A-infinity algebras, strand algebras, and contact categories

Daniel V. Mathews

March 17, 2018

Abstract

In previous work we showed that the contact category algebra of a quadrangulated surface is isomorphic to the homology of a strand algebra from bordered Floer theory. Being isomorphic to the homology of a differential graded algebra, this contact category algebra has an A-infinity structure. In this paper we investigate such A-infinity structures in detail. We give explicit constructions of such A-infinity structures, and establish some of their properties, including conditions for the nonvanishing of A-infinity operations. Along the way we develop several related notions, including a detailed consideration of tensor products of strand diagrams.

Contents

1 Introduction

1.1 Overview

In previous work [21] we demonstrated an isomorphism of two unital $\mathbb{Z}_2$-algebras: the first arising from contact geometry; the second from bordered Floer theory.

$$CA(\Sigma, Q) \cong H(\mathcal{A}(\mathcal{Z})) \tag{1}$$

Here (Σ, Q) is a *quadrangulated surface*, a useful object in TQFT-type structures in contact geometry [19, 20], and $\mathcal{Z}$ is an *arc diagram*, an equivalent object used in bordered sutured Floer theory [28]. The left hand side $CA(\Sigma, Q)$ is the algebra of a *contact category*, with objects and morphisms given by certain contact structures on $\Sigma \times [0,1]$. The right hand side $H(\mathcal{A}(\mathcal{Z}))$ is the homology of the strand algebra $\mathcal{A}(\mathcal{Z})$, a differential graded algebra (DGA) generated by *strand diagrams* on $\mathcal{Z}$, which encode Reeb chords arising as asymptotics of certain holomorphic curves. The isomorphism (1) therefore allows us to interpret (homology classes of) strand diagrams as contact structures.

Of particular interest, (1) expresses the contact category algebra as the homology of a DGA. The homology of a DGA is known to have the structure of an A_∞ algebra. While A_∞ structures are well known to arise in Floer theory (see e.g. [25]), it is perhaps surprising that an A_∞ structure should arise directly out of contact structures.

This paper essentially consists of an investigation of A_∞ structures on this contact category algebra. This investigation is carried out through the use of strand diagrams, which are more general objects, and easier to work with algebraically than contact structures. Therefore, more accurately, this paper consists of an investigation of A_∞ structures on $H(\mathcal{A}(\mathcal{Z}))$, from a contact-geometric perspective.

Throughout this paper we work with $\mathbb{Z}_2$ coefficients; signs are always irrelevant.

1.2 Main results

Our first main result is the explicit construction of A_∞ structures on $H(\mathcal{A}(\mathcal{Z}))$.

Theorem 1.1. *A pair ordering of $\mathcal{Z}$ can be used to define an explicit A_∞ structure X on $H(\mathcal{A}(\mathcal{Z}))$, together with a morphism of A_∞ algebras $f\colon H(\mathcal{A}(\mathcal{Z})) \longrightarrow \mathcal{A}(\mathcal{Z})$. These consist of maps*

$$X_n \colon H(\mathcal{A}(\mathcal{Z}))^{\otimes n} \longrightarrow H(\mathcal{A}(\mathcal{Z})), \quad f_n \colon H(\mathcal{A}(\mathcal{Z}))^{\otimes n} \longrightarrow \mathcal{A}(\mathcal{Z}),$$

where X extends the DGA structure of $H(\mathcal{A}(\mathcal{Z}))$, and $\mathcal{A}(\mathcal{Z})$ is regarded as an A_∞ algebra with trivial n-ary operations for $n \geq 3$.

By (1), theorem 1.1 provides A_∞ structures on the contact category algebra $CA(\Sigma, Q)$.

We will discuss *pair orderings* as we proceed (section 3.6); they consist of a total order on the matched pairs of $\mathcal{Z}$, along with an ordering of the two points in each pair. In fact the full statement (theorem 5.2) allows for a slightly more general A_∞ structures, using certain types of "choice functions" to parametrise the various choices involved in the construction.

The second main result provides necessary conditions under which these A_∞ maps are nontrivial, and under those conditions gives an explicit description of the results. The idea is that certain "local" conditions at the matched pairs of $\mathcal{Z}$ are necessary to obtain nonzero output from the A_∞ maps.

Theorem 1.2. *Let $M = M_1 \otimes \cdots \otimes M_n \in H(\mathcal{A}(\mathcal{Z}))^{\otimes n}$ be a tensor product of nonzero homology classes of strand diagrams. The maps f_n and X_n of theorem 1.1 have the following properties.*

(i) If $\overline{f}_n(M) \neq 0$, then M has l twisted and m critical matched pairs, where $l + m \geq n - 1$ and $m \leq n-2$, and all other matched pairs are tight. In this case $\overline{f}_n(M)$ is a sum of strand diagrams, where each diagram D is tight at all matched pairs where M is critical or tight, and has $n-1-m$ crossed and $l+m-n+1$ twisted matched pairs.

(ii) If $X_n(M) \neq 0$, then M has precisely $n-2$ critical matched pairs, and all other matched pairs tight. In this case, $X_n(M)$ is the unique homology class of tight diagram with the appropriate gradings.

All the terminology will be defined in due course.

The other main results involve the notion of *operation trees*. These will be defined in due course (section 8.1); they encode the way in which contact structures can be combined by the various A_∞ operations. Trees have commonly been used to encode A_∞ operations (e.g. [9, 10, 25]).

Certain trees of this type are required to obtain nonzero output from an A_∞ operation.

Proposition 1.3. *If $X_n(M) \neq 0$ or $\overline{f}_n(M) \neq 0$, there is a valid distributive operation tree for M.*

Our final main result gives sufficient conditions on diagrams and trees which ensure a nonzero result; this result is again described explicitly.

Theorem 1.4.

(i) *Suppose M has no on-on doubly occupied matched pairs. If every valid distributive operation tree for M is strictly f-distributive, and at least one such tree exists, then $\overline{f}_n(M) \neq 0$. Moreover, $\overline{f}_n(M)$ is given by a single diagram D, which is tight at all matched pairs where M is critical or tight, and crossed at all matched pairs where M is twisted.*

(ii) *Suppose M has no twisted or on-on doubly occupied matched pairs. If every valid distributive operation tree for M is strictly X-distributive, and at least one such tree exists, then $X_n(M) \neq 0$. Moreover, $X_n(M)$ is given by the homology class of unique tight diagram with appropriate gradings.*

As we will explain, these results are quite partial. The necessary conditions of theorem 1.2 are far from sufficient, and the sufficient conditions of theorem 1.4 are far from necessary. Since there are many A_∞ structures on $H(\mathcal{A}(\mathcal{Z}))$, we cannot expect a complete characterisation of diagrams which yield zero and nonzero results; still, we hope these results can be improved.

As may already be clear, there is a *lot* of terminology to define. Simply stating these results requires us to describe precisely many aspects of strand diagrams, and their tensor products and homology classes. We must name this world in order to understand it.

1.3 Construction of A-infinity structures

In a certain sense, the A_∞ structures on $CA(\Sigma, Q)$ or $H(\mathcal{A}(\mathcal{Z}))$ are already understood. In the 1980 paper [8], Kadeishvili showed how to define an A_∞ structure on the homology H of any differential graded algebra A (provided H is free, which is always true with $\mathbb{Z}_2$ coefficients). Indeed, in this paper we follow this construction, and theorem 1.1 can be regarded as fleshing out its details when $A = \mathcal{A}(\mathcal{Z})$. The only thing possibly new in theorem 1.1 is the level of explicitness in the construction.

We briefly recall some facts about A_∞ algebra; we refer to Keller's [9] for an introduction to A_∞ algebra, or to Seidel [25] for further details. An A_∞ *structure* m on a $\mathbb{Z}$-graded $\mathbb{Z}_2$-module A is a collection of operations $m_n \colon A^{\otimes n} \longrightarrow A$ for each $n \geq 1$, where each m_n has degree $n-2$. We call m_n the *n-ary* or *level* n operation. The operations m_i satisfy, for each $n \geq 1$,

$$\sum_{i+j+k=n} m_{i+1+k}\left(1^{\otimes i} \otimes m_j \otimes 1^{\otimes k}\right) = 0.$$

This identity for $n=1$ says that $m_1^2 = 0$, so m_1 is a differential; then the identity for $n=2$ is the Leibniz rule, with m_2 regarded as multiplication. Indeed an A_∞ algebra with all $m_n = 0$ for $n \geq 3$ is precisely a DGA. A *morphism* f of A_∞ algebras $A \longrightarrow A'$ (where the operations on A, A' are denoted m_i, m_i' respectively) is a collection of $\mathbb{Z}_2$-module homomorphisms $f_n \colon A^{\otimes n} \longrightarrow A'$, where each f_n has degree $n-1$. We call f_n the *level* n map. The maps f_i satisfy, for each $n \geq 1$,

$$\sum_{i+j+k=n} f_{i+1+k}\left(1^{\otimes i} \otimes m_j \otimes 1^{\otimes k}\right) = \sum_{i_1+\cdots+i_s=n} m_s'\left(f_{i_1} \otimes f_{i_2} \otimes \cdots \otimes f_{i_s}\right).$$

Kadeishvili's construction in [8] produces an A_∞ structure X on H, consisting of operations $X_n \colon H^{\otimes n} \longrightarrow H$, and a morphism f of A_∞ algebras $H \longrightarrow A$, consisting of maps $f_n \colon H^{\otimes n} \longrightarrow A$. The DGA A is regarded as an A_∞ algebra with trivial n-ary operations for $n \geq 3$. The A_∞ structure constructed on H begins with trivial differential $X_1 = 0$, and X_2 is the multiplication on H inherited from A. If H is free then there is a map $f_1 \colon H \longrightarrow A$ (possibly many) which is an isomorphism in

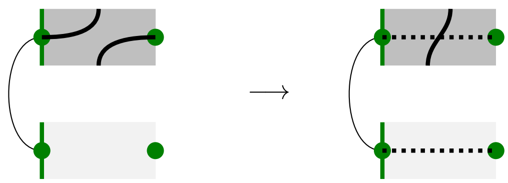

Figure 1: The action of a creation operator.

homology, sending each homology class to a cycle representative. The constructed f_n can be taken to begin with any such f_1.

The construction proceeds inductively, producing maps $U_n \colon H^{\otimes n} \longrightarrow A$ of degree $n-2$ along the way. First, $U_1 = 0$, $X_1 = 0$, and $f_1 \colon H \longrightarrow A$ is given. Once U_i, X_i, f_i are defined for $i < n$, we define U_n by

$$U_n(a_1 \otimes \cdots \otimes a_n) = \sum_{j=1}^{n-1} m_2\left(f_j(a_1 \otimes \cdots \otimes a_j) \otimes f_{n-j}(a_{j+1} \otimes \cdots \otimes a_n)\right) + \sum_{k=0}^{n-2} \sum_{j=2}^{n-1} f_{n-j+1}\left(a_1 \otimes \cdots \otimes a_k \otimes X_j(a_{k+1} \otimes \cdots \otimes a_{k+j}) \otimes \cdots \otimes a_n\right), \quad (2)$$

and X_n is then simply the homology class of U_n,

$$X_n = [U_n]. \quad (3)$$

Since f_1 selects cycles, $f_1 X_n$ and U_n differ by a boundary; f_n is then defined by

$$f_1 X_n - U_n = \partial f_n. \quad (4)$$

It is then shown that such f_n and X_n have the desired properties.

Clearly, in this construction there is a choice for f_n at each stage, but no choice for U_n or X_n. This choice amounts to a choice of inverse for the differential ∂.

Our construction, detailed in section 5, gives an explicit way to choose an f_n at each stage. This choice is made by maps which we call *creation operators*. We regard the differential in $\mathcal{A}(\mathcal{Z})$ as an "annihilation operator", destroying crossings between strands by resolving them. Creation operators, on the other hand, insert crossings in a controlled way. The idea is shown in figure 1. We introduce creation operators in section 3. Creation operators satisfy Heisenberg relations (proposition 3.18); this amounts to a chain homotopy from the identity to zero. In a certain sense, creation operators are the only operators obeying such Heisenberg relations (proposition 3.19); however they only form a very small subspace of the space of operators inverting the differential as required in Kadeishvili's construction (proposition 3.13). Similar creation operators have been put to use elsewhere in contexts related to contact geometry and Floer homology [22, 23].

However, there is still choice involved in where to apply creation operators, i.e. where to insert crossings. There is also a choice for the initial cycle selection homomorphism f_1. We parametrise such choices through notions of *creation choice functions* and *cycle choice functions* respectively. Our construction in general (theorem 5.2) produces an A_∞ structure on $H(\mathcal{A}(\mathcal{Z}))$ or $CA(\Sigma, Q)$ from a given cycle choice function and creation choice function. A pair ordering can be used to obtain such choice functions, leading to the formulation of theorem 1.1.

In order to define the A_∞ structure X on $H(\mathcal{A}(\mathcal{Z}))$, it turns out to be sufficient to work in a particular *quotient* of $\mathcal{A}(\mathcal{Z})$. This simplifies details considerably. We define a two-sided ideal $\mathcal{F}$ in section 2.13. The maps $\overline{f}_n$ appearing in theorems 1.2 and 1.4 are the images of f_n in the quotient by $\mathcal{F}$. Related ideas appeared in [13].

Algorithmically, the calculation of an A_∞ map $X_n(M_1 \otimes \cdots \otimes M_n)$, where $M_1, \ldots, M_n$ are homology classes of strand diagrams — or contact structures — by the method described above requires the computation of each $f_{j-i+1}(M_i \otimes \cdots \otimes M_j)$ and $X_{j-i+1}(M_i \otimes \cdots \otimes M_j)$, for $1 \leq i \leq j \leq n$. The algorithm therefore has complexity $O(n^2)$ (where we regard each computation of expressions such as (2) as constant time, and the complexity of the arc diagram $\mathcal{Z}$, as constant). The contrapositive of theorem 1.2 provides a set of conditions which imply $X_n(M) = 0$, which are easily checked in constant time. On the other hand, proposition 1.3 and theorem 1.4 provide conditions which are much more difficult to check, as the number of operation trees grows much faster with n. We regard these results as interesting not because of algorithmic usefulness, but because they perhaps provide some insight into A_∞ operations.

1.4 Classifications of diagrams, and the many types of twisted

As mentioned above, there are many features of strand diagrams which are relevant for our purposes, but which have not been given names in the existing literature. Large parts of this paper, especially sections 2 and 4, are devoted to defining and classifying these features, and establishing some of their properties. These are all required for our main theorems.

Therefore, some of the work here is an exercise in taxonomy. We briefly explain what we need to define and why, and the resulting classifications.

Contact structures naturally come in two types: tight and overtwisted. This dichotomy goes back to Eliashberg's work in the 1980s [2]. In the present work, consideration of the relationship between strand diagrams and contact structures naturally leads to further distinctions. Roughly speaking, when we look at strand diagrams from a contact-geometric perspective, there are many types of twisted.

According to the isomorphism (1) of [21], tight contact structures correspond to strand diagrams which are nonzero in homology. Such diagrams are characterised by certain conditions; roughly speaking, they must have an appropriate grading, no crossings, and must not have any matched pair that looks like the left of figure 1. A strand diagram which fails one or more of these conditions can be regarded as "overtwisted" in some sense.

The simplest way for a diagram to fail to represent a tight contact structure is by grading: it may lie in a summand of $\mathcal{A}(\mathcal{Z})$ which has no homology. This leads to the notion of *viability* (section 2.3). Only viable strand diagrams can possibly represent contact structures.

It is essential for our purposes to have precise terminology relating to these gradings and summands. We introduce a notion of *H-data*, which combines homological grading and idempotents (section 2.1). We also introduce notions of *on/off* or *1/0* to describe idempotents locally, and *occupation* of various parts of a strand diagram, such as *places* and *steps*, to describe homological grading locally (section 2.5). Some of this terminology was used in [21].

A viable diagram can still fail to represent a tight contact structures for multiple reasons. In the mildest case, shown in figure 2, a strand diagram is the product of strand diagrams corresponding to tight contact structures (and in fact stacking the two relevant cubes yields a tight contact structure), but the full contact structure is overtwisted. We define such "minimally overtwisted" diagrams as *twisted* in section 2.10.

Viable strand diagrams can also fail to represent tight contact structures because they have *crossings*. Thus, the natural tight/overtwisted classification of contact structures naturally becomes a 3-fold classification of viable strand diagrams into tight/twisted/crossed. This classification is, in a precise sense, (lemma 2.10), in ascending order of degeneracy.

When we proceed to *tensor products* of diagrams in section 4, there is again a natural notion of viability (section 4.1). Diagrams can represent contact structures, and their tensor products can be regarded as "stacked" contact structures on $\Sigma \times [0, 1]$. Viability then incorporates the natural contact-geometric condition that such stacked structures agree along their common boundaries.

Tensor products of diagrams again have a natural "tight/twisted" classification, (section 4.2), but now there are *six* types, which we call *tight*, *sublime*, *twisted*, *crossed*, *critical*, and *singular*, again in an ascending scale of degeneracy.

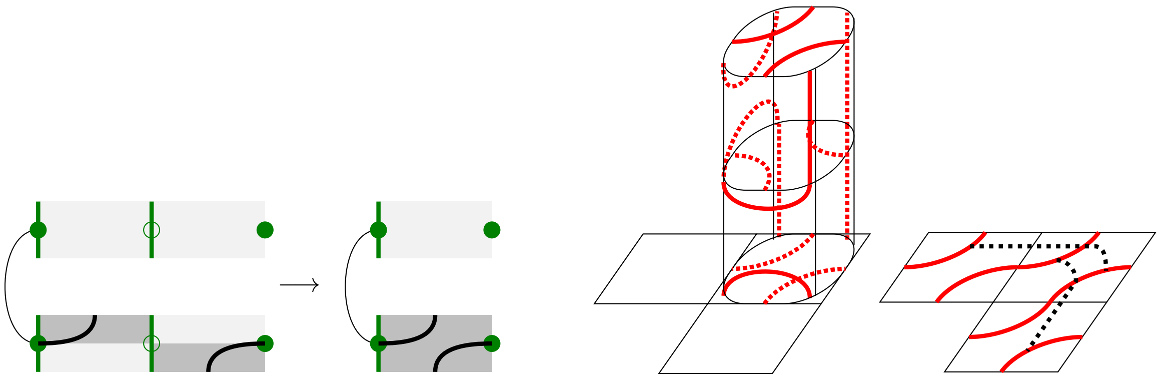

Figure 2: Left: A twisted diagram, the product of two tight diagrams. Right: the corresponding contact geometry. Stacking the two contact cubes yields a tight cube, but bypasses from adjacent squares produce an overtwisted contact structure.

When we then arrive at tensor products of *homology classes* of diagrams in section 4.7, only *four* types of tightness/twistedness remain.

Throughout, it is necessary to consider strand diagrams *locally* at matched pairs; this corresponds to considering contact structures locally at individual cubes of a cubulated contact structure. Indeed, we show that $H(\mathcal{A}(\mathcal{Z}))$ decomposes into a *tensor product* over matched pairs; and we have *local* strand algebras, each with their *local* homology at each matched pair.

Here again, we encounter phenomena not yet given a name in the literature. The observed local diagrams, described as "fragments" of strand diagrams in [21], are not strand diagrams in the usual sense of bordered Floer theory (e.g. [14]) or bordered sutured Floer theory (e.g. [28]), since strands may "run off the top of an arc". Therefore, before we can even start our investigations, we must broaden the usual definition of strand diagrams. In section 2 we therefore introduce a notion of *augmented strand diagram*.

Tensor products of strand diagrams (i.e. elements of $\mathcal{A}(\mathcal{Z})^{\otimes n}$), or their homology classes (i.e. elements of $H(\mathcal{A}(\mathcal{Z}))^{\otimes n}$) thus have a tensor decomposition over matched pairs of $\mathcal{Z}$, into local diagrams, in addition to their obvious decomposition into tensor factors. We regard these two types of decomposition as "vertical" and "horizontal" respectively, and draw pictures accordingly. Contact-geometrically these two types of tensor decomposition correspond to two types of geometric decompositions of stacked contact structures. An element of $CA(\Sigma, Q)^{\otimes n} \cong H(\mathcal{A}(\mathcal{Z}))^{\otimes n}$ can be regarded as a stacking of n cubulated contact structures on $\Sigma \times [0,1]$: this can be cut "horizontally" into n slices, each containing a contact structure on $\Sigma \times [0,1]$; or it can be cut "vertically" to obtain stacked contact structures on $\square \times [0,1]$, over each square $\square$ of the quadrangulation.

We give a complete classification of viable local strand diagrams in section 2.5, summarised in table 1. We show (proposition 4.17) that any viable tensor product of diagrams, observed locally at a single matched pair, must appear as one of the tensor products in the table, up to a notion of *extension* and *contraction*, which provide ways, trivial in a contact-geometric sense, to grow or shrink a tensor product. We are also able to give a complete classification of viable local tensor products of strand diagrams in section 4.5, summarised in table 2. This also yields a complete classification of viable local tensor products of homology classes of strand diagrams, in section 4.7.

Having made such definitions and classifications, we also establish several of their basic properties. In order to prove our main theorems, we need to know facts such as which types of tightness/twistedness can occur within others, as "sub-tensor-products", "vertically" or "horizontally". We consider these and several more questions as we proceed.

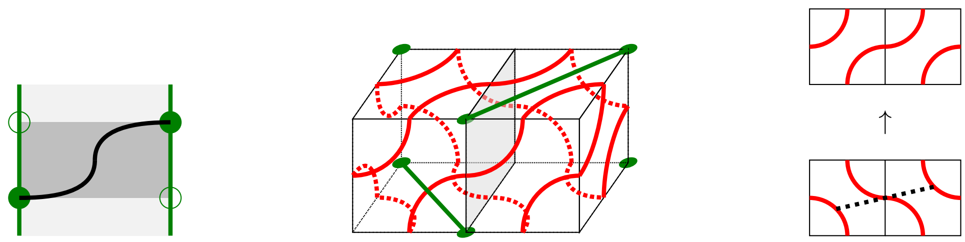

Figure 3: Left: A portion of a strand diagram consisting of a single strand from one place to the next. Centre: The corresponding cubulated contact structure. Right: This contact structure is given by a bypass attachment.

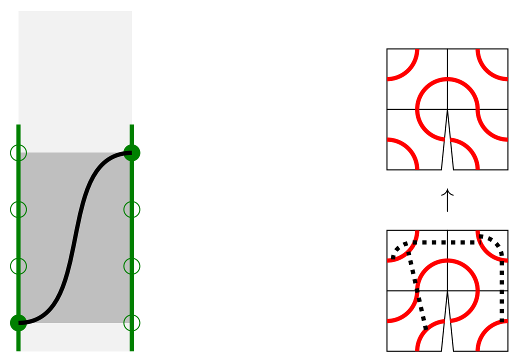

Figure 4: Left: A portion of a strand diagram consisting of a single strand. Right: The corresponding contact structure is given by a sequence of bypass attachments.

1.5 Contact meaning of A-infinity operations

We now attempt to give some idea of what the A_∞ operations X_n "mean" in terms of contact geometry. For details and background on the precise correspondence between contact structures and strand diagrams, we refer to our previous paper [21].

As discussed in [21], a strand diagram D on an arc diagram $\mathcal{Z}$ with appropriate grading (each step of $\mathcal{Z}$ covered at most once; no crossings) can be interpreted as a contact structure on $\Sigma \times [0,1]$. Each matched pair of $\mathcal{Z}$ corresponds to a square of the quadrangulation Q, or a cube in the cubulation $Q \times [0,1]$ of $\Sigma \times [0,1]$.

A strand diagram D containing a single moving strand going from one point ("place" in our terminology) of $\mathcal{Z}$ to the next can be regarded as a *bypass*: in passing from one strand to the next, the strand affects two places, and the corresponding contact structure is a *bypass addition*, where the bypass is placed along the two cubes. Bypasses addition is a basic operation in 3-dimensional contact geometry [4], and in a certain sense is the "simplest" modification one can make to a contact manifold [5]. The result is shown in figure 3.

A strand diagram consisting of a longer strand can (usually; unless other restrictions get in the way! See e.g. the example of figure 11 of [13]) be regarded as a product of diagrams with shorter strands, each covering a single step of $\mathcal{Z}$ as above. The corresponding contact structure is given a sequence of bypass additions very closely related to the *bypass systems* of [15, 16]. See figure 4.

However, when we have a *tensor product* of strand diagrams corresponding to contact structures, the various steps of $\mathcal{Z}$ may not be covered in the order in which they would be covered by single strands. If the various diagrams in the tensor product cover the various steps in a matched pair in a "correct" order, the factors in the tensor product may multiply (using the standard multiplication in

Figure 5: Left: This tensor product (tight, in our classification) has a nonzero product in $\mathcal{A}(\mathcal{Z})$ of $H(\mathcal{A}(\mathcal{Z}))$. Right: This tensor product (critical in our classification) covers the same steps in a different order, and has zero product in $\mathcal{A}(\mathcal{Z})$ or $H(\mathcal{A}(\mathcal{Z}))$, but an A_∞ operation may reorder the bypasses and give a nonzero result.

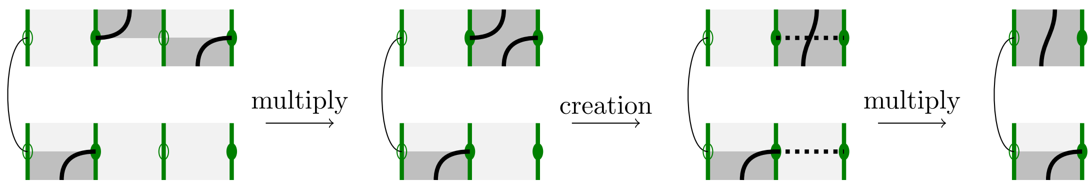

Figure 6: Mechanics of A_∞ operations, effectively reordering bypasses. Multiplying the last two factors of a critical tensor product yields a twisted diagram. A creation operator turns the twisted diagram into a crossed one, and the tensor product becomes sublime. Multiplication then yields a tight diagram.

$\mathcal{A}(\mathcal{Z})$) to give a diagram which is nonzero in homology. This corresponds to a contact structure built out of bypasses as described above. But if the various diagrams cover the various steps in a different order, then they will not multiply to give something nonzero in homology. Moreover, the *Maslov index* at the matched pair will be lower by 1 from the "correct" order.

The simplest example of this phenomenon is shown in figure 2. The product of two diagrams, corresponding to tight contact structures, gives an overtwisted contact structure. But if they were multiplied in the opposite order, the result would be tight. For a slightly more complicated example, still "localised" at a single matched pair, see figure 5.

In general, the A_∞ operation X_n, when it produces a nonzero result, will effectively reorder the bypasses at $n-2$ matched pairs (since it has grading $n-2$) so as to make their product tight. This is the rough meaning of theorem 1.2; the statement is simply an elaboration of this idea, being precise about the various types of tightness/twistedness at each matched pair.

We can say also say a little about how this "reordering" is achieved. As mentioned above, our construction of the A_∞ structure on $CA(\Sigma, Q)$ or $H(\mathcal{A}(\mathcal{Z}))$, following Kadeishvili's method, uses creation operators, whose operation is described locally by figure 1. As seen in figure 2, a creation operator acts on a local diagram which is *twisted*, i.e. represents a "minimally overtwisted" contact structure, and makes it *crossed*.

We may then observe a phenomenon which is rather curious from a contact-geometric point of view. Starting from a tensor product which is twisted (or worse), applying a creation operator yields a tensor product of diagrams including crossings — the most degenerate type of "twistedness". Yet multiplying out this tensor product may yield a diagram corresponding to a tight contact structure! After multiplication, no crossings remain, nor any twistedness. The result is as if the original diagrams were reordered into the "correct" order at that matched pair. See figure 6 for an example based on the "badly ordered" tensor product of figure 5(right).

In this way, strand diagrams may pass from being crossed to tight without being twisted along the way. We call this process *sublimation* because of the "phase-skipping" behaviour analogous to the physical process. We call a tensor product in which the diagrams are not all tight, but their product is tight, *sublime*.

However, it is not the case that X_n *always* performs reorderings and sublimations in this way; it simply *may* do so. Depending on the various choices involved in the construction, the result may or may not be nonzero on various tensor products. Theorem 1.2 tells us what the answer must be, if it exists; and gives necessary conditions for it to be nonzero. Theorem 1.4 does however provide a guarantee that for any A_∞ structures produced by our construction, certain (highly restricted) tensor products always yield a nonzero result.

For lower-level operations, we can say more. We know $X_1 = 0$ and X_2 is just multiplication, and we can in fact give an explicit description of X_3 (proposition 6.9). Beyond that, the multiplicity of choices makes specific statements unwieldy, and theorem 1.4 is the strongest guarantee of nonzero results that we could find, for now.

For the rest of this paper, we work primarily with strand diagrams. But our approach is heavily influenced by contact geometry, and we regularly comment on the contact-geometric significance of our definitions and results. For these comments, we assume some familiarity with the correspondence between strand algebras and contact structures in [21], and refer there for further details.

1.6 Relationship to other work

The strands algebra is a crucial object in bordered Floer theory, appearing in [11, 12, 13, 14]. The slightly more general arc diagrams we use here appeared in Zarev's work [28, 29]. Its homology was explicitly computed in section 4 of [13]. This description was reformulated in [21], where the isomorphism (1) was proved.

The general construction of A_∞ structures on differential graded algebras by Kadeishvili in [8] is part of a much larger subject, not one to which the author claims much expertise. There are other methods, such as those of Kontsevich-Soibelman [10], Nikolov–Zahariev [24] and Huebschmann [7]. We do not know of examples where Kadeishvili's construction has been made as absolutely explicit as by the "creation" operators here.In previous work we have found several roles for objects like creation and annihilation operators in contact geometry [15, 16, 17, 18, 19, 20, 22, 23].

The various contact-geometric interpretations appearing here derive not only from our previous work [21] but also from work on quadrangulated surfaces and their connections to contact geometry, Heegaard Floer theory and TQFT [19, 20]. Some of these ideas are also implicit in Zarev's work cited above. Constructions with bypasses go back to Honda's [4].

The contact category was introduced by Honda in unpublished work. It has been studied by Cooper [1]. Related categorifications have been studied by Tian [26, 27]. The case of discs was considered in our [15] and in detail by Honda-Tian in [6].

1.7 Structure of this paper

As discussed above, there is some work required before we can even properly state our main theorems. We must define the relevant notions and establish their properties as we need them.

We begin in section 2 by considering the algebra and anatomy of strand diagrams. We recall existing definitions in section 2.1, and generalise them to *augmented* diagrams in section 2.2. We can then define the notion of viability in section 2.3. We consider how augmented diagrams can be cut into local diagrams, and the associated algebra, in section 2.4. In section 2.5 we establish terminology for strand diagrams, including occupation of places and pairs for homological grading, and on/off or 1/0 for idempotents; then (section 2.6) we define the types of tightness of local diagrams and (section 2.7) tabulate the various possibilities. In section 2.8 we discuss local strand algebras and their homology, and in section 2.9 the homology of strand algebras in general. In section 2.10 we define and study the types of tightness for viable augmented diagrams. In section 2.11 we consider the set of diagrams representing a homology class, and in section 2.12 we calculate the dimensions of various vector spaces related to strand algebras, which we need later. In section 2.13 we introduce the ideal $\mathcal{F}$ and the quotient algebra to simplify our calculations.

In section 3 we then consider objects parametrising the choices involved in constructing A_∞ structures. We discuss cycle selection homomorphisms in section 3.1. We discuss how different cycle selection maps can differ in section 3.2. We then introduce creation operators in section 3.3, and discuss how they can invert the differential in section 3.4. We put them together into global creation operators in section 3.5, and discuss how they can be obtained from a pair ordering in section 3.6.

We then turn to tensor products of strand diagrams in section 4. We extend the "anatomical" notions and terminology for gradings, viability, occupation and idempotents in section 4.1. We introduce the six types of tightness/twistedness in section 4.2. We discuss sub-tensor-products, and the associated notions of extension and contraction, in section 4.3. We consider the two most curious types of tightness, sublime and singular, in section 4.4. We can then give a full enumeration of all possible viable local tensor products in section 4.5. In section 4.6 we consider how tightness of tensor products and sub-tensor-products are related. We may then consider tensor products of homology classes of diagrams in section 4.7, discussing their tightness, enumerating the possible local tensor products, and establishing some of their properties. In section 4.8 we consider a generalised notion of contraction for tensor products of homology classes.

We then have everything we need to construction A_∞ structures explicitly in section 5. The construction itself is given in section 5.1, proving theorem 1.1. In section 5.2 we establish a shorthand notation for tensor products of strand diagrams. In section 5.3 we calculate some examples at low levels of the A_∞ structure.

In section 6 we then discuss some properties of the A_∞ structures constructed, and in fact slightly more general A_∞ structures from Kadeishvili's construction. In section 6.1 we discuss how A_∞ operations relate to viability. In section 6.2 we discuss how the various choices made in Kadeishvili's construction affect the result. Then in section 6.3 we establish some of the elementary properties of the constructed A_∞ operations, and in section 6.4 prove some necessary conditions for nontrivial A_∞ operations, including those of theorem 1.2. In section 6.5 we establish general properties of the A_∞ maps at levels up to 3.

In a brief section 7 we calculate some further examples, at levels 3 (section 7.1) and 4 (section 7.2), illustrating some of the complexities which arise.

Finally in section 8 we consider higher A_∞ operations and when they are nontrivial. We introduce the notion of operation trees in section 8.1, and notions of *validity* and *distributivity* in section 8.2. In section 8.3 we discuss some constructions we need on trees (*joining* and *grafting*). Then in section 8.4 we show how certain trees are required for nonzero results, proving proposition 1.3. In section 8.5 we discuss the operation trees local to a matched pair, and classify them in section 8.6. In section 8.7 we introduce a stronger notion of validity necessary for our results, and after discussing some further operations on trees of *transplantation* and *branch shifts* in section 8.8, and introducing a stronger notion of distributivity in section 8.9, we prove theorem 1.4 in section 8.10.

1.8 Acknowledgments

The author thanks Robert Lipshitz for posing the questions that sparked this work, and for helpful conversations. He also thanks Anita for putting up with him. He is supported by Australian Research Council grant DP160103085.

2 Algebra and anatomy of strand diagrams

2.1 Strand diagrams

We recall the definition of strand diagrams, before proceeding in section 2.2 to augment them. We follow our previous paper [21], which in turn is based on Zarev [28], as well as Lipshitz–Ozsváth–Thurston [11, 13]. We refer to these papers for further details.

An *arc* diagram consists of a triple $\mathcal{Z} = (\mathbf{Z}, \mathbf{a}, M)$, where $\mathbf{Z} = \{Z_1, \ldots, Z_l\}$ is a set of oriented line segments (intervals), $\mathbf{a} = (a_1, \ldots, a_{2k})$ is a sequence of distinct points in the interior of the line segments of $\mathbf{Z}$, ordered along the intervals, and $M\colon \mathbf{a} \longrightarrow \{1, 2, \ldots, k\}$ is a 2-to-1 function. Performing oriented surgery on $\mathbf{Z}$ at all the 0-spheres $M^{-1}(i)$ is required to yield an oriented 1-manifold consisting entirely of arcs (no circles).

We call the points of $\mathbf{a}$ *places*. We call two places a_i, a_j matched by M (i.e. such that $M(a_i) = M(a_j)$) *twins*. There is a partial order on $\mathbf{a}$ where $a_i \leq a_j$ if a_i, a_j lie on the same oriented interval, and are in order along it. The function M partitions $\mathbf{a}$ into k *matched pairs* or just *pairs*.

An *unconstrained strand diagram* over $\mathcal{Z}$ is a triple $\mu = (S, T, \phi)$ where $S, T \subseteq \{a_1, \ldots, a_{2k}\}$ with $|S| = |T|$ and $\phi\colon S \longrightarrow T$ is a bijection, which is increasing with respect to the partial order on $\mathbf{a}$ in the sense that $\phi(x) \geq x$ for all $x \in S$. There is a standard way to draw an unconstrained strand diagram in the plane (in fact in $[0,1] \times \mathbf{Z}$), with $|S| = |T|$ *strands*. The strands *begin* at S (drawn at $\{0\} \times S$), *end* at T (drawn at $\{1\} \times T$), and move to the right (in the positive direction along $[0, 1]$), never going down, and meeting efficiently without triple crossings. We say μ *goes from S to T*. The points of $\mathbf{a}$ split $\mathbf{Z}$ into intervals called *steps*, of two types: *interior* to an interval Z_i, and *exterior*, i.e. at the boundary of a Z_i. The product $\mu\nu$ of two strand diagrams $\mu = (S, T, \phi)$ and $\nu = (U, V, \psi)$ is given by $(S, V, \psi \circ \phi)$, provided that $T = U$ and the composition $\psi \circ \phi\colon S \to V$ satisfies $\mathrm{inv}(\psi \circ \phi) = \mathrm{inv}(\phi) + \mathrm{inv}(\psi)$; otherwise it is zero. Here $\mathrm{inv}(\mu)$ is the number of inversions, or crossings, in μ. Equivalently, the product $\mu\nu$ is given by concatenating strand diagrams, provided that there are no "excess inversions", i.e. crossings which can be simplified by a Reidemeister II-type isotopy of strands relative to endpoints. There is a differential ∂ which resolves crossings in strand diagrams; $\partial\mu$ is the sum of all strand diagrams obtained from μ by resolving a crossing so that the number of crossings decreases by exactly 1. This structure makes the free $\mathbb{Z}_2$-module on strand diagrams over $\mathcal{Z}$ into a differential graded algebra over $\mathbb{Z}_2$, which we denote $\widetilde{\mathcal{A}}(\mathcal{Z})$. For each subset $S \subseteq \mathbf{a}$ there is an idempotent $I(S)$.

A *$\mathcal{Z}$-constrained*, or just *constrained*, strand diagram takes into account also the matching M of $\mathcal{Z}$. For each $s \subseteq \{1, \ldots, k\}$ we define $I(s) = \sum_S I(S)$, where the sum is over sections S of s under M. Here a *section* of s means an $S \subseteq \mathbf{a}$ such that $M|_S$ is a bijection. The $I(s)$ generate a $\mathbb{Z}_2$-subalgebra of $\widetilde{\mathcal{A}}(\mathcal{Z})$. A strand diagram which begins at a section of s and ends at a section of t, for $s, t \subseteq \{1, \ldots, k\}$, is said to be *$\mathcal{Z}$-constrained*. We say it *begins* at s and *ends* at t, or *goes from s to t*; we also say $I(s)$, or by abuse of notation just s, is the *initial idempotent*, and $I(t)$ or t is the *final idempotent*. Thus a constrained strand diagram begins and ends at subsets of $\mathbf{a}$ which contain at most one place of each matched pair. If $I(s)\widetilde{\mathcal{A}}(\mathcal{Z})I(t)$ is nonzero then $|s| = |t| = i$, in which case it is freely generated as a $\mathbb{Z}_2$-module by $\mathcal{Z}$-constrained strand diagrams of i strands from s to t.

Finally, we symmetrise our strand diagrams with respect to the matched pairs. If $\mu = (S, T, \phi)$ is an unconstrained strand diagram on $\mathcal{Z}$ without horizontal strands (i.e. ϕ has no fixed points) then we consider adding horizontal strands to μ at some places $U \subseteq \mathbf{a} \backslash (S \cup T)$, i.e. adding fixed points to ϕ to obtain a function $\phi_U\colon S \cup U \longrightarrow T \cup U$, which is still a bijection with $\phi(x) \geq x$. We define $a(\mu)$ to be the sum of all strand diagrams that can be obtained from μ by adding horizontal strands,

$$a(\mu) = a(S, T, \mu) = \sum_U (S \cup U, T \cup U, \phi_U) \in \widetilde{\mathcal{A}}(\mathcal{Z}),$$

and then for each $s, t \subseteq \{1, \ldots, k\}$, $I(s)a(\mu)I(t)$ is either zero, or the sum of all $\mathcal{Z}$-constrained strand diagrams from s to t obtained from μ by adding horizontal strands. (Left-multiplying by $I(s)$ filters for diagrams which start at s; right-multiplying by $I(t)$ filters for diagrams which start at t; multiplying by both ensures the result is $\mathcal{Z}$-constrained.) Note that if it is possible to add a horizontal strand to μ at a place a of a matched pair $\{a, a'\}$ to obtain a strand diagram in $I(s)a(\mu)I(t)$, then it is also possible to add a horizontal strand at the twin place a'. In this case every diagram in $I(s)a(\mu)I(t)$ contains a horizontal strand at precisely one of a or a'; further, for every diagram with a horizontal strand at a appearing in $I(s)a(\mu)I(t)$, the corresponding diagram with a horizontal strand at a' and otherwise identical will also appear. If there are j such pairs $\{a, a'\}$, then $I(s)a(\mu)I(t)$ is a sum of 2^j terms, one for each choice of a or a' in each pair.

We denote such a sum $I(s)a(\mu)I(t)$ as a single diagram D by drawing the $2j$ horizontal strands involved as dotted and call it a *symmetrised $\mathcal{Z}$-constrained strand diagram*. In such a diagram, every horizontal strand is dotted, and dotted horizontal strands come in pairs. So a symmetrised $\mathcal{Z}$-constrained strand diagram with j pairs of dotted horizontal strands is in fact a sum of 2^j $\mathcal{Z}$-constrained strand diagrams.

The *strand algebra* $\mathcal{A}(\mathcal{Z})$ is the subalgebra of $\widetilde{\mathcal{A}}(\mathcal{Z})$ generated by symmetrised $\mathcal{Z}$-constrained strand diagrams. It is preserved by ∂ and hence forms a differential graded algebra. This algebra has several gradings.

The *homological* (also known as the *spin-c grading* or *Alexander* grading), which we abbreviate to H-grading, is valued in $H_1(\mathbf{Z}, \mathbf{a})$. Given a strand map $\mu = (S, T, \phi)$ on $\mathcal{Z}$, for each $a \in S$, the oriented interval $[a, \phi(a)]$ from a to $\phi(a)$ gives a homology class in $H_1(\mathbf{Z}, \mathbf{a})$, and the H-grading of μ, denoted $h(\mu)$ or just h, is the sum of such intervals $[a, \phi(a)]$ over all $a \in S$. In other words, h counts how often each step of $\mathbf{Z}$ is covered. Since horizontal strands cover no steps, a symmetrised constrained diagram D has a well-defined H-grading $h(D)$. The H-grading is additive under multiplication of strand diagrams, and ∂ preserves H-grading. We denote by $\mathcal{A}(\mathcal{Z}; h)$ the $\mathbb{Z}_2$-submodule of $\mathcal{A}(\mathcal{Z})$ generated by diagrams with H-grading h, so we have a direct sum decomposition $\mathcal{A}(\mathcal{Z}) = \bigoplus_h \mathcal{A}(\mathcal{Z}; h)$.

If D is a (symmetrised constrained) diagram from s to t (here $s, t \subseteq \{1, \ldots, k\}$) with H-grading $h \in H_1(\mathbf{Z}, \mathbf{a})$, we define the *H-data* of D to be the triple (h, s, t). In other words, the H-data of D consists of its H-grading, i.e. how it covers the steps of $\mathcal{Z}$, together with its initial and final idempotents. Note that h does not in general determine s or t. By inspecting h we may deduce that certain strands must start or end at certain points in $\mathcal{Z}$: in particular, when the local multiplicity of h changes as we pass from one step to the next, necessarily by 1 or -1, a strand must respectively begin or end. But when the local multiplicity of h does not change from one step to the next, we cannot tell whether strands begin or end. In particular, h does not give any information about horizontal strands. Nonetheless $\mathcal{A}(\mathcal{Z})$ naturally decomposes as a direct sum of $\mathbb{Z}_2$-modules over H-data: we define $\mathcal{A}(\mathcal{Z}; h, s, t) = I(s)\mathcal{A}(\mathcal{Z}; h)I(t)$, and then

$$\mathcal{A}(\mathcal{Z}) = \bigoplus_{h,s,t} \mathcal{A}(\mathcal{Z}; h, s, t) = \bigoplus_{h,s,t} I(s)\mathcal{A}(\mathcal{Z}; h)I(t).$$

This decomposition includes the decomposition over idempotents $I(s), I(t)$, and also over H-grading.

