Moduli spaces in bordered Heegaard Floer homology

Jessica J. Zhang

A thesis presented to the Department of Mathematics
in partial fulfillment of the requirements for the degree of
Bachelor of Arts with Honors

Harvard University
Cambridge, Massachusetts

March 24, 2025

Abstract

Bordered Heegaard Floer homology is a powerful invariant of bordered 3-manifolds, i.e., 3-manifolds with a single boundary component that is parameterized by a combinatorial object called a "pointed matched circle." It can be used, for example, to compute Heegaard Floer homology by a so-called "pairing theorem." Though beyond the scope of this thesis, bordered Heegaard Floer homology also has applications to knot theory (e.g., to compute the Ozsváth–Szabó spectral sequence between Khovanov homology and Heegaard Floer homology, or to give a knot Floer homology of tangles), 4-dimensional topology (e.g., to show that there are knots in homology 3-balls which do not bound piecewise linear disks in any homology 4-balls, or to find various examples of exotic phenomena), and contact geometry (e.g., by proving a pairing theorem for contact invariants, or by showing that there are computable, though not yet geometrically understood, $\mathcal{A}_\infty$-style maps on the set of contact structures).

The construction of bordered Heegaard Floer modules $\widehat{CFA}(Y)$ and $\widehat{CFD}(Y)$ involves counting holomorphic curves on $\Sigma \times [0,1] \times \mathbb{R}$, where Σ is the Heegaard surface of some Heegaard diagram representing Y. This is reminiscent of Lipshitz's cylindrical formulation of Ozsváth and Szabó's Heegaard Floer homology, which we also briefly sketch. As in many Floer homologies, we define a moduli space of such holomorphic curves, and that it is contained in a compact manifold whose dimension is given by some index formula. We focus particularly on the proof that this moduli space may be compactified, and prove results from symplectic field theory in doing so. We conclude with the definition of the invariants $\widehat{CFA}(Y)$ and $\widehat{CFD}(Y)$ as a certain count of points in this moduli space and with a proof of the pairing theorem.

This thesis can be thought of as a supplement to Lipshitz, Ozsváth, and Thurston's monograph *Bordered Heegaard Floer homology*, which details the construction and invariance of the bordered Heegaard Floer modules $\widehat{CFA}$ and $\widehat{CFD}$. We focus primarily on their geometric and analytic insights, while being somewhat lighter on the algebra. Instead, whenever possible, we illustrate algebraic results geometrically. We hope that this thesis will be helpful to anyone trying to learn the details of bordered Heegaard Floer homology, as well as to anyone trying to understand many of the compactness proofs that appear in invariants which arise from counting points in moduli spaces (e.g., Gromov–Witten invariants). While readers with the latter aim will not find anything specifically geared toward their interests, we hope that the focus on compactness results will still provide some useful perspective.

Contents

List of Figures

Acknowledgments

> I wrote, "Dear Toad, I am glad that you are my best friend. Your best friend, Frog."
>
> Arnold Lobel, *Frog and Toad are Friends*

It takes a village to raise a child. Admittedly, I did not raise a child. But, as it turns out, it takes, if not a village, then at least a hamlet to write a thesis. Indeed, this thesis would not exist but for the support of many people in my life.

To Professor Denis Auroux, thank you, of course, for advising this thesis. It is only with your patient and enthusiastic mentorship that any of this has been possible. You answered every question I had, explained every proof I didn't understand, read every draft I produced. Thank you, too, for advising readings I did on pseudoholomorphic curves in my junior year. This thesis is, in many ways, an outgrowth of those readings. Learning with you—from my first class in freshman year to readings in junior year to now, finally, this thesis today—has been one of the most exciting and mathematically rewarding experiences I have had, and I will be sorry to no longer have an excuse to come by your office every week with a new list of questions to ask.

I would also like to thank Professor Peter Kronheimer, for teaching me new ideas in math and new ways to think about math. The reading and research I did with you has been a highlight of my past four years, and have formed the basis for my academic interests today. Thanks also goes to Professor Ciprian Manolescu, who agreed to advise a research project last summer without even knowing me and who has been an invaluable source of mathematical advice. And thank you to Professor Mihnea Popa, who taught algebraic geometry with such passion that I was all but forced to like the subject myself.

Thank you to Maxim Jeffs, who introduced me to Floer homology and holomorphic curve theory back in my freshman year of college. And thank you to Zhenkun Li, who introduced me to symplectic geometry and low-dimensional topology. You both planted the seeds for this thesis, well before I even knew what a senior thesis was.

Thank you to my friends, who have filled my life with laughter and love. None of this would have meant anything without you all. I will miss the late nights, and the late lunches, and even the late mornings. I will not miss the early mornings, seeing as we had very few of those anyway.

To my parents, Wentao Zhang and Jihong Yang, thank you for your endless support and your boundless love. And to my sister, Cindy Zhang, thank you for having always been there for me. There is no better friend than a sister, and no better sister in the world than you.

Chapter 1

Introduction

> And I feel like this year is really about, like, the year of realizing stuff. And everyone around me, we're all just, like, realizing things.
>
> Kylie Jenner

In the past few decades, there has been an explosion of interest in low-dimensional topology, owing in part to the discovery of new invariants. A topological invariant is some quantity (e.g., a number, polynomial, or homology group) which does not change under some equivalence, often continuous or smooth deformation. Thus an invariant helps us classify and distinguish manifolds.

Floer homology is a family of such invariants. Perhaps more accurately, we may call Floer homology a "technique," from which many invariants may be created. Loosely speaking, Floer homology is an infinite-dimensional analogue of Morse homology.

In the Morse case, one considers a Morse function on a manifold M and considers the vector space $CM(M)$ spanned by its critical points. If one equips M with a Riemannian metric, then one can compute the gradient of the Morse function. This gradient flows from higher-index critical points to lower-index critical points. One can count the number of flowlines $n_{p,q}$ between critical points p and q of index i and $i-1$, respectively. Then we define the differential as

$$\partial p := \sum_{\mathrm{ind}(q)=\mathrm{ind}(p)-1} n_{p,q} q.$$

Then $(CM(M), \partial)$ is a chain complex, and its homology $HM(M)$ is an invariant of M, i.e., independent of the choice of Morse function and Riemannian metric.

Floer homology does the same thing—considers a functional with nondegenerate critical points, defines the chain complex to be freely generated by these critical points, and then counts flowlines to define the differential—but with a functional that is now defined on an infinite-dimensional space. For example, the earliest version of Floer homology, due to Andreas Floer, is known as Hamiltonian (or symplectic) Floer homology and gives an invariant of symplectic manifolds [Flo88a, Flo88c, Flo89a, Flo89b]. It is defined by doing Morse homology on a function defined on the loop space of the manifold.

Other versions of Floer homology give 3-manifold invariants which are intimately connected to invariants coming from mathematical gauge theory, i.e., coming from solutions to partial differential equations arising in the study of connections on principal bundles. In particular, in the 1980s, Simon Don-

aldson defined the gauge-theoretic Donaldson invariants, which were invariants of 4-manifolds. (These are invariants coming from the so-called "anti-self-dual equations.") Instanton Floer homology gave a 3-manifold version of the Donaldson invariants [Flo88b, DK90]. Subsequently, monopole Floer homology (or Seiberg–Witten Floer homology) gave a 3-dimensional analogue of the Seiberg–Witten invariants, which were similar to but often easier to compute than Donaldson invariants [KM07].

There is another family of 3-manifold invariants known as Heegaard Floer homology, first discovered by Peter Ozsváth and Zoltán Szabó [OS04c, OS04b]. These invariants were conjectured to be isomorphic to monopole Floer homology, and this equivalence was proved in a series of papers by Çağatay Kutluhan, Yi-Jen Lee and Clifford Taubes [KLT20]. It is worth noting, furthermore, that monopole and Heegaard Floer homology are conjecturally isomorphic to instanton Floer homology.) This is an invariant of 3-manifolds, and may be used to construct a 4-manifold invariant as in [OS06], in contrast to how the invariants described above began as 4-manifold invariants whose 3-manifold analogues were discovered later on. (The 4-manifold invariant arising from $\widehat{HF}$, which is the variant of Heegaard Floer homology which will be most relevant to this thesis, is not actually an interesting one. Instead, to obtain an interesting 4-manifold invariant, one must use more refined versions of Heegaard Floer homology, denoted HF^+ and HF^-.)

Informally, the relationship between the 4-dimensional invariants and their 3-dimensional counterparts is as follows: The Floer homology associates a graded abelian group to a 3-manifold Y. If W is a 4-manifold with $\partial W = Y$, then we associate to it a homology class in the Floer homology of Y. Now if X is a *closed* 4-manifold which decomposes as $X = X_1 \cup_\partial X_2$, then the 4-manifold invariant is a number coming from the pairing of the homology classes of X_1 and X_2.

In the Heegaard Floer case, which is the one most relevant to this thesis, this looks like the following. Every closed and oriented smooth 3-manifold Y is assigned an $\mathbb{F}_2$-vector space $\widehat{HF}(Y)$ which is obtained as the homology of some chain complex. Then a (4-dimensional) cobordism W between Y_1 and Y_2 is associated a linear map $F_W : \widehat{HF}(Y_1) \to \widehat{HF}(Y_2)$. Furthermore, this assignment fits into the framework of a $(3+1)$-dimensional topological quantum field theory (TQFT) as in [Ati88]. In particular, we have the following. Let $\mathrm{Cob}(3)$ be the category whose objects are closed, oriented, smooth 3-manifolds and whose morphisms are smooth, oriented, and connected 4-manifold cobordisms. Then $\widehat{HF}$ is a functor from $\mathrm{Cob}(3)$ to the category of $\mathbb{F}_2$-vector spaces. On objects, this is the map $Y \mapsto \widehat{HF}(Y)$, while on morphisms it is the map $W \mapsto F_W$.

It is natural to ask whether a $(3+1)$-TQFT may be extended to a $(2+1+1)$-TQFT, i.e., if we can extend to a map taking 2-manifolds to algebras over $\mathbb{F}_2$ in such a way that the composition axioms of a 2-category are satisfied. If so, one could recover Heegaard Floer invariants of 3-manifolds with boundary by considering them to be cobordisms between 2-manifolds.

Bordered Heegaard Floer homology, an invariant due to Robert Lipshitz, Peter Ozsváth, and Dylan Thurston [LOT18] and the subject of this thesis, is a step in this direction. The 2-manifold invariant it defines ends up being a differential graded algebra which depends on some extra data, namely a choice of parameterization, and thus is not a genuine topological invariant. Similarly, bordered Heegaard Floer homology does extend Heegaard Floer homology to manifolds with one boundary component, but is only an invariant of so-called "bordered 3-manifolds." These are manifolds whose boundary component have been parameterized by a choice of handle decomposition. Furthermore, we may recover Heegaard Floer homology by decomposing a closed manifold into two manifolds with boundary and computing the tensor product of their bordered Heegaard Floer invariants.

1.1 The bordered Heegaard Floer package

Consider a closed, oriented 2-manifold F which has been equipped with a certain parameterization. This parameterization is given by a *pointed matched circle* $\mathcal{Z}$. One may associate $\mathcal{Z}$ to a differential graded algebra $\mathcal{A}(\mathcal{Z})$. This algebra is not itself an invariant of the surface F, but gives rise to bordered Heegaard Floer homology invariants.

In particular, if Y is a 3-manifold equipped with an orientation-preserving diffeomorphism $\phi : F \to \partial Y$, then one can associate two algebraic objects to Y. (One calls such a manifold (Y, ϕ) a *bordered 3-manifold*.) The first is the *type A module* $\widehat{CFA}(Y)$, which is a right $\mathcal{A}_\infty$ module over $\mathcal{A}(\mathcal{Z})$. (An $\mathcal{A}_\infty$ module is like a module, but might not satisfy associativity on the nose. Instead, it satisfies associativity up to a homotopy which is encoded by some higher multiplication m_3. This map m_3 itself satisfies an associativity-type requirement only up to homotopy as well, and so on.) The second module associated to Y is the *type D module* $\widehat{CFD}(Y)$, which is a differential graded left $\mathcal{A}(-\mathcal{Z})$-module. This is an on-the-nose differential module, rather than an $\mathcal{A}_\infty$ module. However, it has a further algebraic structure called a *type D structure*.

This type D structure gives rise to a pairing result. If Y_1 and Y_2 are two bordered 3-manifolds with $\partial Y_1 = -\partial Y_2$, then $\widehat{HF}(Y_1 \cup_\partial Y_2) = \widehat{CFA}(Y_1) \widetilde{\otimes} \widehat{CFD}(Y_2)$. (Note that we write $\widetilde{\otimes}$ instead of the typical tensor product $\otimes$. This is because the pairing theorem involves the $\mathcal{A}_\infty$ tensor product instead.) In theory, this gives a more efficient way to compute the Heegaard Floer homology of a closed 3-manifold.

We briefly remark on the construction of the bordered Heegaard Floer invariants $\widehat{CFA}$ and $\widehat{CFD}$. One defines them on a Heegaard diagram $\mathcal{H}$ representing a bordered 3-manifold Y. This diagram comprises a surface $\overline{\Sigma}$ with boundary $\partial\overline{\Sigma} = S^1$ parameterized by a pointed matched circle $\mathcal{Z}$, along with some curves and arcs with boundary on $\partial\overline{\Sigma}$. These curves and arcs tell us how to attach 2-handles to $\overline{\Sigma}$ to obtain Y. Both $\widehat{CFA}$ and $\widehat{CFD}$ are generated by $g(\overline{\Sigma})$-tuples of intersection points of the curves and arcs. (In particular, the generators of $\widehat{CFA}$ and $\widehat{CFD}$ are finite subsets of $\overline{\Sigma}$.) The multiplication maps in $\widehat{CFA}$ and the differential in $\widehat{CFD}$ are defined by counting the number of elements in a certain moduli space of holomorphic curves in $\overline{\Sigma} \times [0,1] \times \mathbb{R}$. Just as how the differential in Morse homology counted flowlines between generators (i.e., critical points), the maps in bordered Heegaard Floer homology count the number of holomorphic curves which connect generators of $\widehat{CFA}$ and $\widehat{CFD}$.

While bordered Heegaard Floer homology involves somewhat sophisticated algebraic constructions, most notably $\mathcal{A}_\infty$ modules (which are not particularly common in low-dimensional topology, though they come up quite often in symplectic geometry) and type D structures (which were created expressly for bordered Heegaard Floer homology), we focus in this thesis on the geometry of the subject instead. In particular, these holomorphic curves between generators form the geometric heart of bordered Heegaard Floer homology, as well as the technical heart of this thesis. As such, we discuss holomorphic curve theory—and, even more specifically, compactness results—at some length.

Because of this, beyond the standard undergraduate curriculum, it is useful for a reader to have familiarity with basic algebraic topology and symplectic geometry. It may also be helpful, though not necessary, for a reader to have seen some pseudoholomorphic curve theory, e.g., some results from [MS12]. It would also be helpful to not be colorblind, particularly not red-blue colorblind (which I have recently learned does, in fact, exist). I am sorry that this thesis, or at least the figures therein, will not be particularly readable otherwise.

1.2 Organization

We begin in Chapter 2 by introducing the geometric objects which will be used define our bordered Heegaard Floer invariants, namely pointed matched circles and Heegaard diagrams. (It turns out that the definition is independent of the choice of Heegaard diagram which induce the same pointed matched circle, so that $\widehat{CFA}$ and $\widehat{CFD}$ are genuine invariants of 3-manifolds with parameterized boundary.) We will also define the set $\mathfrak{S}(\mathcal{H})$; its elements will eventually be the generators of our type A and type D modules.

In Chapter 3, we give a description of the moduli spaces whose curves we will count in the definition of the maps in $\widehat{CFA}$ and $\widehat{CFD}$. We begin by defining these moduli spaces. To prove that $\widehat{CFA}$ and $\widehat{CFD}$ have the desired algebraic properties, we need to prove a few compactness and gluing results. We omit the latter (as is perhaps standard), but spend some time discussing the former. To do so, we bring in ideas from symplectic field theory.

To emphasize the focus on the geometric aspect of bordered Heegaard Floer homology, we postpone nearly all the algebra until the very end, in Chapter 4. This is where we define $\mathcal{A}_\infty$ modules, as well as the algebra $\mathcal{A}(\mathcal{Z})$ associated to a pointed matched algebra. We then define $\widehat{CFA}$ and $\widehat{CFD}$, and conclude with a brief sketch of the pairing theorem, which reconstructs closed Heegaard Floer homology $\widehat{HF}$ from the type A and type D modules. We focus throughout primarily on examples, and on how the algebraic features of the type A and type D modules correspond to elements of the moduli spaces discussed in the previous chapter.

We are obliged at this point to mention that the curious reader would find a more complete treatment of this subject in the original monograph by Lipshitz, Ozsváth, and Thurston [LOT18]. We do not aim to replace their exposition, but instead hope to serve as a useful and streamlined supplement, filling in certain gaps while omitting some details which are less relevant to the actual geometry of the definition.

Chapter 2

Heegaard diagrams

> We knew hitherto only a superficial image; behold it has gained depth, it extends into three dimensions, it moves.
>
> Marcel Proust, *The Guermantes Way*

Heegaard diagrams, first defined by Poul Heegaard in [Hee98], are a way to encode 3-manifolds one dimension down. They are the essential objects of study in the definition of Heegaard Floer homology. In particular, Heegaard Floer homology is defined for a Heegaard diagram. (It turns out, of course, to be an invariant of the 3-manifold represented by the diagram.)

As such, we begin in Section 2.1 by defining Heegaard diagrams for closed 3-manifolds. Bordered Heegaard Floer homology generalizes Heegaard Floer homology to so-called "bordered 3-manifolds." We discuss these objects in Section 2.2. Roughly speaking, bordered manifolds are just manifolds with one parameterized boundary component. In analogy to Heegaard diagrams, we may define bordered Heegaard diagrams of a bordered 3-manifold. We do so in Section 2.3. In Section 2.4, we define the generators of both the Heegaard Floer and bordered Heegaard Floer complexes. We also introduce certain admissibility conditions which our (bordered) Heegaard diagrams must fulfill in order for them to be well-suited for defining (bordered) Heegaard Floer homology. In particular, these admissibility requirements ensure that certain sums which appear in the definition of the differential remain finite. Finally, in Section 2.5, we give an informal description of Lipshitz's cylindrical reformulation of Heegaard Floer homology [Lip06a]. This introduces the moduli spaces which we will define more rigorously and study in depth in the following chapter.

2.1 Heegaard diagrams of closed 3-manifolds

A Heegaard diagram $\mathcal{H}$ is a way to represent a closed 3-manifold Y via a surface and some collection of curves. Roughly speaking, these curves tell us how to attach 3-balls to the surface, and hence how to construct a 3-manifold. More precisely, a Heegaard diagram encodes a so-called "handle decomposition" of Y; the attached 3-balls are known as "handles."

More precisely, for $0 \leq k \leq n$, recall that an **n-dimensional k-handle** is a copy of $D^n = D^k \times D^{n-k}$ which is attached to an n-dimensional manifold M by some embedding $\partial D^k \times D^{n-k} \hookrightarrow \partial M$. There is a canonical way to smooth corners, so the resulting object may be considered as an n-dimensional manifold

as well. It is homotopically the same as attaching a k-cell; indeed, a k-handle may be thought of as a "thickened" k-cell. Note that a 0-handle simply looks like an n-ball, without any attaching map. With this setup, we call the image of $\partial D^k \times \{0\}$ in ∂M the **attaching sphere**.

If we may write M as the union of some handles, glued to each other via attaching maps, then we say that we have a **handle decomposition** of M. Any handle decomposition must have at least one 0-handle, as all other handles must attach to a preexisting manifold. In fact, any closed n-manifold M admits a handle decomposition into n-dimensional k-handles. (The same actually holds for manifolds with boundary, but we focus for now on the closed case.)

To see this fact, recall that a Morse function on M is a smooth map $f : M \to \mathbb{R}$ whose critical points are nondegenerate. It is a well-known fact that any M admits a Morse function. The topology of the sublevel sets $\{x : f(x) \leq c\}$ changes by attaching an n-dimensional k-handle when c passes through a critical value corresponding to a critical point of index k. Thus a Morse function defines a handle decomposition for M, as in Figure 2.1. In fact, any M actually admits a *self-indexing* Morse function, i.e., a Morse

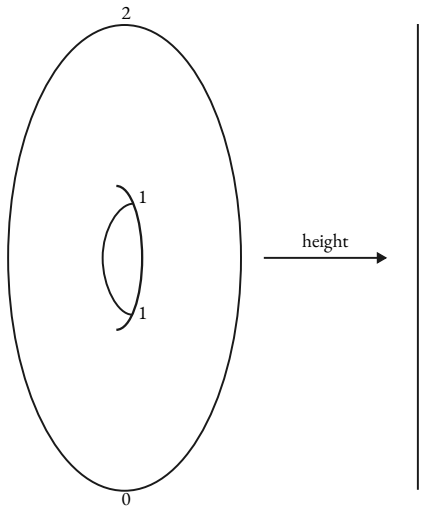

Figure 2.1: The height function above is a Morse function, and gives a decomposition of $S^1 \times S^1$ into one 0-handle, two 1-handles, and one 2-handle. Note that this function isn't self-indexing.

function f such that $f(p) = k$ for any index-k critical point p of f. This means that we may attach all the 0-handles first, then all the 1-handles at the same time, and so on. Details about handle decompositions and Morse functions may be found in [GS99, Chapter 4] and [Mil63, Part 1].

Now let us restrict our attention to the $n = 3$ case. Consider a particular handle decomposition of M. We may require that there is only one 0-handle and one 3-handle. Let U be the union of the 0- and 1-handles, and V the union of the 2- and 3-handles. We call U and V "handlebodies."

In particular, for our purposes, a **handlebody** is a single 3-ball (i.e., 0-handle) with some number of 1-handles attached. For example, the genus g handlebody is simply the 3-manifold with boundary which is bounded by the usual genus g surface. With this definition of a handlebody, U is clearly a handlebody. Turning V upside down, we see that V is also composed of a single 0-handle and several 1-handles, hence is also a handlebody. Thus we may write $M = U \cup_\Sigma V$, where Σ is the closed surface which is the common boundary of U and V. Such a decomposition of M into two handlebodies is known as a **Heegaard decomposition** or a **Heegaard splitting**.

Let $M = U \cup_\Sigma V$ be a Heegaard decomposition. Let $\boldsymbol{\alpha} = \{\alpha_1, \ldots, \alpha_g\}$ and $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_g\}$ be

the collections of attaching curves for the 1-handles in U and V, respectively. (Note that these 1-handles are being attached to Σ as 2-handles, as their attaching spheres are circles.) These are curves living on Σ. Given only Σ, $\boldsymbol{\alpha}$, and $\boldsymbol{\beta}$, we can recover M by attaching 1-handles to the curves and capping off the boundary components with 3-balls. With this in mind, we may make the following definition.

Definition 2.1. Consider a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ consisting of a compact oriented genus-g surface Σ without boundary, along with two sets $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ of g many disjoint closed curves. Suppose the surfaces $\Sigma \setminus \boldsymbol{\alpha}$ and $\Sigma \setminus \boldsymbol{\beta}$ are both connected. Suppose furthermore that the $\boldsymbol{\alpha}$- and $\boldsymbol{\beta}$-curves intersect transversely. Then $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a **(closed) Heegaard diagram**. Furthermore, if the closed manifold M is obtained from $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ as described above, then we say that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a Heegaard diagram **representing** M.

In general, there are many Heegaard decompositions for M, and many Heegaard diagrams for a given Heegaard decomposition (Σ, U, V). For example, we may write S^3 as the union of two 3-balls. Alternatively, we may write it as the union of two solid tori. After all, we know that $S^3 = \partial D^4 = \partial(D^2 \times D^2) = (S^1 \times D^2) \cup (S^1 \times D^2)$. Geometrically, this splitting is seen in Figure 2.2. The Heegaard diagram of this

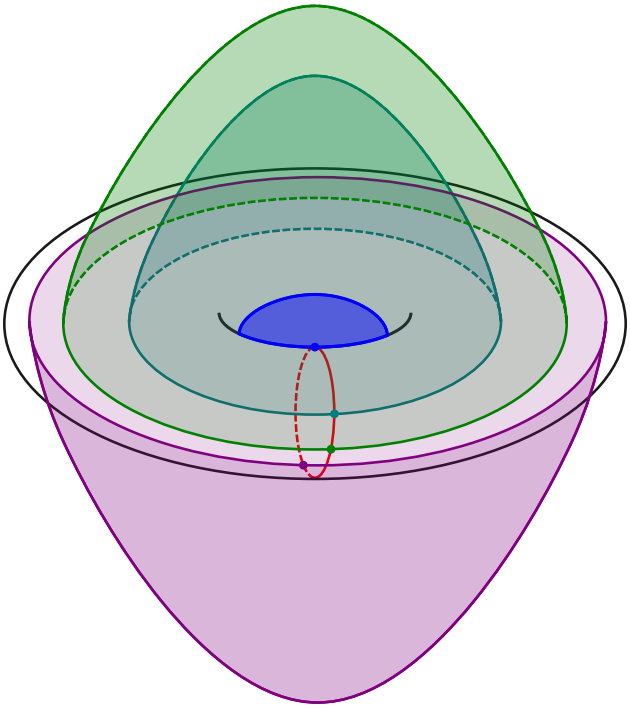

Figure 2.2: A Heegaard splitting of $S^3 = (S^1 \times D^2) \cup (S^1 \times D^2)$. One copy of $S^1 \times D^2$ is indicated by the black torus. The other copy is given by slowly blowing the blue disk up; each "blown-up" disk intersects the red circle at one point. Thus the colorful region is another copy of $S^1 \times D^2$.

splitting is seen in the right side of Figure 2.3. Another example of a Heegaard diagram, this time for the 3-manifold $S^1 \times S^2$, is $(S^1 \times D^2, \boldsymbol{\beta}, \boldsymbol{\beta})$, where $\boldsymbol{\beta}$ refers to the meridian of $S^1 \times D^2$. This Heegaard diagram corresponds to the decomposition of $S^1 \times S^2$ as $(S^1 \times D^2) \cup (S^1 \times D^2)$.

Heegaard Floer homology is an invariant of 3-manifolds which is defined through data in a Heegaard diagram. For this to make sense, we need the following proposition.

Proposition 2.2. *Any two Heegaard diagrams for M may be related by a sequence of Heegaard moves, namely isotopy, handleslides, and (de)stabilizations.*

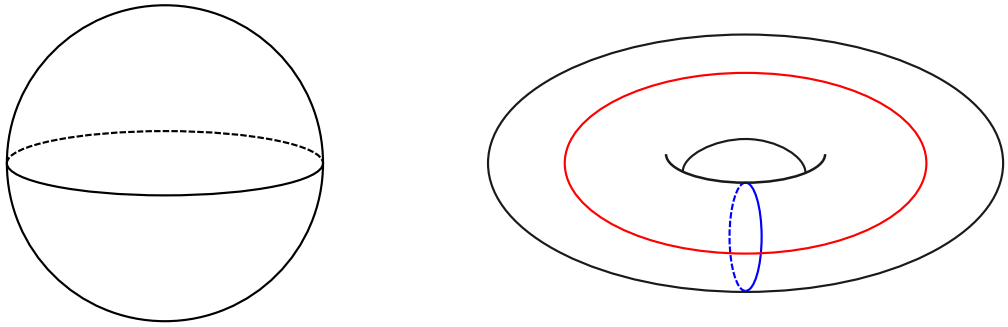

Figure 2.3: Heegaard diagrams of two splittings of S^3, namely the trivial $S^3 = D^3 \cup D^3$ and the splitting depicted above in Figure 2.2. In the first case, we have $\boldsymbol{\alpha} = \boldsymbol{\beta} = \emptyset$. In the second case, $\boldsymbol{\alpha}$ is the red longitude, while $\boldsymbol{\beta}$ is the blue meridian.

The first two Heegaard moves mentioned above relate Heegaard diagrams of the same Heegaard decomposition, while stabilizations relate Heegaard decompositions with a surface Σ of genus g to decompositions with a surface of genus $g + 1$. It is not too important for our purposes to know exactly what these moves are. The point is simply that an invariant of Heegaard diagrams which is preserved by all three Heegaard moves is also an invariant of 3-manifolds.

We will actually work with pointed Heegaard diagrams to define Heegaard Floer homology.

Definition 2.3. A **pointed Heegaard diagram** $\mathcal{H}$ is a quadruple $(\Sigma, \alpha, \beta, z)$ where (Σ, α, β) is a Heegaard diagram, and z is a point in $\Sigma \setminus (\alpha \cup \beta)$.

Any two pointed Heegaard diagrams for the same Heegaard decomposition may be connected by a sequence of pointed Heegaard moves, namely pointed isotopies, pointed handleslides, and (de)stabilizations. The first two are simply pointed generalizations of the Heegaard moves in Proposition 2.2. In particular, pointed isotopies are isotopies which do not cross the basepoint z, and similarly for pointed handleslides.

2.2 Pointed matched circles and bordered 3-manifolds

Our main object of study will be bordered 3-manifolds. These are 3-manifolds with connected boundary whose boundary contains some extra information parameterized by a so-called "pointed matched circle." This pointed matched circle encodes a handle decomposition of a surface in much the same way as how a Heegaard diagram encodes a handle decomposition of a 3-manifold.

Consider a closed orientable surface F of genus g. Then it admits a self-indexing Morse function with one index-0 critical point, $2g$ index-1 critical points, and one index-2 critical point. To specify F, it is enough to specify how the $2g$ 1-handles attach to the 0-handle. After all, there is then a unique way to glue the index-2 critical point, namely in such a way that the result has no boundary. But to specify how a 1-handle is attached to a 0-handle D^2, it suffices to specify an embedding of $S^0 = \partial D_1$ into the boundary of the 0-handle. Thus it suffices to specify $4g$ points on $\partial D^2 = S^1$, as well as a "matching" which tells us which pairs of points belonged to the same 1-handle. Finally, we need the result of attaching 1-handles to have connected boundary, so that we may glue a single 2-handle to obtain a closed surface.

Thus we make the following definition.

Definition 2.4. Consider a triple $\mathcal{Z} = (Z, \mathbf{a}, M)$ comprising an oriented circle Z, a set $\mathbf{a} = \{a_1, \dots, a_{4k}\}$ of $4k$ points in Z, and a 2-to-1 function $M : \mathbf{a} \to \{1, \dots, 2k\}$ for some $k \geq 1$. We call M a **matching**. It defines $2k$ pairs of a_i's. If the 1-dimensional manifold obtained by surgery along each of these pairs, thought of as a 0-sphere, is connected, then we say that $\mathcal{Z}$ is a **matched circle**. If $(Z, \mathbf{a}, M)$ is a matched circle and $z \in Z \setminus \mathbf{a}$, then we call $(Z, \mathbf{a}, M, z)$ a **pointed matched circle**.

If F is obtained from $\mathcal{Z}$ by gluing 1-handles to the pairs specified by M, then we write $F = F(\mathcal{Z})$. Note that, if $\mathcal{Z} = (Z, \mathbf{a}, M)$ has $4k$ points in $\mathbf{a}$, then F has genus $2k$. Figure 2.4 shows an example and a non-example of a matched circle.

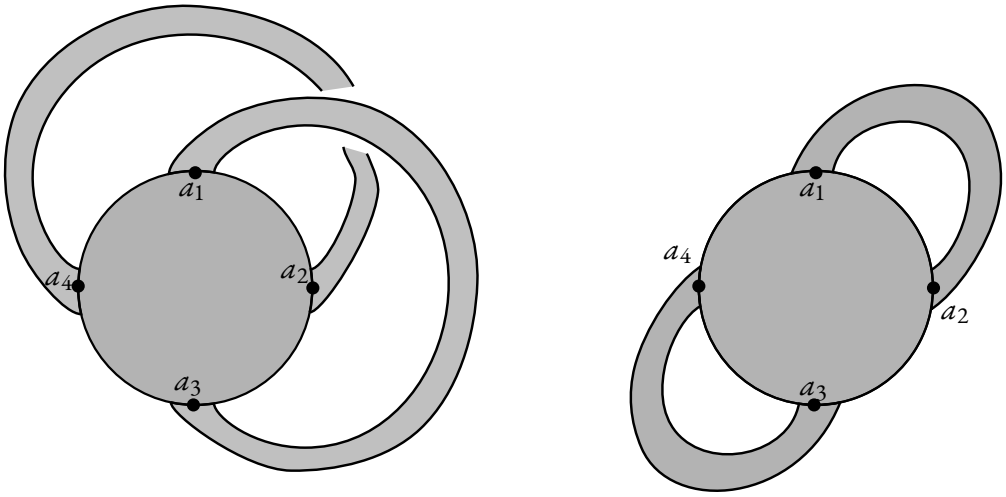

Figure 2.4: The left side, namely $\mathcal{Z} = (Z, \mathbf{a}, M)$ with matching $M(a_1) = M(a_3) = 1$ and $M(a_2) = M(a_4) = 2$, is a matched circle representing the genus 1 surface. The right side, with matching $M(a_1) = M(a_2) = 1$ and $M(a_3) = M(a_4) = 2$, is not a matched circle, since there is more than one boundary component (i.e., surgery on the pairs (a_1, a_2) and (a_3, a_4) produces a disconnected 1-manifold).

Definition 2.5. A **bordered 3-manifold** is a triple $(Y, \mathcal{Z}, \phi)$ where Y is a compact oriented 3-manifold with one boundary component, $\mathcal{Z}$ is a pointed matched circle, and $\phi : F(\mathcal{Z}) \to \partial Y$ is an orientation-preserving homeomorphism.

We close this section with a discussion of Reeb chords. These will be important in our definition of (bordered) Heegaard Floer homology since the differentials will be defined by moduli spaces consisting of maps which, among other things, converge to tuples of Reeb chords.

Definition 2.6. A **Reeb chord** ρ in $(Z \setminus z, \mathbf{a})$ is an embedded arc in $Z \setminus z$ whose endpoints are points in $\mathbf{a}$ and whose orientation is induced by the orientation on Z.

We call these chords "Reeb chords" because we may think of Z as a contact 1-manifold, and the points $\mathbf{a}$ as a Legendrian submanifold. The chords in question are then Reeb chords under the usual definition.

We denote the initial and terminal point of ρ as ρ^- and ρ^+, respectively. A set of Reeb chords is **consistent** if none of the Reeb chords share either an initial point or a terminal point. That is, we call $\boldsymbol{\rho} = \{\rho_1, \dots, \rho_n\}$ consistent if $\rho_i^- \neq \rho_j^-$ and $\rho_i^+ \neq \rho_j^+$ for any $i \neq j$.

Consider the points in $\mathbf{a}$ to be in increasing order from the basepoint z as we go around Z (respecting the orientation on Z). Then we call two Reeb chords ρ and σ **nested** if either $\rho^- < \sigma^- < \sigma^+ < \rho^+$ or $\sigma^- < \rho^- < \rho^+ < \sigma^+$. We call them **interleaved** if σ^+ and ρ^+ are swapped, i.e., if either $\rho^- < \sigma^- < \rho^+ < \sigma^+$ or $\sigma^- < \rho^- < \sigma^+ < \rho^+$. See Figure 2.5.

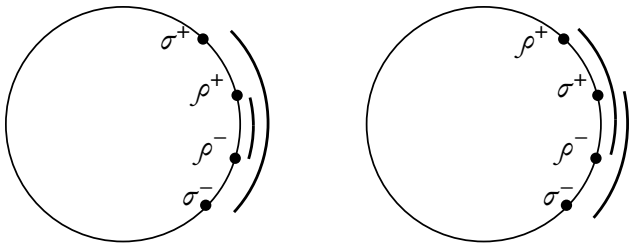

Figure 2.5: The left side shows two nested Reeb chords, while the right side shows two interleaved Reeb chords. Note that the circles are oriented counterclockwise here.

Finally, we may define an operation on Reeb chords as follows. If ρ and σ are **abutting** Reeb chords in the sense that $\rho^+ = \sigma^-$, then their **join** $\rho \uplus \sigma$ is the concatenation, i.e., the Reeb chord from ρ^- to σ^+.

2.3 Bordered Heegaard diagrams

A bordered Heegaard diagram is analogous to a Heegaard diagram, only for bordered manifolds. In particular, the main difference between a closed Heegaard diagram and a bordered one is that the surface in the bordered case has a boundary component, and the α-curves are allowed to be arcs now. Thus we make the following definition.

Definition 2.7. A **(pointed) bordered Heegaard diagram** is a quadruple $\mathcal{H} = (\overline{\Sigma}, \boldsymbol{\overline{\alpha}}, \boldsymbol{\beta}, z)$ where

- $\overline{\Sigma}$ is a compact oriented surface with one boundary component and genus g;
- $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_g\}$ is a g-tuple of pairwise-disjoint circles in the interior of $\overline{\Sigma}$;
- $\boldsymbol{\overline{\alpha}} = \{\alpha_1^c, \dots, \alpha_{g-k}^c, \overline{\alpha}_1^a, \dots, \overline{\alpha}_{2k}^a\}$ is a set of $g+k$ many pairwise-disjoint curves in $\overline{\Sigma}$ where the α_i^c's are circles in the interior and the α_i^a's are arcs with boundary in $\partial\overline{\Sigma}$ which are transverse to $\partial\overline{\Sigma}$; and
- z is a point in $\partial\overline{\Sigma} \setminus (\boldsymbol{\overline{\alpha}} \cap \partial\overline{\Sigma})$,

such that $\boldsymbol{\overline{\alpha}}$ and $\boldsymbol{\beta}$ intersect transversely, $\overline{\Sigma} \setminus \boldsymbol{\overline{\alpha}}$ is connected, and $\overline{\Sigma} \setminus \boldsymbol{\beta}$ is connected.

An example of a bordered Heegaard diagram for the genus 1 handlebody is depicted in Figure 2.6. We may obtain a bordered Heegaard diagram for the genus-2 handlebody by taking the boundary connect sum of two copies of the genus-1 diagram, as shown in Figure 2.7. A different bordered Heegaard diagram for the genus 2 handlebody is shown in Figure 2.8 by picking a different parameterization of the boundary. In general, a bordered Heegaard diagram without α-circles necessarily represents a genus g handlebody for some g.

Notice that the boundary of a bordered Heegaard diagram may be thought of as a pointed matched circle. See [LOT18, Lemma 4.4].

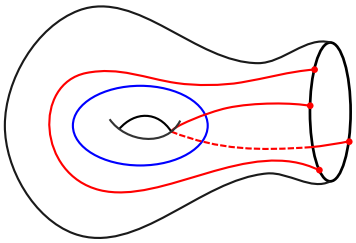

Figure 2.6: Here $\overline{\boldsymbol{\alpha}} = \{\overline{\alpha}_1^a, \overline{\alpha}_2^a\}$ consists of the two red arcs and $\boldsymbol{\beta} = \{\beta_1\}$ consists of the single blue circle. In particular, there are no α-circles α_i^c. This is a bordered Heegaard diagram for a genus 1 handlebody. (We may let z be any point of $\partial\overline{\Sigma}$ which is not on the red α-arcs.)

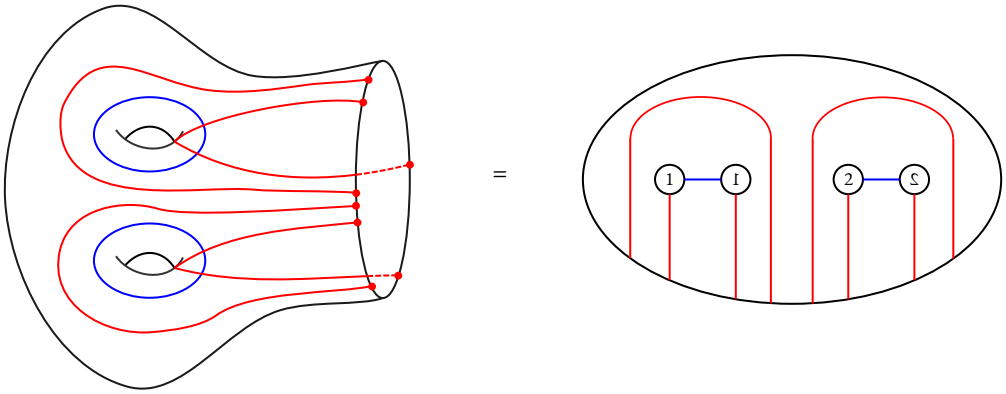

Figure 2.7: A bordered Heegaard diagram for the genus-2 handlebody. Here (and later on), we use circles with numbers to indicate where a 1-handle $D^1 \times D^2$ is attached. In particular, we imagine the plane of this sheet of paper to be the boundary $S^2 = \mathbb{R}^2 \cup \{\infty\}$ of a 0-handle D^3. We actually lop off half of this 0-handle, since we have a manifold with boundary. The black oval denotes this boundary circle, so the interior denotes the boundary of this lopped-off 0-handle. Then attaching a 1-handle corresponds to embedding $\partial D^1 \times D^2 = D^2 \amalg D^2$ into this oval.

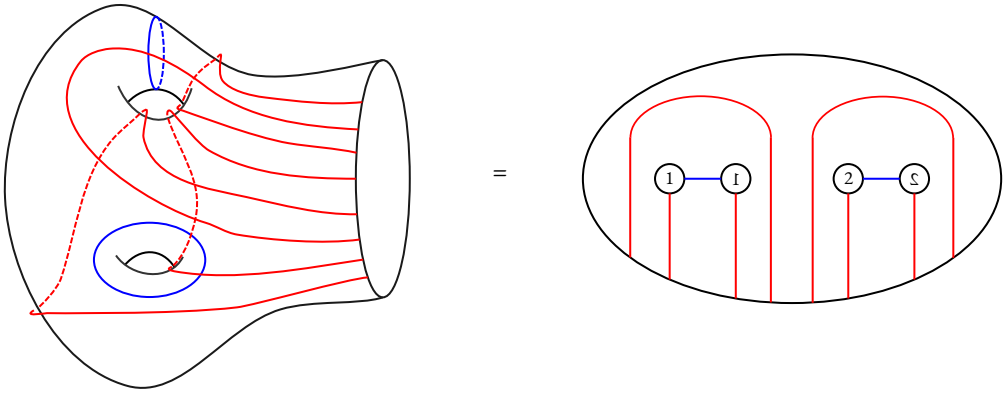

Figure 2.8: A different bordered Heegaard diagram for the genus 2 handlebody.

Lemma 2.8. *Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a bordered Heegaard diagram. Let $Z = \partial\Sigma$ and $\mathbf{a} = \boldsymbol{\alpha} \cap Z$. Consider the matching $M : \mathbf{a} \to \{1, \dots, 2k\}$ which takes $\alpha_i^a \cap Z$ to i. Then $(Z, \mathbf{a}, M, z)$ is a pointed matched circle.*

We denote the pointed matched circle from Lemma 2.8 by $\partial\mathcal{H}$, and call it the **boundary of** $\mathcal{H}$. For example, if $\mathcal{H}$ denotes the bordered Heegaard diagram from Figure 2.6, then we have that $\partial\mathcal{H}$ is the pointed matched circle in the first row of Figure 2.4.

A bordered Heegaard diagram gives rise to a bordered 3-manifold much as how a Heegaard diagram gives rise to a closed 3-manifold, namely by indicating attaching spheres for the 1- and 2-handles. In particular, consider the thickening $\Sigma \times [0,1]$ of the Heegaard surface Σ. Attach a 3-dimensional 2-handle to each $\alpha_i^c \times \{0\}$ and to each $\beta_i \times \{1\}$, and call the resulting manifold Y. Originally, the thickening had boundary $\Sigma \times \{0,1\} \cup \partial\Sigma \times [0,1]$. The result of surgering out the β-circles from $\Sigma \times \{1\}$ is a disk, while the result of surgering out the α-circles α_i^c from $\Sigma \times \{0\}$ is the genus-k surface with one boundary component. Thus ∂Y is exactly a closed genus-k surface.

In fact, there is a canonical identification of ∂Y with $F(\partial\mathcal{H})$. (Recall that $\partial\mathcal{H}$ is a pointed matched circle.) After all, $F(\partial\mathcal{H})$ is formed by a disk with boundary $\partial\Sigma$, 1-handles whose attaching spheres are the endpoints of the α-arcs, and a disk to close up the boundary component. The union of $\Sigma \times \{1\}$ with β-circles surgered out and the annulus $\partial\Sigma \times [0,1]$ is a disk in ∂Y with boundary $\partial\Sigma \times \{0\}$; this is the 0-handle in $F(\mathcal{Z})$. Considering the α-arcs to be the cores of the 2-dimensional 1-handles which are attached to the disk to form $F(\partial\mathcal{H})$, it follows that $\partial Y \cong F(\partial\mathcal{H})$.

An example of this is shown in Figure 2.9.

