Title
Morse-Bott Embedded Contact Homology

Author
Yao, Yuan

Publication Date
2023

Peer reviewed|Thesis/dissertation

Morse-Bott Embedded Contact Homology

by

Yuan Yao

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael Hutchings, Chair
Professor David Nadler
Professor Martin White

Spring 2023

Morse-Bott Embedded Contact Homology

Abstract

Morse-Bott Embedded Contact Homology

by

Yuan Yao

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Michael Hutchings, Chair

This dissertation constructs a Morse-Bott version of embedded contact homology (ECH). The dissertation is comprised of two parts, corresponding to the two papers written by the author as a graduate student at UC Berkeley. The first part explains how to compute ECH in the Morse-Bott setting when certain transversality conditions are met and provided a certain correspondence theorem is true; and gives a large class of examples where the transversality conditions are satisfied. The second part provides the analytic foundations of the first part by giving a proof of the correspondence theorem.

Contents

Acknowledgments

First and foremost I would like thank my advisor Michael Hutchings for his constant guidance and support throughout my PhD. I would like to thank my friends both in Berkeley and outside Berkeley for many times making the more difficult periods of graduate school bearable. I would like to thank the administrative staff at the math department, in particular Jon Philips, Isabel Seneca, and Vicky Lee for replying to many of my emails titled "Urgent! Please respond ASAP!" in a timely fashion. Finally I would like to thank Evans Hall for holding itself together and not collapsing onto its inhabitants during my stay there.

Chapter 1

Introduction

1.1 Structure of the Dissertation

This dissertation is composed of two parts. The first part (Chapter 2) describes how to compute ECH in certain Morse-Bott settings by using intersection theory. Chapter 2 assumes certain transversality conditions and assumes an analytic result about correspondences between holomorphic curves and cascades (2.6.34). Chapter 2 then presents a large class of examples where the transversality condition is always satisfied. The second part (Chapter 3) presents the analytic foundations of Morse-Bott ECH by proving the correspondence between holomorphic curves and cascades assumed in Chapter 2 (3.4.5). The results in Chapter 3 should generalize to other Floer theories and should be of independent interest.

The two parts of the thesis are self contained and can be read independently of each other. Below we give an introduction to ECH and present some context for our results.

1.2 Introduction to ECH

Let (Y^3, λ) be a contact 3-manifold. We assume that λ is nondegenerate, i.e. for any periodic orbit γ of the Reeb vector field, the linearized return map $d\gamma_T : \mathbb{C} \to \mathbb{C}$ around each Reeb orbit does not have 1 as eigenvalue.

Consider the symplectization of Y^3, which we write as

$$(\mathbb{R} \times Y^3, \omega := d(e^s\lambda))$$

where s is the coordinate on $\mathbb{R}$ and ω is the symplectic form. We further choose a compatible almost complex structure on $\mathbb{R} \times Y^3$ (See Definition 2.3.2).

Then ECH is a homology group associated to the above data. The generators of the chain complex are collections of Reeb orbits, and the differential counts J-holomorphic curves of "ECH index one" (see Chapter 2 for a brief review) in the symplectization.

ECH is only dependent on the contact structure $\xi = \ker\lambda$, and is independent of the contact form and the choice of almost complex structure. In particular it is isomorphic to a version of monopole Floer homology (see [58, 59, 60, 61, 62]),

$$\mathrm{ECH}(Y, \xi, \Gamma) \cong \hat{\mathrm{HM}}(Y, \mathfrak{s}_\xi + PD(\Gamma)).$$

Here Γ denotes the homology class in $H_1(Y)$ represented by the orbit sets in the ECH generators; $PD(\Gamma)$ is the Poincare dual; and $\mathfrak{s}_\xi$ is a spin-c structure on Y associated to the 2-plane field ξ.

Being isomorphic to a version of monopole Floer thoery has given ECH a myriad of applications in symplectic geometry, such as symplectic embedding obstructions [49, 30], refinements of the Weinstein conjecture [16, 15], closing lemmas for Reeb vector fields in contact 3-manifolds [44], and many others.

There has been many direct computations of ECH via enumeration of J-holomorphic curves in the literature, see for example [37, 9, 52]. However in many of these examples, the most naturally associated contact form λ (and the one for which we can enumerate the J-holomorphic curves) is degenerate with a specific type of degeneracy called Morse-Bott degeneracy. The Reeb orbits, instead of being isolated, appear in families. Many of these computations (e.g. [37, 9]) are done with certain assumption of how ECH in the Morse-Bott setting should be defined in general. The purpose of this thesis is to give a rigorous foundation to Morse-Bott ECH where the Reeb orbits are allowed to appear in S^1 families.

The idea of computing ECH in Mose-Bott settings is that instead of counting J holomorphic curves of ECH index one, one should count ECH index one J-holomorphic cascades (See definition 3.3.7), which one can think of as J-holomorphic curves connected to each other by gradient flow lines (see 2.1 for an illustration). In Chapter 2 we give a definition of what it means to have a cascade of ECH index one, and under certain transversality conditions characterize the cascades of ECH index one using intersection theory. We further find large families of examples of contact 3-manifolds where our transversality conditions always hold. In chapter 3 we establish a correspondence between cascades of ECH index one and J-holomorphic curves of ECH index one, and hence show the chain complex constructed by counting ECH index one cascades actually compute the embedded contact homology of the underlying contact 3-manifold.

Chapter 2

Computing ECH in Morse-Bott settings

2.1 Abstract

Given a contact three manifold Y with a nondegenerate contact form λ, and an almost complex structure J compatible with λ, its embedded contact homology $ECH(Y, \lambda)$ is defined ([31]) and only depends on the contact structure. In this paper we explain how to compute ECH for Morse-Bott contact forms whose Reeb orbits appear in S^1 families, assuming the almost complex structure J can be chosen to satisfy certain transversality conditions (this is the case for instance for boundaries of concave or convex toric domains, or if all the curves of ECH index one have genus zero). We define the ECH chain complex for a Morse-Bott contact form via an enumeration of ECH index one cascades. We prove using gluing results from Chapter 3 [67] that this chain complex computes the ECH of the contact manifold. This paper and [67] fill in some technical foundations for previous calculations in the literature ([9], [37]).

2.2 Introduction

Embedded contact homology

In this article we develop some tools to compute the embedded contact homology (ECH) of contact 3-manifolds in Morse-Bott settings.

ECH is a Floer theory defined for a pair (Y, λ), where Y is a three dimensional contact manifold with nondegenerate contact form λ (for an introduction see [31]). The ECH chain complex is generated by orbit sets of the form $\alpha = \{(\gamma_i, m_i)\}$. Here γ_i are distinct simply covered Reeb orbits of λ; and the m_i is a positive integer which we call the multiplicity of γ_i. To describe the differential, consider the symplectization $(\mathbb{R} \times Y, d(e^s\lambda))$ of Y with almost complex structure J. Here s denotes the variable in the $\mathbb{R}$ direction; and J is a generic λ-compatible almost complex structure (see Definition 2.3.2). The differential of ECH, which we write as ∂, is defined by counting holomorphic currents of ECH index $I = 1$ in the symplectization. More precisely, the coefficient $\langle \partial\alpha, \beta\rangle$ is defined by counts of J-holomorphic currents that approach α as $s \to \infty$ and β as $s \to -\infty$, where convergence to α, β is in the sense of currents. The resulting homology, which we write as $ECH(Y, \xi)$, is an invariant of the contact structure $\xi = \ker\lambda$. See Section 2.3 below for a more precise review of ECH and the ECH index.

In part due to its gauge theoretic origin, ECH has had spectacular applications to understanding symplectic problems and dynamics in low dimensions; for instance sharp symplectic embedding obstructions of four dimensional symplectic ellipsoids ([49]), closing lemmas for Reeb flows on contact 3-manifolds ([44]), the Arnold chord conjecture ([42, 43]), and quantitative refinements of the Weinstein conjecture [16]. Several computations (e.g. [37, 9, 46]) and applications (e.g. [30]) of ECH have assumed results from its Morse-Bott version, which we develop in detail in this paper.

Morse-Bott theory

The original definition of ECH requires we use non-degenerate contact forms. However, in practice many contact forms we encounter carry Morse-Bott degeneracies, for which the Reeb orbits are no longer isolated but instead show up in families with weaker non-degeneracy conditions imposed (for a more precise description, see Definition 3.2 in [53]). Although all Morse-Bott contact forms can be perturbed to non-degenerate ones, it is often useful to be able to compute ECH directly in the Morse-Bott setting, where often the enumeration of J-holomorphic curves is easier.

For ECH, since we only consider 3-manifolds, the two Morse-Bott cases are either when the Reeb orbits come in a two dimensional family, or come in one dimensional families. For the first case it then follows that the entire contact manifold is foliated by periodic Reeb orbits. ECH with this kind of Morse-Bott degeneracy has been computed in many cases by [52], see also [18].

The other case is when Reeb orbits show up in one dimensional S^1 families, i.e. we see tori foliated by Reeb orbits. We shall call these tori Morse-Bott tori. It is with this case we concern ourselves in this paper (for a description of what the contact form looks like, see Proposition 2.4.2). Examples of this include boundaries of toric domains, and torus bundles over the circle see [25, 10, 13, 46].

For now we consider (Y^3, λ) a contact 3-manifold where λ is a Morse-Bott contact form all of whose Reeb orbits appear in S^1 families. Later for the case of boundary of convex or concave toric domains (Sections 2.10,2.11) we allow the case of both nondegenerate Reeb orbits and S^1 families of Reeb orbits. We consider the symplectization with a generic λ compatible almost complex structure J (see Definition 2.3.2)

$$(\mathbb{R} \times Y^3, d(e^s\lambda)).$$

Following the recipe described in [5], to compute ECH in the Morse-Bott setting we shall count holomorphic cascades of ECH index one. The philosophy behind this is as follows: given λ, a Morse-Bott contact form with Reeb orbits in Morse-Bott tori, we can perturb

$$\lambda \longrightarrow \lambda_\delta$$

where λ_δ with $\delta > 0$ is a nondegenerate contact form up to a certain action level $L >> 0$. This perturbation requires the following information. For each circle of orbits parameterized by S^1, choose a Morse function f on S^1 with two critical points. The effect of this perturbation is so that each Morse-Bott torus splits into two nondegenerate Reeb orbits (corresponding to the critical points of f): one is an elliptic orbit and the other is a hyperbolic orbit. We also need to perturb the λ-compatible almost complex structure on the symplectization into a λ_δ compatible almost complex structure, J_δ. Since λ_δ is nondegenerate up to action L, we can define the ECH chain complex up to action L in this case by counting ECH index one J_δ-holomorphic curves. The idea is to take $\delta \to 0$ and see what these ECH index one holomorphic curves degenerate into.

By a compactness theorem in [6] (see also [5, 67]), such J_δ-holomorphic curves degenerate into J-holomorphic cascades. For a definition of J-holomorphic cascade, see [67]. Roughly speaking, a J-holomorphic cascade, which we shall write as $u^{\lightning}$, consists of a sequence of J-holomorphic curves $\{u^1, .., u^n\}$ that have ends on Morse-Bott tori. We think of the curves u^i as living on different levels, with u^i one level above u^{i+1}. Between adjacent levels there is the data of a single number $T_i \in [0, \infty]$ described as follows. Suppose a positive end of u^{i+1} is asymptotic to a simply covered Reeb orbit γ with multiplicity n. This γ corresponds to a point on S^1 (the S^1 that parameterizes the family of Morse-Bott Reeb orbits). Then if we follow the upwards gradient flow of f for time T_i starting at the point corresponding to the Reeb orbit γ, we arrive at a Reeb orbit $\tilde{\gamma}$, and a negative end of u^i is asymptotic to $\tilde{\gamma}$ with the same multiplicity n. We assume all positive ends of u^{i+1} and negative ends of u^i are matched up in this way. For an illustration of a cascade[1], see Figure 1.

Main results

The Morse-Bott ECH chain complex which we write as $(C_*^{MB}, \partial_{MB})$ (see section 2.8) can be described as follows. Its generators are collections of Morse-Bott tori, equipped with a multiplicity and additional data, which we write as $\alpha = \{(\mathcal{T}_j, \pm, m_j)\}$. Here $\mathcal{T}_j$ denotes a Morse-Bott torus; we call m_j the multiplicity; and a choice of $+$ or $-$. See Section 2.6 for a description. Suppose we can choose a λ compatible almost complex structure J which is "good" (see definition 2.5.3), meaning certain transversality conditions (Definition 2.5.5) are satisfied. The differential in the Morse-Bott chain complex ∂_{MB} counts ECH index one cascades between Morse-Bott ECH generators. The ECH index of a cascade is described in Section 2.6. We describe what it means for an cascade to be asymptotic to a Morse-Bott

[1]This figure and the accompanying explanations are taken from Figure 1 in [67].

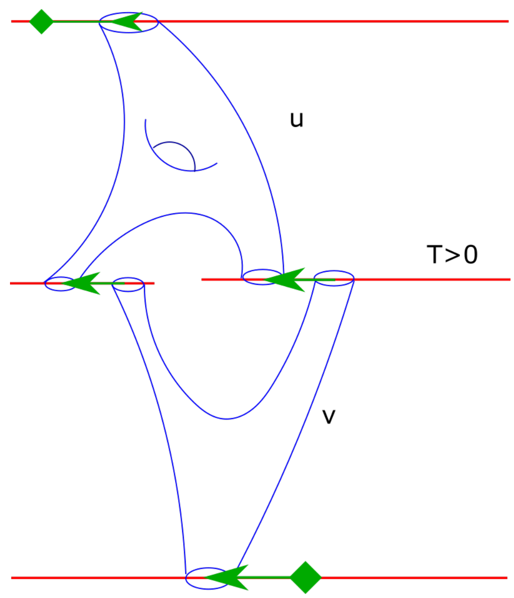

Figure 2.1: A schematic picture of a cascade: the cascade $u^{\natural}$ consists of two levels, u and v. Horizontal lines correspond to Morse-Bott tori. Moving in the horizontal direction along these horizontal lines corresponds to moving to different Reeb orbits in the same S^1 family. Arrows correspond to gradient flows, and diamonds correspond to critical points of Morse functions on S^1 families of Reeb orbits. Between the holomorphic curves u and v, there is a single parameter T that tells us how long positive ends of v must follow the gradient flow to meet a negative end of u.

ECH generator in Section 2.6. For a description of what ECH index one cascades look like, see Corollary 2.6.29, Prop. 2.6.33. We prove that

Theorem 2.2.1. *Let λ be a Morse-Bott contact form on the contact 3-manifold Y whose Reeb orbits all appear in S^1 families. Assuming the almost complex structure J is good (see Definition 2.5.3), the homology of the Morse-Bott ECH chain complex computes the ECH of the contact manifold $ECH(Y,\xi)$.*

A slightly more precise version of this theorem that we prove is Theorem 2.8.1.

We next find some instances there is enough transversality to compute ECH using the Morse-Bott chain complex.

Theorem 2.2.2. *Let λ be a Morse-Bott contact form on the contact 3-manifold Y whose Reeb orbits all appear in S^1 families. We can choose a generic J so that*

- *Every reduced cascade (See Definition 2.4.13) of ≤ 3 levels is transversely cut out (see Definition 2.5.5).*

- *Every reduced cascade where all of the (nontrivial) J-holomorphic components of the reduced cascade (in all of its levels) are distinct up to translation in the symplectization direction is transversely cut out (see Definition 2.5.5).*

If we can show through some other means that we can choose a small perturbation of J to J_δ satisfying conditions of Theorem 2.8.3 so that for small enough δ, all ECH index one curves degenerate into cascades whose reduced version satisfy either of the above conditions, then consider the Morse-Bott ECH chain complex $(C_*^{MB}, \partial_{MB})$ as described more precisely in Section 2.8. For the differential ∂_{MB}, if we restrict to "good" cascades (see Sections 2.6, 2.8 for the notion of "good") of ECH index one whose reduced versions are of the above form, the differential is well defined and the chain complex $(C_*^{MB}, \partial_{MB})$ computes $ECH(Y, \xi)$.

For a discussion how these conditions arise and a proof of this theorem, see the Appendix. This list is by no means exhaustive. We expect there are many other situations where transversality can be achieved; the particulars will depend on the specific details of the contact manifold for which we are computing the ECH chain complex. In particular, for the case relevant for boundaries of convex and concave toric domains, we have the following:

Theorem 2.2.3. *Let λ be a contact form on the contact 3-manifold Y whose Reeb orbits apppear either in Morse-Bott S^1 families or are non-degenerate. Let $\delta > 0$, and λ_δ be the nondegenerate perturbation of λ that perturbs each S^1 family of Reeb orbits into two nondegenerate ones. If for $\delta > 0$ small enough, all J_δ holomorphic curves of ECH index one in $\mathbb{R} \times Y^3$ have genus zero, then the embedded contact homology of Y can be computed from the Morse-Bott chain complex $(C_*^{MB,tree}, \partial_{MB}^{tree})$ (see Section 2.9) using an enumeration of tree like cascades.*

To be more precise, for the above theorem we need to use a slightly different description of cascades which we call "tree like" cascades, which is explained in Sections 2.9, 2.10, 2.11. Consequently, we can prove

Theorem 2.2.4. *For boundaries of concave toric domains or convex toric domains, in the nondegenerate case after a choice of generic almost complex structure all curves of ECH index one have genus zero. Therefore the ECH of boundaries of concave/convex toric domains can be computed using the Morse-Bott ECH chain complex $(C_*^{MB,tree}, \partial_{MB}^{tree})$, via counts of tree-like ECH index one cascades.*

For a definition of convex and concave toric domains, see Sections 2.10, 2.11.

We mention some previous computations of ECH that have assumed Morse-Bott theory of the flavour we develop in this paper, notably in [37] for the case of T^3, and [9] for certain toric contact 3-manifolds, and [46] for the case of T^2 bundles over S^1. This paper and the gluing paper [67] fill in the foundations for these results.

Remark 2.2.5. The above theorems say for genus zero curves we have all the transversality we need by simply restricting to cascades of ECH index one and choosing a generic J;

however this result is not strict, there could well be other scenarios where transversality can be achieved. For instance we expect with some more care we can show the moduli space of cascades of ECH index one and genus one can be shown to be transverse. For discussion of general difficulties see the Appendix.

Some technical details

For ECH in the nondegenerate setting (see [31]), as we review in Section 2.3, the Fredholm index of a somewhere injective curve is bounded from above by its ECH index. Further, the ECH index is superadditive under unions of J-holomorphic curves in symplectizations. Using the fact that after choosing a generic almost complex structure, all somewhere injective curves are transversely cut out, it follows that by restricting to only ECH index one curves we do not need to consider multiply covered nontrivial curves. With this, one defines the ECH differential in the nondegenerate setting via counts of ECH index one J-holomorphic curves.

Parts of the above story continue to hold in the case of cascades if we assume can choose J to be good (Definition 2.5.3), as we explain below.

We first note that the notion of an ECH index continues to make sense for cascades, as we explain in Section 2.6. The case of cascades, however, is more complicated, in two directions.

- During the degeneration process for λ_δ as $\delta \to 0$, simple curves may degenerate into cascades that have multiply covered components;
- For generic J, and even if we restrict to cascades all of whose curves are somewhere injective, the cascade need not be transversely cut out.

The second bullet point is the most problematic. This happens because by requiring there is a single parameter between adjacent levels, we are imposing restrictions on the evaluation maps on the ends of the curves in a cascade. Hence a cascade lives in a fiber product, which need not be transversely cut out even if we restrict to only somewhere injective curves. For an explanation of this, see the Appendix.

However, if we take as an *assumption* that J is good (which isn't always possible, it will depend on the specific contact manifold), then all cascades built out of somewhere injective curves that we consider are transversely cut out. Then we can address the first bullet point by using a version of the ECH index inequality for cascades .

To explain the ECH index inequality for cascades, consider the following. Given a cascade, we can pass to a reduced cascade, which means we replace all multiply covered curves with the underlying simple curves. See Section 2.4 for a precise description of this process. The reduced cascade also lives in a fiber product because of the conditions we imposed on its ends. By the assumption that J is good (and consequently transversality assumptions in Definition 2.5.5 are satisfied), the reduced cascade is transversely cut out. To each reduced cascade we can associate to it a virtual dimension, which is the dimension of the moduli space of curves that lies in the same configuration as the reduced cascade. We prove that the ECH index of the cascade bounds the Fredholm index of the reduced cascade from above; and that equality holds only if the original cascade had no multiply covered components (and is well behaved in various ways, see Section 2.6).

In [67], we proved a correspondence theorem between certain cascades and J-holomorphic curves.

Theorem 2.2.6 ([67])**.** *Given a "transverse and rigid" (see Definition 3.4 in [67]) height one J-holomorphic cascade $u^{\natural}$, it can be glued to a rigid J_δ-holomorphic curve u_δ for $\delta > 0$ sufficiently small. The construction is unique in the following sense: if $\{\delta_n\}$ is a sequence of numbers that converge to zero as $n \to \infty$, and $\{u'_{\delta_n}\}$ is sequence of J_{δ_n}-holomorphic curves converging to $u^{\natural}$, then for large enough n, the curves u'_{δ_n} agree with u_{δ_n} up to translation in the symplectization direction.*

In this paper, using index calculations, we show that if J is good (some instances of which are described in Theorems 2.2.2), then essentially all ECH index one cascades are transverse and rigid[2]. Thus the gluing theorem above is then used to show the Morse-Bott chain complex computes $ECH(Y, \lambda)$. In the cases where we use "tree like" cascades, for instance for boundaries of convex or concave toric domains, the definitions are slightly different, but essentially the same story holds true and we can always choose a generic J so that the Morse-Bott chain complex computes $ECH(Y, \lambda)$.

Finally in the Appendix we explain why the usual techniques for achieving transversality fails for cascades.

Acknowledgements I would like to thank my advisor Michael Hutchings for his consistent help and support throughout this project. I would like to thank Morgan Weiler for helpful comments. I would like to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), PGSD3-532405-2019. Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), PGSD3-532405-2019.

2.3 ECH review

For a thorough introduction to ECH see [31]. We will summarize much of the material from [31] and [29] for convenience of the reader.

Let (Y^3, λ) be a contact 3 manifold with nondegenerate contact form λ. The generator of ECH are collections Θ, where each Θ is a set of Reeb orbits with multiplicities

$$\Theta := \{(\gamma_i, m_i) | \gamma_i \text{ are pairwise distinct simply covered Reeb orbits, } m_i \in \mathbb{Z}_+\}.$$

[2]Technically we need to restrict ourselves to *good* ECH index one cascades. This is a fairly minor point, but see Proposition 5.32 and surrounding discussion.

We require $m_i = 1$ if γ_i is a hyperbolic orbit. Then the chain for ECH are just

$$C_*(\lambda') := \bigoplus_{\Theta_i} \mathbb{Z}_2\langle\Theta_i\rangle.$$

Remark 2.3.1. There is a decomposition of ECH according to homology class of Θ_i in $H_1(Y)$. ECH can also be defined using $\mathbb{Z}$ coefficients. We will not address these issues here.

Let α, β be ECH generators. Consider the symplectization of Y, defined as the symplectic manifold $(\mathbb{R} \times Y, \omega := d(e^a\lambda))$, where a denotes the $\mathbb{R}$ coordinate. Equip it with a generic λ compatible almost complex structure J. By compatible we mean the following

Definition 2.3.2. *Let λ be a contact form (not necessarily nondegenerate) on a contact 3-manifold. Let J be a almost complex structure on the symplectization $(\mathbb{R} \times Y, \omega := d(e^a\lambda))$. We say J is compatible with λ if*

a. *J is invariant in the $\mathbb{R}$ direction;*

b. *Let R denote the Reeb vector field, then $J\partial_s = R$;*

c. *Let ξ denote the contact structure, then $J\xi = \xi$ and $d\lambda(\cdot, J\cdot)$ defines a metric on ξ.*

Then the coefficient $\langle\partial\alpha, \beta\rangle$ is defined by

$$\langle\partial\alpha,\beta\rangle := \left\{\begin{array}{l}\mathbb{Z}_2 \text{ count of holomorphic currents } \mathcal{C} \text{ of ECH index } I=1, \\ \text{so that as } s\to+\infty,\ \mathcal{C} \text{ approaches } \alpha \text{ as a current, and as } s\to-\infty, \\ \mathcal{C} \text{ approaches } \beta \text{ as a current.}\end{array}\right\} \tag{2.1}$$

A holomorphic current $\mathcal{C}$ is by definition a collection $\{(C_i, m_i)\}$ where each C_i is a somewhere injective J holomorphic curve and $m_i \in \mathbb{Z}_{>0}$ accounts for the multiplicity of this curve. The ECH index I of a holomorphic curve C (or more generally a relative 2 homology class in $H_2(\alpha, \beta, Y)$, see section below for a definition) is defined by

$$I(C) := Q_\tau(C) + c_\tau(C) + CZ^I(C) \tag{2.2}$$

where $Q_\tau(C)$ is the relative intersection number, $c_\tau(C)$ is the relative Chern class, and CZ is a sum of Conley Zehnder indices used in ECH. We will review these terms in the upcoming subsections.

Relative first Chern class

Let α, β be orbit sets. We define the relative homology group $H_2(\alpha, \beta, Y)$ to be the set of 2-chains Σ with

$$\partial\Sigma = \alpha - \beta$$

modulo boundary of 3 chains. This is an affine space over $H_2(Y)$, and each J holomorphic curve defines a relative homology class.

We fix trivializations τ of the contact structure ξ over each Reeb orbit in Y. We then define the relative first Chern class c_τ with respect to this choice of trivialization. For a given homology class in $H_2(\alpha,\beta,Y)$, choose a representative $Z \in H_2(\alpha,\beta,Y)$ that is embedded near its boundaries α,β. We assume Z is a smooth surface. Let $\iota : Z \to Y$ be the inclusion. Then consider the bundle $\iota^*\xi$ over Z. Let ψ be a section of this bundle that is constant with respect to the trivialization τ near each of the Reeb orbits, and perturb ψ so that all of its zeroes are transverse. Then $c_\tau(Z)$ is defined to be the algebraic count of zeroes of ψ. See [31] for a more thorough explanation and that this is well defined.

Writhe

Let C be a somewhere injective J holomorphic curve in the symplectization of Y, $(\mathbb{R} \times Y, d(e^a\lambda))$ (with generic λ-compatible complex structure J) that is asymptotic to α as $s \to +\infty$ and β as $s \to -\infty$. For simplicity we focus on $s \to +\infty$ end. It is known (see for example [57]) that for s sufficiently large, $C \cap \{s\} \times Y$ is a union of embedded curves near each orbit of α. For each orbit γ_i of α, the curves $C \cap \{s\} \times Y$ forms a braid ξ_i^+. We use the trivialization τ to identify the braids ξ_i^+ with braids in $S^1 \times D^2$. We can define the writhe of ξ_i^+ by identifying $S^1 \times D^2$ with an annulus times an interval, projecting ξ_i^+ to the annulus, and counting crossings with signs. The same sign convention is clearly explained in [33].

Then given a somewhere injective J-holomorphic curve C that is not the trivial cylinder, with braids ζ_i^+ associated to the i-th Reeb orbit it approaches as $s \to +\infty$ and braids ζ_j^- associated to the jth Reeb orbit it approaches as $s \to -\infty$ we define its writhe to be

$$w_\tau(C) := \sum_i w_\tau(\zeta_i^+) - \sum_j w_\tau(\zeta_j^-).$$

We also recall the writhe of the braid ζ_i^+ can be bounded by expressions in terms of the Conley-Zehnder indices.

Proposition 2.3.3. *Let C be a somewhere injective holomorphic curve that is not a trivial cylinder which is asymptotic to γ_i with total multiplicity n_i. Suppose there are k_i distinct ends of C that are asymptotic to γ_i, with covering multiplicities q_i^j. Then the writhe associated to the braid ζ_i^+ corresponding to Reeb orbit γ_i is bounded above by*

$$w_\tau(\zeta_i^+) \leq \sum_j^{n_i} CZ(\gamma_i^j) - \sum_j^{k_i} CZ(\gamma_i^{q_i^j}) \tag{2.3}$$

A similar bound holds for braids at $s \to -\infty$ with signs reversed.

We will derive an analogue of this bound for the Morse-Bott case. For now we recall another definition:

Definition 2.3.4. *Let C be a somewhere injective J-holomorphic curve that is not a trivial cylinder. For each γ_i that C is asymptotic to as $s \to +\infty$, form the sum $CZ^I(\gamma_i) := \sum_{j=1}^{n_i} CZ(\gamma_i^j)$ as above, and for each γ_i' that C is asymptotic to as $s \to -\infty$, we form an analogous sum, then we define*

$$CZ^I(C) := \sum_{\substack{\gamma_i,\\ C \text{ is asymptotic to } \gamma_i,\\ \text{as } s\to+\infty}} CZ^I(\gamma_i) - \sum_{\substack{\gamma_i',\\ C \text{ is asymptotic to } \gamma_i',\\ \text{as } s\to-\infty}} CZ^I(\gamma_i'). \tag{2.4}$$

This is the Conley-Zehnder index term that appears in the definition of ECH index.

Relative adjunction formula

In this section we review the relative adjunction formula (see [31, 29]). We first review the notion of relative intersection pairing, which is a map depending on the trivialization τ:

$$Q_\tau : H_2(\alpha, \beta, Y) \times H_2(\alpha, \beta, Y) \to \mathbb{Z}$$

as follows. Let S and S' be surfaces representing relative homology classes in $H_2(\alpha, \beta, Y)$. If we identify $\mathbb{R} \times Y$ with $(-1, 1) \times Y \subset [-1, 1] \times Y$, then we have by definition

$$\partial S = \partial S' = \sum_i m_i\{1\} \times \alpha_i - \sum_i n_i\{-1\} \times \beta_i$$

We make the following requirements on the representatives S and S':

a. The projections to Y of the intersections of S and S' with $(1-\epsilon, 1] \times Y$ and $[0, \epsilon) \times Y$ are embeddings.

b. Each end of S or S' covers Reeb orbits α_i (resp β_i) with multiplicity 1.

c. The image of S (after projecting to Y in a neighborhood $S^1 \times D^2$ of α_i determined by the trivialization τ) do not intersect, and do not rotate with respect to the chosen trivialization τ as one goes around α_i. Further, the image of different ends of S approaching α_i lie on distinct rays in a neighborhood of α_i. More concretely using trivialization τ to identify a neighborhood of α_i with $S^1 \times \mathbb{R}^2$, ends of S approach α_i along different rays in $\mathbb{R}^2$. We make a similar requirement for β_i. We make a similar requirement for S'.

d. All interior intersections between S and S' are transverse.

Representatives satisfying all of the above conditions are called τ -representatives in [29], which is a definition we will adopt. Then given τ representatives as listed above, $Q_\tau(S, S')$ is defined to be the algebraic count of intersections between S and S'.

We are now ready to state the relative adjunction formula, see also [29].

Proposition 2.3.5. *If C is a somewhere injective J holomorphic curve,*

$$c_\tau(C) = \chi(C) + Q_\tau(C) + w_\tau(C) - 2\delta(C) \tag{2.5}$$

where $\delta(C) \geq 0$ is defined to be an algebraic count of singularities of C. Each singularity is positive due to the fact C is J-holomorphic.

ECH index inequality

We have now defined all of the terms that appear in the ECH index inequality. We compare this with the Fredholm index. Let C be a somewhere injective J-holomorphic curve, let $Ind(C)$ denote the Fredhom index of C, which in this case is given by

$$-\chi(C) + 2c_\tau(C) + CZ^{Ind}(C).$$

Here $CZ^{Ind}(C)$ is defined as follows. If C is positively asymptotic to γ with k ends, each of multiplicity q_k, then the contribution to $CZ^{Ind}(C)$ from γ is given by $\sum_k CZ(\gamma^{q_k})$. Similarly if C is asymptotic to γ at the negative ends, then its contribution to $CZ^{Ind}(C)$ is $-\sum_k CZ(\gamma^{q_k})$.

Theorem 2.3.6. *Let C denote a somewhere injective J-holomorphic curve as above, then we have the following inequality*

$$Ind(C) \leq I(C) - 2\delta(C). \tag{2.6}$$

An immediate corollary of the above is

Corollary 2.3.7. *Let $\mathcal{C}$ be a J-holomorphic current of $I(\mathcal{C}) = 1$. Then for generic J, the current $\mathcal{C}$ must satisfy*

a. *It contains an unique connected embedded curve C of multiplicity one that is not a trivial cylinder. The ends of C approach Reeb orbits according to partition conditions. (See [31, Section 3] for a discussion of partition conditions). We will review the relevant partition conditions in the Morse-Bott setting later).*

b. *The other components of $\mathcal{C}$ are trivial cylinders with multiplicities.*

Convention 2.3.8. *In this paper we describe a correspondence between ECH index 1 currents in the nondegenerate setting and ECH index 1 cascades in the Morse-Bott setting. We will only care about the nontrivial part of the ECH index 1 current, as the trivial cylinders correspond trivially in the non-degenerate and Morse-Bott situations. Hence when we say cascade, or a sequences of ECH index one curves/currents degenerating into a cascade, unless stated otherwise, we will always be considering what happens to the nontrivial part of the ECH index one current, and what cascade it corresponds to.*

J_0 index and finiteness

We recall (without proof) the following proposition (see [29],[31]):

Proposition 2.3.9. *Let α, β be ECH generators. We choose a generic J, and let $\mathcal{M}^{I=1}(\alpha,\beta)/\mathbb{R}$ denote the moduli space of ECH index $=1$ currents from α to β modulo the action of $\mathbb{R}$. Then $\mathcal{M}^{I=1}(\alpha,\beta)/\mathbb{R}$ is a finite collection of points.*

We will mention two results that go into this proof, for we will need analogous constructions in the Morse-Bott context.

Definition 2.3.10. *Let $\alpha = \{(\alpha_i, m_i)\}, \beta = \{(\beta_i, n_i)\}$ be ECH generators, let $Z \in H_2(\alpha,\beta,Y)$ be a relative homology class. We define:*

$$J_0(\alpha,\beta,Z) = -c_\tau(Z) + Q_\tau(Z) + CZ^{J_0}(\alpha,\beta) \tag{2.7}$$

where

$$CZ^{J_0}(\alpha,\beta) := \sum_i \sum_{k=1}^{m_i-1} CZ(\alpha_i^k) - \sum_i \sum_{k=1}^{n_i-1} CZ(\beta_i^k) \tag{2.8}$$

We have the following proposition bounding the topological complexity of holomorphic curves counted by ECH index 1 conditions:

Proposition 2.3.11. *Let $\mathcal{C} \in \mathcal{M}^{I=1}(\alpha,\beta)$, which decomposes as $\mathcal{C} = C_0 \cup C$ where C_0 is a union of trivial cylinders, and C is somewhere injective and nontrivial. Let n_i^+ denote the number of positive ends C has at α_i, plus 1 if C_0 includes cylinders of the form $\mathbb{R}\times\alpha_i$, define n_j^- analogously for β and negative ends of C then*

$$-\chi(C) + \sum_i (n_i^+ - 1) + \sum_j (n_j^- - 1) \leq J_0(C). \tag{2.9}$$

Finally we state the version of Gromov compactness for currents. Let α, β be orbit sets, we define a broken holomorphic current from α, β to be a finite sequence of J nontrivial holomorphic currents $(\mathcal{C}^0, .., \mathcal{C}^k)$ in $\mathbb{R}\times Y$ such that there exists orbit sets $\alpha = \gamma^0, \gamma^1, .., \gamma^{k+1} = \beta$ so that $\mathcal{C}^i \in \mathcal{M}(\gamma^i, \gamma^{i+1})$ (this notation means $\mathcal{C}^i$ is a current from the orbit set γ^i to γ^{i+1}). By nontrivial we mean a current is not entirely composed of unions of trivial cylinders. We say a sequence of holomorphic currents $\{\mathcal{C}_{v\geq 1}\} \in \mathcal{M}(\alpha,\beta)$ converges to $(\mathcal{C}^0, .., \mathcal{C}^k)$ if for each $i = 0, .., k$ there are representatives $\mathcal{C}_\nu^i$ of $\mathcal{C}_\nu \in \mathcal{M}(\alpha,\beta)/\mathbb{R}$ such that the sequence $\{\mathcal{C}_{v\geq 1}\}$ converges as a current and as a point set on compact sets to $\mathcal{C}^i$.

Proposition 2.3.12. *([31], [63] Prop 3.3) Any sequence $\{\mathcal{C}_v\}$ of holomorphic currents in $\mathcal{M}(\alpha,\beta)/\mathbb{R}$ has a subsequence which converges to a broken holomorphic current $(\mathcal{C}^0, .., \mathcal{C}^k)$. Further if we denote $\{\mathcal{C}_v\}$ the convergent subsequence, we have the equality*

$$[\mathcal{C}_v] = \sum_{i=0}^{k} [\mathcal{C}^i] \in H_2(\alpha,\beta,Y) \tag{2.10}$$

2.4 Morse-Bott setup and SFT type compactness

Let (Y, λ) be a contact 3 manifold with Morse-Bott contact form λ. Throughout we assume the Morse-Bott orbits come in families of tori.

Convention 2.4.1. *Throughout this paper we fix action level $L > 0$ and only consider ECH generators of action level up to L. This is implicit in all of our constructions and will not be mentioned further. We construct Morse-Bott ECH up to action level L, and the full ECH is recovered by taking $L \to \infty$.*

The following theorem, which is a special case of a more general result in [53], gives a characterization of the neighborhood of Morse-Bott Tori. Let λ_0 denote the standard contact form on $(z, x, y) \in S^1 \times S^1 \times \mathbb{R}$ of the form

$$\lambda_0 = dz - ydx.$$

Proposition 2.4.2. *[53] Let (Y, λ) be a contact 3 manifold with Morse-Bott contact form λ. We assume the Morse-Bott orbits come in families of tori $\mathcal{T}_i$ with minimal period T_i. Then we can choose coordinates around each Morse-Bott torus so that a neighborhood of $\mathcal{T}_i$ is described by $S^1 \times S^1 \times (-\epsilon, \epsilon)$, and the contact form λ in this coordinate system looks like:*

$$\lambda = h(x, y, z)\lambda_0$$

where $h(x, y, z)$ satisfies:

$$h(x, 0, z) = 1, dh(x, 0, z) = 0$$

Here we identify $z \in S^1 \sim \mathbb{R}/2\pi T_i\mathbb{Z}$

See [67] Theorem Proposition 2.2 for a sketch of the proof. By the Morse-Bott assumption there are only finitely many such tori up to fixed action L. We assume we have chosen such neighborhoods around all Morse Bott Tori $\mathcal{T}_i$. Next we shall perturb them to nondegenerate Reeb orbits by perturbing the contact form in a neighborhood of each torus as described below. This is the same perturbation as in [67].

Let $\delta > 0$, let $f : x \in \mathbb{R}/\mathbb{Z} \to \mathbb{R}$ be a smooth Morse function with maximum at $x = 1/2$ and minimum $x = 0$. Let $g(y) : \mathbb{R} \to \mathbb{R}$ be a bump function that is equal to 1 on $[-\epsilon_{\mathcal{T}_i}, \epsilon_{\mathcal{T}_i}]$ and zero outside $[-2\epsilon_{\mathcal{T}_i}, 2\epsilon_{\mathcal{T}_i}]$. Here $\epsilon_{\mathcal{T}_i}$ is a small number chosen for each $\mathcal{T}_i$ small enough so that the normal form in the above theorem applies to all Morse-Bott tori of action $< L$, and that all such chosen neighborhoods these Morse-Bott tori are disjoint. Then in neighborhood of the Morse-Bott tori $\mathcal{T}_i$ we perturb the contact form as

$$\lambda \longrightarrow \lambda_\delta := e^{\delta g f}\lambda.$$

We can describe the change in Reeb dynamics as follows:

Proposition 2.4.3. *For fixed action level $L > 0$ there exists $\delta > 0$ small enough so that the Reeb dynamics of λ_δ can be described as follows. In the trivialization specified by Proposition 2.17, each Morse-bott torus splits into two non-degenerate Reeb orbits corresponding to the two critical points of f. One of them is hyperbolic of index 0, the other is elliptic with rotation angle $|\theta| < C\delta << 1$ and hence its Conley-Zehnder index is ± 1. There are no additional Reeb orbits of action $< L$.*

For proof see [5].

Remark 2.4.4. Later when we define various terms in the ECH index, they will depend on the choice of trivializations of the contact structure on the Reeb orbits. We will always choose the trivialization specified by Proposition 2.4.2. For convenience of notation we will call this trivialization τ and write for example c_τ or Q_τ for the definition of relative Chern class or intersection form with respect to this trivialization.

We also observe that after iterating the Reeb orbit in the Morse-Bott tori, their Robbin-Salamon index stays the same ([24]). So up to action L, in the nondegenerate picture, we will only see Reeb orbits of Conley-Zehnder index $-1, 0, 1$.

Definition 2.4.5. *We say a Morse Bott torus is positive if the elliptic Reeb orbit has Conley-Zehnder index 1 after perturbation; otherwise we say it is negative Morse Bott torus. This condition is intrinsic to the Morse-Bott torus itself, and is independent of trivializations or our choice of perturbations.*

We recall our goal is to define the ECH chain complex up to filtration L, and then take $L \to \infty$ to recover the entire ECH chain complex. Hence, let us consider for small $\delta > 0$ the symplectization

$$(M^4, \omega_\delta) := (\mathbb{R} \times Y^3, d(e^s \lambda_\delta))$$

We equip (M, ω_δ) with a λ_δ compatible almost complex structure J_δ, and $(M, \omega) := (\mathbb{R} \times Y^3, d(e^s \lambda_\delta))$ with λ-compatible almost complex structure J. Both J and J_δ should be chosen generically, with genericity condition specified in Definition 2.5.5 and Theorem 2.8.3. In particular J_δ should be a small perturbation of J, i.e. the C^∞ norm difference between J_δ and J should be bounded above by $C\delta$. For fixed L and small enough and generic choice of δ, the ECH of (Y^3, λ_δ) is defined for generators of action less than L via counts of embedded J-holomorphic curves of ECH index 1. To motivate our construction, we next take $\delta \to 0$ to see what kinds of objects these J holomorphic curves degenerate into. By a theorem of that first appeared in Bourgeois' thesis [5] and also stated in [6] (for a proof see the Appendix of [67]), they degenerate into J-holomorphic cascades. (For a more careful definition of cascades see the appendix of [67] that takes into account of stability of domain and marked points, but the definition here suffices for our purposes).

Definition 2.4.6 ([5], See also definition 3.3.7 in [67])**.** *Let Σ be a punctured (nodal) Riemann surface, potentially with multiple connected components. A cascade of height 1, which we will denote by $u^{\sharp}$, in $(\mathbb{R} \times Y^3, d(e^s\lambda)$ consists of the following data :*

- *A labeling of the connected components of* $\Sigma^* = \Sigma \setminus \{nodes\}$ *by integers in* $\{1, ..., l\}$*, called sublevels, such that two components sharing a node have sublevels differing by at most 1. We denote by* Σ_i *the union of connected components of sublevel* i*, which might itself be a nodal Riemann surface.*

- $T_i \in [0, \infty)$ *for* $i = 1, ..., l-1$*.*

- J*-holomorphic maps* $u^i : (\Sigma_i, j) \to (\mathbb{R} \times Y^3, J)$ *with* $E(u_i) < \infty$ *for* $i = 1, ..., l$*, such that:*

  - *Each node shared by* Σ_i *and* Σ_{i+1}*, is a negative puncture for* u^i *and is a positive puncture for* u^{i+1}*. Suppose this negative puncture of* u^i *is asymptotic to some Reeb orbit* $\gamma_i \in \mathcal{T}$*, where* $\mathcal{T}$ *is a Morse-Bott torus, and this positive puncture of* u^{i+1} *is asymptotic to some Reeb orbit* $\gamma_{i+1} \in \mathcal{T}$*, then we have that* $\phi_f^{T_i}(\gamma_{i+1}) = \gamma_i$*. Here* $\phi_f^{T_i}$ *is the upwards gradient flow of* f *for time* T_i *lifted to the Morse-Bott torus* $\mathcal{T}$*. It is defined by solving the ODE*

$$\frac{d}{ds}\phi_f(s) = f'(\phi_f(s)).$$

  - u^i *extends continuously across nodes within* Σ_i*.*
  - *No level consists purely of trivial cylinders. However we will allow levels that consist of branched covers of trivial cylinders.*

Convention 2.4.7. *We fix our conventions as in [67].*

- *We say the punctures of a* J*-holomorphic curve that approach Reeb orbits as* $s \to \infty$ *are positive punctures, and the punctures that approach Reeb orbits as* $s \to -\infty$ *are negative punctures. We will fix cylindrical neighborhoods around each puncture of our* J*-holomorphic curves, so we will use "positive/negative ends" and "positive/negative punctures" interchangeably. By our conventions, we think of* u^1 *as being a level above* u^2 *and so on.*

- *We refer to the Morse-Bott tori* $\mathcal{T}_j$ *that appear between adjacent levels of the cascade* $\{u^i, u^{i+1}\}$ *as above, where negative punctures of* u^i *are asymptotic to Reeb orbits that agree with positive punctures from* u^{i+1} *up to a gradient flow, intermediate cascade levels.*

- *We say that the positive asymptotics of* $u^{\natural}$ *are the Reeb orbits we reach by applying* ϕ_f^{∞} *to the Reeb orbits hit by the positive punctures of* u^1*. Similarly, the negative asymptotics of* $u^{\natural}$ *are the Reeb orbits we reach by applying* $\phi_f^{-\infty}$ *to the Reeb orbits hit by the negative punctures of* u^l*. They are always Reeb orbits that correspond to critical points of* f*. We note if a positive puncture (resp. negative puncture) of* u^1 *(resp.* u^l*) is asymptotic to a Reeb orbit corresponding to a critical point of* f*, then applying* $\phi_f^{+\infty}$ *(resp.* $\phi_f^{-\infty}$*) to this Reeb orbit does nothing.*

Definition 2.4.8 ([5], Chapter 4, See also definition 3.3.9 in Chapter 3 [67])**.** *A cascade of height k consists of k height 1 cascades, $u_k^{\lightning} = \{u^{1\lightning}, ..., u^{k\lightning}\}$ with matching asymptotics concatenated together.*

By matching asymptotics we mean the following. Consider adjacent height one cascades, $u^{i\lightning}$ and $u^{i+1\lightning}$. Suppose a positive end of the top level of $u^{i+1\lightning}$ is asymptotic to the Reeb orbit γ (not necessarily simply covered). Then if we apply the upwards gradient flow of f for infinite time we arrive at a Reeb orbit reached by a negative end of the bottom level of $u^{i\lightning}$. We allow the case where γ is at a critical point of f, and the flow for infinite time is stationary at γ. We also allow the case where γ is at the minimum of f, and the negative end of the bottom level of $u^{i\lightning}$ is reached by following an entire (upwards) gradient trajectory connecting from the minimum of f to its maximum. If all ends between adjacent height one cascades are matched up this way, then we say they have matching asymptotics.

We will use the notation $u_k^{\lightning}$ to denote a cascade of height k. We will mostly be concerned with cascades of height 1 in this article, so for those we will drop the subscript k and write $u^{\lightning} = \{u^1, ..., u^l\}$.

Remark 2.4.9. As mentioned in Chapter 3 ([67]), we can also think of a cascade of height k as a cascade of height 1 where $k-1$ of the intermediate flow times are infinite.

We now state a SFT style compactness theorem relating non-degenerate J_δ holomorphic curves to cascades. However, the precise statement is rather technical and requires us to take up Deligne-Mumford compactifications of the moduli space of Riemann surfaces. The full version is stated in [6] (see also the Appendix of Chapter 3, [67], where we also sketch a proof). For our purposes it will be sufficient to state the theorem informally as below.

Theorem 2.4.10. *(See [6]) Let u_{δ_n} be a sequence of J_{δ_n}-holomorphic curves with uniform upper bound on genus and energy, then a subsequence of u_{δ_n} converges to a cascade of J-holomorphic curves (which can be apriori of arbitrary height).*

Since ECH is really a theory of holomorphic currents, we find it also useful to define a *cascade of holomorphic currents*, which is what we shall primarily work with.

Definition 2.4.11. *A height 1 holomorphic cascade of currents $\mathbf{u}^{\lightning} = \{u^1, .., u^n\}$ consists of the following data:*

- *Each u^i consists of holomorphic currents of the form (C_j^i, d_j^i). Each C_j^i is a somewhere injective holomorphic curve with $E(C_j^i) < \infty$. The positive integer d_j^i is then the multiplicity.*
- *Numbers $T_i \in [0, \infty), i = 1, .., n-1$*
- *Let γ_i be a simply covered Reeb orbit that is approached by the negative end of some component of u^i, say the components $C_{j_1}^i, ..., C_{j_k}^i$ (such curves have associated multiplicity $d_{j_1}^i, ..., d_{j_k}^i$). Each $C_{j_*}^i$ approaches γ_i with a covering multiplicity n_{j_*}, which is how many times γ_i is covered by $C_{j_*}^i$ as currents. Then the total multiplicity of γ_i as*

covered by u^i is given by $\sum_{*=1,..k} d^i_{j_*} n_{j_*}$. Then consider $\phi_f^{T_i}(\gamma_{i+1}) := \gamma_i$. Then u^{i+1} is asymptotic to γ^{i+1} in its positive end with total multiplicity $\sum_{*=1,..k} d^i_{j_*} n_{j_*}$ also.

- *No level consists of purely of trivial cylinders (even if they have higher multiplicities).*

We define the positive asymptotics of $\mathbf{u}^{\lightning} := \{u^1, .., u^n\}$ as before, except we only care about Reeb orbits up to multiplicity. Then we can similarly define a cascade of currents of height k by stacking together cascades of currents of height 1.

We will refer to ordinary cascade a "cascade of curves" when we wish to distinguish it from a cascade of currents. Then given a cascade of curves, we can pass it to a cascade of currents by using the following procedure:

Procedure 2.4.12.

- *Replace every multiple covered non-trivial curve with a current of the form (C, m) where C is a somewhere injective curve, and we translate all m copies along $\mathbb{R}$ to make the entire collection somewhere injective.*
- *If we see a multiply covered trivial cylinder we replace it with (C, m) where m is the multiplicity and C is a trivial cylinder.*
- *If we see a nodal curve in one of the levels, we separate the node and apply the above process to each of the separated components of the nodal curve.*
- *We remove all levels that only have currents made out of trivial cylinders. Suppose u^i is a level only consisting of trivial cylinders to be removed, and suppose the $s \to +\infty$ end is a intermediate cascade level with flow time T_{i-1}, and the $s \to -\infty$ end of u^i has associated flow time T_i, after the removal of u^i level, the newly adjacent levels u^{i-1} and u^{i+1} have flow time between them equal to $T_i + T_{i-1}$.*

In passing from cascades of curves to currents we have lost some information, but we shall see currents are the natural settings to talk about ECH index.

We later wish to make sense of the Fredholm index of a cascade of currents. To this end we make the definition of *reduced cascade of currents.*

Definition 2.4.13. *Given a cascade of currents $\mathbf{u}^{\lightning}$, for components within it of the form (C, m) where $m > 1$ and C is a nontrivial holomorphic curve, we then replace (C, m) with just $(C, 1)$. After we perform this operation we obtain another cascade of currents, which we label $\tilde{\mathbf{u}}^{\lightning}$, which we call the reduced cascade of currents.*

2.5 Index calculations and transversality

The heart of the calculation that underlies ECH is this: the ECH index bounds from the above the Fredholm index, and if there are curves of ECH index one with bad behaviour (singularities, multiply covers), this would imply the existence of somewhere injective curves of Fredholm index less than 1, which cannot happen for generic J. In this section we take up the issue of establishing Fredholm index for J holomorphic cascades, and explain the transversality issue we encounter.

Given a reduced cascades of currents, $\tilde{\mathbf{u}}^{\prime} = \{\tilde{u}^1, ..., \tilde{u}^n\}$, we would like to assign to it a Fredholm index. Ideally this Fredholm index measures geometrically the dimension of the moduli space this particular cascade lives in. We note that by passing to the reduced cascade the multiplicities associated to ends of adjacent levels, $\tilde{u}^i$ and $\tilde{u}^{i+1}$ do not necessarily match up, but by imposing there is a single flow time parameter T_i between adjacent levels still means we can think of $\mathbf{u}^{\prime}$ as living in a fiber product with virtual dimension.

To this end we first recall some conventions when it comes to J-holomorphic curves with ends on Morse-Bott critical submanifolds (in this case, tori). Consider $\tilde{u}^i$, for simplicity suppose its domain $\dot{\Sigma}_i$ is a punctured Riemann surface that is connected. Let $p_j^\pm$ label the positive/negative punctures, and the map $\tilde{u}^i$ is asymptotic to Reeb orbits (of some multiplicity) on Morse-Bott tori at each of its punctures. We wish to associate to $\tilde{u}^i$ a moduli space of curves that contain $\tilde{u}^i$ as an element and contains curves that are "close" to $\tilde{u}^i$. To this end we recall some conventions.

To each puncture $p_j^\pm$ of $\tilde{u}^i$, we can designated it as "fixed" or "free", and each choice of these designations leads to a different moduli space. The designation "free" means we consider J-holomorphic maps from $\dot{\Sigma}_i$ so that $p_j^\pm$ can land on any Reeb orbit with the same multiplicity on the same Morse-Bott torus at the corresponding end of $\tilde{u}^i$. For a puncture to be considered "fixed", we consider moduli space of J-holomorphic curves from $\dot{\Sigma}_i$ so that $p_j^\pm$ lands on a fixed Reeb orbits on a Morse-Bott torus with fixed multiplicity (the same Reeb orbit as $\tilde{u}^i$). Given a designation of "fixed" or "free" on punctures of $\tilde{u}^i$, we can then consider the moduli space of J holomorphic curves from $\dot{\Sigma}_i$ into $\mathbb{R} \times Y$ with the same asymptotic constraints as $\tilde{u}^i$ and living in the same relative homology class. We shall denote this moduli space as $\mathcal{M}_{\mathbf{c}}(\tilde{u}^i)$, using $\mathbf{c}$ to denote our choice of fixed/free ends. This moduli space has virtual dimension given by:

$$Ind(\tilde{u}^i) := -\chi(\tilde{u}^i) + 2c_1(\tilde{u}^i) + \sum_{p_j^+} \mu(\gamma^{q_{p_j^+}}) - \sum_{p_j^-} \mu(\gamma^{q_{p_j^-}}) + \frac{1}{2}\#\text{free ends} - \frac{1}{2}\#\text{fixed ends} \quad (2.11)$$

where χ is the Euler characteristic, c_1 the relative first Chern class, $\mu(-)$ is the Robbin Salamon index for path of symplectic matrices with degeneracies defined in [24]. We use the symbol γ to denote the Reeb orbit the end $p_j^\pm$ is asymptotic to, with multiplicity $q_{p_j^\pm}$.

Given a reduced cascade of currents, $\tilde{\mathbf{u}}^{\prime}$, let α denote the designation of "free"/"fixed" ends of $\tilde{u}^1$ at the $s \to +\infty$ end, and let β denote the "fixed"/"free" designation of $\tilde{u}^n$ at the $s \to -\infty$ end. Later we will see we can replace α and β with Morse-Bott ECH generators. In order to define the Fredholm index we need to assign free/fixed ends to the rest of the ends.

Convention 2.5.1. *If a non trivial curve u^i has an end landing on a critical point of f, then we consider that end to be fixed. If a trivial cylinder has one end on critical point of f, the other end must also land on the same critical point. We allow trivial cylinders with* *both ends free. If the trivial cylinder is at a critical point of f, we take the convention we can only designate one of its ends as fixed.*

Definition 2.5.2. *Let $\tilde{\mathbf{u}}^{\not{t}} = \{u^1, .., u^{n-1}\}$ denote a reduced cascade of currents of height 1. Let $ind(u^i)$ denote the Fredholm index of each of u^i. Note this makes sense since we have assigned free/fixed ends to all ends of u^i by our conventions above.*

Suppose there are $R_2, ..., R_{n-1} \in \mathbb{Z}$ distinct Reeb orbits approached by free ends as $s \to -\infty$ at each intermediate cascade level. Let us denote k_i and k_i' the number of free ends in each intermediate cascade level. e.g. elements in u^1 has k_2 free ends as $s \to -\infty$, and u^2 has k_2' free ends as $s \to +\infty$. Both counts of k_i and k_i', as well as R_i ignores "free" ends of fixed trivial cylinders, as such "free" ends are artificial to our convention. Now we define the cascade dimension

$$\begin{aligned}Ind(\tilde{\mathbf{u}}^{\not{t}}) :=& Ind(u^1) + .. + Ind(u^{n-1}) \\ &- [k_2'... + k_{n-1}'] - [k_2 + ... + k_{n-1}] + [R_2 + .. + R_{n-1}] + (n-2) - (n-1) - L\end{aligned}$$

where L is the number of intermediate cascade levels without free ends plus the number of intermediate cascade levels whose flow time is zero. Again in the count of L we ignore "free" ends coming from fixed trivial cylinders.

Observe for (reduced) cascades of height 1, we always have $k_i \geq R_i$ and $k_i' \geq R_i$.

We next explain how to define/compute the dimension of height k cascades. Let $\tilde{\mathbf{u}}^{\not{t}} = \{u^1, .., u^{n-1}\}$ denote a reduced cascade of currents of height N. We recall the difference between height one and height N cascade is that between cascade levels u^i and u^{i+1} we allow flow times $T_i = \infty$. We assign the free/fixed ends to u^i depending on whether they land on critical points of f as before. We can split a height N cascade into N height 1 cascades by partitioning the levels where the flow times are infinite. In particular we write $\tilde{\mathbf{u}}^{\not{t}} = \left\{\tilde{\mathbf{v}^1}^{\not{t}}, ..., \tilde{\mathbf{v}^N}^{\not{t}}\right\}$. Then the index of $\tilde{\mathbf{u}}^{\not{t}}$ is given by the sum of the indices of $\tilde{\mathbf{v}^i}^{\not{t}}$.

Here we come to the key transversality assumption of this paper. We first make sense of the notion of transversality.

Definition 2.5.3. *Let λ be a Morse-Bott contact form, whose Reeb orbits come in S^1 families. We say a λ compatible almost complex structure J is **good** if all reduced cascades of height one are tranversely cut out, which is defined below.*

Remark 2.5.4. We note the transversality conditions needed to count cascades given below are quite natural. However, since cascades have many parts the notation is bit complicated.

Definition 2.5.5. *Let $\tilde{\mathbf{u}}^{\not{t}} = \{u^1, .., u^{n-1}\}$ denote a reduced cascade of currents of height 1. We say $\tilde{\mathbf{u}}^{\not{t}} = \{u^1, .., u^{n-1}\}$ is **transversely cut out** if the conditions below are met.*

- *Each moduli space $\mathcal{M}_c(u^i)$ is transversely cut out with dimension given by the Fredholm index formula. Here the subscript c implicitly denotes the assignments of fixed and free ends we assigned to each end of u^i according to Convention 2.5.1. Note fixed trivial cylinders are assigned index zero.*

Suppose there are $R_2, ..., R_{n-1} \in \mathbb{Z}$ distinct Reeb orbits reached by free ends at each intermediate cascade level. We label them by $\gamma(i,j)$ where $j = 1, ..., R_i$, and i indexes which level we are referring to. For each $\gamma(i,j)$, we choose a negative puncture of u^{i-1} that is asymptotic to $\gamma(i,j)$. We call this puncture $p^-(i-1,j)$. The other negative ends of u^{i-1} that are asymptotic to $\gamma(i,j)$ are labelled $p^-(i-1,j,c,l)$, where $l = 1, 2.., n(\gamma(i,j), -)$. Next consider $\phi^{-T_{i-1}}(\gamma(i,j)))$. They are approached by positive punctures of u^i. For each $\phi^{-T_{i-1}}(\gamma(i,j)))$, we pick out a special free puncture $p^+(i,j)$. The remaining free positive ends of u^i that are asymptotic to $\phi^{-T_{i-1}}(\gamma(i,j)))$ are labelled $p^+(i,j,c,l)$ for $l = 1, ..., n(\gamma(i,j), +)$.

We next consider the evaluation maps. Given the collection of flow times $T_1, ..., T_{n-1}$. Let $\mathfrak{I} \subset \{1, .., n-1\}$ denote the subset for which $T_i > 0$, we consider the evaluation map

$$EV^- : \mathcal{M}(u^1) \times \mathcal{M}(u^2) \times ... \times \mathcal{M}(u^{n-2}) \to (S^1)^{R_2} \times (S^1)^{R_3} \times ... \times (S^1)^{R_{n-1}} \tag{2.12}$$

given by

$$(u'^1, ..., u'^{n-2}) \to (ev_1^-(u'^1), ev_2^-(u'^2), ..., ev_{n-2}^-(u'^{n-2})) \tag{2.13}$$

Here the evaluation is at the $p^-(i-1,j)$ puncture of u^{i-1}. We also consider the map

$$EV^+ : \mathcal{M}(u^2) \times \mathcal{M}(u^3) ... \times \mathcal{M}(u^{n-1}) \to (S^1)^{R_2} \times (S^1)^{R_3} \times ... \times (S^1)^{R_{n-1}} \tag{2.14}$$

given by:

$$(u'^2, ..., u'^{n-1}) \to (ev_2^+(u'^2), ..., ev_{n-1}^+(u'^{n-1})) \tag{2.15}$$

where the evaluation is at $p^+(i,j)$ of u^i. We consider the flow map

$$\Phi_f : (S^1)^{R_2} \times \mathbb{R}^* \times .. \times (S^1)^{R_{n-1}} \times \mathbb{R}^* \to (S^1)^{R_2} \times (S^1)^{R_3} \times ... \times (S^1)^{R_{n-1}}.$$

The notation $\mathbb{R}^*$ means the following: if $i \in \mathfrak{I}$ then we include a factor of $\mathbb{R}$ in the above product, otherwise we omit the factor. For $x_i \in S^1$ (i.e. a copy of S^1 among the product $(S^1)^{R_i}$), if $i \in \mathfrak{I}$ then the image of x_i under Φ_f is given by $\phi_f^{T_i'}(x_i)$. If the index i is not in $\mathfrak{I}$, then the image under Φ_f is x_i. We use the notation $\Phi_f \circ EV^+$ to denote the composition of the two maps, with domain $\mathcal{M}(u^2) \times \mathbb{R}^* \times \mathcal{M}(u^2) ... \times \mathcal{M}(u^{n-1}) \times \mathbb{R}^*$ and image $(S^1)^{R_2} \times (S^1)^{R_3} \times ... \times (S^1)^{R_{n-1}}$.

Let $\mathcal{K}_-$ denote the subset of $\mathcal{M}(u^1) \times \mathcal{M}(u^2) \times ... \times \mathcal{M}(u^{n-2})$ so that the ends $p^-(i,j)$ and $p^-(i,j,c,l)$ approach the same Reeb orbit. Let $\mathcal{K}_+$ denote the subset of $\mathcal{M}(u^2) \times \mathcal{M}(u^3) ... \times \mathcal{M}(u^{n-1})$ where $p^+(i,j)$ and $p^+(i,j,c,l)$ are asymptotic to the same Reeb orbit. Then

- *Near $\tilde{\mathbf{u}}^{\not{f}}$, both $\mathcal{K}_\pm$ are transversly cut out submanifolds.*

Then we can restrict $EV^\pm$ to $\mathcal{K}_\pm$, in particular the map $\Phi_f \circ EV^+$ admits a natural restriction to $\mathcal{K}_- \times \mathbb{R}^{|\mathfrak{I}|}$, our final condition is:

- *$\Phi_f \circ EV^+$ and EV^-, when restricted to $\mathcal{K}_+ \times (\mathbb{R})^{|\mathfrak{I}|}$ and $\mathcal{K}_-$ respectively are transverse at $\tilde{\mathbf{u}}^{\not{f}} = \{u^1, .., u^{n-1}\}$*

Assumption 2.5.6. *We assume we can choose J to be good so that all reduced cascades of current we encounter satisfy the transversality condition above.*

In particular, this implies all reduced cascades of currents live in a moduli space whose dimension is given by the index formula, and if such index is less than zero, then such cascades cannot exist.

We note that in general the transversality assumption is not automatic. In a reduced cascades of currents, all our curves are somewhere injective, but this is not enough. The issue lies in the fact that the fiber product that defines cascade can fail to have enough transversality. This is because all different levels of the cascade have the same J, and this J cannot be perturbed independently in each level. When the cascade is complicated enough, the same curve can appear multiple times in different levels, and this causes difficulty with the evaluation map. Consequently when there is not enough transversality for the naive definition of the universal moduli space of reduced cascades to be a Banach manifold, one usually needs some additional arguments.

However in simple enough cases we can still achieve the above transversality condition. This is the content of Theorem 2.2.2, which is proved in the Appendix.

2.6 ECH Index of Cascades

In this section we develop the analogue of ECH index one condition for cascades of currents. We shall see this will impose severe limits on currents that can appear in a cascade, provided transversality can be achieved.

To start the definition, we first consider one-level cascades, i.e. holomorphic curves from Morse-Bott tori to Morse-Bott tori. We want to define an index I so that for somewhere injective curves:

$$I(C) \geq dim\mathcal{M}(C) + 2\delta(C)$$

where $\mathcal{M}(C)$ denotes the moduli space of holomorphic curves C belongs in. Note the definition of $dim\mathcal{M}$ is ambiguous, because we need to specify which ends are "fixed" and which are "free". Our definition of I will depend on the type of end conditions imposed on our curve. The key to our construction will be the relative adjunction formula.

Relative adjunction formula in the Morse-Bott setting

Here we clarify what we mean by the intersection form Q. We first provide a provisional definition that is very much similar to regular ECH, then we show this definition descends to a more natural definition adapted to the Morse-Bott setting.

Let α, β be orbit sets. Observe here this means that they pick out discrete Reeb orbits (potentially with multiplicity) among the S^1 family of Reeb orbits. Then we can define the relative intersection formula as:

Definition 2.6.1. *We fix trivializations of Morse-Bott tori as we have specified, and denote it by τ. Given α, β orbit sets, given $Z, Z' \in H_2(\alpha, \beta, Y)$ we choose τ representatives S S' as before, then $Q_\tau(Z, Z')$ is defined as before as the algebraic count of intersections between S and S'.*

Because τ here provides a global trivialization of all Reeb orbits in a given Morse-Bott torus, the intersection Q doesn't depend on *which* specific Reeb orbit α or β picks out in a given Morse-Bott torus. We state the phenomenon in terms of a proposition:

Proposition 2.6.2. *Given orbit sets α, β and relative homology classes $Z, Z' \in H_2(\alpha, \beta)$. For definiteness let γ be a Reeb orbit in the $s \to +\infty$ end of α, let γ' be any translation of γ in its Morse-Bott torus, then using γ' to replace γ defines another orbit set α'. There exists corresponding relative homology classes $\hat{Z}, \hat{Z}' \in H_2(\alpha', \beta, Y)$ obtained by attaching a cylinder that connects between γ to γ' to ends of S and S' that are asymptotic to γ, then*

$$Q_\tau(Z, Z') = Q_\tau(\hat{Z}, \hat{Z}')$$

Proof. Choose τ representatives for Z, Z' which we write as S, S', then attach a cylinder connecting between γ to γ' to S and S'. In our trivialization the resulting surfaces are still τ representatives, and this process does not introduce additional intersections. □

The above proposition suggests Q_τ in the Morse-Bott case descends to a intersection number whose input is not $H_2(\alpha, \beta, Y)$ but a more general relative homology group adapted to the Morse-Bott setting.

Definition 2.6.3. *We define the relative homology classes $\mathcal{H}_2(\alpha, \beta, Y)$. Here α, β are collections of Morse-Bott tori, and multiplicities. For instance we can write $\alpha := \{(\mathcal{T}_i, m_i) | m_i \in \mathbb{Z}_{\geq 0}\}$ where $\mathcal{T}_i$ are Morse-Bott tori, and m_i are multiplicities. A element $Z \in \mathcal{H}_2(\alpha, \beta, Y)$ is a 2-chain in Y so that*

$$\partial Z = \alpha - \beta.$$

The above equality means the boundary (which includes orientation) of Z consists of Reeb orbits on Morse-Bott tori $\{\mathcal{T}_i\}$, and each $\mathcal{T}_i \in \alpha$ has a total of m_i Reeb orbits (counted with multiplicity) to which the ends of Z are asymptotic. Likewise for β. We define a equivalence relation on $\mathcal{H}_2(\alpha, \beta, Y)$, which we write as $Z \sim Z'$ as follows: Z and Z' are equivalent if there is a 3-chain W whose boundary takes the following form:

$$\partial W = Z - Z' + \{I \times S^1\}$$

where the collection $\{I \times S^1\}$ consists of 2 chains on Morse-Bott tori that appear in either α or β. We think of these 2-chains as an Reeb orbit (which we think of S^1) times an interval, I.

The idea is we consider 2-chains but allow their ends to slide along the Reeb orbits in the Morse-Bott family. The next proposition proves the relative intersection Q remains well defined.

Proposition 2.6.4. *Q_τ as defined above descends into a intersection form:*

$$Q_\tau : \mathcal{H}_2(\alpha,\beta,Y)\times \mathcal{H}_2(\alpha,\beta,Y)\to \mathbb{Z}.$$

Proof. For clarity we use $\hat{Q}_\tau$ to denote the intersection form defined in Definition 2.6.1. Suppose $Z, Z' \in \mathcal{H}_2(\alpha,\beta,Y)$, and suppose $Z'' \sim Z$. We pick a distinguished Reeb orbit γ_i for each Morse-Bott torus that appears in α,β, and chosen so that γ_i does not appear as a Reeb orbit in Z, Z' and Z''. We connect Reeb orbits in Z, Z' and Z'' to $\{\gamma_i\}$ counted with multiplicities using cyliners along each Morse-Bott tori to obtain $\hat{Z}, \hat{Z}', \hat{Z}''$. We then define

$$Q_\tau(Z,Z') := \hat{Q}_\tau(\hat{Z},\hat{Z}'').$$

Observe in the above $\hat{Q}_\tau$ is an intersection form defined on $H_2(\alpha',\beta',Y)$ where α' and β' are collections of Reeb orbits of the form $\{(\gamma_i,n_i)\}$. It suffices to prove $Q_\tau(Z'',Z') = Q_\tau(Z,Z'')$. To do this note the fact $Z\sim Z''$ in $\mathcal{H}_2(\alpha,\beta,Y)$ extends to an equivalence of $\hat{Z}\sim\hat{Z}''$ in $H_2(\alpha',\beta',Y)$, hence $\hat{Q}_\tau(\hat{Z}'',\hat{Z}') = \hat{Q}_\tau(\hat{Z},\hat{Z}')$, and hence the proof. □

We observe using the above reasoning the relative Chern class also descends to $\mathcal{H}_2(\alpha,\beta,Y)$. We state this in the form of a definition:

Definition 2.6.5. *Given $Z \in \mathcal{H}_2(\alpha,\beta,Y)$, we define the relative Chern class $c_\tau(Z)$ the same way as before: choose a representative S of Z that is embedded near the boundary. Let $\iota : Z \to Y$ be the inclusion, then consider the pullback of the contact structure $\iota^*\xi$ to Z, pick a section ψ of $\iota^*\xi$ that does not rotate with respect to τ near the end points and has transverse zeroes, then $c_\tau(Z)$ is the signed count of zeroes of ψ.*

Finally we define writhe the same way as before:

Definition 2.6.6. *Let C be a somewhere injective curve that is not a trivial cylinder. We assume at $s\to+\infty$ (resp. $-\infty$) it is asymptotic to orbit set α (resp. β). The trivialization specified in Theorem 2.17 gives an identification a neighborhood of each Reeb $\gamma\in\alpha,\beta$ with $S^1\times\mathbb{R}^2$, then using this we can define writhe of C as we had before in section 2.3.*

Remark 2.6.7. The definition of writhe depends crucially on the fact C is a holomorphic curve, and does not admit constructions as before where we can slide the Reeb orbits of α,β around and obtains a surface with same relative intersection number/Chern class.

Hence we are ready to state the relative adjunction formula.

Theorem 2.6.8. *If C is a simple J-holomorphic curve, then*

$$c_\tau(C) = \chi(C) + Q_\tau(C) + w_\tau(C) - 2\delta(C)$$

with the definition of relative chern class, relative intersection number, and writhe given above.

Proof. This is a purely topological formula. The same proof as in [29] follows through. □

Hence we would like to define a version of ECH index by applying the relative adjunction formula to the Fredholm index formula of holomorphic curves as in [29]. Recall then the proof of index inequality boils down to bounding the writhe of the J holomorphic curve in terms of various algebraic expressions involving the Conley Zehnder indices that the curve is asymptotic to. We turn to this writhe bound in the next subsection.

Writhe Bound

We recall the Fredholm index of a somewhere injective curve u depends on which end is free and which end is fixed. Hence we anticipate that the ECH index we assign to a holomorphic curve u will depend on which end is fixed and which end is free. The writhe inequality we prove shall take into account of the assignment of free and fixed ends. We note that this assignment of an index to a curve that depends on which end is free/fixed is somewhat artificial, but it will be less artificial once we use this index to define the ECH index of an entire cascade.

First we fix some conventions on Conley Zehnder indices. For a given Morse-Bott Torus $\mathcal{T}$ assume the J holomorphic curve has N ends that are positively (resp. negatively) asymptotic to Reeb orbits on this torus. They are asymptotic to the individual Reeb orbits labelled $R_1, .., R_n$. Writhe bound is a local computation so we only consider a particular Reeb orbit, called R_1. Assume k ends of C are asymptotic to R_1. They have multiplicity $q_1, .., q_k$. We adopt the following convention on Conley Zehnder indices.

Convention 2.6.9. *Recall for positive Morse-Bott torus $\mu = 1/2$. We declare $\mu_+ = 1$, $\mu_- = 0$. For negative Morse-Bott torus we declare $\mu_+ = 0$, $\mu_- = -1$.*

This has the following significance: for a curve with free end as $s \to +\infty$ landing in a Morse-Bott torus (regardless of whether it is positive or negative torus), the Conley Zehnder index term in the Fredholm index formula associated to this end is μ_+ (the specific value depends on the positive/negative Morse-Bott torus as above), and the Conley Zehnder index term assigned to fixed end is μ_-. Conversely, at the $s \to -\infty$ end we assign μ_- to free ends and μ_+ to fixed ends.

Using the above conventions given a somewhere injective holomorhic curve u, we assign its total Conley-Zehnder index denoted by $CZ^{Ind}(u)$ according to the convention above. The goal of the writhe inequality is to come up with another Conley-Zehnder index term $CZ^{ECH}(u)$ so that the total writhe of u is bounded above by

$$wr_\tau(u) \leq CZ^{ECH}(u) - CZ^{Ind}(u) \tag{2.16}$$

By way of convention we will use $CZ^*(R_1, \pm\infty)$ where $* = Ind, ECH$ to denote the Conley Zehnder index we should assign to the free/fixed ends approaching R as $s \to \pm\infty$

Positive Morse-Bott tori

Theorem 2.6.10. *In the case of positive Morse-Bott torus, $s \to -\infty$, if ξ_i is an end of u with covering multiplicity q_i and u is not the trivial cylinder, we have the following inequality*

$$\eta(\xi_i) \geq 1 \quad \text{(single end winding number)}.$$

For single end writhe, we have:

$$w(\xi_i) \geq \eta(\xi_i)(q_i - 1).$$

Note this holds true for trivial cylinders (as long as it's somewhere injective).

Let ξ_1 and ξ_2 be two braids that correspond to two distinct ends of u that approach the same Reeb orbit, with multiplicities q_i and winding numbers η_i, then:

$$l(\xi_1, \xi_2) \geq min(q_1\eta_2, q_2\eta_1)$$

Note this holds if one of the ends ξ_i came from a trivial cylinder.

And finally to calculate the writhe of all ends approach the same Reeb orbit, $w(\xi)$, let ξ denote the total braid and ξ_i the various components coming from incoming ends of u (this holds for both $s = \pm\infty$):

$$w(\xi) = \sum_i w(\xi_i) + \sum_{i \neq j} l(\xi_i, \xi_j)$$

In the case of $s \to +\infty$, using the exactly the same notation, we have the following inequalities:

$$\eta(\xi_i) \leq 0$$

$$w(\xi_i) \leq \eta(\xi_i)(q_i - 1) \quad \text{for single end writhe}$$

$$l(\xi_1, \xi_2) \leq max(q_1\eta_2, q_2\eta_1)$$

Proof. (Sketch) The proof constitutes an amalgamation of existing results in the literature. The key result is an description of asymptotics of ends of holomoprhic curves on Morse-Bott torus [57]. Namely, near the $s \to +\infty$ end of u, the s constant slice of $\{s\} \times Y$ of u can be described as follows. We can choose a neighborhood of trivial cylinder $\mathbb{R} \times \gamma$ as $\mathbb{R} \times S^1 \times \mathbb{R}^2$ where s is the symplectization direction, t is the variable along the Reeb orbit and $\{0\} \times \mathbb{R}^2$ is the contact structure along the Reeb orbit, then we can write an end ξ_i of u as

$$u(s,t) = (qs, qt, \sum_{i=1}^{n} e^{\lambda_i s} e_i(t)) \tag{2.17}$$

where λ_i and e_i are respectively the (negative) eigenvalues and corresponding eigenfunctions of the operator $A(t) : L^2(S^1, \mathbb{R}^2) \to L^2(S^1, \mathbb{R}^2)$ coming from the linearization of the Cauchy Riemann operator, which can be written as

$$A(t) = -J\partial_t - S$$

With this normal form, the winding number bound comes from combining the results in [24] about the meaning of Robbin-Salamon index and results in [27] relating Conley-Zehnder indices to crossing of eigenvalues. The relations on writhe and linking number come from direct modifications from the proofs in [29], once we realize that locally the braids can be described by Equation 2.17. □

Next we move to use these relations to prove writhe bound. As in the case of ECH, equality of the writhe bound implies certain partition conditions, which we will carefully state.

Proposition 2.6.11 (link, $-\infty$, positive Morse Bott torus)**.** *Consider the J holomorphic curve u with negative ends on a Reeb orbit γ. We have k_{free} free ends of multiplicity q_i^{free}, and k_{fixed} fixed ends with multiplicity q_i^{fixed} and of total multiplicity N_{fixed}. The writhe bound reads*

$$w(\xi) \geq -\sum_{i=1}^{k_{free}+k_{fixed}} \eta_i + \sum_{i,j}^{k_{free}+k_{fixed}} min(\eta_i q_j, \eta_j q_i) \geq (N_{free} - 1 + N_{fixed}) - (k_{fixed})$$

with equality holding implying there can be only free/fixed ends at this Reeb orbit. If there are only fixed ends the partition conditions is (n), and if there are only free ends the partition condition is (n) or $(1, n-1)$.

Proof. We have the respective bounds

$$-k_{free} + \sum_{i}^{k_{free}} min(\eta_i q_j, \eta_j q_i) \geq N_{free} - 1$$

and

$$-k_{fix} + \sum_{i,j}^{k_{fix}} min(\eta_i q_j, \eta_j q_i) \geq N_{fix} - k_{fixed}$$

and cross terms will imply strict inequality, hence only free or fixed term appears. In the case of only fixed points, we see the only way equality can hold is with partition condition (n). Similar considerations produces the partition conditions for free ends. □

Proposition 2.6.12 (link, ∞, positive Morse Bott Torus)**.** *In the $s \to +\infty$ end, consider the J holomorphic curve u with ends on a Reeb orbit γ. We have k_{free} free ends of multiplicity q_i^{free}, and k_{fixed} fixed ends q_i^{fixed} of total multiplicity N_{fixed}:*

$$w(\xi) \leq -\sum_{i=1}^{k_{free}+k_{fix}} \eta_i + \sum_{i,j}^{k_{free}+k_{fix}} max(q_j\eta_i, q_i, \eta_j) \leq N_{free} - (k_{free}).$$

The partition condition implies $(1, ..., 1)$ on the free ends.

Proof. We see that $lhs \leq 0$, and $RHS = 0$ iff the free end satisfies partition conditions $(1, ...1)$; there are no requirements on fixed ends. □

Negative Morse-Bott tori

In this subsection we take up the analogous writhe bounds for negative Morse-Bott tori.

Theorem 2.6.13. *In the case of negative Morse Bott torus, $s \to -\infty$, we have the following inequalities:*

If ξ_i is an end of u and u is not the trivial cylinder, we have the following inequality:

$$\eta(\xi_i) \geq 0$$

For writhe of a single end, with covering multiplicity q_i, we have:

$$w(\xi_i) \geq \eta(\xi_i)(q_i - 1)$$

Note this holds for the case of a trivial cylinder.

Let ξ_1 and ξ_2 be two braids that correspond to two distinct ends of u that approach the same Reeb orbit, with multiplicities q_i and winding numbers η_i, then:

$$l(\xi_1, \xi_2) \geq min(q_1\eta_2, q_2\eta_1)$$

Note this holds if one of the ends ξ_i came from a trivial cylinder.

