

NOTES ON KHOVANOV HOMOLOGY

MELISSA ZHANG

CONTENTS

1. Introduction	2
1.1. Master list of topics	2
1.2. Acknowledgements	2
2. Knots and Topology	2
2.1. Knots and links	2
2.2. Link diagrams and Reidemeister moves	4
2.3. Jones polynomial via skein relation	5
2.4. Jones polynomial from Kauffman bracket	6
2.5. Jones polynomial via the Khovanov bracket	7
3. Khovanov homology	9
3.1. The Khovanov chain complex $\mathrm{CKh}(D)$	9
3.2. An underlying TQFT	14
3.3. Bar-Natan's tangle categories	19
3.4. Adding crossings : Boosting to chain complexes	21
3.5. Computational tools	22
3.6. Reidemeister invariance of Khovanov homology	25
3.7. (Projective) Functoriality	27
4. Applications of Khovanov homology	30
4.1. Surfaces in B^4	30
4.2. Obstruction to ribbon concordance	33
4.3. Lee homology and the s invariant	35
4.4. Non-isotopic (smooth) surfaces in the 4-ball; exotic disks	49
5. Legendrian, transverse, and annular links	52
5.1. Knots in the standard contact \mathbb{R}^3 and S^3	52
5.2. Plamenevskaya's transverse knot invariant	56
5.3. Annular Khovanov homology	57
5.4. Ng's Thurston-Bennequin bound for Legendrian knots	61
6. Spectral sequences, Kh 's relation to gauge / Floer theories	67
6.1. Spectral sequences from filtered complexes	67
6.2. Spectral sequences from bicomplexes	70
6.3. Khovanov-to-Floer spectral sequences	78
7. Khovanov stable homotopy type	91
7.1. The idea of Morse homology	92
7.2. The spatial refinement problem	95
7.3. Cohen–Jones–Segal construction	96
7.4. A framed flow category for Khovanov homology	105
7.5. The Burnside functor approach	110
8. Homological Skein Invariants for 4-manifolds	116
8.1. Skein Lasagna modules from Khovanov homology	116
8.2. Cabled Khovanov homology	116
8.3. Ian's lecture	116
9. Categorical and algebraic background	116
9.1. Mapping cones	116
References	118

1. INTRODUCTION

These notes will be changing during the course. They will likely start out quite informal and become more formal with iteration.

The goal is to take a journey through areas in/around Khovanov homology, to see some open problems, and also to see the algebraic and topological tools in action. For me, Khovanov homology has been a great way to learn a little bit from a lot of different areas, including homological algebra, algebraic geometry, symplectic geometry, homotopy theory, and of course low-dimensional topology.

1.1. Master list of topics. Here's a list of topics. We will not get to all of these; this is just a master list, which will help guide the pace of our lectures and also provide potential final project ideas.

- Jones polynomial to Khovanov homology
- TQFT, Frobenius algebra, Bar-Natan category, functoriality
- computation*; torsion*
- deformations, web-foam generalizations
- Rasmussen-Lee spectral sequence, [Karoubi envelope](#); Piccirillo and the Conway knot
- ribbon cobordisms
- reduced theory, relation to gauge/Floer theories, knot detection
- odd khovanov homology; Szabó geometric spectral sequence
- tangles and representations (Khovanov, Chen-Khovanov)
- annular, Legendrian, and transverse topology applications
- symplectic Khovanov homology; Cautis-Kamnizer, Anno, Anno-Nandakumar
- physical interpretations* (Witten, Gukov, Aganagic)
- spectrification (Lipshitz-Sarkar, Lawson-Lipshitz-Sarkar, Hu-Kriz-Kriz*)
- skein lasagna modules (Morrison-Walker-Wedrich, Manolescu-Neithalath)
- immersed curves (Kotelskiy-Watson-Zibrowius)

* Topics marked with an asterisk (*) are those that I will almost definitely not be lecturing about, or will only touch upon briefly.

Aside 1.1.1. This list has so far been populated from the following sources:

- (1) the topics I personally feel are fundamental
- (2) Khovanov-Lipshitz's recent survey [\[KL23\]](#)
- (3) chatting with you all

Remark 1.1.2. Turner's [A Hitchhiker's Guide to Khovanov Homology](#) contains exposition for a few of the topics above. [\(add: citation\)](#)

1.2. Acknowledgements. Ian Sullivan - discussions. Mat 280 grad topics course - feedback.

2. KNOTS AND TOPOLOGY

2.1. Knots and links. If you want to learn more about knots in S^3 , Rolfsen is a great resource. There is also a little green book by Livingston that also seems really useful. [\(add: references\)](#)

A knot is sometimes defined as a smooth embedding $S^1 \hookrightarrow S^3$.

Notice that we can

- reparametrize the embedding, preserving the image setwise
- perform an *ambient isotopy* on the knot ('isotop' the knot)

Definition 2.1.1. Let $K, K' : S^1 \hookrightarrow S^3$ be two smooth embeddings. We say K and K' are (smoothly, ambiently) *isotopic* if there exists a smooth family of diffeomorphisms

$$\{\varphi_t : S^3 \rightarrow S^3\}_{t \in [0,1]}$$

such that $\varphi_0 = \text{id}$ and $\varphi(1) \circ K = K'$.

Remark 2.1.2. Convince yourself of the following:

- (1) “Ambiently isotopic” is an equivalence relation. The equivalence classes are called *isotopy classes* of knots.
- (2) In other words, an ambient isotopy smoothly morphs the embedding K into the embedding K' , through a family of diffeomorphisms of S^3 .
- (3) An order-preserving reparametrization is an ambient isotopy.

Remark 2.1.3. In practice, when I say *knot*, I’m probably referring to either (1) the image of the knot in S^3 or (2) an entire isotopy class.

Example 2.1.4. ‘The’ *unknot* U is the equivalence class of the standard embedding $S^1 \hookrightarrow \mathbb{R}^2 \hookrightarrow \mathbb{R}^3 \hookrightarrow S^3$.

Definition 2.1.5. Any diffeomorphic copy of S^1 can have two possible orientations. The *orientation* of a knot is given by the direction $\frac{dK}{d\theta}$.¹

Definition 2.1.6. The *reverse* K^r (also denoted \bar{K}) of a knot K is the² knot obtained by precomposing K with an order-reversing diffeomorphism $\rho : [0, 1] \rightarrow [0, 1]$.

Let τ be an orientation-reversing diffeomorphism of S^3 . The *mirror* of a knot K is the knot $m(K) = \tau \circ K$.

Remark 2.1.7. An *unoriented knot* is just the knot after you forget about the orientation. You can think of this as the union of the isotopy classes of K and \bar{K} .

Exercise 2.1.8. First observe that the orientation of S^3 does not reverse under isotopy.

Prove that

- In general, $m(K)$ is not necessarily isotopic to K , even if we treat them as unoriented knots. You can prove this via a (counter)example. Proving this directly is hard. Use the Jones polynomial, introduced in Section 2.3.
- In general, K^r is not necessarily isotopic to K . There is no general algorithm for determining when $K \not\sim K^r$! For this problem, do a search online and see if you can find an example in the literature, as well as the relevant terms for describing this kind of symmetry.

Then, find examples of knots K that do happen to satisfy the following, and exhibit an explicit isotopy:

- $K \sim K^r$
- $K \sim m(K)$

Diagrammatically exhibit the isotopy by applying Reidemeister moves; see Section 2.2. If there’s an obvious part of the isotopy that’s easy to describe but takes a lot of Reidemeister moves, you can just note what you’re doing between two pictures.

Exercise 2.1.9. (Important) A *link* with $\ell \in \mathbb{Z}_{\geq 0}$ components is a smooth embedding

$$L : \coprod_{i=1}^{\ell} S^1 \hookrightarrow S^3$$

The link with zero components is called the *empty link*. The i th *component* is the embedding of the i th copy of S^1 into S^3 .

All of the definitions above can be generalized to links. Write down definitions for the following:

- (1) oriented link $L = (L, o)$ If we want to be explicit about the orientation of a link, we use the letter o for the orientation information.
- (2) unoriented (isotopy class of) L
- (3) $m(L)$, the mirror of a link L

¹This is terrible notation and should not ever appear again, because of how we use the term ‘knot’; see Remark 2.1.3.

²By using the article ‘the’, I’m using the term ‘knot’ in the sense of Remark 2.1.3 (2).

(4) L^r , the reverse of an oriented link (L, o)


How many orientations does a link with ℓ components have?

Remark 2.1.10. There are also other categories of links, such as *topological links* and *wild links*. We won't talk about these until, maybe, far later in the course.


2.2. Link diagrams and Reidemeister moves. We will now start abusing notation and language without comment, per Remark 2.1.3.

We've so far been implicitly using link diagrams to draw links. More formally, a *link diagram* is a compact projection of a [representative of a](#) link L onto the xy -plane such that the only intersections are transverse double points.

For example, here are some bad singularities:



Theorem 2.2.1 (Reidemeister, [\(add: cite\)](#), 1930s). If D and D' are two diagrams of the same link, then they are related by a finite sequence of the following moves:



Remark 2.2.2. (1) These are local pictures. You can rotate them.

(2) Note that there are technically many cases within each 'move', if you consider all the possible orientations of the interacting strands in these local pictures. This is important when producing invariants for oriented links.

(3) If D and D' are diagrams for the same link, then we will sometimes write $D \sim D'$. In other words, "there exists a sequence of Reidemeister moves between" is an equivalence relations on the set of link diagrams.

Exercise 2.2.3. (Unimportant) Explain why we don't need to include the following local move:



Definition 2.2.4. A *link invariant* F valued in a category \mathcal{C} is a machine that

- takes in a link diagram D

- and outputs $F(D) \in \text{Ob}(\mathcal{C})$

such that

$$D \sim D' \implies F(D) = F(D').$$

Remark 2.2.5. Note that a link invariant is not necessarily a functor. We will define the categories **Link** and **LinkDiag**, and define *functorial* link invariants later.


2.3. Jones polynomial via skein relation. Khovanov homology is a categorical lift of the Jones polynomial, so we will focus on this invariant first.

Remark 2.3.1. The first polynomial invariant of links was the Alexander polynomial, introduced in the 1920s. We will not discuss the Alexander polynomial until needed later.

The Jones polynomial was discovered by Vaughan Jones in the 1980s, and arose from his work in statistical mechanics. We will not study the original definition of his invariant, but you are welcome to look into it if it interests you. Perhaps a final project idea? (add: citation)

Definition 2.3.2. (/ Theorem / Algorithm / Conventions) The Jones polynomial is a link invariant valued in Laurent polynomials $\mathbb{Z}[q, q^{-1}]$ that is uniquely determined by the following recursion:

- base case: $J(\bigcirc) = 1$
- skein relation: $q^{-2}J(L_+) - q^2J(L_-) = (q^{-1} - q)J(L_0)$ where



Remark 2.3.3. The diagrams for L_+ , L_- , and L_0 above show a *positive crossing*, a *negative crossing*, and an *oriented resolution*, respectively. The singularity at the crossing can be resolved in two ways; L_0 is the only resolution that maintains the orientation of all the strands.

Remark 2.3.4. In class, we essentially used Jones' conventions; these produce a Laurent polynomial in the variable \sqrt{t} :

- base case: $V(\bigcirc) = 1$
- skein relation: $t^{-1}V(L_+) - tV(L_-) = (t^{1/2} - t^{-1/2})V(L_0)$.

If you look at a database of Jones polynomials, you'll likely find this convention used.

For the purposes of this course, we will stick with Khovanov's original conventions for the first few weeks of class.

Example 2.3.5. Here we use the skein relation to compute the Jones polynomial (with the conventions set in Definition 2.3.2) of the *unlink* of two components, $\bigcirc\bigcirc$.

$$q^{-2} \left(\text{positive crossing of two strands} \right) - q^2 \left(\text{negative crossing of two strands} \right) = (q^{-1} - q) \left(\text{two separate strands} \right)$$

$$q^{-2} \cdot 1 - q^2 \cdot 1 = (q^{-1} - q)J(\bigcirc\bigcirc)$$

$$J(\bigcirc\bigcirc) = \frac{q^2 - q^{-2}}{q - q^{-1}} = q + q^{-1}.$$


Exercise 2.3.6. Compute the Jones polynomials of the following links:

- (1) (positively linked) Hopf link
- (2) right-handed trefoil

Exercise 2.3.7. (Important)

- (1) Prove that for any link L , $J(L) = J(L^r)$.

- (2) Let K_1 and K_2 be knots, and let $K_1 \# K_2$ denote their *connected sum*³:



Prove that $J(K_1 \# K_2) = J(K_1)J(K_2)$. What is $J(K_1 \sqcup K_2)$?

- Exercise 2.3.8.** (a) Show why the Jones polynomial is invariant under an R1 move.
 (b) Prove that for any link L , $J(\bigcirc L) = (q + q^{-1})J(L)$.

Question 2.3.9. (Open) Does the Jones polynomial *detect* the unknot U ? In other words, if $J(D) = 1$, is D necessarily a diagram for the unknot?

2.4. Jones polynomial from Kauffman bracket. In this section, we generally follow [BN02], certainly with the same conventions. However, I may word things differently or rotate some pictures, for our future benefit.

Definition 2.4.1. The *Kauffman bracket* $\langle \cdot \rangle$ is defined by the recursion

- $\langle \emptyset \rangle = 1$
- $\langle \bigcirc L \rangle = (q + q^{-1})\langle L \rangle$
- $\langle \nearrow \rangle = \langle \searrow \rangle - q\langle \times \rangle$

Note that these local pictures are *unoriented*, unlike those in the skein relation for the Jones polynomial.

Remark 2.4.2. Each crossing has two possible smoothings, which we will name the 0-resolution and the 1-resolution, as shown:

$$\times \xleftarrow{0} \times \xrightarrow{1} \times.$$

- The 0-resolution is the one that you would naturally draw if you started at an over-strand and drew a smile. It gets no coefficient in the Kauffman bracket.
- The 1-resolution is the one that you would draw if you started at the over-strand and drew a frown. It gets a coefficient of $-q$ in the Kauffman bracket.

Note that I draw my smileys like this: ☺

The Kauffman bracket is *not* a link invariant. To see this, compare $\langle \bigcirc \rangle$ and $\langle \bigcirc \bigcirc \rangle$. It is, however, a *framed* invariant. May talk more about this later.

To fix this, we have to take into account the *writhe* of a diagram, which measures how ‘twisty’ the choice of diagram is. 🐻

Definition 2.4.3. Let D be an oriented link diagram.

- Let n be the number of crossings in D .
- Let n_+ be the number of positive crossings, and let n_- be the number of negative crossings.

The *writhe* of D is $\text{wr}(D) = n_+ - n_-$.

³I have not provided a precise definition here. This is a good time to Google the term yourself to see some examples of the connected sum operation.

Definition 2.4.4. The (*unnormalized*) Jones polynomial is defined by

$$\hat{J}(L) = (-1)^{n_-} q^{n_+ - 2n_-} \langle L \rangle$$

Note that we are treating L as both the link diagram and the link it represents. Equivalently, $\hat{J}(L) = (-q)^{-n_-} \cdot q^{\text{wr}(L)} \langle L \rangle$.

Remark 2.4.5. Alternatively, we can build the overall shifts into the bracket, using the following recursion:

- $\langle \emptyset \rangle_o = 1$
- $\langle \bigcirc L \rangle_o = (q + q^{-1}) \langle L \rangle_o$ Recall that the two orientations on the unknot are isotopic.
- (positive crossing) $\langle \nearrow \searrow \rangle_o = q \langle \searrow \rangle_o - q^2 \langle \nearrow \rangle_o$
- (negative crossing) $\langle \nwarrow \nearrow \rangle_o = q^{-2} \langle \nwarrow \rangle_o - q^{-1} \langle \nearrow \rangle_o$

The notation ' $\langle \cdot \rangle_o$ ' is not standard, and is used here only to distinguish it from the bracket used in [BN02].

Exercise 2.4.6. (1) How does the writhe of a diagram change under the Reidemeister moves?
 (2) How does the Kauffman bracket $\langle \cdot \rangle$ of a diagram change under the Reidmeister moves?
 (3) Verify that the bracket $\langle \cdot \rangle_o$ in Remark 2.4.5 computes \hat{J} , i.e. $\hat{J}(L) = \langle L \rangle_o$.

Example 2.4.7. Let H denote (oriented) Hopf link shown below, where $n = 2$, $n_+ = 2$, and $n_- = 0$.

First, we compute the Kauffman bracket of the shown diagram:

$$\begin{aligned}
 & \langle \text{Hopf link} \rangle \\
 &= \langle \text{Hopf link with crossing} \rangle - q \langle \text{Hopf link with crossing} \rangle \\
 &= \langle \text{Hopf link with crossing} \rangle - q \langle \text{Hopf link with crossing} \rangle \\
 &= \langle \text{Hopf link with crossing} \rangle - q \left(\langle \text{Hopf link with crossing} \rangle - q \langle \text{Hopf link with crossing} \rangle \right) \\
 &= (q + q^{-1})^2 - 2q(q + q^{-1}) + q^2(q + q^{-1})^2 \\
 &= q^4 + q^2 + 1 + q^{-2}.
 \end{aligned}$$

Therefore

$$\begin{aligned}
 \hat{J}(H) &= (-1)^{n_-} q^{n_+ - 2n_-} \langle H \rangle \\
 &= q^2(q^4 + q^2 + 1 + q^{-2}) = q^6 + q^4 + q^2 + 1.
 \end{aligned}$$

The normalized Jones polynomial, using these conventions, would be $\frac{\hat{J}(H)}{q + q^{-1}} = q^5 + q$.

2.5. Jones polynomial via the Khovanov bracket. Khovanov homology $\text{Kh}(\cdot)$ is a $(\mathbb{Z} \oplus \mathbb{Z})$ -graded (co)homology theory whose graded Euler characteristic recovers the Jones polynomial. The extra grading is usually called the *quantum grading* or the *internal grading*.⁴

⁴Sometimes I may say 'degree' instead of grading. I will likely pontificate on the terms *grading* and *degree* later in the course.

Definition 2.5.1. Let $C = (\bigoplus_{i,j \in \mathbb{Z}} C^{i,j}, d)$ be a bigraded chain complex of \mathbb{Z} -modules where, for each j , $\bigoplus_{i \in \mathbb{Z}} C^{i,j}$ is finite rank.

The *graded Euler characteristic* of C is

$$\chi_q(C) = \sum_{i,j \in \mathbb{Z}} (-1)^i q^j \text{rank } H^{i,j}(C).$$

Recall that the Euler characteristic of a CW complex can be computed using any CW decomposition; in particular, you do not need to compute the differentials in the CW chain complex to determine the Euler characteristic.

Similarly, we do not need to know the differential d to compute $\chi_q(C)$. So, for now, we will define Bar-Natan's *Khovanov bracket*, which lift the recursion in Definition 2.4.1 to the level of bigraded chain complexes:

Definition 2.5.2. The *Khovanov bracket* is defined by the axioms

- $\llbracket \emptyset \rrbracket = (0 \rightarrow \mathbb{Z} \rightarrow 0)$
- $\llbracket \bigcirc L \rrbracket = V \otimes \llbracket L \rrbracket$ where $V = \mathbb{Z}v_+ \oplus \mathbb{Z}v_-$ of graded dimension $q + q^{-1}$.
We will discuss the actual definition later.
- $\llbracket \bowtie \rrbracket = \text{Tot} \left(0 \rightarrow \llbracket \triangleright \rrbracket \xrightarrow{d} \llbracket \bowtie \rrbracket \{1\} \rightarrow 0 \right)$, where $\{1\}$ means ' q -grading shift of 1' (see Notation 2.5.5).

The terms of each chain complex at homological grading 0 are underlined. The *totalization* functor flattens a multi-dimensional complex into a one-dimensional chain complex. Since we are not considering the differential right now, we will not carefully define this here. If the tensor product of chain complexes looks unfamiliar to you, this is something you should look up.

Remark 2.5.3. (add: Maybe add a remark about Grothendieck groups.)

Theorem 2.5.4 (Khovanov, as interpreted by Bar-Natan). Given a link diagram L with n_{\pm} crossings of sign \pm respectively, the associated Khovanov chain complex is given by

$$\text{CKh}(L) = \llbracket L \rrbracket [n_-] \{n_+ - 2n_-\}.$$

The square brackets indicate a shift in the homological grading (see Notation 2.5.5). Then

$$\hat{J}(L) = \chi_q(\text{CKh}(L)) = \chi_q(\text{Kh}(L)).$$


Notation 2.5.5. We set our conventions for the grading shift functors $[n]$ and $\{m\}$ to agree with those appearing in [BN02]. Let $A = \bigoplus A^{\bullet, \bullet}$ be a bigraded \mathbb{Z} -module. Then $A[n]\{m\}$ is the bigraded \mathbb{Z} -module where

$$(A[n]\{m\})^{i,j} = A^{i-n, j-m}.$$


Visually, if A is plotted on the the $\mathbb{Z} \oplus \mathbb{Z}$ bigrading lattice, $A[n]\{m\}$ is obtained by grabbing A and moving it by the vector $ne_1 + me_2$.

Example 2.5.6. We now reorganize the computation in Example 2.4.7, as a primer for our formal introduction to Khovanov homology in the next section.

Let H denote both the following diagram as well as the underlying oriented Hopf link:



To compute $\llbracket H \rrbracket$, we draw the *cube of complete resolutions*:



We associate $V^{\otimes c}$ to each complete resolution containing c closed components. Each complete resolution corresponds to a binary string $u \in \{0, 1\}^2$, and we perform a quantum grading shift of $|u|$ on associated \mathbb{Z} -module, where $|u|$ is the number of 1's appearing in the bitstring u .

This yields the Khovanov bracket

$$\llbracket H \rrbracket = (\underline{V \otimes V} \rightarrow V \oplus V\{1\} \rightarrow V \otimes V\{2\}).$$

(Recall that the underline indicates the chain group at homological grading 0.)

The Khovanov chain complex (without specifying differentials) is therefore

$$\begin{aligned} \text{CKh}(H) &= (\underline{V \otimes V} \rightarrow V \oplus V\{1\} \rightarrow V \otimes V\{2\}) [n_-] \{n_+ - 2n_-\} \\ &= (\underline{V \otimes V} \rightarrow V \oplus V\{1\} \rightarrow V \otimes V\{2\}) [0] \{2\} \\ &= (\underline{V \otimes V\{2\}} \rightarrow V \oplus V\{3\} \rightarrow V \otimes V\{4\}). \end{aligned}$$

Schematically, can visualize the bigraded chain groups as follows, where each \bullet represents a copy of \mathbb{Z} :

				χ :
$\text{gr}_q = 6$			\bullet	$0 - 0 + 1 = 1$
$\text{gr}_q = 4$	\bullet	$\bullet\bullet$	$\bullet\bullet$	$1 - 2 + 2 = 1$
$\text{gr}_q = 2$	$\bullet\bullet$	$\bullet\bullet$	\bullet	$2 - 2 + 1 = 1$
$\text{gr}_q = 0$	\bullet			$1 - 0 + 0 = 1$
	$\text{gr}_h = 0$	$\text{gr}_h = 1$	$\text{gr}_h = 2$	

Here gr_h and gr_q denote the homological and quantum grading, respectively.

We conclude that, indeed, $\chi_q(\text{CKh}(H)) = q^6 + q^4 + q^2 + 1 = \hat{J}(H)$.

Example 2.5.7. For a very comprehensive computation of the Jones polynomial of the trefoil using the Kauffman bracket, see Equation (1) of [BN02].

3. KHOVANOV HOMOLOGY

3.1. The Khovanov chain complex $\text{CKh}(D)$. Given an oriented link diagram D representing a link L , the *Khovanov chain complex* $\text{CKh}(D)$ is a $(\mathbb{Z} \oplus \mathbb{Z})$ -graded chain complex of abelian groups. We define this chain complex throughout this section. Throughout, we will use the term *bigraded* in place of ‘ $(\mathbb{Z} \oplus \mathbb{Z})$ -graded’.

3.1.1. Cube of resolutions. The cube of resolutions will set up the topological story that the chain complex tells. Afterwards, we will replace each part of the cube of resolution with algebraic objects, to fully define the Khovanov chain complex.

Let n be the number of crossings in the diagram D , and pick an ordering for the crossings. Let c_i denote the i th crossing.

The n -dimensional *binary cube* $\{0,1\}^n$ is the poset of *binary strings* (a.k.a. *bitstrings*) of length n , with partial order given by $0 \prec 1$ and lexicographic order.

Notation 3.1.1. Recall some terms and notations used when discussing partially ordered sets:

- If $u \prec v$, then we say u *precedes* v (and v *succeeds* u).
- If $u \prec v$ and they differ at exactly one bit, then we say u is an *immediate predecessor* of v (and v is an *immediate successor* of u). In this case, we write $u \prec_1 v$.

We may think of a poset as a directed graph, with an edge $u \rightarrow v$ if $u \prec_1 v$.

Notation 3.1.2. For convenience, if $u \prec_1 v$ and they differ at the i th bit, then we use the following length- n string in $\{0,1,*\}^n$ to denote the edge $u \rightarrow v$:

$$u_1 u_2 \cdots u_{i-1} * u_{i+1} \cdots u_n = v_1 v_2 \cdots v_{i-1} * v_{i+1} \cdots v_n$$

For $u \in \{0,1\}^n$, let D_u denote the complete resolution of the diagram D where crossing c_i is smoothed according to the bit u_i :

$$\times \times \xleftarrow{0} \times \xrightarrow{1} \times$$

If $u \prec_1 v$, then D_u and D_v differ only in the neighborhood of one crossing, called the *active crossing* of the edge $u \prec_1 v$. All other crossings are *passive*.

Also, $u \prec_1 v$ represents either

- a *merge* of two *circles* (i.e. closed components) of D_u into one circle in D_v or
- a *split* of one circle of D_u into two in D_v .

Any of these circles are called *active circles* of the edge $u \prec_1 v$. All other circles are *passive*; they look the same in both D_u and D_v .

Finally, when assigning gradings, we will need the *Hamming weight* of bitstrings:

$$|u| = \sum_i u_i.$$

3.1.2. *Chain groups, distinguished generators, gradings.* We now describe the distinguished generators of the Khovanov chain complex.

Our chain groups will be bigraded by $\text{gr} = (\text{gr}_h, \text{gr}_q)$:

- The first grading is called the *homological grading*, denoted gr_h . Its shift functor is denoted by $[\cdot]$.
- The second grading is called the *quantum* (or *internal*) grading, denoted gr_q . Its shift functor is denoted by $\{\cdot\}$.

Let V denote the bigraded \mathbb{Z} -module $\mathbb{Z}v_+ \oplus \mathbb{Z}v_-$ with generators v_{\pm} in bigradings $(0, \pm 1)$.

Let $|D_u|$ denote the number of circles in the resolution D_u . The chain group lying above vertex u of the cube is $V^{\otimes |D_u|}[[u]]\{|u|\}$ generated by the 2^n length- n pure tensors

$$\{v_{\pm} \otimes \cdots \otimes v_{\pm}\}.$$

These are called the *distinguished* (i.e. chosen) generators of the chain group at this vertex.

Remark 3.1.3. The bigrading gr for a distinguished generator in $V^{\otimes k}$ is determined by the bigrading on V :

$$\text{gr}(a \otimes v_{\pm}) = \text{gr}(a) + \text{gr}(v_{\pm}).$$

For example, $\text{gr}(v_+ \otimes v_- \otimes v_-) = (0, -1)$.

Remark 3.1.4. Note that when we actually write down a chain complex, we implicitly homologically shift the chain groups. For example, the following chain complex is acyclic⁵:

$$\underline{V} \xrightarrow{1} V.$$

If we view the chain complex as a graded module C , and the differential as an endomorphism d , then

- $C = V \oplus V[1]$ and
- $d = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$

This *mapping cone* point of view will become useful later in the course. Now is a good time to review/read about *mapping cones* in homological algebra.

3.1.3. Differentials. To each edge of the cube $u \rightarrow v$, we assign a map according to whether the edge represents a merging of two circles or the splitting of one circle. We casually call these the ‘edge maps’.

- If $u \rightarrow v$ represents a merge, the map on tensor components corresponding to active circles is given by:


$$\begin{aligned} m : V \otimes V &\rightarrow V \\ v_+ \otimes v_+ &\mapsto v_+ \\ v_+ \otimes v_-, v_- \otimes v_+ &\mapsto v_- \\ v_- \otimes v_- &\mapsto 0 \end{aligned}$$

- If $u \rightarrow v$ represents a split, the map on active components is given by:


$$\begin{aligned} \Delta : V &\rightarrow V \otimes V \\ v_+ &\mapsto v_+ \otimes v_- + v_- \otimes v_+ \\ v_- &\mapsto v_- \otimes v_- \end{aligned}$$

The map is identity on the passive components of the tensor product.

Example 3.1.5. Here is a diagram for the right-handed trefoil, with a choice of ordering on the three crossings (in blue):




In the cube of resolutions, the edge corresponding to $101 \rightarrow 111$ is a merging of two circles into one, with the active crossing circled in dotted cyan:



We have chosen an ordering on the set of circles (in red) at each resolution so that we can identify the copies of V in the tensor product at each vertex. For example, the distinguished generator $v_- \otimes v_+$ at resolution D_u labels the smaller circle v_- and the larger v_+ .

The linear map d_{uv} is given by the bundle of arrows shown below, where we use shorthand notation for compactness (e.g. $v_{+++} := v_+ \otimes v_+ \otimes v_+$):

⁵i.e. homology vanishes



Observe that d_{uv} is identity on the passive circle labeled ‘1’ in both D_u and D_v .

Exercise 3.1.6. (Important)

- (a) Verify that the merge and split maps, as written, decrease the quantum grading by 1.

So, as part of the differential, we modify the merge and split maps to be gr_q -preserving maps

$$d_{uv} : V \otimes V \rightarrow V\{1\} \quad \text{or} \quad d_{uv} : V \rightarrow V \otimes V\{1\},$$

depending on whether the edge $u \rightarrow v$ corresponds to a merge or a split, respectively.

- (b) Verify that along each 2D face of the binary cube, the edge maps commute.

Therefore, in order to get a chain complex (where $d^2 = 0$), we will need to add some signs so that the faces instead *anticommute*.

- (c) For an edge $u \rightarrow v$ in the cube with active crossing c_i , associate the following sign:

$$s_{uv} = (-1)^{\sum_{j=1}^{i-1} u_j}.$$

In other words, s_{uv} measures the parity of the number of 1’s appearing before the $*$ in the label given to the edge in the binary cube (see Notation 3.1.2).

Verify that, for any face of the binary cube, and odd number of the four edges bounding that face will have sign assignment -1 .

This is not the only possible sign assignment. (add: Discussion about cochains on cube.)

3.1.4. *Global grading shifts and homology.* To compute Khovanov homology of an oriented link L using diagram D :

- (1) Draw the cube of resolutions.
- (2) Associate modules $V^{\otimes |D_u|}$ to each vertex u of the cube.
- (3) Associate linear maps $s_{uv}d_{uv}$ to each edge $u \rightarrow v$.
- (4) Flatten the complex, by taking direct sums along Hamming weights; the resulting complex is the Khovanov bracket, $\llbracket D \rrbracket$.
- (5) Add in the global bigrading shift $\{-n_-\}[n_+ - 2n_-]$ to get the *Khovanov chain complex*:

$$\text{CKh}(D) = \llbracket D \rrbracket \{-n_-\}[n_+ - 2n_-].$$

- (6) Take homology to get *Khovanov homology*:

$$\text{Kh}(L) = H^*(\text{CKh}(D)).$$

The Khovanov chain complex $\text{CKh}(D)$ is bigraded, and its differential d_{Kh} is a bidegree $(1, 0)$ endomorphism. Therefore $\text{CKh}(D)$ is really the direct sum of many chain complexes, one for each quantum grading:


$$\text{CKh}(D) = \bigoplus_{j \in \mathbb{Z}} \text{CKh}^{\bullet, j}(D).$$

The homology is bigraded by the homological and quantum gradings (indexed below by i and j , respectively):

$$\text{Kh}(L) = \bigoplus_{i, j \in \mathbb{Z}} \text{Kh}^{i, j}(L).$$


Example 3.1.7. We now finally compute the Khovanov homology of the Hopf link H from Examples 2.4.7 and 2.5.6.

In order to uniquely identify the generators at the different resolutions in the cube, we label the resolutions u, v, w, x and use these letters to denote the distinguished generators at each resolution:



- The vertices of the binary cube are labeled in **purple**.
- The ordering of the tensor factors (i.e. circles) at each resolution is shown in **pink**. Below each arrow, the **pink** text indicates the active circles in the source and target of that edge.
- The only edge with sign assignment -1 is the split map shown in bold **red**.

In shorthand, the Khovanov chain complex $\text{CKh}(H)$ (i.e. with global shifts incorporated) is the following chain complex (direct sums are taken vertically):



The bigradings of the distinguished generators of $\text{CKh}(H)$ are shown in **pink**.

The reader may now verify that the Khovanov homology of the positively-linked Hopf link H is the bigraded \mathbb{Z} -module

$\text{gr}_q = 6$			\mathbb{Z}
$\text{gr}_q = 4$			\mathbb{Z}
$\text{gr}_q = 2$	\mathbb{Z}		
$\text{gr}_q = 0$	\mathbb{Z}		
	$\text{gr}_h = 0$	$\text{gr}_h = 1$	$\text{gr}_h = 2$

generated by the homology classes

$\text{gr}_q = 6$			$[x_{++}]$
$\text{gr}_q = 4$			$[x_{+-}] = [x_{-+}]$
$\text{gr}_q = 2$	$[u_{+-} - u_{-+}]$		
$\text{gr}_q = 0$	$[u_{--}]$		
	$\text{gr}_h = 0$	$\text{gr}_h = 1$	$\text{gr}_h = 2$

Exercise 3.1.8. The diagrams D and D' below both represent the unknot, U .



- Compute the Khovanov chain complex CKh for both, and then compute homology to verify that they indeed agree.
- Can you prove that Khovanov homology is invariant under the following Reidemeister 1 move?



Bar-Natan's introductory paper does cover this, but try first to use intuition from your solution to part (a) to find the appropriate chain homotopy equivalence.

3.2. An underlying TQFT. The algorithm/definition in the previous section is concrete, but you might be wondering where these m and Δ maps come from. The answer lies in the fact that morphisms in the category of bigraded \mathbb{Z} -modules are images of cobordisms under a functor from a more topologically defined category. In this section, we describe an⁶ underlying TQFT that determines the maps m and Δ from the previous section.

3.2.1. TQFTs. QFTs (quantum field theories) and TQFTs (topological quantum field theories), have a rich history in mathematical physics. The purposes of this course, we will use the following simplified (i.e. vague) definition, adapted from [Ati88]:

Definition 3.2.1. (Vague⁷) Let R be a commutative ground ring. An $(n + 1)$ -dimensional TQFT is a functor Z from a category of closed n -dimensional manifolds and $(n + 1)$ -dimensional cobordisms⁸ between them to a category of finitely generated (see Remark 3.2.3) R -modules, such that

- Z is *multiplicative*: $Z(Y \sqcup Y') = Z(Y) \otimes Z(Y')$.
- Z is *involutive*: If \bar{Y} is Y with the opposite orientation, then $Z(\bar{Y}) = Z(Y)^*$ (the dual module).

⁶We will see variations later, but focus first on the one that's easiest to work with.

⁷We are more interested in specific TQFTs, so will not emphasize the details here.

⁸up to some equivalence relation

We also naturally would like Z to send the identity cobordism $Y \times I : Y \rightarrow Y$ to the identity map.

Functoriality implies that if $C_{01} : Y_0 \rightarrow Y_1$ and $C_{12} : Y_1 \rightarrow Y_2$ are cobordisms, then

$$Z(C_{12} \circ C_{01}) = Z(C_{12}) \circ Z(C_{01}) : Z(Y_0) \rightarrow Z(Y_2).$$


Make sure you know what *category* and *functor* mean. Things will get confusing from here on out if these terms aren't clear.

Remark 3.2.2. To learn more about physical origins of the term *topological quantum field theory*, take a look at the [ncatlab page](#).

Atiyah's *Topology quantum field theories* [Ati88] provides a set of precise axioms for $(n+1)$ -TQFTs and surveys some prominent examples.

Remark 3.2.3. Since any (orientable) manifold only has two orientations, the category of cobordisms is *pivotal*, i.e. $A \cong (A^*)^*$ for any object A .

- (1) Thus the target category of the functor needs to be pivotal as well; this is why we require *finitely generated* R -modules.
- (2) Since our categories are pivotal we can identify any cobordism $C : K \rightarrow L$ with various 'bent' versions of C , shown in the schematics below: (add: discussion of boundary orientation?)



We thus get equivalences

$$\mathrm{Hom}(K, L) \cong \mathrm{Hom}(K \otimes \bar{L}, \emptyset) \cong \mathrm{Hom}(\emptyset, \bar{K} \otimes L).$$

- (3) We can view the equivalence

$$\mathrm{Hom}(Z(K) \otimes Z(\bar{L}), Z(\emptyset)) \cong \mathrm{Hom}(Z(K), Z(L))$$

as an instance of the Tensor-Hom Adjunction:

$$\mathrm{Hom}(M \otimes N, P) \cong \mathrm{Hom}(M, \mathrm{Hom}(N, P))$$

by setting $M = Z(K)$, $N = Z(\bar{L}) = Z(L)^*$, and $P = Z(\emptyset) = R$. (Note that $\mathrm{Hom}(N, P) = \mathrm{Hom}(Z(L)^*, R) = (Z(L)^*)^* \cong Z(L)$.)

3.2.2. Bar-Natan's dotted cobordism category and TQFT. In this section, we will focus on one particular TQFT, which we will denote $\mathcal{F}_{\mathrm{BN}}$. This material is adapted from [BN05].

First, we need to understand the source category.

