

HIGHER DIFFERENTIALS ON KHOVANOV
HOMOLOGY

COTTON SEED

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
MATHEMATICS
ADVISER: ZOLTÁN SZABÓ

JUNE 2014

© Copyright by Cotton Seed, 2014.

All Rights Reserved

Abstract

In this thesis, we study the structure and geometric content of Khovanov homology using higher differentials. We study the Szabó geometric spectral sequence and conjecture that it agrees with the spectral sequence from Khovanov homology to the Heegaard Floer homology of the double-branched cover of a knot. We define a twisted variant of the geometric spectral sequence, connect it to Baldwin-Ozsváth-Szabó homology, and outline a strategy towards the above conjecture. We construct a new spectral sequence that begins at the Khovanov homology of a link and converges to the Khovanov homology of the disjoint union of its components. The page at which the spectral sequence collapses gives a lower bound on the splitting number of the link, the minimum number of times its components must be passed through one another in order to completely separate them. In addition, we build on work of Kronheimer-Mrowka and Hedden-Ni to show that Khovanov homology detects the unlink.

Acknowledgements

I would like to express my heartfelt thanks to the community that supported this work and, more broadly, my mathematical and intellectual development. First and foremost, I would like to thank my advisor Zoltán Szabó for his constant advice, encouragement and support. I would like to thank Josh Batson for his friendship and the fruitful collaboration that produced the material of Chapter 5. I found the community of mathematicians working on low-dimensional topology remarkably welcoming and supportive, and special thanks go to Jon Bloom, John Baldwin, Dave Gabai, Eli Grigsby, Mikhail Khovanov, Robert Lipshitz, Yi Ni, Krzysztof Putyra, Sucharit Sarkar, Alexander Shumakovtich and Bohua Zhan. I would like to thank Will Cavendish and Sam Lewallen for their friendship, support and for many interesting conversations. Also, for the beer and cheese. I am very grateful to Jill LeClair, the graduate program administrator, for her invaluable help. I would also like to thank Pete Clark, John Duncan, Dennis Gaitsgory and, especially, John Francis. If it hadn't been for their early influence and inspiration, I never would have found the path to academic mathematics. Finally, I want to thank Stephanie Frampton for being there throughout.

The author was partially supported by NSF grants DMS-0704053, DMS-0603940, DMS-1006006, DMS-1258274 and EMSW21-RTG during the writing of this thesis.

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Results	4
2 Background	12
2.1 A review of Khovanov homology	12
2.2 A review of the geometric spectral sequence	16
2.3 A review of Baldwin-Ozsváth-Szabó homology	21
3 Computing the Geometric Spectral Sequence	27
3.1 Computing the geometric spectral sequence	27
3.2 Results	30
3.3 Conjectures	41
4 Twisting the Geometric Spectral Sequence	45
4.1 The twisted chain complex	45
4.2 Dependence on the decoration	53
4.3 d is a differential	56
4.4 Total homology	69

5	The Link Splitting Spectral Sequence	72
5.1	Construction	73
5.1.1	Chain complex	73
5.1.2	Change of sign assignment	79
5.1.3	Total homology	80
5.2	Reidemeister invariance	83
5.3	Properties of spectral sequences	94
5.3.1	Endomorphisms of spectral sequences	97
5.4	The splitting number	98
5.5	Detecting unlinks	102
5.5.1	The Khovanov module	103
5.5.2	Proof of Theorem 1.1.10	105
5.6	Sample computations	108
	Bibliography	112

Chapter 1

Introduction

The purpose of this thesis is to investigate the structural properties and geometric content of Khovanov's knot homology theory using higher differentials. We begin with a brief history of Khovanov homology.

In 1984, while studying von Neumann algebras, Jones [26] discovered a novel polynomial knot invariant, the Jones polynomial. In the late 19th century, Tait initiated the mathematical study of knot theory inspired by the conjecture of Thomson (aka Lord Kelvin) [28] that atoms were topologically knotted tubes of aether. Tait, after tabulating knots with at most 10 crossings, made several important conjectures [67] concerning alternating knot diagrams. As an illustration of the power of the Jones polynomial, within a few years of its introduction and nearly a century after Tait, it was used to prove the Tait conjectures. The first two Tait conjectures were established independently by Kauffman [27], Murasugi [51, 52] and Thistlethwaite [68] and the third, the flyping conjecture, by Menasco and Thistlethwaite [50].

Khovanov [29], in 1999, gave a vast generalization of the Jones polynomial by categorifying it. Specifically, he defined a bigraded knot homology $Kh_{i,j}(L)$ whose

graded Euler characteristic recovers the Jones polynomial:

$$J_L(q) = \sum_{i,j} (-1)^i (\text{rank } H_{i,j}(L)) q^j.$$

Roughly speaking, categorification gives categorical analogies for set-theoretic constructions, e.g., the category of bigraded abelian groups for polynomials. For more details, see [3]. Khovanov's construction is a strictly stronger invariant than the Jones polynomial: there exist knots with the same Jones polynomial but distinct Khovanov homologies. More importantly, categorification extends the dimension up: there are no maps between polynomials, but there are maps between abelian groups! Jacobsson [24] showed that Khovanov homology is, in fact, functorial: cobordisms between links give homomorphisms of their Khovanov homologies (up to sign). Loosely speaking, Khovanov homology is a topological quantum field theory (TQFT), that is, a functor from a topological cobordism category to an algebraic category. For more on TQFTs, see [2].

Lee [40] constructed a spectral sequence from Khovanov homology to the degenerate Lee homology, which has rank 2^m , where m is the number of components of the link. Rasmussen [56] used Lee's theory to define the s -invariant, a concordance invariant which also gives a lower bound on the 4-ball genus of a knot. Rasmussen used the s -invariant to give an elementary proof of the Milnor conjecture on the slice genus of torus knots, previous proofs of which had relied on gauge theoretic techniques [33]. Freedman et al. [21] suggested that the s -invariant could potentially be used to obstruct the 4-dimensional smooth Poincaré conjecture. Kronheimer and Mrowka [38] later showed that the s -invariant carries insufficient information for the task; however, several refinements and variants of the s -invariant have been defined [48, 59], so there is still hope Khovanov homology could contribute to our understanding of the smooth 4-dimensional Poincaré conjecture.

Around the time of the discovery of the Jones polynomial, in the early 1980s, Donaldson [17] revolutionized the study of 4-manifold topology using techniques from gauge theory. He showed the moduli space of connections on a $SU(2)$ -bundle over a 4-manifold modulo gauge gives rise to invariants of its smooth structure. This was the beginning of a period of rapid development of gauge theoretic techniques in low-dimensional topology. Floer [19] developed instanton Floer homology, the relative invariant for Donaldson theory. Seiberg and Witten [71] introduced new 4-manifold invariants which were analytically more tractable than Donaldson's theory. Kronheimer and Mrowka [34] defined monopole Floer homology, a rigorous $(3 + 1)$ -dimensional TQFT realizing the Seiberg-Witten invariants. Kronheimer and Mrowka [35] also developed knot invariants from instantons.

Ozsváth and Szabó, starting in 2001, defined another $(3 + 1)$ -dimensional TQFT, Heegaard Floer homology [54]. Heegaard Floer homology associates to a closed 3-manifold Y a family of groups $\widehat{HF}(Y, \mathfrak{s})$, $HF^+(Y, \mathfrak{s})$ and $HF^-(Y, \mathfrak{s})$ graded by Spin^c -structures \mathfrak{s} and to a null-homologous knot K in Y the group $\widehat{HFK}(K, Y, \mathfrak{s})$. The construction of Heegaard Floer homology differs from the preceding gauge-theoretic invariants; instead of studying the moduli space of connections modulo gauge, Heegaard Floer homology starts with a Heegaard diagram: a surface with certain handle attachment curves and basepoints. It then applies Lagrangian Floer homology [20] to the symmetric product of the surface. The Heegaard Floer homology 3-manifold invariants are known to agree with the corresponding monopole Floer homology invariants [39].

One of the driving questions in Khovanov homology has been to understand its relationship with gauge-theoretic invariants. Ozsváth and Szabó [55], exploring an analogy between the skein exact sequence in Khovanov homology and the long exact sequence in Heegaard Floer homology for manifolds related by certain surgeries, discovered a spectral sequence $E_k^{HF}(L)$ from the reduced Khovanov homology $\widetilde{Kh}(L)$ of

a link to the hat-variant of Heegaard Floer homology $\widehat{HF}(-\Sigma_L)$ of the double-branch cover of S^3 along the link. Baldwin [4] showed the higher pages of this spectral sequence are link invariants. Bloom [14] constructed an analogous spectral sequence in monopole Floer homology. Rasmussen [57] observed some curious (and still unexplained) structural similarities and rank inequalities between Khovanov homology and knot Floer homology. Kronheimer and Mrowka [35], inspired by the construction of Ozsváth and Szabó, defined a similar spectral sequence in knot instanton homology. They then used the geometric information in knot instanton homology to prove that Khovanov homology detects the unknot [36], a question that remains open for the Jones polynomial.

1.1 Results

The Geometric Spectral Sequence. In [66], Szabó introduced an explicitly computable spectral sequence $E_k^{\text{gss}}(L)$ whose E_2 -page is the the Khovanov homology $Kh(L)$ and whose higher pages are link invariants. Like $E_k^{HF}(L)$, the spectral sequence is constructed by associating to higher faces of the cube of resolutions certain terms in the differential of a filtered chain complex. However, whereas the terms in the Heegaard Floer, monopole and instanton spectral sequences involve certain analytic considerations that make computation difficult, Szabó's construction is purely combinatorial. The author has developed software to compute the spectral sequence E_k^{gss} [60]. Our first set of results concern computations of Szabó's geometric spectral sequence.

Lipshitz, Ozsváth, and Thurston, using techniques from bordered Floer homology [43], gave an algorithm to compute $\widehat{HF}(Y)$ [46] and, more generally, the spectral sequence $E_k^{HF}(L)$ [45], [41]. Lipshitz developed a program in Sage [1] to compute $\widehat{HF}(-\Sigma(L))$ [42]. Zhan ported this program to C++ and added support for comput-

ing the full spectral sequence $E_k^{HF}(L)$ [72].

Zhan computed $\widehat{HF}(-\Sigma(K))$ for all knots with at most 14 crossings and the spectral sequence $E_k^{HF}(K)$ for all knots with at most 12 crossings. In each case, our computation of $\widetilde{E}_\infty^{gss}$ and \widetilde{E}_k^{gss} matched his, respectively. This evidence, along with the further computations and conjectures presented in Chapter 3, suggest the following two conjectures:

Conjecture 1.1.1. *Let K be a knot in S^3 . The spectral sequence $\widetilde{E}_k^{gss}(K)$ collapsed onto the homological grading and $E_k^{HF}(K)$ are isomorphic as graded vector spaces.*

which would imply the weaker conjecture:

Conjecture 1.1.2. *Let K be a knot in S^3 . The rank of $\widetilde{E}_\infty^{gss}(K)$ is equal to the rank of $E_\infty^{HF}(-\Sigma(K)) = \widehat{HF}(-\Sigma(K))$.*

It is natural to expect these conjectures to hold for links also; however, at this point we have only limited computational evidence for the more general case.

It is well-known that Khovanov homology with coefficients modulo 2 decomposes as a direct sum of two identical copies isomorphic to the reduced Khovanov homology $\widetilde{Kh}(L)$. Based on computational evidence, the same appears to be true of the geometric spectral sequence.

Conjecture 1.1.3 (Twin Arrows). *Let L be a link. For $k \geq 2$, the page $E_k^{gss}(L)$ is isomorphic to two copies of $\widetilde{E}_k^{gss}(L)$. Specifically,*

$$E_k^{gss}(L) \cong \widetilde{E}_k^{gss}(L)\{-1\} \oplus \widetilde{E}_k^{gss}(L)\{1\}.$$

Our next result is towards the twin-arrows conjecture. We show that the two ways of defining the reduced geometric spectral sequence, as a sub- or quotient complex, agree. Let \overline{E}_k^{gss} denote the spectral sequence induced by the quotient by the reduced complex. We prove the following proposition.

Proposition 1.1.4 (Baby twin arrows). *Let L be a link. Then we have*

$$\widetilde{E}_k^{gss}(L) \cong \overline{E}_k^{gss}(L)$$

for all $k \geq 2$.

The above results are presented in Chapter 3 and largely taken from the paper [62].

Twisting the Geometric Spectral Sequence. The Jones polynomial can be defined in terms of the combinatorics of spanning trees of the black graph of a link [68]. Champanerkar and Kofman [15] and Wehrli [70] independently showed Khovanov homology admits a spanning tree model; however, their constructions did not give an explicit complex generated by spanning trees. In 2011, inspired by unpublished ideas of Ozsváth and Szabó on “twisting” chain complexes [53], Roberts defined totally twisted Khovanov homology which admits a complex generated by the spanning trees [58]. Jaeger, using a basepoint moving argument, showed that totally twisted Khovanov homology agrees with Khovanov homology for knots collapsed onto a single grading [25]. Using the idea of twisted coefficients, Baldwin-Levine gave an explicit spanning tree model of δ -graded knot Floer homology [5]. In 2013, Kriz-Kriz showed the spanning tree complex identified by Baldwin, Ozsváth and Szabó for the E_3 -page of the twisted variant of the spectral sequence to $\widehat{HF}(\Sigma_L)$ is in fact a link invariant $H_\delta^{\text{BOS}}(L)$ [32].

Our second result is to connect Szabó’s geometric spectral sequence with $H_\delta^{\text{BOS}}(L)$. This connection suggests a strategy for proving the geometric spectral sequence in fact computes $\widehat{HF}(-\Sigma_L)$. The main theorem is the following.

Theorem 1.1.5. *Let L be a link and R a ring of characteristic 2. Choose a diagram \mathcal{D} for L , let p be a point on \mathcal{D} away from the crossings, choose weights $w(e) \in R$ for each edge e of \mathcal{D} not meeting p , and fix a decoration \mathbf{t} of \mathcal{D} . There exists a*

chain complex $\tilde{C}^{tw}(\mathcal{D}, \mathbf{t}, w, p)$ which is a deformation between the geometric spectral sequence and $H_\delta^{BOS}(L)$ in the following sense: If $w \equiv 0$, then we have the following equality of chain complexes:

$$\tilde{C}^{tw}(\mathcal{D}, \mathbf{t}, w, p) = \tilde{C}^{gss}(\mathcal{D}, \mathbf{t}, p),$$

where \tilde{C}^{gss} is the filtered complex for the reduced geometric spectral sequence. If R is a field and the weights $w(e)$ are algebraically independent, then

$$H_\delta(\tilde{C}^{tw}(\mathcal{D}, \mathbf{t}, w, p)) \cong H_\delta^{BOS}(L).$$

There is also an unreduced complex $C^{tw}(\mathcal{D}, \mathbf{t}, w)$ where weights are chosen on all edges and their sum is zero. Its homology is two copies of the reduced homology.

It is natural to ask, what is the relationship between the geometric spectral sequence and the twisted version with non-trivial twisting, e.g., BOS homology? Based on computations of all knots through 12 crossings, we make the following conjecture.

Conjecture 1.1.6. *Let K be a knot. Fix a diagram \mathcal{D} , a point p on \mathcal{D} away from the crossings and weights $w(e) \in R$ for edges e of \mathcal{D} not meeting p . The homology of the twisted geometric spectral sequence does not depend on the choice of diagram, marked point or weights. In particular,*

$$\tilde{H}_\delta^{gss}(K) \cong H_\delta(\tilde{C}^{tw}(\mathcal{D}, \mathbf{t}, w, p)) \cong H_\delta^{BOS}(K).$$

In fact, we suspect something slightly stronger: that there is a basepoint moving homotopy as in Jaeger's proof. This would imply that the twisted geometric spectral sequence for links only depends on the marked component and the sum of the weights for the remaining components.

For generic twisting, the twisted variant of the spectral sequence to the Heegaard

Floer homology of the double-branched cover for knots converges to $\widehat{HF}(-\Sigma_K)$. Thus, Conjecture 1.1.6 would imply the following rank inequality for knots:

$$\begin{aligned} \text{rank } \widehat{HF}(-\Sigma_K) &= \text{rank } \underline{\widehat{HF}}(-\Sigma_K) \\ &\leq \text{rank } E_3^{\underline{HF}}(K) = \text{rank } H^{\text{BOS}}(K) = \text{rank } H(\tilde{C}^{\text{tw}}) = \text{rank } \tilde{H}^{\text{gss}}(K). \end{aligned}$$

If, in addition, the twisted spectral sequence to the double-branched cover for knots collapses on the E_3 -page, then the total homology of the geometric spectral sequence is equal to the Floer Homology of the double-branched cover of K , which would establish Conjecture 1.1.2.

The twisted complex $C^{\text{tw}}(\mathcal{D}, \mathbf{t}, w, p)$ is a deformation of Szabó's complex for the geometric spectral sequence and the proof that d^{tw} is a differential is primarily an adaptation of his original argument to our situation.

The twisted variant of the geometric spectral sequence is presented in Chapter 4 and largely taken from the paper [61].

The Link Splitting Spectral Sequence. Chapter 5 discusses the link splitting spectral sequence, is largely taken from the paper [11] and is joint work with Joshua Batson. We use Khovanov homology to bound a simple topological invariant of a link, its splitting number. Roughly, the splitting number of a link is the minimum number of times its components must be passed through one another in order to completely separate them. For example, the two-component link in Figure 1.1 can be split into an unknot and a trefoil by changing three crossings. We show that this is best possible. Our bound comes from a new spectral sequence beginning at the Khovanov homology of a link and converging to the Khovanov homology of the disjoint union of its components.

Theorem 1.1.7. *Let L be a link and R a ring. Choose weights $w_c \in R$ for each*

component c of L . Then there is a spectral sequence with pages $E_k^{lsss}(L, w)$, and

$$E_1^{lsss}(L, w) \cong Kh(L; R).$$

If the difference $w_c - w_d$ is invertible in R for each pair of components c and d with distinct weights, then the spectral sequence converges to

$$Kh\left(\coprod_{r \in R} L^{(r)}; R\right),$$

where $L^{(r)}$ denotes the sub-link of L consisting of those components with weight r .

Each choice of weights for a link L gives a lower bound on the splitting number.

Theorem 1.1.8. *Let L be a link and let $w_c \in R$ be a set of component weights such that $w_c - w_d$ is invertible for each pair of components c and d . Let $b(L, w)$ be largest k such that $E_k^{lsss}(L, w) \neq E_\infty^{lsss}(L, w)$. Then $b(L, w) \leq \text{sp}(L)$.*

A simple corollary of Theorem 1.1.7 is a rank inequality in ordinary Khovanov homology.

Corollary 1.1.9. *Let \mathbb{F} be any field, and let L be a link with components K_1, \dots, K_m . Then*

$$\text{rank } Kh(L; \mathbb{F}) \geq \text{rank } \otimes_{c=1}^m Kh(K_c; \mathbb{F}).$$

We also use the spectral sequence to show that the Poincaré polynomial of Khovanov homology detects the unlink. In contrast, there are infinite families of links with the same Jones polynomial as the unlink [18, 69].

Theorem 1.1.10. *Let L be an m -component link, and U^m the m -component unlink. If*

$$\text{rank } Kh^{i,j}(L; \mathbb{F}_2) = \text{rank } Kh^{i,j}(U^m; \mathbb{F}_2)$$

for all i, j , then L is the unlink.

The proof of Theorem 1.1.10 depends on two earlier spectral sequences that relate Khovanov homology to more manifestly geometric invariants coming from Floer homology. The first, constructed by Ozsváth and Szabó, begins at the Khovanov homology of a link and converges to the Heegaard Floer homology of its branched double cover [55]. The second, constructed by Kronheimer and Mrowka, begins at the Khovanov homology of a knot and converges to its instanton knot Floer homology [37]. The latter group is nontrivial for nontrivial knots, so:

Theorem 1.1.11 (Kronheimer-Mrowka). *Let K be a knot, and U the unknot. If*

$$\text{rank } Kh(K) = \text{rank } Kh(U),$$

then K is the unknot.




Figure 1.1: The link 2L13n3752 has splitting number 3.

The Khovanov homology groups contain more information than their ranks alone—there is a natural action of the algebra

$$A_m = \mathbb{F}_2[X_1, \dots, X_m]/(X_1^2, X_2^2, \dots, X_m^2)$$

on the homology of an m -component link. Hedden and Ni [22] showed that the entire spectral sequence of Ozsváth and Szabó admits a compatible A_m action. They then

used of Floer homology to detect $S^1 \times S^2$ summands in the branched double cover of the link, and showed:

Theorem 1.1.12 (Hedden-Ni). *Let L be an m -component link, and U^m the m -component unlink. If there is an isomorphism of A_m modules*

$$Kh(L; \mathbb{F}_2) \cong Kh(U^m; \mathbb{F}_2),$$

then L is the unlink.

To prove Theorem 1.1.10, we apply our spectral sequence with component weights in a suitably large finite field \mathbb{F} of characteristic 2. We lift the A_m -module structure from the abutment of our spectral sequence, which turns out to be isomorphic to $Kh(U^m; \mathbb{F})$, to the first page, $Kh(L; \mathbb{F})$, and then to $Kh(L; \mathbb{F}_2)$, where we apply Theorem 1.1.12.

Outline. In Chapter 2, we review the construction of Khovanov homology, the geometric spectral sequence and Baldwin-Ozsváth-Szabó homology. In Chapter 3, we present a number of conjectures and results of computations of the Szabó geometric spectral sequence. In Chapter 4, we define the twisted geometric chain complex and establish Theorem 1.1.5. Finally, in Chapter 5 we define the link splitting spectral sequence, study its properties and give several applications.

Chapter 2

Background

In this chapter, we review the construction of Khovanov homology, Szabó's geometric spectral sequence, and Baldwin-Ozsváth-Szabó homology.

2.1 A review of Khovanov homology

(Following [29] and [7].)

Let L be an oriented link and \mathcal{D} a diagram for L . A link diagram crossing can be resolved in one of two ways, called the 0-resolution and 1-resolution in Figure 2.1. A resolution of \mathcal{D} , also called a complete resolution of \mathcal{D} , is a choice of resolution at each crossing. Number the crossings of \mathcal{D} from 1 to n so we can index complete resolutions by vertices in the hypercube $\mathcal{R} = \{0, 1\}^n$. An edge in the cube connects a pair of resolutions (I, J) , where J is obtained from I by changing the i^{th} digit from 0 to 1. A complete resolution I yields a finite collection of circles in the plane, which we may also call I . An edge (I, J) yields a cobordism from I to J , given by the natural saddle cobordism from the 0- to the 1-resolution in a neighborhood of the changing crossing and the product cobordism elsewhere.

A $(1 + 1)$ -dimensional TQFT is determined by a commutative Frobenius algebra [31]. We fix a ring of coefficients R , and let \mathcal{A} be the TQFT associated to the Frobenius



Figure 2.1: The 0 and 1 resolutions associated to a crossing.

algebra $V = H^*(S^2; R) = R[x]/(x^2)$. The diagonal map $i : S^2 \hookrightarrow S^2 \times S^2$ induces the multiplication $i^* : H^*(S^2 \times S^2) \rightarrow H^*(S^2)$ and, by conjugation by Poincaré duality, induces the comultiplication $PD \circ i_* \circ PD : H^*(S^2) \rightarrow H^*(S^2 \times S^2)$. More explicitly, the multiplication $m : V \otimes V \rightarrow V$ is given by

$$\begin{aligned} m(1 \otimes 1) &= 1 & m(x \otimes 1) &= x \\ m(1 \otimes x) &= x & m(x \otimes x) &= 0, \end{aligned}$$

and the comultiplication $\Delta : V \rightarrow V \otimes V$ is given by

$$\Delta(1) = 1 \quad \Delta(x) = 1 \otimes x + x \otimes 1.$$

The TQFT \mathcal{A} associates to a circle the R -module V and takes disjoint unions to tensor products. The pair of pants cobordism that merges two circles into one induces the multiplication map m , and the pair of pants cobordism that splits one circle into two induces the comultiplication map Δ .

Let $S = (x_1, \dots, x_p)$ be a collection of circles. To simplify notation, we note that

$$\begin{aligned} \mathcal{A}(S) &= \bigotimes_{i=1}^p V \\ &= R[x_1, \dots, x_p]/(x_1^2, \dots, x_p^2). \end{aligned}$$

We will write elements of $V(S)$ as (commutative) products of the circles x_i rather than elements of the tensor product. Such a product of circles is called a monomial of S .

Applying the TQFT \mathcal{A} to the cube of resolutions, we obtain a cube-graded complex of R -modules. For each resolution I , we have an R -module $\mathcal{A}(I)$, and for each edge (I, J) , we have a homomorphism $\mathcal{A}(I, J) : \mathcal{A}(I) \rightarrow \mathcal{A}(J)$. Khovanov's complex is obtained by collapsing the cube-graded complex. We set

$$C_{\mathcal{D}} = \bigoplus_{\text{resolutions } I} V(I).$$

The differential $d : C_{\mathcal{D}} \rightarrow C_{\mathcal{D}}$ is given by

$$d = \sum_{\text{edges } (I, J)} (-1)^{n(I, J)} \mathcal{A}(I, J),$$

where, if (I, J) differ at i ,

$$n(I, J) = \#\{I(k) = 1 \mid 1 \leq k < i\}.$$

We define five related gradings on $C_{\mathcal{D}}$ as follows. Let $x \in V(I)$. The homological or h grading is given by

$$h(x) = |I| - n_-(\mathcal{D}),$$

where $|I|$ number of 1 digits in I and $n_-(\mathcal{D})$ is the number of negative crossings in \mathcal{D} . Monomials in $V^{\otimes p}$ have a natural grading induced by

$$\text{gr}(1) = 0 \text{ and } \text{gr}(x_i) = 2.$$

The quantum or q grading is given by

$$q(x) = \text{gr}(x) + |I| + n_+(\mathcal{D}) - 2n_-(\mathcal{D}),$$

where $n_+(\mathcal{D})$ is the number of positive crossings in \mathcal{D} . There are several standard

ways to collapse the q and h bigrading onto a single grading. The first is to simply discard the q grading. Next, there is the δ grading, given by

$$\delta(x) = \frac{q(x)}{2} - h(x).$$

Finally, internal or ℓ grading is given by

$$\ell(x) = \text{gr}(x) - p(I) - \text{writhe}(\mathcal{D}),$$

where $p(I)$ is the number of circles in the resolution I . Finally, we define the g grading, a normalization of the q grading, by

$$g(x) = \frac{q(x) - m}{2},$$

where m is the number of components of L . (It turns out that g is always an integer [29, §6.1].) The h grading will induce the filtration in the definition of the geometric spectral sequence and the g grading will induce the filtration on $C_{\mathcal{D}}$ in the definition of the link splitting spectral sequence.