The *Maslov* grading of $\mathcal{A}(\mathcal{Z})$ is valued in $\frac{1}{2}\mathbb{Z}$. If μ is a $\mathcal{Z}$-constrained strand diagram (not yet symmetrised) from S to T with H-grading h, then its Maslov grading is

$$\iota(\mu) = \mathrm{inv}(\mu) - m\,(h, S)\,,$$

where the function

$$m \colon H_1(\mathbf{Z}, \mathbf{a}) \times H_0(\mathbf{a}) \longrightarrow \frac{1}{2}\mathbb{Z}$$

counts local multiplicities of strand diagrams around places. Specifically, for a place a and $h \in H_1(\mathbf{Z}, \mathbf{a})$, $m(h, a)$ is the average of the local multiplicities of h on the steps after and before a. Thus $\iota(\mu)$ has a contribution of $+1$ from each crossing; and then non-positive contributions from each place of S, depending on the multiplicity of h near that place. It is not difficult to check that all the constrained diagrams in a symmetrised $\mathcal{Z}$-constrained diagram D have the same Maslov grading, so the Maslov grading $\iota(D)$ of D is the grading of any of the constrained diagrams in it.

The differential ∂ does not affect H-grading or idempotents, but lowers the number of crossings in a diagram by 1 (if the result is nonzero), hence has Maslov degree -1. The Maslov index does not respect multiplication in $\mathcal{A}(\mathcal{Z})$; rather, for symmetrised $\mathcal{Z}$-constrained strand diagrams D, D' with H-gradings h, h' we have

$$\iota(DD') = \iota(D) + \iota(D') + m(h', \partial h).$$

The homology of $\mathcal{A}(\mathcal{Z})$ was described by Lipshitz–Ozsváth–Thurston [13, thm. 9]. As the differential respects H-data (h, s, t), the decomposition $\mathcal{A}(\mathcal{Z}) = \bigoplus_{h,s,t} \mathcal{A}(\mathcal{Z}; h, s, t)$ descends to homology:

$$H(\mathcal{A}(\mathcal{Z})) = \bigoplus_{h,s,t} H(\mathcal{A}(\mathcal{Z}; h, s, t)).$$

Lipshitz–Ozsváth–Thurston showed that the summand $H(\mathcal{A}(\mathcal{Z}; h, s, t))$ is nontrivial if and only if there exists a symmetrised $\mathcal{Z}$-constrained strand diagram D with H-data (h, s, t) *without crossings*, satisfying two conditions:

(i) the multiplicity of h on every step of $\mathbf{Z}$ is 0 or 1; and

(ii) if $\{a, a'\}$ is a matched pair with a in the interior of the support of h, and a' not in the interior of the support of h, then a does not lie in both s and t.

Such a D, having no crossings, is obviously a cycle and in fact the homology class of any such D generates $H(\mathcal{A}(\mathcal{Z}; h, s, t)) \cong \mathbb{Z}_2$. We use property (i) extensively as a notion of *viability* in the sequel, from section 2.3 onwards. We discuss and reformulate the second condition further in section 2.8 below; see also [21, sec. 3.5–3.7]. For such D, the expression for the Maslov index simplifies to $\iota(D) = -m(h, S)$, where S is the initial idempotent of any $\mathcal{Z}$-constrained strand diagram in D.

Since we usually work with a single arc diagram $\mathcal{Z}$, we often leave $\mathcal{Z}$ implicit and use the shorthand

$$\mathcal{A} = \mathcal{A}(\mathcal{Z}), \quad \mathcal{A}(h, s, t) = \mathcal{A}(\mathcal{Z}; h, s, t), \quad \mathcal{H} = H(\mathcal{A}), \quad \mathcal{H}(h, s, t) = H(\mathcal{A}(h, s, t)).$$

The homology $\mathcal{H}$ inherits multiplication from $\mathcal{A}$ and becomes a differential graded algebra with trivial differential. The point of this paper is to extend this DGA structure to A_∞-structures.

2.2 Augmented strand diagrams

In a symmetrised $\mathcal{Z}$-constrained strand diagram, strands run between places in $\mathbf{a} = (a_1, \ldots, a_{2k})$. Since the places of $\mathbf{a}$ lie in the interior of the intervals Z_i of $\mathbf{Z}$, no strand ever reaches an endpoint of any interval Z_i. In other words, strand diagrams only cover interior steps of $\mathbf{Z}$.

In the sequel however we need to consider strand diagrams where strands cover exterior steps of $\mathbf{Z}$ and reach endpoints of the intervals Z_i. We describe this as *flying off* an interval. *Augmented* strand diagrams, which we define presently, extend strand diagrams to allow such behaviour.

To define augmented diagrams formally we again use non-decreasing bijections, but now on sets including the endpoints of each interval. Let the endpoints of the interval Z_i be $-\infty_i$ and $+\infty_i$, at the start and end respectively. A strand flies off the top end of an interval Z_i if some $x \neq +\infty_i$ is sent to $+\infty_i$, and a strand flies off the bottom if $x \neq -\infty_i$ is satisfies $-\infty_i \mapsto x$. A strand may fly off both ends of an interval if $-\infty_i \mapsto +\infty_i$. We also allow horizontal strands at $\pm\infty_i$, but these present a slight subtlety, discussed below: they simply exist for technical reasons.

Let $\mathbf{a}_{\pm\infty} = \mathbf{a} \cup \{-\infty_1, \ldots, -\infty_l\}_{i=1}^{l} \cup \{+\infty_1, \ldots, +\infty_l\}_{i=1}^{l}$. The points of $\mathbf{a}_{\pm\infty}$ are naturally partially ordered by the total order along each interval, extending the partial order on $\mathbf{a}$.

An *unconstrained augmented strand diagram* over $\mathcal{Z}$ is a triple (S, T, ϕ), where $S, T \subseteq \mathbf{a}_{\pm\infty}$ and $\phi \colon S \longrightarrow T$ is a bijection which is increasing with respect to the partial order on $\mathbf{a}_{\pm\infty}$ in the sense that $\phi(x) \geq x$ for all $x \in S$. Again we say μ *goes from* S to T. The product of two such diagrams is given by composition of bijections ϕ, if such a composition exists and has no excess inversions, otherwise it is zero. Equivalently, the product is given by by concatenating diagrams, if a concatenation exists and has no excess crossings. A differential is again defined by resolving crossings so that the number of crossings decreases by exactly 1. We define $\widetilde{\mathcal{A}}^{aug}(\mathcal{Z})$ to be the free $\mathbb{Z}_2$-module on unconstrained augmented strand diagrams over $\mathcal{Z}$; it is a differential graded algebra over $\mathbb{Z}_2$ with an idempotent $I(S)$ for each $S \subseteq \mathbf{a}_{\pm\infty}$.

A subtlety arises here because if an (unconstrained) augmented strand diagram μ has a strand (say) flying off the end of an interval to $+\infty_i$, it should still be able to give a nonzero result when composed with another diagram on the right, which does not have any strand at $+\infty_i$. We extend our notion of matching to achieve this effect, but it is no longer a function; rather it is a *partial function* (i.e. partially defined).

To this end, we extend the matching $M \colon \mathbf{a} \longrightarrow \{1, \ldots, k\}$ to the partial function $M^{aug} \colon \mathbf{a}_{\pm\infty} \longrightarrow \{1, \ldots, k\}$, which is equal to M on $\mathbf{a}$, and is *not defined* on each $-\infty_i$ or $+\infty_i$.

Given a set $s \subseteq \{1, \ldots, k\}$, a *section* of s under M^{aug} is then any set $S \subseteq \mathbf{a}_{\pm\infty}$ such that the restriction of M^{aug} to S is a (possibly partially defined) function mapping surjectively and injectively to s. Thus a section of s under M^{aug} consists of a section of s under M, together with any subset of $\{-\infty_i, +\infty_i\}_{i=1}^{l}$.

For $s \subseteq \{1, \ldots, k\}$, define $I(s) = \sum_S I(S)$, the sum over sections S of s under M^{aug}. The $I(s)$ again generate a subalgebra of $\widetilde{\mathcal{A}}^{aug}(\mathcal{Z})$. An augmented strand diagram which begins at a section of s and ends at a section of t, for $s, t \subseteq \{1, \ldots, k\}$, is *$\mathcal{Z}$-constrained*; we say it *goes from s to t*. If $I(s)\widetilde{\mathcal{A}}^{aug}(\mathcal{Z})I(t)$ is nonzero, then there is at least one section S of s under M^{aug}, and at least one section T of t under M^{aug}, such that there exists an (unconstrained) augmented strand diagram from S to T. Note then that s and t need not have the same size, because S and T can also contain points of the form $+\infty_i$ or $-\infty_i$.

If $\mu = (S, T, \phi)$ is an unconstrained augmented strand diagram on $\mathcal{Z}$ without horizontal strands, we again consider adding horizontal strands to μ at places $U \subseteq \mathbf{a}_{\pm\infty} \backslash (S \cup T)$ (there can be horizontal strands at $\pm\infty_i$), extending ϕ by the identity to $\phi_U \colon S \cup U \longrightarrow T \cup U$, and defining $a(\mu) = \sum_U (S \cup U, T \cup U, \phi_U) \in \widetilde{\mathcal{A}}^{aug}(\mathcal{Z})$. For each $s, t \subseteq \{1, \ldots, k\}$ then $I(s)a(\mu)I(t)$ is the sum of all $\mathcal{Z}$-constrained augmented strand diagrams obtained from μ by adding horizontal strands, possibly at interval endpoints $\pm\infty_i$. As in the non-augmented case, if one such diagram has j horizontal strands at the places of $\mathbf{a}$, these horizontal strands can be swapped with their twins, resulting in 2^j possible arrangements of horizontal strands at these places. Unlike the non-augmented case, for any point of the form $-\infty_i$ or $+\infty_i$ not in $S \cup T$, a horizontal strand can be added at this point. Thus if $|\bigcup_{i=1}^{l} \{-\infty_i, +\infty_i\} \backslash (S \cup T)| = n$, then there are 2^n possible arrangements of horizontal strands at these endpoints (one for each subset of $|\bigcup_{i=1}^{l} \{-\infty_i, +\infty_i\} \backslash (S \cup T)|$). Thus $I(s)a(\mu)I(t)$ is a sum of 2^{j+n} $\mathcal{Z}$-constrained augmented strand diagrams. We can draw such a sum as a single diagram D with $2j$ dotted horizontal strands (leaving the possible horizontal strands at $\pm\infty_i$ implicit) and we call it a *symmetrised $\mathcal{Z}$-constrained augmented strand diagram* or just *diagram*.

Multiplication of diagrams is described as follows. Consider the product of two (symmetrised $\mathcal{Z}$-constrained augmented strand) diagrams D, D'. If no strand in D or D' flies off an interval, then their product DD' as augmented diagrams is given by concatenating strands, just as for non-augmented diagram. Formally the symmetrised augmented diagram is a sum of 2^n diagrams, involving possible horizontal strands at $\pm\infty_i$, but the augmented diagram DD' is drawn identically to the diagram of the the product of non-augmented diagrams. Thus, at least at the level of drawing diagrams, multiplication of augmented diagrams is an extension of multiplication of (non-augmented) diagrams.

If on some interval Z_i, both D and D' fly off the top end, then $DD' = 0$. This is because, for any $\mathcal{Z}$-constrained augmented strand diagram (S, T, ϕ) in D, and any such diagram (S', T', ϕ') in D', ϕ has $+\infty_i$ in its image, but ϕ' does not have $+\infty_i$ in its domain, so the functions cannot be composed. Similarly, if both D, D' fly off the negative end, then $DD' = 0$. If D flies off the top end of Z_i but D' does not, then the composition is well defined there: each (S, T, ϕ) in D has $+\infty_i$ in the image of ϕ; and half of the constrained augmented diagrams (S', T', ϕ') in D' have ϕ' mapping $+\infty_i \mapsto +\infty_i$ (i.e. a horizontal strand at $+\infty_i$), so such ϕ' compose with ϕ at $+\infty_i$. If D' flies off the top end of Z_i but D does not, then again composition is well defined: each (S', T', ϕ') in D' has $+\infty_i$ in its image, but not in its domain; and half the (S, T, ϕ) in D do not have $+\infty_i$ in the domain or image; and these ϕ and ϕ' compose without any problems at $+\infty_i$. Thus, if one of D, D' flies off the top end of Z_i and the other does not, then the product DD' is well defined there. Similarly, if one of D, D' flies off the bottom end of Z_i and the other does not, then the product DD' is well defined there. Thus, roughly, if we can concatenate strands of D and D' into another augmented diagram, with at most one strand flying off any end of any interval, then the product DD' is given by concatenating strands, just as for (non-augmented) strand diagrams. Some examples are shown in figure 7 .

The *augmented strand algebra* $\mathcal{A}^{aug}(\mathcal{Z})$ is the subalgebra of $\widetilde{\mathcal{A}}^{aug}(\mathcal{Z})$ generated by (symmetrised $\mathcal{Z}$-constrained augmented strand) diagrams. It is preserved by ∂ and forms a differential graded algebra. There is again an *H-grading* h given by taking the sum of oriented intervals $[a, \phi(a)]$, and regarding it as an element of a relative first homology group of the intervals $\mathbf{Z}$. However now the endpoints of $[a, \phi(a)]$

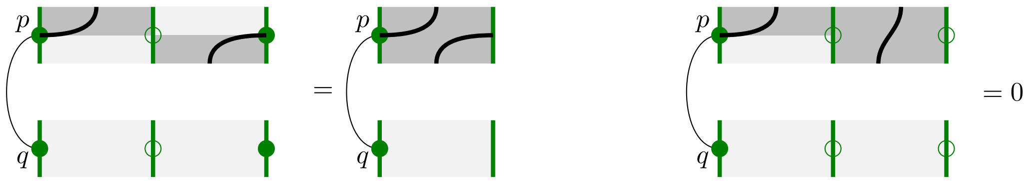

Figure 7: Multiplication of augmented diagrams.

may include the $\pm\infty_i$, so $h \in H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$. Note that $H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$ naturally contains $H_1(\mathbf{Z}, \mathbf{a})$ as a subgroup, and we will always regard it as such: $H_1(\mathbf{Z}, \mathbf{a}) \subset H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$. So we can regard H-grading for augmented diagrams as an extension of H-grading for (non-augmented) diagrams. Again we write $\mathcal{A}^{aug}(\mathcal{Z}; h)$ for the submodule of $\mathcal{A}^{aug}(\mathcal{Z})$ with H-grading h and have a direct sum decomposition $\mathcal{A}^{aug}(\mathcal{Z}) = \bigoplus_h \mathcal{A}^{aug}(\mathcal{Z}; h)$. Again if D is an augmented diagram from s to t with H-grading h, the *H-data* of D is the triple (h, s, t). Again we write $\mathcal{A}^{aug}(\mathcal{Z}; h, s, t) = I(s)\mathcal{A}^{aug}(\mathcal{Z}; h)I(t)$ and then $\mathcal{A}^{aug}(\mathcal{Z}) = \bigoplus_{h,s,t} \mathcal{A}^{aug}(\mathcal{Z}; h, s, t)$.

A $\mathcal{Z}$-constrained augmented diagram μ (not yet symmetrised) from S to T with H-grading h has *Maslov* grading again given by $\iota(\mu) = \iota(\mu) = \text{inv}(\mu) - m\,(h, S) \in \frac{1}{2}\mathbb{Z}$, where inv counts inversions/crossings, and $m\colon H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty}) \times H_0(\mathbf{a}) \longrightarrow \frac{1}{2}\mathbb{Z}$ counts local multiplicities of augmented diagrams around places a_i in S. (We use $H_0(\mathbf{a})$ rather than $H_0(\mathbf{a}_{\pm\infty})$ so that Maslov grading is additive when we glue arc diagrams together. The points $\pm\infty_i$ are not places like the a_i.) Again all the diagrams in a symmetrised diagram have the same Maslov grading. When we add a horizontal strand at a $\pm\infty_i$, the fact that we can add the strand means that there is no strand at $\pm\infty_i$ for the horizontal strand to cross; moreover the horizontal strand at $\pm\infty_i$ does not contribute to $m(h, S)$). Thus Maslov grading is well defined on symmetrised $\mathcal{Z}$-constrained augmented diagrams.

Again the differential ∂ respects H-data but has Maslov degree -1. Maslov index behaves under multiplication as in the non-augmented case. When we have $h \in H_1(\mathbf{Z}, \mathbf{a}) \subset H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$ then strands do not fly off intervals and we have an isomorphism of differential graded algebras

$$\mathcal{A}(\mathcal{Z}; h, s, t) \cong \mathcal{A}^{aug}(\mathcal{Z}; h, s, t).$$

The isomorphism takes a symmetrised diagram $D \in \mathcal{A}(\mathcal{Z}; h, s, t)$ (formally a sum of 2^j constrained diagrams) to the element of $\mathcal{A}^{aug}(\mathcal{Z}; h, s, t)$ represented by the same diagram (formally a sum of 2^{j+2l} constrained diagrams, where l is the number of intervals in $\mathbf{Z}$; all possible horizontal strands at $\pm\infty_i$ are now included). We draw the same diagrams and treat them the same way in both cases.

Accordingly, throughout this paper we regard augmented diagrams as a generalisation of non-augmented diagrams, even though the definition is not formally a generalisation. Alternatively we can regard non-augmented diagrams as augmented diagrams with H-grading zero on exterior steps, in which case augmented diagrams do become a generalisation in a formal sense.

Thus, we drop the "aug" from our notation and simply write $\mathcal{A}(\mathcal{Z})$ or $\mathcal{A}$ for the augmented strand algebra. When referring to specific H-data (h, s, t) we do not distinguish between $\mathcal{A}$ and $\mathcal{A}^{aug}$, since the two summands are isomorphic when both defined. We therefore write $\mathcal{A}(h, s, t)$ for either $\mathcal{A}(\mathcal{Z}; h, s, t)$ or $\mathcal{A}^{aug}(\mathcal{Z}; h, s, t)$, and $\mathcal{H}(h, s, t)$ for either $H(\mathcal{A}(\mathcal{Z}; h, s, t))$ or $H(\mathcal{A}^{aug}(\mathcal{Z}; h, s, t))$.

To summarise: (symmetrised constrained augmented strand) diagrams are a generalisation of symmetrised constrained strand diagrams — generalising the full differential graded algebra structure of strand diagrams, as well as all gradings and idempotents.

2.3 Viability

As mentioned above in section 2.1, the following is a necessary condition for $\mathcal{H}(h, s, t)$ to be nontrivial.

Definition 2.1. *Let $\mathcal{Z} = (\mathbf{Z}, \mathbf{a}, M)$ be an arc diagram.*

(i) *An element* $h \in H_1(\mathbf{Z}, \mathbf{a})$ *or* $H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$ *is* viable *if* h *has multiplicity* 0 *or* 1 *on each step of* $\mathbf{Z}$.

(ii) *A set of H-data* (h, s, t) *is* viable *if* h *is viable.*

(iii) *A* ($\mathcal{Z}$*-constrained augmented strand) diagram* μ *is* viable *if its H-grading is viable.*

(iv) *A summand* $\mathcal{A}(h, s, t)$ *or* $\mathcal{H}(h, s, t)$ *of* $\mathcal{A}$ *or* $\mathcal{H}$ *is* viable *if* h *is viable.*

(v) *An element of* $\mathcal{A}$ *or* $\mathcal{H}$ *is* viable *if it lies in a viable summand* $\mathcal{A}(h, s, t)$ *or* $\mathcal{H}(h, s, t)$.

Thus a diagram $D \in \mathcal{A}$ (or its homology class in $\mathcal{H}$) is viable iff its H-grading is viable.

Note that a set of H-data (h, s, t) may be viable, even though no augmented diagram exists with that H-data! This subtle point is important in the sequel, from section 4.2 onwards.

Lemma 2.2. *In a viable augmented diagram, every crossing is at a horizontal strand.*

Proof. Any diagram with a crossing involving two non-horizontal strands has a step covered with multiplicity at least two. □

Thus, whenever we apply the differential ∂ to a viable diagram, any crossing resolved involves a dotted horizontal strand at some place p of a pair $P = \{p, q\}$. Resolving that crossing affects the strands at p, and makes the dotted horizontal strand at q disappear. Thus the differential acts "locally" on viable diagrams, each resolution at a specific matched pair. We discuss this idea of "locality" next.

2.4 Local diagrams

In the arc diagram $\mathcal{Z} = (\mathbf{Z}, \mathbf{a}, M)$, consider cutting the intervals $Z_1, \ldots, Z_l$ of $\mathbf{Z}$ into sub-intervals, each containing precisely one place. This cuts $\mathcal{Z}$ into disconnected arc diagrams. There is one connected arc diagram for each matched pair. We call the connected arc diagram so obtained containing the matched pair P the *fragment* of $\mathcal{Z}$ at P, denoted $\mathcal{Z}_P$. We can cut $\mathbf{Z}$ at different points between places, but the results are homeomorphic, so $\mathcal{Z}_P$ is well defined up to homeomorphism.

A fragment $\mathcal{Z}_P$ contains just one matched pair consisting of two places; it is the smallest possible nontrivial arc diagram, the only arc diagram up to homeomorphism with one matched pair.

Under the correspondence between arc diagrams $\mathcal{Z}$ and quadrangulated surfaces (Σ, Q) of [21], cutting $\mathcal{Z}$ into fragments corresponds to cutting Σ into squares.

Let now D be a (symmetrised $\mathcal{Z}$-constrained augmented strand) diagram. When we cut $\mathcal{Z}$ into fragments, we would like to cut D into fragments also. Note that even if D is not an augmented diagram, strands may fly off the ends of intervals in fragments, so after cutting into fragments the resulting strand diagram may be augmented.

If D has a crossing involving two non-horizontal strands, then problems arise. For one thing, D can be drawn in various ways, so that the crossing appears at various possible locations in $[0, 1] \times \mathbf{Z}$; after cutting there is then no well-defined fragment in which the crossing appears. For another, after cutting, several strands may fly off the same end of a fragment, which is not permitted in augmented diagrams.

However, if D is *viable* these problems disappear. By lemma 2.2 all crossings occur at horizontal strands, so are localised to specific places. Viability also ensures that each interior step of $\mathbf{Z}$ is covered with multiplicity at most 1, so after cutting $\mathcal{Z}$, at most one strand flies off each end of an interval. We therefore obtain a well-defined augmented diagram on each fragment.

Definition 2.3. *Let* $P = \{p, q\}$ *be a matched pair of the arc diagram* $\mathcal{Z}$*, and let* D *be a viable diagram on* $\mathcal{Z}$*. The* local diagram D_P *of* D *at* P *is the diagram obtained on* $\mathcal{Z}_P$ *after cutting* $\mathcal{Z}$ *into fragments. It lies in the* local strand algebra $\mathcal{A}(\mathcal{Z}_P)$.

Note that since a symmetrised diagram D may contain pairs of dotted horizontal arcs at matched pairs, the local diagram D_P may contain a pair of dotted horizontal arcs.

When the larger arc diagram $\mathcal{Z}$ is understood, we can make the following abbreviations for various augmented strand algebras and summands:

$$\mathcal{A}_P = \mathcal{A}(\mathcal{Z}_P), \quad \mathcal{A}_P(h,s,t) = \mathcal{A}(\mathcal{Z}_P; h,s,t).$$

We observe that the H-data of each D_P is just a restriction of the H-data of D. Maslov gradings are related by $\iota(D) = \sum_P \iota(D_P)$. We write (h,s,t) for the H-data of D, and (h_P, s_P, t_P) for the H-data of D_P.

In addition to cutting diagrams into fragments, we can glue fragments together into diagrams. From augmented diagrams D_P on each $\mathcal{Z}_P$, which fit together in the sense that strands flying off intervals connect, we can glue them together to obtain a viable diagram on $\mathcal{Z}$. Thus, studying viable diagrams locally is equivalent to studying augmented diagrams on a fragment.

For viable H-data (h,s,t) on $\mathcal{Z}$, we thus have

$$\mathcal{A}(h,s,t) \cong \bigotimes_{\text{matched pairs } P} \mathcal{A}_P(h_P, s_P, t_P).$$

Since the differential of a diagram is the sum of its various resolutions at its crossings, this is an isomorphism of complexes, or differential $\mathbb{Z}_2$-modules. (Note that the isomorphism holds even if (h,s,t) is not the H-data of any diagram! In this case both sides are zero.) It is also an isomorphism of differential graded algebras: multiplying two diagrams D, D' on $\mathcal{Z}$, and then cutting into fragments, yields the same result as cutting D, D' into fragments, and then multiplying the local diagrams — provided that it makes sense to cut all the diagrams D, D' and DD' into fragments, i.e. that they are all viable. We prove this now.

Lemma 2.4. *Let D, D' be viable augmented diagrams on $\mathcal{Z}$, with local diagrams D_P, D'_P on each fragment $\mathcal{Z}_P$. Then the product DD' is nonzero and viable if and only if each $D_P D'_P$ is nonzero. In this case $(DD')_P = D_P D'_P$.*

Thus, if under the isomorphism $\mathcal{A}(h,s,t) \cong \bigotimes_P \mathcal{A}_P(h_P, s_P, t_P)$ we have $D = \bigotimes_P D_P$ and $D' = \bigotimes_P D'_P$, then $DD' = \bigotimes_P D_P D'_P$.

Proof. Recall the description of multiplication of augmented diagrams in section 2.2. We examine the products DD' and $D_P D'_P$ on fragments $\mathcal{Z}_P$. If no strand of D or D' flies off the fragment $\mathcal{Z}_P$, then DD' (if nonzero) at P is clearly given by $D_P D'_P$. If strands of both D, D' fly off the top of $\mathcal{Z}_P$, then DD' is not viable, and $D_P D'_P = 0$; similarly if strands fly off the bottom. If a strand of D but not D' flies off the top (resp. bottom) of $\mathcal{Z}_P$, then as discussed in section 2.2, $D_P D'_P$ is well defined, with a single strand flying off the top (resp. bottom) of $\mathcal{Z}_P$, as also does DD' at P. The case where a strand of D' but not D flies off $\mathcal{Z}_P$ is similar. Gluing together these local results at each matched pair gives the desired result. □

Now for any chain complexes A, B over $\mathbb{Z}_2$ we have $H(A \otimes B) \cong H(A) \otimes H(B)$. (See e.g. [22, sec. 3.7] or [3, thm. V.2.1].) Thus the homology $H(\mathcal{A}(h,s,t))$ is isomorphic to the tensor product of the $H(\mathcal{A}_P(h_P, s_P, t_P))$, and in fact this isomorphism preserves all gradings. We use the shorthand

$$\mathcal{H}_P = H(\mathcal{A}_P), \quad \mathcal{H}_P(h_P, s_P, t_P) = H(\mathcal{A}_P(h_P, s_P, t_P))$$

We call $\mathcal{H}_P$ the *local homology at* P. So we have the following isomorphism, which we often use implicitly in the sequel.

Lemma 2.5. *For viable (h,s,t) there is an isomorphism of graded $\mathbb{Z}_2$-algebras, respecting H-grading and Maslov grading, induced by cutting diagrams into fragments.*

$$\mathcal{H}(h,s,t) \cong \bigotimes_{\text{matched pairs } P} \mathcal{H}_P(h_P, s_P, t_P)$$

□

Any fragment of any arc diagram is homeomorphic to any other, so all local strand algebras are isomorphic. Hence we may speak of *the* local arc diagram or *the* local strand algebra, without reference to any specific matched pair or arc diagram. They are abusively denoted $\mathcal{Z}_P$ and $\mathcal{A}_P$.

2.5 Terminology for local strand diagrams

We now develop systematic terminology to describe diagrams locally. Throughout this section, $P = \{p, q\}$ is a matched pair of an arc diagram $\mathcal{Z} = (\mathbf{Z}, \mathbf{a}, M)$, $h \in H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$ (which includes $H_1(\mathbf{Z}, \mathbf{a})$ as a subgroup), and D is a diagram with H-data (h, s, t).

Definition 2.6 (Occupation of places)**.**

(i) If h has multiplicity 0 *on the steps before and after p, then p is* unoccupied *by h.*

(ii) If h has multiplicity 1 *on the step before p, and* 0 *on the step after p, then p is* pre-half-occupied *by h.*

(iii) If h has multiplicity 0 *on the step before p, and* 1 *on the step after p, then p is* post-half-occupied *by h.*

(iv) If h is pre-half-occupied or post-half-occupied, then p is half-occupied *by h.*

(v) If h has multiplicity 1 *on both steps before and after p, then p is* fully occupied *by h.*

Although this definition applies to the H-grading h, we may equally apply it to a *diagram D*. We say that p is unoccupied (resp. pre-half-occupied, post-half-occupied, half-occupied, fully occupied) by D, if p is so occupied by its H-grading h.

Definition 2.7 (Occupation of pairs)**.**

(i) If both p, q are unoccupied by h, then P is unoccupied *by h.*

(ii) If p is half-occupied, and q is unoccupied by h, then P is one-half-occupied at p *by h. Accordingly as p is pre- or post-half-occupied, P is* pre-one-half-occupied *or* post-one-half-occupied.

(iii) If both p, q are half-occupied by h, then P is alternately occupied *by h.*

(iv) If p is fully occupied, and q is unoccupied by h, then P is once occupied at p.

(v) If p is fully occupied and q is half-occupied by h, then P is sesqui-occupied at p. *Accordingly as p is pre- or post-half-occupied, P is* pre-sesqui-occupied *or* post-sesqui-occupied.

(vi) If p, q are both fully occupied by h, then P is doubly occupied *by h.*

Again, we can extend this definition to diagrams: P is unoccupied (resp. sesqui-, pre/post-sesqui-, alternately , one-half-, pre/post-one-half-, doubly occupied) by the diagram D, if P is so occupied by its H-grading $h(D)$.

The following definition applies to any set of H-data (h, s, t), i.e. to $h \in H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$ and idempotents s, t.

Definition 2.8 (Idempotent terminology)**.**

(i) If $P \notin s$ and $P \notin t$ we say P is off-off *or* all-off *or* 00.

(ii) If $P \notin s$ and $P \in t$ we say P is off-on *or* 01.

(iii) If $P \in s$ and $P \notin t$ we say P is on-off *or* 10.

(iv) If $P \in s$ and $P \in t$ we say P is on-on *or* all-on *or* 11.

(We find this terminology awkward, hence we offer several equally awkward alternatives.) Again, this definition extends to diagrams. We say D is all-off, off-on, etc., at P if its H-data (h, s, t) is all-off, off-on, etc., at P.

We often combine this terminology, and simultaneously discuss occupation and on/off properties with H-data: for instance, we may say that a pair P is is all-on doubly occupied by (h, s, t), or equivalently that (h, s, t) is all-on doubly occupied at P. We hope the meaning is clear.

Any H-data (h, s, t), including the H-data of a diagram D, can be described completely by the terminology of occupation and on/off idempotents. The H-grading h described precisely by the occupation of the various matched pairs. The idempotents s, t are described precisely by the on/off data at each pair.

We can often deduce properties of D simply from its occupation of places, or its on/off/etc properties. For instance, if p is pre-half-occupied in a diagram D, then a strand of D must end at p, and no strand can begin at p; so D is off-on at P. Such reasoning is often useful.

2.6 Tightness of a diagram

We now define the tightness or twistedness of a diagram. Throughout this section D is a viable diagram on an arc diagram $\mathcal{Z}$, which lies in $\mathcal{A} = \mathcal{A}(\mathcal{Z})$. If D has no crossings then $\partial D = 0$ so D represents a homology class in $\mathcal{H}$.

Thus the following definition makes sense.

Definition 2.9 (Tightness of a diagram)**.** *A viable diagram $D \in \mathcal{A}$ is:*

(i) tight *if it has no crossings and is nonzero in homology;*

(ii) twisted *if it has no crossings but is zero in homology*

(iii) crossed *if it has crossings.*

The significance of the tightness of a diagram will become clearer as we proceed. Tight diagrams describe tight contact structures; twisted diagrams describe "minimally" overtwisted contact structures; crossed diagrams are more degenerate.

Lemma 2.10 (Local-global tightness of diagram)**.** *Let D be a viable diagram.*

(i) If D is tight iff for each matched pair P, D_P is tight.

(ii) If D is twisted iff for each matched pair P, D_P is tight or twisted, and at least one D_P is twisted.

(iii) If D is crossed iff for some matched pair P, D_P is crossed.

We could equivalently say that D is tight iff D is tight at all matched pairs; and that D is tight or twisted iff D is tight or twisted at all matched pairs. Thus tightness is a "local-to-global" property, as is "(tight or twisted)-ness".

This statement indicates an increasing order of degeneracy: tightness means tight everywhere; being twisted somewhere makes the diagram twisted; and then, being crossed somewhere makes the whole diagram crossed.

Proof. If D has a crossing, then it is local to some P (lemma 2.2), hence D_P is crossed; the converse is clear, proving (iii). So now assume D and all D_P are crossingless. By lemma 2.5 $\mathcal{H}(h, s, t) \cong \bigotimes_P \mathcal{H}_P(h_P, s_P, t_P)$, which is a tensor product of $\mathbb{Z}_2$ vector spaces. So D is nonzero in homology iff all D_P are nonzero in homology; (i) and (ii) follow. $\square$

2.7 Classification of local strand diagrams

The local arc diagram $\mathcal{Z}_P$ is very simple. There is only a small set of possible H-data; and given H-data on a fragment, the set of possible diagrams is even smaller. In this section we explicitly list out the possibilities.

There are only 4 steps in $\mathcal{Z}_P$, all exterior, and h determines which are covered by strands. The idempotents s and t determine whether a strand begins or ends at P — though whether the strand lies specifically at p or q (or at both, with dotted strands) may be ambiguous. In most cases, but not all, this is enough to determine the diagram completely.

In table 1 we draw all augmented diagrams on $\mathcal{Z}_P$ — or equivalently, all possible local diagrams of a viable diagram. For each set of H-data, described in terms of occupation and on/off terminology, there are no more than two viable local diagrams, and specifying tightness then determines a unique diagram, up to relabelling the twin places. This gives the following proposition, which will be proved in section 2.8.

Proposition 2.11 (Classification of viable local diagrams)**.** *Let D be a diagram on $\mathcal{Z}_P$. Then the H-data and tightness of D determines D up to relabelling twins, and D is as shown in table 1.*

Since table 1 shows all diagrams on $\mathcal{Z}_P$, together with H-data, proposition 2.11 is proved except for the classification into tight, twisted and crossed diagrams.

For those H-data admitting more than one diagram, a diagram can alternatively be specified by Maslov grading (rather than tightness), as shown in table 1. In all such cases, there is a crossed and a non-crossed diagram; the Maslov grading of the former is greater by 1 than the latter.

Hence, if we fix viable H-data (h, s, t) on an arc diagram $\mathcal{Z}$, using the isomorphism $\mathcal{A}(h, s, t) \cong \bigotimes_P \mathcal{A}_P(h_P, s_P, t_P)$ of section 2.4, then up to a constant, the Maslov grading of a diagram D is given by the number of matched pairs at which D is crossed.

2.8 Local algebras and homology

We now describe the local strand algebra $\mathcal{A}_P$, and its homology, explicitly, using the terminology of section 2.5 and the classification of section 2.7. Throughout this section, we consider diagrams on the local arc diagram $\mathcal{Z}_P$, where $P = \{p, q\}$. Let (h, s, t) be the H-data of a diagram on $\mathcal{Z}_P$.

Table 1 shows that (h, s, t) determines a diagram in all cases except two: when (h, s, t) is all-on doubly occupied or all-on once occupied. If P is all-on doubly occupied by (h, s, t), then 3 diagrams are possible; if P is all-on once occupied by (h, s, t), then 2 diagrams are possible.

These diagrams are important in the sequel, and so we name then. (Our choice of symbols may seem arbitrary, but there is method in the madness: c stands for "Crossed", g stands for "tiGht", and w stands for "tWisted").

Definition 2.12.

(i) If P is all-on doubly occupied by (h, s, t):

- *(a) c_P is the unique crossed diagram;*
- *(b) g_p is the unique crossingless diagram with strands beginning and ending at p*
- *(c) g_q is the unique crossingless diagram with a strand beginning and ending at q.*

(ii) If P is all-on once occupied at p by (h, s, t):

- *(a) c_p is the unique crossed diagram;*
- *(b) w_p is the unique crossingless diagram.*

(iii) For any other H-data, denote the unique diagram by u_P.

H-data		Tight	Twisted	Crossed
Unoccupied	all-off	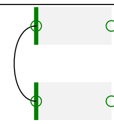 0		
	all-on	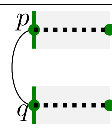 0		
One-half-occupied	Pre-	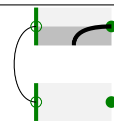 0		
	Post-	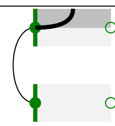 $-\frac{1}{2}$		
Alternately occupied	all-on	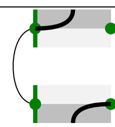 $-\frac{1}{2}$		
Once occupied	all-off	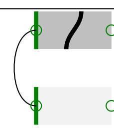 0		
	all-on		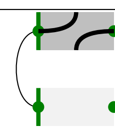 w_p -1	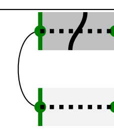 c_p 0
Sesqui-occupied	Pre-	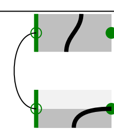 0		
	Post-	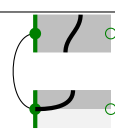 $-\frac{1}{2}$		
Doubly occupied	all-off	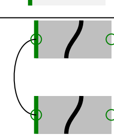 0		
	all-on	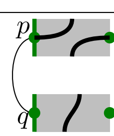 g_p -1 or 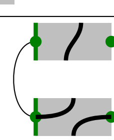 g_q -1		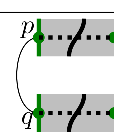 c_P 0

Table 1: Possible local diagrams, classified by H-data and tightness. Maslov indices are shown, and some diagrams are named.

Define chain complexes C''_P, C'_P, C_P by

$$\begin{array}{lll} C''_P: & 0 \longrightarrow \mathbb{Z}_2\langle c_P\rangle \longrightarrow \mathbb{Z}_2\langle g_p, g_q\rangle \longrightarrow 0 & \text{where } \partial c_P = g_p + g_q \text{ and } \partial g_p = \partial g_q = 0 \\ C'_P: & 0 \longrightarrow \mathbb{Z}_2\langle c_p\rangle \longrightarrow \mathbb{Z}_2\langle w_p\rangle \longrightarrow 0, & \text{where } \partial c_p = w_p \text{ and } \partial w_p = 0. \\ C_P: & 0 \longrightarrow \mathbb{Z}_2\langle u_P\rangle \longrightarrow 0. & \end{array}$$

As a chain complex, up to a shift giving the correct Maslov grading, each summand $\mathcal{A}_P(h,s,t)$ of $\mathcal{A}_P$ is isomorphic to C''_P, C'_P or C_P, accordingly as (h,s,t) is all-on doubly occupied, all-on once occupied, or anything else.

Calculating the homology of these complexes is straightforward.

- $H(C''_P) \cong \mathbb{Z}_2$, generated by the homology class of g_p or g_q (equal in homology since $\partial c_P = g_p + g_q$).
- $H(C'_P) = 0$.
- $H(C_P) \cong \mathbb{Z}_2$, generated by the homology class of u_P.

We can now complete the proof of proposition 2.11.

Proof of proposition 2.11. As observed above, it remains to classify diagrams by tightness. Crossed diagrams are clear. Our calculations now show that the only crossingless diagram which is zero in homology is w_p; all other crossingless diagrams are tight. □

2.9 Homology of strand algebras

We can now compute the homology of $\mathcal{A}(h,s,t)$ directly, for any viable H-data (h,s,t) on an arc diagram $\mathcal{Z}$.

Let (h,s,t) be the H-data of a viable diagram. From section 2.8 and the isomorphism of DGAs $\mathcal{A}(h,s,t) \cong \bigotimes_P \mathcal{A}_P(h_P,s_P,t_P)$ (section 2.4), we have the following isomorphism of chain complexes or differential $\mathbb{Z}_2$-modules:

$$\mathcal{A}(h,s,t) \cong \bigotimes_{P \text{ all-on doubly occupied}} C''_P \otimes \bigotimes_{P \text{ all-on once occupied}} C'_P \otimes \bigotimes_{\text{other } P} C_P$$

Homology decomposes as $\mathcal{H}(h,s,t) \cong \bigotimes_P \mathcal{H}_P(h_P,s_P,t_P)$ (lemma 2.5), and the calculations of section 2.8 now give

$$\mathcal{H}(h,s,t) \cong \begin{cases} 0 & \mathcal{Z} \text{ has an all-on once occupied pair,} \\ \mathbb{Z}_2 & \text{otherwise.} \end{cases}$$

Moreover, when there are no all-on once occupied pairs, $\mathcal{H}(h,s,t) \cong \mathbb{Z}_2$ is generated by the homology class of any crossingless diagram, which are therefore all unique in homology.

This calculation recovers the homology calculation of Lipshitz–Ozsváth–Thurston [13], for viable H-data, and extends it to augmented diagrams. Their calculation says that $\mathcal{H}(h,s,t)$ is nontrivial iff there exists a crossingless diagram D satisfying two conditions (stated above in section 2.1), which we can now translate into our terminology. Condition (i) is that D be viable. Condition (ii) is that if p is fully occupied and q is not fully occupied (i.e. P is once occupied at p, or sesqui-occupied), then P is not all-on. But table 1 shows that sesqui-occupied local diagrams are never all-on, so condition (ii) simply rules out crossingless all-on once occupied local diagrams; this is equivalent to ruling out all-on once occupied pairs.

Because of this calculation, the following definition makes sense.

Definition 2.13. *Let (h,s,t) be viable H-data.*

(i) The homology class *of (h,s,t), denoted $M_{h,s,t}$, is the unique nonzero homology class in $H(h,s,t)$, if it exists; otherwise $M_{h,s,t} = 0$.*

(ii) The local homology class *of (h,s,t) at P, denoted $M^P_{h,s,t}$, is the unique nonzero local homology class in $H_P(h,s,t)$, if it exists; otherwise $M^P_{h,s,t} = 0$.*

This definition applies to any viable (abstract) set of H-data on $\mathcal{Z}$, i.e. (h,s,t) where $h \in H_1(\mathbf{Z}, \mathbf{a}_{\pm\infty})$ covers each step at most once, and s,t are idempotents. Note that (h,s,t) need not be the H-data of a diagram; in this case $\mathcal{A}(h,s,t) = \mathcal{H}(h,s,t) = 0$.

The following statement encapsulates the above calculations and discussion.

Proposition 2.14. *Let (h,s,t) be viable H-data. Then precisely one of the following is true.*

(i) There is a tight diagram with H-data (h,s,t); (h,s,t) is the H-data of a diagram with no all-on once occupied pairs; $\mathcal{H}(h,s,t) \cong \mathbb{Z}_2$, generated by the homology class $M_{h,s,t}$ of any crossingless diagram with H-data (h,s,t).

(ii) There is a twisted diagram with H-data (h,s,t); (h,s,t) is the H-data of a diagram with an all-on once occupied pair; $\mathcal{H}(h,s,t) = 0$ but $\mathcal{A}(h,s,t) \neq 0$.

(iii) There is no diagram with H-data (h,s,t); $\mathcal{A}(h,s,t) = 0$.

Proof. Clearly if there is no diagram with H-data (h,s,t) then $\mathcal{A}(h,s,t) = 0$; so suppose (h,s,t) is the H-data of a diagram, and hence $\mathcal{A}(h,s,t) \neq 0$. In this case a crossingless diagram D with H-data (h,s,t) can always be drawn. If there is an all-on once occupied pair P, we calculated $\mathcal{H}(h,s,t) = 0$, so D is twisted. If there are no all-on once occupied pairs then we calculated $\mathcal{H}(h,s,t) \cong \mathbb{Z}_2$, so D is tight and generates $\mathcal{H}(h,s,t)$. □

Proposition 2.14 allows us to define the tightness of H-data as follows.

