Monopole Floer Homology, Link Surgery, and Odd Khovanov Homology

Jonathan Michael Bloom

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2011

ABSTRACT

Monopole Floer Homology, Link Surgery, and Odd Khovanov Homology

Jonathan Michael Bloom

We construct a link surgery spectral sequence for all versions of monopole Floer homology with mod 2 coefficients, generalizing the exact triangle. The spectral sequence begins with the monopole Floer homology of a hypercube of surgeries on a 3-manifold Y, and converges to the monopole Floer homology of Y itself. This allows one to realize the latter group as the homology of a complex over a combinatorial set of generators. Our construction relates the topology of link surgeries to the combinatorics of graph associahedra, leading to new inductive realizations of the latter.

As an application, given a link L in the 3-sphere, we prove that the monopole Floer homology of the branched double-cover arises via a filtered perturbation of the differential on the reduced Khovanov complex of a diagram of L. The associated spectral sequence carries a filtration grading, as well as a mod 2 grading which interpolates between the delta grading on Khovanov homology and the mod 2 grading on Floer homology. Furthermore, the bigraded isomorphism class of the higher pages depends only on the Conway-mutation equivalence class of L. We constrain the existence of an integer bigrading by considering versions of the spectral sequence with non-trivial $U_\dagger$ action, and determine all monopole Floer groups of branched double-covers of links with thin Khovanov homology.

Motivated by this perspective, we show that odd Khovanov homology with integer coefficients is mutation invariant. The proof uses only elementary algebraic topology and leads to a new formula for link signature that is well-adapted to Khovanov homology.

Table of Contents

List of Figures

Acknowledgments

I am deeply grateful to my advisor, Peter Ozsváth, for his unwavering and enthusiastic support over the last five years, and for providing wonderful problems to explore. Even before I arrived at Columbia, he arranged for me to attend the 2006 Summer School at the Park City Mathematical Institute, a formative experience that kindled my interest in low-dimensional topology and geometry. I started graduate school with a focus on Heegaard Floer homology and Khovanov homology. At Professor Ozsváth's suggestion, and with his invaluable guidance, I begin to pursue connections with monopole Floer theory. This has taken me down an exciting and rewarding path.

I am indebted to Tim Perutz and Tom Mrowka for generous mentoring in monopole Floer theory. I have also benefitted from the expertise of Adam Knapp, Peter Kronheimer, and Max Lipyanskiy. I want to thank Peter Kronheimer and Tom Mrowka for investing the time and energy to write their foundational book *Monopoles and Three-Manifolds*. My copy has surely seen better days.

It has been a privilege to be a member of the mathematical community at Columbia. I am grateful for the friendship and support of John Baldwin, Ben Elias, Allison Gilmore, Josh Greene, Eli Grigsby, Adam Levine, Tom Peters, Ina Petkova, Thibaut Pugin, Vera Vertesi, and Rumen Zarev. My research has benefitted from discussions with Professors Dave Bayer, Mikhail Khovanov, Robert Lipshitz, and Dylan Thurston at Columbia, and Professors Satyan Devadoss, Danny Ruberman, James Stasheff, and Zoltán Szábo abroad, among many others.

I had the good fortune of spending the spring of 2010 at the Mathematical Sciences Research Institute for the program *Homology Theories of Knots and Links*. I want to thank my uncle Steven Greenfield for providing a warm home away from home during my stay in Berkeley. I have spent my final year of graduate school as an exchange scholar at MIT, and would like to thank the department for welcoming me. I am also grateful for support from the National Science Foundation (grant DMS-0739392).

I want to express my appreciation to several people who nurtured my interest in mathematics in my youth: my mentor Greg Bachelis, my calculus teacher Joe Brandell, and my cousin David Goss, who inspired me to attend the Ross Program. Through it all, my parents, sister, and grandparents have been a constant source of love and support. Finally, I want to thank Masha for her endless encouragement and so much more.

Chapter 1

Introduction

Monopole Floer homology is a gauge-theoretic invariant defined via Morse theory on the Chern-Simons-Dirac functional. As such, the underlying chain complex is generated by Seiberg-Witten monopoles over a 3-manifold, and the differential counts monopoles over the product of the 3-manifold with $\mathbb{R}$. We review this construction in Section 1.1.

In [27], a surgery exact triangle is associated to a triple of surgeries on a knot in a 3-manifold (for a precursor in instanton Floer homology, see [11], [18]). In Chapter 2, we construct a link surgery spectral sequence in monopole Floer homology, generalizing the exact triangle. This is a spectral sequence which starts at the monopole Floer homology of a hypercube of surgeries on Y along L, and converges to the monopole Floer homology of Y itself. The differentials count monopoles on 2-handle cobordisms equipped with families of metrics parameterized by polytopes called permutohedra. Those metrics parameterized by the boundary of the permutohedra are stretched to infinity along collections of hypersurfaces representing surgered 3-manifolds. The monopole counts satisfy identities obtained by viewing the map associated to each polytope as a null-homotopy for the map associated to its boundary. Note that this can be seen as analogue of Ozsváth and Szabó's link surgery spectral sequence for Heegaard Floer homology [36]. There, the differentials count pseudo-holomorphic polygons in Heegaard multi-diagrams, and they satisfy A_∞ relations which encode degenerations of conformal structures on polygons.

Our construction introduces a number of techniques that we hope will be of more general use. In Sections 2.1 and 2.3, we couple the topology of 2-handle cobordisms arising from link surgeries to the combinatorics of polytopes called graph associahedra [12]. For the chain-level Floer maps induced by 2-handle cobordisms, these polytopes encode a mixture of commutativity and associativity up to homotopy. We hope this coupling, and its relationship to finite product lattices, will be of independent interest to algebraists and combinatorialists. As one application, in Chapter 3 we obtain a simple, recursive construction of realizations of certain graph associahedra (Theorem 3.0.7). This specializes to give realizations of permutohedra as refinements of associahedra, which in turn refine hypercubes (see Figures 2.13 through 3.4). Curiously, these realizations are predicted by the "sliding-the-point" proof of the naturality of the $U_\dagger$ action in Floer theory.

Our construction of polytopes of metrics was inspired by the pentagon of metrics in the proof of the surgery exact triangle [27]. However, to make use of more general polytopes, we must effectively organize the mix of irreducible and reducible moduli spaces in monopole Floer theory. To this end, we systematize the construction of maps associated to cobordisms equipped with certain polytopes of metrics, as well as the identities which count ends of 1-dimensional moduli spaces. This includes the construction of the usual monopole Floer differentials, cobordism maps, and homotopies as special cases, as well as the operators used in the proof of the surgery exact triangle, which we reorganize in Section 2.4. We also prove that the filtered homotopy type of the link surgery spectral sequence is independent of analytic choices, which may be viewed as a gauge-theoretic analogue of the invariance of A_∞ homotopy type in symplectic geometry [40]. In particular, the higher pages are themselves invariants of a framed link in a 3-manifold.

In addition, we equip the spectral sequence with an absolute mod 2 grading, which coincides with the absolute mod 2 grading on monopole Floer homology on the E^∞ page. Furthermore, the spectral sequence is defined for all three of the primary versions of monopole Floer homology, to be reviewed momentarily. In Section 2.6, we introduce a fourth version $\widetilde{HM}_\bullet$, analogous to $\widehat{HF}$, before extending the spectral sequence to this version as well.

We now briefly describe the remaining chapters, referring the reader to the start of each for detailed background and precise statements of theorems. Chapters 4 and 5 are concerned with Khovanov homology, a bigraded invariant of links in the 3-sphere which categorifies the Jones polynomial. In Chapter 4, we give an elementary proof that odd

Khovanov homology is invariant under Conway mutation. In Chapter 5, we derive a new formula for link signature that is well-adapted to Khovanov homology, and use it to recover a simple formula for the signature of an alternating link. We also give a new proof that the homological width of a k-almost alternating link is bounded above by $k+1$.

In Chapter 6, we apply the link surgery spectral sequence to relate the Khovanov homology of a link $L \subset S^3$ to the monopole Floer homology of the branched-double cover with reversed orientation, $-\Sigma(L)$. In particular, we prove that $\widetilde{HM}_\bullet(-\Sigma(L))$ arises via a filtered perturbation of the differential on the reduced Khovanov complex of a diagram of L. The associated spectral sequence carries a filtration grading, as well as a mod 2 grading which interpolates between the δ grading on Khovanov homology and the mod 2 grading on Floer homology. Furthermore, the bigraded isomorphism class of the higher pages depends only on the Conway-mutation equivalence class of L.

In Chapter 7, we discuss the relationship between Donaldson's TQFT, Khovanov homology, and monopole Floer homology, from both an algebraic and geometric point of view. By relating the module structure on Donaldson's TQFT to that on monopole Floer homology, we pin down the monopole maps associated to certain 0-framed 2-handle cobordisms between positive scalar curvature 3-manifolds. These cobordisms include those arising in the context of the spectral sequence from Khovanov homology to monopole Floer homology.

In Chapter 8, we use these maps to relate Khovanov homology to the other three versions of monopole Floer homology with non-trivial $U_\dagger$ action. This relationship is shown to constrains the existence of an integer bigrading and determine all monopole Floer groups of branched double-covers of links with thin Khovanov homology. We also reuse our proof of a bound on homological width to show that, in a sense, the differentials on the $\widetilde{HM}_\bullet$ spectral sequence decrease the δ grading. In the final section, we explain how the link surgery spectral sequence allows one to realize the monopole Floer homology of any 3-manifold Y as the homology of a complex over a combinatorial set of generators.

In Chapter 9, which serves as an Appendix, we review the model case of Morse homology on a manifold with boundary. In particular, we introduce a path algebra formalism to organize the contributions of interior and boundary trajectories. This formalism carries over to monopole Floer theory and motivates many of the constructions in Chapter 2.

Earlier versions of parts of this work appeared in [8] and [9].

1.1 Background on monopole Floer homology

In this section, we review those aspects most relevant to the construction and intuition in subsequent chapters. We refer the reader to [24] for the full construction of the monopole Floer groups (see also [27]) for an efficient survey). We will always work over the 2-element field $\mathbb{F}_2$.

Formal structure. Let COB be the category whose objects are compact, connected, oriented 3-manifolds and whose morphisms are isomorphism classes of connected cobordisms. Then the monopole Floer homology groups define covariant functors from the oriented cobordism category COB to the category $\mathrm{MOD}_\dagger$ of modules over $\mathbb{F}_2[[U_\dagger]]$, the ring of power series in a formal variable $U_\dagger$:

$$\widetilde{HM}_\bullet : \mathrm{COB} \to \mathrm{MOD}_\dagger$$
$$\widehat{HM}_\bullet : \mathrm{COB} \to \mathrm{MOD}_\dagger$$
$$\overline{HM}_\bullet : \mathrm{COB} \to \mathrm{MOD}_\dagger\,.$$

The module structure may be extended over the ring $\Lambda^*(H_1(Y)/\mathrm{torsion}) \otimes \mathbb{F}_2[[U_\dagger]]$. These modules fit into a long exact sequence

$$\cdots \xrightarrow{j_*} \widehat{HM}_\bullet(Y) \xrightarrow{p_*} \overline{HM}_\bullet(Y) \xrightarrow{i_*} \widetilde{HM}_\bullet(Y) \xrightarrow{j_*} \cdots \tag{1.1}$$

which is natural with respect to the maps induced by cobordisms. For $Y = S^3$, the map j_* is zero and the resulting short exact sequence of $\mathbb{F}_2[[U_\dagger]]$-modules is isomorphic to:

$$0 \longrightarrow \mathbb{F}_2[[U_\dagger]] \longrightarrow \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}] \longrightarrow \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]/\mathbb{F}_2[[U_\dagger]] \longrightarrow 0. \tag{1.2}$$

The monopole equations. We now describe the monopole equations underlying the construction of these groups, following [24]. Let Y be a closed, oriented Riemannian 3-manifold. A spin^c structure $\mathfrak{s}$ on Y is a pair (S, ρ) consisting of a unitary rank-2 vector bundle $S \to Y$ and a Clifford multiplication:

$$\rho : TY \to \mathrm{Hom}(S, S).$$

This map ρ identifies TY isometrically with the subbundle $\mathfrak{su}(S)$ of traceless, skew-adjoint endomorphisms equipped with the inner product $\frac{1}{2}\mathrm{tr}(a^*b)$, and satisfies

$$\rho(e_1)\rho(e_2)\rho(e_3) = 1$$

whenever the e_i form an oriented basis. The action of ρ extends to cotangent vectors using the metric, and to real (and complex) forms using the rule:

$$\rho(\alpha \wedge \beta) = \frac{1}{2}(\rho(\alpha)\rho(\beta) + (-1)^{\deg(\alpha)\deg(\beta)}\rho(\beta)\rho(\alpha)).$$

The set of isomorphism classes of spin^c structures on Y admits a free, transitive action of $H^2(Y;\mathbb{Z})$.

A unitary connection B on S is a *spin^c connection* if ρ is parallel. The space of spin^c connections is an affine space over $\Omega^1(Y;i\mathbb{R})$. In particular, the difference between two spin^c connections, regarded as 1-forms with values in the endomorphisms of S, has the form $a \otimes 1_S$ with $a \in \Omega^1(Y;i\mathbb{R})$. A section $\Psi \in \Gamma(S) = C^\infty(Y;S)$ is called a *spinor*. Let

$$\mathcal{C}(Y,\mathfrak{s}) = \{(B,\Psi) \mid B \text{ is a spin}^c \text{ connection and } \Psi \in \Gamma(S)\}.$$

The gauge gauge group $\mathcal{G} = C^\infty(Y;S^1)$ acts on this space by conjugation and multiplication:

$$u \cdot (B,\Psi) = (B - u^{-1}du \otimes 1_S,\ u\Psi).$$

Given a spin^c connection B, let $D_B : \Gamma(S) \to \Gamma(S)$ denote the associated Dirac operator:

$$\Gamma(S) \xrightarrow{\nabla_B} \Gamma(T^*Y \otimes S) \xrightarrow{\rho} \Gamma(S).$$

Let B^t denote the associated connection on the complex line bundle $\Lambda^2 S$, with curvature F_{B^t} regarded as an imaginary-valued 2-form. In particular, $\rho(F_{B^t})$ represents a trace-free Hermitian endomorphism. Fix a reference connection $B_0 \in \mathcal{A}$. The *Chern-Simons-Dirac functional* $\mathcal{L} : \mathcal{C}(Y,\mathfrak{s}) \to \mathbb{R}$ is defined by

$$\mathcal{L}(B,\Psi) = \frac{1}{8}\int_Y (B^t - B_0^t) \wedge (F_{B^t} + F_{B_0^t}) + \frac{1}{2}\int_Y \langle D_B\Psi,\Psi\rangle d\,\mathrm{vol}.$$

The domain of $\mathcal{L}$ is an affine space over the vector space

$$T_{(B,\Phi)}\mathcal{C}(Y,\mathfrak{s}) = C^\infty(Y; iT^*Y \oplus S).$$

The formal gradient of $\mathcal{L}$ with respect to the L^2 inner product vanishes precisely when the following equations are satisfied:

$$\frac{1}{2}\rho(F_{B^t}) - (\Psi\Psi^*)_0 = 0$$
$$D_B\Psi = 0$$

Here $(\Psi\Psi^*)_0 \in \Gamma(i\,\mathfrak{su}(S))$ denotes the trace-free part of the Hermitian endomorphism $\Psi\Psi^*$. These are the 3-dimensional *monopole equations*, or *Seiberg-Witten equations*, on Y for the spin^c structure $\mathfrak{s}$. The solutions, regarded as critical points of $\mathcal{L}$, are called *monopoles*, and the action of the gauge group sends monopoles to monopoles.

Reducibles and the blow-up. A configuration (B, Ψ) is *reducible* if Ψ is zero. If $(B, 0)$ is a solution to the monopole equations, then B^t is flat and $c_1(\mathfrak{s})$ is torsion. Conversely, if $c_1(\mathfrak{s})$ is torsion then there exists a reducible solution $(B_1, 0)$, and all others are of the form $(B, 0)$ with $B = B_1 + b \otimes 1_S$ and b a closed element of $\Omega^1(Y; i\mathbb{R})$. The action of the gauge group changes b by representatives of elements of $2\pi i H^1(Y;\mathbb{Z})$. In particular, the quotient of the set of reducible solutions by the action of the gauge group is identified with the torus $\mathbb{T} = H^1(Y; i\mathbb{R})/(2\pi i H^1(Y;\mathbb{Z}))$, and consists of a single point when $b_1(Y) = 0$.

The constant elements of the gauge group fix the reducible configurations. To obtain a free action, we blow-up the configuration space $\mathcal{C}(Y, \mathfrak{s})$ along the reducible locus to obtain

$$\mathcal{C}^\sigma(Y, \mathfrak{s}) = \{(B, r, \psi) \mid B \text{ is a spin}^c \text{ connection, } r \geq 0, \text{ and } \|\psi\|_{L^2} = 1\}$$

where blow-down sends (B, r, ψ) to $(B, r\psi)$. Here r is a real number and ψ is a spinor. As discussed in Section 9 of [24], the completion (which we suppress) of $\mathcal{C}^\sigma(Y, \mathfrak{s})$ with respect to suitable Sobolev norms L^2_k has the structure of a Hilbert manifold with boundary. The same is true of the quotient

$$\mathcal{B}^\sigma(Y, \mathfrak{s}) = \mathcal{C}^\sigma(Y, \mathfrak{s})/\mathcal{G}(Y).$$

with the boundary consisting of (equivalence classes of) configurations of the form $(B, 0, \psi)$. This quotient has the homotopy type of $\mathbb{T} \times \mathbb{CP}^\infty$, and there is a canonical identification of cohomology rings

$$H^*(\mathcal{B}^\sigma(Y, \mathfrak{s})) = H^*(\mathbb{T}) \otimes \mathbb{F}_2[U].$$

giving rise to the module structure on monopole Floer homology, as we describe in Section 7.2.

The Chern-Simons-Dirac functional $\mathcal{L}$ is invariant under the identity component of the gauge group (and the full gauge group when $c_1(\mathfrak{s})$ is torsion). Its gradient gives rise to a vector field $(\text{grad } \mathcal{L})^\sigma$ on $\mathcal{B}^\sigma(Y, \mathfrak{s})$. The configuration (B, r, ψ) is a critical point of $(\text{grad } \mathcal{L})^\sigma$ if and only if one of the following conditions holds:

(i) $r \neq 0$ and $(B, r\psi)$ is a critical point of grad $\mathcal{L}$; or

(ii) $r = 0$, the point $(B, 0)$ is a critical point of grad $\mathcal{L}$, and ϕ is an eigenvector or D_B.

Critical points of type (i) are called *irreducible*, while those of type (ii) are called *reducible*. A reducible critical point is *boundary stable* (resp., *boundary unstable*) if the corresponding eigenvalue is positive (resp., negative).

In finite-dimensional Morse homology, one may achieve the transversality needed to apply Sard's theorem by perturbing the Morse function. In the monopole setting, one may similarly achieve the transversality necessary for Sard-Smale by perturbing the functional $\mathcal{L}$ by a function $q : \mathcal{C}^\sigma(Y, \mathfrak{s}) \to \mathbb{R}$ which is invariant under the full gauge group. Kronheimer and Mrowka define a Banach space of perturbations q, a residual subset of which force all critical points and moduli spaces of gradient trajectories between them to be regular in an appropriate sense. Such perturbations are called admissible. In particular, for an admissible perturbation, zero does not arise as the eigenvalue of a reducible critical point.

Monopole Floer complex. The construction of monopole Floer homology is modeled on that of Morse homology for a manifold with boundary. The latter is described in the Appendix and in Section 2 of [24]. In place of the downward gradient flow of a Morse function on finite-dimensional manifold with boundary, we have the downward gradient flow of the perturbed Chern-Simons-Dirac functional on the configuration space $\mathcal{B}^\sigma(Y, \mathfrak{s})$ whose boundary consists of reducible configurations. Having chosen an admissible perturbation, let $\mathfrak{C}(Y, \mathfrak{s})$ denote the set of critical points in $\mathcal{B}^\sigma(Y, \mathfrak{s})$. We may express this set as a disjoint union

$$\mathfrak{C}(Y, \mathfrak{s}) = \mathfrak{C}^o(Y, \mathfrak{s}) \cup \mathfrak{C}^s(Y, \mathfrak{s}) \cup \mathfrak{C}^u(Y, \mathfrak{s})$$

where $\mathfrak{C}^o(Y, \mathfrak{s})$ is the set of irreducible critical points, and $\mathfrak{C}^s(Y, \mathfrak{s})$ and $\mathfrak{C}^u(Y, \mathfrak{s})$ are the sets of boundary-stable and boundary-unstable critical points, respectively. We set

$$\check{\mathfrak{C}}(Y,\mathfrak{s}) = \mathfrak{C}^o(Y,\mathfrak{s}) \cup \mathfrak{C}^s(Y,\mathfrak{s})$$
$$\hat{\mathfrak{C}}(Y,\mathfrak{s}) = \mathfrak{C}^o(Y,\mathfrak{s}) \cup \mathfrak{C}^u(Y,\mathfrak{s})$$
$$\bar{\mathfrak{C}}(Y,\mathfrak{s}) = \mathfrak{C}^s(Y,\mathfrak{s}) \cup \mathfrak{C}^u(Y,\mathfrak{s})$$

The monopole Floer complex $\check{C}(Y,\mathfrak{s})$ is the $\mathbb{F}_2$-vector space over the basis $e_{\mathfrak{a}}$ indexed by (irreducible or boundary-stable) monopoles[1] $\mathfrak{a} \in \check{\mathfrak{C}}(Y,\mathfrak{s})$. Given two such critical points $\mathfrak{a}$ and $\mathfrak{b}$, we may consider the moduli space $\breve{M}_z(\mathfrak{a},\mathfrak{b})$ of unparameterized (downward) gradient trajectories (mod gauge) from $\mathfrak{a}$ to $\mathfrak{b}$ in the relative homotopy class z of path from $\mathfrak{a}$ to $\mathfrak{b}$ in $\mathcal{B}^\sigma(Y,\mathfrak{s})$. The differential $\check{\partial}$ is defined to count isolated trajectories in such moduli spaces. In particular, when $\mathfrak{a}$ is irreducible, the coefficient of $e_{\mathfrak{b}}$ in $\check{\partial}(e_{\mathfrak{a}})$ is the number of trajectories in $\breve{M}_z(\mathfrak{a},\mathfrak{b})$, summed over all z such that $\breve{M}_z(\mathfrak{a},\mathfrak{b})$ is 0-dimensional. When $\breve{M}_z(\mathfrak{a},\mathfrak{b})$ is 1-dimensional, it has a compactification $\breve{M}_z^+(\mathfrak{a},\mathfrak{b})$ formed by considering broken trajectories as well. The composition $\check{\partial}^2$ then counts the (even) number of boundary points, proving that $\check{\partial}$ is a differential. The full construction of $\check{\partial}$, which is complicated by the presence of reducible critical points, is given in Section 2.2 as the simplest case of a more general construction.

We now set

$$\widecheck{HM}_*(Y,\mathfrak{s}) = H_*(\check{C}(Y,\mathfrak{s}),\check{\partial})$$

and

$$\widecheck{HM}_*(Y) = \bigoplus_{\mathfrak{s}} \widecheck{HM}_*(Y,\mathfrak{s})$$

where the sum is over all spin^c structures on Y. The group $\widecheck{HM}_*(Y)$ is graded by the set of homotopy classes of oriented 2-plane fields on Y. This set admits a natural action of $\mathbb{Z}$, the orbits of which correspond to the different spin^c structures. The group $\widecheck{HM}_\bullet(Y)$ is defined as the completion of $\widecheck{HM}_*(Y)$ with respect to a decreasing filtration defined using this $\mathbb{Z}$ action (see Definition 3.1.3 in [24] for details). The groups $\widehat{HM}_\bullet(Y)$ and $\overline{HM}_\bullet(Y)$

[1] In [24], the notation $[\mathfrak{a}]$ is used to denote the gauge equivalence class of the configuration $\mathfrak{a} \in \mathcal{C}^\sigma(Y,\mathfrak{s})$. We will always consider critical points on the level of the quotient $\mathcal{B}^\sigma(Y,\mathfrak{s})$ and have dropped the brackets to simplify notation.

are defined similarly using $\hat{\mathfrak{C}}$ and $\bar{\mathfrak{C}}$. Of the three versions, the group $\overline{HM}_{\bullet}(Y)$ is both the simplest to define and the best understood (see Section 35 of [24], especially Proposition 35.1.5). As all critical points and trajectories for $\bar{C}(Y,\mathfrak{s})$ are taken to be reducible — that is, in $\partial\mathcal{B}^{\sigma}(Y,\mathfrak{s})$ — the model case reduces to Morse homology on a closed manifold (namely, the boundary of a manifold with boundary).

Chain maps. There is a fundamental correspondence between gradient trajectories of functional $\mathcal{L}$ in $\mathcal{B}^{\sigma}(Y,\mathfrak{s})$ and solutions (mod gauge) to the 4-dimensional monopole equations on $Y\times\mathbb{R}$ for the corresponding spin^c structure. The latter 4-dimensional interpretation of trajectories underlies the construction of the chain map $\check{m}(W):\check{C}(Y_0)\to\check{C}(Y_1)$ associated to a general cobordism W. Having chosen a metric on W which is cylindrical near the boundary, we denote by W^* the Riemannian manifold built by attaching the half-infinite cylinders $(-\infty,0]\times Y_0$ and $[0,\infty)\times Y_1$ to the ends of W. For monopoles $\mathfrak{a}\in\check{\mathfrak{C}}(Y_0)$ and $\mathfrak{b}\in\check{\mathfrak{C}}(Y_1)$, and a relative homotopy class z from $\mathfrak{a}$ to $\mathfrak{b}$ in the configuration space $\mathcal{B}^{\sigma}(W)$, we consider the moduli space $M_z(\mathfrak{a},W^*,\mathfrak{b})$ of trajectories (mod gauge) on W^* asymptotic to $\mathfrak{a}$ and $\mathfrak{b}$ and in the class z. The map $\check{m}(W)$ is defined to count isolated trajectories in such moduli spaces. In particular, when $\mathfrak{a}$ is irreducible, the coefficient of $e_{\mathfrak{b}}$ in $\check{m}(W)(e_{\mathfrak{a}})$ is the number of trajectories in $M_z(\mathfrak{a},W^*,\mathfrak{b})$, summed over all z such that $M_z(\mathfrak{a},W^*,\mathfrak{b})$ is 0-dimensional. When $M_z(\mathfrak{a},W^*,\mathfrak{b})$ is 1-dimensional, it has a compactification $M_z^+(\mathfrak{a},W^*,\mathfrak{b})$ formed by considering broken trajectories as well. The composite maps $\check{\partial}\check{m}(W)$ and $\check{m}(W)\check{\partial}$ then count the (even) number of boundary points, so

$$\check{\partial}\check{m}(W)+\check{m}(W)\check{\partial}=0,$$

and we conclude that $\check{m}(W)$ is a chain map.

Families of metrics. More generally, suppose we have a family of metrics on W, smoothly parameterized by a closed, oriented[2] manifold P. The map $\check{m}(W)_P:\check{C}(Y_0)\to\check{C}(Y_1)$ is defined to count isolated trajectories in the union

$$M(\mathfrak{a},W^*,\mathfrak{b})_P=\bigcup_z M_z(\mathfrak{a},W^*,\mathfrak{b})_P \tag{1.3}$$

[2]The orientation is irrelevant when working over $\mathbb{F}_2$.

of fiber products

$$M_z(\mathfrak{a}, W^*, \mathfrak{b})_P = \bigcup_{p\in P} M_z(\mathfrak{a}, W(p)^*, \mathfrak{b}), \tag{1.4}$$

where $W(p)$ denotes W with the metric over p. The compact fiber product $M_z^+(\mathfrak{a}, W^*, \mathfrak{b})_P$ is defined similarly. By counting boundary points of $M_z(\mathfrak{a}, W^*, \mathfrak{b})_P$, we again conclude

$$\check{\partial}\check{m}(W)_P + \check{m}(W)_P\check{\partial} = 0.$$

On the other hand, if P is a compact manifold with boundary Q, then $\check{m}(W)_P$ is no longer a chain map, because the boundary of $M_z(\mathfrak{a}, W^*, \mathfrak{b})_P$ now includes the fibers over Q. Including these contributions, we have

$$\check{\partial}\check{m}(W)_P + \check{m}(W)_P\check{\partial} = \check{m}(W)_Q. \tag{1.5}$$

Thus, $\check{m}(W)_Q$ is null-homotopic and $\check{m}(W)_P$ provides the chain homotopy. That $\widecheck{HM}_\bullet(Y)$ is independent of the choice of metric and perturbation follows by letting P be the interval $[0,1]$ parameterizing a path between two such choices.

Composing cobordisms. If $W : Y_0 \to Y_2$ is the composition of cobordisms $W_1 : Y_0 \to Y_1$ and $W_2 : Y_1 \to Y_2$, then the corresponding maps satisfy the composition law

$$\widecheck{HM}_\bullet(W) = \widecheck{HM}_\bullet(W_2) \circ \widecheck{HM}_\bullet(W_1). \tag{1.6}$$

Indeed, this is part of what it means for $\widecheck{HM}_\bullet$ to be a functor. The composition law follows from a "stretching the neck" argument, as do many of the results in this paper, so we now take a moment to review the proof (see Proposition 4.16 of [27] for details over $\mathbb{F}_2$, and Proposition 26.1.2 of [24] for details over $\mathbb{Z}$). Keep in mind that the full argument is complicated by the presence of reducibles. We deal with this issue in Section 2.2.

Returning to the composite cobordism

$$W : Y_0 \xrightarrow{W_1} Y_1 \xrightarrow{W_2} Y_2,$$

fix a metric on W which is cylindrical near each Y_i. For each $T \geq 0$, we construct a new Riemannian cobordism $W(T)$ by cutting W along Y_1 and splicing in the cylinder $[-T,T]\times Y_1$ with the cylindrical metric. We also define $W(\infty)$ as the disjoint union $W_1 \coprod W_2$. In this way, $P = [0, \infty]$ parameterizes a family of metrics on W, where the metric degenerates on Y_1 at infinity. In other words, as T increases, the cylindrical neck stretches, and when $T = \infty$, it breaks.

We again define $\check{m}(W)_P$ to count isolated trajectories in the fiber products $M_z(\mathfrak{a}, W^*, \mathfrak{b})_P$ of (1.3), where now

$$M_z(\mathfrak{a}, W(\infty)^*, \mathfrak{b}) = \bigcup_{\mathfrak{c} \in \check{C}(Y_1)} \bigcup_{z_1, z_2} M_{z_1}(\mathfrak{a}, W_1^*, \mathfrak{c}) \times M_{z_2}(\mathfrak{c}, W_2^*, \mathfrak{b}), \tag{1.7}$$

and the inner union is taken over homotopy classes z_1 and z_2 which concatenate to give z. The compact fiber product $M^+(\mathfrak{a}, W^*, \mathfrak{b})_P$ is defined similarly. By counting boundary points, we conclude

$$\check{\partial}\check{m}(W)_P + \check{m}(W)_P\check{\partial} = \check{m}(W) + \check{m}(W_2)\check{m}(W_1). \tag{1.8}$$

Here $\check{m}(W)$ and $\check{m}(W_2)\check{m}(W_1)$ count trajectories in the fibers over 0 and ∞, respectively. Viewing $\check{m}(W)_P$ as a chain homotopy, the composition law now follows. Note that, while formally similar, (1.5) does not imply (1.8) because the latter involves a degenerate metric. The key analytic machinery behind this generalization consists of compactness and gluing theorems for moduli spaces on cobordisms with cylindrical ends, as developed in [24] and [27]. Our workhorse version of this machinery is Lemma 2.2.3 in Section 2.2.

Canonical gradings. Recall that the group $\widetilde{HM}_\bullet(Y)$ is naturally graded by the set of homotopy classes of oriented 2-plane fields. We will make use of two numerical gradings which factor through this set. The first is an absolute mod 2 grading $\mathrm{gr}^{(2)}$, as explained in Sections 22.4 and 25.4 of [24]. If W is a cobordism from Y_0 to Y_1, then the degree of the map $\widetilde{HM}_\bullet(W)$ with respect to $\mathrm{gr}^{(2)}$ is given by[3]

$$\iota(W) = \frac{\chi(W) + \sigma(W) - b_1(Y_1) + b_1(Y_0)}{2} \tag{1.9}$$

[3]This agrees with [27], but in [24] the signs on $b_1(Y_0)$ and $b_1(Y_1)$ are switched. The value of $-\iota(W)$ should be the index of the operator $d^* \oplus d^+$ acting on weighted Sobolev spaces (see Section 25.4 of [24]). The two formulas for ι correspond to the two choices for the sign of this weight δ. Different conventions lead to mirror theories. We believe the formula (1.9) corresponds to the choice of a small, positive weight. In any case, we take whichever convention is consistent with (1.9) and use it consistently throughout.

where χ is the Euler characteristic and σ is the signature of the intersection form on

$$I^2(W) = \mathrm{Im}\left(H^2(W, \partial W) \to H^2(W)\right).$$

If P parameterizes an n-dimensional family of metrics on W, then the map $\check{m}(W)_P$ shifts $\mathrm{gr}^{(2)}$ by $\iota(W) + n$.

The second numerical grading is only defined if $c_1(\mathfrak{s}) \in H^2(Y;\mathbb{Z})$ is torsion. In this case $\widetilde{HM}_\bullet(Y, \mathfrak{s})$ is also endowed with an absolute grading $\mathrm{gr}^{\mathbb{Q}}$ which takes values in a $\mathbb{Z}$ coset of $\mathbb{Q}$. If $(W, \mathfrak{t}) : (Y_0, \mathfrak{t}|_{Y_0}) \to (Y_1, \mathfrak{t}|_{Y_1})$ is a spin^c cobordism with $c_1(\mathfrak{t})|_{\partial W}$ torsion, then the degree of $\widetilde{HM}_\bullet(W, \mathfrak{t})$ with respect to $\mathrm{gr}^{\mathbb{Q}}$ is given by

$$d(W, \mathfrak{t}) = \frac{c_1^2(\mathfrak{t}) - \sigma(W)}{4} - \iota(W).$$

By (1.9) we may also express this degree as

$$d(W, \mathfrak{t}) = \frac{c_1^2(\mathfrak{t}) - 2\chi(W) - 3\sigma(W)}{4} + \frac{b_1(Y_1) - b_1(Y_0)}{2}. \tag{1.10}$$

Here the rational number $c_1^2(\mathfrak{t})$ is defined by

$$c_1^2(\mathfrak{t}) = (\tilde{c} \cup \tilde{c})[W, \partial W]$$

where $\tilde{c}$ is any class in $H^2(W, \partial W; \mathbb{Q})$ whose image in $H^2(W;\mathbb{Q})$ is the same as the image of $c_1(\mathfrak{t})$. If P parameterizes an n-dimensional family of metrics on W, then the map $\check{m}(W, \mathfrak{t})_P$ shifts $\mathrm{gr}^{\mathbb{Q}}$ by $d(W, \mathfrak{t}) + n$.

The gradings $\mathrm{gr}^{(2)}$ and $\mathrm{gr}^{\mathbb{Q}}$ are also defined on $\widehat{HM}_\bullet(Y, \mathfrak{s})$, and there are modified versions $\bar{\mathrm{gr}}^{(2)}$ and $\bar{\mathrm{gr}}^{\mathbb{Q}}$ defined on $\overline{HM}_\bullet(Y, \mathfrak{s})$. In each case, the degree of a cobordism map is given by the above formulas. With respect to $\mathrm{gr}^{(2)}$ and $\mathrm{gr}^{\mathbb{Q}}$ (when the latter is defined), the monopole Floer groups are graded modules over the graded ring $\mathbb{F}_2[[U_\dagger]]$, with $U_\dagger$ in degree -2. In the exact sequence (1.1), the maps i_* and j_* have degree 0, while p_* has degree -1. For S^3, the associated short exact sequence of $\mathrm{gr}^{\mathbb{Q}}$-graded $\mathbb{F}_2[[U_\dagger]]$-modules is isomorphic to

$$0 \longrightarrow \mathbb{F}_2[[U_\dagger]]\{-1\} \longrightarrow \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]\{-2\} \longrightarrow \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]/\mathbb{F}_2[[U_\dagger]]\{-2\} \longrightarrow 0$$

where $\{-k\}$ shifts the degree of each generator down by k. In particular, the "top" generator of $\widehat{HM}_\bullet(S^3)$, represented by 1, lies in degree -1, while the "bottom" generator of $\widetilde{HM}_\bullet(S^3)$, represented by $U_\dagger^{-1}$, lies in degree 0.

Chapter 2

The link surgery spectral sequence

In order to motivate the statement of the link surgery spectral sequence, we first recall the surgery exact triangle. Let Y be a closed, oriented 3-manifold, equipped with a knot K with framing λ and meridian μ. Orient λ and μ as simple closed curves on the torus boundary of the complement of a neighborhood of K, so that the algebraic intersection numbers of the triple $(\lambda, \lambda+\mu, \mu)$ satisfy

$$(\lambda \cdot (\lambda+\mu)) = ((\lambda+\mu)\cdot \mu) = (\mu \cdot \lambda) = -1.$$

Let $Y(0)$ and $Y(1)$ denote the result of surgery on K with respect to λ and $\lambda+\mu$, respectively. In [27], Kronheimer, Mrowka, Ozsváth, and Szabó prove that the mapping cone

$$\check{C}(Y(0)) \xrightarrow{\check{m}(W(01))} \check{C}(Y(1))$$

is quasi-isomorphic to the monopole Floer complex $\check{C}(Y)$, where $\check{m}(W(01))$ is the chain map induced by the elementary 2-handle cobordism $W(01)$ from $Y(0)$ to $Y(1)$. The associated long exact sequence on homology is known as the surgery exact triangle. However, we can also frame the result in another way. As in [36], if we filter by the index I in $Y(I)$, then the mapping cone induces a spectral sequence with

$$E^1 = \widetilde{HM}(Y(0)) \bigoplus \widetilde{HM}(Y(1))$$

and

$$d^1 = \widetilde{HM}_\bullet(W(01)),$$

which converges by the E^2 page to $\widetilde{HM}(Y)$.

The link surgery spectral sequence generalizes this interpretation of the exact triangle to the case of an l-component framed link $L \subset Y$. For each $I = (m_1, \ldots, m_l)$ in the hypercube $\{0,1\}^l$, let $Y(I)$ denote the result of performing m_i-surgery on the component K_i. For $I < J$, let $W(IJ)$ denote the associated cobordism, composed of $(w(J) - w(I))$ 2-handles. The (iterated) mapping cone now takes the form of a hypercube complex

$$X = \bigoplus_{I \in \{0,1\}^l} \check{C}(Y(I))$$

with differential $\check{D}$ given by the sum of components $\check{D}^I_J : \check{C}(Y(I)) \to \check{C}(Y(J))$ for all $I \leq J$. We filter X by vertex weight $w(I)$, defined as the sum of the coordinates of I. The component $\check{D}^I_I$ is the usual differential on $\check{C}(Y(I))$, whereas for $I < J$, the component $\check{D}^I_J$ counts monopoles on $W(IJ)$ over a family of metrics parametrized by the permutohedron of dimension $w(J) - w(I) - 1$. We define this family in Section 2.1 and construct $(X, \check{D})$ in Section 2.2. In Section 2.5, we complete the proof of:

Theorem 2.0.1. *The filtered complex $(X, \check{D})$ induces a spectral sequence with E^1 page given by*

$$E^1 = \bigoplus_{I \in \{0,1\}^l} \widetilde{HM}_\bullet(Y(I))$$

and d^1 differential given by

$$d^1 = \sum_{\substack{I < J \in \{0,1\}^l \\ w(J) - w(I) = 1}} \widetilde{HM}_\bullet(W(IJ)).$$

The spectral sequence converges by the E^{l+1} page to $\widetilde{HM}_\bullet(Y)$ and comes equipped with an absolute mod 2 grading $\check{\delta}$ which coincides on E^∞ with that of $\widetilde{HM}_\bullet(Y)$. In addition, each page has an integer grading $\check{t}$ induced by the filtration. The differential d^k shifts $\check{\delta}$ by one and increases $\check{t}$ by k.

The complex $(X, \check{D})$ depends on a family of metrics and an admissible family of perturbations on the full cobordism from $Y(\{0\}^l)$ to $Y(\{1\}^l)$. For any two such choices, we produce a homotopy equivalence which induces a graded isomorphism between the associated E^1 pages.

Theorem 2.0.2. *For each $i \geq 1$, the $(\check{t}, \check{\delta})$-graded vector space E^i is an invariant of the framed link $L \subset Y$.*

Remark 2.0.3. While the preceding theorems are stated for $\widetilde{HM}_\bullet$, they hold for $\widehat{HM}_\bullet$ and $\overline{HM}_\bullet$ as well, just like the underlying surgery exact triangle. In the case of $\overline{HM}_\bullet$, we must also replace $\mathrm{gr}^{(2)}$ with $\bar{\mathrm{gr}}^{(2)}$ and similarly $\check{\delta}$ with its analogue $\bar{\delta}$ defined using $\bar{\mathrm{gr}}^{(2)}$.

In Section 2.6, we introduce another version of monopole Floer homology, pronounced "H-M-tilde" and denoted $\widetilde{HM}_\bullet$. By analogy with $\widehat{HF}$ in Heegaard Floer homology, we define $\widetilde{HM}_\bullet(Y)$ as the homology of the mapping cone of $U_\dagger : \check{C}(Y) \to \check{C}(Y)[1]$, where $U_\dagger$ is the even endomorphism on $\widetilde{HM}_\bullet(Y)$ given by the module structure. It follows that $\widetilde{HM}_\bullet(Y)$ inherits a mod 2 grading, and we prove a version of Theorem 2.0.1 for $\widetilde{HM}_\bullet$ as well.

In fact, the group $\widetilde{HM}_\bullet(Y, \mathfrak{s})$ agrees with the sutured monopole Floer homology group $SFH(Y - B^3, \mathfrak{s})$ relative to the equatorial suture. The latter is defined in [25] as

$$\widetilde{HM}_\bullet(Y\#(S^1 \times F), \mathfrak{s}\#\mathfrak{s}_c),$$

where F is an orientable surface of genus $g \geq 2$ and $\mathfrak{s}_c$ is the canonical spin^c-structure with $\langle c_1(\mathfrak{s}_\mathfrak{c}), [F]\rangle = 2g - 2$. This equivalence follows from a Künneth formula for connected sum in monopole Floer homology, to appear in joint work with Tomasz Mrowka and Peter Ozsváth [10].

2.1 Hypercubes and permutohedra

This section involves no Floer homology whatsoever, but rather surgery theory and Kirby calculus as described in Part 2 of [19]. In particular, with respect to a 2-handle $D^2 \times D^2$, the terms *core*, *cocore*, and *attaching region* will refer to the subsets $D^2 \times \{0\}$, $\{0\} \times D^2$, and $\partial D^2 \times D^2$, respectively.

Let Y be a closed, oriented 3-manifold, equipped with an l-component, framed link $L = K_1 \cup \cdots \cup K_l$, and let Y' denote the result of (integral) surgery on L. There is a standard oriented cobordism $W : Y \to Y'$, built by thickening Y to $[0, 1] \times Y$ and attaching 2-handles h_i to $\{1\} \times Y$ by identifying the attaching region of h_i with a tubular neighborhood $\nu(K_i)$ in accordance with the framing. The diffeomorphism type of W is insensitive to whether the handles are attached simultaneously as above, or instead in a succession of batches which express W as a composite cobordism. Our goal in this section is to construct a family of metrics on W, parameterized by the permutohedron P_l, which smoothly interpolates between all ways of expressing W as a composite cobordism.

In order to keep track of the $l!$ ways to build up W one handle at a time, we introduce the hypercube poset $\{0,1\}^l$, with $I = (m_1, \ldots, m_l) \leq J = (m'_1, \ldots, m'_l)$ if and only if $m_i \leq m'_i$ for all $1 \leq i \leq l$. J is called an *immediate successor* of I if there is a k such that $m_k = 0$, $m'_k = 1$, and $m_i = m'_i$ for all $i \neq k$. We define a *path* of length k from I to J to be a sequence of immediate successors $I = I_0 < I_1 < \cdots < I_k = J$. The *weight* of a vertex I is given by $w(I) = \sum_{i=1}^{l} m_i$. We use *0* and *1* as shorthand for the initial and terminal vertices of $\{0,1\}^l$, which we call *external*. The other $2^l - 2$ vertices will be called *internal*. A totally ordered subset of a poset is called a *chain*. A chain is *maximal* if it is not properly contained in any other chain. In $\{0,1\}^l$, the maximal chains are precisely the paths from *0* to *1*, with each such path determined by its internal vertices.

To each vertex I, we associate the 3-manifold Y_I obtained by surgery on the framed sublink

$$L(I) = \bigcup_{\{i \,|\, m_i = 1\}} K_i$$

in Y. Note that the remaining components of L constitute a framed link in Y_I.

Remark 2.1.1. The 3-manifold denoted $Y(I)$ in the introduction and in [36] is obtained from Y_I by shifting forward one frame in the surgery exact triangle for each component of L. We will use Y_I throughout and address this discrepancy in Remark 2.2.12.

We regard $\{Y_I \,|\, I \in \{0,1\}^l\}$ as a poset isomorphic to $\{0,1\}^l$, with Y_{0} and Y_{1} external and the rest internal. To a pair of vertices (I, J) with $I < J$, we associate the 2-handle cobordism

$$W_{IJ} = Y_I \times [0,1] \cup \bigcup_{\{i \,|\, m_i = 0,\, m'_i = 1\}} h_i$$

from Y_I to Y_J. In particular, if J is an immediate successor of I, then W_{IJ} is an *elementary cobordism*, given by a single 2-handle addition. More generally, W_{IJ} will be the composition of $w(J) - w(I)$ elementary cobordisms.

In order to quantify how far two vertices are from being ordered, we define a symmetric function ρ on pairs of vertices by

$$\rho(I,J) = \min\left\{\left|\{i \,|\, m_i > m_i'\}\right|, \left|\{i \,|\, m_i' > m_i\}\right|\right\}.$$

Note that $\rho(I,J) = 0$ if and only if I and J are ordered. In this case, Y_I and Y_J are disjoint:

Lemma 2.1.2. *The full set of $2^l - 2$ internal hypersurfaces Y_I can be simultaneously embedded in the interior of the cobordism W so that the following conditions hold:*

(i) the hypersurfaces in any subset are pairwise disjoint if and only if they form a chain. In this case, cutting on $Y_{I_1} < Y_{I_2} < ... < Y_{I_k}$ breaks W into the disjoint union

$$W_{0I_1} \coprod W_{I_1 I_2} \coprod \cdots \coprod W_{I_k 1}.$$

(ii) distinct hypersurfaces Y_I and Y_J intersect in exactly $\rho(I,J)$ disjoint tori.

Remark 2.1.3. The reader who is convinced by Figure 2.1 may safely skip the proof.

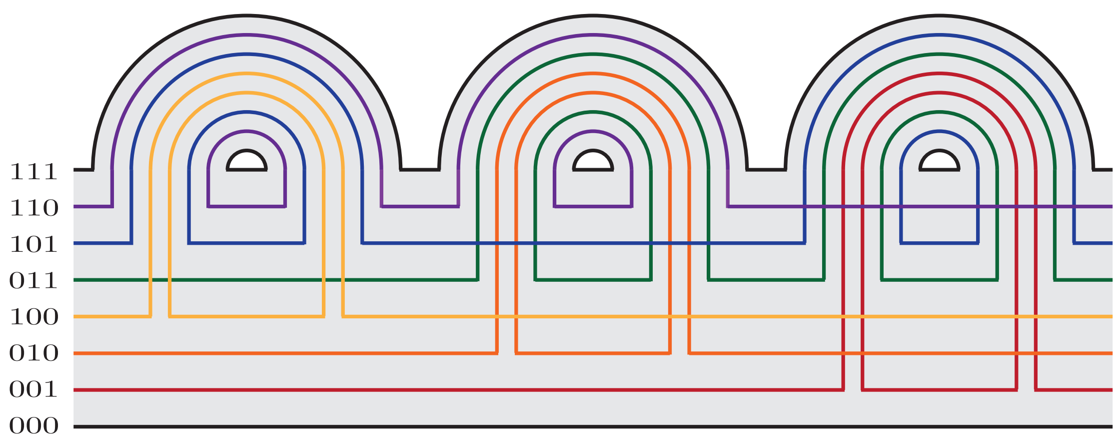

Figure 2.1: Half-dimensional diagram of the cobordism W for the hypercube $\{0,1\}^3$.

Proof. We list all of the vertices as $I_0, I_1, ..., I_{2^l-1}$, first in order of increasing weight and then numerically within each weight class. We express the full cobordism as

$$W = [0, 2^l - 1] \times Y \cup \bigcup_{i=1}^{l} h_i$$

and embed Y_0 and Y_1 as the boundary. We then embed the interior hypersurfaces as follows. For $1 \leq q \leq 2^l - 2$, define a slimmer 2-handle h_i^q as the image of $D^2 \times D_q^2$ in h_i, where D_q^2 is the disk of radius $\frac{q}{2^l}$. Let $\nu_q(K_i)$ be the region to which h_i^q is attached, considered as a subset of Y. Then we may regard

$$\tilde{h}_i^q = [q, 2^l - 1] \times \nu_q(K_i) \cup \bigcup_{\{2^l-1\}\times\nu_q(K_i)} h_i^q$$

as a longer 2-handle which tunnels through $[q, 2^l - 1] \times Y$ in order to attach to $[0, q] \times Y$ along $\{q\} \times v_q(K_i)$. In this way, we embed W_{0I_q} in W as

$$W_{0I_q} = [0, q] \times Y \cup \bigcup_{\{i \,|\, m_i = 1\}} \tilde{h}_i^q$$

and Y_{I_q} as a component of the boundary.

Now consider two vertices $I_q = (m_1, ..., m_l)$ and $I_{q'} = (m_1', ..., m_l')$ and assume without loss of generality that $q < q'$. By construction, $Y_{I_q} \cap Y_{I_{q'}}$ is confined to the union of the thickened attaching regions $[q, q'] \times \nu(K_i)$ in $[q, q'] \times Y$ with $m_i = 1$. If $m_i' = 1$ as well, then $\tilde{h}_i^q$ is contained in the interior of $W_{0q'}$. On the other hand, if $m_i > m_i'$ then $\tilde{h}_i^q$ and $\partial W_{0q'}$ intersect in the solid torus $\{q'\} \times \nu_q(K_i)$. It follows that Y_{I_q} and $Y_{I_{q'}}$ intersect in one torus for each i such that $m_i > m_i'$. With $q < q'$, the number of such i is exactly $\rho(I_q, I_q')$, verifying (ii). The first part of (i) immediately follows, since a subset of $\{0, 1\}^l$ forms a chain if and only if ρ vanishes on every pair of vertices in the subset. In this case, W decomposes as claimed by construction. $\square$

We are now ready to build a special family of Reimannian metrics on the cobordism W. We first construct an initial metric g_0 on W that is cylindrical near every Y_I simultaneously (for a less restrictive approach, allowing one to define metrics on each hypersurface independently, see Remark 2.5.6). We build g_0 inductively on strata, starting with an arbitrary metric on each (transverse) intersection $Y_I \cap Y_J$. We then use a partition of unity to piece together a metric on the union of the Y_I that is locally cylindrical near each intersection $Y_I \cap Y_J$. Finally, we build a metric g_0 on W that is cylindrical near each Y_I. In particular, a neighborhood $\nu(T^2) \subset W$ of a torus $T^2 \subset Y_I \cap Y_J$ is metrically modeled on $T^2 \times (-\varepsilon, \varepsilon) \times (-\varepsilon, \varepsilon)$, with $Y_I \cap \nu(T^2) = T^2 \times (-\varepsilon, \varepsilon) \times \{0\}$ and $Y_J \cap \nu(T^2) = T^2 \times \{0\} \times (-\varepsilon, \varepsilon)$.

Now fix a path γ from 0 to 1. By Lemma 2.1.2, γ corresponds to a maximal subset of *disjoint* internal hypersurfaces $Y_{I_1} < Y_{I_2} < ... < Y_{I_{l-1}}$ in W. So for each point $(T_1, \ldots, T_{l-1}) \in [0, \infty)^{l-1}$, we may insert necks to express W as the Riemannian cobordism $W_\gamma(T_1, \ldots, T_{l-1})$ given by

$$W_{0 I_1} \bigcup_{Y_{I_1}} ([-T_1, T_1] \times Y_{I_1}) \bigcup_{Y_{I_1}} W_{I_1 I_2} \bigcup_{Y_{I_2}} \cdots \bigcup_{Y_{I_{l-1}}} ([-T_{l-1}, T_{l-1}] \times Y_{I_{l-1}}) \bigcup_{Y_{I_{l-1}}} W_{I_{l-1} 1}. \tag{2.1}$$

We then extend this family to the cube $[0, \infty]^{l-1}$ by degenerating the metric on Y_j when $T_j = \infty$. As in the proof of the composition law in Seiberg-Witten theory (see Section 1.1 below), when T_j grows, the Y_{I_j}-neck stretches, and when $T_j = \infty$, it breaks. In particular, $W_\gamma(0, \ldots, 0)$ has the metric g_0, while $W_\gamma(\infty, \ldots, \infty)$ is the disjoint union of l elementary cobordisms which compose to give W with the metric g_0.

In this way, we obtain $l!$ families of metrics on W, each parameterized by a cube C_γ. The facets of each cube fall evenly into two types. A facet is *interior* if it is specified by fixing a coordinate at 0, and *exterior* if it is specified by fixing a coordinate at ∞. Note that each almost-maximal chain $Y_{I_1} < \cdots < \widehat{Y_{I_j}} < \cdots < Y_{I_{l-1}}$ can be completed to a maximal chain in exactly two ways. It follows that each internal facet has a twin on another cube, in the sense that the twins parameterize identical families of metrics on W. By gluing the cubes together along twin facets, we can build a single family of metrics which interpolates between the various ways of expressing W as a composite cobordism. In fact, this construction realizes the cubical subdivision of the following ubiquitous convex polytope (see [50] for more background).