Khovanov's differential d increases the h, δ and ℓ gradings by 1 and it preserves the q and g gradings. Khovanov homology is

$$Kh(L) = H^*(C_{\mathcal{D}}, d),$$

and has a bigrading given by (h, q) .

There is a reduced version of Khovanov homology. We will continue as before, but now suppose L has a marked component c and choose a point p on \mathcal{D} on the component c away from the crossings. Let I be a resolution and let $x_p = x_p(I)$ denote the circle in resolution I meeting p . Let $\tilde{C}_{\mathcal{D}}$ denote the submodule of $C_{\mathcal{D}}$ generated by

the monomials in each resolution I divisible by $x_p(I)$. It is straightforward to verify that this is a subcomplex. The homology of this subcomplex is denoted $\widetilde{Kh}(L, c)$ and depends only on the marked component but not the choice of marked point p . If R has characteristic 2, then the reduced homology doesn't even depend on the marked component.

A choice of marked point on the diagram \mathcal{D} also induces an endomorphism of Khovanov homology [30]. Let p be a marked point on \mathcal{D} away from the double points. As above, let $x_p = x_p(I)$ denote the circle of I meeting p . Define a map $X_p : C_{\mathcal{D}} \rightarrow C_{\mathcal{D}}$ by

$$X_p(x) = x_p x$$

for $x \in V(I)$. The map X_p is a chain map and shifts the (h, q) bigrading by $(0, -2)$. The map induced on homology, which we also call X_p , depends only on the marked component, and not on the choice of marked point.

2.2 A review of the geometric spectral sequence

We briefly review the construction of Szabó's geometric spectral sequence [66]. Let L be an oriented link and R a ring of characteristic 2. Fix a diagram \mathcal{D} for L in S^2 with n crossings. The geometric spectral sequence $E_k^{\text{gss}}(L)$ is constructed by defining higher differentials associated to the faces of the cube of resolutions on the Khovanov chain complex. These higher differentials depend on a certain choice of orientation at each crossing, denoted \mathbf{t} . The spectral sequence $E_k^{\text{gss}}(L)$ is induced from a filtered chain complex $C^{\text{gss}}(\mathcal{D}, \mathbf{t}) = (C_{\mathcal{D}}, d^{\text{gss}}(\mathbf{t}))$. The module $C_{\mathcal{D}}$ is the same module underlying the Khovanov complex.

We now recall the construction of the differential $d^{\text{gss}}(\mathbf{t})$. In the 0-resolution, there is an arc between the segments of the resolution such that surgery along this arc gives the 1-resolution. Let \mathbf{t} denote a choice of orientation of the 0-resolution surgery arc

for each crossing. The pair $(\mathcal{D}, \mathbf{t})$ is called a decorated diagram.

The differential is defined in terms of configurations. A k -dimensional configuration $\mathcal{C} = (x_1, \dots, x_t, \gamma_1, \dots, \gamma_k)$ is a collection of embedded circles (x_1, \dots, x_t) in S^2 together with k embedded, oriented arcs $\gamma_1, \dots, \gamma_k$ such that the circles and the interior of the arcs are all disjoint and the endpoints of the arcs lie on the circles. (Compare the definition of resolution configuration in [47, Definition 2.1].) Recall the following operations on configurations:

- *undecorated configuration.* The undecorated configuration $\bar{\mathcal{C}}$ is obtained from \mathcal{C} by forgetting the orientation of the arcs.
- *dual configuration.* The dual configuration $\mathcal{C}^* = (y_1, \dots, y_s, \gamma_1^*, \dots, \gamma_k^*)$ is the configuration obtained from \mathcal{C} by performing surgery along the arcs γ_i . The dual arcs γ_i^* are obtained by rotating the arcs γ_i counterclockwise by 90 degrees.
- *reverse configuration.* The reverse configuration $r(\mathcal{C})$ is obtained from \mathcal{C} by reversing the orientation of the arcs γ_i .
- *mirror configuration.* The mirror configuration $m(\mathcal{C})$ is obtained from \mathcal{C} by reversing the orientation of the ambient 2-sphere.

Let (I, J) be a face of the hypercube \mathcal{R} , that is, $I(i) \leq J(i)$ for all $1 \leq i \leq k$. We also define the *restricted dual configuration* $\mathcal{C}(I, J)$ by modifying the arcs γ_i as follows:

- if $I(i) = J(i) = 0$, discard arc γ_i ,
- if $I(i) = 0, J(i) = 1$, leave γ_i unchanged, and
- if $I(i) = 1, J(i) = 1$, perform surgery along the arc γ_i and discard the dual arc γ_i^* .

The circles x_i are called the starting circles of \mathcal{C} . The circles y_i of \mathcal{C}^* are called the ending circles of \mathcal{C} . Set

$$V_0(\mathcal{C}) = V(x_1, \dots, x_t), \quad V_1(\mathcal{C}) = V(y_1, \dots, y_s).$$

Let $P(\mathcal{C}) = (x_{i_1}, \dots, x_{i_k})$ be the collection of circles of \mathcal{C} which are disjoint from the arcs; they are called the passive circles of \mathcal{C} . Let \mathcal{C}_0 denote the configuration obtained by deleting the passive circles; this is called the active part of \mathcal{C} . A configuration with no passive circles is called purely active. There are decompositions

$$V_0(\mathcal{C}) = V_0(\mathcal{C}_0) \otimes P(\mathcal{C}), \quad V_1(\mathcal{C}) = V_1(\mathcal{C}_0) \otimes P(\mathcal{C}).$$

A configuration \mathcal{C} is disconnected if the graph with vertices circles of \mathcal{C}_0 and edges arcs γ_i is disconnected.

Let (I, J) be a k -face of the cube of resolutions. The k -dimensional configuration $\mathcal{C}(I, J, \mathbf{t}) = (x_1, \dots, x_t, \gamma_1, \dots, \gamma_k)$ consists of the circles (x_1, \dots, x_t) of I together with the 0-resolution surgery arcs $\gamma_{i_1}, \dots, \gamma_{i_k}$ corresponding to the coordinates i_1, \dots, i_k where I and J differ and with orientation given by \mathbf{t} . Note, the the restricted dual configuration is defined so $\mathcal{C}(I, J, \mathbf{t}) = \mathcal{C}(I, J)$ when $\mathcal{C} = \mathcal{C}(\mathbf{0}, \mathbf{1}, \mathbf{t})$.

A family of maps $G_{\mathcal{C}} : V_0(\mathcal{C}) \rightarrow V_1(\mathcal{C})$, one for each configuration \mathcal{C} , determines a map $G : C_{\mathcal{D}} \rightarrow C_{\mathcal{D}}$ given by

$$G = \sum_{k=1}^n G_k$$

where

$$G_k = \sum_{k\text{-faces } (I, J)} G_{\mathcal{C}(I, J, \mathbf{t})}.$$

We will write $G(\mathbf{t})$ and $G_k(\mathbf{t})$ when we want to make the dependence on the decoration explicit.

The differential $d^{\text{gss}}(\mathbf{t})$ is defined in this way: by giving a family of maps $d_{\mathcal{C}}^{\text{gss}}$.




Figure 2.2: Examples of the five families of configurations.

Note, Szabó denoted this family of maps $F_{\mathcal{C}}$ in [66]. There are five families A_k , B_k , $C_{p,q}$, $D_{p,q}$ and $E_{p,q}$, $p + q = k$ of k -dimensional configurations for which $d_{\mathcal{C}}^{\text{gss}}$ is non-trivial. Examples of those five families are given in Figure 2.2. We will not review the definition of $d_{\mathcal{C}}^{\text{gss}}$ here, but give a generalized definition in Chapter 4 for the differential of the twisted chain complex.

The family of maps $d_{\mathcal{C}}^{\text{gss}}$ satisfies several nice properties. We recall their definition now. These properties will apply to other maps we consider in the sequel, so we will defined them in terms of a generic family of maps $G_{\mathcal{C}}$.

Definition 2.2.1 (Extension Formula). Let \mathcal{C} be a configuration. We have the compositions

$$V_0(\mathcal{C}) = V_0(\mathcal{C}_0) \otimes P(\mathcal{C})$$

and

$$V_1(\mathcal{C}) = V_1(\mathcal{C}_0) \otimes P(\mathcal{C}).$$

We say $G_{\mathcal{C}}$ satisfies the extension formula if

$$G_{\mathcal{C}}(a \otimes b) = G_{\mathcal{C}_0}(a) \otimes b$$

for all pairs of monomials $a \in V_0(\mathcal{C}_0)$ and $b \in P(\mathcal{C})$.

Remark 2.2.2. If $G_{\mathcal{C}}$ satisfies the extension formula, it is enough to define it for purely active configurations and extend it to all configurations by the extension formula.

Definition 2.2.3 (Disconnected Rule). We say $G_{\mathcal{C}}$ satisfies the disconnected rule if

$$G_{\mathcal{C}} \equiv 0$$

for all disconnected configurations \mathcal{C} .

Definition 2.2.4 (Conjugation Rule). We say $G_{\mathcal{C}}$ satisfies the conjugation rule if

$$G_{\mathcal{C}} \equiv G_{r(\mathcal{C})}$$

for all configurations \mathcal{C} .

Definition 2.2.5 (Duality Rule). We say $G_{\mathcal{C}}$ satisfies the duality rule if, for every configuration \mathcal{C} and every pair of monomials $a \in V_0(\mathcal{C})$ and $b \in V_1(\mathcal{C})$, the coefficient of $G_{\mathcal{C}}(a)$ at b is equal to the coefficient of $G_{m(\mathcal{C})}(b^*)$ at a^* .

Definition 2.2.6 (Filtration Rule). Let \mathcal{C} be a configuration and p a point on the union of the starting circles away from the endpoints of the arcs. Let x_p denote the starting circle meeting p and similarly y_p denote the ending circle meeting p . We say that $G_{\mathcal{C}}$ satisfies the filtration rule if, for all configurations \mathcal{C} , points p as above and all monomials $a \in V_0(\mathcal{C})$ divisible by x_p and all monomials $b \in V_1(\mathcal{C})$, if the coefficient of $G_{\mathcal{C}}(a)$ at b is non-zero, then b is divisible by y_p .

Definition 2.2.7 (Grading Rule). We say $G_{\mathcal{C}}$ satisfies the grading rule if, for all configurations \mathcal{C} and pairs of monomials $a \in V_0(\mathcal{C})$ and $b \in V_1(\mathcal{C})$, if the coefficient of $G_{\mathcal{C}}(a)$ at b is non-zero, then

$$gr(b) - gr(a) = k - 2.$$

The family of maps $d_{\mathcal{C}}^{\text{gss}}$ defining the differential $d^{\text{gss}}(\mathbf{t})$ satisfies the extension formula, disconnected rule, conjugation rule, duality rule and grading rule.

The component $d_k^{\text{gss}}(\mathbf{t})$ has homogeneous (h, q) -degree $(k, 2k - 2)$ and δ degree -1 . Thus, the total differential $d^{\text{gss}}(\mathbf{t})$ has δ degree -1 . The homological grading induces a filtration on $C_{\mathcal{D}}$.

Theorem 2.2.8 (Szabó [66]). *The map $d^{\text{gss}}(\mathbf{t})$ is a differential, that is,*

$$d^{\text{gss}}(\mathbf{t}) \circ d^{\text{gss}}(\mathbf{t}) = 0.$$

The spectral sequence $E_k^{\text{gss}}(L)$ induced from the filtration coming from the homological grading is an invariant of the oriented link L for $k \geq 2$ and $E_2^{\text{gss}}(L) = Kh(L)$.

As with Khovanov homology, there is a reduced version of the geometric spectral sequence. Suppose L has a marked component c and choose a point p on \mathcal{D} on the component c away from the crossings. By the filtration rule for $d_{\mathcal{C}}^{\text{gss}}$, the submodule $\tilde{C}_{\mathcal{D}}$ induced by p is a subcomplex. We denote the induced spectral sequence $\tilde{E}_k^{\text{gss}}(L, c)$.

2.3 A review of Baldwin-Ozsváth-Szabó homology

In this subsection, we review the definition of Baldwin-Ozsváth-Szabó homology [53, 32], hereafter referred to as BOS homology.

Let L be link and F a field of characteristic 2. Fix a diagram \mathcal{D} for L and a marked point p on \mathcal{D} away from the crossings. Pick algebraically independent weights




Figure 2.3: Positive and negative 2-dimensional configurations between connected resolutions with marked point p .

$w(e) \in F$ for each edge e of \mathcal{D} such that

$$\prod_e w(e) = 1.$$

For a connected resolution I , let the grading $\delta(I)$ be given by

$$\delta(I) = \frac{1}{2}(n_+ - |I|).$$

Let C_{BOS} be the δ -graded vector space generated by the connected resolutions.

Now we define the differential d_{BOS} . Let I be a connected resolution and consider a connected resolution I' obtained from I by changing 0-resolutions to 1-resolutions at two crossings j and k .

Up to isotopy in the sphere, the configuration $\mathcal{C}(I, I')$ together with the marked point p has one of the two forms shown in Figure 2.3. We call the configuration on left negative, and the one on the right positive. Let the constant $c_{\mathcal{C}(I, I')}$ be given by

$$c_{\mathcal{C}(I, I')} = \begin{cases} 1 & \mathcal{C}(I, I') \text{ negative} \\ 0 & \mathcal{C}(I, I') \text{ positive} \end{cases}$$

Let J and K be resolutions obtained from I by changing crossing j and k , respec-




Figure 2.4: Definition of height $h(e)$ for an edge e of the black graph $B(\mathcal{D})$.

tively, from a 0-resolution to a 1-resolution. The configurations J and K each consist of two circles. Let x_J and x_K denote the circle in J and K , respectively, not meeting the marked point p . Then set

$$d_{\text{BOS}}(I) = \sum_{I'} \left(c_{\mathcal{C}(I, I')} + \frac{1}{1 + w(x_J)} + \frac{1}{1 + w(x_K)} \right) I' \quad (2.1)$$

where the summation is over all such I' as defined above.

Kriz-Kriz [32] prove the following.

Theorem 2.3.1 (Kriz-Kriz). *The map d_{BOS} is a differential and, under the assumption the weights $w(e)$ are algebraically independent, the homology*

$$H_{\text{BOS}, \delta}(L) = H_{\delta}(C_{\text{BOS}}, d_{\text{BOS}})$$

is an invariant of L and does not depend on \mathcal{D}, p or the choice of weights $w(e)$.

Remark 2.3.2. It appears the complex we have defined here differs from that defined in [32], but in fact they are the same.

First, we briefly review the definition of the complex $(C_{\text{KK}}, d_{\text{KK}})$ defined in [32, (2.11)]. Let $B(\mathcal{D})$ and $W(\mathcal{D})$ denote the black and white graph, respectively, of \mathcal{D} . The graph $B(G)$ induces a CW decomposition $B(S^2)$ of the sphere. The marked point p determines a black basepoint. Let the height $h(e) \in \{0, 1\}$ of a black edge be defined by the conventions shown in Figure 2.4. Let T be a spanning tree of $B(G)$.

Set

$$h(T) = \sum_{e \in T} h(e).$$

Define the grading $d(T)$ by

$$d(T) = \frac{1}{2}(h(T) - n_-).$$

Pick algebraically independent weights $z_v \in F$ for each vertex of $B(\mathcal{D})$ such that

$$\prod_v z_v = 1$$

and algebraically independent weights $u_f \in F$ for each face of $B(\mathcal{D})$ such that

$$\prod_f u_f = 1.$$

Let T, T' be spanning trees of $B(\mathcal{D})$ such that

$$T' = (T \setminus \{e\}) \cup \{f\}$$

with $h(e) = 0$ and $h(f) = 1$. The intersection $T \cap T'$ has two components. Let C denote the component not containing the black basepoint, and C' the component containing the basepoint. The union $T \cap T'$ contains a unique cycle c . Orient it so that f is oriented in the intersection from C to C' . Write $[c] = dx$ for a 2-chain x of $B(S^2)$. Define

$$\alpha(T, T') = \prod_f (u_f)^{x_f}$$

and

$$\beta(T, T') = \prod_{v \in C} z_v.$$

The d -graded vector space C_{KK} is generated over F by the spanning trees of $B(G)$.




Figure 2.5: First case in showing two definitions of $H_{\text{BOS}}(L)$ complex agree.

The differential is then given by

$$d_{\text{KK}}(T) = \sum_{T'} \left(\frac{1}{1 + \alpha(T, T')} + \frac{1}{1 + \beta(T, T')} \right) T'.$$

Next, we roughly sketch the connection between our definition and the Kriz-Kriz definition. The spanning trees and connected resolutions are in one-to-one correspondence. More generally, subsets $S \subset E$ of black edges correspond to resolutions I_S by

$$I_S(i) = \begin{cases} 1 & (h(i) = 1) = (i \in S) \\ 0 & \text{otherwise.} \end{cases}$$

Then the d and δ gradings agree. Fix an CCW orientation of some connected resolution of \mathcal{D} . This corresponds to having the point at infinity contained in a black face. This induces a natural orientation, the resolution orientation, on the circles of each resolution. The circles x of I_E correspond to the vertices v_x of $B(\mathcal{D})$ and the circles y of E_\emptyset correspond to the faces f_y of $B(\mathcal{D})$. Orient the faces of $B(\mathcal{D})$ to agree with the orientation of the corresponding circle. Set

$$z_{v_x} = w(x) \text{ and } u_{f_y} = w(y).$$

Let T, T' be spanning trees as above. Set $I = I_T$ and $I' = I_{T'}$. The edge f is on the

outside in I . Set $J = I_{T \setminus \{e\}}$ and $K = I_{T \cup \{f\}}$. It is immediate that

$$\alpha(T, T') = w(x_J).$$

Finally, we claim

$$\frac{1}{1 + \beta(T, T')} = c_{C(I, I')} + \frac{1}{1 + w(x_K)}.$$

There are four cases to consider, depending on the edge which meets p . The first case is shown in Figure 2.5. We find

$$\begin{aligned} \frac{1}{1 + \beta(T, T')} &= 1 + \frac{1}{1 + w(x_1)^{-1}} \\ &= \frac{1}{1 + w(x_1)} \end{aligned}$$

since we are working modulo 2.

Chapter 3

Computing the Geometric Spectral Sequence

The author has implemented software to compute the geometric spectral sequence as part of the `knotkit` software package [60]. In Section 3.1, we give a brief overview of the methods used to compute the spectral sequence. In Section 3.2, we prove several results inspired by computations, including Proposition 1.1.4, baby twin arrows. Finally, in Section 3.3, we make a number of conjectures based on computations of E_k^{gss} , extending Conjecture 1.1.1 that the geometric spectral sequence in fact computes the spectral sequence in Heegaard Floer homology to the double-branched cover.

Throughout this chapter, $E_k = E_k^{\text{gss}}$ will be used to denote the geometric spectral sequence, \tilde{E}_k the reduced version, $C(\mathcal{D}, \mathbf{t}) = (C_{\mathcal{D}}, d(\mathbf{t}))$ filtered chain complex inducing the spectral sequence, etc.

3.1 Computing the geometric spectral sequence

We now describe the algorithm used for computing the geometric spectral sequence. We begin with an decorated planar link diagram $(\mathcal{D}, \mathbf{t})$.

The first step is to build the chain complex $C(\mathcal{D}, \mathbf{t})$. As with Khovanov homology,

Algorithm 1 Algorithm to construct the chain complex $(C_{\mathcal{D}}, d(\mathbf{t}))$.

```

 $d(\mathbf{t}) \leftarrow 0$ 
for each face  $(I, J)$  of  $\mathcal{R}$  do
   $\mathcal{C} \leftarrow \mathcal{C}(I, J, \mathbf{t})$ 
  if  $\mathcal{C}_0$  has type  $A, B, C, D$  or  $E$  then
    for each generator  $a$  of  $V(\mathcal{C}_0)$  and  $p$  of  $P(\mathcal{C})$  do
       $d(\mathbf{t})(a \cdot p) \leftarrow d(\mathbf{t})(a \cdot p) + d_{\mathcal{C}_0}(a) \cdot p$ 
    end for
  end if
end for
return  $d(\mathbf{t})$ 

```

the combinatorial description of the chain group $C_{\mathcal{D}}$ and differential $d(\mathbf{t})$ is amenable to direct computation. In our implementation, we naively follow the combinatorial description. The generators of $C_{\mathcal{D}}$ are represented by pairs $I:m$ with $I \in \mathbf{2}^n$ and $m \in \mathbf{2}^t$ where $t = t(I)$ is the number of circles in the resolution I . To calculate $d(\mathbf{t})$ we simply sum the contributions $d_{\mathcal{C}(I,J,\mathbf{t})}$ for each face (I, J) of \mathcal{R} . Pseudocode is given in Algorithm 1.

After building the chain complex, the next step is compute the spectral sequence. We use repeated application of the cancellation lemma as outlined by Baldwin in [4, Section 4].

Lemma 3.1.1 (Cancellation Lemma). *Let (C, d) be a chain complex freely generated by $\{x_i\}$. Let $d(x_i, x_j)$ denote the coefficient of x_j in $d(x_i)$. Suppose $d(x_k, x_\ell) = 1$. Let (C', d') be the complex where C' is generated by $\{x_i | i \neq k, \ell\}$ and the differential d' given by*

$$d'(x_i) = d(x_i) + d(x_i, x_\ell)d(x_k).$$

The (C, d) is chain homotopy equivalent to (C', d') .

We say that (C', d') is obtained from (C, d) by canceling the term $d(x_k, x_\ell)$. The cancellation lemma admits a refinement for filtered complexes. This refinement, together with the mapping lemma for spectral sequences, establishes the following pro-

cess for computing the spectral sequence. The pair (E_0, d_0) is simply $(C_{\mathcal{D}}, d_0(\mathbf{t}))$. Then we cancel the terms of $d(\mathbf{t})$ that preserve the homological grading, to obtain a new complex which, by abuse of notation, we also call $(C_{\mathcal{D}}, d(\mathbf{t}))$. The pair (E_1, d_1) is then $(C_{\mathcal{D}}, d_1(\mathbf{t}))$, where $d_1(\mathbf{t})$ denotes the terms of $d(\mathbf{t})$ which increase the homological grading by 1. Then we cancel the terms of $d(\mathbf{t})$ which shift the homological grading by 1, and $(E_2, d_2(\mathbf{t}))$ is $(C_{\mathcal{D}}, d_2(\mathbf{t}))$. This process terminates when all the terms of the differential are canceled and $d(\mathbf{t}) = 0$. This will always happen since the homological degree of $C_{\mathcal{D}}$ has bounded support. Pseudocode for the process of canceling terms to compute the spectral sequence is given in Algorithm 2.

Algorithm 2 Algorithm to compute the spectral sequence E_k associated to the filtered chain complex $(C_{\mathcal{D}}, d(\mathbf{t}))$.

```

i ← 0
while  $d(\mathbf{t}) \neq 0$  do
   $(E_i, d_i) \leftarrow (C_{\mathcal{D}}, d_i(\mathbf{t}))$ 
  while  $d(\mathbf{t})(x_k, x_\ell) = 1$  for some  $k, \ell$  with  $h(x_\ell) - h(x_k) = i$  do
    cancel  $d(\mathbf{t})(x_k, x_\ell)$  in  $(C_{\mathcal{D}}, d(\mathbf{t}))$ 
  end while
  i ← i + 1
end while
return  $E_k$ 

```

Both the time and space complexity of the naive algorithm are exponential in the number of crossings of the diagram. In practice, it is feasible to compute the E_k for knots with 18–19 crossings on a computer with 12Gb of RAM. The author has developed a second program to compute E_k based on a partial construction of the spectral sequence for tangles obtained by adapting algebraic techniques from bordered Floer homology [44]. We plan to describe this construction in a subsequent paper.

In [8], [9], Bar-Natan introduced a fast divide-and-conquer algorithm for computing Khovanov homology. Although there is no analysis of its algorithmic complexity, its running time appears to depend heavily on the girth of the knot, with girth 14 the upper end of the feasible range [21]. It would also be interesting to see if Bar-Natan’s

formulation could be extended to yield a fast algorithm for computing Szabó's spectral sequence.

3.2 Results

Recall from the introduction, $\overline{C}(\mathcal{D}, \mathbf{t}, p)$ is the quotient complex

$$C(\mathcal{D}, \mathbf{t}) / \tilde{C}(\mathcal{D}, \mathbf{t}, p)$$

and $\overline{E}_k(L, c)$ is the spectral sequence induced from $\overline{C}(\mathcal{D}, \mathbf{t}, p)$. Towards the twin arrows conjecture, we prove our main result, Proposition 1.1.4.