Definition 2.15 (Tightness of H-data)**.** *Let (h,s,t) be viable H-data on $\mathcal{Z}$.*

(i) (h,s,t) is tight *if there is a tight diagram with H-data (h,s,t). We denote the set of all viable tight H-data by $\mathbf{g}(\mathcal{Z})$.*

(ii) (h,s,t) is twisted *if there is a twisted diagram with H-data (h,s,t). We denote the set of all viable twisted H-data on $\mathcal{Z}$ by $\mathbf{w}(\mathcal{Z})$.*

(iii) Otherwise, (h,s,t) is singular.

When the arc diagram is understood we simply write $\mathbf{g}$ or $\mathbf{w}$ rather than $\mathbf{g}(\mathcal{Z})$ or $\mathbf{w}(\mathcal{Z})$.

Proposition 2.14 in fact gives several equivalent characteristations of tight, twisted or singular H-data. For instance, (h,s,t) is twisted iff it is the H-data of a diagram with an all-on once occupied pair.

Just as for tightness of diagrams, tightness of H-data obeys a “local-to-global” principle.

Lemma 2.16 (Local-global tightness of H-data)**.** *Again let (h,s,t) be viable H-data on $\mathcal{Z}$.*

(i) (h,s,t) is tight iff for all matched pairs P, (h_P, s_P, t_P) is tight .

(ii) (h,s,t) is twisted iff for each matched pair P, (h_P, s_P, t_P) is tight or twisted, and at least one (h_P, s_P, t_P) is twisted.

(iii) (h,s,t) is singular iff for some matched pair P, (h_P, s_P, t_P) is singular.

Proof. First, (h,s,t) is non-singular iff it is the H-data of a diagram iff each (h_P, s_P, t_P) is the H-data of a diagram, proving (iii). So assume (h,s,t), and hence all (h_P, s_P, t_P), are non-singular, hence tight or twisted. Then (h,s,t) is twisted iff there is an all-on once-occupied pair P, in which case this (h_P, s_P, t_P) is twisted. Otherwise, (h,s,t) and all (h_P, s_P, t_P) are tight. □

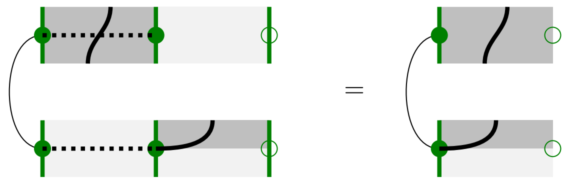

Figure 8: Sublimation.

2.10 Properties of twisted and crossed diagrams

We now consider some properties of twisted and crossed diagrams. We first consider crossed diagrams.

Lemma 2.17 (Products of crossed and crossingless diagrams)**.** *If two diagrams D_1 and D_2 are crossingless then $D_1 D_2$ is zero or crossingless. Hence the submodule of $\mathcal{A}$ generated by crossingless diagrams forms a subalgebra.*

Note this lemma applies to crossingless diagrams in general (not just viable ones). The product of two crossingless viable diagrams, even though nonzero and (by the lemma) crossingless, may not be viable.

Proof. If the product $D_1 D_2$ has a crossing, then one strand starts below and ends above another. The two strands must change their order either in D_1 or D_2 (even if they are dotted strands); so D_1 or D_2 has a crossing. This proves the first statement; applying it to linear combinations of diagrams proves the second. □

Note that the converse of the above lemma is not true: it is possible to have $D_1 D_2$ crossingless, but D_1 or D_2 crossed. In fact the product of a crossed diagram with another diagram may be tight. We call this phenomenon *sublimation*; see figure 8. As mentioned in the introduction (section 1.5), sublimation occurs repeatedly in higher A-infinity operations.

Turning to twisted diagrams, we observe that they are characterised by a specific local diagram w_p (of definition 2.12) at an all-on once occupied pair $P = \{p, q\}$.

Indeed, a diagram D is twisted iff each local diagram D_P is tight or twisted, and at least one D_P is twisted (by the "local-global" property of tightness, lemma 2.10); and by the classification of viable local diagrams in proposition 2.11 and table 1, the only twisted local diagram is w_p.

At the pair P, one place p must be occupied and its twin q unoccupied, so we can speak of a twisted diagram being twisted not just at a matched pair, but *at a specific place*.

Definition 2.18 (Diagram twisted at a place)**.** *A viable diagram D such that D_P is twisted at a pair $P = \{p, q\}$, with p occupied, is* twisted at the place *p.*

In terms of contact geometry, a contact structure which is "minimally overtwisted" at a particular square has two bypasses from adjacent squares; these adjacent squares lie around a particular corner of the square, as in figure 2.

Note that if D, D' are viable crossingless diagrams, at least one of which is twisted, then their product DD' (if nonzero and viable) is twisted: DD' is crossingless by lemma 2.17, and in homology at least one of D or D' is zero, hence DD' is zero in homology. This corresponds to the contact-geometric phenomenon that gluing two tight manifolds can yield an overtwisted manifold.

2.11 Diagrams representing a homology class

Let D be a tight diagram with H-data (h, s, t) on $\mathcal{Z}$, hence (definition 2.9) nonzero in homology, with homology class $M_{h,s,t}$ (definition 2.13), a generator of $\mathcal{H}(h, s, t)$ (proposition 2.14).

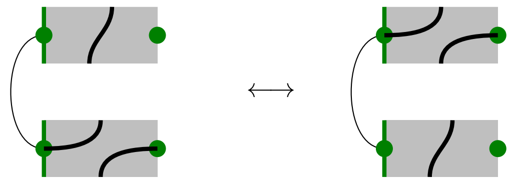

Figure 9: Switching strands.

While D determines $M_{h,s,t}$, the diagram D cannot always be recovered from $M_{h,s,t}$. There is ambiguity at all-on doubly occupied pairs, where the local diagrams g_p, g_q (definition 2.12, or see table 1) are both tight. We say g_p and g_q are related by *strand switching* at the pair $\{p, q\}$; see figure 9.

In fact, $\mathcal{H}(h, s, t)$ is isomorphic, as an abelian group, to the abelian group generated by tight diagrams with H-data (h, s, t), modulo the subgroup generated by sums of diagrams related by strand switching (these are the boundaries: $\partial c_P = g_p + g_q$).

We now determine the number of tight diagrams with given viable H-data (h, s, t). There are none when (h, s, t) is twisted or singular, since then $\mathcal{H}(h, s, t) = 0$ (proposition 2.14). When (h, s, t) is tight, at each pair we have a unique choice of tight local diagram, except at all-on doubly occupied pairs, where we have two tight local diagrams, related by strand switching. For any such choices we can glue local diagrams together into a tight diagram, giving the following statement.

Lemma 2.19. *Let (h, s, t) be tight viable H-data on the arc diagram $\mathcal{Z}$. Let L be the number of pairs all-on doubly occupied by (h, s, t). Then there are precisely 2^L tight diagrams with H-data (h, s, t). Any two such diagrams are related by a sequence of switchings of strands.* □

These 2^L tight diagrams are precisely the diagrams in $\mathcal{A}$ representing their common homology class $M_{h,s,t}$ in $\mathcal{H}(h, s, t)$.

2.12 Dimensions of strand algebras

It will be useful to know the dimension of $\mathcal{A}(h, s, t)$ as a $\mathbb{Z}_2$ vector space, as well as its subspaces of cycles and boundaries.

Throughout this section let $\mathcal{Z}$ be an arc diagram and (h, s, t) be viable non-singular H-data on $\mathcal{Z}$, with L all-on doubly occupied pairs, and N all-on once occupied pairs. Dimension always refer to dimensions of $\mathbb{Z}_2$ vector spaces.

Lemma 2.20. *The dimension of $\mathcal{A}(h, s, t)$ is $3^L 2^N$.*

Proof. A basis is given by the diagrams with H-data (h, s, t); these may be specified locally at each matched pair P. Table 1 shows that if P is all-on once occupied, then there are 3 choices at P; if P is all-on doubly occupied, then there are 2 choices; in any other case there is a unique choice. □

Now we refine $\mathcal{A}(h, s, t)$ by Maslov grading. As discussed in section 2.7, with H-data fixed, the Maslov grading of a diagram D is given, up to a constant, by the number of matched pairs at which D is crossed.

We denote by $\mathcal{A}_n(h, s, t)$ the $\mathbb{Z}_2$ vector subspace of $\mathcal{A}(h, s, t)$ spanned by diagrams with crossings at precisely n matched pairs. Then $\mathcal{A}(h, s, t) = \bigoplus_n \mathcal{A}_n(h, s, t)$ and this is the decomposition by Maslov grading.

Lemma 2.21. *The dimension of $\mathcal{A}_n(h, s, t)$ is given by*

$$\dim \mathcal{A}_n(h, s, t) = \sum_i 2^{L-i} \binom{L}{i} \binom{N}{n-i} = \sum_k \binom{L}{k} \binom{N+k}{n}.$$

For integers a, b, we regard $\binom{a}{b}$ as zero when $b < 0$ or $b > a$. The summations are over all integers i or j; each has only finitely many nonzero terms.

Proof. Let i be the number of crossed all-on doubly occupied pairs. Then there are $\binom{L}{i}$ choices for the all-on doubly occupied pairs which are to be crossed. At each crossed all-on doubly occupied pair there is a unique diagram, but at the $L-i$ non-crossed all-on doubly occupied pairs, there are 2 possible diagrams, giving 2^{L-i} choices at non-crossed all-on doubly occupied pairs. The other $n-i$ pairs with crossings must be all-on once occupied pairs, and there are $\binom{N}{n-i}$ choices for which all-on once occupied pairs will have crossings. Once these are chosen, all local diagrams are uniquely determined, and all such local diagrams glue into a diagram in $\mathcal{A}_n(h,s,t)$. Summing over all i gives the first equality.

For the second equality, fix a reference diagram D_0 with H-data (h,s,t) and no crossings. (Such a diagram always exists locally, by table 1, and the local diagrams glue together.) Consider a diagram D in $\mathcal{A}_n(h,s,t)$ and let k be the number of all-on doubly occupied pairs at which D and D_0 differ. There are $\binom{L}{k}$ ways in which we can choose these k pairs. Now the n pairs with crossings must come from the k all-on doubly occupied pairs just chosen, together with the N all-on once occupied pairs. There are $\binom{N+k}{n}$ ways to choose which of these $N+k$ pairs will be crossed. We now observe that once such choices are made, the diagram D is uniquely determined. For at all-on doubly occupied pairs, the diagram either coincides with D_0 (and has no crossing), or differs from D_0, and if it differs from D_0 then we have chosen it to be crossed or not; these choices correspond to the 3 possible local diagrams at an all-on doubly occupied pair. At all-on once occupied pairs, the diagram either coincides with D_0 (and is uncrossed), or is selected to be crossed; these choices correspond to the 2 possible local diagrams at all-on once occupied pairs. At any other pair there is a unique local diagram. All such choices determine a diagram in $\mathcal{A}_n(h,s,t)$, and all such diagrams are constructed uniquely by such choices. Thus $\dim \mathcal{A}_n(h,s,t) = \sum_k \binom{L}{k}\binom{N+k}{n}$. □

We remark that it is also possible to prove directly that the two summations are equal. Combining lemmas 2.20 and 2.21, we have

$$\sum_n \sum_i 2^{L-i} \binom{L}{i} \binom{N}{n-i} = \sum_n \sum_k \binom{L}{k} \binom{N+k}{n} = 3^L 2^N. \tag{5}$$

Next, we consider the dimension of the spaces of *boundaries* and *cycles* in $\mathcal{A}_n(h,s,t)$. Let $B_n(h,s,t)$ and $Z_n(h,s,t)$ respectively denote the $\mathbb{Z}_2$ vector subspaces of $\mathcal{A}_n(h,s,t)$ generated by boundaries and cycles. In other words, for any $n \geq 0$, $\partial\colon \mathcal{A}_{n+1} \longrightarrow \mathcal{A}_n$ has image B_n and kernel Z_{n+1}. When (h,s,t) is twisted, $\mathcal{A}(h,s,t)$ has trivial homology, so $B_n(h,s,t) = Z_n(h,s,t)$ for all n.

Lemma 2.22. *The dimension of $B_n(h,s,t)$ is given by*

$$\dim B_n(h,s,t) = \sum_i 2^{L-i} \binom{L}{i} \binom{N-1}{n-i} = \sum_k \binom{L}{k} \binom{N+k-1}{n}$$

Again, summations are over all integers.

Proof. We know that nonzero homology only arises from diagrams with no crossings, i.e. with $n=0$. Hence for all $n \geq 0$ we have $Z_{n+1} = B_{n+1}$, so $\dim B_n = \dim \mathcal{A}_{n+1} - \dim B_{n+1}$. Applying this repeatedly we obtain

$$\dim B_n = \dim \mathcal{A}_{n+1} - \dim \mathcal{A}_{n+2} + \dim \mathcal{A}_{n+3} - \cdots = \sum_{k=1}^{\infty} (-1)^{k+1} \dim \mathcal{A}_{n+k}.$$

From lemma 2.21 we have $\dim \mathcal{A}_{n+k} = \sum_i 2^{L-i} \binom{L}{i} \binom{N}{n+K-i}$, and hence the above is equal to

$$\begin{aligned}\sum_{k=1}^{\infty}(-1)^{k+1} \sum_i 2^{L-i} \binom{L}{i} \binom{N}{n+K-i} &= \sum_i 2^{L-i} \binom{L}{i} \sum_{k=1}^{\infty} (-1)^{k+1} \binom{N}{n+k-i} \\ &= \sum_i 2^{L-i} \binom{L}{i} \binom{N-1}{n-i},\end{aligned}$$

where we used the fact that for any integers a, b, $\sum_{k=1}^{\infty} (-1)^{k+1} \binom{a}{b+k} = \binom{a}{b+1} - \binom{a}{b+2} + \cdots = \binom{a-1}{b}$ (easily established, for instance, by induction). We have now proved the first claimed equality; the second follows from the identity of lemma 2.21. □

2.13 An ideal in the strand algebra

We now introduce an ideal $\mathcal{F}$ in $\mathcal{A} = \mathcal{A}(\mathcal{Z})$, which will be important in the sequel.

Definition 2.23. *The $\mathbb{Z}_2$-submodule of $\mathcal{A}$ generated by diagrams which are not viable, or have at least one doubly occupied crossed pair, is denoted $\mathcal{F}$.*

In other words, a diagram D lies in $\mathcal{F}$ iff it has some step covered by two or more strands, or has a doubly occupied crossed pair.

Lemma 2.24. *$\mathcal{F}$ is a two-sided ideal of $\mathcal{A}$.*

Proof. First we claim that if D, D' are diagrams where D is not viable, then DD' and $D'D$ are zero or non-viable. For D then has some step covered by two or more strands, so DD' is either zero, or has a step covered by two or more strands, hence is not viable; similarly for $D'D$.

Now suppose D lies in $\mathcal{F}$ and D' is another diagram. If D is not viable, then the previous paragraph shows $DD', D'D \in \mathcal{F}$. So we may assume D is viable and has a doubly occupied crossed pair P. After multiplication on the right (resp. left) by D' the result may become non-viable, in which case DD' (resp. $D'D$) is in $\mathcal{F}$. If the result is viable, then DD' (resp. $D'D$) still has a doubly occupied crossed pair at P. □

The quotient $\mathcal{A}/\mathcal{F}$ is generated as a $\mathbb{Z}_2$-module by viable diagrams without crossed doubly occupied pairs. Indeed, an element of $\mathcal{A}/\mathcal{F}$ is represented uniquely by a sum of viable diagrams without crossed doubly occupied pairs. Products can then be taken as in $\mathcal{A}$, unless the result is non-viable or has a crossed doubly occupied pair, in which case the result is zero.

The decomposition $\mathcal{A} \cong \bigoplus_{h,s,t} \mathcal{A}(h, s, t)$ also descends to the quotient $\mathcal{A}/\mathcal{F}$. Hence we may make the following definitions.

Definition 2.25.

(i) The $\mathbb{Z}_2$-algebra $\overline{\mathcal{A}}$ is the quotient algebra $\mathcal{A}/\mathcal{F}$.

(ii) The $\mathbb{Z}_2$ vector space $\overline{\mathcal{A}}(h, s, t)$ is the (h, s, t) graded summand of $\overline{\mathcal{A}}$.

(iii) For $x \in \mathcal{A}$, we denote by $\overline{x}$ its image in $\overline{\mathcal{A}}$ under the quotient map $\mathcal{A} \longrightarrow \overline{\mathcal{A}}$.

(iv) For a homomorphism f with image in $\mathcal{A}$, we denote by $\overline{f}$ the homomorphism obtained by composing f with the quotient map $\mathcal{A} \longrightarrow \overline{\mathcal{A}}$.

(v) The standard form representative *$x \in \mathcal{A}$ of an $\overline{x} \in \overline{\mathcal{A}}$ is the sum of viable diagrams without crossed doubly occupied pairs whose image under the quotient map $\mathcal{A} \longrightarrow \overline{\mathcal{A}}$ is $\overline{x}$.*

The quotient $\overline{\mathcal{A}}$ is useful for our needs. Non-viable diagrams cannot contribute to homology, and although some crossed diagrams can be "salvaged" into tight diagrams (thus contributing to homology) via sublimation, sublimation does not apply to crossed doubly occupied pairs. Thus $\overline{\mathcal{A}}$ is generated by diagrams which are "salvageable" in this sense.

3 Cycle selection and creation operators

3.1 Cycle selection homomorphisms

Throughout this section fix an arc diagram $\mathcal{Z}$.

As discussed in section 1.3 that construction of an A-infinity structure on $\mathcal{H}$ begins from a map $f_1 \colon \mathcal{H} \longrightarrow \mathcal{A}$ as follows.

Definition 3.1. *A* cycle selection map *is a $\mathbb{Z}_2$-module homomorphism $f \colon \mathcal{H} \longrightarrow \mathcal{A}$ which preserves H-grading and Maslov grading, and sends each homology class $x \in \mathcal{H}$ to a cycle in $\mathcal{A}$ which represents x.*

Defining a cycle selection map requires finding diagrams representing each homology class — as discussed in section 2.11, but now for all homology classes.

The following constraint is a natural one to make, avoiding a proliferation of diagrams.

Definition 3.2 (Diagrammatically simple homomorphisms)**.**

(i) A $\mathbb{Z}_2$-module homomorphism $f \colon \mathcal{A} \longrightarrow \mathcal{A}$ is diagrammatically simple *if each for each diagram D, $f(D)$ is zero, or a single diagram.*

(ii) A $\mathbb{Z}_2$-module homomorphism $f \colon \mathcal{H} \longrightarrow \mathcal{A}$ is diagrammatically simple *if for all $M \in H$ that can be represented by a single diagram, $f(M)$ is zero, or a single diagram.*

Recall (proposition 2.14) that an H-summand $\mathcal{H}(h,s,t)$ of $\mathcal{H}$ is nonzero precisely when (h,s,t) is tight, in which case $\mathcal{H}(h,s,t) \cong \mathbb{Z}_2$, generated by $M_{h,s,t}$ (definition 2.13). To define a diagrammatically simple $f \colon \mathcal{H} \longrightarrow \mathcal{A}$, we simply need to select, for each tight H-data $(h,s,t) \in \mathbf{g}$ (definition 2.15), a tight diagram with H-data (h,s,t). For such (h,s,t), by lemma 2.19 there are precisely 2^L diagrams representing $M_{h,s,t}$, where L is the number of pairs all-on doubly occupied by (h,s,t). Making any of these 2^L choices for $f_1(M_{h,s,t})$, and repeating for each $(h,s,t) \in \mathbf{g}$, results in a diagrammatically simple cycle selection map; and all diagrammatically simple cycle selection maps are of this form.

We will systematically describe all diagrammatically simple cycle selection homomorphisms. To this end we make some definitions.

As is standard in set theory, given a set S whose elements are sets, a *set choice function* for S assigns to each $x \in S$ an element of x. The set of set choice functions for S is naturally in bijection with the direct product of S (i.e. the direct product of the elements of S), denoted $\prod S$, and we regard an element of $\prod S$ as a choice function for S. If S is empty, we regard S as having a unique choice function, which is the null function.

Definition 3.3 (Pair choice function)**.** *For $(h,s,t) \in \mathbf{g}$, let $\mathbf{P}_{h,s,t}$ be the set of all-on doubly occupied pairs of (h,s,t). A* pair choice function *for (h,s,t) is a set choice function for $\mathbf{P}_{h,s,t}$.*

Note $\mathbf{P}_{h,s,t}$ is a set of sets (each with two elements), so this definition makes sense. A pair choice function for (h,s,t) assigns to each doubly occupied all-on pair $P = \{p,q\}$ of (h,s,t) one of its places p or q. If $|\mathbf{P}_{h,s,t}| = L$ then the number of pair choice functions for $(h,s,t) \in \mathbf{g}$ is $|\prod \mathbf{P}_{h,s,t}| = 2^{|\mathbf{P}_{h,s,t}|} = 2^L$. If $L = 0$, then (h,s,t) has a unique (null) pair choice function.

Given a pair choice function $\mathcal{C}_{h,s,t}$ for $(h,s,t) \in \mathbf{g}$, we draw a tight diagram with H-data (h,s,t), denoted $D_{\mathcal{C}_{h,s,t}}$, as follows. At a matched pair P which is not all-on doubly occupied, we draw the unique tight local diagram with H-data (h_P, s_P, t_P). At an all-on doubly occupied matched pair $P = \{p,q\}$, $\mathcal{C}_{h,s,t}$ selects one of the places p or q; we denote it $\mathcal{C}_{h,s,t}(P)$. There are two possible tight local diagrams g_p, g_q (definition 2.12) with the required H-data at P, and we draw $g_{\mathcal{C}_{h,s,t}(P)}$, the diagram with strands beginning and ending at $\mathcal{C}_{h,s,t}(P)$. Putting these local diagrams together gives $D_{\mathcal{C}_{h,s,t}}$.

Definition 3.4 (Cycle choice function)**.** *A* cycle choice function *for $\mathcal{Z}$ is a function which assigns to each $(h,s,t) \in \mathbf{g}(\mathcal{Z})$ a pair choice function for (h,s,t).*

A cycle choice function can be regarded as an element of the set $\prod_{(h,s,t)\in\mathbf{g}}\prod \mathbf{P}_{h,s,t}$.

If $\mathcal{C}$ is a cycle choice function, then we write $\mathcal{C}(h,s,t)$ for the pair choice function assigned to $(h,s,t)\in\mathbf{g}$; this pair choice function for (h,s,t) then determines a tight diagram $D_{\mathcal{C}(h,s,t)}$ with H-data (h,s,t) as described above.

Now a cycle choice function $\mathcal{C}$ determines a map $f^{\mathcal{C}}\colon \mathcal{H}\longrightarrow\mathcal{A}$ as follows. For $(h,s,t)\in\mathbf{g}$, $\mathcal{H}(h,s,t)$ contains a single nonzero homology class $M_{h,s,t}$, and we define $f^{\mathcal{C}}(M_{h,s,t}) = D_{\mathcal{C}(h,s,t)}$. Combining such maps over all $(h,s,t)\in\mathbf{g}$ yields a diagrammatically simple cycle selection map $f^{\mathcal{C}}\colon \mathcal{H}\longrightarrow\mathcal{A}$. It is not difficult to see that all diagrammatically simple cycle selection maps are of this form, giving the following statement.

Lemma 3.5. *Let $f\colon\mathcal{H}\longrightarrow\mathcal{A}$ be a cycle selection homomorphism. Then f is diagrammatically simple iff $f = f^{\mathcal{C}}$ for some cycle choice function $\mathcal{C}$.* □

In other words, there is a bijective correspondence between diagrammatically simple cycle selection maps, and cycle choice functions.

We also consider a cycle selection map f in general (not necessarily diagrammatically simple). For $(h,s,t)\in\mathbf{g}$ we have $\mathcal{H}(h,s,t)$ generated by $M_{h,s,t}$. Then $f(M_{h,s,t})$ need not be a single diagram, but must be a sum of diagrams representing $M_{h,s,t}$, all of the same H-grading and Maslov grading, hence tight diagrams representing $M_{h,s,t}$. Moreover, as $f(M_{h,s,t})$ represents $M_{h,s,t}$, $f(M_{h,s,t})$ must be the sum of an *odd* number of distinct diagrams. Conversely, if for each $(h,s,t)\in\mathbf{g}$ we define $f(M_{h,s,t})$ to be the sum of an odd number of distinct tight diagrams representing $M_{h,s,t}$, we obtain a cycle selection homomorphism.

3.2 Differences in cycle selection

We now show how the different choices available in cycle selection are related to the ideal $\mathcal{F}$. Fix an arc diagram $\mathcal{Z}$ throughout this section.

Lemma 3.6. *Let $D_1,\ldots,D_{2n}\in\mathcal{A}$ be an even number of tight diagrams, each representing a homology class $M\in\mathcal{H}$. Then we have the following.*

(i) $D_1+\cdots+D_{2n}\in\partial\mathcal{F}$

(ii) *If $g\in\mathcal{A}$ is an element homogeneous in Maslov grading and H-data, satisfying $\partial g = D_1+\cdots+D_{2n}$, then $g\in\mathcal{F}$.*

Proof. We first prove (i) when $n=1$, so take diagrams D, D' which differ by switching strands at some all-on doubly occupied pairs $P_1,\ldots,P_k$ (lemma 2.19). We proceed by induction on k. When $k=1$, let F_1 be the diagram all-on doubly occupied crossed at P_1, and equal to D and D' elsewhere.

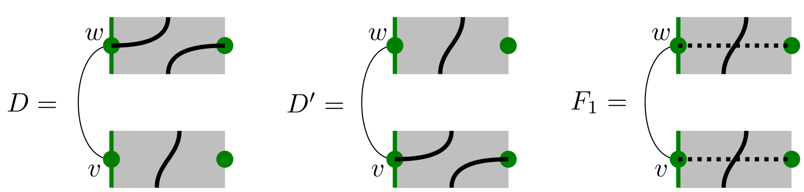

Then F_1 is viable, crossed at P_1, hence lies in $\mathcal{F}$, and is tight at all other pairs; so $\partial F_1 = D + D'$, as desired.

Now consider D, D' differing at k pairs. Let D'' be obtained from D by switching strands at P_1. By induction

$$D + D'' = \partial F_1, \quad D'' + D' = \partial\left(F_2+\cdots+F_k\right)$$

for some viable diagrams $F_1,\ldots,F_k$, with each F_i crossed at P_i (hence in $\mathcal{F}$) and tight elsewhere.Thus $D+D' = \partial(F_1+\cdots+F_k)$, proving (i) when $n=1$. For general n, simply split the diagrams $D_1,\ldots,D_{2n}$ into pairs and apply the $n=1$ case.

If g is homogeneous in Maslov grading and H-data and $\partial g = D_1 + \cdots + D_{2n}$, then every diagram G in g has the same Maslov grading and tight H-data as each F_i considered above. Thus G has precisely one pair with a crossed local diagram. From table 1 we see that crossings can only occur in viable diagrams at pairs which are all-on once occupied or all-on doubly occupied. But having the same viable tight H-data as the F_i, G has no all-on once occupied pairs (proposition 2.14). So G has a crossing at an all-on doubly occupied pair. Thus $G \in \mathcal{F}$, and hence $g \in \mathcal{F}$. $\square$

3.3 Creation operators

Let (h, s, t) be viable H-data on an arc diagram $\mathcal{Z}$, which is all-on once occupied at a pair $P = \{p, q\}$, occupied at p. We saw in section 2.8 that $\mathcal{A}_P(h_P, s_P, t_P)$ is, as a chain complex, given by C'_P (definition 2.12), which has trivial homology:

$$0 \longrightarrow \mathbb{Z}_2\langle c_p \rangle \xrightarrow{\partial} \mathbb{Z}_2\langle w_p \rangle \longrightarrow 0, \quad \text{where } \partial c_p = w_p \text{ and } \partial w_p = 0.$$

Here c_p is the unique local crossed diagram, w_p is the unique local twisted diagram.

There is a chain homotopy $A^* \colon C'_P \longrightarrow C'_P$ from the identity to 0; in fact, as C'_P is so simple, there is a unique such A^*, given as follows.

Definition 3.7 (Local creation operator)**.** *The* creation operator $A^* \colon C'_P \longrightarrow C'_P$ *is the $\mathbb{Z}_2$-module homomorphism given by $A^*(w_p) = c_p$ and $A^*(c_p) = 0$.*

In other words, A^* inserts a crossing, as in figure 1. The name A^* references creation operators in physics. We have $A^*\partial + \partial A^* = 1$ a "Heisenberg relation" or a chain homotopy from the identity to 0.

Still assuming (h, s, t) is viable H-data on $\mathcal{Z}$, all-on once occupied at P, consider $\mathcal{A}(h, s, t) \cong \bigotimes_{P'} \mathcal{A}_P(h_{P'}, s_{P'}, t_{P'})$ (section 2.4). We can rewrite this as

$$\mathcal{A}(h, s, t) \cong \mathcal{A}_P(h_P, s_P, t_P) \otimes \bigotimes_{P' \neq P} \mathcal{A}_{P'}(h_{P'}, s_{P'}, t_{P'}). \tag{6}$$

The first factor contains local diagrams at P, and is isomorphic to C'_P; the second factor contains local diagrams everywhere else. A diagram $D \in \mathcal{A}(h, s, t)$ is then written as $x \otimes y$, where $x \in \mathcal{A}_P(h_P, s_P, t_P) \cong C'_P$ and $y \in \bigotimes_{P' \neq P} \mathcal{A}_{P'}(h_{P'}, s_{P'}, t_{P'})$.

Definition 3.8 (Creation operator)**.** *Let P be an all-on once occupied pair of the viable H-data (h, s, t). The* creation operator $A^*_P \colon \mathcal{A}(h, s, t) \longrightarrow \mathcal{A}(h, s, t)$ *is given by $A^*_P = A^* \otimes 1$, in the tensor decomposition (6) above.*

In other words, A^*_P inserts a crossing at P. Clearly A^*_P is diagrammatically simple (definition 3.2). Note that if the diagram D lies in $\mathcal{F}$ (i.e. has a crossed doubly occupied pair: definition 2.23), then $A^*_P D \in \mathcal{F}$ also. So A^*_P descends to a map $\overline{A}^*_P \colon \overline{\mathcal{A}}(h, s, t) \longrightarrow \overline{\mathcal{A}}(h, s, t)$.

Lemma 3.9. *With P and (h, s, t) as above, $A^*_P \partial + \partial A^*_P = 1$ on $\mathcal{A}(h, s, t)$.*

Proof. Take a diagram in $\mathcal{A}(h, s, t)$ and write it as $x \otimes y$ according to the decomposition (6) above, so $x = c_p$ or w_p. Recalling that $\partial c_p = w_p$, $\partial w_p = 0$, $A^* w_p = c_p$, $A^* c_p = 0$, we have

$$\begin{aligned} (A^*_P \partial + \partial A^*_P)(w_p \otimes y) &= A^*_P (w_p \otimes \partial y) + \partial (c_p \otimes y) = c_p \otimes \partial y + w_p \otimes y + c_p \otimes \partial y = w_p \otimes y, \\ (A^*_P \partial + \partial A^*_P)(c_P \otimes y) &= A^*_P (w_p \otimes y + c_p \otimes \partial y) + 0 = c_p \otimes y. \end{aligned}$$

$\square$

This chain homotopy from the identity to 0 shows directly that $\mathcal{H}(h, s, t) = 0$ when there is an all-on once occupied pair (proposition 2.14).

In fact, creation operators are the *only* way to obtain a *diagrammatically simple* (definition 3.2) chain homotopy to the identity on a summand $\mathcal{A}(h, s, t)$

Lemma 3.10. *Suppose (h, s, t) is viable and non-singular, and $\int : \mathcal{A}(h, s, t) \longrightarrow \mathcal{A}(h, s, t)$ is a diagrammatically simple $\mathbb{Z}_2$-module homomorphism which has pure Maslov degree, satisfying*

$$\int \partial + \partial \int = 1.$$

Then $\int = A_P^*$, for some all-on once occupied matched pair P of (h, s, t).

Proof. The existence of the chain homotopy $\int$ implies $\mathcal{H}(h, s, t) = 0$; being non-singular then (h, s, t) is twisted, so there is an all-on once occupied pair. With (h, s, t) fixed, Maslov degree is given, up to a constant, by the number of pairs at which a diagram is crossed (section 2.7). Since $\int \partial + \partial \int = 1$ and ∂ has Maslov degree -1, $\int$ has Maslov degree 1.

We will use the decomposition $\mathcal{A}(h, s, t) \cong \bigotimes_P \mathcal{A}_P(h_P, s_P, t_P)$, noting (section 2.8) that each $\mathcal{A}_P(h_P, s_P, t_P)$ isomorphic (as a chain complex) to C_P, C'_P or C''_P (definition 2.12).

Take an arbitrary crossingless diagram D_0 with H-data (h, s, t). Then D_0 is twisted at each all-on once occupied pair. As $\int$ is diagrammatically simple, $\int D_0$ is a single diagram or 0. Since $\partial D_0 = 0$ and $\int \partial + \partial \int = 1$ we obtain $\partial \int D_0 = D_0$. Thus $\int D_0$ is a single diagram whose differential is D_0. The only such diagrams are those obtained from D_0 by inserting a crossing at an all-on once occupied pair $P = \{p, p'\}$ (say occupied at p). That is, $\int D_0 = A_P^* D_0$.

We claim that for *any* diagram D with H-data (h, s, t), $\int D = A_P^* D$. The proof is by induction on the number k of pairs at which D is crossed (i.e., up to a constant, Maslov grading).

Suppose D, D' are distinct crossingless diagrams with H-data (h, s, t), which differ by switching strands at a single all-on doubly occupied pair $Q = \{q, q'\}$. The argument above shows that $\int D = A_R^* D$ for some all-on once occupied pair $R = \{r, r'\}$ (occupied at r), and similarly that $\int D' = A_V^* D'$ for some all-on once occupied pair $V = \{v, v'\}$ (occupied at v). We claim $R = V$. To see why, suppose $R \neq V$ and consider $\mathcal{A}(h, s, t)$ as a tensor product. We may write

$$D = g_q \otimes w_r \otimes w_v \otimes z, \quad D' = g_{q'} \otimes w_r \otimes w_v \otimes z,$$
$$\int D = g_q \otimes c_r \otimes w_v \otimes z, \quad \int D' = g_{q'} \otimes w_r \otimes c_v \otimes z,$$

where the first tensor factor is C''_Q, the second is C'_R, the third is C'_V, and the last is the tensor product given by all other matched pairs. Now we consider the diagram $E = c_Q \otimes w_r \otimes w_v \otimes z$, obtained from either D or D' by inserting a pair of crossing dotted horizontal strands at Q. We compute

$$\int \partial E = \int (D + D') = g_q \otimes c_r \otimes w_v \otimes z + g_{q'} \otimes w_r \otimes c_v \otimes z,$$

and hence $\int E$ is a single diagram (by diagrammatic simplicity) whose differential is

$$\partial \int E = \left(\int \partial + 1 \right) E = g_q \otimes c_r \otimes w_v \otimes z + g_{q'} \otimes w_r \otimes c_v \otimes z + c_Q \otimes w_r \otimes w_v \otimes z.$$

The three diagrams on the right respectively have crossings at R, V and Q. Hence $\int E$ must have crossings at R, V and Q, contradicting the fact that $\int$ has Maslov degree 1. We conclude that $R = V$.

Thus, if D, D' are two crossingless diagrams with H-data (h, s, t), which differ by switching strands at a single all-on doubly occupied pair, then $\int D$ and $\int D'$ are both given by applying a creation operator A_P^* at the same matched pair P. Now any two crossingless diagrams with H-data (h, s, t) can be related by switching strands at some all-on doubly occupied pairs. Repeatedly applying this fact, we see that for any crossingless diagram D with H-data (h, s, t), $\int D = A_P^* D$. This proves the result when $k = 0$.

Now take a $k \geq 0$ and suppose that, for all diagrams D with H-data (h, s, t) and crossings at $\leq k$ pairs, $\int D = A_P^* D$. Consider a diagram D with H-data (h, s, t), crossed at $k + 1$ pairs. Then $D = w_p \otimes x$ or $c_p \otimes x$, where the first tensor factor refers to the complex C'_P for the pair P, and the second factor refers to the tensor product given by all other matched pairs.

If $D = w_p \otimes x$ then $\partial D = w_p \otimes \partial x$, which contains diagrams crossed at k pairs. By induction then

$$\int \partial D = A_P^* \partial D = A_P^* \left(w_p \otimes \partial x \right) = c_p \otimes \partial x.$$

It follows that $\int D$ is a single diagram (by diagrammatic simplicity) whose differential is

$$\partial \int D = \int \partial D + D = c_p \otimes \partial x + w_p \otimes x.$$

There is only one such diagram, namely $c_p \otimes x$. Thus $\int D = c_p \otimes x = A_P^* \left(w_p \otimes x \right) = A_P^* D$.

If $D = c_p \otimes x$ then $\partial D = w_p \otimes x + c_p \otimes \partial x$ and so by induction $\int \partial D = A_P^* \partial D = c_P \otimes x = D$. We then have $\partial \int D = \int \partial D + D = 0$, so $\int D$ is a single diagram crossed at $k+1 \geq 1$ pairs, or zero, whose differential is zero. Thus $\int D = 0 = A_P^* D$.

Thus, in any case, $\int D = A_P^* D$. By induction then $\int = A_P^*$. □

It is not difficult to see that, if we drop the requirement that $\int$ be diagrammatically simple, the result no longer holds: there are many $\mathbb{Z}_2$-module homomorphisms $\mathcal{A}(h,s,t) \longrightarrow \mathcal{A}(h,s,t)$ of pure Maslov degree satisfying $\partial \int + \int \partial = 1$, which are not creation operators. (For instance, any sum of an odd number of creation operators gives such an $\int$.)

3.4 Inverting the differential

The following straightforward lemma shows how a creation operator A_P^* "integrates", i.e. finds partial inverses of the differential ∂ (hence the notation $\int$ of section 3.3), as is required in constructing an A_∞ structure (section 1.3).

Lemma 3.11. *Suppose the viable H-data (h,s,t) contains an all-on once occupied pair P. If $x \in \mathcal{A}(h,s,t)$ is a cycle, then $x = \partial A_P^* x$.*

Proof. As x is a cycle, $\partial x = 0$. Hence $x = (A_P^* \partial + \partial A_P^*) x = \partial A_P^* x$. □

Recall from section 2.12 the decomposition $\mathcal{A}(h,s,t) = \bigoplus_n \mathcal{A}_n(h,s,t)$ over Maslov grading, where $\mathcal{A}_n(h,s,t)$ contains diagrams with crossings at n pairs, and the subspaces $Z_n(h,s,t)$ of cycles and $B_n(h,s,t)$ of boundaries. We are interested in maps obeying the following property.

Definition 3.12 (Inverting differential)**.** *A $\mathbb{Z}_2$-module homomorphism $\int \colon Z_n(h,s,t) \longrightarrow \mathcal{A}_{n+1}(h,s,t)$* inverts the differential *if $x = \partial \int x$.*

If we have maps inverting the differential on $Z_n(h,s,t)$ for all n, of course these can be combined into a map $Z(h,s,t) \longrightarrow \mathcal{A}(h,s,t)$ of Maslov degree 1 such that $\partial \int = 1$.

Lemma 3.11 says that A_P^* (more precisely, its restriction to $Z_n(h,s,t)$) inverts the differential.

When we have viable H-data (h,s,t) with several all-on once occupied matched pairs $P_1, P_2, \ldots$, there are several creation operators $A_{P_1}^*, A_{P_2}^*, \ldots$ on $\mathcal{A}(h,s,t)$, and hence many ways to invert the differential. However, not every operator which inverts the differential is a creation operator.

For one thing, we can simply choose different creation operator on each Maslov summand. For another, we can also replace a creation operator with a sum of an odd number of creation operators. More fundamentally, however, not every operator $\int \colon Z_n(h,s,t) \longrightarrow \mathcal{A}_{n+1}(h,s,t)$ inverting the differential is a sum of creation operators.

Proposition 3.13. *Let (h,s,t) be twisted H-data with $L \geq 0$ all-on doubly occupied pairs and $N \geq 1$ all-on once occupied pairs. Then the set of $\mathbb{Z}_2$-module homomorphisms $\int \colon Z_n(h,s,t) \longrightarrow \mathcal{A}_{n+1}(h,s,t)$ inverting the differential is an affine $\mathbb{Z}_2$ vector space of dimension*

$$\left(\sum_k \binom{L}{k} \binom{N+k-1}{n} \right) \left(\sum_k \binom{L}{k} \binom{N+k-1}{n+1} \right), \tag{7}$$

while the set of linear combinations of creation operators is a $\mathbb{Z}_2$ vector space of dimension N.

Clearly the expression 7 is in general much larger than N, so there exist many more maps inverting the differential than linear combinations of creation operators, as mentioned in section 1.3.

For instance, taking $N = 1, L = 1, n = 0$ the dimensions are 1 and 2; taking $N = 4, L = 0, n = 1$, the dimensions are 4 and 9; and so on.

Proof. Since (h, s, t) is twisted, $\mathcal{H}(h, s, t) = 0$ (proposition 2.14), so $Z_n(h, s, t) = B_n(h, s, t)$ for all n. Let P be an arbitrarily chosen all-on once occupied pair.

Let $\mathcal{S}$ be the set of $\mathbb{Z}_2$-module homomorphisms $Z_n(h, s, t) \longrightarrow \mathcal{A}_{n+1}(h, s, t)$ that invert the differential. Denoting the restriction of A^*_P to $Z_n(h, s, t)$ (somewhat abusively) again by A^*_P, by lemma 3.11 $A^*_P \in \mathcal{S}$.

Let $\mathcal{T}$ be the set of $\mathbb{Z}_2$-module homomorphisms $Z_n(h, s, t) \longrightarrow \mathcal{A}_{n+1}(h, s, t)$ with image in $Z_{n+1}(h, s, t)$. A $\mathbb{Z}_2$-module homomorphism $T \colon Z_n(h, s, t) \longrightarrow \mathcal{A}_{n+1}(h, s, t)$ lies in $\mathcal{T}$ if and only if $\partial T = 0$.

We now observe that $T \in \mathcal{T}$ iff $A^*_P + T \in \mathcal{S}$. Indeed, since $\partial A^*_P = 1$ we have $\partial T = 0$ iff $\partial(A^*_P + T) = 1$. Thus $\mathcal{S} = A^*_P + \mathcal{T}$.

Now the space of all linear combinations of creation operators has dimension N: it has basis given by operators $A^*_{P'}$, over each of the N all-on once occupied pairs P'. (They are linearly independent as they affect distinct matched pairs.)

On the other hand, $\mathcal{T}$ is nothing but the set of $\mathbb{Z}_2$-module homomorphisms $Z_n(h, s, t) \longrightarrow Z_{n+1}(h, s, t)$. Hence $\mathcal{T}$ is a $\mathbb{Z}_2$ vector space of dimension $\dim Z_n(h, s, t) \dim Z_{n+1}(h, s, t)$, and $\mathcal{S}$ is an affine vector space of the same dimension. Lemma 2.22 gives the dimension of $Z_n(h, s, t) = B_n(h, s, t)$ as $\sum_k \binom{L}{k}\binom{N+k-1}{n}$, which yields the desired result. $\square$

Question 3.14. *For twisted H-data (h, s, t), is every diagrammatically simple $\mathbb{Z}_2$-module homomorphism $Z_n(h, s, t) \longrightarrow \mathcal{A}_{n+1}(h, s, t)$ inverting the differential a creation operator?*

The above deals with inverting the differential when (h, s, t) is twisted, i.e. there is at least one all-on once occupied pair. When there are *no* all-on once occupied pairs, i.e. (h, s, t) is tight, we have the following.

Lemma 3.15. *Suppose (h, s, t) is tight. If $x \in \mathcal{A}(h, s, t)$ has pure Maslov grading, and $x = \partial f$ for some $f \in \mathcal{A}(h, s, t)$ also of pure Maslov grading, then $f \in \mathcal{F}$.*

Proof. Since f has pure Maslov grading and $\partial f = x$, f is a sum of viable diagrams, each crossed at one more matched pair than x. From proposition 2.11 and table 1, crossings can only occur at all-on once or doubly occupied pairs. But tight (h, s, t) have none of the former (proposition 2.14), so any crossing occurs at a doubly occupied pair. Hence all diagrams in f lie in $\mathcal{F}$. $\square$

In other words, the H-data only permits crossings in places which immediately land us in $\mathcal{F}$.

