

Contact homology and virtual fundamental cycles

John Pardon*

28 October 2015
(revised 4 February 2019)

Abstract

We give a construction of contact homology in the sense of Eliashberg–Givental–Hofer. Specifically, we construct coherent virtual fundamental cycles on the relevant compactified moduli spaces of pseudo-holomorphic curves.

The aim of this work is to provide a rigorous construction of *contact homology*, an invariant of contact manifolds and symplectic cobordisms due to Eliashberg–Givental–Hofer [Eli98, EGH00]. The contact homology of a contact manifold (Y, ξ) is defined by counting pseudo-holomorphic curves in the sense of Gromov [Gro85] in its symplectization $\mathbb{R} \times Y$. The main problem we solve in this paper is simply to give a rigorous definition of these curve counts, the essential difficulty being that the moduli spaces of such curves are usually not cut out transversally. It is therefore necessary to construct the virtual fundamental cycles of these moduli spaces (which play the same enumerative role that the ordinary fundamental cycles do for transversally cut out moduli spaces). For this construction, we use the framework developed in [Par16]. Our methods are quite general, and apply equally well to many other moduli spaces of interest.

We use a compactification of the relevant moduli spaces which is smaller than the compactification considered in [EGH00, BEHWZ03]. Roughly speaking, for curves in symplectizations $\mathbb{R} \times Y$, we do not keep track of the relative vertical positions of different components (in particular, no trivial cylinders appear). Our compactification is more convenient for proving the master equations of contact homology: the codimension one boundary strata in our compactification correspond bijectively with the desired terms in the “master equations”, whereas the compactification from [EGH00, BEHWZ03] contains additional codimension one boundary strata. If we were to use the compactification from [EGH00, BEHWZ03], we would need to additionally argue that the contribution of each such extra codimension one boundary stratum vanishes.

Remark 0.1 (Historical discussion). The theory of pseudo-holomorphic curves in closed symplectic manifolds was founded by Gromov [Gro85]. Hofer’s breakthrough work on the three-dimensional Weinstein conjecture [Hof93] introduced pseudo-holomorphic curves in symplectizations and their relationship with Reeb dynamics. The analytic theory of such curves was then further developed by Hofer–Wysocki–Zehnder [HWZ96, HWZ98, HWZ95, HWZ99,

*This research was partially conducted during the period the author served as a Clay Research Fellow. The author was also partially supported by a National Science Foundation Graduate Research Fellowship under grant number DGE-1147470.

HWZ02]. On the algebraic side, Eliashberg–Givental–Hofer [EGH00] introduced the theories of contact homology and symplectic field theory, based on counts of pseudo-holomorphic curves in symplectizations and symplectic cobordisms (assuming such counts can be defined). The key compactness results for such curves were established by Bourgeois–Eliashberg–Hofer–Wysocki–Zehnder [BEHWZ03]. Gluing techniques applicable to such pseudo-holomorphic curves have been developed by many authors, notably Taubes, Donaldson, Floer, Fukaya–Oh–Ohta–Ono, and Hofer–Wysocki–Zehnder.

A number of important special cases of contact homology and closely related invariants have been constructed rigorously using generic and/or automatic transversality techniques. Legendrian contact homology in \mathbb{R}^{2n+1} was constructed by Ekholm–Etnyre–Sullivan [EES05]. Eliashberg–Kim–Polterovich [EKP06] defined cylindrical contact homology for fiberwise star-shaped open subsets of certain pre-quantization spaces. Embedded contact homology was introduced and constructed by Hutchings and Hutchings–Taubes [Hut02, Hut09, HT07, HT09]. Cylindrical contact homology was constructed by Hutchings–Nelson [HN16] for dynamically convex contact three-manifolds and by Bao–Honda [BH18] for hypertight contact three-manifolds. Soon after the present paper was released, work of Bao–Honda [BH16] appeared, as well as work of Ishikawa [Ish18].

Remark 0.2 (Virtual moduli cycle techniques). The technique of patching together local finite-dimensional reductions to construct virtual fundamental cycles has been applied to moduli spaces of pseudo-holomorphic curves by many authors, including Fukaya–Ono [FO99] (Kuranishi structures), Li–Tian [LT98a], Liu–Tian [LT98b], Ruan [Rua99], Fukaya–Oh–Ohta–Ono [FOOO09a, FOOO09b, FOOO12, FOOO16b, FOOO15, FOOO16a, Fuk17, FOOO17, FOOO18a, FOOO18b], McDuff–Wehrheim [MW17b, MW18, MW17a, McD19] (Kuranishi atlases), Joyce [Joy15, Joy14, Joy12] (Kuranishi spaces and d-orbifolds), and [Par16] (implicit atlases). Features of the framework from [Par16] include that (1) it requires only topological (as opposed to smooth) gluing theorems as input and (2) it applies to moduli spaces with rather general corner structures which we term “cell-like” (more general than manifold-with-corners).

More recently, the theory of polyfolds developed by Hofer–Wysocki–Zehnder [HWZ07, HWZ09b, HWZ09a, HWZ10a, HWZ10b, HWZ14a, HWZ14b, HWZ17, FH18b, FH18c, FH18a] promises to provide a robust new infinite-dimensional context in which all reasonable moduli spaces of pseudo-holomorphic curves may be perturbed “abstractly” to obtain transversality.

Any one of the above theories, once sufficiently developed, could be used to prove the main results of this paper. Although these theories vary in their approach to the myriad of technical issues involved, they are expected to give rise to completely equivalent virtual fundamental cycles.

Acknowledgements: The author thanks Yasha Eliashberg for introducing the author to this problem and for many useful conversations. The author thanks the referee for their careful reading and extensive comments which have improved this paper significantly. The author is also grateful for useful discussions with Mohammed Abouzaid, Frédéric Bourgeois, Roger Casals, Vincent Colin, Tobias Ekholm, Kenji Fukaya, Eleny Ionel, Patrick Massot, Rafe Mazzeo, Peter Ozsváth, Mohan Swaminathan, and Chris Wendl. A preliminary version of this paper appeared as part of the author’s 2015 Ph.D. thesis at Stanford University supervised by Ya. Eliashberg.

1 Statement of results

We now state our main results while simultaneously reviewing the definition of contact homology as given in [Eli98, EGH00, Eli07]. The definition of contact homology involves counting pseudo-holomorphic curves in various different (but closely related) settings, which we distinguish with roman numerals (I), (II), (III), (IV). *We strongly recommend that the reader restrict their attention to (I) on a first reading.* The generalization from (I) to (II), (III), (IV) is mostly straightforward, and may be saved for a second reading.

1.1 Conventions

Everything is in the smooth category unless stated otherwise.

To keep track of signs, everything is $\mathbb{Z}/2$ -graded, and \otimes always denotes the *super tensor product*, meaning that the isomorphism $A \otimes B \xrightarrow{\sim} B \otimes A$ is given by $a \otimes b \mapsto (-1)^{|a||b|} b \otimes a$ and that $(f \otimes g)(a \otimes b) := (-1)^{|g||a|} f(a) \otimes g(b)$. Complexes are $(\mathbb{Z}, \mathbb{Z}/2)$ -bigraded, and differentials are always odd. For specificity, let us declare that $\text{Hom}(A, B) \otimes A \rightarrow B$ be given by $f \otimes a \mapsto f(a)$, though it won't really matter for our arguments.

An *orientation line* is a $\mathbb{Z}/2$ -graded free \mathbb{Z} -module of rank one. Note that for an odd orientation line \mathfrak{o} , there is no nonzero symmetric pairing $\mathfrak{o} \otimes \mathfrak{o} \rightarrow \mathbb{Z}$, so we must be careful to distinguish between \mathfrak{o} and its dual \mathfrak{o}^\vee .

1.2 (I) The differential

Let Y^{2n-1} be a closed manifold, and let λ be a *contact form* on Y , namely a 1-form such that $\lambda \wedge (d\lambda)^{n-1}$ is non-vanishing, denoting by $\xi := \ker \lambda$ the induced co-oriented contact structure. Contact homology can be regarded as an attempt to define the S^1 -equivariant Morse–Floer homology of the loop space $C^\infty(S^1, Y)$ with respect to the action functional $a(\gamma) := \int_{S^1} \gamma^* \lambda$.

Denote by R_λ the *Reeb vector field* associated to λ , defined by the properties $\lambda(R_\lambda) = 1$ and $d\lambda(R_\lambda, \cdot) = 0$. Denote by $\mathcal{P} = \mathcal{P}(Y, \lambda)$ the collection of (unparameterized) *Reeb orbits*, namely closed trajectories of R_λ , not necessarily embedded. A Reeb orbit is called *non-degenerate* iff the associated linearized Poincaré return map $\xi_p \rightarrow \xi_p$ has no fixed vector. A contact form is called non-degenerate iff all its Reeb orbits are non-degenerate, which implies that there are at most finitely many Reeb orbits below any given action threshold. Suppose now that λ is non-degenerate.

There is a natural partition $\mathcal{P} = \mathcal{P}_{\text{good}} \sqcup \mathcal{P}_{\text{bad}}$, and for each good Reeb orbit $\gamma \in \mathcal{P}_{\text{good}}$, there is an associated orientation line \mathfrak{o}_γ with parity $|\gamma| = \text{sign}(\det(I - A_\gamma)) \in \{\pm 1\} = \mathbb{Z}/2$, where A_γ denotes the linearized Poincaré return map of γ acting on ξ (see §2.13). We set $\mathfrak{o}_\Gamma := \bigotimes_{\gamma \in \Gamma} \mathfrak{o}_\gamma$ and $|\Gamma| := \sum_{\gamma \in \Gamma} |\gamma|$ for any finite set of good Reeb orbits $\Gamma \rightarrow \mathcal{P}_{\text{good}}$. For a given Reeb orbit $\gamma \in \mathcal{P}$, let $d_\gamma \in \mathbb{Z}_{\geq 1}$ denote its covering multiplicity.

Let

$$CC_\bullet(Y, \xi)_\lambda := \bigoplus_{n \geq 0} \text{Sym}_{\mathbb{Q}}^n \left(\bigoplus_{\gamma \in \mathcal{P}_{\text{good}}} \mathfrak{o}_\gamma \right) \quad (1.1)$$

denote the free supercommutative (meaning $ab = (-1)^{|a||b|}ba$) unital $\mathbb{Z}/2$ -graded \mathbb{Q} -algebra generated by \mathfrak{o}_γ for $\gamma \in \mathcal{P}_{\text{good}}$. This is the contact homology chain complex, which we now

equip with a differential. To define this differential, fix an almost complex structure $J : \xi \rightarrow \xi$ which is *compatible* with $d\lambda$, meaning that $d\lambda(\cdot, J\cdot)$ is a positive definite symmetric pairing on ξ .

Let $\hat{Y} := \mathbb{R} \times Y$ (with coordinate $s \in \mathbb{R}$) denote the *symplectization*¹ of Y , which is equipped with the Liouville form $\hat{\lambda} := e^s \lambda$. Now J induces an \mathbb{R} -invariant almost complex structure \hat{J} on \hat{Y} defined by the property that $\hat{J}(\partial_s) = R_\lambda$ and $\hat{J}|_\xi = J$. Given a Reeb orbit $\gamma^+ \in \mathcal{P}$ and a finite set of Reeb orbits $\Gamma^- \rightarrow \mathcal{P}$, let

$$\pi_2(Y, \gamma^+ \sqcup \Gamma^-) := [(S, \partial S), (Y, \gamma^+ \sqcup \Gamma^-)] / \text{Diff}(S, \partial S), \quad (1.2)$$

where S is any compact connected oriented surface of genus zero with boundary, equipped with a homeomorphism between ∂S and $\gamma^+ \sqcup \Gamma^-$ (preserving orientation on γ^+ and reversing orientation on Γ^-), and $\text{Diff}(S, \partial S)$ denotes diffeomorphisms of S fixing ∂S pointwise.

Let $\overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J$ denote the compactified moduli space of connected \hat{J} -holomorphic curves of genus zero in \hat{Y} modulo \mathbb{R} -translation, with one positive puncture asymptotic to γ^+ and negative punctures asymptotic to Γ^- , in the homotopy class β , along with asymptotic markers on the domain mapping to fixed basepoints on γ^+ and Γ^- (see §§2.1–2.9). We denote by $\mu(\gamma^+, \Gamma^-; \beta) \in \mathbb{Z}$ the *index* of this moduli problem; we have $\mu(\gamma^+, \Gamma^-; \beta) \equiv |\gamma^+| - |\Gamma^-| \in \mathbb{Z}/2$ (see §2.12).

It is shown in [BEHWZ03] that $\bigsqcup_{(\Gamma^-, \beta)} \overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J$ is compact for any fixed γ^+ (see §2.10).

We say that $\overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J$ is *regular* (or *cut out transversally*, or just *transverse*) iff the relevant linearized operator (which takes into account both variations in the map and variations in the almost complex structure of the domain) is everywhere surjective (see §2.11). In this case, $\overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J$ is a manifold with corners of dimension $\mu(\gamma^+, \Gamma^-; \beta) - 1$, whose orientation local system is naturally isomorphic to $\mathfrak{o}_{\gamma^+} \otimes \mathfrak{o}_{\Gamma^-}^\vee \otimes \mathfrak{o}_{\mathbb{R}}^\vee$ (see §2.14); there are no orbifold points due to our use of asymptotic markers and to the fact that the symplectization is exact so there can be no nodes. In particular, if $\mu(\gamma^+, \Gamma^-; \beta) = 1$ and $\overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J$ is cut out transversally, the moduli count

$$\#\overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J \in \mathfrak{o}_{\gamma^+}^\vee \otimes \mathfrak{o}_{\Gamma^-} \otimes \mathfrak{o}_{\mathbb{R}} \quad (1.3)$$

is well-defined.