Proof of Proposition 1.1.4. We will define a chain map

$$P(\mathbf{t}) : C(\mathcal{D}, \mathbf{t}) \rightarrow C(\mathcal{D}, \mathbf{t})$$

that induces the desired chain homotopy equivalence. The construction of $P(\mathbf{t})$ and proof that it is a chain map closely mimics the the construction of $d(\mathbf{t})$ and proof that it is a differential in [66]. We define a family of maps

$$P_{\mathcal{C}} : V_0(\mathcal{C}) \rightarrow V_1(\mathcal{C})$$

for each configuration \mathcal{C} . The map $P_{\mathcal{C}}$ satisfies the extension formula, the disconnected rule, the naturality rule and the rotation rule. The map $P_{\mathcal{C}}$ satisfies an additional property related to the point p . If $a \in V_0(\mathcal{C})$ and x_p divides a , then $P_{\mathcal{C}}(a) = 0$. If $P_{\mathcal{C}}(a) \neq 0$ for some $a \in V_0(\mathcal{C})$, then y_p divides $P_{\mathcal{C}}(a)$. Note, $P_{\mathcal{C}}$ does *not* satisfy the conjugation or duality rules.

Let $X_p : C_{\mathcal{D}} \rightarrow C_{\mathcal{D}}$ be given by the formula:

$$X_p(a) = \begin{cases} x_p a & x_p \text{ does not divide } a \\ 0 & \text{otherwise,} \end{cases}$$

where $a \in V(I)$ and $x_p = x_p(I)$ denotes the circle of I which meets the point p . Set $P_0(\mathbf{t}) = X_p$.

In what follows, we list only the non-zero terms of $P_{\mathcal{C}}$. Since $P_{\mathcal{C}}$ satisfies the extension formula, it suffices to define $P_{\mathcal{C}}$ for purely active configurations. Now we define $P_{\mathcal{C}}$ for 1-dimensional configurations. We make the following definitions:

Definition 3.2.1. Let $\mathcal{C} = (x, \gamma)$ be a purely active 1-dimensional configuration where γ is a split arc. Then \mathcal{C} has one starting circle x and two ending circles. Let y_1 denote the ending circle which meets the tail of γ^* and y_2 the ending circle which meets the head of γ^* . If $y_2 = y_p$, we define

$$P_{\mathcal{C}}(1) = y_2.$$

Definition 3.2.2. Let $\mathcal{C} = (x_1, x_2, \gamma)$ be a purely active 1-dimensional configuration where γ is a join arc. Then \mathcal{C} has two starting circles and one ending circle y . Let x_1 denote the starting circle that meets the tail of γ and x_2 the starting circle that meets the head of γ . If $x_1 = x_p$, we define

$$P_{\mathcal{C}}(x_2) = y.$$

Now we define $P_{\mathcal{C}}$ for purely active 2-dimensional configurations. Up to naturality and rotation, there are three 2-dimensional configurations for which $P_{\mathcal{C}} \neq 0$. They are given in Figure 3.1. We make the following definitions:




Figure 3.1: The three purely active 2-dimensional configurations up to naturality and rotation for which $P_{\mathcal{C}} \neq 0$.

Definition 3.2.3. • For a configuration of type 1, we define

$$P_{\mathcal{C}}(1) = y_p.$$

- For a configuration of type 2, there are two starting circles. Let $x_1 = x_p$ and x_2 be the other starting circle. We define

$$P_{\mathcal{C}}(x_1) = y_p.$$

- For a configuration of type 3, there are three starting circles. Let $x_1 = x_p$ and let x_2 and x_3 denote the other starting circles which meet a single arc. We define

$$P_{\mathcal{C}}(x_2 x_3) = y_p.$$

Finally, for $k > 2$, we make the following definition:

Definition 3.2.4. A k -dimensional purely active configuration $\mathcal{C} = (x_1, \dots, x_s, \gamma_1, \dots, \gamma_k)$ with $p+1$ starting circles and $q+1$ ending circles is said to be of type $P_{p,q}$ if, for each pair (i, j) with $1 \leq i < j \leq k$, the 2-dimensional configuration $(x_1, \dots, x_s, \gamma_i, \gamma_j)$ is of type 1, 2 or 3. There is a unique starting circle $x_1 = x_p$ called the central starting circle. The other starting circles x_i meet a single arc. Similarly there is a unique

ending circle $y_1 = y_p$ that meets all the all the dual arcs γ_i^* . In this case, we define

$$P_{\mathcal{C}}(x_2 x_3 \cdots x_{p+1}) = y_1$$

when $p \geq 1$ and

$$P_{C_0}(1) = y_1$$

when $p = 0$.

Remark 3.2.5. Khovanov's original formulation of the basepoint multiplication map X_p on Khovanov homology is as follows. Let D be a diagram with basepoint p . Place an unknotted circle near the basepoint. The cobordism which merges the unknotted circle with the edge containing the basepoint induces a chain map of $V \otimes C_{\mathcal{D}} \rightarrow C_{\mathcal{D}}$ [30, Section 3]. The restriction of this map to x in the first factor is the map X_p . The map $P(\mathbf{t})$ can be interpreted as the extension of this construction to the geometric spectral sequence. We won't make this statement rigorous, but remark that to do so it is necessary to represent the merge cobordism as an oriented surgery arc between the marked edge and the unknotted circle. We chose the orientation from the marked edge to the unknotted circle.

We now aim to show $P(\mathbf{t})$ is a chain map.

Recall the definition of the edge homotopy maps H_m from [66]. The map $H_{\mathcal{C}}$ which satisfies the extension property. The map $H_{\mathcal{C}}$ is nonzero only for 1-dimensional configurations. There are only two kinds of purely active 1-dimensional configurations: split and join. For a 1-dimensional split configuration \mathcal{C} , $H_{\mathcal{C}}$ is given by

$$H_{\mathcal{C}}(1) = 1.$$

For a 1-dimensional join configuration \mathcal{C} with starting circles x_1 and x_2 and ending

circle y , H_C is given by

$$H_C(x_1x_2) = y.$$

Note that H_C , like the Khovanov differential, does not depend on the orientation of the arcs. Finally, we define

$$H_m = \sum_{(I,J)} H_{C(I,J)}$$

the sum is taken over 1-dimensional faces (I, J) which differ only m^{th} coordinate.

Lemma 3.2.6. *Suppose \mathbf{t} and \mathbf{t}' are decorations of the diagram \mathcal{D} that differ only at the m^{th} crossing. Then $P(\mathbf{t})$ and $P(\mathbf{t}')$ are related by the following formula:*

$$P(\mathbf{t}') = P(\mathbf{t}) + H_m P(\mathbf{t}) + P(\mathbf{t}) H_m.$$

Proof. The proof is analogous to the proof of Theorem 5.4 in [66].

Let δ denote the (unoriented) arc corresponding to the m^{th} crossing. Let (I, J) be a k -face with $I(m) = 0$ and $J(m) = 1$. Set $\mathcal{C} = \mathcal{C}(I, J, \mathbf{t})$ and $\mathcal{C}' = \mathcal{C}(I, J, \mathbf{t}')$. Let I' be obtained from I by changing the m^{th} coordinate to 1 and J' be obtained from J by changing the m^{th} coordinate to 0. It is sufficient to show the following equation holds:

$$P_{\mathcal{C}} - P_{\mathcal{C}'} = P_{\mathcal{C}(I',J)} H_{C(I,I')} + H_{C(J',J)} P_{\mathcal{C}(I,J')}, \quad (3.1)$$

where the decoration is omitted from the notation on the right hand side since \mathbf{t} and \mathbf{t}' agree on (I, J') and (I', J) .

If $\bar{\mathcal{C}}$ is disconnected, then the left hand side of (3.1) is zero. Then either both $\mathcal{C}(I, J')$ and $\mathcal{C}(I', J)$ are disconnected or $\mathcal{C}(I, I')$ and $\mathcal{C}(I', J)$ are disjoint and $\mathcal{C}(I, J')$ and $\mathcal{C}(J, J')$ are disjoint. In both cases, the right hand side is zero; the latter follows from the extension property.

We assume $\bar{\mathcal{C}}(I, J)$ is connected. The case when $\bar{\mathcal{C}}$ is 1- or 2-dimensional is left as a straightforward exercise for the reader. We now assume $k \geq 3$.




Figure 3.2: Possible positions of the arc δ in proof of Lemma 3.2.6 when $H_{C(J',J)}P_{C(I,J')} \neq 0$.




Figure 3.3: Possible positions of the dual arc δ^* in proof of Lemma 3.2.6 when $P_{C(I',J)}H_{C(I,I')} \neq 0$.

First, suppose P_C or $P_{C'}$ is non-zero; we can assume it is P_C . Then $P_{C'} = 0$. If the arc δ is join, then $P_C = H_{C(J',J)}P_{C(I,J')}$ and $P_{C(I',J)}H_{C(I,I')} = 0$. Alternatively, if δ is a split, then $P_C = P_{C(I',J)}H_{C(I,I')}$ and $H_{C(J',J)}P_{C(I,J')} = 0$.

Now, suppose

$$P_C = P_{C'} = 0. \quad (3.2)$$

In what follows, we need to show that

$$P_{C(I',J)}H_{C(I,I')} = H_{C(J',J)}P_{C(I,J')} \quad (3.3)$$

First, suppose $H_{C(J',J)}P_{C(I,J')} \neq 0$. If δ is a split arc, it must split one of circles y_2, \dots, y_m and (3.3) holds. Alternatively, if δ is a join arc, it must connect a new circle w to y_1 . The arc δ cannot meet w and x_1 , since that would contradict (3.2). Therefore, δ meets w and x_i for some $2 \leq i \leq p+1$ and again (3.3) holds. See Figure 3.2.

Alternatively, suppose $P_{C(I',J)}H_{C(I,I')} \neq 0$. If δ is a split arc, it must split a new

circle w off the central circle x_1 . The arc δ cannot split a circle off a portion of x_1 that lies in y_1 , since that would contradict the fact (3.2). Therefore, δ splits a circle off a portion of x_1 that lies in y_i for $i \leq 2 \leq q + 1$ and (3.3) holds. If δ is a join, it must join a new circle w to a circle x_i for $2 \leq i \leq p + 1$ and again (3.3) holds. See Figure 3.3. \square

Lemma 3.2.7. *Let \mathbf{t} and \mathbf{t}' be two decorations of the diagram \mathcal{D} . Then $P(\mathbf{t})$ is a chain map with respect to $d(\mathbf{t})$ if and only if $P(\mathbf{t}')$ is a chain map with respect to $d(\mathbf{t}')$.*

Proof. It is enough to consider the case when \mathbf{t} and \mathbf{t}' differ at a single crossing, say the m^{th} . The claim follows by Lemma 3.2.6, Theorem 5.4 in [66] and the fact that H_m is zero on $V(I)$ when the m^{th} coordinate of I is 1. \square

Lemma 3.2.8. *Let (I, J) be an undecorated k -face. There exists a decoration \mathbf{t} of the diagram \mathcal{D} such that*

$$\sum_K d_{\mathcal{C}(K, J, \mathbf{t})} P_{\mathcal{C}(I, K, \mathbf{t})} = \sum_K P_{\mathcal{C}(K, J, \mathbf{t})} d_{\mathcal{C}(I, J, \mathbf{t})}, \quad (3.4)$$

where the sums are taken over resolutions K such that $I < K < J$.

Proof. Set $\bar{\mathcal{C}} = \bar{\mathcal{C}}(I, J) = (x_1, \dots, x_t, \gamma_1, \dots, \gamma_k)$. By the extension property it suffices to consider only active configurations $\bar{\mathcal{C}}$.

First, we consider the case when $\bar{\mathcal{C}}$ is disconnected. Let \mathbf{t} be arbitrary. If $\bar{\mathcal{C}}$ has more than 2 components, then at least one of $\mathcal{C}(I, K, \mathbf{t})$ or $\mathcal{C}(K, J, \mathbf{t})$ is disconnected. By the disconnected rule, both sides of (3.4) vanish. If $\bar{\mathcal{C}}$ has 2 components, the only non-zero terms in (3.4) will be when $\mathcal{C}(I, K, \mathbf{t})$ is one component and $\mathcal{C}(K, J, \mathbf{t})$ is the other. In this case, the claim holds by the extension property.

If none of the circles of $\bar{\mathcal{C}}$ go through the point P , then both sides of (3.4) necessarily vanish (after choice of some \mathbf{t}).

We assume $\bar{\mathcal{C}}$ is connected and contains the point P . By orienting the arcs that meet x_p appropriately, we can choose a decoration \mathbf{t} such that $P_{\mathcal{C}(I,K,\mathbf{t})} = 0$ for all K . Then the left hand side of (3.4) vanishes. Our goal is to choose the orientation of the remaining arcs appropriately to make the right hand side vanish as well.

Let $x_1 = x_p$ in $\mathcal{C} = \mathcal{C}(I, J, \mathbf{t})$. Suppose there are two circles $x_i, i = 1, 2$ with the property that there are two (or more) arcs which meet x_i and x_1 . Since $P_{\mathcal{C}}$ vanishes on configurations where multiple arcs meeting the same pair of circles, the only non-zero terms on the right hand side can arise from a resolution K such that $\mathcal{C}_1 = \mathcal{C}(I, K, \mathbf{t})$ has type E and includes an arc meeting x_1 and x_2 and an arc meeting x_1 and x_2 . Set $\mathcal{C}_2 = \mathcal{C}(K, J, \mathbf{t})$. The central circle of \mathcal{C}_1 is marked, so $P_{\mathcal{C}(K,J,\mathbf{t})}$ vanishes unless P meets a circle y_i for $2 \leq i \leq q + 1$. However, the orientation of \mathbf{t} the right hand side is zero in this case. Thus, we can assume there is at most one such circle x_2 . Consider the resolution K where the arcs meeting x_1 and x_2 have 1-resolution. Let $x' = x_p$ denote the circle meeting p in K . Orient the remaining arcs (that is, those that did not meet x_1) so that $P_{\mathcal{C}(K,K',\mathbf{t})} = 0$ for all $K < K'$. We claim \mathbf{t} is our desired decoration.

Let γ_i be an arc that meets x_p . Surgery along arcs that do not meet x_p cannot change the type of the 1-configuration $(x_1, \dots, x_t, \gamma_i)$. Consider the left hand summand for some resolution K . Set $\mathcal{C}_1 = \mathcal{C}(I, K, \mathbf{t})$ and $\mathcal{C}_2 = \mathcal{C}(K, J, \mathbf{t})$. If x_p is not among the active circles of \mathcal{C}_1 , then the summand vanishes. It remains to consider K for which \mathcal{C}_1 is one of the five configuration types. Moreover, $P_{\mathcal{C}(K,J,\mathbf{t})}$ vanishes on monomials divisible by x_p , so we need only consider K when y_p does not divide the image of $d_{\mathcal{C}(I,K,\mathbf{t})}$. There are three cases: when \mathcal{C}_1 has type A, C or E .

First, consider the case when \mathcal{C}_1 has type A_k as shown in Figure 3.4. There are two active starting circles, $x_1 = x_p$ which meets the head of each arc of \mathcal{C}_1 and x_2 which meets the tail of each arc of \mathcal{C}_1 . There are k active ending circles. Set $y_1 = y_p$ and let y_2, \dots, y_k denote the other active ending circles. The orientation of arcs that




Figure 3.4: The case when \mathcal{C}_1 has type A_k in proof of Lemma 3.2.8.




Figure 3.5: The case when \mathcal{C}_1 has type $C_{p,q}$ in proof of Lemma 3.2.8.

do not meet x_1 have been chosen so that $P_{\mathcal{C}(K,J,t)} = 0$.

Now consider the case there are no arcs of \mathcal{C}_2 that fail to meet x_1 . Join arcs and arcs that meet a passive circle of \mathcal{C}_1 already have orientation chosen so that $P_{\mathcal{C}(K,J,t)} = 0$. Suppose an arc meets one of the ending circles y_i for $2 \leq i \leq k$. Then again we have $P_{\mathcal{C}(K,J,t)} = 0$ since the image of $d_{\mathcal{C}(I,K,t)}$ is not divisible by y_i . That leaves one remaining case: \mathcal{C}_2 has a single arc the head of which meets x_1 and the tail of which meets x_2 . Such an arc is shown in red in Figure 3.4. Then $\mathcal{C}(I, J, \mathbf{t})$ has type A and explicit computation shows the right hand side is zero.

Next, suppose \mathcal{C}_1 has type $C_{p,q}$ as shown in Figure 3.5. Again, case analysis again shows the that $P_{\mathcal{C}(K,J,t)} = 0$ except when there are (possibly) arcs of the form shown in red in the figure. In either case, $\mathcal{C}(I, J, \mathbf{t})$ has type C and explicit computation shows the right hand side is zero.

Finally, suppose \mathcal{C}_1 has type $E_{p,q}$. The analysis above showed the the point p must be on an edge that meets a circle y_i for $2 \leq i \leq q + 1$, there is only one split arc in

\mathcal{C}_1 and there are no join arcs. Again by case analysis, we see $P_{\mathcal{C}(K,J,\mathbf{t})} = 0$ except for a single edge of the type shown in red. The configuration $\mathcal{C}(I, J, \mathbf{t})$ is 2-dimensional and it is easily seen that the right hand side vanishes. \square

We are now ready to complete the proof of Proposition 1.1.4. The image of $P(\mathbf{t})$ lies in the subcomplex $\tilde{\mathcal{C}}(\mathcal{D}, \mathbf{t}, p)$ and $P(\mathbf{t})$ vanishes on $\tilde{\mathcal{C}}(\mathcal{D}, \mathbf{t}, p)$. Therefore, P descends to a map $\bar{P}(\mathbf{t}) : \bar{\mathcal{C}}(\mathcal{D}, \mathbf{t}, p) \rightarrow \tilde{\mathcal{C}}(\mathcal{D}, \mathbf{t}, p)$. The map X_p induces an isomorphism between $\overline{Kh}(L)$ and $\widetilde{Kh}(L)$; see [64]. Thus, $\bar{P}(\mathbf{t})$ induces an isomorphism between $\bar{E}_2(L, c)$ and $\tilde{E}_2(L, c)$. By the mapping lemma for spectral sequences $\bar{P}(\mathbf{t})$ induces an isomorphism $\bar{E}_k(L, c) \cong \tilde{E}_k(L, c)$ for $k \geq 2$. \square

Let \mathcal{D} be a diagram with a distinguished crossing c . Let \mathcal{D}_0 and \mathcal{D}_1 be the diagrams obtained by replacing c by its 0 or 1 resolution, respectively. Recall [55] the set of quasi-alternating links \mathcal{Q} is the smallest set of links satisfying

- the unknot is in \mathcal{Q} , and
- if L is a link which admits a diagram \mathcal{D} with a distinguished crossing c such that $\mathcal{D}_0, \mathcal{D}_1 \in \mathcal{Q}$ and $\det(\mathcal{D}_0), \det(\mathcal{D}_1) \neq 0$ and $\det(L) = \det(\mathcal{D}_0) + \det(\mathcal{D}_1)$, then $L \in \mathcal{Q}$.

All alternating links are quasi-alternating. Ozsváth and Szabó [55] showed that the Heegaard Floer variant of the spectral sequence E_k^{HF} collapses at the E_2^{HF} -page for quasi-alternating links. A link L is called δ -thin if its reduced Khovanov homology is supported in a single δ -grading. Manolescu and Ozsváth [49] showed that quasi-alternating knots are δ -thin. Since the higher differentials $d_k(\mathbf{t})$ decrease the δ -grading by 2, the reduced spectral sequence \tilde{E}_k necessarily collapses at the E_2 -page for δ -thin links. Thus, we have shown the following proposition:

Proposition 3.2.9. *Let L be a δ -thin knot. Then the reduced spectral sequence $\tilde{E}_k(L)$ collapses at the E_2 -page.*

Combined with the twin arrows conjecture, this would imply the spectral sequence $E_k(L)$ collapses at the E_2 -page for δ -thin links.

Using the transverse invariant, Baldwin [4] computed E_k^{HF} for the torus knots $T(3, 4)$ and $T(3, 5)$ with a well-defined quantum grading. Our computations agree with his. In addition, he showed that E_k^{HF} distributes over connect sum of links. We prove the following analogous result.

Proposition 3.2.10. *Let L and L' be links with distinguished components c, c' , respectively. Let $L\#L'$ denote the connect sum between the distinguished components and let c'' denote the resulting component. Then*

$$\tilde{E}_k(L\#L', c'') \cong \tilde{E}_k(L, c) \otimes \tilde{E}_k(L', c')\{1\}$$

for $k \geq 2$.

Proof. We show the underlying chain complexes are equal. Choose decorated diagrams $(\mathcal{D}, \mathbf{t})$ and $(\mathcal{D}', \mathbf{t}')$ of L, L' , respectively, such that a decorated diagram $(\mathcal{D}'', \mathbf{t}'')$ of $L\#L'$ is obtained from the disjoint union of \mathcal{D} and \mathcal{D}' by surgery along an arc connecting the distinguished components. Choose points p and p' on the edges of $\mathcal{D}, \mathcal{D}'$, respectively, meeting the connect sum surgery arc. Choose a point P'' of \mathcal{D}'' on either edge resulting from the surgery; they will always belong to the same circle. We have $C_{\mathcal{D}} \otimes C_{\mathcal{D}'}\{1\}$ is isomorphic to $C_{\mathcal{D}''}$ as graded vector spaces, the grading shift coming from the fact the connect sum joins two circles which always carry the generator with grading -1 . The tensor product accounts for all faces which resolve only crossings of \mathcal{D} or \mathcal{D}' . We argue $d''(\mathbf{t})$ has no other terms. Inspection shows that only configurations of type E can be expressed as a connect sum of two nontrivial configurations. However, such a decomposition will be along the central circle of the configuration which in this case is marked. Thus, such configurations do not contribute to $d''(\mathbf{t})$. □

Recall, $E'(L)$ denotes the spectral sequence arising from the alternate differential defined using mirror configuration types. We have the following result [12, Corollary 1.7]:

Theorem 3.2.11 (Beier). *Let L be a knot or link. For $k \geq 2$, we have $E_k(L) \cong E'_k(L)$.*

Proof. Let D be a diagram for L . Let D' be the diagram obtained by mirroring D and reversing the strands in each crossing. D and D' are diagrams for the same link. For example, put D in standard position in the xy -plane and rotate by π about the x -axis. Then $C'(D, \mathbf{t}) = C(D', \mathbf{t})$, so by Reidemeister invariance, the theorem holds. □

(Compare [47, Proposition 6.5].)

3.3 Conjectures

In this section, we present results of computations of the spectral sequence E_k . Unless stated otherwise, the conjectures in this section have been verified on all prime links with 12 or fewer crossings, all prime knots with 14 or fewer crossings, and all torus knots $T_{p,q}$ where $(p-1)q \leq 16$. We used knot data from two sources. We extracted the planar diagram (PD) description of the Rolfsen knot tables from Bar-Natan's `KnotTheory` package [6]. In addition, we used the HTW knot tables [23] and the Thistlethwaite link (MT) tables and from SnapPy [16]. The HTW knot tables and MT link tables are encoded with Dowker-Thistlethwaite (DT) codes. The HTW tables included all prime knots through 16 crossings and the MT tables include all prime links through 14 crossings.

In order to simplify presentation of the results of the computations, we begin with the following two conjectures.




Figure 3.6: Diagrammatic representation of the Poincaré polynomials $p_\ell^k(t, q)$.

Conjecture 3.3.1. *The reduced theory does not depend on the choice of distinguished component, that is, if L is a link and c and c' are components of L , then*

$$\tilde{E}_k(L, c) \cong \tilde{E}_k(L, c')$$

for $k \geq 2$.

Thus we will write $\tilde{E}(L)$ for the reduced spectral sequence of a link L .

Conjecture 3.3.2 (Twin Arrows). *Let L be a link. For $k \geq 2$, the page $E_k(L)$ is isomorphic to two copies of $\tilde{E}_k(L)$. Specifically,*

$$E_k(L) \cong \tilde{E}_k(L)\{-1\} \oplus \tilde{E}_k(L)\{1\}.$$

Bloom [14] conjectured the structure of the monopole Floer homology variant of the spectral sequence for torus knots $T(3, 6n \pm 1)$. Our computations agree with his

conjecture for $T(3, 5)$ and $T(3, 7)$. Based on our computations for torus links $T(3, n)$ with $n \leq 9$, we extend his conjecture to all 3-strand torus links as follows.

Conjecture 3.3.3. *Set*

$$f_j(t, q) = \sum_{i=0}^{j-1} t^{8i} q^{12i}$$

$$p_2^0(t, q) = p_0^3(t, q) = p_0^4(t, q) = 1 + tq^2$$

$$p_2^1(t, q) = 1 + tq^2 + t^2q^2 + 2t^2q^4$$

$$p_3^1(t, q) = p_1^4(t, q) = tq^2 + 2t^2q^4$$

$$p_2^2(t, q) = 1 + tq^2 + t^2q^2 + t^3q^6$$

$$p_3^2(t, q) = p_2^4(t, q) = tq^2 + t^3q^6$$

$$p_2^3(t, q) = 1 + tq^2 + t^2q^2 + t^3q^6 + t^4q^6 + t^5q^8$$

$$p_3^3(t, q) = tq^2 + t^4q^6$$

$$p_4^3(t, q) = 0$$

$$p_2^4(t, q) = 1 + tq^2 + t^2q^2 + t^3q^6 + t^4q^6 + t^5q^8 + t^6q^8 + 2t^6q^{10}$$

$$p_3^4(t, q) = tq^2 + t^4q^6 + t^6q^8 + 2t^6q^{10}$$

$$p_4^4(t, q) = t^6q^8 + 2t^6q^{10}$$

$$p_2(t, q) = p_5^2(t, q) = 1 + tq^2 + t^2q^2 + t^3q^6 + t^4q^6 + t^5q^8 + t^6q^8 + t^7q^{12}$$

$$p_3(t, q) = p_5^3(t, q) = tq^2 + t^4q^6 + t^6q^8 + t^7q^{12}$$

$$p_4(t, q) = p_5^4(t, q) = t^6q^8 + t^7q^{12}.$$

For all $n > 1$, the spectral sequence for the torus link $T(3, n)$ collapses at the E_4 -page.