The permutohedron P_l of order l arises as the convex hull of all points in $\mathbb{R}^l$ whose coordinates are a permutation of $(1, 2, 3, \ldots, l)$. These points lie in general position in the hyperplane $x_1 + \cdots + x_l = \frac{l(l-1)}{2}$, so P_l has dimension $l - 1$. The first four permutohedra are the point, interval, hexagon, and truncated octahedron (see Figure 2.3). The 1-skeleton of P_l is the Cayley graph of the standard presentation of the symmetric group on l letters:

$$S_l = \langle \sigma_1, \cdots, \sigma_{l-1} \,|\, \sigma_i^2 = 1,\, \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1},\, \sigma_i \sigma_j = \sigma_j \sigma_i \text{ for } |i - j| > 1 \rangle.$$

More generally, the $(l - d)$-dimensional faces of P_l correspond to partitions of the set

$\{1, \ldots, l\}$ into an ordered d-tuple of subsets $(A_1, \ldots, A_d)$. Inclusion of faces corresponds to merging of neighboring A_j.

The connection between the permutohedron and the hypercube rests on a simple observation: the face poset of P_l is dual to the poset of chains of internal vertices in the hypercube $\{0,1\}^l$. Namely, to each face $(A_1, \ldots, A_d)$, we assign the chain $I_1 < \cdots < I_{d-1}$, where I_j has i^{th} coordinate 1 if and only if $i \in A_1 \cup \cdots \cup A_{j-1}$. For example, in the case of the edges of the hexagon P_3, the correspondence is given by:

$$\begin{array}{cccccc} (\{3\},\{1,2\}) & (\{2,3\},\{1\}) & (\{2\},\{1,3\}) & (\{1,2\},\{3\}) & (\{1\},\{2,3\}) & (\{1,3\},\{2\}) \\ 001 & 011 & 010 & 110 & 100 & 101 \end{array}$$

In particular, each path γ from *0* to *1* corresponds to a vertex V_γ of P_l.

Now in the cubical subdivision of P_l, we may identify the cube containing V_γ with the cube of metrics C_γ so that twin interior facets are identified (see Figures 2.2 and 2.3). In this way, the interior of P_l parameterizes a family of non-degenerate metrics on W, while the boundary parameterizes a family of degenerate metrics. The parameterization can be made smooth on the interior by a slight adjustment of the rate of stretching. We summarize these observations in the following proposition.

Proposition 2.1.4. *The face poset of the permutohedron P_l is dual to the poset of chains of internal hypersurfaces in W. In particular, the facets of P_l correspond to the ways of breaking W into a composite cobordism along a single internal hypersurface. The interior of P_l smoothly parameterizes a family of non-degenerate metrics on W, which extends naturally to the boundary in such a way that the interior of each face parameterizes those metrics which are degenerate on precisely the corresponding chain.*

Remark 2.1.5. We describe an alternative view of the above construction which is not essential, but will be helpful in Section 2.3 when we consider more general lattices than the hypercube. Recall that a directed graph Γ is *transitive* if the existence of edges from I to J and from J to K implies the existence of an edge from I to K. The *transitive closure* of Γ is the directed graph obtained from Γ by adding the fewest number of edges necessary to achieve transitivity. A *clique* in an undirected graph is a subset of nodes with the property that every two nodes in the subset is connected by an edge.

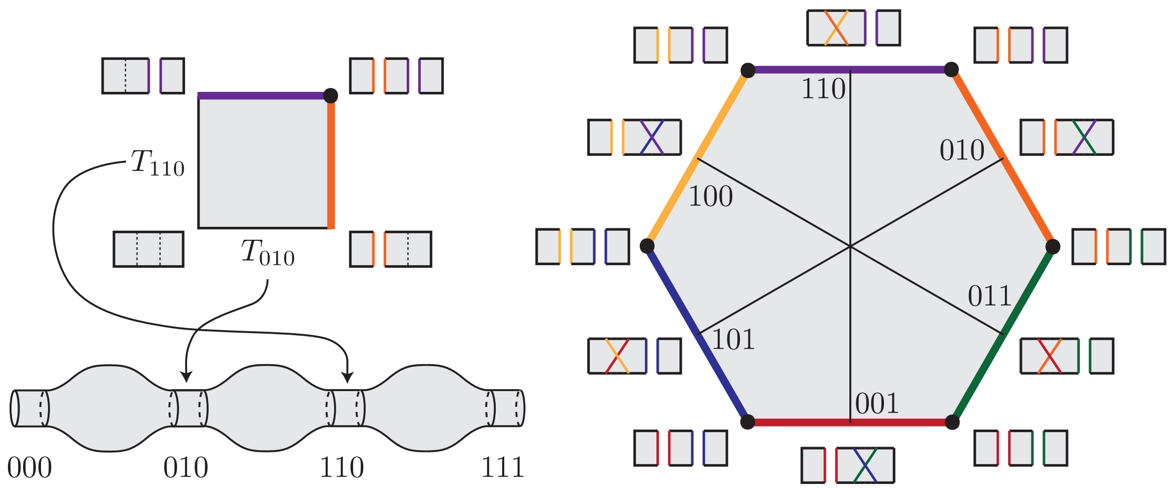

Figure 2.2: At left, we consider the path γ given by $000 < 010 < 110 < 111$ in $\{0,1\}^3$. The corresponding square C_γ with coordinates (T_{010}, T_{011}) parameterizes a family of metrics on the cobordism W^* which stretches at Y_{010} and Y_{110}. We have one square for each non-intersecting pair of hypersurfaces in Figure 2.1. These six squares fit together to form the hexagon P_3 at right. The small figures at the vertices and edges illustrate the metric degenerations on W, read as composite cobordisms from left to right.

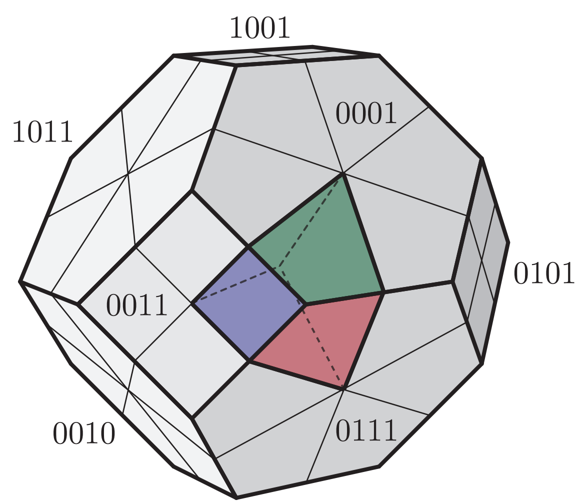

Figure 2.3: The cubical subdivision of the permutohedron P_4 consists of 24 cubes, corresponding to the 4! paths from 0000 to 1111 in $\{0,1\}^4$. Above, the cube corresponding to the path $0000 < 0001 < 0011 < 0111 < 1111$ is shown with its exterior faces in translucent color. Each cube shares one vertex with P_4 and has one vertex at the center.

Consider the directed graph Γ associated to $\{0,1\}^l$, with an edge from I to J whenever J is an immediate successor of I. Let $\bar{\Gamma}$ be the transitive closure of $\Gamma - \{\mathit{0}, \mathit{1}\}$. The nodes of $\bar{\Gamma}$ correspond to internal hypersurfaces, and by Lemma 2.1.2, two nodes are joined by an edge if and only if the corresponding internal hypersurfaces are disjoint. In fact, $\bar{\Gamma}$ is the 1-skeleton of a simplicial complex $\mathcal{C}_l$, whose face poset is isomorphic to the poset of non-empty cliques in $\bar{\Gamma}$ (as an undirected graph) under inclusion. The complex $\mathcal{C}_l$ is dual to the boundary of P_l.

2.2 The link surgery spectral sequence: construction

Let W be the cobordism associated to surgery on a framed link $L \subset Y$. In Section 2.1, we constructed a family of metrics on W, parameterized by a permutohedron P_l and degenerate on the boundary Q_l. We now use such families to define maps between summands in a hypercube complex X associated to the framed link. That these maps define a differential will follow from a generalization of (1.5) similar in spirit to (1.8). The link surgery spectral sequence is then induced by the filtration on the hypercube complex given by vertex weight.

Fix a metric and admissible perturbation on the cobordism W which are cylindrical near every hypersurface Y_I. Let X be the direct sum of the monopole Floer complexes of the hypersurfaces, considered as a vector space over $\mathbb{F}_2$:

$$X = \bigoplus_{I \in \{0,1\}^l} \check{C}(Y_I)$$

We will define a differential $\check{D} : X \to X$ as the sum of maps $\check{D}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ over all $I \leq J$, with $\check{D}^I_I$ the differential on the monopole Floer complex $\check{C}(Y_I)$. We now construct the maps $\check{D}^I_J$ when $I < J$.

Fix vertices $I < J$ and let $k = w(J) - w(I)$. Regarding W_{IJ} as the cobordism arising by surgery on a k-component, framed link in Y_I, with initial metric induced by W, we apply Proposition 2.1.4 to obtain a family of metrics on W_{IJ} parameterized by the permutohedron P_{IJ} of dimension $k-1$. Consider a pair of critical points $\mathfrak{a} \in \mathfrak{C}(Y_I)$ and $\mathfrak{b} \in \mathfrak{C}(Y_J)$, and a relative homotopy class z from $\mathfrak{a}$ to $\mathfrak{b}$ in the configuration space $\mathcal{B}^\sigma(W_{IJ})$. As in (1.7), we must extend the definition of $M_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$ to the degenerate metrics over the boundary of P_{IJ}. If p is in the interior of the face $I_1 < I_2 < \cdots < I_{q-1}$, then an element $\boldsymbol{\gamma}$ of $M_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$ is a q-tuple

$$(\gamma_{01}, \gamma_{12}, \ldots, \gamma_{q-1\,q})$$

where

$$\begin{aligned}\gamma_{j\,j+1} &\in M(\mathfrak{a}_j, W^*_{I_jI_{j+1}}(p), \mathfrak{a}_{j+1})\\ \mathfrak{a}_0 &= \mathfrak{a}\\ \mathfrak{a}_q &= \mathfrak{b}\end{aligned}$$

and the homotopy classes of these elements compose to give z. Here, the metric on $W_{I_jI_{j+1}}(p)$ is the restriction of the metric on $W(p)$. We then define $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ as the fiber product

$$M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = \bigcup_{p \in P} \{p\} \times M_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b}).$$

This space has a reducible analogue $M^{\text{red}}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ which is defined by replacing each moduli space of the form $M_z(\mathfrak{a}, W^*, \mathfrak{b})$ with its reducible locus $M^{\text{red}}_z(\mathfrak{a}, W^*, \mathfrak{b})$.

In order to count the points in these moduli spaces, we define two elements of $\mathbb{F}_2$ by

$$m_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b}) = \begin{cases} |M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}| \bmod 2, & \text{if } \dim M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = 0 \\ 0, & \text{otherwise,} \end{cases} \tag{2.2}$$

$$\bar{m}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b}) = \begin{cases} |M^{\text{red}}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}| \bmod 2, & \text{if } \dim M^{\text{red}}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = 0 \\ 0, & \text{otherwise.} \end{cases} \tag{2.3}$$

Remark 2.2.1. When $I = J$, we replace $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ in (2.2) by the moduli space $\breve{M}_z(\mathfrak{a}, \mathfrak{b})$ of unparameterized trajectories on the cylinder $\mathbb{R} \times Y$ (see the definition below). We similarly replace $M^{\text{red}}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ in (2.3) by $\breve{M}^{\text{red}}_z(\mathfrak{a}, \mathfrak{b})$.

Recall that $C^o(Y)$, $C^s(Y)$, and $C^u(Y)$ are vector spaces over $\mathbb{F}_2$, with bases $e_{\mathfrak{a}}$ indexed by the monopoles $\mathfrak{a}$ in $\mathfrak{C}^o(Y)$, $\mathfrak{C}^s(Y)$, and $\mathfrak{C}^u(Y)$, respectively. We use the above counts to construct eight linear maps $D^o_o\binom{I}{J}$, $D^o_s\binom{I}{J}$, $D^u_o\binom{I}{J}$, $D^u_s\binom{I}{J}$, $\bar{D}^s_s\binom{I}{J}$, $\bar{D}^s_u\binom{I}{J}$, $\bar{D}^u_s\binom{I}{J}$, $\bar{D}^u_u\binom{I}{J}$, where for example,

$$\begin{aligned} &D^u_s\binom{I}{J} : C^u_\bullet(Y_I) \to C^s_\bullet(Y_J) &\qquad D^u_s\binom{I}{J} e_{\mathfrak{a}} = \sum_{\mathfrak{b}\in\mathfrak{C}^u(Y_J)} \sum_z m_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b}) e_{\mathfrak{b}}\,; \\ &\bar{D}^u_s\binom{I}{J} : C^u_\bullet(Y_I) \to C^s_\bullet(Y_J) &\qquad \bar{D}^u_s\binom{I}{J} e_{\mathfrak{a}} = \sum_{\mathfrak{b}\in\mathfrak{C}^u(Y_J)} \sum_z \bar{m}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b}) e_{\mathfrak{b}}\,. \end{aligned} \tag{2.4}$$

Note that the above two maps are distinct. We then define $\check{D}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ by the matrix

$$\check{D}^I_J = \begin{bmatrix} D^o_o\binom{I}{J} & \sum_{I\le K\le J} D^u_o\binom{K}{J}\bar{D}^s_u\binom{I}{K} \\ D^o_s\binom{I}{J} & \bar{D}^s_s\binom{I}{J} + \sum_{I\le K\le J} D^u_s\binom{K}{J}\bar{D}^s_u\binom{I}{K} \end{bmatrix}, \tag{2.5}$$

with respect to the decomposition $\check{C}(Y) = C^o(Y) \bigoplus C^s(Y)$. The motivation behind this definition is explained in the Appendix. Finally, as promised, we let $\check{D} : X \to X$ be the sum

$$\check{D} = \sum_{I\le J} \check{D}^I_J.$$

We now turn to proving that $\check{D}$ is a differential. As in the proof of the composition law, the argument proceeds by constructing an appropriate compactification of $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ and counting boundary points. We first consider the compactification of the space of unparameterized trajectories on Y, repeating nearly verbatim the definitions given in Section 16.1 of [24]. A trajectory γ belonging to $M_z(\mathfrak{a}, \mathfrak{b})$ is *non-trivial* if it is not invariant under the action of $\mathbb{R}$ by translation on the cylinder $\mathbb{R} \times Y$. An *unparameterized trajectory* is an equivalence class of non-trivial trajectories in $M_z(\mathfrak{a}, \mathfrak{b})$. We write $\breve{M}_z(\mathfrak{a}, \mathfrak{b})$ for the space of unparameterized trajectories. An *unparameterized broken trajectory* joining $\mathfrak{a}$ to $\mathfrak{b}$ consists of the following data:

- an integer $n \ge 0$, the *number of components*;
- an $(n+1)$-tuple of critical points $\mathfrak{a}_0, \ldots, \mathfrak{a}_n$ with $\mathfrak{a}_0 = \mathfrak{a}$ and $\mathfrak{a}_n = \mathfrak{b}$, the *restpoints*;
- for each i with $1 \le i \le n$, an unparameterized trajectory $\breve{\gamma}_i$ in $\breve{M}_z(\mathfrak{a}_{i-1}, \mathfrak{a}_i)$, the ith component of the broken trajectory.

The *homotopy class* of the broken trajectory is the class of the path obtained by concatenating representatives of the classes z_i, or the constant path at $\mathfrak{a}$ if $n = 0$. We write $\breve{M}^+_z(\mathfrak{a}, \mathfrak{b})$

for the space of unparameterized broken trajectories in the homotopy class z, and denote a typical element by $\breve{\boldsymbol{\gamma}} = (\gamma_1, \ldots, \gamma_n)$. This space is compact for the appropriate topology (see [24], Section 24.6). Note that if z is the class of the constant path at $\mathfrak{a}$, then $\breve{M}_z(\mathfrak{a}, \mathfrak{a})$ is empty, while $\breve{M}^+_z(\mathfrak{a}, \mathfrak{a})$ is a single point, a broken trajectory with no components.

We are now ready to define the compactification $M^+_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$. If p is in the interior of the face $I_1 < I_2 < \cdots < I_{q-1}$, then an element $\breve{\boldsymbol{\gamma}}$ of $M^+_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$ is a $(2q+1)$-tuple

$$(\breve{\boldsymbol{\gamma}}_0, \gamma_{01}, \breve{\boldsymbol{\gamma}}_1, \gamma_{12}, \ldots, \breve{\boldsymbol{\gamma}}_{q-1}, \gamma_{q-1\, q}, \breve{\boldsymbol{\gamma}}_q)$$

where

$$\begin{aligned}
\breve{\boldsymbol{\gamma}}_j &\in \breve{M}^+(\mathfrak{a}_j, \mathfrak{a}_j) \\
\gamma_{j\, j+1} &\in M(\mathfrak{a}_j, W^*_{I_j I_{j+1}}(p), \mathfrak{a}_{j+1}) \\
\mathfrak{a}_0 &= \mathfrak{a} \\
\mathfrak{a}_q &= \mathfrak{b}
\end{aligned}$$

and $\breve{\boldsymbol{\gamma}}$ is in the homotopy class z. The fiber product

$$M^+_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = \bigcup_{p \in P} \{p\} \times M^+_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$$

is compact for the appropriate topology (see [24], Section 26.1). We also write $M^+_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{Q_{IJ}}$ for the restriction of $M^+_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ to the fibers over the boundary Q_{IJ} . We can similarly define a compactification $M^{\text{red}+}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ of $M^{\text{red}}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ by only considering reducible trajectories. Recall that an unbroken trajectory from $\mathfrak{a}$ to $\mathfrak{b}$ is *boundary-obstructed* if $\mathfrak{a}$ is boundary-stable and $\mathfrak{b}$ is boundary-unstable. Fix a regular choice of metric and perturbation.

Remark 2.2.2. The intuition behind the following classification of ends comes from the model case of Morse homology for manifolds with boundary. We encourage the interested reader to see the Appendix at this time.

Lemma 2.2.3. *If $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ is 0-dimensional, then it is compact. If $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ is 1-dimensional and contains irreducibles, then $M^+_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ is a compact, 1-dimensional space stratified by manifolds. The 1-dimensional stratum is the irreducible part of $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$,* *while the* 0*-dimensional stratum (the boundary) has an even number of points and consists of:*

(A) trajectories with two or three components. In the case of three components, the middle one is boundary-obstructed.

(B) the reducibles locus $M_z^{red}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ *in the case that the moduli space contains reducibles as well (which requires* $\mathfrak{a}$ *to be boundary-unstable and* $\mathfrak{b}$ *to be boundary-stable).*

If $M_z^{red}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ *is* 0*-dimensional, then it is compact. If* $M_z^{red}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ *is* 1*-dimensional, then* $M_z^{red+}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ *is a compact,* 1*-dimensional* C^0*-manifold with boundary. The boundary has an even number of points and consists of:*

(C) trajectories with exactly two components.

Proof. This is essentially Lemma 4.15 of [27], which in turn is a generalization of the gluing theorems in [24] leading up to the proof of the composition law (see Corollary 21.3.2, Theorem 24.7.2, and Propositions 24.6.10, 25.1.1, and 26.1.6). □

Remark 2.2.4. When $I = J$, Lemma 2.2.3 holds with $M_z(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$, $M_z^+(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$, $M_z^{\text{red}}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$, and $M_z^{\text{red}+}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ replaced by $\breve{M}_z(\mathfrak{a}, \mathfrak{b})$, $\breve{M}_z^+(\mathfrak{a}, \mathfrak{b})$, $\breve{M}_z^{\text{red}}(\mathfrak{a}, \mathfrak{b})$, and $\breve{M}_z^{\text{red}+}(\mathfrak{a}, \mathfrak{b})$, respectively.

We obtain a number of identities from the fact that these moduli spaces have an even number of boundary points. We now bundle these identities into a single operator $\check{A}_J^I$, constructed by analogy with $\check{D}_J^I$. Fix a pair of critical points $\mathfrak{a} \in \mathfrak{C}(Y_I)$ and $\mathfrak{b} \in \mathfrak{C}(Y_J)$, and a relative homotopy class z from $\mathfrak{a}$ to $\mathfrak{b}$ in the configuration space $\mathcal{B}^\sigma(W_{IJ})$. We define two elements of $\mathbb{F}_2$ by

$$n_z(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}} = \begin{cases} |\{\text{trajectories in (A) or (B)}\}| \bmod 2, & \text{if } \dim M_z(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}} = 1 \\ 0, & \text{otherwise,} \end{cases}$$

$$\bar{n}_z(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}} = \begin{cases} |\{\text{trajectories in (C)}\}| \bmod 2, & \text{if } \dim M_z^{\text{red}}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}} = 1 \\ 0, & \text{otherwise.} \end{cases}$$

Remark 2.2.5. When $I = J$, we again replace $M_z(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ and $M_z^{\text{red}}(\mathfrak{a}, W_{IJ}^*, \mathfrak{b})_{P_{IJ}}$ by $\breve{M}_z(\mathfrak{a}, \mathfrak{b})$ and $\breve{M}_z^{\text{red}}(\mathfrak{a}, \mathfrak{b})$, respectively.

Remark 2.2.6. Trajectories of type (A) necessarily have at least one irreducible component. It follows that if $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ is 1-dimensional and does not contain irreducibles, then it can only have boundary points in strata of type (C). So the condition "if dim $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = 1$" is equivalent to the usual condition "if dim $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = 1$ and $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ contains irreducibles." A similar remark holds for the definition of $m_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})$.

By Lemma 2.2.3 and the above remark, $n_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ counts the boundary points of $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ when it is 1-dimensional and contains irreducibles, and is zero otherwise. Similarly, $\bar{n}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ counts the boundary points of $M^{\mathrm{red}}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ when it is 1-dimensional, and is zero otherwise. Since the number of boundary points is even, we conclude:

$$n_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} \textit{ and } \bar{n}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} \textit{ vanish for all choices of } \mathfrak{a}, \mathfrak{b}, \textit{ and } z. \tag{2.6}$$

We proceed by analogy with $\check{D}^I_J$, using $n_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ to define linear maps $A^o_o(^I_J)$, $A^o_s(^I_J)$, $A^u_o(^I_J)$, and $A^u_s(^I_J)$, and $\bar{n}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ to define linear maps $\bar{A}^s_s(^I_J)$ and $\bar{A}^s_u(^I_J)$ (we will not need $\bar{A}^u_s(^I_J)$ or $\bar{A}^u_u(^I_J)$). Again, these maps all vanish identically by (2.6). Each of these maps can be expressed as a sum of terms which are themselves compositions of the component maps of $\check{D}^I_J$. Finally, we define the map $\check{A}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ by the matrix

$$\check{A}^I_J = \begin{bmatrix} A^o_o(^I_J) & \sum_{I \le K \le J} \left(A^u_o(^K_J) \bar{D}^s_u(^I_K) + D^u_o(^K_J) \bar{A}^s_u(^I_K) \right) \\ A^o_s(^I_J) & \bar{A}^s_s(^I_J) + \sum_{I \le K \le J} \left(A^u_s(^K_J) \bar{D}^s_u(^I_K) + D^u_s(^K_J) \bar{A}^s_u(^I_K) \right) \end{bmatrix}. \tag{2.7}$$

It follows that $\check{A}^I_J$ vanishes identically as well. The motivation behind the definition of $\check{A}^I_J$ is explained in the Appendix.

Lemma 2.2.7. *$\check{A}^I_J$ is equal to the component of $\check{D}^2$ from $\check{C}(Y_I)$ to $\check{C}(Y_J)$:*

$$\check{A}^I_J = \sum_{I \le K \le J} \check{D}^K_J \check{D}^I_K.$$

Proof. We must show that corresponding matrix entries are equal, that is

$$\begin{aligned}
A_o^o\tbinom{I}{J} &= \sum_{I\le K\le J} D_o^o\tbinom{K}{J} D_o^o\tbinom{I}{K} \\
&+ \sum_{I\le K\le M\le J} D_o^u\tbinom{M}{J}\bar{D}_u^s\tbinom{K}{M} D_s^o\tbinom{I}{K} \\
A_s^o\tbinom{I}{J} &= \sum_{I\le K\le J} D_s^o\tbinom{K}{J} D_o^o\tbinom{I}{K} \\
&+ \sum_{I\le K\le J} \bar{D}_s^s\tbinom{K}{J} D_s^o\tbinom{I}{K} \\
&+ \sum_{I\le K\le M\le J} D_s^u\tbinom{M}{J}\bar{D}_u^s\tbinom{K}{M} D_s^o\tbinom{I}{K} \\
\sum_{I\le K\le J}\left(A_o^u\tbinom{K}{J}\bar{D}_u^s\tbinom{I}{K} + D_o^u\tbinom{K}{J}\bar{A}_u^s\tbinom{I}{K}\right) &= \sum_{I\le L\le K\le J} D_o^o\tbinom{K}{J} D_o^u\tbinom{L}{K}\bar{D}_u^s\tbinom{I}{L} \\
&+ \sum_{I\le K\le M\le J} D_o^u\tbinom{M}{J}\bar{D}_u^s\tbinom{K}{M}\bar{D}_s^s\tbinom{I}{K} \\
&+ \sum_{I\le L\le K\le M\le J} D_o^u\tbinom{M}{J}\bar{D}_u^s\tbinom{K}{M} D_s^u\tbinom{L}{K}\bar{D}_u^s\tbinom{I}{L} \\
\bar{A}_s^s\tbinom{I}{J} + \sum_{I\le K\le J}\left(A_s^u\tbinom{K}{J}\bar{D}_u^s\tbinom{I}{K} + D_s^u\tbinom{K}{J}\bar{A}_u^s\tbinom{I}{K}\right) &= \sum_{I\le L\le K\le J} D_s^o\tbinom{K}{J} D_o^u\tbinom{L}{K}\bar{D}_u^s\tbinom{I}{L} \\
&+ \sum_{I\le K\le J} \bar{D}_s^s\tbinom{K}{J}\bar{D}_s^s\tbinom{I}{K} \\
&+ \sum_{I\le L\le K\le J} \bar{D}_s^s\tbinom{K}{J} D_s^u\tbinom{L}{K}\bar{D}_u^s\tbinom{I}{L} \\
&+ \sum_{I\le K\le M\le J} D_s^u\tbinom{M}{J}\bar{D}_u^s\tbinom{K}{M}\bar{D}_s^s\tbinom{I}{K} \\
&+ \sum_{I\le L\le K\le M\le J} D_s^u\tbinom{M}{J}\bar{D}_u^s\tbinom{K}{M} D_s^u\tbinom{L}{K}\bar{D}_u^s\tbinom{I}{L}.
\end{aligned}$$

After expanding out the A_*^* and distributing, all terms on the right appear exactly once on the left by Lemma 2.2.3 (the terms with four components appear only once since $\bar{D}_u^s D_s^u \bar{D}_u^s$ is not a term of A_u^s). All other terms on the left are of the form $D_o^u\bar{D}_u^u\bar{D}_u^s$, $D_s^u\bar{D}_u^u\bar{D}_u^s$, or $\bar{D}_s^u\bar{D}_u^s$. In the first case, $D_o^u\tbinom{K_2}{J}\bar{D}_u^u\tbinom{K_1}{K_2}\bar{D}_u^s\tbinom{I}{K_1}$ is a term of both $A_o^u\tbinom{K_1}{J}\bar{D}_u^s\tbinom{I}{K_1}$ and $D_o^u\tbinom{K_2}{J}\bar{A}_u^s\tbinom{I}{K_2}$. Similarly, $D_s^u\bar{D}_u^u\bar{D}_u^s$ occurs in $A_s^u\bar{D}_u^s$ and $D_s^u\bar{A}_u^s$, and $\bar{D}_s^u\bar{D}_u^s$ occurs in $A_s^u\bar{D}_u^s$ and A_s^s. Therefore, each of the extra terms occurs twice and we have equality over $\mathbb{F}_2$. $\square$

Remark 2.2.8. An internal restpoint of $\breve{\gamma}$ is called a *break*. A break is *good* if the corresponding monopole is irreducible or boundary-stable. A trajectory $\breve{\gamma} \in M_z^+(\mathfrak{a}_0, W^*, \mathfrak{b}_0)$ occurs in the *extended boundary* of a 1-dimensional stratum if $\breve{\gamma}$ can be obtained by appending (possibly zero) additional components to either end of a boundary point of a 1-dimensinal moduli space $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ or $M_z^{\text{red}}(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$. In these terms, we have shown that among the trajectories counted by $\check{A}^I_J$, those with no good break each occur in the extended boundary of exactly two 1-dimensional strata. The remaining trajectories each have one good break and occur in the extended boundary of exactly one 1-dimensional stratum. In particular, $\check{D}^K_J \check{D}^I_K$ counts those isolated trajectories which break well on Y_K. This remark may also be understood from the perspective of path algebras, as explained in the Appendix.

Remark 2.2.9. A break of $\breve{\gamma} = (\breve{\gamma}_0, \gamma_{01}, \ldots, \breve{\gamma}_q)$ is *central* if it is *not* a restpoint of $\breve{\gamma}_0$ or $\breve{\gamma}_q$. Note that $\breve{\gamma}$ has a central break if and only if it lies over a boundary fiber. We can express $\check{A}^I_J$ as the sum of similarly defined maps $\check{Q}^I_J$ and $\check{B}^I_J$, which count boundary points with and without a central, good break, respectively. It follows from Remark 2.2.8 that

$$\check{B}^I_J = \check{D}^I_J \check{D}^I_I + \check{D}^J_J \check{D}^I_J$$
$$\check{Q}^I_J = \sum_{I<K<J} \check{D}^K_J \check{D}^I_K \,.$$

$\check{B}^I_J$ may be thought of (imprecisely) as an operator associated to the interior of P_{IJ}, while $\check{Q}^I_J$ is (precisely) the operator associated to the boundary Q_{IJ} (in the case $l = 3$ in Figure 2.2, $\check{Q}^{000}_{111}$ is the sum of six composite operators, one for each edge of the hexagon). We can then express $\check{A}^I_J = 0$ as

$$\check{B}^I_J = \check{Q}^I_J,$$

which has the form

$$\check{D}^I_J \check{D}^I_I + \check{D}^J_J \check{D}^I_J = \check{Q}^I_J.$$

This is the sense in which Lemma 2.2.7 should be viewed as a generalization of (1.5). As in that case, $\check{Q}^I_J$ is null-homotopic and $\check{D}^I_J$ provides the chain homotopy.

We now conclude:

Proposition 2.2.10. *$(X, \check{D}, F)$ is a filtered chain complex, where F is the filtration induced by weight, namely*

$$F^i X = \bigoplus_{\substack{I \in \{0,1\}^l \\ w(I) \geq i}} \check{C}(Y_I).$$

Proof. The equation $\check{D}^2 = 0$ holds by Lemma 2.2.7 and the fact that the operators $\check{A}^I_J$ all vanish identically. The differential $\check{D}$ respects the filtration, as $I \leq J$ implies $w(I) \leq w(J)$. □

In order to describe $H_*(X, \check{D})$, we recall some topology. Let Y_0 be a closed, oriented 3-manifold, equipped with an oriented, framed knot K_0, and let Y_1 be the result of surgery on K_0 (this surgery is insensitive to the orientation of K_0). Y_1 comes equipped with a canonical oriented, framed knot K_1, obtained as the boundary of the cocore of the 2-handle in the associated elementary cobordism, and given the -1 framing with respect to the cocore (see Section 42.1 of [24] for details). So we may iterate this surgery process, yielding a sequence of pairs $\{(Y_n, K_n)\}_{n \geq 0}$. It is well-known that this sequence is 3-periodic, in the sense that for each $i \geq 0$, there is an orientation-preserving diffeomorphism

$$(Y_{i+3}, K_{i+3}) \xrightarrow{\cong} (Y_i, K_i)$$

which carries the oriented, framed knot K_{i+3} to K_i. Applying this construction to each component of the link $L \subset Y$, we may extend our collection of surgered 3-manifolds Y_I from the hypercube $\{0,1\}^l$ to the lattice $\{0, 1, \infty\}^l$. We may now state the 2-handle version of the link surgery spectral sequence, which computes $H_*(X, \check{D})$ in stages.

Theorem 2.2.11. *Let Y be a closed, oriented 3-manifold, equipped with an l-component framed link L. Then the filtered complex $(X, \check{D}, F)$ induces a spectral sequence with E^1-term given by*

$$E^1 = \bigoplus_{I \in \{0,1\}^l} \widetilde{HM}_\bullet(Y_I)$$

and d^1 differential given by

$$d^1 = \bigoplus_{w(J) - w(I) = 1} \widetilde{HM}_\bullet(W_{IJ}).$$

The link surgery spectral sequence collapses by stage $l+1$ to $\widetilde{HM}(Y_\infty)$. Each page has an integer grading $\check{t}$ induced by vertex weight, which the differential d^k increases by k.

Remark 2.2.12. The above statement uses different notation than that given in Theorem 2.0.1 in the introduction and in Theorem 4.1 of [36], emphasizing 2-handle addition over surgery. To reconcile the two forms, we describe the 3-periodicity above in the case of a knot $K_0 \subset Y$ from the surgery perspective (see Section 42.1 of [24]). The complements $Y - \nu(K_n)$ are all diffeomorphic, so we may view each of the surgered manifolds Y_n as obtained by gluing a solid torus to the fixed complement $Y_0 - \nu(K_0)$. If we denote the meridian and framing of K_n by μ_n and λ_n, respectively, thought of as curves on the torus $\partial\nu(K_1)$, then we have the relations

$$\mu_{n+1} = \lambda_n$$

$$\lambda_{n+1} = -\mu_n - \lambda_n$$

which correspond to the matrix

$$\begin{bmatrix} 0 & -1 \\ 1 & -1 \end{bmatrix}$$

of order 3. Since the framing is insensitive to the orientation of the curve, we can regard K_0, K_1, and $K_2 = K_\infty$ as having the framings λ_0, $\lambda_0 + \mu_0$, and μ_0, respectively. Therefore, Y_I is shifted one step from $Y(I)$, i.e. $Y_1 = Y(0)$, $Y_\infty = Y(1)$, and $Y_0 = Y(\infty)$. So Theorem 2.0.1 is simply Theorem 2.2.11 applied to $K_1 \subset Y_1$. In the case of a link, the same shift in the 3-periodic sequence occurs in each component.

The first claim of Theorem 2.2.11 follows immediately from the usual construction of the spectral sequence associated to a filtered complex. The $\check{t}$ grading is well-defined since each differential d^k is homogenous with respect to vertex weight. We complete the proof in two stages. First, in Section 2.3, we define a complex $(\widetilde{X}, \check{D})$, modeled on the lattice $\{0, 1, \infty\}^l$, in which $(X, \check{D})$ sits as a quotient complex. Then, in Section 2.4, we use the surgery exact triangle to conclude that $\widetilde{X}$ is null-homotopic. The identity of the E^∞ term quickly follows.

2.3 Product lattices and graph associahedra

Consider the lattice $\{0,1,\infty\}^l$, with the product order induced by the convention $0 < 1 < \infty$. An ∞ digit contributes two to the weight. We will sometimes also use ∞ to denote the final vertex $\{\infty\}^l$, with the meaning clear from context. Consider the full cobordism W from Y_0 to Y_∞, the result of attaching two rounds of l 2-handles:

$$W = \left([0,1] \times Y \cup \bigcup_{i=1}^{l} h_i\right) \cup \bigcup_{j=1}^{l} g_j.$$

Here h_i is attached to the component K_i of $L \subset Y$, and g_j is attached to $K'_j \subset Y_1$, where K'_j denotes the boundary of the co-core of h_j with -1 framing. A valid order of attachment corresponds to a maximal chain in $\{0,1,\infty\}^l$, or equivalently to a path in Γ from 0 to ∞, of which there are $\frac{(2l)!}{2^l}$. For each vertex $I = (m_1, \ldots, m_l)$, we have the hypersurface Y_I, diffeomorphic to a boundary component of

$$W_{0I} = [0,1] \times Y \cup \bigcup_{\{i \,|\, m_i \geq 1\}} h_i \cup \bigcup_{\{i \,|\, m_i = \infty\}} g_i.$$

An ∞ digit corresponds to attaching a stack of two 2-handles to a component of $L \subset Y$.

As in Section 2.1, we will construct a polytope of metrics P_{IJ} on the cobordism W_{IJ} for all pairs of vertices $I < J$. The simplest new case occurs when $l = 1$, $I = 0$, and $J = \infty$. Since $w(\infty) - w(0) = 2$, the polytope $P_{0\infty}$ should be a closed interval with degenerate metrics over the two boundary points. However, we now have only one internal hypersurface, Y_1, on which to degenerate the metric. The solution, as in [27], is to construct an auxiliary hypersurface S_1 as follows. Let E_1 be the 2-sphere formed by gluing the cocore of h_1 to the core of g_1 along their common boundary K'_1. Due to the -1-framing on K'_1, $\nu(E_1)$ is a D^2-bundle of Euler class -1, with E_1 embedded as the zero-section. It follows that

$$\nu(E_1) \cong \overline{\mathbb{CP}}^2 - \mathrm{int}(D^4)$$

and we define the hypersurface S_1 to be the bounding 3-sphere $\partial\nu(K_i)$. $P_{0\infty}$ is then identified with the interval $[-\infty, \infty]$, with the metric degenerating on S_1 at $-\infty$ and Y_1 at ∞.

For the lattice $\{0,1,\infty\}^l$, we will embed l auxiliary 3-spheres $S_1, \ldots, S_l$ in addition to the $3^l - 2$ internal hypersurfaces. We must then construct a family of metrics which interpolates between the $\sum_{i=1}^{l} \binom{l}{i} \frac{(2l-i)!}{2^{l-i}}$ ways to decompose W along $2l-1$ pairwise-disjoint hypersurfaces. As a first step, we generalize Lemma 2.1.2. The following proposition is motivated by a half-dimensional diagram in the spirit of Figures 2.1 and Figure 2.8.

Proposition 2.3.1. *The full set of $3^l - 2$ internal hypersurfaces Y_I and l spheres S_i can be simultaneously embedded in the interior of W so that the following conditions hold:*

(i) The internal hypersurfaces in any subset are pairwise disjoint as submanifolds of W if and only if they form a chain. In this case, cutting along $Y_{I_1} < Y_{I_2} < ... < Y_{I_k}$ breaks W into the disjoint union

$$W_{0I_1} \coprod W_{I_1I_2} \coprod \cdots \coprod W_{I_k\infty}.$$

(ii) Distinct Y_I and Y_J intersect in exactly $\rho(I,J)$ disjoint tori.

(iii) Y_I and S_i intersect if and only if $m_i = 1$, where $I = (m_1, \ldots, m_l)$. In this case, they intersect in a torus.

(iv) The S_i are pairwise disjoint.

Proof. List the vertices as $I_0, I_1, ..., I_{3^l}$, first in order of increasing weight and then numerically within each weight class. We express the full cobordism as

$$W = [0, 3^l] \times Y \cup \bigcup_{i=1}^{l} h_i \cup \bigcup_{i=1}^{l} g_i$$

and embed Y_0 and Y_∞ as the boundary. As in the proof of Lemma 2.1.2, for each $1 \leq q \leq 3^l - 1$, we have slimmer 2-handles h_i^q and g_i^q as the images of $D^2 \times D_q^2$ in h_i and g_i, respectively, where D_q^2 is the disk of radius $\frac{q}{3^n}$. Again, we think of

$$\tilde{h}_i^q = [q, 3^l] \times \nu_q(K_i) \cup \bigcup_{\{3^l\} \times \nu_q(K_i)} h_i^q$$

as a longer 2-handle which tunnels through $[q, 3^l] \times Y$ in order to attach to $[0, q] \times Y$ along $\{q\} \times \nu_q(K_i)$. Let K_i' be the boundary of the cocore of h_i, so that $\nu_q(K_i') = D_q^2 \times \partial D^2$ is the region of h_i to which g_i^q attaches. Let A_q^i be the annulus given in polar coordinates (r, θ) by $[\frac{q}{3^l}, 1] \times S^1$, thought of as sitting in the cocore of h_i. The boundary of $D_q^2 \times A_i^q$ consists of $\nu_q(K_i')$ and a radial contraction of $\nu_q(K_i')$ into the interior of h_i, denoted $\tilde{\nu}_q(K_i')$. So we may regard

$$\tilde{g}_i^q = D_q^2 \times A_i^q \cup \bigcup_{\nu_q(K_i')} g_i^q$$

as a longer 2-handle which tunnels through $D^2 \times A_q^i \subset h_i$ in order to attach to $\tilde{h}_i^q$ along $\tilde{\nu}_q(K_i') \subset \partial \tilde{h}_i^q$. In this way, we embed W_{0I_q} in W as

$$W_{0I_q} = Y \times [0,q] \cup \bigcup_{\{i \,|\, m_i \geq 1\}} \tilde{h}_i^q \cup \bigcup_{\{i \,|\, m_i = \infty\}} \tilde{g}_i^q$$

and Y_{I_q} as a component of its boundary. Here $I_q = (m_1, ..., m_l)$. Next, let the 2-sphere E_i be the result of gluing the cocore of h_i and the core of g_i along their common boundary K_i', and let $\nu(E_i)$ be the result of gluing together the corresponding trivial D^2-bundles of radius $\frac{1}{2 \cdot 3^l}$. Then $\nu(E_i)$ is a D^2-bundle of Euler class -1, and we embed the 3-sphere S_i as its boundary.

Conditions (i) and (ii) now follow from a straightforward generalization of the proof of Lemma 2.1.2. For (iii), note that if $m_i = 1$, then the intersection of Y_I and S_i is the boundary of the restriction of the D^2-bundle $\nu(E_i)$ to K_i'. Finally, the S_i are pairwise disjoint because they live in different pairs of handles. □

For fixed $I < J$, the interval $\{K \mid I \leq K \leq J\}$ takes the form $\{0, 1, \infty\}^m \times \{0, 1\}^k$ for some pair of non-negative integers (m, k) with $m + k = l$. In order to define the maps $\check{D}_J^I$ in general, we need to construct a polytope $P_{m,k}$ of dimension $2m + k - 1$ for each pair (m, k). We define $P_{m,k}$ abstractly to have a face of co-dimension d for every subset of d mutually disjoint hypersurfaces in the interior of W, with inclusion of faces dual to inclusion of subsets. Our definition is justified by Theorem 2.3.3, which realizes $P_{m,k}$ concretely as a convex polytope.

In order to motivate this theorem, we first construct those $P_{m,k}$ of dimension three or less by hand. The polytopes $P_{0,1}$, $P_{0,2}$, $P_{0,3}$, and $P_{0,4}$ are the first few permutohedra of Proposition 2.1.4, namely a point, an interval, a hexagon, and a truncated octahedron (recall Figure 2.3). We saw that $P_{1,0}$ is an interval, and it is easy to see that $P_{1,1}$ is the associahedron K_4, otherwise known as the pentagon. $P_{2,0}$ is more interesting. In Figure 2.4, we use a trick to establish that it is K_5, also known as Stasheff's polytope [42]. See

[13] for an enjoyable, informal introduction to associahedra. Note that K_n has dimension $n-2$ while P_n has dimension $n-1$.

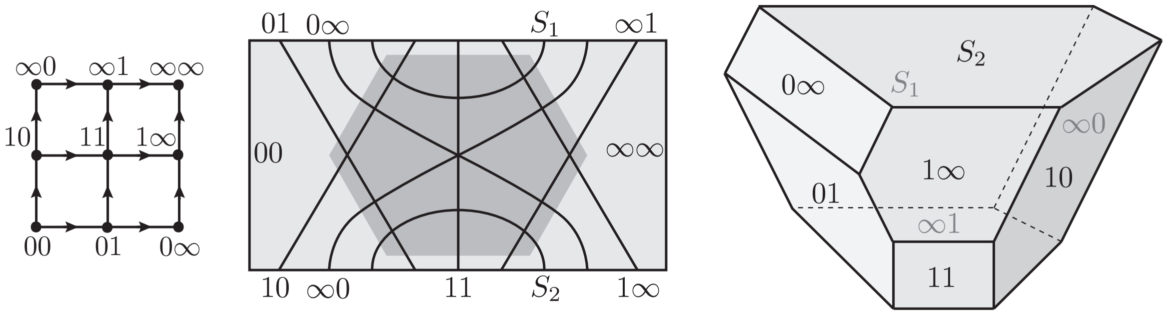

Figure 2.4: Consider the full cobordism W corresponding to the lattice $\{0,1,\infty\}^2$ at left. The seven interior hypersurfaces and two auxiliary 3-spheres are embedded in W in such a way that the diagram at center accurately depicts which pairs intersect (although the triple intersection point is an artifact). The nine internal arcs in the diagram are arranged so that by stretching normal to disjoint subsets, we obtain a parameterization of the space of conformal structures on the hexagon, which is known to compactify to the associahedron K_5 at right. In fact, we can exploit this connection between associahedra and conformal structures on polygons to construct monopole Floer analogues of maps counting psuedoholomorphic polygons in Heegaard Floer homology. We will return to this in a future paper.

At this stage, it may be tempting to conjecture that all the $P_{m,k}$ are permutohedra or associahedra. We check this against the only remaining 3-dimensional case, namely $P_{1,2}$. To build this polyhedron, it is useful to return to the viewpoint of Remark 2.1.5. Let Γ be the oriented graph corresponding to the lattice $\{0,1,\infty\}^m \times \{0,1\}^k$. Let $\bar{\Gamma}$ be the unoriented graph obtained as the transitive closure of Γ with its initial and final nodes removed. We now add l additional nodes I_i' (representing the S_i) to $\bar{\Gamma}$ and connect each I_i' to the others and to those $I = (m_1, \ldots, m_l) \in \bar{\Gamma}$ with $m_i \neq 1$. By Proposition 2.3.1, the nodes of the resulting graph $\bar{\Gamma}'$ are in bijection with the full set of hypersurfaces, with two nodes connected by an edge if and only if the corresponding hypersurfaces are disjoint. The graph $\bar{\Gamma}'$ is the 1-skeleton of a simplical complex $\mathcal{C}_{m,k}$ whose face poset is isomorphic to the poset of non-empty cliques in $\bar{\Gamma}'$ under inclusion. That is, the d-dimensional faces of $\mathcal{C}_{m,k}$ are in bijection with the d-cliques of $\bar{\Gamma}'$ (the fact that this poset defines a simplicial complex will follow from Theorem 2.3.3). The simple polytope dual to $\mathcal{C}_{m,k}$ is then, by definition, the boundary of $P_{m,k}$. In Figure 2.5, we illustrate this process for $P_{1,2}$, concluding that it is indeed something new.

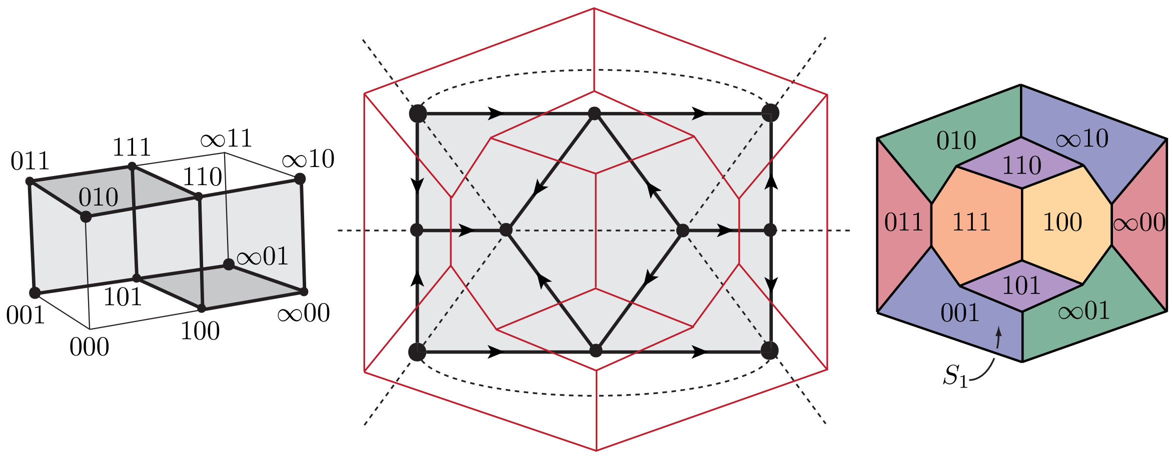

Figure 2.5: We construct the boundary of the polyhedron $P_{1,2}$ as the dual of the simplicial complex $\mathcal{C}_{1,2}$. First, at left, we remove the initial and final nodes from the lattice $\{0,1,\infty\}\times\{0,1\}^2$. We then flatten the shaded region and take the transitive closure to obtain $\bar{\Gamma}$, represented by the shaded rectangle and compact dotted line segments at center. Next we add the vertex I_1' at infinity (not shown) and connect it by dotted lines to the six nodes for which $m_1 \neq 1$. At this stage, we have constructed $\bar{\Gamma}'$, the 1-skeleton of $\mathcal{C}_{1,2}$. The faces of $\mathcal{C}_{1,2}$ are the 3-cliques (triangles). Drawing the dual with thin red lines, we obtain the boundary of $P_{1,2}$. At right, we have redrawn $P_{1,2}$. The face S_1 corresponds to the large hexagonal base under the colorful tortoise shell. The 12 vertices away from S_1 correspond to the 12 paths through the lattice.

The right hand side of Figure 2.5 illustrates $P_{1,2}$ as a convex polytope in $\mathbb{R}^3$. However, our dual-graph perspective does not provide such an explicit realization of $P_{m,k}$ in higher dimensions. While searching for an alternative construction of $P_{1,2}$, the author discovered beautiful illustrations of similar polyhedra in [12] and [15]. Given a connected graph G with n vertices, Carr and Devadoss construct a convex polytope P_G of dimension $n-1$, the *graph-associahedron* of G, using the following notions.

A *tube* of G is a proper, non-empty set of nodes of G whose induced graph is a connected subgraph of G. There are three ways in which tubes t_1 and t_2 can interact:

(1) Tubes are *nested* if $t_1 \subset t_2$ or $t_2 \subset t_1$;

(2) Tubes *intersect* if $t_1 \cap t_2 \neq \emptyset$ and $t_1 \not\subset t_2$ and $t_2 \not\subset t_1$;

(3) Tubes are *adjacent* if $t_1 \cap t_2 = \emptyset$ and $t_1 \cup t_2$ is a tube in G.

Tubes are *compatible* if they do not intersect and they are not adjacent. A *tubing* T of G is a set of tubes of G such that every pair of tubes in T is compatible.

We now define the graph-associahedron of a connected graph G with n nodes. Labelling each facet of the $n-1$ simplex $\triangle_G$ by a node of G, we have a bijection between the faces of $\triangle_G$ and the proper subsets of nodes of G. By definition, P_G is sculpted from $\triangle_G$ by truncating those faces which correspond to a connected, induced subgraph of G (see Figure 2.6). We therefore have a bijection

$$\{\text{facets of } P_G\} \longleftrightarrow \{\text{tubes of } G\}. \tag{2.8}$$

More generally, Carr and Devadoss prove that P_G is a simple, convex polytope whose face poset is isomorphic to the set of valid tubings of G, ordered such that $T < T'$ if T is obtained from T' by adding tubes. Moreover, in [15], Devadoss derives a simple, recursive formula for a set of points with integral coordinates in $\mathbb{R}^n$, whose convex hull realizes P_G.

Remark 2.3.2. Carr and Devadoss trace their construction back to the Deligne-Knudsen-Mumford compactification $\overline{\mathcal{M}}_{0,n}(\mathbb{R})$ of the real moduli space of curves. In this context, the sculpting of P_G is thought of as a sequence of real blow-ups. When G is a Coxeter graph, P_G tiles the compactification of the hyperplane arrangement associated to the corresponding Coxeter system. The n-clique, path, and cycle yield the $(n-1)$-dimensional permutohedron, associahedron, and cyclohedron, respectively. By the n-clique, we mean the complete graph on n nodes.

Comparing Figures 2.5 and 2.6, we see that $P_{1,2}$ is precisely the graph-associahedron of the 3-clique with one leaf. In fact, all of the polytopes $P_{m,k}$ are graph-associahedra:

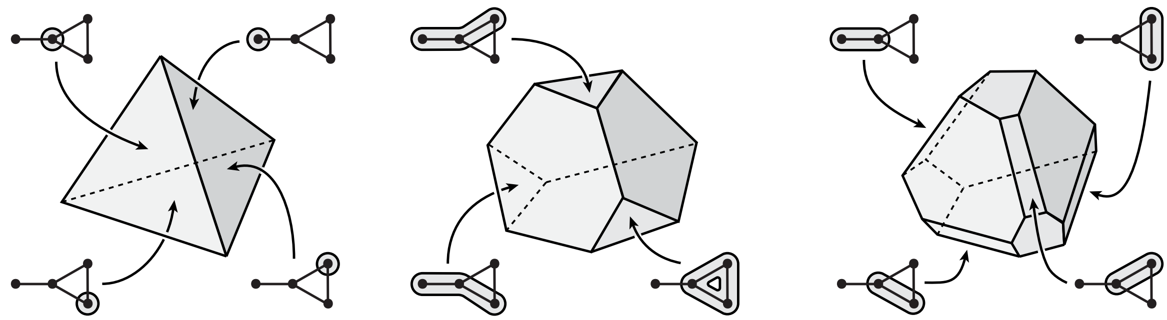

Figure 2.6: We have modified Figure 6 in [15] to illustrate the sculpting of P_G for the graph G given by the 3-clique with one leaf. Each node of G slices out a half-space in $\mathbb{R}^3$, leaving the 3-simplex $\triangle_G$ at left. Next, we shave down those vertices of $\triangle_G$ which correspond to the connected, induced subgraphs of size three. Finally, at right, we shave down those edges of $\triangle_G$ which correspond to the edges of G. This figure also illustrates the bijection (2.8).

