

SPECTRA IN KHOVANOV AND KNOT FLOER THEORIES

MARCO MARENCON, SUCHARIT SARKAR, AND ANDRÁS STIPSICZ

ABSTRACT. These notes provide an introduction to the stable homotopy types in Khovanov theory (due to Lipshitz-Sarkar) and in knot Floer theory (due to Manolescu-Sarkar). They were written following a lecture series given by Sucharit Sarkar at the Rényi Institute during a special semester on “Singularities and low-dimensional topology”, organised by the Erdős Center.

CONTENTS

1. Introduction	1
2. Framed flow categories	2
2.1. A recap of Morse theory	3
2.2. From Morse functions to framed flow categories	4
2.3. The hypercube flow category	4
2.4. The Cohen-Jones-Segal construction	7
3. A Khovanov stable homotopy type	11
3.1. The Khovanov complex	11
3.2. The Lipshitz-Sarkar framed flow category	14
4. A knot Floer stable homotopy type	18
4.1. Grid diagrams	19
4.2. Grid homology	19
4.3. Maslov index 1 domains	21
4.4. Maslov index 2 domains	21
4.5. Maslov index 3 domains	22
4.6. Maslov index k domains	24
4.7. Bubbles	27
References	31

1. INTRODUCTION

Invariants of knots and links witnessed an unparalleled development in the past couple of decades. The subject of knot theory was born at the end of the XIX century, while a systematic mathematical study began with work of Reidemeister in the 1920s [Rei27]. Alexander’s work (and in particular the introduction of the Alexander polynomial invariant [Ale28]) in the late 1920s was a remarkable step in understanding this subject. The next major development was brought by work of Vaughan Jones in the 80s: his introduction of the Jones polynomial [Jon85], and subsequent generalizations, the HOMFLY-PT polynomial and the definition of quantum invariants, reconnected the subject with theoretical physics.

After the turn of the century a new paradigm, *categorification*, provided deep insight into the (still mysterious) world of knots. Khovanov's definition of his homologies [Kho00] refined the Jones polynomial, while work of Ozsváth and Szabó [OS04], and independently of Rasmussen [Ras03] produced a family of homology theories having the Alexander polynomial as their (graded) Euler characteristic. It was a routine exercise to find knots which are not distinguished by their Jones (or Alexander) polynomials, but have distinct Khovanov (or knot Floer) homologies. Indeed, both Khovanov and Heegaard Floer homologies are *unknot detectors* (i.e. the unknot is characterized by its homology in either theory). For comparison, this statement is false for the Alexander polynomial, and an open problem for the Jones polynomial.

Recently a further level of abstraction, *specification*, brought new tools in knot theory. In this context one wants to construct topological spaces (or, more generally, spectra) with the property that these objects are knot invariants, and their singular homology (in the appropriate sense) is the Khovanov homology (or knot Floer homology) of the underlying knot. Having a spectrum provides the opportunity of applying further structures (such as Steenrod operations) and possibly other functors (like K -theory) to find further knot invariants. The program has been completed for Khovanov homology [LS14], and the first step has been taken for knot Floer homology [MS21].

These notes are based on lectures of Sucharit Sarkar given at a focused week of the special semester *Singularities and low-dimensional topology* at the Erdős Center, Budapest in the Spring semester of 2023. We thank Michael Willis for his comments on an earlier version of these notes.

MM acknowledges that: This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 893282. AS was partially supported by NKFIH Grant K146401.

2. FRAMED FLOW CATEGORIES

The basic idea of *specification* via flow categories is to start from a chain complex C together with a given basis and construct a ‘space’ (more precisely a CW spectrum) X with the property that its reduced cellular chain complex $\tilde{C}_*^{\text{cell}}(X)$ is isomorphic to C . The following remark explains why we really need to choose a basis of C for our construction.

Remark 2.1. If we have a chain complex C , which we assume to be finitely generated for simplicity, and we disregard its basis and grading, then we can always make a basis change so that

$$C \cong \mathbb{Z}^k \oplus \bigoplus_i (\mathbb{Z} \xrightarrow{k_i} \mathbb{Z}),$$

where the boundary map on \mathbb{Z}^k is zero and the symbol $\mathbb{Z} \xrightarrow{k_i} \mathbb{Z}$ denotes a two-step complex where the boundary map is the multiplication by k_i , and from here we can construct a CW complex $M(C)$ with $\tilde{C}_*^{\text{cell}}(X) = C$ simply by taking an appropriate wedge of spheres (for the free part) and other simple spaces (for the 2-step complexes) as described in [Hat02, Example 2.40]. The space $M(C)$ is often called a *Moore space*, and its construction can be refined to take gradings in account as well.

The drawback of Moore spaces is that they are not (and they cannot be made) functorial. More precisely, by a result of Carlsson [Car81] the functors

$$H_*, H_n: \text{Top}_* \rightarrow \text{Ab}$$

do not admit functorial sections $M, M_n: \text{Ab} \rightarrow \text{Top}_*$.




FIGURE 1. Schematic of a proper Morse function.

This is in sharp contrast with the case of homotopy groups, where the functor


$$\pi_n: \text{Top}_* \rightarrow \text{Ab},$$

defined for $n > 1$, does admit a functorial section, given by the Eilenberg-MacLane spaces.

2.1. A recap of Morse theory. Let $f: M \rightarrow \mathbb{R}$ be a proper Morse-Smale function with finitely many critical points on a Riemannian manifold M . Each index- k critical point gives rise to a *downward disc* \mathbb{D}^k , which together with all other downward discs presents M as a (based) CW complex.

Remark 2.2. There is a subtlety in the presentation of M as a CW complex. If one tries to define the attaching map between cells just by flowing, then the resulting map would be discontinuous. See [Mil65] for the details on the attaching maps.

The construction of the Morse chain complex for f mirrors the construction of $\tilde{C}_*^{\text{cell}}$ of the associated CW complex. The generators of the Morse complex are the critical points of the Morse function f . We let $|x|$ denote the index of the critical point x . For each $x, y \in \text{Crit } f$ we let $\mathcal{M}(x, y)$ denote the unparameterised moduli space of flowlines from x to y , of dimension $|x| - |y| - 1$.



In the next subsection we will construct framed flow categories, which are ‘equivalent’ to Morse functions with gradient flowlines and to based CW complexes.




FIGURE 2. An illustration of the equality $Tℳ(x, y) \oplus ℝ \oplus T𝔻_y = T𝔻_x$.

2.2. From Morse functions to framed flow categories. This section is included mostly for motivational purposes, we describe ideas and constructions which motivate the definitions and results in Section 3; consequently, no proofs in this section will be given.

Definition 2.3. Let f be a Morse-Smale function. We define the *flow category* C_f as the category whose objects are $\text{Ob } C_f = \text{Crit } f$, and whose non-identity morphisms are given by the unparameterised moduli spaces $\mathcal{M}(x, y)$ of (possibly broken) flowlines from x to y .

In order to promote C_f to a *framed* flow category we need to endow each $\mathcal{M}(x, y)$ with a framing of the stable normal bundle. We do it as follows:

- for every downward disc D^k we choose a trivialisation of TD^k ;
- for each pair of critical points x, y with $|x| > |y|$ and connected by a flowline, there is an identification

$$Tℳ(x, y) \oplus ℝ \oplus T𝔻_y = T𝔻_x,$$


which gives a stable framing of $Tℳ(x, y)$ (see Figure 2);

- a stable framing of $Tℳ(x, y)$ induces a stable framing of the normal bundle too.

In the case when $|x| = |y| + 1$, the choice of a framing of the stable normal bundle is the same as a (coherent) system of orientations.

2.3. The hypercube flow category. This subsection is devoted to a crucial example of flow category, which will be the fundamental building block of the Khovanov stable homotopy type.

The starting point is an innocent-looking function $f: ℝ \rightarrow ℝ$ with only two critical points, at 0 and 1, which are respectively a minimum and a maximum. For example, we can choose $f(x) = x^2 \cdot \left(\frac{3}{2} - x\right)$.



Define the function $f_n: ℝ^n \rightarrow ℝ$ by setting

$$(2.1) \quad f_n(x_1, \dots, x_n) = \sum_{i=1}^n f(x_i).$$




FIGURE 3. Illustration of the flowlines from 11 to 00 and of the resulting moduli space.

