_2 = (C_2, \partial_2, B_2, \iota_2)$ are two ι_K–complexes, then the ι_K–complexes $\mathcal{C}_1 \times_1 \mathcal{C}_2$ and $\mathcal{C}_1 \times_2 \mathcal{C}_2$ are homotopy equivalent.*

Proof. We define the map $F : C_1 \otimes C_2 \to C_1 \otimes C_2$ by

$$F = (\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2).$$

Note that since $\Psi_1^2 \simeq 0$ and $\Phi_2^2 \simeq 0$, and since we are working over $\mathbb{F}_2$, it follows that $F^2 \simeq \mathrm{id}$. We claim that F induces a homotopy equivalence of ι_K–complexes. Note that in the definition of a homotopy equivalence, there are two maps. We use the map F in both directions. We need to show that F intertwines $\iota_1 \times_1 \iota_2$ and $\iota_1 \times_2 \iota_2$. We first compute

$$\begin{aligned} F(\iota_1 \times_1 \iota_2) &\simeq (\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2)(\mathrm{id} \,|\, \mathrm{id} + \Phi_1 | \Psi_2)(\iota_1 | \iota_2) \\ &\simeq (\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2)(\iota_1 | \iota_2)(\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2) \\ &\simeq (\iota_1 \times_2 \iota_2) F. \end{aligned}$$

In a similar manner, we have

$$\begin{aligned} F(\iota_1 \times_2 \iota_2) &\simeq (\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2)(\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2)(\iota_1 | \iota_2) \\ &\simeq (\iota_1 | \iota_2) \\ &\simeq (\iota_1 | \iota_2)(\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2)(\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2) \\ &\simeq (\mathrm{id} \,|\, \mathrm{id} + \Phi_1 | \Psi_2)(\iota_1 | \iota_2)(\mathrm{id} \,|\, \mathrm{id} + \Psi_1 | \Phi_2) \\ &\simeq (\iota_1 \times_1 \iota_2) F, \end{aligned}$$

completing the proof. □

In light of the previous lemma, we will write $\times$ for the multiplication on $\mathfrak{I}_K$, if no confusion will arise. We now show that $\times$ is well defined on local equivalence classes:

Lemma 6.2.15. *The $\times$ operation is a well defined multiplication on $\mathfrak{I}_K$. That is, if $\mathcal{C}_1$ and $\mathcal{C}_1'$ are locally equivalent, and $\mathcal{C}_2$ and $\mathcal{C}_2'$ are locally equivalent, then $\mathcal{C}_1 \times \mathcal{C}_2$ and $\mathcal{C}_1' \times \mathcal{C}_2'$ are locally equivalent.*

Proof. Let us consider changing the first factor, since changing the second factor follows similarly. Suppose that $F : C_1 \to C_1'$ and $G : C_1' \to C_1$ are maps which induce a local equivalence. Since the maps $H_*(C_1) \otimes H_*(C_2) \to H_*(C_1 \otimes C_2)$ and $H_*(C_1') \otimes H_*(C_2) \to H_*(C_1' \otimes C_2)$ are isomorphisms by the same argument as in Lemma 6.2.13, we see that $(F|\operatorname{id}) \circ (G|\operatorname{id}) = (F \circ G)|\operatorname{id}$ and $(G|\operatorname{id}) \circ (F|\operatorname{id})$ are isomorphisms on homology. It remains only to show that $F|\operatorname{id}$ and $G|\operatorname{id}$ intertwine the involutions, up to skew-equivariant, skew-filtered homotopies.

For $F|\operatorname{id}$, this amounts to showing that

$$(F|\operatorname{id})(\iota_1|\iota_2 + \Phi_1\iota_1|\Psi_2\iota_2) \approx (\iota_1'|\iota_2 + \Phi_1'\iota_1'|\Psi_2\iota_2)(F|\operatorname{id}).$$

Rearranging terms reduces this to

$$(F\iota_1 + \iota_1'F)|\iota_2 + (F\Phi_1\iota_1 + \Phi_1'\iota_1'F)|\Psi_2\iota_2 \approx 0.$$

By assumption, $F\iota_1 \approx \iota_1'F$, since F was part of a local equivalence. By Lemma 6.2.8, we also know that $F\Phi_1 \simeq \Phi_1'F$. Combining these facts shows that the above identity is satisfied, so $F|\operatorname{id}$ intertwines the involutions, up to skew-equivariant, skew-filtered chain homotopies. The same argument works for G. The same argument also works for the second factor of the tensor product, and hence the proof is complete. □

Lemma 6.2.16. *The multiplication $\times$ is associative on $\mathfrak{I}_K$.*

Proof. Suppose that $\mathcal{C}_i = (C_i, \partial_i, B_i, \iota_i)$ are ι_K complexes for $i \in \{1, 2, 3\}$. We will construct a homotopy equivalence between the two ι_K complexes $(\mathcal{C}_1 \times \mathcal{C}_2) \times \mathcal{C}_3$ and $\mathcal{C}_1 \times (\mathcal{C}_2 \times \mathcal{C}_3)$. The map on chain complexes will be the identity. The only nontrivial thing to show is that the identity map intertwines the involutions, up to skew-equivariant, skew-filtered homotopy. This is a computation. We have

$$\begin{aligned}(\iota_1 \times_1 \iota_2) \times_1 \iota_3 &= (\iota_1|\iota_2 + \Phi_1\iota_1|\Psi_2\iota_2)|\iota_3 + \big((\Phi_1|\operatorname{id} + \operatorname{id}|\Phi_2)(\iota_1|\iota_2 + \Phi_1\iota_1|\Psi_2\iota_2)\big)|\Psi_3\iota_3 \\ &= \iota_1|\iota_2|\iota_3 + \Phi_1\iota_1|\Psi_2\iota_2|\iota_3 + \Phi_1\iota_1|\iota_2|\Psi_3\iota_3 + \iota_1|\Phi_2\iota_2|\Psi_3\iota_3 + \Phi_1\iota_1|\Phi_2\Psi_2\iota_2|\Psi_3\iota_3 \\ &\quad + \Phi_1^2\iota_1|\Psi_2\iota_2|\Psi_3\iota_3.\end{aligned}$$

On the other hand, we can also compute that

$$\begin{aligned}\iota_1 \times_1 (\iota_2 \times_1 \iota_3) &= \iota_1|(\iota_2|\iota_3 + \Phi_2\iota_2|\Psi_3\iota_3) + \Phi_1\iota_1|\big((\Psi_2|\,\mathrm{id} + \mathrm{id}\,|\Psi_3)(\iota_2|\iota_3 + \Phi_2\iota_2|\Psi_3\iota_3)\big)\\ &= \iota_1|\iota_2|\iota_3 + \iota_1|\Phi_2\iota_2|\Psi_3\iota_3 + \Phi_1\iota_1|\Psi_2\iota_2|\iota_3 + \Phi_1\iota_1|\iota_2|\Psi_3\iota_3 + \Phi_1\iota_1|\Psi_2\Phi_2\iota_2|\Psi_3\iota_3\\ &\quad + \Phi_1\iota_1|\Phi_2\iota_2|\Psi_3^2\iota_3.\end{aligned}$$

The difference is

$$\Phi_1\iota_1|\Phi_2\Psi_2\iota_2|\Psi_3\iota_3 + \Phi_1^2\iota_1|\Psi_1\iota_2|\Psi_3\iota_3 + \Phi_1\iota_1|\Psi_2\Phi_2\iota_2|\Psi_3\iota_3 + \Phi_1\iota_1|\Phi_2\iota_2|\Psi_3^2\iota_3,$$

which is skew-equivariantly, skew-filtered chain homotopic to 0, by Lemmas 6.2.10 and 6.2.11.

□

We now consider inverses on $\mathfrak{I}_K$. We first prove the following:

Lemma 6.2.17. *If $\mathcal{C} = (C, \partial, B, \iota)$ is an ι_K–complex, then the tuple $\mathcal{C}^\vee = (C^\vee, \partial^\vee, B^\vee, \iota^\vee)$ is an ι_K–complex.*

Proof. Most of the axioms are straightforward to verify. To see that $H_*(C^\vee) \cong \mathcal{R}$, one can write down an explicit description of the complex C over $\mathcal{R}$, using the strategy of Lemma 6.2.13 to decompose over Alexander gradings, and then using the classification of chain complexes over the PID $\mathbb{F}_2[\widehat{U}, \widehat{U}^{-1}]$, where $\widehat{U} = UV$. With such description of the dual chain complex, the result follows immediately. □

We now show that $\mathcal{C}^\vee$ is the inverse of $\mathcal{C}$ in $\mathfrak{I}_K$ (compare [HMZ16, Proposition 8.8]):

Lemma 6.2.18. *If $\mathcal{C} = (C, \partial, B, \iota)$ is an ι_K–complex, the product $\mathcal{C} \times \mathcal{C}^\vee$ is locally equivalent to $\mathcal{C}_e = (\mathcal{R}, 0, \{1\}, \iota_e)$, the identity ι_K–complex.*

Proof. The chain complex $C \otimes_{\mathcal{R}} C^\vee$ can be naturally identified with the complex $\mathrm{Hom}_{\mathcal{R}}(C, C)$. We construct two maps,

$$F : \mathcal{R} \to \mathrm{Hom}_{\mathcal{R}}(C, C), \qquad \text{and} \qquad G : \mathrm{Hom}_{\mathcal{R}}(C, C) \to \mathcal{R},$$

as follows. The map F sends $1 \in \mathcal{R}$ to $\mathrm{id}_C \in \mathrm{Hom}_{\mathcal{R}}(C, C)$, and G is defined to be the trace of a matrix. In terms of the basis B, the map F takes the form

$$F(1) = \sum_{\mathbf{x} \in B} \mathbf{x} \otimes \mathbf{x}^\vee$$

and G takes the form

$$G(\mathbf{x} \otimes \mathbf{y}^\vee) = \mathbf{y}^\vee(\mathbf{x}) = \begin{cases} 1 & \text{if} \quad \mathbf{x} = \mathbf{y}, \\ 0 & \text{if} \quad \mathbf{x} \neq \mathbf{y}. \end{cases}$$

The maps F and G are clearly filtered, and respect the grading. We need to show that they are chain maps, intertwine the involutions, and are isomorphisms on homology. That they are chain maps is an easy computation. To see that they are isomorphisms on homology, we note that the composition $G \circ F$ gives the Euler characteristic of $H_*(C)$ over $\mathcal{R}$, modulo 2, which is 1. Since $\mathcal{C}, \mathcal{C}^\vee$ and $\mathcal{C} \times \mathcal{C}^\vee$ are all ι_K–complexes, we know the associated homology groups are all $\mathcal{R}$ so we conclude that F and G must be isomorphisms on homology.

To see that F intertwines the involutions, it is sufficient to show that

$$F\iota_e + (\iota|\iota^\vee + \Phi\iota|\Psi^\vee\iota^\vee)F \backsimeq 0.$$

As both summands on the left are skew-equivariant, it is sufficient to show that

$$F(1) + (\iota|\iota^\vee + \Phi\iota|\Psi^\vee\iota^\vee)F(1) = \partial h, \tag{6.7}$$

for some $h \in C \otimes C^\vee$. Note that $F(1) = \sum_{\mathbf{x} \in B} \mathbf{x} \otimes \mathbf{x}^\vee$. If $L : C \to C$ is an $\mathcal{R}$–equivariant map, using the basis B we can associate a matrix $M(L) \in \mathrm{Mat}_{n \times n}(\mathcal{R})$, and the action of L corresponds to matrix multiplication. If L is skew-equivariant, then we can also associate a matrix $M(L)$. The action of L on an an element $x \in C$ is given by first writing x as an $\mathcal{R}$–linear combination of basis elements in B, switching the powers of U and V, then multiplying by the matrix $M(L)$. If $A : C \to C$ and $B : C^\vee \to C^\vee$ are maps (equivariant or skew equivariant) the element $(A|B)(\sum_{\mathbf{x} \in B} \mathbf{x}|\mathbf{x}^\vee) \in C \otimes C^\vee$ can be identified with the matrix $M(A)M(B)^T \in \mathrm{Mat}_{n \times n}(\mathcal{R})$. As such, the expression on the left side of Equation (6.7) corresponds to the matrix

$$\mathrm{id}_{n \times n} + M(\iota)M(\iota^\vee)^T + M(\Phi)M(\iota)M(\iota^\vee)^T M(\Psi^\vee)^T.$$

Since $\iota^\vee, \Psi^\vee$ and $\Phi^\vee$ have the matrices $M(\iota)^T, M(\Psi)^T$ and $M(\Phi)^T$, respectively, the above matrix can naturally be identified with the homomorphism

$$\mathrm{id}_C + \iota^2 + \Phi\iota^2\Psi \in \mathrm{Hom}_{\mathcal{R}}(C, C),$$

which can be seen to be chain homotopic to zero using the algebraic relation $\iota^2 \simeq \mathrm{id} + \Phi \circ \Psi$, in the definition of an ι_K–complex. A chain homotopy of maps in $\mathrm{Hom}_{\mathcal{R}}(C, C)$ corresponds to a boundary in the complex $\mathrm{Hom}_{\mathcal{R}}(C, C)$, so we conclude that there is an $h \in C \otimes C^\vee$ such that

$$F(1) + (\iota|\iota^\vee + \Phi\iota|\Psi^\vee\iota^\vee)F(1) = \partial h.$$

The proof that G intertwines the involutions follows similarly. □

We can now prove that $\mathfrak{I}_K$ is a well defined abelian group:

Proof of Proposition 6.2.6. Lemma 6.2.13 shows that the product of two ι_K–complexes under either multiplication $\times_1$ and $\times_2$ is an ι_K–complex. Lemma 6.2.14 shows that the two products yield homotopy equivalent (and hence locally equivalent) ι_K–complexes. Similarly, if $\mathcal{C}_1$ and $\mathcal{C}_2$ are two ι_K–complexes, the natural map $T : C_1 \otimes_{\mathcal{R}} C_2 \to C_2 \otimes_{\mathcal{R}} C_1$ induces a homotopy equivalence of ι_K–complexes between $\mathcal{C}_1 \times_1 \mathcal{C}_2$ and $\mathcal{C}_2 \times_2 \mathcal{C}_1$. An application of Lemma 6.2.14 now shows that $\mathcal{C}_1 \times_2 \mathcal{C}_2$ is homotopy equivalent to $\mathcal{C}_1 \times_1 \mathcal{C}_2$, from which we conclude that $\mathcal{C}_1 \times_1 \mathcal{C}_2$ and $\mathcal{C}_2 \times_1 \mathcal{C}_1$ are homotopy equivalent, so $\mathfrak{I}_K$ is abelian. Lemma 6.2.16 shows that multiplication is associative. Lemma 6.2.15 shows that multiplication is well defined on local equivalence classes. Lemmas 6.2.17 and 6.2.18 show that $\mathfrak{I}_K$ contains inverses. □

6.3 Background on link Floer homology and the link Floer TQFT

In this section, we provide background on link Floer homology, and describe some previous results about link cobordisms and link Floer homology. The original constructions of link Floer homology can be found in [OS08], and more detail on the curved variation and the link Floer TQFT can be found in Chapter 3 or Chapter 4.

6.3.1 Preliminaries on the link Floer complexes

Given a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ representing the multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ in Y, we will define the full link Floer complex

$$\mathcal{C}FL^{\circ}(\mathcal{H}, \mathfrak{s}),$$

for $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. We will focus on $\circ \in \{-, \infty, \wedge\}$. If $\mathbf{w} = \{w_1, \ldots, w_n\}$ and $\mathbf{z} = \{z_1, \ldots, z_n\}$ are collections of basepoints and $I = (i_1, \ldots, i_n)$ and $J = (j_1, \ldots, j_n)$ are multi-indices, then we write $U_{\mathbf{w}}^I V_{\mathbf{z}}^J$ for the monomial

$$U_{\mathbf{w}}^I V_{\mathbf{z}}^J := U_{w_1}^{i_1} \cdots U_{w_n}^{i_n} V_{z_1}^{j_1} \cdots V_{z_n}^{j_n}.$$

We consider the tori $\mathbb{T}_\alpha := \alpha_1 \times \cdots \times \alpha_{g+n-1}$ and $\mathbb{T}_\beta := \beta_1 \times \cdots \times \beta_{g+n-1}$ in $\mathrm{Sym}^k(\Sigma)$ (where $g = g(\Sigma)$ and $n = |\mathbf{w}| = |\mathbf{z}|$) and consider the intersection points of $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$. The ∞ flavor of $\mathcal{C}FL^{\circ}(Y, \mathbb{L}, \mathfrak{s})$ is generated over $\mathbb{F}_2$ by monomials of the form $U_{\mathbf{w}}^I V_{\mathbf{z}}^J \cdot \mathbf{x}$, for arbitrary multi-indices I and J, and intersection points $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$. The $-$ flavor is generated over $\mathbb{F}_2$ by monomials with nonnegative multi-indices I and J. The hat flavor, denoted $\widehat{CFL}$, is generated over $\mathbb{F}_2$ by monomials with $I = J = 0$, i.e. setting all of the $U_{\mathbf{w}}$ and $V_{\mathbf{z}}$ variables equal to zero.

After picking an auxiliary almost complex structure, one can define an endomorphism ∂ which counts pseudo-holomorphic disks of Maslov index 1 inside of $\mathrm{Sym}^{g+n-1}(\Sigma)$, keeping track of their multiplicity over both the $\mathbf{w}$ and $\mathbf{z}$ basepoints by increasing the powers of the variables when a disk crosses over a basepoint. Equivalently, one can count holomorphic curves inside of $\Sigma \times [0,1] \times \mathbb{R}$, as in [Lip06].

Counting the ends of the moduli spaces of Maslov index two holomorphic disks yields the formula

$$\partial^2(\mathbf{x}) = \sum_{K \in C(L)} (U_{w_{K,1}} V_{z_{K,1}} + V_{z_{K,1}} U_{w_{K,2}} + U_{w_{K,2}} V_{z_{K,2}} + \cdots V_{z_{K,n_K}} U_{w_{K,1}}) \cdot \mathbf{x}. \tag{6.8}$$

See Lemma 3.2.1 for the proof. Here $C(L)$ denotes the set of components of L, and if $K \in C(L)$, we are writing $w_{K,1}, z_{K,1}, \ldots, z_{K,n_K}$ for the basepoints on K, in the order that

they appear on K. The above formula follows from the fact that the compactification of the moduli spaces of index 2 disks consists of broken strips (corresponding to $\partial^2(\mathbf{x})$) as well as index 2 boundary degenerations (corresponding to the terms on the right side of Equation (6.8)). For $\circ \in \{-, \infty\}$, the modules $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ are filtered by powers of the variables. As pseudo-holomorphic disks have only nonnegative multiplicities on the diagram, the differential is a filtered map.

Any two Heegaard diagrams for a link $\mathbb{L}$ in Y can be connected by a sequence of elementary Heegaard moves. To a sequence of elementary Heegaard moves between two diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ for a link $\mathbb{L}$ in Y, we can define a filtered, equivariant chain homotopy equivalence $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$, as in [OS04e]. A fundamental result is that the map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ is independent of the choice of elementary Heegaard moves, up to equivariant, filtered chain homotopy. See [JT12] for more on naturality in Heegaard Floer homology. We write

$$\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}),$$

to mean the collection of all the chain complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ ranging over strongly $\mathfrak{s}$–admissible diagrams $\mathcal{H}$ for $(Y, \mathbb{L})$, together with the change of diagrams maps. We sometimes call the objects $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ the **coherent, equivariant, filtered chain homotopy types**.

A morphism from $\mathcal{C}FL^\circ(Y_1, \mathbb{L}_1, \mathfrak{s}_1)$ to $\mathcal{C}FL^\circ(Y_2, \mathbb{L}_2, \mathfrak{s}_2)$ is a collection of filtered, equivariant maps $F_{\mathcal{H}_1, \mathcal{H}_2}$ for each pair of diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$ for $(Y_1, \mathbb{L}_1)$ and $(Y_2, \mathbb{L}_2)$, respectively, which commute with the change of diagrams maps up to equivariant, filtered chain homotopies.

For functoriality, we also need to work with collections of formal variables, which we assign to each basepoint:

Definition 6.3.1. A **coloring** of a multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ in Y is a pair $(\sigma, \mathfrak{P})$ where $\mathfrak{P}$ is a finite set (thought of as an index set for a collection of formal variables) and $\sigma : \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ is a function.

If $\mathfrak{P} = \{p_1, \dots, p_n\}$ we will write $\mathbb{F}_2[U_\mathfrak{P}]$ for the ring $\mathbb{F}_2[U_{p_1}, \dots, U_{p_n}]$. If $(\sigma, \mathfrak{P})$ is a coloring of $(L, \mathbf{w}, \mathbf{z})$, then $\mathbb{F}_2[U_\mathfrak{P}]$ becomes a module over $\mathbb{F}_2[U_\mathbf{w}, V_\mathbf{z}]$. If $\mathcal{H}$ is a diagram for $\mathbb{L}$, we can

form the **colored complex**

$$\mathcal{C}FL^-(\mathcal{H},\sigma,\mathfrak{P},\mathfrak{s}) := \mathcal{C}FL^-(\mathcal{H},\mathfrak{s}) \otimes_{\mathbb{F}_2[U_{\mathbf{w}},V_{\mathbf{z}}]} \mathbb{F}[U_{\mathfrak{P}}].$$

We will usually suppress the coloring from the notation. Colored versions of $\mathcal{C}FL^\infty(\mathcal{H},\mathfrak{s})$ can be similarly defined.

Note that by using Equation (6.8), it is straightforward to find many examples of colorings which have $\partial^2 = 0$. For example, if there are only two colors, all $\mathbf{w}$ basepoints are given one color, and all $\mathbf{z}$ are given the other, then $\partial^2 = 0$.

6.3.2 The conjugation action on link Floer homology

In this section we provide some background about the conjugation action on link Floer homology. Recall that Spin^c structures on Y can be described as homology classes of nonvanishing vector fields. Given a nonvanishing vector field v on Y which corresponds to a Spin^c structure $\mathfrak{s}$, the conjugate Spin^c structure $\overline{\mathfrak{s}}$ corresponds to the vector field $-v$.

In [OS04e], Ozsváth and Szabó describe a map

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

The map $\mathfrak{s}_{\mathbf{w}}$ depends on the choice of basepoints $\mathbf{w}$, in the following sense:

Lemma 6.3.2. *If $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z})$ is a diagram for $(Y,L,\mathbf{w},\mathbf{z})$ and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, then the maps $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ are related by the formula*

$$\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) - \mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = PD[L].$$

Proof. See, for example, [Ras07, Equation (1)] or [OS04e, Lemma 2.19]. □

Given a Heegaard diagram $\mathcal{H} = (\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{p},\mathbf{q})$ for $\mathbb{L} = (L,\mathbf{p},\mathbf{q})$, we can form the conjugate diagram $\overline{\mathcal{H}} = (-\Sigma,\overline{\boldsymbol{\beta}},\overline{\boldsymbol{\alpha}},\mathbf{q},\mathbf{p})$ for $\overline{\mathbb{L}} = (L,\mathbf{q},\mathbf{p})$. Here the embedded Heegaard surface is unchanged, though the orientation of the surface is reversed, and the roles of the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$

curves are reversed. Here $\mathbb{L}$ and $\overline{\mathbb{L}}$ have the same underlying oriented link, and both have $\mathbf{p} \cup \mathbf{q}$ as basepoints, but the role of $\mathbf{p}$ and $\mathbf{q}$ as type $\mathbf{w}$ or $\mathbf{z}$ is switched.

There is a tautological chain isomorphism

$$\eta : \mathcal{C}FL^{\infty}_{J_s}(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^{\infty}_{-J_s}(\overline{\mathcal{H}}, \overline{\mathfrak{s}} + PD[L]).$$

The map η sends an intersection point in $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ on $\mathcal{H}$ to the corresponding intersection point $\overline{\mathbf{x}} \in \mathbb{T}_{\overline{\beta}} \cap \mathbb{T}_{\overline{\alpha}}$. Given a coloring $(\sigma, \mathfrak{P})$, the map η induces a map on the colored complexes which is $\mathbb{F}_2[U_{\mathfrak{P}}]$–equivariant and $\mathfrak{P}$–filtered.

Lemma 6.3.3. *If $\mathcal{H}_1$ and $\mathcal{H}_2$ are two diagrams, then the following diagram commutes up to chain homotopy:*

$$\begin{array}{ccc} \mathcal{C}FL^{\infty}(\mathcal{H}_1, \mathfrak{s}) & \xrightarrow{\eta} & \mathcal{C}FL^{\infty}(\overline{\mathcal{H}}_1, \overline{\mathfrak{s}} + PD[L]) \\ \downarrow{\scriptstyle \Phi_{\mathcal{H}_1 \to \mathcal{H}_2}} & & \downarrow{\scriptstyle \Phi_{\overline{\mathcal{H}}_1 \to \overline{\mathcal{H}}_2}} \\ \mathcal{C}FL^{\infty}(\mathcal{H}_2, \mathfrak{s}) & \xrightarrow{\eta} & \mathcal{C}FL^{\infty}(\overline{\mathcal{H}}_2, \overline{\mathfrak{s}} + PD[L]). \end{array}$$

Proof. The map $\Phi_{\mathcal{H}_1 \to \mathcal{H}_2}$ is a computed by picking a sequence of intermediate diagrams which differ each from the previous by a single Heegaard move. The map $\Phi_{\overline{\mathcal{H}}_1 \to \overline{\mathcal{H}}_2}$ can be computed using the same sequence of Heegaard moves, and it is easy to verify that the diagram commutes on the nose for each elementary Heegaard move (though the change of diagrams maps themselves are only defined up to chain homotopy). □

As a consequence of the previous lemma, the map η descends to a map on the coherent, equivariant filtered chain homotopy type

$$\eta : \mathcal{C}FL^{\infty}(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^{\infty}(Y, \overline{\mathbb{L}}, \overline{\mathfrak{s}} + PD[L]),$$

since it commutes up to filtered, equivariant chain homotopy, with the change of diagrams maps.

