

A KHOVANOV STABLE HOMOTOPY TYPE

ROBERT LIPSHITZ AND SUCHARIT SARKAR

CONTENTS

1. Introduction	984
2. Resolution configurations and Khovanov homology	986
3. The Cohen-Jones-Segal construction	991
3.1. $\langle n \rangle$ -manifolds	991
3.2. Flow categories	996
3.3. Framed flow categories to CW complexes	1000
3.4. Covers, subcomplexes, and quotient complexes	1005
4. A framed flow category for the cube	1007
4.1. The cube flow category	1007
4.2. Framing the cube flow category	1009
5. A flow category for Khovanov homology	1013
5.1. Moduli spaces for decorated resolution configurations	1013
5.2. From resolution moduli spaces to the Khovanov flow category	1016
5.3. 0-dimensional moduli spaces	1017
5.4. 1-dimensional moduli spaces	1017
5.5. 2-dimensional moduli spaces	1020
5.6. n -dimensional moduli spaces, $n \geq 3$	1025
6. Invariance of the Khovanov spectrum	1025
7. Skein sequence	1032
8. Reduced version	1032
9. Examples	1033
9.1. An unknot	1034
9.2. The Hopf link	1035
9.3. The Khovanov spectrum of alternating links	1037
10. Some structural conjectures	1039
10.1. Mirrors	1039
10.2. Disjoint unions and connected sums	1039
Acknowledgments	1041
References	1041

Received by the editors August 13, 2012 and, in revised form, May 23, 2013, July 9, 2013, and August 6, 2013.

2010 *Mathematics Subject Classification*. Primary 57M25, 55P42.

The first author was supported by NSF grant number DMS-0905796 and a Sloan Research Fellowship.

The second author was supported by a Clay Mathematics Institute Postdoctoral Fellowship.

1. INTRODUCTION

In [Kho00], Khovanov introduced an elaboration of the Jones polynomial, now generally called Khovanov homology. The Khovanov homology of a link L takes the form of a bigraded abelian group $Kh^{i,j}(L)$, the homology of a bigraded chain complex which we denote $KC^{i,j}(L)$. Khovanov homology relates to the Jones polynomial by taking the graded Euler characteristic:

$$\chi(Kh^{i,j}(L)) = \sum_{i,j} (-1)^i q^j \text{rank } Kh^{i,j}(L) = (q + q^{-1})V(L).$$

In this paper, we give a space-level version of Khovanov homology. That is, we construct a family of (suspension) spectra $\mathcal{X}_{Kh}^j(L)$ so that the Khovanov homology of L is the (reduced) singular cohomology of these spaces. To be precise, we have the following.

Theorem 1.1. *Let L be an oriented link diagram. Let $\mathcal{X}_{Kh}(L) = \bigvee_j \mathcal{X}_{Kh}^j(L)$ be the Khovanov spectrum (Definition 5.5) and let $Kh(L)$ be the Khovanov homology (Definition 2.15). Then the following hold.*

(Θ -1) *The reduced cohomology of $\mathcal{X}_{Kh}^j(L)$ is the Khovanov homology $Kh^{*,j}(L)$:*

$$\tilde{H}^i(\mathcal{X}_{Kh}^j(L)) = Kh^{i,j}(L).$$

(Θ -2) *The stable homotopy type of $\mathcal{X}_{Kh}^j(L)$ is an invariant of the isotopy class of the link corresponding to L (and is independent of all the other choices in its construction).*

In particular, the Euler characteristic of our space $\mathcal{X}_{Kh}^j(L)$ is the coefficient of q^i in the Jones polynomial $(q + q^{-1})V(L)$. (There is also a reduced version of both theories, which eliminates the $q + q^{-1}$, but, as with Khovanov homology itself, the relationship between the reduced and un-reduced homotopy types is not entirely trivial.)

Our work is inspired by a question posed by Cohen-Jones-Segal 15 years ago [CJS95]: Is Floer homology the ordinary (singular) homology of some naturally associated space (or rather, spectrum)? Since then, such a space has been constructed in several interesting cases [Man03, Man07, KM, Kra, Coh10], but the question remains open in general.

In fact, Cohen-Jones-Segal went further: they proposed a construction of such a spectrum using the higher-dimensional Floer moduli spaces. These moduli spaces are encoded in a structure they call a *framed flow category*, and to any framed flow category they showed how to associate a CW complex. A framed flow category is a category \mathcal{C} , together with a \mathbb{Z} -valued grading function $\text{gr}: \text{Ob}(\mathcal{C}) \rightarrow \mathbb{Z}$ such that for $x, y \in \text{Ob}(\mathcal{C})$, $\text{Hom}(x, y)$ is a compact $(\text{gr}(x) - \text{gr}(y) - 1)$ -dimensional manifold with corners, satisfying certain compatibility conditions with respect to the composition maps and corner structure (see Definitions 3.13 and 3.21). The realization $|\mathcal{C}|$ of a framed flow category \mathcal{C} is not the usual geometric realization of a (topological) category, but rather a generalization of the Dold-Thom construction. For example, in the special case that $\mathcal{C} = \{x_1, \dots, x_k, y_1, \dots, y_\ell\}$ with $\text{gr}(x_i) = n + 1$, $\text{gr}(y_j) = n$, the realization $|\mathcal{C}|$ is obtained from a wedge sum of ℓ n -spheres, one for each y_j , by attaching k $(n + 1)$ -disks, one for each x_i . The attaching map $S^n = \partial \mathbb{D}_{x_i}^{n+1} \rightarrow \bigvee_{j=1}^\ell S_{y_j}^n \rightarrow S_{y_j}^n$ is determined by an integer, the degree; this degree is the signed number of points in the framed 0-manifold $\text{Hom}(x_i, y_j)$. One

way to construct the attaching map is to embed $\coprod_j \text{Hom}(x_i, y_j)$ into $\partial\mathbb{D}^{n+1}$, use the framing to identify a neighborhood of $\text{Hom}(x_i, y_j)$ with $\text{Hom}(x_i, y_j) \times \mathbb{D}^n$, project $\text{Hom}(x_i, y_j) \times \mathbb{D}^n \rightarrow S_{y_j}^n$, and collapse the complement of these neighborhoods to the basepoint. The general case of the Cohen–Jones–Segal construction follows along these lines; see Definition 3.24.

So, to construct a Khovanov stable homotopy type, it suffices to construct a flow category for Khovanov homology. Our Khovanov flow category has one object for each generator of the standard Khovanov complex (which is reviewed in Section 2), and the grading is the homological grading on Khovanov homology. Let δ denote the differential on the Khovanov complex. If $\text{gr}(x) - \text{gr}(y) = 1$ and x occurs in $\delta(y)$ with coefficient n_{xy} , then we take $\text{Hom}(x, y)$ to be a disjoint union of $|n_{xy}|$ points, all framed positively or negatively according to the sign of n_{xy} . The spaces $\text{Hom}(x, y)$ with $\text{gr}(x) - \text{gr}(y) > 1$ are constructed inductively. The boundary of $\text{Hom}(x, y)$ is already determined by lower-dimensional Hom spaces, and it turns out that $\partial\text{Hom}(x, y)$ is a disjoint union of spheres. We choose $\text{Hom}(x, y)$ to be a disjoint union of disks. (Keeping track of the corner structure, these disks are, in fact, permutohedra; see Example 3.2 and Lemma 4.3.) If $\text{gr}(x) - \text{gr}(y) > 2$ then there is a unique way to construct such disks starting with their boundaries; if $\text{gr}(x) - \text{gr}(y) = 2$ then the construction depends on a choice, which we call the *ladybug matching* (see Section 5.4).

To prove inductively that it is possible to choose the Hom spaces to be disks, that is, their expected boundaries (unions of products of lower dimensional Hom spaces) are spheres, and to specify the framings of these disks, we use an auxiliary flow category, called the *cube flow category* (Definition 4.1). The cube flow category has one object for each vertex of the cube $\{0, 1\}^N$ and its Hom spaces are permutohedra. During our inductive construction of the Hom spaces for the Khovanov flow category, we impose an additional requirement that these spaces should admit covering maps to the corresponding Hom spaces of the cube flow category. For dimensions up to 2, we check explicitly that it is possible to construct Hom spaces with such covering maps (Sections 5.3–5.5). For the n -dimensional Hom spaces with $n > 2$ (Section 5.6), by the inductive construction, their expected boundaries already admit covering maps to the boundary of the n -dimensional permutohedron, which, being S^{n-1} , is simply connected; therefore, their expected boundaries are disjoint unions of spheres and the covering map is trivial. Then there is a unique way to define the n -dimensional Hom spaces to be disks, and they admit covering maps to the n -dimensional permutohedron. The use of the cube flow category in this construction may be viewed as a space-level analog of the fact that the Khovanov complex lies over the cube $\{0, 1\}^N$.

In particular, our construction gives a CW complex whose cells are in one-to-one correspondence with generators of the standard Khovanov chain complex. Furthermore, it also follows that the Khovanov flow category decomposes as a disjoint union of subcategories, one for each quantum grading, and so its realization decomposes as a wedge sum of spaces, one for each quantum grading. Like the construction of Khovanov homology, our construction of the Khovanov flow category is entirely combinatorial; although we use some language from Morse theory, no Morse theory (not to mention Floer theory) is required.

The cohomology of a spectrum does not carry a cup product, but it does carry stable cohomology operations. In the sequel [LSb] we give an explicit computation

of the operation Sq^2 on Khovanov homology coming from our stable type. This operation is nontrivial for many simple knots, such as the torus knot $T_{3,4}$. This implies, in particular, that our Khovanov stable type is not simply a wedge sum of Moore spaces. Indeed, Seed has extended the results from [LSb] to find pairs of links with isomorphic Khovanov homology but distinct Khovanov homotopy types [See]. So, $\mathcal{X}_{Kh}(L)$ is a strictly stronger invariant than $Kh(L)$.

Rasmussen constructed a slice genus bound, called the s -invariant, using Khovanov homology [Ras10]. Using the Khovanov homotopy type, we produce a family of generalizations of the s -invariant [LSa]; each of them is a slice genus bound, and we show that at least one of them is a stronger bound.

A different kind of Khovanov stable homotopy type has been constructed by Everitt-Turner [ET]. Their construction gives spaces whose homotopy groups (rather than homology groups) are Khovanov homology. It is shown in [ELST] that the spaces constructed in [ET] are products of Eilenberg-MacLane spaces; in particular, they are determined by the Khovanov homology. Since the first version of this paper, Hu-Kriz-Kriz have given another construction of a Khovanov stable homotopy type (in our sense) [HKK]. Although their construction uses quite different techniques, it seems natural to conjecture that the spaces constructed in [HKK] are stably homotopy equivalent to the spaces constructed here. (In particular, their construction also uses the ladybug matching from Section 5.4 as a key choice.)

This paper is organized as follows. The first two sections are background: Section 2 introduces notation related to the Khovanov chain complex, that will be used when constructing the Khovanov flow category, and Section 3 discusses flow categories themselves and how one produces a CW complex from a flow category. Section 4 introduces a simple flow category that will be a kind of local model for the Khovanov flow category. The Khovanov flow category itself is constructed in Section 5. Invariance of the Khovanov homotopy type is proved in Section 6. Like Khovanov homology, the Khovanov homotopy types trivially satisfy an unoriented skein triangle; this is explained in Section 7. The spaces for reduced Khovanov homology are introduced in Section 8. We describe with some examples in Section 9, including the computation of $\mathcal{X}_{Kh}(L)$ for any alternating link L , in terms of the Khovanov homology $Kh(L)$. We conclude in Section 10 with some speculations.

2. RESOLUTION CONFIGURATIONS AND KHOVANOV HOMOLOGY

The original definition of Khovanov homology [Kho00] is in terms of Kauffman states, and transitions between consecutive Kauffman states (with respect to the obvious partial order). Our Khovanov homotopy type will be defined in terms of more complicated families of Kauffman states. This section is devoted to developing terminology to discuss these sequences of Kauffman states. In Definition 2.1, we introduce *resolution configurations*; essentially, a resolution configuration is a Kauffman state together with a marking of the crossings at which one took the 0-resolution. (Resolution configurations have appeared informally elsewhere in the literature—for instance, in [ORS].) We then introduce several operations on resolution configurations, a notion of a labeled resolution configuration, and a partial order on labeled resolution configurations. Most of these concepts are illustrated in Figures 2.1 and 2.2. We conclude the section by restating the definition of Khovanov homology in this language.

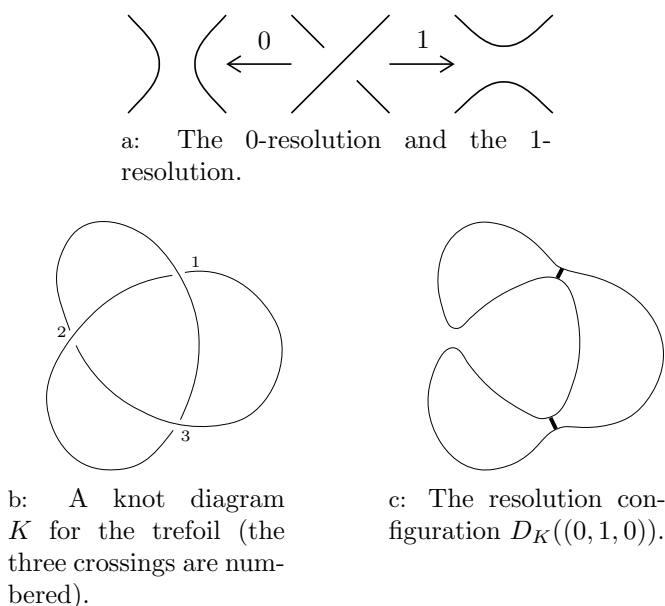


FIGURE 2.1. Resolution configurations from a knot diagram.

Definition 2.1. A resolution configuration D is a pair $(Z(D), A(D))$, where $Z(D)$ is a set of pairwise-disjoint embedded circles in S^2 , and $A(D)$ is a totally ordered collection of disjoint arcs embedded in S^2 , with $A(D) \cap Z(D) = \partial A(D)$.

We call the number of arcs in $A(D)$ the *index* of the resolution configuration D , and denote it by $\text{ind}(D)$.

We sometimes abuse notation and write $Z(D)$ to mean $\bigcup_{Z \in Z(D)} Z$ and $A(D)$ to mean $\bigcup_{A \in A(D)} A$.

Occasionally, we will describe the total order on $A(D)$ by numbering the arcs: a lower numbered arc precedes a higher numbered one.

Definition 2.2. Given a link diagram L with n crossings, an ordering of the crossings in L , and a vector $v \in \{0, 1\}^n$, there is an associated resolution configuration $D_L(v)$ gotten by taking the resolution of L corresponding to v (i.e., taking the 0-resolution at the i^{th} crossing if $v_i = 0$, and the 1-resolution otherwise) and then placing arcs corresponding to each of the crossings labeled by 0's in v (i.e., at the i^{th} crossing if $v_i = 0$); see Figure 2.1.

Therefore, $n - \text{ind}(D_L(v))$ equals $|v| = \sum v_i$, the (Manhattan) norm of v .

Definition 2.3. Given resolution configurations D and E there is a new resolution configuration $D \setminus E$ defined by

$$Z(D \setminus E) = Z(D) \setminus Z(E) \quad A(D \setminus E) = \{A \in A(D) \mid \forall Z \in Z(E) : \partial A \cap Z = \emptyset\}.$$

Let $D \cap E = D \setminus (D \setminus E)$. See Figure 2.2.

Note that $Z(D \cap E) = Z(E \cap D)$ and $A(D \cap E) = A(E \cap D)$; however, the total orders on $A(D \cap E)$ and $A(E \cap D)$ could be different.

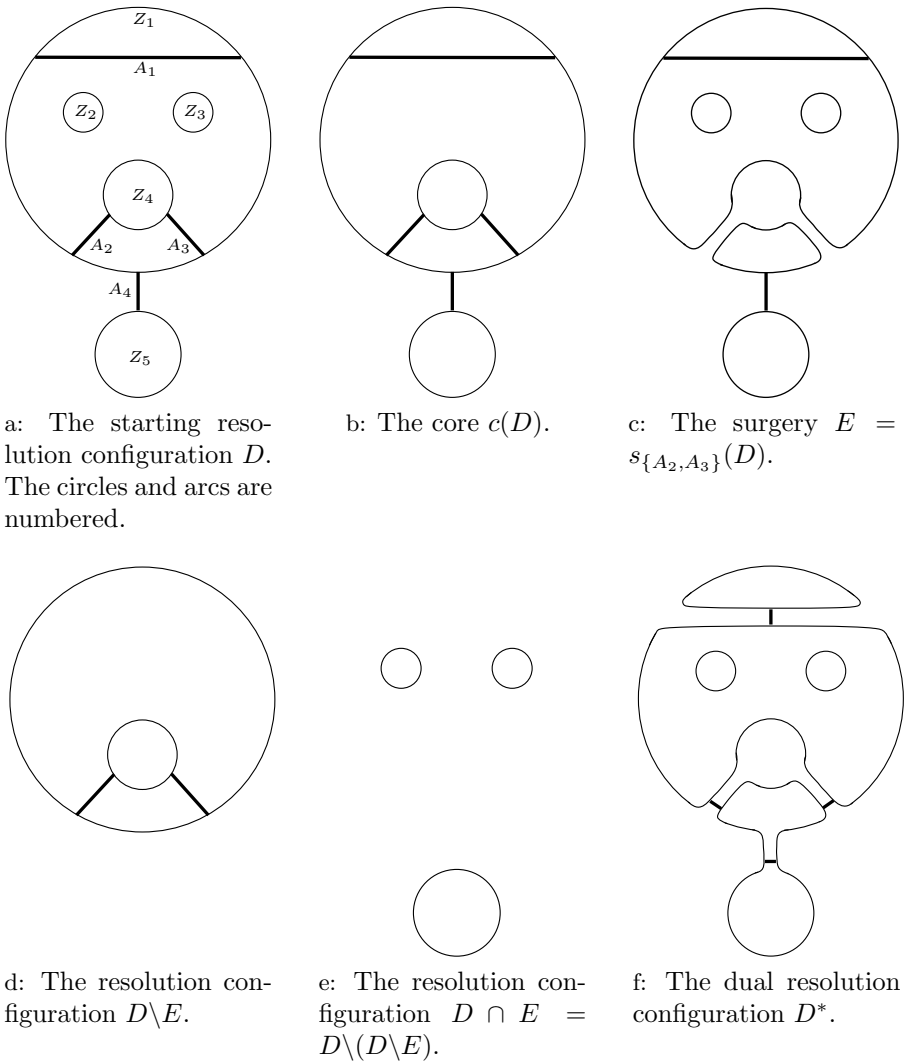


FIGURE 2.2. An illustration of some resolution configuration notations.

Definition 2.4. The *core* $c(D)$ of a resolution configuration D is the resolution configuration obtained from D by deleting all the circles in $Z(D)$ that are disjoint from all the arcs in $A(D)$; see Figure 2.2b.

A resolution configuration D is called *basic* if $D = c(D)$, i.e., if every circle in $Z(D)$ intersects an arc in $A(D)$.

Definition 2.5. Given a resolution configuration D and a subset $A' \subseteq A(D)$ there is a new resolution configuration $s_{A'}(D)$, the *surgery of D along A'* , obtained as follows. The circles $Z(s_{A'}(D))$ of $s_{A'}(D)$ are obtained by performing embedded surgery along the arcs in A' ; in other words, $Z(s_{A'}(D))$ is obtained by deleting a neighborhood of $\partial A'$ from $Z(D)$ and then connecting the endpoints of the result

using parallel translates of A' . The arcs of $s_{A'}(D)$ are the arcs of D not in A' , i.e., $A(s_{A'}(D)) = A(D) \setminus A'$; see Figure 2.2c.

Let $s(D) = s_{A(D)}(D)$ denote the maximal surgery on D .

Lemma 2.6. *If a resolution configuration E is obtained from a resolution configuration D by a surgery, then $D \setminus E$ is a basic resolution configuration, $E \setminus D = s(D \setminus E)$ and $D \cap E = E \cap D$.*

Proof. This is immediate from the definitions. □

Definition 2.7. Associated to a resolution configuration D is a *dual resolution configuration* D^* , defined as follows. The circles $Z(D^*)$ are obtained from $Z(D)$ by performing embedded surgery according to the arcs $A(D)$, i.e., $Z(D^*) = Z(s(D))$. The arcs $A(D^*)$ are duals to the arcs $A(D)$. That is, write $A(D) = A_1 \cup \dots \cup A_n$. For each i , choose an arc A_i^* in a neighborhood of A_i , with boundary on $Z(D^*)$, and intersecting A_i once. Then $A(D^*) = A_n^* \cup \dots \cup A_1^*$; see Figure 2.2f.

There is a forgetful map from resolution configurations to graphs, by collapsing the circles.

Definition 2.8. A resolution configuration D specifies a graph $G(D)$ with one vertex for each element of $Z(D)$ and an edge for each element of $A(D)$. Given an element $Z \in Z(D)$ (resp. $A \in A(D)$) we will write $G(Z)$ (resp. $G(A)$) for the corresponding vertex (resp. edge) of $G(D)$.

A *leaf* of a resolution configuration D is a circle $Z \in Z(D)$ so that $G(Z)$ is a leaf of $G(D)$. A *co-leaf* of D is an arc $A \in A(D)$ so that one endpoint of the dual arc $A^* \in A(D^*)$ is a leaf of the dual configuration D^* . For example, in Figure 2.2a, the circle Z_5 is a leaf of D , while the arc A_1 is a co-leaf of D .

Definition 2.9. A *labeled resolution configuration* is a pair (D, x) of a resolution configuration D and a labeling x of each element of $Z(D)$ by either x_+ or x_- .

Definition 2.10. There is a partial order \prec on labeled resolution configurations defined as follows. We declare that $(E, y) \prec (D, x)$ if:

- (1) The labelings x and y induce the same labeling on $D \cap E = E \cap D$.
- (2) D is obtained from E by surgering along a single arc of $A(E)$. In particular, either:
 - (a) $Z(E \setminus D)$ contains exactly one circle, say Z_i , and $Z(D \setminus E)$ contains exactly two circles, say Z_j and Z_k , or
 - (b) $Z(E \setminus D)$ contains exactly two circles, say Z_i and Z_j , and $Z(D \setminus E)$ contains exactly one circle, say Z_k .
- (3) In Case (2a), either $y(Z_i) = x(Z_j) = x(Z_k) = x_-$ or $y(Z_i) = x_+$ and $\{x(Z_j), x(Z_k)\} = \{x_+, x_-\}$.
 In Case (2b), either $y(Z_i) = y(Z_j) = x(Z_k) = x_+$ or $\{y(Z_i), y(Z_j)\} = \{x_-, x_+\}$ and $x(Z_k) = x_-$.

Now, \prec is defined to be the transitive closure of this relation.

In short, the three conditions of Definition 2.10 say that (D, x) is obtained from (E, y) by applying one of the edge maps in the Khovanov cube; cf. Definition 2.15.

Definition 2.11. A *decorated resolution configuration* is a triple (D, x, y) where D is a resolution configuration and x (resp. y) is a labeling of each component of $Z(s(D))$ (resp. $Z(D)$) by an element of $\{x_+, x_-\}$, such that: $(D, y) \preceq (s(D), x)$.

Associated to a decorated resolution configuration (D, x, y) is the poset $P(D, x, y)$ consisting of all labeled resolution configurations (E, z) with $(D, y) \preceq (E, z) \preceq (s(D), x)$.

Definition 2.12. The *dual* of a decorated resolution configuration (D, x, y) is the decorated resolution configuration (D^*, y^*, x^*) , where D^* is the dual of D , the labelings x and x^* are dual, i.e., they disagree on every circle in $Z(D^*) = Z(s(D))$, and the labelings y and y^* are dual, i.e., they disagree on every circle in $Z(D) = Z(s(D^*))$.

Lemma 2.13. *For any decorated resolution configuration (D, x, y) , the partially ordered set $P(D^*, y^*, x^*)$ is the reverse of the partially ordered set $P(D, x, y)$.*

Proof. When $\text{ind}(D) = 1$, this is immediate from Definition 2.10. The general case reduces to the index 1 case since \prec is defined as a transitive closure. \square

The following is a technical lemma that will be useful to us for a later case analysis.

Lemma 2.14. *Let (D, x, y) be a decorated resolution diagram. Let $Z_1 \in Z(D)$ be a leaf of D , and let $A_1 \in A(D)$ be the arc one of whose endpoints lies on Z_1 . Let $Z_2 \in Z(D)$ be the circle containing the other endpoint of A_1 , $Z_1^* \in Z(s(D))$ be the circle which contains both the endpoints of A_1^* , and $Z_2^* \in Z(s_{A(D) \setminus \{A_1\}}(D)) \setminus \{Z_1\}$ be the unique circle that contains an endpoint of A_1 .*

Let D' be the resolution configuration obtained from D by deleting Z_1 and A_1 . Let (D', x', y') be the following decorated resolution configuration: the labeling y' on $Z(D')$ is induced from the labeling y on $Z(D)$. The labeling x' on $Z(s(D')) \setminus \{Z_2^\}$ is induced from the labeling x on $Z(s(D)) \setminus \{Z_1^*\}$. If $y(Z_1) = x_+$, set $x'(Z_2^*) = x(Z_1^*)$; otherwise, set $x'(Z_2^*) = x_+$.*

Then $P(D, x, y) = P(D', x', y') \times \{0, 1\}$, where $\{0, 1\}$ is the two-element poset with $0 \prec 1$.

Proof. We will produce a map from $P(D, x, y)$ to $P(D', x', y') \times \{0, 1\}$ that will induce an isomorphism of posets. Let $(E, z) \in P(D, x, y)$ with $E = s_A(D)$ for some $A \subseteq A(D)$. We will produce $((E', z'), i) \in P(D', x', y') \times \{0, 1\}$.

There are two cases.

- (1) $A_1 \notin A$. Set $i = 0$ and $E' = s_A(D')$. The labeling z' on $Z(s_A(D'))$ is induced from the labeling z on $Z(s_A(D)) = Z(s_A(D')) \amalg \{Z_1\}$.
- (2) $A_1 \in A$. Set $i = 1$ and $E' = s_{A \setminus \{A_1\}}(D')$. Let $\{Z_A\} = Z(s_A(D)) \setminus Z(s_{A \setminus \{A_1\}}(D'))$ and let $\{Z'_A\} = Z(s_{A \setminus \{A_1\}}(D')) \setminus Z(s_A(D))$. The labeling z' on $Z(s_{A \setminus \{A_1\}}(D')) \cap Z(s_A(D))$ is induced from z . It only remains to define $z'(Z'_A)$. If $y(Z_1) = x_+$, set $z'(Z'_A) = z(Z_A)$; otherwise, set $z'(Z'_A) = x_+$.

It is a straightforward, albeit elaborate, check that this map $P(D, x, y) \rightarrow P(D', x', y') \times \{0, 1\}$ induces an isomorphism of posets. \square

We conclude this section by giving the definition of the Khovanov chain complex from [Kho00] in the language of resolution configurations and the partial order \prec .

Definition 2.15. Given an oriented link diagram L with n crossings and an ordering of the crossings in L , the *Khovanov chain complex* is defined as follows.

The chain group $KC(L)$ is the \mathbb{Z} -module freely generated by labeled resolution configurations of the form $(D_L(u), x)$ for $u \in \{0, 1\}^n$. The chain group $KC(L)$

carries two gradings, a *homological grading* gr_h and a *quantum grading* gr_q , defined as follows:

$$\begin{aligned} \text{gr}_h((D_L(u), x)) &= -n_- + |u|, \\ \text{gr}_q((D_L(u), x)) &= n_+ - 2n_- + |u| + \#\{Z \in Z(D_L(u)) \mid x(Z) = x_+\} \\ &\quad - \#\{Z \in Z(D_L(u)) \mid x(Z) = x_-\}. \end{aligned}$$

Here n_+ denotes the number of positive crossings in L ; and $n_- = n - n_+$ denotes the number of negative crossings.

The differential preserves the quantum grading, increases the homological grading by 1, and is defined as

$$\delta(D_L(v), y) = \sum_{\substack{(D_L(u), x) \\ |u|=|v|+1 \\ (D_L(v), y) \prec (D_L(u), x)}} (-1)^{s_0(\mathcal{C}_{u,v})} (D_L(u), x),$$

where $s_0(\mathcal{C}_{u,v}) \in \mathbb{F}_2$ is defined as follows: if $u = (\epsilon_1, \dots, \epsilon_{i-1}, 1, \epsilon_{i+1}, \dots, \epsilon_n)$ and $v = (\epsilon_1, \dots, \epsilon_{i-1}, 0, \epsilon_{i+1}, \dots, \epsilon_n)$, then $s_0(\mathcal{C}_{u,v}) = (\epsilon_1 + \dots + \epsilon_{i-1})$; see also Definition 4.5.

