

AG
T

*Algebraic & Geometric
Topology*

Volume 25 (2025)

**Holomorphic polygons and the bordered Heegaard Floer homology
of link complements**

THOMAS HOCKENHULL

Holomorphic polygons and the bordered Heegaard Floer homology of link complements

THOMAS HOCKENHULL

We describe the construction of an \mathcal{A}_∞ multimodule in terms of counts of holomorphic polygons in a series of Heegaard multidiagrams. We show that this is quasi-isomorphic to the type-A bordered-sutured invariant of a link complement with a view to calculating, in the sequel, these invariants in terms of the link Floer homology of the corresponding link.

57M27; 57M25

1 Introduction

Suppose that $L \subset S^3$ is a link of h components, and Λ is a choice of framing for L . The purpose of this paper is to introduce an \mathcal{A}_∞ multimodule associated to (L, Λ) over \mathcal{T}^h , where \mathcal{T} is the algebra associated to a parameterised torus, from Lipshitz, Ozsváth and Thurston [14]: we call this $\text{Poly}(L, \Lambda)$. The \mathcal{A}_∞ quasi-isomorphism type of $\text{Poly}(L, \Lambda)$ is an invariant of the pair (L, Λ) , or, indeed, the three-manifold with parameterised boundary $(S^3 - L, \Lambda)$ — and is closely related to the bordered-sutured Heegaard Floer invariants of this three-manifold.

$\text{Poly}(L, \Lambda)$ is defined in terms of *holomorphic polygons*, and is closely related to the invariants $\widehat{\text{CF}}$ for closed three-manifolds, as defined by Ozsváth and Szabó [18]. At its root, the chain complex $\widehat{\text{CF}}(Y)$ is defined as a chain complex arising from counting holomorphic discs in a symplectic manifold associated to a *Heegaard diagram* corresponding to Y ; that is, $\widehat{\text{CF}}(Y)$ is $\widehat{\text{CF}}(\mathcal{D})$ for some Heegaard diagram \mathcal{D} representing Y . More generally, one may consider a *Heegaard n -diagram* for any $n \geq 2$, and count corresponding n -gons in an associated symplectic manifold. A Heegaard n -diagram corresponds with a sequence $Y^{0,1}, Y^{1,2}, \dots, Y^{n-2,n-1}, Y^{n-1,0}$ of three-manifolds, and counting holomorphic n -gons yields maps

$$m_n: \widehat{\text{CF}}(Y^{0,1}) \otimes \dots \otimes \widehat{\text{CF}}(Y^{n-2,n-1}) \rightarrow \widehat{\text{CF}}(Y^{n-1,0}).$$

These are well known to satisfy an \mathcal{A}_∞ relation, which is shown by counting the number of ends of one-dimensional moduli spaces and showing that they are of even parity. These moduli spaces decompose by *homology classes*, and, in fact, the parity of the ends of a one-dimensional moduli space remains even when restricted to curves of a given homology class. We exploit this extra information to define $\text{Poly}(L, \Lambda)$.

Broadly speaking, this multimodule is defined by associating a series of h sequences of elements of the torus algebra \mathcal{T} to each fixed homology class, in such a way that the \mathcal{A}_∞ relations for each homology class are compatible with the algebra — yielding an \mathcal{A}_∞ multimodule over the algebra \mathcal{T}^h .

The main theorem we show is:

Theorem 1.1 *The invariant $\text{Poly}(L, \Lambda)$ is quasi-isomorphic to the bordered-sutured multimodule $\text{BSA}(S^3 - L, \Lambda)$.*

The drawback of the invariants BSA is that they are often rather tricky to compute (see eg Hanselman [6], Levine [9], and Lipshitz, Ozsváth and Thurston [11; 14]). The module $\text{Poly}(L, \Lambda)$, whilst somewhat involved in its definition, is built from objects which have in one guise or other been present from the advent of Heegaard Floer homology for closed three-manifolds. There is a close relationship between the types of polygon maps used to build $\text{Poly}(L, \Lambda)$ and the invariants $\text{CFL}^-(L)$ — see, for instance, Manolescu and Ozsváth [15], Ozsváth and Szabó [17; 21; 22], and Rasmussen [23], or see Section 1.1 below.

Conceptually, our approach is perhaps advantageous to that of [14] in that — besides generalising directly to the case of general link complements — its connections with preexisting gluing type results in closed Heegaard Floer homology are clear; see the discussion in the next section.

1.1 Motivation

The invariant $\text{Poly}(L, \Lambda)$ is rather daunting at first, but has many precursors in Heegaard Floer homology that make it a natural object to study. To orient the reader, we touch upon some of these here — but leave the explicit details to future work.

1.1.1 The mapping cone theorem By the pairing theorems of Lipshitz, Ozsváth and Thurston [14] and Zarev [26], for some choice of framing $\Lambda = (\lambda_1, \dots, \lambda_{|L|})$ the invariant $\text{CFA}(E(L), \Lambda)$ determines the invariant $\widehat{\text{CF}}(S_\Lambda^3, (L)) \otimes \mathbb{F}_2^{\otimes |L|-1}$ for any choice of framing $\Lambda' = (\lambda'_1, \dots, \lambda'_{|L|})$. More specifically, beginning with the invariant $\text{CFA}(E(L), \Lambda)$, one takes a series of box tensor products with type D invariants of appropriately framed solid tori, such that the sutured three-manifold formed by gluing these solid tori to $E(L)$ according to the framing Λ and the framings on the solid tori is the sutured manifold S_Λ^3 , less $|L| - 1$ spheres with standard meridional sutures.

Historically, the invariants $\widehat{\text{CF}}(S_\Lambda^3, \cdot)$ have been computed using a sort of proto-gluing theorem, the *mapping cone theorem* — see [21] for the case of $L = K$ a knot, and [15] for the more general case where L is a general link. This takes the input data of the associated invariants $\text{CFL}^-(L)$, together with some “flip” maps (see [15] for more detail of the precise data required), and writes down the invariants $\widehat{\text{CF}}(S_\Lambda^3, \cdot)$ in terms of a complex built as a mapping cone of associated complexes.

We thus have two ways to compute the same quantity. The optimistic question is: to what extent do the quantities $\text{CFA}(E(L), \Lambda)$ and $\text{CFL}^-(L)$ (together with the “flip” maps) determine one another? More heuristically, does the recipe for $\widehat{\text{CF}}(S_{\Lambda}^3)$ in terms of $\text{CFL}^-(L)$ factor through $\text{CFA}(E(L), \Lambda)$?

The case where $L = K$ is a knot is the most well understood, and, indeed, it turns out not so optimistic:

Theorem 1.2 [14, Chapter 11] *The invariant $\text{CFK}^-(K)$ together with the extra information of the “flip” map determines the invariant $\text{CFD}(E(K), n)$ for any framing n .*

By taking a box tensor product with the invariant $\text{CFAA}(\mathbb{I})$, this gives the desired relationship.

The results of this paper give a concrete reason to expect this to generalise to links. In particular, the proof of the mapping cone theorem hinges upon a detailed understanding of the relationship between $\text{CFL}^-(L)$ and cobordism maps between three-manifolds $S_{\Lambda_0}^3$ and $S_{\Lambda_1}^3$, for $\Lambda_0 \subset \Lambda_1 \subset \Lambda$. These cobordism maps are defined in terms of the holomorphic polygons used in the definition of the invariant $\text{Poly}(L, \Lambda)$ in this paper; thus we have a first step towards associating $\text{CFL}^-(L)$ and $\text{Poly}(L, \Lambda) \cong \text{CFA}(E(L), \Lambda)$.

The invariant $\text{Poly}(L, \Lambda)$, however, tracks extra information about these holomorphic polygons, in the form of an associated set of sequences of algebra elements, which are unseen by the corresponding cobordism maps and higher polygon counting maps in the classic set-up. It transpires, however, that much of this extra information can be packaged in terms of gradings on the chain complexes involved. More about this, and the relationship between the invariants $\text{CFL}^-(L)$, “flip maps”, and bordered Heegaard Floer homology is to be made explicit in [7].

1.1.2 The spectral sequence of the branched double-cover, and an alternative proof In [19, Section 4.2], a chain complex is associated to a framed h -component link. As a set, this complex is given by $X = \bigoplus_{I \in \{m, l\}^h} \widehat{\text{CF}}(Y(I))$, where $Y(I)$ is the manifold formed by surgering Y along L according to the framing induced by Λ and I . The differential D counts a sum of polygon maps that respect a fixed ordering of the vertices of I^h .

The invariant $\text{Poly}(L, \Lambda)$ is defined in a very similar manner, but with two main differences. First, polygon maps in $\text{Poly}(L, \Lambda)$ do not obey the ordering of the vertices: instead, we must consider all possible polygon maps —ie the maps in $\text{Poly}(L, \Lambda)$ split as sets of maps, where each set respects a fixed ordering of the vertices of I^h . The other clear difference is that $\text{Poly}(L, \Lambda)$ is an \mathcal{A}_∞ -module, whereas $X(L, \Lambda)$ is a chain complex. This can be reconciled by considering instead the differential graded module $(P, \partial) := \text{Poly}(L, \Lambda) \otimes \mathcal{T}^*$ (here, we write \mathcal{T}^* to denote the tensor algebra associated to the torus algebra; see [14, page 10]). Both P and $X(L, \Lambda)$ split as a direct sum over idempotents of the torus algebra, and the provision of an extra basepoint for each component of L in the corresponding Heegaard multidiagram gives a filtration upon the differential of $X(L, \Lambda)$, which corresponds with the splitting of ∂ according to corresponding sequences of algebra elements. When L is a knot, this should be compared to Eftekhary [5].

More recently, the complex $X(L, \Lambda)$ has been studied from the perspective of bordered Heegaard Floer homology, to understand the difference between two spectral sequences between the reduced Khovanov homology of the mirror of a link L' and the Heegaard Floer homology $\widehat{\text{HF}}(\Sigma(L'))$, where $\Sigma(L')$ is the branched double cover of L' : the one given by Ozsváth and Szabó [19], and another defined in terms of bordered Heegaard Floer homology by Lipshitz, Ozsváth and Thurston [13]. This relationship follows naturally from Theorem 1.1, but, perhaps more interestingly, the methods of [13] can be applied, together with a model computation, to give a different proof of Theorem 1.1 to the one in the present paper.

We sketch the main idea. The main result of [13] is a pairing theorem for holomorphic polygons in *bordered Heegaard multidiagrams* (see Section 4.6 below). Approximately, this states that if \mathcal{D}_1 and \mathcal{D}_2 are bordered Heegaard k - and l -diagrams respectively (ie \mathcal{D}_1 has β^1, \dots, β^k sets of attaching circles and one set α of attaching arcs and circles, and similar for \mathcal{D}_2), a pair of holomorphic polygons u_1 and u_2 with appropriate boundary conditions (encoded by elements of the torus algebra) corresponds to a holomorphic polygon in a diagram \mathcal{D} . Here, \mathcal{D} corresponds roughly with the diagram formed by gluing \mathcal{D}_1 and \mathcal{D}_2 together along a common boundary; but in order to make sense of attaching circles, it is often necessary to add a series of perturbed “copies” of some β^i to \mathcal{D}_1 or \mathcal{D}_2 , to promote \mathcal{D}_1 to a bordered Heegaard l -diagram (if $k < l$, say).

Consider a cube of framings with vertices given by elements $I \in \{m, l\}^2$. To this we may associate a type D structure given by the sum of the invariants $\text{CFD}(\mathcal{D}(I))$, where each $\mathcal{D}(I)$ is a bordered Heegaard multidiagram for the torus, with attaching circles specified by I . We may promote this type D structure to a type DA -bimodule by assigning to each domain B for $\mathcal{D}(I)$ a sequence of elements of the torus algebra $\vec{\rho}(B)$, and defining $m_{1,k}(x, \vec{\rho})$ to be a map counting holomorphic curves in homology class B with $\vec{\rho}(B) = \vec{\rho}$, such that summing these $m_{1,k}(x, \vec{\rho})$ over all sequences of $\vec{\rho}$ gives the structure map δ^1 for the invariant $\text{CFD}(\mathcal{D}(I))$. By direct calculation, one can see that this DA -bimodule is equivalent to the identity DA -bimodule for the torus, $\text{CFDA}(\mathbb{I})$. As such,

$$\text{CFA}(E(L)) \cong \text{CFA}(E(L)) \boxtimes \text{CFDA}(\mathbb{I}),$$

but, by the pairing theorem and the above discussion, we know that this is equal to the sum

$$\bigoplus_I \text{CFA}(E(L)) \boxtimes \text{CFD}(\mathcal{D}(I)),$$

which is a sum of polygon maps in a bordered Heegaard multidiagram with $|L| - 1$ boundary components, splitting by the algebra element associated to the restriction of their domain to the portion $\mathcal{D}(I)$. Iterating this for all boundary components of $E(L)$ gives an identification of $\text{CFA}(E(L))$ with a module defined in terms of polygon maps in closed Heegaard diagrams, with extra algebra information stemming from how the associated domain behaves in the regions corresponding to $\mathcal{D}(I)$ — this is precisely the way the object $\text{Poly}(L, \Lambda)$ is defined.

This should be compared with the results of Auroux [2, Section 6.2]; in particular, the sketched argument above is approximately equivalent to the sketch proof of Proposition 6.5 given in that paper, but recast in

a setting and language that is perhaps more accessible to the reader with background in Heegaard Floer homology, rather than more broad symplectic geometry.¹

1.2 Organisation

In [Section 2](#) we recall the algebraic definitions of \mathcal{A}_∞ multimodules necessary for the rest of the paper, and some relevant constructions and properties of the torus algebra \mathcal{T} . Then, in [Section 3](#), we consider the close relationship between bordered-sutured Heegaard diagrams which represent the complement of a link, and a corresponding family of “partially splayed” and “splayed” Heegaard diagrams. The invariant $\text{Poly}(L, \Lambda)$ is defined in terms of the splayed diagrams, which are a family of closed Heegaard multidiagrams—but we will need to consider a series of auxiliary modules $\text{Poly}_k(L, \Lambda)$ defined in terms of partially splayed diagrams, which are bordered-sutured k -multidiagrams. These correspond roughly with the final k algebra inputs into a structure map m_n being contributed not by asymptotics at a puncture in the Heegaard diagram, but instead by algebra elements associated to the homology class of a holomorphic curve. In particular, when k is zero, all algebra elements are contributed by asymptotics at a puncture, and so these recover the bordered sutured Floer homology of the link complement; when k is large enough, these modules stabilise and there is no interesting behaviour near the puncture—so these modules may be identified with $\text{Poly}(L, \Lambda)$.

In [Section 4](#), we define a series of moduli spaces of holomorphic polygons which will be instrumental in our definition of the invariants $\text{Poly}(L, \Lambda)$, and the auxiliary modules $\text{Poly}_k(L, \Lambda)$. We carefully analyse the possible boundary of these moduli spaces, thus ensuring that the \mathcal{A}_∞ relations hold for these multimodules. We also lay the foundations of a neck-stretching argument which will allow us, in the final section, to identify all of the modules $\text{Poly}_k(L, \Lambda)$ for each k , thus proving [Theorem 1.1](#).

Acknowledgements

I owe the Heegaard Floer community at large a debt of gratitude for their support and encouragement. In particular, I would like to thank Marco Marengon for many useful and interesting conversations. I am also grateful to Peter Ozsváth for hosting me at Princeton, and the helpful conversations that took place there. I am also grateful to Tom Coates and András Stipsicz for their comments and corrections on the version of this paper which constituted my PhD thesis, along with the referee for such a careful reading and their corrections and comments.

Robert Lipshitz has gone far and beyond in his encouragement and support, and in his tirelessness in answering my questions about the “LOT” oeuvre—he and the University of Oregon were also very gracious hosts to me in the spring of 2016. Richard Thomas has been invaluable in his support, encouragement and cajoling over the past year or so.

¹For full transparency, I learned of Auroux’s results roughly two-thirds of the way through writing this paper.

Finally, I have been constantly humbled by the generosity and kindness of my supervisor Jacob Rasmussen. Without his reassurance, encouragement and patience, I doubt any of this would have been possible.

The author was supported by an EPSRC studentship from 2013–2017, and by Imperial College London mathematics department from 2017–2018. Some of this work was undertaken whilst visiting Princeton University, which was funded by the Doris Chen award.

2 Algebraic background

We spend this section recalling the algebraic background of the bordered sutured invariants — although these are defined in [26, Section 8], the algebraic structure is not discussed at length.

2.1 The torus algebra

We first define the algebras that we will work over for most of this paper. These are the *torus algebra*, \mathcal{T} , and the products

$$\mathcal{T}^k := \underbrace{\mathcal{T} \times \cdots \times \mathcal{T}}_{k \text{ times}}.$$

The torus algebra is generated by eight elements as an \mathbb{F}_2 -module, which are the two idempotents ι_m and ι_l , and the six Reeb chords $\rho_1, \rho_2, \rho_3, \rho_{12}, \rho_{23}$ and ρ_{123} . The differential is trivial, and the idempotents satisfy $\iota_m \iota_l = \iota_l \iota_m = 0$ — we also write 1 for the element $\iota_m + \iota_l$. The Reeb chords satisfy $\rho_1 \rho_2 = \rho_{12}, \rho_2 \rho_3 = \rho_{23}$ and $\rho_1 \rho_{23} = \rho_{12} \rho_3 = \rho_{123}$. The other nonzero products are

$$\begin{aligned} \iota_m \rho_1 &= \rho_1 \iota_l = \rho_1, & \iota_l \rho_2 &= \rho_2 \iota_m = \rho_2, & \iota_m \rho_3 &= \rho_3 \iota_l = \rho_3, \\ \iota_m \rho_{12} &= \rho_{12} \iota_m = \rho_{12}, & \iota_l \rho_{23} &= \rho_{23} \iota_l = \rho_{23}, & \iota_m \rho_{123} &= \rho_{123} \iota_l = \rho_{123}. \end{aligned}$$

These algebra elements may be interpreted in terms of *Reeb chords*. Let Z denote an oriented circle with four marked points $\mathbf{a} = \{a_1, \dots, a_4\} \subset Z$, together with a basepoint $b \in Z \setminus \mathbf{a}$ — we suppose that \mathbf{a} is numbered in such a way that a positive path from b encounters them in order. A *Reeb chord* $\rho \in (Z \setminus b, \mathbf{a})$ is an immersed arc in $Z \setminus b$ with endpoints in \mathbf{a} , and orientation agreeing with that on Z . To each Reeb chord ρ we associate the algebra element $a(\rho) := \rho_{[i,j]}$, where $\partial\rho = a_j - a_i$ and $[i, j]$ is the unique substring of 123 beginning with i and ending with $j - 1$. We will often conflate a Reeb chord with its corresponding algebra element.

We say that two Reeb chords ρ and ρ' are *composable* if the product $\rho\rho'$, which we call their *composition*, is nonzero. More generally, we call sequence (ρ^1, \dots, ρ^k) of Reeb chords *composable* if for every $i = 1, \dots, k - 1$, the chords ρ^i and ρ^{i+1} are composable. If $\vec{\rho}$ is a sequence of Reeb chords, then there is an associated sequence of Reeb chords $\vec{\rho}'$ formed by replacing any composable subsequences of Reeb chords with their composition, which we call the *reduction* of $\vec{\rho}$. There is an equivalence relation upon sequences of Reeb chords, given by saying that $\vec{\rho}$ and $\vec{\rho}'$ are *composable-equivalent* if and only if their reductions agree.

In a similar manner, we say that two chords, ρ and ρ' , are 123-composable if they are composable and their composition is ρ_{123} , and make a similar definition for the 123-reduction of a sequence of Reeb chords. We shall say that a pair of sequences of Reeb chords are 123-equivalent if and only if their 123-reductions agree.

Remark 2.1 Being 123-composable depends only upon pairs of adjacent chords. In particular, the sequence (ρ_1, ρ_2, ρ_3) is 123-reduced, whereas the sequence (ρ_1, ρ_{23}) has 123-reduction ρ_{123} .

One can see from the definition of the torus algebra that some Reeb chords behave rather differently to others with respect to products by the elementary idempotents. If $\iota_\delta \cdot \rho \neq 0$ we will say that $\iota_{\text{start}}(\rho) = \iota_\delta$, and if $\rho \cdot \iota_\delta \neq 0$ we will say that $\iota_{\text{end}}(\rho) = \iota_\delta$.

It is straightforward to see that for a sequence of Reeb chords $\vec{\rho} = (\rho^1, \dots, \rho^m)$, the tensor product $\rho^1 \otimes \dots \otimes \rho^m$ over the ring of idempotents of \mathcal{T} is nonzero if and only if $\iota_{\text{end}}(\rho^i) = \iota_{\text{start}}(\rho^{i+1})$. Accordingly, we make the convention throughout this paper that a sequence of Reeb chords satisfies this condition — ie that it is shorthand for an element of the tensor product $\mathcal{T}^{\otimes m}$.

Definition 2.2 Suppose that $\vec{\rho} = (\rho^1, \dots, \rho^m)$ is a sequence of elements of the torus algebra. To the sequence $\vec{\rho}$ we associate the sequence of its jumping chords. That is, the subsequence

$$\text{jump}(\vec{\rho}) := (\rho^i \in \vec{\rho} \mid \iota_{\text{start}}(\rho^i) \neq \iota_{\text{end}}(\rho^i))$$

of all chords in the sequence that have different starting to ending idempotent.

Example 2.3 For instance, the sequence $\vec{\rho} = (\rho_2, \rho_{12}, \dots, \rho_{12}, \rho_1)$ has $\text{jump}(\vec{\rho}) = (\rho_2, \rho_1)$.

We turn our attention to what we term a splicing of a series of Reeb chords. These should be thought of as something between a set of elements $\vec{b}^1, \dots, \vec{b}^k$ where $\vec{b}^i \in \mathcal{T}^{\otimes h}$, and a set of sequences of elements $\vec{a}_1, \dots, \vec{a}_h$ where each $\vec{a}_j \in \mathcal{T}^{k_j}$. The point is that, later, we will make constructions (of diagrams) which correspond with a given set of sequences of Reeb chords: these depend upon the corresponding set of sequences of jumping chords (and ignore any nonjumping chords), and manifestly depend upon some overall ordering of the jumping chords from all of the sequences — a splicing is a way to specify this overall ordering.

Definition 2.4 Suppose that $\vec{\rho}_1, \dots, \vec{\rho}_h$ are a series of sequences of Reeb chords $\vec{\rho}_j := (\rho_j^1, \dots, \rho_j^{i_j})$ with corresponding jumping sequences $\text{jump}(\vec{\rho}_1), \dots, \text{jump}(\vec{\rho}_j)$.

An m -splicing of $\vec{\rho}_1, \dots, \vec{\rho}_h$ is an ordered length m partition $P_j := (P_j^1, \dots, P_j^m)$ of each $\text{jump}(\vec{\rho}_j)$, allowing for parts P_j^i to be empty, and such that

- $|P_j^i| \leq 1$ for every i and j , and
- for each i , the union $P_1^i \cup \dots \cup P_h^i$ is nonempty.

Every splicing $\vec{\sigma}$ has an underlying set of unordered partitions, which we shall often denote by σ . We often write $|\vec{\sigma}| = m$ to indicate that $\vec{\sigma}$ is an m -splicing. We also write $\text{Col}(\vec{\sigma})$ for the number of elements in the set

$$\{i \mid P_j^i \text{ is nonempty for more than one } j\}.$$

We shall say that $\vec{\sigma}$ is an *interleaving* of $\vec{\rho}_1, \dots, \vec{\rho}_h$, or that $\vec{\sigma}$ is *interleaved* if $\text{Col}(\vec{\sigma}) = 0$; that is, for every i , precisely one P_j^i is nonempty.

Example 2.5 For instance, consider $\vec{\rho}_1 = (\rho_1, \rho_{23}, \rho_2)$ and $\vec{\rho}_2 = (\rho_{12}, \rho_1)$. A noninterleaved 2-splicing of these is given by $P_1 = (\rho_1, \rho_2)$ and $P_2 = (\rho_1, \emptyset)$; an interleaved 3-splicing $\vec{\sigma}_1$ is given by $P_1 = (\rho_1, \emptyset, \rho_2)$ and $P_2 = (\emptyset, \rho_1, \emptyset)$.

If we are given two sets of sequences of Reeb chords $\vec{\lambda} = \vec{\lambda}_1, \dots, \vec{\lambda}_h$ and $\vec{\delta} = \vec{\delta}_1, \dots, \vec{\delta}_h$, where each $\vec{\lambda}_j$ is of length l_j and each $\vec{\delta}_j$ is of length k_j , there is a corresponding set of sequences of Reeb chords $\vec{\lambda} \star \vec{\delta}$ defined by $(\vec{\lambda} \star \vec{\delta})_j := (\vec{\lambda}_j, \vec{\delta}_j)$. Similarly, given splittings $\vec{\sigma}_1 = P_1, \dots, P_h$ and $\vec{\sigma}_2 = Q_1 \cdots Q_h$ for $\vec{\rho}$ and $\vec{\delta}$ respectively, then there is a corresponding splicing $\vec{\sigma}_1 \star \vec{\sigma}_2$ of $\vec{\rho} \star \vec{\delta}$, formed by $(\vec{\sigma}_1 \star \vec{\sigma}_2) = (P_1, Q_1), \dots, (P_h, Q_h)$.

Example 2.6 With $\vec{\rho} = \vec{\rho}_1, \vec{\rho}_2$ as in [Example 2.5](#), and $\vec{\lambda} = \vec{\lambda}_1, \vec{\lambda}_2$ for $\vec{\lambda}_1 = (\rho_{23})$ and $\vec{\lambda}_2 = (\rho_2)$, with corresponding splicing $\vec{\sigma}_2$ given by $Q_1 = (\emptyset)$ and $Q_2 = (\rho_2)$, we have that $(\vec{\rho} \star \vec{\lambda})_1 = (\rho_1, \rho_{23}, \rho_2, \rho_{23})$ and $(\vec{\rho} \star \vec{\lambda})_2 = (\rho_{12}, \rho_1, \rho_2)$. The corresponding splicing is $(\vec{\sigma}_1 \star \vec{\sigma}_2)$, given by the partitions $(\rho_1, \emptyset, \rho_2, \emptyset)$ and $(\emptyset, \rho_1, \emptyset, \rho_2)$.

A splicing naturally gives rise to a sequence of sets of idempotents in the torus algebra:

Definition 2.7 Let $\vec{\rho} := \vec{\rho}_1, \dots, \vec{\rho}_h$ be a series of sequences of algebra elements $\vec{\rho}_j$, and let $\vec{\sigma} = P_1, \dots, P_h$ be an m -splicing for $\vec{\rho}$.

For each j , let $P_j^{i_1}, \dots, P_j^{i_{|\text{jump}(\vec{\rho}_j)|}}$ denote the parts of P_j which are nonempty, and let $\rho_j^{i_k}$ denote the (unique) Reeb chord in each $P_j^{i_k}$. Each $P_j^{i_k}$ has an associated sequence of parts

$$\vec{P}_j^{i_k} := (P_j^{i_k}, P_j^{i_{k+1}}, \dots, P_j^{i_{|\text{jump}(\vec{\rho}_j)|-1}}),$$

its *trailed* sequence, where only $P_j^{i_k} \in \vec{P}_j^{i_k}$ is nonempty. The sequence P_j decomposes as a union

$$P_j = \bigcup_{k=0}^{|\text{jump}(\vec{\rho}_j)|} \vec{P}_j^{i_k},$$

where we define $\vec{P}_j^{i_0}$ to be the union of the empty partitions $(P_j^1, \dots, P_j^{i_1-1})$.

We define a set of idempotents by

$$t_j^i := \begin{cases} t_{\text{start}}(\rho_j^{i_1}) & \text{if } P_j^i \in \vec{P}_j^{i_0} \text{ or } i = 0, \\ t_{\text{end}}(\rho_j^{i_k}) & \text{if } P_j^i \in \vec{P}_j^{i_k}, k \neq 0. \end{cases}$$

We collect these into sequences of idempotents $\vec{\tau}^0, \dots, \vec{\tau}^m$ by $\vec{\tau}^i = \tau_1^i, \dots, \tau_h^i$, and denote by $\iota_{\vec{\sigma}} := (\vec{\tau}^0, \dots, \vec{\tau}^m)$.

Example 2.8 With $\vec{\rho} = \vec{\rho}_1, \vec{\rho}_2$ and $\vec{\sigma}_1$ as in Example 2.5, we have corresponding sequences of idempotents $\vec{\tau}^0 = (\iota_m, \iota_m), \vec{\tau}^1 = (\iota_l, \iota_m), \vec{\tau}^2 = (\iota_l, \iota_l)$ and $\vec{\tau}^3 = (\iota_m, \iota_l)$.

We can also go in the other direction, associating a splicing to a pair of a set of sequences of Reeb chords and a set of sequences of idempotents:

Definition 2.9 Suppose that we are provided with a set of sequences of Reeb chords $\vec{\rho}_1, \dots, \vec{\rho}_h$, and a sequence of sets of idempotents $\vec{\tau} = (\vec{\tau}^0, \dots, \vec{\tau}^m)$ where each $\vec{\tau}^i = \tau_1^i, \dots, \tau_h^i$ satisfies that for $j = 1, \dots, h$, the sequence of idempotents $\vec{\tau}_j := \tau_j^0, \dots, \tau_j^m$ has

$$|\{\tau_j^i \in \vec{\tau}_j \mid \tau_j^i \neq \tau_j^{i-1}\}| = |\text{jump}(\vec{\rho}_j)|,$$

and that for every $i = 1, \dots, m$ there is some j such that $\tau_j^i \neq \tau_j^{i-1}$.

There is a corresponding splicing given by setting $P_j^i = \emptyset$ when $\tau_j^i = \tau_j^{i-1}$, and for each $\tau_j^{i_k} \neq \tau_j^{i_k-1}$ we set P_j^i to be the part $\{\rho_j^{i_k}\}$, where $\rho_j^{i_k}$ is the k^{th} chord in $\text{jump}(\vec{\rho}_j)$ — we call this $\vec{\sigma}(\vec{\rho}, \vec{\tau})$.

Example 2.10 Suppose $\vec{\rho} = \vec{\rho}_1, \vec{\rho}_2$ are as in Example 2.5, and $\vec{\tau} = (\vec{\tau}^0, \dots, \vec{\tau}^m)$ as in Example 2.8. The corresponding splicing is the interleaved splicing given in Example 2.5.

We will also need to consider splitting a sequence of Reeb chords in two — roughly, in what follows, one half $\vec{\rho}_{\Delta}$ corresponds with genuine asymptotic behaviour of a holomorphic curve, whereas $\vec{\rho}_{\pm}$ will be extracted from the topology of the homology class of the curve.

Definition 2.11 Suppose that $\vec{\rho} = \vec{\rho}_1, \dots, \vec{\rho}_h$ is a set of sequences of Reeb chords, and $\vec{\sigma}$ is an interleaved splicing of $\vec{\rho}$. We shall say that a k -shipping of $(\vec{\rho}, \vec{\sigma})$ is a pair $\vec{\rho}_{\Delta}, \vec{\rho}_{\pm}$ (where $\vec{\rho}_{\Delta} = \vec{\rho}_{\Delta,1}, \dots, \vec{\rho}_{\Delta,h}$ and $\vec{\rho}_{\pm} = \vec{\rho}_{\pm,1}, \dots, \vec{\rho}_{\pm,h}$) together with a pair of splittings $\vec{\sigma}_{\Delta}, \vec{\sigma}_{\pm}$ such that

- for some j , we have that $\vec{\rho}_{\pm,j}$ begins with a jumping chord;
- $\vec{\sigma}_{\pm}$ is a k' -splicing of $\vec{\rho}_{\pm}$, where

$$k' = \min\{|\text{jump}(\vec{\rho}_1)| + \dots + |\text{jump}(\vec{\rho}_h)|, k\};$$

- $(\vec{\rho}_{\Delta}, \vec{\sigma}_{\Delta}) \star (\vec{\rho}_{\pm}, \vec{\sigma}_{\pm}) = (\vec{\rho}, \vec{\sigma})$.

If $(\vec{\rho}_{\Delta}, \vec{\sigma}_{\Delta})$ and $(\vec{\rho}_{\pm}, \vec{\sigma}_{\pm})$ form a k -shipping of $(\vec{\rho}, \vec{\sigma})$, and $(\vec{\rho}'_{\Delta}, \vec{\sigma}'_{\Delta})$ and $(\vec{\rho}'_{\pm}, \vec{\sigma}'_{\pm})$ form a $(k+1)$ -shipping of $(\vec{\rho}, \vec{\sigma})$ then we shall say that the pair $(\vec{\rho}'_{\Delta}, \vec{\sigma}'_{\Delta})$ and $(\vec{\rho}'_{\pm}, \vec{\sigma}'_{\pm})$ *splays* $(\vec{\rho}_{\Delta}, \vec{\sigma}_{\Delta})$ and $(\vec{\rho}_{\pm}, \vec{\sigma}_{\pm})$.

Example 2.12 Suppose $\vec{\rho} = \vec{\rho}_1, \vec{\rho}_2$ and $\vec{\sigma}_1$ are as in Example 2.5.

Let

$$\begin{aligned} \vec{\rho}_{\Delta,1} &= (\rho_1, \rho_{23}), & \vec{\rho}_{\pm,1} &= (\rho_2), \\ \vec{\rho}_{\Delta,2} &= (\rho_{12}, \rho_1), & \vec{\rho}_{\pm,2} &= \emptyset, \end{aligned}$$

so that we have

$$\vec{\rho}_{\Delta,1} \star \vec{\rho}_{\ddagger,1} = (\rho_1, \rho_{23}, \rho_2) = \vec{\rho}_1, \quad \vec{\rho}_{\Delta,2} \star \vec{\rho}_{\ddagger,2} = \vec{\rho}_2.$$

Letting

$$\begin{aligned} \vec{\sigma}_{\Delta,1} &= (\rho_1, \emptyset), & \vec{\sigma}_{\ddagger,1} &= (\rho_2), \\ \vec{\sigma}_{\Delta,2} &= (\emptyset, \rho_1), & \vec{\sigma}_{\ddagger,2} &= (\emptyset), \end{aligned}$$

we see that

$$\vec{\sigma}_{\Delta,1} \star \vec{\sigma}_{\ddagger,1} = (\rho_1, \emptyset, \rho_2) = P_1, \quad \vec{\sigma}_{\Delta,2} \star \vec{\sigma}_{\ddagger,2} = (\emptyset, \rho_1, \emptyset) = P_2.$$

As $\vec{\sigma}_{\ddagger} = (\sigma_{\ddagger,1}, \sigma_{\ddagger,2})$ is a 1-splicing of $(\vec{\rho}_{\ddagger,1}, \vec{\rho}_{\ddagger,2})$ it follows that $((\vec{\rho}_{\Delta,1}, \vec{\rho}_{\Delta,2}), (\vec{\sigma}_{\ddagger,1}, \vec{\sigma}_{\ddagger,2}))$ is a 1-shipping of $(\vec{\rho}, \vec{\sigma}_1)$.

Similarly, a 2-shipping of $(\vec{\rho}, \vec{\sigma}_1)$ is given by

$$\begin{aligned} \vec{\rho}_{\Delta,1} &= (\rho_1, \rho_{23}), & \vec{\rho}_{\ddagger,1} &= (\rho_2), \\ \vec{\rho}_{\Delta,2} &= (\rho_{12}), & \vec{\rho}_{\ddagger,2} &= (\rho_1), \end{aligned}$$

with corresponding splittings

$$\begin{aligned} \vec{\sigma}_{\Delta,1} &= (\rho_1), & \vec{\sigma}_{\ddagger,1} &= (\emptyset, \rho_2), \\ \vec{\sigma}_{\Delta,2} &= (\emptyset), & \vec{\sigma}_{\ddagger,2} &= (\rho_1, \emptyset). \end{aligned}$$

Another 2-shipping is given by

$$\begin{aligned} \vec{\rho}_{\Delta,1} &= (\rho_1), & \vec{\rho}_{\ddagger,1} &= (\rho_{23}, \rho_2), \\ \vec{\rho}_{\Delta,2} &= (\rho_{12}), & \vec{\rho}_{\ddagger,2} &= (\rho_1), \end{aligned}$$

with corresponding splittings

$$\begin{aligned} \vec{\sigma}_{\Delta,1} &= (\rho_1), & \vec{\sigma}_{\ddagger,1} &= (\emptyset, \rho_2), \\ \vec{\sigma}_{\Delta,2} &= (\emptyset), & \vec{\sigma}_{\ddagger,2} &= (\rho_1, \emptyset). \end{aligned}$$

2.2 \mathcal{A}_∞ multimodules

We begin with \mathcal{A}_∞ -modules as a warm-up. Let \mathcal{A} be a differential graded algebra over \mathbb{F}_2 , with differential ∂ , multiplication μ and ring of idempotents $\mathcal{I} \subset \mathcal{A}$. A (right) \mathcal{A}_∞ module over \mathcal{A} is a module over \mathbb{F}_2 with right action of \mathcal{I} and maps

$$m_{k+1} : M \otimes_{\mathcal{I}} \mathcal{A} \otimes_{\mathcal{I}} \otimes \cdots \otimes_{\mathcal{I}} \mathcal{A} \rightarrow M,$$

satisfying the \mathcal{A}_∞ relation

$$\begin{aligned} 0 = \sum_{i=0}^n m_{n-i+1}(m_{i+1}(x, a_1, \dots, a_i)a_{i+1}, \dots, a_n) \\ + \sum_{i=1}^{n-1} m_n(x, a_1, \dots, \mu(a_i, a_{i+1}), a_{i+2}, \dots, a_n) + \sum_{i=1}^n m_{n+1}(x, a_1, \dots, \partial a_i, \dots, a_n) \end{aligned}$$

for every $x \in M$ and $a_1, \dots, a_n \in \mathcal{A}$.

More general is the notion of an \mathcal{A}_∞ multimodule. If $\mathcal{A}_1, \dots, \mathcal{A}_h$ are differential graded algebras over \mathbb{F}_2 , each with corresponding ring of idempotents $\mathcal{I}_1, \dots, \mathcal{I}_h$, then a (right) \mathcal{A}_∞ multimodule over $\mathcal{A}_1, \dots, \mathcal{A}_h$ is a module M over \mathbb{F}_2 with a right action of $\mathcal{I} := \mathcal{I}_1 \otimes \dots \otimes \mathcal{I}_h$, and equipped with maps

$$m_{i_1, \dots, i_h} : M \otimes_{\mathcal{I}} \mathcal{A}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{A}_h^{\otimes i_h} \rightarrow M,$$

where $\mathcal{A}_i^{\otimes i_j}$ is the result of taking the tensor product of \mathcal{A}_i with itself i_j times, over the ring \mathcal{I}_j .

