Explicit Koszul-dualizing bimodules in bordered Heegaard Floer homology

BOHUA ZHAN

We give a combinatorial proof of the quasi-invertibility of $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ in bordered Heegaard Floer homology, which implies a Koszul self-duality on the dg-algebra $\mathcal{A}(\mathcal{Z})$, for each pointed matched circle $\mathcal{Z}$. We do this by giving an explicit description of a rank 1 model for $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$, the quasi-inverse of $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$. To obtain this description we apply homological perturbation theory to a larger, previously known model of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$.

1 Introduction

Bordered Heegaard Floer homology, introduced by Robert Lipshitz, Peter Ozsváth, and Dylan Thurston in [1; 4], gives a way of extending the hat version of Heegaard Floer homology to 3–manifolds with boundary. Besides its theoretical interest, it has shown to be an effective computational tool, for example in giving an efficient, algorithmic way to compute the Heegaard Floer homology of any 3–manifold [3].

The theory associates invariants to 3–manifolds with parametrized boundaries. A parametrization of the boundary is a diffeomorphism from the boundary to some standard genus g surface. A standard genus g surface is, in turn, described by a pointed matched circle, which can be considered as a handle decomposition of the surface. For a pointed matched circle $\mathcal{Z}$, we denote by $F(\mathcal{Z})$ the standard surface parametrized by $\mathcal{Z}$. Let $-\mathcal{Z}$ denote $\mathcal{Z}$ with reversed orientation, then $F(-\mathcal{Z})$ is the orientation reversal of $F(\mathcal{Z})$.

To each pointed matched circle $\mathcal{Z}$, the theory associates a combinatorially defined dg-algebra $\mathcal{A}(\mathcal{Z})$. To every 3–manifold Y with boundary parametrized by $\mathcal{Z}$, it associates two invariants: the type A invariant $\widehat{CFA}(Y)_{\mathcal{A}(\mathcal{Z})}$, which is a right A_∞–module over $\mathcal{A}(\mathcal{Z})$, and the type D invariant $^{\mathcal{A}(-\mathcal{Z})}\widehat{CFD}(Y)$, which is a left type D structure over $\mathcal{A}(-\mathcal{Z})$ (we will briefly review these algebraic concepts in the next section). The modules are well-defined up to homotopy equivalence (denoted by $\simeq$).

The type D and type A invariants are related to each other by taking the box tensor product, $\cdot \boxtimes \cdot$, with one of two special bimodules: $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ and $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$. They are called the identity type DD and type AA bimodules, respectively. More explicitly, the relations are

$$\tag{1} \widehat{CFA}(Y)_{\mathcal{A}(\mathcal{Z})} \simeq \widehat{CFAA}(\mathbb{I}_\mathcal{Z})_{\mathcal{A}(-\mathcal{Z}),\mathcal{A}(\mathcal{Z})} \boxtimes {}^{\mathcal{A}(-\mathcal{Z})}\widehat{CFD}(Y),$$

$$\tag{2} {}^{\mathcal{A}(-\mathcal{Z})}\widehat{CFD}(Y) \simeq \widehat{CFA}(Y)_{\mathcal{A}(\mathcal{Z})} \boxtimes {}^{\mathcal{A}(\mathcal{Z}),\mathcal{A}(-\mathcal{Z})}\widehat{CFDD}(\mathbb{I}_\mathcal{Z}).$$

The bimodule $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ is quasi-invertible, with $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$ being its quasi-inverse; that is, we have $\widehat{CFDD}(\mathbb{I}_\mathcal{Z}) \boxtimes \widehat{CFAA}(\mathbb{I}_\mathcal{Z}) \simeq \mathbb{I}$, where $\mathbb{I}$ denotes the identity type DA bimodule. This implies that taking box tensor product with $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ induces an equivalence of categories between the category of right A_∞–modules over $\mathcal{A}(\mathcal{Z})$ and the category of type D structures over $\mathcal{A}(-\mathcal{Z})$. In particular, the invariants $\widehat{CFA}(Y)_{\mathcal{A}(\mathcal{Z})}$ and ${}^{\mathcal{A}(-\mathcal{Z})}\widehat{CFD}(Y)$ actually contain the same information about Y.

Both $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$ and $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ are defined by holomorphic curve counts. While $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ can be described combinatorially (see [2; 3]), the quasi-invertibility of $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ is verified in [4] using holomorphic curve methods.

The aim of this paper is to describe combinatorially the type AA invariant $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$. More precisely, we construct an explicit rank 1 A_∞–bimodule $\mathcal{N}$ with the following property.

Theorem 1.1 $\mathcal{N} \boxtimes \widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ *has rank* 1. *Furthermore,* $\mathcal{N}$ *is quasi-invertible, in the sense that there exists a type DA bimodule* $\mathcal{N}'$, *such that* $\mathcal{N} \boxtimes \mathcal{N}'$ *is homotopy equivalent to the identity bimodule* ($\mathcal{N}'$ *is called the quasi-inverse of* $\mathcal{N}$).

This quickly leads to a combinatorial proof that $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ itself is quasi-invertible. By construction, the bimodule $\mathcal{N}$ is in the homotopy class of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$, so we know from analysis that it is the actual quasi-inverse of $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$. This stronger statement will be proved combinatorially by the author in [5].

Considering $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ as a left–right type DD bimodule ${}^{\mathcal{A}(\mathcal{Z})}\widehat{CFDD}(\mathbb{I}_\mathcal{Z})^{\mathcal{A}(\mathcal{Z})}$, its quasi-invertibility also implies a kind of Koszul self-duality of $\mathcal{A}(\mathcal{Z})$; see [2, Section 8] for details. One consequence of this Koszul duality is the existence of an A_∞ morphism from $\mathcal{A}(\mathcal{Z})$ to the cobar resolution of $\mathcal{A}(\mathcal{Z})$, inducing an isomorphism on homology. We will give some explicit computations of this map as an application.

Besides giving a more concrete understanding of Koszul duality in $\mathcal{A}(\mathcal{Z})$, an explicit description of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$ can be useful in various computations. In addition to $\widehat{CFD}$ and $\widehat{CFA}$, there are also bimodule invariants $\widehat{CFDD}$, $\widehat{CFDA}$ and $\widehat{CFAA}$ associated to any 3–manifold with two boundaries, in particular mapping cylinders of surface

diffeomorphisms. In general, it is easier to compute $\widehat{CFD}$ and $\widehat{CFDD}$, since it involves counting simpler holomorphic curves, and there are known methods to exploit the type DD structure equations. These methods are used to compute $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$ in [2; 3], and $\widehat{CFDD}(\tau)$ for any arcslide τ. With an explicit description of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$, we can then obtain descriptions of type A and type DA invariants with the same number of generators whenever the type D and type DD invariants can be computed. This is used in [5] to give a combinatorial construction and proof of invariance for $\widehat{HF}$ using bordered Floer theory.

The construction of the rank 1 model of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$ begins with a previously known model $\mathcal{M}$ of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$. The chain complex M underlying $\mathcal{M}$ has far more generators than is necessary for $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$. The tool used to reduce the number of generators is *homological perturbation theory*. To use this theory, we find a smaller chain complex N that is homotopy equivalent to M. Since the theory is over $\mathbb{F}_2$, we may take N to be the homology of M. Homological perturbation theory will construct an A_∞–bimodule $\mathcal{N}$ homotopy equivalent to $\mathcal{M}$, whose underlying chain complex is N.

To construct $\mathcal{N}$, we need the additional data that verifies the homotopy equivalence between M and N. That is, morphisms $f\colon M \to N$ and $g\colon N \to M$ such that $g \circ f = \mathbb{I}_N$, and a homotopy map $H\colon M \to M$ such that $\mathbb{I}_M + f \circ g = d \circ H + H \circ d$. Both f and g will become obvious after we describe the chain complex and its homology. So the focus of this paper will be on constructing H and verifying that it is indeed a homotopy.

When $\mathcal{Z}$ is the (unique) genus 1 pointed matched circle, the size of the chain complex M is small enough that H can be found directly. This is done in [3, Section 8.4]. The computation here works for pointed matched circles of any genus, and one can easily check that it agrees with the previous computation in the genus 1 case.

We now describe plans for the rest of this paper. In Section 2, we will review some algebraic concepts, and describe the initial, larger model $\mathcal{M}$ of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$. In Section 3 we describe the homotopy map H. In Section 4 we verify that it satisfies the homotopy equation. In Section 5, we apply homological perturbation theory to describe the rank-1 model $\mathcal{N}$ of $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})$, and prove the quasi-invertibility of $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$. Finally, in Section 6 we give an application calculating some homology classes in the cobar resolution of $\mathcal{A}(\mathcal{Z})$.

2 Algebraic preliminaries

In this section we briefly review some algebraic concepts, and describe the initial model $\mathcal{M}$ of the A_∞–bimodule $\widehat{CFAA}(\mathbb{I}_\mathcal{Z})_{\mathcal{A}(-\mathcal{Z}),\mathcal{A}(\mathcal{Z})}$.

We assume that the reader is familiar with pointed matched circles and the dg-algebra $\mathcal{A}(\mathcal{Z})$ associated to a pointed matched circle $\mathcal{Z}$. In most parts of this paper, we will have in mind some fixed pointed matched circle $\mathcal{Z}$. Then, a generator of $\mathcal{A}(\mathcal{Z})$ will be represented by upward-veering strands, and a generator of $\mathcal{A}(-\mathcal{Z})$ by downward-veering ones. Paired horizontal strands in the generator will be shown using dashed lines. Later on, when we are dealing solely with generators of $\mathcal{A}(\mathcal{Z})$, we will also omit the direction markers on strands.

After fixing a pointed matched circle $\mathcal{Z}$, we will write $\mathcal{A}$ and $\mathcal{A}'$ for the dg-algebras $\mathcal{A}(\mathcal{Z})$ and $\mathcal{A}(-\mathcal{Z})$. For any element $a \in \mathcal{A}$, its corresponding element in $\mathcal{A}'$ is denoted $\bar{a}$. In particular this applies to idempotents. If $i \in \mathcal{A}$ is an idempotent, we also define $o(i) \in \mathcal{A}$ to be the idempotent complementary to i.

Let A be a dg-algebra over a ground ring $\boldsymbol{k}$ (which will be the direct sum of copies of $\mathbb{F}_2$). Recall that a right A_∞–module $\mathcal{M}_A$ over A consists of a module M over $\boldsymbol{k}$, together with multilinear maps

$$m_{1,i}\colon M \otimes A^{\otimes i} \to M$$

for all $i \geq 0$, where $A^{\otimes i}$ denotes the tensor product of i copies of A, and $A^{\otimes 0}$ is just $\boldsymbol{k}$. All tensor products are implicitly taken over $\boldsymbol{k}$. These maps satisfy the A_∞ structure equation

$$(3) \quad 0 = \sum_{i+j=n} m_{1,j}(m_{1,i}(\boldsymbol{x},a_1,\ldots,a_i),a_{i+1},\ldots,a_n)$$

$$+ \sum_{i=1}^{n} m_{1,n}(\boldsymbol{x},a_1,\ldots,da_i,\ldots,a_n) + \sum_{i=1}^{n-1} m_{1,n-1}(\boldsymbol{x},a_1,\ldots,a_i a_{i+1},\ldots,a_n).$$

If we ignore those maps $m_{1,i}$ with $i > 0$, we get a chain complex, which is called the chain complex underlying $\mathcal{M}$.