3. THE COHEN-JONES-SEGAL CONSTRUCTION

In Morse theory, under suitable hypotheses the moduli spaces of flows admit compactifications as manifolds with corners, where the codimension- i boundary consists of i -times broken flows; in Section 5 we will construct a collection of spaces refining the Khovanov complex, satisfying the same property. In [CJS95], Cohen-Jones-Segal proposed a convenient language for encoding this information, which they called a flow category. They went on to explain how one can build a spectrum from a flow category; in the Morse theory case, their construction recovers the suspension spectrum of the base manifold.

This section is devoted to reviewing the Cohen-Jones-Segal construction, and the basic definitions underlying it. In Section 3.1 we review basic notions about manifolds with corners, and in Section 3.2 we recall the definition of a (framed) flow category, and how these arise from Morse theory. Section 3.3 gives a reformulation of the Cohen-Jones-Segal procedure for building a space from a flow category.

The flow category for Khovanov homology is built as a cover of a simpler flow category; this notion of a cover is introduced in Section 3.4. The proof of invariance of the Khovanov stable homotopy type will make extensive use of subcomplexes and quotient complexes, basic results on which are also given in Section 3.4.

3.1. $\langle n \rangle$ -manifolds. We start with some preliminaries on $\langle n \rangle$ -manifolds, following [Lau00] and [Jän68].

Definition 3.1. Let $\mathbb{R}_+ = [0, \infty)$.

A k -dimensional manifold with corners is a (possibly empty) topological space X along with a maximal atlas, where an atlas is a collection of charts (U, ϕ) , where U is an open subset of X and ϕ is a homeomorphism from U to an open subset of $(\mathbb{R}_+)^k$, such that

- $X = \bigcup U$ and
- given charts (U, ϕ) and (V, ψ) , the map $\phi \circ \psi^{-1}$ is a C^∞ diffeomorphism from $\psi(U \cap V)$ to $\phi(U \cap V)$.

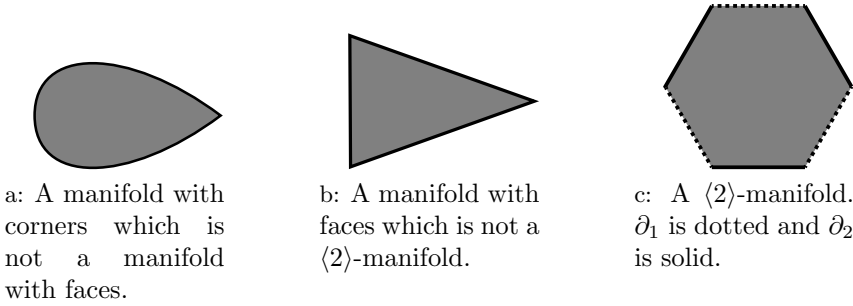


FIGURE 3.1. Manifolds with corners, manifolds with faces, and $\langle n \rangle$ -manifolds.

For $x \in X$, let $c(x)$ denote the number of coordinates in $\phi(x)$ which are 0, for any chart (U, ϕ) with $U \ni x$. The *codimension- i boundary* of X is the subspace $\{x \in X \mid c(x) = i\}$; the usual boundary ∂X is the closure of the codimension-1 boundary. A *connected face* is the closure of a connected component of the codimension-1 boundary of X . A *face* of X is a (possibly empty) union of disjoint connected faces of X .

A *k -dimensional manifold with faces* is a k -dimensional manifold with corners X such that every $x \in X$ belongs to exactly $c(x)$ connected faces of X . It is easy to see that any face of a manifold with faces is also a manifold with faces. A *k -dimensional $\langle n \rangle$ -manifold* is a k -dimensional manifold with faces X along with an ordered n -tuple $(\partial_1 X, \dots, \partial_n X)$ of faces of X such that

- $\bigcup_i \partial_i X = \partial X$ and
- for all distinct i, j , $\partial_i X \cap \partial_j X$ is a face of both $\partial_i X$ and $\partial_j X$;

see Figure 3.1.

Example 3.2. The *permutohedron* P_{n+1} is the following n -dimensional $\langle n \rangle$ -manifold. As a polytope, it is the convex hull in \mathbb{R}^{n+1} of the $(n+1)!$ points $(\sigma(1), \dots, \sigma(n+1))$, where σ varies over the permutations of $\{1, \dots, n+1\}$. (See [Zie95, Example 0.10] for more details.) The faces of P_{n+1} are in bijection with subsets $\emptyset \neq S \subset \{1, \dots, n+1\}$: the face corresponding to such a subset S is the intersection of P_{n+1} and plane $\sum_{i \in S} x_i = \frac{k(k+1)}{2}$ in \mathbb{R}^{n+1} , where k is the cardinality of S . Define $\partial_i P_{n+1}$ to be the union of faces corresponding to subsets of cardinality i .

Definition 3.3. A smooth map f from an $\langle n \rangle$ -manifold X to another $\langle n \rangle$ -manifold Y is said to an *n -map* if $f^{-1}(\partial_i Y) = \partial_i X$ for all i .

Definition 3.4. Let $\underline{2}^n$ be the category which has one object for each element of $\{0, 1\}^n$, and a unique morphism from b to a if $b \leq a$ in the obvious partial order on $\{0, 1\}^n$. For $a \in \{0, 1\}^n$ let $|a|$ denote the sum of the entries in a ; we call this the *weight* of a . Let $\bar{0}$ denote $(0, 0, \dots, 0)$ and let $\bar{1}$ denote $(1, 1, \dots, 1)$.

An *n -diagram* is a functor from $\underline{2}^n$ to the category of topological spaces.

An $\langle n \rangle$ -manifold X can be treated as n -diagram in the following way. For $a = (a_1, \dots, a_n) \in \{0, 1\}^n$, define

$$X(a) = \begin{cases} X & \text{if } a = \bar{1} \\ \bigcap_{i \in \{i \mid a_i = 0\}} \partial_i X & \text{otherwise.} \end{cases}$$

For $b \leq a$ in $\{0, 1\}^n$, the map $X(b) \rightarrow X(a)$ is the inclusion map. It is easy to see that $X(a)$ is an $\langle |a| \rangle$ -manifold; in fact, the $|a|$ -diagram corresponding to $X(a)$ is the functor restricted to the full subcategory of $\underline{2}^n$ corresponding to the vertices less than or equal to a .

By an abuse of terminology, henceforth whenever we say $\langle n \rangle$ -manifold, we will either mean a topological space or an n -diagram, depending on the context.

In this language, products are easy to define: Given $\langle n_i \rangle$ -manifolds X_i for $i \in \{1, 2\}$, the product $X_1 \times X_2$ is an $\langle n_1 + n_2 \rangle$ -manifold defined via the isomorphism $\underline{2}^{n_1+n_2} = \underline{2}^{n_1} \times \underline{2}^{n_2}$.

Since the category $\underline{2}^n$ has a unique maximal vertex $\bar{1}$, and for Y an $\langle n \rangle$ -manifold the maps $Y(a) \rightarrow Y(\bar{1})$ are injective, the manifold $Y = Y(\bar{1})$ is the direct limit (colimit) of Y , thought of as an n -diagram. This observation allows us to talk about the boundary of an $\langle n \rangle$ -manifold Y without explicitly mentioning $Y(\bar{1})$.

Definition 3.5. Let $\underline{2}^n \setminus \bar{1}$ be the full subcategory of $\underline{2}^n$ generated by all the elements except $\bar{1}$. A *truncated n -diagram* is a functor from $\underline{2}^n \setminus \bar{1}$ to the category of topological spaces. A *k -dimensional $\langle n \rangle$ -boundary* is a truncated n -diagram such that the restriction of the functor to each maximal dimensional face of $\underline{2}^n \setminus \bar{1}$ is a k -dimensional $\langle n - 1 \rangle$ -manifold.

Given an $\langle n \rangle$ -boundary Y , let $\text{colim } Y$ denote the pushout of Y . Note that for any vertex $a \in \{0, 1\}^n$, except $\bar{1}$, there is a continuous map

$$\phi_a : Y(a) \rightarrow \text{colim } Y.$$

Therefore, the restriction of a k -dimensional $\langle n \rangle$ -manifold X to $\underline{2}^n \setminus \bar{1}$ produces a $(k - 1)$ -dimensional $\langle n \rangle$ -boundary Y , with $\text{colim } Y = \partial X$.

Lemma 3.6. *Let Y be an $\langle n \rangle$ -boundary, and a any vertex of $\underline{2}^n \setminus \bar{1}$. Then $\phi_a : Y(a) \rightarrow \text{colim } Y$ is injective.*

Proof. This is immediate from the definitions. □

By Lemma 3.6, it makes sense to say that a point $p \in \text{colim } Y$ lies in $Y(a)$ if p is in the image of ϕ_a .

Lemma 3.7. *Given vertices a, b of $\underline{2}^n \setminus \bar{1}$, let c be the greatest lower bound of a and b . (That is, $c \leq a$ and $c \leq b$, and c is maximal with respect to this property; there is a unique such c .) Then $Y(a) \cap Y(b) = Y(c)$.*

Proof. This is immediate from the definition of the colimit (and the fact that, for the cube, any pair of vertices has a unique greatest lower bound). □

Lemma 3.8. *Let Y be an $\langle n \rangle$ -boundary. The following conditions on a point $p \in \text{colim } Y$ are equivalent:*

- (1) *There are exactly k weight $(n - 1)$ vertices $a_1, \dots, a_k \in \{0, 1\}^n$ so that $p \in Y(a_i)$, $i = 1, \dots, k$.*
- (2) *There is a unique weight $(n - k)$ vertex b so that $p \in Y(b)$.*
- (3) *For each weight $(n - 1)$ vertex a so that $p \in Y(a)$, p lies in the codimension- $(k - 1)$ boundary of $Y(a)$.*

If any of these three equivalent conditions hold for p then we say that p has depth k .

Proof. (1) \implies (2): This follows from Lemma 3.7 by induction; the vertex b is the greatest lower bound of a_1, \dots, a_k .

(2) \implies (1): Again, this is immediate from Lemma 3.7: if a_i is a weight $(n - 1)$ vertex so that $b \leq a_i$ then $p \in Y(a_i)$. There are k such vertices. If $p \in Y(a)$ for any other weight $(n - 1)$ vertex a then p would lie in $Y(c)$ for some weight $(n - k - 1)$ vertex c (by the previous step), and hence in $Y(b')$ for another weight $(n - k)$ vertex b' .

(3) \implies (2): Since the face $\{c \mid c \leq a\}$ is an $\langle n - 1 \rangle$ -manifold and p is in the codimension- $(k - 1)$ boundary of $Y(a)$, there must be a weight $(n - k)$ vertex b so that $p \in Y(b)$. Moreover, b is unique, since if b' is another such vertex and c is the greatest lower bound of b and b' then $p \in Y(c)$ (by Lemma 3.7), and the weight of c is less than $(n - k)$, so p lies in the codimension $< (k - 1)$ boundary of $Y(a)$.

(2) \implies (3): By Lemma 3.7, $b \leq a$. Since the face $\{c \mid c \leq a\}$ of $\mathbb{Z}^n \setminus \bar{1}$ is an $\langle n - 1 \rangle$ -manifold, and $p \in Y(b)$, it follows that p lies in the codimension $\geq (k - 1)$ boundary of $Y(a)$. If p lies in the codimension $\geq k$ boundary of $Y(a)$, then $p \in Y(c)$ for some vertex c of weight less than $(n - k)$ (by the previous step), and hence $p \in Y(b')$ for another weight $(n - k)$ vertex b' . \square

Proposition 3.9. *Let Y and X be n -boundaries so that $Y(a)$ and $X(a)$ are compact for each object a of $\mathbb{Z}^n \setminus \bar{1}$. Let $F: Y \rightarrow X$ be a natural transformation such that the restriction of F to each face of Y is an $(n - 1)$ -map. If for every $a \in \text{Ob}_{\mathbb{Z}^n \setminus \bar{1}}$ the map $F(a): Y(a) \rightarrow X(a)$ is a covering map then*

$$G = \text{colim}(F): \text{colim}(Y) \rightarrow \text{colim}(X)$$

is a covering map.

Proof. Fix a point $p \in \text{colim}(X)$ and a point $q \in G^{-1}(p) \subset \text{colim}(Y)$. Let k be the depth of p .

We start by showing that q has depth k . If $a \in \{0, 1\}^n$ is a weight $(n - 1)$ vertex so that $q \in Y(a)$ then $p = G(q) \in X(a)$. By Part (3) of Lemma 3.8, p lies in the codimension- $(k - 1)$ boundary of $X(a)$. Since the restriction of F to the face containing a is an $(n - 1)$ -map, q lies in the codimension- $(k - 1)$ boundary of $Y(a)$, so the depth of q is k .

Now, let a_1, \dots, a_k be the weight $(n - 1)$ vertices so that $q \in Y(a_i)$ (and $x \in X(a_i)$) for $i = 1, \dots, k$. Let $a_{i,j}$ be the unique weight $(n - 2)$ vertex so that $a_{i,j} \leq a_i$ and $a_{i,j} \leq a_j$. Choose a small open neighborhood U_i of p in $X(a_i)$ so that $U_i \cap X(a_{i,j}) = U_j \cap X(a_{i,j})$ for each pair i, j . Then $U = U_1 \cup \dots \cup U_k$ is an open neighborhood of p in $\text{colim} X$. Let $V_i = F^{-1}(U_i) \subset Y(a_i)$. Let V be the connected component of $V_1 \cup \dots \cup V_k$ containing q ; then V is an open neighborhood of q in $\text{colim} Y$, and $G|_V: V \rightarrow U$ is a homeomorphism. \square

Definition 3.10. Given an $(n + 1)$ -tuple $\mathbf{d} = (d_0, \dots, d_n) \in \mathbb{N}^{n+1}$, let

$$\mathbb{E}_n^{\mathbf{d}} = \mathbb{R}^{d_0} \times \mathbb{R}_+ \times \mathbb{R}^{d_1} \times \mathbb{R}_+ \times \dots \times \mathbb{R}_+ \times \mathbb{R}^{d_n}.$$

Treat $\mathbb{E}_n^{\mathbf{d}}$ as an $\langle n \rangle$ -manifold by defining

$$\partial_i(\mathbb{E}_n^{\mathbf{d}}) = \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_{i-1}} \times \{0\} \times \mathbb{R}^{d_i} \times \dots \times \mathbb{R}^{d_n}.$$

For $\mathbf{d} \leq \mathbf{d}'$ in \mathbb{N}^{n+1} (with respect to the obvious partial order), view $\mathbb{E}_n^{\mathbf{d}}$ as the subspace

$$\mathbb{R}^{d_0} \times \{0\}^{d'_0 - d_0} \times \mathbb{R}_+ \times \dots \times \mathbb{R}_+ \times \mathbb{R}^{d_n} \times \{0\}^{d'_n - d_n} \subseteq \mathbb{E}_n^{\mathbf{d}'}$$




FIGURE 3.2. A neat embedding of the 2-dimensional $\langle 3 \rangle$ -manifold “triangle” in $(\mathbb{R}_+)^3$.

A neat immersion ι of an $\langle n \rangle$ -manifold is a smooth immersion $\iota: X \looparrowright \mathbb{E}_n^{\mathbf{d}}$ for some \mathbf{d} , such that:

- (1) ι is an n -map, i.e. $\iota^{-1}(\partial_i \mathbb{E}_n^{\mathbf{d}}) = \partial_i X$ for all i , and
- (2) the intersection of $X(a)$ and $\mathbb{E}_n^{\mathbf{d}}(b)$ is perpendicular, for all $b < a$ in $\{0, 1\}^n$.

A neat embedding is a neat immersion that is also an embedding. (See Figure 3.2 for an example.)

Given $\mathbf{d}' \in \mathbb{N}^{n+1}$ and a neat immersion $\iota: X \looparrowright \mathbb{E}_n^{\mathbf{d}}$, let $\iota[\mathbf{d}']: X \looparrowright \mathbb{E}_n^{\mathbf{d}+\mathbf{d}'}$ be the induced neat immersion.

The existence of neat embeddings is guaranteed by the following version of the Whitney embedding theorem for manifolds with corners.

Lemma 3.11. *Fix a compact $\langle n \rangle$ -manifold X . Let $\iota: \partial X \hookrightarrow \mathbb{E}_n^{\mathbf{d}}$ be an embedding such that $\iota^{-1}(\partial_i \mathbb{E}_n^{\mathbf{d}}) = \partial_i X$ for all i , and the intersection of $X(a)$ and $\mathbb{E}_n^{\mathbf{d}}(b)$ is perpendicular, for all $b < a$ in $\{0, 1\}^n \setminus \{\bar{1}\}$. Then $\iota[\mathbf{d}' - \mathbf{d}]$ can be extended to a neat embedding of X in $\mathbb{E}_n^{\mathbf{d}'}$ for some $\mathbf{d}' \geq \mathbf{d}$.*

Proof. The proof is standard, and is similar to the proof of [Lau00, Proposition 2.1.7], using the fact that faces of manifolds with corners admit collar neighborhoods [Lau00, Proposition 2.1.6]. □

Definition 3.12. Let X be a k -dimensional $\langle n \rangle$ -manifold. Given a neat immersion $\iota: X \looparrowright \mathbb{E}_n^{\mathbf{d}}$, the normal bundle ν_ι is defined as follows. For $a \in \{0, 1\}^n$ and $x \in X(a)$, the bundle ν_ι at x is the normal bundle of the immersion $\iota|_{X(a)}: X(a) \looparrowright \mathbb{E}_n^{\mathbf{d}}(a)$.

Equivalently, after setting $r = |\mathbf{d}| + n - k$, we can view the normal bundle as a map $\nu_\iota: X \rightarrow Gr(r, r + |a|) \rightarrow Gr(r) = BO(r)$, as follows. For $a \in \{0, 1\}^n$ and $x \in X(a)$, $\nu_\iota(x)$ is the orthogonal complement of $T_x(X(a))$ inside $T_x(\mathbb{E}_n^{\mathbf{d}}(a)) = \mathbb{R}^{r+|a|}$. The map is well-defined because for all $b \leq a$ in $\{0, 1\}^n$, the following diagram commutes.

$$\begin{array}{ccc}
 X(a) & \longrightarrow & Gr(r, r + |a|) \\
 \uparrow & & \uparrow \\
 X(b) & \longrightarrow & Gr(r, r + |b|).
 \end{array}$$

(Here, the vertical arrow on the right is induced from the inclusion $\mathbb{E}_n^{\mathbf{d}}(b) \hookrightarrow \mathbb{E}_n^{\mathbf{d}}(a)$.)

3.2. Flow categories. In this subsection, we recall the basic definitions related to flow categories and framed flow categories, and review briefly how they arise in Morse theory. In the next subsection we will show how to associate a CW complex to a framed flow category.

Definition 3.13. A *flow category* is a pair (\mathcal{C}, gr) where \mathcal{C} is a category with finitely many objects $\text{Ob} = \text{Ob}(\mathcal{C})$ and $\text{gr}: \text{Ob} \rightarrow \mathbb{Z}$ is a function, called the grading, satisfying the following additional conditions.

- (M-1) $\text{Hom}(x, x) = \{\text{Id}\}$ for all $x \in \text{Ob}$, and for distinct $x, y \in \text{Ob}$, $\text{Hom}(x, y)$ is a compact $(\text{gr}(x) - \text{gr}(y) - 1)$ -dimensional $\langle \text{gr}(x) - \text{gr}(y) - 1 \rangle$ -manifold (with the understanding that negative dimensional manifolds are empty).
- (M-2) For distinct $x, y, z \in \text{Ob}$ with $\text{gr}(z) - \text{gr}(y) = m$, the composition map

$$\circ: \text{Hom}(z, y) \times \text{Hom}(x, z) \rightarrow \text{Hom}(x, y)$$

is an embedding into $\partial_m \text{Hom}(x, y)$. Furthermore,

$$\circ^{-1}(\partial_i \text{Hom}(x, y)) = \begin{cases} \partial_i \text{Hom}(z, y) \times \text{Hom}(x, z) & \text{for } i < m \\ \text{Hom}(z, y) \times \partial_{i-m} \text{Hom}(x, z) & \text{for } i > m. \end{cases}$$

- (M-3) For distinct $x, y \in \text{Ob}$, \circ induces a diffeomorphism

$$(3.1) \quad \partial_i \text{Hom}(x, y) \cong \coprod_{\substack{z \\ \text{gr}(z)=\text{gr}(y)+i}} \text{Hom}(z, y) \times \text{Hom}(x, z).$$

Therefore, if \mathfrak{D} is the diagram whose vertices are the spaces $\text{Hom}(z_m, y) \times \text{Hom}(z_{m-1}, z_m) \times \cdots \times \text{Hom}(x, z_1)$ for $m \geq 1$ and distinct $z_1, \dots, z_m \in \text{Ob} \setminus \{x, y\}$, and whose arrows correspond to composing a single adjacent pair of Hom's, then

$$\partial \text{Hom}(x, y) \cong \bigcup_i \partial_i \text{Hom}(x, y) \cong \text{colim } \mathfrak{D}.$$

Given objects x, y in a flow category \mathcal{C} , define the *moduli space from x to y* to be

$$\mathcal{M}(x, y) = \begin{cases} \emptyset & \text{if } x = y \\ \text{Hom}(x, y) & \text{otherwise.} \end{cases}$$

Given a flow category \mathcal{C} and an integer n , let $\mathcal{C}[n]$ be the flow category obtained from \mathcal{C} by increasing the grading of each object by n .

Key motivation for Definition 3.13 comes from the following.

Definition 3.14. Given a Morse function $f: M \rightarrow \mathbb{R}$ with finitely many critical points and a gradient-like flow for f , one can construct the *Morse flow category* $\mathcal{C}_{\mathcal{M}}(f)$ as follows.

The objects of $\mathcal{C}_{\mathcal{M}}(f)$ are the critical points of f ; the grading of an object is the index of the critical point. For critical points x, y of relative index $n + 1$, let $\tilde{\mathcal{M}}(x, y)$ be the *parametrized moduli space* from x to y , which is the $(n + 1)$ -dimensional subspace of M consisting of all the points that flow up to x and flow down to y . Let $\mathcal{M}(x, y)$ be the quotient of $\tilde{\mathcal{M}}(x, y)$ by the natural \mathbb{R} -action. Define $\mathcal{M}(x, y)$, the space of morphisms from x to y , to be the *unparametrized moduli space* from x to y , which is obtained by compactifying $\tilde{\mathcal{M}}(x, y)$ by adding all the broken flowlines from x to y , cf. [AB95, Lemma 2.6].

Definition 3.15. Let \mathcal{C} be a flow category. Given an integer i , let $\text{Ob}(i)$ denote the set of all objects of \mathcal{C} in grading i , topologized as a discrete space. Given integers i and j , let

$$\mathcal{M}(i, j) = \coprod_{x \in \text{Ob}(i), y \in \text{Ob}(j)} \mathcal{M}(x, y).$$

There are obvious maps from $\mathcal{M}(i, j)$ to $\text{Ob}(i)$ and $\text{Ob}(j)$. For all m, n, i with $1 \leq i < m - n$, equation (3.1) induces a diffeomorphism between the fiber product $\mathcal{M}(n + i, n) \times_{\text{Ob}(n+i)} \mathcal{M}(m, n + i)$ and $\partial_i \mathcal{M}(m, n)$.

We turn next to (framed) neat embeddings of flow categories; such a framed embedding is part of the input to the Cohen-Jones-Segal construction. Because we will consider covering spaces later on, we will actually work with a slightly more general notion: neat immersions (as we did with $\langle n \rangle$ -manifolds); see also Section 3.4.1. Still, the Cohen-Jones-Segal construction requires an embedding, so we will show in Lemma 3.23 that any framed neat immersion can be perturbed to a framed neat embedding in an essentially unique way.

Definition 3.16. For each integer i , fix a natural number d_i and let \mathbf{d} denote this sequence. For $a < b$, let $\mathbb{E}_{\mathbf{d}}[a : b] = \mathbb{E}_{b-a-1}^{(d_a, \dots, d_{b-1})}$.

A neat immersion (resp. neat embedding) ι of a flow category \mathcal{C} relative \mathbf{d} is a collection of neat immersions (resp. neat embeddings) $\iota_{x,y} : \mathcal{M}(x, y) \looparrowright \mathbb{E}_{\mathbf{d}}[\text{gr}(y) : \text{gr}(x)]$, defined for all objects x, y , such that the following conditions hold.

- (N-1) For all integers i, j , ι induces a neat immersion (resp. neat embedding) $\iota_{i,j}$ of $\mathcal{M}(i, j)$.
- (N-2) For all objects x, y, z , and for all points $(p, q) \in \mathcal{M}(x, z) \times \mathcal{M}(z, y)$,

$$\iota_{x,y}(q \circ p) = (\iota_{z,y}(q), 0, \iota_{x,z}(p)).$$

(See Figure 3.3b for an example.)

Given a neat immersion ι of a flow category \mathcal{C} relative \mathbf{d} , and some other \mathbf{d}' , let $\iota[\mathbf{d}']$ denote the neat immersion relative $\mathbf{d} + \mathbf{d}'$ induced by the inclusions $\mathbb{E}_{\mathbf{d}}[a : b] \hookrightarrow \mathbb{E}_{\mathbf{d}+\mathbf{d}'}[a : b]$, which in turn are induced by the canonical inclusions $\mathbb{R}^{d_i} \cong \mathbb{R}^{d_i} \times \{0\} \hookrightarrow \mathbb{R}^{d_i+d'_i}$.

Lemma 3.17. Any flow category admits a neat embedding relative to some \mathbf{d} .

Proof. This follows from Lemma 3.11, by induction. □

Lemma 3.18. Let \mathbb{D}^k be a k -dimensional disk, and let S^{k-1} denote its boundary. If ι is a smooth S^{k-1} -parameter family of neat immersions of a flow category \mathcal{C} relative to some fixed \mathbf{d} , then for \mathbf{d}' sufficiently large, $\iota[\mathbf{d}']$ can be extended to a smooth \mathbb{D}^k -parameter family of neat immersions ι' of \mathcal{C} . Furthermore, if ι is a neat embedding for some point in S^{k-1} , then we can ensure that ι' is a neat embedding at all points in the interior of \mathbb{D}^k .

Proof. We follow the proof of [Lau00, Proposition 2.2.3]. Define \mathbf{d}' by setting $d'_i = d_i + 1$. Assume \mathbb{D}^k is the unit disk in \mathbb{R}^k . We use ‘spherical coordinates’ (r, θ) to describe points on \mathbb{D}^k , where r is the distance from the origin and θ is a point on S^{k-1} .

Fix objects x, y with $\text{gr}(x) = m$ and $\text{gr}(y) = n$. Recall that

$$\mathbb{E}_{\mathbf{d}'}[n : m] = \mathbb{R}^{d_n} \times \mathbb{R}^{d_m} \times \mathbb{R} \times \mathbb{R}_+ \times \cdots \times \mathbb{R}_+ \times \mathbb{R}^{d_{m-1}} \times \mathbb{R}^{d_{m-1}} \times \mathbb{R}.$$

Hence, write $\iota_{x,y}(\theta)[\mathbf{d}']$ in coordinate form as

$$\iota_{x,y}(\theta)[\mathbf{d}'] = (\iota_n(\theta), 0, 0, \bar{\iota}_{n+1}(\theta), \dots, \bar{\iota}_{m-1}(\theta), \iota_{m-1}(\theta), 0, 0).$$

Let $f: [0, 1] \rightarrow \mathbb{R}_+$ be defined as

$$f(r) = \begin{cases} e^{-1/(r-r^2)} & \text{if } r \neq 0, 1 \\ 0 & \text{otherwise.} \end{cases}$$

Choose some point $\theta_0 \in S^{k-1}$; if possible, choose a θ_0 such that $\iota(\theta_0)$ is a neat embedding. Define the isotopy to be

$$\begin{aligned} \iota'_{x,y}(r, \theta) = & ((1-r)\iota_n(\theta_0) + r\iota_n(\theta), f(r)\iota_n(\theta_0), f(r)\bar{\iota}_{n+1}(\theta_0), \\ & (1-r)\bar{\iota}_{n+1}(\theta_0) + r\bar{\iota}_{n+1}(\theta), \dots, (1-r)\bar{\iota}_{m-1}(\theta_0) + r\bar{\iota}_{m-1}(\theta), \\ & (1-r)\iota_{m-1}(\theta_0) + r\iota_{m-1}(\theta), f(r)\iota_{m-1}(\theta_0), 0). \end{aligned} \quad \square$$

Definition 3.19. Let ι be a neat immersion of a flow category \mathcal{C} relative \mathbf{d} . A *coherent framing* φ for ι is a framing for $\nu_{x,y}$ for all objects x, y , such that the product framing of $\nu_{z,y} \times \nu_{x,z}$ equals the pullback framing of $\circ^* \nu_{\iota_{x,y}}$ for all x, y, z . (Again, see Figure 3.3b for an example.)