Theorem 2.3.3. *The polytope $P_{m,k}$ associated to the lattice $\{0,1,\infty\}^m \times \{0,1\}^k$ is the graph-associahedron of the $(m+k)$-clique with m leaves. More generally, the polytope naturally associated to the lattice $\{0,...,n_1\} \times \cdots \times \{0,...,n_l\}$, with all $n_i \geq 1$, is the graph-associahedron of the l-clique with paths of length $n_1 - 1, \ldots, n_l - 1$ attached.*

Proof. An example is given in Figure 2.7. We first consider the lattice $\{0,1,\infty\}^m \times \{0,1\}^k$. In addition to the $3^m 2^k - 2$ internal hypersurfaces Y_I, we have m auxialliary hypersurfaces S_i. Let G be the complete graph on nodes $v_1, \ldots, v_{m+k}$ with a leaf v_i' attached to v_i for each $i = 1, \ldots, m$. The bijection (2.8) is given by

$$Y_I \mapsto \{v_i \,|\, m_i \geq 1\} \;\cup\; \{v_i' \,|\, m_i = \infty\}$$
$$S_i \mapsto \{v_i'\}.$$

and extends to an isomorphism of posets.

Next, consider the lattice $\{0,1,...,n\}$. The cobordism W is then built by attaching a single stack of handles $h_1 \cup \cdots \cup h_n$ to $[0,1] \times Y$ (the $n = 3$ case is shown in Figure 2.8, though with different notation). In addition to the internal hypersurfaces $Y_1, \ldots, Y_{n-1}$, we include an auxiliary hypersurface S_k^j between each pair of handles (h_j, h_k) with $1 \leq j < k \leq n$, embedded as the boundary of a tubular neighborhood of the union of the intervening 2-spheres E_i. In fact, if $k - j \equiv 2 \pmod 3$, then S_k^j is diffeomorphic to $S^1 \times S^2$. Otherwise,

S_k^j is diffeomorphic to S^3. By a straightforward variation on the theme of Lemma 2.1.2 and Proposition 2.3.1, these $n-1+\binom{n}{2}$ hypersurfaces can all be embedded in W so that

(i) the Y_i are all disjoint;

(ii) Y_i and S_k^j intersect if and only if $j \leq i < k$. In this case, they intersect in a torus.

(iii) $S_{k_1}^{j_1}$ and $S_{k_2}^{j_2}$ intersect if and only if the intervals $\{j_1, \ldots, k_1\}$ and $\{j_2, \ldots, k_2\}$ overlap but are not nested. In this case, they intersect in a torus.

Now let the graph G be the path with nodes $\{v_0, \ldots, v_n\}$. The bijection (2.8) is given by

$$Y_i \mapsto \{v_0, ..., v_i\}$$
$$S_k^j \mapsto \{v_j, ..., v_k\}.$$

and extends to an isomorphism of posets. As remarked above, P_G is then the $(n-1)$-dimensional associahedron K_{n+1}. The result for a lattice consisting of an arbitrary product of chains follows from a straightforward, subscript-heavy amalgamation of the arguments in the above two cases. □

Now consider the lattice $\Lambda = \{0, ..., n_1\} \times \cdots \times \{0, ..., n_l\}$ with the corresponding graph G given by Theorem 2.3.3. Using a formula in [15], we can realize P_G concretely as the convex hull of vertices in general position in $\mathbb{R}^d$, where $d = n_1 \cdots n_l - 1$. Now P_G has one vertex V_γ for every maximal collection γ of disjoint hypersurfaces in the cobordism W with initial metric g_0. As in Section 2.1, we associate to the vertex V_γ a cube of metrics C_γ which stretches on the hypersurfaces in γ. We can then use P_Λ to parameterize a family of metrics on W by identifying each C_γ with the cube containing the vertex V_γ in the cubical subdivision of P_Λ. In particular, $P_{m,k}$ consists of $\sum_{i=0}^{m} \binom{m}{i} \frac{(2m+k-i)!}{2^{m-i}}$ cubes.

Remark 2.3.4. Using these polytopes of metrics, we can define maps $\check{D}_J^I$ associated to any lattice formed as a product of chains of arbitrary length, where $\{0, \ldots, n\}$ has length n. However, we will see that this gives rise to a differential if and only if all the chains have length one or two. When there is a chain of length three or more, additional terms arise from breaks on auxiliary hypersurfaces. We will see this phenomenon explicitly for a single chain of length three in the proof of the surgery exact triangle (see Theorem 2.4.2).

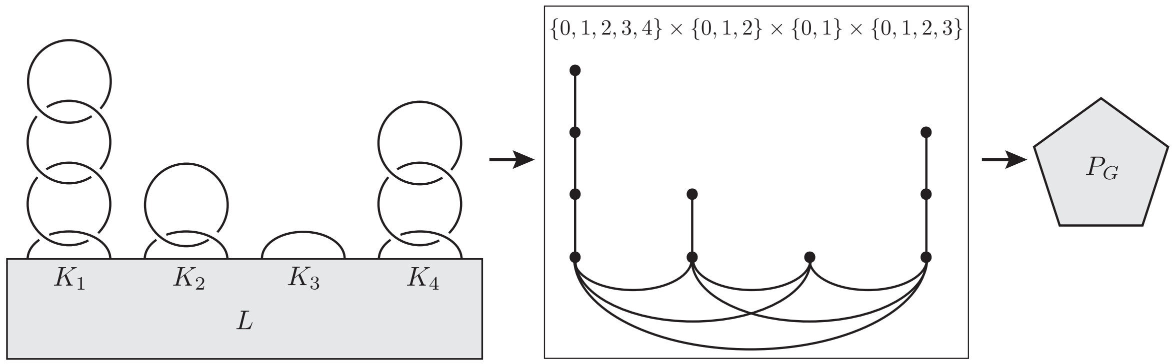

Figure 2.7: The figure at left represents a Kirby diagram arising from the 3-periodic surgery sequence applied to each component of a framed link with four components. The corresponding lattice is the product of four chains, while the graph is obtained by appending paths to the complete graph on four vertices. The pentagon at right represents the corresponding 9-dimensional graph associahedron. The above assignment of a polytope P_G to a finite product lattice generalizes the assignment of the permutohedron to the hypercube described in Section 2.1.

Having constructed polytopes of metrics for all intervals in the lattice $\{0,1,\infty\}^l$, we proceed to define the complex $(\widetilde{X}, \check{D})$. Fix a metric on the cobordism W which is cylindrical near every hypersurface Y_I and auxiliary hypersurface S_i, where each S_i has been equipped with the round metric. We let

$$\widetilde{X} = \bigoplus_{I \in \{0,1,\infty\}^l} \check{C}(Y_I)$$

and define the maps $\check{D}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ by *exactly* the same construction and matrix (2.5) as before, with $\check{D} : \widetilde{X} \to \widetilde{X}$ their sum.

We now prove that $\check{D}$ is a differential by an argument which parallels that in Section 2.2. We first expand our definition of $M^+_z(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$ to intervals of the form $\{0,1,\infty\}^m \times \{0,1\}^k$. Let V_i denote the copy of $\overline{\mathbb{CP}}^2 - \mathrm{int}(D^4)$ cut out by S_i. For each $I < J$, let $S_{j_1}, \dots, S_{j_{n(I,J)}}$ be the spheres completely contained in W_{IJ}. We denote the corresponding cobordism with $n(I,J)+2$ boundary components by

$$U_{IJ} = W_{IJ} - \bigcup_{s=1}^{n(I,J)} \mathrm{int}\left(V_{j_s}\right).$$

If p is in the interior of the face $(I_1 < I_2 < \cdots < I_{q-1}, S_1, ..., S_r)$, then an element $\breve{\boldsymbol{\gamma}}$ of $M_z^+(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$ is a $(2q+2r+1)$-tuple

$$(\breve{\gamma}_0, \gamma_{01}, \breve{\gamma}_1, \gamma_{12}, \ldots, \breve{\gamma}_{q-1}, \gamma_{q-1\,q}, \breve{\gamma}_q, \eta_1, \breve{\boldsymbol{\delta}}_1, \ldots, \eta_r, \breve{\boldsymbol{\delta}}_r) \tag{2.9}$$

where

$$\begin{aligned}
\breve{\gamma}_j &\in \breve{M}^+(\mathfrak{a}_j, \mathfrak{a}_j) \\
\gamma_{j\,j+1} &\in M(\mathfrak{a}_j, \mathfrak{c}_{j_1}, \ldots, \mathfrak{c}_{j_{n(I_j, I_{j+1})}}, U^*_{I_j I_{j+1}}(p), \mathfrak{a}_{j+1}) \\
\breve{\boldsymbol{\delta}}_i &\in \breve{M}^+(\mathfrak{c}_i, \mathfrak{c}_i) \\
\eta_i &\in M(V_i^*(p), \mathfrak{c}_i) \\
\mathfrak{a}_0 &= \mathfrak{a} \\
\mathfrak{a}_q &= \mathfrak{b} \\
\mathfrak{a}_j &\in \mathfrak{C}(Y_{I_j}) \\
\mathfrak{c}_i &\in \mathfrak{C}(S_i)
\end{aligned}$$

and $\breve{\boldsymbol{\gamma}}$ is in the homotopy class z (and similarly when p is in the interior of a face which includes a subset of the S_i other than the first r). The fiber product

$$M_z^+(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}} = \bigcup_{p \in P} \{p\} \times M_z^+(\mathfrak{a}, W_{IJ}(p)^*, \mathfrak{b})$$

is compact. We then define $\check{A}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ by *exactly* the same construction and matrix (2.7) as before.

We will also need the following lemma, consolidated from [27] (see Lemma 5.3 there and the preceding discussion). The essential point is that there is a diffeomorphism of $\overline{\mathbb{CP}}^2 - \mathrm{int}(D^4)$ which restricts to the identity on the boundary and induces a fixed-point-free involution on the set of spin^c structures.

Lemma 2.3.5. *Fix a sufficiently small perturbation on S_i. Then for each $\mathfrak{c}, \mathfrak{c}' \in \mathfrak{C}(S_i)$, $\breve{M}(\mathfrak{c}, \mathfrak{c}') = \emptyset$ and the trajectories in the zero-dimensional strata of $M^+(V_i^*, \mathfrak{c})$ occur in pairs.*

When $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ or $M_z^{\mathrm{red}}(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ is 1-dimensional, the number of boundary points in the corresponding compactification is still even (technically, using a generalization of Lemma 2.2.3 to the case of cobordisms with three boundary components, as done in [24] by introducing doubly boundary obstructed trajectories). By Lemma 2.3.5, the number of boundary points which break on precisely some non-empty, fixed collection $\{S_{i_1}, ...S_{i_r}\}$ of the auxiliary hypersurfaces is a multiple of 2^r, via the pairing of η_{i_j} and η'_{i_j} in (2.9). Therefore, by inclusion-exclusion, the number of boundary points which do not break on any of the S_i is even as well. Since these are precisely the boundary points counted by the matrix (2.7), $\check{A}^I_J$ still vanishes and the proof of Lemma 2.2.7 goes through without change. We conclude:

Proposition 2.3.6. *$(\widetilde{X}, \check{D}, F)$ is a filtered chain complex, where F is the filtration induced by weight, namely*

$$F^i \widetilde{X} = \bigoplus_{\substack{I \in \{0,1,\infty\}^l \\ w(I) \geq i}} \check{C}(Y_I).$$

Remark 2.3.7. While we were compelled to introduce auxiliary hypersurfaces S_i in order to obtain polytopes, the corresponding facets contribute vanishing terms to $\check{Q}^I_J$ by Lemma 2.3.5. We thereby recover

$$\check{Q}^I_J = \sum_{I<K<J} \check{D}^K_J \check{D}^I_K.$$

2.4 The surgery exact triangle

We will identify the E^∞ page of the link surgery spectral sequence by applying the surgery exact triangle to the complex of Proposition 2.3.6. Before stating the surgery exact triangle, we first recall the algebraic framework underlying its derivation in both monopole and Heegard Floer homology (see [27] and [36], respectively).

Lemma 2.4.1. *Let $\{A_i\}_{i=0}^\infty$ be a collection of chain complexes and let*

$$\{f_i : A_i \to A_{i+1}\}_{i=0}^\infty$$

be a collection of chain maps satisfying the following two properties:

(i) $f_{i+1} \circ f_i$ is chain homotopically trivial, by a chain homotopy

$$H_i : A_i \to A_{i+2}$$

(ii) the map

$$\psi_i = f_{i+1} \circ H_i + H_{i+1} \circ f_i : A_i \to A_{i+3}$$

is a quasi-isomorphism.

Then the induced sequence on homology is exact. Furthermore, the mapping cone of f_1 is quasi-isomorphic to A_3 via the map with components H_1 and f_2.

Let Y_0 be a closed, oriented 3-manifold, equipped with a framed knot K_0. Applying the functor $\widetilde{HM}_\bullet$ to the associated 3-periodic sequence of elementary cobordisms

$$\{W_n : Y_n \to Y_{n+1}\}_{n \in \mathbb{Z}/3\mathbb{Z}},$$

we obtain the surgery exact triangle:

Theorem 2.4.2. *With coefficients in $\mathbb{F}_2$, the sequence*

$$\cdots \longrightarrow \widetilde{HM}_\bullet(Y_{n-1}) \xrightarrow{\widetilde{HM}_\bullet(W_{n-1})} \widetilde{HM}_\bullet(Y_n) \xrightarrow{\widetilde{HM}_\bullet(W_n)} \widetilde{HM}_\bullet(Y_{n+1}) \longrightarrow \cdots$$

is exact.

Proof. We reorganize the proof in [27] to fit it within our general framework of polytopes P_{IJ} and identities $\check{A}^I_J$. We use the notation $\{0, 1, \infty, 0'\}$ for the lattice $\{1, 2, 3, 4\}$ considered in [27]. The corresponding graph (in the sense of both Γ and G) is the path of length three, yielding a pentagon of metrics P_G whose sides correspond to Y_1, Y_∞, $S_1 = S^1_\infty$, $S_\infty = S^\infty_{0'}$, and $R_1 = S^1_{0'}$ (where the left-hand notation is shorthand for the right-hand notation in the proof of Theorem 2.3.3). The auxiliary hypersurface R_1 is diffeomorphic to $S^1 \times S^2$ and cuts out $V_1 \cong \overline{\mathbb{CP}}^2 - \mathrm{int}(D^4)$ from W, leaving the cobordism U_1 with three boundary components.

Keeping the 3-periodicity in mind, we prove exactness by applying Lemma 2.4.1 with

$$\begin{array}{llll} A_{1+3j} = \check{C}(Y_0) & f_1 = \check{D}^0_1 & H_1 = \check{D}^0_\infty & \psi_1 = \check{D}^0_{0'} \\ A_{2+3j} = \check{C}(Y_1) & f_2 = \check{D}^1_\infty & H_2 = \check{D}^1_{0'} & \\ A_{3+3j} = \check{C}(Y_\infty) & f_3 = \check{D}^\infty_{0'} & & \end{array}$$

where we have yet to define $\check{D}^0_{0'}$. The first condition of Lemma 2.4.1 is then satisfied by Proposition 2.3.6 with $l = 1$.

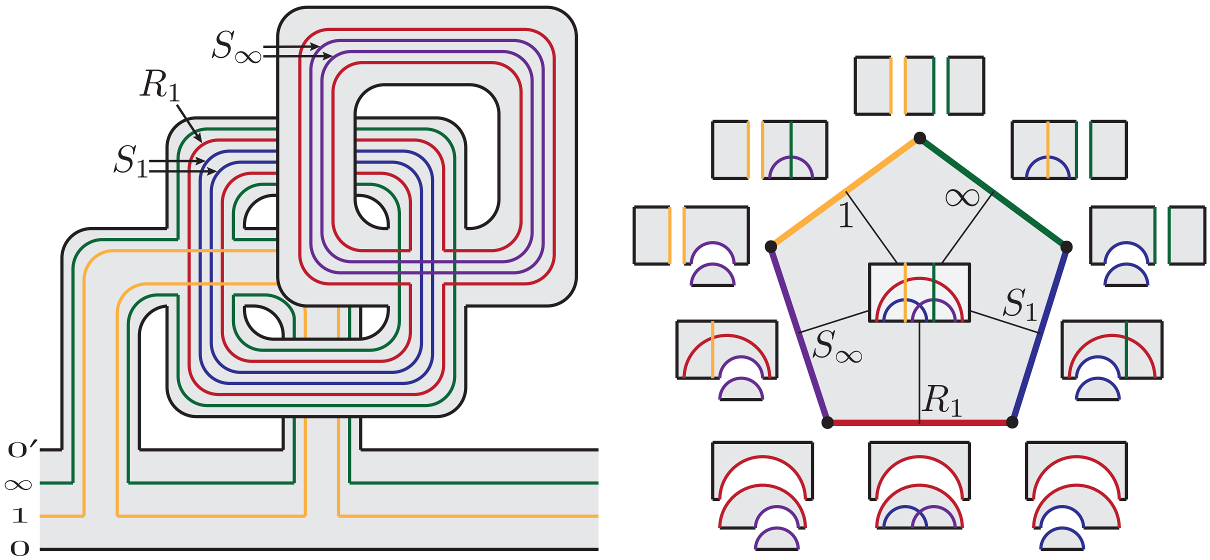

Figure 2.8: At left, the half-dimensional diagram of the cobordism W for the lattice $\{0, 1, \infty, 0'\}$. Note that S_1 is represented by two concentric curves, arising as the boundary of the tubular neighborhood of a circle representing the sphere E_1 (and similarly for S_∞). At right, the pentagon K_4 of metrics, analogous to the hexagon P_3 in Figure 2.2.

Let R denote the edge of the pentagon corresponding to R_1, considered as a one-parameter family of metrics on V_1 stretching from S_1 to S_∞. Viewing V_1 as a cobordism from the empty set to R_1, with the family of metrics R, we have components

$$n_o \in C^o_\bullet(R_1) \qquad n_s \in C^s_\bullet(R_1) \qquad \bar{n}_s \in C^s_\bullet(R_1) \qquad \bar{n}_u \in C^u_\bullet(R_1)$$

In other words, these elements count isolated trajectories in moduli spaces of the form $M_z(V_1^*, \mathfrak{c})_R$ and $M_z^{\mathrm{red}}(V_1^*, \mathfrak{c})_R$. In fact, by Lemma 5.4 of [27], when the perturbation on R_1 is sufficiently small, there are no irreducible critical points and all components of the differential on $\check{C}(R_1)$ vanish, as do n_o and $\bar{n}_s$.

We define the maps $D^*_*\binom{0}{0'}$ exactly as before. We similarly define maps $\bar{D}^{ss}_s\binom{0}{0'}$ and $\bar{D}^{ss}_u\binom{0}{0'}$ which count isolated trajectories in $M_z^{\mathrm{red}}(\mathfrak{a}, \mathfrak{c}, U_1^*, \mathfrak{b})$:

$$\bar{D}^{ss}_s\tbinom{0}{0'} : C^s_\bullet(R_1) \otimes C^s_\bullet(Y_0) \to C^s_\bullet(Y_{0'}) \qquad \bar{D}^{ss}_s\tbinom{0}{0'}(e_\mathfrak{c} \otimes e_\mathfrak{a}) = \sum_{\mathfrak{b} \in \mathfrak{C}^s(Y_{0'})} \sum_z \bar{m}_z(\mathfrak{a}, \mathfrak{c}, U_1^*, \mathfrak{b}) e_\mathfrak{b};$$

$$\bar{D}^{ss}_u\tbinom{0}{0'} : C^s_\bullet(R_1) \otimes C^s_\bullet(Y_0) \to C^u_\bullet(Y_{0'}) \qquad \bar{D}^{ss}_u\tbinom{0}{0'}(e_\mathfrak{c} \otimes e_\mathfrak{a}) = \sum_{\mathfrak{b} \in \mathfrak{C}^u(Y_{0'})} \sum_z \bar{m}_z(\mathfrak{a}, \mathfrak{c}, U_1^*, \mathfrak{b}) e_\mathfrak{b}.$$

We combine these components to define the map $\check{D}^0_{0'} : \check{C}(Y_0) \to \check{C}(Y_{0'})$ by

$$\check{D}^0_{0'} = \begin{bmatrix} D^o_o\binom{0}{0'} & \sum_{0\leq K\leq 0'} D^u_o\binom{K}{0'}\bar{D}^s_u\binom{0}{K} \\ D^o_s\binom{0}{0'} & D^s_s\binom{0}{0'} + \sum_{0\leq K\leq 0'} D^u_s\binom{K}{0'}\bar{D}^s_u\binom{0}{K} \end{bmatrix} + \begin{bmatrix} 0 & D^u_o\binom{0'}{0'}\bar{D}^{ss}_u\binom{0}{0'}(n_s\otimes\cdot) \\ 0 & \bar{D}^{ss}_s\binom{0}{0'}(n_s\otimes\cdot) + D^u_s\binom{0'}{0'}\bar{D}^{ss}_u\binom{0}{0'}(n_s\otimes\cdot) \end{bmatrix}, \tag{2.10}$$

The terms in (2.10) break on a boundary-stable critical point in $\check{C}(R_1)$. Of these, the term $\bar{D}^{ss}_s\binom{0}{0'}(n_s\otimes\cdot)$ is singly boundary-obstructed, while the other two are compositions of a non-boundary obstructed operator and a doubly boundary-obstructed operator (see Definition 24.4.4 in [24]). Finally, we introduce the chain map $\check{L} : \check{C}(Y_0) \to \check{C}(Y_{0'})$ defined by

$$\check{L} = \begin{bmatrix} L^o_o & L^u_o\bar{D}^s_u\binom{0}{0} + D^u_o\binom{0'}{0'}\bar{L}^s_u \\ L^o_s & \bar{L}^s_s + L^u_s\bar{D}^s_u\binom{0}{0} + D^u_s\binom{0'}{0'}\bar{L}^s_u \end{bmatrix} \tag{2.11}$$

where $L^*_* = D^{u*}_*(\bar{n}_u\otimes\cdot)$ and $\bar{L}^*_* = \bar{D}^{u*}_*(\bar{n}_u\otimes\cdot)$. So the coefficient of $\mathfrak{b}$ in $\check{L}(e_\mathfrak{a})$ is a count of the zero-dimensional stratum of $M^+_z(\mathfrak{a},\mathfrak{c},U^*_1,\mathfrak{b})$, over all $\mathfrak{c}$ such that $e_\mathfrak{c}$ is a summand of $\bar{n}_u$.

By Lemma 2.4.4 below, these maps are related by

$$\check{D}^{0'}_{0'}\check{D}^0_{0'} + \check{D}^0_{0'}\check{D}^0_0 = \check{D}^1_{0'}\check{D}^0_1 + \check{D}^\infty_{0'}\check{D}^0_\infty + \check{L}. \tag{2.12}$$

Furthermore, by Proposition 5.6 of [27], $\check{L}$ is a quasi-isomorphism. We conclude that $\check{D}^1_{0'}\check{D}^0_1 + \check{D}^\infty_{0'}\check{D}^0_\infty$ is a quasi-isomorphism as well. This is precisely the second condition of Lemma 2.4.1, which then implies the theorem. □

Remark 2.4.3. In fact, the authors of [27] show that the map induced by $\check{L}$ on $\widetilde{HM}_\bullet(Y_0)$ is given by multiplication by the power series

$$\sum_{k\geq 0} U_\dagger^{k(k+1)/2}.$$

The proof is related to that of the blow-up formula, Theorem 39.3.1 of [24].

Equation (2.12) is proved by counting ends. The maps $A^*_*\binom{0}{0'}$ and $\bar{A}^*_*\binom{0}{0'}$ are defined using the vanishing elements $n_z(\mathfrak{a},W^*,\mathfrak{b})_{P_{IJ}}$ and $\bar{n}_z(\mathfrak{a},W^*,\mathfrak{b})_{P_{IJ}}$ exactly as before. By analogy with the maps D^{ss}_* above, we also define vanishing maps $\bar{A}^{ss}_s\binom{0}{0'}$ and $\bar{A}^{ss}_u\binom{0}{0'}$ which count boundary points of $M_z^{\text{red}+}(\mathfrak{a}, \mathfrak{c}, U_1^*, \mathfrak{b})$. Finally, we define $\check{A}^0_{0'} : \check{C}(Y_0) \to \check{C}(Y_{0'})$ by

$$\begin{aligned}\check{A}^0_{0'} = &\begin{bmatrix} A^o_o\binom{0}{0'} & \sum_{0\le K\le 0'}\left(A^u_o\binom{K}{0'}\bar{D}^s_u\binom{0}{K} + D^u_o\binom{K}{0'}\bar{A}^s_u\binom{0}{K}\right) \\ A^o_s\binom{0}{0'} & A^s_s\binom{0}{0'} + \sum_{0\le K\le 0'}\left(A^u_s\binom{K}{0'}\bar{D}^s_u\binom{0}{K} + D^u_s\binom{K}{0'}\bar{A}^s_u\binom{0}{K}\right)\end{bmatrix} \\ &+ \begin{bmatrix} 0 & A^u_o\binom{0'}{0'}\bar{D}^{ss}_u\binom{0}{0'}(n_s\otimes\cdot) + D^u_o\binom{0'}{0'}\bar{A}^{ss}_u\binom{0}{0'}(n_s\otimes\cdot) \\ 0 & \bar{A}^{ss}_s\binom{0}{0'}(n_s\otimes\cdot) + A^u_s\binom{0'}{0'}\bar{D}^{ss}_u\binom{0}{0'}(n_s\otimes\cdot) + D^u_s\binom{0'}{0'}\bar{A}^{ss}_u\binom{0}{0'}(n_s\otimes\cdot)\end{bmatrix},\end{aligned} \tag{2.13}$$

which therefore vanishes as well. The form of $\check{A}^0_{0'}$ follows from the model case of Morse theory for manifolds with boundary, as described in Appendix I. Note that all the terms in (2.13) break on a boundary-stable critical point in $\check{C}(R_1)$. The term $\bar{A}^{ss}_s\binom{0}{0'}(n_s\otimes\cdot)$ is singly boundary-obstructed, while the other four are compositions of a non-boundary-obstructed operator and a doubly-boundary-obstructed operator.

Lemma 2.4.4. *The map $\check{A}^0_{0'} + \check{L}$ is equal to the component of $\check{D}^2$ from $\check{C}(Y_0)$ to $\check{C}(Y_{0'})$:*

$$\check{A}^0_{0'} + \check{L} = \sum_{0\le K\le 0'} \check{D}^K_{0'}\check{D}^0_K.$$

Proof. As in the proof of Lemma 2.2.7, all terms on the right appear exactly once on the left, with the additional terms on the left being those which do not have a good break on any Y_I. We divide these extra terms into those with

(i) no break on R_1;

(ii) a boundary-stable break on R_1;

(iii) a boundary-unstable break on R_1.

Terms of type (i) can be enumerated just as in the proof of Lemma 2.2.7, so each occurs twice in $\check{A}^0_{0'}$. Dropping indices where it causes no ambiguity, the terms of type (ii) occur in six pairs:

$$\begin{array}{llll}
D^u_o\bar{D}^{su}_u(n_s\otimes\bar{D}^s_u(\cdot)) & \text{in}\ \ A^u_o\bar{D}^s_u & \text{and} & D^u_o\bar{A}^{ss}_u(n_s\otimes\cdot); \\
D^u_o\bar{D}^u_u\bar{D}^{ss}_u(n_s\otimes\cdot) & \text{in}\ \ A^u_o\bar{D}^{ss}_u(n_s\otimes\cdot) & \text{and} & D^u_o\bar{A}^{ss}_u(n_s\otimes\cdot); \\
D^u_s\bar{D}^{su}_u(n_s\otimes\bar{D}^s_u(\cdot)) & \text{in}\ \ A^u_s\bar{D}^s_u & \text{and} & D^u_s\bar{A}^{ss}_u(n_s\otimes\cdot); \\
D^u_s\bar{D}^u_u\bar{D}^{ss}_u(n_s\otimes\cdot) & \text{in}\ \ A^u_s\bar{D}^{ss}_u(n_s\otimes\cdot) & \text{and} & D^u_s\bar{A}^{ss}_u(n_s\otimes\cdot); \\
\bar{D}^{su}_u(n_s\otimes\bar{D}^s_u(\cdot)) & \text{in}\ \ A^u_s\bar{D}^s_u & \text{and} & \bar{A}^{ss}_s(n_s\otimes\cdot); \\
\bar{D}^u_s\bar{D}^{ss}_u(n_s\otimes\cdot) & \text{in}\ \ A^u_s\bar{D}^{ss}_u(n_s\otimes\cdot) & \text{and} & \bar{A}^{ss}_s(n_s\otimes\cdot).
\end{array}$$

Finally, the terms of type (iii) occur in five pairs:

$$
\begin{array}{llrcr}
D^u_o D^{us}_u(\bar{n}_u \otimes \cdot) & \text{in} & D^u_o \bar{A}^s_u & \text{and} & D^u_o \bar{L}^s_u; \\
D^{uu}_o(\bar{n}_u \otimes \bar{D}^s_u(\cdot)) & \text{in} & A^u_o \bar{D}^s_u & \text{and} & L^u_o \bar{D}^s_u; \\
D^u_s D^{us}_u(\bar{n}_u \otimes \cdot) & \text{in} & D^u_s \bar{A}^s_u & \text{and} & D^u_s \bar{L}^s_u; \\
D^{uu}_s(\bar{n}_u \otimes \bar{D}^s_u(\cdot)) & \text{in} & A^u_s \bar{D}^s_u & \text{and} & L^u_s \bar{D}^s_u; \\
\bar{D}^{us}_s(\bar{n}_u \otimes \cdot) & \text{in} & \bar{A}^s_s & \text{and} & \bar{L}^s_s.
\end{array}
$$

We conclude that terms of types (i) and (ii) are double counted by $\check{A}^0_{0'}$ while those of type (iii) are counted once each by $\check{A}^0_{0'}$ and $\check{L}$. We therefore have equality over $\mathbb{F}_2$. □

Remark 2.4.5. If we consider a boundary-unstable break on R_1 to be a good break as well, then Remark 2.2.8 goes through exactly as before. Furthermore, $\check{L}$ counts those trajectories which break well on R_1 (see also the discussion following Proposition 5.5 in [27]).

Remark 2.4.6. For the lattice $\{0, 1, \infty, 0'\}$, we introduced the auxiliary hypersurfaces S_1, S_2, and R_1 in order to build the pentagon of metrics. The S_i edges contribute vanishing terms to $\check{Q}^I_J$ by Lemma 2.3.5, whereas the R_1 edge contributes the term $\check{L}$. Thus,

$$\check{Q}^I_J = \check{D}^1_{0'}\check{D}^0_1 + \check{D}^\infty_{0'}\check{D}^0_\infty + \check{L}$$

and once more we can view (2.12) as a "generalization" of (1.5).

2.5 The link surgery spectral sequence: convergence

We are now positioned to identify the limit of the link surgery spectral sequence.

Proof of Theorem 2.2.11. For $1 \leq k \leq l$, define the map

$$F_k : \bigoplus_{I \in \{\infty\}^{l-k} \times \{0,1\} \times \{0,1\}^{k-1}} \check{C}(Y_I) \longrightarrow \bigoplus_{I \in \{\infty\}^{l-k} \times \{\infty\} \times \{0,1\}^{k-1}} \check{C}(Y_I)$$

as the sum of all compatible components of the differential $\check{D}$ on the subcomplex

$$\bigoplus_{I \in \{\infty\}^{l-k} \times \{0,1,\infty\} \times \{0,1\}^{k-1}} \check{C}(Y_I)$$

of $\widetilde{X}$. Then $\check{D}^2 = 0$ implies that F_k is a chain map. Consider the filtration given by the weight of the last $k-1$ digits of I. By applying the final assertion of Lemma 2.4.1 to the surgery exact triangles arising from the component K_{l-k+1}, we conclude that F_k induces an isomorphism between the E^1 pages of the associated spectral sequences. Therefore, F_k is a quasi-isomorphism, as is the composition

$$F = F_1 \circ F_2 \circ \cdots \circ F_l : X \to \check{C}(Y_\infty). \tag{2.14}$$

□

Remark 2.5.1. The proof of Theorem 2.2.11 hinges on two facts:

(i) lattices of the form $\{0,1\}^k$ and $\{0,1,\infty\} \times \{0,1\}^k$ give rise to filtered complexes;

(ii) the lattice $\{0,1,\infty,0'\}$ gives rise to an exact sequence.

We considered more general lattices in Theorem 2.3.3 and Proposition 2.3.6 in part to make clear how both these facts arise as special cases of the same polytope constructions. The lattice $\{0,1,\infty,0'\} \times \{0,1\}$ will arise naturally in Section 2.6.

2.5.1 Grading

We now introduce an absolute mod 2 grading $\check{\delta}$ on the hypercube complex $(X, \check{D})$ which reduces to $\mathrm{gr}^{(2)}$ in the case $l = 0$. In fact, it will be useful to define $\hat{\delta}$ on the larger complex $(\widetilde{X}, \check{D})$ associated to the lattice $\{0,1,\infty\}^l$. Let $x \in \check{C}(Y_I)$ be homogeneous with respect to the $\mathrm{gr}^{(2)}$ grading. Then we define

$$\begin{aligned}\check{\delta}(x) &= \mathrm{gr}^{(2)}(x) + (\iota(W_{0I}) - w(I)) - (\iota(W_{0\infty}) - 2l) - l \mod 2 \\ &= \mathrm{gr}^{(2)}(x) - (\iota(W_{I\infty}) + w(I)) + l \mod 2.\end{aligned} \tag{2.15}$$

Here the subscripts *0* and ∞ are shorthand for the initial and final vertices of $\{0,1,\infty\}^l$.

Lemma 2.5.2. *The differential $\check{D}$ on $\widetilde{X}$ and X lowers $\hat{\delta}$ by 1.*

Proof. Since $\check{D}^I_J$ is defined using a family of metrics of dimension $w(J) - w(I) - 1$ on W_{IJ}, it shifts $\mathrm{gr}^{(2)}$ by

$$-\iota(W_{IJ}) + (w(J) - w(I) - 1) = (\iota(W_{J\infty}) + w(J)) - (\iota(W_{I\infty}) + w(I)) - 1.$$

The claim now follows from (2.15). □

We now complete the proof of Theorem 2.0.1.

Proposition 2.5.3. *The gradings $\check{\delta}$ and $gr^{(2)}$ coincide under the quasi-isomorphism*

$$F : X \to \check{C}(Y_\infty)$$

defined in (2.14).

Proof. The weight of the vertex $\{\infty\}^l$ is $2l$. Therefore, given $x \in \check{C}(Y_\infty)$, by (2.15) we have

$$\check{\delta}(x) = \mathrm{gr}^{(2)}(x) - l \mod 2.$$

So it suffices to show that the quasi-isomorphism $F : X \to \check{C}(Y_\infty)$ lowers $\check{\delta}$ by l. But F is a composition of l maps F_k, each of which is a sum of components of $\check{D}$. So we are done by the Lemma 2.5.2. □

2.5.2 Invariance

The construction of the hypercube complex

$$X(g, q) = \bigoplus_{I \in \{0,1\}^l} \check{C}(Y_I(g|_I, q|_I))$$

depends on a choice of metric g and admissible perturbation q on the full cobordism W, where the metric is cylindrical near each of the hypersurfaces Y_I. Let (g_0, q_0) and (g_1, q_1) be two such choices.

Theorem 2.5.4. *There exists a $\check{t}$-filtered, $\check{\delta}$-graded chain homotopy equivalence*

$$\phi : X(g_0, q_0) \to X(g_1, q_1),$$

which induces a $(\check{t}, \check{\delta})$-graded isomorphism between the associated E^i pages for all $i \geq 1$.

Proof. We start by embedding a second copy of each Y_I in W as follows (see Figure 2.9 for the case $l = 2$). First, relabel the incoming end Y_0 as $Y_{0\times\{0\}}$ and every other Y_I as $Y_{I\times\{1\}}$. Then embed a second copy of $Y_{0\times\{0\}}$, labeled $Y_{0\times\{1\}}$, just above the original. Finally, embed a second copy of each $Y_{I\times\{1\}}$, labeled $Y_{I\times\{0\}}$, just below the original. We now have an embedded hypersurface $Y_{I\times\{i\}}$ for each $I \times \{i\}$ in the hupercube $\{0,1\}^l \times \{0,1\}$, with diffeomorphisms

$$W_{I\times\{0\},I\times\{1\}} \cong Y_I \times [0,1] \tag{2.16}$$

$$W_{I\times\{i\},J\times\{j\}} \cong W_{IJ} \tag{2.17}$$

where in (2.17) we assume $I < J$. Furthermore, $Y_{I\times\{i\}}$ and $Y_{J\times\{j\}}$ are disjoint if $I \times \{i\}$ and $J \times \{j\}$ are ordered.

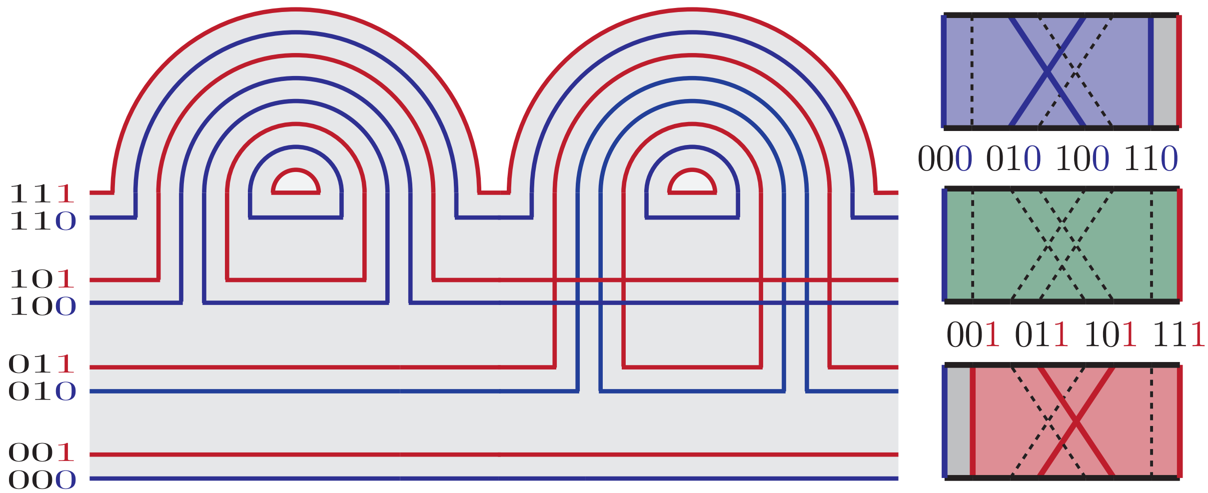

Figure 2.9: At left, we have the half-dimensional diagram of the cobordism W used to prove analytic invariance in the case $l = 2$. For each $I \in \{0,1\}^2$, the hypersurfaces $Y_{I\times\{0\}}$ (in blue) and $Y_{I\times\{1\}}$ (in red) bound a cylindrical cobordism. At right, we can fix the blue metric g_0 on $W_{000,110}$ (top), or the red metric g_1 on $W_{001,111}$ (bottom). The green metric on the middle rectangle represents an intermediate state. To construct the homotopy, we slide the metric from that on the top rectangle to that on the bottom rectangle in a controlled manner, as explained in Figure 2.10.

Our strategy is as follows. We define a complex

$$\underline{\check{X}} = \bigoplus_{I\in\{0,1\}^l, i\in\{0,1\}} \check{C}(Y_{I\times\{i\}}),$$

where the differential $\underline{\check{D}}$ is defined as a sum of components

$$\underline{\check{D}}^{I\times\{i\}}_{J\times\{j\}} : \check{C}(Y_{I\times\{i\}}) \to \check{C}(Y_{J\times\{j\}}).$$

Those components of the form $\check{\underline{D}}^{I\times\{i\}}_{J\times\{i\}}$ are inherited from $X(g_i,q_i)$. So we may view $X(g_0,q_0)$ as the complex over $\{0,1\}^l\times\{0\}$ obtained from quotienting $\check{\underline{X}}$ by the subcomplex $X(g_1,q_1)$ over $\{0,1\}^l\times\{1\}$. The component $\check{\underline{D}}^{I\times\{0\}}_{J\times\{1\}}$ is induced by the cobordism $W_{I\times\{i\},J\times\{j\}}$ over a family of metrics and perturbations parameterized by a permutohedron $\check{\underline{P}}_{I\times\{0\},J\times\{1\}}$, to be defined momentarily. Then $\check{\underline{D}}^2=0$ implies that

$$\phi=\sum_{I\leq J}\check{\underline{D}}^{I\times\{0\}}_{J\times\{1\}}:X(g_0,q_0)\to X(g_1,q_1)$$

is a chain map. If we extend the $\check{\delta}$ grading verbatim to $\check{\underline{X}}$, then ϕ is odd as a map on $\check{\underline{X}}$ by Proposition 2.5.3, and thus even as a map from $X(g_0,q_0)$ and $X(g_1,q_1)$. Thus, ϕ is $\check{\delta}$-graded, and it is clearly $\check{t}$-filtered. By (2.16), the map

$$\check{\underline{D}}^{I\times\{0\}}_{I\times\{1\}}:\check{C}(Y_{I\times\{0\}})\to\check{C}(Y_{I\times\{1\}})$$

induces an isomorphism on homology. Thus, filtering by the horizontal weight $\underline{w}$ defined by $\underline{w}(I\times\{i\})=w(I)$, ϕ induces a $(\check{t},\check{\delta})$-graded isomorphism between the E^1 pages of the corresponding spectral sequences. By Theorem 3.5 of [31], we conclude that ϕ induces a $(\check{t},\check{\delta})$-graded isomorphism between the E^i pages for each $i\geq 1$. Thus, ϕ is a quasi-isomorphism, and therefore (since we are working over a field) a homotopy equivalence.

It remains to construct the family parameterized by each $\check{\underline{P}}_{I\times\{0\},J\times\{1\}}$ and to prove that $\check{\underline{D}}^2=0$. We start by fixing a metric g^I_I on each cylindrical cobordism $W_{I\times\{0\},I\times\{1\}}$ for which $g^I_I(Y_{I\times\{0\}})=g_0(Y_I)$ and $g^I_I(Y_{I\times\{1\}})=g_1(Y_I)$ (we proceed similarly with regard to the perturbations, though we will suppress this). Here the notation $g(Y)$ denotes the restriction of g to Y. The point $\check{\underline{P}}_{I\times\{0\},I\times\{1\}}$ is defined to correspond to the metric g^I_I. Now for each $I\in\{0,1\}^l$, we specify a metric g_I on W by its restriction to each of three pieces:

$$\begin{aligned}g_I(W_{\mathit{0}\times\{0\},I\times\{0\}})&=g_0(W_{\mathit{0}I})\\ g_I(W_{I\times\{0\},I\times\{1\}})&=g^I_I\\ g_I(W_{I\times\{1\},\mathit{1}\times\{1\}})&=g_1(W_{I\mathit{1}}).\end{aligned}$$

We will use these metrics to construct the family parameterized by $\check{\underline{P}}_{\mathit{0}\times\{0\},\mathit{1}\times\{1\}}$ in several stages. The case $l=2$ is illustrated in Figure 2.10.

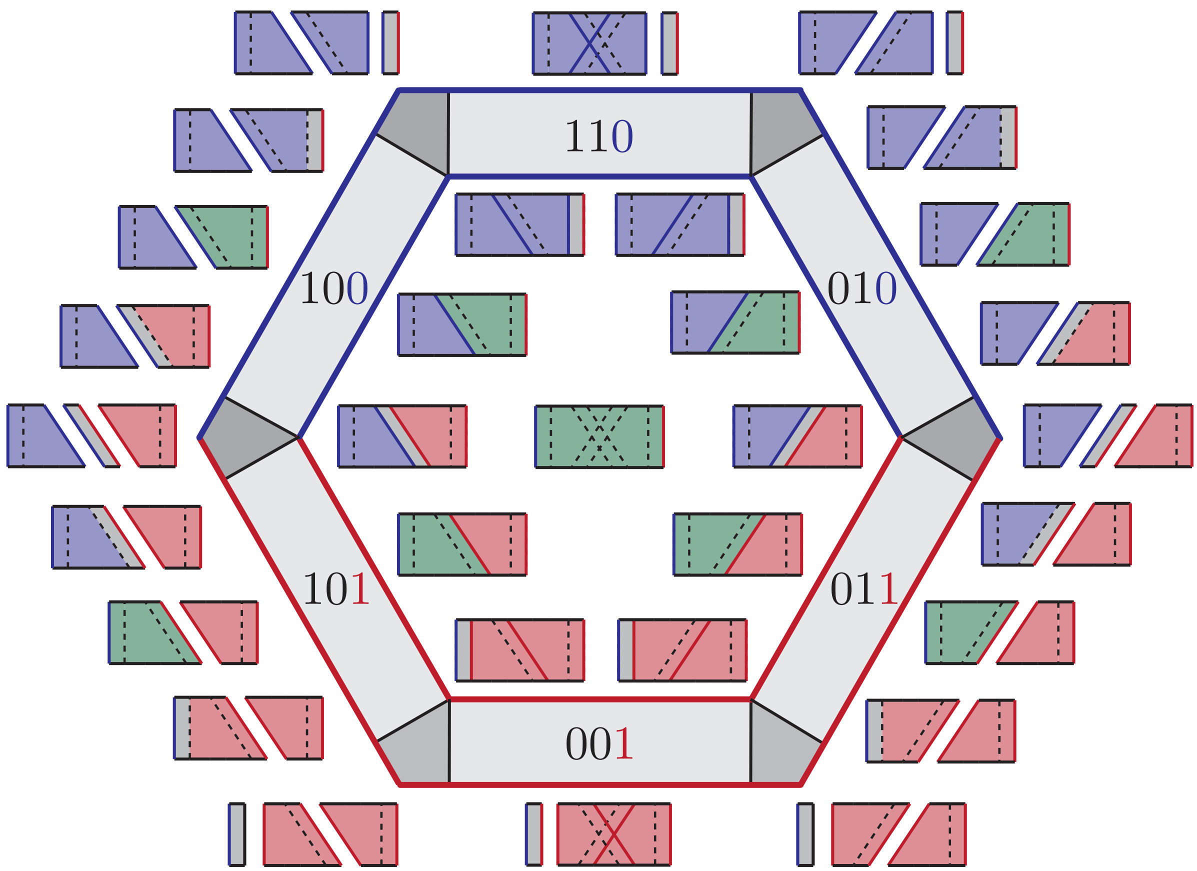

Figure 2.10: The hexagon $\underline{\check{P}}_{111}^{000}$ is drawn so that increasing the vertical coordinate is suggestive of moving from the red metrics to the blue metrics. Gray represents the cylindrical metrics g_I^I, while green represents an intermediate mixture of red, blue, and gray.

We first describe a family $\mathcal{F}$ of non-degerate metrics on W, parameterized by the permutohedron P_{l+1}. Let Q_I denote the facet of P_l corresponding to the internal vertex I. P_{l+1} may be obtained from $P_l \times [0, l]$ by subdividing each facet $Q_I \times [0, 1]$ by the ridge $Q_I \times \{w(I)\}$. (In the $l = 2$ case, this amounts to adding a vertex at the midpoint of each vertical edge in a square. In the $l = 3$ case, shown at right in Figure 3.2, we cross the hexagon with an interval and add an edge to each lateral face. The general case is established in Theorem 3.0.7) We next label the facets $Q_I \times [0, w(I)]$ and $Q_I \times [w(I), l]$ by $I \times \{1\}$ and $I \times \{0\}$, respectively. Furthermore, we label $P_l \times \{0\}$ and $P_l \times \{l\}$ by $\mathit{0} \times \{1\}$ and $\mathit{1} \times \{0\}$, respectively.

We then associate the metric g_I to each vertex of P_{l+1} lying on $Q_I \times \{w(I)\}$. The remaining vertices of P_{l+1} lie on $P_l \times \{0\}$ or $P_l \times \{l\}$. We associate to these vertices the metrics g_{0} and g_{1}, respectively (note that $w(\mathit{0}) = 0$ and $w(\mathit{1}) = l$). At this stage, we have defined $\mathcal{F}$ on the 0-skeleton of P_{l+1}. We proceed inductively: having extended $\mathcal{F}$ to the boundary ∂F of a k-dimensional face F of P_{l+1}, we extend $\mathcal{F}$ to the interior of F, subject to the following constraint:

$$\text{If } \mathcal{F}|_{\partial F} \text{ is constant over some hypersurface or component of } W\text{, then so is } \mathcal{F}|_F. \tag{2.18}$$

In particular, the family $\mathcal{F}$ is constant when restricted to each of the facets $P_l \times \{0\}$ and $P_l \times \{l\}$ and each of the ridges $Q_I \times \{w(I)\}$. Note that the existence of such extensions appeals to the contractibility of the space of metrics on W (or more precisely, the space of metrics on W which extend a fixed metric on a submanifold of W).

The family $\mathcal{F}$ over P_{l+1} slides the metric (and perturbation) on W in stages (in Figure 2.10, P_{2+1} is the inner hexagon). We now extend $\mathcal{F}$ to a family $\mathcal{G}$ which incorporates stretching. To each facet $Q_{I\times\{i\}}$ of P_{l+1}, we glue the polytope $Q_{I\times\{i\}} \times [0, \infty]$ along the facet $Q_{I\times\{i\}} \times \{0\}$ (in Figure 2.10, these are the six lightly shaded rectangles). We extend $\mathcal{G}$ over $Q_{I\times\{i\}} \times [0, \infty]$ by stretching on $Y_{I\times\{i\}}$ in accordance with the latter coordinate (recall that the metric on $Y_{I\times\{i\}}$ is constant over $Q_{I\times\{i\}}$). Next, along each ridge $Q_{I\times\{i\}<J\times\{j\}}$ in P_{l+1}, we glue on the polytope $Q_{I\times\{i\}<J\times\{j\}} \times [0, \infty] \times [0, \infty]$ in the obvious manner (in Figure 2.10, these are the six heavily shaded squares). The first interval parameterizes stretching on $Y_{I\times\{i\}}$ while the second interval parameterizes stretching on $Y_{J\times\{j\}}$. We continue this

process until the last stage, when we glue one cube $[0,\infty]^l$ at each vertex of P_{l+1}, over which $\mathcal{G}$ stretches on the corresponding maximal chain of internal hypersurfaces.

In the end, we have simply thickened the boundary of P_{l+1} to describe a family $\mathcal{G}$ of metrics on W parameterized by the permutohedron $\underline{\check{P}}_{0\times\{0\},1\times\{1\}}$ (the full hexagon in Figure 2.10). This family is degenerate over the boundary of $\underline{\check{P}}_{0\times\{0\},1\times\{1\}}$ precisely as described by Proposition 2.1.4. Now, for each $I \leq J$, we construct a family of metrics $\mathcal{G}_{IJ}$ over $\underline{\check{P}}_{I\times\{0\},J\times\{1\}}$ by restricting the family $\mathcal{G}$ to $W_{I\times\{0\},J\times\{1\}}$ over an appropriate face of $\underline{\check{P}}_{0\times\{0\},1\times\{1\}}$ (here the constraint (2.18) is essential).

The proof that $\underline{\check{D}}^2 = 0$ now lifts directly from the original proof that $\check{D}^2 = 0$, with one new point that we now explain. The component of $\underline{\check{D}}^2$ from $\check{C}(Y_{0\times\{0\}})$ to $\check{C}(Y_{1\times\{1\}})$ vanishes if and only if

$$\begin{aligned}\underline{\check{D}}^{0\times\{0\}}_{1\times\{1\}}\check{D}^{0\times\{0\}}_{0\times\{0\}} + \check{D}^{1\times\{1\}}_{1\times\{1\}}\underline{\check{D}}^{0\times\{0\}}_{1\times\{1\}} &= \underline{\check{D}}^{1\times\{0\}}_{1\times\{1\}}\check{D}^{0\times\{0\}}_{1\times\{0\}} + \check{D}^{0\times\{1\}}_{1\times\{1\}}\underline{\check{D}}^{0\times\{0\}}_{0\times\{1\}} \\ &\quad + \sum_{0<I<1} \underline{\check{D}}^{I\times\{0\}}_{1\times\{1\}}\check{D}^{0\times\{0\}}_{I\times\{0\}} + \check{D}^{I\times\{1\}}_{1\times\{1\}}\underline{\check{D}}^{0\times\{0\}}_{I\times\{1\}}. \end{aligned} \tag{2.19}$$

Consider the composite map corresponding to the family $\mathcal{G}$ over the facet $1 \times \{0\}$ of $\underline{\check{P}}_{0\times\{0\},1\times\{1\}}$. Since the family $\mathcal{F}$ over the corresponding facet of P_{l+1} is constant, the only sections of the facet $1 \times \{0\}$ which contributes non-trivially to this map are those of the form $\{\infty\} \times [0,\infty]^{l-l}$ in the boundary of the cubes $[0,\infty]^l$ (in Figure 2.10, these are the two segments of the top edge of the hexagon which lie in the boundary of the heavily shaded squares). The other sections cannot give rise to 0-dimensional moduli spaces, since they involve at least one parameter which does not change the metric. We can therefore identify the map associated to the facet $1 \times \{0\}$ with $\underline{\check{D}}^{1\times\{0\}}_{1\times\{1\}}\check{D}^{0\times\{0\}}_{1\times\{0\}}$ (in Figure 2.10, we are contracting out the middle segment of the top edge). Similarly, the map associated to the facet $0 \times \{1\}$ coincides with $\underline{\check{D}}^{1\times\{0\}}_{1\times\{1\}}\check{D}^{0\times\{0\}}_{1\times\{0\}}$, and the sum on line (2.19) coincides with the map associated to the remaining lateral facets of $\underline{\check{P}}_{0\times\{0\},1\times\{1\}}$. Thus, the full equation expresses the fact that the map $\underline{\check{D}}^{0\times\{0\}}_{1\times\{1\}}$ associated to the full permutohedron is a null-homotopy for the map associated to its boundary. The other components of $\check{D}^2$ vanish by a completely analogous argument. □

Remark 2.5.5. Recall the top and bottom rectangles at right in Figure 2.9. Suppose that the red and blue metrics agree where they overlap, so that the family $\mathcal{F}$ on P_{l+1} can be made completely constant. Then only the cubes $[0,\infty]^l$ contribute non-trivially to the map $\underline{\check{D}}^{\mathit{0}\times\{0\}}_{\mathit{1}\times\{1\}}$. Discarding the rest of $\underline{\check{P}}_{\mathit{0}\times\{0\},\mathit{1}\times\{1\}}$ and gluing these cubes together, we build a permutohedron giving rise to the same map. This viewpoint highlights the connection between the permutohedra $\underline{\check{P}}_{I\times\{0\},J\times\{1\}}$ and the permutohedra P_{IJ} that we first constructed in Section 2.1, using only cubes which stretch the metric along maximal chains of internal hypersurfaces.