And finally to calculate the writhe of all ends approach the same Reeb orbit, $w(\xi)$, let ξ denote the total braid, and ξ_i the various components coming from incoming ends of u (this holds for both $s = \pm\infty$):

$$w(\xi) = w(\xi_i) + \sum_{i \neq j} l(\xi_i, \xi_j)$$

In the case of $s \to +\infty$, we have the following inequalities

$$\eta(\xi_i) \leq -1$$

$$w(\xi_i) \leq \eta(\xi_i)(q_i - 1) \quad \textit{for single end writhe}$$

$$l(\xi_1, \xi_2) \leq max(q_1\eta_2, q_2\eta_1)$$

Proof. The exact same proof for the positive Morse-Bott torus except we use Robbin-Salamon index $\mu = -1/2$. □

Proposition 2.6.14 (link, $-\infty$,negative Morse Bott torus)**.** *Let u have ends asymptotic to γ on a negative Morse-Bott torus as $s \to -\infty$, suppose there are k_{free} free ends of multiplicity q_i^{free}, of total multiplicity N_{free}; suppose there are k_{fix} fixed ends each of multiplicity q_{fix}, of total multiplicity N_{fix}. Then we have the writhe bound:*

$$w(\xi) \geq -\sum_{i}^{k_{fix}+k_{free}} \eta_i + \sum_{i,j}^{k_{fix}+k_{free}} min(\eta_i q_j, \eta_j q_i) \geq -N_{free} - (-k_{free})$$

with equality enforcing partition condition $(1, .., 1)$ on free ends and no partition condition on fixed ends.

Proof. $\eta \geq 0$ so $lhs \geq 0$, $rhs = k_{free} - N_{free}$ so inequality holds, and equality if free ends has partition conditions $(1, .., 1)$, no restrictions on fixed ends. □

Proposition 2.6.15 (link, $+\infty$,negative Morse-Bott torus)**.** *Let u have ends asymptotic to γ on a negative Morse-Bott torus as $s \to +\infty$, suppose there are k_{free} free ends of multiplicity q_i^{free}, of total multiplicity N_{free}; and suppose there are k_{fix} fixed ends each of multiplicity q_{fix}, of total multiplicity N_{fix}.*

$$w(\xi) \leq -\sum_{i}^{k_{fix}+k_{free}} \eta_i + \sum_{i,j}^{k_{fix}+k_{free}} max(\eta_i q_j, \eta_j q_i) \leq -N_{fix} - N_{free} + 1 + k_{fix}$$

with equality enforcing only free or fixed ends. In the case of only fixed ends the partition condition is (n), and in the case of only free ends the partition condition is either (n) or $(n-1, 1)$.

Proof. We can split the sum into:

$$-\sum_{i}^{k_{free}} \eta_i + \sum_{i,j}^{k_{free}} min(\eta_i q_j, \eta_j q_i) \leq 1 - N_{free}$$

and

$$-\sum_{i}^{k_{fixed}} \eta_i + \sum_{i,j}^{k_{fixed}} min(\eta_i q_j, \eta_j q_i) \leq k_{fix} - N_{fix}.$$

Each of the above inequalities hold individually, and when there are both free and fixed ends, there are cross terms that make the inequality strict. As before, we can deduce the partition conditions directly from imposing the equality condition. □

Morse-Bott tori as ECH generators

Recall that for ECH of nondegenerate contact forms, the generators of the chain complex are orbit sets satisfying the condition that if an orbit is hyperbolic then it can only have multiplicity 1. There are analogues of this in Morse Bott tori. In Morse-Bott ECH, we think of the generators of the chain complex as collections of Morse-Bott tori with additional data, written schematically as:

$$\alpha = \{(\mathcal{T}_j, \pm, m_j)\}$$

and the differential as counting ECH index one height one J holomorphic cascades connecting between chain complex generators as above (which we will also call orbit sets). In the above definition m_j is the total multiplicity, which we think of as total multiplicity of Reeb orbits on $\mathcal{T}_j$ hit by the J holomorphic curves that have ends on this Morse-Bott torus on the top (resp. bottom) level of a (height 1) cascade. $\pm$ is additional information, which specifies how many ends of the J holomorphic curve landing on $\mathcal{T}_j$ are free/fixed. We see that this also depends on whether α appears as the top or bottom level of a J holomorphic cascade, and in context of our correspondence theorem free/fixed ends correspond to elliptic/hyperbolic orbits in the non-degenerate case. We state this explicitly in the next definition in which we also describe the expected correspondence between Morse-Bott ECH generators and nondegenerate ECH generators after the perturbation.

Definition 2.6.16. *We consider the case of positive Morse Bott tori. In the nondegenerate case we let γ_- denote the hyperbolic Reeb orbit that arises from perturbation with Conley Zehnder index 0, and γ_+ the elliptic orbit that arose out of the perturbation with Conley Zehnder index 1. Then the description of our Morse-Bott generator, say $(\mathcal{T}, \pm, m)$ (this is just one Morse-Bott torus, in general α will consist of a collection of such tori, we focus on an example for the sake of brevity) and its correspondence with ECH generators in the perturbed non-degenerate case is given by:*

a. positive side $s \to \infty$,

i) The Morse-Bott generator $(\mathcal{T}, +, m)$ is defined to require all ends on $\mathcal{T}$ are free, with total multiplicity on the torus being m. In the perturbed nondegenerate case, this corresponds to ECH orbit set (γ_+, m). We observe the nondeg partition (θ positive) condition is $(1, .., 1)$, and the Morse-Bott partition condition from the writhe bound is $(1, ..1)$.

By the Conley-Zehnder index convention the ECH conley Zehnder index assigned to $(\mathcal{T}, +, m)$ is given by: $CZ^{ECH}((\mathcal{T}, +\infty, +, m)) = m$

ii) The Morse-Bott generator $(\mathcal{T}, -, m)$ there is one end on $\mathcal{T}$ that is fixed with multiplicity 1, on the critical point of f that corresponds to the hyperbolic orbit. The rest of the ends are free, and the total multiplicity of orbits on $\mathcal{T}$ is m. This corresponds to the orbit set $\{(\gamma_-, 1), (\gamma_+, m-1)\}$ in the nondegenerate case. Note the partition conditions between nondegenerate case and Morse-Bott case agree.

We also have $CZ^{ECH}((\mathcal{T}, +\infty, -, m)) = m - 1$.

b. In the case of negative ends, $s \to -\infty$,

i) The Morse-Bott generator $(\mathcal{T}, +, m)$ is defined to require all ends are fixed and asymptotic to the critical point of f corresponding to the elliptic orbits, and the total multiplicity is m. In the nondegenerate case this correspond to the orbit set (γ_+, m). We observe Morse-Bott and nondegenerate partition conditions agree, both being (m). By our conventions, $CZ^{ECH}(\mathcal{T}, +, m) = m$

ii) The Morse-Bott generator $(\mathcal{T}, -, m)$ requires there is a multiplicity 1 free end landing on $\mathcal{T}$, the remaining ends are fixed and are also required to land on the critical point corresponding to elliptic Reeb orbit. This corresponds to the orbit set $\{(\gamma_+, m-1), (\gamma_-, 1)\}$ in the nondegenerate case, and we have analogous partition conditions for both Morse-Bott and nondegenerate case. $CZ^{ECH}(\mathcal{T}, -, m) = m-1$

We observe $(\mathcal{T}, \pm, m)$ imposes different free/fixed end conditions, depending whether it appears as $s \to \pm\infty$ ends, however we should think of it as being the same generator in the chain complex, as is evidenced by the fact that it is identified to the same nondegenerate orbit set regardless of whether it appears at $+\infty$ or $-\infty$ end.

We also briefly summarize the analogous result for negative Morse-Bott torus.

Definition 2.6.17. *In the case of negative Morse Bott tori, we use γ_- to denote the elliptic Reeb orbit after perturbation of Conley Zehnder index -1, and let γ_+ denote the hyperbolic orbit after perturbation of Conley Zehnder index 0. Let $(\mathcal{T}, \pm, m)$ denote a Morse-Bott generator.*

a. *At the positive end as $s \to \infty$,*

   i) *$(\mathcal{T}, -, m)$ requires all ends fixed at the critical point of f corresponding to γ_-, corresponds to (γ_-, m) in nondegenerate case, both degenerate and nondegenerate case has partition conditions (m). $CZ^{ECH}((\mathcal{T}, -, m)) = -m$*

   ii) *$(\mathcal{T}, +, m)$ requires one end free with multiplicity 1, the rest have multiplicity $m-1$ fixed at the critical point of f corresponding to γ_-. The generator corresponds to $\{(\gamma_+, 1), (\gamma_-, m-1)\}$. $CZ^{ECH}((\mathcal{T}, +, m)) = -m+1$. Partition conditions match.*

b. *Negative end, as $s \to -\infty$,*

   i) *$(\mathcal{T}, -, m)$ has all ends free, of total multiplicity m. This corresponds to (γ_-, m) in the nondegenerate case. Partition conditions match. $CZ^{ECH}((\mathcal{T}, -, m)) = -m$.*

   ii) *$(\mathcal{T}, +, m)$ has one fixed end corresponding to the critical point of f at γ_+ of multiplicity one; the rest are free and of multiplicity $m-1$. This corresponds to the orbit set $\{(\gamma_+, 1), (\gamma_-, m-1)$. The partition conditions correspond, and $CZ^{ECH}((\mathcal{T}, +, m)) = -m+1$.*

We would also like a more general notion of ECH Conley Zehnder index for when there are more free/fixed ends than allowed by ECH generator conditions are above. To keep track of the more refined intersection theory information, we need to make our definition depend slightly on the behaviour of the J-holomorphic curve as its ends approach Reeb orbits on Morse-Bott tori. We consider a nontrivial somewhere injective holomorphic curve $u : \Sigma \to \mathbb{R} \times Y^3$. We isolate this into the following definition.

Definition 2.6.18. *Let $u : \Sigma \to \mathbb{R} \times Y^3$ be a nontrivial somewhere injective holomorphic curve. Let γ be a simple Reeb orbit on a positive Morse-Bott torus.*

a. *At the $s \to \infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}, then $CZ^{ECH}(\gamma) := N_{free}$.*

b. *At the $s \to -\infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}, then $CZ^{ECH}(\gamma) := N_{free}+N_{fixed}-1$.*

Similarly if γ is a simply covered Reeb orbit on a negative Morse-Bott torus.

a. *At the $s \to \infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}, then $CZ^{ECH}(\gamma) := -N_{fix}-N_{free}+1$.*

b. *At the $s \to -\infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}, then $CZ^{ECH}(\gamma) := -N_{free}$.*

Note the above definition agrees with that of the ECH Morse-Bott generator. Then let u be a somewhere injective J holomorphic curve with no trivial cylinder components, and we have chosen which ends of u are fixed/free. Then we define its ECH index using the above notion of ECH Conley-Zehnder index:

Definition 2.6.19. *We define the ECH index of u as:*

$$I(u) := c_\tau(u) + Q_\tau(u) + CZ^{ECH}(u) \tag{2.18}$$

Note the above definition not only depends on the relative homology class of u, it also depends on how the ends of u are distributed among the Reeb orbits (for information of free/fixed beyond that of the Morse-Bott ECH generators)- in particular we have to keep the information of not only how many free/fix ends land on a Morse-Bott torus, we also need to retain the information which ends are asymptotic to which Reeb orbit.

By using the writhe bound we recover directly

Proposition 2.6.20. *Let u be a J-holomorphic map satisfying the conditions above,*

$$Ind(u) \leq I(u) - 2\delta(u). \tag{2.19}$$

with equality enforcing partition conditions described in the writhe bound section.

We next include the case of trivial cylinders in our definition of ECH Conley-Zehnder index.

Definition 2.6.21. *Let γ be a simply covered Reeb orbit on a positive Morse-Bott torus. Let $u : \Sigma \to \mathbb{R} \times Y$ be a J-holomorphic curve with potentially disconnected domain. When we say trivial cylinders below, we allow trivial cylinders with higher multiplicities.*

a. *At the $s \to \infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}, then $CZ^{ECH}(\gamma) := N_{free}$. Here we allow holomorphic curves to be trivial cylinders.*

b. At the $s \to -\infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}. If at least one of the approaching ends is not that of a trivial cylinder, then $CZ^{ECH}(\gamma) := N_{free} + N_{fixed} - 1$. If all the approaching ends are trivial cylinders, then $CZ^{ECH} := N_{fixed}$.

Next let γ be a simply covered Reeb orbit on a negative Morse-Bott torus.

a. At the $s \to \infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}, If at least one of the approaching ends is not that of a trivial cylinder, then $CZ^{ECH}(\gamma) := 1 - N_{free} - N_{fix}$. If there are only trivial cylinders, then $CZ^{ECH} = -N_{fixed}$.

b. At the $s \to -\infty$ end, suppose k_{free} ends approach γ with total multiplicity N_{free}, and k_{fixed} ends approach γ with total multiplcity N_{fixed}. Then we set $CZ^{ECH}(\gamma) := -N_{free}$. This includes the case of trivial cylinders.

Proposition 2.6.22. *Let C be a J holomorphic current which can contain trivial cylinders. Each end in C is implicitly assigned "free" or "fixed", and recall the convention that we can at most designate one end of a trivial cylinder as fixed. With CZ^{ECH} as defined above, we have the inequality:*

$$Ind(C) \leq I(C) - 2\delta(C)$$

Proof. Let C be a J-holomorphic current of the form $\{(C_i, m_i\}$ where C_i are pairwise distinct. If C_i is nontrivial, and $m_i > 1$, then as in [29], we can consider m_i copies of C_i translated by m_i distinct factors in the symplectization direction. Then we can represent (C_i, m_i) as m_i distinct somewhere injective J-holomorphic curves. We do this for all nontrivial components of C. Each resulting end of C_i receives an assignment of "free/fixed", hence both sides of the inequality above are defined. (One can make all the copies of C_i coming from (C_i, m_i) have the same free/fixed assignments at their corresponding ends, but this won't be necessary.)

As before this boils down to writhe bounds at $s = +\infty$ and $s = -\infty$. We first consider γ a Reeb orbit on a positive Morse-Bott torus. We first consider the $s = +\infty$ case. Here for trivial cylinders $q_i = 1$ and the linking number is zero, so the same proof as before produces the writhe bound.

In the case $s \to -\infty$, let $N_{trivial}$ denote the multiplicity of trivial ends. Let $N_{trivial}$ denote the total multiplicity of trivial ends, fixed or free. First assume there is at least one nontrivial end. The apriori bound on writhe is:

$$w(\xi) \geq -\#\text{nontrivial ends} + \sum_{i,j \text{nontrivial ends}} min(q_i, q_j) + N_{trivial} \cdot (\#\text{nontrivial ends}).$$

With our new definition of CZ^{ECH}, we need to establish the writhe bound that

$$-\#\text{nontrivial ends} + \sum_{i,j \text{nontrivial ends}} min(q_i, q_j) + N_{trivial} \cdot (\#\text{nontrivial ends})$$
$$\geq N_{free} + N_{fixed} - 1 - (k_{fixed})$$

We use the superscript ^{T} and ^{NT} to distinguish whether the multiplicity is coming from trivial ends or nontrivial ends. But the writhe bound already established implies

$$-\#\text{nontrivial ends} + \sum_{i,j \text{nontrivial ends}} min(q_i, q_j) \geq N^{NT}_{free} + N^{NT}_{fixed} - 1 - k^{NT}_{fixed}$$

Then it suffices to establish that

$$N_{trivial} \cdot (\#\text{nontrivial ends}) \geq N^T_{free} + N^T_{fixed} - k^T_{fixed}$$

which always holds, hence the writhe bound continues to hold.

When there are only trivial cylinders, the writhe is automatically zero, likewise the writhe bound is trivially satisfied.

We next consider the case γ a Reeb orbit on a negative Morse-Bott torus. We first consider the $s \to -\infty$ case. Since the winding number η in this case is bounded below by zero, the writhe bound continues to hold even in the presence of trivial cylinders.

In the case of $s \to +\infty$, the computation is very much similar to the $-\infty$ end of a positive Morse-Bott torus. Assuming there is at least one nontrivial end

$$\begin{aligned} w(\xi) \leq &+ \#\text{nontrivial ends} + \sum_{i,j \text{nontrivial ends}} max(\eta_i q_j, \eta_j q_i) - N_{trivial} \cdot \#\text{nontrivial ends} \\ \leq &- N_{fix} - N_{free} + 1 + k_{fix}. \end{aligned}$$

With the previous writhe bound we have already proven

$$\#\text{nontrivial ends} + \sum_{i,j \text{nontrivial ends}} max(\eta_i q_j, \eta_j q_i) \leq -N^{NT}_{fix} - N^{NT}_{free} + 1 + k^{NT}_{fix}$$

hence suffices to prove

$$-N_{trivial} \cdot \#\text{nontrivial ends} \leq -N^T_{fix} - N^T_{free} + k^T_{fix}$$

but this follows directly from our assumptions.

In the case there are only trivial ends the total writhe is zero, and the writhe bound is achieved. □

We next establish the subadditivity property of the ECH index.

Proposition 2.6.23. *Let $\mathcal{C}_1 = \{(C_a, m_a)\}$ and $\mathcal{C}_2 = \{(C_b, m_b)\}$ denote two J-holomorphic currents, and C_a is never the same as C_b unless they are both trivial cylinders (they can be $\mathbb{R}$ translates of each other). Then their ECH indices satisfy*

$$I(\mathcal{C}_1 \cup \mathcal{C}_2) \geq I(\mathcal{C}_1) + I(\mathcal{C}_2) + 2\mathcal{C}_1 \cap \mathcal{C}_2. \tag{2.20}$$

In the above $\mathcal{C}_1 \cap \mathcal{C}_2$ counts the intersection with multiplicity of C_a with C_b. Note by intersection positivity each multiplicity is positive. Further by construction the intersection between trivial cylinders is zero.

Proof. We again apply the translation in the symplectization trick to represent nontrival currents (C_a, m_a) (resp. (C_b, m_b)) by m_a (rep. m_b) distinct somewhere injective curves. After relabelling we can also denote them by C_a (resp. C_b). We apply the adjunction inequality as in [29, 33] to obtain

$$\begin{aligned}&I(\mathcal{C}_1 \cup \mathcal{C}_2) - I(\mathcal{C}_1) - I(\mathcal{C}_2) - 2\#\mathcal{C}_1 \cdot \mathcal{C}_2\\ =&CZ^{ECH}(\mathcal{C}_1 \cup \mathcal{C}_2) - CZ^{ECH}(\mathcal{C}_1) - CZ^{ECH}(\mathcal{C}_2) - 2\sum_{a,b} l_\tau(C_a, C_b)\end{aligned}$$

Then this reduces to a local computation relating linking number and our choice of Conley-Zehnder indices. We take this up case by case. First consider γ a Reeb orbit on a positive Morse-Bott torus, consider the $s \to \infty$ end. In this case we have $CZ^{ECH}(\mathcal{C}_1 \cup \mathcal{C}_2) - CZ^{ECH}(\mathcal{C}_1) - CZ^{ECH}(\mathcal{C}_2) = 0$ and $l_\tau(C_a, C_b) \leq 0$. Hence all the contributions from this end is ≥ 0.

We next consider γ on a positive Morse-Bott torus at $s \to -\infty$ ends. Because how we assigned Conley-Zehnder indices depends on whether all the ends are trivial, we split into cases. In the case where all ends of $\mathcal{C}_1$ and $\mathcal{C}_2$ asymptotic to γ as $s \to -\infty$ are trivial, we have again $CZ^{ECH}(\mathcal{C}_1 \cup \mathcal{C}_2) - CZ^{ECH}(\mathcal{C}_1) - CZ^{ECH}(\mathcal{C}_2) = 0$ and the linking number vanishes. If one of them has non-trivial ends approaching γ (WLOG take this to be $\mathcal{C}_1$ and take $\mathcal{C}_2$ consists purely of trivial ends), then we have the Conley Zehnder contribution being

$$N^1_{free} + N^1_{fixed} - 1 + N^2_{fixed} - (N^1_{free} + N^2_{free} + N^1_{fixed} + N^2_{fixed} - 1) = -N^2_{free}$$

where we write N^1_{free} to denote the free ends coming from $\mathcal{C}_1$ etc. The linking number contribution is bounded below by $2(N^2_{fixed} + N^2_{free})$, hence the overall contribution is non-negative. The case where both $\mathcal{C}_1$ and $\mathcal{C}_2$ contains nontrivial ends at γ as $s \to -\infty$, then the Conley-Zehnder difference term is just -1, and the linking number term $2l_\tau(C_a, C_b) \geq 2$, hence once again the overall contribution is non-negative.

We next consider the case γ is a Reeb orbit in a negative Morse-Bott torus. This will be largely analogous to the positive Morse-Bott torus case. For $s \to -\infty$, we have the Coneley-Zehnder indices contribute zero, and $l_\tau(C_a, C_b) \geq 0$ as $s \to \infty$, hence the overall contribution is non-negative. We next consider γ as $s \to +\infty$. Again we break into cases because of trivial cylinders. In the case where all ends approaching γ from $\mathcal{C}_1$ and $\mathcal{C}_2$ are trivial cylinders, the Conley-Zehnder index contribution as well as the linking number is zero. Then in the case $\mathcal{C}_1$ has nontrivial ends but $\mathcal{C}_2$ has all ends trivial, then the Conley-Zehnder index contribution is given by $-N^2_{free}$, and the linking number $\sum 2l_\tau(C_a, C_b) \leq -2(N^2_{free} + N^2_{fixed})$, hence the overall contribution is nonnegative. Similarly in the case where both $\mathcal{C}_1$ and $\mathcal{C}_2$ have nontrivial ends, the difference of Conley-Zehnder index contribution is -1, whereas the linking number $2l_\tau(C_a, C_b) \leq -2$, hence the overall contribution is positive. Hence combining all of the above local inequalities we obtain the overall ECH index inequality. □

Multiple level cascades and ECH index

In this subsection we describe ECH index one cascades. We recall ECH index one cascades should come from degenerations of ECH index one curves, and in particular should respect partition conditions on the end points. In particular we should always keep in mind that ECH index one cascades should flow from a generator Morse-Bott ECH α_1 to another α_n, which includes the information of multiplicities of free/fixed ends that land on Morse-Bott tori.

Given any cascade $u^{\lightning}$ as given in our previous definition, we first turn it into a "cascade of currents": $\mathbf{u}^{\lightning} = \{u^1, .., u^{n-1}\}$. Then we can proceed to define the ECH index of $\mathbf{u}^{\lightning}$. The following is half definition half theorem, as in if this cascade is transverse and rigid and we glued it into a J holomorphic curve the ECH index of its homology class is given by the following calculation. Conversely, if u came from a cascade of curves that came from a degeneration of $I = 1$ holomorphic curve in the λ_δ setting, then our definition of I for the cascade of current will also be one.

Definition 2.6.24. *Let $\mathbf{u}^{\lightning} = \{u^1, ..., u^{n-1}\}$ be a height 1 cascade of currents. Let its positive asymptotics be denoted by α_1 and negative asymptotics be denoted by α_n, both Morse-Bott ECH generators. We can then define the ECH index for the cascade of currents as:*

$$I(\mathbf{u}^{\lightning}) = c_1(\mathbf{u}^{\lightning}) + Q_\tau(\mathbf{u}^{\lightning}) + CZ^{ECH}(\mathbf{u}^{\lightning}). \tag{2.21}$$

The CZ^{ECH} index term for cascade is just the ECH index terms of α_1 and α_n, which corresponds to the nondegenerate ECH Conley Zehnder index once we have identified free/fixed ends with elliptic/hyperbolic orbits. The cascade Chern class and relative intersection terms are just the sum of the Chern class of each of the levels, i.e.

$$c_1(\mathbf{u}^{\lightning}) := c_1(u^1) + ... + c_1(u^{n-1})$$

and

$$Q_\tau(\mathbf{u}^{\lightning}) := Q_\tau(u^1) + ... + Q_\tau(u^{n-1})$$

We would like to compare the ECH index of cascade to the Fredholm index of the reduced version, because then with enough transversality we would be able to rule out certain configurations of cascade of ECH index one by index reasons. To this end, we decompose the ECH index of a cascade into ECH index of its constituents, as follows:

Proposition 2.6.25. *We assume all ends of $u^2, .., u^{n-2}$ are free, and all ends of u^1 and u^{n-1} are considered free except those mandated by α_1 and α_n, and we recall our conventions on trivial cylinders with only one fixed end. Then let $R'_{pos,i+1}$ denote the number of distinct Reeb orbits on positive Morse-Bott tori approached by nontrivial ends of u^i as $s \to -\infty$, and let $V'_{pos,i+1}$ denote the total multiplicity of Reeb orbits on positive Morse-Bott tori approached by u^i at the $s \to -\infty$, so that at these Reeb orbits there are only trivial ends as $s \to -\infty$. Similarly we let $R'_{neg,i}$ denote the number of distinct Reeb orbits on negative Morse-Bott tori* *approached by nontrivial ends of u^i as $s \to +\infty$, and let $V'_{neg,i}$ denote the total multiplicity of Reeb orbits on negative Morse-Bott tori approached by u^i at the $s \to +\infty$, so that at these Reeb orbits there are only trivial ends as $s \to +\infty$. Then we have*

$$\begin{aligned} I(\mathbf{u}^{\lightning}) =& I(u^1) + ... + I(u^{n-1}) \\ & - R'_{pos,2} - ... - R'_{pos,n-1} - V'_{pos,2} - ... - V'_{pos,n-1} \\ & - R'_{neg,2} - .. - R'_{neg,n-1} - V'_{neg,2} - .. - V'_{neg,n-1} \end{aligned}$$

Proof. Follows directly from definition of ECH Conley Zehnder index. □

Remark 2.6.26. Note the assignment of free/fixed end points for calculation of ECH index purposes is different from when we defined free/fixed punctures in the calculation of the Fredholm index.

Remark 2.6.27. We remark the above formula makes sense in the case our cascade consists purely of a chain of cylinder at a critical point. If it started at the minimum of f, the trick is to notice by our convention all trivial cylinders below it are considered free.

In order to compare $I(\mathbf{u}^{\lightning})$ and $Ind(\tilde{\mathbf{u}}^{\lightning})$, we first define

$$\begin{aligned} I(\tilde{\mathbf{u}}^{\lightning}) :=& I(\tilde{u^1}) ... + I(\widetilde{u^{n-1}}) \\ & - R'_{pos,2} - ... - R'_{pos,n-1} - V'_{pos,2} - ... - V'_{pos,n-1} \\ & - R'_{neg,2} - .. - R'_{neg,n-1} - V'_{neg,2} - .. - V'_{neg,n-1} \end{aligned}$$

by removing all multiple covers of nontrivial curves. Note we have

$$I(\tilde{\mathbf{u}}^{\lightning}) \leq I(\mathbf{u}^{\lightning}) \tag{2.22}$$

with equality holding only if $\mathbf{u}^{\lightning}$ is already reduced. Next we compare $Ind(\tilde{\mathbf{u}}^{\lightning})$ and $I(\tilde{\mathbf{u}}^{\lightning})$.

Proposition 2.6.28. $Ind(\tilde{\mathbf{u}}^{\lightning}) \leq I(\tilde{\mathbf{u}}^{\lightning}) - 2\delta(\tilde{\mathbf{u}}^{\lightning}) - 1$

Proof. We make a term-wise comparison, e.g. we compare

$$Ind(\tilde{u}^i) - k'_{i+1} - k_{i+1} + R_{i+1} \tag{2.23}$$

and

$$I(\tilde{u}^i) - 2\delta(\tilde{u}^i) - R'_{pos,i+1} - V'_{pos,i+1} - R'_{neg,i+1} - V'_{neg,i+1}. \tag{2.24}$$

Note there are two different conventions by which we assigned "free" and "fixed" ends to ends of curves appearing in the cascade, we will refer to them respectively as the Fredholm convention and the ECH convention.

We further refine our notation to $k_{pos,i+1}, k_{neg,i+1}, k'_{pos,i+1}, k'_{neg,i+1}$ to denote the number of ends among the k_i and k_{i+1} ends that land on positive/negative Morse-Bott tori, i.e. we have $k_i = k_{pos,i} + k_{neg,i}$.

We first restrict to $1 < i < n-1$ To compare these two terms, we first decompose $\tilde{u^i} = C_i \cup T_{free,i} \cup T_{fixed,i}$, where C_i is a collection of nontrivial somewhere injective curves, $T_{free,i}$ is a collection of free trivial cylinders according to Fredholm convention, and $T_{fixed,i}$ is a collection of fixed cylinder according to the Fredholm index convention. Assume C_i has $l_{free,i}$ free ends, and $l_{fixed,i}$ ends according to Fredholm convention, then we have:

$$Ind(C_i \cup T_{free,i} \cup T_{fixed,i}) + l_{fixed,i} \leq I(C_i \cup T_{free,i} \cup T_{fixed,i}) - 2\delta(C_i \cup T_{free,i} \cup T_{fixed,i}) - |T_{fixed,i}|$$

We may at later points further refine the notation to $l_{fixed,pos/neg,\pm,i}$ to indicate fixed ends at positive/negative Morse-Bott tori, at positive/negative ends. Note $T_{fixed,i}$ is regarded as free cylinders when we measure its ECH index. $|T_{fixed,i}|$ denotes the total number of fixed trivial cylinders that appear in this level.

We will also later refine our notation to distinguish $T_{fixed/free,pos/neg,i}$ for trivial cylinders on positive/negative Morse-Bott tori.

We next consider the case for $i = 1$. We can decompose as before $\tilde{u}^1 = C_1 \cup T_{free,1} \cup T_{fixed,1} \cup T'_{fixed,1}$. We explain the notation. C_1 is a collection of nontrivial somewhere injective holomorphic curves. The information of Morse-Bott generator α_1 tells us which of C_1 should already be considered as fixed as $s \to \infty$. There are additionally l_{fixed} ends of C that we count as fixed when we compute its Fredholm index because they land on critical points of f. $T_{free,1}$ is a collection of free cylinders. $T_{fix,1}$ is a collection of fixed trivial cylinders that come from requirements of α_1. Each positive Morse Bott torus can only have one of these, and they must all be multiplicity 1. $T'_{fixed,1}$ is a collection of trivial cylinders that don't come from requirements of α_1 but also happen to land on a critical point of f. The index inequality we have gives:

$$Ind(C_1 \cup T_{fixed,1} \cup T_{free,1} \cup T'_{fixed,1}) + l_{fixed,1} \leq I(C_1 \cup T_{fixed,1} \cup T_{free,1} \cup T'_{fixed,1}) - 2\delta(\tilde{u}^1) - |T'_{fixed,1}|$$

where for the purpose of computing ECH index we have counted elements of $T'_{fixed,1}$ as free cylinders.

Similarly for the $i = n-1$ level. As before we can decompose $\tilde{u}^{n-1} = C_{n-1} \cup T_{free,n-1} \cup T_{fixed,n-1} \cup T'_{fixed,n-1}$ with the same convention as before. Here we only need to prove:

$$\begin{aligned} &Ind(C_{n-1} \cup T_{free,n-1} \cup T_{fixed,n-1}) + l_{fixed,n-1} \\ \leq &I(C_{n-1} \cup T_{free,n-1} \cup T_{fixed,n-1}) - 2\delta(\tilde{u}^{n-1}) - |T'_{fixed,n-1}| \end{aligned}$$

which holds by the one-level ECH index inequality. When we take the difference between $I(\tilde{u^{\natural}})$ and $Ind(\tilde{u^{\natural}})$, we can break down their difference into the following form:

$$\begin{aligned} I(\tilde{\mathbf{u}}^{\natural}) = &I(\tilde{u^1}) ... + I(\tilde{u}^{n-1}) \\ &- R'_{pos,2} - ... - R'_{pos,n-1} - V'_{pos,2} - ... - V'_{pos,n-1} \\ &- R'_{neg,2} - .. - R'_{neg,n-1} - V'_{neg,2} - .. - V'_{neg,n-1} \end{aligned}$$

and the index term can be re written as

$$Ind = \sum_i ind(\tilde{u}^i) - \sum_{i=2,\dots,n-1}(k_{pos,i}+k'_{pos,i}-R_{pos,i}) - \sum_{i=2,\dots,n-1}(k_{neg,i}+k'_{neg,i}-R_{neg,i}) - 1 - L$$

If we take their difference, and take advantage of the inequalities we proved in the previous paragraphs, we get:

$$\begin{aligned} I - Ind =& \sum I(\tilde{u}^i) - ind(\tilde{u}^i) + \sum_{i=2,\dots,n-1}((k_{pos,i}+k'_{pos,i}-R_{pos,i}-R'_{pos,i}-V'_{pos,i}) \\ &+ \sum_{i=2,\dots,n-1}(k_{neg,i}+k'_{neg,i}-R_{neg,i}-R'_{neg,i}-V'_{neg,i}) + L + 1 \\ \geq& 2\delta(\tilde{\mathbf{u}}^{\not{t}}) + \sum_{i=2,..,n-2}(l_{fixed,i}+|T_{fixed,i}|) + l_{fixed,1} + l_{fixed,n-1} + |T'_{fixed,1}| + |T'_{fixed,n-1}| \\ &+ \sum_{i=2,\dots,n-1}((k_{pos,i}+k'_{pos,i}-R_{pos,i}-R'_{pos,i}-V'_{pos,i}) \\ &+ \sum_{i=2,\dots,n-1}(k_{neg,i}+k'_{neg,i}-R_{neg,i}-R'_{neg,i}-V'_{neg,i}) + L + 1 \end{aligned}$$

It suffices to prove the above expression is bounded below by one. It suffices to prove

$$\begin{aligned} &\sum_{i=2,\dots,n-1} R_{pos,i}+R'_{pos,i}+V'_{pos,i} + \sum_{i=2,\dots,n-1} R_{neg,i}+R'_{neg,i}+V'_{pos,i} \\ \leq& \sum_{i=2,..,n-2}(l_{fixed,i}+|T_{fixed,i}|) + l_{fixed,1} + l_{fixed,n-1} + |T'_{fixed,1}| + |T'_{fixed,n-1}| \\ +& \sum_{i=2,\dots,n-1}(k_{pos,i}+k'_{pos,i}) + \sum_{i=2,\dots,n-1}(k_{neg,i}+k'_{neg,i}) + L \end{aligned}$$

We break down the above inequality into several components. We first observe for $i = 2,..,n-2$ we have

$$R'_{pos,i+1} + V'_{pos,i+1} \leq l_{fixed,pos,-\infty,i} + l_{fixed,pos,+\infty,i} + |T_{fixed,i}| + k_{pos,i+1} + k'_{pos,i+1} - R_{pos,i+1}$$

We first observe the multiplicities counted by $R'_{pos,i+1}$ and $V'_{pos,i+1}$ are disjoint - if a Reeb orbit appear in considerations of $R'_{pos,i+1}$ then it is not considered for $V'_{pos,i+1}$ and vice versa. Multiplicities counted by $V'_{pos,i+1}$ are contained in $k_{pos,-\infty,i+1}$ and $|T_{fixed,i+1}|$, and the Reeb orbits counted by $R'_{pos,i+1}$ are contained in the ends counted by $l_{fixed,pos,-\infty,i+1}$ and $k_{pos,-\infty,i+1}$. We observe for this range of i, we only needed to use the fixed ends of C_i in $l_{fixed,i}$ as $s \to -\infty$ to achieve this inequality, and the prescence of $l_{fixed,i,pos,+\infty}$ will make this inequality strict by that factor. Finally we observe $k'_{pos,i+1} - R_{pos,i+1} \geq 0$. This concludes this inequality.

We next consider the case for $i = 1$ for positive Morse-Bott tori, i.e. we consider the inequality

$$R'_{pos,2} + V'_{pos,2} + R_{pos,2} \leq l_{fixed,pos,-\infty 1} + l_{fixed,pos,+\infty 1} + |T'_{fixed,pos,1}| + k_{pos,2} + k'_{pos,2}$$

This inequality does not hold in general. We first observe $k'_{pos,2} - R_{2,pos} \geq 0$, and the Reeb orbits counted by $R'_{pos,2}$ are included in $k_{pos,2}$ and $l_{fixed,pos,-\infty,1}$. The issue for $V'_{pos,2}$ is slightly more subtle, because each positive Morse-Bott torus can contain one fixed trivial cylinder that is not included in $|T'_{fixed,pos,1}|$, hence a Reeb orbit counted by $V'_{pos,2}$ that does not necessarily appear on the right hand side. If we follow this trivial cylinder downwards, if we encounter an end of a non-trivial J-holomorphic curve that approaches this Reeb orbit at $s \to \infty$, then it will contribute to $l_{fixed,pos,+\infty,i}$ terms in one of the lower levels. And this $l_{fixed,pos,+\infty,i}$ term was not used in our previous computations, so after we add up all the terms in the inequality, the overall inequality will still hold.

If we go downwards and do not see a nontrivial end, then there must be a trivial cylinder at the bottom level of the cascade making a contribution to $T'_{fixed,,pos,n-1}$ located at this specific Reeb orbit on this positive Morse-Bott torus. This cylinder counted by $T'_{fixed,pos,n-1}$ is not used anywhere else in any of our other inequalities, so makes up for the deficit coming from the $i = 1$ inequality.

Finally we consider the terms on the last level concerning the positive Morse-Bott tori contributing to our inequality. This is just

$$|T'_{fixed,pos,n-1}| \geq 0$$

which holds trivially. $|T'_{fixed,pos,n-1}|$ being nonzero does not necessarily mean our inequality is strict, as some of these may be borrowed to make the inequality hold on the $i = 1$ level as per above.

We now repeat the analogous series of inequalities concerning negative Morse-Bott tori. We first prove the inequalities

$$R_{neg,i} + R'_{neg,i} + V'_{neg,i} \leq k_{neg,i} + k'_{neg,i} + l_{fixed,neg,+\infty,i} + |T_{fixed,neg,i}|$$

for i in range $2, ..., n-2$. We have as before that $R_{neg,i} \leq k_{neg,i}$. Similarly the count of orbits in $R'_{neg,i}$ is included $k'_{neg,i}$ and $l_{fixed,neg,+\infty,i}$, and the count of $V'_{neg,i}$ is included among $T_{fixed,neg,i}$ and $k'_{neg,i}$. This concludes the proof of this inequality.

Next we focus on the $i = n-1$ case. We consider the inequality

$$R'_{neg,n-1} + V'_{neg,n-1} + R_{neg,n-1} \leq l_{fixed,neg,+\infty,n-1} + |T'_{fixed,neg,n-1}| + k_{neg,n-1} + k'_{neg,n-1}$$

This does not always hold, as before we first observe $k_{pos,n-1} - R_{neg,n-1} \geq 0$, and $R'_{neg,n-1}$ is included in $l_{fixed,neg,+\infty,n-1}$ and $k_{neg,n-1}$. However each negative Morse-Bott torus can contain one fixed trivial cylinder not included in $T'_{fixed,neg,n-1}$. If we follow this trivial cylinder upwards, if we encounter an end of a non-trivial J-holomorphic curve that approaches this Reeb orbit at $s \to -\infty$, then it will contribute to $l_{fixed,neg,-\infty,i}$ terms in one of the upper levels. And this $l_{fixed,,neg,-\infty,i}$ term was not used in our previous computations, so after we add up the terms in the inequality, the overall inequality will still hold.

If we go upwards and do not see a nontrivial end, then there must be a trivial cylinder contributing to $T'_{fixed,neg,1}$ appearing at the very same Reeb oribt. This cylinder's contribution is not used up by any of our previous inequalities, so makes up for the deficit in the above inequality.

The $i = 1$ level terms for negative Morse-Bott tori is simply $|T'_{fixed,neg,1}| \geq 0$ which holds trivially. This inequality being strict does not necessarily imply the overall inequality is strict, by the mechanism discussed above.

Adding up the above inequalities we get the inequality in the proposition. $\square$

We now state some consequences of the ECH index one condition, assuming transversality can be satisfied.

Corollary 2.6.29. *Assuming J can be chosen to be good, and we have a height one cascade $u^{\lightning}$. Then we pass to cascade of currents $\mathbf{u}^{\lightning}$, the ECH index being one imposes the following conditions:*

a. $\mathbf{u}^{\lightning}$ is reduced.

b. All flow times are strictly positive.

c. All curves are embedded. Curves on the same level are disjoint.

d. Each level only has one nontrivial curve, the rest are trivial cylinders.

e. With the above choice of fixed/free ends, all curves obey partition conditions of free ends for ends that do not land on critical points. They obey the partition conditions for fixed ends for those that land on critical points of f.

f. For any nontrivial curve C appearing in the cascade of currents $\mathbf{u}^{\lightning}$:

- *If C appears in either u^1 or u^{n-1}, then its ends can appear on critical points of f only as mandated by α_1 or α_2. All other ends must avoid critical points of f.*
- *If C appears in a level between u^1 and u^{n-1}, its ends can only end on a critical point of f if this end is then connected by a fixed chain of trivial cylinders to fixed points mandated by α_1 or α_n. All other ends avoid critical points of f, and hence are free.*
- *Further, if we see a chain of fixed trivial cylinders connecting a positive or negative end of C to a critical point of f, suppose the fixed Reeb orbit is called γ. Then no nontrivial end may land on γ on any of the levels of the components of the chain of trivial cylinders in either $s \to +\infty$ or $s \to -\infty$. On the level where C is asymptotic to γ as $s \to \infty$ or $s \to -\infty$, the end of C is the only end that is asymptotic to γ as $s \to +\infty$ and $s \to -\infty$ respectively.*

g. *In particular, if C is a nontrivial curve in the cascade, and an end of C is asymptotic to γ, a Reeb orbit in the $s \to +\infty$ (resp. $-\infty$) end, then no other curve (or other ends of C) in the same level may be asymptotic to γ as $s \to +\infty$ (resp. $-\infty$).*

h. *If an end of a nontrivial curve C is asymptotic to γ with multiplicity > 1, as $s \to \infty$, and if we follow γ upwards, e.g. we consider C' in the level above which is asymptotic to γ as $s \to -\infty$. If all curves above C that are asymptotic to γ are trivial cylinders, then we cannot draw any conclusions aside from partition conditions of C. However, if after some chain of gradient flow lines a nontrivial curve C'' above C is asymptotic to $\phi^f_T(\gamma)$ as $s \to -\infty$ and is connected to the positive end of C at γ via a gradient flow, then by partition conditions both C and C'' can only be asymptotic to γ with multiplicity 1.*

Proof. All statements in the above proposition comes from taking all the inequalities in the previous proposition to be equalities. (a) comes $I(\mathbf{u}^{\mathfrak{t}}) = I(\tilde{\mathbf{u}}^{\mathfrak{t}})$. ($b$) comes from $L = 0$. (c) comes from $\delta(\mathbf{u}^{\mathfrak{t}}) = 0$. ($d$) comes from $Ind = 0$, otherwise the cascade lives in a moduli space of dimension greater than zero. (e) comes from the fact that violations of partition conditions for nontrivial curves would make the inequalities comparing Fredholm index to ECH index strict.

Next consider (f), for the nontrivial curves appearing in u^1 or u^{n-1}. We first consider the case of u^1. We observe all contributions to $l_{fixed,+\infty,1}$ from the $s \to +\infty$ must be zero for equality in 2.6.28 to hold. Similarly we observe that for u^{n-1} all contributions to $l_{fixed,-\infty,n-1}$ from the $s \to -\infty$ must be zero for equality to hold.

If C is a nontrivial curve between u^1 and u^{n-1}, we have to separate this into cases. We first assume it has a negative end landing on a critical point of f on a positive Morse-Bott torus. Then this end makes a contribution to $l_{fixed,pos,-\infty}$, and was used in our computation of inequality. Call this Reeb orbit γ, and consider levels below C that have nontrivial ends asymptotic to γ as $s \to +\infty$. Say this occurs on level i. If there are such curves, and if γ does not appear as a fixed end assigned by α_1 and connected to a trivial cylinder in u^1, then it is a appearance of $l_{fixed,pos,+\infty,i}$ that was not used in our proof of inequality in 2.6.28, hence the inequality is strict.

The case where γ appears in α_1 as a fixed end of a trivial cylinder is handled as follows. In the case there is a contribution to $T'_{fixed,pos,n-1}$ on the u^{n-1} level from a trivial cylinder at γ, then we can use the additional $l_{fixed,pos,+\infty,i}$ at γ to make the inequality strict. In the case $T'_{fixed,n-1}$ does not have a trivial cylinder at γ, then for multiplicity reasons the total multiplicity of nontrivial ends asymptotic to γ as $s \to +\infty$ in the entire cascade must be greater than equal to two. If they come from two different ends (potentially at different levels), then their contribution to $l_{fixed,pos,+\infty,*}$ (of various levels) is at least two, which makes the inequality in proposition 2.6.28 strict. If we only see a single nontrivial end approach γ as $s \to +\infty$ below u^1 level, then this end must have multiplicity ≥ 2, and this violation of writhe inequality also ensures the index inequality is strict.

If no nontrivial curves below C that are positively asymptotic to γ exist, then with the negative puncture of C landing at γ, the negative puncture is connected to the last level u^{n-1}

at γ via a chain of fixed trivial cylinders. If γ is a minimum of f, then this is a contribution to $|T'_{fixed,pos,n-1}|$ that was not considered in the proof of inequality. This will make the overall inequality strict if γ did not appear as a fixed end connected to a trivial cylinder in u^1. If γ did appear (as a fixed end mandated by α_1), then again for multiplicity reasons there is either an additional $l_{fixed,pos,+\infty i}$ contribution from $s \to +\infty$ ending on γ on one of the middle levels, or $|T'_{fixed,n-1}|$ at γ has multiplicity greater than or equal to two. Either case makes the index inequality strict.

However if γ is at a maximum of f, the inequality is not violated if this is a chain of trivial cylinders connecting to a fixed end mandated by α_n. If α_n assigns free ends to this chain of cylinders, then we have extra contributions to $T'_{fixed,pos,n-1}$ which make the index inequality strict (in this case α_1 cannot assign γ as a fixed end). Finally if this is indeed a chain of fixed trivial cylinders connecting to a fixed orbit mandated by α_n, then on the level where C appears no other nontrivial end may be asymptotic to γ as $s \to -\infty$, this is because if this is true, then we consider the inequality for C's level

$$R'_{pos,i+1} + V'_{pos,i+1} \leq l_{fixed,pos,-\infty,i} + |T_{fixed,pos,i}| + k_{pos,i+1} + k'_{pos,i+1} - R_{pos,i+1}$$

Both nontrivial ends at γ are counted once by $R'_{pos,i+1}$, but twice by $l_{fixed,pos,-\infty,i}$, which makes this inequality strict. This automatically imposes the partition condition (n) on this particular negative end of C. Further, down this chain of fixed trivial cylinders, all the way to α_n, no further lower levels may have non-trivial curves whose ends are asymptotic to γ as $\to -\infty$. This is clear for the lowest level u^{n-1}. We already argued $l_{fixed,pos,-\infty,n-1} = 0$, then all fixed ends landing on γ must be fixed ends assigned by α_n, then the partition conditions imposed by ECH index implies we cannot have both trivial and nontrivial ends at γ. On levels above the lowest level and below the level of C, this follows from the inequality

$$R'_{pos,i+1} + V'_{pos,i+1} \leq l_{fixed,pos,-\infty,i} + |T_{fixed,pos,i}| + k_{pos,i+1} + k'_{pos,i+1} - R_{pos,i+1}.$$

If we have both a trivial cylinder and an nontrivial end asymptotic to γ in the negative end, they make an overall contribution of 1 to the left hand side, but make a overall contribution of 2 to the right hand side by increasing $l_{fixed,pos,-\infty,i}$ and $|T_{fixed,pos,i}|$, hence making this inequality strict.

We next consider C has a positive end ending on a critical point of f. Call this Reeb orbit of γ. If γ is not a fixed Reeb orbit mandated by α_1, then this already makes a contribution to $l_{fixed,pos,+\infty,i}$ we did not use in the index inequality, which makes the overall inequality strict. If γ indeed appears in α_1 and is in fact connected to a trivial cylinder, then either this end of C is connected upwards to γ via a sequence of trivial cylinders, or there are more nontrivial ends above C that ends on γ as $s \to +\infty$, but this makes the index inequality strict due to multiplicity reasons (α_1 can only require a fixed end of multiplicity 1 at γ). Hence it must be the case C is connected to γ on the top level via sequence of fixed trivial cylinders, and no level above C have nontrivial ends approaching γ as $s \to +\infty$. If a curve above C has a negative end approaching γ, we are back to the previous case and this also makes the index inequality strict.

The case of negative Morse-Bott tori is similar to positive Morse-Bott tori but with the signs reversed, so we will not repeat it. We remark the proof of Negative Morse-Bott tori is independent of the proof of positive Morse-Bott tori because when we compute $|T'_{fixed,i}|$ the trivial cylinders at negative and positive Morse-Bott tori are independent of each other.

To prove (g) and (h). We already took care of the case a non-trivial curve that is asymptotic to a Reeb orbit corresponding to a critical point of f. We next consider the case of free ends. Let our curve be C in some level of the cascade and consider its $+\infty$ free ends asymptotic to positive Morse-Bott tori. We have $k'_{pos,i+1} = R_{pos,i+1}$, this implies each free Reeb orbit as $s \to +\infty$ is approached by a unique positive end of C. The ECH index also imposes partition conditions of $(1,..,1)$, hence this end is simply covered. Recalling $\mathbf{u}^{\not{t}}$ is reduced, any $s \to -\infty$ free end of curves above C arrived at by following the gradient flow is also simply covered. This proves (g) and (h) for positive Morse-Bott tori. The result for negative Morse-Bott tori holds by considering the negative free ends of C. □

We would also like a way to prove that provided our transversality conditions hold (i.e. J is good), J_δ-holomorphic curves of ECH index one degenerate into cascades of height one, as opposed to cascades of greater height. To do this we need a slight strengthening of the above index inequality where we allow fixed trivial cylinders with higher multiplicities.

Proposition 2.6.30. *Let α_1 and α_n be ECH Morse-Bott generators, except we relax the condition on multiplicities of fixed/free ends - they are allowed to be arbitrary. Let $u^{\not{t}}$ be a cascade of height one connecting from α_1 to α_n. Then we have the inequality*

$$Ind(\tilde{\mathbf{u}}^{\not{t}}) \leq I(\mathbf{u}^{\not{t}}) - 2\delta(\mathbf{u}^{\not{t}}) - 1$$

Proof. We repeat the proof of index inequality in Proposition 2.6.28 and observe the inequalities concerning the intermediate level curves continue to hold. The issue is in allowing fixed trivial cylinders of high multiplicities allowed by α_1 and α_n at the top and bottom levels. We first focus on what happens near positive Morse-Bott tori. For simplicity we fix γ a Reeb orbit corresponding to the hyperbolic orbit in a positive Morse-Bott torus and consider what happens to ends of holomorphic curves with fixed ends at γ. As we have seen above the problematic term comes from the inequality

$$R'_{pos,2} + V'_{pos,2} \leq k_{pos,2} + k'_{pos,2} - R_{pos,2} + l_{fixed,pos,-\infty,1} + l_{fixed,pos,+\infty,1} + |T'_{fixed,1}|,$$

where $V'_{pos,2}$ can contain fixed trivial cylinders mandated by α_1 that appear in $V'_{2,pos}$ but does not appear in $|T'_{fixed,1}|$. For simplicity we consider $T_{\gamma,fixed}$ appearing at γ of multiplicity N. In order for this to make a contribution to $V'_{2,pos}$ instead of $R'_{2,pos}$, we assume that u^1 has no nontrivial end that are asymptotic to γ as $s \to -\infty$. We recall we would like to prove an inequality of the form

$$I(u^{\not{t}}) - 1 \geq Ind(\tilde{\mathbf{u}}^{\not{t}}) + 2\delta(u^{\not{t}})$$

Consider for $i = 2, ..., n-1$, the nontrivial currents $(C_{i,j}, m_{i,j}) \subset u^i$, where we think of $m_{i,j}$ as the multiplicity of $C_{i,j}$ (since we are working in the nonreduced case). We assume each

$C_{i,j}$ has $l_{i,j}$ ends asymptotic to γ as $s \to \infty$, and suppose $C_{i,j}$ has total multiplicity $n_{i,j}$ asymptotic to γ as $s \to \infty$. Finally let $T_{fixed,n-1,\gamma}$ denote the number of trivial cylinders at the last level u^{n-1} at γ. We have the inequality

$$N - \sum_{i,j} m_{i,j} n_{i,j} \leq |T'_{fixed,n-1,\gamma}|.$$

Let's consider $I(C_{i,j})$, by virtue of it being nontrivial and the writhe inequality, $\sum_j I(C_{i,j}) \geq \sum_j (n_{i,j}+1)$. This is coming from the fact in order for the $C_{i,j}$ to exist its Fredholm index must be greater or equal to one, and at the ends of γ the ECH index is treated as free ends whereas the Fredholm index is treated as fixed ends. So in passing from u^i to $\tilde{u}^i$ we decreased the ECH index by at least $\sum_{i,j}(m_{i,j}-1)(n_{i,j}+1)$.

We next compare the ECH index of reduced cascade with its Fredholm index, in particular we consider the inequalities

$$I(\tilde{u^i}) - Ind(\tilde{u^i}) + R'_{pos,i+1} + V'_{pos,i+1} - [l_{fixed,i} + |T_{fixed,i}| + k_{pos,i+1} + k'_{pos,i+1} - R_{pos,i+1}] \geq 0$$

for $i = 2,..,n-2$. We have that by virtue of the writhe inequality occurring at γ across these levels, the γ orbit's contribution is that the left hand side is at least $\sum_j n_{i,j} - l_{i,j}$ bigger than the right hand side.

Finally, on the u^{n-1} level, we originally had the inequality

$$|T'_{fixed,n-1}| \geq 0$$

In the above inequality we have included the $|T'_{fixed,n-1,\gamma}|$ term coming from the last level in our cascade contributed by γ, and the writhe bound for this level also implies this there is also an excess of the index inequality of size $\sum_j n_{n-1,j} - l_{n-1,j}$.

Hence we can think of proving the index inequality as follows: there is a deficit of N at the top level contributed purely by γ, and by making the inequalities of the lower levels strict, we can make up for it. In passing from nonreduced to reduced curve, the "excess" of ECH index is bounded below by $\sum_{i,j}(m_{i,j}-1)(n_{i,j}+1)$. The excess of comparing ECH index of reduced curves $C_{i,j}$ to their Fredholm index coming from writhe inequality is given by $\sum_{i,j} n_{i,j} - l_{i,j}$, and the excess in the index inequality of various levels due to contributions to $l_{fixed,pos,+\infty,i}$ coming from γ is precisely $\sum_{i,j} l_{i,j}$. And on the last level the excess is given simply by $|T'_{fixed,n-1,\gamma}|$ Hence the excess due to γ is bounded below by

$$\sum_{i,j}(m_{i,j}-1)(n_{i,j}+1) + \sum_{i,j} n_{i,j} - l_{i,j} + \sum_{i,j} l_{i,j} + |T'_{fixed,n-1,\gamma}|$$

Using the fact $N - \sum_{i,j} m_{i,j} n_{i,j} \leq |T'_{fixed,n-1,\gamma}|$, we see the excess outweighs the deficit at the top level, so fixed trivial cylinders at γ will keep the overall index inequality intact. We can apply the same reasoning for every γ at positive Morse-Bott tori.

We next consider negative Morse-Bott tori. We assume γ is Reeb orbit on a negative Morse-Bott torus, and α_{n-1} assigns a fixed end of multiplicity N to γ. We consider the overall inequality and show it still holds after we factor in the contributions from other terms. Let $|T'_{fixed,1,\gamma}|$ denote the number of free trivial cylinders located at γ at the u^1 level. For $i = 1, .., n-2$ we consider $(C_{i,j}, m_{i,j}) \subset u^i$ nontrivial curves that asymptote to γ as $s \to -\infty$. We let $l_{i,j}$ denote the number of such ends at each level and $n_{i,j}$ denote the multiplicity. Then the same proof as before will show the inequality continues to hold. $\square$

In fact we have equality of ECH index to Fredholm index also enforces that the cascade is simple.

We now take care of the case of height k cascades.

Proposition 2.6.31. *Consider a sequence of J_{δ_n}-holomorphic ECH index one curves u_n of bounded energy from α_1 to α_n (as nondegenerate ECH generators) converging to a cascade $u^{\natural}$ from α_1 and α_n viewed as Morse-Bott ECH generators, then $u^{\natural}$ has height one.*

Proof. Suppose $u^{\natural}$ is a height k cascade, then it can be written as k height 1 cascades, which we write as $v_1^{\natural}, ..., v_k^{\natural}$. We recall that between cascades $v_i^{\natural}$ and $v_{i+1}^{\natural}$ their end asymptotics are connected by either infinite or semi-infinte gradient flows. We pass each to a cascade of currents, and to each cascade $\mathbf{v}_i^{\natural}$ we assign to it two generalized ECH generators at its topmost and bottom-most level, which we write as α_i and α'_{i+1}. For α_i we assign all the ends approaching the minimum of f as fixed, and all others are free. For α'_{i+1} we consider all ends approaching the maximum of f are fixed, and the rest are free. The exception to this rule is α_1 and α'_{k+1} which we assign Morse-Bott ECH generators corresponding to the degenerating J_δ-holomorphic curve. With this we can assign an ECH index to each cascade $I(v_i^{\natural})$. We can also assign a relative ECH index between the general ECH generators α_i and α'_i, which we write as $I(\alpha'_i, \alpha_i)$. This number is always ≥ 0, and we illustrate it as follows. Let $\mathcal{T}$ be a Morse-Bott torus, and suppose coming from α_i there is multiplicity n_1 at the minimum of f and n_2 away from minimum of f. From α'_i there is n'_1 multiplicity at the maximum of f, and n'_2 away from the maximum of f. Then we have the inequalities

$$n'_1 \geq n_2$$

and

$$n'_2 \leq n_1.$$

Then we say contribution to $I(\alpha'_i, \alpha_i)$ from this Morse-Bott torus is $(n_1 - n'_2) = n'_1 - n_2 \geq 0$. Then we add up this term for each Morse-Bott torus that appears in α_i. Geometrically this is the total mulitplicity of complete gradient trajectories flowing between $v_i^{\natural}$ and $v_{i-1}^{\natural}$ and has potentially nonzero contributions to the ECH index. Then the fact that the cascade came from a ECH index one curve implies

$$I(v_1^{\natural}) + I(\alpha'_2, \alpha_2) + ... + I(v_k^{\natural}) = 1$$

And by previous proposition each $I(v_i) \geq 0$,with equality only if it consisted entirely of fixed trivial cylinders. Hence there is a unique $v_i^{\natural}$ with ECH index 1, the rest have ECH index zero, and all $I(\alpha_i', \alpha_i) = 0$. This means there can only be fixed trivial cylinders above and below $v_i^{\lightning}$ and cannot be infinite gradient flows. This is equivalent to saying the cascade of currents is height one. □

The above gives a description of what ECH index one cascades look like from the perspective of currents, we now reverse the process, and use the above to understand all cascades of curves of ECH index one. We need to add back in the information that was lost from passing from curves to currents. We only care about the cascades of curves that resulted from degeneration of a nondegenerate connected ECH index one curve. Call this curve C_δ. We observe the Fredholm index of C_δ, which we denote by Fred $\operatorname{Ind}(C_\delta)$, is equal to one. We assume as $\delta \to 0$, C_δ degenerates into a cascade of curves $u^{\lightning}$, and denote $\mathbf{u}^{\lightning}$ the resulting cascade of holomorphic currents. From the above we know $\mathbf{u}^{\lightning}$ is a cascade of currents of height one, however $u^{\lightning}$ could apriori be of arbitrary height, and the levels that are removed from $u^{\lightning}$to form $\mathbf{u}^{\lightning}$ must all be branched covers of trivial cylinders occurring at critical points of f.

The first case we need to consider is if $\mathbf{u}^{\lightning}$ is empty, then this implies that $u^{\lightning}$ consists purely of branched covers of trivial cylinders. To be precise $u^{\lightning}$ may contain many levels that consists of branched covers of trivial cylinders, and levels that begin and end on critical point of f, however it may also contain levels where the trivial cylinders (branched covered or not) are away from critical points of f. Here we allow levels where there is only a single unbranched cylinder away from critical points of f. We assume C_δ is connected. If at level i a trivial cylinder is at the critical point of f corresponding to elliptic Reeb orbit (hyperbolic for negative Morse-Bott torus) then all levels above i the trivial cylinders that connected to the original cylinder will be at the same Reeb orbit. Similarly if at level i a trivial cylinder is at the hyperbolic orbit (resp elliptic orbit for negative Morse-Bott torus) then all the trivial cylinders below this level connecting to this original (potentially branched cover of) cylinder will also be at the same Reeb orbits.

If all the levels of $u^{\lightning}$ are at the same Reeb orbit which is also a critical point, then u came from a branched cover of trivial cylinder in the nondegenerate case. If this is not the case, then remove the top most and bottom most levels until none of the trivial cylinders in $u^{\lightning}$ begin/end on critical point of f. Then as currents we don't care where the branched points are, so we can think of u' as a cascade of currents with only 1 level. Then the ECH index of $\mathbf{u}^{\lightning}$ is equal to one, which implies $\mathbf{u}^{\lightning}$ consists of a free trivial cylinder with multiplicity one. Hence the same must be true of $u^{\lightning}$ and there are no top/bottom branch covers.

We now turn our attention to the case where $\mathbf{u}^{\lightning}$ is nonempty. We shall use the fact the Fredholm index of C_δ is one to rule out configurations of height > 1. We observe the trivial cylinders on levels above/below $\mathbf{u}^{\lightning}$ admit the following description:

Proposition 2.6.32. *a. Let $\mathcal{T}$ denote a positive Morse Bott torus contained in the top level of $\mathbf{u}^{\lightning}$. For curves on the top level of $\mathbf{u}^{\lightning}$, as $s \to +\infty$ all free ends have multiplicity one, and avoid critical point of f. The fixed end can only have multiplicity one. Hence all branched covers of trivial cylinders above this level can only happen at* *the critical point of f corresponding to the elliptic orbit. Moreover, because C_δ obeys partition conditions, the top most level in $u^{\lightning}$ of the stack of branched trivial cylinders has partition conditions $(1, .., 1)$.*

b. *Let $\mathcal{T}$ denote a negative Morse Bott torus contained in the top level of $\mathbf{u}^{\lightning}$, as $s \to +\infty$. The positive free end of the top level of $\mathbf{u}^{\lightning}$ has multiplicity 1, so there cannot be branched cover of trivial cylinder at the critical point of f corresponding to the hyperbolic orbit. The fixed end at the critical point of f corresponding to the elliptic orbit can have a stack of branched cover of trivial cylinders on top of it on height levels above u', and again by partition conditions on C_δ the top most level is hit by partition condition (n).*

c. *Let $\mathcal{T}$ denote a positive Morse Bott torus contained in the bottom level of $\mathbf{u}^{\lightning}$. The free end has multiplicity one, so there cannot be branched covers of trivial cylinders at the critical point of f corresponding to the hyperbolic orbit. The fixed end at critical point of f corresponding to elliptic end can have a stack of branched cover of trivial cylinders below it on height levels below u', and again by partition conditions on C_δ the top most level is hit by partition condition (n).*

d. *Let $\mathcal{T}$ denote a negative Morse Bott torus contained in the bottom level of $\mathbf{u}^{\lightning}$. As $s \to -\infty$ all free ends have multiplicity one, and avoid the critical points of f. The fixed end can only have multiplicity one. Hence all branched covers of trivial cylinders above this level can only happen at the critical point of f corresponding to the elliptic orbit. Moreover, because C_δ obeys partition conditions, the bottom most level (in terms of height) of the stack of branched trivial cylinders has partition conditions $(1, .., 1)$.*

In light of the above, we can compute the topological Fredholm index of C_δ via the following procedure:

First consider the height level corresponding to $\mathbf{u}^{\lightning}$, we know all trivial cylinders connecting between nontrivial curves are simply covered, so all the possible branched covers that appear on this height level are chains of trivial branched covers of cylinders that connect to the top and bottom levels of $\mathbf{u}^{\lightning}$. We then create two additional height levels, one above $\mathbf{u}^{\lightning}$, denoted by $\overline{\mathbf{u}^{\lightning}}$ and one below $\mathbf{u}^{\lightning}$, denoted by $\underline{\mathbf{u}^{\lightning}}$, and push all branch points of trivial cylinders that appear in $\mathbf{u}^{\lightning}$ onto these 2 levels $\overline{\mathbf{u}^{\lightning}}$,$\underline{\mathbf{u}^{\lightning}}$, so that all trivial cylinders that appear in $\mathbf{u}^{\lightning}$ have no branch point (though they may be multiply covered), and hence are transversely cut out. We recall we assign $Ind(\mathbf{u}^{\lightning})$ as the dimension of moduli space of $\mathbf{u}^{\lightning}$ lives in, viewed as a cascade of currents

Then the Fredholm index of C_δ is computed as:

$$\begin{aligned} &Ind(C_\delta) = \\ &Ind(\mathbf{u}^{\lightning}) + 1 - \chi(\overline{\mathbf{u}^{\lightning}}) - \chi(\underline{\mathbf{u}^{\lightning}}) \end{aligned}$$

Note by the ECH index assumption $Ind(\mathbf{u}^{\lightning}) = 0$, so it will enforce no branched cover of trivial cylinders appear. Hence we have the proved the following proposition:

Proposition 2.6.33. *Suppose J is chosen to be good, if C_δ is a sequence of connected nontrivial ECH index one curves of bounded energy that converges to a cascade of curves, $u^{\lightning}$, then either*

- *$u^{\lightning}$ is a free cylinder of multiplicity one*
- *$u^{\lightning}$ is the same as a height one cascade of currents of ECH index one, described above, and all trivial cylinders that appear in levels of $u^{\lightning}$ either unbranched chains of fixed trivial cylinders, or trivial cylinders over a Reeb orbit of multiplicity one.*

In the latter case, $u^{\lightning}$ does not contain a sequence of fixed trivial cylinders that do not connect to any nontrivial J holomorphic curve. See Convention 2.3.8.

We call cascades of curves of ECH index one of the form stated in the above theorem *good cascades of ECH index 1.*

Then this is more or less a complete characterization of ECH index one cascades we should count in the Morse-Bott case provided we can achieve enough transversality. Assuming transversality conditions, we quote a theorem from Chapter 3 [67] to show ECH index one cascades can be glued uniquely (up to translation) to ECH index one curves.

Theorem 2.6.34 (3.4.5 Chapter 3[67])**.** *Assuming transversality conditions 2.5.6, any given ECH index one cascades can be glued uniquely to ECH index one J_δ-holomorphic curves for sufficiently small values of $\delta > 0$ up to translation in the symplectization direction.*

The key is to note ECH index one and transversality implies all of the cascades above are transverse and rigid, as in Definition 3.4.4 of Chapter 3[67] and hence can be glued. The final ingredient we need is to show that assuming J is good, the set of good ECH index one cascades is finite. To do this we need the notion of J_0 index for cascades.

2.7 Finiteness

In order to prove the differential in Morse-Bott ECH is well defined we need to prove the for given generators α, β the set of good ECH index one cascades from α to β is finite. For J-chosen to be good, we already know this set is a zero dimensional space, hence it suffices to prove that it is compact. To this end we develop the analogue of J_0 index in the Morse-Bott world. We start with 1-level cascades then build upwards to n level cascades. In this section we assume J is good throughout.

Level 1 cascades

Consider an level 1 cascade of ECH index 1 from generator α to β. In anticipation of multiple level ECH index 1 cascades, here we relax some (but not all) of the conditions on α, β to remove conditions that require certain free/fixed ends (depending on whether we are on a positive/negative Morse-Bott torus) to only have multiplicity 1. This corresponds to relaxing the condition in the nondegenerate case to only allow hyperbolic orbits of multiplicity one (see Theorem 2.6.16). We recall the consequences of generic J:

a. For positive Morse-Bott tori, as $s \to \infty$, all free ends are disjoint and are asymptotic to Reeb orbits in the torus with multiplicity 1. Let $n_+^{pos,free}$ denote the number of such orbits.

b. For positive Morse-Bott tori, the fixed ends at $s \to \infty$ are disjoint from the free ends. They are hit with partition condition (1). Suppose there are $N_+^{pos,fix}$ such ends.

c. For positive Morse-Bott tori, as $s \to -\infty$ all free ends are disjoint and cover the Reeb orbits in the torus with multiplicity 1. Let $n_-^{pos,free}$ denote the number of such orbits.

d. For positive Morse-Bott tori, as $s \to -\infty$, all fixed ends have partition conditions (n). Suppose there are $N_-^{pos,fix}$ such ends, each with multiplicity $n_{-,j}^{pos,fix}$

e. For negative Morse-Bott tori, as $s \to \infty$, all free ends are disjoint and cover the Reeb orbits in the torus with multiplicity 1. Let $n_+^{neg,free}$ denote the number of such orbits.

f. For negative Morse-Bott tori, the fixed ends at $s \to \infty$ are disjoint from the free ends. They are hit with partition conditions (n). Suppose there are $N_+^{neg,fix}$ such ends with multiplicity $n_{+,j}^{neg,fix}$

g. For negative Morse-Bott tori, as $s \to -\infty$ all free ends are disjoint and cover the Reeb orbits in the torus with multiplicity 1. Let $N_-^{neg,free}$ denote the number of such orbits.

h. For negative Morse-Bott tori, as $s \to -\infty$ there is only 1 fixed end for each Morse-Bott tori, and has partition conditions (1). Let there be $N_-^{neg,fix}$ such ends total

Definition 2.7.1. *For a level 1 good ECH index 1 cascade C connecting generator α to β, we define:*

$$J_0(C,\alpha,\beta) := -c_\tau(C) + Q_\tau(C,C) - [\sum_j (n_{-,j}^{pos,fix} - 1)] - [\sum_j (n_{+,j}^{neg,fix} - 1)] \qquad (2.25)$$

We observe that $J_0(C,\alpha,\beta)$ can be computed from the knowledge of α,β and the relative homology class of C alone. We also remark that the J_0 index can be similarly be defined for nontrivial curves of higher ECH index, as long as they satisfy the long list of partition conditions we listed above, and the same genus bounds below holds. We shall have need for this fact for the proof of finiteness below.

Then we have the following genus bound:

Proposition 2.7.2. *Let g denote the genus of a holomorphic curve C. Then we have the upper bound*

$$-\chi(C) \leq J_0(C,\alpha,\beta). \qquad (2.26)$$

Proof. We recall the adjunction formula in our case says

$$c_\tau(C) = \chi(C) + Q_\tau(C) + w_\tau(C) - 2\delta(C)$$

plugging this into J_0 yields

$$J_0(C,\alpha,\beta) = -\chi(C) - w_\tau(C) - [\sum(n_-^{pos,fix}) - 1] - [\sum n_+^{neg,fix} - 1] + 2\delta(C)$$

hence it suffices to prove

$$-w_\tau - [\sum(n_-^{pos,fix}) - 1] - [\sum n_+^{neg,fix} - 1] \geq 0.$$

We break this into cases. If C is a trivial cylinder, then this is trivial. If C has a nontrivial component along with fixed trivial cylinders, we only consider the nontrivial component, also denoted by C. All of the computations below follow from the computations of the writhe bound:

- At a positive Morse-Bott torus
  - $s \to \infty$, free end. $-w_\tau \geq 0$ because the multiplicity is one.
  - $s \to \infty$, fixed end $-w_\tau \geq 0$ because multiplicity is one.
  - $s \to -\infty$, free end. $w_\tau \geq 0$ by multiplicity.
  - $s \to -\infty$, for given fixed end j, the writhe at this end satisfies $w_\tau \geq n_-^{pos,fix} - 1$.
- At a negative Morse-Bott torus
  - $s \to \infty$, free end. $-w_\tau \geq 0$ due to multiplicity constraints.
  - $s \to \infty$, for a single fixed end j, the writhe satisfies $-w_\tau \geq n_+^{neg,fix} - 1$.
  - $s \to -\infty$, free end. $w_\tau \geq 0$ due to multiplicity constraints.
  - $s \to -\infty$, fixed end. $w_\tau \geq 0$ by multiplicity.

combining all of the above we conclude our inequality.

□

Multiple level cascades

We now explain how to generalize the definition of $J_0(C,\alpha,\beta)$ to good ECH index one cascades of arbitrary number of levels. Consider a n level cascade $u^{\lightning} = \{u^1, .., u^n\}$ of ECH index one with input α and output β. Recall we have so called fixed chains of trivial cylinders, i.e. chain of trivial cylinders that all begin/end on a fixed end orbit of either α or β until this chain of trivial cylinders meet an nontrivial holomorphic curve in one of the intermediate levels (which has an fixed end at said Reeb orbit). We remove all of these kinds of trivial cylinders, then the number J_0 is defined for each of the intermediate cascade levels, which we denote by $J_0(u^i)$, then we define the J_0 of the entire cascade as

Definition 2.7.3.

$$J_0(u^{\dagger}) := \sum J_0(u^i) \tag{2.27}$$

We observe this definition also only dependents on the relative homology class and α, β. Recall the Euler characterisitc of the cascade $\chi(u^{\dagger})$ is the Euler characterstic of the surface obtained if we glued a cylinder between each matching end of u^i and u^{i+1}, clearly then the Euler characteristic of the cascade is the sum of the Euler characteristic of each of its components. Applying the proposition for level one cascades we get

Proposition 2.7.4.

$$-\chi(u^{\dagger}) \leq J_0(u^{\dagger}).$$

Finiteness

We finally prove

Theorem 2.7.5. *Given generators α, β, the moduli space of good ECH index 1 cascades from α to β is compact.*

Proof. Let $\{u^{\dagger}_m\}$ be a sequence of good ECH index one cascades from α to β. Each $u^{\dagger}_m$ is a cascade of the form $\{u^n_m\}_n$. We show $\{u^{\dagger}_m\}$ has a convergent subsequence. From the Morse-Bott assumption there is an upper bound to how many cascade levels there are, so we pass to a subsequence where they all have N levels. For each $n = 1, .., N$, we apply the compactness for holomorphic current from [31] to each of u^m_n. To see this, note for fixed n, the energy constraint of $\{u^{\dagger}_m\}$ and Morse-Bott condition implies there are only finitely many possible choices for the positive and negative asymptotics of u^m_n, so we pick a subsequence (also denoted by u^m_n) where the positive and negative asymptotics of u^m_n is independent of m. Here, by positive and negative asymptotics of u^m_n we simply mean the Morse-Bott tori $\mathcal{T}$ that u^m_n are asymptotic to at its positive/negative ends, and the total multiplicity of Reeb orbits at each such Morse-Bott tori.

Then using the Gromov compactness for currents (see [31]) applied to $\{u^m_n\}$ we conclude we can refine a further subsequence of $\{u^m_n\}$ (for all $n = 1, .., N$) with the same relative homology class (our notion of relative homology class here is in $\mathcal{H}_2(-, -, Y)$)). Now for each u^m_n simply the knowledge of its asymptotics (which we can read off directly: by virtue of being part of ECH index one cascade all the ends that avoid the critical points of f are free, and those at critical points of f is fixed) and its relative homology class provides an upper bound on its J_0 index. This upper bound on J_0 then provides a bound the genus of each $u^m_n, n = 1, ..., N$.

With the genus bound we can apply SFT compactness: for fixed n, we observe u^m_n cannot break into a building, for that would yield (if we view $u^{\dagger}_m$ as cascade of currents) an ECH index 1 cascade of currents with $T_i = 0$, which does not exist by genericity conditions. Similarly ruled out by genericity conditions are overlapping free ends and free ends migrating to fixed ends. The u^n_m also cannot converge to a multiple cover of nontrivial curve, for that would yield an ECH cascade of current of index 1 with multiple covers of nontrivial curve, which is ruled out by genericity. Hence we conclude that $\{u_m^{\mathfrak{t}}\}$ has a subsequence that converges to a ECH index 1 cascade, and hence we have compactness. $\square$

2.8 Computing ECH in the Morse-Bott setting using cascades

We now define the Morse-Bott ECH chain complex (over $\mathbb{Z}_2$). We write the chain complex as

$$C_*^{MB}(\lambda, J) := \bigoplus_{\Theta_i} \mathbb{Z}_2\langle \Theta_i \rangle.$$

Here $\Theta_i = \{(\mathcal{T}_j, \pm, m_j)\}$ denotes a collections of Morse-Bott ECH generators. Suppose we can choose our J to be good, the differential, which we write as ∂_{MB} is defined as

$$\langle \partial_{MB}\Theta_1, \Theta_2 \rangle := \left\{ \begin{array}{l} \mathbb{Z}_2 \text{ count of J-holomorphic cascades } \mathcal{C} \text{ of ECH index } I = 1, \\ \text{so that as } s \to +\infty, \mathcal{C} \text{ approaches } \Theta_1 \text{ and as } s \to -\infty, \\ \mathcal{C} \text{ approaches } \Theta_2. \end{array} \right\} \tag{2.28}$$

We clarify that in the above definition the cascade $\mathcal{C}$ must be decomposable into $\mathcal{C}_0 \sqcup \mathcal{C}_1$, where $\mathcal{C}_0$ is a (potentially empty) collection of fixed trivial cylinders with multiplicity, and $\mathcal{C}_1$ is a good ECH index one cascade. We note if (T, n) is an element of $\mathcal{C}_0$, if it is positively asymptotic to Morse-Bott ECH generator $(\mathcal{T}, n, \pm)$, it is also negatively asymptotic to the Morse-Bott ECH generator $(\mathcal{T}, n, \pm)$ (thus far we only considered nontrivial cascades when we talked about their asymptotics).

We note by Theorem 2.7.5 the operator ∂_{MB} is well defined.

Theorem 2.8.1. *Assuming J is good, the chain complex $(C_*^{MB}, \partial_{MB})$ computes $ECH(Y, \xi)$.*

Before we prove this theorem we choose a generic family of almost complex structures J_δ.

Recall that the traditional definition of ECH requires choosing a generic J from a residual subset of almost complex structures. For fixed $\delta > 0$, we say J_δ is ECH adapted if it is an almost complex structure with which the ECH chain complex is well defined.

Definition 2.8.2. *Consider $\delta \in (0, \delta_0]$, we say a path of almost complex structures J_δ, each compatible with λ_δ for any $\delta \in (0, \delta_0]$, is generic if for any collection of Reeb orbits α, β, the moduli space*

$$\mathcal{M}(\alpha, \beta, \delta) := \left\{ (u, \delta) | \bar{\partial}_{J_\delta} u = 0, u \,\text{somewhere injective}, \lim_{s \to +\infty} u \to \alpha, \lim_{s \to -\infty} u \to \beta \right\} \tag{2.29}$$

is cut out transversely.

Theorem 2.8.3. *There is a small enough* $\delta_0 > 0$ *so that there is a generic path of almost complex structures* J_δ, $\delta \in (0, \delta_0]$ *so that:*

- J_{δ_0} *is ECH adapted.*
- $\lim_{\delta\to 0} J_\delta = J$, *where* J *is a generic almost complex structure we have chosen above to count ECH index one cascades.*
- $|J - J_\delta| \le C\delta$ *in* C^k *norm,* $k > 100$, *and* J_δ *take the prescribed form near small fixed neighborhood of Morse-Bott torus described in Section 2.4.*
- *For a residual subset* $S \subset (0, \delta_0]$, *for all* $\delta \in S$, J_δ *is ECH adapted.*

Proof. This is standard application of Sard-Smale theorem. □

Proof of theorem 2.8.1. We observe for fixed $L > 0$, there are only finitely many ECH index 1 cascades of energy $< L$. We fix δ_0 small enough so that for all $\delta \in (0, \delta_0]$ the cascades can be glued (uniquely in our sense specified) to ECH index 1 curves.

We assume δ_0 is such that J_{δ_0} is ECH adapted. We recall we have chosen a generic family $J_\delta, \delta \in [0, \delta_0]$ so that the space:

$$\{(u, J_\delta) \,|\, \delta \in (0, \delta_0] \; u \, J_\delta \, \text{holomorphic, somewhere injective ECH index 1}\}$$

modulo translation is a 1-manifold (not necessarily compact). A SFT compactness theorem ([67, 5, 6]) tells us the $\delta = 0$ ends of this manifold are precisely the good ECH index one cascades.

We recall there is a residual set $A \subset (0, \delta_0]$ so that for all $\delta \in A$, J_δ is ECH adapted and the ECH homology can be computed by counting ECH index one J_δ holomorphic curves for $\delta \in A$.

We make the following observation: if u_δ and v_δ are J_δ-holomorphic curves of ECH index one that converge to the same cascade as $\delta \to 0$, by the gluing theorem, for small enough δ u_δ and v_δ are in fact the same curve up to $\mathbb{R}$ translation.