There is another interpretation, which introduces some canceling handles, which allows one to visualize $F(\mathcal{Z})$ a little more clearly, as opposed to as the union of three pieces (the two layers $\Sigma \times \{0\}$ and $\Sigma \times \{1\}$, with appropriate circles surgered out, and the cylinder $\partial\Sigma \times [0,1]$). See Figure 2.10. In some more detail, the construction is as follows: As before, we consider $\Sigma \times [0,1]$. Consider a collar neighborhood $A_1 := [-\varepsilon, 0] \times Z \subset \Sigma$ such that $\{0\} \times Z$ is identified with $Z = \partial\Sigma$. Recall that $F(\mathcal{Z})$ is obtained by attaching handles to some disk whose boundary is Z. Thus we may consider a tubular neighborhood $A_2 := Z \times [0,1] \subset F(\mathcal{Z})$. Then glue $\Sigma \times [0,1]$ to $[-\varepsilon, 0] \times F(\mathcal{Z})$ with the obvious identification $A_1 \times [0,1] = [-\varepsilon, 0] \times Z \times [0,1] = [-\varepsilon, 0] \times A_2$. Now attach 2-handles to each $\beta_i \times \{1\}$ and $\alpha_i^c \times \{0\}$, as before. We must cancel handles we added by gluing in $[-\varepsilon, 0] \times F(\mathcal{Z})$ In particular, each α-arc $\alpha_i^a \times \{0\}$ has a "counterpart" in $F(\mathcal{Z})$, and so we obtain a circle by taking the union of each such arc with its counterpart. We attach 2-handles along these circles. We finish off by attaching 3-handles (these are the top-

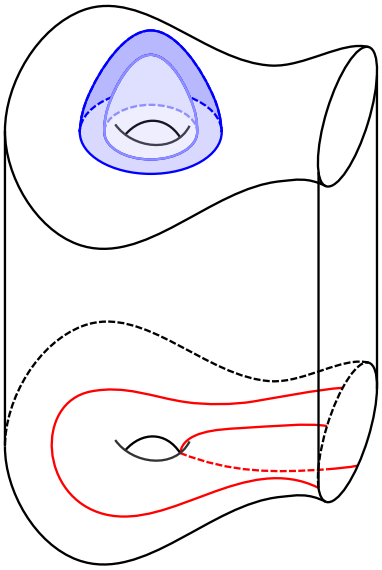

Figure 2.9: The bordered 3-manifold Y associated to the bordered Heegaard diagram in Figure 2.6. Its boundary consists of three parts. First, we have the top layer, namely $\overline{\Sigma} \times \{1\}$ with the β-circle surgered out), which is just a copy of D^2. Second, we have the cylinder (i.e., annulus) $Z \times [0,1]$. Finally, we have the bottom layer, which is identified with the surface $F(\mathcal{Z})$. Thus $\partial Y = F(\mathcal{Z})$.

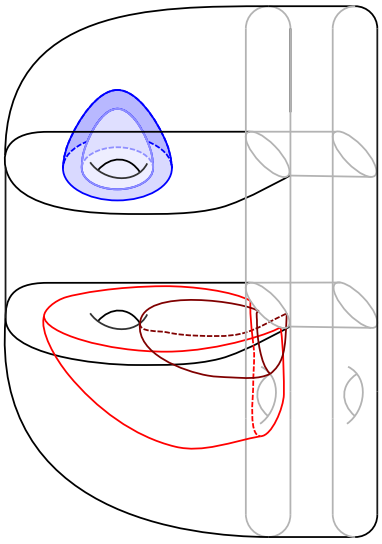

Figure 2.10: Another way, with some canceling handles, to see the bordered 3-manifold associated to a bordered Heegaard diagram. The boundary is now a bit easier to see: It is just the genus-1 handlebody on the very right side of the diagram.

and bottommost black arcs in Figure 2.10). This interpretation will be useful when we discuss the pairing theorem in Section 4.6.

Just as how there is a Morse theoretic picture for 3-manifolds, so too is there one for bordered 3-manifolds which shows that every bordered 3-manifold is represented by some bordered Heegaard diagram. To see this, we must introduce some canceling handles. For details, see [LOT18, Lemma 4.9].

Furthermore, we have the following analogue of Proposition 2.2.

Proposition 2.9 ([LOT18, Lemma 4.10])**.** *Any two bordered Heegaard diagrams for a bordered manifold may be related by a sequence of isotopies of α-curves and β-circles not crossing $\partial\overline{\Sigma}$, handleslides of α-curves over α-circles and β-circles over β-circles, and (de)stabilizations in the interior of $\overline{\Sigma}$.*

2.4 Generators and admissibility

In this section, we describe the generators of the Heegaard Floer chain complexes associated to a pointed or bordered Heegaard diagram, and state an important technical condition ("admissibility") that diagrams must satisfy for Heegaard Floer homology to be defined. We start with the closed case before moving on to the bordered case.

Let $\mathcal{H}$ denote the pointed Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ representing a *closed* 3-manifold Y. Say Σ has genus g.

Definition 2.10. A **generator** of $\mathcal{H}$ is a g-element set $\mathbf{x} = \{x_1, \dots, x_g\}$ of points on $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$ such that each α-circle contains exactly one x_i and similarly for each β-circle.

These points $\mathbf{x}$ will be the generators for the chain complex for Heegaard Floer homology.

In later sections, we will be interested in holomorphic curves with codomain $\Sigma \times [0,1] \times \mathbb{R}$. Let s be the $[0,1]$-coordinate and t the $\mathbb{R}$-coordinate. Then the differential will be obtained by counting curves with boundary on $C_\alpha = \boldsymbol{\alpha} \times \{1\} \times \mathbb{R}$ and $C_\beta = \boldsymbol{\beta} \times \{0\} \times \mathbb{R}$ and with appropriate asymptotics at $t = \pm\infty$. In particular, we consider homology classes in $H_2(\Sigma \times [0,1] \times \mathbb{R})$ Then consider the following definition.

Definition 2.11. Let $\pi_2(\mathbf{x}, \mathbf{y})$ denote the set of *homology* classes in $H_2(\Sigma \times [0,1] \times \mathbb{R}, C_\alpha \cup C_\beta)$ which converge to $\mathbf{x}$ and $\mathbf{y}$ at $t = -\infty$ and $t = \infty$, respectively. We call elements of this set **homology classes connecting x to y**.

Another way of thinking about this set is that, if we denote the closure $[-\infty, \infty]$ of $\mathbb{R}$ by $\overline{\mathbb{R}}$, then elements of $\pi_2(\mathbf{x}, \mathbf{y})$ are (complex) curves in $\Sigma \times [0,1] \times \overline{\mathbb{R}}$ with boundary in the union of C_α, C_β, $\mathbf{x} \times [0,1] \times \{-\infty\}$, and $\mathbf{y} \times [0,1] \times \{\infty\}$.

Definition 2.12. A **region** is a component of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$. Consider the projection of a homology class $B \in \pi_2(\mathbf{x}, \mathbf{y})$ to Σ. This gives a well-defined element of $H_2(\Sigma, \boldsymbol{\alpha} \cup \boldsymbol{\beta})$ which is a linear combination of regions. We call this linear combination the **domain of** B. The **local multiplicity** of B at a point $p \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ is the coefficient for the region containing p of the domain of B, and is denoted $n_p(B)$.

Since we are only concerned with $\widehat{HF}$, we only want to consider homology classes whose domain does not cross z, i.e., has local multiplicity 0 at z. Let $\widehat{\pi}_2(\mathbf{x}, \mathbf{y})$ be the classes B in $\pi_2(\mathbf{x}, \mathbf{y})$ such that this condition is true for the associated domain of B. An element of $\widehat{\pi}_2(\mathbf{x}, \mathbf{x})$ is a **periodic class**, and its

domain is called a **periodic domain**. Finally, we call a homology class B **positive** if all of the coefficients in its corresponding domain are nonnegative.

The differential will involve moduli spaces of curves in some fixed homology class $B \in \pi_2(\mathbf{x}, \mathbf{y})$. In particular, we must be able to count these moduli spaces. To ensure this, we must impose the following admissibility condition.

Definition 2.13. The pointed Heegaard diagram $\mathcal{H}$ is **(weakly) admissible** if every nontrivial periodic domain D which does not cross z has both positive and negative coefficients.

There is another definition of admissibility, known as *strong* admissibility, as in [OS04c]. However, because we only care about $\widehat{HF}$ and not any of the variants, we omit this definition. In particular, "admissibility" always refers to *weak* admissibility for us.

Proposition 2.14 ([OS04c, Lemma 5.8])**.** *A Heegaard diagram is isotopic to an admissible Heegaard diagram, and two admissible Heegaard diagrams may be connected by a sequence of Heegaard moves such that, at every stage, we have an admissible Heegaard diagram.*

We now have the following reformulations of admissibility in terms of area forms. Its proof follows from ideas in linear algebra; see, for example, [OS04c, Lemma 4.12].

Proposition 2.15. *A pointed Heegaard diagram is admissible if and only if there is an area function A on Σ such that $A(P) = 0$ for every periodic domain P.*

The upshot of this is that admissible Heegaard diagrams have finitely many positive homology classes $B \in \widehat{\pi}_2(\mathbf{x}, \mathbf{y})$.

Proposition 2.16. *If $\mathcal{H}$ is admissible, then for any two generators $\mathbf{x}$ and $\mathbf{y}$, there are only finitely many positive homology classes $B \in \widehat{\pi}_2(\mathbf{x}, \mathbf{y})$.*

Proof. Suppose $B, B' \in \widehat{\pi}_2(\mathbf{x}, \mathbf{y})$ are positive homology classes. Then $B - B'$ is a periodic domain. Letting A be an area function, as in Proposition 2.15, we have $A(B) = A(B')$. But there are only finitely many positive domain of a given area. □

This will be useful because the moduli space in Section 3.2 of curves in a given homology class B will be nonempty only if B is positive. The union over homology classes B of the moduli spaces will then be a finite, hence well-defined, union.

The above definitions all have generalizations to the bordered case. In particular, now let $\mathcal{H}$ be the bordered Heegaard diagram $(\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \boldsymbol{\beta}, z)$. In the bordered case, we let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be the result of attaching an infinite cylindrical end $\partial\overline{\Sigma} \times [0, \infty)$ to $\partial\overline{\Sigma}$. In particular, Σ is topologically equivalent to the punctured surface $\overline{\Sigma} \setminus \partial\overline{\Sigma}$, and the α_i's in $\boldsymbol{\alpha}$ are either α_i^c or $\alpha_i^a \setminus \partial\alpha_i^a$.

Definition 2.17. A **generator** of a bordered Heegaard diagram $\mathcal{H}$ is a g-element set $\mathbf{x} = \{x_1, \dots, x_g\}$ of points on $\overline{\boldsymbol{\alpha}} \cap \boldsymbol{\beta}$ such that

- each α-circle contains exactly one x_i;
- each β-circle contains exactly one x_i; and
- each α-arc contains at most one x_i.

Note that the first two conditions make up the definition for the closed case, as there are no α-arcs in that case. Again, our chain complexes will be generated by these intersection points $\mathbf{x}$. We denote the set of all generators by $\mathfrak{S}(\mathcal{H})$. Furthermore, if $\mathbf{x}$ is a generator, then we let $o(\mathbf{x}) := \{i : \mathbf{x} \cap \alpha_i^a \neq \emptyset\}$ denote the subset of $\{1, \dots, 2k\}$ consisting of those arcs which are occupied by some x_i.

In the bordered case, we would like to allow our holomorphic curves to have boundary on $C_\partial = (\partial\overline{\Sigma} \setminus z) \times [0,1] \times \mathbb{R}$. We remove z because the bordered case is in analogy to $\widehat{HF}$; thus our curves may not cross z. Then we may make the following definition.

Definition 2.18. The **set of homology classes connecting x and y**, denoted $\pi_2(\mathbf{x}, \mathbf{y})$, is the set of homology classes in $H_2(\Sigma \times [0,1] \times \mathbb{R}, C_\alpha \cup C_\beta \cup C_\partial)$ which converge to $\mathbf{x}$ and $\mathbf{y}$ at $t = -\infty$ and $t = \infty$, respectively.

In analogy to the above definitions, we may call components of $\overline{\Sigma} \setminus (\overline{\boldsymbol{\alpha}} \cup \boldsymbol{\beta})$ **regions**. Furthermore, if $B \in \pi_2(\mathbf{x}, \mathbf{y})$, then its projection to $\overline{\Sigma}$ gives a linear combination of regions called the **domain of** B. The **local multiplicity** is defined the same way as before.

Because we ask that the boundary of the homology classes (thought of as two-chains) in question avoids $z \in \partial\overline{\Sigma}$, we know that the local multiplicity at z is always 0. Thus we need not define $\widehat{\pi}_2(\mathbf{x}, \mathbf{y})$ in the bordered case, as it is the same as $\pi_2(\mathbf{x}, \mathbf{y})$. Now we call elements of $\pi_2(\mathbf{x}, \mathbf{x})$ **periodic classes**, and their domains **periodic domains**. The definition of a **positive** homology class is the same as before, namely that the coefficients in its corresponding domain are all positive.

We now have the following analogous admissibility condition.

Definition 2.19. A bordered Heegaard diagram is **admissible** if every nontrivial periodic domain has both positive and negative coefficients.

At times, we will use a slightly weaker admissibility condition. To describe it, we must first define a kind of domain called "provincial."

Consider the domain of a class $B \in \pi_2(\mathbf{x}, \mathbf{y})$. Its boundary is composed of three pieces, namely the piece in $\overline{\boldsymbol{\alpha}}$, the piece in $\boldsymbol{\beta}$, and the piece in $\partial\overline{\Sigma}$. Denote the pieces as $\partial^\alpha B$, $\partial^\beta B$, and $\partial^\partial B$, respectively. Furthermore, we orient them such that the domain of B has oriented boundary $\partial^\alpha B + \partial^\beta B + \partial^\partial B$. We may think of $\partial^\partial B$ as an element of $H_1(\partial\overline{\Sigma}, \mathbf{a})$, where $\mathbf{a}$ consists of the $4k$ endpoints of the α-arcs. (This is the same as the set $\mathbf{a}$ of points of the pointed matched circle corresponding to the boundary of $\overline{\Sigma}$.)

Definition 2.20. A class $B \in \pi_2(\mathbf{x}, \mathbf{y})$ is **provincial** if $\partial^\partial B = 0$, i.e., if its domain contains no regions bordering $\partial\overline{\Sigma}$. A bordered Heegaard diagram is **provincially admissible** if every nontrivial provincial periodic domain has both positive and negative coefficients.

Note that provincial admissibility is a weaker condition than admissibility, as we have a condition on provincial periodic domains, rather than on all periodic domains.

We now have analogues to Propositions 2.14 and 2.15, as follows. Their proofs are exactly analogous to the proofs of those statements.

Proposition 2.21. *A bordered Heegaard diagram is isotopic to an admissible (respectively, provincially admissible) bordered Heegaard diagram. Furthermore, two admissible (respectively, provincially admissible) bordered Heegaard diagrams may be connected by a sequence of Heegaard moves such that, at every stage, we have an admissible (respectively, provincially admissible) bordered Heegaard diagram.*

Proposition 2.22. *A bordered Heegaard diagram is admissible (respectively, provincially admissible) if and only if there is an area form A such that $A(P) = 0$ for every periodic (respectively, provincial periodic) domain P.*

The point of these admissibility definitions, as well as the area reformulation of admissibility. is to, as in the closed case, attain some finiteness result. The area reformulation implies the following two results.

Proposition 2.23. *Suppose that $\mathcal{H}$ is a provincially admissible bordered Heegaard diagram. Let $\mathbf{x}$ and $\mathbf{y}$ be two generators. Let $h \in H_1(Z, \mathbf{a})$, where Z and $\mathbf{a}$ are as in Lemma 2.8. Then there are only finitely many positive $B \in \pi_2(\mathbf{x}, \mathbf{y})$ with $\partial^\partial B = h$.*

Proposition 2.24. *Suppose that $\mathcal{H}$ is an admissible bordered Heegaard diagram. Let $\mathbf{x}$ and $\mathbf{y}$ be two generators. Then there are only finitely many positive $B \in \pi_2(\mathbf{x}, \mathbf{y})$.*

Remark 2.25. It is reasonable to ask when $\pi_2(\mathbf{x}, \mathbf{y})$ (or, in the closed case, when $\widehat{\pi}_2(\mathbf{x}, \mathbf{y})$) is nonempty. It turns out that this has to do with spin^c structures on Y. In particular, each generator $\mathbf{x}$ gives rise to a nonvanishing vector field, hence a spin^c structure, on Y. (This uses an interpretation of spin^c structures due to Turaev [Tur97].) We denote this associated spin^c structure as $\mathfrak{s}_z(\mathbf{x})$. Then one can show that $\pi_2(\mathbf{x}, \mathbf{y}) \neq \emptyset$ (or, in the closed case, $\widehat{\pi}_2(\mathbf{x}, \mathbf{y}) \neq \emptyset$) if and only if $\mathfrak{s}_z(\mathbf{x}) = \mathfrak{s}_z(\mathbf{y})$. There is a similar spin^c condition for when there is a provincial domain connecting generators $\mathbf{x}$ and $\mathbf{y}$.

2.5 Heegaard Floer homology

In this thesis, we will focus on the bordered case. But we will sketch out Lipshitz's cylindrical reformulation of Heegaard Floer homology for closed 3-manifolds, which will motivate our next chapter on moduli spaces. From a strictly logical perspective, the following definitions and results would follow from the technical details in Sections 3.1 to 3.5.

Fix a pointed Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ for the (closed) 3-manifold Y. Suppose $\mathcal{H}$ is (weakly) admissible. (Recall that, for our purposes, we always mean *weak* admissibility when we talk about "admissibility" for closed Heegaard diagrams.)

Define the chain complex $\widehat{CF}(\mathcal{H})$ to be $\mathbb{F}_2$-vector space which is freely generated by $\mathfrak{S}(\mathcal{H})$. The differential will be defined by a count of a suitable moduli space of holomorphic curves. For this to make sense, we must begin by picking a generic almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$. Now if we have two generators $\mathbf{x}, \mathbf{y} \in \mathfrak{S}(\mathcal{H})$, then we will consider certain holomorphic curves in $\Sigma \times [0,1] \times \mathbb{R}$ which limit to the g-tuples of chords $\mathbf{x} \times [0,1]$ at $t = -\infty$ and $\mathbf{y} \times [0,1]$ at $t = \infty$. (The exact conditions will be spelled out in Section 3.2.) We may define the moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y})$ to be the space of all such holomorphic curves in the homology class $B \in \widehat{\pi}_2(\mathbf{x}, \mathbf{y})$. There is a (computable) number $\mathrm{ind}(B)$ which is one more than the expected dimension of the moduli space. We will discuss this more in Proposition 3.13.

The upshot, however, is that we may define a map $\partial : \widehat{CF}(\mathcal{H}) \to \widehat{CF}(\mathcal{H})$

$$\partial \mathbf{x} := \sum_{\mathbf{y}} \sum_{\substack{B \in \widehat{\pi}_2(\mathbf{x}, \mathbf{y}) \\ \mathrm{ind}(B) = 1}} \#\left(\mathcal{M}^B(\mathbf{x}, \mathbf{y})\right) \cdot \mathbf{y}.$$

By Proposition 2.16, we have only finitely many positive homology classes $B \in \widehat{\pi}_2(\mathbf{x}, \mathbf{y})$ and which do not cross the basepoint z. There are only finitely many generators, since there are only finitely many α-

and β-curves. Finally, it turns out that $\mathcal{M}^B(\mathbf{x}, \mathbf{y})$ is compact when $\mathrm{ind}(B) = 1$, hence is a finite set of points. Thus $\#(\mathcal{M}^B(\mathbf{x}, \mathbf{y}))$ makes sense, and so this sum is well-defined.

To show that $\partial^2 = 0$ takes a bit of work. It requires that we count the points in index-two moduli spaces. (This corresponds to counting "broken trajectories" in something like Hamiltonian Floer homology, for example.) In particular, the coefficient of $\mathbf{y}$ in $\partial^2\mathbf{x}$ is

$$\sum_{\mathbf{w}} \sum_{\substack{B_1 \in \widehat{\pi}_2(\mathbf{x},\mathbf{w}) \\ \mathrm{ind}(B_1)=1}} \sum_{\substack{B_2 \in \widehat{\pi}_2(\mathbf{w},\mathbf{y}) \\ \mathrm{ind}(B_2)=1}} \#\left(\mathcal{M}^{B_1}(\mathbf{x}, \mathbf{w})\right) \cdot \#\left(\mathcal{M}^{B_2}(\mathbf{w}, \mathbf{y})\right).$$

Showing that this is zero involves counting the ends of the index-two (i.e., one-dimensional) moduli spaces which connect $\mathbf{x}$ and $\mathbf{y}$. Roughly speaking, because this moduli space may be compactified into a one-dimensional compact manifold, it has an even number of ends. (The only compact 1-manifolds are S^1 and $[0, 1]$, both of which have an even number of ends.) Thus the above sum is zero (modulo 2), which shows that $\partial^2 = 0$.

What we have just shown, then, is the following lemma.

Lemma 2.26. *The vector space $(\widehat{CF}(\mathcal{H}), \partial)$ is a chain complex, i.e., $\partial^2 = 0$.*

This is not an invariant of the 3-manifold Y which is represented by the Heegaard diagram $\mathcal{H}$. However, its homology is.

Theorem 2.27. *The homology $\widehat{HF}(\mathcal{H})$ of $H_*(\widehat{CF}(\mathcal{H}), \partial)$ is an invariant of the 3-manifold Y which is represented by the Heegaard diagram $\mathcal{H}$. We denote this invariant, called the* ***Heegaard Floer homology of*** *Y, by $\widehat{HF}(Y)$.*

Showing invariance requires not only invariance of the choice of Heegaard diagram, i.e., invariance under the pointed Heegaard moves of Proposition 2.9, but also invariance of the choice of almost complex structure.

Remark 2.28. Variants of this definition may be obtained by allowing our holomorphic curves to cross z. This leads to versions such as HF^+, HF^-, and HF^∞. See Section 8 of [Lip06a]. Furthermore, we may impose coherent orientations on the moduli spaces so as to get a theory in $\mathbb{Z}$-coefficients, as in Section 6 of the same article. Since the bordered theory deals primarily with the hat-version with coefficients in $\mathbb{F}_2$, however, we do not present the other variants here.

There is also a version which may be applied to knots sitting inside 3-manifolds, known as knot Floer homology. This was first done by Ozsváth and Szabó in [OS04a] and, independently, by Rasmussen in [Ras03]. This is done by equipping the Heegaard diagram with two basepoints. Recall the Morse theory description of Y from $\mathcal{H}$, which has one index-0 critical point and one index-3 critical point. By considering the two flowlines from the index-3 critical point to the two basepoints, and then the two flowlines from the two basepoints to the index-0 critical point, we obtain a knot. (Of course, choosing $2k$ basepoints allows us to obtain a k-component link.) Then a suitable generalization of Heegaard Floer homology recovers the knot invariant $\widehat{HFK}(K)$.

Remark 2.29. As mentioned, the definition of Heegaard Floer homology which we present here follows Lipshitz's presentation in [Lip06a]. The original construction, by Ozsváth and Szabó, involved intersections of the two tori $\mathbb{T}_\alpha := \alpha_1 \times \cdots \times \alpha_g$ and $\mathbb{T}_\beta := \beta_1 \times \cdots \times \beta_g$ in the symmetric product $\mathrm{Sym}^g(\Sigma)$. They counted holomorphic disks with boundary on these tori and which connect intersection points,

i.e., elements of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$. See [OS04c, OS04b]. Since they used holomorphic disks rather than strips, the notion of "connecting generators" is a simpler one; here, it simply means that $u(\pm i)$ are both elements of $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, where u is a holomorphic disk with domain $\{|z| < 1\} \subset \mathbb{C}$. This comes at the expense of having to consider the g-fold symmetric product of Σ, which is a rather more complicated object than $\Sigma \times [0,1] \times \mathbb{R}$.

Chapter 3

Moduli spaces

> The height, in feet, and the number of stories of a building shall be determined based on the type of construction, occupancy classification, and whether there is an automatic sprinkler system installed throughout the building.
>
> 2021 International Building Code, Section 504.1

This chapter contains all the technical results needed to properly define Heegaard Floer homology in both the closed and the bordered case. In particular, the maps are defined by counting certain holomorphic curves with image in $\Sigma \times [0,1] \times \mathbb{R}$. As in many Floer contexts, we want to construct smooth, compact manifolds of dimension 0 and 1 (cf. [Flo88c], [AD14], and [Par16], for example). The former is a finite set of points, and thus may be counted in the definition of a map ∂ in the Floer complex. The endpoints of the latter correspond to terms in ∂^2. Compactness implies an even (i.e., 0 mod 2) number of ends, so that ∂^2 vanishes. Recall Section 2.5.

In general, the argument in that section, and this general Floer argument, is conducted with the following preliminary setup: First, we define a moduli space of certain kinds of holomorphic curves which connect generators of the Floer complex (in this case, these are generators of the Heegaard diagram). We show that this moduli space is generically a manifold by a transversality result. We also compute its expected dimension via an index formula. By allow some kind of "degeneration" or "broken trajectories," we can then compactify this moduli space. Finally, a gluing result implies that this compactification is also a smooth manifold.

We begin this chapter by discussing the moduli space of Riemann surfaces in Section 3.1. In Section 3.2, we then formally define the moduli space whose curves were counted in the definition of $\widehat{HF}$. This moduli space $\mathcal{M}$ comprises holomorphic curves in W with boundary on $C_\alpha \cup C_\beta = (\boldsymbol{\alpha} \times \{1\} \times \mathbb{R}) \cup (\boldsymbol{\beta} \times \{0\} \times \mathbb{R})$ which "connect" generators of the Heegaard diagram. In this section, we also briefly discuss transversality and index results. We spend more time discussing compactification, however, which uses ideas from symplectic field theory. In Section 3.3, we define holomorphic buildings, which roughly correspond to broken trajectories in other Floer theories. Afterwards, we prove compactness of the moduli space of holomorphic buildings in Section 3.4. Finally, by restricting which holomorphic buildings are actually allowed to appear as limits of elements of $\mathcal{M}$ in Section 3.5, we will show that $\mathcal{M}$ is a smooth manifold which is the interior of a compact manifold.

In Sections 3.6 to 3.9, we repeat this construction for the bordered case, beginning by defining the moduli space in Section 3.6. Compactness requires a generalization of holomorphic buildings, known as holomorphic combs, which we introduce in Section 3.7. We show that the moduli space of holomorphic combs is compact in Section 3.8. The compactification does not, unfortunately, produce a honest manifold. However, we can ensure that we will have an even number of ends in the 1-dimensional case, as in Theorem 3.53, which is sufficient for defining the bordered Heegaard Floer invariants. We conclude with Section 3.9.1, in which we provide some examples of degenerations, which will be useful to keep in mind later on in Chapter 4 when we define the bordered invariants $\widehat{CFA}(Y)$ and $\widehat{CFD}(Y)$.

3.1 The Deligne–Mumford moduli space

Our eventual goal is to define and count the points in certain moduli spaces of holomorphic curves. This will allow us to define the differentials of our Floer complexes. But counting points requires that we may compactify these moduli spaces of curves. This involves understanding all possible degenerations which may occur when taking the limit of some sequence of curves in the moduli space. One possible family of degenerations arises by degenerating the domain of these curves. Thus we begin by briefly discussing the compactification of the moduli space of Riemann surfaces.

For our purposes, we allow our Riemann surfaces to have boundary. Furthermore, we allow for punctures and other marked points. Let Z be the set of punctures and M the set of non-puncture marked points. When convenient, we write $\overline{S}$ for $S \cup Z$. Furthermore, when we want to emphasize that we are considering the *punctured* surface S, we sometimes write $\dot{S}$.

We ask that our Riemann surfaces also have the following property which guarantees that they have a finite automorphism group.

Definition 3.1. We say that a Riemann surface is **stable** if

$$2g + \mu + b \geq 3$$

on each component C, where g is the genus of C, μ is the number of points in $(Z \cup M) \cap C$, and b is the number of components in ∂C.

For instance, a sphere with three marked points is stable since Möbius transformations are specified by three points.

Let $\mathcal{M}_{g,\mu,b}$ denote the moduli space of compact connected stable Riemann surfaces with genus g, μ points in $Z \cup M$, and b boundary components.

Roughly speaking, the compactification of this moduli space is given by allowing nodes to form. A nodal Riemann surface is just a smooth Riemann surface with specified double points. More formally, we make the following definition.

Definition 3.2. A **nodal Riemann surface** is a smooth Riemann surface (S, j) equipped with an unordered set D of unordered pairs $\{\{d_1^+, d_1^-\}, \dots, \{d_k^+, d_k^-\}\}$ such that, for each i, the points d_i^+ and d_i^- are either both in the interior or both on the boundary of S. Equivalently, we may think of a nodal Riemann surface as the associated singular surface $S/\{d_i^+ \sim d_i^-\}$. The identified points $d_i^+ \sim d_i^-$ are called **nodes**.

To be completely precise, then, we may specify a nodal Riemann surface as (S, M, Z, D, j). We will almost never do so, however.

The uniformization theorem implies that we may give any stable Riemann surface S a unique complete hyperbolic metric h^S of finite volume which is in the same conformal class as the almost complex structure j on S. We call this metric the Poincaré metric. This detail helps us navigate the later proofs of compactness, where we distinguish between "thick" and "thin" parts of S. We will not always be particularly explicit about it, but when needed we will always assume our Riemann surface comes equipped with the Poincaré metric. For now, we note simply that punctures correspond to cusps (or, in the case of boundary punctures, half-cusps) under h^S.

Note that nodes may be either in the interior or on the boundary. They occur when the complex structures j_n of a sequence $(S_n, j_n) \to (S, j)$ collapses at a geodesic circle or arc with boundary on ∂S_n. The length of a collapsed geodesics goes to zero. Equivalently, the complex structures form infinitely long necks at these geodesics. See, for example, Figures 3.1 and 3.2.

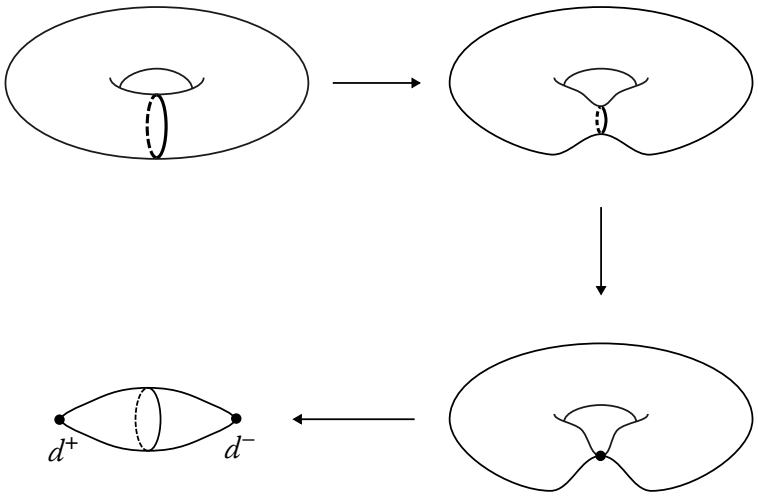

Figure 3.1: Collapsing along a circle. On the top left, we have (S_1, j_1). The top right shows (S_n, j_n) for some n. The geodesic circle has shrunk in length. In the limit (bottom right), the geodesic arc collapses to a point. The bottom left figure gives an interpretation of this nodal surface as a sphere with a double point $d^+ \sim d^-$.

We denote the moduli space of compact connected stable Riemann surfaces with genus g, μ punctures/marked points, and b boundary components by $\overline{\mathcal{M}}_{g,\mu,b}$. This notation is justified by the Deligne–Mumford compactness theorem (Theorem 3.5) below.

Before discussing the compactness theorem below, however, we introduce a few notions. The **compactification** $\widehat{S}$ of a punctured Riemann surface S is obtained by taking the oriented blow-up at the punctures and including the "circle at infinity" (for interior punctures) or the "arc at infinity" (for boundary punctures). This may be seen in Figure 3.3.

Often, we consider the surface S^D obtained by taking the oriented blow-up at the double points $d_i^\pm$ and gluing the boundary circle Γ_i^+ of d_i^+ to the boundary circle Γ_i^- of d_i^- for each i. This surface is called the **deformation** of S. See Figure 3.4.

Remark 3.3. For interior nodes, we may also consider an added piece of data known as a "decoration." This decoration determines the gluing between Γ_i^+ and Γ_i^-. Roughly speaking, this dictates how much we "rotate" one end when gluing it to the other, and is necessary to prove general SFT compactness theorems like Theorem 3.21, which says that a certain family of maps into a cylindrical manifold is compact.

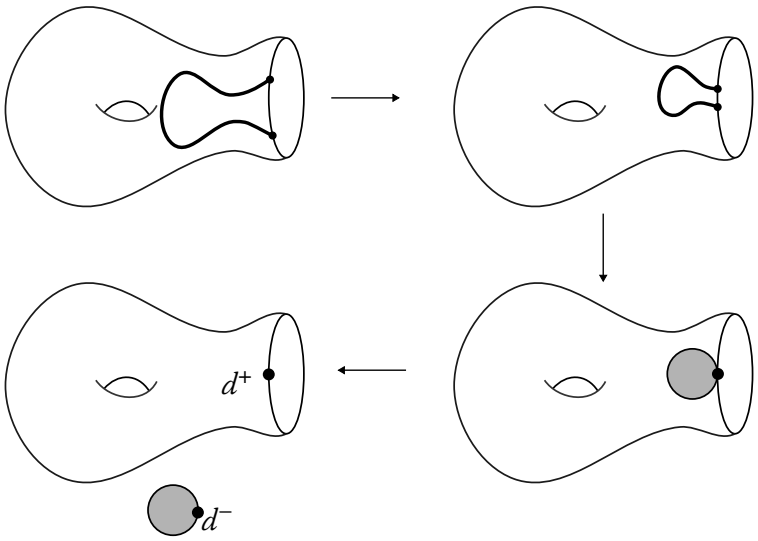

Figure 3.2: Similarly, collapsing along the arc shown above results in a nodal surface in which a disk component is bubbled off.

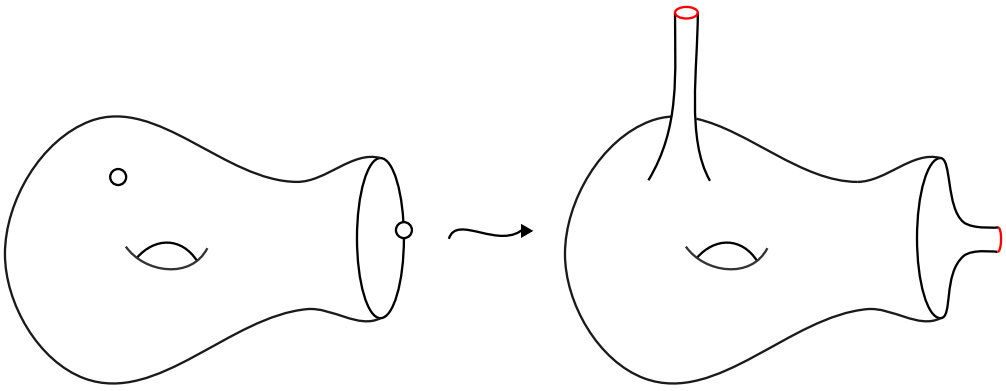

Figure 3.3: The compactification of a twice-punctured Riemann surface. The circle and arc at infinity are colored red.

However, the case which we will be interested in for the remainder of this chapter will not require decorations (see Remark 3.20), as it involves maps into a specific cylindrical manifold. But we may also upgrade everything in this section to moduli spaces of decorated Riemann surfaces. See [BEH+03, Section 3.3] for details.

We now turn to a brief discussion of Deligne–Mumford compactness, first introduced in [DM69], which ensures that any sequence of (smooth or nodal) stable marked Riemann surfaces converge to a nodal Riemann surface S.

Definition 3.4. We say that a sequence $\{(S_n, j_n)\}$ of stable nodal marked Riemann surfaces **converges in the Deligne–Mumford sense** to a limit surface (S, j) if there are diffeomorphisms $\varphi_n : S^D \to S_n^D$ such that

- φ_n takes marked points in S to marked points in S_n;

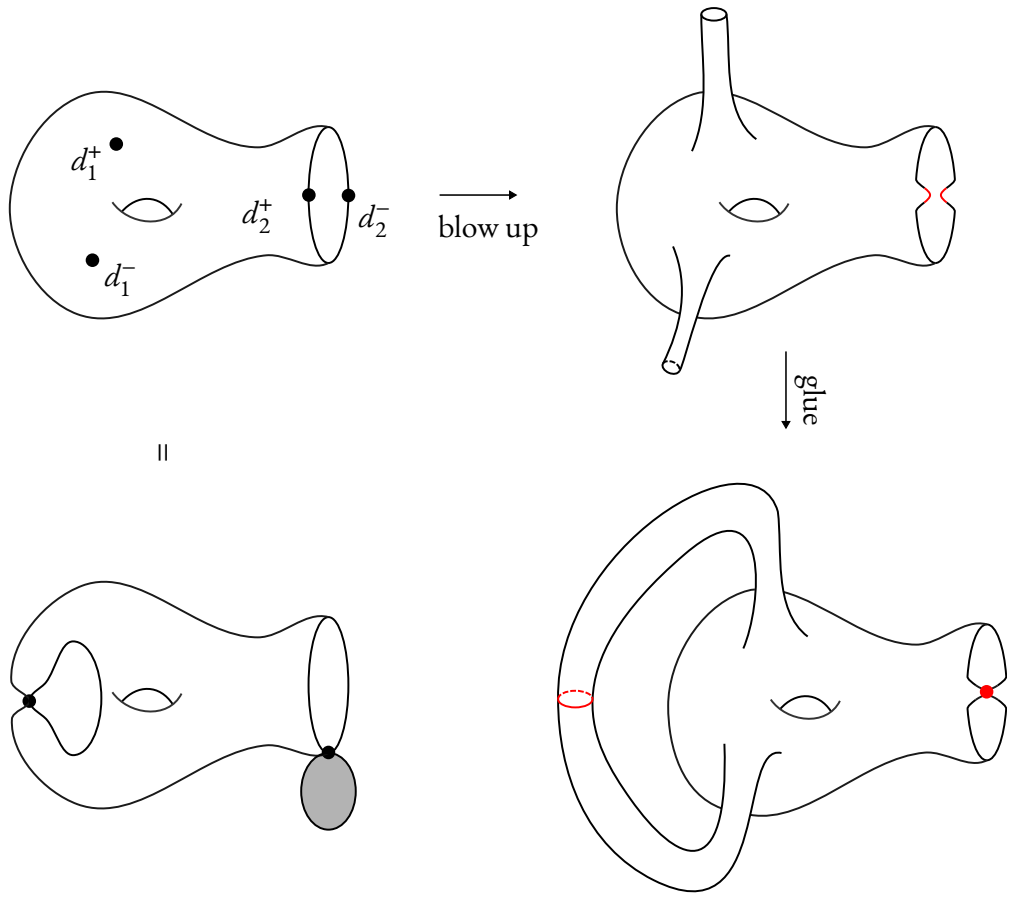

Figure 3.4: The top left shows a Riemann surface with one interior node and one boundary node. (The bottom left diagram shows this same surface with double points identified. This is the "usual" way of visualizing a nodal Riemann surface.) Blowing up gives the top right diagram. The deformation is shown in the bottom right.

- all nodes in S come either from nodes which were already in the S_n's or by degenerating the surfaces S_n along geodesics (either closed geodesics or geodesic arcs with endpoints on ∂S_n);
- all punctures in S come from punctures which were already in the S_n's; and
- the pullback metrics $\varphi_n^* h^{S_n}$ converge to the Poincaré metric h^S on S.

Theorem 3.5 (Deligne–Mumford compactness for surfaces with boundary)**.** *Consider a sequence of stable Riemann surfaces (S_n, j_n) with punctures and marked points in the interior and on the boundary. Suppose the S_n's all have the same topological type and number of punctures/marked points, i.e., belong to the same moduli space $\mathcal{M}_{g,\mu,b}$ for some g, μ, and b. Then we may find subsequence which converges to a stable nodal Riemann surface with marked points.*

This states, more or less, that $\overline{\mathcal{M}}_{g,\mu,b}$ is a compact topological space whose topology is given by Definition 3.4. Thus the only degenerations allowed come from degenerating (or collapsing) at a closed circle or an arc with boundary on the boundary of the Riemann surface. A more formal statement of Deligne–Mumford compactness for Riemann surfaces with boundary may be found in [Wen08, Section 3.3], and a proof may be found in [SS92, Theorem 5.7.1].

For example, if we add at least one marked point to the Riemann surfaces in Figure 3.1, then it would be an example of Deligne–Mumford convergence. Note that the underlying surface, with $g = 1$ and $\mu = b = 0$, is not actually stable. Similarly, if we were to add a marked point to the Riemann surfaces in Figure 3.2, then we would have an instance of convergence in $\overline{\mathcal{M}_{1,1,1}}$. (Indeed, the geodesic arc which is drawn in that figure would only be geodesic with respect to the hyperbolic metric if we already had an interior marked point, or two boundary marked points.)

3.2 The moduli space of holomorphic curves in $\Sigma \times [0,1] \times \mathbb{R}$

We will now define the moduli space which will be used for defining the differential for Heegaard Floer homology. Let $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ be a Heegaard diagram where Σ has genus g. As mentioned before, the differential will be defined by counting certain curves in $\Sigma \times [0,1] \times \mathbb{R}$ with boundary on $C_\alpha \cup C_\beta$ and which converge to generators at $t = \pm\infty$. In this section, we define the relevant moduli space.

We want to only count *holomorphic* curves, so we must put an almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$. To do so, we first make the following definitions.

Let $\pi_{\mathbb{D}} : \Sigma \times [0,1] \times \mathbb{R} \to [0,1] \times \mathbb{R}$ and $\pi_\Sigma : \Sigma \times [0,1] \times \mathbb{R} \to \Sigma$ denote the obvious projections. We let s and t denote the $[0,1]$- and $\mathbb{R}$-coordinates, respectively. Furthermore, fix a point z_i in each component of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$. Let ω_Σ be a symplectic form on Σ, and consider a split symplectic form $\omega = \pi_\Sigma^* \omega_\Sigma + \pi_{\mathbb{D}}^*(ds \wedge dt)$ on $\Sigma \times [0,1] \times \mathbb{R}$. Let j_Σ be an almost complex structure on Σ which is ω_Σ-compatible.

Definition 3.6. An almost complex structure J on $\Sigma \times [0,1] \times \mathbb{R}$ is **admissible** if it satisfies the following requirements:

(J-1) J is tamed by ω.

(J-2) $J = j_\Sigma \times j_{\mathbb{D}}$ is a split almost complex structure in a small cylindrical neighborhood of the fiber $\{z_i\} \times [0,1] \times \mathbb{R}$.

(J-3) The $\mathbb{R}$-action on $\Sigma \times [0,1] \times \mathbb{R}$ defined by translation in the t-coordinate is J-holomorphic.

(J-4) $J(\partial/\partial t) = \partial/\partial s$.

(J-5) J preserves $T(\Sigma \times \{(s,t)\})$ for all $(s,t) \in [0,1] \times \mathbb{R}$.

In the second condition above, we shrink the neighborhoods so that they do not intersect $(\boldsymbol{\alpha} \cup \boldsymbol{\beta}) \times [0,1] \times \mathbb{R}$. Since $J(\partial/\partial t) = \partial/\partial s$, we call the vector field $\partial/\partial s$ the **Reeb vector field**. (Compare this with [BEH+03, Section 2.1]. From now on, we will always assume $\Sigma \times [0,1] \times \mathbb{R}$ is equipped with an admissible almost complex structure, unless otherwise specified.

Since we want our curves to converge to generators at $t = \pm\infty$, we allow their domains to have punctures. Thus we make the following definition.

Definition 3.7. A **source** (S, j) is a Riemann surface with boundary and with finitely many punctures on the boundary such that each puncture is labeled either $+$ or $-$.

We consider two sources to be equivalent if there is an orientation-preserving, label-preserving diffeomorphism between them. Note that S need not be connected.

Fix an admissible J. Then the curves which will factor into our definition of the differential are J-holomorphic maps $u : (S, \partial S) \to (\Sigma \times [0,1] \times \mathbb{R}, C_\alpha \cup C_\beta)$ which satisfy the following:

(M-0) S is smooth, i.e., not nodal.

(M-1) The boundary ∂S is mapped to $C_\alpha \cup C_\beta$.

(M-2) u is an embedding.

(M-3) The energy of u, as defined below, is finite.

(M-4) $\pi_{\mathbb{D}} \circ u$ is nonconstant on every component of S.

(M-5) For every $t \in \mathbb{R}$ and every curve α_i, $u^{-1}(\alpha_i \times \{1\} \times \{t\})$ consists of exactly one point. Similarly, $u^{-1}(\beta_i \times \{0\} \times \{t\})$ consists of exactly one point.

(M-6) For every positive puncture q, we have $\lim_{z \to q}(t \circ u)(z) = \infty$. Here t is the coordinate projection $\Sigma \times [0,1] \times \mathbb{R} \to \mathbb{R}$. Similarly, for every negative puncture q, we have $\lim_{z \to q}(t \circ u)(z) = -\infty$.

We call this last condition **weak boundary monotonicity**. (We will define strong boundary monotonicity later on, in the bordered case.)

Note by [HWZ96, Theorem 2.8] or [Abb04, Proposition 4.5] that a holomorphic map satisfying (M-0)–(M-6) converges to $\mathbf{x} \times [0,1]$ for some generator $\mathbf{x}$ at ∞. After all, each positive puncture should limit to a **characteristic chord**, that is to say, to a trajectory of the Reeb vector field. In this case, since (M-6) implies that $t \to \infty$ near a positive puncture, this means that each positive puncture should limit to some chord $x_i \times [0,1] \times \infty$. (M-5) implies that each x_i should be an element of $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$ and, since there are exactly g positive punctures and g negative punctures, that the set of all x_i's should be a generator of the Heegaard diagram. The same is true at $-\infty$.