We would now like to define a differential

$$d_J : CC_\bullet(Y, \xi)_\lambda \rightarrow CC_{\bullet-1}(Y, \xi)_\lambda \quad (1.4)$$

which satisfies the Leibniz rule ($d(1) = 0$ and $d(ab) = da \cdot b + (-1)^{|a|} a \cdot db$) and is defined by the property that it acts on \mathfrak{o}_{γ^+} by pairing with

$$\sum_{\substack{\Gamma^- \rightarrow \mathcal{P}_{\text{good}} \\ \mu(\gamma^+, \Gamma^-; \beta) = 1}} \frac{1}{|\text{Aut}|} \cdot \#\overline{\mathcal{M}}_1(\gamma^+, \Gamma^-; \beta)_J \cdot \Gamma^- \quad (1.5)$$

¹More intrinsically, the symplectization of a co-oriented contact manifold (Y, ξ) is defined as the total space of the bundle of 1-forms with kernel ξ , namely $\hat{Y} := \ker(T^*Y \rightarrow \xi^*)_+$. The restriction of the tautological Liouville 1-form on T^*Y is a Liouville 1-form $\hat{\lambda}$ on \hat{Y} ; the associated Liouville vector field on \hat{Y} generates an \mathbb{R} -action on \hat{Y} which is simply scaling by e^s . A choice of contact form λ for ξ induces an identification of $(\hat{Y}, \hat{\lambda})$ with $(\mathbb{R} \times Y, e^s \lambda)$.

and contracting on the left with a chosen orientation of \mathbb{R} (where Aut denotes a certain group of automorphisms of the triple $(\gamma^+, \Gamma^-; \beta)$ acting as the identity on γ^+). When the moduli spaces $\overline{\mathcal{M}}_I(\gamma^+, \Gamma^-; \beta)_J$ of index ≤ 2 are cut out transversally, this definition of d_J makes sense, and consideration of the fact that $\#\partial\overline{\mathcal{M}}_I(\gamma^+, \Gamma^-; \beta)_J = 0$ for index two moduli spaces (since they are compact 1-manifolds) shows that d_J squares to zero.

Our main result, Theorem 1.1, allows the above construction of the differential to go through without any transversality assumptions. More precisely, it provides a non-empty set $\Theta_I = \Theta_I(Y, \lambda, J)$, an element of which may be thought of as a specification of “perturbation data” for the moduli spaces $\overline{\mathcal{M}}_I(\gamma^+, \Gamma^-; \beta)_J$, and for each $\theta \in \Theta_I$ it defines *rational* “virtual moduli counts”

$$\#\overline{\mathcal{M}}_I(\gamma^+, \Gamma^-; \beta)_{J,\theta}^{\text{vir}} \in \mathfrak{o}_{\gamma^+}^{\vee} \otimes \mathfrak{o}_{\Gamma^-} \otimes \mathfrak{o}_{\mathbb{R}} \otimes \mathbb{Q}. \quad (1.6)$$

These virtual moduli counts furthermore satisfy $\#\partial\overline{\mathcal{M}}_I(\gamma^+, \Gamma^-; \beta)_{J,\theta}^{\text{vir}} = 0$ and coincide with the ordinary moduli counts when the moduli spaces are transverse. We thus obtain a differential

$$d_{J,\theta} : CC_{\bullet}(Y, \xi)_{\lambda} \rightarrow CC_{\bullet-1}(Y, \xi)_{\lambda}, \quad (1.7)$$

and we denote the resulting homology by $CH_{\bullet}(Y, \xi)_{\lambda, J, \theta}$.

1.3 (II) The cobordism map

Let \hat{X}^{2n} be a manifold and let $\hat{\omega}$ be a symplectic form on \hat{X} , namely a closed 2-form such that $\hat{\omega}^n$ is non-vanishing. Let (Y^{\pm}, λ^{\pm}) be closed manifolds equipped with contact forms, and let

$$([N, \infty) \times Y^+, d\hat{\lambda}^+) \rightarrow (\hat{X}, \hat{\omega}), \quad (1.8)$$

$$((-\infty, -N] \times Y^-, d\hat{\lambda}^-) \rightarrow (\hat{X}, \hat{\omega}), \quad (1.9)$$

(any large $N < \infty$) be diffeomorphisms onto their (disjoint) images, proper, and such that together they cover a neighborhood of infinity. A *symplectic cobordism* is a symplectic manifold $(\hat{X}, \hat{\omega})$ equipped with such markings (1.8)–(1.9), and it is called *exact* when there is a globally defined primitive $\hat{\lambda}$ for $\hat{\omega}$ coinciding with $\hat{\lambda}^{\pm}$ in the ends.

Fix an almost complex structure $\hat{J} : T\hat{X} \rightarrow T\hat{X}$ which is *tamed* by $\hat{\omega}$, meaning $\hat{\omega}(v, \hat{J}v) > 0$ for nonzero $v \in T\hat{X}$, such that outside a compact set, \hat{J} coincides with some \hat{J}^{\pm} for $J^{\pm} : \xi^{\pm} \rightarrow \xi^{\pm}$. Also suppose that λ^{\pm} are non-degenerate.

Let $\overline{\mathcal{M}}_{II}(\gamma^+, \Gamma^-; \beta)_j$ denote the compactified moduli space of connected \hat{J} -holomorphic curves of genus zero in \hat{X} from $\gamma^+ \in \mathcal{P}^+ = \mathcal{P}(Y^+, \lambda^+)$ to $\Gamma^- \rightarrow \mathcal{P}^- = \mathcal{P}(Y^-, \lambda^-)$ in the homotopy class $\beta \in \pi_2(\hat{X}, \gamma^+ \sqcup \Gamma^-)$.

It is shown in [BEHWZ03] that $\bigsqcup_{\langle \hat{\omega}, \beta \rangle < N} \overline{\mathcal{M}}_{II}(\gamma^+, \Gamma^-; \beta)_j$ is compact for any fixed γ^+ , Γ^- , and $N < \infty$. The condition $\langle \hat{\omega}, \beta \rangle < N$ may be interpreted with respect to any fixed null-homology of $[\gamma^+] - [\Gamma^-]$. Note that when the symplectic cobordism is exact, it follows that $\bigsqcup_{(\Gamma^-, \beta)} \overline{\mathcal{M}}_{II}(\gamma^+, \Gamma^-; \beta)_j$ is compact for any fixed γ^+ .

Theorem 1.1 provides virtual moduli counts for the moduli spaces $\overline{\mathcal{M}}_{II}$ in the following form. There is a set $\Theta_{II} = \Theta_{II}(\hat{X}, \hat{\omega}, \hat{J}, (1.8), (1.9))$ together with a surjective forgetful map $\Theta_{II} \twoheadrightarrow \Theta_I^+ \times \Theta_I^-$, where $\Theta_I^{\pm} = \Theta_I(Y^{\pm}, \lambda^{\pm}, J^{\pm})$, along with virtual moduli counts

$$\#\overline{\mathcal{M}}_{II}(\gamma^+, \Gamma^-, \beta)_{j,\theta}^{\text{vir}} \in \mathfrak{o}_{\gamma^+}^{\vee} \otimes \mathfrak{o}_{\Gamma^-} \otimes \mathbb{Q} \quad (1.10)$$

for $\theta \in \Theta_{\text{II}}$ and $\mu(\gamma^+, \Gamma^-, \beta) = 0$.

Now for exact symplectic cobordisms, we may define a unital \mathbb{Q} -algebra map

$$\Phi(\hat{X}, \hat{\lambda})_{\hat{J}, \theta} : CC_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \rightarrow CC_{\bullet}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \quad (1.11)$$

for any $\theta \in \Theta_{\text{II}}$ mapping to $(\theta^+, \theta^-) \in \Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-$, by pairing with the virtual moduli counts as before (exactness ensures the relevant sum is finite). This is a chain map by virtue of the fact that $\#\partial\overline{\mathcal{M}}_{\text{II}} = 0$.

1.4 (III) The deformation homotopy

Let $(\hat{X}, \hat{\omega}^t)_{t \in [0,1]}$ be a one-parameter family of symplectic cobordisms, fixed near infinity (meaning (1.8)–(1.9) are independent of t). Note that the *a priori* more general setup where (1.8)–(1.9) are allowed to vary with t is easily reduced to the case of being fixed near infinity, simply by conjugating by an appropriate family of diffeomorphisms $\varphi^t : \hat{X} \rightarrow \hat{X}$.

Fix a family \hat{J}^t of almost complex structures as above, agreeing with \hat{J}^{\pm} outside of a compact subset independent of t .

We consider moduli spaces $\overline{\mathcal{M}}_{\text{III}}(\{\gamma_i^+, \Gamma_i^-, \beta_i\}_{i \in I})_{\hat{J}^t}$ parameterizing a choice of $t \in [0, 1]$ together with a \hat{J}^t -holomorphic curve from γ_i^+ to Γ_i^- in homotopy class β_i for every $i \in I$. For such moduli spaces, there are two natural notions of transversality: a point is called *regular* iff the associated linearized operator is surjective as before, and a point is called *weakly regular* iff the associated linearized operator becomes surjective once we also take into account variations in $t \in (0, 1)$ (but *not* at the endpoints of this interval).

It is shown in [BEHWZ03] that $\bigsqcup_{\sup_t \langle \hat{\omega}^t, \beta \rangle < N} \overline{\mathcal{M}}_{\text{III}}(\gamma^+, \Gamma^-, \beta)_{\hat{J}^t}$ is compact for any fixed γ^+ , Γ^- , and $N < \infty$.

Theorem 1.1 provides virtual moduli counts

$$\#\overline{\mathcal{M}}_{\text{III}}(\{\gamma_i^+, \Gamma_i^-, \beta_i\}_{i \in I})_{\hat{J}^t, \theta}^{\text{vir}} \in \bigotimes_{i \in I} \mathfrak{o}_{\gamma_i^+}^{\vee} \otimes \mathfrak{o}_{\Gamma_i^-} \otimes \mathfrak{o}_{[0,1]}^{\vee} \otimes \mathbb{Q} \quad (1.12)$$

for $\theta \in \Theta_{\text{III}} := \Theta_{\text{III}}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t \in [0,1]})$ with a surjective map $\Theta_{\text{III}} \twoheadrightarrow \Theta_{\text{II}}^{t=0} \times_{\Theta^+ \times \Theta^-} \Theta_{\text{II}}^{t=1}$.

For families of exact symplectic cobordisms, pairing with these virtual moduli counts (and an orientation of $[0, 1]$) defines a \mathbb{Q} -linear map

$$K(\hat{X}, \hat{\lambda}^t)_{\hat{J}^t, \theta} : CC_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \rightarrow CC_{\bullet+1}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-}. \quad (1.13)$$

Note that, whereas in cases (I) and (II), we defined the differential and the cobordism map on algebra generators and then extended using the multiplicative structure, here in case (III) we define the homotopy on each monomial separately.

Consideration of $\#\partial\overline{\mathcal{M}}_{\text{III}} = 0$ shows that $K(\hat{X}, \hat{\lambda}^t)_{\hat{J}^t, \theta}$ is a chain homotopy between $\Phi(\hat{X}, \hat{\lambda}^0)_{\hat{J}^0, \theta^0}$ and $\Phi(\hat{X}, \hat{\lambda}^1)_{\hat{J}^1, \theta^1}$, implying that the induced maps on homology

$$CH_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \xrightarrow[\Phi(\hat{X}, \hat{\lambda}^1)_{\hat{J}^1, \theta^1}]{\Phi(\hat{X}, \hat{\lambda}^0)_{\hat{J}^0, \theta^0}} CH_{\bullet}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \quad (1.14)$$

coincide.

1.5 (IV) The composition homotopy

Let $(\hat{X}^{01}, \hat{\omega}^{01})$ be a symplectic cobordism with positive end (Y^0, λ^0) and negative end (Y^1, λ^1) , and let $(\hat{X}^{12}, \hat{\omega}^{12})$ be a symplectic cobordism with positive end (Y^1, λ^1) and negative end (Y^2, λ^2) . Given any sufficiently large $t < \infty$, we can form a symplectic cobordism $(\hat{X}^{02,t}, \hat{\omega}^{02,t})$ by truncating the negative end of \hat{X}^{01} to $(-t, 0] \times Y_1$, truncating the positive end of \hat{X}^{12} to $[0, t] \times Y_1$, and identifying these truncated ends by translation by t . Under this identification, the forms $\hat{\omega}^{01}$ and $\hat{\omega}^{12}$ only agree up to a scaling factor of e^t , so their descent $\hat{\omega}^{02,t}$ to $\hat{X}^{02,t}$ is only well-defined up to scale; similarly the natural end markings (1.8)–(1.9) of $\hat{X}^{02,t}$ only respect symplectic forms up to scale (these ambiguities are never problematic, however, so they will be ignored from now on). Let $(\hat{X}^{02,t}, \hat{\omega}^{02,t})_{t \in [0, \infty)}$ denote this one-parameter family of symplectic cobordisms with positive end (Y^0, λ^0) and negative end (Y^2, λ^2) .

Fix almost complex structures \hat{J}^{01} and \hat{J}^{12} on \hat{X}^{01} and \hat{X}^{12} as before, agreeing with fixed $\hat{J}^0, \hat{J}^1, \hat{J}^2$ near infinity. These descend naturally to $\hat{J}^{02,t}$ on $\hat{X}^{02,t}$ for sufficiently large $t < \infty$, and we fix an extension $\hat{J}^{02,t}$ to all $t \in [0, \infty)$, agreeing with \hat{J}^0 and \hat{J}^2 near infinity.

We consider moduli spaces $\overline{\mathcal{M}}_{\text{IV}}(\{\gamma_i^+, \Gamma_i^-; \beta_i\}_{i \in I})_{j^{02,t}}$ parameterizing a choice of $t \in [0, \infty]$ together with a $\hat{J}^{02,t}$ -holomorphic curve in $\hat{X}^{02,t}$ from γ_i^+ to Γ_i^- in homotopy class β_i for every $i \in I$.

Theorem 1.1 provides virtual moduli counts

$$\#\overline{\mathcal{M}}_{\text{IV}}(\{\gamma_i^+, \Gamma_i^-; \beta_i\}_{i \in I})_{j^{02,t}}^{\text{vir}} \in \bigotimes_{i \in I} \mathfrak{o}_{\gamma_i^+}^{\vee} \otimes \mathfrak{o}_{\Gamma_i^-} \otimes \mathfrak{o}_{[0, \infty]}^{\vee} \otimes \mathbb{Q} \quad (1.15)$$

for $\theta \in \Theta_{\text{IV}}$ with a surjective map $\Theta_{\text{IV}} \rightarrow \Theta_{\text{II}}^{02} \times_{\Theta_1^0 \times \Theta_1^2} (\Theta_{\text{II}}^{01} \times_{\Theta_1^1} \Theta_{\text{II}}^{12})$.