Lemma 3.20. Let \mathbb{D}^k be a k -dimensional disk and let \mathbb{D}^{k-1} be a hemisphere on its boundary. Assume ι is a smooth \mathbb{D}^k -parameter family of neat immersions of a flow category \mathcal{C} relative to some fixed \mathbf{d} , and assume φ is a smooth \mathbb{D}^{k-1} -parameter family of coherent framings for $\iota|_{\mathbb{D}^{k-1}}$. Then φ can be extended to a smooth \mathbb{D}^k -parameter family of coherent framings for ι .

Proof. For any x, y , $\iota_{x,y}$ is a \mathbb{D}^k -parameter family of neat immersions of $\mathcal{M}(x, y)$ in $\mathbb{E}_{\mathbf{d}}[\text{gr}(y) : \text{gr}(x)]$. The normal bundle gives a map $\mathcal{M}(x, y) \times \mathbb{D}^k \rightarrow BO(r)$, where we have set $r = d_{\text{gr}(y)} + \cdots + d_{\text{gr}(x)-1}$. The framing φ of $\iota|_{\mathbb{D}^{k-1}}$ produces a map $\mathcal{M}(x, y) \times \mathbb{D}^{k-1} \rightarrow EO(r)$ such that the following commutes:

$$\begin{array}{ccc} \mathcal{M}(x, y) \times \mathbb{D}^{k-1} & \longrightarrow & EO(r) \\ \downarrow & & \downarrow \\ \mathcal{M}(x, y) \times \mathbb{D}^k & \longrightarrow & BO(r) \end{array}$$

We want to produce a map $\mathcal{M}(x, y) \times \mathbb{D}^k \rightarrow EO(r)$ which will commute with the previous diagram, and will induce a coherent framing. We do this by an induction on $\text{gr}(x) - \text{gr}(y)$.

The base case, when $\text{gr}(x) - \text{gr}(y) = 0$, is vacuous. Now assume that we have framed $\nu_{x',y'}(t)$ for all $t \in \mathbb{D}^k$ and all x', y' with $\text{gr}(x') - \text{gr}(y') < \text{gr}(x) - \text{gr}(y)$. Since every point in the boundary of $\mathcal{M}(x, y)$ lies in the product of lower dimensional moduli spaces, the normal bundle $\nu_{\iota_{x,y}}(t)$, restricted to $\partial\mathcal{M}(x, y)$, is already framed. Thus, we have the following commuting diagram; we want to produce the dashed

arrow so that the diagram still commutes.

$$\begin{array}{ccc}
 (\mathcal{M}(x, y) \times \mathbb{D}^{k-1}) \cup (\partial\mathcal{M}(x, y) \times \mathbb{D}^k) & \longrightarrow & EO(r) \\
 \downarrow & \dashrightarrow & \downarrow \\
 \mathcal{M}(x, y) \times \mathbb{D}^k & \longrightarrow & BO(r)
 \end{array}$$

By the collar neighborhood theorem, there is a homeomorphism of pairs

$$\begin{aligned}
 &(\mathcal{M}(x, y) \times \mathbb{D}^k, (\mathcal{M}(x, y) \times \mathbb{D}^{k-1}) \cup (\partial\mathcal{M}(x, y) \times \mathbb{D}^k)) \\
 &\cong (\mathcal{M}(x, y) \times \mathbb{D}^k, \mathcal{M}(x, y) \times \mathbb{D}^{k-1}).
 \end{aligned}$$

So, since $EO(r) \rightarrow BO(r)$ is a fibration, the dashed arrow exists. □

Definition 3.21. A *framed flow category* is a neatly embedded flow category \mathcal{C} , along with a coherent framing for some neat immersion of \mathcal{C} (relative to some \mathbf{d}).

To a framed flow category, one can associate a chain complex $C^*(\mathcal{C})$ as follows. The n^{th} chain group C^n is the \mathbb{Z} -module freely generated by $\text{Ob}(n)$. The differential δ is of degree one. For $x, y \in \text{Ob}$ with $\text{gr}(x) = \text{gr}(y) + 1$, the coefficient of δy evaluated at x is the number of points in $\mathcal{M}(x, y)$, counted with sign (recall, $\mathcal{M}(x, y)$ is a compact framed 0-dimensional $\langle 0 \rangle$ -manifold). We say a framed flow category *refines* its associated chain complex.

Definition 3.22. Let φ be a coherent framing for some neat immersion ι of a flow category \mathcal{C} . Using Lemma 3.17, choose some neat embedding ι' of \mathcal{C} . Using Lemma 3.18, connect $\iota[\mathbf{d}]$ to $\iota'[\mathbf{d}']$ for some \mathbf{d}, \mathbf{d}' . The coherent framing φ induces a coherent framing for $\iota[\mathbf{d}]$, also denoted by φ . Therefore, using Lemma 3.20, produce a coherent framing φ' for $\iota'[\mathbf{d}']$. We say that the framed flow category $(\mathcal{C}, \iota'[\mathbf{d}'], \varphi')$ is a *perturbation* of $(\mathcal{C}, \iota, \varphi)$.

Lemma 3.23. If $(\mathcal{C}, \iota_0, \varphi_0)$ and $(\mathcal{C}, \iota_1, \varphi_1)$ are two perturbations of $(\mathcal{C}, \iota, \varphi)$, then for some $\mathbf{d}_0, \mathbf{d}_1$, there exists a smooth 1-parameter family of framings of \mathcal{C} that connects $(\iota_0[\mathbf{d}_0], \varphi_0)$ to $(\iota_1[\mathbf{d}_1], \varphi_1)$.

Proof. For $i \in \{0, 1\}$, since $(\mathcal{C}, \iota_i, \varphi_i)$ is a perturbation of $(\mathcal{C}, \iota, \varphi)$, for some \mathbf{d}'_i there exists a 1-parameter family $\iota_i(t)$ of neat immersions connecting $\iota[\mathbf{d}'_i]$ to ι_i , along with a 1-parameter family of coherent framings $\varphi_i(t)$ for $\iota_i(t)$.

Using Lemma 3.18, choose a 1-parameter family of neat embeddings $\iota(t)$ connecting $\iota_0[\mathbf{d}''_0]$ to $\iota_1[\mathbf{d}''_1]$ for some \mathbf{d}''_i . Consider a triangle Δ , whose three vertices are v, v_0, v_1 , with opposite edges e, e_0, e_1 , respectively. The three 1-parameter families of neat immersions $\iota(t), \iota_0(t)[\mathbf{d}''_0]$, and $\iota_1(t)[\mathbf{d}''_1]$ together produce a $\partial\Delta$ -parameter family of neat immersions of \mathcal{C} : the restriction to e is $\iota(t)$, and the restriction to e_i is $\iota_i[\mathbf{d}''_i]$.

For \mathbf{d} sufficiently large, using Lemma 3.18 again, choose a Δ -parameter family $\bar{\iota}$ of neat immersions of \mathcal{C} , whose restriction to e is $\iota(t)[\mathbf{d}]$, and whose restriction to e_i is $\iota_i(t)[\mathbf{d}''_i + \mathbf{d}]$. Set $\mathbf{d}_i = \mathbf{d}''_i + \mathbf{d}$. The framings $\phi_0(t)$ and $\phi_1(t)$ induce framings of \mathcal{C} along the edges e_0 and e_1 of Δ . By Lemma 3.20 we can extend these framings to a Δ -parameter family $\bar{\varphi}$ of coherent framings for $\bar{\iota}$. Then $(\bar{\iota}|_e, \bar{\varphi}|_e)$ is a 1-parameter family of framings connecting $\iota_0[\mathbf{d}_0]$ to $\iota_1[\mathbf{d}_1]$. □

3.3. Framed flow categories to CW complexes. We are interested in framed flow categories because one can build a CW complex $|\mathcal{C}|$ from a framed flow category \mathcal{C} in such a way that if \mathcal{C} refines a chain complex C^* then C^* is the cellular cochain complex of $|\mathcal{C}|$. The construction of $|\mathcal{C}|$, which was first given in [CJS95] (partly inspired by [Fra79]), is not the same as the usual geometric realization of a (topological) category. In this subsection we give a slight reformulation of the construction from [CJS95]; our formulation is less elegant but perhaps more concrete.

Definition 3.24. Let $(\mathcal{C}, \iota, \varphi)$ be a framed flow category, where \mathcal{C} is a flow category, ι is a neat embedding relative some \mathbf{d} , and φ is a framing of the normal bundles to ι . Assume that all objects in \mathcal{C} have grading in $[B, A]$ for some fixed $A, B \in \mathbb{Z}$.

We associate a CW complex $|\mathcal{C}|_{\iota, \varphi}$ to $(\mathcal{C}, \iota, \varphi)$ as follows, cf. Figure 3.3. The CW complex $|\mathcal{C}|$ will have one 0-cell, and one cell $\mathcal{C}(x)$ for each object x of \mathcal{C} . If x has grading m then the dimension of $\mathcal{C}(x)$ will be $d_B + \cdots + d_{A-1} - B + m$. For convenience, let $C = d_B + \cdots + d_{A-1} - B$.

Since ι is a neat embedding of \mathcal{C} , for all integers i, j , $\iota_{i,j}$ embeds $\mathcal{M}(i, j)$ in $\mathbb{E}_{\mathbf{d}}[j : i]$. Choose $\epsilon > 0$ sufficiently small so that for all i, j , $\iota_{i,j}$ extends to an embedding of $\mathcal{M}(i, j) \times [-\epsilon, \epsilon]^{d_j + \cdots + d_{i-1}}$ via the framings of the normal bundles. Choose R sufficiently large so that for all i, j , $\iota_{i,j}(\mathcal{M}(i, j) \times [-\epsilon, \epsilon]^{d_j + \cdots + d_{i-1}})$ lies in $[-R, R]^{d_j} \times [0, R] \times \cdots \times [0, R] \times [-R, R]^{d_{i-1}}$.

Suppose that we have defined the $(C + m - 1)$ -skeleton $|\mathcal{C}|^{(C+m-1)}$ of $|\mathcal{C}|$, i.e., the part of $|\mathcal{C}|$ built from objects y with grading $\text{gr}(y) < m$. Fix an object x with $\text{gr}(x) = m$. While describing $\mathcal{C}(x)$, we will define some related spaces $\mathcal{C}_i(x)$. Take a copy $\mathcal{C}_1(x)$ of $\mathbb{R}_+ \times \mathbb{R}^{d_B} \times \cdots \times \mathbb{R}_+ \times \mathbb{R}^{d_{A-1}}$.⁽ⁱ⁾ Define $\mathcal{C}(x)$ to be the following subset:

$$(3.2) \quad \begin{aligned} \mathcal{C}(x) = & [0, R] \times [-R, R]^{d_B} \times \cdots \times [0, R] \times [-R, R]^{d_{m-1}} \times \{0\} \times [-\epsilon, \epsilon]^{d_m} \\ & \times \cdots \times \{0\} \times [-\epsilon, \epsilon]^{d_{A-1}} \subset \mathcal{C}_1(x). \end{aligned}$$

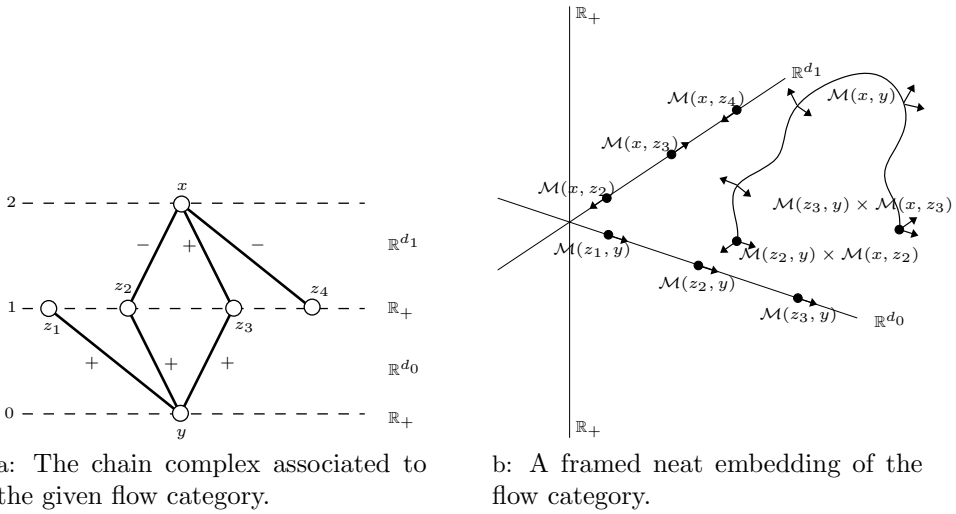
It remains to define the attaching map $\partial\mathcal{C}(x) \rightarrow |\mathcal{C}|^{(C+m-1)}$. Recall that the neat embedding $\iota_{x,y}$, along with the framing φ of $\nu_{x,y}$, identifies $\mathcal{M}(x, y) \times [-\epsilon, \epsilon]^{d_n} \times \{0\} \times \cdots \times \{0\} \times [-\epsilon, \epsilon]^{d_{m-1}}$ with a subset $\mathcal{C}_{y,1}(x)$ of $[-R, R]^{d_n} \times [0, R] \times \cdots \times [0, R] \times [-R, R]^{d_{m-1}}$. Let

$$\begin{aligned} \mathcal{C}_y(x) = & [0, R] \times [-R, R]^{d_B} \times \cdots \times [0, R] \times [-R, R]^{d_{n-1}} \times \{0\} \times \mathcal{C}_{y,1}(x) \\ & \times \{0\} \times [-\epsilon, \epsilon]^{d_m} \times \cdots \times \{0\} \times [-\epsilon, \epsilon]^{d_{A-1}} \subset \partial\mathcal{C}(x). \end{aligned}$$

Now, define the attaching map $\partial\mathcal{C}(x) \rightarrow |\mathcal{C}|^{(C+m-1)}$ as follows: on $\mathcal{C}_y(x) \cong \mathcal{M}(x, y) \times \mathcal{C}(y)$, define it to be the projection map to $\mathcal{C}(y)$; and map $\partial\mathcal{C}(x) \setminus \bigcup_y \mathcal{C}_y(x)$ to the basepoint.


Figure 3.3 illustrates Definition 3.24 for a chain complex with six generators spread over three gradings. All the cells are subsets of $\mathbb{E} = \mathbb{R}_+ \times \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}$. We

⁽ⁱ⁾It is important that we take a copy. For distinct objects x, y , we will take different copies $\mathcal{C}_1(x)$ and $\mathcal{C}_1(y)$, and we will not, *a priori*, identify the corresponding points in $\mathcal{C}_1(x)$ and $\mathcal{C}_1(y)$.



a: The chain complex associated to the given flow category.

b: A framed neat embedding of the flow category.



c: $\mathcal{C}(y)$ is the shaded square with the boundary quotiented to the basepoint.

d: $\mathcal{C}(z_1)$ is the deeply shaded slab with the lightly shaded square on the boundary identified with $\mathcal{C}(y)$ and the rest of the boundary quotiented to the basepoint.

FIGURE 3.3. From a framed flow category to a CW complex. (Part I)

do not draw \mathbb{E} , but we attempt to draw its boundary $\partial E = (\{0\} \times \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}) \cup (\mathbb{R}_+ \times \mathbb{R} \times \{0\} \times \mathbb{R})$. The space ∂E is naturally a 3-dimensional $\langle 2 \rangle$ -manifold. To draw it, we flatten out the codimension-1 corner $\{0\} \times \mathbb{R} \times \{0\} \times \mathbb{R}$, giving a homeomorphism (but not a diffeomorphism) $\partial E \cong \mathbb{R}^3$. In the figures, this (flattened) corner is a horizontal plane, the first \mathbb{R}_+ -factor is drawn below the plane, and the second \mathbb{R}_+ -factor is drawn above the plane.

Lemma 3.25. *Definition 3.24 produces a well-defined CW complex $|\mathcal{C}|$. Furthermore, the chain complex associated to \mathcal{C} is isomorphic to $\tilde{C}^*(|\mathcal{C}|_{\iota, \varphi})[-C]$, where*




FIGURE 3.4. From a framed flow category to a CW complex. (Part II)

$\tilde{\mathcal{C}}^*$ is the reduced cellular cochain complex and $[\]$ is the degree shift operator. (For the signs in the cellular cochain complex, we orient the cells $\mathcal{C}(x)$ from (3.2) by the product orientations.)

Proof. The first part follows from the following elementary check. For arbitrary x, y, z with $\text{gr}(x) \geq \text{gr}(z) \geq \text{gr}(y)$, the dashed arrow in the following diagram exists

such that the diagram commutes:

$$\begin{array}{ccccc}
 \mathcal{C}_y(x) \cap \mathcal{C}_z(x) & \hookrightarrow & \mathcal{C}_z(x) & \longrightarrow & \mathcal{C}(z) \\
 \downarrow & & & \searrow & \downarrow \\
 \mathcal{C}_y(x) & & & & \partial\mathcal{C}(z) \\
 \downarrow & & & & \uparrow \\
 \mathcal{C}(y) & \longleftarrow & & & \mathcal{C}_y(z)
 \end{array}$$

For the second part, fix x, y with $\text{gr}(y) = \text{gr}(x) - 1$. Then the map $\pi: \partial\mathcal{C}(x) \rightarrow \mathcal{C}(y)/\partial\mathcal{C}(y)$ has degree $\#\pi^{-1}(p)$ for any p in the interior of $\mathcal{C}(y)$. However, since $\mathcal{C}_y(x)$ is canonically homeomorphic to $\mathcal{M}(x, y) \times \mathcal{C}(y)$, the number of points in $\pi^{-1}(p)$ equals $(-1)^{d_B + \dots + d_{m-2} + m - B}$ times the number of points in $\mathcal{M}(x, y)$. The sign $(-1)^{d_B + \dots + d_{m-2} + m - B}$ comes from the outward-normal first convention for the boundary orientation—the opposite of the orientation of $0 \in \partial[0, R]$ —and commuting the factor $[0, R]$ in $\mathcal{C}(x)$ to the first position (which requires $d_B + \dots + d_{m-2} + m - B - 1$ exchanges). \square

Lemma 3.26. *For a framed flow category $(\mathcal{C}, \iota, \varphi)$, the isomorphism type of the CW complex $|\mathcal{C}|_{\iota, \varphi}$ from Definition 3.24 is independent of the choice of real numbers ϵ, R . In particular, different choices of ϵ, R give homeomorphic spaces.*

Furthermore, for $t \in [0, 1]$, if $(\iota(t), \varphi(t))$ is a smooth 1-parameter family of framings of a flow category \mathcal{C} , then the CW complexes $|\mathcal{C}|_{\iota(0), \varphi(0)}$ and $|\mathcal{C}|_{\iota(1), \varphi(1)}$ are isomorphic.

Proof. This is clear from the definitions. Namely, changing ϵ or R or deforming the framing or framed embedding has the effect of changing the boxes $\mathcal{C}(x) \subset \mathcal{C}_1(x) = \mathbb{R}^N$ and sub-boxes $\mathcal{C}_y(x)$ by ambient isotopies of \mathbb{R}^N ; and these isotopies are consistent across different x 's in an obvious sense. Following these ambient isotopies gives a homeomorphism between the CW complexes $|\mathcal{C}|$ constructed with the two collections of data, and this homeomorphism takes cells homeomorphically to cells. \square

Lemma 3.27. *Fix a framed flow category $(\mathcal{C}, \iota, \varphi)$, where ι is a neat embedding relative \mathbf{d} . Let $|\mathcal{C}|_{\iota, \varphi, B, A}$ be the CW complex defined in Definition 3.24, where the subscripts also indicate our choice of integers A, B . Let $C_{\mathbf{d}}(B, A) = d_B + \dots + d_{A-1} - B$ denote the grading shift.*

For any $\mathbf{d}' \geq \mathbf{d}$, let $\iota' = \iota[\mathbf{d}' - \mathbf{d}]$ be the induced neat embedding relative \mathbf{d}' ; let φ' denote the induced coherent framing for ι' . For any $A' \geq A$ and $B' \leq B$, there is a homotopy equivalence

$$|\mathcal{C}|_{\iota', \varphi', B', A'} \sim \Sigma^{C_{\mathbf{d}'}(B', A') - C_{\mathbf{d}}(B, A)} |\mathcal{C}|_{\iota, \varphi, B, A}.$$

Proof. For notational convenience, if we denote some object in $|\mathcal{C}|_{\iota, \varphi, B, A}$ by some symbol $*$, then we will denote the corresponding object in $|\mathcal{C}|_{\iota', \varphi', B', A'}$ by the symbol $*'$. It is enough to consider the following three cases.

- (1) $A' = A + 1; B' = B; \mathbf{d}' = \mathbf{d}; d_A = 0$. In this case, from (3.2) it is clear that $\mathcal{C}(x)' = \mathcal{C}(x)$ for all objects x in \mathcal{C} ; furthermore, this identification respects the attaching maps. Therefore, we have a homeomorphism

$$|\mathcal{C}|_{\iota', \varphi', B', A'} \cong |\mathcal{C}|_{\iota, \varphi, B, A} = \Sigma^{C_{\mathbf{d}'}(B', A') - C_{\mathbf{d}}(B, A)} |\mathcal{C}|_{\iota, \varphi, B, A}.$$

- (2) $A' = A; B' = B - 1; \mathbf{d}' = \mathbf{d}; d_{B-1} = 0$. In this case, it follows from (3.2) that for all objects x , $\mathcal{C}(x)' = [0, R] \times \mathcal{C}(x)$; and for all objects x, y , with $\mathcal{C}_y(x)' = [0, R] \times \mathcal{C}_y(x)$. Furthermore, in the attaching maps, the whole of $\{0, R\} \times \mathcal{C}(x)$ is quotiented to the basepoint. Therefore, we have a homeomorphism

$$|\mathcal{C}|_{\nu', \varphi', B', A'} \cong S^1 \wedge |\mathcal{C}|_{\nu, \varphi, B, A} = \Sigma^{C_{\mathbf{d}'}(B', A') - C_{\mathbf{d}}(B, A)} |\mathcal{C}|_{\nu, \varphi, B, A}.$$

- (3) $A' = A; B' = B; d'_m = d_m + 1$ for some $m \in [B, A - 1]$ and $d'_i = d_i$ otherwise. In this case, $C_{\mathbf{d}'}(B', A') - C_{\mathbf{d}}(B, A) = 1$, and we want to produce a homotopy equivalence between $|\mathcal{C}|_{\nu', \varphi', B', A'}$ and $\Sigma |\mathcal{C}|_{\nu, \varphi, B, A}$.

For all objects x with $\text{gr}(x) \leq m$, $\mathcal{C}(x)'$ can be identified with $[-\epsilon, \epsilon] \times \mathcal{C}(x)$, and for all objects x with $\text{gr}(x) > m$, $\mathcal{C}(x)'$ can be identified with $[-R, R] \times \mathcal{C}(x)$. Furthermore, for all objects x, y with $\text{gr}(y) \leq m$, $\mathcal{C}_y(x)'$ can be identified with $[-\epsilon, \epsilon] \times \mathcal{C}_y(x)$, and for all objects x, y with $\text{gr}(y) > m$, $\mathcal{C}_y(x)'$ can be identified with $[-R, R] \times \mathcal{C}_y(x)$.

Isotoping the cells $\mathcal{C}(x) \subset \mathbb{R}^N$ and the subsets $\mathcal{C}_y(x) \subset \mathcal{C}(x)$ (in a consistent way) has the effect of homotoping the attaching maps in $|\mathcal{C}|$ (in a consistent way), and hence does not change the homotopy type of $|\mathcal{C}|$. So, for $t \in [\epsilon, R]$, consider the cells $\mathcal{C}(x)^t$ and subsets $\mathcal{C}_y(x)^t$ defined as follows: for all objects x with $\text{gr}(x) \leq m$, set $\mathcal{C}(x)^t = \mathcal{C}(x)'$, and for all objects x with $\text{gr}(x) > m$, set $\mathcal{C}(x)^t = [-t, t] \times \mathcal{C}(x)$. Similarly, for all objects x, y with $\text{gr}(y) \leq m$, set $\mathcal{C}_y(x)^t = \mathcal{C}_y(x)'$, and for all objects x, y with $\text{gr}(y) > m$, set $\mathcal{C}_y(x)^t = [-t, t] \times \mathcal{C}_y(x)$. This produces a homotopy equivalence between $|\mathcal{C}|^\epsilon$ and $|\mathcal{C}|^R = |\mathcal{C}|_{\nu', \varphi', B', A'}$.

Therefore, in $|\mathcal{C}|^\epsilon$, $\mathcal{C}(x)^\epsilon$ is identified with $[-\epsilon, \epsilon] \times \mathcal{C}(x)$, and $\mathcal{C}_y(x)^\epsilon$ is identified with $[-\epsilon, \epsilon] \times \mathcal{C}_y(x)$ for all x, y . As in Case (2), we get the homotopy equivalence

$$|\mathcal{C}|_{\nu', \varphi', B', A'} = |\mathcal{C}|^R \simeq |\mathcal{C}|^\epsilon \cong S^1 \wedge |\mathcal{C}|_{\nu, \varphi, B, A}. \quad \square$$

Proposition 3.28. *This construction from Definition 3.24 agrees with the construction from [CJS95].*

Proof. To start, we briefly review the Cohen-Jones-Segal formulation, from [CJS95, pp. 309–312]. For any topological space X , let X^+ denote its one-point compactification (with the understanding that $X^+ = X \amalg \text{pt}$ if X is already compact). Once again, choose integers A, B , such that all objects of \mathcal{C} have grading in $[B, A]$. Consider the topological category \mathcal{J}_{B-1}^A with $\text{Ob}_{\mathcal{J}} = \{n \in \mathbb{Z} \mid B - 1 \leq n \leq A\}$ and morphisms

$$\text{Hom}_{\mathcal{J}}(m, n) = \begin{cases} \emptyset & m < n \\ \{\text{Id}\} & m = n \\ (\mathbb{R}^{m-n-1})^+ & m > n. \end{cases}$$

Composition is given by the map

$$\mathbb{R}^{n-p-1} \times \mathbb{R}^{m-n-1} \rightarrow \mathbb{R}^{m-p-1}$$

$$(t_{p+1}, \dots, t_{n-1}) \times (t_{n+1}, \dots, t_{m-1}) \mapsto (t_{p+1}, \dots, t_{n-1}, 0, t_{n+1}, \dots, t_{m-1}).$$

Given a continuous, basepoint-preserving functor Z from \mathcal{J}_{B-1}^A to the category of based topological spaces, one can form the geometric realization $|Z|$ of Z , which

is obtained from

$$(3.3) \quad \coprod_{B-1 \leq m \leq A} \text{Hom}_{\mathcal{J}}(m, B-1) \wedge Z(m) / \sim,$$

where $(s \circ t, x) \sim (s, Z(t)(x))$ for any $s \in \text{Hom}_{\mathcal{J}}(n, B-1)$, $t \in \text{Hom}_{\mathcal{J}}(m, n)$ and $x \in Z(m)$.

So, to produce a space, it suffices to produce a functor from \mathcal{J} to based topological spaces. Cohen-Jones-Segal do this as follows. Define

$$Z(m) = (\text{Ob}_{\mathcal{C}}(m) \times \mathbb{R}^{d_B} \times \cdots \times \mathbb{R}^{d_{m-1}} \times (-\epsilon, \epsilon)^{d_m} \times \cdots \times (-\epsilon, \epsilon)^{d_{A-1}})^+.$$

The map $\text{Hom}_{\mathcal{J}}(m, n) \times Z(m) \rightarrow Z(n)$ is given as follows. There is a proper projection

$$\begin{aligned} &\mathbb{R}^{d_B} \times \cdots \times \mathbb{R}^{d_{n-1}} \times \mathcal{M}(m, n) \times (-\epsilon, \epsilon)^{d_n} \times \cdots \times (-\epsilon, \epsilon)^{d_{A-1}} \\ &\rightarrow \text{Ob}(n) \times \mathbb{R}^{d_B} \times \cdots \times \mathbb{R}^{d_{n-1}} \times (-\epsilon, \epsilon)^{d_n} \times \cdots \times (-\epsilon, \epsilon)^{d_{A-1}}. \end{aligned}$$

The framing identifies the normal bundle to $\mathcal{M}_{\mathcal{C}}(m, n)$ in $\mathbb{R}^{d_n} \times \mathbb{R}_+ \times \cdots \times \mathbb{R}_+ \times \mathbb{R}^{d_{m-1}} \cong \mathbb{R}_+^{m-n-1} \times \mathbb{R}^{d_n} \times \cdots \times \mathbb{R}^{d_{m-1}}$ with $\mathcal{M}_{\mathcal{C}}(m, n) \times (-\epsilon, \epsilon)^{d_n} \times \cdots \times (-\epsilon, \epsilon)^{d_{m-1}}$. So, there is an open inclusion

$$\begin{aligned} &\mathbb{R}^{d_B} \times \cdots \times \mathbb{R}^{d_{n-1}} \times \mathcal{M}(m, n) \times (-\epsilon, \epsilon)^{d_n} \times \cdots \times (-\epsilon, \epsilon)^{d_{A-1}} \\ &\hookrightarrow \mathbb{R}_+^{m-n-1} \times \text{Ob}(m) \times \mathbb{R}^{d_B} \times \cdots \times \mathbb{R}^{d_{m-1}} \times (-\epsilon, \epsilon)^{d_m} \times \cdots \times (-\epsilon, \epsilon)^{d_{A-1}}. \end{aligned}$$

Collapsing everything outside the image of the second map gives an umkehr map the other direction of one-point compactifications. Composing this with the first map gives a map $\text{Hom}_{\mathcal{J}}(m, n) \wedge Z(m) \rightarrow Z(n)$, as desired.