A left type D structure $^A\mathcal{N}$ over A consists of a module N over $\boldsymbol{k}$, together with a map

$$\delta^1\colon N \to A \otimes N,$$

satisfying the type D structure equation

$$(4) \qquad (\mu_2 \otimes \mathbb{I}_N) \circ (\mathbb{I}_A \otimes \delta^1) \circ \delta^1 + (\mu_1 \otimes \mathbb{I}_N) \circ \delta^1 = 0.$$

Here $\mu_1 \colon A \to A$ and $\mu_2 \colon A \otimes A \to A$ denote differential and multiplication on the dg-algebra A (in keeping with the notation for the more general case, where A is an A_∞–algebra). Later on we will also call type D structures *modules*.

In bordered Floer theory, the ground ring $\boldsymbol{k}$ is the direct sum of copies of $\mathbb{F}_2$, one for each indecomposable idempotent in $\mathcal{A}(\mathcal{Z})$. For each dg-algebra $\mathcal{A}(\mathcal{Z})$ and each A_∞– module or type D structure, it is always possible to choose (often canonically) a set of generators over $\mathbb{F}_2$, such that each generator $\boldsymbol{x}$ has an indecomposable idempotent $i(\boldsymbol{x})$ satisfying $\boldsymbol{x} = i(\boldsymbol{x})\boldsymbol{x}$. Intuitively, we can think of algebras and modules as generated over $\mathbb{F}_2$, but each generator has an idempotent, such that idempotents are required to match in algebra actions and in structure equations like (3) and (4).

Given a right A_∞–module $\mathcal{M}_A$ and a left type D structure $^A\mathcal{N}$, the box tensor product $\mathcal{M}_A \boxtimes {}^A\mathcal{N}$ is a chain complex whose underlying vector space is $M \otimes N$, and whose differential is given by

$$(5) \qquad \partial(\boldsymbol{x} \otimes \boldsymbol{y}) = \sum_{k=0}^{\infty} (m_{1,k} \otimes \mathbb{I}_N)(\boldsymbol{x} \otimes \delta^k(\boldsymbol{y})).$$

Here $\delta^k \colon N \to A^{\otimes k} \otimes N$ is given by applying δ^1 repeatedly on the N factor k times, while applying the identity map on all A factors at each step. The sum in (5) is finite under certain boundedness conditions on $\mathcal{M}_A$ and $^A\mathcal{N}$. In bordered Floer theory, these boundedness conditions correspond to admissibility conditions on the Heegaard diagrams; see [1, Sections 2.4 and 4.4].

There are analogous definitions for bimodules, and the box tensor products between them. They are given in detail in [4]. Given dg-algebras A and A', a right A_∞– bimodule $\mathcal{M}_{A,A'}$ over A and A' is a module M over $\boldsymbol{k}$ with structure maps

$$m_{1,i,j} \colon M \otimes A^{\otimes i} \otimes A'^{\otimes j} \to M.$$

A type DA structure (or bimodule) $^A\mathcal{N}_{A'}$ over A and A' is a module N over $\boldsymbol{k}$ with structure maps

$$\delta_i^1 \colon N \otimes A'^{\otimes i} \to A \otimes N.$$

Finally, a type DD structure (or bimodule) $^{A,A'}\mathcal{N}$ over A and A' is a module N over $\boldsymbol{k}$ with structure maps

$$\delta^1 \colon N \to A \otimes A' \otimes N.$$

In each case, the structure maps satisfy a structure equation analogous to Equations (3) and (4). Note each generator of the bimodules has two idempotents, one for each algebra action. Also, we used the notational convention of writing each dg-algebra on

the side it acts on, with superscripts indicating type D action, and subscripts indicating A_∞ (or type A) action.

For any module or bimodule $\mathcal{M}$, we denote its opposite by $\overline{\mathcal{M}}$ (see [4, Definition 2.2.31]). Taking the opposite switches the side of all algebra actions (or equivalently, replaces the action of an algebra A by its opposite A^{opp}).

In all examples of chain complexes and modules that we will encounter in this paper, there is a canonical choice of generators of the underlying vector space. In such cases, we can describe the differentials or structure maps as a sum of *arrows*. For chain complexes, the differential is a sum of arrows $x \to y$, where x and y are generators of the chain complex. The arrow maps x to y, and all other generators to zero. Similarly, the structure map of a left type D module over A is a sum of arrows $x \to a \otimes y$, where x and y are generators of the module, and a is a generator of A. The structure map of a right A_∞–module over A is a sum of arrows of the form $m_{1,i}(x, a_1, \ldots, a_i) \to y$, where x and y are generators of the module, and $a_1, \ldots, a_i$ are generators of A. This terminology extends in a straightforward way to the various types of bimodules.

For dg-algebras A and B with ring of idempotents k, we say a bimodule M (of type DD, DA, or AA) over A and B has *rank* 1 if its underlying module over k is free of rank 1. There is a correspondence between rank-1 type DA bimodules $^B M_A$ with $\delta_1^1 = 0$ and A_∞–morphisms $\phi \colon A \to B$ (see [4, Lemma 2.2.50]). Given $\phi \colon A \to B$, the corresponding bimodule is denoted $^B[\phi]_A$. It has type DA actions $\delta_1^1 = 0$ and

$$\delta_{k+1}^1(\mathbf{1}, a_1, \ldots, a_k) = \phi(a_1, \ldots, a_k) \otimes \mathbf{1}.$$

We now describe in detail the initial model $\mathcal{M}$ of the A_∞–bimodule $\widehat{CFAA}(\mathbb{I}_{\mathcal{Z}})$, using the formula from [4, Proposition 9.2]:

$$
\begin{aligned}
(6) \qquad \widehat{CFAA}(\mathbb{I}_{\mathcal{Z}})_{\mathcal{A}, \mathcal{A}'} &= \mathrm{Mor}^{\mathcal{A}}({}_{\mathcal{A}'}^{\mathcal{A}}\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}}), {}^{\mathcal{A}}\mathbb{I}_{\mathcal{A}}) \\
&= \mathrm{Mor}^{\mathcal{A}}({}_{\mathcal{A}'}\mathcal{A}'_{\mathcal{A}'} \boxtimes {}^{\mathcal{A}, \mathcal{A}'}\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}}), {}^{\mathcal{A}}\mathbb{I}_{\mathcal{A}}) \\
&= \overline{\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})}^{\mathcal{A}', \mathcal{A}} \boxtimes {}_{\mathcal{A}}\mathcal{A}_{\mathcal{A}} \boxtimes {}_{\mathcal{A}'}\overline{\mathcal{A}'}_{\mathcal{A}'},
\end{aligned}
$$

where the second line expands the definition of ${}_{\mathcal{A}'}^{\mathcal{A}}\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$, and the third line uses the identity $\mathrm{Mor}(M, N) = \overline{M} \boxtimes N$. We begin by describing each of the three factors in the last line of (6).

The type DD structure ${}^{\mathcal{A}, \mathcal{A}'}\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$ is computed in [3]. It is generated over $\mathbb{F}_2$ by the set of pairs of complementary idempotents $i \otimes i'$, where $i \in \mathcal{A}$ and $i' = \overline{o(i)} \in \mathcal{A}'$.

The type DD action on $i \otimes i'$ is given by

$$
(7) \qquad \delta^1(i \otimes i') = \sum_{\substack{\xi \in \mathcal{C},\, ia(\xi) = a(\xi)j, \\ i'\overline{a(\xi)} = \overline{a(\xi)}j'}} (a(\xi) \otimes \overline{a(\xi)}) \otimes (j \otimes j').
$$

Here $\mathcal{C}$ is the set of chords on $\mathcal{Z}$, $a(\xi) \in \mathcal{A}$ is the algebra element formed by summing all ways of adding horizontal strands to ξ, and $\overline{a(\xi)}$ is the corresponding element in $\mathcal{A}'$. For example, an arrow in the genus 1 case is

$$
(8) \qquad \delta^1\left(\vcenter{\hbox{(diagram)}} \right) \to \left(\vcenter{\hbox{(diagram)}} \otimes \vcenter{\hbox{(diagram)}} \right) \otimes \vcenter{\hbox{(diagram)}},
$$

where

$$
i = \vcenter{\hbox{(diagram)}}, \quad j = \vcenter{\hbox{(diagram)}}, \quad \text{and} \quad a(\xi) = \vcenter{\hbox{(diagram)}}.
$$

Note that the generators of $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$ are in one-to-one correspondence with idempotents of $\mathcal{A}$. So $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$ has rank 1 as a module over $\boldsymbol{k}$.

On taking the opposite, the directions of the arrows are reversed while the coefficients are kept the same (that is, an arrow $\boldsymbol{x} \to a \otimes \boldsymbol{y}$ in $\mathcal{M}$ corresponds to the arrow $\boldsymbol{y} \to a \otimes \boldsymbol{x}$ in $\overline{\mathcal{M}}$). For later convenience we will also reverse the order of the two algebra actions, so that idempotents and algebra elements in $\mathcal{A}'$ come first.

As an example, the arrow in (8) gives rise to the following arrow in $\overline{\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})}^{\mathcal{A}', \mathcal{A}}$, which is $\delta^1(j' \otimes j) \to (\overline{a(\xi)} \otimes a(\xi)) \otimes (i' \otimes i)$:

$$
(9) \qquad \delta^1\left(\vcenter{\hbox{(diagram)}} \right) \to \left(\vcenter{\hbox{(diagram)}} \otimes \vcenter{\hbox{(diagram)}} \right) \otimes \vcenter{\hbox{(diagram)}}
$$

This is shown graphically in Figure 1.

The bimodule $_{\mathcal{A}}\mathcal{A}_{\mathcal{A}}$ is a left–right A_∞–bimodule with the same generators as $\mathcal{A}$. The A_∞–bimodule actions simply come from the dg-algebra actions on $\mathcal{A}$. That is,

$$
\begin{cases}
m_{0,1,0}(a) \to b & \text{for each generator } b \text{ in } da, \\
m_{1,1,0}(b;a) \to ba & \text{whenever } ba \neq 0, \\
m_{0,1,1}(a;b) \to ab & \text{whenever } ab \neq 0.
\end{cases}
$$

Here $m_{i,1,k}$ is the part of the A_∞–bimodule action with i algebra inputs on the left and k algebra inputs on the right. We write the algebra and module inputs in order, omitting any part with zero elements.
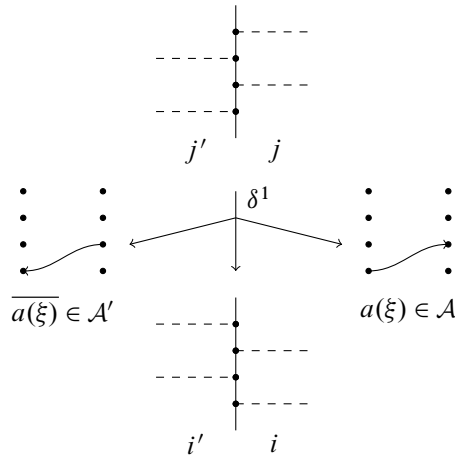
Figure 1: An example of an arrow in $\overline{\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})}^{\mathcal{A}',\mathcal{A}}$. Strands in $\mathcal{A}$ go upward and strands in $\mathcal{A}'$ go downward. For later convenience, we put $\mathcal{A}'$ on the left and $\mathcal{A}$ on the right, but remember that both are right type D actions.