3.5 Global creation operators

For any twisted H-data (h, s, t) on the arc diagram $\mathcal{Z}$ (i.e. $(h, s, t) \in \mathbf{w}(\mathcal{Z})$: definition 2.15), there is an all-on once occupied pair P (proposition 2.14), and hence a creation operator A^*_P on $\mathcal{A}(h, s, t)$ (definition 3.8) which inverts the differential (lemma 3.11).

We now introduce formalism to piece together such operators into a "global" operator on all twisted summands, by choosing an all-on once occupied pair in each twisted summand.

Definition 3.16 (Creation choice function)**.** *A* creation choice function *for $\mathcal{Z}$ assigns to each $(h, s, t) \in \mathbf{w}(\mathcal{Z})$ one of its all-on once occupied matched pairs.*

If $\mathcal{C}$ is a creation choice function and (h, s, t) is twisted, we write $\mathcal{C}(h, s, t)$ for the all-on once occupied matched pair chosen by $\mathcal{C}$. Hence there is a creation operator

$$A^*_{\mathcal{C}(h,s,t)} \colon \mathcal{A}(h, s, t) \longrightarrow \mathcal{A}(h, s, t).$$

Definition 3.17. *Let $\mathcal{C}$ be a creation choice function for $\mathcal{Z}$. The* creation operator *of $\mathcal{C}$ is the $\mathbb{Z}_2$-module homomorphism*

$$A^*_{\mathcal{C}} \colon \bigoplus_{(h,s,t)\in\mathbf{w}} \mathcal{A}(h,s,t) \longrightarrow \bigoplus_{(h,s,t)\in\mathbf{w}} \mathcal{A}(h,s,t) \quad \text{given by} \quad A^*_{\mathcal{C}} = \bigoplus_{(h,s,t)\in\mathbf{w}} A^*_{\mathcal{C}(h,s,t)}.$$

Putting together what we know on each summand, in particular the Heisenberg relation (lemma 3.9) and differential inversion (lemma 3.11), we immediately obtain the following.

Proposition 3.18. *The creation operator $A^*_{\mathcal{C}}$ of a creation choice function $\mathcal{C}$ preserves H-grading, has Maslov degree 1, and satisfies*

$$A^*_{\mathcal{C}}\partial + \partial A^*_{\mathcal{C}} = 1.$$

Moreover, for any cycle $x \in \bigoplus_{(h,s,t)\in\mathbf{w}} \mathcal{A}(h,s,t)$, we have $x = \partial A^*_{\mathcal{C}} x$. □

Moreover, applying lemma 3.10 over all $(h,s,t) \in \mathbf{w}$, we immediately have the following.

Proposition 3.19. *Suppose $\int \colon \bigoplus_{(h,s,t)\in\mathbf{w}} \mathcal{A}(h,s,t) \longrightarrow \bigoplus_{(h,s,t)\in\mathbf{w}} \mathcal{A}(h,s,t)$ is a diagrammatically simple $\mathbb{Z}_2$-module homomorphism which preserves H-data, has pure Maslov degree, and satisfies*

$$\int \partial + \partial \int = 1.$$

Then $\int$ is the creation operator $A^*_{\mathcal{C}}$ of a creation choice function $\mathcal{C}$. □

3.6 Cycle selection and creation operators via ordering

In section 3.1 we defined a diagrammatically simple cycle selection homomorphism $f^{\mathcal{C}} \colon \mathcal{H} \longrightarrow \mathcal{A}$, for any cycle choice function $\mathcal{C}$. Indeed, we saw (lemma 3.5) that any diagrammatically simple cycle selection homomorphism is of this form.

Quite separately, in section 3.5 we defined a creation operator $A^*_{\mathcal{C}}$, for any creation choice function $\mathcal{C}$. Moreover, we showed (proposition 3.19) that any diagrammatically simple chain homotopy from the identity to with appropriate gradings is of the form $A^*_{\mathcal{C}}$ for some creation choice function $\mathcal{C}$.

We now discuss a useful method to obtain such "choice functions", of both types.

Definition 3.20. *Let the pairs of $\mathcal{Z}$ be $P_1 = \{p_1, p_1'\}, \ldots, P_k = \{p_k, p_k'\}$. A* pair ordering *on $\mathcal{Z}$ consists of a total order on each of the sets*

$$\{P_1, \ldots, P_k\}, P_1, P_2, \ldots, P_k.$$

Thus a pair ordering puts the pairs of $\mathcal{Z}$ in some order; and also puts the two places of each pair in some order. We denote a pair ordering by $\preceq$, and use this symbol for each of the total orders involved.

We note that $\mathcal{Z}$ comes with several naturally ordered sets that can be used to give a pair ordering. Recall that $\mathbf{Z}$ consists of l intervals $Z_1, \ldots, Z_l$. Each interval is naturally totally ordered (indeed this order is used in defining strand diagrams). Further, a total ordering of these intervals is implicitly given in listing them as $Z_1, \ldots, Z_l$. Then $\mathbf{Z}$ is totally ordered, and as $\mathbf{a}$ and each P_i are subsets of $\mathbf{Z}$ for each i, they inherit a total order. The ordering on places can also be used to obtain an ordering on the set $\{P_1, \ldots, P_k\}$, in various reasonable ways; for instance if $P_i = \{p_i, p_i'\}$ and $P_j = \{p_j, p_j'\}$, we could define $P_i \prec P_j$ when $\min_{\preceq}\{p_i, p_i'\} \prec \min_{\preceq}\{p_j, p_j'\}$. Thus we obtain a pair ordering. But there is nothing natural about this way to order pairs, just as there is nothing natural about the ordering $Z_1, \ldots, Z_l$ of intervals; reordering the Z_i yields a homeomorphic arc diagram, but an entirely different pair ordering.

Nonetheless, once we have a pair ordering, we naturally obtain a cycle choice function and a creation choice function, as follows.

Definition 3.21. *Let $\preceq$ be a pair ordering on $\mathcal{Z}$.*

(i) The cycle choice function of $\preceq$*, denoted* $\mathcal{CY}^{\preceq}$*, assigns to each set of tight H-data* $(h,s,t) \in \mathbf{g}(\mathcal{Z})$ *the pair choice function on* $\mathbf{P}_{h,s,t}$ *which chooses from each all-on doubly occupied pair its* $\preceq$*-minimal place.*

(ii) The creation choice function of $\preceq$*, denoted* $\mathcal{CR}^{\preceq}$*, assigns to each twisted set of H-data* $(h,s,t) \in \mathbf{w}(\mathcal{Z})$ *its* $\preceq$*-minimal all-on once occupied matched pair.*

Note the definition of $\mathcal{CY}^{\preceq}$ uses the total ordering on the P_i, while the definition of $\mathcal{CR}^{\preceq}$ uses the total ordering on $\{P_1, \ldots, P_k\}$.

Suppose $P_i = \{p_i, p_i'\}$ is an all-on doubly occupied pair for the tight H-data (h,s,t), with $p_i \prec p_i'$. Since $\mathcal{CY}^{\preceq}(h,s,t)$ chooses the $\preceq$-minimal place at each all-on doubly occupied pair, the resulting cycle selection homomorphism $f^{\mathcal{CY}^{\preceq}}$ always chooses a diagram with strands beginning and ending at p_i rather than p_i'.

Now suppose the pairs of $\mathcal{Z}$ are ordered as $P_1 \prec P_2 \prec \cdots \prec P_k$. Since $\mathcal{CR}^{\preceq}$ chooses, for given twisted H-data (h,s,t), the all-on once occupied matched pair which is minimal in the ordering of $\{P_1, \ldots, P_k\}$. So the resulting creation operator $A^*_{\mathcal{CR}^{\preceq}}$ inserts a crossing at P_1, if it is all-on once occupied; otherwise at P_2, if it is all-on once occupied; and so on.

Clearly not every cycle choice function arises from a pair ordering, nor does every creation choice function. Nonetheless pair orderings provide a useful method to construct cycle choice functions and creation choice functions, and thus to construct A_∞ structures on $\mathcal{H}$.

4 Tensor products of strand diagrams

4.1 Anatomy and terminology for tensor products

We now consider tensor powers of $\mathcal{A}$ and $\mathcal{H}$ in some detail. Since $\mathcal{A}$ is freely generated (as a $\mathbb{Z}_2$ vector space) by (symmetrised $\mathcal{Z}$-constrained augmented strand) diagrams on $\mathcal{Z}$, its tensor power $\mathcal{A}^{\otimes n}$ is freely generated by tensor products $D_1 \otimes \cdots \otimes D_n$ of such diagrams.

The Maslov and H-gradings on $\mathcal{A}$ naturally carry over to $\mathcal{A}^{\otimes n}$, in such a way that the gradings of $D_1 \otimes \cdots \otimes D_n$ agree with those of the product $D_1 \cdots D_n$ (when it is nonzero) in $\mathcal{A}$. Recall (section 2.1) that $h(D), \iota(D)$ denote the H- and Maslov gradings of a diagram D, and $m \colon H_1(\mathbf{Z}, \mathbf{a}) \times H_0(\mathbf{a}) \longrightarrow \frac{1}{2}\mathbb{Z}$ counts average local multiplicities around places.

Throughout this section, let $D_1, \ldots, D_n$ be (symmetrised $\mathcal{Z}$-constrained) augmented diagrams with H-data $(h_1, s_1, t_1), \ldots, (h_n, s_n, t_n)$, and homology classes $M_1, \ldots, M_n$. As usual, the non-augmented and augmented cases are similar and we write $\mathcal{A}$ rather than $\mathcal{A}^{aug}$.

We saw in section 2.1 that $\iota(D_1 D_2) = \iota(D_1) + \iota(D_2) + m(h_2, \partial h_1)$ and $h(D_1 D_2) = h(D_1) + h(D_2)$. Applying this result repeatedly shows that

$$h\left(D_1 \cdots D_n\right) = \sum_{i=1}^n h(D_i) = \sum_{i=1}^n h_i \quad \text{and} \quad \iota\left(D_1 \cdots D_n\right) = \sum_{i=1}^n \iota(D_i) + \sum_{1 \leq j < k \leq n} m\left(h_k, \partial h_j\right),$$

motivating the following definition.

Definition 4.1 (Gradings for tensor products)**.**

(i) The H-grading *of* $D_1 \otimes \cdots \otimes D_n$ *or* $M_1 \otimes \cdots \otimes M_n$ *is*

$$h\left(D_1 \otimes \cdots \otimes D_n\right) = h\left(M_1 \otimes \cdots \otimes M_n\right) = \sum_{i=1}^n h_i \in H_1(\mathbf{Z}, \mathbf{a}).$$

(ii) The Maslov grading *of* $D_1 \otimes \cdots \otimes D_n$ *or* $M_1 \otimes \cdots \otimes M_n$ *is*

$$\iota\left(D_1 \otimes \cdots \otimes D_n\right) = \iota\left(M_1 \otimes \cdots \otimes M_n\right) = \sum_{i=1}^n \iota(D_i) + \sum_{1 \leq j < k \leq n} m\left(h_k, \partial h_j\right).$$

The notion of viability (section 2.3) also extends usefully to $\mathcal{A}^{\otimes n}$. We add a requirement that idempotents must match. Thus, while viability of diagrams is fundamentally a property of H-grading, viability of tensor products is fundamentally a property of *H-data*.

The definition uses the decomposition of $\mathcal{A}^{\otimes n}$ over H-data: from $\mathcal{A} = \bigoplus_{(h,s,t)} \mathcal{A}(h,s,t)$, we have

$$\mathcal{A}^{\otimes n} = \bigoplus_{(h_1,s_1,t_1),\ldots,(h_n,s_n,t_n)} \Big(\mathcal{A}(h_1,s_1,t_1) \otimes \mathcal{A}(h_2,s_2,t_2) \otimes \cdots \otimes \mathcal{A}(h_n,s_n,t_n) \Big).$$

There is a similar decomposition of the homology $\mathcal{H}^{\otimes n}$.

Definition 4.2 (Viability for tensor products)**.**

(i) A sequence of H-data $(h_1,s_1,t_1),\ldots,(h_n,s_n,t_n)$ *is* viable *if the following conditions hold:*

(a) for each $1 \leq i \leq n-1$, $t_i = s_{i+1}$*; and*

(b) $h_1 + \cdots + h_n$ *is viable (i.e. has multiplicity* 0 *or* 1 *on each step of* $\mathbf{Z}$*).*

(ii) A summand $\mathcal{A}(h_1,s_1,t_1) \otimes \cdots \otimes \mathcal{A}(h_n,s_n,t_n)$ *of* $\mathcal{A}^{\otimes n}$*, or a summand* $\mathcal{H}(h_1,s_1,t_1) \otimes \cdots \otimes \mathcal{H}(h_n,s_n,t_n)$ *of* $\mathcal{H}^{\otimes n}$*, is* viable *if the sequence of H-data* (h_1,s_1,t_1), …, (h_n,s_n,t_n) *is viable.*

(iii) An element of $\mathcal{A}^{\otimes n}$ *or* $\mathcal{H}^{\otimes n}$ *is* viable *if it lies in a viable summand.*

We refer to the first condition of (i), that all $t_i = s_{i+1}$, as *idempotent matching*. When it fails, we say we have an *idempotent mismatch*.

Thus the tensor product of diagrams $D_1 \otimes \cdots \otimes D_n \in \mathcal{A}^{\otimes n}$ is viable iff their sets of H-data $(h_1,s_1,t_1),\ldots,(h_n,s_n,t_n)$ form a viable sequence. Similarly, a tensor product of homology classes of diagrams is viable iff their sets of H-data form a viable sequence.

We observe that when $n = 1$, the notions of viability discussed above reduce to those previously discussed. The notion of H-data then naturally extends to a tensor product.

Definition 4.3 (H-data of tensor product)**.** *If* $D_1 \otimes \cdots \otimes D_n$ *or* $M_1 \otimes \cdots \otimes M_n$ *is viable, its* H-data *is the triple* $(h_1 + \cdots + h_n, s_1, t_n)$.

The notions of occupation of places (definition 2.6) and pairs (definition 2.7) only depend on H-grading; and on/off terminology (definition 2.8) only depends on idempotents. Hence these notions naturally extend to tensor products.

Definition 4.4 (Occupation of places and pairs by tensor product)**.**

(i) A place is unoccupied, pre-half-occupied, post-half-occupied, half-occupied, or fully occupied by $D_1 \otimes \cdots \otimes D_n$ *(resp.* $M_1 \otimes \cdots \otimes M_n$*) if it is so occupied by* $h(D_1 \otimes \cdots \otimes D_n)$ *(resp.* $h(M_1 \otimes \cdots \otimes M_n)$*).*

(ii) A pair P *is unoccupied, one-half-occupied, pre/post-one-half-occupied, alternately occupied, sesqui-occupied, pre/post-sesqui occupied, or doubly occupied by* $D_1 \otimes \cdots \otimes D_n$ *(resp.* $M_1 \otimes \cdots \otimes M_n$*) if it is so occupied by* $h(D_1 \otimes \cdots \otimes D_n)$ *(resp.* $h(M_1 \otimes \cdots \otimes M_n)$*).*

Definition 4.5 (Idempotent terminology for tensor products)**.** *A viable tensor product* $D_1 \otimes \cdots \otimes D_n$ *or* $M_1 \otimes \cdots \otimes M_n$ *is off-off/00, off-on/01, on-off/10 or on-on/11 at a pair* P *accordingly as the H-data of* $D_1 \otimes \cdots \otimes D_n$ *or* $M_1 \otimes \cdots \otimes M_n$ *is off-off/00, off-on/01, on-off/10 or on-on/11 at* P.

Thus we describe the on/off/0/1 status of $D_1 \otimes \cdots \otimes D_n$ at P via definition 2.8 using $s = s_1$ and $t = t_n$. In other words, we use the idempotents s_1, t_n at the beginning and end of the tensor product; the "interior" idempotents $s_2, \ldots, s_{n-1}$ are irrelevant. (For this reason we avoid the terminology "all-on" or "all-off" in this context, and prefer the more awkward but less misleading "on-on", 11, etc.)

We can also extend the notion of a local diagram to tensor products.

Definition 4.6 (Local tensor product)**.** *Let P be a matched pair of $\mathcal{Z}$ and $D_1 \otimes \cdots \otimes D_n \in \mathcal{A}(\mathcal{Z})^{\otimes n}$ a viable tensor product. The* local tensor product $(D_1 \otimes \cdots \otimes D_n)_P$ *is the tensor product of local diagrams on $\mathcal{Z}_P$*

$$(D_1 \otimes \cdots \otimes D_n)_P = (D_1)_P \otimes \cdots \otimes (D_n)_P \in \mathcal{A}_P^{\otimes n}.$$

Similarly for homology classes,

$$(M_1 \otimes \cdots \otimes M_n)_P = (M_1)_P \otimes \cdots \otimes (M_n)_P \in \mathcal{H}_P^{\otimes n}.$$

An important property of viability is the following.

Lemma 4.7.

(i) If the product $D_1 \cdots D_n$ is a nonzero viable diagram, then $D = D_1 \otimes \cdots \otimes D_n$ is viable.

(ii) If the product $M_1 \cdots M_n$ is a nonzero homology class, then $M = M_1 \otimes \cdots \otimes M_n$ is viable.

Proof. In both cases, if idempotents don't match then the product is zero. In the first case, viability of $D_1 \cdots D_n$ then implies viability of D. In the second case, $\mathcal{H}(h, s, t)$ is only nonzero for viable (h, s, t), so $M_1 \cdots M_n$ nonzero then implies viability of M. □

Note that neither of the converses to lemma 4.7(i) or (ii) is true: there exist viable $D_1 \otimes \cdots \otimes D_n$ with $D_1 \cdots D_n = 0$, and viable $M_1 \otimes \cdots \otimes M_n$ with $M_1 \cdots M_n = 0$. In fact, $D_1 \otimes \cdots \otimes D_n$ and $M_1 \otimes \cdots \otimes M_n$ may be viable, yet there may not exist any strand diagram with its H-data! We will introduce the notions of "critical" and "singular" to describe these phenomena below in section 4.2.

However if $D_1 \otimes \cdots \otimes D_n$ is viable, and the product $D_1 \cdots D_n$ is nonzero, then $D_1 \cdots D_n$ is a viable diagram, so (definition 2.9) is tight, twisted, or crossed.

4.2 Tightness of tensor products of diagrams

As usual, throughout this section $D_1, \ldots, D_n$ are diagrams on $\mathcal{Z}$. We extend our notions of tightness/twistedness to tensor products.

Definition 4.8 (Tightness of tensor product)**.** *Suppose $D = D_1 \otimes \cdots \otimes D_n$ is viable, with H-data (h, s, t).*

(i) If $D_1 \cdots D_n$ is nonzero and tight, and all D_i are tight, then D is tight.

(ii) If $D_1 \cdots D_n$ is nonzero and tight, but not all D_i are tight, then D is sublime.

(iii) If $D_1 \cdots D_n$ is nonzero and twisted, then D is twisted.

(iv) If $D_1 \cdots D_n$ is nonzero and crossed, then D is crossed.

(v) If $D_1 \cdots D_n = 0$, but $\mathcal{A}(h, s, t) \neq 0$, then D is critical.

(vi) If $D_1 \cdots D_n = 0$ and $\mathcal{A}(h, s, t) = 0$, then D is singular.

Definition 4.8 presents tightness as as a list of things that go increasingly wrong. First a diagram is not tight; then the resulting diagram is twisted; then crossed; then zero; and then its existence is nonsensical. It is clear that any viable D falls into precisely one of these types.

Note that this definition applies equally well if we consider D at a single matched pair $P = \{p, p'\}$, or as a whole. So we may say that D is tight (sublime, twisted, etc.) at P, meaning that D_P is tight (sublime, twisted, etc.) on $\mathcal{Z}_P$.

The condition that $\mathcal{A}(h, s, t) \neq 0$ is equivalent to the existence of a diagram with H-data (h, s, t). Thus when D is singular, its H-data (h, s, t) is also singular: no diagram exists with this H-data.

When $n = 1$, the notions of tightness reduce to those of definition 2.9; the sublime, critical and singular cases do not arise.

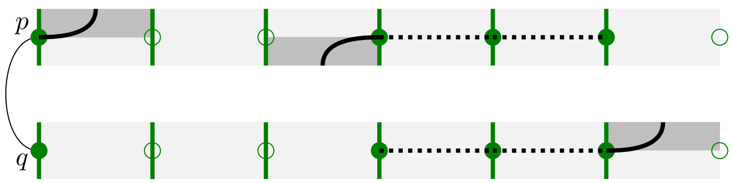

Figure 10: A sesqui-occupied tight tensor product of 6 diagrams.

When D is twisted at a pair $P = \{p, p'\}$, then $D_1 \cdots D_n$ is twisted at one of the places p or p' (definition 2.18), and we say D is twisted at p or p' accordingly.

When D is viable, since idempotents match according to $t_i = s_{i+1}$, we can draw D as a sequence of strand diagrams, side by side, where the right hand side of D_i coincides with the left hand side of D_{i+1}. Thus we regard $\mathcal{A}^{\otimes n}$ as a "horizontal" tensor product, and the local decomposition $\mathcal{A} = \bigotimes \mathcal{A}_P$ as a "vertical" tensor product.

For instance, figure 10 depicts a viable tensor product of 6 diagrams, at a pair which is sesqui-occupied. It may easily be included into a tensor product of non-augmented diagrams on a larger arc diagram.

The following lemma generalises lemma 2.10.

Lemma 4.9 (Local-global tightness of tensor product)**.** *Let $D = D_1 \otimes \cdots \otimes D_n$ be viable.*

(i) D is tight iff D is tight at all matched pairs.

(ii) D is sublime iff D is tight or sublime at all matched pairs, and sublime at ≥ 1 matched pair.

(iii) D is twisted iff D is tight, sublime or twisted at all matched pairs, and twisted at ≥ 1 matched pair.

(iv) D is crossed iff D is tight, sublime, twisted or crossed at all matched pairs, and crossed at ≥ 1 matched pair.

(v) D is critical iff D is tight, sublime, twisted, crossed or critical at all matched pairs, and critical at ≥ 1 matched pair.

(vi) D is singular iff D is singular at ≥ 1 matched pair.

Thus decomposing tensor products of diagrams according to matched pairs, we see that there is an increasing order of degeneracy

$$\text{tight} < \text{sublime} < \text{twisted} < \text{crossed} < \text{critical} < \text{singular},$$

and the tightness of a tensor product of diagrams is given by its "most degenerate" local tensor product.

An equivalent "local-to-global" formulation (as in section 2.6) is that D is X iff D is X at all matched pairs, where X is any of the following ascending properties: tight; tight or sublime; tight, sublime or twisted; etc.

Proof. Let D have H-data (h, s, t).

First, D is singular iff $\mathcal{A}(h, s, t) = 0$, iff (h, s, t) is singular, iff some (h_P, s_P, t_P) is singular (lemma 2.16), iff some $\mathcal{A}_P(h_P, s_P, t_P) = 0$, iff some D_P is singular. So we may now assume all D and D_P are non-singular, hence diagrams exist with H-data (h, s, t).

If D is critical then $D_1 \cdots D_n = \bigotimes_P (D_1 \cdots D_n)_P = 0$ so some $(D_1 \cdots D_n)_P = 0$, hence D is critical at some matched pair. Conversely, if D is critical at P then $(D_1 \cdots D_n)_P = 0$ so $D_1 \cdots D_n = 0$ and D is critical.

If D is crossed then $D_1 \cdots D_n$ is nonzero and crossed. By viability (lemma 2.2), each crossing occurs at some matched pair P, hence D is crossed at P. Moreover, from above no matched pair is critical; hence D is tight, sublime, twisted or crossed at each pair. Conversely, suppose D is crossed at P and is not critical at any pair. Then $D_1 \cdots D_n$ is crossed at P, and nonzero at every matched pair, so $D_1 \cdots D_n$ is nonzero and crossed, hence D is crossed.

If D is twisted then $D_1 \cdots D_n$ is nonzero and twisted; hence (lemma 2.10) $D_1 \cdots D_n$ is twisted at some pair P but not crossed at any pair. Thus D is twisted at P, but from above, D is not critical or crossed at any matched pair; hence D is tight, sublime or twisted at each matched pair. Conversely, suppose D is tight, sublime, or twisted at each matched pair, and twisted at some matched pair P. Then $D_1 \cdots D_n$ is nonzero and crossingless at each matched pair, and twisted at P. Thus $D_1 \cdots D_n$ is nonzero and twisted, so D is twisted.

If D is sublime then $D_1 \cdots D_n$ is nonzero and tight, but some D_i is not tight. Since the product $D_1 \cdots D_n$ is tight, it is tight at each matched pair (lemma 2.10), so each D_P is tight or sublime. But if all D_P are tight, then all $(D_i)_P$ are tight, so all D_i are tight. Hence at least one D_P must be sublime. Conversely, suppose D is tight or sublime at all matched pairs, and sublime at some matched pair. Then $D_1 \cdots D_n$ is tight at each pair, hence tight. However if all D_i are tight, then all $(D_i)_P$ are tight, so all D_P are tight. Thus some D_i is not tight, so D is sublime.

If D is tight then the product $D_1 \cdots D_n$ is tight, and each D_i is tight. Hence $D_1 \cdots D_n$ and each D_i are tight at each matched pair. Thus D is tight at each matched pair (lemma 2.10). Conversely, if D is tight at each matched pair then $D_1 \cdots D_n$ is tight at each matched pair, as is each D_i. Thus $D_1 \cdots D_n$ is tight, as are the D_i. So D is tight. □

4.3 Sub-tensor-products, extension and contraction

It is useful to consider the following notions regarding tensor products.

Definition 4.10. *Let $D_1, \cdots, D_n$ be diagrams, with homology classes $M_1, \cdots, M_n$. Consider the tensor products $D = D_1 \otimes \cdots \otimes D_n \in \mathcal{A}^{\otimes n}$ and $M = M_1 \otimes \cdots \otimes M_n \in \mathcal{H}^{\otimes n}$.*

(i) A sub-tensor-product *of D is a tensor product $D' = D_i \otimes D_{i+1} \otimes \cdots \otimes D_{j-1} \otimes D_j$, where $1 \leq i \leq j \leq n$.*

(ii) A sub-tensor-product *of M is a tensor product $M' = M_i \otimes M_{i+1} \otimes \cdots \otimes M_{j-1} \otimes M_j$, where $1 \leq i \leq j \leq n$.*

Clearly if D (resp. M) is a viable tensor product, then any sub-tensor-product D' (resp. M') is also viable.

Now a diagram is an idempotent iff all its strands are horizontal. Idempotents can be inserted into a tensor product of strand diagrams to "extend" it, as in the following straightforward statement, which also gives a method to "contract" it.

Lemma 4.11 (Extending and contracting tensor products)**.** *Let $D = D_1 \otimes \cdots \otimes D_n$ be a viable tensor product of diagrams. Let M_i be the homology class of D_i, and let $M = M_1 \otimes \cdots \otimes M_n$. Let D_i^* be the unique idempotent diagram consisting of dotted horizontal strands at all places of $t_i = s_{i+1}$, and let M_i^* be its homology class.*

(i) (a) The tensor product $D' = D_1 \otimes \cdots \otimes D_i \otimes D_i^* \otimes D_{i+1} \otimes \cdots \otimes D_n$ is also viable.

(b) Suppose that for some $1 \leq i \leq j \leq n$, the product $D_i D_{i+1} \cdots D_j$ is nonzero. Then $D'' = D_1 \otimes \cdots \otimes D_{i-1} \otimes (D_i \cdots D_j) \otimes D_{j+1} \otimes \cdots \otimes D_n$ is also viable.

(ii) (a) Suppose all M_i are nonzero. Then $M' = M_1 \otimes \cdots M_i \otimes M_i^* \otimes M_{i+1} \otimes \cdots \otimes M_n$ is also viable.

(b) Suppose that for some $1 \leq i \leq j \leq n$, the product $M_i M_{i+1} \cdots M_j$ is nonzero. Then $M'' = M_1 \otimes \cdots \otimes M_{i-1} \otimes (M_i \cdots M_j) \otimes M_{j+1} \otimes \cdots \otimes M_n$ is also viable.

□

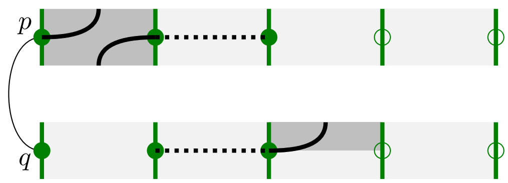

Figure 11: An extension-contraction of figure 10.

Definition 4.12.

(i) In lemma 4.11(i), we say D' is obtained from D by extension *by D^*, and D'' is obtained from D by* contraction *of $D_i \otimes \cdots \otimes D_j$.*

(ii) In lemma 4.11(ii), we say M' is obtained from M by extension *by M^*, and M'' is obtained from M by* contraction *of $M_i \otimes \cdots \otimes M_j$.*

(iii) We say one tensor product (of diagrams, or homology classes) is obtained from another by extension-contraction, *if it is obtained by some sequence of extensions and contractions.*

Observe that extension and contraction of a tensor product preserve H-data and Maslov grading.

Figure 11 shows an extension-contraction of figure 10.

Note that extensions by idempotents may be reversed by contraction, and contractions of idempotents may be reversed by extension. But a contraction involving more than one factor with non-horizontal strands cannot be reversed by extension; hence the following definition.

Definition 4.13. *If two or more of $D_i, D_{i+1}, \ldots, D_j$ (resp. $M_i, M_{i+1}, \ldots, M_j$) contain non-horizontal strands (i.e. are not idempotents), then contraction of $D_i \cdots D_j$ in D (resp. of $M_i \cdots M_j$ in M) is* nontrivial. *Otherwise, the contraction is* trivial.

Thus, in a trivial contraction, either all of $D_i, \ldots, D_j$ are idempotents, as is their product; or precisely one diagram D_k among $D_i, \ldots, D_j$ has non-horizontal strands, in which case $D_i \cdots D_j = D_k$.

4.4 Sublime and singular tensor products

We now collect a couple of useful facts about sublime and singular tensor products.

Lemma 4.14 (Sublime contains crossed)**.** *If the viable tensor product $D = D_1 \otimes \cdots \otimes D_n$ is sublime, then some D_i is 11 once occupied crossed at some matched pair.*

In other words, any sublimation arises by the multiplication of a crossed diagram by another diagram (or diagrams) to undo the crossing, eventually arriving at a tight diagram, as in figure 8.

Proof. Assume for contradiction that all D_i are crossingless. Since all the D_i, as well as $D_1 \cdots D_n$ are crossingless (D sublime implies $D_1 \cdots D_n$ is tight), all these diagrams have homology classes. In homology, $D_1 \cdots D_n$ is nonzero, (definition 2.9), hence so are all D_i. Thus all D_i are tight, contradicting D being sublime; hence some D_i is crossed at some pair P. From table 1, P is 11 once occupied or 11 doubly occupied by D_i. But in the latter case, all steps at P are occupied by D_i, so by viability all other D_j are idempotent at P, so $D_1 \cdots D_n$ cannot be tight, contradicting D being sublime. ☐

A singular tensor product $D = D_1 \otimes \cdots \otimes D_n$ is rather pathological: its H-data, even though arising from a viable tensor product of diagrams, is not the H-data of any single diagram. It may not be clear that singular tensor products exist. Lemma 4.9 says D is singular iff some D_P is singular. As we now show, D_P must take a very specific form.

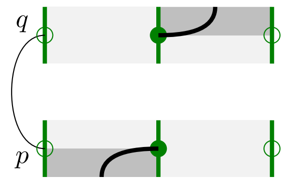

Figure 12: A singular tensor product. By lemma 4.15, up to extension, all local singular tensor products are of this form.

Lemma 4.15 (Structure of singular tensor products)**.** *Let $P = \{p, q\}$ be a matched pair, and let $D = D_1 \otimes \cdots \otimes D_n$ be a viable singular tensor product of local diagrams on $\mathcal{Z}_P$. Then P is 00 alternately occupied by D. Precisely two factors D_i, D_j (with $i < j$) contain non-horizontal strands; both are tight. All other D_k are idempotents. After possibly relabelling p and q, P is pre-one-half-occupied at p by D_i, and post-one-half-occupied at q by D_j.*

In other words, any singular local tensor product is an extension (definition 4.12) of figure 12. It is clear for the diagram $D_1 \otimes D_2$ of the figure that $D_1 D_2 = 0$, and that in fact there is no diagram with its H-data: there is no such thing as a 00 alternately occupied pair in a strand diagram, as proposition 2.11 and table 1 make clear.

Proof. Let D have H-data (h, s, t).

For a viable diagram D_0 on $\mathcal{Z}_P$, we observe that if D_0 covers an even number of the 4 steps of $\mathcal{Z}_P$, then P is 00 or 11; and if D_0 covers an odd number of these steps, then P is 01 or 10. This follows, for instance, from proposition 2.11, from inspecting table 1, or by analysing directly how strands must proceed in such a diagram. Applying this observation to each D_i in D, we see that if h covers an even number of steps of $\mathcal{Z}_P$, then P is 00 or 11; and if h covers an odd number of steps of $\mathcal{Z}_P$, then P is 01 or 10.

Moreover, all steps covered by h cannot be covered by a single D_i. For then all other D_j are idempotents, so D_i has the H-data (h, s, t) of D, contradicting D being singular. In particular, h must cover at least two steps of P.

If h covers 2 steps, then all possible H-data (h, s, t) satisfying the conditions above appear in table 1, hence are non-singular, except if P is 00 alternately occupied. In this case, the 2 steps covered must be covered by two diagrams D_i, D_j, where $i < j$; all other diagrams must be idempotents. Then P must be 01 one-half-occupied by D_i, hence pre-one-half-occupied, and 10 one-half-occupied by D_j, hence post-one-half-occupied. Thus D has the structure claimed.

If h covers 3 steps, then the only possible H-data not appearing in table 1 are where P is 01 pre-sesqui-occupied or 10 post-sesqui-occupied. We consider the first case; the second is similar. Without loss of generality suppose p is pre-half-occupied and q is fully occupied. If the 3 steps are covered by two diagrams D_i, D_j, where D_i covers one step and D_j covers two steps, then P is one-half-occupied by D_i. Moreover, by our initial observation, D_j is 00 or 11, so by viability D_i must be 10, hence P is post-one-half-occupied by D_i. Thus both p and q are pre-half-occupied by D_j, but there is no diagram which does so. If the three steps are covered by three diagrams, then P is pre-one-half-occupied by two diagrams (which must be 01) and post-one-half-occupied by one diagram (which must be 10), and all other diagrams are idempotents. But there is no way to combine the idempotent data 01, 01, 10 of these three diagrams viably so that P is 10 in the tensor product. Hence no such D exists.

If h covers all 4 steps, all possible H-data already appear in table 1 so D cannot be singular. $\square$

4.5 Enumeration of local tensor products

We now enumerate all viable local tensor products of diagrams. So let $D = D_1 \otimes \cdots \otimes D_n$ be a tensor product of diagrams on $\mathcal{Z}_P$, where $P = \{p, q\}$. Viability implies that each of the 4 steps of $\mathcal{Z}_P$ is covered at most once. Hence at most 4 of $D_1, \ldots, D_n$ can contain non-horizontal strands; the rest are idempotents.

Enumerating the possible such D is assisted by the following lemma, describing the tightness of the D_i.

Lemma 4.16 (Tightness of local sub-tensor-products and tensor factors)**.** *Let D be a viable tensor product of local diagrams on $\mathcal{Z}_P$.*

(i) If D is tight, then each D_i is tight.

(ii) If D is sublime, then one D_i is crossed 11 once occupied, and for the remaining factors D_j:

(a) one D_j is twisted, and all other D_j are idempotents; or

(b) one or two D_j are tight with non-horizontal strands, and all other D_j are idempotents.

(iii) If D is twisted, then either:

(a) precisely two D_i are tight, and all other D_j are idempotents; or

(b) precisely one D_i is twisted, and all other D_j are idempotents.

(iv) If D is crossed, then precisely one or two D_i are crossed, and all other factors D_j are idempotents.

(v) If D is critical, then no D_i are crossed, W of the D_i are twisted, G of the D_i are tight with non-horizontal strands, and all other D_i are idempotents, where $2W + G \leq 4$ and $W + G \geq 2$.

(vi) If D is singular, then precisely two D_i are tight, and all other D_j are idempotents.

In the critical case, the inequalities on W and G imply $(W, G) \in \{(2,0), (1,1), (1,2), (0,2), (0,3), (0,4)\}$. (In fact the case $(W, G) = (0, 2)$ never arises; such tensor products turn out to be singular.)

Proof. Part (i) follows by definition 4.8.

If D is sublime then by lemma 4.14 some D_i is crossed 11 once occupied. Each crossed D_i thus covers exactly 2 of the 4 steps in $\mathcal{Z}_P$, so there are at most two crossed D_i. If two factors D_i, D_j are crossed, then by viability P all other factors are idempotents and $D_1 \cdots D_n$ is crossed, contradicting D being sublime. So there is one crossed diagram D_i. Only 2 steps of $\mathcal{Z}_P$ are not covered by D_i, and so there are at most 2 other factors D_j with non-horizontal strands, which are tight or twisted. A twisted D_j would cover both the remaining steps, the only possibilities are (a) and (b) as claimed.

If D is twisted then (definition 4.8) $D_1 \cdots D_n$ is twisted, hence (table 1) only two steps of $\mathcal{Z}_P$ are covered. Thus at most 2 of the D_i are not idempotents. If one D_i is non-idempotent, then D_i is twisted. If two D_i are non-idempotent, then each must cover one step, hence both are tight.

If $D_1 \cdots D_n$ is crossed, then at least one D_i is crossed (lemma 2.17); and by viability, as any crossed diagram covers at least two steps, there are at most two crossed D_i. If there are two crossed factors, then they cover all steps, so all other factors are idempotents. If only one D_i is crossed, we observe that any viable multiplication of D_i with any tight or twisted diagram results in a tight diagram, so all other factors must be idempotents.

For (v), we first claim no D_i is crossed. As noted above at most two D_i are crossed. If two D_i are crossed then all other factors are idempotents, so that $D_1 \cdots D_n$ is nonzero crossed; if one D_i is crossed, then any viable product of D_i with a tight or twisted diagram is nonzero; either way contradicting criticality of D, proving the claim. Hence each non-idempotent D_i is twisted or tight. Each twisted factor covers exactly 2 steps (table 1); each tight factor covers at least 1 step. These factors altogether cover $\geq 2W + G$ steps. Since $\mathcal{Z}_P$ has 4 steps, so $2W + G \leq 4$. On the other hand $W + G \geq 2$ since there must be at least 2 non-idempotent factors; otherwise the single non-idempotent $D_i = D \neq 0$, contradicting criticality.

The final part is just lemma 4.15. □

Using the structure provided by lemma 4.16, or otherwise, we can enumerate viable tensor products of diagrams on $\mathcal{Z}_P$ and obtain the following generalisation of proposition 2.11.

Proposition 4.17 (Classification of viable local tensor products)**.** *Let $D = D_1 \otimes \cdots \otimes D_n$ be a viable tensor product of diagrams on $\mathcal{Z}_P$. Then D is an extension-contraction of a diagram shown in table 2, and its H-data and tightness are as shown.* □

Note that in this proposition, D may be an extension-contraction of more than one of the possibilities: a contraction of a sublime tensor product may coincide with the contraction of a tight tensor product.

Table 2 also shows Maslov gradings with each local tensor product. As mentioned in section 4.3, Maslov grading is preserved under extension and contraction. Observe that, for any given viable H-data, if there is a critical tensor product, then there is also a tight tensor product, and the Maslov grading of the latter is 1 greater than the former.

4.6 Tightness of local and sub-tensor products

In the sequel we need to know about the behaviour of tightness when we decompose or extend or contract tensor products. We saw in lemma 4.9 that when we consider tensor products of diagrams locally (i.e. decompose "vertically"), tightness behaves in an ordered way. However, when we decompose according to the "horizontal" tensor product, tightness is not so well behaved.

By proposition 4.17, a viable local tensor product is an extension-contraction of one shown in table 2. We can then enumerate the tightness of sub-tensor-products in each case, and obtain the following result.

Lemma 4.18 (Tightness of local sub-tensor-products)**.** *Let D be a viable tensor product of diagrams on $\mathcal{Z}_P$, and let D' be a sub-tensor-product of D. Then the possible tightness types of D and D' are as shown in table 3.* □

Thus, for instance, if D is tight then D' is also tight; if D' is sublime then D is also sublime; and if D' is critical then D is also critical.

We also have a similar "global" result about the possible tightness types of a tensor product D and sub-tensor product D', on a general arc diagram.

Lemma 4.19 (Tightness of sub-tensor-products)**.** *Let $D = D_1 \otimes \cdots \otimes D_n$ be a viable tensor product of diagrams on an arc diagram $\mathcal{Z}$, and let D' be a sub-tensor-product of D.*

(i) If D is tight, then D' is tight.

(ii) If D' is critical or singular, then D is critical or singular.

Every combination of tightness types not ruled out by these implications is possible.

Proof. If D is tight, then $D_1 \cdots D_n$ is tight (definition 4.8), hence nonzero in homology (definition 2.9), hence any $D_i \cdots D_j$ is nonzero in homology, hence tight, hence D' is tight.

If D' is critical or singular then $D_i \cdots D_j = 0$ (definition 4.8), so $D_1 \cdots D_n = 0$, so D is critical or singular.

We show some of the remaining possibilities in figures. Figures 13 and 14 show examples of sublime, twisted, crossed and critical tensor products containing many types of sub-tensor-products. The small number remaining are omitted. □

In figures such as 13(left) we only draw the local tensor products at one matched pair; and in figures 13(right) and 14 we only draw the local tensor products at two matched pairs. These can easily be extended to figures of tensor products of non-augmented diagrams on connected arc diagrams if desired.

On the other hand, extension-contraction preserves most, but not all, types of tightness. The only subtlety is sublimation: sublime tensor products may become tight.

H-data	Tight	Sublime	Twisted	Crossed	Critical	Singular
Unoccupied 00	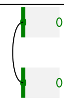 0					
Unoccupied 11	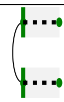 0					
Pre-one-half-occupied (01)	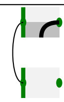 0					
Post-one-half-occupied (10)	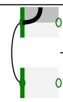 $-\frac{1}{2}$					
Alternately occupied 00						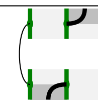 $-\frac{1}{2}$
Alternately occupied 11	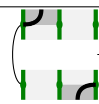 $-\frac{1}{2}$					
Once occupied 00	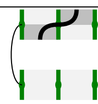 0					
Once occupied 11			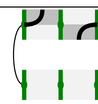 -1	 0		
Pre-sesqui-occupied (01)	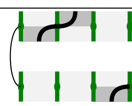 0	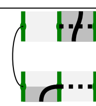 0			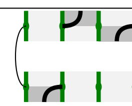 -1	
Post-sesqui-occupied (10)	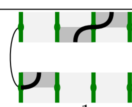 $-\frac{1}{2}$	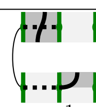 $-\frac{1}{2}$			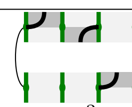 $-\frac{3}{2}$	
Doubly occupied 00	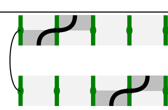 0	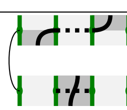 0			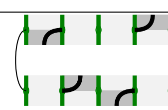 -1	
Doubly occupied 11	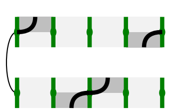 -1	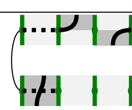 -1 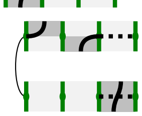		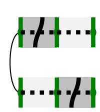 0	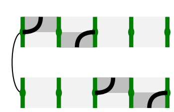 -2	

Table 2: Possible local tensor products, by H-data and tightness. Maslov gradings also shown.

D \ D'	Tight	Sublime	Twisted	Crossed	Critical	Singular
Tight	X					
Sublime	X	X	X	X		
Twisted	X		X			
Crossed	X			X		
Critical	X		X		X	X
Singular	X					X

Table 3: Possible tightness types of a viable local tensor product D and a sub-tensor-product D'.