It remains to identify the realization $|Z|$ of Z with the space $|\mathcal{C}|$ of Definition 3.24. Consider the subset

$$\begin{aligned} \mathcal{C}'(x) = &[0, R) \times (-R, R)^{d_B} \times [0, R) \times (-R, R)^{d_{B+1}} \times \cdots \times [0, R) \times (-R, R)^{d_{m-1}} \\ &\times \{0\} \times (-\epsilon, \epsilon)^{d_m} \times \cdots \times \{0\} \times (-\epsilon, \epsilon)^{d_{A-1}} \subset \mathcal{C}(x). \end{aligned}$$

Then $(\coprod_{\text{gr}(x)=m} \mathcal{C}'(x))^+$ is exactly the space $\text{Hom}_{\mathcal{J}}(m, B) \wedge Z(m)$ appearing in equation (3.3) (except that we have used R instead of ∞).

Under this identification, the umkehr map is the identification of $\mathcal{C}'_y(x) \subset \partial \mathcal{C}'(x)$ with $\mathcal{M}(x, y) \times \mathcal{C}'(y)$ from Definition 3.24. The proper projection corresponds to the projection $\mathcal{M}(x, y) \times \mathcal{C}'(y) \rightarrow \mathcal{C}'(y)$. Thus, the map $\text{Hom}_{\mathcal{J}}(n, B) \wedge \text{Hom}_{\mathcal{J}}(m, n) \wedge Z(m) \rightarrow \text{Hom}_{\mathcal{J}}(n, B) \wedge Z(n)$ is the projection of $\mathcal{C}'_y(x)$ to $\mathcal{C}'(y)$, while the map $\text{Hom}_{\mathcal{J}}(n, B) \wedge \text{Hom}_{\mathcal{J}}(m, n) \wedge Z(m) \rightarrow \text{Hom}_{\mathcal{J}}(m, B) \wedge Z(m)$ is the inclusion of $\mathcal{C}'_y(x)$ in $\partial \mathcal{C}'(x)$. So, the identification \sim in (3.3) agrees with the identification in the CW complex from Definition 3.24. \square

3.4. Covers, subcomplexes, and quotient complexes. We will construct the Khovanov flow category as a kind of cover of the flow category for a cube. In the proof of invariance of the Khovanov homotopy type we will consider certain subcategories of the flow category, corresponding to subcomplexes and quotient complexes of the Khovanov chain complex. We introduce these three notions, of a cover, a downward closed subcategory, and an upward closed subcategory of a flow category here.

3.4.1. *Covers.*

Definition 3.29. A grading-preserving functor \mathcal{F} from a flow category $\widehat{\mathcal{C}}$ to a flow category \mathcal{C} is said to be a *cover* if for all $x, y \in \text{Ob}_{\widehat{\mathcal{C}}}$, either $\mathcal{M}_{\widehat{\mathcal{C}}}(x, y) = \emptyset$, or $\mathcal{F} : \mathcal{M}_{\widehat{\mathcal{C}}}(x, y) \rightarrow \mathcal{M}_{\mathcal{C}}(\mathcal{F}(x), \mathcal{F}(y))$ is a $(\text{gr}(x) - \text{gr}(y) - 1)$ -map, a local diffeomorphism, and a covering map.

We say the cover is *trivial* if for all $x, y \in \text{Ob}_{\widehat{\mathcal{C}}}$, the covering map $\mathcal{F} : \mathcal{M}_{\widehat{\mathcal{C}}}(x, y) \rightarrow \mathcal{M}_{\mathcal{C}}(\mathcal{F}(x), \mathcal{F}(y))$, restricted to each component of $\mathcal{M}_{\widehat{\mathcal{C}}}(x, y)$, is a homeomorphism onto its image.

If $\mathcal{F} : \widehat{\mathcal{C}} \rightarrow \mathcal{C}$ is a cover, and if ι is a neat immersion of \mathcal{C} relative some \mathbf{d} , the “composition” $\iota \circ \mathcal{F}$ is a neat immersion of $\widehat{\mathcal{C}}$. Furthermore, a coherent framing for ι induces a coherent framing for $\iota \circ \mathcal{F}$; in such a situation, one can frame $\widehat{\mathcal{C}}$ by a perturbation (Definition 3.22).

Remark 3.30. In a cover \mathcal{F} , the cardinality of the fiber of

$$\mathcal{F} : \mathcal{M}_{\widehat{\mathcal{C}}}(x, y) \rightarrow \mathcal{M}_{\mathcal{C}}(\mathcal{F}(x), \mathcal{F}(y))$$

may depend on x and y . This is even true for trivial covers, which perhaps should be called “topologically trivial covers,” as they can have interesting combinatorics.

3.4.2. *Subcomplexes and quotient complexes.*

Definition 3.31. A full subcategory \mathcal{C}' of a flow category \mathcal{C} is said to be a *downward closed subcategory* (resp. an *upward closed subcategory*) if for all $x, y \in \text{Ob}_{\mathcal{C}}$ with $\mathcal{M}_{\mathcal{C}}(x, y) \neq \emptyset$, $x \in \text{Ob}_{\mathcal{C}'}$ implies $y \in \text{Ob}_{\mathcal{C}'}$ (resp. $y \in \text{Ob}_{\mathcal{C}'}$ implies $x \in \text{Ob}_{\mathcal{C}'}$). It is clear that downward closed subcategories and upward closed subcategories of a flow category are flow categories.

Lemma 3.32. *Let \mathcal{C}' be a downward closed subcategory (resp. an upward closed subcategory) of a flow category \mathcal{C} . Let \mathcal{C} be framed, with (ι, φ) being the neat embedding and the coherent framing. Let (ι', φ') be the induced neat embedding and coherent framing for \mathcal{C}' . Let $|\mathcal{C}'|'_{(\iota', \varphi')}$ be the subcomplex (resp. the quotient complex) of $|\mathcal{C}'|_{(\iota, \varphi)}$ containing exactly the cells that correspond to the objects that are in \mathcal{C}' . Then*

$$|\mathcal{C}'|'_{(\iota', \varphi')} = |\mathcal{C}'|'_{(\iota, \varphi)}.$$

(Here, we assume the integers A, B of Definition 3.24 are chosen to be the same for the two sides.)

Proof. This follows immediately from Definition 3.24. □

Lemma 3.33. *Let \mathcal{C}' and \mathcal{C}'' be downward closed subcategories of a framed flow category \mathcal{C} , such that any object of \mathcal{C} is in exactly one of \mathcal{C}' or \mathcal{C}'' . Assume that \mathcal{C}' and \mathcal{C}'' are framed by the induced framing. Then $|\mathcal{C}| = |\mathcal{C}'| \vee |\mathcal{C}''|$.*

Proof. Let $|\mathcal{C}'|$ (resp. $|\mathcal{C}''|$) be the subcomplex of $|\mathcal{C}|$ containing precisely the cells that correspond to objects that are in \mathcal{C}' (resp. \mathcal{C}''). Since every object of \mathcal{C} is in exactly one of \mathcal{C}' or \mathcal{C}'' ,

$$|\mathcal{C}'| \cup |\mathcal{C}''| = |\mathcal{C}| \text{ and}$$

$$|\mathcal{C}'| \cap |\mathcal{C}''| = \text{pt.}$$

Therefore, $|\mathcal{C}| = |\mathcal{C}'| \vee |\mathcal{C}''|$. However, from Lemma 3.32, $|\mathcal{C}'| = |\mathcal{C}'|$ and $|\mathcal{C}''| = |\mathcal{C}''|$. □

Lemma 3.34. *Let \mathcal{C}' be a downward closed subcategory of a framed flow category \mathcal{C} , and let \mathcal{C}'' be the complementary upward closed subcategory. Assume \mathcal{C}' and \mathcal{C}'' are framed by the induced framings. Let C^* denote the chain complex associated to a framed flow category (Definition 3.21). Then the following hold.*

- (1) *If $C^*(\mathcal{C}'')$ is acyclic then the inclusion $|\mathcal{C}'| \hookrightarrow |\mathcal{C}|$ induces a stable homotopy equivalence.*
- (2) *If $C^*(\mathcal{C}')$ is acyclic then the quotient map $|\mathcal{C}| \rightarrow |\mathcal{C}''|$ induces a stable homotopy equivalence.*
- (3) *If $C^*(\mathcal{C})$ is acyclic then the Puppe map $|\mathcal{C}''| \rightarrow \Sigma|\mathcal{C}'|$ induces a stable homotopy equivalence.*

Proof. By Lemma 3.32, $|\mathcal{C}'|$ is a subcomplex of $|\mathcal{C}|$ and $|\mathcal{C}''|$ is the corresponding quotient complex. Therefore, there is a short exact sequence of reduced cellular cochain complexes

$$0 \rightarrow \tilde{C}^*(|\mathcal{C}''|) \rightarrow \tilde{C}^*(|\mathcal{C}|) \rightarrow \tilde{C}^*(|\mathcal{C}'|) \rightarrow 0,$$

leading to the long exact sequence of reduced cohomology groups

$$\dots \rightarrow \tilde{H}^{i-1}(|\mathcal{C}'|) \rightarrow \tilde{H}^i(|\mathcal{C}''|) \rightarrow \tilde{H}^i(|\mathcal{C}|) \rightarrow \tilde{H}^i(|\mathcal{C}'|) \rightarrow \tilde{H}^{i+1}(|\mathcal{C}''|) \rightarrow \dots$$

In the long exact sequence, the maps $\tilde{H}^*(|\mathcal{C}|) \rightarrow \tilde{H}^*(|\mathcal{C}'|)$, $\tilde{H}^*(|\mathcal{C}''|) \rightarrow \tilde{H}^*(|\mathcal{C}|)$ and $\tilde{H}^{*-1}(|\mathcal{C}'|) \rightarrow \tilde{H}^*(|\mathcal{C}''|)$ are induced by the inclusion $|\mathcal{C}'| \hookrightarrow |\mathcal{C}|$, the quotient map $|\mathcal{C}| \rightarrow |\mathcal{C}''|$ and the Puppe map $|\mathcal{C}''| \rightarrow \Sigma|\mathcal{C}'|$, respectively.

Therefore, in Case (1) if $C^*(\mathcal{C}'')$ is acyclic, Lemma 3.25 implies that $\tilde{H}^*(|\mathcal{C}''|) = 0$; thus the inclusion $|\mathcal{C}'| \hookrightarrow |\mathcal{C}|$ induces an isomorphism at the level of homology, and hence by Whitehead’s theorem is a stable homotopy equivalence. The other two cases can be dealt with similarly. □

4. A FRAMED FLOW CATEGORY FOR THE CUBE

In Section 5, we will construct a framed flow category refining the Khovanov chain complex. Before doing so, however, we construct a simpler framed flow category which will be used in the construction of the category of interest.

4.1. The cube flow category.

Definition 4.1. Let $f_1: \mathbb{R} \rightarrow \mathbb{R}$ be a Morse function with a single index zero critical point at 0 and a single index one critical point at 1. (For concreteness, we may choose $f_1(x) = 3x^2 - 2x^3$.) Let $f_n: \mathbb{R}^n \rightarrow \mathbb{R}$ be the Morse function

$$f_n(x_1, \dots, x_n) = f_1(x_1) + \dots + f_1(x_n).$$

The n -dimensional cube flow category $\mathcal{C}_C(n)$ is the Morse flow category $\mathcal{C}_M(f_n)$ of f_n .

Let $\mathcal{C}(n) = [0, 1]^n$ be the cube with the obvious CW complex structure. The set of vertices of $\mathcal{C}(n)$ is $\{0, 1\}^n$. We will re-use the notations from Definition 3.4. There is a grading function gr on $\{0, 1\}^n$, defined as $\text{gr}(u) = |u| = \sum_i u_i$. Write $v \leq_i u$ if $v \leq u$ and $\text{gr}(u) - \text{gr}(v) = i$.

We will use the following observations about $\mathcal{C}_C(n)$.

Lemma 4.2. *There is a grading preserving correspondence between the objects of $\mathcal{C}_C(n)$ and the vertices of $\mathcal{C}(n)$. Furthermore:*

- (1) *for $u, v \in \{0, 1\}^n$, the space $\mathcal{M}_{\mathcal{C}_C(n)}(u, v)$ is empty unless $v < u$ and*




FIGURE 4.1. The cube moduli space $\mathcal{M}_{\mathcal{C}(3)}(\bar{1}, \bar{0})$.

(2) if $v < u$, then the space $\mathcal{M}_{\mathcal{C}(n)}(u, v)$ is canonically diffeomorphic to the space $\mathcal{M}_{\mathcal{C}(\text{gr}(u) - \text{gr}(v))}(\bar{1}, \bar{0})$.

Proof. This is clear from the form of the Morse function f_n . □

Lemma 4.3. *The moduli space $\mathcal{M}_{\mathcal{C}(n)}(\bar{1}, \bar{0})$ is diffeomorphic to a single point, an interval, and a hexagon, for $n = 1, 2, 3$, respectively.*

In general, the moduli space $\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ is homeomorphic to \mathbb{D}^n (and its boundary $\partial\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ is homeomorphic to S^{n-1}).

Proof. The structure of $\mathcal{M}_{\mathcal{C}(n)}(\bar{1}, \bar{0})$ for small values of n can be deduced from simple model computations; see Figure 4.1 for the case $n = 3$.

For the general case, recall from Definition 3.14 that $\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ is a certain compactification of the quotient of the parametrized moduli space $\widetilde{\mathcal{M}}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ by a natural \mathbb{R} -action. From the form of the Morse function f_{n+1} , it is clear that $\widetilde{\mathcal{M}}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0}) = \prod_{i=1}^{n+1} \widetilde{\mathcal{M}}_{\mathcal{C}(1)}((1), (0))$; and $\widetilde{\mathcal{M}}_{\mathcal{C}(1)}((1), (0)) = \mathbb{R}$.

Therefore, $\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ is the compactification of the configuration space of $(n + 1)$ points a_1, \dots, a_{n+1} (not necessarily distinct) in \mathbb{R} , up to an overall translation in \mathbb{R} . This moduli space naturally breaks up into cubical chambers indexed by permutations of $\{1, \dots, n + 1\}$ as follows: For σ is a permutation of $\{1, \dots, n + 1\}$, consider the subspace of $\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ corresponding to configurations satisfying $a_{\sigma(1)} \leq \dots \leq a_{\sigma(n+1)}$. Since everything is up to an overall translation, such configurations are uniquely specified by the consecutive distances $(a_{\sigma(2)} - a_{\sigma(1)}, a_{\sigma(3)} - a_{\sigma(2)}, \dots, a_{\sigma(n+1)} - a_{\sigma(n)})$, and hence the compactification of such a chamber is the cube $[0, \infty]^n$.

This identifies $\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ with a cubical complex. However, this cubical complex can also be identified with the cubical subdivision of the permutohedron P_{n+1} —see [Blo11, Section 2 and Figure 4]—thus completing the proof. (Although we do not need it here, but it is not hard to see that $\mathcal{M}_{\mathcal{C}(n+1)}(\bar{1}, \bar{0})$ is diffeomorphic to P_{n+1} as $\langle n \rangle$ -manifolds.) □

Given a pair of vertices $v \leq_m u$ of $\mathcal{C}(n)$, let $\mathcal{C}_{u,v} = \{x \in [0, 1]^n \mid \forall i: v_i \leq x_i \leq u_i\}$ be the corresponding m -cell of $\mathcal{C}(n)$. It will be convenient to describe the “faces” of $\mathcal{C}(n)$ as follows.

Definition 4.4. Assume $\mathcal{C}_{u,v}$ is an m -cell of $\mathcal{C}(n)$. There is a corresponding inclusion functor

$$\mathcal{I}_{u,v}: \mathcal{C}_C(m) \hookrightarrow \mathcal{C}_C(n)$$

defined as follows. Suppose $j_1 < j_2 < \dots < j_m$ are the m indices where u and v differ. Given an object $\mathbf{x} \in \{0, 1\}^m$ in $\mathcal{C}_C(m)$, define a new object $\mathbf{x}' \in \{0, 1\}^n$ by setting the i^{th} coordinate x'_i of \mathbf{x}' to be 1 if $v_i = 1$ and 0 if $u_i = 0$. For the rest, define $x'_{j_i} = x_i$.

Then the full subcategory of $\mathcal{C}_C(n)$ with objects $\{\mathbf{x}' \mid \mathbf{x} \in \text{Ob}_{\mathcal{C}_C(m)}\}$ is canonically isomorphic to $\mathcal{C}_C(m)$ (cf. Lemma 4.2), and $\mathcal{I}_{u,v}$ is defined to be this isomorphism.

4.2. Framing the cube flow category. The main goal of this section is to produce a framing of the cube flow category $\mathcal{C}_C(n)$. This is accomplished in Proposition 4.12, by an obstruction theory argument. (An alternate approach would be to use the fact that any Morse flow category is framed.)

Consider $C^*(\mathcal{C}(n), \mathbb{Z})$, the (cellular) cochain complex of $\mathcal{C}(n)$. Let $\mathcal{C}^{u,v}$ denote the cochain that sends the cell $\mathcal{C}_{u,v}$ to 1, and the rest of the cells to 0. The differentials in $C^*(\mathcal{C}(n), \mathbb{F}_2)$ are fairly easy to write down, viz.

$$\delta \mathcal{C}^{u,v} = \sum_{v' \in \{v' \mid v' \leq_1 v\}} \mathcal{C}^{u,v'} + \sum_{u' \in \{u' \mid u \leq_1 u'\}} \mathcal{C}^{u',v}.$$

To define $C^*(\mathcal{C}(n), \mathbb{Z})$, we need to sign-refine the above cochain complex. Towards this end, we will need the standard sign assignment.

Definition 4.5. Let $1_i \in C^i(\mathcal{C}(n), \mathbb{F}_2)$ be the i -cocycle which assigns 1 to each i -cell. A *sign assignment* is a 1-cochain $s \in C^1(\mathcal{C}(n), \mathbb{F}_2)$, such that $\delta s = 1_2$.

Since $H^2(\mathcal{C}(n), \mathbb{F}_2) = 0$, such a sign assignment exists. Since $H^1(\mathcal{C}(n), \mathbb{F}_2) = 0$, any two such sign assignments s, s' are related by a coboundary, i.e., $s - s' = \delta t$ for some $t \in C^0(\mathcal{C}(n), \mathbb{F}_2)$, or in other words, any two sign assignments are *gauge equivalent*.

One often works with the sign assignment s_0 , called the *standard sign assignment*, which assigns to the edge joining $(\epsilon_1, \dots, \epsilon_{i-1}, 0, \epsilon_{i+1}, \dots, \epsilon_n)$ and $(\epsilon_1, \dots, \epsilon_{i-1}, 1, \epsilon_{i+1}, \dots, \epsilon_n)$, the number $(\epsilon_1 + \dots + \epsilon_{i-1}) \pmod{2} \in \mathbb{F}_2$; see also Definition 2.15.

Definition 4.6. Given a sign assignment s for $\mathcal{C}(n)$, the *n -dimensional cube complex* $C_s^*(n)$ is defined as follows. The group $C_s^*(n)$ is freely generated by $\{0, 1\}^n$, the vertices of $\mathcal{C}(n)$. The differential is given by

$$(4.1) \quad \delta v = \sum_{u \in \{u \mid v \leq_1 u\}} (-1)^{s(\mathcal{C}_{u,v})} u.$$

We will frame the cube flow category so that it refines the cube complex $C_s^*(n)$. To get started, we will need to orient the cube moduli spaces. (The orientations we choose here may not be the ones induced from the framing we will produce.) To do so, first orient each cell of $\mathcal{C}(n)$ by the product orientation. It is easy to see that the differential in $C^*(\mathcal{C}(n), \mathbb{Z})$ can be written as

$$(4.2) \quad \delta \mathcal{C}^{u,v} = \sum_{v' \in \{v' \mid v' \leq_1 v\}} (-1)^{s_0(\mathcal{C}_{u-v', u-v})} \mathcal{C}^{u,v'} - \sum_{u' \in \{u' \mid u \leq_1 u'\}} (-1)^{s_0(\mathcal{C}_{u'-v, u-v})} \mathcal{C}^{u',v}$$

where s_0 is the standard sign assignment.

Definition 4.7. A coherent orientation for $\mathcal{C}_C(n)$ is a choice of orientations for all the moduli space $\mathcal{M}(u, v)$, satisfying the following conditions.

- (1) If $u \leq_1 u'$, then the orientation of $\mathcal{M}(u, v)$, viewed as the subspace $\mathcal{M}(u, v) \times \mathcal{M}(u', u) \subset \partial\mathcal{M}(u', v)$, is $-(-1)^{s_0(\mathcal{C}_{u'-v, u-v})}$ times the boundary orientation.
- (2) If $v' \leq_1 v$, then the orientation of $\mathcal{M}(u, v)$, viewed as the subspace $\mathcal{M}(v, v') \times \mathcal{M}(u, v) \subset \partial\mathcal{M}(u, v')$, is $(-1)^{s_0(\mathcal{C}_{u-v', u-v})}$ times the boundary orientation.

Lemma 4.8. *There exists a coherent orientation for $\mathcal{C}_C(n)$.*

Proof. Orient all the cells of $\mathcal{C}(n)$ so that the differential in $C^*(\mathcal{C}(n), \mathbb{Z})$ satisfies (4.2). Recall from Definition 3.14 that the interior of the cell $\mathcal{C}_{u,v}$ is diffeomorphic to $\overset{\circ}{\mathcal{M}}(u, v) \times \mathbb{R}$, where \mathbb{R} is oriented along the flowlines, i.e., the value of the Morse function decreases along \mathbb{R} . Now orient $\overset{\circ}{\mathcal{M}}(u, v)$ so that the product orientation on $\overset{\circ}{\mathcal{M}}(u, v) \times \mathbb{R}$ agrees with the orientation of $\overset{\circ}{\mathcal{C}}_{u,v}$. □

We digress briefly to talk about orthogonal groups. Treat the orthogonal group $O(n)$ as a subgroup of $O(n + 1)$, by identifying \mathbb{R}^n with the subspace $\mathbb{R}^n \times \{0\}$ in \mathbb{R}^{n+1} . Let $O = \text{colim } O(n)$. Since $O(n)$ is a topological group, the action of $\pi_1(O(n), \text{Id})$ on $\pi_i(O(n), \text{Id})$ is trivial; or in other words, we can omit the base-point while talking about the homotopy groups of $O(n)$. For $n > i + 1$, the map $\pi_i(O(n)) \rightarrow \pi_i(O)$ is an isomorphism.

Given an embedding $(\mathbb{D}^k, \partial\mathbb{D}^k) \hookrightarrow (\mathbb{R}^m \times \mathbb{R}_+, \mathbb{R}^m \times \{0\})$, a framing of the normal bundle of $\partial\mathbb{D}^k$ is said to be *null-concordant* if it extends to a framing of the normal bundle of \mathbb{D}^k .

Definition 4.9. Fix a coherent orientation for $\mathcal{C}_C(n)$ and a neat embedding ι of $\mathcal{C}_C(n)$ relative some \mathbf{d} . A coherent framing of all moduli spaces of dimension less than k produces an *obstruction class* $\mathfrak{o} \in C^{k+1}(\mathcal{C}(n), \pi_{k-1}(O))$ as follows.

Consider some $(k + 1)$ -cell $\mathcal{C}_{u,v}$. The boundary $\partial(\iota_{u,v}\mathcal{M}(u, v))$ is an oriented $(k - 1)$ -sphere S^{k-1} embedded in $\partial\mathbb{E}_{\mathbf{d}}[\text{gr}(v) : \text{gr}(u)]$. Furthermore, since $\partial(\iota_{u,v}\mathcal{M}(u, v))$ is identified with products of lower-dimensional moduli spaces, which are already framed, $\partial(\iota_{u,v}\mathcal{M}(u, v))$ is framed by the product framing, which we denote $\varphi_{u,v}$. Choose some null-concordant framing $\varphi_{u,v}^0$ of S^{k-1} . Thus we get the following map from S^{k-1} to $O(d_{\text{gr}(v)} + \dots + d_{\text{gr}(u)-1})$: send a point $x \in S^{k-1}$ to the element of the orthogonal group that maps $\varphi_{u,v}^0(x)$ to $\varphi_{u,v}(x)$. Different choices for $\varphi_{u,v}^0$ produce homotopic maps;⁽ⁱⁱ⁾ therefore, we get a well-defined element in $\pi_{k-1}(O(d_{\text{gr}(v)} + \dots + d_{\text{gr}(u)-1}))$, and hence an element $\mathfrak{o}(\mathcal{C}_{v,w}) \in \pi_{k-1}(O)$. Collecting all the terms, we get a cochain $\mathfrak{o} \in C^{k+1}(\mathcal{C}(n), \pi_{k-1}(O))$.

By definition, \mathfrak{o} vanishes if and only if the framing of the moduli spaces of dimension less than k extends to a framing of the k -dimensional moduli spaces.

Again, note that the framings of the moduli spaces need not be related to the orientations of the moduli spaces. The framings are needed to construct maps from $(k - 1)$ -spheres to O , while the orientations are needed to orient the spheres, and thereby interpret those maps as elements of $\pi_{k-1}(O)$.

Lemma 4.10. *The obstruction class \mathfrak{o} from Definition 4.9 is a cocycle.*

⁽ⁱⁱ⁾Two different null-concordant frames produce a map from S^{k-1} to the orthogonal group. Their null-concordances produce the null-homotopy of this map.

Proof. It is a standard fact from algebraic topology that obstruction classes are cocycles. The proof adapts easily to our setting; we include it here for completeness.

Fix any $(k + 2)$ -cell $\mathcal{C}_{u,v}$. We want to show $\langle \delta \mathfrak{o}, \mathcal{C}_{u,v} \rangle = 0$.

Let S^k be the oriented k -sphere $\partial(\iota_{u,v}\mathcal{M}(u,v))$. Let $W = \{w \mid v \leq_1 w < u \text{ or } v < w \leq_1 u\}$. For $w \in W$, let $\mathbb{D}_w^k \subset S^k$ be the k -cell $\iota_{u,v}(\mathcal{M}(w,v) \times \mathcal{M}(u,w)) = \iota_{w,v}\mathcal{M}(w,v) \times \iota_{u,w}\mathcal{M}(u,w)$, oriented as a subspace of S^k ; let $S_w^{k-1} = \partial\mathbb{D}_w^k$ be the oriented boundary. Let $B^k = S^k \setminus (\bigcup_{w \in W} \mathring{\mathbb{D}}_w^k)$. Given any map $\Phi: B^k \rightarrow O$, let $\Phi_w \in \pi_{k-1}(O)$ be the element induced by $\Phi|_{S_w^{k-1}}$.

We claim that $\sum_{w \in W} \Phi_w = 0$ in $\pi_{k-1}(O)$. To see this, let \mathbb{D}'_w denote a slightly smaller ball concentric with \mathbb{D}_w^k and let $S'_w = \partial\mathbb{D}'_w$. Let $B' = S^k \setminus (\bigcup_{w \in W} \mathring{\mathbb{D}}'_w)$, so B' is the complement of a collection of balls in S^k with disjoint closures. The map Φ extends to a map $\Phi': B' \rightarrow O$. Let p be a point in the interior of B' and let γ_w be an embedded path from p to S'_w in B' (so that γ_w is disjoint from $\gamma_{w'}$ if $w \neq w'$). Then $B' \setminus \bigcup_w \gamma_w$ is homeomorphic to a ball. It follows that $\sum_{w \in W} \Phi|_{S'_w} = 0$ in $\pi_{k-1}(O)$. But $\sum_{w \in W} \Phi_w = \sum_{w \in W} \Phi|_{S'_w}$.

Next, since all moduli spaces of dimension less than k are framed, we get the product framing φ of the bundle $\nu_{\iota_{u,v}}$ restricted to B^k . Choose some framing φ^0 of the whole normal bundle $\nu_{\iota_{u,v}}$. (We do not impose any coherence conditions; since $\mathcal{M}(u,v)$ is a disk, such a framing exists.) We construct our map $\Phi: B^k \rightarrow O$ as follows: a point $x \in B^k$ is mapped to the element of the orthogonal group that sends $\varphi^0(x)$ to $\varphi(x)$. From the earlier observation, $\sum_{w \in W} \Phi_w = 0$.

