

TOWARDS AN $sl(2)$ ACTION ON THE ANNULAR KHOVANOV SPECTRUM

ROSTISLAV AKHMECHET, VYACHESLAV KRUSHKAL, AND MICHAEL WILLIS

ABSTRACT. Given a link in the thickened annulus, its annular Khovanov homology carries an action of the Lie algebra \mathfrak{sl}_2 , which is natural with respect to annular link cobordisms. We consider the problem of lifting this action to the stable homotopy refinement of the annular homology. As part of this program, the actions of the standard generators of \mathfrak{sl}_2 are lifted to maps of spectra. In particular, it follows that the \mathfrak{sl}_2 action on homology commutes with the action of the Steenrod algebra. The main new technical ingredients developed in this paper, which may be of independent interest, concern certain types of cancellations in the cube of resolutions and the resulting more intricate structure of the moduli spaces in the framed flow category.

CONTENTS

1. Introduction	2
2. Background and conventions	4
2.1. Khovanov homology	4
2.2. Annular Khovanov homology	6
2.3. Arc diagrams	8
2.4. Action of \mathfrak{sl}_2 on annular Khovanov homology	10
3. Khovanov spectra	13
3.1. Using framed flow categories to construct Khovanov spectra	13
3.2. Constructing a map between Khovanov spectra	16
4. New features: cancellations and topology of moduli spaces	17
5. Constructing the framed flow category: The set-up and the base case	21
5.1. Definitions and the overall strategy	21
5.2. 0-dimensional moduli spaces for edges	24
6. Constructing the 1-dimensional moduli spaces for square faces	25
6.1. Square faces involving only the annular Khovanov differential	26
6.2. Square faces involving E or F	27
6.3. Square faces involving H	30
7. Path moduli spaces	32
7.1. Basic definitions and lemmas	33
7.2. Equivalent path moduli spaces and a non-emptiness condition	35
7.3. The E and F maps for connected arc diagrams	36
7.4. The H map for connected arc diagrams	41
8. Constructing the higher dimensional moduli spaces (the inductive step)	44
8.1. Subcubes involving E or F with connected arc diagrams	45
8.2. Subcubes involving E or F with disconnected arc diagrams	46
8.3. Subcubes involving H with connected arc diagrams	48
8.4. Subcubes involving H with disconnected arc diagrams	48

9. Invariance and naturality with respect to link cobordisms	49
9.1. Invariance under choices of cobordisms	50
9.2. Link cobordisms and Reidemeister invariance	51
References	54

1. INTRODUCTION

The structure of link homology theories is closely related to the representation theory of Lie algebras, and to categorification of quantum groups. This paper initiates the study of lifting these relations to the level of stable homotopy types. In the process of doing so, we provide the first example (to our knowledge) of constructing a stable homotopy refinement via framed flow categories of a homological invariant involving signs and cancellations.

We work in the setting of the \mathfrak{sl}_2 homology theory $Kh_{\mathbb{A}}(L)$ for links L in the thickened annulus, sometimes referred to as sutured annular Khovanov homology. Following constructions in [APS04, BN05, Rob13], this triply graded theory may be obtained from the usual Khovanov chain complex [Kho00] of a link diagram D by taking the annular degree zero part of the differential. It was shown in [BPW19] that given a tangle T , the homology $Kh_{\mathbb{A}}(\widehat{T})$ of its annular closure may also be obtained as the Hochschild homology of the complex of bimodules over the Chen-Khovanov algebra [CK14] associated to T .

Motivated in part by results of Auroux, Grigsby and Wehrli [AGW15] and by the work of Lauda [Lau10] on categorified \mathfrak{sl}_2 , Grigsby, Licata and Wehrli showed [GLW18] that the annular Khovanov homology $Kh_{\mathbb{A}}(L)$ carries an action of \mathfrak{sl}_2 . Analogous actions of \mathfrak{sl}_n on annular Khovanov-Rozansky homology were defined in [QR18]. In fact, the \mathfrak{sl}_2 action is constructed on the chain complex level, and it is natural with respect to annular link cobordisms. Here a trivial simple closed curve in the annulus is assigned the trivial representation of \mathfrak{sl}_2 , and essential curves are assigned the fundamental representation or its dual; [GLW18] extended the action to $Kh_{\mathbb{A}}(L)$ for any annular link L and showed that the action is a link invariant; see Section 2.4 for a detailed discussion.

Our results concern the stable homotopy refinement of Khovanov homology, introduced by Lipshitz and Sarkar [LS14a]. This construction, building on the work of Cohen, Jones and Segal [CJS95], associates to a diagram of an oriented link a framed flow category. The resulting suspension spectrum is a stable homotopy invariant of the link L ; its cohomology is isomorphic to the Khovanov homology of L . An alternative, more combinatorial construction using the Burnside category was given by Lawson, Lipshitz and Sarkar in [LLS20b]. The stable homotopy refinement $\mathcal{X}_{\mathbb{A}}(L)$ of an annular link L may be constructed using an analogous method, cf. [AKW, Section 4.3]. Alternatively, lifting the homology isomorphism with the Hochschild homology of Chen-Khovanov bimodules, $\mathcal{X}_{\mathbb{A}}(L)$ may be defined as the topological Hochschild homology of the tangle invariants introduced by Lawson-Lipshitz-Sarkar in [LLS20a].

These new links invariants valued in spectra can be studied using tools of stable homotopy theory. For example, they admit an action of the Steenrod algebra, and they are known to be stronger invariants than the underlying link homology, cf. [See12]. It is an interesting question what features of link homology admit a lift to the level of spectra. We address this question, giving a new structure on the annular Khovanov spectrum and thus providing a

new invariant of annular links and cobordisms. To this end, our main result is the following. Let J denote one of the standard generators E, F, H of \mathfrak{sl}_2 (see Section 2.4).

Theorem 1.1. *Given an annular diagram D of a link L in the thickened annulus, there exists a map of spectra $\mathcal{J} : \mathcal{X}_{\mathbb{A}}(D) \rightarrow \mathcal{X}_{\mathbb{A}}(D)$, whose induced map \mathcal{J}^* on cohomology is equal to the action of J on $Kh_{\mathbb{A}}(L)$. The homotopy class of \mathcal{J} is an invariant of the link L .*

As a consequence, we have the following statement about the \mathfrak{sl}_2 action on annular Khovanov homology.

Corollary 1.2. *The action of \mathfrak{sl}_2 on the annular homology $Kh_{\mathbb{A}}(L)$ of a link L commutes with the Steenrod cohomology operations.*

We carry out the construction in the setting of framed flow categories, recalled in Section 3.1. (An analogous construction likely can be done using a suitably defined “framed Burnside category”; compare the discussion in [LS18, Section 2.5].) The overall strategy is to define a stable homotopy refinement of the cone of the chain map J , see Section 3.2.

An important problem that immediately presents itself is the existence of *cancellations*. That is, within the cube-shaped chain complex $\text{Cone}(J)$, there exist pairs of consecutive edges so that a single generator is sent to a pair of generators, which then cancel by the second edge differential. This does not happen in the usual Khovanov chain complex, since the coefficients of the differential with respect to the canonical generators are either 0 or 1. There are signs but still no cancellations in the setting of odd Khovanov homology, whose stable homotopy refinement was given in [SSS20].

A specific example of cancellations in $\text{Cone}(J)$ is discussed in Section 4. Its manifestation in the context of framed flow categories is the presence of a pair of oppositely framed points, which cobound a framed interval in the 1-dimensional moduli space. In general this leads to moduli spaces of dimensions 2 and higher with non-trivial topology, see Sections 4, 8 and Figure 19.

The construction of the Khovanov homotopy type [LS14a] and other recent results in this setting involved framed flow categories whose n -dimensional moduli spaces are trivial covers of permutohedra, so topologically they are (disjoint unions of) n -dimensional balls. In general, if one attempts to build a framed flow category and the moduli spaces have non-trivial topology in some dimension, the existence of higher-dimensional moduli spaces is not assured. Indeed, there are highly restrictive compatibility conditions on how the moduli spaces of various dimensions fit together in a framed flow category (see Section 3.1), and there are choices involved in their construction. The union of moduli spaces up to a dimension $n - 1$ are required to form the boundary of the moduli spaces of dimension n , forming manifolds with corners. Thus extending the construction to dimension n amounts to finding a collection of null-cobordisms, a problem which a priori might not have a solution.

A key point in our definition of a framed flow category for $\text{Cone}(J)$ in the proof of Theorem 1.1 is that the moduli spaces may be constructed as (trivial covers of) codimension-1 submanifolds of some canonical spaces. The bulk of the paper, Sections 5 - 7, deals with setting up the base cases: 0- and 1-dimensional moduli spaces. The inductive step, using the Pontryagin-Thom construction and propagating the codimension-1 condition, is given in Section 8.

Another important feature in our work, as compared to prior constructions, is the necessity of performing a global combinatorial analysis of configurations of curves and surgery arcs in a link diagram. Recall that the structure of the 1-dimensional moduli spaces used in

the construction of the Khovanov homotopy type is analyzed by hand only for 2- and 3-dimensional faces of the Khovanov cube [LS14a, Sections 5.4,5.5]. These analyses are local in the sense that no more than three surgery arcs need to be considered. The triviality of the permutohedral covers then ensures that all higher dimensional moduli spaces can be built. In contrast, the more intricate structure of moduli spaces in our setting requires a detailed analysis of configurations of an arbitrary number of surgery arcs; see Section 7.

The construction of the map \mathcal{J} in Theorem 1.1 is stated as Corollary 5.9; the main technical result underlying its definition is Theorem 5.8. The invariance of the homotopy class of \mathcal{J} with respect to various choices made in the construction, as well as invariance under Reidemeister moves, is established in Section 9.

Recall that the \mathfrak{sl}_2 action on the annular complex is natural with respect to link cobordisms. However, link cobordism maps on spectra are only defined [LS14b] for link cobordisms which are presented as a sequence of Morse moves and Reidemeister moves. These maps are not known to be well-defined in general, although they are conjectured to be so.

Theorem 1.3. *The map \mathcal{J} constructed in Theorem 1.1 is natural with respect to link cobordisms presented as sequences of Morse moves and Reidemeister moves.*

The results of this paper concern the action of generators of \mathfrak{sl}_2 . The problem of lifting the relations of \mathfrak{sl}_2 is outside the scope of our present work; we hope to address it in a future paper. Another interesting problem concerns quantum annular homology $Kh_{\mathbb{A}_q}(L)$. This theory was defined by Beliakova, Putyra and Wehrli in [BPW19]. Extending the work of [GLW18], they showed that $Kh_{\mathbb{A}_q}(L)$ admits an action of $U_q(\mathfrak{sl}_2)$. Our earlier work [AKW] defined a stable homotopy refinement of quantum annular homology, which takes the form of a $\mathbb{Z}/r\mathbb{Z}$ -equivariant spectrum $\mathcal{X}_{\mathbb{A}_q}^r(L)$, where $r \geq 2$. Conjecture 1.4 in [AKW] states that the action of $U_q(\mathfrak{sl}_2)$ on $Kh_{\mathbb{A}_q}(L)$ can be lifted to an action on $\mathcal{X}_{\mathbb{A}_q}^r(L)$. The methods of the present paper do not immediately extend to the setting of $\mathcal{X}_{\mathbb{A}_q}^r(L)$ which was constructed using an equivariant version of the Burnside category. Nevertheless, our results are consistent with that conjecture, and a reformulation in terms of a “framed Burnside category”, briefly discussed above, is likely to lead to a lift of the action of the generators of $U_q(\mathfrak{sl}_2)$.

Acknowledgements. We are grateful to Sucharit Sarkar for many very helpful discussions on framed flow categories and the Burnside category. MW is grateful to Matt Stoffregen for several helpful conversations.

RA was supported by NSF RTG Grant DMS-1839968 and the Jefferson Scholars Foundation. VK was supported in part by the Miller Institute for Basic Research in Science at UC Berkeley, Simons Foundation fellowship 608604, and NSF Grant DMS-1612159. MW was supported in part by NSF FRG Grant DMS-1563615.

2. BACKGROUND AND CONVENTIONS

2.1. Khovanov homology. We begin with a brief overview of the Bar-Natan category \mathcal{BN} and the construction of the formal chain complex $[[D]]$ assigned to a link diagram D ; for a complete treatment see [BN02].

Recall the (dotted) Bar-Natan category $\mathcal{BN} = \mathcal{BN}(\mathbb{R}^2)$. Let $I := [0, 1]$ denote the unit interval. Objects of \mathcal{BN} are formal direct sums of formally graded collections of simple closed curves in the plane \mathbb{R}^2 . Morphisms are matrices whose entries are formal \mathbb{Z} -linear combinations of dotted cobordisms properly embedded in $\mathbb{R}^2 \times I$. Cobordisms are considered

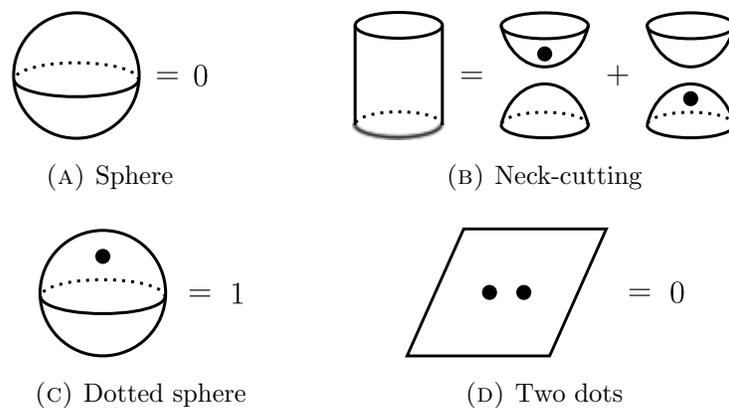


FIGURE 1. Bar-Natan relations

up to isotopy relative to the boundary, and they are subject to the local relations shown in Figure 1.

Let $A = \mathbb{Z}[X]/(X^2)$ denote the Frobenius algebra with comultiplication, unit, and counit described below.

$$\begin{array}{lll}
 \Delta: A \rightarrow A \otimes A & \eta: \mathbb{Z} \rightarrow A & \varepsilon: A \rightarrow \mathbb{Z} \\
 1 \mapsto X \otimes 1 + 1 \otimes X & 1 \mapsto 1 & 1 \mapsto 0 \\
 X \mapsto X \otimes X & & X \mapsto 1
 \end{array}$$

This is the Frobenius algebra underlying \mathfrak{sl}_2 link homology [Kho00]. Define the quantum grading qdeg on A by setting

$$(1) \quad \text{qdeg}(X) = -1 \quad \text{qdeg}(1) = 1.$$

Let $\mathbb{Z}\text{-gmod}$ denote the category of \mathbb{Z} -graded abelian groups and graded maps (of any degree) between them. The Frobenius algebra A defines a $(1+1)$ -dimensional TQFT, which descends to a functor

$$(2) \quad \mathcal{F}_{Kh}: \mathcal{BN} \rightarrow \mathbb{Z}\text{-gmod}.$$

We will refer to \mathcal{F}_{Kh} as the *Khovanov TQFT*. The construction is reviewed below.

On an object \mathcal{C} which is a collection of n disjoint circles in the plane, $\mathcal{F}_{Kh}(\mathcal{C}) = A^{\otimes n}$, generated by elements of the form

$$(3) \quad y = y_1 \otimes \cdots \otimes y_n \in A^{\otimes n}$$

where each $y_i \in \{1, X\}$; this corresponds to a label of each circle in \mathcal{C} by 1 or X . We refer to such generators as *Khovanov generators*.

Morphisms in \mathcal{BN} are generated by so-called *elementary cobordisms* which contain only a single non-degenerate critical point with respect to the height function $\mathbb{R}^2 \times I \rightarrow I$. Index 0 (cup) and 2 (cap) elementary cobordisms are assigned the unit and co-unit maps of A respectively; index 1 (saddle) elementary cobordisms are assigned multiplication or comultiplication maps according to whether the cobordism merged two circles into one, or split one circle into two. The presence of a dot on a cobordism corresponds to multiplication by $X \in A$.

Now let D be a planar diagram for an oriented link $L \subset S^3$. We recall the construction of the formal Bar-Natan chain complex $[[D]]$ from [BN05]. First form the *cube of resolutions*.

Label the crossings of the diagram by $1, \dots, n$. Each crossing may be resolved in one of two ways, called the 0 -smoothing and 1 -smoothing, as in (4).

$$(4) \quad \left. \right) \left(\begin{array}{ccc} & \swarrow & \searrow \\ \xleftarrow{0} & \times & \xrightarrow{1} \\ & \swarrow & \searrow \end{array} \right) \left. \begin{array}{c} \text{---} \\ \text{---} \end{array} \right)$$

For each $u = (u_1, \dots, u_n) \in \{0, 1\}^n$, perform the u_i -smoothing at the i -th crossing. This results in a collection of disjoint simple closed curves in the plane, which we denote D_u . Identifying $\{0, 1\}^n$ with the vertices of an n -dimensional cube, we decorate each vertex u by its corresponding smoothing D_u .

Let $u = (u_1, \dots, u_n)$ and $v = (v_1, \dots, v_n)$ be vertices which differ only in the i -th entry, where $u_i = 0$ and $v_i = 1$. The diagrams D_u and D_v are identical outside of a small disk around the i -th crossing. There is a cobordism from D_u to D_v , which is the obvious saddle near the i -th crossing and the identity (product cobordism) elsewhere. We will call this the *saddle cobordism* from D_u to D_v , and denote it by $d_{u,v}$. Decorate each edge of the n -dimensional cube by the corresponding saddle cobordism. The result is a commutative cube in the category \mathcal{BN} . There is an assignment $s_{u,v} \in \{0, 1\}$ for each edge so that multiplying the edge map $d_{u,v}$ by $(-1)^{s_{u,v}}$ results in an anti-commutative cube (see [BN05, Section 2.7], also [LS14a, Definition 4.5]).

For $u = (u_1, \dots, u_n) \in \{0, 1\}^n$, set $|u| = \sum_i u_i$. Now, form the chain complex $[[D]]$ by setting

$$[[D]]^i = \bigoplus_{|u|=i+n_-} D_u \{i + n_+ - n_-\}$$

where n_- , n_+ are the number of negative and positive crossings in D , and $\{-\}$ denotes the formal grading shift in \mathcal{BN} . The differential is defined on each summand by the edge map $(-1)^{s_{u,v}} d_{u,v}$. That $[[D]]$ is a chain complex follows from anti-commutativity of the cube.

Theorem 2.1. ([BN05, Theorem 1]) *If diagrams D and D' are related by a Reidemeister move, then $[[D]]$ and $[[D']]$ are chain homotopy equivalent.*

Let

$$CKh(D) := \mathcal{F}_{Kh}([[D]])$$

denote the chain complex obtained by applying the Khovanov TQFT to $[[D]]$; it is precisely the invariant defined in [Kho00, Section 7].

2.2. Annular Khovanov homology. Asaeda-Przytycki-Sikora [APS04] introduced a homology theory for links in I -bundles over surfaces. The case of the thickened annulus is known as annular Khovanov homology, sometimes called *sutured* annular Khovanov homology. This section reviews its construction, following [Rob13, GLW18].

Let $\mathbb{A} = S^1 \times I$ denote the annulus. An *annular link* is a link in the thickened annulus $\mathbb{A} \times I$, and its diagram is obtained by projecting onto the first factor \mathbb{A} . Embed \mathbb{A} standardly in \mathbb{R}^2 via

$$\mathbb{A} = \{x \in \mathbb{R}^2 \mid 1 \leq |x| \leq 2\},$$

so that an annular link diagram and all of its smoothings are drawn in the punctured plane $\mathbb{R}^2 \setminus (0, 0)$. We represent the annulus by simply indicating the deleted point using the symbol \times . See Figure 2 for an example of an annular link diagram drawn using this convention. We

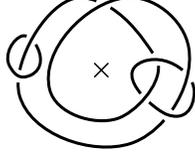


FIGURE 2. An annular link diagram

distinguish two types of circles in \mathbb{A} : *trivial* circles, which are contractible in \mathbb{A} , and *essential* ones, which are not contractible.

Let $\mathcal{BN}(\mathbb{A})$ denote the Bar-Natan category of the annulus. Its objects are formal direct sums of formally \mathbb{Z}^2 -graded collections of simple closed curves in \mathbb{A} . Morphisms are matrices with entries formal \mathbb{Z} -linear combinations of dotted cobordisms properly embedded in $\mathbb{A} \times I$, modulo isotopy relative to the boundary, and subject to the local Bar-Natan relations shown in Figure 1.

Let $\mathbb{Z}\text{-gmod}$ denote the category of \mathbb{Z}^2 -graded abelian groups and graded maps (of any bidegree) between them. We now describe the annular TQFT

$$\mathcal{F}_{\mathbb{A}} : \mathcal{BN}(\mathbb{A}) \rightarrow \mathbb{Z}\text{-gmod},$$

following the construction in [Rob13].

Let $\mathcal{C} \subset \mathbb{A}$ be a collection of n trivial and m essential circles. Viewing \mathbb{A} as a subspace of \mathbb{R}^2 , apply the Khovanov TQFT \mathcal{F}_{Kh} from Section 2.1,

$$\mathcal{F}_{Kh}(\mathcal{C}) = A^{\otimes n} \otimes A^{\otimes m}.$$

Introduce a second grading, called the *annular* grading and denoted adeg , on $\mathcal{F}_{Kh}(\mathcal{C})$ as follows. Every tensor factor corresponding to a trivial circle is concentrated in annular degree zero. For a factor A associated to an essential circle, introduce the notation

$$v_- = X, \quad v_+ = 1$$

to denote a basis for this copy of A , and define adeg by setting

$$\text{adeg}(v_-) = -1 \quad \text{adeg}(v_+) = 1.$$

The underlying abelian group of $\mathcal{F}_{\mathbb{A}}(\mathcal{C})$ is defined to be $\mathcal{F}_{Kh}(\mathcal{C})$, with the bigrading given by $(\text{qdeg}, \text{adeg})$. To define $\mathcal{F}_{\mathbb{A}}$ on a cobordism $S \subset \mathbb{A} \times I$, first view S as a surface in $\mathbb{R}^2 \times I$ and consider the map $\mathcal{F}_{Kh}(S)$ assigned to S by the Khovanov TQFT. It was observed in [Rob13, Section 2] that $\mathcal{F}_{Kh}(S)$ is of the form

$$(5) \quad \mathcal{F}_{Kh}(S) = \mathcal{F}_{Kh}(S)_0 + \mathcal{F}_{Kh}(S)_-$$

where $\mathcal{F}_{Kh}(S)_0$ preserves adeg and $\mathcal{F}_{Kh}(S)_-$ lowers adeg . Set

$$\mathcal{F}_{\mathbb{A}}(S) := \mathcal{F}_{Kh}(S)_0$$

to be the part of $\mathcal{F}_{Kh}(S)$ that preserves adeg . It follows from (5) that $\mathcal{F}_{\mathbb{A}}$ is functorial with respect to composition of cobordisms, and it clearly factors through the relations in Figure 1. We will refer to $\mathcal{F}_{\mathbb{A}}$ as the *annular TQFT*. Note that if S carries d dots, then $\mathcal{F}_{\mathbb{A}}(S)$ is a map of $(\text{qdeg}, \text{adeg})$ bidegree

$$(\chi(S) - 2d, 0).$$

We distinguish the bigraded modules assigned to trivial and essential circles by writing

$$V = \mathcal{F}_{\mathbb{A}}(C)$$

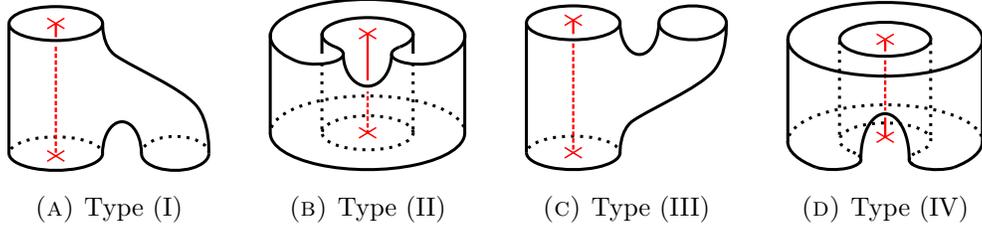


FIGURE 3. Saddles involving essential circles. The vertical red arc is the central axis of $\mathbb{A} \times I \subset \mathbb{R}^2 \times I$.

if C is an essential circle, with basis written as $\{v_-, v_+\}$. The notation A is reserved for the module assigned to a trivial circle, with basis $\{1, X\}$. Then if $\mathcal{C} \subset \mathbb{A}$ consists of n trivial and m essential circles, the module assigned to \mathcal{C} by $\mathcal{F}_{\mathbb{A}}$ is written as

$$\mathcal{F}_{\mathbb{A}}(\mathcal{C}) = A^{\otimes n} \otimes V^{\otimes m}.$$

Cobordisms between trivial circles are assigned the same map by \mathcal{F}_{Kh} and $\mathcal{F}_{\mathbb{A}}$. We record here the maps assigned to the four elementary saddle cobordisms involving at least one essential circle, shown in Figure 3, using the V notation.

$$(6) \quad \begin{aligned} V \otimes A &\xrightarrow{(I)} V \\ v_- \otimes 1 &\mapsto v_- \\ v_+ \otimes 1 &\mapsto v_+ \\ v_- \otimes X &\mapsto 0 \\ v_+ \otimes X &\mapsto 0 \end{aligned}$$

$$(7) \quad \begin{aligned} V \otimes V &\xrightarrow{(II)} A \\ v_- \otimes v_- &\mapsto 0 \\ v_+ \otimes v_- &\mapsto X \\ v_- \otimes v_+ &\mapsto X \\ v_+ \otimes v_+ &\mapsto 0 \end{aligned}$$

$$(8) \quad \begin{aligned} V &\xrightarrow{(III)} V \otimes A \\ v_- &\mapsto v_- \otimes X \\ v_+ &\mapsto v_+ \otimes X \end{aligned}$$

$$(9) \quad \begin{aligned} A &\xrightarrow{(IV)} V \otimes V \\ 1 &\mapsto v_+ \otimes v_- + v_- \otimes v_+ \\ X &\mapsto 0 \end{aligned}$$

Finally, let D be a diagram for an oriented annular link L . The construction of $[[D]]$ as described in Section 2.1 is completely local, and crossings are away from the puncture \times . Thus we may view $[[D]]$ as a chain complex over $\mathcal{BN}(\mathbb{A})$, with the grading shifts $\{-\}$ in \mathcal{BN} rewritten as bigrading shifts $\{-, 0\}$ in $\mathcal{BN}(\mathbb{A})$. Annular link diagrams representing isotopic annular links are related by Reidemeister moves away from the puncture. Therefore the chain homotopy class of $[[D]]$, viewed as a chain complex over $\mathcal{BN}(\mathbb{A})$, is an invariant of L . Applying the annular TQFT $\mathcal{F}_{\mathbb{A}}$, we obtain the annular Khovanov chain complex

$$(10) \quad CKh_{\mathbb{A}}(D) := \mathcal{F}_{\mathbb{A}}([[D]]).$$

It is an invariant of L up to chain homotopy equivalence.