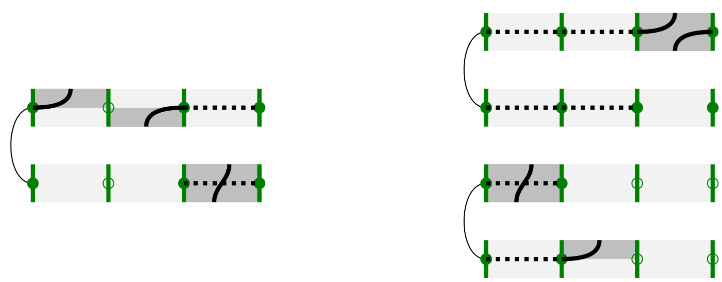

Figure 13: Left: A sublime tensor product $D_1 \otimes D_2 \otimes D_3$ containing tight (D_1, D_2), sublime $(D_2 \otimes D_3, D_1 \otimes D_2 \otimes D_3)$, twisted $(D_1 \otimes D_2)$ and crossed (D_3) sub-tensor-products. Right: A twisted tensor product $D_1 \otimes D_2 \otimes D_3$ containing tight (D_2), sublime $(D_1 \otimes D_2)$, twisted $(D_3, D_2 \otimes D_3, D_1 \otimes D_2 \otimes D_3)$ and crossed (D_1) sub-tensor-products.

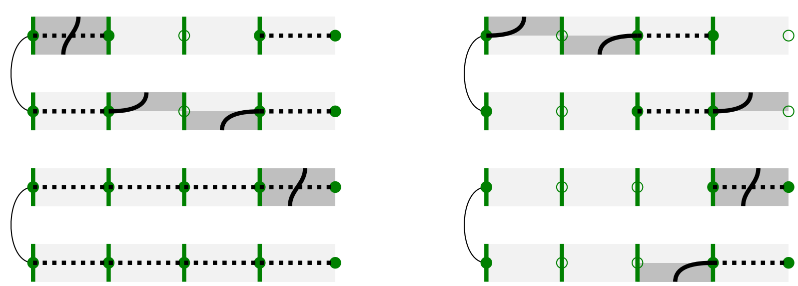

Figure 14: Left: A crossed tensor product $D_1 \otimes D_2 \otimes D_3 \otimes D_4$ containing tight (D_2, D_3), sublime $(D_1 \otimes D_2, D_1 \otimes D_2 \otimes D_3)$, twisted $(D_2 \otimes D_3)$ and crossed $(D_1, D_4, D_3 \otimes D_4, D_2 \otimes D_3 \otimes D_4, D_1 \otimes D_2 \otimes D_3 \otimes D_4)$ sub-tensor-products. Right: A critical tensor product $D_1 \otimes D_2 \otimes D_3 \otimes D_4$ containing tight $(D_1, D_2, D_3, D_2 \otimes D_3)$, sublime $(D_3 \otimes D_4)$, twisted $(D_1 \otimes D_2, D_1 \otimes D_2 \otimes D_3)$, crossed (D_4), critical $(D_1 \otimes D_2 \otimes D_3 \otimes D_4)$ and singular $(D_2 \otimes D_3 \otimes D_4)$ sub-tensor-products.

Lemma 4.20. *Suppose D' is obtained from $D = D_1 \otimes \cdots \otimes D_n$ by extension-contraction. If D, D' do not have the same tightness then D is sublime and D' is tight.*

Proof. Under extension or contraction, the product $D_1 \cdots D_n$ of a tensor product $D_1 \otimes \cdots \otimes D_n$ remains invariant, as does H-data.

Singularity of a tensor product is defined by reference only to H-data; hence hence D is singular iff D' is singular.

Now assume D, D' are not singular. The tightness properties "tight or sublime", "twisted", "crossed" and "critical" of D are defined by the properties of the product $D_1 \cdots D_n$ (i.e. whether $D_1 \cdots D_n$ is tight, twisted, crossed or zero respectively); hence these tightness properties are preserved under extension-contraction.

It remains to prove that if D is tight then D' is tight. In this case, any sub-tensor-product of D' is tight (lemma 4.19), and hence for any $1 \leq i \leq j \leq n$ the product $D_i \cdots D_j$ is tight. Thus in any extension-contraction D' of D, the product of any sub-tensor product is tight; so D' is tight. □

4.7 Tensor products of homology classes

We now return to $\mathcal{H}^{\otimes n}$. Let $M = M_1 \otimes \cdots \otimes M_n$ be a viable tensor product of nonzero homology classes of diagrams on an arc diagram $\mathcal{Z}$, where M_i has H-data (h_i, s_i, t_i) and is represented by a tight diagram D_i.

Since each D_i is tight, no D_i is crossed, so D is neither crossed (lemma 2.17) nor sublime (lemma 4.14). That only leaves the possibilities in the following definition.

Definition 4.21. *Suppose $M = M_1 \otimes \cdots \otimes M_n$ is viable, and let D_i be a diagram representing M_i. Then M is* tight, twisted, critical *or* singular *accordingly as $D = D_1 \otimes \cdots \otimes D_n$ is tight, twisted, critical or singular.*

We can also speak of M being tight, twisted, critical or singular at a matched pair P, or twisted at a place p, as for D.

However, there may be multiple choices for the D_i; so we check that tightness is well defined.

Lemma 4.22. *Let D_i, D_i' be diagrams representing M_i, and let $D = D_1 \otimes \cdots D_n$, $D' = D_1' \otimes \cdots \otimes D_n'$. Then D and D' have the same tightness.*

Proof. Consider a matched pair P. If D_i and D_i' differ at P, then P is all-on doubly occupied by D_i. In this case D_i covers all four steps of the local arc diagram $\mathcal{Z}_P$, as does D_i'; so every D_j and D_j' with $j \neq i$ is idempotent at P. Thus both D and D' are tight at P. Hence D and D' have the same tightness at each matched pair. By lemma 4.9 then D and D' have the same tightness. □

Now we observe from table 2 that any local tight, twisted, critical or singular tensor product, for any viable H-grading, can be constructed using only tight diagrams. So the table of possible local tensor products of homology classes of diagrams is precisely given by the tight, twisted, critical and singular columns of table 2. Hence we have the following proposition, using the notion of extension-contraction from definition 4.12(ii).

Proposition 4.23 (Classification of viable local tensor products of homology classes)**.** *Let $M = M_1 \otimes \cdots \otimes M_n$ be a viable tensor product of nonzero homology classes of diagrams on $\mathcal{Z}_P$. Then M is an extension-contraction of a tensor product of homology classes of diagrams shown in the tight, twisted, critical or singular columns of table 2, and its H-data and tightness are as shown.* □

Strictly speaking, table 2 shows tensor products of diagrams; proposition 4.23 refers to their homology classes. In practice when drawing or referring to homology classes, we draw and refer to diagrams (necessarily tight) representing them.

One immediate observation from the classification of proposition 4.23 and table 2, which will be useful, is the following, allowing us to say something about tightness merely from H-data.

Lemma 4.24. *Let $M = M_1 \otimes \cdots \otimes M_n$ be viable on $\mathcal{Z}_P$, with H-data (h, s, t).*

(i) M is tight or critical at P iff (h, s, t) is tight at P.

(ii) M is twisted at P iff (h, s, t) is twisted at P.

(iii) M is is singular at P iff (h, s, t) is singular at P.

□

We can distinguish tightness in $\mathcal{H}^{\otimes n}$ by the following result.

Lemma 4.25 (Characterising tightness of tensor product of homology classes)**.** *Suppose $M = M_1 \otimes \cdots \otimes M_n$ is a viable tensor product of nonzero homology classes of diagrams on an arc diagram $\mathcal{Z}$, with H-data (h, s, t), and let D_i be a diagram representing M_i.*

(i) M is tight iff $M_1 \cdots M_n \neq 0$.

(ii) M is twisted iff $M_1 \cdots M_n = 0$ but $D_1 \cdots D_n \neq 0$.

(iii) M is critical iff $D_1 \cdots D_n = 0$, but $\mathcal{A}(h, s, t) \neq 0$.

(iv) M is singular iff $\mathcal{A}(h, s, t) = 0$

Like definition 4.8, lemma 4.25 presents tightness as a list of things that go increasingly wrong. Note $\mathcal{A}(h, s, t) = 0$ implies $D_1 \cdots D_n = 0$ implies $M_1 \cdots M_n = 0$, so precisely one of these cases applies.

Recalling the isomorphism between $\mathcal{H}$ and the contact category, $M = M_1 \otimes \cdots \otimes M_n$ describes the stacking of tight cubulated contact structures on a thickened surface $\Sigma \times [0, 1]$. Cases (ii) through (iv) describe overtwisted structures, in increasing order of degeneracy. In case (ii) the stacked contact cubes above each individual square of the quadrangulation remain tight, but the overall contact structure is overtwisted (as in figure 2); in case (iii) the contact cube above some square becomes overtwisted; in case (iv) the contact cube above some square is overtwisted, even when restricted to the boundary of the cube.

Proof. Let $D = D_1 \otimes \cdots \otimes D_n$, so by definition 4.21, M and D have the same tightness, and it is sufficient to consider the tightness of D.

If D is tight then $D_1 \cdots D_n$ is tight, so $M_1 \cdots M_n \neq 0$ (definition 2.9). Conversely, if $M_1 \cdots M_n \neq 0$ then all $M_i \neq 0$, and these nonzero homology classes are represented by the diagrams $D_1 \cdots D_n$ and D_i, which are tight, so D is tight.

If D is twisted then $D_1 \cdots D_n$ is twisted, hence nonzero, but its homology class $M_1 \cdots M_n = 0$ (definition 2.9 again). Conversely, if $M_1 \cdots M_n = 0$ but the diagram $D_1 \cdots D_n$ is nonzero, then $D_1 \cdots D_n$ is not tight; it is also not crossed, since no D_i is crossed (lemma 2.17); hence it is twisted. Thus D is twisted.

The characterisations of critical and singular follow immediately from definition 4.8. □

Lemma 4.9 applied to homology immediately gives the following.

Lemma 4.26. *Let $M = M_1 \otimes \cdots \otimes M_n$ be viable, where M_i is represented by diagram D_i.*

(i) M is tight iff M is tight at all matched pairs.

(ii) M is twisted iff M is tight or twisted at each matched pair, and twisted at ≥ 1 matched pair.

(iii) M is critical iff M is tight, twisted or critical at all matched pairs, and critical at ≥ 1 matched pair.

(iv) M is singular iff M is singular at ≥ 1 matched pair.

□

M \ M'	Tight	Twisted	Critical	Singular
Tight	X			
Twisted	X	X		
Critical	X	X	X	X
Singular	X			X

Table 4: Possible tightness types of a viable local tensor product of homology classes M and a sub-tensor-product M'.

Thus the properties "tight", "'tight or twisted", "tight, twisted or critical" and "not singular" are all local-to-global properties of tensor products of homology classes.

We also consider contractions and extensions.

Lemma 4.27. *Let $M = M_1 \otimes \cdots \otimes M_n$ be a viable tensor product of diagrams on $\mathcal{Z}_P$.*

(i) If M is tight, then for all $1 \leq i \leq j \leq n$, the product $M_i \cdots M_j$ is nonzero, so $M_1 \otimes \cdots \otimes M_{i-1} \otimes (M_i \cdots M_j) \otimes M_{j+1} \otimes M_n$ is a contraction of M.

(ii) If M is twisted, critical or singular, then any contraction of M is trivial. Moreover, M is an extension of a tensor product of homology classes shown in the twisted, critical of singular columns of table 2.

Recall trivial and nontrivial contractions were defined in section 4.3 (definition 4.13).

Proof. If M is tight, then (lemma 4.25) $M_1 \cdots M_n \neq 0$; so any $M_i \cdots M_j \neq 0$.

If M is twisted, critical or singular, then by proposition 4.23, M is an extension-contraction of a tensor product shown in the appropriate column of table 2. We observe that multiplying any two consecutive diagrams in any of these tensor products yields a twisted or zero diagram, which is zero in homology. Thus no nontrivial contraction exists. □

The following fact about critical tensor products will be useful in the sequel.

Lemma 4.28 ("It takes 3 to be critical")**.** *Suppose $M = M_1 \otimes \cdots \otimes M_n$ is viable and critical on an arc diagram $\mathcal{Z}$. Then $n \geq 3$.*

Proof. By lemma 4.26, some local tensor product M_P is critical. By lemma 4.27, M_P is an extension of a critical diagrams in table 2, and all such diagrams have at least 3 factors. □

We also consider how tightness behaves under taking sub-tensor-products of local homology classes. The following lemma is immediate from applying lemma 4.18 to homology (recalling definition 4.21 of tightness). Effectively we simply cross out the sublime and crossed rows and columns of table 3.

Lemma 4.29 (Tightness of local sub-tensor-products of homology classes)**.** *Let M be a viable tensor product of nonzero homology classes of diagrams on $\mathcal{Z}_P$, and let M' be a sub-tensor-product. Then the possible tightness types of M and M' are as shown in table 4.* □

Thus, for instance, if M is tight, then M' is tight; In this case M corresponds to a tight contact manifold and M' to a contact submanifold. Similarly, if M' is critical, then M is critical; and if M' is critical or twisted, then M is critical or twisted.

Considering the tightness of sub-tensor-products globally, we only need the following statement, which is immediate from lemma 4.19.

Lemma 4.30 (Tightness of global sub-tensor-products of homology classes)**.** *Let M be a viable tensor product of nonzero homology classes of diagrams on an arc diagram $\mathcal{Z}$, and let M' be a sub-tensor-product.*

(i) If M is tight, then M' is tight.

(ii) If M' is critical or singular, then M is critical or singular.

□

Note the contrapositive of (ii): if M is tight or twisted, then M' is tight or twisted.

4.8 Generalised contraction

It is useful to generalise the notion of contraction discussed above. Let $M = M_1 \otimes \cdots \otimes M_n$ is a viable tensor product of nonzero homology classes of diagrams. A contraction of M replaces a sub-tensor-product $M' = M_i \otimes \cdots \otimes M_j$ with $M_i \cdots M_j$ provided that this product is nonzero (definition 4.12). Recalling (proposition 2.14) that for any given tight H-data (h, s, t) there is a unique nonzero homology class, we observe $M_i \cdots M_j$ is the unique homology class of diagram with the H-data of M'. This leads to the following generalisation.

Definition 4.31. *Let $M = M_1 \otimes \cdots \otimes M_n$ be a viable tensor product of nonzero homology classes of diagrams on an arc diagram $\mathcal{Z}$. Suppose a sub-tensor-product $M_i \otimes \cdots \otimes M_j$ has tight H-data, and let M^* be the unique nonzero homology class of a diagram with this H-data.*

Then we say $M' = M_1 \otimes \cdots \otimes M_{i-1} \otimes M^* \otimes M_{j+1} \otimes \cdots \otimes M_n$ is obtained from M by H-contraction *of $M_i \otimes \cdots \otimes M_j$.*

Thus H-contraction generalises contraction. Moreover if M' is obtained from M by H-contraction, then M' is viable, and has the same H-data as M.

Tightness locally behaves rather nicely under H-contraction.

Lemma 4.32. *Let M be a viable tensor product of nonzero homology classes of diagrams on $\mathcal{Z}_P$. Suppose M' is obtained from M by H-contraction.*

(i) M is tight or critical iff M' is tight or critical. Moreover,

(a) if M is tight, then M' is tight;

(b) if M' is critical, then M is critical.

(ii) M is twisted iff M' is twisted.

(iii) M is singular iff M' is singular.

Proof. The H-data of M is tight, twisted or singular accordingly as M is respectively (tight or critical), twisted or singular (lemma 4.24). Since H-contraction preserves H-data, these tightness properties are also preserved.

If M is tight, then $M^* = M_i \cdots M_j$ (lemma 4.27), so we have a bona fide contraction (definition 4.12), and M' is tight. (For instance we can note that the product of the factors in both M and M' is $M_1 \cdots M_n$, and apply lemma 4.25.)

If M' is critical, then, by proposition 4.23 and lemma 4.27, M' is an extension of one of the tensor products shown in the critical column of table 2. Thus each tensor factor of M' covers at most one step of $\mathcal{Z}_P$. Since M is obtained from M' by replacing a tensor factor of M' with $M_i \otimes \cdots \otimes M_j$, in a way that preserves H-data, each tensor factor of M also covers at most one step of $\mathcal{Z}_P$; in fact M is an extension of M'. So M is critical. □

5 Constructing A-infinity structures

5.1 The construction

We now describe A_∞ structures on $\mathcal{H}$. We first give an overview of the construction.

Recall the discussion of section 1.3. We regard $\mathcal{A}$ as an A_∞ algebra with trivial n-ary operations for $n \geq 3$. The homology $\mathcal{H}$ can be regarded as a DGA with trivial differential. We construct an A_∞ structure X on $\mathcal{H}$ extending this DGA structure, i.e. with $X_1 = 0$ and X_2 being multiplication, together with a morphism of A_∞ algebras $f\colon \mathcal{H} \longrightarrow \mathcal{A}$, consisting of maps $f_n\colon \mathcal{H}^{\otimes n} \longrightarrow \mathcal{A}$, extending a cycle selection map $f_1\colon \mathcal{H} \hookrightarrow \mathcal{A}$. Kadeishvili's construction [8] builds maps X_n, f_n and auxiliary maps U_n inductively over n. All these maps preserve H-data; X_n and U_n have Maslov grading $n-2$, and f_n has Maslov grading $n-1$. At each stage there we only need to construct f_n explicitly, then U_{n+1} and X_{n+1} are determined by equations (2) and (3).

As it turns out, we only need to construct maps $\overline{f}_n, \overline{U}_n\colon \mathcal{H}^{\otimes n} \longrightarrow \overline{\mathcal{A}}$ (definition 2.25) into the quotient $\overline{\mathcal{A}} = \mathcal{A}/\mathcal{F}$.

To construct the cycle selection homomorphism f_1, as discussed in section 3.1, we use a cycle choice function $\mathcal{CY}$. And as discussed in section 3.6, a cycle choice function can be constructed from a pair ordering on the arc diagram $\mathcal{Z}$.

To construct f_n for $n \geq 2$, we need to solve equation (4): $f_1 X_n - U_n = \partial f_n$. Since all maps preserve H-data, this can be done separately on each H-summand. On twisted summands, defining f_n amounts to inverting the differential, as discussed in section 3.4. We use the creation operators of section 3.3, choosing appropriate creation operators on each summand using a creation choice function, as discussed in section 3.5. On other summands, it turns out that no choice is necessary, once we project to $\overline{\mathcal{A}}$; we can take $\overline{f}_n = 0$ in this case.

Our construction will satisfy the following condition on the maps f_n. The idea is that if $f_1 X_n - U_n = 0$, then it reasonable to say that f_n "ought" to be zero as well. (The constant of integration is most naturally zero!)

Definition 5.1. *Suppose that for all M, if $(f_1 X_n - U_n)(M) = 0$ then $f_n(M) = 0$. In this case we say f_n is* balanced.

We now state and prove the result.

Theorem 5.2. *Let $\mathcal{Z}$ be an arc diagram and let M_i be nonzero homology classes of diagrams on $\mathcal{Z}$. Let $\mathcal{CY}$ be a cycle choice function for $\mathcal{Z}$, and let $\mathcal{CR}$ be a creation choice function. Then there is an A_∞ structure X on $\mathcal{H}(\mathcal{Z})$ extending its DGA structure, and a morphism of A_∞ algebras $f\colon \mathcal{H}(\mathcal{Z}) \longrightarrow \mathcal{A}(\mathcal{Z})$ extending the cycle selection map $f_1 = f^{\mathcal{CY}}$, such that the following conditions hold.*

(i) If $M = M_1 \otimes \cdots \otimes M_n$ is not viable, then $\overline{f}_n(M) = 0$ and $X_n(M) = 0$; and if M has an idempotent mismatch then $f_n(M) = 0$.

(ii) The maps $X_n\colon \mathcal{H}(\mathcal{Z})^{\otimes n} \longrightarrow \mathcal{H}(\mathcal{Z})$ of X, and the maps $f_n\colon \mathcal{H}(\mathcal{Z})^{\otimes n} \longrightarrow \mathcal{A}(\mathcal{Z})$ of f, all preserve H-data.

(iii) Each map f_n is balanced.

(iv) For $n \geq 2$, on each twisted H-summand, $f_n = A^*_{\mathcal{CR}} \circ (f_1 X_n - U_n)$, where U_n is defined by equation (2) and X_n is defined by equation (3):

$$U_n\left(M_1 \otimes \cdots \otimes M_n\right) = \sum_{j=1}^{n-1} f_j\left(M_1 \otimes \cdots \otimes M_j\right) f_{n-j}\left(M_{j+1} \otimes \cdots \otimes M_n\right) + \sum_{k=0}^{n-2} \sum_{j=2}^{n-1} f_{n-j+1}\left(M_1 \otimes \cdots \otimes M_k \otimes X_j\left(M_{k+1} \otimes \cdots \otimes M_{k+j}\right) \otimes \cdots \otimes M_n\right),$$

$$X_n = [U_n].$$

The A_∞-operations X_n satisfying these conditions are unique. The maps f_n are uniquely defined modulo $\mathcal{F}$.

Recall that when M is singular, there are no diagrams with its H-data (h, s, t), and $\mathcal{A}(h, s, t) = 0$ (lemma 4.25). So condition (i), that f_n and X_n preserve H-data, implies that when M is singular, $f_n(M)$ and $X_n(M)$ are both zero.

The uniqueness statement means that, although the f_n are not uniquely determined by the conditions of the theorem, after composing with the quotient map $\mathcal{A} \longrightarrow \overline{\mathcal{A}}$ to obtain $\overline{f}_n : \mathcal{H}^{\otimes n} \longrightarrow \overline{\mathcal{A}}$, the maps $\overline{f}_n$ *are* uniquely determined.

Recall from section 3.6 that a pair ordering $\preceq$ (definition 3.20) determines a cycle choice function $\mathcal{CY}^{\preceq}$ and a creation choice function $\mathcal{CR}^{\preceq}$ (definition 3.21). Hence we immediately obtain the following corollary.

Corollary 5.3. *Let $\preceq$ be a pair ordering on an arc diagram $\mathcal{Z}$. Then there is an A_∞ structure X on $\mathcal{H}$ extending its DGA structure, and a morphism of A_∞ algebras $f : \mathcal{H} \longrightarrow \mathcal{A}$ extending the cycle selection map $f^{\mathcal{CY}^{\preceq}}$, satisfying the conditions of theorem 5.2, such that on twisted summands for $n \geq 2$, $f_n = A^*_{\mathcal{CR}^{\preceq}} \circ (f_1 X_n - U_n)$.* □

Corollary 5.3 is a precise form of theorem 1.1.

Proof of theorem 5.2. We follow the method described above. At level 1, equations (2), (3) and (4) require $U_1 = 0$, $X_1 = 0$ and $\partial f_1 = 0$. The last equation is satisfied by $f_1 = f^{\mathcal{CY}}$. Since diagrams with non-viable H-data are zero in homology, $f_1 = 0$ for such diagrams.

Now suppose we have constructed all operations at level $< n$ as required; we construct U_n, X_n, f_n.

We define U_n by equation 2. Then U_n Maslov grading $n - 2$. As the f_j are not uniquely defined, neither is U_n. However, all the $\overline{f}_j$ are uniquely defined, and since equation 2 expresses U_n as a sum of products of values of f_j, $\overline{U}_n$ is also uniquely defined. Since the f_j (and multiplication in $\mathcal{A}$) preserve H-data, U_n does also.

We define X_n by equation 3; then X_n respects gradings as required. As in [8], U_n is a cycle and X_n is its homology class, so X_n is well defined. Now all diagrams in $\mathcal{F}$ are non-viable or have crossings, and such diagrams do not contribute to homology. Thus $X_n(M)$ is determined completely by $\overline{U}_n(M)$, which is uniquely defined; hence $X_n(M)$ is uniquely defined.

To define f_n, we must solve equation (4) for each viable $M = M_1 \otimes \cdots \otimes M_n$:

$$\partial f_n(M) = (f_1 X_n - U_n)(M).$$

We now consider various cases for M.

First, suppose $M = M_1 \otimes \cdots \otimes M_n$ is non-viable because of an idempotent mismatch. Then each term in $U_n(M)$ from (2) is zero: by induction, f_i and X_i for $i < n$ are zero on tensor products with mismatches, and the product of two mismatched diagrams is zero. Thus $U_n(M) = 0$, and by (3) then $X_n(M) = 0$. We set $f_n(M) = 0$ as required; equation (4) and the balanced condition are then satisfied.

Next, suppose M is non-viable but has no idempotent mismatch, hence has some step covered more than once. As U_n preserves H-data then $U_n(M)$ is a sum of non-viable diagrams, so $U_n(M) \in \mathcal{F}$ and $\overline{U}_n(M) = 0$. Then $X_n(M) = 0$, and equation (4) then requires $\partial f_n(M) = U_n(M)$. If $U_n(M) = 0$ then we set $f_n(M) = 0$, satisfying the balanced condition; otherwise we choose $f_n(M)$ arbitrarily to be any solution to this equation with the same H-data as M, and of pure Maslov grading (necessarily 1 greater than M). Then $f_n(M) \in \mathcal{F}$, being a sum of non-viable diagrams. Thus $f_n(M)$ is not uniquely determined, but $\overline{f}_n(M)$ is uniquely determined, indeed $\overline{f}_n(M) = 0$.

If M is singular, then as there are no diagrams with the H-data of M, $U_n(M)$, $X_n(M)$ and $f_n(M)$ are all zero, and all required conditions are satisfied.

So we may now assume M is viable and non-singular; its H-data (h, s, t) (definition 4.3) is thus tight or twisted (definition 2.15) .

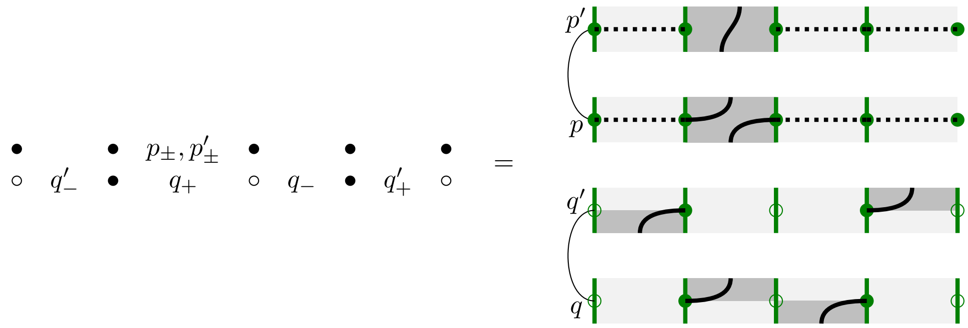

Figure 15: Shorthand notation for viable nonzero tensor products of homology classes of diagrams.

If (h, s, t) is twisted, then as required we take $f_n = A^*_{\mathcal{CR}} \circ (f_1 X_n - U_n)$. Then f_n is balanced. Since $(f_1 X_n - U_n)(M)$ is a boundary, hence a cycle (in fact boundaries and cycles are the same since $\mathcal{H}(h, s, t) = 0$), we have

$$\partial f_n(M) = \partial A^*_{\mathcal{CR}} \left(f_1 X_n(M) - U_n(M) \right) = \left(f_1 X_n - U_n \right)(M),$$

inverting the differential, by lemma 3.11 and the discussion of section 3.5.

If (h, s, t) is tight, by lemma 3.15, any $f_n(M)$ satisfying $\partial f_n(M) = f_1 X_n(M) - U_n(M)$ lies in $\mathcal{F}$. We choose $f_n(M)$ to be zero if $f_1 X_n - U_n = 0$ (satisfying the balanced condition), otherwise to be any solution to this equation with the same H-data as M, and pure Maslov grading. Then $f_n(M)$ is not uniquely determined, but $\overline{f}_n(M) = 0$.

This defines f_n and X_n satisfying the required conditions, with the uniqueness claimed. Having followed Kadeishvili's construction, the X_n form an A_∞ structure on $\mathcal{H}$, and the f_n form a morphism of A_∞ algebras $\mathcal{H} \longrightarrow \mathcal{A}$. □

5.2 Shorthand notation

For convenience, we use some shorthand for viable nonzero tensor products in $\mathcal{A}^{\otimes n}$ and $\mathcal{H}^{\otimes n}$. The shorthand is essentially a stylised version of our previous diagrams.

Let $M = M_1 \otimes \cdots \otimes M_n$ be a viable tensor product of nonzero homology classes of diagrams on an arc diagram $\mathcal{Z}$. A shorthand diagram represents M by an array of data. Each row refers to a matched pair P of $\mathcal{Z}$. The n columns refer to $M_1, \ldots, M_n$. In the row for $P = \{p, p'\}$ and the column of M_i, we write which of the four steps of $\mathcal{Z}_P$ are covered by M_i. Along the row for P, between the columns we draw a hollow or solid circle indicating whether P is contained in the corresponding idempotent ("on or off"). This is well defined since M is viable.

Such notation specifies M completely, since it specifies the H-data of each M_i.

We denote the steps before and after a place p by p_- and p_+ respectively. We write $p_\pm$ to indicate that both p_+ and p_- are covered.

Figure 15 shows an example. As in previous examples, we usually only write the situation at certain matched pairs. It is always possible to include further matched pairs, each with a tight tensor product of homology classes of diagrams, so as to obtain a tensor product of homology classes of non-augmented diagrams on a connected arc diagram. In any case, we always assume the tensor product is tight at any matched pairs not shown.

Occasionally, when the idempotents can be inferred from the H-grading of each M_i, we omit the circles in the notation.

We use a similar notation to write elements of $\overline{\mathcal{A}}^{\otimes n}$. Let $D = D_1 \otimes \cdots \otimes D_n$ be a viable tensor product of diagrams on $\mathcal{Z}$, which does not lie in $\mathcal{F}$. Again we write an array where each row refers to a matched pair P, and each column refers to a D_i, and in each row we have hollow or solid circle to signify idempotents. However, a diagram is not always specified by its H-data, so we use some of the notation of definition 2.12. When there is a unique tight local diagram with the H-data, we simply write which steps are covered. Otherwise, we use the notation g_p, c_p, w_p for 11 doubly occupied tight, 11 once occupied crossed, and 11 once occupied twisted pairs. (As we compute in $\overline{\mathcal{A}}$, all-on doubly occupied crossed diagrams do not arise.)

This shorthand notation is useful for A_∞ operations defined by a pair ordering, as in corollary 5.3. We may then simply order the pairs upwards in our array (just as they are ordered along the intervals of $\mathbf{Z}$). A creation operator then always applies at the all-on once occupied pair which is lowest in our shorthand notation.

We adopt notation where each matched pair is denoted by a capital letter, and its two places by the corresponding lower case letter, the latter under $\preceq$ being primed. Thus we always write pairs as $P = \{p, p'\}$, $Q = \{q, q'\}$, etc., where $p \prec p'$, $q \prec q'$, etc. Then a cycle choice function always selects a cycle with strands at a place with an unprimed label.

When a tensor product $M = M_1 \otimes \cdots \otimes M_n$ is twisted at a place p of a pair $P = \{p, p'\}$, it is 11 once occupied at P, with p fully occupied (proposition 4.23 and table 2), and the two steps p_+, p_- are covered by some M_i and M_j, with $i < j$. Thus, in our shorthand, across the row corresponding to P we see p_+ in one column, then p_- in another column, in that order.

Similarly, if P is critical, then it is sesqui-occupied or doubly occupied. Looking across the row corresponding to P we see one of the following sequences, appearing in order, in distinct columns (possibly after relabelling p and p'):

- Pre-sesqui-occupied: p'_-, p_+, p_-
- Post-sesqui-occupied: p_+, p_-, p'_+
- 00 doubly occupied: p_-, p'_+, p'_-, p_+
- 11 doubly occupied: p'_+, p'_-, p_+, p_-

5.3 Low-level maps

Now that we have constructed A_∞ structures on $\mathcal{H}$, we consider low-level maps explicitly. We assume A_∞ structures are constructed from a cycle choice function $\mathcal{CY}$ and a creation choice function $\mathcal{CR}$, as in theorem 5.2.

Level 1 maps are straightforward ($X_1 = 0$ and $f_1 = f^{\mathcal{CY}}$), as is multiplication X_2. We consider $\overline{f}_2$.

Let $M = M_1 \otimes M_2$ be a tensor product of nonzero homology classes of diagrams. By theorem 5.2, if M is non-viable or singular, then $f_n(M)$ and $X_n(M)$ are both zero; so we assume M is viable and non-singular. The H-data (h, s, t) of M is then tight or twisted.

When (h, s, t) of M is tight, $\overline{f}_2(M) = 0$. So suppose (h, s, t) is twisted, hence has at least one all-on once occupied pair. Clearly M is then not tight; in fact, M cannot be critical either, since it takes 3 to be critical (lemma 4.28). So M is twisted, hence is tight or twisted at each matched pair (lemma 4.26). In particular, M is twisted at each all-on once occupied pair, and tight at each other pair.

Since $U_2(M_1 \otimes M_2) = f_1(M_1) f_1(M_2)$, we have

$$f_2(M_1 \otimes M_2) = A^*_{\mathcal{CR}} \Big(f_1(M_1 M_2) + f_1(M_1) f_1(M_2) \Big). \tag{8}$$

Since M is twisted, $M_1 M_2 = 0$ and $f_1(M_1) f_1(M_2) \neq 0$ (lemma 4.25). The nonzero product $f_1(M_1) f_1(M_2)$ is clearly not tight, and it is also not crossed (being the product of two crossingless diagrams: lemma 2.17), hence it is twisted, hence tight or twisted at each pair (lemma 2.10). Indeed, at each all-on once occupied pair $f_1(M_1) f_1(M_2)$ cannot be tight, so must be twisted; and at each other pair, it must be tight. We then have

$$f_2(M_1 \otimes M_2) = A^*_{\mathcal{CR}^{\preceq}} \left(f_1(M_1) f_1(M_2) \right), \tag{9}$$

so f_2 adds a crossing at the $\preceq$-minimal all-on once occupied pair of $f_1(M_1) f_1(M_2)$.

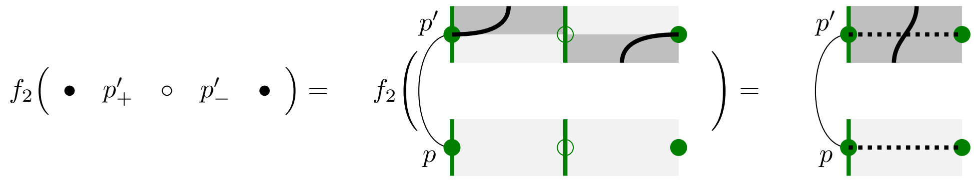

Figure 16: The effect of f_2.

Thus, $f_2(M_1 \otimes M_2)$ is given by a single diagram, which is crossed at the $\preceq$-minimal all-on once occupied pair of $M_1 \otimes M_2$, and is elsewhere given by $f_1(M_1) f_1(M_2)$. The idea is shown in figure 16. In this way, f_2 turns one pair from twisted to crossed.

6 Properties of A-infinity structures

We now consider A_∞ structures on $\mathcal{H}$ constructed by Kadeishvili's method in general — not just those defined by the construction in theorem 5.2, using a cycle choice function and a creation choice function.

Throughout this section, then, we consider an A_∞ structure X on $\mathcal{H}$, with operations $X_n \colon \mathcal{H}^{\otimes n} \longrightarrow \mathcal{H}$, and a morphism of A_∞ algebras $f \colon \mathcal{H} \longrightarrow \mathcal{A}$, with maps $f_n \colon \mathcal{H}^{\otimes n} \longrightarrow A$, constructed by Kadeishvili's method. So there are also auxiliary maps $U_n \colon \mathcal{H}^{\otimes n} \longrightarrow \mathcal{A}$, and we assume the maps U_n, X_n, f_n satisfy equations (2), (3) and (4). We assume all U_n, X_n, f_n preserve H-data and have Maslov degree $n-2$, $n-2$, $n-1$ respectively.

Thus, each f_n is defined by inverting the differential in the equation $\partial f_n = f_1 X_n - U_n$, but not necessarily by inverting it with a creation operator.

As usual, throughout this section, $M = M_1 \otimes \cdots \otimes M_n$ denotes a tensor product of nonzero homology classes of diagrams.

6.1 Non-viable input

We have already seen that, if a tensor product of homology classes of diagrams $M_1 \otimes \cdots \otimes M_n$ is not viable, then their product is zero (lemma 4.7). We now show that other operations are zero as well.

Lemma 6.1. *If all f_n are balanced, and M is not viable, then $\overline{f}_n(M)$ and $X_n(M)$ are both zero. More precisely, we have the following.*

(i) If M is not viable because some step is covered twice, then $\overline{f}_n(M)$ and $X_n(M)$ are both zero.

(ii) If all f_n are balanced, and M is not viable because of an idempotent mismatch, then $f_n(M)$ and $X_n(M)$ are both zero.

Proof. First suppose M has some step covered twice. As X_n preserves H-grading, and there are no tight diagrams with such H-grading, $X_n(M) = 0$. As f_n preserves H-grading, $f_n(M)$ is a sum of non-viable diagrams, hence lies in $\mathcal{F}$, so $\overline{f}_n(M) = 0$.

We now show (ii) by induction on n. When $n = 1$ there is nothing to prove. Suppose the result is true for all f_i with $i < n$; now assume M is mismatched, and we will show $f_n(M) = 0$.

Consider $U_n(M)$. In a term of the form $f_j(M_1 \otimes \cdots \otimes M_j) f_{n-j}(M_{j+1} \otimes \cdots \otimes M_n)$, if the mismatch occurs within $M_1 \otimes \cdots M_j$ or $M_{j+1} \otimes \cdots \otimes M_n$ then by induction the term is zero; otherwise it occurs between M_j and M_{j+1}, in which case the product is zero. In a term of the form $f_{n-j+1}(M_1 \otimes \cdots \otimes M_k \otimes X_j(M_{k+1} \otimes \cdots \otimes M_{k+j}) \otimes \cdots \otimes M_n)$, if the mismatch occurs within $M_{k+1} \otimes \cdots \otimes M_{k+j}$ then the X_j term is zero, hence the whole term is zero by linearity; otherwise it occurs within the f_{n-j+1} term and again we have zero. Thus $U_n(M) = 0$, so $X_n(M) = 0$. Then $\partial f_n(M) = (f_1 X_n - U_n)(M) = 0$, and since f_n is balanced then $f_n(M) = 0$ also. □

6.2 Equivalent choices of maps

In the proof of theorem 5.2, we saw that although there might be many choices available for the f_n on tight summands, such choices had no effect on the resulting X_n.

In a similar vein, we now show that, in applying Kadeishvili's construction in general (i.e. without creation operators), the choices available for the maps $\overline{f}_n$ do not depend on any previous choices.

Lemma 6.2. *Suppose that f_i are defined for all $i < n$, U_i and X_i are defined for all $i \leq n$, and that two functions $a, b \colon H^{\otimes n} \longrightarrow A$ satisfy*

$$\overline{a} = \overline{b} \quad \text{and} \quad \partial a = \partial b = f_1 X_n - U_n.$$

Whether we choose $f_n = a$ or b, for all $N > n$, the choices for each $\overline{f}_N$ are identical.

Let us be more explicit. Taking $f_n = a$ we define U and X maps at level $n+1$, which we denote U^a_{n+1}, X^a_{n+1}. Then we have a set of choices $\mathcal{S}^a_{n+1} = \{\overline{f} \mid \partial f = f_1 X^a_{n+1} - U^a_{n+1}\}$ for $\overline{f}_{n+1}$. On the other hand, taking $f_n = b$ we define U^b_{n+1}, X^b_{n+1} and have another set of choices $\mathcal{S}^b_{n+1} = \{\overline{f} \mid \partial f = f_1 X^b_{n+1} - U^b_{n+1}\}$ for $\overline{f}_{n+1}$. Lemma 6.2 says that $\mathcal{S}^a_{n+1} = \mathcal{S}^b_{n+1}$. Moreover, after taking $f_n = a$ and making arbitrary choices $f^a_{n+1}, \ldots, f^a_{N-1}$ using Kadeishvili's construction, obtaining maps $U^a_{n+1}, X^a_{n+1}, \ldots, U^a_N, X^a_N$, we obtain a set of choices $\mathcal{S}^a_N = \{\overline{f} \mid \partial f = f_1 X^a_N - U^a_N\}$ for $\overline{f}_N$; after taking $f_n = b$ and making arbitrary choices $f^b_{n+1} \ldots, f^b_{N-1}$ and obtaining maps $U^b_{n+1}, X^b_{n+1}, \ldots, U^b_N, X^b_N$, we have choices $\mathcal{S}^b_N = \{\overline{f} \mid \partial f = f_1 X_N - U_N\}$ for $\overline{f}_N$. Lemma 6.2 says, more generally, that these too are equal: $\mathcal{S}^a_N = \mathcal{S}^b_N$.

Proof. Let M have H-data (h, s, t) (definition 4.3).

When (h, s, t) is not viable (i.e. covers some step at least twice: definition 2.1), there is only one choice for $\overline{f}_n(M)$, namely 0, by lemma 6.1. And when (h, s, t) is singular, $\overline{f}_n(M) = 0$ as there are no available diagrams. Hence we need only consider $\overline{f}_n(M)$ when (h, s, t) is viable and non-singular, hence tight or twisted (proposition 2.14, definition 2.15).

Since $\overline{a} = \overline{b}$, $a - b$ takes values in $\mathcal{F}$. As $\mathcal{F}$ is an ideal, $U^a_{n+1}(M)$ and $U^b_{n+1}(M)$ differ by values in $\mathcal{F}$. Diagrams in $\mathcal{F}$ do not contribute to homology, as they have crossings, so $[U^a_{n+1}(M)] = [U^b_{n+1}(M)]$. It follows that $X^a_{n+1}(M) = X^b_{n+1}(M)$; we simply write $X_{n+1}(M)$ in either case. Moreover, $U^a_{n+1}(M) - U^b_{n+1}(M)$ is a boundary; let $U^a_{n+1}(M) - U^b_{n+1}(M) = \partial g_{n+1}$. As $U^a_{n+1}(M) - U^b_{n+1}(M) \in \mathcal{F}$, there is such a g_{n+1} in $\mathcal{F}$: if (h, s, t) is twisted we can use a creation operator; if (h, s, t) is tight then any diagram with crossings lies in $\mathcal{F}$. Then to define $f_{n+1}(M)$ we must solve

$$f_1 X_{n+1}(M) - U^a_{n+1}(M) = \partial f^a_{n+1}(M) \quad \text{or} \quad f_1 X_{n+1}(M) - U^b_{n+1}(M) = \partial f^b_{n+1}(M).$$

Now observe that $f^a_{n+1}(M)$ is a solution of the first equation iff $f^a_{n+1}(M) + g_{n+1}$ is a solution of the second equation. As $g_{n+1} \in \mathcal{F}$, the possible $\overline{f}^a_{n+1}(M)$ and $\overline{f}^b_{n+1}(M)$ are identical.

Thus, the possible choices for f_{n+1} differ by values in $\mathcal{F}$. The possible choices for U_{n+2} then differ by values in $\mathcal{F}$, and the argument proceeds by induction, giving the desired result. $\square$

Thus in the construction of the maps f_n and X_n, it is sufficient to consider $\overline{f}$ rather than f at each level. So we may effectively compute in $\mathcal{A}/\mathcal{F} = \overline{\mathcal{A}}$.

6.3 First properties of A-infinity operations

We now have some preliminary results giving some description of $\overline{f}_n$ and X_n.

Lemma 6.3. *Suppose all f_k are balanced. For any $n \geq 2$ and $M = M_1 \otimes \cdots \otimes M_n$, $f_n(M)$ is a sum of crossed diagrams.*

This includes a sum of no crossed diagrams, when $f_n(M) = 0$.