Definition 3.2.4. A small category \mathcal{C} is *preadditive* (a.k.a. *Mod \mathbb{Z} -enriched*) if for $X, Y \in \mathrm{Ob}(\mathcal{C})$, $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ is an abelian group (i.e. a \mathbb{Z} -module) under composition, and this composition is bilinear (under the action of \mathbb{Z}).

Remark 3.2.5. A category is *additive* if additionally, we have finite co-products. We will boost up to additive categories later in §3.4.

Definition 3.2.6. The preadditive category \mathcal{TL}_0 is defined as follows:

- Objects are closed 1-manifolds with finitely many components embedded smoothly in the plane \mathbb{R}^2 ; we call these *planar circles*. Notice that we are *not* identifying isotopic embeddings, but we *do* ignore parametrization.
- Morphisms are finite sums of *dotted cobordisms*, or smooth surfaces embedded in $\mathbb{R}^2 \times I$,
 - with boundary only in $\mathbb{R}^2 \times \partial I$,
 - possibly decorated with a finite number of dots,
 - up to boundary-preserving isotopy, and
 - subject to the following relations:

(1)

The last relation is often referred to (gruesomely) as the *neck-cutting* relation.

The target category will be bigraded \mathbb{Z} -modules, which we will denote as $gg\text{Mod}_{\mathbb{Z}}$ when we want to emphasize the bigrading, or just $\text{Mod}_{\mathbb{Z}}$ for short. The most important module for us will be the one associated to a single circle in the plane, which we previously called V . We will now rename this module and view it as $V \cong \mathcal{A}\{1\}$, where

$$\mathcal{A} = \mathbb{Z}[X]/(X^2)$$

with $\deg_q(X) = -2$.

Definition 3.2.7. The functor $\mathcal{F}_{\text{BN}} : \mathcal{TL}_0 \rightarrow gg\text{Mod}_{\mathbb{Z}}$ assigns

- to each collection of k planar circles the module $(\mathcal{A}\{1\})^{\otimes k}$ (with $\emptyset \leadsto \mathbb{Z}$)
- to each dotted cobordism a linear map determined by the following assignments:

$$\begin{array}{c} \text{cylinder} \\ \hline \end{array} = \text{id} \qquad \begin{array}{c} \text{cylinder with dot} \\ \hline \end{array} = \bullet X$$

$$\begin{array}{c} \text{cup} \\ \hline \end{array} = \begin{array}{ccc} \iota : \mathbb{Z} & \longrightarrow & \mathcal{A}\{1\} \\ 1 & \longmapsto & 1 \end{array}$$

$$\begin{array}{c} \text{cap} \\ \hline \end{array} = \begin{array}{ccc} \varepsilon : \mathcal{A}\{1\} & \longrightarrow & \mathbb{Z} \\ 1 & \longmapsto & 0 \\ X & \longmapsto & 1 \end{array}$$

Remark 3.2.8. Notice that Definition 3.2.7 is really a Definition-Theorem, since one needs to check that it really is a functor. What does one need to check to make sure that \mathcal{F}_{BN} really is a functor?

Let C_1 and C_2 be cobordisms $\emptyset \rightarrow P$, where P is a collection of planar circles. Then there is a \mathbb{Z} -valued pairing

$$\langle C_1, C_2 \rangle = \mathcal{F}_{\text{BN}}(C_1 \cup_{\partial} \bar{C}_2) \quad (= \mathcal{F}_{\text{BN}}(\bar{C}_1 \cup_{\partial} C_2))$$

because $C_1 \cup_{\partial} \bar{C}_2$ is a closed 2-manifold, and the relations in \mathcal{TL}_0 allow us to evaluate closed surfaces.

In the following exercise, we will use this pairing to recover the linear maps m and Δ from the previous section.

Exercise 3.2.9. (*Not eligible for HW submission; we will complete these in lecture.*)

- In Definition 3.2.7, why did we not need to specify how \mathcal{F}_{BN} assigns higher-genus cobordisms?
- We can identify the distinguished \mathbb{Z} -module generators 1 and X in $\mathcal{A}\{1\}$ with a cup and a dotted cup, respectively. Why does this make sense?
- What is the dual cobordism to the cup? ...the dotted cup?
- Using this basis for $\mathcal{A}\{1\}$ and the pure tensor basis for $\mathcal{A}\{1\}^{\otimes 2}$, determine the matrix (i.e. chart) associated to the merge and split cobordisms by using the pairing.

Aside 3.2.10. (Computing linear maps using pairings)

Here is the linear algebra analogue to the method we are using above to compute the maps m and Δ .

Let $M : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a matrix with respect to the standard basis vectors $\{e_i\}$. We can view M as a pairing, by setting

$$\langle v, w \rangle_M = w^\top M v \in \mathbb{R}.$$

The entries of the matrix are determined by the pairing on basis vectors:

$$M_{ij} = e_i^\top M e_j$$

(Recall i is the row and j is the column.) This is because M_{ij} is the coefficient of e_i in the image vector $M e_j$.

In our setting, we are actually taking one additional step, which is to identify the column vectors v and w with the linear maps $\mathbb{R} \rightarrow \mathbb{R}^n$ that they represent, as $n \times 1$ matrices. (The ‘dual’ vector w^\top is a row vector that represents a linear map $\mathbb{R}^n \rightarrow \mathbb{R}$.)

So a ‘closed surface’ in our setting corresponds to the composition of maps

$$\mathbb{R} \xrightarrow{v} \mathbb{R}^n \xrightarrow{M} \mathbb{R}^n \xrightarrow{w^\top} \mathbb{R},$$

and setting v and w to be basis vectors allows us to compute the entries of M .

(Note that any linear map $\mathbb{R} \rightarrow \mathbb{R}$ is necessarily of the form $\cdot r$ for some $r \in \mathbb{R}$. We are implicitly using $\text{Hom}_{\mathbb{R}}(\mathbb{R}, \mathbb{R}) \cong \mathbb{R}$.)

Observe that even though we defined the target of \mathcal{F}_{BN} to be $gg\text{Mod}_{\mathbb{Z}}$, everything is happening at homological grading 0 at the moment; we don’t yet have chain complexes! However, the quantum degree of morphisms can be determined in the source category \mathcal{TL}_0 , as you’ll discover in the next exercise.

Using Morse theory, we can show that every cobordism C between finite collections of planar circles $P \rightarrow P'$ can be isotoped so that the critical points occur at distinct times $t \in I$. We can slice up the cobordism into pieces

$$C = C_m \circ C_{m-1} \circ \cdots \circ C_2 \circ C_1$$

where each C_i is a disjoint union of some identity cylinders and one of the following four *elementary cobordisms*:


- cup (ι)
- cap (ε)
- merge (m)
- split (Δ)

Exercise 3.2.11. (Important) Prove that for any cobordism $C : P \rightarrow P'$, the bidegree of the associated linear map is

$$\mathrm{gr}(\mathcal{F}_{\mathrm{BN}}(C)) = (0, \chi(C)),$$

where $\chi(C)$ is the Euler characteristic of the surface C .


Remark 3.2.12. Bar-Natan's more general cobordism category does not include dots as decorations. The objects are the same as in \mathcal{TL}_0 , but morphisms are subject to different relations:



These relations are called the S (sphere), T (torus), and $4Tu$ (four tubes) relations.

We will sometimes work with this category, but for now have chosen to start with the dotted category because elementary calculations are easier there.

Exercise 3.2.13. (a) Use the T and $4Tu$ relations to show that the genus-2 orientable surface evaluates to 0:



(b) Use $4Tu$ to show the relation below:



Use this to explain why, morally speaking, ‘dot’ = ‘half of a handle’.

3.2.3. (Aside) Frobenius algebras and $(1+1)$ -TQFTs. We now discuss why we switched from using the \mathbb{Z} -module

$$V = \mathbb{Z}v_+ \oplus \mathbb{Z}v_-$$

to the underlying \mathbb{Z} -module of the \mathbb{Z} -algebra

$$\mathcal{A} = \mathbb{Z}[X]/(X^2).$$

See [Kho06] for a reference.

Definition 3.2.14. A *Frobenius system* is the data $(R, A, \varepsilon, \Delta)$ consisting of


- a commutative ground ring R ;
- an R -algebra A ; in particular:
 - There is a *unit* or inclusion map $\iota : R \rightarrow A$ that sends $1 \mapsto 1$.
 - A has a *multiplication* map $m : A \otimes_R A \rightarrow A$.
- a *comultiplication* map $\Delta : A \rightarrow A \otimes_R A$ that is both coassociative and cocommutative; and
- an R -module *counit* map $\varepsilon : A \rightarrow R$ such that

$$(\varepsilon \otimes \mathrm{id}) \circ \Delta = \mathrm{id}.$$

The algebra A is a *Frobenius algebra*; it is simultaneously both an algebra and a coalgebra, and the following relation holds:

$$(\text{id}_A \otimes m) \circ (\Delta \otimes \text{id}_A) = \Delta \circ m = (m \otimes \text{id}_A) \circ (\text{id}_A \otimes \Delta).$$

Remark 3.2.15. There are many equivalent definitions for the term ‘Frobenius algebra’. The definition we used above is the most topological:



As you might have guessed, there is a strong relationship between Frobenius algebras and TQFTs.

In our case, $\mathcal{A} = \mathbb{Z}[X]/(X^2)$ is a rank 2 Frobenius extension of the ground ring \mathbb{Z} . Rank 2 Frobenius systems yield $(1 + 1)$ -dimensional TQFTs, via the identifications below:

the $(1 + 1)$ -TQFT sends this...	... to this in the Frobenius system
\emptyset	R
S^1	A
cup	ι
merge	m
split	Δ
cap	ε

By modifying the Frobenius algebra, we can get many more flavors of Khovanov homology. For example, if we instead use $R = \mathbb{Z}$, $A = \mathbb{Z}[X]/(X^2 - 1)$, we would build a version of Khovanov homology that is not quantum-graded (because $X^2 = 1$ is not a graded equation). This version of Khovanov homology is called *Lee homology*, and will be discussed later in this course when we talk about topological applications of Khovanov homology.

Warning 3.2.16. Do not use the category \mathcal{TL}_0 from Definition 3.2.6 as the source of the TQFT functor for Lee homology! The category \mathcal{TL}_0 was specifically tailored to the version of Khovanov homology over $\mathcal{A} = \mathbb{Z}[X]/(X^2)$, where ‘two dots = 0’.

3.3. Bar-Natan’s tangle categories. We now return back to dotted cobordisms, but develop the theory for not just planar circles, but also planar tangles.

Just as a link is an embedding of a finite number of circles in \mathbb{R}^3 , a *tangle* is a proper embedding of a finite number of circles and arcs in a 3-ball B^3 . A *planar tangle* is a tangle embedded in a 2-disk D^2 .


Since we will be building complicated categories out of tangles, we want to be very concrete with our definition of tangle categories, and will instead use the following definition.

Definition 3.3.1. An (n, n) -tangle is a 1-manifold with $2n$ boundary components (a.k.a. endpoints), properly embedded in the thickened square $[0, 1] \times [0, 1] \times (-\frac{1}{2}, \frac{1}{2})$, with $2n$ endpoints located at


$$(2) \quad \left\{ \left(\frac{i}{n+1}, 0, 0 \right) \right\}_{i=1}^n \cup \left\{ \left(\frac{i}{n+1}, 1, 0 \right) \right\}_{i=1}^n.$$

An (n, n) -tangle *diagram* is a projection of an (n, n) -tangle to the unit square $[0, 1] \times [0, 1]$ where all singular points are transverse intersections (just as in link diagrams).


A *planar* (n, n) -tangle is an (n, n) -tangle embedded in the square $[0, 1] \times [0, 1] \times \{0\}$. In other words, a planar tangle is a crossingless projection of a (quite untangled) tangle.



Remark 3.3.2. A planar (n, n) -tangle with no closed components is a *crossingless matching*. These are the $\text{Catalan}(n)$ many generators of the Temperley-Lieb algebra $\text{TL}_n(\delta)$ over a field \mathbb{k} , whose composition \otimes is given by stacking squares vertically:



Closed components evaluate to a nonzero element $\delta \in \mathbb{k}$:



We are not thinking about the monoidal structure of these diagrams just yet. But perhaps you can see the analogue between this $(0+1)$ -dimensional cobordism category and the $(1+1)$ -dimensional cobordism categories we're working with.

Definition 3.3.3 (cf. Definition 3.2.6). The preadditive category \mathcal{TL}_n is defined as follows:

- Objects are planar (n, n) -tangles with finitely many components, embedded smoothly in the square I^2 with endpoints at the $2n$ points specified in (2), denoted \mathbf{p} .
- Morphisms are finite sums of dotted cobordisms properly embedded in $I^2 \times [0, 1]$,
 - with vertical boundary (i.e. on $\partial I^2 \times [0, 1]$) consisting only of $2n$ vertical line segments $\mathbf{p} \times [0, 1]$;
 - possibly decorated with a finite number of dots;
 - up to boundary-preserving isotopy;
 - subject to the same local relations (1) as in Definition 3.2.6.

So far, we've only defined a TQFT that can 'evaluate' closed components to a \mathbb{Z} -module. Since tangles in \mathcal{TL}_n will inevitably have non-closed components, i.e. arcs, we will not be specifying a functor to $\text{Mod}_{\mathbb{Z}}$ just yet. Instead, we will bring our algebraic tools to the topological categories \mathcal{TL}_n and do as much homological algebra as we can before passing through

any TQFT. Later on in the course, however, we may discuss how to define a 2-functor from the 2-category of (n points, (n, n) -tangles, (n, n) -tangle cobordisms) to an appropriately rich algebraic 2-category.

In particular, we can still treat \mathcal{TL}_n as a (bi)graded category, where the quantum degree of a dotted cobordism C is *defined* as

$$\deg_q(C) := \chi(C) - n = \chi(C) - \frac{1}{2}(\# \text{ vertical boundary components})$$

(and the homological degree is $\deg_h(C) = 0$).

3.4. Adding crossings : Boosting to chain complexes. The Kauffman bracket showed us how to take a tangle with crossings and express it in terms of planar tangles; the Khovanov bracket tells us that a tangle with crossings is really just a chain complex of planar tangles. In this section, we will boost our categories \mathcal{TL}_n to categories of chain complexes, so that we can capture tangles in general, not just planar tangles.

Throughout this section, we will be working in the more general setting of tangles; recall that \mathcal{TL}_0 is just a special case of the categories \mathcal{TL}_n . Also recall that these are all preadditive categories, by construction.

Definition 3.4.1 ([BN05], Definition 3.2). Let \mathcal{C} be a preadditive category. The *additive closure* of \mathcal{C} , denoted $\text{Mat}(\mathcal{C})$, is the additive category defined as follows:

- Objects are (formal) direct sums of objects of \mathcal{C} . We can represent these as column vectors whose entries are objects of \mathcal{C} .
- Morphisms are matrices of morphisms in \mathcal{C} , which are added, multiplied, and applied to the objects just as matrices are added, multiplied, and applied to vectors.

An additive category is preadditive, by definition (Remark 3.2.5). We can build a category of chain complexes over any preadditive category:

Definition 3.4.2 ([BN05], Definition 3.3). Let \mathcal{C} be a preadditive category. The \star *category of chain complexes* over \mathcal{C} for $\star \in \{b, -, +\}$, denoted $\text{Kom}^\star(\mathcal{C})$, is defined as follows:

- Objects are $\{\text{finite length, bounded above, bounded below}\}^9$ chain complexes of objects and morphisms in \mathcal{C} .
- Morphisms are chain maps between complexes.

Remark 3.4.3. For an example of a category of chain complexes over a preadditive but not additive category, take a look at $\text{Kom}(\mathcal{M}_{\mathbb{Z}[G]})$ in [ILM21, Remark 2.11].

Exercise 3.4.4. compute this composition of morphisms in Mat

We are now ready to define the *Bar-Natan categories*, which are the categories we land in just before applying a TQFT to an algebraic category like chain complex over \mathbb{Z} -modules.

Definition 3.4.5. The Bar-Natan category \mathcal{BN}_n^\star is $\text{Kom}^\star(\text{Mat}(\mathcal{TL}_n))$, for $\star \in \{b, -, +\}$. In practice, I will mostly drop the \star from the notation. The bounded category is a full subcategory of both the $-$ and $+$ categories. We will only be working with bounded complexes at the beginning of the course, so the distinction won't matter.

These categories are rich enough to capture the entirety of the information in Khovanov homology, without ever passing to rings and modules. For example, we can think of Khovanov homology as a functor from links and cobordisms to the Bar-Natan category \mathcal{BN}_0 . We may then choose a TQFT (or, equivalently, a Frobenius system) to apply to the resulting invariant, to obtain many different flavors of Khovanov homology. Moreover, we now have a “Khovanov homology for tangles.”

⁹If the notation feels counterintuitive, just remember that ‘bounded above’ complexes are ‘supported mostly in negative degrees’.

Remark 3.4.6. Since $\mathcal{A} = \text{Mat}(\mathcal{TL}_n)$ is an abelian category, the *homotopy category* $K^*(\mathcal{A})$ of chain complexes over \mathcal{A} (for $\star \in \{b, -, +\}$) is a *triangulated category*. Khovanov homology can be thought of a functor to a triangulated category $K^b(\text{Mat}(\mathcal{TL}_0))$, in which case the quasi-isomorphism class of $\text{Kh}(L)$ is the link invariant. The Kauffman bracket skein relation gives exact triangles in this category. For exercises involving long exact sequences in Khovanov homology, see [Tur16].

To demonstrate the use of these Bar-Natan categories, we will use these categories to prove the invariance of Khovanov homology under some Reidemeister moves. You may find the proofs using the undotted $(S, T, 4Tu)$ theory in [BN05]. We will instead use the dotted theory as part of our demonstration. But first, we need to introduce two important lemmas that are immensely helpful with both by-hand and computer-assisted computations.

3.5. Computational tools. Bar-Natan’s categories are also incredibly useful in computing Khovanov homology (by computer), because they allow for a ‘divide-and-conquer’ approach using tangles. Implemented algorithms typically scan a link diagram, simplifying the homological data at each step of the filtration of the diagram. The main tools used for simplification are *delooping* and *abstract Gaussian elimination*, which we discuss below. This section follows [BN07].

Remark 3.5.1. These algorithms have been immensely important to solving problems in low-dimensional topology. For example, Lisa Piccirillo’s proof that the Conway knot is not slice [Pic20] involves computing the s invariant (see §4.3) from the Khovanov homology of a knot with *a lot (like around 40; I haven’t counted carefully)* of crossings, which is effectively impossible by naive computation.

Let us first write down a concrete formula for the degree of a morphism in \mathcal{TL}_n .

Warmup 3.5.2. As a sanity check, let’s answer some warmup questions. See Notation 2.5.5 for our conventions on shift functors.

- (a) Let A be a graded R -module, with shift functor $[\cdot]$. The set of degree-preserving maps $A[i] \rightarrow A[j]$ correspond to elements of $\text{Hom}^k(A, A)$ for some $k \in \mathbb{Z}$. What is k ? **Answer:** $k = i - j$
- (b) Suppose $\varphi \in \text{Hom}^0(A, A)$. What is the degree of the map induced by φ from $A[i] \rightarrow A[j]$? **Answer:** $j - i$
- (c) Now suppose $\psi \in \text{Hom}^\ell(A, A)$. What is the degree of the map induced by ψ from $A[i] \rightarrow A[j]$? **Answer:** $j - i + \deg(\psi)$

Lemma 3.5.3. Let F be a (possibly dotted) cobordism between planar tangles $T, T' \in \mathcal{TL}_n$. The degree of the morphism

$$F : T\{i\} \rightarrow T'\{j\}$$

is

$$(3) \quad \deg_q(F) = j - i + \chi(F) - n$$

(where n is $\frac{1}{2}$ the number of tangle endpoints, or vertical boundary components).

Verify that this makes sense to you!

From now on, we will treat diagrams and cobordism drawings as the same things as the objects and morphisms they represent in \mathcal{BN} .

Lemma 3.5.4 ([BN07], Lemma 4.1). (Delooping) \bigcirc is chain homotopy equivalent to $\emptyset\{1\} \oplus \emptyset\{-1\}$ via the chain homotopy equivalences

$$\begin{array}{ccc}
 & F = \begin{bmatrix} \text{cup} \\ \text{cap} \end{bmatrix} & \\
 \bigcirc & \xrightarrow{\quad} & \begin{bmatrix} \emptyset \{1\} \\ \emptyset \{-1\} \end{bmatrix} \\
 & \xleftarrow{\quad} & \\
 & G = \begin{bmatrix} \text{cup} & \text{cup} \end{bmatrix} &
 \end{array}$$

Proof. We leave it to the reader to check that F and G are indeed degree-preserving (use Lemma 3.5.3). It suffices to check that $G \circ F \simeq \text{id}_{\bigcirc}$ and $F \circ G \simeq \text{id}_{\emptyset\{1\} \oplus \emptyset\{-1\}}$. Indeed,

$$F \circ G = \begin{bmatrix} \text{cup} \\ \text{cap} \end{bmatrix} \begin{bmatrix} \text{cup} & \text{cup} \end{bmatrix} = \begin{bmatrix} \text{cup} & \text{cup} \\ \text{cap} & \text{cap} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

by the sphere, dotted sphere, and two dots relations, and

$$G \circ F = \begin{bmatrix} \text{cup} & \text{cup} \end{bmatrix} \begin{bmatrix} \text{cup} \\ \text{cap} \end{bmatrix} = \text{cup} + \text{cup} = \text{cylinder}$$

by the neck-cutting relation. (No nontrivial homotopies were needed; these compositions are identity “on the nose”.) \square

Delooping essentially allows us to replace any flat tangle containing a closed component with two (quantum-shifted) copies of that flat tangle with that circle removed.

Our next tool is an abstract version of Gaussian elimination. You can ponder for yourself why this ‘is’ Gaussian elimination. Just remember that, in \mathbb{Q} , any nonzero number is a unit. Row operations are just changes of basis on the target of the linear map. Similarly, column operations are just changes of the basis on the source of the linear map.

Lemma 3.5.5 ([BN07], Lemma 4.2). (Gaussian elimination) Let \mathcal{C} be a pre-additive (add: check why BN said additive instead) category. Suppose in $\text{Kom}(\text{Mat}(\mathcal{C}))$ there is a chain complex segment

$$\dots \xrightarrow{\begin{pmatrix} \alpha \\ \beta \end{pmatrix}} \begin{bmatrix} A \\ B \end{bmatrix} \xrightarrow{\begin{pmatrix} e & g \\ f & h \end{pmatrix}} \begin{bmatrix} C \\ D \end{bmatrix} \xrightarrow{\begin{pmatrix} \gamma & \delta \end{pmatrix}} \dots$$

where e is an isomorphism. Then the chain complex is homotopy equivalent to

$$\dots \xrightarrow{\begin{pmatrix} \beta \end{pmatrix}} [B] \xrightarrow{\begin{pmatrix} h - fe^{-1}g \end{pmatrix}} [D] \xrightarrow{\begin{pmatrix} \delta \end{pmatrix}} \dots$$

Note that

- A, B, C, D are objects in $\text{Mat}(\mathcal{C})$
- $e, f, g, h, \alpha, \beta, \gamma, \delta$ are morphisms in $\text{Mat}(\mathcal{C})$, i.e. these are matrices.

The main idea is that, by row and column operations, we are able to choose a basis so that the first chain complex is the direct sum of the second chain complex and an acyclic complex


$$0 \rightarrow A \xrightarrow{e} C \rightarrow 0.$$

See Bar-Natan’s proof for full details.

Corollary 3.5.6. (Cancellation lemma)¹⁰ Suppose (C, d) is a chain complex freely generated by a distinguished set of generators \mathcal{G} , and we draw it using dots and arrows. For $x, y \in \mathcal{G}$, let $d(x, y)$ denote the coefficient of y in $d(x)$.


Suppose there is an isomorphism arrow $a \xrightarrow{e} c$ between distinguished basis elements $a, c \in \mathcal{G}$, i.e. the coefficient of the arrow is a unit in the ground ring R . Then (C, d) is chain homotopy equivalent to a ‘smaller’ chain complex (C', d') where C' is generated by $\mathcal{G} - \{a, c\}$, and for any $b \in \mathcal{G}$,

$$d'(b) = d(b) - d(b, c)d(a).$$



The new arrows ($fe^{-1}g$ in the figure above) are called *zigzags* for obvious reasons.

Remark 3.5.7. Cancellation provides an especially fast way to compute Khovanov homology over \mathbb{F}_2 . Here is what my calculation for the Khovanov homology of the Hopf link looks like:



¹⁰See [BP10, Lemma 4.1], which directs you to [Ras03, Lemma 5.1], which directs you to [Flo89].

3.6. Reidemeister invariance of Khovanov homology. Bar-Natan's Reidemeister invariance proofs use the more general $S, T, 4Tu$ categories, and gives explicit homotopies where needed.

Here, we will stay with the \mathcal{TL}_n categories that we defined and prove invariance under some (easy) Reidemeister moves, and make use of some facts from homological algebra.

Definition 3.6.1. Let $\mathcal{C}, \mathcal{C}'$ be chain complexes. Let C_n denote the n -th chain group of \mathcal{C} , and let C'_n be defined analogously.

- \mathcal{C}' is a *subcomplex* of \mathcal{C} (written $\mathcal{C}' \subseteq \mathcal{C}$) if each C'_n is a submodule of C_n and the differential on \mathcal{C}' is the restriction of the differential on \mathcal{C} .
- If $\mathcal{C}' \subseteq \mathcal{C}$, then the *quotient complex* \mathcal{C}/\mathcal{C}' has chain groups C_n/C'_n , and the differential is induced by the differential on \mathcal{C} .

Lemma 3.6.2. Let


$$0 \rightarrow \mathcal{C}' \rightarrow \mathcal{C} \rightarrow \mathcal{C}'' \rightarrow 0$$

be an exact sequence of chain complexes, i.e. $\mathcal{C}'' \cong \mathcal{C}/\mathcal{C}'$.

- (1) If $\mathcal{C}' \simeq 0$, then $\mathcal{C} \simeq \mathcal{C}''$.
- (2) If $\mathcal{C}'' \simeq 0$, then $\mathcal{C} \simeq \mathcal{C}'$.


We now show R1 invariance of any flavor of Khovanov homology that factors through the dotted cobordism categories we defined.

Example 3.6.3. Consider the Reidemeister move involving a twist with a negative crossing:




There is only one strand in this local picture, so this crossing will be negative regardless of how you orient the strand.

Using the Khovanov bracket, we resolve the crossing to obtain a two-term chain complex in $\text{Kom}(\text{Mat}(\mathcal{TL}_1))$ representing the twist, and deloop the resolution on the right:



After delooping, our homotopic (actually isomorphic) complex is:




Notice that the downward arrow is an isomorphism in $\text{Mat}(\mathcal{TL}_1)$. The quotient of the highlighted subcomplex is therefore nullhomotopic.

Using Lemma 3.6.2, we conclude that the complex representing the negative R1 twist is chain homotopy equivalent to the highlighted complex.

Remark 3.6.4. If you close up the $(1, 1)$ -tangle in the example above, then we can see exactly why the downward arrow is an isomorphism (red below), and the upward arrow (green below) is not:

After delooping, our complex looks like this:



Notice that the red arrows are isomorphisms. After performing Gaussian elimination on these red arrows, we obtain the homotopy equivalent complex highlighted in yellow.

R3 is always a more complicated move to deal with, because

- there are 3 crossings, so one needs to compare two cubes with eight vertices each, and
- there are three strands, and therefore many possible orientations

Nevertheless, Bar-Natan's proof fits beautifully on just one page; see [BN05, Figure 9] to learn what a 'monkey saddle' is.

- Exercise 3.6.7.** (1) Prove invariance under the other R1 move, where the twist has a positive crossing.
 (2) Prove invariance under the other R2 move, where the strands are antiparallel.

3.7. (Projective) Functoriality. We are now equipped to fully define a functor that allows us to study links and cobordisms by replacing them with chain complexes and chain maps.

Once again, we begin by carefully defining our domain categories.

Definition 3.7.1. The category **Link** is defined as follows:

- Objects are smooth links in \mathbb{R}^3 ¹¹
- Morphisms are cobordisms between links, modulo isotopy rel boundary.

At first glance, **Link** would be the category we would want to define our functor out of. However, recall that we don't actually compute Khovanov homology directly from links; we actually use link diagrams. So, we need to define an intermediate *diagrammatic category* that is equivalent to **Link**.

Definition 3.7.2. The category **LinkDiag** is defined as follows:

- Objects are smooth link diagrams drawn in \mathbb{R}^2 again, NOT up to isotopy!
- Morphisms are *movies* between link diagrams, modulo *movie moves*.

Movies and *movie moves* need to be discussed carefully, analogously to how we defined *link diagrams* and *Reidemeister moves*.

Definition 3.7.3. A *movie* is a finite composition of the following 'movie clips':

- planar isotopy
- Reidemeister moves
- Morse moves: birth of a circle, death of a circle, merging of two circles, splitting of one circle into two

If $F \subset \mathbb{R}^3 \times I$ is a cobordism from $L_0 \subset \mathbb{R}^3 \times \{0\}$ to $L_1 \subset \mathbb{R}^3 \times \{1\}$, then an associated movie M can be thought of as a collection of movie 'frames' $\{M_t \mid t \in [0, 1]\}$ where M_t is a diagrammatic projection of the 'slice' of

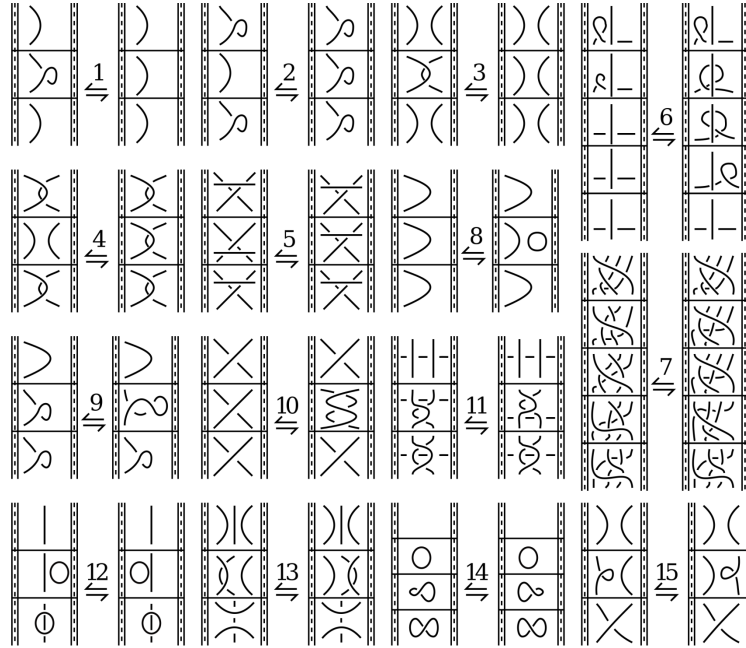
¹¹or S^3 , if preferred, but see Remark 3.7.7

the cobordism at time t , $F \cap (\mathbb{R}^3 \times \{t\})$. Indeed, the four Morse moves correspond to cup, cap, merge, and split cobordisms, respectively.

Just as we required link diagrams to only have transverse intersections, and for crossings to not be on top of each other, our definition for ‘movie’ ensures that our *critical frames* (i.e. frames where the projection $\mathbb{R}^3 \times \{t\} \rightarrow \mathbb{R}^2$ is not a link diagram) are isolated. Reidemeister moves and Morse moves are basically ‘before and after’ pictures of the process of passing through a critical frame.

In the category **Link**, morphisms are considered up to isotopy rel boundary. Just as Reidemeister proved that any diagram isotopy can be described as a finite composition of 3 local Reidemeister moves (and planar isotopy), Carter and Saito showed that any isotopy rel boundary of a cobordism can be captured as a finite composition of 15 local movie moves (and time-preserving isotopy), which are now known as the *Carter-Saito movie moves* [CS93].

Here is a figure ripped from [?]:



Similar definitions can also be made for tangles:

Definition 3.7.4. The category **Tang_n** is defined as follows:

- Objects are smooth (n, n) -tangles in $[0, 1]^2$ with boundary at the $2n$ points \mathbf{p} as in Definition 3.3.3.
- Morphisms are tangle cobordisms whose vertical boundary (i.e. the boundary in $\partial[0, 1]^2 \times I$) consists of the $2n$ line segments $\mathbf{p} \times I$, up to isotopy rel boundary.

The Carter-Saito movie moves are local, and so the definition of the diagrammatic category is essentially the same:

Definition 3.7.5. The category **TangDiag_n** is defined as follows:

- Objects are smooth (n, n) -tangle diagrams drawn in $[0, 1]^2$.
- Morphisms are *movies* between tangle diagrams that preserve the boundary points \mathbf{p} , modulo *movie moves*.

The punchline is that the diagrammatic categories are sufficient for capturing all the information in the topological categories.

Theorem 3.7.6. There is an equivalence of categories between **Link** and **LinkDiag** (and similarly between **Tang_n** and **TangDiag_n**).

We can attribute this theorem to the combined work of Reidemeister and Carter–Saito **but I should check this**. The proof is beyond the scope of this course, but we will make use of these equivalences all the time.

Remark 3.7.7. If you want to work with links in S^3 instead, there are more diagrammatic moves to check. In particular, we need to include the *sweep-around* move, which equates the sweep-around movie (swinging a strand of the diagram around the 2-sphere through the point at infinity) with the identity movie. See [MWW22].

We now know what it means for a link invariant to be *functorial*: it is a functor from the category **Link** to its target category.

For our purposes, we will define the *Khovanov functor*

$$\mathcal{F}_{\text{Kh}} : \mathbf{LinkDiag} \rightarrow gg\text{Mod}_{\mathbb{Z}}$$

to be the composition of the functors

$$\mathbf{LinkDiag} \rightarrow \text{Kom}(\text{Mat}(\mathcal{TL}_0)) \rightarrow gg\text{Mod}_{\mathbb{Z}}.$$

The first functor is the unfortunately unnamed (projective) functor that Bar-Natan gives us in [BN05]¹² The second is the TQFT associated to the Frobenius system $(\mathbb{Z}, \mathcal{A} = \mathbb{Z}[X]/(X^2), \iota, m, \Delta, \varepsilon)$ where the morphisms are given by

$$\begin{aligned} \iota : \mathbb{Z} &\rightarrow \mathcal{A} \\ 1 &\mapsto 1 \\ m : \mathcal{A} \otimes \mathcal{A} &\rightarrow \mathcal{A} \\ 1 \otimes 1 &\mapsto 1 \\ X \otimes 1, 1 \otimes X &\mapsto X \\ X \otimes X &\mapsto 0 \\ \Delta : \mathcal{A} &\rightarrow \mathcal{A} \otimes \mathcal{A} \\ 1 &\mapsto X \otimes 1 + 1 \otimes X \\ X &\mapsto X \otimes X \\ \varepsilon : \mathcal{A} &\rightarrow \mathbb{Z} \\ 1 &\mapsto 0 \\ X &\mapsto 1. \end{aligned}$$

Example 3.7.8. Khovanov homology over \mathbb{F}_2 is functorial; it defines a functor

$$\mathbf{LinkDiag} \xrightarrow{\mathcal{F}_{\text{Kh}}} gg\text{Vect}_{\mathbb{F}}$$

which, when precomposed with the equivalence

$$\mathbf{Link} \xrightarrow{\cong} \mathbf{LinkDiag},$$

defines a functor from the category of links to the category of bigraded vector spaces.


(add: Add MR citations) Jacobsson [?], Khovanov [?], and Bar-Natan [BN05, Section 4.3] each showed that Khovanov homology over \mathbb{Z} is *projectively* functorial, meaning that it's functorial up to a sign. In other words, suppose you take two movies M and M' representing isotopic (rel boundary) cobordisms C and C' . Then the corresponding morphisms in $gg\text{Mod}_{\mathbb{Z}}$ satisfy

$$\text{Kh}(M') = \pm \text{Kh}(M).$$

This sign discrepancy can be fixed, and it has been by a whole host of authors [Cap08, CMW09, Bla10, San21, Vog20, ETW18, BHPW23]. However, projective functoriality is enough for many, many important applications of Khovanov homology, which we will see in the next section.

Exercise 3.7.9. (Highly recommended) The following is a version of Movie Move 14 ('MM14'):

¹²We can't call it \mathcal{F}_{BN} because 'Bar-Natan homology' means something else, which we will discuss later.



Compute the induced map on Khovanov homology for each movie, by using the Reidemeister 1 chain maps given in [BN05, Figure 5]. Show that the resulting morphisms have opposite signs.

You do not need to pass through the TQFT, but you are welcome to; the TQFT is a true (not projective) functor.

Beware: Bar-Natan's cobordisms flow *downward* with time. You can tell by looking at the domain and target objects of the cobordisms.

You can also quickly convince yourself that these movies, when read *backwards*, actually *do* yield morphisms with the same sign.

Maybe add section on general planar algebras, operads

4. APPLICATIONS OF KHOVANOV HOMOLOGY

As a functorial invariant for links in the 3-sphere and cobordisms in the 4-ball, Khovanov homology is a natural tool to study the relationship between links in the context of the surfaces they bound. In this section, we survey some of these applications. We start with some additional background on surfaces properly embedded in B^4 . Once again, everything is smooth.

4.1. Surfaces in B^4 . We have already discussed cobordisms between links in S^3 . In this section, we will use F to denote a cobordism (because they're 2-dimensional, like 'faces'). We reserve C for *concordances*, which are cobordisms that are diffeomorphic to cylinders. *I.e. without the context of their embedding, they are cylinders.* We will use the more precise term *annulus* instead of 'cylinder'.

Definition 4.1.1. Let K_0, K_1 be **knots** in S^3 . A *concordance* C from K_0 to K_1 is an oriented cobordism such that $C \cong S^1 \times I$.