2.3. Arc diagrams. In our resolution diagrams throughout the Khovanov cube for a link diagram D , we will use additional decorations to record whether the 0-smoothing or 1-smoothing was used at a crossing. These form an enhancement of the *resolution configurations* from [LS14a, Section 2], and they will be crucial to our combinatorial analysis in Section 7.

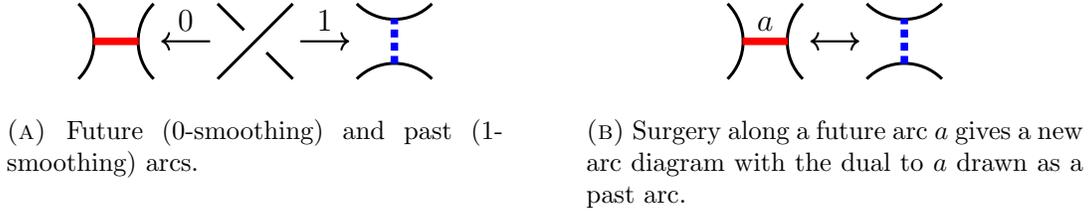


FIGURE 4

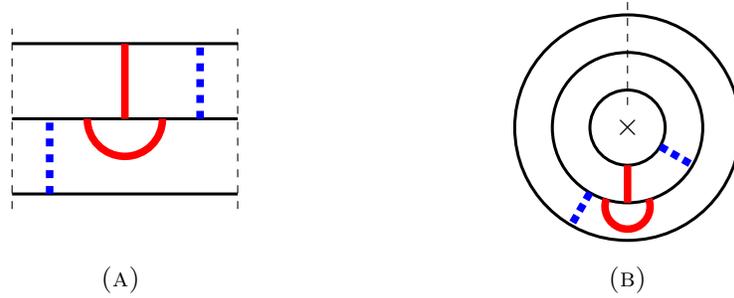


FIGURE 5. A rectangular arc diagram and its annular closure

Definition 2.2. Let D be an n -crossing link diagram and let $u \in \{0, 1\}^n$ be a vertex. The *arc diagram* D_u is obtained from the resolution D_u as follows. Join the two strands at each smoothing by one of two types of arcs, a solid red arc, called a *future arc*, if the 0-resolution was used, or a dashed blue arc, called a *past arc*, if the 1-resolution was used.

The local picture is shown in Figure 4a. To explain the terminology, a 0-smoothing means that the saddle corresponding to that crossing is yet to be performed, while a 1-smoothing means that the associated saddle has already been performed. For an arc a in an arc diagram, we can obtain a new arc diagram by performing surgery along a and then inserting the dual surgery arc to a and switching its type, as shown in Figure 4b. We will refer to this process as performing surgery along a , with the insertion of the dual arc of opposite type implied. An edge $D_u \rightarrow D_v$ in the cube of resolutions of D thus corresponds to performing surgery along a future arc in D_u . More generally, given any pair of vertices $u, v \in \{0, 1\}^n$, the arc diagram D_v can be obtained from D_u by a sequence of such surgeries along past and future arcs in D_u .

In what follows, we will depict annular arc diagrams by drawing planar tangles in the rectangle, with the understanding that the annular arc diagram is obtained by connecting the left endpoints to the right. Thus horizontal intervals are segments of distinct essential circles. See Figure 5 for an example. We will omit the dashed platforms in the future.

An arc diagram is *connected* if it is connected as a subspace of the plane. Similarly, we will refer to the connected components of arc diagrams. Note that, since surgery along an arc preserves connectedness, there is a natural correspondence between the connected components of D_u and the connected components of D_v for any two vertices $u, v \in \{0, 1\}^n$.

Now let $u = (0, \dots, 0)$ be the starting vertex of the cube, and suppose we have a partition of D_u into disjoint closed (not necessarily connected) components

$$D_u = \mathcal{C}_1 \sqcup \dots \sqcup \mathcal{C}_k.$$

(A) The arc diagram at the starting vertex
(0, 0, 0, 0)(B) An arc diagram for the subcube
 $(1, 0, 0, 0) \leq_2 (1, 0, 1, 1)$

FIGURE 6

In a similar fashion, this partition determines such a partition of the arc diagram D_v at every vertex $v \in \{0, 1\}^n$. Abusing notation slightly, we will use the same notation for these closed components at every vertex:

$$D_v = \mathcal{C}_1 \sqcup \cdots \sqcup \mathcal{C}_k.$$

Then for any Khovanov generator $x \in \mathcal{F}_{Kh}(D_v)$, we will write

$$(11) \quad x = x|_{\mathcal{C}_1} \otimes \cdots \otimes x|_{\mathcal{C}_k} \in \mathcal{F}_{Kh}(D_v)$$

where $x|_{\mathcal{C}_i}$ denotes the restriction of the labels in x to the circles in $\mathcal{C}_i \subset D_v$. Note that this is again a slight abuse of notation, since the use of the symbol \otimes in Equation (11) is, strictly speaking, incompatible with the use of the same symbol in Equation (3), but this should not cause any confusion in our arguments.

In many places throughout the paper, we will be concerned with certain *subcubes* within the large cube $\{0, 1\}^n$. A k -dimensional subcube, denoted $u \leq_k v$, is determined by a pair of vertices u, v such that $u_i \leq v_i$ for all i , with precisely k coordinates i where $u_i = 0$ and $v_i = 1$. When focusing on such a subcube, we will draw the arc diagram D_u by including only the k future arcs corresponding to these k coordinates, and likewise for any intermediate vertices $u \leq w \leq v$ within the subcube. See Figure 6 for an example.

2.4. Action of \mathfrak{sl}_2 on annular Khovanov homology. It is shown in [GLW18] that the annular TQFT $\mathcal{F}_{\mathbb{A}}$ can be viewed as taking values in the category $\text{gRep}(\mathfrak{sl}_2)$ of graded \mathfrak{sl}_2 representations. We first recall some basic definitions. Let \mathfrak{sl}_2 be the \mathbb{Z} -span of the 2×2 matrices

$$E = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, F = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, H = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The usual commutator bracket,

$$[x, y] = xy - yx,$$

makes \mathfrak{sl}_2 into a Lie algebra, with the bracket given on generators by

$$(12) \quad [E, F] = H, [H, E] = 2E, [H, F] = -2F.$$

Remark 2.3. The Lie algebra \mathfrak{sl}_2 in [GLW18] is defined over \mathbb{C} . For the purpose of this paper we continue working over \mathbb{Z} , since all desired results already hold integrally.

Recall that if M is an \mathfrak{sl}_2 -module, then its linear dual

$$M^* = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$$

is again an \mathfrak{sl}_2 -module via

$$(xf)(m) := -f(xm)$$

for $x \in \mathfrak{sl}_2$, $f \in M^*$, and $m \in M$. Likewise, if N is another \mathfrak{sl}_2 -module, then the tensor product $M \otimes_{\mathbb{Z}} N$ inherits an \mathfrak{sl}_2 action by setting

$$x(m \otimes n) := xm \otimes n + m \otimes xn$$

for $x \in \mathfrak{sl}_2$, $m \in M$, and $n \in N$.

Let $V_1 = \mathbb{Z} \oplus \mathbb{Z}$ denote the fundamental representation of \mathfrak{sl}_2 with standard basis vectors denoted $\mathbf{v}_0 = (0, 1)$, $\mathbf{v}_1 = (1, 0)$, and let V_1^* denote its dual, with dual basis vectors $\mathbf{v}_0^*, \mathbf{v}_1^* \in V_1^*$. The action of \mathfrak{sl}_2 on V_1 and V_1^* is summarized in (13).

$$(13) \quad \begin{array}{lll} E\mathbf{v}_0 = \mathbf{v}_1 & F\mathbf{v}_0 = 0 & H\mathbf{v}_0 = -\mathbf{v}_0 \\ E\mathbf{v}_1 = 0 & F\mathbf{v}_1 = \mathbf{v}_0 & H\mathbf{v}_1 = \mathbf{v}_1 \end{array}$$

$$\begin{array}{lll} E\mathbf{v}_0^* = 0 & F\mathbf{v}_0^* = -\mathbf{v}_1^* & H\mathbf{v}_0^* = \mathbf{v}_0^* \\ E\mathbf{v}_1^* = -\mathbf{v}_0^* & F\mathbf{v}_1^* = 0 & H\mathbf{v}_1^* = -\mathbf{v}_1^* \end{array}$$

Let $\mathcal{C} \subset \mathbb{A}$ be a collection of n trivial and m essential circles. There is a natural ordering on the essential circles, starting from the innermost circle (closest to the puncture \times). Make

$$\mathcal{F}_{\mathbb{A}}(\mathcal{C}) = A^{\otimes n} \otimes V^{\otimes m}$$

into an \mathfrak{sl}_2 -module as follows. Every factor A corresponding to a trivial circle is assigned the trivial (zero) two-dimensional representation of \mathfrak{sl}_2 . For a factor V corresponding to the i -th essential circle in \mathcal{C} , identify V with V_1 if i is odd via the \mathbb{Z} -linear isomorphism (14) and with V_1^* if i is even via (15).

$$(14) \quad V_1 \rightarrow V, \mathbf{v}_0 \mapsto v_-, \mathbf{v}_1 \mapsto v_+,$$

$$(15) \quad V_1^* \rightarrow V, \mathbf{v}_0^* \mapsto v_+, \mathbf{v}_1^* \mapsto v_-.$$

Thus trivial circles are always assigned the trivial two-dimensional representation, and essential circles are assigned V_1 and V_1^* in an alternating manner, with the convention that the innermost essential circle is assigned V_1 .

Example 2.4. The action of E and F on two essential circles is recorded below.

$$\begin{array}{ll} E(v_- \otimes v_-) = v_+ \otimes v_- - v_- \otimes v_+ & E(v_+ \otimes v_-) = -v_+ \otimes v_+ \\ E(v_- \otimes v_+) = v_+ \otimes v_+ & E(v_+ \otimes v_+) = 0 \end{array}$$

$$\begin{array}{ll} F(v_- \otimes v_-) = 0 & F(v_+ \otimes v_-) = v_- \otimes v_- \\ F(v_- \otimes v_+) = -v_- \otimes v_- & F(v_+ \otimes v_+) = v_- \otimes v_+ - v_+ \otimes v_- \end{array}$$

Remark 2.5. The \mathfrak{sl}_2 representations V_1 and V_1^* are isomorphic via

$$\mathbf{v}_0 \mapsto -\mathbf{v}_1^*, \quad \mathbf{v}_1 \mapsto \mathbf{v}_0^*.$$

However, the map $V_1 \rightarrow V_1^*$ obtained by composing (14) with the inverse of (15) is not \mathfrak{sl}_2 -linear; indeed, the identifications (14) and (15) equip V with two actions of \mathfrak{sl}_2 which differ by a sign.

By [GLW18, Lemma 4], maps assigned to cobordisms in $\mathbb{A} \times I$ commute with the action of \mathfrak{sl}_2 . Therefore the annular TQFT $\mathcal{F}_{\mathbb{A}}$ is upgraded to a functor

$$\mathcal{F}_{\mathbb{A}}: \mathcal{BN}(\mathbb{A}) \rightarrow \text{gRep}(\mathfrak{sl}_2)$$

landing in the category of \mathbb{Z} -graded \mathfrak{sl}_2 representations. The annular grading corresponds to the weight space decomposition, and the quantum grading corresponds to the external grading. It follows that \mathfrak{sl}_2 acts on the annular chain complex (10) and its homology.

Let $\mathcal{C} \subset \mathbb{A}$ be a collection of n trivial and m essential circles. We will always assume that the essential circles are ordered from innermost to outermost, so that the i -th tensor factor of V in

$$\mathcal{F}_{\mathbb{A}}(\mathcal{C}) = A^{\otimes n} \otimes V^{\otimes m}$$

corresponds to the i -th essential circle. Any standard generator in $V^{\otimes m}$ can be written as $v_{s_1} \otimes \cdots \otimes v_{s_m}$, with each $s_i \in \{\pm\}$. Thus any standard generator $x \in \mathcal{F}_{\mathbb{A}}(\mathcal{C})$ is of the form

$$(16) \quad x = y \otimes v_{s_1} \otimes \cdots \otimes v_{s_m}$$

where $y \in A^{\otimes n}$ is a standard generator for the trivial circles in \mathcal{C} .

By construction, \mathfrak{sl}_2 weights correspond to the annular grading, so that for a standard generator $x \in \mathcal{F}_{\mathbb{A}}(\mathcal{C})$, H acts as multiplication by $\text{adeg}(x)$,

$$(17) \quad Hx = \text{adeg}(x)x.$$

Although the formula (17) is straightforward, we can refine it by writing

$$(18) \quad Hx = \sum_{i=1}^m s_i y \otimes v_{s_1} \otimes \cdots \otimes v_{s_m}$$

using the notation (16). The action of E and F can similarly be written as

$$(19) \quad \begin{aligned} Ex &= \sum_{i=1}^m y \otimes v_{s_1} \otimes \cdots \otimes Ev_{s_i} \otimes \cdots \otimes v_{s_m}, \\ Fx &= \sum_{i=1}^m y \otimes v_{s_1} \otimes \cdots \otimes Fv_{s_i} \otimes \cdots \otimes v_{s_m}. \end{aligned}$$

In the formulas (19), Ev_{s_i} and Fv_{s_i} are given by

$$Ev_{s_i} = \begin{cases} (-1)^{i+1}v_+ & \text{if } s_i = - \\ 0 & \text{if } s_i = + \end{cases} \quad Fv_{s_i} = \begin{cases} (-1)^{i+1}v_- & \text{if } s_i = + \\ 0 & \text{if } s_i = - \end{cases}$$

Letting C denote the i -th essential circle in \mathcal{C} , will say that E acts on C to mean that we apply E on the i -th factor of V in $\mathcal{F}_{\mathbb{A}}(\mathcal{C})$,

$$x \mapsto y \otimes v_{s_1} \otimes \cdots \otimes Ev_{s_i} \otimes \cdots \otimes v_{s_m},$$

and likewise for F . If J denotes one of E , F , or H , we will also say the J map to mean the endomorphism of the module assigned to a collection of circles

$$J: \mathcal{F}_{\mathbb{A}}(\mathcal{C}) \rightarrow \mathcal{F}_{\mathbb{A}}(\mathcal{C}), \quad x \mapsto Jx,$$

or on the annular chain complex

$$J: CKh_{\mathbb{A}}(D) \rightarrow CKh_{\mathbb{A}}(D),$$

which is built out of the J map on each smoothing.

The following is a combination of [GLW18, Lemma 2] and [GLW18, Lemma 5].

Lemma 2.6. *Let D be an annular link diagram. For each smoothing D_u , let $\Theta_u: \mathcal{F}_{\mathbb{A}}(D_u) \rightarrow \mathcal{F}_{\mathbb{A}}(D_u)$ be the involution given by $v_{\pm} \mapsto v_{\mp}$ on essential circles and the identity map on the trivial circles. The maps Θ_u assemble into an isomorphism*

$$\Theta: CKh_{\mathbb{A}}(D) \xrightarrow{\sim} CKh_{\mathbb{A}}(D)$$

of chain complexes of abelian groups, which fits into the commutative diagram

$$\begin{array}{ccc} CKh_{\mathbb{A}}(D) & \xrightarrow{\Theta} & CKh_{\mathbb{A}}(D) \\ E \downarrow & & \downarrow F \\ CKh_{\mathbb{A}}(D) & \xrightarrow{\Theta} & CKh_{\mathbb{A}}(D). \end{array}$$

Therefore the action of F is related to E simply by $F = \Theta E \Theta$. Moreover, Θ sends generators to generators bijectively, so in our combinatorial analysis we can focus on one of E or F . Note however that Θ does not preserve $(\text{qdeg}, \text{adeg})$ -bidegree. A modified quantum grading is used in [GLW18], there denoted as j' . In our notation, $j' = \text{qdeg} - \text{adeg}$.

3. KHOVANOV SPECTRA

3.1. Using framed flow categories to construct Khovanov spectra. There is a general method due to Cohen-Jones-Segal [CJS95] for constructing a spectrum \mathcal{X} with cellular cochain complex $C^*(\mathcal{X})$ isomorphic to a given cochain complex C^* . The essential idea is to build a cell complex whose cells correspond to the generators of C^* and whose attaching maps correspond to the differential of C^* . Of course, the various compositions of attaching maps must be homotopically coherent in order for this to make sense. The data needed to guarantee this coherence is organized in a *framed flow category*, which we will review here. Most of this material is based upon the corresponding material in [LS14a, LLS20b] where the authors successfully constructed such a category (and thus a spectrum) lifting the Khovanov chain complex of a link in S^3 .

Definition 3.1. ([LS14a, Definition 3.12]) A *flow category* \mathfrak{F} is a category consisting of the following data.

- A finite set of *objects*, also denoted \mathfrak{F} , together with an indexing map $h: \mathfrak{F} \rightarrow \mathbb{Z}$. In analogy with Morse theory, an object $x \in \mathfrak{F}$ is thought of as a critical point of index $h(x)$ for some Morse function.
- For each $x \neq y \in \mathfrak{F}$, a *morphism space* (usually called a *moduli space*) $\mathcal{M}(y, x)$ which is a $(h(y) - h(x) - 1)$ -dimensional manifold with corners (see [LS14a, Section 3.1]); negative dimensional manifolds are all empty. In analogy with Morse theory, $\mathcal{M}(y, x)$ is often thought of as the moduli space of flow lines from y down to x . Furthermore, the codimension- m boundary $\partial_{[m]}\mathcal{M}(y, x)$ of any such moduli space can be identified with the disjoint union

$$\partial_{[m]}\mathcal{M}(y, x) \cong \coprod_{(z_1, \dots, z_m) \in \mathfrak{F}^m} \mathcal{M}(y, z_m) \times \mathcal{M}(z_m, z_{m-1}) \times \cdots \times \mathcal{M}(z_1, x).$$

In analogy with Morse theory, $\partial_{[m]}\mathcal{M}(y, x)$ is thought of as the m -times broken flow lines from y down to x .

Remark 3.2. In [LS14a, LLS20b], the authors use a notion of an $\langle n \rangle$ -manifold, which is a manifold with corners X whose boundary ∂X is decomposed in a particular manner into n pieces, $\partial X = \partial_1 X \cup \dots \cup \partial_n X$. We use square brackets in $\partial_{[m]} X$ to denote the codimension m boundary of a manifold with corners in order to distinguish from the earlier notation.

Definition 3.3. ([LS14a, Definition 3.20]) A *framed flow category* is a flow category \mathfrak{F} whose moduli spaces come equipped with *framed neat embeddings* into cornered Euclidean space (see [LLS20b, Definition 3.7]) in a manner consistent with the boundary identifications indicated in Definition 3.1.

Definition 3.4. A framed flow category \mathfrak{F} determines a cochain complex $\text{Tot}^*(\mathfrak{F})$, called the *totalization complex* of \mathfrak{F} , defined as follows.

- The generators of $\text{Tot}^*(\mathfrak{F})$ are in bijection with the objects of \mathfrak{F} .
- For any $x \in \mathfrak{F}$, $h(x)$ gives the homological degree of the corresponding generator $x \in \text{Tot}^*(\mathfrak{F})$.
- For any $x \in \text{Tot}^*(\mathfrak{F})$, we have

$$dx = \sum_{\{y \in \mathfrak{F} | h(y) = h(x) + 1\}} |\mathcal{M}(y, x)| y$$

where $|\mathcal{M}(y, x)|$ is the signed count of the framed points comprising the 0-dimensional manifold $\mathcal{M}(y, x)$.

In analogy with Morse theory, $\text{Tot}^*(\mathfrak{F})$ can be viewed as the Morse cochain complex associated to the corresponding Morse function. Then we say \mathfrak{F} *refines* a complex C^* if $\text{Tot}^*(\mathfrak{F}) \cong C^*$.

Remark 3.5. Note that the totalization notation in [LLS20b] is assigned to Burnside functors, while the term *associated cochain complex* is used in [LS14a] to describe the complex we are calling the totalization in Definition 3.4

Theorem 3.6. ([LS14a, Lemmas 3.24 – 3.26]) *A framed flow category \mathfrak{F} gives rise to a cell complex $|\mathfrak{F}|$ whose cellular cochain complex is isomorphic to $\text{Tot}^*(\mathfrak{F})$ up to an overall homological shift. The stable homotopy type of $|\mathfrak{F}|$ is an invariant of the given framed flow category \mathfrak{F} .*

The definitions above are considered in much greater detail in [LS14a, LLS20b], together with the process by which one can build the cell complex $|\mathfrak{F}|$. In short, each object $x \in \mathfrak{F}$ corresponds to a cell of dimension $h(x) + k$ for some fixed offset k . The boundary of such a cell is made up of cornered Euclidean spaces, and the framed embeddings of the moduli spaces into such cornered Euclidean spaces provide homotopically coherent attaching instructions.

We will not need to delve into any greater detail here, because (following [LS14a, LLS20b]) our framed flow category will actually be built upon an already existing one.

Definition 3.7. For any fixed n , the *cube flow category* of dimension n is $\mathfrak{F}_{\mathcal{Q}}$ whose objects are the vertices u of an n -dimensional cube, viewed as elements of $\{0, 1\}^n$, and whose non-empty moduli spaces $\mathcal{M}_{\mathcal{Q}}(v, u)$ are permutohedra. The grading $h: \mathfrak{F} \rightarrow \mathbb{Z}$ is given by $h(u) = \sum_i u_i$. For more precise details, see [LS14a, Definition 4.1] where $\mathfrak{F}_{\mathcal{Q}}$ is also described via a specific Morse function on \mathbb{R}^n , or [LLS20b, Section 3.4] for a combinatorial description.

Remark 3.8. In [LS14a], the notation \mathcal{C} is used for flow categories. We choose to use \mathfrak{F} in order to reserve \mathcal{C} for collections of circles in a diagram.

Proposition 3.9. ([LS14a, Section 4], [LLS20b, Section 3.6]) *The cube flow category $\mathfrak{F}_{\mathcal{Q}}$ is a framed flow category. In particular, its moduli spaces can be neatly embedded (with framing) into cornered Euclidean spaces in a coherent fashion.*

The proof of Proposition 3.9 does involve a choice of *sign assignment* for the edges of the cube which correspond to a choice of signs making a cube-shaped chain complex anti-commute so that $d^2 = 0$ (see [LS14a, Proposition 4.12]). However, different choices of sign assignments give rise to stably equivalent spectra, in a lift of the corresponding argument for chain complexes [LS14a, Proposition 6.1].

In this framework, [LS14a] show how to refine the Khovanov chain complex of a link $L \subset S^3$ with diagram D into a flow category $\mathfrak{F}_{Kh}(D)$ that “trivially covers” $\mathfrak{F}_{\mathcal{Q}}$ in the following sense.

Definition 3.10. ([LLS20b, Definition 3.21]) A *cubical flow category* is a flow category \mathfrak{F} which comes equipped with a graded functor $f : \mathfrak{F} \rightarrow \mathfrak{F}_{\mathcal{Q}}$ to some fixed cube category such that for all $x, y \in \mathfrak{F}$, the map $\mathcal{M}_{\mathfrak{F}}(y, x) \rightarrow \mathcal{M}_{\mathcal{Q}}(f(y), f(x))$ is a trivial covering map.