Proof. Recall f_n preserves H-data and has Maslov grading $n-1 \geq 1$. For fixed H-data, Maslov grading is (up to a constant) given by the number of matched pairs with crossings (section 2.7). So each diagram in $f_n(M)$ must contain at least one crossing. □

Lemma 6.4. *Suppose all f_k are balanced. Then $X_n(M)$ is represented by the sum of all crossingless diagrams in the sum*

$$\sum_{j=1}^{n-1} \overline{f}_j(M_1 \otimes \cdots \otimes M_j)\overline{f}_{n-j}(M_{j+1} \otimes \cdots \otimes M_n),$$

where we write elements of $\overline{\mathcal{A}}$ in standard form.

Recall (definition 2.25(v)) that the standard form of an element of $\overline{\mathcal{A}}$ is a sum of viable diagrams without crossed doubly occupied pairs. For $n = 1$ the result reduces to $X_1 = 0$.

Proof. By construction, $X_n(M)$ is the homology class of $U_n(M)$. Consider the terms of (2) defining $U_n(M)$. Terms of the form $f_\bullet(M_1 \otimes \cdots \otimes X_\bullet(\cdots) \otimes \cdots \otimes M_n)$ only contain crossed diagrams (lemma 6.3) hence do not contribute to homology. Thus, $X_n(M)$ is represented by the sum of tight diagrams of the form $f_j(M_1 \otimes \cdots \otimes M_j) f_{n-j}(M_{j+1} \otimes \cdots \otimes M_n)$. Since diagrams in the ideal $\mathcal{F}$ do not contribute to homology, X_n is represented by the sum claimed in standard form. □

Lemma 6.4 allows us to calculate $X_n(M)$ directly from $\overline{f}_j$ and $\overline{f}_{n-j}$. Diagrams in $\overline{f}_j$ or $\overline{f}_{n-j}$ usually contain crossings (lemma 6.3), but the crossings may multiply out to give a tight result. Such $\overline{f}_j \otimes \overline{f}_{n-j}$ are sublime; sublimation is therefore ubiquitous in the operations X_n, arising in any nonzero $X_n(M)$.

6.4 Conditions for nontrivial A-infinity operations

We now find some necessary conditions for $\overline{f}_n$ or X_n to be nonzero.

Theorem 6.5. *Suppose all f_k are balanced. Let $n \geq 2$, let $M_1, \ldots, M_n$ be nonzero homology classes of tight diagrams, and let $M = M_1 \otimes \cdots \otimes M_n$. If $\overline{f}_n(M) \neq 0$, then the following statements hold.*

(i) In M there are l matched pairs which are twisted, and m matched pairs which are critical, where $l + m \geq n - 1$ and $m \leq n - 2$. All other matched pairs are tight.

(ii) $\overline{f}_n(M)$ is represented by a sum of diagrams, where each diagram D satisfies the following conditions.

(a) All m of the critical matched pairs in M become tight in D.

(b) Precisely $n - m - 1$ of the l twisted matched pairs in M become crossed in D; the other $l + m - n + 1$ twisted matched pairs in M remain twisted in D.

(c) All tight matched pairs in M remain tight in D.

Proof. First, by lemma 6.1, M is viable.

Write $\overline{f}_n(M)$ in standard form (definition 2.25(v)), as a sum of distinct diagrams without crossed doubly occupied pairs. Let D be one of these diagrams. As $\overline{f}_n$ respects H-grading and has Maslov grading $n-1$, D is viable, with $h(D) = h(M)$, and $\iota(D) = \iota(M) + n - 1$. From tables 1 and 2, at each matched pair the Maslov grading can increase by at most 1; hence there are precisely $n-1$ matched pairs at which D has a higher Maslov index than M.

At each matched pair P, D must give a viable local diagram which respects local H-data. There are no such diagrams for singular pairs; hence M has no critical matched pairs. Thus all matched pairs of M are tight, twisted, or critical.

Consider a matched pair P where M is critical. From table 2, P is sesqui-occupied or doubly occupied by M. Every all-on doubly occupied pair must remain uncrossed in D. From table 1, any viable local diagram at a sesqui-occupied or doubly occupied matched pair, which is not crossed all-on doubly occupied, must be tight. So D_P is tight at P, and $\iota(D_P) = \iota(M_P) + 1$.

Now consider a matched pair P where M is tight. Then the local H-data of M at P is tight. We observe from table 1 that, with crossed all-on doubly occupied local diagrams ruled out, any viable local diagram with tight H-data must be tight. Thus D_P is tight, and hence $\iota(D_P) = \iota(M_P)$.

Now $\iota(D) = \iota(M) + n - 1$, and m of this increase is accounted for at critical matched pairs. The remaining increase of $n-1-m$ must arise at the l pairs where M is twisted. From table 2, we observe that these are precisely the pairs where M is all-on once occupied. At such pairs, two viable local diagrams are possible: a tight and a crossed diagram. Crossings can thus be inserted at such pairs to increase the Maslov index; they must be inserted at $n-1-m$ such pairs for D to have the correct Maslov index; so $n-1-m \leq l$. The remaining $l+m-n+1$ pairs must remain twisted in D.

The diagram D thus has precisely $n-m-1$ crossings. But by lemma 6.3, D must have at least one crossing. Thus $n-m-1 \geq 1$. □

Theorem 6.6. *Suppose all f_k are balanced. If $X_n(M) \neq 0$, then the following statements hold.*

(i) M has precisely $n-2$ critical matched pairs, and all other matched pairs are tight.

(ii) $X_n(M)$ is the unique homology class of tight diagram with the H-data of M.

For $n=1$ this result just says $X_1 = 0$; for $n=2$, X_2 being multiplication, it follows from lemmas 4.25 and 4.26.

Proof. By lemma 6.1, M is viable. Since X_n respects H-data, let M and $X_n(M)$ have H-data (h,s,t). Since there is at most one nonzero homology class with fixed H-data, (ii) follows immediately.

As $X_n(M) \in \mathcal{H}(h,s,t)$ is nonzero, the H-data (h,s,t) is tight (proposition 2.14; definition 2.15), and thus tight at each matched pair (lemma 2.16). In particular, M has no 11 once occupied or 00 alternately occupied pairs. From proposition 4.23 and table 2, M is tight or critical at each pair.

When M is tight at a pair P, $M_1 \cdots M_n$ is tight at P (lemma 4.25). As $X_n(M)$ is given at P by the unique tight diagram with the H-data of M then $X_n(M)_P = (M_1 \cdots M_n)_P$. In this case $X_n(M)$ has the same Maslov index as M at P.

On the other hand, when M is critical at P, $X_n(M)$ must still be given at P by the unique tight diagram with the same H-data. Inspecting table 2 (and recalling that multiplication in $\mathcal{H}$ has zero Maslov grading), we observe that $\iota(X_n(M)_P) = \iota(M_P) + 1$.

Since Maslov index is additive over matched pairs (section 2.4), the difference in Maslov indices is given by the number of matched pairs at which M is critical. Since X_n has Maslov index $n-2$, M has precisely $n-2$ critical pairs. □

When the maps f_n and X_n are defined from a pair ordering, as in corollary 5.3 or theorem 1.1, theorems 6.5 and 6.6 respectively yield parts (i) and (ii) of theorem 1.2.

Now we show that it's not possible to have $\overline{f}_n$ nonzero simultaneously with X_n, and more.

Lemma 6.7. *Suppose that for all $k \geq 1$, f_k is balanced. Let $n \geq 2$. If $X_n(M) \neq 0$ or $M_1 \cdots M_n \neq 0$ then $\overline{f}_n(M) = 0$.*

Contrapositively, if $\overline{f}_n(M) \neq 0$ then $X_n(M) = 0$ and $M_1 \cdots M_n = 0$.

Proof. First suppose $X_n(M) \neq 0$. Since X_n respects H-data, $X_n(M)$ is the unique nonzero homology class with the H-data of M. Each of $f_1 X_n(M)$ and $U_n(M)$ is the sum of an odd number of diagrams representing this homology class. Thus $\partial f_n(M) = (f_1 X_n - U_n)(M)$ is the sum of an even number of tight diagrams all representing $X_n(M)$. By lemma 3.6(ii) then $f_n(M) \in \mathcal{F}$, so $\overline{f}_n(M) = 0$.

Now suppose $M_1 \cdots M_n \neq 0$. Then M is tight, hence tight at every matched pair (lemmas 4.25, 4.26). Now $f_n(M)$ has the same H-data as $M_1 \cdots M_n$, and Maslov grading greater by $n-1$. Hence each diagram in $f_n(M)$ has crossings at $n-1 \geq 1$ pairs (section 2.7); but since its H-data is tight, inspecting table 1, each crossing is at an all-on doubly occupied pair. Thus $f_n(M) \in \mathcal{F}$ and $\overline{f}_n(M) = 0$. □

Theorem 6.6 places stringent necessary conditions on a tensor product $M_1 \otimes \cdots \otimes M_n$ to yield a nonzero result under X_n. However, we will see in section 7.2 that these conditions are not sufficient.

6.5 Levels 1, 2 and 3 in general

We now describe some properties of the maps f_n, X_n at levels 1, 2 and 3. Unlike the discussion of section 5.3, we do not assume the A_∞ structure derives from creation operators, as in section 5 and theorem 5.2 and corollary 5.3. We assume, as throughout the present, that Kadeishvili's construction is used, and the maps U_n, X_n, f_n preserve H-data and have appropriate Maslov gradings. We will additionally now assume that the f_n are balanced.

Level 1 is again straightforward. By construction $X_1 = 0$, and f_1 is a cycle selection homomorphism. If f_1 is diagrammatically simple (definition 3.2) then it arises from a cycle choice function (lemma 3.5). In general, for each tight (h, s, t), f_1 maps $M_{h,s,t}$ to the sum of an odd number of tight diagrams representing $M_{h,s,t}$.

Now consider level 2; let $M = M_1 \otimes M_2$. By construction X_2 is multiplication. As for f_2, if $\overline{f}_2(M) \neq 0$, then M is viable (lemma 6.1), and theorem 6.5 says that M has no critical matched pairs (impossible at level 2 in any case: lemma 4.28), and at least one twisted (i.e. all-on once occupied) matched pair. Moreover, $\overline{f}_2(M)$ is represented by a sum of diagrams, each of which has precisely one crossing at a twisted matched pair. Also $M_1 M_2 = 0$ (from the twisted matched pair, or lemma 6.7).

Conversely, suppose $M = M_1 \otimes M_2$ is viable and has at least one twisted matched pair. Then $X_2(M) = 0$, and $U_2(M) = f_1(M_1) f_1(M_2)$ is the sum of an odd number of diagrams, each tight and twisted at the same matched pairs as M; the diagrams differ only by strand switching at all-on doubly occupied pairs (section 2.11). Since $\partial f_2(M) = f_1(M_1) f_1(M_2)$ and f_2 respects H-data and has Maslov grading 1, $f_2(M)$ is a sum of diagrams, each of which has a crossing at precisely one matched pair, and at each other pair is tight or twisted in agreement with M. In standard form (definition 2.25) $\overline{f}_2(M)$ is then given omitting diagrams with crossings at doubly occupied pairs, so that each crossing is at a matched pair where M is twisted. The differential of each remaining diagram is a single diagram, but the differential of each omitted diagram is a sum of two diagrams, so $\overline{f}_2(M)$ in standard form is the sum of an odd number of diagrams. We have now proved the following.

Lemma 6.8. *Suppose that f_1 and f_2 are balanced. Then $\overline{f}_2(M) \neq 0$ if and only if M is viable and has at least one twisted matched pair. Then $\overline{f}_2(M)$ in standard form is the sum of an odd number of diagrams, each with a single crossing at a matched pair where M is twisted, and elsewhere tight or twisted in agreement with M.* □

When the A_∞ structure is defined by creation operators, as in section 5 and theorem 5.2, then any nonzero $f_2(M)$ is a single diagram, given by equation (8) or (9) from section 5.3.

We now consider X_3; the case is illustrative, showing the role of critical and sublime tensor products of diagrams. Let $M = M_1 \otimes M_2 \otimes M_3$ and suppose $X_3(M) \neq 0$. Then M is viable (lemma 6.1) and by theorem 6.6, M has precisely one critical matched pair, and all other matched pairs are tight. By lemma 6.4, $X_3(M)$ is represented by the sum of all crossingless diagrams in

$$\overline{f}_1(M_1)\overline{f}_2(M_2 \otimes M_3) + \overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3).$$

Each diagram in an $\overline{f}_2(M_i \otimes M_{i+1})$ term has a crossing at precisely one matched pair P, where $M_i \otimes M_{i+1}$ is twisted; since such P cannot be tight in M (lemma 4.29, table 4), P is the critical matched pair of M. Multiplying this diagram by the third M_j must then produce a tight diagram. There are two cases for the tightness of the various tensor products:

- $M_1 \otimes M_2$ twisted; each diagram D in $\overline{f}_2(M_1 \otimes M_2)$ crossed; M_3 and each diagram D' in $\overline{f}_1(M_3)$ tight; each $D \otimes D'$ sublime; and $M_1 \otimes M_2 \otimes M_3$ critical.
- $M_2 \otimes M_3$ twisted; each diagram D' in $\overline{f}_2(M_2 \otimes M_3)$ crossed; M_1 and each diagram D in $\overline{f}_1(M_1)$ tight; each $D \otimes D'$ sublime; and $M_1 \otimes M_2 \otimes M_3$ critical.

  In this case, the situation at P is shown in figure 17.

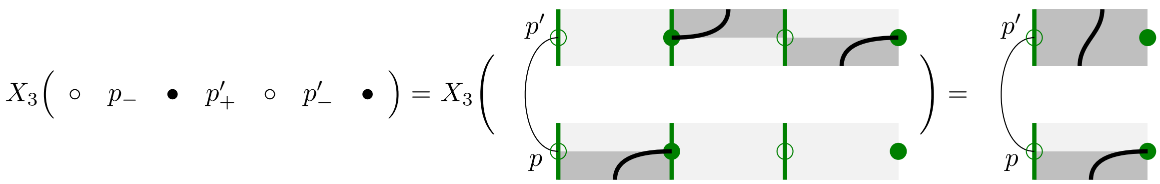

Figure 17: An example of $X_3(M_1 \otimes M_2 \otimes M_3)$, where $M_1 \otimes M_2 \otimes M_3$ is critical, $M_1 \otimes M_2$ is singular, and $M_2 \otimes M_3$ is twisted. Moreover, $\overline{f}_2(M_2 \otimes M_3)$ is crossed, and $\overline{f}_1(M_1) \otimes \overline{f}_2(M_2 \otimes M_3)$ is sublime.

We note that these two cases are mutually exclusive: only one of the terms $f_2(M_1 \otimes M_2)f_1(M_3)$ or $f_1(M_1)f_2(M_2 \otimes M_3)$ can be nonzero. For instance, in the first case $M_2 \otimes M_3$ is singular, and in the second case $M_1 \otimes M_2$ is singular.

Thus, to obtain a nonzero result for X_3, we start with a critical tensor product $M_1 \otimes M_2 \otimes M_3$; then a twisted sub-tensor-product (i.e. $M_2 \otimes M_3$ in figure 17) combines via f_2 into a crossed diagram, yielding a sublime tensor product (i.e. $f_1(M_1) \otimes f_2(M_2 \otimes M_3)$ in figure 17); and then these are multiplied to give a tight result. This process, with tensor products at a matched pair progressing from critical, to twisted, to sublime, to tight, is the process depicted in figure 6; it occurs generally in Kadeishvili's construction, without any need for creation operators.

We can prove a converse, and give necessary and sufficient conditions for $X_3 \neq 0$.

Proposition 6.9. *Suppose f_1 and f_2 are balanced. Then $X_3(M)$ is nonzero if and only if M is viable, critical at precisely one matched pair P, and tight at all other matched pairs.*

Proof. We only need prove that if the conditions on M hold, then $X_3(M) \neq 0$. By lemma 4.27, M_P is an extension of a tensor product shown in the critical column of table 2; but M_P has 3 factors, so M_P is exactly the critical 01 pre-sesqui-occupied or 10 post-sesqui-occupied tensor product shown there.

In the first case $\overline{f}_1(M_1)$ is tight and $\overline{f}_2(M_2 \otimes M_3)$ is a single diagram crossed at P. (From above, $\overline{f}_2(M_1 \otimes M_2)$ is represented by a sum of diagrams, each of which has precisely one crossing at an all-on once occupied matched pair P'. But then P' cannot be tight in M, by lemma 4.29 and table 4. So $P' = P$. Hence only one diagram is possible; so $\overline{f}_2(M_1 \otimes M_2)$ is given by precisely this diagram.) Then $\overline{f}_1(M_1) \otimes \overline{f}_2(M_2 \otimes M_3)$ is sublime, so $\overline{f}_1(M_1)\overline{f}_2(M_2 \otimes M_3)$ is tight. Moreover, in this case $M_1 \otimes M_2$ is singular so $\overline{f}_2(M_1 \otimes M_2) = 0$.

Similarly, in the second case $\overline{f}_2(M_1 \otimes M_2)$ is crossed, $\overline{f}_1(M_1)$ is tight, $\overline{f}_2(M_1 \otimes M_2) \otimes \overline{f}_1(M_3)$ is sublime, and $\overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3)$ is tight. Further $\overline{f}_2(M_2 \otimes M_3) = 0$.

In either case, by lemma 6.4 we have precisely one tight diagram rising in $\overline{f}_1(M_1)\overline{f}_2(M_2 \otimes M_3) + \overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3)$, which must represent $X_3(M)$. Thus $X_3(M) \neq 0$. ☐

Thus, if there are sufficiently few critical matched pairs in M, we may be able to guarantee that $X_n(M) \neq 0$. In section 8 we give some results in this direction, giving sufficient conditions for X_n and $\overline{f}_n$ to be nonzero.

7 Further examples and computations

We now calculate some further examples and prove some further results, for low-level A_∞ maps.

Recall the previous calculations along these lines. In section 5.3 we discussed the operations f_n and X_n for $n \leq 2$, when A_∞ operations are defined by cycle choice and creation choice functions, as in theorem 5.2. And in section 6.5 we again discussed low-level maps, especially $\overline{f}_2$ and X_3, for A_∞ operations obtained more generally using Kadeishvili's method.

In this section we consider A_∞ operations defined by a pair ordering $\preceq$, as in corollary 5.3, and consider maps at level 3 and 4, using the shorthand notation of section 5.2.

$$
\begin{aligned}
&U_3\begin{pmatrix} \bullet & p'_+ & \circ & p'_- & \bullet & & \bullet \\ \bullet & & \bullet & q'_+ & \circ & q'_- & \bullet \end{pmatrix} \\
&\quad = \overline{f}_1\begin{pmatrix} \bullet & p'_+ & \circ \\ \bullet & & \bullet \end{pmatrix} \overline{f}_2\begin{pmatrix} \circ & p'_- & \bullet & & \bullet \\ \bullet & q'_+ & \circ & q'_- & \bullet \end{pmatrix} + \overline{f}_2\begin{pmatrix} \bullet & p'_+ & \circ & p'_- & \bullet \\ \bullet & & \bullet & q'_+ & \circ \end{pmatrix} \overline{f}_1\begin{pmatrix} \bullet & & \bullet \\ \circ & q'_- & \bullet \end{pmatrix} \\
&\quad = \begin{pmatrix} \bullet & p'_+ & \circ \\ \bullet & & \bullet \end{pmatrix}\begin{pmatrix} \circ & p'_- & \bullet \\ \bullet & c_{q'} & \bullet \end{pmatrix} + \begin{pmatrix} \bullet & c_{p'} & \bullet \\ \bullet & q'_+ & \circ \end{pmatrix}\begin{pmatrix} \bullet & & \bullet \\ \circ & q'_- & \bullet \end{pmatrix} \\
&\quad = \begin{pmatrix} \bullet & w_{p'} & \bullet \\ \bullet & c_{q'} & \bullet \end{pmatrix} + \begin{pmatrix} \bullet & c_{p'} & \bullet \\ \bullet & w_{q'} & \bullet \end{pmatrix}
\end{aligned}
$$

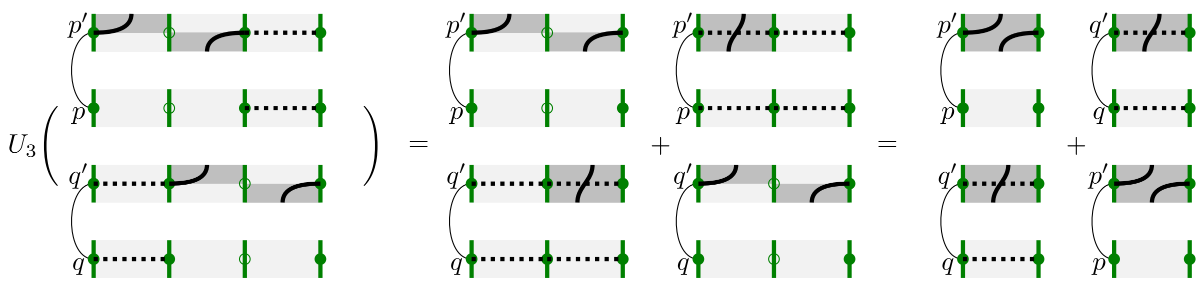

Figure 18: An example of $U_3(M_1 \otimes M_2 \otimes M_3)$, where $M_1 \otimes M_2 \otimes M_3$ is twisted. In this case both $M_1 M_2$ and $M_2 M_3$ are zero, and the two terms $f_1(M_1) f_2(M_2 \otimes M_3)$ and $f_1(M_1 \otimes M_2) f_1(M_3)$ both contribute to U_3. The sub-tensor-products $M_1 \otimes M_2$ and $M_2 \otimes M_3$ are also twisted, and the two nonzero terms of U_3 respectively apply an f_2 to insert a crossing in each.

As always, let $M = M_1 \otimes \cdots \otimes M_n$ denote a tensor product of nonzero homology classes of diagrams. We assume M is viable, necessary for nonzero results (lemma 6.1); let M have H-data (h, s, t). We work with $\overline{f}_n$ and $\overline{U}_n$, which loses no generality for calculating X_n (lemma 6.2).

7.1 Level 3

Consider the operation $\overline{U}_3$, given by

$$\overline{U}_3(M) = \overline{f}_1(M_2)\overline{f}_2(M_2 \otimes M_3) + \overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3) + \overline{f}_2(M_1 M_2 \otimes M_3) + \overline{f}_2(M_1 \otimes M_2 M_3).$$

As in lemma 6.4 (and seen in section 6.5), the last two terms cannot contribute to X_3, since they yield crossed diagrams (lemma 6.3). But in general all four terms can be nonzero; indeed, some terms may be equal and cancel. Some examples are shown in figures 18 and 19. These examples illustrate shorthand calculations alongside the standard notation. Each $\overline{f}_2$ is calculated using section 5.3 and equation 9.

Continuing to $\overline{f}_3$, we know that when $\overline{f}_3(M) \neq 0$ then M is viable (lemma 6.1), $X_3(M) = 0$ and $M_1 M_2 M_3 = 0$ (lemma 6.7). Moreover, M has no singular matched pairs, l twisted matched pairs, and m critical matched pairs, where $m \leq 1$ and $l + m \geq 2$ (theorem 6.5). It follows that $l \geq 1$, so (h, s, t) is twisted, so by theorem 5.2(iv) then $f_3(M) = \overline{A}^*_{\mathcal{CR}^{\preceq}} \circ \overline{U}_3(M)$.

$$\overline{f}_3(M) = \overline{A}^*_{\mathcal{CR}^{\preceq}}\Big(\overline{f}_1(M_1)\overline{f}_2(M_2 \otimes M_3) + \overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3) + \overline{f}_2(M_1 M_2 \otimes M_3) + \overline{f}_2(M_1 \otimes M_2 M_3)\Big). \tag{10}$$

Here $A^*_{\mathcal{CR}^{\preceq}}$ is the creation operator of the creation choice function $\mathcal{CR}^{\preceq}$ (definition 3.17) of the pair ordering $\preceq$ (definition 3.21), and $\overline{A}^*_{\mathcal{CR}^{\preceq}}$ is its image in $\overline{\mathcal{A}}$ (well-defined as discussed in section 3.3).

$$\overline{U}_3 \left(\bullet \quad p'_+ \quad \circ \qquad \circ \quad p'_- \quad \bullet \right) = \overline{f}_2 \left(\bullet \quad p'_+ \quad \circ \qquad \right) + \overline{f}_2 \left(\bullet \quad p'_+ \quad \circ \qquad \right) = 0$$

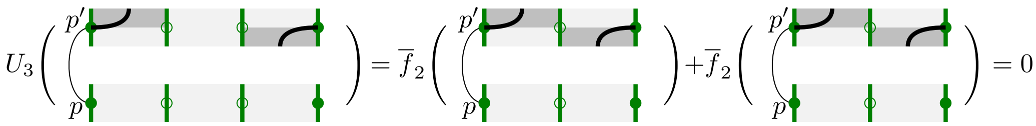

Figure 19: An example of $U_3(M_1 \otimes M_2 \otimes M_3)$. In this case both $f_2(M_1 \otimes M_2)$ and $f_2(M_2 \otimes M_3)$ are zero. The two terms $f_2(M_1 M_2 \otimes M_3)$ and $f_2(M_1 \otimes M_2 M_3)$ cancel out to give zero.

Each of the four terms in equation (10) consists of at most one diagram. Since diagrams in $\overline{U}_3(M)$ may have a crossing, a diagram in $\overline{f}_3(M)$ may have up to two crossings.

Now M has $m \leq 1$ critical matched pairs. If $m = 0$ then all pairs are tight or twisted, and any diagram in $\overline{f}_3$ above has precisely two crossings. If $m = 1$, then the critical pair P must eventually have a tight local diagram to yield a nonzero result, so the diagram at P becomes crossed by an $\overline{f}_2$ and then sublimates; hence any diagram in $\overline{f}_3$ has one crossing.

We find that, in order to obtain a nonzero result for $\overline{f}_3(M)$, the local diagrams at twisted or critical matched pairs must be "distributed" across M_1, M_2, M_3. For twisted pairs we make this precise in the following statement.

Recall (section 4.7) that if M is twisted at a matched pair $P = \{p, p'\}$, then one place p or p' is occupied, and accordingly M is twisted at p or p'. If M is twisted at p, then the two steps p_+, p_- around p are covered by some M_i and M_j respectively, with $i < j$, as in the twisted column of table 2.

Lemma 7.1. *Consider an A_∞ structure defined by a pair ordering $\preceq$.*

Suppose $M = M_1 \otimes M_2 \otimes M_3$ is viable, twisted at precisely two places p, q of matched pairs $P = \{p, p'\}$ and $Q = \{q, q'\}$, with all other pairs tight. Moreover, suppose that p_+, q_+ are both covered by the same M_i, and p_-, q_- are both covered by the same M_j.

Then $X_3(M)$, $\overline{U}_3(M)$ and $\overline{f}_3(M)$ are all zero.

We can denote this result for $\overline{f}_3$ by

$$\overline{f}_3 \begin{pmatrix} \bullet & q_+ & \circ & q_- & \bullet & \bullet \\ \bullet & p_+ & \circ & p_- & \bullet & \bullet \end{pmatrix} = \overline{f}_3 \begin{pmatrix} \bullet & q_+ & \circ & \circ & q_- & \bullet \\ \bullet & p_+ & \circ & \circ & p_- & \bullet \end{pmatrix} = \overline{f}_3 \begin{pmatrix} \bullet & \bullet & q_+ & \circ & q_- & \bullet \\ \bullet & \bullet & p_+ & \circ & p_- & \bullet \end{pmatrix} = 0.$$

Proof. There are three possibilities for i and j: $(i,j) = (1,2), (1,3)$ or $(2,3)$. In all cases $X_3(M) = 0$ as there are no critical matched pairs (theorem 6.6). Suppose without loss of generality that $P \preceq Q$, so creation operators introduce crossings at P in preference to Q.

First suppose $(i,j) = (1,2)$. Then $M_1 M_2 = 0$ (being twisted) and $M_2 M_3 \neq 0$ (being tight), so $\overline{f}_2(M_2 \otimes M_3) = 0$ (lemma 6.7). Thus

$$\overline{U}_3(M) = \overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3) + \overline{f}_2(M_1 \otimes M_2 M_3)$$

Now $\overline{f}_2(M_1 \otimes M_2) = \overline{A}_{\mathcal{CR}^\preceq}\left(\overline{f}_1(M_1)\overline{f}_1(M_2)\right)$ is the diagram obtained from $\overline{f}_1(M_1)\overline{f}_1(M_2)$ by inserting a crossing at P. Similarly $\overline{f}_2(M_1 \otimes M_2 M_3)$ is obtained from $\overline{f}_1(M_1)\overline{f}_1(M_2 M_3)$ by inserting a crossing at P. Since the diagrams $\overline{f}_2(M_1 \otimes M_2 M_3)$ and $\overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3)$ have the same H-data, are crossed at P, twisted at Q, elsewhere tight, and have the same strands at all-on doubly occupied pairs (chosen by the same cycle selection function of $\preceq$), they are equal. Thus $\overline{U}_3(M) = 0$ and $\overline{f}_3(M) = \overline{A}_{\mathcal{CR}^\preceq} \circ \overline{U}_3(M) = 0$.

The case $(i,j) = (2,3)$ is similar.

Finally suppose $(i,j) = (1,3)$. Then $M_1 M_2$ and $M_2 M_3$ are nonzero, so $\overline{f}_2(M_1 \otimes M_2)$ and $\overline{f}_2(M_2 \otimes M_3)$ are zero (lemma 6.7). The remaining two terms of $\overline{U}_3(M)$ are $\overline{f}_2(M_1 M_2 \otimes M_3)$ and $\overline{f}_2(M_1 \otimes M_2 M_3)$, both of which are crossed at p, twisted at q, and equal elsewhere, so again $\overline{U}_3$ and $\overline{f}_3$ are zero. $\square$

The following lemma, together with lemma 7.1 and the general result of theorem 6.5, completely calculates $\overline{f}_3(M)$ when M has two non-tight matched pairs.

Lemma 7.2. *Consider an A_∞ structure defined by a pair ordering $\preceq$.*

Suppose $M = M_1 \otimes M_2 \otimes M_3$ is viable, has two matched pairs $P = \{p, p'\} \prec Q = \{q, q'\}$ which are twisted or critical, and all other matched pairs tight, in one of the arrangements depicted below.

Then $\overline{f}_3(M)$ is zero or nonzero as shown. If nonzero, it is given by a single diagram in $\overline{\mathcal{A}}$, with the H-data of M, which is crossed at each twisted matched pair of M, and elsewhere tight.

Nonzero:

$$\left(\begin{array}{c|c|c} q'_- & q_+ & q_- \\ p_+ & p_- & \end{array}\right) \left(\begin{array}{c|c|c} & q_+ & q_- \\ p_+ & p_- & p'_+ \end{array}\right) \left(\begin{array}{c|c|c} q_+ & & q_- \\ p'_- & p_+ & p_- \end{array}\right) \left(\begin{array}{c|c|c} q_+ & q_- & q'_+ \\ p_+ & & p_- \end{array}\right) \left(\begin{array}{c|c|c} q_+ & q_- & \\ & p_+ & p_- \end{array}\right)$$
$$\left(\begin{array}{c|c|c} q_+ & q_- & \\ p'_- & p_+ & p_- \end{array}\right) \left(\begin{array}{c|c|c} q_+ & q_- & q'_+ \\ & p_+ & p_- \end{array}\right) \left(\begin{array}{c|c|c} q'_- & q_+ & q_- \\ p_+ & & p_- \end{array}\right) \left(\begin{array}{c|c|c} q_+ & & q_- \\ p_+ & p_- & p'_+ \end{array}\right) \left(\begin{array}{c|c|c} & q_+ & q_- \\ p_+ & p_- & \end{array}\right)$$
$$\left(\begin{array}{c|c|c} q_+ & q_- & \\ p_+ & p_- & p'_+ \end{array}\right) \left(\begin{array}{c|c|c} q_+ & q_- & \\ p_+ & & p_- \end{array}\right) \left(\begin{array}{c|c|c} & q_+ & q_- \\ p_+ & & p_- \end{array}\right) \left(\begin{array}{c|c|c} & q_+ & q_- \\ p'_- & p_+ & p_- \end{array}\right)$$

Zero:

$$\left(\begin{array}{c|c|c} q_+ & q_- & q'_+ \\ p_+ & p_- & \end{array}\right) \left(\begin{array}{c|c|c} q_+ & & q_- \\ p_+ & p_- & \end{array}\right) \left(\begin{array}{c|c|c} q_+ & & q_- \\ & p_+ & p_- \end{array}\right) \left(\begin{array}{c|c|c} q'_- & q_+ & q_- \\ & p_+ & p_- \end{array}\right)$$

(Circles denoting idempotents are omitted; they can be inferred since each nontrivial local diagram covers at most one step.)

The conclusion that, if $\overline{f}_3(M)$ is a single diagram, then it is as claimed, follows purely from grading considerations: $\overline{f}_3$ has Maslov grading 2, but the Maslov index can only be increased at non-tight pairs. There are only two non-tight matched pairs, so the Maslov index must be increased by 1 at each. A twisted pair must become crossed, and a critical pair must become tight.

Proof. In the cases depicted in the first four diagrams in the first two rows above, we have a critical and a twisted pair, and $M_1 M_2 = M_2 M_3 = 0$. In each of these cases one of $M_1 \otimes M_2$ or $M_2 \otimes M_3$ is singular, and the other is twisted. Then precisely one of $\overline{f}_2(M_1 \otimes M_2)$ or $\overline{f}_2(M_2 \otimes M_3)$ is nonzero, and $\overline{f}_2$ introduces a crossing at the twisted matched pair. Then the multiplication $\overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3)$ or $\overline{f}_1(M_1)\overline{f}_2(M_2 \otimes M_3)$ is tight at one pair and twisted at the other; and in fact this diagram is $\overline{U}_3(M)$. Applying a creation operator, we obtain $\overline{f}_3(M)$ as a single diagram with a single crossing.

In the cases depicted at the end of the first and second rows, again $M_1 M_2 = M_2 M_3 = 0$, and both $\overline{f}_2(M_1 \otimes M_2)$ and $\overline{f}_2(M_2 \otimes M_3)$ are nonzero, each with a single crossed pair. So $\overline{f}_2(M_1 \otimes M_2)\overline{f}_1(M_3)$ and $\overline{f}_1(M_1)\overline{f}_2(M_2 \otimes M_3)$ are both nonzero, one crossed at p and twisted at q, the other crossed at q and twisted at p. The creation operator $\overline{A}_{\mathcal{CR}\preceq}$ sends the former to zero, and introduces a crossing at p into the latter. Thus $\overline{f}_3(M)$ is given by a single diagram, crossed at both p and q, as desired.

The other cases can be calculated by similar reasoning. □

7.2 Level 4

We now compute two examples at level 4, illustrating some interesting phenomena. As usual, let $M = M_1 \otimes \cdots \otimes M_n$ denotes a tensor product of nonzero homology classes of diagrams, with H-data (h, s, t).

Our first example shows that the necessary conditions for X_n to be nonzero in theorem 6.6 are not sufficient. It is an M with precisely 2 critical matched pairs, and all other matched pairs tight — and in fact one can find a tight diagram with the same H-data — but with $X_4(M) = 0$.

Letting $P = \{p, p'\}$, $Q = \{q, q'\}$ be matched pairs with $P \prec Q$ as usual, we can compute

$$X_4 \begin{pmatrix} \bullet & q_+ & \circ & q_- & \bullet & q'_+ & \circ & \circ \\ \bullet & p_+ & \circ & p_- & \bullet & p'_+ & \circ & \circ \end{pmatrix} = 0,$$

since in this case $\overline{f}_3(M_1 \otimes M_2 \otimes M_3) = 0$ (theorem 6.5; there are two critical pairs), $\overline{f}_3(M_2 \otimes M_3 \otimes M_4) = 0$ (since $M_2 \otimes M_3 \otimes M_4$ is singular), and $\overline{f}_2(M_3 \otimes M_4) = 0$ (lemma 6.7; as $M_3 M_4 \neq 0$).

One can also compute that the following are zero:

$$X_4 \begin{pmatrix} \circ & q'_- & \bullet & q_+ & \circ & q_- & \bullet & q'_+ & \circ \\ \bullet & & \bullet & p_+ & \circ & p_- & \bullet & p'_+ & \circ \end{pmatrix}, \quad X_4 \begin{pmatrix} \bullet & q_+ & \circ & q_- & \bullet & q'_+ & \circ & & \circ \\ \bullet & p_+ & \circ & p_- & \bullet & & \bullet & p'_+ & \circ \end{pmatrix},$$

$$X_4 \begin{pmatrix} \bullet & q_+ & \circ & & \circ & q_- & \bullet & q'_+ & \circ \\ \bullet & p_+ & \circ & p_- & \bullet & & \bullet & p'_+ & \circ \end{pmatrix}, \quad X_4 \begin{pmatrix} \bullet & q_+ & \circ & q_- & \bullet & q'_+ & \circ & & \circ \\ \bullet & p_+ & \circ & p_- & \bullet & p'_+ & \circ & p'_- & \bullet \end{pmatrix}.$$

Our second example shows that $\overline{f}_n$ is not diagrammatically simple (as it might appear from small cases). We have four matched pairs $P \prec Q \prec R \prec S$, with $P = \{p, p'\}$, $Q = \{q, q'\}$, $R = \{r, r'\}$, $S = \{s, s'\}$. We will show that

$$\overline{f}_4 \begin{pmatrix} \bullet & s_+ & \circ & s_- & \bullet & & \bullet & & \bullet \\ \bullet & r_+ & \circ & & \circ & r_- & \bullet & & \bullet \\ \bullet & & \bullet & & \bullet & q_+ & \circ & q_- & \bullet \\ \bullet & & \bullet & p_+ & \circ & & \circ & p_- & \bullet \end{pmatrix} = \begin{pmatrix} c_s \\ w_r \\ c_q \\ c_p \end{pmatrix} + \begin{pmatrix} c_s \\ c_r \\ w_q \\ c_p \end{pmatrix}$$

Observe that, as there are no critical pairs, any X_k term with $k > 2$ is zero (theorem 6.6). Moreover, $M_1 M_2 = M_3 M_4 = 0$. Thus $\overline{f}_4(M) = \overline{A}_P^* \circ \overline{U}_4(M)$, and

$$\begin{aligned} \overline{U}_4(M) = {} & \overline{f}_1(M_1)\overline{f}_3(M_2 \otimes M_3 \otimes M_4) + \overline{f}_2(M_1 \otimes M_2)\overline{f}_2(M_3 \otimes M_4) \\ & + \overline{f}_3(M_1 \otimes M_2 \otimes M_3)\overline{f}_1(M_4) + \overline{f}_3(M_1 \otimes M_2 M_3 \otimes M_4). \end{aligned}$$

Now $M_2 \otimes M_3 \otimes M_4$ is twisted at P and Q, and tight at R and S; $\overline{f}_3(M_2 \otimes M_3 \otimes M_4)$ is then given by lemma 7.2 and is a nonzero diagram. The same applies to $\overline{f}_3(M_1 \otimes M_2 \otimes M_3)$. As $M_1 \otimes M_2$ and $M_3 \otimes M_4$ are twisted at a single matched pair, and tight elsewhere, $\overline{f}_2(M_1 \otimes M_2)$ and $\overline{f}_2(M_3 \otimes M_4)$ are also both nonzero diagrams (section 5.3), each with a single crossing.

We now have all terms in $\overline{U}_4(M)$ except $\overline{f}_3(M_1 \otimes M_2 M_3 \otimes M_4)$. To this end we note that $M_1 M_2 M_3 = M_2 M_3 M_4 = 0$, so $\overline{U}_3(M_1 \otimes M_2 M_3 \otimes M_4) = \overline{f}_1(M_1)\overline{f}_2(M_2 M_3 \otimes M_4) + \overline{f}_2(M_1 \otimes M_2 M_3)\overline{f}_1(M_4)$. Since $M_2 M_3 \otimes M_4$ is twisted at P and Q, the creation operator inserts a crossing at P; and since $M_1 \otimes M_2 M_3$ is twisted at R and S, the creation operator inserts a crossing at R. Hence

$$\begin{aligned} \overline{f}_3(M_1 \otimes M_2 M_3 \otimes M_4) &= \overline{A}_P^* \left(\overline{f}_1(M_1)\overline{f}_2(M_2 M_3 \otimes M_4) + \overline{f}_2(M_1 \otimes M_2 M_3)\overline{f}_1(M_4) \right) \\ &= \overline{A}_P^* \left[\begin{pmatrix} \bullet & s_+ & \circ \\ \bullet & r_+ & \circ \\ \bullet & & \bullet \\ \bullet & & \bullet \end{pmatrix} \begin{pmatrix} \circ & s_- & \bullet \\ \circ & r_- & \bullet \\ \bullet & w_q & \bullet \\ \bullet & c_p & \bullet \end{pmatrix} + \begin{pmatrix} \bullet & w_s & \bullet \\ \bullet & c_r & \bullet \\ \bullet & q_+ & \circ \\ \bullet & p_+ & \circ \end{pmatrix} \begin{pmatrix} \bullet & & \bullet \\ \bullet & & \bullet \\ \circ & q_- & \bullet \\ \circ & p_- & \bullet \end{pmatrix} \right] \\ &= \overline{A}_P^* \left[\begin{pmatrix} \bullet & w_s & \bullet \\ \bullet & w_r & \bullet \\ \bullet & w_q & \bullet \\ \bullet & c_p & \bullet \end{pmatrix} + \begin{pmatrix} \bullet & w_s & \bullet \\ \bullet & c_r & \bullet \\ \bullet & w_q & \bullet \\ \bullet & w_p & \bullet \end{pmatrix} \right] = \begin{pmatrix} \bullet & w_s & \bullet \\ \bullet & c_r & \bullet \\ \bullet & w_q & \bullet \\ \bullet & c_p & \bullet \end{pmatrix}, \\ \overline{U}_4(M) &= \begin{pmatrix} w_s \\ w_r \\ c_q \\ c_p \end{pmatrix} + \begin{pmatrix} c_s \\ w_r \\ c_q \\ w_p \end{pmatrix} + \begin{pmatrix} c_s \\ c_r \\ w_q \\ w_p \end{pmatrix} + \begin{pmatrix} w_s \\ c_r \\ w_q \\ c_p \end{pmatrix} \end{aligned}$$

so that, applying $\overline{A}_P^*$, $\overline{f}_4(M)$ has the claimed form.

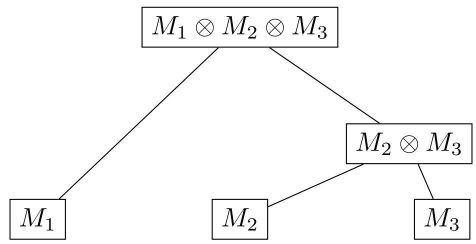

Figure 20: An operation tree.

8 Nontrivial higher operations

In this section we only consider A_∞ structures arising from a pair ordering $\preceq$.

Although we have various necessary conditions for X_n or $\overline{f}_n$ to be nonzero (viability, theorems 6.5 and 6.6, lemma 6.7, lemma 7.1), we do not yet have conditions which are sufficient to ensure X_n or $\overline{f}_n$ are nonzero — whether the operations are defined via a pair ordering, or by Kadeishvili's construction more generally.

We have some results at low levels. For instance, $X_2(M_1 \otimes M_2)$ is nonzero if and only of $M_1 \otimes M_2$ is tight, essentially by definition. Proposition 6.9 shows that the necessary conditions of theorem 6.6 for X_3 to be nonzero are also sufficient. However, the X_4 example of section 7.2 shows that these conditions are not sufficient for X_4 to be zero.

Indeed, the $\overline{f}_3$ examples of section 7.1 (particularly lemma 7.2) show that even the question of whether $\overline{f}_3$ is zero or nonzero can be rather subtle. The $\overline{f}_4$ example of section 7.2 there shows that matters do not get simpler at higher levels.

In this section we prove some sufficient conditions for $\overline{f}_n$ and X_n to be nonzero. They are, however, far from being necessary conditions.

As usual, throughout this section $M = M_1 \otimes \cdots \otimes M_n$ always denotes a tensor product of nonzero homology classes of diagrams

8.1 Operation trees

Lemma 7.1 and some of the level 3 and 4 examples show that, even though a tensor product $M_1 \otimes \cdots \otimes M_n$ might have the right number of critical and twisted matched pairs, the steps of these pairs must be covered by the M_i in a way that is appropriately "horizontally distributed".