Remark 2.5.6. The construction in Section 2.1 starts with an initial metric g_0 which is required be cylindrical near all hypersurfaces Y_I simultaneously. In fact, by proceeding as in Figure 2.10, we can instead start with a finite collection of initial metrics for which each metric need only be cylindrical near a subset of pairwise disjoint hypersurfaces. More precisely, we first fix a cylindrical metric on a neighborhood of each hypersurface Y_I in W (these metrics need *not* be mutually compatible). Then, for each pair of immediate successors $I < J$, we fix a metric on W_{IJ} which extends the corresponding metric near each boundary component. Each vertex of the permutohedron P_l expresses W as a composition of l elementary cobordisms W_{IJ} and therefore determines a metric on W. Inducting up from the 0-skeleton, we define a family of metrics parameterized by all of P_l, imposing condition (2.18) as before. Finally, we enlarge P_l to a family $\underline{P}_l$ which incorporates stretching as well.

Remark 2.5.7. Note that any two filtered chain maps

$$\phi_1, \phi_2 : X(g_0, q_0) \to X(g_1, q_1)$$

constructed using the recipe in the proof of invariance are related by a filtered chain homotopy, and thus induce the same maps on E^k for $k \geq 1$. The filtered chain homotopy is constructed just as in the proof, by building in yet another factor of $\{0, 1\}$ into the lattice (so there are four copies of each hypersurface Y_I embedded in W). In this way, from a framed link $L \subset Y$, we obtain a spectral sequence whose pages E^i for $i \geq 1$ are "groups up to canonical isomorphism", with the later category admitting a functor to GROUP by taking "cross-sections" (see Section 23.1 of [24]). We may therefore regard the pages E^i for $i \geq 1$ as actual groups associated to a framed link L in Y, rather than just isomorphism classes of groups.

2.6 The $U_\dagger$ map and $\widetilde{HM}_\bullet(Y)$

Given a cobordism $W : Y_0 \to Y_1$, Kronheimer and Mrowka construct a map $U_\dagger = \widetilde{HM}_\bullet(u|W) : \widetilde{HM}_\bullet(Y_0) \to \widetilde{HM}_\bullet(Y_1)$. In [24], this map is defined by pairing each moduli space $M_z(\mathfrak{a}, W^*, \mathfrak{b})$ with a representative u of the first Chern class of the natural complex line bundle on $\mathcal{B}^\sigma(W)$. A dual description of the map is given in [27]. We will use the notation $\check{m}(u|W)$ for this map on the chain level.

We will employ a third description which fits in neatly with our previous constructions. We first recall some facts about the monopole Floer homology of the 3-sphere (see Sections 22.7 and 25.6 of [24]). With round metric and small perturbation, the monopoles on the 3-sphere consist of a single bi-infinite tower $\{\mathfrak{c}_i\}_{i\in\mathbb{Z}}$ of reducibles, with $\mathfrak{c}_i$ boundary-stable if and only if $i \geq 0$. As shown in Lemma 27.4.2 of [24], the moduli space $M((D^4)^*, \mathfrak{c}_i)$ is empty for $i \geq 0$ and has dimension $-2i - 2$ for $i \leq 0$. In particular, $M((D^4)^*, \mathfrak{c}_{-1})$ consists of a single point. Furthermore, $U_\dagger$ sends $\mathfrak{c}_{-1}$ to $\mathfrak{c}_{-2}$, so the pairing $\langle u, M((D^4)^*, \mathfrak{c}_{-2})\rangle$ evaluates to 1. This motivates the following reformulation.

Given a cobordism $W : Y_0 \to Y_1$, let W^{**} denote the manifold obtained by removing a ball from the interior of W and attaching cylindrical ends to all three boundary components, with the new $S^3 \times [0, \infty)$ end regarded as incoming. Choose the metric and perturbation on W so that we return to the situation described in the last paragraph over S^3. We define the map $\check{m}(U_\dagger|W) : \check{C}(Y_0) \to \check{C}(Y_1)$ by replacing each moduli space $M_z(\mathfrak{a}, W^*, \mathfrak{b})$ in the definition of $\check{m}(W)$ with the moduli space $M_z(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}, \mathfrak{b})$. In other words,

$$\check{m}(U_\dagger|W) = \begin{bmatrix} m_o^{uo}(\mathfrak{c}_{-2} \otimes \cdot) & m_o^{uu}(\mathfrak{c}_{-2} \otimes \bar{\partial}_u^s(\cdot)) + \partial_o^u \bar{m}_u^{us}(\mathfrak{c}_{-2} \otimes \cdot) \\ m_s^{uo}(\mathfrak{c}_{-2} \otimes \cdot) & \bar{m}_s^{us}(\mathfrak{c}_{-2} \otimes \cdot) + m_s^{uu}(\mathfrak{c}_{-2} \otimes \bar{\partial}_u^s(\cdot)) + \partial_s^u \bar{m}_u^{us}(\mathfrak{c}_{-2} \otimes \cdot) \end{bmatrix}.$$

One sees that $\check{m}(U_\dagger|W)$ is a chain map and well-defined up to homotopy equivalence by the same argument used for $\check{m}(W)$, together with the fact that there are no isolated trajectories from $\mathfrak{c}_{-2}$ to any other $\mathfrak{c}_i \in \mathfrak{C}(S^3)$.

Proposition 2.6.1. *The map $\check{m}(U_\dagger|W)$ is homotopy equivalent to the map $\check{m}(u|W)$.*

Proof. Given a cobordism W, we may assume the cochain u is supported over the configuration space of a small ball. Now the homotopy which stretches the metric normal to the 3-sphere bounding this ball reduces the claim to the local computation that $\langle u, M((D^4)^*, \mathfrak{a}_{-2})\rangle = 1$. □

Remark 2.6.2. We could instead allow for a generic choice of metric and perturbation on the S^3 end, replacing $\mathfrak{a}_{-2}$ by any chain in $\hat{C}(S^3)$ representing the same class in $\widehat{HM}_\bullet(S^3)$. Such an approach would invoke our recent development of monopole Floer maps for cobordisms with multiple ends, as described at end of the Appendix.

We now define a fourth version of monopole Floer homology, denoted $\widetilde{HM}_\bullet(Y)$ and analogous to the group $\widehat{HF}(Y)$ in Heegaard Floer homology. Here we encounter a subtle "point", which arises for $\widehat{HF}(Y)$ as well. Namely, it seems necessary to equip Y with a basepoint y in order to define an actual group $\widetilde{HM}_\bullet(Y, y)$. To a 3-manifold Y without basepoint, we only have an isomorphism class of (graded) group $\widetilde{HM}_\bullet(Y)$. We will elaborate on this point at the end of the section in Remark 2.6.6.

We now proceed to define the (isomorphism class of) group $\widetilde{HM}_\bullet(Y)$. Fix an (open) ball D^4 in Y and equip the cobordism $(Y \times [0,1]) - D^4$ with a metric and perturbation which restrict to the S^3 boundary component as above. We use the shorthand $U_\dagger$ for the map $\check{m}(U_\dagger | Y \times [0,1]) : \check{C}(Y) \to \check{C}(Y)$ induced by this cobordism.

The complex $\tilde{C}(Y)$ is defined to be the mapping cone of $U_\dagger$:

$$\tilde{C}(Y) = \check{C}(Y) \bigoplus \check{C}(Y)\{1\} \qquad\qquad \tilde{\partial} = \begin{bmatrix} \check{\partial} & 0 \\ U_\dagger & \check{\partial} \end{bmatrix}$$

Since $U_\dagger$ is an even map, the differential $\tilde{\partial}$ is odd, and therefore $\mathrm{gr}^{(2)}$ naturally extends to $\tilde{C}(Y)$ (as does $\mathrm{gr}^{\mathbb{Q}}$ for torsion spin^c structures). We then define $\widetilde{HM}_\bullet(Y)$ as the (negative completion[1] of the homology $H_*(\tilde{C}(Y), \tilde{\partial})$. By construction, there is an exact sequence

$$\cdots \xrightarrow{l_*} \widetilde{HM}_\bullet(Y) \xrightarrow{k_*} \widecheck{HM}_\bullet(Y) \xrightarrow{U_\dagger} \widecheck{HM}_\bullet(Y) \xrightarrow{l_*} \cdots \tag{2.20}$$

of $\mathbb{F}_2[[U_\dagger]]$ modules where $U_\dagger$ acts by zero on $\widetilde{HM}_\bullet(Y)$. Here the maps k_*, $U_\dagger$, and l_* have degrees 0, -2, and 1, respectively.

The construction of a chain map $\tilde{m}(W) : \tilde{C}(Y_0) \to \tilde{C}(Y_1)$ from a cobordism W is similar to the $l = 1$ case of the $\widetilde{HM}_\bullet$ spectral sequence. We first relabel the ends of W as Y_{00}

[1] As with $\widecheck{HM}_\bullet$, this completion has no real effect here, but we keep the bullet for notational consistency.

and Y_{11} and embed a second copy of each in the interior of W as follows (see Figure 2.11). Let γ be a path in W from a point y_0 in Y_0 to a point y_1 in Y_1. Fix a small ball inside a tubular neighborhood $\nu(\gamma)$ of γ. The 3-manifold Y_{01} is obtained by taking a parallel copy of Y_{00} just inside the boundary and pushing the region in $\nu(\gamma)$ past the ball, so that cutting along Y_{01} leaves the ball in the first component $W_{00,01} \cong Y_0 \times [0,1]$. Similarly, Y_{10} is obtained by taking a parallel copy of Y_{11} near the boundary and pushing the region inside $\nu(\gamma)$ inward past the ball, so that cutting along Y_{10} leaves the ball in the second component $W_{10,11} \cong Y_1 \times [0,1]$. Note that both $W_{00,10}$ and $W_{01,11}$ are diffeomorphic to W.

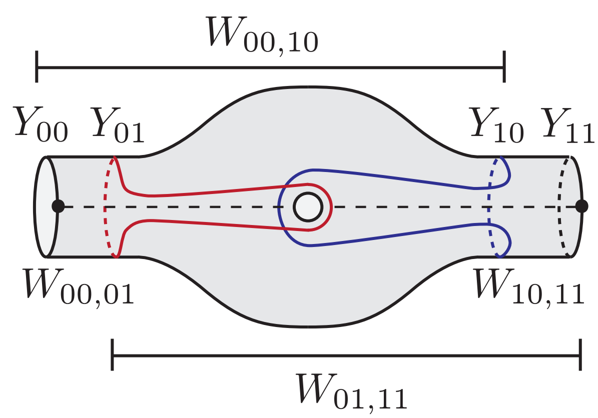

Figure 2.11: The surface above represents the cobordism used to construct the chain map $\tilde{m}(W) : \tilde{C}(Y_0) \to \tilde{C}(Y_1)$. The path γ is depicted as a dotted line.

The intersection of Y_{01} and Y_{10} is modeled on $S^2 \times \{0\} \times \{0\} \subset S^2 \times (-\epsilon, \epsilon) \times (-\epsilon, \epsilon)$, so we can choose the metric on W to be cylindrical near both internal hypersurfaces (or we may choose the metrics independently as in Remark 2.5.6). Consider the interval of metrics $\tilde{P}_{00,11} = [-\infty, \infty]$ which expands a cylindrical neck along Y_{01} as the parameter decreases from 0 and expands a cylindrical neck along Y_{10} as the parameter increases from 0. As in Section 2.1, at $\pm\infty$ we have not metrics but disjoint unions. We use this interval to define eight operators $H^{u*}_{*}(\mathfrak{c}_{-2} \otimes \cdot)$ which count isolated trajectories in the moduli space

$$M(\mathfrak{a}, W^*, \mathfrak{b})_{\tilde{P}_{00,11}} = \bigcup_{p \in \tilde{P}_{00,11}} \bigcup_{z} M_z(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}(p), \mathfrak{b})$$

where

$$M_z(\mathfrak{a}, W(-\infty)^*, \mathfrak{b}) = \bigcup_{\mathfrak{c}\in\check{C}(Y_{01})} \bigcup_{z_1,z_2} M_{z_1}(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}_{00,01}, \mathfrak{c}) \times M_{z_2}(\mathfrak{c}, W^*_{01,11}, \mathfrak{b}),$$

and

$$M_z(\mathfrak{a}, W(\infty)^*, \mathfrak{b}) = \bigcup_{\mathfrak{c}\in\check{C}(Y_{10})} \bigcup_{z_1,z_2} M_{z_1}(\mathfrak{a}, W^*_{00,10}, \mathfrak{c}) \times M_{z_2}(\mathfrak{c}, \mathfrak{c}_{-2}, W^{**}_{10,11}, \mathfrak{b}).$$

We then define $\check{H}(U_\dagger|W_{00,11})$ by the same expression as $\check{D}^{00}_{11}$ in (2.5), except that if I ends in 0 and J ends in 1, then $D^*_*\binom{I}{J}$ is replaced by $D^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2}\otimes\cdot)$. So in full, we have

$$\check{H}(U_\dagger|W_{00,11}) = \begin{bmatrix} H^{uo}_o(\mathfrak{c}_{-2}\otimes\cdot) & H^{uu}_o(\mathfrak{c}_{-2}\otimes\bar{\partial}^s_u(\cdot)) + \partial^u_o\bar{H}^{us}_u(\mathfrak{c}_{-2}\otimes\cdot) \\ H^{uo}_s(\mathfrak{c}_{-2}\otimes\cdot) & \bar{H}^{us}_s(\mathfrak{c}_{-2}\otimes\cdot) + H^{uu}_s(\mathfrak{c}_{-2}\otimes\bar{\partial}^s_u(\cdot)) + \partial^u_s\bar{H}^{us}_u(\mathfrak{c}_{-2}\otimes\cdot) \end{bmatrix}$$
$$+ \begin{bmatrix} 0 & m^u_o\binom{01}{11}\bar{m}^{us}_u\binom{00}{01}(\mathfrak{c}_{-2}\otimes\cdot) + m^{uu}_o\binom{10}{11}(\mathfrak{c}_{-2}\otimes\bar{m}^s_u\binom{00}{10}(\cdot)) \\ 0 & m^u_s\binom{01}{11}\bar{m}^{us}_u\binom{00}{01}(\mathfrak{c}_{-2}\otimes\cdot) + m^{uu}_s\binom{10}{11}(\mathfrak{c}_{-2}\otimes\bar{m}^s_u\binom{00}{10}(\cdot)) \end{bmatrix}$$

From this perspective, the differentials on $\tilde{C}(Y_0)$ and $\tilde{C}(Y_1)$ are

$$\tilde{\partial}(Y_0) = \begin{bmatrix} \check{\partial}(Y_{00}) & 0 \\ \check{m}(U_\dagger|W_{00,01}) & \check{\partial}(Y_{01}) \end{bmatrix} \quad \text{and} \quad \tilde{\partial}(Y_1) = \begin{bmatrix} \check{\partial}(Y_{10}) & 0 \\ \check{m}(U_\dagger|W_{10,11}) & \check{\partial}(Y_{11}) \end{bmatrix},$$

respectively. Finally, the map $\tilde{m}(W) : \tilde{C}(Y_0) \to \tilde{C}(Y_1)$ is defined by

$$\tilde{m}(W) = \begin{bmatrix} \check{m}(W_{00,10}) & 0 \\ \check{H}(U_\dagger|W_{00,11}) & \check{m}(W_{01,11}) \end{bmatrix}. \tag{2.21}$$

The fact that $\tilde{m}(W)$ is a chain map is a special case of the construction of the total complex underlying the $\widetilde{HM}_\bullet$ version of the link surgery spectral sequence, as explained in Remark 2.6.4. We now turn to constructing this total complex, which we will denote here by $(\underline{X}, \underline{D})$ with pages $\underline{E}^i$, though in other sections we may return to the notation $(X, \check{D})$ and E^i when it is clear from context which version is intended.

Given an l-component framed link $L \subset Y$ and a point y in the link complement, we embed a small ball D^4 in the interior of a neighborhood $\nu(\gamma)$ of a path γ with image $\{y\} \times [0,1] \subset (Y - \nu(L)) \times [0,1] \subset W$. Next we relabel the incoming end $Y_{\{0\}^l}$ as $Y_{\{0\}^l\times\{0\}}$ and every other Y_I as $Y_{I\times\{1\}}$. We then embed a second copy of $Y_{\{0\}^l\times\{0\}}$, labeled $Y_{\{0\}^l\times\{1\}}$, just above the first, modified so that it now passes above the ball. Finally, we embed a second copy of each $Y_{I\times\{1\}}$, labeled $Y_{I\times\{0\}}$, just below the first, modified so that it now passes below the ball, using the path γ as a guide. See Figure 2.12 for the case $l = 2$. We now have an embedded hypersurface Y_I for each $I \in \{0,1\}^{l+1}$. Furthermore, the intersection data is precisely what we expect for this hypercube, namely that Y_I intersects Y_J if and only if I and J are not ordered. Therefore, given any $I < J$ we may construct a family of metrics on W_{IJ} parameterized by a permutohedron $\underline{P}^I_J$ of dimension $w(J) - w(I) - 1$.

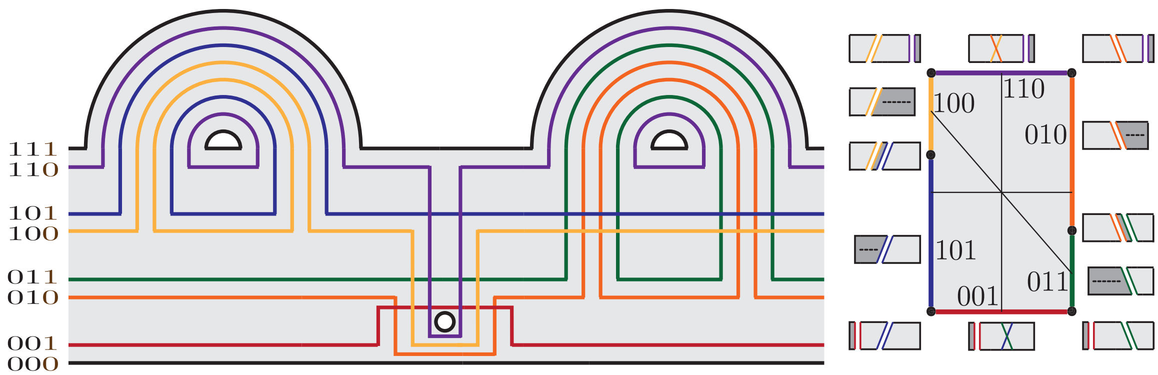

Figure 2.12: At left, we have the half-dimensional diagram of the cobordism W with a small ball removed in the case $l = 2$. For each $I \in \{0,1\}^2$, the pair $Y_{I\times\{0\}}$ and $Y_{I\times\{1\}}$ bound a cylindrical cobordism containing the ball. At right, we have drawn the corresponding hexagon $\underline{P}^{000}_{111}$ so that increasing the vertical coordinate is suggestive of translating the sphere through W. The small figures at the vertices and edges illustrate the metric degenerations, read as composite cobordisms from left to right. In each, the component containing the sphere is more heavily shaded.

Now fix a metric on the cobordism W which is cylindrical near every hypersurface Y_I and round near S^3. We will define a complex

$$\underline{X} = \bigoplus_{I\in\{0,1\}^l\times\{0,1\}} \check{C}(Y_I), \tag{2.22}$$

where the differential $\underline{D} : \underline{X} \to \underline{X}$ is the sum of components $\underline{D}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ over all $I \leq J$. We have set things up so that the ball is contained in W_{IJ} if and only if I ends in 0 and J ends in 1. So when I and J end in the same digit, the operators $D^*_*\binom{I}{J}$ may be defined exactly as before (see (2.4)). In the other case, we construct operators $D^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2}\otimes\cdot)$ using moduli spaces $M_z(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}_{IJ}, \mathfrak{b})_{\underline{P}_{IJ}}$ which are defined by slightly modifying the definition of the moduli spaces $M_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$. Namely, if $p \in \underline{P}^I_J$ is in the interior of the face $I_1 < I_2 < \cdots < I_{q-1}$, with the last digit changing between I_k and I_{k+1}, then an element of $M_z(\mathfrak{a}, \mathfrak{c}_{-2}, W_{IJ}(p)^{**}, \mathfrak{b})$ is a q-tuple

$$(\gamma_{01}, \gamma_{12}, \ldots, \gamma_{q-1\,q})$$

as before except that

$$\gamma_{k\,k+1} \in M(\mathfrak{a}_j, \mathfrak{c}_{-2}, W^{**}_{I_k I_{k+1}}(p), \mathfrak{a}_{k+1}).$$

Let $\underline{D}^*_*\binom{I}{J}$ be synonymous with $D^*_*\binom{I}{J}$ if I and J end in the same digit, and with $D^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2} \otimes \cdot)$ otherwise. Similar remarks apply to $\bar{D}^*_*\binom{I}{J}$ and $\bar{D}^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2} \otimes \cdot)$. We then define $\underline{D}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ by precisely the same expression as (2.5), with each D underlined.

The proof that $\underline{D}^2 = 0$ goes along familiar lines. The operators $A^*_*\binom{I}{J}$ may be defined exactly as before when I and J end in the same digit. When I ends in 0 and J ends in 1, we define operators $A^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2} \otimes \cdot)$ which count ends of 1-dimensional moduli spaces $M^+_z(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}_{IJ}, \mathfrak{b})_{\underline{P}_{IJ}}$, which in turn are defined by slightly modifying the definition of the moduli spaces $M^+_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$ in the same manner as above. As before, these operators all vanish. Now let $\underline{A}^*_*\binom{I}{J}$ be synonymous with $A^*_*\binom{I}{J}$ if I and J end in the same digit, and with $A^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2} \otimes \cdot)$ otherwise. Similar remarks apply to $\bar{A}^*_*\binom{I}{J}$ and $\bar{A}^{u*}_*\binom{I}{J}(\mathfrak{c}_{-2} \otimes \cdot)$. We then define $\underline{A}^I_J : \check{C}(Y_I) \to \check{C}(Y_J)$ by precisely the same expression as (2.7), with each D and A underlined.

Lemma 2.6.3. *$\underline{A}^I_J$ is equal to the component of $\underline{D}^2$ from $\check{C}(Y_I)$ to $\check{C}(Y_J)$:*

$$\underline{A}^I_J = \sum_{I \le K \le J} \underline{D}^K_J \underline{D}^I_K.$$

Thus, $\underline{D}$ is a differential.

Proof. Recall that $\mathfrak{c}_{-2}$ is a boundary-unstable and that there are no isolated trajectories from $\mathfrak{c}_{-2}$ to any other $\mathfrak{c}_i \in \mathfrak{C}(S^3)$. It follows that 1-dimensional moduli spaces $M^+_z(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}_{IJ}, \mathfrak{b})_{\tilde{P}_{IJ}}$ have the same types of ends as $M^+_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{P_{IJ}}$, as described in Lemma 2.2.3. Similarly, $M^{\text{red}+}_z(\mathfrak{a}, \mathfrak{c}_{-2}, W^{**}_{IJ}, \mathfrak{b})_{\underline{P}_{IJ}}$ has the same types of ends as $M^{\text{red}+}_z(\mathfrak{a}, W^*_{IJ}, \mathfrak{b})_{\underline{P}_{IJ}}$. Now we simply repeat the proof of Lemma 2.2.7 with everything underlined. $\square$

Remark 2.6.4. The case $l = 0$ shows that $U_\dagger : \check{C}(Y_0) \to \check{C}(Y_0)$ is a chain map. The case $l = 1$ shows that $\tilde{m}(W) : \tilde{C}(Y_0) \to \tilde{C}(Y_1)$ is a chain map when W is an elementary 2-handle cobordism, and goes through without change for arbitrary cobordisms.

Remark 2.6.5. By the method of Remark 2.5.6, we can choose metric and perturbation data so that the complexes

$$\left(\bigoplus_{I \in \{0,1\}^l \times \{0\}} \check{C}(Y_I), \quad \sum_{I \leq J} \check{D}^{I \times \{0\}}_{J \times \{0\}} \right)$$

and

$$\left(\bigoplus_{I \in \{0,1\}^l \times \{1\}} \check{C}(Y_I), \quad \sum_{I \leq J} \check{D}^{I \times \{1\}}_{J \times \{1\}} \right)$$

are canonically identified with a fixed total complex $(X, \check{D})$. Let

$$\check{U} = \sum_{I \leq J} \underline{D}^{I \times \{0\}}_{J \times \{1\}}.$$

Then Lemma 2.6.3 implies that $\check{U} : (X, \check{D}) \to (X, \check{D})$ is a filtered chain map. We thereby get an induced chain map u^i on each page E^i of the spectral sequence for X, with u^∞ identified with the map $U_\dagger$ under the isomorphism $E^\infty \cong \widetilde{HM}_\bullet(Y)$ induced by (2.14). We will make use of this viewpoint in Chapter 8.

In order to interpret Lemma 2.6.3 as a result in tilde theory, we collapse

$$\underline{X} = \bigoplus_{I \in \{0,1\}^l} \tilde{C}(Y_I)$$

along the final digit, with $\underline{D}$ given by the sum of maps $\tilde{D}^I_J = \tilde{C}(Y_I) \to \tilde{C}(Y_J)$ where

$$\tilde{D}^I_J = \begin{bmatrix} \check{D}^{I \times \{0\}}_{J \times \{0\}} & 0 \\ \underline{D}^{I \times \{0\}}_{J \times \{1\}} & \check{D}^{I \times \{1\}}_{J \times \{1\}} \end{bmatrix}. \tag{2.23}$$

Note that this generalizes (2.21). Define the horizontal weight $\underline{w}(I)$ of a vertex I to be the sum of all but the final digit. Filtering $(\underline{X}, \underline{D})$ by $\underline{w}$, we obtain the $\widetilde{HM}_\bullet$ version of the link surgery spectral sequence. In particular,

$$\underline{E}^1 = \bigoplus_{I \in \{0,1\}^l} \widetilde{HM}_\bullet(Y_I)$$

and the $\underline{d}^1$ differential is given by

$$\underline{d}^1 = \bigoplus_{w(J)-w(I)=1} \widetilde{HM}_\bullet(W_{IJ}).$$

In order to identify $\underline{E}^\infty$ with $\widetilde{HM}_\bullet(Y_\infty)$, we expand to the larger complex

$$\widetilde{\underline{X}} = \bigoplus_{I\in\{0,1,\infty\}^l\times\{0,1\}} \check{C}(Y_I), \tag{2.24}$$

where we have again relabeled the Y_I and embedded a second copy of each which passes on the opposite side of the ball. These hypersurfaces all avoid the auxiliary S^3 hypersurfaces which are confined in the handles. The intersection data is as predicted for the shape of the lattice by Theorem 2.3.3, so we may build families of metrics parameterized by graph associahedra which define maps $\underline{D}^I_J$ and $\underline{A}^I_J$. The auxiliary hypersurfaces still cut out $\overline{\mathbb{CP}}^2 - \text{int}(D^4)$, so the corresponding facets contribute vanishing operators as before. We conclude from 2.6.3 that $(\widetilde{\underline{X}}, \underline{D})$ forms a complex.

Finally, we turn to the surgery exact triangle. Recall the hypersurfaces Y_0, Y_1, Y_∞, $Y_{0'}$ and auxiliary S_1, S_2, and R_1. After relabeling, these become Y_{00}, Y_{10}, $Y_{\infty 0}$, $Y_{0'1}$, S_1, S_2, and R_1, to which we add Y_{01}, Y_{11}, $Y_{\infty 1}$, and $Y_{0'0}$. The nine hypersurfaces in the interior of W intersect as predicted by the shape of the lattice $\{0, 1, \infty, 0'\} \times \{0, 1\}$, yielding the graph associahedron on a chain of length 4, namely, the 3-dimensional associahedron K_5, shown in Figure 2.13.

The map $\underline{L}^{00}_{0'1} : \check{C}(Y_{00}) \to \check{C}(Y_{0'1})$ associated to R_1 is given by the same expression as $\check{L}$, but with $L^*_* = D^{uu*}_*(\mathfrak{c}_{-2} \otimes \bar{n}_u \otimes \cdot)$ and $\bar{L}^*_* = \bar{D}^{uu*}_*(\mathfrak{c}_{-2} \otimes \bar{n}_u \otimes \cdot)$, over the one-parameter family of metrics stretching from Y_{01} to $Y_{0'0}$. This gives the identity

$$\check{D}^{0'1}_{0'1}\underline{L}^{00}_{0'1} + \underline{L}^{00}_{0'1}\check{D}^{00}_{00} = \check{L}^{01}_{0'1}\underline{D}^{00}_{01} + \underline{D}^{0'0}_{0'1}\check{L}^{00}_{0'0}, \tag{2.25}$$

where $\check{L}^{00}_{0'0}$ and $\check{L}^{01}_{0'1}$ are the analogues of $\check{L}$ corresponding to the four hypersurfaces Y_I ending in 0 and 1, respectively, together with S_1, S_2, and R_1. Similarly, $\underline{A}^0_{0'}$ is modeled on the expression (2.13) for $\check{A}^0_{0'}$ and includes terms counting monopoles on cobordisms with four boundary components. By the same reasoning as in the proof of Lemma 2.6.3, the analogue of Lemma 2.4.4 goes through essentially unchanged (although with nearly twice

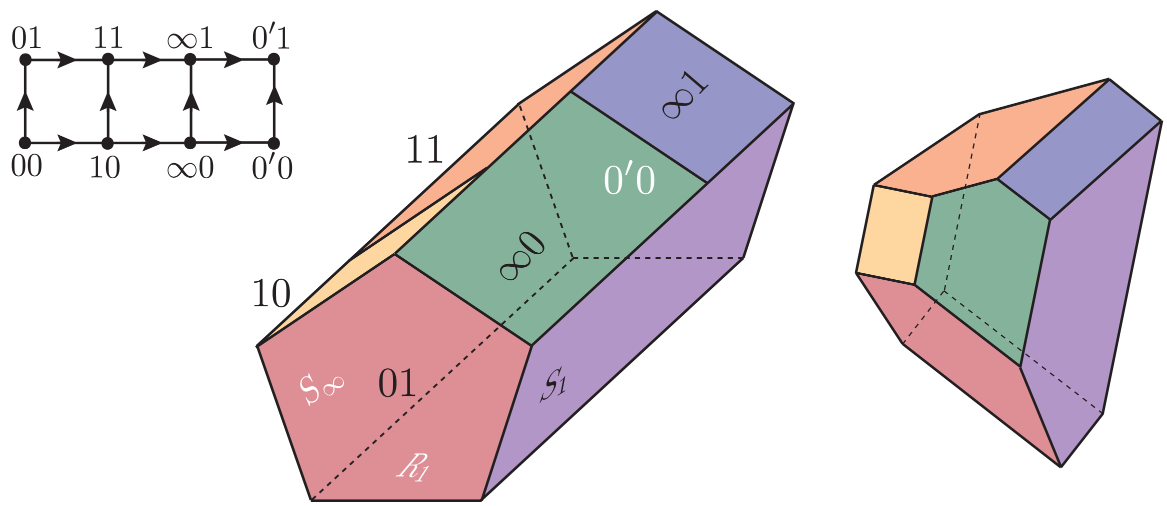

Figure 2.13: The lattice $\{0, 1, \infty, 0'\} \times \{0, 1\}$ corresponds to Stasheff's polytope, drawn at center so that the depth coordinate is suggestive of translating the sphere through W. Recall that the same polytope is associated to the lattice $\{0, 1, \infty\} \times \{0, 1, \infty\}$ in Figure 2.4, redrawn at right. In fact, by Theorem 2.3.3, an associahedron arises whenever the lattice is a product of at most two chains. The map corresponding to K_5 above is a null-homotopy for the sum of the maps associated to the faces. Those associated to S_1 and S_2 vanish as before, leaving a nine-term identity on the chain level.

as many terms), leading to the nine-term identity given by the lower left entry of

$$\begin{bmatrix} \check{D}^{0'0}_{0'0} & 0 \\ \underline{D}^{0'0}_{0'1} & \check{D}^{0'1}_{0'1} \end{bmatrix} \begin{bmatrix} \check{D}^{00}_{0'0} & 0 \\ \underline{D}^{00}_{0'1} & \check{D}^{01}_{0'1} \end{bmatrix} + \begin{bmatrix} \check{D}^{00}_{0'0} & 0 \\ \underline{D}^{00}_{0'1} & \check{D}^{01}_{0'1} \end{bmatrix} \begin{bmatrix} \check{D}^{00}_{00} & 0 \\ \underline{D}^{00}_{01} & \check{D}^{01}_{01} \end{bmatrix}$$
$$= \begin{bmatrix} \check{D}^{10}_{0'0} & 0 \\ \underline{D}^{10}_{0'1} & \check{D}^{11}_{0'1} \end{bmatrix} \begin{bmatrix} \check{D}^{00}_{10} & 0 \\ \underline{D}^{00}_{11} & \check{D}^{01}_{11} \end{bmatrix} + \begin{bmatrix} \check{D}^{\infty 0}_{0'0} & 0 \\ \underline{D}^{\infty 0}_{0'1} & \check{D}^{\infty 1}_{0'1} \end{bmatrix} \begin{bmatrix} \check{D}^{00}_{\infty 0} & 0 \\ \underline{D}^{00}_{\infty 1} & \check{D}^{01}_{\infty 1} \end{bmatrix} + \begin{bmatrix} \check{L}^{00}_{0'0} & 0 \\ \underline{L}^{00}_{0'1} & \check{L}^{01}_{0'1} \end{bmatrix}$$

The upper left and lower right identities are precisely those given by Lemma 2.4.4. Rewriting this identity via (2.23), we have the $\widetilde{HM}_\bullet$ analog of (2.12):

$$\tilde{D}^{0'}_{0'}\tilde{D}^{0}_{0'} + \tilde{D}^{0}_{0'}\tilde{D}^{0}_{0} = \tilde{D}^{1}_{0'}\tilde{D}^{0}_{1} + \tilde{D}^{\infty}_{0'}\tilde{D}^{0}_{\infty} + \tilde{L}.$$

The final map $\tilde{L} : \tilde{C}(Y_0) \to \tilde{C}(Y_{0'})$ is a chain map by (2.25). Filtering the corresponding square complex

$$Z = \bigoplus_{I \in \{00, 0'0, 01, 0'1\}} \check{C}(Y_I)$$

by the second digit, and recalling that $\check{L}^{00}_{0'0}$ and $\check{L}^{01}_{0'1}$ are quasi-isomorphisms, we conclude that $H_*(Z) = 0$. Therefore, $\tilde{L}$ is a quasi-isomorphism as well. Now exactly the same algebraic arguments yield the surgery exact triangle, and more generally the full statement of the link surgery spectral sequence, for $\widetilde{HM}_\bullet$.

The grading results of Section 2.5 readily extend to $\widetilde{HM}_\bullet$ by viewing the underlying complex in terms of $\widetilde{HM}_\bullet$ as in (2.22) and (2.24). In this way, we may extend $\check{\delta}$ to a mod 2 grading on $\underline{\tilde{X}}$ using the same definition. Since $U_\dagger$ cuts down the dimension of moduli spaces by two, the maps $\underline{D}^I_J$ on $\underline{\tilde{X}}$ obey the same mod 2 grading shift formula as the maps $\check{D}^I_J$ on $\widetilde{X}$. In particular, Lemmas 2.5.2 and 2.5.3 still apply, and when $l = 0$, $\check{\delta}$ and $\mathrm{gr}^{(2)}$ coincide on $\widetilde{HM}_\bullet(Y)$.

Remark 2.6.6. We now explain how the extra data of a basepoint $y \in Y$ determines a group $\widetilde{HM}_\bullet(Y, y)$ rather than just an isomorphism class of group, and then discuss invariance of the spectral sequence.

Recall that a cobordism W from Y_0 to Y_1 is a compact, oriented four-manifold with boundary W, together with a diffeomorphism of oriented manifolds

$$\alpha = (\alpha_0 \coprod \alpha_1) : \partial W \to -Y_0 \coprod Y_1.$$

A path γ in W from $y_0 \in Y_0$ to $y_1 \in Y_1$ is a smooth embedding $\gamma : [0,1] \to W$ such that $\alpha_0(\gamma(0)) = y_0$ and $\alpha_1(\gamma(1)) = y_1$. We also require that γ maps the interior of $[0,1]$ to the interior of W. An isomorphism of pairs (W,γ) and (W',γ') is an orientation-preserving diffeomorphism $\phi : W \to W'$ such that $\alpha = \alpha'\phi$ and $\gamma' = \phi\gamma$. The isotopy extension theorem implies that if γ and γ' are paths in W which are isotopic relative to their common endpoints, then (W,γ) is isomorphic to (W,γ'). In particular, the isomorphism class of (W,γ) is independent of the parameterization of γ, and the composition of isomorphism classes of pairs is well-defined[2].

We define a category COB_b, regarded as the based version of COB. An object (Y,y) of COB_b consists of a compact, connected, oriented 3-manifold Y together with a basepoint $y \in Y$. A morphism from (Y_0,y_0) to (Y_1,y_1) is an isomorphism class of pair (W,γ), consisting of a connected cobordism $W : Y_0 \to Y_1$ and a path γ from y_0 to y_1. The identity morphism on (Y,y) is represented by the cylindrical pair $(Y \times [0,1], y \times [0,1])$.

We will define the functor

$$\widetilde{HM}_\bullet : \mathrm{COB}_b \to \mathrm{GROUP} \tag{2.26}$$

in stages. Recall that the chain map $\tilde{m}(W)$ associated to a cobordism $W : Y_0 \to Y_1$ depends on a choice of ball in a neighborhood of a path γ across W, and a family of metrics and perturbations on $W - D^4$ parameterized by an interval (we think of this family as sliding the ball across the cobordism along γ). We constructed such a family by embedding an extra copy of each boundary component of W, parallel to the boundary except for a “finger” of each copy which is routed into the interior of W to “capture” the ball. These fingers are guided by the path γ from a point $y_0 \in Y \times \{0\}$ to a point $y_1 \in Y \times \{1\}$. If we define a second chain map using the same path γ in W as a guide, then we can define a chain homotopy between the first and second chain maps using a family of metrics and perturbations on $W - D^4$ parameterized by a hexagon (this is an amalgamation of the approaches in Figures 2.9 and 2.10 and Figures 2.11 and 2.12).

[2]To construct the composition of pairs, we must fix collar neighborhoods at the interface to extend the smooth structure across, and then smooth the kink in the composite path. So it is only the isomorphism class of the composition that is well-defined, just as in the category COB.

In other words, suppose we are given points $y_i \in Y_i$ and metric and perturbation data on $(Y_i \times [0,1]) - D_i^4$ for some choices of balls around the $y_i \times \{\frac{1}{2}\}$. Then a cobordism $W : Y_0 \to Y_1$ equipped with a path $\gamma : y_0 \to y_1$ induces a chain map up to chain homotopy, and thus a canonical map on homology. Furthermore, given an isomorphic pair (W', γ'), we can use the diffeomorphism ϕ to push families of metrics and perturbations on $W - D^4$ to such families on $W' - \phi(D^4)$. With these choices, the pairs (W, γ) and (W', γ') induce identical chain maps.

These observations are sufficient to apply the same trick used to obtain actual groups for the other versions of monopole Floer homology; we outline the approach here and refer the reader to Section 23.1 of [24] for details. Let $\widetilde{\mathrm{COB}}_b$ be the category in which an object is a quadruple (Y, y, D^4, g, q) where D^4 is ball around y, and g and q are a metric[3] and admissible perturbation on $(Y \times [0,1]) - D^4$. A morphism in $\widetilde{\mathrm{COB}}_b$ is a COB_b-morphism between the underlying based 3-manifolds. As explained in the previous paragraph, we have a functor

$$\widetilde{HM}_\bullet : \widetilde{\mathrm{COB}}_b \to \mathrm{GROUP}. \tag{2.27}$$

If (g, q, D^4) and $(g', q', (D^4)')$ are two sets of data over same pair (Y, y), then the objects (Y, y, D^4, g, q) and $(Y, y, (D^4)', g', q')$ are canonically isomorphic in $\widetilde{\mathrm{COB}}_b$ via the isomorphism $(Y \times [0,1], y \times [0,1])$. Thus, the groups $\widetilde{HM}_\bullet(Y, y, D^4, g, q)$ and $\widetilde{HM}_\bullet(Y, y, (D^4)', g', q')$ are canonically isomorphic. We thereby obtain a functor from COB_b to the category of "groups up to canonical isomorphism". The later admits a functor to GROUP by taking "cross-sections", and composition of functors yields that of (2.26).

We now consider the invariance of the link surgery spectral sequence associated to a framed link $L \subset Y$. The proof of Theorem 2.5.4 in the $\widetilde{HM}_\bullet$ case readily adapts to a version for $\widetilde{HM}_\bullet$ to give a filtered chain map $\underline{\phi}$ inducing a filtered, chain homotopy equivalence between total complexes $\underline{X}$ and $\underline{X}'$ defined using different sets of auxiliary data, including

[3] As before, the metric is required to be cylindrical near the Y_i ends and standard on the S^3 end. More systematically, we could consider the septuple $(Y, y, D^4, g_Y, q_Y, g, q)$ where D^4 is neighborhood of y, g_Y is a metric on Y and g is a metric on $(Y \times [0,1]) - D^4$ extending g_Y, etc. This inductive approach is preferable when we want the flexibility of a family of initial metrics, such as when defining a chain map between $\underline{X}$ and $\underline{X}'$ defined using different analytic data. See Remark 2.5.6.

basepoints y and y'. This chain map is defined using the cobordism W with 2^{l+2} embedded hypersurfaces as guides (that is, two copies of each hypersurfaces in Figure 2.12). If y and y' are distinct, then to specify $\underline{\phi}$ we must also choose a path η between the basepoints in the complement of L in Y, as a tubular neighborhood of $\eta \times [0,1] \subset W$ provides a common tubular neighborhoods of the paths $\{y\} \times [0,1]$ and $\{y'\} \times [0,1]$. We do not know if chain maps defined using different paths η are chain homotopic. However, if y and y' coincide, then we may pin down $\underline{\phi}$ up to filtered chain homotopy by requiring that η be the constant path. The filtered chain homotopy between chain maps is built using the same approach as in the $\widetilde{HM}_\bullet$ case. In this way, from a framed link $L \subset Y$ and a basepoint y in the complement of L, we obtain an $\widetilde{HM}_\bullet$ spectral sequence whose pages $\underline{E}^i$ for $i \geq 1$ are actual groups.

Chapter 3

Realizations of graph associahedra

The purpose of this brief chapter is to unify, under a general theorem, several earlier observations of inductive realizations of graph associahedra.

Recall the realization of Stasheff's polytope K_5 as a refinement of $K_4 \times [0,1]$ at center in Figure 2.13, and the inductive realization of P_{l+1} as a refinement of $P_l \times [0,1]$ in the proof of Theorem 2.5.4. Both of these are motivated by the "sliding-the-point" proof of the naturality of the $U_\dagger$ action in Floer theory To see why, recall that to any product lattice Λ we may associate a map $\check{D}$ whose longest component counts monopoles on W over a family of metrics parameterized by a polytope P_G (see Remark 2.3.4). We then expect the longest component of the homotopy which expresses the naturality of the $U_\dagger$ action with respect to $\check{D}$ to count monopoles over a family of metrics parameterized by $P_G \times [0,1]$, where the latter coordinate slides the point through W (see Figure 3.1).

These considerations led us to the following general theorem.

Theorem 3.0.7. *Let G and G' be the graphs associated to lattices Λ and $\Lambda \times \{0,1\}$, respectively. Given a realization of P_G, the graph associahedron $P_{G'}$ may be realized as a refinement of $P_G \times [0,n]$. Namely, for each internal vertex I of Λ, we refine $P_G \times [0,n]$ by adding the closure of the corresponding facet of $P_G \times \{w(I)\} \subset P_G \times [0,1]$.*

Intuitively, we do not add ridges to $P_G \times [0,n]$ for the auxiliary hypersurfaces because the point never passes through them. This construction is illustrated in Figure 3.2. In the center and righthand realizations, the construction is applied twice, starting from a graph

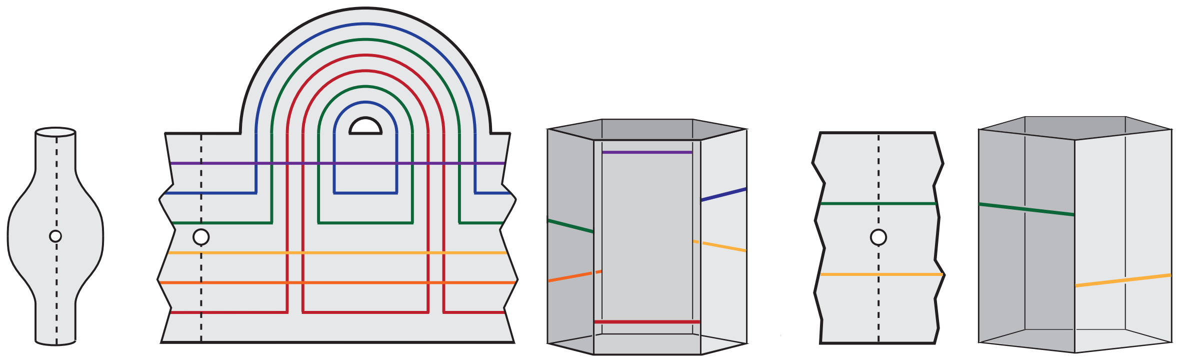

Figure 3.1: At left, we slide the point (or sphere) through the full cobordism from Figures 2.1 and 2.2. Each time the point crosses an internal hypersurface, we add a ridge to the corresponding lateral facet of $P_3 \times [0,1]$. Once the point has completed its journey, we have a realization of P_4. At right, we similarly slide the point through the full cobordism from Figure 2.8, adding ridges to $K_4 \times [0,1]$ corresponding to the two internal hypersurfaces. Once the point has completed its journey, we have K_5.

associahedron which is geometrically an interval.

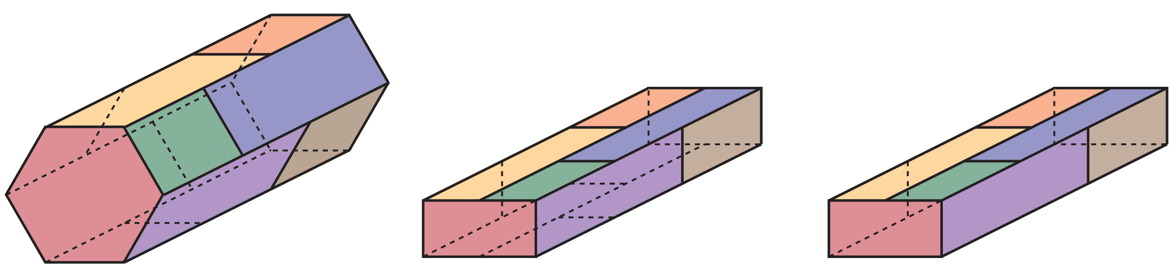

$\{0,1\}^3 \times \{0,1\}$ $\quad$ $\{0,1\}^2 \times \{0,1\} \times \{0,1\}$ $\quad$ $\{0,1,\infty\} \times \{0,1\} \times \{0,1\}$

Figure 3.2: Two alternative realizations of the permutohedron from Figure 2.3, and one of the graph associahedron from Figures 2.5 and 2.6.

We recall the relationship between G and G' as determined by Theorem 2.3.3 and illustrated in Figure 3.3 below. For $\Lambda = \{0, ..., n_1\} \times \cdots \times \{0, ..., n_l\}$, the graph G is the clique on vertices $v_1, \ldots, v_l$ with paths of length $n_1 - 1, \ldots, n_l - 1$ attached. The graph G' is constructed from G by adding a vertex v and connecting it to each of the v_i to form a clique of size $l+1$. Note that internal vertices of Λ correspond to tubes of G which intersect the clique non-trivially.

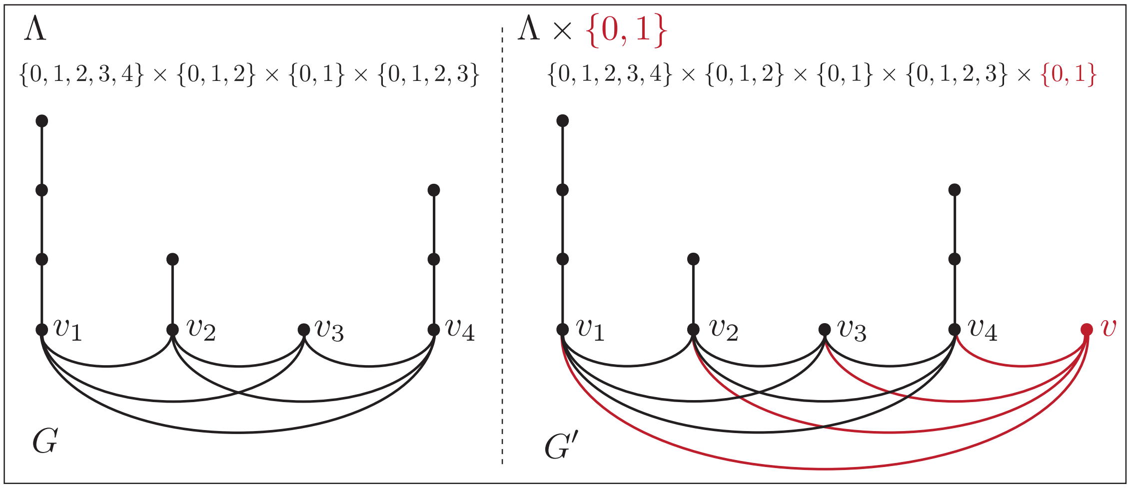

Figure 3.3: The graphs G and G' corresponding to lattices Λ and $\Lambda \times \{0,1\}$, respectively. The additional vertex and edges of G' are in red.

Before proving Theorem 3.0.7, we illustrate two special families of examples. Note that the lattices $\{0,1,\ldots,n\}$ and $\{0,1,\ldots,n-1\} \times \{0,1\}$ both correspond to the graph consisting of a path of length $n-1$. We therefore obtain a realization of the associahedron K_{n+2} by adding ridges[1] to $K_{n+1} \times [0,n]$. Similarly, since $\{0,1\}^n = \{0,1\}^{n-1} \times \{0,1\}$, we obtain a realization of the permutohedron P_{n+1} by adding ridges to all lateral facets of $P_n \times [0,n]$. If we build these realizations inductively, starting from $K_2 = P_1 = \{0\}$, then each is naturally a refinement of the hypercube (see Figure 3.4). We may arrange that P_{n+1} refines K_{n+2} as well. Upon sharing these inductive realizations of K_{n+2} and P_{n+1} with experts, we learned of their discovery a decade earlier[2] [39]. Theorem 3.0.7 may be viewed as a common generalization, motivated by vastly different considerations.

We now turn to its proof. Label the vertices in the original l-clique of G by $v_1,\ldots,v_l$. Let n denote the number of vertices in G. Recalling the graph associahedron construction in Section 2.3, we will need some additional terminology. A tube t in G is *internal* if it contains at least one vertex v_i, or equivalently, if it corresponds to an internal vertex I of Λ.

[1]By the proof of Theorem 3.0.7, one should add ridges to those lateral facets of $K_{n+1} \times [0,n]$ corresponding to tubes in the path of length $n-2$ which include the initial vertex (given by one of the two ends).

[2]For even earlier realizations of K_{n+2} as a convex hull, see [29], which is generalized in [15].

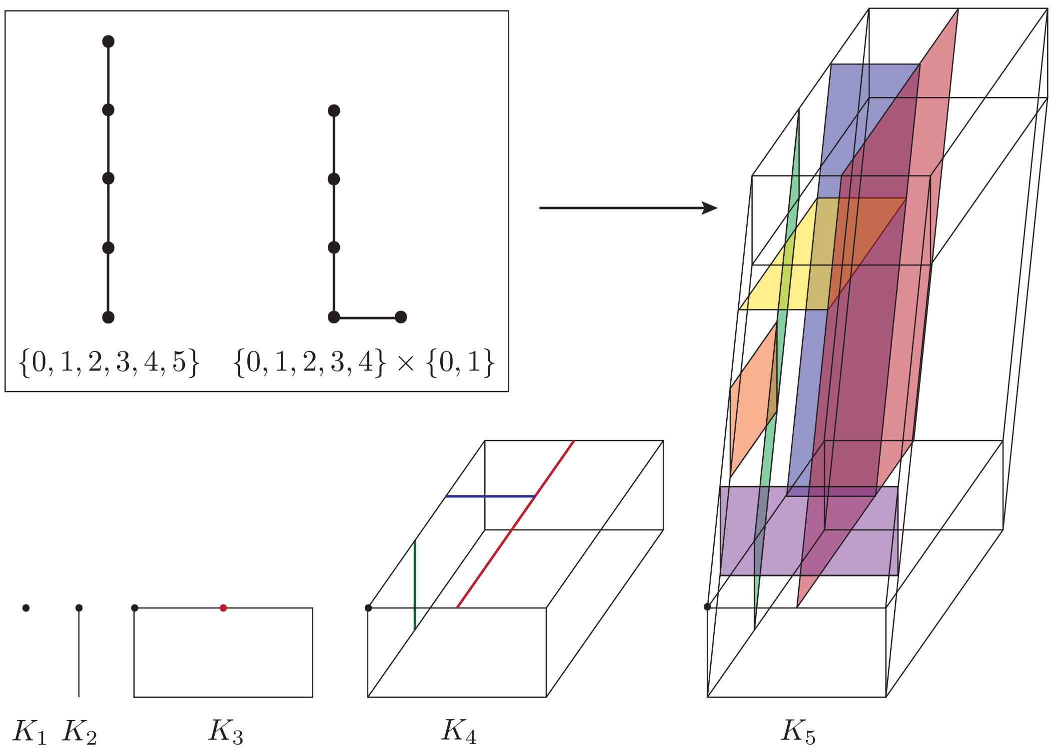

Figure 3.4: We realize K_{n+2} by adding $\binom{n}{2}$ ridges with integral vertices to the hypercube $[0,1] \times [0,2] \times \cdots \times [0,n]$. Similarly, P_{n+1} is obtained by adding $2^{n+1} - 2(n+1)$ ridges to the hypercube, as shown in the center realization of Figure 3.2 for $n = 3$.