The bimodule ${}_{\mathcal{A}'}\mathcal{A}'_{\mathcal{A}'}$ is defined the same way, based on the dg-algebra $\mathcal{A}'$. By taking opposites, the directions of the arrows are reversed. The algebra coefficients stay the same, but acting on the opposite side. So the actions on ${}_{\mathcal{A}'}\overline{\mathcal{A}'}_{\mathcal{A}'}$ can be written as

$$\begin{cases} m_{0,1,0}(b) \to a & \text{for each generator } b \text{ in } da, \\ m_{0,1,1}(ba;b) \to a & \text{whenever } ba \neq 0, \\ m_{1,1,0}(b;ab) \to a & \text{whenever } ab \neq 0. \end{cases}$$

We can now describe the bimodule $\mathcal{M}$ in (6). The underlying vector space is generated by triples $[a', i' \otimes i, a]$ such that $i' = \overline{o(i)}$, the idempotent i' agrees with the right idempotent of a', and i agrees with the left idempotent of a. Since i' and i are determined by a' and a, we will omit them and simply write the pair $[a', a]$. The condition on idempotents becomes that the right idempotent of a' is complementary to the left idempotent of a (using the identification between idempotents of $\mathcal{A}$ and $\mathcal{A}'$).

The differentials in ${}_{\mathcal{A}'}\overline{\mathcal{A}'}_{\mathcal{A}'}$ and ${}_{\mathcal{A}}\mathcal{A}_{\mathcal{A}}$ give rise to the following arrows in $\mathcal{M}$:

$$\begin{cases} [a', a] \to [a', b] & \text{for each generator } b \text{ in } da, \\ [b', a] \to [a', a] & \text{for each generator } b' \text{ in } da'. \end{cases}$$

There is a third type of arrow, coming from the type DD action on $\overline{\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})}^{\mathcal{A}',\mathcal{A}}$ combined with the $m_{1,1,0}$ actions on ${}_{\mathcal{A}}\mathcal{A}_{\mathcal{A}}$ and ${}_{\mathcal{A}'}\overline{\mathcal{A}'}_{\mathcal{A}'}$. The arrows are of the form

$$[a'\overline{a(\xi)}, a] \to [a', a(\xi)a]$$

for each chord ξ, and $a \in \mathcal{A}, a' \in \mathcal{A}'$ such that $a(\xi)a \neq 0$ and $a'\overline{a(\xi)} \neq 0$.

As an example, we show in Figure 2 how one such arrow follows from the definition of $\cdot \boxtimes \cdot$. This example involves the type *DD* arrow in Figure 1. From left to right, the three arrows composing the arrow in Figure 2 come from

- the A_∞–bimodule action $m_{1,1,0}$ on $\overline{\mathcal{A}'}$,
- the type *DD* action δ^1 on $\widehat{CFDD}(\mathbb{I}_\mathcal{Z})$, and
- the A_∞–bimodule action $m_{1,1,0}$ on $\mathcal{A}$.
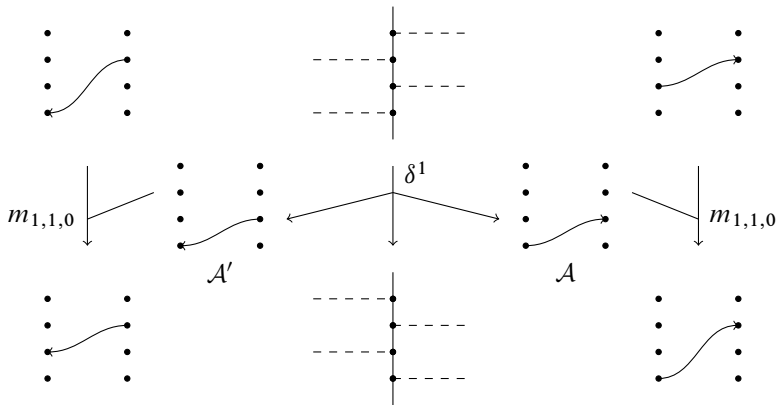
Figure 2: Example of the formation of an arrow in $\mathcal{M}$.

The right A_∞ actions on $\mathcal{M}$ are simply the ones inherited from those on $\mathcal{A}'$ and $\mathcal{A}$ (this is because $_\mathcal{A}\mathcal{A}_\mathcal{A}$ and $_{\mathcal{A}'}\overline{\mathcal{A}'}_{\mathcal{A}'}$ have no A_∞ actions with both left and right algebra inputs). They are given by the arrows

$$\begin{cases} m_{1,1,0}([a',a];b) \to [a',ab] & \text{whenever } ab \neq 0, \\ m_{1,0,1}([b'a',a];b') \to [a',a] & \text{whenever } b'a' \neq 0. \end{cases}$$

Here $m_{1,j,k}$ is the part of the A_∞–bimodule action with j inputs from $\mathcal{A}$ and k inputs from $\mathcal{A}'$. The first of the two equations comes from the action on $\mathcal{A}$, and the second from the action on $\mathcal{A}'$.

For simplicity of the discussion later, we will write everything in terms of elements of $\mathcal{A}$, so that the pair $[a',a]$ is written as the pair $[a_1,a_2]$, where $a_1 = \overline{a'}$ and $a_2 = a$. Translating the differential and the condition on idempotents, we arrive at the following statement.

Proposition 1 *The* A_∞*–bimodule* $\mathcal{M}$ *is generated by pairs* $[a_1, a_2]$*, where* a_1 *and* a_2 *are generators of* $\mathcal{A}$*, such that the left idempotent of* a_1 *is complementary to the left idempotent of* a_2*. There are three types of arrows in the differential:*

$$\begin{cases} [a_1, a_2] \to [a_1, b] & \text{for each generator } b \text{ in } da_2. \\ [b, a_2] \to [a_1, a_2] & \text{for each generator } b \text{ in } da_1. \\ [a(\xi)c, a] \to [c, a(\xi)a] & \text{for each chord } \xi \text{ and } a, c \in \mathcal{A} \text{ such that} \\ & \hspace{5em} a(\xi)c \neq 0 \text{ and } a(\xi)a \neq 0. \end{cases}$$

The A_∞*–bimodule action consists of the following arrows:*

$$\begin{cases} m_{1,1,0}([a_1, a_2]; a) \to [a_1, a_2 a] & \text{whenever } a_2 a \neq 0. \\ m_{1,0,1}([a_1 \overline{a'}, a_2]; a') \to [a_1, a_2] & \text{whenever } a_1 \overline{a'} \neq 0. \end{cases}$$

In the following sections, we will draw a generator of $\mathcal{M}$ by drawing a_1 and a_2 side by side. All strands are then going upward, so we omit directions on strands. Using this convention, examples of the three types of arrows in the differential are shown in Figure 3.
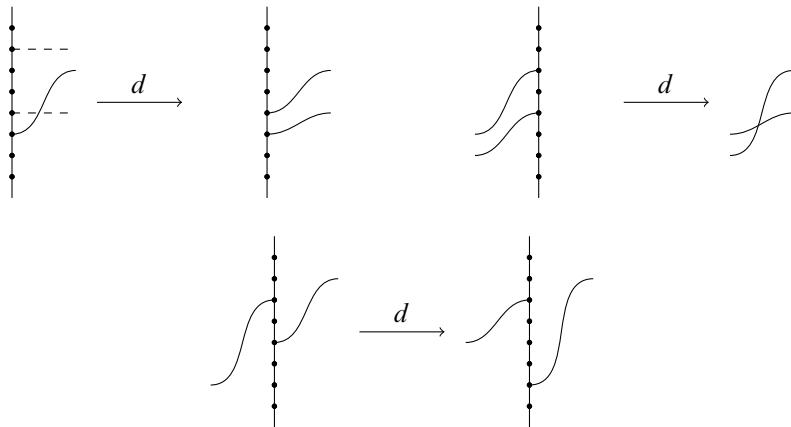
Figure 3: Examples of arrows in the chain complex M underlying $\mathcal{M}$.

3 Description of the homotopy

Let (M, d) be the chain complex underlying $\mathcal{M}$, described in the previous section. There is a distinguished set of generators of M, consisting of $[a_1, a_2]$, where both a_1 and a_2 are idempotents (which are necessarily complementary). By abuse of notation we will also call these generators *idempotents*. It is clear that there are no arrows in the differential in or out of these generators. Let $(N, d' = 0)$ be the subcomplex of M

generated by the idempotents. We will show that N is the homology of M, so that N and M are homotopy equivalent. There are obvious chain maps $f: N \to M$ and $g: M \to N$, where f is the inclusion map and g is the map sending the idempotents to themselves and all other generators to zero. It remains to find a homotopy $H: M \to M$ satisfying

$$d \circ H + H \circ d = \mathbb{I}_M + f \circ g.$$

We set $H(\boldsymbol{x}) = 0$ for all idempotents $\boldsymbol{x}$. Then the above equation holds trivially on idempotents. On the other generators, we have $f \circ g = 0$, so the equation reduces to $d \circ H + H \circ d = \mathbb{I}_M$.

We now describe H on the non-idempotents, and verify the required relation in the next section. From now on all generators that we mention are assumed to be a non-idempotent. We will call the arrows in the differential of M the d–arrows, and the arrows in the map H the H–arrows.

For a generator $[a_1, a_2]$ in C, we define the multiplicity $\text{mult}([a_1, a_2]) \in H_1(Z \setminus z, \boldsymbol{a})$ to be the sum of multiplicities of a_1 and a_2. Note that this is invariant under the differential. So the chain complex splits into disjoint parts according to multiplicity, and to specify the homotopy, it is sufficient to do so on each part. We say an element $[a_1, a_2]$ has multiplicity one if its multiplicity on each component of $Z \setminus \boldsymbol{a}$ is at most one. The construction of H for generators of multiplicity one is different from that for other generators.

3.1 Multiplicity one

We begin with generators of multiplicity one. Let S be the set of components of $Z \setminus \boldsymbol{a}$ not containing the basepoint z (that is, the generators of $H_1(Z \setminus z, \boldsymbol{a})$). There is an ordering $<_{\mathcal{Z}}$ on S, depending on the pointed matched circle $\mathcal{Z}$, given as follows. Attach handles to Z according to the matching on $\boldsymbol{a}$. Then, traverse the boundary of Z with handles attached, starting from the top (with Z oriented so that strands in $\mathcal{A}(\mathcal{Z})$ go upwards). Order S according to when we encounter each element. Note that we will encounter each element of S exactly once before reaching the bottom, due to the condition that must be satisfied for the matching on a pointed matched circle. An example is shown in Figure 4.

Let $(p, p + 1)$ be an element of S (with $p + 1$ above p). The segment immediately before $(p, p + 1)$ in the ordering $<_{\mathcal{Z}}$ can be found as follows. Let q be the point paired with $p + 1$ in $\mathcal{Z}$. If q is the topmost point of $\boldsymbol{a} \subset Z$, then $(p, p + 1)$ is the initial segment. Otherwise $(q, q + 1)$ is the segment prior to $(p, p + 1)$. See Figure 5 for a demonstration.
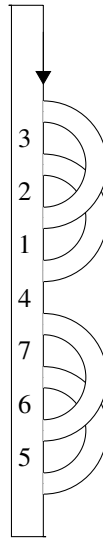
Figure 4: Example of traversal on the split pointed matched circle with genus 2. The numbers denote the ordering $<_{\mathcal{Z}}$ on S. Note that strands in $\mathcal{A}(\mathcal{Z})$ go upwards.
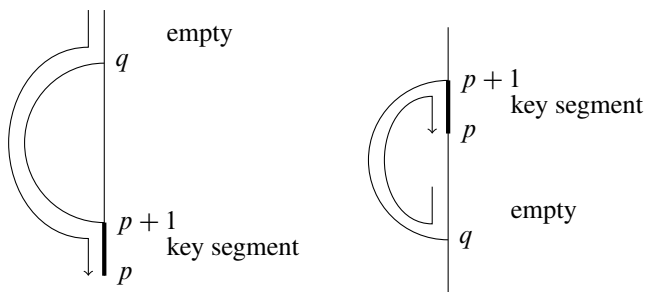

Figure 5: Identifying the segment prior to $(p, p + 1)$. In the first case the segment $(q, q + 1)$ may be the component of Z containing the basepoint.