To this end, we study rooted trees describing the order in which operations are performed.

Definition 8.1 (Operation tree)**.** *An* operation tree *for $\mathcal{H}^{\otimes n}$ is a rooted plane binary tree with n leaves, ordered from left to right, and with each vertex v labelled by a viable tensor product of homology classes of diagrams M_v, so that the following conditions are satisfied.*

(i) Each leaf is labelled with a nonzero homology class of diagram in $\mathcal{H}$.

(ii) Each vertex is labelled with the tensor product of the labels on its ordered children.

If the root vertex is labelled M, we say $\mathcal{T}$ is an operation tree *for M.*

Thus, if the leaves are labelled $M_1, \ldots, M_n$ in order, then the root vertex is labelled $M_1 \otimes \cdots \otimes M_n$. See figure 20 for an example.

It will also be useful to consider a certain type of subtree, as in the following definition.

Definition 8.2 (Subtree below v)**.** *Let $\mathcal{T}$ be an operation tree, and v a vertex of $\mathcal{T}$. The* operation subtree of $\mathcal{T}$ below v *is the subtree $\mathcal{T}_v$ of $\mathcal{T}$ with root vertex v, consisting of all edges and vertices below v, and all vertex labels inherited from $\mathcal{T}$.*

Clearly $\mathcal{T}_v$ is also an operation tree.

8.2 Validity and distributivity

If $\mathcal{T}$ is an operation tree for M, each vertex of $\mathcal{T}$ is labelled by a sub-tensor-product M_v of M. The various labels M_v may have different types of tightness, depending on how the various steps around each matched pair are covered.

Singular tensor products should be avoided, and so we make the following definition.

Definition 8.3. *Let $\mathcal{T}$ be an operation tree for $\mathcal{H}^{\otimes n}$.*

(i) A vertex of $\mathcal{T}$ is valid *if its label is non-singular.*

(ii) The operation tree $\mathcal{T}$ is valid *if it is valid at all of its vertices.*

Thus, in a valid operation tree for M, each vertex label is tight, twisted or critical. (Note M may have singular sub-tensor-products, but they do not appear as vertex labels.) Equivalently, each label is tight, twisted or critical at all matched pairs (lemma 4.26).

Lemma 8.4. *Let $\mathcal{T}$ be a valid operation tree for M, and v a vertex of $\mathcal{T}$. Then the operation subtree $\mathcal{T}_v$ of $\mathcal{T}$ below v is valid.*

Proof. Each label is non-singular in $\mathcal{T}$, hence also non-singular in $\mathcal{T}_v$. □

Nonzero A_∞ operations require carefully regulated numbers of twisted and critical matched pairs, as required by theorems 6.5 and 6.6. Hence we make the following definition.

Definition 8.5. *Let $\mathcal{T}$ be a valid operation tree.*

(i) A vertex of $\mathcal{T}$ with k leaves, labelled M, is distributive *if there are at least $k-2$ matched pairs at which M is twisted or critical.*

(ii) The tree $\mathcal{T}$ is distributive *if every vertex of $\mathcal{T}$ is distributive.*

8.3 Joining and grafting trees

We now consider some methods to combine operation trees into larger trees.

The first operation, *joining*, places two existing operation trees below a new root vertex.

Definition 8.6. *Let $\mathcal{T}', \mathcal{T}''$ be operation trees for M', M'', where $M' \otimes M''$ is viable. Let v', v'' be the root vertices of $\mathcal{T}', \mathcal{T}''$ respectively. The* join *of $\mathcal{T}'$ and $\mathcal{T}''$ is the tree $\mathcal{T}$ obtained by placing $\mathcal{T}'$ and $\mathcal{T}''$ below v_0, so that v', v'' are the left and right children of $\mathcal{T}$. The root vertex v_0 is labelled $M' \otimes M''$, and each other vertex inherits its label from $\mathcal{T}'$ or $\mathcal{T}''$.*

Clearly, the join of two operation trees is again an operation tree; note that this requires the assumption that $M' \otimes M''$ be viable. Figure 21 shows an example.

Under certain circumstances, joining trees preserves validity and distributivity.

Lemma 8.7. *Let $\mathcal{T}', \mathcal{T}''$ be operation trees for $M' = M_1 \otimes \cdots \otimes M_j$ and $M'' = M_{j+1} \otimes \cdots \otimes M_n$, and let $\mathcal{T}$ be their join. Suppose $X_n(M' \otimes M'') \neq 0$ or $\overline{f}_n(M' \otimes M'') \neq 0$.*

If $\mathcal{T}'$ and $\mathcal{T}''$ are valid and distributive, then $\mathcal{T}$ is also valid and distributive.

Note that if $X_n(M' \otimes M'')$ or $\overline{f}_n(M' \otimes M'')$ is nonzero, then $M' \otimes M''$ is certainly viable, so that $\mathcal{T}$ is a well defined operation tree.

Proof. Each non-root vertex of $\mathcal{T}$ retains its label from $\mathcal{T}'$ or $\mathcal{T}''$. So if $\mathcal{T}', \mathcal{T}''$ are valid (resp. distributive), then $\mathcal{T}$ is valid (resp. distributive) at these vertices. So we only need consider the root vertex v_0 of $\mathcal{T}$, which is labelled with $M = M' \otimes M''$.

If $X_n(M) \neq 0$, then by theorem 6.6, M has precisely $n-2$ matched pairs which are critical, and all other matched pairs are tight. If $\overline{f}_n(M) \neq 0$, then by theorem 6.5, M has at least $n-1$ matched pairs which are twisted or critical, and all other matched pairs are tight. Either way M is not singular, and the number of critical or twisted matched pairs is $\geq n-2$. Thus v_0 is valid and distributive, and hence so also is $\mathcal{T}$. □

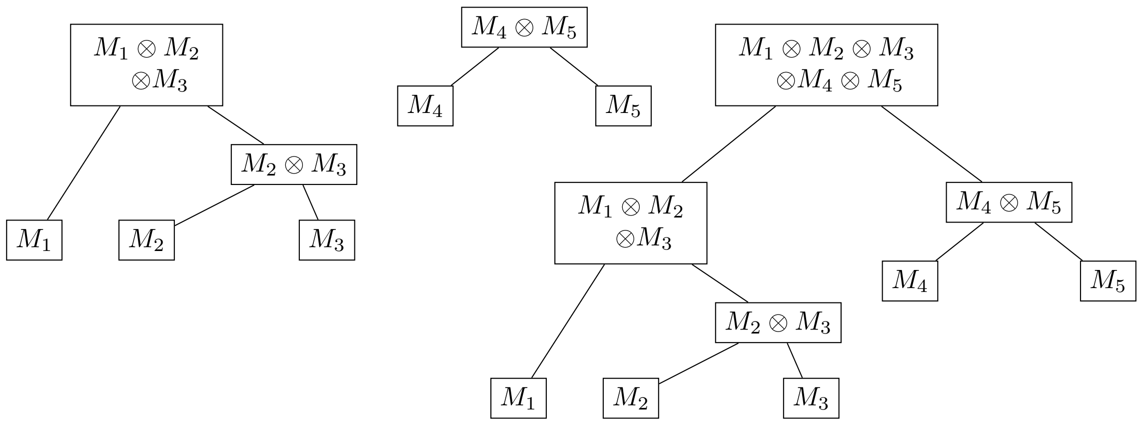

Figure 21: Operation trees $\mathcal{T}', \mathcal{T}'', \mathcal{T}$ (left to right), where $\mathcal{T}$ is the join of $\mathcal{T}'$ and $\mathcal{T}''$.

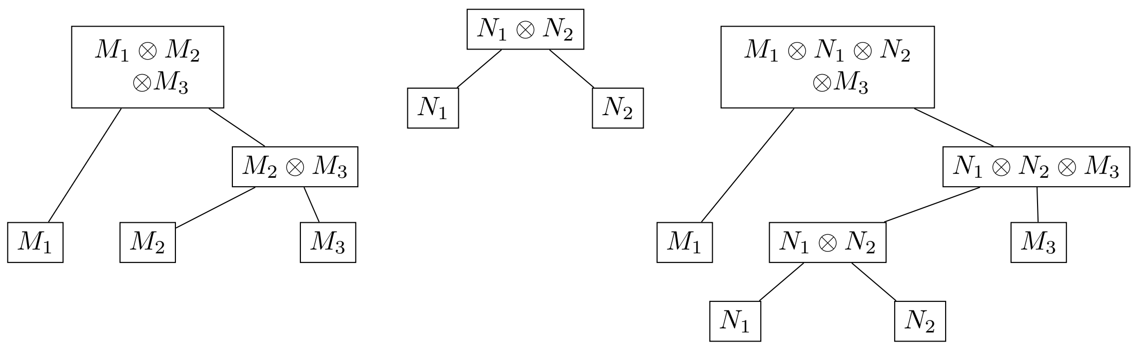

Figure 22: Operation trees $\mathcal{T}', \mathcal{T}'', \mathcal{T}$, where $\mathcal{T}$ is the grafting of $\mathcal{T}''$ onto $\mathcal{T}'$ at position 2.

The second operation, *grafting*, implants a tree at a leaf of an existing tree.

Definition 8.8. *Let $\mathcal{T}', \mathcal{T}''$ be operation trees for $M' = M_1 \otimes \cdots \otimes M_i$ and $N' = N_1 \otimes \cdots \otimes N_j$, and suppose N' and M_k have the same H-data.*

The grafting of $\mathcal{T}''$ onto $\mathcal{T}'$ at position k *is the tree $\mathcal{T}$ obtained by identifying the k'th leaf of $\mathcal{T}'$ with the root vertex of $\mathcal{T}''$. The vertices of $\mathcal{T}'$ are relabelled by replacing every instance of M_k with the tensor product $N_1 \otimes \cdots \otimes N_j$; other labels are inherited from $\mathcal{T}''$.*

Figure 22 shows an example. Thus $\mathcal{T}$ is an operation tree for the tensor product

$$M = M_1 \otimes \cdots \otimes M_{k-1} \otimes N_1 \otimes \cdots \otimes N_j \otimes M_{k+1} \otimes \cdots M_n.$$

The assumption that N' and M_k share the same H-data ensures M is viable.

As with joining, under certain circumstances grafting preserves validity and distributivity.

Lemma 8.9. *Let $\mathcal{T}', \mathcal{T}''$ be operation trees for $M' = M_1 = M_1 \otimes M_i$ and $N' = N_1 \otimes \cdots \otimes N_j$ respectively. Suppose $X_j(N') = M_k$, and let $\mathcal{T}$ be the grafting of $\mathcal{T}''$ onto $\mathcal{T}'$ at position k.*

If $\mathcal{T}'$ and $\mathcal{T}''$ are valid and distributive, then $\mathcal{T}$ is also valid and distributive.

Note $X_j(N') = M_k$ implies N' and M_k have equal H-data, so $\mathcal{T}$ is a well defined operation tree.

Proof. Each vertex of $\mathcal{T}''$ retains its label, hence validity and distributivity are satisfied. At vertices of $\mathcal{T}'$ which retain their label, the same applies. Thus we only need consider vertices of $\mathcal{T}'$ whose labels are changed in $\mathcal{T}$, i.e. those whose label involves M_k.

Let v be a vertex of $\mathcal{T}'$ with l leaves, labelled $M'_v = M_u \otimes \cdots \otimes M_k \otimes \cdots \otimes M_{u+l-1}$ in $\mathcal{T}'$; the label in $\mathcal{T}$ is thus $M_v = M_u \otimes \cdots \otimes (N_1 \otimes \cdots \otimes N_j) \otimes \cdots \otimes M_{u+l-1}$. Since $\mathcal{T}'$ is valid, M'_v is non-singular. Since M_v and M'_v have the same H-data, L is non-singular; so v is valid.

It remains to show that v is distributive. Since $\mathcal{T}'$ is distributive, M'_v has at least $l-2$ matched pairs which are twisted or critical. Now $M_k = X_j(N')$ implies that M_k is the unique nonzero homology class of diagram with the H-data of N', so M'_v is obtained from M_v by an H-contraction (definition 4.31). By lemma 4.32, if M'_v is critical at a matched pair P, then M_v is critical at P; and if M'_v is twisted at P, then M_v is twisted at P. Hence M_v has at least as many twisted and critical matched pairs as M'_v. □

8.4 Nonzero operations require trees

As we now show, a valid distributive operation tree for M is a necessary condition for $X_n(M)$ or $\overline{f}_n(M)$ to be nontrivial.

Proposition 8.10. *Consider an A_∞ structure on $\mathcal{H}$ arising from a pair ordering. If $X_n(M) \neq 0$ or $\overline{f}_n(M) \neq 0$, then there is a valid distributive operation tree for M.*

Proposition 8.10 is a precise version of proposition 1.3.

Proof. First note that as $X_n(M)$ or $\overline{f}_n(M) \neq 0$, M is viable (lemma 6.1).

When $n = 1$, the valid and distributive conditions are vacuous.

Now suppose the statement holds for all X_k and $\overline{f}_k$ for $k < n$, and consider X_n and $\overline{f}_n$.

Suppose $X_n(M) \neq 0$. By lemma 6.4, $X_n(M)$ is represented by the sum of crossingless diagrams in $\overline{f}_j(M_1 \otimes \cdots \otimes M_j)\overline{f}_{n-j}(M_{j+1} \otimes \cdots \otimes M_n)$, so some $\overline{f}_j(M_1 \otimes \cdots \otimes M_j)$ and $\overline{f}_{n-j}(M_{j+1} \otimes \cdots \otimes M_n)$ are nonzero. By induction there are valid distributive operation trees $\mathcal{T}'$ for $M_1 \otimes \cdots \otimes M_i$ and $\mathcal{T}''$ for $M_{i+1} \otimes \cdots \otimes M_n$. Now let $\mathcal{T}$ be the join of $\mathcal{T}'$ and $\mathcal{T}''$; this is well defined as M is viable. Since $\mathcal{T}'$ and $\mathcal{T}''$ are valid and distributive, by lemma 8.7 so is $\mathcal{T}$.

Now suppose $\overline{f}_n(M) \neq 0$. Then $X_n(M) = 0$ (lemma 6.7), and M has all matched pairs tight, twisted or critical, with at least one matched pair twisted (theorem 6.5). Thus $\overline{f}_n(M) = A^*_{\mathcal{CR}\preceq}\overline{U}_n(M)$, and hence $\overline{U}_n(M) \neq 0$. From equation (2) then some term of the form $\overline{f}_j(M_1 \otimes \cdots \otimes M_j)\overline{f}_{n-j}(M_{j+1} \otimes \cdots \otimes a_M)$ or $\overline{f}_{n-j+1}(M_1 \otimes \cdots \otimes M_k \otimes X_j(M_{k+1} \otimes \cdots \otimes M_{k+j}) \otimes \cdots \otimes M_n)$ is nonzero. We consider the two cases separately.

In the first case, by induction, there are operation trees $\mathcal{T}'$ for $M_1 \otimes \cdots \otimes M_j$, and $\mathcal{T}''$ for $M_{j+1} \otimes \cdots \otimes M_n$, which are valid and distributive. Let $\mathcal{T}$ be the join of $\mathcal{T}'$ and $\mathcal{T}''$; as M is viable, $\mathcal{T}$ is well defined. By lemma 8.7 again, $\mathcal{T}$ is valid and distributive.

In the second case induction gives operation trees $\mathcal{T}'$ for $M_1 \otimes \cdots \otimes M_k \otimes X_j(M_{k+1} \otimes \cdots \otimes M_{k+j}) \otimes \cdots \otimes M_n$, and $\mathcal{T}''$ for $M_{k+1} \otimes \cdots \otimes M_{k+j}$, which are valid and distributive. Let $\mathcal{T}$ be the grafting of $\mathcal{T}''$ onto $\mathcal{T}'$ at position $k+1$. This is clearly a well-defined operation tree, and by lemma 8.9, $\mathcal{T}$ is valid and distributive. □

8.5 Local trees

Let $\mathcal{T}$ be an operation tree for M. We now consider M at a single matched pair P, and use this to construct "localised" versions of $\mathcal{T}$. We will define a *local operation tree*, which has the same underlying tree, and a *reduced local operation tree*, whose underlying tree is obtained by contracting "extraneous" vertices.

Recall the local tensor product of $M = M_1 \otimes \cdots \otimes M_n$ at P is given by $M_P = (M_1)_P \otimes \cdots \otimes (M_n)_P$ (definition 4.6).

Definition 8.11. *The* local operation tree $\widetilde{\mathcal{T}_P}$ *of $\mathcal{T}$ at P is obtained from $\mathcal{T}$ by replacing each M_i with $(M_i)_P$ in each vertex label.*

Since each M_i is nonzero, so is each $(M_i)_P$; and since M is viable, so is M_P. Each vertex label remains viable and the tensor product of its children; so $\widetilde{\mathcal{T}_P}$ is indeed an operation tree for M_P.

By proposition 4.23, M_P is an extension-contraction of one of the tensor products shown in table 2. So there are at most 4 tensor factors of M which have non-horizontal strands at P, i.e. which cover one or more of the 4 steps around P.

Definition 8.12. *A tensor factor M_i of M which has a non-horizontal strand at a matched pair P is called P-*active*. The corresponding leaves of $\mathcal{T}$ are called P-*active *leaves.*

For each P, $\mathcal{T}$ has at most 4 P-active leaves. These are precisely the leaves of $\mathcal{T}_P$ labelled by non-idempotent diagrams.

Now we reduce $\widetilde{\mathcal{T}_P}$ to remove non-active leaves and factors. Consider a non-P-active factor M_v of M, and the corresponding leaf v in $\widetilde{\mathcal{T}_P}$. Then $(M_v)_P$ is idempotent, so deleting it as a factor from M_P leaves a tensor product which is still viable. (Indeed, such a deletion is a trivial contraction: definition 4.13). We delete $(M_v)_P$ from all labels on $\widetilde{\mathcal{T}_P}$, and we delete the leaf v and its incident edge. This leaves a degree-2 vertex, which we smooth (i.e. we delete the degree-2 vertex and combine the two adjacent edges into a single edge). We then have a binary planar tree. (If the root vertex is smoothed, precisely one of its children remains; that child becomes the root.) It is an operation tree for $(M_1)_P \otimes \cdots \otimes \widehat{(M_v)_P} \otimes \cdots \otimes (M_n)_P$, where the hat denotes a deleted factor.

Repeating the process for all non-active factors, we obtain an operation tree $\mathcal{T}_P$ for $(M_{i_1})_P \otimes \cdots \otimes (M_{i_k})_P$, where the M_{i_j} are the P-active factors of M. Note $0 \leq k \leq 4$; if $k = 0$, $\mathcal{T}_P$ is the empty tree.

Definition 8.13. *The operation tree $\mathcal{T}_P$ is called the* reduced local operation tree *of $\mathcal{T}$ at P.*

The operation tree $\mathcal{T}_P$ does not depend on the order in which the non-active factors are deleted; in fact it can also be constructed "at once", as follows. The P-active leaves of $\widetilde{\mathcal{T}_P}$ have a lowest common ancestor v_0 in $\mathcal{T}$. Take the edges and vertices along shortest paths in $\mathcal{T}$ from each P-active leaf to v_0. The union of these edges and vertices is a planar subtree of $\widetilde{\mathcal{T}_P}$ with root v_0 and leaves labelled by M_{i_j}. Smoothing degree-2 vertices in this subtree and labelling vertices appropriately yields $\mathcal{T}_P$.

Note that the vertices of $\mathcal{T}_P$ can be regarded as a subset of the vertices of $\mathcal{T}$ or $\widetilde{\mathcal{T}_P}$: namely, those vertices which are not deleted or smoothed as we remove non-P-active factors.

Local operation trees are useful because of the following fact, a "local-to-global" law for validity.

Lemma 8.14. *Let $\mathcal{T}$ be an operation tree. The following are equivalent:*

(i) $\mathcal{T}$ is valid.

(ii) For all matched pairs P, the local operation tree $\widetilde{\mathcal{T}_P}$ is valid.

(iii) For all matched pairs P, the reduced local operation tree $\mathcal{T}_P$ is valid.

Proof. By lemma 4.26, a tensor product of homology classes of diagrams is non-singular if and only if it is non-singular at all its matched pairs. Since the labels on the operation trees $\widetilde{\mathcal{T}_P}$ are precisely the labels on $\mathcal{T}$, localised to P, (i) and (ii) are equivalent.

As mentioned above, deleting a non-P-active leaf from $\widetilde{\mathcal{T}_P}$, corresponding to a non-P-active factor M_i, produces a trivial contraction on vertex labels. Thus if all vertex labels were non-singular in $\widetilde{\mathcal{T}_P}$, then they remain non-singular. Conversely, if all the "new" vertex labels are non-singular after deletion, their "old" labels (being obtained by extension from the "new" ones — even at the smoothed vertices) were also non-singular. The deleted vertex was labelled by a single idempotent diagram, which is non-singular. After deleting all non-P-active leaves, $\widetilde{\mathcal{T}_P}$ is valid if and only if $\mathcal{T}_P$ is valid. □

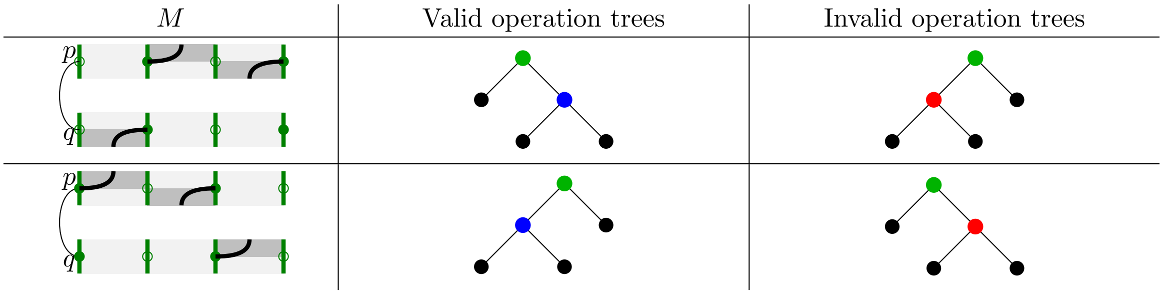

Table 5: Validity of operation trees on sesqui-occupied local critical tensor products. Red, green, blue, black vertices respectively indicate singular, critical, twisted and tight labels.

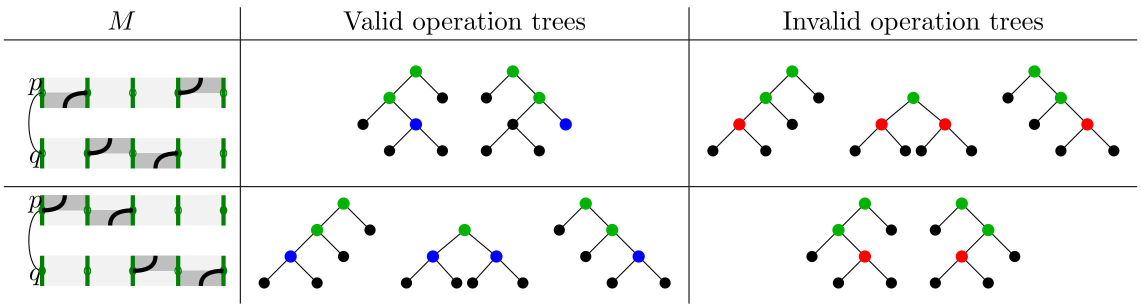

Table 6: Validity of operation trees on doubly occupied local critical tensor products. Red, green, blue, black vertices respectively indicate singular, critical, twisted and tight labels.

8.6 Climbing a tree

Let $\mathcal{T}$ be a reduced local operation tree. Then $\mathcal{T}$ has no more than 4 leaves, so there are not many possible trees. Indeed, the number of rooted planar binary trees with $1, 2, 3, 4, n$ leaves is $1, 1, 2, 5, \frac{1}{n+1}\binom{2n}{n}$.

The tensor products arising in reduced local operation trees are also small in number. If M is the tensor product labelling the root of M, then M is a viable tensor product of homology classes of diagrams on the arc diagram $\mathcal{Z}_P$ consisting of a single matched pair. As $\mathcal{T}$ is a reduced local operation tree, M has no idempotents, i.e. every tensor factor of M has non-horizontal strands. Thus (proposition 4.23) M is one of the tensor products shown in the tight, twisted, critical or singular columns of table 2, or (in the tight case) a contraction thereof.

We ask: for each such tensor product M, which of the possible operation trees on M is valid?

If M is tight or twisted, then any sub-tensor product is tight or twisted (lemma 4.30), and in particular non-singular, so any operation tree for M is valid. And of course if M is singular, then any operation tree for M is invalid, since its root vertex has singular label M.

When M is critical, some but not all operation trees are valid. By examining the possible cases in the critical column of table 2, we observe the following.

- When M is critical and P is sesqui-occupied, precisely 1 of the 2 operation trees are valid.
- When M is critical and P is 00 doubly occupied, precisely 2 of the 5 operation trees are valid.
- When M is critical and P is 11 doubly occupied, precisely 3 of the 5 operation trees are valid.

Tables 5 and 6 illustrate the valid and invalid trees.

Starting from the leaves of $\mathcal{T}$, which are all tight, we can climb $\mathcal{T}$, observing how tightness behaves as the (homology classes of) diagrams labelling the vertices are joined into tensor products.

We observe that whenever there is a singular or twisted vertex label, it occurs when two adjacent diagrams are joined into a singular tensor product. Also, we never see both a a twisted vertex label and a singular vertex label. This leads to the following statement.

Lemma 8.15. *Let $\mathcal{T}$ be an operation tree for a viable tensor product of diagrams M. Then $\mathcal{T}$ is valid if and only if for every non-tight matched pair P of M, $\mathcal{T}_P$ has a twisted vertex label.*

Proof. By lemma 8.14, the validity of $\mathcal{T}$ is equivalent to the the validity of all the $\mathcal{T}_P$. Since M is viable, at each matched pair M is tight, twisted, critical or singular. As discussed above, if M_P is tight at P then $\mathcal{T}_P$ is valid. So it remains to check that when M_P is twisted, critical, or singular, $\mathcal{T}_P$ is valid if and only if $\mathcal{T}_P$ has a twisted vertex label.

If M_P is twisted then again $\mathcal{T}_P$ is valid, and moreover $\mathcal{T}_P$ has a twisted vertex label at the root. If M_P is critical then, from tables 5 and 6, $\mathcal{T}_P$ is valid if and only if there is a twisted vertex label. And if M is singular, then $\mathcal{T}_P$ is invalid, and moreover M_P must be an extension of the singular example in table 2 (lemma 4.27), so $\mathcal{T}_P$ must be the unique rooted binary planar tree with two leaves; the two leaf labels are tight, and the root label is singular, so there is no twisted vertex label. Thus in each case $\mathcal{T}_P$ is valid if and only if it has a twisted vertex label. □

8.7 Strong validity

We saw above that when M is valid, then at every non-tight $P = \{p, p'\}$, the reduced local operation tree $\mathcal{T}_P$ has a twisted vertex label. But in fact, in almost every case, there is *precisely* one twisted vertex label. The only exception is when M_P is 11 doubly occupied and critical (i.e. the second row of table 6), and $\mathcal{T}_P$ is the unique rooted planar binary tree of depth 2 (i.e. the second valid operation tree shown). This particular operation tree can lead to the multiplication of a diagram crossed at p, with a diagram crossed at p', producing a diagram in $\mathcal{F}$. To avoid it, we introduce a "strong" form of validity.

Lemma 8.16. *Let $\mathcal{T}$ be an operation tree for M. The following are equivalent.*

(i) For every non-tight matched pair P of M, there is a unique lowest vertex of $\mathcal{T}$ among those whose label is twisted at P.

(ii) The operation tree $\mathcal{T}$ is valid, and for each non-tight matched pair P of M, there is a unique lowest vertex of $\mathcal{T}$ among those whose label is not tight at P.

(iii) For each non-tight matched pair P of M, there is a unique lowest vertex of $\widetilde{\mathcal{T}_P}$ among those whose label is twisted.

(iv) For every non-tight matched pair P of M, $\mathcal{T}_P$ has a unique twisted vertex label.

Definition 8.17. *The operation tree $\mathcal{T}$ is* strongly valid *if the conditions of lemma 8.16 hold.*

Proof of lemma 8.16. First we show equivalence of (i) and (ii). If $\mathcal{T}$ is not valid, then (i) fails by lemma 8.15, and (ii) obviously fails. So assume $\mathcal{T}$ is valid. We show that a vertex v of $\mathcal{T}$, with label M_v, is lowest among those with labels twisted at P, if and only if it is lowest among those with labels non-tight at P.

If v is lowest among vertices with label twisted at P, then the children of v have labels which are sub-tensor-products of M_v non-twisted at P. Hence by lemma 4.29 and table 4, the labels on these children are tight at P. All descendants of these children have tight labels at P also, again by lemma 4.29 and table 4. So v is lowest among vertices of $\mathcal{T}$ with labels non-tight at P.

Conversely, if v is lowest among those with labels twisted at P, then all descendants of v have tight labels at P. In particular, the children of v both have labels tight at P. Then $(M_v)_P$ is the tensor product of these tight labels: it cannot be critical, by lemma 4.28, and cannot be singular, since $\mathcal{T}$ is valid. So v is a lowest vertex among those with label twisted at P. Thus (i) and (ii) are equivalent.

Condition (iii) is just a reformulation of (i).

To see equivalence of (iii) and (iv), recall how $\mathcal{T}_P$ is obtained from $\widetilde{\mathcal{T}_P}$. If the label M_v on a leaf v of $\widetilde{\mathcal{T}_P}$ is idempotent, then we delete v and its incident edge, delete M_v from all labels, and smooth the resulting degree-2 vertex w. Since M_v has only horizontal strands, deleting M_v from a label yields a trivial contraction (definition 4.13), which does not change the tightness of the label.

Now w has two children v and x in $\widetilde{\mathcal{T}_P}$. Since M_v is tight (being an idempotent), and M_w is an extension of M_x (by the horizontal strands of M_v), so M_w and M_x have the same tightness. In particular, neither v nor w cannot be lowest among those with twisted label. After deleting v and all instances of M_v in labels, the label on w is the same as the label on x. After smoothing w, every remaining vertex has children and descendants with twisted labels if and only if it had them in $\widetilde{\mathcal{T}_P}$. Thus any vertex which was lowest among those with twisted labels was not v or w, so remains as a vertex, and remains lowest among those with twisted labels. So the set of lowest vertices with twisted labels is preserved.

Repeating this process we eventually arrive at $\mathcal{T}_P$. So $\widetilde{\mathcal{T}_P}$ has a unique lowest vertex among those with twisted labels, if and only if the same is true for $\mathcal{T}_P$. From the examination of reduced local operation trees in section 8.6, we observe that a reduced local operation tree has a unique lowest vertex with twisted label if and only if it has a unique vertex with a twisted label. Thus (iii) and (iv) are equivalent. □

The above discussion also immediately implies the following.

Lemma 8.18. *Suppose $\mathcal{T}$ an operation tree for M which is valid but not strongly valid. Then M has a matched pair which is 11 doubly occupied and critical.* □

By lemma 8.16(i), the following map is well defined.

Definition 8.19. *Let $\mathcal{T}$ be a strongly valid operation tree for M. The function*

$$V_{\mathcal{T}} \colon \{\textit{Non-tight matched pairs of } M\} \longrightarrow \{\textit{Non-leaf vertices of } \mathcal{T}\}$$

sends a matched pair P to the lowest vertex of $\mathcal{T}$ whose label is twisted at P.

By the argument in the proof of lemma 8.16 (that (i) and (ii) are equivalent), $V_{\mathcal{T}}(P)$ is also the lowest vertex of $\mathcal{T}$ whose label is not tight at P.

Lemma 8.20. *Let $\mathcal{T}$ be a strongly valid operation tree for M, and let P be a non-tight matched pair of M. Then the vertices of $\mathcal{T}$ whose labels are non-tight at P are precisely $V_{\mathcal{T}}(P)$ and its ancestors.*

Proof. Let the label on $V_{\mathcal{T}}(P)$ be M'. If v is an ancestor of $V_{\mathcal{T}}(P)$, labelled M_v, then M' is a sub-tensor-product of M_v. As M' is not tight at P, by lemma 4.29 M_v is not tight at P.

Conversely, suppose a vertex v_0 of $\mathcal{T}$ has label non-tight at P. Either v_0 is a lowest such vertex, or v_0 has a child v_1 whose label is also not tight at P. If the latter, then v_1 is either a lowest such vertex, or has a child whose label is non-tight at P. In this way, we eventually arrive at a descendant v_* of v_0, which is lowest amongst those whose labels are not tight at P. By the comment after definition 8.19 then $v_* = V_{\mathcal{T}}(P)$, so v_0 is $V_{\mathcal{T}}(P)$ or one of its ancestors. □

Strong validity shares many of the properties of validity. The following results generalise lemmas 8.14 and 8.4.

Lemma 8.21. *Let $\mathcal{T}$ be an operation tree. The following are equivalent:*

(i) $\mathcal{T}$ is strongly valid.

(ii) For all matched pairs P, the local operation tree $\widetilde{\mathcal{T}_P}$ is strongly valid.

(iii) For all matched pairs P, the reduced local operation tree $\mathcal{T}_P$ is strongly valid.

Proof. By definition (i) is equivalent to (iii). As discussed in the proof of lemma 8.16, strong validity is preserved under the deletion operations which transform $\widetilde{\mathcal{T}_P}$ into $\mathcal{T}_P$, hence (ii) and (iii) are equivalent. (This also follows from the equivalence of characterisations (iii) and (iv) in lemma 8.16.) □

Lemma 8.22. *Let $\mathcal{T}$ be a strongly valid operation tree for M, and let v be a vertex of $\mathcal{T}$ labelled by M_v. Let $\mathcal{T}_v$ be the operation subtree of $\mathcal{T}$ below v. Then the following hold.*

(i) $\mathcal{T}_v$ is a strongly valid operation tree for M_v.

(ii) The function $V_{\mathcal{T}_v}$ is a restriction of the function $V_{\mathcal{T}}$.

Note that M_v is a sub-tensor-product of M, so by lemma 4.30, a matched pair which is non-tight in M_v is also non-tight in M. Hence the domain of $V_{\mathcal{T}_v}$ is a subset of the domain of $V_{\mathcal{T}}$, so the assertion of (ii) makes sense.

Proof. Let P be a matched pair, and consider the local operation trees $\widetilde{\mathcal{T}_P}$ for M_P, and $\widetilde{(\mathcal{T}_v)_P}$ for $(M_v)_P$. To prove (i), we show that if $(M_v)_P$ is not tight, then $\widetilde{(\mathcal{T}_v)_P}$ has a unique lowest vertex with twisted label (lemma 8.16(iii)); and to prove (ii), we show that this vertex is also the unique lowest vertex with twisted label in $\widetilde{\mathcal{T}_P}$.

So suppose $(M_v)_P$ is not tight. By lemma 8.4 $\mathcal{T}_v$ is valid, so $(M_v)_P$ is not singular, hence twisted or critical. By lemma 8.15, $\widetilde{(\mathcal{T}_v)_P}$ has a vertex with a twisted label.

Now $\widetilde{(\mathcal{T}_v)_P}$ is the operation subtree of $\widetilde{\mathcal{T}_P}$ below v, with the the same vertex labels, consisting of everything in $\widetilde{\mathcal{T}_P}$ from v down. Thus, any lowest vertex with twisted label in $\widetilde{(\mathcal{T}_v)_P}$ is also a lowest vertex in $\widetilde{\mathcal{T}_P}$ with twisted label. As $(M_v)_P$ is twisted or critical, and is a sub-tensor-product of M_P, then M_P is also twisted or critical (lemma 4.29). By strong validity of $\mathcal{T}$ and lemma 8.16(iii), there is a unique lowest vertex in $\widetilde{\mathcal{T}_P}$ with twisted label. As $\widetilde{(\mathcal{T}_v)_P}$ has a vertex with twisted label, the unique lowest vertex in $\widetilde{\mathcal{T}_P}$ with twisted label lies in $\widetilde{(\mathcal{T}_v)_P}$, and it is also the unique lowest vertex in $\widetilde{(\mathcal{T}_v)_P}$ with twisted label. □

Finally, strong validity implies the following nice separation property of non-tight matched pairs.

Lemma 8.23. *Let $\mathcal{T}$ be a strongly valid operation tree. Let v, w be vertices of $\mathcal{T}$, with labels M_v, M_w respectively, such that the operation subtrees $\mathcal{T}_v, \mathcal{T}_w$ below v, w are disjoint.*

For any matched pair P, at least one of M_v, M_w is tight at P.

The disjointness of $\mathcal{T}_v, \mathcal{T}_w$ is equivalent to neither of v, w being a descendant of the other.

Proof. Suppose to the contrary that both $(M_v)_P, (M_w)_P$ are not tight. By lemma 8.22, $\mathcal{T}_v$ and $\mathcal{T}_w$ are strongly valid, so there is a unique lowest vertex x_v in $\widetilde{(\mathcal{T}_v)_P}$ with twisted label, and a unique lowest vertex x_w in $\widetilde{(\mathcal{T}_w)_P}$ with twisted label. But then x_v, x_w are two distinct vertices of $\widetilde{\mathcal{T}_P}$ which are lowest vertices with twisted labels, contradicting strong validity of $\mathcal{T}$. □

8.8 Transplantation and branch shifts

We now define two further methods to modify operation trees.

The first method, *transplantation*, replaces an operation subtree (definition 8.2) with another tree.

Definition 8.24. *Let $\mathcal{T}$ be an operation tree, and let $\mathcal{T}_v$ be the operation subtree below a non-root vertex v, labelled M'. Let $\mathcal{T}'$ be another operation tree for M'. Then removing $\mathcal{T}_v$ from $\mathcal{T}$ and replacing it with $\mathcal{T}'$ gives an operation tree $\mathcal{U}$. We say $\mathcal{U}$ is obtained from $\mathcal{T}$ by* transplanting *$\mathcal{T}'$ for $\mathcal{T}_v$.*

It is easily verified $\mathcal{U}$ is in fact an operation tree; viability of labels in $\mathcal{T}$ and $\mathcal{T}'$ implies viability of labels in $\mathcal{U}$. If $\mathcal{T}$ is an operation tree for M, then $\mathcal{U}$ is also an operation tree for M. So $\mathcal{T}$ and $\mathcal{U}$ describe operations on the same inputs, but the operations under v are rearranged.

Note that transplantation is quite different from grafting (section 8.3). Grafting adds to an operation tree below a leaf, while transplantation replaces part of an operation tree. Grafting adds new leaves with new labels, requiring relabelling throughout the tree, while leaf labels are unchanged under transplantation.

Lemma 8.25. *Suppose $\mathcal{U}$ is obtained from $\mathcal{T}$ by transplanting $\mathcal{T}'$ for $\mathcal{T}_v$.*

(i) If $\mathcal{T}$ and $\mathcal{T}'$ are valid, then $\mathcal{U}$ is also valid.

(ii) If $\mathcal{T}$ and $\mathcal{T}'$ are strongly valid, then $\mathcal{U}$ is also strongly valid.

Proof. All labels on vertices of $\mathcal{U}$ are inherited from $\mathcal{T}$ or $\mathcal{T}''$. If both $\mathcal{T}$ and $\mathcal{T}'$ are valid, then all labels are non-singular, so $\mathcal{U}$ is valid.

Now suppose $\mathcal{T}$ and $\mathcal{T}'$ are strongly valid. Let M, M' be the labels on the root vertex of $\mathcal{T}$, and v, respectively, and let P be a matched pair at which M is not tight. By strong validity of $\mathcal{T}$, there is a unique vertex w of $\mathcal{T}$ lowest among those with labels non-tight at P (lemma 8.16(ii)). Moreover, the vertices of $\mathcal{T}$ with labels non-tight at P are precisely the ancestors of w (lemma 8.20).

If w is not a vertex of $\mathcal{T}_v$, then v is not an ancestor of w, so the label M' of v is tight at P. Every vertex label in $\mathcal{T}'$ is a sub-tensor-product of M', hence tight at P (lemma 4.29). So the vertices of $\mathcal{U}$ with labels non-tight at P are precisely the vertices of $\mathcal{T}$ with labels non-tight at P, and hence there is a unique lowest such vertex, namely w.

If w is a vertex of $\mathcal{T}_v$, then v is an ancestor of w, so the label M' of v is non-tight at P. Since $\mathcal{T}'$ is strongly valid, there is a unique lowest vertex w' of $\mathcal{T}'$ with label non-tight at P (lemma 8.16(ii) again), and the set of vertices of $\mathcal{T}'$ whose labels are non-tight at P are precisely the ancestors of w' (lemma 8.20 again). Thus in $\mathcal{U}$, the set of vertices whose labels are non-tight at P are the ancestors of w' in $\mathcal{T}'$, together with the ancestors of v in $\mathcal{T}$ — in other words, the ancestors of w' in $\mathcal{U}$.

In any case, there is a unique vertex in $\mathcal{U}$ lowest among those with labels non-tight at P, so by lemma 8.16(ii) once more, $\mathcal{U}$ is strongly valid. ☐

The second method, a *branch shift*, rearranges an operation tree in a way corresponding to a modification $((AB)C) \leftrightarrow (A(BC))$.

Given an operation tree $\mathcal{T}$, denote the left and right children of the root vertex v by v_L and v_R, the left and right children of v_L by v_{LL} and v_{LR}, and generally for any word w in L and R, v_w denotes the descendant of v obtained by successively taking left or right children according to w (if it exists).

Definition 8.26. *The operation tree $\mathcal{T}'$ is defined by*

$$\mathcal{T}'_L = \mathcal{T}_{LL}, \quad \mathcal{T}'_{RL} = \mathcal{T}_{LR}, \quad \mathcal{T}'_{RR} = \mathcal{T}_R.$$

We say the operation trees $\mathcal{T}$ and $\mathcal{T}'$ are related by a branch shift*.*

The vertex labels on $\mathcal{T}'$ are either inherited from $\mathcal{T}$, or determined by the fact that each vertex is labelled with the tensor product of its children's labels.

Let $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$ respectively denote $\mathcal{T}_{LL}, \mathcal{T}_{LR}, \mathcal{T}_R$; let N_1, N_2, N_3 be the vertex labels on v_{LL}, v_{LR}, v_R respectively; let the root vertex of $\mathcal{T}'$ be v', and denote its vertices by v'_w for words w in L and R. Then in $\mathcal{T}$, v_L is labelled $N_1 \otimes N_2$; and in $\mathcal{T}'$, v'_R is labelled $N_2 \otimes N_3$. The viability of labels in $\mathcal{T}$ ensures the viability of labels in $\mathcal{T}'$, so both $\mathcal{T}$ and $\mathcal{T}'$ are both operation trees for $N = N_1 \otimes N_2 \otimes N_3$. Observe that upon reversing left and right, $\mathcal{T}$ is obtained from $\mathcal{T}'$ in the same way. See figure 23.

All labels in $\mathcal{T}'$ appear in $\mathcal{T}$, with one exception. Thus if $\mathcal{T}$ is valid, then we only have one label to check for validity of $\mathcal{T}'$, giving the following.

Lemma 8.27.

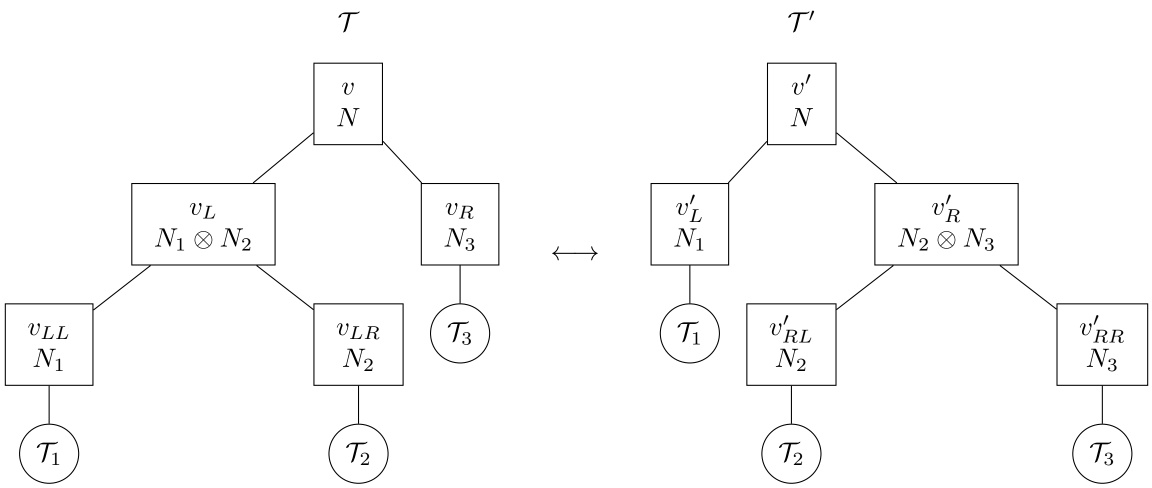

Figure 23: A branch shift.