Equivalently, $C : K_0 \rightarrow K_1$ is a concordance if it is a (smooth, oriented) connected cobordism with $\chi(C) = 0$.

If such a C exists, then we say K_0 and K_1 are *concordant*: $K_0 \sim K_1$. This is an equivalence relation, and the equivalence classes are called *concordance classes*.

In fact, we can turn the set of knots into a group (!) by modding out by concordance:

Definition 4.1.2. The *smooth knot concordance group* \mathcal{C} is the group where

- the elements are the concordance classes knots in S^3 ;
- the binary operation is induced by $\#$ (connected sum);
- the identity element is the class of *slice knots*, or knots that are concordant to the unknot;
- inverses are given by mirroring.

Remark 4.1.3. If you take a knife to a B^4 and cut off a slice, the cut you make is a slice disk. I haven't checked if this is the historical origin of the text.


In general, a surface F properly embedded in B^4 whose boundary is $\partial F = K$ is called a *slice surface* for K .

Remark 4.1.4. We can equivalently say that a knot $K \subset S^3$ is *slice* if it bounds a disk D in B^4 : if we arrange D so that it is in Morse position with respect to the radial function on B^4 , then the boundary of a neighborhood of its lowest 0-handle is an unknot. In other words, an annulus is just a punctured disk.

In this case, we say that D is a *slice disk* for K .

Again, this is actually a definition-theorem, and we will prove the theorem after seeing a quick example.

Example 4.1.5. Consider the right-handed trefoil K :



If we claim that the left-handed trefoil $m(K)$ represents the inverse concordance class, then we must show that $K \# m(K)$ is concordant to the unknot U .

It suffices to show that $K \# m(K)$ bounds a disk D embedded in B^4 . You can see a projection of a slice disk in S^3 in the following picture:



The disk is immersed in S^3 , and the only intersections are of the following form, called *ribbon intersections*:



The interior of the horizontal sheet can be pushed deeper into B^4 , so that the slice disk in B^4 has no self-intersection.

To verify that the concordance group really is a group, we need to check that

- (1) $\#$ is well-defined on equivalence classes
- (2) $\#$ is associative
- (3) the equivalence class $[U]$ acts as the identity
- (4) for any K , $K \# m(K)$ is slice.


You can convince yourself that the binary operation $\#$ on the set of knots is both associative and commutative; if you want to think about this in more detail, see [Ada04]. Once we show that $\#$ is well-defined, showing that $[U]$ is the identity element is also easy.

It remains to show that $\#$ is well-defined, and that $K \# m(K)$ is slice. Both proofs use standard topological arguments.

Claim 4.1.6. $\#$ is a well-defined binary operation on the set of concordance classes.


Proof. To see that $\#$ is a well-defined binary operation on concordance classes, consider knots $K, K', J \subset S^3$, where $K \sim K'$. Then there is some concordance $C : K \rightarrow K'$, a cobordism in $S^3 \times [0, 1]$. Pick basepoints $p \in K$, $p' \in K'$ and isotope K' so that $p' = p$ (as points in S^3). Pick an arc $\gamma : [0, 1] \hookrightarrow C$ such that $\gamma(0) = p$ and $\gamma(1) = p'$.

Perform a boundary-preserving ambient isotopy to ‘straighten out’ γ ; that is, to arrange so that $\gamma(t) \in S^3 \times \{t\}$. (This is possible because the codimension of γ in $S^3 \times [0, 1]$ is 3, and codimension 3 submanifolds can always be unknotted.) We will now assume that C is in such a position so that γ is of the form $p \times [0, 1]$.



Now delete a small neighborhood of γ (i.e. $\nu(p) \times [0, 1] \cap C$) in C . Pick a point $q \in J$ and delete it, forming a $(1, 1)$ tangle $J - \nu(q)$ whose closure is J . Shrink this tangle so that it fits within $\nu(p)$.

Finally, glue $C - (\nu(p) \times [0, 1])$ with $(J - \nu(q)) \times [0, 1]$ to form a concordance from $K \# J$ to $K' \# J$.



□


Aside 4.1.7. *How much should one write for such a proof?* You might notice that, for example, the last sentence of the preceding proof is not super duper precise. However, the figure helps you understand the notation, and also the underlying concept is quite simple. Also, we were careful to define the category **LinkDiag** so that we can make diagrammatic arguments.

As a human who uses language, I don’t have a perfect answer for ‘how much detail to show’ – this comes from getting to know the common vocabulary, techniques, facts, and tricks used in the community you are writing for.

We also give a (more terse) proof sketch that $[m(K)] = [K]^{-1}$:

Claim 4.1.8. For any knot $K \subset S^3$, $K \# m(K)$ is slice.

Proof. The identity cobordism $K \rightarrow K$ is a concordance. Pick a point $p \in K$ and delete $p \times [0, 1]$ from the identity cobordism. The resulting surface can be properly embedded in B^4 and viewed as a slice disk for $K \# m(K)$.



□

Remark 4.1.9. One can also define a notion of concordance between links. However, the notion of a ‘link concordance group’ isn’t obvious, and is an area of active research.


4.2. Obstruction to ribbon concordance. In Example 4.1.5, we envisioned a slice disk by seeing that its projection to S^3 had only *ribbon singularities*. The following definition gives a more general definition of such phenomena.

Definition 4.2.1. Let $F : L_0 \rightarrow L_1$ be a cobordism embedded in $S^3 \times I$ such that the second coordinate gives a Morse function. We say F is *ribbon* if, with respect to the Morse handle decomposition, F contains only 0- and 1-handles.

A ribbon cobordism that $D : \emptyset \rightarrow K$ that is a disk is called a *ribbon disk*, and is a special case of a slice disk.


A knot K that bounds a ribbon disk is called a *ribbon knot*.

Example 4.2.2. Here is a movie for the ribbon disk for the square knot that we visualized in Example 4.1.5:



From the definition, the ‘ribbon concordant’ relation is not symmetric. If we flip a ribbon concordance containing a 0-handle upside-down (by reversing the Morse function), the upside-down concordance \bar{C} clearly has a 2-handle.

Exercise 4.2.3. Draw a movie for the ribbon disk for the Stevedore knot 6_1 implied in the diagram below:



Remark 4.2.4. Unfortunately, we naturally would want to say ‘ 6_1 is concordant to the unknot’, which is true, but not if we add the word ‘ribbon’. Mathematically it makes more sense to say U is ribbon concordant to 6_1 , and this is the language used in [LZ19]. Historically, some people defined ribbon surfaces by taking the descending Morse function and requiring only 1- and 2-handles, so be careful. For my sanity, I will always specify the direction of the cobordism as a morphism.

Remark 4.2.5. In terms of movies, a ribbon disk is a ‘happy movie’, where the only critical moments are births of circles and merges of circles. There are no scenes where circles split or die.

Conjecture 4.2.6. (Open: Slice-Ribbon Conjecture) All slice knots are ribbon.

It is clear that ribbon disks are slice, so any ribbon knot is a slice knot. However, not all slice *disks* are (isotopic to) ribbon disks (see [Aru Ray’s notes](#), Proposition 1.8). So, the conjecture posits that if K bounds a slice disk, it bounds a (potentially non-isotopic) ribbon disk.

Remark 4.2.7. We will see soon that there are knots that bound non-isotopic ribbon disks [HS21].

Levine–Zemke showed that we can use Khovanov homology to obstruct the existence of ribbon cobordisms between two knots:

Theorem 4.2.8 ([LZ19]). If $C : K_0 \rightarrow K_1$ is a ribbon concordance, then the morphism

$$\mathrm{Kh}(C) : \mathrm{Kh}(K_0) \rightarrow \mathrm{Kh}(K_1)$$

is injective, with left inverse $\mathrm{Kh}(\bar{C})$.

Here are some immediate corollaries that follow from basic algebra:

Corollary 4.2.9 ([LZ19]). Suppose $C : K_0 \rightarrow K_1$ is a ribbon concordance.

- (1) At any bigrading, $\mathrm{Kh}^{i,j}(K_0) \hookrightarrow \mathrm{Kh}^{i,j}(K_1)$ as a direct summand.
- (2) If additionally there is a ribbon concordance $C' : K_1 \rightarrow K_0$, then $\mathrm{Kh}(K_0) \cong \mathrm{Kh}(K_1)$.


[Seriously, prove these for yourself.](#) There are four more corollaries in the paper, if you’re interested.

The proof uses a topological lemma first appearing in Zemke’s [Zem19]. This paper sparked a whole series of papers proving similar or related results for other functorial link homology theories ([add: all citations](#)).

The proof of this lemma is embedded in the proof of Zemke’s main theorem:

Lemma 4.2.10 ([Zem19]). Let $C : K_0 \rightarrow K_1$ be a ribbon concordance, and let \bar{C} be the upside-down (and opposite orientation) concordance to C , a morphism $K_1 \rightarrow K_0$.

If a movie presentation for C has n births and n saddles (note that the Euler characteristic of a concordance is 0), then there is a movie presentation for $\bar{C} \circ C$ with n births, n merge saddles, n split saddles that are dual to the merge saddles, and n deaths.



In particular, the lemma tells us that $\bar{C} \circ C$ is isotopic to a cobordism $K_0 \rightarrow K_0$ that looks like the identity cobordism for K_0 with n spheres tubed on. Because of nontrivial knot theory for surfaces in dimension 4, this cobordism might not be isotopic to the identity cobordism, but nevertheless, Khovanov homology can't tell because of the neckcutting relation:

Proof of Theorem 4.2.8. Take C and \bar{C} as in Lemma 4.2.10. In \mathcal{TL}_0 , by neckcutting, we see that the morphism $\bar{C} \circ C$ is equal to the morphism with 2^n summands that are all of the form id_{K_0} disjoint union with n spheres, and all summands have just one dot somewhere. After deleting the summands containing an undotted sphere, the only remaining morphism is the one consisting of id_{K_0} and n dotted spheres, which evaluate to a coefficient of 1. Therefore $\bar{C} \circ C = \text{id}_{K_0}$ as morphisms in \mathcal{TL}_0 .

By (projective) functoriality of Kh , we have $\text{Kh}(\bar{C}) \circ \text{Kh}(C) = \text{Kh}(\bar{C} \circ C)$, and the remainder of the theorem follows. \square

4.3. Lee homology and the s invariant. The most important numerical invariant(s) that you can distill from Khovanov homology techniques is Rasmussen's s invariant, which was introduced in [Ras10]. It is a 'concordance homomorphism,' i.e. a homomorphism of groups

$$s : \mathcal{C} \rightarrow \mathbb{Z}.$$

Rasmussen showed that s gives a lower bound on the *slice genus* of a knot. The following definition defines two knot invariants:

Definition 4.3.1. A *Seifert surface* for a knot $K \subset S^3$ is an oriented surface embedded in S^3 whose boundary is K . The 3-ball genus or *Seifert genus* of a knot K , $g_3(K)$, is the minimal genus of a Seifert surface for K :


$$g_3(K) = \min\{g(F) \mid F \hookrightarrow S^3 \text{ with } \partial F = K\}$$

Analogously, the 4-ball genus or *slice genus* of a knot $K \in S^3$ is

$$g_4(K) = \min\{g(F) \mid F \hookrightarrow B^4 \text{ with } \partial F = K\},$$

the minimal genus of a slice surface for K .

Example 4.3.2. Here is a Seifert surface for the right-handed trefoil, otherwise known as the torus knot $T_{2,3}$:



Since any Seifert surface can be 'punched in' to B^4 to become a slice surface, for any knot K , we have $g_4(K) \leq g_3(K)$, and therefore s also gives a lower bound on the Seifert genus of K .

Rasmussen used the s invariant to give a combinatorial proof of the (topological) Milnor conjecture:

Theorem 4.3.3 ([KM93], [Ras10]). The slice genus of a torus knot $T(p, q)$ is

$$g_4(T(p, q)) = g_3(T(p, q)) = \frac{(p-1)(q-1)}{2}.$$

I bet the first Seifert surface you draw for $T(p, q)$ is the minimal genus (both g_4 and g_3 !) surface!

Another major application of Rasmussen's s invariant in 4D topology is Piccirillo's proof that the Conway knot is not slice [Pic20]; this was a long-standing conjecture until she proved it in less than 8 pages. This is an *Annals of Mathematics* paper. For those of you interested in 4-manifolds and using Kirby calculus, this is a potential final project idea.

The s invariant relies on *Lee homology*, a version of Khovanov homology introduced by E. S. Lee in her study of Khovanov homology of alternating knots [Lee05], very soon after Khovanov homology was first introduced. We now understand Lee homology as the version corresponding to the Frobenius algebra $\mathcal{A}' = \mathbb{Q}[X]/(X^2 - 1)$.

In this section, we will discuss Lee homology and Rasmussen's s invariant, as well various related ideas, including extremely useful algebraic constructions such as filtration spectral sequences and their relationship to torsion order. This will take a while, but these constructions were fundamental to the various applications of Khovanov homology in the following two decades.

4.3.1. Lee homology and quantum filtration. In this section we replace Khovanov's TQFT \mathcal{A} with Lee's, which we denote using primes (e.g. \mathcal{A}'), following Rasmussen's paper [Ras10]. In future sections of these notes we will likely use more specific notation.

Lee's TQFT comes from the Frobenius system

$$(\mathbb{Q}, \mathcal{A}' = \mathbb{Q}[X]/(X^2), \iota', m', \Delta', \varepsilon')$$

where the structure morphisms are given by

$$\begin{aligned} \iota' : \mathbb{Q} &\rightarrow \mathcal{A}' \\ 1 &\mapsto 1 \end{aligned}$$

$$\begin{aligned} m' : \mathcal{A}' \otimes \mathcal{A}' &\rightarrow \mathcal{A}' \\ 1 \otimes 1 &\mapsto 1 \\ X \otimes 1, 1 \otimes X &\mapsto X \\ X \otimes X &\mapsto \mathbf{1} \end{aligned}$$

$$\begin{aligned} \Delta' : \mathcal{A}' &\rightarrow \mathcal{A}' \otimes \mathcal{A}' \\ 1 &\mapsto X \otimes 1 + 1 \otimes X \\ X &\mapsto X \otimes X + \mathbf{1} \otimes \mathbf{1} \end{aligned}$$

$$\begin{aligned} \varepsilon' : \mathcal{A}' &\rightarrow \mathbb{Q} \\ 1 &\mapsto 0 \\ X &\mapsto 1. \end{aligned}$$

The most important thing to notice right now is that these maps are not grading preserving! After all, if you mod out by a non-quantum-homogeneous polynomial $X^2 - 1$, you will not get a graded theory. But since $\deg X^2 = 4$, you *do* get a $\mathbb{Z}/4\mathbb{Z}$ -quantum-graded theory, and this *will* be used in Rasmussen's proofs.

Lee originally described her differentials as a perturbation of Khovanov's. Indeed, we may write

$$d_{\text{Lee}} = d_{\text{Kh}} + \Phi$$

where Φ consists of the red terms in the m' and Δ' morphisms above.

Check by inspection that while d_{Kh} preserves quantum grading, Φ increases quantum grading by 4. So even though d_{Lee} does not preserve grading, the Lee chain complex

$$\text{CLee}(K) = (\text{CKh}(K), d_{\text{Kh}} + \Phi)$$

is *filtered*:

Definition 4.3.4. A chain complex (\mathcal{C}, d) is *filtered* if there exists a

A *filtration* of a chain complex (\mathcal{C}, d) is a sequence of subcomplexes F_\bullet of (\mathcal{C}, d)

$$\cdots \supseteq F_i \supseteq F_{i+1} \supseteq F_{i+2} \supseteq F_{i+3} \supseteq \cdots$$


such that

$$\bigcap_i F_i = \emptyset \quad \text{and} \quad \bigcup_i F_i = \mathcal{C}$$

and where $d(F_i) \subseteq F_i$.

A chain map $f : (\mathcal{C}, d) \rightarrow (\mathcal{C}', d')$ between filtered complexes $\mathcal{C} = \bigcup F_i$ and $\mathcal{C}' = \bigcup F'_i$ is a chain map that respects the filtration: $f(F_i) \subseteq F'_i$.

Example 4.3.5. Here is a cartoon of a filtered chain complex:



Warning 4.3.6. This is only one type of filtration; see a homological algebra book for a more general definition. We will in particular be using *finite length filtrations*, where only finitely many F_i are not \emptyset or \mathcal{C} .

The filtration on the Lee complex is very special, because it is defined by a *filtration grading* gr_q on a free module generated by a distinguished, gr_q -homogeneous basis.

Let $\text{Kg}(L)$ denote the set of pure tensors that generate the Khovanov chain group $\text{CKh}(L)$. Each chain is a finite sum $x = \sum_{i=1}^k c_i g_i$ where $g_i \in \text{Kg}(L)$, and $c_i \in \mathbb{Q}$. The filtration grading is defined for non-homogeneous chains by the following formula:

$$(4) \quad \text{gr}_q\left(\sum_{i=1}^k c_i g_i\right) = \min\{\text{gr}_q(g_i)\}_{i=1}^k.$$

Observe that, by this definition, for all $j \leq \text{gr}_q(x)$, we have $x \in F_j$; meanwhile, for $j \geq \text{gr}_q(x)$, $x \notin F_j$.

Remark 4.3.7. Filtered complexes give rise to *filtration spectral sequences*, which we will likely discuss later on. For now, just know that because the Khovanov differential is the grading-preserving piece of the Lee differential, there is a spectral sequence relating Khovanov homology to Lee homology.

4.3.2. Lee's basis and canonical generators. At first, the Lee complex might look *more* complicated than the Khovanov one, but this is only because we made an inconvenient choice of distinguished basis, because in the Khovanov TQFT, we wanted a homogeneous basis. Note that \mathcal{A}' is still isomorphic, as a non-graded vector space, to what we will still call $V = \mathbb{Q}v_+ + \mathbb{Q}v_-$, *even though we've switched to \mathbb{Q} coefficients, so we're technically working with $V \otimes_{\mathbb{Z}} \mathbb{Q}$.* Lee introduced a much better distinguished basis for her TQFT, where

$$\mathbf{a} = v_- + v_+ = 1 + X \quad \mathbf{b} = v_- - v_+ = 1 - X.$$

Notice that these are the factors of $X^2 - 1$; therefore, you already automatically know that $m'(\mathbf{a} \otimes \mathbf{b}) = m'(\mathbf{b} \otimes \mathbf{a}) = 0$, which is nice. Here are the other morphisms, in this new basis:

Notation 4.3.8.

$$\begin{aligned}
\iota' : \mathbb{Q} &\rightarrow \mathcal{A}' \\
1 &\mapsto \frac{\mathbf{a} - \mathbf{b}}{2} \\
m' : \mathcal{A}' \otimes \mathcal{A}' &\rightarrow \mathcal{A}' \\
\mathbf{a} \otimes \mathbf{a} &\mapsto 2\mathbf{a} \\
\mathbf{a} \otimes \mathbf{b}, \mathbf{b} \otimes \mathbf{a} &\mapsto 0 \\
\mathbf{b} \otimes \mathbf{b} &\mapsto -2\mathbf{b} \\
\Delta' : \mathcal{A}' &\rightarrow \mathcal{A}' \otimes \mathcal{A}' \\
\mathbf{a} &\mapsto \mathbf{a} \otimes \mathbf{a} \\
\mathbf{b} &\mapsto \mathbf{b} \otimes \mathbf{b} \\
\varepsilon' : \mathcal{A}' &\rightarrow \mathbb{Q} \\
\mathbf{a}, \mathbf{b} &\mapsto 1
\end{aligned}$$

You may wish to check that this indeed defines a Frobenius algebra by checking the required relations.

With a Frobenius algebra structure so simple (and with 20 years of reflection and hindsight), one would expect that the homology would be very easy to compute. Indeed, Lee showed that Lee homology is very simple:

Theorem 4.3.9 ([Lee05], Theorem 4.2). For a link L with ℓ components, there is an isomorphism of (ungraded) vector spaces $\text{Lee}(L) \cong (\mathbb{Q} \oplus \mathbb{Q})^\ell$.

Moreover, she knows the homology classes that generate the 2^ℓ -dimensional homology; these correspond to the 2^ℓ orientations of the link! These special homology classes are now known as *Lee's canonical generators*. [And are so, so, so important.](#)

Here is an algorithm used for obtaining Lee's canonical generators.

Algorithm 4.3.10. Suppose we are given a link L , an orientation o on L , and a diagram D of L . Fix a checkerboard coloring of the complement of D in the plane. (A standard convention is to leave the infinite region unshaded.)

Let D_o denote the oriented resolution: that is, resolve each crossing so that the orientations of the arcs are all preserved. The resulting circles have induced orientations. For each circle, draw a small dot slightly to the left of any point on the circle (based on the orientation of the circle).

- If the dot lands in a shaded region, label the circle \mathbf{a} .
- If the dot lands in an unshaded region, label the circle \mathbf{b} .

Let \mathfrak{s}_o denote the chain correspond to the above labeling of the resolution D_o .

Exercise 4.3.11. Prove that \mathfrak{s}_o is closed (i.e. in the kernel of the differential).


Lee's canonical generators are the set of homology classes

$$\{[\mathfrak{s}_o] \mid o \text{ is an orientation for } L\}.$$

In the remainder of this (subsub)section, we give a proof of Lee's Theorem 4.3.9 that is due to Bar-Natan–Morrison [BNM06]. [There are many versions of Lee's proof in the literature. They all involve similar ideas to those in the original proof.](#)

First of all, we need to set up the Bar-Natan category that works for Lee homology. In particular, we know that $X^2 = 1$, so now 'two dots' must evaluate to 1, rather than zero. Below are the Bar-Natan categories we will use. Beware that the notation is not standard, and in fact does not appear anywhere in the literature that I know of ([BNM06] calls these categories ' $\text{Cob}_1(\partial T)$ '). However, we use this primed notation to be consistent with the notation used in Rasmussen's s invariant paper [Ras10].

Definition 4.3.12. The preadditive categories \mathcal{TL}'_n for $n \in \mathbb{N} \cup \{0\}$ have the same objects as \mathcal{TL}_n . The morphisms are also cobordisms, but modulo these slightly modified relations:



Remark 4.3.13. If you wanted to, you could verify that these relations still hold. The dotted identity still means ‘ $\cdot X$ ’. What’s more interesting is coming up with a new decorated cobordism category that innately uses the $\{\mathbf{a}, \mathbf{b}\}$ basis for \mathcal{A}' . (So, define new kinds of dots.) What would your morphism relations look like, with these new decorations? [There is an answer in the literature, and major hints in the remainder of the section.](#) Also, I could have phrased this as an Exercise but I would not want to grade this problem!

We will need an additional algebraic construction; we follow the exposition in [BNM06, Section 3]. Recall that in any category \mathcal{C} , any endomorphism $p \in \text{Hom}_{\mathcal{C}}(X, X)$ satisfying $p^2 = p \circ p = p$ is called a *projection*. We say that p is *idempotent*. The *Karoubi envelope* of \mathcal{C} , denoted $\text{Kar}(\mathcal{C})$, is an enlargement of the category \mathcal{C} that includes all images of idempotent endomorphisms:

Definition 4.3.14. The *Karoubi envelope* (or *idempotent completion*) of \mathcal{C} , denoted $\text{Kar}(\mathcal{C})$, is the category where

- objects are pairs (X, p) where $X \in \text{Ob}(\mathcal{C})$, $p \in \text{Hom}_{\mathcal{C}}(X, X)$, and $p^2 = p$; and
- morphisms from (X, p) to (X', p') are the morphisms $f \in \text{Hom}_{\mathcal{C}}(X, X')$ that respect the projections: $f \circ p = p' \circ f$.

In other words, if you ever find a projection p in (the morphisms of) \mathcal{C} , then its image $\text{im}(p)$ is an object in $\text{Kar}(\mathcal{C})$. In particular, every object in \mathcal{C} also appears in $\text{Kar}(\mathcal{C})$, in the form (X, id_X) .

Example 4.3.15. The Karoubi envelope occurs in the wild! For example, let **Open** be the category consisting of

- objects: open subsets of Euclidean space (all \mathbb{R}^n for $n \in \mathbb{N}$)
- morphisms: smooth maps between these open sets.

The Karoubi envelope of **Open** turns out to be the category of all smooth (closed, finite-dimensional) manifolds **Man**. For example, while S^n is not an open subset of any \mathbb{R}^m , it is the image of the projection

$$p : U \rightarrow U$$

$$x \mapsto \frac{x}{|x|}$$

where U is the open subset $\mathbb{R}^n - \{0\} \subset \mathbb{R}^n$. This example comes from the ncatlab page for ‘Karoubi envelope,’ [nLa24].

Here are some more important algebraic facts that are not hard to prove.

Proposition 4.3.16. Let \mathcal{C} be an additive category.

- If $p \in \text{Hom}(X, X)$ is a projection, then so is $\text{id} - p \in \text{Hom}(X, X)$. Thus in $\text{Kar}(\mathcal{C})$, $X \cong \text{im}(p) \oplus \text{im}(\text{id} - p)$.
- Let A and B be chain complexes in $\text{Kom}(\mathcal{C})$. If A and B are chain homotopy equivalent in $\text{Kom}(\text{Kar}(\mathcal{C}))$, then they were also homotopy equivalent in $\text{Kom}(\mathcal{C})$.

Exercise 4.3.17. Prove Proposition 4.3.16.

Remark 4.3.18. Recall from, say, Mat167 that involutions are related to projections: if P is a projection, then $I - 2P$ is an involution. Geometrically, this involution reflects across the plane that P projects onto, because you (a point in \mathbb{R}^n) are being translated by a vector to the plane, and then you overshoot by traveling by that vector again.

Conversely, given an involution S , the *eigenprojections* of S are given by $\frac{I \pm S}{2}$. If P is an eigenprojection for S , then $SP = \lambda P$ where $\lambda = \pm 1$. These correspond to $\text{im}(P)$ and $\text{im}(P)^\perp$, because $P(I - P) = 0$.

For example, the projection $P = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ projects points in \mathbb{R}^2 to the x -axis. The projection $I - P$ projects the y -axis. The involution (reflection) $I - 2P$ reflects across the x -axis. This fixes the points on the x -axis ($\lambda = 1$) but reflects the points on the y -axis ($\lambda = -1$).

Now consider the additive category $\text{Mat}(\mathcal{TL}_1)$. The object with no closed components



has an interesting involution: the ‘ $\cdot X$ ’ map (however, we can only use this terminology after passing through the TQFT, which we have not). Recall that this morphism is just the dotted identity map. We will call this endomorphism b for `\bullet`, following Bar-Natan–Morrison. The ‘two dots = 1’ relation in \mathcal{TL}' means that $b^2 = \text{id}$. The corresponding idempotents for this involution are

$$r = \frac{\text{id} + b}{2} \quad \text{and} \quad g = \frac{\text{id} - b}{2}$$

and they satisfy $r^2 = r$, $g^2 = g$, $r + g = \text{id}$, $rg = 0$, $br = r$ and $bg = -g$. So, in $\text{Kar}(\text{Mat}(\mathcal{TL}'_1))$, we have a new decomposition of our object

In $\text{Mat}(\mathcal{TL}'_n)$, for every object containing no closed components, we have a similar story; for each strand, the dotted identity (with the dot on the sheet corresponding to that strand) is an involution, and we get a projections r and g for each strand. In short, an n -strand crossingless matching has 2^n colorings that assign either r or g to each strand. The 2^n colors correspond to the 2^n projections $\{r, g\}^{\otimes n}$.

Let’s now study the chain complex associated to a crossing in the category $\text{Kom}(\text{Kar}(\text{Mat}(\mathcal{TL}'_2)))$:

$$\begin{aligned}
\times &= \begin{array}{c} \text{) } (\xrightarrow{\text{X}} \text{) } (\end{array} \\
&= \begin{bmatrix} \text{) } (\\ \text{) } (\\ \text{) } (\\ \text{) } (\end{bmatrix} \begin{pmatrix} \text{X}_r & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \text{X}_g \end{pmatrix} \begin{bmatrix} \text{) } (\\ \text{) } (\\ \text{) } (\\ \text{) } (\end{bmatrix} \\
&= \begin{array}{c} \text{) } (\xrightarrow{\text{X}_r} \text{) } (\\ \text{) } (\xrightarrow{\text{X}_g} \text{) } (\end{array} \oplus \begin{array}{c} \text{) } (\xrightarrow{\text{X}_r} \text{) } (\\ \text{) } (\xrightarrow{\text{X}_g} \text{) } (\end{array} \oplus \begin{array}{c} \text{) } (\xrightarrow{\text{X}_r} \text{) } (\\ \text{) } (\xrightarrow{\text{X}_g} \text{) } (\end{array}
\end{aligned}$$

Above, we claim that first two summands are contractible. To show this, it suffices to see that the red and green saddle maps are chain homotopy equivalences. [To see this, try to come up with inverses for these saddle maps.](#) Therefore the chain complex for the crossing is equivalent to the last summand, which has no differential. The generators of homology are therefore the four objects drawn.

Now consider a link with n crossings. A neighborhood of every crossing (a diagram representing a complex in $\text{Kom}(\text{Kar}(\text{Mat}(\mathcal{TL}'_2)))$) is equivalent to the complex above with 0 differential. The surviving homology classes are therefore obtained by gluing these complexes in the spaghetti-and-meatballs picture; more precisely, by tensoring together these complexes in a planar algebra. Observe that if, upon tensoring, a single strand is colored by red and green, that diagram dies, because $rg = 0$. [Make sure this really actually makes sense to you.](#) The only remaining objects are complete resolutions where each planar circle is colored entirely either red or green.

Aside 4.3.19. We never carefully defined what a planar algebra is, mostly because it is quite intuitive what ‘gluing’ means (and our intuition is correct).

Here is a slightly more precise definition for our particular setting, though we will remain somewhat vague; we will often refer to this whole story as the ‘spaghetti-and-meatballs’ picture.

Our \mathcal{TL}_n categories are categorifications of the Temperley-Lieb algebras $\text{TL}_n(\delta)$. You can look these up, but the main point is that these are generated by crossingless tangles with $2n$ endpoints, and any time you see a closed component, it evaluates to some $\delta \in R$, where R is the base ring.

Consider a planar tangle \mathcal{T} properly embedded in a disk with a finite number of smaller *input disks* removed where each boundary component contains an even number of endpoints.

Suppose the boundary of the input disk D_i contains $2n_i$ tangle endpoints. Then we can input (i.e. glue) any object of TL_{n_i} into this hole. If the

boundary of the big disk (the dinner plate) contains $2n_0$ endpoints of the tangle, then \mathcal{T} gives us a map of R -algebras

$$\mathrm{TL}_{n_1} \otimes \mathrm{TL}_{n_2} \otimes \cdots \otimes \mathrm{TL}_{n_k} \rightarrow \mathrm{TL}_{n_0}.$$

Similar, for the categories \mathcal{TL}_{n_i} , the tangle \mathcal{T} gives us a \mathbb{Z} -linear multifunctor

$$\mathcal{TL}_{n_1} \times \mathcal{TL}_{n_2} \times \cdots \times \mathcal{TL}_{n_k} \rightarrow \mathcal{TL}_{n_0}.$$

Exercise 4.3.20. In this exercise, you will fill in some details in the proof above of Lee’s structure theorem.

- (a) Prove that

$$\left(\text{red circle} \xrightarrow{\text{red crossing}} \text{red circle} \right) \simeq 0$$

- (b) Explain carefully (i.e. at the level of homological algebra, planar algebras, and Karoubi envelopes) why $rg = 0$ implies that gluing a red strand to a green strand results in an object that is equivalent to the zero object in $\mathrm{Kar}(\mathrm{Mat}(\mathcal{TL}'_1))$.

This concludes the proof of Theorem 4.3.9, as we stated it. However, it remains to show that the surviving objects correspond to orientations on the link diagram, as described in Algorithm 4.3.10. We leave this as an exercise for the reader.

Exercise 4.3.21. Let P be a surviving $\{r, g\}$ -colored complete resolution of a link diagram D , in the proof of Lee’s structure theorem.

- (a) Prove that P satisfies the following property: if two circles in P both about (a small neighborhood of) a crossing c in D , then their colors are different (or more precisely, orthogonal!).
- (b) Use this to argue that P corresponds to some \mathfrak{s}_o . *Perhaps use a counting argument.*
- (c) Determine a formula for the homological grading of the canonical generator $[\mathfrak{s}_o]$. *The quantum filtration level is harder to compute, and is captured by a generalization of the s invariant to links.*

Warning 4.3.22. Nowadays, when discussing other TQFTs, we also use the term ‘Lee homology’ to mean a very specific construction that results in a more degenerate homology theory.

Universal Khovanov homology is the theory over the very general Frobenius algebra $\mathbb{k}[h, t, X]/(X^2 - hX - t)$. (add: cite: turner, vaz, mackaay?) The corresponding ‘Lee homology’ (sorry about the confusing language) is obtained by inverting the discriminant $\mathcal{D} = \sqrt{h^2 + 4t}$. (add: cite: link hom and frob ext?) In Lee’s original theory, we set $h = 0$ and $t = 1$, so $\mathcal{D} = \sqrt{4} = 2$. So, we didn’t need to work over $\mathbb{k} = \mathbb{Q}$; we just needed \mathbb{k} to contain $\frac{1}{2}$. Otherwise, we’d be dividing by 0. That is to say, you shouldn’t do Lee homology over \mathbb{F}_2 if you’re not ok with localizing at 0 (I’m not!); use the Frobenius algebra $\mathbb{F}_2[X]/(X^2 - X)$ instead. This theory appears in earlier literature as ‘Bar-Natan homology’; nowadays, ‘Bar-Natan homology’ typically refers to the theory for $\mathbb{F}_2[h, X]/(X^2 - hX)$, which is more powerful.

Finally, also notice that for the original Frobenius algebra, the discriminant is 0 (which happens when the two roots of the quadratic are the same), so we can’t actually invert the discriminant to get a corresponding ‘Lee homology’.

For my sanity, I prefer to call the ‘corresponding Lee homology’ to a flavor of Khovanov homology as the ‘localized homology’ because we localize at the discriminant.

4.3.3. Defining the s invariant. We now return back to [Ras10]. Throughout, let K be a knot. Then $\mathrm{Lee}(K) \cong \mathbb{Q} \oplus \mathbb{Q}$. Rasmussen’s s is defined using the filtration grading gr_q on the Lee complex.

First note that our definition for the filtration grading of a chain $x \in \text{CLee}(K)$ in (4) is equivalent to

$$\text{gr}_q(x) = \max\{j \in \mathbb{Z} \mid x \in F^j(\text{CLee}(K))\}$$

where F^\bullet is the filtration induced by gr_q on the homogeneous distinguished generators.

This induces a filtration grading on the homology class $[x] \in \text{Lee}(K)$ by

$$\text{gr}_q([x]) = \max\{\text{gr}_q(y) \mid y \text{ is homologous to } x\}.$$

Rasmussen defines two knot invariants, s_{\max} and s_{\min} , using this filtration grading on Lee homology ([Ras10, Definition 3.1]):

$$(5) \quad s_{\min}(K) = \min\{\text{gr}_q([x]) \mid [x] \in \text{Lee}(K), [x] \neq 0\}$$

$$(6) \quad s_{\max}(K) = \max\{\text{gr}_q([x]) \mid [x] \in \text{Lee}(K), [x] \neq 0\}$$

Prove to yourself that (1) these two values are always odd integers and (2) these are knot invariants.

The s invariant is defined to be the average of these two odd integers:

Definition 4.3.23. For a knot in $K \subset \mathbb{R}^3$,

$$s(K) = \frac{s_{\max}(K) + s_{\min}(K)}{2}.$$

In fact, $s(K)$ is always an even integer because of the following proposition, whose proof is the focus of the remainder of this section.

Proposition 4.3.24. $s_{\max} = s_{\min} + 2$. Therefore $s(K) = s_{\max} - 1 = s_{\min} + 1$.

For a knot K , the quantum gradings of the homogeneous generators are all odd. (Why? Prove this to yourself! I do have a favorite proof of this.) Now recall that $d_{\text{Lee}} = d_{\text{Kh}} + \Phi$ where Φ raises gr_q by 4, and so we concluded that $\text{CLee}(K)$ is $\mathbb{Z}/4\mathbb{Z}$ -quantum graded. But in fact, we only have generators in degrees $1 \pmod 4$ and $3 \pmod 4$. We get a decomposition of the chain complex

$$(7) \quad \text{CLee}(K) \cong \text{CLee}_1(K) \oplus \text{CLee}_{-1}(K),$$

where $\text{CLee}_i(K)$ is the summand at the $\mathbb{Z}/4\mathbb{Z}$ -quantum degree $i \pmod 4$.

Lemma 4.3.25 (cf. Lemma 3.5 of [Ras10]). Let o be an orientation of K , and fix a diagram D for K . Then the two cycles $\mathfrak{s}_o \pm \mathfrak{s}_{\bar{o}}$ are contained in two different $\mathbb{Z}/4\mathbb{Z}$ -quantum grading summands of $\text{CLee}(D)$.

Proof. The Lee generators \mathfrak{s}_o and $\mathfrak{s}_{\bar{o}}$ both live in the oriented resolution D_o of a diagram D for K . The quantum grading of any homogenous generator g at D_o is given by

$$\text{gr}_q(g) = \text{gr}_h(D_o) + p(g) + n_+ - n_-$$

where

- $\text{gr}_h(D_o)$ is the homological degree at the resolution D_o and
- $p(g) = \#(v_+ \text{ in } g) - \#(v_- \text{ in } g)$ (following notation in [BN02]).

For the purposes of this proof, we only need to worry about $p(g)$, since the actual quantum gradings for all generators appearing in this proof differ from this local p grading only by an overall shift (of $\text{gr}_h(D_o) + n_+ - 2n_-$).