Then $\text{Crit } f_n = \{0, 1\}^n$. The index of a critical point $v \in \{0, 1\}^n$ coincide with its L^1 norm:

$$|v| = ||v||_{L^1} = \sum_{i=1}^n v_i.$$

Moreover, the set $\text{Crit } f_n$ inherits a product partial order by setting $0 < 1$. For two distinct critical points $u, v \in \{0, 1\}^n$, we have that

$$\mathcal{M}(u, v) = \emptyset \quad \text{if } u \not\succ v.$$

In the next examples we will try to identify $\mathcal{M}(u, v)$ when $u > v$.

Example 2.4. The function $f = f_1: \mathbb{R} \rightarrow \mathbb{R}$ has only two critical points, and the unparameterised moduli space $\mathcal{M}(1, 0)$ consists of a single point.

Example 2.5. The function $f_2: \mathbb{R}^2 \rightarrow \mathbb{R}$ has four critical points, namely 00, 01, 10, and 11. We have two kinds of moduli spaces, depending on the index difference of u and v .

Using Example 2.4, one can check that if the index difference of two critical points of f_2 is exactly 1, then the moduli spaces $\mathcal{M}(u, v)$ consists of a single point. We name them

$$\begin{aligned} \mathcal{M}(11, 10) &= \{a\}, & \mathcal{M}(10, 00) &= \{b\}, \\ \mathcal{M}(11, 01) &= \{c\}, & \mathcal{M}(01, 00) &= \{d\}. \end{aligned}$$

The moduli space $\mathcal{M}(11, 00)$ is 1-dimensional, and its boundary consists of the two broken flowlines labelled ab and cd . The moduli space is in fact an interval with boundary $\{ab, cd\}$. One way to visualise $\mathcal{M}(11, 00)$ is to draw the segment from 10 to 01 (in \mathbb{R}^2 this segment lies on the line $x + y = 1$), and noticing that such a segment intersects each flowline from 11 to 00 exactly once. The two endpoints 10 and 01 of the segment correspond to the broken flowlines. See also Figure 3.

Exercise 2.6. Identify the flow category associated with f_3 . Note that $\mathcal{M}(111, 000)$ is a hexagon: this can be visualised as the section of the unit cube given by the plane $x + y + z = \frac{3}{2}$, which intersects each (possibly broken) flowline from 111 to 000 in exactly one point. See Figure 4.

More generally, the moduli space $\mathcal{M}(1 \cdots 1, 0 \cdots 0)$ in the hypercube flow category is a *permutohedron*.



FIGURE 4. The moduli space $\mathcal{M}(111,000)$.

FIGURE 5. A choice of orientation of the 0-dimensional moduli spaces.

Definition 2.7. The n -permutohedron Π^{n-1} is the convex hull of

$$\{(\sigma(1), \dots, \sigma(n)) \in \mathbb{R}^n \mid \sigma \in S_n\}.$$

Note that for every permutation $\sigma \in S_n$, we have $\sum_{i=1}^n \sigma(i) = \frac{n(n+1)}{2}$. Thus, Π^{n-1} is contained in an \mathbb{R}^{n-1} -slice of \mathbb{R}^n , and it is in fact an $(n-1)$ -dimensional manifold with corners.

The hypercube flow category can also be upgraded to a framed flow category. In order to do so, we need to endow the moduli space $\mathcal{M}(u, v)$, which is a permutohedron, with a *coherent framing of the stable normal bundle*. We hint at its construction by one explicit example.

Consider the square flow category, i.e. the one associated with f_2 . A stable framing of the normal bundle of the 0-dimensional moduli space corresponds to a choice of orientation of each of them (either + or -). We need to choose orientations such that exactly one or three of the points a, b, c, d are framed positively; for instance, we may choose orientations as in Figure 5.

We envision the moduli spaces $\mathcal{M}(01,00)$ and $\mathcal{M}(10,00)$ as sitting in one copy of \mathbb{R} , which we denote by $\mathbb{R}_{(i)}$, and the moduli spaces $\mathcal{M}(11,10)$ and $\mathcal{M}(11,01)$ as sitting in another copy of \mathbb{R} , which we denote by $\mathbb{R}_{(ii)}$. Thus, we get a product orientation for the broken flowlines ba and dc , embedded in $\mathbb{R}_{(i)} \times \mathbb{R}_{(ii)}$, as illustrated in the figure below.




FIGURE 6. The product orientation on the broken flowlines in the square flow category.




FIGURE 7. An embedding of the moduli space $\mathcal{M}(11,00)$ into $\mathbb{R}_{(i)} \times \mathbb{R}_{(ii)} \times \mathbb{R}_{\geq 0}$.

We now wish to orient the stable normal bundle of the moduli space $\mathcal{M}(11,00)$, which is an interval with (unoriented) boundary $\{ba, dc\}$. Such a moduli space should come with an embedding

$$\mathcal{M}(11,00) \hookrightarrow \mathbb{R}_{(i)} \times \mathbb{R}_{(ii)} \times \mathbb{R}_{\geq 0}$$

whose restriction to the boundary gives the embedding

$$\{ba, dc\} \hookrightarrow \mathbb{R}_{(i)} \times \mathbb{R}_{(ii)}$$

constructed above. The coherence condition implies that we can find a framed interval embedded in $\mathbb{R}_{(i)} \times \mathbb{R}_{(ii)} \times \mathbb{R}_{\geq 0}$ with boundary the two framed points (see Figure 7). The possible stable framings of such an embedded interval are a torsor over

$$\pi_1(SO) := \varinjlim \pi_1(SO(n)) \cong \mathbb{Z}/2\mathbb{Z}.$$

More generally, using obstruction theory we can frame the stable normal bundle coherently for all moduli spaces in the hypercube flow category.

2.4. The Cohen-Jones-Segal construction. Cohen-Jones-Segal [CJS95] provide a construction to produce a stable based CW complex out of a framed flow category, which we will outline in one particular case in this section. Their construction shares many similarities with the Pontryagin-Thom construction.

Definition 2.8. Two topological spaces X and Y are called *stably homotopy equivalent* (denoted $X \sim_S Y$) if $\Sigma^n X$ and $\Sigma^n Y$ are homotopy equivalent for $n \gg 0$.

Exercise 2.9. $\mathbb{T}^2 \sim_S S^1 \vee S^1 \vee S^2$.

Hint: The attaching map of the 2-cell of \mathbb{T}^2 is a commutator in π_1 of the 1-skeleton.


Exercise 2.10. $\mathbb{CP}^2 \not\sim_S S^2 \vee S^4$.

Hint: Steenrod squares are preserved under stable homotopy equivalence.

We remark that Exercise 2.9 shows that the cup product is not preserved under stable homotopy equivalence.

We will now turn to the construction of a specific example of a stable based CW complex, namely the one associated with the square framed flow category.

Each vertex v of the square will give rise a cell of dimension $|v| + 2$. The resulting stable homotopy type is that of a point.



Given $R \gg \varepsilon > 0$, we parameterise the cells as follows:

$$\begin{aligned} C(11) &:= [0, R] \times [-R, R]^2 \times [0, R]; \\ C(10) = C(01) &:= [0, R] \times [-R, R] \times [-\varepsilon, \varepsilon]; \\ C(00) &:= [-\varepsilon, \varepsilon]^2. \end{aligned}$$

We construct the based CW complex by attaching the cells in order of increasing dimension. We start with the 2-cell $C(00)$, which is attached to the basepoint $*$ via the unique possible attaching map $\varphi_{00}: \partial C(00) \rightarrow *$, see Figure 8. We denote the resulting CW complex by X_0 , i.e.

$$X_0 := C(00) \cup_{\varphi_{00}} \{*\}.$$

We now move to the 3-cells. See Figure 9 for reference. We view the $[-R, R]$ factor as the $\mathbb{R}_{(i)}$ direction, and we recall that we have an embedding of b and d into $\mathbb{R}_{(i)}$, as framed points.

On the boundary of the 3-cell $C(10)$ (resp. $C(01)$), we draw a copy of the $[-\varepsilon, -\varepsilon]^2$ on $\{0\} \times [-R, R] \times [-\varepsilon, \varepsilon]$, centred at $(0, b, 0)$ (resp. $(0, d, 0)$).