This means that $\pi_{\mathbb{D}} \circ u$ is a g-fold branched covering map.

Now we introduce the definition of energy which is used in (M-2), and in the remainder of this work. We are working with holomorphic curves in $\Sigma \times [0,1] \times \mathbb{R}$, which is an instance of a cylindrical symplectic manifold $V \times \mathbb{R}$. (In our case, $V = \Sigma \times [0,1]$.) The pullback of the symplectic form ω_Σ on Σ to V has rank two. Its kernel is generated by $\partial/\partial s$, which is called the Reeb vector field. The condition that $J(\partial/\partial t) = \partial/\partial s$ then says that the almost complex structure maps the vector field which generates the $\mathbb{R}$-translations (i.e., the translations in the cylindrical direction) to the Reeb vector field. This is a standard technical requirement for studying pseudoholomorphic curves on cylindrical manifolds. Details may be found in [BEH+03, Section 2].

Definition 3.8. The **energy** of $u : S \to V \times \mathbb{R} = \Sigma \times [0,1] \times \mathbb{R}$ is given by the formula

$$E(u) := \int_S (\pi_\Sigma \circ u)^* \omega + \sup_\phi \int_S (\phi \circ t \circ u)\, dt \wedge (\pi_V \circ u)^* \lambda,$$

where ω is the symplectic form on Σ and π_V is the projection to $V = \Sigma \times [0,1]$. The supremum is taken over all functions $\phi : \mathbb{R} \to \mathbb{R}_{\geq 0}$ with compact support and integral 1. The first term in the energy is called the ω-**energy**; the second is called the λ-**energy**.

In particular, we "forget" the energy in the $\mathbb{R}$-direction. After all, since u projects to a covering of $[0,1] \times \mathbb{R}$, hence of $\mathbb{R}$, the energy in the $\mathbb{R}$-direction is infinite.

Remark 3.9. We will mostly be concerned with energy only insofar as to whether or not it is bounded. In our case, the λ-energy will always be bounded. After all, roughly speaking, the λ-energy is obtained by computing the maximum "width" of u (in the s-direction) over some interval of $\mathbb{R}$. In particular, since u is a g-fold branched cover of $[0,1] \times \mathbb{R}$, it follows that this width is bounded by g. That is to say, the λ-energy is always bounded by the genus g of the Heegaard diagram.

Finally, we make the following definition.

Definition 3.10. The holomorphic curve u **stable** if

- not all components of the curve are twice-punctured disks which project down to a single point in Σ; and
- every connected component on which u is constant is stable.

Define $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S)$ to be the space of holomorphic curves from a source S which satisfy (M-0)–(M-6), connect the generators $\mathbf{x}$ and $\mathbf{y}$, and belong to the homology class $B \in \pi_2(\mathbf{x}, \mathbf{y})$, all quotiented out by the automorphisms of S. The $\mathbb{R}$-action on $\Sigma \times [0,1] \times \mathbb{R}$, namely the translation action, induces an $\mathbb{R}$-action on this moduli space. This action is free if u is stable.

Thus we may define $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ to be the quotient $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S)/\mathbb{R}$.

Lemma 3.11. *If $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S)$ is nonempty, then B is a positive homology class. That is to say, all of the coefficients in its corresponding domain are nonnegative.*

Proof. Suppose u is an element in $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S)$. Its multiplicity at a region R is exactly equal to the intersection number $u \cdot (\{p\} \times [0,1] \times \mathbb{R})$ for some point $p \in R$. Since $J(\partial/\partial t) = \partial/\partial s$ by (J-4), we know that the fiber $\{p\} \times [0,1] \times \mathbb{R}$ is J-holomorphic. Thus this is the intersection number between two J-holomorphic curves, which is positive by [MW95, Theorem 7.1]. □

Proposition 3.12. *There is a residual set J_{reg} of almost complex structures for which the moduli spaces $\mathcal{M}(\mathbf{x}, \mathbf{y}; S)$ are transversally cut out by the $\overline{\partial}$-equations, hence are smooth manifolds. In particular, by the Baire category theorem, this is actually a dense set of almost complex structures.*

In this context, a residual set is one which contains a countable intersection of open dense sets. The proof is similar to the proof of [MS12, Theorem 3.1.6].

If J is an element of J_{reg}, then we say that it **achieves transversality**. The above proposition states that generic J achieve transversality, so we may always assume this property.

Given that $\mathcal{M}(\mathbf{x}, \mathbf{y}; S)$ are smooth manifolds, we may ask what their expected dimension is. To define this, we define the **Euler measure** of a region R in $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$ to be

$$e(R) := \chi(R) - \frac{1}{4} \cdot \#(\text{acute corners in } R) + \frac{1}{4} \cdot \#(\text{obtuse corners in } R),$$

where $\chi(R)$ is the Euler characteristic and where a "corner" is an intersection of an α- and β-curve. Note that we implicitly choose a Riemannian metric here so that our α- and β-curves meet at right angles. The difference between an acute and obtuse corner is illustrated in Figure 3.5. (The usual definition of the Euler measure is $1/2\pi$ times the integral over R of the curvature. By the Gauss–Bonnet theorem, the definition above is equivalent to this usual, less combinatorial definition.) We extend by linearity to define the Euler measure of any linear combination of regions, i.e., of any domain.

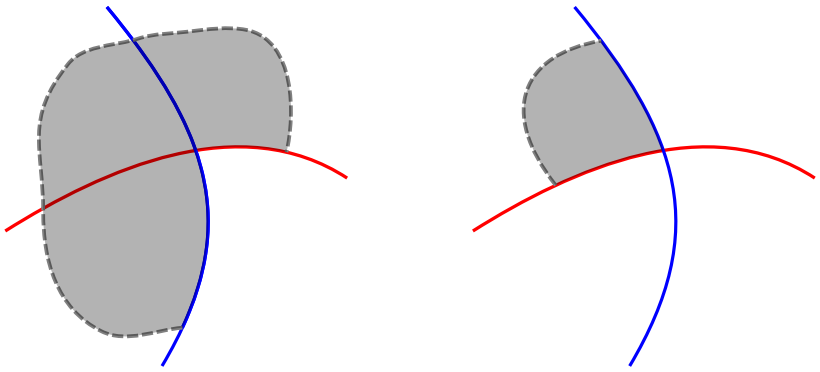

Figure 3.5: The source S appears projects onto some region R in Σ, shown here as the shaded area. The left side shows an obtuse corner, while the right side shows an acute corner. Counting such corners gives the Euler measure.

Proposition 3.13 ([Lip06a, Corollary 4.3])**.** *The expected dimension* $\operatorname{ind}(B,S)$ *of* $\widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y};S)$ *is*

$$\operatorname{ind}(B,S) = g - \chi(S) + 2e(D(B)),$$

where $D(B)$ *is the domain associated to the homology class* B*. Recall we assume that maps in* $\widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y};S)$ *are embedded. If* $\widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y};S)$ *is nonempty, then*

$$\chi(S) = g - \left(\sum_{i=1}^{g} n_{x_i}(B) + n_{y_i}(B)\right) + e(D(B)),$$

where x_i *and* y_i *are the particular points in the generators* $\mathbf{x}$ *and* $\mathbf{y}$*, respectively, and where* $n_p(B)$ *denotes the local multiplicity of* B *at the point* $p \in \Sigma$*. Thus the expected dimension depends only on* B*, and we write*

$$\operatorname{ind}(B) := \dim \widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y};S) = \left(\sum_{i=1}^{g} n_{x_i}(B) + n_{y_i}(B)\right) + e(D(B)).$$

Now we may define a moduli space which is source-independent. In particular, define

$$\chi_{\text{emb}}(B) := g - \left(\sum_{i=1}^{g} n_{x_i}(B) + n_{y_i}(B)\right) + e(D(B)),$$

so that $\chi(S) = \chi_{\text{emb}}(B)$ whenever $\widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y};S)$ is nonempty. Then define the source-independent moduli space

$$\widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y}) := \bigcup_{\chi(S)=\chi_{\text{emb}}(B)} \widetilde{\mathcal{M}}^B(\mathbf{x},\mathbf{y};S).$$

We define $\mathcal{M}(\mathbf{x},\mathbf{y})$ by quotienting out by the translation action in the $\mathbb{R}$-coordinate.

Later on, in the definition of Heegaard Floer homology in Section 2.5, we will define the differential by counting curves in $\mathcal{M}^B(\mathbf{x},\mathbf{y})$ where $B \in \widehat{\pi}_2(\mathbf{x},\mathbf{y})$ ranges over all positive homology classes with $\operatorname{ind}(B) = 1$. We can do this because Proposition 2.16 guarantees that there are only finitely many B to consider.

3.3 Holomorphic buildings

To compactify this moduli space, we must allow two kinds of degeneration: First, we must allow our sources to degenerate into nodal sources. For example, the limit of a sequence of maps $u_n : (S, j_n) \to \Sigma \times [0,1] \times \mathbb{R}$ could have nodal domain if the Riemann surface structure j_n "pinches down" at some point in S. These kinds of degenerations, known as Deligne–Mumford degenerations, were explained in Section 3.1. Second, we must allow degeneration at $t = \pm\infty$. This degeneration occurs when the maps themselves "escape to infinity" near the punctures or nodes, and is indicated in Figure 3.6. In particular,

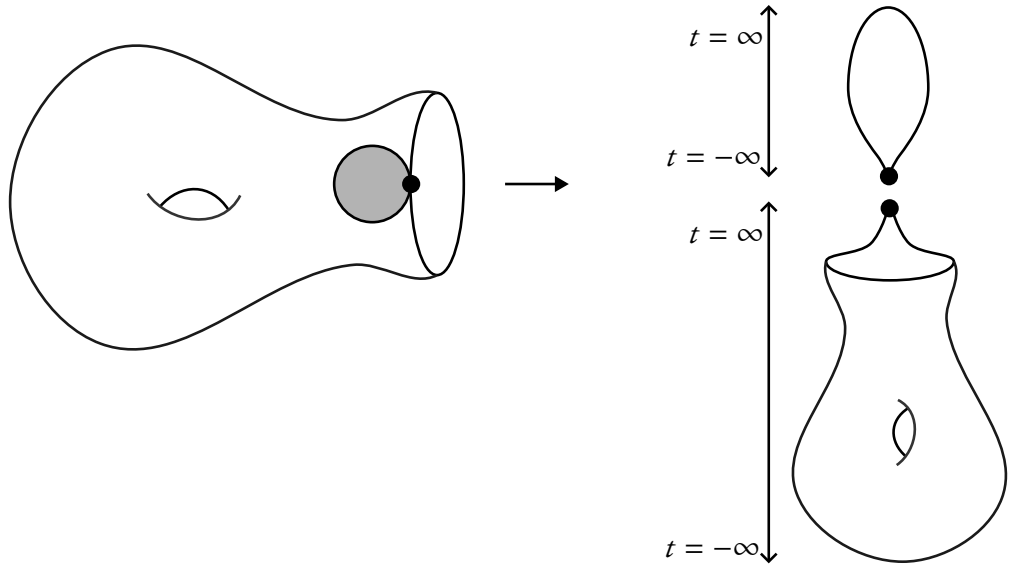

Figure 3.6: Consider the sequence of maps in Figure 3.2. This forms a boundary node, as shown on the left side above. If the t-coordinate approaches ∞ from one side of the node (namely the "genus-1 side," i.e., the side which does not belong to the degenerated disk), then we get the diagram on the right, where the degenerated disk escapes to infinity.

any two points of a single holomorphic curve are necessarily a finite distance apart. Furthermore, the energy is concentrated on a bounded portion of this curve, and thus vanishes as $t \to \pm\infty$. Holomorphic buildings let us consider curves where some parts go to infinite *relative* to other parts and where the energy may accumulate at $t = \pm\infty$.

This second kind of degeneration uses ideas from symplectic field theory, namely the definition of a holomorphic building as introduced in [EGH00] and [BEH+03].

We will only need to use holomorphic buildings in $\Sigma \times [0,1] \times \mathbb{R}$, but in general one may define holomorphic buildings in cylindrical almost complex manifolds $V \times \mathbb{R}$, as well as in manifolds with cylindrical ends. In particular, the domains of our holomorphic maps in this section are once again sources as in Definition 3.7, as opposed to the more general setting of stable nodal Riemann surfaces with punctures and marked points that we considered when discussing Deligne–Mumford compactification. (The main difference is that we do not have marked points or interior punctures.) See Sections 7 and 8 in [BEH+03], or Chapter 3 in [Abb14], for more details about the general case.

Definition 3.14. Consider a sequence of stable maps $u_k \in \mathcal{M}^{B_k}(\mathbf{x}_{k-1}, \mathbf{x}_k; S_k)$ for $k = 1, \dots, N$. Let $\widehat{S_k}$ denote the compactification of S_k and $\widehat{\pi_\Sigma \circ u_k}$ the compactification of $\pi_\Sigma \circ u_k$. We may form the piecewise

smooth surface

$$\widehat{S} := \widehat{S_1} \cup \widehat{S_2} \cup \cdots \cup \widehat{S_k}.$$

If the maps $\widehat{\pi_\Sigma \circ u_k}$ glue to a continuous map on $\widehat{S}$, then we call the sequence $\{u_k\}$ a **holomorphic building of height** N, which we denote U.

In this context, we call the u_k's the **stories** of the building, and call k the **level**. Furthermore, we say that $\widehat{S}$ is the **preglued source** of U. If u_k has homology class B_k, then U has homology class $B_1 * \cdots * B_N \in \pi_2(\mathbf{x}_0, \ldots, \mathbf{x}_N)$.

Remark 3.15. In this context, we allow our sources S_k to be nodal Riemann surfaces. In general, when we want to explicitly include or exclude nodal sources as a possibility, we will say so. However, when the distinction is unimportant, we will sometimes just talk about "holomorphic curves" and "sources" without mentioning whether they are smooth or nodal.

Intuitively, a holomorphic building is a sequence of holomorphic maps such that the asymptotics of u_k at the positive punctures agree with the asymptotics of u_{k+1} at the negative punctures. (Strictly speaking, a holomorphic building also comes with an ordering of all the punctures, but we omit this detail here.) See, for example, Figure 3.7.

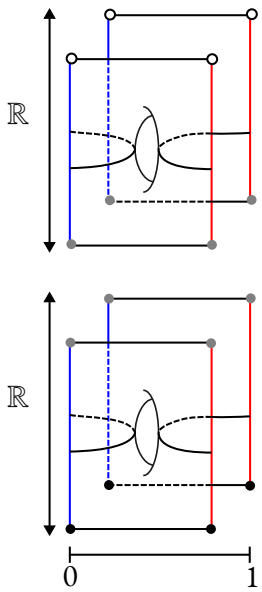

Figure 3.7: The components in blue lie on the cylinder $C_\beta = \beta \times \{0\} \times \mathbb{R}$, while the components in red lie on the cylinder C_α. Let $\mathbf{x}$ be the generator corresponding to the black dots at the bottom. (Note that black dots on the same horizontal line correspond to the same point in Σ; points in $\mathbf{x}$ are elements of $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$.) Let $\mathbf{w}$ be the generator corresponding to the gray dots, and $\mathbf{y}$ the generator corresponding to the white dots. Then this shows a two-story holomorphic building from $\mathbf{x}$ to $\mathbf{y}$ in homology class $B_1 * B_2$, where the first story shows a curve in $\mathcal{M}^{B_1}(\mathbf{x}, \mathbf{w})$ and the second story shows a curve in $\mathcal{M}^{B_2}(\mathbf{w}, \mathbf{y})$.

Another way to think of a holomorphic building is as a nodal curve with some of the nodes removed. (These deleted nodes correspond to the positive/negative ends of consecutive stories which glue together.) See Figure 3.8.

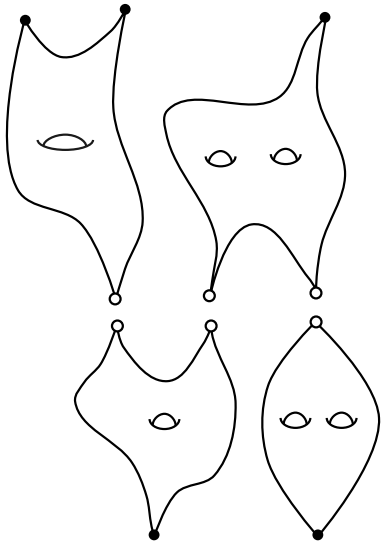

Figure 3.8: This diagram does not show the $[0,1]$-factor. The bottom layer consists has three points x_1, x_2, x_3 in black, while the top layer has three such points y_1, y_2, y_3. Then this diagram goes from $\mathbf{x}$ to $\mathbf{y}$, and we may think of it as a curve whose domain is the nodal curve whose double points are exactly those points labeled with white (unshaded) dots the diagram above. Thus we may think of a holomorphic building as the nodal curve below, with punctures corresponding to the black dots and removed-nodes corresponding to the white dots.

Consider two holomorphic buildings of height 1, i.e., two stable holomorphic curves u and u' with source S and S', respectively. They are called **equivalent** if there is a diffeomorphism between S and S' which preserves the complex structure and such that $u' \circ \phi$ and u are the same up to a translation in the $\mathbb{R}$-direction. In general, if U and U' are height-N holomorphic buildings, then we call them **equivalent** if their respective stories are equivalent and if these equivalences commute with the attaching maps between stories. (These attaching maps were implicit in our definition of a holomorphic building. They are used to glue the $\widehat{S}_k$'s to form the preglued surface $\widehat{S}$.)

We now define moduli spaces of holomorphic buildings, as well as a notion of convergence within these spaces.

Definition 3.16. The moduli space of all (possibly nodal) holomorphic buildings in the homology class B and with asymptotics $\mathbf{x}$ at $-\infty$ and $\mathbf{y}$ at ∞, and whose preglued sources have the same topological type as S, is denoted by $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S)$.

Definition 3.17. Consider a sequence $\{U_n\}$ of holomorphic buildings of height at most N. We say they converge to a building U of height N if the following properties hold.

- The deformations S_i^D of the underlying surfaces S_n of U_n converge to the deformation S^D of the underlying surface S of U in the Deligne–Mumford sense.
- If φ_i is the diffeomorphism from S^D to S_n^D coming from the Deligne–Mumford convergence, then the projection of $u_m \circ \varphi_n$ converges uniformly to the projection of u on $\Sigma \times [0,1]$.
- Let Γ_i be the set of boundary circles of the n-th level of the building U. Let C_ℓ be the union of

the components of $S^{D,r} \setminus \bigcup \Gamma_n$ which are on the ℓ-th level, where $\ell = 1, \dots, N$. Then there exist sequences of real numbers c_n^ℓ for $n \geq 1$ such that $(t \circ u_n \circ \varphi_n - t \circ u - c_n^\ell) \to 0$ in C^0_{loc}.

This last condition more or less says that we can think of the ℓ-th level of U as the limit of $\mathbb{R}$-translates of levels of the U_n's. In the case when the U_n's are all honest curves, and not just buildings, note that $c_n^{\ell+1} - c_n^\ell \to \infty$ as $n \to \infty$. Intuitively, this says that consecutive levels in a holomorphic building are infinitely far apart in the $\mathbb{R}$-direction, since the $\mathbb{R}$-translates which converge to the different levels of the limit building differ by a larger and larger amount.

As the notation suggests, the moduli space $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S)$ is compact; thus holomorphic buildings give a compactification of our moduli space of holomorphic curves.

3.4 Compactification via holomorphic buildings

We now have the following statement of compactness. This is the relative version of [BEH+03, Theorem 10.1], i.e., with sources with boundary. Alternatively, see [Abb14, Theorem 3.20] for a formal proof.

Theorem 3.18 (SFT compactness)**.** *For every E_0, the space of holomorphic buildings in $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S)$ with energy bounded above by E_0 is compact. Since the moduli space has a metric [BEH+03, Appendix B.2], we can restate this as follows: Any sequence $\{U_n\}$ of holomorphic buildings in $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S) \cap \{U : E(U) \leq E_0\}$ has a convergent subsequence, and the limit has homology class B.*

Since we may handle convergence of each level separately, it is enough to prove the above theorem in the case of height-1 holomorphic buildings, that is to say, in the case that each $U_n = u_n$ is a stable holomorphic curve.

Consider sources (S_n, j_n) of u_n with punctures and nodes. (All of the S_n's have the same topological type, namely that of S, but we may vary the Riemann surface structure.) Recall that we think of nodes as identified double points. We may then delete each node of S_n, so that S_n is assumed to be smooth, and carry a set M_n of marked points consisting of the pre-existing punctures and nodes of S_n. This is fine because nodes (and hence all the points in M_n) are treated like punctures, at least from the perspective of hyperbolic metrics and the Deligne–Mumford moduli space.

The upshot of this is that we may prove Theorem 3.18 for stable holomorphic curves whose domains are smooth, punctured (i.e., marked) Riemann surfaces. Note that, even in this case that u_n has smooth source, the limit $u_n \to U$ might have nodal source. Any nodes which form in the limit are still called "nodes."

Thus it is enough to prove the following statement.

Theorem 3.19. *Consider a sequence $\{u_n\}$ of holomorphic curves in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ with energy bounded above by E_0. (As per our remark above, the source S may have some marked points.) Then there is a subsequence which converges to a stable holomorphic building U of finite height N and in the homology class B.*

Proof. We begin with an intuitive idea of the proof. Deligne–Mumford compactness says that the (possibly nodal) limit surface $S_\infty = \lim S_n$ exists. It is possible to obtain a gradient bound which ensures that in the "thick" part of S_∞ Roughly speaking, the "thick" part of S_∞ consists of the points x where there is a positive lower bound on the injectivity radii at x of the metrics h_n. This basically means that we never "pinch" the metric (or the complex structure) at points in the thick part of S_∞.

By contrast, the "thin" part of S_∞ consists exactly of (neighborhoods of) nodes and punctures. A node is adjacent to two components of the thin part. At each component, the limit map U approaches either a point in $\Sigma \times [0,1] \times \mathbb{R}$ or a holomorphic strip modeled over a Reeb chord $x_i \times [0,1]$ at the node. It then suffices to show that U takes the same value on both sides of the node, that is to say, it limits to the same value as it approaches the node from either component. The only possible issue is that there may be some energy which is lost between the two sides. This comes up if there were actually a bubbled-off cylinder or strip which forms at the node. By adding enough marked points to the original surface, we may make sure that this component is seen in the Deligne–Mumford limit S_∞. The same holds true for convergence over punctures. Finally, the level structure of the limit is based on how degenerated components of the Deligne–Mumford limit, which have positive/negative ends at the nodes in the limit surface, attach to other components.

Given this brief overview, we now turn to the actual proof.

Step 1. *The limit surface.* We think of our holomorphic curves u_n as having domain S_n with marked point set M_n. Each S_n has the same topological type S, but the complex structure j_n may vary. Add marked points (as needed) to stabilize the surfaces S_n.

Fix a point p_r in each region r of Σ. We may pick p_r generically, so that they are regular values of $\pi_\Sigma \circ u_n$ for all n. Let $\{q_{i,r,n}\} = (\pi_\Sigma \circ u_n)^{-1}(p_r) \subset S_n$ be the preimages, and add them to the marked point set M_n of S_n. These marked points implicitly keep track of the homology class B, since the number of points $q_{r,i,n}$ for each region r tells us how many times u_n crosses the region.

Deligne–Mumford compactness implies that the surfaces S_n converge to a nodal Riemann surface S_∞. This convergence comes with certain maps $\varphi_n : S_\infty \to S_n$. (Technically, these maps are between the deformations of S_∞ and S_n, but it is more useful for now to think of S_∞ and S_n as punctured surfaces, rather than as their deformations.) We also have a Poincaré metric h_n on each S_n. Recall that the pullback metrics $\varphi_n^* h_n$ converge to the Poincaré metric h_∞ on S_∞. (Note that hyperbolic metrics depend not only on the complex structure j_n, but also the marked point set.)

Step 2. *Thick-thin decomposition.* Let h_n be the Poincaré metrics associated to j_n and our marked point sets M_n. We have a Poincaré metric h_∞ on S_∞ which is the limit of the h_n's. With $\rho(x)$ denoting the injectivity radius at x with respect to this limit metric, we define

$$\mathrm{Thick}_\varepsilon(S_\infty) := \{x \in S_\infty : \rho(x) \geq \varepsilon\} \qquad \mathrm{Thin}_\varepsilon(S_\infty) := \{x \in S_\infty : \rho(x) < \varepsilon\}$$

to be the ε-**thick** and ε-**thin** parts, respectively. It turns out that, for $\varepsilon < \sinh^{-1} 1$, the ε-thin part consists entirely of finite cylinders and punctured disks, as well as finite strips and punctured half-disks. See Figure 3.9 below. Each finite cylinder C has a unique closed geodesic of length $2 \inf_{x \in C} \rho(x)$. We call

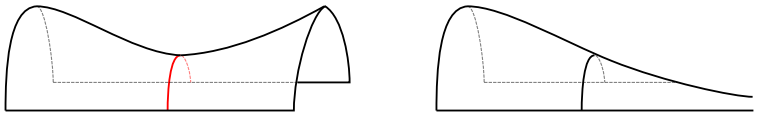

Figure 3.9: These are components of the ε-thin part of S_∞. On the left is a finite strip with short geodesic arc Γ, indicated by the red arc. On the right is a punctured half-disk.

this geodesic a **short geodesic**. Similarly, the finite strips have short geodesic arcs with endpoints on the boundary. There are finitely many of these geodesics, and they are all pairwise disjoint. See [Hum97, Chapter IV.4] for details.

When not specified, we will always assume that $\varepsilon < \sinh^{-1} 1$ when we are discussing ε-thick and ε-thin parts. The upshot is that this implies that every connected component of $\mathrm{Thin}_\varepsilon(S)$ is either a small collar neighborhood of a short geodesic (as in Figure 3.9(a)) or a small cusp neighborhood of a puncture (as in Figure 3.9(b)). Note that we only have punctures on the boundary, since S is a source (cf. Definition 3.7), so our "cusp neighborhoods" are actually half-cusp neighborhoods which, topologically, look like $[0, \infty) \times [0, 1]$ instead of $[0, \infty) \times S^1$.

Step 3. *Convergence over the thick part.* Let Γ_j's denote the short geodesics (corresponding to the nodes and interior punctures) or embedded arcs with endpoints at ∂S (corresponding to the boundary punctures). Strictly speaking, the curves Γ_j, which we call **special curves**, should correspond to the *blown-up* nodes and punctures in the deformation. But these special curves have length 0 in the deformation too, so it doesn't hurt to just think of them as points.

We have uniform bounds on $\|\nabla(u_n \circ \varphi_n)(x)\|$ for all $x \in S_\infty \setminus \bigcup \Gamma_j$. This bound implies convergence over the thick part. In particular, we certainly have a subsequence which converges with all derivatives over the ε-thick part for some fixed $\varepsilon > 0$. (This uses a result of [Gro85]. A proof may also be found in [MS12, Theorem 4.1.1].) After translating in the $\mathbb{R}$-direction, we may apply Arzelà–Ascoli to extract a subsequence which converges in the C^∞_{loc}-topology away from the punctures and the special geodesics, i.e., on $S_\infty \setminus \bigcup \Gamma_j$. Say the limit is some holomorphic map u.

Step 4a. *Convergence over nodes.* Now we must prove convergence on the thin part of S. For small enough ε, this thin part consists of small neighborhoods of two kinds of points. The first kind are nodes in S, which were created by taking the limit $i \to \infty$ by shrinking some of the short geodesics (i.e., the closed geodesic curves in a finite cylinder or the geosdesic arcs in a finite strip) on the sources S_i. The second kind are (boundary) punctures in S, which were already there before taking the limit. In this step, we tackle the first case.

A node q is adjacent to two components C^+ and C^- of $\mathrm{Thick}(S_\infty)$, as seen in Figure 3.10. If u is

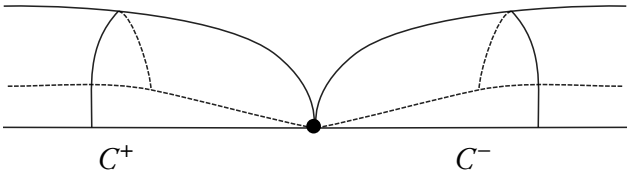

Figure 3.10: The ε-thin part of S_∞ consists of a small neighborhood of the node indicated above. (Topologically, this is just two disks glued together at a point on the boundary.) Outside of this neighborhood is the ε-thick part of S_∞. As $\varepsilon \to 0$, the thick portion approaches the node in two different components, labeled C^+ and C^- above.

bounded near q, thought of as a point of C^+, then the removable singularities theorem (see [MS12, Theorem 4.1.2]) implies that u extends continuously over *q on the C^+ side*. Otherwise, it approaches a holomorphic strip $x_i \times [0, 1] \times [R, \infty)$ or $x_i \times [0, 1] \times (-\infty, R]$ modeled over some characteristic chord [HWZ96, Theorem 2.8]. (In general, it may approach any Reeb orbit; here, since the Reeb vector field is $\partial/\partial s$, this is the only possible limit. In particular, since there are no periodic orbits, interior nodes must converge to a single point.) The same is of course true if we think of the node as a point of C^-.

Let $\gamma^\pm$ be the asymptotic limits over $C^\pm$ of u at the node. We would like to show that $\gamma^+ = \gamma^-$.

This node appeared by degenerating a component of $\mathrm{Thin}_\varepsilon(S_n)$ along a circle (if the node is in the interior) or an arc with boundary on ∂S_n (if the node is on the boundary). In particular, there is a conformal parameterization near q by some interval times either S^1, if q is in the interior, or $[0, 1]$, if q is on the

boundary. See Figure 3.11. More precisely, let T_n^ε be the component of the ε-thin part of S_∞ (equipped

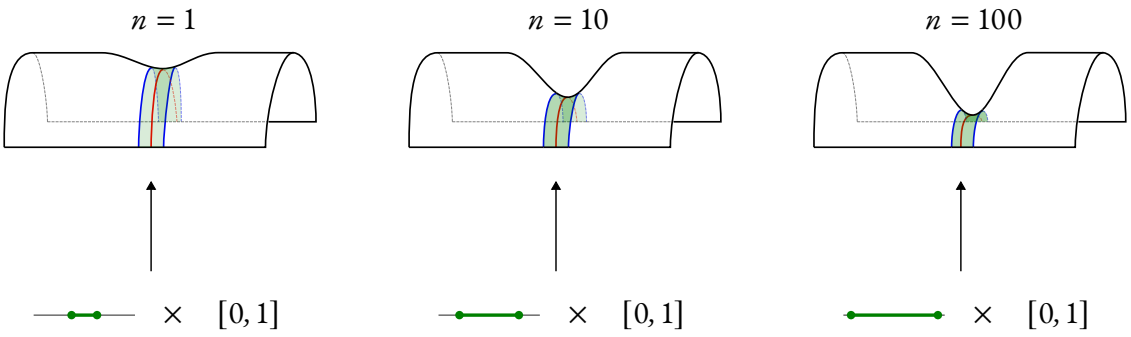

Figure 3.11: This shows how the parameterization $g_n^\varepsilon : [-N_n^\varepsilon, N_n^\varepsilon] \times [0,1] \to T_n^\varepsilon$ changes as n increases. Note that "collapsing" along the red arc is the same as stretching into an arbitrarily long strip. After all, the $[0,1]$-factor corresponds to the direction perpendicular to the red arc. The component $[-N_n^\varepsilon, N_n^\varepsilon]$, which is depicted by the green intervals above, becomes long in comparison to this $[0,1]$-factor, since the red arc becomes short. (In the interior node case, which is less important for our purposes, we stretch into an arbitrarily long cylinder.)

with the metric $\varphi_n^* h_n$, not with the metric h_∞) which contains q. This component is a collar neighborhood of the degenerating (i.e., collapsing) circle or arc. There is a conformal map

$$g_n^\varepsilon : A_n^\varepsilon := [-N_n^\varepsilon, N_n^\varepsilon] \times [0,1] \to T_n^\varepsilon.$$

(Again, we replace S^1 with S^1 in the case that q is in the interior; we will stop repeating this, since this is the less relevant case for us, but it continues to hold true for the rest of this step.) Note that applying u_n, and projecting to $\Sigma \times [0,1]$, limits to our Reeb chords:

$$\lim_{\varepsilon \to 0} \lim_{n \to \infty} (\pi_\Sigma \times s) \circ u_n \circ \varphi_n \circ g_n^\varepsilon|_{\pm N_n^\varepsilon \times [0,1]} = \gamma^\pm.$$

Loosely speaking, after all, we approach q from $C^\pm$ when we evaluate this map at $\pm N_n^\varepsilon \times [0,1]$.

Denote the map $(\pi_\Sigma \times s) \circ u_n \circ \varphi_n \circ g_n^\varepsilon$ by v_n^ε. We may ask that v_n^ε has a uniform gradient bound over A_n^ε. As such, we may choose some $\varepsilon_n \to 0$ and a subsequence (which we still denote with subscript n) such that

$$\lim_{n \to \infty} v_n^{\varepsilon_n} \left(\pm N_n^{\varepsilon_n} \times [0,1] \right) = \gamma^\pm.$$

For large n, the maps $v_n^{\varepsilon_n}$ thus define a homotopy between γ^+ and γ^-. We can assume that the homology class of this homotopy is independent of n, so call this homotopy Φ.

We have two cases now: Either we have lost energy between γ^+ and γ^- (corresponding to energy in the cylinders A_n^ε), or we have not. In fact, it is enough to consider whether we have lost ω-energy, i.e., whether $\int_{[0,1]\times[0,1]} \Phi^* \omega$ is zero or not.

If the ω-energy is also not lost, then $\gamma^+ = \gamma^-$ and our holomorphic strips $[-N_n^\varepsilon, N_n^\varepsilon] \times [0,1]$ have very small ω-energy. In the case that γ^+ is a trajectory, we may use [HWZ02]. If γ^+ is instead a point, the result follows from [BEH+03, Lemma 5.14]. Roughly speaking, the first result tells us that cylinders—and hence, by a doubling argument, strips—with small energy look like trivial cylinders (respectively, strips) over a Reeb trajectory, while the second one says that the diameter of the v_n^ε's approaches 0.

If, on the other hand, we have lost energy, then since ∇v_n^ε was assumed to be bounded (thanks to the gradient bound on g_n^ε), the only possible place for bubbled-off energy is at the node itself. This creates strip-breaking at boundary punctures, or cylinder-breaking at interior punctures. These strips have positive energy corresponding to the lost energy, but no marked points since they did not appear in the Deligne–Mumford limit. But positive energy implies that the projection to Σ is nonconstant, so such a strip must have passed through one of the points p_r chosen in Step 1. This means that there must be a marked point on a bubbled-off strip, a contradiction, so this case never occurs.

Step 4b. *Convergence over punctures.* The idea for punctures is very similar to that of nodes. The only difference is the following: In the case where q was a node, it was adjacent to two components, and we showed that the components approach the same value at q. In the case where q is a puncture, it is only adjacent to a single component. Using the notation $v_n^\varepsilon = u_n \circ \varphi_n \circ g_n^\varepsilon$, but here with g_n^ε the parameterization of a cuspidal neighborhood of q by $[0,\infty)\times[0,1]$, we see that the puncture approaches one limit

$$\gamma_1 \coloneqq \lim_{n\to\infty}\lim_{x\to\infty} v_n^\varepsilon(\{x\}\times[0,1]).$$

(Note that we only consider the case when q is a boundary puncture, since our definition of a source only allows for boundary punctures. Unlike nodes, which can form in the limit by degenerating along curves, the limit surface S_∞ can only have the punctures which were already found in S, i.e., in the topological type of the S_n's.) There is another way to get the limit of the puncture, however, namely by translating v_n^ε in the $\mathbb{R}$-direction, and then taking the limit γ_2. That is to say, the maps $v_n^\varepsilon - (t\circ v_n^\varepsilon)(0,0)$ are also asymptotic to some Reeb trajectory. To show that these are the same, we again split into cases depending on whether any energy is lost.

Step 5. *Obtaining the level structure.* Since we will almost never need the exact details of the level structure, we will be brief in this step. Label the components of $\mathrm{Thick}(S_\infty) = S_\infty \setminus \bigcup \Gamma_j$ as $C_1,\dots,C_N$. We say that $C_i \le C_j$ if, for points $x_i \in C_i$ and $x_j \in C_j$, we have

$$\limsup_{n\to\infty}[(t\circ u_n)(x_i) - (t\circ u_n)(x_j)] < \infty.$$

If $C_i \le C_j$ and $C_j \le C_i$, then we say that $C_i \sim C_j$.

This produces an ordering on the components of $\mathrm{Thick}(S_\infty)$. We say that the first level consists of those components which are minimal with respect to this ordering. The second level consists of the next-smallest components, and so on. It is possible that a node "jumps levels," so that C^+ (to use the notation from Step 5a) has level N but C^- has level $N+5$, for example. In this case, we add the appropriate number of vertical cylinders/strips between these components.

We may now remove the extra marked points $\{q_{i,r,n}\}$ which we added in Step 1. This is fine since these marked points only lie on components whose projection to Σ is nonconstant, so removing these points will not create any unstable components. Thus we have a holomorphic building U which is the limit of the curves u_i.

Finally, note that this limit curve U belongs to the homology class B. This is because the marked points $q_{r,i,n}$ determine the domain of our holomorphic building. Since the u_n's all had homology class B, so too does U. □

Remark 3.20. Recall in Remark 3.3 that we said that we would not need decorations. To prove SFT compactness in the more general setting, i.e., for holomorphic buildings in a general cylindrical manifold $V\times\mathbb{R}$ (see Theorem 3.21), however, we would need decorations. Decorations would be necessary to ensure that $\gamma^\pm$ (in the nodal case) and γ_i (in the puncture case) have the same parameterizations, and

thus glue together properly. But because there are no periodic Reeb orbits, there is only one preferred parameterization of our Reeb orbits, namely once along the chord $[0, 1]$.

This theorem not only implies that our moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ can be compactified by considering holomorphic buildings, but also tells us what the possible degenerations are.

Finally, we note that the proof of Theorem 3.18 may be generalized to the following statement about cylindrical manifolds $V \times \mathbb{R}$.

Theorem 3.21 ([BEH+03, Theorem 10.1],[Abb14, Theorem 3.20]). *If $W = V \times \mathbb{R}$ is a cylindrical manifold with a totally real submanifold L, then for every E_0, the space*

$$\overline{\mathcal{M}}_{S,\mu}(V \times \mathbb{R}, L, J) \cap \{E(U) \leq E\}$$

is compact. The moduli space above is the space of holomorphic buildings in $V \times \mathbb{R}$ with boundary in L, domain of topological type S, and exactly μ marked points.

3.5 Restricting degenerations

Ultimately, our goal is to define some kind of Floer homology using the holomorphic curves in the moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$. As in many Floer contexts, we will show that our moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ is in fact a manifold of dimension $\operatorname{ind}(B, S) - 1$, at least in the case that $\operatorname{ind}(B) \leq 2$.

Our first step will be to limit the possible kinds of degeneration which may occur. It is worth summarizing which degenerations are even allowed. In particular, the proof of the compactness theorem (Theorem 3.18) implies that we only have the following types of degeneration:

- The source can degenerate into a point on the boundary of the moduli space of Riemann surfaces. This happens when we pinch the conformal structure of the S_i's along some circles and/or arcs, which causes nodes to form.
  - The map could extend continuously over a node, thus sending it to a point in $\Sigma \times [0, 1] \times \mathbb{R}$. This results in a nodal surface.
  - If the $\mathbb{R}$-coordinate approaches $\pm\infty$ near a node, then we obtain a level splitting.
- It is also possible that, instead of the source degenerating, the map degenerates by becoming singular at some points. (This happens, in particular, if the gradient blows up at a marked point, e.g., a puncture.)
  - If the derivative blows up at a puncture of S, thus causing another type of level splitting at $\pm\infty$. Whether the splitting occurs at ∞ or $-\infty$ depends on whether the puncture is marked with a $+$ or with a $-$.
  - If the derivative blows up at an interior point, then we bubble off a holomorphic sphere.
  - Finally, the derivative blows up at a boundary point, in which case we bubble off a holomorphic disk.

We begin by restricting which degenerations are allowed to appear.

Lemma 3.22. *The only codimension-one degenerations which may occur from a sequence of J-holomorphic maps $u_n \in \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ converging to a holomorphic building $U = \{U_j\}$ are level splittings. In particular, the limit surface has smooth underlying source.*

Proof. First, observe that we cannot bubble of holomorphic spheres and disks. After all, a holomorphic sphere would correspond to a nontrivial element of $\pi_2(\Sigma \times [0,1] \times \mathbb{R}) = 0$, while a holomorphic disk would correspond to a nontrivial element of $\pi_2(\Sigma \times [0,1] \times \mathbb{R}, C_\alpha \cup C_\beta) = 0$. (Recall that $C_\alpha = \boldsymbol{\alpha} \times \{1\} \times \mathbb{R}$, while $C_\beta = \boldsymbol{\beta} \times \{0\} \times \mathbb{R}$.)

It therefore suffices to show that the only nodes which may form come from level splittings, i.e., we can only pinch the conformal structure along arcs with one boundary point on the α-curves and one on the β-curves. An interior Deligne–Mumford degeneration is obtained by pinching some number of circles, and results in an interior node. This is a codimension-two degeneration.

Now we show that we cannot form "cusp degenerations," that is to say, we cannot pinch the conformal structure along arcs whose boundary is on ∂S. (These degenerations have codimension one.)

Suppose we have a cusp degeneration. Let S' be the nodal surface which is obtained from S by collapsing along the arcs where the almost complex structure degenerates. Denote one of the collapsed arcs by A. Weak boundary monotonicity, i.e., Property (M-6), implies that the components of ∂S are all mapped to a different cylinder in $\Sigma \times [0,1] \times \mathbb{R}$ by each u_n. Thus the two endpoints of A lie on the same component C of ∂S. Without loss of generality we may say that $u_n(C) = \alpha_1 \times \{1\} \times \mathbb{R}$. Let $\bigcup C'$ be the union of components which are mapped to this cylinder $\alpha_1 \times \{1\} \times \mathbb{R}$ by the limit curve F.

Now consider the restriction of $\pi_{\mathbb{D}} \circ U$ to $\bigcup C'$, where $\pi_{\mathbb{D}}$ is the projection onto $[0,1] \times \mathbb{R}$. By the open mapping theorem, we know that $\pi_{\mathbb{D}} \circ U$ is constant near one of the components C' in $\bigcup C'$, and thus is constant on the component of S' which contains C'. Thus there is a component of S' whose boundary is mapped by U onto the union of the cylinders with $s = 1$, i.e., onto $\boldsymbol{\alpha} \times \{1\} \times \mathbb{R}$. From this and the fact that the α-circles are nonseparating, it follows by index calculations in [Lip06a, Section 4] that the rest of U is made up of g trivial disks, so that $B = [\Sigma]$. But this is impossible because we assume that B does not cross z.

We conclude that cusp degenerations cannot form either, so the only codimension-one degeneration left comes from level splitting. This could occur either when the derivative blows up at a puncture, or if the domain degenerates along an arc which connects a curve in $\boldsymbol{\alpha}$ to a curve in $\boldsymbol{\beta}$. Note that the former level splitting results in strip breaking; the degenerated strips are unstable and thus are not seen by the Deligne–Mumford limit. The latter case, which results in level splitting as well, is a Deligne–Mumford degeneration. Either way, the only possible degenerations are level splittings. This is exactly what we claimed. □

Proposition 3.23. *Fix an admissible almost complex structure which achieves transversality. Let u_i be a sequence of J-holomorphic maps in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ for some smooth source S which converges to a holomorphic building U. Suppose* $\operatorname{ind}(B) \leq 2$. *Then each story U_j of U satisfies (M-0)–(M-6).*

Proof. The previous lemma implies that (M-0) is satisfied. Furthermore, since boundary components of the sources for the U_j's are limits of boundary components of the sources for the u_n's, (M-1) is automatically satisfied. If $\pi_{\mathbb{D}} \circ U$ is constant on some component of S, then bubbling must have occurred. Theorem 3.18 implies (M-3), which states that energy is bounded, directly. Furthermore, we already ruled out bubbling, so (M-4) is satisfied.

Showing (M-2) takes the most work. We only sketch it out here, but refer the reader to [Lip06a, Proposition 7.1]. First, note that $\sum \operatorname{ind} U_j = \operatorname{ind} u_n$ for all n. This follows by unpacking the formula for

$\operatorname{ind}(B, S)$ given in Proposition 3.13. Index formulas analogous to the definition of $\operatorname{ind}(B)$ in the same proposition imply that, near each immersed curve with k double points, there is a $2k$-dimensional family of embedded curves. Since we assume $\operatorname{ind}(B) \leq 2$, it follows that the dimension of the resulting moduli would be too high. (Recall that $\operatorname{ind}(B) = \dim \widetilde{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S) = \dim \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S) + 1$.)

To prove (M-5), it suffices to show that boundary components can neither form nor disappear in the limit curve. But this follows from the open mapping theorem and the maximum modulus principle, as well as the fact that the projection to $[0, 1] \times \mathbb{R}$ is holomorphic.