Given a generator $\boldsymbol{x} = [a_1, a_2]$ with multiplicity one, we define the *key segment* of $\boldsymbol{x}$ to be the first segment in S, according to the ordering $<_{\mathcal{Z}}$, at which $\mathrm{mult}(\boldsymbol{x})$ is one. Let $(p, p + 1)$ be the key segment, and q be the point paired with $p + 1$. Define the pair $\{p + 1, q\}$ to be the *key pair* of $\boldsymbol{x}$. The main property that results from this construction is that while $\boldsymbol{x}$ has multiplicity 1 at $(p, p + 1)$, it must have multiplicity zero at $(q, q + 1)$ (one possibility is that q is the topmost point).

In the following, given a generator $[a_1, a_2]$, we say that the key segment is occupied on the left if it is covered by a_1, and on the right if it is covered by a_2. Since the left

idempotents of a_1 and a_2 are complementary, one of them contains the key pair. We say the key pair is occupied on the left if it is contained in the left idempotent of a_1, and on the right if it is contained in the left idempotent of a_2. We also distinguish whether the algebra element realizes the key pair as a double-horizontal or as the start of a non-horizontal (moving) strand. Note that in the latter case, since the multiplicity at $(q, q + 1)$ is zero, the strand must start at $p + 1$.

The above classification divides all (non-idempotent) generators into eight types. We will define $H(x)$ based on the type of x.

In the following, let $a \to b$ denote the strand starting at a and ending at b (so we always have $a < b$). By *moving the strand $a \to b$ to the left*, we mean that starting from $[a_1, a_2]$, we construct a new generator $[a'_1, a'_2]$, with the strands in a'_2 obtained from those in a_2 by either removing the strand $a \to b$, or by factoring $a \to b$ from the end of a longer strand in a_2; and with the strands in a'_1 obtained from those in a_1 by adding the strand $a \to b$. If rather than adding the strand $a \to b$, we wish to attach it to an existing strand in a_1, we will state so explicitly. Similarly we have the notion of moving the strand $a \to b$ to the right. Sometimes we will do two moves at the same time (with the intermediate state possibly not valid strand diagrams). The location of double horizontals on a'_1 and a'_2 is usually clear (noting that the left idempotents of a'_1 and a'_2 must be complementary). We will clarify it when it is ambiguous.

In all figures illustrating multiplicity one cases, we will use a dotted parenthesis to denote the key pair. We will also show q to be above rather than below $p + 1$, so the key segment is just below the dotted parenthesis, but the definition of H is the same in both cases.

For exactly four types of x, we have $H(x) \neq 0$. The values of $H(x)$ in these four types are as follows (see Figure 6):

(1) If the key segment is occupied on the left and the key pair is occupied on the left as a double-horizontal, then there is a strand $a \to b$ on the left with $a < p+1 < b$. Set H to resolve the crossing involving the horizontal strand at $p + 1$.

(2) If the key segment is occupied on the right and the key pair is occupied on the right by a moving strand, then there must be strands $i \to p + 1$ and $p + 1 \to j$ on the right. Set H to replace these two strands with the strand $i \to j$ and the double-horizontal at $\{p + 1, q\}$.
If there is a strand ending at q on the left, then there is an additional term in H, moving the strand ending at q to the right and the strand $i \to p + 1$ to the left.

(3) If the key segment is occupied on the right by a strand $i \to j$, and the key pair is occupied on the left as a double-horizontal, then H factors the strand $i \to p + 1$ from the right and moves it to the left.

(4) If the key segment is occupied on the right, and the key pair is occupied on the left by a moving strand, then there must be a strand $p + 1 \to j$ on the left and a strand $i \to p + 1$ on the right. Set H to move the strand $i \to p + 1$ to the
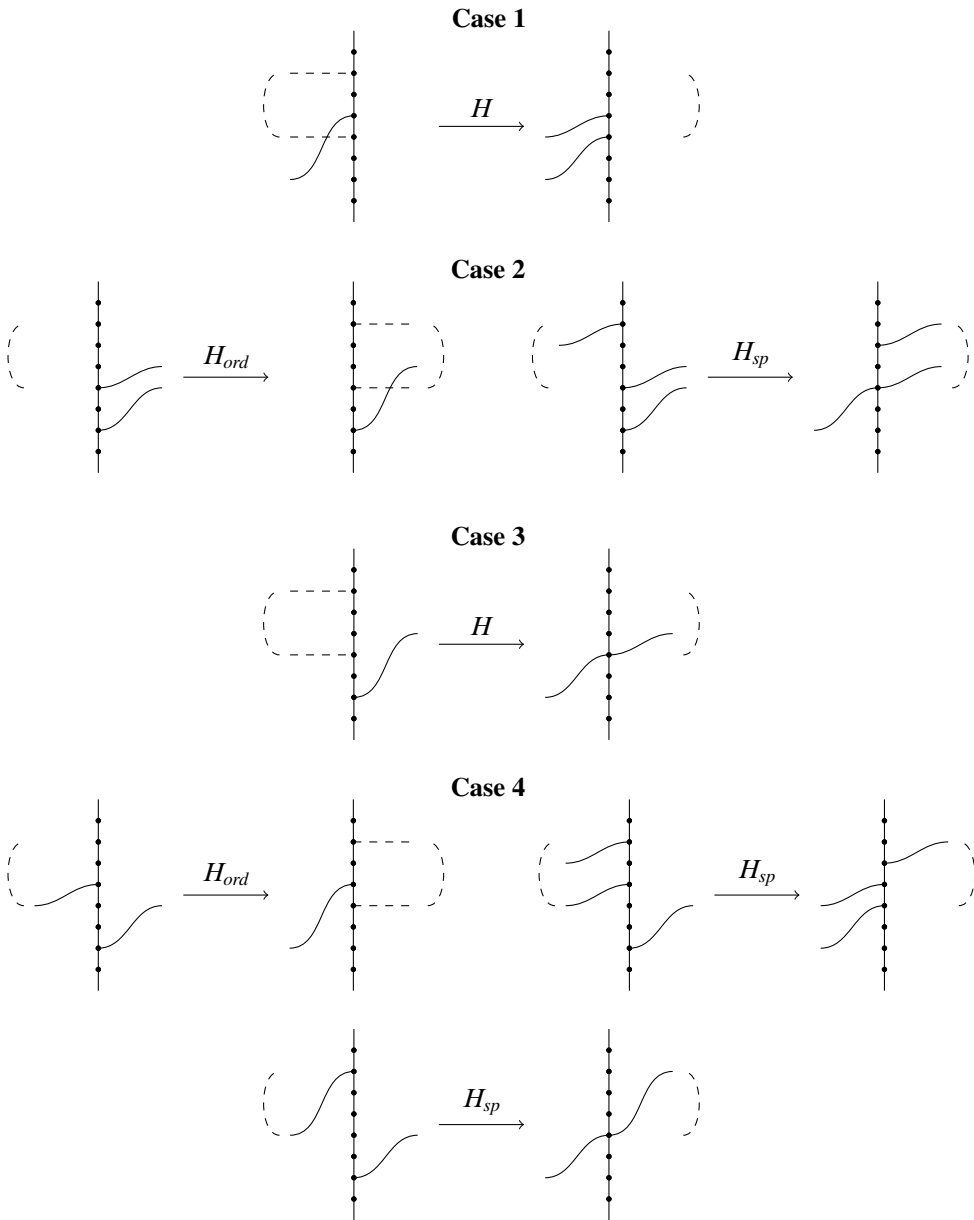

Figure 6: Diagrams for the homotopy map in multiplicity-one cases.

left attaching it to the $p + 1 \to j$ strand, and leaving a double horizontal at $\{p + 1, q\}$ at right.

There are two special cases: first, if there is a strand $j \to q$ with $j \neq p + 1$ on the left, then H contains an additional term moving $j \to q$ to the right and $i \to p + 1$ to the left. Second, if there is a strand $p + 1 \to q$ on the left (possible only if $q > p + 1$), then H contains an additional term moving $p + 1 \to q$ to the right and $i \to p + 1$ to the left.

The overall picture is as follows. We partition all generators of multiplicity one into ordered pairs, such that for each ordered pair (x_i, y_i), there is a d–arrow $x_i \to y_i$. The part of H not including the three special cases maps each y_i to x_i and x_i to zero. In verifying the relation $d \circ H + H \circ d = \mathbb{I}_M$, the compositions $d \circ H \colon y_i \to x_i \to y_i$ and $H \circ d \colon x_i \to y_i \to x_i$ account for the identity. If (x_i, y_i) is an ordered pair, we say x_i is on the d–side and y_i is on the H–side. The three special cases in H are additional maps from the H–side to the d–side (that is, mapping y_i to x_j for some $i \neq j$). Intuitively, the part of H mapping each y_i to x_i performs the inverse of a d–arrow around $p + 1$. The special arrows perform a different move around $p + 1$, along with moving a strand ending at q to the right.

Whether a generator is on the d–side or the H–side depends solely on the type of the generator. The eight types are summarized in Table 1. The numbers after d or H indicate which case they will be covered under in the proof in the next section, and if labeled H, also the case of H in the above description.

key seg. ╲ key pair	left, double-hor.	right, double-hor.	left, moving	right, moving
left	H–1	d–7	d–5	d–8
right	H–3	d–6	H–4	H–2

Table 1: Summary of cases for multiplicity one generators.

It would be simpler if there were no special cases in H; that is, if the only H–arrows were those from the H–side to the d–side of the same pair. However, this is impossible as demonstrated in Figures 7 and 8. Each figure shows four generators x, y, z, w such that

$$dx = y + w \quad \text{and} \quad dz = w,$$

with no other d–arrows involving these generators. This forces the homotopy to be

$$Hy = x + z \quad \text{and} \quad Hw = z.$$
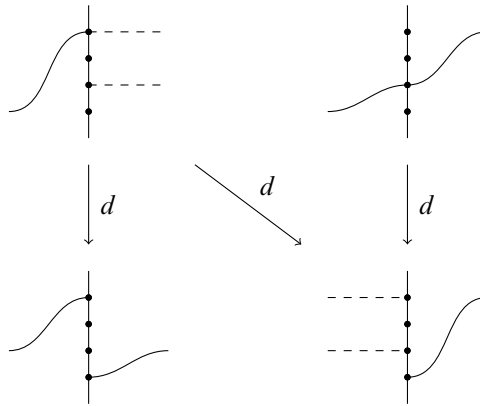
Case 4, second special H.

Figure 7: Examples of why special cases of homotopy are needed, part 1.

3.2 Other generators

We now consider generators that do not satisfy the multiplicity one condition. Let $x = [a_1, a_2]$ be such a generator. This means mult(x) is greater than 1 at some segment in S. Let $[i, i+1]$ be the bottom-most segment with multiplicity greater than 1 (note that we are no longer using the ordering $<_{\mathcal{Z}}$ on S). Since two strands cannot start from the same point i, this segment must have multiplicity 2, and the segment $[i-1, i]$ must have multiplicity 1. There is a unique strand starting at i (call it the i strand) and
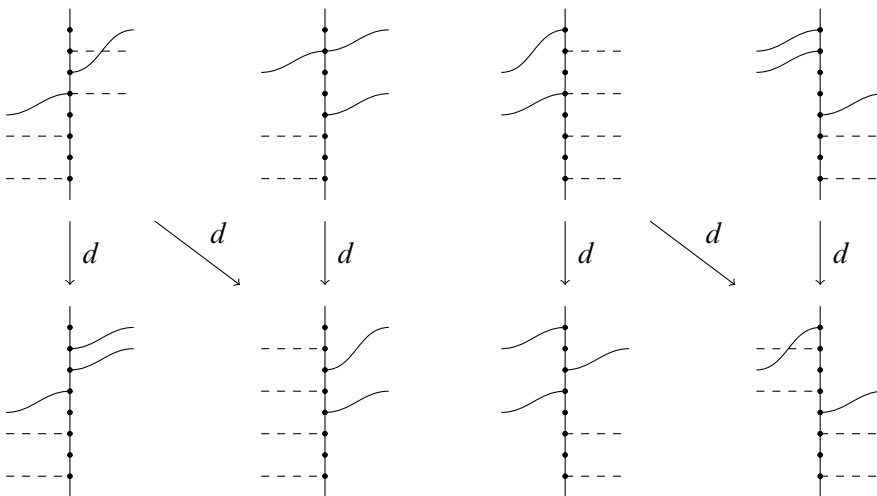
Case 2, special H Case 4, first special H.