The maps

$$\begin{aligned} \varphi_{10}: \partial C(10) &\rightarrow C(00) \cup_{\varphi_{00}} \{*\} \\ \varphi_{01}: \partial C(01) &\rightarrow C(00) \cup_{\varphi_{00}} \{*\} \end{aligned}$$

are defined as follows:

- on the embedded copies of $[-\varepsilon, \varepsilon]^2$ the maps restrict to degree ± 1 maps to $C(00)$, with the sign of the degree depending on the framing of b or d ;




FIGURE 8. An illustration of the CW complex X_0 built as part of the process of constructing the stable homotopy type associated with the square framed flow category.

- on the complement of the embedded copy of $[-\varepsilon, \varepsilon]^2$, the maps send everything to the basepoint $\{*\}$.

The maps so defined descend to continuous maps to $C(00) \cup_{\varphi_{00}} \{*\}$. Using the attaching maps, we construct the CW complex

$$X_1 := (C(10) \cup C(01) \cup X_0) / \sim.$$

We now turn to the last remaining cell, namely $C(11)$. Recall that

$$C(11) = [0, R] \times [-R, R]^2 \times [0, R].$$

We need to define an attaching map

$$\varphi_{11}: \partial C(11) \rightarrow X_1.$$

We view $\partial C(11)$ as the union of three pieces:

- $\{0\} \times [-R, R]^2 \times [0, R]$;
- $[0, R] \times [-R, R]^2 \times \{0\}$;
- the part of $\partial C(11)$ with 1st and 4th coordinate both non-zero.

Note that the first and the second piece intersect at $\{0\} \times [-R, R]^2 \times \{0\}$.

We consider a copy of $C(01)$ and a copy of $C(10)$ sitting in

$$[0, R] \times [-R, R]^2 \times \{0\} \subset \partial C(11)$$

as the slices

$$[0, R] \times [-R, R] \times [a - \varepsilon, a + \varepsilon] \times \{0\} \quad \text{and} \quad [0, R] \times [-R, R] \times [c - \varepsilon, c + \varepsilon] \times \{0\}$$

respectively, as shown in Figure 10.

We define the map φ_{11} separately on the three pieces of $\partial C(11)$. We start from the blue piece $[0, R] \times [-R, R]^2 \times \{0\}$: with the above picture in mind, it is natural to map the embedded cells $C(01)$ and $C(10)$ inside $\partial C(11)$ to the respective abstract ones via degree ± 1 maps, depending on the framings of a and c , and to extend this to a map from $[0, R] \times [-R, R]^2 \times \{0\}$ to X_1 by sending everything in the complement to $*$.

Then, a naïve approach to extend φ_{11} to all of $\partial C(11)$ would be to map the purple and the black pieces to the basepoint $*$. However, the map so defined would not be continuous, because the whole purple piece of $\partial C(11)$ would be mapped to $*$, except the two red squares in $\{0\} \times [-R, R]^2 \times \{0\}$, which lie in $\partial(\{0\} \times [-R, R]^2 \times [0, R])$. Thus, we modify the map in the purple portion of $\partial C(11)$ by using the framed interval from ab to cd . We envision the framed interval constructed above as sitting inside $\{0\} \times [-R, R]^2 \times [0, R]$, where the




FIGURE 9. An illustration of the CW complex X_0 built as part of the process of constructing the stable homotopy type associated with the square framed flow category.

last coordinate corresponds to $\mathbb{R}_{\geq 0}$. A regular neighbourhood of it can be parameterised by $[-\varepsilon, \varepsilon]^2 \times I$. We can then define the map φ_{11} on $\{0\} \times [-R, R]^2 \times [0, R]$ by setting it on the neighbourhood of the framed interval as

$$[-\varepsilon, \varepsilon]^2 \times I \xrightarrow{pr_1} [-\varepsilon, \varepsilon]^2 \xrightarrow{\sim} C(00) \rightarrow X_0 \subset X_1,$$

and by sending the complement to $*$. See Figure 11.

Thus, we have so far defined φ_{11} on $\{0\} \times [-R, R]^2 \times [0, R]$ and $[0, R] \times [-R, R]^2 \times \{0\}$. We extend it to $\partial C(11)$ by sending the rest of $\partial C(11)$ to $*$.




FIGURE 10. The “blue” piece $[0, R] \times [-R, R]^2 \times \{0\}$ of $\partial C(11)$.

3. A KHOVANOV STABLE HOMOTOPY TYPE

Our next goal is to construct a Khovanov homotopy type starting from the Khovanov complex. We quickly recall the construction of the Khovanov complex. For a more detailed exposition we refer the reader to [Kho00, BN02].

3.1. The Khovanov complex. Given a link diagram and a crossing, we define the 0-smoothing and the 1-smoothing of the crossing as in Figure 12. Let D be a link diagram with n crossings, labelled in some order from 1 to n . See Figure 13 for an example. Each $v \in \{0, 1\}^n$ determines a *complete resolution* D_v by smoothing crossing i according to the i -th entry of v . Topologically, each complete resolution consists of an unlink. We can organise the complete resolutions on a hypercube of dimension n , so that each v corresponds to a vertex of the hypercube. See Figure 14.

Let the *weight* or *index* $|v|$ of v be the number of coordinates of v with value 1. Two vertices u and v are connected by an edge e of the hypercube if and only if u and v differ only at one entry, and in such a case the edge is oriented from the vertex with *higher* weight to the vertex with *lower* weight. Note that this convention is opposite to the standard one in the literature (e.g., [Kho00, BN02]); equivalently, in these notes we consider Khovanov homology as opposed to the more popular Khovanov cohomology.




FIGURE 11. The “purple” piece $\{0\} \times [-R, R]^2 \times [0, R]$ of $\partial C(11)$.




FIGURE 12. The 0-smoothing and the 1-smoothing of a crossing.




FIGURE 13. A non-minimal diagram of the 2-component unlink, together with a labelling of the crossings.

From the cube of resolution Khovanov constructed a complex as follows. Let $V = \mathbb{Z}\langle 1, x \rangle$ be the \mathbb{Z} -module spanned by two elements 1 and x . To each $v \in \{0, 1\}^n$ associate the vector space

$$V_v = \bigotimes_{\text{components of } D_v} V,$$




FIGURE 14. The cube of resolutions of the diagram for the unlink in Figure 13. We have also fixed a total ordering on the circles at each resolution: the resolutions, indicated by the labels I and II for the circles at the resolutions 10 and 01. (The other resolutions have only one circle, so there is only one total order there.)

where each factor V is considered as associated to a link component of D_v . The module underlying the Khovanov complex is defined as

$$\text{CKh}(D) = \bigoplus_{v \in \{0,1\}} V_v.$$

Recall that in the construction of framed flow category *spacifying* a given chain complex C one needs to specify a basis of C as well, see Remark 2.1. Luckily for us, the \mathbb{Z} -module $\text{CKh}(D)$ comes with a preferred basis. For each $v \in \{0,1\}$, the \mathbb{Z} -module V_v has a basis consisting of all possible tensor products of the generators 1 and x for V ; we call them the *standard generators* of V_v , and we denote them with a v subscript, e.g. if D_v has four components, after totally ordering them, the element $1x11_v$ is the Khovanov generator that labels the second component by x and other components by 1. As v varies in $\{0,1\}$, all standard generators give a basis for $\text{CKh}(D)$.

In order to define the differential on $\text{CKh}(D)$, Khovanov defined two maps m and Δ , associated to the merge and split cobordisms respectively:

$$\begin{array}{ll} m: V \otimes V \rightarrow V & \Delta: V \rightarrow V \otimes V \\ 1 \otimes 1 \mapsto 1 & 1 \mapsto 1 \otimes x + x \otimes 1 \\ 1 \otimes x \mapsto x & x \mapsto x \otimes x \\ x \otimes 1 \mapsto x & \\ x \otimes x \mapsto 0 & \end{array}$$

If there is an edge from u to v , then there is an elementary cobordism between D_u and D_v , which is a split saddle or a merge saddle. We define a map

$$\partial_{u,v}: V_u \rightarrow V_v$$

as the map m or Δ on the V factors associated with the components of D_u and D_v on which the saddle cobordism happens, and as the identity map on all other V factors.