Now fix some $w \in W$; we will analyze Φ_w . Let φ_w^0 denote the restriction of φ^0 to S_w^{k-1} . We have the following two cases.

- (1) $v \leq_1 w < u$. In this case, $\iota_{u,v}\mathcal{M}(w,v) = \text{pt}$, and the restriction of φ to $S_w^{k-1} = \text{pt} \times \partial(\iota_{u,w}\mathcal{M}(u,w)) \subset \mathbb{R}^{d_{\text{gr}(v)}} \times \partial E_{\mathbf{d}}[\text{gr}(w) : \text{gr}(v)]$ is a product framing, say $\varphi_{w,v} \wedge \varphi_{u,w}$.

Since all of the moduli spaces are coherently oriented, the orientation of $\partial(\iota_{u,w}\mathcal{M}(u,w))$ is $(-1)^{s_0(\mathcal{C}_{u-v,u-w})}$ times the orientation of S_w^{k-1} . Choose some null-concordant framing $\varphi_{u,w}^0$ of $\partial\mathcal{M}(u,w)$; $\mathfrak{o}(\mathcal{C}_{u,w})$ relates $\varphi_{u,w}^0$ to $\varphi_{u,w}$ on $\partial(\iota_{u,w}\mathcal{M}(u,w))$. Therefore, $(-1)^{s_0(\mathcal{C}_{u-v,u-w})}\mathfrak{o}(\mathcal{C}_{u,w})$ relates $\varphi_{w,v} \wedge \varphi_{u,w}^0$ to $\varphi_{w,v} \wedge \varphi_{u,w}$ on S_w^{k-1} .

However, Φ_w relates φ_w^0 to $\varphi_{w,v} \wedge \varphi_{u,w}$ on S_w^{k-1} . Since both φ_w^0 and $\varphi_{w,v} \wedge \varphi_{u,w}^0$ are null-concordant, i.e., they extend to \mathbb{D}_w^k , we have $\Phi_w = (-1)^{s_0(\mathcal{C}_{u-v,u-w})}\mathfrak{o}(\mathcal{C}_{u,w})$.

- (2) $v < w \leq_1 u$. An argument similar to the one in the previous case shows $\Phi_w = -(-1)^{s_0(\mathcal{C}_{u-v,w-v})}\mathfrak{o}(\mathcal{C}_{w,v})$.

Since $\sum_{w \in W} \Phi_w = 0$, we get

$$\sum_{v \leq_1 w < u} (-1)^{s_0(\mathcal{C}_{u-v,u-w})}\mathfrak{o}(\mathcal{C}_{u,w}) - \sum_{v < w \leq_1 u} (-1)^{s_0(\mathcal{C}_{u-v,w-v})}\mathfrak{o}(\mathcal{C}_{w,v}) = 0.$$

Using (4.2) we conclude $\langle \delta \mathfrak{o}, \mathcal{C}_{u,v} \rangle = 0$. □

Lemma 4.11. *Fix a coherent orientation for $\mathcal{C}_C(n)$, a neat embedding ι of $\mathcal{C}_C(n)$ relative some large enough \mathbf{d} , a coherent framing of all moduli spaces of dimension less than k , and a cochain $\mathfrak{p} \in C^k(\mathcal{C}(n), \pi_{k-1}(O))$.*

For each k -cell $\mathcal{C}_{u,v}$, choose a small oriented disk \mathbb{D}^{k-1} in the interior of the oriented manifold $\iota_{u,v}\mathcal{M}(u,v)$, and a map $f_{u,v}: (\mathbb{D}^{k-1}, \partial\mathbb{D}^{k-1}) \rightarrow (O(d_{\text{gr}(v)} + \cdots + d_{\text{gr}(u)-1}), \text{Id})$ that induces $\mathfrak{p}(\mathcal{C}_{u,v}) \in \pi_{k-1}(O)$; then change the framing on the disk \mathbb{D}^{k-1} of the framed manifold $\iota_{u,v}\mathcal{M}(u,v) \subset \mathbb{E}_{\mathbf{d}}[\text{gr}(v) : \text{gr}(u)]$ by $f_{u,v}$.

Let \mathfrak{o} (resp. \mathfrak{o}') be the obstruction class from Definition 4.9 before (resp. after) this modification. Then, $\mathfrak{o}' = \mathfrak{o} + \delta\mathfrak{p}$

Proof. This is clear. The signs work out the same way as in the proof of Lemma 4.10. □

Proposition 4.12. *Given a sign assignment s for $\mathcal{C}(n)$, the cube flow category $\mathcal{C}_{\mathcal{C}}(n)$ can be framed in a way so that it refines the cube complex $C_s^*(n)$.*

Proof. Fix a coherent orientation of $\mathcal{C}_{\mathcal{C}}(n)$ and a neat embedding ι of $\mathcal{C}_{\mathcal{C}}(n)$ relative to some large enough \mathbf{d} . We will produce a coherent framing φ for ι so that the framed flow category $(\mathcal{C}, \iota, \varphi)$ refines the cube complex $C_s^*(n)$.

First, we frame the 0-dimensional moduli spaces. Given $v \leq_1 u$, frame the point $\iota_{u,v}\mathcal{M}(u,v) \in \mathbb{E}_{\mathbf{d}}[\text{gr}(v) : \text{gr}(u)] = \mathbb{R}^{d_{\text{gr}(v)}}$ so that the frame is positive in the tangent space $T\mathbb{R}^{d_{\text{gr}(v)}}$ if and only if $s(\mathcal{C}_{u,v}) = 0$.

Next, we frame the 1-dimensional moduli spaces. Fix $v \leq_2 u$. We are trying to frame the normal bundle of $\iota_{u,v}\mathcal{M}(u,v) \subset \mathbb{E}_{\mathbf{d}}[\text{gr}(v) : \text{gr}(u)]$. Observe that its boundary sphere S^0 is already framed by the product framing. Furthermore, since s is a sign assignment, the framing on S^0 is null-concordant. Therefore, we can extend it to a framing of $\mathcal{M}(u,v)$. Choose some extension (which we might need to change later).

We do the rest inductively. For some $k \geq 2$, assume that we have framed all moduli spaces of dimension less than k . Further assume that the framings of the zero dimensional moduli spaces come from the sign assignment s . After changing the framings in the interior of the $(k-1)$ -dimensional moduli spaces if necessary, we will extend the framings to all k -dimensional moduli spaces.

Definition 4.9 furnishes us with an obstruction class $\mathfrak{o} \in C^{k+1}(\mathcal{C}(n), \pi_{k-1}(O))$ to extending the framing. We know from Lemma 4.10 that \mathfrak{o} is a cocycle. Since $\mathcal{C}(n)$ is acyclic, \mathfrak{o} is a coboundary as well. Choose a cocycle $\mathfrak{p} \in C^k(\mathcal{C}(n), \pi_{k-1}(O))$ such that $\delta\mathfrak{p} = -\mathfrak{o}$. Modify the framings in the interior of the $(k-1)$ -dimensional moduli spaces using the cocycle \mathfrak{p} , as in Lemma 4.11. After the modification, the new obstruction class vanishes, and hence we can extend the framings to the k -dimensional moduli spaces. □

Lemma 4.13. *For any fixed sign assignment s for $\mathcal{C}(n)$, let (ι_0, φ_0) and (ι_1, φ_1) be two framings of $\mathcal{C}_{\mathcal{C}}(n)$ that refine the cube complex $C_s^*(n)$. Then for some $\mathbf{d}_0, \mathbf{d}_1$, $(\iota_0[\mathbf{d}_0], \varphi_0)$ can be connected to $(\iota_1[\mathbf{d}_1], \varphi_1)$ by a smooth 1-parameter family of framings.*

Proof. Choose \mathbf{d}_0 and \mathbf{d}_1 large enough; and then by using Lemma 3.18, for $t \in [0, 1]$, choose a 1-parameter family $\iota(t)$ of neat embeddings of $\mathcal{C}_{\mathcal{C}}(n)$ with $\iota(0) = \iota_0[\mathbf{d}_0]$ and $\iota(1) = \iota_1[\mathbf{d}_1]$. We would like to extend φ_0 and φ_1 to a 1-parameter family of coherent framings φ_t for $\iota(t)$.

We will frame the moduli spaces $\iota_{u,v}(t)(\mathcal{M}(u,v))$ inductively on $|u-v|$. When $v \leq_1 u$, both the points $\iota_{u,v}(0)(\mathcal{M}(u,v))$ and $\iota_{u,v}(1)(\mathcal{M}(u,v))$ are framed with the same sign: positively if $s(\mathcal{C}_{u,v}) = 0$ and negatively if $s(\mathcal{C}_{u,v}) = 1$. Therefore, there

exists a 1-parameter families of framings of $\iota_{u,v}(t)$ connecting the given framings at the endpoints. This takes care of the base case.

Now assume that for some $k \geq 1$, we have produced 1-parameter family of framings φ_t of $\iota_{u,v}(t)$ for all u, v with $|u - v| \leq k$. After changing φ_t relative endpoints for the $(k - 1)$ -dimensional moduli spaces if necessary, we will extend φ_t to all k -dimensional moduli spaces.

The proof is similar to the proof of Proposition 4.12. Fix a coherent orientation of $\mathcal{C}_C(n)$. For any $(k + 1)$ -cell $\mathcal{C}_{u,v}$, consider the k -sphere $S^k = \partial(\mathcal{M}(u, v) \times [0, 1])$, oriented as the boundary of a product. We already have a framing on S^k : the framing on $\mathcal{M}(u, v) \times \{0\}$ is the pullback of the framing φ_0 on $\iota_{u,v}$; the framing on $\mathcal{M}(u, v) \times \{1\}$ is the pullback of the framing φ_1 on $\iota_{u,v}$; and the framing on $\partial\mathcal{M}(u, v) \times [0, 1]$ is the pullback of the product of framings of lower dimensional moduli spaces. Comparing this framing with any other null-concordant framing of S^k produces an element of $\pi_k(O)$. Collecting all these elements, we get a cochain $\mathfrak{o} \in C^{k+1}(\mathcal{C}(n), \pi_k(O))$.

Arguments similar to the ones in the proof of Lemma 4.10 show that this obstruction class \mathfrak{o} is a cocycle, and hence a coboundary. Choose $\mathfrak{p} \in C^k(\mathcal{C}(n), \pi_k(O))$ such that $\delta\mathfrak{p} = -\mathfrak{o}$. For k -cells $\mathcal{C}_{u,v}$, modify the framings for $\mathcal{M}(u, v) \times [0, 1]$ in the interior by $\mathfrak{p}(\mathcal{C}_{u,v})$. After the modification, the new obstruction class vanishes, thus completing the proof. □

5. A FLOW CATEGORY FOR KHOVANOV HOMOLOGY

We return to the notations from Section 2, where D is a link diagram. The goal of this section is to associate a framed flow category to D ; feeding this framed flow category into the Cohen-Jones-Segal construction, as described in Section 3.3, then gives a suspension spectrum.

5.1. Moduli spaces for decorated resolution configurations. The main focus of this section will be to associate to each index n basic decorated resolution configuration (D, x, y) an $(n - 1)$ -dimensional $\langle n - 1 \rangle$ -manifold $\mathcal{M}(D, x, y)$ together with an $(n - 1)$ -map

$$\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}_C(n)}(\bar{1}, \bar{0}).$$

These spaces and maps will be built inductively, satisfying the following properties:

- (RM-1) Let (D, x, y) be a basic decorated resolution configuration, and $(E, z) \in P(D, x, y)$. Let $x|$ and $y|$ denote the induced labelings on $s(E \setminus s(D)) = s(D) \setminus E^{(iii)}$ and $D \setminus E$, respectively; by an abuse of notation, let $z|$ denote the induced labelings on both $s(D \setminus E) = E \setminus D$ and $E \setminus s(D)$. Then there is a *composition map*

$$\circ: \mathcal{M}(D \setminus E, z|, y|) \times \mathcal{M}(E \setminus s(D), x|, z|) \rightarrow \mathcal{M}(D, x, y),$$

⁽ⁱⁱⁱ⁾ $s(D) = s(E)$; therefore, $s(E \setminus s(D)) = s(D) \setminus E$ using Lemma 2.6.

respecting the map \mathcal{F} in the sense that the following diagram commutes:

$$\begin{array}{ccc}
 \mathcal{M}(D \setminus E, z|, y|) \times \mathcal{M}(E \setminus s(D), x|, z|) & \xrightarrow{\circ} & \mathcal{M}(D, x, y) \\
 \mathcal{F} \times \mathcal{F} \downarrow & & \downarrow \mathcal{F} \\
 \mathcal{M}_{\mathcal{E}_C(n-m)}(\bar{1}, \bar{0}) \times \mathcal{M}_{\mathcal{E}_C(m)}(\bar{1}, \bar{0}) & & \\
 \mathcal{I}_{v, \bar{0}} \times \mathcal{I}_{\bar{1}, v} \downarrow & & \downarrow \\
 \mathcal{M}_{\mathcal{E}_C(n)}(v, \bar{0}) \times \mathcal{M}_{\mathcal{E}_C(n)}(\bar{1}, v) & \xrightarrow{\circ} & \mathcal{M}_{\mathcal{E}_C(n)}(\bar{1}, \bar{0}).
 \end{array}
 \tag{5.1}$$

Here, n is the index of D ; m is the index of E ; and the vector $v = (v_1, \dots, v_n)$ has $v_i = 0$ if the i^{th} arc in the totally ordered set $A(D)$ is in $A(E)$, and $v_i = 1$ otherwise. (The maps \mathcal{I} are given in Definition 4.4.)

(RM-2) The faces of $\mathcal{M}(D, x, y)$ are given by

$$\begin{aligned}
 \partial_i \mathcal{M}(D, x, y) &= \partial_{\text{exp}, i} \mathcal{M}(D, x, y) \\
 \tag{5.2} \quad &:= \coprod_{\substack{(E, z) \in P(D, x, y) \\ \text{ind}(D \setminus E) = i}} \circ(\mathcal{M}(D \setminus E, z|, y|) \times \mathcal{M}(E \setminus s(D), x|, z|)).
 \end{aligned}$$

(Here, *exp* stands for “expected.”)

(RM-3) The map \mathcal{F} is a covering map, and a local diffeomorphism. In fact:

(RM-4) The covering map \mathcal{F} is trivial on each component of $\mathcal{M}(D, x, y)$.

Conditions (RM-1)–(RM-3) together imply the analog of Conditions (M-2) and (M-3) of Definition 3.13. In particular, the boundary of $\mathcal{M}(D, x, y)$ is given by a pushout, similarly to Condition (M-3). Let \mathfrak{D} denote the diagram whose vertices are the spaces

$$\mathcal{M}(D \setminus E_k, z_k|, y|) \times \mathcal{M}(E_k \setminus E_{k-1}, z_{k-1}|, z_k|) \times \cdots \times \mathcal{M}(E_1 \setminus s(D), x|, z_1|)$$

where $k \geq 1$ and $(D, y) \prec (E_k, z_k) \prec \cdots \prec (E_1, z_1) \prec (s(D), y)$ is a chain in $P(D, x, y)$; and whose edges correspond to applying the composition map \circ to an adjacent pair of spaces. Let

$$\partial_{\text{exp}} \mathcal{M}(D, x, y) = \text{colim } \mathfrak{D} = \bigcup_i \partial_{\text{exp}, i} \mathcal{M}(D, x, y).$$

Note that $\partial_{\text{exp}} \mathcal{M}(D, x, y)$ is naturally an $\langle n - 1 \rangle$ -boundary (Definition 3.5), whose top-dimensional faces are exactly the $\partial_{\text{exp}, i} \mathcal{M}(D, x, y)$.

Lemma 5.1. $\partial_{\text{exp}} \mathcal{M}(D, x, y) = \partial \mathcal{M}(D, x, y)$.

Proof. This is immediate from (5.2) and the fact that $\mathcal{M}(D, x, y)$ is an $\langle n \rangle$ -manifold. □

In the inductive construction of the $\mathcal{M}(D, x, y)$ we will start by constructing $\mathcal{M}(D, x, y)$ when $\text{ind}(D) = 1$; this moduli space will always be a single point. Next, suppose that $\mathcal{M}(D, x, y)$ has been defined whenever $\text{ind}(D) < k$. We will build the moduli spaces $\mathcal{M}(D, x, y)$ for $\text{ind}(D) = k$ by filling in $\partial_{\text{exp}} \mathcal{M}(D, x, y)$ (which has already been defined) subject to Conditions (RM-1)–(RM-4). The main work will be in checking that the restriction of \mathcal{F} to $\partial_{\text{exp}} \mathcal{M}(D, x, y)$ is a trivial covering.

More formally, we have the following.

Proposition 5.2. *Suppose the moduli space $\mathcal{M}(D, x, y)$ and map $\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}_C(\text{ind}(D))}(\bar{1}, \bar{0})$ has already been constructed for each basic decorated resolution configuration (D, x, y) with $\text{ind}(D) \leq n$, satisfying Conditions (RM-1)–(RM-4). Then, given a decorated resolution configuration (D, x, y) with $\text{ind}(D) = n + 1$:*

(E-1) *The maps \mathcal{F} already defined assemble to give a continuous map*

$$\mathcal{F}|_{\partial}: \partial_{\text{exp}}\mathcal{M}(D, x, y) \rightarrow \partial\mathcal{M}_{\mathcal{C}_C(n+1)}(\bar{1}, \bar{0}).$$

Furthermore, this map respects the $\langle n \rangle$ -boundary structure of the two sides, i.e. $\mathcal{F}|_{\partial}^{-1}(\partial_i\mathcal{M}(\bar{1}, \bar{0})) = \partial_{\text{exp}, i}\mathcal{M}(D, x, y)$ and $\mathcal{F}|_{\partial_{\text{exp}, i}\mathcal{M}(D, x, y)}$ is an n -map.

(E-2) *The map $\mathcal{F}|_{\partial}$ is a covering map.*

(E-3) *There is an n -dimensional $\langle n \rangle$ -manifold $\mathcal{M}(D, x, y)$ and an n -map*

$$\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}_C(n+1)}(\bar{1}, \bar{0})$$

satisfying Conditions (RM-1)–(RM-4) if and only if the map $\mathcal{F}|_{\partial}$ is a trivial covering map.

In particular, if $n \geq 3$ then the space $\mathcal{M}(D, x, y)$ and map \mathcal{F} necessarily exist.

(E-4) *If the space $\mathcal{M}(D, x, y)$ and map $\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}_C(n+1)}(\bar{1}, \bar{0})$ exist, and if $n \geq 2$, then $\mathcal{M}(D, x, y)$ and \mathcal{F} are unique up to diffeomorphism (fixing the boundary).*

Proof. For notational convenience, let $\mathcal{C}_C = \mathcal{C}_C(n + 1)$.

We start with Part (E-1). The fact that $\mathcal{F}|_{\partial}$ is well defined follows from commutativity of diagram (5.1). In more detail, let \mathfrak{E} denote the diagram whose vertices are the spaces

$$\mathcal{M}_{\mathcal{C}_C}(v_k, \bar{0}) \times \mathcal{M}_{\mathcal{C}_C}(v_{k-1}, v_k) \times \cdots \times \mathcal{M}_{\mathcal{C}_C}(\bar{1}, v_1),$$

where $v_i \in \{0, 1\}^{n+1}$, $\bar{0} < v_k < \cdots < v_1 < \bar{1}$, and $k \geq 1$; and whose edges correspond to the composition map \circ for \mathcal{C}_C , applied to an adjacent pair of spaces. Since \mathcal{C}_C is a flow category, it follows from Condition (M-3) of Definition 3.13 that composition induces a homeomorphism

$$\text{colim } \mathfrak{E} \cong \partial\mathcal{M}_{\mathcal{C}_C}(\bar{1}, \bar{0}).$$

The maps \mathcal{F} already defined induce a map of diagrams $\mathcal{G}: \mathfrak{D} \rightarrow \mathfrak{E}$, where \mathcal{G} is given on $\mathcal{M}(D \setminus E_k, z_k, |y|) \times \mathcal{M}(E_k \setminus E_{k-1}, z_{k-1}, |z_k|) \times \cdots \times \mathcal{M}(E_1 \setminus s(D), x, |z_1|)$ by

$$(\mathcal{I}_{v_k, \bar{0}} \times \mathcal{I}_{v_{k-1}, v_k} \times \cdots \times \mathcal{I}_{\bar{1}, v_1}) \circ (\mathcal{F} \times \cdots \times \mathcal{F}),$$

where v_m has i^{th} entry 0 if the i^{th} arc of $A(D)$ is in $A(E_m)$ and 1 otherwise. It follows from diagram (5.1) that \mathcal{G} defines a map of diagrams. Let $\mathcal{F}|_{\partial} = \text{colim } \mathcal{G}$. Since $\partial_i\mathcal{M}(\bar{1}, \bar{0}) = \coprod_{|v|=i} \mathcal{M}(v, \bar{0}) \times \mathcal{M}(\bar{1}, v)$, we have $\partial_{\text{exp}, i}\mathcal{M}(D, x, y) = \mathcal{F}|_{\partial}^{-1}(\partial_i\mathcal{M}(\bar{1}, \bar{0}))$ and the map $\mathcal{F}|_{\partial_{\text{exp}, i}\mathcal{M}(D, x, y)}$ is an n -map by the inductive hypothesis.

Part (E-2) follows from the previous part, Proposition 3.9, and the inductive hypothesis that $\mathcal{F}: \mathcal{M}(D', x', y') \rightarrow \mathcal{M}_{\mathcal{C}_C(\text{ind}(D'))}(\bar{1}, \bar{0})$ is a covering map whenever $\text{ind}(D') \leq n$.

For Part (E-3), recall from Lemma 4.3 that $\mathcal{M}_{\mathcal{C}_C}(\bar{1}, \bar{0})$ is topologically an n -ball, so $\partial\mathcal{M}_{\mathcal{C}_C}(\bar{1}, \bar{0})$ is topologically an $(n - 1)$ -sphere. Hence, by the previous part, $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ is a disjoint union of $(n - 1)$ -spheres. If the covering map $\mathcal{F}|_{\partial}: \coprod_{i=1}^n S^{n-1} \rightarrow S^{n-1}$ is a homeomorphism on each component then it extends in an obvious way to a trivial covering map $\mathcal{F}: \coprod_{i=1}^n \mathbb{D}^n \rightarrow \mathbb{D}^n$. Conversely, if $\mathcal{F}|_{\partial}$

extends then, since any covering space of a disk is trivial, the covering space $\mathcal{F}|_{\partial}$ must have been trivial.

If $\mathcal{F}|_{\partial}$ extends to a covering map \mathcal{F} then requiring \mathcal{F} to be an n -map and a local diffeomorphism defines an $\langle n \rangle$ -manifold structure on $\mathcal{M}(D, x, y)$. Conditions (RM-3) and (RM-4) are vacuously satisfied by the extension \mathcal{F} . The composition map \circ from Condition (RM-1) is induced by the canonical maps

$$\mathcal{M}(D \setminus E, z|, y|) \times \mathcal{M}(E \setminus s(D), x|, z|) \rightarrow \text{colim } \mathfrak{D} = \partial_{\text{exp}} \mathcal{M}(D, x, y) = \partial \mathcal{M}(D, x, y)$$

(since the left-hand side is a vertex in \mathfrak{D}). Since the restriction of \mathcal{F} to $\partial \mathcal{M}(D, x, y)$ is by definition $\mathcal{F}|_{\partial}$, diagram (5.1) commutes.

Note that if $n \geq 3$ then any cover of S^{n-1} is trivial, so the hypothesis is vacuously satisfied.

For Part (E-4), it follows from Lemma 5.1 that for any other $\mathcal{M}'(D, x, y)$ satisfying Conditions (RM-1)–(RM-4), $\partial \mathcal{M}'(D, x, y)$ is identified with $\partial \mathcal{M}(D, x, y)$. So, by the argument in Part (E-3), $\mathcal{M}(D, x, y)$ and $\mathcal{M}'(D, x, y)$ are disjoint unions of disks with the same boundary, and hence diffeomorphic (via a diffeomorphism commuting with \mathcal{F} and fixing the boundary). □

5.2. From resolution moduli spaces to the Khovanov flow category. Suppose we have spaces $\mathcal{M}(D, x, y)$ and maps \mathcal{F} as above, satisfying Conditions (RM-1)–(RM-3). Fix an oriented link diagram L with n crossings, and an ordering of the crossings of L . We will use this data to construct the Khovanov flow category $\mathcal{C}_K(L)$ as a cover of the cube flow category $\mathcal{C}_C(n)$.

Definition 5.3. The *Khovanov flow category* $\mathcal{C}_K(L)$ has one object for each of the standard generators of Khovanov homology, cf. Definition 2.15. That is, an object of $\mathcal{C}_K(L)$ is a labeled resolution configuration of the form $\mathbf{x} = (D_L(u), x)$ with $u \in \{0, 1\}^n$. The grading on the objects is the homological grading gr_h from Definition 2.15; the quantum grading gr_q is an additional grading on the objects. We need the orientation of L in order to define these gradings, but the rest of the construction of $\mathcal{C}_K(L)$ is independent of the orientation.

Consider objects $\mathbf{x} = (D_L(u), x)$ and $\mathbf{y} = (D_L(v), y)$ of $\mathcal{C}_K(L)$. The space $\mathcal{M}_{\mathcal{C}_K(L)}(\mathbf{x}, \mathbf{y})$ is defined to be empty unless $\mathbf{y} \prec \mathbf{x}$ with respect to the partial order from Definition 2.10. So, assume that $\mathbf{y} \prec \mathbf{x}$. Let $x|$ denote the restriction of x to $s(D_L(v) \setminus D_L(u)) = D_L(u) \setminus D_L(v)$ and let $y|$ denote the restriction of y to $D_L(v) \setminus D_L(u)$. Therefore, $(D_L(v) \setminus D_L(u), x|, y|)$ is a basic decorated resolution configuration. Define

$$\mathcal{M}_{\mathcal{C}_K(L)}(\mathbf{x}, \mathbf{y}) = \mathcal{M}(D_L(v) \setminus D_L(u), x|, y|)$$

as smooth manifolds with corners. The composition maps for the resolution configuration moduli spaces (see Condition (RM-1), above) induce composition maps

$$\mathcal{M}_{\mathcal{C}_K(L)}(\mathbf{z}, \mathbf{y}) \times \mathcal{M}_{\mathcal{C}_K(L)}(\mathbf{x}, \mathbf{z}) \rightarrow \mathcal{M}_{\mathcal{C}_K(L)}(\mathbf{x}, \mathbf{y}).$$

The Khovanov flow category $\mathcal{C}_K(L)$ is equipped with a functor \mathcal{F} to $\mathcal{C}_C(n)[-n_-]$, which is a cover in the sense of Definition 3.29. (Henceforth, we will usually suppress the grading shifts from the notation.) On the objects, $\mathcal{F}: \text{Ob}_{\mathcal{C}_K(L)} \rightarrow \text{Ob}_{\mathcal{C}_C(n)}$ is defined as

$$\mathcal{F}(D_L(u), x) = u.$$

For the morphisms,

$$\mathcal{F}: \mathcal{M}_{\mathcal{C}_K(L)}((D_L(u), x), (D_L(v), y)) \rightarrow \mathcal{M}_{\mathcal{C}_C(n)}(u, v)$$

is defined to be composition

$$\mathcal{M}(D_L(v) \setminus D_L(u), x|, y|) \xrightarrow{\mathcal{F}} \mathcal{M}_{\mathcal{C}_C(|u|-|v|)}(\bar{1}, \bar{0}) \xrightarrow{\mathcal{I}_{u,v}} \mathcal{M}_{\mathcal{C}_C(n)}(u, v).$$

Lemma 5.4. *The space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ is empty unless $\text{gr}_q(\mathbf{x}) = \text{gr}_q(\mathbf{y})$ and $\text{gr}_h(\mathbf{y}) < \text{gr}_h(\mathbf{x})$.*

Proof. The moduli space $\mathcal{M}(\mathbf{x}, \mathbf{y})$ is empty unless $\mathbf{y} \prec \mathbf{x}$; however, it is quite clear from Definitions 2.10 and 2.15 that the latter happens only if $\text{gr}_q(\mathbf{x}) = \text{gr}_q(\mathbf{y})$ and $\text{gr}_h(\mathbf{y}) < \text{gr}_h(\mathbf{x})$. \square

Definition 5.5. Fix an oriented link diagram L with n crossings and an ordering of the crossings in L . The Khovanov flow category $\mathcal{C}_K(L)$ (Definition 5.3) is a cover (Definition 3.29) of the cube flow category $\mathcal{C}_C(n)$ (Definition 4.1). Choose a sign assignment s for $\mathcal{C}_C(n)$ (Definition 4.5) and framing of $\mathcal{C}_C(n)$ relative s (Proposition 4.12). The cover \mathcal{F} and the framing of $\mathcal{C}_C(n)$ together produce a neat immersion of $\mathcal{C}_K(L)$ along with a coherent framing for the immersion (Section 3.4.1). Therefore, after choosing a perturbation (Definition 3.22), we get a framing of $\mathcal{C}_K(L)$. The *Khovanov space* is the CW complex $|\mathcal{C}_K(L)|$ (Definition 3.24).