For a doubly based knot, the conjugation map ι_K is defined as the composition

$$\iota_K := \tau_K \circ \eta,$$

where τ_K denotes a half twist in the direction of the links orientation, switching w and z. The map η is filtered and equivariant, while τ_K is skew-filtered and skew-equivariant since it switches the two basepoints.

6.3.3 Maps associated to decorated link cobordisms

In Chapter 4, the author describes Spin^c functorial maps associated to decorated link cobordisms. To a decorated link cobordism $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ in the sense of Definition 6.1.2 (together with a coloring of the components of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, which we suppress in the notation), there is a map

$$F_{W,F,\mathfrak{s}} : \mathcal{C}FL^\infty(Y_1, \mathbb{L}_1, \mathfrak{s}|_{Y_1}) \to \mathcal{C}FL^\infty(Y_2, \mathbb{L}_2, \mathfrak{s}|_{Y_2}),$$

which is an invariant of the decorated link cobordism, up to equivariant, filtered chain homotopy. The identity cobordism induces the identity map, and together the maps satisfy the Spin^c composition law, i.e. if $(W, F) = (W_2, F_2) \circ (W_1, F_1)$ and $\mathfrak{s}_i \in \text{Spin}^c(W_i)$ (and the surfaces are decorated with compatible colorings) one has

$$F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1,\mathfrak{s}_1} \simeq \sum_{\substack{\mathfrak{s} \in \text{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F_{W,F,\mathfrak{s}}.$$

If we are working over a set of colors $\mathfrak{P}$, the colored chain complexes are $\mathbb{F}_2[U_{\mathfrak{P}}]$–modules, and the maps are morphisms of $\mathbb{F}_2[U_{\mathfrak{P}}]$–modules. The colors of components of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ can be arbitrary, though the links forming the boundary of the surface must be colored compatibly with $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$.

In this section, we briefly describe the construction of the link cobordism maps from Chapter 4. More details (such as explicit formulas) can be found in our proof of Theorem 6.1.3, below. Further details can be found in Chapter 4.

The cobordism maps are defined by taking a decomposition of the link cobordism into simple pieces, and then defining a map for each piece. One defines maps for 4–dimensional handles of W added away from $\mathbb{L}$, maps associated to band surgery on the link $\mathbb{L}$, as well as cylindrical cobordisms with simple but nontrivial dividing sets.

The 1–handle and 3–handle maps are approximately the same as in [OS06] (the difference being now that we allow 1–handles or 3–handles which connect different components or separate a component of a 3–manifold).

There are also new 0–handle and 4–handle maps, corresponding to a 4–ball containing a 2–dimensional disk in a standard way.

The 2–handle maps are very similar to those defined by Ozsváth and Szabó in [OS06], which count holomorphic triangles in a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w})$ which represents surgery on a framed 1–dimensional link $\mathbb{S}_1 \subseteq Y$. In [OS06], Ozsváth and Szabó define a map

$$\mathfrak{s}_{\mathbf{w}} : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathrm{Spin}^c(W(\mathbb{S}_1)), \tag{6.9}$$

and the map associated to a Spin^c structure $\mathfrak{s}$ from [OS06] counts holomorphic representatives of homology classes of triangles ψ which have $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$.

Similarly the 2–handle maps from Chapter 4 are also defined by counting holomorphic triangles on a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ representing homology classes of triples ψ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$.

An important subtlety of the link Floer cobordism maps is the dependence of $\mathfrak{s}_{\mathbf{w}}$ on the basepoints $\mathbf{w}$. To see the dependence of the map $\mathfrak{s}_{\mathbf{w}}$ on the basepoints $\mathbf{w}$, we will consider more generally the situation of an arbitrary, multi-based Heegaard triple. We will call a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ a **doubly multi-pointed Heegaard triple** if each component of $\Sigma \setminus \boldsymbol{\tau}$ has exactly one $\mathbf{w}$ basepoint and one $\mathbf{z}$ basepoint, for each $\boldsymbol{\tau} \in \{\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}\}$. Using the construction from [OS04e], to a general Heegaard triple we can construct a 4–manifold $X_{\alpha\beta\gamma}$ with three ends. In Chapter 5, the author describes an oriented, properly embedded surface with boundary $\Sigma_{\alpha\beta\gamma}$ inside of $X_{\alpha\beta\gamma}$ for a doubly multi-pointed Heegaard triple. The construction of Ozsváth and Szabó from [OS06] yields a map

$$\mathfrak{s}_{\mathbf{w}} : \pi_2(\mathbf{x}, \mathbf{y}, \mathbf{z}) \to \mathrm{Spin}^c(X_{\alpha\beta\gamma}), \tag{6.10}$$

analogous to the map in Equation (6.9). One can similarly define a map $\mathfrak{s}_{\mathbf{z}}$. The set $\mathrm{Spin}^c(X_{\alpha\beta\gamma})$ is an affine space over $H^2(X_{\alpha\beta\gamma}; \mathbb{Z}) \cong H_2(X_{\alpha\beta\gamma}, \partial X_{\alpha\beta\gamma}; \mathbb{Z})$. The two maps $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ are related by the following formula:

Lemma 6.3.4 (Lemma 5.3.3)**.** *If* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ *is a doubly multi-pointed Heegaard triple, and* ψ *is a homotopy class of triangles, then*

$$\mathfrak{s}_{\mathbf{w}}(\psi) - \mathfrak{s}_{\mathbf{z}}(\psi) = PD[\Sigma_{\alpha\beta\gamma}].$$

We note that if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ represents surgery on a framed 1–dimensional link, then the map $\mathfrak{s}_{\mathbf{w}}$ from Equation (6.9) can be viewed as a special case of the map $\mathfrak{s}_{\mathbf{w}}$ from Equation (6.10). This follows since by [OS06, Proposition 4.3] the 4–manifold $W(Y, \mathbb{S}_1)$ is obtained from $X_{\alpha\beta\beta'}$ by filling in one of the boundary components with 3–handles and 4–handles, and hence the inclusion of $X_{\alpha\beta\beta'}$ into $W(Y, \mathbb{S}_1)$ induces an isomorphism $\mathrm{Spin}^c(X_{\alpha\beta\beta'}) \cong \mathrm{Spin}^c(W(Y, \mathbb{S}_1))$.

Another important component of the construction of the maps in Chapter 4 are maps for adding or removing basepoints on a link component. These are called the "quasi-stabilization maps". The quasi-stabilization operation was first considered in [MO10]. The author considered them further in Chapter 3, proving several useful results. Notably, according to Theorem A, if $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link such that $w \notin \mathbf{w}$ and $z \notin \mathbf{z}$, then there are chain maps[1]

$$S^+_{w,z},\, T^+_{w,z} : \mathcal{C}FL^\infty(Y, L, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\infty(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \mathfrak{s})$$
$$S^-_{w,z},\, T^-_{w,z} : \mathcal{C}FL^\infty(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \mathfrak{s}) \to \mathcal{C}FL^\infty(Y, L, \mathbf{w}, \mathbf{z}, \mathfrak{s}),$$

which are natural, i.e. commute with the change of diagrams maps, appropriately.

The quasi-stabilization maps are the link cobordism maps for the decorated link cobordisms shown in Figure 6.3. We write $S^\pm_{w,z}$ and $T^\pm_{w,z}$ when w comes after z with respect to the links orientation. If z comes after w, we write instead $S^\pm_{z,w}$ and $T^\pm_{z,w}$. Since the orientation of a link is fixed, it doesn't make sense to compare $S^\pm_{w,z}$ and $S^\pm_{z,w}$ for a single pair of basepoints w and z. Nonetheless, quasi-stabilizations of the form $S^\pm_{w,z}$ can be related

[1]For $S^\pm_{w,z}$ and $T^\pm_{w,z}$ to be chain maps, there is an additional requirement on the coloring. If z and z' are the $\mathbf{z}$ basepoints adjacent to w, by looking at the quasi-stabilized differential (Equation (6.16), below) one can verify that $S^\pm_{w,z}$ is a chain map iff the coloring identifies the variables V_z and $V_{z'}$. Note that this is equivalent to the condition that the coloring on the complexes is induced by the coloring of the decorated link cobordism shown in Figure 6.3. Similarly if z is adjacent to the basepoints w and w', then $T^\pm_{w,z}$ will be chain maps iff U_w and $U_{w'}$ are identified by the coloring. See Chapter 4 for more details

to quasi-stabilizations of the form $S^{\pm}_{z',w}$ in a natural way (they are related by a basepoint moving map; see Chapter 4). The identification is precisely what one would expect from the interpretation of the maps in terms of the surfaces with divides shown in Figure 6.3.

The final ingredient of the link cobordism maps are the band surgery maps. If $\mathbb{L}$ is an oriented link, and B is an oriented band which is appropriately positioned with respect to the basepoints, then we can perform band surgery to obtain a new link $\mathbb{L}(B)$ inside of Y. There are two band maps defined in Chapter 4 for band surgery, corresponding to adding a band which forms a region of type $\mathbf{z}$ or of type $\mathbf{w}$. They are denote $F^{\mathbf{w}}_B$ and $F^{\mathbf{z}}_B$ and they determine maps

$$F^{\mathbf{w}}_B,\ F^{\mathbf{z}}_B : \mathcal{C}FL^{\infty}(Y,\mathbb{L},\mathfrak{s}) \to \mathcal{C}FL^{\infty}(Y,\mathbb{L}(B),\mathfrak{s}).$$

As with the quasi-stabilization maps, there is a requirement on the coloring of the link $\mathbb{L}$ for the maps $F^{\mathbf{w}}_B$ and $F^{\mathbf{z}}_B$ to be well defined chain maps. The coloring requirement for $F^{\mathbf{w}}_B$ is simply that the $\mathbf{w}$ basepoints adjacent to the ends of B are given the same color, and the requirement for the map $F^{\mathbf{z}}_B$ is simply that the $\mathbf{z}$ basepoints adjacent to the ends of B are given the same color. Both maps are defined by counting holomorphic triangles in a Heegaard triple which respects the band.

6.3.4 Maslov gradings, Alexander gradings and link cobordisms

There are three gradings on link Floer homology, denoted

$$\mathrm{gr}_{\mathbf{w}}, \quad \mathrm{gr}_{\mathbf{z}}, \quad \text{and} \quad A.$$

They are described in [OS04c] and [OS08], using slightly different notation than we use here. See Chapter 5 for a description of the gradings in terms of link cobordisms, which is a convenient description for our purposes as it allows us to compute grading changes associated to decorated link cobordisms. The gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ are called the Maslov gradings, and A is the Alexander grading. If $\mathbf{x}$ and $\mathbf{y}$ are intersection points, then as relative gradings they

are defined by picking a homology class of disks $\phi \in \pi_2(\mathbf{x}, \mathbf{y})$ and using one of the formulas

$$\begin{aligned} \mathrm{gr}_{\mathbf{w}}(\mathbf{x}, \mathbf{y}) &= \mu(\phi) - 2n_{\mathbf{w}}(\phi) \\ \mathrm{gr}_{\mathbf{z}}(\mathbf{x}, \mathbf{y}) &= \mu(\phi) - 2n_{\mathbf{z}}(\phi) \\ A(\mathbf{x}, \mathbf{y}) &= n_{\mathbf{z}}(\phi) - n_{\mathbf{w}}(\phi). \end{aligned}$$

Here $n_{\mathbf{w}}(\phi)$ is defined to be the sum of the multiplicities of ϕ at all of the $\mathbf{w}$ basepoints, and $n_{\mathbf{z}}(\phi)$ is defined similarly. If $\mathfrak{s}$ is torsion, then $\mathrm{gr}_{\mathbf{w}}$ is independent of the disk ϕ, and if $\mathfrak{s} - PD[L]$ is torsion, $\mathrm{gr}_{\mathbf{z}}$ is independent of ϕ. If $[L] = 0 \in H_1(Y; \mathbb{Z})$, then A is independent of the disk ϕ. Absolute lifts of $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ can be specified as in [OS06] by using the absolute gradings on $CF^\infty(Y, \mathfrak{s})$ and $CF^\infty(Y, \mathfrak{s} - PD[L])$. With respect to $\mathrm{gr}_{\mathbf{w}}$, the $U_{\mathbf{w}}$ variables are -2 graded, and the $V_{\mathbf{z}}$ variables are 0 graded. With respect to $\mathrm{gr}_{\mathbf{z}}$, the $U_{\mathbf{w}}$ variables are 0 graded, and the $V_{\mathbf{z}}$ variables are -2 graded. With respect to the Alexander grading, the $U_{\mathbf{w}}$ variables are -1 graded, and the $V_{\mathbf{z}}$ variables are $+1$ graded.

As in [OS04c] or [OS08], for knots in integer homology spheres, an absolute lift of the Alexander grading can be specified by a symmetry requirement analogous to the Alexander polynomial of a knot. For links, there is an Alexander multi-grading, which yields the above Alexander grading after collapsing the multi-grading into a single grading. In general, it is possible to define an absolute lift of the Alexander grading under only the assumption the L is null-homologous, even when $\mathfrak{s}$ is non-torsion. However, when $\mathfrak{s}$ is non-torsion, an absolute lift of the grading requires a choice of Seifert surface S for L, and we write A_S for the Alexander grading using the Seifert surface S (see Chapter 5 for a description).

When $\mathfrak{s}$ is torsion and L is null-homologous, all three gradings are defined, and are related by the formula

$$A(\mathbf{x}) = \frac{1}{2}(\mathrm{gr}_{\mathbf{w}}(\mathbf{x}) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x})).$$

In Chapter 5, the author shows that if $(W, F) : (Y_1, K_1) \to (Y_2, K_2)$ is a decorated link cobordism and $\mathbf{x}$ is a homogeneously graded element, then if S_1 and S_2 are Seifert surfaces

for K_1 and K_2, then

$$A_{S_2}(F_{W,F,\mathfrak{s}}(\mathbf{x})) - A_{S_1}(\mathbf{x}) = \frac{\langle c_1(\mathfrak{s}), [\widehat{\Sigma}]\rangle - [\widehat{\Sigma}]\cdot[\widehat{\Sigma}]}{2} + \frac{\chi(\Sigma_{\mathbf{w}}) - \chi(\Sigma_{\mathbf{z}})}{2}, \tag{6.11}$$

$$\mathrm{gr}_{\mathbf{w}}(F_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{w}}(\mathbf{x}) = \frac{c_1(\mathfrak{s})^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{w}}) \tag{6.12}$$

and

$$\mathrm{gr}_{\mathbf{z}}(F_{W,F,\mathfrak{s}}(\mathbf{x})) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}) = \frac{c_1(\mathfrak{s} - PD[\Sigma])^2 - 2\chi(W) - 3\sigma(W)}{4} + \widetilde{\chi}(\Sigma_{\mathbf{z}}). \tag{6.13}$$

Here

$$\widehat{\Sigma} = (-S_1) \cup \Sigma \cup S_2,$$

$\chi(\Sigma_{\mathbf{w}})$ denotes the Euler characteristic, and $\widetilde{\chi}(\Sigma_{\mathbf{w}})$ denotes the reduced Euler characteristic:

$$\widetilde{\chi}(\Sigma_{\mathbf{w}}) = \chi(\Sigma_{\mathbf{w}}) - \tfrac{1}{2}(|\mathbf{w}_1| + |\mathbf{w}_2|).$$

The conjugation map ι_K is graded, in the following sense:

Lemma 6.3.5. *The map ι_K satisfies*

$$\mathrm{gr}_{\mathbf{w}}(\iota_K(\mathbf{x})) = \mathrm{gr}_{\mathbf{z}}(\mathbf{x}), \qquad \mathrm{gr}_{\mathbf{z}}(\iota_K(\mathbf{x})) = \mathrm{gr}_{\mathbf{w}}(\mathbf{x}), \qquad \textit{and} \qquad A(\iota_K(\mathbf{x})) = -A(\mathbf{x}),$$

for homogeneously graded $\mathbf{x}$.

Proof. As $A(\mathbf{x}) = \frac{1}{2}(\mathrm{gr}_{\mathbf{w}}(\mathbf{x}) - \mathrm{gr}_{\mathbf{z}}(\mathbf{x}))$, the third relation follows from the first two, so let us consider only the relations involving $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. The map ι_K is defined as a composition $\tau_K \circ \eta$. Since τ_K is a diffeomorphism map, it is grading preserving, so it remains only to consider the map η. The result then follows easily from the link cobordism description of the absolute gradings (see Proposition 5.10.1). □

6.4 Further properties of the link Floer TQFT

In this section we describe several important properties of the link cobordism maps from Chapter 4. Firstly, we prove Theorem 6.1.3, i.e. conjugation invariance for the link Floer

TQFT. Then we prove several important relations for the link cobordism maps, such as the bypass relation, as well as an interpretation of the maps Φ_w and Ψ_z in terms of dividing sets on cylindrical link cobordisms.

6.4.1 Conjugation invariance of the link cobordism maps

We now consider how the link cobordism maps defined in Chapter 4 interact with the conjugation map η, described in Section 6.3.2, and prove Theorem 6.1.3:

Theorem 6.1.3. *Suppose that $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a link cobordism with coloring $(\sigma, \mathfrak{P})$. Let $(W, \overline{F}) : (Y_1, \overline{\mathbb{L}}_1) \to (Y_2, \overline{\mathbb{L}}_2)$ denote the conjugate decorated link cobordism (but with the same coloring). The following diagram commutes up to filtered, $\mathbb{F}_2[U_{\mathfrak{P}}]$–equivariant chain homotopy:*

$$\begin{array}{ccc} \mathcal{C}FL^\infty(Y_1, \mathbb{L}_1, \mathfrak{s}_1) & \xrightarrow{\eta} & \mathcal{C}FL^\infty(Y_1, \overline{\mathbb{L}}_1, \overline{\mathfrak{s}}_1 + PD[L_1]) \\ \downarrow{\scriptstyle F_{W,F,\mathfrak{s}}} & & \downarrow{\scriptstyle F_{W,\overline{F},\overline{\mathfrak{s}}+PD[\Sigma]}} \\ \mathcal{C}FL^\infty(Y_2, \mathbb{L}_2, \mathfrak{s}_2) & \xrightarrow{\eta} & \mathcal{C}FL^\infty(Y_2, \overline{\mathbb{L}}_2, \overline{\mathfrak{s}}_2 + PD[L_2]). \end{array}$$

Proof. Using the construction from Chapter 4, the theorem is nearly a tautology.

Let us first consider the case that each component of Σ intersects both Y_1 and Y_2.

The construction from Chapter 4 involves taking a decomposition of the link cobordism (W, F), into elementary link cobordisms $(W_i, F_i) : (Y_i, L_i) \to (Y_i, L_{i+1})$. Each (W_i, F_i) satisfies one of the following:

1. W_i has a Morse function f such that f and $f|_{F_i}$ have no critical points, and the divides go from the incoming end to the outgoing end;

2. W_i has a Morse function with one index 1 or 3 critical point, or an arbitrary number of index 2 critical points, such that $f|_{F_i}$ has no critical points and the divides all go from the incoming end to the outgoing end;

3. W_i has a Morse function function f such that f and $f|_{F_i}$ have no critical points, all

dividing arcs go from the incoming boundary to the outgoing boundary, except for one arc, which has two ends on the incoming link, or two ends on the outgoing link;

4. W_i has a Morse function which has an index 1 critical point, which is an index 1 critical point of both f and $f|_{F_i}$. The dividing arcs go from the incoming to the outgoing boundary.

At most one elementary cobordism in the decomposition is allowed to have a Morse function f with index 2 critical points.

In the language of Chapter 4, this is essentially a decomposition into elementary parametrized cobordisms. However, in Chapter 4 we also considered a "parameterization", i.e. a diffeomorphism from $W(Y_i, \mathbb{S}_i)$ to W_i for a framed i–dimensional link in the incoming end of W_i. See Chapter 4 for a precise description. We suppress the parameterizations in our notation here. As the link cobordism maps are defined as a composition of maps associated to elementary cobordisms, as above, it is sufficient to show the claim for each of the cobordisms of one of the above types.

The maps associated to link cobordisms of type (1) correspond to the maps induced by diffeomorphisms, so for such link cobordisms the claim follows from naturality of the map η from Lemma 6.3.3.

The maps for the cobordisms of type (2) correspond to handle attachment maps. Consider first the 1–handle and 3–handle maps. Suppose that a diagram $\mathcal{H}$ has been surgered along a framed 0–sphere $\mathbb{S}_0$ to yield a diagram $\mathcal{H}(\mathbb{S}_0)$. There are two intersection points, θ^+ and θ^-, on $\mathcal{H}(\mathbb{S}_0)$ which are distinguished by the relative homological grading. In our notation, there are two homological gradings, $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$, however the designation of higher and lower relative grading is the same with respect to both $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. It is easy to see that η preserves this decomposition into higher and lower generators. In such a way we see that η commutes with the 1–handle and 3–handle maps.

The interaction of η with the 2–handle maps is slightly more subtle. Suppose $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a 2–handle cobordism for a framed 1–dimensional link $\mathbb{S}_1$ in Y_1.

The map $F_{W,F,\mathfrak{s}}$ is defined by counting holomorphic triangles in a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ which is subordinate to an α-bouquet[2] for $\mathbb{S}_1$. More precisely, for such a Heegaard triple, one defines

$$F_{W,F,\mathfrak{s}}(\mathbf{x}) = F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^+_{\alpha'\alpha}, \mathbf{x}),$$

the map which counts holomorphic triangles in $\Sigma \times \Delta$ (for an almost complex structure J) for representing Maslov index zero homology classes ψ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$ on the triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$. Here $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ is a diagram for an unlink in $(S^1 \times S^2)^{\#k}$ where each component has exactly two basepoints, and $\Theta^+_{\alpha'\alpha}$ is the element of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ of top $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ grading.

There is a bijection between holomorphic triangles on the triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ with the almost complex structure J, and holomorphic triangles on the conjugate triple $\overline{\mathcal{T}} = (-\Sigma, \overline{\boldsymbol{\beta}}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\alpha}}')$ with complex structure $-J$. Note also that $\overline{\mathcal{T}}$ is now subordinate to a β-bouquet for the framed link $\mathbb{S}_1$, whereas $\mathcal{T}$ was subordinate to a α-bouquet. We note that, as in [OS06, Lemma 5.2], we can use either Heegaard triples which are subordinate to an α-bouquet or a β-bouquet to compute the cobordism maps for surgery on a framed 1–dimensional link (see also Lemma 4.4.8).

However, there is a distinction in how Spin^c structures are assigned to homology classes of triangles on the triples $\mathcal{T}$ and $\overline{\mathcal{T}}$, since the roles of the $\mathbf{w}$ basepoints and the $\mathbf{z}$ basepoints has been switched. Note that the 4–manifolds $X_{\alpha'\alpha\beta}$ and $X_{\overline{\beta}\overline{\alpha}\alpha'}$, are canonically diffeomorphic, and diffeomorphism restricts to a diffeomorphism between the embedded surfaces $\Sigma_{\alpha'\alpha\beta}$ and $\Sigma_{\overline{\beta}\overline{\alpha}\alpha'}$. Write $\iota : X_{\alpha'\alpha\beta} \to W(Y_1, \mathbb{S}_1)$ for the embedding, which is well defined up to isotopy (see [OS06, Proposition 4.3] or [Juh16, Proposition 6.6]). Since $W(Y_1, \mathbb{S}_1)$ is formed by attaching 3–handles and 4–handles to $X_{\alpha'\alpha\beta}$, it follows that ι induces an isomorphism on $H^2(-;\mathbb{Z})$ and on the set of Spin^c structures. If $\mathfrak{t} \in \mathrm{Spin}^c(W)$, the map $F_{W,\overline{F},\mathfrak{t}}$ can be computed using the triple $\overline{\mathcal{T}}$ by counting holomorphic triangles representing homology classes $\overline{\psi}$ with

[2]A α-bouquet is a collections of arcs, one for each component of $\mathbb{S}_1$, each with one end on $\mathbb{S}_1$, and one end on a component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which is oriented to be going from a $\mathbf{w}$ basepoint to a $\mathbf{z}$ basepoint. See [OS06] for the original construction of Heegaard triples subordinate to a bouquet of a framed link.