Fix a framing of all the 0-dimensional moduli spaces in the hypercube flow category associated to the function f_n defined in Equation (2.1) at page 4, such that on every square




FIGURE 15. The Khovanov complex associated with the diagram for the 2-component unlink in Figure 13. To simplify the notation, we suppressed the tensor product symbol \otimes and the subscripts specifying the vertex of the hypercube. The relative location of the standard generators indicates the vertex of the hypercube to which it belongs. Compare with Figure 14. For example, the four standard generators located at the left are in $V_{(10)}$; among them, the generator $1x$ is the one that labels the first circle (marked I) by 1 and the second circle (marked II) by x .

face, exactly one or three of the points are framed positively. The differential is then defined by summing over all edges:

$$\partial := \sum_{\substack{|u-v|=1 \\ |u|=|v|+1}} (-1)^{s(u,v)} \partial_{u,v},$$

where $s(u, v) = 0$ if $\mathcal{M}(u, v)$ is framed positively and $s(u, v) = 1$ if $\mathcal{M}(u, v)$ is framed negatively.

As an illustration, Figure 15 shows the Khovanov complex for the diagram of the 2-component unlink in Figure 13. The figure shows the 12 standard generators of the Khovanov complex, organised to mirror the structure of the cube of resolutions (Figure 14). The arrows encode the differential: if there is an arrow from y to z , then z appears in the expansion of ∂y with coefficient ± 1 depending on the sign label on the arrow. (For the moment disregard the fact that some arrows are coloured in red and some of the generators are circled in green.) The homology of $(\text{CKh}(D), \partial)$ is the Khovanov homology of the link. However, we do not take homology, but instead construct a framed flow category from the complex.

3.2. The Lipshitz-Sarkar framed flow category. We will use the example from Figures 13 and 14 to describe the framed flow category $\mathcal{X}(D)$ from [LS14].

We start from the set of objects, $\text{Ob } \mathcal{X}(D)$, which consists of all standard Khovanov generators. Thus, for our recurring example of the diagram of the 2-component unlink in Figure 13, the **objects** are the 12 generators that are shown in Figure 15.

The **0-dimensional moduli spaces** are framed points, given by the Khovanov differential. Following the convention in Figure 15, the elements of 0-dimensional moduli spaces are the arrows shown in the figure, together with a framing specified by the sign label: for example,

$$\mathcal{M}(x_{(11)}, xx_{(01)}) = \{p_-\}, \quad \mathcal{M}(x_{(11)}, x1_{(01)}) = \emptyset,$$




FIGURE 16. The two possibilities for $\mathcal{M}(1_{(11)}, x_{(00)})$ from the example in Figure 15. The *right choice* is shown in red, and the *left choice* is shown in blue. The two possible matchings are called *ladybug* matchings, and a choice between the two of them must be done each time a ladybug configuration as in Figure 17 appears.

where p_- is a negatively framed point. More precisely, given two objects / standard generators \mathbf{y} and \mathbf{z} with $|\mathbf{y}| = |\mathbf{z}| + 1$, we represent $\partial \mathbf{y} \in CKh(D)$ as a linear combination of standard generators, and we take the coefficient $\langle \partial \mathbf{y}, \mathbf{z} \rangle$ of \mathbf{z} in such an expansion. Note that by the definition of the Khovanov differential ∂ , the only values $\langle \partial \mathbf{y}, \mathbf{z} \rangle$ can take are $+1, -1$, and 0 . Then, we set

$$\mathcal{M}(\mathbf{y}, \mathbf{z}) = \begin{cases} \{\text{a positively framed point}\} & \text{if } \langle \partial \mathbf{y}, \mathbf{z} \rangle = +1 \\ \{\text{a negatively framed point}\} & \text{if } \langle \partial \mathbf{y}, \mathbf{z} \rangle = -1 \\ \emptyset & \text{if } \langle \partial \mathbf{y}, \mathbf{z} \rangle = 0 \end{cases}$$

We will construct the higher moduli spaces inductively as disjoint unions of permutohedra, using the fact that $\partial \mathcal{M}$ is the union of products of lower-dimensional moduli spaces already constructed.

Most of the higher moduli spaces are constructed without making choices, but there is one interesting case, which appears during the construction of the **1-dimensional moduli spaces**, where a choice is needed. This is the case of $\mathcal{M}(1_{(11)}, x_{(00)})$ illustrated in Figure 15: $1_{(11)}$ and $x_{(00)}$ are the two standard generators circled in green; the broken flowlines between them, which constitute the boundary of the moduli space, are obtained by following the red arrows. Thus, we can see that

$$\partial \mathcal{M}(1_{(11)}, x_{(00)}) = \{ea, fb, gc, hd\},$$

where a, b, c, d, e, f, g , and h are the red arrows in Figure 15. Of the four broken flowlines, two (namely ea and fb) factor through the vertex 10, and the other two (namely gc and hd) factor through the vertex 01. Note that the vertex through which the broken flowline factors determines its framing.

Thus, if we want it to be a union of permutohedra, $\mathcal{M}(1_{(11)}, x_{(00)})$ should consist of two intervals, which can be chosen in two sensible ways: either an interval connecting ea to gc and one connecting fb to hd (the *right choice*), or an interval connecting ea to hd and one connecting fb to gc (the *left choice*). See Figure 16. In a bombastic moment of self-confidence, we claim that we always make the right choice! The 1-dimensional moduli spaces so defined can be upgraded to *framed* moduli spaces, using a framing on the cube flow category.

Remark 3.1. In the language of [LS14], we just made a choice of a *ladybug matching*. The name comes from the fact that the diagram D can be alternatively specified by its all-1 resolution, together with a collection of embedded arcs that record the band surgeries needed to turn a 1-smoothing into a 0-smoothing. Figure 17 shows the all-1 resolution and the surgery arcs for the diagram in Figure 13. The shape resembling a ladybug motivates the name *ladybug configuration*.




FIGURE 17. A ladybug matching. The figure represents an all-1 resolution of a link diagram (in fact, of the diagram in Figure 13). The red arcs, one for each crossing of the link diagram, trace the band surgeries that turn each 1-smoothing into a 0-smoothing.




FIGURE 18. The figure on the left shows a link diagram for the 2-component unlink. The figure on the right represents its all-1 resolution together with the arcs that turn each 1-smoothing into a 0-smoothing.

Each time a ladybug configuration appears we must choose a ladybug matching. The fact that we always make the right choice will be important for the next step, namely the construction of the 2-dimensional moduli spaces.

Surprisingly, we will see that the choice of the ladybug matchings on the 1-dimensional moduli spaces is *the only choice we need* to make, under the assumption that each moduli space $\mathcal{M}(y, z)$ is a disjoint union of permutohedra.


We now turn to the construction of the **2-dimensional moduli spaces**. We start from the example in Figure 18, whose associated cube of resolutions is shown in Figure 19. One can check that the 3-dimensional cube of resolutions has 2 ladybug faces, which are shaded in light red.

We consider the moduli space

$$\mathcal{M}(\text{Diagram}^1, \text{Diagram}^x).$$


Its boundary (already defined, since it is the union of products of lower-dimensional moduli spaces) contains 12 fully broken flowlines:

- 4 fully broken flowlines for each ladybug face;
- 2 fully broken flowlines factoring through



namely the ones illustrated in Figure 20;

- 2 fully broken flowlines factoring through



which the reader is encouraged to find by themselves.

The boundary of such a moduli space is therefore a 1-dimensional manifold with 12 corners (one for each fully broken trajectory). It turns out that it must be either a dodecagon




FIGURE 19. The cube of resolutions for the link diagram in Figure 18. The faces shaded in red correspond to ladybug configurations.

(i.e., 12-gon) or the disjoint union of two hexagons (i.e., 6-gon). Lipshitz and Sarkar [LS14] showed that if we choose all the ladybug matchings to be right (or all of them to be left), then we get two hexagons. Thus, we can define the moduli space

$$\mathcal{M}(\text{Diagram}_1, \text{Diagram}_2).$$

as the disjoint union of two filled hexagons, which are permutohedra (see Exercise 2.6). Once again, these inherit framings from a framing of the cube flow category.

Finally, we turn to the **moduli spaces of dimension ≥ 3** . Their definition is actually easier than in the cases of lower dimension.