Figure 8: Examples of why special cases of homotopy are needed, part 2.

a unique strand that covers $[i-1, i]$. Let j be the starting point of the strand covering $[i-1, i]$, and call that strand the j strand. Note that i, but not j, is unchanged in any d–arrow.

The definition of H consists of the following cases (see Figure 9):

(1) If both the i strand and the j strand are on the left, and they cross each other, the H–arrow uncrosses the two strands.

(2) If both the i strand and the j strand are on the right, and they do not cross, the H–arrow crosses the two strands.

(3) If the i strand is on the left and the j strand is on the right, the H–arrow factors the strand $j \to i$ from the right and moves it to the left attaching to the i strand.
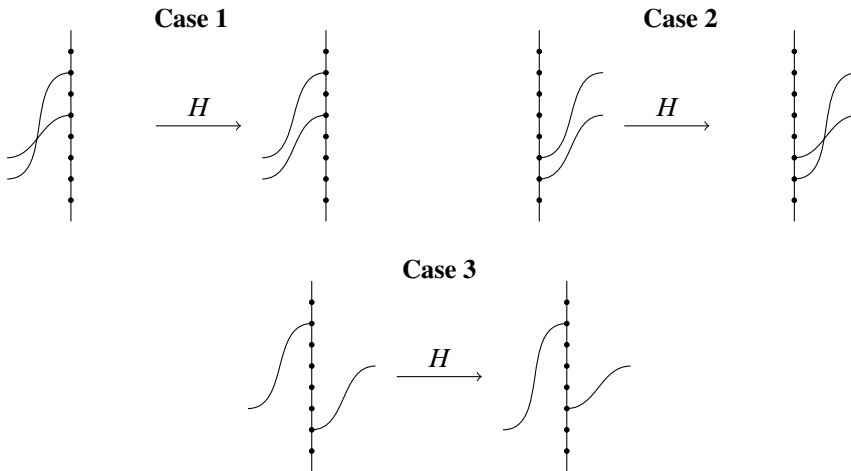
Figure 9: Diagrams for the homotopy map in the non-multiplicity-one cases.

On all other generators H is zero. The cases are summarized in Table 2. There is again a pairing of generators $(\boldsymbol{x}_i, \boldsymbol{y}_i)$, so that for each pair there is a d–arrow $\boldsymbol{x}_i \to \boldsymbol{y}_i$. In this case, there are no special cases in H, so H is exactly the map sending $\boldsymbol{y}_i$ to $\boldsymbol{x}_i$ and $\boldsymbol{x}_i$ to 0 in each pair $(\boldsymbol{x}_i, \boldsymbol{y}_i)$.

4 Verification of the homotopy

4.1 Multiplicity one

In this section we verify that the map H defined above satisfies the equation $d \circ H + H \circ d = \mathbb{I}_M$ on the non-idempotent part of M, beginning with the multiplicity one

i strand $\diagdown$ j strand	left	right
left	(d,H)-1	H-3
right	d-4	(d,H)-2

Table 2: Summary of cases for generators not of multiplicity one.

case. We do so by checking the equation

$$d \circ H + H \circ d = \mathbb{I}_M \tag{10}$$

on each of the eight types of generators. Depending on the type, the generator appears on one side of some pair (x_i, y_i). For every such pair, there is a d–arrow $x_i \to y_i$ and an H–arrow $y_i \to x_i$. So there is a natural identity term in $d \circ H + H \circ d$, and it remains to check that all other terms sum to zero.

Write $H = H_{ord} + H_{sp}$, where H_{sp} contains arrows coming from the three special cases of H, and H_{ord} is the remaining part of H. So for each pair (x_i, y_i) we have $H_{ord} y_i = x_i$. It is also not difficult to determine y_i from x_i: simply take the obvious differential around $p + 1$. We call a generator *special* if it appears on either side of a pair (x_i, y_i) such that $H_{sp} y_i \neq 0$. One easily checks that a generator is special if and only if there is a strand ending at q on the left. Generators of type H-1, H-3, d-5 and d-8 can never be special (the first two due to a double horizontal on the left, the last two due to a strand ending at $p + 1$ on the left).

The strategy for verifying (10) is different for generators on the H side and the d side, so we will describe them separately. For a generator y on the H–side, we define three sets of generators of M. Let G_1 be the set of terms in dy that are on the H side, G_2 be the set of all terms in $d(H_{ord} y)$ excluding y, and G_3 be the set of all terms in $d(H_{sp} y)$. For a non-special y, G_3 is of course empty. We will show that in this case all generators in G_1 are also non-special, and that H_{ord} induces a bijection between G_1 and G_2. For special y we will show the following:

- There is a generator y_o common to both G_2 and G_3.
- All generators in G_1 are also special.
- The map H_{ord} induces a bijection between G_1 and $G_2 \setminus \{y_o\}$.
- The map H_{sp} induces a bijection between G_1 and $G_3 \setminus \{y_o\}$.

These imply that the terms in $d \circ H + H \circ d$ other than the natural identity term sum to zero.

A large part of the bijection can be considered "trivial", since it involves d–arrows that do not modify strands around $p+1$ or q. They can be carried out in the "same" way on either side of an H–arrow. So a term in $d\,y$ coming from a d–arrow away from $p+1$ and q corresponds to the term obtained by performing the "same" move in $d(H_{ord}\,y)$ or $d(H_{sp}\,y)$. Since H–arrows can introduce double horizontals at $\{p+1,q\}$ on the right side, it can affect whether a strand ending at q on the left can be moved to the right in a d–arrow. This is essentially the reason why special cases are needed.

The main checks are therefore for differentials involving strands around $p+1$ and q. The cases are described below, and the reader is advised to follow along using the illustrations in the online supplement [6]. In each cancellation diagram in the supplement, the top-left generator is y, the top-right generator is either $H_{ord}\,y$ or $H_{sp}\,y$ (as indicated on the top edge). The bottom-left generator is an element in G_1, and the bottom-right generator is the corresponding element in G_2 or G_3. For generators that are special, there is one more diagram showing the cancellation involving y_o. In these diagrams, y is on the top left, $H_{ord}\,y$ is on the bottom left, $H_{sp}\,y$ is on the top right, and $y_\mathbf{o}$ is on the bottom right.

We now turn to the case of a generator x on the d–side. The first term in $d\circ H + H\circ d$ is zero. There is a distinguished term y in dx: the generator paired with x. If x is non-special, then y is also non-special, and we show that this is the only term in dx on the H side. If x is special, then y is also special, and we show there is exactly one other, non-special, term y' in dx on the H side, and that $H_{sp}\,y = H\,y'$. This verifies the equation for x. In the supplement, we will show a diagram for each type of the last cancellation, with x on the top-left, y on the bottom-left, y' on the top-right, and $H_{sp}\,y = H\,y'$ on the bottom-right.

We will assume throughout that in the key pair $\{p+1,q\}$ the point q is above $p+1$. The reader may check that in the other case the situation is the same or simpler.
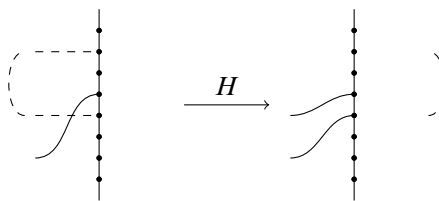
Figure 10: Homotopy, Case 1.

4.1.1 Case 1 In this case we consider generators y with key segment at left and key pair at left as a double horizontal. The H–arrow $y \to x$ is shown in Figure 10. There

are no special generators. The d–arrows starting at y involving the one strand around $p + 1$ include:

(Case 1.1) Merging with a strand above.

(Case 1.2) Merging with a strand below.

(Case 1.3) Moving a part not containing $p + 1$ to the right.

(Case 1.4) Moving a part containing $p + 1$ to the right.

The d–arrows starting at x involving the two strands shown include:

(Case 1.1) Merging upper strand with a strand above.

(Case 1.2) Merging lower strand with a strand below.

(Case 1.3) Moving a part of the lower strand.

(Case 1.4) Moving a part of the upper strand.

Crossing the two strands gets back to y.
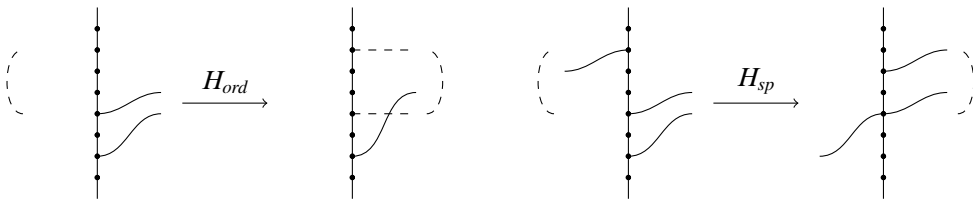
Figure 11: Homotopy, Case 2.

4.1.2 Case 2 In this case we consider generators y with key segment at right and key pair at right as the start of a moving strand. The H_{ord}–arrow $y \to x$ and the possible H_{sp}–arrow $y \to x'$ are shown in Figure 11. First suppose the generator y is non-special (that is, no strand of y ends at q). The d–arrows starting at y involving the two strands shown include:

(Case 2.1) Splitting the lower strand.

(Case 2.2) Splitting the upper strand.

(Case 2.3) Moving a strand from left to attach to lower strand.

The d–arrows starting at x involving the single strand shown include:

(Case 2.1) Splitting the single strand below $p + 1$.

(Case 2.2) Splitting the single strand above $p + 1$.

(Case 2.3) Moving a strand from left to attach to the single strand.

Uncrossing the single strand gets back to y.

Now suppose y is special. There is one more d–arrow starting at x. This is because the homotopy produces a double-horizontal at $\{p + 1, q\}$, so a strand ending at q on the left can now move to the right. This term cancels against a special d–arrow starting at x', as shown in Case 2.4.

For the case when y is special, there is one more relevant strand: the strand ending at q. The d–arrows starting at y involving this strand include:

(Case 2.5) Merging with a strand below.

(Case 2.6) Moving part of the strand to the right.

The d–arrows starting at x involving this strand are exactly the same. The cancellations are shown in Cases 2.5a and 2.6a. The d–arrows starting at x' involving this strand include:

(Case 2.5) Attaching with a strand moved from the left.

(Case 2.6) Splitting the strand.

The cancellations are shown in Cases 2.5b and 2.6b. The d–arrows starting at x' involving the two strands around $p + 1$ include:

(Case 2.7b) Moving part of the lower strand to the right.

(Case 2.8b) Splitting the upper strand.

(Case 2.9b) Merging the lower strand with a strand below.

These cancel against d–arrows starting at y as in Cases 2.1, 2.2, and 2.3, respectively.