We use the open mapping theorem again to prove (M-6), as well as the implicit fact that there are exactly $2g$ punctures in each story of the limit curve. In particular, the open mapping theorem on the holomorphic maps $\pi_{\mathbb{D}} \circ u_n$ implies that the $\mathbb{R}$-coordinate of each u_n must be monotone on any component of $\partial \dot{S}$. (This notation is to emphasize that we mean ∂S, minus the $2g$ boundary punctures.) Since each story U_j of the limit curve U satisfies (M-5), we know that it is a g-fold covering map of $[0, 1] \times \mathbb{R}$. Finally, because U_j is the limit of (translates of) the curves u_n, it follows that there are exactly g positive punctures and g negative punctures on U_j. □

To get the desired result that $\mathcal{M}^B(\mathbf{x}, \mathbf{y})$ is a smooth manifold, we still need one more lemma. We state this roughly, and without some of the technical conditions. The details may be found in [Lip06a, Appendix A].

Proposition 3.24. *Consider a height-two holomorphic building $(u_1, u_2) \in \mathcal{M}^{B_1}(\mathbf{x}, \mathbf{w}) \times \mathcal{M}^{B_2}(\mathbf{w}, \mathbf{y})$. Consider small neighborhoods U_1 of u_1 and U_2 of u_2 inside their respective moduli spaces. Then there is an open neighborhood of (u_1, u_2) in $\overline{\mathcal{M}^{B_1 * B_2}}(\mathbf{x}, \mathbf{y})$ which is homeomorphic to $U_1 \times U_2 \times [0, 1)$.*

Loosely speaking, this "gluing lemma" gives us a converse to our compactness theorem. In particular, the compactness theorem says that a sequence of holomorphic curves will converge to a holomorphic building. The gluing lemma above, on the other hand, says that , with certain conditions, we can reverse this process: Any holomorphic building may be "surrounded" by an affine neighborhood of holomorphic curves. In particular, every holomorphic building can be obtained as the limit of some sequence of holomorphic curves.

Putting everything together, we conclude the following.

Theorem 3.25. *Consider the moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y})$ of holomorphic curves satisfying (M-0)–(M-6) for some smooth source S with $\chi(S) = \chi_{\text{emb}}(B)$, where we identify translated curves. If $\operatorname{ind}(B) \leq 2$ and $B \neq [\Sigma]$, then $\mathcal{M}^B(\mathbf{x}, \mathbf{y})$ is a smooth manifold of dimension $\operatorname{ind}(B) - 1$. Furthermore, it is the interior of the compact manifold $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y})$ which comprises all of the holomorphic buildings whose stories satisfy (M-0)–(M-6).*

3.6 The moduli space in the bordered case

To set up the bordered case, with $\mathcal{H} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, z)$, let Σ be the interior of $\overline{\Sigma}$. We view it as a Riemann surface with a puncture p. Alternatively, we may view it as having a cylindrical end $\partial\overline{\Sigma} \times \mathbb{R}$. In particular, we choose a symplectic structure ω_Σ on Σ with respect to which $\partial\overline{\Sigma}$ is an infinite cylindrical end. Let j_Σ be, as before, an ω_Σ-compatible almost complex structure. We ask, furthermore, that the α-arcs are **cylindrical at** p in the following sense: For a fixed (punctured) neighborhood U of p and identification $\phi : U \to S^1 \times (0, \infty) \subset T^*S^1$, we have that both $j_\Sigma|_U$ and $\phi(\alpha_i^a \cap U)$ are invariant with respect to

$\mathbb{R}$-translation. Finally, let $\Sigma_{\bar{e}}$ be the result of filling in the puncture p of Σ. Note that j_Σ induces an almost complex structure on $\Sigma_{\bar{e}}$.

Definition 3.26. An **admissible** almost complex structure J satisfies (J-1)–(J-5), as well as the following additional requirement:

(J-6) J splits as $J = j_\Sigma \times j_{\mathbb{D}}$ in a fixed $\mathbb{R}$-invariant neighborhood of the fiber $\{p\} \times [0,1] \times \mathbb{R}$ of the puncture p.

In the closed case, $\Sigma \times [0,1] \times \mathbb{R}$ had two ends: one at ∞ and one at $-\infty$. Now, we have a third end, namely the puncture (equivalently, cylindrical end) of Σ. Thus we must allow a wider range of sources.

Definition 3.27. A **decorated source** $S^\triangleright$ consists a smooth (not nodal) Riemann surface S with boundary and with finitely many punctures on the boundary such that each puncture is labeled $+$, $-$, or e. Furthermore, each e puncture is also labeled by a Reeb chord in $(Z \setminus z, \mathbf{a})$, where $(Z, \mathbf{a}, M)$ is the pointed matched circle associated to $\partial\overline{\Sigma}$.

We also refer to the e punctures as "east punctures," since we visualize the cylindrical end of Σ as stretching out into the east.

For our holomorphic curves $u : (S, \partial S) \to (\Sigma \times [0,1] \times \mathbb{R}, C_\alpha \cup C_\beta)$, we now have the following conditions. Note that some of them are similar to conditions (M-0)–(M-6) above. However, note that we drop the condition that u must be an embedded curve. This slightly larger moduli space is more easily compactified, so using this as our definition will help us in Section 3.8.

(M-1) u is J-holomorphic.

(M-2) u is proper and extends to a proper map $u_{\bar{e}} : S_{\bar{e}} \to \Sigma_{\bar{e}} \times [0,1] \times \mathbb{R}$.

(M-3) $u_{\bar{e}}$ has finite energy, again in the symplectic field theory sense.

(M-4) $\pi_{\mathbb{D}} \circ u_{\bar{e}}$ is a g-fold branched cover. In particular, $\pi_{\mathbb{D}} \circ u$ is nonconstant on every component of $S^\triangleright$.

(M-5) For every positive puncture q, we have $\lim_{z \to q}(t \circ u)(z) = \infty$. Here t is the coordinate projection $\Sigma \times [0,1] \times \mathbb{R} \to \mathbb{R}$. Similarly, for every negative puncture q, we have $\lim_{z \to q}(t \circ u)(z) = -\infty$.

(M-6) At each east puncture q, $\lim_{z \to q}(\pi_\Sigma \circ u)(z)$ is the Reeb chord labeling z.

(M-7) $\pi_\Sigma \circ u$ does not cover the region of Σ which is adjacent to z.

(M-8) For every $t \in \mathbb{R}$ and every curve α_i^c, $u^{-1}(\alpha_i^c \times \{1\} \times \{t\})$ consists of exactly one point. Similarly, $u^{-1}(\beta_i \times \{0\} \times \{t\})$ consists of exactly one point.

Again, we call the last condition **weak boundary monotonicity**. Sometimes, we impose the condition of **strong boundary monotonicity**.

(M-9) For every $t \in \mathbb{R}$ and every arc α_i^c, $u^{-1}(\alpha_i^a \times \{1\} \times \{t\})$ consists of *at most* one point.

A holomorphic map satisfying (M-1)–(M-8) converges to a g-tuple of chords $\mathbf{x} \times [0,1]$ at $-\infty$, and similarly at ∞. Here, $\mathbf{x} = \{x_i\}$ is a **generalized generator**; it is like a generator, but more than one x_i may lie on the same α-arc because we do not impose (M-9).

Now for $B \in \pi_2(\mathbf{x}, \mathbf{y})$ and $S^{\triangleright}$ a decorated source, we may define $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ in the same way as before. Once again, if $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ is nonempty, then B is positive.

In the bordered case, we cut down our moduli space by imposing certain time requirements on the east punctures. Let $E(S^{\triangleright})$ be the set of east punctures of $S^{\triangleright}$. In particular, for $u \in \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ and a $q \in E(S^{\triangleright})$, define the evaluation map $\mathrm{ev}_q(u) = (t \circ u_{\vec{e}}(q)$ from $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ to $\mathbb{R}$. We may put all the evaluation maps together to get the map

$$\mathrm{ev} = \prod_{q \in E(S^{\triangleright})} \mathrm{ev}_q : \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}) \to \mathbb{R}^{E(S^{\triangleright})}.$$

Let $P = \{P_i\}$ be a partition of $E = E(S^{\triangleright})$. Define the **partial diagonal** Δ_P to be the subspace of $\mathbb{R}^E$ such that $x_p = x_q$ whenever p and q are east punctures in the same part P_i.

Definition 3.28. Let $\mathbf{x}$ and $\mathbf{y}$ be generalized generations, $B \in \pi_2(\mathbf{x}, \mathbf{y})$ a homology class, $S^{\triangleright}$ a decorated source, and P a partition of E. Then define

$$\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; P) \coloneqq \mathrm{ev}^{-1}(\Delta_P) \subset \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}).$$

Intuitively, a holomorphic curve in $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ goes off to east ∞ (i.e., toward the puncture/cylindrical end) as it approaches any east puncture. However, it may go toward east infinity at different times for different punctures. The partition P dictates which east punctures must go off to east ∞ at the same time. Note that the discrete partition, consisting of $\#E(S^{\triangleright})$ many singleton sets, results in no extra time conditions for the east punctures; thus $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; P)$ is just $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ in this case.

Proposition 3.29. *There is a residual set J_{reg} of almost complex structures for which the moduli spaces $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ are transversally cut out by the $\overline{\partial}$-equations, hence are smooth manifolds. For any countable set $\{M_i\}$ of submanifolds of $\mathbb{R}^E$, there is a residual set of admissible J which, furthermore, satisfy the property that* $\mathrm{ev} : \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}) \to \mathbb{R}^E$ *is transverse to each submanifold M_i.*

Thus we may always choose J which **achieves transversality** in this context, that is to say, a J such that the moduli spaces $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; P)$ are transversely cut out for all choices of $\mathbf{x}$, $\mathbf{y}$, B, $S^{\triangleright}$, and P. From now on, when we are in the bordered case, we will always assume that J achieves transversality unless otherwise specified.

We may compute the expected dimension of the smooth manifold $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; P)$.

Proposition 3.30. *The expected dimension* $\mathrm{ind}(B, S^{\triangleright}, P)$ *of $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; P)$ is*

$$\mathrm{ind}(B, S^{\triangleright}, P) \coloneqq g - \chi(S) + 2e(D(B)) + |P|,$$

where $e(D(B))$ once again denotes the Euler measure of the domain associated to B. Here $|P|$ denotes the number of parts in the partition P.

We may now divide $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ up into different strata based on the time order in which each partition class of east puncture goes off to infinity.

Definition 3.31. Let $\vec{P}$ be an ordered partition of the e punctures of $S^\triangleright$, where P is its associated unordered partition. Let $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ be the open subset of $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ comprising those holomorphic curves such that the ordering of P induced by t agrees with the ordering in $\vec{P}$. In other words, we say that $u \in \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ if $\mathrm{ev}_q(u) < \mathrm{ev}_{q'}(u)$ for all $q \in P_i$ and $q' \in P_{i'}$ with $i < i'$ in $\vec{P}$.

There is an $\mathbb{R}$-action on $\Sigma \times [0,1] \times \mathbb{R}$, namely translation in the $\mathbb{R}$-coordinate. As long as $S^\triangleright$ is **stable**, i.e., $S^\triangleright$ is not the trivial collection of g disks with two boundary punctures each and $B = 0$, this action is free. In this case, we are most interested in the reduced moduli spaces

$$\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P) := \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P)/\mathbb{R} \quad \text{and} \quad \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}) := \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})/\mathbb{R}.$$

Outside of the trivial case $B = 0$, these moduli spaces have dimension $\mathrm{ind}(B, S^\triangleright, P) - 1$.

Since the action is translation in the t-coordinate, the evaluation maps ev_q do not descend to the quotient. However, the difference $\mathrm{ev}_p - \mathrm{ev}_q$ does, for any two e punctures p and q. Thus we define

$$\mathrm{ev}_{p,q} = \mathrm{ev}_p - \mathrm{ev}_q : \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}) \to \mathbb{R}.$$

Furthermore, we can combine the evaluation maps into a single map

$$\mathrm{ev} : \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}) \to \mathbb{R}^E/\mathbb{R}.$$

Here $\mathbb{R}$ acts by translation on each coordinate of $\mathbb{R}^E$.

An unordered partition P of E gives rise to a set $[P]$ of multi-sets of Reeb chords which is defined by replacing the punctures in P with the associated Reeb chords. Similarly, an ordered partition $\vec{P}$ gives rise to a sequence of multi-sets of Reeb chords, which we denote $[\vec{P}]$.

Recall that in Proposition 3.13 we were able to give a source-independent formula for the dimension of the moduli space. Similarly, when $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ has an embedded representative, there is a domain-invariant definition for the index. (We only need this for $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$, and not for $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$, because that is the moduli space whose elements our bordered Heegaard Floer modules will count.)

We require a few definitions first.

Recall that a curve u satisfying both (M-8) and (M-9) is said to satisfy strong boundary monotonicity, and thus is asymptotic to a genuine generator. Maps which are only weakly boundary monotonic are only asymptotic to generalized generators. In fact, the notion of strong boundary monotonicity is purely combinatorial and only depends on the asymptotics of the curve.

Definition 3.32. Let $\mathbf{s}$ be a k-element multi-set of $[2k]$, i.e., a formal linear combination of elements of $[2k]$ whose coefficients are in $\mathbb{N} \cup \{0\}$ and sum to k. Let $\vec{\boldsymbol{\rho}} = (\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_n)$ be a sequence of nonempty multi-sets of Reeb chords. Define

$$o(\mathbf{s}, \vec{\boldsymbol{\rho}}) := \left[\mathbf{s} \cup \left(\textstyle\bigcup_i M(\boldsymbol{\rho}_i^+)\right)\right] \setminus \left(\textstyle\bigcup_i M(\boldsymbol{\rho}_i^-)\right),$$

where the union and difference operations are taken as multi-sets (i.e., by adding or subtracting linear combinations to get linear combinations with coefficients in $\mathbb{Z}$). Recall that M is the matching of the pointed matched circle $\mathcal{Z} = (Z, \mathbf{a}, M)$. Then we say that the pair $(\mathbf{s}, \vec{\boldsymbol{\rho}})$ is **strongly boundary monotone** if the following two conditions are satisfied for each $i = 0, \dots, n$:

(1) The multi-set $o(\mathbf{s}, (\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_i))$ is a k-element subset of $[2k]$ with no repeated elements (hence is a genuine set).

(2) The multi-sets $M(\boldsymbol{\rho}_i^+)$ and $M(\boldsymbol{\rho}_i^-)$ have no elements with multiplicity greater than 1.

As suggested by the notation, this is related to the strong boundary monotonicity of a curve.

Proposition 3.33 ([LOT18, Lemma 5.53]). *A curve $u \in \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ is strongly boundary monotone if and only if $(\mathbf{x}, [\vec{P}])$ is strongly boundary monotone, where $[\vec{P}]$ is the set of multi-sets of Reeb chords which is given by replacing each puncture of $S^\triangleright$ with its associated Reeb chord.*

A sequence $\boldsymbol{\rho} = (\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_n)$ of multi-sets of Reeb chords is **compatible** with a homology class $B \in \pi_2(\mathbf{x}, \mathbf{y})$ if $\partial^\partial B = [\vec{\boldsymbol{\rho}}]$ in $H_1(Z, \mathbf{a})$, and $(\mathbf{x}, \vec{\boldsymbol{\rho}})$ is strongly boundary monotone.

Finally, one may define a so-called Maslov index $\iota(\vec{\boldsymbol{\rho}})$. We do not go into the definition, but note simply that it is an integer depending only on the sequence $\vec{\boldsymbol{\rho}}$ of sets of Reeb chords. See, for example, [LOT18, Section 3.3] or [Lin12, Section 1.5].

Proposition 3.34 ([LOT18, Proposition 5.69]). *Let $u \in \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$. If x_i and y_i are the components in the generators $\mathbf{x}$ and $\mathbf{y}$, respectively, then*

$$\chi(S) = g - \left(\sum_{i=1}^{g} n_{x_i}(B) + n_{y_i}(B) \right) + e(D(B)) - \iota(\vec{\boldsymbol{\rho}})$$

if and only if u is embedded, i.e., if and only if $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ has an embedded holomorphic representative. From this and Proposition 3.30, we conclude that the expected dimension of $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ is

$$\operatorname{ind}(B, [\vec{P}]) := \left(\sum_{i=1}^{g} n_{x_i}(B) + n_{y_i}(B) \right) + e(D(B)) + |\vec{P}| - \iota[\vec{P}]).$$

Thus we may define a moduli space of embedded curves which connect generators $\mathbf{x}$ and $\mathbf{y}$, belong to the homology class B, and have prescribed asymptotics $\vec{\boldsymbol{\rho}}$ as east infinity. In particular, for compatible pairs $(B, \vec{\boldsymbol{\rho}})$, define

$$\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}}) := \bigcup_{\substack{\chi(S^\triangleright) = \chi_{\mathrm{emb}}(B, \vec{\boldsymbol{\rho}}) \\ [\vec{P}] = \vec{\boldsymbol{\rho}}}} \widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}).$$

Here $\chi_{\mathrm{emb}}(B, \vec{\boldsymbol{\rho}})$ is exactly the formula for $\chi(S)$ in Proposition 3.34.

In general, we will always be discussing moduli spaces where $(B, \vec{\boldsymbol{\rho}})$ are compatible, even when we do not specify it.

3.7 Holomorphic combs

Recall that the compactification of the moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ in Section 3.2 requires that we add holomorphic buildings. This allowed the curves to "break" at $t = \pm\infty$. To compactify our moduli spaces $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ and $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ of holomorphic curves connecting $\mathbf{x}$ to $\mathbf{y}$, we must allow the curves to break as well. This breaking can happen at $t = \pm\infty$, or at east ∞, i.e., at the infinite cylindrical end corresponding to $\partial\overline{\Sigma}$.

3.7.1 Holomorphic curves in $\mathbb{R} \times Z \times [0,1] \times \mathbb{R}$

The breaking at $\pm\infty$ is similar to the breaking in Section 3.3 via holomorphic buildings. The breaking at east ∞, on the other hand, requires that we discuss holomorphic curves in $\mathbb{R} \times Z \times [0,1] \times \mathbb{R}$, which we equip with the standard split symplectic form and a fixed split almost complex structure $J = j_\Sigma \times j_\mathbb{D}$.

Such a curve differs from a holomorphic curve in $\Sigma \times [0,1] \times \mathbb{R}$ because it may potentially have up to four different types of ends: two at $\pm\infty$ in the first coordinate, and two more at $\pm\infty$ in the last coordinate. The first coordinate is our "east–west" direction; we call $-\infty$ and $+\infty$ in the first coordinate west ∞ and east ∞, respectively.

We again have projection maps $\pi_\Sigma : \mathbb{R} \times Z \times [0,1] \times \mathbb{R} \to \mathbb{R} \times Z$ and $\pi_\mathbb{D} : \mathbb{R} \times Z \times [0,1] \times \mathbb{R} \to [0,1] \times \mathbb{R}$.

Definition 3.35. A **bidecorated source** $T^\diamond$ is a smooth (not nodal) Riemann surface T with boundary and with finitely many punctures on the boundary. Each puncture is labeled with either e or w, as well as with a Reeb chord in $(Z, \mathbf{a})$.

Definition 3.36. Let $T^\diamond$ be a bidecorated source. Define $\widetilde{N}(T^\diamond)$ to be the moduli space of proper holomorphic maps

$$v : (T, \partial T) \to (\mathbb{R} \times (Z \setminus z) \times [0,1] \times \mathbb{R}, \mathbb{R} \times \mathbf{a} \times \{1\} \times \mathbb{R})$$

which extend to east and west ∞ as dictated by the labelings of the punctures. That is, if the west puncture q is labeled by the Reeb chord ρ, then $\lim_{z \to q}(\pi_\Sigma \circ u)(z) = \{-\infty\} \times \rho$, and similarly if q is an east puncture.

Note that the maps in $\widetilde{N}(T^\diamond)$ take all of ∂T to $s = 1$, where s is the $[0,1]$-coordinate. As such, by the open mapping theorem, every component of T maps to a single point in $[0,1] \times \mathbb{R}$. After all, proper maps are closed, so if $\pi_\mathbb{D} \circ v$ were nonconstant, then it would be surjective. Note that this means that we do not have punctures at $t = \pm\infty$, so the curves in $\widetilde{N}(T^\diamond)$ only have east and west punctures.

We again have an evaluation map $\mathrm{ev}_q : \widetilde{N}(T^\diamond) \to \mathbb{R}$ for each puncture q of $T^\diamond$ given by $v \mapsto \lim_{z \to q}(t \circ v)(z)$. (This limit is trivial since the t-coordinate is constant on each connected component of $T^\diamond$.) We define west and east evaluation maps

$$\mathrm{ev}_w := \prod_{q \in W(T^\diamond)} \mathrm{ev}_q : \widetilde{N}(T^\diamond) \to \mathbb{R}^{W(T^\diamond)}$$

and

$$\mathrm{ev}_e := \prod_{q \in E(T^\diamond)} \mathrm{ev}_q : \widetilde{N}(T^\diamond) \to \mathbb{R}^{E(T^\diamond)}.$$

Here W and E are the sets of west and east punctures, respectively.

As in the context of $\widetilde{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright)$, we may use these evaluation maps to cut down $\widetilde{N}$, as follows.

Definition 3.37. Let P_w and P_e be partitions of the west and east punctures, respectively. Define

$$\widetilde{N}(T^\diamond; P_w, P_e) := (\mathrm{ev}_w \times \mathrm{ev}_e)^{-1}(\Delta_{P_w} \times \Delta_{P_e}).$$

When P_w is the discrete (trivial) partition, we denote $\widetilde{N}(T^\diamond; P_w, P_e)$ by $\widetilde{N}(T^\diamond; P_e)$.

In general, the moduli spaces $\widetilde{N}(T^\diamond; P_w, P_e)$ are not transversally cut out, hence not manifolds. There is a special case, however, when all of the components of T are topological disks.

Proposition 3.38 ([Lip06b, Lemma 4.1.2]). *Suppose $T^\diamond$ is a bidecorated source whose components are all topological disks. Then, if $\mathbb{R}\times Z\times[0,1]\times\mathbb{R}$ is equipped with a split almost complex structure, the associated moduli space $\widetilde{N}(T^\diamond)$ is transversally cut out by the $\overline{\partial}$-equation.*

We will mainly be interested in the reduced moduli space. In particular, there is an $(\mathbb{R}\times\mathbb{R})$-action by translation on the two $\mathbb{R}$-coordinates of $\mathbb{R}\times Z\times[0,1]\times\mathbb{R}$. This induces an $(\mathbb{R}\times\mathbb{R})$-action on $\widetilde{N}(T^\diamond)$. This action is usually free. In particular, we call a holomorphic map v **stable** if

- at least one component of its source $T^\diamond$ is not a twice-punctured disk; and
- if $\pi_\Sigma\circ v$ is constant on some component C of $T^\diamond$, then C has no nontrivial automorphisms.

If every $v\in\widetilde{N}(T^\diamond)$ is stable, then the $(\mathbb{R}\times\mathbb{R})$-action on $\widetilde{N}(T^\diamond)$ is free. Thus we can make the following definition in this case.

Definition 3.39. If $T^\diamond$ is a bidecorated source and $\widetilde{N}(T^\diamond)$ is stable, then

$$\mathcal{N}(T^\diamond) := \widetilde{N}(T^\diamond)/(\mathbb{R}\times\mathbb{R}).$$

3.7.2 Some examples

We will now give some names to a few particularly useful holomorphic curves in $\mathbb{R}\times Z\times[0,1]\times\mathbb{R}$.

Definition 3.40. A **trivial component** is a twice-punctured topological disk where one puncture is labeled e, one is labeled w, and both are labeled by the same Reeb chord.

Holomorphic maps on trivial components are not particularly interesting, as they are preserved under translation of the first $\mathbb{R}$-coordinate.

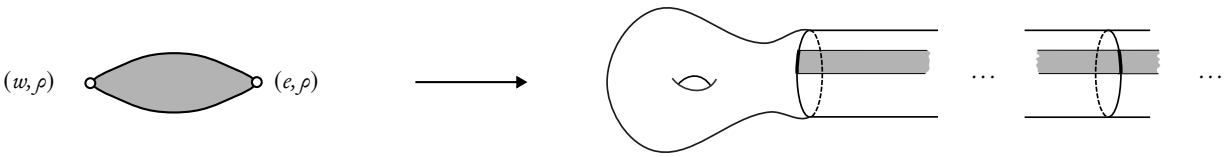

Figure 3.12: A trivial component. Note that the image is the projection to $\Sigma\cup(\mathbb{R}\times Z)$, which is topologically equivalent to Σ. The thickened black chord is the Reeb chord ρ which is determined by the labelings of the punctures. (In later figures, we will only draw the cylinder.)

Definition 3.41. A **join component** is a topological disk with two west punctures and one east puncture. Similarly, a **split component** is a topological disk with two east punctures and one west puncture. A stable curve which is entirely made up on join (respectively, split) components is called a **join** (respectively, **split**) **curve**.

Examples are shown in Figures 3.13 and 3.14. Using the same notation, there exists a holomorphic map v in the moduli space for a join component if and only if $\rho_e = \rho_1 \uplus \rho_2$. Similarly, there exists a holomorphic map v whose source is a split component if and only if $\rho_w = \rho_1 \uplus \rho_2$. Such maps v are unique up to translation in the t-coordinate.

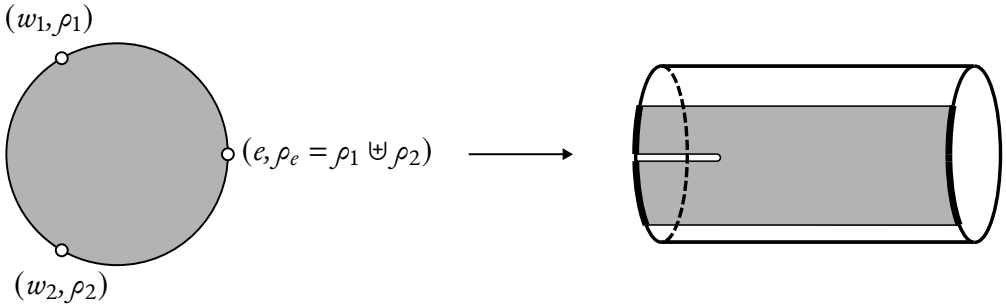

Figure 3.13: A join component. Note that $\rho_e = \rho_1 \uplus \rho_2$ because the boundary of the component must be mapped to the discrete set $\mathbf{a} = \overline{\boldsymbol{\alpha}} \cap \boldsymbol{\beta}$.

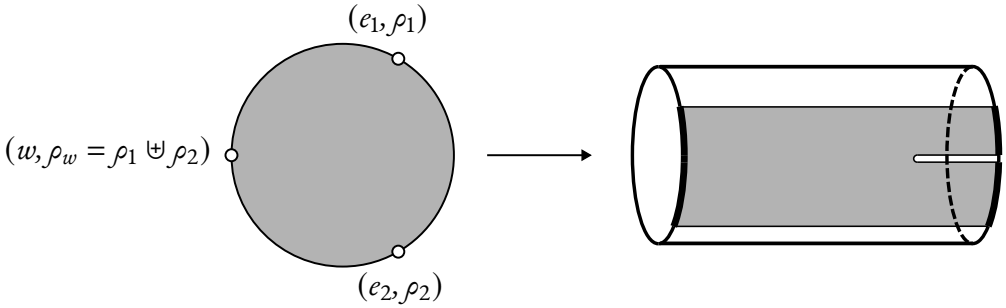

Figure 3.14: A split component. Similarly, we require $\rho_w = \rho_1 \uplus \rho_2$.

Definition 3.42. A **shuffle component** is a topological disk with four punctures, two east and two west, which are ordered east, west, east, west around the boundary. If the two Reeb chords associated to the west punctures are interleaved and the Reeb chords associated to the east punctures are nested, then we call it an **odd shuffle component**. On the flip side, if the two Reeb chords associated to the west punctures are *nested* and the Reeb chords associated to the east punctures are *interleaved*, then we call it an **even shuffle component**.

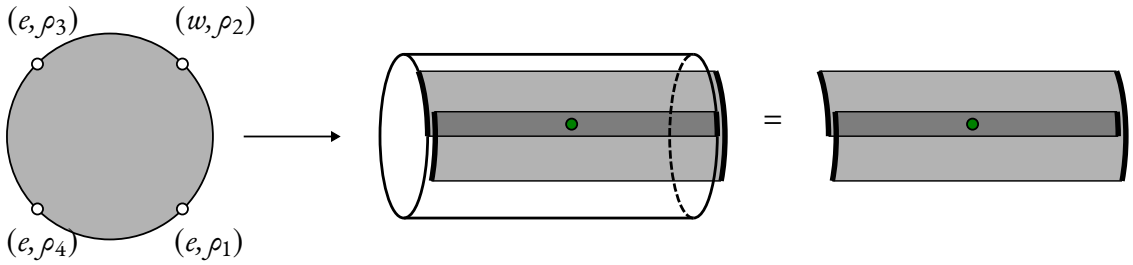

Figure 3.15: An (odd) shuffle component. Its mirror image would be an even shuffle component. The green dot denotes a branch point.

With the notation in Figure 3.15, there exists a holomorphic map whose source is a shuffle component if and only if $\rho_1^+ = \rho_2^+, \rho_2^- = \rho_3^-, \rho_3^+ = \rho_4^+$, and $\rho_4^- = \rho_1^-$. Thus if the associated moduli space is nonempty, then the shuffle component is either odd or even.

Definition 3.43. A holomorphic curve composed of exactly one shuffle component and some number of trivial components is called a **shuffle curve**. The shuffle curve may be **odd** or **even**, depending on the parity of the shuffle component.

3.7.3 Holomorphic combs

Returning now to our original goal of defining a way for our holomorphic curves in $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; P)$ and $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright}; \vec{P})$ to break at east ∞, we will define a generalization of holomorphic buildings known as *holomorphic combs*. As in the case of holomorphic buildings, we allow holomorphic combs with nodal sources. Roughly speaking, the difference between holomorphic combs and holomorphic buildings is that combs allow for a degeneration at east ∞ by breaking off holomorphic curves in $\mathbb{R} \times Z \times [0,1] \times \mathbb{R}$—hence our discussion of $\widetilde{N}(T^{\diamond})$ earlier. (In fact, holomorphic buildings may be even further generalized. Exploded manifolds, for example, give a way for such degeneration to occur in directions like "northeast infinity." [Par12])

Definition 3.44. A **holomorphic story** is a sequence $(u, v_1, \dots, v_k)$ of **horizontal levels**, where $k \geq 0$, such that

- $u \in \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$ for some B and $S^{\triangleright}$;
- $v_i \in \widetilde{N}(T_i^{\diamond})$ for some $T_i^{\diamond}$;
- there is a one-to-one correspondence between $E(S^{\triangleright})$ and $W(T_1^{\diamond})$, as well as between $E(T_i^{\diamond})$ and $W(T_{i+1}^{\diamond})$ for $i = 1, \dots, k-1$, which preserves the labelings by Reeb chords; and
- $\mathrm{ev}(u) = \mathrm{ev}_w(v)$ and $\mathrm{ev}_e(v_i) = \mathrm{ev}_w(v_{i+1})$ for $i = 1, \dots, k-1$.

Intuitively, then, a holomorphic story allows our holomorphic curves to break off at east ∞. The requirements that $\mathrm{ev}(u) = \mathrm{ev}_w(v_1)$ in $\mathbb{R}^{E(S^{\triangleright})}/\mathbb{R} \cong \mathbb{R}^{W(T^{\diamond})}/\mathbb{R}$ and $\mathrm{ev}_e(v_i) = \mathrm{ev}_w(v_{i+1})$ in $\mathbb{R}^{E(T_i^{\diamond})}/\mathbb{R} \cong \mathbb{R}^{W(T_{i+1}^{\diamond})}/\mathbb{R}$ are so that the breaking-off at each east–west puncture pair happens at a well-defined time.

Finally, like with holomorphic buildings, we must allow degeneration at $\pm\infty$.

Definition 3.45. A **holomorphic comb** of height N is a sequence $(u_j, v_{j,1}, \dots, v_{j,k_j})$ for $j = 1, \dots, N$ of holomorphic stories. In general, we use the notation that u_j is a stable curve in $\mathcal{M}^{B_j}(\mathbf{x}_j, \mathbf{x}_{j+1}; S_j^{\triangleright})$, where $\mathbf{x}_1, \dots, \mathbf{x}_{N+1}$ are generalized generators. The index j is the **vertical level**.

The **trivial holomorphic comb** has $N = 0$ and corresponds to a trivial (unstable) holomorphic curve. We call a holomorphic comb **simple** if it has one level, and if that single holomorphic story is (u, v), i.e., has $k = 1$. We call a holomorphic comb **toothless** if it has no components at east ∞. (Note that a toothless comb is basically a holomorphic building.) Finally, the **spine** of a holomorphic comb is the part of it which is toothless, i.e., the sub-comb of components which map to $\Sigma \times [0,1] \times \mathbb{R}$.

A height-N holomorphic comb U naturally represents a homology class $B \in \pi_2(\mathbf{x}_1, \mathbf{x}_{N+1})$. In particular, if B_j is the domain of u_j, then $B = B_1 * \cdots * B_N$. Furthermore, at (far) east ∞, we see that U has the asymptotics of the east punctures of the v_{j,k_j}.

A schematic representation of a holomorphic comb may be seen in Figure 3.16. Note that each component at east infinity occurs at a fixed point $(s, t) \in \mathbb{R}$ within its level in the holomorphic building. Furthermore, components at east infinity must be topological disks.

Now we define what it means for a sequence of holomorphic combs to converge to another holomorphic comb. This will define a topology on the moduli space of holomorphic combs, and hence allow us to compactify $\widetilde{\mathcal{M}}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright})$.

For simplicity of notation, we will define convergence of a sequence of holomorphic *curve* to a holomorphic comb. The general definition is not much different.

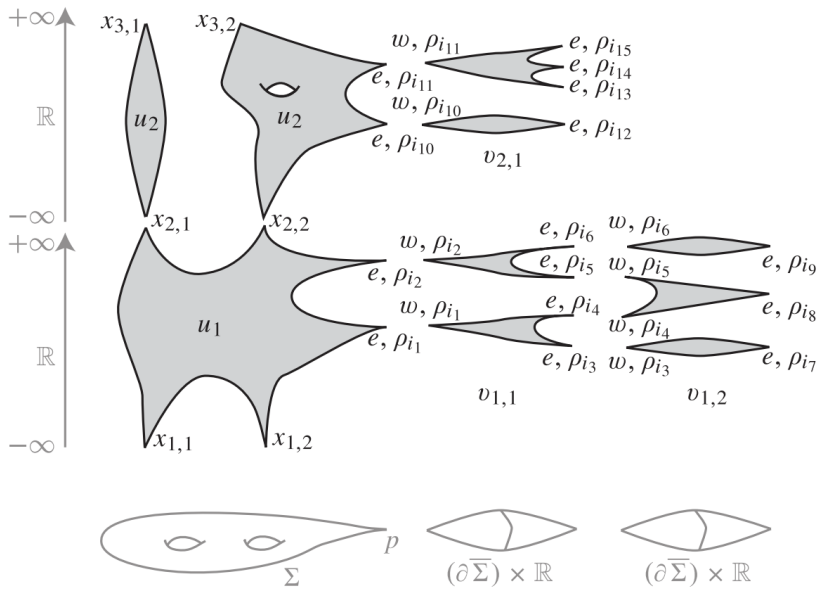

Figure 3.16: A schematic diagram of a two-story holomorphic comb whose first story contains components $(u_1, v_{1,1}, v_{1,2})$ and whose second story is $(u_2, v_{2,1})$ [LOT18]. Note that the components at east infinity all occur at a single time t (and, indeed, at a single point $s \in [0,1]$ as well).

Definition 3.46. A sequence $\{u_n\}$ of holomorphic curves in $\Sigma \times [0,1] \times \mathbb{R}$ converges to a holomorphic comb U if the following conditions hold.

- Let S_Σ be the result of collapsing the components C of the preglued domain of U for which $(\pi_\Sigma \circ U)|_C$ is an unstable map. Then $\{\pi_\Sigma \circ u_n\}$ converges to $\{\pi_\Sigma \circ U|_{S_\Sigma}\}$ as a holomorphic building.
- Similarly, if $S_\mathbb{D}$ is the result of collapsing the components C for which $(\pi_\mathbb{D} \circ U)|_C$ is unstable, then $\{\pi_\mathbb{D} \circ u_n\}$ converges to $\pi_\mathbb{D} \circ U|_{S_\mathbb{D}}$ as a holomorphic building.
- Let τ_t be the translation of $\Sigma \times [0,1] \times \mathbb{R}$ by t units in the $\mathbb{R}$ direction. Let q be a smooth point in the spine of U. There is a neighborhood V of q, as well as a sequence of points q_n in the source of u_n with neighborhoods V_n diffeomorphic to V, such that there are numbers $t_n \in \mathbb{R}$ so that $\tau_{t_n} \circ u|_{V_n}$ converges to $U|_V$ in the C^∞_{loc} topology.
- For sufficiently large n, the maps u_n represent the same homology class B which is represented by the holomorphic comb U.

Recall the idea of preglued surfaces, which allow us to consider the domain of a holomorphic building as a single Riemann surface. We may do the same thing for both nodal or smooth holomorphic combs now. This allows us to define the following moduli spaces.

Definition 3.47. The moduli space of all (possibly nodal) holomorphic combs in the homology class B, whose preglued surfaces are $S^\triangleright$, and with asymptotics $\mathbf{x}$ at $-\infty$ and $\mathbf{y}$ at ∞ is denoted $\overline{\overline{\mathcal{M}^B}}(\mathbf{x}, \mathbf{y}; S^\triangleright)$. The closure of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright)$ within this space is denoted $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright)$.

To define the compactified moduli space when we partition our punctures, note that for any punctures $p, q \in E(S^\triangleright)$, we may extend $\mathrm{ev}_{p,q}$ to

$$\overline{\mathrm{ev}}_{p,q} : \overline{\overline{\mathcal{M}^B}}(\mathbf{x}, \mathbf{y}; S^\triangleright) \to [-\infty, \infty].$$

Then we can make the following definitions.

Definition 3.48. The moduli space of all holomorphic combs in $\overline{\overline{\mathcal{M}^B}}(\mathbf{x}, \mathbf{y}; S^\triangleright)$ which respect the partition P is given by

$$\overline{\overline{\mathcal{M}^B}}(\mathbf{x}, \mathbf{y}; S^\triangleright; P) := \bigcap_{\substack{P_i \in P \\ p,q \in P_i}} \overline{\mathrm{ev}}_{p,q}^{-1}(0).$$

The closure of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ within this space is denoted $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$. Similarly, the closure of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ within $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ is denoted $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$.

Difficulties with transversality at east ∞ make it so that we must define our moduli spaces as closures within the space of all holomorphic combs, rather than simply as the space of all holomorphic combs (cf. Definition 3.16. An example of when $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ is a proper subset of $\overline{\overline{\mathcal{M}^B}}(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ is given in [LOT18, Example 5.23].

3.8 Compactification via holomorphic combs

Much as in the case of holomorphic buildings, we now have a compactness result with holomorphic combs. In particular, we have the following theorem.

Theorem 3.49. *The spaces $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright)$ are compact. If $\{U_n\}$ is a sequence of holomorphic combs in a fixed homology class B and fixed (topological) preglued source $S^\triangleright$, then there is a subsequence converging to a (possibly nodal) holomorphic comb $U \in \overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright)$. Similarly, the moduli spaces $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ and $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ are compact.*

We will only prove the first statement, namely that $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright)$ is compact. It is sufficient, as in Section 3.4, to prove the following.

Theorem 3.50. *Any sequence $\{u_n : S_n \to \Sigma \times [0,1] \times \mathbb{R}\}$ of holomorphic curves in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright)$ has a subsequence which converges to a holomorphic comb $U \in \overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright)$.*

To prove it, we need another version of SFT compactness for maps in manifolds with cylindrical ends. In particular, recall that Theorem 3.18 was only for cylindrical manifolds of the form $V \times \mathbb{R}$. Because Σ has a cylindrical end at the puncture p, however, we will need this more general version of SFT compactness.

Theorem 3.51 ([LOT18, Theorem 5.29])**.** *If (W, j) is a punctured Riemann surface with a Lagrangian submanifold L which is cylindrical near the punctures of W and which is embedded away from finitely many transverse self-intersections, then the space*

$$\overline{\mathcal{M}}_{S,\mu}(W, L, J) \cap \{E(F) \leq E\}$$

is compact. The notation $\overline{\mathcal{M}}_{S,\mu}(W, L, J)$ refers to holomorphic buildings in W with boundary in L, domain of topological type S, and exactly μ marked points.

Proof of Theorem 3.50. We will define the limit holomorphic comb U by defining $\pi_{\mathbb{D}} \circ U$ and $\pi_\Sigma \circ U$ separately. The former defines the "vertical" level structure (i.e., the part of the holomorphic comb which is a building in $[0,1]\times\mathbb{R}$) while the latter defines the "horizontal" level structure (i.e., the part of the comb which is a building in Σ, thought of as a manifold with a cylindrical end at east infinity).

Step 1. *Obtaining the vertical level structure.* Pick generic points p_r in each region r of Σ which are regular values of $\pi_\Sigma \circ u_n$ for all n. Let $\{q_{r,i,n}\} = (\pi_\Sigma \circ u_n)^{-1}(p_r)$ be the preimages of p_r. Note that $q_{r,i,n}$ is a point in the source S_n of u_n. Think of these points as marked points in S_n. These implicitly keep track of the homology class B which is represented by each holomorphic curve u_n, since the number of points $q_{r,i,n}$ for each region r tells us how many times u_n crosses the r.

We may use Theorem 3.21 to extract a convergent subsequence of $\{u_n\}$, where we have added the points $q_{r,i,n}$ to the marked point sets of S_n. SFT compactness applies because the maps u_n are holomorphic maps with bounded energy: The ω-energy

$$\int_{S_n \cup_\partial (-S_n)} u_n^* \omega_{[0,1]}$$

vanishes since there is no nontrivial 2-form on $[0,1]$. The λ energy is the degree of the map, which in this case is $g = g(\Sigma)$. (See Remark 3.9.)

Relabel so that $\{v_n\}$ refers to the convergent subsequence. Then $\{\pi_{\mathbb{D}} \circ u_n\}$ also converges, and thus has a limit $\{\pi_{\mathbb{D}} \circ u_n\} \to \pi_{\mathbb{D}} \circ U$. (Note that we simply use $\pi_{\mathbb{D}} \circ U$ to denote this limit. In particular, we have not yet defined U itself.) This limit is a (vertical) holomorphic building in $[0,1] \times \mathbb{R}$. Note that we have used (J-2) to ensure that $\pi_{\mathbb{D}} \circ u_n$ is holomorphic. Say it has source S_∞.

This gives a vertical level structure on the limit. It remains to find a horizontal level structure, i.e., to understand the components at east ∞.

Let $V_p \subset \Sigma$ be a closed disk neighborhood of p, i.e., of east ∞. (Recall that we think of Σ as the interior of $\overline{\Sigma}$, i.e., as a Riemann surface with a cylindrical end at the east puncture p.) For small V_p, we know by (J-6) that the almost complex structure J splits:

$$J|_{V_p \times [0,1] \times \mathbb{R}} = j_\Sigma \times j_{\mathbb{D}}.$$

Let W_p be the complement of a closed disk around p which is slightly smaller than V_p, so that Σ is a union of the interiors of V_p and W_p. We will find convergent subsequences over V_p and W_p, and then show that they agree on the overlap.

Step 2. *Convergence over V_p.* First, we tackle convergence near east infinity. Define $T_n := (\pi_\Sigma \circ u_n)^{-1}(V_p)$, so that our maps restrict to holomorphic maps

$$(\pi_\Sigma \circ u_n)|_{T_n} : (T_n, \partial T_n) \to (V_p, \partial V_p \cup \boldsymbol{\alpha}).$$

(The $\boldsymbol{\alpha}$ is there because α-arcs go through p, and so $\pi_\Sigma \circ u_n$ might map a point in ∂T_n to these arcs. Since there are no β-curves near p, we need not include $\boldsymbol{\beta}$.) These maps have a finite energy bound, i.e., "area" bound, since they are restrictions of maps which all represent the same fixed homology class $B \in H_2(\Sigma)$. Thus Theorem 3.51 applies.

Letting $\{u_n\}$ be this convergent subsequence, we have a holomorphic building $(\pi_\Sigma \circ U)|_T$ defined as the limit of $(\pi_\Sigma \circ u_n)|_{T_n}$. Note that this holomorphic building is "horizontal." In particular, whereas $\pi_{\mathbb{D}} \circ U$ had stories at $\pm\infty$, this building $(\pi_\Sigma \circ U)|_T$ has stories at east ∞. It is a building in $\mathbb{R} \times Z$.

This might introduce more components to the (presumptive) source of U, namely components which are not seen by $\pi_{\mathbb{D}} \circ U$. Let S_∞ thus denote this new source. Define T' to be the parts of this new S_∞ which are mapped away from east ∞, i.e.,

$$T' := S_\infty \setminus [(\pi_\Sigma \circ U)|_T]^{-1}(W_p^c).$$

Step 3. *Convergence over W_p.* Now we show that, over W_p, we converge to a holomorphic curve over T'. This will define $\pi_\Sigma \circ U$ on T'. (Note that $\pi_{\mathbb{D}} \circ U|_{T'}$ was already defined in Step 1.)

First, we do this for smooth points. In particular, consider some smooth point $q \in T'$, and let X be a neighborhood of q which is contained in T' and which does not touch any nodes or punctures in S_∞. We would like to define $\pi_\Sigma \circ U$ on X by finding a subsequence $\{\pi_\Sigma \circ u_{n,X}\}$ which converges to a holomorphic curve $(\pi_\Sigma \circ U)|_X : X \to W_p$. Hence $\{u_{n,X}\}$ would converge to $U|_X : X \to W_p \times [0,1] \times \mathbb{R}$.