Assume that \mathbf{y} and \mathbf{z} are standard generators of the Khovanov chain complex, such that $|\mathbf{y}| > |\mathbf{z}| + 3$. By induction, we assume that the boundary of the moduli space $\mathcal{M}(\mathbf{y}, \mathbf{z})$ is already defined. It turns out that $\partial \mathcal{M}(\mathbf{y}, \mathbf{z})$ is a regular cover of the moduli space in the cube flow category, i.e. a permutohedron. If $n := |\mathbf{y}| - |\mathbf{z}| > 3$, then the permutohedron Π^{n-1} has simply connected boundary, and therefore every regular cover of it is trivial. Thus, there is




FIGURE 20. The figure shows two fully broken flowlines in the boundary of a 2-dimensional moduli space arising from the cube of resolutions associated to the link diagram in Figure 18.

no obstruction to defining the moduli space $\mathcal{M}(\mathbf{y}, \mathbf{z})$ as a disjoint union of permutohedra of dimension $n - 1$, which also inherit framings from a framing of the cube flow category.

4. A KNOT FLOER STABLE HOMOTOPY TYPE

Once we have the ‘specification’ of Khovanov homology, and witness its success, it is natural to ask a similar construction for invariants of knots defined in the Heegaard Floer context. As the construction requires a good understanding of ‘moduli spaces’, it is instructive to start the process with the version of knot Floer homology where the most control on the moduli spaces is available: grid homology. For an introduction to grid homology see [MOS09, MOST07, OSS15]. For general (pointed or doubly pointed) Heegaard diagrams, the determination of the moduli spaces required for describing the flow category is analytically extremely challenging, and indeed in general these spaces depend on further choices (such as an almost complex structure, or perturbations), which are hard to incorporate into the theory. These difficulties can be bypassed by the combinatorial description provided by grid homology.

Despite the simplicity of the description of the chain complex involved in the definition of grid homology, the presence of ‘bubbles’ provides serious difficulties for setting up the stable homotopy type. These bubbles do arise already on the homological level, but they cancel out, hence in grid homology they are ‘invisible’ to the boundary map, resulting in a (relatively) simple theory. In the construction of the framed flow category, however, the bubbles cannot be ignored anymore, bringing in an extra level of complexity in the theory. In this section (following [MS21]) the construction of a spectrum with stable homology giving grid homology is outlined. The appropriate topological invariance of the spectrum, that is, the proof of the fact that the result is a knot-link invariant is, however, still open. This invariance would provide the real power of the theory, since it would open the way to construct further knot invariants, which might go beyond the mere homology theory given by grid homology.




FIGURE 21. The left hand side shows a grid diagram for the left-handed trefoil knot illustrated on the right. The full and the hollow dots determine two grid states x and y respectively. A small square (out of the 25 ones for this grid) is shaded. The shaded small square is also a rectangle from x to y .

4.1. Grid diagrams. Recall that every knot (and even link) in the standard three-sphere S^3 can be presented by a grid diagram, a combinatorial object with the following properties:

Consider the grid of n^2 small squares in an $n \times n$ square in the plane for some $n \in \mathbb{N}$. Suppose that each column and each row has exactly one marking X and exactly one marking O in one of the small squares in the column or row at hand. By connecting the X in a column to the O in the same column, and the O in the row to the X in the same row (oriented by this order), and with the convention that a vertical segment always passes over a horizontal one, the grid diagram provides the diagram of a (PL embedded) link in S^3 . It is not hard to see that every knot and link can be presented in this way, and a method similar to the proof of Reidemeister's fundamental result in knot theory provides two types of moves with the property that grids defining isotopic knots/links can be transformed into each other by a sequence of these 'grid moves'.

Indeed, by identifying the top and bottom edges of the big square, and doing the same with the left-most and right-most edges, we get a grid diagram on the standard torus of S^3 . The two grid moves of commutation and stabilization then provide a similar theory for knots in this context.

4.2. Grid homology. The chain complex of grid homology is generated by grid states: a grid state x is simply a bijection between the horizontal and vertical circles in the toroidal grid. A geometric presentation of such an object can be given by choosing an intersection point of the vertical and horizontal circles in the grid in such a way that each vertical and each horizontal circle contains exactly one chosen point. The differential in the chain complex is given by counting "empty rectangles", cf. Figure 22. In more detail, to a pair of two distinct grid states x, y we associate 0 if x and y differ in more than 2 coordinates. If the difference is in exactly two circles, then the circles of these coordinates (there are two vertical and two horizontal of those) partition the torus into four rectangles, and by choosing an appropriate orientation convention, two of them will point from x to y and two from y to x . Now we count the rectangles from x to y which are *empty* (i.e., which do not contain any further coordinate of x) and do not contain X -marking. The actual set-up of

the homology theory is slightly more involved (requires a system of weights determined by the O -markings, and consequently is defined over a more elaborate ring); we will not go into the details of the theory here, and advise the interested reader to consult [OSS15] for a more complete account. We just mention here that a crucial step in building the theory is to show that the boundary map ∂ (counting empty rectangles with some weights) has square zero, and therefore gives rise to a chain complex, and hence to homologies. The geometric input in the proof of $\partial^2 = 0$ is an incarnation of Gromov's compactness result, showing that the compositions of two rectangles (counted by ∂^2) come in pairs. (Over the field of two elements this fact is sufficient, over the integers a sign assignment should be also fixed so these pairs contribute zero to the final count.) In turn, this pairing follows from a simple planar geometric observation decomposing the union of two rectangles in two different ways — this is depicted in Figures 23 and 24. The situation is somewhat more complicated in further versions of the theory (where we allow the empty rectangle to contain X -markings), as the concatenation of two rectangles might add up to a strip around the torus; this contribution will be cancelled by another such strip, as shown in Figure 33. The significance of this last small problem becomes more crucial in the specification process, as it will be hinted later on.

The specification of grid homology will once again be done by constructing a framed flow category C associated with a grid diagram. The objects of the category are the grid states, i.e. the n -tuples of intersection points of the vertical and horizontal circles such that each vertical and each horizontal circle contains exactly one chosen point. Given two grid states \mathbf{x} and \mathbf{y} , the moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ will be a smooth manifold with corners. By contrast with the Khovanov homotopy type, the construction is more involved and the moduli spaces will not always be disjoint unions of permutohedra. In the most general case, the moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ will be obtained by gluing together the moduli spaces $\mathcal{M}(D)$ associated with all the *domains* from \mathbf{x} to \mathbf{y} with Maslov index $|\mathbf{x}| - |\mathbf{y}|$. Recall from [OSS15] that a domain from \mathbf{x} to \mathbf{y} in a given grid diagram is a formal linear combination

$$D = \sum n_{i,j} S_{i,j}$$

of (the closures) of the small squares $\{S_{i,j}\}_{i,j=1,\dots,n}$ with the following properties:

- the horizontal part of the boundary (i.e. the part of the boundary ∂D as a 2-chain on the horizontal circles) is a 1-chain $\partial_h D$ with $\partial(\partial_h D) = \mathbf{y} - \mathbf{x}$, while
- the vertical part of the boundary $\partial_v D$ satisfies $\partial(\partial_v D) = \mathbf{x} - \mathbf{y}$.

The domain $D = \sum n_{i,j} S_{i,j}$ is *positive* if all $n_{i,j} \geq 0$.

The Maslov index $\mu(D)$ of a domain D is the formal dimension of the space of holomorphic representatives of the domain (when considered as a 2-chain); it admits a rather simple formula, which will be sufficient in our subsequent discussion. For a domain $D = \sum n_{i,j} S_{i,j}$ from \mathbf{x} to \mathbf{y} we define the point measure $p(D)$ as follows: for each $x_i \in \mathbf{x}$ (i.e. an intersection point of the horizontal and vertical lines which belongs to \mathbf{x}) we take the average of the four numbers $n_{i,j}$ corresponding to the four small squares meeting at x_i . Then $p_{\mathbf{x}}(D)$ is the sum of these local contributions for all $x_i \in \mathbf{x}$, and $p(D) = p_{\mathbf{x}}(D) + p_{\mathbf{y}}(D)$. The Maslov index of D is then

$$\mu(D) = p(D).$$

Remark 4.1. The general formula for the Maslov index in Heegaard Floer theory also involves a term originating from the geometry of D (called the Euler measure of D); as this term is additive and vanishes for the small squares, in a grid diagram all domains have vanishing Euler measure.