Moreover, writing $n = 2 + 6j + \ell$ with $j \geq 0$ and $0 \leq \ell < 6$, the Poincaré polynomial

for $\tilde{E}_k(T(3, n))$, $k = 2, 3, 4$ is given by

$$\tilde{P}_k(T(3, n))(t, q) = q^{2j-3}(1 + t^2q^4(f_j(t, q)p_k(q, t) + t^{7j}q^{12j}p_k^\ell(q, t))).$$

Diagrammatic representations of the Poincaré polynomials $p_k^\ell(t, q)$ are given in Figure 3.6.

Bloom [13] showed that odd Khovanov homology is mutation invariant. In addition, his argument showed that both the Heegaard Floer and monopole Floer variants of the spectral sequence are mutation invariant. Using the mutant knot tables compiled by Stoimenow [65], we verified that E_k is invariant under mutation for all (5300) mutant knot groups through 14 crossings. Thus we make the following conjecture.

Conjecture 3.3.4. *Let K be a knot. For $k \geq 2$, $E_k(K)$ is invariant under mutation.*

Chapter 4

Twisting the Geometric Spectral Sequence

In this chapter, we define the twisted geometric chain complex $C^{\text{tw}}(\mathcal{D}, \mathbf{t}, w)$ and study its properties. In Section 4.1 we define the chain complex. In Section 4.2, we study how $d^{\text{tw}}(\mathbf{t})$ transforms under change of decoration \mathbf{t} . In Section 4.3, we prove $d^{\text{tw}}(\mathbf{t})$ is a differential. Finally, in Section 4.4, we study the total homology of the complex and establish Theorem 1.1.5.

Throughout this chapter, $C(\mathcal{D}, \mathbf{t}, w) = C^{\text{tw}}(\mathcal{D}, \mathbf{t}, w)$ will denote the twisted geometric chain complex, $d(\mathbf{t}) = d^{\text{tw}}(\mathbf{t})$ the differential, etc.

4.1 The twisted chain complex

The goal of this section is to define the unreduced twisted chain complex

$$C(\mathcal{D}, \mathbf{t}, w) = (C_{\mathcal{D}}, d(\mathbf{t})).$$

Let L be a link and $(\mathcal{D}, \mathbf{t})$ a decorated diagram for L . The module underlying the twisted chain complex is the same as for Khovanov homology. Our twisted complex

will depend on a set of weights $w(e)$, one for each edge e of \mathcal{D} , chosen such that

$$\prod_e w(e) = 1.$$

For a circle x in a resolution I , let $w(x)$ denote the product of the weights of the edges contained in x :

$$w(x) = \prod_{e \subset x} w(e).$$

To define the twisted differential $d(\mathbf{t})$, we will proceed as in the definition of the geometric spectral sequence: we will give a family of maps $d_{\mathcal{C}}$ which will determine the differential $d(\mathbf{t})$. This family satisfies a slightly different set of properties than $d_{\mathcal{C}}^{\text{gss}}$; we give their definitions now. As before, let $G_{\mathcal{C}} : V_0(\mathcal{C}) \rightarrow V_1(\mathcal{C})$ be a family of maps.

Definition 4.1.1 (Weak Conjugation Rule). We say $G_{\mathcal{C}}$ satisfies the weak conjugation rule if, for every configuration \mathcal{C} and every pair of monomials $a \in V_0(\mathcal{C})$ and $b \in V_1(\mathcal{C})$, the coefficient of $G_{\mathcal{C}}(a)$ at b is non-zero if and only if the coefficient of $G_{r(\mathcal{C})}(a)$ at b is non-zero.

Definition 4.1.2 (Weak Duality Rule). We say $G_{\mathcal{C}}$ satisfies the weak duality rule if, for every configuration \mathcal{C} and every pair of monomials $a \in V_0(\mathcal{C})$ and $b \in V_1(\mathcal{C})$, the coefficient of $G_{\mathcal{C}}(a)$ at b is non-zero if and only if the coefficient of $G_{m(\mathcal{C})}(b^*)$ at a^* is non-zero.

Definition 4.1.3 (Conjugation Duality Rule). We say $G_{\mathcal{C}}$ satisfies the conjugation duality rule if, for every configuration \mathcal{C} and every pair of monomials $a \in V_0(\mathcal{C})$ and $b \in V_1(\mathcal{C})$, the coefficient of $G_{\mathcal{C}}(a)$ at b is equal to if the coefficient of $G_{r(m(\mathcal{C}))}(b^*)$ at a^* .

Definition 4.1.4 (Rotation Rule). We say $G_{\mathcal{C}}$ satisfies the rotation rule if, for all pairs configurations $\mathcal{C}, \mathcal{C}'$ which differ by one of the local rotation moves shown in




Figure 4.1: Local rotation moves.




Figure 4.2: Decorated local trading move.

Figure 4.1, we have that

$$G_{\mathcal{C}} \equiv G_{\mathcal{C}'}$$

under the natural identification between the starting and ending circles.

Definition 4.1.5 (Trading Rule). Consider a pair of configurations \mathcal{C} and \mathcal{C}' that differ by one of the local rotation moves shown in Figure 4.2. Suppose \mathcal{C} contains the degree 1 starting circle which we denote by x_1 and \mathcal{C}' contains the dual degree 1 circle y_1 . We have the decompositions:

$$V_0(\mathcal{C}) = V_0(\mathcal{C}') \otimes V(x_1) \text{ and } V_1(\mathcal{C}') = V_1(\mathcal{C}) \otimes V(y_1).$$

We say $G_{\mathcal{C}}$ satisfies the trading rule if, for all such pairs $(\mathcal{C}, \mathcal{C}')$ we have that

$$G_{\mathcal{C}}(a \cdot b) = G_{\mathcal{C}'}(a) \cdot b^*.$$

The family of maps $d_{\mathcal{C}}$ satisfies the weak conjugation rule, weak duality rule, conjugation duality rule, filtration rule, grading rule, rotation rule and trading rule.

Moreover, it satisfies the extension formula and disconnected rule for configurations with dimension at least 1. It does not satisfy the duality or conjugation rules individually. In the remainder of this section, we define the family $d_{\mathcal{C}}$. To begin, we define $d_{\mathcal{C}}$ for 0-dimensional configurations.

Definition 4.1.6. Let $\mathcal{C} = (x_1, \dots, x_t)$ be a 0-dimensional configuration and $a \in V_0(\mathcal{C})$ a monomial. Then $d_{\mathcal{C}}$ is given by

$$d_{\mathcal{C}}(a) = \sum_{i=1}^t (1 + w(x_i)) a \cdot x_i.$$

By the extension formula and disconnected rule for configurations of dimension at least one, it suffices to define $d_{\mathcal{C}}$ for connected, purely active configurations. Next, we define $d_{\mathcal{C}}$ for 1-dimensional configurations.

Definition 4.1.7. Let \mathcal{C} be a purely active (connected) 1-dimensional configuration. There are two cases: join and split.

In the join case, there are two starting circles. Suppose the arc goes from x_1 to x_2 . There is a unique ending circle; denote it by y . Then $d_{\mathcal{C}}$ is given by

$$\begin{aligned} d_{\mathcal{C}}(1) &= 1 \\ d_{\mathcal{C}}(x_1) &= y \\ d_{\mathcal{C}}(x_2) &= w(x_1)y \\ d_{\mathcal{C}}(x_1x_2) &= 0. \end{aligned}$$

In the split case, there is a unique starting circle x and suppose the dual arc goes from y_1 to y_2 . Then $d_{\mathcal{C}}$ is given by

$$\begin{aligned} d_{\mathcal{C}}(1) &= w(y_1)y_1 + y_2 \\ d_{\mathcal{C}}(x) &= y_1y_2. \end{aligned}$$




Figure 4.3: The 14 2-dimensional purely active, connected configurations up to rotation and isotopy in the sphere.

Definition 4.1.8. Let \mathcal{C} be an purely active, connected 2-dimensional configuration. Up to isotopy and rotation, there are 14 configurations in the sphere as shown in Figure 4.3.

We now define $d_{\mathcal{C}}$ for these configurations. We only list the non-zero terms of $d_{\mathcal{C}}$. If a type of configuration is not mentioned, $d_{\mathcal{C}}$ is identically zero on configurations of that type.

If \mathcal{C} has type 1, then

$$d_{\mathcal{C}}(1) = 1.$$

If \mathcal{C} has type 2, then

$$d_{\mathcal{C}}(x_1x_2) = y_1y_2.$$

If there is a starting circle which meets both arcs, let it be denoted by x_1 . If there is an ending circle which meets both dual arcs, let it be denoted by y_1 .

If \mathcal{C} has type 3, then

$$d_{\mathcal{C}}(x_2x_3) = w(x_1)y_1.$$

If \mathcal{C} has type 6, then

$$\begin{aligned} d_{\mathcal{C}}(x_2) &= \frac{w(x_1)}{w(y_2)}y_1 \\ &= \frac{w(y_1)}{w(x_2)}y_1. \end{aligned}$$

If \mathcal{C} has type 10, then

$$d_{\mathcal{C}}(1) = w(y_1)y_2y_3.$$

If \mathcal{C} has type 5, then

$$d_{\mathcal{C}}(x_2x_3) = y_1.$$

If \mathcal{C} has type 8, then

$$d_{\mathcal{C}}(x_2) = y_1.$$

If \mathcal{C} has type 11, then

$$d_{\mathcal{C}}(1) = y_1.$$

If \mathcal{C} has a type 13, then

$$d_{\mathcal{C}}(1) = 1 \text{ and } d_{\mathcal{C}}(x_1) = y_1.$$

For $k > 2$, we distinguish six kinds of k -dimensional configurations for which $d_{\mathcal{C}}$ is non-zero. The map $d_{\mathcal{C}}$ is defined as in [66, Section 4] for configurations of type $A_k, B_k, C_{p,q}$ and $D_{p,q}$. We review the definition in those cases now.

Definition 4.1.9. A k -dimensional purely active, connected configuration

$$\mathcal{C} = (x_1, \dots, x_s, \gamma_1, \dots, \gamma_k)$$

has type A_k if, for each pair (i, j) with $1 \leq i < j \leq k$, every 2-dimensional subconfiguration $(x_1, \dots, x_s, \gamma_i, \gamma_j)$ has type 1. It follows that \mathcal{C} has two active starting circles

and k active ending circles. In this case, we define

$$d_{\mathcal{C}}(1) = 1.$$

Definition 4.1.10. A k -dimensional purely active, connected configuration \mathcal{C} has type B_k if $m(\mathcal{C}^*)$ has type A_k . It follows that \mathcal{C} has k active ending circles. In this case, we define

$$d_{\mathcal{C}}(x_1 x_2 \cdots x_k) = y_1 y_2.$$

Definition 4.1.11. Let \mathcal{C} be a purely active, connected k -dimensional configuration with the property that it has a unique starting circle x_1 . This separates the sphere into two regions, and the set of arcs decomposes as

$$\{\gamma_1, \dots, \gamma_k\} = \{e_1, \dots, e_p\} \cup \{f_1, \dots, f_q\},$$

where the e_i lie on one side of x_1 and the f_j the other. We say \mathcal{C} has type $C_{p,q}$ if $p, q \geq 1$ and for each pair (i, j) with $1 \leq i \leq p$ and $1 \leq j \leq q$, the 2-dimensional configuration (x_1, e_i, f_j) has type 11. In this case, we define

$$d_{\mathcal{C}}(1) = 1.$$

Definition 4.1.12. A k -dimensional purely active, connected configuration \mathcal{C} has type $D_{p,q}$ if $m(\mathcal{C}^*)$ has type $C_{p,q}$. In this case, we define

$$d_{\mathcal{C}}(x_1 x_2 \cdots x_{k-1}) = y_1.$$

Besides the contribution for 0-dimensional configurations, our map $d_{\mathcal{C}}$ differs from Szabó's on configurations of type $E_{p,q}$. We break configurations of type $E_{p,q}$ into two cases: $E_{p,q}^{\text{out}}$ and $E_{p,q}^{\text{in}}$.

Definition 4.1.13. An k -dimensional purely active, connected configuration $\mathcal{C} = (x_1, \dots, x_s, \gamma_1, \dots, \gamma_k)$ with $p+1$ active starting and $q+1$ active ending circles has type $E_{p,q}^{\text{out}}$ if, for every distinct pair $1 \leq i < j \leq k$, the subconfiguration $(x_1, \dots, x_s, \gamma_i, \gamma_j)$ has type 3, 6 or 10. A type $E_{p,q}$ configuration has a unique active starting circle x_1 , called the central starting circle, with the property that it meets all the arcs γ_i . All the other starting circles have degree 1. Similarly, there is a unique ending circle y_1 , called the central ending circle, with the property that it meets all the dual arcs γ_i^* . The other active ending circles have degree 1. In this case, we define

$$\begin{aligned} d_{\mathcal{C}}(x_2 x_3 \cdots x_{p+1}) &= \frac{w(y_1)}{w(x_2)w(x_3) \cdots w(x_{p+1})} y_1 \\ &= \frac{w(x_1)}{w(y_2)w(y_3) \cdots w(y_{q+1})} y_1. \end{aligned}$$

Definition 4.1.14. An k -dimensional purely active, connected configuration $\mathcal{C} = (x_1, \dots, x_s, \gamma_1, \dots, \gamma_k)$ with $p+1$ active starting and $q+1$ active ending circles has type $E_{p,q}^{\text{in}}$ if, for every distinct pair $1 \leq i < j \leq k$, the subconfiguration $(x_1, \dots, x_s, \gamma_i, \gamma_j)$ has type 5, 8 or 11. As in the previous definition, there is a central starting circle x_1 and central ending circle y_1 , and all other active starting and ending circles have degree 1. In this case, we define

$$d_{\mathcal{C}}(x_2 x_3 \cdots x_{p+1}) = y_1.$$

Our $d_{\mathcal{C}}$ agrees with Szabó's $d_{\mathcal{C}}^{\text{gss}}$ on configurations of type $E_{p,q}^{\text{in}}$ but not $E_{p,q}^{\text{out}}$. In the latter case, our map differs by a scalar coefficient.

Definition 4.1.15. We say a k -dimensional configuration \mathcal{C} has type $E_{p,q}$ if it has type $E_{p,q}^{\text{out}}$ or $E_{p,q}^{\text{in}}$.

We say a configuration \mathcal{C} has type A_k , etc., if its active part \mathcal{C}_0 has type A_k , etc.

We are now ready to define the differential $d(\mathbf{t})$.

Definition 4.1.16. Let

$$d(\mathbf{t}) = \sum_{k=0}^n d_k(\mathbf{t}),$$

where

$$d_k(\mathbf{t}) = \sum_{\text{faces } (I, J)} d_{\mathcal{C}(I, J, \mathbf{t})}$$

and the sum is taken over all k -dimensional faces (I, J) .

4.2 Dependence on the decoration

In this section we discuss the dependence of $d(\mathbf{t})$ on the decoration \mathbf{t} . We now recall the definition of Szabó's change of decoration homotopy H_m .

Definition 4.2.1. Let \mathcal{C} be a 1-dimensional configuration. We define

$$H_{\mathcal{C}} : V_0(\mathcal{C}) \rightarrow V_1(\mathcal{C})$$

by defining $H_{\mathcal{C}_0}$ for active configurations and extending it to all of \mathcal{C} by the extension formula of Definition 2.2.1. There are two cases: If \mathcal{C}_0 is a join, we define

$$H_{\mathcal{C}_0}(x_1 x_2) = y_1.$$

If \mathcal{C}_0 is a split, we define

$$H_{\mathcal{C}_0}(1) = 1.$$

Definition 4.2.2. Define $H_m : C_{\mathcal{D}} \rightarrow C_{\mathcal{D}}$ by

$$H_m = \sum_{(I, J)} H_{\mathcal{C}(I, J)}$$

where the summation is over 1-dimensional edges (I, J) with $I(m) = 0$ and $J(m) = 1$.

Theorem 4.2.3. *Suppose \mathbf{t} and \mathbf{t}' are decorations of \mathcal{D} that differ only at the m^{th} crossing. Then we have that*

$$d(\mathbf{t}') = d(\mathbf{t}) + H_m d(\mathbf{t}) + d(\mathbf{t}) H_m. \quad (4.1)$$

Proof. Let (I, J) be a k -dimensional face. Set $\mathcal{C} = \mathcal{C}(I, J, \mathbf{t})$ and $\mathcal{C}' = \mathcal{C}(I, J, \mathbf{t}')$. If $I(m) = J(m)$, including the case $k = 0$, then H_m is zero in the component of (4.1) from $V(I)$ to $V(J)$ is zero. Moreover, \mathbf{t} and \mathbf{t}' agree on $\mathcal{C}(I, J)$, so (4.1) holds.

Next, suppose $k \geq 1$, $I(m) = 0$ and $J(m) = 1$. Let I' and J' be the resolutions that differ with I and J , respectively, only at m . Set $\mathcal{C} = \mathcal{C}(I, J, \mathbf{t})$ and $\mathcal{C}' = \mathcal{C}(I', J', \mathbf{t}')$. To establish the lemma, it suffices to show that for all such (I, J) :

$$d_{\mathcal{C}} - d_{\mathcal{C}'} = d_{\mathcal{C}(I', J)} H_{\mathcal{C}(I, I')} + H_{\mathcal{C}(J', J)} d_{\mathcal{C}(I, J')}. \quad (4.2)$$

Note that \mathbf{t} and \mathbf{t}' agree on the $(k-1)$ -dimensional configurations $\mathcal{C}(I, J')$ and $\mathcal{C}(I', J)$, so we can safely omit it from the notation.

Next, suppose $k = 1$. We have $I' = J$ and $J' = I$. The terms of d_0 that multiply by circles that do not meet δ commute with $H_{\mathcal{C}(I, J)}$, so it suffices to consider the active part \mathcal{C}_0 . There are two cases: δ is a join arc or a split arc. Both cases follow by direct computation.

Next, suppose $k \geq 2$. By the extension formula for $d(\mathbf{t})$ and H_m it is enough to verify (4.2) when \mathcal{C} is purely active. By conjugation duality for $d_{\mathcal{C}}$ and $H_{\mathcal{C}}$, if the pair (\mathcal{C}, δ) satisfies (4.2), then so does the conjugate dual $(r(m(\mathcal{C}^*)), \delta^*)$.

Next, suppose \mathcal{C} is disconnected. By the disconnected rule we have that $d_{\mathcal{C}} = d_{\mathcal{C}'} = 0$. If the arcs other than δ lie in different components, then $d_{\mathcal{C}} = d_{\mathcal{C}'} = 0$. Otherwise, δ lies in a different component from the other arcs, in which case

$$d_{\mathcal{C}(I', J)} H_{\mathcal{C}(I, I')} = H_{\mathcal{C}(J', J)} d_{\mathcal{C}(I, J')}$$

by the extension formula for $d(\mathbf{t})$ and H_m .

We leave the case $k = 2$ as a routine check for the interested reader. The list of purely active, connected 2-dimensional configurations up isotopy and rotation are shown in Figure 4.3.

Next, assume $k \geq 3$ and $d_{\mathcal{C}}$ or $d_{\mathcal{C}'}$ is non-zero. Without loss of generality, assume $d_{\mathcal{C}}$ is non-zero. That means \mathcal{C} is of the six non-trivial type A - D , E^{in} or E^{out} . In each case, changing the orientation of a single arc gives a configuration with trivial contribution, so $d_{\mathcal{C}'} = 0$. When \mathcal{C} has type A , C , E^{in} or E^{out} , it is straightforward to verify one of the following two pairs of equations hold:

$$H_{\mathcal{C}(J',J)}d_{\mathcal{C}(I,J')} = d_{\mathcal{C}} \text{ and } d_{\mathcal{C}(I',J)}H_{\mathcal{C}(I,I')} = 0. \quad (4.3)$$

or

$$H_{\mathcal{C}(J',J)}d_{\mathcal{C}(I,J')} = 0 \text{ and } d_{\mathcal{C}(I',J)}H_{\mathcal{C}(I,I')} = d_{\mathcal{C}}. \quad (4.4)$$

The case \mathcal{C} has type B or D follows by the conjugation duality argument above.

Finally, suppose $k \geq 3$ and $d_{\mathcal{C}} = 0$ and $d_{\mathcal{C}'} = 0$. By conjugation duality of $d_{\mathcal{C}}$ and $H_{\mathcal{C}}$, it suffices to consider the case that $H_{\mathcal{C}(J',J)}d_{\mathcal{C}(I,J')}$ is non-zero. There are three cases for δ in \mathcal{C} :

- δ joins two active ending circles of $\mathcal{C}(I, J')$,
- δ joins an active ending circles of $\mathcal{C}(I, J')$ to a passive circle w of $\mathcal{C}(I, J')$, or
- δ splits an active ending circle of $\mathcal{C}(I, J')$.

For each of the six non-trivial types A - D , E^{in} and E^{out} for $\mathcal{C}(I, J')$ and the three cases for δ , it is straightforward to verify that

$$H_{\mathcal{C}(J',J)}d_{\mathcal{C}(I,J')} = d_{\mathcal{C}(I',J)}H_{\mathcal{C}(I,I')}. \quad (4.5)$$

□

4.3 d is a differential

The goal of this section is to prove the map $d(\mathbf{t})$ is a differential. A rough outline of the strategy is as follows. We will use an inductive argument to show that the square does not depend on the decoration \mathbf{t} . Thus, it will be enough to show for undecorated configurations $\bar{\mathcal{C}}$ that there is some decoration \mathbf{t} that makes each term of the square zero. The key lemma is the following.

Lemma 4.3.1. *Let $\bar{\mathcal{C}}$ be an undecorated k -dimensional configuration, $k \geq 1$ and $a \in V_0(\bar{\mathcal{C}})$ and $b \in V_1(\bar{\mathcal{C}})$ be monomials. For every such triple $(\bar{\mathcal{C}}, a, b)$ there exists a decoration \mathbf{t} of $\bar{\mathcal{C}}$ so that the coefficient of*

$$\sum_{I \in \mathcal{R}_k} d_{\mathcal{C}(\mathbf{0}, I)} d_{\mathcal{C}(I, \mathbf{1})}(a) \tag{4.6}$$

at b is zero, where $\mathcal{C} = (\bar{\mathcal{C}}, \mathbf{t})$.

Then we will simplify the cases to consider by using the properties satisfied by $d_{\mathcal{C}}$ and studying the behavior of the square on certain local configurations. Finally, using these simplifications, it will be possible to analyze the general case when some summand $d_{k-i} d_i$ of the square is non-zero.

We now state the main theorem and reduce it to Lemma 4.3.1.

Theorem 4.3.2. *For every decorated diagram $(\mathcal{D}, \mathbf{t})$ and choice of weights $w(e)$ we have that*

$$\sum_{i=0}^k d_i(\mathbf{t}) d_{k-i}(\mathbf{t}) = 0. \tag{4.7}$$

Proof. We will proceed by induction on k . For $k = 0$, (4.7) follows from the fact the action of $d_0(\mathbf{t})$, multiplying a circle, is commutative.

Next, we argue that (4.7) does not depend on \mathbf{t} . Let \mathbf{t} and \mathbf{t}' be two decorations that differ at a single crossing m and suppose (4.7) holds for all $(k - 1)$ -dimensional

faces. By Theorem 4.2.3, we have that

$$d_j(\mathbf{t}') = d_j(\mathbf{t}) + d_{j-1}(\mathbf{t})H_m + H_md_{j-1}(\mathbf{t}).$$

With this and the trivial observations

$$H_mH_m = 0 \text{ and } H_md_jH_m = 0,$$

it is straightforward to verify that

$$\sum_{i=0}^k d_i(\mathbf{t})d_{k-i}(\mathbf{t}) = \sum_{i=0}^k d_i(\mathbf{t}')d_{k-i}(\mathbf{t}').$$

Finally, fix a k -dimensional face (I, J) . The component of (4.7) corresponding to (I, J) is equal to

$$\sum_{I \in \mathcal{R}_k} d_{\mathcal{C}(\mathbf{0}, I)} d_{\mathcal{C}(I, \mathbf{1})}$$

where $\mathcal{C} = \mathcal{C}(I, J, \mathbf{t})$. The theorem now follows from Lemma 4.3.1. \square

The remainder of the section is devoted to proving Lemma 4.3.1.

Proof of Lemma 4.3.1. We begin with some terminology. We will think of I as being both an element of \mathcal{R}_k and the set $\{i \mid I(i) = 1\}$. We will denote by I' the complement of I . We call the configurations $\mathcal{C}(\mathbf{0}, I)$ and $\mathcal{C}(I, \mathbf{1})$ corresponding to I a decomposition of \mathcal{C} . If, for a given $\bar{\mathcal{C}}$, the statement of the Lemma 4.3.1 holds for all $a \in V_0(\bar{\mathcal{C}})$ and $b \in V_1(\bar{\mathcal{C}})$, we say that Lemma 4.3.1 holds for $\bar{\mathcal{C}}$. Similarly, if, for a given pair $(\bar{\mathcal{C}}, a)$ the statement holds for all $b \in V_1(\bar{\mathcal{C}})$, we say that Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a)$.

Lemma 4.3.3. *If Theorem 4.3.1 holds for $\bar{\mathcal{C}}_0$, it holds for $\bar{\mathcal{C}}$.*

Proof. This follows from the extension formula for $d_{\mathcal{C}}$ and the fact that

$$d_0(a \otimes p) = d_0(a) \otimes p + a \otimes d_0(p). \quad \square$$



Figure 4.4: Local “stalk” configuration: a join arc that meets a degree 1 circle x_2 .