In the exact setting, pairing with these virtual moduli counts and using $\#\partial\overline{\mathcal{M}}_{\text{IV}} = 0$ produces a chain homotopy which shows that the following diagram commutes:

$$\begin{array}{ccc} & CH_{\bullet}(Y^1, \xi^1)_{\lambda^1, J^1, \theta^1} & \\ \Phi(\hat{X}^{01}, \hat{\lambda}^{01})_{j^{01}, \theta^{01}} \nearrow & & \searrow \Phi(\hat{X}^{12}, \hat{\lambda}^{12})_{j^{12}, \theta^{12}} \\ CH_{\bullet}(Y^0, \xi^0)_{\lambda^0, J^0, \theta^0} & \xrightarrow{\Phi(\hat{X}^{02}, \hat{\lambda}^{02})_{j^{02}, \theta^{02}}} & CH_{\bullet}(Y^2, \xi^2)_{\lambda^2, J^2, \theta^2} \end{array} \quad (1.16)$$

for any $(\theta^{0,1,2}, \theta^{01,12,02}) \in \Theta_{\text{II}}^{02} \times_{\Theta_1^0 \times \Theta_1^2} (\Theta_{\text{II}}^{01} \times_{\Theta_1^1} \Theta_{\text{II}}^{12})$.

1.6 Main result

We now give a precise statement of our main result, Theorem 1.1, which has already been alluded to in §§1.2–1.5 above. This result takes as input a datum \mathcal{D} (as in any of Setups I–IV below) and produces a set $\Theta(\mathcal{D})$ together with virtual moduli counts $\#\overline{\mathcal{M}}_{\theta}^{\text{vir}}$ for $\theta \in \Theta(\mathcal{D})$ satisfying certain properties.

Setup I consists of a closed manifold Y equipped with a non-degenerate contact form λ and a $d\lambda$ -compatible almost complex structure $J : \xi \rightarrow \xi$.

Setup II consists of a symplectic cobordism $(\hat{X}, \hat{\omega})$ with positive end (Y^+, λ^+) and negative end (Y^-, λ^-) , together with an $\hat{\omega}$ -tame almost complex structure $\hat{J} : T\hat{X} \rightarrow T\hat{X}$ agreeing with \hat{J}^{\pm} outside a compact set, where $(Y^{\pm}, \lambda^{\pm}, J^{\pm})$ are as in Setup I.

Setup III consists of a one-parameter family of symplectic cobordisms $(\hat{X}, \hat{\omega}^t)_{t \in [0,1]}$ (fixed near infinity) with positive/negative ends (Y^\pm, λ^\pm) , together with $\hat{\omega}^t$ -tame almost complex structures \hat{J}^t agreeing with \hat{J}^\pm outside a compact set, where $(Y^\pm, \lambda^\pm, J^\pm)$ are as in Setup I.

Setup IV consists of a one-parameter family of symplectic cobordisms $(\hat{X}^{02,t}, \hat{\omega}^{02,t})_{t \in [0,\infty)}$ with ends (Y^0, λ^0) and (Y^2, λ^2) , which for sufficiently large t agrees with the t -gluing of a symplectic cobordism $(\hat{X}^{01}, \hat{\omega}^{01})$ with ends (Y^0, λ^0) and (Y^1, λ^1) and a symplectic cobordism $(\hat{X}^{12}, \hat{\omega}^{12})$ with ends (Y^1, λ^1) and (Y^2, λ^2) , together with an appropriately compatible family of almost complex structures $\hat{J}^{02,t}$, with (Y^i, λ^i, J^i) as in Setup I.

Theorem 1.1. *Let \mathcal{D} be a datum as in any of Setups I–IV. There exists a set $\Theta(\mathcal{D})$ along with rational virtual moduli counts $\#\overline{\mathcal{M}}_\theta^{\text{vir}}$ satisfying the following properties:*

- (i) $\Theta(\mathcal{D})$ and $\#\overline{\mathcal{M}}_\theta^{\text{vir}}$ are functorial in \mathcal{D} , meaning that an isomorphism $i : \mathcal{D} \rightarrow \mathcal{D}'$ induces an isomorphism $i_* : \Theta(\mathcal{D}) \xrightarrow{\sim} \Theta(\mathcal{D}')$ such that $\text{id}_* = \text{id}$, $(i \circ j)_* = i_* \circ j_*$, and $\#\overline{\mathcal{M}}_\theta^{\text{vir}} = \#\overline{\mathcal{M}}_{i_*\theta}^{\text{vir}}$.
- (ii) There are natural surjective forgetful maps

$$\Theta_{\text{I}} \rightarrow *, \quad (1.17)$$

$$\Theta_{\text{II}} \rightarrow \Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-, \quad (1.18)$$

$$\Theta_{\text{III}} \rightarrow \Theta_{\text{II}}^0 \times_{\Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-} \Theta_{\text{II}}^1, \quad (1.19)$$

$$\Theta_{\text{IV}} \rightarrow \Theta_{\text{II}}^{02} \times_{\Theta_{\text{I}}^0 \times \Theta_{\text{I}}^2} (\Theta_{\text{II}}^{01} \times_{\Theta_{\text{I}}^1} \Theta_{\text{II}}^{12}), \quad (1.20)$$

covering the natural forgetful maps of data \mathcal{D} .

- (iii) $\#\overline{\mathcal{M}}_\theta^{\text{vir}}$ is nonzero only for moduli spaces of virtual dimension zero.
- (iv) $\#\overline{\mathcal{M}}_\theta^{\text{vir}} = \#\overline{\mathcal{M}}$ for (weakly) regular moduli spaces of dimension zero. In particular, if a moduli space is empty then its virtual count is zero.
- (v) The virtual moduli counts satisfy the “master equation”

$$\#\partial\overline{\mathcal{M}}_\theta^{\text{vir}} = 0, \quad (1.21)$$

where the left hand side denotes the sum over all codimension one boundary strata of the relevant products of virtual moduli counts and inverse covering multiplicities of intermediate orbits (this sum is finite by compactness).

1.7 The contact homology functor

Contact homology can be viewed as a functor (1.22), whose (mostly formal) construction from Theorem 1.1 we now describe.

Let $(\mathbf{Contact}, \mathbf{Exact})_n$ denote the category whose objects are closed co-oriented contact manifolds (Y^{2n-1}, ξ) and whose morphisms are deformation classes of exact symplectic cobordisms $(\hat{X}^{2n}, \hat{\lambda})$ between contact manifolds. Let $\mathfrak{Ring}_{\mathbb{Q}}^{\mathbb{Z}/2}$ denote the category whose objects are supercommutative $\mathbb{Z}/2$ -graded unital \mathbb{Q} -algebras and whose morphisms are graded unital \mathbb{Q} -algebra homomorphisms. Contact homology is a symmetric monoidal functor

$$CH_\bullet : (\mathbf{Contact}, \mathbf{Exact})_n^{\sqcup} \rightarrow (\mathfrak{Ring}_{\mathbb{Q}}^{\mathbb{Z}/2})^\otimes. \quad (1.22)$$

The symmetric monoidal structure on $(\mathbf{Contact}, \mathbf{E}act)_n$ is disjoint union \sqcup , and the symmetric monoidal structure on $\mathfrak{Ring}_{\mathbb{Q}}^{\mathbb{Z}/2}$ is the super tensor product \otimes ($A \otimes B$ is endowed with the multiplication $(a \otimes b)(a' \otimes b') := (-1)^{|a'| |b|} aa' \otimes bb'$, and the isomorphism $A \otimes B \xrightarrow{\sim} B \otimes A$ is given by $a \otimes b \mapsto (-1)^{|a| |b|} b \otimes a$).

Concretely, to define (1.22) we should specify:

- For every co-oriented contact manifold (Y, ξ) , a supercommutative $\mathbb{Z}/2$ -graded unital \mathbb{Q} -algebra $CH_{\bullet}(Y, \xi)$.
- For every exact symplectic cobordism $(\hat{X}, \hat{\lambda})$ from (Y^+, ξ^+) to (Y^-, ξ^-) , a graded unital \mathbb{Q} -algebra map $\Phi(\hat{X}, \hat{\lambda}) : CH_{\bullet}(Y^+, \xi^+) \rightarrow CH_{\bullet}(Y^-, \xi^-)$.
- Isomorphisms $CH_{\bullet}(Y, \xi) \otimes CH_{\bullet}(Y', \xi') = CH_{\bullet}(Y \sqcup Y', \xi \sqcup \xi')$.

such that:

- The morphism $\Phi(\hat{X}, \hat{\lambda})$ depends only on the deformation class of $(\hat{X}, \hat{\lambda})$.
- The morphism associated to the identity cobordism is the identity map.
- The morphism $\Phi(\hat{X}^{02}, \hat{\lambda}^{02})$ associated to a composition of exact symplectic cobordisms $\hat{X}^{02} = \hat{X}^{01} \# \hat{X}^{12}$ coincides with the composition $\Phi(\hat{X}^{12}, \hat{\lambda}^{12}) \circ \Phi(\hat{X}^{01}, \hat{\lambda}^{01})$.
- The isomorphisms $CH_{\bullet}(Y, \xi) \otimes CH_{\bullet}(Y', \xi') = CH_{\bullet}(Y \sqcup Y', \xi \sqcup \xi')$ are commutative, associative, and compatible with the cobordism maps.

The construction is as follows.

Theorem 1.1 applied to Setup I provides a supercommutative $\mathbb{Z}/2$ -graded unital \mathbb{Q} -algebra

$$CH_{\bullet}(Y, \xi)_{\lambda, J, \theta} \tag{1.23}$$

for any co-oriented contact manifold (Y, ξ) with non-degenerate contact form λ , admissible almost complex structure J , and $\theta \in \Theta_I(Y, \lambda, J)$.

Theorem 1.1 applied to Setup II provides a graded unital \mathbb{Q} -algebra map

$$CH_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \xrightarrow{\Phi(\hat{X}, \hat{\lambda})_{\hat{J}, \hat{\theta}}} CH_{\bullet}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \tag{1.24}$$

for any exact symplectic cobordism $(\hat{X}, \hat{\lambda})$ with λ^{\pm} non-degenerate, admissible almost complex structure \hat{J} coinciding with \hat{J}^{\pm} near infinity, and $\hat{\theta} \in \Theta_{II}(\hat{X}, \hat{\lambda}, \hat{J})$ mapping to $\theta^{\pm} \in \Theta_I^{\pm}$.

Theorem 1.1 applied to Setup III shows that the following two maps coincide:

$$CH_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \xrightarrow[\Phi(\hat{X}, \hat{\lambda}^1)_{\hat{J}^1, \hat{\theta}^1}]{\Phi(\hat{X}, \hat{\lambda}^0)_{\hat{J}^0, \hat{\theta}^0}} CH_{\bullet}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-}. \tag{1.25}$$

Note that this immediately implies that $\Phi(\hat{X}, \hat{\lambda})_{\hat{J}, \hat{\theta}}$ is independent of \hat{J} and $\hat{\theta}$, and depends only on the deformation class of $(\hat{X}, \hat{\lambda})$. Thus we may rewrite (1.24) as

$$CH_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \xrightarrow{\Phi(\hat{X}, \hat{\lambda})} CH_{\bullet}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-}. \tag{1.26}$$

Theorem 1.1 applied to Setup IV shows that the following diagram commutes:

$$\begin{array}{ccc}
& & CH_{\bullet}(Y^1, \xi^1)_{\lambda^1, J^1, \theta^1} & & \\
& \nearrow^{\Phi(\hat{X}^{01}, \hat{\lambda}^{01})} & & \searrow_{\Phi(\hat{X}^{12}, \hat{\lambda}^{12})} & \\
CH_{\bullet}(Y^0, \xi^0)_{\lambda^0, J^0, \theta^0} & \xrightarrow{\Phi(\hat{X}^{02}, \hat{\lambda}^{02})} & & \xrightarrow{\Phi(\hat{X}^{12}, \hat{\lambda}^{12})} & CH_{\bullet}(Y^2, \xi^2)_{\lambda^2, J^2, \theta^2}.
\end{array} \tag{1.27}$$

Lemma 1.2. *Let (Y, ξ) be a co-oriented contact manifold with two non-degenerate contact forms λ^+, λ^- . Let $(\hat{X}, \hat{\lambda})$ denote the trivial exact symplectic cobordism with positive end (Y, λ^+) and negative end (Y, λ^-) (namely, \hat{X} is simply the symplectization \hat{Y} marked appropriately). The map*

$$CH_{\bullet}(Y, \xi)_{\lambda^+, J^+, \theta^+} \xrightarrow{\Phi(\hat{X}, \hat{\lambda})} CH_{\bullet}(Y, \xi)_{\lambda^-, J^-, \theta^-} \tag{1.28}$$

is an isomorphism for any J^{\pm} and θ^{\pm} .

Proof. In view of the commutativity of (1.27), it suffices to treat the case $\lambda^+ = \lambda^- = \lambda$ and $J^+ = J^- = J$.

Choose the \mathbb{R} -invariant almost complex structure $\hat{J} = \hat{J}^{\pm}$ on \hat{X} , and choose any $\theta \in \Theta_{\text{II}}$ mapping to θ^{\pm} . We will show that the map on chains

$$CC_{\bullet}(Y, \xi)_{\lambda, J, \theta^+} \xrightarrow{\Phi(\hat{X}, \hat{\lambda})_{\hat{J}, \theta}} CC_{\bullet}(Y, \xi)_{\lambda, J, \theta^-} \tag{1.29}$$

is an isomorphism, which is clearly sufficient.

We consider the ascending filtration on both sides of (1.29) whose $\leq^{(a,k)}$ filtered piece is the \mathbb{Q} -subspace generated by all monomials of Reeb orbits with total action $< a$ or total action $= a$ and degree $\geq k$. We claim that the map (1.29) as well as the differentials on both its domain and codomain all respect this filtration. Indeed, the integral of $d\lambda$ over any \hat{J} -holomorphic curve is ≥ 0 , with equality iff the curve is a branched cover of a trivial cylinder (note the difference between $d\lambda$ and the symplectic form $d\hat{\lambda} = d(e^s \lambda) = e^s(d\lambda + ds \wedge \lambda)$). Every branched cover of a trivial cylinder has at least one negative end, which proves the claim. Since the filtration is well-ordered, to show that (1.29) is an isomorphism, it suffices to show that the induced map on associated gradeds is an isomorphism.