$\mathfrak{s}_{\mathbf{z}}(\overline{\psi}) = \mathfrak{t}$. From Lemma 6.3.4, we have

$$\mathfrak{s}_{\mathbf{z}}(\overline{\psi}) = \mathfrak{s}_{\mathbf{w}}(\overline{\psi}) + PD[\Sigma_{\overline{\beta}\overline{\alpha}\alpha'}].$$

Since $\iota^* PD[\Sigma] = PD[\Sigma_{\overline{\beta}\overline{\alpha}\alpha'}]$, we conclude that

$$\eta \circ F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^+_{\alpha'\alpha}, \mathbf{x}) = F_{\overline{\beta}\overline{\alpha}\alpha',\overline{\mathfrak{s}}+\iota^* PD[\Sigma]}(\eta(\mathbf{x}), \Theta^+_{\overline{\alpha}\alpha'}).$$

Hence, for 2–handle maps we conclude that

$$\eta \circ F_{W,F,\mathfrak{s}} \simeq F_{W,\overline{F},\overline{\mathfrak{s}}+PD[\Sigma]} \circ \eta.$$

We now consider link cobordisms of type (3). The maps correspond to quasi-stabilization maps $S^+_{w,z}, S^-_{w,z}, T^+_{w,z}$ or $T^-_{w,z}$, depending on the configuration of $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$ regions with respect to the dividing set. The configurations of the regions, as well as the maps induced, are shown in Figure 6.3.

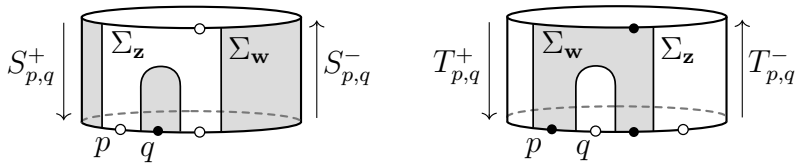

Figure 6.3: **Decorated link cobordisms for the quasi-stabilization maps.** The underlying undecorated link cobordisms consist of $L \times [0,1]$ inside of $Y \times [0,1]$.

Since the type of each region of $\Sigma \setminus \mathcal{A}$ is changed (from $\mathbf{w}$ to $\mathbf{z}$, or vice versa), it is sufficient to show that

$$S^{\pm}_{p,q} \circ \eta \simeq \eta \circ T^{\pm}_{p,q}, \qquad \text{and} \qquad T^{\pm}_{p,q} \circ \eta \simeq \eta \circ S^{\pm}_{p,q},$$

where $S^{\pm}_{p,q}$ and $T^{\pm}_{p,q}$ are the quasi-stabilization maps defined in Chapter 3 and Chapter 4. This is essentially a tautology from the definition. For example, the maps $S^+_{p,q}$ and $S^-_{p,q}$ are defined by the formulas

$$S^+_{p,q}(\mathbf{x}) = \mathbf{x} \times \theta^+_{\mathbf{w}}, \qquad S^-_{p,q}(\mathbf{x} \times \theta^+_{\mathbf{w}}) = 0, \qquad \text{and} \qquad S^-_{p,q}(\mathbf{x} \times \theta^-_{\mathbf{w}}) = \mathbf{x}, \tag{6.14}$$

where $\theta^+_{\mathbf{w}}$ is the top degree generator, with respect to the grading $\mathrm{gr}_{\mathbf{w}}$ induced by the $\mathbf{w}$ variables. Analogously the maps $T^+_{p,q}$ and $T^-_{p,q}$ are defined by the formulas

$$T^+_{p,q}(\mathbf{x}) = \mathbf{x} \times \theta^+_{\mathbf{z}}, \qquad T^-_{p,q}(\mathbf{x} \times \theta^+_{\mathbf{z}}) = 0, \qquad \text{and} \qquad T^-_{p,q}(\mathbf{x} \times \theta^-_{\mathbf{z}}) = \mathbf{x}. \tag{6.15}$$

Since η switches the type designation between the $\mathbf{w}$ and $\mathbf{z}$ basepoints, it sends $\theta^+_{\mathbf{w}}$ to $\theta^+_{\mathbf{z}}$. This is shown in Figure 6.4. Hence, we have $T^+_{p,q} \circ \eta \simeq \eta \circ S^+_{p,q}$. A similar argument shows that $T^-_{p,q} \circ \eta \simeq \eta \circ S^-_{p,q}$.

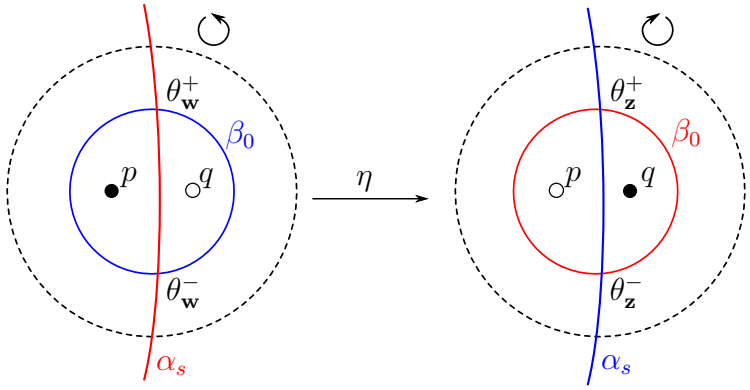

Figure 6.4: **The effect of conjugation on a quasi-stabilized diagram.** The map η sends the top degree intersection point with respect to $\mathrm{gr}_{\mathbf{w}}$ to the top degree intersection points with respect to $\mathrm{gr}_{\mathbf{z}}$. The outer dashed circle represents where the connected sum with the rest of the diagram takes place. The circular arrows indicate the orientation of the Heegaard surface.

We now consider the final type of elementary parametrized cobordism, type (4), corresponding to a framed 0–sphere along the link L. Such a map is a composition of a 1–handle map, a band map, and some diffeomorphism maps. We already know the 1–handle maps commute with η, so let's consider only the band maps. As defined in Chapter 4, there are two band maps, $F^{\mathbf{w}}_B$ and $F^{\mathbf{z}}_B$, both defined by counting holomorphic triangles. Supposing that B is an α–band (i.e. the ends of B are in segments of $\mathbb{L}$ going from $\mathbf{w}$ basepoints to $\mathbf{z}$ basepoints) then we can pick a Heegaard triple $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ representing band surgery of the link using the band. The band maps $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ are defined by

$$F^{\mathbf{z}}_B(\mathbf{x}) = F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{w}}_{\alpha'\alpha}, \mathbf{x}), \qquad \text{and} \qquad F^{\mathbf{w}}_B(\mathbf{x}) = F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{z}}_{\alpha'\alpha}, \mathbf{x})$$

for $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Here $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ is the top degree generator with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading on $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$, and $\Theta^{\mathbf{z}}_{\alpha'\alpha}$ is the top degree generator with respect to $\mathrm{gr}_{\mathbf{z}}$ (see Lemma 4.3.2). The map $F^{\mathbf{w}}_B$ corresponds to the case that $B \subseteq \Sigma_{\mathbf{w}}$ and the map $F^{\mathbf{z}}_B$ corresponds to the case that $B \subseteq \Sigma_{\mathbf{z}}$. Analogous to the quasi-stabilization maps, it is easy to see that η maps $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ to $\Theta^{\mathbf{z}}_{\overline{\alpha}\overline{\alpha}'}$. Paying attention to the change of Spin^c structures, as we did with the 2–handle maps, we see that

$$\eta(F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{w}}_{\alpha'\alpha}, \mathbf{x})) = F_{\overline{\beta}\overline{\alpha}\overline{\alpha}', \overline{\mathfrak{s}}+PD[\Sigma_{\overline{\beta}\overline{\alpha}\overline{\alpha}'}]}(\eta(\mathbf{x}), \Theta^{\mathbf{z}}_{\overline{\alpha}\overline{\alpha}'}).$$

However we note that the 4–manifold $X_{\alpha'\alpha\beta}$ is diffeomorphic to $Y \times [0,1]$ once we fill in $Y_{\alpha'\alpha}$ with some 3-handles and 4-handles. Hence

$$\mathrm{Spin}^c(X_{\alpha'\alpha\beta}) \cong \mathrm{Spin}^c(Y).$$

It is not hard to see that the cohomology class $PD[\Sigma_{\overline{\beta}\overline{\alpha}\overline{\alpha}'}]$ is equal to $\iota^* PD[L \times [0,1]]$, where ι is the inclusion $\iota : X_{\alpha'\alpha\beta} \hookrightarrow Y \times [0,1]$. Hence, under the above isomorphism of Spin^c structures $PD[\Sigma_{\overline{\beta}\overline{\alpha}\overline{\alpha}'}]$ acts by $PD[L]$. This is, however, just the correction in Spin^c structures due to η (from Lemma 6.3.2), and hence we conclude that

$$\eta \circ F^{\mathbf{w}}_B \simeq F^{\mathbf{z}}_B \circ \eta \qquad \text{and} \qquad F^{\mathbf{w}}_B \circ \eta \simeq \eta \circ F^{\mathbf{z}}_B.$$

Notice too that the α–band has turned into a β–band after conjugating, analogously to how the 2–handle maps change after conjugation in [OS06, Lemma 5.2]. However the α–band maps and the β–band maps are related by a basepoint moving map, according to Proposition 4.7.8, and hence either can be used to compute the link cobordism maps.

Combining these observations with the composition law, we conclude that

$$\eta \circ F_{W,F,\mathfrak{s}} \simeq F_{W,\overline{F},\overline{\mathfrak{s}}+PD[\Sigma]} \circ \eta,$$

for a link cobordism where each component of Σ intersects both Y_1 and Y_2 nontrivially.

For a link cobordism where some components of Σ don't intersect one or both of Y_1 or Y_2, the link cobordism maps are defined by puncturing the link cobordism at points along the

dividing set, adding extra copies of (S^3, U, w, z) to $(Y_1, \mathbb{L}_1)$ and $(Y_2, \mathbb{L}_2)$ using the 0–handle and 4–handle maps. As the conjugation map η commutes with the 0–handle and 4–handle maps, the statement thus follows for general link cobordisms. □

6.4.2 The maps Φ and Ψ

In this section we describe the maps Φ and Ψ which feature in Theorem 6.1.1, and prove that they are induced by two link cobordisms with relatively simple dividing sets.

In general, if $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multi-based link, we can define endomorphisms Φ_w and Ψ_z of $\mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$ for each $w \in \mathbf{w}$ and each $z \in \mathbf{z}$. They are defined by counting holomorphic disks on a diagram using the formulas

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

and

$$\Psi_z(\mathbf{x}) = V_z^{-1} \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} n_z(\phi) \# \widehat{\mathcal{M}}(\phi) U^{n_w(\phi)} V^{n_z(\phi)} \cdot \mathbf{y}.$$

When $\mathbb{L} = (K, w, z)$ is a doubly based knot, as in Theorem 6.1.1, we will often write Φ for Φ_w and Ψ for Ψ_z. In the case of a doubly based knot, these are the same maps considered in Section 6.2, in the context of ι_K–complexes, for the basis B consisting of the set of intersection points $\mathbf{x}$ with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$ on a given Heegaard diagram.

The maps Φ_w and Ψ_z, described above are induced by two relatively simple dividing sets on $(Y \times [0,1], L \times [0,1])$:

Lemma 6.4.1. *The link cobordism maps for $(Y \times [0,1], L \times [0,1])$ with the dividing sets shown in Figure 6.1 are filtered, chain homotopy equivalent to Φ_w and Ψ_z, respectively. This holds without any assumption on the number of basepoints on the link component, though if there are more than two basepoints, the remaining arcs of the dividing set (not shown in Figure 6.1) are vertical arcs from $L \times \{0\}$ to $L \times \{1\}$.*

Proof. Let us consider only the map Φ_w, as the map Ψ_z is handled in an analogous fashion. Let (W, F) denote the link cobordism on the left side of Figure 6.1. If there are more than two basepoints on the link component, then the link cobordism is a composition of two quasi-stabilization maps, and we can use the computation of the quasi-stabilized differential from Chapter 3 to compute Φ_w directly. To make this explicit, suppose that z is the basepoint immediately preceding w, and suppose $\mathcal{H}$ is a Heegaard diagram for the link with w and and z removed. Let $\widehat{\mathcal{H}}$ be a diagram obtained from $\mathcal{H}$ by quasi-stabilizing the link component at w and z, as in Figure 6.4. Write w' and z' for the basepoints adjacent to w and z, which are already on $\mathcal{H}$. By Proposition 3.5.3, the differential on the quasi-stabilized diagram $\widehat{\mathcal{H}}$ can be written as

$$\partial_{\widehat{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H}} \end{pmatrix}, \tag{6.16}$$

under the identification (of modules)

$$\mathcal{C}FL^\infty(\widehat{\mathcal{H}}) = \left(\mathcal{C}FL^\infty(\mathcal{H}) \otimes_{\mathbb{F}_2} \mathbb{F}_2[U_w, V_z, U_w^{-1}, V_z^{-1}]\right) \otimes_{\mathbb{F}_2} \langle \theta^+_{\mathbf{w}}, \theta^-_{\mathbf{w}} \rangle.$$

The above matrix notation is with respect to writing $\theta^+_{\mathbf{w}}$ as the first basis element, and $\theta^-_{\mathbf{w}}$ as the second. Using this, we see that only the upper right component has any terms involving nonzero powers of U_w, and furthermore

$$F_{W,F,\mathfrak{s}} = S^+_{w,z} S^-_{w,z} = \Phi_w,$$

using the formulas for $S^+_{w,z}$ and $S^-_{w,z}$ from Equation (6.14).

In the case where there are exactly two basepoints on the link component, to show the formula, our strategy is to move the dividing set around, so that $L \times \{t\}$ intersects the dividing set non-trivially for each $t \in [0, 1]$. Let ϕ be the diffeomorphism of (Y, L) which fixes all of the basepoints except w, and which moves w to a nearby point w', which is not a basepoint of the link. Let z' be between w and w', and assume that the ordering of these points is w then z' then w'. We can move the dividing set around, as in Figure 6.5.

Note that $F_{W,F,\mathfrak{s}}$ is chain homotopic to the composition $(\phi^{-1})_* \phi_* F_{W,F,\mathfrak{s}}$, while the composition $\phi_* F_{W,F,\mathfrak{s}}$ is the link cobordism map for the link cobordism consisting of the bottom two

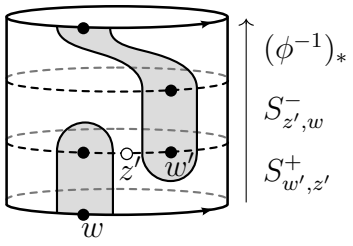

Figure 6.5: **Computing the link cobordism map for** Φ by manipulating the dividing set for the link cobordism from Figure 6.1. The induced link cobordism map in Lemma 6.4.1 is denoted $F_{W,F,\mathfrak{s}}$.

portions of the link cobordism in Figure 6.5. By construction, the induced link cobordism map is a quasi-stabilization followed by a quasi-destabilization. Hence

$$F_{W,F,\mathfrak{s}} \simeq (\phi^{-1})_*\phi_* F_{W,F,\mathfrak{s}} \simeq (\phi^{-1})_* S^-_{z',w} S^+_{w',z'}.$$

The diffeomorphism map $(\phi^{-1})_*$ also admits a convenient description in terms of quasi-stabilization maps. According to Section 4.3.5 we have $(\phi^{-1})_* \simeq T^-_{w',z'} S^+_{z',w}$. Putting this together, and using the relations from Section 4.3.4, we have

$$F_{W,F,\mathfrak{s}} \simeq (\phi^{-1})_* S^-_{z',w} S^+_{w',z'} \simeq T^-_{w',z'} S^+_{z',w} S^-_{z',w} S^+_{w',z'} \simeq T^-_{w',z'} \Phi_w S^+_{w',z'} \simeq T^-_{w',z'} S^+_{w',z'} \Phi_w \simeq \Phi_w.$$

The same strategy works for the map Ψ_z and the corresponding link cobordism. □

6.4.3 The bypass relation

In this section we prove Lemma 6.1.4, the bypass relation. The bypass relation encodes several relations between maps which appear on link Floer homology. As an application, we use it to give a quick proof of the formula for the action of a diffeomorphism corresponding to a Dehn twist around a link component.

Lemma 6.1.4. *If $(W,F), (W,F')$ and (W,F'') are decorated link cobordisms which fit into a bypass triangle, i.e. the underlying undecorated link cobordisms are equal, and the decorations differ only inside a disk on the surface, as shown in Figure 6.2, then*

$$F_{W,F,\mathfrak{s}} + F_{W,F',\mathfrak{s}} + F_{W,F'',\mathfrak{s}} \simeq 0.$$

Proof. We will interpret the relation

$$T^+_{w,z}S^-_{w,z} + S^+_{w,z}T^-_{w,z} + \mathrm{id} \simeq 0,$$

in terms of surfaces with divides. The above relation is proven in Lemma 4.3.13, and is immediate from the formulas for the maps (see Equations (6.14) and (6.15), and note that $\theta^{\pm}_{\mathbf{w}} = \theta^{\mp}_{\mathbf{z}}$, by definition). The dividing sets for these link cobordisms form a bypass triple on the undecorated link cobordism $(Y \times [0,1], L \times [0,1])$. Combining the above result with the composition law, the bypass relation follows for general link cobordisms. □

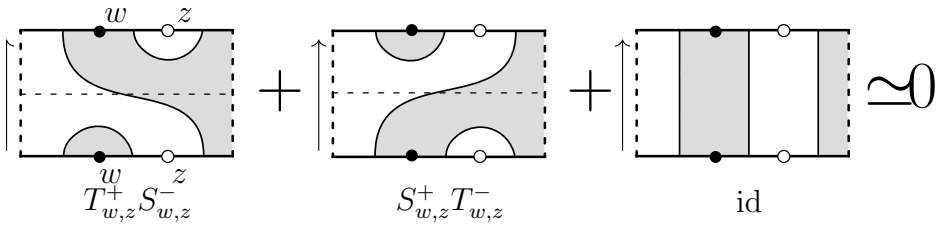

Figure 6.6: **The bypass triple formed by the dividing sets for the three summands in the relation** $T^+_{w,z}S^-_{w,z} + S^+_{w,z}T^-_{w,z} + \mathrm{id} \simeq 0$**.**

As a corollary, we can give a quick proof of Sarkar's formula for the diffeomorphism map induced by twisting a link component in one full twist, in the case that a link component has exactly two basepoints (see [Sar15] and Chapter 3). A similar pictorial argument using more bypass triples could be used in the case that the link component has more basepoints, but we will not consider that case.

Corollary 6.4.2. *Suppose* $\mathbb{L}$ *is a multi-based link in* Y*, and* K *is a component of* $\mathbb{L}$ *with exactly two basepoints. If* ρ *denotes the diffeomorphism resulting from twisting* K *in one full twist, in the direction of its orientation, then the induced map*

$$\rho_* : \mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^\infty(Y, \mathbb{L}, \mathfrak{s})$$

has the filtered chain homotopy type

$$\rho_* \simeq \mathrm{id} + \Phi_w \circ \Psi_z,$$

where w *and* z *are the two basepoints on* K*.*

Proof. The key is to use the bypass triple shown in Figure 6.7, on decorated link cobordisms with underlying undecorated link cobordism $(Y \times [0,1], L \times [0,1])$. We use Lemma 6.4.1 to interpret the resulting dividing sets as maps, and the formula follows.

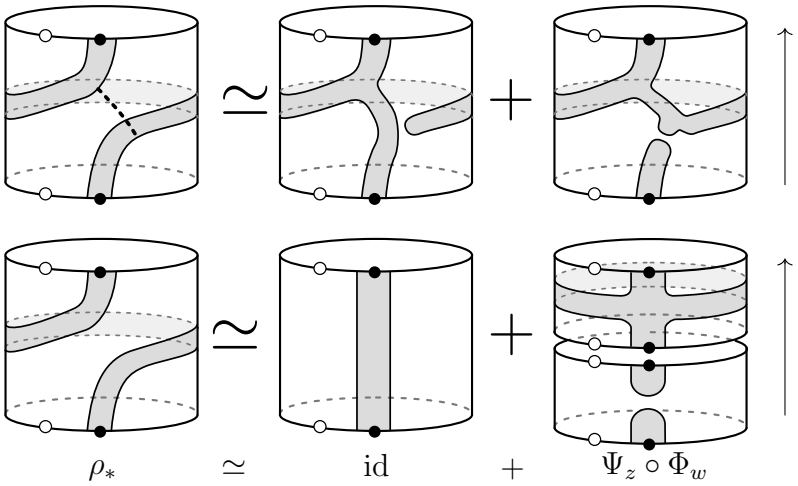

Figure 6.7: **Obtaining the formula $\rho_* \simeq \mathrm{id} + \Psi \circ \Phi$ from the bypass relation.** The bypass region is a neighborhood of the dashed arc in the top left. On the bottom row the same dividing sets have been manipulated to demonstrate the desired relation. We view the dividing sets as cobordisms from the bottom to the top. These are dividing sets on the surfaces $L \times [0,1]$ inside of $Y \times [0,1]$.

□

6.5 Link cobordisms and connected sums

In this section, we show that the link cobordism maps from Chapter 4 yield chain homotopy equivalences between the link Floer homology of a disjoint union and the link Floer homology of the connected sum (providing an alternate proof of the connected sum formula from [OS04c]).

If $\mathbb{K}_1 = (K_1, w_1, z_1)$ is a null-homologous knot in Y_1, and $\mathbb{K}_2 = (K_2, w_2, z_2)$ is a null-homologous knot in Y_2, we will define two link cobordism maps

$$G_1, G_2 : \mathcal{C}FL^\infty(Y_1 \sqcup Y_2, \mathbb{K}_1 \sqcup \mathbb{K}_2, \mathfrak{s}_1 \sqcup \mathfrak{s}_2) \to \mathcal{C}FL^\infty(Y_1 \# Y_2, \mathbb{K}_1 \# \mathbb{K}_2, \mathfrak{s}_1 \# \mathfrak{s}_2)$$

as well as two maps, E_1 and E_2, in the opposite direction. Here $\mathbb{K}_1\#\mathbb{K}_2 = (K_1\#K_2, w, z)$ is the connected sum, with exactly two basepoints. The main goal of this section is to show that $E_i \circ G_i \simeq \mathrm{id}$ and $G_i \circ E_i \simeq \mathrm{id}$ through filtered, equivariant, chain homotopies.

6.5.1 The decorated link cobordisms for connected sum maps

We now describe the link cobordisms which we will use to induce the chain homotopy equivalences E_1, G_1, E_2 and G_2. Briefly, the link cobordisms are simply obtained by adding a 1–handle and a band across the 1–handle, but there is an interesting ambiguity in the dividing set which is important for our connected sum formula, leading to four maps, instead of two. The four maps are depicted in Figure 6.8.

More explicitly, the link cobordisms are constructed as follows. Pick points $p_1 \in K_1 \setminus \{w_1, z_1\}$ and $p_2 \in K_2 \setminus \{w_2, z_2\}$ where the connected sum will take place. Also pick coordinate balls $D^3 \times \{1\}$ and $D^3 \times \{-1\}$, centered at p_1 and p_2 respectively. Assume further that

$$K_1 \cap (D^3 \times \{1\}) = \{(0, 0, t) : t \in [-1, 1]\} \times \{1\} \subseteq D^3 \times \{1\},$$

and similarly for K_2. Here we view D^3 as the closed unit ball in $\mathbb{R}^3$. We can form $Y_1\#Y_2$ by removing interior of the balls and gluing in $S^2 \times [-1, 1]$. Inside of $Y_1\#Y_2$ we can form the connected sum $K_1\#K_2$ by gluing in $\{(0, 0, \pm 1)\} \times [-1, 1]$. We assume further that the embeddings of $D^3 \times \{\pm 1\}$ are chosen so that $Y_1\#Y_2$ and $K_1\#K_2$ can be oriented compatibly with the orientations of Y_1 and Y_2.

Write W for the 1–handle cobordism obtained by gluing $D^3 \times [-1, 1]$ to $(Y_1 \sqcup Y_2) \times [0, 1]$ using the embeddings of $D^3 \times \{\pm 1\}$ into Y_i. Inside of W, we construct the surface

$$\Sigma = ((K_1 \sqcup K_2) \times [0, 1]) \cup B,$$

where B is the band

$$B = \{(0, 0, t) : t \in [-1, 1]\} \times [-1, 1] \subseteq D^3 \times [-1, 1].$$

There is also a 3–handle cobordism in the opposite direction, for which we write W'.

We now decorate Σ with a dividing set. However the decoration depends on the positions of p_1 and p_2 with respect to the basepoints. There are two natural configurations which we will consider:

($C1$) p_1 is in a component of $K_1 \setminus \{w_1, z_1\}$ going from w_1 to z_1, while p_2 is in a component of $K_2 \setminus \{w_2, z_2\}$ going from z_2 to w_2;

($C2$) p_1 is in a component of $K_1 \setminus \{w_1, z_1\}$ going from z_1 to w_1, while p_2 is in a component of $K_2 \setminus \{w_2, z_2\}$ going from w_2 to z_2.

Given a pair of points p_1 and p_2 which together satisfy either ($C1$) or ($C2$), we can decorate the surface Σ by adding three dividing arcs, as follows. One arc goes from $\{p_1\} \times \{0\}$ to $\{p_1\} \times \{1\}$ along the path $\{p_1\} \times \{t\}$, then travels across the core of the band, and then from $\{p_2\} \times \{1\}$ to $\{p_2\} \times \{0\}$ along $\{p_2\} \times \{t\}$. The other two arcs of the dividing set are obtained by picking points $p_i' \in K \setminus \{w_i, z_i\}$ in the components opposite to p_i, and defining the arc to be $\{p_i'\} \times [0,1]$. Writing $\mathcal{A}$ for the dividing set consisting of these three divides, the set $\Sigma \setminus \mathcal{A}$ has two components. To $K_1 \# K_2 \subseteq Y_1 \# Y_2$ we add two basepoints, w and z, in the complement of the divides, in such a way which is compatible with the basepoints w_1, w_2, z_1 and z_2.