Then we claim we can find small enough $\delta' \in A$ so that the corbordism from $\delta = 0$ to δ' built by $\{(u, J_\delta) \,|\, \delta \in (0, \delta'] \; u \, J_\delta \, \text{holomoprhic, somewhere injective ECH index 1}\}$ is the trivial cobordism. Suppose not, then for arbitrarily small δ we can find u_δ a ECH index one somewhere injective curve that does not come from gluing, take $\delta \to 0$ and after taking a subsequence, u_δ degenerates into a good ECH index one cascade, but by our observation must have come from a curve obtained by gluing together an ECH index one cascade, contradiction. □

2.9 ECH index one curves of genus zero

We showed in the previous section that when there is enough transversality for cascades, the cascades of ECH index one take a particularly nice form. However this is not always achievable, except in special circumstances. In this section and the next we outline some special circumstances in which transversality can always be achieved. Here we consider the case where all ECH index one curves in the perturbed picture must have genus zero. This is the case for T^3 and some toric domains.

We shall use a slightly different description of cascades that do not allow for the presence of trivial cylinders. We will call this description "tree-like" cascades and will be described below. The reason we can use this description is that if the curve has genus zero, we can do the gluing without requiring that between each adjacent cascade levels there is a *single* flow time parameter; instead we can assign a different flow time between each pair of adjacent nontrivial curves.

We use the following convention to represent our holomorphic curves. We use a vertex to represent a J holomorphic curve of genus zero, and use directed edges to denote the positive and negative punctures of the curve. Edges directed away from the J-holomorphic curve correspond to positive punctures, and edges directed towards the vertex correspond to negative punctures. The figure below illustrates how we go from J-holomorphic curve to vertex with directed edges.

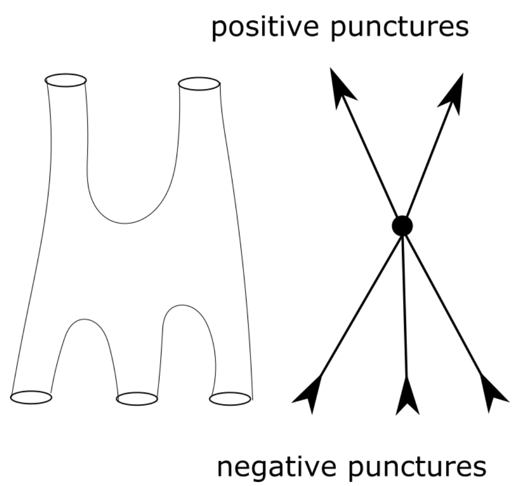

Figure 2.2: Passing from genus zero curve to vertex with edges

Then a *height one cascade with tree-like compactifications* from Morse-Bott ECH generator consists of the following data:

a. A collection of vertices $\{v_1, .., v_n\}$ each equipped with the data of inward and outward pointing edges. Each vertex has at least one outgoing edge. Each edge is also equipped with the information of which Reeb orbit it lands on.

b. Given two vertices v_i and v_j, if we can find a Morse-Bott torus $\mathcal{T}$ so that a positive puncture of v_i lands on γ, and if we follow the gradient flow for time $T_{i,j} \in [0, \infty)$ along γ we arrive at a negative puncture of v_j landing on the corresponding orbit, then we say it is possible to connect v_i and v_j via the given pair of edges. The data of a height one cascade in this compactification consists of choices of connections between the vertices of $\{v_1, .., v_n\}$, so that after we connect the edges, we obtain a connected tree. See figure below for an example. We call these connections internal connections.

c. The positive punctures of $\{v_1, .., v_n\}$ that are not assigned internal connections are assigned free/fixed as per required by ECH generator α_1, and likewise for negative punctures and α_n.

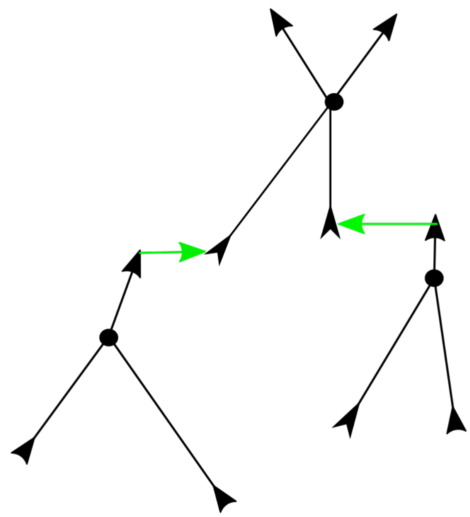

Figure 2.3: Cascade with tree like compactification. The green arrow denote finite gradient flow lines.

For genus zero J_δ-holomorphic curves degenerating into a cascade with our previous compactification, we can easily pass to a tree like compactification by removing all the trivial cylinders.

Given a cascade of height one with tree like compactification, which we write as $u^{\lightning} = \{v_1, .., v_n\}$. We can compute its ECH index as follows: we treat all edges participating in internal connections as free, then the ECH index is simply given by

$$I(u^{\lightning}) = I(v_1) + ... + I(v_n) - n + 1.$$

In order to talk about Fredholm index we also need to pass to the reduced cascade $\tilde{u}^{\lightning}$ consisting of curves $\{\tilde{v_1}, .., \tilde{v_n}\}$. If in our tree like compactification all free ends assigned by

α_1 and α_n as well as all internal connections avoided critical points of f, then the reduced cascade lives in a transversely cut out moduli space of dimension

$$\sum_i Ind(\tilde{v}_i) - 1$$

since being tree like removes the condition of needing to have the same flow time between adjacent cascade levels.

Hence to achieve the necessary transversality conditions to count ECH index one cascades, we choose a generic J so that

a. For any punctured sphere that is the domain of a J-holomorphic curve, we endow it with an assignment of incoming and outgoing punctures, and for each end we assign a free/fixed end; and if an end is assigned fixed it must land on a Reeb orbit corresponding to a critical point of f under the J-holomorphic map; and if an end is free it must avoid critical points of f. Then all moduli spaces of somewhere injective J holomorphic curves with the above information are transversely cut out with dimension given by the index formula.

b. For any two curves v_1 and v_2 satisfying the above condition and both rigid, if their free ends land on the same Morse-Bott torus from opposite sides (one as a positive puncture the other as a negative puncture), then they do not land on the same Reeb orbit in the Morse-Bott family (we only care about where they land on the Morse-Bott torus and ignore information of multiplicity, i.e. even if they cover the same Reeb orbit of different multiplicity on their free ends, this is prohibited).

The above conditions are easily achieved by choosing a generic J by classical transversality methods. We next consider cascades of height one. We observe we have the inequality (if we treat all internal connections as free for both ECH index and Fredholm index)

$$I(u^{\lightning}) - n \geq \sum Ind(\tilde{v}_i) - 1 \geq 0$$

since each $\tilde{v}_i$, by virtue of it existing and transversality conditions, must have Fredholm index ≥ 0. ECH index one implies $Ind(\tilde{v}_i) = 1$, hence all these curves are rigid, and embedded. By the above genericity of J all flow times are nonzero, and the cascade itself is already reduced. All free ends and ends coming from internal connections avoid critical points of f. Also observe that by partition conditions derived previous sections that between internal connections, the participating edges can only over Reeb orbits with multiplicity one.

Then suppose a sequence of genus zero ECH index one J_δ holomorphic curves from α_1 to α_n degenerates into a cascade with tree like compactification for arbitrary height. This just means we allow internal connections adjoint to each other with semi-infinite or infinite gradient trajectories. Then for each internal connection whose flow time is infinite, we separate them into two different cascades. Then we get a collection of height one cascades each of which is tree like. We write them as $u_1^\dagger, ..., u_k^\dagger$. Then we can assign generalized ECH generators to ends of $u_i^\dagger$ as before, and the ECH index one condition imposes

$$I(u_1^\dagger) + I(u_2^\dagger) + \cdots + I(u_k^\dagger) + \text{relative difference between ECH generators} = 1$$

By relative difference between ECH generators we mean the same construction as proposition 2.6.31. We have for all height one cascades that

$$I(u_i^\dagger) - 1 \geq Ind(\tilde{u}_i^\dagger) \geq 0$$

Hence there is either a unique cascade $u_i^\dagger$ of index zero, or the entire cascade is just one gradient flow line. By considerations of topological Fredholm index we also rule out additional branched cover of trivial cylinders at the top/bottom level of the cascade with tree-like compactifications. Hence using the above description we have the following proposition.

Proposition 2.9.1. *In the nondegenerate case, ECH index one curves of genus zero degenerate into ECH index one tree like cascades that are reduced and transversely cut out.*

We call the type of cascades of the above proposition "good ECH index one tree like cascades", because we eliminated branched covers of trivial cylinders via topological Fredholm index.

As in the previous section we choose J_δ to be a generic family of almost complex structures satisfying the same conditions as Theorem 2.8.3.

We then quote a gluing theorem from Chapter 3 [67].

Theorem 2.9.2. *Let $u^\dagger$ be a good ECH index one cascade of genus zero as per above, then for small enough $\delta > 0$ there exists a unique (up to translation) J_δ-holomorphic curve in an ϵ neighborhood of this cascade.*

Proof. The main difference is that because the whole curve is genus zero, we no longer need to make sure the pregluing is well defined by restricting our choice of asymptotic vectors to $\hat{\Delta}$, as in proposition 8.28 in [67]. □

We define a chain complex as before. We We write the chain complex as

$$C_*^{MB,tree}(\lambda, J) := \bigoplus_{\Theta_i} \mathbb{Z}_2\langle\Theta_i\rangle.$$

We use the superscript "tree" to denote the fact we are counting tree like cascades. As before $\Theta_i = \{(\mathcal{T}_j, \pm, m_j)\}$ denotes a collections of Morse-Bott ECH generators. After we choose a generic J, all good tree like cascades are transversely cut out. Then we define the differential ∂_{MB}^{Tree} to be

$$\langle\partial_{MB}^{tree}\Theta_1, \Theta_2\rangle := \left\{\begin{array}{l} \mathbb{Z}_2 \text{ count of tree like J-holomorphic cascades } \mathcal{C} \text{ of ECH index } I = 1, \\ \text{so that as } s \to +\infty, \mathcal{C} \text{ approaches } \Theta_1 \text{ and as } s \to -\infty, \\ \mathcal{C} \text{ approaches } \Theta_2. \end{array}\right\} \tag{2.30}$$

As before, we clarify in the cascade $\mathcal{C}$ must be decomposable into $\mathcal{C}_0 \sqcup \mathcal{C}_1$, where $\mathcal{C}_0$ is a (potentially empty) collection of fixed trivial cylinders with multiplicity, and $\mathcal{C}_1$ is a good ECH index one tree like cascade.

Theorem 2.9.3. *Suppose J is chosen to be generic so that all ECH index one good tree like cascades are transversely cut out, and we can choose a generic family of perturbations to J, which we write as J_δ that meets the conditions of Theorem 2.8.3. We further for small enough $\delta > 0$, all J_δ-holomorphic curves of ECH index one are genus zero. Then the chain complex $(C_*^{MB,tree}, \partial_{MB}^{Tree})$ computes $ECH(Y,\xi)$.*

Proof of Theorem 2.9.3. The same proof as in Theorem 2.8.1 works. □

2.10 Applications to concave toric domains

As an application of our methods we show that for concave toric domains, ECH can be computed via enumeration of ECH index one cascades. By what we proved above, it suffices to show all ECH index one holomorphic curves after the Morse-Bott perturbation have genus zero.

We recall the definition of a concave toric domain. Consider $\mathbb{C}^2$ equipped with the standard symplectic product symplectic form. Consider the diagonal S^1 action on $\mathbb{C}^2$, and the associated moment map $\mu : \mathbb{C}^2 \to \mathbb{R}^2$ given by

$$\mu(z_1, z_2) = (\pi|z_1|^2, \pi|z_2|^2).$$

Let $\Omega \subset \mathbb{R}^2$ be a domain in the first quadrant of $\mathbb{R}^2$, we define the domain X_Ω to be

$$X_\Omega := \{(z_1, z_2)|\mu(z_1, z_2) \in \Omega\}.$$

Suppose Ω is a domain bounded by the horizontal segment from $(0, 0)$ to $(a, 0)$, the vertical segment from $(0, 0)$ to $(0, b)$ and the graph of a convex function $f : [0, a] \to [0, b]$ so that $f(0) = b$ and $f(a) = 0$. We further assume f is smooth, $f'(0)$ and $f'(a)$ are irrational, $f'(x)$ is constant near 0 and a, and $f''(x) > 0$ whenever $f'(x)$ is rational, then we say X_Ω is a **concave toric domain**. Note our definition is slightly more restrictive than that of [10], because we are not interested in capacities; we need the boundary of X_Ω to be well behaved enough to define ECH.

For a concave toric domain X_Ω, its boundary ∂X_Ω is a contact 3-manifold diffeomorphic to S^3. We now describe the Reeb orbits that appear in ∂X_Ω. We also note their Conley Zehnder indices, having chosen the same trivializations as in [10].

a. $\gamma_1 = \{(z_1, 0) \in \partial X_\Omega\}$. The orbit γ_1 is elliptic with rotation angle $-1/f'(a)$, hence $CZ(\gamma_1^k) = 2\lfloor -k/f'(a) \rfloor + 1$

b. $\gamma_2 = \{(0, z_2) \in \partial X_\Omega\}$. The orbit γ_2 has rotation angle $-f'(0)$, hence $CZ(\gamma_2^k) = 2\lfloor -kf'(0) \rfloor + 1$.

c. Let $x \in (0, a)$ be such that $f'(x)$ is rational. Then the torus described by $\{(z_1, z_2) | \mu(z_1, z_2) = (x, f(x))\}$ is a (negative) Morse-Bott torus. Each Reeb orbit has Robbin-Salamon index $-1/2$.

We say a bit more about the Reeb dynamics for the third case. Consider the point $(x, f(x))$ so that $f'(x)$ is rational. We set $f'(x) = tan(\phi), \phi \in (-\pi/2, 0)$. Then the Reeb vector field is given by (see [50])

$$R = \frac{2\pi}{-x\sin(\phi) + f(x)\cos(\phi)}(-\sin\phi\partial_{\theta_1} + \cos(\phi)\partial_{\theta_2}).$$

For large action $L > 0$, we perturb each Morse-Bott torus to a pair of orbits, one elliptic, the other hyperbolic. Then an ECH generator $\alpha = \{\alpha_i, m_i\}$ is a collection of nondegenerate Reeb orbits with multiplicities. We associate to each ECH generator a **combinatorial generator**.

Definition 2.10.1. *(see [10]) A combinatorial generator is a quadruple $\tilde{\Lambda} = (\Lambda, \rho, m, n)$ where*

a. Λ is a concave integral path from $(0, B)$ to $(A, 0)$ such that the slope of each edge is in the interval $[f'(0), f'(a)]$.

b. ρ is a labeling of each edge of Λ by e or h.

c. m and n are nonnegative integers.

Let $\Lambda_{m,n}$ denote the concatenation of the following sequence of paths:

a. The highest polygonal path with vertices at lattice points from $(0, B + n + \lfloor -mf'(0) \rfloor)$ to $(m, B + n)$ which is below the line through $(m, B + n)$ with slope $f'(0)$.

b. The image of Λ under the translation $(x, y) \mapsto (x + m, y + n)$.

c. The highest polygonal path with vertices at lattice points from $(A + m, n)$ to $(A + m + \lfloor -n/f'(a) \rfloor, 0)$ which is below the line through $(A + m, n)$ with slope $f'(a)$.

Let $\mathcal{L}(\Lambda_{m,n})$ denote the number of lattice points bounded by the axes and $\Lambda_{m,n}$, not including the lattice points on the image of Λ under the translation $(x, y) \mapsto (x + m, y + n)$. We then define

$$I^{comb}(\Lambda_{m,n}) = 2\mathcal{L}(\Lambda_{m,n}) + h(\Lambda)$$

where $h(\Lambda)$ is the number of edges in Λ labelled by h. To each ECH generator $\alpha = \{(\alpha_i, m_i)\}$ we associate a combinatorial ECH generator (Λ, m, n) as follows. The number m is the multiplicity of γ_2 as it appears in α, and the integer n is the multiplicity of γ_1 as it appears in α. For other (nondegenerate) Reeb orbits of α, they all come from small perturbations of Morse-Bott tori. If $\gamma \in \alpha$ is a Reeb orbit that comes from breaking the degeneracy of a Morse-Bott torus at $(x, f(x))$, then let v_1 be the smallest positive integer so that $v_2 = f'(x)v_1 \in \mathbb{Z}$.

Let v denote the vector $v = (v_1, v_2)$. The path is obtained by taking each Reeb orbit γ in α that come from Morse-Bott tori, associating to it the vector that is v multiplied by the multiplicity of γ as it appears in α, and concatenating these vectors in order of increasing slope. The labelling ρ is obtained by labelling the vector associated to γ the letter h if γ is hyperbolic, and e if γ is elliptic.

Proposition 2.10.2. *([10]) If C is a current from α to β, its ECH index is given by $I^{comb}(\alpha) - I^{comb}(\beta)$.*

For future usage, we also record how the Chern class is computed (see [10]). Let α denote a ECH generator, we associate to it the combinatorial generator (Λ, ρ, m, n), then we take

$$c_\tau(\alpha) = A + B + m + n.$$

Then if we have a J-holomorphic curves from ECH generator α to β, then its relative first Chern class is calculated by $c_\tau(\alpha) - c_\tau(\beta)$.

We need a version of the local energy inequality, which we take up presently. Versions of this inequality have appeared in [38, 68, 14, 9]. Consider the boundary of Ω with its intersections with the two coordinate axes removed, then its preimage under the moment map is an interval times a two torus. We write the two torus as $(x_1, x_2) \in S^1_1 \times S^1_2$, where the first S^1_1 is the S^1 coming from rotation in the first complex plane $\mathbb{C}$, and the second S^1 comes from the second copy of $\mathbb{C}$. We use $\mathbb{Z} \oplus \mathbb{Z}$ to denote the lattice of first homology with $\mathbb{Z}$ coefficients. Consider a Morse-Bott torus at $(x, f(x))$ with $f'(x) = v_2/v_1$ as before, then the homology class of the Reeb orbit is given by the pair $(-v_2, v_1) \in \mathbb{Z}^2$ (this is true before or after the Morse-Bott perturbation).

Consider $F_{[x_0,x_1]}$, by which we denote the preimage of the graph $\{(x, f(x)) | x \in [x_0, x_1]\}$ under the moment map. We similarly consider F_x, which is the preimage of $(x, f(x))$ under the moment map. Let C be a somewhere injective J holomorphic curve, we consider $C \cap F_{x_0}$ (we choose x_0 generically so this intersection is transverse). We orient this intersection using the boundary orientation of $C \cap F_{[x_0-\epsilon,x_0]}$. Its homology class in $\mathbb{Z}^2$ we write as $[F_x]$.

Proposition 2.10.3. *Let $(p, q) \in \mathbb{Z}^2$ denote the homology of $C \cap F_{x_0}$, then we have the inequality*

$$p + f'(x)q \geq 0.$$

We further observe equality holds only if C is a trivial cylinder.

Proof. We consider $C \cap F_{[x_1,x_2]}$, and observe with our conventions $\partial(C \cap F_{[x_0,x_1]}) = C \cap F_{x_1} - C \cap F_{x_0}$. We next consider

$$\begin{aligned}
\int_{C\cap F_{[x_1,x_2]}} d\lambda &= \int_{C\cap F_{x_1}} \lambda - \int_{C\cap F_{x_0}} \lambda \\
&= \int_{C\cap F_{x_1}} r_1 d\theta_1 + r_2 d\theta_2 - \int_{C\cap F_{x_0}} r_1 d\theta_1 + r_2 d\theta_2 \\
&= (x_1 - x_0)p + (f(x_1) - f(x_0))q \geq 0.
\end{aligned}$$

By taking the limit $x_0 \to x_1$, we conclude the proof. □

Suppose the J-holomorphic C current connects from α_+ to α_- and has ECH index one. Suppose C does not contain trivial cylinder components, hence it is embedded. Let α_+ contain γ_1 with multiplicity n_+, the orbit γ_2 with multiplicity m_+, and contains e_+ distinct elliptic orbits and h_+ hyperbolic orbits. Suppose further C has k_m^+ ends at γ_2, with multiplicities m_+^i, and C has k_n^+ ends at γ_1 with multiplicities n_+^i Likewise we use m_-, n_-, e_-, h_- and $k_m^-, m_-^i, k_n^-, n_-^i$ to denote the respective quantities in α_-, except here e_- denotes the number of elliptic Reeb orbits counted with multiplicity. Then the key is the following proposition (similar proofs have appeared in [14, 38, 9])

Proposition 2.10.4. *For the case of concave toric domains, after a small perturbation away from the Morse-Bott degeneracies, all ECH index one curves have genus zero.*

Proof. **Step 1** We know that the integers $m_\pm^i$ and $n_\pm^i$ satisfy partition conditions because C has ECH index one. Recall that for an elliptic Reeb orbit of rotational angle θ, suppose C is asymptotic to this Reeb orbit at its positive ends with multiplicity m. Consider the line $y = \theta x$ on the $x - y$ plane, then draw the maximal concave polygonal path connecting lattice points beneath $y = \theta x$. This polygonal path $\mathcal{P}$ starts at the origin and connects to $(m, \lfloor m\theta \rfloor)$. The horizontal displacements of the edges in this path we will write as (m_i) and take the convention that if $i < j$, then m_i is the segment before m_j if we count starting from the origin. This gives an integer partition of m, which is the partition conditions for positive ends of C that are asymptotic to this Reeb orbit.

We observe that $\sum_i \lfloor m_i\theta \rfloor = \lfloor \theta m \rfloor$. To see this, first it follows from the properties of the floor function that

$$\sum_i \lfloor m_i\theta \rfloor \le \lfloor m\theta \rfloor.$$

For the converse inequality, consider the polygonal path $\mathcal{P}$ with vertices at $(\sum_i^k m_i, \lfloor \sum_i^k m_i\theta \rfloor)$. It suffices to show

$$\lfloor m_k\theta \rfloor \ge \lfloor \sum_i^k m_i\theta \rfloor - \lfloor \sum_i^{k-1} m_i\theta \rfloor.$$

This follows from the fact that

$$\theta \ge \frac{\lfloor \sum_i^k m_i\theta \rfloor - \lfloor \sum_i^{k-1} m_i\theta \rfloor}{m_k}$$

which is a consequence of the fact that $\mathcal{P}$ is maximally concave.

We next recall the partition conditions for negative ends of C asymptotic to the Reeb orbit with rotation angle θ. Consider the line $y = \theta x$, and the minimal convex path above $y = \theta x$ that connects between $(0, 0)$ and $(m, \lceil m\theta \rceil$ through lattice points. The horizontal displacements of the edges of of this path are labelled (in order) m_i, and form the partition conditions for ends of C. Using a very similar proof as before, we can show

$$\sum \lceil m_i\theta \rceil = \lceil m\theta \rceil.$$

Then we can compute the Fredholm index of C as

$$\begin{aligned}Ind(C) =&2g-2+(e_+ + h_+ + k_m^+ + k_n^+) + (e_- + h_- + k_m^- + k_n^-)\\&+2(A_+ + B_+ + m_+ + n_+ - A_- - B_- - m_- - n_-)\\&-e_+ + e_-\\&+(k_n^+ + k_m^+ + k_m^- + k_n^-)\\&+\sum_{i=1}^{k_n^+} 2\lfloor -n_+^i/f'(a)\rfloor + \sum_{i=1}^{k_m^+} 2\lfloor -m_+^i f'(0)\rfloor - \sum_{i=1}^{k_n^-} 2\lceil -n_-^i/f'(a)\rceil - \sum_{i=1}^{k_m^-} 2\lceil -m_-^i f'(0)\rceil.\end{aligned}$$

Step 2 To analyze the above equation further, we first note that

$$A_+ + n_+ + \sum_{i=1}^{k_m^+} \lfloor -m_+^i f'(0)\rfloor - A_- - n_- - \sum_{i=1}^{k_m^-} \lceil -m_-^i f'(0)\rceil \geq 0 \tag{2.31}$$

This is accomplished by considering the interior intersections of C with $\gamma_2 \times \mathbb{R}$. All such intersection points are positive, by positivity of intersections. The count of interior intersections is given by (see [32])

$$l_+(C, \gamma_2) - l_-(C, \gamma_2)$$

where l_+ denotes the linking number of positive ends of C with γ_2, and l_- is the linking of negative ends of C with γ_2. We note the linking numbers in a concave toric domain are calculated as follows ([10]):

$$lk(\gamma_1, \gamma_2) = 1, \quad lk(\gamma_1, o_v) = -v_2, \quad lk(\gamma_2, o_v) = v_1, \quad lk(o_v, o_w) = \min\{-v_1 w_2, -v_2 w_1\}.$$

Here we use o_v to denote nondegenerate orbits that come from perturbing a Morse-Bott torus at $(x, f(x))$, with $f'(x) = v_2/v_1$.

From this we see that $lk_+ = A_+ + n_+ + \sum_{i=1}^{k_m^+} \lfloor -m_+^i f'(0)\rfloor$, and $lk_- = A_- + n_- + \sum_{i=1}^{k_m^-} \lceil -m_-^i f'(0)\rceil$. The $A_\pm$ terms come from ends of C asymptotic to o_v, the $n_\pm$ term comes from ends of C asymptotic to γ_1, and the floor and ceiling terms come from ends of C asymptotic to γ_2 and the fact that C has ECH index one. From the partition conditions we see that $\sum_{i=1}^{k_m^+} \lfloor -m_+^i f'(0)\rfloor = \lfloor -m_+ f'(0)\rfloor$. Likewise we can show

$$B_+ + m_+ + \sum_{i=1}^{k_n^+} 2\lfloor -n_+^i/f'(a)\rfloor - B_- - m_- - \sum_{i=1}^{k_n^-} 2\lfloor -n_-^i/f'(a)\rfloor \geq 0$$

Hence we conclude from the Fredholm index formula that if C has ends at γ_+ or γ_-, then it must have genus 0.

Step 3 Next we consider the case where C has no ends at γ_+ or γ_-. We assume C has genus one. Then $A_+ = A_-, B_+ = B_-$ from Fredholm index considerations. Let $\Lambda_\pm$ denote the polygonal paths associated to generators $\alpha_\pm$. We first show Λ_+ lies outside Λ_-. By the above we already know they agree at end points.

As a preamble, we consider the homology classes $F_x \cap C$. First for x very close to zero, say equal to $\epsilon > 0$, let $[F_\epsilon] = (p, q)$. Then we have $p + f'(0)q \geq 0$. Similarly consider $[F_{1-\epsilon}] = (-p, -q)$. We have $-p - f'(a)q \geq 0$. Adding these inequalities to get $(f'(0) - f'(a))q \geq 0$ from which we deduce $q \leq 0$. Then we have $-f'(a)q \geq p \geq -f'(0)q$, which implies $p = q = 0$. Incidentally this implies a kind of maximal principle for holomorphic curves. Note $p + f'(x)q = 0$ only if the curve is a branched cover of a trivial cylinder. This implies for our curves they are confined to have $x \in (0, 1)$.

Next we compute $[F_x]$ for any x irrational and $\epsilon > 0$ sufficiently small. We have

$$\begin{aligned}&[F_x] - [F_\epsilon] + \text{homology class of Reeb orbits in } [\epsilon, x] \text{ approached by positive ends of } C\\ &- \text{homology class of Reeb orbits in } [\epsilon, x] \text{ approached by negative ends of } C = 0.\end{aligned}$$

Next we consider the no crossing of polygonal paths.

Suppose the no crossing result does not hold, since we know $\Lambda_\pm$ have the same beginning and end points, there must exists two intersection points which we call (a, b) and (c, d), with $a < c$. Then on the interval (a, c) the path Λ_- is strictly above Λ_+ except at end points where they overlap. Form the line connecting (a, c) and (b, d), we can find $x_0 \in (a, c)$ such that $f'(x_0) = \frac{d-b}{c-a}$. We compute $[F_{x_0-\epsilon}]$ and apply the local energy inequality to it. We use $x_0 - \epsilon$ to avoid the case where x_0 is the x coordinate of lattice points in $\Lambda_\pm$, practically this will not make a difference.

Let the lattice point (p, q) have the following property: it is a vertex on Λ_+, the edge to the left of this lattice point has slope less than $f'(x_0)$, and the edge to the right of this vertex has slope greater than equal to $f'(x_0)$. Then the contribution to $[F_{x_0-\epsilon}]$ from Λ_+ is simply $(-(B - q), -p)$. We also consider the contribution of $F_{x_0-\epsilon}$ from Λ_-, which takes the form $(B - q', p')$. The lattice point (p', q') on Λ_- is chosen the same way as (p, q). If no such vertex exists, then Λ_- must overlap with the line segment connecting (a, b) and (c, d). Then the point (p', q') is still the lattice point on Λ_- which corresponds to the left most end point of where Λ_- overlaps with the line connecting (a, b) to (c, d). In either case the local energy inequality says that

$$(q - q') + \frac{d-b}{c-a}(p' - p) \geq 0$$

We first assume (p', q') is not on the line connecting (a, b) and (c, d), then this means that the point (p, q) is further away from the line connecting (a, b) to (c, d) than (p', q'). Geometrically this is described by

$$(b - d)(p - p') + (c - a)(q - q') < 0.$$

which is impossible. Now assume (p', q') is on the line connecting (a, b) to (c, d), then since we have chosen $[F_{x_0-\epsilon}]$, we must have $p' < p$. The energy inequality implies

$$\frac{q - q'}{p - p'} > \frac{d-b}{c-a}$$

contradicting the geometric picture.

Step 4. After we proved no-crossing in the previous step, we show there cannot be a genus one curve satisfying the assumptions of the previous step. The Fredholm index formula tells us that (recall we are assuming $g = 1$)

$$1 = h_+ + h_- + 2e_-$$

which means $e_- = 0$ and at most one of h_+ and h_- is one. If $h_+ = 1$, and $h_- = 0$, then $\alpha_- = \emptyset$. By inspection C cannot have ECH index one.

On the other hand, if $h_+ = 1$ and $h_- = 1$, then Λ_- consists of a single line segment. Λ_+ has the same end points as Λ_- and is concave, hence must also agree with Λ_- as polygonal paths. One checks easily that in this case the ECH index cannot be one.

This concludes the proof that all ECH index one curves have genus zero. □

After we have proved all ECH index one curves have genus zero, we can then use the tree like compatification to describe the moduli space of cascades. However there is the complication that there are two nondegenerate orbits, γ_+ and γ_-. So in the tree like compactification, we allow the ends of J-holomorphic curves to land on nondegenerate orbits. Furthermore, connecting between two nontrivial curves, instead of a gradient trajectory, it could be that adjacent ends of J-holomorphic curves land on the same non-degenerate orbits and no gradient trajectories connect between them. See figure 2.4.

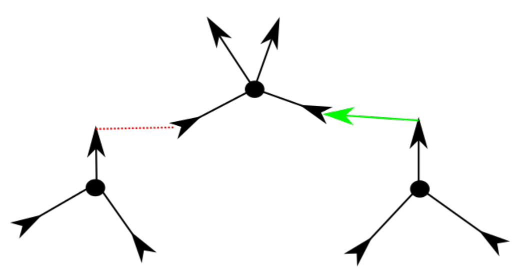

Figure 2.4: Cascade with tree like compactification for concave toric domains. The unconnected ends of holomorphic curves can land on either Morse-Bott tori or nondegenerate Reeb orbits. The green arrow denotes a finite gradient flow line connecting between two adjacent ends that land on Morse-Bott tori. The dashed line is used to indicate the adjacent ends land on non-degenerate Reeb orbits, and there is no need for gradient trajectories to connect between them.

Given such a cascade of ECH index one, we can cut it into subtrees along each matching pair of nondegenerate orbits, see figure 2.5.

The ECH index is additive with respect to concatenation of sub-trees. So the ECH index one conditions implies there are no matching along nondegenerate orbits, and we can use the correspondence theorem 2.9.3 as before.

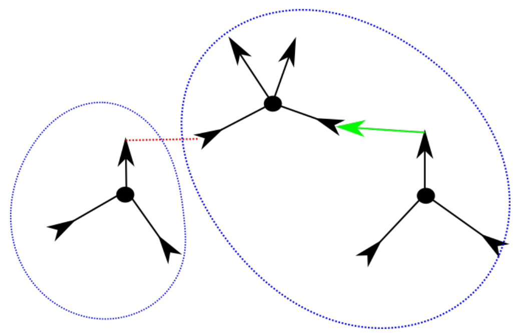

Figure 2.5: We cut along the red dashed lines to sub trees of cascades. For this figure each subtree is circled by dashed blue lines. The ECH index is additive along concatenation of such sub trees.

2.11 Convex Toric Domains

In this section we show we can compute the ECH chain complex of convex toric domains via enumeration of J-holomorphic cascades. As there are many similarities with the case of concave toric domains, we will be brief in its treatment.

Suppose Ω is a domain bounded by the horizontal segment from $(0,0)$ to $(a,0)$, the vertical segment from $(0,0)$ to $(0,b)$ and the graph of a concave function $f:[0,a]\to[0,b]$ so that $f(0)=b$ and $f(a)=0$. We further assume f is smooth, $f'(0)$ and $f'(a)$ are irrational, $f'(x)$ is constant near 0 and a, and $f''(x)<0$ whenever $f'(x)$ is rational, then we say X_Ω is a **convex toric domain**.

As in the case of a concave toric domain, the boundary of X_Ω, written as ∂X_Ω, is a contact 3-manifold diffeomorphic to S^3. We now describe the Reeb orbits that appear in ∂X_Ω. We also note their Conley Zehnder indices, having chosen the same trivializations as in [30]

a. $\gamma_1=\{(z_1,0)\in\partial X_\Omega\}$. The orbit γ_1 is elliptic with rotation angle $-1/f'(a)$, hence $CZ(\gamma_1^k)=2\lfloor -k/f'(a)\rfloor+1$

b. $\gamma_2=\{(0,z_2)\in\partial X_\Omega\}$. The orbit γ_2 has rotation angle $-f'(0)$, hence $CZ(\gamma_2^k)=2\lfloor -kf'(0)\rfloor+1$.

c. Let $x\in(0,a)$ be such that $f'(x)$ is rational. Then the torus described by $\{(z_1,z_2)|\mu(z_1,z_2)=(x,f(x))\}$ is a (positive) Morse-Bott torus. Each Reeb orbit has Robbin-Salamon index $+1/2$.

Definition 2.11.1. *A combinatorial generator is a quadruple* $\tilde{\Lambda}=(\Lambda,\rho,m,n)$ *where*

a. Λ *is a convex integral path from* $(0,B)$ *to* $(A,0)$ *such that the slope of each edge is in the interval* $[f'(0),f'(a)]$.

b. ρ is a labeling of each edge of Λ by e or h.

c. m and n are nonnegative integers.

Let $\Lambda_{m,n}$ denote the concatenation of the following sequence of paths:

a. The highest polygonal path with vertices at lattice points from $(0, B+n+\lfloor -mf'(0)\rfloor)$ to $(m, B+n)$ which is below the line through $(m, B+n)$ with slope $f'(0)$.

b. The image of Λ under the translation $(x, y) \mapsto (x+m, y+n)$.

c. The highest polygonal path with vertices at lattice points from $(A+m, n)$ to $(A+m+\lfloor -n/f'(a)\rfloor, 0)$ which is below the line through $(A+m, n)$ with slope $f'(a)$.

Let $\mathcal{L}(\Lambda_{m,n})$ denote the number of lattice points bounded by the axes and $\Lambda_{m,n}$, including the lattice points on the edges of $\Lambda_{m,n}$. We then define

$$I^{comb}(\Lambda_{m,n}) = 2(\mathcal{L}(\Lambda_{m,n}) - 1) - h(\Lambda)$$

And the Chern class of $\Lambda_{m,n}$ is given by

$$c_\tau(\Lambda_{m,n}) = A + B + m + n.$$

Theorem 2.11.2. *The ECH index of a holomorphic curve between two ECH generators is the difference of the I^{comb} we associate to their corresponding combinatorial ECH generators.*

Proof. The proof is a generalization of the computation in [30, 10]. We briefly summarize this below. Let α denote a ECH orbit set. We consider $I(\alpha, \emptyset, Z)$ where Z is the unique relative homology class that is represented by discs with boundary α. Let m, n denote the multiplicity of γ_2, γ_1 respectively in α, and let Λ be the resulting convex integral path defined by associating Reeb orbit sets to integral paths as in [30]. Then it suffices to show $I(\alpha, \emptyset, Z) = I^{comb}(\Lambda_{m,n})$. The computation is the same as the one in [30], except the Conley-Zehnder index terms arising from γ_1 and γ_2 may not just be 1 due to the fact their rotation angles θ need not be very close to zero. This is accounted for by the polygonal paths we append to image of Λ under the translation $(x, y) \mapsto (x+m, y+n)$.

□

Theorem 2.11.3. *A nontrivial J_δ-holomorphic curve in a convex toric domain of ECH index one has genus zero. Here we use J_δ to mean we have perturbed away all Morse-Bott degeneracies.*

Proof. We borrow the notation of the previous section, except here e_+ denotes the total multiplicity of elliptic Reeb orbits in α_+ arising from Morse-Bott tori and e_- denotes the total number of distinct elliptic Reeb orbits in α_- arising from perturbations of Morse-Bott tori. The Fredholm index of a connected J-holomorphic curve C between two orbit sets α_+ and α_- is given by

$$\begin{aligned}
Ind(C) =& 2g-2+(e_+ + h_+ + k_m^+ + k_n^+) + (e_- + h_- + k_m^- + k_n^-)\\
&+2(A_+ + B_+ + m_+ + n_+ - A_- - B_- - m_- - n_-)\\
&+e_+ - e_-\\
&+(k_n^+ + k_m^+ + k_m^- + k_n^-)\\
&+\sum_{i=1}^{k_n^+} 2\lfloor -n_+^i/f'(a)\rfloor + \sum_{i=1}^{k_m^+} 2\lfloor -m_+^i f'(0)\rfloor - \sum_{i=1}^{k_n^-} 2\lceil -n_-^i/f'(a)\rceil - \sum_{i=1}^{k_m^-} 2\lceil -m_-^i f'(0)\rceil.
\end{aligned}$$

The same linking number relations as in 2.10.4 holds in the case of convex toric domains; so similarly by considering the intersections of C with the trivial cylinders at γ_1 and γ_2, we conclude

$$A_+ + n_+ + \sum_{i=1}^{k_m^+} \lfloor -m_+^i f'(0)\rfloor - A_- - n_- - \sum_{i=1}^{k_m^-} \lceil -m_-^i f'(0)\rceil \geq 0$$

and

$$B_+ + m_+ + \sum_{i=1}^{k_n^+} 2\lfloor -n_+^i/f'(a)\rfloor - B_- - m_- - \sum_{i=1}^{k_n^-} 2\lfloor -n_-^i/f'(a)\rfloor \geq 0.$$

Hence for C to have genus nonzero it must not have any ends at γ_1 and γ_2.

The local energy inequality holds as before, to prove the no-crossing lemma, we can associate two polygonal paths Λ_+ and Λ_- to ECH generators α_+ and α_- respectively. As before from index considerations the x and y intercepts of Λ_+ and Λ_- agree. Hence as before we can choose points (a,b) and (c,d) where Λ_+ and Λ_- intersect, and between these two points Λ_- is strictly above Λ_+. As before we may choose $x_0 \in (a,c)$ so that $f'(x_0) = \frac{d-b}{c-a}$. Let the lattice point (p',q') have the following property: it is a vertex on Λ_-, the edge to the left of this lattice point has slope greater than or equal to $f'(x_0)$, and the edge to the right of this vertex has slope less than $f'(x_0)$. Let (p,q) denote a vertex of Λ_+ with the same property. We assume such a vertex (p,q) exists and leave the case where such a vertex does not exist to later. Then consider $[F_{x_0+\epsilon}] = (q-q', p'-p)$. Now again the energy inequality says

$$(q-q') + \frac{d-b}{c-a}(p'-p) \geq 0$$

In this case, the point (p,q) is closer to the line connecting (a,b) and (c,d) than (p',q'), but this time on the other side of the line. This means that

$$(p-p')(b-d) + (c-a)(q-q') < 0$$

Comparing with the energy inequality we see a contradiction. Now if (p,q) is in fact on the line connecting (a,b) and (c,d), then since we are computing $[F_{x_0+\epsilon}]$, we must have $p > p'$,

from which we have
$$\frac{d-b}{c-a} > \frac{q-q'}{p-p'}$$
which is a contradiction.

With the no-crossing result at hand, we turn to the index formula. If C had genus one, then
$$1 = 2e_+ + h_+ + h_-.$$
As before we break this into cases. We must have $e_+ = 0$. If $h_+ = 1$ then Λ_+ consists of a single edge, by no-crossing Λ_- is either an identical edge or empty. We check either case cannot produce an ECH index 1 curve. h_+ cannot equal zero because then $\Lambda_+ = \emptyset$.

□

Hence we concluded all ECH index one curves are index zero, a similar description of tree-like cascades shows we can use them to compute the ECH chain complex.

2.12 Appendix: Transversality Issues

In this Appendix we describe some the transversality difficulties in the moduli space of cascades, even if all the appearing curves are somewhere injective. Note we are not claiming transversality is impossible, we are simply saying there are issues with the standard universal moduli space approach of transversality. We give some simple examples below to illustrate this.

Consider the universal moduli space of somewhere injective cascades, written as
$$\mathcal{B} := \{(u^{\lightning}, J)|\, u^{\lightning} \text{ is a } J\text{-holomorphic cascade, and that all curves appearing in } u^{\lightning} \text{ are simple}\}.$$
We explain why the standard proof that $\mathcal{B}$ is a Banach manifold does not necessarily work. Given a cascade $u^{\lightning} \in \mathcal{B}$, there are two evaluation maps EV^+ and EV^- that map into a product of S^1, as in Definition 2.5.5. The usual procedure to show that $\mathcal{B}$ is a Banach manifold is to show the maps $EV^\pm$ are transverse to each other. However in complicated enough cascades, the same curve can appear in multiple different levels. An illustration is given in the figure below. Here we have a cascade of 5 levels. The red curve is a map $u : \Sigma \to \mathbb{R} \times Y^3$, and the blue curve is a map $v : \Sigma' \to \mathbb{R} \times Y^3$. Green horizontal arrows denote the upwards gradient flow, and the black horizontal lines denote Morse-Bott tori. Diamonds denote the critical points of f on the Morse-Bott tori. For instance, one of the positive ends of the black curve ends on a critical point of f, and there is a chain of fixed trivial cylinders atop this end. This is an illustration of how the same curves can happen in the same cascade. To illustrate the transversality issue, we assume that the configuration consisting the red and blue curves (which we labelled u and v) in figure 2.7 happens n times in a cascade $u^{\lightning}$. We assume both u and v are rigid (we are allowed since we are working in the universal moduli space, in general more complicated things can still happen but the principle is the same). We label the n identical copies of u and v as u_i, v_i with $i = 1, ..., n$.

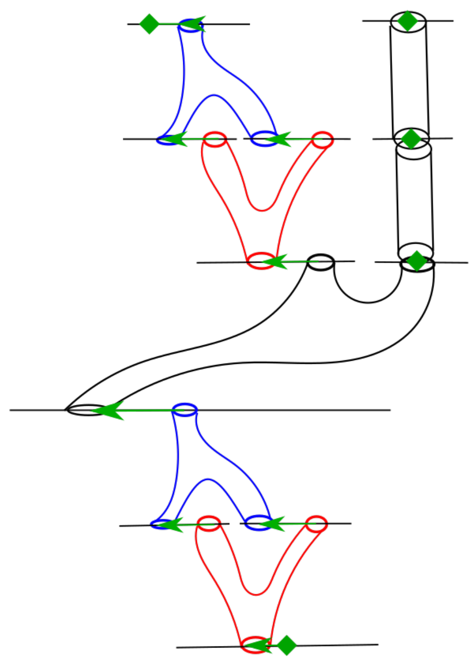

Figure 2.6: Cascade with 5 levels

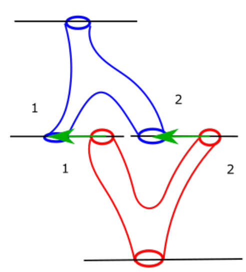

Figure 2.7: A repetitive pattern that can appear multiple times in a cascade.

The two negative ends of u_i and the two positive ends of v_i are labelled by $1, 2$, as shown in the figure. The remaining end of u_i and v_i is labelled 3. We denote their evaluation maps by $ev(u_i, k)$ and $ev(v_i, k)$ where $k = 1, 2, 3$. As a necessary condition for the EV^+ and EV^- to be transverse, we must have

$$\bigoplus (dev(u_i,1) + dev(v_i,1) + t_i, dev(u_i,2) + dev(v_i,1) + t_i) : T\mathcal{W}_u \oplus T\mathcal{W}_v \bigoplus_{i=1,..,n} \mathbb{R}$$

$$\longrightarrow \bigoplus_{i=1,\ldots,n} (TS^1 \oplus TS^1)$$

is surjective. Note $(t_1, ..., t_n) \in \bigoplus_{i=1,..,n} \mathbb{R}$. The vector space $T\mathcal{W}_u$ has the following description. Recall a neighborhood of (not necessarily J holomorphic) curves near u can be represented by $W^{2,p,d}(u^*TM) \oplus T\mathcal{J} \oplus V_1 \oplus V_2 \oplus V_3$. Here $W^{2,p,d}(u^*TM)$ is the Sobolev space of vector fields on u with exponential weight $e^{d|s|}$ near the cylindrical ends. $T\mathcal{J}$ is a finite dimensional Teichmuller slice, and the vector spaces V_i consist of asymptotically constant vectors near each of the cylindrical ends, which we labelled $1, 2, 3$ (see [67, 65]). Recalling the coordinate choices of Section 2.4 near Morse-Bott tori, the V_i is spanned by vector fields of the form

$$\beta\partial_z, \quad \beta\partial_a, \quad \beta\partial_x.$$

β here is a cutoff function that is one near a cylindrical neighborhood of a puncture and zero elsewhere. We denote a triple of these vector fields in V_i as $(r, a, p)_i$.

Then the vector space $\mathcal{W}_u$ is given by

$$\{(\xi, (r, a, p)_i, Y) \in W^{2,p,d}(u^*TM) \oplus T\mathcal{J} \oplus V_1 \oplus V_2 \oplus V_3 \oplus T\mathcal{I} | D\bar{\partial}_J(\xi + \sum_i (r, a, p)_i) + Y \circ Tu \circ j = 0\}$$

$D\bar{\partial}_J$ is the linearization of Cauchy Riemann operator along u that includes deformation of the domain complex structure of u. Here $T\mathcal{I}$ denotes the Sobolev space that is the tangent space of all λ compatible almost complex structures (we should choose a Sobolev space for this but that is unimportant for now). A similar expression holds for $T\mathcal{W}_v$. We note the same $Y \in T\mathcal{I}$ appears in the definition of $T\mathcal{W}_v$ as well. Now since u is rigid for given Y there exists a unique tuple $(\xi, (r, a, p)_i)$ (up to translation in the symplectization direction) so that $(\xi, (r, a, p)_i, Y) \in T\mathcal{W}_u$. A similar statement holds for $\mathcal{W}_v$. Conversely, given two tuples $(p_1(u), p_2(u), p_3(u))$ and $(p_1(v), p_2(v), p_3(v))$ (we use brackets to denote whether the vector field is living over u or v, we can find $Y \in T\mathcal{I}$ and $(\xi(u), (r(u), a(u))_i)$ and $(\xi(v), (r(v), a(v))_i)$ so that the tuples $(\xi(u), (r(u), a(u), p(u))_i, Y) \in T\mathcal{W}_u$, and similarly for $T\mathcal{W}_v$. Hence we can think of the map described in Equation **??** as the following. Its imagine is spanned by vector fields of the form

$$\bigoplus_i (x_1 + y_1 + t_i, x_2 + y_2 + t_i)$$

where (x_1, y_1) and (x_2, y_2) are arbitrary real numbers. We think of x_1 as $p_1(u)$, x_2 corresponding to $p_2(u)$, and likewise for y and $p(v)$. For given n the domain has $2+n$ independent variables, but the target is $2n$ dimensional. Hence for large values of n this space cannot be transverse.

Proof of Theorem 2.2.2. We note if the above situation does not happen, then the usual proof that $\mathcal{B}$ is a Banach manifold follows through. To be precise, if we let $\tilde{B}$ denote the universal moduli space so that

$$\tilde{B} := \left\{ (u^{\natural}, J) \;\middle|\; \begin{array}{l} u^{\natural} \text{ is a reduced } J\text{-holomorphic cascade as in Definition 2.5.5;} \\ \text{in addition, either all nontrivial curves} \\ \text{are distinct, or the cascade has less than or equal to 3 levels} \end{array} \right\} \tag{2.32}$$

Then $\tilde{B}$ is a Banach manifold, and for generic J, cascades satisfying the extra hypothesis of $\tilde{B}$ are transversely cut out living in moduli spaces given by the virtual dimension. In particular if we take as assumption after we perturb away the Morse-Bott degeneracy, all ECH index one curves degenerate (as reduced cascades) to reduced cascades of the form specified in $\tilde{B}$, then we can choose a J so that the conditions 2.5.5 are satisfied for these cascades. A straightforward modification of the proofs in Sections 2.7, 2.8 shows the Morse-Bott chain complex $(C_*^{MB}, \partial_{MB})$ when we further restrict the differential to only consider cascades whose reduced versions can appear in $\tilde{B}$ is well defined and computes $ECH(Y,\xi)$. The only different part is showing the cascades counted by ∂_{MB} is finite. Consider the following. Suppose $u_n^{\lightning}$ is a sequence of cascades of the form allowed in $\tilde{B}$ and $u_n^{\lightning} \to u^{\lightning}$. Then for each $u_n^{\lightning}$ there is a sequence of $J_{\delta_n^m}$-holomorphic curves v_n^m of ECH index one that converges to $u_n^{\lightning}$ as $m \to \infty$. We pass to a diagonal subsequence, which we denote by v_n, of ECH index one J_{δ_n}-holomorphic curves that degenerate into u. By assumption, then the reduced version of $u^{\lightning}$ must be of the form allowed in $\tilde{B}$, and this concludes the proof of finiteness. $\square$

Chapter 3

From cascades to J-holomorphic curves and back

3.1 Abstract

This paper develops the analysis needed to set up a Morse-Bott version of embedded contact homology (ECH) of a contact three-manifold in certain cases. In particular we establish a correspondence between "cascades" of holomorphic curves in the symplectization of a Morse-Bott contact form, and holomorphic curves in the symplectization of a nondegenerate perturbation of the contact form. The cascades we consider must be transversely cut out and rigid. We accomplish this by studying the adiabatic degeneration of J-holomorphic curves into cascades and establishing a gluing theorem. We note our gluing theorem satisfying appropriate transversality hypotheses should work in higher dimensions as well. The applications to ECH is worked out in Chapter 2 [66].

3.2 Introduction

Let (Y^3, λ) be a contact 3-manifold. We assume the Reeb orbits of λ are Morse-Bott and come in S^1 families, i.e. we have tori foliated by Reeb orbits of equal period, which we call Morse-Bott tori. Examples of this include the standard contact structure on the 3-torus, and boundaries of toric domains. See [37], [10]. (Toric domains are also called Reinhardt domains in [25].)

In this setup, one would like to make sense of Floer theoretic invariants constructed via counting J-holomorphic curves in the symplectization of our contact manifold, which we write as

$$\left(\mathbb{R} \times Y^3, d(e^a \lambda)\right).$$

In the above a is the variable in the $\mathbb{R}$ direction, $d(e^a\lambda)$ is the symplectic form. We also fix J to be a (generic) almost complex structure compatible with λ.

However, most versions of Floer homology require the contact form to be non-degenerate. One way to get around this is as follows. We first fix a very large number $L > 0$, and consider the action filtered version of our Floer theory up to action L. We will have embedded contact homology (ECH) in mind when we describe this process, but it also applies to other types of Floer theories assuming suitable transversality. For a Morse-Bott torus with action less than L, which we write as $\mathcal{T}$, we perform a small perturbation of the contact form λ written as

$$\lambda \longrightarrow \lambda_\delta$$

for $\delta > 0$ small, in a small fixed neighborhood of $\mathcal{T}$. Such perturbation requires the information of a Morse function $f : S^1 \to \mathbb{R}$, with two critical points. After this perturbation we also need to change the almost complex structure to J_δ to make it compatible with the new contact form λ_δ.

The effect of this perturbation is so that the Morse-Bott torus $\mathcal{T}$ splits into two nondegenerate Reeb orbits corresponding to the critical points of f, one elliptic and one hyperbolic, and that no other Reeb orbits of action less than L are introduced. We perform this perturbation for all Morse-Bott tori of action less than L. Then in this case, for at least up to action L, we can define our Floer theory with generators as collections of non-degenerate Reeb orbits with total action $< L$ and the differential as counts of J_δ-holomorphic curves connecting between our generators (the details of which Reeb orbits/holomorphic curves to consider depend on whichever Floer theory we choose to work with.)

However, it is often desirable to be able to compute our Floer theory purely in the Morse-Bott setting, in part because often the count of J-holomorphic curves is easier in the Morse-Bott setting. To this end, in order to find out what kind of objects that ought to be counted in the Morse-Bott setting, one can imagine turning the above process around. For given $\delta > 0$, we know how to compute our Floer theory up to action L with the contact form λ_δ via counts of a collection of J_δ-holomorphic curves; then we take the limit of $\delta \to 0$, and see what kind of objects our J_δ-holomorphic curves degenerate into. It turns out in this process J-holomorphic curves degenerate into cascades [5], [7],[20], [6]. See Definition 3.3.7 for the definition of a (height one) cascade, and Definition 3.3.9 for the more general case. See Section 3.3 and the Appendix for a fuller explanation of setup and more precise definition of degeneration of J_δ-holomorphic curves into cascades.

Roughly speaking, a cascade $u^{\cancel{t}} = \{u^1, ..., u^n\}$ consists of a sequence of (not necessarily connected) J-holomorphic curves with ends on Morse-Bott tori. We think of the curves u^i as living on different levels. Between adjacent levels, say u^i and u^{i+1}, there is also the data of a number $T_i \in [0, \infty]$. The negative ends of u^i and positive ends of u^{i+1} are connected by gradient flow segments of length T_i. Said differently, recall each S^1 family of Reeb orbits is equipped with a Morse function f on S^1, and if we start at a Reeb orbit reached by a positive puncture of u^{i+1}, follow the upwards gradient flow of f on S^1 (this S^1 means the S^1 family of Reeb orbits) for time T_i, we will arrive at a Reeb orbit hit by a negative puncture of u^i. The Reeb orbits hit by the positive punctures of u^1 and negative punctures of u^n are connected to Reeb orbits on the Morse-Bott tori corresponding to critical points of f via the upwards gradient flow. The cascade $u^\natural = \{u^1, ..., u^n\}$ has n **levels**, and if in addition no flow time T_i between adjacent curves u^i and u^{i+1} is infinite, we say the cascade has **height** one. See Definitions 3.3.7 and 3.3.9 for slightly more precise definitions. In particular a height one cascade can have arbitrary number of levels, and we mostly concern ourselves with height one cascades in this paper. A schematic picture of a height one cascade (of two levels) is given in Figure 3.1.

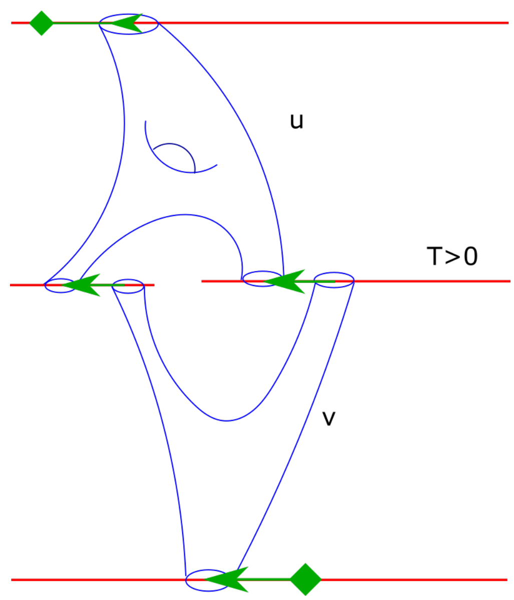

Figure 3.1: A schematic picture of a height one 2-level cascade: the cascade $u^\natural$ consists of two levels, u and v. Horizontal lines correspond to Morse-Bott tori. Moving in the horizontal direction along these horizontal lines corresponds to moving to different Reeb orbits in the same S^1 family. Arrows correspond to gradient flows, and diamonds correspond to critical points of Morse functions on S^1 families of Reeb orbits. Between the holomorphic curves u and v, there is a single parameter T that tells us how long positive ends of v must follow the gradient flow to meet a negative end of u.

We would then like a way to compute Floer homology purely in the Morse-Bott setting via an enumeration of cascades. To prove that the enumeration of cascades recovers the enumeration of J_δ-holomorphic curves in the non-degenerate setting, we would require a correspondence theorem between the two types of objects. The correspondence theorem will of course then involve gluing cascades into J_δ-holomorphic curves. We remark that we currently do not have the technology to glue together all cascades; there are issues pertaining to transversality: curves could be multiply covered, and even if they are somewhere injective and even after a generic choice of J, there could still be non-transverse cascades because we required all negative ends of u^i meet positive ends of u^{i+1} after flowing for a single time length, T_i. In general it is convenient to think of a cascade as existing in a fiber product, and we require the fiber product to be transverse. Also we only concern ourselves with rigid cascades and their correspondences with rigid holomorphic curves. For a more precise definition of transverse and rigid, as well as the description of this fiber product, see Definition 3.4.4. Our version of the gluing theorem should work for gluing higher index (transversely cut out) cascades, but making sense of a correspondence between two high dimensional moduli spaces could be much trickier. With the above preamble we state in a slightly imprecise way our main theorem:

Theorem 3.2.1. *Given a transverse and rigid height one J-holomorphic cascade $u^{\natural}$, it can be glued to a rigid J_δ-holomorphic curve u_δ for $\delta > 0$ sufficiently small. The construction is unique in the following sense: if $\{\delta_n\}$ is a sequence of numbers that converge to zero as $n \to \infty$, and $\{u'_{\delta_n}\}$ is sequence of J_{δ_n}-holomorphic curves converging to $u^{\natural}$, then for large enough n, the curves u'_{δ_n} agree with u_{δ_n} up to translation in the symplectization direction.*

See Definition 3.4.4 for the description of "transverse and rigid". See Theorem 3.4.5 for a more precise formulation of this theorem.

Remark 3.2.2. The purpose of Morse-Bott theory is usually that J-holomorphic curves are often more easily enumerated in the Morse-Bott setting due to presence of symmetry. While cascades are built out of J-holomorphic curves in the Morse-Bott situation, counting them explicitly can be difficult in its own way. Even though rigid and transverse cascades are themselves discrete, they may be built out of curves that live in high dimensional moduli spaces. Since in principle arbitrarily high dimensional moduli spaces can show up, one usually needs some extra simplifications for the enumeration of cascades to be tractable.

Remark 3.2.3. Since we will have future applications to ECH in mind, we make some comments about our "transverse and rigid" condition versus the ECH index one condition:

- In general restricting to cascades that have ECH index one (of course one first needs to extend the notion of ECH index one to cascades) and choosing a generic J does not necessarily imply the cascades we get are transversely cut out. However there are special cases where transversality can be achieved by restricting to ECH index one cascades, and the correspondence theorem (Theorem 3.4.5) would allow us to compute ECH using an enumeration of J-holomorphic cascades. For details of computing ECH using cascades, see Chapter 2 [66].
- If we already had cascades that are transverse and rigid, from a gluing point of view, further restricting to the cascades that have ECH index one does not change very much: it just implies all the curves in the cascade are embedded (with the exception of unbranched covers of trivial cylinders) and distinct curves within each level do not intersect each other. We further have some partition conditions on the ends of holomorphic curves in the cascade, but again from a gluing point of view this does not make a difference.

Relations to other work

The idea of doing Morse-Bott homology certainly isn't new. Methods of working with Morse-Bott homology predate the construction of cascades, and were described in [2], [21]. The construction of cascades was discovered independently in [5] and [20]. There were then a plethora of constructions of Floer-type theories using cascades (or in many cases, constructions very similar to cascades). For Lagrangian Floer theory, in addition to [20], there was also [4]. For symplectic homology, see [7]. See also [54]. For Morse homology, see [3], [28]. For special cases of contact homology, see [34],[51]. For the special case of ECH where the cascades can only have one *level* and in context of stable Hamiltonian structures, see [12]. For abstract perspectives on Morse-Bott theory, see [69], [35]. Finally, the gauge theory analogue of ECH, monopole Floer homology, has a Morse-Bott version constructed in [47], though there they do not use a cascade model.

For cascades there are two general approaches to show the Morse-Bott homology theory constructed agrees with the original homology theory. One way is to show the differential obtained via counts of cascades squares to zero, hence one has *some* homology theory. Then one shows that this homology theory is isomorphic to the original by constructing a cobordism interpolating the Morse-Bott geometry and the non-degenerate geometry. For standard Floer theory reasons this cobordism induces a cobordism map between the two homology groups. Also for standard Floer theory reasons we could show this cobordism map induces an isomorphism on homology. This is the approach taken in [20], [4].

The other approach is to directly show that non-degenerate holomorphic curves degenerate into cascades in the $\delta \to 0$ limit, and there is a correspondence between cascades and holomorphic curves. This degeneration of holomorphic curves into cascades is also sometimes called the adiabatic limit. This approach of computing Morse-Bott homology is taken in [7] [5] [3] [54]. This is also the approach we take here. We prove the correspondence theorem under transversality assumptions (Definition 3.4.4), and the applications to ECH are in a separate paper (Chapter 2 [66]).

The reason we take the latter approach is that in ECH, which is the application we have in mind, everything except transversality is very hard. That the differential squares to zero requires 200 pages of obstruction bundle gluing calculations [40] [41], and a similar story must be repeated in the Morse-Bott case for showing the count of ECH index one cascades defines a chain complex. Constructing cobordism maps in ECH is even harder, and generally relies on passing to Seiberg-Witten theory. Cobordism maps on ECH defined purely using holomorphic curves techniques have only been worked out for very special cases [56], [8],[23],[22]. Hence in light of these difficulties, it would seem the path of least resistance would be to prove a correspondence theorem and do the adiabatic limit analysis for ECH, despite this being a generally difficult approach.

We must highlight the relation of our work with [7], which produces a correspondence theorem in the case of symplectic homology. We borrowed heavily the techniques of that paper in the areas of analysis of linear operators over gradient flow trajectories (most notably the construction of uniformly bounded right inverses in the $\delta \to 0$ limit), as well as the degeneration of holomorphic curves into gradient trajectories near Morse-Bott tori. Both of these ideas have previously appeared in [5] but were worked out in more detail in [7]. However, our construction of gluing is markedly different from [7], as we were unfortunately unable to adopt their approach. Instead, our approach of both gluing and proving the gluing procedure produces a bijection between cascades and holomorphic curves mirrors the approach of [40] [41], the two papers where Hutchings and Taubes show the differential in ECH squares to zero using obstruction bundle gluing. In particular, our approach can in fact be rephrased in terms of obstruction bundle gluing, see Remark 3.9.23, though in our case the obstruction bundle gluing is particularly simple and can be thought of as an application of the intermediate value theorem. For a formulation of this kind of gluing results in a simpler case in ECH where there is only one level in the cascades using the language of obstruction bundle gluing, see the Appendix of [12], which we wrote jointly with Colin, Ghiggini and Honda.

We briefly outline the differences between our approach to gluing compared to those in [7], [54]. In [7], [54], the gradient trajectories connecting different levels of the cascade are preglued to the J-holomorphic curves in the cascade; they consider the deformations of the entire preglued curve, and use the implicit function theorem to obtain gluing results. In our approach, in following the approach of [40], [41], the condition that a cascade can be glued to a J_δ-holomorphic curve is translated into a *system* of coupled nonlinear PDEs, which we loosely write as $\{\boldsymbol{\Theta}_i = 0\}$. Gluing is established by systematically solving this system of PDEs. How this is accomplished is explained first in a simplified setting in Section 3.7, then in the general case in Section 3.9.

To briefly explain the difficulties of the gluing problem, we first observe that here we are gluing with less transversality than the usual gluing problems. In the usual case when ends of pseudo-holomorphic curves meet along a critical submanifold (for instance the gluing problem described in Chapter 10 of [48]), in order to glue we require the evaluation maps of the respective moduli spaces of pseudo-holomorphic curves to be transverse to each other. Here, however, we require the evaluation maps to be transverse up to a flow time parameters T_i (see Definition 3.4.4 for the precise transversality conditions). Hence we adopt the obstruction bundle gluing setup of Hutchings and Taubes ([40, 41]) to find a gluing with less transversality by solving a sytem of PDEs (We do not explicitly use obstruction bundles, but they are implicit in our setup. See remark 3.9.23 for an explanation).

To say a bit more about this system of PDEs, we note that in this system there is a PDE for each J-holomorphic curve that appears in the cascade, and a PDE for each upwards gradient trajectory. In some sense this allows us to think about deformations of J-holomorphic curves and deformations of gradient trajectories separately from each other (of course in the end the equations are coupled, so this is only metaphorically true). The main technical difficulty this addresses is that in considering cascades, gradient trajectories and J-holomorphic curves are in some sense different kinds of objects. For small values of $\delta > 0$, measured in a suitable norm, J-holomorphic curves in the cascade are very close to being J_δ-holomorphic curves in the perturbed picture, but the gradient flow trajectories in the cascade come from very long gradient flow cylinders (their lengths go to ∞ as $\delta \to +\infty$)

that follow a very slow gradient flow (the gradient flow of δf). For a description of these cylinders see Section 3.5. The consequence of this is that deformations that appear to be small from the perspective of J-holomorphic curves in the cascade can be extremely large from the perspective of gradient flow cylinders in the cascade. See Figure 3.2 and the accompanying explanations. Hence a sizable portion of the work after writing down a system of equations $\{\boldsymbol{\Theta}_i = 0\}$ is to keep track of which deformations are very large from the perspective of gradient flow trajectories, and to understand the effects of these deformations on the equations in $\{\boldsymbol{\Theta}_i = 0\}$ and the way different equations in the system $\{\boldsymbol{\Theta}_i = 0\}$ are coupled to each other.

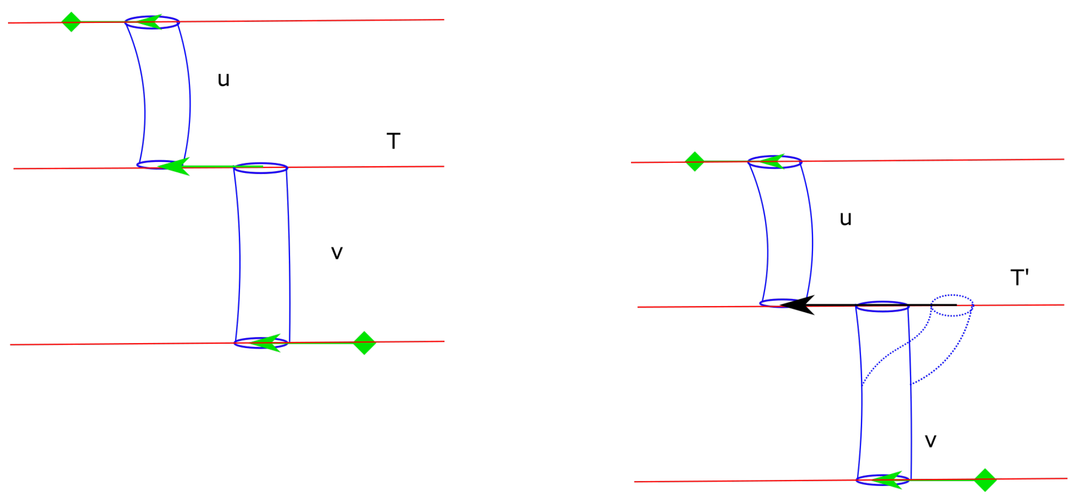

Figure 3.2: On the left we see a cascade of two levels consisting of the curves $\{u, v\}$. They meet along a Morse-Bott torus in the middle, and their ends are connected by a gradient flow trajectory of length T, shown by the green arrow. On the right, we imagine slightly displacing the end of the curve v along the Morse-Bott torus, shown in dashed blue lines. From the perspective of v, measured with appropriate norms this is a small deformation of v. The flow time from this deformed v to u is given by $T' = T + \Delta T$, which is a slightly longer flow time, indicated by the black arrow. However the picture is deceiving, because for small values of δ, the gradient flow cylinder corresponding to the black arrow is significantly longer than the gradient cylinder of the green arrow in the original picture, by an additional length of order $\Delta T/\delta$. This is because for small values of δ, the gradient flow is very slow (it follows the gradient of δf), hence it needs to flow for very long to cover that extra distance. This is what we mean when we say deformations that can seem very small from the perspective of J-holomorphic curves in the cascade can be arbitrarily large from the perspective of gradient trajectories.

Finally, we remark that despite only working with Morse-Bott tori, we expect our approach to work for most Floer theories based on counts of holomorphic curves that do not have multiple covers or issues with transversality (both in the non-degenerate setting and the Morse-Bott setting). We expect the generalization from Reeb orbits showing up in S^1 families to higher dimensional families to be straightforward, and the rest of the analysis should carry over directly. However, we do not know how our analysis or proof of correspondence theorem interact with the virtual techniques that are often used to define Floer theories when classical transversality methods fail.

Applications to ECH

As mentioned above the main application we have in mind of this work is the computation of embedded contact homology in the Morse-Bott setting. Previously several computations of ECH (or its related cousin periodic Floer homology) have assumed results about Morse-Bott theory and cascades, for instance computations in [39], [37], [9].

This paper does not contain the full construction of Morse-Bott ECH, but the analysis done here lays the groundwork for constructing a correspondence theorem for ECH index one cascades and ECH index one holomorphic currents. This is worked out in detail in Chapter 2 [66].

In particular, the machinery developed here and in Chapter 2 [66] fill in the foundations for Morse-Bott ECH for the computations in [39], [37], [9].

For contact 3-manifolds, Morse-Bott degeneracy might also mean the Morse-Bott critical manifold is two dimensional, which means the manifold itself is foliated by Reeb orbits. The computation of ECH in that case was done using different techniques, see [52], [18]. However, our methods (suitably extended to allow for the case where Reeb orbits can come in higher dimensional families) could potentially be applied to ECH computations in these cases as well.

Applications beyond ECH

Morse-Bott situations are ubiquitous in contact geometry. In many instances, to conclude results about the behaviour J-holomorphic curves in nondegenerate settings, a correspondence in the style of Theorem 3.4.5 has been cited without complete proofs. Examples include the computation of the contact invariant in ECH in the presence of planar torsion in [64, Theorem 2], and the computations of algebraic torsion in [45]. We expect this paper (see also the appendix of [12] for cascades of one level in stable Hamiltonian settings) to fill in many gaps in the existing literature. In particular we note that our gluing construction can be modified in a straightforward manner to glue together cascades in higher dimensional contact manifolds (and higher dimensional Morse-Bott critical submanifolds). Hence we expect our main theorem to be a useful tool to understanding the behaviour of J-holomorphic curves in many applications to come.

Outline

The paper is organized as follows. After some quick descriptions of the geometric setup, we describe in Section 3.3 how holomorphic curves in the non-degenerate case degenerate into objects we call cascades, and introduce a version of SFT type compactness, already introduced in [6], [5]. We relegate the more technical definitions of convergence and the proof of degeneration into cascades to the Appendix for the sake of exposition.

In Section 3.4 we establish what we mean by a generic choice of J, the definition of transversality, and in particular describe the set of cascades we will be able to glue into J-holomorphic curves.

Then we get to the most technical part of the paper, in which we prove transverse and rigid J-holomorphic cascades can be glued to J_δ-holomorphic curves as we perturb the contact form. We first start with some preamble on differential geometry in Section 3.5, and describe the gradient trajectories that arise from perturbing the contact form. We then find a suitable Sobolev space for the gradient trajectories which we will use for our gluing, and prove some nice properties of the linearized Cauchy Riemann operator in this Sobolev space for later use in Section 3.6.

To initiate the gluing, first as a warm up we explain how to glue a semi-infinite gradient trajectory to a J-holomorphic curve in Section 3.7. This corresponds to gluing 1-level cascades to J_δ-holomorphic curves as we perturb the contact form from λ to λ_δ, which is also done in [12]. We then prove an important property of the curve we constructed in this process, i.e. the solution exponentially decays along the gradient trajectory. This is done in Section 3.8, and will be crucial for gluing together multiple level cascades.

Section 3.9, we complete the gluing construction. We first consider the simplified case of gluing together 2-level cascades, which will contain the heart of the construction and is markedly different from gluing semi-infinite trajectories. As before we first do some basic Sobolev space setup. The key idea is to first preglue, then use the solution constructed for semi-infinite trajectories to construct another pregluing on top of the original pregluing with substantially smaller pregluing error, and then use the contraction mapping principle one last time to turn the second pregluing into a genuine gluing. As illustrated in Figure 3.2, during the $\delta \to 0$ degeneration the gradient flow cylinders correspond to very long necks, and when we try to preglue a cascade, deformations that appear small from the perspective of J-holomorphic curves in the cascade can become very large from the perspective of the preglued curve when we try to fit a gradient flow cylinder between adjacent levels of the cascade during the pregluing. So all of the complications in the gluing we mentioned above arise from trying to keep track of these deformations and finding a setup where all of the vectors that we see are sufficiently small, so that the contraction mapping principle can be applied.

After this, the generalization to multiple level cascades is mostly a matter of keeping track of notation.

In anticipation of proving bijectivity of gluing, we deduce some analytic estimates of how J_δ-holomorphic curves behave near Morse-Bott tori as we degenerate the contact form λ_δ.

This is done in Section 3.10. Much of this analysis is taken from the appendix of [7] where they work out a very similar case in symplectic homology. This kind of analysis has also appeared in [5].

Finally we take up the bijectivity of gluing; for this step we largely follow the footsteps of [40] [41]. This is taken up in Section 3.11.

The appendix contains the necessary background to state the SFT compactness theorem required for our kind of degenerations, which was stated in [6] and proved in [5]. A similar result also appears in [7]. We also provide a proof for completeness, which relies also on the analysis done in Section 3.10.

Acknowledgements

First and foremost I would like to thank my advisor Michael Hutchings for his consistent help and support throughout this project. I would also like to thank Ko Honda, Jo Nelson, Alexandru Oancea, Katrin Wehrheim, Morgan Weiler and Chris Wendl for helpful discussions and comments.

I would like to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), PGSD3-532405-2019. Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), PGSD3-532405-2019.

3.3 Morse-Bott setup and SFT type compactness

Let (Y^3, λ) be a contact 3-manifold with Morse-Bott contact form λ. Throughout we assume all Reeb orbits come in S^1 families; hence we have tori foliated by Reeb orbits, which we call Morse-Bott tori.

Convention 3.3.1. *Throughout this paper we fix a large number $L > 0$, and only consider collections of Reeb orbits that have total action less than L. This is implicit in all of our constructions and will not be mentioned further. We prove the correspondence theorem between cascades and J_δ-holomorphic curves up to action level L, and then when we need to apply this construction to Floer theories we can take $L \to \infty$.*

The following theorem, which is a special case of a more general result in [53], gives a characterization of the neighborhood of Morse-Bott tori. Let λ_0 denote the standard contact form on $(z, x, y) \in S^1 \times S^1 \times \mathbb{R}$ of the form

$$\lambda_0 = dz - ydx.$$

Proposition 3.3.2. (M^3, λ) *be a contact 3 manifold with Morse-Bott contact form λ. We assume the Morse-Bott Reeb orbits come in families of tori, which we write as $\mathcal{T}_i$, with minimal period T_i. Then we can choose coordinates around each Morse-Bott torus so that* *a neighborhood of $\mathcal{T}_i$ is described by $(z,x,y) \in S^1 \times S^1 \times (-\epsilon,\epsilon)$, and the contact form λ in this coordinate system looks like*

$$\lambda = h(x,y,z)\lambda_0$$

where $h(x,y,z)$ satisfies

$$h(x,0,z) = 1, \quad dh(x,0,z) = 0.$$

Here we identify $z \in S^1 \sim \mathbb{R}/2\pi T_i\mathbb{Z}$.

Proof. This is implicitly in [53]. We need to apply the setup of [53] Theorem 4.7 to [53], Theorem 5.1.

In [53] Theorem 4.7, in their notation we have $E = 0$, $Q = S^1 \times S^1$, with coordinates (z,x), and $\theta = dz$. The foliation $\mathcal{N}$ is given by $\{z\} \times S^1$. The contact form on the total space of the fiber bundle $F = T^*\mathcal{N} = \mathbb{R} \times T^2$ is given by $dz + ydx$. Our proposition then follows from Theorem 5.1 in [53] (we need to take another transformation $y \to -y$ to get our specific choice of contact form, our sign conventions for the contact form are different from those of [5].) $\square$

We assume we have chosen above neighborhoods around all Morse-Bott tori $\mathcal{T}_i$ with action less than L. By the Morse-Bott assumption there are only finitely many such tori up to fixed action L. Next we perturb them to nondegenerate Reeb orbits by perturbing the contact form in a neighborhood of each torus.

Let $\delta > 0$, let $f : x \in \mathbb{R}/\mathbb{Z} \to R$ be a smooth Morse function with max at $x = 1/2$ and minimum $x = 0$. Let $g(y) : \mathbb{R} \to \mathbb{R}$ be a bump function that is equal to 1 on $[-\epsilon_{\mathcal{T}_i}, \epsilon_{\mathcal{T}_i}]$ and zero outside $[-2\epsilon_{\mathcal{T}_i}, 2\epsilon_{\mathcal{T}_i}]$. Here $\epsilon_{\mathcal{T}_i}$ is a number chosen for each $\mathcal{T}_i$ small enough so that the normal form in the above theorem applies, and that all such chosen neighborhoods of Morse-Bott tori of action $< L$ are disjoint. Then in a neighborhood of the Morse Bott torus $\mathcal{T}_i$, we perturb the contact form as

$$\lambda \longrightarrow \lambda_\delta := e^{\delta g f}\lambda.$$

We can describe the change in Reeb dynamics as follows:

Proposition 3.3.3. *For fixed action level $L > 0$ there exists $\delta > 0$ small enough so that the Reeb dynamics of λ_δ can be described as follows. In the neighborhood specified by Proposition 3.3.2, each Morse-Bott torus splits into two non-degenerate Reeb orbits corresponding to the two critical points of f. One of them is hyperbolic of index 0, the other is elliptic with rotation angle $|\theta| < C\delta << 1$ and hence its Conley-Zehnder index is ± 1. There are no additional Reeb orbits of action $< L$.*

Definition 3.3.4. *We say an Morse-Bott torus is **positive** if the elliptic Reeb orbit has Conley Zehnder index 1 after perturbation, otherwise we say it is **negative** Morse Bott torus. This condition is intrinsic to the Morse-Bott torus itself, and is independent of perturbations.*

Proof of Proposition 3.3.3. After we have fixed our local neighborhood near a Morse-Bott torus from Proposition 3.3.2, we get natural trivializations of the contact plane along the Morse-Bott torus given by the $x-y$ plane. With this trivialization in mind, the linearized return map takes either of the following forms [1]

- Positive Morse-Bott Torus: $\phi(t) = \begin{bmatrix} 1 & -t \\ 0 & 1 \end{bmatrix}$.
- Negative Morse-Bott torus: $\phi(t) = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}$.

They are degenerate, but they admit a Robbin-Salamon index, see Section 4 of [24]. The positive Morse-Bott torus has Robbin-Salamon index $1/2$ and the negative Morse-Bott torus has Robbin-Salamon index $-1/2$ (see [24], Proposition 4.9). Then the claims behaviour of Reeb orbits follow from Lemmas 2.3 and 2.4 in [5]. □

Remark 3.3.5. Later when we define various terms in the Fredholm index, they will depend on choices of trivializations of the contact structure along the Reeb orbits. We will always choose the trivializations specified by Proposition 3.3.2, and where the return maps take the form specified above. For notational convenience we will call this trivialization τ.

We also observe that after iterating the Reeb orbits in the Morse-Bott tori, their Robbin-Salamon indices stay the same. So up to action L, in the nondegenerate picture, we will only see Reeb orbits of Conley-Zehnder indices in the set $\{-1,0,1\}$.