Proposition 3.11. *Any cubical flow category \mathfrak{F} can be upgraded to a framed flow category.*

Proof. The covering maps $\mathcal{M}_{\mathfrak{F}} \rightarrow \mathcal{M}_{\mathcal{Q}}$ can be composed with the embeddings of $\mathcal{M}_{\mathcal{Q}}$ into cornered Euclidean space, and all of the framings are then inherited. These framings also provide a consistent manner in which to separate the components of the trivial covers. In [LLS20b] such embeddings are called *cubical neat embeddings*; see Section 3.6 in that reference for more details. \square

Remark 3.12. Alternatively, one may follow the arguments in Sections 3.1 and 3.2 of [LS14a] to upgrade a cubical flow category into a framed flow category. Our compositions of covering maps and embeddings give rise to *neat immersions* with coherent framings in the language employed in that reference, which can be perturbed to give *neat embeddings*. See also Section 3.4.1 in that reference.

Theorem 3.13. ([LS14a, LLS20b]) *Given a link $L \subset S^3$ with diagram D , there is a cubical flow category $\mathfrak{F}_{Kh}(D)$ which refines the Khovanov complex $CKh^*(D)$. The resulting suspension spectrum $\mathcal{X}_{Kh}(D) = \Sigma^\infty |\mathfrak{F}_{Kh}(D)|$ is an invariant of the link L up to stable homotopy equivalence, allowing the notation $\mathcal{X}_{Kh}(L)$ for the spectrum.*

One crucial feature of the Khovanov chain complex used in the proof of Theorem 3.13 is that, before sign assignment, all signs in the differential are positive. Thus the commutation of any square face in the Khovanov cube corresponds to a *bijection* of 0-dimensional moduli spaces coming from the edges of the square. Such bijections correspond to trivial covers of the 1-dimensional permutohedron (which is an interval), allowing \mathfrak{F}_{Kh} to satisfy the requirements of Definition 3.10 for the 1-dimensional moduli spaces. In our situation however, this feature will not be present.

Remark 3.14. It should be noted that the odd Khovanov differential used to define the odd Khovanov complex $CKh_o^*(D)$ in [ORS13] does include signs within the cube before the sign assignment. In [SSS20], the authors used the language of signed correspondences in a suitable Burnside category to resolve this issue. In terms of our language here, \mathfrak{F}_{Kh_o} could be thought of as a “signed” cubical flow category, where the various covering maps carry signs which indicate the need to reverse framing on certain components of the given cover. Despite this need for signs, it is shown in [SSS20] that the square faces in the odd Khovanov cube still correspond to (signed) bijections, which give (signed) trivial covers of intervals as needed.

The rest of the arguments in [LS14a, LLS20b] used to construct \mathfrak{F}_{Kh} are less immediately relevant, so we simply summarize. The moduli spaces are built inductively. The 1-dimensional moduli spaces in \mathfrak{F}_{Kh} corresponding to square faces involve a choice of bijection between the compositions of 0-dimensional moduli spaces coming from the two ways to traverse the edges in a square. For all but one case, these two compositions are either both empty or both consist of one point, so the bijection is determined. The remaining nontrivial case is known as the *ladybug configuration* [LS14a, Figure 5.1], where the two 0-dimensional moduli spaces each have two points. In this situation, the choice is the *ladybug matching* ([LS14a, Section 5.4]). The 2-dimensional \mathcal{M}_{Kh} are then bounded by some cover of $\partial_{[1]}\mathcal{M}_{\mathcal{Q}}$ (a hexagon, topologically a circle S^1). The hexagon relation ([LS14a, Section 5.5]) is the check that this cover is indeed the trivial cover, so that the 2-dimensional \mathcal{M}_{Kh} can be chosen to be a trivial cover of $\mathcal{M}_{\mathcal{Q}}$. The higher dimensional $\partial\mathcal{M}_{\mathcal{Q}}$ are all topologically spheres of higher dimension, and so have only trivial covers, allowing the construction of all higher dimensional \mathcal{M}_{Kh} from here.

If we instead let $L \subset \mathbb{A} \times I$ be an annular link with diagram D , then the methods in [LS14a, LLS20b] extend in a straightforward way to build a stable homotopy refinement of annular Khovanov homology. A formulation using the language of Burnside functors is given in [AKW, Section 4.3]. The correspondence between Burnside functors and cubical flow categories [LLS20b, Section 4.3] yields a flow category $\mathfrak{F}_{\mathbb{A}}(D)$ refining $CKh_{\mathbb{A}}(D)$. The associated suspension spectrum

$$(20) \quad \mathcal{X}_{\mathbb{A}}(D) = \Sigma^{\infty}|\mathfrak{F}_{\mathbb{A}}(D)|$$

is an invariant of L up to stable homotopy equivalence. An alternative construction can be found in [LLS20a].

Remark 3.15. In [LS14a], there is a global choice between two types of ladybug matchings, the so-called *left pair* and *right pair*. By [LS14a, Proposition 6.5], the two spectra associated to these choices are stably equivalent. To the authors' knowledge, it is not known if the annular spectrum is independent of the choice of ladybug matching; in particular the argument in [LS14a, Proposition 6.5] does not hold for links in $\mathbb{A} \times I$. For the remainder of this paper we fix one of the two choices.

3.2. Constructing a map between Khovanov spectra. Given an annular link diagram D , each of the generators E, F , and $H \in \mathfrak{sl}_2$ yield chain maps $CKh_{\mathbb{A}}(D) \rightarrow CKh_{\mathbb{A}}(D)$. The goal of this paper is to lift these endomorphisms to maps of spectra. This section discusses our general strategy for lifting maps of cochain complexes to maps of spectra.

Definition 3.16. Let \mathfrak{F} be a flow category with grading $h: \mathfrak{F} \rightarrow \mathbb{Z}$. The *suspension* of \mathfrak{F} , denoted $\Sigma\mathfrak{F}$, is the flow category whose underlying category is identical to \mathfrak{F} but whose grading Σh is shifted up by one,

$$\Sigma h(x) = h(x) + 1.$$

It is straightforward to verify that if \mathfrak{F} is a framed flow category, then $\Sigma\mathfrak{F}$ inherits a natural framing, and moreover there is a homeomorphism

$$|\Sigma\mathfrak{F}| \cong \Sigma|\mathfrak{F}|.$$

Consider a map $A^* \xrightarrow{f} B^*$ of cochain complexes, where A^* and B^* are refined by framed flow categories $\mathfrak{F}(A)$ and $\mathfrak{F}(B)$ respectively. In order to lift f to a map of cell complexes

$|\mathfrak{F}(B)| \xrightarrow{F} |\mathfrak{F}(A)|$ satisfying $F^* = f$ on cohomology (up to a shift), the strategy is to refine the complex $\text{Cone } f$ into a framed flow category $\mathfrak{F}(\text{Cone } f)$. The key point is that, by construction, $\mathfrak{F}(\text{Cone } f)$ will contain complementary downward and upward closed subcategories ([LS14a, Section 3.4]) corresponding to $\mathfrak{F}(A)$ and $\Sigma\mathfrak{F}(B)$ respectively. The downward closed subcategory yields a subcomplex $|\mathfrak{F}(A)| \subset |\mathfrak{F}(\text{Cone } f)|$ whose quotient complex is canonically identified with $|\Sigma\mathfrak{F}(B)| = \Sigma|\mathfrak{F}(B)|$. The Puppe sequence yields

$$|\mathfrak{F}(A)| \hookrightarrow |\mathfrak{F}(\text{Cone } f)| \twoheadrightarrow |\mathfrak{F}(\text{Cone } f)|/|\mathfrak{F}(A)| = \Sigma|\mathfrak{F}(B)| \xrightarrow{\star} \Sigma|\mathfrak{F}(A)| \rightarrow \dots$$

where \star is the Puppe map. After shifting degrees, the Puppe map is the desired F .

Now we note that, if $A^* = B^* = CKh^*(D)$ for a link diagram D having n crossings and f is assembled from maps $\mathcal{F}_{\mathbb{A}}(D_u) \rightarrow \mathcal{F}_{\mathbb{A}}(D_u)$ on each smoothing, then the complex $\text{Cone } f$ takes the shape of an $(n+1)$ -dimensional cube \mathcal{Q} . If in addition the map f contains no signs, then one can seek to build a cubical flow category $\mathfrak{F}_{Kh}(\text{Cone } f)$ covering $\mathfrak{F}_{\mathcal{Q}}$ as before. Assuming this can be done, the resulting Puppe map F induces a map of spectra $\mathcal{X}(D) \xrightarrow{\Sigma^\infty F} \mathcal{X}(D)$ which lifts f as desired. All of this is illustrated by the exact triangle for a crossing ([LS14a, Theorem 2]) which is also used in [LS14b, Section 3.3] to define the map induced by a saddle cobordism between links.

In our situation, D is an annular link and f is one of E, F , or H . Moreover, as illustrated in Section 4, these maps contain signs leading to cancellations, which prevents $\mathfrak{F}(\text{Cone } f)$ from being made into a cubical flow category. Nevertheless, we will refine $\text{Cone } f$ and lift E, F , and H by using the above procedure.

Remark 3.17. All of the constructions thus far can be converted into the language of functors from cube categories to Burnside categories. This is the framework introduced in [LLS20b] to investigate disjoint unions and connect sums of links, and it provides a convenient combinatorial approach to constructing stable homotopy types, as in [SSS20] for the odd Khovanov spectrum. In this language, chain maps are lifted as natural transformations of Burnside functors. In a previous paper on annular Khovanov spectra [AKW], we used (an equivariant version of) this approach to refine the quantum annular Khovanov homology constructed in [BPW19] and the action coming from the generators $K, K^{-1} \in U_q(\mathfrak{sl}_2)$ on quantum annular Khovanov homology. However, as we shall see, this viewpoint is not sufficient to refine the action of the remaining generators of $U_q(\mathfrak{sl}_2)$. See [AKW, Section 8] for a related discussion.

4. NEW FEATURES: CANCELLATIONS AND TOPOLOGY OF MODULI SPACES

This section considers a particular example to illustrate new complexities that appear in the analysis of the \mathfrak{sl}_2 action. These features motivate the introduction of new techniques in follow-up sections.

Consider the 2-crossing annular knot diagram and the associated surgery arc diagram (see Section 2.3) in Figure 7. The cube of resolutions for this link diagram (to be more precise, a square since there are two crossings) is given in Figure 8; the coordinate directions are labeled in accordance with the numbering of the surgery arcs in Figure 7b.

The cube-shaped chain complex $\text{Cone}(E)$ is shown in Figure 9. This cube consists of two copies (pictured horizontally) of the cube of resolutions from Figure 8; the action of E takes place along the vertical axis. Each diagram in the cube is endowed with past/future arcs as in Definition 2.2.

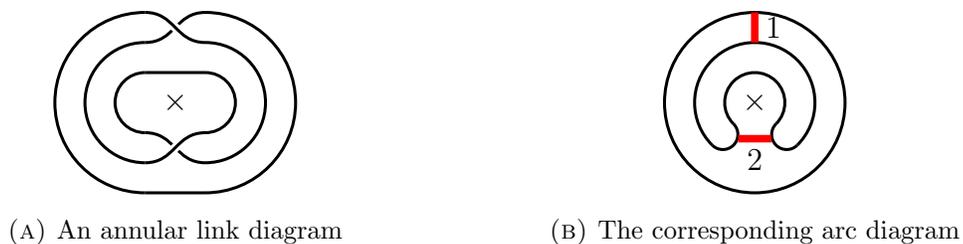


FIGURE 7

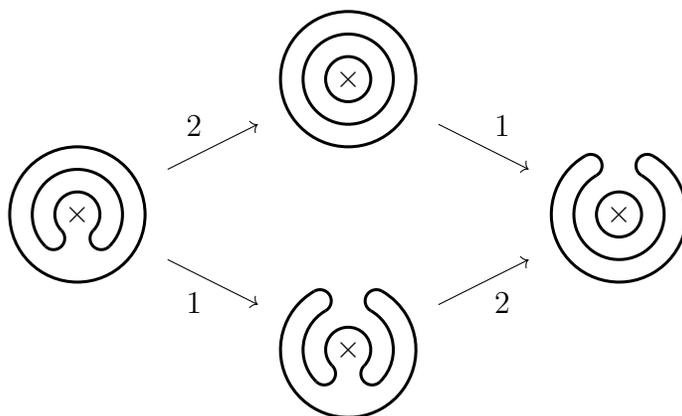


FIGURE 8. The cube of resolutions for the link in Figure 7a

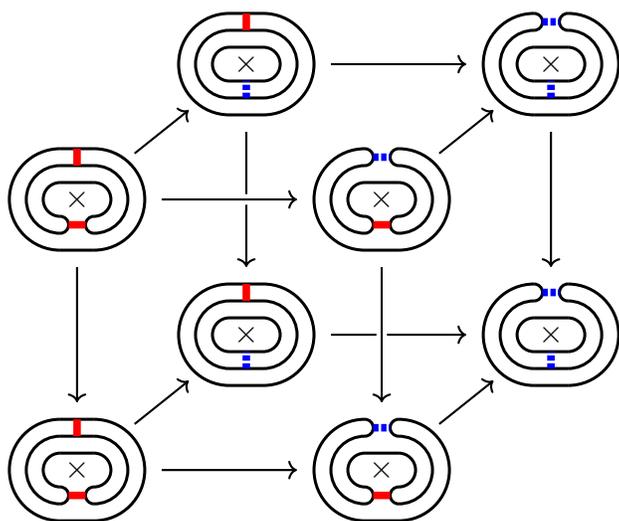


FIGURE 9. The cube-shaped complex associated to E

Consider a particular generator, $1 \otimes v_-$, at the upper left vertex of the cube. We use the computations from Section 2.4 to arrive at the following diagram which summarizes the cube edge maps (we suppress the tensor product notation for the v factors, for brevity).

$$(21) \quad \begin{array}{ccccc} & & v_+v_-v_- + v_-v_+v_- & \xrightarrow{1} & X \otimes v_- \\ & \nearrow 2 & \downarrow E & \nearrow 2 & \downarrow E \\ 1 \otimes v_- & \xrightarrow{1} & v_- & & \\ \downarrow E & & \downarrow E & & \\ & \nearrow 2 & v_+v_-v_+ + v_-v_+v_+ & \xrightarrow{1} & X \otimes v_+ \\ 1 \otimes v_+ & \xrightarrow{1} & v_+ & & \end{array}$$

A crucial feature in this example is *cancellation*. It is not seen in the result of the algebraic calculation in the cubical diagram above, and we will explain it now in more detail. The generator $1 \otimes v_-$ is sent by the edge map labeled 2 to $v_+v_-v_- + v_-v_+v_-$. The value of the E map on these two summands is given as follows:

$$(22) \quad E(v_+v_-v_-) = v_+v_-v_+ - v_+v_+v_-$$

$$(23) \quad E(v_-v_+v_-) = v_-v_+v_+ + v_+v_+v_-$$

Considering the composition $E \circ 2$, the terms $-v_+v_+v_-$, $v_+v_+v_-$ in equations (22), (23) cancel; this is an instance of cancellation that is the root of the complexity of this problem as we explain next.

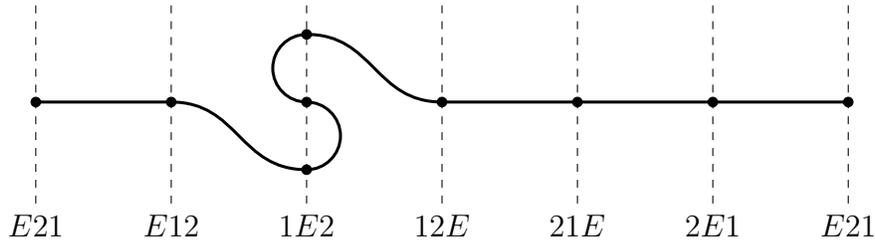


FIGURE 10

A relevant part of the moduli spaces is illustrated in Figure 10. The vertical dashed lines are labeled by compositions of edges of length three in the cube, and the corresponding 1-morphisms are indicated by thickened dots. For example, the effect of the edge path $E21$ on the generator $1 \otimes v_-$ is

$$1 \otimes v_- \xrightarrow{1} v_- \xrightarrow{2} X \otimes v_- \xrightarrow{E} X \otimes v_+$$

The resulting single flow line is represented by a dot above the label $E21$ in Figure 10. There is a single flow line associated to the edge path $E12$ as well, since the summand $v_+v_-v_-$ is sent to zero by the edge map labeled 1. The interesting case in this example is the edge path $1E2$: there are three flow lines

$$\begin{aligned} 1 \otimes v_- &\longmapsto v_+v_-v_- \longmapsto v_+v_-v_+ \longmapsto X \otimes v_+ \\ 1 \otimes v_- &\longmapsto v_+v_-v_- \longmapsto -v_+v_+v_- \longmapsto -X \otimes v_+ \end{aligned}$$

$$1 \otimes v_- \longmapsto v_-v_+v_- \longmapsto v_+v_+v_- \longmapsto X \otimes v_+$$

represented by the three dots above the label $1E2$. (Note that $v_-v_+v_+$ is sent to zero by the map 1, so this term does not contribute a flow line.) The rest of the edge paths give rise to a single flow line.

The vertical strips between dashed lines in Figure 10 represent 2-morphisms, or in other words 1-dimensional moduli spaces corresponding to square faces in the cube $\text{Cone}(E)$. Restricting to the square face in the back of the cube (21), the flow line

$$v_-v_+v_- \longmapsto X \otimes v_- \longmapsto X \otimes v_+$$

along the edge path $E1$ corresponds to the flow line

$$v_-v_+v_- \longmapsto v_+v_+v_- \longmapsto X \otimes v_+$$

along $1E$. This explains the matching of the flow lines in the 2-morphism $E12 - 1E2$, i.e. the interval that goes across that vertical strip in Figure 10.

On the other hand, $v_+v_-v_-$ is sent to zero by the map 1, and thus it does not contribute to a flow line along $E1$. The same generator $v_+v_-v_-$ is sent by the map E to a pair of generators in equation (22) which then cancel under the map 1. Thus the 2-morphism $E12 - 1E2$ has to “match” the empty set at $E12$ to two canceling flow lines at $1E2$. It is natural to represent this pair of flow lines by oppositely framed points in the interval labeled by $1E2$, and to connect them by a framed “turnback” in the strip $E12 - 1E2$.

A similar analysis along the square face on the left of the cube forces the matching of the flow lines under the 2-morphism $1E2 - 12E$, creating the “cubic” shape of the 1-dimensional moduli space indicated in Figure 10.

This example shows that bijections across square faces, used in [LS14a] in the setting of Khovanov homology, are insufficient in the presence of cancellations. On the other hand, 2-dimensional faces of the cube force the matching of flow lines associated to edge paths by framed 1-dimensional moduli spaces. As indicated in the introduction, in general building higher-dimensional moduli spaces is a rather non-trivial problem. Our solution in this paper is based on representing the moduli spaces as codimension-1 submanifolds, as illustrated in Figure 10. A subtlety in this approach is that such a codimension-1 position of the flow lines (0-dimensional moduli spaces) in the interval has to be prescribed for all edge paths in a given cube-shaped chain complex $\text{Cone}(E)$. Indeed, these flow lines for the edge paths of arbitrary lengths are a part of the inductive construction of higher-dimensional moduli spaces, based on Definition 3.1. Moreover, given any two generators at two arbitrary vertices u, v in the cube, the positions in the interval for all edge paths connecting u, v in the cube have to be consistent, in the sense that they can be connected by the prescribed framed 1-dimensional moduli spaces in codimension-1 without intersections. The analysis that enables a solution to this problem is the content of Sections 5 - 7.

The construction of higher-dimensional moduli spaces in Section 8 proceeds by induction. For example, in the basic case considered above, the 2-dimensional moduli space has to be constructed for the two chosen generators $1 \otimes v_-$, $X \otimes v_+$ in the cube (21). Identifying the leftmost and rightmost vertical intervals in Figure 10 gives a boundary condition for the 2-dimensional moduli space in a hexagon times an interval, see Figure 19. In general, the higher-dimensional inductive step amounts to finding a null-cobordism in a permutohedron times an interval I , as explained in Section 8.

5. CONSTRUCTING THE FRAMED FLOW CATEGORY: THE SET-UP AND THE BASE CASE

We begin by fixing an n -crossing annular link diagram D , and we let J denote one of the maps E , F , or H . As discussed in Section 3.2, the chain map J on the annular Khovanov complex $CKh_{\mathbb{A}}(D)$ gives rise to Cone J , a complex in the shape of an $(n + 1)$ -dimensional cube \mathcal{Q} . Our goal is to build a framed flow category $\mathfrak{F}_J(D)$ (we will often omit the diagram D from the notation) that refines Cone J with the help of $\mathfrak{F}_{\mathcal{Q}}$, the framed flow category for the cube \mathcal{Q} .

5.1. Definitions and the overall strategy. To begin with we set some notation and basic definitions. The vertices of the cube \mathcal{Q} will use labels u, v, w , with

$$u = (u', i) \in \{0, 1\}^n \times \{0, 1\}.$$

Here we are viewing the first n coordinates u' as indicating where in $CKh_{\mathbb{A}}(D)$ the corresponding vertex lies, while the last coordinate i indicates whether we are in the domain or range of the J map.

Definition 5.1. Given a vertex $u = (u', i) \in \mathcal{Q}$, the *resolution diagram* D_u is the resolution of our link diagram D at the vertex u' in the $CKh_{\mathbb{A}}(D)$. We further define $\Gamma_{Kh}(u)$ to be the set of Khovanov generators for D_u ; see Equation (3).

Definition 5.2. An m -dimensional sub-cube is a choice of two vertices

$$u \leq_m v$$

such that u and v match in all but m indices where u has entry zero and v has entry one. A 1-dimensional sub-cube is also called an *edge*. We say that a subcube $u \leq_m v$ *involves* J if $u = (u', 0)$ and $v = (v', 1)$ (so that some edges corresponding to the J -map are included).

Definition 5.3. For any sub-cube $u \leq_m v$, the *cube moduli space* $\mathcal{M}_{\mathcal{Q}}(v, u)$ is by definition the $(m - 1)$ -dimensional permutohedron (see Definition 3.7). This is a manifold with corners (for example, if $m = 3$ this is a hexagon) that is topologically equivalent to D^{m-1} , whose codimension i boundary (for all $i = 1, \dots, m - 1$) is identified with i -fold “composition” moduli spaces

$$\partial_{[i]} \mathcal{M}_{\mathcal{Q}}(v, u) \cong \coprod \mathcal{M}_{\mathcal{Q}}(v, w_i) \times \cdots \times \mathcal{M}_{\mathcal{Q}}(w_1, u),$$

where the disjoint union is taken over all sequences of vertices $u \leq w_1 \leq \cdots \leq w_i \leq v$. For more details, see [LS14a, Definition 4.1] and [LLS20b, Section 3.4].

Remark 5.4. Note the reversal of direction. For vertices u and v with $u \leq_1 v$, there is an edge from u to v in the cube of resolutions. However, moduli spaces are directed from v to u .

As described in Section 3.1, the original construction [LS14a] used trivial coverings of the moduli spaces in the cube flow category (see Definition 3.10). The example of Section 4 shows that this framework alone is insufficient for building our desired \mathfrak{F}_J . Instead, we will need to keep track of signs of the J map with the help of framings. First we give ourselves an extra dimension to allow for meaningful framings to occur.

Definition 5.5. For any sub-cube $u \leq_m v$, the *cube ambient space* $\mathcal{E}(v, u)$ is defined depending on whether or not the J map was involved in $u \leq_m v$.

$$(24) \quad \mathcal{E}(v, u) := \begin{cases} \mathcal{M}_{\mathcal{Q}}(v, u) & \text{if } J \text{ is not involved} \\ \mathcal{M}_{\mathcal{Q}}(v, u) \times (0, 1) & \text{if } J \text{ is involved} \end{cases}$$

In either case, $\mathcal{E}(v, u)$ is a manifold with corners whose boundary structure is determined by that of $\mathcal{M}_{\mathcal{Q}}(v, u)$. In our figures throughout the paper, we will draw the $(0, 1)$ -direction horizontally.

Next we introduce the types of maps that we will be using in place of the trivial covering maps utilized in [LS14a] via Definition 3.10.

Definition 5.6. A *thick embedding* is a framed map $f = e \circ t$ where t is a trivial covering and e is an embedding. We will denote thick embeddings with the symbol \hookrightarrow . If t is the trivial identity cover (i.e. the identity map), we call f a *thin embedding* and denote with the usual symbol \hookrightarrow . Thus a thin embedding is just a framed embedding.

Remark 5.7. Note that according to the definition, different sheets in a thick embedding may have different framings. In codimension 1, such framings are either identical or are reversed. As we show later in the paper, the former case occurs for $J = E, F$ and the latter case occurs for $J = H$.

With all of this notation in place, we can state our main construction as a theorem.