Write $F_1 = (\Sigma, \mathcal{A})$ for a surface with divides inside of the 1–handle cobordism W, using connected sum points $p_1 \in K_1$ and $p_2 \in K_2$ satisfying the configuration condition ($C1$). Write F_1' for the surface with divides in the 3–handle cobordism W' constructed, using connected sum points p_1 and p_2 which satisfy ($C1$). The link cobordism (W', F_1') is obtained by turning around and reversing the orientation of (W, F_1). We then define

$$G_1 = F_{W,F_1,\mathfrak{s}}, \qquad \text{and} \qquad E_1 = F_{W',F_1',\mathfrak{s}},$$

where $\mathfrak{s}$ is the unique Spin^c structure on the 1–handle or 3–handle cobordism which extends $\mathfrak{s}_1 \sqcup \mathfrak{s}_2$.

Naturally, we would like to define the maps E_2 and G_2 to be the link cobordism maps for the link cobordism maps constructed when the points p_1 and p_2 satisfy configuration

condition $(C2)$. This however, makes for an awkward comparison with the maps E_1 and G_2, since moving the connected sum point changes both the 3–manifold $Y_1\#Y_2$ and the knot $K_1\#K_2$. Instead we define

$$G_2 = \tau_{K_1\#K_2}^{-1} \circ F_{W,\overline{F}_1,\mathfrak{s}} \circ (\tau_{K_1}|\tau_{K_2}), \qquad \text{and} \qquad E_2 = (\tau_{K_1}^{-1}|\tau_{K_2}^{-1}) \circ F_{W',\overline{F}'_1,\mathfrak{s}} \circ \tau_{K_1\#K_2} \quad (6.17)$$

where τ_K denotes the half twist diffeomorphism of the knot K (in the direction of the knots orientation, switching the two basepoints). Also $\overline{F}_1$ and $\overline{F}'_1$ denote the conjugate link cobordisms of F_1 and F'_1 (obtained by switching the designation of regions as type $\mathbf{w}$ or type $\mathbf{z}$).

Note that (W, F_1) is not diffeomorphic as a decorated link cobordism to (W, F_2), since the order that Y_1, Y_2 and $Y_1\#Y_2$ appear along the boundary of $\Sigma_{\mathbf{w}}$ is reversed for (W, F_1) and (W, F_2).

The link cobordisms for E_1, G_1, E_2 and G_2 are shown in Figure 6.8.

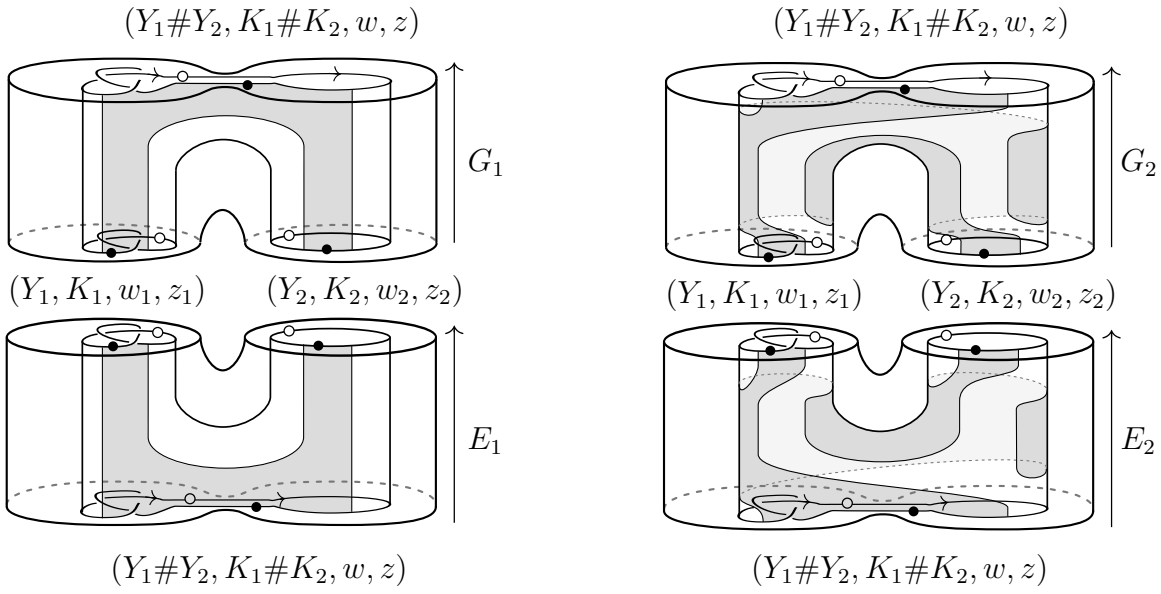

Figure 6.8: **The link cobordisms used to define the maps E_1, G_1, E_2 and G_2.** See Equation (6.17) for the definition of G_2 and E_2.

Proposition 6.5.1. *As endomorphisms of $\mathcal{C}FL^\infty(Y_1, \mathbb{K}_1) \otimes_{\mathcal{R}} \mathcal{C}FL^\infty(Y_2, \mathbb{K}_2)$, one has*

$$E_1 \circ G_1 \simeq E_2 \circ G_2 \simeq \mathrm{id}$$

through filtered, equivariant chain homotopies. As endomorphisms of $\mathcal{C}FL^\infty(Y_1\#Y_2, \mathbb{K}_1\#\mathbb{K}_2)$*, one has*

$$G_1 \circ E_1 \simeq G_2 \circ E_2 \simeq \mathrm{id},$$

through filtered, equivariant chain homotopies.

The proof of the above proposition will take several steps. The key observation is that can perform a 4–dimensional surgery operation to the link cobordisms inducing the compositions $E_1 \circ G_1$ and $G_1 \circ E_1$ which yield the identity link cobordism. Thus our strategy will be to prove two 4–dimensional surgery operations and their effect on the link cobordism maps.

6.5.2 Surgering a link cobordism on a 0–sphere

In this section, we describe the effect of taking the connected sum of two link cobordisms at two points along the dividing sets. We phrase this in terms of surgering the link cobordism $(W_1 \sqcup W_2, F_1 \sqcup F_2)$ on an embedded 0–sphere.

Let us define the link cobordism $(D^4 \times S^0, N_0)$ where D^4 is the unit ball in $\mathbb{R}^4$, $S^0 = \{\pm 1\}$ (oriented as the boundary of $[-1,1]$) and $N_0 = (\Sigma_0, \mathcal{A}_0)$ is the surface with divides defined by

$$\Sigma_0 = \{(x,y,0,0) : x^2 + y^2 \leq 1\} \times S^0, \qquad \text{and} \qquad \mathcal{A}_0 = \{(0,y,0,0) : -1 \leq y \leq 1\} \times S^0,$$

with designation of type **w** and type **z** regions given by

$$\Sigma_{0,\mathbf{w}} = (\Sigma_0 \cap \{x \geq 0\}) \times S^0, \qquad \text{and} \qquad \Sigma_{0,\mathbf{z}} = (\Sigma_0 \cap \{x \leq 0\}) \times S^0.$$

Also define the decorated link cobordism $(S^3 \times [-1,1], N_1)$ where $N_1 = (\Sigma_1, \mathcal{A}_1)$ is the surface with divides

$$\Sigma_1 = \{(x,y,0,0) : x^2 + y^2 = 1\} \times [-1,1]\}, \qquad \text{and} \qquad \mathcal{A}_1 = \{(0,\pm 1,0,0)\} \times [-1,1],$$

with designation of type **w** and type **z** regions given by

$$\Sigma_{1,\mathbf{w}} = \Sigma_1 \cap \{x \geq 0\}, \qquad \text{and} \qquad \Sigma_{1,\mathbf{z}} = \Sigma_1 \cap \{x \leq 0\}.$$

Given an orientation preserving embedding of $(D^4 \times S^0, N_0)$ into $(W_1 \sqcup W_2, F_1 \sqcup F_2)$ such that one component of $D^4 \times S^0$ is mapped into W_1 and the other component is mapped into W_2, we can form the connected sum, by removing $(D^4 \times S^0, N_0)$ and gluing in $(S^3 \times [-1,1], N_1)$, according to the embedding of $(D^4 \times S^0, N_0)$. We write $(W_1 \# W_2, F_1 \# F_2)$ for the surgerd link cobordism.

We will prove the following result about the effect of 4–dimensional surgery:

Proposition 6.5.2. *Suppose that (W_1, F_1) and (W_2, F_2) are two link cobordisms, with chosen points $p_1 \in \mathcal{A}_1 \subseteq W_1$ and $p_2 \in \mathcal{A}_2 \subseteq W_2$, as well as coordinate balls centered at p_1 and p_2, as above. If F_1 and F_2 are colored so that the regions of $\Sigma_{\mathbf{w}}$ adjacent to p_1 and p_2 are given the same color, and the regions of $\Sigma_{\mathbf{z}}$ adjacent to p_1 and p_2 are given the same color, one has*

$$F_{W_1 \# W_2, F_1 \# F_2, \mathfrak{s}_1 \# \mathfrak{s}_2} \simeq F_{W_1 \sqcup W_2, F_1 \sqcup F_2, \mathfrak{s}_1 \sqcup \mathfrak{s}_2}.$$

The proof is based on the fact that the maps are invariant under puncturing a link cobordism at a point along the dividing set, as well as a simple model computation which shows that we can connect up the punctures without changing the map.

Lemma 6.5.3. *Let $(W_0, F_0) : (S^3, \mathbb{U}) \to (S^3, \mathbb{U}_1) \sqcup (S^3, \mathbb{U}_2)$ be the link cobordism obtained by splitting an unknot in S^3 into two unknots, and then adding a 3–handle which separates the two unknots (corresponding to the map E_i shown in Figure 6.8, when $Y_1 = Y_2 = S^3$ and $K_1 = \mathbb{U}_1$ and $K_2 = \mathbb{U}_2$). The link cobordism (W_0, F_0) induces a map*

$$F_{W_0, F_0, \mathfrak{s}_0} : \mathcal{C}FL^\infty(S^3, \mathbb{U}) \to \mathcal{C}FL^\infty(S^3 \sqcup S^3, \mathbb{U}_1 \sqcup \mathbb{U}_2)$$

which is filtered chain homotopic to

$$1 \mapsto 1,$$

under the identification of both the domain and range as $\mathbb{F}_2[U, V, U^{-1}, V^{-1}]$.

Proof. It is straightforward to explicitly compute the map $F_{W_0, F_0, \mathfrak{s}_0}$ in terms of a quasi-stabilization, followed by a band map, and finally followed by a 3–handle map. It is easier, however, to simply use formal properties of the link Floer TQFT. Note that both the domain

and range of $F_{W_0,F_0,\mathfrak{s}_0}$ are filtered chain homotopy equivalent to $\mathcal{R} = \mathbb{F}_2[U, V, U^{-1}, V^{-1}]$, with vanishing differential. The cobordism map, being an equivariant map, is thus determined by its value on $1 \in \mathcal{R}$. However we can cap off either copy of $(S^3, \mathbb{U}_i)$ with a 4–ball, and the result must be the identity map on $\mathcal{R}$. This forces the image of $1 \in \mathcal{C}FL^\infty(S^3, \mathbb{U})$ in $\mathcal{C}FL^\infty(S^3 \sqcup S^3, \mathbb{U}_1 \sqcup \mathbb{U}_2)$ to also be 1, completing the proof. □

We can now prove the connected sum formula for the link cobordism maps:

Proof of Proposition 6.5.2. Using the Spin^c composition law for link cobordisms, it is sufficient to show the claim in the case that (W_1, F_1) and (W_2, F_2) are the identity cobordisms. Decompose the link cobordism representing $(W_1 \# W_2, F_1 \# F_2)$ as a 0–handle, followed by the cobordism considered in Lemma 6.5.3, followed by the link cobordisms (W_1, F_1) and (W_2, F_1), each with a puncture removed at a point along the divides. This is shown in Figure 6.9.

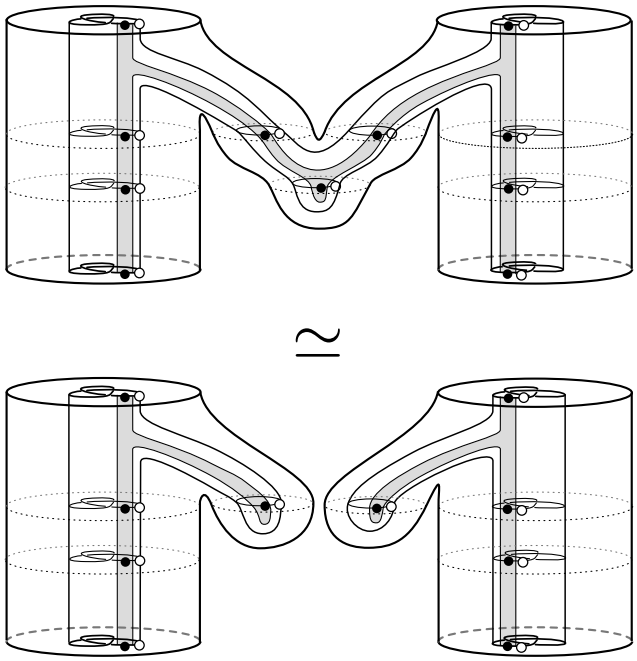

Figure 6.9: **The link cobordisms** $(W_1 \sqcup W_2, F_1 \sqcup F_2)$ **and** $(W_1 \# W_2, F_1 \# F_2)$ **considered in Proposition 6.5.2.** Using Lemma 6.5.3 and the composition law, we see that the two cobordism maps are equal.

Using the computation from Lemma 6.5.3, the middle portion of the cobordism

$(W_1 \# W_2, F_1 \# F_2)$ obtained by attaching a 0–handle, quasi-stabilizing, attaching a band, and attaching a 3–handle can be replaced with two 0–handle attachments, without changing the cobordism map. By Lemma 4.12.3, removing a ball from W which intersects F along an arc of the dividing set (and adding the new copy of $(S^3, \mathbb{U})$ to one of the ends of W) doesn't affect the cobordism map. Hence the two cobordism maps are equal. □

6.5.3 Surgering a link cobordism on a 1–sphere

We now describe a second surgery relation for the link cobordism maps, for surgering on framed 1–spheres in the 4–manifold W.

If γ is an embedded 1–manifold in the interior of W and we are given an identification of a regular neighborhood of γ with $S^1 \times D^3$, we can remove a neighborhood $S^1 \times D^3$ of γ and glue in a copy of $D^2 \times S^2$. Let us call the surgered 4–manifold $W(\gamma)$. Our goal is to relate the invariants for W to those for the surgered 4–manifold $W(\gamma)$. To do this, we need to assume that γ is embedded in the surface Σ, and in fact that γ is one of the divides in $\mathcal{A}$. In this case, we can simultaneously surger both W and F to get a new decorated link cobordism. We describe precisely the procedure, below.

We define two link cobordisms, $(S^1 \times D^3, M_1)$ and $(D^2 \times S^2, M_2)$, both from $(\varnothing, \varnothing)$ to a link in $S^1 \times S^2$. The decorated surfaces M_1 and M_2 are shown in Figure 6.10. Write $M_i = (\Sigma_i, \mathcal{A}_i)$. The underlying surface Σ_1 is the annulus

$$\Sigma_1 = S^1 \times \{(y, 0, 0) : -1 \leq y \leq 1\}.$$

The underlying surface Σ_2 is the pair of disks

$$\Sigma_2 = D^2 \times \{(\pm 1, 0, 0)\}.$$

Notice that

$$\partial \Sigma_1 = \partial \Sigma_2 = S^1 \times \{(\pm 1, 0, 0)\} \subseteq S^1 \times S^2.$$

Write $\mathbb{L}_0$ for this oriented link, with two basepoints per component, oriented as the boundary

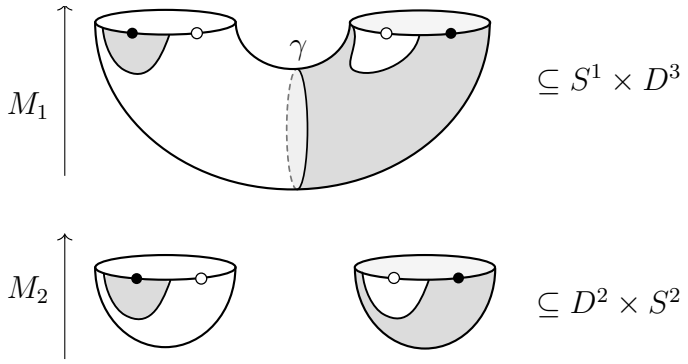

Figure 6.10: **The decorated surfaces M_1 and M_2.** The maps f and g in Lemma 6.5.7 are the link cobordism maps induced by $(S^1 \times D^3, M_1)$ and $(D^2 \times S^2, M_2)$, respectively, viewed as link cobordisms from $(\varnothing, \varnothing)$ to $(S^1 \times S^2, \mathbb{L}_0)$. The closed curve γ (which we think of as the curve which we perform surgery on) is shown on M_1.

of the surfaces Σ_1 and Σ_2. Note that $\mathbb{L}_0$ is null-homologous in the sense that the total homology class is zero, since the two components go in opposite directions around $S^1 \times S^2$.

Given a Spin^c structure on W, the restriction to $W \setminus N(\gamma)$ is torsion on $\partial N(\gamma) = S^1 \times S^2$. By considering the Mayer-Vietoris long exact sequence for cohomology, the obstruction to gluing a Spin^c structure on $W \setminus N(\gamma)$ and a Spin^c structure on $D^2 \times S^2$ lies in $H^2(S^1 \times S^2; \mathbb{Z})$ (i.e. the restriction of the two Spin^c structures to $S^1 \times S^2$). There is a unique Spin^c structure on $D^2 \times S^2$ which is torsion on $S^1 \times S^2$, which we can thus glue to the Spin^c structure $\mathfrak{s}|_{W \setminus N(\gamma)}$. The ambiguity in gluing lies in $\delta H^1(S^1 \times S^2; \mathbb{Z}) \subseteq H^2(W(\gamma); \mathbb{Z})$. By exactness of the Mayer-Vietoris sequence, the property that $\delta H^1(S^1 \times S^2; \mathbb{Z})$ vanishes is equivalent to the property that the restriction map $H^1(W \setminus N(\gamma); \mathbb{Z}) \to H^1(S^1 \times S^2; \mathbb{Z})$ is surjective. Surjectivity of the previous map is equivalent to the image of $[\gamma]$ in $H_1(W; \mathbb{Z})/\operatorname{Tors} \cong H_1(W \setminus N(\gamma); \mathbb{Z})/\operatorname{Tors}$, being non-divisible, i.e., having the property that if $n \cdot [\gamma'] = [\gamma]$ for some $[\gamma'] \in H_1(W; \mathbb{Z})/\operatorname{Tors}$, then $n = \pm 1$.

When $[\gamma]$ is non-divisible in $H_1(W; \mathbb{Z})/\operatorname{Tors}$, we will write $\mathfrak{s}(\gamma)$ for the Spin^c structure on $W(\gamma)$, as above.

Proposition 6.5.4. *Suppose (W, F) is a link cobordism with $F = (\Sigma, \mathcal{A})$ and γ is a closed curve on Σ such that the image $[\gamma] \in H_1(W; \mathbb{Z})/\operatorname{Tors}$ is non-divisible, and suppose that γ*

is one of the divides of $\mathcal{A}$. Further, suppose we are given a diffeomorphism of ϕ between a regular neighborhood $N(\gamma)$ of γ and the decorated link cobordism $(S^1 \times D^3, M_1)$. Writing $(W(\gamma), F(\gamma))$ for the surgered link cobordism (using ϕ to glue in $(D^2 \times S^2, M_2)$), we have

$$F_{W,F,\mathfrak{s}} \simeq F_{W(\gamma),F(\gamma),\mathfrak{s}(\gamma)},$$

where $\mathfrak{s}(\gamma)$ is as above.

Remark 6.5.5. One way to think of the above proposition is that it gives a way of exchanging handles of different indices. The 4–manifold $S^1 \times D^3$ can be decomposed as a 0–handle followed by a 1–handle, while $D^2 \times S^2$ can be decomposed by a 0–handle followed by a 2–handle. We can think of the above proposition as a way of exchanging a 1–handle for a 2–handle, or vice-versa, assuming that we also surger the surface with divides as well.

Remark 6.5.6. The requirement that $[\gamma] \in H_1(W;\mathbb{Z})/\operatorname{Tors}$ is non-divisible can be relaxed, though $\mathfrak{s}(\gamma)$ is no longer uniquely specified. Instead, the following relation holds:

$$F_{W,F,\mathfrak{s}} \simeq \sum_{\substack{\mathfrak{s}' \in \mathrm{Spin}^c(W(\gamma)) \\ \mathfrak{s}'|_{W \setminus N(\gamma)} = \mathfrak{s}|_{W \setminus N(\gamma)}}} F_{W(\gamma),F(\gamma),\mathfrak{s}'}.$$

The main ingredient of the proof of the previous proposition is the following lemma:

Lemma 6.5.7. *For the decorated surfaces with divides M_1 and M_2, described above, the two maps*

$$F_{S^1 \times D^3, M_1, \mathfrak{t}}, \quad F_{D^2 \times S^2, M_2, \mathfrak{t}_0} : \mathcal{C}FL^\infty(\varnothing, \varnothing) \to \mathcal{C}FL^\infty(S^1 \times S^2, \mathbb{L}_0)$$

are filtered chain homotopy equivalent. Here $\mathfrak{t}$ and $\mathfrak{t}_0$ are unique Spin^c *structures on $S^1 \times D^3$ and $D^2 \times S^2$ which are torsion on $S^1 \times S^2$. Also the links and surfaces are colored so that there is one U variable for all of the $\mathbf{w}$ basepoints and regions, and one V variable for all of the $\mathbf{z}$ basepoints and regions.*

Proof. We will use formal properties of the link Floer TQFT. See Appendix 6.9 for an alternate, explicit proof by decomposing the link cobordisms into pieces, and counting holomorphic curves. A diagram for $(S^1 \times S^2, \mathbb{L}_0)$ is shown in Figure 6.11. Using the intersection

points labeled in that diagram, write

$$\mathcal{C}FL^{\infty}(S^1 \times S^2, \mathbb{L}_0, \mathfrak{s}_0) = \langle \theta_1^+\theta_2^+, \theta_1^+\theta_2^-, \theta_1^-\theta_2^+, \theta_1^-\theta_2^- \rangle \otimes_{\mathbb{F}_2} \mathbb{F}_2[U, V, U^{-1}, V^{-1}],$$

where the designation into higher and lower intersection points is with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. Note that we are using a coloring which identifies the two $U_{\mathbf{w}}$–variables to U, and the two $V_{\mathbf{z}}$–variables to a single V.

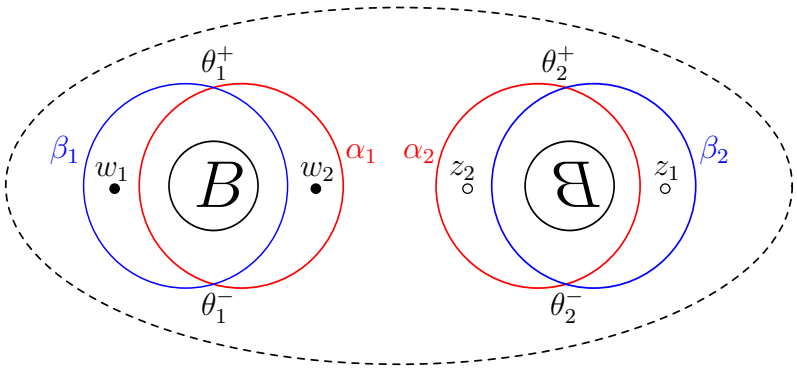

Figure 6.11: **A genus one diagram for** $(S^1 \times S^2, \mathbb{L}_0)$**.** The designation into higher and lower intersection points is with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

Using the grading change formulas from Equations (6.11), (6.12) and (6.13), one sees that both $F_{S^1\times D^3, M_1, \mathfrak{t}}$, $F_{D^2\times S^2, M_2, \mathfrak{t}_0}$ induce $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ grading changes of 0. For example, to compute the $\mathrm{gr}_{\mathbf{w}}$ grading changes of the two maps, one computes that

$$\widetilde{\chi}(\Sigma_{1,\mathbf{w}}) = 0, \qquad \text{and} \qquad \chi(S^1 \times D^3) = 0,$$

while

$$\widetilde{\chi}(\Sigma_{2,\mathbf{w}}) = 1, \qquad \text{and} \qquad \chi(D^2 \times S^2) = 2,$$

and that the other terms in the grading change formula vanish. Here $\Sigma_{i,\mathbf{w}}$ denotes the type $\mathbf{w}$ regions of Σ_i, the underlying surface of M_i.