(i) Suppose $\mathcal{T}$ is valid. Then $\mathcal{T}'$ is valid if and only if $N_2 \otimes N_3$ is non-singular.

(ii) Suppose $\mathcal{T}'$ is valid. Then $\mathcal{T}$ is valid if and only if $N_1 \otimes N_2$ is non-singular.

□

8.9 Strict distributivity

We now strengthen our notion of distributivity (definition 8.5).

Definition 8.28. *Let $\mathcal{T}$ be a valid operation tree for $M = M_1 \otimes \cdots \otimes M_n$.*

(i) Let v be a vertex of $\mathcal{T}$ with k leaves, labelled M_v. Then v is strictly distributive *if there are exactly $k-1$ matched pairs at which M_v is twisted or critical.*

(ii) The tree $\mathcal{T}$ is strictly f-distributive *if it is strictly distributive at each vertex.*

(iii) The tree $\mathcal{T}$ is strictly X-distributive *if it is strictly distributive at each non-root vertex, and there are precisely $n-2$ matched pairs at which M is twisted or critical.*

Recall distributivity (definition 8.5) at v requires at least $k-2$ twisted or critical matched pairs at v; the strict requirement is that there are precisely $k-1$ such pairs. Note that definition 8.28 requires $\mathcal{T}$ to be valid, so no labels are singular.

An obvious but useful property is the following.

Lemma 8.29. *Let $\mathcal{T}$ be a valid strictly f- or X-distributive operation tree, and let v be a non-root vertex. Then the operation subtree $\mathcal{T}_v$ of $\mathcal{T}$ below v is strictly f-distributive.*

Proof. By lemma 8.4 $\mathcal{T}_v$ is valid, and every vertex of $\mathcal{T}_v$, being a non-root vertex of $\mathcal{T}$, is strictly distributive. □

Strict distributivity imposes strong conditions on the function $V_{\mathcal{T}}$ (definition 8.19) of section 8.7.

Lemma 8.30. *Let $\mathcal{T}$ be an operation tree for $M = M_1 \otimes \cdots \otimes M_n$.*

(i) If $\mathcal{T}$ is strongly valid and strictly f-distributive, then $V_{\mathcal{T}}$ is a bijection between non-tight matched pairs of M and non-leaf vertices of $\mathcal{T}$.

(ii) If $\mathcal{T}$ is strongly valid and strictly X-distributive, then $V_{\mathcal{T}}$ is a bijection between non-tight matched pairs of M and non-leaf non-root vertices of $\mathcal{T}$.

Since M has n tensor factors, $\mathcal{T}$ has n leaves, hence $n-1$ non-leaf vertices and $n-2$ non-leaf non-root vertices. Strict f-distributivity (resp. X-distributivity) requires that M has precisely $n-1$ (resp. $n-2$) non-tight matched pairs. So in each case the claimed bijective sets have the same size.

Proof. When $n=1$, if $\mathcal{T}$ is strictly f-distributive, then M has no twisted or critical matched pairs (i.e. is tight), and $\mathcal{T}$ has no non-leaf vertices. When $n=2$, if $\mathcal{T}$ is strictly X-distributive, again M is tight, and $\mathcal{T}$ has no non-leaf non-root vertices. In both cases $V_{\mathcal{T}}$ is a bijection between empty sets.

We now proceed by induction on n. So suppose the result is true for operation trees for $M = M_1 \otimes \cdots \otimes M_k$, where $k < n$, and consider an operation tree $\mathcal{T}$ for $M = M_1 \otimes \cdots M_n$ which is strongly valid and strictly f-distributive or strictly X-distributive.

Let v_0 be the root vertex of $\mathcal{T}$, and let v_L and v_R be its left and right children; let their labels be $M_L = M_1 \otimes \cdots \otimes M_i$ and $M_R = M_{i+1} \otimes \cdots \otimes M_n$ respectively. Let $\mathcal{T}_L, \mathcal{T}_R$ be the operation subtrees of $\mathcal{T}$ below v_L and v_R (definition 8.2). Now $\mathcal{T}_L$ and $\mathcal{T}_R$ are strongly valid (lemma 8.22) and strictly f-distributive (lemma 8.29), so by induction we have bijections

$$V_L \colon \{\text{non-tight matched pairs of } M_L\} \longrightarrow \{\text{non-leaf vertices of } \mathcal{T}_L\}$$
$$V_R \colon \{\text{non-tight matched pairs of } M_R\} \longrightarrow \{\text{non-leaf vertices of } \mathcal{T}_R\},$$

which by lemma 8.22(ii) are restrictions of $V_{\mathcal{T}}$. Moreover, since $\mathcal{T}_L$ and $\mathcal{T}_R$ are disjoint, and $\mathcal{T}$ is strongly valid, lemma 8.23 says that the domains of V_L and V_R are disjoint. It's also clear that the ranges of V_L and V_R are disjoint; their union consists of all non-leaf non-root vertices of $\mathcal{T}$. The domains (and ranges) of V_L and V_R have cardinalities $i-1$ and $n-i-1$ respectively.

Since non-tight matched pairs in M_L or M_R are non-tight in M (lemma 4.30), the non-tight matched pairs in M_L and M_R form precisely $(i-1)+(n-i-1) = n-2$ non-tight matched pairs of M.

If M is strictly X-distributive, then these are all the matched pairs in M, and $V_{\mathcal{T}}$ is the disjoint union of V_L and V_R, hence a bijection as claimed.

If M is strictly f-distributive, then M has precisely $n-1$ non-tight matched pairs. So there is precisely one non-tight matched pair P_0 in M which is tight in M_L and M_R. Since P is tight in both M_L and M_R, but non-tight in M, v_0 is the lowest vertex of $\mathcal{T}$ whose label is non-tight at P_0, so $V_{\mathcal{T}}(P_0) = v_0$. This, together with V_L and V_R, defines $V_{\mathcal{T}}$; we conclude V is a bijection. □

8.10 Guaranteed nonzero results

We now show that, in certain cases, X_n and $\overline{f}_n$ must be nonzero, and compute their values.

Theorem 8.31. *Consider an A_∞ structure on $\mathcal{H}$ arising from a pair ordering $\preceq$. Suppose M is viable and satisfies the following conditions.*

(i) Every valid and distributive operation tree for M is strictly f-distributive, and such a tree exists.

(ii) No matched pair of M is on-on doubly occupied.

Then $\overline{f}_n(M) \neq 0$. Moreover $\overline{f}_n(M)$ is given by a single diagram D, which is tight at all matched pairs where M is tight or critical, and crossed at all matched pairs where M is twisted.

Theorem 8.31 is a precise version of theorem 1.4(i). It explicitly describes $\overline{f}_n(M) = D$, which is determined by its H-data, and tightness at each matched pair. Even when D is tight, there is no choice since there are no 11 doubly occupied pairs.

In fact, the description of D follows entirely from Maslov index considerations. The existence of a valid and strictly f-distributive tree for M implies that M is twisted or critical at precisely $n-1$ matched pairs, and tight at all other matched pairs. The Maslov index can only increase by 1 at each non-tight matched pair. Since $\overline{f}_n$ has Maslov grading $n-1$, Maslov grading must increase at every non-tight matched pair: from twisted to crossed, and from critical to tight.

When $n = 1$, condition (i) says that $M = M_1$ is tight (there is only one possible operation tree), and the conclusion is that $\overline{f}_1(M)$ is a tight diagram representing M_1

When $n = 2$, condition (i) says that $M = M_1 \otimes M_2$ has precisely one non-tight matched pair P (again there is only one possible operation tree), which must be twisted (lemma 4.28), and all other matched pairs tight. The conclusion is that $\overline{f}_2(M)$ is a single diagram D twisted at P and elsewhere tight, in agreement with the discussion of section 5.3.

Theorem 8.32. *Consider an A_∞ structure on $\mathcal{H}$ arising from a pair ordering $\preceq$. Suppose M is viable and satisfies the following conditions.*

(i) Every valid and distributive operation tree for M is strictly X-distributive, and such a tree exists.

(ii) No matched pair of M is twisted or on-on doubly occupied.

Then $X_n(M)$ is nonzero, and is the homology class of the unique tight diagram with the H-data of M.

Theorem 8.32 is a precise version of theorem 1.4(ii). The description of $X_n(M)$ follows entirely from the fact that X_n preserves H-data. The uniqueness claim in the theorem makes sense: since M has no 11 doubly occupied pairs, there is only one tight diagram with the same H-data as M.

The exclusion of twisted matched pairs is necessary, since they preclude the existence of a tight diagram (or by theorem 6.6). The exclusion of 11 doubly occupied pairs is purely to facilitate the proof, but it is straightforward to check that with our shorthand, the example $X_5(M_1 \otimes \cdots \otimes M_5)$ below, with a 11 doubly occupied matched pair $P = \{p, p'\}$, has two valid distributive operation trees, both of which are strongly valid and strictly X-distributive, and in fact $\overline{f}_1 \overline{f}_4 = \overline{f}_4 \overline{f}_1 \neq 0$, so $X_5 = 0$. So without this exclusion, the theorem fails.

$$X_5 \left(\begin{array}{c|c|c|c|c} p_+ & p_- & & p'_+ & p'_- \\ q'_- & & q_+ & q_- & \\ & r_+ & r_- & & r'_- \end{array} \right) = 0.$$

While the conditions of theorems 8.31 and 8.32 may seem rather restrictive, they do show that $\overline{f}_n$ and X_n are nonzero in many cases. For instance, in the $\overline{f}_3$ examples of lemma 7.2, the first two lines (i.e. 10 out of 14 examples) can be shown to be nonzero directly from theorem 8.31. It follows from theorem 8.32 that X_4 of all the following tensor products are nonzero.

$$\left(\begin{array}{c|c|c|c} p_+ & p_- & p'_+ & \\ q_+ & & q_- & q'_+ \end{array} \right) \quad \left(\begin{array}{c|c|c|c} p_+ & p_- & p'_+ & \\ & q_+ & q_- & q'_+ \end{array} \right) \quad \left(\begin{array}{c|c|c|c} p_+ & p_- & & p'_+ \\ & q_+ & q_- & q'_+ \end{array} \right) \quad \left(\begin{array}{c|c|c|c} p_+ & & p_- & p'_+ \\ q_+ & q_- & q'_+ & \end{array} \right)$$

$$\left(\begin{array}{c|c|c|c} p'_- & p_+ & p_- & p'_+ \\ q_+ & q_- & & q'_+ \end{array} \right) \quad \left(\begin{array}{c|c|c|c} p'_- & p_+ & p_- & p'_+ \\ q_+ & & q_- & q'_+ \end{array} \right) \quad \left(\begin{array}{c|c|c|c} p'_- & p_+ & p_- & p'_+ \\ q'_- & q_+ & & q_- \end{array} \right) \quad \left(\begin{array}{c|c|c|c} p'_- & p_+ & p_- & p'_+ \\ q'_- & & q_+ & q_- \end{array} \right)$$

The hypotheses of theorems 8.31 and 8.32 essentially mandate that in each operation described by an operation tree, only one matched pair can be affected.

We first need a preliminary lemma.

Lemma 8.33 (Plenty of trees)**.** *Consider an A_∞ structure on $\mathcal{H}$ defined by a pair ordering. Suppose $M = M_1 \otimes \cdots \otimes M_n$ is viable. Further suppose that every valid and distributive operation tree for M is strongly valid and strictly f-distributive, and at least one such tree exists.*

Let $P_0 = \{p_0, p'_0\}$ be a matched pair at which M is twisted. Then there exists a strongly valid, strictly f-distributive operation tree $\mathcal{T}$ for M such that $V_{\mathcal{T}}(P)$ is the root vertex v_0 of $\mathcal{T}$.

Let us say something about what lemma 8.33 means. At P_0, M is twisted and hence an extension of the twisted tensor product of table 2 (lemma 4.27). So two steps of P_0 are covered, say p_{0+} and

p_{0-}, by some M_i and M_j respectively for some $i < j$. Now a sub-tensor-product M' of M labelling a non-root vertex of $\mathcal{T}$ is twisted or tight accordingly as M' contains both M_i and M_j, or does not. Lemma 8.33 guarantees the existence of a tree such that all labels on non-root vertices are tight at P_0. In other words, M_i and M_j never appear together in any label in $\mathcal{T}$ except at the root vertex v_0; as we work our way up the tree, combining tensor factors, M_i and M_j are only combined at the final step, at v_0. Since P_0 only becomes twisted at v_0, v_0 is the lowest vertex of $\mathcal{T}$ whose label is twisted at P, and $V_{\mathcal{T}}(P_0) = v_0$.

Since we can find such a tree for each all-on once occupied pair P, this gives us "plenty of trees", which we need for the proof of theorem 8.31.

Note the hypotheses of lemma 8.33 are weaker than those of theorem 8.31. If M satisfies the hypotheses of theorem 8.31, then every valid distributive operation tree for M is strictly f-distributive; but as there are no 11 doubly occupied pairs, any such tree is strongly valid (lemma 8.18), so M satisfies the hypotheses of lemma 8.33.

The following lemma captures an argument we will use repeatedly.

Lemma 8.34. *Let M be a viable tensor product of nonzero homology classes of diagrams, which has one of the following two properties (where the terms in square brackets may be included or not):*

(i) Every valid and distributive operation tree for M is [strongly valid and] strictly f-distributive, and such a tree exists.

(ii) Every valid and distributive operation tree for M is [strongly valid and] strictly X-distributive, and such a tree exists.

Let $\mathcal{T}$ be an operation tree for M of the type guaranteed by the condition, and let v be a non-root vertex of $\mathcal{T}$, with label M_v. Then M_v satisfies condition (i).

Proof. Let $\mathcal{T}'$ be a valid distributive operation tree for M_v. Then we can transplant $\mathcal{T}'$ for the operation subtree $\mathcal{T}_v$ of $\mathcal{T}$ below v to obtain an operation tree $\mathcal{U}$ for M, which is valid (lemma 8.25) and distributive (since distributive at each vertex: definition 8.5). By assumption then $\mathcal{U}$ is [strongly valid and] strictly f- or X-distributive, so its subtree $\mathcal{T}'$ is also [strongly valid (lemma 8.22) and] strictly f-distributive (lemma 8.29). Finally, $\mathcal{T}_v$ demonstrates that such a tree exists. □

Proof of lemma 8.33. When $n = 1$ the statement is vacuous: $M = M_1$ is tight, the unique operation tree is strongly valid and strictly f-distributive, and $V_{\mathcal{T}}$ is a bijection between empty sets. Proceeding by induction on n, consider a general n, and suppose the result holds for all smaller values of n.

Let $\mathcal{T}$ be a strongly valid and strictly f-distributive operation tree for M, which exists by assumption. By strict f-distributivity at v_0, there are precisely $n - 1$ matched pairs at which M is non-tight (i.e. twisted or critical). Let the two children of v_0 be v_L and v_R, with labels $M_L = M_1 \otimes \cdots \otimes M_m$ and $M_R = M_{m+1} \otimes \cdots \otimes M_n$ respectively. Let $\mathcal{T}_L, \mathcal{T}_R$ be the operation subtrees of $\mathcal{T}$ below v_L, v_R respectively. Then $\mathcal{T}_L, \mathcal{T}_R$ are strongly valid (lemma 8.22) and strictly f-distributive (lemma 8.29) operation trees for M_L, M_R respectively.

By lemma 8.30, $V_{\mathcal{T}}$, $V_{\mathcal{T}_L}$ and $V_{\mathcal{T}_R}$ are all bijections, between sets of size $n-1$, $m-1$ and $n-m-1$, respectively; moreover $V_{\mathcal{T}_L}$ and $V_{\mathcal{T}_R}$ are restrictions of $V_{\mathcal{T}}$ (lemma 8.22) with disjoint domains (lemma 8.23). Hence there is a unique matched pair P_1 such that $V_{\mathcal{T}}(P_1) = v_0$. Then P_1 is twisted in M (definition 8.19), but tight in every other tensor product labelling a vertex.

If $P_1 = P_0$ then we are done; so suppose that P_1 and P_0 are distinct. Then $V_{\mathcal{T}}(P_0) \neq V_{\mathcal{T}}(P_1) = v_0$, so $V_{\mathcal{T}}(P_0)$ is a vertex of $\mathcal{T}_L$ or $\mathcal{T}_R$. Suppose $V_{\mathcal{T}}(P_0)$ lies in $\mathcal{T}_L$; the $\mathcal{T}_R$ case is similar.

By lemma 8.34, M_L satisfies the hypotheses of this lemma. By induction there then exists a strongly valid, strictly f-distributive operation tree $\mathcal{T}'_L$ for M_L, such that $V_{\mathcal{T}'_L}(P_0) = v_L$. Transplanting this $\mathcal{T}'_L$ for $\mathcal{T}_L$ yields a strongly valid (by lemma 8.25) and strictly f-distributive (since strictly distributive at each vertex: definition 8.28) operation tree $\mathcal{T}'$ for M. Moreover, $V_{\mathcal{T}'_L}$ is a restriction of $V_{\mathcal{T}'}$ (lemma 8.22), so $V_{\mathcal{T}'}(P_0) = v_L$, and since P_1 is tight in M_L and M_R, $V_{\mathcal{T}'}(P_1) = v_0$.

Let the children of v_L be v_{LL} and v_{LR}, and let their labels in $\mathcal{T}'$ be $M'_{LL} = M_1 \otimes \cdots \otimes M_k$ and $M'_{LR} = M_{k+1} \otimes \cdots \otimes M_m$. Denote the operation subtrees of $\mathcal{T}'$ (or $\mathcal{T}'_L$) below v_{LL}, v_{LR} respectively by $\mathcal{T}'_{LL}, \mathcal{T}'_{LR}$. These are again strongly valid and strictly f-distributive (lemmas 8.22 and 8.29).

By strict f-distributivity, M'_{LL}, M'_{LR}, M_R, M have precisely $k-1$, $m-k-1$, $n-m-1$, $n-1$ non-tight matched pairs respectively. But since $\mathcal{T}'_{LL}, \mathcal{T}'_{LR}$ and $\mathcal{T}_R$ are disjoint subtrees (below v_{LL}, v_{LR}, v_R) of the strongly valid $\mathcal{T}'$, the sets of matched pairs at which M'_{LL}, M'_{LR}, M_R are non-tight are also disjoint (lemma 8.23). Their union consists of $(k-1)+(m-k-1)+(n-m-1) = n-3$ matched pairs, which remain non-tight in M (lemma 4.30). The two remaining non-tight matched pairs of M are P_0 and P_1; these two pairs are tight in each of M'_{LL}, M'_{LR}, M_R since $V_{\mathcal{T}'}(P_0) = v_L$ and $V_{\mathcal{T}'}(P_1) = v_0$.

Now perform a branch shift on $\mathcal{T}'$ (definition 8.26) to obtain an operation tree $\mathcal{T}''$ for M. Its root has children $v''_L = v_{LL}$ and v''_R, and the children of v''_R are $v''_{RL} = v_{LR}$ and $v''_{RR} = v_R$. Below $v''_L, v''_{RL}, v''_{RR}$ respectively we have $\mathcal{T}''_L = \mathcal{T}'_{LL}$, $\mathcal{T}''_{RL} = \mathcal{T}'_{LR}$, and $\mathcal{T}''_{RR} = \mathcal{T}_R$. The labels on $\mathcal{T}''$ are inherited from $\mathcal{T}'_{LL}, \mathcal{T}'_{LR}, \mathcal{T}'_R$, except that v''_0 is labelled M and v''_R is labelled with $M''_R = M_{k+1} \otimes \cdots \otimes M_n = M'_{LR} \otimes M_R$. In particular, $v''_L, v''_{RL}, v''_{RR}$ are respectively labelled with $M''_L = M'_{LL}$, $M''_{RL} = M'_{LR}$ and $M''_{RR} = M_R$.

We claim $\mathcal{T}''$ is valid. If P is a matched pair non-tight in M, other than P_0 or P_1, then P is twisted in the label of $V_{\mathcal{T}'}(P)$ (definition 8.19), which is a vertex of one of $\mathcal{T}'_{LL} = \mathcal{T}''_L$, $\mathcal{T}'_{LR} = \mathcal{T}''_{RL}$, or $\mathcal{T}_R = \mathcal{T}''_{RR}$. And P_0, P_1 are twisted in M, which is the label of the root. Thus for every matched pair P, there is a vertex of $\mathcal{T}''$ whose label is twisted at P. By lemma 8.15 then $\mathcal{T}''$ is valid.

We also claim $\mathcal{T}''$ is distributive. Each vertex of $\mathcal{T}''$ which shares a label with a vertex of distributive tree $\mathcal{T}'$ is distributive. The only remaining vertex is v''_R, which has label $M''_R = M_{k+1} \otimes \cdots \otimes M_n = M'_{LR} \otimes M_R$. Each of the $(m-k-1)+(n-m-1) = n-k-2$ matched pairs P such that $V_{\mathcal{T}'}(P)$ is a vertex of $\mathcal{T}'_{LR}$ or $\mathcal{T}_R$ is non-tight in M'_{LR} or M_R, hence also in $M''_R = M'_{LR} \otimes M_R$ (lemma 4.30). Since there are $n-k$ leaves below v''_R, and there are at least $n-k-2$ matched pairs at which M''_R is twisted or critical, v''_R is distributive, and the claim follows.

Since $\mathcal{T}''$ is valid and distributive, by assumption then $\mathcal{T}''$ is strongly valid and strictly f-distributive. Now P_0 is twisted in M and satisfies $V_{\mathcal{T}'}(P_0) = v_L$, so P_0 is twisted in $M_L = M_1 \otimes \cdots \otimes M_m$, but tight in $M'_{LL} = M_1 \otimes \cdots \otimes M_k$ and $M'_{LR} = M_{k+1} \otimes \cdots \otimes M_m$. Supposing without loss of generality that P_0 is twisted at p_0 in M, then the step p_{0+} must be covered by one of $M_1, \ldots, M_k$, and the step p_{0-} must be covered by one of $M_{k+1}, \ldots, M_m$, with no steps of P covered by any of $M_{m+1}, \ldots, M_n$. Thus P_0 is tight in $M''_R = M_{k+1} \otimes \cdots \otimes M_n$, and in $M''_L = M_1 \otimes \cdots \otimes M_k$, the labels of v''_L and v''_R; but P_0 is twisted in M, the label of v_0. So $V_{\mathcal{T}''}(P_0) = v_0$, and $\mathcal{T}''$ is the desired tree. By induction, the proof is complete. □

Proof of theorem 8.31. We have verified the theorem in small cases, so suppose it is true for all $\overline{f}_k$ with $k < n$, and consider $\overline{f}_n$.

By lemma 8.18, since there are no 11 doubly occupied pairs in M, validity and strong validity are equivalent; we use this fact repeatedly.

Our strategy is to compute $\overline{U}_n(M)$ explicitly, and then compute $\overline{f}_n$, using the construction of corollary 5.3. Recall $\overline{U}_n(M)$ is a sum of terms of the form $\overline{f}_i(M_1 \otimes \cdots \otimes M_i)\overline{f}_{n-i}(M_{i+1} \otimes \cdots \otimes M_n)$, and $\overline{f}_{n-j+1}(M_1 \otimes \cdots \otimes M_k \otimes X_j(M_{k+1} \otimes \cdots \otimes M_{k+j}) \otimes \cdots \otimes M_n)$.

The latter type of term is easiest to deal with: we claim they are all zero. Suppose to the contrary that $\overline{f}_{n-j+1}(M_1 \otimes \cdots \otimes M_k \otimes X_j(M_{k+1} \otimes \cdots \otimes M_{k+j} \otimes \cdots \otimes M_n) \neq 0$. Then by proposition 8.10 there are valid distributive operation trees $\mathcal{T}_X$ for $M_{k+1} \otimes \cdots \otimes M_{k+j}$ and $\mathcal{T}_f$ for $M_1 \otimes \cdots \otimes M_k \otimes X_j(M_{k+1} \otimes \cdots \otimes M_{k+j}) \otimes \cdots \otimes X_n$. Grafting $\mathcal{T}_X$ onto $\mathcal{T}_f$ at position $k+1$ yields (lemma 8.9) a valid distributive operation tree $\mathcal{T}_{fX}$ for M. By assumption then $\mathcal{T}_{fX}$ is strictly f-distributive. Applying strict distributivity to the vertex of $\mathcal{T}_{fX}$ corresponding to the root of $\mathcal{T}_X$, then $M_{k+1} \otimes \cdots \otimes M_{k+j}$ is twisted or critical at precisely $j-1$ matched pairs. But by theorem 6.6, $X_j(M_{k+1} \otimes M_{k+j}) \neq 0$ implies that there are precisely $j-2$ such pairs. This gives a contradiction, so all such terms are zero.

We now consider terms of the form $\overline{f}_i(M_1 \otimes \cdots \otimes M_i)\overline{f}_{n-i}(M_{i+1} \otimes \cdots \otimes M_n)$ which are nonzero. We will associate to them matched pairs at which M is twisted and eventually obtain a bijection $F \colon A \longrightarrow B$ where

$$A = \{i \mid \overline{f}_i(M_1 \otimes \cdots \otimes M_i)\overline{f}_{n-i}(M_{i+1} \otimes \cdots \otimes M_n) \neq 0\}, \quad B = \{P \mid M \text{ is twisted at } P\}.$$

So let $M' = M_1 \otimes \cdots \otimes M_i$ and $M'' = M_{i+1} \otimes \cdots \otimes M_n$ and suppose $\overline{f}_i(M')\overline{f}_{n-i}(M'') \neq 0$. By proposition 8.10 there are valid (hence strongly valid) distributive trees $\mathcal{T}'$ for M' and $\mathcal{T}''$ for M''. Joining these trees yields an operation tree $\mathcal{T}$ for M (definition 8.6), which is valid (hence strongly valid) and distributive (lemma 8.7), hence by hypothesis strictly f-distributive. We then have a bijection $V_\mathcal{T}$ between non-tight matched pairs of M and non-leaf vertices of $\mathcal{T}$ (lemma 8.30). Moreover, since $\mathcal{T}', \mathcal{T}''$ are subtrees of $\mathcal{T}$, they are also strictly f-distributive (lemma 8.29). Thus M, M', M'' are twisted or critical at $n-1$, $i-1$, $n-i-1$ matched pairs respectively, and tight elsewhere.

By lemma 8.34, any valid distributive tree for M' is strongly valid and strictly f-distributive; and similarly for M''. And since M has no 11 doubly occupied matched pairs, neither do the sub-tensor-products M' or M''. So the hypotheses of the theorem apply to M' and M''. By induction then $\overline{f}_i(M')$ and $\overline{f}_{n-i}(M'')$ are given by single diagrams as described in the statement. Moreover, as $\mathcal{T}', \mathcal{T}''$ are disjoint subtrees of the strongly valid $\mathcal{T}$, the matched pairs at which M', M'' are non-tight are disjoint (lemma 8.23). This yields $(i-1)+(n-i-1) = n-2$ matched pairs at which M' or M'' is non-tight; such pairs are also non-tight in M (lemma 4.30). So there is precisely one matched pair P_i at which M is non-tight but M' and M'' are tight. Then $V_\mathcal{T}(P_i)$ is the root vertex v_0, and P_i is twisted in M (by the comment after definition 8.19, or lemma 4.28). Indeed, $V_\mathcal{T}(P_i)$ is the root vertex, for any $\mathcal{T}$ arising as the join of valid distributive operation trees for M' and M''. Define the function $F \colon A \longrightarrow B$ by $F(i) = P_i$.

We now describe the diagram representing $\overline{f}_i(M')\overline{f}_{n-i}(M'')$, at each matched pair P.

First, suppose P is critical in M. Then $P \neq P_i$, so P is non-tight in precisely one of M' or M''. Considering the various cases in the critical column of table 2, of which M is an extension (lemma 4.27), and how the P-active factors can be distributed across M' and M'', we observe that in every case $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ is tight at P.

Second, suppose P is a matched pair at which M is twisted, other than P_i. Then P is non-tight in precisely one of M' or M''. Indeed, there are two P-active factors and they are both in M', or both in M''. So $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ at P is the product of an all-on once occupied crossed diagram, and an idempotent, hence is crossed.

Third, suppose P is tight in M. Then P is also tight in M' and M'' (lemma 4.29), hence also in $\overline{f}_i(M')$ and $\overline{f}_{n-i}(M'')$ (by inductive assumption). So $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ at P is given by multiplying factors in a tight tensor product, hence is tight.

Finally, at P_i, M' and M'' are both tight, but M is twisted. Hence P_i is 11 once occupied by M, with one step covered by M', and the other by M''; by inductive assumption then $\overline{f}_i(M')$ and $\overline{f}_{n-i}(M'')$ are both tight at P_i, so $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ is twisted at P_i.

To summarise: when $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ is nonzero, there is a unique matched pair P_i which is non-tight (in fact twisted) in M but tight in M' and M''; $V_\mathcal{T}(P_i)$ is the root vertex of $\mathcal{T}$; and $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ is given by a single diagram which is twisted at P_i, crossed at all other matched pairs which are twisted in M, and tight at all other matched pairs. We set $F(i) = P_i$.

Now we claim that F is injective. Consider another nonzero term $\overline{f}_j(M_1 \otimes \cdots \otimes M_j)\overline{f}_{n-j}(M_{j+1} \otimes \cdots \otimes M_n)$, where $i \neq j$. We consider the case $i < j$; the case $i > j$ is similar. Applying the same argument as above, we obtain strongly valid and strictly f-distributive trees $\mathcal{T}_j, \mathcal{T}'_j, \mathcal{T}''_j$ for M, $M'_j = M_1 \otimes \cdots \otimes M_j$ and $M''_j = M_{j+1} \otimes \cdots \otimes M_n$ respectively. The bijection $V_{\mathcal{T}_j}$ between non-tight matched pairs of M and non-leaf vertices of $\mathcal{T}_j$. The matched pair P_j has $V_{\mathcal{T}_j}(P_j)$ as the root of $\mathcal{T}_j$, and $F(j) = P_j$. We will show that $P_i \neq P_j$.

Now in the valid distributive tree $\mathcal{T}$ constructed above for $\overline{f}_i(M')\overline{f}_{n-i}(M'')$, let v be the lowest common ancestor of the leaves labelled M_j and M_{j+1}. Let P be the matched pair such that $V_\mathcal{T}(P) = v$ (well-defined since $V_\mathcal{T}$ is bijective). Since $i < j$, v is a vertex of $\mathcal{T}''$, hence not the root, so $P \neq P_i$. The label M_v of v is then twisted at P (definition 8.19), say at the place p. So some M_a, with $a \leq j$, covers the step p_+, and some M_b, with $j+1 \leq b$, covers p_-. As M contains no 11 doubly occupied pairs, any sub-tensor-product of M which is twisted at P must contain M_a and M_b.

Now consider $V_{\mathcal{T}_j}(P)$, a vertex of $\mathcal{T}_j$; call its label $M_\#$. Then $M_\#$ is twisted at P (definition 8.19), so $M_\#$ contains M_a and M_b as tensor factors. But since $a \leq j$ and $b \geq j+1$, $M_\#$ cannot be a sub-tensor-product of M'_j or M''_j; thus $M_\# = M$ and $V_{\mathcal{T}_j}(P)$ is the root vertex. Thus $P = P_j$. As

$P \neq P_i$ then $P_i \neq P_j$. Thus F is injective.

We now show F is surjective. Take a matched pair P at which M is twisted; we will show $P = P_i$ for some i. By lemma 8.33 (which, as discussed above, has weaker hypotheses than the present theorem) there is a strongly valid, strictly f-distributive operation tree $\mathcal{T}^*$ for M such that $V_{\mathcal{T}^*}(P)$ is the root vertex v_0^* of $\mathcal{T}^*$. Let the children of v_0^* be v_L^*, v_R^*, with labels $M_L^* = M_1 \otimes \cdots \otimes M_i$ and $M_R^* = M_{i+1} \otimes \cdots \otimes M_n$ respectively. Then by definition of $V_{\mathcal{T}^*}$, P is tight in M_L^* and M_R^*. By lemma 8.34, M_L^* and M_R^* satisfy condition (i) of the present theorem; and as M_L^* and M_R^* are sub-tensor-products of M, which has no 11 doubly occupied pairs, they satisfy condition (ii) also. So by induction $\overline{f}_i(M_L^*)$ and $\overline{f}_{n-i}(M_R^*)$ are both nonzero, given by single diagrams as described in the statement. By lemma 8.23 they are non-tight at disjoint matched pairs. Examining the various possible cases at each matched pair, we conclude that $\overline{f}_i(M_L^*)\overline{f}_{n-i}(M_R^*) \neq 0$. Since P is non-tight in M but tight in M_L^* and M_R^*, we have $P = P_i$. So F is surjective, hence a bijection.

Returning to $\overline{U}_n(M)$, we now see that each nonzero terms of $\overline{U}_n(M)$ is of the form $\overline{f}_i(M')\overline{f}_{n-i}(M'')$, and these terms correspond bijectively to the matched pairs P_i at which M is twisted. In fact $\overline{f}_i(M')\overline{f}_{n-i}(M'')$ is twisted at P_i, and crossed at all other matched pairs where M is twisted.

We also observe that $X_n(M) = 0$, since M has precisely $n-1$ non-tight matched pairs, by theorem 6.6. Thus, following the construction of corollary 5.3 and the discussion of section 3.3,

$$\overline{f}_n(M) = \overline{A}^*_{\mathcal{CR}^{\preceq}} \overline{U}_n(M).$$

By definition 3.21, $A^*_{\mathcal{CR}^{\preceq}}$ applies a creation operator at P_{min}, where P_{min} is the $\preceq$-minimal matched pair among pairs where M is twisted.

We observe that there is precisely one diagram in $\overline{U}_n(M)$ which is twisted at P_{min}, namely $\overline{f}_i \overline{f}_{n-i}$ where $P_i = P_{min}$. Applying $\overline{A}^*_{\mathcal{CR}^{\preceq}} = \overline{A}^*_{P_{min}}$ inserts a crossing at P_{min} to this diagram. All the other diagrams in $\overline{U}_n(M)$ are crossed at P_{min}, and applying the creation operator gives zero.

We conclude that $\overline{f}_n(M)$ is given by a single diagram, crossed at all matched pairs where M is twisted, and tight elsewhere, as desired. □

Proof of theorem 8.32. As there are no 11 occupied matched pairs, by lemma 8.18, validity and strong validity are equivalent.

By lemma 6.4 (since all the maps f_k in the pair ordering construction are balanced), $X_n(M)$ is represented by the sum of all terms of the form $\overline{f}_i(M_1 \otimes \cdots \otimes M_i)\overline{f}_{n-i}(M_{i+1} \otimes \cdots \otimes M_n)$.

Let $\mathcal{T}$ be a valid and strictly X-distributive operation tree for M, which exists by hypothesis. Let its root vertex be v_0, with children v_L, v_R labelled $M_L = M_1 \otimes \cdots \otimes M_i$, $M_R = M_{i+1} \otimes \cdots \otimes M_n$. Let $\mathcal{T}_L, \mathcal{T}_R$ be the subtrees below v_L, v_R respectively.

By lemma 8.34, M_L and M_R satisfy condition (i) of theorem 8.31; and being sub-tensor-products of M, which has no 11 doubly occupied pairs, M_L and M_R also satisfy condition (ii). So by theorem 8.31, $\overline{f}_i(M_L)$ and $\overline{f}_{n-i}(M_R)$ are both nonzero, given by single diagrams. Since $\mathcal{T}$ is strictly X-distributive, M_L and M_R have $i-1, n-i-1$ non-tight matched pairs respectively. These sets of non-tight matched pairs are distinct by lemma 8.23, and also non-tight in M (lemma 4.30); hence they provide $n-2$ non-tight matched pairs in M. Again by strict X-distributivity, M has precisely $n-2$ non-tight matched pairs, so each non-tight matched pair of M is non-tight in precisely one of M_L or M_R.

By theorem 8.31, $\overline{f}_i(M_L)$ (resp. $\overline{f}_{n-i}(M_R)$) is crossed at every matched pair where M_L (resp. M_R) is twisted, and elsewhere tight. Thus at every non-tight (hence critical) matched pair of M, precisely one of M_L, M_R is non-tight (twisted or critical), and the other is tight. If one of M_L, M_R is critical and the other is tight, then $\overline{f}_i(M_L)$ and $\overline{f}_i(M_R)$ are tight, and by reference to table 2 or otherwise, $\overline{f}_i(M_l) \otimes \overline{f}_{n-i}(M_R)$ is tight. If one of M_L, M_R is twisted and the other is tight, then one of $\overline{f}_i(M_L), \overline{f}_{n-i}(M_R)$ is crossed, and the other is tight, so again by reference to table 2 or otherwise, $\overline{f}_i(M_L) \otimes \overline{f}_{n-i}(M_R)$ is sublime. Either way, $\overline{f}_i(M_L)\overline{f}_{n-i}(M_R)$ is tight at each non-tight matched pair of M. At tight matched pairs of M, $\overline{f}_i(M_L)$ and $\overline{f}_i(M_R)$ are both tight, with tight product. So $\overline{f}_i(M_L)\overline{f}_{n-i}(M_R)$ is the unique tight diagram with the same H-data as M.

Now let P be a non-tight matched pair of M. By assumption, P is critical, but not 11 doubly occupied. Thus, by reference to table 2, P is sesqui-occupied or 00 doubly occupied and M_P is an extension of one of the corresponding critical diagrams shown there (proposition 4.27). In particular, there is precisely one place p of P such that the steps p_+ and p_- are covered by some M_a and M_b respectively, where $a < b$. We call these the *principal factors* of P. Now if $a \leq i < i+1 \leq b$, then considering the various cases of table 2, P is singular in M_L or M_R, contradicting validity of $\mathcal{T}$. Thus a, b are both $\leq i$, or both $\geq i+1$. In other words, for any non-tight matched pair of M, its principal factors have positions which are both $\leq i$, or both $\geq i+1$; they do not cross the i'th position.

On the other hand, for any $1 \leq j \leq n-1$ with $j \neq i$, the least common ancestor w of the leaves labelled M_j and M_{j+1} lies in $\mathcal{T}_L$ or $\mathcal{T}_R$, accordingly as $i > j$ or $i < j$. We suppose $i < j$, so $w \in \mathcal{T}_R$; the $\mathcal{T}_L$ case is similar. Clearly w is neither a leaf nor root, so by lemma 8.30, there is a unique matched pair P such that $V_{\mathcal{T}}(P) = w$. Letting M_w denote the label of w, then M_w is twisted at P. Letting w_L, w_R denote the children of v, their labels are tight at P. The label on w_L contains M_a, so by construction $a \leq j$. Similarly the label of w_R contains M_b, and $j+1 \leq b$. So the two principal factors have positions with are $\leq j$ and $\geq j+1$ respectively.

Thus, for any $j \neq i$, there is a non-tight matched pair of P whose principal factors have positions $\leq j$ and $\geq j+1$.

It follows that for any $j \neq i$, we must have $\overline{f}_j(M_1 \otimes \cdots \otimes M_j)\overline{f}_{n-j}(M_{j+1} \otimes \cdots \otimes M_n) = 0$. For if this product were nonzero, then we could repeat the argument above and find that no non-tight matched pair of M has principal factors whose positions cross the j'th position, contradicting the previous paragraph.

We conclude that $X_n(M)$ is the homology class of the single diagram $\overline{f}_i(M_L)\overline{f}_{n-i}(R)$, which has the desired properties. □

References

[1] Benjamin Cooper, *Formal Contact Categories*, (2015), arXiv: 1511.04765.

[2] Y. Eliashberg, *Classification of overtwisted contact structures on 3-manifolds*, Inventiones Mathematicae **98** (1989), no. 3, 623–637.

[3] Peter John Hilton and Urs Stammbach, *A course in homological algebra*, Springer-Verlag, New York-Berlin, 1971.

[4] Ko Honda, *On the classification of tight contact structures. I*, Geometry and Topology **4** (2000), 309–368.

[5] ______, *Gluing tight contact structures*, Duke Mathematical Journal **115** (2002), no. 3, 435–478.

[6] Ko Honda and Yin Tian, *Contact categories of disks*, (2016), arXiv: 1608.08325.

[7] Johannes Huebschmann, *On the construction of A-infinity structures*, Georgian Mathematical Journal **17** (2010), no. 1, 161–202.

[8] T. V. Kadeishvili, *On the homology theory of fibre spaces*, Russian Mathematical Surveys **35** (1980), no. 3, 231 (en).

[9] Bernhard Keller, *Introduction to A-infinity algebras and modules*, (1999), arXiv:math/9910179.

[10] Maxim Kontsevich and Yan Soibelman, *Homological mirror symmetry and torus fibrations*, Symplectic geometry and mirror symmetry (Seoul, 2000), World Sci. Publ., River Edge, NJ, 2001, pp. 203–263.

[11] Robert Lipshitz, Peter Ozsvath, and Dylan Thurston, *Bordered Heegaard Floer homology: Invariance and pairing*, (2008), arXiv: 0810.0687.

[12] ______, *Slicing planar grid diagrams: a gentle introduction to bordered Heegaard Floer homology*, (2008), arXiv: 0810.0695.

[13] Robert Lipshitz, Peter S. Ozsvath, and Dylan P. Thurston, *Bimodules in bordered Heegaard Floer homology*, Geometry & Topology **19** (2015), no. 2, 525–724.

[14] Robert Lipshitz, Peter Ozsvth, and Dylan Thurston, *Notes on bordered Floer homology*, (2012), arXiv: 1211.6791.

[15] Daniel Mathews, *Chord diagrams, contact-topological quantum field theory and contact categories*, Algebraic & Geometric Topology **10** (2010), no. 4, 2091–2189 (en).

[16] Daniel V. Mathews, *Sutured Floer homology, sutured TQFT and non-commutative QFT*, Algebraic & Geometric Topology **11** (2011), no. 5, 2681–2739.

[17] ———, *Sutured TQFT, torsion and tori*, International Journal of Mathematics **24** (2013), no. 05, 1350039.

[18] ———, *Contact topology and holomorphic invariants via elementary combinatorics*, Expositiones Mathematicae **32** (2014), no. 2, 121–160.

[19] ———, *Itsy Bitsy Topological Field Theory*, Annales Henri Poincaré **15** (2014), no. 9, 1801–1865 (en).

[20] ———, *Twisty itsy bitsy topological field theory*, International Journal of Mathematics **25** (2014), no. 10, 1450097 (en).

[21] ———, *Strand algebras and contact categories*, (2016), arXiv: 1608.02710.

[22] ———, *Strings, fermions and the topology of curves on annuli*, Journal of Symplectic Geometry **15** (2017), no. 2, 421–506 (EN).

[23] Daniel V. Mathews and Eric Schoenfeld, *Dimensionally reduced sutured Floer homology as a string homology*, Algebraic & Geometric Topology **15** (2015), no. 2, 691–731.

[24] Nikolay M. Nikolov and Svetoslav Zahariev, *Curved A-infinity algebras and Chern classes*, Pure and Applied Mathematics Quarterly **9** (2013), no. 2, 333–369.

[25] Paul Seidel, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zrich, 2008.

[26] Yin Tian, *A categorification of $U_q sl(1,1)$ as an algebra*, (2012), arXiv: 1210.5680.

[27] ———, *A categorification of $U_T sl(1,1)$ and its tensor product representations*, Geometry & Topology **18** (2014), no. 3, 1769–1863.

[28] Rumen Zarev, *Bordered Floer homology for sutured manifolds*, (2009), arXiv: 0908.1106.

[29] ———, *Joining and gluing sutured Floer homology*, (2010), arXiv: 1010.3496.