The relations these maps must satisfy are more difficult to write down; we roughly follow Hanselman’s presentation in [6, Section 2.1], with some modifications. If $\vec{a}_j := (a_j^1, \dots, a_j^{i_j}) \in \mathcal{A}_j^{\otimes i_j}$ is a sequence of elements of \mathcal{A}_j , then for every $0 \leq k \leq i_j$ we define the truncations

$$T_k(\vec{a}_j) := (a_j^1, \dots, a_j^k), \quad T^k(\vec{a}_j) := (a_j^k, \dots, a_j^{i_j}),$$

and the maps

$$\begin{aligned} \bar{\mu}^k(\vec{a}_j) &:= (a_j^1, \dots, a_j^{k-1}, \mu(a_j^k, a_j^{k+1}), a_j^{k+2}, \dots, a_j^{i_j}), \\ \bar{\partial}(\vec{a}_j) &:= \sum_{k=1}^{i_j} (a_j^1, \dots, a_j^{k-1}, \partial a_j^k, a_j^{k+1}, \dots, a_j^{i_j}). \end{aligned}$$

This allows us to define the sequence

$$\bar{\mu}_j^k(\vec{a}_1, \dots, \vec{a}_h) := (\vec{a}_1, \dots, \bar{\mu}^k(\vec{a}_j), \dots, \vec{a}_h).$$

With this in mind, and omitting subscripts, the relations which the multiplication maps must satisfy is given by

$$\begin{aligned} 0 = \sum_{j_1, \dots, j_h} m(m(x, T_{j_1}(\vec{a}_1), \dots, T_{j_h}(\vec{a}_h)), T^{j_1}(\vec{a}_1), \dots, T^{j_h}(\vec{a}_h)) \\ + \sum_{j=1}^h \sum_{k=1}^{i_j-1} m(x, \bar{\mu}_j^k(\vec{a}_1, \dots, \vec{a}_h)) + \sum_{k=1}^h m(x, \vec{a}_1, \dots, a_{k-1}^{\vec{a}_k}, \bar{\partial}(\vec{a}_k), a_{k+1}^{\vec{a}_k}, \dots, \vec{a}_h). \end{aligned}$$

When $\mathcal{A}_1, \dots, \mathcal{A}_h$ are all copies of the torus algebra, we will have cause to consider decompositions of these maps according to the extra data of a splicing of $\vec{a}_1, \dots, \vec{a}_h$ — the easiest way to define the \mathcal{A}_∞ structure maps we will be interested in later in the paper is in terms of a series of auxiliary maps which sum to give the structure maps, and themselves satisfy a relation that implies the \mathcal{A}_∞ relations for their sums. First, we need to discuss operations upon interleaved splittings which are compatible with the operations T_* , T^* and $\bar{\mu}_j^i$.

Fix an interleaved splicing $\vec{\sigma}$ of a set of sequences of Reeb chords $\vec{\rho} = \vec{\rho}_1, \dots, \vec{\rho}_h$. We shall say a pair of indices (i, j) is *compatible* with $(\vec{\rho}, \vec{\sigma})$ if either

- (1) $\rho_j^i \rho_j^{i+1}$ is nonzero and precisely one of ρ_j^i or ρ_j^{i+1} is jumping (in part P_j^k , say), or
- (2) $\rho_j^i \rho_j^{i+1}$ is nonzero, both of $\rho_j^i \in P_j^{k_i}$ and $\rho_j^{i+1} \in P_j^{k_i+1}$ are jumping, and we have that they are in consecutive parts — that is, $\rho_j^i \in P_j^{k_i}$ and $\rho_j^{i+1} \in P_j^{k_i+1}$ for some k_i .

If (i, j) are compatible with $\vec{\rho}, \vec{\sigma}$, then there is an associated splicing of $\vec{\mu}_j^i(\vec{\rho})$ which we call $\vec{\sigma}_j^i$. In case (1), we set $\vec{\sigma}_j^i$ to be the splicing of $\vec{\mu}_j^i(\vec{\rho})$ obtained from $\vec{\sigma}$ by replacing P_j^k with the part $\{\rho_j^i \rho_j^{i+1}\}$. In case (2), we set $\vec{\sigma}_j^i$ to be the splicing given by omitting, for each l , the parts $P_l^{k_i}$ and $P_l^{k_i+1}$ from P_l .

If $\vec{\sigma}$ is interleaved, we shall say that an index $0 < k < m$ is *collidable* for $(\vec{\rho}, \vec{\sigma})$ if for every j we have that only one of P_j^k and P_j^{k+1} is nonempty. In this circumstance, we can define an associated $(m-1)$ -splicing of $\vec{\rho}$ associated to $\vec{\sigma}$ and k which we call $\vec{\sigma}(k)$, the *collision of $\vec{\sigma}$ at k* . $\vec{\sigma}(k)$ consists of the partitions Q_j , where each Q_j is formed by

$$Q_j^i := \begin{cases} P_j^i & \text{if } i < k, \\ P_j^i \cup P_j^{i+1} & \text{if } i = k, \\ P_j^{i+1} & \text{if } i > k + 1. \end{cases}$$

In words, we amalgamate the k and $(k+1)^{\text{st}}$ parts of each partition. Colliding does not preserve being interleaved.

Example 2.13 Suppose that $\vec{\rho} = \vec{\rho}_1 = (\rho_2, \rho_{12}, \rho_3, \rho_2, \rho_3)$, with splicing $\vec{\sigma} = (\rho_2, \rho_3, \rho_2, \rho_3)$. The indices $(2, 1)$ are compatible, in case (1), as $\rho_{12} \cdot \rho_3 = \rho_{123} \neq 0$ and ρ_3 is the only jumping chord of the two—here, the associated splicing of $\vec{\mu}_1^2(\vec{\rho})$ is given by $(\rho_2, \rho_{123}, \rho_2, \rho_3)$. The indices $(4, 1)$ are compatible, in case (2), as (trivially) the final ρ_2 and ρ_3 belong to consecutive parts of $\vec{\sigma}$ —the associated splicing of $\vec{\mu}_1^4(\vec{\rho})$ is given by (ρ_2, ρ_3) .

Suppose that $\vec{\rho}_2 = (\rho_{12}, \rho_1, \rho_{23}, \rho_2)$, and let

$$\vec{\sigma}_2 := (\emptyset, \emptyset, \rho_1, \emptyset, \rho_2, \emptyset), \quad \vec{\sigma}_1 := (\rho_2, \rho_3, \emptyset, \rho_2, \emptyset, \rho_3).$$

Together, $(\vec{\sigma}_1, \vec{\sigma}_2)$ form an interleaved splicing of $(\vec{\rho}_1, \vec{\rho}_2)$ and we have that the index 1 is *not* collidable as P_1^1 and P_1^2 are both nonempty. The index 3, on the other hand, is collidable: the collision of $(\vec{\sigma}_1, \vec{\sigma}_2)$ at 3 is given by $((\rho_2, \rho_3, \rho_2, \emptyset, \rho_3), (\emptyset, \emptyset, \rho_1, \rho_2, \emptyset))$ which is no longer interleaved.

Definition 2.14 Let $\vec{\sigma}(i_1, \dots, i_h)$ denote the set of all splittings of all sets of sequences of Reeb chords in $\mathcal{A}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{A}_h^{\otimes i_h}$. We shall say a sequence of maps

$$n_{i_1, \dots, i_h} : M \otimes_{\mathcal{I}} \mathcal{A}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{A}_h^{\otimes i_h} \times \vec{\sigma}(i_1, \dots, i_h) \rightarrow M$$

satisfies a *partial \mathcal{A}_∞ relation* if, for every $x \in M$, set of sequences of Reeb chords $\vec{\rho}$, and interleaving $\vec{\sigma}$ of $\vec{\rho}$, we have that

$$0 = \sum_{\substack{\vec{\lambda} \star \vec{\delta} = \vec{\rho} \\ \vec{\sigma}_1 \star \vec{\sigma}_2 = \vec{\sigma}}} n(n(x, \vec{\lambda}, \vec{\sigma}_1), \vec{\delta}, \vec{\sigma}_2) + \sum_{i, j \text{ compatible}} n(x, \vec{\mu}_j^i(\vec{\rho}), \vec{\sigma}_j^i) + \sum_{k \text{ collidable}}^{|\vec{\sigma}|-1} n(x, \vec{\rho}, \vec{\sigma}(k)),$$

Proposition 2.15 Let n_{i_1, \dots, i_h} be a sequence of maps satisfying the partial \mathcal{A}_∞ relation. Then the maps

$$m_{i_1, \dots, i_h} : M \otimes_{\mathcal{I}} \mathcal{A}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{A}_h^{\otimes i_h} \rightarrow M$$

defined by

$$m_{i_1, \dots, i_h}(x, \vec{\rho}) := \sum_{\vec{\sigma} \text{ an interleaving for } \vec{\rho}} n_{i_1, \dots, i_h}(x, \vec{\rho}, \vec{\sigma})$$

satisfy the \mathcal{A}_∞ relations.

Proof This follows readily by fixing x and $\vec{\rho}$, and summing the partial \mathcal{A}_∞ relations over all interleavings for $\vec{\rho}$. The first two terms in the partial relation sums to the first two terms in the \mathcal{A}_∞ relations, and the final terms cancel in pairs where $\vec{\sigma}$ and $\vec{\sigma}'$ are the two unique interleavings which satisfy $\vec{\sigma}(k) = \vec{\sigma}'(k)$. \square

3 Diagrams

In this section, we will be concerned with *diagrams* of various sorts. The reader is encouraged to read Sections 3.0.1 and 3.0.3 before this section, for some precise cases of the construction and for some examples. We make a broad but vague definition here to try and minimise repetition of associated definitions later:

Definition 3.1 A *diagram* is a tuple $\mathcal{D} = (\Sigma, \boldsymbol{\gamma}; \mathbf{b}, \mathbf{p})$ consisting of the following data:

- (1) a connected, oriented surface Σ of genus g , possibly with boundary;
- (2) an ordered tuple of *attaching curves* $\boldsymbol{\gamma} = \{\boldsymbol{\gamma}^1, \dots, \boldsymbol{\gamma}^m\}$, where each $\boldsymbol{\gamma}^i$ is a set of embedded curves and arcs in Σ ;
- (3) a finite set of marked points $\mathbf{b} \subset \partial\Sigma$;
- (4) a finite set of distinguished points $\mathbf{p} \subset \text{Int}(\Sigma)$.

We make the following broad definition:

Definition 3.2 Let \mathcal{D} be a diagram with two sets of attaching curves and arcs, $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$. A *generator* for $(\boldsymbol{\gamma}, \boldsymbol{\gamma}')$ is a set of intersection points $\mathbf{x} = \{x_1, \dots, x_\kappa\}$ where the x_i are members of $\boldsymbol{\gamma} \cap \boldsymbol{\gamma}'$, such that

- every closed curve in $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ contains precisely one x_i , and
- exactly half of the arcs in each of $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ contain precisely one x_i .

We will denote the set of generators for $(\boldsymbol{\gamma}, \boldsymbol{\gamma}')$ by $\mathcal{G}(\boldsymbol{\gamma}, \boldsymbol{\gamma}')$. With the genericity of the above definition of a diagram, of course, it is possible that $\mathcal{G}(\boldsymbol{\gamma}, \boldsymbol{\gamma}')$ is empty.

In more generality, we will say a *sequence of generators* for a diagram with more than two sets of attaching curves $\boldsymbol{\gamma}^1, \dots, \boldsymbol{\gamma}^m$ is a sequence $\mathbf{x}^{1,2}, \dots, \mathbf{x}^{m,1}$ where $\mathbf{x}^{i,i+1} \in \mathcal{G}(\boldsymbol{\gamma}^i, \boldsymbol{\gamma}^{i+1})$ (here, the superscripts are taken modulo m).

If \mathcal{D} is a diagram with m sets of attaching curves, then for any subset $U \subset \{1, \dots, m\}$ of size l , there is an associated set of attaching curves $\boldsymbol{\gamma}(U) := \{\boldsymbol{\gamma}^i \mid i \in U\}$, and the data $\mathcal{D}|_U := (\Sigma; \boldsymbol{\gamma}(U); \mathbf{p}; \mathbf{b})$ also constitutes a diagram, which we call an *l -subdiagram of \mathcal{D}* .




Figure 1: A neighbourhood of an attaching curve γ , shown in red, and the images of the Hamiltonian isotopy $H(\gamma, t)$ for various $t \in [0, \epsilon]$, where the curves get progressively more orange as t nears ϵ .

We will later be interested in perturbing some of the curves in a diagram. A way to do this is given in [10, Lemma 11.8]. Namely, Lipshitz describes a Hamiltonian $H: \Sigma \rightarrow \mathbb{R}$ which is supported within a tubular neighbourhood of a set of attaching curves, such that the corresponding Hamiltonian isotopy $H(x, t)$ behaves as in Figure 1 (θ refers to the S^1 coordinate of the tubular neighbourhood, identified with $S^1 \times [0, 1]$).

Definition 3.3 Suppose γ is a collection of curves in some diagram. Then we define the ϵ -small m -fold approximation of γ to be the set of sets of arcs and curves $\gamma(\epsilon, m) = \{\gamma^1(\epsilon, m), \dots, \gamma^j(\epsilon, m)\}$, where $\gamma^i(\epsilon, m) = H(\gamma, i\epsilon/m)$. We shall sometimes write $\gamma^{\epsilon/m;i}$ for $\gamma^i(\epsilon, m)$ when it is more convenient.

If $\gamma^1, \dots, \gamma^m$ is a sequence of curves that may be written as $\gamma(\epsilon, m)$ for some γ, ϵ and m , we will often say that $\gamma^1, \dots, \gamma^m$ are *small approximations of one another*, and, when $m = 2$, we shall write $\gamma^1 \approx \gamma^2$. If $\gamma^1, \gamma^2 \subset \Sigma$ are small approximations of one another for some ϵ , we will say that the subset of Σ given by $\bigcup_{t \in [0, \epsilon]} H(\gamma^1, t)$ is the *approximation region* for γ^1, γ^2 . More generally, if $\gamma^1, \dots, \gamma^m$ are all small approximations of one another, we say the approximation region for $\gamma^1, \dots, \gamma^m$ is the union of the approximation regions for each $\gamma^i \approx \gamma^{i+1}$.

Note that in an ϵ -small m -fold approximation, any consecutive set of arcs and curves γ^i, γ^{i+1} constitute the ϵ/m -small 2-fold approximation of γ . For any 2-fold approximation γ^i, γ^{i+1} , we can choose a *distinguished generator* $\theta^{i,i+1}$: a choice of, for every $i = 1, \dots, g + k$, one of the two elements of $\gamma \cap \gamma'$. Hence, given an ϵ -small m -fold approximation, we may choose a distinguished generator for each consecutive pair γ^i, γ^{i+1} to give a set of distinguished generators $\vec{\theta} = \{\theta^{1,2}, \dots, \theta^{m-1,m}\}$.

We will be largely concerned with a canonically defined set of distinguished generators. Namely:

Example 3.4 Consider the intersection points $\theta_{+,j}^{i,i+1} \in \gamma_j^i \cap \gamma_j^{i+1}$ as shown in Figure 1. The collection of $\theta_+^{i,i+1} = \bigcup_j \theta_{+,j}^{i,i+1}$ for every i gives a set of distinguished generators $\vec{\theta}_+$.

Sometimes, we refer to the other choice of intersection point in $\gamma_j^i \cap \gamma_j^{i+1}$ as $\theta_{-,j}^{i,i+1}$.

Definition 3.5 Suppose γ and γ' are two collections of arcs and curves in some diagram, and let $x \in \mathcal{G}(\gamma, \gamma')$ be a generator.

If $\gamma^{\varepsilon/m;1}, \dots, \gamma^{\varepsilon/m;m}$ is the ε -small m -fold approximation to γ , then, provided ε is sufficiently small, there are a series of canonical *nearest point maps*, which are given by, for each i , the map

$$\varphi_i: \mathcal{G}(\gamma, \gamma') \rightarrow \mathcal{G}(\gamma^{\varepsilon/m;i}, \gamma')$$

which sends $x_j \in \gamma_j \cap \gamma'_{\sigma(j)}$ to the canonical nearest point $x_j^{\varepsilon/m;i} \in \gamma_j^{\varepsilon/m;i} \cap \gamma'_{\sigma(j)}$.

If we are given a set of distinguished generators $\vec{\theta}$ for the approximation, then this allows us to associate to any pair of generators $x, y \in \mathcal{G}(\gamma, \gamma')$ a sequence of generators

$$(x', \theta^{1,2}, \dots, \theta^{m-1,m}, y')$$

in the set

$$\mathcal{G}(\gamma', \gamma^{\varepsilon/m;1}) \times \mathcal{G}(\gamma^{\varepsilon/m;1}, \gamma^{\varepsilon/m;2}) \times \dots \times \mathcal{G}(\gamma^{\varepsilon/m;m-1}, \gamma^{\varepsilon/m;m}) \times \mathcal{G}(\gamma^{\varepsilon/m;m}, \gamma'),$$

where $x' = x$, and $y' = \varphi_m(y)$. We call this sequence the ε -small m -fold *splaying* of (x, y) , and shall often abuse notation and write $(x, \theta^{1,2}, \dots, \theta^{m-1,m}, y)$ for the relevant generator.

3.0.1 Toroidal bordered Heegaard diagrams One of the main instances of a diagram we shall consider is when \mathcal{D} specifies a three-manifold with boundary that is a disjoint union of tori. The definition we shall use is:

Definition 3.6 Suppose that $\mathcal{D} = (\Sigma; \gamma; \mathbf{b}; \mathbf{p})$ is a diagram. Then we shall say that \mathcal{D} is a *toroidal bordered Heegaard diagram* if:

- (1) Σ is a genus g surface with h boundary components.
- (2) γ consists of only two sets of pairwise disjoint curves α and β , which intersect transversely, such that β consists of $g - 1 + h$ simple closed curves, and α can be partitioned into two sets, $\alpha = \alpha^a \cup \alpha^c$, where α^a consists of $2h$ arcs and α^c consists of $g - 1$ closed curves.
- (3) The set α^a can be written $\alpha^a = \{\alpha_1^m, \alpha_1^l, \dots, \alpha_h^m, \alpha_h^l\}$, where α_i^m and α_i^l intersect $\partial\Sigma_i$ relative to the point $b_i \in \partial\Sigma_i$ as shown in Figure 2, and is disjoint from $\partial\Sigma_j$ for $i \neq j$.
- (4) Every region of the surfaces $\Sigma - \alpha$ and $\Sigma - \beta$ meets some basepoint $b_i \in \mathbf{b}$ (we sometimes call this *homological linear independence*).
- (5) The set of points \mathbf{p} is empty (we therefore usually omit this from the notation).




Figure 2: Labelling conventions for α_i^m and α_i^l .

An example of a toroidal bordered Heegaard diagram with 2 boundary components can be found on the left of Figure 8.

In the case where \mathcal{D} is a toroidal bordered Heegaard diagram with $|\partial\Sigma| = 1$, this constitutes a bordered Heegaard diagram compatible with a manifold with torus boundary in the sense of [14]. When $|\partial\Sigma| > 1$, the data corresponds to a *bordered sutured Heegaard diagram* in the sense of [26]. The data of $\mathbf{b} = \{b_1, \dots, b_h\}$ corresponds to a series of arcs $\Gamma_{\mathbf{b}} \subset \partial\Sigma$, so that α^a and $\Gamma_{\mathbf{b}}$ specify an embedding of the disjoint union of h arc diagrams for the torus into Σ ; that is, \mathcal{D} is equivalent to a bordered sutured Heegaard diagram for a three-manifold with h torus boundary components. We will tend to ignore the distinction between \mathcal{D} and the corresponding bordered sutured diagram.

Convention 3.7 To avoid proliferation of adjectives, we will often refer to a toroidal bordered Heegaard diagram as simply a *bordered diagram*.

Note that when \mathcal{D} is a bordered diagram, a generator for \mathcal{D} is a set of intersection points x_1, \dots, x_{g+h-1} , where $x_1, \dots, x_h \in \alpha^a \cap \beta$ and $x_{h+1}, \dots, x_{g+h-1} \in \alpha^c \cap \beta$. Without loss of generality, for $i = 1, \dots, h$, we have that $x_i \in \alpha_i^{\delta_i} \cap \beta_{\sigma(i)}$, where $\delta_i \in \{m, l\}$ and σ is some permutation of $\{1, \dots, g\}$. If $\delta = (\delta_1, \dots, \delta_h)$ is a sequence of elements of $\{m, l\}$, then we can associate the idempotent $\iota(\delta) = (\iota_{\delta_1}, \dots, \iota_{\delta_h})$ in \mathcal{T}^h . Accordingly, we may associate to any generator \mathbf{x} the corresponding idempotent $\iota(\mathbf{x}) := \iota(\delta)$ — this is called $I_A(\mathbf{x})$ in [14].

It will be of use to us later to note that although the set \mathbf{p} is empty if \mathcal{D} is a bordered diagram, there is a set of points $\mathbf{p}(\mathbf{b}) \subset \text{Int}(\Sigma)$ given by pushing each $b_i \in \mathbf{b}$ slightly off of the boundary $\partial\Sigma_i$ into the region of $\Sigma - \alpha - \beta$ adjacent to b_i .

3.0.2 Link diagrams Another important instance of a diagram is what we term a *2h-pointed link diagram*; cf [20, Definition 3.1] where these are called *2l-pointed Heegaard diagrams*; these are

Definition 3.8 Let $\mathcal{D} = (\Sigma; \boldsymbol{\gamma}; \mathbf{b}; \mathbf{p})$ be a diagram. Then we say \mathcal{D} is a *2h-pointed link diagram* if:

- (1) Σ is a closed surface of genus g .

- (2) $\boldsymbol{\gamma}$ consists of two sets of pairwise disjoint curves $\boldsymbol{\gamma}^1$ and $\boldsymbol{\gamma}^2$, which intersect transversely, such that each is a set of $g + h - 1$ simple closed curves (ie none of them are arcs).
- (3) Each set $\boldsymbol{\gamma}^i$ spans a g -dimensional sublattice of $H_1(\Sigma)$. Equivalently, for each i , the surface $\Sigma - \boldsymbol{\gamma}^i$ has precisely h components A_1^i, \dots, A_h^i .
- (4) The set \boldsymbol{b} is empty (this is implied by Σ being closed, and we will often omit it from the notation).
- (5) The set \boldsymbol{p} can be ordered as $\boldsymbol{p} = (w_1, z_1, \dots, w_h, z_h)$, such that $w_j, z_j \in A_j^1 \cap A_j^2$. We will often write $\boldsymbol{p} = \boldsymbol{w} \cup \boldsymbol{z}$ for the partition of \boldsymbol{p} into (w_1, \dots, w_h) and (z_1, \dots, z_h) respectively.

Often, we write $\boldsymbol{\gamma}^1 := \boldsymbol{\beta}$ and $\boldsymbol{\gamma}^2 := \boldsymbol{\alpha}$. For an example of a 4-pointed link diagram with $g = 0$ and $h = 2$, please see Figure 3 — at this stage, the purple and orange curves are to be ignored. Another example of a link diagram with basepoints omitted can be found on the right of Figure 8.

3.0.3 Heegaard multidigraphs Another important instance of a diagram is a *Heegaard multidigraph*. We follow approximately the treatment given in [13], though the notion goes back as far as [18, Section 8.1.5].

Definition 3.9 Let $\mathcal{D} = (\Sigma; \boldsymbol{\gamma}; \boldsymbol{b}; \boldsymbol{p})$ be a diagram. Then we say \mathcal{D} is a *Heegaard m -diagram* if:

- (1) Σ is a closed surface of genus g .
- (2) $\boldsymbol{\gamma}$ consists of m sets of $g + h - 1$ pairwise disjoint simple closed curves $\boldsymbol{\gamma}^i$, such that each $\boldsymbol{\gamma}^i$ intersects $\boldsymbol{\gamma}^j$ transversally for $i \neq j$.
- (3) For each i , the surface $\Sigma - \boldsymbol{\gamma}^i$ has precisely h components A_1^i, \dots, A_h^i .
- (4) The set \boldsymbol{b} is empty (this is implied by Σ being closed, and we will often omit it from the notation).
- (5) The set \boldsymbol{p} consists of h points, $\boldsymbol{p} = (z_1, \dots, z_h)$, such that (without loss of generality) $z_j \in \bigcap_{i=1}^m A_j^i$.

Examples of 3- and 4-diagrams with basepoints omitted can be found on the left and right (respectively) of Figure 10.

3.0.4 Bordered multidigraphs We roughly generalise both a toroidal bordered diagram and a Heegaard m -diagram:

Definition 3.10 Let $\mathcal{D} = \mathcal{D} = (\Sigma; \boldsymbol{\gamma}^0, \boldsymbol{\gamma}^1, \dots, \boldsymbol{\gamma}^{m-1}; \boldsymbol{b}; \boldsymbol{p})$ be a diagram. Then we say \mathcal{D} is a *toroidal bordered multidigraph* if:

- (1) Σ is a surface with h boundary components $\partial\Sigma_1, \dots, \partial\Sigma_h$.
- (2) $(\Sigma, \boldsymbol{\gamma}^0, \boldsymbol{\gamma}^1, \boldsymbol{b})$ is a toroidal bordered Heegaard diagram (such that $\boldsymbol{\gamma}^0$ are closed curves, and $\boldsymbol{\gamma}^1 = \boldsymbol{\gamma}^{1,a} \cup \boldsymbol{\gamma}^{1,c}$ splits as a series of $2h$ arcs and $g - 1$ closed curves).
- (3) For $i \geq 2$ each of the sets of curves $\boldsymbol{\gamma}^i$ is a set of $g + h - 1$ pairwise disjoint simple closed curves such that each $\boldsymbol{\gamma}^i$ intersects $\boldsymbol{\gamma}^j$ transversally for $i \neq j$.

- (4) For each $i > 1$, the surface $\Sigma - \gamma^i$ has precisely h components A_1^i, \dots, A_h^i .
- (5) The set $\mathbf{p} = (z_1, \dots, z_h)$ is given by the push-off $\mathbf{p}(\mathbf{b})$, and (without loss of generality) $z_j \in \bigcap_{i=1}^m A_j^i$.

Given a toroidal bordered multidiagram, there are two associated surfaces which we will have cause to mention later. The first is the associated compact surface $\bar{\Sigma}$, which comes with associated arcs $\bar{\gamma}^1$ — we may view Σ as $\bar{\Sigma} - \partial\bar{\Sigma}$. The second is the surface $\Sigma_{\bar{e}}$ which is the result of collapsing each boundary circle $\partial\bar{\Sigma}_i$ of $\bar{\Sigma}$ down to a point e_j^∞ — thus we can also view Σ as a punctured surface $\Sigma_{\bar{e}} \setminus \{e_1^\infty, \dots, e_h^\infty\}$. This comes endowed with natural closed curves $\gamma_{\bar{e}}$ which are the closures of γ under this identification, and basepoints $\mathbf{p}_{\bar{e}}$ such that $\mathcal{D}_{\bar{e}} = (\Sigma_{\bar{e}}; \gamma_{\bar{e}}; \mathbf{p}_{\bar{e}})$ is a diagram, which we sometimes call the *east compactification* of \mathcal{D} .

3.0.5 Converting between diagrams Most of the content of this paper is in relating counts of certain moduli spaces related to bordered diagrams with those related to multipointed Heegaard multidiagrams. As such, we will need to have a dictionary between concepts for the above two diagrams.

We will work with a special class of diagrams in both cases. We first need to introduce the notion of a *handle*. Please see [Figure 6](#) for a picture of the closed -6 -twisted handle with the curves $\beta_{\mathfrak{H}}$ and $\alpha_{\mathfrak{H}}(l)$ pictured; and [Figure 5](#) for a picture of a closed handle with the curves $\beta_{\mathfrak{H}}$ and $\alpha_{\mathfrak{H}}(m)$ shown. A bordered -6 -twisted handle is also shown in [Figure 4](#).

Definition 3.11 For $n \in \mathbb{Z}$, the *closed n -twisted handle* is a tuple

$$\mathfrak{H}(n) = (H; \alpha_{\mathfrak{H}}(m), \beta_{\mathfrak{H}}, \alpha_{\mathfrak{H}}(l); \mathbf{p}_{\mathfrak{H}}),$$

where

- H is the tube $S^1 \times [-1, 1]$, with left and right ends $\mathfrak{H}_L := S^1 \times \{-1\}$ and $\mathfrak{H}_R := S^1 \times \{1\}$ respectively;
- $\alpha_{\mathfrak{H}}(m)$ is the curve $S^1 \times \{0\}$;
- $\beta_{\mathfrak{H}}$ is the arc $\{0\} \times [-1, 1]$;
- $\alpha_{\mathfrak{H}}(l)$ is the arc $\{(e^{ni\pi t}, t) \mid t \in [-1, 1]\}$.
- $\mathbf{p}_{\mathfrak{H}}$ is the pair of basepoints $(w_{\mathfrak{H}}, z_{\mathfrak{H}})$ as defined below.

For the definition of $\mathbf{p}_{\mathfrak{H}}$, we choose a disc neighbourhood of the point $\alpha_{\mathfrak{H}}(m) \cap \alpha_{\mathfrak{H}}(l)$ which is sufficiently small that the triple $(D, \alpha_{\mathfrak{H}}(m) \cap D, \alpha_{\mathfrak{H}}(l) \cap D)$ is modelled on the unit disc $\mathbb{D} \subset \mathbb{C}$, together with the imaginary and real axes respectively. Under this identification, we put $w_{\mathfrak{H}} = \frac{1}{2}e^{3\pi i/4}$ and $z_{\mathfrak{H}} = \frac{1}{2}e^{-3\pi i/4}$.

The *bordered n -twisted handle* is a tuple $\mathfrak{H}(n) = (H; \alpha_{\mathfrak{H}}(m), \beta_{\mathfrak{H}}, \alpha_{\mathfrak{H}}(l); b)$, where H , $\alpha_{\mathfrak{H}}(m)$, $\beta_{\mathfrak{H}}$ and $\alpha_{\mathfrak{H}}(l)$ are the punctured tube and arcs induced from H , $\alpha_{\mathfrak{H}}(m)$, $\beta_{\mathfrak{H}}$ and $\alpha_{\mathfrak{H}}(l)$ by removing D from H . Under the identification of \mathbb{D} with D , the boundary of H consists of three components: \mathfrak{H}_L , \mathfrak{H}_R and a new circle \mathfrak{H}_∂ modelled upon $S^1 = \partial\mathbb{D}$. We define b to be the point in η_∂ that corresponds to $e^{-\pi i/4}$.

It will be useful for us to note here that there is a clear bijection between the sets of curves $\alpha_{\mathfrak{H}}(m)$, $\beta_{\mathfrak{H}}$, $\alpha_{\mathfrak{H}}(l)$ and $\alpha_{\mathfrak{H}}(m)$, $\beta_{\mathfrak{H}}$, $\alpha_{\mathfrak{H}}(l)$ respectively, given by the map $\cap_D(\gamma) = \gamma \cap (\Sigma \setminus D)$.

We will shortly construct some diagrams by gluing handles to a multipointed Heegaard multidiagram in a way specified by some fixed data. First, we consider sequences of idempotents:

Definition 3.12 An *h-splaying sequence of idempotents* is a sequence $\vec{\iota}$ of $m + 1$ sequences of elements in the subring of idempotents of the h -torus algebra; $\vec{\iota} = \vec{\iota}^0, \dots, \vec{\iota}^m$, where $\vec{\iota}^i = (\iota_1^i, \dots, \iota_h^i)$ and each ι_j^i is one of ι_m or ι_l .

Note that if $\vec{\rho} = \vec{\rho}_1, \dots, \vec{\rho}_h$ is a set of sequences of Reeb chords and $\vec{\sigma}$ is a splicing for $\vec{\rho}$, the associated sequence $\vec{\iota}_\sigma$ is an *h-splaying sequence of idempotents*. In general, the data we shall require is as follows.

Definition 3.13 Suppose \mathcal{D} is a link diagram $(\Sigma; \alpha, \beta; \mathbf{p})$, where $|\mathbf{p}| = 2h$. By definition, we can number the regions A of $\Sigma - \alpha$ and B of $\Sigma - \beta$ by the numbers $1, \dots, h$ so that the points $\mathbf{p} = \{w_1, z_1, \dots, w_h, z_h\}$ satisfy $(w_i, z_i) \in A_i \cap B_i$.

A set of *stabilisation data for \mathcal{D}* is a tuple $d = (\lambda_{w,z}, \lambda_{z,w}, \mathbf{n}, \varepsilon, \vec{\iota})$, where:

- (1) $\lambda_{w,z} = \{\lambda_{w_1,z_1}, \dots, \lambda_{w_h,z_h}\}$ and $\lambda_{z,w} = \{\lambda_{z_1,w_1}, \dots, \lambda_{z_h,w_h}\}$ are oriented arcs in Σ chosen such that
 - $\partial\lambda_{w_i,z_i} = z_i - w_i$,
 - $\partial\lambda_{z_i,w_i} = w_i - z_i$,
 - $\lambda_{w,z} \cap \alpha = \emptyset$, and
 - $\lambda_{z,w} \cap \beta = \emptyset$.
- (2) $\mathbf{n} = (n_1, \dots, n_h)$ is a sequence of integers.
- (3) ε is a small approximation parameter.
- (4) $\vec{\iota}$ is an *h-splaying sequence of idempotents*.

This is demonstrated in Figure 3 for a specific choice of link diagram. (This is a 4-pointed Heegaard diagram for the Hopf link in S^3 .)

We now describe the process of gluing handles to a link diagram. The reader is strongly encouraged to refer to Figure 4 either whilst or instead of reading the following definition.

Definition 3.14 Let $\mathcal{D} = (\Sigma; \alpha, \beta; \mathbf{p})$ be a link diagram, where $|\mathbf{p}| = 2h$, and let $d = (\lambda_{w,z}, \lambda_{z,w}, \mathbf{n}, \varepsilon, \vec{\iota})$ be a set of stabilisation data for \mathcal{D} . Suppose that $(\mathfrak{H}_1, \dots, \mathfrak{H}_h)$ is a set of bordered handles.

We label the elements of \mathbf{p} by $(w_1, z_1, \dots, w_h, z_h)$, where w_i and z_i both belong to the same component of $\Sigma - \alpha$ and $\Sigma - \beta$. Choose a pair of small neighbourhoods N_i^w and N_i^z around w_i and z_i which are holomorphically identified with the unit disc $\mathbb{D} \subset \mathbb{C}$ in such a way that the basepoint z_i is identified with the origin and the arcs $\lambda_{z_i,w_i} \cap N_i^z$ and $\lambda_{w_i,z_i} \cap N_i^z$ are identified with the positive and negative sides of the real axis respectively; and that $\lambda_{w_i,z_i} \cap N_i^w$ and $\lambda_{w_i,z_i} \cap N_i^w$ are identified with the positive and negative sides of the real axis respectively.

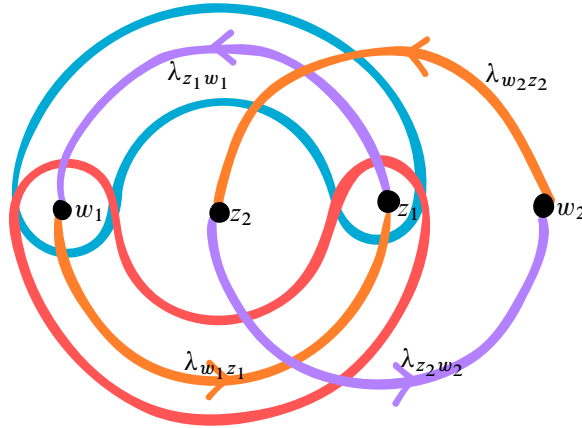


Figure 3: Stabilisation data for a 4-pointed link diagram. Here, the α curves are shown in red and the β in blue.

Now choose homeomorphisms h_i^L (resp. h_i^R) from \mathfrak{S}_i^L (resp. \mathfrak{S}_i^R) to ∂N_i^w (resp. ∂N_i^z), in such a way that $h_i^L(\alpha_{\mathfrak{S}_i}(l) \cap \mathfrak{S}_i^L) = \lambda_{z_i, w_i} \cap N_i^w$ and $h_i^L(\beta_{\mathfrak{S}_i} \cap \mathfrak{S}_i^L) = \lambda_{w_i, z_i} \cap N_i^w$ (resp. $h_i^R(\alpha_{\mathfrak{S}_i}(l) \cap \mathfrak{S}_i^R) = \lambda_{z_i, w_i} \cap N_i^w$ and $h_i^R(\beta_{\mathfrak{S}_i} \cap \mathfrak{S}_i^R) = \lambda_{w_i, z_i} \cap N_i^w$).

Form the surface

$$\Sigma(d) := (\Sigma - N_1^w - N_2^z - \dots - N_h^w - N_h^z) \cup_{h^L, h^R} (H_1, \dots, H_h)$$

by gluing the handles to Σ via the homeomorphisms h_i^L and h_i^R . We denote by **glue** the composition of the projection map from Σ to $(\Sigma - N_1^w - N_2^z - \dots - N_h^w - N_h^z)$ with the gluing map; that is, it is a map

$$\mathbf{glue}: \Sigma \sqcup H_1 \sqcup \dots \sqcup H_h \rightarrow \Sigma(d).$$

There is a set of h curves $\beta_{\mathfrak{S}} := (\lambda_{w_1, z_1} \cup_{h_1^L, h_1^R} \beta_{\mathfrak{S}_1}, \dots, \lambda_{w_h, z_h} \cup_{h_h^L, h_h^R} \beta_{\mathfrak{S}_h})$ contained in $\Sigma(d)$, and also a sequence of arcs $\alpha_{\mathfrak{S}}^a$ given by $\alpha^a(m_i) := \mathbf{glue}(\alpha_{\mathfrak{S}_i}(m))$ and $\alpha^a(l_i) := \alpha_{\mathfrak{S}_i}(l) \cup_{h_h^L, h_h^R} \lambda_{z_1, w_1}$, for $i = 1, \dots, h$.




Figure 4: The bordered stabilisation of a link diagram at one pair of basepoints.

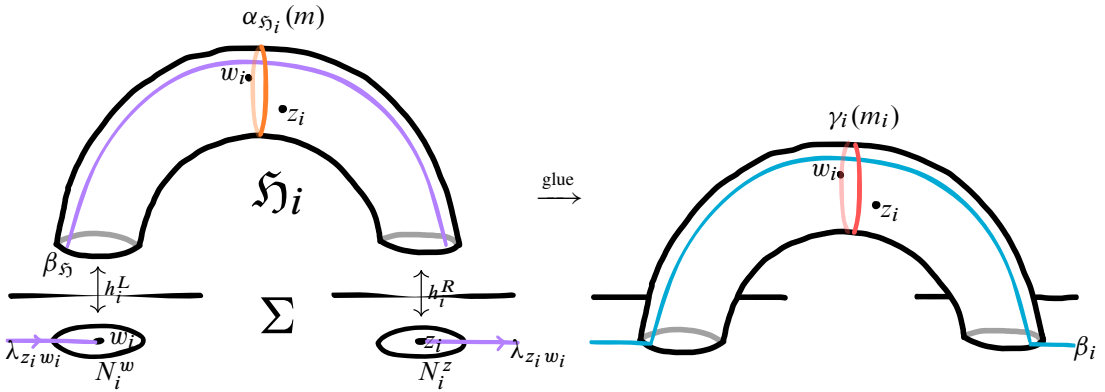


Figure 5: The closed stabilisation of a link diagram at one pair of basepoints w_i, z_i , with meridional corresponding idempotent.