Lemma 4.3.4. *If $\bar{\mathcal{C}}$ is purely active and disconnected, Theorem 4.3.1 holds for $\bar{\mathcal{C}}$.*

Proof. If $\bar{\mathcal{C}}$ has more than two components, $d_{\mathcal{C}(0,I)}$ or $d_{\mathcal{C}(I,1)}$ is zero by the disconnected rule. If $\bar{\mathcal{C}}$ has exactly two components, then there are two potentially non-zero decompositions, but they cancel by the extension formula. \square

From now on we will assume $\bar{\mathcal{C}}$ is purely active and connected.

Lemma 4.3.5. *Lemma 4.3.1 holds for the triple $(\bar{\mathcal{C}}, a, b)$ if and only if it holds for $(m(\bar{\mathcal{C}}^*), b^*, a^*)$.*

Proof. Let \mathbf{t} be a decoration for $(\bar{\mathcal{C}}, a, b)$ satisfying Lemma 4.3.1. Set $\mathcal{C} = (\bar{\mathcal{C}}, \mathbf{t})$. Then \mathbf{t}' in $r(m(\mathcal{C}^*)) = (m(\bar{\mathcal{C}}^*), \mathbf{t}')$ is the desired decoration by the conjugation duality rule. \square

Lemma 4.3.6. *Suppose $\bar{\mathcal{C}}$ and $\bar{\mathcal{C}}'$ are purely active and connected undecorated configurations which differ by a rotation move shown in Figure 4.1. Then Theorem 4.3.1 holds for $\bar{\mathcal{C}}$ if and only if it holds for $\bar{\mathcal{C}}'$.*

Proof. This is an immediate consequence of the rotation rule for $d_{\mathcal{C}}$. \square

Lemma 4.3.7. *Suppose $\bar{\mathcal{C}}$ contains the local “stalk” configuration shown on the left side of Figure 4.4 where there is a degree 1 circle meeting an arc γ_1 . Let $a \in V_0(\bar{\mathcal{C}})$ be a monomial and denote the other circle meeting γ_1 by x_2 . If x_2 divides a or x_1 does not divide a , then Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a)$.*

Proof. Let \mathbf{t} be a decoration of $\bar{\mathcal{C}}$ that locally agrees with the right side of Figure 4.4. When x_1 does not divide a , there are only two possible non-trivial decompositions




Figure 4.5: Local trading move.




Figure 4.6: The 3 purely active, connected 2-dimensional undecorated configurations up to rotation, trading and isotopy.

corresponding to $I = \{1\}$ and $I' = \{1\}$. The two terms cancel since the complementary configurations are equivalent. When x_1 and x_2 both divide a , there are no such decompositions. \square

Lemma 4.3.8. *Suppose $\bar{\mathcal{C}}$ and $\bar{\mathcal{C}}'$ differ by the trading move shown in Figure 4.5. Then Lemma 4.3.1 holds for $\bar{\mathcal{C}}$ if and only if it holds for $\bar{\mathcal{C}}'$.*

Proof. This follows immediately from the trading rule for $d_{\mathcal{C}}$, modifying the decoration according to the decorated trading moves shown in Figure 4.2. \square

Lemma 4.3.9. *Let $\bar{\mathcal{C}}$ be a k -dimensional undecorated configuration with $k \in \{1, 2, 3\}$. Then Lemma 4.3.1 holds for $\bar{\mathcal{C}}$.*




Figure 4.7: The 11 purely active, connected 3-dimensional undecorated configurations up to rotation, trading and isotopy.

Proof. We explicitly verify the claim on every purely active, connected k -dimensional undecorated configuration modulo rotation and trading with some decoration \mathbf{t} . There is only one such 1-dimensional configuration. All 3 such 2-dimensional configurations are shown in Figure 4.6. All 10 such 3-dimensional configurations are shown in Figure 4.7. We leave these as a check for the interested reader. \square

Lemma 4.3.10. *Suppose there is a decoration \mathbf{t} so that $\mathcal{C} = (\bar{\mathcal{C}}, \mathbf{t})$ has type A_k , B_k , $C_{p,q}$, $D_{p,q}$, $E_{p,q}^{in}$ or $E_{p,q}^{out}$. Then Lemma 4.3.1 holds for $\bar{\mathcal{C}}$.*

Proof. Suppose \mathcal{C} has type A_k . There are $k + 1$ non-trivial contributions when $a = 1$ and $I' = \emptyset$ or $I' = \{i\}$ for $1 \leq i \leq k$. In the second case, I corresponds to a configuration of type A_{k-1} . Denote the ending circles of \mathcal{C} by y_1, \dots, y_k . The first case and the sum of the contributions in the second case both equal

$$\sum_i (1 + w(y_i)) y_i.$$

Next, suppose \mathcal{C} has type $C_{p,q}$ with $p + q = k$. There are $k + 1$ non-trivial contributions when $a = 1$ and $I' = \emptyset$ or $I' = \{i\}$ for $1 \leq i \leq k$. In the latter case, I corresponds to a configuration type $C_{p-1,q}$ (or $C_{p,q-1}$). Denote the ending circles of \mathcal{C} by y_1, \dots, y_{k-1} . The first case and the sum of the contributions in the second case both equal

$$\sum_i (1 + w(y_i)) y_i.$$

The case \mathcal{C} has types B_k or $D_{p,q}$ follow from the previous two cases and by the conjugate duality argument of Lemma 4.3.5.

Next, suppose \mathcal{C} has type $E_{p,q}^{in}$ with $p + q = k$. There are non-trivial contributions when $a = x_2 \cdots x_{p+1}$ or $a = x_2 \cdots x_{i-1} x_{i+1} \cdots x_{p+1}$. In the first case, there are $q + 1$ non-trivial contributions corresponding to $I' = \emptyset$ and $I' = \{i\}$ for $p + 1 \leq i \leq p + q$,




Figure 4.8: Local double arc configuration considered in Lemma 4.3.12.




Figure 4.9: Decorations of local double arc configuration considered in Lemma 4.3.12.

where γ_i is a split arc. Decompositions of each type sum to

$$\sum_{i=2}^{q+1} (1 + w(y_i)) y_1 y_i.$$

In the second case, there are two non-trivial contributions corresponding to $I' = \emptyset$ and $I' = \{i - 1\}$, where γ_{i-1} connects x_1 and x_i . Each contribution is equal to

$$(1 + w(x_i)) y_1.$$

The case \mathcal{C} has type E^{out} follows similarly and we leave it for the interested reader. □

Remark 4.3.11. We now assume there is no decoration \mathbf{t} such that \mathcal{C} has one of the six types A - D , E^{in} or E^{out} . By the preceding lemma, we can assume $d_{\mathcal{C}}(\mathbf{0}, \mathbf{1}) = 0$ in Lemma 4.3.1 and it remains to prove the lemma holds summing over I such that $1 \leq |I| \leq k - 1$.

Lemma 4.3.12. *Suppose $\bar{\mathcal{C}}$ contains the local configuration shown in Figure 4.8, where x and x' might be the same circle. If a is not divisible by x_1 or a is divisible by*

x, x' and x_1 , then Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a)$.

Proof. For the first case, we use the first decoration in Figure 4.9 if a is not divisible by x and the second if it is not divisible by x' . In either case, there are potentially two decompositions with non-trivial contributions corresponding to $I_1 = \{1\}$ and $I_2 = \{2\}$. Then $\mathcal{C}(I_1, \mathbf{1})$ and $\mathcal{C}(I_2, \mathbf{1})$ are of type A_{k-1} or $E_{p,q}^{\text{in}}$ with $p + q = k - 1$ and equivalent up to weights. We have $d_{\mathcal{C}(\mathbf{0}, I_1)}(a) = d_{\mathcal{C}(\mathbf{0}, I_2)}(a)$, so the contributions of the two decompositions cancel.

If a is divisible by x and x' , we use either of the first two decorations in Figure 4.9 which are now symmetric. We are in the same situation as before: there are two compositions and $\mathcal{C}(I_1, \mathbf{1})$ and $\mathcal{C}(I_2, \mathbf{1})$ are equivalent up to weights. Set $z = d_{\mathcal{C}(\mathbf{0}, I_1)}(a)$ so that $d_{\mathcal{C}(\mathbf{0}, I_2)}(a) = w(x_1)z$. If the image of z is non-zero, $\mathcal{C}(I_1, \mathbf{1})$ must have type B or D . However, then \mathcal{C} would type B or D , contrary to assumption.

For the second case, we use the fourth decoration in Figure 4.9. Then there are no non-trivial contributions with the given a . □

Lemma 4.3.13. *Suppose there is more than one active starting circle and that $a \in V_0(\mathcal{C})$ is the product of all of them. Then Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a)$.*

Proof. Let W denote the index set of join arcs in $\bar{\mathcal{C}}$ that meet x_1 . The set W is non-empty because $\bar{\mathcal{C}}$ is connected and has more than one active starting circle. Choose a decoration \mathbf{t} that orients all these arcs away from x_1 . Suppose $d_{\mathcal{C}(\mathbf{0}, I)}(a)$ is non-zero for some I . Then $\mathcal{C}(\mathbf{0}, I)$ is a 1-dimensional split configuration or has type B or D . The chosen decoration implies that W is disjoint from I . If the coefficient of $d_{\mathcal{C}(\mathbf{0}, I)}(a)$ at b is non-zero, b is a product of starting circles of $\mathcal{C}(I, \mathbf{1})$. Then $d_{\mathcal{C}(I, \mathbf{1})}(b) = 0$ by the chosen orientation. □

Lemma 4.3.14. *Suppose $\bar{\mathcal{C}}$ contains at least one of the local configurations M_1, \dots, M_9 shown in Figure 4.10. Then $\bar{\mathcal{C}}$ satisfies Theorem 4.3.1.*

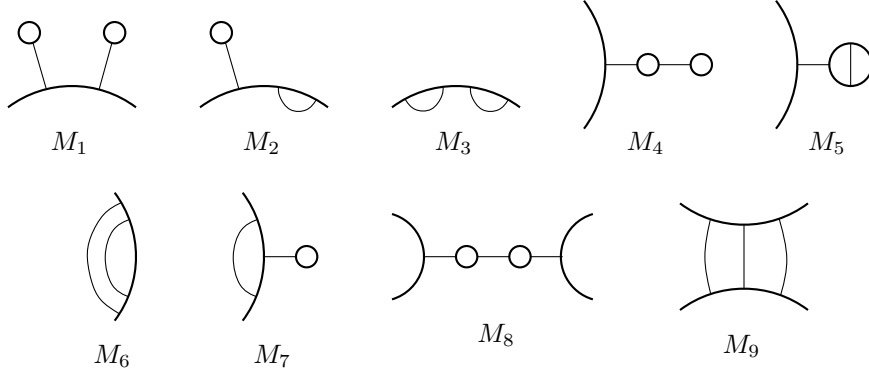


Figure 4.10: Nine local configurations M_1 - M_9 considered in Lemma 4.3.14.




Figure 4.11: Decorations for M_1 and M_8 considered in Lemma 4.3.14.

Proof. By trading and rotation M_2 and M_3 can be reduced to M_1 , M_5 can be reduced to M_4 and M_7 can be reduced to M_6 . By the conjugate duality argument of Lemma 4.3.5, M_6 can be reduced to M_4 and M_9 can be reduced to M_8 . It remains to establish the theorem for M_1, M_4 and M_8 .

For M_1 , we use the decoration on the left side of Figure 4.11. By Lemma 4.3.7, we only need to consider the case when a is divisible by x_1 and x_2 . Then by the filtration rule, no decompositions contribute.

For M_4 , if a is divisible by x_2 , then Lemma 4.3.7 applies. Otherwise, Lemma 4.3.12 applies.

For M_8 , we use the decoration in Figure 4.11. By Lemma 4.3.12, it is enough to consider the case when a is divisible by x_1x_2 . Then there are no non-trivial contributions by case analysis on I . □

Lemma 4.3.15. *Suppose there is a decoration \mathbf{t} of $\bar{\mathcal{C}}$ and I with $|I| = k - 1$ such that*




Figure 4.12: Maximal configuration when $\mathcal{C}(\mathbf{0}, I)$ has type $C_{p,q}$, case (ii), in Lemma 4.3.15.

the coefficient of

$$d_{\mathcal{C}(I, \mathbf{1})} d_{\mathcal{C}(\mathbf{0}, I)}(a)$$

at b is non-trivial. Then Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a, b)$.

Proof. The undecorated configuration $\bar{\mathcal{C}}$ is active and connected. By Lemma 4.3.9, it suffices to consider the case when $k \geq 4$. By assumption, there is a monomial $z \in V_1(\mathcal{C}(I, \mathbf{1}))$ such that the coefficient of $d_{\mathcal{C}(\mathbf{0}, I)}(a)$ at z and the coefficient of $d_{\mathcal{C}(I, \mathbf{1})}(z)$ at b is non-trivial. The active part $\mathcal{C}(\mathbf{0}, I)_0$ is a $(k-1)$ -dimensional configuration of type A_{k-1} , B_{k-1} , $C_{p,q}$ and $D_{p,q}$, $E_{p,q}^{\text{in}}$ or $E_{p,q}^{\text{out}}$ with $p+q = k-1$. The 1-dimensional configuration $\mathcal{C}(I, \mathbf{1})$ has a single arc δ . There are three cases for δ to consider:

1. δ joins two active ending circles of $\mathcal{C}(I, \mathbf{1})$,
2. δ joins an active ending circles of $\mathcal{C}(I, \mathbf{1})$ to passive circle w of $\mathcal{C}(I, \mathbf{1})$, or
3. δ splits an active ending circles of $\mathcal{C}(I, \mathbf{1})$.

Suppose $\mathcal{C}(\mathbf{0}, I)$ has type A_{k-1} . In case (i), \mathcal{C} and $\mathcal{C}(\mathbf{0}, I)$ have the same starting circles, so $b = z = a = 1$. Since b^* is a product of active starting circles, we can use Lemmas 4.3.5 and 4.3.13. In cases (ii) and (iii), δ only meets on one of the $k-1$ ending circles of $\mathcal{C}(\mathbf{0}, I)$, so $\bar{\mathcal{C}}$ contains M_9 and we use Lemma 4.3.14.

Suppose $\mathcal{C}(\mathbf{0}, I)$ has type B_{k-1} . Cases (i) and (iii) follow from Lemma 4.3.13. Case (ii) follows from Lemma 4.3.7.




Figure 4.13: Configuration when $\mathcal{C}(\mathbf{0}, I)$ has type $E_{p,q}$, case (i), in Lemma 4.3.15.




Figure 4.14: Maximal configurations when $\mathcal{C}(\mathbf{0}, I)$ has type $E_{p,q}$, case (ii), in Lemma 4.3.15.

Suppose $\mathcal{C}(\mathbf{0}, I)$ has type $C_{p,q}$ with $p + q = k - 1$. In case (i), we have $b = a = 1$. Since b^* is a product of ending circles and there are $p + q - 2 = k - 1 \geq 2$ ending circles, we use Lemmas 4.3.5 and 4.3.13.

In case (ii), if $\bar{\mathcal{C}}$ does not contain a M_9 , then the cases to check are (possibly sub-) configurations of the maximal configuration shown in Figure 4.12. By direct calculation Lemma 4.3.1 holds for $\bar{\mathcal{C}}$.

In case (iii), if δ is parallel an arc of $\mathcal{C}(\mathbf{0}, I)$ then $\bar{\mathcal{C}}$ has type $C_{p+1,q}$ and Lemma 4.3.10 holds. Otherwise, we can reduce to case (ii) by Lemma 4.3.8.

Suppose $\mathcal{C}(\mathbf{0}, I)$ has type $D_{p,q}$ with $p + q = k - 1$. Cases (i) and (ii) follow from Lemma 4.3.13 and case (ii) follows from Lemma 4.3.7.

Finally, suppose $\mathcal{C}(\mathbf{0}, I)$ has type $E_{p,q}$ with $p + q = k - 1$. In case (i), δ must split the central ending circle of $\mathcal{C}(\mathbf{0}, I)$ by connecting the central starting circle x_1 to a degree 1 starting circle $x_i, i \geq 2$. After trading and rotation, if $\bar{\mathcal{C}}$ does not contain a local configuration of type M_1 , there is one case to consider. It is shown in Figure 4.13.

In case (ii), δ must connect a degree 1 ending circles of $\mathcal{C}(\mathbf{0}, I)$ to either the another degree 1 ending circle or the central ending circle. In the second case, the edge of the central ending circle can correspond either to the starting circle or a degree 1 starting circle. After trading and rotation, the cases to check are (possibly sub-)




Figure 4.15: Configuration for case $Q = (D, D)$ in Lemma 4.3.17.




Figure 4.16: Configuration for case $Q = (A, E)$ in Lemma 4.3.17.

configurations of two maximal configurations show in Figure 4.14. The result follows by direct computation. □

Lemma 4.3.16. *Suppose there is a decoration \mathbf{t} of $\bar{\mathcal{C}}$ and I with $|I| = 1$ such that the coefficient of*

$$d_{\mathcal{C}(I,1)}d_{\mathcal{C}(\mathbf{0},I)}(a)$$

at b is non-trivial. Then Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a, b)$.

Proof. This follows from the previous lemma, Lemma 4.3.15, and the conjugate duality argument of Lemma 4.3.5. □

Lemma 4.3.17. *Suppose there is a decoration \mathbf{t} of $\bar{\mathcal{C}}$ and I such that $2 \leq |I| \leq k - 2$ such that the coefficient of*

$$d_{\mathcal{C}(I,1)}d_{\mathcal{C}(\mathbf{0},I)}(a)$$

at b is non-trivial. Then Lemma 4.3.1 holds for $(\bar{\mathcal{C}}, a, b)$.




Figure 4.17: Configuration for case $Q = (B, E)$ in Lemma 4.3.17.




Figure 4.18: Configuration for case $Q = (E, E)$ in Lemma 4.3.17.

Proof. By assumption, there is a monomial $z \in V_0(\mathcal{C}(I, \mathbf{1}))$ such that the coefficient of $d_{\mathcal{C}(\mathbf{0}, I)}(a)$ at z and the coefficient of $d_{\mathcal{C}(I, \mathbf{1})}(z)$ at b are both non-trivial. We will proceed by case analysis on the element Q in $\{A, B, C, D, E\}^2$ given by the types of $\mathcal{C}(\mathbf{0}, I)$ and $\mathcal{C}(I, \mathbf{1})$, respectively. A 2-dimensional configuration of type 13, which we think of as having types $C_{1,1}$ and $D_{1,1}$ will appear twice in the case analysis.

By the definition of $d_{\mathcal{C}}$, $Q \in \{A, C\} \times \{B, D\}$ and $Q \in \{B, D\} \times \{A, C\}$ are not possible unless \mathcal{C} is disconnected, contrary to assumption.

If $Q \in \{B, D\}^2$, then a must be a product of the starting circles. By Lemma 4.3.13, we need only consider the case when $\bar{\mathcal{C}}$ has a single starting circle. The only possibility is that $\mathcal{C}(\mathbf{0}, I)$ and $\mathcal{C}(I, \mathbf{1})$ both have type $D_{1,1}$. There are five cases to consider, shown in Figure 4.15.

The case $Q \in \{A, C\}^2$ follows from the previous case $Q \in \{B, D\}^2$ and Lemma 4.3.5 (conjugate duality).

The remaining cases contain a configuration of type E . We reduce each case to a finite (possibly empty) list of configurations to verify.

$\mathbf{Q} = (\mathbf{A}, \mathbf{E})$. The s -dimensional configuration $\mathcal{C}(\mathbf{0}, I)$ has s ending circles y_1, \dots, y_s , none of which divide z . Therefore, one of the ending circles, say y_1 , is the central starting circle for $\mathcal{C}(I, \mathbf{1})$ and the others are passive starting circles. If $\bar{\mathcal{C}}$ doesn't contain M_8 , then $\mathcal{C}(\mathbf{0}, I)$ is 2-dimensional. Assuming $\bar{\mathcal{C}}$ does not contain M_1, M_2 and M_3 , up to trading and rotation, one configurations remains, see Figure 4.16.

$\mathbf{Q} = (\mathbf{B}, \mathbf{E})$. The configuration $\mathcal{C}(\mathbf{0}, I)$ has two ending circles y_1 and y_2 , both of which divide z . Therefore, the central starting circle y_3 of $\mathcal{C}(I, \mathbf{1})$ is a passive circle of $\mathcal{C}(\mathbf{0}, I)$,

and there is a single arc between y_3 and one or both of y_1 and y_2 . That means $\mathcal{C}(\mathbf{0}, I)$ is 2-dimensional. Since $k \geq 4$ and up to trading and rotation $\bar{\mathcal{C}}$ does not contain M_4 , y_3 must connect to both y_1 and y_2 . Two possible configurations remain, see Figure 4.17.

$\mathbf{Q} = (\mathbf{E}, \mathbf{C})$. The active starting circle of $\mathcal{C}(I, \mathbf{1})$ agrees with one of the degree 1 ending circles of $\mathcal{C}(\mathbf{0}, I)$. It follows that if the dimension of $\mathcal{C}(\mathbf{0}, I)$ is greater than 2, then $\bar{\mathcal{C}}$ contains M_1 , M_2 or M_3 after rotations. Furthermore, if $\mathcal{C}(\mathbf{0}, I)$ is 2-dimensional, then $\bar{\mathcal{C}}$ contains an M_6 or M_7 after rotations.

$\mathbf{Q} = (\mathbf{D}, \mathbf{E})$. The single active ending circle y_1 of $\mathcal{C}(\mathbf{0}, I)$ must be a degree 1 active starting circle for $\mathcal{C}(I, \mathbf{1})$. If $\mathcal{C}(I, \mathbf{1})$ has degree three or more, then $\bar{\mathcal{C}}$ contains M_1 after rotation and trading. If $\mathcal{C}(I, \mathbf{1})$ has degree 2, then $\bar{\mathcal{C}}$ contains M_4 after rotation and trading.

$\mathbf{Q} = (\mathbf{E}, \mathbf{E})$. There are three cases to consider. First, the central starting circle of $\mathcal{C}(I, \mathbf{1})$ is a passive circle of $\mathcal{C}(\mathbf{0}, I)$. Then the central ending circle of $\mathcal{C}(\mathbf{0}, I)$ is a degree 1 starting circle for $\mathcal{C}(I, \mathbf{1})$, the other starting circles are passive circles for $\mathcal{C}(I, \mathbf{1})$ and $\bar{\mathcal{C}}$ contains a configuration of type M_1 or M_4 up to rotation and trading.

Second, the central starting circle of $\mathcal{C}(I, \mathbf{1})$ is a dual degree 1 active ending circle for $\mathcal{C}(\mathbf{0}, I)$ and the central ending circle for $\mathcal{C}(\mathbf{0}, I)$ is a passive circle for $\mathcal{C}(I, \mathbf{1})$. Then, again, \mathcal{C} contains a M_1 or M_4 up to trading and rotation.

Finally, the central starting circle of $\mathcal{C}(I, \mathbf{1})$ is a dual degree 1 active ending circle for $\mathcal{C}(\mathbf{0}, I)$ and the central ending circle for $\mathcal{C}(\mathbf{0}, I)$ is degree 1 starting circle for $\mathcal{C}(I, \mathbf{1})$. Then $\bar{\mathcal{C}}$ is of one of the two configurations shown in Figure 4.18, or a sub-configuration where some arcs meeting degree 1 degree one circles (local “stalk” configurations) have been removed.

The other five cases follow by the conjugate duality argument of Lemma 4.3.5. \square

Lemma 4.3.1 now follows directly from Lemmas 4.3.15, 4.3.16 and 4.3.17. \square

4.4 Total homology

In this section, we study homology of the reduced complex $(\tilde{C}_{\mathcal{D}}, d)$ and, when the weights are chosen generically, identify it with $H_{\text{BOS}}(L)$.

Proof of Theorem 1.1.5. The first part of the theorem follows directly from the definition of $d(\mathbf{t})$. When the twisting is zero, it reduces directly to the definition of $d_{\text{gss}}(\mathbf{t})$.

We now consider the second part. Recall, we are working over a field F and the weights $w(e)$ for edges e of \mathcal{D} not meeting p are chosen to be algebraically independent.

Fix a resolution I . If \mathcal{C}_I has more than one circle, then the $V(I)$ summand of the complex is a tensor product of the form

$$\bigotimes_i (0 \rightarrow F \cdot 1_{x_i} \rightarrow F \cdot x_i \rightarrow 0),$$

where the center arrow is multiplication by $1 + w(x_i)$. Since the weights $w(e)$ are algebraically independent, $1 + w(x_i)$ is non-trivial. This complex is well-known to be acyclic. Our complex is “totally twisted” in the sense of Roberts [58]. When we simplify the complex using the Cancellation Lemma [4, Lemma 4.1] by cancelling terms internal to the vertices of the cube of resolutions, the only generators that remain in the simplified complex $C'_{\mathcal{D}}$ correspond to connected resolutions. Our goal is now to determine the differential d' after cancellation.

The terms of $d(\mathbf{t})$ have (h, q) -grading $(k, 2k - 2)$. This is not changed by cancellation. For connected resolutions \mathcal{C}_I , the q grading of the single generator x is determined by the h -grading by

$$q(x) = h(x) + n_+ - n_-.$$

In particular, after cancellation, terms of the simplified differential must shift the




Figure 4.19: Cases to consider when computing simplified differential d' .





Figure 4.20: Detail of case 2 in computing simplified differential d' .

quantum and homological gradings by equal amounts. This means the only terms that remain must have grading $(2, 2)$.

Let (I, I') be a square of the cube of resolutions where \mathcal{C}_I and $\mathcal{C}_{I'}$ are connected resolutions. Let J and K be the other two vertices of the square. Let x and y be generators for $V(I)$ and $V(I')$, respectively. The terms in d' from x to y come from two sources: the original terms of $d(\mathbf{t})$ having grading $(2, 2)$ and the terms introduced by cancellation.