In the next subsections we will give examples of the moduli spaces associated with some domains, and in Subsection 4.6 we describe an inductive way (based on the Maslov index) for constructing these spaces.

4.3. Maslov index 1 domains. These are the domains we need to encounter in the boundary map of grid homology. When \mathbf{x} and \mathbf{y} can be connected by a rectangle (that is, they differ in two circles), the moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ contains two, one or zero points depending on the number of empty rectangles connecting \mathbf{x} and \mathbf{y} . An example of an empty rectangle is shown by Figure 22.




FIGURE 22. A rectangle connecting the generators \mathbf{x} (full circle) and \mathbf{y} (hollow circle). Note that the two generators have all further coordinates equal, symbolized by the crossing on the right.

This case corresponds to the connecting domains with Maslov index $\mu(D) = 1$. The space associated to such a domain is a single point, hence the moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ (the morphisms between the two objects \mathbf{x} and \mathbf{y} in the flow category) consists of two, one or zero points.

4.4. Maslov index 2 domains. In the next step we consider domains with Maslov index $\mu(D) = 2$ connecting the grid states \mathbf{x} and \mathbf{y} . These domains are relatively easy to understand: they are concatenations of two rectangles. A simple analysis shows that this can happen in two ways. In the first case the two rectangles ‘move’ different coordinate pairs (so the two rectangles move four coordinates in total). Geometrically in this case the domain D is the union of two rectangles (which might not be disjoint, but they do not share sides), see Figure 23.

In the second case one of the moving coordinates of the first rectangle is also a moving coordinate of the second one. In this case, in geometric terms the domain D is an L-shaped region in the grid torus, see Figure 24.

The domain D can be decomposed into the concatenation of two empty rectangles in two different ways, once as a rectangle D_1 from \mathbf{x} to a grid state \mathbf{w} composed with a rectangle D_2 from \mathbf{w} to \mathbf{y} , and also as the composition of D'_1 from \mathbf{x} to \mathbf{z} with D'_2 from \mathbf{z} to \mathbf{y} . When the two rectangles move four coordinates, the geometric presentation is quite obvious: D_1 (as a subset of the grid torus) agrees with D'_2 and D_2 with D'_1 . When the rectangles share moving coordinates, the decomposition is given by the two ways an L-shaped domain decomposes as the union of two rectangles. This is indicated in Figure 24.

Indeed, these decompositions correspond to the broken flowline picture, hence we associate to such a domain an interval as moduli space, with the two endpoints corresponding to the two products of one-point moduli spaces, as it is shown in Figure 25. The moduli space in this case associated to the domain D is a manifold with boundary (an interval, really), and $\mathcal{M}(\mathbf{x}, \mathbf{y})$ will be (as always) the union of all the spaces associated to Maslov index 2 domains from \mathbf{x} to \mathbf{y} .




FIGURE 23. Each of the two figures show a Maslov index 2 domain from the (black) full circle to the hollow circle. Geometrically, the domain consists of two rectangles, which can be disjoint (as in the left) or overlapping (as on the right). In either case, the Maslov index 2 domain can be expressed as the concatenation of two rectangles (i.e., doing one rectangle after the other) in two different ways, depending on which rectangle comes first. One such concatenation factors through the grid state given by red circles, and the other one factors through the grid state given by blue circles. The resulting moduli space is a segment, as in Figure 25, with the endpoints (broken flowlines) given by the two concatenations just described.




FIGURE 24. The L -shaped region decomposes into two rectangles in two different ways. Such a region is a domain from \mathbf{x} (full circles) to \mathbf{y} (hollow circles); the two decompositions factor through \mathbf{w} (red circles) or \mathbf{z} (blue circles).

4.5. Maslov index 3 domains. The same scheme applies if the domain D has Maslov index $\mu(D) = 3$, and D moves 6 coordinates, so it can be decomposed into three ‘independent’ rectangles, each moving different pairs of coordinates, as shown in Figure 26. The by now customary argument suggests that we need to associate to D a permutohedron, i.e. a hexagon in this dimension.




FIGURE 25. The figure on the left shows the broken flowlines for a domain D of Maslov index 2 (such as those in Figures 23 and 24). The figure on the right show the moduli space $\mathcal{M}(D)$.




FIGURE 26. A domain with Maslov index 3 (on the left), and its moduli space (on the right). The boundary of the moduli space consists of all possible (partially or fully) broken flowlines: in particular, the vertices of the hexagon correspond to doubly broken flowlines, whereas the edges consist of once-broken flowlines.

There are many more possibilities for a domain D to have Maslov index equal to 3, and to still have a hexagon as its moduli space $\mathcal{M}(D)$; see Figure 27 for some examples.

The situation, however, can be much more complicated, and further configurations lead to other natural choices: the domain shown by Figure 28, for example, provide 4 possible starts from x when decomposed into rectangles, and altogether there are 8 broken trajectories, hence the moduli space is an octagon (see Figure 29).

A further example is provided by a rectangle which is not empty (contains a further coordinate of x and therefore of y). This example has Maslov index 3, and there are four broken trajectories in the decomposition, so the natural choice for the moduli spaces is a rectangle, depicted in Figure 30.




FIGURE 27. Other domains with Maslov index 3.




FIGURE 28. Both sides show a more involved domain with Maslov index 3. The 4 rectangles shown within the given domain are the 4 possible starts of the 8 decompositions of the domain into rectangles. Each such decomposition corresponds to a fully broken flowlines.




FIGURE 29. The structure of the eight fully broken flowlines and the moduli space associated with the domain of Maslov index 3 from Figure 28. The reader is invited to identify the eight fully broken flowlines (corresponding to the eight vertices of the octagon) and the partially broken flowlines (corresponding to the sides of the octagon).

4.6. Maslov index k domains. The above *ad hoc* arguments is extended to any Maslov index in a more systematic way by induction, using obstruction theory. There are several complications along the way, though.




FIGURE 30. A domain with Maslov index 3 (on the left) from \mathbf{x} (full circles) to \mathbf{y} (hollow circles), the fully broken flowlines (in the middle), and the moduli space (on the right). The reader is invited to work out the boundary of the moduli space.

First of all, remember that in the flow category we require the moduli spaces to be framed manifolds (indeed, manifolds with corners). Because of the presence of ‘bubbles’ in Heegaard Floer homology, however, the natural choices for spaces associated to (positive) domains are not always manifolds with corners — we need to consider more complicated stratified spaces. On the positive side, the bubbles will guide us to glue the stratified spaces associated to the domains connecting two grid states \mathbf{x} and \mathbf{y} to get a framed manifold with corner, which will then be our choice for the morphism space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ in the flow category. To explain the main ideas, first we will ignore the complications brought by the bubbles, and describe the naïve inductive idea to define the manifold with corner associated to a domain D of Maslov index $k+1$. We will return to the discussion about the role of bubbles in Subsection 4.7.

Our inductive hypothesis is therefore that we have already constructed the framed moduli spaces for all (positive) domains of Maslov index $\leq k$; consider now a domain D with Maslov index $\mu(D) = k+1$. By the inductive hypothesis, the framed boundary $\partial\mathcal{M}(D)$ of the moduli space $\mathcal{M}(D)$ is an already defined $(k-1)$ -dimensional framed manifold with corners, and hence provides an element $o(D)$ of the framed cobordism group Ω_{k-1}^{fr} . The latter is the group of $(k-1)$ -dimensional framed cornered manifolds taken up to framed cornered cobordism, with group operation given by the disjoint union.

Remark 4.2. In reality the boundary defines an element in a more complicated group, which is denoted by $\bar{\Omega}_{k-1}^{\text{fr}}$ in [MS21]: while its definition differs from that of Ω_{k-1}^{fr} , it turns out to be isomorphic to it. Once again, we will ignore this subtlety here.

Thus, for each (positive) domain D with Maslov index $\mu(D) = k+1$, we get an element $o(D) \in \Omega_{k-1}^{\text{fr}}$ (the notation indicates that it is an obstruction class). Putting these classes together, we get a k -cochain on a chain complex generated by the domains, with values in Ω_{k-1}^{fr} .