If we can do this, then we may choose a countable collection $\{X_n\}$ of X which covers the smooth part of T' and take the diagonal subsequence $\{u_{n,X_n}\}$. In particular, the resulting subsequence, which we may denote as $\{u_n\}$, converges in C^∞_{loc} to U away from the collapsed curves in the source. (Recall that $S_n \to S_\infty$ in the Deligne–Mumford sense. These collapsed curves are where the complex structures form infinitely long necks, thus producing nodes in S_∞.)

Notice that

$$\pi_2(\Sigma \times [0,1] \times \mathbb{R}) = \pi_2(\Sigma \times [0,1] \times \mathbb{R}, C_\alpha \cup C_\beta) = 0,$$

so no spheres may bubble off. In particular, since bubbling occurs when the gradient approaches infinity at a point, it follows that $\|du_n\|$ is bounded on X. Since we once again have an energy bound, we may apply [MS12, Theorem 4.1.1], for example, to obtain a subsequence of the u_n which converges to a holomorphic curve $(\pi_\Sigma \circ U)|_X$ over X. Note that this limit may escape toward $\pm\infty$, as detailed in Step 1, but has neither bubbling phenomena nor any escape toward east infinity.

This defines the holomorphic comb on T' away from any nodes and punctures. The cases of nodes and punctures may largely be solved by using existing compactness and convergence theorems. For nodes, we use the energy bound on the $\pi_\Sigma \circ u_n$'s again, this time apply the removable singularities theorem (e.g., [MS12, Theorem 4.1.2]). This means that we may extend U across the nodes. To see that it approaches the same value from both sides of the nodes involves an argument similar to the argument in the proof of convergence in the thin part in Theorem 3.18. One may also refer to [MS12, Section 4.7] for a similar argument in a slightly different setting.

For punctures of $U|_{T'}$, we would like to show that the map $\pi_\Sigma \circ U$ approaches points in $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$. (Recall that generators $\mathbf{x} \in \mathfrak{S}(\mathcal{H})$ are g-element subsets of $\boldsymbol{\alpha} \cap \boldsymbol{\beta}$.) But this follows from [Flo88d, Theorem 2].

Step 4. *Piecing together T and T'.* We have, at this point, defined $\pi_\Sigma \circ U$ on both T and T'. On their intersection $T \cap T'$, we have defined the holomorphic comb U twice. These coincide because C^∞_{loc} is Hausdorff, and thus the sequence $\{u_n|_{T \cap T'}\}$ may only have one limit. Thus we may glue $U|_T$ and $U|_{T'}$ to get a holomorphic comb U in $\Sigma \times [0,1] \times \mathbb{R}$ which is defined on all of S_∞.

We claimed that this map U should still belong to the homology class B. This follows from the marked points $q_{r,i,n}$ which we added and which determine the domain of our holomorphic comb.

Step 5. *Defining $\pi_{\mathbb{D}} \circ U$ on components at east infinity.* We have defined π_Σ on $U_p \cup V_p$, hence on S_∞. But because we added some components to S_∞ at Step 2, we have not yet defined $\pi_{\mathbb{D}} \circ U$ on the components which only appear at east infinity. In other words, we have not defined $\pi_{\mathbb{D}} \circ U$ on the components of S_∞ which map to $\mathbb{R} \times Z \times [0,1] \times \mathbb{R}$ instead of to $\Sigma \times [0,1] \times \mathbb{R}$.

We will show that $\pi_{\mathbb{D}} \circ U$ is constant on each such component. Thus each escape to east infinity happens at a fixed point $(s,t) \in [0,1] \times \mathbb{R}$.

Let C be a fixed component of S_∞ which is mapped to east ∞. If $(\pi_{\mathbb{D}} \circ U)|_C$ is stable, then this component was already defined before, namely when we were defining the vertical level structure of the limit. We have thus already defined a map $C \to [0,1] \times \mathbb{R}$. The surface S_∞ was created by collapsing some arcs in S, the topological type of the decorated sources of the curves u_n, so let C_0 be the preimage of $C \subset S_\infty$ in S. Since C is mapped to east ∞, it follows that ∂C_0 consists of collapsing arcs and arcs in S which map to the α-arcs under $\pi_\Sigma \circ u_n$. Thus the limit $\pi_{\mathbb{D}} \circ U$ takes ∂C to $\{1\} \times \mathbb{R}$, so the open mapping principle tells us that $(\pi_{\mathbb{D}} \circ U)|_C$ is constant.

Otherwise, the map $(\pi_{\mathbb{D}} \circ U)|_C$ is unstable and thus has not yet appeared in the limit. The only unstable components at east infinity are constant in $\mathbb{D}$, so it suffices to figure out which constant this should be. To do so, we simply add marked points so that C appears in the limit. This tells us which constant $(\pi_{\mathbb{D}} \circ U)|_C$ should be.

This defines U completely. □

3.9 Codimension-one degenerations

The compactness theorem in the previous section ensures that $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; P)$ contains only a few kinds of degenerations. We are allowed to have all the degenerations which were permitted in the closed case (see Section 3.5). The bordered case, however, gives rise to one other degeneration, occurring at east infinity. In particular, if the source degenerates to a nodal source (i.e., to a point on the boundary of the Deligne–Mumford moduli space of Riemann surfaces), then the holomorphic map may converge to the puncture p in Σ on one side of the node. This corresponds to a level splitting. However, instead of this splitting occurring at $\pm\infty$, as in the case where the $\mathbb{R}$-coordinate of the map converges to $\pm\infty$, this splitting occurs at east infinity instead.

To summarize, then, we have the following types of degeneration: (1) becoming nodal; (2) level splitting at either $\pm\infty$ or $e\infty$; (3) level splitting with an unstable source when the derivative blows up at a puncture; and (4) bubbling of a holomorphic sphere or disk.

Note that Case (4) does not occur. The proof is similar to the argument that bubbling cannot occur in Lemma 3.22. In particular, neither $\Sigma \times [0,1] \times \mathbb{R}$ nor $\mathbb{R} \times Z \times [0,1] \times \mathbb{R}$ have holomorphic spheres, as their π_2's both vanish. Similarly, since $\pi_2(\Sigma, \boldsymbol{\alpha})$ and $\pi_2(\Sigma, \boldsymbol{\beta})$ both vanish, so too do the relative homotopy groups with respect to the Lagrangians $C_\alpha \cup C_\beta$. Thus there are no bubbled-off disks.

Proposition 3.52. *Define the boundary of* $\widetilde{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P})$ *to be*

$$\partial\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}) := \overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}) \setminus \mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^\triangleright; \vec{P}).$$

If $(\mathbf{x}, \vec{\boldsymbol{\rho}})$ *is strongly boundary monotone and* $\operatorname{ind}(B, \vec{\boldsymbol{\rho}}) \leq 2$*, then for generic* J*, every holomorphic comb in this boundary may be written in one of the following forms:*

(1) a toothless height-2 holomorphic comb (u_1, u_2)*;*

(2) a simple holomorphic comb (u, v) *where* v *is a join curve;*

(3) a simple holomorphic comb (u, v) *where* v *is a shuffle curve; or*

(4) a height-1 holomorphic comb $(u, v_1, \dots, v_k)$ *such that each* v_i *is a split curve and the preglued surface of the* v_i*'s is also a split curve.*

Degenerations of the first form are called **two-story ends**. Formally, they are elements of

$$\mathcal{M}^{B_1}(\mathbf{x}, \mathbf{w}; S_1^{\triangleright}; \vec{P}_1) \times \mathcal{M}^{B_2}(\mathbf{w}, \mathbf{y}; S_2^{\triangleright}; \vec{P}_2),$$

where $B_1 * B_2 = B$ and $S^{\triangleright} = S_1^{\triangleright} \natural S_2^{\triangleright}$ is a splitting of $S^{\triangleright}$ which divides the ordered partition $\vec{P}$ into two parts $\vec{P}_1 < \vec{P}_2$. (Thus all of the degeneration to east infinity which occurs at the first level also occurs earlier, i.e., at a smaller time t, than the degeneration to east infinity which occurs at the second level.)

Those of the second form are **join curve ends**. Such ends are elements of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright\prime}; \vec{P}')$, where $S^{\triangleright\prime}$ and $\vec{P}'$ are obtained as follows: First, pick some east puncture q from the i-th part of $\vec{P}$ which is labeled by the Reeb chord ρ_q. (Here i is some fixed number.) Decompose this chord as $\rho_q = \rho_a \uplus \rho_b$. Then let $(S^{\triangleright})'$ be a decorated source with east punctures a and b, labeled by ρ_a and ρ_b, respectively, such that we may recover $S^{\triangleright}$ by pregluing a join component to the punctures a and b. (Thus the join component has west punctures labeled by ρ_a and ρ_b, and an east puncture labeled by ρ_q.) The partition $\vec{P}'$ is obtained by replacing q with $\{a, b\}$ in the i-th part of $\vec{P}$. In this case, we say that the join curve end occurs **at level** i.

We call degenerations which take the third form **shuffle curve ends**. Similarly to the join curve ends, these are elements of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright\prime}; \vec{P}')$ where $S^{\triangleright}$ may be recovered by pregluing a shuffle curve to $S^{\triangleright\prime}$ and $\vec{P}'$ is obtained by replacing two punctures from its i-th part with a different two punctures which belong to $S^{\triangleright\prime}$. We say that the shuffle curve occurs **at level** i. We may distinguish between odd and even shuffle curve ends, depending on whether the degenerated shuffle curve is odd or even.

Finally, the degenerations which are formed by degenerating split curves are a special case of **collisions of levels**. A collision of levels i and $i+1$ is an element of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S^{\triangleright\prime}; \vec{P}')$ where $\vec{P}' = (P_1, \dots, P_i \uplus P_{i+1}, \dots, P_n)$. The decorated source $S^{\triangleright\prime}$ is obtained as follows: Contract arcs on $\partial S^{\triangleright}$ which connect punctures in P_i and P_{i+1} which are labeled by abutting Reeb chords. Replace these pairs of Reeb chords by their join, and let $P_i \uplus P_{i+1}$ comprise these joins. In general, we may have collisions of levels $i, \dots, i+j$. These correspond to degenerating several split curves.

We do not prove Proposition 3.52, and instead refer the reader to Proposition 5.43 and Section 5.6.3 of [LOT18]. The former restricts our degenerations to the four types highlighted above, as well as some potential nodal degenerations. The latter proves that in the strongly boundary monotone case, there are no nodal degenerations at all.

Unfortunately, we do not have the necessary transversality or gluing results to conclude that the moduli space $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^{\triangleright}; \vec{P})$ is a genuine manifold. This is owing to difficulties when we degenerate shuffle curves or split curves. There is a one-dimensional moduli space of shuffle curves based on where the branch point occurs, as in Figure 3.17, so shuffle curve ends are not isolated. If $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^{\triangleright}; \vec{P})$ were a manifold with $\operatorname{ind}(B, \vec{\boldsymbol{\rho}}) = 2$, then its ends, which would be ends of a 1-dimensional manifold, would be isolated.

Though we do not have the result that $\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^{\triangleright}; \vec{P})$ is a compact 1-dimensional manifold, we do have the following result, which is sufficient for defining our bordered invariants.

Theorem 3.53. *When* $\operatorname{ind}(B, \vec{\boldsymbol{\rho}}) = 2$*, the total number of two-story ends, join curve ends, shuffle curve ends, and collisions of levels is even. (In fact, there are an even number of even shuffle curve ends, so the total number of two-story ends, join curve ends, odd shuffle curve ends, and collisions of levels is zero.) In particular, the boundary* $\partial\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; S^{\triangleright}; \vec{P})$ *has an even number of points.*

Finally, it is useful to highlight the index-one case as well. The proof is very similar to the closed case.

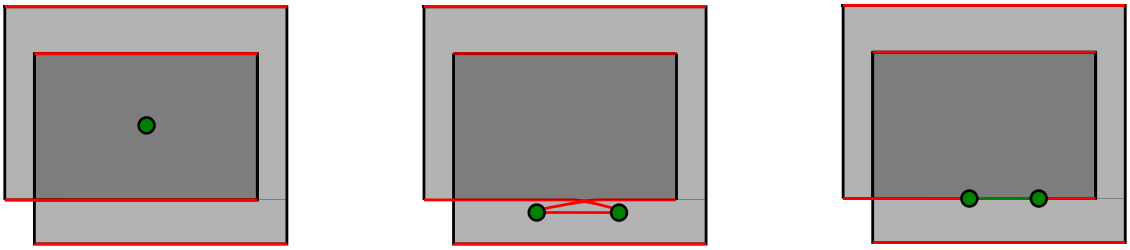

Figure 3.17: The dots in the diagram above depict branch points. By moving the branch points, we obtain a one-dimensional moduli space of shuffle curves. When the branch point is on the boundary, there is a branch cut (the green segment) between the two branch points. We can also approach the other red boundary component.

Proposition 3.54. *With a generic choice of almost complex structure, the moduli space* $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{P})$ *is a compact 0-dimensional manifold if* $\mathrm{ind}(B, \vec{\boldsymbol{\rho}}) = 1$.

3.9.1 Examples of degenerations

Example 3.55. Consider Figure 3.18 below. Here, the black arc is part of the boundary $\partial\overline{\Sigma}$, the red lines

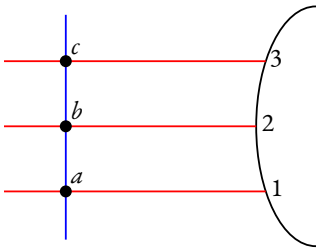

Figure 3.18: A small neighborhood of a Heegaard diagram, where red lines denote α-arcs and blue lines denote β-curves, as usual. (Note that, in principle, some of these α-arcs could be the same.)

are parts of α-arcs, and the blue line is part of some β-circle. We draw our basepoint z somewhere on the part of $\partial\overline{\Sigma}$ which is not depicted, i.e., between c and a.

There is a 1-parameter family connecting the generator $\{a\}$ to the generator $\{c\}$ with asymptotics at east infinity given by $(\{\rho_{12}\}, \{\rho_{23}\})$. This family may be seen in Figure 3.19. Thus east punctures must

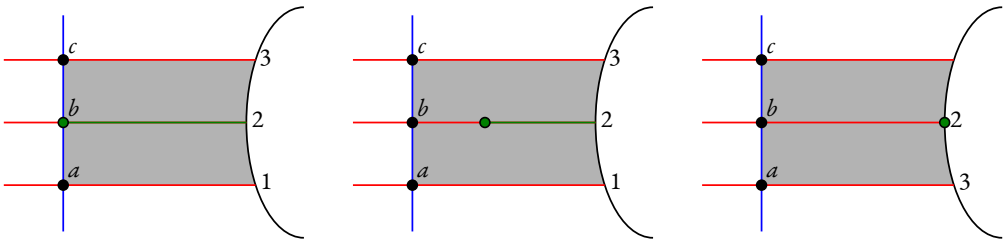

Figure 3.19: A 1-parameter family in $\mathcal{M}^B(\{a\}, \{c\}; \{\rho_{12}\}, \{\rho_{23}\})$. Here B is the shaded region, i.e., the sum of the rectangles $B_1 = 12ba$ and $B_2 = 23cb$. We have drawn the projection $\pi_\Sigma \circ u$, where green dots and lines denote branch points and cuts.

approach ρ_{12} before ρ_{23}.

This moduli space has index two, and is parameterized by the branch point. As such, one end occurs when the branch point approaches b. This corresponds to a curve which maps to the two rectangles B_1 and B_2 at increasingly distant times. In other words, we have a holomorphic building at this end, as on the left side of Figure 3.19. This is an element of

$$\mathcal{M}^{B_1}(\{a\}, \{b\}; (\rho_{12})) \times \mathcal{M}^{B_2}(\{b\}, \{c\}; (\rho_{23})),$$

and is thus a two-story end.

On the other hand, if the branch point approaches the point labeled 2 on $\partial\overline{\Sigma}$, then we obtain another end. This end corresponds to degenerating a split curve whose west puncture is associated to the Reeb chord ρ_{13}. In particular, we see the split curve in Figure 3.20. In our previous language, this is a collision

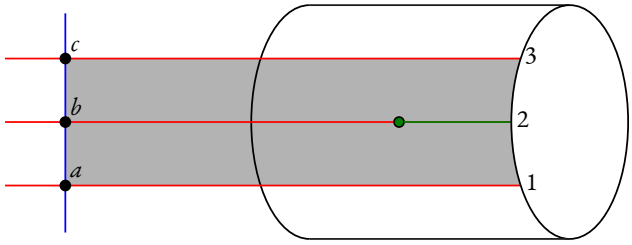

Figure 3.20: Letting the branch point escape to east infinity degenerates a split component.

of levels 1 and 2, and is thus an element of $\mathcal{M}^B(\{a\}, \{c\}; (\rho_{13}))$.

Thus we have exactly two ends in this case, which makes sense in light of Theorem 3.53.

Example 3.56. Now consider Figure 3.21. This has four generators, namely $\{a, c\}$, $\{a, d\}$, $\{b, c\}$, and

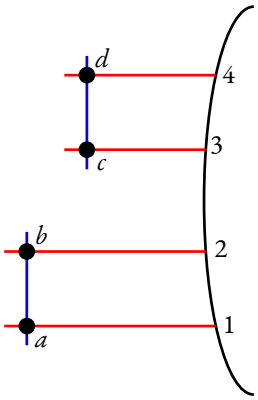

Figure 3.21: In this local picture of a Heegaard diagram, let B_1 be the lower rectangle $12ba$ and B_2 the upper rectangle $34dc$.

$\{b, d\}$. There is a moduli space of curves from $\mathbf{x} = \{a, c\}$ to $\mathbf{y} = \{b, d\}$ with asymptotics at east infinity given by the sequence $(\{\rho_{12}\}, \{\rho_{34}\})$. This moduli space is index 2 and is parameterized by $\mathrm{ev}_{34} - \mathrm{ev}_{12}$, i.e., by the difference in evaluations between the two east punctures.

From this description, we see that this moduli space has two ends. One is a two-story end, which occurs when we encounter the Reeb chord ρ_{34} infinitely far away from the Reeb chord ρ_{12}. It belongs

to the product $\mathcal{M}^{B_1}(\mathbf{x}, \{b, c\}; \{\rho_{12}\}) \times \mathcal{M}^{B_2}(\{b, c\}, \mathbf{y}; \{\rho_{34}\})$. The other is a collision of levels, which occurs when $\mathrm{ev}_{34} = \mathrm{ev}_{12}$. This is an example of a collision of levels which does not degenerate a split curve. Instead, this end is a genuine holomorphic curve, but simply belongs to another moduli space, namely $\mathcal{M}^{B_1 * B_2}(\mathbf{x}, \mathbf{y}; \{\rho_{12}, \rho_{34}\})$ instead of $\mathcal{M}^{B_1 * B_2}(\mathbf{x}, \mathbf{y}; \{\rho_{12}\}, \{\rho_{34}\})$.

Example 3.57. The generators of Figure 3.22 are $\{a, d\}$, $\{a, e\}$, $\{b, c\}$, and $\{b, e\}$. Consider the moduli

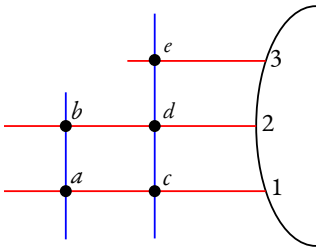

Figure 3.22: A third local picture.

space connecting generators $\{a, d\}$ and $\{b, e\}$ with the Reeb chord ρ_{13}. This is an index-two moduli

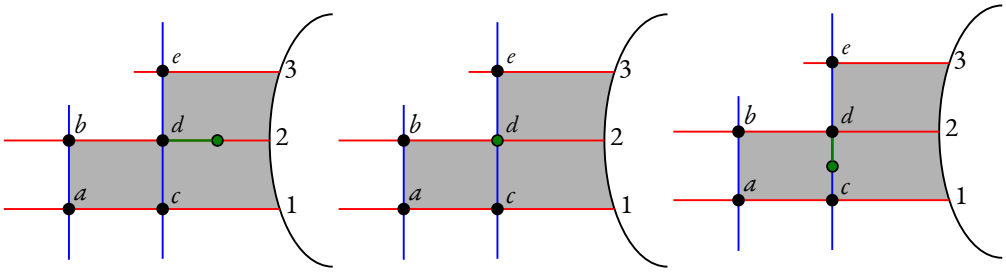

Figure 3.23: A 1-parameter family of holomorphic maps in the moduli space. The green point is the branch point, while the green line denotes a branch cut.

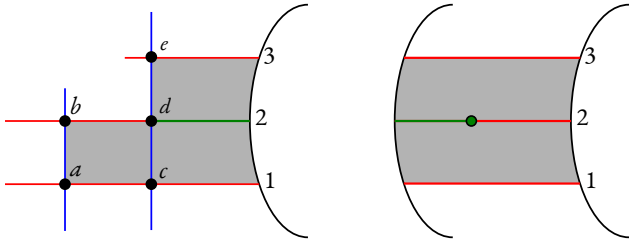

Figure 3.24: Degeneration occurs when the branch point approaches either 2 or c. When the branch point approaches 2, we degenerate off a join curve.

space. There is a 1-parameter family of holomorphic maps, as shown in Figure 3.23, which belong to this moduli space. As the branch point approaches c, however, we degenerate a two-story holomorphic building with one level going from $\{a, d\}$ to $\{b, c\}$ with no east punctures, and one level going from $\{b, c\}$ to $\{b, e\}$ with east puncture labeled by the Reeb chord ρ_{13}. As the branch point approaches the point 2 on $\partial\overline{\Sigma}$, the curve degenerates a join curve with west punctures labeled by ρ_{12} and ρ_{23} and east puncture labeled by ρ_{13}. See Figure 3.24.

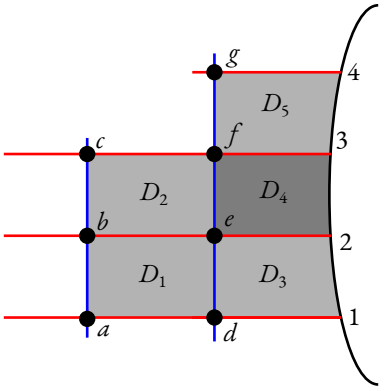

Figure 3.25: A more complicated example which shows how a shuffle curve may degenerate. Note that D_4 is shaded more darkly, to denote that B has two copies of D_4.

Example 3.58. We now give an example in which a shuffle curve may degenerate. In Figure 3.25, we show the projection $\pi_\Sigma \circ u$ of some $u \in \mathcal{M}^B(\{a,e\},\{c,g\};\{\rho_{23},\rho_{14}\})$, where $B = D_1 + D_2 + D_3 + 2D_4 + D_5$. Note that the partition dictates that we approach ρ_{23} and ρ_{14} at the same time.

In Figure 3.26, we show one possible degeneration, which occurs when the interior branch point approaches the boundary. In particular, on the left, we have a "typical" element of the moduli space.

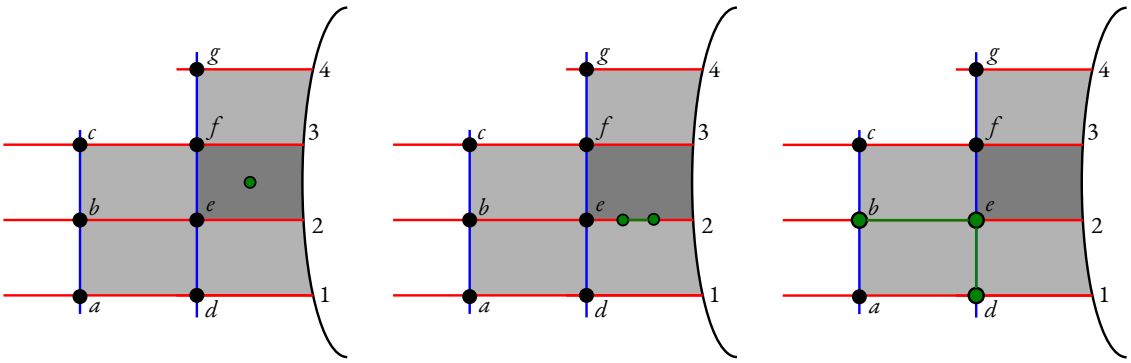

Figure 3.26: Elements of the moduli space $\mathcal{M}^B(\{a,e\},\{c,g\};\{\rho_{23},\rho_{14}\})$, where B is as indicated by the shading.

As the branch point nears the segment between e and 2, it forms a branch cut with two boundary branch points (cf. Figure 3.17). We may move these branch points, and the corresponding branch cuts, such that they pass through the concave corner between b, e, and d, as in Figure 3.27.

This results in a degeneration into a two-story holomorphic building. In particular, we have one story belonging to $\mathcal{M}^{B_1}(\{a,e\},\{b,d\};\emptyset)$, where $B_1 = D_1$ and we have no east punctures. The second story belongs to $\mathcal{M}^{B_2}(\{b,d\},\{c,g\};\{\rho_{14},\rho_{23}\})$, where $B_2 = D_2 + D_3 + 2D_4 + D_5$.

There is another type of degeneration, which occurs when an interior branch point (or two boundary branch points) approach $\partial\overline{\Sigma}$. This degenerates an odd shuffle curve, as shown in Figure 3.28. In particular, this is an element of $\mathcal{M}^B(\{a,e\},\{c,g\};\{\rho_{13},\rho_{24}\})$.

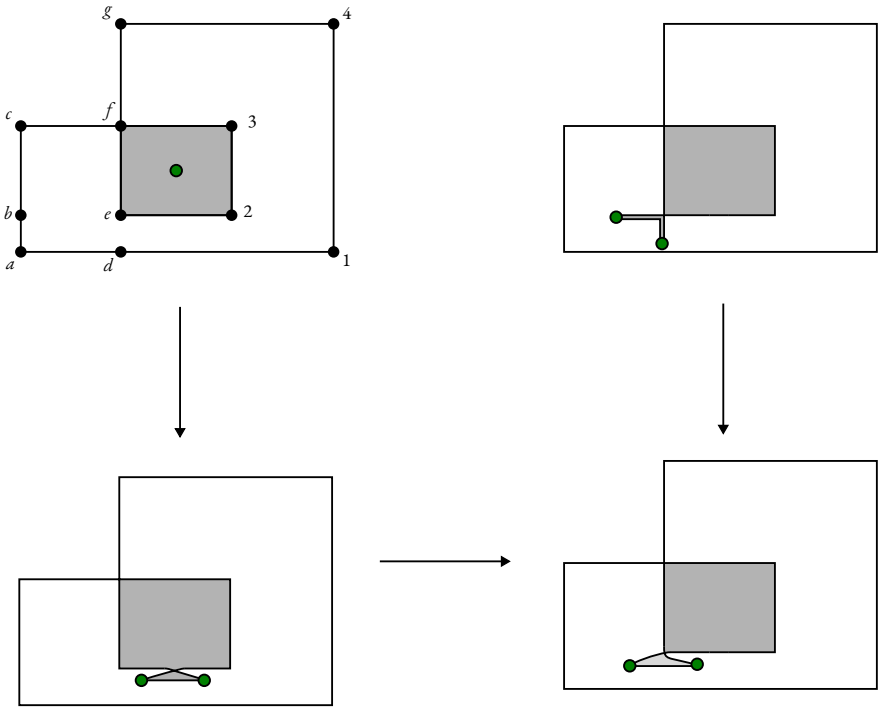

Figure 3.27: A schematic depicting how the branch cut may be made to look like the right side of Figure 3.25. We have drawn slits of nonzero width, as opposed to simply indicating cuts.

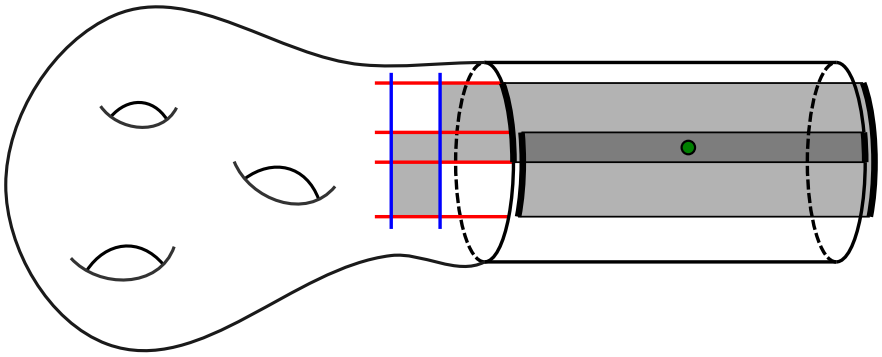

Figure 3.28: A shuffle curve degenerates in the cylinder $\mathbb{R} \times Z \times [0,1] \times \mathbb{R}$.

Chapter 4

Bordered Heegaard Floer homology

Like an old stone wall that'll never fall
Some things are always true
Some things never change

"Some Things Never Change" from *Frozen 2*

With all the geometric details about moduli spaces out of the way, we can now—finally!—define bordered Heegaard Floer homology. There are two modules associated to the bordered Heegaard Floer package, namely the "type A module" $\widehat{CFA}(Y)$ and the "type D module" $\widehat{CFD}(Y)$. The former is an object known as an "$\mathcal{A}_\infty$ module," while the latter has an associated "type D structure." Both bordered Heegaard Floer modules are defined, in some sense, by counting curves in the moduli spaces $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{P})$ defined in the previous chapter (cf. Section 2.5, which informally discusses this idea in the closed case). Exactly which curves they count, however, differs, and corresponds to their differing algebraic properties.

This chapter will be mostly concerned with defining the relevant algebraic notions, and showing how they correspond to the geometric objects (i.e., the moduli spaces) we have seen before. We begin in Section 4.1 with a somewhat preparatory section in which we discuss $\mathcal{A}_\infty$ structures. These structures will form the algebraic framework for our type A module, which is an $\mathcal{A}_\infty$ module over a certain differential algebra. In Section 4.2, we define this differential algebra. This algebra, denoted $\mathcal{A}(\mathcal{Z})$, is associated to the pointed matched circle $\mathcal{Z} = \partial\mathcal{H}$ associated to a bordered Heegaard diagram. With these ideas in place, we turn in Section 4.3 to the definition of the type A module $\widehat{CFA}(\mathcal{H})$ associated to a bordered Heegaard diagram. It turns out that this module is a 3-manifold invariant.

In Section 4.4, we define type D structures. These are a somewhat more novel algebraic notion than $\mathcal{A}_\infty$ structures, and were in fact defined expressly for bordered Heegaard Floer theory. We define $\widehat{CFD}(\mathcal{H})$, which is also an invariant of the bordered manifold represented by $\mathcal{H}$, in Section 4.5. The type D module will be a *left* $\mathcal{A}(\mathcal{Z})$ module, equipped with a type D structure. We conclude in Section 4.6 with a discussion of the pairing theorem. This combines Heegaard Floer homology with both bordered Heegaard Floer objects. In particular, it says that a suitable pairing of $\widehat{CFA}(Y_1)$ with $\widehat{CFD}(Y_2)$, where $\partial Y_1 = \partial Y_2$, recovers the Heegaard Floer homology $\widehat{HF}(Y)$ of their union $Y_1 \cup_\partial Y_2$.

4.1 $\mathcal{A}_\infty$ algebras and modules

The bordered invariant $\widehat{CFA}(Y) = \widehat{CFA}(\mathcal{H})$ is an $\mathcal{A}_\infty$ module over a differential graded (dg) algebra $\mathcal{A}(\mathcal{Z})$ associated to the pointed matched circle $\mathcal{Z} = \partial\mathcal{H}$. As such, before defining $\widehat{CFA}$, we will introduce in this section some details about $\mathcal{A}_\infty$ structures, which were first introduced in 1963 by John Stasheff [Sta63a, Sta63b]. Since then, $\mathcal{A}_\infty$ algebras, categories, and tensor products have all become part of the typical algebraic toolbox in symplectic geometry [Fuk93, Sei08, FOO+09]. A more detailed exposition of this algebraic background may be found in [Kel01].

Let $\mathbf{k}$ denote some fixed commutative ring of characteristic two. (Usually, this will be some direct sum of copies of $\mathbb{F}_2 = \mathbb{Z}/2\mathbb{Z}$.)

Sometimes, an associative $\mathbf{k}$-algebra comes as the (co)homology of a chain complex (A, ∂) that one expects to have the structure of a differential algebra, i.e., to admit some multiplication $\mu : A \otimes A \to A$ which is associative and which satisfies the Leibniz rule. (Unless otherwise specified, all tensor products are over $\mathbf{k}$.) But it sometimes occurs that μ is only associative up to a homotopy, i.e., there exists a map μ_3 such that when $a_1, a_2, a_3 \in A$ are closed (i.e., $\partial a_i = 0$), we have

$$\mu(\mu(a_1 \otimes a_2) \otimes a_3) + \mu(a_1 \otimes \mu(a_2 \otimes a_3)) = \partial\mu_3(a_1 \otimes a_2 \otimes a_3).$$

If μ_3 now only satisfies associativity up to homotopy, then we get further maps μ_4, and so on.

Definition 4.1. An $\mathcal{A}_\infty$ **algebra over** $\mathbf{k}$ is a $\mathbf{k}$-module A, along with multiplication maps $\mu_i : A^{\otimes i} \to A$ for $i \geq 1$. The multiplication maps satisfy the following compatibility condition:

$$\sum_{i+j=n+1} \sum_{\ell=1}^{n-j+1} \mu_i(a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_j(a_\ell \otimes \cdots \otimes a_{\ell+j-1}) \otimes a_{\ell+j} \otimes \cdots \otimes a_n) = 0.$$

We denote the $\mathcal{A}_\infty$ algebra by $\mathcal{A}$, and call its underlying $\mathbf{k}$-module A.

Note that an $\mathcal{A}_\infty$ algebra comes with a map $\mu_1 : A \to A$ as well. This is the same as the map ∂ in the case that A comes from a chain complex (A, ∂). When working over general $\mathbf{k}$, we actually require that A is a *graded* $\mathbf{k}$-module. The compatibility condition then comes with certain signs. In this context, μ_1 is a differential, so that an $\mathcal{A}_\infty$ algebra with $\mu_i = 0$ for $i \geq 2$ is simply a chain complex. Similarly, if our only nontrivial multiplication maps are μ_1 and μ_2, then $\mathcal{A}$ is a differential graded algebra.

It turns out that the algebraic structures defined below in the bordered Heegaard Floer package may all be equipped with a grading by a certain noncommutative group. We will not spend time defining this grading. As such, we define our $\mathcal{A}_\infty$ objects without grading, noting simply that this is okay because $\mathbf{k}$ has characteristic two.

We say that $\mathcal{A}$ is **strictly unital** if there is an element $1 \in \mathcal{A}$ such that $\mu_2(a, 1) = \mu_2(1, a) = a$ and $\mu_i(a_1 \otimes \cdots \otimes a_i) = 0$ if $i \neq 2$ and $a_j = 1$ for some j. We say that it is **operationally bounded** if $\mu_i = 0$ for all but finitely many i.

There is a graphical representation of the compatability condition in Definition 4.1. In particular, think of $A^{\otimes i}$ as being denoted by i parallel, downward-oriented strands. Let μ_i be represented by the i strands of $A^{\otimes i}$ meeting at a "vertex," which we label by μ_i, and exiting as one strand below. For example,

μ_3 may be denoted by

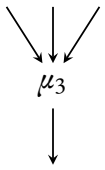

The compatibility condition on $A^{\otimes i}$ is obtained by summing over all ways to use two μ_j's to get from i parallel strands to one single strand. For example, the compatibility condition for μ_3 says that

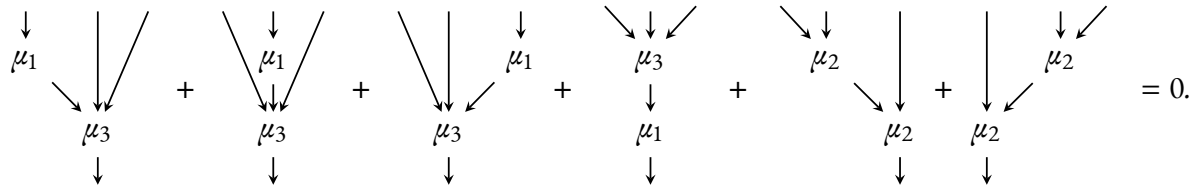

The last two terms are exactly the usual associativity terms, while the first four are the homotopy up to which μ_2 is associative.

A more concise way to draw these diagrams is to combine all the multiplication maps into a single map

$$\mu : \mathcal{T}^*(A) \coloneqq \bigoplus_{n=0}^{\infty} A^{\otimes n} \to A$$

on the tensor algebra. By convention, $\mu_0 = 0$. There is now an endomorphism $\overline{D}$ on $\mathcal{T}^*(A)$ defined by

$$\overline{D}(a_1 \otimes \cdots \otimes a_n) = \sum_{j=1}^{n} \sum_{\ell=1}^{n-j+1} a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_j(a_\ell \otimes \cdots \otimes a_{\ell+j-1}) \otimes \cdots \otimes a_n. \tag{4.1}$$

Then the compatibility condition for $\mathcal{A}_\infty$ algebras may be written as $\mu \circ \overline{D} = 0$ or, equivalently, as $\overline{D} \circ \overline{D} = 0$. Graphically, we use doubled arrows to indicate elements of $\overline{D}$ and single arrows, as before, to indicate elements of A. Thus we may depict this relation as

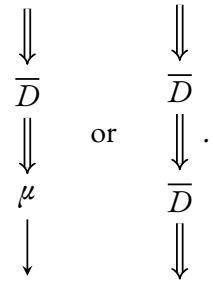

We have discussed $\mathcal{A}_\infty$ algebras, but in fact the invariant $\widehat{CFA}$ is an $\mathcal{A}_\infty$ module over some algebra. (In fact, it is a module over a differential graded algebra. In particular, we do not need all the details of $\mathcal{A}_\infty$ algebras.) We define this next.

Definition 4.2. A **(right)** $\mathcal{A}_\infty$ **module** $\mathcal{M}$ over an $\mathcal{A}_\infty$ algebra $\mathcal{A}$ is a $\mathbf{k}$-module M with operations

$$m_i : M \otimes A^{\otimes(i-1)} \to M$$

for all $i \geq 1$. We ask that these maps satisfy the following compatibility condition:

$$\begin{aligned} 0 = &\sum_{i+j=n+1} m_i(m_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}) \\ &+ \sum_{i+j=n+1} \sum_{\ell=1}^{n-j} m_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_j(a_\ell \otimes \cdots \otimes a_{\ell+j-1}) \otimes a_{\ell+j-1} \otimes \cdots \otimes a_{n-1}). \end{aligned}$$

If $\mathcal{M}$ is an $\mathcal{A}_\infty$ module over a strictly unital $\mathcal{A}_\infty$ algebra $\mathcal{A}$, then we call $\mathcal{M}$ **strictly unital** if, for every $\mathbf{x} \in \mathcal{M}$, we have $m_2(\mathbf{x} \otimes 1) = \mathbf{x}$ and $m_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{i-1}) = 0$ if $i \neq 2$ and some $a_j = 1$. Furthermore, we say that $\mathcal{M}$ is **bounded** if $m_i = 0$ for all but finitely many i.

We may draw a similar graphical representation for $M \times A^{\otimes(i-1)}$ as for $A^{\otimes i}$. The only difference is that we now distinguish the leftmost strand, which represents our factor of M in the module, by drawing it as a dashed line instead. The output of m_i is an element of M, hence is also "colored" by M, i.e., drawn as a dashed line. For example, m_3 may be drawn as

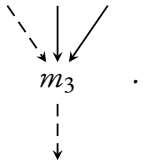

In our context, where $\mu_i = 0$ for all $i > 2$, the module compatibility condition may be written as

$$\begin{aligned} 0 = &\sum_{i+j=n+1} m_i(m_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}) \\ &+ \sum_{\ell=1}^{n-1} m_n(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_1(a_\ell) \otimes a_{\ell+1} \otimes \cdots \otimes a_{n-1}) \\ &+ \sum_{\ell=1}^{n-2} m_{n-1}(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_2(a_\ell \otimes a_{\ell+1}) \otimes a_{\ell+2} \otimes \cdots \otimes a_{n-1}). \end{aligned} \tag{4.2}$$

For m_3, this is drawn as

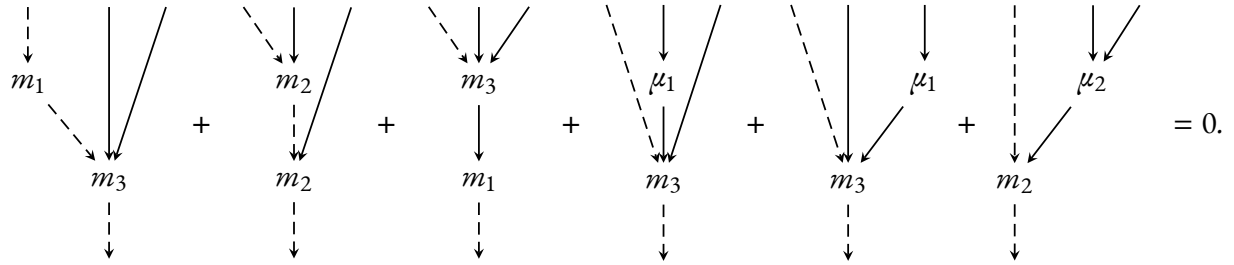

We may again write this in terms of the tensor algebra. Let $\Delta : \mathcal{T}^*(A) \to \mathcal{T}^*(A) \otimes \mathcal{T}^*(A)$ be the comultiplication map

$$\Delta(a_1 \otimes \cdots \otimes a_n) = \sum_{m=0}^{n} (a_1 \otimes \cdots \otimes a_m) \otimes (a_{m+1} \otimes \cdots \otimes a_n).$$

Then we may draw the compatibility condition for $\mathcal{A}_\infty$ modules as follows:

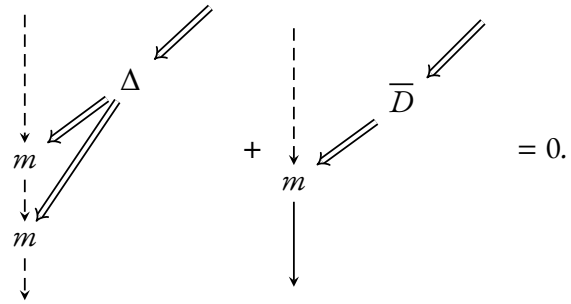

It will turn out that $\widehat{CFA}(Y)$ is defined up to homotopy equivalence. To understand the $\mathcal{A}_\infty$ version of (chain) homotopy equivalence, we must first understand what a homomorphism of $\mathcal{A}_\infty$ modules looks like.

Definition 4.3. A **(strictly unital)** $\mathcal{A}_\infty$ **homomorphism** between the strictly unital right $\mathcal{A}_\infty$ modules $\mathcal{M}$ and $\widetilde{N}$ over a strictly unital $\mathcal{A}_\infty$ algebra $\mathcal{A}$ is a collection of maps $f_i : M \otimes A^{i-1} \to N$ for $i \geq 1$ which satisfies the compatibility condition

$$\begin{aligned} 0 = \sum_{i+j=n+1} & n_i(f_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}) \\ & + \sum_{i+j=n+1} f_i(m_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}) \\ & + \sum_{i+j=n+1} \sum_{\ell=1}^{n-j} f_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_j(a_\ell \otimes \cdots \otimes a_{\ell+j-1}) \otimes a_{\ell+j} \otimes \cdots \otimes a_{n-1}), \end{aligned}$$

where n_i are the multiplication maps of $\widetilde{N}$, as well as the unital conditions $f_1(1_M) = 1_N$ and

$$f_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{i-1}) = 0$$

if $i > 1$ and some $a_j = 1_A$.

Graphically, this compatibility condition may be drawn as

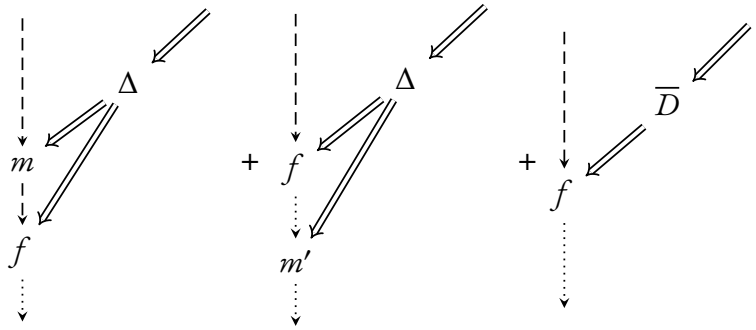

Here, we use dashed lines to represent elements of $\mathcal{M}$. Dotted lines represent elements of $\mathcal{M}'$.

Consider, for example, the identity homomorphism $\mathbb{I}_{\mathcal{M}}$ of a strictly unital $\mathcal{A}_\infty$ module $\mathcal{M}$. It is defined by setting $(\mathbb{I}_{\mathcal{M}})_1(\mathbf{x}) := \mathbf{x}$ and $(\mathbb{I}_{\mathcal{M}})_i(\mathbf{x} \otimes A^{\otimes(i-1)}) = 0$ for all $i > 1$. This is an $\mathcal{A}_\infty$ homomorphism.