Up to isotopy, there are four cases to consider depending on the orientation of the arcs $\mathcal{C}(I, I')$ and the placement of p , as shown in Figure 4.19. The second case is

shown in Figure 4.20. The complex before cancellation is:



After cancellation, we get

$$\begin{aligned}
 d'(x) &= \left(1 + \frac{w(r)}{1+w(r)} + \frac{1}{1+w(t)} \right) y \\
 &= \left(\frac{1}{1+w(s)} + \frac{1}{1+w(t)} \right) y
 \end{aligned}$$

since we are working modulo 2. This agrees with (2.1) since the configuration $\mathcal{C}(I, I')$ is negative. We leave the other three cases as a check for the interested reader. \square

Chapter 5

The Link Splitting Spectral Sequence

In this chapter, we define the the link splitting spectral sequence and give the applications mentioned in the introduction: we give a bound on the link splitting number and prove that Khovanov homology detects the unlink. In Section 5.1, we define the filtered chain complex inducing the link splitting spectral sequence and establish some basic properties. In Section 5.2 we prove that the link splitting spectral sequence is a link invariant. In Section 5.3, we give some background on spectral sequences which we apply in the following section, Section 5.4, where we give the link splitting bound. In Section 5.5, we prove that Khovanov homology detects the unlink. Finally, in Section 5.6, we give some sample computations of the link splitting spectral sequence and the link splitting bound. The material in this chapter is joint work with Joshua Batson.

Throughout this chapter, $E_k = E_k^{\text{lsss}}$ will denote the link splitting spectral sequence, $C(\mathcal{D}, w) = C^{\text{lsss}}(\mathcal{D}, w)$ will denote the filtered chain complex inducing the spectral sequence, $d = d^{\text{lsss}}$ the differential, etc.

5.1 Construction

Khovanov's construction begins with a diagram \mathcal{D} for a link L . He builds a cube of resolutions for \mathcal{D} and applies a $(1+1)$ -dimensional TQFT \mathcal{A} to produce a cube-graded complex. A sprinkling of signs yields a chain complex $(C(\mathcal{D}), d_0)$, with homology $Kh(L)$. We will give another differential d on the same chain complex, but first we must set some notation.

5.1.1 Chain complex

We begin by describing our construction in the case of a two-component link L with coefficients in \mathbb{F}_2 . The link splitting spectral sequence will be g -graded, so we denote the Khovanov differential by d_0 . Khovanov's construction assigns a bigraded chain complex $(C(\mathcal{D}), d_0)$ to a planar diagram \mathcal{D} for L . We will give an endomorphism d_1 of $C(\mathcal{D})$ with the following properties:

1. $d := d_0 + d_1$ is a differential, which increases the ℓ -grading by 1.
2. d_1 lowers the g -grading by 1, making $(C(\mathcal{D}), d)$ a g -filtered complex.
3. If i is a crossing in \mathcal{D} involving strands from different components of L (a *mixed* crossing), and \mathcal{D}' is the diagram for a link L' produced by changing over-strand to under-strand at i , then $(C(\mathcal{D}), d)$ and $(C(\mathcal{D}'), d')$ are isomorphic chain complexes (with different g -filtrations).

The new endomorphism is

$$d_1 = \sum_{\text{mixed edges } (I, J)} \mathcal{A}(J, I), \quad (5.1)$$

where an edge in the cube of resolutions is mixed if the I and J differ at a mixed crossing, and (J, I) denotes the cobordism (I, J) viewed backwards as a cobordism

from J to I . The total differential

$$d = \sum_{\text{non-mixed edges } (I, J)} \mathcal{A}(I, J) + \sum_{\text{mixed edges } (I, J)} \mathcal{A}(I, J) + \mathcal{A}(J, I)$$

is manifestly unchanged if we swap a mixed crossing. The square d^2 can have a component from $V(I)$ to $V(J)$ only when I and J differ at 2 crossings or when $I = J$. The former vanish because they come in commuting squares (all maps are induced by cobordisms, and those commute due to the TQFT). The latter will vanish too, essentially because each circle in a complete resolution must have an even number of mixed crossings. We will establish that $d^2 = 0$ more carefully in Proposition 5.1.2, where we also handle multi-component links and other rings of coefficients.

To define the endomorphism d_1 when there are more than two components, or over bigger rings, we need some additional data. First, we must weight each component by an element of the coefficient ring R : component c has weight w_c . Then we must construct a sign assignment so that d^2 will be zero, not just even. As usual, different choices of sign assignment will produce isomorphic complexes.

We now define a sign assignment. The shadow of the diagram \mathcal{D} in the plane gives a CW decomposition X of S^2 : the 0-cells are the double points of the diagram, the 1-cells are the $2n$ edges between the crossings (oriented by the orientation of the link), and the 2-cells are the remaining regions (with the natural orientation induced from S^2). For each 1-cell e , let $e(0)$ denote the initial vertex and $e(1)$ denote the final vertex.

Let

$$h(e, i) = \begin{cases} 1 & e \text{ is an upper strand at } e(i) \\ -1 & e \text{ is a lower strand at } e(i), \end{cases}$$

where $i \in \{0, 1\}$. There is a natural 1-cochain $\beta : X^1 \rightarrow \mathbb{Z}/2$, where $\mathbb{Z}/2 = \{\pm 1\}$ is written multiplicatively, given by

$$\beta(e) = \begin{cases} -1 & h(e, 0) = h(e, 1) \\ 1 & \text{otherwise.} \end{cases}$$

A *sign assignment* is a 0-cochain $s : X^0 \rightarrow \{\pm 1\}$ such that

$$s(e(0))s(e(1)) = \beta(e), \tag{5.2}$$

for all 1-cells e . This is equivalent to $\delta s = \beta$. Note that if \mathcal{D} is an alternating diagram, then $s \equiv 1$ is a legal sign assignment. In the definition of d_1 , we will use s to sign the weight of the top strand at each crossing; the bottom strand will get the opposite sign. The condition $\delta s = \beta$ means that at adjacent crossings, connected by a strand in component c of the link, the weight w_c will appear with opposite signs in the contributions from each.

We now define the endomorphism d_1 of $C(\mathcal{D})$ as

$$d_1 = \sum_{\text{edges } (I, J)} (-1)^{n(I, J)} s(i) (w_{\text{over}}^i - w_{\text{under}}^i) \mathcal{A}(J, I),$$

where I and J differ at the i^{th} crossing, and w_{over}^i and w_{under}^i are the weights of the over- and under-strands at the i^{th} crossing. Only the differences of weights appear, so shifting all the weights by some $r \in R$, leaves the complex invariant.

In particular, the complex for a two-component link is determined by the choice of a single value $w_1 - w_2 \in R$. If that difference is $1 \in \mathbb{F}_2$, then this definition of d_1 reduces to (5.1).

The complex $(C(\mathcal{D}), d = d_0 + d_1)$ now satisfies properties (1) and (2) from the beginning of this section. Both d_0 and d_1 increase the (internal) l -grading by 1. The




Figure 5.1

differential d_0 preserves the g grading and d_1 decreases the g grading by 1. So we have a g -filtration on $(C(\mathcal{D}), d)$ given by

$$\mathcal{F}^p C(\mathcal{D}) := \{x \mid x \in C(\mathcal{D}), g(x) \leq p\}.$$

Moreover, the spectral sequence associated to this filtration has E_1 page given by $H^*(C(\mathcal{D}), d_0) \cong Kh(L)$.

We now show it is always possible to choose a sign assignment.

Proposition 5.1.1. *Let \mathcal{D} be a connected diagram. There are precisely two sign assignments s_1 and s_2 for \mathcal{D} , and $s_1 = -s_2$.*

Proof. By (5.2), a choice of sign at one crossing determines the sign assignment for a connected diagram, if one exists. Existence is a simple cohomological argument. Since a sign assignment is just a cochain $s \in C^0(S^2)$ with $\delta s = \beta$, such an s exists if and only if $\beta \in C^1(S^2)$ is exact, and is unique up to multiplication by an element of $H^0(S^2) = \{\pm 1\}$. Since $H^1(S^2) = 0$, β is exact if and only if it is closed.

We now show that β is closed. Let f be a 2-cell with the incident 0- and 1-cells numbered counterclockwise v_1, \dots, v_n and e_1, \dots, e_n , respectively; see Figure 5.1. Each vertex v_i is incident to two edges, e_{i-1} and e_i (where we set $e_0 = e_n$). For one of those edges, v_i is an over-crossing, and for the other v_i is an under-crossing. More




Figure 5.2: We choose marked points p_i and q_i on the understrands at each crossing i (left) and a marked point p_e on each edge e (right).

formally, if $v_i = e_{i-1}(a_i) = e_i(b_i)$ for some $a_i, b_i \in \{0, 1\}$, then

$$h(e_{i-1}, a_i)h(e_i, b_i) = -1.$$

By definition, $\beta(e_i) = -h(e_i, 0)h(e_i, 1)$. We then have

$$\begin{aligned} (\delta\beta)(f) &= \prod_{i=1}^n \beta(e_i) \\ &= \prod_{i=1}^n -h(e_i, 0)h(e_i, 1) \\ &= \prod_{i=1}^n -h(e_{i-1}, a_i)h(e_i, b_i) \\ &= 1. \end{aligned} \quad \square$$

For a split diagram, sign assignments can be chosen on each connected component independently.

Property (3) does not hold on the nose. If \mathcal{D} and \mathcal{D}' are related by changing a crossing, then the associated differentials d and d' are not identical—they differ by elements of R . We will investigate this in Subsection 5.1.3 after verifying that our new differential squares to zero and showing the filtered chain homotopy type of the $(C(\mathcal{D}), d)$ does not depend on the choice of sign assignment.

Proposition 5.1.2. *We have that $d^2 = 0$.*

Proof. Fix a resolution I and let $x \in V(I)$. The terms of $d^2(x)$ lie in $V(K)$ where K differs from I in exactly two positions or $K = I$ itself. We study these two cases.

Case 1. Let K be a resolution that differs from I in exactly two positions i, j with $i < j$. Let J differ from I at i , and J' differ from I at j . Then I, J, J' and K are the four vertices of a face of the hypercube of resolutions. By functoriality of \mathcal{A} , we have that $\mathcal{A}(J, K)\mathcal{A}(I, J) = \mathcal{A}(J', K)\mathcal{A}(I, J')$. The endomorphism d_1 uses the usual Khovanov sign assignments, so the two paths around the face have different signs. Namely, we have that $n(I, J) = n(J', K)$ and $n(J, K) = -n(I, J')$. The weights on the cobordism maps in d_0 and d_1 depend only on which crossing is changed, not the edge of the cube. Denote the weights involved by $c(k)$, where

$$c(k) = \begin{cases} 1 & I(k) = 0 \\ s(k)(w_{\text{over}}^k - w_{\text{under}}^k) & I(k) = 1. \end{cases}$$

The terms of $d^2(x)$ in $V(K)$ are

$$\begin{aligned} & c(i)c(j)((-1)^{n(I,J)+n(J,K)}\mathcal{A}(J, K)\mathcal{A}(I, J)(x) \\ & \quad + (-1)^{n(I,J')+n(J',K)}\mathcal{A}(J', K)\mathcal{A}(I, J')(x) \\ & = c(i)c(j)(-1)^{n(I,J)+n(J,K)}(\mathcal{A}(J, K)\mathcal{A}(I, J)(x) - \mathcal{A}(J', K)\mathcal{A}(I, J')(x)) \\ & = 0. \end{aligned}$$

Case 2. The terms of $d^2(x)$ in $V(I)$ are

$$\sum_{i=1}^n s(i)(w_{\text{over}} - w_{\text{under}})\mathcal{A}(J_i, I)\mathcal{A}(I, J_i)(x),$$

where J_i is the resolution which differs from I solely at the position i . We choose marked points on the under-strands at each crossing and on each edge, see Figure 5.2. Straightforward computation shows that $\mathcal{A}(J_i, I)\mathcal{A}(I, J_i) = X_{p_i} + X_{q_i}$. We can rewrite

the above sum as

$$\begin{aligned}
& \sum_{i=1}^n s(i)(w_{\text{over}}^i - w_{\text{under}}^i)(X_{p_i} + X_{q_i}) \\
&= \sum_{e \in X^1} s(e(0))h(e, 0)w_e X_{p_e} + s(e(1))h(e, 1)w_e X_{p_e} \\
&= \sum_{e \in X^1} (s(e(0))h(e, 0) + s(e(1))h(e, 1))w_e X_{p_e} \\
&= 0,
\end{aligned}$$

where w_e denotes the weight of the component containing the edge e , the first equality follows from indexing the sum by edges, and the second equality follows from the definition of a sign assignment. \square

5.1.2 Change of sign assignment

While finding a sign assignment s is crucial for defining the complex over rings where $2 \neq 0$, different choices produce isomorphic complexes. Indeed, consider a connected diagram D , weight w , and sign assignment s producing the complex $(C(D), d = d_0 + d_1)$. Then taking the other sign assignment, $-s$, yields the differential $d' = d_0 - d_1$ on the same group of chains $C(D)$. Since d_0 fixes g -grading, and d_1 lowers it by 1, the endomorphism

$$\begin{aligned}
\phi : C(D) &\rightarrow C(D) \\
x &\mapsto (-1)^{g(x)}x
\end{aligned}$$

has the property that $d\phi = \phi d'$. That is, ϕ is an invertible chain map between $(C(D), d)$ and $(C(D), d')$.

Next, consider the case when D is possibly split and s and s' are two sign assignments. Then, since \mathcal{A} is a monoidal functor, the complexes $(C(D), d)$ and $(C(D), d)$




Figure 5.3: The crossing change move C.

each decomposes into a tensor product of complexes indexed over the components of D . The above analysis gives a chain equivalence ϕ for each component, and their tensor product gives an invertible chain map between $(C(D), d)$ and $(C(D), d')$.

Henceforth, we will often suppress the choice of a sign assignment, writing $C(D, w)$ to indicate one of the two possible complexes.

5.1.3 Total homology

We now show that changing a crossing doesn't affect the total homology of $(C(D), d)$, so long as the relevant weight $w_{\text{over}} - w_{\text{under}}$ is invertible. Of course, changing the crossing does *not* preserve the g -filtration on C .

Proposition 5.1.3. *Let D and D' be diagrams for links L and L' related by changing a crossing i between components c and d . Let w be a weighting for L , and write w' for the induced weighting on L' . Then if $w_c - w_d$ is invertible in R , the complexes $C(D, w)$ and $C(D', w')$ are isomorphic as relatively ℓ -graded chain complexes.*

Proof. Let s be a sign assignment for D . A sign assignment s' for D' is given by $s'(j) = s(j)$ for $j \neq i$ and $s'(i) = -s(i)$. Let (C, d) be the complex $C(D, w, s)$, and let (C', d') be the complex $C(D', w, s')$. Let C_0 be the summand of C consisting of complete resolutions which include the 0 resolution at crossing i , and let C_1, C'_0 and C'_1 be defined analogously. Note that C_0 and C'_1 are identical as relatively ℓ -graded complexes; similarly for C_1 and C'_0 . (The writhes of the diagrams differ by 2, which will contribute a global shift between their ℓ -gradings.)

The crossing change exchanges over-strand for under-strand, so $(w_{\text{over}}^i - w_{\text{under}}^i) = -(w_{\text{over}}^i - w_{\text{under}}^i)$. This means that

$$s(i)(w_{\text{over}}^i - w_{\text{under}}^i) = s'(i)(w_{\text{over}}^i - w_{\text{under}}^i).$$

Before giving the chain map $f : C \rightarrow C'$, we must first introduce some notation. Let I be a resolution of D . We write I' to denote the same element of $\{0, 1\}^n$ interpreted as a resolution of D' . We write I_i for the resolution of D that differs with I solely at crossing i . Note that I and I_i are canonically isomorphic resolutions. Let J denote a resolution of D that differs from I at some crossing $j \neq i$. Finally, let

$$a(I, i) = \#\{I(k) = 1 \mid i < k \leq n\}$$

be the number of one digits in I above i .

We define the map $f : C \rightarrow C'$ as follows.

$$x \mapsto \begin{cases} (-1)^{a(I, i)} x \in C'_1 \text{ if } x \in V(I) \subset C_0 \\ (-1)^{a(I_i, i)} s(i)(w_{\text{over}}^i - w_{\text{under}}^i)x \in C'_0 \text{ if } x \in V(J) \subset C_1, \end{cases}$$

To verify f is a chain map, we use two easily verifiable facts about the signs:

$$(-1)^{a(I, i)} = (-1)^{a(I_i, i)}$$

and

$$(-1)^{n(I, J)} (-1)^{a(J, i)} = (-1)^{n(I'_i, J'_i)} (-1)^{a(I, i)}.$$

Consider $x \in V(I) \subset C_0$. The image of x under fd or $d'f$ has components in $V(I')$ and $V(J'_i)$, for the resolutions J differing from I at one crossing.

First, consider the $V(I')$ -component of the image. We have

$$\begin{aligned}
fd(x)|_{V(I')} &= f((-1)^{n(I, I_i)} \mathcal{A}(I, I_i)(x)) \\
&= (-1)^{a(I, i)} (-1)^{n(I, I_i)} s(i) (w_{\text{over}}^i - w_{\text{under}}^i) \mathcal{A}(I, I_i)(x) \\
&= (-1)^{a(I, i)} (-1)^{n(I'_i, I')} s'(i) (w_{\text{over}}^{I'_i} - w_{\text{under}}^{I'_i}) \mathcal{A}(I'_i, I')(x) \\
&= d'((-1)^{a(I, i)} x)|_{V(I')} \\
&= d'f(x)|_{V(I')}.
\end{aligned}$$

Next, consider the image in $V(J'_i)$ for some J which differs from I at crossing j .

Let

$$c(j) = \begin{cases} 1 & I(j) = 0 \\ s(i)(w_{\text{over}}^j - w_{\text{under}}^j) & I(j) = 1 \end{cases}$$

denote the coefficient of $\mathcal{A}(I, J)$ in d . It is the same as the coefficient of $\mathcal{A}(I'_i, J'_i)$ in d' . We have

$$\begin{aligned}
fd(x)|_{V(J'_i)} &= f((-1)^{n(I, J)} c(j) \mathcal{A}(I, J)(x)) \\
&= (-1)^{a(J, i)} (-1)^{n(I, J)} c(j) \mathcal{A}(I, J) \\
&= (-1)^{a(I, i)} (-1)^{n(I'_i, J'_i)} c(j) \mathcal{A}(I'_i, J'_i) \\
&= d'((-1)^{a(I, i)} x)|_{V(J'_i)} \\
&= d'f(x)|_{V(J'_i)}.
\end{aligned}$$

A similar analysis shows that $fd(x) = d'f(x)$ for $x \in C_1$.

Let $f' : C' \rightarrow C$ be the chain map produced by reversing the roles of D and D' .

The composition

$$ff' = f'f = s(i)(w_{\text{over}}^i - w_{\text{under}}^i)$$

is an isomorphism if $(w_{\text{over}}^i - w_{\text{under}}^i)$ is invertible for all i . In that case, f is an

isomorphism. □

5.2 Reidemeister invariance

The proof that the filtered chain homotopy type of $C(D, w, s)$ is invariant under the Reidemeister moves parallels the standard proof that the Khovanov chain complex is invariant. We divide the complex into the summands corresponding to the 2, 4, or 8 ways of resolving the crossings involved in the move, and cancel isomorphic summands along components of the differential. This is complicated slightly by the d_1 terms which prevent the natural summands of C from being subcomplexes; the post-cancellation differential is not merely a restriction of the original one. The new differential is provided by the following standard cancellation lemma.

Lemma 5.2.1. *Let (C, d) be a chain complex. Suppose that C , viewed as an R -module, splits as a direct sum $V \oplus W \oplus C'$. Let d_{WV} denote the component of d mapping from V to W , and similarly for other components. If d_{WV} is an isomorphism, then (C, d) is chain homotopy equivalent to (C', d') with*

$$d' = d_{C'C'} - d_{C'V}d_{WV}^{-1}d_{WC'}.$$

Proof. Let $f : C' \rightarrow C$, $g : C \rightarrow C'$, and $h : C \rightarrow C$ be defined by

$$f = \iota_{C'} - d_{WV}^{-1}d_{WC'}, \quad g = \pi_{C'} - d_{C'V}d_{WV}^{-1}, \quad \text{and} \quad h = d_{WV}^{-1},$$

where ι and π denote inclusion and projection with respect to the direct sum decomposition of C . The map f is an isomorphism onto its image, since the second term in f merely adds a V -component. The image of f turns out to be a subcomplex, and the new differential d' is merely the pullback of d along f .

We claim that f and g are mutually inverse chain homotopy equivalences between

(C, d) and (C', d') . Specifically, the following four equations hold:

$$fd' = df \quad gd = d'g \quad \mathbb{I}_{C'} = gf \quad \mathbb{I}_C = fg + hd + dh$$

Verifying these is a routine exercise in applying the identities contained in the equation $d^2 = 0$, such as

$$d_{WV}d_{VV} + d_{WC'}d_{C'V} + d_{WW}d_{WV} = 0. \quad \square$$

If the complex (C, d) has a filtration induced by a grading g and the cancelled map, d_{WV} above, preserves g -degree, then d' will respect the induced filtration on C' and the maps f and g will be filtered chain homotopy equivalences. This will be our situation in each of the Reidemeister moves below.

Proposition 5.2.2. *Let D and D' be two diagrams for a link L related by a Reidemeister move of type I, II, or III. Fix an R -weighting w for L and a sign assignment s for the diagram D . Then there exists a sign assignment s' for the diagram D' which agrees with the sign assignment for D at all crossings uninvolved in the Reidemeister move, and the complexes $C(D, w, s)$ and $C(D', w, s')$ are chain homotopy equivalent as ℓ -graded, q -filtered complexes.*

In Section 5.1, we saw that different sign assignments produce isomorphic complexes. Since any two diagrams for a link are related by a sequence of Reidemeister moves, this proposition implies that the ℓ -graded q -filtered chain homotopy type of the complex $C(D, w, s)$ is also independent of the choice of planar diagram, and hence an invariant of the R -weighted link (L, w) . This establishes that the associated spectral sequence, called $E_k(L, w)$ in Theorem 1.1.7, is an invariant of (L, w) .

Proof. The proof for each of the three Reidemeister moves is similar. We first decompose the complex into summands sitting over each of the 2^k different resolutions of the crossings implicated in the k -th move. One of these resolutions contains an isolated



Figure 5.4: Left is the first Reidemeister move R1. Right is chain complex for the diagram D , split into two summands corresponding to the two resolutions of the pictured crossing.

circle, and we split the complex over that resolution further according to whether or not the monomial contains that circle. We then identify two summands V and W for which d_{WV} is a q -grading-preserving isomorphism, and apply the cancellation lemma.

R1 Consider two diagrams D and D' for a link L in Figure 5.4. Let s be a sign assignment for D . It can be verified easily that the restriction of s to the vertices of the diagram for D' yields a valid sign assignment s' .

Let (C, d) be the complex $C(D, w, s)$, and let (C', d') be the complex $C(D', w, s')$. Let C_0 be the summand of C corresponding to complete resolutions which include the 0-resolution at the pictured crossing, and let C_1 be the summand of C corresponding to complete resolutions which include the 1-resolution at the pictured crossing. Let C_0^- and C_0^+ be the summands of C_0 spanned by monomials divisible and not divisible, respectively, by the variable x_r corresponding to the pictured circle.

Since the component of d mapping from C_0^+ to C_1 is just merging in the 1 on the pictured circle, it is an isomorphism. Hence we may apply the cancellation lemma with $V = C_0^+$ and $W = C_1$. Since C_0^+ and C_0^- have the same resolution at the pictured crossing, there is no component of d mapping from one to the other. Hence the new complex is just C_0^- with the restriction of the original differential. (This cancellation preserves the filtration, since the cancelled part of the differential is a component of the ordinary Khovanov d_0 , which preserves q - and g - degree.) Since the extra circle never interacts with the remainder of the diagram for L , this complex (C_0^-, d) is isomorphic to the post-move complex (C', d) . That isomorphism also respects the gradings, as can be verified from $n_+(D) = n_+(D') + 1$, $n_-(D) = n_-(D')$, and the




Figure 5.5: Left is the second Reidemeister move $R2$. Right is chain complex for the diagram D , split into four summands corresponding to the resolutions of the pictured crossings.

definitions of ℓ and h .

R2 Consider two diagrams D and D' for a link L in Figure 5.5. Let s be a sign assignment for D . It can be verified easily that the restriction of s to the vertices of the diagram for D' yields a valid sign assignment s' .

Let D_{ij} with $i, j \in \{0, 1\}$ denote the diagrams obtained by resolving the crossings involved in the Reidemeister move in D . Let $C_{ij} = C(D_{ij}, w, s)$. Let C_{01}^- and C_{01}^+ be the summands of C_{01} spanned by generators divisible and not divisible, respectively, by the variable x_r corresponding to the pictured circle. The four summands C_{00}, C_{11}, C_{01}^+ and C_{01}^- are all naturally isomorphic, and the summand C_{10} is isomorphic to the post-move complex $C' = C(D', w, s')$.

We will apply the cancellation lemma with $V = C_{00} \oplus C_{01}^+$ and $W = C_{01}^- \oplus C_{11}$. The component of d from V to W is just the original Khovanov differential d_0 , and it is block diagonal: C_{01}^+ maps to C_{11} isomorphically (merging in a 1) and C_{00} maps to C_{01}^- isomorphically (splitting of an x).

The cancelled complex is just C_{10} , with differential

$$d_{C_{10}C_{10}} - d_{C_{10}V}d_{WV}^{-1}d_{WC_{10}}.$$




Figure 5.6: Left is the third Reidemeister move R3. Right is chain complex for the diagram D , split into eight summands corresponding to the resolutions of the pictured crossings.