Let us consider for small $\delta > 0$ the symplectization

$$(M^4, \omega_\delta) := (\mathbb{R} \times Y^3, de^a \lambda_\delta).$$

We also consider the symplectization in the Morse-Bott case

$$(M^4, \omega_0) := (\mathbb{R} \times Y^3, de^a \lambda).$$

We fix our conventions for almost complex structures for the rest of the article as follows:

Convention 3.3.6. *We equip (M, ω_0) with λ compatible almost complex structure J (for purposes of tranversality Definition 3.4.4 we may want to take J to be generic). We restrict J to take the following form near a neighborhood of each Morse-Bott torus (if we are using the action filtration we can only require this condition for Morse-Bott tori up to action L).*

[1] This fact is referenced in Section 4 of [12], and Section 5 of [30]. Section 3 of the paper [26] works out the detailed computation leading up to this result. We remark that in all of these three papers the linearized return map is lower triangular. This is because we have chosen different conventions. For instance in [26] they chose their y coordinate to denote the S^1 family of Reeb orbits, and their x coordinate to denote the normal direction to their Morse-Bott torus. Hence their contact form is written as $dz + xdy$. Our linearized return maps agree with theirs after we switch to their coordinate system.

Recall each Morse-Bott torus has neighborhood described by $(a, z, x, y) \in \mathbb{R} \times S^1 \times S^1 \times \mathbb{R}$, then on the surface of the Morse-Bott torus, i.e. $y = 0$, we require

$$J\partial_x = \partial_y.$$

Our requirement for J_δ is that it is λ_δ compatible, and in a neighborhood of each Morse-Bott torus (resp. Morse-Bott tori up to action L), its restriction to the contact distribution agrees with the restriction of J. See Remark 3.4.6 for additional comments for genericity.

For fixed $L > 0$ large and $\delta > 0$ small enough, all collections of orbits with total action less than L are non-degenerate, and hence there are corresponding J-holomorphic curves with energy less than L with non-degenerate asymptotics. To motivate our construction, we next take $\delta \to 0$ to see what these J-holomorphic curves degenerate into. By a theorem that first appeared in Bourgeois' thesis [5] (Chapter 4) and also stated in [6] (Theorem 11.4), they degenerate into J-holomorphic cascades. (For a more careful definition see the appendix that takes into account of stability of domain and marked points, but the definition here suffices for our purposes).

Definition 3.3.7. *[5] Let Σ be a punctured (nodal) Riemann surface, potentially with multiple components. A cascade of height 1, which we will denote by $u^{\natural}$, in $(\mathbb{R} \times Y^3, \lambda, J)$ consists of the following data :*

- *A labeling of the connected components of $\Sigma^* = \Sigma \setminus \{nodes\}$ by integers in $\{1, ..., l\}$, called levels, such that two components sharing a node have levels differing by at most 1. We denote by Σ_i the union of connected components of level i, which might itself be a nodal Riemann surface.*
- *$T_i \in [0, \infty)$ for $i = 1, ..., l-1$.*
- *J-holomorphic maps $u^i : (\Sigma_i, j) \to (\mathbb{R} \times Y^3, J)$ with $E(u_i) < \infty$ for $i = 1, ..., l$, such that:*
  - *Each node shared by Σ_i and Σ_{i+1}, is a negative puncture for u^i and is a positive puncture for u^{i+1}. Suppose this negative puncture of u^i is asymptotic to some Reeb orbit $\gamma_i \in \mathcal{T}$, where $\mathcal{T}$ is a Morse-Bott torus, and this positive puncture of u^{i+1} is asymptotic to some Reeb orbit $\gamma_{i+1} \in \mathcal{T}$, then we have that $\phi_f^{T_i}(\gamma_{i+1}) = \gamma_i$. Here $\phi_f^{T_i}$ is the upwards gradient flow of f for time T_i. It is defined by solving the ODE*
  $$\frac{d}{ds}\phi_f(s) = f'(\phi_f(s)).$$
  - *u^i extends continuously across nodes within Σ_i.*
  - *No level consists purely of trivial cylinders. However we will allow levels that consist of branched covers of trivial cylinders.*

With $u^\dagger$ defined as above, we will informally write $u^\dagger = \{u^1, .., u^l\}$.

Convention 3.3.8. *We fix our conventions as follows.*

- *We say the punctures of a J-holomorphic curve that approach Reeb orbits as $a \to \infty$ are positive punctures, and the punctures that approach Reeb orbits as $a \to -\infty$ are negative punctures. We will fix cylindrical neighborhoods around each puncture of our J-holomorphic curves, so we will use "positive/negative ends" and "positive/negative punctures" interchangeably. By our conventions, we think of u^1 as being a level above u^2 and so on.*

- *We refer to the Morse-Bott tori $\mathcal{T}_j$ that appear between adjacent levels of the cascade $\{u^i, u^{i+1}\}$ as above, where negative punctures of u^i are asymptotic to Reeb orbits that agree with positive punctures from u^{i+1} up to a gradient flow, intermediate cascade levels.*

- *We say that the positive asymptotics of $u^\dagger$ are the Reeb orbits we reach by applying ϕ_f^∞ to the Reeb orbits hit by the positive punctures of u^1. Similarly, the negative asymptotics of $u^\dagger$ are the Reeb orbits we reach by applying $\phi_f^{-\infty}$ to the Reeb orbits hit by the negative punctures of u^l. We note if a positive puncture (resp. negative puncture) of u^1 (resp. u^l) is asymptotic to a Reeb orbit corresponding to a critical point of f, then applying $\phi_f^{+\infty}$ (resp. $\phi_f^{-\infty}$) to this Reeb orbit does nothing.*

Definition 3.3.9 ([5], Chapter 4)**.** *A cascade of height k consists of k height 1 cascades, $u_k^\dagger = \{u^{1\dagger}, ..., u^{k\dagger}\}$ with matching asymptotics concatenated together. By matching asymptotics we mean the following. Consider adjacent height one cascades, $u^{i\dagger}$ and $u^{i+1\dagger}$. Suppose a positive end of the top level of $u^{i+1\dagger}$ is asymptotic to the Reeb orbit γ (not necessarily simply covered). Then if we apply the upwards gradient flow of f for infinite time we arrive at a Reeb orbit reached by a negative end of the bottom level of $u^{i\dagger}$. We allow the case where γ is at a critical point of f, and the flow for infinite time is stationary at γ. We also allow the case where γ is at the minimum of f, and the negative end of the bottom level of $u^{i\dagger}$ is reached by following an entire (upwards) gradient trajectory connecting from the minimum of f to its maximum. If all ends between adjacent height one cascades are matched up this way, then we say they have matching asymptotics.*

We will use the notation $u_k^\dagger$ to denote a cascade of height k. We will mostly be concerned with cascades of height 1 in this article, so for those we will drop the subscript k and write $u^\dagger = \{u^1, ..., u^l\}$.

Remark 3.3.10. In this paper our families of Reeb orbits are parameterized by S^1, and in particular there are no broken gradient flow lines on S^1. In general, when the critical manifold (the manifold that parameterizes the Morse-Bott family of Reeb orbits) is more complicated, the notion of matching asymptotics between height one cascades mentioned in the above definition involves going from a Reeb orbit hit by a positive puncture of the top level of $u^{i+1\sharp}$ to a Reeb orbit hit by a negative puncture of the bottom level of $u^{i\sharp}$ via broken Morse trajectories on the critical manifold.

Remark 3.3.11. Once we have given the definition of cascades, we must then describe what it means for two cascades to be equivalent to each other. The precise definition of when two cascades are equivalent to one another can only be more precisely stated after we have given the more precise definition of cascades in the Appendix, where we keep track of all of the marked points and punctures of each level. Essentially we simply need to adapt the definition of when SFT buildings are equivalent to one another as stated in [6] Section 7.2 by viewing gradient flow trajectories in cascades as extra levels. Here we just remark that for our gluing purposes this is not really an issue for us, all of the cascades we care about (see Definition 3.4.4) will have $u^i : \Sigma_i \to \mathbb{R} \times Y$ be somewhere injective J-holomorphic curves, with the possible exception of unbranched covers of trivial cylinders, hence for us it will be obvious when two cascades are equivalent to one another.

Now we state informally our version of the SFT compactness theorem, the full version with a precise definition of convergence is stated in the Appendix.

Theorem 3.3.12. *A sequence of J_δ-holomorphic curves $\{u_{\delta_n}\}$ that have fixed genus, are asymptotic to the same non-degenerate Reeb orbits, and $\delta_n \to 0$, has a subsequence that converges to a J-holomorphic cascade of height k.*

Remark 3.3.13. It is apparent, with the definition of convergence outlined in the Appendix, that if u_{δ_n} converges to a cascade $u_k^\sharp$ of height k, and all the curves in the cascade are somewhere injective (except unbranched covers of trivial cylinders), then this limit $u_k^\sharp$ is unique up to equivalence.

3.4 Transversality

In this section we describe the necessary transversality hypothesis we need for gluing and the correspondence theorem.

We fix a metric g that is invariant under $\mathbb{R}$ which we shall use for linearization purposes. We require that it is of the form

$$g = da^2 + dx^2 + dy^2 + dz^2$$

in a neighborhood of each Morse-Bott torus.

We also note the following convention that will be followed throughout this paper:

Convention 3.4.1. *Since we will be doing a lot of gluing in the paper there is a lot of demand for various cut off functions. We fix once and for all our convention for cut off functions. We use the notation $\beta_{a;b,c;d} : \mathbb{R} \to \mathbb{R}$ to denote a function with support in (b, c), all of its derivatives are also supported in this interval. $\beta_{a;b,c;d}$ is equal to 1 on the interval*

$(b+a, c-d)$, and over the interval $(a, b+a)$ it satisfies a derivative bound of the form $|\beta'(s)| \leq C/(a)$, and likewise for the interval $(c-d, c)$.

If we would want cut off functions that are equal to 1 at either $\pm\infty$, we will write $\beta_{-\infty,c;d}$ or $\beta_{a;b,\infty}$. The behaviour of the cut off function on intervals $(c-d, c)$ (resp. $(a, b+a)$) is the same as the above paragraph.

Let $u : \dot{\Sigma} \to (Y \times \mathbb{R}, \lambda)$ denote a holomorphic curve from a punctured Riemann surface $\dot{\Sigma}$ with $N_{\pm}$ positive (resp. negative) punctures labeled $p_j^{\pm}$, the collection of which we denote by $\Gamma_{\pm}$. For each puncture $p_j^{\pm}$ we fix cylindrical neighborhoods around the puncture of the form $(s,t) \in [0, \pm\infty) \times S^1$. The punctures of $\dot{\Sigma}$ are asymptotic to Reeb orbits on Morse-Bott tori. There are moduli spaces, which we generally write as $\mathcal{M}$, of J-holomorphic maps from $\dot{\Sigma} \to (Y \times \mathbb{R}, \lambda)$, which can be specified as follows. For given puncture $p_j^{\pm}$, we first specify which Morse-Bott torus $\mathcal{T}_j^{\pm}$ that it lands on and with what multiplicity it covers the Reeb orbits on that Morse-Bott torus. Then we have the option of specifying whether this end is "free" or "fixed", and each choice will lead to a different moduli space. By "free" end we mean elements in the moduli space can have their $p_j^{\pm}$ puncture be asymptotic to any Reeb orbit on $\mathcal{T}_j^{\pm}$ with given multiplicity. By "fixed" end we mean elements in the moduli space must have their $p_j^{\pm}$ end land on a specific Reeb orbit in $\mathcal{T}_j^{\pm}$ with given multiplicity. With this designation it will enough to specify a moduli space of J-holomorphic curves. The virtual dimension of this moduli space with the above specifications, is given by (See Section 3 of [65] or Corollary 5.4 of [5])

$$Ind(u) := -\chi(u) + 2c_1(u) + \sum_{p_j^+} \mu\left(\gamma^{q_{p_j^+}}\right) - \sum_{p_j^-} \mu\left(\gamma^{q_{p_j^-}}\right) + \frac{1}{2}\#\text{free ends} - \frac{1}{2}\#\text{fixed ends} \quad (3.1)$$

where χ is the Euler characteristic, c_1 the first Chern class, $\mu(-)$ is the Robbin Salamon index for path of symplectic matrices with degeneracies defined in [24]. The letter γ denotes the embedded Reeb orbit the end $p_j^{\pm}$ is asymptotic to, with covering multiplicity $q_{p_j^{\pm}}$.

To explain the notation we think of u as being an element of the moduli space $\mathcal{M}$, and $Ind(u)$ is the Fredholm index of u. Implicitly when we write $Ind(u)$ we are including the information of which punctures of $\dot{\Sigma}$ are considered free/fixed. We also note in constructing this moduli space the complex structure of $\dot{\Sigma}$ is allowed to vary.

This moduli space can be viewed as the zero set of a Fredholm map. We borrow the set up as explained in Section 3.2 of [65]. To this end, consider the space of vector fields $W^{2,p,d}(u^*TM)$ with exponential weights at the cylindrical ends of the form $e^{d|s|}$. We consider the map, following [65][2] Section 3.2

$$D_J : W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J} \longrightarrow W^{1,p,d}(\overline{\operatorname{Hom}}(T\dot{\Sigma}, u^*TM)) \quad (3.2)$$

where $V_\Gamma := \oplus V_j^{\pm}$ is a direct sum of vector spaces for each puncture $p_j^{\pm}$. For a positive puncture at a fixed end, it is 2 dimensional vector space spanned by vector fields

$$\beta_{1;0,\infty}\partial_t, \quad \beta_{1;0,\infty}\partial_a$$

[2]In [65] this operator is denoted by $D\bar{\partial}_J$.

where $a \in \mathbb{R}$ is the symplectization coordinate. For negative punctures we use instead cut off functions $\beta_{-\infty,0;1}$.

For free ends we additionally include another asymptotic vector that displaces the ends along the Morse-Bott torus

$$\beta_* \partial_x$$

where β_* is as above, depending on whether this end is at a positive or negative puncture.

$T\mathcal{J}$ is a finite dimensional vector space corresponding to the variation of complex structure (in [65] Section 3.1 it is called a Teichmuller slice). We note we have chosen the variation of complex structure to be supported away from the fixed cylindrical neighborhoods.

Remark 3.4.2. It will later turn out very important to us we work with $W^{2,p}$ as our domain instead of $W^{1,p}$. The reason for this is the analytical fact that product of L^p functions is generally not in L^p for $p > 2$, but products of $W^{1,p}$ functions remain in $W^{1,p}$. In particular in Equation 3.19 we took one more s derivative than usual due to translations of terms, and if we used $W^{1,p}$ spaces we would have ended up with products of L^p functions.

Another possibility is working with the Morrey spaces in Section 5.5 of [41], where all products are allowed and the space has an L^2-type inner product. In fact this is the approach taken in the Appendix of [12], and if we did this we might be able to avoid the awkward exponential factors of $2/p$ that appear in our subsequent exponential decay estimates.

If $\dot{\Sigma}$ is stable, u is somewhere injective, not a trivial cylinder, and of positive index, then for generic J, the operator D_J is surjective and its index is equal to the dimension of the moduli space $\mathcal{M}$, which is given by $Ind(u)$.

If $\dot{\Sigma}$ is not stable and u is not a trivial cylinder, u still lives in a moduli space of dimension calculated by the index, after we quotient out by automorphisms of the domain. As an analytic matter we address this by adding some marked points to make the domain stable and make the appropriate modifications to Sobolev spaces, in the following convention:

Convention 3.4.3 (Stabilization of Domain)**.** *Given a cascade $u^{\lightning}$, each of the u^i may have components that are unstable, i.e. holomorphic curves whose domain are cylinders or planes. A main source of example is trivial cylinders. Since in this paper we are gluing curves as opposed holomorphic submanifolds, we stabilize these domains following Section 5 of [55] (see also [11] Section 4). For each J-holomorphic curve whose domain is a cylinder, we first fix a surface Σ that intersects the J-holomorphic curve transversely at one point. We endow the J-holomorphic curve with an additional marked point on its domain and require this marked point passes through Σ. For a J-holomorphic curve whose domain is a plane, we fix two disjoint surfaces Σ_1, Σ_2, each of which intersects the J-holomorphic curve transversely at a single point. We add two marked points p_1, p_2 to the domain and require the J-holomorphic curve maps them to Σ_1 and Σ_2 respectively.*

The effect of this is that we eliminate the reparametrization symmetry of the domain. This makes (subsequent) uniqueness statements unambiguous. We note here during the gluing construction, we will be performing large scale symplectization direction translations of each of the u^i. We translate the surfaces Σ_i along with u^i in these large scale symplectization *direction translations. We shall make no further remark on this point and henceforth assume all u^i have domains that are stable.*

For trivial cylinders there is a tad more to be said. If both ends are free then the moduli space is transversely cut out of index 1, where the one dimension of freedom is moving the the trivial cylinder along the Morse-Bott torus. With one end fixed the other free the moduli space is still transversely cut out of index zero. However with both ends fixed the D_J operator is of index -1, yet obviously such trivial cylinders still exist. In this discussion we will only talk about trivial cylinders with at most one fixed ends.

We now come to the definition of what we call *transverse and rigid* cascade. It is these cascades that we will eventually glue.

Definition 3.4.4. *Suppose $u^{\lightning} = \{u^1, .., u^n\}$ is a height 1 cascade that satisfies the following properties:*

a. *All curves u^i are somewhere injective, except trivial cylinders, which can be unbranched covers.*

b. *The[3] numbers T_i satisfy $T_i \in (0, \infty)$.*

c. *Given u^i and $i > 1$, the $s \to \infty$ ends of u^i approach distinct Reeb orbits. For u^i, and $i < n$, the $s \to -\infty$ ends of u^i approach distinct Reeb orbits.*

d. *No end of u^i land on critical points of f, with the following exceptions:*

   i) *If a positive end of u^i lands on a Reeb orbit corresponding to a critical point of f in the intermediate cascade levels, it must then be the minimum of f. Suppose this orbit is γ. Furthermore, for all $j < i$, u^j has a trivial cylinder (potentially unbranched cover) asymptotic to γ and no other curves asymptotic to γ.*

   ii) *If a negative end of u^i is asymptotic to a Reeb orbit corresponding to a critical point of f in the intermediate cascade levels, it must be the maximum of f. Call this orbit γ. For all $j > i$, u^j has a trivial cylinder (potentially unbranched cover) asymptotic to γ and no other curves asymptotic to γ.*

   We call the chain of trivial cylinders all at a critical point of f ***a chain of fixed trivial cylinders****.*

e. *We remove all chains of fixed trivial cylinders from $u^{\lightning}$. For the remaining curves, if an end of u^i lands on a critical point of f, we designate it as fixed, and if an end of u^i avoids critical points of f, we designate it as free. Then each u^i can be thought of as living in a moduli space of J-holomorphic curves. If the domain of u^i is written as $\dot{\Sigma}$, this is the moduli space of J-holomorphic maps from $(\dot{\Sigma}, j) \to (Y \times \mathbb{R}, J)$ where we*

[3]We only consider $T_i > 0$, the case of $T_i = 0$ requires different transversality assumptions and is handled by standard gluing methods.

allow the complex structure to vary, but impose the same fixed/free conditions on the punctures as u^i. We denote this moduli space by $\mathcal{M}(u^i)$. Then $\mathcal{M}(u^i)$ is transversely cut out, and its dimension is given by $Ind(u^i)$.

f. *Again we assume we have removed all chains of fixed trivial cylinders and $u^{\lightning}$ satisfy all of the previous conditions. Each $\mathcal{M}(u^i)$ comes with two evaluation maps, $ev_i^{\pm} : \mathcal{M}(u^i) \to (S^1)^{k_i^{\pm}}$ where $k_i^{\pm}$ refers to how many Reeb orbits are hit by free ends of u^i at $\pm\infty$. Note $k_i^- = k_{i+1}^+$. The evaluation map simply outputs the location of the Reeb orbit that an end of u^i is asymptotic to on the Morse-Bott torus. If we let $u'^i \in \mathcal{M}(u^i)$, and the map*

$$EV^+ : \mathcal{M}(u^2) \times \mathbb{R} \times \mathcal{M}(u^3) \times \mathbb{R} \times \ldots \times \mathcal{M}(u^n) \times \mathbb{R} \tag{3.3}$$

$$\longrightarrow (S^1)^{k_2^+} \times (S^1)^{k_3^+} \times \ldots \times (S^1)^{k_n^+} \tag{3.4}$$

given by

$$(u'^2, T'_1, \ldots, u'^n, T'_{n-1}) \longrightarrow \left(\phi_f^{T'_1}(ev_2^+(u'^2)), \phi_f^{T'_2}(ev_3^+(u'^3)), \ldots, \phi_f^{T'_{n-1}}(ev_n^+(u'^n))\right) \tag{3.5}$$

and the map

$$EV^- : \mathcal{M}(u^1) \times \mathcal{M}(u^2) ... \times \mathcal{M}(u^{n-1}) \longrightarrow (S^1)^{k_1^-} \times (S^1)^{k_2^-} \times ... \times (S^1)^{k_{n-1}^-} \tag{3.6}$$

given by

$$(u'^1, ..., u'^n) \longrightarrow \left(ev_2^-(u'^1), ..., ev_{n-1}^-(u'^{n-1})\right) \tag{3.7}$$

are transverse at $u^{\lightning}$, then we say the cascade $u^{\lightning}$ is transversely cut out.

g. *In particular if $u^{\lightning}$ is transversely cut out, it lives in a moduli space that is a manifold. The manifold has dimension given by the following formula. Assuming again we have removed all chains of fixed trivial cylinders, then the dimension is given by*

$$Ind(u^{\lightning}) := Ind(u^1) + \ldots + Ind(u^n) - k_1^- - \cdots - k_{n-1}^- + (n-1) - n$$

If $u^{\lightning}$ is transversely cut out and $Ind(u^{\lightning}) = 0$, then we say $u^{\lightning}$ is rigid. Note here we have quotiented by the $\mathbb{R}$ action on each level.

h. *(Asymptotic matchings).* [4]

[4] We describe the analogue of this construction in the nondegenerate case. Suppose u and v are both nontrivial somewhere injective transversely cut out rigid holomorphic cylinders in $\mathbb{R} \times Y^3$, and the negative end of u approaches an embedded (nondegenerate) Reeb orbit γ with multiplicity n, and the positive end of v also approaches γ with multiplicity n. Then there are n distinct ways to glue u and v together, and we use asymptotic markers to keep track of this. This is explained in Lemma 4.3 of [36]. Our definition of matching is an analogue of this phenomenon for cascades, except we fit a gradient trajectory (or several segments of gradient trajectories connected to each other by trivial cylinders) between two non-trivial curves. Our notation for asymptotic markers is taken from Section 1 of [40].

*Suppose C_i is a nontrivial curve that appears on the ith level, i.e. it is a component of u^i, and a negative end of C_i is asymptotic to a Reeb orbit γ. Suppose γ is the m times multiple cover of an embedded Reeb orbit, γ'. Consider the preimage of a point w in γ' of the covering map from γ to γ', which is a set of multiplicity m that we write as $\{1, ..., m\}$. Consider the smallest $j > i$ such that u^j contains a nontrivial curve C_j that has a positive end that is asymptotic to a Reeb orbit $\tilde{\gamma}$ that is connected to γ by the upwards gradient flow segments in $u^{\sharp}$. (If $j > i+1$ we allow several segments of gradient flow concatenated together with trivial cylinders in the middle[5].) The orbit $\tilde{\gamma}$ covers some embedded Reeb orbit $\tilde{\gamma}'$ with multiplicity m. Let $\tilde{w}$ denote the point in $\tilde{\gamma}$ that corresponds to w under the gradient flow, i.e. in the neighborhoods we have chosen for Morse-Bott tori, w and $\tilde{w}$ have the same z coordinate. Then the preimage of $\tilde{y}$ under the covering map $\tilde{\gamma} \to \tilde{\gamma}'$ is a set with m elements, $\{\tilde{1}, .., \tilde{m}\}$. A **matching**, which is part of the data of the cascade $u^{\sharp}$, is a choice of a function from $\{1\}$ to $\{\tilde{1}, .., \tilde{m}\}$.*

With the above conditions all satisfied, we say $u^{\sharp}$ is a transverse and rigid height 1 cascade.

Now we are in a position to state our correspondence theorem:

Theorem 3.4.5. *For all $\delta > 0$ sufficiently small, a rigid transverse cascade $u^{\sharp}$ can be uniquely glued to a J_δ-holomorphic curve u_δ with non-degenerate ends.*

The free ends in the u^1 level correspond to ends of u_δ that are asymptotic to Reeb orbits corresponding to the maximum of f. The fixed positive ends of u^1 correspond to positive ends of u_δ that are asymptotic to the Reeb orbits at the minimum of f. Similarly the free negative ends of u^n correspond to negative ends of u_δ that are asymptotic to the Reeb orbits at the minimum of f, and the fixed negative ends of u^n correspond to negative ends asymptotic to Reeb orbits corresponding to the maximum of f. The curve u_δ also has Fredholm index 1.

By uniqueness we mean that if $\{\delta_n\}$ is a sequence of numbers that converge to zero as $n \to \infty$, and u'_{δ_n} is a sequence of J_{δ_n}-holomorphic curves converging to $u^{\sharp}$, then for large enough n, the curves u'_{δ_n} agree with u_{δ_n} up to translation in the symplectization direction.

Remark 3.4.6. In Section 3.5 we will see that in perturbing from λ to λ_δ, the almost complex structure J will need to be perturbed to J_δ to ensure it is compatible with λ_δ. We will specify how to perturb J into J_δ near each Morse Bott torus in Section 3.5. We can in fact perturb J_δ to be different from J away from the Morse-Bott tori as well. Our construction works as long as in C^∞ norm the difference between J and J_δ is bounded above by $C\delta$. We bring this up because we can choose a generic path between J and J_δ as $\delta \to 0$ so that for generic $\delta > 0$, the glued curve u_δ is also transversely cut out. This will be useful for Floer theory constructions in Chapter 2 [66], and will be explained in more detail there.

[5]As explained in the proof of Theorem 3.12.6 in the Appendix, we should really think of collections of trivial cylinders connected by finite gradient flow segments between them as being a *single* gradient flow segment that flows across different cascade levels.

3.5 Differential geometry

In this section we work out the differential geometry surrounding the Morse-Bott tori. We first work out the Reeb dynamics, then we show two gradient flow trajectories of f correspond to J_δ-holomorphic cylinders.

Reeb dynamics

We recall the local neighborhoood near a Morse-Bott torus: if (Y^3, λ) is a contact 3 manifold with Morse-Bott degenerate contact form, near a Morse-Bott torus we have coordinates $(z, x, y) \in S^1 \times S^1 \times \mathbb{R}$. Let

$$\lambda_0 = dz - ydx$$

denote the standard contact form, then by Theorem 3.3.2 λ looks like

$$\lambda = h(x, y, z)\lambda_0$$

where $h(x, y, z)$ satisfies

$$h(x, 0, z) = 0, \quad dh(x, 0, z) = 0.$$

Next we perturb the contact form to

$$\lambda \longrightarrow \lambda_\delta := e^{\delta g f}\lambda$$

We assume we are working in a small enough neighborhood so that $g = 1$. We are interested in the Reeb dynamics on the torus $y = 0$.

Proposition 3.5.1. *On the torus $y = 0$, let R_δ denote the Reeb vector field of λ_δ. We write it in the form $R_\delta = R + X$ where $R = \partial/\partial_z$ is the Reeb vector field of λ. Then the following equaitons are satisfied*

$$\iota_X\lambda = \frac{1 - e^{\delta f}}{e^{\delta f}}$$

$$\iota_X d\lambda = \frac{de^{\delta f}}{e^{2\delta f}}$$

and these two equations completely characterize the the behaviour of R_δ on the $y = 0$ surface.

Proof. From definition

$$\iota_{R_\delta}\lambda_\delta = 1$$

hence we have

$$\iota_X\lambda = \frac{1 - e^{\delta f}}{e^{\delta f}}.$$

For the second equation,

$$\begin{aligned}
&\iota_{R+X} d\left(e^{\delta f}\lambda\right) \\
&= \iota_{R+X}\left(e^{\delta f} d\lambda + \delta e^{\delta f} df \wedge \lambda\right) \\
&= \iota_R(...) + \iota_X(...)
\end{aligned}$$

If we look at the first term we see

$$\begin{aligned}
&\iota_R(e^{\delta f}d\lambda + \delta e^{\delta f}df \wedge \lambda)\\
&= \iota_R e^{\delta f}d\lambda + \delta e^{\delta f}(\iota_R df) \wedge \lambda - \delta e^{\delta f}df\,\iota_R\lambda\\
&= 0 + 0 - \delta e^{\delta f}df.
\end{aligned}$$

Next we look at the second term

$$\begin{aligned}
&\iota_X d\lambda_\delta\\
&= \iota_X(e^{\delta f}d\lambda + \delta e^{\delta f}df \wedge \lambda)\\
&= \iota_X e^{\delta f}d\lambda + \delta e^{\delta f}(\iota_X df)\lambda - \delta e^{\delta f}df\,\iota_X\lambda\\
&= e^{\delta f}\iota_X d\lambda + \delta e^{\delta\lambda}(\iota_X df)\lambda - \delta e^{\delta f}df\,\iota_X\lambda.
\end{aligned}$$

Combining the above two equations we get

$$e^{\delta f}\iota_X d\lambda + \delta e^{\delta f}(\iota_X df)\lambda - \delta e^{\delta f}df\,\iota_X\lambda = \delta e^{\delta f}df.$$

Evaluate both sides with ι_R we see that

$$\iota_X df = 0$$

so we get

$$e^{\delta f}\iota_X d\lambda = de^{\delta f}\left(1 + \frac{1 - e^{\delta f}}{e^{\delta f}}\right)$$

$$\iota_X d\lambda = de^{\delta f}/e^{2\delta f}.$$

□

In particular on the $y = 0$ surface we can write

$$X = \frac{1 - e^{\delta f}}{e^{\delta f}}\partial_z - \frac{\delta e^{\delta f}f'(x)\partial_y}{e^{2\delta f}}$$

Almost complex structures and gradient flow lines

For $(\mathbb{R} \times Y^3, \lambda)$ we choose a generic almost complex structure J that is standard on the surface of the Morse-Bott torus, i.e. $J\partial_x = \partial_y$. After we perturb to λ_δ, we must perturb J to J_δ to make the complex structure compatible with the new contact form. However we keep the same complex structure on the contact distribution, i.e.

$$J_\delta\partial_x = \partial_y.$$

We wish to understand what J_δ does to the Reeb vector field and the vector field in the symplectization direction. By definition

$$J_\delta(R+X) = -\partial_a$$

$$J_\delta \partial_a = R + X.$$

From the above we deduce

$$J_\delta R = -\partial_a - J_\delta X = -\partial_a - J_\delta \left(\frac{1-e^{\delta f}}{e^{\delta f}} \partial_z - \frac{\delta e^{\delta f} f'(x) \partial_y}{e^{2\delta f}} \right)$$

$$J_\delta \partial_z = e^{\delta f} \left(-\partial_a - \frac{\delta e^{\delta f} f'(x) \partial_x}{e^{2\delta f}} \right) = -e^{\delta f} \partial_a - \delta f'(x) \partial_x.$$

Next we consider J_δ-holomorphic curves constructed by lifting gradient flows of δf. Consider maps

$$v : (s,t) \longrightarrow (a(s), z(t), x(s), y(s)) \in \mathbb{R} \times S^1 \times S^1 \times \mathbb{R}$$

defined by

$$\partial_s a = e^{\delta f(x(s))}$$

$$\partial_t z(t) = R$$

$$\partial_s x(s) = +\delta f'(x)$$

$$y = 0$$

and initial conditions

$$a(0,t) = 0, x(0) = \text{constant}.$$

Proposition 3.5.2. *The map v as defined above is a J_δ-holomorphic curve.*

Proof. Let $\hat{v} := (z(t), x(s), y(s))$. We apply the J_δ holomorphic curve equation to v

$$\begin{aligned} &\partial_s \hat{v} + \partial_s a \frac{\partial}{\partial a} + J \partial_t \hat{v} \\ &= e^{\delta f(x(s))} \frac{\partial}{\partial a} + \delta f'(x) \frac{\partial}{\partial x} - (e^{\delta f}) \partial_a - \delta f'(x) \partial_x = 0. \end{aligned}$$

□

We observe since there are two gradient flow lines on S^1, there are two J_δ-holomorphic curves as above corresponding to their lifts. Further:

Proposition 3.5.3. *The curve v is transversely cut out. The same is true for unbranched covers of v by cylinders.*

Proof. We use Theorem 1 from Wendl's paper on automatic transversality [65]. In the language of Theorem 1, $\operatorname{Ind}(v) = 1$, $\Gamma_0 = 1$ (only one end is asymptotic to Reeb orbits with even Conley-Zehnder index), there are no boundary components, and $c_N = 0$, hence

$$\operatorname{Ind}(v) = 1 > c_N + Z(du) = 0.$$

The same proof works for unbranched covers of v as well. □

For future references, we record the form of the vector field

$$v_*\partial_s = e^{\delta f(x(s))}\partial_a + \delta f'(x)\partial_x.$$

3.6 Linearization of $\bar{\partial}_{J_\delta}$ over v

In this section we define the linearization of the Cauchy Riemann operator $\bar{\partial}_{J_\delta}$ over v, the holomorphic cylinder constructed in the above section that corresponds to a gradient flow of f. We also equip it with an appropriate Sobolev space on which the linearized operator is Fredholm. This is preparation for the gluing construction.

Convention 3.6.1. *For this point onward in the paper we will assume all gradient trajectories are simply covered for ease of notation. In practice they can be (unbranched) multiply covered. For any of the analysis we are doing this will not make any difference.*

The point to note here is that if we see any finite gradient cylinders (or chains of finite gradient cylinders connected to each over by trivial cylinders) that are multiply covered connecting between two non-trivial curves in the cascade, the number of ways to glue is counted precisely by the number of different matchings (see Definition 3.4.4) we can assign to such a segment.

Fix a holomorphic cylinder v_δ (we make the δ dependence explicit), consider the space of vector fields over v_δ,

$$\Gamma(v_\delta^* TM).$$

We take a weighted Sobolev space

$$W^{2,p,d}(v_\delta^* TM)$$

which is the $W^{2,p}(v_\delta^* TM)$ with exponential weight $e^{w(s)} = e^{ds}$, where $d > 0$ is a small fixed number that only depends on the Morse-Bott torus. Here we can also use e^{-ds}.

Note as given, these are vector fields with exponential decay as $s \to \infty$ and exponential growth as $s \to -\infty$. The end with exponential growth is not suited for nonlinear analysis of the Cauchy Riemann equation, but we will find them useful as a formal device so all our linear operators have the right Fredholm index and uniformly bounded right inverse. It will be apparent from our gluing construction that vector fields with exponential growth will not cause any difficulty. This is also the approach taken in [12]. The main result of the section is the following:

Proposition 3.6.2. *Let D_{J_δ} denote the linearization of $\bar{\partial}_{J_\delta}$ along v_δ using metric g. Then the operator*

$$D_{J_\delta} : W^{2,p,d}(v_\delta^* TM) \longrightarrow W^{1,p,d}(v_\delta^* TM)$$

is a Fredholm operator of index 0. In particular it is an isomorphism. Further it has right (and left) inverse Q_δ whose operator norm is uniformly bounded as $\delta \to 0$.

The proof will occupy the rest of this section. The idea is for sufficiently small $\delta > 0$ the J_δ-holomorphic curve v_δ is nearly horizontal, and hence can be approximated by a finite collection of trivial cylinders glued together. But the linearization of $\bar{\partial}$ over a trivial cylinder is an isomorphism with inverse independent of δ, and by standard gluing theory of operators the operator glued from linearizations of $\bar{\partial}$ over trivial cylinders has the properties described in the theorem.

Linearizations over trivial cylinders

Fix x, which corresponds to fixing a Reeb orbit in the Morse-Bott torus. Consider the trivial cylinder C_x at x. The Cauchy Riemann operator $\bar{\partial}_J$ (with unperturbed complex structure J) has linearization D_x of the form

$$\partial_s + J_0 \partial_t + S_x(t).$$

The matrix J_0 is the standard complex structure on $\mathbb{R}^4$, and $S_x(t)$ is a symmetric matrix. Considered as an operator, we have

$$D_x : W^{2,p,d}(C_x^* TM) \longrightarrow W^{1,p,d}(C_x^* TM)$$

with exponential weight e^{ds} on both sides.

Lemma 3.6.3. *D_x is an isomorphism.*

Proof. We consider this operator defined on $W^{2,p}(C_x^* TM)$ instead of $W^{2,p,d}(C_x^* TM)$ by using the isometry

$$e^{-ds} : W^{2,p}(C_x^* TM) \longrightarrow W^{2,p,d}(C_x^* TM).$$

The effect of this on the operator D_x is

$$e^{ds} D_x e^{-ds} : W^{2,p}(C_x^* TM) \longrightarrow W^{1,p}(C_x^* TM)$$

$$e^{ds} D_x e^{-ds} = \partial_s + J_0 \partial_t + S_x(t) - d.$$

The operator $A(t) : W^{2,p}(S^1) \to W^{1,p}(S^1)$ given by $A = -J_0 \partial_t - S_x(t) + d$ has eigenfunctions $\{e_n\}$ with eigenvalues $\{\lambda_n\}$, and no eigenvalue λ_n is equal to zero. This shows D_x is index 0 because there is no spectral flow. An element in the kernel of $e^{ds} D_x e^{-ds}$ can be written in the form

$$\sum c_n e^{\lambda_n s} e_n(t)$$

but all c_n must equal to zero because terms like $e^{\lambda_n s}$ have exponential growth on one end hence cannot live in $W^{2,p}(C_x^*TM)$. This implies D_x is an isomorphism hence has an inverse, which we denote by Q_x. Note this inverse does not depend on δ. □

Observe since x varies in a S^1 family, there exists C such that

$$\|Q_x\| \leq C$$

in operator norm for all $x \in S^1$.

Uniformly bounded inverse for D_{J_δ}

In this subsection we prove the main theorem of this section. This is inspired by analogous constructions in Proposition 4.9 in [7] and Proposition 5.14 in [5].

Proof of Proposition 3.6.2. We identify S^1, the circle of Morse-Bott orbits, with $x \in [0,1]/\sim$, and we recall f has critical points at $x = 0$ and $x = 1/2$. WLOG we consider the $v_\delta(s,t)$ corresponding flow from with $-\infty$ end at $x = 0$, towards $x = 1/2$ as $s \to +\infty$ and take $s_0 = -\infty$.

Fix N large, let $x_i = 1/2N$, $i = 1,..,N$ denote Reeb orbits on the Morse-Bott torus. Let $s_i \in \mathbb{R}$ denote the time it takes for v_δ to flow to x_i, i.e. when x component of $v_\delta(s_i,\cdot) = x_i$. We implicitly take $s_N = +\infty$. We observe s_i implicitly depends on δ and

$$s_{i+1} - s_i \geq C/(\delta N).$$

We let $D_i := \partial_s + J_0\partial_t + S_{x_i}(t)$ denote the linearization of the $\bar{\partial}_J$ operator at a trivial cylinder at x_i. We define the parameter

$$R := \frac{1}{5d}\log\frac{1}{\delta}$$

Let $\beta_o(s)$ be a cut off function equal to 1 for $s \geq 1$ and 0 for $s \leq 0$. we define the "glued" operator

$$\#_N D_i := \partial_s + J_0\partial_t + \sum_{i=0}^{N-1}(1-\beta_o(s-s_{i+1}))\beta_o(s-s_i)S_{x_{i+1}}(t).$$

So we have $\#_N D_i = D_i$ on the interval $[s_{i-1}+1, s_i]$ by construction. Viewed as operators

$$W^{2,p,d}(v_\delta^*TM) \longrightarrow W^{1,p,d}(v_\delta^*TM)$$

we have

$$\|D_{J_\delta} - \#_N D_i\| \leq C(1/N + \delta)$$

in operator norm with constant C independent of δ or N. It follows from the same spectral flow argument as above that $\#_N D_i$ is Fredholm of index 0. We now proceed to construct a uniformly bounded (as $\delta \to 0$) right inverse Q_N for it. Let Q_i denote inverses to D_i, we first construct approximate inverse Q_R using the following commutative diagram

$$\begin{tikzcd} W^{1,p,d}(v_\delta^* TM) \arrow[r, "Q_R"] \arrow[d, "s_R"] & W^{2,p,d}(v_\delta^* TM) \\ \bigoplus_i [W^{1,p,d}(v_\delta^* TM)]_i \arrow[r, "\bigoplus Q_i"] & \bigoplus_i [W^{2,p,d}(v_\delta^* TM)]_i \arrow[u, "g_R"] \end{tikzcd}$$

with splitting maps s_R and gluing maps g_R defined as follows: if $\eta \in W^{1,p,d}(v_\delta^* TM)$, $s_R(\eta) = (\eta_i, .., \eta_N)$ where

$$\eta_i := \eta(1 - \beta_o(s - s_i))\beta_o(s - s_{i-1}).$$

We see immediately s_R has uniformly bounded operator norm as $\delta \to 0$, and that its norm is also bounded above independently of N. Let $\gamma_R(s)$ be a cut off function $\gamma_R(s) = 1$ for $s < 1$ and $\gamma_R(s) = 0$ for $s > R/2$ and $\gamma'(s) \leq C/R$. If $(\xi_1, .., \xi_N) \in \bigoplus_i W^{2,p,d}(v_\delta^* TM)$ we define

$$g_R(\xi_1, .., \xi_N) = \sum_i \xi_i \gamma_R(s - s_i)\gamma_R(s_{i-1} - s).$$

We also see that g_R is an uniformly bounded operator as $\delta \to 0$ and its upper bound on norm is independent of N. We conclude Q_R has uniformly bounded norm as $\delta \to 0$. We next show it is an approximate inverse to $\#_N D_i$.

If we start with $\eta \in W^{1,p,d}(v_\delta^* TM)$, with $Q_R(\eta) = \sum_i \xi_i \gamma_R(s - s_i)\gamma_R(s_{i-1} - s)$. We apply $\#_N D_i$ to it and observe away from the intervals of the form $\bigcup_i [s_i - R, s_i + R]$ - which we think of the region where gluing happens,

$$\#_N D_i Q_R \eta = D_i Q_i \eta_i = \eta$$

so we focus our attention to an interval of the form $[s_i - R, s_i + R]$, in which $Q_R(\eta) = \gamma_R(s - s_i)\xi_i + \gamma_R(s_i - s)\xi_{i+1}$.

We observe over intervals of this form $\|D_i - \#_N D_i\| \leq C/N$ in operator norm, so when we apply $\#_N D_i$ to $Q_R(\eta)$ we get

$$\begin{aligned} \#_N D_i Q_R \eta =& \#_N D_i(\gamma_R(s - s_i)\xi_i + \gamma_R(s_i - s)\xi_{i+1}) \\ =& \gamma_R'(s - s_i)\xi_i - \gamma_R'(s_i - s)\xi_{i+1} \\ & + \gamma_R(s - s_i)\#_N D_i \xi_i + \gamma_R(s_i - s)\#_N D_i \xi_{i+1}. \end{aligned}$$

In light of the above, in this region we have

$$\begin{aligned} \#_N D_i \xi_i &= D_i \xi_i + (\#_N D_i - D_i)\xi_i \\ &= \beta_o(s - s_i)\eta + (\#_N D_i - D_i)\xi_i \end{aligned}$$

with

$$\|(\#_N D_i - D_i)\xi_i\| \leq C/N\|\xi_i\| \leq C/N\|\eta\|$$

and likewise for the ξ_{i+1} term in weighted Sobolev norm. We also note $\gamma_R' \leq C/R$, so we can write

$$\#_N D_i Q_R \eta = \gamma_R(s-s_i)\beta_o(s-s_i)\eta + \gamma_R(s_i-s)(1-\beta_o(s-s_i))\eta + \text{error}$$

for $s \in [s_i - R, s_i + R]$. But by the construction of β_o and γ_R, we have $\gamma_R(s-s_i)\beta_o(s-s_i)\eta + \gamma_R(s_i-s)(1-\beta_o(s-s_i))\eta = \eta$ in $[s_i - R, s_i + R]$, so we have

$$\|\#_N D_i Q_R \eta - \eta\| \leq C/N \|\eta\|$$

in weighted Sobolev norm. So for sufficiently large values of N, the operator Q_R is an approximate right inverse. Then we can define a true right inverse $\#_N D_i$ by

$$Q_N := Q_R(\#_N D_i Q_R)^{-1}$$

which also has uniformly bounded norm as $\delta \to 0$. This in particular implies $\#_N D_i$ is surjective.

Finally using

$$\|D_{J_\delta} - \#_N D_i\| \leq C(1/N + \delta)$$

in operator norm we see that Q_N is an approximate right inverse to D_{J_δ} because:

$$\begin{aligned}
\|D_{J_\delta} Q_N \eta - \eta\| &= \|(D_{J_\delta} - \#_N D_i)Q_N + \#_N D_i Q_N \eta - \eta\| \\
&= \|(D_{J_\delta} - \#_N D_i)Q_N \eta\| \\
&\leq \|Q_N\| \cdot \|(D_{J_\delta} - \#_N D_i)\| \cdot \|\eta\| \\
&\leq C/N \|\eta\|
\end{aligned}$$

and hence D_{J_δ} has uniformly bounded right inverse as $\delta \to 0$. □

Remark 3.6.4. We proved for given D_{J_δ} acting on $W^{2,p,d}(v_\delta^* TM)$ over fixed v_δ it has uniformly bounded right inverse. For our proof we assumed the exponential weight is of the form e^{ds}, but it should be apparent from our proof even as we translate the weight profile from e^{ds} to e^{ds-T} for any $T \in \mathbb{R}$, the same proof goes through. Said another way, for any sufficiently small δ and any T, the operator D_{J_δ} defined over $W^{2,p,d}(v_\delta^* TM)$ with weight $e^{\pm ds + T}$ has a uniformly bounded inverse.

Remark 3.6.5. In the above construction we implicitly fixed a parametrization of v_δ with respect to the t variable, i.e. we picked out which point on the Reeb orbit corresponds to $t = 0$. We could also have changed this, resulting in a reparametrization of v_δ, of the form $t \to t + c$. For all such reparametrizations it is obvious D_{J_δ} continues to have uniformly bounded right inverse, and this upper bound is uniform across all possible reparametrizations in the t variable.

3.7 Gluing a semi-infinite gradient trajectory to a holomorphic curve

In this section we glue a J-holomorphic curve u to a semi-infinite gradient trajectory v. This is a simpler case of gluing for multi-level cascades, and properties of this gluing developed here and in the following sections will be used extensively in gluing together multiple level cascades. The novel feature of this gluing construction, which separates it from standard types of gluing constructions, is that we will make the pregluing dependent on asymptotic vectors. The general setup will follow that of Section 5 in [41], and in a sense we are doing obstruction bundle gluing, see also Remark 3.9.23. This approach to gluing has appeared in the Appendix of [12].

The section is organized as follows: in subsection 1 we first introduce the gluing setup. In subsection 2 we do the pregluing. In subsection 3 we take care of the differential geometry/estimates needed to deform the pregluing. Further, we write down the J-holomorphic curve equation we need to solve, and split it into two different equations as was done in Section 5 of [41]. And finally in subsection 4 we solve both of these equations. We do not yet say anything about surjectivity of gluing and save it for the end when we discuss surjectivity of gluing in the general case.

Gluing setup

Let $u : \dot{\Sigma} \to M$ be a J-holomorphic curve with only one positive puncture which is free, asymptotic to a Morse-Bott torus with multiplicity 1 (higher multiplicities are handled similarly). We choose local coordinates on u around the puncture given by $(s,t) \in [0,\infty) \times S^1$. We also assume $\dot{\Sigma}$ is stable. Our assumptions are purely a matter of convenience since it will be apparent from our construction how to glue semi-infinite gradient trajectories with arbitrary number of positive/negative ends. We also assume (purely as a matter of notational convenience) that we have shifted our coordinates so that $\lim_{s\to\infty} u(s,t)$ converges to the Reeb orbit at $x = 0$, and the critical points of f are at $x = \pm 1/4$ with max at $x = 1/4$ and min at $x = -1/4$. We assume u is rigid, i.e. the operator

$$D_J : W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J} \longrightarrow W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u^*TM))$$

is surjective of index 1. It has a right inverse Q_u. Here $V_\Gamma := span\{\beta_{1;0,\infty}\partial_z, \beta_{1;0,\infty}\partial_a, \beta_{1;0,\infty}\partial_x\}$. This is a 3 dimensional vector space with a given basis, we denote elements of this space by triples (r,a,p). The norm of elements $(r,a,p) \in V_\Gamma$ is simply $|r| + |a| + |p|$. We will often write $|(r,a,p)|$ to mean this norm.

Convention 3.7.1. *We will generally use the symbol (r,a,p) as a shorthand for the asymptotically constant vector field*

$$r\beta_{1;0,\infty}\partial_z + a\beta_{1;0,\infty}\partial_a + p\beta_{1;0,\infty}\partial_x.$$

This is generally the case when we use (r, a, p) to deform curves, and the case later where the symbol (r, a, p) appears in the equations $\Theta_\pm$. We will also sometimes to use the symbol (r, a, p) to simply denote the tuple of numbers, r, a, p. It will be clear from context what we mean.

We observe by definition $D_J(r, a, p)$ decays exponentially (at a rate faster than e^{-ds}, which we denote by e^{-Ds}, $D >> d$) as $s \to \infty$.

Convention 3.7.2. *We use the following convention regarding d and D. The symbol D, when written as e^{-Ds} will always be used to denote a rate of exponential decay that only depends on the background geometry, say the local geometry around the Morse-Bott torus. An example will be the rate of exponential convergence to a trivial cylinder of a J-holomorphic curve asymptotic to a Reeb orbit. The lower case d will be chosen to be independent of δ, $d << D$ and as usual much smaller than the distance between the nonzero eigenvalues of operator $A(t)$ and 0. This is the exponential weight we will use in our weighted Sobolev spaces.*

The rest of the section is devoted to proving the following:

Proposition 3.7.3. *For every $\delta > 0$ sufficiently small, there is a J_δ-holomorphic curve $u_\delta : \dot{\Sigma} \to M$ that is positively asymptotic to the Reeb orbit $x = 1/4$ obtained by gluing a semi-infinite gradient trajectory along the Morse-Bott torus to u.*

Pregluing

We make the pregluing dependent on the triple of asymptotic vectors (r, a, p). We first describe the neighborhood of $u|_{[0,\infty)\times S^1}$. Recall we are working in a neighborhood of the Morse-Bott torus whose local coordinates in the symplectization are given by

$$\mathbb{R} \times S^1 \times S^1 \times \mathbb{R} \ni (a, z, x, y)$$

where x is displacement across Morse-Bott torus direction, y is the vertical direction, a symplectization direction, and z Reeb direction. At the surface of $y = 0$, J is the standard complex structure. The metric here is the flat metric, so we will simply "add" vectors together as opposed to taking the exponential map. The map u comes in the form

$$u(s, t) = (s + \epsilon_s, t + \epsilon_t, \eta_x, \eta_y)$$

where

$$\lim_{s \longrightarrow \infty} \eta_* = 0$$

of order e^{-Ds}, where D is some fixed constant specific to Morse-Bott torus ($d << D$). We also have

$$\epsilon_* \approx O(e^{-Ds}).$$

Then

$$u(s,t) + (r,a,p) = (s + \epsilon_s + a\beta_{1;0,\infty}, t + \epsilon_t + r\beta_{1;0,\infty}, \eta_x + \beta_{1;0,\infty} p, \eta_y).$$

Recalling the important parameter R:

$$R = \frac{1}{5d} \log(1/\delta)$$

which we will take to be our gluing parameter, we cut off $u + (r,a,p)$ at $s = R$ and glue in a gradient trajectory $v_{r,a,p}(s,t)$ satisfying

$$v_{r,a,p}(R,t) = (R + a, t + r, p, 0).$$

We observe that since $\delta << R$, in the range of $s \in [R, 5R]$, the map $v_{r,a,p}(s,t)$ remains almost a trivial cylinder, which can make precise by noting

$$|v_{r,a,p} - (R + s, t, p, 0)|_{C^k} \le CR\delta, s \in [R, 5R].$$

We are now ready to define the pregluing. We define

$$u_{r,a,p}(s,t) := \begin{cases} u(s,t) + (r,a,p), s < R - 1 \\ v_{r,a,p}, s \ge R - 1/2 \\ \text{smooth, bounded interpolation between } u + (r,a,p) \text{ and } v_{r,a,p} \text{ for} \\ s \in [R-1, R-1/2]. \end{cases}$$

The interpolation above should be chosen so that the difference between $u_{r,a,p}$ and the trivial cylinder of the form $(s,t) \to (R + a, t + r, p, 0)$ should be bounded by e^{-DR} in C^k norm.

We first observe the preglued curve is still defined on the same domain $\dot{\Sigma}$. It still has the same coordinate neighborhood $[0, \infty) \times S^1$ near the unique positive puncture. As a warm up to considering the deformations of this preglued curve, we next measure how non-holomorphic this preglued curve is by applying $\bar{\partial}_{J_\delta}$ to it.

Remark 3.7.4. Here in constructing the domain for the pregluing we "rotated" our gradient trajectory $v_{r,a,p}$ (denoted by v in Section 3.6) by r to match $u + (r,a,p)$. It is also possible to instead glue $u_{r,a,p}$ with $v_{0,a,p}$ by making the identification $t \sim t + r$ at $s = R$. In this case we get back the same surface, however when we later glue over finite cylinders this will make a difference, as it corresponds to the same topological surface but a new complex structure on the preglued domain.

Convention 3.7.5. *We adopt the convention that for terms that are supposed to be small, e.g. uniformly bounded by $C\epsilon$ (in say C^k norms or any norm we care about), we just write the upper bound $C\epsilon$ instead of the specific term in its entirety.*

Proposition 3.7.6. *After we apply the $\bar{\partial}_{J_\delta}$ operator to the preglued curve $u_{r,a,p}$ over the interval $(s,t) \in [0,\infty) \times S^1$ we get terms of the form*

$$\begin{aligned}
&[D_J(r,a,p) + C(s,t)|(r,a,p)|^2 e^{-Ds}]\beta_{[0,R+1;1]} \\
&+ C[\delta(1+(r,a,p))]\beta_{[0,R+1;1]} \\
&+ C[e^{-DR}(1+|(r,a,p)|) + C\delta(1+|(r,a,p)|)]\beta_{[1;R-2,R+2;1]}.
\end{aligned}$$

By $C(s,t)$ or oftentimes C, we mean a function of (s,t) and occasionally also including the variables (r,a,p), whose derivatives are uniformly bounded. When we write $|r,a,p|$ we mean the absolute value of the numbers $|r|, |a|, |p|$.

Note the term $D_J(r,a,p)$, which is the only term in this expression that is not "small". Figuratively we can write this as

$$D_J(r,a,p) + \mathcal{F}(r,a,p) + \mathcal{E}(r,a,p)$$

where

$$\mathcal{F}(r,a,p) = C(s,t)|(r,a,p)|^2 e^{-Ds}$$

and

$$\begin{aligned}
\mathcal{E}(r,a,p) =& C\,[\delta(1+(r,a,p))]\,\beta_{[0,R+1;1]} \\
&+ C\left[e^{-DR}(1+(r,a,p)) + C\delta(1+(r,a,p))\right]\beta_{[1;R-2,R+2;1]}
\end{aligned}$$

where we think of $\mathcal{F}(r,a,p)$ as a quadratic order term and $\mathcal{E}(r,a,p)$ as an error term.

Remark 3.7.7. We first note that u is holomorphic with respect to J, but in the above theorem we applied the $\bar{\partial}$ operator with respect to J_δ, which is responsible for the appearance of several error terms. Further since u is not holomorphic with respect to J_δ, there is another error term that appears in the interior of u, i.e. $\dot{\Sigma} \setminus [0,\infty) \times S^1$ of size $C\delta$. Note no such error term appears in the interior region of $v_{r,a,p}$ This term is not very important because by our metric it is (uniformly) small, we will include it when we solve for the equation more globally.

Proof. We first consider downwards of the pregluing region, in the region $s \in [0, R-1]$, the pregluing is simply consider $u + (r,a,p)$, then after applying $\bar{\partial}_J$ we get

$$\partial_s[u+(r,a,p)] + J(u+(r,a,p))\partial_t(u+(r,a,p)) = \beta'_{1;0,\infty}(r\partial_z + a\partial_a + p\partial_x) + \partial_s u + J(u+(r,a,p))\partial_t u.$$

To this end, observe $\partial_s u + J(u)\partial_t u = 0$ so we get an expansion of the form $D_J(r,a,p) + \sum C|(r,a,p)|^{n\geq 2}\partial^n_{r,a,p}J(u)\partial_t u$. This is a C^0 bound, we will need a better bound since eventually the size of the vector will be measured with respect to weighted Sobolev norms. Observe $\partial_t u$ is of the form

$$(\partial_t\epsilon_s, 1 + \partial_t\epsilon_t, \partial_t\eta_x, \partial_t\eta_y).$$

All η_* terms decay like e^{-Ds}, except $(0,1,0,0)$. But we observe by compatibility of J, the term $(\partial^n_{r,a,p}J(u(\infty,t)))(0,1,0,0) = 0$. Hence overall the second term $C|(r,a,p)|^{n\geq 2}\partial^n_{r,a,p}J(u)\partial_t u$ is of the form

$$C|r,a,p|^2 e^{-Ds}.$$

Next let's include the effect of J_δ, now we have

$$(\partial_{J_\delta} - \partial_J)(u + (r,a,p)) = (J_\delta(u + (r,a,p)) - J(u + (r,a,p))\partial_t(u + (r,a,p)).$$

This term has size δC and it only exists for length $s \in [0,R]$ and disappears after the pregluing region. We clarify its dependence on various variables: it is of the form

$$C\delta(1 + |r,a,p|)\beta_{[0,R+1;1]}$$

and this is everything in the region $s \in [0,R-1]$. We observe by definition $u_{r,a,p}$ is J_δ-holomorphic in the region $s > R$ so we only need to look at the pregluing region to find rest of the pregluing error. It follows from the uniform boundedness of our interpolation construction in the pregluing that this error is of the form $C[e^{-DR}(1 + |r,a,p|) + C\delta(1 + |r,a,p|)]\beta_{[1;R-2,R+2;1]}$, whence we complete our proof. □

Remark 3.7.8. The reason we are painstakingly computing all of these terms carefully (and in our subsequent computations) is because later we will be differentiating this entire expression with respect to (r,a,p) so we must take note how our expressions depend on these asymptotic vectors.

Deforming the pregluing

Now that we have constructed $u_{r,a,p}$, we deform it to try to make it J_δ-holomorphic. We recall a neighborhood of u is given by: $W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J}$. We recall for $[0,\infty) \times S^1$ there is an exponential weight e^{ds}. We already explained how to construct the pregluing with asymptotic vector fields (r,a,p). We fix $\psi \in W^{2,p,d}(u^*TM)$, $\delta j \in T\mathcal{J}$. Recall deformations of complex structure of the domain δj is away from the cylindrical neighborhood so does not affect our gluing construction for the most part, so unless it is explicitly needed for rest of this section we will drop it from our notation. Now for $v_{r,a,p}$ fix $\phi \in W^{2,p,w}(v^*_{r,a,p}TM)$. Note this choice of Sobolev space is dependent on the asymptotic vectors (r,a,p). We equip the space $W^{2,d,w}(v^*_{r,a,p}TM)$ with weighted Sobolev norm $e^{w(s)} = e^{ds}$.

We fix cut off functions

$$\beta_u := \beta_{[-\infty,2R;R/2]}$$

and

$$\beta_v := \beta_{[R/2;R,+\infty]}.$$

We deform the pregluing $u_{r,a,p}$ via

$$(u_{r,a,p}, j_0) \longrightarrow (u_{r,a,p} + \beta_u\psi + \beta_v\phi, j_0 + \delta j). \tag{3.8}$$

The next proposition describes what happens to the deformed curve when we apply $\bar{\partial}_{J_\delta}$ to it.

Proposition 3.7.9. *The deformed curve $(u_{r,a,p} + \beta_u \psi + \beta_v \phi, j_0 + \delta j)$ is J_δ-holomorphic if and only if the equation*

$$\beta_u \Theta_u + \beta_v \Theta_v = 0$$

is satisfied. Θ_u and Θ_v are equations depending on $\psi_u, \psi_v, \delta j$, and they take the following form

$$\Theta_u = D_J \psi + \beta_v' \phi + \mathcal{F}_u(\psi, \phi) + \mathcal{E}_u(\psi, \phi)$$

and

$$\Theta_v = \beta_u' \psi + D_{J_\delta} \phi + \mathcal{F}_v(\phi, \psi).$$

The forms of functionals $\mathcal{F}_*, \mathcal{E}_*$ are given in the course of the proof.

Remark 3.7.10. We will write the equation Θ_u and Θ_v in two different forms, one form will make it easy to apply elliptic regularity, the other makes it easy to use the contraction mapping principle. It will be later crucial for us to use elliptic regularity, as when we do finite trajectory gluing we will lose one derivative by lengthening/shortening the domain of the neck, and we will use elliptic regularity to gain one derivative to make up for this. The key ingredient is to arrange things so that Θ_v does not contain derivatives of ψ, and Θ_u does not contain derivatives of ϕ. We shall see that this requires some careful differential geometry to achieve.

Proof. **Step 0** We first prepare to write our equation in a way that makes apparent the elliptic regularity in the equation, then we will linearize everything to make linear operators appear. We first consider

$$\bar{\partial}_{J_\delta}(u_{r,a,p} + \beta_u \psi + \beta_v \phi)$$

in the region $s > R$. We recall over in this region $u_{r,a,p} = v_{r,a,p}$. Let's use u_* to denote $u_{r,a,p}$ for short. Then we are looking at the equation

$$\partial_s(u_* + \beta_u \psi + \beta_v \phi) + J_\delta(u_* + \beta_u \psi + \beta_v \phi)\partial_t(u_* + \beta_u \psi + \beta_v \phi) = 0.$$

We rewrite this in the following fashion

$$\begin{aligned}
&\partial_s u_* + \beta_u' \psi + \beta_v' \phi + \beta_u \partial_s \psi + \beta_v \partial_s \phi + J_\delta(u_* + \beta_u \psi + \beta_v \phi)\partial_t(u_* + \beta_u \psi + \beta_v \phi) \\
=&\beta_v(\partial_s \phi + J_\delta(u_* + \beta_u \psi + \beta_v \phi)\partial_t \phi + \beta_u' \psi) + \beta_v' \phi + \partial_s u_* + J_\delta(u_* + \beta_u \psi + \beta_v \phi)\partial_t u_* \\
&+ \beta_u(\partial_s \psi + J_\delta(u_* + \beta_u \psi + \beta_v \phi)\partial_t \psi) \\
=&0.
\end{aligned}$$

Recalling that $\partial_s u_* + J_\delta(u_*)\partial_t u_* = 0$ in this region, we can write
$\partial_s u + J_\delta(u + \beta_u \psi + \beta_v \phi)\partial_t u$ as

$$\partial_s u_* + J_\delta(u_*)\partial_t u_* + \partial_{\beta_u \psi} J_\delta(u_*)\partial_t u_* + \partial_{\beta_v \phi} J_\delta(u_*)\partial_t u_* + G(\beta_v \psi, \beta_u \phi) = 0$$

where we can further write

$$G(\beta_v \psi, \beta_u \phi) = \beta_v \phi g_v(\beta_u \psi, \beta_v \phi) + \beta_u \psi g_u(\beta_u \psi, \beta_v \phi)$$

where $g_u(x,y)$ and $g_v(x,y)$ are smooth functions so that pointwise

$$|g_*(x,y)| \leq C(|x|+|y|) \tag{3.9}$$

and g_* have uniformly bounded derivatives. If we introduce the modified cutoff functions,

$$\beta_{ug} := \beta_{[1/2;R-1/2,2R;R/2]}$$

$$\beta_{vg} := \beta_{[R;R/2,2R+2;1/2]}.$$

Then we define Θ_v to be

$$\begin{aligned}\Theta_v :=& \partial_s\phi + J_\delta(v_{r,a,p} + \beta_{ug}\psi + \beta_{[1;R-2,\infty]}\beta_v\phi)\partial_t\phi + \partial_\phi J_\delta(v_{r,a,p})\partial_t v_{r,a,p} & (3.10)\\ &+\beta_{[1;R-2,\infty]}\phi g_v(\beta_{ug}\psi, \beta_v\phi) + \beta_u'\psi. & (3.11)\end{aligned}$$

We make a few remarks about the important features of our definition of Θ_v. We first remark by our cut off function $\beta_{[1;R-2,\infty]}$, the equation becomes linear for $s < R-1$, as all the quadratic terms have disappeared. This is desirable as we will be solving Θ_v with $W^{2,p}(v^*TM)$ with exponential weight e^{ds}. Usually having vector fields that grow exponentially is undesirable for doing analysis, but in our case where the vector field grows exponentially the equation is linear, and hence poses no problem for the solution for our equation. The deformed preglued curve also doesn't see the segments of ϕ that grows exponentially by our choice of cut off functions.

We also remark that Θ_v appears in a form that allows us to apply elliptic regularity as stated in Theorem 3.7.11, which we will need much later on.

The definition of Θ_u is slightly more involved. From now on we think of (s,t) as coordinates in the cylindrical ends of u. Let $\overline{u}$ denote the interpolation from $v_{r,a,p}$ to $u+(r,a,p)$ that starts at $s=2R+1$ and finishes the interpolation process at $s=2R+2$. The difference between $v_{r,a,p}$ and $u+(r,a,p)$ in this interval is uniformly bounded in C^k norm by $C(e^{-2DR}+R\delta)$ over $s\in[2R+1,2R+2]$. Note also where β_u is nonzero and $s>R$, $\overline{u}$ agrees with u_*. Let us also consider

$$\partial_s\overline{u} + J_\delta(\overline{u} + \beta_u\psi + \beta_{vg}\phi)\partial_t\overline{u}$$

which we expand as

$$\partial_s\overline{u} + J_\delta(\overline{u})\partial_t\overline{u} + \partial_{\beta_u\psi}J_\delta(\overline{u})\partial_t\overline{u} + \partial_{\beta_{vg}\phi}J_\delta(\overline{u})\partial_t\overline{u} + \overline{G}(\beta_u\psi, \beta_{vg}\phi)$$

where the definition of $\overline{G}$ is analogous to that of G. We recognize that the first term $\partial_s\overline{u} + J_\delta(\overline{u})\partial_t\overline{u}$ is supported for $s\in[2R+1,\infty]$ whenever $s>R+1$ and is of size (in C^k norm) $C(e^{-2DR}+R\delta)$. The term $\overline{G}(\beta_u\psi,\beta_{vg}\phi)$ admits a similar expansion as G above that gives

$$\overline{G}(\beta_u\psi, \beta_{vg}\phi) = \beta_{vg}\phi\overline{g}_v(\beta_u\psi, \beta_{vg}\phi) + \beta_u\psi\overline{g}_u(\beta_u\psi, \beta_{vg}\phi)$$

with $\overline{g}_*$ satisfying the same norm bound as before. Then for $s > R$ we define Θ_u to be

$$\Theta_u := \beta_v'\phi + \partial_s\psi + J_\delta(\overline{u} + \beta_{vg}\phi + \beta_u\psi)\partial_t\psi + \partial_\psi J_\delta(\overline{u}) + \psi\overline{g}_u(\beta_u\psi, \beta_{vg}\phi).$$

Note that we choose $\overline{g}_u$ to agree with g_u for $s < 2R$. Then we observe by this construction over $s \in [R, \infty)$, the equation:

$$\beta_u\Theta_u + \beta_v\Theta_v = 0$$

implies directly that the deformation of the pregluing u_* under $\beta_u\psi + \beta_v\phi$ is J_δ holomorphic.

The definition of Θ_u extends also naturally to $s \in [0, R]$ as

$$\partial_s(u_* + \psi) + J_\delta(u_* + \psi)\partial_t(u + \psi) = 0$$

as in this region the effect of ϕ vanishes. The extension of Θ_u to the interior of u is standard, albeit one also needs to take into account of deformation of complex structure δj in the interior of u.

As promised the derivatives of ψ does not appear in Θ_v and vice versa. As written it is manifest that solutions of Θ_u and Θ_v satisfy elliptic regularity. We next rewrite them into a form that makes the linearizations of operators appear, and hence more amendable to fixed point techniques.

Step 2 We now establish an alternative form of Θ_v, namely we take Equation 3.10 and expand the nonlinear terms. We get

$$\Theta_v = D_{J_\delta}\phi + \beta_u'\psi + \beta_{[1;R-2,\infty]}\phi g_{v1}(\beta_{ug}\psi, \beta_v\phi) + \partial_t\phi g_{v2}(\beta_{ug}\psi, \beta_{[1;R-2,\infty]}\beta_v\phi)$$

where g_{v*} have the same properties as g_*. Even though they are different functions, we will sometimes just write $g_v(\phi + \partial_t\phi)$ for convenience. We then can take

$$\mathcal{F}_v := \beta_{[1;R-2,\infty]}\phi g_{v1}(\beta_{ug}\psi, \beta_v\phi) + \partial_t\phi g_{v2}(\beta_{ug}\psi, \beta_{[1;R-2,\infty]}\beta_v\phi)$$

which we think of a quadratic term. There is no error term.

Step 3 The analogous expression for Θ_u is more complicated, in part because we have to deal with asymptotic vectors and have to pull back everything to $W^{2,p,d}(u^*TM) \oplus V_\Gamma$. We first focus on $s > R$ part of Θ_u from which we can write this as

$$\beta_v'\phi + [(\partial_s + J_\delta(\overline{u})\partial_t)\psi + \partial_\psi J_\delta(\overline{u})] + \psi\overline{g}_u(\beta_u\psi, \beta_{vg}\phi) + (J_\delta(\overline{u} + \beta_{vg}\phi + \beta_u\psi) - J_\delta(\overline{u}))\partial_t\psi = 0.$$

We loosely think of $[(\partial_s + J_\delta(\overline{u})\partial_t)\psi + \partial_\psi J_\delta(\overline{u})]$ as the linearization, and the rest of the terms as quadratic perturbation. The quadratic terms

$$\psi\overline{g}_u(\beta_u\psi, \beta_{vg}\phi) + (J_\delta(\overline{u} + \beta_{vg}\phi + \beta_u\psi) - J_\delta(\overline{u}))\partial_t\psi$$

generally take the form: $\psi \cdot g(\psi, \phi) + \partial_t\psi g(\psi, \phi)$ where g is the function having the property of Equation 3.9 and uniformly bounded derivatives.

Next we consider what happens for $s \leq R$, where Θ_u takes the form

$$\partial_s(\overline{u} + \psi) + J_\delta(\overline{u} + \psi)\partial_t(\overline{u} + \psi) = 0$$

which we can rewrite as

$$\partial_s\psi + J_\delta(\overline{u})\partial_t\psi + \partial_\psi J_\delta(\overline{u})\partial_t\overline{u} + g(\psi)\partial_t\psi + \partial_s\overline{u} + J_\delta(\overline{u})\partial_t\overline{u}.$$

We think of $\partial_s\psi + J_\delta(\overline{u})\partial_t\psi + \partial_\psi J_\delta(\overline{u})\partial_t\overline{u}$ as the linear term, $g(\psi)\partial_t\psi$ as the quadratic correction (g is just some function satisfying property of Equation 3.9), and $\partial_s\overline{u} + J_\delta(\overline{u})\partial_t\overline{u}$ the pregluing error, which was already estimated in the previous proposition.

We next wish to understand how the linear terms in the various segments of Θ_u compare with the linearization of $\bar{\partial}_J$ along u, which we turn to in the next step.

Step 4 We first focus on $s < R$. We are trying to compare the linearization term in Θ_u to D_J, which can be written as

$$\partial_s\psi + J(u)\partial_t\psi + \partial_\psi J(u)\partial_t u.$$

We compare their difference. We first consider the linear term in Θ_u with J instead of J_δ, and in taking their difference we see terms of the form

$$(J(u) - J(\overline{u}))\partial_t\psi + \partial_\psi J(u)\partial_t u - \partial_\psi J(\overline{u})\partial_t\overline{u}.$$

In the first term above the difference is of the form
$C(s,t)(r,a,p)\partial_t\psi + C(s,t)\beta_{[1;R-2.R+2;1]}e^{-DR}\partial_t\psi$ where $e^{-DR}\partial_t\psi$ is coming from pregluing error. In the second term above we can write it as:

$$C(s,t)(r,a,p)\psi + C\psi\beta_{[1;R-2.R+2;1]}e^{-DR}$$

the second term coming from the difference between $\partial_t u$ and $\partial_t\overline{u}$.

Then we must take into accout the difference between J_δ and J, this introduces terms of the form

$$C\delta\psi + C\delta\partial_t\psi.$$

This concludes our computations for the $s < R$ region. For $s > R$, we repeat a similar procedure, we recall the linear term in Θ_u in this region takes the form

$$\partial_s\psi + J_\delta(\overline{u})\partial_t\psi + \partial_\psi J_\delta(\overline{u})\partial_t\overline{u}.$$

As before to understand this difference we first replace instances of J_δ with J, and get

$$\begin{aligned}
&D_J\psi - (\partial_s\psi + J(\overline{u})\partial_t + \partial_\psi J(\overline{u})\partial_t\overline{u})\\
=&C(s,t)\{((r,a,p) + C(\delta + e^{-Ds}))\beta_{[1;R-1,2R+2;1]} + \beta_{[1;2R-2,2R+2;1]}(e^{-DR} + \delta)\}\partial_t\psi\\
&+ C(s,t)\{((r,a,p) + C(\delta + e^{-Ds}))\beta_{[1;R-1,2R+2;1]} + \beta_{[1;2R-2,2R+2;1]}(e^{-DR} + \delta)\}\psi
\end{aligned}$$

where the terms of the form $C(\delta + e^{-Ds})\beta_{[1;R-1,2R+2;1]}$ and $\beta_{[1;2R-2,2R+2;1]}(e^{-DR} + \delta)$ came from the difference between $v_{r,a,p}$ and $\overline{u}$. Finally the effect of putting J_δ is to add a term of size:

$$C\delta\psi + C\delta\partial_t\psi.$$

Hence collecting all of the above computations, we can write

$$\Theta_u = \beta_v'\phi + D_J\psi + \mathcal{F}_u + \mathcal{E}_u$$

where we think of $\mathcal{F}_u$ as the quadratic term and $\mathcal{E}_u$ as the error term. They take the following forms:

$$\mathcal{F}_u = \begin{cases} g(\psi)\partial_t\psi + C(s,t)(r,a,p)\psi + C(s,t)(r,a,p)\partial_t\psi + C(s,t)(r,a,p)^2 e^{-Ds}, s < R \\ \\ \psi\overline{g}_u(\beta_u\psi, \beta_{vg}\phi) + (J_\delta(\overline{u} + \beta_{vg}\phi + \beta_u\psi) - J_\delta(\overline{u}))\partial_t\psi + C(s,t)(r,a,p)\psi \\ +C(s,t)(r,a,p)\partial_t\psi \text{ for } s \geq R \end{cases}$$

and for $s < R$

$$\begin{aligned} \mathcal{E}_u =& C\delta\psi + C\delta\partial_t\psi + \beta_{[1;R-2.R+2;1]}e^{-DR}\partial_t\psi + \beta_{[1;R-2.R+2;1]}e^{-DR}\partial_t\psi \\ &+ C\delta(1 + |r,a,p|)\beta_{[0,R+1;1]} \\ &+ C[e^{-DR}(1 + |r,a,p|) + C\delta(1 + |r,a,p|)]\beta_{[1;R-2,R+2;1]}. \end{aligned}$$

For $s \geq R$ we have

$$\begin{aligned} \mathcal{E}_u =& C\{(\delta + e^{-Ds})\beta_{[1;R-1,2R+2;1]} + (\beta_{[1;2R-2,2R+2;1]}(e^{-DR} + \delta))\}\partial_t\psi \\ &+ C\{(\delta + e^{-Ds})\beta_{[1;R-1,2R+2;1]} + (\beta_{[1;2R-2,2R+2;1]}(e^{-DR} + \delta))\}\psi \\ &+ C\delta\psi + C\delta\partial_t\psi. \end{aligned}$$

□

We also need to version of elliptic regularity given in Proposition B.4.9 in [48], which we reproduce here.

Theorem 3.7.11. *Let $\Omega' \subset \Omega$ be open domains in $\mathbb{C}$ so that $\overline{\Omega}' \subset \Omega$. Let l be a positive integer and $p > 2$. Assume $J \in W^{l,p}(\Omega, \mathbb{R}^{2n})$ satisfies $J^2 = -1$, and $\|J\|_{W^{l,p}} \leq c_0$, then:*

a. If $u \in L^p_{loc}(\Omega, \mathbb{R}^{2n})$, $\eta \in W^{l,p}_{loc}(\Omega, \mathbb{R}^{2n})$, and u weakly solves

$$\partial_s u + J\partial_t u = \eta. \tag{3.12}$$

Then $u \in W^{l+1,p}_{loc}(\Omega, \mathbb{R}^{2n})$, and satisfies this equation almost everywhere.

b.

$$\|u\|_{W^{l+1,p}(\Omega',\mathbb{R}^{2n})} \leq c(\|\partial_s u + J\partial_t u\|_{W^{l,p}(\Omega,\mathbb{R}^{2n})} + \|u\|_{W^{l,p}(\Omega,\mathbb{R}^{2n})}) \tag{3.13}$$

where c only depends on $c_0, \Omega,$ and Ω'.

Remark 3.7.12. In what follows, ignoring for now our choice of cut off functions, we will think of $\mathcal{F}_v$ in the following form:

$$\mathcal{F}_v(\phi, \psi) = g(\phi, \psi)\phi + h(\phi, \psi)\partial_t(\phi)$$

where measured in C^0 norm we have,

$$|g(x, y)| \leq C(|x| + |y|)$$

$$|h(x, y)| \leq C(|x| + |y|)$$

We also have g, h are both smooth functions whose derivatives are uniformly bounded, which in particular implies that the $W^{k,p}$ norm of $g(\phi, \psi)$ and $h(\phi, \psi)$ are bounded above by the $W^{k,p}$ norm of ϕ and ψ.

In comparison with Section 5 of [41], our conditions on $\mathcal{F}_v$ are slightly stronger than the condition in there called *quadratic of type 2* because only the derivative of ϕ is allowed to appear. We will think of $\mathcal{F}_u$ in the following form (note this is slightly different from above conventions):

$$\mathcal{F}_u = g(\beta_v\phi, \psi, r, a, p) + h(\beta_v\phi, \psi, r, a, p)\partial_t(\psi)$$

Ignoring the precise details of cut off functions, we have (all norms below are the C^0 norm)

$$\|g\| \leq C(\|\phi\| \cdot \|\psi\| + \|\psi\|^2 + |(r, a, p)|^2 e^{-Ds} + |(r, a, p)|(\|\phi\| + \|\psi\|))$$

$$\|h\| \leq C(|(r, a, p)| + \|\phi\| + \|\psi\|)$$

Analogous expressions for pointwise bounds for higher derivatives of $\mathcal{F}_u$ also hold, essentially because $\mathcal{F}_u$ comes from expanding a smooth function. For most of our purposes the bounds above will suffice.

Remark 3.7.13. The terms $\mathcal{F}_*$ and $\mathcal{E}_*$ are viewed as error terms, so what is important is their relative sizes and not the constants appearing in front of them. In what follows we will not be too careful to distinguish $+\mathcal{F}_*$ and $-\mathcal{F}_*$ and similarly for $\mathcal{E}_*$.

Solution of $\Theta_u = 0, \Theta_v = 0$

In this subsection we will finally solve the system of equations $\Theta_u = 0$, $\Theta_v = 0$. We will adopt the following strategy:

- Given fixed $(r, a, p), \psi$, construct our lift of gradient trajectory, $v_{r,a,p}$, which we preglue to $u + (r, a, p)$.
- For this fixed choice, we solve $\Theta_v(\phi) = 0$ over $v^*_{r,a,p}TM$ using the contraction mapping principle to obtain an unique solution, $\phi(r, a, p, \psi)$.

- Then we try to solve $\Theta_u = 0$ over u^*TM. We do this via another contraction mapping principle with input variables $(\psi, r, a, p, \delta j)$. The function ϕ enters the equation, but as a dependent on these variables. As such, we need to understand how ϕ varies when we change ψ, r, a, p. This is made non-trivial by the fact when we change variables r, a, p, the deformation is not local, we are twisting/moving an entire semi-infinite cylinder. We will need to understand under these changes, how the ϕ terms that enter Θ_u change. Hence to keep track of these changes, we will make certain identifications of bundles $v^*_{r,a,p}TM$ and $v^*_{r',a',p'}TM$ so we can compare the solutions of different equations over the same space. Then from that we get from the perspective of the equation Θ_u over u^*TM, $\phi_{r,a,p}$ depends nicely on the variables $(r, a, p, \psi, \delta j)$.

- We apply the contraction mapping principle over u^*TM to solve Θ_u.