The *weight* of an internal tube t is the number of vertices it contains, or equivalently, the weight of the corresponding internal vertex. Note that if t_1 and t_2 are distinct internal tubes in a tubing T, then they must be nested (as they cannot be non-adjacent), and thus have distinct weights. Given a tubing T with k internal tubes, we will always index the internal tubes t_i in order of increasing weight w_i. Let $[0,n]_T$ denote the partition of the interval $[0,n]$ into $k+1$ edges by adding a vertex at each integer w_i. Recall that the face poset of P_G corresponds to the poset of tubings of G. Let F_T denote the face of P_G corresponding to the tubing T. Let $\mathcal{T}_G$ denote the set of tubings of G.

Theorem 3.0.7 may now be restated as the assertion that P_G is realized by the polytope

$$P = \bigcup_{T \in \mathcal{T}_G} F_T \times [0,n]_T.$$

Proof. For a tubing $T \in \mathcal{T}_G$ with k internal tubes, the partitioned interval $[0,n]_T$ consists of $k+2$ vertices labeled w_i and $k+1$ edges labeled (w_i, w_{i+1}), where we set $w_0 = 0$ and $w_{k+1} = n$. We write $f \in [0,n]_T$ to denote that f is a face of $[0,n]_T$. Tracing through the definitions, the face poset of P is isomorphic to the poset

$$\mathcal{S}_G = \{(T,f) \,|\, T \in \mathcal{T}_G, f \in [0,n]_T\}$$

with $(T_1, f_1) \leq (T_2, f_2)$ if and only if $T_1 \leq T_2$ and $f_1 \subset \overline{f_2}$ as subsets of $[0,n]$. Here $\overline{f_2}$ denotes the closure of f_2, and the poset structure on tubings is by reverse inclusion. On the other hand, the face poset of $P_{G'}$ is isomorphic to $\mathcal{T}_{G'}$.

We define a map $\sigma : \mathcal{S}_G \to \mathcal{T}_{G'}$ as follows:

(1) $\sigma(T,0) = T \cup \{\{v\}\}$.

(2) $\sigma(T,w_i) = T \cup \{t_i \cup \{v\}\}$, where $t_i \in T$ has weight w_i.

(3) $\sigma(T,n) = T \cup \{V\}$.

(4) $\sigma(T,(w_i,w_{i+1})) = (T - \{t_j \,|\, j \geq i+1\}) \cup \{t_j \cup \{v\} \,|\, j \geq i+1\}$.

We will show that σ is an isomorphism of posets. The inverse map $\tau : \mathcal{T}_{G'} \to \mathcal{S}_G$ is defined as follows. Regard a tubing $T' \in \mathcal{T}_{G'}$ as a multi-set of sets of vertices of G'. Delete the vertex v from each set in T' where it appears to obtain a new multi-set $\mathbf{T}$ of sets of vertices of G. The structure of G' forces $\mathbf{T}$ to satisfy *exactly one*[3] of the following four properties:

(1) $\emptyset \in \mathbf{T}$. This occurs when $\{v\} \in T'$.

(2) $\{t_i, t_i\} \subset \mathbf{T}$ for some (unique) internal tube t_i. This occurs when $t_i, t_i \cup \{v\} \in T'$.

(3) $\{V\} \in \mathbf{T}$, where V is the set of all vertices of G. This occurs when $V \in T'$.

(4) $\mathbf{T}$ is a valid tubing of G. In this case, let $i+1$ be the smallest index such that t_{i+1} contains v. If there is no such tube, let $i+1 = n$.

In each case, we define $\tau(T')$ as follows:

(1) $\tau(T') = (\mathbf{T} - \{\emptyset\}, 0)$.

(2) $\tau(T') = (\mathbf{T} - \{t_i\}, w_i)$.

(3) $\tau(T') = (\mathbf{T} - \{V\}, n)$.

(4) $\tau(T') = (\mathbf{T}, (w_i, w_{i+1}))$.

It is immediate that σ and τ are inverse maps.

Finally, we must verify that σ is order preserving: if $(T_1, f_1) \leq (T_2, f_2)$ then $T_1' \leq T_2'$, where $T_1' = \sigma(T_1, f_1)$ and $T_2' = \sigma(T_2, f_2)$. Since the poset structure on tubings is by reverse inclusion, this is equivalent to the following:

Lemma 3.0.8. *If $T_1 \supset T_2$ and $\overline{f_1} \subset \overline{f_2}$ then $T_1' \supset T_2'$.*

Proof. If f_2 is a vertex then $f_1 = f_2$, and from the definition of σ we see that $T_1 \supset T_2$ implies $T_1' \supset T_2'$. We are left to consider the case that f_2 is an edge:

Let $s_1, \ldots, s_l$ denote the internal tubes of T_1 with weights $u_1 < \cdots < u_l$. Let $t_1, \ldots, t_k$ denote the internal tubes of T_2, with weights $w_1 < \cdots < w_k$. As before, set $u_0 = w_0 = 0$ and $u_{l+1} = w_{k+1} = n$. Now $f_2 = (w_i, w_{i+1})$ for some i, and $T_1 \supset T_2$ implies that $u_{i'} = w_i$ and $u_{i''} = w_{i+1}$ for some $i \leq i' < i''$. Furthermore, $f_1 \subset \overline{f_2} = [w_i, w_{i+1}]$.

[3] Using the structure of G' in Figure 3.3, the reader may verify that these four properties are mutually exclusive, and that if none of the first three properties hold, then $\mathbf{T}$ is a valid tubing of G.

We first consider the case $f_1 = u_m$ with $u_m \neq 0, n$. Then

$$T_1' = T_1 \cup \{s_m \cup \{v\}\}$$
$$T_2' = (T_2 - \{t_j \,|\, j \geq i+1\}) \cup \{t_j \cup \{v\} \,|\, j \geq i+1\}$$

Since T_1' is a valid tubing, we know that s_j contains the vertex v for all $j \geq m+1$, and in particular for all $j \geq i''+1$. Since $T_1 \supset T_2$, we have $t_j \cup \{v\} \in T_1'$ for all $j \geq i+1$, and thus $T_1' \supset T_2'$. The argument is similar when $u_m = 0$ or $u_m = n$.

The remaining case is $f_1 = (u_m, u_{m+1})$ for some $i' \leq m < i''$. Then

$$T_1' = (T_1 - \{s_{j'} \,|\, j' \geq m+1\}) \cup \{s_{j'} \cup \{v\} \,|\, j' \geq m+1\}$$
$$T_2' = (T_2 - \{t_j \,|\, j \geq i+1\}) \cup \{t_j \cup \{v\} \,|\, j \geq i+1\}$$

Since $T_1 \supset T_2$, each t_j with $j \geq i+1$ corresponds to some $s_{j'}$ with $j' \geq i'' \geq m+1$, so $t_j \cup \{v\} \in T_2'$. On the other hand, for each j' such that $i'' > j' \geq m+1$, we have $u_{j'} \neq w_j$ for any j, so $s_{j'} \notin T_2$. Thus $T_1' \supset T_2'$. $\square$

We conclude that σ yields an isomorphism between the face poset of P and the face poset of $P_{G'}$. So P is indeed a realization of $P_{G'}$, proving the theorem. $\square$

Chapter 4

Odd Khovanov homology and Conway mutation

In the chapter alone, chain complexes and homology groups are presumed to have $\mathbb{Z}$ coefficients unless otherwise specified.

To an oriented link $L \subset S^3$, Khovanov associates a bigraded homology group $Kh(L)$ whose graded Euler characteristic is the unnormalized Jones polynomial [22]. This invariant also has a reduced version $\overline{Kh}(L, K)$, which depends on a choice of marked component K. While the Jones polynomial itself is insensitive to Conway mutation, Khovanov homology generally detects mutations that swap strands between link components [49]. Whether the theory is invariant under component-preserving mutation, and in particular for knots, remains an interesting open question, explored in [7], [23], and [48]. No counterexamples exist with fewer than 14 crossings, although Khovanov homology does distinguish knots related by genus 2 mutation [17], whereas the (colored) Jones polynomial does not.

In 2003, Ozsváth and Szabó introduced a link surgery spectral sequence whose E^2 term is $\overline{Kh}(L; \mathbb{F}_2)$ and which converges to $\widehat{HF}(-\Sigma(L))$, the Heegaard Floer homology of the branched double-cover with reversed orientation [36]. In search of a candidate for the E^2 page over $\mathbb{Z}$, Ozsváth, Rasmussen and Szabó developed odd Khovanov homology $Kh'(L)$, a theory whose mod 2 reduction coincides with that of Khovanov homology [34]. While the reduced version $\widetilde{Kh}(L)$ also categorifies the Jones polynomial, it is independent of the choice of marked component and determines $Kh'(L)$ according to the equation

$$Kh'_{t,q}(L) \cong \widetilde{Kh}_{t,q-1}(L) \oplus \widetilde{Kh}_{t,q+1}(L). \tag{4.1}$$

In contrast to Khovanov homology for links, we prove:

Theorem 4.0.9. *Odd Khovanov homology is mutation invariant. Indeed, connected mutant link diagrams give rise to isomorphic odd Khovanov complexes.*

Corollary 4.0.10. *Khovanov homology over $\mathbb{F}_2$ is mutation invariant.*

It is not known if these results extend to genus 2 mutation [17]. Wehrli announced a proof of Corollary 4.0.10 for component-preserving Conway mutation in 2007, using an approach outlined by Bar-Natan in 2005 [7]. Shortly after our paper first appeared, Wehrli posted his proof, which is completely independent and extends to the case of Lee homology over $\mathbb{F}_2$ [48].

Mutant links in S^3 have homeomorphic branched double-covers. It follows that the E^∞ page of the $\widehat{HF}$ link surgery spectral sequence is also mutation invariant. Building on work of Roberts [37], Baldwin has shown that all pages E^i with $i \geq 2$ are link invariants, as graded vector spaces [5]. To this, we add:

Theorem 4.0.11. *The E^i page of the $\widehat{HF}$ link surgery spectral sequence is mutation invariant for $i \geq 2$. Indeed, connected mutant link diagrams give rise to isomorphic filtered complexes.*

Our proof will apply equally well to the monopole version of the spectral sequence developed in Chapter 6. Note that Khovanov homology, even over $\mathbb{F}_2$, is not an invariant of the branched double-cover itself [46].

This chapter is organized as follows. In Section 4.1, to a connected, decorated link diagram $\mathcal{D}$ we associate a set of numerical data that is determined by (and, in fact, determines) the equivalence class of $\mathcal{D}$ modulo diagrammatic mutation (and planar isotopy). From this data alone, we construct a complex $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$ which is *a priori* invariant under mutation of $\mathcal{D}$. In Section 4.2, we recall the construction of odd Khovanov homology and prove:

Proposition 4.0.12. *The complex $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$ is canonically isomorphic to the reduced odd Khovanov complex $(\overline{C}(\mathcal{D}), \bar{\partial}_\epsilon)$.*

This establishes Theorem 4.0.9 and verifies that our construction leads to a well-defined (in fact, previously-defined) link invariant.

In Section 4.3, we consider a surgery diagram for the branched double-cover of $\mathcal{D}$ given by a framed link $\mathbb{L} \subset S^3$ with one component for each crossing. Ozsváth and Szabó associate a filtered complex to $\mathcal{D}$ by applying the Heegaard Floer functor to a hypercube of 3-manifolds and cobordisms associated to various surgeries on $\mathbb{L}$ (see [36] for details). The link surgery spectral sequence is then induced by standard homological algebra. In Proposition 4.3.1, we prove that the framed isotopy type of $\mathbb{L}$ is determined by the mutation equivalence class of $\mathcal{D}$, establishing Theorem 4.0.11. Note that Corollary 4.0.10 may be viewed as the E^2 page of Theorem 4.0.11.

We conclude Section 4.3 with a remark on our original motivation for the construction of the complex $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$. An essential observation in [36] is that one recovers the reduced Khovanov complex over $\mathbb{Z}/2\mathbb{Z}$ by first branched double-covering the Khovanov hypercube of 1-manifolds and 2-dimensional cobordisms and then applying the Heegaard Floer TQFT to the resulting hypercube of 3-manifolds and 4-dimensional cobordisms. As we establish starting in Chapter 6, a similar relationship holds between reduced Khovanov homology and the monopole Floer TQFT. From this perspective, our results may be viewed as immediate consequences of the topological fact that branched double-covering destroys all evidence of mutation.

4.1 A thriftier construction of reduced odd Khovanov homology

Given an oriented link L, fix a connected, oriented link diagram $\mathcal{D}$ with crossings $c_1, \ldots, c_n$. Let n_+ and n_- be the number of positive and negative crossings, respectively. We use $\mathcal{V}(\mathcal{D})$ and $\mathcal{E}(\mathcal{D})$ to denote the sets of vertices and edges, respectively, of the hypercube $\{0,1\}^n$, with edges oriented in the direction of increasing weight. Decorate each crossing c_i with an arrow x_i, which may point in one of two parallel directions and appears in each complete resolution $\mathcal{D}(I)$ as an oriented arc between circles according to the conventions in Figure 4.1.

Recall that a planar link diagram $\mathcal{D}$ admits a checkerboard coloring with white exterior, as illustrated at left in Figure 4.3. The black graph $\mathcal{B}(\mathcal{D})$ is formed by placing a vertex in each black region and drawing an edge through each crossing. The edge through c_i connects the vertices in the black regions incident to c_i. Given a spanning tree $\mathcal{T} \subset \mathcal{B}(\mathcal{D})$, we may form a resolution of $\mathcal{D}$ consisting of only one circle by merging precisely those black regions which are incident along $\mathcal{T}$. In particular, all connected diagrams admit at least one such resolution.

With these preliminaries in place, we now give a recipe for associating a bigraded chain complex $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$ to the decorated diagram $\mathcal{D}$. The key idea is simple. Think of each resolution of $\mathcal{D}$ as a connected, directed graph whose vertices are the circles and whose edges are the oriented arcs. While the circles merge and split from one resolution to the next, the arcs are canonically identified throughout. So we use the exact sequence

$$\{\text{cycles}\} \hookrightarrow \mathbb{Z}\langle \text{arcs} \rangle \xrightarrow{d} \mathbb{Z}\langle \text{circles} \rangle \to \mathbb{Z} \to 0$$

of free Abelian groups to suppress the circles entirely and instead keep track of the cycles in each resolution, thought of as relations between the arcs themselves.

We begin the construction by fixing a vertex $I^* = (m_1^*, \ldots, m_n^*) \in \mathcal{V}(\mathcal{D})$ such that the resolution $\mathcal{D}(I^*)$ consists of only one circle S. To each pair of oriented arcs (x_i, x_j) in $\mathcal{D}(I^*)$ we associate a linking number $a_{ij} \in \{0, \pm 1\}$ according to the symmetric convention in Figure 4.2. We set $a_{ii} = 0$. Note that arcs on the same side of S cannot link. For each $I \in \mathcal{V}(\mathcal{D})$, we have an Abelian group

$$\widetilde{V}(\mathcal{D}(I)) = \mathbb{Z}\langle x_1, \ldots, x_n \,|\, r_1^I, \ldots, r_n^I \rangle$$

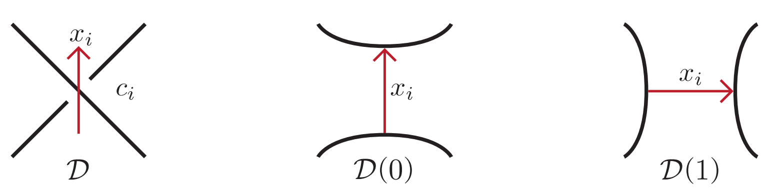

Figure 4.1: **Oriented resolution conventions.** The arrow x_i at crossing c_i remains fixed in a 0-resolution and rotates 90° clockwise in a 1-resolution. To see the other choice for the arrow at c_i, rotate the page 180°. Mutation invariance of the Jones polynomial follows from the rotational symmetries of the above tangles in $\mathbb{R}^3$.

presented by relations

$$r_i^I = \begin{cases} x_i - \displaystyle\sum_{\{j \mid m_j \neq m_j^*\}} (-1)^{m_j^*} a_{ij} x_j & \text{if} \quad m_i = m_i^* \\ -\displaystyle\sum_{\{j \mid m_j \neq m_j^*\}} (-1)^{m_j^*} a_{ij} x_j & \text{if} \quad m_i \neq m_i^*. \end{cases} \tag{4.2}$$

Indeed, these relations generate the cycles in the graph of circles and arcs at $\mathcal{D}(I)$ (see Lemma 4.2.1).

To an edge $e \in \mathcal{E}(\mathcal{D})$ from I to J given by an increase in resolution at c_i, we associate a map

$$\tilde{\partial}_J^I : \Lambda^* \widetilde{V}(\mathcal{D}(I)) \to \Lambda^* \widetilde{V}(\mathcal{D}(J))$$

of exterior algebras, which is defined in the case of a split and a merge, respectively, by

$$\tilde{\partial}_J^I(u) = \begin{cases} x_i \wedge u & \text{if} \quad x_i = 0 \in \widetilde{V}(\mathcal{D}(I)) \\ u & \text{if} \quad x_i \neq 0 \in \widetilde{V}(\mathcal{D}(I)). \end{cases}$$

Extending by zero, we may view each of these maps as an endomorphism of the group

$$\widetilde{C}(\mathcal{D}) = \bigoplus_{I \in \mathcal{V}(\mathcal{D})} \Lambda^* \widetilde{V}(\mathcal{D}(I)).$$

Consider a 2-dimensional face of the hypercube from I to J corresponding to an increase in resolution at c_i and c_j. The two corresponding composite maps in $\widetilde{C}(\mathcal{D})$ commute up to sign, and they vanish identically if and only if the arcs x_i and x_j in $\mathcal{D}(I)$ are in one of the two configurations in Figure 4.2, denoted Type X and Type Y. Note that we can distinguish a Type X face from a Type Y face without reference to the diagram by checking

Type $X = -1$ Type $Y = +1$

Figure 4.2: **Linking number conventions.** Two arcs are linked in $\mathcal{D}(I^*)$ if and only if their endpoints are interleaved on the circle. Otherwise, $a_{ij} = 0$. Any linked configuration is isotopic to one of the above on the 2-sphere $\mathbb{R}^2 \cup \{\infty\}$.

which of the relations $x_i \pm x_j = 0$ holds in $\widetilde{C}(\mathcal{D})$ over the two vertices of the face that are strictly between I and J.

A Type Y edge assignment on $\widetilde{C}(\mathcal{D})$ is a map $\epsilon : \mathcal{E}(\mathcal{D}) \to \{\pm 1\}$ such that the product of signs around a face of Type X or Type Y agrees with the sign of the linking convention in Figure 4.2 and such that, after multiplication by ϵ, every face of $\widetilde{C}(\mathcal{D})$ anticommutes. Such an assignment defines a differential $\tilde{\partial}_\epsilon : \widetilde{C}(\mathcal{D}) \to \widetilde{C}(\mathcal{D})$ by

$$\tilde{\partial}_\epsilon(v) = \sum_{\{e \in \mathcal{E}(\mathcal{D}), J \in \mathcal{V}(\mathcal{D}) \,|\, e \text{ goes from } I \text{ to } J\}} \epsilon(e) \cdot \tilde{\partial}^I_J(v)$$

for $v \in \Lambda^* \widetilde{V}(\mathcal{D}(I))$. Type Y edge assignments always exist and any two yield isomorphic complexes, as do any two choices for the initial arrows on $\mathcal{D}$. We can equip $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$ with a bigrading that descends to homology and is initialized using $n_\pm$ just as in [34]. The bigraded group $\widetilde{C}(\mathcal{D})$ and maps $\tilde{\partial}^I_J$ are constructed entirely from the numbers a_{ij}, m^*_i, and $n_\pm$. Thus, up to isomorphism:

Proposition 4.1.1. *The bigraded complex $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$ is determined by the linking matrix A of any one-circle resolution of $\mathcal{D}$, the vertex of this resolution, and the number of positive and negative crossings.*

Proof of Theorem 4.0.9. The following argument is illustrated in Figure 4.3. Given oriented, mutant links L and L', fix a corresponding pair of oriented, connected diagrams $\mathcal{D}$ and $\mathcal{D}'$ for which there is a circle C exhibiting the mutation. This circle crosses exactly two black regions of $\mathcal{D}$, which we connect by a path Γ in $\mathcal{B}(\mathcal{D})$. To simplify the exposition, we will assume there is a crossing between the two strands of $\mathcal{D}$ in C, so that Γ may be chosen in C. Extend Γ to a spanning tree $\mathcal{T}$ to obtain a resolution $\mathcal{D}(I^*)$ with one circle. The natural pairing of the crossings of $\mathcal{D}$ and $\mathcal{D}'$ induces an identification $\mathcal{V}(\mathcal{D}) \cong \mathcal{V}(\mathcal{D}')$. The resolution $\mathcal{D}'(I^*)$ may be obtained directly from $\mathcal{D}(I^*)$ by mutation and also consists of one circle S. We can partition the set of arcs inside C into those which go across S (in dark blue) and those which have both endpoints on the same side of S (in red). Since this division is preserved by mutation, the mod 2 linking matrix is preserved as well.

Arrows at the crossings of $\mathcal{D}$ orient the arcs of $\mathcal{D}(I^*)$, which in turn orient the arcs of $\mathcal{D}'(I^*)$ via the mutation. To preserve the linking matrix at I^* with sign, we modify the arcs of $\mathcal{D}'(I^*)$ as follows. Let $A = \{\text{arcs in } C \text{ and in } S\}$ and $B = \{\text{arcs in } C \text{ and not in } S\}$. We reverse those arcs of $\mathcal{D}'(I^*)$ that lie in A, B, or $A \cup B$, according to whether the mutation is about the z-, y-, or x-axis, respectively (as represented at right in Figure 4.3). We then select a corresponding set of arrows on $\mathcal{D}'$. Note that it may also be necessary to switch the orientations of both strands inside C so that $\mathcal{D}'$ will be consistently oriented. In any case, the number of positive and negative crossings is unchanged. Propositions 4.0.12 and 4.1.1 now imply the theorem for $(\overline{C}(\mathcal{D}), \bar{\partial}_\epsilon)$. The unreduced odd Khovanov complex is isomorphic to two copies of the reduced complex, just as in (4.1). □

Remark 4.1.2. We have seen that if connected diagrams $\mathcal{D}$ and $\mathcal{D}'$ are related by diagrammatic mutation (and planar isotopy), then there is an identification of their crossings and a vertex I^* such that $\mathcal{D}(I^*)$ and $\mathcal{D}'(I^*)$ have the same mod 2 linking matrix. Remarkably, the converse holds as well, i.e., I^* and A mod 2 together determine $\mathcal{D}$ up to diagrammatic mutation. This follows from a theorem of Chmutov and Lando: *Chord diagrams have the same intersection graph if and only if they are related by mutation* [14]. Here we view a one-circle resolution as a bipartite chord diagram, so that its mod 2 linking matrix is precisely the adjacency matrix of the corresponding intersection graph. Note that in the bipartite case, any combinatorial mutation as defined in [14] can be realized by a finite sequence of our diagrammatic ones.

Chmutov and Lando apply their result to the chord-diagram construction of finite type invariants. All finite type invariants of order ≤ 10 are insensitive to Conway mutation, whereas there exists an invariant of order 11 that distinguishes the knots in Figure 4.3 and one of order 7 that distinguishes genus 2 mutants (see [32] and [14]).

4.2 The original construction of reduced odd Khovanov homology

We now recall the original construction of reduced odd Khovanov homology, following [34]. Given an oriented link $L \subset S^3$, we fix a decorated, oriented diagram $\mathcal{D}$ as before, though now it need not be connected. For each vertex $I \in \mathcal{V}(\mathcal{D})$, the resolution $\mathcal{D}(I)$ consists of a

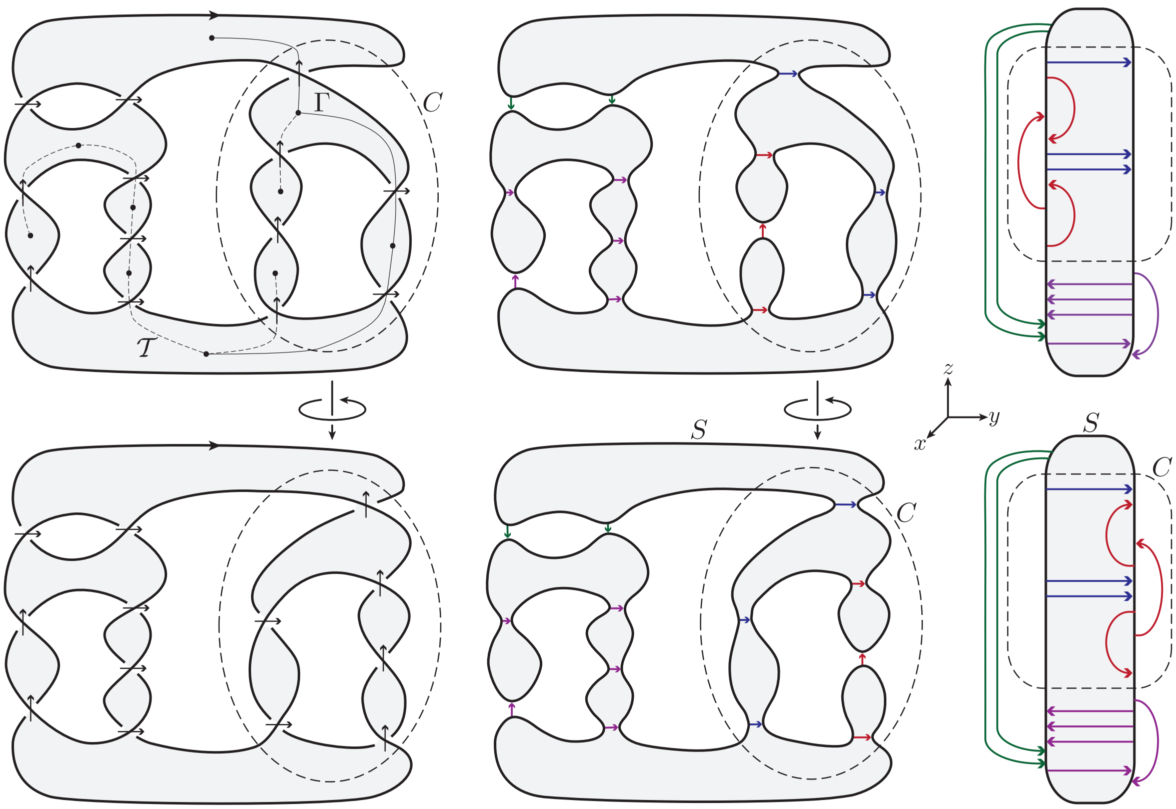

Figure 4.3: **The Kinoshita-Terasaka and Conway knots.** Orientations on the arcs in the upper-right resolution induce orientations on the arcs in the lower-right resolution via the mutation. In order to obtain the same linking data with sign, we have reversed the five arcs in the lower-right resolution that lie inside both S and C. We then work backwards to select arrows in the lower-left knot diagram.

set of circles $\{S_i^I\}$. Let $V(\mathcal{D}(I))$ be the free Abelian group generated by these circles. The reduced group $\overline{V}(\mathcal{D}(I))$ is defined to be the kernel of the augmentation $\eta : V(\mathcal{D}(I)) \to \mathbb{Z}$ given by $\sum a_i S_i^I \mapsto \sum a_i$.

Now let $\mathbb{Z}\langle x_1, \ldots, x_n\rangle$ denote the free Abelian group on n generators. For each $I \in \mathcal{V}(\mathcal{D})$, we have a boundary map

$$d^I : \mathbb{Z}\langle x_1, \ldots, x_n\rangle \to V(\mathcal{D}(I))$$

given by $d^I x_i = S_j^I - S_k^I$, where x_i is directed from S_j^I to S_k^I in $\mathcal{D}(I)$. Consider an edge $e \in \mathcal{E}(\mathcal{D})$ from I to J corresponding to a increase in resolution at c_i. If two circles merge as we move from $\mathcal{D}(I)$ to $\mathcal{D}(J)$, then the natural projection map $\{S_i^I\} \twoheadrightarrow \{S_i^J\}$ induces a morphism of exterior algebras. Alternatively, if a circle splits into two descendants, the two reasonable inclusion maps $\{S_i^I\} \hookrightarrow \{S_i^J\}$ induce equivalent morphisms on exterior algebras after wedging with the ordered difference of the descendents in $\mathcal{D}(J)$. In other words, we have a well-defined map

$$\bar{\partial}_J^I : \Lambda^* \overline{V}(\mathcal{D}(I)) \to \Lambda^* \overline{V}(\mathcal{D}(J))$$

given by

$$\bar{\partial}_J^I(v) = \begin{cases} d^J x_i \wedge v & \text{if} \quad d^I x_i = 0 \in \overline{V}(\mathcal{D}(I)) \\ v & \text{if} \quad d^I x_i \neq 0 \in \overline{V}(\mathcal{D}(I)) \end{cases}$$

in the case of a split and a merge, respectively, along x_i.

As in Section 4.1, we now form a group $\overline{C}(\mathcal{D})$ over the hypercube and choose a Type Y edge assignment to obtain a differential $\bar{\partial}_\epsilon : \overline{C}(\mathcal{D}) \to \overline{C}(\mathcal{D})$. The reduced odd Khovanov homology $\widetilde{Kh}(L) \cong H_*(\overline{C}(\mathcal{D}), \bar{\partial}_\epsilon)$ is independent of all choices and comes equipped with a bigrading that is initialized using $n_\pm$. The unreduced version is obtained by replacing $\overline{V}(\mathcal{D}(I))$ with $V(\mathcal{D}(I))$ above.

Proof of Proposition 4.0.12. Suppose that $\mathcal{D}$ is connected. Then for each $I \in \mathcal{V}(\mathcal{D})$, the image of $d^I : \mathbb{Z}\langle x_1, \ldots, x_n\rangle \to V(\mathcal{D}(I))$ is precisely $\overline{V}(\mathcal{D}(I))$. In fact, by Lemma 4.2.1 below, d^I induces an isomorphism $\widetilde{V}(\mathcal{D}(I)) \cong \overline{V}(\mathcal{D}(I))$. The collection of maps d^I therefore induce a group isomorphism $\widetilde{C}(\mathcal{D}) \cong \overline{C}(\mathcal{D})$ which is immediately seen to be equivariant with respect to the edge maps $\tilde{\partial}_J^I$ and $\bar{\partial}_J^I$. After fixing a common Type Y edge assignment, the proposition follows. $\square$

Lemma 4.2.1. *The relations r_i^I generate the kernel of the map $d^I : \mathbb{Z}\langle x_1, \ldots, x_n\rangle \to V(\mathcal{D}(I))$.*

Proof. To simplify notation, we assume that $m_i \neq m_i^*$ if and only if $i \leq k$, for some $1 \leq k \leq n$. Consider the $n \times n$ matrix M^I with column i given by the coefficients of $(-1)^{m_i^*} r_i^I$. Let A^I be the leading $k \times k$ minor, a symmetric matrix. We build an orientable surface F^I by attaching k 1-handles to the disk D^2 bounded by S so that the cores of the handles are given by the arcs $x_1, \ldots, x_k$ as they appear in $\mathcal{D}(I^*)$. We obtain a basis for $H_1(F^I)$ by extending each oriented arc to a loop using a chord through D^2. The cocores of the handles are precisely $x_1, \ldots, x_k$ as they appear in $\mathcal{D}(I)$, so these oriented arcs form a basis for $H_1(F^I, \partial F^I)$. With respect to these bases, the homology long exact sequence of the pair $(F^I, \partial F^I)$ includes the segment

$$H_1(F^I) \xrightarrow{A^I} H_1(F^I, \partial F^I) \xrightarrow{d^I|_{\mathbb{Z}\langle x_1,\ldots,x_k\rangle}} H_0(\partial F^I) \xrightarrow{\eta} \mathbb{Z} \to 0. \tag{4.3}$$

Furthermore, for each $i > k$, the oriented chord in D^2 between the endpoints of x_i is represented in $H_1(F^I, \partial F^I)$ by the first k entries in column i of M^I. We can therefore enlarge (4.3) to an exact sequence

$$\mathbb{Z}\langle x_1, \ldots, x_n\rangle \xrightarrow{M^I} \mathbb{Z}\langle x_1, \ldots, x_n\rangle \xrightarrow{d^I} V(\mathcal{D}(I)) \xrightarrow{\eta} \mathbb{Z} \to 0,$$

which implies the lemma. □

We can reduce the number of generators and relations in our construction by using the smaller presentation in (4.3). Namely, for each $I = (m_1, \ldots, m_n) \in \mathcal{V}(\mathcal{D})$, we let $\widehat{V}(\mathcal{D}(I))$ be the group generated by $\{x_j \,|\, m_j \neq m_j^*\}$ and presented by A^I. By (4.2), the edge map $\tilde{\partial}_J^I$ at c_i is replaced by

$$\hat{\partial}_J^I(u) = \begin{cases} x_i \wedge u & \text{if} \quad x_i = 0 \in \widehat{V}(\mathcal{D}(I)) \ \text{ and } \ m_i = m_i^* \\ -r_i^I \wedge u & \text{if} \quad x_i = 0 \in \widehat{V}(\mathcal{D}(I)) \ \text{ and } \ m_i \neq m_i^* \\ u & \text{if} \quad x_i \neq 0 \in \widehat{V}(\mathcal{D}(I)), \end{cases}$$

where it is understood that $x_i \mapsto 0$ when $m_i \neq m_i^*$. While the definition of $\hat{\partial}_J^I$ is more verbose, the presentations A^I are simply the 2^n principal minors of a single, symmetric matrix: $\{(-1)^{m_i^*+m_j^*} a_{ij}\}$. The resulting complex $(\widehat{C}(\mathcal{D}), \hat{\partial}_\epsilon)$ sits in between $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$ and $(\overline{C}(\mathcal{D}), \bar{\partial}_\epsilon)$ and is canonically isomorphic to both.

4.3 Branched double-covers, mutation, and link surgery

To a one-circle resolution $\mathcal{D}(I^*)$ of a connected diagram of a link $L \subset S^3$, we associate a framed link $\mathbb{L} \subset S^3$ that presents $-\Sigma(L)$ by surgery (see also [21]). Figure 4.4 illustrates the procedure starting from each resolution at right in Figure 4.3. We first cut open the circle S and stretch it out along the y-axis, dragging the arcs along for the ride. We then slice along the Seifert surface $\{x = 0, z < 0\}$ for S and pull the resulting two copies up to the xy-plane as though opening a book. This moves those arcs which started inside S to the orthogonal half-plane $\{z = 0, x > 0\}$, as illustrated in the second row. The double cover of S^3 branched over S is obtained by rotating a copy of the half-space $\{z \geq 0\}$ by 180° about the y-axis and gluing it back onto the upper half-space. The arcs x_i lift to circles $\mathbb{K}_i \subset S^3$, which comprise $\mathbb{L}$. We assign $\mathbb{K}_i$ the framing $(-1)^{m_i^*}$.

If $\mathcal{D}$ is decorated, then $\mathbb{L}$ may be oriented by the direction of each arc in the second row of Figure 4.4. The linking matrix $\mathbb{A}$ of $\mathbb{L}$ then coincides with A off the diagonal, with the diagonal itself encoding I^*. In fact, the geometric constraints on $\mathbb{L}$ are so severe that it is determined up to framed isotopy by $\mathbb{A}$. This seems to follow intuitively from hanging $\mathbb{L}$ on a wall, and is rigorously true by:

Proposition 4.3.1. *The isotopy type of $\mathbb{L} \subset S^3$ is determined by the intersection graph of $\mathcal{D}(I^*)$, whereas the framing of $\mathbb{L}$ is determined by I^*.*

Proof. Suppose that $\mathcal{D}(I^*)$ and $\mathcal{D}'(J^*)$, thought of as bipartite chord diagrams, have the same intersection graph. Then by [14], $\mathcal{D}(I^*)$ is connected to $\mathcal{D}'(J^*)$ by a sequence of mutations (see Remark 4.1.2). Each mutation corresponds to a component-preserving isotopy of $\mathbb{L}$ modeled on a half-integer translation of a torus $\mathbb{R}^2/\mathbb{Z}^2$ embedded in S^3 (see Figure 4.4). Therefore, the associated links $\mathbb{L}$ and $\mathbb{L}'$ are isotopic. The second statement is true by definition. $\square$

Proof of Theorem 4.0.11. From the construction of the spectral sequence in [36], it is clear that the filtered complex associated to a connected diagram $\mathcal{D}$ only depends on $\mathcal{D}$ through the framed isotopy type of the link $\mathbb{L}$ associated to a one-circle resolution. We conclude that each page (E^i, d^i) for $i \geq 1$ is fully determined by the mutation invariant data in Proposition 4.3.1 (for further details, see [5] and [37]). $\square$

We finally come to the surgery perspective that first motivated the construction of the complex $(\widetilde{C}(\mathcal{D}), \tilde{\partial}_\epsilon)$. For each $I = (m_1, \ldots, m_n) \in \mathcal{V}(\mathcal{D})$, let $\mathbb{L}_I$ consist of the same underlying link as $\mathbb{L}$, but with framing modified to

$$\lambda_i^I = \begin{cases} \infty & \text{if } m_i = m_i^* \\ 0 & \text{if } m_i \neq m_i^* \end{cases}$$

on $\mathbb{K}_i$. Then $\mathbb{L}_I$ is a surgery diagram for $\Sigma(\mathcal{D}(I)) \cong \#^{k_I} S^1 \times S^2$, where $\mathcal{D}(I)$ consists of $k_I + 1$ circles. The linking matrix of $\mathbb{L}_I$ then presents $H_1(\Sigma(\mathcal{D}(I)))$ with respect to fixed meridians $\{x_i \,|\, m_i \neq m_i^*\}$. By identifying $H_1(\Sigma(\mathcal{D}(I)))$ with $\widehat{V}(\mathcal{D}(I))$, we may construct $(\widehat{C}(\mathcal{D}), \hat{\partial}_\epsilon)$, and therefore $\widetilde{Kh}(L)$, completely on the level of branched double-covers. We elaborate on this perspective, and its relationship to the monopole Floer homology and Donaldson TQFT's in Chapter 6.

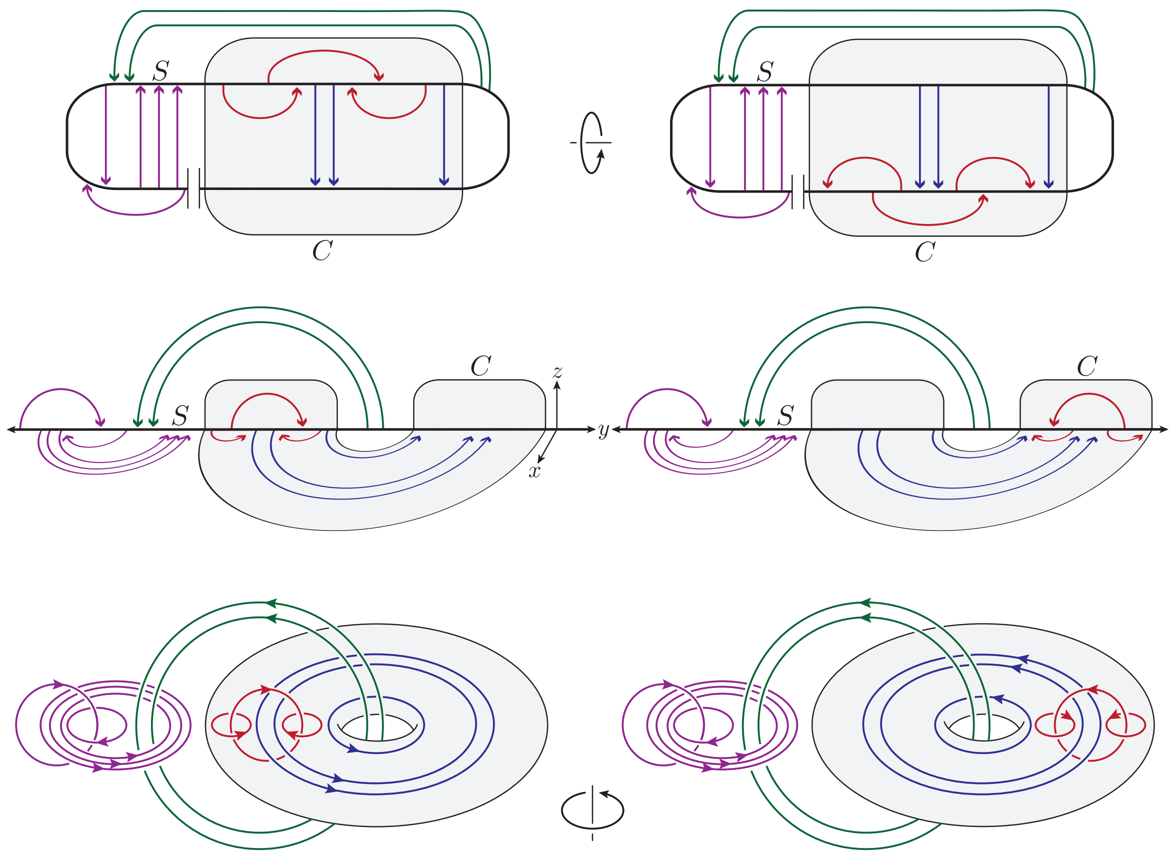

Figure 4.4: **Constructing a surgery diagram for the branched double-cover.** The resolutions in the first row are related by mutation along the Conway sphere formed by attaching disks to either side of C. The double cover of S^2 branched over its intersection with S is represented by each torus in the third row. Rotation of the torus about the z-axis yields a component-preserving isotopy from $\mathbb{L}$ to $\mathbb{L}'$.

Chapter 5

Link signature and homological width

Let $\mathcal{D}$ be a decorated, connected diagram of an oriented link $L \subset S^3$. In Section 4.3, given a one-circle resolution $\mathcal{D}(I^*)$, we constructed a surgery diagram $\mathbb{L} \subset S^3$ for $-\Sigma(L)$ with linking matrix $\mathbb{A}$. By Remark 4.1.2, this linking matrix contains sufficient information to recover any invariant of $\mathcal{D}$ which is unchanged by mutation. In particular:

Theorem 5.0.2. *The signature, determinant, and nullity of L are given by*

$$\sigma(L) = \sigma(\mathbb{A}) + w(I^*) - n_-$$
$$\det(L) = |\det(\mathbb{A})|$$
$$\nu(L) = \nu(\mathbb{A})$$

Proof. The last two equations follow from the fact that $\mathbb{A}$ presents $H_1(-\Sigma(L))$. The first equation follows from the 4-dimensional viewpoint of Gordon and Litherland [20]. Let $F \subset S^3$ be a spanning surface for L. Let L' be the push-off of L into F with parallel orientation. Finally, let $X = \Sigma(D^4, F)$ be the branched double-cover of F pushed into D^4 with $\partial F \subset \partial D^4$. The central result of [20] says that

$$\sigma(L) = -\sigma(X) + \frac{1}{2}\mathrm{lk}(L, L').$$

The linking number at right may be thought of as correcting for the fact that F may be non-orientable. For the combinatorial Gordon-Litherland signature formula, F is the black surface in the checkerboard coloring, and the intersection form of X is presented by removing a row and column from the Goeritz matrix of the black graph.

By constrast, for our signature formula, F is the spanning surface obtained from the disk bounded by $\mathcal{D}(I^*)$ by a attaching a twisted band along each chord. Based on a standard diagram of F, the correction term is given by

$$\frac{1}{2}\mathrm{lk}(L, L') = w(I^*) - n_-.$$

Furthermore, the branched double-cover X is obtained by attaching 2-handles along $\mathbb{L}$ and then reversing the orientation, so the intersection form if X is presented by $-\mathbb{A}$. □

A less direct proof of this signature formula, routed through the combinatorial Gordon-Litherland signature formula by comparing surgery diagrams, appears in Chapter 6.

Example 5.0.3. Consider the resolution $\mathcal{D}(010)$ of the right-handed trefoil $\mathcal{T}$ in Figure 6.2, with $\mathbb{A}$ given by the linking matrix at right. The signature formula gives

$$\sigma(\mathcal{T}) = \sigma(\mathbb{A}) + w(010) - n_-(\mathcal{D}) = 1 + 1 - 0 = 2.$$

For the mirrored diagram $\overline{\mathcal{D}}$ representing the left-handed trefoil $\overline{\mathcal{T}}$, consider the mirrored resolution $\mathcal{D}(101)$. Now the signature formula gives

$$\sigma(\overline{\mathcal{T}}) = \sigma(-\mathbb{A}) + w(101) - n_-(\overline{\mathcal{D}}) = -1 + 2 - 3 = -2.$$

This signature formula makes certain properties of link signature quite transparent:

Corollary 5.0.4. *Link signature is invariant under Conway mutation.*

Proof. Mutant diagrams have the same number of negative crossings, and give rise to the same matrix $\mathbb{A}$ at corresponding one-circle resolutions I^* (see Figure 4.3). □

In fact, link signature is known to be invariant under genus 2 mutation [17], a stronger result. As another example, we also quickly recover a result of Pawel Traczyk from 2004 [44].

Corollary 5.0.5. *If $\mathcal{D}$ is a connected alternating diagram of a link L, then*

$$\sigma(L) = n_+ + 1 - c$$

where c is the number of circles in the initial resolution of $\mathcal{D}$.

Proof. Fix a one-circle resolution $\mathcal{D}(I^*)$. Since $\mathcal{D}$ is alternating, its initial resolution is the black (or white) resolution with respect to the checkerboard coloring. So $\mathcal{D}(I^*)$ is obtained by resolving along a spanning tree of the black graph (or white graph). Thus

$$w(I^*) = c - 1$$

and the 0-resolution chords lie on the opposite side of the circle $\mathcal{D}(I^*)$ from the 1-resolution chords. So if we order the 0-resolution chords before the 1-resolution chords, then the linking matrix $\mathbb{A}$ has the block form

$$\mathbb{A} = \begin{bmatrix} I_{n-w(I^*)} & B \\ B^T & -I_{w(I^*)} \end{bmatrix}$$

where I_k denotes the $k \times k$ identity matrix and B^T is the transpose of B. Since $\mathbb{A}$ is positive definite on the subspace spanned by the first $n - w(I^*)$ basis vectors and negative definite on the subspace spanned by the last $w(I^*)$ basis vectors, we have $\sigma(\mathbb{A}) = n - 2w(I^*)$. From Theorem 5.0.2 we conclude

$$\sigma(L) = n - 2w(I^*) + w(I^*) - n_- = n_+ - w(I^*) = n_+ - c + 1.$$

□

Let $\widetilde{CKh}(\mathcal{D})$ denote the reduced Khovanov complex of a link diagram $\mathcal{D}$, with $\mathbb{F}_2$ coefficients. Let $\widetilde{Kh}(L)$ denote the reduced Khovanov homology of L with $\mathbb{F}_2$ coefficients. The homological width $w_{\widetilde{Kh}}(L)$ is defined to be one more than the difference between the maximal and minimal δ-gradings over which $\widetilde{Kh}(L)$ is supported (for example, alternating links have width one). We now turn to a new proof of the following proposition regarding homological width, originally proven in [3]. Our method underlies the proof of Theorem 8.1.8, which restricts the shape of the differentials on the monopole link surgery spectral sequence with respect to the δ grading.

Proposition 5.0.6. *If L is non-split and k-almost alternating, then $w_{\widetilde{Kh}}(L) \leq k + 1$.*

In preparation for the proof, we recall some terminology. A diagram $\mathcal{D}$ is *almost alternating* if it is not alternating, but can be made alternating by reversing one crossing, called a *dealternator* [2]. More generally, a diagram $\mathcal{D}$ is *k-almost alternating* if $\mathcal{D}$ can be made alternating by reversing k crossings, but not by reversing only $k-1$ crossings. A link is *k-almost alternating* if it has a k-almost alternating diagram, but not a $(k-1)$-almost alternating diagram. If $\mathcal{D}$ has l crossings, then $\mathcal{D}$ is k-almost alternating for some $k \leq l/2$. In particular, every link is k-almost alternating for some k.

For the proof of Proposition 5.0.6, we will need a refinement of the notion of almost alternating, also from [2]. A diagram $\mathcal{D}$ is called *dealternator connected, k-almost alternating* if $\mathcal{D}$ is k-almost alternating for some choice of k dealternators, such that the corresponding 2^k alternating resolutions are all connected (see Figure 5.1). A link is *dealternator connected, k-almost alternating* if it has a dealternator connected, k-almost alternating diagram, but not a dealternator connected, $(k-1)$-almost alternating diagram.

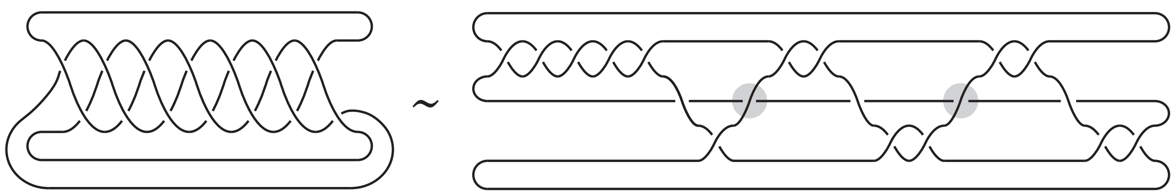

Figure 5.1: A standard diagram of $T(3,7)$, as well as a dealternator connected, 2-almost alternating diagram, constructed using an algorithm in [1].

Two questions arise. Does every link have a dealternator connected, k-almost alternating diagram for some k? When is a k-almost alternating link also dealternator connected, k-almost alternating? Note that the 2-component unlink is 0-almost alternating, but dealternator connected, 1-almost alternating. On the other hand:

Proposition 5.0.7. *A non-split link is dealternator connected, k-almost alternating if and only if it is k-almost alternating. Furthermore, every link is dealternator connected, k-almost alternating for some k.*

Proof. It suffices to show that a connected and k-almost alternating diagram $\mathcal{D}$ may be modified to a dealternator connected, k-almost alternating diagram $\mathcal{D}'$ of the same link. To obtain $\mathcal{D}'$, we modify $\mathcal{D}$ near each dealternator by the local move in Figure 5.2. $\square$

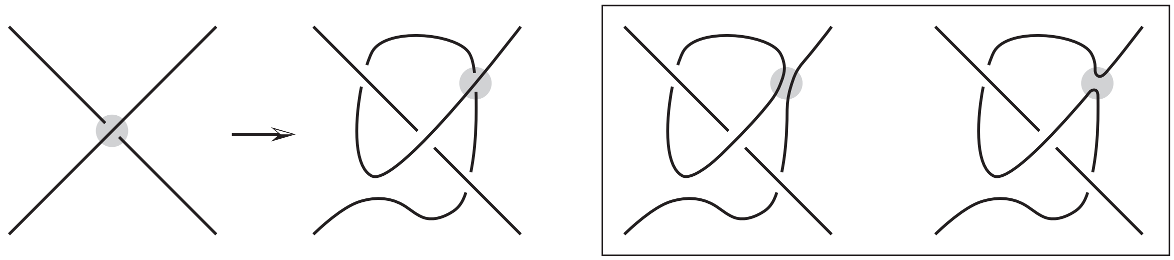

Figure 5.2: At left, the local move (compare with Figure 1 of [1]). At right, we see that resolving at the new dealternators does not disconnect the diagram.

Proof of Proposition 5.0.6. Fix a dealternator connected, k-almost alternating diagram $\mathcal{D}$ for L, and number the crossings so that the first k of l are the dealternators. Let $\mathcal{D}_{I'}$ denote the diagram which results from resolving the dealternators according to $I' \in \{0,1\}^k$. Note that each $\mathcal{D}_{I'}$ is connected and alternating, and therefore non-split by a theorem of Menasco.

Let $x(I')$ be the generator in lowest quantum grading in $\widetilde{CKh}(D)|_{I'\times\{0\}^{l-k}}$. For immediate successors $I' < J'$ in $\{0,1\}^k$, let e be the edge in the hypercube from $I' \times \{0\}^{l-k}$ to $J' \times \{0\}^{l-k}$. A short calculation shows that

$$\delta(x(J')) = \begin{cases} \delta(x(I')) & \text{if } e \text{ merges two circles,} \\ \delta(x(I')) - 1 & \text{if } e \text{ splits a circle in two.} \end{cases}$$

So for any $I' \leq J'$, we have

$$\delta(x(I')) \geq \delta(x(J')), \tag{5.1}$$

with the largest span given by

$$\delta(x(\mathit{0}')) - \delta(x(\mathit{1}')) \leq k. \tag{5.2}$$

Now define the modified weight $\overline{w}(K)$ to be the sum of the first k digits of $K \in \{0,1\}^l$. Filtering $\widetilde{CKh}(\mathcal{D})$ by $\overline{w}(K)$, we obtain a spectral sequence with pages $\overline{E}^i$, converging to $\widetilde{Kh}(L)$. The page $\overline{E}^1$ is the direct sum of the groups $\widetilde{Kh}(\mathcal{D}_{I'})$, each of which is supported on a single diagonal $\delta = \delta(x(I'))$, by Lemmas 5.0.8 and 5.0.9 below. Therefore, by (5.2) this page has width at most $k+1$, and the same must be true of the final page $\widetilde{Kh}(L)$. $\square$

Lemma 5.0.8. *If L is non-split and alternating, then $\widetilde{Kh}(L)$ is thin.*

Proof. We proceed by induction on the crossing number of L. The lemma clearly holds for the unknot. Fix a reduced, alternating diagram $\mathcal{D}$ of L, which necessarily exhibits the crossing number. Since width is preserved under mirroring, without loss of generality, we may assume that the initial complete resolution has at least two circles, and therefore that all edges leaving it are merges. Now fix any crossing in $\mathcal{D}$. The corresponding resolutions have lower crossing number, and are also connected and alternating, thus non-split, and thin by induction. Since the edge e connecting these resolutions is a merge, the page $\overline{E}^1$ is thin as well, as is the final page $\overline{E}^2 \cong \widetilde{Kh}(L)$. $\square$

Let $\widetilde{CKh}(\mathcal{D})|_I$ denote the summand of $\widetilde{CKh}(\mathcal{D})$ supported over resolution I.