Define an involution $\tau : V \rightarrow V$ by $v_- \mapsto v_-$, $v_+ \mapsto -v_+$. Extend this to tensor products $V^{\otimes k}$, and in particular, $V^{\otimes |D_o|}$, which is isomorphic to $\text{CLee}(D_o)$ up to the global shift we discussed.

Observe that

- (1) τ acts as identity on generators (pure tensors) containing an even number of instances of v_+ ; call the span of these generators $\text{CLee}_{\text{even}}(D_o)$.
- (2) τ acts as -1 on generators with an odd number of instances of v_+ ; call the span of these generators $\text{CLee}_{\text{odd}}(D_o)$.

Since the number of tensor factors ($|D_o|$) is fixed, these two summands of $\text{CLee}(D_o)$ are at two different $\mathbb{Z}/4\mathbb{Z}$ quantum gradings.

Since $\mathbf{a} = X - 1$ and $\mathbf{b} = X + 1$, we have $\tau(\mathbf{a}) = \mathbf{b}$ and $\tau(\mathbf{b}) = \mathbf{a}$; therefore

$$\tau(\mathfrak{s}_o) = \mathfrak{s}_{\bar{o}}.$$

Therefore

- τ acts as identity on the sum $\mathfrak{s}_o + \mathfrak{s}_{\bar{o}}$, so we must have $\mathfrak{s}_o + \mathfrak{s}_{\bar{o}} \in \text{CLee}_{\text{even}}(D_o)$.
- τ acts as -1 on the difference $\mathfrak{s}_o - \mathfrak{s}_{\bar{o}}$, so we must have $\mathfrak{s}_o - \mathfrak{s}_{\bar{o}} \in \text{CLee}_{\text{odd}}(D_o)$.

Switching p out for gr_q (by applying the quantum shift), we still have that the cycles $\mathfrak{s}_o \pm \mathfrak{s}_{\bar{o}}$ are supported in different $\mathbb{Z}/4\mathbb{Z}$ -quantum grading summands. \square

Remark 4.3.26. In general, if L has n components, then the chains $\text{CKh}(L)$ are all supported in quantum gradings $n \bmod 2$. Lemma 3.5 of [Ras10] gives a more general statement than our discussion above. The definition for the s -invariant can be extended to oriented links, though the filtration gradings of the canonical classes behave more complexly; see (add: cite Beliakova... putyra?).

Lemma 4.3.25 actually tells us more. Because $\text{Lee}(K)$ is 2-dimensional, we now know that we can choose the basis $\{\mathfrak{s}_o \pm \mathfrak{s}_{\bar{o}}\}$, and that these generate the summands of the direct sum decomposition on homology

$$\text{Lee}(K) \cong \text{Lee}_1(K) \oplus \text{Lee}_{-1}(K).$$

arising from the decomposition of the chain complex in (7). We conclude that

$$(8) \quad s_{\max}(K) \geq s_{\min}(K).$$

Furthermore, we must have

$$(9) \quad s_{\min}(K) = \text{gr}_q([\mathfrak{s}_o]) = \text{gr}_q([\mathfrak{s}_{\bar{o}}])$$

because both have components in the $\mathbb{Z}/4\mathbb{Z}$ -quantum grading summand corresponding to $s_{\min}(K)$.

To finish the proof of Proposition 4.3.24, we will need another lemma, which will let us bound the difference between s_{\max} and s_{\min} .

Definition 4.3.27. A map $f : C \rightarrow C'$ between filtered chain complexes (of the form in Definition 4.3.4 is *filtered of degree k* if $f(F_i) \subseteq F'_{i+k}$).

Lemma 4.3.28 ([Ras10], Lemma 3.8). Let K_1, K_2 be knots. There is a short exact sequence

$$0 \rightarrow \text{Lee}(K_1 \# K_2) \xrightarrow{p^*} \text{Lee}(K_1) \otimes \text{Lee}(K_2) \xrightarrow{m^*} \text{Lee}(K_1 \# K_2) \rightarrow 0$$

where p^* and m^* have filtered degree -1 .

Proof. Let D_1, D_2 be diagrams for K_1, K_2 , respectively. We have the following mapping cone:

$$\begin{array}{c} \text{K}_1 \# \text{K}_2^r \\ \underbrace{\quad \quad \quad} \\ \boxed{\mathcal{D}_1} \quad \boxed{\mathcal{D}_2} = \\ \underbrace{\quad \quad \quad} \\ \left(\underbrace{\boxed{\mathcal{D}_1} \quad \boxed{\mathcal{D}_2}}_{\text{K}_1 \sqcup \text{K}_2} \xrightarrow{m} \underbrace{\boxed{\mathcal{D}_1} \quad \boxed{\mathcal{D}_2}}_{\text{K}_1 \# \text{K}_2} \{i\} \right) \end{array}$$

Note that the first diagram isn't exactly $D_1 \# D_2^r$, but you can visualize an isotopy that spins the D_2^r around the x -axis so that it looks like a disjoint union. So, the Lee complex is chain homotopy equivalent to $\text{CLee}(D_1 \# D_2^r)$, which is in turn obviously isomorphic to $\text{CLee}(D_1 \# D_2)$. This yields a short exact sequence (see Section 9.1):

$$0 \rightarrow \text{CLee}(K_1 \# K_2) \{1\} \xrightarrow{i} \text{CLee}(K_1 \# K_2^r) \xrightarrow{p} \text{CLee}(K_1 \sqcup K_2) \rightarrow 0$$

where $\text{CLee}(K_1 \# K_2^r)$ is the chain homotopy representative given by the specific diagram as shown above, and the quantum shift comes from the internal quantum shift in the Khovanov chain complex construction. This gives you a long exact sequence on Lee homology. But because of Lee's structure theorem, we know the dimensions of the \mathbb{Q} vector spaces in the sequence. Since $2 + 2 = 4$, the long exact sequence splits (yay vector spaces!), i.e. $i^* = 0$.

We conclude that we have a short exact sequence

$$0 \rightarrow \text{Lee}(K_1 \# K_2) \xrightarrow{p^*} \text{Lee}(K_1) \otimes \text{Lee}(K_2) \xrightarrow{m^*} \text{Lee}(K_1 \# K_2) \rightarrow 0,$$

but we've kind of discarded the internal quantum degree shifts when we write this down. If we're careful and keep track of everything (i.e. keep track of orientations), we can figure out that both p^* and m^* are filtered of degree -1 , as follows.

Without loss of generality, we can choose diagrams D_1, D_2 so that our crossing is positive. Then the actual mapping cone would be

$$\begin{aligned} \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \nearrow \\ \searrow \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] &= \left(\left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \xrightarrow{m} \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \text{---} \\ \text{---} \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \{1\} \right) \{1\} \\ &= \left(\left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \{1\} \xrightarrow{m} \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \text{---} \\ \text{---} \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \{2\} \right) \end{aligned}$$

By construction, m (secretly m' because we are doing Lee homology)

$$m: \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \{1\} \longrightarrow \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \text{---} \\ \text{---} \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \{2\}$$

is filtered of degree 0. Also, tautologically, p

$$p: \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \nearrow \\ \searrow \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \longrightarrow \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \{1\}$$

is filtered of degree 0. Therefore the *actual* SES maps

$$\begin{aligned} \hat{m}: \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] &\longrightarrow \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \text{---} \\ \text{---} \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \\ \hat{p}: \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \nearrow \\ \searrow \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] &\twoheadrightarrow \left[\begin{array}{|c|} \hline \mathcal{D}_1 \\ \hline \end{array} \right] \begin{array}{c} \circlearrowleft \\ \circlearrowright \end{array} \left[\begin{array}{|c|} \hline \mathcal{D}_2 \\ \hline \end{array} \right] \end{aligned}$$

(that we meant to use to obtain our LES) are actually filtered of degree -1 .

Alternatively, Rasmussen suggests that we first note that the maps must be filtered of some degree, because, without shifts, m^* is induced by a saddle (filtered degree -1) and p^* is induced by a quotient (a filtered degree 0 operation). So, he instead just uses the case $K_1, K_2 = U$ to figure out the filtered degree of these maps, and concludes that they are indeed both -1 .

□

Proof of Proposition 4.3.24. Now let's apply Lemma 4.3.28 to $K_1 = K$ and $K_2 = U$; we get a short exact sequence

$$0 \rightarrow \text{Lee}(K) \xrightarrow{p^*} \text{Lee}(K) \otimes \text{Lee}(U) \xrightarrow{m^*} \text{Lee}(K) \rightarrow 0$$

where each nontrivial map has filtered degree -1 . We know $\text{Lee}(U) = \mathbb{Q}\mathbf{a} \oplus \mathbb{Q}\mathbf{b}$. We also know that one of $\{\mathfrak{s}_o, \mathfrak{s}_{\bar{o}}\}$ has label \mathbf{a} on the component where the connected sum operation occurs; let's call this generator \mathfrak{s}_a , and call the other \mathfrak{s}_b .

We know that $\text{gr}_q([\mathfrak{s}_a + \varepsilon \mathfrak{s}_b]) = s_{\max}$, where ε is either 1 or -1 . In any case, we have

$$m^*([\mathfrak{s}_a + \varepsilon \mathfrak{s}_b] \otimes \mathbf{a}) = [\mathfrak{s}_a].$$

Because m^* is filtered of degree -1 , it can only decrease the filtration degree by at most 1, i.e.

$$\text{gr}_q([\mathfrak{s}_a]) \geq \text{gr}_q([\mathfrak{s}_a + \varepsilon \mathfrak{s}_b] \otimes \mathbf{a}) - 1,$$

i.e.

$$\text{gr}_q([\mathfrak{s}_a + \varepsilon \mathfrak{s}_b] \otimes \mathbf{a}) \leq \text{gr}_q([\mathfrak{s}_a]) + 1.$$

By (9), we have

$$s_{\max}(K) - 1 \leq s_{\min}(K) + 1,$$

and by (8), we conclude that

$$s_{\max}(K) = s_{\min}(K) + 2.$$

□

4.3.4. s is a concordance homomorphism. We begin this section with the properties of s that show it is a concordance homomorphism.

These are proven in [Ras10], but that I leave for you to ponder on your own. If you aren't able to prove the statements on your own (but with the hints provided), feel free to look at Rasmussen's proof (and understand it) and then put it into your own words.

Exercise 4.3.29 ([Ras10], Proposition 3.9). (Maybe hard?) Let \bar{K} denote the mirror of K . Prove that $s(\bar{K}) = -s(K)$.

Hint: Let D be a diagram for K , and let \bar{D} be the diagram where all the crossing information is reversed. Explain briefly why, in the category of chain complexes of filtered vector spaces over \mathbb{Q} , we have

$$\text{CLee}(\bar{D}) \cong \text{CLee}(D)^*$$

where $(\cdot)^*$ is the dualizing functor that

- sends $v_{\pm} \mapsto v_{\mp}^*$; here $\text{gr}_q(v_{\mp}^*) = \mp 1$; I don't like this notation... The point is that under the degree-preserving isomorphism $V^* \cong V$, we have $v_{\pm}^* \mapsto v_{\mp}$. You can therefore think of the functor $(\cdot)^*$ as taking $v_{\pm} \mapsto v_{\mp}$.
- reverses all quantum shifts (in fact, all quantum gradings);
- reverses all homological gradings and is contravariant on morphisms.

Remark 4.3.30. In Khovanov's notation, the dualizing functor is denoted $(\cdot)^!$. In other contexts, you may also see it written as $(\cdot)^{\vee}$, though the diagram mirroring rule might be different (e.g. reflect across x -axis). Nevertheless, however you mirror a diagram, there is a clear identification of chains for D and for \bar{D} .

Exercise 4.3.31 ([Ras10], Proposition 3.11). (Maybe hard?) For knots K_1, K_2 , prove that $s(K_1 \# K_2) = s(K_1) + s(K_2)$.

Hint: Use Lemma 4.3.28 for bounds on the filtration degree of morphisms, and then flip your morphisms upside down and use the result of Exercise 4.3.29 too.

4.3.5. *Behavior of canonical generators under cobordisms.* In the proof of Proposition 4.3.24, we saw how the filtration degree of maps helped us bound the behavior of s . Now, we will study the behavior of s under cobordisms, by using the fact that every cobordism can be decomposed into elementary cobordisms, and we know the filtration degree of the Lee map associated to each elementary cobordism F : it's $\chi(F)$!

But this is not enough to know what happens to s : after all, the induced map could be 0.

Remark 4.3.32. Our discussion here will differ from Rasmussen's original paper, mainly because it was written before Bar-Natan's categories were introduced. (Rasmussen posted [Ras10] in February 2004, whereas Bar-Natan posted [BN05] in October 2004.) Rasmussen carefully proves that the Reidemeister maps in Lee homology are filtered of degree 0. For us, this fact is almost immediate because Bar-Natan's proof of Reidemeister invariance (in the most general $S, T, 4Tu$ cobordism categories) is grading-preserving. Instead, we can imagine passing first through a TQFT for the Frobenius algebra $\mathbb{Z}[T, X]/(X^2 = T)$, where $\deg_q(T) = -4$, and then setting $T = 1$ (and then localize at $\sqrt{D} = 2$, or just $\otimes \mathbb{Q}$). This preserves the *filtration* grading, because now all the arrows that used to have a T coefficient (and were grading-preserving) are now increase filtration degree by 4.

We know the Reidemeister moves are chain homotopy equivalences in Bar-Natan's category, so these maps are nonzero. We are more worried about compositions of Morse moves. Luckily, we have the following lemma:

Lemma 4.3.33 (see proof of Proposition 4.1 in [Ras10]). Let $S : L \rightarrow L'$ be a cobordism with no closed components. Let o be an orientation on L . Then

$$\phi_S([\mathfrak{s}_o]) = \sum_I a_I [\mathfrak{s}_{o'_I}]$$

where $\{o'_I\}$ is the set of all orientations on L' that are compatible with (L, o) on the bottom boundary of S , and where all coefficients a_I are nonzero.

Let us state this another way, using Rasmussen's vocabulary. There are two types of connected components F of the cobordism S :

- F is of the *first type* if it has a boundary component in L . In this case, the chosen orientation on the bottom boundary of F determines the orientation on the rest of the boundary components on L' .
- F is of the *second type* if it only has boundary on L' . Then either orientation of F is compatible with (L, o) . This yields two different possible sets of orientations on the components of L' at the (top) boundary of F .

An orientation on F is *permissible* (with respect to (L, o)) if the bottom boundary of F agrees with (L, o) . The lemma therefore states that the set of orientations $\{o'_I\}$ is the set of top boundary orientations of the set of permissible orientations on S .

Remark 4.3.34. Requiring that S contain no closed components obstructs the situation where you have, say, a closed undotted sphere that makes the entire cobordism evaluate to 0.

Remark 4.3.35. Clarify this on Friday add remark on cobordism boundary orientation (in and out).

Before we prove Lemma 4.3.33, let's state the important corollaries:

Corollary 4.3.36 (Proposition 4.1 of [Ras10]). If S is an oriented cobordism from (L, o) to (L', o') that is *weakly connected*, i.e. every component of S has a boundary component in L (i.e. S has no components of the second type), then $\phi_S([\mathfrak{s}_o])$ is a *nonzero* multiple of $\mathfrak{s}_{o'}$.

Corollary 4.3.37 (Corollary 4.2 of [Ras10]). If S is a connected cobordism between knots, then ϕ_S is an isomorphism.

This corollary is particularly important because we can then study the filtration degree of ϕ_S to understand the behavior of s .

Proof of Lemma 4.3.33. To set up, first isotope S so that you may select regular values $0 < t_1 < t_2 < \dots < t_k = 1$ such that, between any t_i and t_{i+1} , there is at most one critical value of S , and there are no critical values between 0 and t_1 . **This hopefully feels like a familiar procedure by now!**

Let $L_i = S \cap \mathbb{R}^3 \times \{t_i\}$, and let S_i be the cobordism $S \cap [0, t_i]$, a cobordism $L \rightarrow L_i$.

The proof basically proceeds by induction on the heights of elementary cobordisms. The cobordism S_1 is a planar isotopy; this is the base case, and the proposition holds because $\phi_{S_1} = \text{id}$.

Assume the proposition holds for S_i . Then $S_e := S_{i+1} - S_i$ consists of a single elementary cobordism (union identity everywhere else, of course). The induction step holds if S_e is a Reidemeister move, so we only need to consider the cases of the four possible Morse moves; see Notation 4.3.8. **We will abuse notation a bit below and let ϕ_{S_e} also denote the chain map in $\text{Kom}(gg\text{Vect}_{\mathbb{Q}})$, rather than just the map on homology. These are individually not hard to prove; you should actually think through these yourself.**

- (0) If S_e is a 0-handle attachment (birth), then $\phi_{S_e}(\mathfrak{s}(o_I)) = \mathfrak{s}(o_I) \otimes \frac{1}{2}\mathbf{a} \otimes \mathbf{b}$. (This is a disjoint union, so the differentials and homology respect the tensor product decomposition.) **Done!**

add pic

- (1) If S_e is a 1-handle attachment (i.e. a merge or a split), once again you can check by inspection that the induction step holds. **Verify it! Here you actually need to think about orientations in the plane.**

add pic

- (2) If S_e is a 2-handle attachment (death), then there's only one possible compatible orientation for L_{i+1} ; since $\varepsilon'(\mathbf{a}) = \varepsilon'(\mathbf{b}) = 1$, you basically just erase the component that died. But we need to check that only one permissible orientation on S_i (and thus on L_i) extends to this unique orientation on L_{i+1} . **Here, argue why if two different permissible orientations on S_i extended to the orientation on L_{i+1} , then S overall would have a closed component.**

□

Exercise 4.3.38. Fill in the details in the proof of Lemma 4.3.33 for the four Morse moves, by supplying pictures working out the possible cases and explaining why the induction step holds in each case.

Remark 4.3.39. **This is a very subtle comment. Feel free to ignore.** Notice that while we are using the Lee maps associated to cobordisms, nowhere did we actually use the fact that this assignment of cobordisms to filtered chain maps actually defines a functor. In other words, Rasmussen's invariant can be defined in settings where you have something weaker than an actual functor; we only needed the fact that isotopic cobordisms have the same Euler characteristic, and that Euler characteristics add when you glue cobordisms along collections of circles.

4.3.6. *Profit!* Now that we understand how Lee's canonical generators behave under cobordisms, we reap the benefits and prove a bunch of facts about the behavior of s under cobordisms embedded in 4D. Clever mathematicians (**like you!**) can then use these facts to do interesting topology.

Theorem 4.3.40 ([Ras10], Theorem 1).

$$|s(K)| \leq 2g_4(K).$$

Proof. Suppose $K \subset S^3$ bounds a genus- g slice surface $F \subset B^4$. Let S be the cobordism $K \rightarrow U$ obtained by deleting a 0-handle from F and flipping

it upside down. Then $\chi(S) = -2g$. Let $[x]$ be a homology class where $\text{gr}_q([x]) = s_{\max}(K)$, and further pick $x \in [x]$ so that $\text{gr}_q(x) = \text{gr}_q([x]) = s_{\max}(K)$. **Note that none of these classes or cycles are 0.**

Since $s_{\max}(U) = 1$, we have

$$\text{gr}_q(\phi_S(x)) \leq 1.$$

Since ϕ_S is filtered of degree $-2g$, we have

$$\text{gr}_q(\phi_S(x)) \geq \text{gr}_q(x) - 2g$$

and so

$$s(K) + 1 = s_{\max}(K) = \text{gr}_q(x) \leq \text{gr}_q(\phi_S(x)) + 2g \leq 1 + 2g$$

so we conclude that

$$s(K) \leq 2g.$$

To get that $s(K) \geq -2g$, we use the same argument but for \bar{K} , which is cobordant to the unknot via the surface \bar{S} , and we similarly conclude that $s(\bar{K}) \geq 2g$. But $s(\bar{K}) = -s(K)$, so we multiply the whole inequality by -1 to get $s(K) \geq -2g$. The theorem follows. \square

Exercise 4.3.41. Suppose knots K and K' differ by just one crossing change.

- (a) Prove that $|s(K) - s(K')| \leq 1$.
- (b) From this, obtain a bound on the *unknotting number* $u(K)$ of knot K , where $u(K)$ is the minimal number of self-intersections needed in a homotopy of $K \hookrightarrow S^3$ that changes K into an unknot.

Theorem 4.3.42. If D is a positive diagram for a positive knot K , then

$$s(K) = s_{\min}(K) + 1 = \text{gr}_q(\mathfrak{s}_o) + 1.$$

Proof. The oriented resolution D_o is the unique resolution at the left-most vertex of the cube of resolutions. There are no differentials into this homological grading, so \mathfrak{s}_o and $\mathfrak{s}_{\bar{o}}$ are alone in their respective homology classes. \square

Exercise 4.3.43 ([Ras10], Corollary 1). (The Milnor Conjecture) Let $\text{gcd}(p, q) = 1$. For the braid-closure diagram of the torus knot $T_{p,q} = \hat{\beta}$ for $\beta = (\sigma_1 \sigma_2 \cdots \sigma_{p-1}^q)$, compute $\text{gr}_q(\mathfrak{s}_o)$. Deduce the Milnor Conjecture:

$$g_4(T_{p,q}) = g_3(T_{p,q}) = \frac{(p-1)(q-1)}{2}.$$

Remark 4.3.44. Rasmussen also shows that the s invariant is not so useful for alternating knots, as it gives the same bound as the classical knot signature. Manolescu–Ozsváth later proved that, in fact, Khovanov homology is not very interesting for a larger class of recursively defined knots called *quasialternating knots* [MO08]. **We probably won't cover this; good idea for a final project.**

4.4. Non-isotopic (smooth) surfaces in the 4-ball; exotic disks.

Here we give some more examples where Khovanov homology was used to study surfaces embedded in 4D.

Recall that an $(n+1)$ -dimensional TQFT will give an invariant of $(n+1)$ -dimensional *closed* manifolds: a cobordism $F : \emptyset \rightarrow \emptyset$ yields map $Z(F) : R \rightarrow R$, which is determined by $Z(F)(1) \in R$. Rasmussen and Tanaka (add: citation) showed that this invariant for (vanilla) Khovanov homology is simply determined by the genus of the surface F .

However, if we instead consider a relative version of such an invariant, we do get a rich invariant; this is possible because the target of the map $\text{Kh}(K)$, is big.


Definition 4.4.1. Let F be a slice surface for a knot K . The *Khovanov-Jacobsson class* of F is

$$\text{Kh}(F)(1) \in \text{Kh}(K).$$

Example 4.4.2. Sundberg–Swann [SS22] showed that the slice knot $K = 9_{46}$ has a pair of non-isotopic slice disks D, D' by showing that their *Khovanov–Jacobsson classes* are unequal:

$$\text{Kh}(D)(1) \neq \text{Kh}(D')(1).$$

Here are the disks:



(Image taken from their paper.) For the pink disk, surgery the knot at the pink arc; you will find the result unravels to a two-component unlink. Fill these in and run your movie backwards. Do the same for the orange disk.


Remark 4.4.3. Sundberg–Swann’s paper is much more general; they show that for any pretzel knot $K = P(p, -p, p)$, you can use the same trick to get two different slice disks. Furthermore, they show that you get the same result if you take arbitrary connected sums of such K .

When a topologist sees this kind of result, they will naturally wonder, “Is this because the disks were *topologically* non-isotopic in the first place, or are these disks actually an *exotic pair*?” And a certain type of topologist will get very excited if the latter is true.


Let’s casually explain the difference. (All embeddings are proper.)

Let F be a surface homeomorphically embedded in B^4 , without any further regularity conditions. The embedding is *locally flat* if, in any neighborhood of any point on $F \subset B^4$, we can choose coordinates so that F looks flat, in a smooth sense.

For example, using the Jordan curve theorem, you can convince yourself that any embedding $S^1 \hookrightarrow \mathbb{R}^2$ is locally flat. However, as we saw on the first day of class, this is no longer true for $S^1 \hookrightarrow \mathbb{R}^3$:



or $S^2 \hookrightarrow \mathbb{R}^3$:



(Image taken from wolfram.com’s article.)

Definition 4.4.4. Let F and F' be two *smoothly* embedded surfaces in B^4 that share the same boundary in $S^3 = \partial B^4$.

- We say F and F' are *topologically (locally flatly) isotopic* if there is an ambient isotopy taking (B^4, F) to (B^4, F') through locally flat(ly embedded) surfaces, rel boundary.
- If F and F' are topologically isotopic but *not smoothly isotopic*, then we say that they are an *exotic pair* (or just *exotic*, for short).

Remark 4.4.5. Sometimes I cook food from my ancestral culture and it's referred to as 'exotic'. I like to think that the people who say this really mean to say that my dish and their dish are an 'exotic pair'. After all, being 'exotic' is a symmetric relation.


One particularly famous question about exotic phenomena in dimension 4 is THE:

Conjecture 4.4.6 (Smooth Poincaré Conjecture in dimension 4 (SPC4)). (Open) If X is homeomorphic to S^4 , then X is diffeomorphic to S^4 .

As a smooth invariant, Khovanov homology can help with questions of this type. We will see some results about disks here, and then later about 4-manifolds when we talk about skein lasagna modules.

Here is an example, again only a specific one taken from a larger class described in the paper:

Example 4.4.7. Hayden–Sundberg showed that the following knot J bounds slice disks D and D' that are topologically isotopic (by work of Akbulut) but not smoothly isotopic, by using Khovanov homology.



(Image taken from their paper.) They view the two disks upside-down, and study the maps $\text{Kh}(D)$ and $\text{Kh}(D')$ by identifying what happens to a class called *Plamenevskaya's class* $\phi \in \text{Kh}(J)$ under these maps. For one of these disks, $\text{Kh}(D)(\phi) = 0$ but for the other, $\text{Kh}(D')(\phi) \neq 0$.

Plamenevskaya's class will show up again very very soon!

That excited topologist from before will probably then also ask, “Do these surfaces remain exotic after *stabilization*? How many times do I need to *stabilize* before the surfaces become smoothly isotopic?”

In this case, *stabilization* refers to connect summing with a standard (small, boundary-parallel) torus embedded in B^4 . In 4D topology, one expects that exoticness between two objects is eventually lost after stabilizations, because, roughly speaking, there is more room for maneuvering. *We'll talk a little more about this later on, but not about the details.*

It turns out that if we work with a TQFT over a polynomial ring, such as $\mathbb{F}_2[H]$, then Khovanov homology ‘sees’ stabilization (‘adding a handle’), e.g. as multiplication by $2X - H$ in the specific case of $\mathcal{A} = \mathbb{F}_2[H, X]/(X^2 = HX)$ (*this theory is usually called Bar-Natan homology*).

Example 4.4.8. Hayden [Hay23] gave an example (again, one of many) where (reduced) Bar-Natan homology can show us two disks remain smoothly non-isotopic even after one stabilization. He first finds two disks slice disks for a knot K that are not topologically isotopic, and therefore also not smoothly isotopic. Then, he computes that

$$H \cdot \left(\widetilde{\text{BN}}(D)(1) - \widetilde{\text{BN}}(D')(1) \right) \neq 0$$

and therefore

$$H \cdot \widetilde{\text{BN}}(D)(1) \neq H \cdot \widetilde{\text{BN}}(D')(1)$$

as classes in $\text{BN}(K)$. *Where did the $2X$ part of the handle $2X - H$ go? We're over characteristic 2!*

Remark 4.4.9. We did not talk about reduced homology, but will when we start talking about relationships between Khovanov homology and Floer homologies. For now, just know that it's a slimmer version of full Khovanov homology, so if it can tell you something, so can the full theory.

5. LEGENDRIAN, TRANSVERSE, AND ANNULAR LINKS

5.1. Knots in the standard contact \mathbb{R}^3 and S^3 . In this section, we give just enough exposition on the standard contact structure on S^3 as well as Legendrian and transverse knots in this setting to proceed with our applications of Khovanov homology. See [John Etnyre's notes \[Etn05\]](#) for a fairly standard introductory reference on this material, and much more context.


5.1.1. Legendrian knots in $(\mathbb{R}^3, \xi_{std})$. Roughly speaking, a *contact structure* on a 3-manifold is a 2-plane field that is *maximally non-integrable*, i.e. not tangent to any surface on any neighborhood of any point on the surface.

Definition 5.1.1. The *standard contact structure* on \mathbb{R}^3 is

$$\xi_{std} = \ker(dz - ydx).$$

Here $\alpha = dz - ydx$ is the *standard contact 1-form* on \mathbb{R}^3 .

The *contact planes* are invariant under x and z translation (i.e. parallel transport) and look like this at (the tangent spaces of the points on) the plane $z = 0$:



(Image from [\[Etn05\]](#), where he says the images are courtesy of S. Schönenberger.)

Along each line parallel to the y -axis, you can think of these planes as the tangent spaces to a belt that is overall a half-twisted band, but whose ends are at infinity. [Let's call this an 'infinite half-twisted band'](#). In particular, the planes are *nearly* vertical (i.e. yz -planes) when $|y|$ is large.

While embedded 2-manifolds can't be tangent to ξ_{std} at more than isolated points, embedded 1-manifolds can. For example, notice that the x -axis is everywhere tangent to the contact planes.

Definition 5.1.2. A *Legendrian knot (or link)* Λ in $(\mathbb{R}^3, \xi_{std})$ is a smoothly embedded knot (or link) that is everywhere tangent to ξ_{std} .

[In other words, \$T_*L \subset \xi_{std}|_L\$, i.e. the tangent bundle of \$L\$ is contained in the contact plane bundle over \$L \subset \mathbb{R}^3\$.](#)


Again, we will consider Legendrian knots only up to *Legendrian isotopy*, i.e. a smooth isotopy only through Legendrian knots.

The x - and z -translation invariance of ξ_{std} allows us to draw Legendrian knots very easily, by projecting away the y -coordinate. The only thing we need to remember is that, if you see a crossing, the more negatively sloped strand is closer to you.

A *front projection* of a Legendrian knot (or link) $\Lambda \subset (\mathbb{R}^3, \xi_{std})$ is the (image of the) projection of Λ onto the xz -plane. [The viewer is at either \$\pm\infty\$ on the \$y\$ -axis. It doesn't matter which one, as long as we're all at the same place.](#)

If you stare at the contact form again, you will notice that the front projection will satisfy these properties:

- There cannot be any vertical tangencies.
- Consequently, there will be left and right cusps:




- At any crossing, the more negatively sloped strand is in front (i.e. closer to the viewer):



There are of course *Legendrian Reidemeister moves* that relate front diagrams of the same Legendrian knot. See Figure 8 of [Etn05].

Remark 5.1.3. Each Legendrian knot of course also represents a smooth knot. But because of the stricter geometric conditions on the equivalence between Legendrian knots, the set of Legendrian knots is a refinement of the set of smooth knots. For example, the Chekanov pair is a pair of Legendrian *representatives* of the knot $m(5_2)$ that are famously hard to tell apart, but were proven to be in fact non-Legendrian-isotopic (add: citation):



(K_1 is a sleeping cat with a long torso but short neck. K_2 is a sleeping cat with a short torso but a long neck.)

5.1.2. *Transverse knots in $(\mathbb{R}^3, \xi_{rot})$.* We can also look at a less restrictive class of knots that are everywhere *transverse* to the contact planes.

Definition 5.1.4. A *transverse knot (or link)* in $(\mathbb{R}^3, \xi_{std})$ is a smoothly embedded knot (or link) \mathcal{T} where at any point $p \in \mathcal{T}$, the tangent space is transverse to the contact planes. In other words, at all $p \in \mathcal{T}$,

$$T_p \mathcal{T} \pitchfork (\xi_{std})_p.$$

Since our contact structure is orientable (it's defined globally by a single 1-form, and therefore we can talk about the positive and negative sides of the contact planes), we will also require that our transverse knots intersect the contact planes *positively*. This allows us to fix an orientation on transverse knots (and links) without having to draw any orientation arrows. (Technically, these are *positive transverse knots*, and we could also study the entirely analogous set of *negative transverse knots*.)

Remark 5.1.5. You can also describe front diagrams of transverse knots. The geometry will give you the following constraints:

- no downward pointing vertical tangencies (but upward is ok):
- no downward pointing positive crossings:

There are also transverse Reidemeister moves for front diagrams; see Figure 14 of [Etn05].

Remark 5.1.6. Transversality is a generic condition, so transverse knots are a less fine refinement of the set of smooth knots than Legendrian knots. In fact, by perturbing a Legendrian knot, you will land in a well-defined transverse knot class.

The front projection gives us a diagrammatic way to make sense of the statement

$$\{\text{transverse knots}\} \leftrightarrow \{\text{Legendrian knots}\} / \sim$$


(which we will not discuss but that you can easily look up).

Two contact structures ξ_1, ξ_2 on the same 3-manifold M are *contactomorphic* if there is a diffeomorphism that sends $(M, \xi_1) \rightarrow (M, \xi_2)$. In order to study a different type of knot living in contact 3-space, we sometimes prefer to instead work with a more rotationally symmetric version of the standard contact structure, by using cylindrical coordinates (r, θ, z) for \mathbb{R}^3 instead of the Cartesian (x, y, z) .

Definition 5.1.7. The *rotationally symmetric contact structure on \mathbb{R}^3* is

$$\xi_{rot} = \ker(dz + r^2 d\theta) \quad (= \ker(dz + xdy - ydx)).$$

The contact plane for ξ_{rot} are now z - and θ -translation invariant. The contact planes look like this:




Take the infinite half-twisted band from before, and now lay them out radially so that their flat planes are identified at the z -axis. Copy this to each z -coordinate. Take the tangent planes.

In particular, the contact planes are *nearly* vertical when $|r|$ is large. This is particularly exciting because this means that if you take a braid, wrap it around the z axis and close it up, and then push it radially far enough away from the origin, the result is a transverse link!

The more surprising thing is that this characterizes *all* transverse links. In order to appreciate this fact, we need to take a detour to talk about braid closures.

Remark 5.1.8. There are moves that relate braid closure representatives of \mathcal{T} . This will be discussed in the next section (the Transverse Markov Theorem).

5.1.3. Braid closures and annular links. Let $\beta \in B_n$ be a braid on n strands. The *braid closure* of β , denoted $\hat{\beta}$, is the link obtained by wrapping β around the z -axis and gluing the strands:



This is also referred to as taking the *trace*. Prove to yourself that indeed, this operation satisfies the trace condition $\text{Tr}(\beta_1 \beta_2) = \text{Tr}(\beta_2 \beta_1)$.

A braid closure naturally lives in a *thickened annulus* $A \times I$, which is diffeomorphic to the complement of a neighborhood of the z -axis $\cup \{\infty\}$ in S^3 .

Remark 5.1.9. The thickened annulus is also diffeomorphic to the solid torus. Also, if you're not too careful about details about open and closed sets, you can also think of this ambient space as

$$\mathbb{R}^3 - \{z\text{-axis}\} = (\mathbb{R}^2 - \{0\}) \times \mathbb{R}.$$


In any case, the braid closure is properly embedded and compact, so any of these descriptions is valid. They all appear in the literature. The important thing is that you can draw your diagrams on an annulus by projecting away the 'thickening' direction.

Theorem 5.1.10. Here are some classical theorems about knots in S^3 .


Alexander's theorem (proved by Alexander) showed that every knot or link in S^3 can be braided with respect to the z -axis. The *Yamada-Vogel* algorithm (discovered by Vogel building on work of Yamada) gives a hands-on way to turn a link diagram into a braided diagram (through valid isotopies, of course).

The *Markov theorem* (originally proven by Markov, but now there are many, many proofs) states that, for any two braids $\beta_1 \in B_{n_1}$ and $\beta_2 \in B_{n_2}$ whose closures represent a link L , β_1 and β_2 are related by a finite sequence of the following moves:

- (Markov I) conjugation, i.e. replacing $\beta \in B_n$ with $\omega\beta\omega^{-1}$ where ω is any braid in B_n :



- (Markov II) positive or negative stabilization / destabilization, i.e. Reidemeister 1 of either type across the braid axis):



These are meant to be construed in addition to braid isotopy in the annulus (i.e. braid-like Reidemeister 2 and 3, i.e. braid relations in the braid group).

Remark 5.1.11. Together, these two theorems allow us to study knots in terms of their braid representatives. This can be very useful. For example, at the moment, *triple-graded Khovanov-Rozansky homology* (a.k.a. HOM-FLYPT homology) can only be computed from a braided diagram of a link, even though it's an invariant for links in S^3 .

We have similar theorems for braid representatives of transverse links. Bennequin proved the analogue to Alexander's theorem:

Theorem 5.1.12 ([Ben83]). (Bennequin's theorem) For any transverse knot in $(\mathbb{R}^3, \xi_{rot})$, there exists a representative that is braided with respect to the z -axis.

Braids are also naturally oriented; a (positive) transverse link will be oriented counterclockwise.

Nancy Wrinkle proved the following in her PhD thesis (!):

Theorem 5.1.13 ([Wri02], [OS03]). (Transverse Markov theorem) Any two braid representatives of a transverse link are related by conjugation and *positive* Markov stabilization / destabilization.

These theorems allow us to study transverse knots via their braid representatives.

Finally, note that you can also have other non-braided links living in the solid annulus; braid (conjugacy classes) are just a special type of *annular link*.

An *annular link* L is a link $L \subset A \times I$. These links can be projected onto the annulus; these are called *annular diagrams*. We usually draw them on the plane and put a marking where the axis is. In other words, an annular link is really just the data of a link L and an unknotted axis U inside S^3 . We will talk more about annular links later.

5.1.4. *Summary.* Let's summarize the relationships among the different types of links in \mathbb{R}^3 that we've introduced.

We have the following surjections of sets.

- (1) Thinking about the front projection and $(\mathbb{R}^3, \xi_{std})$, we have

Legendrian links \twoheadrightarrow transverse links \twoheadrightarrow smooth links

- (2) Thinking about braids and $(\mathbb{R}^3, \xi_{rot})$, we have

braids \twoheadrightarrow braid conjugacy classes \twoheadrightarrow transverse links \twoheadrightarrow smooth links.