Note that since $\mathcal{C}FL^{\infty}(\varnothing, \varnothing)$ is by definition equal to $\mathcal{R} = \mathbb{F}_2[U, V, U^{-1}, V^{-1}]$, it is sufficient to show that the value of the two maps on $1 \in \mathcal{R}$ are equal. Since the maps are filtered, they map $1 \in \mathcal{R}$ into the subcomplex $\mathcal{C}FL^{-}(S^1 \times S^2, \mathbb{L}_0, \mathfrak{s}_0)$. The set of elements of $\mathcal{C}FL^{-}(S^1 \times$

$S^2, \mathbb{L}_0, \mathfrak{s}_0)$ which are in both $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ grading zero is equal to the two dimensional vector space

$$W = \mathrm{Span}_{\mathbb{F}_2}(\theta_1^+\theta_2^-, \theta_1^-\theta_2^+).$$

As such, it is sufficient to restrict to the hat flavor, where we set $U = V = 0$. Let us write

$$f = F_{S^1\times D^3, M_1, \mathfrak{t}}, \qquad \text{and} \qquad g = F_{D^2\times S^2, M_2, \mathfrak{t}_0}.$$

Note that the differential on $\widehat{CFL}(S^1 \times S^2, \mathbb{L}_0, \mathfrak{s}_0)$ vanishes. Also, by using the bypass relation, we can write

$$f = A + B,$$

where A and B are the link cobordism maps for the link cobordisms shown in Figure 6.12, which both have the same underlying, undecorated link cobordism as $(S^1 \times D^3, M_1)$.

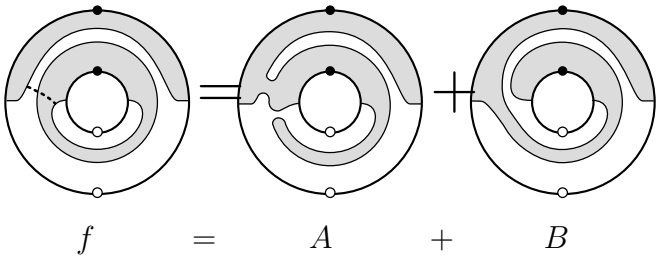

Figure 6.12: **The link cobordisms inducing the maps f, A and B, and a bypass relation between all three.** The bypass region is a neighborhood of the dashed line on the left. All three have underlying undecorated link cobordism $(W, \Sigma) = (S^1 \times D^3, S^1 \times \{(t, 0, 0) : -1 \leq t \leq 1\})$.

Let A' and B' be the link cobordism maps obtained by turning around and reversing the orientation of the link cobordisms for A and B (i.e. viewing them as link cobordisms from $(S^1 \times S^2, \mathbb{L}_0)$ to $\varnothing$). We claim that we have the following relations:

$$(A' \circ A)(1) = (B' \circ B)(1) = 0 \tag{6.18}$$

$$(A' \circ B)(1) = (B' \circ A)(1) = 1 \tag{6.19}$$

$$(A' \circ f)(1) = (A' \circ g)(1) = 1 \tag{6.20}$$

$$(B' \circ f)(1) = (B' \circ g)(1) = 1. \tag{6.21}$$

Before we prove the above relations, we will show that they are sufficient to show that $f = g$. Equations (6.18) and (6.19) imply that $A'|_W$ and $B'|_W$ are linearly independent and hence form a basis of the 2–dimensional vector space $W^\vee = \mathrm{Hom}_{\mathbb{F}_2}(W, \mathbb{F}_2)$. Equations (6.20) and (6.21) then imply that their values on $f(1)$ and $g(1)$ agree, so $f = g$.

Hence it remains only to establish Equations (6.18) – (6.21). This is done by interpreting each of the compositions as the link cobordism map for a link cobordism from $(\varnothing, \varnothing)$ to $(\varnothing, \varnothing)$. Notice that all compositions represent either a decorated S^2 which is unknotted in S^4, or a copy of $S^1 \times U$ inside of $S^1 \times S^3$ (where U denotes the unknot). These are shown in Figure 6.13. In Figure 6.13, the link cobordism on the right side of the top row is just a 0–handle followed by a 4–handle (each containing standard disks). The composition induces the identity map. The maps associated to the two link cobordisms on left side of the top row can be computed by using the map from Lemma 6.5.3 (splitting $(S^3, \mathbb{U})$ into two copies of itself), composed with the diffeomorphism associated to twisting one copy of the unknot in a full twist, composed with the link cobordism map from Lemma 6.5.3, turned upside down. This is clearly the identity. Finally the link cobordism shown on the bottom row, which induces the maps $A' \circ A$ and $B' \circ B$, induces the zero map, since by moving the dividing sets around we can factor the induced map through the map $\Phi|\Psi$ on $\widehat{CFL}(S^3, \mathbb{U}) \otimes_{\mathbb{F}_2} \widehat{CFL}(S^3, \mathbb{U})$, which trivially vanishes, since each of Φ and Ψ vanish for an unknot in S^3 with exactly two basepoints. Hence Equations (6.18)–(6.21) hold, completing the proof. □

Proof of Proposition 6.5.4. Note that the composition law implies that

$$F_{W,F,\mathfrak{s}} \simeq F_{W\setminus N(\gamma),F\setminus N(\gamma),\mathfrak{s}|_{W\setminus N(\gamma)}} \circ F_{S^1\times D^3,M_1,\mathfrak{t}}$$

which is equal by Lemma 6.5.7 to

$$F_{W\setminus N(\gamma),F\setminus N(\gamma),\mathfrak{s}|_{W\setminus N(\gamma)}} \circ F_{D^2\times S^2,M_2,\mathfrak{t}_0}.$$

The condition that $[\gamma] \in H_1(W;\mathbb{Z})/\operatorname{Tors}$ is non-divisible is equivalent to the condition that there is a unique Spin^c structure $\mathfrak{s}(\gamma)$ on $W(\gamma)$ which restricts to $\mathfrak{s}|_{W\setminus N(\gamma)}$ on $W \setminus N(\gamma)$ and

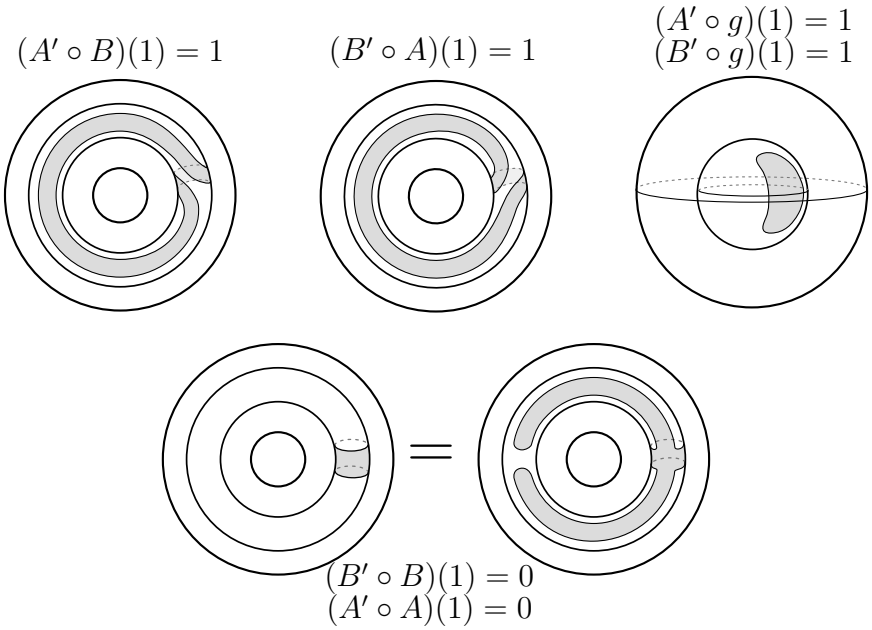

Figure 6.13: **Link cobordisms representing compositions appearing in Equations** (6.18)**,** (6.19)**,**(6.20) **and** (6.21)**.** The link cobordism in the top right consists of an unknotted S^2 inside of S^4. The remaining represent decorations of $S^1 \times U \subseteq S^1 \times S^3$.

to $\mathfrak{t}_0$ on $D^2 \times S^2$. Hence the composition law yields that the above composition is exactly $F_{W(\gamma),F(\gamma),\mathfrak{s}(\gamma)}$. □

6.5.4 Proof of Proposition 6.5.1

We now show that $E_i \circ G_i$ and $G_i \circ E_i$ are chain homotopic to the identity as consequences of the 4–dimensional surgery relations proven in the previous sections:

Proof of Proposition 6.5.1. The fact that $E_1 \circ G_1 \simeq \mathrm{id}$ follows from Proposition 6.5.2, since the composition of the link cobordisms defining E_1 and G_1 is a connected sum of two identity link cobordisms (in the sense of Proposition 6.5.2).

The relation $G_1 \circ E_1 \simeq \mathrm{id}$ is similarly proven using Proposition 6.5.4. Write (W, F) for the link cobordism inducing the map $G_1 \circ E_1$. There is a single divide γ which is a closed curve. After performing an isotopy of the other two dividing arcs, we can arrange so that a neighborhood of this dividing arc looks like the surface M_1 in Figure 6.10. After replacing

this copy of $(S^1 \times D^3, M_1)$ with $(D^2 \times S^2, M_2)$ we are left with the identity link cobordism from $(Y_1 \# Y_2, K_1 \# K_2)$ to itself. We briefly describe the diffeomorphism between $W(\gamma)$ and $(Y_1 \# Y_2) \times [0,1]$. Note that W has a handle decomposition given by a 3–handle followed by a "dual" 1–handle. We pull the 1–handle below the 3–handle. The handle decomposition of $W(\gamma)$ is obtained by exchanging the 1–handle for a 2–handle attached along a 0-framed unknot. However, the original 3–handle cancels this 2–handle, and we are left with a handle decomposition for $W(\gamma)$ which has no 4–dimensional handles.

Showing that this diffeomorphism restricts to a diffeomorphism between the embedded surfaces is a straightforward. To describe this in terms of handle operations, one picks a handle decomposition of the surface in the link cobordism for $G_1 \circ E_1$ with a band B which disconnects $K_1 \# K_2$, and a dual band B' which connects $K_1 \sqcup K_2$ to form $K_1 \# K_2$. One trades B for a disk D and two bands, B_1 and B_2. The band B' has both ends on D. The effect of the 4–dimensional surgery operation is to trade B' and D for two disks, D_1 and D_2. After removing the canceling 2–handle and 3–handle of the surgered 4–manifold, as described in the previous paragraph, the disks D_1 and D_2 then cancel B_1 and B_2, and we are left with the identity link cobordism.

Finally we note that the curve γ which we surger on represents a non-divisible element of $H_1(W;\mathbb{Z})/\operatorname{Tors}$ since W has a handle decomposition as a 1–handle followed by a 3–handle, and γ goes around the 1–handle exactly once. Hence by Proposition 6.5.4, the composition $G_1 \circ E_1$ is chain homotopy equivalent to the link cobordism map for the identity link cobordism from $(Y_1 \# Y_2, K_1 \# K_2)$ to itself.

Finally, to see that $G_2 \circ E_2$ and $E_2 \circ G_2$ are chain homotopy equivalent to the identity, we note by definition (see Equation (6.17)) the link cobordism maps for G_2 and E_2 are defined as a composition involving the conjugates of the link cobordisms used to define G_1 and E_1, as well as the half twist diffeomorphisms τ_{K_1}, τ_{K_2} and $\tau_{K_1 \# K_2}$. One can first cancel the adjacent pair of half twists in the middle of either composition $E_2 \circ G_2$ or $G_2 \circ E_2$, since they go in opposite directions. Then one simply uses Proposition 6.5.2 or 6.5.4 to surger the cobordism, leaving only a composition of two half twists on the identity cobordism, which go in opposite

directions. After canceling these, one is left with the identity cobordism. □

Remark 6.5.8. The 4–dimensional surgery formulas from Propositions 6.5.2 and 6.5.4 cannot be applied to the maps $E_i \circ G_j$ and $G_j \circ E_i$ if $i \neq j$. For example, the dividing set of the link cobordism corresponding to the compositions $G_2 \circ E_1$ does not contain a closed curve γ. Similarly the link cobordism corresponding to the composition $E_1 \circ G_2$ is not diffeomorphic to a connected sum of two link cobordisms (in the sense of Proposition 6.5.2), since in a connected sum, at most two arcs of the dividing set (the ones along which we perform the connected sum) can have an endpoint in both an end of W_1 and an end of W_2, while in the link cobordism for $E_1 \circ G_2$, all four divides have one endpoint in an end of W_1 and one endpoint in an end of W_2.

6.6 The map ι_K on connected sums

In this section, we prove Theorem 6.1.1:

Theorem 6.1.1. *Suppose $\mathbb{K}_1 = (K_1, w_1, z_1)$ and $\mathbb{K}_2 = (K_2, w_2, z_1)$ are two doubly based knots in Y_1 and Y_2 respectively, and let $\mathbb{K}_1 \# \mathbb{K}_2 = (K_1 \# K_2, w, z)$ denote their connected sum in $Y_1 \# Y_2$, with two basepoints. The two filtered chain homotopy equivalences*

$$G_1, G_2 : \mathcal{C}FL^\infty(Y_1 \sqcup Y_2, \mathbb{K}_1 \sqcup \mathbb{K}_2, \mathfrak{s}_1 \sqcup \mathfrak{s}_2) \to \mathcal{C}FL^\infty(Y_1 \# Y_2, \mathbb{K}_1 \# \mathbb{K}_2, \mathfrak{s}_1 \# \mathfrak{s}_2),$$

constructed in the last section, satisfy

$$G_1(\iota_{K_1}|\iota_{K_2} + \Phi_{w_1}\iota_{K_1}|\Psi_{z_2}\iota_{K_2}) \simeq (\iota_{K_1 \# K_2})G_1,$$

and

$$G_2(\iota_{K_1}|\iota_{K_2} + \Psi_{z_1}\iota_{K_1}|\Phi_{w_2}\iota_{K_2}) \simeq (\iota_{K_1 \# K_2})G_2.$$

Proof of Theorem 6.1.1. Let τ_{K_i} denote the half twist diffeomorphism on $(Y_i, \mathbb{K}_i)$, and let $\tau_{K_1 \# K_2}$ denote the half twist diffeomorphism on $(Y_1 \# Y_2, \mathbb{K}_1 \# \mathbb{K}_2)$. Let

$$(W, F_1) : (Y_1 \sqcup Y_2, \mathbb{K}_1 \sqcup \mathbb{K}_2) \to (Y_1 \# Y_2, \mathbb{K}_1 \# \mathbb{K}_2)$$

be the link cobordism described in Section 6.5.1. The decorations on the surface are redrawn in Figure 6.14. Recall that we defined

$$G_1 := F_{W,F_1,\mathfrak{s}}$$

to be the link cobordism map for the unique Spin^c structure $\mathfrak{s}$ on W which extends $\mathfrak{s}_1 \sqcup \mathfrak{s}_2$ on $Y_1 \sqcup Y_2$. Using the composition law, the composition $\tau_{K_1\#K_2} \circ G_1$ is also a link cobordism map, as shown in Figure 6.15.

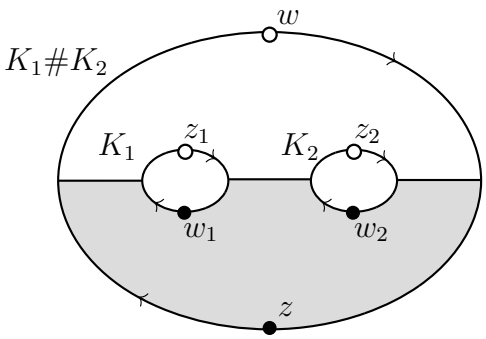

Figure 6.14: **The decoration on the surface used in the cobordism defining the map G_1.** This is the surface inside of the link cobordism from $K_1 \sqcup K_2$ to $K_1\#K_2$ shown in Figure 6.8. The arrows indicate the orientation of the knot components.

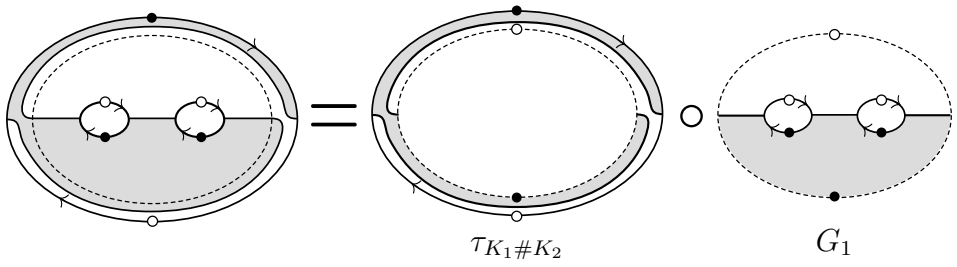

Figure 6.15: **The decoration on a link cobordism whose cobordism map is $\tau_{K_1\#K_2} \circ F_{W,F_1}$.** The arrows indicate the orientation of the knots.

We now use the bypass relation from Lemma 6.1.4 on the link cobordism for $\tau_{K_1\#K_2} \circ G_1$. The disk we pick for the bypass relation is a regular neighborhood of the dashed arc shown in Figure 6.16. The resulting three dividing sets, in a neighborhood of the dashed arc, are also shown in that figure.

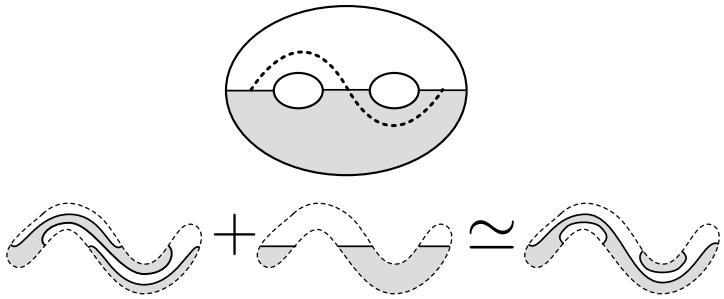

Figure 6.16: **A bypass relation.** We perform a bypass in a neighborhood of the dashed curve on top. The resulting bypass triple of dividing sets are drawn on the line below.

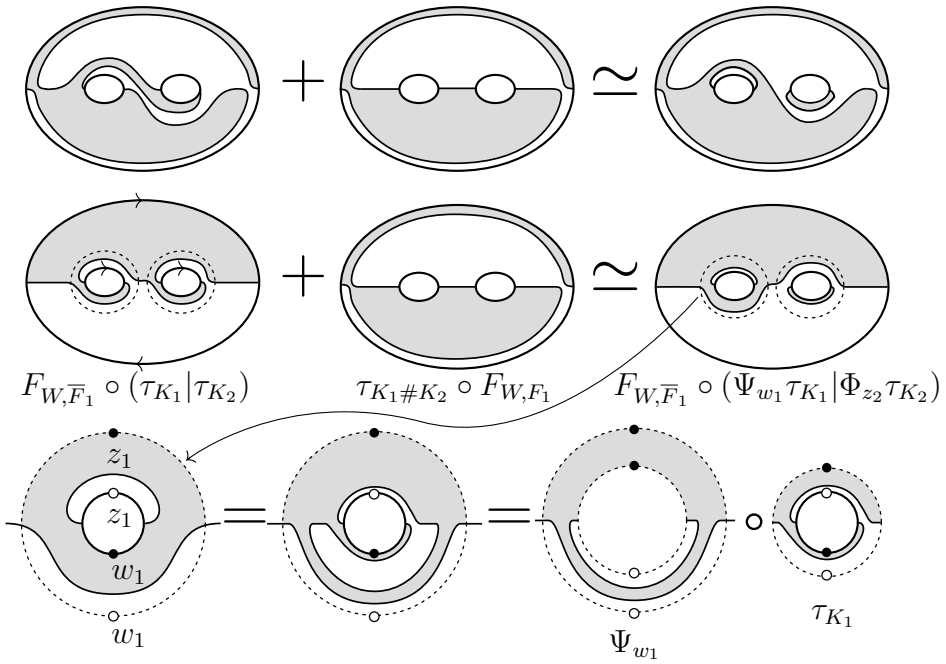

Figure 6.17: **Isotoping the dividing sets in the bypass relation, to derive the connected sum formula.** Composing with $\tau_{K_1\#K_2}$ and using the bypass relation for the bypass arc in Figure 6.16 yields the top row. Performing some manipulations yields the middle row. Here $(W, \overline{F}_1)$ is the link cobordism obtained by taking (W, F_1) and switching the type of each region from **w** to **z** and vice-versa. The cobordisms goes from the inner two circles to the outer circle. The bottom row is a zoomed-in manipulation of the one of the regions in the middle row, to illustrate the induced map.

In Figure 6.17, we compose with the half twist map $\tau_{K_1\#K_2}$, and perform some manipulations. From Figure 6.17, and Lemma 6.4.1 we see that

$$F_{W,\overline{F}_1} \circ \tau_{K_1 \sqcup K_2} + \tau_{K_1\#K_2} \circ F_{W,F_1} \simeq F_{W,\overline{F}_1} \circ (\Psi_{w_1}\tau_{K_1} | \Phi_{z_2}\tau_{K_2}).$$

Here F_1 is the surface with divides used to define G_1, and $\overline{F}_1$ is the conjugate surface with divides (i.e. the one resulting from changing the types of regions from $\mathbf{w}$ to $\mathbf{z}$, and vice versa). We now compose with the conjugation map, η, on the left, which yields

$$\eta \circ F_{W,\overline{F}_1} \circ \tau_{K_1 \sqcup K_2} + \eta \circ \tau_{K_1\#K_2} \circ F_{W,F_1} \simeq \eta \circ F_{W,\overline{F}_1} \circ (\Psi_{w_1}\tau_{K_1} | \Phi_{z_2}\tau_{K_2}).$$

Using Theorem 6.1.3, we have $\eta \circ F_{W,\overline{F}_1} \simeq F_{W,F_1} \circ \eta$. Since $\iota_K = \tau_K \circ \eta$, by definition, we arrive at

$$F_{W,F_1} \circ (\iota_{K_1} | \iota_{K_2}) + \iota_{K_1\#K_2} \circ F_{W,F_1} \simeq F_{W,F_1} \circ (\Phi_{w_1}\iota_{K_1} | \Psi_{z_2}\iota_{K_2}).$$

Since $G_1 = F_{W,F_1}$ by definition, the first statement of the theorem now follows.

The second statement, involving the map G_2, follows from an easy modification of the above argument. □

6.7 The homomorphism $\mathcal{C} \to \mathfrak{I}_K$

We finally put together the pieces of our theory and describe the homomorphism $\mathcal{C} \to \mathfrak{I}_K$. For a knot $K \in S^3$, let $[\mathcal{C}FL^\infty(S^3, K)]$ denote the local equivalence class in $\mathfrak{I}_K$ determined by the tuple

$$(\mathcal{C}FL^\infty(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z), \partial, \mathbb{T}_\alpha \cap \mathbb{T}_\beta, \iota_K)$$

for a diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, w, z)$ of (S^3, K, w, z). Note that with no additional effort, we can extend this homomorphism to the group $\mathcal{C}_3^{\mathbb{Z}}$ consisting of pairs of integer homology spheres with knots, modulo integer homology concordance (i.e. integer homology cobordisms with a smoothly embedded, genus zero surface). The homomorphism $\mathcal{C} \to \mathfrak{I}_K$ factors through the natural map $\mathcal{C} \to \mathcal{C}_3^{\mathbb{Z}}$.

Theorem 6.1.5. *The map*

$$K \mapsto [\mathcal{C}FL^\infty(S^3, K)],$$

defines a homomorphism from the smooth concordance group $\mathcal{C}$ to $\mathfrak{I}_K$.

Proof. We first claim that the map $K \mapsto [\mathcal{C}FL^\infty(S^3, K)]$ is well defined on the smooth concordance group. If (W, F) is a concordance from K to K' (or more generally a homology concordance) then we consider the map $F_{W,F,\mathfrak{s}}$ for the unique Spin^c structure on $W = S^3 \times [0, 1]$. Note that by reversing the orientation and turning around (W, F), we get a map $F_{-W,-F,\mathfrak{s}}$ in the opposite direction. We claim that $F_{W,F,\mathfrak{s}}$ and $F_{-W,-F,\mathfrak{s}}$ induce local equivalences between $\mathcal{C}FL^\infty(S^3, K)$ and $\mathcal{C}FL^\infty(S^3, K')$.

Note that $\mathcal{C}FL^\infty(S^3, K_i)$ decomposes over Alexander gradings as $\bigoplus_{k\in\mathbb{Z}} CFK^\infty(S^3, K_i)\{k\}$ where we think of $CFK^\infty(S^3, K_i)$ as being concentrated in Alexander grading zero, and $\{k\}$ denotes a shift in the Alexander grading. Hence the homology of each $\mathcal{C}FL^\infty(S^3, K_i)$ decomposes as

$$\bigoplus_{k\in\mathbb{Z}} H_*(CFK^\infty(S^3, K_i))\{k\} \cong \bigoplus_{k\in\mathbb{Z}} HF^\infty(S^3)\{k\}.$$

The Alexander grading change of $F_{W,F,\mathfrak{s}}$ is equal to zero, by the Alexander grading change formula from Equation (6.11). The map $F_{W,F,\mathfrak{s}}$ on $\mathcal{H}FL^\infty$ thus becomes simply the regular cobordism map $F^\infty_{W,\mathfrak{s}}$ in each Alexander grading. By [OS03b], the map $F^\infty_{W,\mathfrak{s}} : HF^\infty(Y_1, \mathfrak{s}_1) \to HF^\infty(Y_2, \mathfrak{s}_2)$ is an isomorphism, and induces Maslov grading change zero by the formula from [OS06]. We note that the link cobordism map for a concordance intertwine the ι_K maps on the two ends, using Theorem 6.1.3 and the fact that the genus of the surface in a concordance is zero, so the half twist in the divides can be pulled from one end to the other. The map $F_{-W,-F,\mathfrak{s}}$ intertwines the involutions and is an isomorphism on homology by the same reasoning. This implies that the map $\mathcal{C} \to \mathfrak{I}_K$ is well defined.