Proposition 3.7.14. *Let $\epsilon > 0$ be fixed and sufficiently small (small relative to the constants C that describe the local geometry of Morse-Bott torus but fixed with respect to $\delta > 0$). Let the tuple $(\psi, r, a, p, \delta j)$ be fixed and in an ϵ ball around zero. Then we can view $\Theta_v = 0$ as an equation with input $\phi \in W^{2,p,d}(v^*_{r,a,p}TM)$. This equation has an unique solution $\phi \in W^{2,p,d}(v^*_{r,a,p}TM)$ whose norm is bounded by*

$$\|\phi\| \le \epsilon/R.$$

Furthermore, this solution ϕ is actually in $W^{3,p,d}(v^*_{r,a,p}TM)$, its $W^{3,p,d}$ norm is likewise bounded by $C\epsilon/R$.

Proof. The equation we need to solve is

$$D_v\phi + \mathcal{F}_v(\phi, \psi) + \beta'_u\psi = 0$$

where D_v is the linearization of $\bar{\partial}_{J_\delta}$ along $v_{r,a,p}$, which we previously denoted by D_{J_δ}. We dropped the subscript (r, a, p) to make the notation manageable.

Let Q denote the inverse of D_v. Now consider the map
$I : W^{2,p,d}(v^*_{r,a,p}TM) \to W^{2,p,d}(v^*_{r,a,p}TM)$ defined by

$$\phi \longrightarrow Q(-\beta'_v\psi - \mathcal{F}_v(\phi, \psi))$$

We note a solution to $\Theta_v = 0$ is equivalent to I having a fixed point. We demonstrate a fixed point exists via the contraction mapping principle. Since $\psi \in W^{2,p,d}(u^*TM)$ has norm $\le \epsilon$, the function $\beta_u\psi$ viewed as a element in $W^{2,p,d}(v^*_{r,a,p}TM)$ also has norm bounded above by ϵ, hence $\|\beta'_u\psi\| \le \epsilon/R$. Also we note for $\|\phi\| \le \epsilon$, $\|Q \circ \mathcal{F}_v\| \le C\epsilon^2$, both of these being measured in $W^{2,p,d}(v^*_{r,a,p}TM)$ norm.

If we let $B_\epsilon(0)$ denote the ϵ ball in $W^{2,p,d}(v^*_{r,a,p}TM)$ then by the above, we see I sends $B_\epsilon(0)$ to itself. We also see it satisfies the contraction mapping property. If $\phi, \phi' \in B_\epsilon(0)$ then

$$\begin{aligned}\|I(\phi) - I(\phi')\| &\le \|\mathcal{F}_v(\phi, \psi) - \mathcal{F}_v(\phi', \psi)\| \\ &\le Cmax\{\|\phi\|_{W^{2,p,d}}, \|\phi'\|_{W^{2,p,d}}, \|\psi\|_{W^{2,p,d}}\}\|\phi - \phi'\| \\ &\le C\epsilon\|\phi - \phi'\|\end{aligned}$$

Hence for small enough ϵ the conditions for contraction principle is satisfied, the map I has a unique fixed point. Since $D_{v_{r,a,p}}$ is invertible, this is equivalent to Θ_v having a unique solution.

We can estimate the size of this fixed point. Consider the equation

$$\phi = Q(-\beta_v'\psi - \mathcal{F}_v(\phi,\psi))$$

If we measure the size of both sides in $W^{2,p,d}(v_{r,a,p}^*TM)$ we get

$$\|\phi\| \leq C\epsilon/R + C\epsilon\|\phi\|$$

hence we get

$$\|\phi\| \leq C\epsilon/R.$$

The fact we can improve the regularity and bound the $W^{3,p,d}$ norm of ϕ follows directly from Theorem 3.7.11. □

We next need to solve Θ_u. As we mentioned in the introduction to this subsection, we think of this equation taking place over $W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J}$, with input variables $(r,a,p,\psi,\delta j)$ in an ϵ ball. We think of the $\phi(r,a,p,\psi,\delta j)$ term that appears as a dependent variable. From above we know that for each tuple $(r,a,p,\psi,\delta j)$ there exists a unique solution $\phi(r,a,p,\psi,\delta j)$ of small norm. To apply the contraction mapping principle we need to see the *derivative* of $\phi(r,a,p,\psi,\delta j)$ with respect to the tuple $(r,a,p,\psi,\delta j)$ behaves nicely as well. This is made nontrivial by the fact when we vary (r,a,p) we are pregluing different gradient trajectories, hence the solution of $\Theta_{v_{r,a,p}} = 0$ takes place in different spaces. We take the approach of identifying all the solutions into one space, and that as (r,a,p) vary the equation over the *same* space changes, and by understanding this change, we understand how the terms in Θ_u change.

To this end we let the pair (D_v, W) denote the vector space

$$\left\{W^{2,p,d}(v_{0,0,0}^*TM), e^{ds}\right\}$$

with operator D_v given by $D_{v_{0,0,0}}$-the linearization of $\bar\partial_{J_\delta}$ over $v_{0,0,0}$. We first consider varying r,a, and keeping $p=0$. Let $\phi_{r,a,0} \in W^{2,p,d}(v_{r,a,0}^*TM)$, then there is an obvious parallel transport map $PT: W^{2,p,d}(v_{r,a,0}^*TM) \to W$ that sends

$$\phi_{r,a,0} \longrightarrow \phi_{r,a,0} - r - a \in W$$

which is an isometry. Here we are using additive notation for parallel transport maps because the metric is flat. We denote its image by $\hat\phi_{r,a,0}$. Under this identification $\hat\phi_{r,a,0}$ satisfies a different equation, which we denote by $\hat\Theta_v$. This equation is of the form

$$\hat D_{r,a,0}\hat\phi_{r,a,0} + \hat{\mathcal{F}}_v(\hat\phi_{r,a,0},\psi) = 0$$

If we write $D_v = \partial_s + J_\delta(s,t)\partial_t + S(s,t)$, then $\hat D_{r,a,0}$ is given by

$$\partial_s + J_\delta(s+a,t+r)\partial_t + S(s+a,t+r)$$

and the term $\hat{\mathcal{F}}_v$ has some mild dependence on r, a depending on the local geometry of the Morse-Bott torus. The point here is that the term $\phi(r, a, p, \psi, \delta j)$ that enters directly Θ_u can be identified with $s \leq 5R$ portion of $\hat{\phi}_{r,a,0}$ solving $\hat{\Theta}_v = 0$. Hence to understand how $\phi(r, a, p, \psi, \delta j)$ feeds back into Θ_u as we vary (r, a) we only need to understand how the parametrized solutions $\hat{\phi}_{r,a,0}$ solving the (r, a) parametrized family of PDEs $\hat{\Theta}_v$ changes.

We would like to extend the previous discussion to include variations of p. To do that we need the next lemma which concerns the differential geometry of the situation.

Lemma 3.7.15. *In our chosen coordinate system, let $v_p(s,t) = (a_p(s), t, x_p(s), 0)$ be a lift of gradient trajectory satisfying*

$$a_p(0) = 0, x_p(0) = p$$

and let $v_{p'} = (a_{p'}(s), t, x_{p'}(s), 0)$ be another lift satisfying

$$a_{p'}(0) = 0, x_{p'}(0) = p'$$

with $|p - p'| \leq \epsilon$ then there exists a C independent of δ such that:

$$|v_p(s,t) - v_{p'}(s,t)| \leq C|p - p'|$$

for all s, t.

Proof. This is a fundamentally a statement about gradient flows. We recall x_p satisfies the equation $x_p(s)_s = \delta f'(x_p)$. If we reparametrize $x_p(s)$ to be $x_p(\frac{s}{\delta})$ then it is simply a gradient flow of f, then we have for all δ and all s

$$\left| x_p\left(\frac{s}{\delta}\right) - x_{p'}\left(\frac{s}{\delta}\right)\right| \leq C|p - p'|$$

where C is independent of δ, and the claim follows. To verify the claim about a_p and $a_{p'}$, we need to be a bit more careful. Assume we have re-chosen coordinates $[U, V]$ around p so that on $x \in [U, V]$, we have $f(x) = Mx$ and on $x \in [0, U + \epsilon]$ $f(x) = M'x^2/2$.

This fixed choice of coordinate is independent of δ so our quantitative conclusions drawn from this coordinate system continues to hold in our original coordinate system up to a change of constant. Then we analyze the behaviour of $a_p(s)$ and $a_{p'}(s)$ as $s \to -\infty$, with the positive end being similar. For s so that $x_p(s) \in [U, V]$, $x_p(s) = \delta M s + p$ and $x_{p'}(s) = \delta M s + p'$, since $a_p' = e^{\delta f(x_p)}$ this is equivalent to

$$a_p(s)' = e^{\delta(\delta M s + p)}$$

and

$$a_{p'}(s)' = e^{\delta(\delta M s + p')}.$$

If we take T large enough so that $x_p(-T), x_{p'}(-T) \in [U, U + \epsilon]$, we have the upper bound:

$$T < C(\max(p, p') - U)/\delta.$$

Hence we have a uniform upper bound in the integral of the form:

$$\begin{aligned}|a_p(-T) - a_{p'}(-T)| &\leq \left| \int_0^{A/\delta} e^{\delta^2 Ms+\delta p} - e^{\delta^2 Ms+\delta p'} ds \right| \\ &\leq \frac{1}{\delta^2 M}(e^{\delta^2 MA/\delta} - 1)(e^{\delta p} - e^{\delta p'}) \\ &\leq C|p-p'|\end{aligned}$$

where A is just some constant. Next for $s < -T$, the curves $x_p(s)$ and $x_{p'}(s)$ enter the region where $f(x) = M'x^2/2$, and they satisfy the differential equation

$$x_p(s)' = \delta M' x$$

so they satisfy

$$x_p(s) = x_p(-T)e^{\delta M' s}$$

and likewise for $x_{p'}(s)$, hence the difference between $a_p(s)$ and $a_{p'}(s)$ satisfies

$$\partial_s(a_p - a_p') = e^{\delta x_p(-T)e^{\delta M's}} - e^{\delta x_{p'}(-T)e^{\delta M's}}.$$

Since we are taking $s \to -\infty$ the right hand side can be bounded by

$$C\delta(x_p(-T) - x_{p'}(-T))e^{\delta M's} \leq \delta C|p-p'|e^{\delta M's}.$$

Hence for $s < -T$

$$|a_p - a_{p'}| \leq C|p-p'|\delta \int_{-\infty}^{-T} e^{\delta M's} ds \leq C|p-p'|$$

and hence our conclusion follows. □

With the above calculations we have can extend the parallel transport map

$$PT : W^{2,p,d}(v_{0,0,p}^* TM) \longrightarrow W$$

defined by

$$PT(\psi) = \psi - (v_{0,0,p} - v_{0,0,0})$$

which is also an isometry. The $\Theta_v = 0$ equation also pulls back to W, the linearized operator $\hat{D}_{0,0,p}$ given by

$$\partial_s \hat{\phi}_{0,0,p} + J_\delta(p)\partial_t \hat{\phi}_{0,0,p} + S(p,s,t)\hat{\phi}_{0,0,p}.$$

This is what $D_{v_{0,0,p}}$ would look like in the coordinates we chose for $v_{0,0,p}$. The full equation $\hat{\Theta}_v = 0$ looks like

$$\hat{D}_{0,0,p}\hat{\phi}_{0,0,p} + \hat{\mathcal{F}}_v(p, \hat{\phi}_{0,0,p}, \psi) = 0.$$

The previous lemma ensures the coefficient matrices J_δ, $S(p,s,t)$, as well as $\hat{\mathcal{F}}_v$ are uniformly well behaved (say in C^k norm) as we vary p as $\delta \to 0$. And as before the components of $\phi(r,a,p,\psi,\delta j)$ that enters Θ_u can be identified with the $s > 5R$ component of $\hat{\phi}_{0,0,p}$. Combining the previous discussion, we have:

Proposition 3.7.16. *The derivative of the operator $\hat{D}_{v_{r,a,p}} : W \to W^{1,p,d}(v_0^*TM)$ with respect to (r,a,p) is well defined, and satisfies for a fixed constant C*

$$\|\partial_a D_{v_{r,a,p}}\| = 0$$

$$\|\partial_r D_{v_{r,a,p}}\| \leq C$$

$$\|\partial_p D_{v_{r,a,p}}\| \leq C.$$

Further we have the bound:

$$\|\partial_* \hat{\mathcal{F}}_v\| \leq C.$$

In the above, $\partial_* D_{v_{r,a,p}}$ is viewed as an operator $W \to W^{1,p,d}(v_0^*TM)$, and $\hat{\mathcal{F}}_v$ is viewed as a map from $W \to W^{1,p,d}(v_0^*TM)$ and its derivative is a map over the same space.

The next step is to understand how $\hat{\phi}_{r,a,p}$ varies with respect to the variables (r,a,p,ψ).

Proposition 3.7.17. *Fix (r,a,p,ψ), let $\hat{\phi}(r,a,p,\psi,\delta j)$ (which for the purpose of this proof we abbreviate $\hat{\phi}$) denote the solution of $\hat{\Theta}_v = 0$ viewed as an element of W. We can take the derivative of $\partial_*\hat{\phi}(r,a,p,\psi)$ where $* = r,a,p,\psi$. They satisfy the equations*

$$\|\partial_*\hat{\phi}\| \leq C\epsilon, \quad * = r,a,p,\psi,\delta j$$

where the norm is measured in $W^{2,p,d}(v_0^*TM)$ for $* = r,a,p$ and in $Hom(W^{2,p,w}(u^*TM), W^{2,p,w}(v_0^*TM))$ for $* = \psi$, and $Hom(T\mathcal{J}, W^{2,p,w}(v_0^*TM))$ for $* = \delta j$. See Remark 3.7.18 for the interpretation of terms $\partial_*\hat{\phi}$.

Proof. The fixed point equation looks like

$$\hat{\phi} = \hat{Q}_{r,a,p}(-\beta_u'\psi - \hat{\mathcal{F}}_v(\hat{\phi},\psi,r,a,p)).$$

This is really a family of equations over ψ, r, a, p.

We first differentiate w.r.t. ψ. Note unlike differentiating w.r.t. r,a,p, the term $\frac{d\phi}{d\psi}$ is a Frechet derivative which should be viewed as a linear operator

$$\frac{d\hat{\phi}}{d\psi} : W^{2,p,w}(u^*TM) \longrightarrow W^{2,p,w}(v_0^*TM)$$

and when we measure its norm it is the operator norm. When we write β_u' below we mean the operator defined by multiplication by β_u' etc. Differentiating both sides of the fixed point equation we have

$$\left\|\frac{d\hat{\phi}}{d\psi}\right\| \leq C\left\|Q \circ \left(1/R + \psi\frac{d\hat{\phi}}{d\psi} + (\partial_t\hat{\phi})\cdot\frac{d\hat{\phi}}{d\psi} + \hat{\phi}\partial_t\frac{d\hat{\phi}}{d\psi}\right)\right\|$$

where the norms of both sides are operator norms. To make sense of $Q\hat{\phi}\partial_t\frac{d\hat{\phi}}{d\psi}$ see Remark 3.7.18. We recall the C^1 norm of is ψ bounded above by $C\epsilon$ via the Sobolev embedding theorem, so the above equation can be rearranged to be

$$(1-C\epsilon)\|d\hat{\phi}/d\psi\| \leq C(1/R) \leq \epsilon$$

and this proves the claim for $* = \psi$. We next consider the case $* = p$, the cases for $* = r, a$ are analogous. We consider the equation

$$\hat{D}_{r,a,p}\hat{\phi} + \hat{\mathcal{F}}_v(\psi, r, a, p, \hat{\phi}) + \beta_u'\psi = 0$$

and differentiate both sides w.r.t. p:

$$\frac{d\hat{D}_{r,a,p}}{dp}\hat{\phi} + \hat{D}_{r,a,p}\frac{d\hat{\phi}}{dp} = -\frac{d}{dp}\mathcal{F}_v(r, a, p, \psi, \hat{\phi}) \tag{3.14}$$

rearranging to get

$$\frac{d\hat{\phi}}{dp} = \hat{Q}_{r,a,p}\left[-\frac{d\hat{D}_{r,a,p}}{dp}\hat{\phi} - \frac{d}{dp}\mathcal{F}(r, a, p, \psi, \hat{\phi})\right].$$

The only new thing we need to estimate is $\frac{d}{dp}\hat{\mathcal{F}}(r, a, p, \psi, \hat{\phi})$ which we calculate as

$$\begin{aligned}
&\left\|\hat{Q}_{r,a,p} \circ \frac{d}{dp}\hat{\mathcal{F}}_v(r, a, p, \psi, \hat{\phi})\right\| \\
&= \left\|\hat{Q}_{r,a,p} \circ \left\{\partial_{\hat{\phi}}g(\hat{\phi}, \psi)\frac{d\hat{\phi}}{dp} + \partial_{\hat{\phi}}h(\hat{\phi}, \psi)\partial_t\hat{\phi}\frac{d\hat{\phi}}{dp} + h(\hat{\phi}, \psi)\partial_t\frac{d\hat{\phi}}{dp} + \partial_p g(\hat{\phi}, \psi) + \partial_p h(\hat{\phi}, \psi)\partial_t(\hat{\phi})\right\}\right\| \\
&\leq C(\|\psi\| \cdot \|\hat{\phi}\| + \|\hat{\phi}\|^2) + \epsilon\left\|\frac{d\hat{\phi}}{dp}\right\|
\end{aligned}$$

where $\partial_p h$ and $\partial_p g$ refer to how the functions h and g themselves depend on p. The norms above are all in $W^{2,p,d}$ norm in the domain. Combining the above inequalities we conclude

$$\|d\hat{\phi}/dp\| \leq C\epsilon.$$

The same proof works for $* = r, a, \delta j$.

□

Remark 3.7.18. As we mentioned in the course of the proof, the attentive reader might feel uneasy about the appearance of the term $Q \circ \partial_t\frac{d\hat{\phi}}{d\psi}$. The proper way to take this Frechet derivative is explained in Proposition 5.6 in [41]. The idea is to take the fixed point equation

$$\hat{\phi} = \hat{Q}_{r,a,p}(-\beta_u'\psi - \hat{\mathcal{F}}(r, a, p, \hat{\phi}, \psi)).$$

Let $\mathcal{D}_\psi$ and $\mathcal{D}_{\hat{\phi}}$ denote the derivative of the right hand side of the above equation with respect to $\psi, \hat{\phi}$ respectively, then the derivative $\frac{d\hat{\phi}}{d\psi}$ is defined to be $(1-\mathcal{D}_{\hat{\phi}})^{-1}\mathcal{D}_\psi$, and sends $W^{2,p,w}(u^*TM) \to W^{2,p,w}(v_0^*TM)$. In light of this the composition $Q \circ \partial_t \frac{d\hat{\phi}}{d\psi}$ is not at all problematic, and our estimates of norms continue to hold.

We take our approach to things because in the case of differentiation with respect to the parameters $* = r, a, p$ (say for p for definiteness), the resulting derivative is a linear map from $\mathbb{R} \to W^{2,p,w}(v_0^*TM)$, and hence has a right to be viewed as an honest function in $W^{2,p,w}(v_0^*TM)$. It further satisfies an elliptic PDE as in Equation 3.14, which gives us estimates on its norms and exponential decay properties, which will be essential for our later purposes.

Remark 3.7.19. We can actually show $\frac{d\hat{\phi}}{dp}$ (and likewise for $* = r, a$) belongs in $W^{3,p,d}(v_0^*TM)$ with the help of elliptic regularity. If we recall the form of $\Theta_v = 0$ written in Equation 3.10, we differentiate with respect to p to see $\frac{d\hat{\phi}}{dp}$ weakly satisfies an equation of the form

$$\partial_s \frac{d\hat{\phi}}{dp} + J_\delta(v_{r,a,p} + \beta_{ug}\psi + \beta_{[1;R-2,\infty]}\beta_v\hat{\phi})\partial_t \frac{d\hat{\phi}}{dp} + \mathcal{F}(\hat{\phi}, \partial_t\hat{\phi}, \frac{d\hat{\phi}}{dp}, \psi) = 0$$

where $\mathcal{F}$ is a smooth function. Using elliptic regularity we see that $\frac{d\hat{\phi}}{dp}$ is in $W^{3,p,d}(v_0^*TM)$ with norm bounded above by $C\epsilon$.

Finally we turn our attention to solving $\Theta_u(\psi, \phi, r, a, p) = 0$.

Proposition 3.7.20. *The equation $\Theta_u = 0$ has a solution.*

Proof. Recall we view $\Theta_u(r, a, p, \psi, \phi(r, a, p, \psi))$ as an equation with independent variables $(r, a, p, \psi, \delta j) \in W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J}$, and we have the surjective linearized operator

$$D_J : W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J} \longrightarrow W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u^*TM)).$$

The equation $\Theta_u = 0$ over the entire domain $\dot{\Sigma}$ can be written to be of the form

$$D_J(r, a, p, \psi, \phi) + \mathcal{F}_u + \mathcal{E}_u + \mathcal{F}_{int}(\psi, \delta j) = 0 \tag{3.15}$$

where $\mathcal{F}_u$ is supported in the cylindrical neighborhood $[0, \infty) \times S^1$ and the term $\mathcal{F}_{int}$ is quadratic in $\psi, \delta j$ and supported in $\dot{\Sigma} \setminus [0, \infty) \times S^1$ and $\mathcal{E}_u$ is described by Proposition 3.7.9 and supported in the cylindrical neck. We let Q_u denote a right inverse to D_J. Then to find a solution to Θ_u it suffices to find a fixed point of the map

$$I : (r, a, p, \psi, \delta j) \longrightarrow Q_u(-\mathcal{F}_u - \mathcal{E}_u - \mathcal{F}_{int}).$$

Let B_ϵ denote the ϵ ball in $W^{2,p,d}(u^*TM) \oplus V_\Gamma \oplus T\mathcal{J}$. It follows from the fact $\mathcal{F}_{int}$ is quadratic, our previous derived expressions for $\mathcal{F}_u$, $\mathcal{E}_u$, and our size estimate $\|\phi\| \leq C\epsilon$ that for $\delta << \epsilon$ small enough, I maps B_ϵ to itself. It further follows from the fact that $\|\partial_* \phi\| \leq C\epsilon$ and the explicit expressions of $\mathcal{F}_u, \mathcal{F}_{int}, \mathcal{E}_u$ that I is a contraction mapping, and hence a solution to $\Theta_u = 0$ exists. ∎

Remark 3.7.21. The contraction mapping principle actually says the fixed point of I is unique. However this does not mean the solution to Θ_u is unique. If D_J is not injective (which it never is if the curve is nontrivial due to translations in the symplectization direction, and if the curve is a free trivial cylinder there is also a global translation along the Morse-Bott torus), we could have chosen a different right inverse Q' which leads to a (presumably) different solution of Θ_u. We leave discussions of uniqueness of gluing to after when we glued together general cascades.

We note that even though the previous construction was only for one end, the construction works for arbitrary number of free ends.

Corollary 3.7.22. *Given a transversely cut out J-holomorphic curve u with free-end Morse-Bott asymptotics, the ends can be glued with semi-infinite gradient trajectories into J_δ-holomorphic curves.*

Remark 3.7.23. In the above we only glued gradient flows to free ends. We could have also glued in trivial cylinders to fixed ends. The only difference is instead of V_Γ being spanned by r, a, p it is only spanned by r, a. The rest of the argument follows exactly the same way.

3.8 Exponential decay for solution of Θ_v

Consider, in notation of previous section, $\Theta_v = 0$ for $s > 3R$, then it is an equation of the form

$$D_v\phi + \mathcal{F}(\phi) = 0$$

where D_v denotes the linearization of ∂_{J_δ} operator, and $\mathcal{F}$ we loosely think of as an quadratic expression in ϕ and $\partial_t\phi$, see Remark 3.7.12. In this section we study the properties of this solution, in particular, it exhibits exponential decay as $s \to \infty$ for δ sufficiently small (exponential decay beyond what is imposed by the exponential weight e^{ds}). This property will be crucial for our gluing construction for multiple level cascades. The idea why ϕ undergoes exponential decay is the following: for $\delta > 0$ sufficiently small, the gradient flow cylinder $v_{r,a,p}$ flows so slowly that locally the geometry resembles that of a trivial cylinder, and the usual proof that J-holomorphic curve decays exponentially along asymptotic ends can be applied.

The section is organized as follows: We first remove the exponential weights from our Sobolev spaces and work instead in $W^{2,p}(v^*TM)$. Next we follow the strategy of [27] Section 2(this strategy as far as we know also dates back at least to [19], see Section 4, and is used frequently in various kinds of Floer homologies to prove exponential convergence), using second derivative estimates to derive the exponential decay, and finally we show the various derivatives of ϕ also decay exponentially.

Exponential decay for solutions of Θ_v

We begin by studying the exponential decay of ϕ, then move on to study the exponential decay of its derivatives. First some setup that will be used for both cases.

Change of coordinates and setup

We study $\Theta_v = 0$ for $s > 3R$, which takes the form

$$D_v\phi + \mathcal{F}_v(\phi) = 0.$$

WLOG we assume $(r, a, p) = (0, 0, 0)$ and write v instead of $v_{r,a,p}$. It will be clear our analysis holds for any value of (r, a, p) and later we will identify sections of $v_{r,a,p}^*TM$ with sections of v_0^*TM via parallel transport.

Recall $\mathcal{F}_v(\phi)$ takes the form

$$\mathcal{F}_v(p, \phi) = g(p, \phi)\phi + h(p, \phi)\partial_t\phi.$$

Here we have made explicit the dependence of this term on the p, which controls the background geometry. It also implicitly depends on (r, a), which we suppress from our notation. Here the functions satisfy (uniformly in p) $\|g(\phi)\|_{C^0} \leq \|\phi\|_{C^0}$, $\|\nabla g(\phi)\|_{C^0} \leq \|\nabla\phi\|_{C^0}$. For $h(\phi)$ we have $\|h(\phi)\|_{C^0} \leq \|\phi\|_{C^0}$, $\|\nabla h(\phi)\|_{C^0} \leq \|\nabla\phi\|_{C^0}$. These bounds will be important to us in the subsequent estimates.

Next we change variables to $W^{2,p}(v^*TM)$ i.e. by conjugation we remove the exponential weights on our space. We use the following diagram:

$$\begin{tikzcd}
W^{2,p}(v^*TM) \arrow[r, "\Theta_v'"] \arrow[d, "e^{-d(s)}"] & W^{2,p}(v^*TM) \\
W^{2,p,d}(v^*TM) \arrow[r, "\Theta_v"] & W^{2,p,d}(v^*TM). \arrow[u, "e^{d(s)}"']
\end{tikzcd}$$

We use this to define the operator Θ_v'. In terms of actual equations it looks like this: if ζ is the corresponding element of ϕ without exponential decay (i.e. $\zeta = e^{ds}\phi$), then Θ_v' is the same as

$$D_{J_\delta}'\zeta + e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta) = 0$$

where $D_{J_\delta}' = e^{ds}D_{J_\delta}e^{-ds}$. We decompose D_{J_δ}' as follows

$$D_{J_\delta}' = d/ds - A(s) - \delta A$$

where by $A(s)$ we denote self adjoint operator associated with linearizing $\bar{\partial}_J$ along Morse-Bott orbit plus d due to the exponential conjugation

$$A = -J_0\partial_t - S + d.$$

Consequently the eigenvalues of $A(s)$ are bounded away from zero, say by a factor of $\lambda > 0$.

Remark 3.8.1. We will often change the value of λ from one line to another, as long as it is bounded away from zero. The choice of λ above depends somewhat on the choice of d, because the operator $-J_0\frac{d}{dt} - S$ has zero as an eigenvalue. With more careful estimates we can make the decay rate only depend on local geometry, but this won't be necessary for us for purpose of gluing.

In the Section 3.11 we also use a λ to describe exponential decay behaviour of J_δ-holomorphic curve near a Morse-Bott torus, there the λ is genuinely independent of d and only dependent on the local geometry, as will be apparent from our proofs.

δA is the perturbed correction to A due to the fact we are using J_δ instead of J. It has the form

$$\delta A = \delta M d/dt + \delta N$$

where we use M, N to denote matrices whose entries are uniformly bounded in C^k (by abuse of notation we will later use them to denote other matrices where each of the coefficient terms is uniformly bounded).

Exponential decay estimates

Let us define

$$g(s) := \int_{S^1} \langle \zeta(s,t), \zeta(s,t) \rangle dt.$$

We shall show:

Proposition 3.8.2.

$$g''(s) \geq \lambda^2 g. \tag{3.16}$$

This proposition combined with the following proposition, will imply exponential decay:

Proposition 3.8.3 (Lemma 8.9.4 in [1])**.** *If $g''(s) \geq \lambda^2 g(s)$ for $s > s_0$, then either:*

- $g(s) \leq g(s_0)e^{-\lambda(s-s_0)}$,
- $g(s) \to \infty$ *as* $s \to \infty$.

Proof of Proposition 3.8.2.

$$g''(s) = 2(\langle \zeta_s, \zeta_s \rangle + \langle \zeta_{ss}, \zeta \rangle)$$

where when we write $\langle \cdot, \cdot \rangle$ we implicitly take the S^1 integral over t. The proof is long and we separate it into steps.

Step 1 Let us first determine $\langle \zeta_s, \zeta_s \rangle$. This is given by

$$\begin{aligned}\langle \zeta_s, \zeta_s \rangle =& \langle (A+\delta A)\zeta + e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta), (A+\delta A)\zeta + e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle \\ =& \langle A\zeta, A\zeta\rangle \\ &+ \langle A\zeta, \delta A\zeta\rangle \\ &+ \langle A\zeta, e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle \\ &+ \langle \delta A\zeta, \delta A\zeta\rangle \\ &+ \langle \delta A\zeta, e^{d(s)}\mathcal{F}_v(e^{-w(s)}\zeta)\rangle \\ &+ \langle e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta), e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta)\rangle.\end{aligned}$$

We recall we are not tracking the signs in front of the term $\mathcal{F}_v$ since it will eventually be upperbounded. We look at the **six** terms, which we label by **T1-T6**, in the above expression one by one, we use **bold** to remind the reader which term we are referring to since the computation gets very long. Also when we upper bound terms from **T1-T6** we are implicitly taking the absolute value of terms. We shall keep this convention for all proofs involving exponential decay.

T1 gives

$$\langle A\zeta, A\zeta\rangle \geq \lambda^2 \langle \zeta, \zeta\rangle.$$

This is because if we expand $\zeta = \sum a_n e_n$, with the collection $\{e_n(s)\}$ of eigenbasis for $A(s)$, we have $Aa_ne_n = \sum \lambda_n a_n e_n$. We see this is greater than $\lambda^2\langle \zeta, \zeta\rangle$.

T2 is given by

$$\begin{aligned}\langle A\zeta, \delta A\zeta\rangle &= \langle A\zeta, \delta(MA+N)\zeta\rangle \\ &= \delta\langle A\zeta, MA\zeta\rangle + \delta\langle A\zeta, N\zeta\rangle.\end{aligned}$$

The first term above is bounded in absolute value by

$$|\delta\langle A\zeta, MA\zeta\rangle| \leq \delta(\|A\zeta\|^2 + \|MA\zeta\|^2) \leq C\delta\langle A\zeta, A\zeta\rangle.$$

The second term satisfies

$$|\delta\langle A\zeta, N\zeta\rangle| \leq \delta(\langle A\zeta, A\zeta\rangle + \langle N\zeta, N\zeta\rangle).$$

Hence for the **T2** term we have the overall bound by

$$\langle A\zeta, \delta A\zeta\rangle \leq C\delta(\langle A\zeta, A\zeta\rangle + \langle \zeta, \zeta\rangle).$$

T3 satisfies

$$\begin{aligned}&|\langle A\zeta, e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle| \\ =&\left|\langle \sqrt{\epsilon}A\zeta, \frac{1}{\sqrt{\epsilon}}e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle\right| \\ \leq& \epsilon\langle A\zeta, A\zeta\rangle + \frac{1}{\epsilon}\langle e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta), e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta)\rangle \\ \leq& \epsilon\langle A\zeta, A\zeta\rangle + \epsilon\langle \zeta, \zeta\rangle.\end{aligned}$$

In the last line we used the fact $\partial_t \zeta$ and ζ have C^0 norm uniformly bounded above by $C\epsilon$. **T4** satisfies

$$\begin{aligned}\langle \delta A\zeta, \delta A\zeta\rangle =&\delta^2 \langle MA + N\zeta, MA + N\zeta\rangle\\ =&\delta^2\langle MA\zeta, MA\zeta\rangle + \langle N\zeta, N\zeta\rangle + \langle MA\zeta, N\zeta\rangle\\ \leq& 2\delta^2\langle MA\zeta, MA\zeta\rangle + \langle N\zeta, N\zeta\rangle\\ \leq& C\delta^2(\langle A\zeta, A\zeta\rangle + \langle \zeta,\zeta\rangle).\end{aligned}$$

T5 satisfies

$$\begin{aligned}&\langle \delta A\zeta, e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle\\ =&\delta\left\langle MA + N\zeta, e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle\right\rangle\\ \leq&\delta\left[C\langle A\zeta, e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta)\rangle + \langle N\zeta, e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)\rangle\right]\\ \leq&\delta\left[C\epsilon\langle A\zeta, A\zeta\rangle + (1+\epsilon)\langle\zeta,\zeta\rangle\right]\end{aligned}$$

as the second last line follows from previous computation. **T6** satisfies

$$\langle e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta), e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta)\rangle \leq C\epsilon\langle\zeta,\zeta\rangle$$

simply because $\mathcal{F}_v$ is quadratic. Putting all these terms together we conclude

$$\langle\zeta_s,\zeta_s\rangle \geq (\lambda^2 - C\epsilon)\langle\zeta,\zeta\rangle$$

and this concludes the first step.
Step 2 We next compute

$$\begin{aligned}&\langle\zeta_{ss},\zeta\rangle\\ =&\langle\frac{d}{ds}[(A+\delta A)\zeta + e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta)],\zeta\rangle\\ =&\langle d/ds(A+\delta A)\zeta,\zeta\rangle\\ &+\langle (A+\delta A)\zeta_s,\zeta\rangle\\ &+\langle\frac{d}{ds}e^{d(s)}\mathcal{F}_v(e^{-d(s)})\zeta,\zeta\rangle.\end{aligned}$$

We will need to dissect these terms one by one. We label them **T1-T3**. For **T1** recall

$$A = -J_0 d/dt - S(s,t) + d,$$

hence its s derivative is a uniformly bounded matrix $\frac{dS}{ds}$ of norm $\leq C\delta$. Here we are using the fact J is the standard almost complex structure along the surface of the Morse-Bott torus. We also recall

$$\delta A = \delta(M\partial_t + N).$$

When we take its s derivative we get

$$\frac{d}{ds}\delta A = \delta\frac{dM}{ds}\partial_t + \delta\frac{dN}{ds}.$$

Again we have

$$\partial_s(\delta A) = \delta^2(MA + N)$$

where M, N denotes matrices with bounded entries. So the **T1** term is given by:

$$\begin{aligned}
&|\langle \partial_s(A+\delta A)\zeta, \zeta\rangle| \\
&\leq C\delta\langle\zeta,\zeta\rangle + \langle \delta^2 A\zeta, \zeta\rangle \\
&\leq C\delta\langle\zeta,\zeta\rangle + \delta^2\langle A\zeta, A\zeta\rangle.
\end{aligned}$$

The **T2** term looks like

$$\begin{aligned}
&\langle (A+\delta A)\zeta_s, \zeta\rangle \\
=&\langle \zeta_s, (A+\delta A^T)\zeta\rangle \\
=&\langle (A+\delta A)\zeta + e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta), (A+\delta A^T)\zeta\rangle \\
=&\langle A\zeta, A\zeta\rangle \\
&+ \langle \delta A\zeta, A\zeta\rangle + \langle A\zeta, \delta A^T\zeta\rangle \\
&+ \langle e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta), A+\delta A^T\zeta\rangle.
\end{aligned}$$

We can estimate the above using the same analysis as before (δA^T behaves really similarly to δA since we only care about the δ factor in front, technically we will need to take a t derivative of M but in our case this is still upper bounded by δ multiplied by a uniformly bounded matrix). This shows all these terms combine to make the **T2** term satisfy

$$\geq \lambda^2/2g(s).$$

Finally we look at the **T3** term

$$\begin{aligned}
&\left\langle \frac{d}{ds}e^{d(s)}\mathcal{F}_v(e^{-d(s)}\zeta), \zeta\right\rangle \\
=&\left\langle \frac{d}{ds}[(g(e^{-ds}\zeta)\zeta + h(e^{-ds}\zeta)\zeta_t)], \zeta\right\rangle \\
=&\left\langle \frac{d}{ds}(g(e^{-ds}\zeta)\zeta, \zeta\rangle + \langle\frac{d}{ds}h(e^{-ds}\zeta)\zeta_t), \zeta\right\rangle \\
\leq&\epsilon\langle\zeta,\zeta\rangle + \epsilon\langle\zeta_s,\zeta\rangle
\end{aligned}$$

where we used the elliptic estimate $\|\zeta_{st}\|_{C^0} \leq C\epsilon$ and $\|\zeta_t\|_{C^0} \leq C\epsilon$ (technically the version of elliptic regularity in [1] or [48] only applies to $\phi = e^{-ds}\zeta$, but seeing everything we used above is local, that this implies corresponding bounds on ζ is immediate). As before we need to estimate

$$\begin{aligned}&\epsilon\langle\zeta_s,\zeta\rangle\\&=\epsilon\langle A+\delta A\zeta+e^{d(s)}\mathcal{F}_v(e^{-d(s)})\zeta,\zeta\rangle.\end{aligned}$$

The third term in the equation above is easily bounded above by

$$\epsilon\langle\zeta,\zeta\rangle.$$

The first term is bounded by

$$\epsilon(\langle A\zeta,A\zeta\rangle+\langle\zeta,\zeta\rangle).$$

The second term is similarly bounded by

$$\epsilon\delta C(\langle A\zeta,A\zeta\rangle+\langle\zeta,\zeta\rangle)$$

thus the entire **T3** term satisfies

$$\leq C\epsilon(\langle A\zeta,A\zeta\rangle+\langle\zeta,\zeta\rangle)$$

then putting all of these terms together globally we have

$$g''\geq\lambda^2 g$$

for small enough $\epsilon>0$ and this concludes the proof. □

It still requires some work to go from this to exponential decay in the Sobolev spaces we want. The easiest way to do this is to realize our solution ζ has C^0 norm bounded above by $C\epsilon/R$. Hence for small enough ϵ, its C^0 norm always below 1. This means L^2 norm bounds over intervals of form $[k,k+1]\times S^1$ gives rise to L^p norm bounds over this interval. Using a version of elliptic regularity found in Theorem 12.1.5 in [1] (there's a typo in this version) or appendix B of [48], reproduced in Theorem 3.7.11, we conclude the exponential decay bounds can be improved to $W^{k,p}$, which we can then turn to pointwise bounds. We summarize this in the following theorem:

Proposition 3.8.4. *For $s>3R$, $j=0,1,..k$,*

$$\|\nabla^j\zeta\|(s,t)\leq C\|\zeta\|^{\frac{2}{p}}_{W^{2,p}(S^1\times[3R,\infty))}e^{-\lambda(s-3R)}.$$

Note the λ here is not the same as λ from before.

More details of the elliptic bootstrapping argument is written up in Corollary 3.8.7.

Exponential decay w.r.t. p

In this subsection we show the derivative of ζ with respect to p also decays exponentially. To explain the notation, we recall for each p we can use the parallel transport map to transport $\phi(r, a, p, \psi)$ to W. We remove the exponential weights to view them as vector fields:

$$\zeta(p) \in W^{2,p}(v_0^* TM)$$

(we suppress the dependence on r, a, ψ), and for $s > 3R$ they satisfy equations

$$D'(p)\zeta + e^{d(s)}\mathcal{F}_v(p, e^{-d(s)}\zeta(p)) = 0$$

where $D'(p)$ is of the form

$$D'(p) = d/ds - (A(p) + \delta A(p)).$$

As before $A(p)$ and $\delta A(p)$ take the form

$$A(p) = -J_0 d/dt - S(p) + d$$

$$\delta A(p) = \delta M A - N.$$

The nonlinear term takes the form

$$\mathcal{F}_v(\phi) = g(p,\phi)\phi + h(p,\phi)\partial_t \phi$$

where g and h and their p derivatives (uniformly with respect to p) satisfy the assumptions listed in Remark 3.7.12 as well as Proposition 3.3.2.

We know from the above subsection that for each fixed p, the vector field $\zeta(p)$ is exponentially suppressed as $s \to \infty$. In this subsection we show the derivative of this family of vector fields

$$\frac{d}{dp}\zeta(p)$$

is exponentially suppressed as $s \to \infty$, as this will be crucial for our applications in gluing together multiple level cascades. In this subsection we use $\zeta(p)$ to make explicit the dependence on p, and use subscripts ζ_p to denote the partial derivative with respect to p. For this subsection we define

$$p' = p/\epsilon$$

for $\epsilon > 0$ small enough. This ϵ is comparable to the ϵ balls we have chosen (we can take them to be the same), and depends only on the local geometry near the Morse-Bott torus, and is in particular independent of δ. We write everything in terms of p' instead of p. We next differentiate the defining equation for $\zeta(p)$ w.r.t to p':

$$\frac{dD'}{dp'}\zeta(p) + D'(p)\frac{d\zeta(p)}{dp'} = \frac{de^{d(s)}\mathcal{F}_v(p, e^{-d(s)}\zeta(p))}{dp'}.$$

By elliptic regularity we can assume ζ in this region is infinitely differentiable in s, t, and its p' derivative is also infinitely differentiable in s, t. Further the s, t derivatives of $\zeta_{p'}$ are bounded in $W^{1,p}$ norm by $W^{2,p}$ norms of $\zeta_{p'}$ and ζ. Now we observe that

$$\frac{dD'(p)}{dp}\zeta(p) = (MA - N)\zeta(p)$$

because when we are differentiating $D'(p)$ w.r.t. p we are really looking at how the coefficient matrices J_δ, $S(p)$ behave w.r.t. p, and this is determined by the local geometry and hence their variation is uniformly bounded. Hence by our definition of p' we have

$$\frac{dD(p')}{dp'} = \epsilon(MA - N) =: \epsilon B\phi.$$

Recalling the form of $\mathcal{F}_v$:

$$\mathcal{F}_v(p, \phi) = g(p, \phi) + h(p, \phi)\partial_t\phi.$$

Here we have a p dependence on both g and h since we are shifting the local geometry when we change p. Thus

$$e^{ds}\mathcal{F}_v(p, e^{-ds}\zeta) = g(p, e^{-ds}\zeta)\zeta + h(p, e^{-ds}\zeta)\partial_t\zeta$$

Hence the p' derivative of $e^{ds}\mathcal{F}_v(p, e^{-ds}\zeta)$ looks like

$$\begin{aligned}
&\frac{d}{dp'}e^{ds}\mathcal{F}_v(p, e^{-ds}\zeta)\\
=&\epsilon g_1(p, e^{-ds}\zeta)\zeta + \epsilon h_1(p, e^{-ds}\zeta)\partial_t\zeta\\
&+ g(p, e^{-ds}\zeta)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}\\
&+ g_2(p, e^{-ds}\zeta)e^{-ds}\zeta_{p'}\zeta + h_2(p, e^{-ds}\zeta)\partial_t\zeta\zeta_{p'}
\end{aligned}$$

where g_1 and h_1 denote the derivative with respect to its first variable, namely p. The ϵ appears because we are differentiating with p' instead of p. The functions g_1 and h_1 have the same properties as g and h, i.e.

$$g_1(x, y) \leq |x| + |y|$$

and g_1 has uniformly bounded derivatives with respect to each of its variables; similarly for h_1.

g_2 and h_2 are the derivatives of g and h on their second variable. They are just bounded functions whose derivatives are also bounded.

Hence we can write

$$\frac{d}{dp'}e^{ds}\mathcal{F}_v(p, e^{-ds}\zeta) = F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}$$

where

$$F = \epsilon g_1(p, e^{-ds}\zeta)\zeta + \epsilon h_1(p, e^{-ds}\zeta)\partial_t \zeta$$

which essentially behaves like $e^{ds}\mathcal{F}_v(p, e^{-ds\zeta})$, and

$$G(\zeta, \zeta_t) = g(p, e^{-ds}\zeta) + g_2(p, e^{-ds}\zeta)\zeta + h_2(p, e^{-ds}\zeta)\zeta_t.$$

Hence G is obviously bounded pointwise by $C(|\zeta| + |\zeta_t|)$, and the derivatives of G w.r.t s, t are also bounded by the corresponding derivatives of ζ and ζ_t.

So the equation satisfied by $\zeta_{p'}$ is

$$\frac{d}{ds}\zeta_{p'} = A\zeta_{p'} + \delta A\zeta_{p'} + \epsilon B\zeta + F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t \zeta_{p'}$$

The idea is to let

$$g(s) := \langle \zeta, \zeta \rangle + \langle \zeta_{p'}, \zeta_{p'} \rangle$$

and repeat the proof of the previous subsection to show:

Proposition 3.8.5. *$g''(s) \geq \lambda^2 g(s)$.*

Proof. The term involving $\langle \zeta, \zeta \rangle$ behaves exactly the same way. So let's examine

$$\langle \zeta_{p'}, \zeta_{p'} \rangle'' = 2(\langle \zeta_{p's}, \zeta_{p's} \rangle + \langle \zeta_{p'}, \zeta_{p'ss} \rangle).$$

Step 1: The first term looks like

$$\langle \zeta_{p's}, \zeta_{p's} \rangle.$$

This is equal to

$$\langle A\zeta_{p'}, A\zeta_{p'} \rangle + [...].$$

We have as before $\langle A\zeta_{p'}, A\zeta_{p'} \rangle \geq \lambda^2 \langle \zeta_{p'}, \zeta_{p'} \rangle$, we think of the terms in [...] as error terms. We will introduce them one by one and show they are bounded. The first few are of the form (the list continues)

$$\langle A\zeta_{p'}, \epsilon B\zeta \rangle, \quad \langle A\zeta_{p'}, F \rangle, \quad \langle A\zeta_{p'}, G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t \zeta_{p'} \rangle.$$

We shall resume our convention of using **bold** face letters **T1-T3** to refer to the above terms. **T1** can be bounded

$$\begin{aligned}
&\leq \langle \sqrt{\epsilon} A\zeta_{p'}, \sqrt{\epsilon} B\zeta \rangle \\
&\leq \epsilon \langle A\zeta_{p'}, A\zeta_{p'} \rangle + \langle \sqrt{\epsilon} B\zeta, \sqrt{\epsilon} B\zeta \rangle \\
&\leq \epsilon \langle A\zeta_{p'}, A\zeta_{p'} \rangle + \epsilon(\langle A\zeta, A\zeta \rangle + \langle \zeta, \zeta \rangle).
\end{aligned}$$

For **T2**

$$\begin{aligned}
&\leq \langle A\zeta_{p'}, \epsilon\zeta \rangle + \langle A\zeta_{p'}, \epsilon h_1(p, e^{-ds}\zeta)\partial_t \zeta \rangle \\
&\leq \epsilon(\langle A\zeta_{p'}, A\zeta_{p'} \rangle + \langle \zeta, \zeta \rangle) + \epsilon(\langle A\zeta_{p'}, A\zeta_{p'} \rangle + \langle h_1, h_1 \rangle \\
&\leq C\epsilon(\langle A\zeta_{p'}, A\zeta_{p'} \rangle + \langle \zeta, \zeta \rangle).
\end{aligned}$$

T3 is bounded by

$$\begin{aligned}&\le \langle A\zeta_{p'}, \epsilon\zeta_{p'}\rangle + \langle A\zeta_{p'}, h(p, e^{-ds}\zeta)\zeta_{p't}\rangle\\ &\le \epsilon(\langle A\zeta_{p'}, A\zeta_{p'}\rangle + \langle \zeta_{p'}, \zeta_{p'}\rangle + \langle \zeta, \zeta\rangle).\end{aligned}$$

There are actually several ways to bound this term. The easiest way as above is to observe ζ_p has $W^{2,p}$ norm $\le C\epsilon$, hence ζ_{pt} has C^0 norm bounded by $C\epsilon$, and hence $\zeta_{p't}$ has C^0 norm bounded by $C\epsilon$, hence the second term is bounded above by $\epsilon(\langle A\zeta_{p'}, A\zeta_{p'}\rangle + \langle h, h\rangle)$ which implies the overall bound by the form of h.

More terms that also appear in $\langle \zeta_{p's}, \zeta_{p's}\rangle$ are given below:

$$\begin{aligned}&\langle \delta A\zeta_{p'}, \epsilon B\zeta\rangle, \quad \langle \delta A\zeta_{p'}, F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}\rangle,\\ &\langle \epsilon B\zeta, \epsilon B\zeta\rangle, \quad \langle \epsilon B\zeta, F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}\rangle,\\ &\langle F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}, F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}\rangle.\end{aligned}$$

The common feature with all of the above terms is that both inputs into the inner product are small, hence we can bound all of the terms above by

$$\begin{aligned}&\langle \delta A\zeta_{p'}, \delta A\zeta_{p'}\rangle, \quad \langle \epsilon B\zeta, \epsilon B\zeta\rangle, \quad \langle F, F\rangle, \quad \langle G(\zeta, \zeta_t)\zeta_{p'}, G(\zeta, \zeta_t)\zeta_{p'}\rangle,\\ &\langle h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}, h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}\rangle.\end{aligned}$$

Using techniques already established when we considered exponential decay in the previous subsection, we can bound each of these above terms by (respectively)

$$\begin{aligned}&C\delta(\langle A\zeta_{p'}, A\zeta_{p'}\rangle + \langle \zeta_{p'}, \zeta_{p'}\rangle), \quad C\epsilon(\langle A\zeta, A\zeta\rangle + \langle \zeta, \zeta\rangle), \quad C\epsilon(\langle A\zeta, A\zeta\rangle + \langle \zeta, \zeta\rangle),\\ &\epsilon\langle \zeta_{p'}, \zeta_{p'}\rangle, \quad C\epsilon\langle \zeta, \zeta\rangle.\end{aligned}$$

This concludes the first step, in which we bounded all terms appearing in $\langle \zeta_{p's}, \zeta_{p's}\rangle$.

Step 2 We next compute

$$\begin{aligned}&\langle \zeta_{p'ss}, \zeta_{p'}\rangle\\ &= \langle \partial_s((A + \delta A)\zeta_{p'} + \epsilon B\zeta + F + G(\zeta, \zeta_t)\zeta_{p'} + h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}), \zeta_{p'}\rangle\\ &= \left\langle (A' + \delta A')\zeta_{p'} + \epsilon B'\zeta + \frac{d}{ds}F + \frac{d}{ds}(G(\zeta, \zeta_t)\zeta_{p'}) + \frac{d}{ds}(h(p, e^{-ds}\zeta)\partial_t\zeta_{p'}), \zeta_{p'}\right\rangle\\ &+ \langle (A + \delta A)\zeta_{p's} + \epsilon B\zeta_s, \zeta_p\rangle.\end{aligned}$$

We label the above two terms by **T1** and **T2** respectively. We first examine **T1**. In order to make the sizes of various terms more apparent, we shall replace $\frac{d}{ds}F$ with

$$C\zeta^2 + C\zeta\partial_t\zeta + +C\zeta\zeta_s + C\zeta_s\zeta_t + C\zeta\zeta_{ts}$$

where the C as it appears in each of the above terms may be different, but they are all uniformly bounded smooth functions of (s, t) with uniformly bounded derivatives.

Similarly we shall replace $\frac{d}{ds}G(\zeta,\zeta_t)\zeta_{p'}$ with

$$(C\zeta + C\zeta_s + C\zeta^2 + C\zeta\zeta_s + C\zeta\zeta_t + C\zeta_s\zeta_t + C\zeta\zeta_{ts})\zeta_{p'} + (C\zeta + C\zeta^2 + C\zeta\zeta_t)\zeta_{p's}$$

with the same convention on C as before. Finally we shall replace $\frac{d}{ds}h(p,e^{-ds}\zeta)\zeta_{tp'}$ with

$$(C\zeta + \zeta_s)\zeta_{tp'} + C\zeta\zeta_{p'ts}.$$

We examine various components of the **T1** term, starting with

$$\langle (A' + \delta A')\zeta_{p'} + \epsilon B'\zeta, \zeta_{p'}\rangle.$$

The operator $A' + \delta A'$ for our purposes looks like $\epsilon(A+N)$, since the derivatives of the coefficient matrices with respect to p' are bounded by ϵ. Similarly $\epsilon B'$ behaves like ϵB so we have, using these estimates

$$\langle (A' + \delta A')\zeta_{p'} + \epsilon B'\zeta, \zeta_{p'}\rangle \leq C\epsilon(\langle A\zeta_{p'}, A\zeta_{p'}\rangle + \langle \zeta_{p'},\zeta_{p'}\rangle) + \epsilon(\langle \zeta_{p'},\zeta_{p'}\rangle + \langle\zeta,\zeta\rangle + \langle A\zeta, A\zeta\rangle).$$

We next estimate

$$\begin{aligned}
&\left\langle \frac{d}{ds}F, \zeta_{p'}\right\rangle \\
\leq &\langle C\zeta^2 + C\zeta\partial_t\zeta + +C\zeta\zeta_s + C\zeta_s\zeta_t + C\zeta\zeta_{ts}, \zeta_{p'}\rangle \\
\leq &C\epsilon(\langle\zeta,\zeta\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle) + C\epsilon\langle\zeta_t,\zeta_{p'}\rangle
\end{aligned}$$

where we used the fact C^0 norm of $\zeta, \partial_t\zeta, \partial_s\zeta, \zeta_{st}$ are all uniformly bounded by $C\epsilon$ using elliptic regularity. The term $C\epsilon\langle\zeta_t,\zeta_{p'}\rangle$ is bounded by

$$\begin{aligned}
&C\epsilon\langle\zeta_t,\zeta_{p'}\rangle \\
\leq &C\epsilon(\langle\zeta_t,\zeta_t\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle \\
\leq &C\epsilon(\langle A\zeta, A\zeta\rangle + \langle\zeta,\zeta\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle).
\end{aligned}$$

which concludes the estimates for $\langle\frac{d}{ds}F,\zeta_{p'}\rangle$.

We next examine $\left\langle\frac{d}{ds}G(\zeta,\zeta_t)\zeta_{p'},\zeta_{p'}\right\rangle$, which we can bound by

$$\leq \epsilon\langle\zeta_{p'},\zeta_{p'}\rangle + \epsilon\langle\zeta_{p's},\zeta_{p'}\rangle.$$

The second term in the above inequality is in turn bounded by

$$\begin{aligned}
&\leq \epsilon\langle\zeta_{p's},\zeta_{p'}\rangle \\
&\leq \epsilon\langle A\zeta_{p'} + \delta A\zeta_{p'} + \epsilon B\zeta + F + G(\zeta,\zeta_t)\zeta_{p'} + h(p,e^{-ds}\zeta)\partial_t\zeta_{p'},\zeta_{p'}\rangle \\
&\leq \epsilon(\langle A\zeta_{p'},A\zeta_{p'}\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle + \langle\zeta,\zeta\rangle + \langle A\zeta,A\zeta\rangle)
\end{aligned}$$

using techniques of the previous step. This concludes all bounds for $\langle\frac{d}{ds}G(\zeta,\zeta_t)\zeta_{p'},\zeta_{p'}\rangle$.

We next turn to $\left\langle \frac{d}{ds}h(p,e^{-ds}\zeta)\zeta_{tp'},\zeta_{p'}\right\rangle$, which we bound by

$$\begin{aligned}
&\langle (C\zeta + C\zeta_s)\zeta_{tp'} + C\zeta\zeta_{p'ts},\zeta_{p'}\rangle\\
\leq&\epsilon\langle\zeta_{tp'},\zeta_{p'}\rangle + \epsilon\langle\zeta_{p'ts},\zeta_{p'}\rangle\\
\leq&\epsilon(\langle A\zeta_{p'},A\zeta_{p'}\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle + \epsilon\langle\zeta_{p'ts},\zeta_{p'}\rangle.
\end{aligned}$$

To bound $\epsilon\langle\zeta_{p'ts},\zeta_{p'}\rangle$, we use

$$\begin{aligned}
&\epsilon\langle\zeta_{tp's},\zeta_{p'}\rangle\\
&\leq \epsilon\langle\zeta_{p's},\zeta_{p't}\rangle\\
&\leq \epsilon\langle A\zeta_{p'} + \delta A\zeta_{p'} + \epsilon B\zeta + F + G(\zeta,\zeta_t)\zeta_{p'} + h(p,e^{-ds}\zeta)\partial_t\zeta_{p'}),(MA+N)\zeta_{p'}\rangle\\
&\leq \epsilon[\langle A\zeta_{p'},A\zeta_{p'}\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle + \langle\zeta,\zeta\rangle + \langle A\zeta,A\zeta\rangle].
\end{aligned}$$

This concludes the bounds for $\langle\frac{d}{ds}h(p,e^{-ds}\zeta)\zeta_{tp'},\zeta_{p'}\rangle$, and consequently all of **T1**.

We now turn to **T2**. We first examine $\langle\epsilon B\zeta_s,\zeta_{p'}\rangle$. It can be rewritten as

$$\langle\zeta_s,\epsilon B^T\zeta_{p'}\rangle = \langle(A+\delta A)\zeta + e^{d(s)}\mathcal{F}(e^{-d(s)}\zeta),\epsilon B^T\zeta_{p'}\rangle.$$

We recall that $\epsilon B = \epsilon MA + N$. Now in taking the adjoint we had to differentiate the coefficient matrices w.r.t the variable t, but in our case ϵB^T would still take the same form. Hence these terms can be handled by entirely similar techniques as before, giving

$$\leq \epsilon(\langle A\zeta_{p'},A\zeta_{p'}\rangle + \langle A\zeta,A\zeta\rangle + \langle\zeta,\zeta\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle).$$

We consider the remaining term $\langle A + \delta A\zeta_{p's},\zeta_{p'}\rangle$. We can rewrite it as

$$\langle\zeta_{p's},(A+\delta A^T)\zeta_{p'}\rangle.$$

Noting that δA^T essentially takes the same form as δA, the above term will resemble the terms we computed in step 1. Hence it is equal to

$$\langle A\zeta_{p'},A\zeta_{p'}\rangle$$

plus an error term which is uniformly bounded by

$$\epsilon(\langle A\zeta_{p'},A\zeta_{p'}\rangle + \langle\zeta_{p'},\zeta_{p'}\rangle + \langle A\zeta,A\zeta\rangle + \langle\zeta,\zeta\rangle).$$

This gives bounds on all of the terms appearing in $g''(s)$, from which we conclude that

$$g''(s) \geq \lambda^2 g(s)$$

for $\epsilon > 0$ sufficiently small. □

We now switch to trying to understand $\langle\zeta_p,\zeta_p\rangle$, we can get this simply by rearranging terms in $g(s)$ and realizing derivatives w.r.t p versus p' differ by a factor of ϵ.

Corollary 3.8.6.

$$\langle \zeta_p, \zeta_p \rangle_{L^2(S^1)}(s) \leq C \frac{\langle \zeta, \zeta \rangle_{L^2(S^1)}(s_0) + \epsilon^2 \langle \zeta_p, \zeta_p \rangle_{L^2(S^1)}(s_0)}{\epsilon^2} e^{-\lambda(s-s_0)}.$$

for $s > s_0$ (in our case we can take $s_0 = 3R$, we are just stating the corollary more generallly to indicate the decay starts at s_0.)

It might seem unpleasant we are dividing by ϵ^2, but in practice by elliptic regularity (and the ζ term we will be working with) we will have $\langle \zeta, \zeta \rangle \sim C\epsilon^2/R^2$, so the decay really is of the form $C(\epsilon^2 + \frac{1}{R^2})e^{-\lambda(s-s_0)}$. Also in the cases that interest us the decay will be so large factors of size $1/\epsilon^2$ will become irrelevant.

Using the same argument as before ζ_p has $W^{2,p}$ norm of size $C\epsilon$ so our previous strategy of bounding L^p norm with L^2 norm continues to work, so we obtain the bound:

Corollary 3.8.7. *For $s > s_0 > 3R$, we have*

$$|\zeta_p(s,t)| \leq C \left[\frac{(\|\zeta\|^2_{W^{2,p}} + \|\zeta_p\|^2_{W^{2,p}})}{\epsilon^2} \right]^{\frac{1}{p}} e^{-\lambda(s-s_0)}.$$

Here λ is different from the λ we chose previously. We will abbreviate this by writing $|\zeta_p(s,t)| \leq Ce^{-\lambda(s-s_0)}$ as some more careful estimates can show the coefficient in front to be of order $O(1)$, similarly using elliptic regularity we can bound

$$|\zeta_{p*}(s,t)| \leq Ce^{-\lambda(s-s_0)}, \quad * = s, t \text{ and higher derivatives.}$$

Proof. For completeness we explain how elliptic regularity is used. First using $W^{2,p} \hookrightarrow C^0$, we have

$$\langle \zeta_p, \zeta_p \rangle_{L^2(S^1)}(s) \leq C \frac{\|\zeta\|^2_{W^{2,p}} + \|\zeta_p\|^2_{W^{2,p}}}{\epsilon^2} e^{-\lambda(s-s_0)}.$$

Using the fact C^0 norm of ζ_p is < 1, we have

$$\|\zeta_p\|^p_{L^p([k-1,k+2]\times S^1)} \leq C \int_{k-1}^{k+2} \langle \zeta_p, \zeta_p \rangle_{L^2(S^1)}(s) ds \leq \frac{\|\zeta\|^2_{W^{2,p}} + \|\zeta_p\|^2_{W^{2,p}}}{\epsilon^2} e^{-\lambda k}.$$

Given this L^p norm bound, we can use elliptic regularity and the fact ζ_p satisfies an equation of the form

$$D'\zeta_p = B\zeta + F + G(\zeta, \zeta_t)\zeta_p + h(p, e^{-ds}\zeta)\partial_t \zeta_p.$$

Here we are differentiating with respect to p instead of p' so we are rescaling some of the above terms so that they have norm $O(1)$ instead of $O(\epsilon)$.

From elliptic bootstrapping we have

$$\|\zeta_p\|_{W^{1,p}([k,k+1]\times S^1)} \leq C\|\zeta_p\|_{L^p([k-1,k+2]\times S^1)} + \|\zeta\|_{W^{1,p}([k-1,k+2]\times S^1)} \leq Ce^{-\lambda k}$$

where we used the exponential decay estimate of ζ. Note we have slightly shrunk the domain to $[k, k+1] \times S^1$ to use elliptic regularity. We can iterate this argument to show

$$\|\zeta_p\|_{W^{l,p}([k,k+1]\times S^1)} \leq C_l e^{-\lambda k}$$

and use Sobolev embedding theorems to obtain pointwise bounds as in the proposition. $\square$

We also note the could have used the exact same techniques when applied to the r asymptotic vector. There we need to identify $r \in S^1 = [0,1]/\sim$, and $r' := r/\epsilon \in S^1 = [0, 1/\epsilon]/\sim$. The result is very similar: we can obtain exponential decay bounds on ζ_r, given as:

Corollary 3.8.8. *For $s > s_0 > 3R$, we have*

$$|\zeta_r(s,t)| \leq C \left[\frac{(\|\zeta\|^2_{W^{2,p}} + \|\zeta_r\|^2_{W^{2,p}})}{\epsilon^2}\right]^{\frac{1}{p}} e^{-\lambda(s-s_0)}$$

$$|\zeta_{r*}(s,t)| \leq Ce^{-\lambda(s-s_0)}, \quad * = s,t \textit{ and higher derivatives.}$$

We do not need such a result for the a-asymptotic vector since the geometry is invariant in the a direction.

3.9 Gluing multiple-level cascades

We have assembled all the technical ingredients we need to do gluing, which we take up in this section. We note gluing together cascades with finite gradient trajectories is substantially harder than semi-infinite gradient trajectories. We start with a simplified setup of gluing together 2-level cascades, which captures most of the technical difficulty. The generalization to n level cascades is then a problem of linear algebra.

Our simplified setup is this: let $u_\pm : \Sigma_\pm \to (M, J)$ be two rigid (nontrivial) J-holomorphic curves. u_+ has one negative end, asymptotic to Reeb orbit γ_+; and u_- has one positive end, asymptotic to Reeb orbit γ_-. Both of these ends are on the same Morse-Bott torus, and in fact they are connected by a gradient trajectory of length T. We will perturb J to J_δ near this Morse-Bott torus (and nowhere else), and glue $u_\pm$ along with the finite gradient trajectory into a J_δ-holomorphic curve. This construction ignores the other ends of $u_\pm$, which we assume to remain on other Morse-Bott tori, and we only have two levels. The reason for this is that the process of gluing together two J-holomorphic curves along a finite gradient trajectory is very technical, and we would like to carry out the heart of the technical construction with as little extra baggage as possible.

This section is organized as follows: we first introduce the general setup and the process of pregluing. We, as before, show gluing can be realized by solving a system of three equations. We then proceed to discuss the linear theory required to describe the linearization of $\bar{\partial}_{J_\delta}$ over the finite gradient trajectory. After that we show the feedback terms coming from Θ_v

and going into $\Theta_\pm$ (defined in the subsection below) depend nicely on the input - this is the most technical step and will take some careful estimates. Finally after this we will be able to solve the three equations as we did in the previous section. Finally, we explain the generalization to n-level transverse index one cascades.

Setup and pregluing

Recall near the Morse-Bott torus we have coordinates $(z, x, y) \in S^1 \times S^1 \times \mathbb{R}$. For definiteness we assume γ_+ is the Reeb orbit with x coordinate x_+ and γ_- is at x_-. To simplify notation we assume we have rescaled the x coordinate so that $f(x) = x + C$ over the interval on S^1 connecting x_- and x_+.

We recall for each $u_\pm$ we choose a cylindrical neighborhood around each of its punctures $(s, t) \in S^1 \times (0, \pm\infty)$. We also recall near our punctures $u_\pm$ has the coordinate form

$$u_\pm = (a_\pm, z_\pm, x_\pm, y_\pm).$$

We assume $u_+(s = -\infty, t) \to (-\infty, t, x_+, 0)$ and $u_-(s = \infty, t) \to (\infty, t, x_-, 0)$. For each $u_\pm$ we describe a neighborhood of this map as

$$W^{2,p,d}(u_\pm^* TM) \oplus T\mathcal{J}_\pm \oplus V'_\pm \oplus V_\pm$$

where $W^{2,p,d}(u_\pm^* TM)$ is the weighted vector space of vector fields with weight $e^{\pm ds}$ at positive/negative punctures. We use $T\mathcal{J}_\pm$ to denote a Teichmuller slice. We use $V'_\pm$ to denote asymptotic vectors at other ends of $u_\pm$, and $V_\pm$ is the end that we are considering, being a 3 dimensional space consisting of vectors $(r, a, p)_\pm$.

We recall the important gluing constant

$$R := \frac{1}{5d}\log(1/\delta)$$

which we think of our gluing parameter.

Let v_δ be a gradient trajectory suitably translated so that over the interval $s \in [0, T/\delta]$, the map v_δ corresponds to the gradient flow that connects $\gamma_\pm$, in particular this means the x component of v_δ satisfies

$$x \text{ component of } v_\delta(R) = x_-$$

$$x\,\text{component of}\,v_\delta(T/\delta - R) = x_+.$$

We next construct our preguling, similar to the semi-infinte case our pregluing will depend on our asymptotic vectors $(r, a, p)_\pm$.

Given fixed $(r, a, p)_\pm$, let $v_{r,a,p} = (a_v(s), t_v(t), x_v(s), 0)$ (we suppress the $\pm$ that should appear in the subscript to ease the notation) denote the suitably translated gradient trajectory, so that when restricted to $s \in [0, T_p/\delta]$ satisfies

$$v_{r,a,p}(T_p/\delta - R, t) = (a_+(-R, 0) + a_+, t + r_+, x_+ + p_+, 0)$$

and
$$v_{r,a,p}(R,t) = (a_v(R), t + r_+, x_- + p_-, 0).$$
We observe that due to the form of f in this region, we have $T_p = T + (p_+ - p_-)$. We preglue this gradient trajectory to the deformed curve $u_+ + (r,a,p)_+$ at $s = T_p/\delta - R$ of $v_{r,a,p}$. This value of s over $v_{r,a,p}$ is identified with $s_+ = -R$ over u_+. At the other end we consider u_- translated in a direction so that $a_-(R) = a_v(R) - a_-$. Then we would like to preglue $v_{r,a,p}$ at $s = R$ to $u_- + (r,a,p)_-$ at $s_- = R$, except there is an issue that since r_+ is in general different from r_-, the curve $v_{r,a,p}(-,t)$ has z component $t + r_+$, while $u_-(s,t) + (r,a,p)_-$ has t component roughly equal to $t + r_-$. To remedy this we need to preglue with a different domain $\Sigma_{r,a,p}$ so that at $s = T_p/\delta - R$ we do our usual pregluing (as in the semi-infinite gradient trajectory case), but at $s_- = s = R$ we glue with a twist: recall $(s_-, t_-) \in \mathbb{R} \times S^1$ is a cylindrical neighborhood on u_- and (s,t) is the usual coordinate on $v_{r,a,p}$, then we construct the domain $\Sigma_{r,a,p}$ by identifying $t_- + r_- \sim t + r_+$ at $s = s' = R$. Then we can construct a preglued map
$$u_{r,a,p} : \Sigma_{r,a,p} \longrightarrow (M, J)$$
that depends on the asymptotic vectors $(r,a,p)_\pm$.

Remark 3.9.1. We first observe that here the domain depends non-trivially on the asymptotic vectors $(r,a,p)_\pm$, in fact in the case where the domains for $u_\pm$ are stable, changing the pregluing in $r_\pm$, i.e. "twisting", or changing the length of the cylindrical neck by changing $p_\pm$ or $a_\pm$ correspond to changing the complex structure of the domain curve.
We also observe here that if we change $a_\pm$ by size ϵ, then the length of cylindrical length changes by size ϵ. Similarly if we change $r_\pm$ by ϵm in some appropriate sense the complex structure changes within an ϵ neighborhood. However when we change $p_\pm$ by ϵ, the length of the neck changes by ϵ/δ. This is in some sense the main source of difficulty in studying this degeneration. Since $\delta << \epsilon$ they operate on different scales, and care must be taken to ensure all the vectors we encounter have the right sizes.

Remark 3.9.2. Here because we only have one end the pregluing is rather simple, when there are multitple ends and/or when we talk about degeneration into cascades more care must be taken to into pregluing, which we defer to subsection 3.9.

Linear theory over $v_{r,a,p}$

In this subsection we take a detour to study the linearization of $\bar{\partial}_{J_\delta}$ over $v_{r,a,p}$. In particular we find a suitable Sobolev space with suitable exponential weights so that for given $(r,a,p)_\pm$, the said linearization denoted by D_{J_δ} is surjective with uniformly bounded right inverse as $\delta \to 0$.
After fixing $(r,a,p)_\pm$, we consider
$$D_{J_\delta} : W^{2,p,w_p}(v_{r,a,p}^* TM) \longrightarrow W^{1,p,w_p}(v_{r,a,p}^* TM).$$

Here w_p is a piecewise linear function that is zero at $s = 0$ and $s = T_p/\delta$, has a peak at $s = T_p/2\delta$, and has slope $\pm d$. Explicitly it is given by

$$w_p = -|d(s - T_p/2\delta)| + dT_p/2\delta.$$

It looks like an inverted V. The space $W^{2,p,w_p}(v^*_{r,a,p}TM)$ is a weighted Sobolev space with exponential weight $e^{w_p(s)}$. As is with the case for semi-infinite ends these vector fields have exponential growth as $s \to \pm\infty$, but we do not care about them because those regions do not make an appearance in our construction.

Remark 3.9.3. Observe with our choice of $w_p(s)$, which we sometimes denote by $w(s)$ for brevity, over the preglued curve $u_{r,a,p}$, the pregluing takes place at $s = R$ and $s = T/\delta - R$, and at these two values of s where the pregluing takes place, the exponential weight profile of $v_{r,a,p}$ agrees with the exponential weight profile over $u_\pm$.

Theorem 3.9.4. *D_{J_δ} as defined above is surjective of index 3. It has a uniformly bounded right inverse as $\delta \to 0$.*

Proof. We can view D_{J_δ} as the gluing of two operators D_1 and D_2. The operators D_i are both defined over $W^{2,p,w_i}(v^*_{r,a,p}TM)$, except they use different exponential weights. We let $w_1(s) = d(T_p/\delta - s)$, and $w_2(s) = ds$ We glue D_1 and D_2 together at $s = T_p/2\delta$ to recover D_{J_δ}. By results in Section 3.6, D_i are both surjective with uniformly bounded right inverse Q_i, hence as before we can construct approximate right inverse of D_{J_δ} via $Q_1\#Q_2$, hence D_{J_δ} is surjective with uniformly bounded right inverse as $\delta \to 0$.

The index computation is done by conjugating to $W^{2,p}(v^*_{r,a,p}TM)$ via multiplication by $e^{w_p(s)}$. There we observe by shape of $w_p(s)$ there are 3 eigenvalues that cross 0 as s goes from $-\infty$ to ∞, hence by spectral flow this operator has index 3. □

We now proceed to describe the kernel of D_{J_δ} and a codimensional 3 subspace H_0 of its domain so that $D_{J_\delta}|_{H_0}$ is an isomorphism with uniformly bounded inverse as $\delta \to 0$. This will be crucial for us when we try to solve equations over $W^{2,p,w}(v^*_{r,a,p}TM)$.

Consider the vector fields

$$\partial_z, \partial_a \in W^{2,p,w_p}(v^*_{r,a,p}TM).$$

They are asymptotically constant, but they live in $W^{2,p,w_p}(v^*_{r,a,p}TM)$ because as $|s| \to \infty$ the Sobolev norm is exponentially suppressed (written as is they still have very large norm, of order $e^{\frac{dT_p}{2\delta}}$.) Also observe they live in the kernel of D_{J_δ}. Recall from the differential geometry section

$$v_*\partial_s = e^{\delta f(x(s))}\partial_a + \delta f'(x)\partial_x.$$

This vector field also lives in the kernel of D_{J_δ}, and is linearly independent of $\{\partial_z, \partial_a\}$, we modify it to have more palatable form. Consider

$$\frac{v_*\partial_s - \partial_a}{\delta} = [e^{\delta f(x(s))} - 1]/\delta\partial_a + f'(x)\partial_x.$$

This still lives in the kernel of D_{J_δ}, and we see from Taylor expansion that the coefficient in front of ∂_a is bounded above as $\delta \to 0$. We defined the vector field ∂_v to be $a\frac{v_*\partial_s - \partial_a}{\delta} + b\partial_s$ where a, b are constants (both of order 1, bounded above and away from 0) chosen so that $\partial_v(s = T_p/2\delta, t) = \partial_x$. Thus the kernel of D_{J_δ} is spanned by $\{\partial_z, \partial_a, \partial_v\}$. We construct a complement of this space. Consider the linear functionals $L_*, * = z, a, v : W^{2,p,w}(v^*_{r,a,p}TM) \to \mathbb{R}$ defined by

$$L_* : \phi \in W^{2,p,w}(v^*_{r,a,p}TM) \longrightarrow \int_0^1 \langle \phi(s,t), \partial_* \rangle dt \in \mathbb{R}.$$

We define the complement subspace of $kerD_{J_\delta}$, which we write as H_0, via

$$H_0 := \{\phi \in W^{2,p,w}(v^*_{r,a,p}TM) | L_*(\phi) = 0, * = z, a, v\}.$$

We next show:

Proposition 3.9.5. *The projection map*

$$\Pi : W^{2,p,w}(v^*_{r,a,p}TM) \longrightarrow H_0$$

has uniformly bounded norm as $\delta \to 0$. The map Π also commutes with D_{J_δ}.

Proof. We first observe Π is defined by

$$\Pi(\phi) = \phi - \sum_* L(\phi)\partial_*.$$

We now estimate the norm of this operator. By the Sobolev embedding theorem

$$W^{2,p,w}(v^*_{r,a,p}TM) \hookrightarrow C^0(v^*_{r,a,p}TM).$$

In view of the fact we have exponential weights, we have the upper bound

$$L_*(\phi) \leq Ce^{-dT_p/2\delta} \|\phi\|_{W^{2,p,w}}.$$

Hence to estimate the norm of Π it suffices to calculate

$$\begin{aligned}
&\frac{\|L_*(\phi)\partial_*\|}{\|\phi\|} \\
\leq &Ce^{-dT_p/2\delta}\|\partial_*\| \\
\leq &Ce^{-dT_p/2\delta}\left[\int_0^{T_p/2\delta} e^{ds} ds + \int_{-\infty}^0 e^{ds} ds\right] \\
\leq &Ce^{-dT_p/2\delta}\frac{(e^{dT_p/2\delta})}{d} \leq C.
\end{aligned}$$

We only integrated from $(-\infty, T/2\delta)$ because the integral over $(T/2\delta, \infty)$ takes the same form. And hence we see readily the operator norm of Π is uniformly bounded above independently of δ.

The fact that Π commutes with D_{J_δ} follows from the fact Π subtracts off elements that are in the kernel of D_{J_δ}.

Hence we conclude $\Pi \circ Q$ is a uniformly bounded inverse to D_{J_δ} restricted to H_0. □

Deforming the pregluing

Recall that given a pair of asymptotic vectors over $u_\pm$, which we denote by $(r,a,p)_\pm$, we constructed a preglued map $u_{r,a,p} : \Sigma_{r,a,p} \to M$. Next given vector fields with exponential decay, $\psi_\pm \in W^{2,p,d}(u_\pm^* TM)$, and $\phi \in W^{2,p,w}(v_{r,a,p}^* TM)$, we use them to deform $u_{r,a,p}$. Technically the space of deformations of $u_\pm$ also includes $T\mathcal{J}_\pm \oplus V'_\pm$, but we suppress them from our notation because these deformations happen away from the region where the pregluing takes place. For $s \in [R, T_p/\delta - R]$ considered over $v_{r,a,p}$, we define the cut off functions

$$\beta_- = \beta_{[-\infty,2R;R/2]}$$

$$\beta_+ := \beta_{[R/2;T_p/\delta-2R,\infty]}$$

$$\beta_v := \beta_{[R/2;R,T_p/\delta-R;R/2]}.$$

We would like to deform $u_{r,a,p}$ by $\beta_+\psi_+ + \beta_-\psi_- + \beta_v\phi$, however there is one subtlety that when we constructed $\Sigma_{r,a,p}$ there was a twist at $s = R$ when we identified $t_- + r_- \sim t + r_+$ when we glued $v_{r,a,p}$ with u_-. Since β_- cuts off ψ_- within the interior of $v_{r,a,p}$, the only effect of this is that when we view the equation over $v_{r,a,p}$ instead of seeing $\psi_-(s,t)$, the term we see is $\psi(s, t + (r_+ - r_-))$. Aside from this point, as before we can add the vector field $\beta_+\psi_+ + \beta_-\psi_- + \beta_v\phi$ to $u_{r,a,p}$, and apply the $\bar\partial_{J_\delta}$ operator. Using the same "splitting up the equations" trick as we did for semi-infinite trajectories we get:

Proposition 3.9.6. *The deformed curve $u_{r,a,p} + \beta_+\psi_+ + \beta_-\psi_- + \beta_v\phi$, where $\psi_\pm \in W^{2,p,d}(u_\pm^* TM) \oplus T\mathcal{J} \oplus V'_\pm$ implicitly includes the variations of complex structure away from the gluing region, is J_δ-holomorphic iff the following 3 equations are satisfied*

$$\Theta_v(\phi, \psi_\pm) = 0$$

$$\Theta_\pm(\phi, \psi_\pm) = 0$$

where Θ_v is of the form

$$D_{J_\delta}\phi + \beta'_\pm\psi_\pm + \mathcal{F}_v(\phi, \psi_\pm).$$

Here $\mathcal{F}_v$ is of the same form as semi-infinite case (except at the end near $s = R$ we see effects of ψ_- and near $s = T_p/\delta - R$ we see the effect of ψ_+). The equations $\Theta_\pm$ take the form

$$\Theta_+ = D_J\psi_+ + D_J(r,a,p)_+ + \mathcal{F}_+ + \mathcal{E}_+ + \beta'_v\phi$$

$$\Theta_- = D_J\psi_- + D_J(r,a,p)_- + \mathcal{F}_- + \mathcal{E}_- + \beta'_v\phi$$

where the scripted expressions $\mathcal{F}_\pm, \mathcal{E}_\pm$ taking the same form as they did in the semi-infinite case. Implicit in the above notation is also the variation of the domain complex structure $u_\pm$, which we denote by $\delta j_\pm$ when we need to make them explicit.