- (3) Finally, we obviously have

annular links \twoheadrightarrow smooth links

by just forgetting the unknotted axis.

5.2. **Plamenevskaya's transverse knot invariant.** The reference for this section is [Pla06]

Definition 5.2.1. Let $\beta \in B_n$ be a braid representative for $\mathcal{T} = \hat{\beta}$.

- *Plamenevskaya's cycle*, which we will denote by \vec{v}_- , is the chain in $\text{Kh}(\hat{\beta})$ at the oriented resolution where all circles are labeled v_- .
- *Plamenevskaya's invariant*, denoted ψ , is the homology class in $\text{Kh}(\mathcal{T})$ of \vec{v}_- .

Once again, this is a theorem-definition. We need to convince ourselves of the following facts:

- (1) \vec{v}_- is indeed a cycle. In a previous exercise, you proved that there are no split maps out of an oriented resolution. Observe that any merge map out of D_o sends $\vec{v}_- \mapsto 0$.
- (2) ψ is invariant under the transverse Markov moves, so that it is actually an invariant of

Moreover, we also naturally want to know the answers to the following questions:

- (1) Is ψ actually a transverse knot invariant, or is it just a smooth knot invariant? In other words, are there smoothly isotopic transverse knots $\mathcal{T} \neq \mathcal{T}'$ such that $\psi(\mathcal{T}) \neq \psi(\mathcal{T}')$? Note that '1' is a knot invariant, but it's not very useful.
- (2) How does ψ behave under cobordisms? Is it *functorial*, i.e. if $F : \mathcal{T} \rightarrow \mathcal{T}'$ is a cobordism, is $\text{Kh}(F)(\psi(\mathcal{T})) = \psi(\mathcal{T}')$? You can restrict to symplectic cobordisms F , but since Khovanov homology is a smooth invariant, $\text{Kh}(F)$ only takes into account the underlying smooth cobordism.
- (3) Is ψ *effective*? In other words, can ψ distinguish pairs of transverse knots that existing invariants can't?

The first question is easy to answer, after we have the following proposition.

Proposition 5.2.2. If \mathcal{T} is a negative stabilization of some \mathcal{T}' , then $\psi(\mathcal{T}) = 0$.

Thus it suffices to find $\hat{\beta}$ such that $\psi(\hat{\beta}) \neq 0$, since the negative stabilization of $\hat{\beta}$ represents the same *smooth* knot type.

Example 5.2.3. Let $\mathbb{1}_n$ denote the identity braid in the braid group B_n . The transverse (fillable, max tb) unknot satisfies $\psi(\hat{\mathbb{1}}_1) \neq 0$ but its stabilization $\psi(\hat{\sigma}_1^{-1}) = 0$.

To answer the last question, we need to know what the *classical* transverse invariants are.

Definition 5.2.4. Let \mathcal{T} be a transverse knot.

- Let $\beta \in B_b$ be a braid representative for \mathcal{T} . The *self-linking number* of \mathcal{T} is

$$sl(\mathcal{T}) = sl(\hat{\beta}) = -b + \text{wr}(\hat{\beta}).$$

In terms of a front diagram D for \mathcal{T} , the self-linking number is just $\text{wr}(D)$. There is a more geometric definition of self-linking, by thinking of \mathcal{T} as a framed knot.

You can quickly check for yourself that self-linking is preserved by the transverse Markov moves.

- The *smooth knot type* of \mathcal{T} is simply the equivalence class of the underlying smooth knot.

A transverse invariant is called *effective* if it can distinguish transverse knots better than these classical knots can.

In other words, in practice we just need the invariant to tell us more than just the number $sl(\mathcal{T})$.

Proposition 5.2.5. The quantum grading of the class $\psi(\mathcal{T})$ is the self-linking number of \mathcal{T} .

Exercise 5.2.6. Prove Propositions 5.2.2 and 5.2.5.

Question 5.2.7. (Open) Is Plamenevskaya's invariant effective?

In fact, there have been many refinements of Plamenevskaya's invariant defined throughout the years (e.g. [LNS15]). It is not known if these are effective either!

5.3. Annular Khovanov homology. Annular Khovanov homology is the a version of Khovanov homology for links in the thickened annulus. We will first give a definition for AKh as a tri-graded homology theory, analogously to how we defined Khovanov homology. Then, in §5.3.2, we give a more nuanced interpretation of AKh as it relates to Kh.

5.3.1. Annular Khovanov homology. Let $L \subset A \times I$ be a link in the thickened annulus, and let D be a diagram for L drawn on the annulus A . Instead of drawing the annulus, we typically mark the position of the deleted z -axis with an asterisk or other marking.

The annular Khovanov chain complex $\text{CAKh}(D)$ is generated by the same distinguished generators as for Khovanov homology, but we now distinguish between homologically trivial and nontrivial circles. We assign them *triply* graded modules, graded by gr_h , gr_q , and a new *winding number grading*, which we interpret below in Remark 5.3.2.



trivial (W)



nontrivial / essential
(V)

- A circle Z is *trivial* if the marked point $*$ and ∞ are in the same region of $\mathbb{R}^2 \setminus Z$. To these circles, we associate the tri-graded module

$$\mathbb{W} := \mathbb{Z}\langle w_+, w_- \rangle$$

where

$$(\text{gr}_h, \text{gr}_q, \text{gr}_k)(w_{\pm}) = (0, \pm 1, 0).$$

- A circle Z is *non-trivial* or *essential* if the marked point $*$ and ∞ are in different regions of $\mathbb{R}^2 \setminus Z$. To these circles, we associate the tri-graded module

$$\mathbb{V} := \mathbb{Z}\langle v_+, v_- \rangle$$


where

$$(\text{gr}_h, \text{gr}_q, \text{gr}_k)(v_{\pm}) = (0, \pm 1, \pm 1).$$

The tri-grading for the Khovanov generators is then defined by extension to tensor products, and allowing for homological and quantum shifts.

The annular Khovanov differential d_{AKh} is comprised of all the components of the Khovanov differential that preserve gr_k grading. One can work out all six possible interactions between trivial and essential circles to determine the tri-grading-preserving differential.

Instead of re-drawing these, here's stuff grad-school-me TeXed up [Zha18]:



The ‘type’ names are not used in the literature; I just needed them for a proof later in the paper.

The homology of $(\text{CAKh}(D), d_{\text{AKh}})$ is an invariant of the smooth isotopy type of the annular link.

Remark 5.3.1. One can also define the annular Bar-Natan category, typically denoted $\text{BN}(\mathbb{A})$ in the literature. We will not dig into this in this course.

- Remark 5.3.2.** (1) Pick a generic arc from $*$ to ∞ . For a given *oriented* circle ζ in the plane missing the marked point $*$, γ intersects ζ transversely (otherwise, perturb). The signed intersection (algebraic intersection) $I(\gamma, \zeta)$ is 0 if ζ is homologically trivial in the annulus, and is ± 1 if γ is essential. If we associate counter-clockwise circles with the ‘plus’ labeling (and clockwise with ‘minus’), then $I(\gamma, \zeta) = \text{gr}_k(\zeta)$, (interpreted appropriately).
- (2) Note that this association of counter-clockwise/clockwise with plus/minus is a matter of convention and bookkeeping. This ‘orientation’ on ζ is *not* meant to be interpreted topologically in the context of cobordisms and such!
- (3) Grigsby–Licata–Wehrli show that $\text{AKh}(L)$ has a $\mathfrak{sl}_2(\mathbb{C})$ -representation, where gr_k is the weight-space grading. From this point of view \mathbb{W} is two copies of the trivial representation, and \mathbb{V} is one copy of the defining representation.

So, if you’re working with \mathbb{C} coefficients, it’s also fair to call gr_k the weight-space grading. I have in the past, quite unfortunately, also used this terminology even while working over \mathbb{F}_2 before really absorbing the ‘winding number grading’ terminology, which is more intuitive.

5.3.2. Summary for AKh. Given the data of a link diagram D in \mathbb{R}^2 and a basepoint $*$ in the complement of D in \mathbb{R}^2 , the winding number grading gr_k defines a filtration on the Khovanov chain complex $(\text{CKh}, d_{\text{Kh}})$. Thus (CKh) is gr_h and gr_q graded, and gr_k filtered.

The annular Khovanov homology of the annular link represented by $(D, *)$ is the homology of the *associated graded object* to the gr_k -induced filtration on CKh . The Khovanov differential decomposes into two homogeneous-degree components:

$$d_{\text{Kh}} = d_{(1,0,0)} + d_{(1,0,-2)}$$

Definition 5.3.3. Let F_\bullet be a \mathbb{Z} -filtration on a chain complex (\mathcal{C}, d) , with $F_i \supseteq F_{i+1}$. The situation with inclusion the other way is basically identical. The *associated graded object* is the chain complex

$$(\bigoplus F_i/F_{i+1}, \bar{d})$$

where \bar{d} is the induced differential on the quotients. If \mathcal{C} is generated by distinguished generators that are homogeneous with respect to a filtration grading, then \bar{d} is comprised precisely of the components of the differential between generators with the same filtration grading.


So, for an annular link $L \subset A \times I$, the annular Khovanov homology $\text{AKh}(L)$ is a triply-graded invariant of annular links.

- Exercise 5.3.4.** (a) Compute the annular Khovanov differentials for all six merge/split interactions between pairs of circles. Confirm that d_{Kh} only either preserves gr_k or decreases it by 2.
- (b) Since gr_k is defined on the level of Khovanov generators, this also gives a filtration on the *Lee differential*. Identify all the $(\text{gr}_h, \text{gr}_q, \text{gr}_k)$ homogeneous-degree components of the Lee differential.

Remark 5.3.5. Grigsby–Licata–Wehrli defined an annular version of Rasmussen’s invariant by studying the annular Khovanov–Lee complex, which is gr_h graded and $(\text{gr}_q, \text{gr}_k)$ -bifiltered. By taking a linear combinations of gr_q (denoted gr_j in their paper) and gr_k , they obtain a 1-parameter family

$\{d_t\}$ of *annular* concordance invariants [GLW17]. This type of construction (mixing two filtration gradings) was popularized by Ozsváth–Stipsicz–Szabó’s Υ invariant from knot Floer homology.

Here is a cartoon:



The blue lines are level sets in the $\text{gr}_t = \text{gr}_q - t\text{gr}_k$ filtration grading. For any $t \in [0, 2]$, gr_t is a valid filtration grading on the annular-Khovanov–Lee complex.

Grigsby–Licata–Wehrli then apply these d_t invariants to measure some contact topological quantities, such as *right-veeringness* of the mapping class group element for the b -pointed disk given by a braid β .

Remark 5.3.6. We probably won’t say much more about the mapping class group interpretation of braids in this course, but for those of you interested, here’s an intuitive definition of the property of ‘right-veeringness’.

The braid group B_n can be thought of as the mapping class group of a disk with n (indistinguishable) marked points. A braid $\beta \in B_n$ is then a dance of these n marked points around each other, ending with the points returning to the original marked positions, but possibly permuted. This gives a sort of stirring action on the particles around them.

Draw an oriented arc γ from basepoint on the boundary of the disk to one of the marked points, then apply the stirring action by β . If the resulting arc γ' veers right upon leaving the boundary, then β is right-veering.

We can use mapping class groups to describe the construction of contact 3-manifolds by Giroux’s *open-book decompositions*. Honda–Kazez–Matić (add: cite) showed that a contact structure (M, ξ) is *tight* if and only if *all* its open books have right-veering monodromies.

5.3.3. *Hubbard–Saltz’s braid conjugacy class invariant κ .* As I mentioned previously **but possibly only verbally in class**, Plamenevskaya’s cycle $\tilde{\psi}$ / class ψ gives us a homology class to measure, and thereby obtain invariants from in different settings. This is analogous to how Lee’s canonical generators are measured to obtain Rasmussen’s invariant.

As originally defined, ψ is inherently annular. Hubbard and Saltz use the gr_k from AKh to define a braid conjugacy class invariant κ by measuring the filtration level of ψ [HS16]:

Definition 5.3.7 ([HS16], Definition 1). Let $\beta \in B_n$. Let $\tilde{\psi}(\hat{\beta}) \in \text{CAKh}(\hat{\beta})$ denote Plamenevskaya’s cycle in the annular Khovanov complex for the closure of β . Let $\{F_i\}$ denote the gr_k -induced filtration on $\text{CAKh}(\hat{\beta})$.

- If $\psi(\hat{\beta}) \in \text{im}(d_{\text{Kh}})$, let

$$\kappa(\beta) = n + \min\{i \mid [\psi(\hat{\beta})] = 0 \in H(F_i)\}$$

In other words, if ψ dies in homology, κ measures the filtration level at which it dies.

- If $\psi(\hat{\beta}) \neq 0 \in \text{Kh}(\hat{\beta})$, then let $\kappa(\beta) = \infty$. If ψ never dies, then κ is infinite.

Remark 5.3.8. In [HS16, Theorem 2], Hubbard–Saltz show that κ sometimes increases by 2 under positive (Markov) stabilization, and is therefore not a transverse invariant.

However, braid conjugacy class invariant can also be useful. For example, they use κ to show the following:

Theorem 5.3.9 ([HS16], Corollary 17). Let $\beta \in B_n$. If $\kappa(\beta) \neq 2$ and $\kappa(m(\beta)) \neq 2$, then $\beta = 1 \in B_n$.

In short, satisfying both parts of the hypothesis tell us that the mapping class group element β is both right-veering and left-veering, and by work of Baldwin-Grigsby [BG15], the only such braid is the identity braid.

Remark 5.3.10. All of these ‘filtration level at which ...’ definitions can equivalently be made in terms of spectral sequences. For example, κ gives a lower bound on the length of the spectral sequence $\mathrm{AKh}(\hat{\beta}) \Rightarrow \mathrm{Kh}(\hat{\beta})$.

Since many of you might not have worked with spectral sequences before, I will give an introduction to this algebraic tool next Wednesday before we talk about these in class.

Remark 5.3.11. Lipshitz, Ng, and Sarkar also give various filtered refinements of Plamenevskaya’s invariant in [LNS15], which are more likely to be effective than ψ (because they are refinements). [If you are interested, this could be a good final project.](#)

5.4. Ng’s Thurston-Bennequin bound for Legendrian knots. Recall that a knot $\Lambda \subset (\mathbb{R}^3, \xi_{std})$ is *Legendrian* if it is always tangent to the contact planes $\xi_{std} = \ker(dz - ydx)$. We draw front diagrams to describe Legendrian knots, where the more negatively sloped strand is in front, and the diagram has no vertical tangents, only cusps.

We have talked about classical invariant for transverse knots: the topological (smooth) knot type, and the self-linking number. We now discuss classical invariants for Legendrian knots, and how smooth topology plus framing information can tell us about Legendrian representatives of knots K .

Remark 5.4.1. Everything we discuss in this section is for oriented links. For knots, the orientation turns out not to matter. I will however write ‘knot’ everywhere for simplicity.

Definition 5.4.2. Let Λ be a Legendrian knot. We have the following classical invariants:

- (0) Topological knot type, i.e. $K = [\Lambda]$, the equivalence class of Λ under smooth isotopy
- (1) *Thurston-Bennequin number*, which measures how much the contact planes twist as you follow the knot. Given a front diagram F of Λ ,

$$\mathrm{tb}(\Lambda) = \mathrm{wr}(F) - c(F)$$


where

$$c(F) = \frac{1}{2}(\#\text{cusps}).$$

A zigzag in the front projection represents the knot traveling like a spiral staircase. If you think of this as the path of a car traveling up to down levels in a multi-story parking garage, $c(F)$ basically measures how many times you drive around in circles (ignoring the z -coordinate).

- (1) *Rotation number*, which measures (roughly speaking) how many spiral staircase floors you travel up and down (counted with sign). Given a front diagram F for Λ ,

$$\mathrm{rot}(\Lambda) = \#(\text{down cusps} - \text{up cusps}).$$



In my head, I think of this in terms of a kind of energy. Suppose you are running around a building with a lot of spiral staircases. You gain energy by descending staircases but lose energy by climbing staircases. In this Escher-esque world, you can travel the path of a

knot and end up with more or less energy than you started with!
(Sounds freaky, but this is just... monodromy.)

Remark 5.4.3. Recall that for a transverse knot \mathcal{T} , the self-linking number is just the writhe of a front projection. So, if \mathcal{T} is a transverse push-off of Λ , then

$$sl(\mathcal{T}) = \text{tb}(\Lambda) + c(F).$$


Our motivating question today is the following:

Question 5.4.4. Given a smooth knot K , what is the maximal $\text{tb}(\Lambda)$ for Legendrian representatives Λ of K ?


The *maximal Thurston-Bennequin number* achievable by a Legendrian representative of K is called the “max tb” of K , and is denoted $\overline{\text{tb}}(K)$.

Remark 5.4.5. (1) Given K , we can study the Legendrian *geography* and *botany* of K . If you plot all Legendrian representatives of K on a lattice with coordinates (rot, tb) , the support will be a “mountain range” (add: add citation) (see [?]). The $\overline{\text{tb}}$ is the height of the tallest mountain. *Geography* is the study of the support, and *botany* is the study of the flora (i.e. Legendrian representatives) at each particular lattice point.


(2) The tallest mountain might not be unique. For example, the mountain range for positive torus knots looks like this:



while the mountain range for negative torus knots looks like this:



(3) It is easy to descend the mountain, i.e. decrease tb , simply by negative and positive stabilization:



However, it is not always possible to climb the mountain.

Remark 5.4.6. (1) In the drawings above, the thicker strands in the ‘S’ and ‘Z’ are closer to you. The ‘view from above’ is a cartoon only. This would normally be called a Lagrangian projection, but technically then one needs to be careful about the areas enclosed

by curves in the projection. We will not comment on this further in this course.

- (2) Notice that the labels ‘positive’ and ‘negative’ for these Legendrian stabilizations have more to do with whether the rotation number increases or decreases, respectively. From the point of view of smooth topology, these are both negative R1 moves.
- (3) On the other hand, the legal Legendrian R1 moves



(and the 180° rotation of the above) are smoothly positive R1 moves. This jives with the idea that, for the transverse pushoff of this Legendrian, positive stabilizations preserve transverse link type, but negative stabilization change the transverse link type (and therefore definitely also the Legendrian link type).

I really ought to give citations for all this; I will need to spend some time finding the right references. Nevertheless, you do not need these references for the remainder of this course.

Bennequin (add: cite) discovered the first bound on $\overline{\text{tb}}$; we have since found stronger and stronger bounds. Here’s a sampling:

Proposition 5.4.7. Let Λ be a Legendrian representative for K . Then $\text{tb}(\Lambda) \leq \overline{\text{tb}}(K)$, which is in turn bound by the following values: (add: all citations)

- (1) (Bennequin) $\overline{\text{tb}}(K) \leq 2g_3(K) - 1$ Impactful result – was used to prove that there are contact structures on \mathbb{R}^3 that are homotopic but not contactomorphic to $(\mathbb{R}^3, \xi_{std})$; do not ask me for details. Maybe ask Orsola instead. :)
- (2) (Slice-Bennequin bound)

$$\overline{\text{tb}}(K) \leq 2g_4(K) - 1$$

So the original Bennequin bound was probably only detecting the 3D shadow of the real bound, which comes from 4D behavior! This makes sense because $(\mathbb{R}^3, \xi_{std})$ is the boundary of the standard symplectic 4-ball, which comes from the complex structure of the unit ball in \mathbb{C}^2 .

- (3) (s -Bennequin bound) [Pla06, Shu07]

$$\overline{\text{tb}}(K) \leq s(K) - 1$$

Here $s(K)$ is Rasmussen’s invariant.

- (4) (Ng’s Khovanov bound)

$$\overline{\text{tb}}(K) \leq \min\{\delta \mid \text{Kh}^\delta(K) \neq 0\}$$

where $\text{gr}_\delta = \text{gr}_h - \text{gr}_q$. We will discuss the δ grading and prove this bound today.

Remark 5.4.8. Observe that (4) \implies (3) \implies (2) \implies (1). Such is the progress of humanity.

Exercise 5.4.9. Prove the s -Bennequin inequality.

Hint: First relate $\overline{\text{tb}}$ with sl . Then compare the quantum grading of Plamenevskaya’s cycle \vec{v}_- in the Khovanov complex with the quantum filtration of Lee’s canonical classes in the Lee homology.

Let’s state Ng’s ‘strong Khovanov bound’ again the way he phrased it:

Theorem 5.4.10 ([Ng05], Theorem 1). For any link K ,

$$\overline{\text{tb}}(K) \leq \min\{k \mid \bigoplus_{i-j=k} \text{Kh}^{i,j}(K) \neq 0\}.$$

Remark 5.4.11. For alternating knots, the Khovanov homology (over \mathbb{Q}) is rather simple; it is supported on two diagonals on the $(\text{gr}_h, \text{gr}_q)$ lattice. This was first shown by Lee in [Lee05], and was extended to *quasi-alternating links* by Manolescu and Ozsváth [MO08]. Consequently, the Khovanov homology of a quasi-alternating knot is entirely determined by two ‘classical’ invariants, the knot signature and the Jones polynomial. I was at a conference in Berkeley a few summers back where Vaughn Jones objected to his polynomial being called ‘classical’.

Ng proves that his bound is *sharp* for alternating knots.

Remark 5.4.12. The δ grading here is significant because it is also the grading used by Seidel–Smith in their *symplectic Khovanov homology* [SS06], Kh_{symp} , which is a Lagrangian intersection Floer homology interpretation of Khovanov homology that is isomorphic to (combinatorial, δ -graded) Khovanov homology over fields of characteristic 0 [AS19].

This means that Ng’s bound can equivalently be viewed as a $\overline{\text{tb}}$ bound coming from symplectic Khovanov homology!

In addition to working with δ -graded Khovanov homology, Ng also makes use of a ‘shifted’ version of Khovanov homology, which depends on the diagram; you may think of this as a version of Khovanov homology for *framed links*, i.e. links made out of ribbon (or pappardelle) rather than string (or spaghetti).

Observe that, for a distinguished generator g for $\text{CKh}(D)$, the δ -grading is given by

$$\text{gr}_\delta(g) = p(g) + \text{wr}(D).$$

Ng’s *shifted Khovanov homology* Kh_{sh} is related to regular (δ -graded) Khovanov homology by

$$(10) \quad \text{Kh}^*(D) = \text{Kh}_{sh}^{*- \text{wr}(D)}(D)$$

so that the grading for g in CKh_sh is simply $p(v)$.

Remark 5.4.13. This terminology (‘shifted Kh’) is entirely local to this current section, and is not used in the literature, as far as I’m aware. However, the homology theory absolutely does appear elsewhere in the literature in various forms (and slightly different conventions...). Here’s how I think about Kh vs Kh_{sh} .

If we were to define a bracket for Kh so that we didn’t have to include global shifts at the end, we would use the bracket relations

$$\begin{aligned} \left[\begin{array}{c} \nearrow \searrow \\ \nwarrow \nearrow \end{array} \right]_{\text{Kh}} &= \left(\underbrace{\left[\begin{array}{c}) (\end{array} \right]_{\{1\}}} \longrightarrow \left[\begin{array}{c} \smile \frown \end{array} \right]_{\{2\}} \right) \\ \left[\begin{array}{c} \nearrow \searrow \\ \nwarrow \nearrow \end{array} \right]_{\text{Kh}} &= \left(\left[\begin{array}{c} \smile \frown \end{array} \right]_{\{-2\}} \longrightarrow \underbrace{\left[\begin{array}{c}) (\end{array} \right]_{\{-1\}}} \right) \end{aligned}$$

For the framed link invariant Kh_{sh} , the bracket relation would instead be

$$\begin{aligned} \left[\begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} \right]_{\text{Kh}_{sh}} &= \left(\underline{\left[\begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} \right]_{\{0\}}} \longrightarrow \left[\begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} \right]_{\{1\}} \right) \\ \left[\begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} \right]_{\text{Kh}_{sh}} &= \left(\left[\begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} \right]_{\{-1\}} \longrightarrow \underline{\left[\begin{array}{c} \diagup \diagdown \\ \diagdown \diagup \end{array} \right]_{\{0\}}} \right) \end{aligned}$$

The remainder of this section will be dedicated to proving the following proposition, which will immediately imply Theorem 5.4.10.

Proposition 5.4.14 ([Ng05], Proposition 5). Let D be a front diagram for a Legendrian knot Λ with smooth knot type K . Then

$$\text{Kh}_{sh}^*(D) = 0$$

for all $*$ $< -c(D)$, and therefore $\text{Kh}^*(D) = 0$ for all $*$ $< \text{tb}(D)$.

The ‘therefore’ part is clear after you remember that $\text{tb}(D) = \text{wr}(D) - c(D)$.

We will need the following lemma, which should hopefully feel somewhat similar to Lemma 4.3.28. *If you want practice with long exact sequences coming from mapping cones, try proving this lemma!*

Lemma 5.4.15 ([Ng05], Lemma 6). Let D be a front diagram, and let D_0, D_1 denote the (Legendrian) 0- and 1- resolutions at a specific crossing:

$$\smile \xleftarrow{0} \times \xrightarrow{1} \succ \prec$$

There is a long exact sequence

$$\begin{array}{ccc} \text{Kh}_{sh}(D_0) & \xrightarrow{(-1)} & \text{Kh}_{sh}(D_1) \\ & \nwarrow \quad \nearrow & \\ & \text{Kh}_{sh}(D) & \end{array}$$


where the (-1) indicates a shift in δ -grading.

Because this triangle is exact, we immediately have the following corollary.

Corollary 5.4.16 ([Ng05], Lemma 7). If $\text{Kh}_{sh}(D_0)$ and $\text{Kh}_{sh}(D_1)$ are both supported only in δ gradings $\geq n$, then so is $\text{Kh}_{sh}(D)$.

As you can imagine, the proof of Proposition 5.4.14 will proceed by induction on the number of crossings (since D_0 and D_1 both have one fewer crossing than D).


(Base case.) If D has no crossings, then D is an unlink (call it U^n) of n components, for some n . Since $\text{gr}_\delta(g) = p(g)$ for all generators in $\text{CKh}_{sh}(U^n)$, the complex (and homology) are supported on δ gradings $\{-n, \dots, n\}$. The Legendrian representative consisting only of max-tb unknots




has $c(D) = n$ for the diagram above (and is in fact the max-tb representative of U^n). So, indeed, the $\text{Kh}_{sh}(U^n)$ is supported only in gradings $\geq c(D)$.

(Induction step.) Now assume that for all diagrams with fewer crossings than D , the claim holds.

For the sake of understanding Ng's argument quickly, we are going to draw a concrete example to follow. But note that the proof is completely general.




In D , pick a right-most crossing. There are two possible scenarios:



- (1) The two strands emanating rightward out of the crossing meet at a cusp, or
- (2) the two strands emanating rightward out of the crossing do not meet at a cusp, and in fact go to two separate cusps.

We study the two cases separately. In our example, Case 1 appears first, and then after reducing to D_0 , we need to use Case 2, when the diagram looks like this:



Case 1. Resolving the crossing in Case 1, we have

$$\begin{array}{c} \text{Diagram with crossing and blue circle} \\ \mathcal{D} \end{array} = \left(\begin{array}{c} \text{Diagram with crossing and blue circle} \\ \mathcal{D}_0 \end{array} \longrightarrow \begin{array}{c} \text{Diagram with crossing and blue circle} \\ \mathcal{D}_1 \end{array} \right)$$

By (smooth) R1 invariance, $\text{Kh}(D) \cong \text{Kh}(D_0)$. Note that $\text{wr}(D_0) = \text{wr}(D) - 1$. Combined, this tells us that $\text{Kh}_{sh}^*(D) \cong \text{Kh}_{sh}^{*-1}(D_0)$.

On the other hand, $c(D_0) = c(D)$. By the induction hypothesis, $\text{Kh}_{sh}(D_0)$ is supported on δ -gradings $\geq c(D_0) = c(D)$. Therefore $\text{Kh}_{sh}^*(D) \cong \text{Kh}_{sh}^{*-1}(D_0)$ must also be supported in these gradings as well.

Case 2. Resolving the crossing in Case 2, we have

$$\begin{array}{c} \text{Diagram with crossing and blue circle} \\ \mathcal{D} \end{array} = \left(\begin{array}{c} \text{Diagram with crossing and blue circle} \\ \mathcal{D}_0 \end{array} \longrightarrow \begin{array}{c} \text{Diagram with crossing and blue circle} \\ \mathcal{D}_1 \end{array} \right)$$

By Lemma 5.4.15, it suffices to show that both D_0 and D_1 are supported in degrees $\geq c(D)$.

- Note that $\text{wr}(D_0) = \text{wr}(D) - 1$, but $c(D) = c(D_0)$; by a similar argument as in Case 1, D_0 satisfies the induction hypothesis.
- While $c(D_1)$ is different from $c(D)$, D_1 is clearly isotopic to a diagram where you remove the stabilization on the far right strand. We can then apply the same argument as before.

This concludes the proof of Proposition 5.4.14!

6. SPECTRAL SEQUENCES, Kh'S RELATION TO GAUGE / FLOER THEORIES

Spectral sequences are an algebraic tool in homological algebra that computes homology by taking successively more accurate approximations.

In our text, the spectral sequence itself can be a useful object to study, e.g. as a knot invariant.

We will start out with the most general version of a spectral sequence that you will see in this course; these come from filtered complexes. We'll then discuss more specific settings, like spectral sequences coming from bicomplexes.

6.1. Spectral sequences from filtered complexes. Let's begin by recalling some basics of filtered complexes in a different light, and in the process, set some notation (slightly different from before).

A *chain complex* (of, say, \mathbb{Z} -modules), in its most general form, is just a pair (\mathcal{C}, d) where \mathcal{C} is a \mathbb{Z} -module and d is an endomorphism of \mathcal{C} satisfying $d^2 = 0$.

In the situations we're usually most familiar with, there is a homological grading, and $\mathcal{C} = \bigoplus C_i$ where C_i are the homologically graded pieces.

A (\mathbb{Z}) -*filtration* on the complex (\mathcal{C}, d) is a sequence of subcomplexes $\{F_j \mathcal{C}\}_{j \in \mathbb{Z}}$ such that

$$F_j \mathcal{C} \supseteq F_{j+1} \mathcal{C}$$

and $d(F_j \mathcal{C}) \subseteq F_j \mathcal{C}$. (In other words, d preserves the filtration level.)

In our settings, we also require that there are levels m and M where $F_m \mathcal{C} = \mathcal{C}$ and $F_M \mathcal{C} = \emptyset$. This makes our filtration *finite-length*.

The filtration could also be indexed such that $F_j \mathcal{C} \subseteq F_{j+1} \mathcal{C}$. We will see examples of both.

In our setting (i.e. Khovanov homology), \mathcal{C} is generated by a set of *distinguished generators* that are homogeneous with respect to some *filtration grading* gr_f .

In this case, we can decompose the differential as

$$d = d_0 + d_1 + d_2 + \cdots$$

where d_i is the component of the differential that is gr_f graded degree i . Because our filtrations are finite-length, we know that eventually, for k large enough, $d_k = 0$, so this is a finite sum.

Remark 6.1.1. Notice that the associated graded chain complex can be written as just (\mathcal{C}, d_0) in this case.

Remark 6.1.2. Observe that since $d^2 = 0$, we know that the graded pieces of the endomorphism d^2 must be 0, so we know

$$\begin{aligned} d_0^2 &= 0 \\ d_1 d_0 + d_0 d_1 &= 0 \\ d_2 d_0 + d_0 d_2 + d_1^2 &= 0 \\ &\vdots \end{aligned}$$

and so on.


Definition 6.1.3. The *filtration spectral sequence* for the filtered complex (\mathcal{C}, d) with filtration $\{F_j \mathcal{C}\}$ is a sequence of chain complexes, successively defined as follows:

- $(E^0, d^{(0)}) = (\mathcal{C}, d_0)$; note that this is just the associated graded complex
- $(E^1, d^{(1)})$:
 - E^1 is the homology of E^0
 - $d^{(1)}$ is the graded degree 1 piece of the induced differential \bar{d} , viewed as a map on homology; this is the lowest graded piece of the induced map on homology
- $(E^2, d^{(2)})$ is obtained the same way: E^2 is the homology of E^1 , and $d^{(2)}$ is the lowest degree graded piece of the (now doubly induced map) \bar{d} , which is in fact graded degree 2.

Remark 6.1.4. First of all, if this definition is hard to parse, don't worry – we will see an example and it will all be much clearer.

- The notation in the definition above isn't really standard. For example, I added the parentheses around the superscript to remind you that, in general, $d^{(i)}$ is *not* d_i . (It wouldn't even make sense since they're maps on completely different algebraic objects.)
- The E^i are called the *pages* of the spectral sequence. Every time you turn the page, you are taking a subquotient (i.e. homology).
- Subject to some boundedness conditions, these spectral sequences will *collapse* (or *terminate*) on a finite page. This means that, eventually, for $k \geq N$ for some large N , $E^k = E^N$, and all the $d^{(k)} = 0$. We say that the spectral sequence *abuts* to the homology on page E^N , and we call this page the “infinity page” E^∞ .

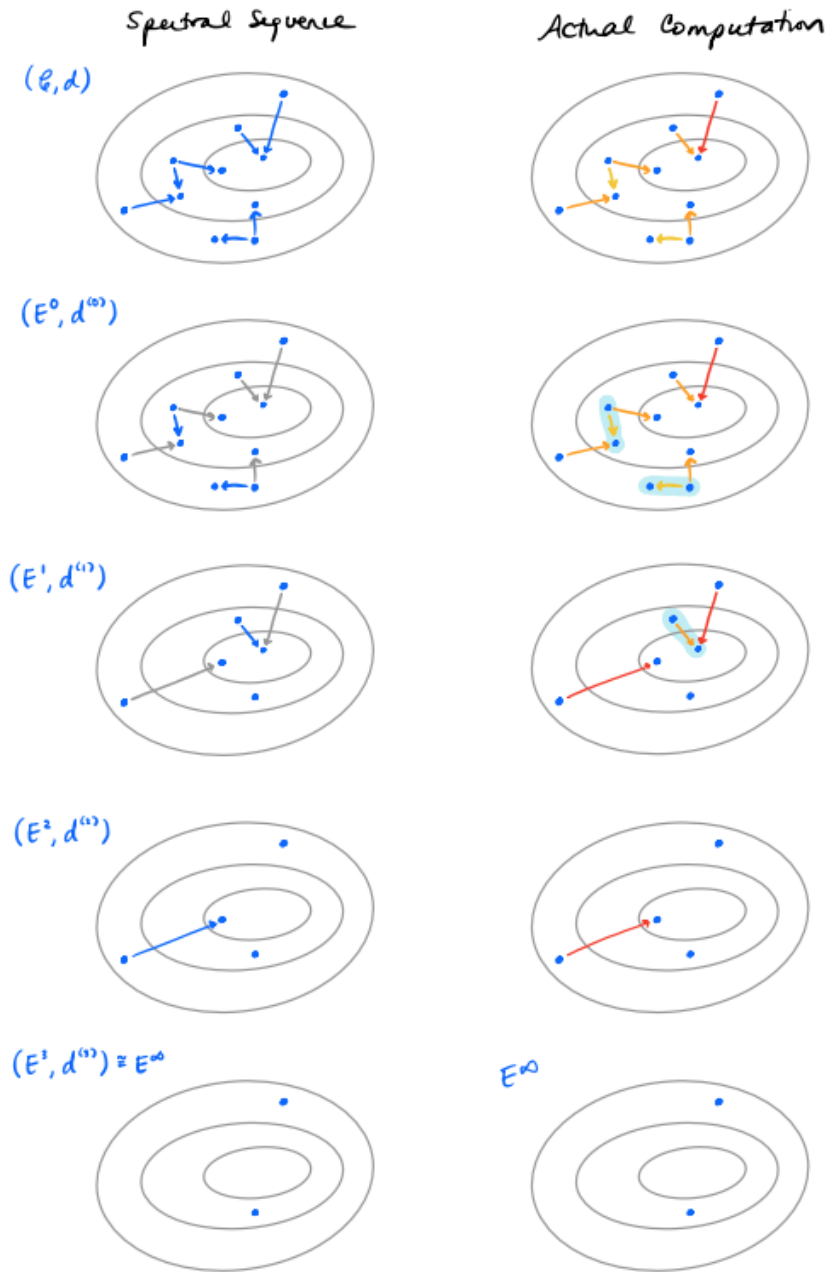
For instructional reasons, I will work over \mathbb{F}_2 today, so that we don't have to complicate the discussion with signs. Our main tool is Gaussian elimination (see Corollary 3.5.6), which works over \mathbb{F}_2 as follows:



If you wish to work over characteristic 0, then you have to add a sign to your zigzag differentials.


Example 6.1.5. Here is an illustrative example of an abstract filtered complex and the pages of the associated spectral sequence.

- On the left are the pages of the spectral sequence; induced maps that are not to be considered on the present page are grayed out.
- On the right is the actual computation one might do, where one keeps track of the longer induced differentials while working through the calculations of the pages.



Exercise 6.1.6. The Lee complex is gr_q -filtered, and the homology of the associated graded complex is Khovanov homology. This gives a spectral sequence, known as the Rasmussen-Lee spectral sequence, or the Khovanov-to-Lee spectral sequence. Compute the Khovanov-to-Bar-Natan spectral sequence over \mathbb{F}_2 for the trefoil appearing in [BN02].

Recall from Warning 4.3.22 that Bar-Natan homology corresponds to the Frobenius algebra $\mathbb{F}_2[X]/(X^2 - X)$. The merge and split maps are shown below:




The solution to this exercise is available on the class website. It will later probably be included as an appendix.


To see how the Kh-to-Lee spectral sequence can be used to bound some topological quantities, see [AD19] and [CGL⁺21].