Finally, we claim that the map $\mathcal{C} \to \mathfrak{I}_K$ is a homomorphism. By Theorem 6.1.1, the local equivalence class $[\mathcal{C}FL^\infty(S^3, K_1\#K_2)]$ is equal to the product $[\mathcal{C}FL^\infty(S^3, K_1)] \times [\mathcal{C}FL^\infty(S^3, K_2)]$ in $\mathfrak{I}_K$. □

6.8 Computations and examples

In this section we apply Theorem 6.1.1 to compute involutive concordance invariants $\overline{V}_0$ and $\underline{V}_0$ from [HM15] of connected sums of various knots. We compute the invariants for $T_r \# T_r$ "by hand" to demonstrate the formula, and then describe the results of some computer computations.

We recall the definition of $V_0, \overline{V}_0$ and $\underline{V}_0$. If (C, ∂) is a $\mathbb{Q}$–graded chain complex, freely generated over $\mathbb{F}_2[U]$, with a homotopy involution ι, we consider the complex

$$I^-(C, \iota) = \mathrm{Cone}(C \xrightarrow{Q\cdot(1+\iota)} Q \cdot C),$$

a chain complex freely generated over the ring $\mathbb{F}_2[U, Q]/(Q^2)$. The grading on the cone is obtained by shifting the domain of the map $Q \cdot (1 + \iota)$ up by one. Similarly one can define complexes $I^\infty(C, \iota)$ and $I^+(C, \iota)$. There is a long exact sequence

$$\cdots \to H_*(I^-(C, \iota)) \xrightarrow{i} H_*(I^\infty(C, \iota)) \xrightarrow{\pi} H_*(I^+(C, \iota)) \to \cdots .$$

Assuming that $H_*(C \otimes \mathbb{F}_2[U, U^{-1}])$ is isomorphic to $\mathbb{F}_2[U, U^{-1}]$, we can define $d(C)$ as the maximal grading of a non-torsion element of $H_*(C)$. Similarly, following [HM15], we define

$$\underline{d}(C, \iota) := \max\{r \equiv d(C) + 1 (\mathrm{mod}\, 2) \mid H_r(I^-(C, \iota)) \to H_r(I^\infty(C, \iota)) \text{ is nontrivial}\} - 1,$$

and

$$\overline{d}(C, \iota) := \max\{r \equiv d(C) (\mathrm{mod}\, 2) \mid H_r(I^-(C, \iota)) \to H_r(I^\infty(C, \iota)) \text{ is nontrivial}\}.$$

Note that unlike [HM15] and [HMZ16], the convention in this paper is that $HF^-(S^3)$ has top degree generator in grading 0.

To define the concordance invariants $\overline{V}_0$ and $\underline{V}_0$, one first defines $A_0^-(K)$ as the subcomplex of $CFK^\infty(S^3, K) = \mathcal{C}FL^\infty(Y, K)_0$ generated by terms $[\mathbf{x}, i, j]$ with $i \leq 0$, $j \leq 0$ and $A(\mathbf{x}) + (i - j) = 0$. In the notation used elsewhere in this paper, one would write this as the module generated by monomials $\mathbf{x} \cdot U^i V^j$ with $i, j \geq 0$ and $A(\mathbf{x}) + j - i = 0$. The map ι_K on

$CFK^\infty(K)$ induces a homotopy involution ι_0 on $A_0^-(K)$. The complex $A_0^-(K)$ inherits a homological grading from the Maslov grading on $CFK^\infty(S^3, K)$. One then defines

$$V_0(K) = -\frac{1}{2}d(A_0^-(K)), \qquad \underline{V}_0(K) = -\frac{1}{2}\underline{d}(A_0^-(K), \iota_0),$$

$$\text{and} \qquad \overline{V}_0(K) = -\frac{1}{2}\overline{d}(A_0^-(K), \iota_0).$$

6.8.1 A simple example: $T_r \# T_r$

We consider the connected sum of the right handed trefoil with itself. The complexes $\mathcal{C}FL^\infty(T_r)$ and $CFK^\infty(T_r)$, as well as the maps Φ and Ψ are shown in Figure 6.18. In [HM15, Section 7] it is shown that the knot involution ι_K is homotopic to reflection across the dashed line. It's easily computed that $\overline{V}_0(T_r) = V_0(T_r) = \underline{V}_0(T_r) = 1$.

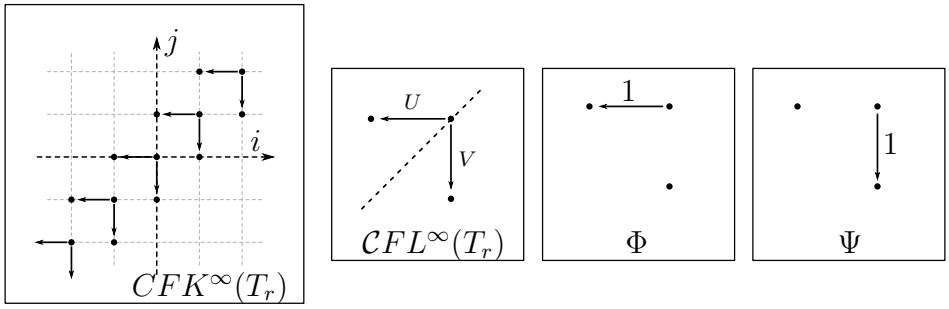

Figure 6.18: **The complexes for T_r, and the maps Φ and Ψ.** On the left is the standard picture for $CFK^\infty(T_r)$ for the right handed trefoil, drawn as a complex over $\mathbb{F}_2$. The right three boxes show the complex $\mathcal{C}FL^\infty(T_r)$, a module over $\mathbb{F}_2[U, V, U^{-1}, V^{-1}]$, as well as the maps Φ and Ψ. Reflection across the dashed line yields ι_K. Note that Φ and Ψ individually do not define maps on $CFK^\infty(T_r)$ since they shift Alexander grading.

In Figure 6.19 we see the conjugation action on $\mathcal{C}FL^\infty(T_r \# T_r)$. Note that ι_K on the tensor product complex appears somewhat asymmetric: the extra arrow on the right could instead point to the other generator nearby. This reflects the two formulas for $\iota_{T_r \# T_r}$ from Theorem 6.1.1. The arrow shown in Figure 6.18 corresponds to a summand of $\iota_1\Phi_1|\iota_2\Psi_2$. If we replaced that summand by $\iota_1\Psi_1|\iota_2\Phi_2$, we'd replace the arrow with one pointing to the

nearby generator. Note that the two ι_K–complexes are homotopy equivalent, by Lemma 6.2.14.

In computing the invariants $\underline{d}$ and $\overline{d}$, the following lemma is convenient. It is modified slightly from [HMZ16], because of our grading convention:

Lemma 6.8.1 ([HMZ16, Lemma 2.12])**.** *Suppose ι is a homotopy involution on a $\mathbb{Z}$–graded complex (C, ∂) which is freely generated over $\mathbb{F}_2[U]$.*

(a) The quantity $\underline{d}(C, \iota)$ is the maximum grading of a homogeneous $v \in C$ such that $\partial v = 0$, $[U^n v] \neq 0$ for all $n \geq 0$ and there exists a $w \in C$ such that $\partial w = (\mathrm{id} + \iota)v$.

(b) The quantity $\overline{d}(C, \iota)$ can be computed by considering triples (x, y, z) of homogeneous elements in C with at least one of x and y nonzero such that $\partial y = (\mathrm{id} + \iota)x, \partial z = U^m x$ for some $m \geq 0$ and $[U^n(U^m y + (\mathrm{id} + \iota)z)] \neq 0$ for all $n \geq 0$. The quantity $\overline{d}(C, \iota)$ is equal to the maximum of $\mathrm{gr}(x) + 1$ (ranging over triples with $x \neq 0$) and $\mathrm{gr}(y)$ (ranging over triples where $x = 0$).

We now compute the involutive invariants of $T_r \# T_r$. We note that $T_r \# T_r$ is alternating, and hence by [HM15, Proposition 8.1] the map ι_K is uniquely determined by its grading and filtration properties, and by the relation $\iota_K^2 \simeq \mathrm{id} + \Phi \circ \Psi$. Using this, the large surgery invariants can be computed for alternating knots [HM15, Theorem 1.7]. Nonetheless, we include this example as a demonstration of the formula in a simple case.

Proposition 6.8.2. *The invariants for $T_r \# T_r$ from large surgeries are*

$$\overline{V}_0(T_r \# T_r) = V_0(T_r \# T_r) = 1 \qquad \text{and} \qquad \underline{V}_0(T_r \# T_r) = 2.$$

Proof. The complex $\mathcal{C}FL^\infty(T_r \# T_r)$ and the involution $\iota_{T_r \# T_r}$ are shown in Figure 6.19. The complex $A_0^-(T_r \# T_r)$ is shown in Figure 6.20.

Note that in the complex A_0^-, clearly $[b]$ is an element such that $U^n[b] \neq 0$, and $[b]$ has maximal grading in the complex. Hence

$$V_0 = -\frac{1}{2} d(A_0^-) = -\frac{1}{2} \mathrm{gr}([b]) = 1.$$

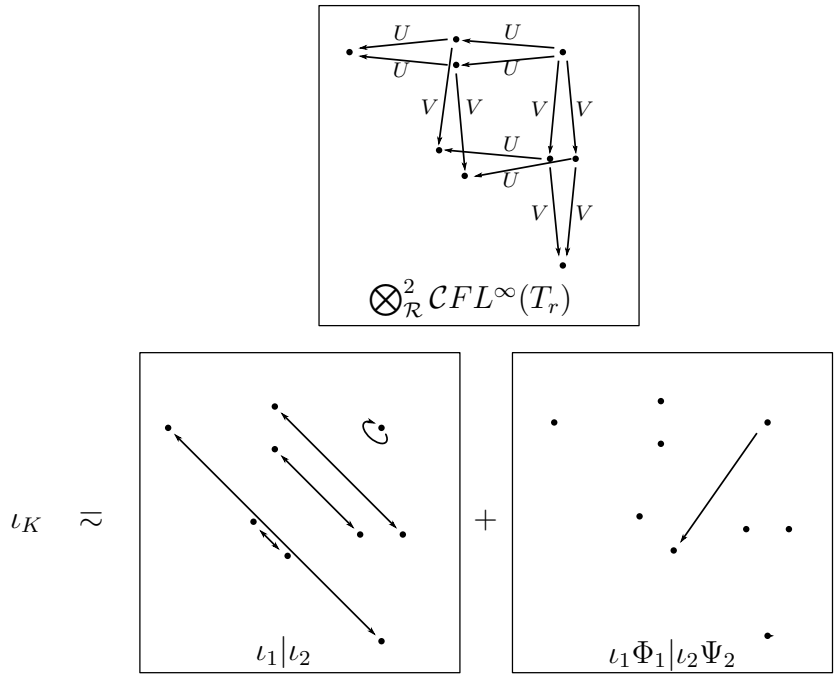

Figure 6.19: **The complex $\mathcal{C}FL^\infty(T_r\#T_r)$, as well the conjugation map $\iota_K \simeq \iota_1|\iota_2 + \iota_1\Phi_1|\iota_2\Psi_2$.**

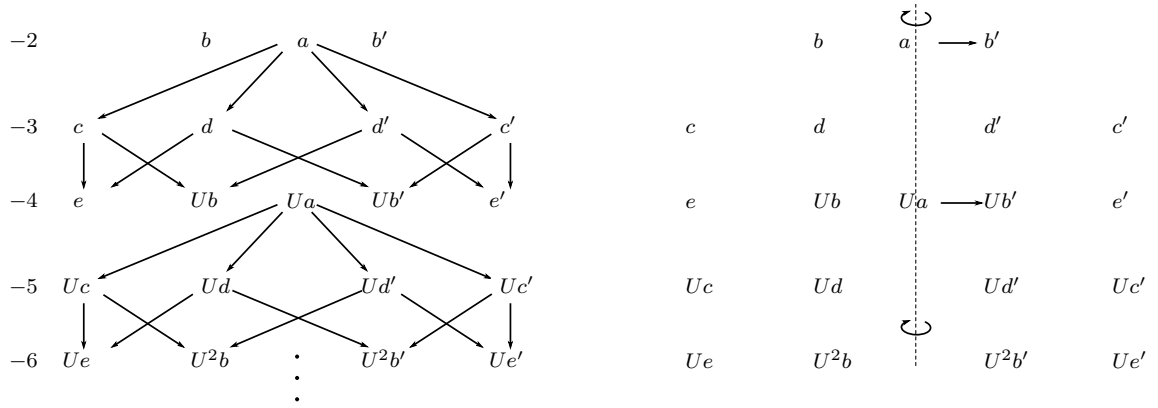

Figure 6.20: **The complex $A_0^-(T_r\#T_r)$ on the left and the involution ι_0, on the right.** The involution on the right is equal to reflection across the dashed line plus the arrows.

To compute $\overline{V}_0 = -\frac{1}{2}\overline{d}_0(AI_0^-)$, we note that $(0, b, 0)$ is a triple as in part (b) of Lemma 6.8.1, so $\overline{d}(AI_0^-) \geq -2$. For this not to be an equality, we need there to be a triple (x, y, z) satisfying condition (b) and with $\mathrm{gr}(x) \geq -2$ and $x \neq 0$. This implies that $\mathrm{gr}(x) = -2$, as that is the top grading of elements in the chain complex. Similarly, grading considerations

imply that such a triple must have $y = 0$ and $x = \iota(x)$. There is only one nonzero element of grading -2 with $x = \iota(x)$, namely $x = b+b'$ (note that a has grading -2 but is not preserved by ι). The elements z such that $\partial z = U^m(b + b')$ are all of the form $z = U^{m-1}(c + d)$ or $z = U^{m-1}(d' + c')$. But

$$[U^n U^{m-1}(\mathrm{id} + \iota)(c + d)] = [U^{n+m-1}(c + d + d' + c')] = [0]$$

so there are no triples (x, y, z) satisfying the criteria of part (b) of Lemma 6.8.1 with $x = b + b'$ and $z = U^{m-1}(c + d)$. The same argument rules out triples with $x = b + b'$ and $z = U^{m-1}(c' + d')$, so we conclude that

$$\overline{V}_0 = -\frac{1}{2}\,\mathrm{gr}([b]) = 1.$$

Finally we consider $\underline{V}_0 = -\frac{1}{2}\underline{d}(AI_0^-)$ and part (a) of Lemma 6.8.1. Note that $[b]$ and $[b']$ are the only elements x in grading -2 such that $[U^n x] \neq 0$. Since neither is fixed by the involution and there are no nonzero boundaries of degree -2, we see that neither can satisfy part (a) of Lemma 6.8.1. Hence $\underline{d} < -2$, and since it is an even integer, we must have $\underline{d} \leq -4$. On the other hand, $x = Ub$ satisfies $U^n[Ub] \neq 0$ for all $n \geq 0$ and $\partial(d' + c') = Ub + Ub' = (\mathrm{id} + \iota)(Ub)$, so $\underline{d} = -4$ and

$$\underline{V}_0 = 2.$$

□

6.8.2 Further computer based computations

The procedure for computing $V_0, \overline{V}_0$ and $\underline{V}_0$ is amendable to computer computations when $CFK^\infty(K)$ and the involution ι_K are known. For knots with relatively simple $CFK^\infty(K)$ complex, such as thin knots, L-space knots and mirrors of L-space knots, it is shown in [HM15] that the involution ι_K is determined up to chain homotopy by the fact that $\iota_K^2 \simeq \mathrm{id} + \Phi \circ \Psi$. Using the formula from Theorem 6.1.1, we can compute the involution for any connected sum of such knots. In this section we provide some example computations using the assistance of *Macaulay2* [GS].

Using our computations, we can prove the following:

Proposition 6.1.7. *None of the knots in Figure 6.21 is concordant to a thin knot, an L-space knot, or the mirror of an L-space knot.*

Proof. According to [HM15, Proposition 8.2], if K is an thin knot, an L-space knot, or the mirror of an L-space knot, then the triple $(\overline{V}_0, V_0, \underline{V}_0)$ satisfies one of the two following conditions:

1. All three of $\overline{V}_0, V_0$, and $\underline{V}_0$ are nonnegative and $0 \leq \underline{V}_0 - \overline{V}_0 \leq 1$;
2. $\overline{V}_0 \leq 0$ while $V_0 = \underline{V}_0 = 0$.

The complexes for $T_{2,5}, T_{3,4}, T_{4,5}, T_{5,6}$ and $T_{6,7}$ are staircase complexes and are shown in Figure 6.22. As shown in [HM15, Section 7], the involution is uniquely determined on these complexes, and is reflection across the diagonal in the (i, j) plane. Using *Macaulay2*, we can compute the involutive complexes associated to the connected sums. The relevant correction terms from large surgeries are shown in Figure 6.21. Note that for the knots shown in that figure, none of them have the pattern listed above for the $\overline{V}_0, V_0$ and $\underline{V}_0$ invariants of thin knots, L–space knots or mirrors of L–space knots. □

Knot	$\overline{V}_0$	V_0	$\underline{V}_0$
$T_{4,5}\#T_{4,5}$	4	4	6
$T_{4,5}\#T_{4,5}\#T_{5,6}$	7	7	9
$T_{6,7}\#T_{6,7}$	9	9	12
$T_{4,5}\#T_{6,7}$	7	7	9
$T_{3,4}^{-1}\#T_{4,5}^{-1}\#T_{5,6}$	-1	1	1

Figure 6.21: **Some knots and their involutive invariants of large surgeries.**

Some caution is warranted making conclusions about connected sums of torus knots from these computations. For example $T_{5,6}\#T_{5,6}$ has $(\overline{V}_0, V_0, \underline{V}_0)$ equal to $(6, 6, 6)$.

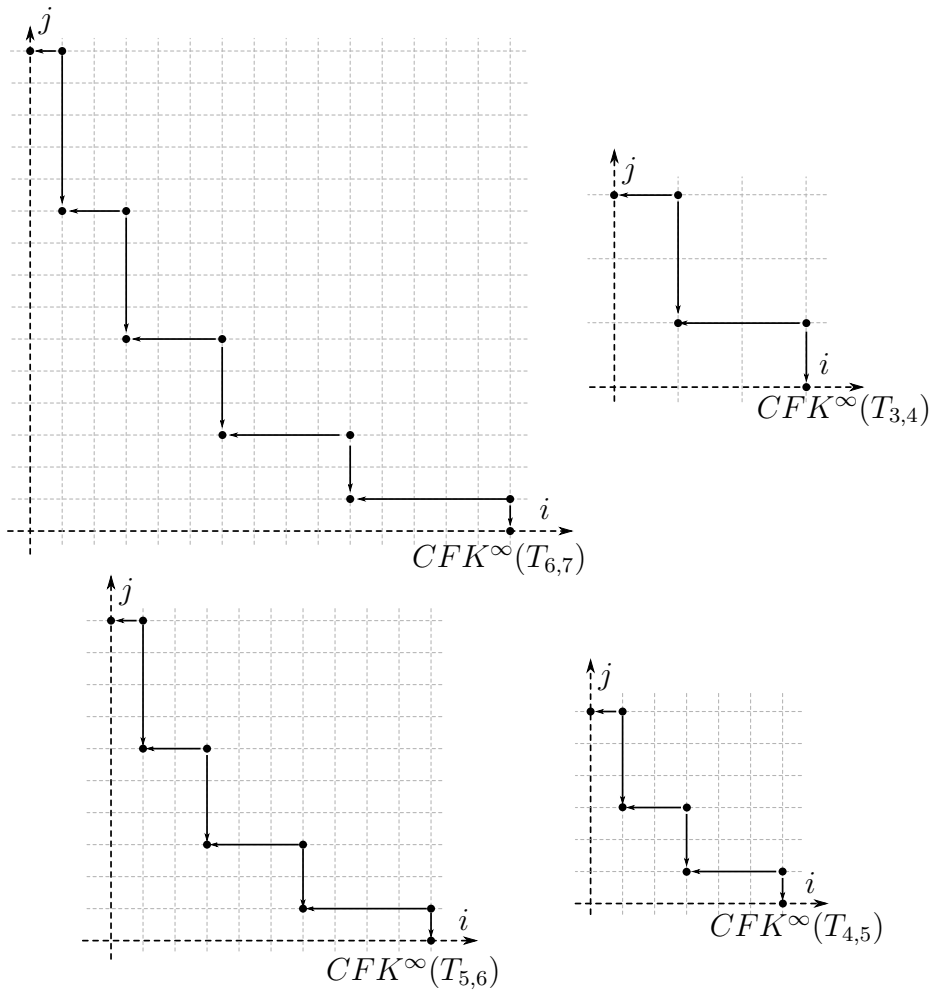

Figure 6.22: **Model complexes for $T_{3,4}, T_{4,5}, T_{5,6}$ and $T_{6,7}$.** The CFK^∞ complexes are formed by tensoring with $\mathbb{F}_2[U, U^{-1}]$.

Question 6.8.3. Are there knots in S^3 where all three $\overline{V}_0, V_0$ and $\underline{V}_0$ are all different?

The above question seems likely to have an affirmative answer, in light of the examples from [HMZ16].

6.9 Alternate proofs that G_i and E_i are homotopy inverses

In this section, we sketch two alternate proofs that G_i and E_i are homotopy inverses by direct computation. Both of these proofs highlight different features of the link Floer TQFT, and

involve directly counting holomorphic disks. The proof we gave earlier was based on formal properties of the link Floer TQFT.

The strategy described in Section 6.9.1 is based on the observation that holomorphic disk and triangle counts can be simplified in the presence of a 2–sphere embedded in the 3–manifold (e.g. in the connected sum of two 3–manifolds), as long as there are enough basepoints on the diagram.

In Section 6.9.2 we provide an alternate proof of Lemma 6.5.7, a key computation in our proof of the 4–dimensional surgery relations, by explicitly examining the holomorphic disks and triangles which contribute to the two maps.

6.9.1 Computing $E_i \circ G_i$ and $G_i \circ E_i$ by analyzing connected sum regions

We sketch a strategy for proving that E_1 and G_1 are homotopy inverses by using the 2–sphere separating the two summands of a connected sum to simplify the counts of holomorphic curves appearing in the link cobordism maps.

Proposition 6.5.1. *The connected sum maps E_1, E_2, G_1 and G_2 satisfy*

$$E_i \circ G_i \simeq \mathrm{id}, \qquad \textit{and} \qquad G_i \circ E_i \simeq \mathrm{id}\,.$$

Proof. Let us consider only E_1 and G_1. The strategy is to cut the link cobordisms for $E_1 \circ G_1$ and $G_1 \circ E_1$ into pieces, and compute some of the maps involved in the composition, explicitly, on certain diagrams. We can write $G_1 \circ E_1$ as the composition

$$G_1 \circ E_1 \simeq S^-_{w_1,z_1} \Psi_{z_1} F^{\mathbf{z}}_{B'} F_1 F_3 F^{\mathbf{z}}_{B} \Psi_{z_1} S^+_{w_1,z_1} \tag{6.22}$$

where $S^{\pm}_{w_1,z_1}$ are the quasi-stabilization maps, F_1 and F_3 denote the 3–handle and 1–handle maps, respectively, and B and B' are bands which disconnect and reconnect $K_1 \# K_2$, respectively. This is shown in Figure 6.23. Similarly, the composition $E_1 \circ G_1$ can be written as

$$E_1 \circ G_1 \simeq F_3 F^{\mathbf{z}}_{B} \Psi_{z_1} S^+_{w_1,z_1} S^-_{w_1,z_1} \Psi_{z_1} F^{\mathbf{z}}_{B'} F_1. \tag{6.23}$$

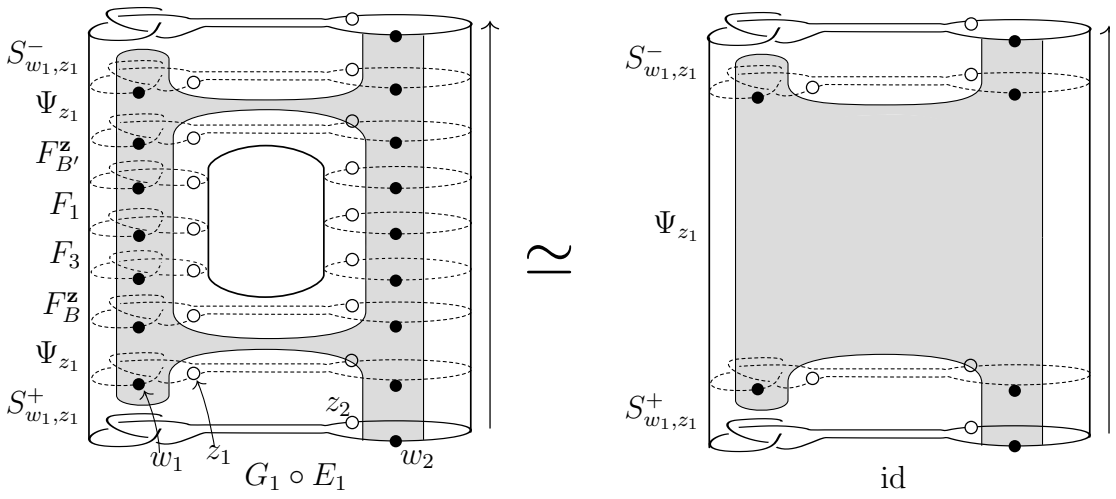

Figure 6.23: **A decomposition of $G_1 \circ E_1$ (left) and a decomposition of the identity link cobordism (right).** Here we have drawn only the surfaces with divides, and none of the 4–dimensional handles attached away from the link. The 4–manifold on the left consists of a 3–handle followed by a 1–handle. Equation (6.26), which allows us to relate the two maps in the figure, is an example of the surgery relation from Proposition 6.5.4.