The curves contributing to the action of (1.29) on associated gradeds are the branched covers of trivial cylinders with exactly one negative end, and such curves are themselves necessarily trivial cylinders by Riemann–Hurwitz. Since there is exactly one such trivial cylinder for every Reeb orbit, it suffices to show that trivial cylinders are cut out transversally. This is a standard fact, whose proof we recall in Lemma 2.40. \square

Now for a contact manifold (Y, ξ) , all groups $CH_{\bullet}(Y, \xi)_{\lambda, J, \theta}$ are canonically isomorphic via the morphisms $\Phi(\hat{X}, \hat{\lambda})$ associated to the trivial cobordisms (by Lemma 1.2 and the commutativity of (1.27)). Thus we get a well-defined object

$$CH_{\bullet}(Y, \xi), \tag{1.30}$$

independent of (λ, J, θ) . Formally speaking, $CH_{\bullet}(Y, \xi)$ is the limit (and the colimit) of $\{CH_{\bullet}(Y, \xi)_{\lambda, J, \theta}\}_{\lambda, J, \theta}$, which is attained at any particular triple (λ, J, θ) . Note that for any

contact structure, the set of non-degenerate contact forms is generic (and in particular non-empty).

The commutativity of (1.27) also implies that a deformation class of exact symplectic cobordism $(\hat{X}, \hat{\lambda})$ from (Y^+, ξ^+) to (Y^-, ξ^-) induces a well-defined graded unital \mathbb{Q} -algebra map

$$\Phi(\hat{X}, \hat{\lambda}) : CH_{\bullet}(Y^+, \xi^+) \rightarrow CH_{\bullet}(Y^-, \xi^-), \quad (1.31)$$

and that $\Phi(\hat{X}^{02}, \hat{\lambda}^{02}) = \Phi(\hat{X}^{12}, \hat{\lambda}^{12}) \circ \Phi(\hat{X}^{01}, \hat{\lambda}^{01})$ for $\hat{X}^{02} = \hat{X}^{01} \# \hat{X}^{12}$.

To construct the symmetric monoidal structure on CH_{\bullet} , it suffices to observe that (as we shall prove in Proposition 4.36) the sets Θ_I, Θ_{II} are themselves weakly symmetric monoidal in a manner which preserves the virtual moduli counts. This completes the construction of the contact homology functor (1.22) in terms of Theorem 1.1.

1.8 Applications and extensions

There is quite some literature devoted to properties and applications of contact homology and symplectic field theory which relies, explicitly or implicitly, on the existence of “suitable” virtual curve counting foundations. An exhaustive investigation of the applicability and extendibility of our results to these myriad of settings is beyond the scope of this paper. We nevertheless include some brief indications in this direction.

- (Grading by $H_1(Y)$) Contact homology $CH_{\bullet}(Y, \xi)$ has a grading by $H_1(Y)$ (the grading of a given monomial in Reeb orbits equals its total homology class).
- (Refinement of $\mathbb{Z}/2$ -grading) Contact homology $CH_{\bullet}(Y, \xi)$ has a relative grading by $\mathbb{Z}/2c_1(\xi) \cdot H_2(Y)$, which is absolute over the $0 \in H_1(Y)$ graded piece. The homological grading of γ is given by $|\gamma| = CZ(\gamma) + n - 3$.
- (Action filtration) If we equip (Y, ξ) with a contact form λ , then for $a \in \mathbb{R}$, there is an invariant $CH_{\bullet}(Y, \lambda)^{<a}$ which is equipped with functorial maps $CH_{\bullet}(Y, \lambda)^{<a} \rightarrow CH_{\bullet}(Y, \lambda')^{<a'}$ for $\frac{\lambda}{a} \geq \frac{\lambda'}{a'}$ (pointwise) such that

$$CH_{\bullet}(Y, \xi) = \varinjlim CH_{\bullet}(Y, \lambda)^{<a}. \quad (1.32)$$

Namely, for λ non-degenerate, $CH_{\bullet}(Y, \lambda)^{<a}$ is defined as the homology of the subcomplex $CC_{\bullet}(Y, \xi)_{\lambda, J, \theta}^{<a} \subseteq CC_{\bullet}(Y, \xi)_{\lambda, J, \theta}$ spanned by those monomials of total action $< a$, and for general λ we define $CH_{\bullet}(Y, \lambda)^{<a}$ as the direct limit of $CH_{\bullet}(Y, \lambda')^{<a}$ over non-degenerate $\lambda' > \lambda$. This invariant $CH_{\bullet}^{<a}$ may be constructed out of Theorem 1.1 as in §1.7.

- (Coefficients in $\mathbb{Q}[H_2(Y)]$) Contact homology $CH_{\bullet}(Y, \xi)$ has a natural lift $\overline{CH}_{\bullet}(Y, \xi)$ to the group ring $\mathbb{Q}[H_2(Y; \mathbb{Z})]$. More intrinsically, \overline{CH}_{\bullet} may be thought of as a local system over the space of 1-cycles in Y , namely $\tau_{\geq 0}C_{\bullet+1}(Y)$. Contact homology with group ring coefficients $\overline{CH}_{\bullet}(Y, \xi)$ has a relative \mathbb{Z} -grading, where $\mathbb{Q}[H_2(Y)]$ is \mathbb{Z} -graded by $2c_1(\xi) : H_2(Y) \rightarrow \mathbb{Z}$. This invariant \overline{CH}_{\bullet} may be constructed out of Theorem 1.1 as in §1.7.

- (Contact homology of contractible orbits) An invariant $CH_{\bullet}^{\text{contr}}(Y, \xi)$ is obtained from the chain complex $CC_{\bullet}^{\text{contr}}(Y, \lambda)$ generated as an algebra by contractible Reeb orbits (with a differential which counts curves whose asymptotic orbits are all contractible). There is also an invariant $CH_{\bullet}^{\alpha}(Y, \xi)$ obtained from the chain complex $CC_{\bullet}^{\alpha}(Y, \lambda)$ generated as a module over $CC_{\bullet}^{\text{contr}}(Y, \lambda)$ by Reeb orbits in a fixed nontrivial homotopy class α (with differential counting curves whose asymptotic orbits are either all contractible or all contractible except for the positive end and one negative end both in class α). These invariants $CH_{\bullet}^{\text{contr}}$ and CH_{\bullet}^{α} may be constructed out of Theorem 1.1 as in §1.7.
- (Cylindrical contact homology) If (Y, ξ) is *hypertight* (admits a contact form with no contractible Reeb orbits) then there is an invariant $CH_{\bullet}^{\text{cyl}}(Y, \xi)$ defined as follows. If (Y, ξ) admits a *non-degenerate* contact form with no contractible Reeb orbits, then $CH_{\bullet}^{\text{cyl}}(Y, \xi)$ is defined as the homology of the complex $CC_{\bullet}^{\text{cyl}}(Y, \lambda) := \bigoplus_{\gamma \in \mathcal{P}_{\text{good}}} \sigma_{\gamma}$ with the differential which counts pseudo-holomorphic cylinders. If this is not the case, then one must first define $CH_{\bullet}^{\text{cyl}}(Y, \lambda)^{<a}$ for non-degenerate contact forms λ with no contractible Reeb orbits of action $< a$, and then let $CH_{\bullet}^{\text{cyl}}(Y, \xi) := \varinjlim CH_{\bullet}^{\text{cyl}}(Y, \lambda)^{<a}$. As this definition makes clear, to define $CH_{\bullet}^{\text{cyl}}(Y, \xi)$ we actually only need to assume that (Y, ξ) is *asymptotically hypertight*, namely that it admits a sequence of contact forms λ_i converging uniformly to zero, each with no contractible Reeb orbits of action < 1 . This invariant $CH_{\bullet}^{\text{cyl}}$ may be constructed out of Theorem 1.1 as in §1.7.

Eliashberg–Givental–Hofer [EGH00] assert that $CH_{\bullet}^{\text{cyl}}(Y, \xi)$ can be defined assuming only that (Y, ξ) admits a non-degenerate contact form with no contractible Reeb orbits of index 1, 0, or -1 . We have nothing positive to say about the applicability of our methods to this question.

- (Even contact forms) It is trivial to calculate contact homology given a non-degenerate contact form all of whose Reeb orbits are even, since the differential then vanishes for index reasons. For examples of such situations, we refer the reader to Ustilovsky [Ust99] and Abreu–Macarini [AM12].
- (Overtwisted contact manifolds) A given (connected, non-empty) contact manifold is either *tight* or *overtwisted*. Overtwisted contact structures are classified completely by an h -principle due to Eliashberg [Eli89] in dimension three and Borman–Eliashberg–Murphy [BEM15] in general.

Contact homology (even with group ring coefficients) vanishes on any overtwisted contact manifold. In dimension three, this is a result of Eliashberg [Eli98, p334, Theorem 3.5(2)] (a proof is given in Yau [Yau06] and the appendix by Eliashberg). In all dimensions, this follows from the result of Bourgeois–van Koert [BvK10, Theorem 1.3] that contact homology vanishes for any contact manifold admitting a negatively stabilized open book, together with the result of Casals–Murphy–Presas [CMP19, Theorem 1.1] that a contact manifold admits a negatively stabilized open book iff it is overtwisted. These vanishing results are proved by exhibiting a contact form with a non-degenerate Reeb orbit bounding exactly one pseudo-holomorphic plane in the symplectization

(which is cut out transversally); in particular, they are valid for the contact homology we construct here.

An *a priori* weaker notion of a contact manifold being *PS-overtwisted* was introduced by Niederkrüger [Nie06] and Massot–Niederkrüger–Wendl [MNW13]. An argument, due to Bourgeois–Niederkrüger, that contact homology vanishes for PS-overtwisted contact manifolds is sketched in [Bou09]; the point is to count pseudo-holomorphic disks with boundary on the plastikstufe/bLob, one boundary marked point constrained to map to a fixed curve on the plastikstufe/bLob from its core/binding to its boundary, and an arbitrary number of negative punctures. The proof of Theorem 1.1 generalizes easily to provide sufficient virtual moduli counts for this argument.

- (Filling and cobordism obstructions) A contact manifold with vanishing contact homology is not symplectically fillable, and more generally if the positive end of a symplectic cobordism has vanishing contact homology then so does the negative end. Indeed, the existence of a unital ring map $0 \rightarrow R$ implies $R = 0$, and the cobordism in question induces a unital map $CH_\bullet(Y^+, \xi^+) \rightarrow CH_\bullet(Y^-, \xi^-)$ (to obtain such a map in the non-exact case, one can count curves which represent zero in $H_2(\hat{X}, Y^+ \sqcup Y^-)$). The curve counts produced by Theorem 1.1 are sufficient for this argument.

It is an observation of Niederkrüger–Wendl [NW11] and Latschev–Wendl [LW11] that a similar implication holds for “stable symplectic cobordisms” and contact homology with twisted coefficients. (What we call a symplectic cobordism is often called a “strong symplectic cobordism”, and the notion of a stable symplectic cobordism is more general; see also Massot–Niederkrüger–Wendl [MNW13].) The literal statement of Theorem 1.1 does not cover stable symplectic cobordisms, however its proof applies to them without modification, as the differences between the two settings are simply not relevant to any part of the argument (except for compactness, which is proved in the requisite generality in [BEHWZ03]).

- (Invariants of contactomorphisms) There is a natural homomorphism

$$\pi_0 \text{Cont}(Y, \xi) \rightarrow \text{Aut}_{\mathbb{Q}}(CH_\bullet(Y, \xi)), \quad (1.33)$$

namely the tautological action of $\text{Cont}(Y, \xi)$ on $CH_\bullet(Y, \xi)$. This action admits the following description in terms of cobordism maps which shows that it descends to π_0 . For any $\varphi \in \text{Cont}(Y, \xi)$, denote by X_φ the exact symplectic cobordism from Y to itself obtained from the trivial cobordism by changing the marking on the negative end by φ . The action of φ on $CH_\bullet(Y, \xi)$ clearly coincides with the cobordism map $\Phi(X_\varphi)$. On the other hand, $\Phi(X_\varphi)$ only depends on the class of φ in π_0 since cobordism maps are invariant under deformation. It is also clear that $\varphi \mapsto \Phi(X_\varphi)$ is a group homomorphism since $X_\varphi \# X_\psi = X_{\varphi\psi}$.

As pointed out by P. Massot, any contactomorphism which is *symplectically pseudo-isotopic* to the identity (a notion due to Cieliebak–Eliashberg [CE12, §14.5]) lies in the kernel of (1.33); indeed, φ is symplectically pseudo-isotopic to the identity iff X_φ is isomorphic to X_{id} as symplectic cobordisms from Y to itself. Thus a contactomorphism which acts nontrivially on contact homology cannot be symplectically pseudo-isotopic to the identity.

- (Invariants of families of contactomorphisms) The above construction generalizes to give a natural homomorphism

$$H_k(\text{Cont}(Y, \xi)) \rightarrow \text{Hom}_{\mathbb{Q}}(CH_{\bullet}(Y, \xi), CH_{\bullet+k}(Y, \xi)), \quad (1.34)$$

as introduced by Bourgeois [Bou06] (more precisely, Bourgeois introduced the pre-compositions of (1.33)–(1.34) with the map $\Omega_{\xi}\Xi(Y) \rightarrow \text{Cont}(Y, \xi)$ coming from Gray’s fibration sequence $\text{Cont}(Y, \xi) \rightarrow \text{Diff}(Y, \xi) \rightarrow \Xi(Y)$, where $\Xi(Y)$ denotes the space of contact structures on Y and Ω_{ξ} denotes the space of loops based at $\xi \in \Xi(Y)$). Namely, a family of $\varphi \in \text{Cont}(Y, \xi)$ gives rise to a family of cobordisms X_{φ} , and case (III) of Theorem 1.1 (and its proof) generalize immediately to such higher-dimensional families of cobordisms, in the following form. For any family of cobordisms $(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t \in \Delta^n})$, there is a set $\Theta_{\text{III}(n)} = \Theta_{\text{III}(n)}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t \in \Delta^n})$ (specializing to $\Theta_{\text{III}(0)} = \Theta_{\text{II}}$ and $\Theta_{\text{III}(1)} = \Theta_{\text{III}}$) along with surjective maps

$$\Theta_{\text{III}(n)}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t \in \Delta^n}) \twoheadrightarrow \lim_{\Delta^k \subsetneq \Delta^n} [\Theta_{\text{III}(k)}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t \in \Delta^k}) \rightarrow (\Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-)] \quad (1.35)$$

(where the limit is in the category of sets over $\Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-$), and there are associated virtual moduli counts satisfying the natural master equation, thus giving rise (in the exact case) to “higher homotopies” $K(\hat{X}, \hat{\lambda}^t)_{\hat{J}^t, \theta}$ satisfying

$$\begin{aligned} d_{J^-, \theta^-} K(\hat{X}, \hat{\lambda}^t)_{\hat{J}^t, \theta} - (-1)^n K(\hat{X}, \hat{\lambda}^t)_{\hat{J}^t, \theta} d_{J^+, \theta^+} \\ = \sum_{\substack{k=0 \\ i: \Delta^{[0 \dots \hat{k} \dots n]} \hookrightarrow \Delta^{[0 \dots n]}}}^n (-1)^k K(\hat{X}, \hat{\lambda}^{i(t)})_{\hat{J}^{i(t)}, i^* \theta}. \end{aligned} \quad (1.36)$$

This defines (1.34).