Theorem 5.8. *Fix an n -crossing annular link diagram D and an \mathfrak{sl}_2 generator $J = E, F$, or H giving rise to a chain map on annular Khovanov complexes with mapping cone $\text{Cone } J$. Let \mathcal{Q} denote the $(n + 1)$ -dimensional cube, with corresponding framed flow category $\mathfrak{F}_{\mathcal{Q}}$. Then there exists a flow category $\mathfrak{F}_J = \mathfrak{F}_J(D)$ together with a “thick embedding functor” $f : \mathfrak{F}_J \rightarrow \mathfrak{F}_{\mathcal{Q}}$ satisfying the following properties.*

- (1) *The objects of \mathfrak{F}_J are in bijection with the Khovanov generators throughout the complex $\text{Cone } J$:*

$$\text{Ob}(\mathfrak{F}_J) := \coprod_{u \in \mathcal{Q}} \Gamma_{Kh}(u),$$

with grading given by homological degree in $\text{Cone } J$, and graded functor $f : \text{Ob}(\mathfrak{F}_J) \rightarrow \text{Ob}(\mathfrak{F}_{\mathcal{Q}})$ being the obvious covering map of sets

$$f(x) = u \quad \Leftrightarrow \quad x \in \Gamma_{Kh}(u).$$

- (2) *For any generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$ such that $u \not\leq v$, the moduli space $\mathcal{M}_J(y, x)$ is the empty set.*
- (3) *Given a sub-cube $u \leq_m v$ and a pair of generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$, the moduli space $\mathcal{M}_J(y, x)$ is an $(m - 1)$ -dimensional manifold with corners, whose codimension i boundary is identified with i -fold compositions*

$$\partial_{[i]} \mathcal{M}_J(y, x) \cong \coprod \mathcal{M}_J(y, z_i) \times \cdots \times \mathcal{M}_J(z_1, x)$$

where the disjoint union is taken over all tuples of generators $(z_1, \dots, z_i) \in \Gamma_{Kh}(w_1) \times \cdots \times \Gamma_{Kh}(w_i)$ for sequences of vertices $u \leq w_1 \leq \cdots \leq w_i \leq v$ in the subcube $u \leq_m v$.

- (4) *Each moduli space $\mathcal{M}_J(y, x)$ of item (3) comes equipped with a thick embedding*

$$f : \mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, u)$$

such that, for any tuple $(z_1, \dots, z_i) \in \Gamma_{Kh}(w_1) \times \cdots \times \Gamma_{Kh}(w_i)$ giving rise to a component of $\partial_{[i]} \mathcal{M}_J(y, x)$, the following diagram of thick embeddings (using the product framing as necessary) commutes

$$(25) \quad \begin{array}{ccc} \mathcal{M}_J(y, x) & \xleftarrow{f} & \mathcal{E}(v, u) \\ \uparrow & & \uparrow \\ \partial_{[i]} \mathcal{M}_J(y, x) & \xleftarrow{\quad} & \partial_{[i]} \mathcal{E}(v, u) \\ \uparrow & & \uparrow \\ \mathcal{M}_J(y, z_i) \times \cdots \times \mathcal{M}_J(z_1, x) & \xleftarrow{f \times \cdots \times f} & \mathcal{E}(v, w_i) \times \cdots \times \mathcal{E}(w_1, u) \end{array}$$

(5) In the case $m = 1$, the 0-dimensional moduli spaces correspond to the components of the differential in Cone J .

Item (3) ensures that the moduli space boundaries of \mathfrak{F}_J are arranged in the same way as those of \mathfrak{F}_Q , while item (4) ensures that this arrangement is respected by the thick embeddings as well. Meanwhile item (5) refers to the edges of the cube, where the moduli spaces are all thickly embedded points in either D^0 or $D^0 \times (0, 1)$ corresponding to either the annular Khovanov differential in $CKh_{\mathbb{A}}(D)$ or the J map, respectively.

Recall the annular Khovanov spectrum $\mathcal{X}_{\mathbb{A}}(D)$ from Equation (20) and the surrounding discussion.

Corollary 5.9. *The flow category $\mathfrak{F}_J(D)$ defined above can be upgraded to a framed flow category which refines the complex Cone J , giving rise to a map of spectra $\mathcal{J} : \mathcal{X}_{\mathbb{A}}(D) \rightarrow \mathcal{X}_{\mathbb{A}}(D)$ whose induced map \mathcal{J}^* on cohomology is equal to the action of J .*

Proof. The cube flow category \mathfrak{F}_Q is a framed flow category, so each $\mathcal{M}_Q(v, u)$ comes equipped with a framed neat embedding $\mathcal{M}_Q(v, u) \hookrightarrow \mathbb{E}_Q(v, u)$ where $\mathbb{E}_Q(v, u)$ denotes a suitable cornered Euclidean space (see [LLS20b, Section 3.6]). As in Definition 5.5, we set

$$(26) \quad \mathbb{E}(v, u) := \begin{cases} \mathbb{E}_Q(v, u) & \text{if } J \text{ is not involved} \\ \mathbb{E}_Q(v, u) \times \mathbb{R} & \text{if } J \text{ is involved} \end{cases}$$

We then have embeddings $\mathcal{E}(v, u) \hookrightarrow \mathbb{E}(v, u)$, where $(0, 1)$ embeds into \mathbb{R} in the standard way when J is involved. We then compose with our thick embeddings provided by Item (4) to give

$$\mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, u) \hookrightarrow \mathbb{E}(v, u).$$

Just as in Proposition 3.11 for the cubical flow categories of [LLS20b], the framing inherited by this composition provides us with a consistent way to separate the various components of each trivial cover, giving us framed neat embeddings $\mathcal{M}_J(y, x) \hookrightarrow \mathbb{E}(v, u)$ which respect the necessary boundary conditions as required (note that crossing with \mathbb{R} in Equation (26) does not affect the cornered structure of $\mathbb{E}_Q(v, u)$). Item (5) then ensures our resulting framed flow category $\mathfrak{F}_J(D)$ has refined Cone J . This gives rise to the desired map as described in Section 3.2. \square

The rest of this section together with Sections 6, 7, and 8 are devoted to proving Theorem 5.8. Items (3) and (4) of the theorem are the most crucial aspects of the construction, and the various thickly embedded moduli spaces will be built inductively. If the sub-cube $u \leq_m v$ does not involve the J map, then it corresponds to a sub-cube in the cube of resolutions of D . In this case the thick embedding $f : \mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, u) = \mathcal{M}_Q(v, u)$ will be a trivial cover built exactly as in [LS14a] (based on the annular differential rather than

the traditional Khovanov differential). If $u \leq_m v$ involves the J map, then $\mathcal{M}_J(y, x)$ will be thickly embedded as a codimension 1 submanifold of $\mathcal{E}(v, u)$. We handle the base case below.

5.2. 0-dimensional moduli spaces for edges. Fix $u \leq_1 v$, $x \in \Gamma_{Kh}(u)$, and $y \in \Gamma_{Kh}(v)$. We begin by defining the 0-dimensional moduli space $\mathcal{M}_J(y, x)$.

Suppose first that $u \leq_1 v$ is not the J edge. Then it corresponds to an edge $u \rightarrow v$ in the cube of resolutions of D . Let

$$\psi_{u,v}: \mathcal{F}_{\mathbb{A}}(D_u) \rightarrow \mathcal{F}_{\mathbb{A}}(D_v)$$

denote the map assigned to the saddle cobordism $D_u \rightarrow D_v$, and write

$$\psi_{u,v}(x) = \sum_{z \in \Gamma_{Kh}(v)} k_z z$$

where each $k_z \in \{0, 1\}$. Say that y appears in $\psi_{u,v}(x)$ if the above coefficient k_y is nonzero. If y appears in $\psi_{u,v}(x)$, then $\mathcal{M}_J(y, x)$ is defined to consist of a single element, which we denote

$$x \rightarrow y.$$

If y does not appear in $\psi_{u,v}(x)$, set $\mathcal{M}_J(y, x) = \emptyset$.

Now suppose that $u \leq_1 v$ is the J map. Let s denote the number of essential circles in $D_u = D_v$. We distinguish two cases. First, if J is one of E or F , again write

$$Jx = \sum_{z \in \Gamma_{Kh}(v)} \ell_z z$$

where now each $\ell_z \in \{-1, 0, 1\}$, and say that y appears in Jx if $\ell_y \neq 0$. If y appears in Jx , then it is obtained from x by the action of J on an essential circle in D_u . Let $1 \leq i \leq s$ denote the position of this essential circle, and define $\mathcal{M}_J(y, x)$ to consist of a single element, which we denote

$$x \xrightarrow{i} y.$$

If y does not appear in Jx , then set $\mathcal{M}_J(y, x) = \emptyset$.

Finally, if $J = H$, set $\mathcal{M}_H(y, x) = \emptyset$ unless the labels $(v_{\pm}, 1, X)$ on each circle in x and y match, in which case $\mathcal{M}_H(y, x)$ contains a point for every essential circle in $D_u = D_v$. These s points in $\mathcal{M}_H(y, x)$ will be denoted

$$x \xrightarrow{i} y$$

for $1 \leq i \leq s$.

We now describe the thick embeddings, which again depend on the type of edge $u \leq_1 v$. If $u \leq_1 v$ is not the J map, then $\mathcal{E}(v, u)$ is a point, so the map is unique.

Suppose now that $u \leq_1 v$ corresponds to the J map. Let s denote the number of essential circles in $D_u = D_v$. Recall that $\mathcal{E}(v, u) = \mathcal{M}_{\square}(v, u) \times (0, 1) \cong (0, 1)$. We define our thick embedding $\mathcal{M}_J(y, x) \xrightarrow{f} \mathcal{E}(v, u)$ by setting¹

$$(27) \quad f(x \xrightarrow{i} y) := \frac{i}{s+1} \in (0, 1),$$

¹The important data here is the ordering of points in the interval; a specific formula is given just to fix the convention.

with framing determined as follows. If J is either E or F , then the framing is given by the sign of $(-1)^{i+1}$. If $J = H$, then the framing is ± 1 if x (and thus y) is labeled by v_{\pm} on the i -th essential circle in D_u .

Note that all of these are actually thin embeddings when the subcube is 1-dimensional. It follows from the construction of the spaces and thin embeddings that item (5) of Theorem 5.8 is clearly satisfied.

Remark 5.10. Recall from (17) that $Hx = \text{adeg}(x)y$, where the circle labels in y match those in x . It may seem natural to define the 0-dimensional moduli space $\mathcal{M}_H(y, x)$ to contain $|\text{adeg}(x)|$ points, and define the thick embedding to be constant, with framing given by the sign of $\text{adeg}(x)$. In the above definition of $\mathcal{M}_H(y, x)$ we are motivated by the refined formula (18) in which H acts on each essential circle.

Example 5.11. We illustrate the 0-dimensional moduli spaces for an edge corresponding to various J maps when $D_u = D_v$ consists of two essential circles. The thin horizontal segment is the interval $\mathcal{E}(v, u) \cong (0, 1)$, the dot \bullet is the image of $\mathcal{M}_J(y, x)$, and the arrows indicate the framing.

Moduli spaces $\mathcal{M}_E(y, x)$ are depicted below.

$$\begin{array}{cccc}
 \begin{array}{c} \bullet \rightarrow \\ \hline \end{array} & \begin{array}{c} \leftarrow \bullet \\ \hline \end{array} & \begin{array}{c} \bullet \rightarrow \\ \hline \end{array} & \begin{array}{c} \leftarrow \bullet \\ \hline \end{array} \\
 v_- \otimes v_- \xrightarrow{1} v_+ \otimes v_- & v_- \otimes v_- \xrightarrow{2} v_- \otimes v_+ & v_- \otimes v_+ \xrightarrow{1} v_+ \otimes v_+ & v_+ \otimes v_- \xrightarrow{2} v_+ \otimes v_+
 \end{array}$$

For $J = F$ the situation is similar; see also Lemma 2.6 and the discussion surrounding it.

Moduli spaces $\mathcal{M}_H(y, x)$ are shown below. There are four generators of $\mathcal{F}_{\mathbb{A}}(D_u)$, hence four nonempty moduli spaces, each containing two elements of the form $x \xrightarrow{i} y$ for $i = 1, 2$. The numbers above the dots \bullet in the following diagram indicate the image of the point $x \xrightarrow{i} y$.

$$\begin{array}{cccc}
 \begin{array}{c} 1 \quad 2 \\ \leftarrow \bullet \leftarrow \bullet \\ \hline \end{array} & \begin{array}{c} 1 \quad 2 \\ \leftarrow \bullet \rightarrow \bullet \\ \hline \end{array} & \begin{array}{c} 1 \quad 2 \\ \bullet \rightarrow \leftarrow \bullet \\ \hline \end{array} & \begin{array}{c} 1 \quad 2 \\ \bullet \rightarrow \bullet \rightarrow \\ \hline \end{array} \\
 v_- \otimes v_- \xrightarrow{1} v_- \otimes v_- & v_- \otimes v_+ \xrightarrow{1} v_- \otimes v_+ & v_+ \otimes v_- \xrightarrow{1} v_+ \otimes v_- & v_+ \otimes v_+ \xrightarrow{1} v_+ \otimes v_+ \\
 v_- \otimes v_- \xrightarrow{2} v_- \otimes v_- & v_- \otimes v_+ \xrightarrow{2} v_- \otimes v_+ & v_+ \otimes v_- \xrightarrow{2} v_+ \otimes v_- & v_+ \otimes v_+ \xrightarrow{2} v_+ \otimes v_+
 \end{array}$$

The moduli spaces for pairs (x, y) which are not depicted above are all empty.

6. CONSTRUCTING THE 1-DIMENSIONAL MODULI SPACES FOR SQUARE FACES

We now define the 1-dimensional moduli spaces and their thick embeddings, with an eye towards the inductive step in Section 8. Although the combinatorial analysis in this section can be simplified by the more general considerations in Section 7, we present a ‘‘hands-on’’ approach to illustrate the notions in Section 5.1.

Fix $u \leq_2 v$, $x \in \Gamma_{Kh}(u)$, and $y \in \Gamma_{Kh}(v)$. The moduli space $\mathcal{M}_J(y, x)$ will be a disjoint union of closed intervals, so we need to specify which points in $\partial\mathcal{M}_J(y, x)$ bound an interval in $\mathcal{M}_J(y, x)$. Item (3) of Theorem 5.8 requires that these boundary points come from products corresponding to composition along the edges of the subcube. Letting w_1, w_2 denote the two

vertices with $u \leq_1 w_1, w_2 \leq_1 v$, we introduce the notations

$$(28) \quad \begin{aligned} \partial_{w_1} \mathcal{M}_J(y, x) &:= \prod_{z \in \Gamma_{Kh}(w_1)} \mathcal{M}_J(y, z) \times \mathcal{M}_J(z, x), \\ \partial_{w_2} \mathcal{M}_J(y, x) &:= \prod_{z \in \Gamma_{Kh}(w_2)} \mathcal{M}_J(y, z) \times \mathcal{M}_J(z, x) \end{aligned}$$

for these two composition moduli spaces - one for the edge-path passing through w_1 , and the other for the edge-path passing through w_2 . Thus we have

$$\partial \mathcal{M}_J(y, x) = \partial_{w_1} \mathcal{M}_J(y, x) \prod \partial_{w_2} \mathcal{M}_J(y, x).$$

We write elements of each $\partial_{w_i} \mathcal{M}_J(y, x)$ as

$$(29) \quad (x \rightarrow z \rightarrow y) := (z \rightarrow y, x \rightarrow z) \in \mathcal{M}_J(y, z) \times \mathcal{M}_J(z, x)$$

for various $z \in \Gamma_{Kh}(w_i)$. If one of the edges of $u \leq_m v$ corresponds to the J map, we decorate the corresponding arrow in $(x \rightarrow z \rightarrow y)$ according to which essential circle was acted upon as in Section 5.2. Each composition moduli space comes equipped with a thick product embedding

$$(30) \quad \partial_{w_i} \mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, w_i) \times \mathcal{E}(w_i, u) \cong \begin{cases} D^0 & \text{if } J \text{ is not involved} \\ (0, 1) & \text{if } J \text{ is involved} \end{cases}$$

induced by the definitions in Section 5.2. Our goal now is to construct a moduli space with thick embedding

$$\mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, u) \cong \begin{cases} D^1 & \text{if } J \text{ is not involved} \\ D^1 \times (0, 1) & \text{if } J \text{ is involved} \end{cases}$$

which ‘fills’ the thick embeddings (30) in the sense that Equation (25) is satisfied. Naturally, we split into cases depending on whether or not J was involved.

6.1. Square faces involving only the annular Khovanov differential. A subcube $u \leq_2 v$ that does not involve the J edge corresponds to a square face in the cube of resolutions of D , and in this case we proceed exactly as in [LS14a]. From (28) and the analysis in [LS14a], we see that

$$|\partial_{w_1} \mathcal{M}_J(y, x)| = |\partial_{w_2} \mathcal{M}_J(y, x)| =: N \in \{0, 1, 2\}$$

where the number of points N depends on the combinatorics of the square face $D_u \rightarrow D_v$.

Following [LS14a], we define

$$\mathcal{M}_J(y, x) := \prod_S \mathcal{M}_Q(v, u)$$

where S is a set of N elements. If $N \leq 1$, then there is no ambiguity in the above definition. If $N = 2$, then the ladybug matching [LS14a, Section 5.5] is used to determine the pairing of the two pairs of boundary points in $\partial \mathcal{M}_J(x, y)$. To define the thick embedding, recall that

$$\mathcal{E}(v, u) = \mathcal{M}_Q(v, u)$$

in this case, so we let

$$\mathcal{M}_J(y, x) \xrightarrow{\prod \text{id}} \mathcal{E}(v, u)$$

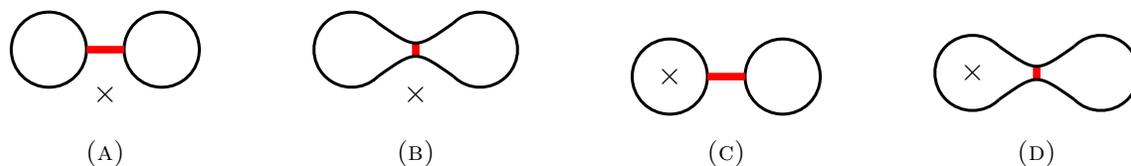


FIGURE 11. Connected arc diagrams where the number of essential circles is unchanged

be a disjoint union of the identity map. In particular, the above thick embedding is a trivial cover requiring no framing data, and this is exactly how the cubical flow category \mathfrak{F}_{Kh} is constructed in [LS14a] (the framing data in [LS14a] is entirely inherited by the neat embedding of the cube moduli spaces $\mathcal{M}_{\mathcal{Q}}(v, u) \hookrightarrow \mathbb{E}_{\mathcal{Q}}(v, u)$; see Section 3.1 and Definition 5.5 there).

6.2. Square faces involving E or F . Now we suppose that the subcube $u \leq_2 v$ involves the $J = E$ map; the case involving $J = F$ is nearly identical (see Lemma 2.6). Without loss of generality, let the edges $u \rightarrow w_1$ and $w_2 \rightarrow v$ be saddle maps S , while $u \rightarrow w_2$ and $w_1 \rightarrow v$ are the E map. In particular, we have equality of resolutions $D_u = D_{w_2}$, $D_{w_1} = D_v$, and the saddles $D_u \rightarrow D_{w_1}$ and $D_{w_2} \rightarrow D_v$ are identical.

$$\begin{array}{ccc} D_u & \xrightarrow{S} & D_{w_1} \\ E \downarrow & & \downarrow E \\ D_{w_2} & \xrightarrow{S} & D_v \end{array}$$

We focus first on connected arc diagrams D_u , and we begin by considering those for which the saddle S does not change the number of essential circles. There are four of these, as shown in Figure 11.

For Figures 11a and 11b, the E map is the zero map, and all of the moduli spaces are empty.

In the case of Figure 11c, we see that all of the moduli spaces are empty unless $x = 1 \otimes v_-$ and $y = v_+$, in which case we have single point composition moduli spaces

$$\partial_{w_1} \mathcal{M}_E(y, x) = \{(1 \otimes v_- \rightarrow v_- \xrightarrow{1} v_+)\},$$

$$\partial_{w_2} \mathcal{M}_E(y, x) = \{(1 \otimes v_- \xrightarrow{1} 1 \otimes v_+ \rightarrow v_+)\}$$

as illustrated in Figure 12a.

Figure 11d is dual to Figure 11c, so the analysis is similar. We have empty moduli spaces unless $x = v_-$ and $y = X \otimes v_+$, in which case we have

$$\partial_{w_1} \mathcal{M}_E(y, x) = \{(v_- \rightarrow X \otimes v_- \xrightarrow{1} X \otimes v_+)\},$$

$$\partial_{w_2} \mathcal{M}_E(y, x) = \{(v_- \xrightarrow{1} v_+ \rightarrow X \otimes v_+)\}$$

In all of these cases where S does not change the number of essential circles, we see that there is a bijection of one-point composition moduli spaces that commutes with the thin

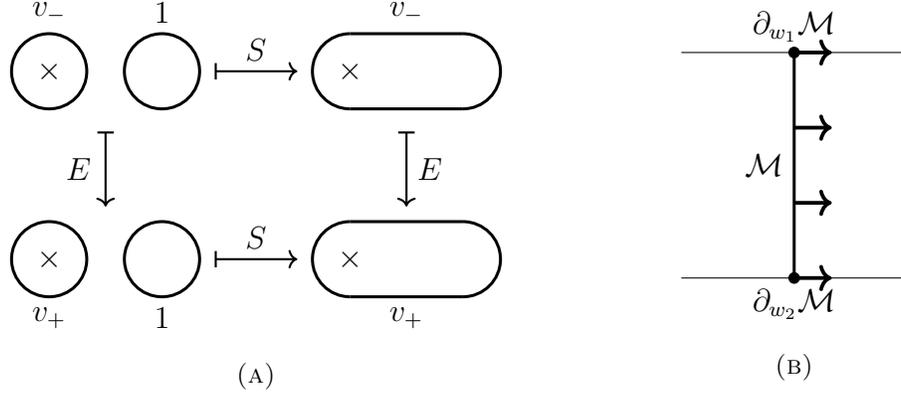


FIGURE 12. The relevant generators in the subcube of resolutions for the case of Figure 11c, together with the corresponding moduli space embedded and framed into $D^1 \times (0, 1)$.



FIGURE 13. Connected arc diagrams where the number of essential circles is changed.

embeddings

$$(31) \quad \begin{array}{ccc} \partial_{w_1}\mathcal{M}_E(y, x) & \hookrightarrow & \mathcal{E}(v, w_1) \times \mathcal{E}(w_1, u) \\ \cong \downarrow & & \parallel \\ & & \mathbb{R} \\ & & (0, 1) \\ & & \parallel \\ \partial_{w_2}\mathcal{M}_E(y, x) & \hookrightarrow & \mathcal{E}(v, w_2) \times \mathcal{E}(w_2, u) \end{array}$$

This commuting bijection allows us to define $\mathcal{M}_E(y, x)$ as a single closed interval D^1 which is thinly embedded as an identity

$$\mathcal{M}_E(y, x) := D^1 \xrightarrow{\text{id} \times \{p_1\}} D^1 \times (0, 1) = \mathcal{E}(v, u)$$

with the constant framing. See Figure 12b for a depiction of the case where S is the saddle in Figure 11c.

Next, we consider the case where S changes the number of essential circles. For this there are two possible connected arc diagrams, illustrated in Figure 13.

In Figure 13a, we see that all moduli spaces are empty unless $x = 1$ and $y = v_+ \otimes v_+$, in which case we have $\partial_{w_1}\mathcal{M}_E(y, x)$ containing two points while $\partial_{w_2}\mathcal{M}_E(y, x)$ is empty (see Figure 14a).

$$\begin{aligned} \partial_{w_1}\mathcal{M}_E(y, x) &= \{(1 \rightarrow v_- \otimes v_+ \xrightarrow{1} v_+ \otimes v_+), (1 \rightarrow v_+ \otimes v_- \xrightarrow{2} v_+ \otimes v_+)\}, \\ \partial_{w_2}\mathcal{M}_E(y, x) &= \emptyset. \end{aligned}$$

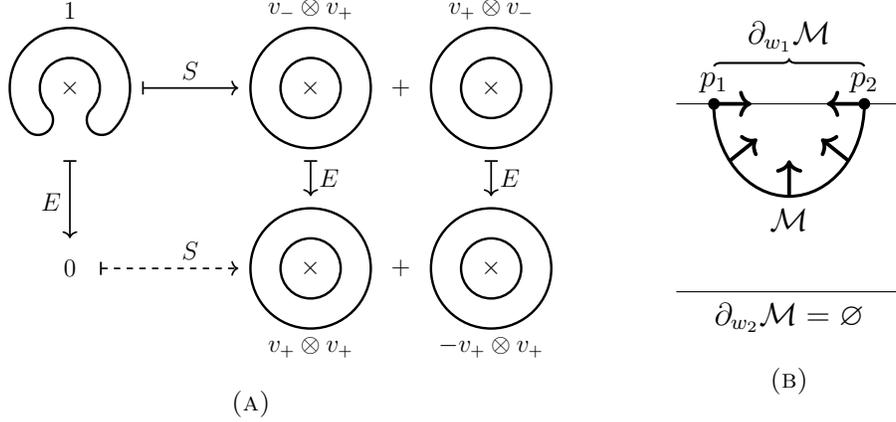


FIGURE 14. The relevant generators in the subcube of resolutions for the case of Figure 13a, together with the corresponding moduli space embedded and framed into $\mathcal{E}(v, u) \cong D^1 \times (0, 1)$.

Following the embedding instructions from Section 5.2, we see that our two points in $\partial_{w_1} \mathcal{M}_E(y, x)$ are embedded in $\mathcal{E}(v, w_1) \times \mathcal{E}(w_1, u) \cong (0, 1)$ with framing shown below.



We then define $\mathcal{M}_E(y, x)$ in this setting to be a single interval which is thinly embedded into $\mathcal{E}(v, u) \cong D^1 \times (0, 1)$ as a framed turnback as indicated in Figure 14b.