The **composite** of $\mathcal{A}_\infty$ homomorphisms $f : \mathcal{M} \to \mathcal{N}$ and $g : \mathcal{N} \to \mathcal{P}$ is given by

$$(g \circ f)_n(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{n-1}) = \sum_{i+j=n+1} g_i(f_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}).$$

A strictly unital homomorphism is **bounded** if $f_i = 0$ for all but finitely many i.

Recall that, in the typical (i.e., not $\mathcal{A}_\infty$) setting, a chain homotopy between f and g is a family of maps h_i such that $f_i - g_i = dh_{i-1} + h_{i+1}d$. (The differential d is similar to the map μ_1.) The $\mathcal{A}_\infty$ version of this should be this equality, up to a certain homotopy involving higher multiplication maps. Indeed, we may make the following definition.

Definition 4.4. Let $\mathcal{M}$ and $\mathcal{M}'$ be strictly unital $\mathcal{A}_\infty$ modules over the $\mathcal{A}_\infty$ algebra $\mathcal{A}$. Consider a collection of maps

$$h_i : M \otimes A^{\otimes(i-1)} \to M'$$

with $h_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{i-1}) = 0$ whenever $i > 1$ and $a_j = 1$ for some j. Then define f_n by

$$\begin{aligned} f_n(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{n-1}) = &\sum_{i+j=n+1} h_i(m_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}) \\ &+ \sum_{i+j=n+1} m'_i(h_j(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes a_j \otimes \cdots \otimes a_{n-1}) \\ &+ \sum_{i+j=n+1} \sum_{\ell=1}^{n-j} h_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{\ell-1} \otimes \mu_j(a_\ell \otimes \cdots \otimes a_{\ell+j-1}) \otimes \cdots \otimes a_{n-1}). \end{aligned}$$

This map f is an $\mathcal{A}_\infty$ homomorphism, and we call it **nullhomotopic**. Two maps $f, g : \mathcal{M} \to \mathcal{M}'$ are **homotopic** if $f - g$ is nullhomotopic. Finally, if there are maps $f : \mathcal{M} \to \mathcal{M}'$ and $g : \mathcal{M}' \to \mathcal{M}$ such that $g \circ f$ and $f \circ g$ are homotopic to the identity homomorphisms $\mathbb{I}_{\mathcal{M}}$ and $\mathbb{I}_{\mathcal{M}'}$, respectively.

This definition is somewhat convoluted, but is related to the typical definition of chain homotopy as follows: One may upgrade the maps $m = \{m_i\}$ to an endomorphism $\overline{m}$ of $M \otimes \mathcal{T}^*(A)$. This definition is analogous to the definition of $\overline{D}$ in the $\mathcal{A}_\infty$ algebra case. Similarly, we may promote the homomorphisms f and g, as well as the homotopy h, to maps $\overline{f}, \overline{g}, \overline{h} : M \otimes \mathcal{T}^*(A) \to M \otimes \mathcal{T}^*(A)$. Then the condition in Definition 4.4 may be written as

$$\overline{h} \circ \overline{m} + \overline{m'} \circ \overline{h} = \overline{f} + \overline{g}.$$

This looks exactly like the condition for h to be a chain homotopy between chain maps f and g.

The condition that $f, g : \mathcal{M} \to \mathcal{M}'$ are homotopic is represented graphically as follows:

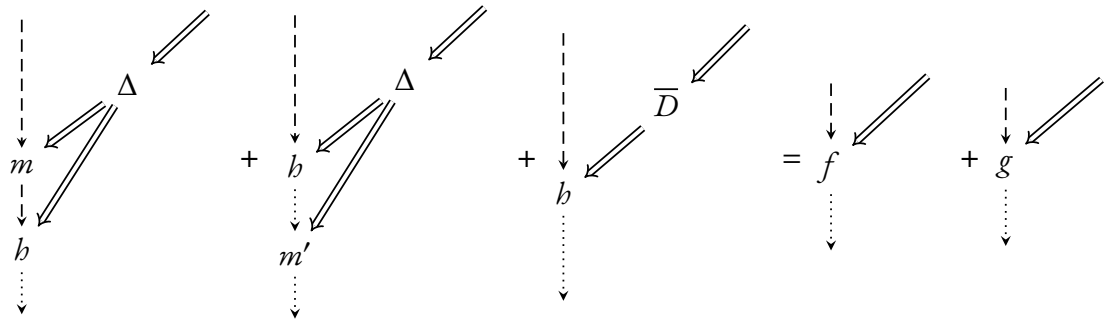

We conclude with the definition of the $\mathcal{A}_\infty$ tensor product, which will be useful for stating the pairing theorem (Theorem 4.44).

Definition 4.5. Let $\mathcal{A}$ be an $\mathcal{A}_\infty$ algebra over $\mathbf{k}$, $\mathcal{M}$ a right $\mathcal{A}_\infty$ module over $\mathcal{A}$, and $\mathcal{N}$ a left $\mathcal{A}_\infty$ module over $\mathcal{A}$. Then their $\mathcal{A}_\infty$ **tensor product** is the chain complex

$$\mathcal{M}\widetilde{\otimes}\mathcal{N} := M \otimes \mathcal{T}^* A \otimes N$$

with differential ∂ defined by

$$\begin{aligned}\partial(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_n \otimes \mathbf{y}) := &\sum_{i=1}^{n+1} m_i(\mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{i-1}) \otimes a_i \otimes \cdots \otimes a_n \otimes \mathbf{y} \\ &+ \sum_{i=1}^{n} \sum_{\ell=1}^{n-i+1} \mathbf{x} \otimes a_1 \otimes \cdots \otimes \mu_i(a_i \otimes \cdots \otimes a_{\ell+i-1}) \otimes \cdots \otimes a_n \otimes \mathbf{y} \\ &+ \sum_{i=1}^{n+1} \mathbf{x} \otimes a_1 \otimes \cdots \otimes a_{n-i+1} \otimes m_i(a_{n-i+2} \otimes \cdots \otimes a_n \otimes \mathbf{y}).\end{aligned}$$

4.2 The algebra associated to a pointed matched circle

In this section, we define the algebra $\mathcal{A}(\mathcal{Z})$ which is associated to a pointed matched circle $\mathcal{Z}$. In fact, this algebra is a differential graded algebra, but, as usual, we will not discuss the grading. The bordered invariants will be modules over this algebra.

To define $\mathcal{A}(\mathcal{Z})$, we must discuss the strands algebra $\mathcal{A}(n, k)$. A **strand diagram with n places and k strands** is created as follows: First, on both the left and right side, one draws n dots numbered 1 (at the bottom) to n (at the top). Then a strand diagram is obtained by drawing a set of k strands going up and to the right such that no two strands cross more than once and no two strands can have either the same start or end. Consider, for example, the diagram

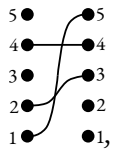

which is an element of $\mathcal{A}(5, 3)$. Note that horizontal strands are permitted, but not strands which go down and to the right.

Definition 4.6. The **strands algebra with n places and k strands**, denoted $\mathcal{A}(n, k)$, is the $\mathbb{F}_2$-vector space generated by these strand diagrams. The **strands algebra $\mathcal{A}(n)$ with n places** is the direct sum

$$\mathcal{A}(n) = \bigoplus_{k=0}^{n} \mathcal{A}(n, k).$$

That this is an algebra, and not merely a vector space, requires that we define a multiplication on the strands algebra. When the concatenation of a and b is defined (i.e., when the right side of a matches the

left side of b) and when the juxtaposition of the two diagrams has no twice-crossing strands, we declare $a \cdot b$ to be this juxtaposition. Otherwise, we declare the product to be zero. In particular, if the fragment

appears, then the product is zero.

There is a boundary operator defined on the strands algebra as follows: There is a unique way to smooth a single crossing in a strand diagram, namely by

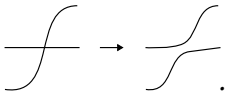

The differential of a strand diagram is then the sum of all ways to smooth a single crossing, where we declare terms with twice-crossing strands to be zero. For example, we have

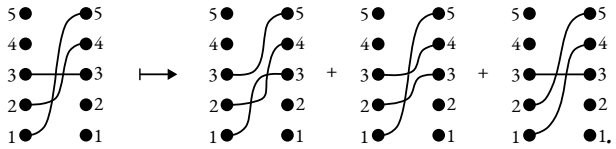

The first term on the right-hand side vanishes. Note that strand diagrams without crossings have differential zero.

Lemma 4.7. *With the multiplication and boundary operators as defined above, the set $\mathcal{A}(n, k)$, and hence also the direct sum $\mathcal{A}(n) = \bigoplus \mathcal{A}(n, k)$, is a differential algebra.*

Proof. Consider the algebra $\overline{\mathcal{A}}(n, k) \supset \mathcal{A}(n, k)$ which is the $\mathbb{F}_2$-vector space generated by strand diagrams where we do not set diagrams with twice-crossing strands to be zero. (We still do not allow isotoping away such double crossings.) Then $\overline{\mathcal{A}}(n, k)$ is a differential algebra, and the sub-vector space $\overline{\mathcal{A}}_d(n, k)$ which is generated by all terms with at least one double crossing is a differential ideal. (See Lemma 3.1 in [LOT18] for details.) Since $\mathcal{A}(n, k) = \overline{\mathcal{A}}(n, k)/\overline{\mathcal{A}}_d(n, k)$, the result follows. □

To a subset $S \subset \{1, \dots, n\}$, we may associate an idempotent $I(S)$ consisting of a horizontal strand at each $i \in S$. For example, the following is the idempotent $I(\{1, 3, 4\})$ for $n = 5$:

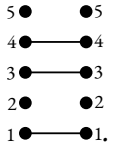

Remark 4.8. There is an equivalent way of thinking about the strands algebra. In particular, we may think of a strand diagram in $\mathcal{A}(n, k)$ as representing a partial permutation $\phi : S \to T$ between two k-element subsets of $\{1, \dots, n\}$ such that $\phi(i) \geq i$ for every $i \in S$. For instance, the element of $\mathcal{A}(5, 3)$ above may be thought of as the permutation $1 \mapsto 5$, $2 \mapsto 3$, and $4 \mapsto 4$. We may denote this by $\left\langle \begin{smallmatrix} 1 & 2 & 4 \\ 5 & 3 & 4 \end{smallmatrix} \right\rangle$. We will sometimes use this notation, as well as the notation (S, T, ϕ), for convenience. Idempotents, for example, take the form (S, S, id_S).

Now we may define the algebra associated to a pointed matched circle $\mathcal{Z} = (Z, \mathbf{a}, M, z)$. Recall from Section 2.2 that a pointed matched circle comprises a circle Z, some set of points $\mathbf{a} = \{a_1, \dots, a_{4k}\}$ on Z, a 2-to-1 matching $M : \mathbf{a} \to \{1, \dots, 2k\}$, and a basepoint $z \in Z \setminus \mathbf{a}$. Furthermore, recall from Lemma 2.8 that the boundary $\partial\mathcal{H}$ of a bordered Heegaard diagram is a pointed matched circle whose matching is determined by which α-arc an intersection point $\boldsymbol{\alpha} \cap \partial\overline{\Sigma}$ belongs to.

Fix a pointed matched circle $\mathcal{Z}$. Without loss of generality, say that the basepoint z lies between a_{4k} and a_1 and Z is oriented to go from a_i to a_{i+1}, so that we may think of $Z \setminus \mathbf{a}$ as a line with $4k$ points labeled, in order, by $a_1, \dots, a_{4k}$. These will be the dots on either side of our strand diagram; strands in the diagram will then represent Reeb chords.

Recall our terminology of nested, interleaved, and abutting Reeb chords from Section 2.2. Recall also our definition of a consistent set of Reeb chords, i.e., a set $\boldsymbol{\rho} = \{\rho_1, \dots, \rho_j\}$ of Reeb chords such that the set $\boldsymbol{\rho}^- := \{\rho_1^-, \dots, \rho_j^-\}$ of initial endpoints and the set $\boldsymbol{\rho}^+ := \{\rho_1^+, \dots, \rho_j^+\}$ of final endpoints both have exactly j elements.

Definition 4.9. Let $\boldsymbol{\rho}$ be a set of Reeb chords. It may be considered as a strand diagram with strands from $\boldsymbol{\rho}^-$ to $\boldsymbol{\rho}^+$. The **strands algebra element associated to** $\boldsymbol{\rho}$, denoted $a_0(\boldsymbol{\rho}) \in \mathcal{A}(n)$ is defined to be the formal sum of all the ways to add horizontal strands to this diagram such that the result is still a strand diagram. For example, the if $\boldsymbol{\rho}$ is the one-element set consisting of the Reeb chord from a_2 to a_4, then its associated strands algebra element is

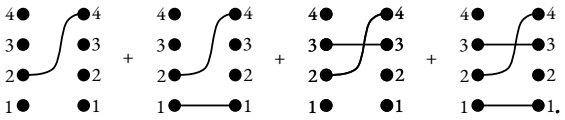

If $\boldsymbol{\rho}$ is not consistent, then $a_0(\boldsymbol{\rho}) = 0$.

These elements $a_0(\boldsymbol{\rho})$, as well as the idempotents $I(S)$ with $S \subset \mathbf{a}$, generate $\mathcal{A}(4k)$. In particular, the strands algebra $\mathcal{A}(4k)$ has basis as an $\mathbb{F}_2$-vector space given by terms of the form $I(S)a_0(\boldsymbol{\rho})$.

The algebra $\mathcal{A}(\mathcal{Z})$ associated to a pointed matched circle $\mathcal{Z}$ with $\mathbf{a} = \{a_1, \dots, a_{4k}\}$ is a subalgebra of $\mathcal{A}(4k)$. A **section** of the 2-to-1 matching M over a subset $\mathbf{s} \subset \{1, \dots, 2k\}$ is a subset of $\mathbf{a}$ which maps bijectively to $\mathbf{s}$ under M. In other words, because M associates each intersection point $a_i \in \boldsymbol{\alpha} \cap \overline{\Sigma}$ with the arc α_j^a which contains a_i, a section over a subset $\mathbf{s}$ of α-arcs consists of an endpoint of each arc in $\mathbf{s}$. The idempotent associated to $\mathbf{s}$ is the sum

$$I(\mathbf{s}) := \sum_{S \text{ is a section over } \mathbf{s}} I(S).$$

For example, if $k = 1$ and M is the matching associated to the Heegaard diagram of the genus-1 handlebody, i.e., $M(a_1) = M(a_3) = 1$ and $M(a_2) = M(a_4) = 2$, then the idempotent associated to $\{1, 2\}$ is

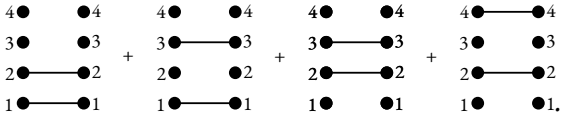

The **ring of idempotents associated to** $\mathcal{Z}$ is generated by $I(\mathbf{s})$ for all subsets $\mathbf{s}$ of $\{1, \dots, 2k\}$. This ring is denoted $\mathcal{I}(\mathcal{Z})$ and has unit

$$\mathbf{I} := \sum_{\mathbf{s} \subset \{1, \dots, 2k\}} I(\mathbf{s}).$$

Another way to draw the idempotent $I(\mathbf{s})$ is to draw the matching on the side of the strand diagram, and to draw dashed horizontal lines at each a_i with $M(i) \in \mathbf{s}$. The example above can be written as

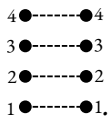

The dashed horizontal lines indicate that the chord appears in exactly half of the terms in the sum.

Definition 4.10. The **algebra associated to a pointed matched circle** $\mathcal{Z}$, denoted $\mathcal{A}(\mathcal{Z})$, is the subalgebra of $\bigoplus_{i=0}^{2k} \mathcal{A}(4k, i)$ which is generated (as an algebra) by $\mathcal{I}(\mathcal{Z})$ and $\mathbf{I}a_0(\boldsymbol{\rho})\mathbf{I}$ for every (consistent) set of Reeb chords $\boldsymbol{\rho}$. The **algebra element associated to** ρ is $\mathbf{I}a_0(\boldsymbol{\rho})\mathbf{I}$. It is the projection of $a_0(\boldsymbol{\rho})$ to $\mathcal{A}(\mathcal{Z})$, and is denoted by $a(\boldsymbol{\rho})$. Finally, define the parts of **weight** i to be

$$\mathcal{A}(\mathcal{Z}, i) \coloneqq \mathcal{A}(\mathcal{Z}) \cap \mathcal{A}(4k, k+i) \quad \text{and} \quad \mathcal{I}(\mathcal{Z}, i) \coloneqq \mathcal{I}(\mathcal{Z}) \cap \mathcal{I}(4k, k+i).$$

Thus our algebra decomposes into the parts of weight i for $-k \leq i \leq k$.

It is not too hard to show that $\mathcal{A}(\mathcal{Z})$ is closed under multiplication.

In the previous section, we considered algebras and modules over a ground ring $\mathbf{k}$. We consider $\mathcal{A}(\mathcal{Z})$ to be an algebra over the characteristic-two ring $\mathbf{k} = \mathcal{I}(\mathcal{Z})$.

The algebra $\mathcal{A}(\mathcal{Z})$ has a (vector space) basis over $\mathbb{F}_2$ given by all nonzero elements $I(\mathbf{s})a(\boldsymbol{\rho})$. Using our dashed-line strand diagrams for $I(\mathbf{s})$ from above, we may draw these elements as the strand diagram for $\boldsymbol{\rho}$, along with $2|\mathbf{s}|$ dashed horizontal lines representing $\mathbf{s}$. We may also write them in two-line notation as

$$\left[\begin{smallmatrix} x_1 & \dots & x_m & z_1 & \dots & z_\ell \\ y_1 & \dots & y_m \end{smallmatrix}\right] \coloneqq I\left(\{M(x_i), M(z_j)\}\right) a(\boldsymbol{\rho})$$

where $\boldsymbol{\rho}$ consists of the m Reeb chords which begin at x_i and end at y_i.

4.3 The type A module $\widehat{CFA}$

Fix a bordered Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ which is provincially admissible as in Definition 2.20. Let $\mathcal{Z} = \partial\mathcal{H}$ be a pointed matched circle (cf. Lemma 2.8).

Then, in this section, we will finally define our first invariant, namely the type A module $\widehat{CFA}(Y)$. We begin in Section 4.3.1 with the definition of this module. Then in Section 4.3.2, we provide many examples and eventually prove the fact that this module is in fact an $\mathcal{A}_\infty$ module over $\mathcal{A}(\mathcal{Z})$. Finally, we state an invariance result in Section 4.3.3, which tells us that we the right $\mathcal{A}_\infty$ module $\widehat{CFA}(\mathcal{H})$ over $\mathcal{A}(\mathcal{Z})$ is, in fact, a bordered 3-manifold invariant.

Remark 4.11. As a brief aside, it turns out that $\mathcal{A}(\mathcal{Z})$ is not an invariant of the surface $F(\mathcal{Z})$ which is specified by the pointed matched circle. In fact, if $\mathcal{Z}$ and $\mathcal{Z}'$ are pointed matched circles representing diffeomorphic surfaces, then in general $\mathcal{A}(\mathcal{Z})$ and $\mathcal{A}(\mathcal{Z}')$ do not have the same rank. But there is a derived equivalence between the module categories of $\mathcal{A}(\mathcal{Z})$ and $\mathcal{A}(\mathcal{Z}')$. See Theorems 1 and 9 in [LOT15].

4.3.1 Definition of the type A module

The type A invariant $\widehat{CFA}(\mathcal{H})$ is generated, as a vector space over $\mathbb{F}_2$, by the generators $\mathfrak{S}(\mathcal{H})$ of the bordered Heegaard diagram. Recall that $o(\mathbf{x}) := \{i : \mathbf{x} \cap \alpha_i^a \neq \emptyset\} \subset \{1, \dots, 2k\}$ is the set of α-arcs which are occupied by some x_i. (Since the $\mathbf{x}$'s here are generators, and not merely generalized generators, we know that each α-arc contains at most one x_i.)

For each $\mathbf{x} \in \mathfrak{S}(\mathcal{H})$, define $I_A(\mathbf{x}) := I(o(\mathbf{x})) \in \mathcal{A}(\mathcal{Z}, 0)$. Then we may define a right action of $\mathcal{I}(\mathcal{Z})$ on $\widehat{CFA}(\mathcal{H})$ as follows:

$$\mathbf{x} \cdot I(\mathbf{s}) := \begin{cases} I_A(\mathbf{x}) & \text{if } \mathbf{s} = o(\mathbf{x}), \\ 0 & \text{otherwise.} \end{cases}$$

This is only nontrivial if $\mathbf{s}$ has k elements, i.e., if $I(\mathbf{s}) \in \mathcal{A}(\mathcal{Z}, 0)$. This action extends to an action of $\mathcal{A}(\mathcal{Z})$ on $\widehat{CFA}(\mathcal{H})$ which is trivial on summands $\mathcal{A}(\mathcal{Z}, i)$ with $i \neq 0$.

We would like to define multiplication maps

$$m_{n+1} : \widehat{CFA}(\mathcal{H}) \otimes \mathcal{A}(\mathcal{Z})^n \to \widehat{CFA}(\mathcal{H}).$$

The tensor products above are all over $\mathbf{k} = \mathcal{I}(\mathcal{Z})$. It turns out to be sufficient to define m_{n+1} only on basis elements $\mathbf{x} \otimes a(\boldsymbol{\rho}_1) \otimes \cdots \otimes a(\boldsymbol{\rho}_n)$ such that, with $\vec{\boldsymbol{\rho}}$ denoting the sequence $(\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_n)$, the pair $(\mathbf{x}, \vec{\boldsymbol{\rho}})$ is strongly boundary monotone as in Definition 3.32. This is thanks to the following lemma.

Lemma 4.12 ([LOT18, Lemma 7.2])**.** *Let $\mathbf{x} \in \mathfrak{S}(\mathcal{H})$ and $\vec{\boldsymbol{\rho}} = \{\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_n\}$ be a sequence of nonempty sets of Reeb chords. Then $(\mathbf{x}, \vec{\boldsymbol{\rho}})$ is strongly boundary monotone if and only if*

$$I_A(\mathbf{x}) \otimes a(\boldsymbol{\rho}_1) \otimes \cdots \otimes a(\boldsymbol{\rho}_n) \neq 0.$$

This means that, in order to define $m_{n+1}(\mathbf{x}, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_n))$, we may use the associated source-independent moduli spaces $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ from Section 3.6. (Here, and later on, we use commas to separate the tensor factors which are inputted into m_{n+1}.)

Definition 4.13. Let J be an almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$ which is admissible, in the sense of Definition 3.26, and which achieves transversality, so that Proposition 3.29 holds. Let $\vec{\boldsymbol{\rho}} = (\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_n)$ be a sequence of nonempty sets of Reeb chords. If $(\mathbf{x}, \vec{\boldsymbol{\rho}})$ is strongly boundary monotone, then make the following definitions:

$$\begin{aligned} m_{n+1}(\mathbf{x}, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_n)) &:= \sum_{\mathbf{y} \in \mathfrak{S}(\mathcal{H})} \sum_{\substack{B \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mathrm{ind}(B, \vec{\boldsymbol{\rho}}) = 1}} \#\left(\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})\right) \mathbf{y} \\ m_2(\mathbf{x}, \mathbf{I}) &:= \mathbf{x} \\ m_{n+1}(\mathbf{x}, a(\boldsymbol{\rho}_1) \dots, \mathbf{I}, \dots, a(\boldsymbol{\rho}_n)) &:= 0 \quad \text{if } n > 1. \end{aligned}$$

Here $\mathbf{I}$ is the unit of the ring $\mathcal{I}(\mathcal{Z})$.

For convenience, we will sometimes denote $m_1(\mathbf{x})$ by $\partial \mathbf{x}$ and $m_2(\mathbf{x}, a)$ by $\mathbf{x} \cdot a$. Note that this ∂ is different from the differential ∂ which makes $\mathcal{A}(n)$ into a differential algebra. Note also that nonzero coefficients $m_1(\mathbf{x})$ come from curves in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ where $\vec{\boldsymbol{\rho}} = \emptyset$. Since $\vec{\boldsymbol{\rho}}$ is a partition of the east punctures, this means that $m_1(\mathbf{x})$ counts curves which do not have any east punctures, i.e., which do not approach

$\partial\overline{\Sigma}$. The coefficient of a generator $\mathbf{y}$ in $m_1(\mathbf{x})$ is given by the number of provincial holomorphic curves connecting $\mathbf{x}$ to $\mathbf{y}$.

With these multiplication maps, $\widehat{CFA}(\mathcal{H})$ becomes a strictly unital right $\mathcal{A}_\infty$ module over $\mathcal{A}(\mathcal{Z})$. Before we show why these maps m_{n+1} fulfill the compatibility condition for an $\mathcal{A}_\infty$ module, we verify that the definitions make sense at all.

Lemma 4.14. *If $\mathcal{H}$ is provincially admissible, then the sum appearing in the definition of m_{n+1} is finite. If $\mathcal{H}$ is in fact admissible, then $m_{n+1} = 0$ for all but finitely many n. (Thus the type A structure will be operationally bounded.)*

Proof. There are only finitely many generators $\mathbf{y} \in \mathfrak{S}(\mathcal{H})$ in the first place. By Lemma 3.11, we know that any B with $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}}) \neq \emptyset$ must have B positive. But provincial admissibility implies by Proposition 2.23 that there are only finitely many positive classes $B \in \pi_2(\mathbf{x}, \mathbf{y})$ whose boundary on $\partial\overline{\Sigma}$ is given by some predetermined Reeb chords $\vec{\boldsymbol{\rho}}$. Thus the sum is finite. The number of elements in each $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ is finite thanks to Proposition 3.54. (Recall that compact 0-dimensional manifolds are just finite sets of points, after all.)

To show the second statement, let $|B|$ be the sums of all the local multiplicities of the regions in Σ. Recall that m_{n+1} is only nonzero when $\vec{\boldsymbol{\rho}}$ consists of n *nonempty* sets of Reeb chords. Then the only nonzero terms in m_{n+1} must involve homology classes B with $|B| \geq n$. After all, the sum of the local multiplicities of such B at the regions adjacent to $\partial\overline{\Sigma}$ should be n. But recall by Proposition 2.24 that, for any two generators $\mathbf{x}$ and $\mathbf{y}$, there are only finitely many positive domains which connect them. Again, there are only finitely many generators, so for all $|B| > N$ for some large N, the moduli space $\mathcal{M}^B$ is empty. Thus $m_{n+1} = 0$ for all $n > N$. □

Remark 4.15. Recall from Remark 2.25 that $\pi_2(\mathbf{x}, \mathbf{y})$ is nonempty if and only if $\mathbf{x}$ and $\mathbf{y}$ induce the same spin^c structure. This means that the moduli space $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ is nonempty only if $\mathfrak{s}_z(\mathbf{x}) = \mathfrak{s}_z(\mathbf{y})$, so we may decompose $\widehat{CFA}(\mathcal{H})$ into $\bigoplus \widehat{CFA}(\mathcal{H}, \mathfrak{s})$, where $\mathfrak{s}$ ranges over all spin^c structures on the 3-manifold Y represented by $\mathcal{H}$ and where $\widehat{CFA}(\mathcal{H}, \mathfrak{s})$ is the part of $\widehat{CFA}$ which only involves generators that induce the spin^c class $\mathfrak{s}$. This is useful for constructing a grading on $\widehat{CFA}$, but we will not spend time on this detail here.

4.3.2 Compatibility with the algebra

We have thus defined $\widehat{CFA}(\mathcal{H})$, though we have not actually shown that it is actually an $\mathcal{A}_\infty$ module over $\mathcal{A}(\mathcal{Z})$. In fact, we have the following statement.

Theorem 4.16. *If $\mathcal{H}$ is provincially admissible and $\mathcal{Z} = \partial\mathcal{H}$, then $\left(\widehat{CFA}(\mathcal{H}), \{m_i\}\right)$ is a (right) $\mathcal{A}_\infty$ module over $\mathcal{A}(\mathcal{Z})$.*

To prove this theorem, it suffices to prove Equation (4.2). Roughly speaking, this amounts to counting the ends of the index-2 moduli spaces. In particular, the terms in the $\mathcal{A}_\infty$ compatibility equation correspond to a given type of end, so Theorem 3.53 implies the compatibility.

Before explaining which terms correspond to which end, we provide a few examples which suggest the more general argument for compatibility and give geometric intuition for the $\mathcal{A}_\infty$ compatibility equation. Compare these examples with Examples 3.55 to 3.58.

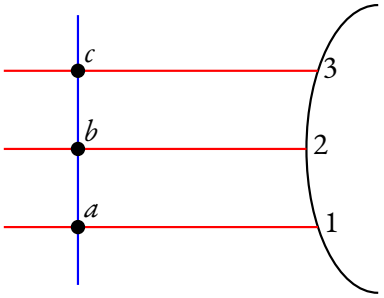

Figure 4.1: The local picture in Example 4.17.

Example 4.17. Consider Figure 4.1. The $\widehat{CFA}$ module of this local picture has three generators, namely $\{a\}$, $\{b\}$, and $\{c\}$. The relevant generators of the strands algebra are $\left[\frac{1}{2}\right]$, $\left[\frac{2}{3}\right]$, and $\left[\frac{1}{3}\right]$. (In Example 3.55, we denoted the Reeb chord $\left[\frac{1}{2}\right]$, for example, by ρ_{12}.)

There are two regions in Figure 4.1, namely the rectangle $12ba$ and the rectangle $23cb$. Both regions touch $\partial\overline{\Sigma}$ and hence are not provincial. Since $m_1 = \partial$ counts provincial domains between generators, this means that ∂ is trivial: $\partial\{a\} = \{a\}$, for example.

The compatibility condition for m_2 says that

$$\partial\mathbf{x} \cdot \boldsymbol{\rho} + \mathbf{x} \cdot \partial\boldsymbol{\rho} + \partial(\mathbf{x} \cdot \vec{\boldsymbol{\rho}}) = 0$$

for any generator $\mathbf{x}$ and Reeb chord $\boldsymbol{\rho}$. (Note that $\partial\boldsymbol{\rho}$ is the boundary operator on the differential algebra $\mathcal{A}(\mathcal{Z})$, while $\partial\mathbf{x}$ refers to the map m_1 which counts provincial domains between generators.) Each of these terms vanishes: The differential on the strands algebra vanishes, since there are no crossing in the strand diagrams of $\left[\frac{1}{2}\right]$, $\left[\frac{2}{3}\right]$, and $\left[\frac{1}{3}\right]$, while the differential on $\widehat{CFA}$ vanishes, since there are no provincial domains.

In fact, in this case, we have $m_i = 0$ for all $i \geq 3$. In particular, we have genuine associativity, instead of just associativity up to homotopy, so that $\widehat{CFA}$ is actually a differential module. To see this, it is enough to show that

$$(\mathbf{x} \cdot \boldsymbol{\rho}_1) \cdot \boldsymbol{\rho}_2 + \mathbf{x} \cdot (\boldsymbol{\rho}_1 \cdot \boldsymbol{\rho}_2) = 0.$$

Because our only nontrivial multiplications are

$$\begin{aligned} a \cdot \left[\tfrac{1}{2}\right] &= b \\ b \cdot \left[\tfrac{2}{3}\right] &= c \\ a \cdot \left[\tfrac{1}{3}\right] &= c, \end{aligned}$$

the only case to check is that

$$\left(a \cdot \left[\tfrac{1}{2}\right]\right) \cdot \left[\tfrac{2}{3}\right] + a \cdot \left(\left[\tfrac{1}{2}\right] \cdot \left[\tfrac{2}{3}\right]\right) = 0.$$

But both terms equal c.

One way to understand this associativity geometrically is as follows. Recall from Example 3.55 that the index-two moduli space from $\{a\}$ to $\{c\}$ with asymptotics at east infinity given by $\left\{\left[\frac{1}{2}\right], \left[\frac{2}{3}\right]\right\}$ has two ends: a two-story end and a split curve (i.e., collision of levels) end. The two-story end corresponds to the term $\left(a \cdot \left[\frac{1}{2}\right]\right) \cdot \left[\frac{2}{3}\right]$: The product $a \cdot \left[\frac{1}{2}\right]$ counts curves which begin at a and end at b via the Reeb

chord ρ_{12}, while the subsequent product with $\left[\begin{smallmatrix}2\\3\end{smallmatrix}\right]$ counts the second story of the holomorphic building, which ends at $\{c\}$ via the Reeb chord ρ_{23}. On the other hand, the term $a \cdot \left(\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right] \cdot \left[\begin{smallmatrix}2\\3\end{smallmatrix}\right]\right)$ counts curves which converge to the join of $\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right]$ and $\left[\begin{smallmatrix}2\\3\end{smallmatrix}\right]$, and hence which represent a collision of levels. Thus showing associativity in this context amounts to counting ends of a given index-two moduli space, which is 0 mod 2 by Theorem 3.53.

Example 4.18. Now consider Figure 4.2. This has four generators, namely $\{a, c\}$, $\{a, d\}$, $\{b, c\}$, and

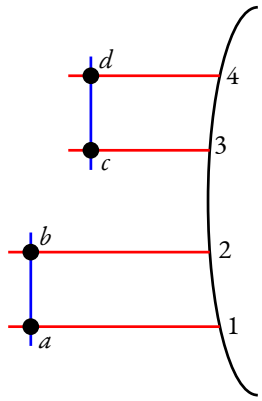

Figure 4.2: The local picture in Example 4.18.

$\{b, d\}$. Since none of the regions shown in this local picture are provincial, we again have that ∂ is trivial. The nontrivial m_2 operations are

$$\begin{aligned}
\{a, c\} \cdot \left[\begin{smallmatrix}1 & 3\\2\end{smallmatrix}\right] &= \{b, c\}\\
\{a, c\} \cdot \left[\begin{smallmatrix}3 & 1\\4\end{smallmatrix}\right] &= \{a, d\}\\
\{a, c\} \cdot \left[\begin{smallmatrix}1 & 3\\2 & 4\end{smallmatrix}\right] &= \{b, d\}\\
\{b, c\} \cdot \left[\begin{smallmatrix}3 & 2\\4\end{smallmatrix}\right] &= \{b, d\}\\
\{a, d\} \cdot \left[\begin{smallmatrix}1 & 4\\2\end{smallmatrix}\right] &= \{b, d\}
\end{aligned}$$

Note that, here, we have strands algebra elements like $\left[\begin{smallmatrix}1 & 3\\2\end{smallmatrix}\right]$ to indicate that the second sheet of the covering over $[0, 1] \times \mathbb{R}$ is provincial, with boundary on the third α-arc (or, rather, the α-arc corresponding to the point 3 on the boundary).

This multiplication, like in the previous example, is associative. The only equation to check is that

$$\left(\{a, c\} \cdot \left[\begin{smallmatrix}1 & 3\\2\end{smallmatrix}\right]\right) \cdot \left[\begin{smallmatrix}3 & 2\\4\end{smallmatrix}\right] + \{a, c\} \cdot \left(\left[\begin{smallmatrix}1 & 3\\2\end{smallmatrix}\right] \cdot \left[\begin{smallmatrix}3 & 2\\4\end{smallmatrix}\right]\right) = 0. \tag{4.3}$$

But $\left[\begin{smallmatrix}1 & 3\\2\end{smallmatrix}\right] \cdot \left[\begin{smallmatrix}3 & 2\\4\end{smallmatrix}\right] = \left[\begin{smallmatrix}1 & 3\\2 & 4\end{smallmatrix}\right]$, and so both terms are $\{a, c\} \cdot \left[\begin{smallmatrix}1 & 3\\2 & 4\end{smallmatrix}\right] = \{b, d\}$.

Geometrically, this equation is given by counting ends of the moduli space from $\mathbf{x} = \{a, c\}$ to $\mathbf{y} = \{b, d\}$ with asymptotics at east infinity given by $\left[\begin{smallmatrix}1 & 3\\2 & 4\end{smallmatrix}\right]$ such that the Reeb chord from 3 to 4 occurs before the Reeb chord from 1 to 2. The first term in Equation (4.3) represents the holomorphic building end, in which we approach ρ_{12} infinitely far before we approach ρ_{34}. In the language of Example 3.56, this corresponds to the two-story end which occurs when $\mathrm{ev}_{34} - \mathrm{ev}_{12} = \infty$. The second term occurs at the collision of levels end which occurs when $\mathrm{ev}_{34} - \mathrm{ev}_{12} = 0$. Thus the two terms correspond to the two ends of the moduli space from Example 3.56. Proving that the $\widehat{CFA}$ algebra of Figure 4.2 is compatible

with the strands algebra is thus equivalent to stating that a certain moduli space has an even number of ends.

Example 4.19. The generators of Figure 4.3 are $\{a, d\}$, $\{a, e\}$, $\{b, c\}$, and $\{b, e\}$. We now have the following nontrivial multiplications in $\widehat{CFA}$:

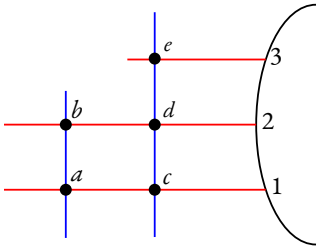

Figure 4.3: The local picture in Example 4.19.

$$\begin{aligned}
\partial\{a, d\} &= \{b, c\} \\
\{a, e\} \cdot \left[\begin{smallmatrix} 1 & 3 \\ 2 & \end{smallmatrix}\right] &= \{b, e\} \\
\{b, c\} \cdot \left[\begin{smallmatrix} 1 & 2 \\ 3 & \end{smallmatrix}\right] &= \{b, e\} \\
\{a, d\} \cdot \left[\begin{smallmatrix} 1 & 2 \\ 2 & 3 \end{smallmatrix}\right] &= \{b, e\} \\
\{a, d\} \cdot \left[\begin{smallmatrix} 2 & 1 \\ 3 & \end{smallmatrix}\right] &= \{a, e\}
\end{aligned}$$

The only nontrivial $\mathcal{A}_\infty$ relation is

$$(\partial\{a, d\}) \cdot \left[\begin{smallmatrix} 1 & 2 \\ 3 & \end{smallmatrix}\right] + \{a, d\}\partial\left[\begin{smallmatrix} 1 & 2 \\ 3 & \end{smallmatrix}\right] = 0.$$

Recall that $\partial\left[\begin{smallmatrix} 1 & 2 \\ 3 & \end{smallmatrix}\right]$ refers to the boundary operator on the strands algebra, which is itself a differential algebra, whereas the ∂ in $\partial\{a, d\}$ refers to the map m_1 on $\widehat{CFA}$.

The first term in the $\mathcal{A}_\infty$ relation above corresponds to the two-story end from Example 3.57. The first story is represented by $\partial\{a, d\}$, while the second story is represented by the product with $\left[\begin{smallmatrix} 1 & 2 \\ 3 & \end{smallmatrix}\right]$. The second term corresponds to the join curve end when the branch point escapes to east infinity.

So far, all of our type A modules have been genuine differential modules. If one checks $\widehat{CFA}$ for the more complicated diagram for Example 3.58, one obtains a differential module again. In general, however, this is not the case, and our higher multiplications $m_3, m_4, \ldots$ are, in fact, necessary.

Example 4.20. Consider the local Heegaard diagram and corresponding shaded domain $D_1 + D_2$ in Figure 4.4. The type A module is generated by $\{a, c\}$, $\{b, c\}$, $\{b, d\}$, and $\{a, d\}$. The disk D_1 from d to c is counted in the moduli spaces involved in the definition of ∂. Thus

$$\partial\{a, d\} = \{a, c\} \quad \text{and} \quad \partial\{b, d\} = \{b, c\}.$$

We may choose J so that there is a holomorphic map whose projection to Σ looks like D_2. Thus

$$\{a, d\} \cdot \left[\begin{smallmatrix} 1 \\ 3 \end{smallmatrix}\right] = \{b, d\}.$$

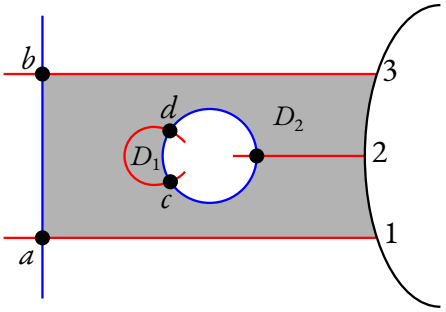

Figure 4.4: A local picture which demonstrates that $\widehat{CFA}(\mathcal{H})$ need not be associative. The unlabeled point in $\overline{\boldsymbol{\alpha}} \cap \boldsymbol{\beta}$ is never part of a generator. After all, a generator must have either c or d, since they are the only generators lying on the α-curve bounding D_1. (While this curve is drawn as an arc, we imagine it closing off somewhere away from this local picture.) But since only one element in a generator can lie on the β-circle in the middle, no generator can include this unlabeled point.

To understand this, consider the rightmost diagram in Figure 4.5. The green slit goes from d to some point on the red circle. This point is the unique point such that the slit domain (shaded in gray) is conformally equivalent to an annulus which is a branched double cover of $[0,1] \times \mathbb{R}$.

But note that associativity fails because

$$\{a,d\} \cdot \left(\left[\tfrac{1}{2}\right] \cdot \left[\tfrac{2}{3}\right]\right) = \{a,d\} \cdot \left[\tfrac{1}{3}\right] = \{b,d\} \neq 0 = \left(\{a,d\} \cdot \left[\tfrac{1}{2}\right]\right) \cdot \left[\tfrac{2}{3}\right]. \tag{4.4}$$

(The last equality follows from the fact that $\{a,d\} \cdot \left[\tfrac{1}{2}\right] = 0$.) The term on the left-hand side of Equation (4.4) corresponds to the rightmost figure in Figure 4.4, in which a split curve degenerates. In particular, the split curve end corresponds to an element of $\mathcal{M}^{D_1+D_2}(\{a,d\},\{b,d\},\{\rho_{13}\})$. This is indeed an end of the moduli space $\mathcal{M}^{D_1+D_2}(\{a,d\},\{b,d\},\{\rho_{12},\rho_{23}\})$.

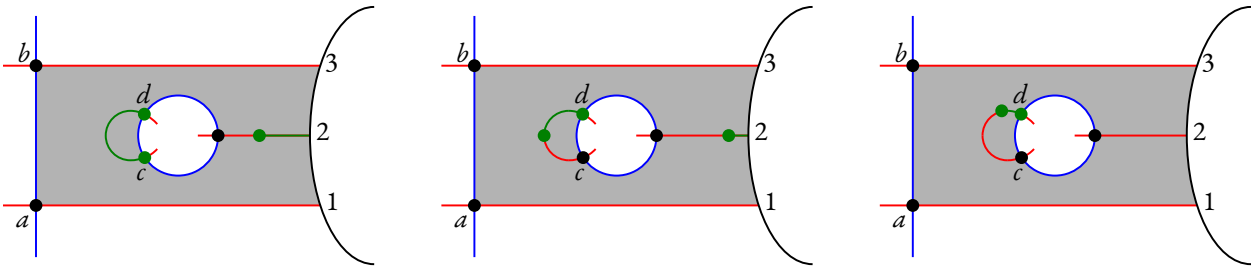

Figure 4.5: Ends of the moduli space in question.

But this time, because the right-hand side of Equation (4.4) is zero, we have not yet accounted for the other end of this moduli space. In fact, there is also a two-story end, seen on the left side of Figure 4.5. This building is an element of $\mathcal{M}^{D_1}(\{a,d\},\{a,c\}) \times \mathcal{M}^{D_2}(\{a,c\},\{b,d\},\{\rho_{12},\rho_{23}\})$. That is, the first story, which comes from $\partial\{a,d\}$, corresponds to the region D_1. The second story occurs by approaching the Reeb chords $\left[\tfrac{1}{2}\right]$ and $\left[\tfrac{2}{3}\right]$ at the same time. Algebraically, then, this term corresponds to

$$m_3\left(\partial\{a,d\},\left[\tfrac{1}{2}\right],\left[\tfrac{2}{3}\right]\right) = m_3\left(\{a,c\},\left[\tfrac{1}{2}\right],\left[\tfrac{2}{3}\right]\right) = \{b,d\}.$$

Thus we see that the $\mathcal{A}_\infty$ associativity relation, which says that

$$\{a,d\} \cdot \left(\left[\tfrac{1}{2}\right] \cdot \left[\tfrac{2}{3}\right]\right) + \left(\{a,d\} \cdot \left[\tfrac{1}{2}\right]\right) \cdot \left[\tfrac{2}{3}\right] + m_3(\partial\{a,d\},\left[\tfrac{1}{2}\right],\left[\tfrac{2}{3}\right]) = 0,$$

holds, since the middle term is zero and the other two terms correspond to ends of an index-two moduli space. (Note that we do not have any other terms since $\partial\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right] = \partial\left[\begin{smallmatrix}2\\3\end{smallmatrix}\right] = 0$ and $\partial\{b,d\} = 0$.)

Example 4.21. The examples we have discussed thus far have involved local pictures of bordered Heegaard diagrams. We now do an example in full. Recall the example of a bordered Heegaard diagram $\mathcal{H}$ for the genus-1 handlebody in Figure 2.6. It will be helpful for our purposes to draw this diagram as in Figure 4.6. There is only one generator in this example, namely $\{x_0\}$. Note that there are only two re-

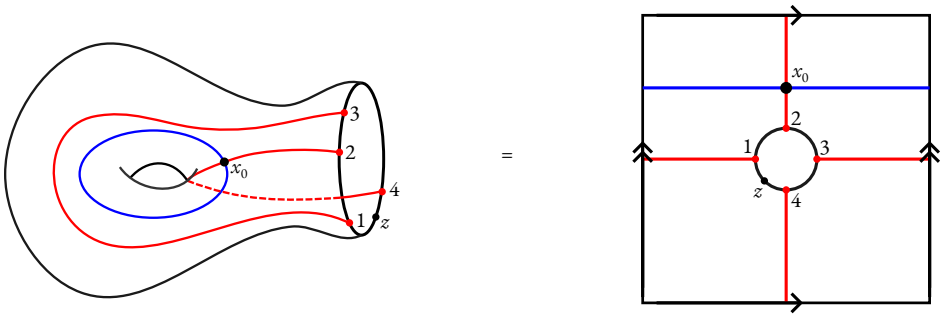

Figure 4.6: A bordered Heegaard diagram for the genus-1 handlebody.