Lemma 5.0.9. *Given a connected, alternating diagram $\mathcal{D}$, let x be the generator in lowest quantum grading in $\widetilde{CKh}(\mathcal{D})|_{\mathit{0}}$, and let y be the generator in highest quantum grading in $\widetilde{CKh}(\mathcal{D})|_{\mathit{1}}$. Then $\delta(x) = \delta(y) = \sigma(L)/2$, and at least one of x and y is a non-trivial generator of $\widetilde{Kh}(L)$.*

Proof. Let l be the number of crossings, and let $c(I)$ be the number of circles in resolution $\mathcal{D}(I)$. With respect to the checkerboard coloring, the extremal resolutions of an alternating diagram are the black and white resolutions. Since the black graph and white graph each have l edges and are dual on the sphere, we may express the Euler characteristic of either by

$$c(\mathit{0}) - l + c(\mathit{1}) = 2 \tag{5.3}$$

The quantum and homological gradings of x and y are given by

$$\begin{aligned}
q(x) &= -(c(\mathit{0}) - 1) + n_+ - 2n_- \\
q(y) &= c(\mathit{1}) - 1 + l + n_+ - 2n_- \\
t(x) &= -n_- \\
t(y) &= l - n_-
\end{aligned}$$

Thus, the difference in δ grading is

$$\begin{aligned}\delta(y)-\delta(x) &= \frac{1}{2}(c(\mathit{1})-1+l+c(\mathit{0})-1)-l \\ &= \frac{1}{2}(c(\mathit{0})-l+c(\mathit{1})-2)=0\end{aligned}$$

where the last equality follows from (5.3). Furthermore,

$$2\delta(x)=-(c(\mathit{0})-1)+n_+=-w(I^*)+n_+=\sigma(L)$$

by Corollary 5.0.5.

If $\mathcal{D}$ has no crossings, then x and y coincide and generate $\widetilde{Kh}(L)$. Otherwise, at least one of $\mathcal{D}(\mathit{0})$ and $\mathcal{D}(\mathit{1})$ has two or more circles. If $c(\mathit{0})>1$, then every component of d leaving $\widetilde{CKh}(D)|_{\mathit{0}}$ corresponds to a merging of circles, so x is in the kernel. If $c(\mathit{1})>1$, then every component of d ending at $\widetilde{CKh}(D)|_{\mathit{1}}$ corresponds to a splitting of circles, so y is not in the image. The second claim immediately follows. □

Note that the two lemmas imply a result from [30]:

Corollary 5.0.10. *If L is a non-split alternating link, then $\widetilde{Kh}(L)$ is thin and supported on the diagonal $\delta=\frac{1}{2}\sigma(L)$.*

Chapter 6

From Khovanov homology to monopole Floer homology

Given an oriented link $L \subset S^3$, let $\widetilde{Kh}(L)$ denote the reduced Khovanov homology of L with $\mathbb{F}_2$ coefficients. To a diagram of L, we will associate framed link $\mathbb{L} \subset -\Sigma(L)$. This link is closely related to framed link $\mathbb{L} \subset S^3$ defined in Chapter 4, which we will instead denote by $\mathbb{L}'$ from now on. In this chapter, we apply the $\widetilde{HM}_\bullet$ version of Theorem 2.0.1 to $\mathbb{L} \subset -\Sigma(L)$ to prove:

Theorem 6.0.11. *The link surgery spectral sequence for $\mathbb{L} \subset -\Sigma(L)$ has E^2 page isomorphic to $\widetilde{Kh}(L)$ and converges by the E^{l+1} page to $\widetilde{HM}_\bullet(-\Sigma(L))$.*

While the construction of this spectral sequence depends on a choice of diagram for L, as well as analytic data, Theorem 6.0.11 implies that the E^2 and E^∞ pages are actually link invariants. These pages are also insensitive to Conway mutation, since this is true of Khovanov homology over $\mathbb{F}_2$ as well as branched double covers. More generally, we prove:

Theorem 6.0.12. *For each $k \geq 2$, the isomorphism class of the $(\check{t}, \check{\delta})$-graded vector space E^k depends only on the mutation equivalence class of L.*

The analytic invariance described in Theorem 2.0.2 is crucial here. As explained in Section 6.1, Reidemeister invariance is then an immediate consequence of Baldwin's proof in the Heegaard Floer case [5], whereas mutation invariance follows from our proof in the

Heegaard and monopole Floer cases. Note that both Heegaard Floer proofs, in turn, depend on Roberts' work on invariance with respect to isotopy, handleslides, and stabilization in Heegaard multi-diagrams [37], and Baldwin's work on invariance with respect to almost complex data [5].

Recall that Khovanov homology is graded by two integers, the homological grading t and the quantum grading q. We may repackage this as a rational (t, δ)-bigrading, where

$$\delta = q/2 - t$$

marks the diagonals of slope two in the (t, q)-plane. On the other hand, monopole Floer homology has a canonical mod 2 grading and decomposes over the set of spin^c structures. Using the $\check{\delta}$ grading on the spectral sequence, we derive the first result relating these finer features of monopole or Heegaard Floer homology to those of Khovanov homology, leading to a refinement of the rank inequality

$$\text{rk } \widetilde{Kh}(L) \geq \text{rk } \widetilde{HM}_{\bullet}(-\Sigma(L)) \geq \det(L).$$

Let $\widetilde{HM}^0_{\bullet}(Y)$ and $\widetilde{HM}^1_{\bullet}(Y)$ denote the even and odd graded pieces of $\widetilde{HM}_{\bullet}(Y)$, respectively. Let $\widetilde{Kh}^0(L)$ and $\widetilde{Kh}^1(L)$ denote the even and odd graded pieces of $\widetilde{Kh}(L)$ with respect to the integer grading $\delta - (\sigma(L) - \nu(L))/2$. The terms $\sigma(L)$ and $\nu(L)$ refer to the signature and nullity of L, respectively. Our convention is that the signature of the right-handed trefoil is $+2$ (that is, minus the signature of a Seifert matrix). Recall that $\nu(L) = b_1(\Sigma(L))$.

Theorem 6.0.13. *The $\check{\delta}$ grading on the spectral sequence coincides with*

$$\delta - \frac{1}{2}(\sigma(L) - \nu(L)) \mod 2$$

on the E^2 page. Thus, the rank inequality may be refined to

$$\begin{aligned} \text{rk } \widetilde{Kh}^0(L) &\geq \text{rk } \widetilde{HM}^0_{\bullet}(-\Sigma(L)) \geq \det(L) \\ \text{rk } \widetilde{Kh}^1(L) &\geq \text{rk } \widetilde{HM}^1_{\bullet}(-\Sigma(L)). \end{aligned}$$

Furthermore, the $\check{\delta}$ Euler characteristic of each page is given by $\det(L)$. *The $\check{t}$ grading on the spectral sequence coincides with $t + n_-$ on the E^2 page.*

In particular, all the differentials on the spectral sequence shift $\delta + 2\mathbb{Z}$ by one. Note also that $\widetilde{HM}_\bullet(-\Sigma(L), \mathfrak{s})$ is non-trivial[1] whenever $c_1(\mathfrak{s})$ is torsion, as follows from Corollary 35.1.4 in [24]. We conclude:

Corollary 6.0.14. *If $\widetilde{Kh}(L)$ is supported on a single diagonal, then the spectral sequence collapses at the E^2 page. In particular, $\widetilde{HM}_\bullet(-\Sigma(L))$ is supported in even grading and has rank* $\det(L)$*, with one generator in each spin^c structure.*

In fact, $\widetilde{Kh}(L)$ is supported on the single diagonal $\delta = \sigma(L)/2$ whenever L is quasi-alternating [30]. We proved this fact when L is (non-split) alternating in Corollary 5.0.10. This is consistent with Theorem 6.0.13, since quasi-alternating links have non-zero determinant, and therefore vanishing nullity.

By modifying the proof of Theorem 6.0.13, which invokes the combinatorial Gordon-Litherland signature formula (see [20]), we obtain a (closely) related proof of the signature formula in Theorem 5.0.2. Remarkably, a deep structure theorem in graph theory due to W. H. Cunningham implies that the linking matrix $\mathbb{A}$ alone determines the mutation equivalence class of L, the framed isotopy type of $\mathbb{L}$, and therefore E^i for all $i \geq 1$ (see Remark 6.1.2 and [14] for more details).

Since our results first appeared, Kronheimer and Mrowka have established a similar connection between Khovanov homology and a version of instanton Floer homology for links [26]. As a corollary, they conclude that Khovanov homology detects the unknot.

6.1 Khovanov homology and branched double-covers

We now construct the spectral sequence in Theorem 6.0.11. The key new ingredient is the following $\widetilde{HM}_\bullet$ analog of Proposition 6.2 in [36] for $\widehat{HF}$. We delay the proof to Chapter 7, as it follows from the more general Theorem 7.3.3, as explained in Remark 7.3.6.

Proposition 6.1.1. *Let $Y = \#^k(S^1 \times S^2)$. Then, $\widetilde{HM}_\bullet(Y)$ is a rank one, free module over the ring $\Lambda^* H_1(Y)$ generated by some class Θ, and entirely supported over the torsion spin^c structure. Moreover, if $K \subset Y$ is a curve which represents one of the circles in one of the*

[1]Alternatively, one can show that the $\mathrm{gr}^{(2)}$ Euler characteristic of $\widetilde{HM}_\bullet(Y, \mathfrak{s})$ is 1 if $b_1(Y) = 0$.

$S^1 \times S^2$ summands, then the three-manifold $Y' = Y_0(K)$ is diffeomorphic to $\#^{k-1}(S^1 \times S^2)$, with a natural identification

$$\pi : H_1(Y)/[K] \longrightarrow H_1(Y').$$

Under the cobordism W induced by the two-handle, the map

$$\widetilde{HM}_\bullet(W) : \widetilde{HM}_\bullet(Y) \longrightarrow \widetilde{HM}_\bullet(Y')$$

is specified by

$$\widetilde{HM}_\bullet(W)(\xi \cdot \Theta) = \pi(\xi) \cdot \Theta', \tag{6.1}$$

where here Θ' is the generator of $\widetilde{HM}_\bullet(Y')$, and ξ is any element of $\Lambda^* H_1(Y)$. Dually, if $K \subset Y$ is an unknot, then $Y'' = Y_0(K) \cong \#^{k+1}(S^1 \times S^2)$, with a natural inclusion

$$i : H_1(Y) \longrightarrow H_1(Y'').$$

Under the cobordism W' induced by the two-handle, the map

$$\widetilde{HM}_\bullet(W) : \widetilde{HM}_\bullet(Y) \longrightarrow \widetilde{HM}_\bullet(Y'')$$

is specified by

$$\widetilde{HM}_\bullet(W')(\xi \cdot \Theta) = (\xi \wedge [K'']) \cdot \Theta'', \tag{6.2}$$

where here $[K''] \in H_1(Y'')$ is a generator in the kernel of the map $H_1(Y'') \to H_1(W')$.

We now construct the $\widetilde{HM}_\bullet$ spectral sequence associated to a link $L \subset S^3$. We first fix a diagram $\mathcal{D}$ with l crossings. Following Section 2 of [36], we associate to $\mathcal{D}$ a framed link $\mathbb{L} \subset -\Sigma(L)$ to which we will apply the link surgery spectral sequence. First, in a small ball B_i about the crossing c_i, place an arc with an end on each strand as shown in the ∞ resolution of Figure 6.1. Each of these arcs lifts to a closed loop $\mathbb{K}_i$ in the branched double cover $\Sigma(L)$, giving the components of a link $\mathbb{L} \subset \Sigma(L)$. Note that all of the resolutions of $\mathcal{D}$ agree outside of the union of the B_i. Furthermore, the branched double cover of B_i over the two unknotted strands of $\mathcal{D}(I) \cap B_i$ is a solid torus, with meridian given by the preimage of either of the two strands pushed out to the boundary of B_i. So for each $I \in \{0, 1, \infty\}^l$, we may identify $\Sigma(\mathcal{D}(I)) - \nu(\mathcal{D}(I))$ with $\Sigma(L) - \nu(\mathbb{L})$.

In this way, for each crossing c_i, we obtain a triple of curves $(\lambda_i, \lambda_i + \mu_i, \mu_i)$ in the corresponding boundary component of $\Sigma(L) - \nu(\mathbb{L})$, which represent meridians of the fillings giving the branched covers of the 0, 1, and ∞ resolutions at c_i, respectively. In this cyclic order, the curves may be oriented so that the algebraic intersection number of consecutive curves is $+1$. We change this to -1 by flipping the orientation on the branched double cover (whereas in [36] this is done by replacing L with its mirror). In the language of [36], each triple $(\lambda_i, \lambda_i + \mu_i, \mu_i)$ forms a triad. From our 4-manifold perspective, this is precisely the condition that each 2-handle in a stack is attached to the previous 2-handle using the -1 framing with respect to the cocore (see the discussion preceding Theorem 2.2.11). From either perspective, framing $\mathbb{K}_i$ by λ_i, we are in precisely the setup of the link surgery spectral sequence, with $Y_I = -\Sigma(\mathcal{D}(I))$ for all $I \in \{0, 1, \infty\}^l$. Now, using Proposition 6.1.1, the argument in [36] may be repeated verbatim to show that the complexes (E^1, d^1) and $\widetilde{CKh}(L)$ are isomorphic, and therefore that E^2 is isomorphic to the reduced Khovanov homology $\widetilde{Kh}(L)$.

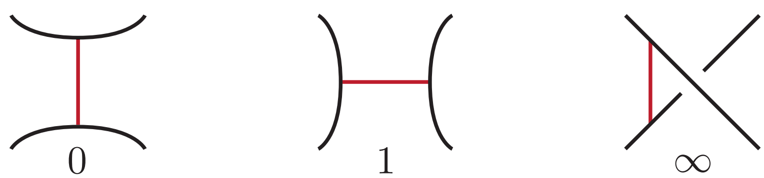

Figure 6.1: We have one short arc between the two strands near each crossing, in both the original diagram and its resolutions.

Alternative proof of Theorem 6.0.11. As an alternative to the argument in [36], we now present a more global description of the isomorphism $E^1 \cong \widetilde{CKh}(L)$, taking advantage of the construction in Section 4.3 of a framed link $\mathbb{L}' \subset S^3$ (there denoted $\mathbb{L}$) which gives a surgery diagram for the branched double cover with reversed orientation $-\Sigma(L)$. This construction of $\mathbb{L}'$ is illustrated in Figure 6.2 for the standard diagram of the right-handed trefoil $\mathcal{T}$. We first fix a vertex $I^* = (m_1^*, \ldots, m_n^*) \in \{0,1\}^l$ for which the resolution $\mathcal{D}(I^*)$ consists of one circle. The link $\mathbb{L}'$ is obtained as the preimage of the corresponding red arcs in Figure 6.1, where the component $\mathbb{K}_i'$ is given the framing $\lambda_i = (-1)^{m_i^*}$. Note that the link $\mathbb{L} \subset -\Sigma(L)$ is represented in this surgery diagram by the framed push-off of $\mathbb{L}'$.

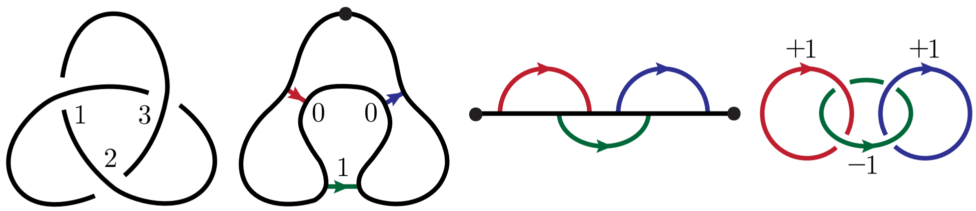

$$\begin{bmatrix} 1 & -1 & 0 \\ -1 & -1 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

Figure 6.2: At left, we number the crossings in a diagram of the right-handed trefoil $\mathcal{T}$. The resolution $\mathcal{D}(010)$ has one circle, and one (arbitrarily-oriented) arc for each crossing. We then cut the circle open at the dot and stretch it out to a line, dragging the arcs along for the ride. Reflecting each arc under the line yields the framed link $\mathbb{L}' \subset S^3$ and linking matrix at right. Surgery on $\mathbb{L}'$ gives $-\Sigma(\mathcal{T})$, which is the lens space $-L(3,1)$.

For each $I = (m_1, \ldots, m_n) \in \{0,1\}^l$, let $\mathbb{L}'_I$ be the link $\mathbb{L}'$ with framing modified to

$$\lambda_i = \begin{cases} \infty & \text{if } m_i = m_i^* \\ 0 & \text{if } m_i \neq m_i^* \end{cases}$$

on $\mathbb{K}'_i$. Then $\mathbb{L}'_I$ gives a surgery diagram for $Y_I = -\Sigma(\mathcal{D}(I)) \cong \#^k(S^1 \times S^2)$, where the resolution $\mathcal{D}(I)$ consists of $k+1$ circles. This is illustrated in Figure 6.3 for the trefoil $\mathcal{T}$. Furthermore, each elementary 2-handle cobordism $W_{IJ} = -\Sigma(F_{IJ}) : -\Sigma(\mathcal{D}(I)) \to -\Sigma(\mathcal{D}(J))$ is given explicitly by 0-surgery on either $\mathbb{K}_i$ or its meridian x_i, in the case where $m_i^* = 0$ or 1, respectively. Let $K_{IJ} \subset -\Sigma(\mathcal{D}(I))$ denote the knot on which we are doing 0-surgery, and let $K''_{IJ} \subset -\Sigma(\mathcal{D}(J))$ be the boundary of the co-core of the corresponding 2-handle.

By Proposition 6.1.1, the (E^1, d^1) page of the link surgery spectral sequence for $\mathbb{L} \subset -\Sigma(L)$ is isomorphic to a complex complex with underlying $\mathbb{F}_2$-vector space

$$\widehat{C}(\mathcal{D}) = \bigoplus_{I \in \{0,1\}^l} \Lambda^* H_1(-\Sigma(\mathcal{D}(I)))$$

and differential $\hat{\partial}$ given by the sum of maps $\hat{\partial}^I_J$ over all immediate successors $I < J$ in $\{0,1\}^l$. These in turn are defined by

$$\hat{\partial}^I_J(\xi) = \begin{cases} \pi(\xi) & \text{if } K_{IJ} \text{ represents a circle factor,} \\ [K''_{IJ}] \wedge i(\xi) & \text{if } K_{IJ} \text{ is an unknot,} \end{cases}$$

On the other hand, $(\widehat{C}(\mathcal{D}), \hat{\partial})$ is precisely the version of $\widetilde{CKh}(D)$ defined at the end of Section 4.2, using the identification of $H_1(-\Sigma(\mathcal{D}(I)))$ with $\widehat{V}(\mathcal{D}(I))$ as explained in 4.3. We conclude that (E^1, d^1) and $\widetilde{CKh}(L)$ are isomorphic complexes, and therefore that E^2 is isomorphic to $\widetilde{Kh}(L)$ as an $\mathbb{F}_2$-vector space. □

Remark 6.1.2. In Chapter 4, we showed that the framed isotopy type of $\mathbb{L} \subset -\Sigma(L)$ is completely determined by the linking matrix $\mathbb{A}$ of $\mathbb{L}' \subset S^3$. If follows that the pages E^i for $i \geq 1$ are determined by $\mathbb{A}$ as well (up to an overall shift in bigrading that depends on $n_\pm$). In fact, since we are working with $\mathbb{F}_2$ coefficients, the orientations of the arcs (and corresponding components of $\mathbb{L}'$) are extraneous as well. We need only record which pairs of arcs in $\mathcal{D}(I^*)$ are linked, as well as I^* itself. On the other hand, the matrix $\mathbb{A}$ with signs fully encodes the odd Khovanov homology of L with $\mathbb{Z}$ coefficients (again, up to an overall shift in bigrading that depends on $n_\pm$), and should also encode a lift of the spectral sequence to $\mathbb{Z}$.

6.1.1 Grading

For the duration of this paragraph, we return to the notation $\underline{E}^i$ to distinguish the $\widetilde{HM}_\bullet$ version of the spectral sequence from the $\widetilde{HM}_\bullet$ version. Our goal is to relate the mod 2 grading $\check{\delta}$ on $\underline{E}^1$ to the integer grading δ on $\widetilde{CKh}(L)$. Since $U_\dagger$ is surjective on $\widetilde{HM}(\#^k S^1 \times S^2)$, $\underline{E}^1$ may be identified as a $\check{\delta}$-graded vector space with the kernel of the map

$$\sum_{I \in \{0,1\}^l} \widetilde{HM}_\bullet(U_\dagger \mid Y_I \times [0,1]) : E^1 \to E^1.$$

This permits us to work with the $\check{\delta}$ grading on E^1 instead.

Proof of Theorem 6.0.13. Let $L \subset S^3$ be an oriented link and fix a diagram $\mathcal{D}$ with n crossings. Let n_+ and n_- denote the number of positive and negative crossings, respectively. Consider the hypercube complex $(X, \check{D})$ given by surgeries on $\mathbb{L}' \subset -\Sigma(L)$. Recall that $Y_I = -\Sigma(\mathcal{D}(I)) \cong \#^k(S^1 \times S^2)$ when the resolution $\mathcal{D}(I)$ consists of $k+1$ circles.

We may think of a generator $x \in \Lambda^r(H_1(Y_I))$ as an element of either $\underline{E}^1$ or $\widetilde{CKh}(\mathcal{D})$. The group $\widetilde{HM}_\bullet(Y_I)$ is supported over the torsion spin^c structure, and Proposition 7.3.7 shows

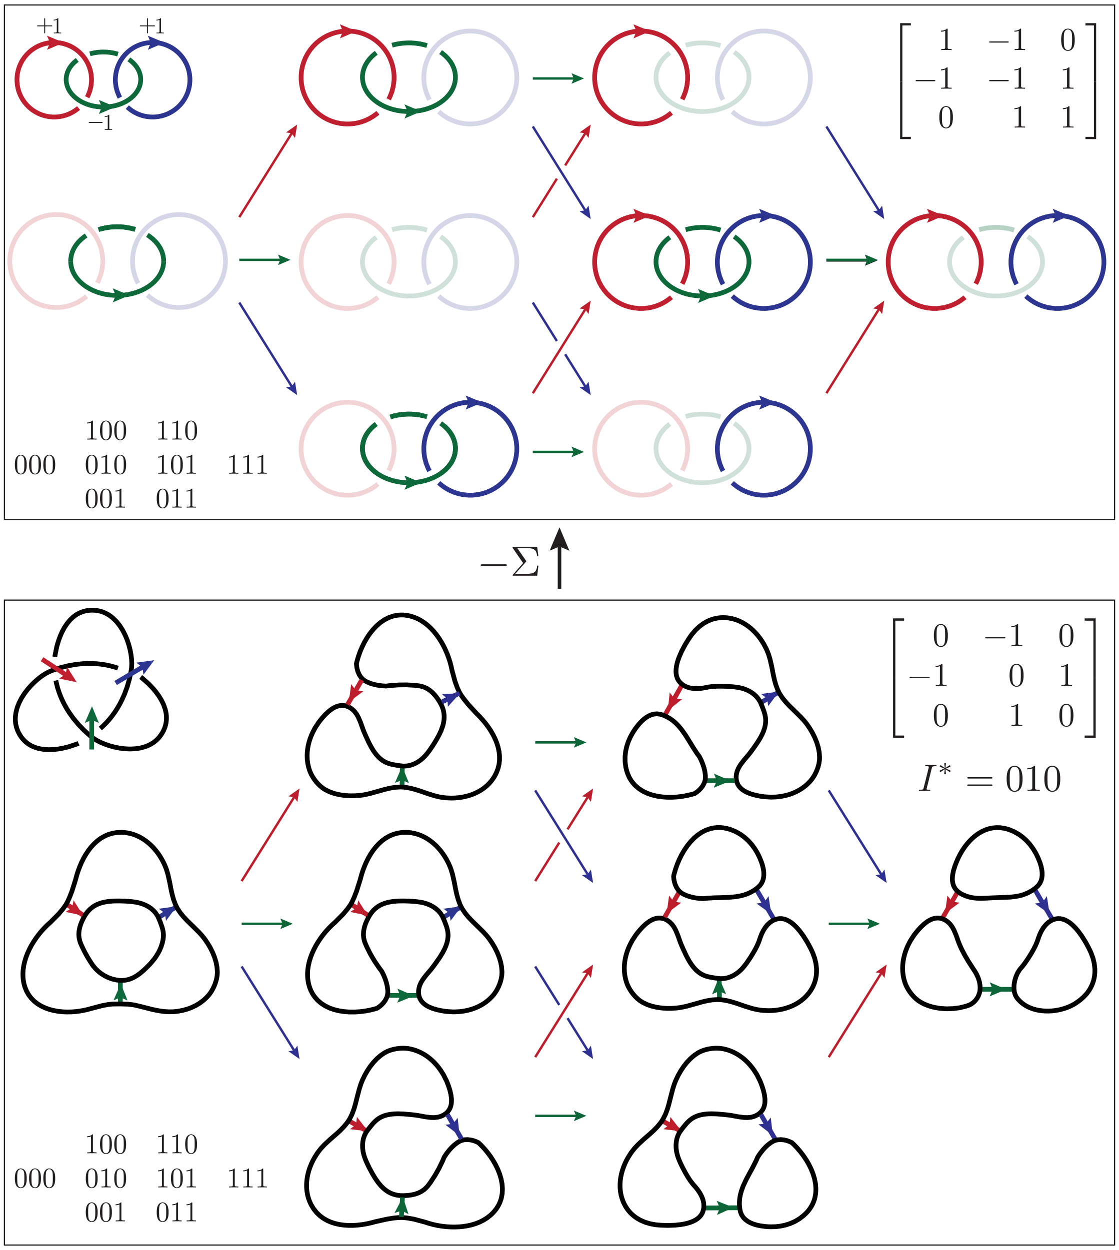

Figure 6.3: Continuing from Figure 6.2, above we obtain the cube of surgery diagrams $\mathbb{L}'_I$ for the right-handed trefoil $\mathcal{T}$. All solid components are 0-framed, while all faded components are ∞-framed. The link surgery cube of 3-manifolds Y_I and 4-dimensional 2-handle cobordisms W_{IJ} above is the branched double cover of the Khovanov cube of 1-manifolds $\mathcal{D}(I) \subset S^3$ and 2-dimensional 1-handle cobordisms $F_{IJ} \subset S^3 \times [0,1]$ below.

that $\mathrm{gr}^{\mathbb{Q}}(x) = k - r$, where $\mathrm{gr}^{\mathbb{Q}}(x)$ is the rational grading over the torsion spin^c-structure, reviewed at the end of Section 1.1. Moreover, on Y_I we have

$$\mathrm{gr}^{\mathbb{Q}}(x) \equiv \mathrm{gr}^{(2)}(x) \mod 2. \tag{6.3}$$

Recall that $\widetilde{CKh}(\mathcal{D})$ has a quantum grading q and a homological grading t. The δ-grading is defined as the linear combination $\delta = \frac{1}{2}q - t$. Translating from the definitions in [34], we may express these gradings as

$$\begin{aligned}
q(x) &= 2\mathrm{gr}^{\mathbb{Q}}(x) - b_1(Y_I) + w(I) + n_+ - 2n_- \\
t(x) &= w(I) - n_- \\
\delta(x) &= \mathrm{gr}^{\mathbb{Q}}(x) - \frac{1}{2}b_1(Y_I) - \frac{1}{2}w(I) + \frac{1}{2}n_+
\end{aligned}$$

Here n_+ and n_- denote the number of positive and negative crossings in $\mathcal{D}$. The final formula defines a function $\delta : X \to \mathbb{Q}$.

By (6.3), we may define a lift $\check{\delta}^{\mathbb{Q}} : X \to \mathbb{Z}$ of the mod 2 grading $\check{\delta}$ from (2.15) by

$$\check{\delta}^{\mathbb{Q}}(x) = \mathrm{gr}^{\mathbb{Q}}(x) - (\iota(W_{I\infty}) + w(I)) + l. \tag{6.4}$$

Note that all cobordisms W_{IJ} over the hypercube satisfy $\sigma(W_{IJ}) = 0$. This is easily seen for an elementary cobordism, and follows in general from signature additivity. So using (1.9), we may expand $\check{\delta}^{\mathbb{Q}}(x)$ as

$$\check{\delta}^{\mathbb{Q}}(x) = \mathrm{gr}^{\mathbb{Q}}(x) - \frac{1}{2}b_1(Y_I) - \frac{1}{2}w(I) - \frac{1}{2}\sigma(W_{0\infty}) + \frac{1}{2}b_1(\Sigma(L)).$$

Finally, we compare $\delta(x)$ with $\check{\delta}^{\mathbb{Q}}$:

$$\begin{aligned}
\delta(x) - \check{\delta}^{\mathbb{Q}}(x) &= \frac{1}{2}(\sigma(W_{0\infty}) + n_+ - b_1(\Sigma(L))). \\
&= \frac{1}{2}(\sigma(L) - \nu(L))
\end{aligned}$$

The last line follows from Lemma 6.1.3 below. Reducing mod 2, we have the first claim of Theorem 6.0.13.

For the remaining claim about the determinant, note that when $\nu(L) = 0$ we have

$$\chi_{\check{\delta}}(E^2) = (-1)^{\sigma(L)/2}\chi_{\delta}(\widetilde{Kh}(L)) = (-1)^{\sigma(L)/2}V_L(-1) = \det(L),$$

where $V_L(q)$ denotes the Jones polynomial of L, and when $\nu(L) > 0$ all of the above terms vanish. Alternatively, one can show that the number of spin^c structures on $-\Sigma(L)$ is $\det(L)$ and that

$$\chi_{\mathrm{gr}^{(2)}}(\widetilde{HM}_\bullet(-\Sigma(L), \mathfrak{s}))$$

is one when $\nu(L) = 0$, and vanishes otherwise. □

Lemma 6.1.3. *The signature and nullity of L are given by*

$$\sigma(L) = \sigma(W_{0\infty}) + n_+$$
$$\nu(L) = b_1(\Sigma(L))$$

Proof. The nullity $\nu(L)$ is sometimes defined as the nullity of any symmetric Seifert matrix S for L, and sometimes as $b_1(\Sigma(L))$. These definitions are equivalent since S presents $H_1(\Sigma(L))$.

We will prove the formula for $\sigma(L)$ by relating $\sigma(W_{0\infty})$ to the signature of a certain 4-manifold X_L bounding $\Sigma(L)$. Recall that the diagram $\mathcal{D}$ has a checkerboard coloring with infinite region in white. The black area forms a spanning surface F for L with one disk for each black region, and one half-twisted band for each crossing. View L as in the boundary of D^4, and push F into the interior. We then define X_L as the branched double cover of D^4 over F. In [20], Gordon and Litherland show that the intersection form of X_L is the Goeritz form G associated to $\mathcal{D}$, and that

$$-\sigma(L) = \sigma(G) - \mu(\mathcal{D}),$$

where $\mu(\mathcal{D}) = c - d$ (see Figure 6.4). The minus sign in front of $\sigma(L)$ is due to the fact that the signature convention in [20] is the opposite of ours. Using the relations

$$w(B) = b + c$$
$$n_- = b + d$$

we can also express $\mu(\mathcal{D})$ as

$$\mu(\mathcal{D}) = w(B) - n_-.$$

Therefore,

$$\sigma(L) = -\sigma(X_L) + w(B) - n_-. \tag{6.5}$$

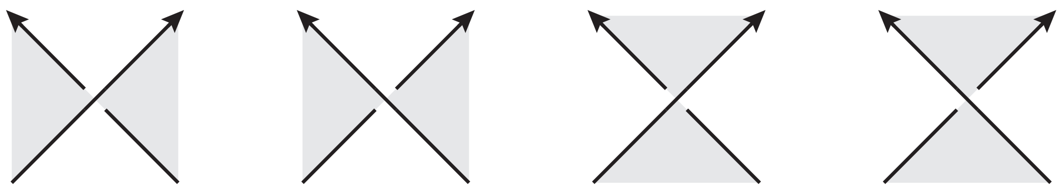

a b c d

Figure 6.4: Four types of crossings in an oriented diagram with checkerboard coloring. The letters a, b, c, and d denote the number of crossings of each type.

We now construct a Kirby diagram for X_L (see Section 3 of [36] and Section 3 of [21] for similar constructions). First, form the black graph resolution $\mathcal{D}(B)$ by resolving each crossing so as to separate the black regions into islands (that is, 1-resolve a crossing if and only if it is of type b or c in Figure 6.4). Draw a 1-handle in dotted circle notation along the boundary of each black region in $\mathcal{D}(B)$. Next, add a 2-handle clasp at each crossing, with framing $+1$ if the crossing is of type b or c, and -1 otherwise. Finally, delete one of the 1-handles.

Next we construct a relative Kirby diagram for the cobordism $W_{B\infty}$. First turn all but one of the circles in $\mathcal{D}(B)$ into 1-handles to get a surgery diagram for $Y_B = \Sigma(\mathcal{D}(B))$, regarded as the incoming end of $W_{B\infty}$. Next, introduce a 0-framed clasp at each of the $n - w(B)$ crossings corresponding to 0 digits of B. This gives a relative Kirby diagram for the cobordism W_{B1}. Finally, introduce -1 framed clasps at each of the remaining crossings, and -1 framed meridians on each of the 0-framed clasps. This gives a relative Kirby diagram for the cobordism $W_{B\infty}$. After pulling off and blowing down all $n - w(B)$ of the -1 framed meridians, and filling in the incoming end with a boundary connect sum of copies of $S^1 \times D^3$, we recover the Kirby diagram for $-X_L$. Therefore,

$$\sigma(W_{B\infty}) = -\sigma(X_L) - (n - w(B)).$$

Combined with (6.5), we conclude

$$\begin{aligned}\sigma(L) &= -\sigma(X_L) + w(B) - n_- \\ &= \sigma(W_{B\infty}) + (n - w(B)) + w(B) - n_- \\ &= \sigma(W_{0\infty}) + n_+.\end{aligned}$$

For the last line, note that since $\sigma(W_{0B})$ vanishes, signature additivity implies that $\sigma(W_{0\infty}) = \sigma(W_{B\infty})$. □

Modifying the above proof, we obtain the signature formula in Section 5:

Alternative proof of the signature formula in Theorem 5.0.2. Recall the construction of the surgery diagram $\mathbb{L}'$ for $-\Sigma(L)$, as in Figure 6.2. Let Z_L be the 4-manifold obtained by attaching 2-handles along $\mathbb{L}'$. By construction, Z_L bounds $-\Sigma(L)$. Just as in the above proof, a Kirby diagram argument shows that

$$\sigma(W_{0\infty}) = \sigma(W_{I^*\infty}) = \sigma(Z_L) - (n - w(I^*)) = \sigma(\mathbb{A}) - (n - w(I^*)),$$

where $\mathbb{A}$ is the linking matrix of $\mathbb{L}$, which is congruent to the linking matrix of the arcs in Theorem 5.0.2. From Lemma 6.1.3, we arrive at the formula

$$\begin{aligned}\sigma(L) &= \sigma(W_{0\infty}) + n_+ \\ &= \sigma(\mathbb{A}) + w(I^*) - n_-.\end{aligned}$$

□

6.1.2 Invariance

We now turn to the proof of Theorem 6.0.12, which describes the extent to which the spectral sequence depends on the choice of diagram for the link L.

Proof of Theorem 6.0.12. Let $\mathcal{D}_1$ and $\mathcal{D}_2$ be two diagrams of the oriented link $L \subset S^3$. Let $\underline{X}(\mathcal{D}_i)$ represent the hypercube complex associated to diagram $\mathcal{D}_i$, for some choice of analytic data (which we may suppress by Theorem 2.0.2). The goal is to construct a filtered chain map

$$\phi : \underline{X}(\mathcal{D}_1) \to \underline{X}(\mathcal{D}_2)$$

which induces an isomorphism on the $\underline{E}^2$ page, and therefore on all higher pages as well. It suffices to consider the case where $\mathcal{D}_1$ and $\mathcal{D}_2$ differ by a single Reidemeister move.

In [5], Baldwin defines such a map ϕ for each of the three Reidemeister moves. While he was considering the Heegaard Floer version of the spectral sequence, his maps have direct analogues in the monopole Floer case. The difficult part is proving that ϕ induces a homotopy equivalence from $\widetilde{CKh}(\mathcal{D}_1)$ to $\widetilde{CKh}(\mathcal{D}_2)$ on the $\underline{E}^1$ page. However, this argument only involves properties of the Khovanov differential, drawing heavily on the proof that Khovanov homology is a bigraded link invariant (see [34]). It is also clear from the construction that ϕ preserves the bigrading on Khovanov homology, and therefore $\check{\delta}$.

Now, suppose links L_1 and L_2 are related by a mutation. Fix diagrams $\mathcal{D}_1$ and $\mathcal{D}_2$ for L_1 and L_2 which exhibit the mutation. Let $\mathbb{L}_1'$ and $\mathbb{L}_2'$ be the associated framed links in S^3, and let $\mathbb{L}_1$ and $\mathbb{L}_2$ be the associated framed links in $-\Sigma(L_1)$ and $-\Sigma(L_2)$. In Section 4.3, we proved that there is an orientation-preserving diffeomorphism

$$\psi' : S^3 \to S^3$$

for which $\psi(\mathbb{L}_1')$ and $\mathbb{L}_2'$ are isotopic as framed links. This implies that there is an orientation-preserving diffeomorphism

$$\psi : -\Sigma(L_1) \to -\Sigma(L_2)$$

for which $\psi(\mathbb{L}_1)$ and $\mathbb{L}_2$ are isotopic as framed links. Appealing to Theorem 2.0.2, we conclude that the $\underline{E}^i$ pages of the spectral sequences associated to $\mathcal{D}_1$ and $\mathcal{D}_2$ are isomorphic for all $i \geq 1$. □

6.2 The spectral sequence for a family of torus knots?

In order to illustrate the spectral sequence in action, we now present an example which is both speculative and, we hope, compelling. Consider the family of torus knots given by

$$\{T(3, 6n \pm 1) \,|\, n \geq 1\}.$$

For this family, the unreduced Khovanov homology with coefficients in $\mathbb{Q}$ takes the form of repeating blocks, and is stable, up to a shift in quantum grading, as n grows [41],

[45]. Watson recently deduced a similar structure over $\mathbb{F}_2$ up to a repeating indeterminate summand [47]. Computing $\widetilde{Kh}(T(3,6n\pm 1))$ explicitly for several values of n using Bar-Natan's *KnotTheory* package [4], a consistent pattern emerges, as shown in Figure 6.5. Assuming this pattern persists, up to a shift in quantum grading, we then have inclusions

$$\widetilde{Kh}(T(3,5)) \subset \widetilde{Kh}(T(3,7)) \subset \widetilde{Kh}(T(3,11)) \subset \widetilde{Kh}(T(3,13)) \subset \cdots.$$

Conjecture 6.2.1. *For each such torus knot, and some choice of analytic data and diagram, the higher differentials are as shown in Figure 6.5. In particular, the spectral sequence converges at the E^4 page, and the above inclusions on the E^2 page extend to the E^3 and E^4 pages.*

One intriguing way to frame this conjecture is as follows. Using the (t,q)-bigrading in Figure 6.5, we may define *higher δ-polynomials* by

$$U^k_{T(3,6n\pm1)}(\delta) = \sum_{i,j}(-1)^i \operatorname{rk} E^k_{i,j}(T(3,6n\pm 1))\, \delta^{j/2-i}$$

for each $k \geq 2$. Then Conjecture 6.2.1 implies that the δ-polynomials on the E^2 and E^3 pages are monic monomials, while

$$U^4_{T(3,6n\pm1)}(\delta) = \delta^{\sigma/2} - \delta^{\sigma/2+1} + \delta^{\sigma/2+2} - \cdots + \delta^{s/2}. \tag{6.6}$$

Here s denotes Rasmussen's s-invariant. We emphasize that these polynomials will in general depend on the branch set (not only on the branched double-cover), even on the E^∞ page, as we show in Proposition 8.2.5. Still, it would be interesting to compare (6.6) with the polynomials arising from Greene's conjectured δ-grading on $\widehat{HF}(-\Sigma(T(3,6n\pm 1)),$ which may also depend on the branch set. This grading is defined in Section 8 of [21] using a special Heegaard diagram associated to the link diagram.

Our primary evidence for Conjecture 6.2.1 comes from [5], where Baldwin deduces that the Heegaard Floer spectral sequence for $T(3,5)$ is as shown in Figure 6.5. His argument uses the Khovanov and Heegaard Floer contact invariants to show that the lower left generator survives to E^∞ for every torus knot. This is the only survivor in the case of $T(3,5)$, as the branched double cover is the Poincaré homology sphere. We have not rigorously computed

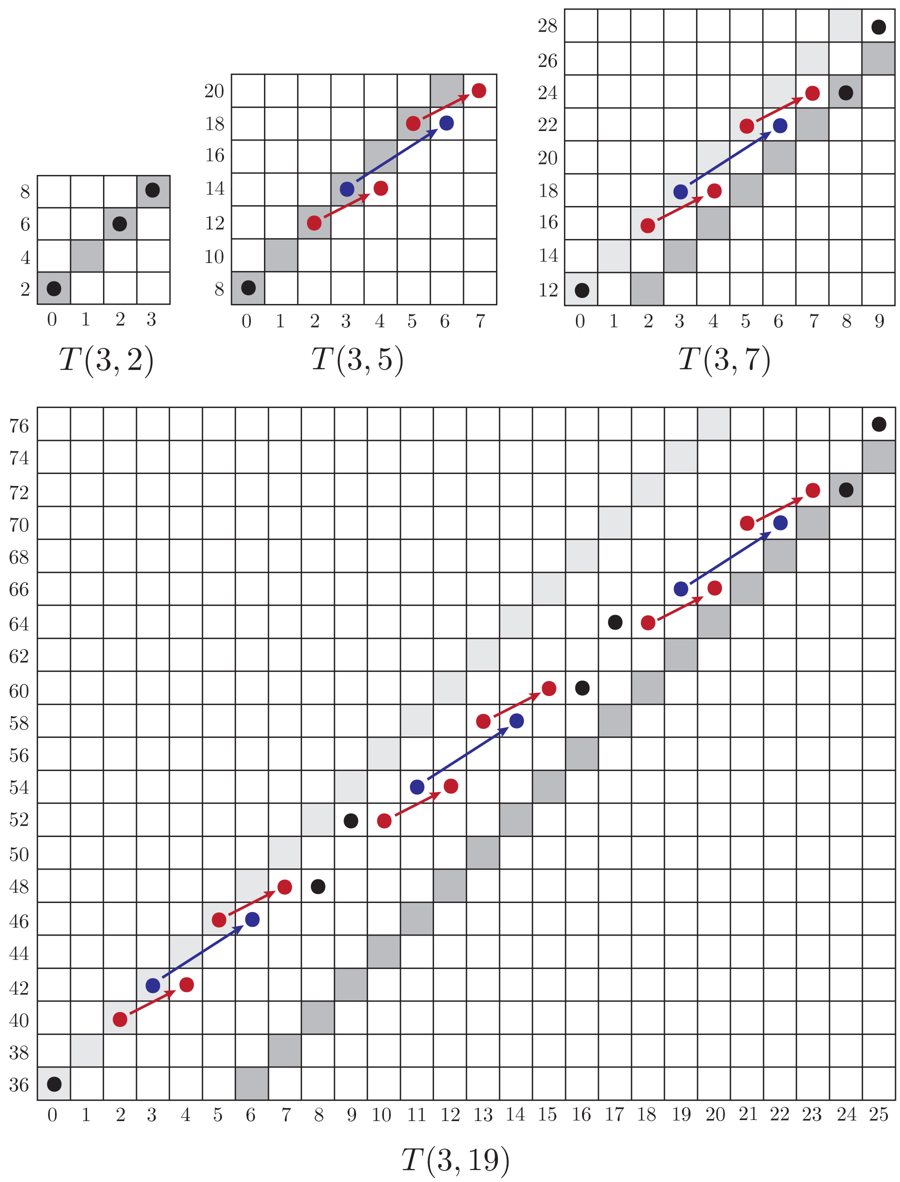

Figure 6.5: Each dot represents an $\mathbb{F}_2$ summand of $\widetilde{Kh}(T(3,6n \pm 1))$ in the (t,q)-plane. The diagonal $\delta = \sigma/2$ is heavily shaded and the diagonal $\delta = s/2$ is lightly shaded (unless $s = \sigma$). The d^2 and d^3 differentials are in red and blue, respectively, as are their victims. The surviving (black) dots generate $\widetilde{HM}_{\bullet}(-\Sigma(2,3,6n \pm 1))$. The shaded diagonals also correspond to $\check{\delta}^{\mathbb{Q}} = 0$ and $\check{\delta}^{\mathbb{Q}} = (2n-1) \pm 1$. For $n \geq 1$, there is precisely one black dot on each diagonal in this range, giving the expected rank of E^{∞} in each $\mathrm{gr}^{(2)}$ grading.

the monopole Floer spectral sequence even in this case, since we lack an analogous contact invariant.

As further evidence, we cite the compatibility of $\widetilde{HM}_\bullet(-\Sigma(T(3,6n-1)))$ with the E^∞ page implied by our conjecture. The branched double cover of $T(3,6n\pm 1)$ is the Brieskorn integer homology sphere

$$\Sigma(2,3,6n\pm 1),$$

which arises by $1/(6n\pm 1)$ Dehn surgery on a trefoil knot (the right-handed one for $6n+1$ and the left-handed one for $6n-1$). Using the surgery exact triangle, gradings, and module structure, the Heegaard Floer groups $HF^+(\Sigma(2,3,6n\pm 1))$ are explicitly calculated in [35]. The same techniques should apply in the monopole case to show directly, using (2.20), that

$$\begin{aligned}\widetilde{HM}_\bullet(-\Sigma(2,3,6n-1)) &= \mathbb{Z}_{(-2)} \oplus \left(\mathbb{Z}_{(-2)} \oplus \mathbb{Z}_{(-1)}\right)^{n-1}\\ \widetilde{HM}_\bullet(-\Sigma(2,3,6n+1)) &= \mathbb{Z}_{(0)} \oplus \left(\mathbb{Z}_{(0)} \oplus \mathbb{Z}_{(1)}\right)^{n}\end{aligned}$$

where the subscript denotes $\mathrm{gr}^{\mathbb{Q}}$ grading (see also [33]). Indeed, we carry this out for $n=1$ in Section 8.2. For general n, we may also appeal to the equivalence of Heegaard and monopole Floer groups announced in [28]. In particular,

$$\operatorname{rk}\widetilde{HM}_\bullet(-\Sigma(2,3,6n\pm 1)) = 2n\pm 1$$

and so the $\check{\delta}$-graded ranks on our conjectural E^4 page are compatible with the $\mathrm{gr}^{(2)}$-graded ranks on $\widetilde{HM}_\bullet(-\Sigma(2,3,6n\pm 1))$, as required by Theorem 6.0.13.

For links such that the (t,q) bigrading on the higher pages is well-defined, we may encode the higher pages of the spectral sequence in the form of a 2-variable *higher Khovanov polynomial*, given by

$$E^k_L(t,q) = \sum_{i,j} \operatorname{rk} E^k_{i,j}(L)\, t^i q^j$$

for each $k \geq 2$. We then obtain *higher Jones polynomials*, given by

$$V^k_L(q) = E^k_L(-1, q^{1/2})$$

for each $k \geq 2$. The ordinary Jones polynomial $V_L(q)$ coincides with $V^2_L(q)$. Furthermore, if L is quasi-alternating, then $V^k_L(q) = V_L(q)$ for all $k \geq 2$.

We now record in full the various higher polynomials associated with the differentials in Figure 6.5, in order to provide a (conjectural) data-set to aid in the search for a combinatorial description (in fact, one possible description has recently been suggested by Szabó [43], see the discussion following Corollary 8.0.12). For comparison, we include the polynomials on the E^2 page as well. Note that the optimal input for an algorithm may not be a diagram of the link itself, but rather the arc-linking data which encodes the mutation equivalence class, as described in Remark 6.1.2.

Conjecture 6.2.2. *Let $\mathcal{S}^n = T(3, 6n+1)$ and let $\mathcal{T}^n = T(3, 6n-1)$. Set*

$$f_n(t,q) = \sum_{k=0}^{n-1} t^{8k} q^{12k} \qquad\qquad f_n(q) = \sum_{k=0}^{n-1} q^{6k}$$

For each $n \geq 1$, the higher Khovanov polynomials are given by

$$q^{-s} E^2_{\mathcal{S}^n}(t,q) = 1 + \big((t^8q^{12} + t^5q^{16}) + (t^3q^6 + t^6q^{10}) + (t^2q^4 + t^4q^6 + t^5q^{10} + t^7q^{12})\big)\, f_n(t,q)$$

$$q^{-s} E^3_{\mathcal{S}^n}(t,q) = 1 + \big((t^8q^{12} + t^5q^{16}) + (t^3q^6 + t^6q^{10})\big)\, f_n(t,q)$$

$$q^{-s} E^4_{\mathcal{S}^n}(t,q) = 1 + (t^8q^{12} + t^5q^{16}) f_n(t,q)$$

$$q^{-s} E^2_{\mathcal{T}^n}(t,q) = 1 + (t^8q^{12} + t^5q^{16}) f_{n-1}(t,q) + \big((t^3q^6 + t^6q^{10}) + (t^2q^4 + t^4q^6 + t^5q^{10} + t^7q^{12})\big)\, f_n(t,q)$$

$$q^{-s} E^3_{\mathcal{T}^n}(t,q) = 1 + (t^8q^{12} + t^5q^{16}) f_{n-1}(t,q) + (t^3q^6 + t^6q^{10}) f_n(t,q)$$

$$q^{-s} E^4_{\mathcal{T}^n}(t,q) = 1 + (t^8q^{12} + t^5q^{16}) f_{n-1}(t,q)$$

The higher Jones polynomials are given by

$$q^{-s/2}\, V^2_{\mathcal{S}^n}(q) = 1 + q^2 - q^{6n+2}$$

$$q^{-s/2}\, V^3_{\mathcal{S}^n}(q) = 1 + \big((q^6 - q^8) + (-q^3 + q^5)\big)\, f_n(q)$$

$$q^{-s/2}\, V^4_{\mathcal{S}^n}(q) = 1 + (q^6 - q^8) f_n(q)$$

$$q^{-s/2}\, V^2_{\mathcal{T}^n}(q) = 1 + q^2 - q^{6n}$$

$$q^{-s/2}\, V^3_{\mathcal{T}^n}(q) = 1 + (q^6 - q^8) f_{n-1}(q) + (-q^3 + q^5) f_n(q)$$

$$q^{-s/2}\, V^4_{\mathcal{T}^n}(q) = 1 + (q^6 - q^8) f_{n-1}(q)$$

The higher δ-polynomials are given by

$$\delta^{-\sigma/2}\,U^2_{\mathcal{S}^n}(\delta) = 1$$
$$\delta^{-\sigma/2}\,U^3_{\mathcal{S}^n}(\delta) = 1$$
$$\delta^{-\sigma/2}\,U^4_{\mathcal{S}^n}(\delta) = 1 - \delta + \delta^2 - \cdots + \delta^{2n}$$
$$\delta^{-\sigma/2}\,U^2_{\mathcal{T}^n}(\delta) = \delta^{-1}$$
$$\delta^{-\sigma/2}\,U^3_{\mathcal{T}^n}(\delta) = \delta^{-1}$$
$$\delta^{-\sigma/2}\,U^4_{\mathcal{T}^n}(\delta) = 1 - \delta + \delta^2 - \cdots + \delta^{2n-2}$$

Here $s(\mathcal{S}^n) = 12n$, $s(\mathcal{T}^n) = 12n - 4$, and $\sigma(\mathcal{S}^n) = \sigma(\mathcal{T}^n) = 8n$. So both $U^4_{\mathcal{S}^n}(\delta)$ and $U^4_{\mathcal{T}^n}(\delta)$ may be expressed as

$$\delta^{\sigma/2} - \delta^{\sigma/2+1} + \delta^{\sigma/2+2} - \cdots + \delta^{s/2}.$$

Chapter 7

Donaldson's TQFT

7.1 The algebraic perspective

The identity of the E^2 page in Theorem 6.0.11 may be understood as follows. To a diagram of a link $L \subset S^3$, we associate a framed link $\mathbb{L} \subset -\Sigma(L)$. With respect to $\mathbb{L}$, the link surgery hypercube of 3-manifolds $Y(I)$ and 4-dimensional 2-handle cobordisms $W(IJ)$ is precisely the branched double cover of the Khovanov hypercube of 1-manifolds $\mathcal{D}(I) \subset S^3$ and 2-dimensional 1-handle cobordisms $F(IJ) \subset S^3 \times [0,1]$, as illustrated for the trefoil knot in Figures 6.2 and 6.3. Furthermore, the functor $\widetilde{HM}_\bullet$ and the functor CKh underlying Khovanov's *unreduced* theory over $\mathbb{F}_2$ fit into the commutative square of functors in Figure 7.1.

$$\begin{tikzcd}
\boxed{Y(I) \xrightarrow{W(IJ)} Y(J)} \arrow[r, "\widetilde{HM}_\bullet"] & \boxed{V(I) \xrightarrow{T(IJ)} V(J)} \arrow[d] \\
\boxed{\mathcal{D}(I) \xrightarrow{F(IJ)} \mathcal{D}(J)} \arrow[u, "-\Sigma"] \arrow[r, "CKh"] \arrow[ur, phantom, "\circlearrowright"] & \boxed{U(I) \xrightarrow{S(IJ)} U(J)}
\end{tikzcd}$$

Figure 7.1: Commutative diagram relating the functors $\widetilde{HM}_\bullet$ and CKh.

Here $S(IJ) : U(I) \to U(J)$ and $T(IJ) : V(I) \to V(J)$ represent the induced maps of $\mathbb{F}_2$-vector spaces with respect to each theory. If we replace CKh with the *reduced* Khovanov functor $\widetilde{CKh}$ over $\mathbb{F}_2$, then the vertical arrow at right induces an equivariant isomorphism of vector spaces. Consequently, we may identify the complex (E^1, d^1) with $\widetilde{CKh}(L)$, and hence E^2 with $\widetilde{Kh}(L)$.