These allow us to form the bordered diagram $\mathcal{D}(d) = (\Sigma(d), \alpha(d)^a, \alpha(d)^c, \beta(d), \mathbf{b}(d))$, where

- $\Sigma(d)$ is as above,
- $\alpha(d)^a$ are the arcs $\alpha_{\mathfrak{H}_i}^a$,
- $\alpha(d)^c$ are the images $\mathbf{glue}(\alpha)$,
- $\beta(d)$ is the union of $\beta_{\mathfrak{H}_i}$ and the images $\mathbf{glue}(\beta)$, and
- $\mathbf{b}(d)$ is the image $\mathbf{glue}(\mathbf{b}(\mathfrak{H}_1) \cup \dots \cup \mathbf{b}(\mathfrak{H}_h))$.

We call this diagram the *bordering* of \mathcal{D} ; any bordered diagram that arises as $\mathcal{D}(d)$ for some (D, d) we will call *stable*.

Note that in this construction, the elements of the approximation data ε and \vec{l} were redundant. A small adaptation of the above also gives us a multipointed Heegaard multidiagram (illustrated in Figures 5 and 6):

Definition 3.15 Let $\mathcal{D} = (\Sigma; \alpha, \beta; \mathbf{p})$ be a link diagram, where $|\mathbf{p}| = 2h$, and let $d = (\lambda_{w,z}, \lambda_{z,w}, \mathbf{n}, \varepsilon, \vec{l})$ be a set of stabilisation data, such that \vec{l} consists of a single idempotent of \mathcal{T}^h , $\iota = (\iota_{\delta_1}, \dots, \iota_{\delta_h})$ where each $\delta_i \in \{m, l\}$. Suppose further that $(\mathfrak{H}_1, \dots, \mathfrak{H}_h)$ is a set of closed handles.

One can imitate the construction of the surface $\Sigma(d)$ from the previous definition, but with $(\mathfrak{H}_1, \dots, \mathfrak{H}_h)$ in place of $(\mathfrak{H}_1, \dots, \mathfrak{H}_h)$ to yield a surface $\Sigma(d, \iota) := \Sigma \cup_{h^L, h^R} (H_1, \dots, H_h)$ by gluing the handles to Σ via homeomorphisms h_i^L and h_i^R . We denote by \mathbf{glue} the corresponding gluing map

$$\mathbf{glue}: \Sigma \sqcup H_1 \sqcup \dots \sqcup H_h \rightarrow \Sigma(d, \iota).$$

Again, there is a set of h curves $\beta_{\mathfrak{H}_i} = (\lambda_{w_1, z_1} \cup_{h_i^L, h_i^R} \beta_{\mathfrak{H}_{i1}}, \dots, \lambda_{w_n, z_n} \cup_{h_n^L, h_n^R} \beta_{\mathfrak{H}_{in}})$ contained in $\Sigma(d, \iota)$. There is also a sequence of curves contingent on ι , given by the set $\alpha_{\mathfrak{H}_i}(d, \iota) = \{\gamma_i(\delta_i)\}$, where

$$\gamma_i(\delta_i) = \begin{cases} \mathbf{glue}(\alpha_{\mathfrak{H}_i}(m)) & \text{if } \delta_i = m, \\ \alpha_{\mathfrak{H}_i}(l) \cup_{h_i^L, h_i^R} \lambda_{z_i, w_i} & \text{if } \delta_i = l. \end{cases}$$




Figure 6: The closed stabilisation of a link diagram at one pair of basepoints w_i, z_i , with longitudinal corresponding idempotent.

These allow us to form the diagram $\mathcal{D}(d, \iota) = (\Sigma(d, \iota), \alpha(d, \iota), \beta(d, \iota), \mathbf{p}(d, \iota))$, where

- $\Sigma(d, \iota)$ is as above,
- $\alpha(d, \iota)$ is the union of $\alpha_{\mathfrak{S}_i}(d, \iota)$ and the images $\text{glue}(\alpha)$,
- $\beta(d, \iota)$ is the union of $\beta_{\mathfrak{S}_i}$ and the images $\text{glue}(\beta)$, and
- $\mathbf{p}(d, \iota)$ is the image $\text{glue}(\mathbf{p}_{\mathfrak{S}_1} \cup \dots \cup \mathbf{p}_{\mathfrak{S}_h})$.

Note that $\cap_D : \mathfrak{S} \rightarrow \mathfrak{S}$ extends to a map $\cap_D : \mathcal{D} \rightarrow \mathcal{D}(d, \iota)$, which yields a bijection $\cap_D : \beta(d, \iota) \rightarrow \beta(d)$. It no longer yields a bijection of the sets $\alpha(d, \iota)$; rather, it gives an injective map sending $\text{glue}(\alpha) \mapsto \mathbf{glue}(\alpha)$ bijectively, and each element $\gamma_i(\delta_i) \in \alpha_{\mathfrak{S}_i}(d, \iota)$ to the corresponding arc $\cap_D(\gamma_i(\delta_i)) = \alpha^a(\delta_i)$. It follows that to every generator $x \in \mathcal{G}(\mathcal{D})$ we may associate a generator $\cap_D^{-1}(x) \in \mathcal{D}(d, \iota(x))$ (we often abuse notation for this generator and also denote it by x).

The diagram $\mathcal{D}(d, \iota)$ as defined above constitutes a multipointed Heegaard multidigraph (in fact a simple Heegaard diagram), but does not use the information ε from the approximation data d , and is restricted to the case where $\vec{\iota}$ is a single set of idempotents. In more generality, we make a similar construction:

Definition 3.16 Let $\mathcal{D} = (\Sigma; \alpha, \beta; \mathbf{p})$ be a link diagram, where $|\mathbf{p}| = 2h$, and let $d = (\lambda_{w,z}, \lambda_{z,w}, \mathbf{n}, \varepsilon, \vec{\iota})$ be a set of stabilisation data for \mathcal{D} such that $\vec{\iota} = \iota^1, \dots, \iota^m$.

The *splayed diagram* associated to the pair (\mathcal{D}, d) is the tuple

$$\mathcal{D}(d) := (\Sigma(d), \alpha^1(d), \dots, \alpha^m(d), \beta(d), \mathbf{p}(d))$$

given by:

- $\Sigma(d) = \Sigma(d, \iota^1)$ (this is the same as $\Sigma(d, \iota^i)$ for any i).
- $\alpha^i(d)$ is given by the ε -small approximation $H(\alpha(d, \iota(i)), i\varepsilon/m)$ to the set of curves $\alpha(d, \iota(i))$ for each $i = 1, \dots, m$, as described in Definition 3.15.




Figure 7: A splayed diagram near e_j^∞ .

Here, for each $i = 1, \dots, h$, if we denote by $\alpha(d, \iota, \mathfrak{H}_i)$ the curve of $\alpha(d, \iota)$ which intersects \mathfrak{H}_i , then we chose the identification of the tubular neighbourhood $N(\alpha, i)$ of $\alpha(d, \iota, \mathfrak{H}_i)$ with $S^1 \times [0, 1]$ so that the intersection of the complementary curve $\alpha(d, \iota^c, \mathfrak{H}_i)$ with $N(\alpha, i)$ coincides with the arc $\pi \times [0, 1]$. (See Figure 7.)

- $\beta(d)$ is given by the set $\beta(d, \iota_1)$ (this, again, is really independent of ι_1).
- $p(d)$ is given by the set $\bigcup_{i=1}^m p_{\mathfrak{H}_i}$.

If \mathcal{D} is a stable bordered diagram $\mathcal{D} = \mathcal{D}(d)$ for some d and \mathcal{D} , then we will refer to the corresponding splayed diagram $\mathcal{D}(d)$ as the *splaying* of \mathcal{D} . We will often denote this simply by $\mathcal{D}(\varepsilon, \vec{\iota})$, as these are the only parts of the splaying data d not determined by the bordered diagram $\mathcal{D}(d)$.

There is one further variant which we shall need to consider:

Definition 3.17 Suppose that \mathcal{D} is a link diagram, and that d is a set of stabilisation data for \mathcal{D} . Consider the corresponding splayed diagram

$$\mathcal{D}(d, \vec{\iota}) := (\Sigma(d), \alpha^1(d), \dots, \alpha^m(d), \beta(d), p(d)).$$

We note that each $\alpha^i(d)$ is disjoint from the points e^∞ . It follows that there is an induced set of sets of curves in the surface $\Sigma(d)$, given by the images of the $\alpha^i(d)$ under the map $\cap_{\mathcal{D}}$. As such, there is a corresponding bordered multidiagram $\mathcal{D}^\partial(d, \vec{\iota})$ given by

$$\mathcal{D}^\partial(d, \vec{\iota}) := (\Sigma(d, \vec{\iota}), \alpha^a(d, \vec{\iota}) \cup \alpha^c(d, \vec{\iota}), \alpha^1(d, \vec{\iota}), \dots, \alpha^m(d, \vec{\iota}), \beta(d, \vec{\iota}), p(d, \vec{\iota})).$$

(We often relabel so that $\Sigma(d, \vec{\iota}) = \Sigma(d)$, etc.)

We call this the *partial splaying* of \mathcal{D} associated to the data $d, \vec{\iota}$.

Later, we will need to perturb subdiagrams into a more standard form:

Definition 3.18 Suppose that \mathcal{D} is a splayed or partially splayed diagram, and consider a subdiagram \mathcal{D}' with attaching curves $\boldsymbol{\gamma} = \boldsymbol{\gamma}^{i_1}, \dots, \boldsymbol{\gamma}^{i_k}$. Let $\boldsymbol{\gamma}' = \boldsymbol{\gamma}^{i'_1}, \dots, \boldsymbol{\gamma}^{i'_{k'}}$ denote the subsequence of these attaching curves obtained by omitting any $\boldsymbol{\gamma}^{i_j}$ such that $\boldsymbol{\gamma}^{i_{j-1}} \equiv \boldsymbol{\gamma}^{i_j}$.

Each set of curves $\boldsymbol{\gamma}^{i'_j}$ is given either by $H(\boldsymbol{\alpha}(d, \vec{l}^{i'_j}, i'_j \varepsilon / m))$ for some sequence of idempotents $\vec{l}^{i'_j}$ or is equal to $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ — say there are m' curves of the first sort. We let $\boldsymbol{\gamma}_{\text{reg}}$ denote the sequence of attaching curves obtained from $\boldsymbol{\gamma}'$ by replacing each $\boldsymbol{\gamma}^{i'_j}$ of the first sort with the perturbation

$$\boldsymbol{\gamma}_{\text{reg}}^{i'_j} := H(\boldsymbol{\alpha}(d, \vec{l}^{i'_j}, j \varepsilon / m')),$$

and $\mathcal{D}'_{\text{reg}}$ be the diagram obtained from \mathcal{D}' by replacing the attaching curves used with $\boldsymbol{\gamma}_{\text{reg}}$.

We shall call this diagram the *regularisation* of \mathcal{D}' .

The bordered stabilisation and splaying of a Heegaard diagram exhibit a number of similarities.

Let $\mathcal{D} := (\Sigma; \boldsymbol{\alpha}, \boldsymbol{\beta}; \boldsymbol{p})$ be a link diagram, and $d = (\boldsymbol{\lambda}_{w,z}, \boldsymbol{\lambda}_{z,w}, \boldsymbol{n}, \varepsilon, \vec{l})$ some stabilisation data for \mathcal{D} . Recall that associated to the curves $\boldsymbol{\alpha}$ are the m -fold ε approximations $\boldsymbol{\alpha}^1, \dots, \boldsymbol{\alpha}^m$.

The image of $\boldsymbol{\alpha}^i$ under the map glue coincides with a set of $g - 1$ of the closed stabilisation curves $\boldsymbol{\alpha}^i(d) \subset \mathcal{D}(d) = (\Sigma(d); \boldsymbol{\alpha}^1(d), \dots, \boldsymbol{\alpha}^m(d); \boldsymbol{\beta}(d); \boldsymbol{p}(d))$, or similarly in the partial splaying of \mathcal{D} . We number the $\boldsymbol{\alpha}^i(d)$ so that these $g - h$ curves are given by the set $\{\boldsymbol{\alpha}_{h+1}^i(d), \dots, \boldsymbol{\alpha}_{g+h-1}^i(d)\} \subset \boldsymbol{\alpha}^i(d)$. It follows that to any distinguished generators $\vec{\boldsymbol{\theta}}_{\text{approx}} = \boldsymbol{\theta}_{\text{approx}}^{1,2}, \dots, \boldsymbol{\theta}_{\text{approx}}^{m-1,m}$ for the m -fold ε -approximation of $\boldsymbol{\alpha}$ we may associate a corresponding set of partial generators $\text{glue}(\vec{\boldsymbol{\theta}}_{\text{approx}})$ for the splayed or partially splayed diagram $\mathcal{D}(d)$, which we also denote by $\vec{\boldsymbol{\theta}}_{\text{approx}}$.

The curves in each $\boldsymbol{\alpha}^i(d)$ not in the image of $\text{glue}(\boldsymbol{\alpha}^i)$, which we assume are labelled by $\boldsymbol{\alpha}_1^i, \dots, \boldsymbol{\alpha}_h^i$ such that, for each i either:

- $\boldsymbol{\alpha}_j^i \cap \boldsymbol{\alpha}_j^{i+1}$ is a single point $e_j^{\infty,i}$ (this corresponds with the idempotents l_j^i and l_j^{i+1} being different).
- $\boldsymbol{\alpha}_j^i$ and $\boldsymbol{\alpha}_j^{i+1}$ are Hamiltonian isotopic (in fact, they are ε/m -small approximations of one another).

In the second case, there are two points in $\boldsymbol{\alpha}_j^i \cap \boldsymbol{\alpha}_j^{i+1}$. For every such i, j , we choose an element $\boldsymbol{\theta}_j^{i,i+1} \in \boldsymbol{\alpha}_j^i \cap \boldsymbol{\alpha}_j^{i+1}$. We term such a choice a *splaying extension* of the distinguished generators $\vec{\boldsymbol{\theta}}_{\text{approx}}$. For every fixed i , we collect the intersection points $e_j^{\infty,i,i+1}$ and $\boldsymbol{\theta}_j^{i,i+1}$ into a set $\boldsymbol{\theta}_{\text{ext}}^{i,i+1}$, and set $\boldsymbol{\theta}^{i,i+1} = \boldsymbol{\theta}_{\text{ext}}^{i,i+1} \cup \boldsymbol{\theta}_{\text{approx}}^{i,i+1}$ — this is a bona fide set of generators dependent upon both the choice of distinguished generators and the choice of splaying extension.

Example 3.19 Suppose that $\vec{\boldsymbol{\theta}}_{\text{approx}}$ is the distinguished set of generators $\vec{\boldsymbol{\theta}}_+$ described in [Example 3.4](#). Then we shall call the corresponding set of $\boldsymbol{\theta}^{i,i+1}$ by the same name $\boldsymbol{\theta}_+^{i,i+1}$ — and hope that this does not cause confusion.

If we wish to draw attention to the number of generators in $\vec{\boldsymbol{\theta}}_+$, we will write $\vec{\boldsymbol{\theta}}_+(k)$, where $k = |\vec{\boldsymbol{\theta}}_+(k)|$.




Figure 8: The splings of the diagram in Figure 3 with splaying data d as stated in Example 3.21. On the left is shown the bordered stabilisation with a pair of generators x (violet) and y (lime green); on the right is shown the splaying by the length one sequence of pairs of idempotents $((l_m, u_l))$ together with the corresponding splayed generator.

Recall there are corresponding generators (x', y') in the splings $\mathcal{D}(0, \iota(x))$ and $\mathcal{D}(0, \iota(y))$, given by the map $\cap_{\mathcal{D}}^{-1}$ (see the discussion following Definition 3.15). The nearest point map gives a generator \bar{y} corresponding with y' in the set of generators $\mathcal{G}(\alpha^m, \beta)$; it follows that the tuple $(x, \theta^{1,2}, \dots, \theta^{m-1,m}, \bar{y})$ constitutes a sequence of generators for the splayed or partially splayed diagram $\mathcal{D}(\varepsilon, \bar{l})$. We call this sequence the *splaying* of (x, y) .

Remark 3.20 For sufficiently small ε , for each j all of the points $e_j^{\infty,i}$ may all be assumed to be contained in a small disc neighbourhood N_j of e_j^{∞} which is disjoint from the curves $\beta(d)$.

Example 3.21 Figures 8, 9 and 10 show the bordered stabilisation of the diagram in Figure 3 by the data $d = (\lambda_{w,z}, \lambda_{z,w}, \mathbf{n}, \varepsilon, \bar{l})$, where $\lambda_{w,z}$ and $\lambda_{z,w}$ are as given in Figure 3, $\mathbf{n} = (0, 0)$, ε isn't pictured, and \bar{l} is stated below each figure. Also demonstrated are a series of generators, and their respective splings.




Figure 9: The splings of the diagram in Figure 3 with splaying data d as stated in Example 3.21. On the left is shown the bordered stabilisation with a pair of generators x (violet) and y (lime green); on the right is shown the splaying by the length two sequence of pairs of idempotents $((l_m, l_m), (l_l, l_m))$, together with the corresponding splayed generators. Here, the distinguished generator is shown in dark grey.




Figure 10: The splayings of the diagram in Figure 3 with splaying data d as stated in Example 3.21. On the left is shown the bordered stabilisation with β in blue and α in red, with a pair of generators x (violet) and y (lime green). On the right is shown the splaying by the length three sequence of pairs of idempotents $((\iota_m, \iota_l), (\iota_l, \iota_l), (\iota_l, \iota_m))$, where β is blue, α_1 is red, α_2 is orange and α_3 is yellow, together with the corresponding splayed generators. Here, the distinguished generators are shown in dark grey and light grey respectively.

Convention 3.22 From here on, we will label the attaching curves of a splayed or partially splayed diagram $\mathcal{D}(\varepsilon, \vec{\iota})$ or $\mathcal{D}^\partial(\varepsilon, \vec{\iota})$ by $\gamma^0, \dots, \gamma^m$, where $\gamma^0 = \beta$, and the curves $\gamma^1, \dots, \gamma^m$ are given by the splayed curves $\gamma^i = \alpha^i(d)$, or, for a partially splayed diagram, $\gamma^1 = \alpha = \alpha^a \cup \alpha^c$ and the subsequent curves are the closed curves $\gamma^i = \alpha^i(d)$.

3.1 Domains

Suppose that $\mathcal{D} = (\Sigma; \gamma^0, \dots, \gamma^m; \mathbf{b}; \mathbf{p})$ is a diagram. In the next chapter, we will be interested in maps

$$u: (S, \partial S) \rightarrow (\Sigma \times D_{m+1}, \gamma^0 \times e^0 \cup \dots \cup \gamma^m \times e^m)$$

where S is a Riemann surface with boundary and punctures upon its boundary, and D_{m+1} is an $(m+1)$ -gon—that is, a disc with $m + 1$ boundary punctures $v^{0,1}, \dots, v^{m,0}$ (see Section 4.1)—with sides e^0, \dots, e^m . Given any such u , one can postcompose with the projection maps $\pi_\Sigma: \Sigma \times D_{m+1} \rightarrow \Sigma$ and $\pi_D: \Sigma \times D_{m+1} \rightarrow D_{m+1}$ to yield corresponding maps u_Σ and u_D .

Near any boundary puncture q of S that is mapped to $v^{i,i+1}$ under u_D , the maps u which we study will be asymptotic to a point $I^{i,i+1}(q)$ of the form

$$I^{i,i+1}(q) := x^{i,i+1}(q) \times v^{i,i+1},$$

for some point $x^{i,i+1}(q) \in \gamma^i \cap \gamma^{i+1}$, such that the set

$$\mathbf{x}^{i,i+1} := \bigcup_{q \in u_D^{-1}(v^{i,i+1})} x^{i,i+1}(q)$$

constitutes a generator for $(\Sigma, \gamma^i, \gamma^{i+1})$ —that is, the sequence $\mathcal{G}(u) := (\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m,0})$ is a sequence of generators for \mathcal{D} .

We label the following objects:

$$C^i := \boldsymbol{\gamma}^i \times e^i, \quad I^{i,i+1} := \bigcup_{q \in u_{\mathcal{D}}^{-1}(v^{i,i+1})} I^{i,i+1}(q), \quad S^\partial := (\partial\Sigma \setminus \mathbf{b}) \times e^0.$$

It will follow from our assumptions on u that there is an associated map \bar{u} which represents some class in the relative homology group

$$H_2(\bar{\Sigma} \times I_s \times I_t, (C^0 \cup \dots \cup C^m) \cup (I^{0,1}(\mathbf{x}^{0,1}) \cup \dots \cup I^{m,0}(\mathbf{x}^{m,0})) \cup S^\partial).$$

We denote by $\pi_2(\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m,0})$ the set of all elements of the above group which map to the relative fundamental class of $I^{0,1}(\mathbf{x}^{0,1}) \cup \dots \cup I^{m,0}(\mathbf{x}^{m,0})$ under the composition of the boundary homomorphism and collapsing the remainder of the boundary.

A related notion is that of a *domain*. Namely, if \mathcal{D} is a diagram, then a *domain* is an element of the homology group

$$\text{Dom}(\mathcal{D}) := H_2(\bar{\Sigma}, \boldsymbol{\gamma}^1 \cup \dots \cup \boldsymbol{\gamma}^m \cup \partial\bar{\Sigma}).$$

This homology group is generated by the regions R_i of $\Sigma - \boldsymbol{\gamma}$, and thus any domain is given by a formal linear combination $\sum r_i R_i$.

A *periodic* domain is a domain with zero multiplicity (coefficient) at all regions R_i where $R_i \cap (z \cup \mathbf{b}) \neq \emptyset$ and with boundary equal to a linear combination of (full) curves and arcs from $\boldsymbol{\gamma}$. A *provincial periodic domain* is a domain which further has zero multiplicity at all regions R_i such that $R_i \cap \partial\Sigma \neq \emptyset$.

We say a diagram \mathcal{D} is *admissible* if every periodic domain in \mathcal{D} has both positive and negative coefficients; we say it is *provincially admissible* if every provincial periodic domain in \mathcal{D} has both positive and negative coefficients.

Any homology class B in the group $\pi_2(\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m,0})$ can be projected under π_Σ to an element of $\text{Dom}(\mathcal{D})$, which we call its *domain*. One can see that a homology class is uniquely determined by the corresponding domain; thus we often conflate the two notions.

Given a sequence of Reeb chords (ρ^1, \dots, ρ^k) of the torus algebra, there is an associated homology class $[\bar{\rho}] = [\rho^1] + [\rho^2] + \dots + [\rho^k]$, where $[\rho^i]$ is the image of the Reeb chord $\bar{\rho}^i$ in the relative homology group $H_1(S^1, \mathbf{a})$. If \mathcal{D} is a toroidal Heegaard multidiagram with h boundary components, each of the relative homology groups $H_1(\partial\Sigma_j, \mathbf{a}_j)$ are identified with $H_1(S^1, \mathbf{a})$.

Definition 3.23 Suppose that $B \in \pi_2(\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m,0})$ is a domain. We shall say that B is $\bar{\rho}$ -compatible for $\bar{\rho} = (\bar{\rho}_1, \dots, \bar{\rho}_h)$ if, for each j , we have that the image of B in $H_1(\partial\Sigma_j, \mathbf{a}_j)$ is equal to $[\bar{\rho}_j]$.

Convention 3.24 When \mathcal{D} is the splaying of a diagram \mathcal{D} , if $\mathbf{x}, \mathbf{y} \in \mathcal{G}(\mathcal{D})$ are generators, we shall often abbreviate the corresponding group

$$\pi_2(\mathbf{x}', \boldsymbol{\theta}_+^{1,2}, \dots, \boldsymbol{\theta}_+^{m-1,m}, \mathbf{y}')$$

as $\pi_2^{\boldsymbol{\theta}^+}(\mathbf{x}, \mathbf{y}) \subset \text{Dom}(\mathcal{D}(d))$.

We end this subsection with a brief discussion of generators and domains in east-compactified diagrams, which will be useful in Section 4.6.1. Recall that to any diagram \mathcal{D} is associated a corresponding east-compactified diagram $\mathcal{D}_{\bar{e}}$. Generators for \mathcal{D} have obvious counterparts in $\mathcal{D}_{\bar{e}}$, and we generally abuse notation and call these by the same label; any homology class B for \mathcal{D} has a corresponding east-compactified class $B_{\bar{e}}$ in the relative homology group

$$H_2(\Sigma_{\bar{e}} \times D_n, ((\gamma_{\bar{e}}^1 \times e^1) \cup \dots \cup (\gamma^m \times e^m)) \cup (I^{0,1} \times \dots \times I^{m-1,0})).$$

3.1.1 Splaying domains We would like to understand the relationship between domains in a bordered diagram $\mathcal{D} = (\Sigma; \alpha, \beta; \mathbf{b})$ and a corresponding (partial) splaying $\mathcal{D}^\partial(\varepsilon, \vec{l})$. First, we set up some notation for describing different qualities of sequences of Reeb chords.

Fix a stable bordered Heegaard diagram \mathcal{D} , generators $\mathbf{x}, \mathbf{y} \in \mathcal{G}(\mathcal{D})$, a set of sequences of Reeb chords $\vec{\rho}_{\Delta} = \vec{\rho}_{\Delta,1}, \dots, \vec{\rho}_{\Delta,h}$, and a sequence of idempotents \vec{l} . Consider the partially splayed diagram $\mathcal{D}^\partial(\varepsilon, \vec{l})$ and the splayed diagram $\mathcal{D}(\varepsilon, \vec{l})$. We shall associate to any $\vec{\rho}_{\Delta}$ -compatible domain $B \in \pi_2^{\theta+}(\mathbf{x}, \mathbf{y}) \subset \text{Dom}(\mathcal{D}^\partial(\varepsilon, \vec{l}))$, or any domain $B \in \pi_2^{\theta+}(\mathbf{x}, \mathbf{y}) \subset \text{Dom}(\mathcal{D}(\varepsilon, \vec{l}))$ a sequence of Reeb chords $\vec{\rho}(B)$.

Convention 3.25 Recall from the beginning of the section that any homology class in $\text{Dom}(\mathcal{D})$ corresponds with a homology class in $\text{Dom}(\mathcal{D}_{\bar{e}})$ — we make this identification for the rest of this section, and consider domains in diagrams where the underlying surface is closed.

Consider a small neighbourhood N_j of $e_j^\infty \in \Sigma_{\bar{e}}$, which is disjoint from β and such that for every i , $e_j^{\infty,i}$ is contained in N_j , as in Remark 3.20. Suppose that $B \in \pi_2^{\theta+}(\mathbf{x}, \mathbf{y})$ is a domain, and that u is a map from a surface with boundary and punctures upon the boundary to $\Sigma \times D_m$ or $\Sigma_{\bar{e}} \times D_m$, with homology class B . If S_j denotes the restricted source $(\pi_\Sigma \circ u)^{-1}(N_j)$, then the map $u'_j = u|_{S_j}$ takes the form of a map

$$u'_j : (S_j, \partial S_j) \rightarrow (N_j \times D_{m+1}, \partial N_j \cup \gamma_j^1 \times e^1 \cup \dots \cup \gamma_j^m \times e^m).$$

Maps of this form are in bijection with maps

$$u_j : (S_j, \partial S_j) \rightarrow (N_j \times D_{m+1}, \partial N_j \cup \delta_j^1 \times e^1 \cup \dots \cup \delta_j^{m'} \times e^{m'}),$$

where $\delta_j^1, \dots, \delta_j^{m'}$ is the subsequence of $\gamma_j^1, \dots, \gamma_j^m$ formed by deleting every arc γ_j^i such that γ_j^i approximates γ_j^{i-1} . On the level of domains, the map u_j has an associated domain

$$B_j \in H_2(N_j, \partial N_j \cup (\delta_j^1 \cup \dots \cup \delta_j^{m'}))$$

which can be described by choosing a point x_R for every region of $N_j - (\gamma_j^1 \cup \dots \cup \gamma_j^m)$ not contained in the approximation regions between any two adjacent curves $\gamma_j^i, \gamma_j^{i+1}$ which approximate one another: B_j is the domain given by requiring that the multiplicity of B_j at each x_R is the same as the multiplicity of $B|_{N_j}$ at each x_R .

In what follows, it will help us to have an understanding of the homology groups $H_2(N_j, \partial N_j \cup \delta_j^i)$ and $H_2(N_j, \partial N_j \cup \delta_j^i \cup \delta_j^{i+1})$.

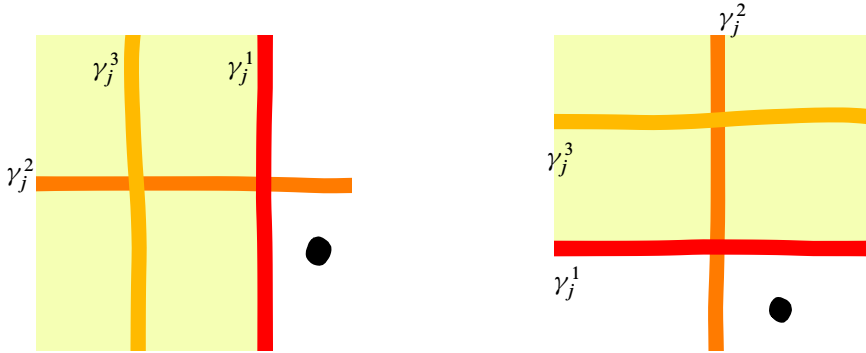


Figure 11: A series of examples of various domains (shaded green) in the neighbourhood N_j of e_j^∞ . On the left is an example of $D^1(\rho_{12})$, and on the right is $D^1(\rho_{23})$.

The pair (N_j, δ_j^i) is modelled upon either the pair $(\mathbb{D}, \mathfrak{R})$ or $(\mathbb{D}, \mathfrak{S})$ depending upon whether δ_j^i is meridional or longitudinal. In these cases, the groups $H_2(\mathbb{D}, \mathfrak{R})$ and $H_2(\mathbb{D}, \mathfrak{S})$ are generated by the half-discs $\mathbb{D} - \mathfrak{R}$ and $\mathbb{D} - \mathfrak{S}$ respectively, only one of which has multiplicity zero at the basepoint p_j in each case. We label these homology classes $D(\rho_{12})$ and $D(\rho_{23})$ respectively, and the corresponding domain under the identification with (N_j, δ_j^i) by $D_j^i(\rho_{12})$ and $D_j^i(\rho_{23})$ respectively (see Figure 11). We generalise this to length n_i repeated sequences of these chords, by setting $D_j^i(\rho_{12}, \dots, \rho_{12}) := n_i \cdot D_j^i(\rho_{12})$, and similar for ρ_{23} — in both instances, we allow the degenerate length 0 sequence with associated domain equal to zero. If δ_j^i and δ_j^{i+1} are not small approximations of one another, there are two models for the ordered triple $(N_j, \delta_j^i, \delta_j^{i+1})$ depending upon the configuration of the curves:

- $(\mathbb{D}, \mathfrak{S}, \mathfrak{R})$ if δ_j^i is meridional and δ_j^{i+1} longitudinal;
- $(\mathbb{D}, \mathfrak{R}, \mathfrak{S})$ if δ_j^i is longitudinal and δ_j^{i+1} is meridional.

$H_2(\mathbb{D}, \mathfrak{R} \cup \mathfrak{S})$ is generated by four quadrants $B(\rho_0)$, $B(\rho_1)$, $B(\rho_2)$ and $B(\rho_3)$, where $B(\rho_0)$ is the quadrant of $(\mathbb{D}, \mathfrak{R} \cup \mathfrak{S})$ containing the point p_j , and the others are given in the order they are encountered clockwise from $B(\rho_0)$ (see Figure 12). There is a natural identification $\varphi: H_1(S^1, \mathbf{a}_j) \rightarrow H_2(\mathbb{D}, \mathfrak{R} \cup \mathfrak{S})$, and we denote $B(\rho) := \varphi(\rho)$. We denote by $B_j^i(\rho)$ the corresponding domain in $H_2(N_j, \partial N_j \cup (\delta_j^i \cup \delta_j^{i+1}))$. (Please see Figure 13 for some examples.)




Figure 12: The four quadrants of \mathbb{D} .




Figure 13: A series of examples of various domains (shaded green) in the neighbourhood N_j of e_j^∞ . On the left is an example of $B^1(\rho_1)$, and on the right is $B^2(\rho_2)$. We note how the sum of these domains has an obtuse corner at the intersection of γ_j^1 and γ_j^2 .

We also extend this assignment to the case where \mathcal{D} is partially splayed and $\boldsymbol{\gamma}^1 \cap N_j = (\gamma_j^{1,m} \cup \gamma_j^{1,l}) \cap N_j$. Noting that $(N_j, \boldsymbol{\gamma}^1 \cap N_j)$ may be identified with $(\mathbb{D}, \mathfrak{S} \cup \mathfrak{R})$, we may associate to a sequence of Reeb chords $\vec{\rho}$ the corresponding domain $B(\vec{\rho})$ identified with $\varphi(\vec{\rho})$, in the homology group

$$H_2(N_j, \partial N_j \cup (\boldsymbol{\gamma}^1 \cap N_j)).$$

Recall that two sequences of Reeb multichords, $\vec{\rho}$ and $\vec{\rho}'$, are 123-equivalent if their corresponding 123-composed sequences agree. We introduce a related notion for domains:

Definition 3.26 We shall say that a sequence of domains

$$S = (D_j^1, B_j^1, D_j^2, B_j^2, \dots, D_j^k, B_j^k, D_j^{k+1}),$$

is a *sequence of Reeb domains* if each D_j^i is equal to $D_j^i(\vec{\rho}^i)$ for some repeated sequence $\vec{\rho}^i := (\rho^i, \rho^i, \dots, \rho^i)$ of a nonjumping chord $\rho^i \in \mathcal{T}$, and each B_j^i is equal to $B_j^i(\underline{\rho}^i)$ for some jumping chord $\underline{\rho}^i \in \mathcal{T}$. For such a sequence, denote by $\vec{\rho}(S)$ the sequence of Reeb chords

$$(\vec{\rho}^1 \cup \underline{\rho}^1 \cup \vec{\rho}^2 \cup \dots \cup \underline{\rho}^k \cup \vec{\rho}^k).$$

If S' is another sequence of Reeb domains, we shall say that S and S' are *composable-equivalent* if the homology classes in $H_2(N_j, \partial N_j \cup \delta_j^1 \cup \dots \cup \delta_j^{m'})$ formed by summing S and S' are the same.

As the name suggests, this relation is closely related to the composable-equivalence of corresponding Reeb chords:

Proposition 3.27 *Suppose S and T are sequences of Reeb domains. Then they are composable-equivalent if and only if the sequences $\vec{\rho}(S)$ and $\vec{\rho}(T)$ are 123-equivalent.*

Proof Suppose that S is a sequence of Reeb domains and consider the corresponding sequence of Reeb chords $\vec{\rho}(S)$. One can see there is a unique sequence $U(S)$ of Reeb domains that is composable-equivalent to S , and such that $\vec{\rho}(U(S))$ is the 123-reduction of $\vec{\rho}(S)$.

Indeed, $\vec{\rho}(S)$ is not 123-reduced if and only if there are two adjacent chords $\rho_k, \rho_{k+1} \in \vec{\rho}(S)$ that are either equal to the torus algebra chords ρ_1, ρ_{23} or ρ_{12}, ρ_3 . In the first case, if $B_j^i = B^i(\rho_1)$ and $D_j^{i+1} = n_{i+1} B^{i+1}(\rho_{23})$ are the corresponding domains in S then one can see that the sequence S' given by replacing B_j^i by $B^i(\rho_{123}) = B^i(\rho_1) + B^i(\rho_{23})$ and D_j^{i+1} by $(n_{i+1} - 1) B^{i+1}(\rho_{23})$ is composable-equivalent to S , and the corresponding $\vec{\rho}(S')$ differs from $\vec{\rho}(S)$ by replacing ρ_k, ρ_{k+1} with $\rho_k \cdot \rho_{k+1}$. Similar holds when $(\rho_k, \rho_{k+1}) = (\rho_{12}, \rho_3)$. Repeating this process for any 123-composable chords yields the $U(S)$ required.

It follows that S and T are composable-equivalent if and only if $U(S) = U(T)$, which holds if and only if the 123-reductions of $\vec{\rho}(S)$ and $\vec{\rho}(T)$ agree, and so the result follows. \square

Proposition 3.28 *Suppose that $B \in \pi_2^{\theta+}(\mathbf{x}, \mathbf{y})$ is a domain which carries a holomorphic representative, and is $\vec{\rho}_{\Delta}$ -compatible. Then for every $j = 1, \dots, h$, the domain $B_j - B(\vec{\rho}_{\Delta, j})$ may be written uniquely, up to composable-equivalence, as a sequence of Reeb domains.*

Proof We construct an explicit sequence of Reeb domains S inductively. If there is only one set of γ -curves, B_j is an element of the group $H_2(N_j, \partial N_j \cup \delta_j^1)$, which, as discussed, is equal to $n_1 B_j^1(\rho)$ for some nonjumping $\rho \in \mathcal{T}$.

Consider a small disc neighbourhood M_j^1 of $e_j^{\infty, 1}$, disjoint from all other $e_j^{\infty, i}$ and e_j^{∞} . The restriction of the domain of B to M_j^1 is modelled upon $n_1 B(\rho_1) + n_2 B(\rho_2) + n_3 B(\rho_3)$ for some $n_1, n_2, n_3 \geq 0$. There are three cases:

- If γ_j^2 is meridional, write $D_j^1 := n_3 B_j^1(\rho_{23})$ and $B_j^1 := B_j^1(\rho_2)$.
- If γ_j^2 is longitudinal and $n_1 > n_2$, write $D_j^1 := n_2 B_j^1(\rho_{12})$ and $B_j^1 := B_j^1(\rho_1)$.
- If γ_j^2 is longitudinal and $n_1 \leq n_2$, write $D_j^1 := n_1 B_j^1(\rho_{12})$ and $B_j^1 := B_j^1(\rho_3)$.