We expect that it is possible to upgrade (1.33)–(1.34) into the statement that $CC_{\bullet}(Y, \xi)$ is an A_{∞} -module over $C_{\bullet}(\text{Cont}(Y, \xi))$ (i.e. that contact chains of (Y, ξ) is a “derived local system” over $B\text{Cont}(Y, \xi)$), by proving a more intricate higher-dimensional version of case (IV) of Theorem 1.1. To realize this rigorously, it would be sufficient to specify the correct compactified moduli spaces to consider and to show that their stratifications are locally cell-like as in §§2.5, 3.2 (these are essentially combinatorial questions).

- (Chain homotopy vs DGA homotopy) Given a one-parameter family of exact symplectic cobordisms, case (III) of Theorem 1.1 shows that the two cobordism maps

$$CC_{\bullet}(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \begin{array}{c} \xrightarrow{\Phi(\hat{X}, \hat{\lambda}^0)_{J^0, \theta^0}} \\ \xrightarrow{\Phi(\hat{X}, \hat{\lambda}^1)_{J^1, \theta^1}} \end{array} CC_{\bullet}(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \quad (1.37)$$

are *chain homotopic*, by constructing virtual moduli counts for moduli spaces of *disconnected curves* (and analogously for case (IV)). For certain purposes (such as defining linearized contact homology), one may need to know that such cobordism maps are homotopic in a stronger sense, namely one which takes into account the algebra

structure. The construction of such homotopies is sketched in Eliashberg–Givental–Hofer [EGH00, §2.4] and in Ekholm–Oancea [EO17, §§5.4–5.5] in more detail (also see Ekholm–Honda–Kálmán [EHK16, Lemma 3.13] for the related, but easier, case of Legendrian contact homology). We have nothing positive to say about the applicability of our methods to the rigorous construction of such homotopies. One can generalize case (III) of Theorem 1.1 (and its proof) a little by counting curves occurring at times $0 \leq t_1 \leq \dots \leq t_k \leq 1$ and obtain an A_∞ -homotopy, though it is not clear if this is any improvement.

- (Symplectic field theory) To construct the symplectic field theory “homology groups” as described in Cieliebak–Latschev [CL09] and Latschev–Wendl [LW11] (specializing work of Eliashberg–Givental–Hofer [EGH00]) using the methods of this paper, it would be sufficient to specify the correct compactified moduli spaces to consider and to show that their stratifications are locally cell-like as in §§2.5,3.2 (these are essentially combinatorial questions).
- (Morse–Bott contact forms) It should be possible to define contact homology using Morse–Bott contact forms, by counting appropriate pseudo-holomorphic cascades as in Bourgeois [Bou02, Bou03]. To realize this rigorously, the main tasks would be to specify the correct compactified moduli spaces to consider, to show that their stratifications are locally cell-like as in §§2.5,3.2 (these are essentially combinatorial questions), and to establish the relevant gluing result as in §5.
- (Integer coefficients) It is an interesting open question (promoted by Abouzaid) how to naturally lift contact homology from \mathbb{Q} to \mathbb{Z} .
- (Non-equivariant contact homology) We conjecture that contact homology can be lifted to a framed E_2 -algebra, namely an algebra over the operad whose space of n -ary operations is the space of embeddings $(D^2)^{\sqcup n} \hookrightarrow D^2$. More precisely, the free symmetric \mathbb{Q} -algebra on \mathfrak{o}_γ for $\gamma \in \mathcal{P}_{\text{good}}$ in (1.1) should be replaced with the free framed E_2 -algebra on \mathfrak{o}_γ for $\gamma \in \mathcal{P}$ subject to the relations that rotating any d -fold multiply covered orbit by $2\pi/d$ acts on \mathfrak{o}_γ by the corresponding sign (± 1 according to whether the Reeb orbit is good or bad). The differential of a given orbit γ^+ should lift naturally to an element of this (almost) free framed E_2 -algebra by taking into account the conformal types of the domains of the pseudo-holomorphic curves in question. This lift can be viewed as the “non-equivariant” version of contact homology (which should itself be viewed as an S^1 -equivariant theory).

2 Moduli spaces of pseudo-holomorphic curves

In this section, we define the moduli spaces of pseudo-holomorphic curves which we will use to define contact homology.

2.1 Categories of strata \mathcal{S}

We begin by introducing, for any datum \mathcal{D} from any of Setups I–IV, a collection $\mathcal{S} = \mathcal{S}(\mathcal{D})$ of labeled trees which index the strata of the compactified moduli spaces of pseudo-holomorphic curves (which of Setups I–IV is being considered is often indicated with a subscript, as in \mathcal{S}_I , \mathcal{S}_{II} , etc.). A labeled tree describes the “combinatorial type” of a pseudo-holomorphic curve: the tree is (almost) the dual graph of the domain, and it is labeled with the homotopy class and asymptotics of the map.

Gluing pseudo-holomorphic curves corresponds to contracting a subset of the edges of a tree and updating the labels accordingly. We regard these collections \mathcal{S} as categories, with such edge contractions as morphisms. Morphisms of labeled trees will induce maps of compactified moduli spaces.

The categories \mathcal{S} carry some additional (vaguely monoidal) structure, corresponding to the fact that boundary strata in compactified moduli spaces can be expressed as products of other “smaller” moduli spaces. The relevant operation on trees is that of “concatenation”, where we take some collection of trees and identify some pairs of input/output edges with matching labels to produce a new tree.




Figure 1: A tree. The edges are all directed downwards.




Figure 2: A contraction of trees. The edges which have been contracted are marked bold.

The category \mathcal{S}_I is defined as follows, with respect to some fixed datum of Setup I. An object of \mathcal{S}_I consists first of a finite directed tree T (for us, the word “tree” entails being non-empty) in which every vertex has a unique incoming edge (see Figure 1). Edges with

missing source (called input edges) and edges with missing sink (called output edges) are allowed, together these are called external edges, and the remaining edges will be called interior edges. The tree T is further equipped with decorations as follows:

- (i) For each edge $e \in E(T)$, a Reeb orbit $\gamma_e \in \mathcal{P}$.
- (ii) For each vertex $v \in V(T)$, a homotopy class $\beta_v \in \pi_2(Y, \gamma_{e^+(v)} \sqcup \{\gamma_{e^-}\}_{e^- \in E^-(v)})$, where we denote by $e^+(v) \in E(T)$ the unique incoming edge at v , and by $E^-(v)$ the set of outgoing edges at v .
- (iii) For each external edge $e \in E^{\text{ext}}(T)$, a basepoint $b_e \in |\gamma_e|$, where $|\gamma_e|$ denotes the underlying simple orbit of γ_e .

A morphism $\pi : T \rightarrow T'$ in \mathcal{S}_I consists first of a contraction of underlying trees. By this we mean that some collection of interior edges of T is specified, and T' is identified with the result of contracting these edges (see Figure 2). This identification must be compatible with the decorations in the following sense:

- (i) For each non-contracted edge $e \in E(T)$, we have $\gamma_{\pi(e)} = \gamma_e$.
- (ii) For each vertex $v' \in V(T')$, we have $\beta_{v'} = \#_{\pi(v)=v'} \beta_v$.

Finally, we must specify for each external edge $e \in E^{\text{ext}}(T) = E^{\text{ext}}(T')$ a path along $|\gamma_e|$ between the basepoints b_e and b'_e , modulo the relation that identifies two such paths iff their “difference” lifts to γ_e (i.e. iff it has degree divisible by the covering multiplicity of γ_e). Such paths will correspond, in the realm of pseudo-holomorphic curves, to ways of moving asymptotic markers. Conceptually, we should really regard the input/output edges of T as being labeled by objects of the category $\tilde{\mathcal{P}}$ whose objects are Reeb orbits $\gamma \in \mathcal{P}$ together with a basepoint $b \in |\gamma|$ and whose morphisms are paths as just described (the set of isomorphism classes in this category is $|\tilde{\mathcal{P}}| = \mathcal{P}$, and the automorphism group of an object in the isomorphism class of $\gamma \in \mathcal{P}$ is (canonically) \mathbb{Z}/d_γ).

We now introduce the notion of a concatenation in \mathcal{S}_I . A concatenation in \mathcal{S}_I consists of a finite non-empty collection of objects $T_i \in \mathcal{S}_I$ along with a matching between some pairs of external edges of the T_i 's (with matching γ_e) such that the resulting gluing is a tree, along with a choice of paths between the basepoints for each pair of matched edges. Given a concatenation $\{T_i\}_i$ in \mathcal{S}_I , there is a resulting object $\#_i T_i \in \mathcal{S}_I$. A morphism of concatenations $\{T_i\}_i \rightarrow \{T'_i\}_i$ shall mean a collection of morphisms $T_i \rightarrow T'_i$ covering a bijection of index sets; a morphism of concatenations $\{T_i\}_i \rightarrow \{T'_i\}_i$ induces a morphism $\#_i T_i \rightarrow \#_i T'_i$. If $\{T_i\}_i$ is a concatenation in \mathcal{S}_I and $T_i = \#_j T_{ij}$ for some concatenations $\{T_{ij}\}_j$, there is a resulting composite concatenation $\{T_{ij}\}_{ij}$ with natural isomorphisms $\#_{ij} T_{ij} = \#_i \#_j T_{ij} = \#_i T_i$.

Continuing on to Setup II, the category \mathcal{S}_{II} (again for a particular choice of datum of Setup II) is defined as follows. An object of \mathcal{S}_{II} is a decorated tree T as before, but with additional labels which specify the target of the pseudo-holomorphic curve corresponding to each vertex. For ease of notation, we set $\hat{X}^{00} = \hat{Y}^0 := \hat{Y}^+$, $\hat{X}^{01} := \hat{X}$, and $\hat{X}^{11} = \hat{Y}^1 := \hat{Y}^-$. The additional labels (which in turn determine the allowable decorations) are as follows:

- (i) For each edge $e \in E(T)$, a symbol $*(e) \in \{0, 1\}$ such that input edges are labeled with 0 and output edges are labeled with 1.
- (ii) For each vertex $v \in V(T)$, a pair of symbols $*^\pm(v) \in \{0, 1\}$ such that $*^+(v) \leq *^-(v)$ and $*(e^\pm(v)) = *^\pm(v)$. If $*^+(v) = *^-(v)$, we call v a symplectization vertex.

See Figure 3. These labels determine the allowable decorations: $\gamma_e \in \mathcal{P}(Y^{*(e)}, \lambda^{*(e)})$ and $\beta_v \in \pi_2(\hat{X}^{*(v)}, \gamma_{e^+(v)} \sqcup \{\gamma_{e^-}\}_{e^- \in E^-(v)})$. A morphism $\pi : T \rightarrow T'$ in \mathcal{S}_{II} consists first of a




Figure 3: An object of \mathcal{S}_{II} , with vertex labels as shown. Note that the vertex labels determine the edge labels uniquely (and conversely except for vertices with no outgoing edges).




Figure 4: A morphism in \mathcal{S}_{II} . A morphism in \mathcal{S}_{II} is determined uniquely by the set of contracted edges and the set of vertices whose label changes from 00 to 01 (only needed for vertices with no outgoing edges).




Figure 5: This object of \mathcal{S}_{II} cannot be contracted along (only) the marked edge (there is no way to consistently label the result).

contraction of underlying trees, required to satisfy $*^+(\pi(v)) \leq *^+(v)$, $*^-(\pi(v)) \geq *^-(v)$, and $*(\pi(e)) = *(e)$ for non-contracted edges e . Given these, it makes sense to require compatibility with the decorations and to specify paths between basepoints as before. See Figures 4 and 5.

A concatenation $\{T_i\}_i$ in \mathcal{S}_{II} is defined as before, except that each T_i may be an object of either \mathcal{S}_{II} , \mathcal{S}_{I}^+ , or \mathcal{S}_{I}^- (where $\mathcal{S}_{\text{I}}^\pm = \mathcal{S}_{\text{I}}(Y^\pm, \lambda^\pm)$). Given a concatenation $\{T_i\}_i$ in \mathcal{S}_{II} , there is a resulting object $\#_i T_i \in \mathcal{S}_{\text{II}}$. If $\{T_i\}_i$ is a concatenation in \mathcal{S}_{II} and $T_i = \#_j T_{ij}$ for some concatenations $\{T_{ij}\}_j$ (in whichever of \mathcal{S}_{II} , \mathcal{S}_{I}^+ , \mathcal{S}_{I}^- contains T_i), there is a resulting *composite concatenation* $\{T_{ij}\}_{ij}$ with natural isomorphisms $\#_{ij} T_{ij} = \#_i \#_j T_{ij} = \#_i T_i$.

Continuing on to Setup III, the category \mathcal{S}_{III} is defined as follows. An object of \mathcal{S}_{III} is a forest (meaning connected and non-empty assumptions are dropped) with labels and decorations as in the case of \mathcal{S}_{II} , together with a choice of set $\mathfrak{s} \in \{\{0\}, \{1\}, (0, 1)\}$. A morphism $T \rightarrow T'$ in \mathcal{S}_{III} is defined as for \mathcal{S}_{II} , with the additional requirement that $\mathfrak{s}(T) \subseteq \mathfrak{s}(T')$.