Before proceeding any further, let us discuss the chain complex generated by the positive domains — this chain complex will be called the *obstruction chain complex* and denoted by CD_* .

Definition 4.3. The chain complex $CD_* = CD_*(\mathbb{G})$ is freely generated over \mathbb{Z} by the positive domains and graded by their Maslov indices:

$$CD_k(\mathbb{G}) = \mathbb{Z}\langle(\mathbf{x}, \mathbf{y}, D) \mid D \in \mathcal{D}^+(\mathbf{x}, \mathbf{y}), \mu(D) = k\rangle.$$

Moreover, CD_* is equipped with the endomorphism ∂ , given by

$$\partial(D) = \sum_{\substack{(R, E) \in \mathcal{R}(\mathbf{x}, \mathbf{w}) \times \mathcal{D}^+(\mathbf{w}, \mathbf{y}) \\ R*E=D}} s(R)E + (-1)^k \sum_{\substack{(E, R) \in \mathcal{D}^+(\mathbf{x}, \mathbf{w}) \times \mathcal{R}(\mathbf{w}, \mathbf{y}) \\ E*R=D}} s(R)E.$$

Here $\mathcal{R}(\mathbf{x}, \mathbf{y})$ denotes the rectangles from \mathbf{x} to \mathbf{y} , and s denotes a sign assignment.

It is not hard to see that

Lemma 4.4. *The pair (CD_*, ∂) is a chain complex, that is, $\partial^2 = 0$.*

On the other hand, it requires a rather tedious calculation to show that the homology of this complex is rather simple:

Lemma 4.5. *The homology of (CD_*, ∂) is isomorphic to \mathbb{Z} , supported in grading 0.*

Because of the presence of possible bubbles (a.k.a. *boundary degenerations*), however, later we are forced to consider a more complicated chain complex, an enhancement CDP_* of CD_* where also vectors of partitions and further data describing the positions of the bubbles (and their relative distances) are added to the generators. These extra data will be essential in locating possible bubbles, hence play a crucial role in gluing arguments. (We will omit the details of this extension here.)

The ‘earlier defined’ cochains (based on the definition of the obstruction $o(D)$) naturally live in a subcomplex of the dual CDP^* of CDP_* , and an important technical step in the construction is the calculation showing that this subcomplex is acyclic.

By associating the obstruction class $o(D)$ to a Maslov index $k + 1$ positive domain, we get a $(k + 1)$ -cochain, which turns out to be a cocycle, that is

$$\delta(o) = 0,$$

where δ is the dual coboundary map. This is proved by showing that $o(\partial D) = 0$ for all positive domains D with $\mu(D) = k + 2$. As the cochain complex is acyclic, o is a coboundary, that is, there is a map f on index k domains associating

$$E \mapsto f(E) \in \Omega_{k-1}^{\text{fr}}$$

to the domain E , satisfying

$$o = \delta(f).$$

Now we ‘change’ the framed manifold $\mathcal{M}(E)$ (that was already constructed) to $\mathcal{M}(E) \amalg -f(E)$. With respect to these new moduli spaces, the obstruction class vanishes. In other words, for each positive domain D with $\mu(D) = k + 1$, its boundary $\partial\mathcal{M}(D)$ is framed null-cobordant, and therefore, can be filled as a framed manifold, which we define to be the moduli space $\mathcal{M}(D)$. Then the induction proceeds.

In this naïve approach we would define the moduli spaces $\mathcal{M}(\mathbf{x}, \mathbf{y})$ in the flow category by taking the disjoint union for all domains of the given Maslov index connecting the two grid states \mathbf{x} and \mathbf{y} .

The potential (and actual) presence of bubbles, however, prevent us to proceed as outlined above. We need to allow the spaces $\mathcal{M}(D)$ to be more general than framed manifolds with corners; and these stratified spaces will be glued together in an intricate way to form $\mathcal{M}(\mathbf{x}, \mathbf{y})$, which will ultimately be a framed manifold with corners. We will discuss some of the complications caused by the bubbles in the next subsection through some examples; for the full theory the interested reader is advised to turn to [MS21].




FIGURE 31. Two rectangles forming a width-one strip around the torus.

4.7. Bubbles. Recall that for two grid states \mathbf{x} and \mathbf{y} which differ in exactly two coordinates, we associated four rectangles, two from \mathbf{x} to \mathbf{y} and two from \mathbf{y} to \mathbf{x} . Let A be a rectangle from \mathbf{x} to \mathbf{y} and pick a rectangle B from \mathbf{y} to \mathbf{x} . The union of A and B will comprise a positive domain D , which geometrically is a strip running around the torus (the direction depending on which rectangle from \mathbf{y} to \mathbf{x} is chosen). See Figure 31. If $\mu(D) = 2$, then the strips will have height or width equal to 1 (otherwise they will contain entire circles, which carry coordinates, driving the Maslov index up).

The composition D of these domains, seen as a domain from \mathbf{x} to \mathbf{x} , has a moduli space which is an interval, with one end being the broken flowline given by A (from \mathbf{x} to \mathbf{y}) followed by B (from \mathbf{y} to \mathbf{x}), and the other end comprising a bubble, see Figure 32.




FIGURE 32. The other end of this moduli space can be viewed as a boundary degeneration. We follow the convention that horizontal bubbles will be drawn on the left, while vertical bubbles on the right.

Remark 4.6. In the holomorphic picture the other end of the moduli space, the ‘bubble’, corresponds to a geometric picture: it is given by a disk which has its entire boundary only on one of the sets of curves (in the grid language either only on the horizontal or on the vertical ones). This explains the pictorial presentation of Figure 32 on the right.

The contribution from the bubbles does not create any confusion in grid *homology*, as these bubbles come in pairs: a horizontal bubble will be cancelled by a vertical one. To exhibit an explicit pairing, we use the O -markings: we take the strips passing through a fixed O -marking to cancel each other, see Figure 33.

In constructing the flow category, we use these pairs to glue the moduli spaces $\mathcal{M}(A+B)$ and $\mathcal{M}(C+D)$ together, along the portion of their ‘bubble-type’ boundary, as depicted in Figure 34. In this simple example therefore both moduli spaces $\mathcal{M}(A+B)$ and $\mathcal{M}(C+D)$ are intervals, and both admit two different types of boundaries: one corresponding to a broken flowline, and the other to a bubble. The gluing is performed on the bubble-type boundary (where some extra data should be recorded to get a framed manifold with boundary at the end). A crucial feature of the end-result is that the glued-up space no longer has any strata consisting only of bubbles.




FIGURE 33. The two thin strips around the torus, passing through the fixed O -marking. Each can be decomposed into two rectangles, as shown. (The x coordinates are shown as solid dots.)




FIGURE 34. Gluing the two moduli spaces along the two (vertical and horizontal) boundary degenerations.

The situation encountered above is the first instance when bubbles should be taken into account in our constructions. In grid homology we only need Maslov index 1 domains (to define the boundary map in the chain complex) and Maslov index 2 domains (to prove that the square of the boundary map is zero, i.e. we do have a chain complex). In the spacification process, however, domains with higher Maslov indices should be also taken into account — indeed, we need to work with all positive domains connecting grid states. Therefore higher dimensional bubbles should be examined and handled.

Next we describe a Maslov index 3 example. Continuing from the previous example, consider the domains $2A + B$ and $A + C + D$ from x to y . Both the moduli spaces $\mathcal{M}(2A + B)$ and $\mathcal{M}(A + C + D)$ should have bubble type boundaries, but we will glue them carefully so that the glued-up space has no strata consisting only of bubbles. For $\mathcal{M}(2A + B)$, the only completely broken flowline is ABA , and after a little analysis, we see that $\mathcal{M}(2A + B)$ should be a triangle. A similar analysis shows that $\mathcal{M}(A + C + D)$ should be a pentagon. Each of these spaces have one edge on the boundary corresponding to bubbling, and the edge is parametrized by the *height* of the bubble. (We will not discuss the concept of the height of a bubble, or the gluing parameters appearing later — these technical data play crucial role in the gluing theory.) We then glue the two edges together — by matching their heights — and the new moduli space is a smooth manifold with corners, with no stratum consisting only of bubbles. This process is illustrated in Figure 35.