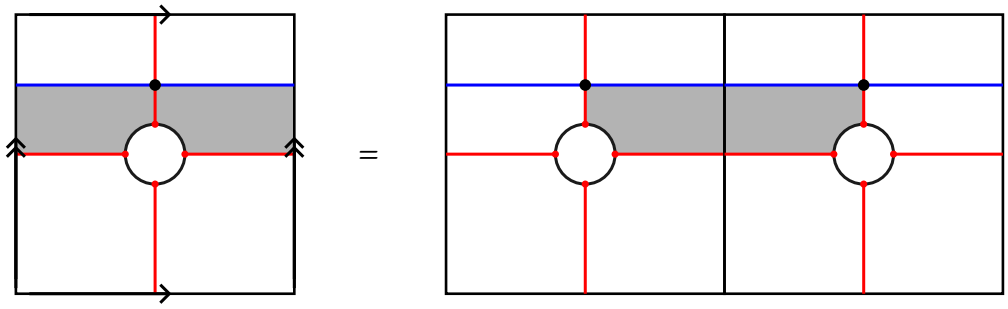

Figure 4.7: The shaded region indicates a curve in $\mathcal{M}^B(x_0, x_0; \rho_{23}, \rho_{12})$ which is counted in the x_0-coefficient of $m_3(x_0, \rho_{23}, \rho_{12})$.

gions: One is the shaded domain $D(B)$ on the left side of Figure 4.7, while the other is its complement. The latter crosses z, though, so our holomorphic curves must project only to the region $D(B)$. As shown in Figure 4.7 and Figure 4.8, then, we have

$$m_3(x_0, \rho_{23}, \rho_{12}) = x_0 \quad \text{and} \quad m_4(x_0, \rho_{23}, \rho_{13}, \rho_{12}) = x_0.$$

In general, our only nontrivial multiplications are these higher multiplications of the form

$$m_{n+2}(x_0, \rho_{23}, \rho_{13}, \dots, \rho_{13}, \rho_{12}) = x_0,$$

where there are n copies of $\rho_{13} = \rho_{12} \uplus \rho_{23}$. It is straightforward to verify that this obeys our $\mathcal{A}_\infty$ compatibility relations. (Roughly speaking, the second term in Equation (4.2) vanishes, while the first and third terms cancel out.)

In this case, not only is $\widehat{CFA}(\mathcal{H})$ not a differential module, but it is also, in fact, an *unbounded* $\mathcal{A}_\infty$ module. Thus $\mathcal{H}$ could not be an admissible Heegaard diagram, thanks to Lemma 4.14. Indeed, the domain $D(B)$ is a periodic domain, since $B \in \pi_2(\mathbf{x}, \mathbf{x})$, but has only positive coefficients.

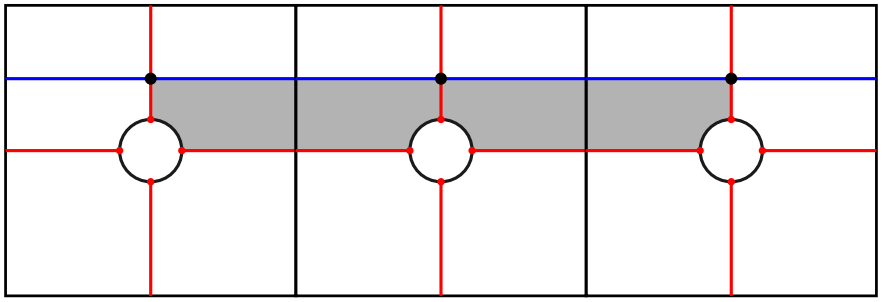

Figure 4.8: By doubling the domain, we obtain a curve in $\mathcal{M}^{2B}(x_0, x_0; \rho_{23}, \rho_{13}, \rho_{12})$ which is counted in the x_0-coefficient of $m_4(x_0, \rho_{23}, \rho_{13}, \rho_{12})$.

Example 4.22. The bordered Heegaard diagram $\mathcal{H}'$ in Figure 4.9, like that of Figure 4.6, represents the genus-1 handlebody. There is a provincial domain between y and w, so $\partial y = w$. Our only nontrivial

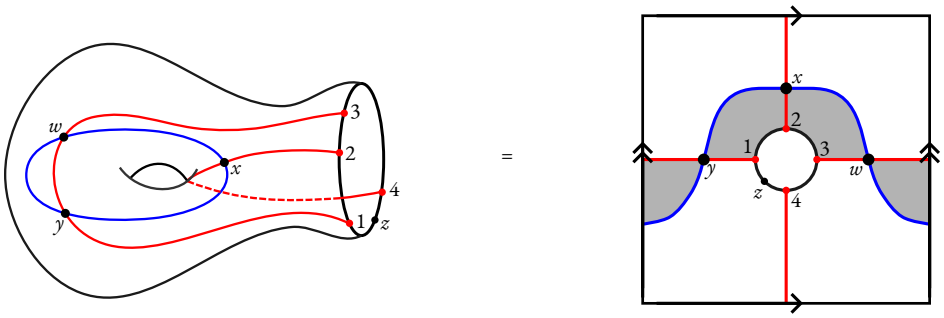

Figure 4.9: This is a different bordered Heegaard diagram representing the standard genus-1 handlebody. There are three regions not adjacent to z. (Note that the region to the left of y is the same as the region to the right of w.)

multiplications are

$$y \cdot \left[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}\right] = x$$
$$x \cdot \left[\begin{smallmatrix} 2 \\ 3 \end{smallmatrix}\right] = w$$
$$y \cdot \left[\begin{smallmatrix} 1 \\ 3 \end{smallmatrix}\right] = w.$$

Higher multiplications $m_3, m_4, \ldots$ vanish identically. This is thus an honest differential module which satisfies associativity on the nose, as opposed to up to homotopy.

Remark 4.23. One might wonder whether there is some condition on $\mathcal{H}$ which makes $\widehat{CFA}(\mathcal{H})$ a genuine differential module. A bordered Heegaard diagram is called **nice** if each region which does not touch the basepoint $z \in \partial\mathcal{H}$ is a topological disk with at most four corners. In this case, the associated type A module does not have higher differentials. Loosely speaking, nice diagrams do not have higher multiplications because any holomorphic curve $u \in \mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ whose source goes to east infinity must escape toward east infinity at a single time. Thus $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ is empty if $\vec{\boldsymbol{\rho}}$ has more than one part. Note that the regions D_1 in Example 4.20 and B in Example 4.21 are both annuli, so the Heegaard diagrams in those examples are not nice. This is why the corresponding type A modules are not differential modules.

Furthermore, one may always obtain a nice Heegaard diagram by performing "finger moves." These moves look like the move between Examples 4.21 and 4.22. See [LOT18, Chapter 8] and [SW10] for more details on nice diagrams.

We now prove that $\widehat{CFA}(\mathcal{H})$ is indeed an $\mathcal{A}_\infty$ module over $\mathcal{A}(\mathcal{Z})$.

Proof of Theorem 4.16. We would like to show that

$$\begin{aligned}0 = \sum_{i+j=n+1} & m_i(m_j(\mathbf{x}, a_1, \dots, a_{j-1}), \dots, a_{n-1}) \\ & + \sum_{\ell=1}^{n-1} m_n(\mathbf{x}, a_1, \dots, \partial a_\ell, \dots, a_{n-1}) \\ & + \sum_{\ell=1}^{n-2} m_{n-1}(\mathbf{x}, a_1, \dots, a_\ell a_{\ell+1}, \dots, a_{n-1}).\end{aligned}$$

Consider the $\mathbf{y}$-coefficient of each term above, where $\mathbf{y} \in \mathfrak{S}(\mathcal{H})$ is arbitrary. The first term corresponds to two-story ends in all moduli spaces $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ for $B \in \pi_2(\mathbf{x}, \mathbf{y})$, $\mathrm{ind}(B, \vec{\boldsymbol{\rho}}) = 1$, and $(\mathbf{x}, \vec{\boldsymbol{\rho}})$ strongly boundary monotone.

The second term corresponds to all join curve ends and odd shuffle curve ends. Showing this takes a little bit of work: First, note that $\partial a(\boldsymbol{\rho})$ is the sum of all $a(\boldsymbol{\rho}')$ where $\boldsymbol{\rho}'$ is obtained from $\boldsymbol{\rho}$ in one of two ways: (1) by replacing some chord $\rho_1 \in \boldsymbol{\rho}$ with a splitting $\{\rho_2, \rho_3\}$, i.e., with a pair such that $\rho_1 = \rho_2 \uplus \rho_3$, such that the result $\boldsymbol{\rho}'$ is consistent and has no double crossings; or (2) by replacing a nested pair of Reeb chords in $\boldsymbol{\rho}$ by its corresponding interleaved pair such that no double crossings are introduced. A priori, however, the east asymptotics of join curve ends and shuffle curve ends may either be inconsistent or have double crossings. It is easy to show, however, that such cases never occur as elements of the boundary $\partial\overline{\mathcal{M}^B}(\mathbf{x}, \mathbf{y}; \boldsymbol{\rho}')$.

The final term corresponds to the collisions of levels. We must argue analogously to the second term to verify this. See [LOT18, Section 7.2] for details.

Note that Theorem 3.53 implies the result. In particular, the sum in the $\mathcal{A}_\infty$ compatibility equation is exactly equal to the total number of ends of all moduli spaces $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$ for suitable B and $\vec{\boldsymbol{\rho}}$, minus any even shuffle curve ends. But there are an even number of even shuffle curve ends, proving the result. □

4.3.3 Invariance

It turns out that, with this definition, the $\mathcal{A}_\infty$ module $\widehat{CFA}(\mathcal{H})$ is actually an invariant of the bordered 3-manifold defined by $\mathcal{H}$. This is thus our first bordered Heegaard Floer invariant.

Theorem 4.24. *The $\mathcal{A}_\infty$ module $\widehat{CFA}(\mathcal{H})$ is independent, up to $\mathcal{A}_\infty$ homotopy equivalence, of the choice of admissible almost complex structure which achieves transversality. Furthermore, if $\mathcal{H}$ and $\mathcal{H}'$ are provincially admissible bordered Heegaard diagrams for the same bordered 3-manifold $(Y, \mathcal{Z}, \phi : F(\mathcal{Z}) \to \partial Y)$, then the $\mathcal{A}_\infty$ $\mathcal{A}(\mathcal{Z})$-modules $\widehat{CFA}(\mathcal{H})$ and $\widehat{CFA}(\mathcal{H}')$ are homotopy equivalent.*

This justifies the notation $\widehat{CFA}(Y)$ for the type A module of the bordered 3-manifold Y. (It would perhaps be more accurate to write the type A module as $\widehat{CFA}(Y, \mathcal{Z}, \phi)$, or as $\widehat{CFA}(Y, \phi : F(\mathcal{Z}) \to \partial Y)$, but this is a bit unwieldy.)

We do not prove this theorem, whose proof is somewhat distinct from the primary narrative thus far of moduli spaces of holomorphic cures and the degenerations which appear in the compactification. See [LOT18, Section 7.3] for a full proof.

Example 4.25. Recall the diagrams $\mathcal{H}$ and $\mathcal{H}'$ from Examples 4.21 and 4.22. They clearly represent the same bordered 3-manifold (namely the genus-1 handlebody), since their respective curves only differ by an isotopy. The $\mathcal{A}_\infty$ modules $\widehat{CFA}(\mathcal{H})$ and $\widehat{CFA}(\mathcal{H}')$ are not equal, as the former is unbounded while the latter is operationally bounded. But they are homotopy equivalent.

First, we construct a homomorphism $f : \widehat{CFA}(\mathcal{H}) \to \widehat{CFA}(\mathcal{H}')$. The most natural such homomorphism should have $f_1(x_0) = x$. To build this into an $\mathcal{A}_\infty$ homomorphism, we need

$$f_1(x_0) \cdot \left[\tfrac{2}{3}\right] = \partial f_2(x_0, \left[\tfrac{2}{3}\right]).$$

Note that, technically, there are three other terms, namely $f_2(x_0, \partial\left[\tfrac{2}{3}\right]), f_2(\partial x_0, \left[\tfrac{2}{3}\right])$, and $f_1(x_0 \cdot \left[\tfrac{2}{3}\right])$. But because $\partial\left[\tfrac{2}{3}\right] = 0$ and $\mathcal{H}$ has trivial m_1 and m_2, it follows that these terms vanish. Since $f_1(x_0) = x$, it follows that $f_2(x_0, \left[\tfrac{2}{3}\right]) = y$. Note that the corresponding $\mathcal{A}_\infty$ compatibility equation for any other basis element $a(\rho)$ of the strands algebra $\mathcal{A}(\mathcal{Z})$ is trivially satisfied by setting $f_2(x_0, a(\rho)) = 0$, since all the terms vanish. This completely defines $f_2 : \widehat{CFA}(\mathcal{H}) \otimes \mathcal{A}(\mathcal{Z}) \to \widehat{CFA}(\mathcal{H}')$.

Now the compatibility equations for f_3 and f_4 say that

$$\begin{aligned} f_1\left(m_3'\left(x_0, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{2}\right]\right)\right) &= m_2\left(f_2\left(x_0, \left[\tfrac{2}{3}\right]\right), \left[\tfrac{1}{2}\right]\right) + \partial f_3\left(x_0, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{2}\right]\right) \\ f_1\left(m_4'\left(x_0, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{3}\right], \left[\tfrac{1}{2}\right]\right)\right) &= m_2\left(f_3\left(x, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{3}\right]\right), \left[\tfrac{1}{2}\right]\right) + \partial f_4\left(x, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{3}\right], \left[\tfrac{1}{2}\right]\right). \end{aligned}$$

Again, there are other terms, but they vanish. (Note that the strand algebra elements do not compose.) The left-hand side of the f_3-compatibility equation is x, as is the first term on the right-hand side, so it follows that $\partial f_3(x_0, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{2}\right]) = 0$. It turns out that we may set $f_3(x_0, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{2}\right]) = 0$. On the other hand, the compatibility equation for f_4 is satisfied if we let $f_3(x, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{3}\right]) = y$ and $f_4(x, \left[\tfrac{2}{3}\right], \left[\tfrac{1}{3}\right], \left[\tfrac{1}{2}\right]) = 0$.

This suggests a general formula for f:

$$\begin{aligned} f_1(x_0) &= x, \\ f_k(x_0, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_{k-1})) &= y \text{ when } a(\boldsymbol{\rho}_1) = \left[\tfrac{2}{3}\right] \text{ and } a(\boldsymbol{\rho}_i) = \left[\tfrac{1}{3}\right], \\ f_k(x_0, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_{k-1})) &= 0 \text{ otherwise.} \end{aligned}$$

We may verify that this gives an $\mathcal{A}_\infty$ homomorphism $\widehat{CFA}(\mathcal{H}) \to \widehat{CFA}(\mathcal{H}')$.

Before even trying to show that this is a homotopy equivalence, note that it is a quasi-isomorphism, i.e., gives an isomorphism on cohomology. After all, the cohomology is generated by $\{x\}$ (respectively, $\{x_0\}$) in $\widehat{CFA}(\mathcal{H}')$ (respectively, $\widehat{CFA}(\mathcal{H})$) with trivial differential. The homomorphism f is simply the isomorphism $x_0 \mapsto x$ on cohomology, then.

We briefly indicate how to construct the homotopy equivalence between the type A modules. The details are very similar to that of the construction of f above. Define $g : \widehat{CFA}(\mathcal{H}') \to \widehat{CFA}(\mathcal{H})$ by

$$x_0 = g_1(x) = g_2\left(w, \left[\tfrac{1}{2}\right]\right) = g_3\left(w, \left[\tfrac{1}{3}\right], \left[\tfrac{1}{2}\right]\right) = g_4\left(w, \left[\tfrac{1}{3}\right], \left[\tfrac{1}{3}\right], \left[\tfrac{1}{2}\right]\right) = \dots.$$

(Everything else gets mapped to zero.)

Notice that $(g \circ f)(x_0, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_{k-1})) = x_0$ if $a(\boldsymbol{\rho}_1) = \left[\tfrac{2}{3}\right], a(\boldsymbol{\rho}_{k-1}) = \left[\tfrac{1}{2}\right]$, and the intermediate algebra elements $a(\boldsymbol{\rho}_i) = \left[\tfrac{1}{3}\right]$. Otherwise, we have $(g \circ f)(x_0, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_{k-1})) = 0$. One may verify

that this is homotopic to the identity homomorphism on $\widehat{CFA}(\mathcal{H})$ via the homotopy with $h_1(x_0) = x_0$ and higher h_i's vanishing everywhere.

On the other hand, we have $(f \circ g)(x) = x$ and $(f \circ g)(w, a(\boldsymbol{\rho}_1), \dots, a(\boldsymbol{\rho}_{k-1})) = y$ if the $\boldsymbol{\rho}_i$'s are, in order, some number of $\left[\frac{1}{3}\right]$'s, then a $\left[\frac{1}{2}\right]$, followed by a $\left[\frac{2}{3}\right]$, and some number of $\left[\frac{1}{3}\right]$'s. Otherwise, we define $f \circ g$ to be 0. This is homotopic to the identity homomorphism on $\widehat{CFA}(\mathcal{H}')$, this time by the homotopy with $h_1 = 0$ and $h_n\left(w, \left[\frac{1}{3}\right], \dots, \left[\frac{1}{3}\right]\right) = y$ for $n \geq 1$.

4.4 Type D structures

The other bordered invariant which we will define, denoted $\widehat{CFD}(Y)$, is a somewhat stranger algebraic creature, namely a module associated to a so-called "type D structure."

Let $\mathcal{A}$ be an $\mathcal{A}_\infty$ algebra with underlying module A, as usual. Let N be a left $\mathbf{k}$-module and let $\delta_N^1 : N \to A \otimes N$, with tensor product taken over $\mathbf{k}$, as usual. We may now construct a sequence of maps $\delta_N^i : N \to A^{\otimes i} \otimes N$ given by

$$\begin{aligned} \delta_N^0 &= \mathrm{id}_N \\ \delta_N^i &= (\mathrm{id}_{A^{\otimes(i-1)}} \otimes \delta_N^1) \circ \delta_N^{i-1} \end{aligned}$$

If $\delta_N^i = 0$ for sufficiently large i, then we say that (N, δ_N^1) is **bounded** and we may promote our maps δ_N^i to a map on the tensor algebra:

$$\begin{aligned} \delta : N &\to \mathcal{T}^*(A) \otimes N \\ \mathbf{x} &\mapsto \sum_{i=0}^{\infty} \delta_N^i(\mathbf{x}). \end{aligned}$$

Almost by definition, we have the condition that $(\mathrm{id}_{A^{\otimes j}} \otimes \delta^i) \circ \delta^j = \delta^{i+j}$ for all $i, j \geq 0$. Here, and later on, we omit the N in the subscript when N is clear from the context. Graphically, we may depict this relationship as follows:

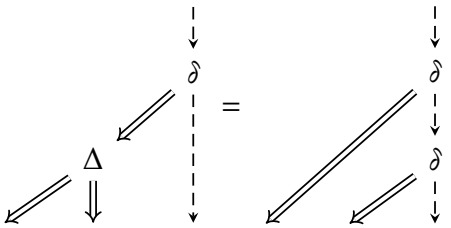

As in our diagrams from Section 4.1, we use dashed lines to "color" elements of the module, which in this case is N. We put the dashed lines on the right side this time to indicate that N is a left module.

Remark 4.26. If (N, δ_N^1) is not bounded, we may complete the tensor algebra into $\overline{\mathcal{T}}^*(A) := \prod_{i=0}^{\infty} A^{\otimes i}$. Then we may still put our maps δ^i together into a map δ.

Definition 4.27. Let $\mathcal{A}$ be an $\mathcal{A}_\infty$ algebra and (N, δ_N^1) a pair as above. In particular, we let N be a left $\mathbf{k}$-module and $\delta^1 : N \to A \otimes N$ is a map such that either $\mathcal{A}$ is operationally bounded or the pair (N, δ_N^1) is

bounded. Recall our definition of $\overline{D} : \mathcal{T}^*(\mathcal{A}) \to \mathcal{T}^*(\mathcal{A})$ from Equation (4.1). Then we say that (N, δ^1_N) is a **type D structure over** $\mathcal{A}$ if

$$(\overline{D} \otimes \mathrm{id}_N) \circ \delta = 0.$$

Graphically, we write this condition as

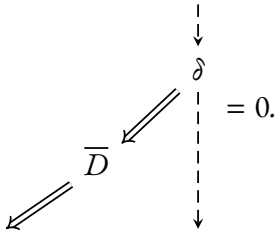

In the contexts which we will be concerned about, we will always have $\mathcal{A}$ operationally bounded, so all of our definitions will work for unbounded type D structures too.

Now suppose we have two bounded type D structures $(N_1, \delta^1_{N_1})$ and $(N_2, \delta^1_{N_2})$. Let $\psi^1 : N_1 \to \mathcal{A} \otimes N_2$ be a map, and define

$$\begin{aligned} \psi^k : N_1 &\to \mathcal{A}^{\otimes k} \otimes N_2 \\ \mathbf{x} &\mapsto \sum_{i+j=k-1} (\mathrm{id}_{\mathcal{A}^{\otimes(i+1)}} \otimes \delta^j_{N_2}) \circ (\mathrm{id}_{\mathcal{A}^{\otimes i}} \otimes \psi^1) \circ \delta^i_{N_1}. \end{aligned}$$

Since N_1 and N_2 are bounded, so too is ψ^k in the sense that $\psi^k = 0$ for sufficiently large k. Then $\psi = \sum \psi^k$ is a map $N_1 \to \mathcal{T}^*(\mathcal{A}) \otimes N_2$. If we remove the condition that N_1 and N_2 are bounded, then ψ simply has codomain $\overline{\mathcal{T}}^*(\mathcal{A}) \otimes N_2$.

Definition 4.28. A map $\psi^1 : N_1 \to \mathcal{A} \otimes N_2$ is a **type D homomorphism** if $(\overline{D} \otimes \mathrm{id}_{N_2}) \circ \psi = 0$. With dashed lines representing elements of N_1 and dotted lines representing elements of N_2, this may be represented graphically as

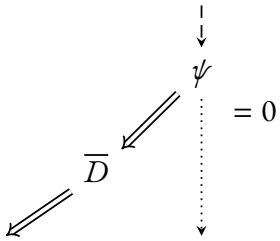

Definition 4.29. Similarly, two type D homomorphisms $\psi^1, \phi^1 : N_1 \to \mathcal{A} \otimes N_2$ are **(type D) homotopic** if there is an analogously constructed $h^1 : N_1 \to \mathcal{A} \otimes N_2$ such that

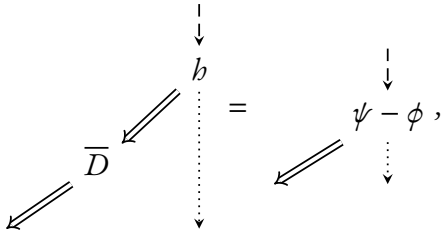

i.e., such that $(\overline{D} \otimes \mathrm{id}_{N_2}) \circ h = \psi - \phi$.

Unraveling these definitions in the case that $\mathcal{A}$ is a differential algebra instead of a general $\mathcal{A}_\infty$ algebra, we see that the compatibility condition is equivalent to the condition that

$$(\mu_2 \otimes \mathrm{id}_N) \circ (\mathrm{id}_A \otimes \delta^1_N) \circ \delta^1_N + (\mu_1 \otimes \mathrm{id}_N) \circ \delta^1_N : N \to A \otimes N \tag{4.5}$$

vanishes. One may rewrite the conditions for homomorphisms and homotopies in a similar way. See [LOT18, Definition 2.18].

The case in which we are interested is actually even more restricted than simply asking that $\mathcal{A}$ is a differential algebra. In particular, we are interested only in type D structures arising from the following example.

Example 4.30. Let $\mathcal{A}$ be a differential algebra and M a differential module which is free as an A-module. Consider some basis of M over A and let X denote the span of this basis, so that $M = \mathcal{A} \otimes X$. Restricting the boundary operator on M to X gives a map

$$\delta^1 : X \to \mathcal{A} \otimes X = M.$$

The pair (X, δ^1) is a type D structure, and restrictions of module maps are type D homomorphisms.

We may go the other direction, and obtain a left module from a type D structure. One may thus think of a type D structure as an additional combinatorial piece of data on top of the differential module structure. In particular, we have the following proposition.

Proposition 4.31. *Let $\mathcal{A}$ be a differential algebra and (N, δ^1_N) a type D structure. Then we may define an associated differential module $\mathcal{N}$ over $\mathcal{A}$. In particular, $\mathcal{N}$ has underlying module $A \otimes N$, which may then be given the structure of a (differential) left $\mathcal{A}$-module with maps*

$$\begin{aligned} m_1(a \otimes \mathbf{x}) &:= \left[(\mu_2 \otimes \mathrm{id}_N) \circ (\mathrm{id}_A \otimes \delta^1_N) + \mu_1 \otimes \mathrm{id}_N\right](a \otimes \mathbf{x}) \\ m_2(a \otimes (b \otimes \mathbf{x})) &:= \mu_2(a \otimes b) \otimes \mathbf{x}. \end{aligned}$$

Moreover, if we have a type D homomorphism $\psi^1 : N_1 \to A \otimes N_2$ between two type D structures, then there is an induced (chain) map of differential modules $A \otimes N_1 \to A \otimes N_2$ defined by

$$a \otimes \mathbf{x} \mapsto (m_2 \otimes \mathrm{id}_{N_2}) \circ (\mathrm{id}_A \otimes \psi^1).$$

Similarly, homotopies between type D homomorphisms induce chain homotopies between the corresponding chain maps.

The proof simply involves unwinding the definitions, and we omit it here.

We conclude this somewhat tedious section with the following proposition, which will be useful for showing that $\widehat{CFD}(Y)$ is invariant up to homotopy.

Proposition 4.32. *If N_1 and N_2 are two type D structures over a differential algebra A, with associated differential modules $\mathcal{N}_1$ and $\mathcal{N}_2$, then the correspondence in the previous proposition gives an identification between a type D homomorphism from N_1 to N_2 and a homomorphism of the differential modules $\mathcal{N}_1$ and $\mathcal{N}_2$. Moreover, two type D homomorphisms are homotopic if and only if the corresponding differential module homomorphisms are A-equivariant homotopic.*

4.5 The type D module $\widehat{CFD}$

As in the definition of $\widehat{CFA}$, let $\mathcal{H} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, z)$ be a provincially admissible Heegaard diagram. This time, let $\mathcal{Z} = -\partial\mathcal{H}$ be the *reverse* of the pointed matched circle. In particular, we let $\mathcal{Z} = (-\partial\overline{\Sigma}, \overline{\boldsymbol{\alpha}} \cap \partial\overline{\Sigma}, M, z)$, where M is the usual matching coming from the α-arcs. Let $\mathcal{A}(\mathcal{Z})$ be the associated algebra of this orientation-reversed pointed matched circle.

4.5.1 Definition of the type D module

The type D module, like the type A module, is defined by counting holomorphic curves satisfying certain asymptotic conditions at east infinity. As an $\mathbb{F}_2$-vector space, the two modules are defined identically, but the $\mathcal{A}(\mathcal{Z})$-module action is different. It will be a genuine differential module, instead of an $\mathcal{A}_\infty$ module. Furthermore, it will be a left module, rather than a right module.

Let $X(\mathcal{H})$ denote the $\mathbb{F}_2$-vector space generated by $\mathfrak{S}(\mathcal{H})$. We define the idempotent $I_D(\mathbf{x})$ to be exactly the "opposite" of $I_A(\mathbf{x})$ in the sense that $I_D(\mathbf{x}) := I(\{1, \dots, 2k\} \setminus o(\mathbf{x})) \in \mathcal{A}(\mathcal{Z}, 0)$. The *left* action of $\mathcal{I}(\mathcal{Z})$ on $X(\mathcal{H})$ is given by

$$I(\mathbf{s}) \cdot \mathbf{x} := \begin{cases} I_D(\mathbf{x}) & \text{if } \mathbf{s} = \{1, \dots, 2k\} \setminus o(\mathbf{x}), \\ 0 & \text{otherwise.} \end{cases}$$

Then we define $\widehat{CFD}(\mathcal{H})$ by

$$\widehat{CFD}(\mathcal{H}) = \mathcal{A}(\mathcal{Z}) \otimes X(\mathcal{H}),$$

where the tensor product is again over $\mathbf{k} = \mathcal{I}(\mathcal{Z})$. In particular, $\widehat{CFD}(\mathcal{H})$ is essentially free over $\mathcal{A}(\mathcal{Z})$ and has a very simple module structure:

$$a \cdot (b \otimes \mathbf{x}) = (ab) \otimes \mathbf{x}.$$

As in the $\widehat{CFA}$ case, the only summand $\mathcal{A}(\mathcal{Z}, i)$ which acts nontrivially on $\widehat{CFD}$ is the summand $i = 0$.

Recall that $\widehat{CFA}$ was an $\mathcal{A}_\infty$ module whose multiplication maps were defined by counting holomorphic curves in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\boldsymbol{\rho}})$, where $\vec{\boldsymbol{\rho}}$ was a sequence of nonempty sets of Reeb chords. For $\widehat{CFD}$, we only need to define the differential, as there are no higher multiplications. To define the differential, we only count curves in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\rho})$, where $\vec{\rho}$ is a sequence of *one-element* sets of Reeb chords.

If $\vec{\rho} = (\{\rho_1\}, \dots, \{\rho_n\})$ is a sequence of one-element sets, then let $a(\vec{\rho}) := a(\rho_1) \dots a(\rho_n)$ denote the product of the algebra elements associated to the one-element sets $\{\rho_i\}$. Note that this is, in general, not equal to $a(\vec{\boldsymbol{\rho}}) = a(\{\rho_1, \dots, \rho_n\})$. Furthermore, let $-\vec{\rho} = (\{-\rho_1\}, \dots, \{-\rho_n\})$ denote the sequence of chords with reversed orientation. Recall that if ρ_i is a Reeb chord of $\partial\mathcal{H}$, then $-\rho_i$ is a Reeb chord of $\mathcal{Z} = -\partial\mathcal{H}$.

Definition 4.33. As in the $\widehat{CFA}$ case, let J be a fixed almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$ which is admissible and which achieves transversality. Let $\mathbf{x}, y \in \mathfrak{S}(\mathcal{H})$ be generators, and let $B \in \pi(\mathbf{x}, \mathbf{y})$ connect the two. Define the coefficient

$$a^B_{\mathbf{x},\mathbf{y}} := \sum_{\substack{\operatorname{ind}(B,\vec{\rho})=1 \\ (B,\vec{\rho}) \text{ compatible}}} \#\left(\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\rho})\right) a(-\vec{\rho}) \in \mathcal{A}(\mathcal{Z}).$$

Then the differential on $\widehat{CFD}(\mathcal{H})$ is defined on basis elements by

$$\partial(\mathbf{I}\otimes\mathbf{x}) \coloneqq \sum_{\mathbf{y}\in\mathfrak{S}(\mathcal{H})}\sum_{B\in\pi_2(\mathbf{x},\mathbf{y})} a^B_{\mathbf{x},\mathbf{y}}\otimes\mathbf{y}.$$

For convenience, we often write $\partial\mathbf{x}$ instead of $\partial(\mathbf{I}\otimes\mathbf{x})$. Extending by linearity and the Leibniz rule

$$\partial(a\otimes\mathbf{x}) = (\partial a)\otimes\mathbf{x} + a\otimes(\partial\mathbf{x})$$

gives a map $\partial : \widehat{CFD}(\mathcal{H}) \to \widehat{CFD}(\mathcal{H})$.

Thus far we have defined $\widehat{CFD}(\mathcal{H})$ only as a differential module. (In fact, we have not yet proven that $\partial^2 = 0$. See Section 4.5.2.) In fact, the $\mathcal{I}(\mathcal{Z})$-module $X(\mathcal{H})$ comes with a map

$$\begin{aligned}\delta^1_{X(\mathcal{H})} : X(\mathcal{H}) &\to \mathcal{A}(\mathcal{Z})\otimes X(\mathcal{H}) = \widehat{CFD}(\mathcal{H})\\ \mathbf{x} &\mapsto \partial(\mathbf{I}\otimes\mathbf{x}).\end{aligned}$$

Since $\widehat{CFD}(\mathcal{H})$ is a differential module, Example 4.30 implies that $(X(\mathcal{H}), \delta^1)$ defines a type D structure over $\mathcal{A}(\mathcal{Z})$ with base ring $\mathcal{I}(\mathcal{Z})$. Then Proposition 4.31 implies that $\widehat{CFD}(\mathcal{H})$ is the differential module defined by this type D structure. It turns out that $\widehat{CFD}(\mathcal{H})$ is invariant up to homotopy not only as a module, but in fact as a type D structure.

Lemma 4.34. *If $\mathcal{H}$ is provincially admissible, then the boundary map ∂ is well defined. In particular, the sum in the definition of ∂ is finite for every $\mathbf{x}\in\mathfrak{S}(\mathcal{H})$. Furthermore, if $\mathcal{H}$ is in fact admissible, then the map δ^1 is bounded.*

Proof. First, observe by Lemma 3.11 that any B with $\mathcal{M}^B(\mathbf{x},\mathbf{y};\vec{\rho})$ must have B positive. As in the type A case, we may now apply Proposition 2.23, which says that there are only finitely many classes $B\in\pi_2(\mathbf{x},\mathbf{y})$ for a given $\vec{\rho}$. Since there are only finitely many generators $\mathbf{y}\in\mathfrak{S}(\mathcal{H})$ and algebra elements $a\in\mathcal{A}(\mathcal{Z})$, it suffices for the first part to simply check that for any given a, there are only finitely many ways to write it as the product of Reeb chords, so that we sum over finitely many $\vec{\rho}$. But this is certainly true, and follows from the strand diagram interpretation of elements of $\mathcal{A}(\mathcal{Z})$. This proves the first part of the lemma.

Note that the coefficient of $\mathbf{y}\in\mathfrak{S}(\mathcal{H})$ in $\delta^k(\mathbf{x})$ counts elements in all products of the form

$$\prod_{i=1}^{k}\mathcal{M}^{B_i}(\mathbf{x}_i,\mathbf{x}_{i+1};\vec{\rho}_i),$$

where $\mathbf{x}_1=\mathbf{x}$ and $\mathbf{x}_{k+1}=\mathbf{y}$. Thus any nonzero $\delta^k(\mathbf{x})$ term corresponds to a sequence of generators $\{\mathbf{x}_1,\dots,\mathbf{x}_{k+1}\}$ and, by Lemma 3.11 again, *positive* homology classes $B_i\in\pi_2(\mathbf{x}_i,\mathbf{x}_{i+1})$. Let $B\coloneqq\sum B_i\in\pi_2(\mathbf{x},\mathbf{y})$. This is a positive homology class. The sum $|B|$ of its coefficients is at least k. Since there are only finitely many positive homology classes, thanks to Proposition 2.24, it follows that there is a maximum $|B|$ over all positive B, hence an upper bound on k. Thus $\delta^k=0$ for all but finitely many k. □

Remark 4.35. As with type A modules, we may decompose the type D module along spin^c structures on Y. Thus we may write $\widehat{CFD}(\mathcal{H}) = \bigoplus \widehat{CFD}(\mathcal{H},\mathfrak{s})$, which in theory would help us obtain a grading of $\widehat{CFD}(\mathcal{H})$. Indeed, with some more work, we can show that $\widehat{CFD}(\mathcal{H})$ is a differential graded module, and not just a differential module.

4.5.2 $\partial^2 = 0$

While we have been calling $\widehat{CFD}(\mathcal{H})$ the type D module, we have not yet even shown that it is actually a differential module. We will show the following statement in this section.

Theorem 4.36. *The boundary operator ∂ on $\widehat{CFD}(\mathcal{H})$ satisfies $\partial^2 = 0$.*

As with type A modules, we first discuss a few illustrative examples. Before doing so, we briefly spell out what exactly $\partial^2 = 0$ means in this context. We have explicitly defined ∂ on a generator $\mathbf{x}$ (or, more precisely, on $\mathbf{I} \otimes \mathbf{x}$). If $a \in \mathcal{A}(\mathcal{Z})$ is an arbitrary algebra element, then

$$\partial(a \otimes \mathbf{x}) = (\partial a) \otimes \mathbf{x} + a \otimes \partial(\mathbf{I} \otimes \mathbf{x}) = (\partial a) \otimes \mathbf{x} + a \cdot \left(\sum_{\mathbf{y} \in \mathfrak{S}(\mathcal{H})} a_{\mathbf{x},\mathbf{y}} \otimes \mathbf{y} \right),$$

where $a_{\mathbf{x},\mathbf{y}} \coloneqq \sum_{B \in \pi_2(\mathbf{x},\mathbf{y})} a^B_{\mathbf{x},\mathbf{y}}$. In particular, it follows that

$$\begin{aligned} \partial^2(a \otimes \mathbf{x}) &= (\partial^2 a) \otimes \mathbf{x} + 2(\partial a) \otimes \partial(\mathbf{I} \otimes \mathbf{x}) + a \otimes \partial^2(\mathbf{I} \otimes \mathbf{x}) \\ &= (\partial^2 a) \otimes \mathbf{x} + 2(\partial a)(\partial \mathbf{x}) + a \cdot \left(\sum_{\mathbf{y} \in \mathfrak{S}(\mathcal{H})} \left(\partial a_{\mathbf{x},\mathbf{y}} \otimes \mathbf{y} + \sum_{\mathbf{w} \in \mathfrak{S}(\mathcal{H})} (a_{\mathbf{x},\mathbf{y}} \cdot a_{\mathbf{y},\mathbf{w}}) \otimes \mathbf{w} \right) \right). \end{aligned}$$

The first term vanishes because $\mathcal{A}(\mathcal{Z})$ is a differential module. The second term vanishes because we are working in characteristic two. As such, to show $\partial^2 = 0$, it is sufficient to show that

$$\partial a_{\mathbf{x},\mathbf{y}} + \sum_{\mathbf{w} \in \mathfrak{S}(\mathcal{H})} a_{\mathbf{x},\mathbf{w}} a_{\mathbf{w},\mathbf{y}} = 0 \tag{4.6}$$

for all $\mathbf{x}, \mathbf{y} \in \mathfrak{S}(\mathcal{H})$.

Example 4.37. In Figure 4.10, we have drawn the Heegaard diagram to the right of the pointed matched circle, unlike in Examples 3.55 and 4.17. This is because $\widehat{CFD}$ is defined via an orientation reversal of the

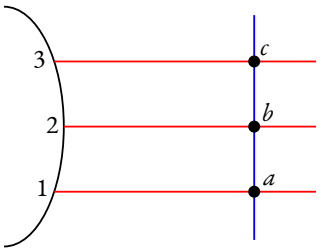

Figure 4.10: The local picture for Example 4.37.

pointed matched circle. (And, in fact, later on in Section 4.6, we will primarily be interested in cutting a closed 3-manifold into two bordered manifolds, and will compute $\widehat{CFA}$ of the left half and $\widehat{CFD}$ of the right half.)

We have, as before, three generators $\{a\}$, $\{b\}$, and $\{c\}$. Remember that we have orientation-reversed Reeb chords. Thus our boundary maps are

$$\partial\{c\} = \left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right]\{b\} + \left[\begin{smallmatrix} 3 & 2 \\ 1 & \end{smallmatrix}\right]\{a\}$$
$$\partial\{b\} = \left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right]\{a\}$$

Note that we omit tensor products. In particular, technically we should write $\partial\{b\} = \left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right] \otimes \{a\}$. Furthermore, note that, unlike in the $\widehat{CFA}$ case, the only nontrivial multiplications occur when the strands algebra element we keep track of which α-arcs are *not* occupied by the generator. We may verify directly that $\partial^2 = 0$: Note that

$$\partial^2\{c\} = \left(\partial\left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right]\right)\{b\} + \left(\partial\left[\begin{smallmatrix} 3 & 2 \\ 1 & \end{smallmatrix}\right]\right)\{a\} + \left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right]\partial\{b\} + \left[\begin{smallmatrix} 3 & 2 \\ 1 & \end{smallmatrix}\right]\partial\{a\}.$$

This last term vanishes because $\partial\{a\} = 0$. The first term vanishes because the strand diagram corresponding to $\left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right]$ has no crossings. Finally, because

$$\partial\left[\begin{smallmatrix} 3 & 2 \\ 1 & \end{smallmatrix}\right] = \left[\begin{smallmatrix} 2 & 3 \\ 1 & 2 \end{smallmatrix}\right] = \left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right] \cdot \left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right],$$

it follows that the two middle terms cancel out.

As in the $\widehat{CFA}$ case, there is a geometric interpretation of this in terms of ends of an index-two moduli space. Consider the same moduli space as in Examples 3.55 and 4.17, namely $\mathcal{M}^B(\{a\}, \{c\}; \{\rho_{12}\}, \{\rho_{23}\})$. Recall that this moduli space has a two-story end and a split curve end. The two-story end corresponds to the third term in $\partial^2\{c\}$, namely

$$\left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right]\partial\{b\} = \left[\begin{smallmatrix} 3 & 1 \\ 2 & \end{smallmatrix}\right] \cdot \left(\left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right]\{a\}\right).$$

Similarly, the split curve end (i.e., collision of levels) corresponds to the second term in $\partial^2\{c\}$, i.e., when the curve approaches ρ_{12} and ρ_{23} at the same time. In particular, the component $\left[\begin{smallmatrix} 3 & 2 \\ 1 & \end{smallmatrix}\right]\{a\}$ of $\partial\{c\}$ corresponds to the projection onto Σ, which approaches the Reeb chord ρ_{13} at east infinity. The differential $\partial\left[\begin{smallmatrix} 3 & 2 \\ 1 & \end{smallmatrix}\right]$ in the strands algebra tells us that we have degenerated a split component with west puncture ρ_{13} and east punctures ρ_{12} and ρ_{23}. (Note that, in the diagram, east infinity is at the west, because we have flipped the diagram around.)

Example 4.38. The type D module of the bordered Heegaard diagram in Figure 4.11 has boundary maps

$$\partial\{a, d\} = \left[\begin{smallmatrix} 4 & 2 \\ 3 & \end{smallmatrix}\right]\{a, c\}$$
$$\partial\{b, c\} = \left[\begin{smallmatrix} 2 & 4 \\ 1 & \end{smallmatrix}\right]\{a, c\}$$
$$\partial\{b, d\} = \left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right]\{a, d\} + \left[\begin{smallmatrix} 4 & 1 \\ 3 & \end{smallmatrix}\right]\{b, c\}.$$

To show $\partial^2 = 0$, it suffices to check that

$$0 = \partial^2\{b, d\} = \partial\left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right]\{a, d\} + \left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right] \cdot \left(\left[\begin{smallmatrix} 4 & 2 \\ 3 & \end{smallmatrix}\right]\{a, c\}\right) + \partial\left[\begin{smallmatrix} 4 & 1 \\ 3 & \end{smallmatrix}\right]\{b, c\} + \left[\begin{smallmatrix} 4 & 1 \\ 3 & \end{smallmatrix}\right] \cdot \left(\left[\begin{smallmatrix} 2 & 4 \\ 1 & \end{smallmatrix}\right]\{a, c\}\right). \quad (4.7)$$

Note that the first and third term vanish. Thus $\partial^2 = 0$ follows from the fact that $a(\rho_{34})$ and $a(\rho_{12})$ commute in the strands algebra, i.e., that

$$\left[\begin{smallmatrix} 2 & 3 \\ 1 & \end{smallmatrix}\right]\left[\begin{smallmatrix} 4 & 2 \\ 3 & \end{smallmatrix}\right] = \left[\begin{smallmatrix} 2 & 4 \\ 1 & 3 \end{smallmatrix}\right] = \left[\begin{smallmatrix} 4 & 1 \\ 3 & \end{smallmatrix}\right]\left[\begin{smallmatrix} 2 & 4 \\ 1 & \end{smallmatrix}\right].$$

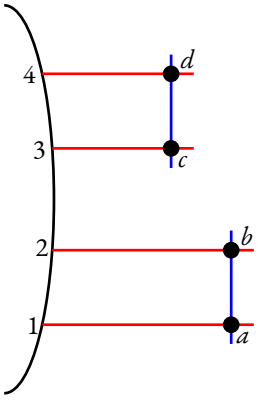

Figure 4.11: The local picture for Example 4.38.