It follows that the domain $B'_j = B_j - (D_j^1 + B_j^1)$ has the same multiplicity upon either side of γ_j^2 . As such, B'_j is an element of $H_2(N^j, \partial N^j \cup (\delta_j^3 \cup \dots \cup \delta_j^k))$, and the result follows by induction. \square

With this in mind, we may make a definition:

Definition 3.29 Let $B \in \pi_2^{\theta+}(\mathbf{x}, \mathbf{y})$ be a domain which carries a holomorphic representative u , and let S_j be a decomposition of each B_j into Reeb domains given by the construction in Proposition 3.28. We denote by $[\vec{\rho}(B)_j]$ the 123-equivalence class of Reeb multichords defined by $[\vec{\rho}(B)_j] := [\vec{\rho}(S_j)]$. We often collect these together, writing $[\vec{\rho}(B)] = ([\vec{\rho}(B)_1], \dots, [\vec{\rho}(B)_h])$.

By construction, $[\vec{\rho}(B)_1], \dots, [\vec{\rho}(B)_h]$ satisfies that for any set of representatives $\vec{\rho}(B)_1, \dots, \vec{\rho}(B)_h$ of each corresponding 123-equivalence class, the sequence of idempotents $\vec{l}_j := \iota_j^1, \dots, \iota_j^{m+1}$ has

$$|\{\iota_j^i \in \vec{l}_j \mid \iota_j^i \neq \iota_j^{i-1}\}| = |\text{jump}(\vec{\rho}_j)|,$$

as in Definition 2.9. As such:

Definition 3.30 If $B \in \text{Dom}(\mathcal{D}(\varepsilon, \vec{\iota}))$ is a domain, then there is a well-defined splicing of $\vec{\rho}(B)_1, \dots, \vec{\rho}(B)_h$ given by $\vec{\sigma}(B) := \vec{\sigma}(\vec{\rho}, \vec{\iota})$.

We shall later be interested in domains which look similar between splayed diagrams. A useful construction to bear in mind (cf [13, Definition 3.34]) is:

Definition 3.31 Fix an h -splaying sequence of idempotents $\vec{\iota}$, and consider an additional tuple of idempotents $\iota = (\iota_1, \dots, \iota_h)$. Let $\vec{\iota}' = (\iota, \vec{\iota})$ and consider the partially splayed diagrams $\mathcal{D}^\partial(\varepsilon, \vec{\iota})$ and $\mathcal{D}^\partial(\varepsilon, \vec{\iota}')$; we label the curves in each by $\gamma^0, \dots, \gamma^m$ and $\delta^0, \dots, \delta^{m+1}$ respectively. For every $i > 1$, the curves γ^i and δ^{i+1} are small approximations of one another (with associated approximation region R^i , and the curves γ^2 are small approximations of some subset of the curves γ_ε^1 in the east-compactification $\mathcal{D}^\partial(\varepsilon, \vec{\iota})|_{\bar{\varepsilon}}$ (with associated approximation region R^1).

For every curve $\gamma_j^i \in \gamma^i$, choose a pair of basepoints on either side of γ_j^i such that the basepoints are disjoint from each of the approximation regions R^i . For any

$$B' \in \pi_2(x, \theta^{1,2}, \dots, \theta^{m-1,m}, y) \subset \text{Dom}(\mathcal{D}^\partial(\varepsilon, \vec{\iota}')),$$

we let $\phi(B')$ denote the unique domain

$$B \in \pi_2(x, \theta^{2,3}, \dots, \theta^{m-1,m}, y) \subset \text{Dom}(\mathcal{D}^\partial(\varepsilon, \vec{\iota}))$$

with the same multiplicities at each of these basepoints as B' (here we abuse notation and write the same symbols for generators which are identified by the nearest-point map). We sometimes call $\phi(B')$ the domain obtained from B' by *taking multiplicities away from the approximation region*, and shall say that B' *splays* B if $\phi(B') = B$.

4 Moduli spaces of holomorphic curves and polygons

In this section, we define a series of moduli spaces which are required in the definition of the invariants $\text{Poly}(\mathcal{D})$; we also examine the structure of these moduli spaces and establish results which restrict the possible elements which may occur in the boundaries of their compactifications, in order to show that $\text{Poly}(\mathcal{D})$ satisfies the \mathcal{A}_∞ relations, and also to understand the behaviour of $\text{Poly}(\mathcal{D})$ under a neck-stretching argument which we use in Section 5 to identify $\text{Poly}(\mathcal{D})$ with the bordered-sutured Heegaard Floer invariants of the bordered Heegaard diagram.

4.1 Complex structures on polygons

Before we describe the moduli spaces that the results of the paper will be based upon, we need to fix notation for moduli spaces of polygons. Most of this section follows [24, Chapter 9].

Let D_m denote a disc, with m punctures upon the boundary (we often view these as marked vertices). Choose a distinguished edge $e^0 \subset \partial D_m$, and label the remaining edges e^1, \dots, e^{m-1} anticlockwise

from e^0 ; also label the vertices by $p^{i,i+1}$, so that the vertex $p^{i,i+1}$ lies between e^i and e^{i+1} (here, superscripts are taken modulo m).

Let $\text{Conf}(D_m)$ denote the moduli space of equivalence classes of positively oriented complex structures upon the disc D_m , up to biholomorphism that preserves the labelling of the edges (and vertices). The moduli space $\text{Conf}(D_m)$ has a natural compactification by trees of holomorphic polygons which we summarise now.

Definition 4.1 An m -leaved tree is a properly embedded planar tree $T \subset \mathbb{R}^2$ with $m + 1$ semi-infinite edges, one of which is distinguished and called the *root* of T , and the others we call the *leaves*. We denote by $V(T)$ the set of vertices of T , and Ed^{int} the set of interior (finite) edges of T .

A *flag* in T is an oriented pair of a vertex and adjacent edge, (V, e) , oriented so that the edge e begins at V . We orient T so that every path from the root to a leaf is positively oriented. To every vertex v of T we can associate a series of flags, given by the pair (V, e) and oriented accordingly. One of these flags is *negative* with respect to V , in the sense that the orientation on e induced by T differs from that stemming from the flag (ie the edge points “into” V), and the others are *positive*. It follows that we can order the flags originating at any given vertex V by saying $f_0(V)$ is the unique negative flag at V , and ordering the remainder $f_1(V), \dots, f_{|V|-1}(V)$ so that they are encountered in order from $f_0(V)$, anticlockwise. We call this the *root-ordering* of the flags.

Suppose we are given an m -leaved tree T , and a collection of Riemann surfaces $S = \{S(V)\}_{V \in V(T)}$ with boundary and punctures on the boundary, such that the boundary punctures of $S(V)$ are in one-to-one correspondence with the flags $f(V)$. To each interior edge e in T connecting V_1 and V_2 there are thus two related boundary punctures: one in $S(V_1)$ corresponding to the flag (V_1, e) , and one in $S(V_2)$ corresponding to the flag (V_2, e) . There is an operation upon the surface S given by deleting a small open neighbourhood of each of these punctures, and identifying the two new resultant boundary arcs to yield a new surface. Performing this operation iteratively over all interior edges gives a connected (topological) surface, $S^\cup(T)$.

For disks, we require more predictable behaviour near the punctures.

Definition 4.2 Let D be a representative for some equivalence class in $\text{Conf}(D_n)$. A choice of *strip-like ends* for D is a choice of open neighbourhood $N^{i,i+1}$ of each boundary puncture $p^{i,i+1}$ — such that $N^{i,i+1} \cap N^{j,j+1} = \emptyset$ for $i \neq j$ — together with a choice of a biholomorphism between each $N^{i,i+1}$ and the strip $\mathbb{R}_+ \times [0, 1]$ for every $i \neq 0$, and between $N^{0,1}$ and the strip $\mathbb{R}_- \times [0, 1]$.

Suppose T is a d -leaved tree, and that we are given a collection of Riemann surfaces $S = \{S(V)\}_{V \in V(T)}$, where each $S(V)$ represents some equivalence class in $\text{Conf}(D_{|V|})$ and is endowed with a choice of strip-like ends (ie a tree of holomorphic discs with strip-like ends). We shall say that a map $\rho : \text{Ed}^{\text{int}}(T) \rightarrow (0, \infty)$ is a *length assignment* for T — thus giving T the structure of a metric graph. We often omit ρ from the notation.

We refine the construction of $S^\cup(T)$. There is a natural correspondence between the flags of each vertex V and the punctures of $D_{|V|}$, between the puncture $p^{i,i+1} \in D_{|V|}$ and the flag $f_i(V)$. When we delete neighbourhoods of punctures corresponding with a negative flag (V_1, e) and positive flag (V_2, e) , we take these to be the open neighbourhoods identified by the choices of strip-like ends with $(-\infty, -\log(\rho(e)) \times [0, 1] \subset \mathbb{R}_- \times [0, 1]$ and $(\log(\rho(e), \infty) \times [0, 1] \subset \mathbb{R}_+ \times [0, 1]$ respectively.

This construction yields an element D_T in some equivalence class of $\text{Conf}(D_m)$ — an m -pointed disc with Riemann surface structure, with vertices labelled $p^{0,1}, \dots, p^{m-1,0}$ starting from the one corresponding with the root of T (and edges labelled accordingly). Moreover, there is a natural choice of strip-like ends upon D_T induced by those upon the punctures corresponding with the semi-infinite flags of T .

Allowing edge lengths to degenerate and take the value 0 allows us to describe the compactification of the moduli space $\text{Conf}(D_m)$. If a given edge length is zero, we omit this edge from the cutting and pasting procedure described above — effectively cutting the edge into two semi-infinite edges, and thus the tree into more components, upon which we perform the gluing process individually. The resulting collection of discs is described by a tree \bar{T} , which is given by collapsing all internal edges of nonzero length, and assigning to each vertex \bar{v} of \bar{T} a $|\bar{v}|$ -pointed disc in $\text{Conf}(D_{|\bar{v}|})$.

There is a canonical way to label the edges of \bar{T} in a way compatible with those of T . We first label the semi-infinite flags of \bar{T} by $(0, 1), (1, 2), \dots, (m, 0)$ anticlockwise from the root of \bar{T} . The task now is to label the interior flags of \bar{T} by labels compatible with these; in such a way that the edges of the corresponding disks D_V can be labelled by $0, 1, \dots, m$, so that an edge between punctures (i, k) and (k, j) is labelled by k . The way to do this is to first label the vertices which just have semi-infinite flags and one in-coming flag — in this case, the semi-infinite flags are already labelled anticlockwise by some contiguous subsequence $(j, j+1), (j+1, j+2), \dots, (j+k-1, j+k)$, so the root flag is forced to be labelled by $(j+k, j)$. We can then delete this vertex from \bar{T} and replace the root flag with a semi-infinite flag labelled by $(j+k, j)$, and repeat the process until all edges of \bar{T} are labelled.

Definition 4.3 Suppose that \mathcal{D} is a diagram with m sets of attaching curves $\gamma^1, \dots, \gamma^m$, and let T be a (possibly degenerate) m -leaved metric tree. Then a *generator assignment* for T is an assignment of a generator $x^{i,j} \in \mathcal{G}(\gamma^i, \gamma^j)$ to each edge of \bar{T} labelled by (i, j) . If \vec{x} is a generator assignment for T , we refer to the sequence of generators corresponding to the edges adjacent to a vertex $V \in \bar{T}$, ordered by root-ordering, by $\vec{x}(V)$. Note that a generator assignment for T has a well-defined overall sequence of generators $\vec{x}(T) \in \mathcal{G}(\gamma^1, \dots, \gamma^m)$, given by the set of generators labelling the exterior edges of T .

Given a generator assignment \vec{x} , a *family of compatible domains* is a set $\{B_V\}_{V \in V(\bar{T})}$ of domains such that B_V is a member of the group $\pi_2(\vec{x}(V))$. The juxtaposition $B(T) = \sum_{V \in V(\bar{T})} B_V$ is, accordingly, a member of the group $\pi_2(\vec{x}(T))$.

We would like to work with bona fide representatives of each equivalence class for $\text{Conf}(D_m)$, with strip-like ends. Once these are chosen, the compactification of the moduli space $\text{Conf}(D_m)$ is then described by

the set $\bigcup_{T,\rho} D_T$, indexed over all d -leaved trees and all (possibly degenerate) length assignments ρ . As edge lengths in this space degenerate to zero, the surface D_T degenerates to a disjoint union of surfaces $\{D_V\}_{V \in V(\bar{T})}$.

To make this make sense, one must choose representatives D and choices of strip-like ends for every $d \leq m$ and each equivalence class of $\text{Conf}(D_d)$, so that at these boundary strata the strip-like ends chosen agree with those induced by those chosen upon D_d . This is possible; see, for instance, [24, Section 9g].

From this point, we fix such a set of choices of representatives for equivalence classes with strip-like ends such that this works (ie we have consistency on the boundary strata), and abuse notation by identifying elements of $\text{Conf}(D_m)$ with these representatives. We also fix a choice of symplectic forms $\{\omega_j\}_{j \in \text{Conf}(D_m)}$ where ω_j is compatible with j , so that if $(j_n)_{n=1}^\infty \in \text{Conf}(D_m)$ is a sequence of complex structures converging to the collection $\{j_{\bar{v}}\}_{\bar{v} \in V(\bar{T})} \in \partial \overline{\text{Conf}(D_m)}$, then $\omega_{j_n} \rightarrow \{\omega_{j_{\bar{v}}}\}_{\bar{v} \in V(\bar{T})}$.

We will also have need to fix identifications η^i of a small neighbourhood $N(e^i)$ of each edge e^i in the polygons $(D_m, j) \in \text{Conf}(D_m)$ with $\mathbb{R} \times [0, \varepsilon)$, so that on each of these $N(e^i)$ the symplectic form ω_j agrees with the pullback of the standard symplectic form on \mathbb{R}^2 . To make this make sense, we require that these identifications are continuous, consistent with the choice of strip-like ends, and consistent across boundary strata of $\text{Conf}(D_m)$ (see [13, Section 3.3; 24, Section 10e], for instance).

We will later need to single out some vertices of our trees:

Definition 4.4 Let T be a tree. Then the set of *comb vertices* of \bar{T} is the set of vertices of \bar{T} visited by the path in \bar{T} connecting the exterior edge labelled by $(0, 1)$ (ie the root of \bar{T}) with the exterior edge labelled by $(1, 2)$. The set of *spinal vertices* of \bar{T} is the set of vertices of \bar{T} which have a flag f_V labelled by (i, j) where one of i or j is equal to 0. The set of *easterly vertices* of \bar{T} is the set of nonspinal vertices of \bar{T} .

If V is spinal, we associate to it its *storied tree* $T' \subset T$, given by the maximal connected subtree of T which contains V and no other spinal vertices (so, it contains V and a series of easterly vertices).

4.2 Sources

In this section, we define what shall be the sources of the various types of holomorphic curves considered in the remainder of this chapter. We first fix positive integers m and h , and a disc $D_m \in \text{Conf}(D_m)$.

Definition 4.5 A *topological (h, m) -source* is a surface S , with boundary and punctures $\text{punc}(S)$ together with a labelling $\lambda: \text{punc}(S) \rightarrow \text{punc}(D_m) \cup \mathcal{T}$ of the boundary punctures of S by either a puncture in ∂D_m or a Reeb chord in \mathcal{T} . We will write $E(S)$ for the subset of “east” punctures $\lambda^{-1}(\mathcal{T})$, and also require that $E(S)$ is partitioned into h subsets $E_1(S), \dots, E_h(S)$.

An *(h, m) -source* is a topological (h, m) -source together with a choice of complex structure j_S which gives S the structure of a smooth (not nodal) Riemann surface.

It will be useful to have some more adjectives to describe various types of source.

Definition 4.6 We give names to the following specific varieties of (h, m) source:

- A (topological) *bordered* source is a (topological) $(h, 2)$ -source.
- A (topological) *polygonal m -source* is a (topological) $(0, m)$ -source such that $E(S) = \emptyset$.
- A (topological) *tooth* source is a (topological) (h, m) -source together with a further labelling of the punctures $E(S)$ by either “ e ” or “ w ”. We will usually denote these by T , and shall abuse notation to write $E_j(S)$ (resp. $W_j(S)$) for the set of punctures labelled by “ e ” (resp. “ w ”) in part $E_j(S)$.
- In the degenerate case where $m = 1$, ie $D_m = \mathbb{H}$ is the upper-half plane, with one boundary puncture at ∞ , we call a (topological) $(h, 1)$ source an *anchor source*.

Later, we will need to fix the behaviour of holomorphic curves near a certain subset of noneast punctures. The relevant kind of source is:

Definition 4.7 A (topological) *splayed (h, m) -source* is a (topological) (h, m) source together with a distinguished set $\rho_{123}(S) \subset \text{punc}(S) - E(S)$, and a partition Q of $\rho_{123}(S)$ into three sets, $Q = (Q_{123}, Q_{1,23}, Q_{12,3})$.

If we are given a metric tree T , then we can associate to every vertex V of the corresponding tree \bar{T} a (topological) $(h, |V|)$ source S_V . If, for each V and V' connected by an edge of \bar{T} , we are given a bijection between the punctures of S_V and $S_{V'}$ corresponding with this edge, we can associate to T a corresponding topological *preglued (h, m) source* S_T , where m is the number of noninterior edges of \bar{T} . This is formed by deleting small semi-infinite strip neighbourhoods of punctures of S corresponding with an edge of T , and gluing the resultant ends together according to the chosen bijection. There are natural decorations of S_T induced by those upon S_V : each east puncture in a source S_V is preserved by the gluing process, and thus has a corresponding puncture in S_T ; we label the resultant punctures by the corresponding element of \mathcal{T} , and partition them according to $E_j(S_T) = \bigcup_{V \in V(T)} E_j(S_V)$. If some sources S_V are also endowed with the extra data of a distinguished set $\rho_{123}(S_V)$ and partition Q_V as in [Definition 4.7](#), the resultant source has a natural distinguished set $\rho_{123}(S_T)$ given by the union of all of the $\rho_{123}(S_V)$ after the gluing procedure — where we omit any elements of $\rho_{123}(S_V)$ which are glued along, and regard a (topological) (h, m) source as a splayed source with empty distinguished set.

Remark 4.8 The definition of the source S_T differs from the definition of the surface $S^\cup(T)$ defined in [Section 4.1](#) in that (in addition to the extra data of labellings, etc) each of the sources $S(V)$ can have multiple punctures corresponding to the same edge of T . It is also worth noting that we have not specified any complex structure on the surface underlying S_T ; the construction gives a purely topological source.

If \bar{T} only has two vertices, V_1 and V_2 , then we sometimes denote the corresponding source by $S_{V_1} \natural S_{V_2}$. A similar construction can also be made for tooth sources: if S is a bordered source and T is a tooth source, with a bijection between $E(S)$ and $W(T)$, we can form a glued source $S \natural T$ given by deleting small neighbourhoods of $E(S)$ and $W(T)$ and identifying the resultant “stubs” via the bijection.

4.3 Moduli spaces of holomorphic curves

We will now define a series of moduli spaces of holomorphic curves. Suppose that $\mathcal{D} = (\Sigma; \boldsymbol{\gamma}; \mathbf{b}; \mathbf{p})$ is a diagram, where $\boldsymbol{\gamma} = \{\boldsymbol{\gamma}^1, \dots, \boldsymbol{\gamma}^m\}$. We shall be concerned with holomorphic maps from an m -source into the four-manifold $X_{\mathcal{D},m} := \Sigma \times D_m$. We assume that Σ is endowed with a symplectic form ω_Σ , and choose a complex structure j_Σ upon Σ which is compatible with ω_Σ .

We also assume that j_Σ is cylindrical near the boundary, in that we fix small neighbourhoods $N_j \subset \Sigma_{\bar{e}}$ of each e_j^∞ , together with a symplectic identification of each $U_j = N_j - e_j^\infty$ with $S^1 \times (0, 1)$. Under this identification we assume that the complex structures $(j_\Sigma)|_{U_j}$ and arcs $\boldsymbol{\gamma} \cap U_j$ are also invariant in the $(0, 1)$ -direction. It follows from this assumption that the complex structure j_Σ can be extended to a complex structure $j_{\bar{e}}$ upon $\Sigma_{\bar{e}}$.

We denote by

$$\pi_\Sigma: \Sigma \times D_m \rightarrow \Sigma \quad \text{and} \quad \pi_{\mathcal{D}}: \Sigma \times D_m \rightarrow D_m$$

the relevant projection maps.

Definition 4.9 An *admissible* family of almost-complex structures for \mathcal{D} is a collection $\mathbf{J} = \bigcup_{m \geq 1} \mathbf{J}_m$ of families of almost-complex structures $\mathbf{J}_m = \{J_j\}_{j \in \text{Conf}(D_m)}$ upon $X_{\mathcal{D},m}$ satisfying:

- (1) The projection map $\pi_{\mathcal{D}}: \Sigma \times D_m \rightarrow D_m$ is (J_j, j) -holomorphic for each j .
- (2) The fibres of the maps $\pi_{\mathcal{D}}$ are J_j -holomorphic for every j .
- (3) Each J_j is adjusted to the split symplectic form $\omega_\Sigma \oplus \omega_j$ on $\Sigma \times D_m$.
- (4) Near the punctures $e_1^\infty \times D_m, \dots, e_h^\infty \times D_m$, the almost complex structures J_j are split.

For fixed m , we shall study moduli spaces of maps into $X_{\mathcal{D},m}$ which are $J \in \mathbf{J}_m$ -holomorphic. We will require some form of consistency for the family \mathbf{J} , so that for any family of elements of $\text{Conf}(D_m)$ converging to some point in $\overline{\partial \text{Conf}(D_m)}$, the corresponding family of almost-complex structures upon $X_{\mathcal{D},m}$ converges to the correct tree of almost-complex structures. Namely:

Definition 4.10 Suppose \mathbf{J} is a family of admissible almost-complex structures. Recall that for any m and sequence of m -leaved metric trees $(T_n)_{n \geq 0}$ converging to some degenerate tree T , the chosen family of equivalence classes $\text{Conf}(D_m)$ satisfies that the corresponding almost-complex structures $j_n := j_{T_n} \in \text{Conf}(D_m)$ converge to a tree of complex structures $\{j_V\}_{V \in \mathcal{V}(T)}$, where each $j_V \in \text{Conf}(D_{|V|})$.

We say that \mathbf{J} is *consistent* if the family J_{j_n} converges to the corresponding tree $\{J_{j_V}\}_{V \in \mathcal{V}(T)}$ — where each J_{j_V} is an almost-complex structure on $X_{\mathcal{D},|V|}$.

Given a diagram \mathcal{D} , there are corresponding Lagrangian “cylinders” inside the four-manifold $X_{\mathcal{D}}$. At their most basic, these are given by the sets of cylinders and half-cylinders $C^i := \boldsymbol{\gamma}^i \times e_i \subset \Sigma \times e_i$. In Section 4.7, we will perturb these cylinders slightly, replacing C^i with cylinders $C^i(\varepsilon)$ associated to some Hamiltonian perturbation of Σ . In this case, the corresponding definition is to be made by replacing C^i with $C^i(\varepsilon)$ below.

With these definitions in mind, we specify the types of holomorphic curves we shall be interested in.

Definition 4.11 Let $\mathcal{D} = (\Sigma; \boldsymbol{\gamma}; \mathbf{b}; \mathbf{p})$ be a diagram with attaching curves $\boldsymbol{\gamma} = \boldsymbol{\gamma}^1, \dots, \boldsymbol{\gamma}^m$, and write $\boldsymbol{\gamma}^i = \boldsymbol{\gamma}^{i,a} \cup \boldsymbol{\gamma}^{i,c}$ for the partition of each $\boldsymbol{\gamma}^i$ into arcs and closed curves, respectively. Suppose also that we are given a partition of the points $\mathbf{p} = \mathbf{w} \cup \mathbf{z}$, where either part may be empty.

Fix a consistent, admissible family \mathbf{J} of almost complex structures for \mathcal{D} , and a source S , endowed with a complex structure j_S . We consider pairs (u, j) , where $j \in \text{Conf}(D_m)$ and

$$u: (S, \partial S) \rightarrow (X_{\mathcal{D},m}, C^1 \cup \dots \cup C^m)$$

is a map that satisfies the following conditions:

- (M-1) u is (j_S, J_j) -holomorphic.
- (M-2) $u: S \rightarrow \Sigma \times D_n$ is proper.
- (M-3) u has finite energy.
- (M-4) The image of u is disjoint from the sets $\mathbf{b} \times D_m$ and $\mathbf{z} \times D_m$.
- (M-5) u extends to a proper map $u_{\bar{e}}: S_{\bar{e}} \rightarrow \Sigma_{\bar{e}} \times D_n$. (This is only meaningful when $\Sigma_{\mathcal{D}}$ has boundary.)
- (M-6) $\pi_{\mathbb{D}} \circ u_{\bar{e}}$ is a g -fold branched cover.
- (M-7) u is an embedding.
- (M-8) At each puncture q of S labelled by $p^{i,i+1}$, $\lim_{z \rightarrow q} (\pi_{\mathcal{D}} \circ u)(z) = p^{i,i+1}$.

When $\Sigma_{\mathcal{D}}$ has boundary, we require three further conditions:

- (BOR-1) At each puncture $q \in E(S)_j$, $\pi_{\Sigma} \circ u$ is asymptotic to the Reeb chord labelling q upon the boundary component $\partial \Sigma_j$.
- (BOR-2) For every $i = 1, \dots, m$ and $j = 1, \dots, g+h-1$, and every $t \in e^i$, we require that $u^{-1}(\boldsymbol{\gamma}_j^{i,c} \times \{t\})$ consists of a single point.
- (BOR-3) **Strong boundary monotonicity** For every $i = 1, \dots, m$ and $j = 1, \dots, 2k_i$, and every $t \in e^i$, we require that $u^{-1}(\boldsymbol{\gamma}_j^{i,a} \times \{t\})$ consists of at most one point.

From these conditions, it is clear that near each puncture q of S which is labelled by $p^{i,i+1}$, the map u is asymptotic to

$$I^{i,i+1}(q) := x^{i,i+1}(q) \times p^{i,i+1},$$

where $x^{i,i+1}(q)$ is a point in $\boldsymbol{\gamma}^i \cap \boldsymbol{\gamma}^{i+1}$. We collect these into sets of points

$$\mathbf{x}^{i,i+1} := \{x^{i,i+1}(q) \mid \lambda(q) = p_{i,i+1}\},$$

for each i , and collect the points $I^{i,i+1}(q)$ into corresponding sets $I^{i,i+1}(\mathbf{x}^{i,i+1})$.

We will be concerned with a few specific instances of this construction, which we spend the next sections giving more thorough names to.

4.4 For bordered diagrams

This section briefly recounts [14, Section 5.2].

Suppose that \mathcal{D} is a toroidal bordered diagram, $\mathcal{D} = (\Sigma; \alpha, \beta; \mathbf{b})$. We relabel $C^1 := C^\alpha$ and $C^2 := C^\beta$ for clarity's sake, and consider the case when S is a bordered source.

Given generators \mathbf{x} and \mathbf{y} , a choice of almost-complex structure J upon $X_{\mathcal{D},2}$, a fixed decorated source S and a homology class $B \in \pi_2(\mathbf{x}, \mathbf{y})$, we shall write $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S)$ for the moduli space of curves u with source S which satisfy (M-1)–(M-8) and (BOR-1)–(BOR-3) with respect to J , in the homology class B , with asymptotics (\mathbf{x}, \mathbf{y}) . Sometimes, to make explicit which almost-complex structure we are using, we will write $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; J)$.

If $u \in \mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S)$ is a holomorphic curve, then there is a corresponding *evaluation map* for each puncture q of S labelled by a Reeb chord, $ev_q(u) := t \circ u_{\bar{e}}(q)$. Sometimes, if u is implied, we refer to $ev_q(u)$ as the *height* of q . There are evaluation maps

$$ev_j = \prod_{q \in E_j(S)} ev_q: \mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; J) \rightarrow \mathbb{R}^{|E_j(S)|},$$

which fit together into a single map $ev: \mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; J) \rightarrow \mathbb{R}^{|E(S)|}$ given by $ev := ev_1 \times \cdots \times ev_h$.

If we fix a partition P_j of $E_j(S)$, there is a corresponding subspace $\Delta_P \subset \mathbb{R}^{|E_j(S)|}$ given by

$$\{x_p = x_q \mid P_j^i \in P_j, p, q \in P_j^i\}.$$

This allows us to further split the moduli space above by requiring that the heights of punctures in each part coincide; namely:

Definition 4.12 Let \mathbf{x}, \mathbf{y} be generators, $B \in \pi_2(\mathbf{x}, \mathbf{y})$, S be a decorated source, and $P := (P_1, \dots, P_h)$ partitions of the set $E_1(S), \dots, E_h(S)$. Then we write

$$\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; P_1, \dots, P_h) := ev^{-1}(\Delta_{P_1} \times \cdots \times \Delta_{P_h}) \subset \mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S).$$

Again, we sometimes write $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; P_1, \dots, P_h; J)$ to make explicit which almost-complex structure we are using, and often abbreviate to $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; P)$.

A further refinement is obtained by endowing each partition P_j with an ordering of its parts, \vec{P}_j . We can then write $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; \vec{P})$ for the subset of the set of curves $u \in \mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; P_1, \dots, P_h)$ for which the ordering of each \vec{P}_j agrees with that given by dictating $P_j^i < P_j^{i'}$ if $ev_{q^i}(u) < ev_{q^{i'}}(u)$ for all (any) $q^i \in P_j^i$ and $q^{i'} \in P_j^{i'}$. Often we abbreviate $\vec{P} := (\vec{P}_1, \dots, \vec{P}_h)$.

If we are given a set of sequences of Reeb chords $\vec{\rho}_{\Delta} = (\vec{\rho}_{\Delta,1}, \dots, \vec{\rho}_{\Delta,h})$, then we shall say the pair (\vec{P}, S) is *compatible* with $\vec{\rho}_{\Delta}$ if the sequence of sets of Reeb chords gotten from \vec{P}_j by replacing each part P_j^i with the set of Reeb chords labelling the punctures $q \in P_j^i \subset E_j(S)$ is equal to $\rho_{\Delta,j}^i$, and we

shall often write $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\triangle})$ for the moduli space $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{P})$, where (\vec{P}, S) are compatible with $\vec{\rho}_{\triangle}$. If (\vec{P}, S) are $\vec{\rho}_{\triangle}$ -compatible, for any $\rho \in \vec{\rho}_{\triangle, j} \in \vec{\rho}_{\triangle}$ we shall write $q(\rho)$ for the corresponding puncture of S .

Remark 4.13 The moduli space $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\triangle})$ has an action by \mathbb{R} , induced by translating in the \mathbb{R} -direction of $D_2 = [0, 1] \times \mathbb{R}$. Unless S is a disjoint union of g strips, each with two boundary punctures, and the map u has trivial homology class, this translation is free — and we say u is *stable*. We will have cause to consider the quotient moduli space $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\triangle})/\mathbb{R}$ which we term $\tilde{\mathcal{M}}^B(x, y; \vec{P})$.

In [14], what we call $\mathcal{M}_{\leftrightarrow}$ is called $\tilde{\mathcal{M}}$, and the moduli space we call here $\tilde{\mathcal{M}}$ is simply called \mathcal{M} . We reserve the latter for certain moduli spaces of bordered polygon with constraints on behaviour near certain corners; however, in the case where the bordered polygon is in fact a bigon, this does not present a further constraint and thus $\mathcal{M} = \tilde{\mathcal{M}}$, agreeing with the previous notation in the literature.

In [14], a number of properties of these moduli spaces are explored. We presently restate those that shall be of use to us.

Proposition 4.14 [14, Proposition 5.6] *For each x, y, S , there is a dense set of admissible J such the moduli spaces $\mathcal{M}_{\leftrightarrow}^B(x, y; S; J)$ are transversally cut out by the $\bar{\partial}$ -equations. Moreover, for any countable set $\{M_i\}$ of submanifolds of $\mathbb{R}^{|E(S)|}$, there is a dense set of admissible J which satisfy that $\text{ev}: \mathcal{M}_{\leftrightarrow}^B(x, y; S; J) \rightarrow \mathbb{R}^{|E(S)|}$ is transverse to all of the M_i ; ie there is a dense set of admissible J such that the moduli spaces $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\triangle}; J)$ are also transversely cut out.*

*In particular, the set of **sufficiently generic** admissible almost complex structures — those for which the moduli spaces $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\triangle}; J)$ are transversally cut out for **all** choices of $x, y, S, \vec{\rho}_{\triangle}$ — is nonempty.*

By the implicit function theorem, it follows that for sufficiently generic J the moduli space of $u \in \mathcal{M}_{\leftrightarrow}^B(x, y; S; J; \vec{\rho}_{\triangle})$ satisfying $\text{ind}(D\bar{\partial}(u)) = k$ is a k -dimensional manifold. The index at any given u satisfies a formula in terms of the topology of S and B , contingent on the *Euler measure* of a domain:

Definition 4.15 Let $B \in \pi_2(\mathbf{x}^{0,1}, \dots, \mathbf{x}^{k,0})$ be a domain. Then its *Euler measure*, $e(B)$, is given by

$$e(B) := \chi(B) - \frac{1}{4}\text{ac}(B) + \frac{1}{4}\text{ob}(B),$$

where $\chi(B)$ denotes the Euler characteristic, and $\text{ac}(B)$ and $\text{ob}(B)$ are the number of acute and obtuse corners in B respectively.

The relevant index formula is:

Proposition 4.16 [14, Proposition 5.8] *The expected dimension of $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\triangle}; J)$ is given by*

$$\text{ind}(B, S, P) = g - \chi(S) + 2e(B) + |\vec{\rho}_{\triangle, 1}| + \dots + |\vec{\rho}_{\triangle, h}|,$$

where $|\vec{\rho}_{\triangle, j}|$ is the number of Reeb chords in $\vec{\rho}_{\triangle, j}$, and $e(B)$ is the Euler measure of B .

It follows from this that for fixed B and S , the dimension of the moduli space $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; \vec{\rho}_{\Delta}; J)$ is well defined. There is a further interpretation of this formula solely in terms of data about the domain $[u]$, though we do not pursue this here; the important related fact for us will be:

Proposition 4.17 [14, Proposition 5.69] *Whether or not a holomorphic curve $u \in \mathcal{M}_{\leftrightarrow}(\mathbf{x}, \mathbf{y}; S; \vec{\rho}_{\Delta}; J)$ is embedded is determined solely by the Euler characteristic of the source S : the curve u is embedded if and only if $\chi(S) = \chi_{\text{emb}}(S)$ for some well-defined quantity $\chi_{\text{emb}}(S)$.*

Though the moduli spaces $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; \vec{\rho}_{\Delta}; J)$ form manifolds, they are not compact except when the moduli space $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; \vec{\rho}_{\Delta}; J)$ is zero-dimensional. They admit certain compactifications in terms of *holomorphic combs*, which we consider in Section 4.6.3.

4.4.1 Splicings and moduli spaces Suppose that S is a bordered source and \vec{P} is a sequence of partitions $\vec{P}_1, \dots, \vec{P}_h$, such that (S, \vec{P}) is compatible with a sequence of sets of Reeb chords $\vec{\rho} = (\vec{\rho}_1, \dots, \vec{\rho}_h)$.

The definitions of the moduli spaces $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; P)$ and $\mathcal{M}_{\leftrightarrow}^B(\mathbf{x}, \mathbf{y}; S; \vec{P})$ stipulate relations on the relative heights of punctures from each set $E_j(S)$, as compatible with the sequence $\vec{\rho}$. More generally, we could consider constraints upon relative heights of punctures which belong to $E_j(S)$ and $E_{j'}(S)$ for $j \neq j'$.

Suppose we are given an m -splicing $\vec{\sigma} = \vec{\sigma}_1, \dots, \vec{\sigma}_h$ of $\vec{\rho}$, and that (S, \vec{P}) is $\vec{\rho}$ -compatible. There are m corresponding sets

$$Q^i := \{q(\rho) \mid \rho \in \sigma_1^i \cup \dots \cup \sigma_h^i\}$$

of punctures. In turn, these yield corresponding submanifolds of $\mathbb{R}^{|E(S)|}$ specified by

$$\Delta_{\vec{\sigma}}^i := \{x_p = x_q \mid p, q \in Q^i\}.$$

This allows us to define moduli spaces $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; P, \sigma)$ and $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; \vec{P}, \sigma)$, given by

$$\text{ev}^{-1}(\Delta_{\vec{\sigma}}^1, \dots, \Delta_{\vec{\sigma}}^m)$$

for the evaluation maps from $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; P)$ and $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; \vec{P})$, respectively, to $\mathbb{R}^{|E(S)|}$. These are transversally cut out for sufficiently generic J by Proposition 4.14.

The corresponding index formula is given by

$$\text{ind}(B, S, P, \sigma) = g - \chi(S) + 2e(B) + |\vec{\rho}_{\Delta,1}| + \dots + |\vec{\rho}_{\Delta,h}| - \text{Col}(\sigma),$$

and is proved in a similar manner as [14, Proposition 5.8] — each height constraint coming from σ drops the expected dimension by one.

These moduli spaces only depend upon the unordered partitions in σ . The extra information of the order of the partitions in $\vec{\sigma}$ allows us to define even more restricted moduli spaces $\tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; \vec{P}, \vec{\sigma})$ as the set of $u \in \tilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S; \vec{P}, \sigma)$ such that the ordering of the parts in σ according to $\vec{\sigma}$ agrees with the ordering of the parts in σ by dictating that $Q^i < Q^{i'}$ if $\text{ev}_{q_i}(u) < \text{ev}_{q_{i'}}(u)$ for every (any) $q^i \in Q^i$ and $q^{i'} \in Q^{i'}$.

4.5 Heegaard multidiagrams

We now turn our attention to the relevant moduli spaces when \mathcal{D} is a Heegaard m -diagram, $\mathcal{D} = (\Sigma, \gamma^0, \dots, \gamma^{m-1}, \mathbf{p})$, where \mathbf{p} splits as $\mathbf{w} \cup \mathbf{z}$.

Fix a collection of admissible compatible almost-complex structures \mathbf{J} . As Σ is closed, the conditions (BOR-1)–(BOR-3) do not apply, and we are concerned with pairs (u, j) where $j \in \text{Conf}(D_m)$ and u is a (j_S, J_j) -holomorphic curve satisfying (M-1)–(M-8) disjoint from the set $\mathbf{z} \times D_m$.

The conditions imply that $\mathbf{x}^{i,i+1}$ is a generator for the pair (γ^i, γ^{i+1}) . In particular, the sequence $\vec{\mathbf{x}} = (\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m-1,0})$ constitutes a sequence of generators for \mathcal{D} .

For $m > 2$, we collect the holomorphic curves (u, j) with source S and given asymptotics $\vec{\mathbf{x}}$ in homology class $B \in \pi_2(\vec{\mathbf{x}})$ satisfying (M-1)–(M-8) into moduli spaces $\tilde{\mathcal{M}}^B(\vec{\mathbf{x}}; S)$. If $m = 2$, as in the previous section the relevant moduli space has a free \mathbb{R} -action, and so we denote by $\mathcal{M}^B(\vec{\mathbf{x}}; S)$ the result of taking the quotient by this. We call an element of one of these moduli spaces a *holomorphic m -gon*. As in the previous section, we sometimes include the family \mathbf{J} in the notation to draw attention to it; writing things like $\mathcal{M}^B(\vec{\mathbf{x}}; S; \mathbf{J})$, for instance.