Note that the formula for $E_1 \circ G_1$ in Equation (6.23) can be simplified since $S^+_{w_1,z_1} S^-_{w_1,z_1} \simeq \Phi_{w_1}$, which implies

$$\begin{aligned} \Psi_{z_1} S^+_{w_1,z_1} S^-_{w_1,z_1} \Psi_{z_1} &\simeq \Psi_{z_1} \Phi_{w_1} \Psi_{z_1} \\ &\simeq \Psi_{z_1} + \Phi_{w_1} \Psi^2_{z_1} \\ &\simeq \Psi_{z_1}. \end{aligned}$$

Hence Equation (6.23) simplifies to

$$E_1 \circ G_1 \simeq F_3 F^{\mathbf{z}}_B \Psi_{z_1} F^{\mathbf{z}}_{B'} F_1. \tag{6.24}$$

A key observation to this approach is that most of the maps appearing in Equations (6.22) and (6.24) can be computed using Heegaard diagrams or triples where there is a pair or triple of curves in the connected sum region. The curves appearing in the connected sum region can be taken to be isotopic to each other, though possibly not isotopic relative to the basepoints. All of the maps appearing in Equation (6.24) can be computed using such

diagrams or triples, and we will directly compute that

$$F_3 F^{\mathbf{z}}_B \Psi_{z_1} F^{\mathbf{z}}_{B'} F_1 \simeq \text{ id} \tag{6.25}$$

In the description of $G_1 \circ E_1$ in Equation (6.22), the only maps which can't be computed in diagrams like this are the quasi-stabilization maps $S^{\pm}_{w_1,z_1}$. However, a miracle occurs, which allows us to avoid computing transition maps between diagrams where $S^{\pm}_{w,z}$ can be computed, and diagrams where the 1–handle maps can be computed. Namely, we will see that

$$\Psi_{z_1} F^{\mathbf{z}}_{B'} F_1 F_3 F^{\mathbf{z}}_B \Psi_{z_1} \simeq \Psi_{z_1}. \tag{6.26}$$

Combining this with Equation (6.22), it will follow that

$$G_1 \circ E_1 \simeq S^-_{w_1,z_1} \Psi_{z_1} S^+_{w_1,z_1},$$

which is chain homotopic to the identity, by Lemma 3.9.6. Note that the "miracle" which occurs in Equation (6.26) is actually a special case of our formula for 4–dimensional surgery on an embedded curve from Proposition 6.5.4, as shown in Figure 6.23.

Thus, it remains to describe the diagrams we will use and prove Equations (6.25) and (6.26). Let $\mathcal{H}_{K_1\#K_2}$ and $\mathcal{H}_{K_1\sqcup K_2}$ denote the two diagrams shown in Figure 6.24, for $(Y_1\#Y_2, K_1\#K_2)$ and $(Y_1\#Y_2, K_1 \sqcup K_2)$, respectively.

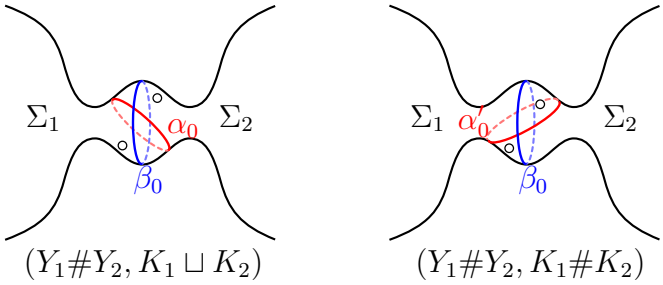

Figure 6.24: **The connected sum region of the diagrams $\mathcal{H}_{K_1\#K_2}$ and $\mathcal{H}_{K_1\sqcup K_2}$ for $(Y_1\#Y_2, K_1\#K_2)$ and $(Y_1\#Y_2, K_1 \sqcup K_2)$.** Both diagrams have two $\mathbf{w}$ basepoints which are not shown. Outside of this region, the diagrams coincide.

We now list the main subclaims which we will use to verify Equations (6.25) and (6.26). Note that in all of the following, we are setting $U_{w_1} = U_{w_2} = U$ and $V_{z_1} = V_{z_2} = V$:

Claim (1): Writing $\theta^-_{\mathbf{w}}$ and $\theta^+_{\mathbf{w}}$ for the intersection points of both $\alpha_0 \cap \beta_0$ and $\alpha'_0 \cap \beta_0$ (with the $+$ and $-$ generator determined by the $\mathrm{gr}_{\mathbf{w}}$ grading), there is an identification of both $\mathcal{C}FL^\infty(\mathcal{H}_{K_1\#K_2})$ and $\mathcal{C}FL^\infty(\mathcal{H}_{K_1\sqcup K_2})$ as

$$\left(\mathcal{C}FL^\infty(Y_1, K_1, w_1, z_1) \otimes_{\mathcal{R}} \mathcal{C}FL^\infty(Y_2, K_2, w_2, z_2)\right) \otimes_{\mathbb{F}} \langle \theta^+_{\mathbf{w}}, \theta^-_{\mathbf{w}} \rangle$$

and under this identification, for appropriately chosen almost complex structure, the differentials on both $\mathcal{C}FL^\infty(\mathcal{H}_{K_1\#K_2})$ and $\mathcal{C}FL^\infty(\mathcal{H}_{K_1\sqcup K_2})$ are identified with

$$\begin{pmatrix} \partial_1 |\, \mathrm{id} + \mathrm{id} \,| \partial_2 & 0 \\ 0 & \partial_1 |\, \mathrm{id} + \mathrm{id} \,| \partial_2 \end{pmatrix}.$$

Note that this implies that two maps on the complexes are chain homotopic iff each of the four entries of their matrices are individually chain homotopic over $\mathcal{C}FL^\infty(Y_1, K_1, w_1, z_1) \otimes_{\mathcal{R}} \mathcal{C}FL^\infty(Y_2, K_2, w_2, z_2)$.

Claim (2): On $\mathcal{H}_{K_1\#K_2}$, for appropriately chosen almost complex structure, the map Ψ_{z_1} is chain homotopic to

$$\Psi_{z_1} \simeq \begin{pmatrix} \Psi_{z_1} |\, \mathrm{id} & 0 \\ \mathrm{id} & \Psi_{z_1} |\, \mathrm{id} \end{pmatrix}.$$

Claim (3): For appropriately chosen almost complex structure, the band map $F^{\mathbf{z}}_B$ for splitting $K_1\#K_2$ into $K_1 \sqcup K_2$ using a type $\mathbf{z}$ band is chain homotopic to

$$F^{\mathbf{z}}_B \simeq \begin{pmatrix} V \cdot \mathrm{id} & * \\ 0 & \mathrm{id} \end{pmatrix}.$$

Claim (4): For appropriately chosen almost complex structure, the band map $F^{\mathbf{z}}_{B'}$ for connecting $K_1 \sqcup K_2$ to form $K_1\#K_2$ using a type $\mathbf{z}$ band is chain homotopic to

$$F^{\mathbf{z}}_{B'} \simeq \begin{pmatrix} \mathrm{id} & * \\ 0 & V \cdot \mathrm{id} \end{pmatrix}.$$

Claim (5): Using this matrix notation, the 1–handle and 3–handle maps take the form

$$F_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad \text{and} \qquad F_3 = \begin{pmatrix} 0 & 1 \end{pmatrix}.$$

Claim (6): An almost complex structure can be chosen such that all of the above statements hold.

Claim (1) follows from a standard neck stretching argument (see [HMZ16, Lemma 5.7]).

Claim (2) follows from sending the connected sum points towards the α_0' curve when we do the neck stretching (see [HMZ16, Lemma 5.7]).

Claims (3) and (4) follow since the band maps $F_B^{\mathbf{z}}$ and $F_{B'}^{\mathbf{z}}$ can both be computed by counting holomorphic triangles in a Heegaard triple with three curves in the connected sum region, as shown in Figure 6.25. By applying Theorem 2.7.8, we see that the counts of the holomorphic triangles contributing to the diagonal components of the maps $F_{B'}^{\mathbf{z}}$ and $F_B^{\mathbf{z}}$ can be obtained from the counts of triangles on the diagram with the 1–handle removed, and the lower left entry of both maps is zero. The Heegaard triple with the connected sum region removed just represents a change of diagrams map, for changing the $\boldsymbol{\alpha}$ curves. The only difference from the argument in Chapter 2 is that we also have to pay attention to the multiplicity over the $\mathbf{z}$ basepoints in the connected sum region. When we do this, we get the extra power of V in one of the diagonal components of each matrix, but not the other.

Note also that Claims (3) and (4) are written in terms of matrices with id or $V \cdot \mathrm{id}$ along the diagonal. However using a single triangle map to compute the band maps forces us to change the $\boldsymbol{\alpha}$ curves outside of the connected sum region slightly. To get back to the original $\boldsymbol{\alpha}$ curves outside of the connected sum region we would have to compute each of the band maps as a composition of two triangle maps. One Heegaard triple would be as in Figure 6.25 for band surgery, and the second Heegaard triple would be a change of diagrams map to move the $\boldsymbol{\alpha}$ curves back to their original position outside of the connected sum region. The triangle counts from Theorem 2.7.8 can be used to compute the diagonal and lower left

components of matrices for both maps. By naturality, the diagonal entries of the composition are chain homotopic to id or $V \cdot \mathrm{id}$.

Claim (5) is the definition of the 1–handle and 3–handle maps, put into matrix notation.

Claim (6) follows since Claims (3),(4) and (5) hold whenever the two connected sum necks are sufficiently stretched. Claims (1) and (2) require (by the argument in [HMZ16, Lemma 5.7]) that the connected sum points are additionally degenerated towards the α_0' curve, which can be arranged throughout.

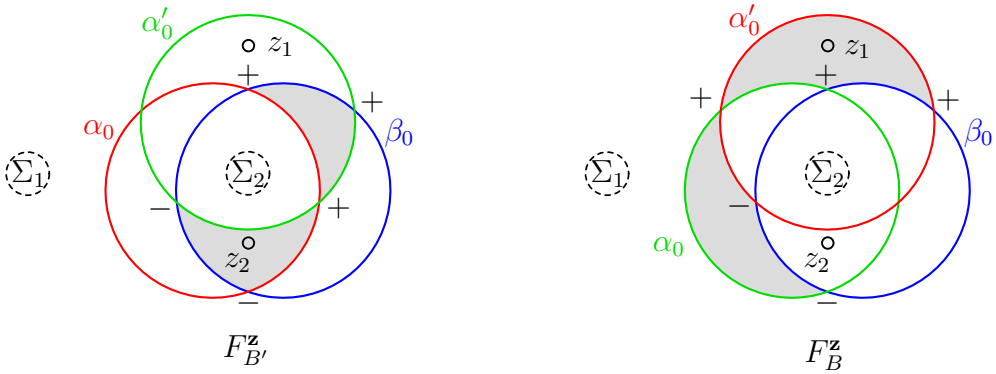

Figure 6.25: **The connected sum region of two triples which can compute the band maps $F^{\mathbf{z}}_{B}$ and $F^{\mathbf{z}}_{B'}$.** The shaded regions on each triple denote two classes of triangles which could be counted by the band maps.

Using Claims (1) – (6), Equations (6.25) and (6.26) follow from matrix multiplication, completing the proof. □

6.9.2 Direct proof of the 4–dimensional surgery relation in Proposition 6.5.4

In this section, we give a direct verification of Lemma 6.5.7, a key step in the proof of Proposition 6.5.4. Our previous argument followed from formal properties of the link Floer TQFT. We have included this proof in the appendix since it gives a different perspective of the maps, but involves some subtle counts of holomorphic curves.

Lemma 6.5.7. *For the decorated surfaces with divides M_1 and M_2 in Section 6.5.3, the two*

maps

$$F_{S^1\times D^3,M_1,\mathfrak{t}},\quad F_{D^2\times S^2,M_2,\mathfrak{t}_0}:\mathcal{C}FL^\infty(\varnothing,\varnothing)\to\mathcal{C}FL^\infty(S^1\times S^2,\mathbb{L}_0)$$

are filtered chain homotopy equivalent.

Proof. First, note that by grading considerations, as in the argument from Section 6.5.3, it is sufficient to restrict to $\widehat{CFL}$, the version obtained from $\mathcal{C}FL^-$ by setting all of the $U_{\mathbf{w}}$ and $V_{\mathbf{z}}$ variables equal to zero. On $\widehat{CFL}$, the maps Φ_w and Ψ_z are defined by counting disks going over w or z exactly once.

The argument follows from the following sublemma:

Sublemma 6.9.1. *In terms of the intersection points shown in Figure 6.11, the following hold:*

1. *One has*

$$F_{S^1\times D^3,M_3,\mathfrak{s}}(1)=\theta_1^+\theta_2^-+(C_2+C_3)\cdot\theta_1^-\theta_2^+$$

   for two constants $C_2,C_3\in\mathbb{F}_2$;

2. *One has*

$$F_{D^2\times S^2,M_2,\mathfrak{s}_0}(1)=\theta_1^+\theta_2^-+(C_2'+C_3')\cdot\theta_1^-\theta_2^+$$

   for two constants $C_2',C_3'\in\mathbb{F}_2$;

3. *For any choice of almost complex structure, the constants C_2,C_3,C_2', and C_3' satisfy*

$$C_2+C_3=C_2'+C_3'.$$

Once we complete the proof of the above sublemma, it will follow that $F_{S^1\times D^3,M_3,\mathfrak{s}}=F_{D^2\times S^2,M_2,\mathfrak{s}_0}$. Below, we prove each part of the sublemma separately. □

Proof of Part (1) *of Sublemma 6.9.1.* We now compute the map $F_{S^1\times D^3,M_1,\mathfrak{s}}$. The four dimensional cobordism is simply obtained by adding a 0–handle and a 1–handle. A decomposition of the decorated surface M_1 is shown in Figure 6.26. The decorated link cobordism $(S^1\times D^3,M_1)$ is constructed by adding a 4–dimensional 0–handle with a 2–dimensional disk

inside of it (the 2–dimensional disk is split into two pieces by the dividing set), then performing a $T^+_{w_1,z_1}$ quasi-stabilization, attaching a 4–dimensional 1–handle, then attaching a type $\mathbf{z}$ band, which goes around the 1–handle once, then attaching a cylindrical link cobordism with dividing set corresponding to the composition $\Phi_{w_1}\Psi_{z_1}$ for w_1 and z_1 on one of the link components of $\mathbb{L}_0$.

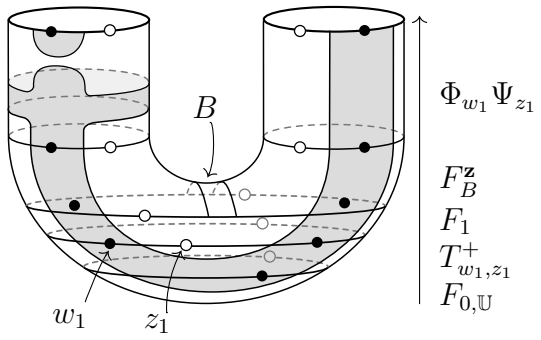

Figure 6.26: **A decomposition of the link cobordism** $(S^1 \times D^3, M_1)$**.** Here we have drawn only the surface (and not the 4–manifold $S^1 \times D^3$, obtained by adding a 0–handle and a 1–handle). Also F_1 and $F_{0,\mathbb{U}}$ denote the 1–handle and 0–handle maps, respectively.

The map $F_{0,\mathbb{U}}$ is the tautological map, and the composition $F_1 \circ T^+_{w,z}$ of the 1–handle map and the quasi-stabilization map send $F_{0,\mathbb{U}}(1)$ to $\Theta^{\mathbf{z}}_{\alpha\beta}$, the top $\mathrm{gr}_{\mathbf{z}}$ graded intersection point. The intersection point $\Theta^{\mathbf{z}}_{\alpha\beta}$ can be seen in the subdiagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ in the Heegaard triple shown at the top of Figure 6.27. The map $F^{\mathbf{z}}_B$ can be computed by using the Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ from Figure 6.27, using the formula

$$F^{\mathbf{z}}_B(\mathbf{x}) = F_{\alpha'\alpha\beta}(\Theta^{\mathbf{w}}_{\alpha'\alpha}, \mathbf{x}).$$

It is straightforward to verify that, ranging over all intersection points $\mathbf{y}$, there is only one homology class in any $\pi_2(\Theta^{\mathbf{w}}_{\alpha'\alpha}, \Theta^{\mathbf{z}}_{\alpha\beta}, \mathbf{y})$ which is nonnegative and has Maslov index zero. This class is shown in Figure 6.27 and has $\mathbf{y} = \theta^-_1\theta^+_2$. It has a unique holomorphic representative by the Riemann mapping theorem. Hence $F^{\mathbf{z}}_B(\Theta^{\mathbf{z}}_{\alpha\beta}) = \theta^-_1\theta^+_2$.

It is easy to see that there is exactly one nonnegative Maslov index one homology class ϕ_0 which contributes to $\Psi_{z_1}(\theta^-_1\theta^+_2)$ on $\widehat{CFL}$. This is shown in Figure 6.27. Since $\#\widehat{\mathcal{M}}(\phi_0) = 1$,

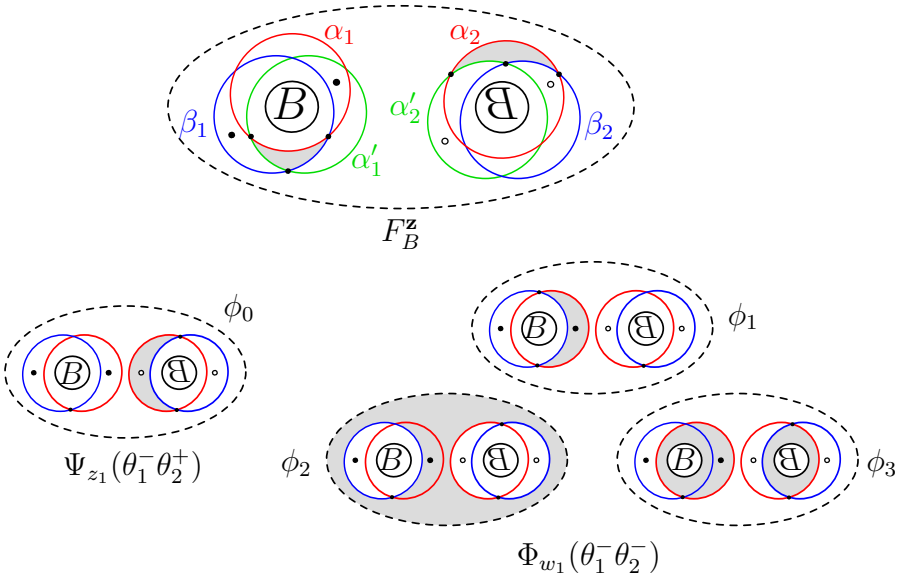

Figure 6.27: **Computing some terms of the maps $F_B^{\mathbf{z}}$ and Ψ_{z_1} and Φ_{w_1}.** On the top is a Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta})$ which can be used to compute $F_B^{\mathbf{z}}$. On the bottom left is the only Maslov index zero homology class of disks which contribute to $\Psi_{z_1}(\theta_1^-\theta_2^+)$. On the bottom right are the three homology classes which contribute to $\Phi_{w_1}\Psi_{z_1}(\theta_1^-\theta_2^+)$

have

$$\Psi_{z_1}(\theta_1^-\theta_2^+) = \theta_1^-\theta_2^-.$$

On the other hand, there are three nonnegative Maslov index one homology classes which potentially contribute to $\Phi_{w_1}(\theta_1^-\theta_2^-)$. Clearly $\widehat{\mathcal{M}}(\phi_1) = 1$. Write $C_2 = \#\widehat{\mathcal{M}}(\phi_2)$ and $C_3 = \#\widehat{\mathcal{M}}(\phi_3)$. Then

$$\Phi_{w_1}(\theta_1^-\theta_2^-) = \theta_1^+\theta_2^- + (C_2 + C_3)\cdot\theta_1^-\theta_2^+,$$

and hence,

$$F_{S^1\times D^3, M_1, \mathfrak{s}}(1) = \theta_1^+\theta_2^- + (C_2 + C_3)\cdot\theta_1^-\theta_2^+. \tag{6.27}$$

□

Proof of Part (2) *of Sublemma 6.9.1* . The map $F_{D^2\times S^2, M_2, \mathfrak{s}_0}$ can be computed by starting with a 0–handle map, which sends 1 to the generator $\widehat{CFL}(S^3, \mathbb{U})$ (where $\mathbb{U}$ is a doubly pointed unknot). Then one adds an additional unknotted component $\mathbb{U}'$ using the "disk

stabilization maps" from Chapter 2. This maps the generator to $\Theta^+_{\alpha\beta}$, the top graded element of $\widehat{CFL}(S^3, \mathbb{U} \cup \mathbb{U}')$ (the gradings $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ for this link). Finally, this is composed with a 2–handle map for a surgery on a 0–framed unknot which has linking number $+1$ with $\mathbb{U}$ and -1 with $\mathbb{U}'$. The 2–handle map is computed by counting holomorphic triangles in the Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ shown in Figure 6.28. Shown also in that Figure are three homology classes of triangles. One can easily check that these are the only homology classes with Maslov index zero in any $\pi_2(\Theta^+_{\alpha'\alpha}, \Theta^+_{\alpha\beta}, \mathbf{y})$ which have only nonnegative multiplicities.

Note that the maps

$$\mathfrak{s}_{\mathbf{w}}, \mathfrak{s}_{\mathbf{z}} : \pi_2(\Theta^+_{\alpha'\alpha}, \Theta^+_{\alpha\beta}, \mathbf{y}) \to \mathrm{Spin}^c(D^2 \times S^2)$$

are equal, since Spin^c structures on $D^2 \times S^2$ are uniquely determined by their restriction to $S^1 \times S^2$ and $\mathfrak{s}_{\mathbf{w}}(\mathbf{y}) = \mathfrak{s}_{\mathbf{z}}(\mathbf{y})$ for an intersection point $\mathbf{y}$ on $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta})$ by Lemma 6.3.2 since $\mathbb{L}_0$ is null-homologous. Furthermore, all three classes in Figure 6.28 represent $\mathfrak{s}_0$, since they restrict to the torsion Spin^c structure on $S^1 \times S^2$.

The homology class ψ_1 is in $\pi_2(\Theta^+_{\alpha'\alpha}, \Theta^+_{\alpha\beta}, \theta_1^+\theta_2^-)$, while ψ_2 and ψ_2 are in $\pi_2(\Theta^+_{\alpha'\alpha}, \Theta^+_{\alpha\beta}, \theta_1^-\theta_2^+)$. The homology class ψ_1 consists of only small triangles, and has a unique holomorphic representative by the Riemann mapping theorem. The homology classes ψ_2 and ψ_3 each consist of annular domains. Write $C_2' = \#\mathcal{M}(\psi_2)$ and $C_3' = \#\mathcal{M}(\psi_3)$.