FIGURE 35. The moduli space $\mathcal{M}(2A+B)$ (the triangle) is glued to the moduli space $\mathcal{M}(A+B+C+D)$ (the pentagon) along their bubble-type boundaries. (Note that the glued-up moduli space is a rectangle, not a hexagon.)

Pushing one dimension further, and continuing from the above examples, consider the domains $2A+2B$, $A+B+C+D$, and $2C+2D$, all connecting \mathbf{x} to \mathbf{x} . As it was noted earlier, this time the three moduli spaces $\mathcal{M}(2A+2B)$, $\mathcal{M}(A+B+C+D)$, and $\mathcal{M}(2C+2D)$ will not be manifolds with corners, but will be Whitney stratified spaces. Each will have bubble-type boundaries, but they will be glued together carefully to produce a smooth manifold with corners, with no stratum consisting only of bubbles. We will not describe the full gluing, but rather the gluing near the most complicated portions of their boundaries, namely the portion of the boundary with two bubbles. Each of these three moduli spaces has an interval worth of such ‘double’ bubbling as shown in Figure 36.




FIGURE 36. The three intervals, I_1 , I_2 , I_3 , of double-bubbles in $\mathcal{M}(2A+2B)$, $\mathcal{M}(A+B+C+D)$, $\mathcal{M}(2C+2D)$, respectively. The intervals are parametrized by d , the height difference between the bubbles, and we have $I_1 = I_3 = [0, \infty]$ and $I_2 = [-\infty, \infty]$.




FIGURE 37. The Whitney umbrella stratification of \mathbb{R}^3 , seen as $\text{Sym}^2(\mathbb{C})/\mathbb{R}$. Let $\alpha_1 + i\beta_1$ and $\alpha_2 + i\beta_2$ denote the two complex numbers, which are well-defined up to the \mathbb{R} -action $(\alpha_1, \alpha_2) \mapsto (\alpha_1 + t, \alpha_2 + t)$. There is one 0-dimensional stratum, one 1-dimensional stratum, two 2-dimensional strata, and three 3-dimensional strata, which are shown in different colors and are labelled.

As we got used to gluing horizontal bubbles to vertical bubbles, we might expect that the ‘double horizontal bubbling’ of I_1 be glued to the ‘double vertical bubbling’ of I_3 , and the double ‘mixed’ bubbling of I_2 (one horizontal bubble and one vertical bubble) be glued to itself, hence giving rise to a quotient of I_2 in the gluing. What really happens is, in fact, more complicated, because all the above portions of the moduli spaces corresponding to double bubblings are glued together, as we shall see in Figure 37. Nonetheless, I_2 will be glued to itself, and this is why in what follows we consider the quotient $\mathcal{M}(A+B+C+D)/\sim$, obtained by quotienting $I_2 \subset \mathcal{M}(A+B+C+D)$ by $d \sim -d$.

We now describe how $\mathcal{M}(2A+2B)$, $\mathcal{M}(A+B+C+D)/\sim$, and $\mathcal{M}(2C+2D)$ are glued in a neighborhood of I_1 , I_2/\sim , and I_3 . Each of the intervals I_1 , I_2/\sim , and I_3 is parametrized by $[0, \infty]$, and they are glued together respecting this parametrization. To extend this gluing to a neighborhood, we first need to understand the neighborhoods of these three intervals in $\mathcal{M}(2A+2B)$, $\mathcal{M}(A+B+C+D)/\sim$, and $\mathcal{M}(2C+2D)$.

Towards that end, consider the stratification of \mathbb{R}^3 given by the Whitney umbrella: the ambient space is $\text{Sym}^2(\mathbb{C})/\mathbb{R} \cong \mathbb{R}^3$ — where the \mathbb{R} -action is by translating the real parts — stratified by the signs of the imaginary parts. (There is an extra 0-dimensional stratum when the real parts are equal and both the imaginary parts are 0.) See Figure 37.

Given $(\alpha_1 + i\beta_1, \alpha_2 + i\beta_2) \in \text{Sym}^2(\mathbb{C})$, the two real parts encode the heights of the two bubbles, while the imaginary parts encode the gluing parameters for gluing the two bubbles, and their signs encode whether the bubbles are glued in as horizontal bubbles or as vertical bubbles.

The closure of the red interval (described by the equations $\beta_1 = \beta_2 = 0$) corresponds to double bubbling, and is parametrized by $d = |\alpha_1 - \alpha_2| \in [0, \infty]$. The interval $I_2 \subset \mathcal{M}(A+B+C+D)$ locally looks like $I_2 \subset I_2 \times \mathbb{R}_{\geq 0}^2$, where the $\mathbb{R}_{\geq 0}$ are gluing parameters for gluing the two bubbles. Therefore, after quotienting, $I_2/\sim \subset \mathcal{M}(A+B+C+D)/\sim$ locally looks like the red stratum inside the orange stratum. (The orange stratum has one imaginary part positive and one imaginary part negative, which corresponds to gluing one

bubble in as the horizontal bubble H and one bubble in as the vertical bubble V .) Similarly, $I_1 \subset \mathcal{M}(2A + 2B)$ locally looks like the red stratum inside the light green stratum and $I_3 \subset \mathcal{M}(2C + 2D)$ locally looks like the red stratum inside the light blue stratum. The Whitney umbrella witnesses how the light green, orange, and light blue strata glue together to produce the smooth manifold \mathbb{R}^3 , and that is the local model of how the neighborhoods of $I_1, I_2/\sim, I_3$ are glued together.

REFERENCES

- [Ale28] James W Alexander. Topological invariants of knots and links. *Transactions of the American Mathematical Society*, 30(2):275–306, 1928.
- [BN02] Dror Bar-Natan. On Khovanov’s categorification of the Jones polynomial. *Algebraic & Geometric Topology*, 2(1):337–370, 2002.
- [Car81] Gunnar Carlsson. A counterexample to a conjecture of Steenrod. *Invent. Math.*, 64(1):171–174, 1981.
- [CJS95] Ralph L Cohen, John DS Jones, and Graeme B Segal. Floer’s infinite dimensional Morse theory and homotopy theory. *The Floer memorial volume*, pages 297–325, 1995.
- [Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, 2002.
- [Jon85] Vaughan F. R. Jones. A polynomial invariant for knots via von Neumann algebras. *Bull. Amer. Math. Soc. (N.S.)*, 12(1):103–111, 1985.
- [Kho00] Mikhail Khovanov. A categorification of the Jones polynomial. *Duke Mathematical Journal*, 101(3):359 – 426, 2000.
- [LS14] Robert Lipshitz and Sucharit Sarkar. A Khovanov stable homotopy type. *Journal of the American Mathematical Society*, 27(4):983–1042, 2014.
- [Mil65] John Milnor. *Lectures on the h-cobordism theorem*. Princeton University Press, Princeton, NJ, 1965. Notes by L. Siebenmann and J. Sondow.
- [MOS09] Ciprian Manolescu, Peter Ozsváth, and Sucharit Sarkar. A combinatorial description of knot Floer homology. *Ann. of Math. (2)*, 169(2):633–660, 2009.
- [MOST07] Ciprian Manolescu, Peter Ozsváth, Zoltán Szabó, and Dylan Thurston. On combinatorial link Floer homology. *Geom. Topol.*, 11:2339–2412, 2007.
- [MS21] Ciprian Manolescu and Sucharit Sarkar. A knot Floer stable homotopy type. *arXiv:2108.13566*, 2021.
- [OS04] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and knot invariants. *Adv. Math.*, 186(1):58–116, 2004.
- [OSS15] Peter S. Ozsváth, András I. Stipsicz, and Zoltán Szabó. *Grid homology for knots and links*, volume 208 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.
- [Ras03] Jacob Andrew Rasmussen. *Floer homology and knot complements*. Harvard University, 2003.
- [Rei27] Kurt Reidemeister. Elementare Begründung der Knotentheorie. In *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, volume 5, pages 24–32. Springer, 1927.

HUN-REN ALFRÉD RÉNYI INSTITUTE OF MATHEMATICS, BUDAPEST, HUNGARY
Email address: marengon@renyi.hu

UCLA, LOS ANGELES, CA
Email address: sucharit@math.ucla.edu

HUN-REN ALFRÉD RÉNYI INSTITUTE OF MATHEMATICS, BUDAPEST, HUNGARY
Email address: stipsicz@renyi.hu