For any subset $U \subset \{0, \dots, m - 1\}$ with $|U| = l$, we can apply the above construction to $\mathcal{D}|_U$ to define corresponding moduli spaces $\mathcal{M}^B(\mathbf{x}^{i_0,i_1}, \dots, \mathbf{x}^{i_{l-1},i_0}; S; J)$.

The relevant transversality result for holomorphic m -gons is:

Proposition 4.18 [13, Proposition 3.8] *There is an admissible, compatible family of almost-complex structures \mathbf{J} for \mathcal{D} , such that for every source, every l -subdiagram of \mathcal{D} , every sequence of generators for such a subdiagram, and every almost-complex structure $J \in \mathbf{J}_l$, the corresponding moduli spaces are all transversally cut out.*

There is a corresponding index formula:

Proposition 4.19 *Suppose that $(u, j) \in \mathcal{M}^B(\mathbf{x}^{0,1}, \dots, \mathbf{x}^{k-1,k}, \mathbf{x}^{k,0}; S)$ is a holomorphic polygon. Then the index of the $D\bar{\partial}$ operator at (u, J) is given by*

$$\text{ind}(B, S) = \frac{1}{2}(3 - k)g - \chi(S) + 2e(B).$$

Once more, it follows from this and the implicit function theorem that these moduli spaces are manifolds of dimension $\text{ind}(B, S) + k - 2$ (the $k - 2$ stems from the space $\text{Conf}(D_k)$ having moduli), though in general noncompact.

4.6 Bordered polygons

In this section, we construct relevant moduli spaces of polygons for bordered multidiagrams, generalising both of the previous two sections. For our purposes, it will be enough to consider bordered polygons that only have east asymptotics during the first edge of D_n , which we have called e^1 above.

We fix a toroidal bordered Heegaard m -diagram \mathcal{D} . Supposing that \mathbf{J} is a collection of consistent, admissible almost-complex structures, we again consider pairs (u, j) of $j \in \text{Conf}(D_m)$ and (j_s, J_j) -holomorphic curve u with source S satisfying (M-1)–(M-8) and (BOR-1)–(BOR-3), mapping into

$$(\Sigma \times D_m, C^0 \cup \dots \cup C^{m-1}).$$

The conditions imply that the sequence of points $\vec{x} = (x^{0,1}, \dots, x^{m-1,0})$ form a sequence of generators for \mathcal{D} . As per usual, such (u, j) have a homology class B in the group $\pi_2(\vec{x})$, and so we collect these maps into moduli spaces $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S)$.

Once again, if we wish to draw attention to the family \mathbf{J} in question, we include it in the notation; eg $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \mathbf{J})$. We often refer to a holomorphic curve (u, j) by solely u , and hope that this does not cause undue confusion.

As in Section 4.4, there is an evaluation map at each east puncture q of S , $\text{ev}_q(u): \mathcal{M}_{\leftrightarrow}^B(\vec{x}; S) \rightarrow e^1 \subset \partial D_m$. These fit together to form evaluation maps

$$\text{ev}_j: \mathcal{M}_{\leftrightarrow}^B(\vec{x}; S) \rightarrow (e^1)^{|E_j(S)|}$$

for each $j = 1, \dots, h$, which in turn fit together to give a map

$$\text{ev}: \mathcal{M}_{\leftrightarrow}^B(\vec{x}; S) \rightarrow (e^1)^{|E(S)|}.$$

As such, for any partitions $P := (P_1, \dots, P_h)$ of $E_1(S), \dots, E_h(S)$ there is once more a corresponding moduli space

$$\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; P) := \text{ev}^{-1}(\Delta_{P_1} \times \dots \times \Delta_{P_h}) \subset \mathcal{M}_{\leftrightarrow}^B(\vec{x}; S).$$

Furthermore, ordering the parts of the P_j as $\vec{P}_j = (P_j^1, \dots, P_j^{k_j})$ yields a corresponding moduli space

$$\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{P}),$$

where we further require that the orderings of the punctures induced by the evaluation maps ev_j agree with that of \vec{P}_j .

Now consider the set of sequences of Reeb chords $\vec{\rho}_{\Delta} = \vec{\rho}_1, \dots, \vec{\rho}_k$. We will say that a map u is compatible with $\vec{\rho}_{\Delta}$ if for each i and j , we have that $\{\rho(q)\}_{q \in P_j^i} = \rho_j^i$. Sometimes we shall write

$$\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{\rho}_{\Delta})$$

for the moduli space of curves in $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{P})$ which is compatible with $\vec{\rho}_{\Delta}$.

As in Section 4.4.1, if we are given an m -splicing $\vec{\sigma} = \vec{\sigma}_1, \dots, \vec{\sigma}_h$, there are corresponding sub-manifolds $\Delta_{\vec{\sigma}}^1, \dots, \Delta_{\vec{\sigma}}^m$ of $\mathbb{R}^{|E(S)|}$, and thus we may define the moduli spaces $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; P; \sigma)$ and $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{P}; \sigma)$ by $\text{ev}^{-1}(\Delta_{\vec{\sigma}}^1, \dots, \Delta_{\vec{\sigma}}^h)$ for the evaluation maps from $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; P)$ and $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{P})$ to $\mathbb{R}^{|E(S)|}$ respectively. We also write $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{P}; \vec{\sigma})$ for the subset of $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{P}; \sigma)$ such that the ordering upon jumping punctures induced by the evaluation maps agrees with the ordering dictated by $\vec{\sigma}$.

As the group of Möbius transformations acts three-transitively upon boundary points of the disc, we can assume that the greatest height of an east puncture of S labelled by a jumping chord of \mathcal{T} is 0. When $m \neq 2$, this does not alter the moduli spaces we study; when $m = 2$, these moduli spaces are precisely the bordered moduli spaces $\mathcal{M}_{\leftrightarrow}^B(x, y; S; \vec{\rho}_{\Delta}; \vec{\sigma})$ described in Section 4.4, and the effect of fixing the height to be zero is to instead consider the quotient moduli spaces $\tilde{\mathcal{M}}^B(x, y; S; \vec{\rho}_{\Delta}; \vec{\sigma})$. In general, we denote the relevant moduli space of curves with greatest height of an east puncture fixed to be zero by replacing any instances of $\mathcal{M}_{\leftrightarrow}$ with $\tilde{\mathcal{M}}$ in the above definitions.

There is a corresponding transversality result:

Proposition 4.20 *There is an admissible, compatible family of almost-complex structures \mathbf{J} for \mathcal{D} , such that for every l -fold subdiagram of \mathcal{D} , every source, every sequence of generators for the subdiagram, every set of sequences of Reeb chords $\vec{\rho}_{\Delta}$ and every corresponding splicing $\vec{\sigma}$, the corresponding moduli spaces $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}; \vec{\sigma})$ are all transversally cut out by the $\bar{\partial}$ equations. In this instance, we shall say that \mathbf{J} is **generic**.*

Proof This is a minor modification of [13, Proposition 4.18]. □

In fact, the proof of Proposition 4.20 contains a little more information. Namely:

Proposition 4.21 *Let \mathcal{U} denote the space of all admissible, compatible families almost-complex structures, and suppose that \mathbf{J} is a generic family of admissible, compatible almost-complex structures. For any path $\mathbf{J}_t : [0, 1] \rightarrow \mathcal{U}$ which satisfies that $\mathbf{J}_0 = \mathbf{J}$, there is some ε such that \mathbf{J}_t is generic for all $t \in [0, \varepsilon)$.*

This should be compared with, for instance, [16, Remark 3.2.8].

The common generalisation of Theorems 4.16 and 4.19 is below, and proved in a very similar manner; cf [13, Lemma 4.7].

Proposition 4.22 *Suppose that $(u, j) \in \tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, \vec{\sigma})$ is a bordered holomorphic polygon. Then the index of the $D\bar{\partial}$ operator at u is given by*

$$\text{ind}(B, S, \vec{\rho}_{\Delta}, \sigma) := \frac{1}{2}(3 - m)g - \chi(S) + 2e(B) + |\vec{\rho}_{\Delta, 1}| + \dots + |\vec{\rho}_{\Delta, h}| - \text{Col}(\sigma).$$

And, much like Proposition 4.17, we have the following:

Proposition 4.23 [13, Proposition 4.9] *Whether or not a holomorphic curve in $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, \vec{\sigma})$ is embedded is determined solely by the Euler characteristic of the source S .*

4.6.1 Compactifying at east infinity We now consider the effect of compactifying a toroidal bordered diagram at east infinity upon the corresponding moduli spaces. If $\mathcal{D} = (\Sigma; \gamma^1, \dots, \gamma^m; \mathbf{p}; \mathbf{b})$ is a toroidal bordered multidigraph, recall that there is a corresponding compactified diagram $\mathcal{D}_{\bar{e}}$, where $\partial\Sigma_j$ is

replaced by a single point e_j^∞ , and completing the arcs $\gamma^{1,a}$ to curves $\gamma_\varepsilon^{1,a}$. There is a corresponding complex structure j_{Σ_ε} (resp. symplectic form) upon Σ_ε which agrees with j_Σ (resp. ω_Σ) outside of arbitrarily small neighbourhoods of e_j^∞ , and is such that the compactified $\gamma_i^{1,a}$ are smooth curves (ie the angles between them are correct). Similarly, to every family of (generic, admissible, consistent) almost-complex structures $\mathbf{J} = \bigcup_m \mathbf{J}^m$ there is a corresponding family $\mathbf{J}_\varepsilon = \bigcup_m \mathbf{J}_\varepsilon^m$, compatible with the corresponding split symplectic form $\omega_{\Sigma_\varepsilon} \times \omega_{D_n}$. Any sequence of generators $\vec{x} \in \mathcal{G}(\mathcal{D})$ has an obvious counterpart in \mathcal{D}_ε , which we call \vec{x}' here but usually abuse notation and also call \vec{x} . Similarly each domain $B \in \pi_2(\vec{x})$ has a corresponding domain $B_\varepsilon \in \pi_2(\vec{x}')$.

To any map $u: S \rightarrow X_{\mathcal{D},m} \times D_m$ there is a corresponding map $u': S \rightarrow X_{\mathcal{D}_\varepsilon} \times D_m$. We can do slightly better, though: because of the conformal structure upon $X_{\mathcal{D}_\varepsilon}$ near east infinity, punctures corresponding to nonjumping chords can be removed completely from S . To any bordered polygonal source S we may associate the source S_ε given by compactifying S at its east punctures labelled by ρ_{12} or ρ_{23} . We then consider the space of pairs (u, J_ε) where $J_\varepsilon \in \mathbf{J}_\varepsilon$ and u is a J_ε -holomorphic curve satisfying analogues of properties (M-1)–(M-8) and (BOR-1)–(BOR-3).

We denote the moduli space of such pairs in homology class B_ε with asymptotics described by $\vec{x} = (\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m-1,0})$ by $\mathcal{M}_{\leftrightarrow \varepsilon}^{B_\varepsilon}(\mathbf{x}^{0,1}, \dots, \mathbf{x}^{m-1,0}; S_\varepsilon; \mathbf{J}_\varepsilon)$.

This is closely related to the moduli space $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \mathbf{J})$. In particular, if u is a holomorphic curve in $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \mathbf{J})$, there is a corresponding map $u_\varepsilon \in \mathcal{M}_{\leftrightarrow \varepsilon}^{B_\varepsilon}(\vec{x}; S_\varepsilon; \mathbf{J}_\varepsilon)$ which can be constructed as follows.

In a small neighbourhood of an east puncture of S the map u takes the form $u = (u_\Sigma, u_{D_n})$, as the almost-complex structure J is split near east infinity. By the Schwarz reflection principle and Riemann removable singularities theorem, the map u_Σ can be extended holomorphically across the puncture. By the uniqueness in Schwarz and Riemann's theorems, this identifies the moduli space $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \mathbf{J})$ with $\mathcal{M}_{\leftrightarrow \varepsilon}^{B_\varepsilon}(\vec{x}; S_\varepsilon; \mathbf{J}_\varepsilon)$, and we denote the image of $\mathcal{M}_{\leftrightarrow}^B(\vec{x}; S; \vec{\rho}_\Delta, \vec{\sigma}; \mathbf{J})$ under this identification by

$$\mathcal{M}_{\leftrightarrow \varepsilon}^{B_\varepsilon}(\vec{x}; S_\varepsilon; \vec{\rho}_\Delta, \vec{\sigma}; \mathbf{J}_\varepsilon).$$

If $u \in \mathcal{M}_{\leftrightarrow \varepsilon}^{B_\varepsilon}(\vec{x}; S_\varepsilon; \mathbf{J}_\varepsilon)$, then for $j = 1, \dots, h$, there is a sequence of partitions P_j of $\bar{E}_j(S_\varepsilon) := (\pi_\Sigma \circ u)^{-1}(e_j^\infty)$, determined by dictating that q and q' belong to the same part if and only if $\text{ev}(q) = \text{ev}(q')$. These may be ordered as \vec{P}_j , where for parts P_j and P'_j we say that $P_j < P'_j$ if and only if for some $q \in P_j$ and $q' \in P'_j$ we have that $\text{ev}_j(q) < \text{ev}_j(q')$.

If $q \in \bar{E}_j(S_\varepsilon)$, then there is an associated Reeb chord $\rho(q)$ defined by first choosing a small disc neighbourhood M_j of e_j^∞ such that $S_q := (\pi_\Sigma \circ u)^{-1}(M_j)$ is a small half-disc neighbourhood of $q \in S$ which is disjoint from all other elements of $\bar{E}_j(S_\varepsilon)$. The homology class of $(\pi_\Sigma \circ u)(S_q)$ is given by $B(\rho(q))$ for $\rho(q)$ some element of the torus algebra, as described in Section 3.1.1. It follows that u has a corresponding set of sequences of algebra elements $\vec{\rho}(u) := \vec{\rho}_1, \dots, \vec{\rho}_h$, given by replacing each $q \in \vec{P}_j$ by $\rho(q)$.

The subspace

$$\mathcal{M}_{\leftrightarrow \bar{e}}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; \bar{\rho}_{\Delta}, \bar{\sigma}; \mathbf{J}_{\bar{e}}) \subset \mathcal{M}_{\leftrightarrow \bar{e}}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; \mathbf{J}_{\bar{e}})$$

can be characterised as the set of u such that $\bar{\rho}(u) = \bar{\rho}_{\Delta}$, and such that the ordered partition of the punctures of $S_{\bar{e}}$ (which correspond with jumping chords in $\bar{\rho}_{\Delta}$) induced by the evaluation maps agrees with the splaying $\bar{\sigma}$.

As with the moduli spaces $\mathcal{M}_{\leftrightarrow}^B(\bar{x}; S; \bar{\rho}_{\Delta}; \bar{\sigma})$, we may assume that the highest height of an east puncture of S is 0—calling the corresponding moduli spaces $\tilde{\mathcal{M}}_{\bar{e}}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; \mathbf{J}_{\bar{e}})$ and $\tilde{\mathcal{M}}_{\bar{e}}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; \bar{\rho}_{\Delta}, \bar{\sigma}; \mathbf{J}_{\bar{e}})$.

If $\bar{\rho}_{\Delta}$ contains no jumping chords, then along the edge $\Sigma_{\bar{e}} \times e^1$, $u(\partial S_{\bar{e}})$ is contained in a set of $g + h - 1$ cylinders which are a subset of $C_{\bar{e}}^1 := \gamma_{\bar{e}}^1 \times e^1$. It follows that:

Proposition 4.24 *If $\bar{\rho}_{\Delta}$ has no jumping chords, then the moduli space $\tilde{\mathcal{M}}_{\bar{e}}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; \bar{\rho}_{\Delta}; \emptyset; \mathbf{J}_{\bar{e}})$ defined by counting (bordered) polygons associated with the partially splayed diagram $\mathcal{D}^{\partial}(\varepsilon, \bar{\iota})$ is identified with the moduli space $\tilde{\mathcal{M}}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; \mathbf{J}_{\bar{e}})$, defined by counting polygons with respect to the (not partially) splayed Heegaard multidiagram $\mathcal{D}(\varepsilon, \bar{\iota})$.*

4.6.2 Fixing cut behaviour Now we consider the case where \mathcal{D} is a splayed or partially splayed diagram $\mathcal{D}'(\varepsilon, \bar{\iota})$ or $\mathcal{D}^{\partial}(\varepsilon, \bar{\iota})$, and suppose that $B \in \text{Dom}(\mathcal{D})$ is a domain where $\bar{\sigma}(B) = \bar{\sigma}^1, \dots, \bar{\sigma}^k$ (see Definition 3.30).

If $\bar{\rho}(B) = ([\bar{\rho}_1(B)], \dots, [\bar{\rho}_h(B)])$, let $\tilde{\rho}_j = (\tilde{\rho}_j^1, \dots, \tilde{\rho}_j^{i_j})$ be the 123-reduced representative of $[\bar{\rho}_j(B)]$, and denote by $\tilde{\rho}_{\pm} := (\tilde{\rho}_1, \dots, \tilde{\rho}_h)$. If S is a decorated k -source, we denote by $\rho_{123}(S, u)$ the set of punctures of S which are mapped by u to a point $e_j^{\infty, i}$ such that the i^{th} chord $\tilde{\rho}_j^{k_i}$ in $\text{jump}(\tilde{\rho}_j)$ satisfies $\tilde{\rho}_j^{k_i} = \rho_{123}$.

We will impose extra conditions upon the moduli spaces $\tilde{\mathcal{M}}^B(\bar{x}; S; \bar{\rho}_{\Delta}; \bar{\sigma})$ based upon the behaviour of u near any $q \in \rho_{123}(S, u)$. Namely, fix such a q which is mapped to $e_j^{\infty, i}$. We shall say that u has a *cut* at q if $\pi_{\Sigma} \circ u$ has a boundary branch point along one of the arcs γ_j^i or γ_j^{i+1} . There are three cases:

- (1) If there is a boundary branch point upon $\gamma_j^{i, i+1}$, we will say that q is of (ρ_{12}, ρ_3) -type.
- (2) When there is a boundary branch point upon $\gamma_j^{i, i+1}$, we will say that q is of (ρ_1, ρ_{23}) -type.
- (3) When there is no boundary branch point we shall say that q is of ρ_{123} -type.

This has an equivalent description: if there is a branch point along γ_j^i then there is some small disc neighbourhood $S(q)$ of $e_j^{\infty, i}$ such that when restricted to the component $S(q)$ of $u|_{S(q)}$ containing q , the homology class of $\pi_{\Sigma} \circ u(S(q))$ is modelled upon $B(\rho_3)$. Similarly, if the branch point appears upon γ_j^{i+1} , there is a corresponding $S(q)$ such that the homology class of $\pi_{\Sigma} \circ u(S(q))$ is modelled upon $B(\rho_1)$. Finally, if there is no cut at q , every sufficiently small disc neighbourhood M_j^i of q satisfies that, restricted to the component $S(q)$ containing q , the map $\pi_{\Sigma} \circ u$ is modelled upon $B(\rho_{123})$.

We would like to impose conditions upon our moduli spaces that dictates which of these cases holds for each q . We first choose sensible notation, and then rewrite the conditions in a way for which transversality and index formulae are more apparent.

Definition 4.25 Fix $B \in \pi_2(\bar{x})$ with associated sets of sequences of Reeb chords $\vec{\rho}_\pm$, and fix also a splayed source S with corresponding partition of $\rho_{123}(S)$ into $Q = (Q_{123}, Q_{1,23}, Q_{12,3})$.

We shall write $\tilde{\mathcal{M}}^B(\bar{x}; S; \vec{\rho}_\pm, Q_{123}; \vec{\sigma})$ for the subset of $u \in \tilde{\mathcal{M}}^B(\bar{x}; S; \vec{\rho}_\pm; \vec{\sigma})$ where $\rho_{123}(S, u) = \rho_{123}(S)$ and each $q \in Q_{123}$ is of ρ_{123} -type (with respect to u). Similarly, we shall write $\mathcal{M}^B(\bar{x}; S; \vec{\rho}_\pm, Q; \vec{\sigma})$ for the subset of $\tilde{\mathcal{M}}^B(\bar{x}; S; \vec{\rho}_\pm, Q_{123}; \vec{\sigma})$ such that each $q \in Q_{1,23}$ is of (ρ_1, ρ_{23}) -type, and each $q \in Q_{12,3}$ is of (ρ_{12}, ρ_3) -type.

We often use a shorthand for these moduli spaces:

Definition 4.26 Write $\vec{\rho}_\pm(Q, u)$ for the representative for the 123-equivalence class $[\vec{\rho}(B)]$ given by replacing each $\rho \in \vec{\rho}_\pm$ which is equal to ρ_{123} by ρ_{12}, ρ_3 if $q(\rho) \in Q_{12,3}$; by ρ_1, ρ_{23} if $q(\rho) \in Q_{1,23}$; and, finally, leaving it as ρ_{123} if $q(\rho) \in Q_{123}$.

If $\vec{\rho}_\pm$ is a representative of $[\vec{\rho}(B)]$, define by

$$\mathcal{M}^B(x, S, \vec{\rho}_\pm, \vec{\rho}_\pm; \vec{\sigma})$$

the moduli space

$$\mathcal{M}^B(\bar{x}; S; \vec{\rho}_\pm, Q; \vec{\sigma}),$$

where $\vec{\rho}_\pm = \vec{\rho}_\pm(Q, u)$.

We call an element of $\mathcal{M}^B(x, S, \vec{\rho}_\pm, \vec{\rho}_\pm; \vec{\sigma})$ a *splayed bordered holomorphic polygon*.

To ensure the corresponding moduli spaces are transversally cut out, we rephrase this condition in a way more compatible with the language of almost-complex structures (cf [18, Lemma 9.4; 1, Chapter 6]). First, a definition:

Definition 4.27 Let $u \in \mathcal{M}^B(\bar{x}; S; \vec{\rho}_\pm, Q; \vec{\sigma})$, and let $q \in Q_{1,23}$ or $Q_{12,3}$. Denote the corresponding boundary branch point by b , and the arc of ∂S connecting q and $(\pi_\Sigma \circ u)^{-1}(b)$ by $\lambda_{q,b}$. We define the *length of the cut at q* to be the length of the arc $\bar{\lambda}_{q,b} := (\pi_\Sigma \circ u)(\lambda_{q,b})$ with respect to the metric induced upon Σ by the uniformisation theorem.

We focus initially upon curves where the cut behaviour near each $q \in \rho_{123}(S)$ is fixed so that either there is no cut, or such that the length of the cut at q is short enough that the corresponding boundary branch point b is contained in a disc neighbourhood N_j^i of $e_j^{\infty,i}$ which is holomorphically identified with the unit disc in \mathbb{C} . We also require that this disc is small enough that for each q , the component $S(q)$ of $(\pi_\Sigma \circ u)^{-1}(N_j^i)$ containing q is a half-disc with q its only boundary puncture. Denote by (u_j^i, j') the restriction of (u, j') to $(\pi_\Sigma \circ u)^{-1}(N_j^i)$ (we recall that we have been writing u for the pair (u, j') where $j' \in \text{Conf}(D_m)$).

We can modify the complex structure upon N_j^i so that the angle between the arcs in $\partial B_j^i(\rho_{123})$ is π (by composing with the map $\varphi: z \mapsto z^{2/3}$, say), and modify the (split) almost-complex structure $J|_{N_j^i \times j'}$

over $N_j^i \times D_m$ accordingly, to a corresponding almost-complex structure $J_{N_j^i}^0 \times j'$. To u_j^i we may associate a corresponding v_j^i with respect to this new almost-complex structure: moreover, if u_j^i had no boundary branch point then this new map does not either, and so there is a corresponding holomorphic map from the half-disk $S(q)$ with boundary puncture q removed to the manifold $X_{\mathcal{D},m}$ (by the removable singularities theorem and Schwarz reflection principles).

Similarly, if u_j^i had a boundary branch point then so does v_j^i . We further modify the complex structure by composing the chart on N_j^i with the Riemann mapping which takes the open disc with slit of length t along $\varphi(\gamma_j^{i+1})$ (resp. slit of length $-t$ along $\varphi(\gamma_j^i)$) to the open half-disc. These Riemann mappings yield a family of almost-complex structures $J_{N_j^i}^t$ upon N_j^i for each $t \in (-1, 1)$, such that for every $(J_{N_j^i}^t \times j')$ -holomorphic curve u_j^i with a cut of length t along γ_j^{i+1} (resp. length $-t$ along γ_j^i), the corresponding curve (which we also denote by v_j^i) with respect to $J_{N_j^i}^t \times j'$ has no cut, and thus can be completed to a holomorphic map with source given by the half-disk.

We would like to modify the almost-complex structures $J_{j'}$ to agree with this $J_{N_j^i}^t \times j'$ in a neighbourhood of $u(q)$ and then apply the transversality result of Proposition 4.20, but this does not make sense with respect to other components of $u^{-1}(N_j^i)$. The solution, as in [14, Proposition 6.36] (see also [25]) is to instead consider, for $\vec{t} = \{t_q\}_{q \in \rho_{123}(S)}$, related moduli spaces of $\mathbf{J}^{\vec{t}}$ -holomorphic sections $u: S \rightarrow S \times \Sigma_{\vec{e}} \times D_n$, where each $\mathbf{J}^{\vec{t}}$ is an almost-complex structure on $S \times \Sigma_{\vec{e}} \times D_n$ equal upon each (s, z, w) to some $j_S \times \mathbf{J}_s^{\vec{t}}$. We choose $\mathbf{J}_s^{\vec{t}}$ so that for s near any puncture $q \in \rho_{123}(S)$, the almost-complex structure $\mathbf{J}_s^{\vec{t}}|_{N_j \times D_m}$ agrees with the split almost-complex structure $J_{N_j^i}^t \times j'$ described in the previous paragraph, and away from this agrees with \mathbf{J} .

The proof of transversality for such families follows closely along the lines of [10, Chapter 3]. As such, the moduli spaces $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, Q_{123}; \vec{\sigma})$ are transversally cut out. Moreover the moduli space

$$\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, \rho_{123}(S); \vec{\sigma}) \subset \tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, Q_{123}; \vec{\sigma})$$

of curves with no cut at any member of $\rho_{123}(S)$ is also transversally cut out and, along the lines of Proposition 4.21, for some ε , the moduli spaces defined with respect to $\mathbf{J}^{\vec{t}}$ for $\vec{t} \in (-\varepsilon, \varepsilon)^{|\rho_{123}(S) - Q_{123}|}$ are also transversally cut out — thus each curve $u \in \tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, \rho_{123}(S), \vec{\sigma})$ has a neighbourhood in $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, Q_{123}; \vec{\sigma})$ homeomorphic to $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, Q_{123}; \vec{\sigma}) \times (-\varepsilon, \varepsilon)^{\rho_{123}(S) - Q_{123}}$.

Suppose that $u, u' \in \tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, Q_{123}; \vec{\sigma}_{\Delta})$, and that with respect to u a puncture $q \in \rho_{123}(S) - Q_{123}$ is of (ρ_{12}, ρ_3) -type, and with respect to u' the same puncture is of (ρ_1, ρ_{23}) -type. The topology upon the moduli spaces described in Definition 4.34 is metrizable (cf [4, Appendix B]), and for each $q \in \rho_{123}(S) - Q_{123}$ the length of the curve $\bar{\lambda}_{q,b}$ is continuous with respect to this metric. As such (by the intermediate value theorem), the curves u and u' belong to different path-components of

$$\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, Q_{123}; \vec{\sigma}) - \tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_{\Delta}, \rho_{123}(S); \vec{\sigma}).$$

It follows that:

Proposition 4.28 *There is an admissible, compatible family of almost-complex structures \mathbf{J} for \mathcal{D} , such that for all choices of argument, the moduli spaces $\mathcal{M}^B(\vec{x}; S; P, Q; \vec{\sigma}_\Delta)$ are all transversally cut out, and of expected dimension*

$$\text{ind}(B, S, P, Q, \sigma) := \text{ind}(B, S, P, \sigma) - \#_{123}(\vec{\rho}_\pm).$$

Proof Transversality is explained in the lead-up to the statement of the proposition. The index formula follows along similar lines to [10, Section 4.1]: we do not count “corners” stemming from punctures of ρ_{123} -type, as with respect to the relevant \mathbf{J}^i they are not corners at all. \square

Convention 4.29 In the course of this paper, we will be considering a number of different partial splittings of a bordered Heegaard diagram, which has the potential to become confusing. By Definition 2.7, a splicing $\vec{\sigma}$ of a set of sequences of Reeb chords $\vec{\rho}_\pm$ determines a sequence of sets of idempotents $\iota_{\vec{\sigma}}$; as such, we may indicate the diagram being considered for a specific moduli space by using the notation

$$\mathcal{M}^B(\mathbf{x}; S; \vec{\rho}_\Delta, \vec{\rho}_\pm; \vec{\sigma}_\Delta, \vec{\sigma}_\pm)$$

to refer to the moduli space

$$\mathcal{M}^B(\mathbf{x}; S; \vec{\rho}_\Delta, \vec{\rho}_\pm; \vec{\sigma}_\Delta)$$

defined using the diagram $\mathcal{D}^\partial(\varepsilon, \iota_{\vec{\sigma}_\pm})$.

Similarly, to keep notation compact, when it is clear from the context we will write

$$\mathcal{M}^B(\mathbf{x}; S; \vec{\rho}_\pm; \vec{\sigma}_\pm)$$

for the moduli space

$$\mathcal{M}^B(\mathbf{x}; S; \emptyset, \vec{\rho}_\pm; \emptyset, \vec{\sigma}_\pm)$$

defined using the closed multidiagram $\mathcal{D}(\varepsilon, \iota_{\vec{\sigma}})$, and

$$\text{ind}(B, S, Q_\pm) := \text{ind}(B, S, \emptyset, Q_\pm, \emptyset),$$

where B is a domain in $\mathcal{D}(\varepsilon, \iota_{\vec{\sigma}})$.

4.6.3 Compactifications As in [14], the aforementioned moduli spaces admit natural compactifications in terms of *holomorphic combs*.

To discuss these, we need the notion of a holomorphic tooth; this is a map

$$v: (T, \partial T) \rightarrow (\mathbb{R} \times (\partial\Sigma - \mathbf{b}) \times D_n, \mathbb{R} \times \mathbf{a} \times e^1)$$

where T is a tooth source, and the codomain is endowed with the obvious split symplectic form and the split complex structure $J = j_\Sigma \times j_{D_n}$. We denote by $\mathcal{N}(T)$ the moduli space of such maps satisfying:

- (1) v is (j, J) -holomorphic for some almost complex structure j upon T .
- (2) v is proper.
- (3) $\pi_D \circ v$ maps ∂T to e^1 .

(4) Near each west puncture $q \in W_j(T)$ labelled by a Reeb chord ρ ,

$$\lim_{z \rightarrow q} \pi_\Sigma \circ v(z) = \rho \subset \{-\infty\} \times \partial \Sigma_j.$$

(5) Near each east puncture $q \in E_j(T)$ labelled by a Reeb chord ρ ,

$$\lim_{z \rightarrow q} \pi_\Sigma \circ v(z) = \rho \subset \{+\infty\} \times \partial \Sigma_j.$$

These moduli spaces can once more be refined by requiring the evaluation maps $\text{ev}_q := \lim_{z \rightarrow q} t \circ v(z)$ partition the east and west punctures of T in a prescribed manner. Explicitly, if $P_w := (P_{w,1}, \dots, P_{w,h})$ and $P_e := (P_{e,1}, \dots, P_{e,h})$ are partitions of the sets $W(T) = (W_1(T), \dots, W_h(T))$ and $E(T) = (E_1(T), \dots, E_h(T))$ respectively, denoting

$$\begin{aligned} \text{ev}_{w,j} &:= \prod_{q \in W_j(T)} \text{ev}_q : \mathcal{N}(T) \rightarrow (e^1)^{|W_j(T)|}, \\ \text{ev}_{e,j} &:= \prod_{q \in E_j(T)} \text{ev}_q : \mathcal{N}(T) \rightarrow (e^1)^{|E_j(T)|}, \end{aligned}$$

and $\text{ev}_w := \text{ev}_{w,1} \times \dots \times \text{ev}_{w,h}$ and $\text{ev}_e := \text{ev}_{e,1} \times \dots \times \text{ev}_{e,h}$, we write $\mathcal{N}(T; P_w, P_e)$ for the moduli space

$$(\text{ev}_w \times \text{ev}_e)^{-1}(\Delta_{P_{w,1}} \times \dots \times \Delta_{P_{w,h}} \times \Delta_{P_{e,1}} \times \dots \times \Delta_{P_{e,h}});$$

omitting P_w from the notation if it is the discrete partition. If we are given a splicing σ_e of the corresponding sequence of Reeb chords given by replacing each $q \in P_e$ by the Reeb chord labelling it, we can restrict these moduli spaces even further, writing $\mathcal{N}(T; P_w, P_e; \sigma_e)$ for the subset $\text{ev}^{-1}(\Delta_\sigma^1, \dots, \Delta_\sigma^h)$ for the evaluation maps from $\mathcal{N}(T; P_w, P_e; \sigma_e) \rightarrow e^1$.

Transversality is known in a simple case:

Proposition 4.30 [14, Proposition 5.16] *If all components of T are topological disks, then $\mathcal{N}(T)$ is transversally cut out by the $\bar{\partial}$ -equations for any choice of split almost-complex structure J .*

This will allow us to describe the compactifications of our moduli spaces, as in [14, Section 5.4]. We firstly consider small constituent pieces:

Definition 4.31 A simple holomorphic m -comb is a pair (u, v) of a bordered holomorphic polygon $(u, j) \in \mathcal{M}_{\leftrightarrow}^B(\vec{x}; S)$ and holomorphic teeth $v \in \mathcal{N}(T)$ for some bordered polygon source S and tooth source T , together with a one-to-one correspondence between $E(S)$ and $W(T)$ that preserves labelling by Reeb chords, such that $\text{ev}_w(v) = \text{ev}(u)$.

In more generality, we must consider curves with even more degenerate behaviour at east infinity:

Definition 4.32 A holomorphic m -storey is a sequence $U = (u, v_1, \dots, v_k)$ for some $k \geq 0$, where, for some B , bordered polygonal source S and collection of tooth sources T_1, \dots, T_k ,

- $u \in \mathcal{M}_{\leftrightarrow}^B(\vec{x}; S)$;

- $v_i \in \mathcal{N}(T_i)$;
- (u, v_1) is a simple holomorphic m -comb;
- there is a correspondence between $E(T_i)$ and $W(T_{i+1})$ for $i = 1, \dots, k - 1$ which preserves labelling by Reeb chords; and
- $\text{ev}_e(v_i) = \text{ev}_w(v_{i+1})$ for every i .

If $k = 0$, we say that U is *toothless*.

In full generality, we need to consider trees of holomorphic stories.

Definition 4.33 A tree of holomorphic combs is:

- (1) A metric tree (T, ρ) with corresponding reduced tree \bar{T} .
- (2) A generator assignment $\vec{x}(\bar{T})$ for \bar{T} ; a tree of $\vec{x}(\bar{T})$ -compatible domains $B(\bar{T})$; a comb-tree of sources $S(\bar{T})$; and a tree of admissible, generic, compatible tree of almost-complex structures $J_T \subset J$.
- (3) A set of holomorphic stories $\{U_V\}_{V \in \mathcal{V}(\bar{T})}$, such that each U_V is a $J(V)$ -holomorphic $|V|$ -storey in the moduli space $\mathcal{M}_{\leftrightarrow}^{B(V)}(\vec{x}(V); S(V))$, and for every noncomb vertex $V \in \bar{T}$, the storey U_V is toothless.

If $|V(\bar{T})| = k$, we often say that U has *height* k , and we often refer to a height k tree of holomorphic combs by its vertices: ie $T = (u_1, \dots, u_k)$.

A priori, we must allow holomorphic curves to have nodal sources to ensure compactness holds. As in [14, Section 5.4], we refer to a tree of elementary holomorphic combs as *nodal* or *smooth* if we want to explicitly allow or forbid nodes.

Recall that to each spinal vertex V_1, \dots, V_k of \bar{T} we may associate its storied tree $T(V_i)$. Each of these has an associated domain $B(V_i)$, equal to the sum of the domains $\{B(V')\}_{V' \in T(V_i)}$, and so we may associate to any T -tree U the domain B given by the sum of the $B(V_i)$. We can also associate to U a set of well-defined Reeb chords which is given by the Reeb chords of the right-most tooth of each storey in the spine.

As in [14, Section 5.4], we call a component C of a T -tree U Σ -stable (resp. D -stable) if $(\pi_\Sigma \circ U)|_C$ (resp. $(\pi_D \circ U)|_C$) is a stable map, and Σ -unstable (resp. D -unstable) otherwise. We modify the definition of convergence from [14] to general T -trees of elementary combs:

Definition 4.34 We shall say that a sequence (u_n, j_n) of holomorphic n -gons in $\Sigma_{\bar{e}} \times D_m$ converges to a holomorphic tree U if:

- (1) Let S_Σ denote the result of collapsing all Σ -unstable components of U . Then the map $\{\pi_\Sigma \circ u_n\}$ converges to $\pi_\Sigma \circ U|_{S_\Sigma}$, in the sense of [4, Section 7.3].

- (2) The complex structures j_n on D_m converge to the tree of almost complex structures $\{j_v\}_{v \in T}$ over $\{\Sigma_{\vec{e}} \times D_{|v|}\}_{v \in T}$.
- (3) Let S_D be the result of collapsing all D -unstable components of U . Then $\{\pi_D \circ u_n\}$ converges to $\pi_D \circ U|_{S_D}$ in the sense of [4, Section 7.3].
- (4) For n sufficiently large, all u_n represent the same homology class $B = [U]$.

As is tradition, to avoid drowning in notation, we leave the more general definition of convergence of T -trees of elementary holomorphic combs to the reader.

With this in hand, we make the following definition:

Definition 4.35 Denote by

- $\overline{\overline{\mathcal{M}}}^B(\vec{x}; S)$ the space of all trees of elementary holomorphic combs with corresponding preglued source S in homology class B , asymptotic to the sequence of generators $\vec{x} = (x^{0,1}, \dots, x^{m-1,0})$, compatible with the almost complex structures \mathbf{J} ;
- $\overline{\mathcal{M}}^B(\vec{x}; S)$ the closure of the moduli space $\mathcal{M}^B(\vec{x}; S)$ in $\overline{\overline{\mathcal{M}}}^B(\vec{x}; S)$;
- $\overline{\overline{\mathcal{M}}}^B(\vec{x}; S; P; \sigma)$ the space of all relevant holomorphic trees respecting the partition P ;
- $\overline{\mathcal{M}}^B(\vec{x}; S; P; \sigma)$ the closure of $\mathcal{M}^B(\vec{x}; S; P; \sigma)$ in $\overline{\overline{\mathcal{M}}}^B(\vec{x}; S; P)$.