In fact, when I and J are immediate successors, the entire commutative diagram admits a more elementary and unified description. Both horizontal arrows may be regarded as an instance of a TQFT described by Donaldson in [16]. The algebraic basis for his construction is as follows. To an $\mathbb{F}_2$-vector space U, we associate the exterior algebra $\Lambda^* U$. To a linear map $i : \Gamma \to U_0 \oplus U_1$, we associate a map $|\Gamma| : \Lambda^* U_0 \to \Lambda^* U_1$ defined as follows. Let k and n be the dimensions of Γ and U_0, respectively. By taking the exterior product of the images of the elements in any basis of Γ, we obtain an element of $\Lambda^k(U_0 \oplus U_1)$, which may be regarded as a map via the series of isomorphisms

$$\begin{aligned}\Lambda^k(U_0 \oplus U_1) &\cong \textstyle\bigoplus_{i=0}^k \Lambda^i U_0 \otimes \Lambda^{k-i} U_1 \\ &\cong \textstyle\bigoplus_{i=0}^k (\Lambda^{n-i} U_0)^* \otimes \Lambda^{k-i} U_1 \\ &\cong \textstyle\bigoplus_{i=0}^k \mathrm{Hom}(\Lambda^{n-i} U_0, \Lambda^{k-i} U_1).\end{aligned}$$

Given $i_1 : \Gamma_1 \to U_0 \oplus U_1$ and $i_2 : \Gamma_2 \to U_1 \oplus U_2$, we define the composition $i : \Gamma \to U_0 \oplus U_2$ by setting

$$\Gamma = \{(x, z) \in \Gamma_1 \oplus \Gamma_2 \,|\, i_1(x)|_{U_1} = i_2(z)|_{U_1}\}$$

and $i(x, z) = (i_1(x)|_{U_0}, i_2(z)|_{U_2})$. If i_1 and i_2 are transverse as maps into U_1, then indeed $|\Gamma| = |\Gamma_2| \circ |\Gamma_1|$.

To a manifold M, Donaldson assigns the exterior algebra $\Lambda^* H^1(M)$. To a cobordism $N : M_0 \to M_1$, he assigns the map

$$|H^1(N)| : \Lambda^* H^1(M_0) \to \Lambda^* H^1(M_1),$$

obtained from the restriction $H^1(N) \to H^1(\partial N) \cong H^1(M_0) \oplus H^1(M_1)$. We refer to these assignments as Donaldson's TQFT, even though $\Lambda^* H^1$ is not a functor in general, as the transversality condition may fail to hold (see Remark 7.1.1 below). However, for I and J immediate successors, the above commutative diagram may be rewritten in terms of $\Lambda^* H^1$ as shown in Figure 7.2, as explained below. Since CKh and $\widetilde{HM}_\bullet$ are indeed functors, the first commutative diagram then holds for all $I < J$, not only immediate successors.

$$
\begin{tikzcd}
\boxed{Y(I) \xrightarrow{W(IJ)} Y(J)} \arrow[r, "\Lambda^* H^1"] \arrow[dr, phantom, "\circlearrowright"] & \boxed{V(I) \xrightarrow{T(IJ)} V(J)} \arrow[d, "(\Lambda^* \rho)*"'] \\
\boxed{\mathcal{D}(I) \xrightarrow{F(IJ)} \mathcal{D}(J)} \arrow[u, "-\Sigma"] \arrow[r, "\Lambda^* H^1"] & \boxed{U(I) \xrightarrow{S(IJ)} U(J)}
\end{tikzcd}
$$

Figure 7.2: For immediate successors $I < J$, the commutative diagram in Figure 7.1 is identified with the commutative diagram above, with both rows given by Donaldson's TQFT.

We explain the vertical map $(\Lambda^*\rho)*$ at right in the diagram. For any link L in S^3, the exact sequence

$$H_1(S^3, L) \to H_0(L) \to H_0(S^3) \to 0$$

gives an isomorphism $H_1(S^3, L) \cong \widetilde{H}_0(L)$. There is a natural map $\widetilde{H}_0(L) \cong H_1(S^3, L) \to H_1(\Sigma(L))$ which takes a relative 1-cycle to its preimage. Dually, there is a map

$$\rho : H^1(\Sigma(L)) \to H^1(S^3, L) \cong \widetilde{H}^0(L) \subset H^0(L) \cong H^1(L),$$

which induces a map on exterior algebras

$$\Lambda^*\rho : \Lambda^* H^1(\Sigma(L)) \to \Lambda^* H^1(L).$$

The vertical map is given by this map precomposed with the Hodge dual $*$ on $\Lambda^* H^1(\Sigma(L))$. Here L equals $\mathcal{D}(I)$ or $\mathcal{D}(J)$.

The equivalence of the two commutative diagrams when I and J are immediate successors may be understood as follows. The manifold $Y(I)$ admits a metric of positive scalar curvature, so it follows from Proposition 36.1.3 of [24] that $\widetilde{HM}_\bullet(Y(I))$ is the cohomology of the torus $\mathbb{T}(Y(I)) = H^1(Y(I); \mathbb{R})/H^1(Y(I); \mathbb{Z})$, parameterizing flat $U(1)$-connections on $Y(I)$ modulo gauge. The cobordism $W(IJ)$ also admits a metric of positive scalar curvature, and indeed, we establish in Section 7.3 below that for elementary 2-handle cobordisms, the map $\widetilde{HM}_\bullet(W(IJ))$ coincides with the map on cohomology induced by the correspondence between tori defined by flat connections over $W(IJ)$. As Donaldson observes, the map on cohomology induced by such a correspondence is encoded in the above TQFT. Along the bottom row, note that Donaldson's TQFT is a bona fide functor in $1+1$ dimensions, and one may easily check that the same Frobenius algebra underlies both $\Lambda^* H^1$ and CKh with $\mathbb{F}_2$ coefficients.

Remark 7.1.1. The transversality needed for the functoriality in Donaldson's TQFT fails to hold in $3+1$ dimensions. The simplest such failure is that of a single circle splitting and re-merging. Then $F(IJ)$ is the complement of two disks in the 2-torus, and $W(IJ)$ is the (simply-connected) complement of two balls in $S^2 \times S^2$. Tracing definitions, we see $\Lambda^* H^1(F(IJ)) = 0$ whereas $\Lambda^* H^1(W(IJ)) = \mathrm{Id}$. More generally, if $F(IJ)$ has positive genus, then the restriction map from $H^1(F(IJ))$ has non-trivial kernel, and thus $\Lambda^* H^1(F(IJ)) = 0$, whereas $\Lambda^* H^1(W(IJ))$ is non-zero even though $\widetilde{HM}_\bullet(W(IJ))$ vanishes[1]. Conversely, if the genus of $F(IJ)$ is zero, then indeed $\Lambda^* H^1(W(IJ)) = \widetilde{HM}_\bullet(W(IJ))$ and the two commutative diagrams coincide.

Remark 7.1.2. The arrows on a decorated link diagram may be used to specify homology orientations[2] on both $F(IJ)$ and $W(IJ)$. In this way, the entire commutative diagram in Figure 7.1 lifts to $\mathbb{Z}$ coefficients, with $\widetilde{HM}_\bullet$ along the top row and the odd Khovanov TQFT, or equivalently Donaldson's TQFT, along the bottom row. We will return to this in future work concerning a lift of the spectral sequence to $\mathbb{Z}$ coefficients.

7.2 The geometric perspective

In the previous section, we described Donaldson's TQFT, there denoted $\Lambda^* H^1$, from an algebraic standpoint. We now elaborate on its geometric interpretation in order to make explicit its relationship to monopole Floer homology. The group $\Lambda^* H^1(Y)$ arises geometri-

[1] Note that $b_2^+(W(IJ))$ coincides with the genus of $F(IJ)$. So when this genus is positive, we may directly deduce that $\widetilde{HM}_\bullet(W(IJ))$ vanishes from the fact that $W(IJ)$ admits a positive scalar curvature metric with cylindrical ends.

[2] In Donaldson's TQFT, a homology orientation on the cobordism $N : M_0 \to M_1$ is an orientation of the line $\Lambda^{\max} H^1(N; \mathbb{R}) \otimes \Lambda^{\max} H^1(M_0; \mathbb{R})$. This determines an overall sign on the map $\Lambda^* H^1(N)$. The definition of homology orientation in monopole Floer homology reduces to this one for cobordisms with $b_2^+ = 0$ (see Definition 3.4.1 of [24] for details).

cally as the homology group $H_*(\mathbb{T}(Y))$ of the torus

$$\mathbb{T}(Y) = H^1(Y;\mathbb{R})/H^1(Y;\mathbb{Z})$$

parameterizing flat $U(1)$-valued connections on Y mod gauge. A cobordism $W : Y_0 \to Y_1$ gives rise to a correspondence

$$\mathbb{T}(Y_0) \xleftarrow{r_0} \mathbb{T}(W) \xrightarrow{r_1} \mathbb{T}(Y_1) \tag{7.1}$$

by restriction of connections. The Donaldson map $\Lambda^* H^1(W)$ then arises as the map on homology groups

$$f_W : H_*(\mathbb{T}(Y_0)) \to H_*(\mathbb{T}(Y_1))$$

given by "pull-up/push-down". More precisely,

$$f_W = (r_1)_* \, \omega_W^{-1} \, r_0^* \, \omega_{Y_0}$$

where ω_{Y_0} and ω_W denote the two Poincaré duality isomorphisms[3]

$$\begin{aligned} \omega_{Y_0} : & H_*((\mathbb{T}(Y_0)) \to H^*((\mathbb{T}(Y_0)) \\ \omega_W : & H_*((\mathbb{T}(W)) \to H^*((\mathbb{T}(W)). \end{aligned}$$

The degree of f_W is given by $b_1(W) - b_1(Y_0)$. The homology group $H_*(\mathbb{T}(Y_i))$ is a module over the cohomology group $H^*(\mathbb{T}(Y_i))$ via the cap product. The map f_W is natural with respect to the cap product in the following sense: if $\gamma_0 \in H^*(\mathbb{T}(Y_0))$, $\gamma_1 \in H^*(\mathbb{T}(Y_1))$, and

$$r_0^*(\gamma_0) = r_1^*(\gamma_1) \tag{7.2}$$

then

$$f_W(\gamma_0 \cap x) = \gamma_1 \cap f_W(x) \tag{7.3}$$

for each $x \in H_*(\mathbb{T}(Y_0))$.

---

[3]Unless otherwise specified, we continue to use $\mathbb{F}_2$ coefficients, so $H_*(\mathbb{T}(Y)) = H_*(\mathbb{T}(Y);\mathbb{F}_2)$ and Poincaré duality is well-defined without additional orientation data. See Section 2.8 of [24] for an equivalent Morse theoretic construction of f_W.

Remark 7.2.1. In particular, if $\mathrm{Im}(\mathrm{r}_0^*) \subset \mathrm{Im}(\mathrm{r}_1^*)$, then f_W is determined by its value on the fundamental class of $\mathbb{T}(Y_0)$. This is the case, for example, whenever W is the trace of surgery on a framed link.

Kronheimer and Mrowka develop a related module structure in monopole Floer homology (see Sections 3.2 and 23.2 of [24]). Indeed, the group $\widetilde{HM}_\bullet(Y)$, as well as its relatives, is a module over the ordinary cohomology ring of the ambient space $\mathcal{B}^\sigma(Y)$ via a cap product

$$\cap : H^*(B^\sigma(Y)) \times \widetilde{HM}_\bullet(Y) \to \widetilde{HM}_\bullet(Y)$$

which is natural with respect to cobordism maps, as we now explain. The space $\mathcal{B}^\sigma(Y)$ is a union over spin^c structures $\mathfrak{s}$ of components $\mathcal{B}^\sigma(Y, \mathfrak{s})$ of homotopy type

$$\mathcal{B}^\sigma(Y, \mathfrak{s}) \cong \mathbb{T}(Y) \times \mathbb{CP}^\infty,$$

and there is a canonical identification[4] of cohomology rings

$$H^*(\mathcal{B}^\sigma(Y, \mathfrak{s})) = H^*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[U].$$

Similarly, given a cobordism $W : Y_0 \to Y_1$ equipped with a spin^c structure $\mathfrak{t}$, we have

$$H^*(\mathcal{B}^\sigma(W, \mathfrak{t})) = H^*(\mathbb{T}(W)) \otimes \mathbb{F}_2[U].$$

Letting $\mathfrak{s}_i$ denote the restriction of $\mathfrak{t}$ to Y_i, there is a (partially-defined) correspondence

$$\mathcal{B}^\sigma(Y_0, \mathfrak{s}_0) \leftarrow \mathcal{B}^\sigma(W, \mathfrak{t}) \to \mathcal{B}^\sigma(Y_1, \mathfrak{s}_1)$$

which induces bona fide maps on cohomology rings of the form

$$H^*(\mathbb{T}(Y_0)) \otimes \mathbb{F}_2[U] \xrightarrow{r_0^* \otimes \mathrm{Id}} H^*(\mathbb{T}(W)) \otimes \mathbb{F}_2[U] \xleftarrow{r_1^* \otimes \mathrm{Id}} H^*(\mathbb{T}(Y_1)) \otimes \mathbb{F}_2[U].$$

where r_0^* and r_1^* are the same maps that appear in (7.2). The map $\widetilde{HM}_\bullet(W, \mathfrak{t})$ is natural with respect to the cap product in the following sense: if $\gamma_0 \in H^*(\mathbb{T}(Y_0))$, $\gamma_1 \in H^*(\mathbb{T}(Y_0))$, and

$$r_0^*(\gamma_0) = r_1^*(\gamma_1)$$

[4]The configuration space $\mathcal{B}^\sigma(Y, \mathfrak{s})$ also depends on a choice of metric, but the cohomology ring of the space is independent of the metric up to canonical isomorphism.

then

$$\widetilde{HM}_\bullet(W,\mathfrak{t})((\gamma_0 \otimes U_\dagger^n) \cap x) = (\gamma_1 \otimes U_\dagger^n) \cap \widetilde{HM}_\bullet(W,\mathfrak{t})(x) \tag{7.4}$$

for each $x \in \widetilde{HM}_\bullet(Y_0,\mathfrak{s}_0)$. Here we continue to adopt the convention in [24] of using the subscript $\dagger$ on U to emphasize the module structure with respect to cap product. After completion, $\widetilde{HM}_\bullet(Y)$ and its companions are modules over the completed, graded ring $H^*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger]]$, with $U_\dagger$ in degree -2 and $\gamma \in H^d(\mathbb{T}(Y))$ in degree $-d$. The group $\widetilde{HM}_\bullet(Y)$ inherits the structure of a module over $H^*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U)_\dagger]]$ from $\widetilde{HM}_\bullet(Y)$, in which $U_\dagger$ acts by zero. For this reason, we will consider $\widetilde{HM}_\bullet(Y)$ as a module over $H^*(\mathbb{T}(Y))$.

7.3 Monopole Floer homology and positive scalar curvature

In the presence of positive scalar curvature, this module structure takes a particularly explicit form:

Theorem 7.3.1 (Kronheimer, Mrowka)**.** *Suppose Y admits a metric of (strictly) positive scalar curvature. For each spin^c-structure $\mathfrak{s}$ with $c_1(\mathfrak{s})$ torsion, we have canonical isomorphisms*[5]

$$\begin{aligned}
\widehat{HM}_\bullet(Y,\mathfrak{s}) &= H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger]] \\
\overline{HM}_\bullet(Y,\mathfrak{s}) &= H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}] \\
\widetilde{HM}_\bullet(Y,\mathfrak{s}) &= H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]/\mathbb{F}_2[[U_\dagger]]
\end{aligned}$$

as relatively $\mathbb{Q}$-graded modules over $H^*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger]]$, where the action on the right-hand side is by cap product on the first factor and by multiplication on the second factor. If $c_1(\mathfrak{s})$ is not torsion, then the Floer groups are zero.

Proof. We outline the argument given in Section 36 of [24] (see Proposition 36.1.3 in particular). The structure of $\overline{HM}_\bullet(Y,\mathfrak{s})$ as a module follows from the general structure theorem

[5] Note that $\widehat{HM}_\bullet(Y,\mathfrak{s})$ and $H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger]]$ are rank one, free modules over the same ring, each with a unique generator, and thus canonically isomorphic. The two remaining isomorphisms are specified by insisting on naturality with respect to the short exact sequences of modules associated to the left- and right-hand sides of the three equations.

of Section 35, using the fact the that triple cup product on Y is trivial. Considering the underlying complex $\bar{C}(Y,\mathfrak{s})$, we may regard the perturbation of the equations as a choice of Morse function f on the torus $\mathbb{T}(Y)$. If we replace f with ϵf for sufficiently small ϵ then all critical points are reducible and $\bar{\partial}^u_s = 0$, using the Weitzenböck formula. This formula also rules out spectral flow, which implies that $\bar{\partial}^s_u = 0$ for grading reasons. In particular,

$$\check{C}(Y,\mathfrak{s}) = C^s(Y,\mathfrak{s}), \qquad \check{\partial} = \bar{\partial}^s_s$$

$$\bar{C}(Y,\mathfrak{s}) = C^s(Y,\mathfrak{s}) \oplus C^u(Y,\mathfrak{s}), \qquad \bar{\partial} = \begin{bmatrix} \bar{\partial}^s_s & 0 \\ 0 & \bar{\partial}^u_u \end{bmatrix}$$

$$\hat{C}(Y,\mathfrak{s}) = C^u(Y,\mathfrak{s}), \qquad \hat{\partial} = \bar{\partial}^u_u$$

and the module structures on $\widecheck{HM}_\bullet(Y,\mathfrak{s})$ and $\widehat{HM}_\bullet(Y,\mathfrak{s})$ are inherited via the decomposition

$$\bar{C}(Y,\mathfrak{s}) = \check{C}(Y,\mathfrak{s}) \oplus \hat{C}(Y,\mathfrak{s})\{-1\} \tag{7.5}$$

where $\{-1\}$ shifts the degree of each element of $\hat{C}(Y,\mathfrak{s})$ down by one. □

Remark 7.3.2. Theorem 7.3.1 says that, over a torsion spin^c structure, each version of the monopole Floer homology of Y reduces to Donaldson's TQFT tensor the appropriate $\mathbb{F}_2[[U]]$-module. Theorem 7.3.3 below says that, for certain cobordisms, the Floer maps similarly reduce to those of Donaldson's TQFT as well.

We continue to suppose Y is an oriented, closed, connected 3-manifold with a positive scalar curvature metric. By work of Schoen and Yau, the manifold Y decomposes as the sum of copies of $S^1 \times S^2$ and a rational homology sphere[6]. Let

$$Y_0 = Y$$

and

$$Y_1 = Y\#(S^1 \times S^2).$$

Consider the cobordism

$$W' : Y_0 \to Y_1$$

[6] By Thurston's geometrization conjecture, this rational homology sphere is a sum of spherical space forms. Also note that the sum of positive scalar curvature manifolds admits such a metric as well.

given by 0-surgery on an unknot, and the dual cobordism

$$W'' : Y_1 \to Y_0$$

given by 0-surgery on the circle factor in $S^1 \times S^2$. Let $k = b_1(Y)$, so that the dimensions of $\mathbb{T}(Y_0)$ and $\mathbb{T}(Y_1)$ are k and $k+1$, respectively. The Donaldson maps

$$f_{W'} : H_*(\mathbb{T}(Y_0)) \to H_*(\mathbb{T}(Y_1))$$

and

$$f_{W''} : H_*(\mathbb{T}(Y_1)) \to H_*(\mathbb{T}(Y_0))$$

are non-trivial and determined via the cap product by their values on the fundamental classes Φ_k and Φ_{k+1} of $\mathbb{T}(Y_0)$ and $\mathbb{T}(Y_1)$, respectively (see Remark 7.2.1). Since $f_{W''}$ has degree -1, we have

$$f_{W''}(\Phi_{k+1}) = \Phi_k. \tag{7.6}$$

Let γ' be the generator of the kernel of $r_1^* : H^1(\mathbb{T}(Y_1)) \to H^1(\mathbb{T}(W))$. By (7.3), we have

$$\gamma' \cap f_{W'}(\Phi_k) = f_{W'}(0 \cap \Phi_k) = 0.$$

Since $f_{W'}$ has degree 0, the element $f_{W'}(\Phi_k)$ must generate the 1-dimensional kernel of γ' acting on $H_k(\mathbb{T}(Y_1))$. Thus

$$f_{W'}(\Phi_k) = \gamma' \cap \Phi_{k+1}. \tag{7.7}$$

By a parallel argument, we now determine the Floer maps associated to W' and W''.

Theorem 7.3.3. *Let W be either W' or W''. Let $\mathfrak{t}$ be a spin^c structure on W with $c_1(\mathfrak{t})$ torsion*[7]*. Then with respect to the canonical isomorphisms in Theorem 7.3.1, we have:*

$$\widehat{HM}_\bullet(W, \mathfrak{t}) = f_W \otimes \mathrm{Id}_{\mathbb{F}_2[[\mathrm{U}_\dagger]]} \tag{7.8}$$

$$\overline{HM}_\bullet(W, \mathfrak{t}) = f_W \otimes \mathrm{Id}_{\mathbb{F}_2[[\mathrm{U}_\dagger, \mathrm{U}_\dagger^{-1}]} \tag{7.9}$$

$$\widecheck{HM}_\bullet(W, \mathfrak{t}) = f_W \otimes \mathrm{Id}_{\mathbb{F}_2[[\mathrm{U}_\dagger, \mathrm{U}_\dagger^{-1}]/\mathbb{F}_2[[\mathrm{U}_\dagger]]} \tag{7.10}$$

[7]Note that the topology of W is such that $c_1(\mathfrak{t})$ is torsion if and only if $c_1(\mathfrak{t})|_{\partial W}$ is torsion. In fact, restriction of torsion spin^c structures on W gives a 1-to-1 correspondence between torsion spin^c structures on the incoming and outgoing ends.

Proof. We first establish (7.8). By Theorem 7.3.1, the element $\Phi_{k+i}\otimes 1$ generates $\widehat{HM}_\bullet(Y_i,\mathfrak{s}_i)$ as a rank one, free module over $H_*(\mathbb{T}(Y_i))\otimes\mathbb{F}_2[[U_\dagger]]$. By the compatibility of (7.3) and (7.4), and Remark 7.2.1, it therefore suffices to show:

$$\widehat{HM}_\bullet(W',\mathfrak{t})(\Phi_k\otimes 1) = f_{W'}(\Phi_k)\otimes 1, \tag{7.11}$$

$$\widehat{HM}_\bullet(W'',\mathfrak{t})(\Phi_{k+1}\otimes 1) = f_{W''}(\Phi_{k+1})\otimes 1. \tag{7.12}$$

We will prove (7.12) first, and then use it to prove (7.11). The map $\widehat{HM}_\bullet(W'',\mathfrak{t})$ is surjective since the identity cobordism on Y_0 factors as

$$Y_0\times[0,1] = W''\circ W_1$$

where W_1 is a 1-handle attachment. In particular, $\Phi_{k+1}\otimes 1$ has non-zero image. Since $\Phi_{k+1}\otimes 1$ and $\Phi_k\otimes 1$ each generate the top-most graded piece of their respective Floer groups, and $\widehat{HM}_\bullet(W'',\mathfrak{t})$ has a well-defined degree, it must be the case that

$$\widehat{HM}_\bullet(W'',\mathfrak{t})(\Phi_{k+1}\otimes 1) = \Phi_k\otimes 1, \tag{7.13}$$

which implies (7.12) by (7.6).

Since both W' and W'' have unit Euler characteristic and vanishing intersection form, the associated degrees are $d(W',\mathfrak{t}) = 0$ and $d(W'',\mathfrak{t}) = -1$ by (1.10). So (7.13) implies that

$$\mathrm{gr}^{\mathbb{Q}}(\Phi_k\otimes 1) = \mathrm{gr}^{\mathbb{Q}}(\Phi_{k+1}\otimes 1) - 1$$

and thus

$$\widehat{HM}_\bullet(W',\mathfrak{t})(\Phi_k\otimes 1)\in H_k(\mathbb{T}(Y_1))\otimes 1.$$

Furthermore, the map $\widehat{HM}_\bullet(W',\mathfrak{t})$ is injective since the identity cobordism on Y_0 factors as

$$Y_0\times[0,1] = W_3\circ W'$$

where W_3 is a 3-handle attachment. Thus $\Phi_k\otimes 1$ has non-zero image. Since (7.4) implies

$$(\gamma'\otimes 1)\cap\widehat{HM}_\bullet(W',\mathfrak{t})(\Phi_k\otimes 1) = \widehat{HM}_\bullet(W',\mathfrak{t})((0\otimes 1)\cap(\Phi_k\otimes 1) = 0,$$

the image of $\Phi_k\otimes 1$ must generate the 1-dimensional kernel of $\gamma'\otimes 1$ acting on $H_k(\mathbb{T}(Y_1))\otimes 1$. Thus

$$\widehat{HM}_\bullet(W',\mathfrak{t})(\Phi_k\otimes 1) = (\gamma'\cap\Phi_{k+1})\otimes 1,$$

which implies (7.11) by (7.7).

We now turn to proving (7.9). We may assume that W is equipped with a metric and perturbation that are compatible with the choices made in the proof of Theorem 7.3.1, so that $\bar{\partial}^s_u = \bar{\partial}^u_s = \partial^u_s = 0$ on both Y_0 and Y_1. In this case, the Floer chain maps associated to W only involve reducibles:

$$\hat{m} = \bar{m}^u_u$$

$$\bar{m} = \begin{bmatrix} \bar{m}^s_s & \bar{m}^u_s \\ \bar{m}^s_u & \bar{m}^u_u \end{bmatrix}$$

$$\check{m} = \bar{m}^s_s.$$

Furthermore, in sufficiently low $\mathrm{gr}^{\mathbb{Q}}$ grading, we have

$$\bar{m}(W, \mathfrak{t}) = \begin{bmatrix} 0 & 0 \\ 0 & \bar{m}^u_u(W, \mathfrak{t}) \end{bmatrix}$$

simply because all critical points on either end $(Y_i, \mathfrak{s}_i)$ in sufficiently low $\mathrm{gr}^{\mathbb{Q}}$ grading are boundary unstable and $\bar{m}(W, \mathfrak{t})$ has a well-defined, finite degree. Therefore 7.8 implies (7.9) in sufficiently low $\mathrm{gr}^{\mathbb{Q}}$ grading. Since $U_\dagger$ is invertible on $\overline{HM}_\bullet(Y_1, \mathfrak{s}_1)$, we conclude from the naturality expressed in (7.4) that (7.9) holds in all gradings.

Now (7.9) similarly implies (7.10) in sufficiently high $\mathrm{gr}^{\mathbb{Q}}$ grading. Since $U_\dagger$ is surjective on $\widecheck{HM}_\bullet(Y_0, \mathfrak{s}_0)$, we conclude from (7.4)that (7.10) holds in all gradings. □

Remark 7.3.4. The above argument applies essentially verbatim to the case of $\mathbb{Z}$ coefficients (up to an overall sign determined by homology orientations). Note that (7.11) may also be deduced from (7.12) using a form of duality for cobordism maps (see Proposition 25.5.3 of [24]).

With Y as above, the map $U_\dagger$ is surjective on $\widetilde{HM}_\bullet(Y)$. The exact sequence

$$0 \longrightarrow \widetilde{HM}_\bullet(Y) \longrightarrow \widecheck{HM}_\bullet(Y) \xrightarrow{U_\dagger} \widecheck{HM}_\bullet(Y) \longrightarrow 0$$

then identifies $\widetilde{HM}_\bullet(Y)$ with the submodule $\mathrm{Ker}(U_\dagger) \subset \widecheck{HM}_\bullet(Y)$. From Theorems 7.3.1 and 7.3.3, we have:

Corollary 7.3.5. *Suppose Y admits a metric of (strictly) positive scalar curvature. For each spin^c-structure $\mathfrak{s}$ with $c_1(\mathfrak{s})$ torsion, we have a canonical isomorphism*

$$\widetilde{HM}_\bullet(Y, \mathfrak{s}) = H_*(\mathbb{T}(Y)) \tag{7.14}$$

as relatively $\mathbb{Q}$-graded, rank one, free modules over $H^*(\mathbb{T}(Y))$ acting by cap product. If $c_1(\mathfrak{s})$ is not torsion, then the Floer group is zero.

Let the cobordism W be either W' or W''. Let $\mathfrak{t}$ be a spin^c structure on W with $c_1(\mathfrak{t})$ torsion. Then with respect to (7.14)*, we have*

$$\widetilde{HM}_\bullet(W, \mathfrak{t}) = f_W.$$

So for these cobordisms, the functor $\widetilde{HM}_\bullet$ is canonically identified with Donaldson's TQFT.

Remark 7.3.6. Proposition 6.1.1 is a reinterpretation of Corollary 7.3.5, with the generator of $\Lambda^0 H_1(Y)$ identified with the fundamental class of $H_*(\mathbb{T}(Y))$.

We now specialize to the case of $Y = \#^k(S^1 \times S^2)$. The Floer groups of Y are supported over the unique torsion spin^c structure, so we are free to omit it from the notation. Furthermore, the gradings $\mathrm{gr}^{\mathbb{Q}}$ and $\bar{\mathrm{gr}}^{\mathbb{Q}}$ take values in $\mathbb{Z}$:

Proposition 7.3.7. *Let $Y = \#^k(S^1 \times S^2)$. Then then we have*

$$\widetilde{HM}_\bullet(Y) = H_*(\mathbb{T}(Y))$$

as absolutely $\mathbb{Z}$-graded modules over $H^*(\mathbb{T}(Y))$. More generally, we have

$$\begin{aligned}
\widehat{HM}_\bullet(Y) &= H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger]]\{-1\} \\
\overline{HM}_\bullet(Y) &= H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]\{-2\} \\
\widetilde{HM}_\bullet(Y) &= H_*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]/\mathbb{F}_2[[U_\dagger]]\{-2\}
\end{aligned}$$

as absolutely $\mathbb{Z}$-graded modules over $H^*(\mathbb{T}(Y)) \otimes \mathbb{F}_2[[U_\dagger]]$.

Proof. Let $x \in \widetilde{HM}_\bullet(Y)$ be the generator of lowest degree. We first prove that $\mathrm{gr}^{\mathbb{Q}}(x) = 0$ by induction on $k \geq 0$. This is true by definition in the base case $Y = S^3$. For the induction step, consider the cobordism

$$W' : \#^k(S^1 \times S^2) \to \#^{k+1}(S^1 \times S^2)$$

given by 0-surgery on an unknot. The map $\widetilde{HM}_\bullet(W')$ has degree 0, and sends the lowest degree generator of $\widetilde{HM}_\bullet(\#^k(S^1 \times S^2))$ to that of $\widetilde{HM}_\bullet(\#^{k+1}(S^1 \times S^2))$ by Theorem 7.3.3. This completes the induction.

Let x_0 be the generator of $H_0(\mathbb{T}(Y))$. With respect to the isomorphism in Theorem 7.3.1, the generator $x \in \widetilde{HM}_\bullet(Y)$ corresponds to $x_0 \otimes U_\dagger^{-1}$. The latter has degree 2, which accounts for the shift $[-2]$ in the case of $\widetilde{HM}_\bullet(Y)$. The shifts for $\overline{HM}_\bullet(Y)$ and $\widehat{HM}_\bullet(Y)$ now follow from (7.5). There is no shift for $\widetilde{HM}_\bullet(Y)$ because its lowest degree generator is also x and the degree of x_0 alone is 0. $\square$

Chapter 8

Khovanov homology and $U_\dagger$

In Chapter 6, we considered the $\widetilde{HM}_\bullet$ version of the spectral sequence associated to a diagram of an oriented link $L \subset S^3$. In this chapter, we investigate the other versions of the spectral sequence for L. From the associated framed link $\mathbb{L} \subset -\Sigma(L)$, we may construct a filtered complex $(X, \check{D})$ whose homology is isomorphic to $\widetilde{HM}_\bullet(-\Sigma(L))$. As explained in Remark 2.6.5, we may also construct an even, filtered chain map $\check{U} : X \to X$, which is a sum of components

$$\check{U} = U_0 + U_1 + U_2 + \cdots U_l$$

where U_i shifts $\check{t}$ by i. Let $u^i : E^i \to E^i$ denote the induced filtered chain map on the E^i page. Under the isomorphism $E^\infty \cong \widetilde{HM}_\bullet(-\Sigma(L))$, the map u^∞ is identified with $U_\dagger$. Let $u_\dagger : E^2 \to E^2$ denote the filtration-preserving component of u^2.

Theorem 8.0.8. *For the $\widetilde{HM}_\bullet$ version of the link surgery spectral sequence of $\mathbb{L} \subset -\Sigma(L)$, we have*

$$E^2 \cong \widetilde{Kh}(L) \otimes \mathbb{F}_2[[u_\dagger, u_\dagger^{-1}]/\mathbb{F}_2[[u_\dagger]]$$

as a module over $\mathbb{F}_2[[u_\dagger]]$. The spectral sequence converges by the E^{l+1} page to $\widetilde{HM}_\bullet(-\Sigma(L))$. For each $k \geq 2$, the $(\check{t}, \check{\delta})$-graded vector space E^k depends only on the mutation equivalence class of L. The $\check{\delta}$ grading on the spectral sequence coincides with

$$\delta - \frac{1}{2}(\sigma(L) - \nu(L)) \mod 2$$

on the E^2 page and $gr^{(2)}$ on the E^∞ page. The $\check{t}$ grading on the spectral sequence coincides with $t + n_-$ on the E^2 page. The d^k differential shifts $\check{\delta}$ by 1 and increases $\check{t}$ by k.

Remark 8.0.9. Thereom 8.0.8 holds for $\overline{HM}_\bullet$ and $\widehat{HM}_\bullet$ upon substituting the modules $\mathbb{F}_2[[u_\dagger, u_\dagger^{-1}]$ and $\mathbb{F}_2[[u_\dagger]]$, respectively. For $\overline{HM}_\bullet$, we must also replace $\mathrm{gr}^{(2)}$ with $\bar{\mathrm{gr}}^{(2)}$ and similarly $\check{\delta}$ with its analogue $\bar{\delta}$ defined using $\bar{\mathrm{gr}}^{(2)}$. For $\widehat{HM}_\bullet$, the $\check{\delta}$ grading coincides with

$$\delta + 1 - \frac{1}{2}(\sigma(L) - \nu(L)) \mod 2$$

on the E^2 page. We expect the three versions of the spectral sequence to fit into a long exact sequence of spectral sequences, though have yet to work out the details.

Proof. Let $u_\dagger^1 : E^1 \to E^1$ denote the filtration-preserving component of u^1, a chain map in its own right. By construction, the map $u_\dagger^1$ acts as $U_\dagger$ on each summand $\widetilde{HM}_\bullet(Y_I)$ of the E^1 page. By Theorem 7.3.1, we have

$$E^1 = \widetilde{CKh}(L) \otimes \mathbb{F}_2[[u_\dagger^1, (u_\dagger^1)^{-1}]/\mathbb{F}_2[[u_\dagger^1]]$$

as an $\mathbb{F}_2[[u_\dagger^1]]$-module, and by Theorem 7.3.3 we have

$$d^1 = \partial_{\widetilde{Kh}} \otimes \mathrm{Id}.$$

The chain map $u_\dagger^1 : E^1 \to E^1$ induces the map $u_\dagger$ on $E^2 = H_*(E^1, d^1)$. We conclude that

$$E^2 = \widetilde{Kh}(L) \otimes \mathbb{F}_2[[u_\dagger, u_\dagger^{-1}]/\mathbb{F}_2[[u_\dagger]]$$

as an $\mathbb{F}_2[[u_\dagger]]$-module. The remaining statements follow from the $\widetilde{HM}_\bullet$ version. $\square$

We would expect a Heegaard Floer analogue of $\check{D}$ to be U_0-equivariant by definition, allowing us to set $U_i = 0$ for $i \geq 1$ and replace $u_\dagger$ with u^2 in the theorem. We do not see a way to force this in the monopole setting, although we can still prove the following analog of Corollary 6.0.14:

Corollary 8.0.10. *If $\widetilde{Kh}(L)$ is supported on a single diagonal, then all versions of the spectral sequence collapse at the E^2 page. In this case, for each spin^c structure $\mathfrak{s}$ on $-\Sigma(L)$, we have*

$$\widetilde{HM}_\bullet(-\Sigma(L), \mathfrak{s}) = \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]/\mathbb{F}_2[[U_\dagger]]$$
$$\overline{HM}_\bullet(-\Sigma(L), \mathfrak{s}) = \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]]$$
$$\widehat{HM}_\bullet(-\Sigma(L), \mathfrak{s}) = \mathbb{F}_2[[U_\dagger]]\{1\}$$

as absolutely $\mathbb{Z}/2\mathbb{Z}$-graded modules over $\mathbb{F}_2[[U_\dagger]]$.

Proof. As before, the spectral sequence collapses at E^2 because this page is supported in a single $\check{\delta}$ grading (note that $u_\dagger$ is an even map). The Lee spectral sequence implies that $\widetilde{Kh}(L)$ always has positive rank. So if $\det(L)$ vanishes, then $\widetilde{Kh}^0(L)$ and $\widetilde{Kh}^1(L)$ must have equal, positive rank. From our hypothesis on $\widetilde{Kh}(L)$, we conclude that $\det(L)$ is non-zero, which implies that $-\Sigma(L)$ is a rational homology sphere. Therefore, all spin^c structures $\mathfrak{s}$ on $-\Sigma(L)$ are torsion and the above structure of $\overline{HM}_\bullet(-\Sigma(L), \mathfrak{s})$ is asserted by Proposition 35.3.1 of [24].

In the $\widehat{HM}_\bullet$ case, the map $u_\dagger : E^2 \to E^2$ is injective. As $u_\dagger$ is the filtration-preserving component of the filtered map u^2, the latter map is injective as well. Since $u^2 = u^\infty$, we conclude that $U_\dagger$ is injective on $\widehat{HM}_\bullet(-\Sigma(L))$.

Now the odd map p_* in the exact sequence (1.1) identifies $\widehat{HM}_\bullet(-\Sigma(L), \mathfrak{s})$ with $\mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]\{1\}$ in sufficiently low $\mathrm{gr}^{\mathbb{Q}}$ grading (where $\widetilde{HM}_\bullet(-\Sigma(L), \mathfrak{s})$ vanishes). Let $j_0 \in \mathbb{Q}$ be such a grading. Since $\widehat{HM}_\bullet(-\Sigma(L), \mathfrak{s})$ vanishes in sufficiently high $\mathrm{gr}^{\mathbb{Q}}$ grading and $U_\dagger$ is injective and of degree -2, there must be some even integer k_0 such that

$$\widehat{HM}_{j_0+j}(-\Sigma(L), \mathfrak{s}) = \begin{cases} \mathbb{F}_2 & \text{if } j \leq k_0 \text{ and } j \text{ even} \\ 0 & \text{otherwise} \end{cases}$$

with $U_\dagger$ sending the generator of each $\mathbb{F}_2$ summand down to the next generator. This completes the proof for $\widehat{HM}_\bullet(-\Sigma(L), \mathfrak{s})$. Finally, the exact sequence (1.1) forces the structure of $\widehat{HM}_\bullet(-\Sigma(L), \mathfrak{s})$. □

Recall that the differential $\check{D}$ on X is a sum of components

$$\check{D} = \check{D}^0 + \check{D}^1 + \check{D}^2 + \cdots \check{D}^l$$

where $\check{D}_i$ shifts $\check{t}$ by i. The complex $(X, \check{D}^0)$ is identified with (E^0, d^0). If $\check{D}^0 = 0$, then the complex $(X, \check{D}^1)$ is identified with (E^1, d^1). In fact, we can always guarantee this:

Lemma 8.0.11. *We can choose metric and perturbation data on the spectral sequence in Theorem 8.0.8 so that the d^0 differential vanishes.*

Proof. By Remark 2.5.6, it suffices to show we can choose a metric and perturbation for $\#^k(S^1 \times S^2)$ such that the differential vanishes. This is accomplished by choosing a positive scalar curvature metric and a sufficiently small perturbation modeled on a perfect Morse function on the torus $\mathbb{T}(\#^k(S^1 \times S^2))$. For details, see the proof of Theorem 7.3.1. $\square$

Recall from Remark 2.6.5 that we can arrange for the total complex $(\underline{X}, \underline{D})$ which computes $\widetilde{HM}_\bullet(-\Sigma(L))$ to be identified with the mapping cone of $\check{U} : X \to X$. We may further suppose that $d^0 = 0$ on X, with d^1 identified with $\partial_{\widetilde{Kh}} \otimes \mathrm{Id}$ on

$$E^1 = \widetilde{CKh}(L) \otimes \mathbb{F}_2[[u_\dagger^1, (u_\dagger^1)^{-1}]/\mathbb{F}_2[[u_\dagger^1]].$$

Since $u_\dagger^1 : E^1 \to E^1$ is surjective, the total complex that remains from $\underline{X} = \underline{E}^1$ after applying cancellation to the $u_\dagger^1$ components is identified with $\widetilde{CKh}(\mathcal{D})$, or equivalently $\mathrm{Ker}(\check{U})$ regarded as a quotient complex of X. We conclude:

Corollary 8.0.12. *Let $\mathcal{D}$ be a connected diagram of a link $L \subset S^3$ with n crossings, with mod 2 reduced Khovanov complex $(\widetilde{CKh}(\mathcal{D}), \partial_{\widetilde{Kh}})$. There exists t-filtered total differential*

$$\tilde{D} = \partial_{\widetilde{Kh}} + \tilde{D}^2 + \tilde{D}^3 + \cdots + \tilde{D}^n \tag{8.1}$$

on the vector space $\widetilde{CKh}(\mathcal{D})$ such that we have an isomorphism

$$H_*(\widetilde{CKh}(\mathcal{D}), \tilde{D}) = \widetilde{HM}_\bullet(-\Sigma(L)).$$

This formulation is particularly striking in relation to recent work of Zoltan Szabó [43]. Given a decorated link diagram $\mathcal{D}$, Szabó *combinatorially* defines a differential D_{Sz} of the form (8.1) on $\widetilde{CKh}(\mathcal{D})$. The filtered chain homotopy type of the complex $(\widetilde{CKh}(\mathcal{D}), D_{\mathrm{Sz}})$ is independent of the choice of diagram, and thus provides a new link invariant refining Khovanov homology. Furthermore, this complex has been implemented in *C++* by Cotton Seed, and the resulting spectral sequence coincides precisely with that conjectured in Figure 6.5 for $T(3,5)$ and $T(3,7)$. In fact, all computations to date are consistent with the possibility that the total homology $H_*(\widetilde{CKh}(\mathcal{D}), D_{\mathrm{Sz}})$ has the same rank as $\widetilde{HM}_\bullet(-\Sigma(L))$ and $\widehat{HF}(-\Sigma(L))$. If so, we obtain a new combinatorial algorithm for computing the ranks of these groups.

This leaves open the intriguing possibility that D_{Sz} is an instance of the monopole differential $\tilde{D}$ in Corollary 8.0.12. In fact, Szabó's formulation of D_{Sz} was motivated by this perspective, with the decorations on $\mathcal{D}$ thought of as specifying Morse perturbations in the sense of the proof of Theorem 7.3.1. Note that Lemma 8.0.11 above tells us that there exists a differential

$$\check{D} = (\partial_{\widetilde{Kh}} \otimes \mathrm{Id}) + \check{D}^2 + \cdots \check{D}^l$$

on $\widetilde{CKh}(L) \otimes \mathbb{F}_2[[u_\dagger^1, (u_\dagger^1)^{-1}]/\mathbb{F}_2[[u_\dagger^1]]$ which computes $\widetilde{HM}_\bullet(-\Sigma(L))$. This gives hope, and perhaps guidance, for extending D_{Sz} to incorporate $U_\dagger$ as well.

It therefore seems worthwhile to compare the properties of D_{Sz} with those of $\tilde{D}$ and $\check{D}$. The differential D_{Sz} shifts δ by -1, and thus induces a spectral sequence with a well-defined, invariant integer bigrading (t, δ) on each page. In the next section, we investigate whether $\tilde{D}$ and $\check{D}$ share this property. That is, can we lift our $\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$ bigrading $(\check{t}, \check{\delta})$ to an integer bigrading?

8.1 An integer bigrading?

We continue to consider the total complex $(X, \check{D})$ underlying the link surgery spectral sequence converging to $\widetilde{HM}_\bullet(-\Sigma(L))$. Recall the grading $\check{\delta}^{\mathbb{Q}} : X \to \mathbb{Z}$ defined in (6.4) by

$$\check{\delta}^{\mathbb{Q}}(x) = \mathrm{gr}^{\mathbb{Q}}(x) - (\iota(W_{I\infty}) + w(I)) + l \tag{8.2}$$

which lifts the mod 2 grading $\check{\delta}$. To examine how $\check{D}$ shifts this grading, we must decompose $\check{D}$ over spin^c structures:

$$\check{D} = \sum_{I \leq J} \sum_{\mathfrak{s} \in \mathrm{Spin}^c(W_{IJ})} \check{D}^I_J(\mathfrak{s}).$$

Proposition 8.1.1. *Let $\mathfrak{s}$ be a spin^c structure on W_{IJ}. Then the map $\check{D}^I_J(\mathfrak{s})$ has degree*

$$-1 + \frac{c_1^2(\mathfrak{s})}{4}$$

with respect to the grading $\check{\delta}^{\mathbb{Q}}$ on X.

Proof. The case $I = J$ is clear, so we suppose $I < J$. Since $\sigma(W_{IJ}) = 0$ and the family of metrics P_{IJ} has dimension $w(J) - w(I) - 1$, the map $\check{D}^I_J(\mathfrak{s})$ shifts $\mathrm{gr}^{\mathbb{Q}}$ by

$$\left(\frac{c_1^2(\mathfrak{s})}{4} - \iota(W_{IJ})\right) + w(J) - w(I) - 1 = \left(-1 + \frac{c_1^2(\mathfrak{s})}{4}\right) + (\iota(W_{J\infty}) + w(J)) - (\iota(W_{I\infty}) + w(I)).$$

The claim now follows from 8.2. $\square$

Corollary 8.1.2. *The grading $\check{\delta}^{\mathbb{Q}}$ is well-defined on E^1 and E^2.*

Proof. For $I = J$ or $I < J$ immediate successors, the map $\check{D}^I_J$ is supported over the unique spin^c structure with $c_1(\mathfrak{s}) = 0$. $\square$

This grading would extend to the higher pages if either of the following conjectures held:

Conjecture 8.1.3. *If $c_1^2(\mathfrak{s}) \neq 0$, then $\check{D}^I_J(\mathfrak{s}) = 0$.*

Conjecture 8.1.4. *If $\mathfrak{s}$ and $\bar{\mathfrak{s}}$ are conjugate spin^c structures on W_{IJ}, then $\check{D}^I_J(\mathfrak{s}) = \check{D}^I_J(\bar{\mathfrak{s}})$. Consequently, only the unique spin^c structure with $c_1(\mathfrak{s}) = 0$ gives a net contribution to $\check{D}^I_J$.*

In fact, both of these conjectures, which we entertained for a time, are FALSE. In Section 8.2, we show that there must be a non-trivial component of a higher differential d^i (and thus a component of $\check{D}$) that does not shift $\check{\delta}^{\mathbb{Q}}$ by -1 in the case of the $\widetilde{HM}_\bullet$ spectral sequence for the torus knot $T(3,7)$.

There are two sensible extensions δ' and δ'' of the δ-grading on $\widetilde{CKh}(L)$, depending on whether one takes into account the power of $U_\dagger$:

$$\delta', \delta'' : \widetilde{CKh}(D) \otimes \mathbb{F}_2[[U_\dagger, U_\dagger^{-1}]/\mathbb{F}_2[[U_\dagger]] \to \mathbb{Q}$$
$$\delta'(x \otimes U_\dagger^n) = \delta(x) - 2n - 2$$
$$\delta''(x \otimes U_\dagger^n) = \delta(x) - 2$$

The extra shift of -2 is accounted for by Proposition 7.3.7, which also allows us to regard $x \otimes U_\dagger^n$ as an element of E^1. With respect to this identification, we have

$$\delta'(x \otimes U_\dagger^n) - \check{\delta}^{\mathbb{Q}}(x \otimes U_\dagger^n) = \frac{1}{2}(\sigma(L) - \nu(L))$$
$$\delta''(x \otimes U_\dagger^n) - \check{\delta}^{\mathbb{Q}}(x \otimes U_\dagger^n) = \frac{1}{2}(\sigma(L) - \nu(L)) + 2n$$

by the same argument as in Section 6.1.1, and thus:

Proposition 8.1.5. *With respect to the isomorphism*

$$E^2 \cong \widetilde{Kh}(L) \otimes \mathbb{F}_2[[u_\dagger, u_\dagger^{-1}]/\mathbb{F}_2[[u_\dagger]],$$

the grading $\check{\delta}^{\mathbb{Q}}$ corresponds to the grading

$$\delta' - \frac{1}{2}(\sigma(L) - \nu(L)).$$

While the example of $T(3,7)$ rules out the possibility that $\check{D}$ shifts δ' by -1, this example (as well as all others known to the author) is consistent with:

Conjecture 8.1.6. *The differential $\check{D}$ shifts δ'' by -1. Equivalently (with $D^0 = 0$), the differential $\check{D}^I_J(\mathfrak{s})$ maps $\widetilde{CKh}(\mathcal{D}) \otimes u_\dagger^n$ into $\widetilde{CKh}(\mathcal{D}) \otimes u_\dagger^{n+c_1^2(\mathfrak{s})/8}$.*

In particular, this would imply that $\tilde{D}$ has the same bidegree as ∂_{Sz}:

Conjecture 8.1.7. *The differential $\tilde{D}$ on $\widetilde{CKh}(\mathcal{D})$ shifts δ by -1.*

If these conjectures are indeed simply waiting for a counterexample, the following theorem and corollary may explain why one has yet to be found:

Theorem 8.1.8. *Let $\mathcal{D}$ be a dealternator connected diagram. Then the complex $(\widetilde{CKh}(\mathcal{D}), \tilde{D})$ is t-filtered chain-homotopy equivalent to a t-filtered complex $(\widetilde{Kh}(\mathcal{D}), \tilde{D}')$ with $\tilde{D}'$ a sum of components that shift δ by* negative*, odd integers.*

Proof. The key idea is from the proof of Proposition 5.0.6. We continue from the notation there, setting $E^{1.5} = \overline{E}^1$. We refine the filtration on $(\widetilde{CKh}(\mathcal{D}), \tilde{D})$ so as to visit the page $E^{1.5}$ in between E^1 and E^2. The total complex $(\widetilde{CKh}(\mathcal{D}), \tilde{D})$ is t-filtered chain-homotopy equivalent to the t-filtered total complex $(E^{1.5}, \tilde{D}^{1.5})$, via a sequence of cancellations of bi-degree $(1,-1)$ with respect to (t,δ). Now $E^{1.5}$ is the direct sum of groups $\widetilde{Kh}(\mathcal{D}_{I'})$, each of which is supported in a single δ-grading of $E^{1.5}$. If $I \leq J$, then $I' \leq J'$ and the δ grading of $\widetilde{Kh}(\mathcal{D}_{J'})$ is less than or equal to the δ grading of $\widetilde{Kh}(\mathcal{D}_{I'})$ by (5.1). So the fact that $\tilde{D}^{1.5}$ is positively t-filtered implies that $\tilde{D}^{1.5}$ is negatively δ-filtered. In fact, $\tilde{D}^{1.5}$ is strictly negatively δ-filtered, as the δ-shift is odd by Lemma 2.5.2. Now $(E^{1.5}, \tilde{D}^{1.5})$ is t-filtered chain-homotopy equivalent to the t-filtered total complex $(\widetilde{Kh}(\mathcal{D}), \tilde{D}')$, again via a sequence of cancellations of bi-degree $(1,-1)$, so $\tilde{D}'$ shares the same properties with respect to the δ grading on $\widetilde{Kh}(\mathcal{D})$.

□

Corollary 8.1.9. *If $w_{\widetilde{Kh}}(L) \leq 3$, then there exists a t-filtered differential $\tilde{D}'$ on $\widetilde{Kh}(L)$ which (i) shifts δ by precisely -1 and (ii) yields the $(\check{t}, \check{\delta})$-graded pages E^k associated to L for $k \geq 2$.*

Remark 8.1.10. This theorem and corollary readily extend to the complex $\widetilde{CKh}(\mathcal{D}) \otimes \mathbb{F}_2[[u_\dagger, u_\dagger^{-1}]/\mathbb{F}_2[[u_\dagger]]$ with differential $\check{D}$ and grading δ'', by performing the above cancellations $u_\dagger$-equivariantly in parallel.

So for links with $w_{\widetilde{Kh}}(L) \leq 3$, we will not be able to force a counterexample to Conjecture 8.1.7 using ranks alone. That is, for such links there will always exist higher differentials on $\widetilde{Kh}(L)$ which shift δ by -1 and are compatible with collapsing the rank to that of $\widetilde{HM}_\bullet(-\Sigma(L))$. Note that the first torus link of width four has 18 crossings, while the first torus knot of width four has 20 crossings. This suggests that almost all links with fewer than about 18 crossings have width three or less. The width bound holds for all non-split almost alternating and 2-almost alternating links by Proposition 5.0.6, as well as for links with Turaev genus at most two and thus all pretzel knots, which have Turaev genus at most one.

This rigidity also suggests it may be necessary to compute Szabó's spectral sequence on rather complicated links to find a counterexample to

$$\operatorname{rk} H_*(\widetilde{CKh}(\mathcal{D}), D_{\mathrm{Sz}}) = \operatorname{rk} \widetilde{HM}_\bullet(-\Sigma(L)),$$

if one exists at all.