4.1.3 Case 3

In this case we consider generators y with key segment at right and key pair at left as a double horizontal. The H–arrow $y \to x$ is shown in Figure 12. Note if the single strand on the H side ends exactly at $p + 1$, then a double horizontal is produced on the right side (however, since no strand can end at q on the left, there are no special cases). The d–arrows starting at y involving the single strand include:

(Case 3.1) Splitting at a point above $p + 1$.
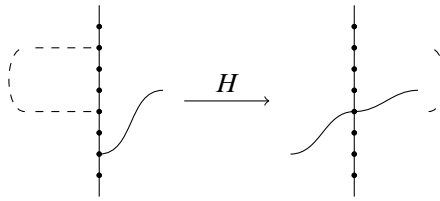
Figure 12: Homotopy, Case 3.

(Case 3.2) Splitting at a point below $p + 1$.

(Case 3.3) Attaching with a strand moved from the left.

The d–arrows starting at x include:

(Case 3.1) Splitting the right strand.

(Case 3.2) Moving part of the left strand to the right.

(Case 3.3) Merging the left strand with a strand below.

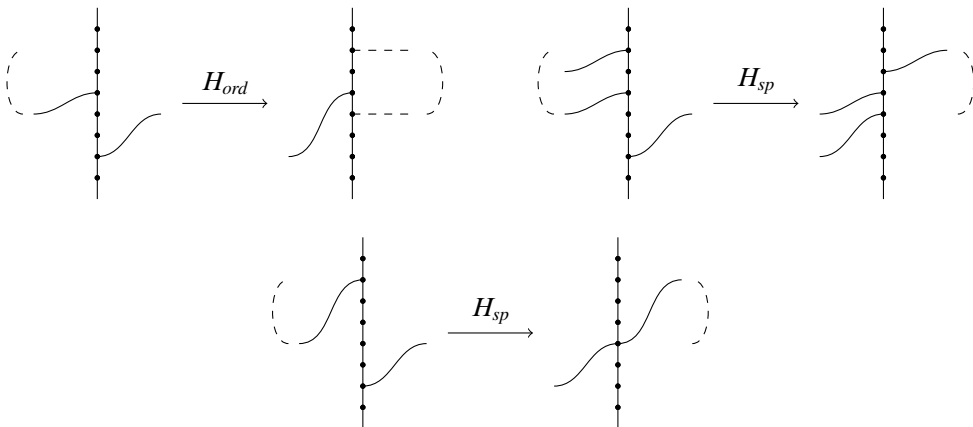Moving all of the left strand to the right gets back to y.

Figure 13: Homotopy, Case 4.

4.1.4 Case 4 In this case we consider generators y with key segment at right and key pair at left as the start of a moving strand. The H_{ord}–arrow $y \to x$ and the two possible H_{sp}–arrows $y \to x'$ are shown in Figure 13. Note that the conditions for the two possible H_{sp}–arrows are mutually exclusive, so there is at most one H_{sp}–arrow in any given case.

We first consider the case where y is non-special. The d–arrows starting at y involving the two strands shown include:

(Case 4.1)   Splitting the right strand.

(Case 4.2)   Attaching the right strand with a strand moved from the left.

(Case 4.3)   Moving part of the left strand to the right.

(Case 4.4)   Merging the left strand with a strand above.

The d–arrows starting at x involving the single strand shown include:

(Case 4.1)   Moving to the right a part not containing $p+1$.

(Case 4.3)   Moving to the right a part strictly containing $p+1$.

(Case 4.2)   Merging with a strand below.

(Case 4.4)   Merging with a strand above.

Moving a part containing $p+1$ on the boundary gets back to y.

Now suppose y is special. Since H creates a double horizontal at $\{p+1, q\}$, there is one more d–arrow starting at x, moving the strand ending at q to the right. This strand may or may not extend to $p+1$, giving the two special cases. This cancels against a special d–arrow starting at x' in both cases, as shown in Cases 4.5 and 4.6.

For the other cases, there is again one more relevant strand, ending at q. The d–arrows starting at y involving this strand include:

(Case 4.7)   Merging with a strand below, resulting in a strand starting above $p+1$.

(Case 4.8)   Moving part of the strand to the right.

(Case 4.9)   Merging with a strand below, resulting in the strand $p+1 \to q$.

The d–arrows starting at x are exactly the same for Cases 4.7 and 4.8. For Case 4.9 the strand ending at q is merged with a longer strand (shown in Case 4.9a). The d–arrows starting at x' involving this strand include:

(Case 4.7b)   Attaching a strand moved from the left, resulting in a strand starting above $p+1$.

(Case 4.8b)   Splitting the strand.

(Case 4.9b)   Attaching a strand moved from the left, resulting in the strand $p+1 \to q$.

Finally, we consider d–arrows starting at x' involving the two strands around $p+1$. They include:

(Case 4.10b,c)   Moving part of the lower strand to the right.

(Case 4.11b,c)   Joining the lower strand with a strand below.

(First special case, Case 4.12b)   Moving part of the upper strand to the right.

(Second special case, Case 4.12c)   Splitting the upper strand.

(First special case, Case 4.13b)   Joining the upper strand with a strand above.

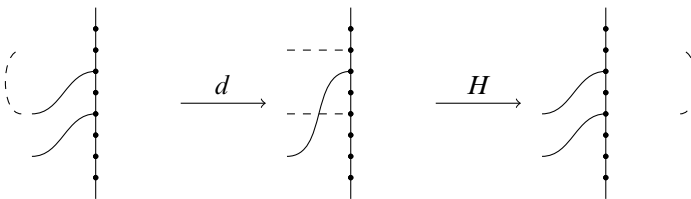These cancel against d–arrows starting at y as in Cases 4.1 through 4.4.

Figure 14: Identity term in Case 5.

4.1.5   Case 5   In this and the next three cases we consider generators on the d side. In this case, the generator x has key segment at left and key pair at left as the start of a moving strand. There are no special generators in this case (due to a strand ending at $p+1$ on the left). The only d–arrow $x \to y$ to the H side and the H–arrow $y \to x$, giving the identity term, are shown in Figure 14.

The only other d–arrow starting at x that changes the type of the generator is shown in Figure 15, but we are still on the d side.
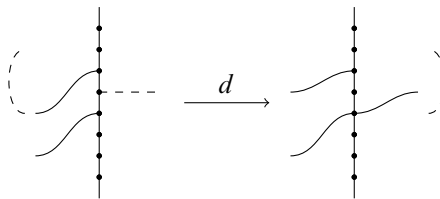
Figure 15: A d–arrow that moves to a different type, but still on the d side.

4.1.6   Case 6   In this case, the generator x has key segment at right and key pair at right as a double horizontal. If x is not special, Figure 16 shows the only d–arrow $x \to y$ to the H side. If x is special — that is, if there is a strand ending at q on the left — then the d–arrow moving that strand to the right will move the double horizontal to the left, making a generator on the H side. This corresponds to the special case in Case 2 (Case 6.1).
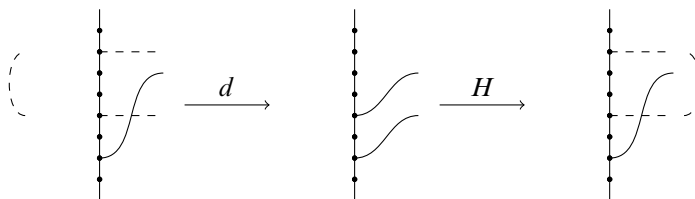
Figure 16: Identity term in Case 6.
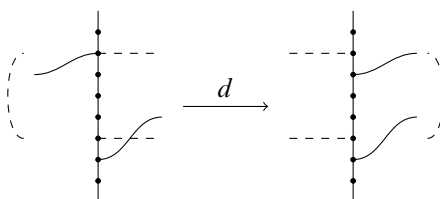
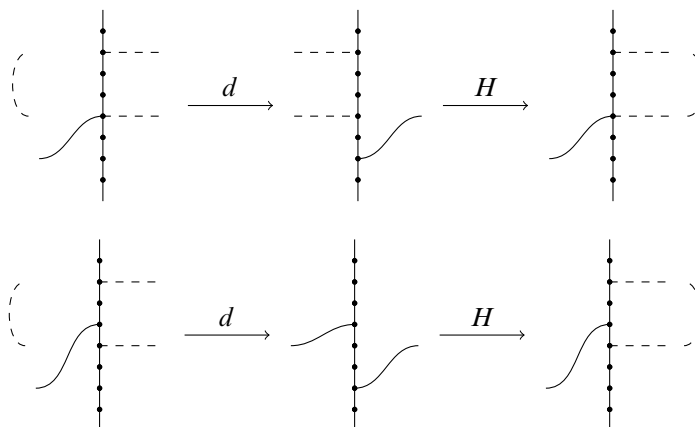Figure 17: A d–arrow that maps to the H side in Case 6, forcing a special case.

Figure 18: Identity term in Case 7.

4.1.7 Case 7 In this case, the generator x has key segment at left and key pair at right as a double horizontal. Figure 18 shows the obvious d–arrow $x \to y$ to the H side. If there is a strand ending at q (only possible in the second of the two cases in Figure 18), this strand can be moved to the right, moving the double horizontal to the left and changing the type of the generator. There are two ways this can happen, as shown in Figure 19. They correspond to the special cases in Case 4 (Cases 7.1 and 7.2).
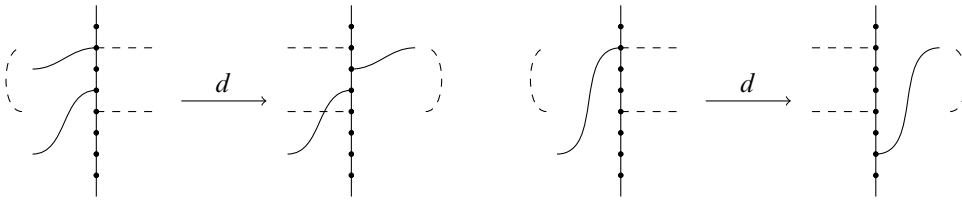
Figure 19: d–arrows that map to the H side in Case 7, forcing two special cases.

4.1.8 Case 8 In this case, the generator x has key segment at left and key pair at right as the start of a moving strand. The only d–arrow $x \to y$ to the H side is shown in Figure 20. There are no special generators.

This concludes the check for generators of multiplicity one.
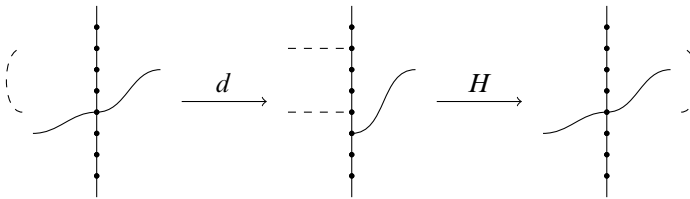
Figure 20: Identity term in Case 8.

4.2 Other generators

We now consider generators that do not have multiplicity one. They can be classified into four types as in Table 2. We will again check each type separately. The strategy is even simpler than before, as there are no special cases. For a generator on the H side, we check that there is a bijection between the corresponding sets G_1 and G_2. For a generator on the d side, we check that there is exactly one differential to the H side, carrying it to the generator in the same pair.

The checks for Cases 1 and 2 in the definition of H have a common ingredient: checking those terms where both the d–arrow and the H–arrow are restricted to the same side. For Case 2, the argument is the same as the one used to show $H_*(\mathcal{A}(\mathcal{Z})) = 0$ on generators with multiplicity greater than one. This is given in [4, Section 4]. While the proof there uses a slightly different homotopy, the argument easily adapts to the homotopy we give here. For Case 1 we use the dual of that argument. With these cases covered, we will only consider terms in $d \circ H + H \circ d$ that involve both sides.
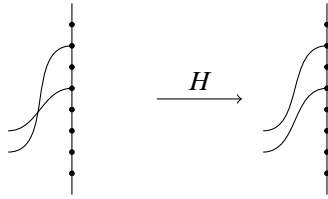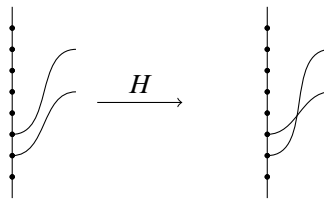
Figure 21: Homotopy in Case 1.