But $d_{WC_{10}}$ lands on C_{11} , which is carried to C_{01}^+ by d_{WV}^{-1} , and d has no component from C_{01}^+ to C_{10} . Hence the new differential is just the restriction of the old, and we have

$$(C, d) \cong (C_{10}, d|_{C_{10}}) \cong (C', d').$$

Again, the cancelled pieces of the differential come from Khovanov's d_0 , which preserves g and q , so the first isomorphism preserves the filtration. The second isomorphism also preserves the bigrading, as can be verified from $n_{\pm}(D) = n_{\pm}(D') + 1$.

R3 Reidemeister 3 is more complicated, and we must keep track of the signs in Khovanov's cube, the sign assignment s , and the weights.

Consider the diagrams D and D' in Figure 5.6. Label the strands i , j , and k , from left to right along the top of D . Denote by w_i, w_j , and w_k the weights of their components. Order the crossings up the page 1, 2, and 3. Using Khovanov's sign assignment, the edges in the cube of resolutions for D labeled -1 in the figure have

a negative sign in the differential: $(-1)^{n(I,J)} = -1$. (100 \leftrightarrow 110, 100 \leftrightarrow 101, 010 \leftrightarrow 011, 101 \leftrightarrow 111.)

Choose a sign assignment s for D such that

$$s(1) = s(3) = 1 \text{ and } s(2) = -1.$$

A choice of sign at one crossing determines the sign assignment on that component of the diagram by (5.2). Take the sign assignment s' for D' which agrees with s on the crossings not involved in the Reidemeister move. Again, (5.2) implies

$$s'(1) = s'(3) = -1 \text{ and } s'(2) = 1.$$

Let $(C, d) := C(D, w, s)$ and $(C', d') := C(D', w, s')$. The weights $c(j) = s(j)(w_{\text{over}} - w_{\text{under}})$ of the reverse edge maps in d_1 evaluate to

$$c(1) = w_j - w_k$$

$$c(2) = w_k - w_i$$

$$c(3) = w_i - w_j,$$

at the three pictured crossings, and the weights $c'(j)$ in d'_1 are

$$c'(1) = w_j - w_i$$

$$c'(2) = w_i - w_k$$

$$c'(3) = w_k - w_j.$$

First, we will simplify the complex (C, d) . As in the previous parts, let C_{010}^- and C_{010}^+ be the summands of C_{010} spanned by monomials divisible and not divisible, respectively, by the variable x_r corresponding to the pictured circle.

We apply the cancellation lemma with

$$V = C_{000} \oplus C_{010}^+ \quad W = C_{010}^- \oplus C_{011}.$$

The component of d from V to W is just the Khovanov differential d_0 , and it is block diagonal: C_{000} maps to C_{010}^- isomorphically (splitting off an x) and C_{010}^+ maps to C_{011} isomorphically (merging in a 1, with a minus sign from the cube). The reduced complex will have underlying abelian group

$$C_{\text{red}} = C_{100} \oplus C_{001} \oplus C_{110} \oplus C_{101} \oplus C_{111}.$$

After chasing the diagram to find the maps into V and the maps out of W , you will find that the correction term $d_{C_{\text{red}}} d_{WV}^{-1} d_{WC_{\text{red}}}$ has four components.

$$\begin{aligned} C_{001} & \xrightarrow{-1} C_{110} \\ C_{111} & \xrightarrow{w_k - w_j} C_{110} \\ C_{110} & \xrightarrow{w_j - w_k} C_{100} \\ C_{110} & \xrightarrow{w_j - w_k} C_{001}. \end{aligned}$$

Each map is induced by the obvious cobordism relating the resolutions, weighted by some element of R indicated by the label on the arrow. Subtracting these from the restriction of the original differential d to C_{red} yields the complex pictured in Figure 5.7. Here, the edge labels give the total coefficient of the forward or reverse edge maps in d_{red} . The absence of a label on a forward edge maps the coefficient is $+1$. The label $i - j$, for example, denotes the coefficient $w_i - w_j$.

The complex (C', d') can be simplified using a similar cancellation. The relevant

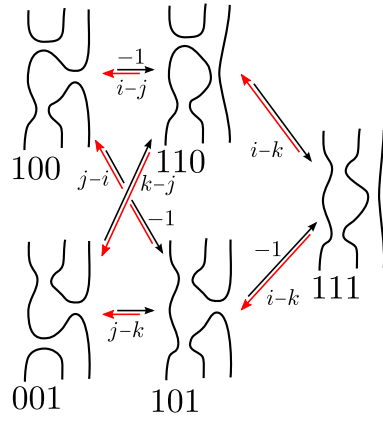


Figure 5.7: A reduced chain complex for the diagram D .




Figure 5.8: The chain complex for the post-R3 diagram D' .

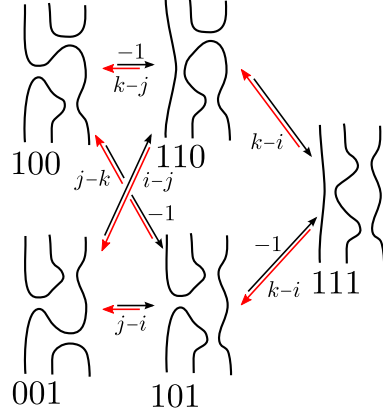


Figure 5.9: A reduced chain complex for the post-R3 diagram D' .

resolutions are drawn in Figure 5.8. Apply the cancellation lemma with

$$V = C'_{000} \oplus C'_{010} \quad W = C'_{010} \oplus C'_{011}.$$

The resulting complex, C'_{red} is pictured in Figure 5.9. It contains all the same resolutions as C_{red} ; the only difference is that all of the maps between pictured summands have reversed signs. The map $\phi : C_{\text{red}} \rightarrow C'_{\text{red}}$, defined by

$$\begin{aligned} C_{100} &\xrightarrow{1} C'_{001} \\ C_{001} &\xrightarrow{1} C'_{100} \\ C_{110} &\xrightarrow{-1} C'_{101} \\ C_{101} &\xrightarrow{-1} C'_{110} \\ C_{111} &\xrightarrow{1} C'_{111} \end{aligned}$$

is an invertible chain map. The sequence $C \cong C_{\text{red}} \cong C''_{\text{red}} \cong C'$ yields the desired isomorphism for diagrams related by Reidemeister 3. The first and third isomorphism preserve the filtration, since the cancelled components of the differential all preserved q . The second isomorphism preserves the bigrading, since the diagrams

satisfy $n_{\pm}(D) = n_{\pm}(D')$ and the map ϕ preserves the norm $|I|$ of each resolution I . □

We can now prove that the total homology of the complex for a link is just the Khovanov homology of the disjoint union of its components. This completes the proof of Theorem 1.1.7.

Theorem 5.2.3. *Let (L, w) be an R -weighted link, and suppose that for each pair of components i and j with distinct weights, the difference $w_i - w_j$ is invertible in R . Let D be any diagram for L . Let $L^{(r)}$ denote the sublink of L consisting of those components with weight r . Then the spectral sequence converges to*

$$H^*(C(D, w)) \cong Kh^* \left(\coprod_{r \in R} L^{(r)}; R \right)$$

Proof. Choose an arbitrary ordering \succ on the set $w_1, \dots, w_n \subset R$ of weights. By Proposition 5.1.3, changing a crossing between components with distinct weights will produce a chain complex $C(D', w)$ with the same l -graded total homology. So we may change crossings until each component i lies entirely over component j whenever $w_i \succ w_j$. This produces a diagram D' for some link L' , whose sublinks are still the $L^{(r)}$, now completely unlinked from one another. By repeated application of Reidemeister moves 1 and 2, we may slide these components off of one another until we get a diagram D'' for L' with no crossings between $L^{(r)}$ and $L^{(r')}$ for $r \neq r'$. The differential for $C(D'', w)$ is the same as Khovanov's differential, since $d_1 = 0$, and L' is just the disjoint union of the sublinks $L^{(r)}$. □

We can now give a stronger version of the rank inequality Corollary 1.1.9.

Corollary 5.2.4. *Let \mathbb{F} be any field, and let L be a link with components K_1, \dots, K_m . Then*

$$\text{rank}^{\ell} Kh^*(L; \mathbb{F}) \geq \text{rank}^{\ell+t} \otimes_{c=1}^m Kh^*(K_c; \mathbb{F}),$$

where each side is ℓ -graded and the shift t is given by

$$t = \sum_{c < d} 2\text{lk}(L_c, L_d).$$

Proof. Assume for the moment that the field \mathbb{F} has more elements than L has components, so we can weight each component by a different element $w_c \in \mathbb{F}$. Then all differences will be invertible, so the above theorem characterizing the abutment of the spectral sequence applies. That would give an inequality of total ranks. To see the ℓ -gradings, we need to compute the grading shift in the isomorphism relating $C(D, w)$ and $C(D', w)$ (we use the same notation as in the above proof). Recall the formula for the ℓ grading:

$$\ell(x) = \deg(x) - p(I) - \text{writhe}(D).$$

For a fixed monomial x over a fixed resolution I , the terms $\deg(x)$ and $p(I)$ are the same before and after a crossing change; only the writhe differs. Each time we change a crossing between components c and d , the writhe will shift by ± 2 and the linking number $\text{lk}(L_c, L_d)$ will shift by ± 1 . (The linking numbers with other components remain unchanged.) Thus

$$\ell(x) + \sum_{c < d} 2\text{lk}(L_c, L_d) = \ell(x') + \sum_{c < d} 2\text{lk}(L'_c, L'_d),$$

where x' is the same monomial viewed as a generator of $C(D', w)$, and L'_c is the component of L' which L_d turns into. But the components of L' are unlinked, so we ultimately have

$$\ell(x') = \ell(x) + \sum_{c < d} 2\text{lk}(L_c, L_d).$$

Now we address the size of \mathbb{F} . Since the differential in the chain complex computing $Kh(L)$ uses only ± 1 coefficients, its rank is the same after a field extension. We may take a suitably large extension \mathbb{F}' of \mathbb{F} , run the above argument for some choice of

weights, and then note that $\text{rank}'_{\mathbb{F}} Kh(L; \mathbb{F}') = \text{rank}_{\mathbb{F}} Kh(L; \mathbb{F})$. \square

5.3 Properties of spectral sequences

We offer a quick review of spectral sequences, following Serre [63]. Let (C, d) be a finitely generated chain complex. A filtration \mathcal{F} on C is an assignment to each element $x \in C$ a filtration degree $p(x) \in \mathbb{Z} \cup \{-\infty\}$ such that $p(x - y) \leq \max(p(x), p(y))$ and $p(dx) \leq p(x)$. We will occasionally write C^k for the k^{th} piece of the filtration $\mathcal{F}^k C = \{x \in C \mid p(x) \leq k\}$. Homological algebra usually concerns cycles and boundaries. The filtration provides notions of approximate cycles and early boundaries:

$$\begin{aligned} Z_r^k &= \{x \in C^k \mid dx \in C^{k-r}\} \\ B_r^k &= \{dy \in C^k \mid y \in C^{k+r}\}. \end{aligned}$$

The spectral sequence corresponding to the filtration is a sequence of chain complexes (E_r^k, d_r) , called *pages*, defined by

$$E_r^k = Z_r^k / Z_{r-1}^{k-1} + B_{r-1}^k.$$

If x is in Z_r^k , then dx is in Z_{r-1}^{k-1} : by definition $dx \in C^{k-r}$, and $d(dx) = 0$. The differential on E_r^k is then given by taking the equivalence class: $d_r[x] := [dx]$. The remarkable property of this sequence is that each page is the homology of the previous one: $E_{r+1}^k = H_*(E_r^k, d_r)$.

A spectral sequence is said to *collapse on page l* if $d_r = 0$ for all $r \geq l$.

Since C is finitely generated, there is some integer N such that, for all $r > N$, Z_r^k just consists of all cycles in degree $\leq k$ and B_r^k consists of all boundaries in degree $\leq k$ (that is, $Z_r^k = Z^k$ and $B_r^k = B^k$). The quotient Z^k / B^k is not the homology of the k^{th} filtered piece C^k , because B^k consists of elements of C^k which are boundaries

in C , not just boundaries of elements in C^k . In fact, the quotient is

$$Z^k/B^k \cong i_*H_*(C)$$

where $i : C^{k-1} \hookrightarrow C$ denotes the inclusion of the k^{th} filtered piece into the total complex. For all $r > N$, then, we have

$$\begin{aligned} E_r^k &= Z^k/Z^{k-1} + B^k \\ &= (Z^k/B^k) / (Z^{k-1}/B^{k-1}) \\ &= i_*H_*(C^k)/i_*H_*(C^{k-1}). \end{aligned}$$

We denote this stable page by E_∞^k , and observe that it is the associated graded group of the total homology $H_*(C)$ by the filtration

$$\mathcal{G}^k H_*(C) = i_*H_*(C^k).$$

In particular, the total rank of the E_∞ page is independent of the choice of filtration:

$$\sum_k \text{rank } E_\infty^k = \text{rank } H_*(C).$$

In contrast, the time of collapse does depend on the choice of filtration, though in a controlled way. (We doubt that the following proposition is original, but were unable to find it in the literature.)

Proposition 5.3.1. *Let (C, d) be a finitely generated chain complex, with two different filtrations \mathcal{F} and \mathcal{F}' which are close in the following sense: for any $x \in C$, the difference in filtration degree $p'(x) - p(x)$ is either 0 or 1. Then the p -spectral sequence collapses at most one page after the p' -spectral sequence does.*

Proof. Say that the p' -spectral sequence has collapsed by the $(r-1)^{\text{st}}$ page. We want

to show that any class $[x] \in E_r^k$ must have $d_r[x] = 0 \in E_r^{k-r}$, for then the p -spectral sequence will have collapsed on page r .

Suppose for the sake of contradiction that there is some $x \in Z_r$ such that $[x] \in E_r$ has nonzero differential. Without loss of generality, we may take the chain x with minimal $p'(x) + p'(dx)$. Let k be the degree $p(x)$, so $x \in Z_r^k$. If $p(dx) < k - r$, then $dx \in Z_{r-1}^{k-r-1}$ and $[dx]_r$ would represent 0 in E_r^{k-r} . Since $d_r[x] = [dx]$ is nontrivial, we must have $p(dx) = k - r$.

We now consider the p' -degrees of all the elements. Let $k' = p'(x)$ and $r' = p'(x) - p'(dx)$. Note that

$$\begin{aligned} r' &= p'(x) - p'(dx) \\ &= p'(x) - p(x) - (p'(dx) - p(dx)) + p(x) - p(dx) \\ &\in \{0, 1\} - \{0, 1\} + r \\ &\geq r - 1 \end{aligned}$$

Since the p' -spectral sequence, whose pages we will denote $E_*^*(p')$, has collapsed by page $r - 1$, it has also collapsed by page r' . And by construction, x represents a class in $E_{r'}^{k'}(p')$. Post-collapse, the differential is identically zero, so $d_{r'}[x]_{p'}$ must represent zero in $E_{r'}^{k'-r'}(p')$. In terms of chains, this means that

$$dx = w + dz$$

for some $w \in Z_{r'-1}^{k'-r'-1}$ with $p'(w) \leq k' - r' - 1$ and some $dz \in B_{r'-1}^{k'-r'}$ with $p'(z) \leq k' - 1$. Since p -gradings are at most one less than p' -gradings, $p(x) \geq k' - 1$ and $p(dx) \geq k' - r' - 1$. Consequently, $p(z) \leq p(x)$ and $p(w) \leq p(dx)$.

Since $dw = ddx - ddz = 0$, we have that $w \in Z_r^{k-r}$. Since $dz = dx - w$, we have $p(dz) \leq \max(p(dx), p(w))$, and $z \in Z_r^k$.

We break into two cases.

Case 1: $[w] = 0 \in E_r^{k-r}$.

Set $\bar{x} = z$. Then $[\bar{x}]$ is a class in E_r^k with

$$d_r[\bar{x}] = [dz] = [dz] + [w] = [dx] \neq 0.$$

But $p'(\bar{x}) = p'(z) < p'(x)$ and $p'(d\bar{x}) = p'(dx - w) = p'(dx)$, violating minimality.

Case 2: $[w] \neq 0 \in E_r^{k-r}$.

Set $\bar{x} = x - z$. Then $[\bar{x}]$ is a class in E_r^k with

$$d_r[\bar{x}] = [dx - dz] = [w] \neq 0.$$

But $p'(\bar{x}) = p'(x - z) \leq p'(x)$ and $p'(d\bar{x}) = p'(w) < p'(dx)$, violating minimality. \square

5.3.1 Endomorphisms of spectral sequences

Suppose that f is an endomorphism of the filtered chain complex C which shifts filtration degree by l ,

$$p(fx) = p(x) - l \quad \forall x \in C.$$

Then f acts on the spectral sequence the following sense

1. There is an endomorphism f_r of the r^{th} page given by

$$\begin{aligned} f_r : E_r^k &\rightarrow E_r^{k-l} \\ [x] &\mapsto [fx] \end{aligned}$$

This is well-defined: since f shifts $p(dx)$ by the same amount that it shifts $p(x)$, it takes Z_r^k into Z_r^{k-l} and B_r^k into B_r^{k-l} .

2. The action of f_{r+1} on E_{r+1} is the same as the one induced by f_r on the homology

of (E_r, d_r) .

3. The action of f_∞ on E_∞ is the associated graded action of

$$f_* : H_*(C) \rightarrow H_*(C)$$

with respect to the filtration \mathcal{G} above. That is, if $[x] \in \mathcal{G}^k = i_*H_*(C^k)$ is represented by $x \in Z^k$, then $fx \in Z^{k-l}$ and the image $f_*[x] = [fx]$ lies in \mathcal{G}^{k-l} . Moreover, x also serves as a representative of the equivalence class of $[x]_\infty \in E_\infty^k = \mathcal{G}^k/\mathcal{G}^{k-1}$ and $f_\infty[x]_\infty = [fx]_\infty$.

We will later encounter a spectral sequence where we know the action of an endomorphism X on $H_*(C)$ and investigate the possible associated graded actions on the E_∞ page.

5.4 The splitting number

The unknotting number of a knot is the minimum number of times the knot must be passed through itself to untie it. It is an intuitive measure of the complexity of a knot, though strikingly difficult to compute. We would like to suggest a similar number measuring the complexity of the linking *between* the components of a link, unrelated to the knotting of the individual components.

Definition 5.4.1. The splitting number of a link L , written $\text{sp}(L)$, is the minimum number of times the different components of the link must be passed through one another to completely split the link. Equivalently, $\text{sp}(L)$ is the minimum over all diagrams for L of the number of between-component crossings changes required to produce a completely split link.




Figure 5.10: The Whitehead link has splitting number 2.

A completely split link has splitting number 0. The Hopf link has splitting number 1, as demonstrated by the standard diagram. In general, any diagram for a link L gives an upper bound on $\text{sp}(L)$, as one may change crossings until the components of the link are layered one atop the next.

The Whitehead link L_W has splitting number 2—change two diagonally opposite crossings in the standard diagram (Figure 5.10). While changing the crossing in the center would split the link, that crossing is internal to one component so not allowed. To see that $\text{sp}(L_W) \neq 1$, note that a crossing change between components K_c and K_d of a link L changes the linking number $\text{lk}(K_c, K_d)$ by ± 1 . Since the Whitehead link has linking number 0, an even number of crossing changes will be required.

If L is a two-component link with components K_1 and K_2 , then the quantity

$$b_{\text{lk}}(L) := \begin{cases} |\text{lk}(K_1, K_2)| & \text{if } L \text{ is non-split and } \text{lk}(K_1, K_2) > 0 \\ 2 & \text{if } L \text{ is non-split and } \text{lk}(K_1, K_2) = 0 \\ 0 & \text{if } L \text{ is split} \end{cases}$$

provides a lower bound on $\text{sp}(L)$. If L has many components, we define

$$b_{lk}(L) := \sum_{c < d} b_{lk}(L_{cd}),$$

where L_{cd} denotes the sublink consisting of the c^{th} and d^{th} components. Since splitting a link certainly requires that one change enough crossings to split each pair of components (and each crossing implicates only one two-component sublink), we conclude that

$$\text{sp}(L) \geq b_{lk}(L).$$

Our spectral sequence provides less obvious lower bound for the splitting number: the splitting number plus one is at least the index of the page on which the spectral sequence collapses.

Theorem 1.1.8. *Let L be a link and let $w_c \in R$ be a set of component weights such that $w_c - w_d$ is invertible for each pair of components c and d . Let $b(L, w)$ be the largest integer k such that $E_k(L, w) \neq E_\infty(L, w)$. Then $b(L, w) \leq \text{sp}(L)$.*

$$\text{sp}(L) \geq b(L, w)$$

Proof. We proceed by induction on splitting number. If L is a split link, then there is a diagram in which $d_1 = 0$ so the spectral sequence collapses immediately: $E_1 = E_\infty$ and $b(L) = 0$.

If L is non-split, then there is a diagram D in which changing exactly $k = \text{sp}(L)$ crossings produces a diagram for a split link. Consider the diagram D' resulting by changing just one of those crossings, say i ; the link L' depicted will have splitting number $k - 1$.

In the proof of Proposition 5.1.3, we constructed an isomorphism $f : C(D, w) \rightarrow C(D', w)$ of ℓ -graded chain complexes. To compare the filtrations and the spectral

sequences, we pull back the g -grading on $C(D', w)$ to a grading on $C(D, w)$, which we write g' . In an abuse of notation, we will write g for the original g -grading on $C(D, w)$. The two corresponding filtrations on $C(D, w)$ differ in a controlled way.

Recall that for a generator x of $C(D, w)$, the relevant gradings are

$$q(x) = \deg(x) + p(I) + |I| + n_+(D) - 2n_-(D) \quad \text{and} \quad g(x) = \frac{q(x) - |L|}{2}.$$

The monomial degree $\deg(x)$ and circle count $p(I)$ are the same in both D and D' . If x sits over the 0-resolution of the crossing i in D , then it sits over the 1-resolution of i in D' , and vice versa. So the value of $|I|$ differs by ± 1 between the two complexes. Finally, the difference $n_+(D) - 2n_-(D)$ decreases (increases) by 3 if i is a positive (negative) crossing in D . Thus the difference in filtration degree $g'(x) - g(x)$ is in $\{-1, -2\}$ if the crossing is positive and $\{1, 2\}$ if the crossing is negative.

Since a global shift in filtration degree does not affect the page at which the corresponding spectral sequences collapses, Proposition 5.3.1 applies. We conclude that the spectral sequence for L collapses at most one page after the spectral sequence for L' , so

$$b(L, w) \leq b(L', w) + 1 \leq \text{sp}(L') + 1 = \text{sp}(L). \quad \square$$

An interesting example is the link $L = {}^2L13n3752$ shown in Figure 1.1. The two components are a trefoil and the unknot, and they have linking number 1. There is an obvious way to split the L by changing three crossings, say, pulling the red component on top of the green one. The spectral sequence with the nontrivial \mathbb{F}_2 weighting w , shown in Table 5.1, collapses on the E_3 page, so $\text{sp}(L) \geq b(L, w) = 2$. Since $\text{sp}(L)$ must have the same parity as the linking number, we have that $\text{sp}(L) = 3$.

The calculation of the spectral sequence for 2L13n3752 and many other links is discussed in Section 5.6.

Table 5.1: $E_k(^2L13n3752, w)$ over \mathbb{F}_2 with non-trivial weight function w . $E_1(L, w) = Kh(L)$ omitted.

Link L	E_k	rank E_k	$P_k(q, t) = \sum_{i,j} (\text{rank } E_k^{ij}) t^i q^j$
2L13n3752	E_2	20	$t^{-2}q^{-2} + t^{-2} + q^2 + q^4 + t^1q^2 + t^1q^4 + t^2q^4 + t^2q^6 + t^3q^6 + t^3q^8 + 2t^4q^8 + 2t^4q^{10} + t^5q^{10} + t^5q^{12} + t^6q^{12} + t^6q^{14} + t^7q^{14} + t^7q^{16}$
	E_3	12	$t^2q^4 + t^2q^6 + 2t^4q^8 + 2t^4q^{10} + t^5q^{10} + t^5q^{12} + t^6q^{12} + t^6q^{14} + t^7q^{14} + t^7q^{16}$

5.5 Detecting unlinks

In this section, we work over a field \mathbb{F} of characteristic 2. Since our construction relies on choosing different weights for different components, \mathbb{F}_2 itself is not large enough to accommodate many-component links. The specific choice of a larger field is unimportant, so we will write \mathbb{F} for some finite field of characteristic 2 with more elements than there are components of the link under consideration. Since $Kh(L; \mathbb{F}) \cong Kh(L; \mathbb{F}_2) \otimes \mathbb{F}$, the rank of Khovanov homology is independent of the choice of \mathbb{F} . For this reason, we will often write $Kh(L)$ for $Kh(L; \mathbb{F})$.

Kronheimer and Mrowka have shown that Khovanov homology detects the unknot. That is, if a knot K has $Kh(K)$ of rank 2, then K is the unknot.

Corollary 1.1.9 provides an immediate upgrade.

Proposition 5.5.1. *Let L be an m -component link, and suppose that the rank of $Kh(L)$ is 2^m . Then each component of L is an unknot.*

Proof. Let K_1, \dots, K_m be the components of L . By Corollary 1.1.9, we have a rank inequality

$$\text{rank } Kh(L) \geq \text{rank } Kh(K_1) \times \text{rank } Kh(K_2) \times \cdots \times \text{rank } Kh(K_m).$$



Figure 5.11: A forest F gives rise to a link L_F whose Khovanov homology has the same rank as that of the unlink.