8.2 The Brieskorn sphere $-\Sigma(2,3,7)$

On the other hand, we now show that the 2-almost alternating knot $T(3,7)$ serves as counterexample to Conjectures 8.1.3 and 8.1.4. Let p, q, and r be pairwise relatively prime positive integers. The Brieskorn integer homology sphere $\Sigma(p,q,r)$ is the intersection of the complex surface $\{x^p + y^q + z^r = 0\} \subset \mathbb{C}^3$ with the unit 5-sphere. The Poincaré homology sphere arises as $\Sigma(2,3,5)$, and more generally the Brieskorn sphere $\Sigma(2,q,r)$ arises as the branched double-cover of the torus knot $T(q,r)$. Note that for integer homology spheres, the rational grading $\mathrm{gr}^{\mathbb{Q}}$ lifts $\mathrm{gr}^{(2)}$ over the unique spin^c structure.

Proposition 8.2.1. *We can choose a basis for $\widetilde{HM}_*(-\Sigma(2,3,7))$ over $\mathbb{F}_2$ such that, with* *respect to the grading $gr^{\mathbb{Q}}$, we have*

$$\widetilde{HM}_j(-\Sigma(2,3,7)) = \begin{cases} \mathbb{F}_2\langle x_0\rangle \oplus \mathbb{F}_2\langle y_0\rangle & \text{if } j=0 \\ \mathbb{F}_2\langle x_j\rangle & \text{if } j = 2,4,6,\dots \\ 0 & \text{otherwise} \end{cases}$$

with $U_\dagger x_j = x_{j-2}$ for $j \neq 0$, and $U_\dagger x_0 = U_\dagger y_0 = 0$. In particular, we have

$$\widetilde{HM}_j(-\Sigma(2,3,7)) = \begin{cases} \mathbb{F}_2\langle x_0\rangle \oplus \mathbb{F}_2\langle y_0\rangle & \text{if } j=0 \\ \mathbb{F}_2\langle y_0'\rangle & \text{if } j = 1 \\ 0 & \text{otherwise} \end{cases}$$

where y_0' represents the generator of the cokernel of $U_\dagger$. Furthermore,

$$\widehat{HM}_j(-\Sigma(2,3,7)) = \begin{cases} \mathbb{F}_2\langle v_0\rangle & \text{if } j=0 \\ \mathbb{F}_2\langle v_j\rangle & \text{if } j = -1,-3,-5,\dots \\ 0 & \text{otherwise} \end{cases}$$

with $U_\dagger v_j = v_{j-2}$ for $j \neq 0$, and $U_\dagger v_0 = 0$. Finally, with respect to the grading $\bar{gr}^{\mathbb{Q}}$, we have

$$\overline{HM}_j(-\Sigma(2,3,7)) = \begin{cases} \mathbb{F}_2\langle u_j\rangle & \text{if } j \text{ is even} \\ 0 & \text{otherwise} \end{cases}$$

with $U_\dagger u_j = u_{j-2}$ for all j.

Proof. The corresponding Heegaard Floer computation is done in Section 8.1 of [35], relying entirely on structural properties (the surgery triangle, rational gradings, and module structure) that carry over to monopole Floer homology. The manifolds $\Sigma(2,3,5)$ and $-\Sigma(2,3,7)$ arise via integral surgery on the left-handed trefoil knot $\overline{T}$:

$$\Sigma(2,3,5) = S^3_{-1}(\overline{T})$$
$$-\Sigma(2,3,7) = S^3_{1}(\overline{T})$$

In the monopole setting, we may also exploit the fact that the Poincaré homology sphere $\Sigma(2,3,5)$ admits a metric of positive scalar curvature. By Theorem 7.3.1, we then have

$$\widetilde{HM}_j(\Sigma(2,3,5)) = \begin{cases} \mathbb{F}_2\langle w_j\rangle & \text{if } j = 2,4,6,\dots \\ 0 & \text{otherwise} \end{cases}$$

with $U_\dagger w_j = w_{j-2}$ for $j \neq 2$, where the grading is pinned down by viewing $\Sigma(2,3,5)$ as the oriented boundary of the negative-definite $E8$ plumbing.

It follows from the surgery exact triangle

$$\cdots \to \widetilde{HM}_\bullet(S^3_{-1}(\overline{T})) \to \widetilde{HM}_\bullet(S^3_0(\overline{T})) \to \widetilde{HM}_\bullet(S^3) \to \cdots$$

that

$$\widetilde{HM}_j(S^3_0(\overline{T})) = \begin{cases} \mathbb{F}_2\langle z_j \rangle & \text{if } j = 1, 2, 3, \ldots \\ 0 & \text{otherwise} \end{cases}$$

with $U_\dagger z_j = z_{j-2}$ for $j \geq 3$, and $U_\dagger z_1 = U_\dagger z_2 = 0$. Now the surgery triangle

$$\cdots \to \widetilde{HM}_\bullet(S^3) \to \widetilde{HM}_\bullet(S^3_0(\overline{T})) \to \widetilde{HM}_\bullet(S^3_1(\overline{T})) \to \cdots$$

determines $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$. In each of the above triangles, the map out of the middle term has degree -1 when the others have degree zero. The first sequence is split, while in the second sequence the map into $\widetilde{HM}_\bullet(S^3)$ identifies y_0 with the bottom generator of $\widetilde{HM}_\bullet(S^3)$. The structure of $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$ now follows from the exact sequence (2.20). The structure of $\overline{HM}_\bullet(-\Sigma(2,3,7))$ is as claimed for all integer homology spheres, and that of $\widehat{HM}_\bullet(-\Sigma(2,3,7))$ now follows from the exact sequence (1.1). □

Remark 8.2.2. Since $\overline{T}$ is a genus 1, fibered knot, the 3-manifold $S^3_0(\overline{T})$ is a bundle over S^1 with fiber T^2. The monodromy is an automorphism of T^2 fixing a point and having order 3. In fact, the monopole Floer homology of this flat 3-manifold is computed directly in Section 37.4 of [24].

Proposition 8.2.3. *In the spectral sequence for $T(3,7)$ converging to $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$ with total complex $(X, \check{D})$, the differential $\check{D}$ does not decreases $\check{\delta}^{\mathbb{Q}}$ (or equivalently, δ') uniformly by -1.*

Proof. First note that since $\sigma(T(3,7)) = 8$, we have $\check{\delta}^{\mathbb{Q}} = \delta' - 4$ on the E^2 page, and both $\check{\delta}^{\mathbb{Q}}$ and δ' mod 2 coincide with $\mathrm{gr}^{(2)}$ via the identification $E^\infty = \widetilde{HM}_\bullet(-\Sigma(2,3,7))$.

We now suppose by contradiction that $\check{D}$ shifts δ' uniformly by -1. Consider the upper left diagram in Figure 8.1, which shows the E^2 page of the $\widetilde{HM}_\bullet$ spectral sequence (ignore

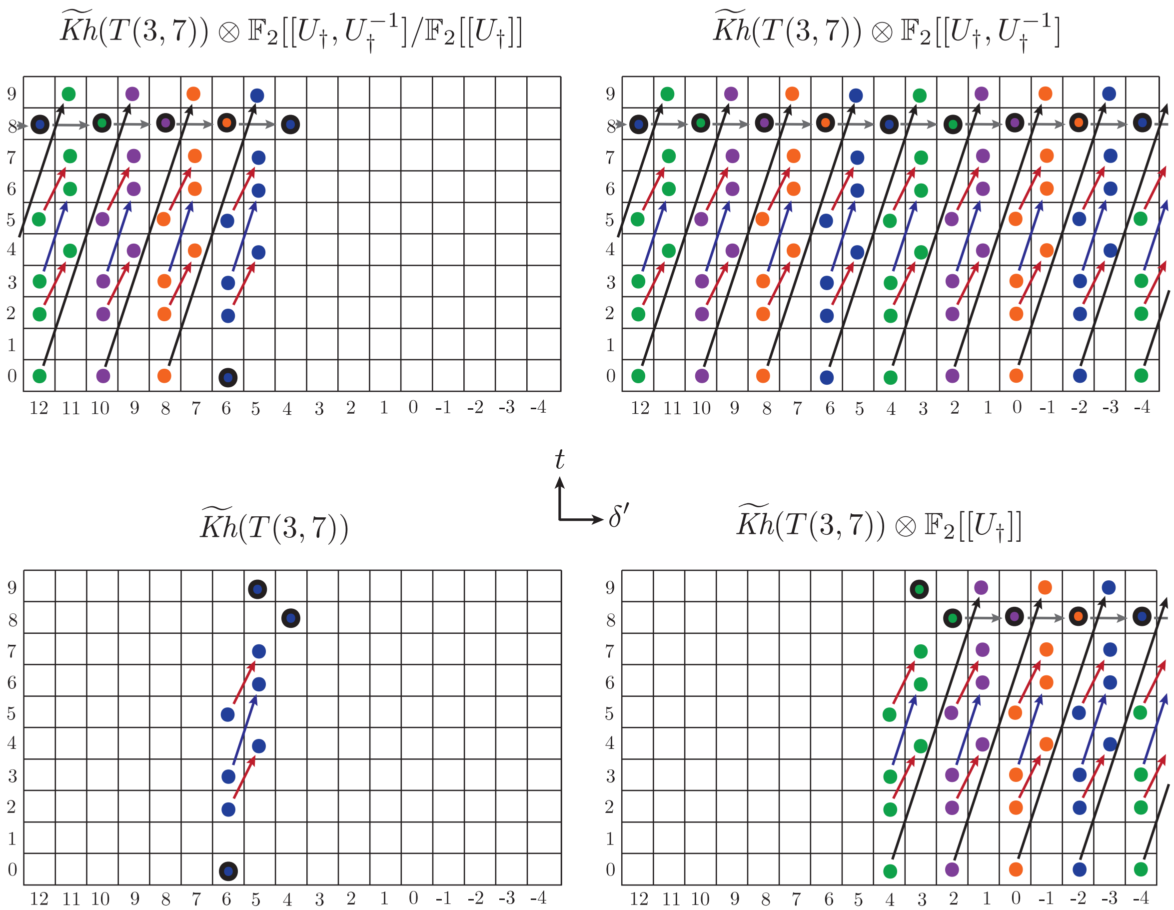

Figure 8.1: Each version of the E^2 page for $T(3,7)$ is shown above, with adjacent copies of $\widetilde{Kh}(T(3,7))$ distinguished by color. We measure the t grading vertically and the δ' grading horizontally (with values *decreasing* from left to right). Clockwise from the upper left, these spectral sequences converge to $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$, $\overline{HM}_\bullet(-\Sigma(2,3,7))$, $\widehat{HM}_\bullet(-\Sigma(2,3,7))$, and $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$. Since $\sigma(T(3,7)) = 8$, the $\mathrm{gr}^{(2)}$ grading on $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$ coincides with δ' mod 2 on the E^∞ page, while that on $\widehat{HM}_\bullet(-\Sigma(2,3,7))$ coincides with $\delta' + 1$ mod 2. We have conjectured differentials so that, in each case, the bolded E^∞ page is consistent with the module structure and mod 2 grading given in Proposition 8.2.1. The residual $U_\dagger$ action on E^∞ is indicated by horizontal gray arrows. The long, black components of the differentials run between adjacent copies of $\widetilde{Kh}(T(3,7)$ and shift δ' by -3. If the horizontal axes measured δ'' in place of δ', then all copies of $\widetilde{Kh}(T(3,7))$ would be superimposed. Indeed, all components of the above differentials shift δ'' by -1.

the conjectural arrows). Since the remaining differentials d^i all increase t by at least 2, there can be no further interaction between distinct copies of $\widetilde{CKh}(\mathcal{D})$. On the other hand, by the argument preceding Corollary 8.0.12, the higher differentials d^i on the rightmost blue copy (regarded as a quotient complex) compute $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$. The latter has rank 1 in odd $\mathrm{gr}^{(2)}$ grading, so one generator in grading $\delta' = 5$ survives to E^∞. This contradicts the fact that $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$ is supported in even grading. □

Remark 8.2.4. If we use a dealternator-connected diagram for $T(3,7)$, such as the right-hand diagram in Figure 5.1, then Theorem 8.1.8 implies that there must be a non-trivial component of a higher differential d^i that decreases $\check{\delta}^{\mathbb{Q}}$ by at least 3. There is therefore some non-trivial component $\check{D}^I_J(\mathfrak{s})$ of $\check{D}$ with $c_1^2(\mathfrak{s}) = -8k < 0$. We have indicated a plausible candidate for each version of the spectral sequence in Figure 8.1. We have also conjectured the shape of the $\widehat{HM}_\bullet$ spectral sequence for torus knots of the form $T(3, 6n \pm 1)$ in Figures 8.2 and 8.3.

The $\check{\delta}$-graded E^∞ page of the spectral sequence is an invariant of the 3-manifold $-\Sigma(L)$. On the other hand:

Proposition 8.2.5. *Suppose that Conjecture 8.1.7 holds, so that the (t, δ) ranks are well-defined invariants of the oriented link L. These bigraded ranks are* not, *in general, invariants of the 3-manifold $-\Sigma(L)$. Nor are the graded ranks of $\widetilde{HM}_\bullet(-\Sigma(L))$ with respect to t, q, or δ individually.*

Proof. Both $T(3,7)$ and the pretzel knot $P(-2,3,7)$ have double branched double-cover $\Sigma(2,3,7)$. While $\widetilde{Kh}(P(-2,3,7))$ and $\widetilde{Kh}(T(3,7))$ both have rank 9, the support of $\widetilde{Kh}(P(-2,3,7))$ only overlaps the support of $\widetilde{Kh}(T(3,7))$ in a single bigrading, where they each have rank 1 (see Figure 8.4). Since $\widetilde{HM}_\bullet(-\Sigma(2,3,7))$ has rank 3, the bigraded ranks of the E^∞ pages must differ in either four or six bigradings.

For a more extreme example, consider Watson's 17-crossing knot $\tau(-1)$ in [46] with $\Sigma(\tau(-1)) = \Sigma(T(5,11)) = \Sigma(2,5,11)$. The rank of $\widetilde{Kh}(\tau(-1))$ is 17 while the rank of $\widetilde{Kh}(T(5,11))$ is 73. The support of $\widetilde{Kh}(\tau(-1))$ is separated from the support of $\widetilde{Kh}(T(5,11))$ by, for example, both the $\{q = 30\}$ horizontal river and the $\{\delta = 10\}$ diagonal river in the (t, q) plane. Furthermore, there is not sufficient overlap after projecting to the t grading to

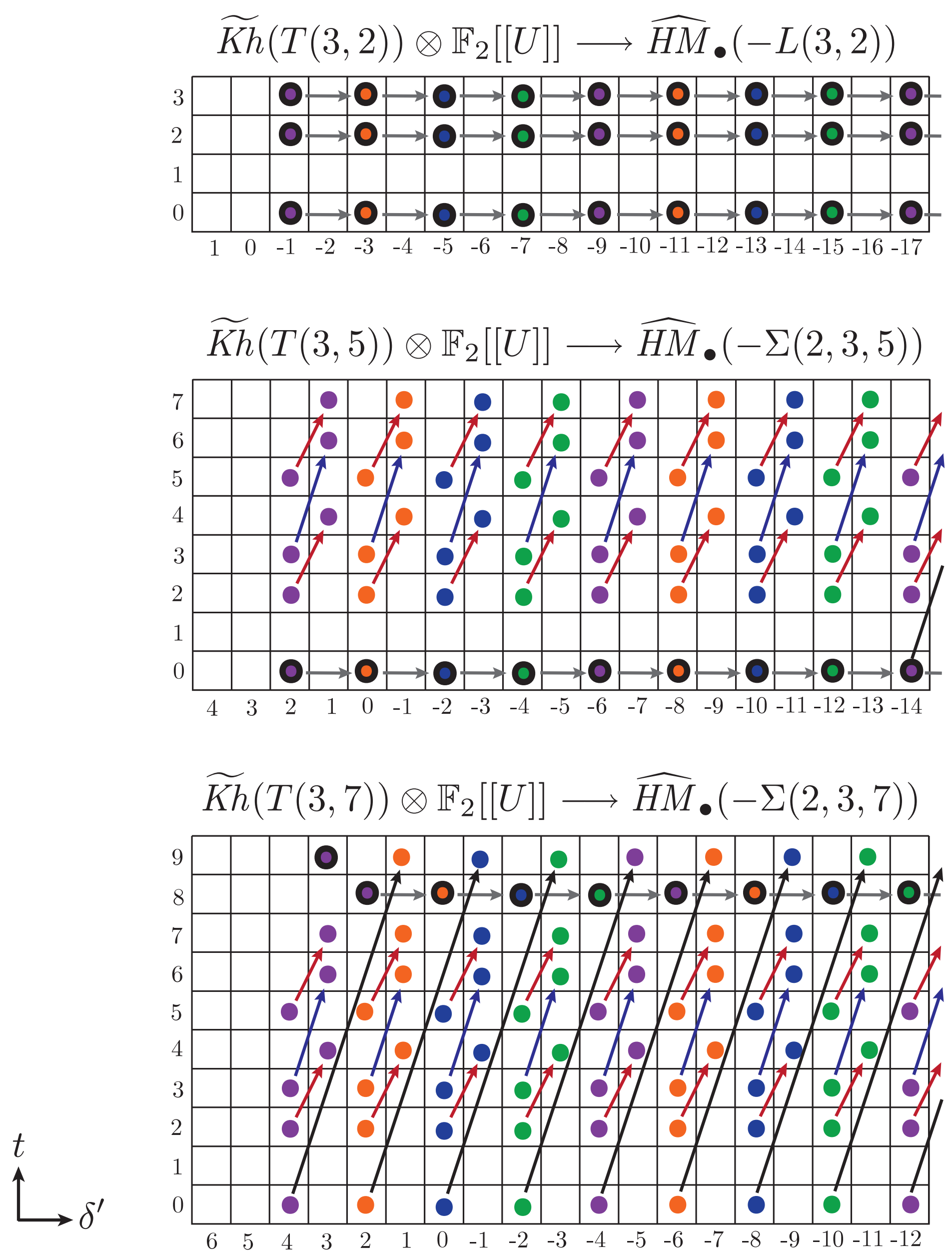

Figure 8.2: Conjectural $\widehat{HM}_\bullet$ spectral sequence for torus knots of the form $T(3, 6n \pm 1)$. These are consistent with Figure 6.5 (keeping in mind the different axes). Versions for $\widetilde{HM}_\bullet$ and $\overline{HM}_\bullet$ may be inferred from Figure 8.1.

$$\widetilde{Kh}(T(3,19)) \otimes \mathbb{F}_2[[U]] \longrightarrow \widehat{HM}_\bullet(-\Sigma(2,3,19))$$

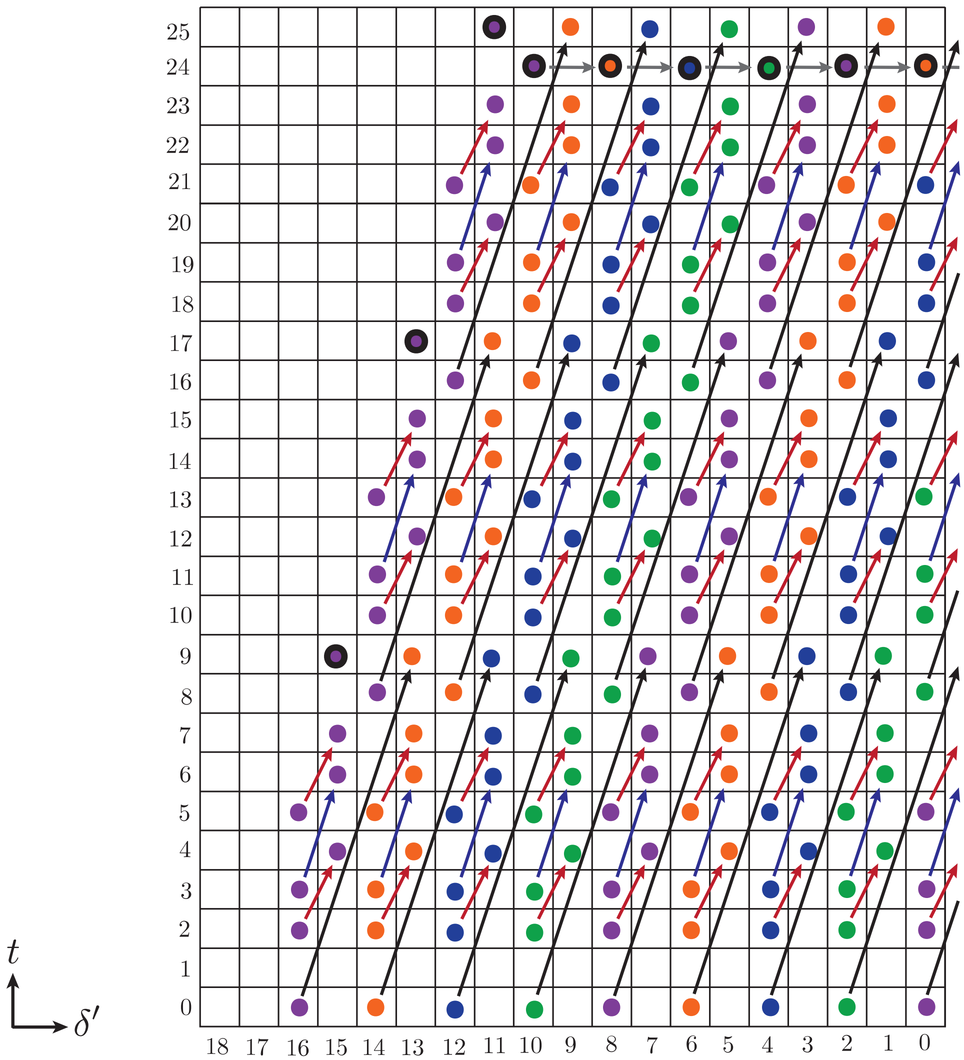

Figure 8.3: Conjectural $\widehat{HM}_\bullet$ spectral sequence for $T(3,19)$, indicating the general pattern for $T(3,6n \pm 1)$ consistent with Conjectures 6.2.1 and 6.2.2. Versions for $\widetilde{HM}_\bullet$ and $\overline{HM}_\bullet$ may be inferred from Figure 8.1.

fit the E^∞ page, which has rank 7 in the Heegaard Floer case by [38] and therefore in the monopole case by [28]. □

Remark 8.2.6. Both $P(-2,3,7)$ and $T(3,7)$ are 2-almost alternating, so by Thereom 8.1.8, we can choose diagrams such that $\check{D}$ shifts δ by -1 uniformly, and thus E^∞ inherits an integer bigrading (although we have not proven that this bigrading is independent of the choice of diagram). Note also that for any knot K, the support of $\widetilde{Kh}(K)$ always intersects the grading $t=0$ due to the s invariant.

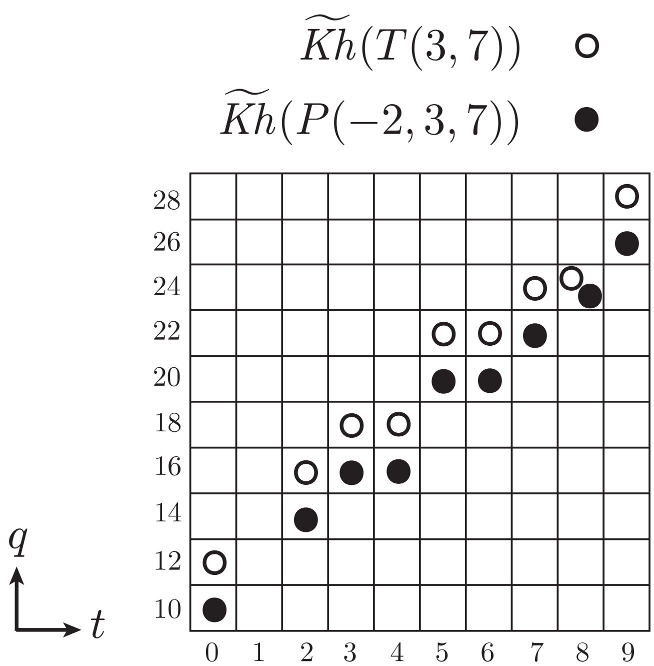

Figure 8.4: Each white dot represents an $\mathbb{F}_2$ summand of $\widetilde{Kh}(T(3,7))$, while each black dot represents and $\mathbb{F}_2$ summand of $\widetilde{Kh}(P(-2,3,7))$. The support of $\widetilde{Kh}(T(3,7))$ overlaps the support of $\widetilde{Kh}(P(-2,3,7))$ in the bigrading $(t,q)=(8,24)$ alone.

8.3 Beyond branched double-covers

We can transport certain observations of a combinatorial nature from branched-double covers to general 3-manifolds using a theorem of Baldwin, which says that, in the sense of the link surgery spectral sequence, all 3-manifolds sit "kitty-corner" from a hypercube of sums of $S^1 \times S^2$. In fact, Baldwin proves a stronger statement:

Lemma 8.3.1 (Baldwin)**.** *Let Y be any closed, connected, oriented 3-manifold. Then there exists a framed link $\mathbb{L} \subset Y$ such that:*

(i) *$Y(I)$ is a connect sum of $S^1 \times S^2$ for all $I \in \{0,1\}^n$.*

(ii) *$W(IJ)$ is the trace of 0-surgery on an unknot or a circle factor in $Y(I)$ for all $I < J$ immediate successors.*

The notation $Y(I)$ and $W(IJ)$ above is as defined prior to Theorem 2.0.1. Baldwin constructs $\mathbb{L}$ from an open book decomposition of Y, with monodromy expressed as a composition of Dehn twists along the generating curves in Figure 5 of [6]. The link $\mathbb{L}$ has one component for each Dehn twist. Such an open book decomposition is associated to the branched double-cover of a braid B given by a diagram with n crossings; in this case, we recover the same framed link $\mathbb{L} \subset -\Sigma(B)$ used to define the spectral sequence for B.

By combining Baldwin's observation with the link surgery spectral sequence from Theorem 2.0.1 with $d^0 = 0$ as in Lemma 8.0.11, we obtain two corollaries:

Corollary 8.3.2. *For general Y, the group $\widetilde{HM}_\bullet(Y)$ arises as the homology of a filtered complex whose underlying vector space E^1 and first non-trivial differential d^1 are defined combinatorially from an open book decomposition of Y. The same is true for $\widehat{HM}_\bullet(Y)$ and $\overline{HM}_\bullet(Y)$.*

Corollary 8.3.3. *For general Y, the group $\widetilde{HM}_\bullet(Y)$ is finitely-generated. In particular, it arises as the homology of a filtered complex $Ker(u^1_\dagger)$ whose underlying vector space is finitely-generated and combinatorially-defined. The first non-trivial differential $\tilde{d}^1$ is combinatorial as well.*

Remark 8.3.4. There is an alternative definition of $\widetilde{HM}_\bullet$ from the perspective of sutured Floer homology [25], as explained at the end of Section 2. In this case, the underlying spin^c structures are all non-torsion, so the complex is a priori finitely-generated, although not combinatorially-defined.

Particularly in light of Szabo's combinatorial spectral sequence, we hope that these observations will yield new approaches to computing monopole Floer homology and insights into its axiomatization and relationship with Heegaard Floer homology.

Chapter 9

Appendix: Morse homology with boundary via path algebras

Monopole Floer homology may be viewed as an infinite dimensional version of Morse homology for manifolds with boundary. For a beautiful treatment of the finite dimensional model, see Section 2 of [24]. We now give a brief presentation of its essential features, assuming familiarity with Morse homology for closed manifolds. By recasting the combinatorics in terms of path algebras, we hope to illuminate the classification of ends in Lemma 2.2.3 and the form of the matrices (2.5), (2.7), (2.10), (2.11), and (2.13) used to define $\check{D}^I_J$, $\check{L}$, and $\check{A}^I_J$.

Consider a manifold M with boundary ∂M, equipped with a (sufficiently generic) Morse function and metric that extend equivariantly to the double. In particular, the gradient vector field is tangent along ∂M. The critical points in the boundary are classified as stable or unstable, according to whether the flow in the normal direction is toward or away from the boundary, respectively. We denote interior, boundary-stable, and boundary-unstable critical points by o, s, and u, respectively. Note that interior gradient trajectories always flow from o or u to o or s, whereas boundary trajectories flow from s or u to s or u. We distinguish between interior and boundary trajectories from u to s, so there are eight types in all.

On the surface M in Figure 9.1, we have marked one isolated gradient trajectory for

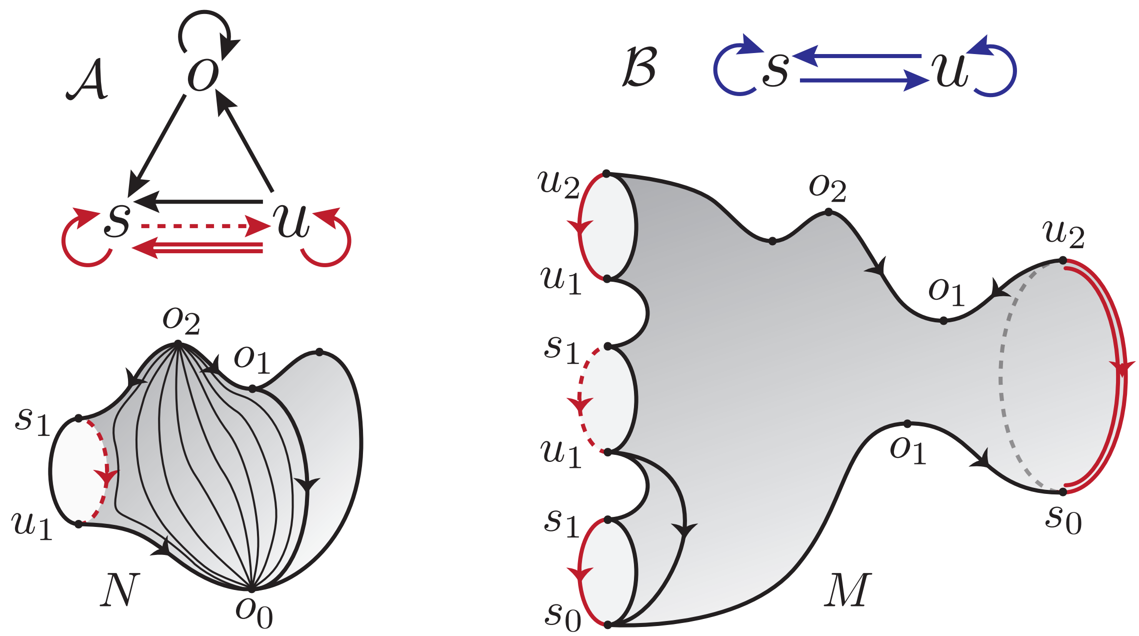

Figure 9.1: Path algebras and Morse homology for manifolds with boundary.

each of these eight types, where those in ∂M (in red) are isolated with respect to ∂M. The subscripts on the critical points denote Morse index with respect to M. While most types of isolated trajectories lower Morse index by 1, there are two exceptions. The doubled trajectory from u to s in ∂M lowers Morse index by 2, while the dashed trajectory from s to u in ∂M fixes Morse index. This last type is called *boundary-obstructed*.

All of this information may be neatly encoded in a path algebra over $\mathbb{F}_2$, denoted $\mathcal{A}$ (this was first pointed out to the author by Dylan Thurston). As an $\mathbb{F}_2$-vector space, $\mathcal{A}$ has a basis given by the set of all paths in the directed graph at left in Figure 9.1. The product of two paths is given by concatenation if the target of the first coincides with the source of the second, and is zero otherwise. The weight of a path is the sum of the weights of its edges, where the dashed, single, and doubled edges have weights 0, 1, and 2, respectively.

If we consider ∂M as a closed manifold in its own right, then the Morse index of each boundary-unstable critical point is one less. So now all four types of isolated trajectories in ∂M lower Morse index by 1, as encoded in the path algebra $\mathcal{B}$ in Figure 9.1.

The groups $H_*(\partial M)$, $H_*(M)$, and $H_*(M, \partial M)$ arise from the Morse complex generated by critical points of types $\{s, u\}$, $\{o, s\}$, and $\{o, u\}$, respectively. The correspondence with the monopole Floer groups is reflected by the exact sequences

$$\cdots \longrightarrow H_*(\partial M) \longrightarrow H_*(M) \longrightarrow H_*(M, \partial M) \longrightarrow \cdots$$
$$\cdots \longrightarrow \overline{HM}_\bullet(Y) \longrightarrow \widetilde{HM}_\bullet(Y) \longrightarrow \widehat{HM}_\bullet(Y) \longrightarrow \cdots .$$

Since we are primarily concerned with $\widetilde{HM}_\bullet(Y)$, we focus on the absolute case $H_*(M)$. The Morse complex then has the form

$$C(M) = C^o(M) \oplus C^s(M).$$

The differential ∂ may be thought of as an element of $\mathcal{A}$, given by the sum of all weight 1 paths from $\{o, s\}$ to $\{o, s\}$, as depicted in Figure 9.2. In matrix form, this becomes

$$\partial = \begin{bmatrix} \partial^o_o & \partial^u_o \bar{\partial}^s_u \\ \partial^o_s & \bar{\partial}^s_s + \partial^u_s \bar{\partial}^s_u \end{bmatrix}.$$

We introduce an ideal $\mathcal{I}$ of $\mathcal{A}$, generated by the eight elements in the second and third rows of Figure 9.2. We have one relation for each interior (black) generator of $\mathcal{A}$, given by the sum all paths of weight 2 between its ends. We similarly have one relation for each boundary (red) generator of $\mathcal{A}$, given by the sum all paths of weight 2 between the ends of the corresponding (blue) generator in $\mathcal{B}$. These relations correspond precisely to the maps counting the ends of 1-dimensional moduli spaces, and can be expressed in that form as

$$\begin{aligned} A^o_o &= \partial^o_o \partial^o_o + \partial^u_o \bar{\partial}^s_u \partial^o_s & \bar{A}^s_s &= \bar{\partial}^s_s \bar{\partial}^s_s + \bar{\partial}^u_s \bar{\partial}^s_u \\ A^o_s &= \partial^o_s \partial^o_o + \bar{\partial}^s_s \partial^o_s + \partial^u_s \bar{\partial}^s_u \partial^o_s & \bar{A}^s_u &= \bar{\partial}^s_u \bar{\partial}^s_s + \bar{\partial}^u_u \bar{\partial}^s_u \\ A^u_o &= \partial^o_o \partial^u_o + \partial^u_o \bar{\partial}^u_u + \partial^u_o \bar{\partial}^s_u \partial^u_s & \bar{A}^u_s &= \bar{\partial}^s_s \bar{\partial}^u_s + \bar{\partial}^u_s \bar{\partial}^u_u \\ A^u_s &= \bar{\partial}^u_s + \partial^o_s \partial^u_o + \bar{\partial}^s_s \partial^u_s + \partial^u_s \bar{\partial}^u_u + \partial^u_s \bar{\partial}^s_u \partial^u_s & \bar{A}^u_u &= \bar{\partial}^s_u \bar{\partial}^u_s + \bar{\partial}^u_u \bar{\partial}^u_u \end{aligned}$$

We have illustrated two broken trajectories counted by the map A^o_o on the surface N in Figure 9.1. The 1-dimensional family of interior trajectories from o to o has one end with two components and another end with three components, where the middle component is boundary-obstructed. Note that the terms in the above relations correspond precisely to those described in Lemma 2.2.3.

We next define a coboundary operator $\delta : \mathcal{A} \to \mathcal{A}$ which acts on edges by sending ∂^*_* to A^*_* and $\bar{\partial}^*_*$ to $\bar{A}^*_*$. We extend δ to paths by the Leibniz rule and to $\mathcal{A}$ linearly. Note that $\mathcal{I}$

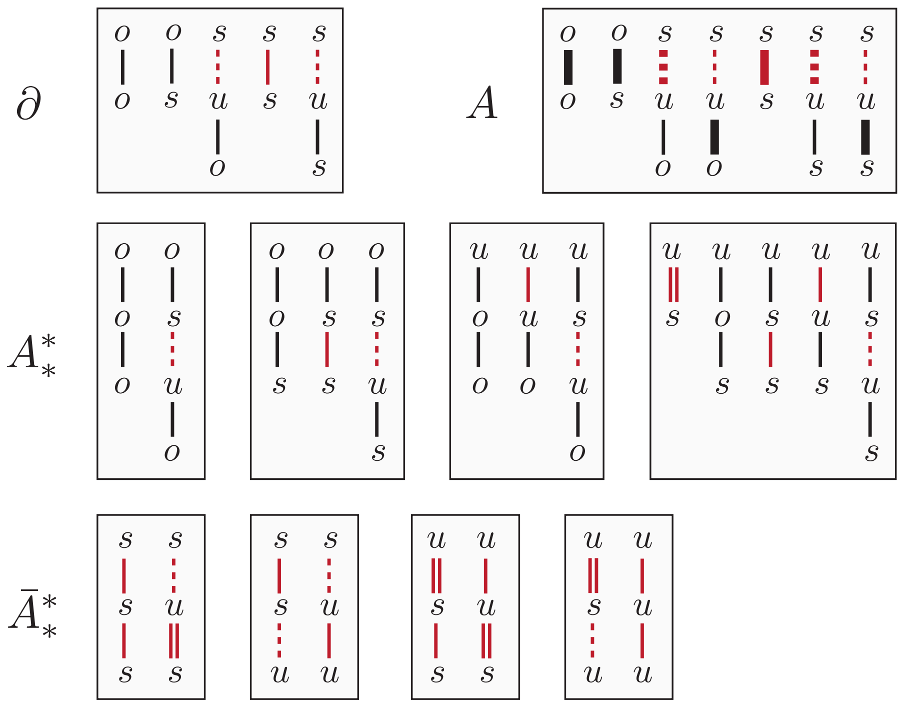

Figure 9.2: The differential ∂ and identity A as elements of $\mathcal{A}$. The bold line from o to o is shorthand for A_o^o, and similarly for the other bold lines. Thus, A lies in the ideal of $\mathcal{A}$ generated by the eight relations of the form A_*^* and $\bar{A}_*^*$.

is generated by the image of δ. Let $A \in \mathcal{A}$ be the image of ∂ under δ. In other words, A is the sum of seven elements, each the result of replacing one edge in ∂ with the corresponding relation. This is illustrated at the top right of Figure 9.2, with the relations bolded. As a map, A is given by

$$A = \begin{bmatrix} A_o^o & A_o^u \bar{\partial}_u^s + \partial_o^u \bar{A}_u^s \\ A_s^o & \bar{A}_s^s + A_s^u \bar{\partial}_u^s + \partial_s^u \bar{A}_u^s \end{bmatrix}$$

Now it is easy to check that A and ∂^2 coincide as elements of the path algebra $\mathcal{A}$. One observes cancellation of precisely those paths with no interior o or s (i.e., no good break). Thus ∂^2 is in the ideal generated by the relations as well, with the implication being that ∂ is a differential on the Morse complex.

Remark 9.0.5. The author and Dave Bayer wrote a program in *Haskell* which formally implements the path algebra associated to a cobordism equipped with a permutohedron of metrics as in Section 2.1. Indeed, the program verifies Lemma 2.2.7 in this language. The case of a single metric is illustrated in Figure 9.3. The maps $\check{\partial}_0^0 = \check{D}_0^0$, $\check{m}_1^0 = \check{D}_1^0$, and $\check{\partial}_1^1 = \check{D}_1^1$ may be thought of as elements of the weighted path algebra $\mathcal{A}_1^0$ over the red and black graph with 24 edges and 6 vertices. There is one relation for each black edge in $\mathcal{A}_1^0$ and one relation for each blue edge in $\mathcal{B}_1^0$, each consisting of all paths of weight 2 between the corresponding ends. The coboundary map $\delta : \mathcal{A}_1^0 \to \mathcal{A}_1^0$ has image lying in the ideal $\mathcal{I}_1^0$ generated by these relations. Setting $\check{A}_1^0 = \delta \check{m}_1^0$, the computation $\check{m}_1^0 \check{\partial}_0^0 + \check{\partial}_1^1 \check{m}_1^0 = \check{A}_1^0 \in \mathcal{I}_1^0$ verifies that $\check{m}_1^0 : \check{C}(Y_0) \to \check{C}(Y_1)$ is a chain map. Similarly, $(\check{\partial}_i^i)^2 = \delta \check{\partial}_i^i \in \mathcal{I}_1^0$ for $i = 0, 1$.

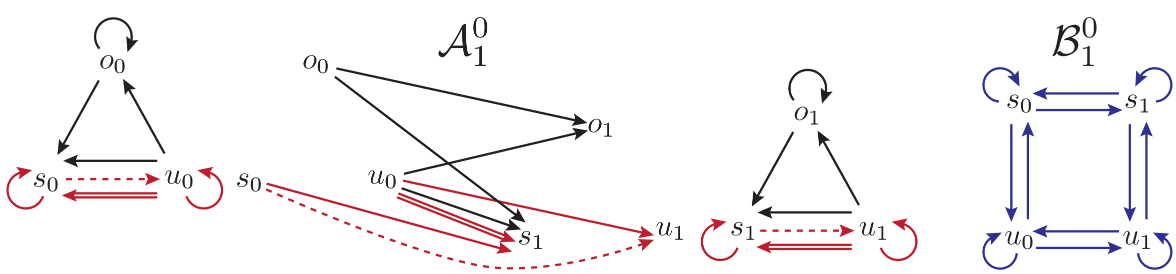

Figure 9.3: The path algebra of a cobordism $W : Y_0 \to Y_1$ with fixed metric.

Remark 9.0.6. Kronheimer and Mrowka discuss functoriality in Morse homology in Section 2.8 of [24]. The above path algebra interpretation generalizes to this setting and indeed organizes the combinatorics necessary to define the Morse category of a manifold with boundary. In the monopole Floer setting, we may extend the path algebra interpretation to cobordisms equipped with polytopes of metrics and multiple incoming and outgoing ends, by including one copy of o, s, and u for each end or interior hypersurface. A map which counts unbroken trajectories on such a cobordism is represented by an "edge" whose source and target are subsets of vertices. The notion of boundary-obstructedness generalizes in a natural manner to determine the weight of such an edge (see *boundary-obstructed of corank* c in Section 24.4 of [24]). It turns out that in the "to" (resp., "from") theory, we must restrict to cobordisms with exactly one incoming (resp., outgoing) end. We will elaborate on this construction in future work, but as a prelude, we give an example with two incoming ends and one outgoing end in Figure 9.4.

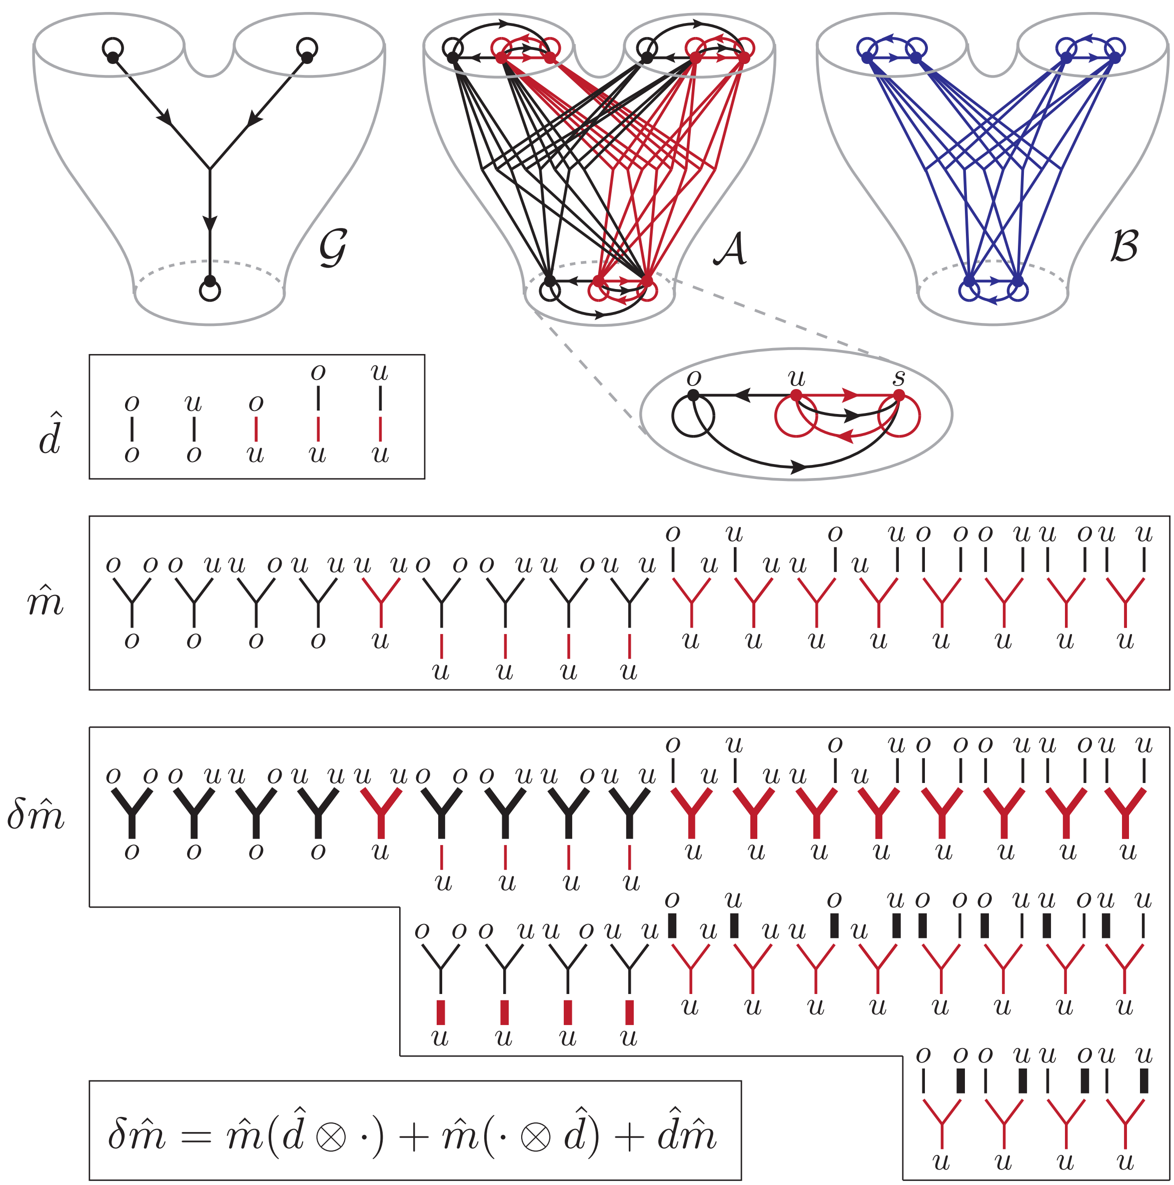

Figure 9.4: Here we represent the components of the chain map $\hat{m}$ associated to a cobordism with two incoming ends and one outgoing end as an element of the "path algebra" on the weighted directed hypergraph $\mathcal{A}$. The left-hand side of the above equation expands to 197 terms, of which the 62 terms with no good break cancel in pairs, leaving the 135 terms which arise on the right-hand side.

Bibliography

[1] T. Abe and K. Kishimoto. The dealternating number and the alternation number of a closed 3-braid. 2008, arXiv:0808.0573v2.

[2] C. Adams and J. Brock et. al. Almost alternating links. *Topology Appl.*, 46:151–165, 1992.

[3] M. Asaeda and J. Przytyck. Khovanov homology: torsion and thickness. In J.Bryden, editor, *New Techniques in Topological Quantum Field Theory*, Calgary/Kananaskis, Canada, August 2001. arXiv:math/0402402v2.

[4] The Knot Atlas. http://katlas.org/.

[5] J. Baldwin. On the spectral sequence from Khovanov homology to Heegaard Floer homology. 2008, arXiv:0809.3293.

[6] J. Baldwin and O. Plamenevskaya. Khovanov homology, open books, and tight contact structures. Preprint (2008), http://arxiv.org/abs/0808.2336.

[7] D. Bar-Natan. Mutation invariance of Khovanov homology. *The Knot Atlas (2005).* `http://katlas.org/drorbn/index.php?title=Mutation_Invariance_of_Khovanov_Homology`.

[8] J. Bloom. A link surgery spectral sequence in monopole Floer homology. To appear in Advances in Mathematics, preprint (2009), http://arxiv.org/abs/0909.0816.

[9] J. Bloom. Odd Khovanov homology is mutation invariant. *Math. Res. Lett.*, 17(1):1–10, 2010, arXiv:0903.3746.

[10] J. Bloom, T. S. Mrowka, and P. Ozsváth. The Kunneth principle in Floer homology. In preparation.

[11] P.J. Braam and S. K. Donaldson. Floer's work on instanton homology, knots, and surgery. In H. Hofer, editor, *The Floer Memorial Volume*, volume 133 of *Progress in Math.*, pages 195–256. Birkhüauser, 1995.

[12] M. Carr and S. Devadoss. Coxeter complexes and graph-associahedra. *Topology Appl.*, 153:2155–2168, 2006.

[13] W. Casselman. Strange associations. www.ams.org/featurecolumn/archive/associahedra.html, 2006.

[14] S. V. Chmutov and S. K. Lando. Mutant knots and intersection graphs. 2007, arXiv:0704.1313v1.

[15] S. Devadoss. A realization of graph-associahedra. *Discrete Math.*, 309:271–276, 2009.

[16] S. K. Donaldson. Topological field theories and formulae of Casson and Meng-Taubes. In *Proceedings of the Kirbyfest*, volume 2 of *Geometry and Topology Monographs*, pages 87–102, 1999.

[17] N. Dunfield, S. Garoufalidis, A. Shumakovitch, and M. Thistlethwaite. Behavior of knot invariants under genus 2 mutation. Preprint (2006), http://arxiv.org/abs/math/0607258.

[18] A. Floer. Instanton homology, surgery, and knots. In S. K. Donaldson and C.B. Thomas, editors, *Geometry of Low-Dimensional Manifolds: Gauge Theory and Algebraic Surface*, volume 150 of *London Mathematical Society Lecture Note Series*, pages 97–114. Cambridge University Press, 1990.

[19] R. Gompf and A. Stipsicz. *4-Manifolds and Kirby Calculus*. Graduate Studies in Math. AMS, 1999.

[20] C. McA. Gordon and R. A. Litherland. On the signature of a link. *Invent. Math.*, 47:53–69, 1978.

[21] J. Greene. A spanning tree model for the Heegaard Floer homology of a branched double-cover. 2008, arxiv/0805.1381v1.

[22] M. Khovanov. A categorification of the Jones polynomial. *Duke Math. J.*, 101(3):359–426, 2000.

[23] I. Kofman and A. Champanerkar. On mutation and Khovanov homology. *Communications in Contemporary Mathematics*, 10:973–992, 2008.

[24] P. Kronheimer and T. Mrowka. *Monopoles and Three-Manifolds*. New Mathematical Monographs. Cambridge University Press, 2007.

[25] P. Kronheimer and T. Mrowka. Knots, sutures and excision. 2008, arXiv:0807.4891v2.

[26] P. Kronheimer and T. Mrowka. Khovanov homology is an unknot-detector. 2010, arXiv:1005.4346v1.

[27] P. Kronheimer, T. Mrowka, P. Ozsváth, and Z. Szabó. Monopoles and lens space surgeries. *Ann. of Math. (2)*, 165(2):457–546, 2007, math.GT/0310164.

[28] C. Kutluhan, Yi-Jen Lee, and C. H. Taubes. HF = HM I : Heegaard Floer homology and Seiberg-Witten floer homology. 2010, arxiv.org/abs/1007.1979.

[29] J. Loday. Realization of the Stasheff polytope. *Arch. Math. (Basel)*, 83(3):267–278, 2004.

[30] C. Manolescu and P. Ozsváth. On the Khovanov and knot Floer homologies of quasi-alternating links. In *Proceedings of the 14th Gökova Geometry-Topology Conference*, pages 60–81, 2007.

[31] J. McCleary. *A User's Guide to Spectral Sequences*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2 edition, 2000.

[32] H. R. Morton and N. Ryder. Invariants of genus 2 mutants. *J. Knot Theory Ramifications*, 18:1423–1438, 2009.

[33] T. Mrowka, P. Ozsváth, and B. Yu. Seiberg-Witten monopoles on Seifert fibered spaces. *MSRI 1996-093*, 1996, arXiv:math/9612221v1.

[34] P. Ozsváth, J. Rasmussen, and Z. Szabó. Odd Khovanov homology. 2007, math.QA/0710.4300.

[35] P. Ozsváth and Z. Szabó. Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary. *Adv. Math.*, 173:179–261, 2003, math.SG/0110170.

[36] P. Ozsváth and Z. Szabó. On the Heegaard Floer homology of branched double-covers. *Adv. Math.*, 194(1):1–33, 2005, math.SG/0309170.

[37] L.P. Roberts. Notes on the Heegaard-Floer link surgery spectral sequence. 2008, math.GT/0808.2817.

[38] R. Rustamov. Calculation of Heegaard Floer homology for a class of Brieskorn spheres. 2003, arXiv:math/0312071v1 [math.SG].

[39] S. Saneblidze and R. Umble. Diagonals on the permutahedra, multiplihedra, and associahedra. *Homology, Homotopy Appl.*, 6(1):363–411, 2004.

[40] P. Seidel. *Fukaya categories and Picard-Lefschetz theory*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, 2008.

[41] M. Stošíc. Homology of torus links. 2007, arXiv:math/0606656v2.

[42] J. Stasheff. Homotopy associativity of H-spaces I. *Trans. Amer. Math. Soc.*, 108:275–292, 1963.

[43] Z. Szábo. A geometric spectral sequence in Khovanov homology. 2010, arXiv:1010.4252v1 [math.GT].

[44] P. Traczyk. A combinatorial formula for the signature of alternating diagrams. *Fund. Math.*, 184:311–316, 2004.

[45] P. Turner. A spectral sequence for Khovanov homology with an application to (3,q)-torus links. 2006, arXiv:math/0606369v1.

[46] L. Watson. A remark on Khovanov homology and two-fold branched covers. 2008, arXiv:0808.2797.

[47] L. Watson. New proofs of certain finite filling results. 2010, arXiv:1010.3051v1 [math.GT].

[48] S. Wehrli. Mutation invariance of Khovanov homology over $\mathbb{F}_2$. Preprint (2009), http://arxiv.org/abs/0904.3401.

[49] S. Wehrli. Khovanov homology and Conway mutation. 2003, math.GT/0409328.

[50] G. Ziegler. *Lectures on polytopes.* Graduate Texts in Mathematics. Springer, 2010.