There are three types of concatenations in \mathcal{S}_{III} . The first type consists of objects $T_i \in \mathcal{S}_{\text{I}}^+ \sqcup \mathcal{S}_{\text{II}}^{t=0} \sqcup \mathcal{S}_{\text{I}}^-$ with the usual matching data, producing an object $\#_i T_i \in \mathcal{S}_{\text{III}}$ with $\mathfrak{s}(\#_i T_i) = \{0\}$. The second type consists of objects $T_i \in \mathcal{S}_{\text{I}}^+ \sqcup \mathcal{S}_{\text{II}}^{t=1} \sqcup \mathcal{S}_{\text{I}}^-$ with matching data, producing an object $\#_i T_i \in \mathcal{S}_{\text{III}}$ with $\mathfrak{s}(\#_i T_i) = \{1\}$. The third type consists of objects $T_i \in \mathcal{S}_{\text{I}}^+ \sqcup \mathcal{S}_{\text{III}} \sqcup \mathcal{S}_{\text{I}}^-$ (exactly one of which lies in \mathcal{S}_{III}) with matching data, producing an object $\#_i T_i \in \mathcal{S}_{\text{III}}$ (where $\mathfrak{s}(\#_i T_i) = \mathfrak{s}(T_i)$ for the unique $T_i \in \mathcal{S}_{\text{III}}$). A composition of concatenations is defined as before.

The reason for the restriction that concatenations in \mathcal{S}_{III} of the third type must contain exactly one object of \mathcal{S}_{III} is explained as follows. As mentioned earlier, a concatenation $\{T_i\}_i$ is supposed to induce an identification of $\overline{\mathcal{M}}(\#_i T_i)$ with $\prod_i \overline{\mathcal{M}}(T_i)$ (modulo the symmetry coming from multiply covered orbits). The moduli spaces $\overline{\mathcal{M}}_{\text{III}}$ parameterize pseudo-holomorphic curves at varying times $t \in [0, 1]$, and thus they come with a forgetful map to $[0, 1]$. Corresponding to a concatenation in \mathcal{S}_{III} of the third type without the requirement of containing exactly one object of \mathcal{S}_{III} , there is not an identification of $\overline{\mathcal{M}}(\#_i T_i)$ with $\prod_i \overline{\mathcal{M}}(T_i)$, but rather with the corresponding *fiber product* over $[0, 1]$. We do not discuss such identifications in this paper because we do not know how to prove (or even formulate precisely) the corresponding expected compatibility of virtual fundamental cycles with such fiber products (this is related to the problem of showing compatibility of the homotopy with the algebra structure on contact homology, see §1.8). The same remark will apply to concatenations in \mathcal{S}_{IV} .

The category \mathcal{S}_{IV} is defined as follows. An object in \mathcal{S}_{IV} is a forest T with the following labels and decorations. Each edge shall be labeled with a symbol $*(e) \in \{0, 1, 2\}$, such that input edges are labeled with 0 and output edges are labeled with 2. Each vertex shall be labeled with a pair of symbols $*^\pm(v) \in \{0, 1, 2\}$ such that $*^+(v) \leq *^-(v)$ and $*(e^\pm(v)) = *^\pm(v)$. There shall also be decorations $\gamma_e \in \mathcal{P}(Y^{*(e)})$ and $\beta_v \in \pi_2(X^{*(v)}, \gamma_{e^+(v)} \sqcup \{\gamma_{e^-}\}_{e^- \in E^-(v)})$ with basepoints $b_e \in |\gamma_e|$ for input/output edges as before. Finally, we specify a set $\mathfrak{s} \in \{\{0\}, \{\infty\}, (0, \infty)\}$ and we require that if $\mathfrak{s} \in \{\{0\}, (0, \infty)\}$ then $*(v) \in \{00, 02, 22\}$ for all v , and if $\mathfrak{s} = \{\infty\}$ then $*(v) \in \{00, 01, 11, 12, 22\}$ for all v . A morphism $T \rightarrow T'$ in \mathcal{S}_{IV} is defined as for \mathcal{S}_{II} , with the additional requirement that $\mathfrak{s}(T) \subseteq \mathfrak{s}(T')$.

There are three types of concatenations in \mathcal{S}_{IV} . The first type consists of objects $T_i \in \mathcal{S}_{\text{I}}^0 \sqcup \mathcal{S}_{\text{II}}^{02} \sqcup \mathcal{S}_{\text{I}}^2$ with matching data, producing an object $\#_i T_i \in \mathcal{S}_{\text{IV}}$ with $\mathfrak{s}(\#_i T_i) = \{0\}$. The second type consists of objects $T_i \in \mathcal{S}_{\text{I}}^0 \sqcup \mathcal{S}_{\text{II}}^{01} \sqcup \mathcal{S}_{\text{I}}^1 \sqcup \mathcal{S}_{\text{II}}^{12} \sqcup \mathcal{S}_{\text{I}}^2$ with matching data, producing an

object $\#_i T_i \in \mathcal{S}_{\text{IV}}$ with $\mathfrak{s}(\#_i T_i) = \{\infty\}$. The third type consists of objects $T_i \in \mathcal{S}_1^0 \sqcup \mathcal{S}_{\text{IV}} \sqcup \mathcal{S}_1^2$ (exactly one of which lies in \mathcal{S}_{IV}) with matching data, producing an object $\#_i T_i \in \mathcal{S}_{\text{IV}}$ (where $\mathfrak{s}(\#_i T_i) = \mathfrak{s}(T_i)$ for the unique $T_i \in \mathcal{S}_{\text{IV}}$). A composition of concatenations is defined as before.

2.2 Properties of categories \mathcal{S}

The categories \mathcal{S} together with their concatenation structure satisfy some basic properties which will be used frequently. We begin with some basic definitions.

Definition 2.1 (Slice categories). For any category \mathcal{T} and an object $T \in \mathcal{T}$, denote by $\mathcal{T}/_T$ the “over-category” whose objects are morphisms $T' \rightarrow T$. Similarly define the “under-category” $\mathcal{T}_{/T}$ (objects are morphisms $T \rightarrow T'$) and $\mathcal{T}_{T//T'}$ (objects are factorizations $T \rightarrow T'' \rightarrow T'$ of a fixed morphism $T \rightarrow T'$).

Definition 2.2 (Poset). A *poset* is a category \mathcal{T} for which for every pair of objects $x, y \in \mathcal{T}$, there is at most one morphism $x \rightarrow y$. The *cardinality* of a poset is its number of isomorphism classes of objects.

Both $\mathcal{S}_{T/}$ and $\mathcal{S}_{T//T'}$ are finite posets for any T or $T \rightarrow T'$, respectively. On the other hand, $\mathcal{S}_{/T}$ is almost never finite nor a poset (see Figure 6).




Figure 6: A map of trees $T' \rightarrow T$ for which $\text{Aut}(T'/T) = \mathbb{Z}/2$, corresponding to exchanging the lower two vertices of T' .

Definition 2.3 (Automorphism groups). Given a map $T' \rightarrow T$, we denote by $\text{Aut}(T'/T)$ the subgroup of $\text{Aut}(T')$ compatible with the map $T' \rightarrow T$ (i.e. the automorphism group of $(T' \rightarrow T) \in \mathcal{S}_{/T}$). Given a concatenation $\{T_i\}_i$, we denote by $\text{Aut}(\{T_i\}_i/\#_i T_i)$ the group of automorphisms of $\{T_i\}_i$ acting trivially on $\#_i T_i$ (i.e. the product $\prod_e \mathbb{Z}/d_{\gamma_e}$ over junction edges, acting diagonally).

Definition 2.4. An object $T \in \mathcal{S}$ is called *maximal* iff the only morphism out of T is the identity map (i.e. $\mathcal{S}_{T/}$ consists of a single object with trivial endomorphism monoid).

Maximality can be characterized concretely as follows:

- (i) $T \in \mathcal{S}_1$ is maximal iff it has exactly one vertex.
- (ii) $T \in \mathcal{S}_{\text{II}}$ is maximal iff it has exactly one vertex and this vertex has $*(v) = 01$.
- (iii) $T \in \mathcal{S}_{\text{III}}$ is maximal iff every component of T has exactly one vertex, all these vertices have $*(v) = 01$, and $\mathfrak{s}(T) = (0, 1)$.
- (iv) $T \in \mathcal{S}_{\text{IV}}$ is maximal iff every component of T has exactly one vertex, all these vertices have $*(v) = 02$, and $\mathfrak{s}(T) = (0, \infty)$.

Lemma 2.5. *Every object $T \in \mathcal{S}$ can be expressed as a concatenation of maximal T_i , uniquely up to isomorphism. In particular, $T \in \mathcal{S}$ is maximal iff it cannot be expressed nontrivially as a concatenation.* \square

2.3 Moduli spaces $\overline{\mathcal{M}}(T)$

We now define the compactified moduli spaces of pseudo-holomorphic curves $\overline{\mathcal{M}}(T)$ relevant for contact homology. We first define the strata $\mathcal{M}(T)$ for each $T \in \mathcal{S}$, and then we define the compactified moduli spaces $\overline{\mathcal{M}}(T)$ as unions of strata $\overline{\mathcal{M}}(T')$ for $T' \rightarrow T$.

Equip $\mathbb{R} \times S^1$ with coordinates (s, t) and with the standard almost complex structure $j(\partial_s) = \partial_t$, i.e. $z = e^{s+it}$.

Definition 2.6. Fix (Y, λ, J) as in Setup I, and let $u : [0, \infty) \times S^1 \rightarrow \hat{Y}$ be a smooth map. We say that u is C^0 -convergent to (the positive trivial cylinder over) a Reeb orbit $\gamma \in \mathcal{P}$ iff

$$u(s, t) = (Ls + b, \tilde{\gamma}(t)) + o(1) \quad \text{as } s \rightarrow \infty, \quad (2.1)$$

for some $b \in \mathbb{R}$ and some $\tilde{\gamma} : S^1 \rightarrow Y$ with $\partial_t \tilde{\gamma} = L \cdot R_\lambda(\tilde{\gamma})$ parameterizing γ , where $o(1)$ is meant with respect to any \mathbb{R} -invariant Riemannian metric. When the error $o(1)$ decays like $O(e^{-\delta|s|})$ in all derivatives for some $\delta > 0$, we say u is C^∞ -convergent with weight δ . Similarly, we define the notion of $u : (-\infty, 0] \times S^1 \rightarrow \hat{Y}$ being convergent to a negative trivial cylinder by considering (2.1) in the limit $s \rightarrow -\infty$.

It is straightforward to check that for any Riemann surface C and $p \in C$, it makes sense to say that $u : C \setminus p \rightarrow \hat{Y}$ is C^0 - or C^∞ -convergent of weight $\delta \in (0, 1)$ to a trivial cylinder at p , i.e. that this notion is independent of the choice of coordinates $[0, \infty) \times S^1 \rightarrow C \setminus p$ near p (though note the importance of $\delta < 1$ in this statement). This notion also generalizes immediately to maps to any \hat{X} from Setups II, III, IV, by virtue of the markings (1.8)–(1.9).

If u is \hat{J} -holomorphic, then C^0 -convergence to a trivial cylinder implies C^∞ -convergence of every weight $\delta < \delta_\gamma$, where $\delta_\gamma > 0$ denotes the smallest nonzero absolute value among eigenvalues of the linearized operator of γ as in Definition 2.29 (by Hofer–Wysocki–Zehnder [HWZ02, Theorems 1.1, 1.2, and 1.3] which we restate as Proposition 2.20).

If u converges to a trivial cylinder over $\gamma \in \mathcal{P}$ at $p \in C$, then it induces a well-defined constant speed parameterization of γ denoted $u_p : S_p C \rightarrow Y$, where $S_p C := (T_p C \setminus 0)/\mathbb{R}_{>0}$ is the tangent sphere at p , which is a $U(1)$ -torsor due to the complex structure on C .

Definition 2.7 (Moduli space $\mathcal{M}_I(T)$). A *pseudo-holomorphic building of type $T \in \mathcal{S}_I$* consists of the following data (see Figure 7):

- (i) For every vertex v , a closed, connected, nodal Riemann surface of genus zero C_v , along with distinct points $\{p_{v,e} \in C_v\}_e$ indexed by the edges e incident at v .
- (ii) For every vertex v , a smooth map $u_v : C_v \setminus \{p_{v,e}\}_e \rightarrow \hat{Y}$.
- (iii) We require that u_v be C^0 -convergent to the positively trivial cylinder over $\gamma_{e^+(v)}$ at $p_{v,e^+(v)}$ and to the negative trivial cylinders over γ_{e^-} at p_{v,e^-} for $e^- \in E^-(v)$, and be in the homotopy class β_v .
- (iv) For every input/output edge e , an “asymptotic marker” $\tilde{b}_e \in S_{p_{v,e}} C_v$ which is mapped to the basepoint $b_e \in |\gamma_e|$ by $(u_v)_{p_{v,e}}$.




Figure 7: A stable pseudo-holomorphic building and the corresponding tree.

- (v) For every interior edge $v \xrightarrow{e} v'$, a “matching isomorphism” $m_e : S_{p_{v,e}} C_v \rightarrow S_{p_{v',e}} C_{v'}$ intertwining $(u_v)_{p_{v,e}}$ and $(u_{v'})_{p_{v',e}}$.
- (vi) We require that u_v be \hat{J} -holomorphic, i.e. $(du)_j^{0,1} = 0$.

An *isomorphism* $(\{C_v\}, \{p_{v,e}\}, \{u_v\}, \{\tilde{b}_e\}, \{m_e\}) \rightarrow (\{C'_v\}, \{p'_{v,e}\}, \{u'_v\}, \{\tilde{b}'_e\}, \{m'_e\})$ between pseudo-holomorphic buildings of type T consists of isomorphisms $\{i_v : C_v \rightarrow C'_v\}$ and real numbers $\{s_v \in \mathbb{R}\}$ such that $u_v = \tau_{s_v} \circ u'_v \circ i_v$ ($\tau_s : \hat{Y} \rightarrow \hat{Y}$ denotes translation by s), $i_v(p_{v,e}) = p'_{v,e}$, $i_v(\tilde{b}_e) = \tilde{b}'_e$, and $i_v \circ m_e = m'_e \circ i_v$. A pseudo-holomorphic building is called *stable* iff its automorphism group is finite. We denote by $\mathcal{M}_1(T)$ the set of isomorphism classes of stable pseudo-holomorphic buildings of type T . Note the tautological action of $\text{Aut}(T)$ on $\mathcal{M}(T)$ by changing the marking.