To save on ink and patience, we observe the convention that for all other moduli spaces $\mathcal{M}^B(*)$, we denote by $\overline{\mathcal{M}}^B(*)$ the closure of $\mathcal{M}^B(*)$ in the corresponding moduli space $\overline{\overline{\mathcal{M}}}^B(\vec{x}; S; P; \sigma)$.

Indeed, these are the required notions for compactifying the moduli spaces:

Proposition 4.36 *The spaces above are compact; ie if $\{U_n\}$ is a tree of holomorphic polygon combs in a fixed homology class with fixed preglued topological source S , then $\{U_n\}$ has a subsequence which converges to a tree of holomorphic polygon combs U in the same homology class.*

Proof The argument from [14, Proposition 5.24] carries over readily to this more general case. In particular, [14, Propositon 5.27] applies to a general symmetric cylindrical almost-complex manifold, and doubling both source and $X_{\mathcal{D},m}$ is still a manifold of this type. □

4.7 Holomorphic anchors and perturbations

We discuss modifying the Lagrangian cylinders used in the definition of the moduli spaces

$$\mathcal{M}^B(\vec{x}; S; \vec{\rho}_{\Delta}, \vec{\rho}_{\ddagger}; \vec{\sigma}_{\Delta}, \vec{\sigma}_{\ddagger}).$$

4.7.1 Interpolating cylinders Fix a partially splayed bordered diagram $\mathcal{D}' = \mathcal{D}^\partial(\varepsilon, \vec{\tau})$ with attaching curves $\gamma^0, \gamma^1, \dots, \gamma^m$, and recall the east-compactified curves $\gamma_{\vec{e}}^1 = \gamma_{\vec{e}}^{1,a} \cup \gamma_{\vec{e}}^{1,c}$ with corresponding cylinders in the east-compactification of $X_{\mathcal{D}',m}$ given by

$$C_j^1 := \gamma_{\vec{e}}^1 \times e^1 \subset \Sigma_{\vec{e}} \times e^1.$$

We can perturb these cylinders using the Hamiltonian H as described in Section 3. Namely, we put

$$C^1(\varepsilon) = \begin{cases} H(\boldsymbol{\gamma}_{\bar{e}}^1, 0) \times t & \text{for } t \in (-\infty, -\frac{\varepsilon}{2}], \\ H(\boldsymbol{\gamma}_{\bar{e}}^1, \frac{2t+\varepsilon}{2(m+1)}) \times t & \text{for } t \in [-\frac{\varepsilon}{2}, \frac{\varepsilon}{2}], \\ H(\boldsymbol{\gamma}_{\bar{e}}^1, \frac{\varepsilon}{m+1}) \times t & \text{for } t \in [\frac{\varepsilon}{2}, \infty). \end{cases}$$

(recall that we fixed an identification of a neighbourhood of e^1 with $\mathbb{R} \times [0, \varepsilon)$). For $t \leq -\frac{\varepsilon}{2}$, these cylinders coincide with the cylinders $C^1 = \boldsymbol{\gamma}_{\bar{e}}^1 \times e^1$, and for $t \geq \frac{\varepsilon}{2}$, they coincide with the cylinders $(C^1)' = H(\boldsymbol{\gamma}_{\bar{e}}^1, \varepsilon/(m+1)) \times e^1$. They are Lagrangian with respect to a perturbed symplectic form upon the east compactification of $X_{\mathcal{D}',m}$, as described in [13, Equation 3.25].

4.7.2 Moduli spaces with perturbed cylinders Consider again a partially splayed diagram $\mathcal{D}^\partial(\varepsilon, \vec{\iota})$, with attaching curves labelled $\boldsymbol{\gamma}^0, \dots, \boldsymbol{\gamma}^m$. We consider the moduli space of holomorphic curves

$$\mathcal{M}_{\bar{e},\varepsilon}^{B_{\bar{e}}}(\vec{x}; S_{\bar{e}}; \vec{\rho}_{\blacktriangle}, \vec{\rho}_{\blacktriangle\uparrow}; \vec{\sigma}; \mathbf{J}_{\bar{e}}),$$

which are defined by replacing the unperturbed cylinders $C_{\bar{e}}^i = C^i(\boldsymbol{\gamma}_{\bar{e}}^i)$ in the definition of the east-compactified moduli space

$$\mathcal{M}_{\bar{e}}^{B_{\bar{e}}}(\vec{x}; S_{\bar{e}}; \vec{\rho}_{\blacktriangle}, \vec{\rho}_{\blacktriangle\uparrow}; \vec{\sigma}; \mathbf{J}_{\bar{e}})$$

with cylinders such that

- $C^0(\varepsilon)$ are the cylinders $\boldsymbol{\gamma}_{\bar{e}}^0 \times e^0$,
- $C^1(\varepsilon)$ are the perturbed cylinders $C^1(\varepsilon)$, and
- $C^i(\varepsilon)$ are the cylinders $H(\boldsymbol{\gamma}_{\bar{e}}^i, \frac{\varepsilon(m-i)}{m(m+1)}) \times e^i$ for all $i \geq 2$.

Note that if $\iota = (\iota_1, \dots, \iota_h)$ is a set of idempotents, then for the corresponding partially splayed diagram $\mathcal{D}^\partial(\varepsilon, \iota \cup \vec{\iota})$ with attaching curves $\boldsymbol{\delta}^0, \boldsymbol{\delta}^1, \dots, \boldsymbol{\delta}^{m+1}$, we have that

- $C^0(\varepsilon)$ agree with $\boldsymbol{\delta}_{\bar{e}}^0 \times e^0$,
- $C^1(\varepsilon)$ agrees with $\boldsymbol{\delta}_{\bar{e}}^1$ for $t \leq -\frac{\varepsilon}{2}$ and with $\boldsymbol{\delta}_{\bar{e}}^2$ for $t \geq \frac{\varepsilon}{2}$, and
- for every $i \geq 2$, $C^i(\varepsilon)$ agrees with the cylinders $\boldsymbol{\delta}^{i+1} \times e^i$.

Proposition 4.37 For small enough ε , all of the moduli spaces of the form

$$\mathcal{M}(\mathbf{J}_{\bar{e}}) := \mathcal{M}_{\bar{e},\varepsilon}^{B_{\bar{e}}}(\vec{x}; S_{\bar{e}}; \vec{\rho}_{\blacktriangle}, \vec{\rho}_{\blacktriangle\uparrow}; \vec{\sigma}; \mathbf{J}_{\bar{e}})$$

are identified with the corresponding moduli space

$$\mathcal{M}_{\varepsilon}(\mathbf{J}_{\bar{e}}) := \mathcal{M}_{\bar{e},\varepsilon}^{B_{\bar{e}}}(\vec{x}; S_{\bar{e}}; \vec{\rho}_{\blacktriangle}, \vec{\rho}_{\blacktriangle\uparrow}; \vec{\sigma}; \mathbf{J}_{\bar{e}}).$$

Proof For every $m \geq 2$, consider a Lagrangian isotopy $\phi_t = \phi(x, t): X_{\mathcal{D},m} \times [0, 1] \rightarrow X_{\mathcal{D},m}$ which is supported in a ε -neighbourhood of each $\Sigma_{\bar{e}} \times e^i$, and satisfies that $\phi_t(C_{\bar{e}}^i) = C_{\bar{e}}^i(\varepsilon/t)$ for every $t \in [0, 1]$.

A holomorphic curve in the moduli space $\mathcal{M}_{t\varepsilon}(\mathbf{J}_{\bar{\varepsilon}})$ is the same as a holomorphic curve in the moduli space $\mathcal{M}(\mathbf{J}_{\bar{\varepsilon}}^t)$, where $\mathbf{J}_{\bar{\varepsilon}}^t$ is the family of perturbed almost-complex structures $\phi_t^* \mathbf{J}$. Thus we may identify these two moduli spaces. Moreover, the family of almost-complex structures $\mathbf{J}_{\bar{\varepsilon}}^t$ constitutes a smooth path in the space of families of admissible almost-complex structures. As such, by [Proposition 4.21](#), there is some $T \in (0, 1)$ such that

$$\bigcup_{t \in [0, T]} \mathcal{M}_{t\varepsilon}(\mathbf{J}_{\bar{\varepsilon}}) = \bigcup_{t \in [0, T]} \mathcal{M}(\mathbf{J}_{\bar{\varepsilon}}^t)$$

provides a cobordism between $\mathcal{M}(\mathbf{J}_{\bar{\varepsilon}}) = \mathcal{M}_0(\mathbf{J}_{\bar{\varepsilon}})$ and $\mathcal{M}_{T\varepsilon}(\mathbf{J}_{\bar{\varepsilon}})$; thus choosing $\varepsilon = T\varepsilon$ provides the result. □

Remark 4.38 There is a similar construction for subdiagrams: if $\mathcal{D}' := \mathcal{D}^\partial(\varepsilon, \bar{\iota})|_S$ is a subdiagram of $\mathcal{D}^\partial(\varepsilon, \bar{\iota})$, and $\bar{\mathbf{x}}'$ is a set of generators for \mathcal{D}' , there is a corresponding perturbed moduli space

$$\mathcal{M}_{\bar{\varepsilon}, \varepsilon}^{B_{\bar{\varepsilon}}}(\bar{\mathbf{x}}'; S'_{\bar{\varepsilon}}; \bar{\rho}_{\blacktriangle}; \mathbf{J}_{\bar{\varepsilon}})$$

formed by replacing the cylinders $\{C^i \mid i \in S\}$ used in the definition of

$$\mathcal{M}_{\bar{\varepsilon}}^{B_{\bar{\varepsilon}}}(\bar{\mathbf{x}}'; S'_{\bar{\varepsilon}}; \bar{\rho}_{\blacktriangle}; \mathbf{J}_{\bar{\varepsilon}})$$

with the corresponding perturbed cylinders $\{C^i(\varepsilon) \mid i \in S\}$. The proof of [Proposition 4.37](#) carries over to this case to identify these two moduli spaces.

Similarly, for regularisations of diagrams ([Definition 3.18](#)) we also have:

Proposition 4.39 *Let $\mathcal{D}|_S$ be a subdiagram of \mathcal{D} , and $(\mathcal{D}|_S)_{\text{reg}}$ be its regularisation. Then the moduli space $\mathcal{M}_{\text{reg}} := \mathcal{M}_{\text{reg}}^B(\bar{\mathbf{x}}; S; \bar{\rho}_{\blacktriangle}, \bar{\rho}_{\blacktriangle}; \bar{\sigma})$ defined using the diagram $(\mathcal{D}|_S)_{\text{reg}}$ agrees with the moduli space $\mathcal{M}|_S := \mathcal{M}(\bar{\mathbf{x}}; S; \bar{\rho}_{\blacktriangle}, \bar{\rho}_{\blacktriangle}; \bar{\sigma})$ defined using $\mathcal{D}|_S$.*

Proof This follows closely along the lines of the proof of [Proposition 4.37](#) — one perturbs the cylinders for $\mathcal{D}|_S$ to agree with those for $(\mathcal{D}|_S)_{\text{reg}}$. When $\boldsymbol{\gamma}^i$ and $\boldsymbol{\gamma}^{i+1}$ are small approximations of one another, this identifies $\mathcal{M}|_S$ with a moduli space of curves with boundary contained in Lagrangians $\boldsymbol{\gamma}^i \times e^i$ and $\boldsymbol{\gamma}^i \times e^{i+1}$; applying the removable singularities theorem to extend any u' in this moduli space to a bona fide map in \mathcal{M}_{reg} yields the result. □

4.7.3 Holomorphic anchors and stretching the neck Later on, we will stretch the neck of the almost-complex structures \mathbf{J} along a hypersurface, forcing degenerations of moduli spaces. In this section, we fix notation in preparation for this; although we later largely restrict possible degenerations, we begin here in full generality.

Fix an almost-complex structure J_{\blacktriangle} upon $\Sigma_{\bar{\varepsilon}} \times \mathbb{H}$, where \mathbb{H} denotes the upper-half plane $\{z \in \mathbb{C} \mid \Im(z) \geq 0\}$. Consider the moduli space of maps

$$u_{\blacktriangle} : (S_{\blacktriangle}, \partial S_{\blacktriangle}) \rightarrow (\Sigma_{\bar{\varepsilon}} \times \mathbb{H}, C^1(\varepsilon) \times \mathbb{R}),$$

where S_{\blacktriangle} is an anchor source.

We require these to satisfy:

- (1) u_{\pm} is (j, J_{\pm}) -holomorphic for some almost complex structure j upon S_{\pm} ;
- (2) u_{\pm} is proper;
- (3) u_{\pm} has finite energy;
- (4) u_{\pm} is disjoint from \mathbf{p}, \mathbf{b} ;
- (5) the map $\pi_{\mathbb{H}} \circ u_{\pm}$ is a branched covering;
- (6) u_{\pm} is an embedding;
- (7) At each puncture q of S_{\pm} labelled by ∞ , $\lim_{z \rightarrow q} (\pi_D \circ u)(z) = \infty$.
- (8) At each puncture q of S_{\pm} labelled by a Reeb chord ρ , we have that upon some neighbourhood $\eta(q)$ of q , $(\pi_{\Sigma} \circ u)(\eta)$ is modelled upon $B(\rho)$.
- (9) For every $t \in \mathbb{R}$ and each $j = 1, \dots, h$, we have that $u^{-1}(C_j^1|_{\Sigma \times \{t\}})$ consists of at most one point, and for each $j = h + 1, \dots, g + h - 1$, the set $u^{-1}(C_j^1|_{\Sigma \times \{t\}})$ consists of exactly one point.

From these conditions, it follows that near each puncture q of S_{\pm} labelled by ∞ , the map u_{\pm} is asymptotic to a chord of the form $\theta^{1,2}(q) \times \infty$, such that $\theta^{1,2} = \bigcup_q \theta^{1,2}(q)$ is a generator for the pair $(\gamma_{\bar{e}}^1, H(\gamma_{\bar{e}}^1, \varepsilon/(m+1)))$, or, more transparently, for the pair (δ^1, δ^2) described in the previous subsection. Each such map has an associated homology class B_{\pm} in the relative homology group $H_2(\Sigma_{\bar{e}} \times \mathbb{H}, C^1(\varepsilon))$. We call such curves *holomorphic anchors* and collect them into moduli spaces $\mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm})$. We shall later see that in fact, for the moduli spaces to be nonempty, the generator $\theta^{1,2}$ must be the distinguished generator $\theta_{+}^{1,2}$.

As with all of our other moduli spaces, we can cut these down according to various partitions of the east punctures of S_{\pm} . If P_{\pm} is a set of partitions $P_{\pm} = (P_{\pm,1}, \dots, P_{\pm,h})$, where each $P_{\pm,j}$ is a partition of $E_j(S_{\pm})$, there is a corresponding evaluation map

$$\text{ev} := \prod_{q \in E(S_{\pm})} \text{ev}_q : \mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}) \rightarrow (e^1)^{|E(S_{\pm})|},$$

and we write

$$\mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; P_{\pm})$$

for the moduli space $(\text{ev})^{-1}(\Delta_{P_{\pm}})$. If we are further provided with a splicing σ_{\pm} of the set of punctures of S_{\pm} labelled with a jumping chord, we can define the moduli spaces $\mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; P_{\pm}; \sigma_{\pm})$ for the subset $\text{ev}^{-1}(\Delta_{\sigma_{\pm}}^1, \dots, \Delta_{\sigma_{\pm}}^h)$, where $\text{ev} : \mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; P_{\pm}) \rightarrow e^1$.

It transpires that for particularly simple sources S_{\pm} , the moduli spaces $\mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S)$ are always transversally cut out:

Proposition 4.40 *Let S_{\pm} be a disjoint union of topological discs. Then the moduli spaces $\mathcal{M}_{\pm}(\theta^{1,2}, S_{\pm})$ are always transversally cut out by the $\bar{\partial}$ -equations, for any choice of almost-complex structure J_{\pm} .*

Proof This follows along the lines of [14, Proposition 5.16] — namely, any $u = (u_\Sigma, u_{\mathbb{H}})$ doubles to a pair of maps from a punctured CP^1 to a punctured CP^1 . The first Chern class of these doubled maps is equal to the Maslov index of u_Σ and $u_{\mathbb{H}}$, which is -1 in both cases (it is straightforward to see that each is a small anticlockwise disc). Thus by [16, Lemma 3.3.1], $D\bar{\partial}$ is surjective for this doubled map, which implies surjectivity for the original map — see, for instance, [8, page 158]. \square

These anchors will arise as the result of “stretching the neck” of bordered holomorphic polygons. Explicitly, for every m fix an ε -radius half-disc neighbourhood η_ε^m of the point $\{0\} \in e^1 \subset \partial D_m$. We consider the family J^t obtained from J by inserting a neck of length t along the boundary of each η_ε^m , as described in [10, Section A.2].

Suppose that $\text{ind}(B, S, \vec{\rho}_\Delta, \vec{\rho}_\ddagger, \sigma_\Delta) = 0$, and let

$$\mathcal{M}_{\geq T}^B(x; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J)$$

denote the set $\bigcup_{t \geq T} \mathcal{M}^B(\vec{x}; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$. By a standard transversality argument, for generic choice of J and generic T , the moduli spaces

$$\mathcal{M}_{\geq T}^B := \mathcal{M}_{\geq T}^B(\vec{x}; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$$

are one-dimensional manifolds. Moreover, for small enough ε , the moduli spaces

$$\mathcal{M}_{\geq T, \bar{\varepsilon}, \varepsilon}^{B_{\bar{\varepsilon}}}(\vec{x}; S_{\bar{\varepsilon}}; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$$

defined to be equal to $\bigcup_{t \geq T} \mathcal{M}_{\bar{\varepsilon}, \varepsilon}^{B_{\bar{\varepsilon}}}(\vec{x}; S_{\bar{\varepsilon}}; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$ are also transversally cut out, and form one-manifolds homeomorphic to $\mathcal{M}_{\geq T}^B$ (cf Proposition 4.37). By Gromov compactness, in both instances, there is some T such that for every $t > T$, each constituent moduli space $\mathcal{M}^B(\vec{x}; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$ (resp. $\mathcal{M}_{\bar{\varepsilon}, \varepsilon}^{B_{\bar{\varepsilon}}}(\vec{x}; S_{\bar{\varepsilon}}; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$) is transversally cut out.

The perturbed moduli spaces admit compactifications similar to the aforementioned trees of holomorphic polygons; but there is one more possible type of degeneration which we force in the neck-stretching process. This is an *anchored holomorphic curve*:

Definition 4.41 An *anchored holomorphic polygon* is a pair (u, u_\ddagger) where

$$u \in \mathcal{M}^B(\vec{x}, \theta^{0,1}, \theta^{1,2}, \dots, \theta^{m-1,0}, y; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \sigma_\Delta, \sigma_\ddagger; J)$$

is a bordered holomorphic polygon and $u_\ddagger \in \mathcal{M}_\ddagger(\theta^{1,2}; S_\ddagger; P; \sigma_\ddagger)$ is a holomorphic anchor. (See Figure 14 for a schematic.)

We write

$$\overline{\mathcal{M}}_{\geq T, \bar{\varepsilon}, \varepsilon}^{B_{\bar{\varepsilon}}}(\vec{x}; S_{\bar{\varepsilon}}; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t)$$

for the compactification of

$$\mathcal{M}_{\geq T, \bar{\varepsilon}, \varepsilon}^{B_{\bar{\varepsilon}}}(\vec{x}; S_{\bar{\varepsilon}}; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger; J^t),$$

which can be defined in much the same way as in Definition 4.35, but allowing for the bordered polygons in each definition to be anchored holomorphic polygons.

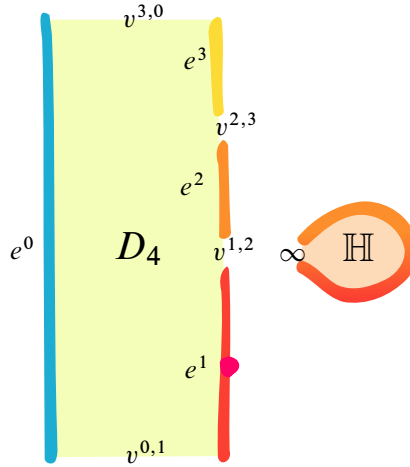


Figure 14: A schematic of an anchored holomorphic polygon. The pink dot indicates $\pi_D \circ u$ of an east puncture of the source S .

4.8 Gluing for trees of elementary holomorphic combs and anchors

We now quote the relevant gluing results, and then explore their consequences.

Proposition 4.42 *Let $U = (u_1, u_2)$ be a height 2 tree of (bordered) holomorphic polygons where $u_1 \in \tilde{\mathcal{M}}^{B_1}(\vec{x}(V_1), S_1; P_1, Q_1; \sigma_1)$ and $u_2 \in \tilde{\mathcal{M}}^{B_2}(\vec{x}(V_2), S_2; P_2, Q_2; \sigma_2)$ are spinal.*

*Let $B = B_1 * B_2$, S denote the splayed source $S_T = S_1 \natural S_2$, with induced distinguished set of noneasterly punctures $Q = Q_1 \cup Q_2$, and finally let $\sigma = \sigma_1 * \sigma_2$.*

*For some sufficiently small open neighbourhoods U_1 of u_1 and U_2 of u_2 , there is an open neighbourhood of (u_1, u_2) in $\bar{\mathcal{M}}^B(\vec{x}(T), S; P_1 * P_2, Q; \sigma)$ which is homeomorphic to $U_1 \times U_2 \times [0, 1)$.*

We also require gluing for two-storey combs with one easterly vertex:

Proposition 4.43 *Let $U = (u_1, u_2)$ be a height 2 tree of (bordered) holomorphic polygons where $u_1 \in \tilde{\mathcal{M}}^{B_1}(\vec{x}(V_1), S_1; P_1, Q; \sigma_1)$ is spinal and $u_2 \in \tilde{\mathcal{M}}^{B_2}(\vec{x}(V_2), S_2; P_2; \sigma_2)$ is easterly.*

*Let $B = B_1 * B_2$, S denote the splayed source $S_T = S_1 \natural S_2$, with induced distinguished set of noneasterly punctures Q' , and finally let $\sigma = \sigma_1 * \sigma_2$.*

*For some sufficiently small open neighbourhoods U_1 of u_1 and U_2 of u_2 , there is an open neighbourhood of (u_1, u_2) in $\bar{\mathcal{M}}^B(\vec{x}(T), S; P_1 * P_2, Q; \sigma)$ which is homeomorphic to $U_1 \times U_2 \times [0, 1)$.*

We also require gluing for simple combs:

Proposition 4.44 *Let (u, v) be a simple m -comb, with $u \in \tilde{\mathcal{M}}^B(\vec{x}; S; Q; J)$, and $v \in \mathcal{N}(T, P_e; \sigma)$, and suppose the corresponding moduli spaces are transversally cut out at u and v respectively, and the evaluation maps $\text{ev}: \tilde{\mathcal{M}}^B(\vec{x}; S; Q; J) \rightarrow e^1$ and $\text{ev}_w: \mathcal{N}(T; P_e; \sigma) \rightarrow e^1$ are transverse at (u, v) .*

If we denote by $S' = S \natural T$, and Q' the corresponding set of distinguished noneasterly punctures induced by S , then for sufficiently small open neighbourhoods U_u of u and U_v of v , there is an open neighbourhood of (u, v) in $\overline{\mathcal{M}}^B(\vec{x}; S'; P, Q; \sigma)$ homeomorphic to $(U_u \times_{\text{ev}} U_v) \times [0, 1)$.

Finally, we require the following result for anchored holomorphic curves:

Proposition 4.45 *Let u' be a splayed holomorphic polygon in*

$$\mathcal{M}_{\vec{e}}^{B'}(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S'; P', Q'; \sigma'; \mathbf{J}),$$

and $u_{\natural} \in \mathcal{M}_{\natural}^{B_{\natural}}(\theta^{1,2}; S_{\natural}; P_{\natural}; \sigma_{\natural}; J_{\natural})$ a holomorphic anchor, and suppose that these moduli spaces are transversally cut out at u and u_{\natural} respectively.

Denote by $B = B' * B_{\natural}$, S the surface formed by gluing S' and S_{\natural} at punctures labelled by $v^{1,2}$, $P = P' \star P_{\natural}$, $\sigma = \sigma' \star \sigma_{\natural}$, and Q the set of distinguished noneast punctures of S induced by Q' . Then, for sufficiently small open neighbourhoods U and U_{\natural} containing u and u_{\natural} respectively, there is an open neighbourhood of (u, u_{\natural}) in the moduli space

$$\overline{\mathcal{M}}_{\geq T}(\mathbf{x}, \vec{\theta}_+(k-1), \mathbf{y}; S; P, Q; \sigma)$$

homeomorphic to $U \times U_{\natural} \times [0, 1)$.

Proof of Propositions 4.42, 4.43, 4.44 and 4.45 A straightforward combination of [14, Proposition 5.30] and [13, Proposition 4.21] gives the first three of these results when dropping the relevant instances of the letter Q from the statements, and the last follows similarly. In each case, as the punctures of each relevant source are mapped either to east infinity or to some boundary puncture of the relevant disc, the fact that the manifolds $X_{\mathcal{D},m}$ have two types of infinity is irrelevant, so the result follows by standard gluing techniques (the provision of splittings behaves no differently than the other partitions of east punctures). To include the provision of cut behaviour, one sees that the nearby families of holomorphic curves constructed by the implicit function theorem in the proof of [14, Proposition 5.30] and [13, Proposition 4.21] must have small cuts near the relevant punctures themselves — thus, by the discussion in Section 4.6.2 there are nearby families of curves with no cut at these punctures.

For example, for Proposition 4.42, one constructs preglued maps $u_1 \natural_R u_2$ with sources $S_1 \natural S_2$ which converge to the two-storey building. One shows that $D\bar{\partial}(u_1 \natural_R u_2) \rightarrow 0$ as $R \rightarrow \infty$, and applies the implicit function theorem to find a neighbourhood $N_{\delta} = N_{\delta}(u_1 \natural_R u_2)$ such that the set of zeroes of the $\bar{\partial}$ -equations in N_{δ} form a smooth manifold of the correct dimension. As such, choosing R sufficiently large finds families of solutions to the $\bar{\partial}$ -equations near to $u_1 \natural_R u_2$.

If S_1 and S_2 have distinguished sets of punctures Q_1 and Q_2 respectively, where $Q_1 \cup Q_2 = Q$, it follows that the preglued maps $u_1 \natural_R u_2$ have no cut at the corresponding punctures of $S_1 \natural S_2$. As such, any holomorphic curve in the neighbourhood N_{δ} has (a priori) at least a very small cut at the relevant punctures. By the discussion in Section 4.6.2, it follows that the set $\bar{\partial}^{-1}(0) \cap N_{\delta}$ can be parameterised



Figure 15: A schematic picture of a one-dimensional moduli space (left), with two types of spinal ends shown (centre and right). The breaks in the boundary indicate punctures of D_n , whereas the dots indicate east punctures — with punctures corresponding to one east infinity shown in grey, and a puncture corresponding to another shown in white.

by cut length as $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_\Delta, Q; \vec{\sigma}) \times (-\varepsilon, \varepsilon)^k$ for some k . It follows by the triangle inequality that the set $\tilde{\mathcal{M}}^B(\vec{x}; S; \vec{\rho}_\Delta, Q; \vec{\sigma})$ intersects B_δ , and thus provides a family of holomorphic curves with the correct cut behaviour near to $u_1 \natural_R u_2$ and thus near to (u_1, u_2) , as in [10, Section A.4] (see also [3, Section 5.3.3]). \square

4.9 Restrictions upon easterly curves and the boundary of one-dimensional moduli spaces

In this section, we prove the main result which we shall use to show that maps we define in the next section satisfy the partial \mathcal{A}_∞ relations. We will need to understand the modulo-two count of the ends of moduli spaces of expected dimension 1 in more detail. As in [14, Definition 5.58], we give names to the types of ends that occur. Some of these are illustrated in Figures 15, 16 and 17 — recall that we distinguished some special types of vertices of a tree in Definition 4.4.




Figure 16: A schematic picture of two types of mixed ends of the one-dimensional moduli space on the left of Figure 15.




Figure 17: A schematic picture of a composition end (left) and a collision end (right) of the moduli space on the left of Figure 15.

Definition 4.46 Let $\mathcal{M} := \mathcal{M}^B(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S; \vec{\rho}_\Delta, \vec{\rho}_\pm; \sigma_\Delta, \sigma_\pm)$ be a one-dimensional moduli space. Then:

(1) A *spinal two-storey end* of \mathcal{M} is a height two spinal tree of prescribed-cut polygons (u, v) , where

$$u \in \mathcal{M}^{B_1}(\mathbf{x}, \vec{\theta}_+, \mathbf{w}; S_1; \vec{\rho}_\Delta^1, \vec{\rho}_\pm^1; \sigma_\Delta^1, \sigma_\pm^1)$$

and

$$v \in \mathcal{M}^{B_2}(\mathbf{w}, \vec{\theta}_+, \mathbf{y}; S_2; \vec{\rho}_\Delta^2, \vec{\rho}_\pm^2; \sigma_\Delta^2, \sigma_\pm^2)$$

are spinal, and such that $B_1 \in \pi_2^{\vec{\theta}_+}(\mathbf{x}, \mathbf{w})$ and $B_2 \in \pi_2^{\vec{\theta}_+}(\mathbf{w}, \mathbf{z})$ are such that $B_1 * B_2 = B$, $S_1 \natural S_2 = S$, $\vec{\rho}_\Delta^1 \star \vec{\rho}_\Delta^2 = \vec{\rho}_\Delta$, $\vec{\rho}_\pm^1 \star \vec{\rho}_\pm^2 = \vec{\rho}_\pm$, $\sigma_\Delta^1 \star \sigma_\Delta^2 = \sigma_\Delta$ and $\sigma_\pm^1 \star \sigma_\pm^2 = \sigma_\pm$.

(2) A *mixed two-storey end* of \mathcal{M} is a height two tree (u, v) such that

$$u \in \mathcal{M}^{B'}(\mathbf{x}, \theta_+^{1,2}, \dots, \theta_+^{\kappa-1,\kappa}, \theta^{\kappa,\kappa+\lambda}, \theta_+^{\kappa+\lambda,\kappa+\lambda+1}, \dots, \theta_+^{\kappa-1,k}, \mathbf{y}; S'; \vec{\rho}'_\Delta, \vec{\rho}_\pm; \sigma'_\Delta, \sigma_\pm)$$

is spinal and

$$v \in \mathcal{M}^{B_E}(\theta_+^{\kappa,\kappa+1}, \dots, \theta_+^{\kappa+\lambda-1,\kappa+\lambda}, \theta^{\kappa+\lambda,\kappa}; S_E; \vec{\rho}_E)$$

is easterly, and such that $B' * B_E = B$ and $S' \natural S_E = S$.

(3) An (i, j) -*composition end* is an element of $\mathcal{M}^B(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S'; \vec{\rho}'_\Delta, \vec{\rho}_\pm; \sigma'_\Delta, \sigma_\pm)$, where $\vec{\rho}'_\Delta$ is given by $\bar{\mu}_j^i(\vec{\rho}_\Delta)$ for some i, j ; σ'_Δ is given by σ_j^i ; and S' is obtained from S by contracting the arc of ∂S connecting the punctures corresponding with ρ_j^i and ρ_j^{i+1} , and labelling the corresponding puncture $\rho_j^i \cdot \rho_j^{i+1}$.

(4) A κ -*collision end* is an element of $\mathcal{M}^B(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S; \vec{\rho}_\Delta, \vec{\rho}_\pm; \sigma_\Delta(\kappa), \sigma_\pm)$.

(5) A *cut-vanishing end* of \mathcal{M} is a prescribed-cut polygon u in the moduli space

$$\mathcal{M}^B(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S'; \vec{\rho}'_\Delta, \vec{\rho}'_\pm; \sigma'_\Delta, \sigma'_\pm)$$

such that $\vec{\rho}'_{\Delta} \star \vec{\rho}'_{\pm}$ is a 123-reduction of $\vec{\rho}_{\Delta} \star \vec{\rho}_{\pm}$ formed by replacing an adjacent pair ρ_1, ρ_{23} or ρ_{12}, ρ_3 with ρ_{123} , where ρ_1 (resp. ρ_3) is an element of $\vec{\rho}_{\pm}$. Here, S' is given by contracting the arc in ∂S between the punctures corresponding with ρ_1, ρ_{23} (resp. ρ_{12}, ρ_3) and labelling the resultant puncture by ρ_{123} .

(The condition on a cut-vanishing end essentially is the same as $\vec{\rho}'_{\pm}$ being a 123-reduction of $\vec{\rho}_{\pm}$, except there is a possibility for the first chord of $\vec{\rho}_{\pm, j}$ to have been ρ_3 and the final chord of $\vec{\rho}_{\pm, j}$ to have been ρ_{12} .)

With this in mind, we state the first main theorem of this section:

Theorem 4.47 *Suppose that*

$$\mathcal{M} = \mathcal{M}^B(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S; \vec{\rho}_{\Delta}, \vec{\rho}_{\pm}; \sigma_{\Delta}, \sigma_{\pm})$$

is one-dimensional. Then the total number of spinal two-storey ends of \mathcal{M} , mixed two-storey ends of \mathcal{M} , (i, j) -composition ends of \mathcal{M} , κ -collision ends of \mathcal{M} , and cut-vanishing ends of \mathcal{M} is even.

Proof This follows readily along the same lines as the proof of [14, Theorem 5.61]. In particular, trees of height greater than two are ruled out in a very similar way to [14, Proposition 5.43]. In our situation, however, there are two types of two-storey ends. Height one combs still degenerate, but due to the simplicity of the diagram near the boundary, join or shuffle curves cannot appear, just sequences of split curves (corresponding with composition and collision ends of \mathcal{M}). A priori, nodal curves can still degenerate, but are ruled out along the lines of [14, Section 5.6.3].

One may truncate $\bar{\mathcal{M}}$ near composition and collision ends, as in the proof of [14, Theorem 5.61] to give a moduli space $\bar{\mathcal{M}}_{\text{cropped}}$ which is a compact 1-manifold with boundary. This has spinal two-storey ends if the complex structure on D_{k+2} degenerates along an arc with one end on the edge e^0 , mixed two-storey ends if an arc without an end upon e^0 degenerates, a composition or collision end if a holomorphic tooth degenerates at east infinity, and finally a cut-vanishing end if the length of a cut at some corner of $\rho_{123}(S)$ degenerates to zero. Ends where more than one cut degenerates, or a cut degenerates at the same time as another type of degeneration, are ruled out by the index formula in Proposition 4.28 — each degenerated length subtracts one from the expected dimension; thus if more than one cut degenerated the corresponding moduli spaces would have negative dimension and so be empty. □

We now examine the types of mixed two-storey ends which may occur. Fix a partially splayed diagram $\mathcal{D} = \mathcal{D}'(\varepsilon, \vec{\iota})$ with attaching curves $\boldsymbol{\gamma}^0, \boldsymbol{\gamma}^1, \dots, \boldsymbol{\gamma}^m$, where we recall that $\boldsymbol{\gamma}^1$ splits as $\boldsymbol{\gamma}^1 = \boldsymbol{\gamma}^{1,a} \cup \boldsymbol{\gamma}^{1,c}$. Fix also a subdiagram $\mathcal{D}|_S$ where S is a contiguous subsequence of $\{1, \dots, m\}$. For clarity's sake, relabel the attaching curves of this subdiagram by $\delta^1, \dots, \delta^\lambda$, and the corresponding sequence of generators by $(\boldsymbol{\theta}_+^{\kappa, \kappa+1}, \dots, \boldsymbol{\theta}_+^{\kappa+\lambda-1, \kappa+\lambda}) := \vec{\boldsymbol{\theta}}_+(\lambda)$ and $\boldsymbol{\theta} = \boldsymbol{\theta}^{\kappa+\lambda, \kappa} \in \mathcal{G}(\delta^\lambda, \delta^1)$. We call the resultant generator for $\mathcal{D}|_S$ by $\vec{\boldsymbol{\theta}}(\lambda) := (\vec{\boldsymbol{\theta}}_+(\lambda), \boldsymbol{\theta})$. We first restrict B_E to some fairly manageable possibilities:

Proposition 4.48 *Let $B_E \in \pi_2(\vec{\boldsymbol{\theta}}(\lambda)) \subset \text{Dom}(\mathcal{D}|_S)$. Then it is supported within the approximation region for the curves $\delta^1, \dots, \delta^\lambda$.*

Proof This follows from the homological linear independence of the bordered diagram \mathcal{D}' . The corresponding domain $B'_E \in \text{Dom}(\mathcal{D}')$ given by taking multiplicities away from the approximation region has boundary contained solely in α and multiplicity zero at \mathbf{b} — contradicting homological linear independence unless all multiplicities away from the approximation region are zero. \square

It follows from this that the domain B_E decomposes as a sum $B_E = B_1 + \dots + B_{g+h-1}$, where each B_j is supported in the approximation region for $\delta_j^i, \delta_j^{i+1}$.

The next step is to separate the curves into the more widely studied case where they are each approximations of one another, and the case where they change in a given index. As such, we let

$$\Gamma := \{j \in \{1, \dots, h\} \mid \text{jump}(\vec{\rho}_{E,j}) = \emptyset \text{ and } \delta_j^1 \approx \dots \approx \delta_j^l\} \cup \{h+1, \dots, g+h-1\},$$

and $\bar{\Gamma} := \{1, \dots, h\} - (\Gamma \cap \{1, \dots, h\})$.

We now show that the domains of the less widely studied case are supported within the handles.

Proposition 4.49 *Suppose that $B_E \in \pi_2(\vec{\theta}(\lambda))$, and let $j \in \bar{\Gamma}$. Then the domain B_j is supported within the handle \mathfrak{H}_j .*

Proof Any domain B_j not supported within \mathfrak{H}_j but supported within the approximation region can only have corners upon points of the form $\theta_{+,j}^{i_k, i_{k+1}}$ where $\delta_j^{i_k} \approx \delta_j^{i_{k+1}}$. As $j \in \bar{\Gamma}$, either one coordinate $\theta(\lambda)_j^i$ is equal to some $e_j^{\infty, i}$, or B_j is required to have a corner at e_j^∞ — neither of which are of this form. \square

For each $j \in \Gamma$, then, if δ^1 is a set of $2h$ arcs, any map $u \in \mathcal{M}^{B_E}(\theta(\lambda); S_E; \vec{\rho}_E)$ has boundary in just one of the cylinders $\delta_j^{1,m} \times e^1$ or $\delta_j^{1,l} \times e^1$. We label the corresponding curve δ_j^1 . Let $\dot{\Sigma}_\Gamma$ denote the surface given by removing the handles $\bigcup_{j \in \bar{\Gamma}} \mathfrak{H}_j$ from the east-compactified splayed surface $\Sigma(\varepsilon, \vec{l})_{\bar{e}}$, and Σ_Γ the result of collapsing the remaining boundary components $\mathfrak{H}_{L,j}$ and $\mathfrak{H}_{R,j}$ in $\dot{\Sigma}_\Gamma$ to $2|\bar{\Gamma}|$ single points $\mathbf{p}_{\bar{\Gamma}}$. This contains the curves $\delta_\Gamma^i = \{\delta_{e,j}^i \mid j \in \Gamma\}$, which approximate one another and basepoints $\mathbf{p}_\Gamma := \{p_j \mid j \in \Gamma\}$ and thus we form the closed multidigraph \mathcal{D}_Γ given by the tuple $(\Sigma_\Gamma; \delta_\Gamma^1, \dots, \delta_\Gamma^l; \mathbf{p}_\Gamma \cup \mathbf{p}_{\bar{\Gamma}})$.