The Khovanov homology (Definition 2.15) is the reduced cohomology of the Khovanov space shifted by $(-C)$ for some positive integer C (Lemma 3.25). The *Khovanov spectrum* $\mathcal{X}_{Kh}(L)$ is the suspension spectrum of the Khovanov space, de-suspended C times.

By Lemmas 5.4 and 3.33, the Khovanov spectrum $\mathcal{X}_{Kh}(L)$ decomposes as a wedge sum over the quantum gradings,

$$\mathcal{X}_{Kh}(L) = \bigvee_j \mathcal{X}_{Kh}^j(L).$$

So, all that remains is to define the moduli spaces associated to resolution configurations; we turn to this next.

5.3. 0-dimensional moduli spaces. In this subsection we define $\mathcal{M}(D, x, y)$ where (D, x, y) is an index 1 basic decorated resolution configuration. Unsurprisingly, we always define $\mathcal{M}(D, x, y)$ to consist of a single point. Since $\mathcal{M}_{\mathcal{C}_C(1)}(\bar{1}, \bar{0})$ also consists of a single point, there is an obvious map

$$\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}_C(1)}(\bar{1}, \bar{0}).$$

So far, Conditions (RM-1)–(RM-4) are vacuously satisfied.

5.4. 1-dimensional moduli spaces. In this subsection we define $\mathcal{M}(D, x, y)$ where (D, x, y) is an index 2 basic decorated resolution configuration. As we will see, this depends on a (global) choice.

Definition 5.6. An index 2 basic resolution configuration D is said to be a *ladybug configuration* if the following conditions are satisfied:

- $Z(D)$ consists of a single circle, which we will abbreviate as Z and
- the endpoints of the two arcs in $A(D)$, say A_1 and A_2 , alternate around Z (i.e., ∂A_1 and ∂A_2 are linked in Z).

See Figure 5.1a. (Ladybug configurations are the same as “configurations of types X and Y ” from [ORS].)

An index 2 basic decorated resolution configuration (D, x, y) is said to be a ladybug configuration if D is a ladybug configuration. For grading reasons, in such

configurations, $y(Z) = x_+$, and the label in x of the (unique) circle in $Z(s(D))$ is x_- .

Now, fix an index 2 basic decorated resolution configuration (D, x, y) . Write $A(D) = \{A_1, A_2\}$, and assume that A_1 precedes A_2 in the total order on $A(D)$. Let $(D_1, z_1), \dots, (D_m, z_m)$ be the labeled resolution configurations such that $(D, y) \prec (D_i, z_i) \prec (s(D), x)$. It is clear that $\partial_{exp}\mathcal{M}(D, x, y)$ consists of m points.

Lemma 5.7. *For any index 2 basic decorated resolution configuration (D, x, y) , the number m of labeled resolution configurations between (D, y) and $(s(D), x)$ is either 2 or 4. Moreover, m equals 4 if and only if (D, x, y) is a ladybug configuration.*

Proof. If D has a leaf or a co-leaf then Lemma 2.14, in conjunction with Lemma 2.13, implies that $P(D, x, y)$ is isomorphic to $\{0, 1\}^2$, and hence $m = 2$.

However, it follows from a straightforward case analysis that the only index 2 basic decorated resolution configuration without leaves or co-leaves is a ladybug configuration, in which case $m = 4$; compare [ORS, Figure 2]. \square

5.4.1. $m = 2$. The map

$$\mathcal{F}|_{\partial}: \partial_{exp}\mathcal{M}(D, x, y) \rightarrow \partial\mathcal{M}_{\mathcal{C}(2)}(\bar{1}, \bar{0})$$

is a trivial covering space, so by Part (E-3) of Proposition 5.2, we can define $\mathcal{M}(D, x, y)$ and $\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}(2)}(\bar{1}, \bar{0})$ satisfying Conditions (RM-1)–(RM-4). More concretely, $\mathcal{M}(D, x, y)$ is defined to be a single interval with boundary $\partial_{exp}\mathcal{M}(D, x, y)$. Assume without loss of generality that D_i contains the arc A_i . Then the diffeomorphism $\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{C}(2)}(\bar{1}, \bar{0})$ is characterized (up to isotopy) by the condition that

$$\mathcal{F}(\mathcal{M}(D \setminus D_1, z_1, y) \times \mathcal{M}(D_1, x, z_1)) = \mathcal{M}_{\mathcal{C}(2)}((0, 1), (0, 0)) \circ \mathcal{M}_{\mathcal{C}(2)}((1, 1), (0, 1))$$

and


$$\mathcal{F}(\mathcal{M}(D \setminus D_2, z_2, y) \times \mathcal{M}(D_2, x, z_2)) = \mathcal{M}_{\mathcal{C}(2)}((1, 0), (0, 0)) \circ \mathcal{M}_{\mathcal{C}(2)}((1, 1), (1, 0)).$$

(This is equivalent to diagram (5.1).) Note that, in this case, $\mathcal{M}(D, x, y)$ and the map \mathcal{F} were uniquely determined, even though Part (E-4) of Proposition 5.2 did not apply.

5.4.2. $m = 4$. We need to decide which pairs of points in $\partial_{exp}\mathcal{M}(D, x, y)$ bound intervals, and for this we need some more notation.


Let Z denote the unique circle in $Z(D)$. The surgery $s_{A_1}(D)$ (resp. $s_{A_2}(D)$) consists of two circles; denote these $Z_{1,1}$ and $Z_{1,2}$ (resp. $Z_{2,1}$ and $Z_{2,2}$); i.e., $Z(s_{A_i}(D)) = \{Z_{i,1}, Z_{i,2}\}$. Our main goal is to find a bijection between $\{Z_{1,1}, Z_{1,2}\}$ and $\{Z_{2,1}, Z_{2,2}\}$; this bijection will then tell us which points in $\partial_{exp}\mathcal{M}(\mathbf{x}, \mathbf{y})$ to identify.

As an intermediate step, we distinguish two of the four arcs in $Z \setminus (\partial A_1 \cup \partial A_2)$. Assume that the point $\infty \in S^2$ is not in D , and view D as lying in the plane $S^2 \setminus \{\infty\} \cong \mathbb{R}^2$. Then one of A_1 or A_2 lies outside Z (in the plane) while the other lies inside Z . Let A_i be the inside arc and A_o the outside arc. The circle Z inherits an orientation from the disk it bounds in \mathbb{R}^2 . With respect to this orientation, each component of $Z \setminus (\partial A_1 \cup \partial A_2)$ either runs from the outside arc A_o to an inside arc A_i or vice-versa. The *right pair* is the pair of components of $Z \setminus (\partial A_1 \cup \partial A_2)$ which run from the outside arc A_o to the inside arc A_i . The other pair of components is the *left pair*; see Figure 5.1.



a: A ladybug configuration. The right pair is marked with small, numbered squares. (The numbering is arbitrary.)

b: The pairs of circles $Z_{1,i}$ and $Z_{2,i}$; the bijection induced by the right pair is indicated.



c: The image of Z under three isotopies in S^2 ; note that the old right pair maps to the new right pair.

FIGURE 5.1. A ladybug configuration.

Lemma 5.8. *Any isotopy of $Z \cup A_1 \cup A_2$ in S^2 takes the right pair to the right pair.*

Proof. See Figure 5.1c. This also follows from the following local description of the right pair: The two arcs in the right pair are precisely the arcs that can be reached by traveling along A_1 or A_2 until Z , and then taking a right turn. \square

The rest of the construction of the flow category is based on choosing either the right pair or the left pair. Both are equally natural choices and we will choose the right one. To keep notation simple, the moduli spaces that we will construct while working with the right pair will be denoted $\mathcal{M}(D, x, y)$. Replacing the right pair by the left pair throughout gives a different set of moduli spaces, which we denote $\mathcal{M}_*(D, x, y)$.

So, let $\{P, Q\}$ be the right pair. After renumbering, we can assume that

$$P \subset Z_{1,1}, \quad Q \subset Z_{1,2}, \quad P \subset Z_{2,1}, \quad Q \subset Z_{2,2}.$$

Thus, P, Q determine a bijection between $\{Z_{1,1}, Z_{1,2}\}$ and $\{Z_{2,1}, Z_{2,2}\}$ by

$$Z_{1,1} \longleftrightarrow Z_{2,1} \quad Z_{1,2} \longleftrightarrow Z_{2,2}.$$

We call this bijection the *ladybug matching*.

The four points in $\mathcal{M}(D, x, y)$ are given by the four maximal chains in $P(D, x, y)$:

$$\begin{aligned} a &= [(D, x_+) \prec (s_{A_1}(D), x_-x_+) \prec (s(D), x_-)] \\ b &= [(D, x_+) \prec (s_{A_1}(D), x_+x_-) \prec (s(D), x_-)] \\ c &= [(D, x_+) \prec (s_{A_2}(D), x_-x_+) \prec (s(D), x_-)] \\ d &= [(D, x_+) \prec (s_{A_2}(D), x_+x_-) \prec (s(D), x_-)] \end{aligned}$$

Here, the expression x_-x_+ in a , say, means that $Z_{1,1}$ is labeled by x_- and $Z_{1,2}$ is labeled by x_+ . So, define $\mathcal{M}(D, x, y)$ to consist of two intervals, one with boundary $a \amalg c$ and the other with boundary $b \amalg d$. Note that this depends on the bijection between $\{Z_{1,1}, Z_{1,2}\}$ and $\{Z_{2,1}, Z_{2,2}\}$, but not the ordering between $Z_{1,1}$ and $Z_{1,2}$.

Next, we verify that we can define the map $\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{G}_C(2)}(\bar{1}, \bar{0})$ compatibly with our choices, in an essentially unique way. The space $\mathcal{M}_{\mathcal{G}_C(2)}(\bar{1}, \bar{0})$ is a single interval. Diagram 5.1 implies that, under the map $\mathcal{F}|_{\partial}$, the points a and b in $\partial\mathcal{M}(D, x, y)$ map to one endpoint of $\mathcal{M}_{\mathcal{G}_C(2)}(\bar{1}, \bar{0})$, while c and d map to the other endpoint. Thus, $\mathcal{F}|_{\partial}$ extends to a covering map $\mathcal{F}: \mathcal{M}(D, x, y) \rightarrow \mathcal{M}_{\mathcal{G}_C(2)}(\bar{1}, \bar{0})$, uniquely up to isotopy. (This is also why we could not simply match a with b , and c with d .)

In summary, we have the following.

Proposition 5.9. *There are spaces $\mathcal{M}(D, x, y)$ for $\text{ind}(D) \leq 2$ and maps \mathcal{F} satisfying Conditions (RM-1)–(RM-4).*

Proof. This follows from Proposition 5.2 and the discussion in Section 5.3 and the rest of this subsection. □

5.5. 2-dimensional moduli spaces. In Section 5.4, we constructed the spaces $\mathcal{M}(D, x, y)$ when $\text{ind}(D) = 2$ or, in other words, $\mathcal{M}(D, x, y)$ is 1-dimensional. The goal of this subsection is to analyze

$$\mathcal{F}|_{\partial}: \partial_{\text{exp}}\mathcal{M}(D, x, y) \rightarrow \partial\mathcal{M}_{\mathcal{G}_C(3)}(\bar{1}, \bar{0})$$

when (D, x, y) is a basic index 3 decorated resolution configuration. By Part (E-2) of Proposition 5.2, we know that $\mathcal{F}|_{\partial}$ is a covering map. By Part (E-3) of Proposition 5.2, in order to construct $\mathcal{M}(D, x, y)$, it suffices to show that the map $\mathcal{F}|_{\partial}$ is a trivial covering map on each component. The spaces $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ and $\partial\mathcal{M}_{\mathcal{G}_C(3)}(\bar{1}, \bar{0})$ are 2-boundaries, so we treat them as graphs. From Lemma 4.3, $\partial\mathcal{M}_{\mathcal{G}_C(3)}(\bar{1}, \bar{0})$ is a 6-cycle. Since the map $\mathcal{F}|_{\partial}$ respects the graph structure, it suffices to show that $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ is a disjoint union of 6-cycles.

The proof that $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ is a disjoint union of 6-cycles is a somewhat involved case analysis. To reduce the number of cases, we will use the notion of dual resolution configurations (Definitions 2.7 and 2.12) as well as the spaces $\mathcal{M}_*(D, x, y)$ constructed analogously to $\mathcal{M}(D, x, y)$ but using the left pair in the ladybug matching (see Section 5.4).

Lemma 5.10. *The graph $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ is a disjoint union of cycles, and the number of vertices in each cycle is divisible by 6. In particular, if $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ has exactly 6 vertices, then it is a 6-cycle.*

Proof. This is an immediate consequence of the fact that $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ is a cover of $\partial\mathcal{M}_{\mathcal{G}_C(3)}(\bar{1}, \bar{0})$, which is a 6-cycle. □

Lemma 5.11. *The graphs $\partial_{\text{exp}}\mathcal{M}(D, x, y)$ and $\partial_{\text{exp}}\mathcal{M}_*(D^*, y^*, x^*)$ are isomorphic.*

Proof. For any decorated resolution configuration (E, u, v) , Lemma 2.13 gives a natural isomorphism between $P(E, u, v)$ and the reverse of $P(E^*, v^*, u^*)$. Let $(f_E(F), w^*) \in P(E^*, v^*, u^*)$ be the labeled resolution configuration corresponding to the labeled resolution configuration $(F, w) \in P(E, u, v)$.

It is easy to see that for any basic index 2 decorated resolution configuration (E, u, v) , the vertices $[(E, v) \prec (F_1, w_1) \prec (s(E), u)]$, and $[(E, v) \prec (F_2, w_2) \prec (s(E), u)]$ are matched (bound an interval) in $\mathcal{M}(E, u, v)$ if and only if corresponding vertices $[(f_E(s(E)), u^*) \prec (f_E(F_1), w_1^*) \prec (f_E(E), v^*)]$ and $[(f_E(s(E)), u^*) \prec (f_E(F_2), w_2^*) \prec (f_E(E), v^*)]$ are matched in the dual $\mathcal{M}_*(E^*, v^*, u^*)$.

Now, in the graphs $\partial_{exp}\mathcal{M}(D, x, y)$ and $\partial_{exp}\mathcal{M}_*(D^*, y^*, x^*)$, the vertices correspond to the maximal chains of $P(D, x, y)$ and $P(D^*, y^*, x^*)$, respectively. Therefore, f_D induces a bijection of the vertices. The edges correspond to the matchings of the maximal chains in index 2 resolution configurations. By the above discussion, f_D induces a bijection of the edges as well. \square

Lemma 5.12. *The graphs $\partial_{exp}\mathcal{M}(D, x, y)$ and $\partial_{exp}\mathcal{M}_*(D, x, y)$ have the same number of vertices and same parity of the number of cycles. In particular, if $\partial_{exp}\mathcal{M}(D, x, y)$ has at most 12 vertices, then the 2 graphs are isomorphic.*

Proof. The vertices in either graph correspond to the maximal chains in $P(D, x, y)$. Therefore, the only thing that we have to prove is that the parity of the number of cycles in the two graphs is the same.

By [ORS, Lemma 2.1], or a straightforward case analysis of configurations, it follows that an even number of faces in the cube of resolutions corresponding to (D, x, y) come from ladybug configurations. (We, in fact, carry out this case analysis: the proofs of Lemmas 5.14 and 5.17 imply that this even number is either 0 or 2.)

Now orient the 6-cycle $\partial\mathcal{M}_{\mathcal{C}(3)}(\bar{1}, \bar{0})$ arbitrarily. This induces orientations on the graphs $\partial_{exp}\mathcal{M}(D, x, y)$ and $\partial_{exp}\mathcal{M}_*(D, x, y)$, and hence we can think of them as oriented 1-manifolds. If a face in the cube of resolutions of (D, x, y) comes from a ladybug configuration, then changing the ladybug matching just for that face corresponds to a single oriented local saddle move for the underlying graphs; see Figure 5.2. Each such oriented saddle move changes the number of components of the graph by 1. Since an even number of faces in the cube of resolutions of (D, x, y) come from ladybug configurations, we are done. \square

Corollary 5.13. *The graph $\partial_{exp}\mathcal{M}(D^*, y^*, x^*)$ is a 6-cycle (resp. a disjoint union of two 6-cycles) if and only if the graph $\partial_{exp}\mathcal{M}(D, x, y)$ is a 6-cycle (resp. a disjoint union of two 6-cycles).*




FIGURE 5.2. An oriented saddle move. Changing the ladybug matching on exactly one face corresponds to a single oriented saddle move for the underlying graphs.

Proof. This is immediate from Lemmas 5.11 and 5.12. □

Lemma 5.14. *Suppose that D has a leaf. Then $\partial_{exp}\mathcal{M}(D, x, y)$ is either a 6-cycle or a disjoint union of two 6-cycles.*

Proof. By Lemma 2.14, there is a naturally associated index 2 decorated resolution configuration (D', x', y') such that $P(D, x, y)$ is naturally isomorphic to $P(D', x', y') \times \{0, 1\}$. If (D', x', y') is not a ladybug configuration, then $P(D', x', y')$ is isomorphic to $\{0, 1\}^2$; in that case, $P(D, x, y)$ has exactly 6 maximal chains, and hence by Lemma 5.10, $\partial_{exp}\mathcal{M}(D, x, y)$ is a 6-cycle.

Now assume that (D', x', y') is a ladybug configuration. An easy enumeration of the possible graphs $G(D)$ (Definition 2.8) tells us that, after reordering the arcs in $A(D)$ if necessary, D is isotopic to one of the three resolution configurations (5.3a)–(5.3c) of Figure 5.3. Furthermore, due to grading reasons, in configurations (5.3a) and (5.3b) (resp. in configuration (5.3c)), the label in y of the leaf (resp. one of the leaves) is x_+ .

Assume the four maximal chains in $P(D', x', y')$ are $d_i = [b \prec c_i \prec a]$ for $1 \leq i \leq 4$. Assume further that the ladybug matching is $d_1 \longleftrightarrow d_2$ and $d_3 \longleftrightarrow d_4$. The 12 vertices in $P(D, x, y) = P(D', x', y') \times \{0, 1\}$ can be represented by the following maximal chains:

$$\begin{aligned} u_i &= [(b, 0) \prec (c_i, 0) \prec (a, 0) \prec (a, 1)] \\ v_i &= [(b, 0) \prec (c_i, 0) \prec (c_i, 1) \prec (a, 1)] \\ w_i &= [(b, 0) \prec (b, 1) \prec (c_i, 1) \prec (a, 1)] \end{aligned}$$

for $1 \leq i \leq 4$. The following pairs are not parts of ladybug configurations, and so the following edges exist independently of our choice of ladybug matching:

$$\begin{array}{cccccc} v_1 - u_1 & v_1 - w_1 & v_2 - u_2 & v_2 - w_2 & v_3 - u_3 & v_3 - w_3 \\ v_4 - u_4 & v_4 - w_4 & & & & \end{array}$$

Since D contains a leaf where the label in y is x_+ , it is clear that if the ladybug matching matches d_i with d_j , then there are edges joining u_i to u_j and w_i to w_j . Therefore, the following edges come from the ladybug matching:

$$u_1 - u_2 \quad u_3 - u_4 \quad w_1 - w_2 \quad w_3 - w_4.$$


So, the components of $\partial_{exp}\mathcal{M}(D, x, y)$ are:

$$v_1 - u_1 - u_2 - v_2 - w_2 - w_1 - v_1 \quad \text{and} \quad v_3 - u_3 - u_4 - v_4 - w_4 - w_3 - v_3. \quad \square$$


Corollary 5.15. *Suppose that D has a co-leaf. Then $\partial_{exp}\mathcal{M}(D, x, y)$ is a disjoint union of 6-cycles.*

Proof. This is immediate from Corollary 5.13 and Lemma 5.14. □

Lemma 5.16. *Up to isotopy in S^2 and reordering of the arcs, the four resolution configurations (5.3d)–(5.3g) of Figure 5.3 are the only ones with no leaves or co-leaves (Definition 2.8). Moreover, configurations (5.3d) and (5.3f) are dual, and configurations (5.3e) and (5.3g) are dual.*



Up to isotopy in S^2 and reordering of the arcs, these are the only basic index 3 resolution configurations with leaves and ladybugs.



Up to isotopy in S^2 and reordering of the arcs, these are the only basic index 3 resolution configurations without leaves or co-leaves. Configurations (5.3f) and (5.3g) are dual to configurations (5.3d) and (5.3e), respectively.

FIGURE 5.3. Some index 3 resolution configurations.

Proof. This again follows from an easy enumeration of the possible graphs $G(D)$. Compare [ORS, Figure 4]. □

Next, we check two cases directly.

Lemma 5.17. *If D is of the type (5.3d), then $\partial_{exp} \mathcal{M}(D, x, y)$ is a 6-cycle; and if D is of the type (5.3e), then $\partial_{exp} \mathcal{M}(D, x, y)$ is a disjoint union of two 6-cycles.*

Proof. For the configuration (5.3d), it is fairly easy to see that $P(D, x, y)$ has exactly 6 maximal chains. In any maximal chain

$$v = [(D, y) \prec (E_2, z_2) \prec (E_1, z_1) \prec (s(D), x)],$$

E_2 is obtained from D by a merge; therefore, the labeling z_2 is determined entirely by D , E_2 and y (cf. Definition 2.10), and hence, there are exactly three choices for (E_2, z_2) . However, it is easy to see that for each of those three choices, $E_2 \setminus s(D)$ is not a ladybug configurations, and hence there are exactly two maximal chains in $P(E_2 \setminus s(D), x, z_2)$. Therefore, $\partial_{exp} \mathcal{M}(D, x, y)$ has exactly 6 vertices and hence, by Lemma 5.10, we are done.

For the configuration (5.3e), for grading reasons, $x = x_- x_-$ and $y = x_+$. The corresponding cube of resolutions is shown in Figure 5.4. The vertices of




FIGURE 5.4. The cube of resolutions for configuration (5.3e).The individual circles at each vertex are numbered.

$\partial_{exp} \mathcal{M}(D, x, y)$ are the maximal chains of $P(D, x, y)$:

- $v_1 = [(D, x_+) \prec (s_{\{A_1\}}(D), x_+x_-) \prec (s_{\{A_1, A_2\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_2 = [(D, x_+) \prec (s_{\{A_1\}}(D), x_+x_-) \prec (s_{\{A_1, A_3\}}(D), x_-x_-x_+) \prec (s(D), x_-x_-)]$
- $v_3 = [(D, x_+) \prec (s_{\{A_1\}}(D), x_-x_+) \prec (s_{\{A_1, A_2\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_4 = [(D, x_+) \prec (s_{\{A_1\}}(D), x_-x_+) \prec (s_{\{A_1, A_3\}}(D), x_-x_+x_-) \prec (s(D), x_-x_-)]$
- $v_5 = [(D, x_+) \prec (s_{\{A_2\}}(D), x_+x_-) \prec (s_{\{A_1, A_2\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_6 = [(D, x_+) \prec (s_{\{A_2\}}(D), x_+x_-) \prec (s_{\{A_2, A_3\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_7 = [(D, x_+) \prec (s_{\{A_2\}}(D), x_-x_+) \prec (s_{\{A_1, A_2\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_8 = [(D, x_+) \prec (s_{\{A_2\}}(D), x_-x_+) \prec (s_{\{A_2, A_3\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_9 = [(D, x_+) \prec (s_{\{A_3\}}(D), x_+x_-) \prec (s_{\{A_2, A_3\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_{10} = [(D, x_+) \prec (s_{\{A_3\}}(D), x_+x_-) \prec (s_{\{A_1, A_3\}}(D), x_-x_+x_-) \prec (s(D), x_-x_-)]$
- $v_{11} = [(D, x_+) \prec (s_{\{A_3\}}(D), x_-x_+) \prec (s_{\{A_2, A_3\}}(D), x_-) \prec (s(D), x_-x_-)]$
- $v_{12} = [(D, x_+) \prec (s_{\{A_3\}}(D), x_-x_+) \prec (s_{\{A_1, A_3\}}(D), x_-x_-x_+) \prec (s(D), x_-x_-)]$.

The faces of the cube with vertices $(000, 100, 010, 110)$ and $(000, 001, 010, 011)$ come from ladybug configurations. The ladybug matchings for these arrangements are:

Face	Circle matching
$(000, 100, 010, 110)$	$1_{100} \longleftrightarrow 2_{010}, 2_{100} \longleftrightarrow 1_{010}$
$(000, 010, 001, 011)$	$1_{010} \longleftrightarrow 1_{001}, 2_{010} \longleftrightarrow 2_{001}$

The following pairs are not parts of ladybug configurations, and so the following edges exist independently of our choice of ladybug matching:

$$v_1 - v_2 \quad v_3 - v_4 \quad v_5 - v_6 \quad v_7 - v_8 \quad v_9 - v_{10} \quad v_{11} - v_{12}$$

$$v_2 - v_{12} \quad v_4 - v_{10}.$$

The following edges do depend on our choice of ladybug matching:

$$v_1 - v_7 \quad v_3 - v_5 \quad v_6 - v_9 \quad v_8 - v_{11}.$$

So, the components of $\partial_{exp} \mathcal{M}(D, x, y)$ are:

$$v_1 - v_2 - v_{12} - v_{11} - v_8 - v_7 - v_1 \quad \text{and} \quad v_3 - v_4 - v_{10} - v_9 - v_6 - v_5 - v_3. \quad \square$$

Corollary 5.18. *In cases (5.3f) and (5.3g) of Figure 5.3, the graph $\partial_{exp} \mathcal{M}(D, x, y)$ is a disjoint union of 6-cycles.*

Proof. This is immediate from Lemma 5.17 and Corollary 5.13. □

In summary, we have the following.

Proposition 5.19. *There are spaces $\mathcal{M}(D, x, y)$ for $\text{ind}(D) \leq 3$ and maps \mathcal{F} satisfying Conditions (RM-1)–(RM-4).*

Proof. This follows from Proposition 5.2, Proposition 5.9, and the discussion in the rest of this subsection. Proposition 5.9 takes care of the case that $\text{ind}(D) \leq 2$. For the case $\text{ind}(D) = 3$, by Proposition 5.2 it suffices to verify that the map $\partial_{exp} \mathcal{M}(D, x, y) \rightarrow \partial \mathcal{M}_{\mathcal{C}(3)}(\overline{1}, \overline{0})$ is a trivial covering map. As discussed in the beginning of the section, this is equivalent to verifying that each component of $\partial_{exp} \mathcal{M}(D, x, y)$ is a 6-cycle. Lemma 5.14 and Corollary 5.15 take care of the cases where D has a leaf or co-leaf. By Lemma 5.16, there are four remaining cases. Lemma 5.17 takes care of two of these four cases and Corollary 5.18 polishes off the last two cases. □

5.6. *n*-dimensional moduli spaces, $n \geq 3$.

Proposition 5.20. *There are spaces $\mathcal{M}(D, x, y)$ and maps \mathcal{F} satisfying Conditions (RM-1)–(RM-4).*

Proof. Proposition 5.19 takes care of the case that $\text{ind}(D) \leq 3$. The case $\text{ind}(D) \geq 4$ follows from Part (E-3) of Proposition 5.2. □

6. INVARIANCE OF THE KHOVANOV SPECTRUM

We devote this section to proving our main theorem, Theorem 1.1. The main work is in proving Part $(\Theta-2)$. In the construction of $\mathcal{X}_{Kh}^j(L)$ we made the following choices:

- (1) A choice of ladybug matching: left or right.
- (2) An oriented link diagram L with n crossings.

- (3) An ordering of the crossings of L .
- (4) A sign assignment s of the cube $\mathcal{C}(n)$.
- (5) A neat embedding ι and a framing φ for the cube flow category $\mathcal{C}_C(n)$ relative s .
- (6) A framed neat embedding of the Khovanov flow category $\mathcal{C}_K(L)$ relative to some \mathbf{d} . This framed neat embedding is a perturbation of (ι, φ) .
- (7) Integers A, B and real numbers ϵ, R used in the construction of the CW complex (Definition 3.24).

We will prove independence of these choices in reverse order.

Proposition 6.1. *Modulo the choice of a ladybug matching, the stable homotopy type of $\mathcal{X}_{Kh}^j(L)$ is an invariant of the link diagram.*

Proof. We must verify independence of choices (3)–(7). Throughout the proof we will suppress the decomposition along quantum gradings, but it will be clear that this decomposition is preserved.