Remark 2.8. For topological reasons (exactness of \hat{Y}), the domains C_v of every pseudo-holomorphic building as in Definition 2.7 are in fact smooth (that is, are without nodes). We find it convenient to set up all the definitions in the generality of nodal domains (which do occur in non-exact cobordisms).

Remark 2.9. Concretely, a pseudo-holomorphic building is stable iff it does not contain either (1) a trivial cylinder at a symplectization vertex of T , or (2) an unstable irreducible component of C_v over which the map u_v is constant.

Definition 2.10 (Moduli space $\mathcal{M}_{\text{II}}(T)$). A *pseudo-holomorphic building of type $T \in \mathcal{S}_{\text{II}}$* consists of the following data:

- (i) Domains C_v and punctures $p_{v,e}$ as in Definition 2.7(i).
- (ii) For every vertex v , a smooth map $u_v : C_v \setminus \{p_{v,e}\}_e \rightarrow \hat{X}^{*(v)}$.
- (iii) Asymptotic conditions on u_v as Definition 2.7(iii).
- (iv) Asymptotic markers \tilde{b}_e as in Definition 2.7(iv).
- (v) Matching isomorphisms m_e as in Definition 2.7(v).
- (vi) \hat{J} -holomorphicity condition on u_v as Definition 2.7(vi).

An *isomorphism* between pseudo-holomorphic buildings of type T is defined as in Definition 2.7, except that there is a translation $s_v \in \mathbb{R}$ only if v is a symplectization vertex. We denote by $\mathcal{M}_{\text{II}}(T)$ the set of isomorphism classes of stable pseudo-holomorphic buildings of type T .

Definition 2.11 (Moduli space $\mathcal{M}_{\text{III}}(T)$). For $T \in \mathcal{S}_{\text{III}}$, denote by $\mathcal{M}_{\text{III}}(T)$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable pseudo-holomorphic buildings of type T (as in Definition 2.10) in (\hat{X}, \hat{J}^t) .

Definition 2.12 (Moduli space $\mathcal{M}_{\text{IV}}(T)$). For $T \in \mathcal{S}_{\text{IV}}$, denote by $\mathcal{M}_{\text{IV}}(T)$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable pseudo-holomorphic buildings of type T (as in Definition 2.10) in $(\hat{X}^{02,t}, \hat{J}^{02,t})$.

Definition 2.13 (Moduli spaces $\overline{\mathcal{M}}(T)$). For $T \in \mathcal{S}$, we define

$$\overline{\mathcal{M}}(T) := \bigsqcup_{T' \rightarrow T} \mathcal{M}(T') / \text{Aut}(T'/T). \quad (2.2)$$

The union is over the set of isomorphism classes $|\mathcal{S}_{/T}|$.

For a morphism $T' \rightarrow T$, there is a natural map

$$\overline{\mathcal{M}}(T') / \text{Aut}(T'/T) \rightarrow \overline{\mathcal{M}}(T), \quad (2.3)$$

and for a concatenation $\{T_i\}_i$, there is an induced identification

$$\prod_i \overline{\mathcal{M}}(T_i) / \text{Aut}(\{T_i\}_i / \#_i T_i) \xrightarrow{\sim} \overline{\mathcal{M}}(\#_i T_i). \quad (2.4)$$

We wish to warn the reader that (2.3) is *not* always an inclusion! The reason for this is explained in §2.4, in particular Figure 8.

2.4 Stratifications of moduli spaces

We now discuss the tautological stratifications of the moduli spaces $\overline{\mathcal{M}}(T)$ by $\mathcal{S}_{/T}$. This discussion is (slightly) less trivial than one might first expect, since $\mathcal{S}_{/T}$ is not a poset (see the discussion following Definition 2.2). A simple nontrivial example of a stratification by a category is given in Figure 8.

Definition 2.14. A *stratification* of a topological space X by a poset \mathcal{T} is a map $X \rightarrow \mathcal{T}$ for which $X^{\geq \mathfrak{t}}$ (the inverse image of $\mathcal{T}^{\geq \mathfrak{t}} = \mathcal{T}_{\geq \mathfrak{t}}$) is open for all $\mathfrak{t} \in \mathcal{T}$.

Definition 2.15. A *stratification* of a topological space X by a category \mathcal{T} (with the property that $\mathcal{T}_{\mathfrak{t}/}$ is a poset for all $\mathfrak{t} \in \mathcal{T}$) consists of the following data:

- (i) For every $x \in X$, we specify an object $\mathfrak{t}_x \in \mathcal{T}$.
- (ii) For every $x \in X$ and every y sufficiently close to x , we specify a morphism $f_{yx} : \mathfrak{t}_x \rightarrow \mathfrak{t}_y$ in \mathcal{T} (this data is regarded as a germ near x).

subject to the following requirements:

- (iii) For every $x \in X$, every y sufficiently close to x , and every z sufficiently close to y , we have $f_{zy} f_{yx} = f_{zx}$.
- (iv) For every $x \in X$, the induced germ of a map $X \rightarrow \mathcal{T}_{\mathfrak{t}_x/}$ near x is a stratification in the sense of Definition 2.14.

Given a stratification by a poset $X \rightarrow \mathcal{T}$, for any element $\mathfrak{t} \in \mathcal{T}$ there is a “locally closed stratum” $X^{\mathfrak{t}} \subseteq X$ and a “closed stratum” $X^{\leq \mathfrak{t}} \subseteq X$ equipped with a stratification $X^{\leq \mathfrak{t}} \rightarrow \mathcal{T}^{\leq \mathfrak{t}} = \mathcal{T}_{/\mathfrak{t}}$. Given a stratification by a category $X \rightarrow \mathcal{T}$, one can similarly define for any $\mathfrak{t} \in \mathcal{T}$ topological spaces $X_{\mathfrak{t}}$ and $X_{/\mathfrak{t}}$, the latter equipped with a stratification by $\mathcal{T}_{/\mathfrak{t}}$. The points of $X_{/\mathfrak{t}}$ are pairs $(x \in X, f_x : \mathfrak{t}_x \rightarrow \mathfrak{t})$, and a neighborhood of such a pair (x, f_x) consists of those (y, f_y) where $y \in X$ is close to x and $f_y f_{yx} = f_x$; the subspace $X_{\mathfrak{t}} \subseteq X_{/\mathfrak{t}}$ consists of those pairs (x, f_x) for which f_x is an isomorphism. In general, both natural maps $X_{\mathfrak{t}} \rightarrow X$ and $X_{/\mathfrak{t}} \rightarrow X$ may fail to be injective.




Figure 8: Example of a stratification by a category. The two maps out of the tree with three vertices collapse each of the two interior edges, respectively. The closed stratum associated to the tree with a single vertex is simply the entire space, homeomorphic to D^2 . The closed stratum associated to the tree with two vertices is homeomorphic to the interval $[0, 1]$. The closed stratum associated to the tree with three vertices is two points. Evidently, the latter two closed strata are “non-embedded”.

There is a tautological stratification

$$\overline{\mathcal{M}}(T) \rightarrow \mathcal{S}_{/T} \quad (2.5)$$

in the sense of Definition 2.15 (this follows immediately from the definition of the Gromov topology given in §2.9 below). The (not necessarily embedded!) closed stratum associated to $T' \rightarrow T$ is naturally isomorphic to $\overline{\mathcal{M}}(T')$. To understand this stratification, the reader should satisfy themselves that they understand the example in Figure 8.

2.5 Local models G

We introduce spaces $G_{T/}$ for $T \in \mathcal{S}$ which parameterize maps $\pi : T \rightarrow T'$ together with a collection of maps

$$\{\pi_* : \hat{X}_v \rightarrow \hat{X}_{\pi(v)}\}_{v \in V(T)} \quad (2.6)$$

(often defined only away from a small neighborhood of infinity in \hat{X}_v), where \hat{X}_v denotes the target of the v -component of a pseudo-holomorphic building of type T (despite the notation,

π_* is not determined uniquely by π). These spaces have a natural stratification $G_{T'} \rightarrow \mathcal{S}_{T'}$, and we denote by $G_{T'//T'}$ the inverse image of $\mathcal{S}_{T'//T'}$. There is a unique basepoint $0 \in G_{T'}$ corresponding to the identity map $T \rightarrow T$, and we only really care about a neighborhood of this basepoint in $G_{T'}$.

The definition of the Gromov topology on the moduli spaces $\overline{\mathcal{M}}(T)$ is based on these spaces $G_{T'//T}$, and we will see in Theorem 2.54 that the regular loci of $\overline{\mathcal{M}}(T)$ are in fact locally modelled on $G_{T'//T}$. To extract the virtual fundamental cycles of $\overline{\mathcal{M}}(T)$, a crucial step is to show that the stratifications $G_{T'//T} \rightarrow \mathcal{S}_{T'//T}$ are *cell-like* (see §3.2).

Remark 2.16. The spaces $G_{T'}$ considered here are model real toric singularities, as studied in Joyce [Joy16], Kottke–Melrose [KM15], and Gillam–Molcho [GM15] (though we do not appeal to any of their results). This is most apparent after the change of variables $h = e^{-g}$.

In Setup I, all targets \hat{X}_v are simply \hat{Y} , and we allow the map $\pi_* : \hat{Y} \rightarrow \hat{Y}$ to be any choice of \mathbb{R} -translation. However, since overall translations of \hat{Y} for $v \in V(T')$ induce isomorphisms of pseudo-holomorphic buildings, it is only the differences of the translation parameters which really matter. The result is that for $T \in \mathcal{S}_I$, we define

$$(G_I)_{T'} := (0, \infty]^{E^{\text{int}}(T)}. \quad (2.7)$$

There is a natural stratification $(G_I)_{T'} \rightarrow (\mathcal{S}_I)_{T'}$, sending $\mathbf{g} = \{\mathbf{g}_e\}_e$ to the map $\pi : T \rightarrow T'$ which contracts those edges $e \in E^{\text{int}}(T)$ for which $\mathbf{g}_e < \infty$. We denote by $0 \in (G_I)_{T'}$ the point corresponding to all gluing parameters equal to ∞ (i.e. corresponding to no gluing at all).

In Setup II, the targets \hat{X}_v are either \hat{X} or one of \hat{Y}^\pm , and the possible maps π_* are $\hat{Y}^\pm \rightarrow \hat{Y}^\pm$, $\hat{X} \rightarrow \hat{X}$, and $\hat{Y}^\pm \rightarrow \hat{X}$. Maps of the first type $\hat{Y}^\pm \rightarrow \hat{Y}^\pm$ are again allowed to be any \mathbb{R} -translation. Maps of the second type $\hat{X} \rightarrow \hat{X}$ must be the identity. Maps of the third type $\hat{Y}^\pm \rightarrow \hat{X}$ are the pre-composition of the relevant boundary collar (1.8)–(1.9) with any \mathbb{R} -translation of \hat{Y}^\pm . For $T \in \mathcal{S}_{II}$, the space of such collections of maps π_* is given by

$$(G_{II})_{T'} := \left\{ (\{\mathbf{g}_e\}_e, \{\mathbf{g}_v\}_v) \in (0, \infty]^{E^{\text{int}}(T)} \times (0, \infty]^{V_{00}(T)} \mid \mathbf{g}_v = \mathbf{g}_e + \mathbf{g}_{v'} \text{ for edges } v \xrightarrow{e} v' \text{ with } v \in V_{00}(T) \right\}. \quad (2.8)$$

We interpret $\mathbf{g}_{v'} = 0$ if $v' \notin V_{00}(T)$. Here $V_{ij}(T) \subseteq V(T)$ is the subset of v with $(\ast^+(v), \ast^-(v)) = (i, j)$. There is a natural stratification $(G_{II})_{T'} \rightarrow (\mathcal{S}_{II})_{T'}$, sending $\mathbf{g} = (\{\mathbf{g}_e\}_e, \{\mathbf{g}_v\}_v)$ to the map $\pi : T \rightarrow T'$ which contracts those edges e for which $\mathbf{g}_e < \infty$, and changes $\ast(v) = 00$ to $\ast(v) = 01$ for those vertices v with $\mathbf{g}_v < \infty$. We denote by $0 \in (G_{II})_{T'}$ the point corresponding to all gluing parameters equal to ∞ (i.e. corresponding to no gluing at all).

In Setup III, the situation is identical to that of Setup II, except that we also include variation in $t \in [0, 1]$. Namely, for $T \in \mathcal{S}_{III}$, we set

$$(G_{III})_{T'} := \left\{ (\{\mathbf{g}_e\}_e, \{\mathbf{g}_v\}_v) \in (0, \infty]^{E^{\text{int}}(T)} \times (0, \infty]^{V_{00}(T)} \mid \mathbf{g}_v = \mathbf{g}_e + \mathbf{g}_{v'} \text{ for } v \xrightarrow{e} v' \text{ with } v \in V_{00}(T) \right\} \times \begin{cases} [0, 1) & \mathfrak{s}(T) = \{0\} \\ (0, 1) & \mathfrak{s}(T) = (0, 1) \\ (0, 1] & \mathfrak{s}(T) = \{1\} \end{cases} \quad (2.9)$$

(the first factor is identical to (2.8)). There is a natural stratification $(G_{\text{III}})_{T'} \rightarrow (\mathcal{S}_{\text{III}})_{T'}$, sending $\mathbf{g} = (\{\mathbf{g}_e\}_e, \{\mathbf{g}_v\}_v, \mathbf{g}_t)$ to the map $\pi : T \rightarrow T'$ which contracts those edges e for which $\mathbf{g}_e < \infty$, changes $*(v) = 00$ to $*(v) = 01$ for those vertices v with $\mathbf{g}_v < \infty$, and has $\mathbf{g}_t \in \mathfrak{s}(T')$.