The last connected case is dual to the previous one, where the moduli spaces are empty unless $x = v_- \otimes v_-$ and $y = X$, and the roles of $\partial_{w_1} \mathcal{M}_E(y, x)$ and $\partial_{w_2} \mathcal{M}_E(y, x)$ are reversed. The details here are left to the reader.

Finally, we consider the situation when the diagram D_u is disconnected. We let $\mathcal{C} \subset D_u$ denote the connected component of the saddle arc S within D_u , while \mathcal{C}' denotes the complement of \mathcal{C} . As in Section 2.3, these components determine connected components in the resolutions at every vertex, which we also denote by \mathcal{C} and \mathcal{C}' . Recall from Section 2.3 the notion of restriction of generators. This allows us to write generators as $x = x|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}$ and $y = y|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}$, and we use this breakdown to aid in our analysis.

Lemma 6.1. *Suppose some $\partial_{w_i} \mathcal{M}_E(y, x)$ is nonempty. Then precisely one of the two following statements are true.*

- (1) $\text{adeg}(y|_{\mathcal{C}}) = \text{adeg}(x|_{\mathcal{C}}) + 2$ and $\text{adeg}(y|_{\mathcal{C}'}) = \text{adeg}(x|_{\mathcal{C}'})$. This corresponds to E acting on an essential circle in \mathcal{C} .
- (2) $\text{adeg}(y|_{\mathcal{C}}) = \text{adeg}(x|_{\mathcal{C}})$ and $\text{adeg}(y|_{\mathcal{C}'}) = \text{adeg}(x|_{\mathcal{C}'}) + 2$. This corresponds to E acting on an essential circle in \mathcal{C}' .

When $J = F$, the value 2 above is replaced by -2 .

Proof. Saddle maps preserve annular degree, but the E map increases annular degree by two. Similarly, the F map decreases annular degree by two. \square

If $\partial_{w_1} \mathcal{M}_E(y, x) = \partial_{w_2} \mathcal{M}_E(y, x) = \emptyset$, we define $\mathcal{M}_E(y, x) := \emptyset$ as well. Otherwise, we use Lemma 6.1 to consider the two cases.

In the case (1) where E was acting upon the component \mathcal{C} containing S , we must have the labels of $y|_{\mathcal{C}'}$ matching the labels of $x|_{\mathcal{C}'}$, in which case these components are irrelevant

for the analysis of the moduli spaces and we may define

$$\mathcal{M}_E(y, x) := \mathcal{M}_E(y|_{\mathcal{C}}, x|_{\mathcal{C}})$$

as defined and embedded earlier. This case is illustrated for a specific example in Figure 15b.

Meanwhile, in the case (2) where E was acting on some essential circle in \mathcal{C}' disconnected from the saddle S , we see that the generator at w_1 must be $y|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}$ while the generator at w_2 must be the ‘opposite’ $x|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}$. This allows us to show that the maps S and E (and their corresponding edge moduli spaces) commute in the following sense:

$$\begin{aligned} \partial_{w_1} \mathcal{M}_E(y, x) &= \mathcal{M}_E(y|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}, y|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}) \times \mathcal{M}_E(y|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}, x|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}) \\ &\cong \mathcal{M}_E(y|_{\mathcal{C}'}, x|_{\mathcal{C}'}) \times \mathcal{M}_E(y|_{\mathcal{C}}, x|_{\mathcal{C}}) \\ (32) \quad &\cong \mathcal{M}_E(y|_{\mathcal{C}}, x|_{\mathcal{C}}) \times \mathcal{M}_E(y|_{\mathcal{C}'}, x|_{\mathcal{C}'}) \\ &\cong \mathcal{M}_E(y|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}, x|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}) \times \mathcal{M}_E(x|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}, x|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}) \\ &= \partial_{w_2} \mathcal{M}_E(y, x). \end{aligned}$$

The framed embedding data for such moduli spaces is determined by which circle is being acted on by E , and this is maintained throughout all of the equivalences above. Therefore we again have a (one point) bijection $\partial_{w_1} \mathcal{M}_E(y, x) \cong \partial_{w_2} \mathcal{M}_E(y, x)$ that commutes with the thin embeddings as in Equation (31), allowing us to once again define $\mathcal{M}_E(y, x)$ as

$$\mathcal{M}_E(y, x) := D^1 \xrightarrow{I \times \{p_1\}} D^1 \times (0, 1) = \mathcal{E}(v, u).$$

This case is illustrated for a specific example in Figure 15c.

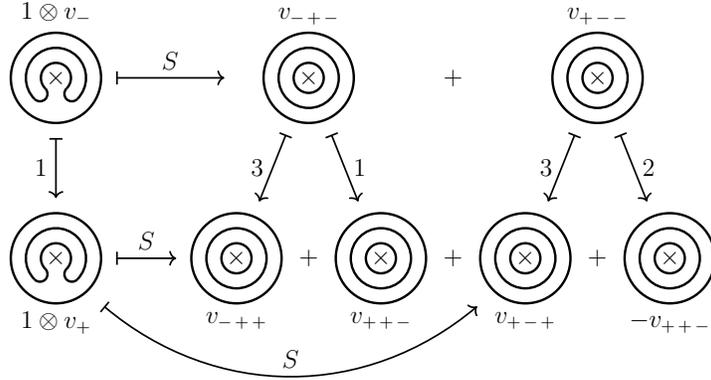
6.3. Square faces involving H . Finally, we turn to the case where $J = H$. As before, we distinguish between two cases.

Suppose first that the saddle S does not change the number of essential circles in D_u . Let s denote the number of essential circles in D_u (equivalently, in any of D_u, D_{w_1}, D_{w_2} , or D_v). If $s = 0$, then $\partial_{w_1} \mathcal{M}_H(y, x) = \partial_{w_2} \mathcal{M}_H(y, x) = \emptyset$, and we define $\mathcal{M}_H(y, x) = \emptyset$. Otherwise, there are unique generators in $z \in \Gamma_{Kh}(w_1)$ and $z' \in \Gamma_{Kh}(w_2)$ such that

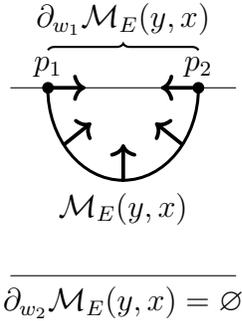
$$\begin{aligned} \partial_{w_1} \mathcal{M}_H(y, x) &= \mathcal{M}_H(y, z) \times \mathcal{M}_H(z, x) \\ \partial_{w_2} \mathcal{M}_H(y, x) &= \mathcal{M}_H(y, z') \times \mathcal{M}_H(z', x). \end{aligned}$$

In fact, $z = y$ and $z' = x$ as generators of $D_u = D_{w_2}$ and $D_v = D_{w_1}$, respectively. The edge moduli spaces $\partial_{w_1} \mathcal{M}_H(y, x)$ and $\partial_{w_2} \mathcal{M}_H(y, x)$ each have s elements, thinly embedded in $(0, 1)$ in order according to Equation (27). From the formulas (6) and (8), we see that the essential circles in $D_u = D_{w_2}$ and $D_v = D_{w_1}$ have the same labels in $x = z'$ and $y = z$, respectively, so the i -th point in $\partial_{w_1} \mathcal{M}_H(y, x)$ has the same framing as the i -th point in $\partial_{w_2} \mathcal{M}_H(y, x)$. Thus we have a bijection of s -point composition moduli spaces that commutes with the thin embeddings precisely as in Equation (31). We therefore define $\mathcal{M}_H(y, x)$ to be a disjoint union of s intervals, embedded in $\mathcal{E}(v, u)$ as a framed identity cobordism from $\partial_{w_1} \mathcal{M}_H(y, x)$ to $\partial_{w_2} \mathcal{M}_H(y, x)$ connecting the i -th point of $\partial_{w_1} \mathcal{M}_H(y, x)$ to the i -th point of $\partial_{w_2} \mathcal{M}_H(y, x)$.

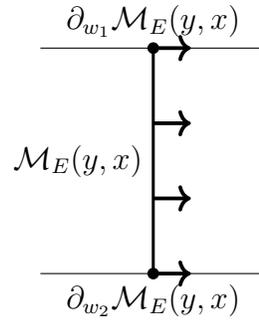
Suppose now that S changes the number of circles. Suppose that $D_v = D_{w_1}$ has s circles and $D_u = D_{w_2}$ has $s + 2$, so that S is the Type II saddle from Figure 3b. If both



(A) The generators and differentials within the subcube $u \leq_2 v$. The numbers on maps indicate which essential circle is being acted on by E . Certain tensor products are written with subscripts to avoid clutter.



(B) The moduli space for $x = 1 \otimes v_-$ and $y = v_+ \otimes v_+ \otimes v_-$, embedded and framed in $\mathcal{E}(v, u)$.



(C) The moduli space for $x = 1 \otimes v_-$ and either $y = v_- \otimes v_+ \otimes v_+$ or $y = v_+ \otimes v_- \otimes v_+$, embedded and framed in $\mathcal{E}(v, u)$.

FIGURE 15. Some examples of $\mathcal{M}_E(y, x) \hookrightarrow \mathcal{E}(v, u) \cong D^1 \times (0, 1)$ for a subcube $u \leq_2 v$ involving E for a disconnected arc diagram. For the fixed $x = 1 \otimes v_-$, different values of y correspond to different cases of Lemma 6.1.

$\partial_{w_1} \mathcal{M}_H(y, x) = \partial_{w_2} \mathcal{M}_H(y, x) = \emptyset$, then define $\mathcal{M}_H(y, x) = \emptyset$. Otherwise, there are generators $z \in \Gamma_{Kh}(w_1)$, $z' \in \Gamma_{Kh}(w_2)$ such that

$$\begin{aligned} \partial_{w_1} \mathcal{M}_H(y, x) &= \mathcal{M}_H(y, z) \times \mathcal{M}_H(z, x) \\ \partial_{w_2} \mathcal{M}_H(y, x) &= \mathcal{M}_H(y, z') \times \mathcal{M}_H(z', x). \end{aligned}$$

The edge moduli spaces $\mathcal{M}_H(z, x)$ and $\mathcal{M}_H(y, z')$ each have one element, while $\mathcal{M}_H(y, z)$ and $\mathcal{M}_H(z', x)$ have s and $s + 2$ elements, respectively, corresponding to the number of essential circles in D_{w_1} and D_{w_2} . Note that the two essential circles merged by S are consecutive in D_{w_2} , and their labels in z' are different. Define $\mathcal{M}_H(y, x)$ to consist of $s + 1$ intervals, with s of them thinly embedded in $\mathcal{E}(v, u)$ matching the s points in $\mathcal{M}_H(y, z)$ with the s points in $\mathcal{M}_H(z', x)$ that correspond to essential circles not participating in the saddle. The remaining interval in $\mathcal{M}_H(y, x)$ is thinly embedded as a turnback with boundary the two remaining unmatched points in $\partial_{w_2} \mathcal{M}_H(y, x)$ corresponding to the essential circles created by S . See Figure 16 for an example. Finally, the case when S increases the number of essential circles by two is dual to the above discussion.

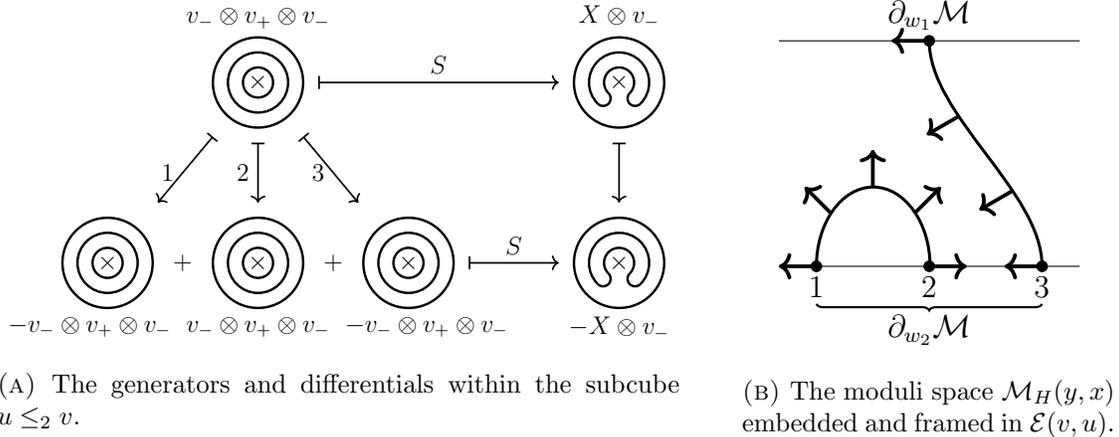


FIGURE 16. An example of $\mathcal{M}_H(y, x) \hookrightarrow \mathcal{E}(v, u) \cong D^1 \times (0, 1)$ for a subcube $u \leq_2 v$ involving an H edge and a saddle S which decreases the number of essential circles by two. The two merged circles must be consecutive, allowing a framed turnback to ‘cancel’ the corresponding points in $\partial_{w_2} \mathcal{M}_H(y, x)$.

7. PATH MODULI SPACES

This section contains a detailed combinatorial analysis that will be crucial for completing our inductive construction in Section 8. For an m -dimensional subcube $u \leq_m v$ and generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$, Item (3) of Theorem 5.8 states that the codimension $m - 1$ boundary of $\mathcal{M}_J(y, x)$ is built out of products of the moduli spaces associated to a sequence of edges from u to v . Such a sequence of edges is a *path* from u to v , and each edge moduli space, along with its thick embedding, was defined in Section 5.2. We analyze these *path moduli spaces* and their thick embeddings, obtained as the product of the thick embeddings for each edge moduli space. While each edge moduli space is thinly embedded, taking the disjoint union over intermediate generators as in Theorem 5.8 (3) may produce a thick embedding.

We focus on connected arc diagrams in this section. The key results are Proposition 7.10 and Proposition 7.17. In Proposition 7.10 we show that for $J = E, F$ all path moduli spaces involving J are in fact thinly embedded. For $J = H$, path moduli spaces are thinly embedded except for one type of arc diagram, for which an explicit description is provided, see Proposition 7.17. The outline is as follows. In Section 7.1 we introduce path moduli spaces, their thick embeddings, and present a helpful lemma which bounds the number of points in some path moduli spaces. Section 7.2 establishes two key lemmas. First, Lemma 7.7 allows us to assume that a path has the special property that the arc diagram just before the J edge (the so-called *J resolution*) consists entirely of essential circles. Second, in Lemma 7.8 we give restrictions on the type of arc configurations that may appear in such a diagram.

To establish thinness we focus on the J -resolution arc diagram. While the two cases $J = E, F$ and $J = H$ require different arguments, the overall strategy is the following. Connectivity and the formulas (7) and (9) imply that picking an essential circle for J to act on and a label v_{\pm} on that circle determines labels on all the other circles. Lemma 7.3 ensures thinness ‘before and after’ the J edge. With further analysis, depending on $J = E, F$ or $J = H$, we arrive at our desired results. Note that connectivity of the arc diagram is critical:

disconnected arc diagrams may yield non-thinly embedded path moduli spaces. See also the discussion at the beginning of Section 8.2.

7.1. Basic definitions and lemmas.

Definition 7.1. A *path* P through a subcube $u \leq_m v$ is a sequence of edges, denoted as

$$P = (u \rightarrow w_1 \rightarrow \cdots \rightarrow w_{m-1} \rightarrow v).$$

Fixing generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$, any such path determines a *path moduli space*

$$\partial_P \mathcal{M}_J(y, x) := \coprod \mathcal{M}_J(y, z_{m-1}) \times \mathcal{M}_J(z_{m-1}, z_{m-2}) \times \cdots \times \mathcal{M}_J(z_1, x),$$

where the disjoint union is taken over all sequences (z_1, \dots, z_{m-1}) with $z_i \in \Gamma_{Kh}(w_i)$ for each $1 \leq i \leq m-1$. As in Section 6, points in $\partial_P \mathcal{M}_J(y, x)$ will be denoted as

$$(x \rightarrow z_1 \rightarrow \cdots \rightarrow z_{m-1} \rightarrow y) \in \partial_P \mathcal{M}_J(y, x),$$

with a label on an arrow corresponding to the J -edge indicating which essential circle is being acted on by J if necessary (see Equation (35) for an example of such a label).

Recall from Definition 3.1 that we use the notation $\partial_{[i]} \mathcal{M}$ to indicate the codimension- i portion of the boundary of a moduli space. In order to satisfy the demands of Theorem 5.8, we define

$$(33) \quad \partial_{[m-1]} \mathcal{M}_J(y, x) \cong \coprod_{\text{paths } P} \partial_P \mathcal{M}_J(y, x)$$

where the disjoint union is taken over paths P from u to v ; to make the notation less cumbersome we omit reference to u, v .

Furthermore, every $\partial_P \mathcal{M}_J(y, x)$ inherits a thick embedding

$$(34) \quad \partial_P \mathcal{M}_J(y, x) \xrightarrow{\varphi_P} \mathcal{E}(P) := \mathcal{E}(v, w_{m-1}) \times \cdots \times \mathcal{E}(w_1, u)$$

via the disjoint union of products of the thin embeddings

$$\mathcal{M}_J(z_{m-i}, z_{m-i-1}) \hookrightarrow \mathcal{E}(w_{m-i}, w_{m-i-1})$$

defined in Section 5.2. If the subcube $u \leq_m v$ does not involve J , then $\mathcal{E}(P)$ is a point and the thick embedding is constant. Otherwise, J is involved and there is a natural identification $\mathcal{E}(P) \cong (0, 1)$. Then the thick embedding (34) sends a point

$$(35) \quad (x \rightarrow z_1 \rightarrow \cdots \rightarrow z_i \xrightarrow{j} z_{i+1} \rightarrow \cdots \rightarrow z_{m-1} \rightarrow y) \in \partial_P \mathcal{M}_J(y, x),$$

to the fixed j -th point in $(0, 1)$ with framing determined by the sign $(-1)^{j+1}$ if $J = E, F$ and otherwise given by the label of the j -th circle in z_i if $J = H$, as described in Section 5.2.

We may also consider *sub-paths* of a path P . If w_k is any vertex along a path

$$P = (u \rightarrow w_1 \rightarrow \cdots \rightarrow w_k \rightarrow \cdots \rightarrow v),$$

we may define $P_{\leq w_k} := (u \rightarrow \cdots \rightarrow w_k)$ and $P_{\geq w_k} := (w_k \rightarrow \cdots \rightarrow v)$. This then allows us to split the path moduli space as a product

$$\partial_P \mathcal{M}_J(y, x) = \coprod_{z \in \Gamma_{Kh}(w_k)} \partial_{P_{\geq w_k}} \mathcal{M}_J(y, z) \times \partial_{P_{\leq w_k}} \mathcal{M}_J(z, x)$$

embedded as a product with the product framing.

Definition 7.2. For a fixed path P through a sub-cube $u \leq_m v$ that involves J , we single out the vertices along P directly before and after the J map has been applied:

- $e_0 = e_0(P)$ denotes the source vertex of the J edge in P , and
- $e_1 = e_1(P)$ denotes the target vertex of the J edge in P .

The arc diagram $D_e := D_{e_0} = D_{e_1}$ being acted on by J will be referred to as the J -resolution of P .²

Note that the J -resolution of a path P depends on the path, although we do not include P in the notation to avoid clutter.

Now if $u \leq_m v$ does not involve the J map, then any path P from u to v represents a path in the cube of resolutions of D . For each such path there is an associated cobordism $S: D_u \rightarrow D_v$.

Lemma 7.3. *Let P be a path from u to v which does not involve the J map, and let $S: D_u \rightarrow D_v$ denote its corresponding surface cobordism. Let g denote the maximal genus among all components of S .*

- (1) *If $g = 0$, then for every $x \in \Gamma_{Kh}(u)$ and $y \in \Gamma_{Kh}(v)$, the moduli space $\partial_P \mathcal{M}_J(y, x)$ contains at most one element.*
- (2) *If $g \geq 1$ and one of D_u or D_v consists only of essential circles, then $\mathcal{F}_{\mathbb{A}}(S) = 0$. In particular in this case, for every $x \in \Gamma_{Kh}(u)$ and $y \in \Gamma_{Kh}(v)$, the moduli space $\partial_P \mathcal{M}_J(y, x)$ is empty.*

Proof. For (1), first view S as a cobordism in $\mathbb{R}^2 \times I$ using the inclusion $\mathbb{A} \times I \hookrightarrow \mathbb{R}^2 \times I$. Since $\mathcal{F}_{\mathbb{A}}$ is the adeg-preserving part of \mathcal{F}_{Kh} , the cardinality of $\partial_P \mathcal{M}_J(y, x)$ is bounded above by the coefficient of y in $\mathcal{F}_{Kh}(S)(x)$. Represent the generators x and y using (dotted) cup cobordisms, denoted Σ_x and Σ_y respectively. Let Σ_y^* denote the dual cap cobordism to Σ_y , so that the closed surface

$$\Sigma_y^* \circ \Sigma_y$$

evaluates to 1 under \mathcal{F}_{Kh} . The coefficient of y in $\mathcal{F}_{Kh}(S)(x)$ is then the value of \mathcal{F}_{Kh} on the closed surface

$$S' := \Sigma_y^* \circ S \circ \Sigma_x.$$

Since S has genus zero, every component of S' is a (possibly dotted) sphere. Then $\mathcal{F}_{Kh}(S') \leq 1$, which completes the proof of statement (1).

We now address (2). Let $\bar{S} \subset S$ be a component of positive genus. It suffices to show $\mathcal{F}_{\mathbb{A}}(\bar{S}) = 0$. By neck-cutting, Figure 1b, we can write \bar{S} as a sum

$$S_1 + \cdots + S_k,$$

where each component of each S_i is either an annulus with essential boundary or a disk, possibly carrying dots. We may assume that each component carries at most one dot. By assumption, each S_i contains an annulus. We will show that each S_i has a dotted annulus, which implies $\mathcal{F}_{\mathbb{A}}(\bar{S}) = 0$ since X acts by zero on essential circles, see formula (6).

Fix $1 \leq i \leq k$. Recall that the Bar-Natan relations are homogeneous, so $\deg(\bar{S}) = \deg(S_i)$. Let n denote the number of (possibly dotted) disks in S_i . If none of the annuli in S_i are dotted, then

$$\deg(S_i) \geq -n.$$

²To relate this notation to equation (35), $e_0 = w_i$, $e_1 = w_{i+1}$.

On other other hand, considering the connected surface \bar{S} ,

$$\deg(\bar{S}) = \chi(\bar{S}) < -n,$$

since $g(\bar{S}) \geq 1$ and the number of boundary components of \bar{S} is strictly greater than n . Therefore S_i contains a dotted annulus, which completes the proof. \square

7.2. Equivalent path moduli spaces and a non-emptiness condition. In this section we analyze the manner in which two paths P, P' can give rise to equivalent path moduli spaces $\partial_P \mathcal{M}_J(y, x) \cong \partial_{P'} \mathcal{M}_J(y, x)$. This analysis will allow us to choose certain preferred paths through our subcubes. It will also give rise to an important non-emptiness condition on path moduli spaces.

Definition 7.4. Let P and P' be two paths in a subcube $u \leq_m v$. Write

$$\begin{aligned} P &= (u = w_0 \rightarrow w_1 \rightarrow \cdots \rightarrow w_m = v) \\ P' &= (u = w'_0 \rightarrow w'_1 \rightarrow \cdots \rightarrow w'_m = v). \end{aligned}$$

We say that P and P' are *related by a square face* if for some $1 \leq j \leq m - 1$, we have $w_j \neq w'_j$ while $w_i = w'_i$ for all indices $i \neq j$. We say the square face involves J if one of the edges in the 2-dimensional subcube $w_{j-1} \leq_2 w_{j+1}$ is the J map. In this case, two of the opposite edges are J , and the remaining two opposite edges are a saddle.

Lemma 7.5. *Let P and P' be two paths in a subcube $u \leq_m v$ that are related by a square face $w_{j-1} \leq_2 w_{j+1}$ as above, and let $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$. If either*

- (1) *the square face does not involve the J map, or*
- (2) *the square face involves the J map, and the number of essential circles in the resolutions appearing in $w_{j-1} \leq_2 w_{j+1}$ is the same,*

then there is a bijection

$$\partial_P \mathcal{M}_J(y, x) \cong \partial_{P'} \mathcal{M}_J(y, x)$$

such that the diagram (36) commutes.

$$(36) \quad \begin{array}{ccc} \partial_P \mathcal{M}_J(y, x) & \xhookrightarrow{\quad} & \mathcal{E}(P) \\ \cong \downarrow & & \cong \downarrow \\ \partial_{P'} \mathcal{M}_J(y, x) & \xhookrightarrow{\quad} & \mathcal{E}(P') \end{array}$$

Proof. The desired bijection is assembled from the bijections in Section 6 on the factors of $\partial_P \mathcal{M}_J(y, x)$ and $\partial_{P'} \mathcal{M}_J(y, x)$ involving the square face and the identity on all other factors. Commutativity of the diagram is immediate for item (1), since $\mathcal{E}(P) = D^0 = \mathcal{E}(P')$ in this case. For item (2) it follows from commutativity of the diagram (31) and the discussion below Lemma 6.1. \square

In the case that the number of essential circles in the J -resolution has changed, we instead see turnbacks as in Figure 14b, but this fact does not concern us at the moment.