4.2.1 Case 1 In this case both the i strand and the j strand are on the left. The H–arrow $y \to x$ is shown in Figure 21.

We begin by verifying the equation for y. The d–arrows starting at y involve moving a part of the j strand to the right. The end of that part can be:

   (Case 9.1)   Lower than i.

   (Case 9.2)   Higher than the endpoint of the i strand.

   (Case 9.2')   Coinciding with the endpoint of the i strand.

   (Case 9.3)   In between i and the endpoint of the i strand.

Note that if part of the shorter strand is moved, we remain on the d side (Figure 22). As stated in the beginning of this section, we may ignore d–arrows within the left side.
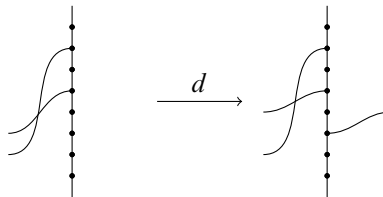
Figure 22: A term in the differential that remains on the d side.

The d–arrows starting at $Hy = x$ include:

- (Case 9.1)   Moving part of the j strand not containing i.

- Moving part of the i strand. The end of that part can be:

   (Case 9.2)   Above the endpoint of the j strand.
   (Case 9.2')   Coinciding with the endpoint of the j strand.
   (Case 9.3)   Below the endpoint of the j strand.

Note that it is impossible to move part of lower strand containing j, due to the multiplication rule that double crossing gives zero.

Now we verify the equation for x on the d side. Crossing the two strands shown is the only way to go to a generator of type 1 on the H side (note that we already covered the case where the d–arrow happens within the left side). It is impossible to reach type 2 (both strands on the right). To reach type 3, we need to arrange that the new j strand is at right. This means we will need to move a part of the current j strand containing i to the right. However, this will produce a double crossing.

Figure 23: Homotopy in Case 2.

4.2.2 Case 2 In this case, both the i strand and the j strand are on the right. The H–arrow $y \to x$ is shown in Figure 23.

For generators on the H side, the d–arrows that involve only the right side are again covered by the computation of $H_*(A(\mathcal{Z}))$. There is only one other case that involves the displayed strands: moving a strand from the left to attach to the j strand (for both dy and $d(Hy)$). This is shown in Case 10.1. Note that the interval $[j, i]$ cannot be occupied on the left side.

For generators on the d side, there is no way to reach the H side other than by d–arrows involving only the right side.
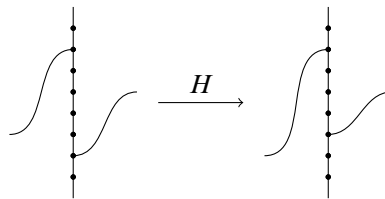

Figure 24: Homotopy in Case 3.

4.2.3 Case 3 In this case, the i strand is on the left and the j strand is on the right. The H–arrow $y \to x$ is shown in Figure 24. We consider the generator y in this case.

The d–arrows starting at y include:

(Case 11.1) Merging the i strand with a strand above.

(Case 11.2) Splitting the j strand below i.

(Case 11.3) Splitting the j strand above i.

(Case 11.4) Attaching the j strand with a strand moved from left.

(Case 11.5) Moving part of the i strand to the right, so that the new strand on the right does not cross the j strand.

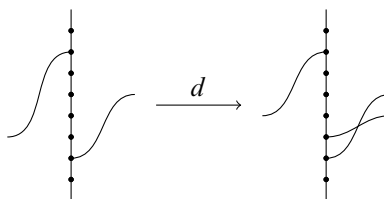If the new strand on the right does cross the i strand, we remain on the d side (Figure 25).

Figure 25: A term in the differential that remains on the d side.

The d–arrows starting at $Hy = x$ include:

(Case 11.1) Merging the j strand with a strand above.

(Case 11.4) Merging the j strand with a strand below.

(Case 11.3) Splitting the i strand.

(Case 11.2) Moving part of j strand not containing i to the right.

(Case 11.5) Moving part of j strand to the right, so that the new strand on the right crosses the i strand.

Note that if the new strand on the right interleaves, and does not cross the i strand, the term is zero by double crossing (Figure 26).

4.2.4 Case 4 In this case we consider the generator x in the H–arrow in the previous case, with the i strand on the right and the j strand on the left. Besides the obvious d–arrow to y, the only other possible way to reach the H side is by moving part of the j strand containing i to the right. However, this is impossible due to double crossing (same example as in Figure 26).

This concludes the check for generators with multiplicity greater than one, and the verification that H is indeed a homotopy.
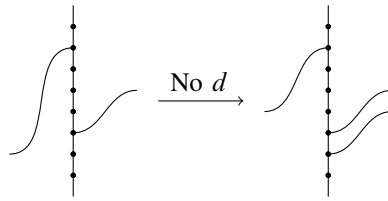
Figure 26: An example of a term that is not in the differential of $H\mathbf{x}$ due to double crossing.

5   A rank-1 model of $\widehat{CFAA}\,(\mathbb{I})_{\mathcal{Z}}$

Having found a homotopy H on the chain complex underlying the A_∞–bimodule $\mathcal{M}$, we can use homological perturbation theory to describe the smaller bimodule $\mathcal{N}$. We refer the reader to [3, Section 8.2] on how to construct an A_∞–bimodule action using homological perturbation theory. The result is summarized in the following theorem.

Theorem 5.1   *The A_∞–bimodule $\mathcal{N}_{\mathcal{A}',\mathcal{A}}$, described below, is homotopy equivalent to the bimodule $\mathcal{M}$ given in (6), and is therefore a model of the invariant $\widehat{CFAA}\,(\mathbb{I}_{\mathcal{Z}})_{\mathcal{A}',\mathcal{A}}$.*

The vector space N underlying $\mathcal{N}$ is generated over $\mathbb{F}_2$ by the indecomposable idempotents of $\mathcal{A}$ (using the definitions in Section 2, $\mathcal{N}$ is a rank-1 bimodule). The arrows in the A_∞–bimodule action are in one-to-one correspondence with sequences $[a_{1,1}, a_{1,2}], \ldots, [a_{2n,1}, a_{2n,2}]$ of generators of $\mathcal{A}$ that satisfy these three conditions:
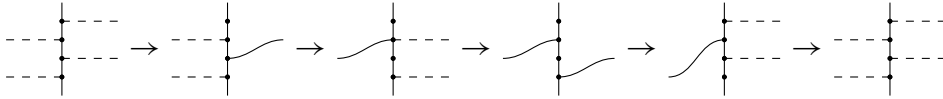
  (i) *$[a_{1,1}, a_{1,2}] = [i', i]$ and $[a_{2n,1}, a_{2n,2}] = [j', j]$ for some idempotents $i, j \in A$, and $i' = o(i)$, $j' = o(j)$.*

 (ii) *Each $[a_{2k,1}, a_{2k,2}]$ is obtained from $[a_{2k-1,1}, a_{2k-1,2}]$ by either factoring out some $\overline{b'} \in \mathcal{A}$ from $a_{2k-1,1}$ on the right (so that $a_{2k-1,1} = a_{2k,1}\overline{b'}$ and $a_{2k-1,2} = a_{2k,2}$), or by multiplying $a_{2k-1,2}$ with some $b \in \mathcal{A}$ on the right (so that $a_{2k,2} = a_{2k-1,2}b$ and $a_{2k,1} = a_{2k-1,1}$). The elements b' and b are not necessarily equal between steps.*

(iii) *Each $[a_{2k+1,1}, a_{2k+1,2}]$ is obtained from $[a_{2k,1}, a_{2k,2}]$ by one of the H arrows described in Section 3.*

Let $b'_1, \ldots, b'_p$ be the sequence of $b' \in \mathcal{A}'$ and let $b_1, \ldots, b_q$ be the sequence of $b \in \mathcal{A}$ used in condition (ii). *The arrow corresponding to this sequence is*

$$m_{1,p,q}([i]; b'_1, \ldots, b'_p; b_1, \ldots, b_q) \to [j].$$

Here $m_{1,p,q}$ is the part of the A_∞–bimodule action on $\mathcal{N}$ with p inputs on the $\mathcal{A}'$ side and q inputs on the $\mathcal{A}$ side. The idempotent of a generator $[i]$ of this A_∞–bimodule is i on the $\mathcal{A}$ side and $\overline{o(i)}$ on the $\mathcal{A}'$ side.

For example, let $\mathcal{Z}$ be the genus 1 pointed matched circle. The sequence of $[a_{i,1}, a_{i,2}]$ given by

gives rise to the following arrow in the action:

where $[i]$ for an idempotent $i \in \mathcal{A}$ denotes the corresponding generator of $\mathcal{N}$.

Since $\mathcal{A}'$ and $\mathcal{A}$ are opposite algebras, we can also consider $\mathcal{N}$ as a left-right A_∞–bimodule $_{\mathcal{A}}\mathcal{N}_{\mathcal{A}}$. We will use this form of $\mathcal{N}$ in the next section.

The bimodule $\mathcal{N}$ inherits a relative grading from the larger model $\mathcal{M}$. Since the generators in $\mathcal{N}$ are those in $\mathcal{M}$ composed of idempotents, we may choose the relative grading so that the grading of all generators in $\mathcal{N}$ are zero (this is also the grading one would obtain from the standard Heegaard diagram for the identity diffeomorphism, starting with a zero grading for any one of the generators).

We now prove the main results of this chapter.

Proof of Theorem 1.1 Let

$$(11) \qquad {}^{\mathcal{A}}\mathcal{L}_{\mathcal{A}} = \widehat{CFAA}(\mathbb{I}_{\mathcal{Z}})_{\mathcal{A}',\mathcal{A}} \boxtimes {}^{\mathcal{A},\mathcal{A}'}\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}}) \simeq \mathcal{N}_{\mathcal{A}',\mathcal{A}} \boxtimes {}^{\mathcal{A},\mathcal{A}'}\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}}).$$

Since the vector spaces underlying both $\mathcal{N}$ and $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$ are generated by idempotents, the vector space underlying $\mathcal{L}$ is also generated by idempotents. So $\mathcal{L}$ is a rank-1 bimodule. Furthermore, the type DA–bimodule action in $\mathcal{L}$ satisfies $\delta_1^1 = 0$ since all arrows in $\mathcal{N}$ involve non-idempotent algebra inputs from both $\mathcal{A}'$ and $\mathcal{A}$. Hence [4, Lemma 2.2.50] applies, showing ${}^{\mathcal{A}}\mathcal{L}_{\mathcal{A}} = {}^{\mathcal{A}}[\phi]_{\mathcal{A}}$ for some A_∞–algebra morphism $\phi \colon \mathcal{A} \to \mathcal{A}$.