It follows from the above proposition that for $j \in \Gamma$ the domains B_j are contained in $\dot{\Sigma}_\Gamma$, and thus have a corresponding domain $B_\Gamma \in \text{Dom}(\mathcal{D}_\Gamma)$, given by the sum of their images in Σ_Γ . Moreover, the domain B_Γ lives in the homology group $\pi_2(\vec{\theta}(\lambda)|_\Gamma)$, where $\vec{\theta}(\lambda)|_\Gamma$ is the restriction of $\vec{\theta}(\lambda)$ to the curves $\delta_\Gamma^1, \dots, \delta_\Gamma^l$.

Considering the surface $\dot{\Sigma}_\Gamma$ as a subset of both surfaces $\Sigma(\varepsilon, \vec{l})_{\bar{e}}$ and Σ_Γ , we choose a family of almost-complex structures \mathbf{J}' upon $\Sigma_\Gamma \times D_l$ such that their restriction to $\dot{\Sigma}_\Gamma \times D_l$ agrees with the restriction of \mathbf{J} .

For every $j \in \bar{\Gamma}$, we identify the handle \mathfrak{H}_j with the punctured torus $\Sigma_j := T^2 - e_j^\infty$ formed by gluing the ends $\mathfrak{H}_{L,j}$ and $\mathfrak{H}_{R,j}$ (more properly, we identify the handle with its end circles removed with an open subset of the torus). The curves and arcs $\delta_j^1 \cap \mathfrak{H}_j, \dots, \delta_j^l \cap \mathfrak{H}_j$ have natural images in T^2 , for

which we abuse notation and call $\delta_j^1, \dots, \delta_j^\lambda$. Together with the image of the basepoint b_j , the tuple $\mathcal{D}_j := (\Sigma_j; \delta_j^1, \dots, \delta_j^\lambda; b_j)$ forms a toroidal bordered Heegaard multidiagram. The domain B_j has a natural image in $\text{Dom}(\mathcal{D}_j)$ which we also denote by B_j , and lives in the group $\pi_2(\theta(\lambda)_j^1, \dots, \theta(\lambda)_j^\lambda)$. We choose a family of almost-complex structures \mathbf{J}_j upon $\Sigma_j \times D_\lambda$ which agree with the almost-complex structure \mathbf{J} upon $\mathcal{H}_j \times D_\lambda$.

It follows that the moduli space $\mathcal{M}^{B_E}(\vec{\theta}(\lambda); S_E; \vec{\rho}_\Delta)$ splits as a fibred product

$$\mathcal{M}^{B_E}(\vec{\theta}(\lambda); S_E; \vec{\rho}_E) = \mathcal{M}_\Gamma \times_{\text{Conf}(D_\lambda)} \mathcal{M}_1 \times_{\text{Conf}(D_\lambda)} \cdots \times_{\text{Conf}(D_\lambda)} \mathcal{M}_{|\bar{\Gamma}|},$$

where $\mathcal{M}_\Gamma = \mathcal{M}^{B_\Gamma}(\vec{\theta}(\lambda)|_\Gamma; S_\Gamma; \mathbf{J}')$ is the moduli space defined by counting holomorphic polygons according to \mathcal{D}_Γ as defined in Section 4.5, and each \mathcal{M}_j is the moduli space

$$\mathcal{M}^{B_j}(\theta(\lambda)_j^{1,2}, \dots, \theta(\lambda)_j^{\lambda,1}; S_j; \vec{\rho}_{\Delta,j}; \mathbf{J}_j)$$

defined by counting bordered polygons according to \mathcal{D}_j , as defined in Section 4.6.

Proposition 4.50 *There is a unique homology class $D_\Gamma \in \text{Dom}(\mathcal{D}_\Gamma)$ which carries a holomorphic representative in the moduli space \mathcal{M}_Γ . This lives in the homology group $\pi_2(\vec{\theta}_+(\lambda)|_\Gamma)$. Moreover, the projection of the moduli space \mathcal{M}_Γ onto $\text{Conf}(D_\lambda)$ has degree one for this homology class; thus $\mathcal{M}^{B_E}(\theta(\lambda); S_E; \vec{\rho}_\Delta)$ is nonempty only if the corresponding $B_\Gamma = D_\Gamma$, and $\vec{\theta}(\lambda)|_\Gamma = \vec{\theta}_+(\lambda)|_\Gamma$. In this case, $\mathcal{M}^{B_E}(\vec{\theta}(\lambda); S_E; \vec{\rho}_E)$ has the same number of elements as*

$$\mathcal{M}_1 \times_{\text{Conf}(D_\lambda)} \cdots \times_{\text{Conf}(D_\lambda)} \mathcal{M}_{|\bar{\Gamma}|}.$$

Proof This is well known; see, for example, [13, Lemma 3.50] or [15, Lemma 6.17]. □

Proposition 4.51 *Let B_j, S_j and $\vec{\rho}_{\Delta,j}$ be such that $\text{ind}(B_j, S_j, \vec{\rho}_{\Delta,j}) = 2 - \lambda$. For each j , the moduli space $\mathcal{M}^{B_j}(\theta(\lambda)_j^{1,2}, \dots, \theta(\lambda)_j^{\lambda,1}; S_j; \vec{\rho}_{\Delta,j}; \mathbf{J}_j)$ (of expected dimension zero) is nonzero only if either*

- $\lambda = 2$ and $\vec{\rho}_{\Delta,j}$ contains a single jumping chord, or
- $\lambda = 3$ and $\vec{\rho}_{\Delta,j}$ is empty.

In these cases there is only a single domain for which the moduli space is nonempty, and in these cases it contains a unique element.

Proof We assume that $\vec{\rho}_{\Delta,j}$ is empty — the case where it is nonempty follows similarly, requiring only the minor modification of changing which index formulae are used. Suppose initially that $\lambda > 3$.

If, for some (noncyclically) consecutive $i, i + 1$, the curves $\delta_j^i, \delta_j^{i+1}$ are approximations of one another, then by Proposition 4.45 and [10, Lemma 11.8] (cf Proposition 4.61 below), there is a suitable holomorphic anchor which may be glued to any curve $u \in \mathcal{M}^{B_j}(\theta(\lambda)_j^{1,2}, \dots, \theta(\lambda)_j^{\lambda,1}; S_j; \vec{\rho}_{\Delta,j})$ of index $2 - \lambda$ to yield an element u' of the moduli space

$$\mathcal{M}' = \mathcal{M}^{B'_j}(\theta(\lambda)_j^{1,2}, \dots, \theta(\lambda)^{i-1,i}, \theta(\lambda)^{i+1,i+2}, \dots, \theta(\lambda)_j^{\lambda,1}; S'_j; \vec{\rho}_{\Delta,j})$$




Figure 18: Ruling out $\lambda > 3$. On the left is shown the case where $\lambda = 4$, and the only possible domain connecting the intersection points shown has the incorrect orientation. On the right is shown the case where $\lambda = 5$: the domain shown has index -3 ; thus the corresponding moduli space is of negative expected dimension.

where the domain B'_j has Euler measure $e(B'_j) = e(B_j) - 1/4$. It follows from the index formula that

$$\text{ind}(B'_j, S'_j, \vec{\rho}_{\Delta,j}) = \frac{3 - (\lambda - 1)}{2} - \chi(S'_j) + 2e(B'_j) = \text{ind}(B_j, S_j, \vec{\rho}_{\Delta,j}),$$

which implies that the moduli space \mathcal{M}' has negative expected dimension. As we are working with a generic family of almost-complex structures, this implies that such a u' (and thus such a u) does not exist — thus if any (noncyclically) adjacent $\delta_j^i, \delta_j^{i+1}$ approximate one another, the corresponding moduli space is empty.

So we may assume that the family $\delta_j^1, \dots, \delta_j^\lambda$ satisfies:

- (1) For every $i = 1, \dots, \lambda - 1$, δ_j^i and δ_j^{i+1} intersect transversally in a single point.
- (2) For every $i = 1, \dots, \lambda - 2$, the curves $\delta_j^i, \delta_j^{i+2}$ are small approximations of one another.
- (3) If λ is even, then the curves $\delta_j^\lambda, \delta_j^1$ intersect transversally in a single point; if λ is odd, then the curves $\delta_j^\lambda, \delta_j^1$ are small approximations of one another.

If λ is even, then it follows that $\theta_j \in \delta_j^\lambda \cap \delta_j^1$ is uniquely determined and it is straightforward to see that the corresponding homology group is empty (see Figure 18 for a picture when $\lambda = 4$).

If $\lambda > 3$ is odd, then θ_j is also uniquely determined for nonempty moduli spaces; if $\theta_j = \theta_{j,+}$ then a similar gluing argument to above shows that the corresponding moduli space is empty. We can see that for general odd λ , there is a unique element of $\pi_2(\theta(\lambda)_j^{1,2}, \dots, \theta(\lambda)_j^{\lambda,1})$ — an embedded λ -gon with $\frac{1}{2}(\lambda + 3)$ acute corners and $\frac{1}{2}(\lambda - 3)$ obtuse corners. It follows from the index formula that the expected




Figure 19: Three possible configurations for the curves, and the corresponding domains are shown (assuming that $\lambda = 3$) — the other three are obtained by rotating these diagrams by 90 degrees clockwise. On the left is a (ρ_1, ρ_2) -composition piece, and on the right are two collision pieces, a precollision and postcollision respectively. Rotating these gives a (ρ_2, ρ_3) -composition piece and two collision pieces respectively.

dimension of $\mathcal{M}_e^{B_j}((\theta(\lambda)_j^{1,2}, \dots, \theta(\lambda)_j^{\lambda,1}; S_j)$ is equal to $\frac{1}{2}(3 - \lambda)$. As such, this domain has index $3 - \lambda$ if and only if $\lambda = 3$, as claimed in the proposition. (See Figure 18 for the case where $\lambda = 5$.)

There are six different possible configurations for the curves, with corresponding unique domain — as shown in Figure 19. □

We give things names:

Definition 4.52 We separate the possible domains near e_j^∞ described in Proposition 4.51 into types, as indicated in Figure 19: (ρ_1, ρ_2) -composition pieces, (ρ_2, ρ_3) -composition pieces and collision pieces.

Suppose that $\lambda = 3$ and $B_E \in \pi_2(\vec{\theta}(\lambda))$ is a domain with no Reeb chords. We say B_E is a (ρ_1, ρ_2, j) -composition domain (resp. (ρ_2, ρ_3, j) -composition domain) if the set $\bar{\Gamma}$ has two elements, j and j' , and B_j is a (ρ_1, ρ_2) -composition piece (resp. (ρ_2, ρ_3) -composition piece) and $B_{j'}$ is a collision piece.

Similarly, we say B_E is a collision domain if $\bar{\Gamma}$ has two elements, j and j' , and B_j and $B_{j'}$ are both collision pieces.

A similar definition may be made for the case $\lambda = 2$, where the corresponding domain has at most one jumping Reeb chord upon each boundary component.

It follows almost immediately that these are all of the possibilities for easterly domains:

Proposition 4.53 Suppose B_E is such that $\mathcal{M}^{B_E}(\vec{\theta}(\lambda); S; \vec{\rho}_\Delta)$ is nonempty. Then $\lambda = 3$ or 2 , and either

- B_E is a (ρ_1, ρ_2, j) -composition domain for some j , and $\theta_j^{\kappa, \kappa+\lambda} = \theta_{-,j}^{\kappa, \kappa+\lambda}$ and $\theta_{j+1}^{\kappa, \kappa+\lambda} = \theta_{+,j+1}^{\kappa, \kappa+\lambda}$;
- B_E is a (ρ_2, ρ_3, j) -composition domain for some j , and $\theta_j^{\kappa, \kappa+\lambda} = \theta_{-,j}^{\kappa, \kappa+\lambda}$ and $\theta_{j+1}^{\kappa, \kappa+\lambda} = \theta_{+,j+1}^{\kappa, \kappa+\lambda}$;
- B_E is a collision domain, and we have that $\theta_j^{\kappa, \kappa+\lambda} = \theta_{+,j}^{\kappa, \kappa+\lambda}$ and $\theta_{j+1}^{\kappa, \kappa+\lambda} = \theta_{+,j+1}^{\kappa, \kappa+\lambda}$ for some j .

(and similar for $\lambda = 2$). In these cases, $\mathcal{M}^{B_E}(\vec{\theta}(\lambda); S; \vec{\rho}_\Delta)$ contains a unique holomorphic curve.

Proof By Proposition 4.50, the domain is uniquely determined by its restriction to handles \mathfrak{H}_j where $j \in \bar{\Gamma}$, and by Proposition 4.51 we know that $\lambda = 3$ or $\lambda = 2$.

As we assumed $\sigma_{\Delta} \star \sigma_{\ddagger}$ was interleaved, the set $\bar{\Gamma}$ having more than two elements is impossible: only one curve in δ^2 is not an approximation of one in δ^1 , and only one curve in δ^3 is not an approximation of one in δ^2 . As such, we only need check the possibilities for B_j near each element j, j' of $\bar{\Gamma}$. There are two cases which may occur for this: either (without loss of generality), B_j is a composition piece and $B_{j'}$ is forced to be a collision piece (else too many curves from δ^1 are not approximations of curves from δ^2); or B_j is a precollision piece, forcing $B_{j'}$ to be a postcollision piece (for the same reason). The case where one piece is a composition piece splits depending upon whether δ_j^1 is meridional or longitudinal.

A similar argument applies for the case where $\lambda = 2$. □

We use this to give names to types of mixed two-storey trees of holomorphic curves:

Definition 4.54 Let

$$\mathcal{M} := \mathcal{M}^B(\mathbf{x}, \vec{\theta}_+(k), \mathbf{y}; S; \vec{\rho}_{\Delta}, \vec{\rho}_{\ddagger}; \sigma_{\Delta}, \sigma_{\ddagger})$$

be a one-dimensional moduli space, and $U = (u, v)$ be a height two tree of prescribed-cut polygons in a mixed end of the boundary of $\bar{\mathcal{M}}$, where

$$u \in \mathcal{M}^{B'}(\mathbf{x}, \vec{\theta}_+(k - \lambda), \mathbf{w}; S'; \vec{\rho}_{\Delta}, \vec{\rho}'_{\ddagger}; \sigma_{\Delta}, \sigma'_{\ddagger})$$

is spinal and $v \in \mathcal{M}^{BE}(\vec{\theta}(\lambda); S_E; \vec{\rho}_E)$ is easterly.

We shall say that U is

- a (κ, j) -composition end if B_E is a (ρ_1, ρ_2, j) -composition or (ρ_2, ρ_3, j) -composition domain, or
- a κ -collision end if B_E is a collision domain.

By Proposition 4.53, any mixed $U \in \partial\mathcal{M}$ is either an (i, j) -composition end or a κ -collision end.

We conclude this section by identifying composition and collision ends with recognisable moduli spaces:

Proposition 4.55 Suppose that

$$\mathcal{M}' := \mathcal{M}^{B'}(\mathbf{x}, \theta_+^{1,2}, \dots, \theta_+^{\kappa-1,\kappa}, \theta_+^{\kappa,\kappa+3}, \theta_+^{\kappa+3,\kappa+4}, \dots, \theta_+^{\kappa-1,k}, \mathbf{y}; S'; \vec{\rho}_{\Delta}, \vec{\rho}'_{\ddagger}; \sigma_{\Delta}, \sigma'_{\ddagger}),$$

$$\mathcal{M}_E := \mathcal{M}^{BE}(\theta_+^{\kappa,\kappa+1}, \theta_+^{\kappa+1,\kappa+2}, \theta_+^{\kappa+2,\kappa}; S_E; \vec{\rho}_E),$$

are such that $\mathcal{M}' \times \mathcal{M}_E \subset \partial\bar{\mathcal{M}}$.

If B_E is a (ρ_1, ρ_2, j) -composition domain (resp. (ρ_2, ρ_3, j) -composition domain), then $\mathcal{M}' \times \mathcal{M}_E$ contains the same number of elements modulo two as the moduli space

$$\mathcal{M}_{\text{comp}} := \mathcal{M}^{B'}(\mathbf{x}, \vec{\theta}_+(k - 1), \mathbf{y}; S'; \vec{\rho}_{\Delta}, \vec{\mu}_j^{\kappa}(\vec{\rho}_{\ddagger}); \sigma_{\Delta}; \sigma_{\ddagger}, (\sigma_{\Delta})^{\kappa}).$$

Similarly, if B_E is a collision domain, then the moduli space $\mathcal{M}' \times \mathcal{M}_E$ contains the same number of elements modulo two as

$$\mathcal{M}_{\text{coll}} := \mathcal{M}^{B'}(\mathbf{x}, \vec{\theta}_+(k - 1), \mathbf{y}; S'; \vec{\rho}_{\Delta}, \vec{\rho}_{\ddagger}; \sigma_{\Delta}, \sigma_{\ddagger}(\kappa)).$$

Proof By Proposition 4.53, the projection of each $\mathcal{M}' \times \mathcal{M}_E$ onto \mathcal{M}' has degree one; hence we only have to show that, in each case, \mathcal{M}' has the same number of elements as the moduli spaces in question.

By Proposition 4.39, as the diagram $\mathcal{D}_{\text{comp}}$ used for the definition of $\mathcal{M}_{\text{comp}}$ (resp. $\mathcal{D}_{\text{coll}}$ and $\mathcal{M}_{\text{coll}}$) is the regularisation of the diagram \mathcal{D}' used in the definition of \mathcal{M}' , we know that if B_E is a composition domain, \mathcal{M}' is identified with a subset of the moduli space

$$\mathcal{M}^{B'}(x, \theta_+^{1,2}, \dots, \theta_+^{\kappa-1,\kappa}, \theta_+^{\kappa+3,\kappa+4}, \dots, \theta_+^{\kappa-1,\kappa}, y; S'; \vec{\rho}_\Delta, Q; (\sigma_\Delta)_j^\kappa),$$

and it remains to be seen that this subset is precisely $\mathcal{M}_{\text{comp}}$.

Assume first that B_E is a (ρ_1, ρ_2, j) -composition domain, and suppose that for each $j' = 1, \dots, h$, the sequences

$$R_{j'} := (D_{j'}^1, B_{j'}^1, D_{j'}^2, B_{j'}^2, \dots, D_{j'}^k, B_{j'}^k, D_{j'}^{k+1})$$

form a sequence of Reeb domains for B . For every $j' \neq j$, the corresponding domain $B_{j'}$ (as discussed in Section 3.1.1) is equal to $B_{j'}$ — it follows that in this case $R_{j'}$ is also a sequence of Reeb domains for B_j . For $j' = j$, we have that B_j^κ decomposes as the sum $C_j^\kappa + \widehat{B}_j^\kappa$, where C_j^κ is a (ρ_1, ρ_2) -composition domain. It follows that, by definition, $B_j^\kappa = B_j^\kappa(\rho_1)$, and that $\widehat{B}_j^\kappa = B_j^\kappa(\rho_1) - C_j^\kappa$ has multiplicity one in the south-west region adjacent to $e_j^{\infty, \kappa+1}$.

The domain B' has no corner at the intersection between γ_j^κ and $\gamma_j^{\kappa+1}$ — it has the same multiplicity upon either side of $\gamma_j^{\kappa+1}$. As such, $\widehat{B}_j^\kappa + B_j^{\kappa+1}$ has no corners at $e_j^{\infty, \kappa+1}$ and so $B_j^{\kappa+1} = B_j^{\kappa+1}(\rho_2)$ — and we have that $D_j^{\kappa+1} = D_j^{\kappa+1}(\emptyset)$. It follows (by comparing multiplicities away from the approximation region between γ_j^κ and $\gamma_j^{\kappa+2}$) that if $D_j^i := n_i B^i(\rho_i)$, $B_j^i := B^i(\rho'_i)$ and we denote by

$$E_j^i := \begin{cases} n_i B^i(\rho_i), & \text{for } i = 1, \dots, \kappa - 1, \\ (n_\kappa + n_{\kappa+2} + 1) B^i(\rho_i), & \text{for } i = \kappa, \\ n_{i+2} B^i(\rho_{i+2}), & \text{for } i = \kappa + 1, \dots, k - 1, \end{cases}$$

and

$$F_j^i := \begin{cases} B^i(\rho'_i) & \text{for } i = 1, \dots, \kappa - 1, \\ B^i(\rho'_{i+2}) & \text{for } i = \kappa, \dots, k - 2, \end{cases}$$

then

$$(E_j^1, F_j^1, E_j^2, \dots, E_j^{k-2}, F_j^{k-2}, E_j^{k-1})$$

is a sequence of Reeb domains for B'_j and the result follows. Similar holds for B_E a (ρ_2, ρ_3, j) -composition domain.

If B_E is a collision domain then it is completely supported within the approximation regions between γ^l and γ_k^{l+1} . As such, subtracting the domain B_E has no effect upon the sequence $\vec{\rho}_\Delta(B)$, and thus $\vec{\rho}_\Delta(B'') = \vec{\rho}_\Delta(B)$. The jumping chords previously in partitions P_j^κ and $P_{j'}^{\kappa+1}$ are now determined by corners at the copies of east infinity, e_j^κ and $e_{j'}^\kappa$, and thus the corresponding partition is given by $\sigma_\Delta(\kappa)$, as required. □

4.10 Behaviour of holomorphic anchors and stretching the neck

In this section, we will try and understand the moduli spaces $\mathcal{M}^{B_{\pm}}(\theta^{1,2}; S_{\pm})$.

Proposition 4.56 *Suppose that $(u_n) \in \mathcal{M}_{\geq T, \bar{e}, \varepsilon}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; P, Q; \sigma_{\Delta}, \sigma_{\pm}; J')$, is a sequence of splayed holomorphic polygons, converging to some two-storey building*

$$(u', u_{\pm}) \in \mathcal{M}^{B'}(x, \theta^{0,1}, \dots, \theta^{m-1,0}, y; S'; P', Q', \sigma'_{\Delta}, \bar{\sigma}'_{\pm}; J) \times \mathcal{M}^{B_{\pm}}(\bar{\theta}(m_{\pm}); S_{\pm}; P_{\pm}; \sigma).$$

Then u_{\pm} is a holomorphic anchor.

Proof For notational convenience, we let $g' = g + h - 1$. We note that $\chi(S') + \chi(S_{\pm}) - g' = \chi(S)$ and $m_{\pm} + m' - 1 = m$, so $5 - m = 6 - (m_{\pm} + m')$. Denoting by $e(B_{\pm}) := e(B) - e(B')$, $\text{Col}_{\pm} := \text{Col}(\sigma_{\Delta}) - \text{Col}(\sigma'_{\Delta})$ and $\#_{\pm} := |Q| - |Q'|$, we have that

$$\begin{aligned} \text{ind}(B, S, P, Q, \sigma_{\Delta}) &= \frac{1}{2}(3 - m)g' - \chi(S) + 2e(B) + |P| - \text{Col}(\sigma_{\Delta}) - |Q| \\ &= \frac{1}{2}(3 - m)g' - (\chi(S') + \chi(S_{\pm}) - g') + 2e(B) + |P| - \text{Col}(\sigma_{\Delta}) - |Q| \\ &= \frac{1}{2}(5 - m)g' - \chi(S') - \chi(S_{\pm}) + 2e(B) + |P| - \text{Col}(\sigma_{\Delta}) - |Q| \\ &= \frac{1}{2}(3 - m')g' - \chi(S') + 2e(B') + |P'| - \text{Col}(\sigma'_{\Delta}) - |Q'| \\ &\quad + \frac{1}{2}(3 - m_{\pm})g' - \chi(S_{\pm}) + 2e(B_{\pm}) + |P_{\pm}| - \text{Col}_{\pm} - \#_{\pm} \\ &= \text{ind}(B', S', P', Q', \sigma'_{\Delta}) + \frac{1}{2}(3 - m_{\pm})g' - \chi(S_{\pm}) + 2e(B_{\pm}) + |P_{\pm}| - \text{Col}_{\pm} - \#_{\pm}. \end{aligned}$$

Denote by ind' the first term in this final sum, and ind_{\pm} the remainder. If u_{\pm} were not an anchor (ie $m_{\pm} > 1$), by [Proposition 4.37](#) we can unperturb the Lagrangian $C^i(\varepsilon)$ to give a corresponding bordered holomorphic polygon u'_{\pm} which lives in a moduli space of curves of index ind'_{\pm} .

It follows that $2 - m - \text{ind}' = \text{ind}_{\pm}$. Noting that $\text{ind}' \geq 2 - m'$, it follows that

$$\text{ind}_{\pm} \leq 2 - m - (2 - m') = m' - m = 1 - m_{\pm},$$

and so u'_{\pm} lives in a moduli space of negative expected dimension, so does not exist — and thus nor does u_{\pm} . Therefore, u_{\pm} is a holomorphic anchor as required. □

Proposition 4.57 *The domain B_{\pm} is supported within the approximation region for the curves, and B' splays B .*

Proof The first half of the statement follows along similar lines to [Proposition 4.48](#). If we fix points away from the approximation region between γ^1 and γ^2 in the diagram corresponding with the moduli space for u' , then each curve u_n has the same multiplicity at these points and thus so does u' — so B' has the same multiplicity as B everywhere away from this approximation region, which is precisely the definition of B' splaying B . □

We need a lemma about the acute corners of anchored curves:

Lemma 4.58 *Suppose that*

$$(u', u_{\pm}) \in \mathcal{M}^{B'}(x, \theta^{0,1}, \dots, \theta^{m-1,0}, y; S'; P', Q', \sigma'_{\pm}, \bar{\sigma}'_{\pm}; J) \times \mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; P_{\pm}, \sigma)$$

is an anchored holomorphic curve in the boundary of the moduli space

$$\mathcal{M}_{\geq T, \bar{e}, \varepsilon}^{B_{\bar{e}}}(\bar{x}; S_{\bar{e}}; P, Q; \sigma_{\pm}, \sigma_{\pm}).$$

If A denotes the number of acute corners of B' at the points $e^{\infty,1}$, then the difference between the number of acute corners of B and B' satisfies

$$\text{ac}(B') - \text{ac}(B) \geq A - 2(|P| - |P'|),$$

with equality holding only if every component of S_{\pm} has at most one east puncture.

Proof The acute corners of B and B' consist of two types: *east acute corners* and *main acute corners*. Here, a main acute corner is an acute corner corresponding with a noneast puncture of S (resp. S'), and east acute corners correspond with intersections of arcs in \mathcal{Y}^1 with $\partial\Sigma$. The main corners of B' split into two sets: those A at $e^{\infty,1}$ and all others. As B' approximates B , it follows that the number of main acute corners of B' not at $e^{\infty,1}$ is precisely the number of main acute corners of B . As such,

$$\text{ac}(B') - \text{ac}(B) = A + A'_E - A_E,$$

where A'_E and A_E are the number of easterly acute corners of B' and B respectively.

The quantities A_E and A'_E may be read from the multiplicity of B and B' in a neighbourhood of each e_j^{∞} . Indeed, each puncture in $E_j(S)$ (resp. S') contributes two acute corners (at the start and end of the corresponding Reeb chord) unless there is some other puncture labelled with the subsequent chord. If we choose a small neighbourhood N_j of east infinity and three points p_1, p_2, p_3 in the corresponding regions of $N_j - \mathcal{Y}_j^1$, this translates as

$$\begin{aligned} A_E &= n_{p_1}(B) + |n_{p_1}(B) - n_{p_2}(B)| + |n_{p_2}(B) - n_{p_3}(B)| + n_{p_3}(B), \\ A'_E &= n_{p_1}(B') + |n_{p_1}(B') - n_{p_2}(B')| + |n_{p_2}(B') - n_{p_3}(B')| + n_{p_3}(B'). \end{aligned}$$

We extend this to B_{\pm} , setting

$$A_{\pm, E} := n_{p_1}(B_{\pm}) + |n_{p_1}(B_{\pm}) - n_{p_2}(B_{\pm})| + |n_{p_2}(B_{\pm}) - n_{p_3}(B_{\pm})| + n_{p_3}(B_{\pm}).$$

This final quantity is less than or equal to $2 \cdot |E_j(S_{\pm})|$, with equality holding only if $|E_j(S_{\pm})| = 1$. By the triangle inequality, $A'_E + A_{\pm, E} \geq A_E$ and hence $A_E - A'_E \leq A_{\pm, E} \leq 2 \cdot |E_j(S_{\pm})|$, with equality holding only if $|E_j(S_{\pm})| = 1$.

Summing over all j , it follows that

$$\text{ac}(B') - \text{ac}(B) = A + A'_E - A_E \geq A - 2(|P| - |P'|)$$

as required; with equality only if $|E_j(S_{\pm})| = 1$ for each j . □

Proposition 4.59 Suppose that $u_n \in \mathcal{M}_{\geq T}^B(x; S; P, Q; \sigma_{\triangle}, \vec{\sigma}_{\triangle})$ is a sequence of splayed holomorphic polygons, converging to some anchored holomorphic curve

$$(u', u_{\triangle}) \in \mathcal{M}^{B'}(x, \theta^{0,1}, \dots, \theta^{m-1,0}, y; S'; P', Q', \sigma'_{\triangle}, \sigma'_{\triangle}; \mathbf{J}) \times \mathcal{M}_{\triangle}^{B_{\triangle}}(\theta^{1,2}; S_{\triangle}; P_{\triangle}, \sigma).$$

Then S_{\triangle} a disjoint union of topological discs, where $|E_j(S_{\triangle})| \leq 1$ for every j , and either $A = 1$ and Q' consists of the punctures of S' induced by Q , or $A = 0$ and Q' consists of the punctures of S' induced by Q , together with a single puncture labelled by $v^{1,2}$.

Proof Let $g' = g + h - 1$, $e(B_{\triangle}) := e(B) - e(B')$, $\text{Col}_{\triangle} := \text{Col}(\sigma_{\triangle}) - \text{Col}(\sigma'_{\triangle})$ and $\#_{\triangle} := |Q| - |Q'|$ as in the proof of Proposition 4.56, and observe that the quantities Col_{\triangle} and $\#_{\triangle}$ are negative. Note also that $\chi(S_{\triangle}) \leq g'$, with equality if and only if S_{\triangle} is a disjoint union of topological disks.

The argument in Proposition 4.56 tells us that $\text{ind}(B, S, P, Q, \sigma_{\triangle}) - \text{ind}'$ is equal to

$$(4.60) \quad \frac{1}{2}(3 - m_{\triangle})g' - \chi(S_{\triangle}) + 2e(B_{\triangle}) + |P_{\triangle}| - \text{Col}_{\triangle} - \#_{\triangle}.$$

As u_{\triangle} is an anchor, it follows that $\text{ind}' \geq 2 - (m + 1) = 1 - m$. By assumption, $\text{ind}(B, S, P, Q, \sigma_{\triangle}) = 2 - m$ and so the formula (4.60) is less than or equal to 1.

We examine the quantity $2(e(B) - e(B')) = 2e(B_{\triangle})$. It is clear that $\chi(B) = \chi(B')$, and thus

$$2e(B_{\triangle}) = \frac{1}{2}(\text{ob}(B) - \text{ob}(B') - \text{ac}(B) + \text{ac}(B')).$$

If we let A denote the number of acute corners at the vertices $e_1^{\infty,1}, \dots, e_h^{\infty,1}$, then as no corners at east infinity are obtuse we can see that $\text{ob}(B) - \text{ob}(B') = (1 - g + (A - h)) = -g' + A$. By the previous lemma, we know that $\text{ac}(B') - \text{ac}(B) \geq A - 2(|P| - |P'|)$ and as such $2(e(B) - e(B')) \geq -g' + A - |P_{\triangle}|$. Hence

$$1 \geq \frac{1}{2}(3 - m_{\triangle})g' - \chi(S_{\triangle}) + 2e(B_{\triangle}) + |P_{\triangle}| - \text{Col}_{\triangle} - \#_{\triangle} \geq 0$$

and at most one of the terms A , $g' - \chi(S_{\triangle})$, Col_{\triangle} or $\#_{\triangle}$ is positive (and equal to one).

Assume first that $\chi(S_{\triangle}) = g'$, so that S_{\triangle} is a disjoint union of topological disks. There are a series of cases for this:

- If $A = 1$, then $\text{Col}_{\triangle} = \#_{\triangle} = 0$ (so $|Q| = |Q'|$) and each component of S_{\triangle} has at most one east puncture.
- If $A = 0$, then either:
 - $\#_{\triangle} = 1$, so $|Q| = |Q'| - 1$, $\text{Col}_{\triangle} = 0$ and each component of S_{\triangle} has at most one east puncture.
 - $\text{Col}_{\triangle} = 1$, $\#_{\triangle} = 0$ and each component of S_{\triangle} has at most one east puncture: this does not occur as gluing (u', u_{\triangle}) gives a nearby curve in some $\mathcal{M}^B(\vec{x}; S; P, Q; \sigma'_{\triangle}, \vec{\sigma}_{\triangle}; \mathbf{J}')$, where $\text{Col}(\sigma'_{\triangle}) > 0$ — such moduli spaces have negative expected dimension.
 - Or, finally, $A_E - A_{E'} = 2(|P| - |P'|) - 1$: this does not occur as the left hand side of this equation is even.

If $\chi(S_{\triangle}) = g' - 1$, then one component of S_{\triangle} is a topological annulus without an “ ∞ ” puncture on one boundary component. Any nearby curve u in $\mathcal{M}^B(x; S; P, Q; \sigma'_{\triangle}, \vec{\sigma}_{\triangle}; \mathbf{J}')$ has a boundary component

with no noneast puncture upon it, so that $\pi_D \circ u$ is constant on this component or violates the maximum modulus principle—hence this case does not occur (compare [14, Proposition 5.43]). \square

Proposition 4.56 tells us that if $(u', u_{\pm}) \in \partial \overline{\mathcal{M}}_{\geq T}^B(x; S; \vec{\rho}_{\pm}, \vec{\rho}_{\pm}; \vec{\sigma}_{\pm}, \vec{\sigma}_{\pm}, \mathbf{J}^t)$, each component of the source S_{\pm} has at most one puncture which is mapped to e_j^{∞} for each j . By definition, each of these punctures is labelled by some element of the torus algebra ρ_j , and we shall write $\mathcal{M}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; \vec{\rho}, \sigma)$ to refer to the moduli space of curves in $\mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; P_{\pm}, \sigma)$ where P_{\pm} is the discrete partition and $\vec{\rho} = \{\rho_j \mid \rho_j \in E_j(S_{\pm})\}$.

In a similar manner to Section 4.9, we can define

$$\overline{\Gamma} = \{j \in 1, \dots, h \mid E_j(S_{\pm}) \neq \emptyset\}$$

and show that the moduli space $\mathcal{M}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; \vec{\rho}, \sigma)$ splits as a product

$$\mathcal{M}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; \vec{\rho}, \sigma) = \mathcal{M}_{\pm} \times \mathcal{M}_{\pm,1} \times \dots \times \mathcal{M}_{\pm,|\overline{\Gamma}|}$$

Here \mathcal{M}_{\pm} is a moduli space of one-gons in a closed Heegaard multidiagram \mathcal{D}'_{\pm} , and each $\mathcal{M}_{\pm,j}$ is a moduli space of one-gons defined using a bordered multidiagram $\mathcal{D}_{\pm,j} = (\Sigma_{\pm,j}; \gamma_j^1; b_j)$, where $\Sigma_{\pm,j}$ is a torus.

Proposition 4.61 *There is a unique homology class B'_{\pm} which carries a holomorphic representative in the moduli space \mathcal{M}_{\pm} . Moreover, $B'_{\pm} \in \pi_2(\theta_{\text{res},+}^{1,2})$, where $\theta_{\text{res},+}^{1,2}$ is the restriction of the distinguished generator $\theta_{+}^{1,2}$ to the diagram \mathcal{D}'_{\pm} . As such, the moduli space $\mathcal{M}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; \vec{\rho}, \sigma)$ is nonempty only if $B_{\pm} = B'_{\pm} + \sum_{j \in \overline{\Gamma}} B_{\pm,j}$, where each $B_{\pm,j}$ is supported within the handle \mathfrak{H}_j , and $\theta^{1,2} = \theta_{+}^{1,2}$, in which case it has the same number of elements as $\sum_{j \in \overline{\Gamma}} \mathcal{M}_{\pm,j}$.*

Proof This is considered in [10, Lemma 11.8], though as stated there the proof is incorrect; see the corrections. \square

As such, we need only to understand the moduli spaces $\mathcal{M}_{\pm,j}$.

Proposition 4.62 *Let $S_{\pm,j}$ be an anchor source with underlying surface a disc, and with unique east puncture q . For the evaluation map $\text{ev}_q: \mathcal{M}^{B_{\pm,j}}(\theta_j^{1,2}; S_{\pm,j}; \rho_j) \rightarrow \mathbb{R}$, the moduli space $\mathcal{M}_{\pm,j}(0) = \text{ev}^{-1}(0)$ contains a unique element for every ρ_j .*

Proof This follows by a model computation for each choice of ρ_j , and Proposition 4.45. For instance, consider the bordered Heegaard diagram and Heegaard multidiagrams as shown in Figure 20.

The moduli space

$$\mathcal{M}^B(x, y; S; \rho_3, \emptyset; \{\rho_3\}, \emptyset)$$

clearly contains a unique element, as does the moduli space

$$\mathcal{M}^{B'}(x, \theta_{+}^{1,2}, y; S'; \emptyset, \rho_3; \emptyset, \{\rho_3\}).$$

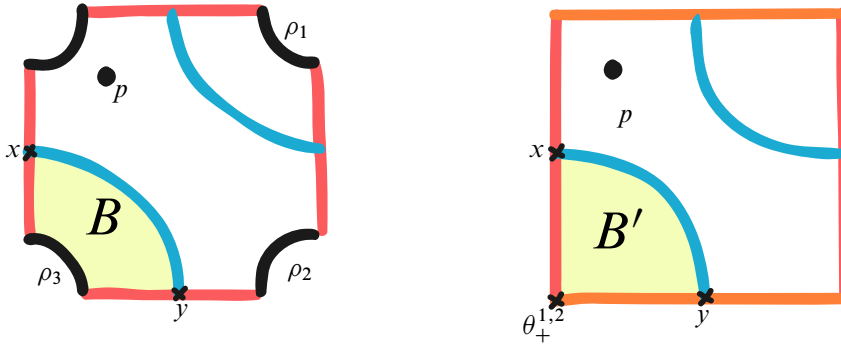


Figure 20: A model computation for Proposition 4.62.