The left-hand-side is 2^m . Since every knot has Khovanov homology of rank at least two, the right-hand side is at least 2^m . Hence every one of the components K_i must have $\text{rank}(Kh(K_i)) = 2$. By Kronheimer and Mrowka's result, each of those components is an unknot. \square

Equality is possible: the Hopf link has rank four, just like the two-component unlink. This generates a family of such examples: iterated connect-sums and disjoint unions of Hopf links and unknots. The resulting links can be described as forests of unknots: given a (planar) forest F , form a link L_F by placing an unknot at each vertex then clasping them together along each edge (Figure 5.11). By [64], we have $\text{rank } Kh(L_F) = \text{rank } Kh(U^m)$.

Question 5.5.2. *Are forests of unknots the only m -component links with Khovanov homology of rank 2^m over \mathbb{F}_2 ?*

None of these nontrivial links have the same bigradings at the unlink. As we will show later, this is no coincidence.

5.5.1 The Khovanov module

Khovanov homology is not just an abelian group: $Kh(L)$ is a module over the component algebra

$$A_m = \mathbb{F}_2[X_1, \dots, X_m]/(X_1^2, X_2^2, \dots, X_m^2),$$




Figure 5.12: Moving a marked point across a crossing.

see [22]. The module structure is defined by choosing marked points p_c on each component c of L . Then X_c acts by X_{p_c} .

In fact, this module structure extends to all the pages E_k . The map X_{p_c} shifts the g gradings by -1 , so it preserves the filtration \mathcal{F}^p . It remains to show that X_p for a marked point p is a chain map with respect to the total differential d and that the module structure induced on E_k is independent of the choice of marked points.

Proposition 5.5.3. *Let p be a marked point on D away from the double points. Then we have that $dX_p = X_p d$.*

Proof. X_p commutes the Khovanov edge maps; this is the standard proof that it commutes with d_0 . The deformation d_1 is also a sum of edge maps, so the proposition follows. \square

Proposition 5.5.4. *Let p and q be marked points on either side of crossing i as shown in Figure 5.12. Then X_p and X_q are chain homotopic.*

Proof. We use the same chain homotopy H as in the proof that the Khovanov module is well-defined [22]:

$$H = \sum_{\substack{\text{resolutions } I \\ I(i)=1}} \mathcal{A}(I, J_i),$$

where J_i differs from I solely at i . Hedden and Ni show $X_p + X_q = Hd_0 + d_0H$. It remains for us to show that $Hd_1 + d_1H = 0$. This is an immediate consequence of the facts that H and d_1 both decrease homological grading and that the reverse edge maps commute. \square

Since we use \mathbb{F} coefficients to define the complex $C(D, w)$, we will first prove results regarding the action of $A_m^{\mathbb{F}} := A_m \otimes \mathbb{F}$. The Khovanov module of the unknot is just a copy of $A_1^{\mathbb{F}}$ viewed as a module over itself: $Kh(U) \cong \mathbb{F}[X]/(X^2)$. Disjoint union of links gives tensor products of modules, over the tensor product algebra; in particular, $Kh(U^m) \cong A_m^{\mathbb{F}}$.

Proposition 5.5.5. *Let L be a m -component link with $\text{rank } Kh(L) = 2^m$. If D is any diagram for L and w any set of distinct weights for the components, then the total homology $H_*(C(D, w))$ is a free rank-one module over the algebra $A_m^{\mathbb{F}}$.*

Proof. By Proposition 5.5.1, the components of L are all unknots. Number the components from 1 to m . Since we are only interested total homology and its module structure, we can ignore the g -filtration on $C(D, w)$. We can produce a diagram D' for U^m by swapping mixed crossings in D so that at each crossing, the under-strand has lower index than the over-strand. As we saw in the proof of Proposition 5.1.3, $C(D, w)$ and $C(D', w)$ differ only by rescaling generators by elements of \mathbb{F} . The action of X_{c_p} commutes with rescaling generators. The total homology and $A_m^{\mathbb{F}}$ action are also invariant under Reidemeister and marked point moves. By such moves, D' can be transformed into D'' , the standard diagram for U^m with no crossings: a disjoint collection of circles with marks. The complex for D'' has vanishing differential, and $H_*(C(D'', w))$ is manifestly a free rank-one $A_m^{\mathbb{F}}$ -module, as desired. \square

5.5.2 Proof of Theorem 1.1.10

Hedden and Ni have shown that the module structure of Kh detects the unlink [22].

Theorem 5.5.6 (Hedden-Ni). *Let L be an m -component link. If there is an isomorphism of A_m modules*

$$Kh(L; \mathbb{F}_2) \cong A_m,$$

then L is the unlink.

We can deduce the module structure from the bigradings.

Theorem 1.1.10. *Let L be an m -component link, and U^m the m -component unlink.*

If

$$\text{rank } Kh^{i,j}(L; \mathbb{F}_2) = \text{rank } Kh^{i,j}(U^m; \mathbb{F}_2)$$

for all i, j , then L is the unlink.

Proof. The Khovanov homology of the unlink is supported entirely in homological grading 0, where it has rank $\binom{m}{r}$ in quantum grading $2r - m$. Since our spectral sequence is graded by $g = (q - m)/2$, the group

$$E_1^{-k}(L, w) \cong Kh^{0, m-2k}(L)$$

has rank $\binom{m}{k}$ for $0 \leq k \leq m$.

As described in Section 5.3.1, there is a filtration \mathcal{G} on the total homology $H = H_*(C(D, w))$ with respect to which

$$E_\infty^{-k} \cong \mathcal{G}^{-k} H / \mathcal{G}^{-k-1} H .$$

Since the spectral sequence collapses with $E_1 = E_\infty$, this determines the rank of each filtered piece,

$$\text{rank}_{\mathbb{F}} \mathcal{G}^{-k} H = \binom{m}{k} + \binom{m}{k+1} + \cdots + \binom{m}{m}$$

Let I be the (maximal) ideal in $A_m^{\mathbb{F}}$ generated by the X_i . The top nonvanishing power of the ideal is I^m , which is spanned by the element $X_1 X_2 \cdots X_m$, and we have $I^{m+1} = 0$. Consider the filtration

$$0 \subset I^m \subset I^{m-1} \subset \cdots \subset I \subset A_m^{\mathbb{F}}.$$

By Proposition 5.5.5, the total homology is a free rank-one module over $A_m^{\mathbb{F}}$,

generated by some $e \in \mathcal{G}^0 H \cong H$. Moreover, since each endomorphism X_i lowers the g -grading by 1, it takes \mathcal{G}^{-k} into \mathcal{G}^{-k-1} . Hence

$$I^k e \subset \mathcal{G}^{-k} H$$

for every $0 \leq k \leq m$.

Since e is the generator of a free $A_m^{\mathbb{F}}$ -module, we know that $I^k e$ actually has the same rank as I^k itself, which is the same as the rank of $\mathcal{G}^{-k} H$ computed above. Hence $I^k e = \mathcal{G}^{-k} H$.

The associated graded module is

$$\bigoplus_k I^k e / I^{k+1} e \cong \bigoplus_k A_m^{\mathbb{F}}[k]e,$$

where $A_m^{\mathbb{F}}[k]$ denotes the linear span of the monomials of degree k in the X_i . This is isomorphic to $A_m^{\mathbb{F}}$ itself, viewed as an $A_m^{\mathbb{F}}$ -module. But $E_{\infty}(L, w) \cong E_1(L, w) \cong Kh(L)$. So $Kh(L)$ is a free, rank-one $A_m^{\mathbb{F}}$ module.

More precisely, $Kh(L; \mathbb{F})$ is a free, rank-one $\mathbb{F}[X_1, \dots, X_n]/X_i^2$ -module. To apply Hedden-Ni, and conclude that L is the unlink, we need to show that $Kh(L; \mathbb{F}_2)$ is a free, rank-one $\mathbb{F}_2[X_1, \dots, X_n]/X_i^2$ -module. In general, extending the ground field can make a free module out of a non-free one [10]. This cannot happen for A_m -modules, essentially because A_m is a local ring.

Indeed, suppose that M is a module over A_m such that $M_{\mathbb{F}} = M \otimes \mathbb{F}$ is a free rank-one module over $A_m \otimes \mathbb{F}$. Let $a \in M_{\mathbb{F}}$ be a generator, so the \mathbb{F} -span of $A_m a$ is all of $M_{\mathbb{F}}$. Now pick some element b of the original module M such that $b \notin I \cdot M_{\mathbb{F}}$. Then $b = \alpha(1 + X)a$ where $\alpha \in \mathbb{F}$ and $X \in I$. Because I is nilpotent of order $m + 1$, the coefficient $\alpha(1 + X)$ is a unit with inverse $\alpha^{-1}(1 - X + X^2 + \dots \pm X^m)$. Thus b is also a generator for $M_{\mathbb{F}}$ as a free $A_m \otimes \mathbb{F}$ -module. In particular, b is not annihilated

by any element of A_m . This means that

$$\text{rank}_{\mathbb{F}_2} A_m b = 2^m = \text{rank}_{\mathbb{F}} M_{\mathbb{F}} = \text{rank}_{\mathbb{F}_2} M.$$

Hence $A_m b$ is all of M , and M is a free A_m -module. □

Example. It is instructive to see where this argument breaks down for the Hopf link. There, $Kh(L) = E_{\infty}$ has total rank four, with rank-one summands in g -degrees $-1, 0, 1, 2$. Thus the filtration of the rank-1 free A_n -module $H^*(C(D, w))$ has ranks

$$1 < 2 < 3 < 4.$$

Let e be a generator, and write $x_i \dots x_j$ for $X_i \dots X_j e$. The filtration is

$$\langle x_1 x_2 \rangle \subset \langle x_1 x_2, x_1 + x_2 \rangle \subset \langle x_1 x_2, x_1, x_2 \rangle \subset \langle x_1 x_2, x_1, x_2, 1 \rangle$$

The associated graded has a nonstandard module structure:

$$\langle a, b | X_1 a = X_2 a, X_1 b = X_2 b \rangle.$$

In contrast, the two component unlink has a filtration of ranks $1 < 3 < 4$, and the associated graded is isomorphic A_2 itself.

5.6 Sample computations

The combinatorial definition of the spectral sequence makes it amenable to computer calculation. We use `knotkit`, a C++ software package for knot homology computations written by the second author, to compute the spectral sequence for thousands of links [60].

These computations show that the spectral sequence is not determined by the Khovanov homology of the links involved. The links 2L12n817 and 2L14n38362 have the same Khovanov homology, and each has two unknot components (see Figure 5.13). Yet the spectral sequences collapse on different pages (E_2 vs E_3).




Figure 5.13: Links 2L12n817 (left) and 2L14n38362 (right).

Table 5.2: The link splitting spectral sequence $E_k(L)$ over $R = \mathbb{Z}_2(x_1, \dots, x_m)$ with weight function $w_c = x_c$ for examples in this section. $E_1(L) = Kh(L)$ omitted.

Link L	E_k	rank E_k	$P_k(q, t) = \sum_{i,j} (\text{rank } E_k^{ij}) t^i q^j$
2L12n817	E_2	4	$q^{-2} + 2 + q^2$
2L14n38362	E_2	68	$t^{-5}q^{-10} + t^{-5}q^{-8} + 2t^{-4}q^{-8} + 2t^{-4}q^{-6} + 2t^{-3}q^{-6} + 2t^{-3}q^{-4} + t^{-2}q^{-6} + 4t^{-2}q^{-4} + 3t^{-2}q^{-2} + 2t^{-1}q^{-4} + 5t^{-1}q^{-2} + 3t^{-1} + 3q^{-2} + 6 + 3q^2 + 3t^1 + 5t^1q^2 + 2t^1q^4 + 3t^2q^2 + 4t^2q^4 + t^2q^6 + 2t^3q^4 + 2t^3q^6 + 2t^4q^6 + 2t^4q^8 + t^5q^8 + t^5q^{10}$
	E_3	4	$q^{-2} + 2 + q^2$

We have computed splitting number bounds for all links with 12 or fewer crossings in the Morwen hyperbolic link tables from SnapPy [16]. Some choices and approximations must be made, which we describe before giving the results.

We use two coefficient rings, $\mathbb{P} = \mathbb{Z}/2(x)$ and \mathbb{Q} . For the former, we weight component c by $w_c = x^c$, and for the latter, we weight component c by the integer c

itself.

Since `knotkit` is not currently able to detect split links, we need an approximation to the bound coming from the linking number, b_{lk} . Since the link table contains only non-split links, there is no problem for two-component links. But non-split links with more than two components, such as the Borromean rings, may have split sublinks. We define b'_{lk} as follows: If L has two components and is known to be non-split, we set $b'_{lk}(L) = b_{lk}(L)$. If $L = K_1 \cup K_2$ may be split, then we define

$$b'_{lk}(L) = \begin{cases} |lk(K_1, K_2)| & lk(K_1, K_2) \neq 0 \\ 2 & Kh(L; \mathbb{Z}/2) \not\cong Kh(K_1 \amalg K_2; \mathbb{Z}/2) \cdot \\ 0 & \text{otherwise} \end{cases}$$

If L has more than two components and is non-split, we define

$$b'_{lk}(L) = \max\left(\sum_{i < j} b'_{lk}(L_{ij}), 2\right),$$

where L_{ij} is the sublink of L consisting of the i^{th} and j^{th} components.

Any diagram D for a link L gives an upper bound on the splitting number. Number the components of L from 1 to m . Let $\sigma \in S_m$ be a permutation of the components. We can produce a diagram D' for a split link by swapping the $u(D, \sigma)$ crossings of D where $\sigma(c_{\text{upper}}) < \sigma(c_{\text{lower}})$. Let $u(D)$ be the minimum of $u(D, \sigma)$ over all σ , so $u(D)$ is a upper bound for $\text{sp}(L)$.

We computed

$$b'_{lk}(L), b^{\mathbb{Q}}(L, w^{\mathbb{Q}}), b^{\mathbb{P}}(L, w^{\mathbb{P}}), u(D),$$

for all 5698 links in the Morwen link table with 12 or fewer crossings, where D is the tabulated minimal diagram. Of those links, 4770 (83.7%) have non-trivial lower bounds for $\text{sp}(L)$ and the lower bound is known to be tight for 3587 (63%) links. Our

upper bound is very rough, so the lower bound is likely to be tight in many more cases. The bound coming from the spectral sequence is stronger than the linking number bound for 17 of those links and equal to it for 2421. The examples with $b > b'_{lk}$ are shown Table 5.3. For those 17 examples, we verified by hand that $b'_{lk} = b_{lk}$.

The paper [11] tabulates all links with 10 or fewer crossings for which any of b'_{lk} , $b^{\mathbb{Q}}$ or $b^{\mathbb{P}}$ give a nontrivial lower bound on $\text{sp}(L)$.

Table 5.3: Knots with 12 or fewer crossings for which $b'_{lk}(L) < b(L)$. $u(D)$ gives the upper bound on $\text{sp}(L)$.

Link L	$b_{lk}(L)$	$b^{\mathbb{P}}(L)$	$b^{\mathbb{Q}}(L)$	$\text{sp}(L)$
2L12n1342	1	0	2	3
2L12n1350	1	1	2	3
2L12n1357	1	1	2	3
2L12n1363	1	1	2	3
2L12n1367	1	1	2	3
2L12n1374	1	1	2	3
2L12n1404	1	1	2	3
2L11a372	1	1	2	3-5
2L12a1233	1	1	2	3-5
2L12a1264	1	1	2	3-5
2L12a1384	1	1	2	3-5
2L12n1319	1	2	2	3-5
2L12n1320	1	2	2	3-5
2L12n1321	1	2	2	3-5
2L12n1323	1	2	2	3-5
2L12n1326	1	2	2	3-5
3L12a1622	1	1	2	3-5

Bibliography

- [1] *Sage*, URL: <http://www.sagemath.org/>.
- [2] M. Atiyah, *Topological quantum field theories*, Inst. Hautes Études Sci. Publ. Math. (1988), no. 68, 175–186 (1989).
- [3] J. Baez and J. Dolan, *Categorification*, arXiv.org (1998), arXiv:math/9802029.
- [4] J. Baldwin, *On the spectral sequence from Khovanov homology to Heegaard Floer homology*, Int. Math. Res. Not. IMRN (2011), no. 15, 3426–3470.
- [5] J. Baldwin and A. Levine, *A combinatorial spanning tree model for knot Floer homology*, arXiv.org (2011), arXiv:1105.5199.
- [6] D. Bar-Natan, *KnotTheory*[‘], URL: http://katlas.math.toronto.edu/wiki/The_Mathematica_Package_KnotTheory[‘].
- [7] ———, *On Khovanov’s categorification of the Jones polynomial*, Algebraic & Geometric Topology **2** (2002), 337–370.
- [8] ———, *Khovanov’s Homology for Tangles and Cobordisms*, Geometry & Topology **9** (2005), no. 33, 1443–1499.
- [9] ———, *Fast Khovanov Homology Computations*, arXiv.org (2006), arXiv:math/0606318.

- [10] J. Batson, *If M is not a free A -module, can tensoring with a bigger field make it free?*, MathOverflow (version: 2013-01-15), URL: <http://mathoverflow.net/questions/118939>.
- [11] J. Batson and C. Seed, *A link splitting spectral sequence in Khovanov homology*, arXiv.org (2013), arXiv:1303.6240v1.
- [12] S. Beier, *An integral lift, starting in odd Khovanov homology, of Szabó's spectral sequence*, arXiv.org (2012), arXiv:1205.2256.
- [13] J. Bloom, *Odd Khovanov homology is mutation invariant*, Math. Res. Lett. **17** (2010), no. 1, 1–10.
- [14] ———, *A link surgery spectral sequence in monopole Floer homology*, Advances in Mathematics **226** (2011), no. 4, 3216–3281.
- [15] A. Champanerkar and I. Kofman, *Spanning trees and Khovanov homology*, Proc. Amer. Math. Soc. **137** (2009), no. 6, 2157–2167.
- [16] M. Culler, N. Dunfield, and J. Weeks, *Snappy, a computer program for studying the geometry and topology of 3-manifolds*, URL: <http://snappy.computop.org>.
- [17] S. Donaldson, *An application of gauge theory to four-dimensional topology*, J. Differential Geom. **18** (1983), no. 2, 279–315.
- [18] S. Eliahou, L. Kauffman, and M. Thistlethwaite, *Infinite families of links with trivial Jones polynomial*, Topology **42** (2003), no. 1, 155–169.
- [19] A. Floer, *An instanton-invariant for 3-manifolds*, Communications in Mathematical Physics **118** (1988), no. 2, 215–240.
- [20] ———, *The unregularized gradient flow of the symplectic action*, Comm. Pure Appl. Math. **41** (1988), no. 6, 775–813.

- [21] M. Freedman, R. Gompf, S. Morrison, and K. Walker, *Man and machine thinking about the smooth 4-dimensional Poincaré conjecture*, Quantum Topol. **1** (2010), no. 2, 171–208.
- [22] M. Hedden and Y. Ni, *Khovanov module and the detection of unlinks*, Geometry & Topology **17** (2013), 3027–3076.
- [23] J. Hoste, M. Thistlethwaite, and J. Weeks, *The first 1,701,936 knots*, Math. Intelligencer **20** (1998), no. 4, 33–48.
- [24] M. Jacobsson, *An invariant of link cobordisms from Khovanov homology*, Algebraic & Geometric Topology **4** (2004), 1211–1251.
- [25] T. Jaeger, *A Remark on Roberts’ Totally Twisted Khovanov Homology*, arXiv.org (2011), arXiv:1109.1805.
- [26] V. Jones, *A polynomial invariant for knots via von Neumann algebras*, Bull. Amer. Math. Soc. (N.S.) **12** (1985), no. 1, 103–111.
- [27] L. Kauffman, *State models and the Jones polynomial*, Topology **26** (1987), no. 3, 395–407.
- [28] W. Thomson (Kelvin), *On Vortex Motion*, Trans. R. Soc. Edin. **25** (1867), 217–260.
- [29] M. Khovanov, *A categorification of the Jones polynomial*, Duke Math. J. **101** (2000), no. 3, 359–426.
- [30] ———, *Patterns in knot cohomology I*, Experiment. Math. **12** (2003), 365–374.
- [31] J. Kock, *Frobenius algebras and 2-d topological quantum field theories*, Cambridge University Press, 2003.

- [32] D. Kriz and I. Kriz, *A Spanning Tree Cohomology Theory for Links*, arXiv.org (2011), arXiv:1109.0064.
- [33] P. Kronheimer and T. Mrowka, *Gauge theory for embedded surfaces, I*, *Topology* **32** (1993), no. 4, 773–826.
- [34] ———, *Monopoles and three-manifolds*, Cambridge University Press, 2007.
- [35] ———, *Knot homology groups from instantons*, arXiv.org (2008), arXiv:0806.1053v2.
- [36] ———, *Khovanov homology is an unknot-detector*, *Publications mathématiques de l’IHES* (2011), 1–112.
- [37] ———, *Khovanov homology is an unknot-detector*, *Publ. Math. Inst. Hautes Études Sci* **113** (2011), no. 1, 97–208.
- [38] ———, *Gauge theory and Rasmussen’s invariant*, *J. Topol.* **6** (2013), no. 3, 659–674.
- [39] C. Kutluhan, Y.-J. Lee, and C. Taubes, *HF=HM I: Heegaard Floer homology and Seiberg-Witten Floer homology*, arXiv.org (2010), arXiv:1007.1979.
- [40] E. Lee, *An endomorphism of the Khovanov invariant*, *Adv. Math.* **197** (2005), no. 2, 554–586.
- [41] R. Lipshitz, P. Ozsváth, and D. Thurston, *Bordered Floer homology and the branched double cover II*, in preparation.
- [42] ———, *BordProg*, URL: <http://www.math.columbia.edu/~lipshitz/research.html>.
- [43] ———, *Bordered Heegaard Floer Homology: Invariance and Pairing*, arXiv.org (2008), arXiv:0810.0687.

- [44] ———, *Slicing planar grid diagrams: a gentle introduction to bordered Heegaard Floer homology*, Proceedings of Gökova Geometry-Topology Conference 2008, Gökova Geometry/Topology Conference (GGT), Gökova, 2009, pp. 91–119.
- [45] ———, *Bordered Floer homology and the branched double cover I*, arXiv.org (2010), arXiv:1011.0499.
- [46] ———, *Computing \widehat{HF} by factoring mapping classes*, arXiv.org (2010), arXiv:1010.2550.
- [47] R. Lipshitz and S. Sarkar, *A khovanov homotopy type*, arXiv.org (2012), arXiv:1112.3932.
- [48] ———, *A refinement of Rasmussen's s -invariant*, arXiv.org (2012), arXiv:1206.3532.
- [49] C. Manolescu and P. Ozsváth, *On the Khovanov and knot Floer homologies of quasi-alternating links*, arXiv.org (2007), arXiv:0708.3249.
- [50] W. Menasco and M. Thistlethwaite, *The tait flyping conjecture*, Bulletin (New Series) of the American Mathematical Society **25** (1991), no. 2, 403–412.
- [51] K. Murasugi, *Jones polynomials and classical conjectures in knot theory*, Topology **26** (1987), no. 2, 187–194.
- [52] ———, *Jones polynomials and classical conjectures in knot theory. II*, Math. Proc. Cambridge Philos. Soc. **102** (1987), no. 2, 317–318.
- [53] P. Ozsváth and Z. Szabó, *Heegaard Floer homology of double-covers, Kauffman states, and Novikov rings*, in preparation.
- [54] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, The Annals of Mathematics (2004), 1027–1158.

- [55] ———, *On the Heegaard Floer homology of Branched Double-covers*, *Advances in Mathematics* **194** (2005), no. 1, 1–33.
- [56] J. Rasmussen, *Khovanov homology and the slice genus*, *Inventiones Mathematicae* (2004), 1–29.
- [57] ———, *Knot polynomials and knot homologies*, *Geometry and topology of manifolds*, *Fields Inst. Commun.*, vol. 47, Amer. Math. Soc., Providence, RI, 2005, pp. 261–280.
- [58] L. Roberts, *Totally Twisted Khovanov Homology*, *arXiv.org* (2011), arXiv:1109.0508.
- [59] S. Sarkar, C. Seed, and Z. Szabó, *A perturbation of the geometric spectral sequence in Khovanov homology*, in preparation.
- [60] C. Seed, *knotkit*, URL: <http://math.princeton.edu/~cseed/>.
- [61] C. Seed, *Twisting Szabó's geometric spectral sequence*, in preparation.
- [62] ———, *Computations of Szabó's geometric spectral sequence in Khovanov homology*, *arXiv.org* (2011), arXiv:1110.0735.
- [63] J. P. Serre, *Homologie singulière des espaces fibrés, applications*, *Ann. of Math.* **54** (1951), 425–505.
- [64] A. Shumakovitch, *Torsion of the Khovanov homology*, *arXiv.org* (2004), arXiv:math/0405474.
- [65] A. Stoimenow, *Knot data tables*, URL: <http://stoimenov.net/stoimeno/homepage/ptab/index.html>.
- [66] Z. Szabó, *A geometric spectral sequence in Khovanov homology*, *arXiv.org* (2010), arXiv:1010.4252.

- [67] P. Tait, *On knots I, II, III*, Scientific Papers, Volume 1, Cambridge University Press, London, 1900, pp. 273–347.
- [68] M. Thistlethwaite, *A spanning tree expansion of the Jones polynomial*, *Topology* **26** (1987), no. 3, 297–309.
- [69] ———, *Links with trivial Jones polynomial*, *J. Knot Theory Ramifications* **10** (2001), no. 4, 641–643.
- [70] S. Wehrli, *A spanning tree model for Khovanov homology*, *J. Knot Theory Ramifications* **17** (2008), no. 12, 1561–1574.
- [71] E. Witten, *Monopoles and four-manifolds*, *Math. Res. Lett.* **1** (1994), no. 6, 769–796.
- [72] B. Zhan, Personal communication.