Example 6.1.7. Another filtration spectral sequence comes from the fact that AKh is the homology of the associated graded complex for Kh with respect to the gr_k grading. This filtration grading satisfies $F_k \subseteq F_{k+1}$. Since $d_{\text{Kh}} = d_0 + d_{-2}$ where $d_0 = d_{\text{AKh}}$, our spectral sequence first computes differentials that preserve gr_k , then those that *decrease* gr_k by 2, then those induced differentials that *decrease* gr_k by 4, and so on.

In the following annular Hopf link example, we don't have longer differentials, but you can see the difference between $\text{AKh}(D)$ and $\text{Kh}(D)$ on the different pages of the spectral sequence nevertheless.



6.2. Spectral sequences from bicomplexes. Roughly speaking, a *bicomplex* is a 2D complex with two differentials, one given by vertical arrows, and one given by horizontal arrows, as shown in the cartoon below:



It is automatically bigraded by $\mathbb{Z} \oplus \mathbb{Z}$ by the row number and the column number.


Here is a more formal definition with more details, now that you have some context.

Definition 6.2.1. A *bigraded complex* or *bicomplex* is the data of

- chain groups $C_{i,j}$ for $i, j \in \mathbb{Z}$
- two gradings gr_r and gr_c , where $(\text{gr}_r, \text{gr}_c)(C_{i,j}) = (i, j)$
- two commuting differentials d_h and d_v such that
 - the $(\text{gr}_r, \text{gr}_c)$ bidegree of d_h is $(0, 1)$ and
 - the $(\text{gr}_r, \text{gr}_c)$ bidegree of d_v is $(1, 0)$.

Remark 6.2.2. The *total complex* of the bicomplex is a flattening of the above data into a single (one-dimensional) complex, with homological grading given by $\text{gr}_r + \text{gr}_c$ (whose level sets are lines of slope -1 in the cartoon).


However, since d_h and d_v commute, their sum is not a differential. So, we need to introduce a minus sign to every other row of d_h differentials in order to make each square *anticommute*:



Then, we get the *total complex* of the bicomplex, which is a flattening of the bicomplex into a 1-dimensional, linear complex. If \mathcal{B} denotes the bicomplex, we denote the total complex (or *totalization*) by $\text{Tot}(\mathcal{B})$.

We again work over \mathbb{F}_2 in this section, in order to keep things as simple as possible while we are learning a new algebraic gadget. In this case, there is no difference between commuting and anticommuting differentials.

Suppose we have a bicomplex $(\mathcal{C} = \bigoplus C_{i,j}, d_h, d_v)$ of \mathbb{F}_2 vector spaces. The total complex clearly has two filtration gradings: gr_r and gr_c . We call these the *row-wise* and *column-wise* filtrations:



There are therefore two associated filtration spectral sequences:


- The filtration spectral sequence for the row-wise filtration (for filtration grading gr_r) first computes homology with respect to d_h on page 0. On page 1, it proceeds to compute homology with respect to the chain maps induced by d_v . So, we denote this spectral sequence by $^{hv}E_\bullet$.
- The filtration spectral sequence for the columnwise-wise filtration (for filtration grading gr_c) first computes homology with respect to d_v on page 0. On page 1, it proceeds to compute homology with respect to the chain maps induced by d_h . So, we denote this spectral sequence by $^{vh}E_\bullet$.

Both spectral sequences ultimately compute the homology of the total complex, so they abut to isomorphic vector spaces: $^{hv}E_\infty \cong ^{vh}E_\infty$.

Remark 6.2.3. *Beware*, though, that because the total differential is filtered with respect to two different filtration gradings in these spectral sequence computations, the gradings will be messed up. That is, the two infinity pages might have support on different lattice points! So all you can really say is that the ungraded homologies are isomorphic.

However, if you have some third unrelated grading gr_u (u for *unrelated* to the row and column gradings) hanging around, then obviously the spectral sequences will not affect that grading. In that case, you really have many spectral sequences, one for each value of gr_u , and they don't talk to each other at all.

The longer differentials (i.e. differentials on pages ≥ 2 are induced by cancellation of differentials in the previous pages. For example, the 'length-2' differentials on page 2 of the two spectral sequences travel like this:



The length- n differentials would similarly travel the path determined by zigzags caused by Gauss elimination.

Remark 6.2.4. I’m using the terms ‘Gauss elimination’ and ‘taking homology’ sort of interchangeably, because we are working over a field. Allow me to explain:

Recall that ‘taking homology’ is essentially the process of finding the smallest chain homotopy representative in a chain homotopy equivalence class; this is the one with zero differential, and we call it the homology.

Gaussian elimination (a.k.a. cancellation) is an algorithm for finding this smallest chain homotopy representative. Given a chain complex (C, d) , we can Gaussian eliminate away components of the differential d , *in some order*, until we end up at a complex with zero differential. Because we are **working over a field**, we are able to cancel along any component of the differential, because every nonzero coefficient is a unit.

This process is equivalent to changing the basis for our chain groups so that our original chain complex becomes a direct sum of the homology and a bunch of acyclic complexes, which we then discard.

As a demonstration of the power of these two spectral sequences ${}^{hv}E_\bullet$ and ${}^{vh}E_\bullet$ computing the total homology, we will now discuss examples coming from equivariant homology. The goal is to obtain a spectral sequence relating two interesting objects X and Y via the following strategy.

- (1) Set up a bicomplex using the topological data relating X and Y .
- (2) Arrange so that ${}^{vh}E_1$ is easy to compute, and describes X .
- (3) Arrange so that ${}^{hv}E_\bullet$ is easy to compute, so that you know ${}^{hv}E_\infty$, which describes Y .
- (4) Assuming the bicomplex is sufficiently bounded, we know that ${}^{vh}E^\infty$ must then also describe Y , and we are done. (The ${}^{vh}E_\bullet$ spectral sequence therefore relates X and Y .)

6.2.1. Classical Smith inequality. Let p be a prime. Let X be a topological space with a $\mathbb{Z}/p\mathbb{Z}$ action generated by $\sigma : X \rightarrow X$ (that is, $\sigma^p = \text{id}_X$). Let X^{fix} denote the fixed-point set, which itself is also a topological space.


The *classical Smith*¹³ *inequality* gives an inequality on the total dimensions of the singular homologies of X and X^{fix} , over \mathbb{F}_p :

$$\dim H_*(X; \mathbb{F}_p) \geq \dim H_*(X^{\text{fix}}; \mathbb{F}_p).$$

Remark 6.2.5. The inequality requires us to work over \mathbb{F}_p . (add: explain why need \mathbb{F}_p not \mathbb{Z})

(add: citations, maybe) We give Borel’s proof of the classical Smith inequality for $p = 2$, working over \mathbb{F}_2 (as required). The proof for $p \geq 3$ is similar. Also, Borel’s proof uses singular homology; we will instead demonstrate the proof using CW-homology, which means it’s really only a proof for CW-complexes. Also, if you look at the pictures, this is kind of a ‘proof’ by example; but you can piece together the real proof from the written text if you wanted to. We basically follow the proof found on pages 5-6 of [LT16].

Our example to keep in mind is the following S^2 with the reflection involution τ :



We’ve chosen a cell decomposition that plays well with this involution. The fixed point set is a copy of S^1 with the cell structure shown on the right of the figure.

¹³P.A.Smith, back in the 1920s, extensively studied $\mathbb{Z}/p\mathbb{Z}$ actions on topological spaces.

The *Tate bicomplex* is the bi-infinite bicomplex obtained by basically tensoring the complex $(C_*(X), \partial)$ with the localized polynomial ring $\mathbb{F}_p[\theta, \theta^{-1}]$, and then adding horizontal differentials $1 + \tau$:

$$\begin{array}{ccccccc}
 \xrightarrow{1+\tau} & C_2(X) & \xrightarrow{1+\tau} & C_2(X) & \xrightarrow{1+\tau} & C_2(X) & \xrightarrow{1+\tau} \\
 & \downarrow \partial & & \downarrow \partial & & \downarrow \partial & \\
 \xrightarrow{1+\tau} & C_1(X) & \xrightarrow{1+\tau} & C_1(X) & \xrightarrow{1+\tau} & C_1(X) & \xrightarrow{1+\tau} \\
 & \downarrow \partial & & \downarrow \partial & & \downarrow \partial & \\
 \xrightarrow{1+\tau} & C_0(X) & \xrightarrow{1+\tau} & C_0(X) & \xrightarrow{1+\tau} & C_0(X) & \xrightarrow{1+\tau}
 \end{array}$$

The action of θ moves you to the next column over (to the right).

I'm abusing notation and letting τ denote both the topological involution as well as the involution on chains. This is because I want to use τ_* to denote the induced map on homology at some point.

The $^{vh}E^\bullet$ spectral sequence is very hard to compute, but it's easy to compute up to page 1:

- $(^{vh}E^0, d^0) =$

$$\begin{array}{ccccccc}
 \xrightarrow{1+\tau} & C_2(X) & \xrightarrow{1+\tau} & C_2(X) & \xrightarrow{1+\tau} & C_2(X) & \xrightarrow{1+\tau} \\
 & \downarrow \partial & & \downarrow \partial & & \downarrow \partial & \\
 \xrightarrow{1+\tau} & C_1(X) & \xrightarrow{1+\tau} & C_1(X) & \xrightarrow{1+\tau} & C_1(X) & \xrightarrow{1+\tau} \\
 & \downarrow \partial & & \downarrow \partial & & \downarrow \partial & \\
 \xrightarrow{1+\tau} & C_0(X) & \xrightarrow{1+\tau} & C_0(X) & \xrightarrow{1+\tau} & C_0(X) & \xrightarrow{1+\tau}
 \end{array}$$

The homology is just $H_*(X) \otimes \mathbb{F}_2[\theta, \theta^{-1}]$, which leads us to ...


- $(^{vh}E^1, d^1) =$

$$\begin{array}{ccccccc}
 \xrightarrow{(1+\tau)_*} & H_2(X) & \xrightarrow{(1+\tau)_*} & H_2(X) & \xrightarrow{(1+\tau)_*} & H_2(X) & \xrightarrow{(1+\tau)_*} \\
 & \nearrow (d^2) & & \nearrow (d^2) & & \nearrow (d^2) & \\
 \xrightarrow{(1+\tau)_*} & H_1(X) & \xrightarrow{(1+\tau)_*} & H_1(X) & \xrightarrow{(1+\tau)_*} & H_1(X) & \xrightarrow{(1+\tau)_*} \\
 & \nearrow (d^2) & & \nearrow (d^2) & & \nearrow (d^2) & \\
 \xrightarrow{(1+\tau)_*} & H_0(X) & \xrightarrow{(1+\tau)_*} & H_0(X) & \xrightarrow{(1+\tau)_*} & H_0(X) & \xrightarrow{(1+\tau)_*}
 \end{array}$$

I could think hard about what $(1+\tau)_*$ does, but this page is already interesting – it describes X . In any case, $^{vh}d^2$ looks pretty bad, so I'm going to stop computing this spectral sequence here.

The $^{hv}E^\bullet$ spectral sequence turns out to be easy to compute:

- $(^{vh}E^0, d^0) =$




The kernel of $d^0 = 1 + \tau$ consists of precisely two types of chains:

- cells that are literally fixed by τ (p, q, e, f), and
- orbits of cells that are not fixed by τ (A, B).

The second class are the image of $1 + \tau$. Therefore the homology of this page can be canonically identified with $C_*(X^{\text{fix}})$.


- Now what are the induced differentials ∂_* ? These are just the restriction of ∂ to the fixed cells, of course – the gluing maps definitely haven't changed! Let ∂^{fix} denote the CW differentials for $C_*(X^{\text{fix}})$.

So, $({}^{hv}E^1, d^1) =$



which is just $(C_*(X^{\text{fix}}), \partial^{\text{fix}}) \otimes \mathbb{F}_2[\theta, \theta^{-1}]$.

- Hence ${}^{hv}E^2 \cong H_*(X^{\text{fix}}) \otimes \mathbb{F}_2[\theta, \theta^{-1}]$. We now need to reckon with the d^2 differentials shown in gray in the previous figure. These would have been induced by zigzags of the form



For such a length-2 arrow to exist (i.e. have nonzero coefficient) on page 2, then both w and z needed to have survived to page 2.

A subtlety here is that, technically speaking, we are thinking of w as a homology class in $H_*(X^{\text{fix}})$. However, because we have been Gaussian eliminating this whole time, it's also ok to think of w as a generator (CW cell) that survived to page 2; to be more precise, a chain representative \tilde{w} of w survived to page 1. But we will abuse notation and just call this chain w , since Gaussian elimination picks out a (noncanonical) representative / generator in a new basis.

Since w survived to page 2, we know that it's (1) a fixed cell under τ and (2) in the kernel of ∂^{fix} . But then $\partial w = 0$ as well, so the first step of this zigzag is 0.

There are no more longer differentials!

- So the spectral sequence collapses, and we know that

$${}^{hv}E^2 \cong H_*(X^{\text{fix}}; \mathbb{F}_2) \otimes \mathbb{F}_2[\theta, \theta^{-1}].$$

Furthermore, this is also isomorphic to both ${}^{hv}E^\infty$ and the homology of the total complex, $H_*(\text{Tot}(\mathcal{C}, \partial + (1 + \tau)))$.

Since ${}^{vh}E^\bullet$ also computes the homology of the total complex, we now know that this is a spectral sequence

$$H_*(X; \mathbb{F}_2) \otimes \mathbb{F}[\theta, \theta^{-1}] \rightrightarrows H_*(X^{\text{fix}}; \mathbb{F}_2) \otimes \mathbb{F}[\theta, \theta^{-1}].$$

Now since every page turn of spectral sequence is a homology computation, the rank of the $\mathbb{F}_2[\theta, \theta^{-1}]$ -module on each page must be monotone decreasing.

After convincing ourselves that the rank of $H_*(Y; \mathbb{F}_2) \otimes \mathbb{F}[\theta, \theta^{-1}]$ as an $\mathbb{F}[\theta, \theta^{-1}]$ -module is the same as the dimension of $H_*(Y; \mathbb{F}_2)$ over \mathbb{F}_2 , we recover the classical Smith inequality.

6.2.2. Khovanov homology and periodic links. We now use the same technique to prove a similar Smith inequality for the Khovanov homology of knots with a type of $\mathbb{Z}/2\mathbb{Z}$ symmetry. The reference for this section is [Zha18]; a more general reference is [SZ24], but this uses techniques we have yet to discuss in this course.

Let p be a prime. Consider a link $\tilde{L} \subset S^3$ together with an orientation-preserving diffeomorphism $\sigma : S^3 \rightarrow S^3$ of order p that preserves \tilde{L} set-wise.

The *Smith conjecture* (named after P.A.Smith, and now a theorem), proven by (add: citations), tells us that the fixed point set σ must either be (1) empty or (2) an unknotted fixed-point axis.

If σ has an unknotted axis and \tilde{L} is disjoint from this axis, then we say that \tilde{L} is a *p-periodic link*.

Here are some examples of different p -periodic links:

(add: pictures)

In 2010, Seidel and Smith¹⁴ proved the following Smith-like inequality for symplectic Khovanov homology:

Theorem 6.2.6 ([SS10]). Let \tilde{L} be a 2-periodic link, and let L denote the quotient link. Then

$$\dim \text{Kh}_{\text{symp}}(\tilde{L}; \mathbb{F}_2) \geq \text{Kh}_{\text{symp}}(L; \mathbb{F}_2).$$

While Abouzaid–Smith proved that Kh_{symp} agrees with δ -graded Kh over characteristic 0, we still don’t know if there is an isomorphism over \mathbb{F}_2 . So this result does not tell us if there’s a similar dimension inequality for regular Kh .


Let’s see what happens when we use Borel’s technique. Our goal is to relate $\text{Kh}(\tilde{L})$ and $\text{Kh}(L)$ (all over \mathbb{F}_2), by hoping to see $\text{Kh}(L)$ on some page of a spectral sequence built from $\text{CKh}(\tilde{L})$.

Here is the Khovanov–Tate complex for a two-periodic diagram \tilde{D} :

$$\begin{array}{ccccccc} & \uparrow & & \uparrow & & \uparrow & \\ \rightarrow & \text{CKh}^{i+1}(\tilde{\mathfrak{b}}) & \xrightarrow{1+\tau} & \text{CKh}^{i+1}(\tilde{\mathfrak{b}}) & \xrightarrow{1+\tau} & \text{CKh}^{i+1}(\tilde{\mathfrak{b}}) & \rightarrow \\ & \uparrow d_{\text{Kh}} & & \uparrow d_{\text{Kh}} & & \uparrow d_{\text{Kh}} & \\ \rightarrow & \text{CKh}^i(\tilde{\mathfrak{b}}) & \xrightarrow{1+\tau} & \text{CKh}^i(\tilde{\mathfrak{b}}) & \xrightarrow{1+\tau} & \text{CKh}^i(\tilde{\mathfrak{b}}) & \rightarrow \\ & \uparrow d_{\text{Kh}} & & \uparrow d_{\text{Kh}} & & \uparrow d_{\text{Kh}} & \\ \rightarrow & \text{CKh}^{i-1}(\tilde{\mathfrak{b}}) & \xrightarrow{1+\tau} & \text{CKh}^{i-1}(\tilde{\mathfrak{b}}) & \xrightarrow{1+\tau} & \text{CKh}^{i-1}(\tilde{\mathfrak{b}}) & \rightarrow \\ & \uparrow & & \uparrow & & \uparrow & \end{array}$$

For concreteness, let’s consider a specific example. Let \tilde{D} denote the following 2-periodic diagram for the Hopf link, along with its cube of resolutions:

¹⁴a different, much younger Smith



Observe that there is a discrepancy! The dashed component does not correspond to an $^{hv}d^2$ differential, whereas the other two do.

So what we *actually* have is the *annular* Khovanov differential, and you can check this in the following exercise.

- Exercise 6.2.7.** (a) For each of the six possible Khovanov differentials, in the presence of the axis of symmetry ($\mathbb{V} \otimes \mathbb{V} \rightarrow \mathbb{W}$, $\mathbb{V} \otimes \mathbb{W} \rightarrow \mathbb{V}$, etc.), compute the induced $^{hv}d^2$ differentials, and verify that they correspond to the annular Khovanov differential, under the identification of generators $^{hv}E^2 \cong \text{CKh}(D) \otimes \mathbb{F}_2[\theta, \theta^{-1}]$ discussed above.
- (b) Verify the same is true if you build the *annular Khovanov-Tate* complex, i.e. use d_{AKh} as the vertical differentials in the bicomplex.

Using grading arguments, it's possible to show that for the annular Khovanov-Tate complex, the $^{hv}E^\bullet$ spectral sequence collapses at page 3, and obtain the following theorem:

Theorem 6.2.8 ([Zha18], Theorem 1). There is a spectral sequence

$$\text{AKh}(\tilde{L}; \mathbb{F}_2) \otimes \mathbb{F}_2[\theta, \theta^{-1}] \rightrightarrows \text{AKh}(L; \mathbb{F}_2) \otimes \mathbb{F}_2[\theta, \theta^{-1}]$$

which implies the Smith-type inequality

$$\dim \text{AKh}(\tilde{L}; \mathbb{F}_2) \geq \dim \text{AKh}(L; \mathbb{F}_2).$$

The actual theorem is more refined, and includes some grading information.

However, the same technique does not work for the Khovanov-Tate complex. In order to prove the corresponding theorem for $\text{Kh}(\tilde{L})$, and also to generalize this to all prime periodicities, we needed to use the Lipshitz–Sarkar Khovanov stable homotopy type, which we will discuss later in the course.

6.3. Khovanov-to-Floer spectral sequences. While this is not a course about Floer homologies, the relationship between categorified quantum invariants and Floer-type invariants should not be understated, for at least two reasons. First of all, the group of people studying these two types of invariants often overlap; these tools are often used in conjunction in low-dimensional topology, for example. Second, both ultimately have a homological flavor, and algebraic tricks can often be emulated on either side.

Example 6.3.1. Recall that earlier in the course we mentioned the following wide-open problem:

Does the Jones polynomial detect the unknot?

Kronheimer–Mrowka proved that *Khovanov homology* in fact *does* detect the unknot [?]. The proof has two steps:

- (1) Find a spectral sequence from Khovanov homology to instanton Floer homology
- (2) Prove that instanton Floer homology detects the unknot.

Prove to yourself that this completes the proof!

The paper is of course very involved and very long. The two steps are very hard to prove!

Ozsváth–Szabó’s *Heegaard Floer homology* (add: cite original holo disks 3-manifolds), and invariant of 3-manifolds, is most closely intertwined with Khovanov homology; both were discovered around the same time and were studied and expanded upon by the same cohort of grad students! (These grad students later became PhD advisors to my generation!)

In this section, we will focus on the first concrete connection between the two theories, a spectral sequence discovered by Ozsváth–Szabó

$$(11) \quad E^2 \cong \widetilde{\text{Kh}}(\bar{L}) \rightrightarrows E^\infty \cong \widehat{\text{HF}}(\Sigma_2(S^3, L))$$


relating the *reduced Khovanov homology* of (the mirror of) a link L to the Heegaard Floer homology of the branched double cover of S^3 along L [OS05].

We will explain all the components of this statement at various levels of detail.

We first give a very, very, almost criminally elementary introduction to Heegaard Floer homology. See [?], [?], and [?] for good introductory readings on Heegaard Floer homology. (add: add also) <https://users.math.msu.edu/users/heddenma/Lectures%20on%20Heegaard%20Floer%20homology.pdf>

6.3.1. *Heegaard diagrams for closed, connected, orientable 3-manifolds.* We will work with the most basic flavor of Heegaard Floer homology and only concern ourselves with the simplest case of 3-manifolds.

Notation 6.3.2. This is how I draw a 2-sphere (left) vs a 3-ball (right):




Let Y be a closed, connected, orientable 3-manifold. Let $f : Y \rightarrow \mathbb{R}$ be a self-indexing Morse function on Y ; that is, the index i critical points are all in the critical level set $f^{-1}(i)$.

Observe the following:

- We can arrange so that there are unique index-0 and index-3 critical points.
- By Poincaré duality,

$$\# \text{ index-1 critical points} = \# \text{ index-2 critical points}.$$
- If there are g index-1 critical points, (and g index-2 critical points), then the level set $f^{-1}(3/2)$ is a closed, connected, orientable surface with genus g .


In this context, the surface $f^{-1}(3/2)$ splits Y into two handlebodies. This decomposition is called a *Heegaard splitting* for Y , and the surface $f^{-1}(3/2)$ is called the *Heegaard surface* for this splitting.



A *Heegaard diagram* \mathcal{H} for a Y is a combinatorial model for a Heegaard splitting, and consists of the following data:

- a genus g surface Σ_g ;
- a set of g **red** α -curves that each bound disks in the $f^{-1}([0, 3/2])$ handlebody, and such that after cutting along these disks, the result is a 3-ball;
- a set of g **blue** β -curves that each bound disks in the $f^{-1}([3/2, 3])$ handlebody, and such that after cutting along these disks, the result is a 3-ball.


Our running example will be this genus-2 Heegaard diagram for S^3 :





Given a Heegaard diagram \mathcal{H} for Y , we know how to build Y :

- (1) Start with a thickening of Σ_g .
- (2) On inside, glue on thickened red disks, along the α curves.
- (3) On the outside, glue on thickened blue disks, along the β curves.
- (4) Fill in the remaining two ball-shaped holes on each side with two 3-balls.


Starting with a thickened version of Σ_2 , we can build the red handlebody by gluing in the disks on the inside with the red curves as boundary:



Then, we can build the remainder of S^3 by attaching the blue disks on the outside, whose boundaries are the blue curves to get a ball, and then gluing to this B^3 another B^3 :



We can relate this construction (which started with the thickened Σ_2 with the handle decomposition given by our Morse function as follows:




Remark 6.3.3. This corresponds to a handle decomposition of Y with

- one 0-handle: the ball on the inside
- g 1-handles: thickened version of the part of the ascending manifold for each index-1 critical point, in $f^{-1}([1, \frac{3}{2}])$
- g 2-handles: thickened version of the part of the descending manifold for each index-2 critical point, in $f^{-1}([\frac{3}{2}, 2])$
- one 3-handle: the ball on the outside.

Example 6.3.4. Let's now consider $S^1 \times S^2$. This manifold is basically obtained by gluing two solid tori $S^1 \times D^2$ together so that, for each $\theta \in S^1$, the disks $\{\theta\} \times D^2$ on the two solid tori are glued together along their boundary to form a S^2 .

Here is a genus-1 Heegaard splitting for $S^1 \times S^2$:



The red and blue curves bound disks on the inside and outside. If you slide the blue disk so that its boundary is the same as the red curve, then you'll clearly see an S^2 with a red side and a blue side.

6.3.2. Heegaard Floer homology basics. Heegaard Floer homology was originally defined by Ozsváth and Szabó, in [?]. There are many different version now, but today we will discuss only the most basic version, the 'hat' flavor $\widehat{\text{HF}}$, for closed, connected, orientable 3-manifolds Y .

Here is a summary of the main steps:

- (1) Start with a Heegaard diagram for Y . Pick a basepoint z and perturb the α and β curves so that the diagram is *admissible* (see Remark 6.3.6).
- (2) Build a symplectic manifold from Σ_g , as well as two Lagrangian submanifolds \mathbb{T}_α and \mathbb{T}_β from the data of the α - and β -curves, respectively.
- (3) The chain complex is generated as an \mathbb{F}_2 vector space by the points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$.
- (4) The differential counts moduli spaces of holomorphic disks with boundary on $\mathbb{T}_\alpha \cup \mathbb{T}_\beta$, that do not intersect the divisor defined by $z \in \Sigma_g$.

For a genus g surface Σ_g , $\text{Sym}^g(\Sigma_g)$ is defined as the quotient of the g -fold Cartesian product of copies of Σ_g by the obvious action of the symmetric group S_g :

$$\text{Sym}^g(\Sigma_g) = \Sigma_g \times \cdots \times \Sigma_g / S_g.$$

Observe that the points in $\text{Sym}^g(\Sigma_g)$ are basically *unordered* g -tuples of points in Σ_g .

Now Σ_g has a complex structure, so $\Sigma_g^{\times g}$ does as well. What's surprising is that, while $\text{Sym}^g(\Sigma_g)$ is a priori only an orbifold (the action of S_g is obviously not free), it's in fact smooth, and moreover has an induced complex

structure! This makes $\text{Sym}^g(\Sigma_g)$ into a symplectic manifold as well, with symplectic form induced by the complex structure.

Inside $\text{Sym}^g(\Sigma_g)$ are two totally real tori

$$\begin{aligned}\mathbb{T}_\alpha &= \alpha_1 \times \cdots \times \alpha_g / \sim \\ \mathbb{T}_\beta &= \beta_1 \times \cdots \times \beta_g / \sim\end{aligned}$$


which are Lagrangian with respect to the symplectic structure.

Heegaard Floer homology is then the Lagrangian intersection Floer homology of $(\text{Sym}^g(\Sigma_g), \mathbb{T}_\alpha, \mathbb{T}_\beta)$...

... More concretely, here's how you think about the generators and the differential.

The generators of the chain complex are the points $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. In other words, these are unordered g -tuples of intersection points between α and β curves.

The coefficient of the differential from one intersection point x to another y is determined by counting the size of the moduli space of holomorphic disks from x to y of the following form:




Officially, here is the definition of the differential:

$$\partial x = \sum_{y \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\phi \in \pi_2(x, y), \mu(\phi)=1, n_z(\phi)=0} \# \overline{\mathcal{M}(\phi)} y.$$

- $\pi_2(x, y)$ is the set of homotopy classes of Whitney disks from x to y as shown in the above drawing.
- $\mathcal{M}(\phi)$ is the moduli space of holomorphic representatives of ϕ ; these disks have an \mathbb{R} action by parametrization (complex automorphisms of the unit disk in \mathbb{C} that preserve $\pm i$)
- $\mu(\phi)$ is the Maslov index of ϕ , the ‘expected dimension’ of the smooth, real manifold $\mathcal{M}(\phi)$; [Lipshitz gave more explicit formulas for \$\mu\(\phi\)\$ in his cylindrical reformulation for Heegaard Floer homology.](#)
- $\overline{\mathcal{M}(\phi)} = \mathcal{M}(\phi)/\mathbb{R}$ is a finite set of points, assuming that $\mu(\phi) = 1$.

For our purposes, it is necessary for you to really understand only the following examples.

Example 6.3.5. Our first example is the genus-1 Heegaard diagram for $S^1 \times S^2$ from Example 6.3.4, but now perturbed so that we actually have intersection points between the α and β curves, and with a basepoint slapped on the complement of the α and β curves:



In this $g = 1$ example, we are lucky that $\text{Sym}^g(\Sigma_g) = \Sigma_g$, and the drawn α and β curves are literally Lagrangian. What you see is what you get, so we can explicitly determine the chain complex.

There are two intersection points, x and y , as shown in the Heegaard diagram \mathcal{H} drawn above. Therefore

$$\widehat{\text{CF}}(\mathcal{H}) = \mathbb{F}_2 x \oplus \mathbb{F}_2 y.$$

We now compute the differential; recall that we are working over \mathbb{F}_2 :

- $\partial x = 0$: There are two holomorphic disks from x to y ; they are shaded in purple and yellow.
- $\partial y = 0$: There are no holomorphic disks from y to x (the disks from computing ∂x have the wrong orientation!).

Therefore $\widehat{\text{HF}}(S^1 \times S^2) \cong \mathbb{F}_2^2$.

It turns out that $\widehat{\text{HF}}(Y)$ enjoys an action by $H_1(Y; \mathbb{Z})/\text{tor}$. For $S^1 \times S^2$, either generator θ of $H_1(Y; \mathbb{Z}) \cong \mathbb{Z}$ sends $x \mapsto y$.


Remark 6.3.6. Admissibility of a Heegaard diagram ensures that we can actually build a chain complex, and that the invariant is well-defined. The actual definition of admissibility is somewhat involved, as the requirements really come from studying the spaces of holomorphic disks between generators.

Our original diagram for $S^1 \times S^2$ is *not* admissible:



We needed to perturb the β curve so that, if we cut Σ_1 along both α and β , we get only disks, aside from the piece containing the basepoint.

Example 6.3.7. Let's do another example: $Y = (S^1 \times S^2) \# (S^1 \times S^2)$. Here is a Heegaard diagram:




There are four generators: $\{x_1, y_1\} \times \{x_2, y_2\}$.

Consider first the space of disks from (x_1, x_2) to (y_1, x_2) . There are two, given precisely by the product of the disks from Example 6.3.5 from x_1 to y_1 and the constant disk at x_2 .

After working through all four differentials (and remembering that we're working over \mathbb{F}_2 , you should find that $\partial = 0$, and each of the four generators represent a distinct homology class.

The H_1 action is likewise induced by the H_1 action on the two 'sides' of the Heegaard diagram. Fix generators θ_1 and θ_2 of the H_1 of the left and right copies of $S^1 \times S^2$, respectively. Then the $H_1(Y)$ action is given by



This should start to remind you of Khovanov homology...

Heegaard Floer homology behaves very well under connected sum:

$$\widehat{\text{HF}}(Y \# Y') \cong \widehat{\text{HF}}(Y) \otimes \widehat{\text{HF}}(Y').$$

If \mathcal{H} is a Heegaard diagram for Y , and \mathcal{H}' is a Heegaard diagram for Y' , then the connected sum of \mathcal{H} and \mathcal{H}' in the regions containing the basepoints is a Heegaard diagram for $Y \# Y'$. You can now check that, at least on the level of chain groups, it is perhaps believable that

$$\widehat{\text{CF}}(\mathcal{H} \# \mathcal{H}') \cong \widehat{\text{CF}}(\mathcal{H}) \otimes \widehat{\text{CF}}(\mathcal{H}').$$

And finally, because the basepoint in $\mathcal{H} \# \mathcal{H}'$ is in the connect-sum region, no holomorphic disks pass 'between' the two sides. More precisely, if \mathbf{x} and \mathbf{x}' are tuples of intersection points for \mathcal{H} and \mathcal{H}' , respectively, then $(\mathbf{x}, \mathbf{x}')$ (obtained by concatenating the tuples) is a tuple of intersection points for $\mathcal{H} \# \mathcal{H}'$. Define another tuple $(\mathbf{y}, \mathbf{y}')$ similarly. Then holomorphic disks from $(\mathbf{x}, \mathbf{x}')$ to $(\mathbf{y}, \mathbf{y}')$ are in bijection with

$$\{\text{hol. disks } \phi : \mathbf{x} \rightarrow \mathbf{y}\} \times \{\text{hol. disks } \phi' : \mathbf{x}' \rightarrow \mathbf{y}'\}$$

by taking the diagonal disk in $\phi \times \phi'$.

Moreover, the action of $H_1(Y \# Y'; \mathbb{Z})/\text{tor} \cong H_1(Y; \mathbb{Z})/\text{tor} \otimes H_1(Y'; \mathbb{Z})/\text{tor}$ is induced by the actions of $H_1(Y; \mathbb{Z})/\text{tor}$ and $H_1(Y'; \mathbb{Z})/\text{tor}$ on the tensor factors of $\widehat{\text{HF}}(Y) \otimes \widehat{\text{HF}}(Y')$, respectively.

Example 6.3.8. We now know how to compute $\widehat{\text{HF}}(\#^k S^1 \times S_2)$. And we also understand how H_1 acts.

6.3.3. Branched double covers. In Example 6.3.7, we saw an inkling of how Khovanov homology (for unlinks, at least) might be related to $\widehat{\text{HF}}$ (for connected sums of $S^1 \times S^2$, at least).

We now explain the topological relationship between planar circles and connected sums of $S^1 \times S^2$, by discussing branched double covers.


Let M be a manifold, and let $L \subset M$ be a submanifold. The k -fold branched cover of M along branching locus L is a manifold $\Sigma_k(M)$ with a (continuous) map $\pi : \Sigma_k(M) \rightarrow M$ such that

- for each point $p \in M \setminus L$, $\pi^{-1}(p)$ consists of k copies of p (a covering map), and
- for each point $p \in L$, $\pi^{-1}(p)$ is just one copy of p .

A prototypical example is the branched covering map $\pi : \mathbb{C} \rightarrow \mathbb{C}$ given by $z \mapsto z^2$. In the following example, we describe a homeomorphic version of this branched covering, and how to construct the 2-fold branched cover, also known as the *branched double cover*.


The typical notation for the branched double cover of M along L is $\Sigma_2(M, L)$. However, to avoid notation collision with a genus-2 Heegaard surface, we will instead just write Σ to denote 'branched double cover'.

Example 6.3.9. Let D be a disk, and let p be a point in the interior of the disk:



We wish to build $\Sigma(D, p)$.

We first make a branch cut (right side of above figure), and then take two copies of our snipped disk:




We now glue $1 \rightarrow 2'$ and $2 \rightarrow 1'$ so that the resulting surface F has two lifts of each point other than p (including the points along the branch cut – notice that these were doubled already when we made the cut, so it's ok that we then glued them to their twins on the second copy). There is, however, only one lift of the branch point p .

We can visualize gluing the two shown disks above to get one very floppy disk.


We can also think of F as the ramp in a parking garage where you can climb from the 1st to the 2nd floor, but then when you keep climbing, you end up back at the 1st floor.

Example 6.3.10. Now consider a disk with two branch points, red and blue:




We again make some branch cuts, and make two copies of our snipped disk. (These cuts are just my choice, you could also decide to make a cut that connects the red and blue. However, I find the drawn visualization easier to explain.)

When we glue the two pieces together by $1 \rightarrow 2'$, $2 \rightarrow 1'$; $3 \rightarrow 4'$, $4 \rightarrow 3'$, we get an annulus:




So $\Sigma(D, 2 \text{ interior points}) \cong A$, an annulus.


Example 6.3.11. Now consider the branched double cover of two arcs inside a 3-ball. This is homeomorphic to the thickening of the previous example by an interval I :



We make the same branch cuts, but now thickened by I as well:



Following the previous example, we make two copies and glue them together. This shows that the branched double cover of a sphere with branch locus given by two properly embedded arcs is a solid torus:



Example 6.3.12. Now we are ready to find the branched double cover of a 2-component unlink inside S^3 :

$$\Sigma(S^3, \text{red circle} \cup \text{blue circle})$$

The trick is to first cut S^3 into two balls as shown:



Using the previous example, we have $\Sigma(S^3, \text{red circle} \cup \text{blue circle}) \cong$

$$= \Sigma(S^3, \text{red circle} \cup \text{blue circle}) \cup \Sigma(S^3, \text{red circle} \cup \text{blue circle})$$

inside ball ↗
outside ball ↗

$$= \text{Diagram 1} \cup \text{Diagram 2}$$

Noting that the red endpoints need to glue to the red endpoints, and similarly for the blue endpoints, we reason that the result is

$$\Sigma(S^3, \text{red circle} \cup \text{blue circle}) \cong S^1 \times S^2.$$

6.3.4. Reduced Khovanov homology. In Example 6.3.7, I hinted at the fact that $\widehat{\text{HF}}((S^1 \times S^2) \# (S^1 \times S^2))$ looks similar to $\text{Kh}(\text{red circle} \cup \text{blue circle}) = V \otimes V$. In the previous section, we hinted that the relationship between $\#^k S^1 \times S^2$ and unlinks.

But notice that there is a discrepancy:

$$\text{Kh}(\text{red circle} \cup \text{blue circle}) \cong \widehat{\text{HF}}((S^1 \times S^2) \# (S^1 \times S^2))$$

but

$$\Sigma(\text{red circle} \cup \text{blue circle}) \cong S^1 \times S^2.$$

So, $\text{Kh}(\text{red circle} \cup \text{blue circle}) \not\cong \widehat{\text{HF}}(\Sigma(S^3, \text{red circle} \cup \text{blue circle}))!$

This is one motivation for defining and studying *reduced Khovanov homology*. This is a version of Khovanov homology defined for *based* links, i.e. links with a chosen basepoint:




Definition 6.3.13. Let $(D, *)$ be a based diagram for based link (L, p) .