Solving the equations $\Theta_\pm, \Theta_v$

Preamble

In this very lengthy subsection we show the system $\Theta_\pm = 0, \Theta_v = 0$ has a solution with nice properties. Since this is a long process we give a preamble:

- We first show as before given fixed tuple of input data $(\psi_\pm, (r, a, p)_\pm)$ there exists a unique solution $\phi(r, a, p, \psi_\pm) \in H_0$ to Θ_v.
- Then we verify that when we vary the input, $\psi_\pm$ the solution ϕ behaves nicely (in the sense that its input into the equations into $\Theta_\pm$ varies differentiably, as was the case for the gluing of semi-infinite gradient trajectories.)
- Then we verify as we change $p_\pm$ the solution is well behaved. This is the crux of the matter, because when we vary $p_\pm$ what is actually happening is that we are drastically changing the pregluing by dramatically lengthening/shortening the length of the neck. We do this via the following process:
  - We make sense of what it means for ϕ to be well behaved when we vary $p_\pm$.
  - We translate Θ_v into the vector space $W^{2,p}(v^*TM)$ by removing exponential weights.
  - We write the solution ϕ as a sum of two terms: an approximate solution $\gamma_+\zeta_+ + \gamma_-\zeta_-$ to Θ_v that behaves nicely when we vary $p_\pm$ and a correction to this approximate solution $\delta\zeta$, we show $\delta\zeta$ is extremely small. Here $\gamma_\pm$ are cut off functions (and definitely not Reeb orbits).
  - We consider the behaviour of $\delta\zeta$ as we vary $p_\pm$. We consider two ways $p_\pm$ can vary called "lengthening/stretching" and "translation". We show $\delta\zeta$ varies nicely with $p_\pm$, hence the entire solution ϕ varies nicely with $p_\pm$.
- We finally show as a much easier step ϕ varies nicely with (r, a).
- Using all of the above steps, we solve $\Theta_\pm$ with the contraction mapping principle.

Solution to Θ_v

Proposition 3.9.7. *For $\epsilon > 0$ sufficiently small, for all $\delta > 0$ sufficiently small, for fixed tuple $(\psi_\pm, (r, a, p)_\pm)$ with norm less than $\epsilon > 0$, there exists a unique solution $\phi(\psi_\pm, r, a, p_\pm) \in H_0$ to $\Theta_v = 0$ of size $C\epsilon/R$. Moreover the regularity of ϕ can be improved to $W^{3,p,w}(v^*_{r,a,p}TM)$ with its norm similarly bounded above by $C\epsilon/R$.*

Proof. Let Q denote the uniformly bounded right inverse to D_{J_δ}. Consider $\Pi \circ Q : W^{1,p,w}(v^*_{r,a,p}TM) \to H_0$. We observe this operator has uniformly bounded norm as $\delta \to 0$. Further we claim this is an inverse to $D_{J_\delta}|_{H_0}$. To see this first oberseve $D_{J_\delta}|_{H_0}$ is an isomorphism, as it has the same image as $D_{J_\delta}|_{W^{2,p,w}(v^*_{r,a,p}TM)}$ and has index 0. Hence it suffices to show $\Pi \circ Q$ is a right inverse for $\phi \in H_0$. This follows from

$$D_{J_\delta}\Pi \circ Q\phi = \Pi(\phi) = \phi.$$

Hence we consider the map $I : H_0 \to H_0$ defined by

$$I(\phi) = \Pi \circ Q(-\beta'_\pm \psi_\pm - \mathcal{F}_v(\phi, \psi_\pm)).$$

(For ease of notation we will write $\psi_\pm$ when both ψ_+ and ψ_- appear in similar ways). It is apparent that a solution $\phi \in H_0$ to Θ_v is equivalent to a fixed point of $I(\phi)$. We show that a fixed point in an epsilon ball $B_\epsilon \in H_0$ exists and is unique via the Banach contraction mapping principle. Since $\psi_\pm$ has norm $\leq \epsilon$, we have $I(\phi) \leq C(\epsilon/R + C\epsilon^2)$ hence it sends B_ϵ to itself. That I satisfies the contraction property follows from the fact $\mathcal{F}_v$ is quadratic in $\phi, \psi_\pm, \partial_t\phi, \partial_t\psi_\pm$, as well as the fact $\|\psi_\pm\| \leq \epsilon$. Hence it follows from contraction mapping principle there exists unique $\phi(\psi_\pm, r, a, p_\pm)$ solving Θ_v in B_ϵ. We can use the equation itself to estimate the size of ϕ as before and get the size estimate of $C\epsilon/R$. The improvement to $W^{3,p,w}$ and its norm bound follows from elliptic regularity. □

How $\phi(\psi_\pm, r, a, p)$ varies with $\psi_\pm$

For fixed $(r, a, p)_\pm$ we consider the variation of $\phi(\psi_\pm, (r, a, p)_\pm)$ as above with respect to $\psi_\pm$. As we recall from the above the expression $\frac{d\phi}{d\psi_\pm}$ is a linear operator $W^{2,p,w}(u^*_\pm TM) \to W^{2,p,w}(v^*_{r,a,p}TM)$ and has its sized measured via the operator norm. When we write below $\beta'_\pm$ we really mean the multiplication map operating between Sobolev spaces. As in the case of semi-infinite gradient trajectories, we have:

Proposition 3.9.8. $\left\|\frac{d\phi}{d\psi_\pm}\right\|_{W^{2,p,w}(u^*_\pm TM)\to W^{2,p,w}(v^*_{r,a,p}TM)} \leq C\epsilon.$

Proof. Consider the fixed point equation

$$\phi = \Pi \circ Q(-\beta'_\pm \psi_\pm - \mathcal{F}_v(\phi, \psi_\pm)).$$

We differentiate both sides w.r.t $\psi_\pm$ to get (see Remark 3.7.18 for this kind of operation)

$$\frac{d\phi}{d\psi_\pm} = -\Pi \circ Q(\beta'_\pm + \frac{d}{d\psi_\pm}\mathcal{F}_v).$$

Since we know $\mathcal{F}_v$ is a polynomial expression of $\psi_\pm, \phi, \partial_t\phi, \psi_\pm$, we can bound (norm wise)

$$\left\|\Pi \circ Q\frac{d}{d\psi_\pm}\mathcal{F}_v\right\| \leq C(\|\phi\| + \|\psi_\pm\|) + C\epsilon\left\|\frac{d\phi}{d\psi_\pm}\right\|$$

where in the above equation, the norm for $\|\Pi \circ Q \frac{d}{d\psi_\pm}\mathcal{F}_v\|$ and $\|\frac{d\phi}{d\psi_\pm}\|$ are operator norms, and $\|\phi\|$ and $\|\psi_\pm\|$ are $W^{2,p,w}$ norms.
Since $\psi_\pm$ have C^1 bounds of size $\leq C\epsilon$, we can move the term $d\phi/d\psi_\pm$ to the left and get

$$(1 - C\epsilon)\left\|\frac{d\phi}{d\psi_\pm}\right\| \leq C(1/R)$$

which implies our conclusion. □

Variation of ϕ w.r.t. $p_\pm$

In this subsubsection we study the variation of ϕ w.r.t. $p_\pm$. When we change $p_\pm$, we are considerably changing the pregluing. So we need to make sense of what kind of result that we want. We recall from previous section we already found a solution to Θ_v in H_0 for every choice of $(\psi_\pm, (r, a, p)_\pm)$, so our next order of business is to solve $\Theta_\pm$, and in order to do that we need to show as we vary $p_\pm$, the part of ϕ that enters into equations $\Theta_\pm$ varies nicely w.r.t. $p_\pm$. We recall $\Theta_\pm$ is an equation defined over $u_\pm^* TM$. What is happening is as we vary $p_\pm$, the maps $u_\pm$ are translated further/closer to each other, but since our equations are invariant in the symplectization direction, we can identify all those translates of $u_\pm$ and consider one set of equations Θ_+, Θ_- as we vary $p_\pm$. Thus we need to understand how ϕ behaves near the pregluing region. We make this a definition.

Definition 3.9.9. *Let $s \in [-3R, 3R]$, recall if we let $s_\pm$ denote coordinates near the cylindrical neighborhoods of punctures of $u_\pm$, then we have identified $s \sim s_-$ and $s \sim -s_+ + T_p/\delta$. Then for $s \in [-3R, 3R]$ (resp. $[-3R+T_p/\delta, 3R+T_p/\delta]$), the vector field $\phi(s,t)$ can be viewed as a vector field in $W^{2,p,d}(u_-^*TM)$ (resp. $W^{2,p,d}(u_+^*TM)$), as we noted in the pregluing section. We say $\phi(\psi_\pm, r, a, p)$ is well behaved w.r.t. $p_\pm$ if over $s \in [-3R, 3R]$,*

$$\left\|\frac{d}{dp_\pm}\phi(s + T_p/\delta, t)\right\| \leq C\epsilon$$

and

$$\left\|\frac{d}{dp_\pm}\phi(s, t)\right\| \leq C\epsilon$$

where $\frac{d\phi}{dp_\pm}$ is viewed as a vector field over $W^{2,p,d}(u_\pm^*TM)$, and the norm is the weighted Sobolev norm in $W^{2,p,d}(u_\pm^*TM)$.

Remark 3.9.10. Actually because no derivatives of ϕ appears in $\Theta_\pm$, only the $W^{1,p}$ norm is enough for our purposes.

The main theorem of this subsubsection is then:

Proposition 3.9.11. *ϕ is well behaved w.r.t. $p_\pm$.*

To do this we need to very carefully analyze the solutions to Θ_v. It turns out it is not so convenient to analyze this equation with exponential weights, because the weights themselves depend on $p_\pm$. So we first remove the exponential weights via conjugation. We use the following convention:

$$\zeta := e^{w(s)}\phi,$$

$$\psi'_\pm := e^{w(s)}\psi_\pm.$$

The exponential weights are removed and Θ_v is rewritten using the following diagram:

$$\begin{tikzcd} W^{2,p}(v^*_{r,a,p}TM) \arrow[r, "\Theta'_v"] \arrow[d, "e^{-w(s)}"] & W^{2,p}(v^*TM) \\ W^{2,p,d}(v^*_{r,a,p}TM) \arrow[r, "\Theta_v"] & W^{2,p,d}(v^*_{r,a,p}TM). \arrow[u, "e^{w(s)}"'] \end{tikzcd} \tag{3.17}$$

Then the equation Θ_v can be rewritten as

$$\Theta'_v := D'_{J_\delta}\zeta + \beta'_\pm\psi'_\pm + e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta, e^{-w(s)}\psi'_\pm) = 0$$

where $\zeta \in H'_0 \subset W^{2,p}(v^*_{r,a,p}TM)$. We use H'_0 to denote the subspace in $W^{2,p}(v^*_{r,a,p}TM)$ corresponding to H_0. To better understand ζ, let us focus our attention near $s \in [0, T_p/2\delta]$. For this range of s, the equation Θ_v is exactly the same equation as we had solved for semi-infinite gradient trajectories since we do not see the effects of ψ_+. Then by previous result we have a (uniquely constructed) solution $\phi_- \in W^{2,p,d}(v^*_{r,a,p}TM)$ for $s \in [0, T_p/2\delta]$ subject to exponential weight e^{ds} (which for our range of s agrees with $e^{w(s)}$). Defining

$$\zeta_- := e^{w(s)}\phi_-$$

we see ζ_- is a solution to Θ'_v for $s \in [0, T_p/2\delta]$.

There is a slight subtlety in that near u_- there is a twist in the t coordinate as we constructed the pregluing domian, $\Sigma_{r,a,p}$. By the construction in the semi-infinite gradient trajectory case, ϕ_- should depend on input variables (s_-, t_-), which we write as $\phi_-(s_-, t_-)$, but when we view it as a vector field over $v^*_{r,a,p}TM$, using coordinates (s,t) it should be written as $\phi_-(s, t + (r_+ - r_-))$. This won't make a difference for us as we consider variations in the (p_-, p_+) direction, and for the most part we will suppress the t coordinate for brevity of notation. We will take up variations in the (r_+, r_-) variables after considerations of $p_\pm$.

We similarly construct ζ_+. The point is:

Proposition 3.9.12. *$\zeta_\pm$ is well behaved w.r.t. p. i.e. the part of $\zeta_\pm$ that enters into $\Theta_\pm$ has derivative w.r.t. $p_\pm$ bounded above by $C\epsilon$.*

Proof. This follows from our results on $\phi_\pm$ when we proved this property for semi-infinite trajectories. $\square$

The next step is to actually construct ζ from approximate solutions $\zeta_\pm$. Consider the cut off functions $\gamma_\pm$ defined by

$$\gamma_+ := \beta_{[\infty, T_p/2\delta-1;1]}$$
$$\gamma_+ := \beta_{[-\infty;T_p/2\delta-1;1]}.$$

Then we consider the approximate solution

$$\gamma_+\zeta_+ + \gamma_-\zeta_-.$$

We also observe by construction that $\gamma_+\zeta_+ + \gamma_-\zeta_- \in H_0'$. We plug this into Θ_v', we observe by definition this produces zero for all s except $s \in [T_p/2\delta - 2, T_p/2\delta + 2]$. In this interval the Θ_v' takes the form:

$$D'_{J_\delta}(\gamma_+\zeta_+ + \gamma_-\zeta_-) + e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_\pm\zeta_\pm) \tag{3.18}$$

which equals

$$E := \sum_\pm (\gamma'_\pm\zeta_\pm + \gamma_\pm D'_{J_\delta}\zeta_\pm) + e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_\pm\zeta_\pm).$$

Observe $D'_{J_\delta}\zeta_\pm = -e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta_\pm)$ so the error term takes the form

$$\begin{aligned} E =& \gamma'_+\zeta_+ + \gamma'_-\zeta_- \\ &+ [e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_+\zeta_+) - \gamma_+e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta_+)] \\ &+ [e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_-\zeta_-) - \gamma_-e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta_-)]. \end{aligned}$$

We can estimate the size of this term (say in C^1 norm), by elliptic regularity it is easily bounded by the $W^{2,p}$ norm of $\zeta_\pm$ restricted to $s \in [T_p/2\delta - 2, T_p/2\delta + 2]$. (We actually see t derivatives of $\zeta_\pm$ in $\mathcal{F}_v$ but this is fine, we can bound them by elliptic regularity). We know the norm of $\zeta_\pm$ undergoes exponential decay as s moves into this center region, so the size of the error term is bounded above by

$$C\max\{\|\zeta_+\|, \|\zeta_-\|\}^{2/p}e^{-\lambda(T/2\delta-3R)}$$

where in the above equation $\|\zeta_\pm\|$ denotes the full norm of $\zeta_\pm$ over $W^{2,p}(v^*_{r,a,p}TM)$, or equivalently the norm of $\phi_\pm \in W^{2,p,d}(v^*_{r,a,p}TM)$.
From the above we conclude the error term to the approximate solution $\gamma_+\zeta_+ + \gamma_-\zeta_-$ is very small-exponentially suppressed in fact. We now perturb it by adding a small term $\delta\zeta \in H_0'$ to make it into a solution to Θ_v'. We state this in the form of a proposition:

Proposition 3.9.13. *We can choose $\delta\zeta \in H_0'$ so that $\zeta = \gamma_+\zeta_+ + \gamma_-\zeta_- + \delta\zeta$. Further, the norm of $\delta\zeta$, as measured in $W^{2,p}(v^*_{r,a,p}TM)$ is bounded above by*

$$C\epsilon^{2/p}e^{-\lambda(T_p/2\delta-3R)}.$$

The vector field $\delta\zeta$ also lives in $W^{3,p}(v^*_{r,a,p}TM)$, and its $W^{3,p}(v^*_{r,a,p}TM)$ norm is similarly bounded above by

$$C\epsilon^{2/p}e^{-\lambda(T_p/2\delta-3R)}.$$

Remark 3.9.14. We remark in the term $T_p/2\delta - 3R$, the term $3R$ appears because we can only start the exponential decay after the effects of $\psi'_\pm$ in Θ'_v disappear. (Technically we could have used $2R$ but this will not make a difference).

Proof. We plug $\zeta := \gamma_+\zeta_+ + \gamma_-\zeta_- + \delta\zeta$ into Θ'_v and solve for $\delta\zeta$ using the contraction mapping principle. We are now looking at an equation of the form:

$$D'_{J_\delta}(\gamma_+\zeta_+ + \gamma_-\zeta_- + \delta\zeta) + \beta'_\pm\psi'_\pm + e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta, e^{-w(s)}\psi'_\pm) = 0.$$

We examine the term $e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta, e^{-w(s)}\psi'_\pm)$, recall $\mathcal{F}_v$ generally takes the form:

$$\mathcal{F}_v := \beta_{[1;R-2,\infty]}\phi g_{v1}(\beta_{ug}\psi, \beta_v\phi) + \partial_t\phi g_{v2}(\beta_{ug}\psi, \beta_{[1;R-2,\infty]}\beta_v\phi).$$

Hence our expression can really be expanded as

$$\begin{aligned} e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta, e^{-w(s)}\psi'_\pm) =& e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_+\zeta_+ + \gamma_-\zeta_-, e^{-w(s)}\psi'_\pm) \\ &+ G_1(e^{-w(s)}\zeta_\pm, e^{-w(s)}\partial_t\zeta_\pm, e^{-w(s)}\psi'_\pm, e^{-w(s)}\delta\zeta)\delta\zeta \\ &+ G_2(e^{-w(s)}\zeta_\pm, e^{-w(s)}\psi'_\pm)\partial_t\delta\zeta. \end{aligned}$$

The functions G_* (the functions themselves, ignoring its inputs such as $\zeta_\pm$) have uniformly bounded smooth derivatives and are bounded in the following way:

$$G_*(x_1, \ldots, x_n) \leq |x_1| + \ldots + |x_n|$$

for x_* small. Recalling our choice of cut off functions we always have $w(s) > 1$, so this assumption is always satisfied. Recalling the elliptic regularity results on $\zeta_\pm$ above we can actually bound the $W^{2,p}(v^*_{r,a,p}TM)$ norm of G_1 and G_2 by ϵ. Then our equation for Θ'_v simplifies to

$$D'_{J_\delta}\delta\zeta + G_1\delta\zeta + G_2\partial_t\delta\zeta = E$$

where E was defined as the error term above. We now apply the contraction mapping principle to this equation, let $\Pi' \circ Q$ denote the right inverse to $D'_{J_\delta}|_{H'_0}$ (where Π' corresponds to projection to H'_0 as we have removed exponential weights). Consider the linear functional $I(\delta\zeta)$:

$$\delta\zeta \longrightarrow \Pi' \circ Q(-G_1\delta\zeta - G_2\partial_t\delta\zeta + E).$$

Let $B_\epsilon \subset H'_0$ denote a ball of size ϵ, then it follows from the form of G_* as well as the size estimate of E that I maps B_ϵ to itself. It follows similarly from above that I is a contraction mapping, hence it follows from the contraction mapping principle that such $\delta\zeta$ is unique. It follows from uniqueness of $\zeta \in H'_0$ in previous theorem that this ζ from this contraction mapping is the ζ we constructed earlier.

The norm estimate of $\delta\zeta$ follows directly from the norm estimate of E. The improvement from $W^{2,p}$ to $W^{3,p}$ is as follows: we first realize $\zeta = \gamma_+\zeta_+ + \gamma_-\zeta_- + \delta\zeta$ lives in $W^{3,p}$, the same is true for $\gamma_\pm\zeta_\pm$, hence $\delta\zeta$ also lives in $W^{3,p}$. To get the actual norm estimates, we recall the fixed point equation

$$\delta\zeta = \Pi' \circ Q(-G_1\delta\zeta - G_2\partial_t\delta\zeta - E).$$

We first realize $-G_1\delta\zeta - G_2\partial_t\delta\zeta + E$ actually lives in $W^{2,p}(v^*_{r,a,p}TM)$ by previous elliptic regularity results. We then realize $\Pi' \circ Q$ restricts to a bounded operator from $W^{2,p}(v^*_{r,a,p}TM) \to W^{3,p}(v^*_{r,a,p}TM)$ with image in $H'_0 \subset W^{3,p}(v^*_{r,a,p}TM)$ by applying elliptic regularity to D_{J_δ}. Finally we observe the $W^{2,p}$ norm of E is similarly bounded above by $Cmax\{\|\zeta_+\|, \|\zeta_-\|\}^{2/p}e^{-\lambda(T/2\delta-3R)}$ owing to the fact in the region where E is supported, $\zeta_\pm$ is smooth. Then to get the $W^{3,p}$ norm of $\delta\zeta$ we just measure the $W^{3,p}$ norm of both side of the fixed point equation and conclude. □

We now investigate how $\delta\zeta$ varies w.r.t. $p_\pm$, because we already understand $\zeta_\pm$ is well behaved w.r.t. $p_\pm$. Instead of varying $p_\pm$ individually, we find it is more convenient to change basis and distinguish two kinds of variations. We introduce the new variable p.

- We call the transformation of this type: $(p_-, p_+) \to (p_- - p, p_+ + p)$ a stretch.
- We can transformation of the type $(p_-, p_+) \to (p_- + p, p_+ + p)$ a translation.

We shall vary $\delta\zeta$ w.r.t. p with these kind of transformations. In both cases we shall show $\delta\zeta$ is well behaved w.r.t. differentiating via p.

Stretch

Observe in our region of interest we assumed $f'(x) = 1$, and that $x'(s) = \delta f'(x)$. The effect of stretch will be thought of as keeping the same gradient trajectory $v_{r,a,p}$ prescribed by (p_+, p_-) but lengthen the interval $s \in [0, T_p/\delta]$ to $[-p/\delta, T_p + p/\delta]$ over $v^*_{r,a,p}TM$ with the peak of exponential weight profile $w_p(s)$ still at $s = T_p/2\delta$. We translate u_+ and u_- in opposite directions along symplectization coordinate. We then think of equation Θ_v as taking place over the same gradient cylinder, but various terms like $\psi'_\pm$ being translated as we stretch along p. (There is some abuse in notation here, T_p refers to the gradient flow length for original pair (p_+, p_-), and p is how much we stretched).

The a distance between $a(-p/\delta)$ and $a(T_p + p/\delta)$ also changes but not in a linear fashion since $a'(s) = e^{\delta f(x)}$ but this is fine since none of our operators depend on a.

We make the following important observation about $\phi_\pm$. In our section dealing with semi-infinite trajectories when we moved the asymptotic vector p (p here as in an element among the tuple (r, a, p)) we preglued to a different gradient trajectory. To be specific, let's focus on ϕ_-. In the case of semi-infinite trajectories, after changing gradient trajectories, no matter the value of p_- the pregluing always happened at $s = R$. We denote the resulting function of (s,t) by $\widetilde{\phi}_-(p)$ so that in this system preluing always happened at $s = R$. Now in the stretch picture we are taking a different perspective, that when we deform by p we are pregluing to a different segment of the same gradient trajectory $v_{r,a,p}$, so $\tilde{\phi}_-(p)$ and ϕ_- are related via translation, to be precise

$$\phi_-(s + p/\delta) = \tilde{\phi}_-(p)(s)$$

Here we only consider variations in the $p_\pm$ directions and have suppressed the t variable - there should be some identification of $t + (r_+ - r_-)$ and t_-. Variations in $r_\pm$ will be considered in a subsequent section. The feedback into Θ_- is given precisely by $\tilde{\phi}(p)(s)$ for $s \in [-3R, 3R]$. And we understand how $\tilde{\phi}_-(p)$ depends on p, and by our previous sections its feedback into equation Θ_- is *well behaved* w.r.t. p. A similar relation also holds to ϕ_+, and $\psi_\pm$.

Here we see the advantage of working in $W^{2,p}(v^*_{r,a,p}TM)$ instead of $W^{2,p,w}(v^*_{r,a,p}TM)$ since our norms are independent of p. Observe similarly our definition of H'_0 is independent of p. The only dependence in p comes from terms of the form $e^{w_p(s)}$ (we include, where relevant, the subscript p into our exponential weight profiles), which we will be able to describe explicitly. The formulate the following proposition:

Proposition 3.9.15. *In the case of a stretch,*

$$\left\| \frac{d}{dp}\delta\zeta \right\| \leq \frac{C}{\delta} e^{-\lambda(T_p/2\delta - 3R)}$$

where the norm of $\|d/dp\delta\zeta\|$ is measured w.r.t. $W^{2,p}(v^*_{r,a,p}TM)$. We are taking the derivative at $p = 0$, but it is obvious a similar formula holds for all small values of p uniformly.

Proof. We already know for every p there is a $\delta\zeta$ (we suppress the dependence on p) satisfying

$$D'_{J_\delta}\delta\zeta + G_1\delta\zeta + G_2\partial_t\delta\zeta = E$$

which we may rewrite as

$$\delta\zeta = \Pi' \circ Q(-G_1\delta\zeta - G_2\partial_t\delta\zeta + E).$$

We next proceed to differentiate both sides w.r.t. p. We see that the result is an expression of the form

$$\begin{aligned}
&\frac{d}{dp}(\delta\zeta) \\
=&(\frac{d}{dp}\Pi' \circ Q)(-G_1\delta\zeta - G_2\partial_t\delta\zeta + E) \\
&+ \Pi' \circ Q \cdot (-\frac{dG_1}{dp}\delta\zeta - \frac{dG_2}{dp}\partial_t\delta\zeta) \\
&+ \Pi' \circ Q(-G_1 \cdot \frac{d}{dp}\delta\zeta - G_2\frac{d}{dp}\partial_t\delta\zeta) \\
&+ \Pi' \circ Q\frac{dE}{dp}.
\end{aligned}$$

See Remark 3.7.18 for this kind of differentiation.

Step 1. We first differentiate $\Pi \circ Q$ w.r.t. p. Recall over $W^{2,p,w}(v^*_{r,a,p}TM)$ Π takes the form:

$$\Pi(\phi) = \phi - \sum_* L_*(\phi)\partial_*$$

after we remove the exponential weights the corresponding operator Π' takes the form:

$$\Pi'\zeta = \zeta - \sum_* L_*(e^{-w(s)}\zeta)e^{w(s)}\partial_*.$$

For stretch ∂_* is independent of p, so the only dependence we see is on $w(s)$. We realize $L((e^{-w(s)}\zeta)) = L(\zeta)e^{-w(T_p/2\delta)}$, but we realize that $w(s) - w(T_p/2\delta)$ is independent of p, so we conclude Π' is independent of p.

We next consider $\frac{d}{dp}\Pi' \circ Q$. We observe this is a map from $W^{1,p}(v^*_{r,a,p}TM) \to H_0$. It is the inverse of $D_{J_\delta}|_{H_0'}$, so we can instead differentiate the relation

$$(\Pi' \circ Q) \circ (D'_{J_\delta} \circ \iota) = I.$$

where $\iota : H_0' \to W^{2,p}(v^*_{r,a,p}TM)$, to get

$$\frac{d(\Pi' \circ Q)}{dp}(D'_{J_\delta} \circ \iota) + (\Pi' \circ Q)\frac{(D'_{J_\delta} \circ \iota)}{dp} = 0$$

$$\frac{d(\Pi' \circ Q)}{dp} = -(\Pi' \circ Q)\frac{d(D'_{J_\delta} \circ \iota)}{dp}(\Pi' \circ Q).$$

We already know $\Pi' \circ Q$ is uniformly bounded w.r.t. $\delta \to 0$, now recall that

$$D'_{J_\delta} = D_{J_\delta} + w'(s).$$

Of course we know that $w(s)$ has a “kink” where the absolute value bends (see definition equation for w) but we can smooth it. Noting that D_{J_δ} is independent of p and $w'(s)$ is independent of p, we conclude that

$$\left\| \frac{d(\Pi' \circ Q)}{dp} \right\| \leq C$$

where we use the operator norm. (In this case it's in fact zero).

Step 2 We next examine the term $\frac{dG_*}{dp}$. There are two kinds of dependencies, one on how the function G_* depends on p, which we denote by $\partial G/\partial p$, and second how its arguments $\zeta'_\pm$and $\psi_\pm$ depend on p. We first recall G_* comes about from the expansion

$$\begin{aligned} e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta, e^{-w(s)}\psi'_\pm) =& e^{w(s)}\mathcal{F}_v(e^{-w(s)}(\gamma_+\zeta_+ + \gamma_-\zeta_-), e^{-w(s)}\psi'_\pm) \\ &+ G_1(e^{-w(s)}\zeta_\pm, e^{-w(s)}\partial_t\zeta_\pm, e^{-w(s)}\psi'_\pm, e^{-w(s)}\delta\zeta)\delta\zeta \\ &+ G_2(e^{-w(s)}\zeta_\pm, e^{-w(s)}\psi'_\pm)\partial_t\delta\zeta. \end{aligned}$$

The function $\mathcal{F}_v$ only depends on the geometry, so the only dependence of the function G on p is given by dw/ds:

$$\frac{\partial G_*}{\partial p} \leq C \left| \frac{dw}{dp} \right| \leq C/\delta.$$

Next we try to understand the dependence of dG_*/dp through its dependence on terms like $d\zeta_\pm/dp$ and $d\psi'/dp$. From previous remark by the semi-infinite trajectory case we understand $d\tilde{\zeta}_\pm/dp \leq C\epsilon$, and $\zeta_\pm = \tilde{\zeta}_\pm(s \mp p/\delta)$. Similarly we have $\psi'_\pm = \tilde{\psi}'(s \mp p/\delta)$. Noting $\tilde{\psi}_\pm$ doesn't depend on p, we have

$$\frac{d}{dp}\psi'_\pm = -\frac{1}{\delta}\frac{d}{ds}\tilde{\psi}'_\pm$$

thus estimates:

$$\left\| \frac{d\psi'_\pm}{dp} \right\| \leq C\epsilon/\delta.$$

Note taking the p derivative of ψ' has cost us a derivative, hence the above norm can only be measured in $W^{1,p}$. Thankfully this is enough for our purposes because Q brings back another derivative.

$$\left\| \frac{d\zeta_\pm}{dp} \right\| \leq \left\| \frac{d}{dp}\tilde{\zeta}_\pm \right\| + \frac{1}{\delta}\left\| \frac{d}{ds}\tilde{\zeta}_\pm \right\| \leq C\epsilon(1 + 1/\delta)$$

Now in this computation $\frac{d}{dp}\tilde{\zeta}_\pm$ lives naturally in $W^{2,p}$, by elliptic regularity $\tilde{\zeta}_\pm$ lives in $W^{3,p}$, so its s derivative lives in $W^{2,p}$. Hence the above inequality can at most hold in $W^{2,p}$, which suffices for our purposes.

Next we need to consider the $W^{1,p}$ norm of $\frac{d}{dp}\partial_t\zeta_\pm$, which we re-write as

$$\frac{d}{dp}\partial_t\tilde{\zeta}_\pm(s \mp p/\delta) = \partial_t\partial_p\tilde{\zeta}_\pm(s \mp p/\delta) - \frac{1}{\delta}\partial_s\partial_t\tilde{\zeta}_\pm(s \mp p/\delta).$$

We first make a remark about commutativity of derivatives, e.g. we have commuted $\partial_p\partial_t\tilde{\zeta}_\pm = \partial_t\partial_p\tilde{\zeta}_\pm$. We know $\partial_p\tilde{\zeta}_\pm$ is in $W^{2,p}$, hence we can commute the derivatives using the following version of Clairut's theorem:

Proposition 3.9.16. *If $f : \mathbb{R}^2 \to \mathbb{R}$ is so that $\partial_1 f, \partial_2 f, \partial_{2,1} f$ exists everywhere (here $\partial_1 f$ denotes the partial derivative of f w.r.t the first variable), $\partial_{2,1} f$ is continuous, then $\partial_{1,2} f$ exists and is equal to $\partial_{2,1} f$.*

Hence we can commute the derivative, measure the $W^{1,p}$ norm of $\partial_t\partial_p\tilde{\zeta}_\pm(s \mp p/\delta)$, and bound it by the $W^{2,p}$ norm of $\frac{d}{dp}\tilde{\zeta}_\pm$, which is bounded above by $C\epsilon$. The $W^{1,p}$ norm of $\partial_s\partial_t\tilde{\zeta}_\pm(s \mp p/\delta)$ is bounded by the $W^{3,p}$ norm of $\tilde{\zeta}_\pm$, which is also bounded by $C\epsilon$ by elliptic regularity.

But observe the expression involving $\frac{dG_*}{dp}$ is multiplied by $\delta\zeta$ or $\partial_t\delta\zeta$, so overall we have the estimate:

$$\left\| Q \circ \left\{ \frac{dG_1}{dp} \cdot \delta\zeta + \frac{dG_2}{dp}\partial_t\delta\zeta \right\} \right\|_{W^{2,p}} \leq C\epsilon/\delta e^{-\lambda(T_p/2\delta - 3R)} + \epsilon\|d\delta\zeta/dp\|_{W^{2,p}}. \tag{3.19}$$

The last term coming from the dependence of G_1 on $\delta\zeta$. We remark in the above the term $\frac{dG_1}{dp}\cdot\partial_t\delta\zeta$ we have a product of $W^{1,p}$ functions, which remains in $W^{1,p}$. This is where we justify our use of $W^{2,p}$ instead of $W^{1,p}$. See Remark 3.4.2.

Step 3 The next term is $\Pi'\circ Q(-G_1\cdot\frac{d}{dp}\delta\zeta - G_2\frac{d}{dp}\partial_t\delta\zeta)$. Note we have C^1 bound on G_*, which is bounded by $C\epsilon$, so after we apply $\Pi'\circ Q$ the norm of this term is overall bounded by $C\epsilon\|\frac{d\delta\zeta}{dp}\|_{W^{2,p}}$, and we move this term to the left hand of the equation.

Step 4 We finally estimate how the error term E depends on p, and here we shall use the exponential decay estimates proved in Section 3.8. Recall E takes the form

$$\begin{aligned} E =&\gamma'_{\pm}\zeta_{\pm} + [e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_+\zeta_+) - \gamma_+e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta_+)] \\ &+ [e^{w(s)}\mathcal{F}_v(e^{-w(s)}\gamma_-\zeta_-) - \gamma_-e^{w(s)}\mathcal{F}_v(e^{-w(s)}\zeta_-)]. \end{aligned}$$

The important feature of this expression is that it has support in $s\in[T_p/2\delta-2, T_p/2\delta+2]$, the term E and its derivative over p can be upper bounded by the terms:

$$|E(s,t)| \le C|\zeta_{\pm}(s,t)|(1+|\partial_t\zeta_{\pm}(s,t)|)$$

$$\left|\frac{dE}{dp}(s,t)\right| \le C\left|\frac{d\zeta_{\pm}}{dp}(s,t)\right| + C\,|\zeta_{\pm}|\left(\left|\frac{d}{dp}e^{-w(s)}\zeta_{\pm}\right| + \left|\frac{d}{dp}e^{-w}\partial_t\zeta_{\pm}\right|\right)$$

where for both equations size refers to C^1 norm. Since this is supported over $s\in[T_p/2\delta-2, T_p/2\delta+2]$, bounds on the uniform norm imply bounds on Sobolev norms. Furthermore we know by elliptic regularity $\zeta_{\pm}$ and its p derivative are smooth over this region so it make sense to talk about C^1 norms. We first note

$$\frac{d}{dp}e^{-w} = \frac{C}{\delta}e^{-w}.$$

So the only terms we need to worry about are

$$\frac{d}{dp}\zeta_{\pm}, \quad \frac{d}{dp}\partial_t\zeta_{\pm}$$

over the interval $s\in[T_p/2\delta-2, T_p/2\delta+2]$. We recall by our convention $\zeta_{\pm}(s) = \tilde{\zeta}_{\pm}(s\mp p/\delta)$ so we have

$$\frac{d\zeta_{\pm}}{dp} = \frac{d}{dp}\tilde{\zeta}_{\pm}(s\mp p/\delta) + \frac{1}{\delta}\frac{d}{ds}(\tilde{\zeta}_{\pm}(s\mp p/\delta)).$$

By the constraint that $s\in[T_p/2\delta-2, T_p/2\delta+2]$, the terms on the right hand side have already decayed substantially, hence they are bounded by

$$\frac{C}{\delta}e^{-\lambda(T_p/2\delta-3R)}$$

which quickly decays to zero as $\delta\to 0$. Finally we compute the derivative $\frac{d}{dp}\partial_t\zeta_{\pm}$ for $s\in[T_p/2\delta-1, T_p/2\delta+1]$. We can also break this down into

$$\frac{d\zeta_{\pm,t}}{dp} = \frac{d}{dp}\tilde{\zeta}_{\pm,t}(s\mp p/\delta) + \frac{1}{\delta}\frac{d}{ds}(\tilde{\zeta}_{\pm,t}(s\mp p/\delta)).$$

The exponential decay estimates in Corollary 3.8.7, as well as exponential decay in Proposition 3.8.4, say in the interval $s \in [T_p/2\delta - 1, T_p/2\delta + 1]$ the above is also bounded by

$$\frac{C}{\delta} e^{-\lambda(T_p/2\delta - 3R)}.$$

Step 5 Combining all of the above estimates we see that

$$\left\| \frac{d}{dp} \delta\zeta \right\| \leq \frac{C}{\delta} e^{-\lambda(T_p/2\delta - 3R)}$$

as claimed. □

Now we use the above to show ζ is well behaved in the sense we originally described.

Proposition 3.9.17. *ζ, and hence ϕ is well behaved with respect to p when p controls a stretch.*

Proof. Note what we feed into $\Theta_\pm$ are vector fields with exponential weights, so we put ζ back into $\Theta_\pm$ we need to turn it back to ϕ via

$$\phi = e^{-w(s)}\zeta$$

but note that from above we have

$$\phi = \gamma_+\phi_+ + \gamma_-\phi_- + e^{-w(s)}\delta\zeta.$$

And we know terms like $\gamma_\pm\phi_\pm$ behave nicely with respect to p. So it suffices to understand how $e^{-w(s)}\delta\zeta$ feeds back into $\Theta_\pm$. For simplicity we focus on Θ_-. For fixed p, and for $s \in [-p/\delta, -p/\delta + 3R]$, if we define $\delta\phi := e^{-w(s)}\delta\zeta$, then from the perspective of Θ_-, the vector field we see is $\delta\phi(s' - p/\delta)$ for $s' \in [0, 3R]$ equipped weighted norm $e^{ds'}$. We observe over the region $s' \in [0, 3R]$, the weight function coming from $e^{w(s')} = e^{ds'}$, so when we calculate how the p variation feeds back into Θ_- we are really looking at

$$\left\| \frac{d}{dp} \delta\zeta(s' - p/\delta) e^{-ds'} \right\|$$

for $s' \in [0, 3R]$ with respect to the norm $W^{2,p,d}(u_-^* TM)$, which is equivalent to the expression:

$$\left\| \frac{d}{dp} \delta\zeta(s' - p/\delta) \right\|$$

with the unweighted $W^{2,p}$ norm over the interval $s' \in [0, 3R]$. We observe

$$\left\| \frac{d}{dp} \delta\zeta(s' - p/\delta) \right\|_{W^{2,p}} \leq \frac{C}{\delta} \left\| \frac{d}{ds} \delta\zeta(s' - p/\delta) \right\|_{W^{2,p}} + \left\| \frac{d}{dp} \delta\zeta(s' - p/dt) \right\|_{W^{2,p}}$$

Here we have used elliptic regularity on $\delta\zeta$ to control its $W^{3,p}$ norm by its $W^{2,p}$ norm. The by the preceding proposition both of the above expressions are bounded above by $\frac{C}{\delta} e^{-\lambda(T_p/2\delta - 3R)}$, hence the proof. □

Translation

The case of translation is much easier than the case of stretch, as it bears many similarities with the case of semi-infinite trajectory. We don't even need to remove exponential weights. The only salient difference is we now have to work in a subspace H_0.

Let us first recall/set up some notation. Fix tuples $(r, a, p)_\pm$, and they determine a pregluing between u_+ and u_-. We use $v_{p_\pm}$ to denote the intermediate trajectory that connects between $u_+ + (r, a, p)_+$ and $u_- + (r, a, p)_-$ in the pregluing. As before we define $w_{p\pm}(s)$ as our exponential weight profile, and we have the codimension 3 subspace H_0. We fix (s, t) coordinates over $v_{p_\pm}$, with gluing happening at $s = R$ and $s = T_p/\delta - R$. Let $p \in \mathbb{R}$ be a small number denoting the size of the translation, let $p^*_\pm = p_\pm + p$, and let v_{p^*} denote the gradient trajectory between the pregluing determined by $p^*_\pm$. We equip vector fields over v_{p^*} with Sobolev norms as previous described and it also has a subspace H^*_0. On v_{p^*} we choose coordinates (s^*, t^*) and because we assumed the function $f(x)$ is locally linear (after maybe a change of coordinates) we have that pregluing happens at $s^* = R$ and $s^* = T_p/\delta - R$. Observe there is a parallel transport map using the flat metric

$$PT : W^{2,p,w}(v^*_{p^*}TM) \longrightarrow W^{2,p,w}(v^*_p TM)$$

such that if $\phi^*(s^*, t^*)$ is a vector based at $v_{p^*}(s^*, t^*)$, it is transported to $\phi(s = s^*, t = t^*)$ over $v_p(s, t)$. Note the parallel transport map send H^*_0 to H_0. And the solution $\phi^*_{p^*_\pm}$ to Θ_v over v_{p^*} can be identified with $\phi(p) \in H_0$ to an equation of the form

$$D_{J_\delta}(p)\phi_p + \mathcal{F}_v(p, \psi_\pm, \phi_p) + \beta'_\pm \psi = 0$$

and the feedback term from $\phi^*_{p^*_\pm}$ into $\Theta_\pm(p^*_\pm)$ can be identified with the feedback of ϕ_p which corresponds to regions $s \in [-3R, 3R]$ for Θ_- and $s \in [T_p/\delta - 3R, T_p/\delta + 3R]$ for Θ_+. Then it suffices to calculate the norm of $d\phi_p/dp$ in H_0.

Proposition 3.9.18. $\left\| \frac{d\phi_p}{dp} \right\|_{W^{2,p,w}(v^*_p TM)} \leq C\epsilon.$

Proof. Observe that $\|d/dpD_{J_\delta}\| \leq C$ when measured in the operator norm because the coefficient matrices in this operator only depend on the background geometry. The same is true for $\|\frac{\partial \mathcal{F}_v}{\partial p}(p, -, -)\|_{C^1} \leq C$. We recall $D_{J_\delta}(p)$ is an isomorphism from

$$H_0 \longrightarrow W^{1,p,w}(v^*_p TM)$$

hence has an inverse whose operator norm is uniformly bounded over p and as $\delta \to 0$. The same is true for the derivative in p of this inverse. To see this, we recall $D_{J_\delta}(p) : W^{2,p,w}(v^*_p TM) \to W^{1,p,w}(v^*_p TM)$ has a right inverse uniformly bounded in p and $\delta \to 0$, which we denote by Q. We also recall the inverse for $D_{J_\delta}(p)$ is obtained by $\Pi \circ Q$. Hence it suffices to show Π has uniformly bounded norm as p changes in a translation.
Recall

$$\Pi(\phi_p) = \phi_p - \sum_* L_*(\phi_p)\partial_*$$

as an operator we see that the terms involving $* = z, s$ are independent of p, the vector field $v := a\frac{v_*\partial_s - \partial_s}{\delta} + b\partial_s$ depends on p but we see in C^1 norm that $|\frac{dv}{dp}| \leq C$, so we see Π has uniformly bounded norm as p varies, which in turn implies $\Pi \circ Q$ has uniformly bounded norm. We now investigate $\frac{d}{dp}\Pi \circ Q$, which we can understand by differentiating the expression

$$\Pi \circ Q \circ D_{J_\delta}(p)|_{H_0} = id|_{H_0}$$

w.r.t. p, which yields

$$\frac{d}{dp}\Pi \circ Q = -\Pi \circ Q(\frac{d}{dp}D_{J_\delta}(p)) \circ \Pi \circ Q$$

which implies as an operator $\frac{d}{dp}\Pi \circ Q$ has uniformly bounded norm. Next we recast the equation:

$$D_{J_\delta}(p)\phi_p + \mathcal{F}_v(p, \psi_\pm, \phi_p) + \beta'_\pm \psi = 0$$

as a fixed point equation

$$\phi_p = \Pi \circ Q(-\mathcal{F}_v(p, \psi_\pm, \phi_p) - \beta'_\pm \psi)$$

using the exact same procedure as we did for for semi infinite gradient trajectories, we differentiate this equation in p to show $\|d\phi_p/dp\| \leq C\epsilon$. Observe after parallel transport there was no translation of $\psi_\pm$ involved. □

Since in this case we worked directly with weighted norms we can directly conclude:

Corollary 3.9.19. *With respect to translations, the vector field ϕ_p is well behaved.*

With this and the previous subsection, we conclude that ϕ is well behaved with respect to variations of $p_\pm$. In the next part we examine how ϕ varies when we change $r_\pm, a_\pm$.

Variations in $r_\pm$, $a_\pm$

In this subsection we show that when we vary the parameters $a_\pm$ and $r_\pm$ the solution ϕ is well behaved.

Proposition 3.9.20. *The solution ϕ to Θ_v is well behaved w.r.t. $a_\pm$.*

Proof. Observe that changing $a_\pm$ can also have the effect of lengthening and shortening the gradient trajectory we need to glue between u_+ and u_-, though the process is substantially less dramatic than when we changed $p_\pm$. For instance when we change $a_\pm$ by size ϵ, the connecting gradient trajectory may lengthen/shrink by size $C\epsilon$, instead of $C\epsilon/\delta$. In particular we can redo all of the previous subsection. We separate the change to stretch and translation. We first observe in case of translation the equation Θ_v actually stays invariant, because all of our background geometry is invariant in the a direction. In the case of stretch, we remove the exponential weights, and repeat the above proof. The difference is that no factor of C/δ ever appears, so we don't even need the exponential decay estimates. The rest follows as above. □

Proposition 3.9.21. *$\phi(r, a, p)$ is well behaved as we vary $r_\pm$.*

Proof. Recall for r_+ in the pregluing construction we are rotating the entire gradient trajectory $v_{r,a,p}$ along with it, so we can again use parallel transport in r_+ to turn it into a family of equations over the same space, which we denote by W as before, and the resulting (r_+, r_-) family of PDEs over W by $\hat{\Theta}_v$. We use $\hat{\phi}(r_+, r_-)$ to denote the solution to $\hat{\Theta}_v$. By assumption, the almost complex structure J, when restricted to the surface of the Morse-Bott torus, is r invariant, however, the local geometry is not necessarily invariant. Therefore, the linearized operator as well as nonlinear term picks up a r_+ dependence, so the equation solved by $\phi(r_+, r_-)$ has a linear operator $D_{J_\delta}(r_+)$ and a nonlinear term $\hat{\mathcal{F}}_v(r_+, -)$ with r_+ dependence. Also observe H_0 is invariant under changing $r_\pm$, so we denote it by the same letter when viewed as subspace in W.

We now recall what happens to the pregluing near the u_- end, the domain Riemann surface $\Sigma_{r,a,p}$ is constructed at $s = R$ with the identification $t + r_+ \sim t_- + r_-$. So we see this effect in the equation $\hat{\Theta}_v$ via the dependence of ψ_- on $r_\pm$, in particular the ψ_- term in $\hat{\Theta}_v$ should be instead $\psi_-(s, t + r_+ - r_-)$. Hence after parallel transport we see $\hat{\phi}(r_+, r_-)$ is the unique solution to the equation in H_0:

$$D_{J_\delta}(r_+)\hat{\phi} + \beta'_+\psi_+(s,t) + \beta'_-\psi(s, t + r_+ - r_-) + \hat{\mathcal{F}}_v(r_+, \psi_\pm) = 0.$$

Again, following the same procedure as we did for semi-infinite gradient trajectories we recast this as a fixed point equation

$$\hat{\phi} = \Pi \circ Q(-\beta'_+\psi_+(s,t) - \beta'_-\psi(s, t + r_+ - r_-) - \hat{\mathcal{F}}_v(r_+, \psi_\pm))$$

and differentiate both sides with respect to $r_\pm$, observing that $\frac{\partial \hat{\mathcal{F}}_v(r_+,\psi_\pm)}{\partial r_+} = g(\phi, \psi) + h(\phi, \psi)\partial_t(\phi)$ as in Remark 3.7.12. However, it is important to note that taking an $r_\pm$ derivative of the above equation will produce a t derivative of ψ_-, which will produce a function in $W^{1,p}$ (we neglect any mention of weights for now). But since we are not taking any further derivatives of $\psi_\pm$, this is fine as Q will send this to $W^{2,p}$, then the same argument as before shows that

$$\left\| \frac{d\hat{\phi}}{dr_\pm} \right\|_{W^{2,p,w}(v^*_{r,a,p}TM)} \leq C\epsilon$$

as desired. □

Solution of $\Theta_\pm$

In this subsection we use the results from previous section to finally solve $\Theta_\pm$ and hence conclude gluing exists. Recall deformations of $u_\pm$ are given by the elements $(\psi_\pm, (r, a, p)_\pm, \partial'_\pm, \delta j_\pm) \in W^{2,p,d}(u^*_\pm TM) \oplus V_\pm \oplus V'_\pm \oplus T\mathcal{J}_\pm$, and the linearized Cauchy Riemann operator

$$D\bar{\partial}_{J\pm} : W^{2,p,d}(u^*_\pm TM) \oplus V_\pm \oplus V'_\pm \oplus T\mathcal{J}_\pm \longrightarrow W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u^*_\pm TM))$$

is surjective with right inverse $Q_\pm$. Then $\Theta_\pm$ are equations of the form

$$D\bar{\partial}_{J_\pm}((\psi_\pm,(r,a,p)_\pm,\partial'_\pm,\delta j_\pm)) + \mathcal{F}_\pm((\psi_\pm,(r,a,p)_\pm,\partial'_\pm,\delta j_\pm,\phi) + \mathcal{E}_\pm + \beta'_v\phi = 0$$

where $\mathcal{F}_\pm$ is a quadratic expression in each of its variables, implicit in $\mathcal{F}_\pm$ are quadratic terms depending on $(\delta j_\pm,\psi_\pm)$ responsible for variation of domain complex structure away from the punctures. And implicit in term $\mathcal{E}_\pm$ are error terms uniformly bounded by $C\delta$ in the interior of $u_\pm$ responsible for the fact that $u_\pm$ are J-holomorphic, instead of J_δ-holomorphic.

Theorem 3.9.22. *The system of equations $\Theta_\pm = 0$ has a solution, and hence 2 level cascades with one intermediate end can be glued. Furthermore, for specific choices of $Q_\pm$, which are right inverse to $D\bar{\partial}_{J_\pm}$, there is a unique solution in the image of (Q_+,Q_-).*

Proof. We consider the system of $\Theta_\pm$ as a map from

$$\begin{aligned}(\Theta_+,\Theta_-):&(W^{2,p,d}(u_+^*TM)\oplus V_+\oplus V'_+\oplus T\mathcal{J}_+)\oplus(W^{2,p,d}(u_-^*TM)\oplus V_-\oplus V'_-\oplus T\mathcal{J}_-)\longrightarrow\\ &W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma},u_+^*TM))\oplus W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma},u_-^*TM)).\end{aligned}$$

We solve this via a fixed point theorem by finding a fixed point to the map

$$\begin{aligned}&[(\psi_+,(r,a,p)_+,\partial'_+,\delta j_+),(\psi_-,(r,a,p)_-,\partial'_-,\delta j_-)]\\ &\longrightarrow[Q_+(-\mathcal{F}_+-\mathcal{E}_+-\beta'_v\phi),Q_-(-\mathcal{F}_--\mathcal{E}_--\beta'_v\phi)].\end{aligned}$$

We show it maps the ϵ ball to itself. This follows from the size estimates we had of ϕ relative to $\psi_\pm$, as well as the fact $\mathcal{F}_\pm$ is quadratic, and the size of the terms that appear in $\mathcal{E}_\pm$ are very small. We next argue this map has the contraction property as we vary $\psi_\pm,(r,a,p)_\pm,\partial'_\pm,\delta j_\pm$; this follows directly from the previous subsection in which we showed ϕ is well behaved with respect to these input variables, plus the fact $\mathcal{F}_\pm$ is quadratic (see remark 3.7.12).The sizes of terms that appear in $\mathcal{E}_\pm$ are also uniformly small, as we derived in the pregluing section. Hence the contraction mapping principle shows there is a unique solution in the image of (Q_+,Q_-). □

Remark 3.9.23. Relation to obstruction bundle gluing. We remark we could have proved a gluing exists via obstruction bundle gluing methods in [40],[41]. This is more similar to how the gluing of 1 level cascades was constructed in [12]. We explain this in the simplified setting as above, and the general case of multiple level cascade can be done analogously. Recall $D\bar{\partial}_{J,\pm}$ is index 1 (i.e. $u_\pm$ is rigid), we let $U_\pm$ be a 1 dimensional vector space in $W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma},u_-^*TM))$ spanned by image of asymptotically constant vector field ∂_x under $D\bar{\partial}_{J,\pm}$, and let $U'_\pm$ denote a fixed complement given by the image of $(\psi_\pm,r_\pm,a_\pm,\delta j_\pm)$ under $D\partial_{J,\pm}$. The fact $U'_\pm$ is closed follows from the fact our operators are Fredholm, and it form a complement for index reasons (as long as neither $u_\pm$ is a trivial cylinder).

Then we form the (trivial) obstruction bundle with base $(p_+,p_-)\in[-\epsilon,\epsilon]^2$ and fiber $U_+\oplus U_-$. Then instead of solving $\Theta_\pm$ on the nose we introduce projections $\Pi_{U'_\pm}$ that project to $U'_\pm$. Then for fixed input data $\{(p_+,p_-),(\psi_\pm,r_\pm,a_\pm,\delta j_\pm)\}$ in an epsilon ball, we solve the equation Θ_v for ϕ

$$D_{J_\delta}\phi+\beta'_+\psi_+(s,t)+\beta'_-\psi(s,t+r_+-r_-)+\mathcal{F}_v(r_+,\psi_\pm)=0.$$

Its (unique) solution ϕ, which depends on all input data $\{(p_+,p_-),(\psi_\pm,r_\pm,a_\pm,\delta j_\pm)\}$, will have norm uniformly bounded by $C\epsilon/R$ (the C is uniform as we vary (p_+,p_-)). Then the solution to the system of equations $\Theta_\pm=0$ is equivalent to the solution of the following system of equations

$$\theta_\pm := D\bar\partial_{J,\pm}(\psi_\pm,(r,a)_\pm,\partial'_\pm,\delta j_\pm)+\Pi_{U'_\pm}[(+\mathcal{F}_\pm+\mathcal{E}_\pm+\beta'_v\phi]$$

$$D\bar\partial_{J,+}p_+ + (1-\Pi_{U'_+})[(+\mathcal{F}_+ +\mathcal{E}_+ +\beta'_v\phi)]=0$$

$$D\bar\partial_{J,-}p_- + (1-\Pi_{U'_-})[(+\mathcal{F}_- +\mathcal{E}_- +\beta'_v\phi)]=0.$$

We observe for fixed (p_+,p_-) the equations $\theta_\pm$ can always be solved via contraction mapping principle, essentially because the nonlinear term under the projection $\Pi_{U'_\pm}$ always lands in the image of $D\bar\partial_{J,\pm}$ by construction, and we have estimates $\|\phi\|\le C\epsilon/R$. The other two equations in the language of [41] define an obstruction section to the obstruction bundle, as

$$\mathfrak{s}:=\{p_++(1-\Pi_{U_+})[(\mathcal{F}_++\mathcal{E}_++\beta'_v\phi],p_-+(1-\Pi_{U_-})[(\mathcal{F}_-+\mathcal{E}_-+\beta'_v\phi]\}\in\Gamma(U_+\oplus U_-\longrightarrow[-\epsilon,\epsilon]^2)$$

and the vanishing of $\mathfrak{s}$ corresponds to gluing. In the above expression we think of $p_\pm$ as real numbers (because we have projected to the one dimensional spaces $U_\pm$). But we observe by the size estimates of $\phi,\psi_\pm$, the size of the nonlinear term $(1-\Pi_{U_\pm})[(\mathcal{F}_\pm+\mathcal{E}_+\pm\beta'_v\phi)]$ under $\Pi_{U_\pm}$ is uniformly bounded above by $C\epsilon^2/R$. However the linear term $p_\pm$ varies freely from $-\epsilon$ to ϵ. The nonlinear term is clearly continuous with respect to variations in $p_\pm$. Hence from topological considerations the obstruction section must have at least one zero, hence we have at least one gluing.

The difficulty with the above approach, is of course there is at least one gluing, but it is unclear how many there are in total. One could improve the above conclusion by trying to argue that $\mathfrak{s}$ is not only C^0 close to the (p_+,p_-) but also C^1 close, and this would imply the zero is unique. In fact what we proved about "ϕ being well behaved w.r.t. $p_\pm$" is tantamount to showing $\mathfrak{s}$ is C^1 close to (p_+,p_-). This required we do very careful exponential decay estimates as well as another contraction mapping principle. That we previously proved gluing via contraction mapping and here phrased it here as obstruction bundle gluing is purely a matter of repackaging.

Remark 3.9.24. Another possible approach to obstruction bundle gluing might be to show for generic choice of J_δ we can arrange to have the zeros of the obstruction section be transverse to the zero section. This will show there is only one gluing up to sign. This is more in line with the strategy taken in [41]. However, it's unclear whether we can choose generic enough J_δ since here we have a family of J_δ degenerating as $\delta\to 0$ as opposed to some fixed generic J.

Remark 3.9.25. We shall later prove surjectivity of gluing. The appendix of [12] used a different strategy for surjectivity, hence did not need to prove the solution obtained via obstruction bundle gluing is unique. Conceivably the methods there could also be applied here, but the construction would be difficult for two reasons: one they used stable Hamiltonian structures as opposed to contact structures, therefore their equation is nicer than ours. Two it seems their methods would be difficult to carry out in multiple level cascades where the dimensions of moduli spaces that appear could be very high. Instead in what follows we use an approach in Section 7 of [41].

Gluing multiple level cascades

In this subsection we generalize gluing to multiple cascade levels. Given what we have proved above, this is mostly a matter of linear algebra. However there are still subtle details we need to take care of, we first take care of the simple case where we are still gluing together a 2-level cascades, except now with multiple ends meeting in the middle. This contains all the important features required for the gluing. Then we will simply generalize this situation to n level cascades.

2-level cascade meeting at multiple ends

We consider a 2-level cascade built out of two J-holomorphic curves u_+ and u_- meeting along n free ends along an intermediate Morse-Bott torus. It does not matter how many intermediate Morse-Bott tori are there, so for simplicity we assume there is only one. We assume all ends of u_- and u_+ landing on this Morse-Bott torus avoid critical points of f, and we have chosen coordinates so that the Morse function looks like $f(x) = x$. We assume this cascade is rigid, and $ev_-(u_+)$ and $ev_+(u_-)$ are separated by gradient flow of f for time T. We also assume the x coordinates of the positive asymptotic Reeb orbits of u_- are labelled by $x_1, ..., x_n$.

In this example, for simplicity of exposition, we only focus on gluing finite gradient cylinders, and ignore gluing for semi-infinite trajectories. Hence we assume no positive end of u_+ nor negative end of u_- lands on the Morse-Bott torus that appear in the intermediate cascade level, and we only perturb the contact form to be nondegenerate in a neighborhood of this torus.

The fact the cascade is rigid and transverse implies the following operator is surjective

$$\begin{aligned} D_+ \oplus D_- : & W^{2,p,d}(u_+^* TM) \oplus T\mathcal{J}_+ \oplus V'_+ \oplus V''_+ \\ & \oplus W^{2,p,d}(u_-^* TM) \oplus T\mathcal{J}_- \oplus V'_- \oplus V''_- \oplus (\Delta_t) \longrightarrow \\ & W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u_+^* TM)) \oplus W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u_-^* TM)). \end{aligned}$$

$V'_\pm$ denotes asymptotic vectors associated to ends away from glued ends. $V''_\pm$ denotes asymptotic vector fields at glued ends except they only include $(r,a)_\pm$ components. Δ_t is a $n+1$ dimensional vector space that consists of asymptotic vectors that satisfy relations $p_i^+ - p_i^- = t$, where t is a positive real number that varies freely. $D_\pm$ is our shorthand for the linearization of the Cauchy Riemann operator, which implicitly also depends on the complex structure of the domain.

We next do a much more careful pregluing. The main difficulty is for fixed $\delta > 0$, suppose we start at x_i, and connect to a lift of a gradient trajectory that flows for distance T in the x direction, if we used (s,t) coordinates on this gradient flow cylinder, the s coordinate has range $s \in [0, T/\delta]$ which is independent of i. But the a distance (i.e. distance in the symplectization direction) traveled by this gradient trajectory for the same s from 0 to T/δ varies depending on i, fundamentally this is because the a coordinate satisfies the ODE

$$a'(s) = e^{\delta f(x(s))}$$

which depends on the value of f. Hence in the pregluing, instead of using the vector field Δ_t where $p_i^+ - p_i^- = t$, there would be some nonlinear relations between the asymptotic vectors ∂_s and ∂_x.

Fix cylindrical coordinates (s_i^+, t_i^+) around each of the punctures of u_+ that hits the intermediate cascade level (i.e. the Morse Bott torus) and likewise (s_i^-, t_i^-) for punctures of u_-. Near each of the punctures the maps $u_\pm$ takes the form

$$(a_{i\pm}(s_i^\pm, t_i^\pm), z_{i\pm}(s_i^\pm, t_i^\pm), x_{i\pm}(s_i^\pm, t_i^\pm), y_{i\pm}(s_i^\pm, t_i^\pm)).$$

We use π_* for $* = a, z, x, y$ to denote the relevant component of a map, i.e. $(\pi_a u_+)_i$ denotes the a component of u_+ at its ith end. We use the following notation to denote the various evaluation maps

$$ev_i^-(a)(R) := \int_{S^1} \pi_a(u^+)_i(-R,t)dt, \quad ev_i^+(a)(R) := \int_{S^1} \pi_a(u^-)_i(R,t)dt$$

$$ev_i^-(x)(R) := \int_{S^1} \pi_x(u^+)_i(-R,t), \quad ev_i^+(x)(R) := \int_{S^1} \pi_x(u^+)_i(-R,t)(R,t)dt.$$

We observe the deformation with respect to the asymptotically constant vector field ∂_r is constructed the same way as before, so we focus our attention on the vector spaces $V_+(x) \oplus V_-(x) \oplus V_+(a) \oplus V_-(a)$ consisting of the tuples $(p_i^+, p_i^-, a_i^+, a_i^-)$.

Let $T' > 0$. Consider the submanifold $\hat{\Delta}$ in $V_+(x) \oplus V_-(x) \oplus V_+(a) \oplus V_-(a)$ defined as follows

$$p_1^+ - p_1^- = T'$$

$$p_i^+ - p_i^- = T' + f_i(a_i^\pm, a_1^\pm, p_i^-, R)$$

$$|a_i^\pm|, |p_i^\pm| < \epsilon$$

where f_i is defined as follows: let v_{1p} denote the gradient trajectory connecting the $i = 1$ ends between u_+ and u_-. We endow it with the following specification: its a coordinate at

$s = R$ starts at $ev_1^+(a)(R) + a_i^+$, and its x coordinate at $s = R$ starts at $ev_1^+(x)(R) + p_1^-$. It follows the gradient flow for s length T'. We then translate u_+ in the a direction so that $ev_1^-(a)(R) + a_{1-} = \pi_a(v_{1p}(T'/\delta - R, t))$. Further we have $ev_1^-(x)(R) + p_1^+ = \pi_x(v_{1p}(T'/\delta - R, t))$.

Then for $a_i^\pm, i \geq 2$, we define f_i to be the amount of displacement in the x direction required so that a gradient flow of s-length $(T' + f_i(a_i^\pm, a_1^\pm, p_i^\pm, R))/\delta$ flows from $ev_i^-(a)(R) + a_i^-$ to $ev_i^+(a)(R) + a_i^+$ at the i th end between $u_+ + p_i^+$ and $u_- + p_i^-$. By s-length we mean for a finite segment of gradient cylinder, after having chosen coordinates (s, t) on the gradient cylinder, the amount by which s needs to change to go from one end of the gradient cylinder to the other end. We see immediately that

$$f_i \leq C\delta$$

and

$$\frac{\partial f_i}{\partial a_j^\pm} \leq C\delta \text{ where } j = 1, i$$

$$\frac{\partial f_i}{\partial p_i^-} \leq C\delta.$$

From this it follows immediately that $\hat{\Delta}$ is a submanifold. And that for small enough δ the operator

$$\begin{aligned} D_+ \oplus D_- : & W^{2,p,d}(u_+^* TM) \oplus T\mathcal{J}_+ \oplus V_+' \oplus V_+(r)'' \\ & \oplus W^{2,p,d}(u_-^* TM) \oplus T\mathcal{J}_- \oplus V_-' \oplus V_-''(r) \oplus (\hat{\Delta}) \longrightarrow \\ & W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u_+^* TM)) \oplus W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u_-^* TM)) \end{aligned}$$

is surjective with uniformly bounded right inverse. By $V_\pm''(r)$ we mean the subspace of $V_\pm''$ that only includes the r components of the asymptotic vectors. Then it follows immediately that any element in $\hat{\Delta}$ gives rise to a pregluing, since the a and x components of $u_\pm$ and the intermediate gradient trajectories match.

Remark 3.9.26. Our operator

$$\begin{aligned} D_+ \oplus D_- : & W^{2,p,d}(u_+^* TM) \oplus T\mathcal{J}_+ \oplus V_+' \oplus V_+(r)'' \oplus W^{1,p,d}(u_-^* TM) \\ & \oplus T\mathcal{J}_- \oplus V_-' \oplus V_-''(r) \oplus (\hat{\Delta}) \longrightarrow \\ & W^{1,p,d}(\overline{Hom}(T\dot{\Sigma}, u_+^* TM)) \oplus W^{1,p,d}(\overline{Hom}(T\dot{\Sigma}, u_-^* TM)) \end{aligned}$$

has a two dimensional “kernel”. The kernel is in quotations because $\hat{\Delta}$ is a submanifold instead of a vector subspace, but as we have seen it is exceedingly close to a linear subspace, so we gloss over this point. The two dimensional “kernel” consists of two kinds of elements, they both come from the fact $u_\pm$ are J-holomorphic curves in symplectizations and hence there is a translation symmetry. The first kind of kernel element comes from translating

u_+ and u_- by the same amount in the sympletization direction. This is an genuine kernel element of $D_1 \oplus D_2$. The other kernel element is translation u_+ and u_- in opposite directions, so that they become closer/farther away from each other. This is no longer in the kernel because of the nonlinearities of $\hat{\Delta}$, but as we see the corrections are small. For the purposes of this section choosing a right inverse for $D_1 \oplus D_2$ doesn't matter, since we only need to show a gluing exists. Later when we need to prove surjectivity of gluing we will choose specific right inverses for $D_1 \oplus D_2$, which amounts to saying we consider vector fields where there are approximately no $\mathbb{R}$ translations over the curves u_+ and u_-.

We can now state the gluing construction.

Theorem 3.9.27. *2-level cascades of the form above can be glued. The gluing is unique up to choosing a right inverse for $D_+ \oplus D_1$ when we restrict the allowed asymptotic vectors corresponding to ends that meet on the intermediate cascade level on the Morse-Bott torus to $\hat{\Delta} \oplus V_+(r)'' \oplus V_-''(r)$ as above.*

Proof. Given a tuple of elements $(a_i^\pm, p_i^\pm) \in \hat{\Delta}$, as well as $r_i^\pm \in V_+(r)'' \oplus V_-(r)''$ as twist parameters, we can define a preglued curve u_* by pregluing n gradient trajectories v_i between u_- and a translated u_+. Then, just as how we proved gluing for two curves with a single end meeting at intermediate cascade level, we deform the pregluing with appropriate vector fields, i.e. starting with vector fields $\psi_\pm \in W^{2,p,d}(u_\pm^* TM)$ and $\phi_i \in W^{2,p,w_i}(v_i^* TM)$. We also implicitly deform the domain complex structures of $u_\pm$ using $\delta j_\pm$; we also deform using asymptotic vectors at other ends in $u_\pm$, they live in $V'_\pm$ and we denote them by $\partial'_\pm$; since they are not super relevant to our construction we suppress them from our notation. We construct the perturbation

$$\beta_+\psi_+ + \beta_-\psi_- + \sum \beta_{v_i}\phi_i.$$

And as before, the deformation is holormophic iff the system of equations can be solved:

$$\begin{aligned}
&\Theta_+(\psi_+, (r,a,p)_{\pm i}, \partial'_+, \delta j_+) = 0\\
&\Theta_-(\psi_-, (r,a,p)_{\pm i}, \partial'_-, \delta j_-) = 0\\
&\Theta_i(\psi_\pm, \phi_i) = 0.
\end{aligned}$$

Then we follow the same strategy of proof as before, given the tuples $(\psi_\pm, (r,a,p)_{i\pm}, \partial'_\pm, \delta j_\pm)$ of input along $u_\pm$ we can define subspaces $H_{0i} \subset W^{2,p,w_i}(v_i^* TM)$ such that there exists unique solution to $\mathbf{\Theta}_i = 0$, $\phi_i \in H_{0i}$. It follows immediately from previous theorems that ϕ_i has norm bounded above by $C\epsilon/R$ and is nicely behaved with respect to variations of all input data $(\psi_\pm, (r,a,p)_{i\pm}, \partial'_\pm, \delta j_\pm)$. We then view the system $\Theta_\pm = 0$ as looking for a zero of a map

$$\begin{aligned}
&W^{2,p,d}(u_+^* TM) \oplus T\mathcal{J}_+ \oplus V'_+ \oplus V_+(r)'' \oplus W^{2,p,d}(u_-^* TM) \oplus T\mathcal{J}_- \oplus V'_- \oplus V''_-(r) \oplus (\hat{\Delta}) \longrightarrow\\
&W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u_+^* TM)) \oplus W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u_-^* TM)).
\end{aligned}$$

It follows from our previous calculations of how $\Theta_\pm$ looks like in these coordinates, as well as the fact the operator $D_+ \oplus D_-$ restricted to $W^{2,p,d}(u_+^* TM) \oplus T\mathcal{J}_+ \oplus V'_+ \oplus V_+(a)'' \oplus W^{2,p,d}(u_-^* TM) \oplus T\mathcal{J}_- \oplus V'_- \oplus V''_-(a) \oplus (\hat{\Delta})$ being surjective that $\Theta_\pm$ can be solved simultaneously for $(\psi_\pm, (r,a,p)_{i\pm}, \partial'_\pm, \delta j_\pm)$ via the contraction mapping principle. Such a solution is unique provided we fix a right inverse for $D_+ \oplus D_-$. □

We now turn to gluing n-level cascades. It will follow the same strategy as above as long as we introduce some new notations so we will be brief. The main purpose of the ensuing proof is to introduce some useful notations.

Theorem 3.9.28. *n-level tranverse and rigid cascades can be glued, and the solutions are unique up to choosing a right inverse, as specified in the proof.*

Proof. Let $u^{\natural} = \{u^i\}_{i=1,..,n}$ be an n level cascade that is transverse and rigid. For each u^i we let W_i denote the vector space $W^{2,p,d}(u_i^* TM) \oplus T\mathcal{J}_+ \oplus V'_+ \oplus V_+(r)''$ and L_i the vector space $W^{1,p,d}(\overline{\mathrm{Hom}}(T\dot{\Sigma}, u^{i*}TM))$ and let $\hat{\Delta}_{i,i+1}$ denote the submanifold consisting of asymptotic vectors in a, x directions corresponding to free ends that meet each other between u^i and u^{i+1}, analogous to $\hat{\Delta}$ for the 2 level case, so that pregluing makes sense. Then the fact that the cascade exists, is transversely cut out, and of Fredholm index 0 implies the operator

$$\oplus D_i : W_1 \oplus \hat{\Delta}_{1,2} \oplus \ldots \oplus W_n \longrightarrow L_1 \oplus .. \oplus L_n$$

is surjective with uniformly bounded right inverse. Hence for each element in $\hat{\Delta}_{i,i+1}$ we preglue together u^i and u^{i+1} by inserting a collection of gradient trajectories in the middle. In case u^i and u^{i+1} have components consisting of trivial cylinders that begin and end on critical points, we recall such chains of trivial cylinders will eventually meet a non-trivial J holomorphic curve with fixed end at the critical point. We replace such chains of trivial cylinders with a single fixed trivial cylinder as in the case of gluing fixed trivial cylinders in the case of semi-infinite trajectories. We add marked points to unstable components in cascade levels to make them stable, see Convention 3.4.3. For the positive ends of u^1 and negative ends of u^n, if it a free end we glue in a semi-infinite gradient trajectory, and if it is a fixed end we glue in a trivial cylinder. This constructs for us a preglued curve u_*. Then we deform this preglued curve using vector fields ψ_i over u^i and ϕ_i over the gradient flow lines we preglued. We require that ϕ_i lives in the vector space H_{0i}, which is defined analogously to H_0 in the case of 2 level cascades, if ϕ_i corresponds to a finite gradient flow trajectory, and no such requirement is imposed if ϕ_i is over a semi-infinite gradient trajectory. As before the entire preglued curve can be deformed to be holomorphic iff the system of equations

$$\mathbf{\Theta_i}(\psi_i, (r,a,p)_{\pm i}, \phi_j, \partial'_i, \delta j_i) = 0, \quad \mathbf{\Theta}_{v_i}(\phi_i, \psi_j) = 0$$

can be solved. We use bold to denote a system of equations. $\mathbf{\Theta_i}(\psi_i, (r,a,p)_{\pm i}, \phi_j, \partial'_i, \delta j_i)$ corresponds to equations over u^i, and $\mathbf{\Theta}_{v_i}$ corresponds to equations over gradient flow trajectories, which implicitly includes semi-infinite trajectories. As before for fixed epsilon ball in $\oplus W_i \oplus \hat{\Delta}_{i,i+1}$, the equations Θ_{v_i} have unique solutions in H_{0i} that are well behaved w.r.t.

input. Then the equations $\Theta_i(\psi_i, (r,a,p)_{\pm i}, \phi_j, \partial_i', \delta j_i) = 0$ have unique solutions follow from the fact $\oplus D_i$ is surjective with uniformly bounded inverse and the contraction mapping principle. The solution is unique to a choice of right inverse for the operator $\oplus D_i$. □

Remark 3.9.29. We note here by elliptic regularity all of our solutions are smooth, with their higher $W^{k,p}$ norms bounded by their $W^{1,p}$ norm.

Remark 3.9.30. We note by the additivity of the relative first Chern class and the Euler characteristic, the resulting glued curve has Fredholm index one.

3.10 Behaviour of holomorphic curve near Morse-Bott tori

In this section we prove a series of results concerning how J-holomorphic curves behave near Morse-Bott tori. This is part of the analysis that is needed to prove the degeneration result from J-holmoprhic curves to cascades in Bourgeois' thesis [5]. We redo this part of the analysis, not only to prove the degeneration result in our case in the Appendix, but we will also need them to later show that the gluing we construct is surjective. The analysis here is very similar to the analysis performed in the Appendix of Bourgeois and Oancea's paper [7], the only major difference is we are working in symplectizations where they work near a Hamiltonian orbit. We start with a series of analytical lemmas.

Semi-infinte ends

Recall the neighborhood of Morse-Bott torus we have coordinates $(z, x, y) \in S^1 \times S^1 \times \mathbb{R}$, with J chosen so that at the surface of the Morse-Bott torus $J\partial_x = \partial_y$. The linearized Cauchy Riemann operator along trivial cylinders that land on this Morse-Bott torus takes the form $\partial_s + A$, where

$$-A := J_0(d/dt) + S_0(x, z)$$

$S(x, z)$ is a symmetric matrix that only depends on x and z. The kernel of $A(s)$ is spanned by $\partial_a, \partial_z, \partial_x$. Let P denote the L^2 projection to its kernel, and let Q denote the projection $\ker A^{\perp}$.

Theorem 3.10.1. *Let $u_\delta(s,t) = (a(s,t), z(s,t), x(s,t), y(s,t))$ be a J_δ-holomorphic map that converges to a simply covered Reeb orbit corresponding to a critical point of f as $s \to \infty$. We also assume for $s > 0$, the map u_δ stays away from all other Reeb orbits corresponding to other critical points of f, uniformly as $\delta \to 0$. Assume for $s > 0$ we have*

$$|y|, |z - t|, |\partial_*^{\leq k} x|, |\partial_*^{\leq k} y| \leq \epsilon$$

where $* = s, t$, and $\epsilon > 0$ is sufficiently small (but independent of δ). We also assume all other derivatives are uniformly bounded above by C. There is some $r > 0$ independent of δ

and only depending on the local geometry of Morse-Bott tori so that

$$|y|, |z - t + c| \le C\|Q(Y)(0,t)\|_{L^2(S^1)}^{2/p} e^{-rs}$$

$$|x - x_p(s)| \le C\|Q(Y)(0,t)\|_{L^2(S^1)}^{2/p} e^{-rs}$$

$$\left| a(s,t) - c - \int_{s_0}^{s} e^{\delta f(x_p(s'))} ds' \right| \le C\|Q(Y)(0,t)\|_{L^2(S^1)}^{2/p} e^{-rs}$$

where inheriting previous notation, we use $x_p(s)$ to denote a gradient trajectory of $\delta f(x)$, the definition of Y is given in the proof. Further, inequalities of the above form continue to hold after we differentiate both sides with respect to (s,t), in other words the inequalities hold in the C^k norm.

Proof. In the course of this proof we first perform some important calculations which we will later reuse for decay estimates over finite gradient trajectories.

Step 0 In our coordinates system the equation looks like (we will drop the δ subscript from u)

$$\partial_s u + J_\delta(u)\partial_t u = 0.$$

Following the Appendix of [7], let's change variables

$$Y := (w, v, x, y)$$

where $w := a(s,t) - s$ and $v := z - t$. Then the equation changes to

$$\partial_s(Y) + J_\delta(u)\partial_t Y + (\partial_s + J_\delta(u)\partial_t) = 0.$$

We simplify this as

$$\begin{aligned}
&\partial_s Y + J(u)\partial_t(Y) + \delta J(u)(\partial_t Y) + (\partial_s + J_\delta(u)\partial_t)\\
=&\partial_s Y + J_0\partial_t(Y) + S_1(x,y,z)(\partial_t Y) + (\partial_s + J_\delta(u)\partial_t) + \delta J\partial_t Y\\
=&\partial_s Y + J_0\partial_t(Y) + S_1(x,y,z)(\partial_t Y) + (\partial_s + J(u)\partial_t) + \delta J(u)\partial_t + \delta J\partial_t Y\\
=&\partial_s Y + J_0\partial_t(Y) + S_1(x,y,z)(\partial_t Y + \partial_t) + \delta J(u)\partial_t + \delta J\partial_t Y\\
=&\partial_s Y + J_0\partial_t(Y) + S_0(x,z)y\partial_t\\
&+ S_1(x,y,z)(\partial_t Y) + S_2(x,y,z)\partial_t + \delta J(u)\partial_t + \delta J\partial_t Y
\end{aligned}$$

We clarify $J(u)$ is the Morse-Bott complex structure evaluated at u, and J_0 is the standard complex structure, which coincides with the Morse-Bott almost complex structure on the surface of the Morse-Bott torus. $\delta J := J_\delta - J$ is the difference between the Morse-Bott complex structure and the perturbed almost complex structure, and as a matrix is has norm bounded above by (its derivatives are also bounded above by) the expression $C\delta$. The term S_0 is a 4 by 4 matrix coming from the linearization of $\bar{\partial}_J$ on the surface of the Morse-Bott torus, and hence it only depends on x, z.

In the above expansion, we have the estimates

$$S_1(x,y,z) \leq C(x,y,z)|y|$$

and

$$S_2(x,y,z) \leq C(x,y,z)y^2$$

We implicitly assume we have taken absolute values of both sides. Similar expressions hold for their derivatives (lowering orders in y as we differentiate).
Next consider the term

$$\begin{aligned}
&\delta J(a,z,x,y)\partial_t u \\
=&\delta J(a,z,x,y)\partial_t Y + \partial_t \\
=&\delta J(u,v,x,0)\partial_t u + \delta T(v,x,y)y\partial_t u \\
=&\delta J(u,v,x,0)(\partial_t Y + \partial_t) + \delta T(v,x,y)y(\partial_t Y + \partial_t)
\end{aligned}$$

where δT is some matrix whose C^k norm is bounded above by $C\delta$. We further observe $\delta J(x,0,v)$ doesn't depend on v since at the surface of Morse-Bott torus it is rotationally symmetric.