Independence of ϵ and R is Lemma 3.26, and independence of A, B , and \mathbf{d} is Lemma 3.27. (Note that this does introduce some suspensions and canceling desuspensions.) Next, for a given neat embedding ι and framing φ for the cube flow category $\mathcal{C}_C(n)$, if we consider two framed neat embeddings (ι_0, φ_0) and (ι_1, φ_1) of $\mathcal{C}_K(L)$ which are perturbations of (ι, φ) , by Lemma 3.23, after increasing \mathbf{d} if necessary we can find a one-parameter family of neat embeddings and framings (ι_t, φ_t) connecting (ι_0, φ_0) and (ι_1, φ_1) . Then, it again follows from Lemma 3.26 that the Khovanov spaces $|\mathcal{C}_K|_{\iota_0, \varphi_0}$ and $|\mathcal{C}_K|_{\iota_1, \varphi_1}$ are homotopy equivalent. By Lemma 4.13, after increasing \mathbf{d} if necessary, any two choices of framed neat embeddings of the cube flow category $\mathcal{C}_C(n)$ relative to s can be connected by a 1-parameter family of framed neat embeddings, and so lead to homotopy equivalent space $|\mathcal{C}_K|$, as before.

Next, we turn to independence of the sign assignment. Given two different sign assignments s_0 and s_1 for $\mathcal{C}(n)$ there is a sign assignment s for $\mathcal{C}(n+1)$ whose restriction to the codimension-1 face $\{i\} \times \mathcal{C}(n)$ is s_i for $i \in \{0, 1\}$. Consider the link diagram L' obtained by taking the disjoint union of L with a 1-crossing unknot U , drawn so that the 0-resolution of U consists of two circles (and hence the 1-resolution of U consists of a single circle). Order the crossings of L' so that the crossing in U is the first crossing and the other crossings are ordered as in L . Frame $\mathcal{C}(n+1)$ relative to the sign assignment s , and choose a framed neat embedding of $\mathcal{C}_K(L')$ which is a perturbation.

Consider the full subcategory \mathcal{C} of $\mathcal{C}_K(L')$ generated by objects where the (one or two) circles coming from U are labeled x_+ . In $\mathcal{C}_K(L')$, there are no arrows into or out of \mathcal{C} , so $|\mathcal{C}|$ is a summand of $|\mathcal{C}_K(L')|$. Moreover, $|\mathcal{C}_K(L, s_0)|$ is a subcomplex of $|\mathcal{C}|$, and the suspension $\Sigma|\mathcal{C}_K(L, s_1)|$ is the quotient $|\mathcal{C}|/|\mathcal{C}_K(L, s_0)|$. The reduced cohomology of $|\mathcal{C}|$ is clearly trivial (compare Proposition 6.2). Therefore, using Lemma 3.34, $\Sigma|\mathcal{C}_K(L, s_0)| \simeq \Sigma|\mathcal{C}_K(L, s_1)|$, as desired.

Finally, for independence of the ordering of crossings, observe that changing the ordering of the crossings has the effect of composing the functor $\mathcal{F}: \mathcal{C}_K(L) \rightarrow \mathcal{C}_C(n)$ with an automorphism ϕ of $\mathcal{C}_C(n)$. This is equivalent to replacing the sign assignment s and framed neat embedding ι of $\mathcal{C}_C(n)$ with $\phi^*(s) = s \circ \phi$ and $\phi^*(\iota) = \iota \circ \phi$. So, independence of the ordering of the crossings follows from independence of the sign assignment and neat embedding. \square




FIGURE 6.1. The three Reidemeister moves. The orientations of the strands are arbitrary.




FIGURE 6.2. Invariance under the Reidemeister moves RI. The relevant subcomplexes and quotient complexes are indicated.

We now consider the three Reidemeister moves RI, RII, and RIII of Figure 6.1. For unoriented link diagrams, it is enough to consider just these three moves (and their inverses). The other type of Reidemeister I move can be generated by RI and RII and the usual Reidemeister III move can be generated by this braid-like RIII and RII (see [Bal10, Section 7.3]).

For each of RI, RII, and RIII, the strands are oriented in the same way before and after the move, but are otherwise oriented arbitrarily. The orientations of the strands are only needed in order to count the number of positive and negative crossings, which in turn are only needed in the definition of the gradings. For all consistent orientations, in RI, RII, and RIII, (n_+, n_-) increases by $(1, 0)$, $(1, 1)$, and $(3, 3)$, respectively.

Proposition 6.2. *Let L be a link diagram and suppose that L' is obtained from L by performing the Reidemeister move RI. Then $\mathcal{X}_{Kh}^j(L)$ is stably homotopy equivalent to $\mathcal{X}_{Kh}^j(L')$.*

Proof. We will frame the two flow categories appropriately so that $|\mathcal{C}_K(L)|$ will be homotopy equivalent to $|\mathcal{C}_K(L')|$. It will be clear that the bi-grading shifts will work out so that the homotopy equivalence will induce a homotopy equivalence between $\mathcal{X}_{Kh}(L)$ and $\mathcal{X}_{Kh}(L')$, and the latter will respect the decomposition along the quantum gradings.

Following [BN02, Section 3.5.1], there is a contractible subcomplex C_1 of $KC(L')$ so that the corresponding quotient complex $C_2 = KC(L')/C_1$ is isomorphic to $KC(L)$; to develop some notation, we spell this out. The subcomplex C_1 and the quotient complex C_2 are shown in Figures 6.2a and 6.2b, respectively. That is, the 0-resolution of the new crossing has a small circle; denote it U . Then C_1 consists of all 0-resolutions at the new crossing where U is labeled by x_+ , as well as all 1-resolutions at the new crossing (with any labeling of the components); while C_2 is generated by all 0-resolutions at the new crossing where U is labeled by x_- .

It is immediate from the definitions that C_1 corresponds to a subcomplex of $KC(L')$, and that C_2 is isomorphic to $KC(L)$. In fact, C_1 corresponds to an upward closed subcategory \mathcal{C}_1 of $\mathcal{C}_K(L')$. The complementary downward closed subcategory \mathcal{C}_2 (corresponding to C_2) is obviously isomorphic to $\mathcal{C}_K(L)$. Moreover, the chain complex C_1 is acyclic (the horizontal arrow in Figure 6.2a is an isomorphism). So, by Lemma 3.34, for compatible choices of framings, the inclusion map $|\mathcal{C}_K(L)| \hookrightarrow |\mathcal{C}_K(L')|$ is a homotopy equivalence. \square

Proposition 6.3. *Let L be a link diagram and suppose L' is obtained from L by performing the Reidemeister move RII. Then $\mathcal{X}_{Kh}^j(L)$ is stably homotopy equivalent to $\mathcal{X}_{Kh}^j(L')$.*

Proof. As in the proof of RI invariance, we will suppress the bi-grading information. We will once again collapse acyclic subcomplexes and quotient complexes of $KC(L')$ until we are left with $KC(L)$. One sequence of such subcomplexes and quotient complexes is given in [BN02, Section 3.5.2], which we repeat in Figure 6.3 for the reader's convenience. Specifically, the (01)-resolution (of the two new crossings) has a closed circle U . The (01)-resolutions in which U is labeled by x_+ together with all (11)-resolutions generate an acyclic subcomplex C_1 of $KC(L')$, and a corresponding upward closed subcategory \mathcal{C}_1 of $\mathcal{C}_K(L')$. The objects not in C_1 (resp. \mathcal{C}_1) form a quotient complex C_2 (resp. downward-closed subcategory \mathcal{C}_2), whose objects correspond to the (01)-resolutions in which U is labeled by x_- , the (00)-resolutions, and the (10)-resolutions. The complex C_2 has an acyclic quotient complex C_3 (corresponding to a further downward closed subcategory \mathcal{C}_3) consisting of the (00)-resolutions and the remaining (01)-resolutions. Moreover, the complement, C_4 , corresponds exactly to the (10)-resolutions. The corresponding subcategory \mathcal{C}_4 is obviously isomorphic to $\mathcal{C}_K(L)$. So, two applications of Lemma 3.34 give the result. \square

There are two general principles at work in the proof of Proposition 6.3, which we abstract for use in the RIII case. Let v and u be vertices in a (partial) cube




FIGURE 6.3. Invariance under the Reidemeister moves RII. The relevant subcomplexes and quotient complexes are indicated.

of resolutions, such that there is an arrow from v to u , and one of the following holds.

- The arrow from v to u merges a circle U of the (partial) resolution at v . Let S be the set of all generators that correspond to u , and let T be the set of all generators corresponding to v with the circle U labeled by x_+ .
- The arrow from v to u splits off a circle U in the (partial) resolution at u . Let S be the set of all generators that correspond to u with the circle U labeled by x_- , and let T be the set of all generators that correspond to v .

Let C be the chain complex generated by S and T ; it is an acyclic complex, and, therefore, we can delete it without changing the homology. If, in addition, C is a subcomplex or a quotient complex of the original chain complex, then in deleting C , we do not introduce any new boundary maps.

Proposition 6.4. *Let L be a link diagram and suppose L' is obtained from L by performing the braid-like Reidemeister move RIII. Then $\mathcal{X}_{Kh}^j(L)$ is stably homotopy equivalent to $\mathcal{X}_{Kh}^j(L')$.*

Proof. Once again, we suppress the bi-grading information. The relevant part of the cube of resolutions for L' is shown in Figure 6.4. The top half of the figure indicates a subcomplex of $KC(L')$ (corresponding to an upward closed subcategory of $\mathcal{C}_K(L')$), the bottom half a quotient complex of $KC(L')$ (corresponding to a downward closed subcategory of $\mathcal{C}_K(L')$); the homology is carried by the vertex




FIGURE 6.4. Cube of resolutions for the braid-like RIII.

(000111) (corresponding to a subcategory isomorphic to $\mathcal{C}_K(L)$) on the middle-right of the figure. Note that the division essentially corresponds to the homological grading, with some of the vertices in the middle of the homological grading occurring in both the bottom and top halves (with some circle labeled differently). We will show that the top and bottom halves are acyclic. Assuming this, as in the proof of

Proposition 6.3, two applications of Lemma 3.34 show that $|\mathcal{C}_K(L')|$ is homotopy equivalent to $|\mathcal{C}_K(L)|$.

In order to complete the proof, it only remains to see that the two halves are indeed acyclic chain complexes. We demonstrate this by doing the following two sequences of cancellations (the cancellations are similar to the ones in [Bal10, Section 7.3]).

Top half	1*1111, 1*1110, 1*1101, 1*1011, 110*11, 0111*1, *01111, 1*1100, 1*1010, 1*1001, 0111*0, 110*01, *01110, 01101*, *01011, 110*10, 010*11, 1*0011, 01*101, *01101, 1001*1, 10*101.
Bottom half	0000*0, 1000*0, 0100*0, 0010*0, 00*100, 0*0001, 00001*, 0110*0, 1100*0, 1*1000, 1010*0, 01*100, 1*0001, *01001, 00*101, 01*001, 01001*, 010*10, 0*0110, *01100, *01010, 100*10, 1*0100, 10*100.

Here, for instance, 1*1111 means to cancel along the edge from the 101111 resolution to the 111111 resolution (using the principles enunciated above). It is straightforward to verify that each of these cancellations corresponds to taking either a subcomplex or a quotient complex of the image under the previous cancellations; hence, during the cancellations we do not introduce additional maps. \square

Finally, we show that the space $\mathcal{X}_{Kh}^j(L)$ is independent of the choice of ladybug matching. This argument is inspired by the ideas of O. Viro.

Proposition 6.5. *Up to stable homotopy equivalence, $\mathcal{X}_{Kh}^j(L)$ is independent of the ladybug matching.*

Proof. Fix a link diagram L , and let L' be the result of reflecting L across the y -axis, say, and reversing all of the crossings. (In other words, the corresponding links differ by rotation by π around the y -axis.) The Khovanov homologies of L and L' are isomorphic for two different reasons:

- (1) The diagrams L and L' represent the same link, and so are related by a sequence of Reidemeister moves.
- (2) There is a bijection between states for L and L' , obtained by reflecting across the y -axis.

This observation allows us to prove independence of the ladybug matching, as follows. Let \mathcal{C}_K denote the standard Khovanov flow category, constructed using the usual, right pair in the ladybug matching, and let \mathcal{C}_K^* denote the Khovanov flow category constructed using the left pair in the ladybug matching.

We claim that there is an isomorphism of framed flow categories $\mathcal{C}_K(L) \cong \mathcal{C}_K^*(L')$. Indeed, reflecting across the y -axis gives a bijection between objects of $\mathcal{C}_K(L)$ and objects of $\mathcal{C}_K^*(L')$. This bijection respects the Khovanov differential, and hence induces an identification of the index 1 moduli spaces. The reflection exchanges the right and left pairs in each ladybug configuration, and thus induces an identification between index 2 moduli spaces in $\mathcal{C}_K(L)$ and $\mathcal{C}_K^*(L')$. Since the higher index moduli spaces are built inductively from these, it follows that the reflection induces the claimed isomorphism $\mathcal{C}_K(L) \cong \mathcal{C}_K^*(L')$.

So, we have homotopy equivalences

$$|\mathcal{C}_K(L)| \simeq |\mathcal{C}_K(L')| \cong |\mathcal{C}_K^*(L)|.$$

Here, the first stable homotopy equivalence uses Reidemeister invariance of $|\mathcal{C}_K|$, and the second isomorphism uses the claim above. This proves the result. \square

Proof of Theorem 1.1. The proof of Part (Θ-1) is straightforward. To be precise, choose the standard sign assignment s_0 in Definition 5.5; then it follows from Definitions 2.15 and 5.3 that the subcategory of the Khovanov flow category $\mathcal{C}_K(L)$ in quantum grading j refines $KC^{*,j}(L)$. Hence, Lemma 3.25 implies that the reduced cohomology of $\mathcal{X}_{Kh}^j(L)$ is the Khovanov homology in quantum grading j .

Part (Θ-2) follows from the previous propositions. Proposition 6.1 implies that the stable homotopy type of $\mathcal{X}_{Kh}^j(L)$ depends only on the link diagram L , and not on the other choices made in the construction. Propositions 6.2, 6.3, and 6.4 imply that, if L and L' are link diagrams for isotopic links then, for appropriate auxiliary choices, $\mathcal{X}_{Kh}^j(L) \simeq \mathcal{X}_{Kh}^j(L')$. Finally, Proposition 6.5 tells us that this construction is independent of the choice of ladybug matching as well. \square

7. SKEIN SEQUENCE

Let L be a link diagram with a distinguished crossing c , and let L_0 and L_1 be the links obtained by performing the 0- and 1-resolutions at c , respectively. Then for any choice of orientations of L , L_0 , and L_1 there is an exact triangle on Khovanov homology

$$\dots \rightarrow Kh^{i+a,j+b}(L_1) \rightarrow Kh^{i,j}(L) \rightarrow Kh^{i+c,j+d}(L_0) \rightarrow Kh^{i+a+1,j+b}(L_1) \rightarrow \dots$$

Here, a , b , c , and d are integers which depend on how the orientations were chosen.

Theorem 7.1. *With notation as above, there is a cofibration sequence*

$$\Sigma^{-c} \mathcal{X}_{Kh}^{j+d}(L_0) \hookrightarrow \mathcal{X}_{Kh}^j(L) \twoheadrightarrow \Sigma^{-a} \mathcal{X}_{Kh}^{j+b}(L_1).$$

Proof. This is immediate from Lemma 3.32: the flow category $\mathcal{C}_K(L_0)$ is downward closed subcategory of $\mathcal{C}_K(L)$ with complementary upward closed subcategory $\mathcal{C}_K(L_1)$, so $|\mathcal{C}_K(L_0)|$ is a subcomplex of $|\mathcal{C}_K(L)|$ with quotient complex $|\mathcal{C}_K(L_1)|$. \square

8. REDUCED VERSION

In [Kho03, Section 3], Khovanov introduced a reduced version of Khovanov homology, whose Euler characteristic is the ordinary Jones polynomial $V(L)$. To define it, fix a point $p \in L$, not at one of the crossings. For each resolution $D_L(v)$ of L , p lies in some circle $Z_p(v)$. The generators of $KC(L)$ in which $Z_p(v)$ is labeled by x_- span a subcomplex $\widetilde{KC}_-(L)$ of $KC(L)$, and the generators of $KC(L)$ in which $Z_p(v)$ is labeled by x_+ generate the corresponding quotient complex $\widetilde{KC}_+(L)$. Moreover, it is easy to see that the complexes $\widetilde{KC}_-(L)$ and $\widetilde{KC}_+(L)$ are isomorphic; keeping track of gradings, the isomorphism identifies $\widetilde{KC}_-^{i,j-1}(L)$ and $\widetilde{KC}_+^{i,j+1}(L)$. Let $\widetilde{KC}(L)$ denote either of these complexes, graded so that $\widetilde{KC}^{i,j}(L) = \widetilde{KC}_-^{i,j-1}(L)$. Then there are short exact sequences of chain complexes

$$0 \rightarrow \widetilde{KC}^{*,j+1} \rightarrow KC^{*,j} \rightarrow \widetilde{KC}^{*,j-1} \rightarrow 0,$$

which induce long exact sequences

$$(8.1) \quad \dots \rightarrow \widetilde{Kh}^{i,j+1}(L) \rightarrow Kh^{i,j}(L) \rightarrow \widetilde{Kh}^{i,j-1}(L) \rightarrow \widetilde{Kh}^{i+1,j+1}(L) \rightarrow \dots;$$

compare [Ras05, Proposition 4.3].

We can imitate this construction to produce a spectrification of the reduced Khovanov homology. Again, fix $p \in L$, not at one of the crossings. Let $\widetilde{\mathcal{C}}_{K-}(L)$ (resp. $\widetilde{\mathcal{C}}_{K+}(L)$) denote the full subcategory of $\mathcal{C}_K(L)$ whose objects are labeled resolution configurations $(D_L(v), x)$ so that the labeling in x of the circle $Z_p(v)$ is x_- (resp. x_+); this is an upward closed (resp. downward closed) subcategory. As with the Khovanov chain complex, it is straightforward to verify that there is an isomorphism $\widetilde{\mathcal{C}}_{K-}(L) \cong \widetilde{\mathcal{C}}_{K+}(L)$, commuting with the functor $\mathcal{F}: \mathcal{C}_K(L) \rightarrow \mathcal{C}_C(n)$. Let $\widetilde{\mathcal{C}}_K(L)$ denote either of these (isomorphic) categories. Define the gradings on $\widetilde{\mathcal{C}}_K(L)$ to agree with the gradings on $\widetilde{KC}(L)$; that is, the quantum grading is shifted by ∓ 1 from the quantum grading on $\mathcal{C}_{K\pm}(L)$, and the internal (homological) grading is inherited from $\mathcal{C}_K(L)$. Let

$$\widetilde{\mathcal{X}}_{Kh}(L) = \Sigma^{-C} |\widetilde{\mathcal{C}}_K(L)|,$$

where C is chosen so that the cohomology of $\widetilde{\mathcal{X}}_{Kh}(L)$ agrees with the Khovanov homology; compare Definition 5.5.

Theorem 8.1. *The spectrum $\widetilde{\mathcal{X}}_{Kh}(L) = \bigvee_j \widetilde{\mathcal{X}}_{Kh}^j(L)$ satisfies the following properties:*

- (1) *The reduced cohomology of the spectrum $\widetilde{\mathcal{X}}_{Kh}^j(L)$ is the reduced Khovanov homology $\widetilde{Kh}^{*,j}(L)$.*
- (2) *The stable homotopy type of $\widetilde{\mathcal{X}}_{Kh}^j(L)$ depends only on the isotopy class of L , and not on any of the auxiliary choices made in the construction.*
- (3) *There is a cofibration $\widetilde{\mathcal{X}}_{Kh}^{j-1}(L) \rightarrow \mathcal{X}_{Kh}^j(L)$ so that $\mathcal{X}_{Kh}^j(L)/\widetilde{\mathcal{X}}_{Kh}^{j-1}(L) \simeq \widetilde{\mathcal{X}}_{Kh}^{j+1}(L)$. These maps induce the long exact sequence in equation (8.1).*

Proof. Parts (1) and (3) are immediate from the definitions; for Part (3) see also Section 3.4.2. For Part (2), observe that for Reidemeister moves not crossing the marked point p , the stable homotopy equivalences in Section 6 preserve the subcomplex $|\widetilde{\mathcal{C}}_{K+}(L)| \subset |\widetilde{\mathcal{C}}_K(L)|$. Any two diagrams for isotopic pointed links can be connected by Reidemeister moves not crossing the basepoint p and isotopies of the diagrams in S^2 (but possibly passing over ∞): passing a strand over (or under) p can be exchanged for taking the strand around S^2 the other way (compare [Kho03, Section 3]). □

Remark 8.2. It is not the case that $\mathcal{X}_{Kh}^j(L) \cong \widetilde{\mathcal{X}}_{Kh}^{j-1}(L) \vee \widetilde{\mathcal{X}}_{Kh}^{j+1}(L)$; this can be seen already in the case of the trefoil (see Example 9.4), and follows from the fact that $Kh^{i,j}(L) \not\cong \widetilde{Kh}^{i,j-1}(L) \oplus \widetilde{Kh}^{i,j+1}(L)$.

9. EXAMPLES

In this section, we start by computing the Khovanov spectra associated with a one-crossing unknot and the Hopf link. After the explicit examples, we observe that the Khovanov spectrum for any alternating link contains no more information than the Khovanov homology of the link.




FIGURE 9.1. A 1-crossing unknot U and its resolutions.

9.1. **An unknot.** Consider the 1-crossing unknot U shown in Figure 9.1. For convenience, name the generators of the Khovanov complex (or equivalently the labeled resolution configurations, or equivalently the objects of the Khovanov flow category $\mathcal{C}_K(U)$, or equivalently the cells of the Khovanov space $|\mathcal{C}_K(U)|$) as follows:

Name	Generator	gr_h	gr_q	Name	Generator	gr_h	gr_q
a	$(D_U(1), x_+x_+)$	0	1	d	$(D_U(1), x_-x_-)$	0	-3
b	$(D_U(1), x_+x_-)$	0	-1	x	$(D_U(0), x_+)$	-1	-1
c	$(D_U(1), x_-x_+)$	0	-1	y	$(D_U(0), x_-)$	-1	-3

The only \prec -comparable resolution configurations (Definition 2.10) are:

$$\mathbf{x} \prec \mathbf{b} \quad \mathbf{x} \prec \mathbf{c} \quad \mathbf{y} \prec \mathbf{d}.$$

The moduli spaces $\mathcal{M}(\mathbf{b}, \mathbf{x})$, $\mathcal{M}(\mathbf{c}, \mathbf{x})$, and $\mathcal{M}(\mathbf{d}, \mathbf{y})$ each consist of a single, positively framed point. (The framing comes from the standard sign assignment, as described in Section 4.2).

To choose an embedding of this flow category as in Definition 3.16, choose $d_{-1} = 1$ and $d_i = 0$ for $i \neq -1$. Then, embed the three points $\mathcal{M}(\mathbf{b}, \mathbf{x}) \amalg \mathcal{M}(\mathbf{c}, \mathbf{x}) \amalg \mathcal{M}(\mathbf{d}, \mathbf{y})$ in $\mathbb{E}_{\mathbf{d}}[-1 : 0] = \mathbb{R}^1$; call the resulting real numbers b, c , and d . These points inherit the standard, positive framing; see Figure 9.2.

Next, we turn to the construction of $|\mathcal{C}_K|$ from the framed flow category (as in Definition 3.24). Take $B = -1$ and $A = 0$. The cells $\mathcal{C}(\mathbf{x})$ and $\mathcal{C}(\mathbf{y})$ are given by

$$(\{0\} \times [-\epsilon, \epsilon]^{d_{-1}=1}) / \partial(\{0\} \times [-\epsilon, \epsilon]^{d_{-1}=1}) \cong S^1.$$

So, the 1-skeleton of $|\mathcal{C}_K|$ is $S^1 \vee S^1$. The cell corresponding to \mathbf{b} is

$$\mathcal{C}(\mathbf{b}) = [0, R] \times [-R, R]^{d_{-1}=1} / \sim,$$


where \sim identifies $\{0\} \times [b - \epsilon, b + \epsilon]$ with $\mathcal{C}(\mathbf{x})$ and the rest of the boundary with the basepoint. So, $\mathcal{C}(\mathbf{b})$ consists of a 2-cell whose boundary is attached by a degree-one map to the first S^1 -summand of $S^1 \vee S^1$. The story for $\mathcal{C}(\mathbf{c})$ is exactly the same. The cell $\mathcal{C}(\mathbf{d})$ consists of a 2-cell whose boundary is attached by a degree-one map to the second S^1 -summand of $S^1 \vee S^1$. Finally, all of the boundary of $\mathcal{C}(\mathbf{a})$ is collapsed to the basepoint. So,

$$|\mathcal{C}_K| = S^2 \vee S^2 \vee \mathbb{D}^2,$$

where the first S^2 is $\mathcal{C}(\mathbf{b}) \cup_{\mathcal{C}(\mathbf{x})} \mathcal{C}(\mathbf{c})$; the second one is $\mathcal{C}(\mathbf{a})$; and the \mathbb{D}^2 is $\mathcal{C}(\mathbf{d})$.

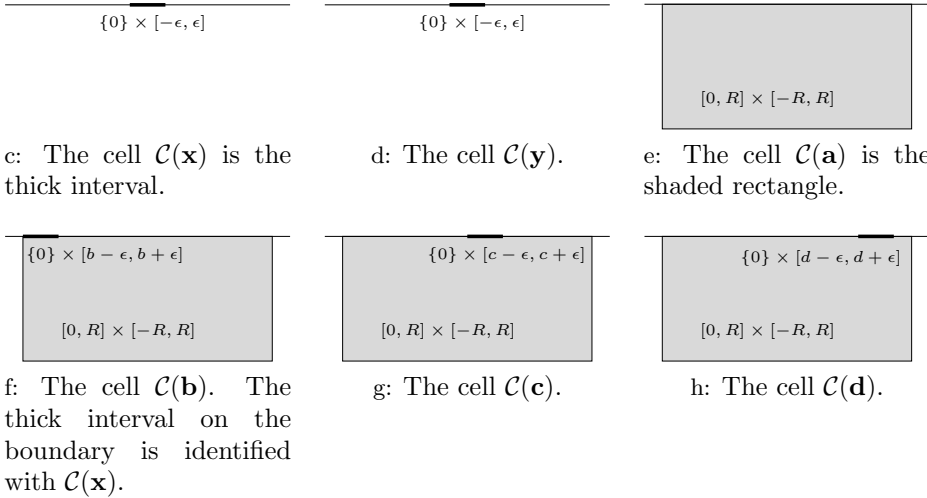
Hence,

$$\begin{aligned} \mathcal{X}_{Kh} &\simeq \Sigma^{-2}(S^2 \vee S^2) \\ &= S_{-1}^0 \vee S_1^0. \end{aligned}$$



a: The Khovanov flow category $\mathcal{C}_K(U)$.

b: A framing of $\mathcal{C}_K(U)$.



c: The cell $\mathcal{C}(\mathbf{x})$ is the thick interval.

d: The cell $\mathcal{C}(\mathbf{y})$.

e: The cell $\mathcal{C}(\mathbf{a})$ is the shaded rectangle.

f: The cell $\mathcal{C}(\mathbf{b})$. The thick interval on the boundary is identified with $\mathcal{C}(\mathbf{x})$.

g: The cell $\mathcal{C}(\mathbf{c})$.

h: The cell $\mathcal{C}(\mathbf{d})$.

FIGURE 9.2. Constructing the Khovanov space for the 1-crossing unknot U of Figure 9.1a.

The first S^0 summands lies in quantum grading -1 and the second in quantum grading 1 . We have indicated this by the subscripts.

9.2. **The Hopf link.** We turn next to the Hopf link H , shown in Figure 9.3. As before, for convenience we name the generators:

Name	Generator	gr_h	gr_q	Name	Generator	gr_h	gr_q
a	$(D_H(00), x_+x_+)$	0	4	c	$(D_H(01), x_+)$	1	4
v	$(D_H(00), x_+x_-)$	0	2	y	$(D_H(01), x_-)$	1	2
w	$(D_H(00), x_-x_+)$	0	2	e	$(D_H(11), x_-x_+)$	2	4
m	$(D_H(00), x_-x_-)$	0	0	z	$(D_H(11), x_-x_-)$	2	2
b	$(D_H(10), x_+)$	1	4	l	$(D_H(11), x_+x_+)$	2	6
x	$(D_H(10), x_+)$	1	2	d	$(D_H(11), x_+x_-)$	2	4

To keep the notation simple, we will discuss each quantum grading separately. In quantum grading 0 (resp. 6) there is a unique generator, and hence the Khovanov spectrum is a sphere in this grading, of virtual dimension 0 (resp. 2) after de-suspending.