The *reduced Khovanov complex* for $(D, *)$ is


$$\widetilde{\text{CKh}}(D, *) := (\text{CKh}(D)/\text{im}(\xi)) \{-1\}$$

where ξ is the chain map that merges a small circle labeled X into the arc containing the basepoint:



In other words, the $\widetilde{\text{CKh}}(D, *)$ complex is the quotient complex you get from $\text{CKh}(D)$ by erasing all the generators that label the based component by X ; the label on the component in a complete resolution D_u containing the based point *must* be 1.

Example 6.3.14. Here is part of the chain complex for the based trefoil:



The highlighted region is part of the subcomplex $\text{im}(\xi)$.

Exercise 6.3.15. (a) Determine the reduced Khovanov homology of the unlink with k components.
 (b) Compute the reduced Khovanov homology of the trefoil. *Feel free to work over \mathbb{F}_2 . Remember to incorporate the quantum shift in the definition!*

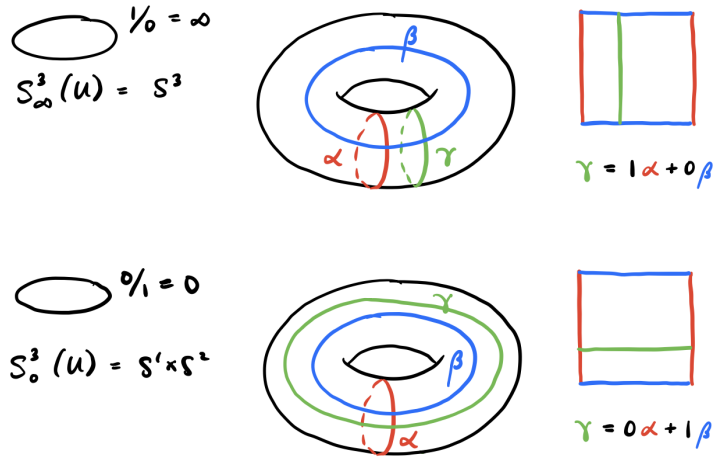
Fact 6.3.16. For knots, the isomorphism class of $\widetilde{\text{Kh}}$ does not depend on the location of the basepoint.

6.3.5. *Dehn surgery.* *Dehn surgery* is a method for building 3-manifolds. Here's the general idea:

- (1) Pick a knot $K \subset S^3$ and a framing of the knot, i.e. a homologically nontrivial circle γ on the boundary of a neighborhood of K (a longitude for a tubular neighborhood of K).
- (2) Drill out the solid donut that is a neighborhood of K . Now, glue it back in but so that the γ bounds a disk for the inside solid torus.

Lickorish [Lic62] and Wallace [Wal60] proved that any closed, connected, orientable 3-manifold can be obtained in this way.

There are varying conventions in the literature, and we will stay agnostic about these conventions. Everyone agrees on the following two examples, though:



In the Heegaard diagrams above, $(\Sigma_1, \alpha, \beta)$ is a Heegaard diagram for S^3 . If we replace α with γ , then $(\Sigma_1, \gamma, \beta)$ is the Heegaard diagram for the manifold we get after surgery.

The ∞ -surgery of S^3 along the unknot U is just S^3 , since γ is parallel to α .

The 0 -surgery instead yields $S^1 \times S^2$, because now, on the inside of the torus, γ is supposed to bound a disk.

6.3.6. The Ozsváth–Szabó spectral sequence. We are now ready to describe, in very broad strokes, the Ozsváth–Szabó spectral sequence (11).

Recall that we have already arranged so that for the k -component unlink U^k ,


$$\widetilde{\text{Kh}}(U^k) \cong \widehat{\text{HF}}(\Sigma(S^3, U^k)).$$

over \mathbb{F}_2 coefficients.

In order to boost this relationship to links other than the unlink, we use the same cube-of-resolutions idea from Khovanov homology: we essentially use an iterated mapping cone construction. The result will be a spectral sequence rather than an immediate identification.


If we see a neighborhood of a crossing in L , we can associate the 0 and 1 resolutions as shown below, with conventions reversed from Khovanov homology (hence the mirroring in (11)).

In the drawing below, imagine that L has the red crossing. Let L_0 be the link with the red crossing replaced with the blue arcs, and let L_1 be the link with the red crossing replaced with the green arcs.



These have a cyclic relationship, because they can be viewed as the three pairs of opposite edges of a tetrahedron.

We want to take the branched double cover of the ball along the red arcs first. To help with visualization, we will first rotate the inscribed tetrahedron (by rotating the ball) so that we see the diagram on the left:




We can then make a branch cut and isotop to get the solid cylinder in the middle, and note that once we build the branched double cover, the red arc will be contained in the red disk shown on the right.

In other words, if we branch along the red arcs, the parallel arcs on the boundary of the ball will, together with their double, bound a disk in the resulting solid torus.


A similar situation would occur if we took the branched double cover along blue or green, of course. On the right, we've faintly highlighted the boundary-parallel copies of the blue and green curves.

After taking the branched double cover of the ball along the red arcs, we have the following solid torus, where we now only draw the boundary-parallel copies of the original red, blue, and green arcs (actually, just one copy):



We only need to keep track of one copy of, say, the red circle, because the other copy will be parallel.

Let's study the Heegaard diagram on the right more carefully. In fact, let's first isotop the green curve so that its intersections with red and blue are clearer:



In the diagram above, we've labeled the red, blue, and green curve as α , β , and γ , and given them orientations so that their pairwise algebraic intersection numbers satisfy the following equation:

$$(12) \quad I(\alpha, \beta) = I(\beta, \gamma) = I(\gamma, \alpha) = -1.$$

A consequence of this cyclic relationship is that the three 3-manifolds $(\Sigma(S^3, L) = \Sigma(S^3, L_\infty), \Sigma(S^3, L_0), \Sigma(S^3, L_1))$ form a *triad*; Ozsváth–Szabó show that there is then a *surgery exact triangle* (i.e. long exact sequence) relating their Heegaard Floer homologies:

$$\begin{array}{ccc} \widehat{\mathrm{HF}}(\Sigma(S^3, L_0)) & \xrightarrow{\quad} & \widehat{\mathrm{HF}}(\Sigma(S^3, L_1)) \\ & \nwarrow \quad \swarrow & \\ & \widehat{\mathrm{HF}}(\Sigma(S^3, L_\infty)) & \end{array}$$

This reveals how we can use the mapping cone construction: to compute the bottom object, we can ‘replace’ it with the cone of the top row. [This is mostly analogous to how the Khovanov bracket works; see Remark 6.3.17.](#)

Remark 6.3.17. Ozsváth–Szabó show that for a triad (Y, Y_0, Y_1) ,


$$\widehat{\mathrm{CF}}(Y) \simeq \mathrm{Cone} \left(\widehat{\mathrm{CF}}(Y_0) \xrightarrow{\hat{f}} \widehat{\mathrm{CF}}(Y_1) \right)$$

where \hat{f} is the appropriate surgery map, and \simeq means ‘quasi-isomorphic’, which is the same as ‘chain homotopic’ over a field. Unlike in Khovanov homology, there will likely be a nontrivial homotopy that realizes this homotopy equivalence.

Ozsváth–Szabó’s key proposition tells us that these surgery maps actually correspond to the (reduced) Khovanov merge and split maps!

So, on E^0 of spectral sequence, we see a cube with $\widehat{\mathcal{CF}}$ of various connected sums of $S^1 \times S^2$ at the vertices. The differential d^0 computes the Heegaard Floer homology for these very simple 3-manifolds; we can identify E^1 with the $\widehat{\text{Kh}}$ cube, with unlinks at each vertex. On the next page, d^1 corresponds to the $\widehat{\text{Kh}}$ differential, by the key proposition alluded to in the previous paragraph. Therefore $E^2 \cong \widehat{\text{Kh}}(\bar{L})$.

We can view the remainder of the spectral sequence as a consequence of Remark 6.3.17. Due to the nontrivial homotopy equivalences, we have to add homotopies of all cube-degree lengths to the components of d^1 in order to truly recover the Heegaard Floer chain complex:



So, the spectral sequence is really a filtration spectral sequence with respect to this cube degree.

7. KHOVANOV STABLE HOMOTOPY TYPE

In this course, we first carefully went through a definition of Khovanov homology and how one works with it. Then, we went on a tour of various applications of Khovanov homology. We now slow down (sort of) once again, to study a topic relatively carefully. But don't be fooled: our four-lecture journey through the construction of a Khovanov stable homotopy type is meant only as a primer, to help you read the papers in case you choose to dig deeper.


We began this course by defining the Jones polynomial, and introduced Khovanov homology as a higher-level version of the invariant: while the Jones polynomial can only describe the knot embedded in S^3 , the Khovanov invariant, as a TQFT, can also be used to describe cobordisms between knots.

In this section, we will see how Lipshitz–Sarkar further boost the Khovanov homology invariant to a ‘space-level’ invariant \mathcal{X}_{Kh} . Roughly speaking, for a link $L \subset S^3$, $\mathcal{X}_{\text{Kh}}(L)$ is a sufficiently high-dimensional space whose singular cohomology agrees with $\text{Kh}(L)$, after an appropriate shift in homological grading.


Just as the Khovanov invariant can be thought of as a chain homotopy equivalence class of chain complexes, the Lipshitz–Sarkar invariant is really a *stable homotopy type*, rather than an actual particular space. (The ‘stable’ part refers to the fact that it is an invariant once you reach a sufficiently high dimension.)

We first begin with some motivation, starting with a brief introduction to / recollection of Morse theory. A standard reference for Morse theory is [Mil63]; if you want a quick overview, I am willing to bet that there are many, many video resources online covering Morse theory.


7.1. The idea of Morse homology.



Recall that we can build the torus $S^1 \times S^1$ using one 0-handle, two 1-handles, and one 2-handle:



As a CW-complex, we can also build it using one 0-cell, two 1-cells, and one 2-cell:



The CW chain complex $\mathbb{Z} \rightarrow \mathbb{Z} \oplus \mathbb{Z} \rightarrow \mathbb{Z}$ is generated by the cells:

$$\langle F \rangle \rightarrow \langle a, b \rangle \rightarrow \langle v \rangle.$$

Based on the attaching maps shown in the drawing, we have

$$\begin{aligned}\partial F &= a + b - a - b = 0 \\ \partial a &= \partial b = v - v = 0\end{aligned}$$

so $\partial = 0$.


We will now describe a third way to compute the same homology, by looking at flow lines dictated by a sufficiently generic height function, called a *Morse function*.

Lean T against a wall, and imagine dripping chocolate from the point at the top:



At most points on the torus, chocolate flows. However, at the four critical points of the height function, the chocolate ‘pools’. Two are shown above; the other two are in the back of the torus, and you can’t see them because chocolate is opaque.

Here is a better picture; the two critical points on the back side of the torus are shown as x ’s:



There are three types of critical points shown.

- At the green critical point, if you shift slightly away in any direction, the chocolate flows away. More formally, we say that the *descending manifold*, a.k.a. *unstable manifold*, of the green critical point is 2-dimensional.
- At each of the red critical points, shifting in two different directions might yield two different results: in one direction, the chocolate will flow back to the critical point; in another, the chocolate will flow away from the critical points. We say these critical points have a 1-dimensional descending / unstable manifold; they also have a 1-dimensional ascending / stable manifold.
- At the blue critical point at the bottom, any two directions you choose to shift in will yield chocolate flowing back into the critical point. So this point has a 0-dimensional descending / unstable manifold, and a 2-dimensional ascending / stable manifold.

The generators for the Morse homology complex are these critical points, and the homological grading of a critical point is given by the dimension of


its descending / unstable manifold. The homological grading of a critical point is usually called the *index* of the critical point.

Remark 7.1.1. Beware: We will unfortunately also use the term ‘index’ later to describe a difference between homological gradings, i.e. a homological degree.

A height function is said to be *Morse-Smale* if


- the critical points are isolated and
- the ascending and descending manifolds of different critical points intersect transversely.

In this case, the Morse chain complex is well-defined and yields the Morse homology. Here are some examples of height functions on the torus that are *not* Morse-Smale:



Our leaning-torus height function is Morse-Smale, so we can describe the differential. The Morse differential counts flows from index i critical points to index $i - 1$ critical points, with sign. This means that the critical points, which are indeed points, actually have a framing. You could imagine them as tiny ball creatures facing a particular direction, which is either $+$ or $-$.

Let's revisit the picture of the flow lines we saw before, but now with the generators labeled:




We see that out of F there are green and teal flow lines to index-1 critical points a and b , and they come in pairs, with opposite signs. For example, the two teal flow lines from F to b clearly oppose each other while entering b . There are ways to make all of this precise, of course! We unfortunately don't have enough time to stray off-course. Similarly, both a and b have pairs of oppositely-oriented flow lines into v . We conclude, just as we did in our CW-homology calculation, that the differential is 0.

Observe that, after choosing our Morse-Smale height function, we only actually needed the data of the 1-skeleton to complete our computation.

This should give you a sense of how much information is being lost when we compute homology from a topological space.

For example, we could have also thought about all the flow lines from F to v . There are four 1-parameter families of these flow lines, one of which is shaded in purple below:



The *moduli space* of flow lines from a critical point e^j of index j and a critical point e^i of index $i < j$ is a space $\mathcal{M}(e^j, e^i)$ where each point represents a single flow line; the identification is continuous, so that flow lines close to each other correspond to points in the moduli space that are close to each other. [Here is where we use the word ‘index’ again: these are flows of index \$j - i\$.](#)

In our situation, $\mathcal{M}(F, v)$ is the disjoint union of four intervals. The boundaries of each interval represent *broken flow lines*, e.g. the concatenation of a flow line from F to a with a flow line from a to v . (See the right-most drawing in the future above.)

Remark 7.1.2. The actual proof that $\partial^2 = 0$ in Morse homology essentially boils down to the fact that the moduli spaces of index 2 (i.e. flows of index 2) are one manifolds, and we really understand 1-manifolds.

All this is to say that, when we compute Morse homology, we are really only concerning ourselves with the data of moduli spaces of flows of index 1 (and you could maybe say 2, since we do want to check that $\partial^2 = 0$).

7.2. The spatial refinement problem. Now suppose I *start* with a chain complex of free \mathbb{Z} -modules (a free resolution!), and I wish to build a space whose Morse homology (or CW, singular, etc. homology) agrees with that of my chain complex.

The obvious way is to compute the homology, and then wedge together a bunch of spheres of the appropriate homological dimension. This method guarantees that your space will be no more interesting than your chain complex.

For example, had we start with the CW chain complex for the torus, we would have ended up with the space $S^0 \vee S^1 \vee S^1 \vee S^2$ (eww) which is definitely not our nice, smooth, beloved (not to mention *connected*) torus $S^1 \times S^1$.

Suppose we want to really build the space up with cells corresponding to the generators of the chain complex. The generators would tell us how many points to start with, and the differentials would tell us how to glue together the 1-skeleton. But at this point, we might still have 63 more dimensions to go — how do we glue on the faces (corresponding to the moduli spaces of index 2 flow lines), or the 3-dimensional pieces? What if we glued together our 2-skeleton so weirdly that it’s impossible to glue in a 3-ball where we want it?

This gluing information is what is lost when we take homology.

Remark 7.2.1. One should expect that this lost information is nontrivial. For example, in §6.2.2 we saw (asserted) that if a CW complex has a $\mathbb{Z}/2\mathbb{Z}$ action, then we can build the Tate bicomplex, and the ${}^{hv}E^\bullet$ spectral sequence is guaranteed to collapse by page 3. Perhaps you could imagine that you could build an abstract chain complex to place in the vertical columns of the Tate bicomplex, in such a way that the spectral sequence *does not*

collapse by page 3; then you'd know that the particular chain complex you chose could not have come from a $\mathbb{Z}/2\mathbb{Z}$ equivariant CW decomposition of a space!

While homology is just a graded \mathbb{Z} -module, a space has a *cohomology ring*, whose multiplication is given by the cup product. This means that two spaces can have the same homology, but different cohomology rings.

Lipshitz–Sarkar's Khovanov stable homotopy type is a great example of this phenomenon, though in the category of suspension spectra rather than topological spaces. The analogous operation to cup product operations $a \smile a$ for suspension spectra are *Steenrod squares* Sq^i . Seed showed by computing Steenrod squares that \mathcal{X}_{Kh} can distinguish between knots and Khovanov homology can't (add: cite)!

So indeed, the space-level refinement of Khovanov homology is a stronger invariant than Khovanov homology.

There are currently three methods for constructing $\mathcal{X}_{\text{Kh}}(L)$, which we outline briefly below. All three were shown to be equivalent (in the appropriate sense) by Lawson–Lipshitz–Sarkar (add: cite).

(1) **Morse flow category approach** by Lipshitz–Sarkar.

This is the original method Lipshitz–Sarkar used to construct the spectrum. The main idea is to use Cohen–Jones–Segal's proposed idea (add: cite) for capturing the full data needed to rebuild a space, by building a *framed flow category* that describes all the moduli spaces of flow lines, with framings, i.e. full gluing instructions.

This is the most hands-on method, and we will study this one in detail.

(2) **Burnside functor approach** by Lawson–Lipshitz–Sarkar.

This method is still quite hands-on, in the sense that some powerful theorems allow you to work with some very simple categories. The idea is to capture the information that a framed flow category would, but in a more abstract way. We can view the Khovanov complex as a functor from category $\underline{2}^n \rightarrow \text{Mod}_{\mathbb{Z}}$. We boost this to a functor to a version of the Burnside category \mathcal{B} , by making some careful choices. Using the Pontrjagin–Thom construction, this gives us a functor $\underline{2}^n \rightarrow \mathbf{Top}_*$ to based topological spaces. We then take the hocolim of this diagram of spaces, which gives us a (very high dimensional) space. Finally, suspend and desuspend to get a suspension spectrum.

We might touch upon this construction if we have time.

(3) **K -theory approach** by Hu–Kriz–Kriz.

This method feels the least hands-on to me, and I am actively trying to learn enough homotopy theory to be able to work with it. The basic idea is to build a permutative category which, again, contains all the framing information that we need to be able to build a space, and then pass this category through the Elmendorf–Mandell K -theory machine, which magically spits out a spectrum.


Finally, to bring us back down to earth, here is a warm-up question from class:

Warmup 7.2.2. In our torus example, we saw that $\mathcal{M}(F, a)$ is two dots (there were two flow lines).

- (a) What does $\mathcal{M}(F, v)$ look like?
- (b) There are no moduli spaces for higher index flows in our example, but what would you guess a moduli space of index-3 flows would look like?


7.3. Cohen–Jones–Segal construction.

7.3.1. *Examples to keep in mind.* Let $f_1(x) = 3x^2 - 2x^3$. This is a Morse function on \mathbb{R} with one index 0 critical point and a one index 1 critical point:



The flows together form the 1-dimensional cube. The Morse homology complex is captured by the flows of index 1 between the critical points.


Remark 7.3.1. The Morse function $f_1 : \mathbb{R} \rightarrow \mathbb{R}$ for \mathbb{R} is indeed Smale. The ascending manifold of e^0 and descending manifold of e^1 are shown below:




Their codimensions in \mathbb{R} are both 0. Their intersection, the red curve, also has codimension 0. Since $0 + 0 = 0$, the intersection is transverse.

(For more on transversality, see [GP10], a standard textbook on differential topology.)

Now let $f_2(x_1, x_2) = f_1(x_1) + f_1(x_2)$; this is a Morse function for \mathbb{R}^2 . Here is the contour (i.e. topographic) map of the graph of this function:



The red flag is at the top of a hill, and the blue dot is at the bottom of a lake. These points are at $(1, 1)$ and $(0, 0)$. The green dots, located at $(0, 1)$ and $(1, 0)$, are saddle points. A skier starting slightly away from the red dot following the path of gradient descent to the lake has a 1-parameter family of possible paths, depending on where they started:




The compactification of this moduli space of flows includes the flows where the skier skis from the red point to a green point, pauses (for eternity) and then skis to the blue point.

We write $\mathcal{M}(e^2, e^0)$ to represent the (compactified) moduli space of flows from e^2 to e^0 , the pink interval. Notice that the flows all together form a 2D cube.


Now let $f_3(x_1, x_2, x_3) = f_1(x_1) + f_1(x_2) + f_1(x_3)$. To visualize the graph of this function, imagine swimming in a pool, where the water temperature varies based on location. There is a local maximum in temperature at $(1, 1, 1)$ and a local minimum in temperature at $(0, 0, 0)$. At $(1, 0, 0)$, traveling away along the x_1 dimension means you'll feel colder; along the x_2 and x_3 dimensions, you'll feel warmer.

The Morse flows form a 3D cube:



The moduli spaces of index 1 and 2 flows are dots and intervals, respectively. The moduli space $\mathcal{M}(e^3, e^0)$ of flows from the unique maximum to the unique minimum is a hexagon, which you can see by taking a cross-section of all the flows along the dotted line shown.

The boundary of this hexagon is *stratified*, i.e. decomposed into submanifolds of various dimensions. Here is the most natural stratification:



This is most natural because, if you wanted to build this hexagon out of gumdrops, dried spaghetti, and cloth, you would build it this way.

- The 6 vertices are *codimension 2 strata*. These correspond to *twice broken flows*.
- The 6 edges are *codimension 1 strata*. These correspond to *once broken flows*. These come in two flavors:
 - index 2 flow followed by index 1 flow, i.e. a composition

$$\mathcal{M}(e_i^1, e^0) \times \mathcal{M}(e^3, e_i^1) \xrightarrow{\circ} \partial\mathcal{M}(e^3, e^0)$$

- index 1 flow followed by index 2 flow, i.e. a composition

$$\mathcal{M}(e_i^2, e^0) \times \mathcal{M}(e^3, e_i^2) \xrightarrow{\circ} \partial\mathcal{M}(e^3, e^0)$$

The codimension 1 strata of the boundary are going to play a very important role in determining how the rest of the boundary looks. We will think of the flavors of codimension 1 strata as colors of paint that we will use to paint the boundary of a $\langle n \rangle$ -manifold, which we define in the following section.

In general, the function

$$(13) \quad f_n(x_1, \dots, x_n) = \sum_{i=1}^n f_1(x_i)$$

is a Morse function for \mathbb{R}^n .

- The flows between critical points glue together to form a cube of dimension n .
- The moduli spaces that appear are all *permutohedra*, which we will study more carefully soon.
- The codimension k strata correspond to k -times broken flows.

The reason why we are so focused on this particular family of Morse functions is because the Khovanov chain complex lies over the n -dimensional


cube, for a diagram with n crossings. Lipshitz–Sarkar’s *Khovanov flow category*, which we will define in a few lectures, will be built by heavily relying on the *cube flow category*, which is what we have been studying carefully.

7.3.2. $\langle n \rangle$ -*manifolds*. The moduli spaces of Morse flows will be very special kinds of manifolds with boundaries of all possible codimension, and that have a certain kind of combinatorial rigidity.

Definition 7.3.2. A k -dimensional *manifold with corners* locally looks like a neighborhood of $(\mathbb{R}_+)^k$, where $\mathbb{R}_+ = [0, \infty)$.

Remark 7.3.3.

- In other words, it can have boundary points belonging to strata of any codimension c where $1 \leq c \leq k$.
- This loose definition is meant to be analogous to how you would explain what a manifold is to a friend, without having to talk about charts and atlases.
- This teardrop-shaped disk is a 2-dimensional manifold with corners:



We will need a very precise definition of the term ‘face’.


Definition 7.3.4. Let X be a manifold with corners.

- A *connected face* is the closure of a connected component of the codimension-1 boundary of X .

For example, the 6 teal arcs in the stratified hexagon (??) are, individually, connected faces of the hexagon.

- A *face* of X is a possibly empty union of **disjoint** connected faces of X .

For example, the pink subspace below is a face of the hexagon, but purple subspace below is *not* a face of the hexagon:




I have been asking you to vibe out the ‘obvious’ stratification of the boundary of a manifold with corners. To be more precise, we can describe the strata as the connected components of codimension i boundary, with i ranging from 1 to k .

For each x in a manifold with corners X , let $c(x)$ denote the codimension of the stratum that x lives in. (If x is in the interior of X , let $c(x) = 0$.)

Definition 7.3.5. A k -dimensional *manifold with faces* is a k -dimensional manifold with corners such that, for each $x \in X$, x belongs to exactly $c(x)$ different connected faces of X .

Remark 7.3.6. (1) The teardrop example from before is *not* a manifold with corners because the vertex v is adjacent to only one distinct edge, even though $c(v) = 2$.

- (2) This triangle is a manifold with faces:




- (3) Any face of a manifold with faces must necessarily also be a manifold with faces. **explain why**


Definition 7.3.7. A k -dimensional $\langle n \rangle$ -manifold is a k -dimensional manifold with faces along with an ordered n -tuple of faces of X where

- $\partial X = \bigcup_i \partial_i X$
- for $i \neq j$, $\partial_i X \cap \partial_j X$ is a face of both $\partial_i X$ and $\partial_j X$.

- Remark 7.3.8.**
- (1) We can think of each $i \in \{1, \dots, n\}$ as a color of paint. We paint a face of X with each color, so that by the time we're done painting with color n , we have covered the entire boundary of X with paint.
 - (2) A point $x \in \partial X$ will be covered with $c(x)$ different layers of paint, all of different colors, corresponding to the codimension-1 faces that it abuts.
 - (3) For this reason, I verbally refer to $\langle n \rangle$ -manifolds as ' n -painted manifolds'.
 - (4) Notice that k and n are different and basically unrelated.
 - (5) The triangle previously shown is cannot possibly be a 2-dimensional $\langle 2 \rangle$ -manifold, but it can be a $\langle 3 \rangle$ -manifold:



Because the colors of the corners (boundary of codimension ≥ 2) are determined by the colors painted on the codimension-1 boundary, it suffices to draw the figure below:



Example 7.3.9. (Permutohedra)

The *permutohedron* P_{n+1} is the most important example of an n -dimensional $\langle n \rangle$ -manifold, for our purposes.

- As a space, P_{n+1} is the convex hull in \mathbb{R}^{n+1} of the $(n+1)!$ points

$$\{(\sigma(1), \dots, \sigma(n+1)) \mid \sigma \in S_{n+1}\}$$

where S_{n+1} is the symmetric group on the letters $[n+1] := \{1, 2, \dots, n+1\}$.


- The connected faces of P_{n+1} are in bijection with proper subsets $\emptyset \neq S \subseteq [n+1]$; for a proper subset S of cardinality k , the corresponding face $F_S \subset \partial P_{n+1}$ is the intersection of P_{n+1} with the plane

$$\sum_{i \in S} x_i = \frac{k(k+1)}{2}.$$

- The $\langle n \rangle$ -manifold structure is given by

$$\partial_i P_{n+1} = \bigcup_{|S|=i} F_S.$$

- Here is the hexagon P_3 again ¹⁵:



Notice that the two ‘flavors’ of once-broken flows on the boundary of $\mathcal{M}(e^3, e^0)$ exactly correspond to the faces $\partial_1 P_3$ and $\partial_2 P_3$.

7.3.3. Flow categories. We are now ready to look at the definition of a *flow category*. If you keep in mind the Morse flows picture, all of these conditions should feel quite natural.

Note that I am using different variables names from the source.

Definition 7.3.10 ([LS14], Definition 3.13). A *flow category* is a category \mathcal{C} consisting of the following data:

- finitely many objects $\text{Ob}(\mathcal{C})$
- a \mathbb{Z} -grading on objects $\text{gr} : \text{Ob}(\mathcal{C}) \rightarrow \mathbb{Z}$
- morphisms between objects $\text{Hom}(x, y)$

subject to the following conditions:

- (1) $\text{Hom}(x, x) = \{\text{id}_x\}$ for all $x \in \text{Ob}(\mathcal{C})$. This is so that we have a category. Note that $\mathcal{M}(x, x) = \emptyset$.
- (2) For distinct objects $x \neq y$, $\text{Hom}(x, y)$ is a compact $(\text{gr}(x) - \text{gr}(y) - 1)$ -dimensional $(\text{gr}(x) - \text{gr}(y) - 1)$ -manifold. This space is denoted $\mathcal{M}(x, y)$.
- (3) Let $x, y, z \in \text{Ob}(\mathcal{C})$ be distinct, with $\text{gr}(x) \geq \text{gr}(y) \geq \text{gr}(z)$, with

$$\begin{aligned} \text{gr}(x) - \text{gr}(y) &= \ell \\ \text{gr}(y) - \text{gr}(z) &= m \\ \text{gr}(x) - \text{gr}(z) &= n = \ell + m. \end{aligned}$$

$$\begin{array}{c} x \\ \downarrow \text{index } \ell \\ y \\ \downarrow \text{index } m \\ z \end{array}$$

- The composition map

$$\circ : \text{Hom}(y, z) \times \text{Hom}(x, y) \rightarrow \text{Hom}(x, z)$$

¹⁵Thanks to Josh Turner for spotting and fixing a previous error in this example!

is an embedding into $\partial_m \text{Hom}(x, z)$. Recall this is the face of $\text{Hom}(x, z)$ painted with the color m .

- For $i < m$, the preimage of $\partial_i \text{Hom}(x, z)$ under the particular composition map we are looking at here is

$$\circ^{-1}(\partial_i \text{Hom}(x, z)) = \partial_i \text{Hom}(y, z) \times \text{Hom}(x, y).$$

- For $j > m$, the preimage of $\partial_j \text{Hom}(x, z)$ under the particular composition map we are looking at here is

$$\circ^{-1}(\partial_j \text{Hom}(x, z)) = \text{Hom}(y, z) \times \partial_{j-m} \text{Hom}(x, y).$$

- (4) If you run over all the y with grading in between that of x and z , you will be able to piece together the whole $\partial \text{Hom}(x, z) = \partial \mathcal{M}(x, z)$:


$$\partial_m \text{Hom}(x, z) \cong \coprod_{y \mid \text{gr}(y) = \text{gr}(z) + m} \text{Hom}(y, z) \times \text{Hom}(x, y).$$

Here \cong means ‘diffeomorphism’.

Example 7.3.11. For a Morse function $f : X \rightarrow \mathbb{R}$, we get the *Morse flow category* for f , exactly the way we have been motivating the definition of ‘flow category’.

7.3.4. *Framing the flow categories.* In this section we give a sense of what it means to upgrade a flow category to a *framed flow category*.

A ‘framing’ typically refers to a continuous choice of orthogonal basis (a ‘frame’) for the fibers of a bundle. When we talk about a framing for a knot K , for instance, we pick a choice of longitude λ for the solid torus $\nu(K)$; this gives a continuous choice of orthogonal bases for the normal bundle of the knot in S^3 :



Our goal is to provide framings for the moduli spaces in a flow category, in such a way that they can be glued together nicely. To do this, we embed the moduli spaces into some high dimensional Euclidean spaces and pick framings for their normal bundles there.

First, recall that for finite dimensional smooth manifolds X , we have the *Whitney embedding theorem*, which guarantees that we can find some large N such that there exists an embedding $X \hookrightarrow \mathbb{R}^N$.

In a way, this theorem tells us that Euclidean space \mathbb{R}^d are the ‘model smooth manifolds’. We already knew that, of course, but this raises the question,

“what is a model $\langle n \rangle$ -manifold?”

We already saw that manifolds with corners are modeled after \mathbb{R}_+^d , so we will certainly need at least that much information. The important difference between $\langle n \rangle$ manifolds and run-of-the-mill manifolds with corners is the painting of the faces; so, we need model spaces that are also $\langle n \rangle$ -manifolds, where it is very clear how to paint the faces with n colors.

Definition 7.3.12. Let $\vec{d} = (d_0, d_1, \dots, d_n) \in \mathbb{N}^{n+1}$ (here $\mathbb{N} = \mathbb{N} \cup \{0\}$), and let

$$\mathbb{E}_n^{\vec{d}} := \mathbb{R}^{d_0} \times (\mathbb{R}_+ \times \mathbb{R}^{d_1}) \times (\mathbb{R}_+ \times \mathbb{R}^{d_2}) \times \cdots \times (\mathbb{R}_+ \times \mathbb{R}^{d_n}).$$

Note that there is a copy of \mathbb{R}_+ for each $i \in \{1, \dots, n\}$. The boundaries of these give us n different faces. The $\langle n \rangle$ -manifold structure is given by

$$\partial_i(\mathbb{E}_n^{\vec{d}}) = \mathbb{R}^{d_0} \times (\mathbb{R}_+ \times \mathbb{R}^{d_1}) \times \cdots \times (\{0\} \times \mathbb{R}^{d_i}) \times \cdots \times (\mathbb{R}_+ \times \mathbb{R}^{d_n}).$$

The appropriate, structure-preserving notion for an embedding of $\langle n \rangle$ -manifolds, for our framing purposes, is the following:

Definition 7.3.13. A *neat embedding* ι of an $\langle n \rangle$ -manifold is a smooth embedding

$$\iota : X \hookrightarrow \mathbb{E}_n^{\vec{d}}$$

for some $\vec{d} \in \mathbb{N}^{n+1}$, where


- $\iota^{-1}(\partial_i \mathbb{E}_n^{\vec{d}}) = \partial_i X$ for all i (respecting the $\langle n \rangle$ -manifold structure) and
- whenever a stratum of X abuts a lower-dimensional (higher-codimensional) stratum of $\mathbb{E}_n^{\vec{d}}$, it does so perpendicularly. This is the *neat part*; see Remark 7.3.15 for a more precise definition.

Example 7.3.14. The triangle, as a 2-dimensional $\langle 3 \rangle$ -manifold, neatly embeds in $\mathbb{E}_3^{(0,0,0,0)} = (\mathbb{R}_+)^3$ as an octant of the unit 2-sphere.

Remark 7.3.15. Lawson–Lipshitz–Sarkar note that a $\langle n \rangle$ -manifold can be viewed as a functor $\underline{2}^n \rightarrow \mathbf{Top}$ by

$$X(u) = \begin{cases} X & \text{if } u = \bar{1} \\ \bigcap_{i \in \{i \mid u_i = 0\}} \partial_i X & \text{otherwise.} \end{cases}$$

For example, here is the hexagon as a cube-shaped diagram of topological spaces:



The second requirement in Definition 7.3.13 becomes:

$$\text{For all } u < v \text{ in } \{0, 1\}^n, X(v) \perp \mathbb{E}_n^{\vec{d}}(u).$$

There is an analogue to the Whitney embedding theorem for $\langle n \rangle$ -manifolds, which we will paraphrase loosely:

Fact 7.3.16 (see Lemma 3.11 of [LS14] for more detail). Let X be a $\langle n \rangle$ -manifold. Given a neat embedding ι of $\partial X \hookrightarrow \mathbb{E}_n^{\vec{d}}$, there exists a neat embedding of $X \hookrightarrow \mathbb{E}_n^{\vec{d}'}$ for some $\vec{d}' \geq \vec{d}$, induced by ι . Here, $\vec{d}' \geq \vec{d}$ means $d'_i \geq d_i$ for all i .

Note that, since we are quite capable of neatly embedding finite sets of points, by induction we know that every $\langle n \rangle$ -manifold can be neatly embedded.

We now return back to framing flow categories, and give very rough definitions.

A *neat embedding* of a flow category \mathcal{C} is a collection of neat embeddings for all the moduli spaces in \mathcal{C} , in a coherent way so that composition of moduli spaces is respected. (See Definition 3.16 of [LS14] for more detail.)

- To make sure all the embeddings are in the same spatial universe (in the non-technical sense), we have to first pick a bi-infinite sequence $\vec{d}: \mathbb{Z} \rightarrow \mathbb{N}$. Don't worry, we'll only use a finite subsequence of this, because \mathcal{C} only has finitely many objects.
- Then, let $\mathcal{M}(i, j)$ be the union of all the moduli spaces $\mathcal{M}(x, y)$ where $\text{gr}(x) = i$ and $\text{gr}(y) = j$; neatly embed $\mathcal{M}(i, j)$ into pieces edgey spaces of the form $\mathbb{E}_{i-j-1}^{(d_j, d_{j+1}, \dots, d_{i-1})}$. Call these embeddings $\iota_{i,j}$.
- If $p \in \mathcal{M}(x, y)$ and $q \in \mathcal{M}(y, z)$, then the composition condition requires that

$$\iota_{x,z}(q \circ p) = (\iota_{y,z}(q), 0, \iota_{x,y}(p)).$$

The most important thing to note here is that $\iota_{x,y}$ and $\iota_{y,z}$ determine $\iota_{x,z}(q \circ p)$.

Observe that at this point, we haven't really actually picked a *framing* for anything. However, we do have moduli spaces embedded in Euclidean spaces, with normal bundles ready to be framed.

We can now give a rough definition of framed flow categories:

Definition 7.3.17. A *framed flow category* is a neatly embedded flow category \mathcal{C} **plus** a *coherent* framing of the normal bundles.

The word 'coherent' is pulling a lot of weight here. This is the point at which we make sure we'll actually be able to glue our moduli spaces together to make a nice space. The main requirement is that the product framing of the normal bundles of two embeddings of moduli spaces $\nu(\iota_{y,z}) \times \nu(\iota_{x,y})$ equals the pullback framing of $\circ^{-1}\nu(\iota_{x,z})$ induced by the composition map.

Finally, Lipshitz–Sarkar give explicit instructions for how to take the data of a framed flow category and transform them into a CW-complex (of some dimension). See Figures 3.3 and 3.4 of [LS14] to get a sense of how to do this. We talked about these pictures in class, but I will not be including my commentary in these notes.