Corollary 7.6. *Let P and P' be paths from u to v such that the J -resolution of P and P' are at the same vertex, $e_0(P) = e_0(P')$. Then there is a bijection*

$$\partial_P \mathcal{M}_J(y, x) \cong \partial_{P'} \mathcal{M}_J(y, x)$$

such that the diagram (36) commutes.

Proof. The paths P and P' can be related by a sequence of square faces that do not involve the J map. The claim follows from Lemma 7.5. \square

Now we consider how the various bijections of Lemma 7.5 can be used to replace a given path through a subcube with one that is easier to analyze.

Lemma 7.7. *Consider a path P through the sub-cube $u \leq_m v$, specifying a J -resolution arc diagram D_e , and fix generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$. Then there exists a path P' through $u \leq_m v$, with J -resolution $D_{e'}$, satisfying the following properties.*

- *Each connected component of $D_{e'}$ contains the same number of essential circles as the corresponding component of D_e .*
- *Any connected component of $D_{e'}$ which contains at least one essential circle also contains no trivial circles.*
- *There is a bijection $\partial_P \mathcal{M}_J(y, x) \cong \partial_{P'} \mathcal{M}_J(y, x)$ such that the diagram (36) commutes.*

Proof. We work one connected component \mathcal{C} of D_e at a time. If \mathcal{C} contains only essential circles or only trivial circles, then we are done. Otherwise, \mathcal{C} must contain a trivial circle that is joined by an arc (either past or future) to an essential circle. This arc specifies a square face in the subcube $u \leq_m v$, and the swap across this square face yields a path P'' in which the corresponding component \mathcal{C}'' in the J -resolution of P'' contains one fewer trivial circle and the same number of essential circles as \mathcal{C} . By Lemma 7.5 (2), the moduli spaces for P and P'' are in bijection with diagram (36) commuting. We continue this process until there are no trivial circles left in any of the connected components with essential circles. \square

Lemma 7.7 ensures that it is enough to understand $\partial_P \mathcal{M}_J(y, x)$ for paths P having the J -resolution D_e consisting of only essential circles, possibly together with a disconnected component consisting of trivial circles. However, in the following sections, we will focus largely on connected arc diagrams; the disconnected cases will follow from the connected ones in a manner very similar to the arguments of Section 6.2.

We end this section with an important consequence of Corollary 7.6. Recall from Section 2.3 that a subcube $u \leq_m v$ gives rise to its own arc diagrams at each vertex including only those arcs that correspond to saddle maps within the subcube. Note further that, for any path P , the arc diagrams D_{e_0} and D_{e_1} directly before and after the J map are the same. The following lemma provides an important restriction on the types of arc diagrams that can appear for non-empty path moduli spaces.

Lemma 7.8. *Let P be a path through $u \leq_m v$ and let w be a vertex in P . Suppose that any of the configurations in Figure 17, as well as their vertical and horizontal reflections, appear for some pair of arcs in the arc diagram D_w . Then $\partial_P \mathcal{M}_J(y, x) = \emptyset$ for every $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$.*

Proof. Surgery arcs can be performed in any order by Corollary 7.6. We assume that the depicted pair of arcs correspond to consecutive edges in P . Direct calculations verify that they all yield the zero map. \square

7.3. The E and F maps for connected arc diagrams. In this section we analyze the path moduli spaces $\partial_P \mathcal{M}_J(y, x)$ for connected arc diagrams when $J = E, F$ and show that the thick embeddings (34) in such cases are in fact thin. (We will see in Section 8.2 how genuinely thick embeddings appear when the arc diagrams are disconnected.) The analysis is stated

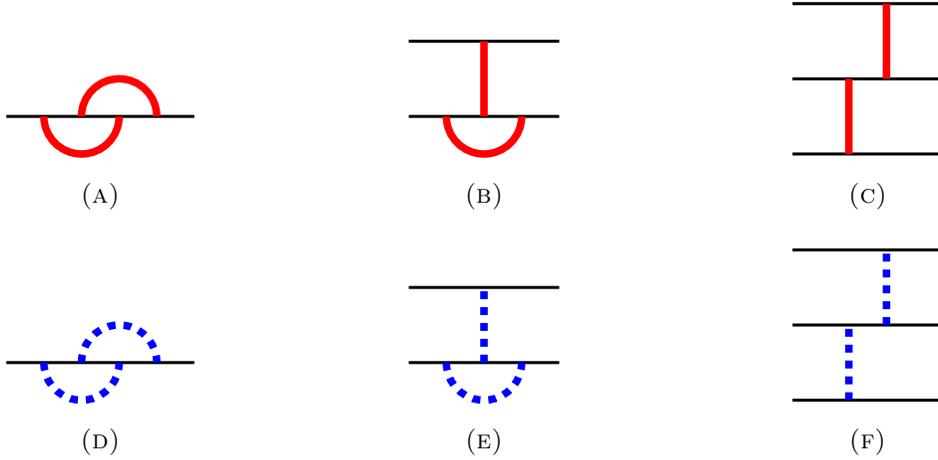


FIGURE 17. Forbidden configurations; only the two relevant arcs are pictured, but there could be an arbitrary number of additional past and future arcs anywhere within such a diagram.

only for the E map, but the F map is nearly identical, see Lemma 2.6; in particular all results here hold with F and E interchanged, though the proofs may need slight modifications.

Fix an annular link diagram D with n -crossings, so that together with the E map we have an $(n + 1)$ -dimensional cube \mathcal{Q} . We also fix a sub-cube $u \leq_m v$. For the remainder of this section, we assume that the sub-cube $u \leq_m v$ contains the E edge and that the arc diagram D_u is connected. Our first lemma rules out a very specific arrangement of arcs for such a diagram.

Lemma 7.9. *Suppose the arc diagram D_u is connected. Let P be a path through $u \leq_m v$ such that the E -resolution D_e contains only essential circles. If the configuration shown in (37) appears for some pair of arcs in the arc diagram D_{e_0} , where the two horizontal intervals are segments of distinct essential circles, then $\partial_P \mathcal{M}_E(y, x) = \emptyset$ for every $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$.*



Proof. We consider two cases. If there are no other essential circles in D_e , then by Corollary 7.6 we may assume that the past arc occurs directly before the E map, and the future arc occurs directly after. It is straightforward to verify that the configuration results in empty moduli spaces. If there are other essential circles, then since the arc diagram D_e is connected, at least one of the two depicted circles must be connected to a third circle. Then one of Figure 17c or Figure 17f appears in the arc diagram, so we are done by earlier considerations. \square

We now present the main technical result of this section which concerns the thinness of path moduli spaces for connected arc diagrams.

Proposition 7.10. *Suppose P is a path through a subcube $u \leq_m v$ having connected arc diagram D_u . Then for any fixed generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$, the thick embedding φ_P from (34) is in fact a thin embedding*

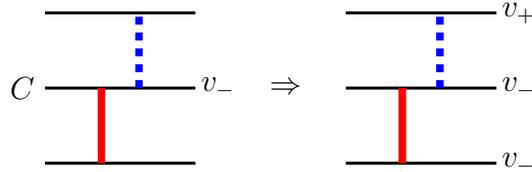
$$\partial_P \mathcal{M}_E(y, x) \xrightarrow{\varphi_P} (0, 1) \cong \mathcal{E}(v, u).$$

Proof. If the path moduli space $\partial_P \mathcal{M}_E(y, x)$ is empty, then the statement is trivially true. Otherwise, by Lemma 7.7, we can assume that our E -resolution consists entirely of essential circles. We fix two points $z, z' \in \partial_P \mathcal{M}_E(y, x)$ and seek out to prove that, if $\varphi_P(z) = \varphi_P(z')$, then in fact we must have had $z = z'$.

We write z, z' as

$$\begin{aligned} z &= (x \rightarrow z_1 \rightarrow \cdots \rightarrow z_i \xrightarrow{j} z_{i+1} \rightarrow \cdots \rightarrow z_{m-1} \rightarrow y) \\ z' &= (x \rightarrow z'_1 \rightarrow \cdots \rightarrow z'_i \xrightarrow{j'} z'_{i+1} \rightarrow \cdots \rightarrow z'_{m-1} \rightarrow y), \end{aligned}$$

where $z_i, z'_i \in \Gamma_{Kh}(e_0)$ are the generators that E acts on. Then by definition of φ_P , we must have $j = j'$. Letting C denote the j -th essential circle in D_e , this implies that the label on C must be v_- in z_i and v_+ in z_{i+1} . Based on the assumption that the arc diagram is connected and using Lemmas 7.8, 7.9, one observes that the curves are connected by past/future arcs in an alternating manner as illustrated in the following figure (also see Figure (38)). The labels discussed above then determine the labels on the circles adjacent to C as shown below.



By propagating this argument we see that the labels on all essential circles in D_e are determined by j , and thus for $j = j'$ we must have $z_i = z'_i$ and $z_{i+1} = z'_{i+1}$ as well.

From this we can conclude that both z and z' belong to the same composition moduli space

$$z, z' \in \partial_{P_{\geq e}} \mathcal{M}_E(y, z_{i+1}) \times \mathcal{M}_E(z_{i+1}, z_i) \times \partial_{P_{\leq e}} \mathcal{M}_E(z_i, x).$$

Then since D_e contains only essential circles, Lemma 7.3 ensures that both $\partial_{P_{\leq e}} \mathcal{M}_E(z_i, x)$ and $\partial_{P_{\geq e}} \mathcal{M}_E(y, z_{i+1})$ are at most single point moduli spaces, while $\mathcal{M}_E(z_{i+1}, z_i)$ is an edge moduli space and so contains at most a single point as indicated in Section 5.2. Thus we have $z = z'$ as desired, and the embedding φ_P must in fact be thin. \square

Although Proposition 7.10 suffices to continue our construction, we can actually prove a stronger statement about our path moduli spaces which makes them easier to understand in concrete examples. We begin with a special case.

Lemma 7.11. *Suppose P is a path through a subcube $u \leq_m v$ with (connected) E -resolution arc diagram D_e consisting entirely of s essential circles, with precisely one arc connecting each consecutive pair. Fix $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$ such that the moduli space $\partial_P \mathcal{M}_E(y, x)$ is nonempty. Then for each $1 \leq j \leq s$, there is at least one point of the form*

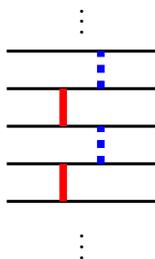
$$(z^j) := (x \rightarrow \cdots \rightarrow z_i \xrightarrow{j} z_{i+1} \rightarrow \cdots \rightarrow y) \in \partial_P \mathcal{M}_E(y, x).$$

Proof. Let $S_{\leq e}: D_u \rightarrow D_e, S_{\geq e}: D_e \rightarrow D_v$ denote the cobordisms obtained by composing the saddle maps in P before the E map and after the E map, respectively.

Fixing $1 \leq j \leq s$, we first note that the j -th essential circle in D_e must be labeled by v_- in z_i . Using the argument in Proposition 7.10, there is a unique choice of generators z_i, z_{i+1} satisfying this condition. To construct (z^j) , we will show that $\mathcal{F}_A(S_{\geq e})(z_{i+1})$ is a generator

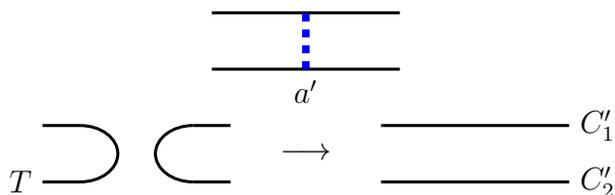
of $\mathcal{F}_{\mathbb{A}}(D_v)$, which we denote y' . Then we show there exists a unique $x' \in \Gamma_{Kh}(u)$ such that z_i appears in $\mathcal{F}_{\mathbb{A}}(S_{\leq e})(x')$. Finally we will show that $x' = x$ and $y' = y$, completing the proof.

By Lemma 7.8, Figures 17c and 17f, the arc diagram for D_e consists of s essential circles connected by past and future arcs in an alternating manner,

(38) 

Let us first show that z_{i+1} is not killed by applying the future arcs in D_e . Note that each circle in D_e is incident to at most one future arc. Pick circles C_1, C_2 in D_e that are connected by a future arc a . Surgery along a results in a trivial circle C' , and by construction of z_{i+1} the labels on C_1 and C_2 are different. Therefore the map assigned to a is nonzero on z_{i+1} , and moreover C' is labeled by X in the image of z_{i+1} . Repeating this for all future arcs, we see that $y' := \mathcal{F}_{\mathbb{A}}(S_{\geq e})(z_{i+1})$ is a generator of D_v . Each trivial circle in D_v is labeled X in y' . Essential circles in D_v are those with no future arc incident on them in D_e , and they have the same label as in z_{i+1} .

Let us now define $x' \in \Gamma_{Kh}(u)$ such that z_i appears in $\mathcal{F}_{\mathbb{A}}(S_{\leq e})(x')$. As before, each circle in D_e is incident to at most one past arc. If a' is a past surgery arc joining the essential circles C'_1 and C'_2 in D_e , then undoing a' we see a trivial circle T splitting into the two essential circles $C'_1 \cup C'_2$,

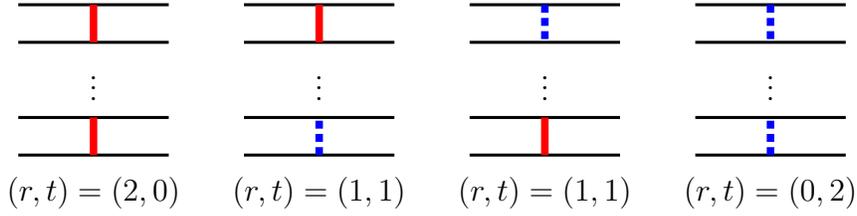


The diagram shows a trivial circle T on the left, which is a single loop. An arrow points to the right, where a past surgery arc a' (represented by a vertical blue dashed line) connects two horizontal lines representing essential circles C'_1 and C'_2 . The arc a' is labeled with a' below it.

We must label T by the generator 1. By construction, C'_1 and C'_2 have opposite labels in z_i . The saddle map corresponding to a' sends $1 \in \mathcal{F}_{\mathbb{A}}(T)$ to a sum of two generators in $\mathcal{F}_{\mathbb{A}}(C'_1 \cup C'_2)$, one of which matches the labels on $C'_1 \cup C'_2$ in z_i . Repeat this for all past arcs in D_e to obtain the desired generator $x' \in \Gamma_{Kh}(u)$. Note that trivial circles in D_u are labeled by 1 in x' . Essential circles in D_u correspond to essential circles in D_e with no past arc incident on them, and they have the same label in x' and z_i .

It remains to show that $x' = x$ and $y' = y$. By the above analysis, x' and x are labeled 1 on all trivial circles in D_u , and likewise y' and y are labeled X on all trivial circles in D_v . We will argue that labels on essential circles in D_u and D_v are determined by the arc diagram D_e .

Let r and t denote the number of essential circles in D_u and D_v , respectively. Since the arc diagram D_e is connected, Lemma 7.8, Figures 17c and 17f imply that all but the first and last essential circles in D_e are connected to adjacent essential circles by both a past and future arc. It follows that the pair (r, t) is one of $(2, 0)$, $(1, 1)$, or $(0, 2)$, which can be read off from the type of arcs incident on the first and last circle in D_e ,



Finally, note that if $x'' \in \Gamma_{Kh}(u)$, $y'' \in \Gamma_{Kh}(v)$ such that $\partial_P \mathcal{M}_E(y'', x'') \neq \emptyset$, then $\text{adeg}(y'') = \text{adeg}(x'') + 2$. For each of the three possibilities for (r, t) , we see that the labels on essential circles in D_u and D_v are always uniquely determined whenever the moduli space is nonempty. Therefore $x = x'$ and $y = y'$. \square

With Lemma 7.11 in place, we can state and prove the following proposition which characterizes all non-empty path moduli spaces in any subcube for a connected diagram which involves the E map.

Proposition 7.12. *Suppose P is a path through a subcube $u \leq_m v$ having connected arc diagram D_u , such that the E -resolution D_e of P contains s essential circles. Then for any fixed generators $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$, if the path moduli space $\partial_P \mathcal{M}_E(y, x)$ is non-empty, then the thick embedding φ_P from (34) is a thin embedding mapping $\partial_P \mathcal{M}_E(y, x)$ bijectively onto $\{\frac{1}{s+1}, \dots, \frac{s}{s+1}\} \subset (0, 1)$.*

Proof. Thinness of the embedding $\partial_P \mathcal{M}_E(y, x) \xrightarrow{\varphi_P} (0, 1)$ was established in Proposition 7.10. By Lemma 7.7, we can assume that our E -resolution consists entirely of s essential circles. It thus remains to show that each essential circle in D_e contributes at least one point to $\partial_P \mathcal{M}_E(y, x)$.

Fix $1 \leq j \leq s$. By Corollary 7.6, the order in which the surgery arcs in the sub-cubes $u \leq e_0$ and $e_1 \leq v$ are performed is irrelevant, and we use this implicitly.

Pick future surgery arcs a_1, \dots, a_f and past surgery arcs b_1, \dots, b_p in D_e , such that consecutive circles in D_e are connected by exactly one of the chosen arcs. We assume that in the path P , the f edges directly after the E edge correspond to performing surgery on a_1, \dots, a_f , and the p edges directly before the E edge correspond to surgery on the dual arcs to b_1, \dots, b_p . In other words, P is of the form

$$u \rightarrow \cdots \rightarrow u_0 \xrightarrow{b_1} u_1 \xrightarrow{b_2} \cdots \xrightarrow{b_p} e_0 \xrightarrow{E} e_1 \xrightarrow{a_1} v_1 \xrightarrow{a_2} \cdots \xrightarrow{a_f} v_f \rightarrow \cdots \rightarrow v.$$

Consider the sub-cube $u_0 \leq v_f$ and the path P^2 through this sub-cube, consisting of all edges in P from u_0 to v_f . Let D'_{u_0} denote the diagram obtained from D_{u_0} by including only arcs in P^2 . The E -resolution of P^2 is connected by construction, so D'_{u_0} is connected as well. Let P^1 and P^3 denote the sub-paths of P consisting of edges from u to u_0 and v_f to v , respectively.

$$\underbrace{u \rightarrow \cdots \rightarrow u_0}_{P^1} \xrightarrow{b_1} u_1 \xrightarrow{b_2} \cdots \xrightarrow{b_p} e_0 \xrightarrow{E} e_1 \xrightarrow{a_1} v_1 \xrightarrow{a_2} \cdots \xrightarrow{a_f} v_f \xrightarrow{\cdots} v \xrightarrow{\cdots} v$$

$\underbrace{\hspace{15em}}_{P^2}$

Since $\partial_P \mathcal{M}_E(y, x)$ is nonempty, there exist generators $x' \in \Gamma_{Kh}(u_0)$, $y' \in \Gamma_{Kh}(v_f)$ such that the moduli spaces

$$\partial_{P^1} \mathcal{M}_E(x', x), \partial_{P^2} \mathcal{M}_E(y', x'), \partial_{P^3} \mathcal{M}_E(y, y')$$

are all nonempty. By Lemma 7.11 applied to D'_{u_0} and P^2 , we can find $z^2 \in \partial_{P^2}\mathcal{M}_E(y', x')$ of the form

$$z^2 = (x' \rightarrow \cdots \rightarrow z_i \xrightarrow{j} z_{i+1} \rightarrow \cdots \rightarrow y') \in \partial_{P^2}\mathcal{M}_E(y', x').$$

Then choosing any $z^1 \in \partial_{P^1}\mathcal{M}_E(x', x)$ and $z^3 \in \partial_{P^3}\mathcal{M}_E(y, y')$ gives rise to at least one point in the composition moduli space

$$z = (z^3, z^2, z^1) \in \partial_{P^3}\mathcal{M}_E(y, y') \times \partial_{P^2}\mathcal{M}_E(y', x') \times \partial_{P^1}\mathcal{M}_E(x', x) \subset \partial_P\mathcal{M}_E(y, x)$$

which φ_P maps to the j -th point in $(0, 1)$ as desired. \square

7.4. The H map for connected arc diagrams. As in the previous section, we continue to fix a subcube $u \leq_m v$ involving $J = H$ such that the initial arc diagram D_u is connected (and hence all intermediate diagrams D_w for $u \leq w \leq v$ are connected as well). The main technical result is a characterization of the path moduli spaces in $u \leq_m v$, Proposition 7.17. We begin with several lemmas.

Lemma 7.13. *Let P be a path through $u \leq_m v$ such that the H -resolution D_e consists entirely of s essential circles. Fix generators $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$ such that $\partial_P\mathcal{M}_H(y, x)$ is nonempty. If every pair of consecutive circles in D_e is joined by exactly one type of arc (either past or future), then the thick embedding φ_P from (34) is in fact thin.*

Proof. Let $z, z' \in \partial_P\mathcal{M}_H(y, x)$ with

$$\begin{aligned} z &= (x \rightarrow \cdots \rightarrow z_i \xrightarrow{j} z_{i+1} \rightarrow \cdots \rightarrow y), \\ z' &= (x \rightarrow \cdots \rightarrow z'_i \xrightarrow{j} z'_{i+1} \rightarrow \cdots \rightarrow y), \end{aligned}$$

for some $1 \leq j \leq s$. Let C denote the j -th circle in D_e . We will first show that $z_i = z'_i$ (and thus $z_{i+1} = z'_{i+1}$ as well).

Choosing a label of v_{\pm} on C uniquely determines the generator z_i (and thus z_{i+1}), as in the proof of Proposition 7.10. Therefore, in order to have $z_i \neq z'_i$, we must have the labels on C in z_i and z'_i be opposites. It follows that z_i is obtained from z'_i by swapping the labels on all circles in D_e , and likewise for z_{i+1} and z'_{i+1} .

On the other hand, Lemma 7.8 (Figures 17c and 17f) implies that the circles in D_e are connected by past and future arcs in an alternating manner as in (38), where there may be more than one past or future arc between any two adjacent circles. Since the arc diagram D_e is connected, all but the first and last essential circles in D_e are connected to adjacent essential circles by both a past and future arc. Therefore at least one of D_u or D_v has an essential circle. Moreover, the labels on the essential circle(s) in x (resp. y) are the same as the labels on the corresponding essential circle(s) in z_i (resp. z_{i+1}). Since x and y are fixed, at least one of these prevents the circles in D_e from having opposite labels for either z_i and z'_i , or z_{i+1} and z'_{i+1} .

Thus we must have that $z_i = z'_i$ and $z_{i+1} = z'_{i+1}$, and the lemma follows by the same argument as in the proof of Proposition 7.10. \square

Lemma 7.13 does not handle all cases where D_e consists of only essential circles. After ruling out the forbidden configurations Figure 17c and Figure 17f from Lemma 7.8, it remains to analyze the situation where D_e consists of exactly two essential circles with both a past

and future arc connecting them, as in (37). The following lemma gives an explicit description of such a path moduli space and its thick embedding.

Lemma 7.14. *Let P be a path through $u \leq_m v$ such that the H -resolution D_e consists entirely of two essential circles, with at least one past and one future arc connecting them. Fix generators $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$ such that $\partial_P \mathcal{M}_H(y, x)$ is nonempty. Then $\partial_P \mathcal{M}_H(y, x)$ consists of four points, and the thick embedding from (34) is a two-to-one map onto the points $\{\frac{1}{3}, \frac{2}{3}\} \subset (0, 1) \cong \mathcal{E}(v, u)$. Moreover, any two points in $\partial_P \mathcal{M}_H(y, x)$ that map to the same point in $\mathcal{E}(v, u)$ are oppositely framed.*

Proof. Let a_p and a_f denote the saddles corresponding to one past and one future surgery arc in D_e , respectively. By Corollary 7.6, we may assume that a_p and a_f are performed directly before and after the H edge, respectively. Let P' denote the 3-edge subpath of P corresponding to a_p , H , and a_f , and note that the diagrams at the start and end of this path both consist of a single trivial circle. Let $P_{\leq p}$ (resp. $P_{\geq f}$) denote the subpath consisting of the edges before a_p (resp. after a_f).

It is a straightforward calculation to check that, in order for $\partial_{P'} \mathcal{M}_H(z_f, z_p)$ to be nonempty, the generator z_p (resp. z_f) must consist of a label 1 (resp. X) on the trivial circle at the start (resp. end) of the path P' . In this unique case one further computes that $\partial_{P'} \mathcal{M}_H(z_f, z_p)$ consists of four points thickly embedding into $(0, 1)$ as in the statement of the lemma.

Then for the full path moduli space we have

$$\partial_P \mathcal{M}_H(y, x) \cong \partial_{P_{\geq f}} \mathcal{M}_H(y, z_f) \times \partial_{P'} \mathcal{M}_H(z_f, z_p) \times \partial_{P_{\leq p}} \mathcal{M}_H(z_p, x).$$

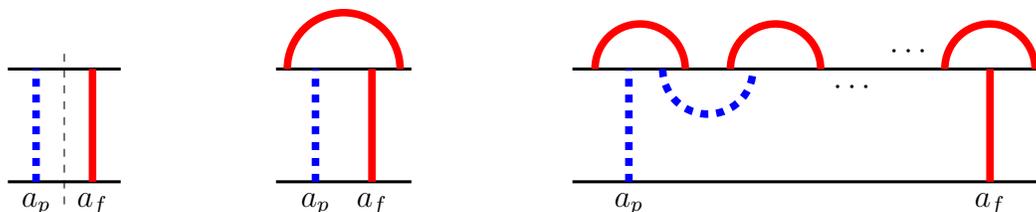
The cobordism corresponding to the path $P_{\leq p}$ (resp. $P_{\geq f}$) must have genus zero, otherwise Lemma 7.3 would cause the path moduli space to be empty after composing with a_p (resp. a_f) and connecting to D_e which consists of all essential circles. Thus both $\partial_{P_{\leq p}} \mathcal{M}_H(z_p, x)$ and $\partial_{P_{\geq f}} \mathcal{M}_H(y, z_f)$ consist of a single trivially embedded point and we are done. \square

Remark 7.15. The two past and future arcs in Lemma 7.14 form a ladybug configuration ([LS14a, Figure 5.1]). Lemma 7.9 disallows this configuration for $J = E, F$, and it is the only case among connected arc diagrams for which the thick embedding (34) is not thin. The configuration in Lemma 7.8, Figure 17a is also a ladybug, which is disallowed for all J .