FIGURE 9.3. The Hopf link H and its resolutions.

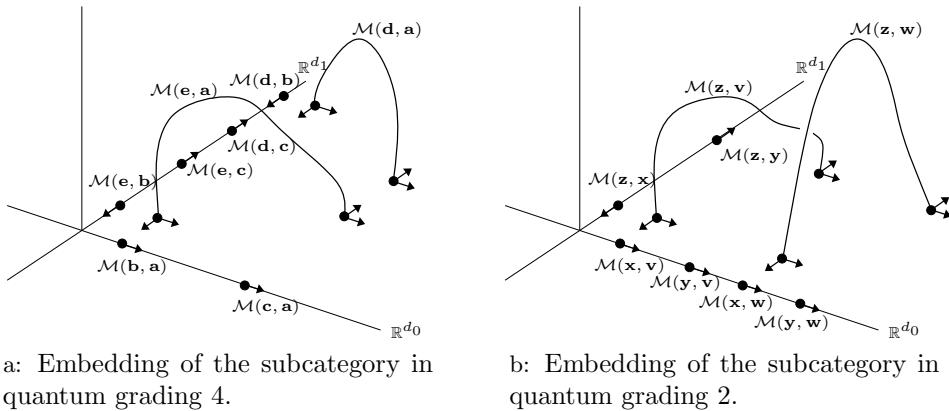



FIGURE 9.4. Embedding of the flow category for the Hopf link.

In quantum gradings 2 and 4 the Khovanov chain complex is given by



respectively.

Consider quantum grading 4. Embed this flow category relative \mathbf{d} where $d_0 = 1$, $d_1 = 1$, and $d_i = 0$ for $i \notin \{0, 1\}$, as shown in Figure 9.4a. Choose $B = 0$ and $A = 2$. The cells of this part of the flow category are given by

$$\begin{aligned}
 \mathcal{C}(\mathbf{a}) &= \{0\} \times [-\epsilon, \epsilon] \times \{0\} \times [-\epsilon, \epsilon] / \sim_a \\
 \mathcal{C}(\mathbf{b}) &= [0, R] \times [-R, R] \times \{0\} \times [-\epsilon, \epsilon] / \sim_b \\
 \mathcal{C}(\mathbf{c}) &= [0, R] \times [-R, R] \times \{0\} \times [-\epsilon, \epsilon] / \sim_c \\
 \mathcal{C}(\mathbf{d}) &= [0, R] \times [-R, R] \times [0, R] \times [-R, R] / \sim_d \\
 \mathcal{C}(\mathbf{e}) &= [0, R] \times [-R, R] \times [0, R] \times [-R, R] / \sim_e .
 \end{aligned}$$

The identification \sim_a identifies all of the boundary to the basepoint. Each of the 3-cells $\mathcal{C}(\mathbf{b})$ and $\mathcal{C}(\mathbf{c})$ has part of its boundary identified to $\mathcal{C}(\mathbf{a})$, by a degree-one map, and the rest collapsed to the basepoint; the picture for this identification is exactly as in Figure 3.3d. The result of this attaching is two 3-balls with their boundaries glued together, i.e., a 3-sphere.

Each of the 4-cells $\mathcal{C}(\mathbf{d})$ and $\mathcal{C}(\mathbf{e})$ has part of its boundary attached to $\mathcal{C}(\mathbf{b})$ (resp. $\mathcal{C}(\mathbf{c})$) via a degree ± 1 map (according to the signs in the Khovanov complex), part attached to $\mathcal{C}(\mathbf{a})$, and the rest collapsed to the basepoint; the picture for this identification is like Figure 3.4d but without the furthest back box (the one corresponding to $\mathcal{C}(z_4)$). Up to homotopy, one can shrink the thickened arc corresponding to $\mathcal{C}(\mathbf{a})$ to become a 2-dimensional rectangle, leaving the boxes corresponding to $\mathcal{C}(\mathbf{b})$ and $\mathcal{C}(\mathbf{c})$ glued together along a rectangle on each of their boundaries. This rectangle is the S^2 along which the 3-balls $\mathcal{C}(\mathbf{b})$ and $\mathcal{C}(\mathbf{c})$ are glued. So, gluing on $\mathcal{C}(\mathbf{d})$ and $\mathcal{C}(\mathbf{e})$ gives two 4-balls glued along their boundary. Thus, in quantum grading 4, the Khovanov space $|\mathcal{C}_K(H)|$ is homotopy equivalent to S^4 . To obtain the Khovanov spectrum, we de-suspend twice, so $\mathcal{X}_{Kh}^4 \simeq S^2$.

Finally, we turn to quantum grading 2. Each of $\mathcal{C}(\mathbf{v})$ and $\mathcal{C}(\mathbf{w})$ are 2-cells; $\mathcal{C}(\mathbf{x})$ and $\mathcal{C}(\mathbf{y})$ are 3-cells; and $\mathcal{C}(\mathbf{z})$ is a 4-cell. In this quantum grading, the 2-skeleton of the Khovanov space $|\mathcal{C}_K|$ is $S^2 \vee S^2$. The attaching map of $\mathcal{C}(\mathbf{x})$ (resp. $\mathcal{C}(\mathbf{y})$) is degree-one onto each S^2 . (Each of $\mathcal{C}(\mathbf{x})$ and $\mathcal{C}(\mathbf{y})$ is a box with two rectangles on top, one identified with $\mathcal{C}(\mathbf{v})$ and the other with $\mathcal{C}(\mathbf{w})$; compare Figure 9.4b.) The result is the suspension of the 2-sphere with the north and south poles identified, and is homotopy equivalent to $S^3 \vee S^2$. The cell $\mathcal{C}(\mathbf{z})$ is attached by a map which is degree -1 to $\mathcal{C}(\mathbf{x})$ and degree 1 to $\mathcal{C}(\mathbf{y})$. To understand this map better, consider the embedding of the moduli spaces from Figure 9.4b. This gives an embedding of $\mathcal{C}(\mathbf{x}) \cup \mathcal{C}(\mathbf{y}) \cup ([0, 1] \times \mathcal{C}(\mathbf{v})) \cup ([0, 1] \times \mathcal{C}(\mathbf{w}))$ in $\partial\mathcal{C}(\mathbf{z}) = S^3$. As in the previous paragraph, collapse the $[0, 1]$ factors to points. We are left with two 3-balls, corresponding to $\mathcal{C}(\mathbf{x})$ and $\mathcal{C}(\mathbf{y})$, glued together along two disks in their boundaries. The complement of these two balls in S^3 is an unknotted solid torus which is quotiented to the basepoint. Therefore, the result of attaching $\mathcal{C}(\mathbf{z})$ according to this embedding is \mathbb{D}^4 with a solid torus on its boundary quotiented to the basepoint. Thus, in quantum grading 4, the Khovanov space $|\mathcal{C}_K|$ is homotopy equivalent to S^2 . To obtain the Khovanov spectrum, we again de-suspend twice, so $\mathcal{X}_{Kh}^2 \simeq S^0$.

In summary,

$$\mathcal{X}_{Kh}(H) = S_0^0 \vee S_2^0 \vee S_4^2 \vee S_6^2.$$

9.3. The Khovanov spectrum of alternating links. Generalizing the previous two examples, we show that for quasi-alternating links the Khovanov spectrum is somewhat uninteresting. We start by collecting some facts about Khovanov homology.

Lemma 9.1. *If L is a quasi-alternating link with signature^(iv) σ then:*

- (1) *The reduced Khovanov homology $\widetilde{Kh}^{i,j}(L)$ is supported on the single diagonal $2i - j = \sigma(K)$.*
- (2) *The reduced Khovanov homology is torsion free.*

^(iv)We use the convention in which a positive knot, such as the right-handed trefoil $T_{2,3}$, has negative signature.

- (3) The un-reduced Khovanov homology is supported on the two diagonals $2i - j = \sigma \pm 1$.
- (4) All torsion in the un-reduced Khovanov homology lies on the diagonal $2i - j = \sigma + 1$.

Proof. Part (1) is exactly [MO08, Theorem 1.1] (Note that their j is one half of our j .); see also [Lee05] in the alternating case.

Part (2) follows from the fact that [MO08, Theorem 1.1] holds over the ring \mathbb{Z}/p and the universal coefficient theorem. (This is not explicitly stated in [MO08], but is clear from their proof.)

To prove Part (3), it follows from the exact sequence in (8.1) and Part (1) that the only cases in which $Kh^{i,j}(L)$ is nonzero are when $j = 2i + \sigma \pm 1$. Further,

$$\begin{aligned}
 Kh^{i,2i-\sigma-1}(L) &= \text{coker}(\widetilde{Kh}^{i-1,2i-\sigma-2}(L) \rightarrow \widetilde{Kh}^{i,2i-\sigma}(L)) \\
 Kh^{i,2i-\sigma+1}(L) &= \text{ker}(\widetilde{Kh}^{i,2i-\sigma}(L) \rightarrow \widetilde{Kh}^{i+1,2i-\sigma+2}(L)).
 \end{aligned}$$

Combined with Part (2), this gives Part (4). □

Proposition 9.2. *Let L be a quasi-alternating link. Then $\mathcal{X}_{Kh}(L)$ and $\widetilde{\mathcal{X}}_{Kh}(L)$ are wedge sums of Moore spaces.*

Proof. We start with the reduced Khovanov homology. Since $\widetilde{\mathcal{X}}_{Kh}(L)$ is a wedge-sum over quantum gradings, it suffices to show that in each quantum grading j , $\widetilde{\mathcal{X}}_{Kh}^j(L)$ is a Moore space. But by Lemma 9.1, the reduced cohomology $\widetilde{H}^*(\widetilde{\mathcal{X}}_{Kh}^j(L))$ is nontrivial in a single grading, and so $\widetilde{\mathcal{X}}_{Kh}(L)$ is by definition a Moore space (and is, in fact, a wedge of spheres).

Next, we turn to the unreduced story. Again, fix a quantum grading j . By Lemma 9.1, the reduced cohomology of $\mathcal{X}_{Kh}^j(L)$ is nontrivial in gradings $(j-1+\sigma)/2$ and $(j+1+\sigma)/2$, and the torsion lies in grading $(j+1+\sigma)/2$. For notational convenience, let $n = (j-1+\sigma)/2$. By the universal coefficient theorem, the reduced homology of $\mathcal{X}_{Kh}^j(L)$ is nontrivial in gradings n and $n+1$, and the torsion lies in grading n . Let $G = H_n(\mathcal{X}_{Kh}^j(L))$ and $K = H_{n+1}(\mathcal{X}_{Kh}^j(L))$ (so K is a free abelian group). Since the homology of \mathcal{X}_{Kh} vanishes in degrees less than n , there is a map f from the Moore space $M(G, n)$ to \mathcal{X}_{Kh} inducing an isomorphism on H_n . Moreover, the map $\pi_{n+1}(\mathcal{X}_{Kh}^j(L)) \rightarrow H_{n+1}(\mathcal{X}_{Kh}^j(L))$ is surjective (see, for instance, [Swi75, Theorem 10.25] or [Hat02, Exercise 4.2.23]). So, there is a map $S^{n+1} \vee \dots \vee S^{n+1} = M(K, n+1) \rightarrow \mathcal{X}_{Kh}^j(L)$ inducing an isomorphism on H_{n+1} . Thus, we have obtained a map $M(G, n) \vee M(K, n+1) \rightarrow \mathcal{X}_{Kh}^j(L)$ inducing an isomorphism on homology; by Whitehead’s theorem, this map is a homotopy equivalence. □

Corollary 9.3. *If L is an alternating link then the Khovanov spectrum (reduced or un-reduced) is determined by its Khovanov homology. In particular, the reduced Khovanov spectrum of an alternating link L is determined by the Jones polynomial and signature of L .*

Example 9.4. The Khovanov homology of the left-handed trefoil 3_1 is given by $Kh^{-3,-9} = Kh^{-2,-5} = Kh^{0,-3} = Kh^{0,-1} = \mathbb{Z}$ and $Kh^{-2,-7} = \mathbb{Z}/2$; and the reduced Khovanov homology is given by $\widetilde{Kh}^{-3,-8} = \widetilde{Kh}^{-2,-6} = \widetilde{Kh}^{0,-2} = \mathbb{Z}$. Thus,

$$\mathcal{X}_{Kh}(3_1) = \Sigma^{-3}S^0_{-9} \vee \Sigma^{-2}S^0_{-5} \vee S^0_{-3} \vee S^0_{-1} \vee \Sigma^{-4}\mathbb{R}P^2_{-7}$$

and

$$\tilde{\mathcal{X}}_{Kh}(3_1) = \Sigma^{-3}S_{-8}^0 \vee \Sigma^{-2}S_{-6}^0 \vee S_{-2}^0,$$

where the subscripts denote the quantum gradings.

Remark 9.5. The proof of Proposition 9.2 was completely independent of our construction of \mathcal{X}_{Kh} ; the corresponding result holds for any Khovanov stable homotopy type of this form (whatever its provenance).

Remark 9.6. We show in [LSb] that for many nonalternating knots, such as $T_{3,4}$, the Khovanov space is not a wedge sum of Moore spaces.

10. SOME STRUCTURAL CONJECTURES

We conclude with a few structural conjectures about the behavior of the Khovanov homotopy type under some basic operations on knots.

10.1. Mirrors. Given a link L , let $m(L)$ denote the mirror of L .

Given a spectrum X , let X^\vee denote the Spanier-Whitehead dual (or S -dual) of X . That is, X^\vee is the dual of X as a module over the sphere spectrum. Equivalently, if X is the suspension spectrum of a finite CW complex (which we also denote X) then X^\vee can be obtained by choosing an embedding $i: X \hookrightarrow S^N$ (for some large N) and formally desuspending the complement $S^N \setminus i(X)$ ($N - 1$) times:

$$X^\vee = \Sigma^{1-N}(S^N \setminus i(X));$$

see [Spa59].

Conjecture 10.1. *Let L be a link. Then*

$$\mathcal{X}_{Kh}^j(m(L)) \simeq \mathcal{X}_{Kh}^{-j}(L)^\vee.$$

Remark 10.2. The analogous result for Manolescu’s Seiberg-Witten stable homotopy type, that $\text{SWF}(-Y)$ is Spanier-Whitehead dual to $\text{SWF}(Y)$, is observed in [Man03, Remark 2], using the corresponding result about Conley indices [Cor00].

10.2. Disjoint unions and connected sums. For Khovanov homology, disjoint unions correspond to tensor products. The obvious analog at the space level is the smash product.

Conjecture 10.3. *Let L_1 and L_2 be links, and $L_1 \amalg L_2$ their disjoint union. Then*

$$\mathcal{X}_{Kh}^j(L_1 \amalg L_2) \simeq \bigvee_{j_1+j_2=j} \mathcal{X}_{Kh}^{j_1}(L_1) \wedge \mathcal{X}_{Kh}^{j_2}(L_2).$$

Moreover, if we fix a basepoint p in L_1 , not at a crossing, and consider the corresponding basepoint for $L_1 \amalg L_2$, then

$$\tilde{\mathcal{X}}_{Kh}^j(L_1 \amalg L_2) \simeq \bigvee_{j_1+j_2=j} \tilde{\mathcal{X}}_{Kh}^{j_1}(L_1) \wedge \mathcal{X}_{Kh}^{j_2}(L_2).$$

Similarly, reduced Khovanov homology should be well behaved for connected sums:

Conjecture 10.4. *Let L_1 and L_2 be based links and $L_1 \# L_2$ the connected sum of L_1 and L_2 , where we take the connected sum near the basepoints. Then*

$$\tilde{\mathcal{X}}_{Kh}^j(L_1 \# L_2) \simeq \bigvee_{j_1+j_2=j} \tilde{\mathcal{X}}_{Kh}^{j_1}(L_1) \wedge \tilde{\mathcal{X}}_{Kh}^{j_2}(L_2).$$

The unreduced Khovanov homology of a connected sum is slightly more complicated. A choice of basepoint makes the Khovanov complex $KC(L)$ into a module over $Kh(U) = \mathbb{Z}\langle x_+, x_- \rangle$. (The multiplication on $Kh(U)$ is defined as follows: x_+ acts as the identity and $x_-^2 = 0$.) At the level of complexes, $KC(L_1 \# L_2) \cong KC(L_1) \otimes_{Kh(U)} KC(L_2)$ [Kho03, Proposition 3.3]. There is also an analogous statement using comultiplication: $Kh(U)$ is a coalgebra with $\Delta(x_-) = x_- \otimes x_-$ and $\Delta(x_+) = x_+ \otimes x_- + x_- \otimes x_+$, and a choice of basepoint makes $KC(L)$ into a comodule over $Kh(U)$.

Lemma 10.5. *Let L_1 and L_2 be based links and $L_1 \# L_2$ the connected sum of L_1 and L_2 , where we take the connected sum near the basepoints. Then $KC(L_1 \# L_2)$ is the cotensor product over $Kh(U)$ of $KC(L_1)$ and $KC(L_2)$:*

$$KC(L_1 \# L_2) = KC(L_1) \square_{Kh(U)} KC(L_2).$$

Proof. Recall that the cotensor product $M \square_R N$ is defined by the short exact sequence

$$0 \longrightarrow M \square_R N \longrightarrow M \otimes N \xrightarrow{\Delta_M \otimes \text{Id}_N - \text{Id}_M \otimes \Delta_N} M \otimes R \otimes N;$$

see, for instance, [BW03, Section 1.10]. In particular, $KC(L_1) \square_{Kh(U)} KC(L_2)$ is the subcomplex of $KC(L_1) \otimes KC(L_2)$ generated by:

- elements $(D_{L_1}(u_1), x_1) \otimes (D_{L_2}(u_2), x_2)$ where x_i labels the marked circle in $D_{L_i}(u_i)$ by x_- , and
- elements $(D_{L_1}(u_1), x_1) \otimes (D_{L_2}(u_2), y_2) + (D_{L_1}(u_1), y_1) \otimes (D_{L_2}(u_2), x_2)$ where x_i (respectively y_i) labels the marked circle in $D_{L_i}(u_i)$ by x_- (respectively x_+).

The split map $\Delta: KC(L_1 \# L_2) \rightarrow KC(L_1) \otimes KC(L_2)$ is injective, and its image is exactly $KC(L_1) \square_{Kh(U)} KC(L_2)$ as just described. □

We can lift this to the level of spaces as follows. The space $\mathcal{X}_{Kh}(U)$ is $S^0 \vee S^0$. Denote the basepoint in $\mathcal{X}_{Kh}(U)$ by $*$, and the other two points by p_+ and p_- . Then

$$\mathcal{X}_{Kh}(U) \wedge \mathcal{X}_{Kh}(U) = \{ * \wedge *, p_- \wedge p_-, p_- \wedge p_+, p_+ \wedge p_-, p_+ \wedge p_+ \}.$$

The comultiplication on $Kh(U)$ dualizes to a basepoint-preserving map $\mathcal{X}_{Kh}(U) \wedge \mathcal{X}_{Kh}(U) \rightarrow \mathcal{X}_{Kh}(U)$ given by

$$\begin{aligned} p_- \wedge p_- &\mapsto p_- & p_+ \wedge p_- &\mapsto p_+ \\ p_- \wedge p_+ &\mapsto p_+ & p_+ \wedge p_+ &\mapsto *. \end{aligned}$$

This makes $\mathcal{X}_{Kh}(U)$ into a (discrete) ring spectrum.

Similarly, if L is a based link then $\mathcal{X}_{Kh}(L) \wedge \mathcal{X}_{Kh}(U) \simeq \mathcal{X}_{Kh}(L) \vee \mathcal{X}_{Kh}(L)$, where the first summand corresponds to p_- (say) and the second to p_+ . There is a map

$$m: \mathcal{X}_{Kh}(L) \wedge \mathcal{X}_{Kh}(U) \simeq \mathcal{X}_{Kh}(L) \vee \mathcal{X}_{Kh}(L) \rightarrow \mathcal{X}_{Kh}(L)$$

given as follows:

- On the first wedge summand $\mathcal{X}_{Kh}(L)$, m is the identity map.
- Recall from Section 8 that $|\mathcal{C}_K(L)|$ has a subcomplex $|\widetilde{\mathcal{C}}_{K^+}(L)|$ and a quotient complex $|\widetilde{\mathcal{C}}_{K^-}(L)| = |\mathcal{C}_K(L)| / |\widetilde{\mathcal{C}}_{K^+}(L)|$, and there is a canonical

identification $|\widetilde{\mathcal{E}}_{K_+}(L)| \cong |\widetilde{\mathcal{E}}_{K_-}(L)|$. On the second wedge summand of $\mathcal{X}_{Kh}(L)$, m collapses $|\widetilde{\mathcal{E}}_{K_+}(L)|$ to the basepoint and takes $|\widetilde{\mathcal{E}}_{K_-}(L)|$ to $|\widetilde{\mathcal{E}}_{K_+}(L)|$ by the canonical identification.

This makes $\mathcal{X}_{Kh}(L)$ into a module spectrum over $\mathcal{X}_{Kh}(U)$.

We can now formulate a space-level version of Lemma 10.5.

Conjecture 10.6. *Let L_1 and L_2 be based links. Then*

$$\mathcal{X}_{Kh}(L_1 \# L_2) \simeq \mathcal{X}_{Kh}(L_1) \otimes_{\mathcal{X}_{Kh}(U)} \mathcal{X}_{Kh}(L_2).$$

We leave the gradings in Conjecture 10.6 as an exercise to the reader.

Remark 10.7. Note that a smash product of Moore spaces is typically not a Moore space. (For example, consider $\mathbb{R}P^2 \wedge \mathbb{R}P^2$.) So, Conjecture 10.3, together with the computation of the Khovanov homology of the trefoil T , implies that $\mathcal{X}_{Kh}(T \amalg T)$ is not a wedge sum of Moore spaces. Similarly, Conjecture 10.4 (respectively Conjecture 10.6) implies that there are non-prime knots K for which $\widetilde{\mathcal{X}}_{Kh}(K)$ (respectively $\mathcal{X}_{Kh}(K)$) is not a wedge sum of Moore spaces.

ACKNOWLEDGMENTS

The authors are indebted to Ralph Cohen, Christopher Douglas and Ciprian Manolescu for many clarifying conversations on the Cohen-Jones-Segal construction. We thank Mikhail Khovanov and Cotton Seed for helpful comments on a previous version. Finally, we thank the referee for many helpful suggestions.

REFERENCES

- [AB95] D. M. Austin and P. J. Braam, *Morse-Bott theory and equivariant cohomology*, The Floer memorial volume, 1995, pp. 123–183. MR1362827 (96i:57037)
- [Bal10] John A. Baldwin, *On the spectral sequence from Khovanov homology to Heegaard Floer homology*, Int. Math. Res. Not. (2010).
- [Blo11] Jonathan M. Bloom, *A link surgery spectral sequence in monopole Floer homology*, Adv. Math. **226** (2011), no. 4, 3216–3281, DOI 10.1016/j.aim.2010.10.014. MR2764887
- [BN02] Dror Bar-Natan, *On Khovanov’s categorification of the Jones polynomial*, Algebr. Geom. Topol. **2** (2002), 337–370 (electronic), DOI 10.2140/agt.2002.2.337. MR1917056 (2003h:57014)
- [BW03] Tomasz Brzezinski and Robert Wisbauer, *Corings and comodules*, London Mathematical Society Lecture Note Series, vol. 309, Cambridge University Press, Cambridge, 2003. MR2012570 (2004k:16093)
- [CJS95] R. L. Cohen, J. D. S. Jones, and G. B. Segal, *Floer’s infinite-dimensional Morse theory and homotopy theory*, The Floer memorial volume, 1995, pp. 297–325. MR1362832 (96i:55012)
- [Coh10] Ralph L. Cohen, *The Floer homotopy type of the cotangent bundle*, Pure Appl. Math. Q. **6** (2010), no. 2, Special Issue: In honor of Michael Atiyah and Isadore Singer, pp. 391–438. MR2761853 (2011m:53173)
- [Cor00] Octavian Cornea, *Homotopical dynamics: suspension and duality*, Ergodic Theory Dynam. Systems **20** (2000), no. 2, 379–391, DOI 10.1017/S0143385700000183. MR1756976 (2002b:37017)
- [ELST] Brent Everitt, Robert Lipshitz, Sucharit Sarkar, and Paul Turner, *Khovanov homotopy types and the Dold-Thom functor*. arXiv:1202.1856.
- [ET] Brent Everitt and Paul Turner, *The homotopy theory of Khovanov homology*. arXiv:1112.3460.
- [Fra79] John M. Franks, *Morse-Smale flows and homotopy theory*, Topology **18** (1979), no. 3, 199–215, DOI 10.1016/0040-9383(79)90003-X. MR546790 (80k:58063)
- [Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002. MR1867354 (2002k:55001)

- [HKK] Po Hu, Daniel Kriz, and Igor Kriz, *Field theories, stable homotopy theory and Khovanov homology*. arXiv:1203.4773.
- [Jän68] Klaus Jänich, *On the classification of $O(n)$ -manifolds*, Math. Ann. **176** (1968), 53–76. MR0226674 (37 #2261)
- [Kho00] Mikhail Khovanov, *A categorification of the Jones polynomial*, Duke Math. J. **101** (2000), no. 3, 359–426, DOI 10.1215/S0012-7094-00-10131-7. MR1740682 (2002j:57025)
- [Kho03] Mikhail Khovanov, *Patterns in knot cohomology. I*, Experiment. Math. **12** (2003), no. 3, 365–374. MR2034399 (2004m:57022)
- [KM] Peter Kronheimer and Ciprian Manolescu, *Periodic Floer pro-spectra from the Seiberg-Witten equations*. arXiv:math/0203243.
- [Kra] Thomas Kragh, *The Viterbo transfer as a map of spectra*. arXiv:0712.2533.
- [Lau00] Gerd Laures, *On cobordism of manifolds with corners*, Trans. Amer. Math. Soc. **352** (2000), no. 12, 5667–5688 (electronic), DOI 10.1090/S0002-9947-00-02676-3. MR1781277 (2001i:55007)
- [Lee05] Eun Soo Lee, *An endomorphism of the Khovanov invariant*, Adv. Math. **197** (2005), no. 2, 554–586, DOI 10.1016/j.aim.2004.10.015. MR2173845 (2006g:57024)
- [LSa] Robert Lipshitz and Sucharit Sarkar, *A refinement of Rasmussen’s s -invariant*. To appear in Duke Math. J., arXiv:1206:3532.
- [LSb] Robert Lipshitz and Sucharit Sarkar, *A Steenrod square on Khovanov homology*. To appear in J. Topology, arXiv:1204.5776.
- [Man03] Ciprian Manolescu, *Seiberg-Witten-Floer stable homotopy type of three-manifolds with $b_1 = 0$* , Geom. Topol. **7** (2003), 889–932 (electronic), DOI 10.2140/gt.2003.7.889. MR2026550 (2005b:57060)
- [Man07] Ciprian Manolescu, *A gluing theorem for the relative Bauer-Furuta invariants*, J. Differential Geom. **76** (2007), no. 1, 117–153. MR2312050 (2008e:57033)
- [MO08] Ciprian Manolescu and Peter Ozsváth, *On the Khovanov and knot Floer homologies of quasi-alternating links*, Proceedings of Gökova Geometry-Topology Conference 2007, 2008, pp. 60–81. MR2509750 (2010k:57029)
- [ORS] Peter Ozsváth, Jacob Rasmussen, and Zoltán Szabó, *Odd Khovanov homology*, Vol. 13, 2013.
- [Ras05] Jacob Rasmussen, *Knot polynomials and knot homologies*, Geometry and topology of manifolds, Fields Inst. Commun., vol. 47, Amer. Math. Soc., Providence, RI, 2005, pp. 261–280. MR2189938 (2006i:57029)
- [Ras10] Jacob Rasmussen, *Khovanov homology and the slice genus*, Invent. Math. **182** (2010), no. 2, 419–447, DOI 10.1007/s00222-010-0275-6. MR2729272 (2011k:57020)
- [See] Cotton Seed, *Computations of the Lipshitz-Sarkar Steenrod square on Khovanov homology*. arXiv:1210.1882.
- [Spa59] E. H. Spanier, *Function spaces and duality*, Ann. Math. (2) **70** (1959), 338–378. MR0107862 (21 #6584)
- [Swi75] Robert M. Switzer, *Algebraic topology—homotopy and homology*, Springer-Verlag, New York, 1975. Die Grundlehren der mathematischen Wissenschaften, Band 212. MR0385836 (52 #6695)
- [Zie95] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995. MR1311028 (96a:52011)

DEPARTMENT OF MATHEMATICS, COLUMBIA UNIVERSITY, 2900 BROADWAY, NEW YORK, NEW YORK 10027

E-mail address: lipshitz@math.columbia.edu

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY 08544

E-mail address: sucharit@math.princeton.edu