We again have a geometric interpretation for this result. Consider the same moduli space as in Examples 3.56 and 4.18, namely $\mathcal{M}^B(\{a,c\},\{b,d\};\rho_{12},\rho_{34})$. The ends appear when $\mathrm{ev}_{34}-\mathrm{ev}_{12}$ approaches 0 or ∞. But now also consider the moduli space where $\mathrm{ev}_{34}<\mathrm{ev}_{12}$, i.e., $\mathcal{M}^B(\{a,c\},\{b,d\};\rho_{34},\rho_{12})$. This has a two-story end when $\mathrm{ev}_{12}-\mathrm{ev}_{34}\to\infty$, as well as a collision of levels end when $\mathrm{ev}_{12}-\mathrm{ev}_{34}\to 0$. The total number of ends of these two moduli spaces should certainly be even. The two collisions of levels cancel each other out, since they are identical curves. In particular, they both have asymptotics $\{\rho_{12},\rho_{34}\}$ occurring at the same time. The two-story ends, on the other hand, occur as the first and third terms in Equation (4.7).

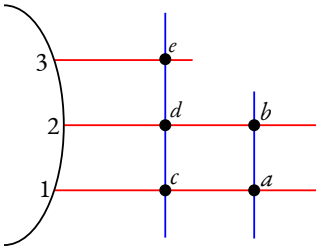

Figure 4.12: The local picture for Example 4.39.

Example 4.39. Finally, consider the bordered Heegaard diagram in Figure 4.12. We have the nontrivial differentials

$$\begin{aligned}\partial\{b,c\}&=\{a,d\}\\ \partial\{b,e\}&=\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]\{a,e\}+\left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]\{b,c\}\\ \partial\{a,e\}&=\left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]\{a,d\}.\end{aligned}$$

Since the algebra elements $\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]$ and $\left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]$ are closed in the strands algebra (i.e., the boundary operator takes them to zero), we know that

$$\partial^2\{b,e\}=\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]\partial\{a,e\}+\left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]\partial\{b,c\}=\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]\cdot\left(\left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]\{a,d\}\right)+\left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]\{a,d\}=0$$

since $\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]\left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]=\left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]$.

Geometrically, this again corresponds to counting the ends of two different moduli spaces. The first is the moduli space from Examples 3.57 and 4.19, namely the one connecting $\{b, e\}$ to $\{a, d\}$ via the Reeb chord ρ_{13}. The second has east asymptotics given by (ρ_{12}, ρ_{23}). It is parameterized by the difference in evaluations $\mathrm{ev}_{23} - \mathrm{ev}_{12}$, and has a collision end and a two-story end. (Note that, this time, we do not have a moduli space with east asymptotics given by (ρ_{23}, ρ_{12}), since the resulting pair is not strongly boundary monotone.) The join curve end of the first family is the same as the collision end of the second family, as they both have asymptotics $\{\rho_{12}, \rho_{23}\}$. In particular, they both approach the Reeb chords ρ_{12} and ρ_{23} at the same time. The two-story end of the first family corresponds to $\left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]\{a, d\}$, since the first story is the provincial rectangle with corners $abcd$, while the second story is the rectangle $13ec$. Finally, the two-story end of the second family corresponds to $\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right] \cdot \left(\left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]\{a, d\}\right)$, with stories projecting to the rectangles $12ba$ and $23ed$.

Example 4.40. Let $\mathcal{H}$ be the Heegaard diagram for the genus-1 handlebody with a single generator x_0, as in Figure 4.13. As before, there is only one region which does not cross the basepoint z. There is a unique

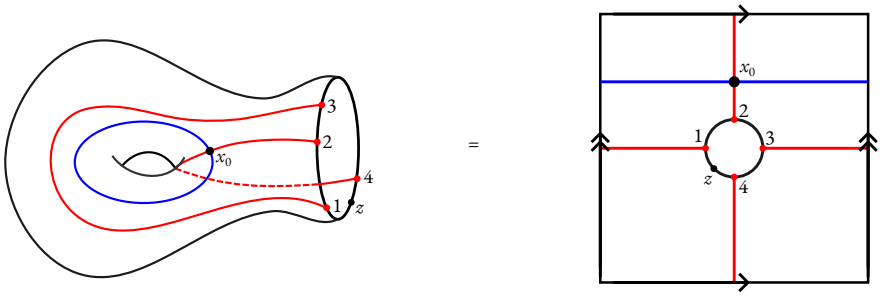

Figure 4.13: The genus-1 handlebody (again).

element of $\mathcal{M}^B(x_0, x_0; \rho_{23}, \rho_{12})$; this is the same as the element which is counted in $m_3(x_0, \rho_{23}, \rho_{12})$ in the $\widehat{CFA}$ case. The associated algebra element is

$$a(\rho_{32})a(\rho_{21}) = \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]$$

Note that the Reeb chords ρ_{32} and ρ_{21} are reversed since the type D module is a module over the strands algebra associated to $\mathcal{Z} = -\partial\mathcal{H}$. Recall that there are also holomorphic curves connecting x_0 to itself with east asymptotics given by $\vec{\rho} = (\rho_{23}, \rho_{13}, \ldots, \rho_{13}, \rho_{12})$. But the associated algebra element $a(-\vec{\rho})$ vanishes because the Reeb chords do not compose. (In fact, this issue arises because $\mathcal{H}$ is not admissible.) Thus we conclude that $\widehat{CFD}(\mathcal{H})$ is generated by x_0 with differential

$$\partial x_0 = \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right] x_0.$$

The associated type D structure is the map δ^1 on the $\mathbb{F}_2$-vector space $X(\mathcal{H})$ generated by $\mathfrak{S}(\mathcal{H}) = \{x_0\}$ which is defined by

$$\delta^1(x_0) = \partial x_0 = \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right] x_0.$$

This map is not bounded since

$$\delta^2(x_0) = \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right] \otimes \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right] \otimes x_0,$$

and so on for higher iterates δ^i. Notice, after all, that Lemma 4.34 only guarantees that δ^1 is bounded if the bordered Heegaard diagram is admissible.

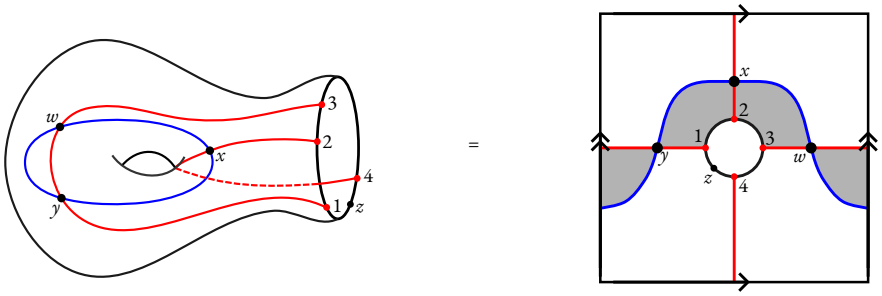

Figure 4.14: The deformed genus 1 handlebody (again).

Example 4.41. Consider the admissible diagram in Figure 4.14. (Note that the previous example was only a *provincially* admissible diagram.) There are no curves which connect w at $t = -\infty$ to another generator, so $\partial w = 0$. The only curve which limits at $t = -\infty$ to x is the one in $\mathcal{M}^B(x, w; \rho_{23})$, and so $\partial x = \left[\begin{smallmatrix} 3 \\ 2 \end{smallmatrix}\right] w$. Note that we have one provincial domain connecting y to w, as well as a domain connecting y to w via the Reeb chord ρ_{13}. Finally, the moduli space $\mathcal{M}^B(y, x; \rho_{12})$ has one element. Putting this together, we have

$$\begin{aligned} \partial w &= 0 \\ \partial x &= \left[\begin{smallmatrix} 3 \\ 2 \end{smallmatrix}\right] w \\ \partial y &= w + \left[\begin{smallmatrix} 3 \\ 1 \end{smallmatrix}\right] w + \left[\begin{smallmatrix} 2 \\ 1 \end{smallmatrix}\right] x. \end{aligned}$$

This time around, the map δ^1 which takes a generator to its differential is bounded. After all, we have

$$\delta^2 x = \left[\begin{smallmatrix} 3 \\ 2 \end{smallmatrix}\right] \partial w = 0.$$

Similarly, we have $\delta^3 y = 0$, and so all higher iterates δ^i for $i \geq 3$ must vanish.

Based on these examples, we see that, unlike in the compatibility equation for the type A module, showing that $\partial^2 = 0$ for the type D module often requires that we count ends of multiple moduli spaces. Furthermore, these ends sometimes cancel with each other. Because of this, the proof of Theorem 4.36 is somewhat more complicated than that of Theorem 4.16. However, the basic idea of counting ends of the compactified moduli spaces remains the same.

Proof of Theorem 4.36. Recall that it is enough to prove Equation (4.6) where $a_{\mathbf{x},\mathbf{y}} = \sum_{B \in \pi_2(\mathbf{x},\mathbf{y})} a^B_{\mathbf{x},\mathbf{y}}$. Fixing a homology class $B \in \pi_2(\mathbf{x}, \mathbf{y})$, then, it is enough to show that

$$\partial a^B_{\mathbf{x},\mathbf{y}} + \sum_{\substack{\mathbf{w} \in \mathfrak{S}(\mathcal{H}) \\ B_1 * B_2 = B}} a^{B_1}_{\mathbf{x},\mathbf{w}} a^{B_2}_{\mathbf{w},\mathbf{y}} = 0.$$

Let a be an algebra element, and let $\vec{\rho}$ range over all vectors with $a(\vec{\rho})$. Then the total number of two-story ends of $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \vec{\rho})$, summed over all choices of $\vec{\rho}$, is equal to the coefficient of a in the second term $\sum a^{B_1}_{\mathbf{x},\mathbf{w}} a^{B_2}_{\mathbf{w},\mathbf{y}}$ above. One may also show (though this takes a little more work) that the coefficient of a in $\partial a^B_{\mathbf{x},\mathbf{y}}$ is given by the total number of split curve ends. Note, however, that collisions of levels which do not result in a split curve end, i.e., elements of

$$\mathcal{M}^B(\mathbf{x}, \mathbf{y}; \rho_1, \dots, \rho_{i-1}, \{\rho_i, \rho_{i+1}\}, \rho_{i+2}, \dots, \rho_n)$$

with $\rho_i^+ \neq \rho_{i+1}^-$, are not counted in this sum.

Thus it remains to show that the total contribution from join curve ends, shuffle curve ends, and collisions of levels which are not split curve ends is zero. Note first that shuffle curve ends do not exist since the partition $\vec{\rho}$ is discrete. One can determine that the only collisions of levels which may appear are the following:

- $\rho_i^- = \rho_{i+1}^+$: In this case, these moduli spaces are exactly the ones which also degenerate off join curve ends via the factorization $a(-\rho_1)\dots a(-(\rho_i \uplus \rho_{i+1}))\dots a(-\rho_n)$ of a. Thus collision ends coming from this case cancel with join curve ends.
- $\{M(\rho_i^-), M(\rho_i^+)\} \cap M(\rho_{i+1}^-), M(\rho_{i+1}^+) = \emptyset$ and the chords ρ_i and ρ_{i+1} are either nested (in either order) or disjoint: There is then another factorization of a given by swapping the order of the $a(-\rho_i)$ and $a(-\rho_{i+1})$ factors, so this case contains collision ends which appear in pairs. Thus this case cancels with itself.

We conclude that there is no contribution from other moduli space ends, and so Theorem 3.53 implies the equation $\partial^2 = 0$, as desired. □

4.5.3 Invariance

We thus far have a differential module $(\widehat{CFD}(\mathcal{H}), \partial)$ associated to a bordered Heegaard diagram. As seen in Examples 4.40 and 4.41, two $\widehat{CFD}$ modules of bordered Heegaard diagrams which represent the same 3-manifold are not necessarily isomorphic as differential modules. But, as in the type A case, the two type D modules are homotopy equivalent. In particular, we have the following statement.

Theorem 4.42. *Let $\mathcal{H}$ be a bordered Heegaard diagram and let $\mathcal{Z} = -\partial\mathcal{H}$. The differential module $\widehat{CFD}(\mathcal{H})$ is independent, up to homotopy equivalence of differential $\mathcal{A}(\mathcal{Z})$-modules, of the choice of admissible, transversality-achieving almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$. Furthermore, if $\mathcal{H}$ and $\mathcal{H}'$ are provincially admissible bordered Heegaard diagrams for the same bordered 3-manifold $(Y, -\mathcal{Z}, \phi : -F(\mathcal{Z}) \to \partial Y)$, then $\widehat{CFD}(\mathcal{H})$ and $\widehat{CFD}(\mathcal{H}')$ are homotopy equivalent.*

As in the case for the type A module, the proof is somewhat technical, but somewhat less related to our discussion of moduli spaces. To show invariance under isotopies and choice of almost complex structure, one uses something called a "continuation map," as in [Flo89a]. Invariance under Heegaard moves, particularly a handleslide of an α-arc over an α-circle, requires the definition of a new "moduli space of triangles." See [LOT18, Section 6.3] for a detailed proof.

We give one example of invariance.

Example 4.43. As usual, let $\mathcal{H}$ and $\mathcal{H}'$ denote the bordered Heegaard diagrams from Examples 4.40 and 4.41, respectively. Again, these differ only by an isotopy, so we would expect a homotopy equivalence between their type D modules.

Constructing this homotopy equivalence is easier than constructing a homotopy equivalence between their type A modules (cf. Example 4.25), since we no longer have higher multiplications to worry about. Consider the map

$$\begin{aligned} f : \widehat{CFD}(\mathcal{H}) &\to \widehat{CFD}(\mathcal{H}') \\ x_0 &\mapsto x + \left[\tfrac{3}{2}\right] y + \left[\tfrac{3}{2}\right] w. \end{aligned}$$

This is a chain map because

$$f(\partial x_0) = \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]x = \partial\left(x + \left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]y + \left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]w\right) = \partial f(x_0).$$

On the other hand, define

$$\begin{aligned} g : \widehat{CFD}(\mathcal{H}') &\to \widehat{CFD}(\mathcal{H}) \\ x &\mapsto x_0 \\ y &\mapsto \left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]x_0 \\ w &\mapsto \left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]x_0. \end{aligned}$$

It is not too hard to verify that this is also a chain map. The most interesting case is

$$g(\partial x) = g\left(\left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]w\right) = \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]x_0 = \partial x_0 = \partial g(x).$$

Clearly $g \circ f$ is the identity on $\widehat{CFD}(\mathcal{H})$. It is not too hard to construct a homotopy equivalence between $f \circ g$ and the identity map on $\widehat{CFD}(\mathcal{H}')$ either. Indeed, the chain homotopy defined by

$$\begin{aligned} x &\mapsto \left[\begin{smallmatrix}3\\1\end{smallmatrix}\right]x + \left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]y + \left[\begin{smallmatrix}3\\2\end{smallmatrix}\right]w \\ y &\mapsto \left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]x + w \\ w &\mapsto \left[\begin{smallmatrix}2\\1\end{smallmatrix}\right]x + y \end{aligned}$$

works.

It follows, then, that $\widehat{CFD}(\mathcal{H}) \simeq \widehat{CFD}(\mathcal{H}')$, as desired.

4.6 The pairing theorem

Notice that slicing a closed 3-manifold Y along a separating surface gives a decomposition $Y = Y_1 \cup_\partial Y_2$ into two bordered manifolds. We might wonder whether there is a relationship between $\widehat{HF}(Y)$ and the bordered Heegaard Floer invariants of Y_1 and Y_2. In fact, bordered Heegaard Floer homology gives a way to compute $\widehat{HF}(Y)$, thanks to the following pairing theorem.

Theorem 4.44. *Let Y_1 and Y_2 be bordered 3-manifolds with $\partial Y_1 = F(\mathcal{Z}) = -\partial Y_2$ for some pointed matched circle $\mathcal{Z}$. If Y is the closed 3-manifold obtained by gluing Y_1 and Y_2 together along $F(\mathcal{Z})$, then $\widehat{CF}(Y)$ is homotopy equivalent to $\widehat{CFA}(Y_1)\widetilde{\otimes}\widehat{CFD}(Y_2)$, where $\widetilde{\otimes}$ denotes the $\mathcal{A}_\infty$ tensor product. In particular, we have the following relationship between the Heegaard Floer homology of Y and the bordered invariants of Y_1 and Y_2:*

$$\widehat{HF}(Y) \cong H_*\left(\widehat{CFA}(Y_1)\widetilde{\otimes}\widehat{CFD}(Y_2)\right).$$

The first step is to note that, in the setup of the pairing theorem, gluing Y_1 and Y_2 into the closed manifold Y is equivalent to gluing their respective bordered Heegaard diagrams. In particular, suppose $\mathcal{H}_1 = (\overline{\Sigma}_1, \overline{\boldsymbol{\alpha}}_1, \boldsymbol{\beta}_1, z)$ and $\mathcal{H}_2 = (\overline{\Sigma}_2, \overline{\boldsymbol{\alpha}}_2, \boldsymbol{\beta}_2, z)$ are bordered Heegaard diagrams for Y_1 and Y_2, respectively, with $\partial\mathcal{H}_1 = \mathcal{Z} = -\partial\mathcal{H}_2$. Let $\Sigma = \overline{\Sigma}_1 \cup_\partial \overline{\Sigma}_2$, and similarly for $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. Then $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ is a closed

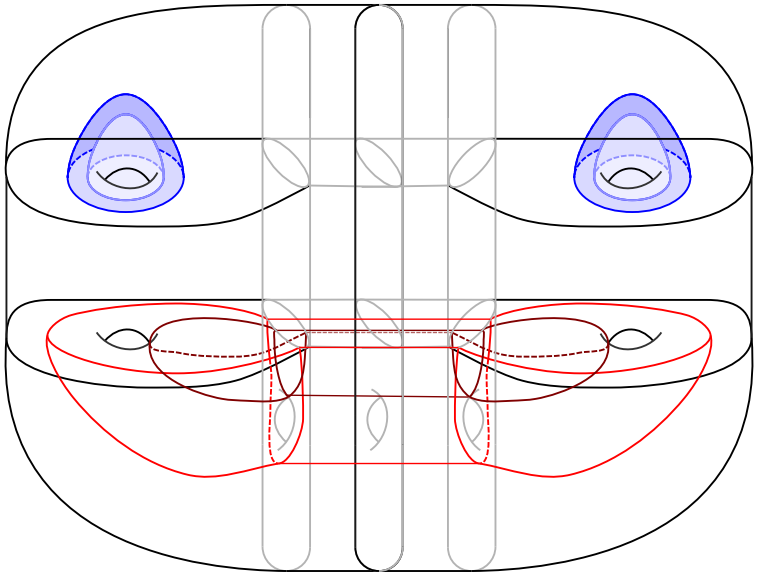

Figure 4.15: Gluing along the shared boundary makes it clear that gluing bordered Heegaard diagrams is the same as gluing the corresponding bordered 3-manifolds.

Heegaard diagram for Y, which we may write as $\mathcal{H} = \mathcal{H}_1 \cup_\partial \mathcal{H}_2$. Recalling the Morse theoretic description of a (bordered) manifold from its (bordered) Heegaard diagram in Section 2.3, this intuitively looks like Figure 4.15. (Similarly, cutting a closed Heegaard diagram at some separating circle $\mathcal{Z}$ corresponds to cutting a closed 3-manifold at the surface $F(\mathcal{Z})$.)

There is a relationship between the generators of these diagrams. Suppose $\mathbf{x}_1 \in \mathfrak{S}(\mathcal{H}_1)$ and $\mathbf{x}_2 \in \mathfrak{S}(\mathcal{H}_2)$. If $o(\mathbf{x}_1) \cap o(\mathbf{x}_2) = \emptyset$, so that $\mathbf{x}_1 \cup \mathbf{x}_2$ is a generator of $\mathcal{H}$, then we call $(\mathbf{x}_1, \mathbf{x}_2)$ a **compatible pair**. Let $\mathfrak{S}(\mathcal{H}_1, \mathcal{H}_2) \subset \mathfrak{S}(\mathcal{H}_1) \times \mathfrak{S}(\mathcal{H}_2)$ be the set of compatible pairs of generators. There is a bijection between $\mathfrak{S}(\mathcal{H}_1, \mathcal{H}_2)$ and $\mathfrak{S}(\mathcal{H})$. Furthermore, if $\mathbf{x}_1 \times \mathbf{x}_2$ and $\mathbf{y}_1 \times \mathbf{y}_2$ are both compatible pairs, then there is a natural identification of $\pi_2(\mathbf{x}_1 \cup \mathbf{x}_2, \mathbf{y}_1 \cup \mathbf{y}_2)$ with the subset of $\pi_2(\mathbf{x}_1, \mathbf{y}_1) \times \pi_2(\mathbf{x}_2, \mathbf{y}_2)$ consisting of pairs (B_1, B_2) with $\partial^\partial B_1 + \partial^\partial B_2 = 0$, i.e., with B_1 and B_2 "matching" at their shared boundary. Finally, it is not difficult to show that, if $\mathcal{H}_1$ is admissible and $\mathcal{H}_2$ is provincially admissible, then $\mathcal{H}$ is admissible as a closed Heegaard diagram. As such, let $\mathcal{H}_1$ be admissible and $\mathcal{H}_2$ provincially admissible, so that we may use $\mathcal{H}$ to compute the Heegaard Floer homology of Y.

Recall now that the Heegaard Floer homology of $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, z)$ counts holomorphic curves in $\Sigma \times [0,1] \times \mathbb{R}$ which connect generators. In theory, we can cut such a holomorphic curve at $\partial\Sigma_i \subset \Sigma$. (From the perspective of the complex structure on Σ, rather than simply cutting at the circle Z, one actually stretches the neck connecting Σ_1 to Σ_2, thus creating a pair of infinite cylindrical ends at east infinity for both Σ_1 and Σ_2. This setup is reminiscent of proofs of some of the product theorems in gauge theory, as in [Don02, KM07]) Thus we may consider a holomorphic curve in $\Sigma \times [0,1] \times \mathbb{R}$ as two maps, one onto $\Sigma_1 \times [0,1] \times \mathbb{R}$ and the other onto $\Sigma_2 \times [0,1] \times \mathbb{R}$. The restriction to Σ_i connects generators $\mathbf{x}_i$ and $\mathbf{y}_i$ of $\mathcal{H}_i$. In particular, if $B \in \pi_2(\mathbf{x}_1 \cup \mathbf{x}_2, \mathbf{y}_1 \cup \mathbf{y}_2)$ may be written $B_1 * B_2$, where B_i is the projection in $\pi_2(\mathbf{x}_i, \mathbf{y}_i)$, then there is a map

$$\mathcal{M}^B(\mathbf{x}_1 \cup \mathbf{x}_2, \mathbf{y}_1 \cup \mathbf{y}_2; S) \to \mathcal{M}^{B_1}(\mathbf{x}_1, \mathbf{y}_1; S_1^\triangleright) \times \mathcal{M}^{B_2}(\mathbf{x}_2, \mathbf{y}_2; S_2^\triangleright)$$

where $S_1^\triangleright$ and $S_2^\triangleright$ are compatible decorated sources which glue to S. More precisely, compatibility means that there is a bijection $\varphi : E(S_1^\triangleright) \to E(S_2^\triangleright)$ such that the Reeb chords which label q_i and $\varphi(q_i)$ are

orientation-reverses for each puncture $q_i \in E(S_1^\triangleright)$. That they glue to S simply means that we may obtain S by gluing $S_1^\triangleright$ to $S_2^\triangleright$ along identified neighborhoods of corresponding punctures. (In other words, S is the preglued surface obtained from $S_1^\triangleright$ and $S_2^\triangleright$.) In this case, we write $S = S_1^\triangleright \natural S_2^\triangleright$.

Of course, two holomorphic curves $u_1 \in \mathcal{M}^{B_1}(\mathbf{x}_1, \mathbf{y}_1; S_1^\triangleright)$ and $u_2 \in \mathcal{M}^{B_2}(\mathbf{x}_2, \mathbf{y}_2; S_2^\triangleright)$ only glue together if the $\mathrm{ev}(u_1) = \mathrm{ev}(u_2)$ under the correspondence defined by φ. In particular, for the curves to glue together, their corresponding east punctures must converge to the Reeb chords at the same time. As such, we may expect some relationship between the moduli space $\mathcal{M}^B(\mathbf{x}_1 \cup \mathbf{x}_2, \mathbf{y}_1 \cup \mathbf{y}_2; S)$ and the **moduli space of matched pairs**, i.e., the pullback

$$\widetilde{\mathcal{MM}}^B(\mathbf{x}_1, \mathbf{y}_1; S_1^\triangleright; \mathbf{x}_2, \mathbf{y}_2; S_2^\triangleright) := \widetilde{\mathcal{M}}^{B_1}(\mathbf{x}_1, \mathbf{y}_1; S_1^\triangleright) \times_{\mathrm{ev}_1 = \mathrm{ev}_2} \widetilde{\mathcal{M}}^{B_2}(\mathbf{x}_2, \mathbf{y}_2; S_2^\triangleright).$$

There is, as usual, an $\mathbb{R}$-action on the moduli space of matched pairs which is free unless both sides of the matching are trivial strips. The action here involves simultaneous translation of both curves along the t-direction. We may drop the tilde to indicate the quotient of $\widetilde{\mathcal{MM}}$ by this translation action. It turns out, by arguments similar to those in Chapter 3, that $\mathcal{MM}$ is a manifold with dimension given by a certain index formula.

Now we restrict, as with the moduli spaces in Chapter 3, to *embedded* curves. As usual, the condition that the curves are embedded is equivalent to a numerical condition on the Euler characteristic. In particular, let $\widetilde{\mathcal{MM}}^B(\mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2)$ be the union over all compatible decorated surfaces $S_1^\triangleright$ and $S_2^\triangleright$ with $\chi(S_1^\triangleright \natural S_2^\triangleright) = \chi_{\mathrm{emb}}(B)$, where $\chi_{\mathrm{emb}}(B)$ is as defined at the end of Section 3.2. Assuming $B \neq 0$, so that the $\mathbb{R}$-action is free, define the corresponding quotient to be $\mathcal{MM}^B(\mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2; \mathbf{y}_2)$.

If $\mathrm{ind}(B, S) = 1$, then it turns out that the number of elements in $\mathcal{M}^B(\mathbf{x}, \mathbf{y}; S)$ is equal (modulo 2) to the number of elements in $\bigcup \mathcal{MM}^B(\mathbf{x}_1, \mathbf{y}_1, S_1^\triangleright; \mathbf{x}_2, \mathbf{y}_2, S_2^\triangleright)$. The union here is taken over all $S = S_1^\triangleright \natural S_2^\triangleright$, $\mathbf{x} = \mathbf{x}_1 \cup \mathbf{x}_2$, and $\mathbf{y} = \mathbf{y}_1 \cup \mathbf{y}_2$. The details of the proof of this statement are reminiscent of our arguments in the previous chapter.

Thus we have the following proposition, which tells us that this moduli space of matched pairs, which is defined by looking at the bordered Heegaard diagrams $\mathcal{H}_1$ and $\mathcal{H}_2$, gives an equivalent characterization for the Heegaard Floer homology of $\mathcal{H}$.

Proposition 4.45. *Let $\widehat{CF}(\mathcal{H}_1, \mathcal{H}_2)$ be generated as an $\mathbb{F}_2$-vector space by the set $\mathfrak{S}(\mathcal{H}_1, \mathcal{H}_2)$ of compatible pairs of generators. Furthermore, define the map ∂ on $\widehat{CF}(\mathcal{H}_1, \mathcal{H}_2)$ by*

$$\partial(\mathbf{x}_1 \times \mathbf{x}_2) := \sum_{\mathbf{y}_1 \times \mathbf{y}_2 \in \mathfrak{S}(\mathcal{H}_1, \mathcal{H}_2)} \sum_{\substack{B \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mathrm{ind}(B) = 1}} \#\left(\mathcal{MM}^B(\mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2)\right) \cdot (\mathbf{y}_1 \times \mathbf{y}_2).$$

Then ∂ is a differential. Furthermore, for a generic choice of almost complex structure, the chain complex $\widehat{CF}(\mathcal{H}_1, \mathcal{H}_2)$ is isomorphic to the Heegaard Floer chain complex $\widehat{CF}(\mathcal{H}) = \widehat{CF}(\mathcal{H}_1 \cup_\partial \mathcal{H}_2)$.

Note that the sum defining $\partial(\mathbf{x}_1, \mathbf{x}_2)$ is finite since $\mathcal{H}$ is admissible, thanks to an argument similar to Lemma 4.14. (Recall that we assumed $\mathcal{H}_1$ was admissible and $\mathcal{H}_2$ was provincially admissible specifically so that $\mathcal{H}$ would be admissible.) Note that the counts in the differentials are exactly the same, modulo 2.

The main issue with this proposition is that the moduli space of matched pairs is a fiber product, which cannot be easily translated into the language of $\mathcal{A}_\infty$ modules, and thus does not fit with our existing algebraic framework. A workaround is to instead introduce a "time-dilated" version of the matching

between $\mathcal{M}^{B_1}(\mathbf{x}_1, \mathbf{y}_1)$ and $\mathcal{M}^{B_2}(\mathbf{x}_2, \mathbf{y}_2)$. In particular, define the **moduli space of T-matched pairs** to be

$$\widetilde{\mathcal{MM}}^B(T; \mathbf{x}_1, \mathbf{y}_1; S_1^{\triangleright}; \mathbf{x}_2, \mathbf{y}_2; S_2^{\triangleright}) := \widetilde{\mathcal{M}}^{B_1}(\mathbf{x}_1, \mathbf{y}_1; S_1^{\triangleright}) \times_{T \cdot \mathrm{ev}_1 = \mathrm{ev}_2} \widetilde{\mathcal{M}}^{B_2}(\mathbf{x}_2, \mathbf{y}_2; S_2^{\triangleright}).$$

In particular, we ask that the t-coordinates for u_1 at the punctures of $S_1^{\triangleright}$ match with those of u_2 at the punctures of $S_2^{\triangleright}$, up to a factor of T. This still has an $\mathbb{R}$-action, though we should translate u_1 by $T \cdot t$ and u_2 by t. We call the quotient by this action $\mathcal{MM}^B(T; \mathbf{x}_1, \mathbf{y}_1; S_1^{\triangleright}; \mathbf{x}_2, \mathbf{y}_2; S_2^{\triangleright})$. We have embedded moduli spaces as well, which we denote

$$\mathcal{MM}^B(T; \mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2) = \widetilde{\mathcal{MM}}^B(T; \mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2)/\mathbb{R}.$$

We may define $\widehat{CF}(T; \mathcal{H}_1, \mathcal{H}_2)$ to be the chain complex which is generated by $\mathfrak{S}(\mathcal{H}_1, \mathcal{H}_2)$ and has differential ∂_T given by counting elements of this moduli space, i.e.,

$$\partial_T(\mathbf{x}_1 \times \mathbf{x}_2) := \sum_{\mathbf{y}_1 \times \mathbf{y}_2 \in \mathfrak{S}(\mathcal{H}_1, \mathcal{H}_2)} \sum_{\substack{B \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mathrm{ind}(B) = 1}} \#\left(\mathcal{MM}^B(T; \mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2)\right) \cdot (\mathbf{y}_1 \times \mathbf{y}_2).$$

Showing that ∂_T is a differential requires that we extend our definition of holomorphic combs to so-called "T-matched combs." These are combs such that, at each story $(u_1, v_1, \dots, v_k, u_2)$, the two eastmost components v_k and u_2 are T-matched, in the sense that $T \cdot \mathrm{ev}_e(v_k) = \mathrm{ev}_w(u_2)$.

Proposition 4.46. *For any $T \in (0, \infty)$, the chain complex $\widehat{CF}(T; \mathcal{H}_1, \mathcal{H}_2)$ is chain homotopy equivalent to $\widehat{CF}(1; \mathcal{H}_1, \mathcal{H}_2) = \widehat{CF}(\mathcal{H}_1, \mathcal{H}_2)$, and hence to $\widehat{CF}(\mathcal{H})$.*

As T grows larger, the moduli space in the differential ∂_T of $\widehat{CF}(T; \mathcal{H}_1, \mathcal{H}_2)$ counts matched pairs where the left side (which will correspond to the $\widehat{CFA}$ side in the pairing theorem) converges to Reeb chords on a smaller and smaller interval of $\mathbb{R}$, as seen from the perspective of the right side (which will correspond to the $\widehat{CFD}$ side in the pairing theorem). That is, in the limit, Reeb chords begin to collide on the left side, while Reeb chords grow infinitely far apart on the right side.

One way to encode this is via simple ideal-matched combs. In particular, consider a simple holomorphic comb U_1 for $\mathcal{H}_1$ (i.e., U_1 has at most one story) and a toothless holomorphic comb U_2 for $\mathcal{H}_2$ (i.e., U_2 has no components at east infinity). Then a **simple ideal-matched comb** may be obtained by allowing each story in U_2 to occur at a single time t in U_1. In other words, if $\varphi : E(U_2) \to E(U_1)$ is the correspondence between punctures (hence between Reeb chords), and if p and q are two punctures on the same story of U_2, then the t-coordinates of $\varphi(p)$ and $\varphi(q)$ are the same. See Figure 4.16 for a schematic representation of this.

Remark 4.47. Technically, there are some more conditions which are needed to define simple ideal-matched combs (see [LOT18, Definition 9.28]), but the details are not too important for us, so we omit them here.

It turns out that T-matched curves converge to simple ideal-matched curves. If we "trim at east infinity," then the resulting moduli space of trimmed simple ideal-matched curves has exactly the same counts as $\mathcal{MM}^B(T; \mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2)$ for sufficiently large T. In particular, the moduli space in the definition of ∂_T may be replaced by the moduli space $\mathcal{MM}^B_{\mathrm{tsic}}(\mathbf{x}_1, \mathbf{y}_1; \mathbf{x}_2, \mathbf{y}_2)$ of these trimmed simple ideal-matched curves.

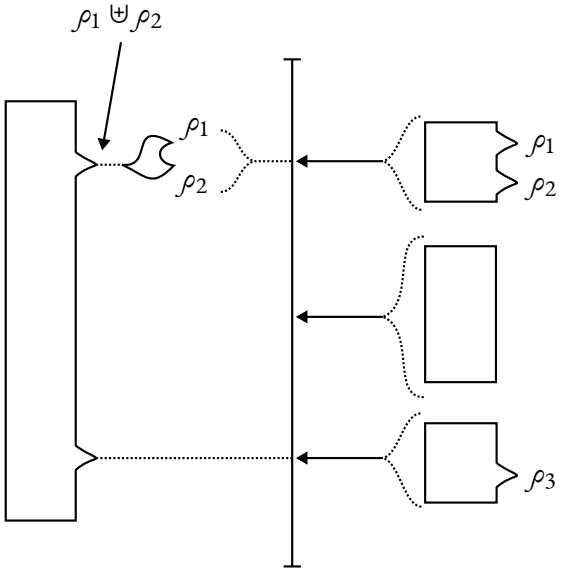

Figure 4.16: The one-story comb on the left is U_1, and is allowed to have components at east infinity. The three-story comb on the right is U_2, and has no components at east infinity. Each story of U_2 is mapped to a single point on the vertical line in the middle, i.e., to a single time t.

Algebraically, what this corresponds to is the following. Recall that we have a map $\delta^1_{X(\mathcal{H}_2)} : X(\mathcal{H}_2) \to \widehat{CFD}(\mathcal{H}_2)$ which is defined by sending a generator $\mathbf{x}_2 \in \mathfrak{S}(\mathcal{H}_2)$ to $\partial(\mathbf{I} \otimes \mathbf{x}_2)$. Graphically, the elements which the differential counts (i.e., the trimmed simple ideal-matched curves) correspond to maps of the form

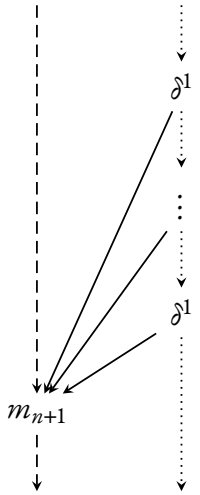

In particular, each story of the right side U_2 of a trimmed simple ideal-matched curve corresponds to taking some repeated differential coming from the type D structure. This outputs an element of $\mathcal{A}(\mathcal{Z})^n \otimes X(\mathcal{H}_2)$, where n is the height of U_2. Now the holomorphic curve U_1 on the left side of the trimmed simple ideal-matched curve corresponds to the higher multiplication m_{n+1} of these n elements of $\mathcal{A}(\mathcal{Z})$ which are outputted by the right side, along with the generator $\mathbf{x}_1 \in \mathfrak{S}(\mathcal{H}_1)$.

Summing over all trimmed simple ideal-matched curves, we see that, for large T, the differential ∂_T

is given by

$$\partial_T(\mathbf{x}_1 \times \mathbf{x}_2) = \sum m_{n+1}(\mathbf{x}_1, a_{i_1}, \dots, a_{i_n})(D_{i_n} \circ \cdots \circ D_{i_1})(\mathbf{x}_2).$$

Here $\{a_i\}$ be basic generators for $\mathcal{A}(\mathcal{Z})$, and $D_i : X(\mathcal{H}_2) \to X(\mathcal{H}_2)$ are operators such that

$$\delta^1_{X(\mathcal{H}_2)}(\mathbf{x}_2) = \sum_i a_i \otimes D_i(\mathbf{x}_2).$$

The sum is taken over all finite sequences $(a_{i_1}, \dots, a_{i_n})$ of basic generators of $\mathcal{A}(\mathcal{Z})$.

It turns out that there is a model of the $\mathcal{A}_\infty$ tensor product, namely the box tensor product, which exactly corresponds to this geometric interpretation. See [LOT18, Section 2.4]. The pairing theorem follows from this. More details may be found in Chapter 9 of [LOT18].

Remark 4.48. There is a more algebraic proof, using Sarkar and Wang's nice diagrams [SW10]. Every Heegaard diagram may be turned into a nice Heegaard diagram, and thus into a diagram whose type A module is a genuine differential module (cf. Remark 4.23). Then the $\mathcal{A}_\infty$ tensor product in Theorem 4.44 coincides with the usual tensor product. Assuming all this, the proof for the pairing theorem becomes quite simple. See [LOT18, Chapter 8] for more.

Bibliography

[Abb04] Casim Abbas. Pseudoholomorphic strips in symplectisations I: Asymptotic behavior. *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, 21(2):139–185, 2004.

[Abb14] Casim Abbas. *An Introduction to Compactness Results in Symplectic Field Theory*. Springer Berlin Heidelberg, 2014.

[AD14] Michèle Audin and Mihai Damian. *Morse Theory and Floer Homology*. French. Translated by Reinie Erné. Universitext. Springer London, London, 2014.

[Ati88] Michael Atiyah. Topological quantum field theories. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 68(1):175–186, 1988.

[BEH+03] Frederic Bourgeois, Yakov Eliashberg, Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. Compactness results in symplectic field theory. *Geometry & Topology*, 7(2):799–888, 2003.

[DK90] Simon K. Donaldson and Peter B. Kronheimer. *The Geometry of Four-Manifolds*. Oxford Mathematical Monographs. Clarendon Press (Oxford University Press), New York, NY, 1990.

[DM69] Pierre Deligne and David Mumford. The irreducibility of the space of curves of given genus. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 36(1):75–109, 1969.

[Don02] Simon K. Donaldson. *Floer Homology Groups in Yang-Mills Theory*, volume 147 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2002. With assistance from Mikio Furuta and Dieter Kotschick.

[EGH00] Yakov Eliashberg, Alexander Givental, and Helmut Hofer. *Introduction to symplectic field theory*. In *Visions in Mathematics*. Birkhäuser Basel, 2000, pages 560–673.

[Flo88a] Andreas Floer. A relative Morse index for the symplectic action. *Communications on Pure and Applied Mathematics*, 41(4):393–407, 1988.

[Flo88b] Andreas Floer. An instanton-invariant for 3-manifolds. *Communications in Mathematical Physics*, 118(2):215–240, 1988.

[Flo88c] Andreas Floer. Morse theory for Lagrangian intersections. *Journal of Differential Geometry*, 28(3), 1988.

[Flo88d] Andreas Floer. The unregularized gradient flow of the symplectic action. *Communications on Pure and Applied Mathematics*, 41(6):775–813, 1988.

[Flo89a] Andreas Floer. Symplectic fixed points and holomorphic spheres. *Communications In Mathematical Physics*, 120(4):575–611, 1989.

[Flo89b] Andreas Floer. Witten's complex and infinite-dimensional Morse theory. *Journal of Differential Geometry*, 30(1), 1989.

[FOO+09] Kenji Fukaya, Young-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian Intersection Floer Theor: Anomaly and Obstruction*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society and International Press of Boston, Providence, RI, 2009.

[Fuk93] Kenji Fukaya. Morse Homotopy, $\mathcal{A}^\infty$-category, and Floer homologies. In *Proceedings of GARC Workshop on Geometry and Topology '93*, volume 18 of *Lecture Note Series*, pages 1–102. Seoul National University, 1993.

[Gro85] Mikhail Gromov. Pseudo holomorphic curves in symplectic manifolds. *Inventiones Mathematicae*, 82(2):307–347, 1985.

[GS99] Robert Gompf and András Stipsicz. *4-Manifolds and Kirby Calculus*, volume 20 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 1999.

[Hee98] Poul Heegaard. *Preliminary studies for a topological theory of the connectivity of algebraic surfaces*. Danish. PhD thesis, University of Copenhagen, 1898.

[Hum97] Christoph Hummel. *Gromov's Compactness Theorem for Pseudo-holomorphic Curves*, volume 151 of *Progress in Mathematics*. Birkhäuser Basel, 1997.

[HWZ02] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. Finite energy cylinders of small area. *Ergodic Theory and Dynamical Systems*, 22(05), 2002.

[HWZ96] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. Properties of pseudoholomorphic curves in symplectisations I: Asymptotics. *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, 13(3):337–379, 1996.

[Kel01] Bernhard Keller. Introduction to A-infinity algebras and modules. *Homology, Homotopy and Applications*, 3(1):1–35, 2001.

[KLT20] Çağatay Kutluhan, Yi-Jen Lee, and Clifford Taubes. $HF = HM$, I: Heegaard Floer homology and Seiberg–Witten Floer homology. *Geometry & Topology*, 24(6):2829–2854, 2020.

[KM07] Peter Kronheimer and Tomasz Mrowka. *Monopoles and Three-Manifolds*. New Mathematical Monographs. Cambridge University Press, Cambridge, 2007.

[Lin12] Francesco Lin. *Bordered Heegaard Floer homology*. Master's thesis, University of Pisa, 2012.

[Lip06a] Robert Lipshitz. A cylindrical reformulation of Heegaard Floer homology. *Geometry & Topology*, 10(2):955–1096, 2006.

[Lip06b] Robert Lipshitz. *A Heegaard-Floer invariant of bordered 3-manifolds*. PhD thesis, Stanford University, 2006.

[LOT15] Robert Lipshitz, Peter Ozsváth, and Dylan Thurston. Bimodules in bordered Heegaard Floer homology. *Geometry & Topology*, 19(2):525–724, 2015.

[LOT18] Robert Lipshitz, Peter Ozsváth, and Dylan Thurston. *Bordered Heegaard Floer homology*, volume 254 of number 1216 in *Memoirs of the American Mathematical Society*. American Mathematical Society, Providence, RI, 2018.

[Mil63] John Milnor. *Morse Theory*, number 51 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1963.

[MS12] Dusa McDuff and Dietmar Salamon. *J-holomorphic Curves and Symplectic Topology*. Colloquium Publications. American Mathematical Society, Providence, RI, 2nd edition, 2012.

[MW95] Mario J. Micallef and Brian White. The structure of branch points in minimal surfaces and in pseudoholomorphic curves. *The Annals of Mathematics*, 141(1):35, 1995.

[OS04a] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and knot invariants. *Advances in Mathematics*, 186(1):58–116, 2004.

[OS04b] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and three-manifolds invariants: Properties and applications. *Annals of Mathematics*, 159(3):1159–1245, 2004.

[OS04c] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and topological invariants for closed three-manifolds. *Annals of Mathematics*, 159(3):1027–1158, 2004.

[OS06] Peter Ozsváth and Zoltán Szabó. Holomorphic triangles and invariants for smooth four-manifolds. *Advances in Mathematics*, 202(2):326–400, 2006.

[Par12] Brett Parker. Exploded manifolds. *Advances in Mathematics*, 229(6):3256–3319, 2012.

[Par16] John Pardon. An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves. *Geometry & Topology*, 20(2):779–1034, 2016.

[Ras03] Jacob Rasmussen. *Floer homology and knot complements*. PhD thesis, Harvard University, 2003.

[Sei08] Paul Seidel. *Fukaya Categories and Picard–Lefschetz Theory*. Zurich Lectures in Advanced Mathematics. European Mathematical Society Press, Zurich, 2008.

[SS92] Mika Seppälä and Tuomas Sorvali. *Geometry of Riemann Surfaces and Teichmüller Spaces*. Leopoldo Nachbin, editor, number 169 in North-Holland Mathematics Studies. Elsevier Science Publishers, Amsterdam, Netherlands, 1992.

[Sta63a] James Dillon Stasheff. Homotopy Associativity of H-Spaces. I. *Transactions of the American Mathematical Society*, 108(2):275, 1963.

[Sta63b] James Dillon Stasheff. Homotopy Associativity of H-Spaces. II. *Transactions of the American Mathematical Society*, 108(2):293, 1963.

[SW10] Sucharit Sarkar and Jiajun Wang. An algorithm for computing some Heegaard Floer homologies. *Annals of Mathematics*, 171(2):1213–1236, 2010.

[Tur97] Vladimir Turaev. Torsion invariants of Spin^c-structures on 3-manifolds. *Mathematical Research Letters*, 4(5):679–695, 1997.

[Wen08] Chris Wendl. Finite energy foliations on overtwisted contact manifolds. *Geometry & Topology*, 12(1):531–616, 2008.