The following important lemma concerns the planar topology of an arc diagram D_e as in the hypothesis of Lemma 7.14.

Lemma 7.16. *Let P be a path through $u \leq_m v$ such that the H -resolution D_e consists entirely of two essential circles, with at least one past and one future arc connecting them. Let $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$ be generators such that $\partial_P \mathcal{M}_H(y, x)$ is nonempty. Then there exists a ray from the puncture \times to the point at infinity which is disjoint from all the arcs in D_e and which intersects the circles in D_e in exactly two points.*

Proof. Let a_p and a_f denote two past and future arcs, respectively, which join the two essential circles in D_e . We assume that a_p and a_f are adjacent, in the sense that they can be connected by an interval whose interior is disjoint from the arc diagram D_e . We will show that there exists such a ray passing between them, as in Figure 18a. Indeed, the only possible obstruction to such a ray would be an arc, or a sequence of interlacing arcs, along one of the essential circles with endpoints ‘outside’ of the arcs a_p, a_f (see Figure 18b with



(A) The desired ray (B) One obstructing arc (C) An obstructing sequence of interlaced arcs

FIGURE 18

single arc case and Figure 18c for the interlacing arc case). By considering the forbidden configurations of Figures 17a and 17d, we see that the arcs on one side of the essential circle must be either all past or all future. In either case, we will arrive at another forbidden configuration as in Figure 17b or Figure 17e. Thus no such arrangement may exist and our ray can be drawn. \square

We are now ready for the main result in this section.

Proposition 7.17. *Consider a subcube $u \leq_m v$ such that the arc diagram D_u is connected, and fix generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$. Then one of the following two holds.*

- (1) *For any path P from u to v such that $\partial_P \mathcal{M}_H(y, x) \neq \emptyset$, the thick embedding from (34) is in fact thin.*
- (2) *For any path P from u to v such that $\partial_P \mathcal{M}_H(y, x) \neq \emptyset$, the moduli space $\partial_P \mathcal{M}_H(y, x)$ consists of four points, and the thick embedding from (34) is a two-to-one map onto the points $\{\frac{1}{3}, \frac{2}{3}\} \subset (0, 1) \cong \mathcal{E}(v, u)$. Moreover, any two points in $\partial_P \mathcal{M}_H(y, x)$ that map to the same point in $\mathcal{E}(v, u)$ are oppositely framed.*

Proof. If all paths yield empty moduli spaces then there is nothing to show. Take a path P from u to v with $\partial_P \mathcal{M}_H(y, x) \neq \emptyset$. By Lemma 7.7, we may assume that the H -resolution D_e of P consists entirely of essential circles. We proceed by considering two cases for this fixed P .

Suppose first that D_e contains a pair of essential circles which are joined by both a past and future arc. We will argue that case (2) of the lemma holds. By Lemma 7.8, Figures 17c and 17f, we know that D_e contains exactly two essential circles. It was shown in Lemma 7.14 that the conclusion of case (2) holds for $\partial_P \mathcal{M}_H(y, x)$, so it remains to show that the same holds for every path from u to v . To that end, let P' be another path from u to v with nonempty moduli space, and denote its H -resolution by $D_{e'}$. As usual, it suffices to assume that $D_{e'}$ contains only essential circles. The arc diagram $D_{e'}$ can be obtained from D_e by performing a sequence of surgeries along the arcs in D_e . Then Lemma 7.16 and the assumption that $\partial_{P'} \mathcal{M}_H(y, x) \neq \emptyset$ guarantees that $D_{e'}$ contains two essential circles. Moreover, the ray constructed in the proof of Lemma 7.16, along with the forbidden configurations of Lemma 17, Figures 17a and 17d, imply that the two circles in $D_{e'}$ are joined by both a past and a future arc. Thus (2) holds by Lemma 7.14.

Now suppose that every pair of consecutive circles in D_e are joined by exactly one type (either past or future) of arc. The above argument guarantees that the H -resolution of every path with nonempty moduli space also has the property that each pair of consecutive circles

are joined by exactly one type of arc. Then we conclude that case (1) holds by Lemma 7.13. \square

Proposition 7.17 suffices to continue our construction. However, as in Section 7.3, we can completely describe the thickly embedded path moduli spaces for $J = H$. In Case (2) above, the moduli spaces are already described completely. Meanwhile, for Case (1), we have the following proposition; compare with Proposition 7.12.

Proposition 7.18. *Let $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$ such that we are in Case (1) of Proposition 7.17. Then for any path P with $\partial_P \mathcal{M}_H(y, x) \neq \emptyset$, the thin embedding φ_P from (34) surjects onto $\{\frac{1}{s+1}, \dots, \frac{s}{s+1}\}$ where s denotes the number of essential circles in D_e .*

Proof. Pick an element

$$(x \rightarrow \dots \rightarrow z_i \xrightarrow{j} z_{i+1} \rightarrow \dots \rightarrow y) \in \partial_P \mathcal{M}_H(y, x)$$

for some $1 \leq j \leq s$. Since H does not change the labels on essential circles, there must be an element of the form

$$(x \rightarrow \dots \rightarrow z_i \xrightarrow{k} z_{i+1} \rightarrow \dots \rightarrow y) \in \partial_P \mathcal{M}_H(y, x)$$

for every $1 \leq k \leq s$. \square

8. CONSTRUCTING THE HIGHER DIMENSIONAL MODULI SPACES (THE INDUCTIVE STEP)

In order to complete the proof of Theorem 5.8, we induct on the dimension of the subcube being considered. That is to say, we fix $m > 2$ and inductively assume that, for all subcubes $a \leq_i b$ with $i \leq m - 1$ and generators $r \in \Gamma_{Kh}(a)$, $s \in \Gamma_{Kh}(b)$, we have constructed moduli spaces $\mathcal{M}_J(s, r)$ together with thick embeddings

$$\mathcal{M}_J(s, r) \hookrightarrow \mathcal{E}(b, a)$$

satisfying all relevant boundary conditions. For any $a \leq_i b$ that does not involve J , this thick embedding is just

$$\mathcal{M}_J(s, r) \cong \coprod \mathcal{M}_Q(b, a) \xrightarrow{\text{Id}} \mathcal{M}_Q(b, a) = \mathcal{E}(b, a).$$

Meanwhile, for any $a \leq_i b$ that does involve J , the thick embedding is of codimension 1 due to the extra factor of $(0, 1)$ in the definition of $\mathcal{E}(b, a)$.

Now we fix an m -dimensional subcube $u \leq_m v$ and seek to build new moduli spaces $\mathcal{M}_J(y, x)$ thickly embedding into $\mathcal{E}(v, u)$ for various $x \in \Gamma_{Kh}(u)$, $y \in \Gamma_{Kh}(v)$. In the case that $u \leq_m v$ does not involve J , none of the subfaces of $u \leq_m v$ involve J either and the entire construction follows from [LS14a]. In this case, the ‘‘trivial cover’’ language of [LS14a] gives a thickened identity embedding

$$\mathcal{M}_J(y, x) \cong \coprod \mathcal{M}_Q(v, u) \xrightarrow{\text{Id}} \mathcal{M}_Q(v, u) = \mathcal{E}(v, u)$$

and we are done. Otherwise, the subcube $u \leq_m v$ involves J and we split into cases.

In each case, the essential strategy is the same. The boundary of $\mathcal{M}_J(y, x)$ is made up of pieces of various codimensional pieces $\partial_{[i]} \mathcal{M}_J(y, x)$, each of which consists of products of lower dimensional moduli spaces, which we inductively assume that we have already constructed and thickly embedded into the relevant ambient spaces. When we assemble these pieces together, we have a thick embedding of $\partial \mathcal{M}_J(y, x)$ into $\partial \mathcal{E}(v, u)$. Completing

the inductive step is then equivalent to proving that this thick embedding can be ‘thickly filled’ by a moduli space embedding $\mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, u)$ that satisfies Equation (25).

8.1. Subcubes involving E or F with connected arc diagrams. In the case that $J = E$ (the case $J = F$ is again dual to this case and left to the reader), we begin by considering subcubes $u \leq_m v$ whose corresponding arc diagram D_u is connected. We fix generators $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(v)$ and consider the codimension $(m - 1)$ (i.e. dimension 0) part of the boundary

$$\partial_{[m-1]}\mathcal{M}_E(y, x) = \coprod_{\text{paths } P} \partial_P \mathcal{M}_E(y, x).$$

If all of these path moduli spaces are empty, then we define $\mathcal{M}_E(y, x) := \emptyset$. Otherwise, Proposition 7.10 shows that any non-empty $\partial_P \mathcal{M}_E(y, x)$ is *thinly* embedded into $\mathcal{E}(P) \cong (0, 1)$.

Now consider $\partial_{[m-2]}\mathcal{M}_E(y, x)$, which Theorem 5.8 demands is of the form

$$\partial_{[m-2]}\mathcal{M}_E(y, x) = \bigcup \left(\coprod_{\substack{r \in \Gamma_{Kh}(a) \\ s \in \Gamma_{Kh}(b)}} \partial_{P^b} \mathcal{M}_E(y, s) \times \mathcal{M}_E(s, r) \times \partial_{P^a} \mathcal{M}_E(r, x) \right),$$

where the outer union is taken over

- all square sub-faces $a \leq_2 b$ in the subcube $u \leq_m v$, and
- all sub-paths P^a from u to a and P^b from b to v .

We will denote the parenthetical term by $\partial_{P^b P^a} \mathcal{M}_E(y, x)$. Then our (inductively defined) lower dimensional moduli space embeddings determine a thick product embedding

$$(39) \quad \partial_{P^b P^a} \mathcal{M}_E(y, x) \hookrightarrow \mathcal{E}(P^b) \times \mathcal{E}(b, a) \times \mathcal{E}(P^a) \cong [0, 1] \times (0, 1)$$

with boundary conditions based on the embeddings of the two path moduli spaces corresponding to composing P^a and P^b with the two paths around the edges of the square face $a \leq_2 b$.

$$(40) \quad \begin{array}{ccccccc} & & & & c & & \\ & & & & \swarrow & & \searrow \\ u & \longrightarrow & \cdots & \longrightarrow & a & & b & \longrightarrow & \cdots & \longrightarrow & v \\ & & & & \searrow & & \swarrow & & & & \\ & & & & c' & & & & & & \end{array}$$

$\underbrace{\hspace{10em}}_{P^a} \qquad \qquad \qquad \underbrace{\hspace{10em}}_{P^b}$

The key point is that, because the path moduli spaces on either boundary of $\partial_{P^b P^a} \mathcal{M}_E(y, x)$ are *thinly* embedded for our connected D_u , we can conclude that our composition moduli space embeddings (39) are thin as well. Thus we see that all of $\partial_{[m-2]}\mathcal{M}_E(y, x)$ is thinly embedded, and we move on to consider $\partial_{[m-3]}\mathcal{M}_E(y, x)$. This is again built from products of lower dimensional moduli spaces that we inductively assume have been constructed and thickly embedded in such a way that their boundaries are built out of parts of the thinly embedded $\partial_{[m-2]}\mathcal{M}_E(y, x)$ - and thus $\partial_{[m-3]}\mathcal{M}_E(y, x)$ must actually have been thinly embedded also.

This reasoning continues all the way until we consider the thinly embedded $\partial_{[1]}\mathcal{M}_E(y, x) = \partial \mathcal{M}_E(y, x)$. This is an $(m - 2)$ -dimensional manifold with corners $\partial \mathcal{M}_E(y, x)$ thinly embedded into the $(m - 1)$ -dimensional ambient space boundary $\partial \mathcal{E}(v, u) \cong S^{m-2} \times (0, 1)$. We may

view this as a codimension-1 framed embedding $\partial\mathcal{M}_E(y, x) \hookrightarrow S^{m-1}$ that misses the poles. Using the Pontryagin-Thom construction, this framed submanifold corresponds to a map $S^{m-1} \rightarrow S^1$. Since $m > 2$, the map is null-homotopic so it extends to a map $D^m \rightarrow S^1$. Therefore $\partial\mathcal{M}_E(y, x) \hookrightarrow S^{m-1}$ bounds a framed embedding of an $(m-1)$ -dimensional manifold into the interior of D^m , which is equivalent to the interior of $D^{m-1} \times (0, 1) \cong \mathcal{E}(v, u)$. It is worth noting that this is the point in the proof that uses (and propagates) the crucial codimension 1 inductive assumption. This is our construction of $\mathcal{M}_E(y, x)$ for connected D_u . See Figure 19.

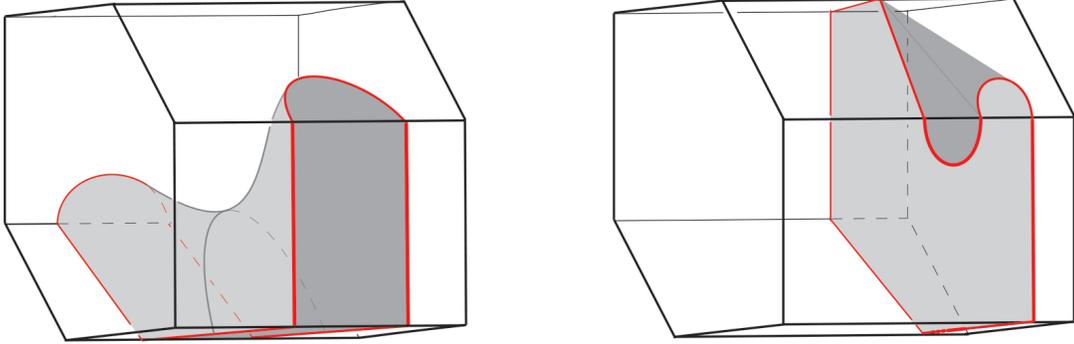


FIGURE 19. Examples of 2-dimensional $\mathcal{M}_E(y, x) \hookrightarrow \mathcal{E}(v, u) \cong \mathcal{M}_Q(v, u) \times (0, 1)$, where $\mathcal{M}_Q(v, u)$ is the 2-dimensional permutohedron (a hexagon, topologically D^2). The Pontryagin-Thom construction ensures that our framed embedded $\partial\mathcal{M}_E(y, x)$ (drawn in red) can be filled as shown, regardless of the presence of turnbacks (framing data is omitted to avoid clutter).

8.2. Subcubes involving E or F with disconnected arc diagrams. Next we consider the situation when $J = E$ and the diagram D_u is disconnected. Unlike in Section 8.1, path moduli spaces for subcubes with disconnected arc diagrams need not be thinly embedded. For example, adding a disjoint, contractible ladybug configuration ([LS14a, Figure 5.1]) has the effect of turning a nonempty thinly embedded path moduli space into one where the thick embedding is not injective. We begin with a simple lemma generalizing Lemma 6.1 from the $m = 2$ construction.

Lemma 8.1. *Suppose there is some path P for which $\partial_P\mathcal{M}_E(y, x) \neq \emptyset$. Then there is a single connected component \mathcal{C} of the arc diagram D_u such that*

$$\text{adeg}(y|_{\mathcal{C}}) = \text{adeg}(x|_{\mathcal{C}}) + 2, \quad \text{adeg}(y|_{\mathcal{C}'}) = \text{adeg}(x|_{\mathcal{C}'}),$$

where \mathcal{C}' is the complement of \mathcal{C} . We say that E is acting on the component \mathcal{C} . When $J = F$, the value 2 above is replaced by -2 .

Proof. As in the proof of Lemma 6.1, saddles maps preserve annular degree, but the E map increases annular degree by two. \square

As before, we consider

$$\mathcal{M}_E(y, x) = \mathcal{M}_E(y|_{\mathcal{C}} \otimes y|_{\mathcal{C}'}, x|_{\mathcal{C}} \otimes x|_{\mathcal{C}'}).$$

We let $\mathcal{M}_E^{\mathcal{C}}(y|_{\mathcal{C}}, x|_{\mathcal{C}})$ denote the moduli space that would be built for the diagram \mathcal{C} for generators $x|_{\mathcal{C}}, y|_{\mathcal{C}}$ if \mathcal{C}' were not present. We also include the E edge in the cube used

to construct $\mathcal{M}_E^{\mathcal{C}}(y|_{\mathcal{C}}, x|_{\mathcal{C}})$, since E is acting on \mathcal{C} . Because \mathcal{C} is connected, we can build $\mathcal{M}_E^{\mathcal{C}}(y|_{\mathcal{C}}, x|_{\mathcal{C}})$ using the techniques of Sections 5.2, 6.2, and 8.1. Furthermore, if \mathcal{C}' contains no saddles, then $y|_{\mathcal{C}'} = x|_{\mathcal{C}'}$ and we can define

$$\mathcal{M}_E(y, x) \cong \mathcal{M}_E^{\mathcal{C}}(y|_{\mathcal{C}}, x|_{\mathcal{C}})$$

as thickly embedded moduli spaces, and so we are done.

Therefore we may assume that \mathcal{C}' contains at least one saddle. In this case we let $\mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'})$ denote the moduli space that would be built for the diagram \mathcal{C}' for generators $x|_{\mathcal{C}'}, y|_{\mathcal{C}'}$ if \mathcal{C} were not present (not including the E edge, since E is not acting on \mathcal{C}'). Note that, because E is not involved, we can build $\mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'})$ as a trivial cover of a suitable permutohedron using the techniques in [LS14a].

Now consider that, in the large subcube $u \leq_m v$, there are many k -dimensional subfaces $a \leq_k b$ whose edges consist of all of the saddles in \mathcal{C}' . Indeed for any path P starting from u which consists of edges coming only from saddles in \mathcal{C} (or the E edge), the ending vertex of P can be taken as the starting vertex a of such a subface.

Lemma 8.2. *Let $k > 0$ denote the number of saddles in \mathcal{C}' and let $a \leq_k b$ denote any subface of the subcube $u \leq_m v$ consisting of only edges corresponding to these k saddles. Then for any choice of generator z on $\mathcal{C} \subset D_a$, we have a commuting diagram*

$$(41) \quad \begin{array}{ccc} \mathcal{M}_E(z \otimes (y|_{\mathcal{C}'}), z \otimes (x|_{\mathcal{C}'})) & \xleftarrow{\quad} & \mathcal{E}(b, a) \\ \parallel & & \\ \mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'}) & & \parallel \\ \parallel & & \\ \coprod \mathcal{M}_{\mathcal{Q}}(b, a) & \xleftarrow{\quad \coprod \text{id} \quad} & \mathcal{M}_{\mathcal{Q}}(b, a) \end{array}$$

Proof. Because the arbitrary generator z is kept constant throughout such a subcube $a \leq_k b$, the definitions make this clear. \square

Implicit in Lemma 8.2 is the fact that the diagram (41) commutes for any portion of the boundaries of the moduli spaces in question. Furthermore, this trivial cover remains consistent across all choices of $a \leq_k b$ and z .

The idea now is to assign a *label* to each component of this trivial cover

$$\mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'}) \xleftarrow{\quad \coprod \text{id} \quad} \mathcal{M}_{\mathcal{Q}}(b, a).$$

Then because \mathcal{C} is disjoint from \mathcal{C}' , any such label will be maintained across the entire subcube $u \leq_m v$ and we are free to build the large moduli space $\mathcal{M}_E(y, x)$ one label at a time.

For each fixed label, we can regard $\mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'})$ as a single copy of $\mathcal{M}_{\mathcal{Q}}(b, a)$ embedding via the identity map. Thus in any product used to build a composition moduli space, any factors coming from subfaces in \mathcal{C}' contribute only trivial identity embeddings. In this way we are free to build $\mathcal{M}_E(y, x)$ in the same manner as in the previous sections, constructing the boundary one codimension at a time. At each stage of this process, products with the trivial identity embeddings maintain thinness and we eventually arrive at a thinly embedded $\partial\mathcal{M}_E(y, x)$ of codimension one which can be filled using the Pontryagin-Thom construction as

before. Since this procedure is identical for each fixed label of component in $\mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'})$, we conclude that our eventual thick embedding

$$\mathcal{M}_E(y, x) \hookrightarrow \mathcal{E}(v, u)$$

contains a covering map of the same order as the trivial cover $\mathcal{M}^{\mathcal{C}'}(y|_{\mathcal{C}'}, x|_{\mathcal{C}'}) \xrightarrow{\prod \text{id}} \mathcal{M}_Q(b, a)$ and we are done.

8.3. Subcubes involving H with connected arc diagrams. In the case that $J = H$ and the subcube $u \leq_m v$ has corresponding arc diagram D_u that is connected, Proposition 7.17 splits us into two further cases. In case (1) of the proposition, we again have that all of our path moduli spaces are thin and the construction of $\mathcal{M}_H(y, x)$ proceeds in precisely the same manner as in Section 8.1 when $J = E$.

Therefore we focus on case (2) of Proposition 7.17, where for fixed $x \in \Gamma_{Kh}(u), y \in \Gamma_{Kh}(y)$, every non-empty path moduli space $\partial_P \mathcal{M}_H(y, x)$ consists of four points thickly embedded in two-to-one fashion onto $\{\frac{1}{3}, \frac{2}{3}\} \subset (0, 1) \cong \mathcal{E}(v, u)$. The proposition also shows that the two points embedded at any single $\frac{i}{3}$ are oppositely framed. Thus for any path P we can split $\partial_P \mathcal{M}_H(y, x)$ into two disjoint sets

$$\partial_P \mathcal{M}_H(y, x) = \partial_P^+ \mathcal{M}_H(y, x) \sqcup \partial_P^- \mathcal{M}_H(y, x),$$

where each $\partial_P^\pm \mathcal{M}_H(y, x)$ consists of two points, one embedded at $\frac{1}{3}$ with framing ± 1 , and the other embedded at $\frac{2}{3}$ with framing ∓ 1 . In particular, each $\partial_P^\pm \mathcal{M}_H(y, x)$ is *thinly* embedded into $\mathcal{E}(v, u)$.

The main idea then is to construct $\mathcal{M}_H(y, x)$ in two ‘signed’ pieces, which we denote by $\mathcal{M}_H^\pm(y, x)$. Each of these is built and thinly embedded into $\mathcal{E}(v, u)$ precisely as in Section 8.1 starting from the thinly embedded

$$\partial_{[m-1]} \mathcal{M}_H^\pm(y, x) := \prod_{\text{paths } P} \partial_P^\pm \mathcal{M}_H(y, x),$$

and ending with a codimension one Pontryagin-Thom argument. Note that the computations in Section 6.3 that are used to build $\partial_{[m-2]} \mathcal{M}_H(y, x)$ respect this decomposition into signed pieces, since any turnback described there must connect two oppositely framed points at $\frac{1}{3}$ and $\frac{2}{3}$. The processes for building $\mathcal{M}_H^+(y, x)$ and $\mathcal{M}_H^-(y, x)$ are thus identical except that all framings are reversed, and we can combine them into a two-to-one thick embedding

$$\mathcal{M}_H(y, x) = \mathcal{M}_H^+(y, x) \sqcup \mathcal{M}_H^-(y, x) \hookrightarrow \mathcal{E}(v, u).$$

Remark 8.3. The process above can be compared to Section 8.2, where identical thin constructions could be done one label at a time leading to thickly embedded moduli spaces. The difference here is that our labels $+$ and $-$ give rise to oppositely framed moduli spaces throughout.

8.4. Subcubes involving H with disconnected arc diagrams. Finally we consider subcubes $u \leq_m v$ involving the $J = H$ edge, where the corresponding arc diagram D_u is disconnected. The construction in this case involves a mixture of the ideas in Sections 8.2 and 8.3, which we will refer to freely without reiterating any details. As always, we begin by fixing $x \in \Gamma_{Kh}(u)$ and $y \in \Gamma_{Kh}(v)$ with some path P through $u \leq_m v$ having non-empty $\partial_P \mathcal{M}_H(y, x)$.

Suppose the arc diagram D_u has k homologically essential connected components, which we denote by $\mathcal{C}_1, \dots, \mathcal{C}_k$, ordered by their nesting in the annulus.