We further examine

$$\begin{aligned}
&\delta J(u,v,x,0)(\partial_t Y + \partial_t) \\
=&\delta J(u,0,x,0)\partial_t(a, t+v, x, y) \\
=&\begin{pmatrix} (1-e^{\delta f})\partial_t(t+v) \\ -\frac{e^{\delta f}-1}{e^{\delta f}}\partial_t a \\ -\delta f'(x)\partial_t(t+v) \\ -\delta f'(x)\partial_t a \end{pmatrix}
\end{aligned}$$

where we used the fact J restricted to the surface of the Morse-Bott torus is invariant in the (x,y) direction.
Now we recall our assumptions about the form of $z(s,t)$. In particular we assume $z(t) = t+v$ with $|v| \leq \epsilon$ (this can always be achieved via a reparametrization of the neighborhood around the puncture), so if we plug that in to the above expression it is equal to:

$$\begin{pmatrix} (1-e^{\delta f})\partial_t v \\ -\frac{e^{\delta f}-1}{e^{\delta f}}\partial_t a \\ \delta f'(x)\partial_t v \\ -\delta f'(x)\partial_t a \end{pmatrix} + \begin{pmatrix} 1-e^{\delta f(x(s,t))} \\ 0 \\ -\partial_x \delta f(x(s,t)) \\ 0 \end{pmatrix}.$$

Having performed these computations we return to the overall equation of the form

$$\begin{aligned}
&\partial_s Y + J_0\partial_t(Y) + S_0(x,z)y \\
&+ S_1(x,y,z)(\partial_t Y) + S_2(x,y,z)\partial_t + \delta J(x)\partial_t + \delta T(x,y,z)y(\partial_t Y + \partial_t) + \delta J(u)\partial_t Y.
\end{aligned}$$

For later elliptic regularity purposes it will be useful to write the above in the following form:

$$\partial_s Y + J_\delta(u)\partial_t Y + \delta T y(\partial_t Y + \partial_t) + S_1 \partial_t Y + \delta J(\partial_t) = 0.$$

Step 1 As before consider the operator

$$-A(s) := J_0(d/dt) + S_0(x, z).$$

Note this operator as it appears in the above equation depends on the $x(s,t), z(s,t)$ coordinates of u, but we observe it remains true there exists a λ so that for all functions $h(t) \in W^{1,2}(S^1)$,

$$\langle Ah, Ah\rangle_{L^2(S^1)} \geq \lambda^2 \langle h, h\rangle_{L^2(S^1)}.$$

As a matter of bookkeeping we observe our vector field Y is smooth, hence Y has a well defined restriction to $\{s\} \times S^1$ for any value of s .
We define

$$g(s) = \langle QY, QY\rangle_{L^2(S^1)}$$

as before for our decay estimates we compute

$$g''(s) = 2\langle \partial_s QY, \partial_s QY\rangle + 2\langle QY, \partial_s^2 QY\rangle.$$

We observe both Q and P commute with $\partial_*, * = s, t$.
Step 2 Examining the first term above

$$\begin{aligned}
&\langle \partial_s QY, \partial_s QY\rangle \\
=&\|Q(AY + S_1(x,y,z)(\partial_t Y) + S_2(x,y,z)\partial_t + \delta J(x)\partial_t \\
&+ \delta T(x,y,z) y(\partial_t Y + \partial_t) + \delta J(u)\partial_t Y)\|^2.
\end{aligned}$$

Let's dissect these terms one by one. First since Q commutes with A we have

$$\|AQY\|^2 \geq \lambda \|QY\|^2$$

for some λ independent of s.
We next consider

$$Q\delta J(x(s,t), z(s,t))\partial_t$$

which warrants special treatment. For fixed s we denote by $\bar{x}$ the average value of $x(s,t)$ over t.
Then we can write terms like

$$f(x(s,t)) = f(x(s,t) - \bar{x} + \bar{x}) = f(\bar{x}) + G_x(x - \bar{x}) \leq f(\bar{x}) + G_x(QY)$$

and for $|Y|_{C^0} \leq C\epsilon$ we have $G_x(x) \leq C|x|$. Therefore we observe $Qf(\bar{x}) = 0$ and hence we have the estimate

$$Qf(x(s,t)) \leq CQY.$$

The same also applies to other functions built out of $f(x)$, hence we have

$$\|Q\delta J(x(s,t))\partial_t\| \leq C\delta\|QY\|.$$

Here we also note that the equation satisfied by QY is of the form

$$\partial_s QY + J_\delta(u)\partial_t QY + \delta T(x,y,z)QY \cdot (\partial_t QY + \partial_t) + S_1\partial_t Y + \delta C(x,y,z)QY = 0 \quad (3.20)$$

where $C(x,y,z)$ is just a function of x,y,z whose derivatives are uniformly bounded.

Aside from the two terms we calculated above, applying Q to Y does not have a major impact on other terms. To consider the rest of the terms appearing in $\langle \partial_s QY, \partial_s QY\rangle$ let's estimate their norms (since later we can just use the triangle inequality to either estimate their cross term with themselves or with terms involving $\|AQY\|^2$).
The norms of the terms below after we apply Q

$$S_1(x,y,z)(\partial_t Y), \quad S_2(x,y,z)\partial_t, \quad \delta J(x)\partial_t, \quad \delta T(x,y,z)y(\partial_t Y + \partial_t), \quad \delta J(u)\partial_t Y)$$

are respectively bounded by the norms

$$\epsilon^2(\|QAY\|^2 + \|QY\|^2), \quad ,\epsilon^2\|QY\|^2, \delta^2(\|QY\|^2), \quad \delta^2\epsilon^2(\|AQY\|^2 + \|QY\|^2) + \delta^2\|QY\|^2,$$
$$\delta^2(\|AQY\|^2 + \|QY\|^2).$$

The key observation is $\partial_*^* y \leq \epsilon$ as part of our assumption, as well as the fact all occurrences of y are upper bounded by QY. Another key observation is $\partial_t = \partial_t Q$, so every time we see $\partial_t Y$ we replace it by $\partial_t Q$ hence the appearance of the many Q in the above expression.

Step 3 We look at the next term

$$\begin{aligned}
&\langle QY, \partial_s^2 QY\rangle \\
=&\langle QY, Q\partial_s(AY + S_1(x,y,z)(\partial_t Y) + S_2(x,y,z)\partial_t + \delta J(x)\partial_t \\
&+ \delta T(x,y,z)y(\partial_t Y + \partial_t) + \delta J(u)\partial_t Y)\rangle.
\end{aligned}$$

Note for terms we think of being small, we are not careful about their signs. We introduce some more convenient notation. We write the J-holomorphic curve equation as

$$\partial_s Y - AY + E(Y) = 0$$

Then we have

$$\begin{aligned}
&\langle QY, Q\partial_s(AY + E(Y))\rangle \\
=&\langle QY, \epsilon QY\rangle + \langle QY, Q[A(AY + E) + \partial_s E]\rangle \\
=&\langle QY, \epsilon QY\rangle + \langle QAY, QAY\rangle + \langle QAY, QE\rangle + \langle QY, Q\partial_s E\rangle.
\end{aligned}$$

To obtain the first term in the above expression we used the fact that

$$\partial_s A$$

is a 4 by 4 matrix whose only nonzero entry is the diagonal entry corresponding to y, so

$$Q(\partial_s AY) = \epsilon y.$$

The only term we don't know how to control is the last one $\langle QY, Q\partial_s E\rangle$, the previous ones follow from computation in previous steps. Let's recall the terms in E:

$$S_1(x,y,z)(\partial_t Y), \quad S_2(x,y,z)\partial_t, \quad \delta J(x)\partial_t, \quad \delta T(x,y,z)y(\partial_t Y+\partial_t), \delta J(u)\partial_t Y.$$

We need to compute the $L^2(S^1)$ norm of these terms after we take their s derivative. We first only consider the s derivatives on $S_1, \delta T, \delta J(u)$, by assumption that $\partial_*^k y \leq \epsilon$, when we take the s derivatives of $S_1, \delta T, \delta J(u)$, they are still operators of the same form. For example $\partial_s S_1$ is of the form $C_1(x,y,z)y + C_2(x,y,z)y_s$, and the norm of each term can be bounded above by ϵ. The same can be said about $\partial_s \delta T$, $\partial_s \delta J(u)$, so by abuse of notation we use the same symbols. Then techniques from previous steps immediately show the norm of these terms are upper bounded by

$$\{\partial_s(S_1)\partial_t Y, \quad \partial_s(\delta T y)\partial_t Y, \quad \partial_s(\delta J)\partial_t Y\} \leq C(\epsilon+\delta)\|Q\partial_t Y\|^2 \leq C(\epsilon+\delta)(\|AQY\|^2 + \|QY\|^2).$$

We next consider the s derivative of $Q\delta J(z,x)\partial_t$. We first observe Q commutes with ∂_s so we are evaluating $Q\partial_s \delta J(z,x)\partial_t$. Recalling the previous form of this vector field, the components are essentially built out of $f(x(s,t))$, so we need to take its s derivative and projection via Q. Using the previous trick of introducing $\bar{x}$

$$\begin{aligned}
&\partial_s f(x(s,t))\\
=&\partial_s f(x(s,t) - \bar{x}(s) + \bar{x}(s))\\
=&f_x(x(s,t) - \bar{x}(s) + \bar{x}(s))(x_s(s,t) - \bar{x}_s + \bar{x}_s)\\
=&f_x(x)(x_s(s,t) - \bar{x}_s) + \bar{x}_s f_x(x(s,t) - \bar{x}(s) + \bar{x}(s))\\
=&f_x(x)(Qx_s) + \bar{x}_s[f_x(\bar{x}_s) + G_x(Qx)].
\end{aligned}$$

Observe $Q(\bar{x}_s f_x(\bar{x}_s)) = 0$ because this term doesn't depend on t. Hence pointwise

$$Q\partial_s f(x_s(t)) \leq C|QY| + |QY_s|.$$

Hence:

$$\|\partial_s Q\delta J\partial_t\|_{L^2(S^1)} \leq C\delta(\|QY\|_{L^2} + \|QY_s\|_{L^2})$$

and we have seen above how to bound the norm of $\|QY_s\|_{L^2}$. Next:

$$\partial_s S_2 \partial_t = Cyy_s\partial_t.$$

We assumed $y_s \leq \epsilon$ this term can be upper bounded by

$$C\epsilon QY.$$

Finally we turn our attention to terms of the form

$$\epsilon \partial_s \partial_t Y$$

which appear once the s derivative hits QY. Here ϵ denotes a matrix whose C^{k-1} norm is uniformly upper bounded by the real number ϵ. We can insert a factor of Q after the t derivative and get

$$\begin{aligned}
&\langle QY, \epsilon \partial_s \partial_t QY \rangle \\
=&\langle \epsilon^T Q\partial_t Y, \partial_s QY \rangle + \langle \epsilon_t^T QY, \partial_s QY \rangle \\
\leq& \epsilon(\|Q\partial_t Y\|^2 + \|\partial_s QY\|^2 + \|QY\|^2).
\end{aligned}$$

The terms in the last line are already well understood by previous computations. In particular $\|\partial_s QY\|^2$ was worked out in the previous step and $\|Q\partial_t Y\|^2$ was worked out in this step. Hence putting all of these terms together we have

$$g''(s) \geq (4\lambda^2 - C\epsilon) g(s)$$

hence from previous lemma we have $g(s) \leq g(0)e^{-\lambda s}$, hence the L^2 norm of QY undergoes exponential decay. That this extends to pointwise C^k norm follows from elliptic regularity, using equation 3.20.

Step 4 In this step we look at what equation PY satisfies. Let's recall the original equation

$$\partial_s(Y) + \partial_s + J_\delta(u)\partial_t Y + J_\delta(u)\partial_t = 0.$$

We split $Y = QY + PY$ and plug into the above equation to get the pointwise bound:

$$|\partial_s PY + \delta J \partial_t| \leq C\|Q(Y)(0,t)\|_{L^2(S^1)}^{2/p} e^{-\lambda s}.$$

where we used the previous bound on the norm of QY. We remind ourselves λ might change from previous parts because of various changes in norm. We can replace $f(x)$ with $f(Px)$ because $\partial_t x$ is bounded by $\partial_t QY$, which decays exponentially, so we can take the error term to the right hand side to get

$$|\partial_s PY + \delta J(PY)\partial_t| \leq C\|Q(Y)(0,t)\|_{L^2(S^1)}^{2/p} e^{-\lambda s}.$$

Observe that the function PY only depends on s, and the above are pointwise inequalities. Differentiability of PY comes from bootstrapping and observing the differentiability of QY in the s variable. The decay estimates of the higher order s derivatives follow as well. We let PY_*, where $* = a, z, x, y$ denote the various components of PY. The equations in these coordinates are

$$PY_y = 0$$

$$|\partial_s PY_z| \leq C\|Q(Y)(0,t)\|_{L^2(S^1)}^{2/p} e^{-\lambda s}$$

$$|\partial_s PY_x - \delta f'(PY_x)| \leq C\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}$$
$$|\partial_s PY_a - e^{\delta f(PY_x)}| \leq C\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}.$$

We now solve the above inequalities. For brevity we denote by $G(s)_*$ the expression on the right hand side for $* = z, x, a$, and the only property we will need about $G(s)$ is that it is asymptotically of the form $e^{-\lambda s}$. The inequality

$$|\partial_s PY_z| \leq C\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}$$

integrates to

$$|PY_z - c| = \left|\int_0^s G_z(s')ds'\right| \leq C\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}.$$

Next $|\partial_s PY_x - \delta f'(PY_x)| \leq C\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}$, we pick a coordinate neighborhood so that $f(x) = \mp\frac{1}{2}x^2 + C$. We can do this because we know u eventually limits to a critical point of f as $s \to \infty$ and stays away from all other critical points of f, the choice of $\mp$ corresponds to whether we are in a neighborhood of maximum or minimum of f. Then this is an equation of the form

$$\partial_s PY_x \pm \delta PY_x = G(s)_x.$$

Then we have

$$(PY_x e^{\pm\delta s})_s = G_x(s)e^{\pm\delta s}.$$

We can write

$$PY_x = c(s)e^{\mp\delta s}$$

where $c(s)$ satisfies the equation

$$\frac{d}{ds}c(s) = G(s)_x e^{\pm\delta s}.$$

Since we known $G(s)_x$ decays quickly when $s \to \infty$, the function c must have a limit as $s \to \infty$, call this limit c_∞. Then we have

$$c(s) = c_\infty + \int_s^\infty G_x(t)_x e^{\pm\delta t} dt$$

hence

$$PY_x(s) = c_\infty e^{\mp\delta s} + e^{\mp\delta s}\int_s^\infty G_x(t)e^{\pm\delta t}dt.$$

We recognize $c_\infty e^{\mp\delta s}$ as the gradient flow $x_p(s)$ we identified earlier and $e^{\mp\delta s}\int_s^\infty G(s)_x e^{\pm\delta t}$ is considered the error term, and by the form of G_x the error term has the decay we needed.

We note in the case $f = -\frac{1}{2}x^2 + C$ the gradient flow converges to zero, and this corresponds to "free" ends converging to the maximum of f on positive punctures. In the case where $f = +\frac{1}{2}x^2 + C$, the gradient flow segment $c_\infty e^{\delta s}$, if we have $c_\infty \neq 0$, will actually flow away from the critical point $x = 0$, so it will eventually leave the neighborhood where the expression $f(x) = \frac{1}{2}x^2 + C$ is valid, and instead flow to the other critical point/maximum of f, for which we can use the above analysis directly. The exception is if $c_\infty = 0$, and this end will converge to the $x = 0$, or the minimum of f. This corresponds to the case of a "fixed" end converging to the minimum of f. Implicit in the above discussion is the assumption that u_δ stays away from all except one critical point of f uniformly as $\delta \to 0$. This, in the language of our equations, means $c_\infty(\delta)$ (this constant implicitly depends on δ), is either bounded away from zero for all δ small enough, or is identically zero for δ small enough. These correspond respectively to the above two cases. The case where $c_\infty(\delta) \to 0$ and $c_\infty(\delta) \neq 0$ as $\delta \to 0$ corresponds to the J_δ-holomorphic curves u_δ breaking into a cascade of height > 1, and is outside the scope of our discussion.

Finally we consider the equation

$$\partial_s PY_a - e^{\delta f(Px)} = G_a(s).$$

Now by the above estimate on $P(x)$, there is a gradient trajectory v whose x component, $\pi_x v$ is approximated by PY_x, in the sense that

$$|PY_x - \pi_x v| \leq C\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}.$$

Then for small enough ϵ, we have the estimate

$$|e^{\delta f(PY_x)} - e^{\delta f(\pi_x v)}| \leq C\delta \|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}$$

hence we can write

$$\partial_s P_a - e^{\delta f(\pi_x v)} = G_a(s)$$

where we absorbed the error term $C\delta\|Q(Y)(0,t)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}$ into $G_a(s)$ since they are of the same form. Then we integrate both sides to get:

$$P_a(s) - \int_0^s e^{\delta f(\pi_x v)} = \int_0^s G_a(s')ds'.$$

Using the same trick as before we write $\int_0^s G_a(s')ds' = c_\infty - \int_s^\infty G_a(s')ds'$, recognizing $c_\infty + \int_0^s e^{\delta f(\pi_x v)}$ is the a component of a lift of a gradient trajectory, we arrive at the desired bound. $\square$

Remark 3.10.2. In the above proof and what follows we assume that u_δ stays uniformly away from all but one critical point of f. The estimates for $Q(Y)$ is largely unaffected by this assumption, the main reason we use this is so that we could have nice exponential decay estimates for PY (this is where we used local form of f). In general (and in manifolds where the critical Morse-Bott manifolds are higher dimensional) we could have u_δ degenerate into a broken trajectory of f along the critical set, and the estimate there is more involved. Fortunately our transverse rigid constraint means our assumption about u_δ being away from except at most one critical point of f will be sufficient for our purposes.

Finite gradient segments

We now extend these exponential decay estimates to finite gradient trajectories.

Theorem 3.10.3. *Consider an interval $I = [s_0, s_1]$ and a J_δ-holomorphic curve u so that when restricted to $s \in I$ the map u is close to the Morse-Bott torus, i.e. in a neighborhood of the Morse-Bott torus u has coordinates (a, z, x, y) and the functions a, z, x, y satisfy*

$$|y|, |\partial_*^{\leq k}(z-t)|, |\partial_*^{\leq k} x|, |\partial_*^{\leq k} y| \leq \epsilon$$

for some $\epsilon > 0$ depending only on the local geometry and independent of δ, then

$$\|QY\|_{C^{k-1}} \leq \max(\|QY(s_0,t)\|_{L^2(S^1)}^{2/p}, \|QY(s_1,t)\|_{L^2(S^1)}^{2/p}) \frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0)/2)}$$

for some $\lambda > 0$ only depending on the local geometry.
If u is uniformly bounded away from all critical points of f except maybe one, there is a lift of a gradient trajectory, which we denote by v, so that

$$\|PY - v\|_{C^{k-1}} \leq \max(\|QY(s_0,t)\|_{L^2(S^1)}^{2/p}, \|QY(s_1,t)\|_{L^2(S^1)}^{2/p}) \frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0)/2)}.$$

Proof. The proof will follow the general thread of the semi-infinite case. We recall our convention λ may change from line to line, but not in a fashion that depends on δ. Recall we defined the function

$$g(s) := \langle QY, QY \rangle_{L^2(S^1)}$$

then we have the inequality

$$g'' \geq \lambda^2 g.$$

We define the auxiliary function

$$k(s) := \max(\|QY(s_0)\|_{L^2(S^1)}^2, \|QY(s_1)\|_{L^2(S^1)}^2) \frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0))}$$

then we have the inequality

$$(g-k)'' \geq \lambda^2 (g-k).$$

Then $g - k$ cannot have positive maximum, and by construction $g - k \leq 0$ at $s = s_0, s_1$. Hence $g \leq k$ globally for $s \in I$.

With elliptic regularity as before, we obtain the pointwise bound

$$|Q(Y)(s,t)| \leq k(s)^{1/p}$$

which by elliptic regularity can be improved to bound the derivatives of QY. Using the inequalities

$$c_1 \cosh(x/p) \leq \cosh(x)^{1/p} \leq c_2 \cosh(x/p).$$

We then obtain inequalities:

$$\|QY\|_{C^{k-1}} \leq \max(\|QY(s_0,t)\|^{2/p}_{L^2(S^1)}, \|QY(s_1,t)\|^{2/p}_{L^2(S^1)})\frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0)/2)}$$

where we have of course changed the definition of λ. We also have

$$|\partial_s PY - J_\delta \partial_t| \leq k_1$$

where for brevity we have defined

$$k_1 = max(\|QY(s_0,t)\|^{2/p}_{L^2S^1}, \|QY(s_1,t)\|^{2/p}_{L^2S^1})\frac{cosh(\lambda(s-(s_0+s_1)/2))}{cosh(\lambda(s_1-s_0)/2)}.$$

We try to integrate this inequality as before: $|\partial_s PY - J_\delta \partial_t| \leq k_1$. There are various components to this equation, which we examine one by one. For the easiest case we have:

$$|\partial_s PY_z| \leq k_1.$$

Integrating both sides we get

$$\begin{aligned}
|PY_z(s) - PY_z((s_0+s_1)/2)| \leq & \int_{(s_0+s_1)/2}^{s} k_1 \\
\leq & C\frac{\max(\|QY(s_0,t)\|^{2/p}_{L^2S^1}, \|QY(s_1,t)\|^{2/p}_{L^2S^1}))}{\cosh(\lambda(s_1-s_0)/2)} \\
& \cdot \int_{(s_0+s_1)/2}^{s} \cosh(\lambda(s'-(s_1+s_0)/2)ds' \\
\leq & C\frac{\max(\|QY(s_0,t)\|^{2/p}_{L^2S^1}, \|QY(s_1,t)\|^{2/p}_{L^2S^1}))}{\cosh(\lambda(s_1-s_0)/2)} \\
& \cdot |\sinh(\lambda(s-(s_0+s_1)/2))| \\
\leq & C\max(\|QY(s_0,t)\|^{2/p}_{L^2S^1}, \|QY(s_1,t)\|^{2/p}_{L^2S^1}) \\
& \cdot \frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0)/2}.
\end{aligned}$$

Identifying $PY_z((s_0+s_1)/2$ as a constant, we obtain the required estimate. We next examine:

$$|\partial_s PY_x - (\partial_x \delta f)(PY_x)| \leq k_1.$$

For segments of gradient flow uniformly away from all critical points of f, then we can choose our coordinates so that locally $f(x) = x + c$. Then the above equation takes the form:

$$|\partial_s PY_x - (\partial_x \delta f)(PY_x)| \leq k_1.$$

Using the exact same techniques as above, we conclude

$$\begin{aligned}&|PY_x(s) - PY_x((s_0+s_1)/2) - (s-(s_0+s_1)/2)|\\ \leq& C\max(\|QY(s_0,t)\|^{2/p}_{L^2S^1}, \|QY(s_1,t)\|^{2/p}_{L^2S^1})\frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0)/2}.\end{aligned}$$

Identifying $PY_x((s_0+s_1/2)) - (s-(s_0+s_1))/2$ as the x component of a lift of the gradient flow, the conclude the required estimate.
If u is uniformly bounded away from all critical points of f except one, then we can only choose coordinates so that $f(x) = \frac{1}{2}x^2$ (the case for $f(x) = -\frac{1}{2}x^2$ is similar), then the above equation takes the form

$$|\partial_s PY_x - \delta PY_x| \leq k_1.$$

Recyling notation from the previous proof we get

$$\partial_s PY_x - \delta PY_x = G_x(s)$$

where $G_x(s) \leq k_1(s)$ Using integration factors as before we obtain

$$(PY_x e^{-\delta s})' = G_x(s)e^{-\delta s}.$$

Integrating both sides, from $(s_0+s_1)/2$ to s

$$PY_x = c(s)e^{\delta(s-(s_0+s_1)/2)}$$

where $c(s)' = G_x(s)e^{-\delta s}$. Then

$$c(s) = c_0 + \int_{s_0+s_1/2}^{s} G_x(s')e^{-\delta s'}ds'$$

$$|PY_x - c_0 e^{\delta M(s-(s_0+s_1)/2)}| \leq e^{\delta(s-(s_0+s_1)/2)}\int_{(s_1+s_0)/2}^{s} G_x(s')e^{-\delta s'}ds'.$$

Here we need be a bit careful about this integral, by our assumptions on $G_x(t)$ it is upper bounded by:

$$G_x(t) \leq C\max(\|QY(s_0,t)\|^{2/p}_{L^2S^1}, \|QY(s_1,t)\|^{2/p}_{L^2S^1})\frac{\cosh(\lambda(s-(s_0+s_1)/2))}{\cosh(\lambda(s_1-s_0)/2}.$$

WLOG we assume $s>0$ and $(s_0+s_1)/2 = 0$, then we have the inequalities:

$$C'\cosh(\lambda s) \leq e^{\lambda s} \leq C\cosh(\lambda s).$$

Then the integral

$$\begin{aligned}
&e^{\delta(s)}\int_0^s G_x(s')e^{-\delta s'}ds'\\
&\leq e^{\delta s}C\frac{max(\|QY(s_0,t)\|^{2/p}_{L^2(S^1)},\|QY(s_1,t)\|^{2/p}_{L^2(S^1)})}{cosh(\lambda(s_1-s_0)/2)}\frac{e^{(\lambda-\delta)s}-1}{\lambda-\delta}\\
&\leq Cmax(\|QY(s_0,t)\|^{2/p}_{L^2(S^1)},\|QY(s_1,t)\|^{2/p}_{L^2(S^1)})\frac{cosh(\lambda s)}{cosh(\lambda(s_1-s_0)/2)}
\end{aligned}$$

which is exactly our estimate. The same works for $s<0$.

Finally we consider $|\partial_s PY_a - e^{\delta f(PY_x)}| \leq k_1(s)$. As before we replace $f(PY_x)$ with $f(\pi_x v)$, which introduces an error of the same form as k_1 due to our above estimate, so we simply absorb it into k_1 on the right hand side, we then integrate both sides to get

$$\begin{aligned}
&\left|PY_a + c - \int_{s_0+s_1/2}^{s} e^{\delta f(\pi_x v)}\right|\\
&\leq C\frac{\max(\|QY(s_0,t)\|^{2/p}_{L^2(S^1)},\|QY(s_1,t)\|^{2/p}_{L^2(S^1)})}{\cosh(\lambda(s_1-s_0)/2)}\cosh(\lambda(s-(s_0+s_1)/2))
\end{aligned}$$

from this we conclude the proof. □

3.11 Surjectivity of gluing

In previous sections we proved every transverse rigid cascade glues to a J_δ-holomorphic curve of Fredholm index 1. In this section we show this gluing is unique, i.e. if a J_δ-holomorphic curve is sufficiently close to the cascade, then it must have come from our gluing construction. The main strategy is to consider a degeneration $u_\delta \to u^{\not{z}} = \{u^i\}$ of a J_δ-holomorphic curve u_δ into a cascade $\{u^i\}$. Using the compactness results stated in Section 11.2 of [6] (See also Chapter 4 of [5]) and proved in our appendix, we know the convergence is C^∞_{loc}, using our local estimates we show u_δ corresponds to a solution of our tuple of equations

$$\mathbf{\Theta}_u = 0, \quad \mathbf{\Theta}_v = 0.$$

Here we use $\mathbf{\Theta}_u = 0$ to denote the system of equations over the J-holomorphic curves in the cascade, and $\mathbf{\Theta}_v = 0$ denotes the system of equations over each gradient trajectory (finite or semi-infinte) that appear in the cascade. Furthermore, we show we can arrange that vector fields among the equations in $\mathbf{\Theta}_v = 0$ that correspond to finite gradient trajectories all live in H_0, the codimension 3 subspace we fixed for each finite gradient trajectory (we abuse notation slightly, there is a different H_0 for each different finite gradient trajectory). We showed in the gluing section such vector fields in H_0 are unique. We also make a choice of right inverse for $\oplus D_i$ for the system $\mathbf{\Theta}_u = 0$, and show we can arrange so that the vector field producing u_δ lands in the image of said right inverse for $\oplus D_i$. Therefore from the uniqueness of our gluing construction there is a 1-1 correspondence between J_δ-holomorphic curves and cascades.

The outline of this section is as follows. We will first focus on the simplest possible case: a two level cascade $u^{\natural} = \{u^1, u^2\}$ meeting along a single Reeb orbit in the intermediate cascade level. Even in this simplified setting there are several stages to our construction: we first use the previous decay estimates to show that u_δ is in an ϵ neighborhood of a preglued curve constructed from the cascade $u^{\natural}$. Then we adjust the pregluing using the asymptotic vectors so that the vector field over the finite gradient trajectory v lives in H_0, and the part of the vector field living over u^i lives in the preimage of our specified right inverse, while maintaining the fact the vector field still lives in the ϵ ball. Finally we extend the vector fields over all of $u^{i*}TM$ and v^*TM so that they become solutions to $\mathbf{\Theta_u} = 0$ and $\mathbf{\Theta_v} = 0$, using tools from Section 7 of [41]. We also develop some properties of linear operators for this purpose.

Then after the 2-level cascades case has been thoroughly analyzed and proper tools developed, we introduce some more elaborate notation to set up the more general n-level cascade case.

Notation and setup, for 2-level cascades

We note here that we are not proving the SFT compactness statement, we are simply using it. For ease of exposition, we first describe the case with u_δ degenerating into a 2 level cascade consisting of u^1 and u^2 and such that they only have 1 intermediate end meeting in the cascade level. We let $\gamma_1 := ev^-(u^1)$, and $\gamma_2 := ev^+(u^2)$ denote the Reeb orbits on the Morse-Bott torus. We fix domains Σ_1 and Σ_2 for u^1 and u^2. We fix cylindrical neighborhoods near punctures of Σ_i, and let $(s,t)_i$ denote coordinates near the puncture that meet along the intermediate Morse-Bott torus. We let also let $(s,t)_i'$ denote the cylindrical coordinates on u^i that are on punctures away from the Morse-Bott torus that appear in the intermediate cascade level. Recall a neighborhood of the maps u^i is given by

$$W^{2,p,d}(u^{i*}TM) \oplus V_i \oplus V_i' \oplus T\mathcal{J}_i.$$

We let Σ_δ denote the domain for u_δ. Then by the analog of SFT compactness, for each δ we can break down the domain Σ_δ into 3 regions,

$$\Sigma_\delta = \Sigma_{\delta+} \cup N_\delta \cup \Sigma_{\delta-}$$

where we think of $\Sigma_{\delta\pm}$ as regions that converge to Σ_i, and N_δ the thin region biholomorphic to a very long cylinder that converges to the finite (yet very long) gradient trajectory connecting u^1 and u^2. To be more precise, we can translate u_δ globally so that over $\Sigma_{\delta+}$ the map u_δ converges in C^∞_{loc} to u^1, and there exists a sequence of a translations, we denote by a_δ, so that after we translate u^2 by a_δ, which we denote by $u^2 + a_\delta$, the map u_δ when restricted to $\Sigma_{\delta-}$ converges in C^∞_{loc} to $u_2 + a_\delta$. Technically the convergences to u^1 and u^2 are over compact subsets of $\Sigma_{\delta\pm}$, near the other punctures $(s,t)_i'$ there are additional convergences to semi-infinite gradient trajectory. Here we only concern ourselves with convergences near N_δ, and worry about semi-infinite gradient trajectories in a later section.

Finding appropriate vector fields

We first consider the degeneration in the intermediate cascade level. We will later consider degeneration to the configuration of a semi-infinite gradient trajectory.

Finding a global vector field

Let $0 < \epsilon' << \epsilon$, the specific size of ϵ' will be specified in the course of the construction. We fix a large real number $K > 0$, then we consider the region $|s_i| \le K, |s_i'| \le K$ as subsets of Σ_i. We denote this compact subset of the domain by Σ_{iK}. We take K large enough so that for $|s_i| \ge K$ the maps u^i are in a small enough neighborhood of γ_i, that up to k derivatives, we can think of u^i as exponentially decaying to trivial cylinders, with exponential decay bounded by e^{-Ds_i}.

This choice of K also determines a decomposition of the domain of u_δ, to wit

$$\Sigma_\delta = \Sigma_{+\delta K} \cup N_{\delta K} \cup \Sigma_{-\delta K}.$$

Then the convergence statement in C^∞_{loc} implies there are vector fields $\zeta_{i\delta} \in u^{i*}TM|_{\Sigma_{iK}}$ of C^1 norm $< \epsilon'$ and variation of complex structure $\delta j_i \in T\mathcal{J}_i$ of size $\le \epsilon'$ so that

$$u_\delta|_{\Sigma_{\delta+K}} = exp_{u^1,\delta j_1}(\zeta_{1\delta})$$

and

$$u_\delta|_{\Sigma_{\delta-K}} = \exp_{u^2,\delta j_2}(\zeta_{2\delta}).$$

We shall for now suppress the variation of complex structure $(u^i, \delta j_i)$ and simply write u^i. When later we want to include it in the notation we shall write $(u^i, \delta j_i)$. We also recall that our metric is flat around Morse-Bott tori, so for small enough $\zeta_{i\delta}$, we have $\exp_{u^i}(\zeta_{i\delta}) = u^i + \zeta_{i\delta}$ near Morse-Bott tori.

We here simply note the $W^{2,p,d}$ norm of $\zeta_{i\delta}$ is then bounded above by $C\epsilon' e^{dK}$. For fixed K, as $\delta \to 0$, by C^∞_{loc} convergence we can take $\epsilon'(\delta) \to 0$ to make this expression as small as we please. We also observe for fixed K and small enough ϵ' the deformations $(\zeta_{i\delta}, \delta j_i)$ are within an ϵ ball of $W^{2,p,d}(u^{i*}TM) \oplus V_i \oplus V_i' \oplus T\mathcal{J}_i$. We next consider the behaviour of u_δ when restricted to the neck region N_δ. We first informally write $N_{\delta K}$ as the cylinder $[0, N_{\delta K}] \times S^1$. We start with the following lemma:

Lemma 3.11.1. *By our assumption as $K \to \infty$ (which would take $\delta \to 0$ with it in order to satisfy our previous assumptions) we have $u_\delta|_{N_{\delta K}}$ converges in C^∞_{loc} to trivial cylinders. This is also true uniformly, i.e. for given $\epsilon'' > 0$, there is a K large enough so that for every small enough values of $\delta > 0$, $u_\delta|_{[k,k+1]\times S^1}$ is within ϵ'' (in the C^k norm) of a trivial cylinder of the form $\gamma \times \mathbb{R}$ for all values of k so that $[k, k+1] \times S^1 \subset N_{\delta K} \times S^1$.*

Proof. **Step 1** We claim for K large enough $|du_\delta| < C$ for all of $N_{\delta K}$. Suppose not, then we can find a sequence (s_δ, t_δ) where $|du_\delta(s_\delta, t_\delta)| \to \infty$, by Gromov compactness a holomorphic plane bubbles off. But a holomorphic plane must have energy bounded below, by the Morse-Bott assumption. However as $K \to \infty$ the energy of $u_\delta|_{N_{\delta K}}$ goes to zero, which in particular is less than the minimum energy required to have a holomorphic plane, this is a contradiction.

Step 2 We argue by contradiction, Suppose for all $K > 0$ there exists an interval $[a_K, a_K + 1] \times S^1$ so that the distance of $u_\delta|_{[a_K, a_K+1]\times S^1}$ and any trivial cylinder is $\geq \epsilon''$. However we observe as $K \to \infty$ the energy of $u_\delta|_{[a_K, a_K+1]\times S^1}$ goes to zero uniformly in K, then by Azerla-Ascoli this converges to a holomorphic curve of zero area, which must be a segment of a trivial cylinder. Hence we have a contradiction. □

Then the previous convergence estimate implies the following:

Proposition 3.11.2. *We take $\epsilon'' > 0$ small enough so that previous convergence estimates near Morse-Bott tori apply. Then there is a large enough K, and small enough ϵ' (which depends on K), so that if we choose small enough $\delta > 0$ (which depends on the choice of ϵ' and K but can always be achieved), there is a gradient trajectory v_K defined over the cylinder $(s_v, t_v) \in [0, N_{\delta K}] \times S^1$ so that there is a vector field ζ_K over v_K so that*

$$u_\delta|_{N_{\delta K}} = \exp_{v_K}(\zeta_K)$$

and the norm of ζ_K satisfies

$$\|\zeta_K\|_{C^{k-1}} \leq C \max(\|\zeta_K(0,-)\|^{2/p}_{L^2(S^1)}, \|\zeta_K(N_{\delta K}, -)\|^{2/p}_{L^2(S^1)}) \frac{\cosh(\lambda(s - N_{\delta K}/2))}{\cosh(\lambda N_{\delta K}/2)}$$

and in particular, if we choose $\delta > 0$ small enough, by C^∞_{loc} convergence

$$\|\zeta_K(0)\|^{2/p}, \|\zeta_K(N_{\delta K})\|^{2/p} \leq \epsilon'.$$

We estimate the norm of ζ_K for later use. With some foresight we realize we need to use a weighted norm $e^{w(s_v)}$ for $s_v \in [0, N_{\delta K}]$, where

$$w(s) = d(N_{\delta K}/2 + K - |s - N_{\delta K}/2|).$$

Then we measure the $W^{1,p}$ norm of ζ_K with respect to $e^{w(s)}$, but by the previous proposition the norm of ζ_K undergoes exponential decay as it enters the interior of $N_{\delta K}$. Hence we have

$$\int_{S^1} \int_0^{N_{\delta K}} \|\zeta_K\|_{C^k} ds dt \leq C\epsilon' e^{dK}.$$

We now come to the main result of this subsection. We combine $\zeta_{i\delta}$ and ζ_K into a vector field over some preglued curve built from $\Sigma_{\delta \pm K}$, the curve v_K and some asymptotic vector fields. We first recall that

$$u^1(-K, t) + \zeta_{1\delta}(-K, t) = u_\delta(N_{\delta K}, t) = v_K + \zeta_K(N_{\delta K}, t).$$

Now we have the C^k norm of $\zeta_{1\delta}(-K,t)$ and $\zeta_K(0,t)$ are both bounded above by ϵ', then we can deform $u^1|_{\Sigma_{1K}}$ by asymptotic vectors r_1, a_1, p_1 all of which are of size $\leq \epsilon'$ so that

$$\|u^1(-K,t_1) + (r_1,a_1,p_1) - v_K(N_{\delta K}, t_v)\|_{C^k} \leq \epsilon'.$$

There are naturally several possible choices possible for (r_1, a_1, p_1). In anticipation of our later constructions, we make the following important specification.
Recall for $(s_1,t_1) \in (-\infty, 0] \times S^1$, for $s_1 << 0$ the map u^1 converges to a parametrized trivial cylinder

$$\tilde{\gamma}_1(s_1,t_1) : \mathbb{R} \times S^1 \longrightarrow M$$

whose image is of course the trivial cylinder $\gamma_1 \times \mathbb{R}$. The key property is that $u^1|_{(-\infty,0]\times S^1}$ decays exponentially to $\tilde{\gamma}_1$:

$$\|u^1 - \tilde{\gamma}_1\|_{C^k} \leq Ce^{-Ds_1}.$$

We also recall properties of v_K, which is the finite gradient trajectory u_δ converges to. For small enough $\delta > 0$, the gradient flow is extremely slow, so for $s_v \in [N_{\delta K} - 2R, N_{\delta K}]$, there is another trivial parametrized cylinder

$$\hat{\gamma}_1(s_v,t_v) : \mathbb{R} \times S^1 \longrightarrow M$$

so that for $s_v \in [N_{\delta K} - 2R, N_{\delta K}]$

$$\|\hat{\gamma}_1 - v_K\|_{C^k} \leq CR\delta$$

which goes to zero as $\delta \to 0$. By the comparison result above there are vectors $(r_1, a_1, p_1) \leq \epsilon'$ so that:

$$\tilde{\gamma}_1 + (r_1,a_1,p_1) = \hat{\gamma}_1.$$

Then we choose this particular choice of (r_1, a_1, p_1). There is some free choice of (r_1, a_1, p_1) up to size $R\delta$, which for our purpose is extremely small. We will always make a choice so that the $s_1 = R$ end of $u^1 + (r_1, a_1, p_1)$ and $s_v = N_{\delta K} - R + K$ of v_K can be preglued together, in the sense we preglued them together in Section 3.9. (Also see below).

Similarly we recall that

$$u^2(K,t) + \zeta_{2\delta}(K,t) = u_\delta(0,t) = v_K + \zeta_K(0,t_v).$$

By the same reasoning there is a parametrized trivial cylinder $\tilde{\gamma}_2 : \mathbb{R} \times S^1 \to M$ that u^2 decays exponentially to:

$$\|u_2 - \tilde{\gamma}_2\|_{C^k} \leq Ce^{-Ds_2}.$$

And we can find parametrized trivial cylinder $\hat{\gamma}_2$ so that for $s_v \in [0, 3R]$ we have

$$\|\hat{\gamma}_2 - v_K\|_{C^k} \leq CR\delta.$$

Hence by comparison we choose asymptotic vectors (r_2, a_2, p_2) of size bounded above by ϵ' over u^2 so that

$$\|u^2(K,t_2) + (r_2,a_2,p_2) - v_K(0,t)\|_{C^k(S^1)} \leq \epsilon'.$$

The trivial cylinders satisfy the relation

$$\tilde{\gamma}_2 + (r_2, a_2, p_2) = \hat{\gamma}_2.$$

Observe since v_K as a parametrized cylinder does not rotate in the z direction, here we have $r_1 = r_2$. We note this here because in our gluing construction earlier where we identified $t_v \sim t_- + r_+ - r_-$. We shall see where this is used in a later section.

Then we construct the preglued domain by gluing together

$$\Sigma_{\delta,K,(r,a,p)_i} := (u^1, \delta j_i) + (r_1, a_1, p_1)|_{\Sigma_{1R}} \cup [R-K, N_{\delta K} + K - R] \times S^1 \cup (u^2, \delta j_2) + (r_2, a_2, p_2)|_{\Sigma_{2R}}$$

by Σ_{1R} we mean the domain of u^1 with $s_1 < -R$ removed. In other words $\Sigma_{1R} := \Sigma_K \cup (s_1, t_1) \in [-R, -K] \times S^1$. (We ignore the ends of u^1 glued to semi-infinite trajectories for now). A similar expression holds for Σ_{2R}. By $(u_i, \delta j_i) + (r_i, a_i, p_i)$ we mean Σ_{iR} with complex structure deformed by δj_i and the cylindrical neck twisted/stretched/translated by asymptotic vector fields (r_i, a_i, p_i). We specify the gluing as follows. We glue together

$$[u^1 + (r_1, a_1, p_1)](s_1 = -R, t_1) \sim v_K(s_v = N_{\delta K} - R + K, t_v).$$

Using the same pregluing interpolation as we did in our pregluing construction. At u^2 end we are making the identification

$$[u^2 + (r_2, a_2, p_2)(s_2 = R, t_2) \sim v_K(s_v = R - K, t_v)$$

and this determines our preglued domain, $\Sigma_{\delta,K,(r,a,p)_i}$. In constructing this preglued domain, we have identified:

$$-s_1 - R \sim s_v - N_{\delta K} - R + K$$

$$s_2 - R \sim s_v - R - K$$

$$t_1 \sim t_v \sim t_2.$$

Since $r_1 = r_2$, here the t_v is identified with t_2 without any twist. It carries a natural preglued map into M by defining it to be $(u^i, \delta j_i) + (r_i, a_i, p_i)$ on Σ_{iR} and v_K on $[R - K, N_{\delta K} - R + K] \times S^1$, and interpolated in the pregluing region the same way we preglued in Section 3.9. We call the preglued map $u_{\delta,K,(r,a,p)_i}$. Then we can form the interpolation of the vector fields $\zeta_{i\delta}$ and ζ_K into a vector field we call $\zeta_{\delta,K,(r,a,p)_i}$ so that

$$u_\delta = u_{\delta,K,(r,a,p)_i} + \zeta_{\delta,K,(r,a,p)_i}$$

We should at this stage measure the size of $\zeta_{\delta,K,(r,a,p)_i}$. We need to measure it with exponential weights. The weight in question takes the form $e^{d|s|}$ over Σ_{iR} and of the form $e^{w(s)}$ over $N_{\delta K}$.

Proposition 3.11.3. *The norm of $\zeta_{\delta,K,(r,a,p)_i}$ measured over $u_{\delta,K,(r,a,p)_i}$ with weights as specified above is bounded above by*

$$C\epsilon'(C + e^{dK}) + CRe^{dR}\delta + Ce^{-D'K}.$$

For small enough δ we can make this bound be as small as we please. For convenience we use another letter $\tilde{\epsilon}$, informally thought of as $\epsilon' << \tilde{\epsilon} << \epsilon$, and say given $\tilde{\epsilon} > 0$, we can take $\delta > 0$ small enough so that norm of $\zeta_{\delta,K,(r,a,p)_i}$ is bounded above by $\tilde{\epsilon}$. With some foresight, we will need to make it a bit smaller than $\tilde{\epsilon}$, we can take δ small enough so that the vector field is bounded above by $\tilde{\epsilon}^2$.

Proof. The norm of $\zeta_{\delta,K,(r,a,p)_i}$ measured over Σ_{iK} is upper bounded by $C\epsilon' e^{dK}$ as we discussed earlier.
Next consider the segment of $\zeta_{\delta,K,(r,a,p)_i}$ over v_K for $s_v \in [R-K, N_{\delta K} - R + K]$. Recall we glued at the end points of this interval, hence by previous estimates the norm of $\zeta_{\delta,K,(r,a,p)_i}$ it is bounded above by $C\epsilon' e^{dK}$. Next we address the remaining region. WLOG we focus on $s_1 \in [K, R]$ for u^2. In this region, the distance between v_K and u_δ is bounded above (even when integrated against weights) by $C\epsilon' e^{dK}$. The distance between v_K and the trivial cylinder $\hat{\gamma}_2$ is bounded above by $Re^{dR}\delta$ after integrating with the exponential weights. The distance between $u^2 + (r_2, a_2, p_2)$ and $\hat{\gamma}_2$ in pointwise C^k norm is bounded above by

$$Ce^{-Ds_1}$$

so when we integrate this pointwise difference over $s_1 \in [K, R-K]$ with weight e^{ds_1}, we have the upper bound

$$Ce^{-(D-d)K}$$

and hence our overall bound on $\zeta_{\delta,K,(r,a,p)_i}$ is as claimed in the proposition.

To explain how we make the vector field smaller than $\tilde{\epsilon}$, we first choose K fixed large enough so that $e^{-dK} << \tilde{\epsilon}$, then by choosing ϵ' small enough we can make $C\epsilon'(C + e^{dK})$ much less than $\tilde{\epsilon}$, and we recall as $\delta \to 0$, $\epsilon' \to 0$. Finally $Re^{dR}\delta \to 0$ as $\delta \to 0$ by the definition of R. □

Separating global vector field into components

After we have obtained the preglued map $u_{\delta,K,(r,a,p)_i}$ and vector field $\zeta_{\delta,K,(r,a,p)_i}$, there are a few more steps to complete our construction. They are:

a. Truncate the vector field $\zeta_{\delta,K,(r,a,p)_i}$ into

$$\zeta_{\delta,K,(r,a,p)_i} = \zeta_{1,\delta,K,(r,a,p)_1} + \zeta_{\delta,K,(r,a,p)_i,v} + \zeta_{2,\delta,K,(r,a,p)_2}$$

where

$$\zeta_{i,\delta,K,(r,a,p)_i} \in u^{i*}(TM)$$
$$\zeta_{\delta,K,(r,a,p)_i,v} \in v_K^*(TM).$$

b. Adjust the asymptotic vectors $(r,a,p)_i$ in the pregluing so that the vector fields $\zeta_{i,\delta,K,(r,a,p)_i}, \zeta_{\delta,K,(r,a,p)_i,v}$ live in images of Q_i and H_0 respectively, where the definition of these conditions will be specified below.

c. Show $\zeta_{i,\delta,K,(r,a,p)_i}$ and $\zeta_{\delta,K,(r,a,p)_i,v}$ can be extended to (unique) solutions of the equations $\mathbf{\Theta_i} = 0$ and $\mathbf{\Theta_v} = 0$, subject to our choice of right inverse in the previous step.

In this subsection we address the first two bullet points, and the third bullet point will conclude the surjectivity of gluing, which we will take up after a technical detour.

To address the first bullet point we introduce the cut off functions over v_K. We define

$$\beta_1(s_v) := \beta_{[R/2;N_{\delta K}-K-2R,\infty]}$$

$$\beta_2(s_v) := \beta_{[-\infty,2R-K;R/2]}$$

$$\beta_v := \beta_{[R/2;R-K,N_{\delta K}+K-R;R/2]}.$$

The obvious inference is that if we imagine we constructed $\Sigma_{\delta,K,(r,a,p)_i}$ from a prelguing construction by deforming u^i with (r_i, a_i, p_i) and gluing to it a finite gradient segment, the cut off functions listed above should correspond to the cut off functions we used for our gluing construction. In fact this is exactly the case, β_i ought to be identified with $\beta_{\pm}$. The only difference is a change in notation where our coordinates are shifted by a factor of K.

Then to address the first bullet point, we take some care to specify what we mean in our definition of $\zeta_{*,\delta,K,(r,a,p)_1}, * = 1, 2, v$ in anticipation of our upcoming proof of surjectivity of gluing. In particular we must define $\zeta_{*,\delta,K,(r,a,p)_i}$ so that they satisfy the following properties:

- $$\zeta_{\delta,K,(r,a,p)_i} = \beta_1\zeta_{1,\delta,K,(r,a,p)_1} + \beta_2\zeta_{\delta,K,(r,a,p)_i,v} + \beta_v\zeta_{v,\delta,K,(r,a,p)_i}.$$

- Their norms satisfy

  $$\|\zeta_{i,\delta,K,(r,a,p)_i}\| \leq C\epsilon'(C + e^{dK}) + CRe^{dR}\delta + Ce^{-D'K}$$

  as measured in $W^{2,p,d}(u^i + (r_i, a_i, p_i)^*TM) \oplus T\mathcal{J}_i$ with weighted norm (we ignore the other ends of u^i for now). As well as the fact

  $$\|\zeta_{v,\delta,K,(r,a,p)_i}\| \leq C\epsilon'(C + e^{dK}) + CRe^{dR}\delta + Ce^{-D'K}$$

  as measured in $W^{k,p,w}(v_K^*TM)$.

- The vector fields $\zeta_{*,\delta,K,(r,a,p)_i}$ have support as follows. Using coordinates $(s_v, t_v) \in [0, N_{\delta K}] \times S^1$ over $\zeta_{2,\delta,K,(r,a,p)_2} = 0$ for $s_v > 3R - K$. The vector field $\zeta_{1,\delta,K,(r,a,p)_1} = 0$ for $s_v < N_{\delta K} - K - 3R$, and $\zeta_{v,\delta,K,(r,a,p)_i} = 0$ for $s_v < R - K$ and $N_{\delta K} - K - R < s_v$. (These supports are not too significant as we will find some other way to extend them later).

We observe such extensions are always possible. We note the previous theorem on the norm of the global vector field $\zeta_{\delta,K,(r,a,p)_i}$ implies analogous statements on the individual vector fields $\zeta_{i,\delta,K,(r,a,p)_i}, \zeta_{\delta,K,(r,a,p)_i,v}$. The bullet point about support follows from choice of cut off functions β_*.

To address the second bullet point, we specify what we mean by the "right spaces". Assume u^i are nontrivial with stable domains, recall for u^1 and u^2 (or rather a suitable translate of u^2 in the symplectization direction) with domain Σ_i the space of deformations is given by

$$W^{2,p,d}(u^{i*}TM) \oplus V_i \oplus V_i' \oplus T\mathcal{J}_i.$$

The operators D_i with domain $W^{2,p,d}(u^{i*}TM) \oplus V_i \oplus V_i' \oplus T\mathcal{J}_i$ are defined as the linearization of $\bar{\partial}_J$ along with variations of complex structure of Σ_i. By assumption D_i have index 1 with kernel ∂_a, i.e. global translation in the a direction. Recall to define the equations $\mathbf{\Theta}_i$ we needed to fix a right inverse Q_i to D_i. We choose Q_i as follows, consider the codimension 1 subspace $W_i \subset W^{2,p,d}(u_i^*TM)$ defined by

$$\zeta \in W' \text{ iff } \int_{\hat{\Sigma}_i} \langle \zeta, \partial_a \rangle = 0 \tag{3.21}$$

where $\hat{\Sigma}_i$ is the compact subset of Σ_i with all cylindrical neighborhood around punctures removed. Then D_i restricted to $W_i' \oplus V_i \oplus V_i' \oplus T\mathcal{J}_i$ is an isomorphism with inverse Q_i, and we take this Q_i to be the right inverse used in the contraction mapping principle we use to solve $\mathbf{\Theta_i} = 0$. In the case where the domain is not stable, we note the following convention.

Convention 3.11.4. *For definiteness say the domain of u^1 is either a plane or a cylinder, then the act of placing marked points in the domain presents us a subspace $W' \subset W^{2,p,d}(u_i^*TM)$ so that the restriction of D_1 to $W' \oplus V_1 \oplus V_1' \oplus T\mathcal{J}_1$ is an isomorphism: we simply take W' to be vector fields that preserve the condition that marked points remain on the auxiliary surfaces we chose in Convention 3.4.3. If u^1 has several connected components, some of which are stable, and some of which are unstable prior to adding marked points, then we impose the integral condition 3.21 for vector fields over the domains that are stable without adding marked points, and the marked point condition in Convention 3.4.3 for domains are stable only after adding marked points. This picks out the subspace W'.*

We now explain our definition of H_0. Recall v_K is a segment of a gradient trajectory that has coordinates (s_v, t_v), with the segment of interest being $[0, N_{\delta K}] \times S^1$, with exponential weight $e^{w(s)}$ with peak at $s_v = N_{\delta K}/2$. We recall the functionals, analogous to previous section: $L_*, * = a, s, v : W^{2,p,w}(v_K^*TM) \to \mathbb{R}$ defined by

$$L_* : \zeta \in W^{2,p,w}(v_K^*TM) \longrightarrow \int_0^1 \langle \zeta(s_v = N_{\delta K}/2, t), \partial_* \rangle dt \in \mathbb{R}$$

and we define

$$H_0 := \{\zeta \in W^{2,p,w}(v_K^*TM) | L_*(\zeta) = 0, * = a, x, z\}.$$

We now deform the pair $(r, a, p)_i$ to ensure our vector fields lie in the correct subspace. We first observe we can ensure $\zeta_{1,\delta,K,(r,a,p)_1}$ lives in the image of Q_1 by using the global a translation of u_δ, i.e. when we first started talking about the degeneration of u_δ into u^1 and u^2, we translate u_δ by a so that u_δ always converges to u^1 in C^∞_{loc} via a vector field that lives in image of Q_1. This degree of freedom is afforded to us by the fact that the problem is $\mathbb{R}$ invariant. Hence we next focus on vector fields $\zeta_{2,\delta,K,(r,a,p)_2}$ and $\zeta_{v,\delta,K,(r,a,p)_i}$. We imagine $u_1|_{\Sigma_{1,K}}$ is fixed in place. To make $\zeta_{2,\delta,K,(r,a,p)_2}$ live in the image of Q_2, we change $p_2 \to p_2+\delta p_2$. This changes the pregluing domain: $\Sigma_{\delta,K,(r,a,p)_1,(r,a,p+\delta p)_2}$ as well as the associated map into M, which is given by

$$u_{\delta,K,(r,a,p)_1,(r,a,p+\delta p)_2}$$

The effect of changing p_2 changes the length of the finite gradient trajectory that is glued between u^1 and u^2. Naturally changing the preglued map also deforms the global vector field $\zeta_{\delta,K,(r,a,p)_1,(r,a,p+\delta p)_2}$ and its cutoffs. (Adjusting our cut off functions accordingly). We observe to make $\zeta_{2,\delta,K,(r,a,p+\delta p)_2}$ live in the image of Q_2 we need to lengthen/shorten the glued cylinder by a-length ϵ', this corresponds to a δp_2 of size $\epsilon'\delta$. After this adjustment, if we take δ sufficiently small, the global vector field still has size bounded above by $\tilde{\epsilon}$ (or in the case we will need, $\tilde{\epsilon}^2$), and hence the same is true of its cut offs.

Finally we turn our attention to $\zeta_{v,\delta,K,(r,a,p)_i}$. To do this we need the following lemma:

Lemma 3.11.5.

$$L_*(\zeta_{v,\delta,K,(r,a,p)_i}) \leq C\epsilon'^{2/p}e^{-\lambda(N_{\delta K}-CR)/2}.$$

for $* = a, z, x$.

Proof. Follows directly from exponential decay estimates ☐

Observe this upper bound is extremely small in the following sense. If we consider the vector field $C\epsilon'^{2/p}e^{-\lambda(N_{\delta K}-CR)/2}\partial_*$ where $* = a, z, x$ and measured the size of this vector field over v_k with domain $s_v \in [0, N_{\delta K}]$, with the exponential weight $e^{w(s)}$, we would still get an extraordinarily small number, of size $C\epsilon'^{2/p}e^{-(\lambda-d)N_{\delta K}/2}e^{C\lambda R}$, which goes to zero as $\delta \to 0$. This means we can apply a constant translation of form $C\epsilon'^{2/p}e^{-\lambda(N_{\delta K}-CR)/2}\partial_*$ over $s_v \in [-K, N_{\delta K} + K]$ to the vector field $\zeta_{v,\delta,K,(r,a,p)_i}$ to try to make it land in H_0 while still keeping its overall norm $< \tilde{\epsilon}$. In practice this is done by further adjusting the pair of asymptotic vectors $(r_1, a_1, p_1), (r_2, a_2, p_2 + \delta p_2)$ used in the pregluing, which we take up in the next lemma.

Lemma 3.11.6. *By adjusting the pair of asymptotic vectors $(r_1, a_1, p_1), (r_2, a_2, p_2 + \delta p_2)$ we can make $\zeta_{v,\delta,K,(r,a,p)_i} \in H_0$ while keeping its norm $\leq \tilde{\epsilon}^2$. We can also maintain the vector fields $\zeta_{i,\delta,K,(r,a,p)_i}$ are still within the image of Q_i.*

Proof. We examine L_* for $* = r, a, z$ one by one, as the different cases are relatively independent of each other.

Let us consider L_x. The idea is to change both p_1 and p_2 in the same direction, which we denote by $p_i + \Delta p$, and preglue to a different gradient trajectory v'_K, but the trouble is as we change p_i to $p_i + \Delta p$, the new gradient trajectory v'_K connecting $p_1 + \Delta p$ to $p_2 + \Delta p$ travels a different amount of a distance as s'_v (the variable for v'_K) ranges from $s_v = 0$ to $s'_v = N_{\delta K}$, hence there must be a corresponding deformation in the pair of asymptotic vectors a_1 and

a_2 to make the curves still match up and glue. Further we must also choose the deformation of a_1 and a_2 so that $\zeta_{i,\delta,K,(r,a,p)_i} \in ImQ_i$. Said differently we deform a_1 and a_2 so that there is no induced global translation of u^1 or u^2 that enters the pregluing. This is always possible. The exact expressions for these quantities are not so important, the important information is their sizes. The size of Δ_p is $L_x(\zeta_{v,\delta,K,(r,a,p)_i}) \leq C\epsilon'^{2/p}e^{-\lambda(N_{\delta K}-CR)/2}$, to make the new $\zeta'_{v,\delta,K,(r,a,p)_i}$ evaluate to 0 under L_x, the corresponding change to a_1, a_2 is also of size $CC\epsilon'^{2/p}e^{-\lambda(N_{\delta K}-CR)/2}$, which we absorb into our notation $(r,a,p)_i$. It is also apparent after this deformation all vector fields are still small. Next we adjust both a_1 and a_2 by $L_a(\zeta_{v,\delta,K,(r,a,p)_i})$ to make $L_a(\zeta_{v,\delta,K,(r,a,p)_i}) = 0$. We adjust a_1 and a_2 by the same amount in the same direction so as to maintain $\zeta_{i,\delta,K,(r,a,p)_i} \in ImQ_i$. It is clear this will land $\zeta_{v,\delta,K,(r,a,p)_i}$ in H_0 and keep the norm of $\zeta_{v,\delta,K,(r,a,p)_i}$ small.

Finally we consider ∂_z. We shift r_1 by size $-L_r(\zeta_{v,\delta,K,(r,a,p)_i})$, and twist the segment of gradient trajectory v_K along with it. But we do not change r_2, hence there is an new identification $t_v + L_r(\zeta_{v,\delta,K,(r,a,p)_i}) \sim t_2$ near the u^2 end. The result is a new $\zeta_{v,\delta,K,(r,a,p)_i}$, denoted by the same symbol by abuse of notation, so that $L_r(\zeta_{v,\delta,K,(r,a,p)_i}) = 0$. We also observe by the previous discussion the norm of $\zeta_{v,\delta,K,(r,a,p)_i}$ changed at most by $C\epsilon'^{2/p}e^{-(\lambda-d)N_{\delta K}/2}e^{C\lambda R}$.

It is also clear that this process will keep $\zeta_{i,\delta,K,(r,a,p)_i}$ in the image of Q_i because the regions in which we are performing these deformations are disjoint. □

To summarize:

Proposition 3.11.7. *We can choose suitable asymptotic vectors $(r,a,p)_1$, $(r,a,p)_2$, from which to construct a preglued domain $\Sigma_{\delta,K,(r,a,p)_i}$ that decomposes as*

$$\Sigma_{\delta,K,(r,a,p)_i} := (u^1, \delta j_i) + (r_1, a_1, p_1)|_{\Sigma_{1R}} \cup [R-K, N_{\delta K}+K-R] \times S^1 \cup (u^2, \delta j_2) + (r_2, a_2, p_2)|_{\Sigma_{2R}}$$

where δj_i represents variation of complex structure on Σ_{iR}, and we let v_K denote the segment of gradient trajectory whose domain is $[R-K, N_{\delta K}+K-R] \times S^1$. There is a preglueing map $u_{\delta,K,(r,a,p)_i} : \Sigma_{\delta,K,(r,a,p)_i} \to M$ that agrees with our prescription for constructing pregluing maps in Section 3.9, so that there exists a vector field $\zeta_{\delta,K,(r,a,p)_i}$ so that

$$u_\delta = u_{\delta,K,(r,a,p)_i} + \zeta_{\delta,K,(r,a,p)_i}.$$

If we split $\zeta_{\delta,K,(r,a,p)_i}$ into components that live over u^1, u^2, v_K using cut off functions β_* as we did in our gluing sections. The resulting vector fields $\zeta_{1,\delta,K,(r,a,p)_1}, \zeta_{\delta,K,(r,a,p)_i,v}, \zeta_{2,\delta,K,(r,a,p)_1}$ satisfy

$$\zeta_{i,\delta,K,(r,a,p)_1} \in \operatorname{Im} Q_i$$

$$\zeta_{\delta,K,(r,a,p)_i,v} \in H_0$$

and they all have norm $< \tilde{\epsilon}$ when measured with exponential weights. For $\zeta_{i,\delta,K,(r,a,p)_1}$ this means $W^{2,p,d}(u_i^*TM)$ and $\zeta_{\delta,K,(r,a,p)_i,v} \in W^{2,p,w(s)}(v_K^*TM)$. We remark here we are bounding the size of our vector fields by $\tilde{\epsilon}$, but it practice we can make them as small as we please, by $\tilde{\epsilon}^2$, for instance.

Now we are in the position to extend $\zeta_{v,\delta,K,(r,a,p)_i}$ and $\zeta_{i,\delta,K,(r,a,p)_i}$ to solutions of Θ_i, Θ_v, but before that we need to take a detour on linear operators.

A detour on linear operators

In this detour of a subsection we prove several key facts about linear operators to be used later. Naturally, very similar lemmas appear in Section 3 of [41] since we are using their strategy for surjectivity of gluing.

We shall first consider the case for semi-infinite trajectories, then we will do the case for finite gradient trajectories.

We shall first work out the case for $p = 2$, then deduce the necessary results $p > 2$ from Morrey's embedding theorem. For this section we shall work with Sobolev regularity $k > 3$, this will not make a difference to us since elliptic regularity will afford us all the regularity we need.

Let

$$v : [0, \infty) \times S^1 \longrightarrow M$$

be a semi-infinite gradient trajectory, equipped with linearized operator

$$D_\delta = \partial_s - (A(s,t) + \delta A) \tag{3.22}$$

where $A(s,t) = -(J_0 \frac{d}{dt} + S)$ corresponds to the linearized operator of the Morse-Bott contact form, and δA is a operator of the form $\delta(M\frac{d}{dt} + N)$ is the correction due to having used the J_δ almost complex structure.

We equip it with the weighted Sobolev space $W^{k,2,w(s)}$ where

$$w(s) = d(s + R).$$

We conjugate this over to $W^{k,2}$ at which point it becomes

$$D'_\delta = \partial_s - (A + \delta A) - d. \tag{3.23}$$

Let's first consider the restriction of $A + d$ to $s = 0$, which we shall denote by A_0. By the spectral theorem there exists an orthonormal basis of $L^2(S^1)$ given by eigenfunctions of A_0, which we write as $\{e_n\}_{n\in\mathbb{Z}}$ with eigenvalue λ_n. By assumption 0 is not an eigenvalue of A_0, and by convention we say $\lambda_n > 0$ for $n > 0$ and vice versa.

Theorem 3.11.8. *There is a continuous trace operator $T : W^{k,2}([0,\infty)\times S^1) \to W^{k-1/2,2}(S^1)$ given as follows, if $f \in W^{k,2}([0,\infty) \times S^1)$:*

$$(Tf)(t) = f(0,t)$$

The norm in $W^{k-1/2,2}(S^1)$is given as follows, every $f(t) \in W^{k-1/2,2}(S^1)$ has a Fourier expansion

$$f(t) = \sum_n a_n e_n(t)$$

then the norm is equivalent to following expression:

$$\|f(t)\|^2 := \sum_n |a_n|^2 \lambda_n^{2k-1}.$$

Proof. This is a standard theorem in analysis, for a description of this see for instance proof of Lemma 3.7 in [41]. □

Then we come to the first main theorem of this detour.

Theorem 3.11.9. *Let $W_-^{k-1/2,2}(S^1)$ denote the subspace of $W^{k-1/2,2}(S^1)$ such that $a_n = 0$ for all $n > 0$, let $\Pi_- : W^{k-1/2,2}(S^1) \to W_-^{k-1/2,2}(S^1)$ denote the projection. Then the map $(\Pi_-, \partial_s - A_0) : W^{k,2}([0,\infty) \times S^1) \to W_-^{k-1/2,2}(S^1) \times W^{k-1,2}([0,\infty) \times S^1)$ taking*

$$f(s,t) \longrightarrow (\Pi_- f(0,t), (\partial_s - A_0) f(s,t))$$

is an isomorphism.

Proof. We can solve this equation explicitly. Given a pair $(g,h) \in W^{k-1/2,2}(S^1) \times W^{k-1,2}([0,\infty) \times S^1)$, we can write

$$g = \sum_{n<0} c_n e_n(t)$$

$$h = \sum_n b_n(s) e_n(t)$$

where

$$\|g\|^2 = \sum_{n<0} |c_n|^2 |\lambda_n|^{2k-1}$$

$$\|h\|^2 = \sum_n \int_0^\infty (|b_n(s)|^2 |\lambda_n|^{2k-2} + |b_n'(s)|^2 |\lambda_n|^{2k-4} + .. + |b_n^{(k-1)}(s)|^2) ds.$$

The usual Sobolev norms are equivalent to the expressions we've written above. Comparing term by term we see that a_n satisfies the following ODE:

$$a_{ns} - \lambda_n a_n = b_n(s)$$

with boundary condition $a_n(0) = c_n$ for all $n < 0$. They have solutions

$$a_n = e^{\lambda_n s} \int_0^s b_n(s') e^{-\lambda_n (s')} ds' + c_n e^{\lambda_n s}$$

where the c_n term only appears for $n < 0$. We need to verify several things:

a. The terms $e_n(t) e^{\lambda_n s} \int_0^s b_n(s') e^{-\lambda_n (s')}$ and $c_n e_n(t) e^{\lambda_n s}$ are in $W^{k,2}([0,\infty) \times S^1)$.

b.

$$\begin{aligned} &\int_0^\infty (|a_n|^2 |\lambda_n|^{2k} + |a_n'(s)|^2 \lambda_n|^{2k-2} + .. + |a_n^{(k)}(s)|^2)| \\ \leq& C(|c_n|^2 |\lambda_n|^{2k-1}) + C \int_0^\infty (|b_n(s)|^2 |\lambda_n|^{2k-2} + |b_n'(s)|^2 |\lambda_n|^{2k-4} + .. + |b_n^{(k-1)}(s)|^2) \end{aligned}$$

The first item says our constructed solution f lives in our Sobolev space, the second item says its norm is upper bounded by our input.

First consider $c_n e_n(t)e^{\lambda_n s}$, its norm in $W^{k,2}(S^1\times[0,\infty))$ is given by

$$\int_0^\infty |c_n|^2|\lambda_n|^{2k}e^{2\lambda_n s}ds \leq C|c_n|^2|\lambda_n|^{2k-1}$$

and that this is finite after we sum over n follows from our assumptions on g. Similarly consider $d_n := e^{\lambda_n s}\int_0^s b_n(s')e^{-\lambda_n(s')}ds'$, its norm as measured in $W^{k,2}([0,\infty)\times S^1)$ is given by

$$\int_0^\infty (|d_n(s)|^2|\lambda_n|^{2k} + |d_n'(s)|^2|\lambda_n|^{2k-4} + .. + |d_n^{(k)}(s)|^2).$$

We have

$$\begin{aligned}|d_n^{(l)}| \leq C(l)e^{\lambda_n s}\Big\{&|\lambda_n|^l\int_0^s |b_n(s')|e^{-\lambda_n(s')}ds' + |\lambda_n|^{l-1}|b_n(s)|e^{-\lambda_n s}\\ &+|\lambda_n|^{l-2}|b_n^{(1)}(s)|e^{-\lambda_n s} + \ldots + |b_n^{(l-1)}|e^{-\lambda_n s}\Big\}\end{aligned}$$

and we need to take derivatives up to $l = 0,...,k$. We remind ourselves we need to place upper bounds on terms of the form $|\lambda_n|^{2k-2l}|d_n^{(l)}|^2$, hence from above it suffices to bound terms of the form

$$\int_0^\infty |\lambda_n|^{2(k-j-1)}|b_n^{(j)}|^2ds, j = 0,...,k-1$$

$$\int_0^\infty e^{2\lambda_n s}|\lambda_n|^{2k}\left(\int_0^s b_n(s')e^{-\lambda_n(s')}ds'\right)^2 ds.$$

The first term is bounded by the norm of g. The second term really is the L^2 norm of a_n multiplied by λ_n^{2k}. We use a technique (probably much more well known) we found in [17], Chapter 3. We first observe by Sobolev embedding our functions are at least C^1, so we can use the fundamental theorem of calculus. We then consider the defining equation for a_n

$$\frac{d}{ds}a_n - \lambda_n a_n = b_n$$

from which we get

$$\left(\frac{d}{ds}a_n\right)^2 + (\lambda_n a_n)^2 = b_n^2 + \lambda_n\frac{d}{ds}(a_n)^2.$$

Integrate both sides from $[0,\infty)$ to get

$$\int_0^\infty \left(\frac{d}{ds}a_n\right)^2 + (\lambda_n a_n)^2 ds = \int_0^\infty b_n^2 ds - \lambda_n|c_n|^2$$

where we used continuity to apply fundamental theorem of calculus. We also used the fact for any fixed b_n, we have $\lim_{s\to\infty} e^{2\lambda_n s}\int_0^s b_n^2(s')e^{-2\lambda_n(s')}ds' \to 0$. Hence we get

$$\int_0^\infty |a_n|^2 ds \le \frac{1}{|\lambda_n|^2}\int_0^\infty |b_n|^2 ds + |\lambda_n|^{-1}|c_n|^2.$$

From which we deduce the second term is also bounded by norm of g and h. Combining the above computations we see that our solution f is indeed in $W^{k,2}([0,\infty)\times S^1)$, and the inequality

$$\|f\| \le C(\|g\| + \|h\|)$$

holds, from which we conclude the theorem. □

Corollary 3.11.10. *Let $D_{\delta'0}$ denote the operator D_δ restricted at $s=0$, i.e. $D_{\delta'0} = \partial_s - A(0,t) - \delta A(0,t) - d$, then for small enough $\delta > 0$, the map $(\Pi_-, D_{\delta'0}) : W^{k,2}([0,\infty)\times S^1) \to W_-^{k-1/2,2}(S^1)\times W^{k-1,2}([0,\infty)\times S^1)$ is an isomorphism with inverse Q_0 whose operator norm is uniformly bounded as $\delta \to 0$.*

Using the above results we come to the theorem we will really need later on:

Theorem 3.11.11. *For small enough $\delta > 0$, the operator $(\Pi_-, D'_\delta) : W^{k,2}([0,\infty)\times S^1) \to W_-^{k-1/2,2}(S^1)\times W^{k-1,2}([0,\infty)\times S^1)$ is an isomorphism whose inverse Q has operator norm uniformly bounded with respect to $\delta \to 0$.*

Proof. The proof is reminiscent of our original proof that D_δ (which we earlier denoted by D_{J_δ}) has uniformly bounded inverse over the entire gradient trajectory, i.e. we approximate it by a sequence of operators over trivial cylinders.

Let N be a large integer, choose x_i for $i = 0, 1, ..N$ so that x_0 is the x coordinate on the Morse-Bott torus of $v(0,t)$, we have $|x_i - x_{i-1}| \le 1/N$, and x_N is distance $< 1/N$ away from the critical point on the Morse-Bott torus corresponding to $v(\infty,t)$. We let $D'_i : W^{k,2}(\mathbb{R}\times S^1) \to W^{k-1,2}(\mathbb{R}\times S^1)$ denote the linearization of $\bar\partial_J$ at the trivial cylinder located at x_i on the Morse-Bott torus, and conjugated by exponential weights to remove exponential weight. In formulas we have

$$D'_i = \partial_s + J\partial_t + S(x_i,t) - d.$$

Uniformly in N and $\delta > 0$ and independently of i, the D'_i are isomorphisms with uniformly bounded inverses Q'_i. Then similar to previous section we construct the glued operator $\#D'_i$ which satisfies

$$\|D'_\delta - \#D_i\| \le C(1/N + \delta).$$

As before we construct an approximate inverse to $\#D'_i$, which we call Q'_R. If we let W_- abbreviate $W^{k-1,2}([0,\infty)\times S^1)\oplus W_-^{k-1/2,2}(S^1)]$, we have via the following diagram:

$$\begin{tikzcd}
W_- \oplus [W^{k-1,2}((-\infty,\infty)\times S^1)]_1 \oplus \dots \arrow[r, "Q'_R"] \arrow[d, "s_R"] & W^{k,2}([0,\infty)\times S^1) \\
W_- \oplus W^{k-1,2}((-\infty,\infty)\times S^1)_1 \oplus .. \arrow[r, "Q_0\oplus Q_1.."] & W^{k,2}([0,\infty)\times S^1))_0 \oplus W^{k,2}((-\infty,\infty)\times S^1)_1.. \arrow[u, "g_R"]
\end{tikzcd}$$

where we clarify

$$s_R|_{W_-^{k-1/2,2}(S^1)} = \mathrm{Id}\,.$$

The subscripts under $W^{k,2}((-\infty,\infty)\times S^1))_i$ denote the copies of Sobolev spaces in the direct sum. And the splitting map s_R and the gluing map g_R are defined exactly the same way we did in section 3.9. We observe as before this Q'_R is uniformly bounded as $\delta \to 0$. Let's verify that this constructs an approximate inverse to $\#D'_i$. We first observe away from the gluing region

$$\#D'_i Q'_R \eta = \eta$$

and near the gluing region as before we have

$$\|\#_N D'_i Q'_R \eta - \eta\| \leq C/N\|\eta\|.$$

Hence we can construct a right inverse of $\#D_i$ with uniformly bounded norm. Next since D'_δ is a uniformly bounded small perturbation of $\#_N D'_i$, it also has a uniformly bounded right inverse.

To see that this operator is injective, since we don't have index calculations (versions of index theorems probably exist but we cannot find an easy reference) we take a more direct approach, in part inspired by the appendix of [12]. Suppose $\zeta_\delta \in Ker(\Pi_-, D'_\delta)$ is of norm 1, consider $s = R$, for definiteness we first assume for all δ the norm of ζ_δ restricted to $0 < s < R$ is $\geq 1/2$. Let $\beta_R := \beta[-\infty, 2R; R]$, and consider $\beta_R\zeta_\delta$. Then we can consider it to lie in the domain of $(\Pi_-, \partial_s - A_0) : W^{k,2}([0,\infty)\times S^1) \to W_-^{k-1/2,2}(S^1)\times W^{k-1,2}([0,\infty)\times S^1)$. To estimate its image under $(\Pi_-, \partial_s - A_0)$, first consider

$$\|D_\delta \beta_R \zeta_\delta\| = \|\beta'_R \zeta_\delta\| \leq C/R.$$

Observing that over $s < 2R$ we have $\|\partial_s - A_0 - D'_\delta\| \leq C\delta$, we have

$$(\Pi_-, \partial_s - A_0)(\beta_R\zeta_\delta) = (0, (\partial_s - A_0)\beta_R\zeta_\delta)$$

where $\|(\partial_s - A_0)\beta_R\zeta_\delta\| \leq C/R$, but then the element

$$\beta_R\zeta_\delta - (\Pi_-, \partial_s - A_0)^{-1}((\Pi_-, \partial_s - A_0)(\beta_R\zeta_\delta)) \in W^{k,2}([0,\infty)\times S^1)$$

has norm $> 1/3$, but lies in the kernel of $(\Pi_-, \partial_s + A_0)$, which is a contradiction. Similarly, if the norm of ζ_δ when restricted to $s > R$ is $\geq 1/2$ for all $\delta > 0$, then we use a similar cut off function $\hat{\beta}_R := \beta_{[R/2;R/2,\infty]}$ to view $\hat{\beta}_R\zeta_\delta$ as element of $(D'_\delta, W^{k,2}(v^*TM))$ and use the same process to produce a nonzero kernel of D'_δ, which cannot exist since D'_δ is an isomorphism. □

We now state the finite interval analogue of the above theorems for later use.

Theorem 3.11.12. *Let v be a gradient trajectory. Let D'_δ be the linearization of $\bar{\partial}_{J_\delta}$ over v with exponential weight removed via conjugation as above. We consider its restriction to $(s,t) \in [0, CR] \times S^1$, and the Sobolev space $W^{k,2}([0,CR] \times S^1, \mathbb{R}^4)$. Consider the two projections $\Pi_\pm$, where they project to the positive/ negative eigenvalues of $A_0 := -A(0,t) - d$. Then the map*

$$\begin{aligned}&(\Pi_-, \Pi_+, D'_\delta) : W^{k,2}([0,CR] \times S^1, \mathbb{R}^4)\\&\longrightarrow W^{k-1/2,2}_-(S^1) \times W^{k-1/2,2}_+(S^1) \times W^{k-1,2}([0,CR] \times S^1, \mathbb{R}^4)\end{aligned}$$

defined by

$$f(s,t) \longrightarrow (\Pi_- f(0,t), \Pi_+ f(CR,t), D'_\delta f)$$

is an isomorphism whose inverse has uniformly bounded norm as $\delta \to 0$.

Proof. As before we first show the map $(\Pi_-, \Pi_+, \partial_s - A_0) : W^{k,2}([0,CR] \times S^1, \mathbb{R}^4) \to W^{k-1/2,2}(S^1) \times W^{k-1/2,2}(S^1) \times W^{k-1,2}([0,CR] \times S^1, \mathbb{R}^4)$ is an isomorphism with uniformly bounded inverse Q_0. This is essentially the same proof as before, i.e. if $(\Pi_-, \Pi_+, \partial_s - A_0) f = (g_-, g_+, h)$ with $f = \sum a_n e_n$, $g_\pm = \sum c_{n\pm} e_n$ and $h = \sum b_n e_n$ then we still have the formulas

$$a_n = e^{\lambda_n s} \int_0^s b_n(s') e^{-\lambda_n(s')} ds' + c_{n-} e^{\lambda_n s}$$

for $n < 0$ and

$$a_n = e^{\lambda_n(s-CR)} \int_{CR}^s b_n(s') e^{-\lambda_n(s')} ds' + c_{n+} e^{\lambda_n(s-CR)}$$

for $n > 0$. This already implies injectivity. The same proof shows Q_0 exists and is uniformly bounded as $\delta \to 0$. To elaborate a bit further, we still need to estimate sizes of 3 kinds of terms. For definiteness we focus on the case $n < 0$. The terms we need to consider are of forms

a. $\int_0^{CR} |c_n|^2 |\lambda_n|^{2k} e^{2\lambda_n s} ds$

b. $\int_0^{CR} |\lambda_n|^{2(k-j-1)} |b_n^{(j)}|^2 ds, j = 0, ..., k-1$

c. $\int_0^{CR} e^{2\lambda_n s} |\lambda_n|^{2k} (\int_0^s b_n(s') e^{-\lambda_n(s')} ds')^2 ds.$

The first two terms work exactly the same way as before with CR replacing ∞. The third term requires a bit more care in that when we tried to estimate the L^2 norm of a_n, the domain of integration is different giving us an extra term via integration by parts. So instead we have

$$\lambda_n^2 \int_0^{CR} |a_n|^2 ds \leq \int_0^{CR} |b_n|^2 ds + \lambda_n \{(a_n(CR))^2 - (a_n(0))^2\}.$$

The additional term we need to estimate is $|\lambda_n| a_n^2(CR)$. This is upper bounded by

$$|\lambda_n| \cdot |c_{n-}|^2 e^{2\lambda CR} + |\lambda_n| e^{2\lambda_n CR} \left(\int_0^{CR} b_n(s') e^{-\lambda_n s'} ds' \right)^2.$$

The first term above, after multiplying by $|\lambda_n|^{2k-2}$, is upper bounded by the norm of g_- with the correct weight of $|\lambda_n|$. To examine the second term note it is bounded above by

$$|\lambda_n| e^{2\lambda_n CR} \int_0^{CR} b_n^2(s')ds' \int_0^{CR} e^{-2\lambda_n s'} ds' \leq C \int_0^{CR} b_n^2(s')ds'$$

by Cauchy-Schwartz, and this has the right weight of $|\lambda_n|$ so that when we multiply by $|\lambda_n|^{2k-2}$ it is upper bounded by the norm of g. This concludes the discussion of the third bullet point. Putting all of these together as in the semi-infinite case we see that the inverse is well defined, and its norm is uniformly bounded above as $\delta \to 0$.