We conclude

$$F_{D^2\times S^2, M_2, \mathfrak{s}_0}(1) = \theta_1^+\theta_2^- + (C_2' + C_3') \cdot \theta_1^-\theta_2^+. \tag{6.28}$$

□

Proof of Part (3) *of Sublemma 6.9.1 .* We now show that

$$C_2 + C_3 = C_2' + C_3'.$$

One could hope to find precise values for C_2, C_3, C_2' and C_3' (which we defined to be the number of representatives of ϕ_2, ϕ_3, ψ_2 and ψ_3, respectively) by using conformal mapping techniques, however this is challenging to make precise. Instead, we will use a (somewhat

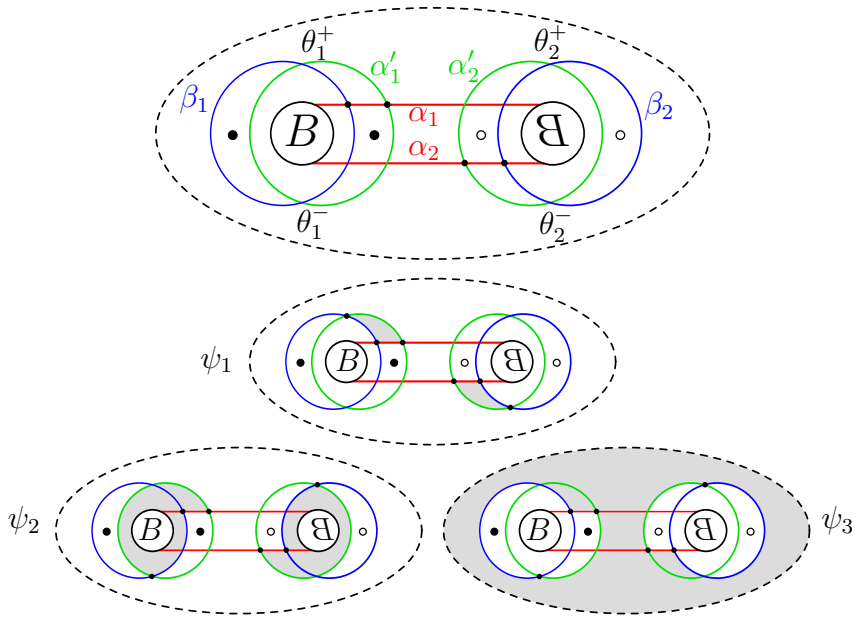

Figure 6.28: **Computing the 2–handle map.** On top is the Heegaard triple for computing the 2–handle map for the cobordism $D^2 \times S^2$. The intersection points $\Theta^+_{\alpha'\alpha}$ and $\Theta^+_{\alpha\beta}$ are labeled with dots, in the top diagram. On the bottom are the three nonnegative Maslov index zero homology classes in $\pi_2(\Theta^+_{\alpha'\alpha}, \Theta^+, \mathbf{x})$ for some $\mathbf{x}$.

indirect) Gromov compactness argument to show that $C_2 + C_3 = C_2' + C_3'$, for any choice of almost complex structure.

We will prove the following subclaims:

Claim (1): The map Φ_{w_1} vanishes on both $\widehat{CFL}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ and $\widehat{CFL}(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$.

Claim (2): $\widehat{\partial}_{\alpha'\beta}, \widehat{\partial}_{\alpha\beta}$ and $\widehat{\partial}_{\alpha'\alpha}$ all vanish on $\widehat{CFL}$.

Claim (3): The map $F_{\alpha'\alpha\beta,\mathfrak{s}_0}$ satisfies

$$F_{\alpha'\alpha\beta,\mathfrak{s}_0}(\Phi_{w_1}|1+1|\Phi_{w_1}) + \Phi_{w_1} F_{\alpha'\alpha\beta,\mathfrak{s}_0} + H_{w_1}(\partial_{\alpha'\alpha}|1+1|\partial_{\alpha\beta}) + \partial_{\alpha'\beta} H_{w_1} = 0$$

on $\widehat{CFL}$, for some map H_{w_1}.

Claim (4): As maps on $\widehat{CFL}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta})$, one has

$$\Phi_{w_1}(\theta_1^- \theta_2^+) = \theta_1^+ \theta_1^+, \qquad \text{and} \qquad \Phi_{w_1}(\theta_1^+ \theta_2^-) = (C_1'' + C_2'') \cdot \theta_1^+ \theta_2^+, \tag{6.29}$$

where $C_1'' = \#\widehat{\mathcal{M}}(\phi_2'')$ and $C_2'' = \#\widehat{\mathcal{M}}(\phi_3'')$ for the classes $\phi_2'', \phi_3'' \in \pi_2(\theta_1^+\theta_2^-, \theta_1^+\theta_2^+)$ shown in Figure 6.29.

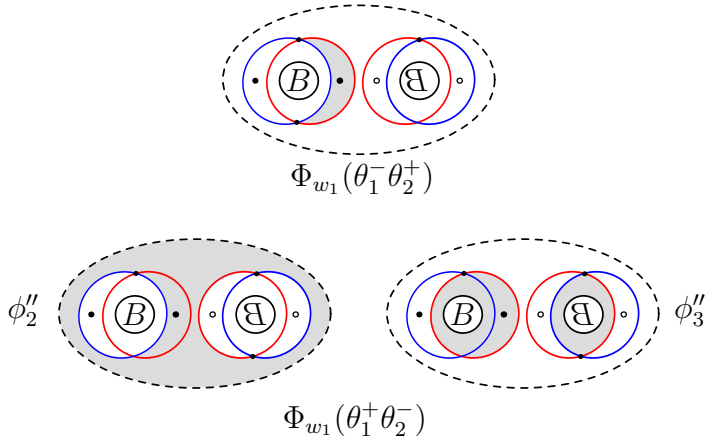

Figure 6.29: **The classes counted by $\Phi_{w_1}(\theta_1^-\theta_2^+)$ ad $\Phi_{w_1}(\theta_1^+\theta_2^-)$.**

We now verify Claims (1)–(4).

Claim (1) is obvious, since on $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ the region with w_1 also has z_1 in it (so no disks are counted by Φ_{w_1} on the hat flavor).

Claim (2) is also straightforward. On $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta})$ there are no nonnegative, Maslov index 1 domains which have zero multiplicity over all of the basepoints. Similarly, on each of the diagrams $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, there are exactly two nonnegative, Maslov index 1 classes which don't go over any basepoints. Each has domain equal to a rectangle, so each class has a unique representative and modulo two, and hence the total count between the two classes cancels, modulo two.

Claim (3) follows by noting that Gromov compactness yields that

$$F_{\alpha'\alpha\beta,\mathfrak{s}_0}(\partial_{\alpha'\alpha}|1+1|\partial_{\alpha\beta}) + \partial_{\alpha'\beta}F_{\alpha'\alpha\beta,\mathfrak{s}_0} = 0,$$

on $\mathcal{C}FL^-$, before coloring the complexes. One then differentiates the above expression with respect to U_{w_1}, and then sets all of the U_{w_i} and V_{z_i} variables to zero. Claim (3) follows.

Claim (4) follows from direct examination of the diagram in Figure 6.29.

Having proven Claims (1)–(4), we proceed with the proof that $C_2 + C_3 = C_2' + C_3'$. Combining Claims (1), (2) and (3), we conclude that

$$\Phi_{w_1}(F_{\alpha'\alpha\beta,\mathfrak{s}_0}(\Theta^+_{\alpha'\alpha}, \Theta^+_{\alpha\beta})) = 0.$$

In particular,

$$\Phi_{w_1}(F_{D^2\times S^2,M_2,\mathfrak{s}_0}(1)) = 0. \tag{6.30}$$

By Part (2) of Sublemma 6.9.1, we have

$$F_{D^2\times S^2,M_2,\mathfrak{s}_0}(1) = \theta_1^+\theta_2^- + (C_2' + C_3') \cdot \theta_1^-\theta_1^+.$$

Composing with Φ_{w_1} and using Equations (6.29) and (6.30), we see that

$$0 = \Phi_{w_1}(F_{D^2\times S^2,M_2,\mathfrak{s}_0}(1)) = (C_2'' + C_3'' + C_2' + C_3') \cdot \theta_1^+\theta_1^+ = 0.$$

Thus we conclude that

$$C_2' + C_3' = C_2'' + C_3''. \tag{6.31}$$

The final step is to show that $C_2 = C_2''$ and $C_3 = C_3''$, which combined (6.31) will finish the proof of Part (3) of Sublemma 6.9.1. To this end, we will show that $C_3 = C_3''$; the argument that that $C_2 = C_2''$ is nearly identical. This will follow from a holomorphic curve chasing argument (i.e. by using Gromov compactness). Recall that $C_3 = \#\widehat{\mathcal{M}}(\phi_3)$ and $C_3'' = \widehat{\mathcal{M}}(\phi_3'')$. Also note that ϕ_3 and ϕ_3'' have the same domain, but they represent different homology classes since they have different incoming and outgoing intersection points. We count the ends of two Maslov index 2 homology classes. This is shown in Figure 6.30. Note that the domains are the same, though the homology classes have different choices of incoming and outgoing intersection points. Here $\#\widehat{\mathcal{N}}^{\alpha}(A)$ denotes the count of holomorphic $\boldsymbol{\alpha}$–boundary degenerations in a class $A \in \pi_2^{\alpha}(\mathbf{x})$ (see [OS08, Section 5.3]). By [OS08, Theorem 5.5], the count of each of the two spaces of boundary degenerations is 1 (modulo 2). Each the two 1–dimensional moduli spaces in Figure 6.30 has boundary components corresponding to a single class of $\boldsymbol{\alpha}$–boundary degenerations, which hence each contribute equally to the boundaries of the two 1–dimensional moduli spaces.

Thus the left columns of Figure 6.30 have moduli spaces contributing 1 to the total numbers of ends. The ends in the middle column feature the same pair of holomorphic disks resulting from strip breaking, just in reversed order. Hence the contributions from the middle columns are equal. In the right column, the end on the top of the Figure contributes $C_3'' = \#\widehat{\mathcal{M}}(\phi_3'')$. On the bottom, the contribution is $C_3 = \#\widehat{\mathcal{M}}(\phi_3)$. Comparing all of these contributions, we conclude that $C_3 = C_3''$, concluding the proof. □

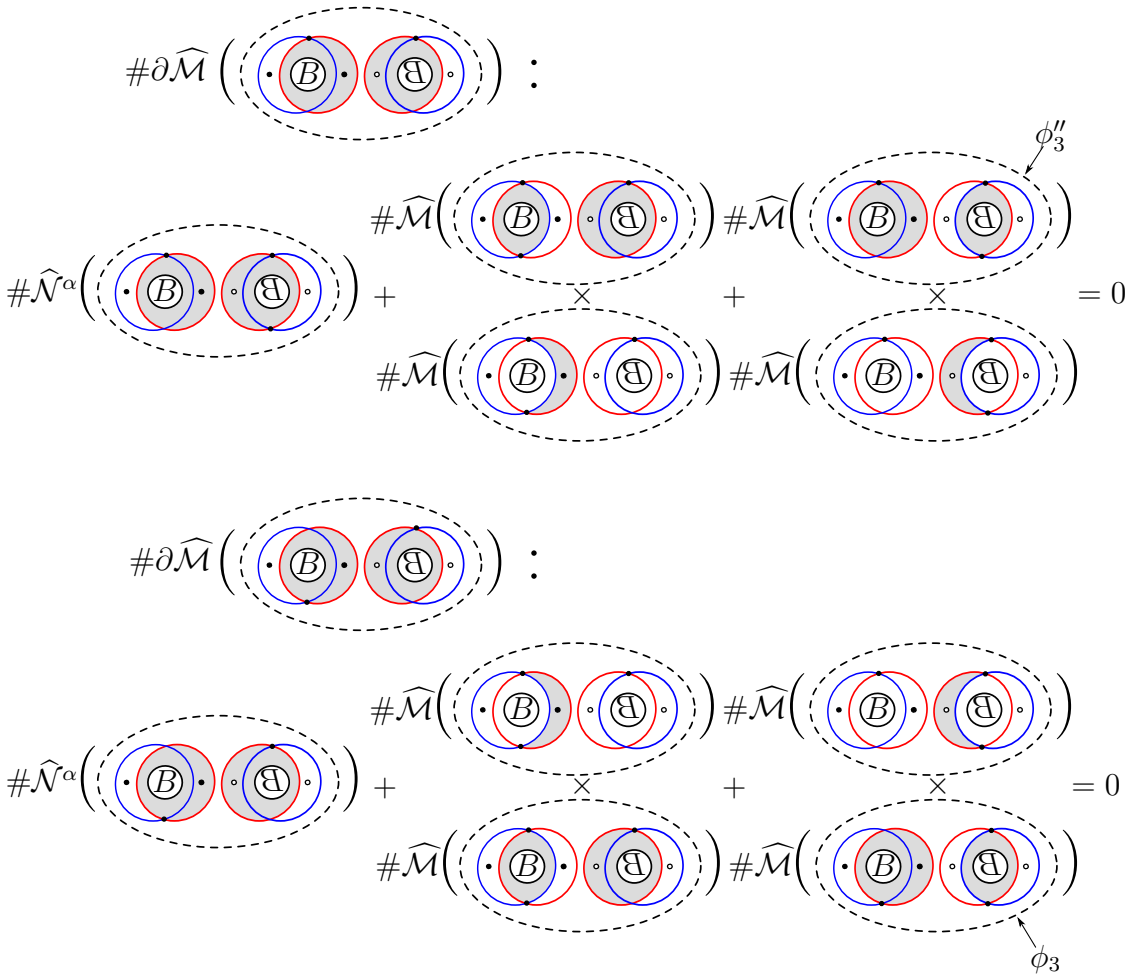

Figure 6.30: **Showing $C_3 = C_3''$ by counting the ends of two 1–dimensional moduli spaces.** There are ends corresponding to $\boldsymbol{\alpha}$ boundary degenerations, as well as strip breaking.

Bibliography

[Abr96] Lowell Abrams, *Two-dimensional topological quantum field theories and Frobenius algebras*, J. Knot Theory Ramifications **5** (1996), no. 5, 569–587.

[AE15] Akram S. Alishahi and Eaman Eftekhary, *A refinement of sutured Floer homology*, J. Symplectic Geom. **13** (2015), no. 3, 609–743.

[AE] Akram Alishahi and Eaman Eftekhary, *Tangle floer homology and cobordisms between tangles.*

[Ati88] Michael Atiyah, *Topological quantum field theories*, Inst. Hautes Études Sci. Publ. Math. **68** (1988), 175–186 (1989).

[BEH+03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888.

[BH16] Maciej Borodzik and Jennifer Hom, *Involutive Heegaard Floer homology and rational cuspidal curves*, 2016. e-print, `arXiv:1609.08303`.

[BT06] Christian. Blanchet and Vladimir. Turaev, *Axiomatic approach to TQFT*, Encyclopedia of Mathematical Physics, Elsevier (2006), 232234.

[Cer70] Jean Cerf, *La stratification naturelle des espaces de fonctions différentiables réelles et le théorème de la pseudo-isotopie*, Inst. Hautes Études Sci. Publ. Math. **39** (1970), 5–173.

[CGH10] Vincent Colin, Paolo Ghiggini, and Ko Honda, *Embedded contact homology and open book decompositions*, 2010. e-print, `arXiv:1008.2734`.

[CGH12a] ———, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions I*, 2012. e-print ,`arXiv:1208.1074`.

[CGH12b] ———, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions II*, 2012. e-print, `arXiv:1208.1077`.

[CGH12c] ———, *The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions III: from hat to plus*, 2012. e-print, `arXiv:1208.1526`.

[Don83] S. K. Donaldson, *An application of gauge theory to four-dimensional topology*, J. Differential Geom. **18** (1983), no. 2, 279–315.

[Elk95] Noam D. Elkies, *A characterization of the Z^n lattice*, Math. Res. Lett. **2** (1995), no. 3, 321–326.

[FK16] Peter Feller and David Krcatovich, *On cobordisms be1602.02637tween knots, braid index, and the Upsilon-invariant*, 2016. e-print, `arXiv:1602.02637`.

[FR79] Roger Fenn and Colin Rourke, *On Kirby's calculus of links*, Topology **18** (1979), no. 1, 1–15.

[FS95] Ronald Fintushel and Ronald J. Stern, *Immersed spheres in* 4*-manifolds and the immersed Thom conjecture*, Turkish J. Math. **19** (1995), no. 2, 145–157.

[GS99] Robert E. Gompf and András I. Stipsicz, 4*-manifolds and Kirby calculus*, Graduate Studies in Mathematics, vol. 20, American Mathematical Society, Providence, RI, 1999.

[GS] Daniel R. Grayson and Michael E. Stillman, *Macaulay2, a software system for research in algebraic geometry.*

[GWW12] David Gay, Katrin Wehrheim, and Christopher Woodward, *Connected cerf theory*, 2012. e-print, `https://math.berkeley.edu/~katrin/papers/cerf.pdf`.

[Hir94] Morris W. Hirsch, *Differential topology*, Graduate Texts in Mathematics, vol. 33, Springer-Verlag, New York, 1994. Corrected reprint of the 1976 original.

[HKM08] Ko Honda, William H. Kazez, and Gordana Matić, *Contact structures, sutured Floer homology and TQFT*, 2008. e-print, `arXiv:0807.2431`.

[HM15] Kristen Hendricks and Ciprian Manolescu, *Involutive Heegaard Floer homology*, 2015. e-print, `arXiv:1507.00383`.

[HMZ16] Kristen Hendricks, Ciprian Manolescu, and Ian Zemke, *A connected sum formula for involutive Heegaard Floer homology*, 2016. e-print, `arXiv:1607.07499`.

[Hom14] Jennifer Hom, *The knot Floer complex and the smooth concordance group*, Comment. Math. Helv. **89** (2014), no. 3, 537–570.

[Hon00] Ko Honda, *On the classification of tight contact structures. II*, J. Differential Geom. **55** (2000), no. 1, 83–143.

[HT09] Michael Hutchings and Clifford Henry Taubes, *The Weinstein conjecture for stable Hamiltonian structures*, Geom. Topol. **13** (2009), no. 2, 901–941.

[HW73] Allen Hatcher and John Wagoner, *Pseudo-isotopies of compact manifolds*, Société Mathématique de France, Paris, 1973. With English and French prefaces, Astérisque, No. 6.

[JM16a] Andás Juhász and Marco Marengon, *Computing cobordism maps in link Floer homology and the reduced Khovanov TQFT*, 2016. e-print, `arXiv:1503.00665`.

[JM16b] András Juhász and Marco Marengon, *Concordance maps in knot Floer homology*, Geom. Topol. **20** (2016), no. 6, 3623–3673.

[JT12] András Juhász and Dylan P. Thurston, *Naturality and mapping class groups in Heegaard Floer homology*, 2012. e-print, `arXiv:1210.4996`.

[Juh06] András Juhász, *Holomorphic discs and sutured manifolds*, Algebr. Geom. Topol. **6** (2006), 1429–1457.

[Juh15] András Juhász, *Defining and classifying TQFTs via surgery*, 2015. e-print, `arXiv:1408.0668v5`.

[Juh16] András Juhász, *Cobordisms of sutured manifolds and the functoriality of link Floer homology*, Adv. Math. **299** (2016), 940–1038.

[Kir78] Robion Kirby, *A calculus for framed links in S^3*, Invent. Math. **45** (1978), no. 1, 35–56.

[KLT10a] Çağatay Kutluhan, Yi-Jen Lee, and Clifford H. Taubes, *HF=HM I : Heegaard Floer homology and Seiberg–Witten Floer homology*, 2010. e-print, arXiv:1007.1979.

[KLT10b] ———, *HF=HM II: Reeb orbits and holomorphic curves for the ECH/Heegaard-Floer correspondence*, 2010. e-print, arXiv:1008.1595.

[KLT10c] ———, *HF=HM III: holomorphic curves and the differential for the ECH/Heegaard Floer correspondence*, 2010. e-print, arXiv:1010.3456.

[KLT11] ———, *HF=HM IV: Seiberg–Witten Floer homology and ECH correspondence*, 2011. e-print, arXiv:1107-2297.

[KLT12] ———, *HF=HM V: Seiberg–Witten Floer homology and handle additions*, 2012. e-print, arXiv:1204.0115.

[KM07] Peter Kronheimer and Tomasz Mrowka, *Monopoles and three-manifolds*, New Mathematical Monographs, vol. 10, Cambridge University Press, Cambridge, 2007.

[Kos93] Antoni A. Kosinski, *Differential manifolds*, Pure and Applied Mathematics, vol. 138, Academic Press, Inc., Boston, MA, 1993.

[KR08] Mikhail Khovanov and Lev Rozansky, *Matrix factorizations and link homology*, Fund. Math. **199** (2008), no. 1, 1–91.

[Lin14] Francesco Lin, *Morse-Bott singularities in monopole Floer homology and the triangulation conjecture*, 2014. e-print, `arXiv:1404.4561`.

[Lin16] ———, *Pin(2)-monopole Floer homology, higher compositions and connected sums*, 2016. e-print, `arXiv:1605.03137`.

[Lip06] Robert Lipshitz, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. **10** (2006), 955–1097.

[LP08] Aaron D. Lauda and Hendryk Pfeiffer, *Open-closed strings: two-dimensional extended TQFTs and Frobenius algebras*, Topology Appl. **155** (2008), no. 7, 623–666.

[Lut77] Robert Lutz, *Structures de contact sur les fibrés principaux en cercles de dimension trois*, Ann. Inst. Fourier (Grenoble) **27** (1977), no. 3, ix, 1–15.

[Man07] Ciprian Manolescu, *An unoriented skein exact triangle for knot Floer homology*, Math. Res. Lett. **14** (2007), no. 5, 839–852.

[Man16] ———, *Pin(2)-equivariant Seiberg-Witten Floer homology and the triangulation conjecture*, J. Amer. Math. Soc. **29** (2016), no. 1, 147–176.

[MO10] Ciprian Manolescu and Peter S. Ozsváth, *Heegaard Floer homology and integer surgeries on links*, 2010. e-print, `arXiv:1011.1317`.

[Ni14] Yi Ni, *Homological actions on sutured Floer homology*, Math. Res. Lett. **21** (2014), no. 5, 1177–1197.

[Nic03] Liviu I. Nicolaescu, *The Reidemeister torsion of 3-manifolds*, De Gruyter Studies in Mathematics, vol. 30, Walter de Gruyter & Co., Berlin, 2003.

[Oh96] Yong-Geun Oh, *Fredholm theory of holomorphic discs under the perturbation of boundary conditions*, Math. Z. **222** (1996), no. 3, 505–520.

[OS00a] Peter Ozsváth and Zoltán Szabó, *Higher type adjunction inequalities in Seiberg-Witten theory*, J. Differential Geom. **55** (2000), no. 3, 385–440.

[OS00b] ———, *The symplectic Thom conjecture*, Ann. of Math. (2) **151** (2000), no. 1, 93–124.

[OS03a] ———, *Knot Floer homology and the four-ball genus*, Geom. Topol. **7** (2003), 615–639.

[OS03b] Peter S. Ozsváth and Zoltán Szabó, *Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary*, Adv. Math. **173** (2003), no. 2, 179–261.

[OS04a] Peter Ozsváth and Zoltán Szabó, *Holomorphic triangle invariants and the topology of symplectic four-manifolds*, Duke Math. J. **121** (2004), no. 1, 1–34.

[OS04b] ———, *Holomorphic triangle invariants and the topology of symplectic four-manifolds*, Duke Math. J. **121** (2004), no. 1, 1–34.

[OS04c] Peter S. Ozsváth and Zoltán Szabó, *Holomorphic disks and knot invariants*, Adv. Math. **186** (2004), no. 1, 58–116.

[OS04d] ———, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math. (2) **159** (2004), no. 3, 1159–1245.

[OS04e] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027–1158.

[OS06] ———, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. Math. **202** (2006), no. 2, 326–400.

[OS08] ———, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, Algebr. Geom. Topol. **8** (2008), no. 2, 615–692.

[OSS14] Peter S. Ozsváth, András I. Stipsicz, and Zoltán Szabó, *Concordance homomorphisms from knot Floer homology*, 2014. e-print, `arXiv:1407.1795`.

[OSS15] Peter S. Ozsváth, András I. Stipsicz, and Zoltán Szabó, *Unoriented knot floer homology and the unoriented four-ball genus*, 2015. e-print, `arXiv:1508.03243`.

[Pal63] Richard S. Palais, *Morse theory on Hilbert manifolds*, Topology **2** (1963), 299–340.

[Ras03] Jacob A. Rasmussen, *Floer homology and knot complements*, Ph.D. Thesis, 2003. `arXiv:math.GT/0306378`.

[Ras04] Jacob Rasmussen, *Lens space surgeries and a conjecture of Goda and Teragaito*, Geom. Topol. **8** (2004), 1013–1031.

[Ras07] Jacob A. Rasmussen, *Lens space surgeries and L-space homology spheres*, 2007. e-print, `arXiv:0710.2531`.

[Rob97] Justin Roberts, *Kirby calculus in manifolds with boundary*, Turkish J. Math. **21** (1997), no. 1, 111–117.

[Sar11a] Sucharit Sarkar, *Grid diagrams and the Ozsváth-Szabó tau-invariant*, Math. Res. Lett. **18** (2011), no. 6, 1239–1257.

[Sar11b] ———, *Maslov index formulas for Whitney n-gons*, J. Symplectic Geom. **9** (2011), no. 2, 251–270.

[Sar15] ———, *Moving basepoints and the induced automorphisms of link Floer homology*, Algebr. Geom. Topol. **15** (2015), no. 5, 2479–2515.

[Saw95] Stephen Sawin, *Direct sum decompositions and indecomposable TQFTs*, J. Math. Phys. **36** (1995), no. 12, 6673–6680.

[Seg88] G. B. Segal, *The definition of conformal field theory*, Differential geometrical methods in theoretical physics (Como, 1987), 1988, pp. 165–171.

[Sto15] Matthew Stoffregen, *Pin(2)-equivariant Seiberg-Witten Floer homology of Seifert fibrations*, 2015. e-print, `arXiv:1505.03234`.

[Wit88] Edward Witten, *Topological quantum field theory*, Comm. Math. Phys. **117** (1988), no. 3, 353–386.

[Zem15] Ian Zemke, *A graph TQFT for hat Heegaard Floer homology*, 2015. e-print, `arXiv:1503.05846`.

[Zem16] ———, *Quasi-stabilization and basepoint moving maps in link Floer homology*, 2016. e-print, `arXiv:1604.04316`.

[Zem] ———, *Duality and mapping tori in Heegaard Floer homology.* (in preparation).