The first idea is to build $\mathcal{M}_H(y, x)$ one homologically essential connected component at a time, in the spirit of Section 8.3. To this end, note that each path moduli space $\partial_P \mathcal{M}_H(y, x)$ can be decomposed as

$$\partial_P \mathcal{M}_H(y, x) = \partial_P^{\mathcal{C}_1} \mathcal{M}_H(y, x) \sqcup \cdots \sqcup \partial_P^{\mathcal{C}_k} \mathcal{M}_H(y, x),$$

where each $\partial_P^{\mathcal{C}_i} \mathcal{M}_H(y, x)$ consists of only the points corresponding to H acting on essential circles in \mathcal{C}_i . This decomposition is clearly respected throughout the entire cube, allowing us to decompose our desired moduli space $\mathcal{M}_H(y, x)$ as

$$\mathcal{M}_H(y, x) = \mathcal{M}_H^{\mathcal{C}_1}(y, x) \sqcup \cdots \sqcup \mathcal{M}_H^{\mathcal{C}_k}(y, x).$$

Then fixing a connected component \mathcal{C}_i , we seek to build and thicken embed the space $\mathcal{M}_H^{\mathcal{C}_i}(y, x)$ starting from its codimension $(m - 1)$ boundary

$$\partial_{[m-1]} \mathcal{M}_H^{\mathcal{C}_i}(y, x) := \coprod_{\text{paths } P} \partial_P^{\mathcal{C}_i} \mathcal{M}_H(y, x)$$

as usual.

In order to do this, we appeal to the arguments in Section 8.2 letting \mathcal{C}_i take the role of \mathcal{C} as notated there. If the component \mathcal{C}_i corresponds to case (1) of Proposition 7.17, then this construction proceeds precisely as before and we produce a trivial cover according to the usual Lipshitz-Sarkar construction for generators on \mathcal{C}'_i , the complement of \mathcal{C}_i . If the component \mathcal{C}_i corresponds to case (2) of Proposition 7.17, then the procedures of Section 8.2 are done one signed label at a time, leading to a trivial cover with twice the order of the covering coming from Lipshitz-Sarkar's construction for generators on \mathcal{C}'_i .

In any case, each moduli space $\mathcal{M}_H^{\mathcal{C}_i}(y, x)$ can be built and thickly embedded into a copy of $\mathcal{E}(v, u)$ which we write as

$$\mathcal{M}_H^{\mathcal{C}_i}(y, x) \hookrightarrow \mathcal{E}(v, u) \cong \mathcal{M}_{\mathcal{Q}} \times (i - 1, i).$$

We then 'stack' all of these thick embeddings together as

$$\mathcal{M}_H(y, x) = \coprod_{i=1}^k \mathcal{M}_H^{\mathcal{C}_i}(y, x) \hookrightarrow \coprod_{i=1}^k \mathcal{M}_{\mathcal{Q}} \times (i - 1, i) \hookrightarrow \mathcal{M}_{\mathcal{Q}} \times (0, k) \cong \mathcal{E}(v, u)$$

to complete our construction. See Figure 20.

9. INVARIANCE AND NATURALITY WITH RESPECT TO LINK COBORDISMS

In order to complete the proof of Theorem 1.1, we must show that our construction of the map $\mathcal{J} : \mathcal{X}_{\mathbb{A}}(D) \rightarrow \mathcal{X}_{\mathbb{A}}(D)$, or equivalently our spectrum $\mathcal{X}_J(D)$ lifting $\text{Cone}(J)$, is invariant under the choices involved. The majority of these choices are considered in [LS14a] where invariance of the similarly constructed $\mathcal{X}_{Kh}(D)$ is shown, and the proofs there apply to our construction in the same manner. However, in Section 8 we made an additional choice of null-cobordism while appealing to the Pontryagin-Thom construction. We show below that $\mathcal{X}_J(D)$ is invariant under this choice as well.

Finally, we will also show that \mathcal{J} commutes, up to homotopy, with the maps on spectra assigned to link cobordisms. Proposition 9.5 below will concern elementary link cobordisms, including Reidemeister moves, and will thus conclude the proofs of both Theorem 1.1 and Theorem 1.3.

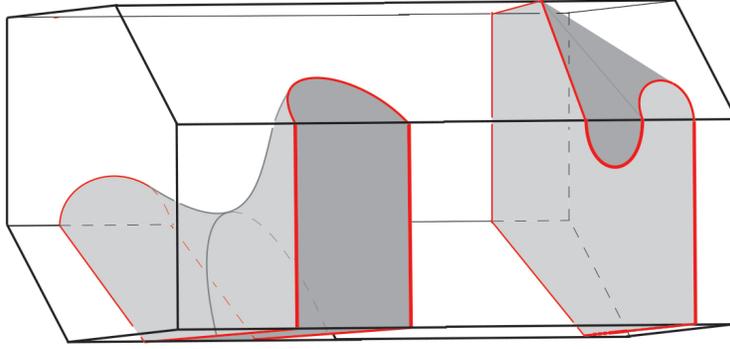


FIGURE 20. A possible embedding for $\mathcal{M}_H(y, x)$ for a disconnected diagram. The two components are constructed and embedded individually for H acting on different connected components of the arc diagram D_u . They are then horizontally ‘stacked’ together as shown.

9.1. Invariance under choices of cobordisms. The inductive step in the construction of higher-dimensional moduli spaces (see Section 8) relied on choosing a codimension-1 framed null-cobordism for the boundary. We will now show that different choices are in turn related by a framed codimension-1 cobordism. (Even though we work with codimension-1 orientable submanifolds where the framing can always be assumed to exist, we choose to explicitly specify the framing since it is central to the construction in the context of framed flow categories.)

Lemma 9.1. *Suppose $\mathcal{M}_J(y, x), \mathcal{M}'_J(y, x)$ are two moduli spaces constructed as in Section 8 using some choices of null-cobordisms. Then there exists a codimension-1 cobordism, that is a thick embedding*

$$\overline{\mathcal{M}}_J(y, x) \hookrightarrow \mathcal{E}(v, u) \times [0, 1]$$

whose boundary image in $\mathcal{E}(v, u) \times \{0\}$ (respectively $\mathcal{E}(v, u) \times \{1\}$) recovers the thick embedding $\mathcal{M}_J(y, x) \hookrightarrow \mathcal{E}(v, u)$ (respectively $\mathcal{M}'_J(y, x) \hookrightarrow \mathcal{E}(v, u)$).

Proof. The construction of $\overline{\mathcal{M}}_J(y, x)$ is done inductively in parallel with that of the moduli spaces. There is no choice involved in the construction of 0- and 1-dimensional moduli spaces, so in these cases $\overline{\mathcal{M}}$ is defined to be the product $\mathcal{M} \times I$. For the inductive step we focus first on the case when $J = E$ (or F) for connected diagrams as in Section 8.1.

Recall that we begin with the $(m-2)$ -manifold with corners $\partial\mathcal{M}_E(y, x)$ (for $m > 2$) thinly embedded into $\partial\mathcal{E}(v, u)$, viewed as S^{m-1} minus the poles. Consider $\partial\mathcal{M}_E(y, x) \subset S^{m-1} \times \{0\}$, $\partial\mathcal{M}'_E(x, y) \subset S^{m-1} \times \{1\}$. Just as in the arguments of Section 8.1, the inductive assumption of this lemma provides the existence of a manifold with corners (denote it $\partial\overline{\mathcal{M}}_E(x, y)$) thinly embedded as a framed cobordism for these boundaries in $S^{m-1} \times [0, 1]$. Assembling this data:

$$\mathcal{M}_E(y, x) \cup \partial\overline{\mathcal{M}}_E(x, y) \cup \mathcal{M}'_E(x, y),$$

we have a codimension-1 framed submanifold of $\partial(D^m \times I)$. The corresponding map $\partial(D^m \times I) \rightarrow S^1$ extends to $D^m \times I$, giving a desired cobordism in $\mathcal{E}(v, u) \times [0, 1]$ between $\mathcal{M}_E(y, x), \mathcal{M}'_E(x, y)$. Note that it extends the cobordism $\partial\overline{\mathcal{M}}_E(x, y)$ in $\partial\mathcal{E}(v, u) \times [0, 1]$ that was given by the inductive assumption.

The extension to the other cases, where the embeddings are indeed thick, is done in direct analogy to the constructions in Sections 8.2, 8.3, and 8.4, where all components of the trivial cover can be considered one at a time. \square

Next we will use Lemma 9.1 to show that stable homotopy types constructed using different choices of moduli spaces are homotopy equivalent.

Proposition 9.2. *Fix an annular link diagram D , and let $\mathcal{X}_J(D), \mathcal{X}'_J(D)$ denote two stable homotopy types built via Theorem 5.8 using different choices of null-cobordisms to construct the moduli spaces as described above. Then $\mathcal{X}_J(D) \simeq \mathcal{X}'_J(D)$.*

Proof. We work in the setting of [LS14a, Section 3.3]. An analogue of Lemma 3.16 in that reference applied to our construction in Lemma 9.1, together with the ability to ‘separate sheets of the cover’ as in the proof of Proposition 3.11, gives a collection of embeddings $\overline{\mathcal{M}}_J(y, x) \subset \mathbb{E}_d[\text{gr}(y) : \text{gr}(x)] \times [0, 1]$ which extend the neat embeddings $\iota_{y,x}, \iota'_{y,x}$ of $\mathcal{M}_J(y, x), \mathcal{M}'_J(y, x)$ into $\mathbb{E}_d[\text{gr}(y) : \text{gr}(x)] \times \{0\}$, respectively $\times \{1\}$. Informally, this may be thought of as a 1-parameter family of neat embeddings of flow categories, with singularities at times corresponding to critical points of the cobordisms $\overline{\mathcal{M}}$.

Consider the construction of the cell complex corresponding to a framed flow category [LS14a, Definition 3.24]. To each object y in the flow category, one associates a cell $\mathcal{C}(y)$. To define the attaching map of $\mathcal{C}(y)$ to a lower dimensional cell $\mathcal{C}(x)$, one uses the neat embedding of the moduli space $\mathcal{M}(y, x)$ to identify a subset $\mathcal{C}_x(y) \cong \mathcal{M}(y, x) \times \mathcal{C}(x) \subset \partial\mathcal{C}(y)$. The map is then defined on $\mathcal{C}_x(y)$ to be the projection to $\mathcal{C}(x)$, while sending all of $\partial\mathcal{C}(y) \setminus \cup_x \mathcal{C}_x(y)$ to the basepoint.

In our setting, the neat embedding of $\overline{\mathcal{M}}_J(y, x)$ into $\mathbb{E}_d[\text{gr}(y) : \text{gr}(x)] \times [0, 1]$ provides an identification of $\overline{\mathcal{M}}_J(y, x) \times \mathcal{C}(x) \subset (\partial\mathcal{C}(y)) \times [0, 1]$, and the corresponding projection $\overline{\mathcal{M}}_J(y, x) \times \mathcal{C}(x) \rightarrow \mathcal{C}(x)$ gives the desired homotopy between attaching maps corresponding to the neat embeddings of $\mathcal{M}_J(y, x)$ and $\mathcal{M}'_J(y, x)$. The constructions of the thickly embedded $\overline{\mathcal{M}}_J(y, x)$ in Lemma 9.1 give rise to homotopies of this sort which are coherent across all attaching maps of all cells, showing that our resulting spectra $\mathcal{X}_J(D)$ and $\mathcal{X}'_J(D)$ are stably homotopy equivalent. \square

9.2. Link cobordisms and Reidemeister invariance. Our goal in this section is to show that the \mathcal{J} map constructed in Corollary 5.9 commutes up to homotopy with the maps assigned to link cobordisms in [LS14b, Section 3]. This will include the case of link isotopies, indicating that \mathcal{J} commutes with the stable equivalences assigned to Reidemeister moves in [LS14a, Section 6]. We begin with some general notions about certain subcategories of flow categories that are used to construct the desired maps.

Let \mathfrak{F} be a framed flow category. Recall from [LS14a, Definition 3.29] the notion of a downwards (resp. upwards) closed subcategory $\mathcal{Z} \subset \mathfrak{F}$, which is a certain full subcategory whose geometric realization $|\mathcal{Z}|$ is naturally a subcomplex (resp. quotient complex) of $|\mathfrak{F}|$. Note that any full subcategory is specified by its objects. Thus if $\mathcal{Z} \subset \mathfrak{F}$ is downward closed, then it has a complementary upward closed subcategory \mathcal{Z}' whose objects are $\text{Ob}(\mathfrak{F}) \setminus \text{Ob}(\mathcal{Z})$. Moreover, upon geometric realization, $|\mathcal{Z}|$ is a subcomplex of $|\mathfrak{F}|$ with quotient complex $|\mathcal{Z}'|$, allowing us to write the cofibration sequence

$$(42) \quad |\mathcal{Z}| \hookrightarrow |\mathfrak{F}| \twoheadrightarrow |\mathcal{Z}'|.$$

We will say a full subcategory $\mathcal{Z} \subset \mathfrak{F}$ is *closed* if it is either upwards or downwards closed.

We now specialize to the situation at hand. Let D be an annular link diagram, and let $\mathfrak{F}_J(D)$ be the flow category refining Cone J (see Theorem 5.8). By definition, the objects of $\mathfrak{F}_J(D)$ consist of two copies of the objects of $\mathfrak{F}_\mathbb{A}(D)$,

$$\text{Ob}(\mathfrak{F}_J(D)) = \text{Ob}(\mathfrak{F}_\mathbb{A}(D))_0 \amalg \text{Ob}(\mathfrak{F}_\mathbb{A}(D))_1,$$

where we use the subscript 0 and 1 to indicate that some objects are before the J map (subscript 0), and some are after (subscript 1). The gradings of objects in $\text{Ob}(\mathfrak{F}_\mathbb{A}(D))_1$ are shifted up by 1 from the gradings in $\text{Ob}(\mathfrak{F}_\mathbb{A}(D))_0$.

Definition 9.3. For a full subcategory $\mathcal{Z} \subset \mathfrak{F}_\mathbb{A}(D)$, let $\overline{\mathcal{Z}}$ denote the full subcategory of $\mathfrak{F}_J(D)$ whose objects are two copies of the objects of \mathcal{Z} ,

$$\text{Ob}(\overline{\mathcal{Z}}) = \text{Ob}(\mathcal{Z})_0 \amalg \text{Ob}(\mathcal{Z})_1,$$

where $\text{Ob}(\mathcal{Z})_i \subset \text{Ob}(\mathfrak{F}_\mathbb{A}(D))_i$. We call $\overline{\mathcal{Z}}$ the *double* of \mathcal{Z} , and we let $\mathcal{Z}_0, \mathcal{Z}_1$ denote the two full subcategories of $\overline{\mathcal{Z}} \subset \mathfrak{F}_J(D)$ defined by

$$\text{Ob}(\mathcal{Z}_i) = \text{Ob}(\mathcal{Z})_i.$$

As with $\mathfrak{F}_J(D)$, we think of \mathcal{Z}_0 and \mathcal{Z}_1 as two parallel copies of \mathcal{Z} within the double $\overline{\mathcal{Z}}$, separated by the J edge direction. In this way, \mathcal{Z}_0 and \mathcal{Z}_1 are complementary downward and upward closed subcategories of $\overline{\mathcal{Z}}$, so that upon geometric realization, there is an inclusion and projection

$$|\mathcal{Z}_0| \hookrightarrow |\overline{\mathcal{Z}}| \twoheadrightarrow |\mathcal{Z}_1|.$$

Combining this with the canonical³ identifications $\mathcal{Z}_0 = \mathcal{Z}, \mathcal{Z}_1 = \Sigma\mathcal{Z}$ (the $\Sigma(-)$ denotes a grading shift as in Definition 3.16) allows us to write, for any full subcategory $\mathcal{Z} \subset \mathfrak{F}_\mathbb{A}(D)$, the cofibration sequence

$$(43) \quad |\mathcal{Z}| \hookrightarrow |\overline{\mathcal{Z}}| \twoheadrightarrow \Sigma|\mathcal{Z}|.$$

Meanwhile, if \mathcal{Z}' denotes the complementary full subcategory of \mathcal{Z} in $\mathfrak{F}_\mathbb{A}(D)$, then the complement of the double is the double of the complement, $\overline{\mathcal{Z}'} = \overline{\mathcal{Z}}$. However, if \mathcal{Z} is closed in $\mathfrak{F}_\mathbb{A}(D)$, then its double $\overline{\mathcal{Z}}$ need not be closed in $\mathfrak{F}_J(D)$, since the J map may take certain generators corresponding to objects of \mathcal{Z} into \mathcal{Z}' , or vice versa.

Lemma 9.4. *Suppose that \mathcal{Z} is a closed subcategory of $\mathfrak{F}_\mathbb{A}(D)$ whose double $\overline{\mathcal{Z}}$ is also a closed subcategory of $\mathfrak{F}_J(D)$. Then the cofibration sequences (42) for complementary pairs $\mathcal{Z}, \mathcal{Z}'$ and $\overline{\mathcal{Z}}, \overline{\mathcal{Z}'}$ commute with the cofibration sequences (43) for complementary copies within a double, as illustrated below for the case that \mathcal{Z} is downward closed.*

$$(44) \quad \begin{array}{ccccc} |\mathcal{Z}| & \hookrightarrow & |\overline{\mathcal{Z}}| & \twoheadrightarrow & \Sigma|\mathcal{Z}| \\ \downarrow & & \downarrow & & \downarrow \\ |\mathfrak{F}_\mathbb{A}(D)| & \hookrightarrow & |\mathfrak{F}_J(D)| & \twoheadrightarrow & \Sigma|\mathfrak{F}_\mathbb{A}(D)| \\ \downarrow & & \downarrow & & \downarrow \\ |\mathcal{Z}'| & \hookrightarrow & |\overline{\mathcal{Z}'}| & \twoheadrightarrow & \Sigma|\mathcal{Z}'| \end{array}$$

³The subcategories $\mathcal{Z}_0, \mathcal{Z}_1$ do not involve the J edge, so they are built in exactly the same way as \mathcal{Z} .

Proof. The inclusion maps, the quotient maps, and the canonical identifications involved are all defined by identifying objects of the categories with cells in the geometric realizations, so the definitions make the commutation clear. \square

Note that, whenever we do have a closed subcategory $\mathcal{Z} \subset \mathfrak{F}_{\mathbb{A}}(D)$ whose double $\overline{\mathcal{Z}}$ remains closed in $\mathfrak{F}_J(D)$ as in Lemma 9.4, each row in (44) gives rise to a ‘horizontal’ Puppe map which can be identified as the \mathcal{J} map restricted to the setting of its row. Similarly, each column in (44) also gives rise to a ‘vertical’ Puppe map.

An annular link cobordism can be decomposed into a sequence of so-called *elementary cobordisms* corresponding to either Morse moves or Reidemeister moves. To prove Theorem 1.3, it suffices to show that the map \mathcal{J} commutes up to homotopy with all maps assigned to such elementary cobordisms.

Proposition 9.5. *Let $\phi : \mathcal{X}_{\mathbb{A}}(D) \rightarrow \mathcal{X}_{\mathbb{A}}(D')$ be the map on annular Khovanov spectra assigned to an elementary annular link cobordism from D' to D as in [LS14b, Section 3]. Then the map \mathcal{J} commutes with ϕ up to homotopy:*

$$\begin{array}{ccc} \mathcal{X}_{\mathbb{A}}(D) & \xrightarrow{\mathcal{J}} & \mathcal{X}_{\mathbb{A}}(D) \\ \phi \downarrow & & \downarrow \phi \\ \mathcal{X}_{\mathbb{A}}(D') & \xrightarrow{\mathcal{J}} & \mathcal{X}_{\mathbb{A}}(D') \end{array}$$

In particular, if D, D' are two diagrams for the same annular link, then the stable equivalence $\mathcal{X}_{\mathbb{A}}(D) \cong \mathcal{X}_{\mathbb{A}}(D')$ commutes with \mathcal{J} up to homotopy.

Proof. Each elementary cobordism map arises from identifying a closed subcategory \mathcal{Z} in the framed flow category $\mathfrak{F}_{\mathbb{A}}(\tilde{D})$ for a suitable diagram \tilde{D} , leading to a cofibration sequence of the form (42) (or in the case of Reidemeister II and III, a finite sequence of such closed subcategories). We claim that each such closed subcategory \mathcal{Z} used in this way gives rise to a closed double $\overline{\mathcal{Z}} \subset \mathfrak{F}_J(\tilde{D})$. Indeed the saddle map arises from a diagram \tilde{D} having a crossing at the place of the saddle, so that \mathcal{Z} and its complement \mathcal{Z}' correspond to 0- and 1-resolutions of this crossing, which clearly remain closed upon doubling. Meanwhile, all of the other elementary cobordisms (cups, caps, and Reidemeister moves) identify closed subcategories \mathcal{Z} by fixing labels on *trivial* circles in their corresponding \tilde{D} . Since the J map does not affect the labels on trivial circles, any such closed subcategory remains closed upon doubling.

Thus any closed subcategory \mathcal{Z} used to define an elementary cobordism map ϕ gives rise to a commuting diagram of cofibrations as in (44). In all such cases, the map being built will be identified with one of the vertical maps in the first and/or last column (or perhaps the corresponding ‘vertical’ Puppe map), while the \mathcal{J} maps will be identified with the ‘horizontal’ Puppe maps, and these two maps will commute due to the naturality of the Puppe construction. \square

Remark 9.6. Conjecturally the maps ϕ assigned to link cobordisms in [LS14b, Section 3] are well-defined and do not depend (up to homotopy) on the decomposition into elementary cobordisms. If this is the case, the statement of Theorem 1.3 can be similarly strengthened as well.

REFERENCES

- [AGW15] Denis Auroux, J. Elisenda Grigsby, and Stephan M. Wehrli. Sutured Khovanov homology, Hochschild homology, and the Ozsváth-Szabó spectral sequence. *Trans. Amer. Math. Soc.*, 367(10):7103–7131, 2015.
- [AKW] Rostislav Akhmechet, Vyacheslav Krushkal, and Michael Willis. Stable homotopy refinement of quantum annular homology. *To appear in Compos. Math.* Preprint: <https://arxiv.org/abs/2001.00077>.
- [APS04] Marta M. Asaeda, Józef H. Przytycki, and Adam S. Sikora. Categorification of the Kauffman bracket skein module of I -bundles over surfaces. *Algebr. Geom. Topol.*, 4:1177–1210, 2004.
- [BN02] Dror Bar-Natan. On Khovanov’s categorification of the Jones polynomial. *Algebr. Geom. Topol.*, 2:337–370, 2002.
- [BN05] Dror Bar-Natan. Khovanov’s homology for tangles and cobordisms. *Geom. Topol.*, 9:1443–1499, 2005.
- [BPW19] Anna Beliakova, Krzysztof K. Putyra, and Stephan M. Wehrli. Quantum link homology via trace functor I. *Invent. Math.*, 215(2):383–492, 2019.
- [CJS95] R. L. Cohen, J. D. S. Jones, and G. B. Segal. Floer’s infinite-dimensional Morse theory and homotopy theory. In *The Floer memorial volume*, volume 133 of *Progr. Math.*, pages 297–325. Birkhäuser, Basel, 1995.
- [CK14] Yanfeng Chen and Mikhail Khovanov. An invariant of tangle cobordisms via subquotients of arc rings. *Fund. Math.*, 225(1):23–44, 2014.
- [GLW18] J. Elisenda Grigsby, Anthony M. Licata, and Stephan M. Wehrli. Annular Khovanov homology and knotted Schur-Weyl representations. *Compos. Math.*, 154(3):459–502, 2018.
- [Kho00] Mikhail Khovanov. A categorification of the Jones polynomial. *Duke Math. J.*, 101(3):359–426, 2000.
- [Lau10] Aaron D. Lauda. A categorification of quantum $\mathfrak{sl}(2)$. *Adv. Math.*, 225(6):3327–3424, 2010.
- [LLS20a] Tyler Lawson, Robert Lipshitz, and Sucharit Sarkar. Chen-Khovanov spectra for tangles. *To appear in Michigan Mathematical Journal*, 2020.
- [LLS20b] Tyler Lawson, Robert Lipshitz, and Sucharit Sarkar. Khovanov homotopy type, Burnside category and products. *Geom. Topol.*, 24(2):623–745, 2020.
- [LS14a] Robert Lipshitz and Sucharit Sarkar. A Khovanov stable homotopy type. *J. Amer. Math. Soc.*, 27(4):983–1042, 2014.
- [LS14b] Robert Lipshitz and Sucharit Sarkar. A refinement of Rasmussen’s s -invariant. *Duke Math. J.*, 163(5):923–952, 2014.
- [LS18] Robert Lipshitz and Sucharit Sarkar. Spatial refinements and Khovanov homology. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pages 1153–1173. World Sci. Publ., Hackensack, NJ, 2018.
- [ORS13] Peter S. Ozsváth, Jacob Rasmussen, and Zoltán Szabó. Odd Khovanov homology. *Algebr. Geom. Topol.*, 13(3):1465–1488, 2013.
- [QR18] Hoel Queffelec and David E. V. Rose. Sutured annular Khovanov-Rozansky homology. *Trans. Amer. Math. Soc.*, 370(2):1285–1319, 2018.
- [Rob13] Lawrence P. Roberts. On knot Floer homology in double branched covers. *Geom. Topol.*, 17(1):413–467, 2013.
- [See12] Cotton Seed. Computations of the Lipshitz-Sarkar Steenrod Square on Khovanov Homology. 2012. Preprint: <https://arxiv.org/abs/1210.1882>.
- [SSS20] Sucharit Sarkar, Christopher Scaduto, and Matthew Stoffregen. An odd Khovanov homotopy type. *Adv. Math.*, 367:107112, 51, 2020.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF VIRGINIA, CHARLOTTESVILLE VA 22904-4137
Email address: ra5aq@virginia.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF VIRGINIA, CHARLOTTESVILLE VA 22904-4137
Email address: krushkal@virginia.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, CA 90095
Email address: msw188@ucla.edu