From the grading on $\mathcal{N}$ and $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$, there is a relative grading on $\mathcal{L}$ with all generators having grading zero. This implies that the A_∞ morphism ϕ respects the $G(\mathcal{Z})$–grading on $\mathcal{A}$. By classifying arrows in $\mathcal{N}$ involving only one length 1 chord on each side, it is clear that $\phi_1(\xi) = \xi$ for any length 1 chord ξ. By [4, Proposition 4.11], we conclude that ϕ_1 induces the identity map on homology. So ϕ is a quasi-isomorphism and $\mathcal{L}$ is quasi-invertible. $\qquad\square$

Corollary 5.2 $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$ *is quasi-invertible, hence the functor $\cdot \boxtimes \widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$ induces an equivalence of categories between the category of right A_∞–modules over $\mathcal{A}(\mathcal{Z})$, and the category of left type D structures over $\mathcal{A}(-\mathcal{Z})$.*

Proof With $\mathcal{L}$ as in (11), there is an $\mathcal{L}'$ such that

$$^{\mathcal{A}}\mathbb{I}_{\mathcal{A}} \simeq \mathcal{L}' \boxtimes \mathcal{L} \simeq \mathcal{L}' \boxtimes \mathcal{N} \boxtimes \widehat{CFDD}(\mathbb{I}_{\mathcal{Z}}),$$

which means $\mathcal{L}' \boxtimes \mathcal{N}$ is the quasi-inverse to $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$. $\qquad\square$

Remark The stronger statement, that $\mathcal{N}$ is the quasi-inverse of $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$, will be proved combinatorially in [5].

6 Examples of Koszul duality

In this section, we use our description of $\mathcal{N} \simeq \widehat{CFAA}(\mathbb{I}_{\mathcal{Z}})$ to give an explicit A_∞ morphism from $\mathcal{A} = \mathcal{A}(\mathcal{Z})$ to $\mathrm{Cob}(\mathcal{A})$, inducing an isomorphism on homology. We will assume in this section that $\mathcal{N}$ is in fact the quasi-inverse of $\widehat{CFDD}(\mathbb{I}_{\mathcal{Z}})$. First, we review some material from [4] and [2, Section 8].

Definition 1 An *augmentation* of a dg-algebra A is a map $\epsilon \colon A \to \boldsymbol{k}$ satisfying $\epsilon(1) = 1$ and $\epsilon(a_1 a_2) = \epsilon(a_1)\epsilon(a_2)$. Given an augmentation, we let $A_+ = \ker(\epsilon)$ be the *augmentation ideal* of A.

The strand algebra $\mathcal{A}$ is augmented with the augmentation map ϵ sending each idempotent to itself and other generators to zero. So the augmentation ideal $\mathcal{A}_+$ is generated by the non-idempotents.

Definition 2 Given an augmented dg-algebra A, the cobar resolution $\mathrm{Cob}(A)$ is defined as $T^*(A_+[1]^*)$, the tensor algebra of the dual of the augmentation ideal. This can be given the structure of a dg-algebra, with product being the one on the tensor algebra, and the differential consisting of the following arrows:

$$a_1^* \otimes \cdots \otimes b^* \otimes \cdots \otimes a_k^* \to a_1^* \otimes \cdots \otimes a_i^* \otimes \cdots \otimes a_k^*$$

for each i and term b in da_i, and

$$a_1^* \otimes \cdots \otimes a_i^* \otimes \cdots \otimes a_k^* \to a_1^* \otimes \cdots \otimes b^* \otimes b'^* \otimes \cdots \otimes a_k^*$$

for each i and generators b, b' such that $bb' = a_i$.

For any augmented dg-algebra A, there is a type DD bimodule $^{\mathrm{Cob}(A)}K^A$, of rank 1 over $\boldsymbol{k}$, and with type DD action given by

$$\delta^1(\mathbf{1}) = \sum_i a_i^* \otimes \mathbf{1} \otimes a_i,$$

where the sum is over a set of generators of A_+ (the sides of the action are reversed in comparison to [2], for ease of computation later). One may check that this satisfies the structure equation for type *DD* bimodules. The fact that this bimodule is always quasi-invertible shows that A is Koszul dual to $\mathrm{Cob}(A)$ (see [2, Proposition 8.12]).

The following is contained in the proof of [2, Proposition 8.11].

Proposition 2 *Let* $_{\mathcal{A}}\mathcal{N}_{\mathcal{A}}$ *be a rank-1 representative of* $_{\mathcal{A}}\widehat{CFAA}(\mathbb{I}_{\mathcal{Z}})_{\mathcal{A}}$. *Then*

$$^{\mathrm{Cob}(\mathcal{A})}K^{\mathcal{A}} \boxtimes {}_{\mathcal{A}}\mathcal{N}_{\mathcal{A}}$$

is of the form $^{\mathrm{Cob}(\mathcal{A})}[\phi]_{\mathcal{A}}$, *where* ϕ *is an* A_∞ *morphism* $\mathcal{A} \to \mathrm{Cob}(\mathcal{A})$ *that induces an isomorphism on homology.*

In particular ϕ_1 (the part of ϕ taking one input) maps representatives of homology classes of $\mathcal{A}$ to representatives of homology classes of $\mathrm{Cob}(\mathcal{A})$. Unwinding all the definitions, we obtain the following rule for computing this map.

Proposition 3 *Let* $a \in \mathcal{A}$ *be a generator of* $\mathcal{A}$. *Then terms in* $\phi_1(a) \in \mathrm{Cob}(A)$ *correspond to arrows in the* A_∞ *action of* $_{\mathcal{A}}\mathcal{N}_{\mathcal{A}}$ *of the form*

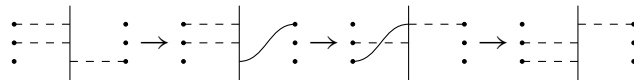$$m_{p,1,1}(a_1, \ldots, a_p, [i], a) \to [j],$$

with the above arrow giving rise to the term

$$a_p^* \otimes \cdots \otimes a_1^*$$

in $\phi_1(a)$.

We now give some examples. In each case, there is exactly one algebra input $a \in \mathcal{A}$ on the right. So in the notation of Theorem 5.1, the sequence must start with $a_{2,2} = a$, then alternate between applying an H–arrow and factoring from $a_{2k-1,1}$.

For a length 1 strand $a \in \mathcal{A}$, we get $\phi_1(a) = a^*$. For the first non-trivial case, we consider a length 2 strand, when none of the three points are paired. There are two possible orderings of the two intervals in $<_{\mathcal{Z}}$. If the upper interval comes first, the sequence
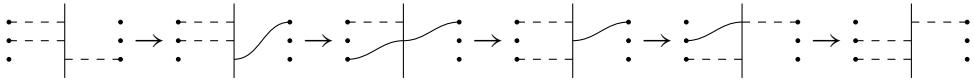

gives

(12)
$$\phi_1\left(\begin{smallmatrix}:\diagup:\end{smallmatrix}\right) = \begin{smallmatrix}:\diagup:\end{smallmatrix},$$

where we used the same diagram to represent a and a^*.

If the lower interval comes first, the sequence

gives:

(13)
$$\phi_1\left(\vphantom{\Big|} \right) = \quad \otimes \quad$$

The right side of Equations (12) and (13) represent the same homology class in $\mathrm{Cob}(\mathcal{A})$, since their difference is given by

$$d\left(\vphantom{\Big|} \right).$$

Now we consider a more complicated example, which shows the possible complications that can arise in such a computation.
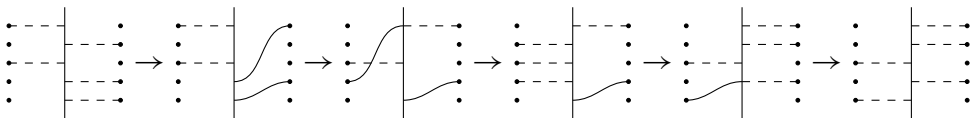
Consider a length 4 interval in a larger pointed matched circle, where the second and fourth of the five points are paired. There are two possibilities for the relative orderings of the four intervals in $<_{\mathcal{Z}}$. This is because the second interval (counting from below) must immediately precede the third interval, and the fourth interval must immediately precede the first interval. The two possible relative orderings are given as follows:
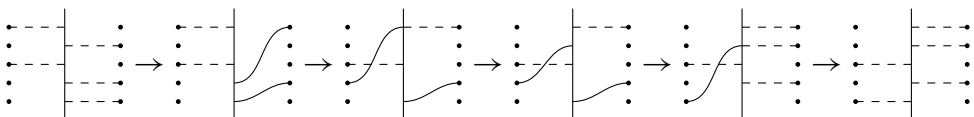
$$
\begin{array}{cc}
\begin{array}{c} 1 \\ 4 \\ 3 \\ 2 \end{array}
&
\begin{array}{c} 3 \\ 2 \\ 1 \\ 4 \end{array}
\end{array}
$$
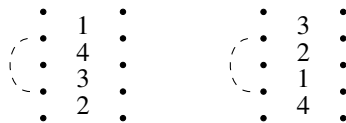
We will simply refer to them as 1-4-3-2 and 3-2-1-4 from now on. The two generators
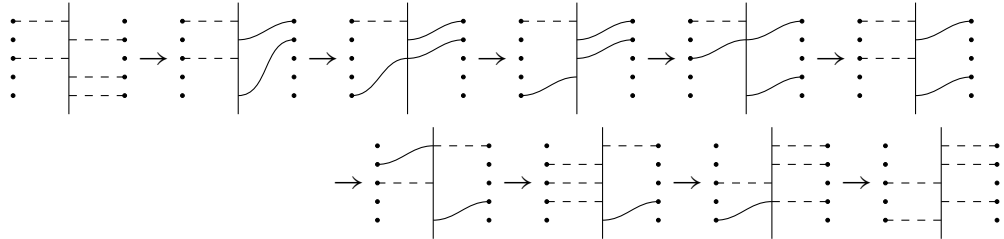
(14)
$$a_1 = \quad, \quad a_2 = \quad$$

represent the same homology class in $\mathcal{A}$. Consider first the ordering 1-4-3-2, using the representative a_1. There are three terms in $\phi_1(a_1)$. The first term comes from the following sequence:

The second term comes from:

The third term comes from:

Note the use of a special H–arrow in the fourth step of the last sequence. This produces:

(15)

If the ordering is 3-2-1-4, a straightforward sequence gives:

(16)

The difference between the right sides of Equations (15) and (16) is:

Now we consider representative a_2 of the same homology class of $\mathcal{A}$. If the ordering is 1-4-3-2, the sequence is straightforward, giving:

(17)

The difference between the right sides of Equations (17) and (15) is

Finally, if the ordering is 3-2-1-4, one can check there is a single sequence, which uses a special H–arrow in the fourth step:

This produces:

$$(18) \qquad \phi_1\left(\begin{matrix} \cdot \\ \vdots \end{matrix} \right) = \cdot \cdot \cdot \otimes \cdot \cdot \cdot \otimes \cdot \cdot \cdot \otimes \cdot \cdot \cdot$$

The right side is the same as that in Equation (16). This finishes the consistency check that the same homology class in $\mathrm{Cob}(\mathcal{A})$ is obtained using any valid local ordering of the intervals, and any representative of the same homology class in $\mathcal{A}$. In general, one can expect the computation to be more complicated as the length of elements in the homology class increases.

References

[1] **R Lipshitz**, **P Ozsváth**, **D Thurston**, *Bordered Heegaard Floer homology: invariance and pairing* `arXiv:0810.0687`

[2] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Heegaard Floer homology as morphism spaces*, Quantum Topol. 2 (2011) 381–449 MR2844535

[3] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Computing $\widehat{HF}$ by factoring mapping classes*, Geom. Topol. 18 (2014) 2547–2681 MR3285222

[4] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Bimodules in bordered Heegaard Floer homology*, Geom. Topol. 19 (2015) 525–724

[5] **B Zhan**, *Combinatorial proofs in bordered Heegaard Floer homology* `arXiv:1405.3004`

[6] **B Zhan**, *Explicit Koszul-dualizing bimodules in bordered Heegaard Floer homology: Cancellation diagrams*, Alg. Geom. Topol 16 (2016) online supplement

Department of Mathematics, Massachusetts Institute of Technology
Building E18, Room 306, 77 Massachusetts Avenue, Cambridge, MA 02139-4307, USA
`bzhan@mit.edu`