Therefore, by Proposition 4.45,

$$|\mathcal{M}^B(x, y; S; \rho_3, \emptyset; \{\rho_3\}, \emptyset)| = |\mathcal{M}^{B'}(x, \theta_+^{1,2}, y; S'; \emptyset, \rho_3; \emptyset, \{\rho_3\})| \cdot |\mathcal{M}_{\pm, j}(0)|,$$

hence $|\mathcal{M}_{\pm, j}(0)| = 1$, as required. □

This lets us rule out anchored curves with holomorphic anchors that have more than one easterly puncture:

Proposition 4.63 *Suppose that*

$$(u', u_{\pm}) \in \mathcal{M}^{B'}(x, \theta^{1,2}, \dots, \theta^{m-1,0}, y; S'; P', Q', \sigma'_{\pm}, \bar{\sigma}'_{\pm}; \mathbf{J}) \times \mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; \bar{\rho}, \bar{\sigma})$$

is an anchored holomorphic curve in the boundary of the moduli space

$$\overline{\mathcal{M}}_{\geq T}^B(x; S; P, Q; \sigma_{\pm}, \bar{\sigma}_{\pm}).$$

Then $\bar{\sigma}'_{\pm}$ is interleaved, and hence S_{\pm} consists of a disjoint union of topological disks, only one of which has an east puncture.

Proof We already know that each component of S_{\pm} is a disk, with at most one east puncture. Suppose that $\bar{\sigma}'_{\pm}$ were not interleaved. By Proposition 4.62, there is a (unique) element v_{\pm} of $\mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; \bar{\rho}, \bar{\sigma})$ where the height of each east puncture is 0. As such, by Proposition 4.45, we could glue u' to v_{\pm} to get an element of a moduli space $\mathcal{M}_{\geq T}^B(x; S; \bar{\rho}_{\pm}, \bar{\rho}_{\pm}; \sigma_{\pm}; \mathbf{J}')$ with $\text{Col}(\sigma_{\pm}) > 0$; but this moduli space has negative expected dimension. □

The final result we need is to understand the implied asymptotics of curves in the boundary:

Proposition 4.64 *Suppose that*

$$(u', u_{\pm}) \in \mathcal{M}^{B'}(x, \theta^{1,2}, \dots, \theta^{m-1,0}, y; S'; P', Q', \sigma'_{\pm}, \bar{\sigma}'_{\pm}; \mathbf{J}) \times \mathcal{M}_{\pm}^{B_{\pm}}(\theta^{1,2}; S_{\pm}; P_{\pm}, \sigma)$$

is an anchored holomorphic curve in the boundary of the moduli space

$$\mathcal{M}_{\geq T}^B(\bar{x}; S; \bar{\rho}_{\pm}, \bar{\rho}_{\pm}; \bar{\sigma}_{\pm}, \bar{\sigma}_{\pm}).$$

Then $u' \in \mathcal{M}^{B'}(x, \theta^{1,2}, \dots, \theta^{m-1,0}, y; S'; \bar{\rho}'_{\pm}, \bar{\rho}'_{\pm}, \bar{\sigma}'_{\pm}, \bar{\sigma}'_{\pm}; \mathbf{J})$, where the pair $(\bar{\rho}'_{\pm}, \sigma'_{\pm})$ and $(\bar{\rho}'_{\pm}, \sigma'_{\pm})$ splay $(\bar{\rho}_{\pm}, \sigma_{\pm})$ and $(\bar{\rho}_{\pm}, \sigma_{\pm})$.

Proof Consider a small disc neighbourhood M_j^1 of $e_j^{\infty,1}$ which satisfies that each component of $\pi_\Sigma \circ u'^{-1}(M_j^1)$ contains a single point of the set $\Xi := \pi_\Sigma \circ u'^{-1}(e_j^{\infty,1})$ (we include the puncture q which is mapped to $e_j^{\infty,1}$ in this set) — label the components by $S(\xi)$ so that $\xi \in \Xi$ is contained in $S(\xi)$. The set Ξ splits into three sets: q , the set Ξ_L of points mapped to e^2 by $\pi_D \circ u'$, and the set Ξ_R of points mapped to e^1 by $\pi_D \circ u'$. There are corresponding restricted sources $S_L := \bigcup_{\xi \in \Xi_L} S(\xi)$ and $S_R := \bigcup_{\xi \in \Xi_R} S(\xi)$.

The homology class of $\pi_\Sigma \circ u'|_{S_L}$ is modelled upon $n \cdot B(\rho)$ for some nonjumping ρ (which is equal to ρ_{12} if γ_j^2 is meridional, and ρ_{23} if γ_j^2 is longitudinal), and the homology class of $\pi_\Sigma \circ u'|_{S(q)}$ is modelled upon $B(\rho')$ for some jumping ρ' . It follows (by considering multiplicities away from the approximation region between γ^1 and γ^2) that if

$$(D_j^1, B_j^1, D_j^2, B_j^2, \dots, D_j^k, B_j^k, D_j^{k+1})$$

were a sequence of Reeb domains for B_j , the sequence obtained by replacing each $D_j^i = n_i B^i(\rho_i)$ with $n_i B^{i+1}(\rho_i)$ and each $B_j^i = B^i(\rho'_i)$ with $B^{i+1}(\rho'_i)$, adjoining $B_j^0 = B_j^1(\rho')$ and replacing $D_j^1 = B^1(\rho_1)$ by $n \cdot B^2(\rho) + n_1 B^2(\rho_1)$ is a sequence of Reeb domains for $B'_j - B(\vec{\rho}'_{\Delta,j})$.

It remains to show that the sequence

$$\vec{\rho} = (\rho', \underbrace{\rho, \dots, \rho}_{n \text{ times}})$$

satisfies that $\vec{\rho}'_{\Delta,j} \star \vec{\rho} = \vec{\rho}_{\Delta,j}$. To see this, we note that the domain $B(\vec{\rho}_{\pm})$ must agree with $B^1(\vec{\rho}'_{\pm}) + B^1(\vec{\rho})$ away from the approximation region. As such, $B(\vec{\rho}) = B(\vec{\rho}_{\pm}) - B(\vec{\rho}'_{\pm})$. This latter quantity, however, may be written in a more useful way.

Let $\mathcal{Q} = \pi_\Sigma \circ u'^{-1}(e_j^\infty)$, and let \mathcal{Q}' denote the subset of \mathcal{Q} given by $\{q \in \mathcal{Q} \mid \text{ev}_j(q) < 0\}$. If we fix a small neighbourhood M_j^∞ of e_j^∞ such that each component of $\pi_\Sigma \circ u'^{-1}(M_j^\infty)$ contains a unique point in \mathcal{Q} , the domains $B(\vec{\rho}_{\pm})$ and $B(\vec{\rho}'_{\pm})$ can be written as a sum over points in \mathcal{Q} or \mathcal{Q}' , respectively, of the quantities $B(\rho(q))$, where $\rho(q)$ is such that $(\pi_\Sigma \circ u')(S(q))$ is modelled upon $B(\rho(q)) \subset M_j^\infty$ — and $S(q)$ denotes the component of $\pi_\Sigma \circ u'^{-1}(M_j^\infty)$ containing q . It follows by considering the multiplicity of u' at basepoints in each of the four quadrants of M_j^∞ that $B(\vec{\rho}_{\pm}) - B(\vec{\rho}'_{\pm})$ coincides with the sum over points q in $\mathcal{Q} - \mathcal{Q}'$ of $B(q)$, and thus is equal to $B(\vec{\rho})$, where $\vec{\rho}$ satisfies $\vec{\rho}'_{\Delta,j} \star \vec{\rho} = \vec{\rho}_{\Delta,j}$. As such, $B(\vec{\rho}) = B(\vec{\rho})$ and so $\vec{\rho}$ must lie in the same 123-equivalence class as $\vec{\rho}$. The cardinality of the sets \mathcal{Q} and $\mathcal{Q}' \cup \Xi_L \cup \{q\}$ agree, so in fact $\vec{\rho} = \vec{\rho}$. □

We finish this section by utilising the results above to understand properly the behaviour of moduli spaces under the neck-stretching operation.

Proposition 4.65 *Let $B, S, \vec{\rho}_{\Delta}, \vec{\rho}_{\pm}$ and $\vec{\sigma}_{\Delta}$ be such that*

$$\mathcal{M}_{\geq T} := \mathcal{M}_{\geq T, \vec{e}, \varepsilon}^{B_{\vec{e}}}(\vec{x}; S_{\vec{e}}; \vec{\rho}_{\Delta}, \vec{\rho}_{\pm}; \vec{\sigma}_{\Delta}, \vec{\sigma}_{\pm})$$

is one-dimensional, and $\vec{\sigma}_{\Delta}$ is interleaved.

Say, as shorthand, that $(B'; \bar{\rho}'_{\Delta}, \sigma'_{\Delta}; \bar{\rho}'_{\pm}, \sigma'_{\pm})$ splays $(B; \bar{\rho}_{\Delta}, \sigma_{\Delta}; \bar{\rho}_{\pm}, \sigma_{\pm})$ if B' splays B , and $(\bar{\rho}'_{\Delta}, \sigma'_{\Delta})$ and $(\bar{\rho}'_{\pm}, \sigma'_{\pm})$ splay $(\bar{\rho}_{\Delta}, \sigma_{\Delta})$ and $(\bar{\rho}_{\pm}, \sigma_{\pm})$.

Denote by

$$\mathcal{M}_{\text{splayed}} := \bigcup_{\substack{(B'; \bar{\rho}'_{\Delta}, \sigma'_{\Delta}; \bar{\rho}'_{\pm}, \sigma'_{\pm}) \\ \text{splays} \\ (B; \bar{\rho}_{\Delta}, \sigma_{\Delta}; \bar{\rho}_{\pm}, \sigma_{\pm})}} \mathcal{M}^{B'}(\mathbf{x}, \boldsymbol{\theta}_+(k-1), \mathbf{y}; S'; \bar{\rho}'_{\Delta}, \bar{\rho}'_{\pm}; \sigma'_{\Delta}, \sigma'_{\pm}; \mathbf{J}).$$

Then $\overline{\mathcal{M}}_{\geq T}$ is a compact one-manifold with boundary given by the disjoint union of

$$\mathcal{M}_{\varepsilon, \varepsilon}^{B_{\varepsilon}}(\mathbf{x}, \boldsymbol{\theta}_+(k-2), \mathbf{y}; S; \bar{\rho}_{\Delta}, \bar{\rho}_{\pm}; \bar{\sigma}_{\Delta}, \bar{\sigma}_{\pm})$$

and $\mathcal{M}_{\text{splayed}} \times \mathcal{M}^{B_{\pm}}(\boldsymbol{\theta}^{1,2}, S_{\pm}, \rho)$, where ρ is the unique jumping chord in the final part of $\bar{\sigma}_{\pm}$.

Therefore, the moduli spaces

$$\mathcal{M}^B(\mathbf{x}, \boldsymbol{\theta}_+(k-2), \mathbf{y}; S; \bar{\rho}_{\Delta}, \bar{\rho}_{\pm}; \bar{\sigma}_{\Delta}, \bar{\sigma}_{\pm}; \mathbf{J}^T)$$

and $\mathcal{M}_{\text{splayed}}$ contain the same number of elements modulo two.

Proof By Proposition 4.64, we know that any curve in $\partial \overline{\mathcal{M}}_{\geq T}$ lives in the moduli space

$$\mathcal{M}_{\text{splayed}} \times \mathcal{M}^{B_{\pm}}(\boldsymbol{\theta}^{1,2}, S_{\pm}, \rho).$$

By Proposition 4.56, we know that S_{\pm} is a disjoint union of topological disks, and thus by Proposition 4.45, $\partial \overline{\mathcal{M}}_{\geq T}$ is a compact one-manifold with boundary given.

The second part follows from the identification of

$$\mathcal{M}_{\varepsilon, \varepsilon}^{B_{\varepsilon}}(\mathbf{x}, \boldsymbol{\theta}_+(k-2), \mathbf{y}; S; \bar{\rho}_{\Delta}, \bar{\rho}_{\pm}; \bar{\sigma}_{\Delta}, \bar{\sigma}_{\pm})$$

with

$$\mathcal{M}^B(\mathbf{x}, \boldsymbol{\theta}_+(k-2), \mathbf{y}; S; \bar{\rho}_{\Delta}, \bar{\rho}_{\pm}; \bar{\sigma}_{\Delta}, \bar{\sigma}_{\pm})$$

given in Propositions 4.6.1 and 4.7.2, and the fact that $\mathcal{M}^{B_{\pm}}(\boldsymbol{\theta}^{1,2}, S_{\pm}; \rho)$ has one element modulo two by Proposition 4.62. □

5 Statement and proof of theorems

After all this formality, we are at long last ready to state some theorems.

Let \mathcal{D} be a bordered diagram with h boundary components. We shall define a right \mathcal{A}_{∞} multimodule over the torus algebra, which we call $\text{Poly}(\mathcal{D})$. We first define a series of auxiliary modules $\text{Poly}_k(\mathcal{D})$.

Consider the module generated over \mathbb{F}_2 by the set of generators $\mathcal{G}(\mathcal{D})$. This has an obvious right action by the ring of idempotents of \mathcal{T}^h , defined by setting

$$\mathbf{x} \cdot \iota = \begin{cases} \mathbf{x} & \text{if } \iota(\mathbf{x}) = \iota, \\ 0 & \text{otherwise.} \end{cases}$$

We define an \mathcal{A}_∞ module structure upon $\text{Poly}_k(\mathcal{D})$ by first defining maps

$$n_{i_1, \dots, i_h}^k : \text{Poly}_k(\mathcal{D}) \otimes_{\mathcal{I}} \mathcal{T}_1^{\otimes i_1} \otimes_{\mathcal{I}} \dots \otimes_{\mathcal{I}} \mathcal{T}_h^{\otimes i_h} \times \vec{\sigma}(i_1, \dots, i_h) \rightarrow \text{Poly}_k(\mathcal{D}),$$

where $\vec{\sigma}(i_1, \dots, i_h)$ denotes the set of all splittings of all sets of sequences of Reeb chords in

$$\mathcal{T}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{T}_h^{\otimes i_h}.$$

Suppose that $\vec{\rho} = \vec{\rho}_1, \dots, \vec{\rho}_h$ is a set of sequences of Reeb multichords in \mathcal{T}^h , where $|\vec{\rho}_j| = i_j$. For every interleaved splicing σ of $\vec{\rho}$, we define the coefficient

$$a^k(x, y, \vec{\rho}, \vec{\sigma}) := \sum_{\substack{(\vec{\rho}_\Delta, \vec{\sigma}_\Delta), (\vec{\rho}_\sharp, \vec{\sigma}_\sharp) \\ \text{is a } k\text{-shipping} \\ \text{of } (\vec{\rho}, \vec{\sigma})}} \sum_{\substack{B \in \pi_2^{\theta^+}(x, y) \\ \chi(S) = \chi_{\text{emb}}(S) \\ \text{ind}(B, S, \vec{\rho}_\Delta, \vec{\rho}_\sharp, \vec{\sigma}_\Delta) = 2 - |\vec{\sigma}_\sharp|}} \#\mathcal{M}^B(x, \theta_+(|\vec{\sigma}_\sharp| - 2), y'; S; \vec{\rho}_\Delta, \vec{\rho}_\sharp; \vec{\sigma}_\Delta, \vec{\sigma}_\sharp)$$

where the nearest-point map takes y' to y , and hence put

$$n_{i_1, \dots, i_h}^k(x, \vec{\rho}, \vec{\sigma}) := \sum_{y \in \mathcal{G}(\mathcal{D})} a^k(x, y, \vec{\rho}, \vec{\sigma}) \cdot y.$$

We define

$$m_{i_1, \dots, i_h}^k : \text{Poly}_k(\mathcal{D}) \otimes_{\mathcal{I}} \mathcal{A}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{A}_h^{\otimes i_h} \rightarrow \text{Poly}_k(\mathcal{D})$$

by

$$m_{i_1, \dots, i_h}^k(x, \vec{\rho}) := \sum_{\sigma \text{ an interleaving for } \vec{\rho}} n_{i_1, \dots, i_h}^k(x, \vec{\rho}, \vec{\sigma}).$$

Proposition 5.1 *For sufficiently generic choice of \mathbf{J} , the definition of $\text{Poly}_0(\mathcal{D})$ above agrees with the definition of the multimodule $\text{BSA}(\mathcal{D})$ described in [26, Section 8.3].*

Proof In the case where $k = 0$, the above moduli spaces are formed of bigons in the bordered diagram \mathcal{D} , with constraints upon evaluation maps given by the interleaving σ . To see that these are precisely the bigons counted in the definition of $\text{BSA}(\mathcal{D})$, we note that for any $(B, S, \vec{\rho}_\sharp)$ such that the moduli space $\mathcal{M}(x, y; S; \vec{\rho}_\sharp)$ is zero-dimensional (as counted in the definition of the BSA multiplication maps), the subspace of $\mathcal{M}(x, y; S; \vec{\rho}_\sharp)$ where the heights of any punctures q, q' labelled by a jumping chord coincide is given by $\mathcal{M}(x, y; S; \vec{\rho}_\sharp; \sigma)$ for σ some noninterleaved splicing of $\vec{\rho}_\sharp$. By the formula at the end of Section 4.4.1 (or Proposition 4.22), $\text{ind}(B, S, \vec{\rho}_\sharp, \sigma) < 1$ and so the corresponding moduli space is empty; ie the multiplication maps in $\text{BSA}(\mathcal{D})$ do not count curves where heights coincide. (This should be compared with [12, Section 2.4.3].)

As such, the moduli space $\mathcal{M}(x, y; S; \vec{\rho}_\sharp)$ splits as a union over all interleaved splittings of moduli spaces $\mathcal{M}(x, y; S; \vec{\rho}_\sharp; \vec{\sigma})$, which is precisely what the maps m_{i_1, \dots, i_h}^0 count. □

Proposition 5.2 (compare [14, Lemma 7.7; 26, Theorem 7.8]) *Suppose that the diagram \mathcal{D} is provincially admissible. For any fixed generators x and y , and a set of nonempty sets of Reeb chords, there are at most finitely many $B \in \pi_2^{\theta^+}(x, y)$ such that $\mathcal{M}^B(x, \theta_+(k - 2), y'; S; \vec{\rho}_\Delta, \vec{\rho}_\sharp)$ are nonempty — ie the sums in the definition of the maps m_n are finite.*

Moreover, if \mathcal{D} is admissible, then there is a constant κ depending only upon \mathcal{D} for which every map m_i^k with $i > \kappa$ is zero, independent of k .

Proof We proceed by induction. In the case $k = 0$, by Proposition 5.1 the proposition coincides with [26, Theorem 7.8.].

For higher k , we prove the first half of the proposition by noting that by Proposition 4.45, if there were an infinite number of $B \in \pi_2(\mathbf{x}, \boldsymbol{\theta}_+(k), \mathbf{y})$ with an element $u_B \in \mathcal{M}^B(\mathbf{x}, \boldsymbol{\theta}_+(k-2), \mathbf{y}'; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \sigma_\Delta, \sigma_\ddagger)$, we can glue u_B to a holomorphic anchor to give a curve $u_{B'} \in \mathcal{M}^{B'}(\mathbf{x}, \boldsymbol{\theta}_+(k-3), \mathbf{y}'; S'; \vec{\rho}'_\Delta, \vec{\rho}'_\ddagger; \vec{\sigma}'_\Delta, \vec{\sigma}'_\ddagger)$, where $(\vec{\rho}'_\Delta, \vec{\sigma}'_\Delta), (\vec{\rho}'_\ddagger, \vec{\sigma}'_\ddagger)$ is a $k - 1$ shipping of $(\vec{\rho}, \vec{\sigma})$. By our induction hypothesis, this is impossible.

For the second half, we let κ be the corresponding constant for m_i^0 . Again, if some $i > \kappa$ does not satisfy that $m_i^k = 0$, we may glue the corresponding polygon to a holomorphic anchor, implying that m_i^{k-1} is also nonzero. By our induction hypothesis, this is, again, impossible. \square

We now turn to the main definition and theorem of the paper. Consider the 2^h elementary splayings $\mathcal{D}(\iota(\delta))$, where $\delta \in \{m, l\}^h$. The module $\text{Poly}(\mathcal{D})$ is generated over the field \mathbb{F}_2 by the sum $\mathcal{G} := \bigoplus_{\delta \in \{m, l\}^h} \mathcal{G}(\mathcal{D}(\iota(\delta)))$. We endow this module with the obvious right action of the ring of idempotents of \mathcal{T}^h by setting

$$\mathbf{x} \cdot \iota(\delta) = \begin{cases} \mathbf{x} & \text{if } \mathbf{x} \in \mathcal{G}(\mathcal{D}(\iota(\delta))), \\ 0 & \text{otherwise.} \end{cases}$$

We define an \mathcal{A}_∞ module structure upon $\text{Poly}(\mathcal{D})$ by first defining maps

$$n_{i_1, \dots, i_h} : \text{Poly}(\mathcal{D}) \otimes_{\mathcal{I}} \mathcal{T}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{T}_h^{\otimes i_h} \times \vec{\sigma}(i_1, \dots, i_h) \rightarrow \text{Poly}(\mathcal{D}),$$

where $\vec{\sigma}(i_1, \dots, i_h)$ denotes the set of all splittings of all sets of sequences of Reeb chords in

$$\mathcal{T}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{T}_h^{\otimes i_h}.$$

Suppose that $\vec{\rho}_\ddagger = \vec{\rho}_1, \dots, \vec{\rho}_h$ is a set of sequences of Reeb multichords in \mathcal{T}^h , where $|\vec{\rho}_j| = i_j$. For every interleaved splicing $\vec{\sigma}_\ddagger$ of $\vec{\rho}_\ddagger$, we define

$$n_{i_1, \dots, i_h}(\mathbf{x}, \vec{\rho}_\ddagger, \vec{\sigma}_\ddagger) := \sum_{\mathbf{y} \in \mathcal{G}(\mathcal{D})} \sum_{\substack{B \in \pi_2^{\theta_+}(\mathbf{x}, \mathbf{y}) \\ \chi(S) = \chi_{\text{emb}}(S) \\ \text{ind}(B, S, \vec{\rho}_\ddagger) = 2 - |\vec{\sigma}_\ddagger|}} \# \mathcal{M}_\ddagger^B(\mathbf{x}, \boldsymbol{\theta}_+(|\vec{\sigma}_\ddagger| - 2), \mathbf{y}'; S; \vec{\rho}_\ddagger; \vec{\sigma}_\ddagger) \cdot \mathbf{y},$$

where the nearest-point map takes \mathbf{y}' to \mathbf{y} .

We define

$$m_{i_1, \dots, i_h} : \text{Poly}(\mathcal{D}) \otimes_{\mathcal{I}} \mathcal{A}_1^{\otimes i_1} \otimes \dots \otimes \mathcal{A}_h^{\otimes i_h} \rightarrow \text{Poly}(\mathcal{D})$$

by

$$m_{i_1, \dots, i_h}(\mathbf{x}, \vec{\rho}) := \sum_{\vec{\sigma} \text{ an interleaving for } \vec{\rho}} n_{i_1, \dots, i_h}(\mathbf{x}, \vec{\rho}, \vec{\sigma}).$$

We first verify that the modules $\text{Poly}_k(\mathcal{D})$ are \mathcal{A}_∞ modules.

Theorem 5.3 *The maps $\{n_{i_1, \dots, i_h}^k\}_{i_1, \dots, i_h}$ satisfy the partial \mathcal{A}_∞ relation. As such, the data*

$$(\text{Poly}_k(\mathcal{D}), \{m_{i_1, \dots, i_h}^k\}_{i_1, \dots, i_h}^\infty)$$

forms an \mathcal{A}_∞ module over \mathcal{T}^h .

Proof Recall the relation which we need to verify is

$$0 = \sum_{\substack{\vec{\lambda} \star \vec{\delta} = \vec{\rho} \\ \vec{\sigma}_1 \star \vec{\sigma}_2 = \vec{\sigma}}} n(n(x, \vec{\lambda}, \vec{\sigma}_1), \vec{\delta}, \vec{\sigma}_2) + \sum_{i, j \text{ compatible}} n(x, \vec{\mu}_j^i(\vec{\rho}), \vec{\sigma}_j^i) + \sum_{k' \text{ collidable}}^{|\vec{\sigma}|-1} n(x, \vec{\rho}, \vec{\sigma}(k')).$$

Consider a one-dimensional moduli space $\mathcal{M} = \mathcal{M}^B(x, \theta_+(k), y; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger)$, where $\vec{\sigma}_\Delta \star \vec{\sigma}_\ddagger = \vec{\sigma}$, and $\vec{\sigma}_\ddagger$ is a k -splicing. By Theorem 4.47, the ends of this moduli space are formed of the following.

- Spinal two-storey ends, corresponding with terms in the first sum in the statement of the theorem.
- Mixed two-storey ends, which, by Proposition 4.53 are either composition ends or collision ends (see Definition 4.54). Composition ends correspond (by Proposition 4.55) with terms in the second sum in the statement where i, j satisfy that $\rho_j^i \in P_j^{i'} \in \vec{\sigma}$ for $i' \geq k$, and ρ_j^i, ρ_j^{i+1} are equal to ρ_1, ρ_2 or ρ_2, ρ_3 ; collision ends correspond with terms in the third sum in the statement where $k' \geq k$.
- (i, j) -composition ends which correspond with terms in the second sum where i, j satisfy that $\rho_j^i \in P_j^{i'}$ for $i' < k$.
- k' -collision ends, which correspond with terms in the third sum where $k' < k$.
- Cut-vanishing ends, corresponding with terms in the second sum in the statement where i, j satisfy that $\rho_j^i \in P_j^{i'} \in \vec{\sigma}$ for $i' \geq k$, and $\rho_j^i \cdot \rho_j^{i+1}$ is equal to ρ_{123} .

By Theorem 4.47, the number of these is zero modulo two; it corresponds that the sum is zero, as required. □

We will not show directly that $\text{Poly}_k(\mathcal{D})$ is an invariant of \mathcal{D} up to equivalence — indeed, we have not discussed any notion of equivalence for diagrams. Instead, we just need to understand how $\text{Poly}_k(\mathcal{D})$ depends upon the choice of generic admissible compatible almost-complex structures \mathbf{J} .

Proposition 5.4 *Let \mathbf{J}_0 and \mathbf{J}_1 be generic, admissible, compatible families of almost-complex structures. Then the \mathcal{A}_∞ modules $\text{Poly}_k(\mathcal{D}, \mathbf{J}_0)$ and $\text{Poly}_k(\mathcal{D}, \mathbf{J}_1)$ are \mathcal{A}_∞ homotopy equivalent.*

Proof This follows by modifying [14, Section 7.3.1] to the case of splayed polygons — combining it with, say, [13, Proposition 3.30], and the proof of Theorem 5.3 above. □

Theorem 5.5 *Let $k \geq 0$ be an integer. Then the \mathcal{A}_∞ modules $\text{Poly}_k(\mathcal{D}, \mathbf{J})$ and $\text{Poly}_{k+1}(\mathcal{D}, \mathbf{J})$ are \mathcal{A}_∞ homotopy equivalent.*

Proof By Proposition 4.65, and the finiteness of the zero-dimensional moduli spaces involved in the definition of the m_i^k and m_i^{k+1} there is some $T > 0$ such that, for every \mathbf{x} , \mathbf{y} , B , S , $\vec{\rho}$ and $\vec{\sigma}$, the moduli spaces $\mathcal{M}^B(\mathbf{x}, \theta_+(k), \mathbf{y}; S; \vec{\rho}_\Delta, \vec{\rho}_\ddagger; \vec{\sigma}_\Delta, \vec{\sigma}_\ddagger, \mathbf{J}^T)$ and

$$\mathcal{M}_{\text{splayed}} := \bigcup_{\substack{(B'; \vec{\rho}'_\Delta, \vec{\sigma}'_\Delta; \vec{\rho}'_\ddagger, \vec{\sigma}'_\ddagger) \\ \text{splays} \\ (B; \vec{\rho}_\Delta, \vec{\sigma}_\Delta; \vec{\rho}_\ddagger, \vec{\sigma}_\ddagger)}} \mathcal{M}^{B'}(\mathbf{x}, \theta_+(k-1), \mathbf{y}; S'; \vec{\rho}'_\Delta, \vec{\rho}'_\ddagger; \vec{\sigma}'_\Delta, \vec{\sigma}'_\ddagger; \mathbf{J}).$$

are identified.

It follows that if we sum the number of elements of moduli spaces of both types over all $B \in \pi_2^{\theta_+}(\mathbf{x}, \mathbf{y})$, we get the same number modulo 2. These are the coefficients of \mathbf{y} in the maps $n_{i_1, \dots, i_h}^k(\mathbf{x}, \vec{\rho})$, defined with respect to \mathbf{J}^T , and of \mathbf{y} in $n_{i_1, \dots, i_h}^{k+1}(\mathbf{x}, \vec{\rho})$, defined with respect to \mathbf{J} , respectively. It follows that the \mathcal{A}_∞ modules $\text{Poly}_k(\mathcal{D}, \mathbf{J}^T)$ and $\text{Poly}_k(\mathcal{D}, \mathbf{J})$ are isomorphic as \mathcal{A}_∞ modules.

By Proposition 5.4, $\text{Poly}_k(\mathcal{D}, \mathbf{J}^T)$ is \mathcal{A}_∞ homotopy equivalent to $\text{Poly}_k(\mathcal{D}, \mathbf{J})$, and so the proposition follows. \square

We finally state and prove the main theorem of this paper:

Theorem 5.6 $\text{Poly}(\mathcal{D})$ is an \mathcal{A}_∞ module, and $\text{Poly}(\mathcal{D})$ is \mathcal{A}_∞ homotopy equivalent to $\text{BSA}(\mathcal{D})$.

Proof Let κ denote the constant defined in Proposition 5.2. The only nonzero maps in Poly_κ are determined by moduli spaces $\mathcal{M}^B(\mathbf{x}, \theta_+(k), \mathbf{y}; S; \emptyset, \vec{\rho}_\ddagger; \emptyset, \vec{\sigma}_\ddagger; \mathbf{J})$, which, by Proposition 4.24, are all identified with the moduli spaces used in the definition of the structure maps of $\text{Poly}(\mathcal{D})$. As such, $\text{Poly}(\mathcal{D})$ is an \mathcal{A}_∞ -module which is \mathcal{A}_∞ isomorphic to $\text{Poly}_\kappa(\mathcal{D})$. Moreover, by Proposition 5.1 and Theorem 5.5,

$$\text{BSA}(\mathcal{D}) \equiv \text{Poly}_0(\mathcal{D}) \cong \text{Poly}_1(\mathcal{D}) \cong \cdots \cong \text{Poly}_{\kappa-1}(\mathcal{D}) \cong \text{Poly}_\kappa(\mathcal{D}) \equiv \text{Poly}(\mathcal{D}). \quad \square$$

References

- [1] L V Ahlfors, *Conformal invariants: topics in geometric function theory*, McGraw-Hill, New York (1973) MR Zbl
- [2] D Auroux, *Fukaya categories of symmetric products and bordered Heegaard–Floer homology*, J. Gökova Geom. Topol. 4 (2010) 1–54 MR Zbl
- [3] F Bourgeois, *A Morse–Bott approach to contact homology*, PhD thesis, Stanford University (2002) Available at <https://www.proquest.com/docview/305591502>
- [4] F Bourgeois, Y Eliashberg, H Hofer, K Wysocki, E Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. 7 (2003) 799–888 MR Zbl
- [5] E Eftekhary, *Floer homology and splicing knot complements*, Algebr. Geom. Topol. 15 (2015) 3155–3213 MR Zbl

- [6] **J Hanselman**, *Bordered Heegaard Floer homology and graph manifolds*, *Algebr. Geom. Topol.* 16 (2016) 3103–3166 [MR](#) [Zbl](#)
- [7] **T Hockenhull**, *Gluing bypass exact triangles*, in preparation
- [8] **H Hofer**, **V Lizan**, **J-C Sikorav**, *On genericity for holomorphic curves in four-dimensional almost-complex manifolds*, *J. Geom. Anal.* 7 (1997) 149–159 [MR](#) [Zbl](#)
- [9] **A S Levine**, *Knot doubling operators and bordered Heegaard Floer homology*, *J. Topol.* 5 (2012) 651–712 [MR](#) [Zbl](#)
- [10] **R Lipshitz**, *A cylindrical reformulation of Heegaard Floer homology*, *Geom. Topol.* 10 (2006) 955–1096 [MR](#) [Zbl](#) Corrections at <http://pages.uoregon.edu/lipshitz/MoreCylErat.pdf>
- [11] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Computing \widehat{HF} by factoring mapping classes*, *Geom. Topol.* 18 (2014) 2547–2681 [MR](#) [Zbl](#)
- [12] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Bimodules in bordered Heegaard Floer homology*, *Geom. Topol.* 19 (2015) 525–724 [MR](#) [Zbl](#)
- [13] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Bordered Floer homology and the spectral sequence of a branched double cover, II: The spectral sequences agree*, *J. Topol.* 9 (2016) 607–686 [MR](#) [Zbl](#)
- [14] **R Lipshitz**, **P S Ozsváth**, **D P Thurston**, *Bordered Heegaard Floer homology*, *Mem. Amer. Math. Soc.* 1216, Amer. Math. Soc., Providence, RI (2018) [MR](#) [Zbl](#)
- [15] **C Manolescu**, **P Ozsváth**, *Heegaard Floer homology and integer surgeries on links* (2010) [arXiv 1011.1317](#)
To appear in *Geom. Topol.*
- [16] **D McDuff**, **D Salamon**, *J-holomorphic curves and symplectic topology*, 2nd edition, Amer. Math. Soc. Colloq. Publ. 52, Amer. Math. Soc., Providence, RI (2012) [MR](#) [Zbl](#)
- [17] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and knot invariants*, *Adv. Math.* 186 (2004) 58–116 [MR](#) [Zbl](#)
- [18] **P Ozsváth**, **Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, *Ann. of Math.* 159 (2004) 1027–1158 [MR](#) [Zbl](#)
- [19] **P Ozsváth**, **Z Szabó**, *On the Heegaard Floer homology of branched double-covers*, *Adv. Math.* 194 (2005) 1–33 [MR](#) [Zbl](#)
- [20] **P Ozsváth**, **Z Szabó**, *Holomorphic disks, link invariants and the multi-variable Alexander polynomial*, *Algebr. Geom. Topol.* 8 (2008) 615–692 [MR](#) [Zbl](#)
- [21] **P S Ozsváth**, **Z Szabó**, *Knot Floer homology and integer surgeries*, *Algebr. Geom. Topol.* 8 (2008) 101–153 [MR](#) [Zbl](#)
- [22] **P S Ozsváth**, **Z Szabó**, *Knot Floer homology and rational surgeries*, *Algebr. Geom. Topol.* 11 (2011) 1–68 [MR](#) [Zbl](#)
- [23] **J A Rasmussen**, *Floer homology and knot complements*, PhD thesis, Harvard University (2003) Available at <https://www.proquest.com/docview/305332635>
- [24] **P Seidel**, *Fukaya categories and Picard–Lefschetz theory*, Eur. Math. Soc., Zürich (2008) [MR](#) [Zbl](#)
- [25] **P Seidel**, *Lectures on four-dimensional Dehn twists*, from “Symplectic 4-manifolds and algebraic surfaces”, *Lecture Notes in Math.* 1938, Springer (2008) 231–267 [MR](#) [Zbl](#)
- [26] **R Zarev**, *Bordered sutured Floer homology*, PhD thesis, Columbia University (2011) Available at <https://www.proquest.com/docview/868276640>

Frome, Somerset, United Kingdom

thomas.hockenhull@gmail.com

Received: 3 April 2019 Revised: 18 July 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Markus Land	LMU München markus.land@math.lmu.de
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Octav Cornea	Université de Montreal cornea@dms.umontreal.ca	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Robert Oliver	Université Paris 13 bobol@math.univ-paris13.fr
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futer	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com


See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US \$760/year for the electronic version, and \$1110/year (+\$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by [Mathematical Reviews](#), [Zentralblatt MATH](#), [Current Mathematical Publications](#) and the [Science Citation Index](#).

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**

nonprofit scientific publishing

<https://msp.org/>

© 2025 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25 Issue 6 (pages 3145–3787) 2025

Holomorphic polygons and the bordered Heegaard Floer homology of link complements	3145
THOMAS HOCKENHULL	
Exact Lagrangian tori in symplectic Milnor fibers constructed with fillings	3225
ORSOLA CAPOVILLA-SEARLE	
A note on embeddings of 3-manifolds in symplectic 4-manifolds	3251
ANUBHAV MUKHERJEE	
A note on knot Floer homology of satellite knots with $(1, 1)$ -patterns	3271
WEIZHE SHEN	
A K -theory spectrum for cobordism cut and paste groups	3287
RENEE S HOEKZEMA, CARMEN ROVI and JULIA SEMIKINA	
The Curtis–Wellington spectral sequence through cohomology	3315
DANA HUNTER	
The slices of quaternionic Eilenberg–Mac Lane spectra	3341
BERTRAND J GUILLOU and CARISSA SLONE	
Cocycles of the space of long embeddings and BCR graphs with more than one loop	3385
LEO YOSHIOKA	
Asymptotic cones of snowflake groups and the strong shortcut property	3429
CHRISTOPHER H CASHEN, NIMA HODA and DANIEL J WOODHOUSE	
Whitney tower concordance and knots in homology spheres	3503
CHRISTOPHER W DAVIS	
The asymptotic behaviors of the colored Jones polynomials of the figure-eight knot, and an affine representation	3523
HITOSHI MURAKAMI	
The Goldman bracket characterizes homeomorphisms between noncompact surfaces	3585
SUMANTA DAS, SIDDHARTHA GADGIL and AJAY KUMAR NAIR	
A geometric computation of cohomotopy groups in codegree one	3603
MICHAEL JUNG and THOMAS O ROT	
Calabi–Yau structure on the Chekanov–Eliashberg algebra of a Legendrian sphere	3627
NOÉMIE LEGOUT	
On the resolution of kinks of curves on punctured surfaces	3679
CHRISTOF GEISS and DANIEL LABARDINI-FRAGOSO	
Weinstein presentations for high-dimensional antisurgery	3707
IPSITA DATTA, OLEG LAZAREV, CHINDU MOHANAKUMAR and ANGELA WU	
Singular Legendrian unknot links and relative Ginzburg algebras	3737
JOHAN ASPLUND	
An RBG construction of integral surgery homeomorphisms	3755
QIANHE QIN	
Powell’s conjecture on the Goeritz group of S^3 is stably true	3775
MARTIN SCHARLEMANN	