7.4. A framed flow category for Khovanov homology. Let L be an oriented link, and let D denote a diagram for L with n crossings. Lipshitz–Sarkar build the framed flow category for Khovanov homology as follows:

- (1) Frame the n -dimensional cube flow category $\mathcal{C}_C(n)$.
- (2) Build a Khovanov flow category $\mathcal{C}_K(D)$ with moduli spaces corresponding to *decorated resolution configurations* for D .
- (3) Make sure that the moduli spaces of $\mathcal{C}_K(D)$ are all *trivial* covers of the corresponding moduli spaces in $\mathcal{C}_C(n)$.
- (4) Conclude that framings can be lifted from $\mathcal{C}_C(n)$ to $\mathcal{C}_K(D)$, thereby boosting $\mathcal{C}_K(D)$ to a *framed* flow category.


There is a lot to say and prove about all of these steps, but we will give an overview of the key philosophical ideas in each step (as opposed to the key technical ideas).

7.4.1. Framing the cube flow category: obstruction theory. Let's first talk a bit about the idea behind *obstruction theory*.


Consider the unknot $U \subset S^3$. A framing λ of U is a continuous choice of frame (orthonormal basis) for the fibers of the normal bundle.



This is equivalent to an integral choice of longitude on the torus $\partial\nu(U)$, the boundary of a tubular neighborhood of U :



Here are two different framings for U .



Notice that the 0-framed annulus can be extended to an embedded disk in S^3 bound by U , whereas the (-1) -framed annulus cannot.

In this sense, the value $\lambda \in \mathbb{Z}$, when nonzero, provides an *obstruction* to extending the framed annulus to a disk.

Pictorially, we can think of λ as a longitude. But another way to think about λ is as an element of $\pi_1(O(2))$, as follows. The columns of a matrix in $O(2)$ give an orthogonal basis for the fiber of the normal bundle over a point $p \in U$. Moving p all the way around U is a loop in $O(2)$, which corresponds to an element of $\pi_1(O(2))$. [You may convince yourself why we only care about these loops up to homotopy.](#)

In order to frame the cube flow category $\mathcal{C}(n)$, Lipshitz–Sarkar first prove that there exists a coherent set of orientations for the moduli spaces [\(we didn't really talk much about this\)](#). Then, suppose we are given a coherent orientation and a neat embedding of the $\mathcal{C}(n)$ (both of which we know exist), we wish to frame the moduli spaces, starting from the 0-dimensional ones, whose framings are given by the orientations chosen, and then building up to framings for the moduli spaces of higher dimension — if there is no obstruction.

Suppose you already have a coherent framing of all the moduli spaces of dimension $< k$. This data produces an obstruction class

$$\mathfrak{o} \in C^{k+1}(\mathcal{C}(n), \pi_{k-1}(\mathbf{O})).$$

- Here, $\mathcal{C}(n)$ is ‘the cube’ of dimension n , i.e. the CW-complex you get from gluing all the flow lines between critical points of $f_n(x_1, \dots, x_n)$ (from (13)).
- $\mathbf{O} = \text{colim}_{k \rightarrow \infty} O(k)$, the colimit of all the orthogonal groups of finite dimension, related by inclusion maps.

While we will not discuss how the obstruction class \mathfrak{o} is obtained, but we can still get an idea of why \mathfrak{o} makes sense as the obstruction to extending the framing to the k -dimensional moduli spaces. [Recall that all our moduli spaces are permutahedra, which are in particular homeomorphic to balls of various dimensions.](#)

- A (component of a) k -dimensional moduli space corresponds to a $k + 1$ -dimensional cell in the CW decomposition of $\mathcal{C}(n)$ given by the spaces of flow lines.
- A framing of a k -dimensional moduli space $\mathcal{M}(x, y)$ is a choice of an orthogonal basis for the fiber of the normal bundle above every point. So, depending on how high-dimensional the space you're embedded in is, your normal bundle is some rank, and you'd be working with that many vectors in your orthogonal bases.
- For the framing along the boundary $\partial\mathcal{M}(x, y) \cong S^{k-1}$ to extend to a framing of $\mathcal{M}(x, y)$, we need the boundary framing to 'not be twisted,' or more precisely, for the map $S^k \rightarrow \mathbf{O}$ to be nullhomotopic.

Lipshitz–Sarkar show that indeed, the cochain \mathfrak{o} is in fact a coboundary, and modifications can be made so that the cochain itself vanishes, so that there is no obstruction.

7.4.2. *Resolution configurations, ladybug configuration.* In the Khovanov flow category, the objects are in correspondence with the distinguished generators of the Khovanov chain complex.


Remark 7.4.1. This time, I really do mean the *chain* complex, rather than the technically cohomologically graded version we've been discussing all quarter. But, since we are just getting a sense of the Khovanov stable homotopy type construction, I will not spend time worrying about things like whether we should be using \mathbb{Z} or \mathbb{Z}^{op} . Perhaps take this as a **warning** that I am not being careful about these conventions here.

The 0-dimensional moduli spaces are in correspondence with the components of the differential between distinguished generators (the 'arrows').

If we think of the 'dots and arrows' picture of the Khovanov chain complex as a direct graph, then the 1-dimensional moduli spaces come from paths of length 2 (these are the analogues to once broken flow lines in the Morse picture). We capture these moduli spaces by drawing *resolution configurations*.

We will now introduce some vocabulary words in context, via the following very important example.

Consider the cube of resolutions for this two-component unlink diagram:




- At vertex 00 is a complete resolution, along with two red arcs indicating that there are two crossings at which we have yet to switch the resolution to the 1 resolution. In other words, we have yet to *surger* along the red arcs. Yes, we are doing surgery on 1-manifolds.
- At the 10 and 01 resolutions, we have already surgered along one red arc, but have yet to surger along the other.
- At the 11 resolution, all red arcs have been surgered, so there is nothing left to do.


These complete-resolutions-with-red-arcs are called *resolution configurations*. These diagrams are useful because they point out 'subcubes', in the sense that the 00 resolution configuration in the above example can represent the whole 2-dimensional cube that you get from surgering the red arcs in the two different orders.

In general, a resolution configuration need not only describe a cube whose extremal vertex is the all-ones resolution. By omitting some red arcs, one can describe a subcube somewhere in the middle of the cube of resolutions. In other words, the 2D cube of resolutions above could appear somewhere within a much bigger cube of resolutions for a much more complicated link diagram.

A *labeled resolution configuration* additionally specifies a particular distinguished generator. For example, for the unlink diagram above, at quantum grading $\text{gr}_q = 0$ we see the following subcomplex:



I can specify this entire subcomplex by drawing




In other words, a labeled resolution configuration is a pair (D, x) where D is a resolution configuration and x is a labeled of all circles in D with v_+ or v_- .

Finally, a *decorated resolution configuration* specifies both the beginning and the end of the path. In the example above, there was a unique distinguished generator at the 11 resolution that length-2 paths from the shown generator map to. In general, this might not be true. So a decorated resolution configuration also specifies which distinguished generator your path should end at.

In other words, a decorated resolution configuration is a triple (D, x, y) where (D, x) and $(s(D), y)$ are both labeled resolution configuration; here, $s(D)$ means ‘result of surgery along the red arcs of D ’.

So, if D is a resolution configuration with k red arcs, we say that (D, x, y) is an *index k* decorated resolution configuration. These should correspond to $(k - 1)$ -dimensional moduli spaces in the Khovanov flow category.

Notice that in our example, there were two possible ways that we could build our 1-dimensional moduli spaces:



A ‘correct choice’ (i.e. a choice that allows us to obtain a framed flow category) will be determined by the constraints in the following section.

7.4.3. Khovanov flow category as a trivial cover.

Definition 7.4.2. A grading-preserving functor $\mathcal{F} : \hat{\mathcal{C}} \rightarrow \mathcal{C}$ between flow categories is a *cover* if, for all $x, y \in \hat{\mathcal{C}}$, whenever $\mathcal{M}_{\hat{\mathcal{C}}}(x, y) \neq \emptyset$, the map of $\langle \text{gr}(x) - \text{gr}(y) - 1 \rangle$ manifolds

$$(14) \quad \mathcal{F} : \mathcal{M}_{\hat{\mathcal{C}}}(x, y) \rightarrow \mathcal{M}_{\mathcal{C}}(\mathcal{F}(x), \mathcal{F}(y))$$

- respects the $\langle \text{gr}(x) - \text{gr}(y) - 1 \rangle$ manifold structures,
- is a local diffeomorphism, and
- is a covering map.


The cover is *trivial* if the map (14) is a trivial covering map.

Note that the term ‘trivial’ makes it sound like the whole covering situation is somehow trivial; this is not true, because moduli spaces can glue together along their boundaries in many interesting ways.

Trivial covers are particularly important to the construction at hand, because Lipshitz–Sarkar essentially prove the following:

Proposition 7.4.3 ([LS14], Proposition 5.2, (E-3), casually stated). If the Khovanov flow category trivially covers the cube flow category $\mathcal{C}_{\mathcal{C}}(n)$, then we can use the framing of $\mathcal{C}_{\mathcal{C}}(n)$ to frame the Khovanov flow category.

We return to the ladybug configuration to see the difference between trivial and nontrivial covers:



We assume that all the 0- and 1-dimensional moduli spaces are trivial covers, as shown.

In order to make sure the covering is trivial for the hexagonal 2-dimensional moduli space, we have to make sure that, on its boundary, we get two 6-cycles (left), rather than a single 12-cycle (right).

Lipshitz–Sarkar cleverly use the orientation of the diagram plane to make a global choice for any ladybug configurations that arise, ensuring that all covering maps for hexagons look like the picture on the left.

7.4.4. Finishing up. It turns out that this is the last important choice to make! Moduli spaces of dimension $k \geq 3$ are copies of the permutahedron P_{k+1} , which is homeomorphic to B^k (recall that it’s the convex hull of a bunch of points). This means the boundary of the moduli space is homeomorphic to S^{k-1} , which is simply connected, and therefore only have trivial covers! Recall that connected covering spaces of X correspond to subgroups of $\pi_1(X, *)$.

To finish up:

- We can build our moduli spaces inductively, and frame them using our framing of the cube flow category. This gives us a framed flow category refining the Khovanov chain complex we started with.

- Next, use the Cohen–Jones–Segal construction to build a high-dimensional CW-complex from the framed flow category, called the *realization*.
- Finally, take the suspension spectrum of the realization (and formally desuspend appropriately to get the correct homological gradings) to get the link invariant \mathcal{X}_{Kh} .

Remark 7.4.4. Actually, we get a spectrum for each quantum grading $\text{gr}_q = j$:

$$\mathcal{X}_{\text{Kh}}(L) = \bigvee_j \mathcal{X}_{\text{Kh}}^j(L)$$

7.5. The Burnside functor approach. The reference for this section is [LLS20].

Lawson–Lipshitz–Sarkar give a different way to build equivalent Khovanov spectra, via the following process:


- Start with the data of the Khovanov complex (but not totalized), drawn as a (commuting) diagram of abelian groups and maps between them. If a link diagram has n crossings, then we can think of this diagram as a functor $\underline{2}^n \rightarrow \text{Mod}_{\mathbb{Z}}$.
- Make some choices (i.e. add information) to lift this to a functor $\underline{2}^n \rightarrow \mathcal{B}$, where \mathcal{B} is the *Burnside 2-category*.
- Use the *Pontrjagin–Thom* construction to turn this data into a functor $(\underline{2}^{\text{op}})^n \rightarrow \mathbf{Top}_*$.
- Check that this diagram of topological spaces (actually, wedges of spheres of some large dimension) is *homotopy coherent*, and then take the homotopy colimit to obtain a space. (Then, as usual, take the suspension spectrum and desuspend to make up for the large dimension of the spheres you used.)

While we sketch out ‘Burnside-functor’ approach to building \mathcal{X}_{Kh} , look out for the following important points we encountered while discussing the flow categories construction method:

- What elements of this construction correspond to the higher-dimensional moduli spaces?
- Where do we make the ladybug matching choice?
- At what index (for the resolution configurations) can we stop making choices, and why?

7.5.1. The Burnside category. The *Burnside category* \mathcal{B} is a 2-category, meaning that there are objects, morphisms, and also morphisms between morphisms (level 2 homomorphisms). In \mathcal{B} :


- Objects are finite sets X
- 1-morphisms are *finite correspondences*: a morphism in $\text{Hom}(X, Y)$ is a diagram of set maps



where s_A and t_A are called the *source* and *target* set maps, respectively.

You can think of the elements of A as ‘arrows’ from points in X to points in Y ; an arrow $a \in A$ has a source point and a target point, and $s_A(a)$ and $t_A(a)$ are just telling us which points these are.

- Composition of 1-morphisms is given by *pullback* (i.e. *fiber product*): the composition of $(A, s_A, t_A) \in \text{Hom}(X, Y)$ and $(B, s_B, t_B) \in \text{Hom}(Y, Z)$ is the finite correspondence $(C, s_A \circ s_C, t_B \circ t_C) \in \text{Hom}(X, Z)$




where

$$C = A \times_Y B = \{(a, b) \in A \times B \mid t_A(a) = s_B(b)\}.$$

The maps s_C and t_C are induced by projection from $A \times B$, i.e. $s_C((a, b)) = a$ and $t_C((a, b)) = b$.

- 2-morphisms are *isomorphisms of correspondences*: given (A, s_A, t_A) and (B, s_B, t_B) in $\text{Hom}(X, Y)$, an element of $2\text{Hom}(A, B)$ is an isomorphism of sets $f : A \rightarrow B$ such that the following diagram commutes:




7.5.2. *Cube-to-Burnside functor.* Given a Khovanov cube-shaped complex $\underline{2}^n \rightarrow \text{Mod}_{\mathbb{Z}}$, we want to build a lift of this functor to $\underline{2}^n \rightarrow \mathcal{B}$. Here is how we define the new functor.

(Objects) At each vertex u of the cube, the distinguished generators (pure tensors in $V^{\otimes |\pi_0(D_u)|}$) become the elements of the finite set at u .

(1-Morphisms) The arrows in the Khovanov differential between distinguished generators determine the finite correspondences.

For example, consider the merge map $m : V \otimes V \rightarrow V$.




On the left is our familiar dots-and-arrows picture, but now each arrow has a name. On the right is the ... *corresponding*... correspondence in the Burnside category.


(2-Morphisms) To define a 2-functor from $\underline{2}^n$, we technically need to think of $\underline{2}^n$ as a 2-category as well. We do this by asserting that $\underline{2}^n$ has no non-identity 2-morphisms.

Defining our 2-morphisms is where new choices are made. For an example, we will once again study the ladybug configuration (see §7.4.2).


Here is the usual dots-and-arrows picture of the chain complex (at the interesting quantum grading), but now the arrows have names:



Translating to a diagram in the Burnside category, we have the following picture, where a blue lines are parts of source maps, and green lines are parts of target maps.



Now let's consider the 1-morphisms we get from length-2 paths through the cube; these are compositions of the arrows in the dots-and-arrows picture.



- Along the top half of the cube, the composition of 1-morphisms $A = \{a_1, a_2\}$ and $B = \{b_1, b_2\}$ yields a 1-morphism $C \in \text{Hom}(X = \{x\}, Z = \{z\})$ containing two elements,

$$c_1 = (a_1, b_1) \quad \text{and} \quad c_2 = (a_2, b_2).$$

- Along the bottom half of the cube, we get a different composite 1-morphism $C' \in \text{Hom}(X, Z)$ also containing two elements,


$$c'_1 = (a'_1, b'_1) \quad \text{and} \quad c'_2 = (a'_2, b'_2)$$

An element of $2\text{Hom}(C, C')$ is a choice of isomorphism between these two 2-element sets. There are exactly two possible isomorphisms (either $c_1 \mapsto c'_1$ or $c_1 \mapsto c'_2$).

Lawson–Lipshitz–Sarkar check through all the combinatorics in the Khovanov cube, and this resolution configuration is, once again, the most interesting.

7.5.3. Burnside functor to homotopy coherent diagram of wedges of spheres. We now use the Burnside functor $\underline{2}^n \rightarrow \mathcal{B}$ to obtain a diagram of topological spaces, using the Pontrjagin–Thom construction. We will stick with our running ladybug example.


First, replace all the objects in the Burnside functor diagram with k -dimensional disks, with the size of the disks increasing with the cube grading $|u| = \sum_i u_i$. In the drawing below, $k = 2$:




The correspondences tell us where to embed smaller disks into larger disks, so that images are disjoint.

Beware: The following picture is just for showing you how we're embedding the smaller disks into bigger ones. For example, I've drawn the image of the embeddings $\mathbb{D}_x \hookrightarrow \mathbb{D}_{y_i}$ ($i = 1, 2$), in order to indicate a *multi-embedding* (see Remark 7.5.1) of \mathbb{D}_x in $\mathbb{D}_{y_1} \sqcup \mathbb{D}_{y_2}$.

Remark 7.5.1. **Beware** that what I've drawn is not a map of topological spaces. This is not even a function, but rather a *multifunction*, where each point maps to multiple points. For this reason, I will be referring to these disk embeddings as 'multi-embeddings'. I don't believe this is standard language. For more technical language, you can look up the *little boxes operad*.



We are now set up to build a functor $(\mathbb{Z}^{op})^n \rightarrow \text{Top}_*$. First, at each vertex of the cube, identify all of the boundaries to a single basepoint $*$:



We can get a map from $S_{y_1}^2 \rightarrow S_x^2$ by collapsing everything outside of the image of $\mathbb{D}_x \hookrightarrow \mathbb{D}_{y_1}$ to the basepoint in S_x^2 . The same process gives us a map $S_{y_2}^2 \rightarrow S_x^2$. Combined, we obtain a map of based topological spaces $S_{y_1}^2 \vee S_{y_2}^2 \rightarrow S_x^2$. The map from $S_z^2 \rightarrow S_{y_1}^2 \vee S_{y_2}^2$ is slightly more interesting, but defined the same way. [Try to visualize this yourself, by imagining that you are, perhaps, making a balloon animal, grabbing a ball of dough from a large mass of dough, or balling up some mozzarella from a vat of curds.](#)

Even more interesting is the composition of embeddings.

- Along the top of the cube in (15), we get two homeomorphic images of \mathbb{D}_x inside \mathbb{D}_z . These are the blue disks inside the orange circles.
The composite multi-embedding of \mathbb{D}_x to \mathbb{D}_z agrees with the collapse map $S_z^2 \rightarrow S_x^2$: we crush everything outside of the two images of \mathbb{D}_x to the basepoint, and get a degree-2 map $S_z^2 \rightarrow S_x^2$.
- Along the bottom of the cube, we get a different set of homeomorphic images of \mathbb{D}_x inside \mathbb{D}_z . These are the blue disks inside the red circles.

This path gives us a different degree-2 map $S_z^2 \rightarrow S_x^2$.

In order to make sure that the two maps of topological spaces $S_z^2 \rightarrow S_x^2$ are homotopic, we need the dimension k to be large enough (we will discuss more about this soon). But because this square is potentially only a small part of a much larger cube, we actually need to *pick specific homotopies* between the two maps. To do this, we use the 2-morphisms we chose when we built the Burnside functor.

From the Burnside functor, we get a particular matching of the orange-blue disks (associated with $C = \{c_1, c_2\}$ in the Burnside functor) with the red-blue disks (associated with $C' = \{c'_1, c'_2\}$). Suppose we chose the matching $c_i \mapsto c'_i$. To choose a homotopy, we just need to pick two paths: one that takes the orange-blue disk associated with c_1 to the red-blue disk associated with c'_1 , and another that takes the orange-blue disk associated with c_2 to the red-blue disk associated with c'_2 .

After choosing these homotopies, we also need to choose homotopies among homotopies when we zoom out and look at an index-3 subcube (or index-3 resolution configuration), and then homotopies of homotopies of homotopies ... etc. If we can successfully find homotopies at all levels all the way to index- n resolution configurations (recall that our link diagram has n crossings), then we have a *homotopy coherent* diagram of topological spaces.

We want to make as few choices as possible. To do this, we just make sure k is sufficiently large, because this will ensure that all possible choices are effectively equivalent (i.e. homotopic). Observe that the data of a multi-embedding $e : \mathbb{D}_x \rightarrow \mathbb{D}_y$ can be captured by just recording the centers of the image disks. Suppose there are m image disks. Then, e is just an

element of $\text{Conf}(\mathbb{D}_y, m \text{ points})$, the configuration space of m points inside the $(k\text{-dimensional})$ disk \mathbb{D}_y .

Fact 7.5.2. The configuration space of finitely many points in a k -dimensional disk (ball) is $(k - 2)$ -connected, i.e., for any $j \leq k - 2$,

$$\pi_j(\text{Conf}(\mathbb{D}^k, m \text{ points})) = 0.$$

Using this fact, we see that as long as $k \geq n + 2$, no matter how we choose to multi-embed our disks, our diagram will be homotopy coherent.

At this point, we can proceed to the last step, where we use this diagram to build a space, by throwing our homotopy coherent diagram into a machine that spits out the *homotopy colimit*.

7.5.4. Homotopy colimit. If you want to see a concrete model of a homotopy colimit of based topological spaces, you can check out the reference for this whole subsection, [LLS20].

Instead, let's talk about the significance of homotopy colimits in general, given that we are currently in possession of a homotopy coherent diagram of topological spaces $(\underline{2}^{op})^n \rightarrow \mathbf{Top}_*$.

When we work with topological spaces, we really consider them up to homeomorphism. When I take a colimit of a diagram of topological spaces, the resulting space should be well-defined up to homeomorphism.

On the other hand, in homotopy theory, we want to consider spaces up to homotopy equivalence. When I take a colimit of a diagram of topological spaces, I should expect that if I replace one of the spaces in the diagram with something homotopy equivalent to it, that I should get the same colimit. But this is not always true, as evidenced by the following classic example.

Consider the two diagrams below (ι_∂ means ‘inclusion of boundary’):

$$\begin{array}{ccc} & S^1 & \\ \swarrow & & \searrow \\ D^2 & \xrightarrow{\iota_\partial} & * \end{array} \qquad \begin{array}{ccc} & S^1 & \\ \swarrow & & \searrow \\ * & & * \end{array}$$

In the second diagram (the one on the right), we've replaced D^2 with a homotopy equivalent space, $*$.

Remark 7.5.3. For this particular shape of diagram, the colimit is called a *pushout*.


Let's briefly recall how pushouts work in the category of sets, for simplicity.

$$\begin{array}{ccc} A & \xrightarrow{\beta} & B \\ \downarrow \gamma & & \\ C & & \end{array} \qquad \text{the pushout is the set } B \sqcup C / \sim$$

where $b \sim c$ iff there exists some $a \in A$ such that $\beta(a) = b$ and $\gamma(a) = c$.


Colimits are a generalization of this. I might add an appendix about colimits.

The colimit of the diagram on the left is S^2 , where as the colimit of the diagram on the right is a one-point space:



These spaces are not homotopy equivalent! The problem is that the colimit on the right is too degenerate.

To fix this, we thicken the pieces before we glue:



This time, for each map $S^1 \rightarrow *$, we first replace S^1 with $S^1 \times I$, and then glue one end to $*$. Then, we glue the two pieces together.

On the right, we see that after gluing, we indeed get an S^2 , which now agrees with the colimit of the first diagram we saw.

This tells us that S^2 is the homotopy colimit (or homotopy pushout) of both diagrams.

The takeaway is that the Burnside functor approach to building \mathcal{K}_{Kh} is meant to produce a *stable homotopy type* rather than a homeomorphism class of topological spaces. This flexibility allows us to actually obtain an invariant of the link (under Reidemeister moves of the diagram). So, when we build the space, we need to treat everything up to homotopy.

8. HOMOLOGICAL SKEIN INVARIANTS FOR 4-MANIFOLDS

8.1. Skein Lasagna modules from Khovanov homology.

8.2. Cabled Khovanov homology.

8.3. Ian's lecture.

9. CATEGORICAL AND ALGEBRAIC BACKGROUND


9.1. Mapping cones. In a category of chain complexes graded cohomologically, let $f : A \rightarrow B$ be a chain map between complexes $A = (\bigoplus_i A^i, d_A^i)$ and $B = (\bigoplus_i B^i, d_B^i)$:

$$\begin{array}{ccccccc}
 \cdots & \rightarrow & A^{-1} & \xrightarrow{d_A^{-1}} & A^0 & \xrightarrow{d_A^0} & A^1 & \xrightarrow{d_A^1} & A^2 & \rightarrow \cdots \\
 & & \downarrow f^{-1} & \curvearrowright & \downarrow f^0 & \curvearrowright & \downarrow f^1 & \curvearrowright & \downarrow f^2 & \\
 \cdots & \rightarrow & B^{-1} & \xrightarrow{d_B^{-1}} & B^0 & \xrightarrow{d_B^0} & B^1 & \xrightarrow{d_B^1} & B^2 & \rightarrow \cdots
 \end{array}$$

The *mapping cone* of f , denoted $C(f)$ or $\text{Cone}(f)$, is defined as

$$C(f) := \left(\bigoplus_i (A^{i+1} \oplus B^i), \quad \bigoplus_i \begin{pmatrix} -d_A^{i+1} & 0 \\ f^{i+1} & d_B^i \end{pmatrix} \right)$$

and can be visualized as a collapse of previous diagram:



In the previous diagram, the squares commuted (because f is a chain map). In order to get $d_{C(f)}^2 = 0$, we instead need the squares to *anticommute*, and so we made the choice to negate all the instances of the d_A^i differentials.

Warning 9.1.1. Recall that the shift functor $[n]$ in Bar-Natan's conventions [BN02] is defined by

$$(A[n])^i = A^{i-n}$$

so that the quantum dimension of $(A[n])^i$ is q^n times the quantum dimension of A^i . In these conventions, we may write

$$C(f) = \left(A[-1] \oplus B, \begin{pmatrix} d_{A[-1]} & 0 \\ f[-1] & d_B \end{pmatrix} \right)$$

where we use the convention that $d_{A[-1]} = -d_A$.

This is not how you'd usually see the definition of a cone in a cochain complex category, so I encourage you to instead carefully understand the cone construction without memorizing the notation. We will, however, continue using these conventions for the sake of having useful notation.

In practice, we usually just write the cone as

$$C(f) = \left(A \xrightarrow{f} B \right),$$

with the parentheses indicating that we are thinking of a single complex, i.e. the whole cone, rather than a map between complexes. We can think of B as 'shifted by 0 homologically', and instead write

$$C(f) = \left(A \xrightarrow{f} \underline{B} \right)$$

to be more clear that A is the one being shifted left.

Notice that $B \hookrightarrow C(f)$ as a subcomplex, and $A[-1] \cong C(f)/B$ is the quotient. Thus there is a short exact sequence

$$0 \rightarrow B \xrightarrow{i} C(f) \xrightarrow{p} A[-1] \rightarrow 0$$

where i is inclusion and p is the quotient map. This gives rise to a long exact sequence on homology

$$\cdots \rightarrow H^j(B) \xrightarrow{i^*} H^j(C(f)) \xrightarrow{p^*} H^{j+1}(A) \xrightarrow{f^*} H^{j+1}(B) \xrightarrow{i^*} H^{j+1}(C(f)) \rightarrow \cdots$$

where the connecting map f^* is induced by f .

Remark 9.1.2. Once again, because of our conventions, we have a somewhat nonstandard-looking exact triangle:

$$A \rightarrow B \rightarrow C(f) \rightarrow A[-1].$$

But all is well; this is just a notational difference from what you might be used to.

REFERENCES

- [AD19] Akram Alishahi and Nathan Dowlin. The Lee spectral sequence, unknotting number, and the knight move conjecture. *Topology Appl.*, 254:29–38, 2019.
- [Ada04] Colin C. Adams. *The knot book*. American Mathematical Society, Providence, RI, 2004. An elementary introduction to the mathematical theory of knots, Revised reprint of the 1994 original.
- [AS19] Mohammed Abouzaid and Ivan Smith. Khovanov homology from Floer cohomology. *J. Amer. Math. Soc.*, 32(1):1–79, 2019.
- [Ati88] Michael Atiyah. Topological quantum field theories. *Inst. Hautes Études Sci. Publ. Math.*, (68):175–186, 1988.
- [Ben83] Daniel Bennequin. Entrelacements et équations de Pfaff. In *Third Schnepfenried geometry conference, Vol. 1 (Schnepfenried, 1982)*, volume 107–108 of *Astérisque*, pages 87–161. Soc. Math. France, Paris, 1983.
- [BG15] John A. Baldwin and J. Elisenda Grigsby. Categorified invariants and the braid group. *Proc. Amer. Math. Soc.*, 143(7):2801–2814, 2015.
- [BHPW23] Anna Beliakova, Matthew Hogancamp, Krzysztof K. Putyra, and Stephan M. Wehrli. On the functoriality of \mathfrak{sl}_2 tangle homology. *Algebr. Geom. Topol.*, 23(3):1303–1361, 2023.
- [Bla10] Christian Blanchet. An oriented model for Khovanov homology. *J. Knot Theory Ramifications*, 19(2):291–312, 2010.
- [BN02] Dror Bar-Natan. On Khovanov’s categorification of the Jones polynomial. *Algebr. Geom. Topol.*, 2:337–370, 2002.
- [BN05] Dror Bar-Natan. Khovanov’s homology for tangles and cobordisms. *Geom. Topol.*, 9:1443–1499, 2005.
- [BN07] Dror Bar-Natan. Fast Khovanov homology computations. *J. Knot Theory Ramifications*, 16(3):243–255, 2007.
- [BNM06] Dror Bar-Natan and Scott Morrison. The Karoubi envelope and Lee’s degeneration of Khovanov homology. *Algebr. Geom. Topol.*, 6:1459–1469, 2006.
- [BP10] John A. Baldwin and Olga Plamenevskaya. Khovanov homology, open books, and tight contact structures. *Adv. Math.*, 224(6):2544–2582, 2010.
- [Cap08] Carmen Livia Caprau. $\mathfrak{sl}(2)$ tangle homology with a parameter and singular cobordisms. *Algebr. Geom. Topol.*, 8(2):729–756, 2008.
- [CGL⁺21] Carmen Caprau, Nicolle González, Christine Ruey Shan Lee, Adam M. Lowrance, Radmila Sazdanović, and Melissa Zhang. On Khovanov homology and related invariants. In *Research directions in symplectic and contact geometry and topology*, volume 27 of *Assoc. Women Math. Ser.*, pages 273–292. Springer, Cham, [2021] ©2021.
- [CMW09] David Clark, Scott Morrison, and Kevin Walker. Fixing the functoriality of Khovanov homology. *Geom. Topol.*, 13(3):1499–1582, 2009.
- [CS93] J Scott Carter and Masahico Saito. Reidemeister moves for surface isotopies and their interpretation as moves to movies. *Journal of Knot Theory and its Ramifications*, 2(03):251–284, 1993.
- [Etn05] John B Etnyre. Legendrian and transversal knots. In *Handbook of knot theory*, pages 105–185. Elsevier, 2005.
- [ETW18] Michael Ehrig, Daniel Tubbenhauer, and Paul Wedrich. Functoriality of colored link homologies. *Proc. Lond. Math. Soc. (3)*, 117(5):996–1040, 2018.
- [Flo89] Andreas Floer. Witten’s complex and infinite-dimensional Morse theory. *J. Differential Geom.*, 30(1):207–221, 1989.
- [GLW17] J. Elisenda Grigsby, Anthony M. Licata, and Stephan M. Wehrli. Annular Khovanov-Lee homology, braids, and cobordisms. *Pure Appl. Math. Q.*, 13(3):389–436, 2017.
- [GP10] Victor Guillemin and Alan Pollack. *Differential topology*. AMS Chelsea Publishing, Providence, RI, 2010. Reprint of the 1974 original.
- [Hay23] Kyle Hayden. An atomic approach to Wall-type stabilization problems. *arXiv preprint arXiv:2302.10127*, 2023.
- [HS16] Diana Hubbard and Adam Saltz. An annular refinement of the transverse element in Khovanov homology. *Algebr. Geom. Topol.*, 16(4):2305–2324, 2016.
- [HS21] Kyle Hayden and Isaac Sundberg. Khovanov homology and exotic surfaces in the 4-ball. *arXiv preprint arXiv:2108.04810*, 2021.
- [ILM21] Damian Iltgen, Lukas Lewark, and Laura Marino. Khovanov homology and rational unknotting. *arXiv preprint arXiv:2110.15107*, 2021.
- [Kho06] Mikhail Khovanov. Link homology and Frobenius extensions. *Fund. Math.*, 190:179–190, 2006.
- [KL23] Mikhail Khovanov and Robert Lipshitz. Categorical lifting of the Jones polynomial: a survey. *Bull. Amer. Math. Soc. (N.S.)*, 60(4):483–506, 2023.
- [KM93] P. B. Kronheimer and T. S. Mrowka. Gauge theory for embedded surfaces. I. *Topology*, 32(4):773–826, 1993.
- [Lee05] Eun Soo Lee. An endomorphism of the Khovanov invariant. *Adv. Math.*, 197(2):554–586, 2005.
- [Lic62] W. B. R. Lickorish. A representation of orientable combinatorial 3-manifolds. *Ann. of Math. (2)*, 76:531–540, 1962.

- [LLS20] Tyler Lawson, Robert Lipshitz, and Sucharit Sarkar. Khovanov homotopy type, Burnside category and products. *Geom. Topol.*, 24(2):623–745, 2020.
- [LNS15] Robert Lipshitz, Lenhard Ng, and Sucharit Sarkar. On transverse invariants from Khovanov homology. *Quantum Topol.*, 6(3):475–513, 2015.
- [LS14] Robert Lipshitz and Sucharit Sarkar. A Khovanov stable homotopy type. *J. Amer. Math. Soc.*, 27(4):983–1042, 2014.
- [LT16] Robert Lipshitz and David Treumann. Noncommutative Hodge-to-de Rham spectral sequence and the Heegaard Floer homology of double covers. *J. Eur. Math. Soc. (JEMS)*, 18(2):281–325, 2016.
- [LZ19] Adam Simon Levine and Ian Zemke. Khovanov homology and ribbon concordances. *Bull. Lond. Math. Soc.*, 51(6):1099–1103, 2019.
- [Mil63] J. Milnor. *Morse theory*, volume No. 51 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1963. Based on lecture notes by M. Spivak and R. Wells.
- [MO08] Ciprian Manolescu and Peter Ozsváth. On the Khovanov and knot Floer homologies of quasi-alternating links. In *Proceedings of Gökova Geometry-Topology Conference 2007*, pages 60–81. Gökova Geometry/Topology Conference (GGT), Gökova, 2008.
- [MW22] Scott Morrison, Kevin Walker, and Paul Wedrich. Invariants of 4-manifolds from Khovanov-Rozansky link homology. *Geom. Topol.*, 26(8):3367–3420, 2022.
- [Ng05] Lenhard Ng. A Legendrian Thurston-Bennequin bound from Khovanov homology. *Algebr. Geom. Topol.*, 5:1637–1653, 2005.
- [nLa24] nLab authors. Karoubi envelope. <https://ncatlab.org/nlab/show/Karoubi+envelope>, October 2024. [Revision 39](#).
- [OS03] S. Yu. Orevkov and V. V. Shevchishin. Markov theorem for transversal links. *J. Knot Theory Ramifications*, 12(7):905–913, 2003.
- [OS05] Peter Ozsváth and Zoltán Szabó. On the Heegaard Floer homology of branched double-covers. *Adv. Math.*, 194(1):1–33, 2005.
- [Pic20] Lisa Piccirillo. The Conway knot is not slice. *Ann. of Math. (2)*, 191(2):581–591, 2020.
- [Pla06] Olga Plamenevskaya. Transverse knots and Khovanov homology. *Math. Res. Lett.*, 13(4):571–586, 2006.
- [Ras03] Jacob Andrew Rasmussen. *Floer homology and knot complements*. ProQuest LLC, Ann Arbor, MI, 2003. Thesis (Ph.D.)—Harvard University.
- [Ras10] Jacob Rasmussen. Khovanov homology and the slice genus. *Invent. Math.*, 182(2):419–447, 2010.
- [San21] Taketo Sano. Fixing the functoriality of Khovanov homology: a simple approach. *J. Knot Theory Ramifications*, 30(11):Paper No. 2150074, 12, 2021.
- [Shu07] Alexander N. Shumakovitch. Rasmussen invariant, slice-Bennequin inequality, and sliceness of knots. *J. Knot Theory Ramifications*, 16(10):1403–1412, 2007.
- [SS06] Paul Seidel and Ivan Smith. A link invariant from the symplectic geometry of nilpotent slices. *Duke Math. J.*, 134(3):453–514, 2006.
- [SS10] Paul Seidel and Ivan Smith. Localization for involutions in Floer cohomology. *Geom. Funct. Anal.*, 20(6):1464–1501, 2010.
- [SS22] Isaac Sundberg and Jonah Swann. Relative Khovanov-Jacobsson classes. *Algebr. Geom. Topol.*, 22(8):3983–4008, 2022.
- [SZ24] Matthew Stoffregen and Melissa Zhang. Localization in Khovanov homology. *Geom. Topol.*, 28(4):1501–1585, 2024.
- [Tur16] Paul Turner. A hitchhiker’s guide to Khovanov homology. In *Lectures on quantum topology in dimension three*, volume 48 of *Panor. Synthèses*, pages 91–129. Soc. Math. France, Paris, 2016.
- [Vog20] Pierre Vogel. Functoriality of Khovanov homology. *J. Knot Theory Ramifications*, 29(4):2050020, 66, 2020.
- [Wal60] Andrew H. Wallace. Modifications and cobounding manifolds. *Canadian Journal of Mathematics*, 12:503–528, 1960.
- [Wri02] Nancy Court Wrinkle. *The Markov theorem for transverse knots*. ProQuest LLC, Ann Arbor, MI, 2002. Thesis (Ph.D.)—Columbia University.
- [Zem19] Ian Zemke. Knot Floer homology obstructs ribbon concordance. *Ann. of Math. (2)*, 190(3):931–947, 2019.
- [Zha18] Melissa Zhang. A rank inequality for the annular Khovanov homology of 2-periodic links. *Algebr. Geom. Topol.*, 18(2):1147–1194, 2018.