To conclude $(\Pi_-, \Pi_+, D'_\delta)$ has uniformly bounded inverse we need to be slightly careful, since as $\delta \to 0$ the domain changes. Since $\delta R \to 0$ the actual operator $(\Pi_-, \Pi_+, D'_\delta)$ is a size $\leq R\delta$ perturbation of $(\Pi_-, \Pi_+, \partial_s - A_0)$, then by the above we can construct a right inverse with uniform bound Q for $(\Pi_-, \Pi_+, D'_\delta)$ and this implies surjectivity. To show injectivity we proceed similarly as before, we assume ζ_δ has norm 1 and lives in the kernel of $(\Pi_-, \Pi_+, D'_\delta)$, then $\|(\Pi_-, \Pi_+, \partial_s - A_0)\zeta_\delta\| \leq CR\delta\|\zeta\|$, then the element

$$\zeta_\delta - Q_0(\Pi_-, \Pi_+, \partial_s - A_0)\zeta_\delta$$

is an element of norm $> 1/2$ in the kernel of $(\Pi_-, \Pi_+, \partial_s - A_0)$, contradiction. □

Surjectivity of gluing

In this subsection we finally prove surjectivity of gluing in our simplified setting. The idea is that we shall extend our vector fields $\zeta_{*,\delta,K,(r,a,p)_i}, i = 1, 2, v$ so that they satisfy the set of equations $\Theta_i = 0, \Theta_v = 0$, subject to our choice of right inverses, which we constructed in the pregluing section. Then this shows our holomorphic curve u_δ can be realized as a solution of $\Theta_i = 0, \Theta_v = 0$. Since we proved such solution is unique, this shows gluing is surjective. We will first focus on extending the vector fields $\zeta_{*,\delta,K,(r,a,p)_i}$ over the intermediate finite gradient trajectory. The extension to semi-infinte trajectories is similar but independent of this process so will be treated separately.

We remark additionally since there are exponential weights in place, we clarify our notation: when we write a vector field ζ_* without $'$, we think of it as living in some exponentially weighted Sobolev space, when we write ζ'_* we think of it as living in an unweighted space where the weight has been removed by multiplication with the exponential weight. When we write $W^{k,2,d}$ we will always mean the exponential weight e^{ds}; we will write $W^{k,2,w}$ if a more complicated weight is used.

Finally we remark that we will work with Sobolev exponent $p = 2$, then extend our result for $p > 2$, since all of our linear theory was only worked out for $p = 2$. We first observe by virtue of u_δ being J_δ-holomorphic, the vector fields $\zeta_{*,\delta,K,(r,a,p)_i}$ already satisfy $\Theta_* = 0$ at most places. We focus on what happens around u^2 and where u^2 is glued to the finite gradient cylinder simply for ease of notation. Entirely analogous statements hold for u^1.

Proposition 3.11.13. *For $(s_v, t_v) \in [3R-K, N_{\delta K}-3R+K] \times S^1$, the vector field $\zeta_{v,\delta,K,(r,a,p)_i}$ satisfies $\Theta_v = 0$.*
For $(s_v, t_v) \in [0, R-K] \times S^1 \cup \Sigma_{2R}$, the vector field $\zeta_{2,\delta,K,(r,a,p)_2}$ satisfies $\Theta_2 = 0$. An entirely analogous statement is true near u_1.

Because of our choice cut off functions, the global vector field $\zeta_{\delta,K,(r,a,p)_i}$ agrees with $\zeta_{2,\delta,K,(r,a,p)_2}$ at $s_v = R-K$ and $\zeta_{v,\delta,K,(r,a,p)_i}$ at $s_v = 3R-K$. Here we use (s_v, t_v) coordinates, and see next proposition for using (s_2, t_2) coordinates.

Proposition 3.11.14. *There exists a unique vector field ξ of norm less than $\tilde{\epsilon}$ over $W^{k,2,w}([R-K, 3R-K] \times S^1, \mathbb{R}^4)$ where $w = d(s+K)$ that satisfies*

$$\Pi_{-}\xi(R-K, t_v) = \Pi_{-}\zeta_{\delta,K,(r,a,p)_i}(R-K, t_v)$$

$$\Pi_{+}\xi(3R-K) = \Pi_{+}\zeta_{\delta,K,(r,a,p)_i}(3R-K, t_v)$$

and $v_K + \xi$ is J_δ-holomorphic. An entirely analogous statement holds near the ends of u^1.

Proof. The J_δ-holomorphicity condition amounts to ξ solving a equation of the form

$$D_\delta \xi + \mathcal{F}(\xi) = 0$$

where $\mathcal{F}$ is an expression bounded above in C^k by $C|\xi|^2 + |\xi||\partial_t \xi|$. We next remove the exponential weights to get an equation

$$D'_\delta \xi' + \mathcal{F}'(\xi') = 0$$

where we also have $\mathcal{F}' \leq C|\xi'|^2 + |\xi'||\partial_t \xi'|$. Then finding a solution to this equation with prescribed boundary conditions amounts to finding a fixed point of the map

$$I : W^{k,2}([R-K, 3R-K] \times S^1, \mathbb{R}^4) \longrightarrow W^{k,2}([R-K, 3R-K] \times S^1, \mathbb{R}^4)$$

defined by

$$I(\xi') = Q(\Pi_{-}\zeta'_{\delta,K,(r,a,p)_i}(R-K, t_v), \Pi_{+}\zeta'_{\delta,K,(r,a,p)_i}(3R-K, t_v), -\mathcal{F}'(\xi'))$$

where Q is inverse of the operator $(\Pi_{-}, \Pi_{+}, D'_\delta)$, and $\zeta'_{\delta,K,(r,a,p)_i}(R-K, t_v)$ is $\zeta_{\delta,K,(r,a,p)_i}(R-K, t_v)$ multiplied with the inverse of the exponential weight. That Q exists, is an isomorphism with uniformly bounded norm follows from previous section on linear analysis. That I is a contraction mapping principle follows the fact $\mathcal{F}$ is quadratic, the images of projection maps $\Pi_{\pm}$ are independent of the input ξ', as well as the fact that the norm of Q is uniformly bounded as $\delta \to 0$. The fact that I sends $\tilde{\epsilon}$ ball to itself is inherited in the fact $\mathcal{F}'$ is quadratic. We also need to recall from previous estimates that the $W^{2,p,w}$ norms of (hence its C^0 norm) $\zeta_{\delta,K,(r,a,p)_i}$ can be made arbitrarily small as we take $\delta \to 0$ and the norms of $\Pi_{\pm}$ and Q are uniformly bounded, which ensure the image of the contraction map I land easily in the $\tilde{\epsilon}$ ball in the codomain (in our previous propositions we used $\tilde{\epsilon}^2$ to bound the norms, and this is where it comes in). The theorem now follows from contraction mapping principle. $\square$

We next extend $\zeta_{2,\delta,K,(r,a,p)_2}$ and $\zeta_{v,\delta,K,(r,a,p)_i}$ to solutions of $\Theta_2 = 0$ and $\Theta_v = 0$ for $s_v < R$ and $s_2 > R$. We recall there is a slight subtlety in that near the pregluing at u_2 there is a twist in the domain, i.e. an identification $t_v = t_2 + (r_1 - r_2)$, and the vector fields over u_2 have coordinates t_2. We will be careful to make this identification, though we remark it doesn't cause any difficulties.

Proposition 3.11.15. *There are vector fields $\hat{\zeta}_{2,\delta,K,(r,a,p)_2}$, $\hat{\zeta}_{v,\delta,K,(r,a,p)_i}$ defined over $W^{k,2,d}([R,\infty) \times S^1, \mathbb{R}^4)$ and $W^{k,2,d}([-\infty, 3R-K) \times S^1, \mathbb{R}^4)$ respectively, both of norm $< \tilde{\epsilon}$ so that*

$$\Theta_v(\hat{\zeta}_{2,\delta,K,(r,a,p)_2}, \hat{\zeta}_{v,\delta,K,(r,a,p)_i}) = 0$$
$$\Theta_2(\hat{\zeta}_{2,\delta,K,(r,a,p)_2}, \hat{\zeta}_{v,\delta,K,(r,a,p)_i}) = 0$$

where the exponential weight looks like e^{ds} over $W^{k,2,d}([R,\infty) \times S^1, \mathbb{R}^4)$ and $e^{d(s+K)}$ over $W^{k,2,d}([-\infty, 3R-K) \times S^1, \mathbb{R}^4)$. Further, we have the boundary conditions that

$$\Pi_-(\hat{\zeta}_{2,\delta,K,(r,a,p)_2}(s_2 = R, t_2)) = \Pi_-(\zeta_{\delta,K,(r,a,p)_i}(s_v = R-K, t_v - (r_1 - r_2))$$

$$\Pi_+(\hat{\zeta}_{v,\delta,K,(r,a,p)_i})(s_v = 3R-K, t_v) = \Pi_+(\zeta_{\delta,K,(r,a,p)_i}(s_v = 3R-K, t_v).$$

Proof. We immediately switch to primed coordinates by removing the weight. In these primed coordinates the equations look like

$$\Theta_v' = D_\delta' + \mathcal{F}_v'$$

where $\mathcal{F}_v'$ can be upper bounded by quadratic expressions of $\hat{\zeta}'_{2,\delta,K,(r,a,p)_2}$, $\hat{\zeta}'_{v,\delta,K,(r,a,p)_i}$, and their t derivatives, as in Remark 3.7.12. Likewise for

$$\Theta_2' = D_\delta' + \mathcal{F}_2' + \mathcal{E}'.$$

We remark for Θ_2', the operator D_δ' is the linearization of the $\bar{\partial}_{J_\delta}$ along u_2, with exponential weight removed via conjugation. The dependence on (r_2, a_2, p_2) of the linearization appears in the quadratic term $\mathcal{F}_2'$. The term $\mathcal{E}'$ is the corresponding error term which takes the form in pregluing section (it is slightly different since we are using D_δ' as the linear operator, but this is of no consequence). For Θ_v', D_δ' is the linearization of $\bar{\partial}_{J_\delta}$ along v with exponential weights removed.

Then finding an solution to the equation is tantamount to finding a fixed point of the operator

$$\begin{aligned} I :& W^{k,2}([R,\infty) \times S^1, \mathbb{R}^4) \oplus W^{k,2}([-\infty, 3R-K) \times S^1, \mathbb{R}^4) \\ &\longrightarrow W^{k,2}([R,\infty) \times S^1, \mathbb{R}^4) \oplus W^{k,2}([-\infty, 3R-K) \times S^1, \mathbb{R}^4) \end{aligned}$$

defined by:

$$\begin{aligned} I(\hat{\zeta}'_{2,\delta,K,(r,a,p)_2}, \hat{\zeta}'_{v,\delta,K,(r,a,p)_i}) =& \{Q_2(\Pi_-(\zeta'_{\delta,K,(r,a,p)_i}(s_v = R, t_v - (r_1 - r_2)), -\mathcal{F}_2' - \mathcal{E}'), \\ & Q_v(\Pi_+(\zeta'_{\delta,K,(r,a,p)_i}(s_v = 3R-K, t_v)), -\mathcal{F}_v')\}. \end{aligned}$$

Where Q_v is the inverse to the pair $(D'_\delta, \Pi_+) : W^{k,2}([-\infty, 3R-K) \times S^1, \mathbb{R}^4) \to W^{k-1,2}([-\infty, 3R-K) \times S^1, \mathbb{R}^4) \oplus W_+^{k-1/2,2}(S^1)$ where Π_+ take place at $s_v = 3R-K$. Q_2 is the inverse to $(D'_\delta, \Pi_-) : W^{k,2}([R,\infty)\times S^1, \mathbb{R}^4) \to W^{k-1,2}([R,\infty) \times S^1, \mathbb{R}^4) \oplus W_-^{k-1/2,2}(S^1)$ where Π_- takes place at $s_2 = R$. It follows as in the previous proposition that I is a contraction, from the $\tilde{\epsilon}$ ball to itself, and translating back to the weighted Sobolev spaces proves our theorem. □

It follows from the above proposition and uniqueness that the extensions extend smoothly past $s_2 = R$ and $s_v = 3R-K$, and they recover u_δ:

Proposition 3.11.16. *The concatenation of $\hat{\zeta}_{2,\delta,K,(r,a,p)_2}$ at $s_2 = R$ with $\zeta_{2,\delta,K,(r,a,p)_2}$ at $s_2 = R$ is of class C^k, we denote the resulting vector field by $\zeta_{2,\delta,K,(r,a,p)_2}$, with slight abuse in notation. A similar story holds for $\zeta_{v,\delta,K,(r,a,p)_i}$. The resulting vector fields $\zeta_{2,\delta,K,(r,a,p)_2}$ and $\zeta_{i,\delta,K,(r,a,p)_i}$ are $\leq \epsilon$ in $W^{2,p,d}(u_2^*TM)$ and $W^{2,p,d}(v_k^*TM)$ respectively, and satisfy the pair of equations $\Theta_2 = 0, \Theta_v = 0$.*

Proof. By Proposition 3.11.14 there exists a unique vector field over $s_v \in [R-K, 3R-K]$ satisfying the boundary conditions imposed by $\zeta_{\delta,K,(r,a,p)_i}$, and by whose deformation of v_K makes the resulting surface J_δ-holomorphic. But observe $\beta_2\hat{\zeta}_{2,\delta,K,(r,a,p)_2} + \beta_v\hat{\zeta}_{v,\delta,K,(r,a,p)_i}$ satisfy these conditions as well by virtue of the defining conditions for the pair $\hat{\zeta}_{2,\delta,K,(r,a,p)_2}$, and $\hat{\zeta}_{v,\delta,K,(r,a,p)_i}$: that they are solutions of the pair of equations $\Theta_2 = 0, \Theta_v = 0$. Hence we conclude $\beta_2\hat{\zeta}_{2,\delta,K,(r,a,p)_2} + \beta_v\hat{\zeta}_{v,\delta,K,(r,a,p)_i}$ agrees with $\zeta_{\delta,K,(r,a,p)_i}$ over $s_v \in [R-K, 3R-K]$, and by our choice of cut off functions this implies the concatenation of $\hat{\zeta}_{2,\delta,K,(r,a,p)_2}$ with $\zeta_{2,\delta,K,(r,a,p)_2}$ is smooth, and likewise for $\hat{\zeta}_{v,\delta,K,(r,a,p)_2}$ with $\zeta_{v,\delta,K,(r,a,p)_2}$. That we can take $p > 2$ when we only constructed $\hat{\zeta}_{2,\delta,K,(r,a,p)_2}$ for $p = 2$ follows from the fact that the vector fields and their first order derivatives have C^0 norm < 1, and in which case we have their $W^{2,p}$ norm bounded above by powers of their $W^{2,2}$ norm. Hence in this case the extended parts $\hat{\zeta}_{2,\delta,K,(r,a,p)_2}$ and $\hat{\zeta}_{v,\delta,K,(r,a,p)_2}$ have their $W^{2,p,d}$ norm (over $u^{2*}TM$ and v_K^*TM respectively) bounded above by $\tilde{\epsilon}^{2/p}$, and for small enough $\tilde{\epsilon}$ this lands in the ϵ ball in $W^{2,p,d}(u^{2*}TM)$ and $W^{2,p,d}(v_K^*TM)$ respectively. □

We make one additional remark that the equations $\Theta_2 = 0$ and $\Theta_v = 0$ also depends on the asymptotic vectors $(r,a,p)_i$, but from our constructions these vectors have norm $< \epsilon'$.

Extension of solutions near semi-infinite gradient trajectories

In this subsection we briefly outline how to carry out the above in the case where u^1 is glued to a semi-infinite gradient trajectory. This is simpler than the finite gradient case because we don't need our vector fields to lie in H_0.

Recall our conventions, we assume u_δ degenerates into the cascade $\{u^1, u^2\}$. We focus on what happens near a positive puncture of u^1, which has coordinate $(s'_1, t'_1) \in [0,\infty) \times S^1$.

For large $K > 0$ we can recall the decomposition of the domain of u^1

$$\{(s_1', t_1') \in [K, \infty) \times S^1\} \cup \Sigma_{1K} \cup \text{other punctures of}\, u^1.$$

We will not worry about the other punctures of u^1 and only talk about $\Sigma_{1K} \cup [K, \infty) \times S^1$. Similarly we can break down the domain of u_δ into

$$\Sigma_{\delta K} \cup [0, \infty) \times S^1 \cup \text{other parts of }\; u_\delta.$$

As above we will only care about $\Sigma_{\delta K} \cup [K, \infty) \times S^1$ and neglect other parts of u_δ.

Following our previous conventions, we take $\epsilon > 0$ to be the ϵ that controls the size of ϵ balls we use in the contraction mapping principle and it is fixed for any choice of $\delta > 0$. Let the parameter ϵ' depend on K, δ and go to zero as $\delta > 0$. We further introduce $\tilde{\epsilon}$ which we consider to be of the form $0 < \epsilon' << \tilde{\epsilon} << \epsilon$ to help bound various norms of vector fields as in the previous discussion.

The convergence to cascade implies for given K, we can choose small enough $\delta > 0$ so that there exists a vector field $\zeta_{1\delta}$ and variation of complex structure of u_1 so that

$$u_\delta|_{\Sigma_{\delta K}} = \exp_{u^1, \delta j_1}(\zeta_{1\delta})$$

and for given K, as $\delta \to 0$, the C^k norm of $\zeta_{1\delta} \to 0$ (we will take δ small enough so that it is bounded by ϵ'). We now turn our attention to the cylindrical end of u_δ, which is of the form $[K, \infty) \times S^1$.

Proposition 3.11.17. *For K large, (which would take $\delta \to 0$ with it in order to satisfy our previous assumptions) we have $u_\delta|_{[K,\infty)\times S^1}$ converges in C^∞_{loc} to trivial cylinders. This is also true uniformly, i.e. for given $\epsilon'' > 0$, there is a K large enough so that for every small enough values of $\delta > 0$, $u_\delta|_{[k,k+1]\times S^1}$ is within ϵ'' (in the C^k norm) of a trivial cylinder of the form $\gamma \times \mathbb{R}$ for all values of k so that $[k, k+1] \times S^1 \subset [K, \infty) \times S^1$.*

Proof. The same proof as Proposition 3.11.1. □

Therefore we can choose a large enough K so that when u_δ is restricted to $[K, \infty) \times S^1$ the conditions for asymptotic estimates are met, namely we have the following:

Proposition 3.11.18. *We take $\epsilon'' > 0$ small enough so that previous convergence estimate near Morse-Bott torus applies. Then there is a large enough K, so that for small enough ϵ' (which depends on K), and for small enough $\delta > 0$ (which depends on ϵ'), there is a gradient trajectory v defined over the cylinder $(s_v, t_v) \in [K, \infty) \times S^1$ so that there is a vector field ζ_v over v_K so that*

$$u_\delta|_{[K,\infty)\times S^1} = \exp_{v_K}(\zeta_v)$$

and the norm of ζ_v measured in C^k satisfies the bound

$$\|\zeta_v(s_1')\| \le \|\zeta_v(K)\|^{2/p}_{L^2(S^1)} e^{-\lambda s}.$$

Retracing our footsteps we now construct a preglued curve $u_{r,a,p}$. There is a trivial cylinder $\gamma \times \mathbb{R}$ so that over the interval $[0, R)$ the difference between v_K and $\gamma \times \mathbb{R}$ is bounded above by $R\delta \to 0$, then choose (r, a, p) so that for $u^1 + (r, a, p)$, the difference between

$$|(u^1 + (r, a, p))(R, t_1') - (\gamma \times \mathbb{R})(R, t_1')| \leq C(e^{-DR} + R\delta)$$

then using this choice of (r, a, p) we construct a pregluing, by gluing together $u^1 + (r, a, p)(R, t_1')$ to $v_K(R - K, t_1')$ as we did in the section for gluing. This constructs for us a preglued map

$$u_{(r,a,p)} : (\Sigma_{K\delta}, \delta j_1) \cup [K, \infty) \times S^1 \longrightarrow M$$

so that there exists a vector field ζ so that over $(\Sigma_{K\delta}, \delta j_1) \cup [K, \infty) \times S^1$

$$u_\delta = \exp_{u_{r,a,p}}(\zeta).$$

We also have estimates of the size of ζ:

Proposition 3.11.19. *Over the semi-infinite interval $[0, \infty) \times S^1 \subset (\Sigma_{K\delta}, \delta j_1) \cup [K, \infty) \times S^1$ we impose the exponential weight e^{ds}, Then with respect to this exponential weight the $W^{k,p,d}$ norm of ζ is bounded above by:*

$$\|\zeta\| \leq C\epsilon'(C + e^{dK}) + CRe^{dR}\delta + Ce^{-D'K}$$

which can be made arbitrarily small by taking K large and $\epsilon' \to 0$ as $\delta \to 0$. In particular for given $\tilde{\epsilon}^2$ we can upper bound its norm by $\tilde{\epsilon}^2$ by taking $\delta \to 0$.

Proof. The same proof as Proposition 3.11.3. □

Then we truncate ζ into $\beta_v\zeta_v + \beta_1\zeta_1$ so that the pair (ζ_v, ζ_1) solves the equations $\Theta_u = 0, \Theta_v = 0$ near the cylindrical end. Here Θ_v is the equation living over the gradient cylinder v_K, and Θ_u lives over u^1. This process is entirely analogous to the previous section, to wit, we apply the contraction mapping principle over domains of the form $[R, \infty) \times S^1$ to show (ζ_v, ζ_1) can be extended to solutions of $\Theta_u = 0, \Theta_v = 0$. We note that we no longer need to worry whether ζ_v lands in H_0 because there is no such requirement over Θ_v. Further ζ_1 already lands inside image of Q_1 because we arranged this when we preglued the finite gradient trajectories, and that conclusion is unaffected by extensions of ζ_1 near the cylindrical neck. Hence we apply the above to each of the ends of u^i and in conjunction with the extension of vector fields along the finite gradient trajectories, we conclude :

Proposition 3.11.20. *The gluing construction is surjective in the case of 2-level cascades with one finite gradient cylinder segment in the intermediate cascade level. To be more precise, suppose u_δ degenerates into a transverse and rigid 2-level cascade $\{u^1, u^2\}$, with only one finite gradient trajectory in the intermediate cascade level, then u_δ corresponds (up to translation) to the unique solution of the system of equations $\mathbf{\Theta_v} = 0, \mathbf{\Theta}_u = 0$ with our given choice of right inverses.*

Multiple level cascades

In this subsection we generalize our result to multiple level cascades. The main subtlety is when two consecutive levels meet along multiple ends on an intermediate cascade level. Hence we take that up first in what follows. The main difficulty will be setting up notation.

2-level cascade meeting along multiple ends

Let u_δ converge to a 2 level cascade $\{u^1, u^2\}$. Each u^i is not necessarily connected. As before we first consider the finite gradient trajectories. The maps u^1 and u^2 meet along $i \in \{1, ..., N\}$ free ends in the middle. Consider the tuple (i,j) where $i \in \{1, .., N\}$ label the specific end, and $j \in \{1,2\}$ denotes whether the end belongs to u^1 or u^2. We fix cylindical ends around each puncture of the form $[0, \pm\infty) \times S^1$ (we won't bother labelling these with (i,j) to avoid further clutter of notation). Recall the vector spaces with asymptotic vectors we associate to each end that meets the intermediate cascade level of u^i, which we denote by $V_{(i,j)}$. Each of these vector spaces are spanned by asymptotic vectors $(\partial_a, \partial_z, \partial_x)$, we denote an element of these vector spaces by triples $(r, a, p)_{(i,j)}$. Recall there is a submanifold

$$\Delta \subset \oplus_{(i,j)} V_{(i,j)}$$

within an ϵ ball of the origin of $\oplus_{(i,j)} V_{(i,j)}$ so that if we used elements in Δ we would be able to construct a pregluing from the domains of u^1 and u^2. Recall the reason we have to do this is that, as we recall, moving each $p_{(i,j)}$ affects the a distance between u^1 and u^2, and we need to make sure that the ends (i,j) can be matched together.

Now given the degeneration of a J_δ-holomorphic curve u_δ to the cascade $u^{\not{t}} = \{u^1, u^2\}$, let $K > 0$ be large enough, for each end i there is a gradient flow trajectory v_i so that when restricted to the segment $[-s_i, s_i] \times S^1$, we have that

$$|v_i(s_i,t) - u^1(-K,t)|, |v_i(-s_i,t) - u^2(K,t)| \leq \epsilon'$$

and u_δ is very close to the gradient flow v_i. Then as before we can constructed a preglued curve $u_{(r,a,p)_{(i,j)}}$ so that over the domain of the preglued curve, we have

$$u_\delta = \exp_{u_{(r,a,p)_{(i,j)}}}(\zeta)$$

for ζ a global vector field whose norm can be taken to be arbitrarily small by picking K large enough and (consequently) ϵ' and δ small enough. Again here we are only worrying about the finite gradient trajectories, we will worry about the semi-infinite trajectories later.

Then we can split ζ into a sum of several other vector fields as before, namely we can write

$$\zeta = \zeta_1 + \zeta_2 + \sum_i \zeta^i$$

where $\zeta_i \in W^{2,p,d}(u^{i*}TM)$ for $i = 1,2$, and $\zeta^i \in W^{2,p,w}(v_i^* TM)$ for $i = 1, ..., N$. Using global a translation of entire cascade we can ensure $\zeta_1 \in ImQ_1$, and using a global increase in $p_{(i,1)} - p_{(i,2)}$ inside Δ we can ensure also $\zeta_2 \in ImQ_2$. Here the definition of Q_i is as before: we take compact neighborhoods of u^i and require the integral of $\langle \zeta, \partial_a \rangle$ over these neighborhoods is zero. This defines a codimension one subspace which we take to be the image of Q_i.

Finally to ensure $\zeta^i \in H_{0i}$. As before by exponential decay estimates the actual size of vector fields to make $\zeta^i \in H_{0i}$ are negligible compared to ϵ'. The difference from the previous case is that now there are multiple ends to worry about. To do this we need some understanding of Δ as a manifold.

Recall for near any point $x \in \Delta$, its tangent space is spanned by

$$\begin{aligned}
&\{r_{(i,1)}, r_{(i,2)}\}, \quad \{a_{(i,j)}\},\\
&\{p_{(1,1)} - p_{(1,2)} = T, p_{(i,1)} - p_{(i,2)} = T + \delta f_i(a_{(1,1)}, a_{(2,1)}, p_{(i,2)}, a_{(i,1)}, a_{(i,2)})\}
\end{aligned}$$

the functions f_i have uniformly bounded C^1 norm. The reason they appear is because ends meeting at different values of f travel different amounts of a distance for the same change of p, so a correction term is needed so the preglued curve can be constructed.
Recall that for $\zeta_i \in H_{0i}$ we must have the functionals

$$L_{i,*}(\zeta_i) = 0, \quad * = r, a, p.$$

For $* = r$, this can be adjusted for each i by a change in $\{r_{(i,1)} = r_{(i,2)}\}$. For $* = p, a$, we first repeat the previous construction for $i = 1$ verbatim to get vector fields $\zeta^1 \in H_{01}$ while keeping $\zeta_i \in \Im Q_i$. I.e. we take $a_{(1,j)}, p_{(1,j)}$ so that it does not induce global translations in a direction of the thick parts of u^1, u^2 as they enter the pregluing to ensure $\zeta^1 \in H_{01}$. For any other $i > 1$, the only constraint is $p_{(i,1)} - p_{(i,2)} = T + f_i(a_{(1,1)}, a_{(2,1)}, p_{i,2}, a_{(i,1)}, a_{(i,2)})$, hence as before we first change $p_{(i,j)}$ simultaneously by Δp_i to make $L_p(\zeta^i) = 0$ and in this process we adjust $a_{(i,j)}$ to make the pregluing condition still hold. Finally we change $a_{(i,j)}$ by the same amount $\Delta a_{i,j}$ to make $L_a(\zeta^i) = 0$ while preserving the previous equalities.

Using the same kind of machinery to extend the vector field ζ to solutions of $\Theta_1 = 0, \Theta_2 = 0, \Theta_{v_i} = 0$, and using the exactly the same set up for semi-infinite gradient trajectories, we arrive at the follow proposition:

Proposition 3.11.21. *If a sequence of J_δ-holomorphic curves u_δ degenerates into a transverse and rigid 2-level cascade $\{u^1, u^2\}$, then u_δ comes from the unique solution to our gluing construction, namely, $\Theta_1 = 0, \Theta_2 = 0, \Theta_{v_i} = 0$, subject to our choice of right inverses.*

General case

The general case proceeds largely analogously to the 2-level case. We shall be very brief in sketching it out. Assuming u_δ degenerates into an n-level transverse and rigid cascade, $u^{\cancel{t}} = \{u^1, .., u^n\}$, then we use the notation $v_{(i,j)}$ to denote a finite gradient trajectory connection between u^i and u^{i+1}, connecting between the jth end in that intermediate cascade level. As before we can find a pregluing $u_{pre} : \Sigma \to M$ depending on the data $(r, a, p)_{(i,j)} \in \oplus V_{(i,j)}$ so that there is a global vector field ζ so that

$$u_\delta = exp_{u_{pre}}(\zeta)$$

where ζ has very small norm. and as before we split

$$\zeta = \sum_i \zeta_i + \sum_{(i,j)} \zeta^{i,j}$$

for the intermediate cascade levels. by adjusting the asymptotic vector fields $p_{i,j}$ we can ensure $\zeta_i \in ImQ_i$, and using the same kind of adjustments as above we make sure $\zeta^{(i,j)} \in H_{0ij}$. Finally using the same analysis we extend them to solutions of $\mathbf{\Theta}_* = 0$ - here we just mean the system of equations we used in the gluing construction, using the same kind of analysis to take case of semi-infinite gradient ends. Hence we have proved:

Theorem 3.11.22. *The gluing construction is surjective in the following sense: if u_δ converges to a n-level transverse and rigid cascade $u^\natural$. Then for each such cascade $u^\natural$ after our choice of right inverses we constructed a unique glued curve for $\delta > 0$ small enough, and u_δ agrees with this glued curve up to translation in the symplectization direction.*

Remark 3.11.23. We note our theorem about correspondence between transverse rigid cascades and rigid J_δ-holomorphic curves studies the correspondence of a *single* cascade and a *single* curve. Usually in Floer theory one needs to show the collection of all transverse and rigid cascades is in bijection with the collection of all rigid holomorphic curves. To apply our results in these circumstances one usually needs some finiteness assumptions on the cascades and the holomorphic curves. For more details see Chapter 2 [66].

3.12 Appendix: SFT compactness for cascades

In this appendix we outline the SFT compactness result required for the degeneration of holomorphic curves to cascades.

We borrow heavily the results and notation from the original SFT compactness paper [6]. In fact our compactness theorem will follow from their setup in combination with our estimates of how J_δ-holomorphic curves behave near Morse-Bott tori. The behaviour of holomorphic curves near a Morse-Bot torus is already discussed in Chapter 4 of [5], and is implicit in [7], for example their Section 4.2 and Appendix. Hence this appendix is more of an expository nature for the sake of completeness, and we will point out the differences and similarities between our results and theirs in the course of proving our version of SFT compactness theorem.

Deligne-Mumford moduli space of Riemann surfaces

We begin with a review of the Deligne- Mumford compactification of stable Riemann surfaces. Most of the material in this section is taken directly from Section 4 of [6], but is repeated for the convenience of the reader.

Let $\mathbf{S} = (S, j, M)$ denote a closed Riemann surface S with complex structure j with marked points set M. The surface is called *stable* if $2g + \mu \geq 3$, where g is the genus and $\mu := |M|$ is the number of marked points. Stability implies the automorphism group of the surface $\mathbf{S}$ is finite.

The uniformization theorem equips $\dot{\mathbf{S}} := (S \setminus M, j)$ with a unique complete hyperbolic metric of constant curvature and finite volume, which we denote by $h^{\mathbf{S}}$. Each puncture in $\dot{\mathbf{S}}$ corresponds to a cusp in the metric. We let $\mathcal{M}_{g,\mu}$ denote the moduli space of Riemann surfaces of signature (g, μ).

Thick-Thin decomposition

Fix $\epsilon > 0$, given a stable Riemann surface $\mathbf{S}$, for $x \in \dot{\mathbf{S}}$ let $\rho(x)$ denote the injectivity radius of $h^{\mathbf{S}}$ at x. As in Section 4 of [6], we denote by $\mathrm{Thin}_\epsilon(\mathbf{S})$ and $\mathrm{Thick}_\epsilon(\mathbf{S})$ its ϵ-thin and thick parts where

$$\mathrm{Thin}_\epsilon(\mathbf{S}) := \overline{\{x \in \dot{\mathbf{S}} | \rho(x) < \epsilon\}}$$

$$\mathrm{Thick}_\epsilon(\mathbf{S}) := \{x \in \dot{\mathbf{S}} | \rho(x) \geq \epsilon\}.$$

It is a fact of hyperbolic geometry that there is a constant $\epsilon_0 = sinh^{-1}(1)$ so that for all $\epsilon < \epsilon_0$ we have each component of $\mathrm{Thin}_\epsilon(\mathbf{S})$ is conformally equivalent to either a finite cylinder of the form $[-L, L] \times S^1$ or semi-infintie cylinder $[0, \infty) \times S^1$. Each compact component of the form $C = [-L, L] \times S^1$ contains a unique closed geodesic of length equal to $2\rho(C)$, which we denote by Γ_C. Here we set $\rho(C) := \inf_{x \in C} \rho(x)$.

Oriented blow up of punctured Riemann surface

Given $\mathbf{S} = (S, j, M)$, let $z \in M$, then as in [6] we can define the oriented blow up S^z as the circle compactification of $S \setminus z$ with boundary $\Gamma_z = T_z S / \mathbb{R}_+^*$. The complex structure j defines an S^1 action on Γ_z. The surface S^z comes equipped with a map $\pi : S^z \to S \setminus \{z\}$ which collapses the blown up circle. Given a finite set $M = \{z_1, ..., z_k\}$ we can similarly define the blown up space S^M with boundary circles $\Gamma_1, .., \Gamma_k$, with projection $\pi : S^M \to S \setminus M$ that collapses the boundary circles.

Stable nodal Riemann surface

See Section 4.4 in [6]. Let $\mathbf{S} = (S, j, M, D)$ be a possibly disconnected Riemann surface, where M, D are both marked points, and the cardinality of D is even. We write $D = \{\overline{d}_1, \underline{d_1}, ..., \overline{d_k}, \underline{d_k}\}$. The *nodal Riemann surface* is the tuple $\mathbf{S} = (S, j, M, D)$ under the additional equivalence relations so that each pair $(\overline{d_i}, \underline{d_i})$ and the set of all such special pairs are unordered.

From a given nodal Riemann surface $\mathbf{S} = (S, j, M, D)$ we can construct the following singular surface

$$\hat{S}_D := S/\{\overline{d_i} \sim \underline{d_i}, i = 1, .., k\}. \tag{3.24}$$

The arithmetic genus of a nodal Riemann surface is defined to be $g = \frac{1}{2}\#D - b_0 + \sum_{i=1}^{b_0} g_i + 1$, where b_0 is the number of connected components of S. The *signature* of a nodal Riemann surface is given by the pair (g, μ), where g is the arithmetic genus and μ is the number of marked points in M.

A stable Riemann surface $\mathbf{S} = (S, j, M, D)$ is called *decorated* if for each pair $(\overline{d_i}, \underline{d_i})$ we include the information of orientation reversing orthogonal map

$$r_i : \overline{\Gamma_i} := (T_{\overline{d_i}} S \setminus 0)/\mathbb{R}_{>0} \longrightarrow \underline{\Gamma_i} := (T_{\underline{d_i}} S \setminus 0)/\mathbb{R}_{>0}. \tag{3.25}$$

We also consider *partially decorated Riemann surfaces* where such r_i maps are only given for a subset $D' \subset D$.

We consider the moduli space of nodal Riemann surfaces $\overline{\mathcal{M}}_{g,\mu}$ and decorated nodal Riemann surface $\overline{\mathcal{M}}^{\$}_{g,\mu}$ of signature (g, μ). The moduli space of smooth Riemann surfaces of signature (g, μ), which we write as $\mathcal{M}_{g,\mu}$, includes naturally in the above spaces. We refer the reader to Section 4.5 in [6] for detailed topologies of these spaces. For us we only need the notion of convergence, which we summarize below.

Given a decorated stable nodal Riemann surface (r denotes the decoration), which we write as $(\mathbf{S}, r) = (S, j, M, D, r)$, we first take its oriented blow up along points of D, to obtain boundary circles $\overline{\Gamma_i}$ and $\underline{\Gamma_i}$ associated to the pair $\{\overline{d_i}, \underline{d_i}\}$, then using the orthogonal maps r_i, we glue the resulting pieces together along $\overline{\Gamma_i}, \underline{\Gamma_i}$ and call the resulting surface $S^{D,r}$. The glued copy of $\overline{\Gamma_i}$ and $\underline{\Gamma_i}$ is called Γ_i. The surface $S^{D,r}$ has the same genus as the arithmetic genus of $(\mathbf{S}, r)$, and inherits a uniformizing metric from $h^{j,M\cup D}$, which we write as $h^{\mathbf{S}}$. The metric $h^{\mathbf{S}}$ is defined away from the Γ_i and points of M. We can talk about the thick/thin components of $\mathbf{S}$ and view them as subsets of $\dot{S}^{D,r}$. Every compact component C of $\overline{Thin_\epsilon(S)} \subset S^{D,r}$ is a compact annulus, it has either a closed geodesic which we denote by Γ_C, or one of the special circles Γ_i constructed above, which we will also denote by Γ_C.

Let $(\mathbf{S}_n, r_n) = \{S_n, j_n, M_n, D_n, r_n\}$ be a sequence of decorated stable nodal Riemann surfaces. We say $(\mathbf{S}_n, r_n)$ converges to a nodal stable Riemann surface $(\mathbf{S}, r) = (S, j, M, D, r)$ if for large enough n there are diffeomorphisms $\phi_n : S^{D,r} \to S_n^{D_n,r_n}$ with $\phi_n(M_n) = M$, and the following conditions hold (Section 4.5 in [6]):

- **CRS1** For all $n \geq 1$, the images $\phi_n(\Gamma_i)$ of the special circles $\Gamma_i \subset S^{D,r}$ for $i = 1, .., k$ are special circles or closed geodesics of the metrics $h^{j_n,M_n\cup D_n}$ on $\dot{S}^{D_n,r_n}$. Moreover, all special circles on S^{D_n,r_n} are among these images.
- **CRS2** $h_n \to h$ in $C^\infty_{loc}(S^{D,r} \setminus (M \cup \bigcup_1^k \Gamma_i))$ where $h_n := \phi_n^* h^{j_n,M_n\cup D_n}$.

- **CRS3** Given a component C of $Thin_\epsilon(\mathbf{S}) \subset \dot{S}^{D,r}$, which contains a special circle Γ_i, and given a point $c_i \in \Gamma_i$, we consider for every $n \geq 1$ the geodesic arc δ_i^n for the induced metric $h^n = \phi_n^* h^{j_n, M_n \cup D_n}$, which intersects Γ_i orthogonally at c_i (even though the distance is infinite it still makes sense to talk about geodesics intersecting orthogonally at infinity), and whose ends are contained in the ϵ-thick parts of h^n. Then $C \cap \delta_i^n$ converges as $n \to \infty$ in C^0 as a continuous geodesic for $h^{\mathbf{S}}$ which passes through the point c_i.

We note that **CRS2** is equivalent to $\phi_n^* j_n \to j$ in $C^\infty_{loc}(S^{D,r} \setminus (M \cup \bigcup_1^k \Gamma_i))$. The topology on $\overline{\mathcal{M}}_{g,\mu}$ is defined to be the weakest topology for which the forgetful map $\overline{\mathcal{M}}^{\$}_{g,\mu} \to \overline{\mathcal{M}}_{g,\mu}$ defined by forgetting the r_i is continuous. Finally the compactness theorem.

Theorem 3.12.1 (Theorem 4.2 in [6]). *The spaces $\overline{\mathcal{M}}_{g,\mu}$ and $\overline{\mathcal{M}}^{\$}_{g,\mu}$ are compact metric spaces that contain $\mathcal{M}_{g,\mu}$, and are equal to the closure of the inclusion of $\mathcal{M}_{g,\mu}$ (i.e. they are compactifications of $\mathcal{M}_{g,\mu}$). As we are in a metric space, sequential compactness suffices.*

We now state a proposition which we will later need to find all components of a holomorphic building/cascade.

Proposition 3.12.2 (Proposition 4.3 in [6]). *Let $\mathbf{S}_n = (S_n, j_n, M_n, D_n)$ be a sequence of smooth marked nodal Riemann surfaces of signature (g, μ) which converges to a nodal curve $\mathbf{S} = (S, j, M, D)$ of signature (g, μ). Suppose for each $n \geq 1$ we are given a pair of points $Y_n = \{y_n^1, y_n^2\} \subset S_n \setminus (M_n \cup D_n)$ so that*

$$dist_n(y_n^1, y_n^2) \longrightarrow 0 \tag{3.26}$$

where $dist_n$ is with respect to the hyperbolic metric $h^{j_n, M_n \cup D_n}$. Suppose in addition there is a sequence $R_n \to +\infty$ such that there exists injective holomorphic maps $\phi_n : D_{R_n} \to S_n \setminus (M_n \cup D_n)$ where D_{R_n} is the disk in $\mathbb{C}$ with radius R_n,satisfying $\phi_n(0) = y_n^1, \phi_n(1) = y_n^2$. Then there exists a subsequence of the new sequence $\mathbf{S}'_n = (S_n, j_n, M_n \cup Y_n, D_n)$ which converges to a nodal curve $\mathbf{S}' = (S', j', M', D')$ of signature $(g, \mu + 2)$, which has one or two additional spherical components. One of these components contains the marked points y^1, y^2, which corresponds to the sequence y_n^1, y_n^2. The possible cases are illustrated in Fig 5 of [6].

Stated in words (and also explained in Section 4 of [6]), the scenarios are as follows. Let r_n and r be decorations on stable nodal Riemann surfaces $\mathbf{S}_n$ and $\mathbf{S}$ respectively, and we have $\mathbf{S}_n \to \mathbf{S}$ in the sense specified above, and $\phi_n : S^{D,r} \to S_n^{D_n, r_n}$ be the corresponding diffeomorphism. Let $\hat{S}_D$ be the singular nodal Riemann surface obtained from $\mathbf{S}$ by gluing together the nodal points D, and $\pi : S^{D,r} \to \hat{S}^D$ the associated projection. Let $Z_n = \pi(\phi^{-1}(Y_n)) \subset \hat{S}_D$. Then the following can happen:

- The points $z_n^1, z_n^2 \in Z_n$ converge to a point z_0, which does not belong to M or D. Then the limit $\mathbf{S}'$ of $\mathbf{S}'_n$ has an extra sphere attached at z_0 on which lie two extra points y^1, y^2.

- The points $z_n^1, z_n^2 \in Z_n$ converge to a marked point $m \in M$. In this case the limit $\mathbf{S}'$ is $\mathbf{S}$ with two extra sphere T_1 and T_2 attached. The sphere T_1 is attached at ∞ to the original m, and has m at its zero. The sphere T_2 has its ∞ point attached to $1 \in T_1$ and y^1, y^2 lie on T_2.

- The points $z_n^1, z_n^2 \in Z_n$ converge to a double point d corresponding to pair of points $\{x, x'\} \in D$. Then we insert a sphere T_1 between nodes x, x', with x attached to ∞ on T_1 and x' attached to 0. We insert a second sphere T_2 whose ∞ point is attached to 1 in T_1, and the two points y_1 and y_2 lie on T_2.

SFT compactness theorem for Morse-Bott degenerations

We are now ready to state the SFT compactness theorem for degeneration of holomorphic curves to cascades. We first state a more careful definition of holomorphic cascades of height 1, taking into account of decorations. This is also taken directly out of [6].

Recall λ is a Morse-Bott contact form, and λ_δ is its perturbation defined by $\lambda_\delta = e^{\delta f}\lambda$. Fix $L >> 0$, then for all $\delta > 0$ small enough all Reeb orbits with action $< L$ come from critical points of f on each Morse-Bott torus.

Definition 3.12.3 (Section 11.2 in [6]). *Suppose we are given:*

- *n nodal stable J-holomorphic curves*

$$u^i := (a^i, \hat{u}^i; S_i, D_i, \overline{Z}^i \cup \underline{Z_i}), i = 1, ..., n \tag{3.27}$$

  where u^i is a J-holomorphic map from S^i to $\mathbb{R} \times Y$. The map a^i goes from S_i to $\mathbb{R}$, the symplectization direction; and $\hat{u}^i$ is the map to Y. The sets $\overline{Z}^i, \underline{Z_i}$ correspond to punctures that are asymptotic to Reeb orbits hit by u^i at $s = +\infty$ and $s = -\infty$ respectively. Let $\Gamma_i^\pm$ denote the corresponding boundary circle after blowing up the marked points $\overline{Z}^i$, $\underline{Z}_i$ respectively.

- *$n+1$ collections of cylinders that are lifts of gradient trajectories of f along the Morse-Bott tori, which we write as*

$$\{G_{j,i,T_i}, j = 1, .., p_i\},\ i = 0, .., n. \tag{3.28}$$

  In the above, i indexes which collection the cylinder is in, and j indexes specific element in that collection. Said another way, i indexes the specific level in the cascade, and j refers to which gradient flow segment in the level. The numbers T_i denotes the flow time along gradient flow of f, with $T_0 = -\infty$, $T_n = \infty$, $0 \leq T_i < \infty$ for $i = 1, .., n$. Denote the domain of the cylinders by $\tilde{S}_i$, and $\tilde{\Gamma}_i^\pm$ their boundary circles corresponding at positive/negative ends. Even though the gradient flows may be finite, we think of these domain cylinders as infinitely long, and will think of them as living in the thin part of the glued domain Riemann surface. We do this even if the flow time T_i is zero.

- *Each positive puncture of u^i (with $i = 1, .., n$) is matched with a negative puncture of u^{i-1}, where they cover Reeb orbits on the same Morse-Bott torus of the same multiplicity. Between these two matched pair of punctures there is a unique gradient trajectory G_{j,i,T_i} that connects between them after gradient flow of time T_i. Then there are orientation reversing diffeomorphisms $\Phi_i : \Gamma_i^+ \to \tilde{\Gamma}_i^-$ and $\Psi_{i-1} : \tilde{\Gamma}_{i-1}^+ \to \Gamma_i^-$ which are orthogonal on each boundary component.*

- *We glue the domains $S_i^{Z_i}$ and $\tilde{S}_i$ via the maps Φ_i and Ψ_i, to obtain a surface*
$$\overline{S} := \tilde{S}_0 \cup_{\Psi_0} S_z^{Z_1} \cup_{\Phi_1} \cup ... \cup_{\Phi_n} \tilde{S}_n. \tag{3.29}$$
The maps u^i and G_{j,i,T_i} fit together to define a continous map from $\overline{u} : S \to \mathbb{R} \times Y$. Here for defining $\overline{u}$, on the gradient segment parts we use the literal gradient flow of f without re-scaling by δ.

- *For the surface $\overline{S}$, we describe its complex structure. The idea is to keep the thin parts corresponding to $\underline{Z}_i, \overline{Z}^{i+1}$, and insert between them an infinite cylinder corresponding to the connecting gradient trajectory (with one marked point added to make it stable), with now $\underline{Z}_i, \overline{Z}^{i+1}$ viewed as nodal points which comes with their own special circles. In our case, two points among $\underline{Z}_i, \overline{Z}^{i+1}$ are viewed as nodal points for each gradient segment we are gluing in. Then the new decorated Riemann surface underlying the cascade can then be written as*
$$(\overline{S}, M = \bigcup M_i \cup \{\text{one for each gradient flow segment}\}, \tag{3.30}$$
$$D = \bigcup_i \overline{Z}^i \cup \underline{Z_i} \cup \{\text{punctures corresponding to gradient flow cylinders}\}) \tag{3.31}$$
We note this does not necessarily guarantee the stability of the underlying domain $\overline{S}$, since the definitions of stability of Riemann surface and J holomorphic curves are distinct (see remark below). However we can always add several marked points M' to make the underlying nodal Riemann surface stable.

Then we say we have defined a n level J holomorphic cascade of curves of height 1.

Remark 3.12.4. In the above definition by stable we mean stable in the sense of J- holomorphic curves, i.e. no level consists purely of trivial cylinders, and if a component of J holomorphic curve is constant, then the underlying domain for that component is stable in the sense of Riemann surfaces. We will treat the issue of stability of domain separately.

The definition of height k holomorphic cascade is very similar, we stack k height 1 cascades on top of one another, and identify the edge punctures with maps like Ψ and Φ. The definition of when two cascades are equivalent to one another is identical to the definition in Section 7.2 of [6] of when two SFT buildings are equivalent to one another, with the addition that we think of gradient flow trajectories in the cascade as extra levels with marked points.

Then we are ready to state the SFT compactness result.

Definition 3.12.5 (Section 11.2 of [6]). *Let $(u_{\delta_n}, S_n, j_n, M_n, D_n, r_n)$ be a sequence of J_{δ_n}-holomorphic curves. And let $u^{\natural} = \{u^1, .., u^m\}$ be a height k holomorphic cascade (we allow k infinite flow times), and let (S, j, M, D, r) be the underlying decorated nodal Riemann surface. We say $(u_{\delta_n}, S_n, j_n, M_n, D_n, r_n)$ converges to $u^{\natural}$ if we can find an extra set of marked points M' on (S, j, M, D, r), and an extra sequence of marked points M'_n on $(u_{\delta_n}, S_n, j_n, M_n, D_n, r_n)$ to make the underlying nodal Riemann surfaces stable, with diffeomorphisms $\phi_n : S^{D,r} \to S^{D_n,r_n}$ with $\phi_n(M) = M_n$ and $\phi_n(M') = M'_n$ satisfying the convergence definition of stable decorated Riemann surfaces in* **CRS1 − 3**, *and suppose in addition the following conditions hold:*

- **CGHC1** *For every component C of $S^{D,r} \setminus \bigcup \Gamma_i$ which is not a cylinder coming from a gradient flow, identify the corresponding component C_n in $(u_{\delta_n}, S_n, j_n, M_n \cup M'_n, D_n, r_n)$, and if we write $u_{\delta_n} = (a_{\delta_n}, \hat{u}_{\delta_n})$ and similarly for $u^{\natural}$. Then $\hat{u}_{\delta_n}|_{C_n}$ converges to $\hat{u^{\natural}}|_C$ in $C^{\infty}_{loc}(Y)$*
- **CGHC2** *If C_{ij} is the union of components of $S^{D,r} \setminus \bigcup \Gamma_i$ which correspond to the same level j of height i of $u^{\natural}$, (recall specifying height i specifies a height 1 cascade, and j labels the level within that height 1 cascade), then there exists sequences c_n^{ij} so that $a_n \circ \phi_n - a - c_n^{ij}|_{C_{ij}} \to 0$ in C^{∞}_{loc}*

Then we say the sequence of $J_{\delta,n}$-holomorphic curves are convergent to the J-holomorphic cascade $u^{\natural}$.

Theorem 3.12.6. *[Theorem 11.4 in [6]] Let $\delta_n > 0$ and $\delta_n \to 0$, let $(u_{\delta_n}, S_n, j_n, M_n \cup M'_n, D_n, r_n)$ denote a sequence of J_{δ_n}-holomorphic curves of fixed signature and asymptotic to the same Reeb orbits (recall as long as $\delta > 0$ and all orbits have energy $< L$, the orbit themselves do not depend on δ), then there exists a subsequence that converges to a J-holomorphic cascade of height k.*

The rest of this appendix is dedicated to the proof of this theorem. First a theorem on gradient bounds:

Theorem 3.12.7. *(Gradient Bounds, Lemma 10.7 in [6]) Let $\delta_n \to 0$ and $(u_{\delta_n}, S_n, j_n, M_n \cup M'_n, D_n, r_n)$ be a sequence of J_δ-holomorphic curves with fixed signature, and the curves u_{δ_n} have a uniform energy bound E. Then by Deligne-Mumford compactness the domain $(S_n, j_n, M_n \cup M'_n, D_n, r_n)$ converges in the sense of* **CRS1 − 3** *to a decorated Riemann surface*

$$(\mathbf{S}, j, M, D, \overline{Z} \cup \underline{Z}, r)$$

Then there exists a constant K which only depends on the upper energy bound E so that if we add to each M_n an additional collection of marked points

$$Y_n = \{y_n^1, w_n^1, ..., y_n^K, w_n^K\} \subset \dot{S}_n = S_n \setminus (M_n \cup \underline{Z_n} \cup \overline{Z_n})$$

we have the following uniform gradient bound

$$\|\nabla u_{\delta_n}(x)\| \leq \frac{C}{\rho(x)} \tag{3.32}$$

Here the gradient $\nabla u_{\delta_n}(x)$ is measured with respect the fixed $\mathbb{R}$ invariant metric in $\mathbb{R} \times Y$ in conjunction with the hyperbolic metric on $\dot{S}_n \setminus Y_n$, and $\rho(x)$ is the injectivity radius of the hyperbolic metric at x.

Proof. The same proof as in for Lemma 10.7 in [6] goes through. The only two observations needed are: first the analogue of lemma 5.11 continues to hold even as we take $J_{\delta_n} \to J$. The second observation is that due to Morse-Bott assumption each plane or sphere that bubbles off also has a lower nonzero bound on energy, so the set of points that bubbles off is finite. □

Proof of Theorem 3.12.6. **Step 1**. We first discuss convergence in the thick parts. The discussion largely mirrors the discussion of [6] Section 10.2.2. Following the setup in Theorem 3.12.7, we assume we have added enough marked points to the converging Riemann surfaces $(S_n, j_n, M_n \cup M_n', D_n, r_n)$ so that the gradient bound holds everywhere away from the marked points. We call the limit of the sequence $(\mathbf{S}, j, M, D, \overline{Z} \cup \underline{Z}, r)$. We let Γ_i denote the special circles on S, then we may assume

$$\|\nabla u_{\delta_n} \circ \phi_n(x)\| \leq \frac{C}{\rho(x)}, x \in S \setminus \bigcup \Gamma_i$$

where ϕ_n is the diffeomorphism from $S \to S_n$ (defined away from the nodes) given by the definition of convergence. Then by Azerla-Ascoli and Gromov-Schwarz we can extract a subsequence that over the thick parts of S_n converges in $C^\infty_{loc}(Y \times \mathbb{R})$ to a J holomorphic map defined on thick parts of $\mathbf{S}$.

Step 2. Next we consider what happens on the thin parts near a node, following [6] Section 10.2.3. Let $C_1, .., C_N$ denote the connected components of $S \setminus \cup \Gamma_i$, we have from the above discussion that $u_{\delta_n} \circ \phi$ converges to J-holomorphic maps in C^∞_{loc} over each of C_i. Call these maps u_i. The point is in this description there may be levels missing near the nodes that connect between C_i, and by examining closely what happens near the nodes we recover the entire cascade.

The first case is if u_i is bounded in $\mathbb{R} \times Y$ near one of the nodes, then by the removal of singularities theorem then u_i extends continuously to the node. If u_i is unbounded near a node then it must converge to a Reeb orbit, and extend continuously to the circle at infinity which compactifies the puncture.

Given a pair of components of $S \setminus \cup \Gamma_i$, call them C_i and C_j, that are adjacent to each other. The behaviour of u_i and u_j could be quite different. The maps u_i and u_j may be asymptotic to either a point or a Reeb orbit at their connecting node, and even if they are both asymptotic to Reeb orbits they might not even be asymptotic to the same one (not even Reeb orbits that land on the same Morse-Bott torus). The reason for this, as explained above, is that there may be further degenerations of the curve u_{δ_n} near this node. To capture this idea, let $\gamma^\pm$ denote the the asymptotic limit of u_i and u_j (which could be either a point or a Reeb orbit), then there is a component T_n^ϵ of the ϵ-thin region of the hyperbolic metric $h^n = \phi_n^* h^{j_n, M_n}$ on $S = S^{D,r}$, with conformal parametrization

$$g_n^\epsilon : A_n^\epsilon := [-N_n^\epsilon, N_n^\epsilon] \times S^1 \longrightarrow T_n^\epsilon$$

such that in $C^\infty(S^1)$

$$\lim_{\epsilon\longrightarrow 0}\lim_{n\longrightarrow\infty} \hat{u}_{\delta_n}\circ\phi_n\circ g_n^\epsilon|_{\pm N_n^\epsilon\times S^1}=\gamma^\pm$$

Note that g_n^ϵ can be chosen to satisfy

$$\|\nabla g_n^\epsilon(x)\|\leq C\rho(g_n^\epsilon(x))$$

where the norm on the left hand side is measured with respect to the flat metric on the source and the hyperbolic metric on the target. Then under this parametrization we have

$$\|\nabla u_{\delta_n}\circ\phi_n\circ g_n^\epsilon\|\leq C$$

and if we take a subsequence of $\epsilon_k\to 0$ (also denoted by ϵ_k), then we have

$$\lim_{k\to+\infty}\hat{u}_{\delta_n}\circ\phi_k\circ g_k^{\epsilon_k}(\pm N_k\times S^1)=\gamma^\pm$$

and hence obtain a homotopically unique map $\Phi:[0,1]\times S^1\to Y$ satisfying $\Phi(0\times S^1)=\gamma^-$ and $\Phi(1\times S^1)=\gamma_+$.

Sinice we have a uniform bound on $\|\nabla u_{\delta_n}\circ\phi_n\circ g_n^\epsilon\|$, by Azerla-Ascoli it converges in C^∞_{loc} to holomorphic curves (which specfic curve it converges to might depend on which shift we are considering on the domain, this is akin to the degeneration of a gradient flow line to a broken gradient flow line in the Morse case). We break it down in to cases:

- $\int_{\gamma^+}\lambda-\int_{\gamma^-}\lambda=0$
- $\int_{\gamma^+}\lambda-\int_{\gamma^-}\lambda>0.$

Case 1: We first consider when $\int_{\gamma^+}\lambda-\int_{\gamma^-}\lambda=0$. If both $\gamma^\pm$ are points, then they are connected by a sequence of J-holomorphic spheres touching each other at nodes, however in symplectizations all J holomorphic sphere are points, so in this case $\gamma^\pm$ are the same point.

If one of the ends (say γ^+) is a Reeb orbit, and Γ^- is a point. The fact that $u_{\delta_n}\circ\phi_n\circ g_n^\epsilon$ converges in C^∞_{loc} implies we can find a J holomorphic plane with γ^+ as its positive puncture. But then this J holomorphic plane must have zero energy, which contradicts the Morse-Bott assumption.

The last case is if both $\gamma^\pm$ are Reeb orbits. Then they must lie on the same Morse-Bott Torus, because the energy of the segment $u_{\delta_n}\circ\phi_n\circ g_n^\epsilon|_{[-N_n,N_n]\times S^1}$ approaches zero as $n\to\infty$, and there is not enough energy to support a cylinder connecting Reeb orbits from one Morse-Bott torus to another, hence the Reeb orbits must lie on the same Morse-Bott torus.

Then by Lemma 3.11.1 for large enough n the derivatives of $u_{\delta_n}\circ\phi_n\circ g_n^\epsilon$ are pointwise bounded by $\epsilon>0$, then by Propositions 3.10.3, 3.10.1, there is a number $T\in[0,\infty]$, a segment of gradient trajectory of f of time T_n, lifted to be a J_{δ_n}-holomorphic curve, which we denote by v_{δ_n}, such that after taking a subsequence, over $[-N_n,N_n]\times S^1$, $u_{\delta_n}\circ\phi_n\circ g_n^\epsilon$ is $C^k([-N_n,N_n]\times S^1)$ close to v_δ.

To elaborate a bit more, we note Propositions 3.10.3, 3.10.1 only apply when we can establish the segment of J_δ-holomorphic cylinder is uniformly bounded away from all except at most one critical point of f. If this is not the case, then necessarily $T_n \to +\infty$. We assume $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon(-N_n \times S^1)$ approaches the minimum of f and $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon(N_n \times S^1)$ approaches the maximum of f. Then we can choose $L_n \in [-N_n, N_n]$ so that $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ restricted to $[-N_n, L_n] \times S^1$ is uniformly bounded away from the maximum of f, and its restriction to $[L_n, N_n] \times S^1$ is uniformly bounded away from the minimum of f. Then we apply Proposition 3.10.1 to find two semi-infinite gradient cylinders $v_{\delta_n -}$ and $v_{\delta_n +}$ to which the restriction of $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ to $[-N_n, L_n] \times S^1$ (resp. $[L_n, N_n] \times S^1$) converges in C^k norm. By local convergence the restriction of $v_{\delta_n -}$ and $v_{\delta_n +}$ to $L_n \times S^1$ are C^k close to each other, so for our purposes[6] [7] we can take v_{δ_n} to be either $v_{\delta_n +}$ or $v_{\delta_n -}$.

The estimate we proved for its local behaviour also tells us how to define the relevant gluing maps Φ_i and Ψ_i. We should also attach a marked point to this cylindrical segment to make the domain stable.

Case 2: We consider the second case $\int_{\gamma^+} \lambda - \int_{\gamma^-} \lambda > 0$. We first observe that there is a lower bound on $\int_{\gamma^+} \lambda - \int_{\gamma^-} \lambda$ by the Morse-Bott assumption. We shall see that over $[-N_n, N_n] \times S^1$ the map $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ converges to a sequence of J holomorphic cylinders (and in the case where γ^- is a point, a collection of cylinders followed by a J-holomorphic plane) connected by gradient cylinders along Morse-Bott tori similar to the previous case. We first observe by the gradient bounds that there is no bubbling off of holomorphic planes, and that over any compact domain of $[-N_n, N_n] \times S^1$ the sequence $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ converges uniformly to a J holomorphic curve. We note this is very similar to the case in Morse theory where a gradient trajectory converges to a broken gradient trajectory.

Let h denote the minimal energy of a nontrivial J holomorphic curve in the Morse-Bott setting, after successively taking subsequences, we pick out numbers $a_n^i, b_n^i \in [-N_n, N_n]$ which partition the interval $[-N_n, N_n]$ so that the following holds:

- $b_n^i - a_n^i \to \infty$, $a_n^{i+1} - b_n^i \to \infty$.
- $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ converges uniformly to a nontrivial J holomorphic curve u^i over $[a_n^i, b_n^i]$.
- $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ restricted to $[b_n^i, a_n^{i+1}]$ has energy $< h/20$. We shall show that in fact the energy goes to zero as $n \to \infty$.

We first observe by our assumptions there must be an interval of the form $[a_n^i, b_n^i]$, because otherwise the entire interval $[-N_n, N_n]$ the curve $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ has energy less than $h/20$,

[6]Establishing exponential decay estimates for gradient flow lines that go from critical point of f to another critical point requires more careful analysis, and is outside the scope of this work. Incidentally this is related to the problem of gluing cascades of height greater than 1 - we need to think more carefully about how we choose our Sobolev spaces and place our exponential weights.

[7]We mention here our work is simplified because our critical manifold (the manifold that parametrizes the space of Reeb orbits) is S^1, hence there are no broken gradient trajectories. In the case where the critical manifold is higher dimensional the analysis near broken trajectories is more delicate, and is outside the scope of the current work. However it is probably within the convex span of current technology.

hence over each compact subset the curve converges to trivial cylinders, then this implies that γ^+ and γ^- are on the same Morse-Bott torus, which is the situation in case 1.

We note in the second bullet point we required *uniform* convergence over the interval $[a_n^i, b_n^i]$, as opposed to the usual convergence over compact set. The reason is that if we had C^∞_{loc} convergence over an interval of the form $[a_n^i, b_n^i]$, and by looking at different compact subsets in the domain we got convergence in C^∞_{loc} into two different curves, we could have inserted more partitions into the interval $[a_n^i, b_n^i]$ until the three bullet points above are achieved.

We also observe that the evaluation maps $ev^+(u^i)$ and $ev^-(u^{i+1})$ land in the same Morse-Bott torus, since over $[b_n^i, a_n^{i+1}]$ the energy is too small to cross from one Morse-Bott torus to the next, hence in fact the energy of $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ over $[b_n^i, a_n^{i+1}]$ converges to zero. As in the proof of the previous case, over the interval $[b_n^i, a_n^{i+1}]$, the map $u_{\delta_n} \circ \phi_n \circ g_n^\epsilon$ converges uniformly to a gradient flow trajectory v^i, as was shown in the previous case. As a technical point, once we have found u^i and u^{i+1}, we should identify the region when they first enter the Morse-Bott torus, and perform the analysis as we did in propositions 3.10.3, 3.10.1 to identify the correct length of the gradient trajectory (this may result in us moving the partition points b_n^i, a_n^i to lengthen the segment that we think of as being the gradient trajectory). We add marked points to both domains of v^i and u^i to make the domain stable, and the gluing map Φ and Ψ are naturally supplied by considerations of convergence.

Step 3: We remark that the behaviour of u_{δ_n} near a puncture (either symptotic to a Reeb orbit or to a point), around which the hyperbolic metric produces another thin region, is entirely analogous to the analysis we performed above: we can choose a cylindrical neighborhood of the form $[0,\infty) \times S^1$ or $(-\infty, 0] \times S^1$, and along this neighborhood the curve u_{δ_n} reparametrized as above degenerates into a cascade of cylinders connected by Morse flow lines. The only additional piece of information which follows readily is that if in the original u_{δ_n} is asymptotic to γ near this puncture, then the end of the chain of holomorphic cylinders and gradient trajectories also is also asymptotic to γ.

Step 4: Finally we discuss the level structure.

Recall that after the previous modifications the domain of u_{δ_n} converges to a stable Riemann surface $(S, j, M, D, \overline{Z} \cup \underline{Z}, r)$ so that each connected component of $S \setminus D_i$ is assigned either a J-holomorphic curve u, or a gradient cylinder v. We label the components of the domain associated with J-holomorphic curves C_i and those labeled with gradient flow cylinders $\tilde{C}_i$. Now for each C_i we pick a point $x_i \in C_i$ and define an ordering that

$$C_i \leq C_j$$

if

$$a_n(\phi_n(x_i)) - a_n(\phi_n(x_j)) < \infty$$

and if $C_i \leq C_j$ and $C_j \leq C_i$, we say $C_i \sim C_j$. This ordering defines a level structure as in the SFT picture, then we add in the gradient flow v_j by hand at each of the levels. We note that if a gradient flow flows across multiple levels of holomorphic curves, then it will appear at these levels as a trivial cylinder. With this convention we see that then the flow time at each intermediate cascade level is the same for all Morse-Bott tori on that level (if a gradient flow needed to flow longer it would simply appear as a trivial cylinder). Then we have the SFT compactness result as desired. □

Bibliography

[1] Michèle Audin and Mihai Damian. *Morse theory and Floer homology*. Universitext. Translated from the 2010 French original by Reinie Erné. Springer, London; EDP Sciences, Les Ulis, 2014, pp. xiv+596.

[2] D. M. Austin and P. J. Braam. “Morse-Bott theory and equivariant cohomology”. In: *The Floer memorial volume*. Vol. 133. Progr. Math. Birkhäuser, Basel, 1995, pp. 123–183.

[3] Augustin Banyaga and David E. Hurtubise. “Cascades and perturbed Morse-Bott functions”. In: *Algebr. Geom. Topol.* 13.1 (2013), pp. 237–275.

[4] Paul Biran and Octav Cornea. *Quantum Structures for Lagrangian Submanifolds*. 2007. arXiv: `0708.4221 [math.SG]`.

[5] Frederic Bourgeois. “A Morse-Bott approach to contact homology”. PhD thesis. Stanford University, 2002.

[6] Frederic Bourgeois et al. “Compactness results in symplectic field theory”. In: *Geometry & Topology* 7.2 (2003), pp. 799–888.

[7] Frédéric Bourgeois and Alexandru Oancea. “Symplectic homology, autonomous Hamiltonians, and Morse-Bott moduli spaces”. In: *Duke Math. J.* 146.1 (2009), pp. 71–174.

[8] Guanheng Chen. “On cobordism maps on periodic Floer homology”. In: *Algebr. Geom. Topol.* 21.1 (2021), pp. 1–103.

[9] Keon Choi. *Combinatorial embedded contact homology for toric contact manifolds*. 2016. arXiv: `1608.07988 [math.SG]`.

[10] Keon Choi et al. “Symplectic embeddings into four-dimensional concave toric domains”. In: *J. Topol.* 7.4 (2014), pp. 1054–1076.

[11] Kai Cieliebak and Klaus Mohnke. “Symplectic hypersurfaces and transversality in Gromov-Witten theory”. In: *J. Symplectic Geom.* 5.3 (2007), pp. 281–356.

[12] Vincent Colin, Paolo Ghiggini, and Ko Honda. *Embedded contact homology and open book decompositions*. 2021. arXiv: `1008.2734 [math.SG]`.

[13] Dan Cristofaro-Gardiner. “Symplectic embeddings from concave toric domains into convex ones”. In: *J. Differential Geom.* 112.2 (2019). With an appendix by Cristofaro-Gardiner and Keon Choi, pp. 199–232.

[14] Dan Cristofaro-Gardiner, Vincent Humilière, and Sobhan Seyfaddini. *Proof of the simplicity conjecture*. 2020. URL: `https://arxiv.org/abs/2001.01792`.

[15] Dan Cristofaro-Gardiner, Michael Hutchings, and Daniel Pomerleano. "Torsion contact forms in three dimensions have two or infinitely many Reeb orbits". In: *Geom. Topol.* 23.7 (2019), pp. 3601–3645.

[16] Daniel Cristofaro-Gardiner and Michael Hutchings. "From one Reeb orbit to two". In: *J. Differential Geom.* 102.1 (2016), pp. 25–36.

[17] S. K. Donaldson. *Floer homology groups in Yang-Mills theory*. Vol. 147. Cambridge Tracts in Mathematics. With the assistance of M. Furuta and D. Kotschick. Cambridge University Press, Cambridge, 2002, pp. viii+236.

[18] David Farris. "The embedded contact homology of nontrivial circle bundles over Riemann surfaces". PhD thesis. University of California, Berkeley, 2011.

[19] Andreas Floer. "The unregularized gradient flow of the symplectic action". In: *Comm. Pure Appl. Math.* 41.6 (1988), pp. 775–813. ISSN: 0010-3640.

[20] Urs Frauenfelder. "The Arnold-Givental conjecture and moment Floer homology". In: *Int. Math. Res. Not.* 42 (2004), pp. 2179–2269.

[21] Kenji Fukaya. "Floer homology of connected sum of homology 3-spheres". In: *Topology* 35.1 (1996), pp. 89–136.

[22] Chris Gerig. *Seiberg-Witten and Gromov invariants for self-dual harmonic 2-forms.* eprint: `1809.03405`.

[23] Chris Gerig. "Taming the pseudoholomorphic beasts in $\mathbb{R} \times (S^1 \times S^2)$". In: *Geom. Topol.* 24.4 (2020), pp. 1791–1839.

[24] Jean Gutt. "Generalized Conley-Zehnder index". eng. In: *Annales de la faculté des sciences de Toulouse Mathématiques* 23.4 (2014), pp. 907–932.

[25] David Hermann. *Non-equivalence of Symplectic Capacities for Open Sets with Restricted Contact Type Boundary.* 1998. URL: `www.math.u-psud.fr/~biblio/pub/1998/abs/ppo1998_32.html`.

[26] H. Hofer, K. Wysocki, and E. Zehnder. "Properties of pseudoholomorphic curves in symplectisation. IV. Asymptotics with degeneracies". In: *Contact and symplectic geometry (Cambridge, 1994)*. Vol. 8. Publ. Newton Inst. Cambridge Univ. Press, Cambridge, 1996, pp. 78–117.

[27] Helmut Hofer, K. Wysocki, and E. Zehnder. "Properties of pseudoholomorphic curves in symplectisations. I : asymptotics". In: *Annales de l'I.H.P. Analyse non linéaire* 13.3 (1996), pp. 337–379.

[28] David E. Hurtubise. "Three approaches to Morse-Bott homology". In: *Afr. Diaspora J. Math.* 14.2 (2013), pp. 145–177.

[29] Michael Hutchings. “An index inequality for embedded pseudoholomorphic curves in symplectizations”. eng. In: *Journal of the European Mathematical Society* 004.4 (2002), pp. 313–361. URL: `http://eudml.org/doc/277717`.

[30] Michael Hutchings. “Beyond ECH capacities”. In: *Geom. Topol.* 20.2 (2016), pp. 1085–1126.

[31] Michael Hutchings. “Lecture notes on embedded contact homology”. In: *Contact and symplectic topology*. Vol. 26. Bolyai Soc. Math. Stud. János Bolyai Math. Soc., Budapest, 2014, pp. 389–484.

[32] Michael Hutchings. “Mean action and the Calabi invariant”. In: *J. Mod. Dyn.* 10 (2016), pp. 511–539.

[33] Michael Hutchings. “The embedded contact homology index revisited”. In: *New perspectives and challenges in symplectic field theory*. Vol. 49. CRM Proc. Lecture Notes. Amer. Math. Soc., Providence, RI, 2009, pp. 263–297.

[34] Michael Hutchings and Jo Nelson. *S^1-equivariant contact homology for hypertight contact forms*. 2019. eprint: `1906.03457`.

[35] Michael Hutchings and Jo Nelson. “Axiomatic S^1 Morse-Bott theory”. In: *Algebr. Geom. Topol.* 20.4 (2020), pp. 1641–1690. ISSN: 1472-2747.

[36] Michael Hutchings and Jo Nelson. “Cylindrical contact homology for dynamically convex contact forms in three dimensions”. In: *J. Symplectic Geom.* 14.4 (2016), pp. 983–1012.

[37] Michael Hutchings and Michael Sullivan. “Rounding corners of polygons and the embedded contact homology of T^3”. In: *Geom. Topol.* 10 (2006), pp. 169–266.

[38] Michael Hutchings and Michael Sullivan. “The periodic Floer homology of a Dehn twist”. In: *Algebr. Geom. Topol.* 5 (2005), pp. 301–354.

[39] Michael Hutchings and Michael Sullivan. “The periodic Floer homology of a Dehn twist”. In: *Algebr. Geom. Topol.* 5 (2005), pp. 301–354.

[40] Michael Hutchings and Clifford Henry Taubes. “Gluing pseudoholomorphic curves along branched covered cylinders. I”. In: *J. Symplectic Geom.* 5.1 (2007), pp. 43–137.

[41] Michael Hutchings and Clifford Henry Taubes. “Gluing pseudoholomorphic curves along branched covered cylinders. II”. In: *J. Symplectic Geom.* 7.1 (2009), pp. 29–133.

[42] Michael Hutchings and Clifford Henry Taubes. “Proof of the Arnold chord conjecture in three dimensions 1”. In: *Math. Res. Lett.* 18.2 (2011), pp. 295–313.

[43] Michael Hutchings and Clifford Henry Taubes. “Proof of the Arnold chord conjecture in three dimensions, II”. In: *Geom. Topol.* 17.5 (2013), pp. 2601–2688.

[44] Kei Irie. “Dense existence of periodic Reeb orbits and ECH spectral invariants”. In: *J. Mod. Dyn.* 9 (2015), pp. 357–363.

[45] Janko Latschev and Chris Wendl. "Algebraic torsion in contact manifolds". In: *Geom. Funct. Anal.* 21.5 (2011). With an appendix by Michael Hutchings, pp. 1144–1195.

[46] Eli Lebow. "Embedded Contact Homology of 2-Torus Bundles over the Circle". PhD thesis. University of California, Berkeley, 2007.

[47] Francesco Lin. "A Morse-Bott approach to monopole Floer homology and the triangulation conjecture". In: *Mem. Amer. Math. Soc.* 255.1221 (2018), pp. v+162.

[48] Dusa McDuff and Dietmar Salamon. *J-holomorphic curves and symplectic topology.* Second. Vol. 52. American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, 2012, pp. xiv+726.

[49] Dusa McDuff and Felix Schlenk. "The embedding capacity of 4-dimensional symplectic ellipsoids". In: *Ann. of Math. (2)* 175.3 (2012), pp. 1191–1282.

[50] Mihai Munteanu. "Noncontractible loops of symplectic embeddings between convex toric domains". In: *J. Symplectic Geom.* 18.4 (2020), pp. 1169–1196.

[51] Jo Nelson. "Automatic transversality in contact homology II: filtrations and computations". In: *Proc. Lond. Math. Soc. (3)* 120.6 (2020), pp. 853–917.

[52] Jo Nelson and Morgan Weiler. *Embedded contact homology of prequantization bundles.* 2020. arXiv: `2007.13883 [math.SG]`.

[53] Yong-Geun Oh and Rui Wang. "Analysis of contact Cauchy–Rieman maps II: canonical neighborhoods and exponential convergence for the Morse-Bott case". In: *Nagoya Mathematical Journal* 231 (2018), pp. 128–223.

[54] Yong-Geun Oh and Ke Zhu. *Partial collapsing degeneration of Floer trajectories and adiabatic gluing.* eprint: `1103.3525v3`.

[55] John Pardon. "Contact homology and virtual fundamental cycles". In: *J. Amer. Math. Soc.* 32.3 (2019), pp. 825–919.

[56] Jacob Rooney. *Cobordism maps in embedded contact homology.* 2020. arXiv: `1912.01048 [math.SG]`.

[57] Richard Siefring. "Relative asymptotic behavior of pseudoholomorphic half-cylinders". In: *Communications on Pure and Applied Mathematics* 61.12 (), pp. 1631–1684.

[58] Clifford Henry Taubes. "Embedded contact homology and Seiberg-Witten Floer cohomology I". In: *Geom. Topol.* 14.5 (2010), pp. 2497–2581.

[59] Clifford Henry Taubes. "Embedded contact homology and Seiberg-Witten Floer cohomology II". In: *Geom. Topol.* 14.5 (2010), pp. 2583–2720.

[60] Clifford Henry Taubes. "Embedded contact homology and Seiberg-Witten Floer cohomology III". In: *Geom. Topol.* 14.5 (2010), pp. 2721–2817.

[61] Clifford Henry Taubes. "Embedded contact homology and Seiberg-Witten Floer cohomology IV". In: *Geom. Topol.* 14.5 (2010), pp. 2819–2960.

[62] Clifford Henry Taubes. “Embedded contact homology and Seiberg-Witten Floer cohomology V”. In: *Geom. Topol.* 14.5 (2010), pp. 2961–3000.

[63] Clifford Henry Taubes. “The structure of pseudo-holomorphic subvarieties for a degenerate almost complex structure and symplectic form on $S^1 \times B^3$”. In: *Geom. Topol.* 2 (1998), pp. 221–332.

[64] Chris Wendl. “A hierarchy of local symplectic filling obstructions for contact 3-manifolds”. In: *Duke Math. J.* 162.12 (2013), pp. 2197–2283.

[65] Chris Wendl. “Automatic transversality and orbifolds of punctured holomorphic curves in dimension four”. In: *Commentarii Mathematici Helvetici* (2010), pp. 347–407.

[66] Yuan Yao. *Computing Embedded Contact Homology in Morse-Bott Settings*. arXiv: 2211.13876 [math.SG].

[67] Yuan Yao. *From Cascades to J-holomorphic Curves and Back*. 2022. DOI: 10.48550/ARXIV.2206.04334. URL: https://arxiv.org/abs/2206.04334.

[68] Yuan Yao and Ziwen Zhao. *Product and Coproduct on Fixed Point Floer Homology of Positive Dehn Twists*. 2022. URL: https://arxiv.org/abs/2205.14516.

[69] Zhengyi Zhou. *Morse-Bott Cohomology from Homological Perturbation Theory*. arXiv: 1902.06587 [math.SG].