In Setup IV, if $T \in \mathcal{S}_{\text{IV}}$ satisfies $\mathfrak{s}(T) = \{0\}$ or $\mathfrak{s}(T) = (0, \infty)$, then the situation is identical to that of Setup III, so we set

$$(G_{\text{IV}})_{T'} := \left\{ (\{\mathbf{g}_e\}_e, \{\mathbf{g}_v\}_v) \in (0, \infty]^{E^{\text{int}}(T)} \times (0, \infty]^{V_{00}(T)} \mid \right. \\ \left. \mathbf{g}_v = \mathbf{g}_e + \mathbf{g}_{v'} \text{ for } v \xrightarrow{e} v' \text{ with } v \in V_{00}(T) \right\} \times \begin{cases} [0, \infty) & \mathfrak{s}(T) = \{0\} \\ (0, \infty) & \mathfrak{s}(T) = (0, \infty) \end{cases} \quad (2.10)$$

(the first factor is identical to (2.8)). The case $\mathfrak{s}(T) = \{\infty\}$ (i.e. the ‘‘split cobordism’’) is more complicated. The possible targets \hat{X}_v for a pseudo-holomorphic building of type T with $\mathfrak{s}(T) = \{\infty\}$ are $\hat{Y}^0, \hat{X}^{01}, \hat{Y}^1, \hat{X}^{12}, \hat{Y}^2$. If $\mathfrak{s}(T') = \{\infty\}$ as well, then the allowable maps are the same as in Setup II. If $\mathfrak{s}(T') = (0, \infty)$, then the targets \hat{X}_v are $\hat{Y}^0, \hat{X}^{02,t}, \hat{Y}^2$, for a choice of $t \in (0, \infty)$. The possible maps π_* are \mathbb{R} -translations $\hat{Y}^0 \rightarrow \hat{Y}^0$ and $\hat{Y}^2 \rightarrow \hat{Y}^2$, the natural maps $\hat{X}^{01} \rightarrow \hat{X}^{02,t}$ and $\hat{X}^{12} \rightarrow \hat{X}^{02,t}$ defined away from a neighborhood of infinity in the negative and positive ends, respectively, and the natural (up to \mathbb{R} -translation) maps $\hat{Y}^0 \rightarrow \hat{X}^{02,t}, \hat{Y}^1 \rightarrow \hat{X}^{02,t}, \hat{Y}^2 \rightarrow \hat{X}^{02,t}$. We conclude that for $T \in \mathcal{S}_{\text{IV}}$ with $\mathfrak{s}(T) = \{\infty\}$, we should set

$$(G_{\text{IV}})_{T'} := \left\{ (\{\mathbf{g}_e\}_e, \{\mathbf{g}_v\}_v, \mathbf{g}_t) \in (0, \infty]^{E^{\text{int}}(T)} \times (0, \infty]^{V_{00}(T) \sqcup V_{11}(T)} \times (0, \infty] \mid \right. \\ \left. \begin{array}{l} \mathbf{g}_v = \mathbf{g}_e + \mathbf{g}_{v'} \text{ for } v \xrightarrow{e} v', *(e) = 0, *(v) = 00 \\ \mathbf{g}_v = \mathbf{g}_e + \mathbf{g}_{v'} \text{ for } v \xrightarrow{e} v', *(e) = 1, *(v) = 11 \\ \mathbf{g}_t = \mathbf{g}_e + \mathbf{g}_{v'} \text{ for } v \xrightarrow{e} v', *(e) = 1, *(v) = 01 \end{array} \right\}. \quad (2.11)$$

We interpret $\mathbf{g}_{v'} = 0$ if it is undefined. There is a natural stratification $(G_{\text{IV}})_{T'} \rightarrow (\mathcal{S}_{\text{IV}})_{T'}$, which in the case $\mathfrak{s}(T) = \{\infty\}$ sends $\mathbf{g} = (\{\mathbf{g}_e\}, \{\mathbf{g}_v\}, \mathbf{g}_t)$ to the map $\pi : T \rightarrow T'$ which contracts those edges e for which $\mathbf{g}_e < \infty$, increments $*^-(v)$ for those vertices v with $\mathbf{g}_v < \infty$, and has $\mathbf{g}_t \in \mathfrak{s}(T')$.

2.6 Deformation theory of Riemann surfaces

Given a closed Riemann surface C (without nodes), infinitesimal deformations of the almost complex structure are given by $C^\infty(C, \text{End}^{0,1}(TC))$. An infinitesimal diffeomorphism of C , that is a vector field X , gives rise to an infinitesimal deformation $\mathcal{L}_X j$, where \mathcal{L} denotes the Lie derivative. Noting that $\text{End}^{0,1}(TC) = TC \otimes_{\mathbb{C}} \Omega_C^{0,1}$, we are thus lead to consider the two-term complex

$$C^\infty(C, TC) \xrightarrow{X \mapsto \mathcal{L}_X j} C^\infty(C, TC \otimes \Omega_C^{0,1}), \quad (2.12)$$

which is simply the Dolbeaut complex calculating $H^\bullet(C, TC)$. Thus $H^0(C, TC)$ classifies infinitesimal automorphisms of C (though this is clear *a priori*) and $H^1(C, TC)$ classifies infinitesimal deformations of C .

Now suppose C has (a finite number of distinct) marked points p_i . Infinitesimal deformations of the almost complex structure and the marked points are given by $C^\infty(C, \text{End}^{0,1}(TC)) \oplus \bigoplus_i T_{p_i}C$, and the infinitesimal deformation associated to a vector field X is $\mathcal{L}_X j \oplus \bigoplus_i (-X(p_i))$. This gives rise to the two-term complex

$$C^\infty(C, TC) \xrightarrow{X \mapsto \mathcal{L}_X j \oplus \bigoplus_i (-X(p_i))} C^\infty(C, TC \otimes \Omega_C^{0,1}) \oplus \bigoplus_i T_{p_i}C, \quad (2.13)$$

which calculates $H^\bullet(C, TC(-P))$, where $P = \sum_i p_i$ denotes the divisor of marked points. It is somewhat more convenient for us to consider just those variations in the complex structure which are supported away from the marked points. That is, we consider the subcomplex

$$\left\{ X \in C^\infty(C, TC) \left| \begin{array}{l} X \text{ holomorphic near } p_i \\ X(p_i) = 0 \end{array} \right. \right\} \xrightarrow{X \mapsto \mathcal{L}_X j} C_c^\infty(C \setminus \{p_i\}, TC \otimes \Omega_C^{0,1}). \quad (2.14)$$

The inclusion of this subcomplex into (2.13) is a quasi-isomorphism (this is obvious except for surjectivity of the map on cokernels, which follows from elliptic regularity).

Now suppose C has both marked points p_i and “doubly marked points” q_i , meaning q_i are marked points equipped with markings $\mathbb{C} \xrightarrow{\sim} T_{q_i}C$ of their tangent fibers. The analogue of (2.14) is now

$$\left\{ X \in C^\infty(C, TC) \left| \begin{array}{l} X \text{ holomorphic near } p_i, q_i \\ X(p_i) = 0 \\ X(q_i) = dX(q_i) = 0 \end{array} \right. \right\} \xrightarrow{X \mapsto \mathcal{L}_X j} C_c^\infty(C \setminus \{p_i\} \cup \{q_i\}, TC \otimes \Omega_C^{0,1}), \quad (2.15)$$

which calculates $H^\bullet(C, TC(-P - 2Q))$, where $P = \sum_i p_i$ and $Q = \sum_i q_i$.

Finally, when C has nodes, we simply treat these as pairs of marked points on its normalization \tilde{C} and consider the deformation theory of \tilde{C} instead.

Definition 2.17 (Normalization). For any nodal Riemann surface C , we denote by \tilde{C} the *normalization* of C , i.e. the unique smooth Riemann surface equipped with a map $\tilde{C} \rightarrow C$ which identifies points in pairs to form the nodes of C .

2.7 Moduli of Riemann surfaces

We denote by $\overline{\mathcal{M}}_{0,n}$ the Deligne–Mumford moduli space [DM69] of stable genus zero compact nodal Riemann surfaces with n marked points. We regard $\overline{\mathcal{M}}_{0,n}$ as a complex manifold (which is compact). It comes equipped with a universal curve $\overline{\mathcal{C}}_{0,n} \rightarrow \overline{\mathcal{M}}_{0,n}$ (which is secretly just the map forgetting the last marked point $\overline{\mathcal{M}}_{0,n+1} \rightarrow \overline{\mathcal{M}}_{0,n}$). The pair $\overline{\mathcal{C}}_{0,n} \rightarrow \overline{\mathcal{M}}_{0,n}$ can be defined as the solution to a moduli problem, and it is thus unique up to unique isomorphism.

There is a stratification of $\overline{\mathcal{M}}_{0,n}$ by locally closed analytic submanifolds $\overline{\mathcal{M}}_{0,n}^{\#\text{nodes}=r}$ indexed by integers $r \geq 0$ consisting of those curves with exactly r nodes. The open stratum of smooth curves $\overline{\mathcal{M}}_{0,n}^{\#\text{nodes}=0}$ is usually denoted $\mathcal{M}_{0,n}$. Its complement $\overline{\mathcal{M}}_{0,n}^{\#\text{nodes} \geq 1}$ is a divisor with normal crossings, namely the pair $\overline{\mathcal{M}}_{0,n}^{\#\text{nodes} \geq 1} \subseteq \overline{\mathcal{M}}_{0,n}$ is locally biholomorphic to $\{z_1 \cdots z_k = 0\} \subseteq \mathbb{C}^N$ (any $k \leq N$), and in such local coordinates, the “number of nodes” is simply the number of coordinates z_1, \dots, z_k which vanish.

If C is any point of $\mathcal{M}_{0,n}$ (i.e. any stable smooth Riemann surface of genus zero with n marked points), then there is a canonical isomorphism

$$T_C\mathcal{M}_{0,n} = H^1(C, TC(-P)). \quad (2.16)$$

The map $T_C\mathcal{M}_{0,n} \rightarrow H^1(C, TC(-P))$ is the Kodaira–Spencer map, which can be defined by (smoothly) trivializing the universal curve in a neighborhood of C and invoking the deformation theory reviewed in §2.6. This map $T_C\mathcal{M}_{0,n} \rightarrow H^1(C, TC(-P))$ is well-known to be an isomorphism (see, for instance, Wendl [Wen14, Theorem 4.30]). For $C \in \overline{\mathcal{M}}_{0,n}$ with r nodes, by considering the normalization \tilde{C} equipped with the inverse images of the marked points $\tilde{P} \subseteq \tilde{C}$ and the inverse images of the nodes $\tilde{N} \subseteq \tilde{C}$, we deduce a similar isomorphism

$$T_C\mathcal{M}_{0,n}^{\#\text{nodes}=r} = H^1(\tilde{C}, T\tilde{C}(-\tilde{P} - \tilde{N})). \quad (2.17)$$

2.8 Elliptic *a priori* estimates on pseudo-holomorphic curves

We record here some fundamental *a priori* estimates which guarantee the regularity of pseudo-holomorphic curves. These estimates play a fundamental role in establishing the basic local properties of moduli spaces of pseudo-holomorphic curves.

The first result implies that for pseudo-holomorphic curves, C^0 -close implies C^∞ -close.

Lemma 2.18 (Gromov [Gro85]). *Let $u : D^2(1) \rightarrow (B^{2n}(1), J)$ be J -holomorphic, where J is tamed by $d\lambda$. For every $k < \infty$, we have*

$$\|u\|_{C^k(D^2(1-s))} \leq M \cdot (1 + s^{-k}) \quad (2.18)$$

for some constant $M = M(k, J, \lambda) < \infty$.

Proof. The Gromov–Schwarz Lemma (see Gromov [Gro85, 1.3.A] or Muller [Mul94, Corollary 4.1.4]) is the case $k = 1$. Standard elliptic bootstrapping allows one to upgrade this to bounds on all higher derivatives (see [Par16, Lemma B.11.4]). Note that it is enough to bound $D^k u$ at $0 \in D^2(1)$; the full result is then recovered by restricting to the maximal disk centered at a given point in $D^2(1)$. \square

The next two “long cylinder estimates” are crucial for proving compactness of moduli spaces (which we do not discuss) and surjectivity of the gluing map (which we discuss in detail). Note that both of these next results imply corresponding results for semi-infinite cylinders $[0, \infty) \times S^1$, simply by applying the results as stated to $[0, N] \times S^1 \subseteq [0, \infty) \times S^1$ and letting $N \rightarrow \infty$.

Lemma 2.19 (Well-known). *Let $u : [0, N] \times S^1 \rightarrow (B^{2n}(1), J)$ be J -holomorphic, where J is tamed by $d\lambda$. For every $k < \infty$ and $\varepsilon > 0$, we have*

$$\|u\|_{C^k([s, N-s] \times S^1)} \leq M \cdot (1 + s^{-k})e^{-(1-\varepsilon)s} \quad (2.19)$$

for some constant $M = M(k, \varepsilon, J, \lambda) < \infty$.

Proof. See [MS04] or [Par16, Proposition B.11.1]. \square

Proposition 2.20 (Hofer–Wysocki–Zehnder [HWZ02, Theorems 1.1, 1.2, and 1.3]). *Let (Y, λ, J) be as in Setup I. Suppose all periodic orbits of R_λ of action $\leq E_0 < \infty$ are non-degenerate, let $\sigma > 0$ denote the minimum gap between any two elements of $\{0\} \cup a(\mathcal{P}) \cap [0, E_0]$, and let $\varepsilon > 0$. For any \hat{J} -holomorphic map $u : [0, N] \times S^1 \rightarrow \hat{Y}$ satisfying*

$$\sup_{\substack{\varphi: \mathbb{R} \rightarrow [0, 1] \\ \varphi'(s) \geq 0}} \int_{[0, N] \times S^1} u^* d(\varphi \lambda) \leq E_0, \quad (2.20)$$

$$\int_{[0, N] \times S^1} u^* d\lambda \leq \sigma - \varepsilon, \quad (2.21)$$

at least one of the following holds for $H = H(Y, \lambda, J, E_0, \varepsilon) < \infty$:

- $u([H, N - H] \times S^1) \subseteq B_\varepsilon(u(\frac{N}{2}, 0))$.
- There exists $\gamma \in \mathcal{P}$ with $a(\gamma) = L \leq E_0$ such that for $s \geq H$,

$$\|u(s, t) - (Ls + b, \tilde{\gamma}(t))\|_{C^0([s, N-s] \times S^1)} \leq H \cdot e^{-(\delta_\gamma - \varepsilon) \min(s-H, N-H-s)} \quad (2.22)$$

for some $b \in \mathbb{R}$ and some $\tilde{\gamma} : S^1 \rightarrow Y$ with $\partial_t \tilde{\gamma} = L \cdot R_\lambda(\tilde{\gamma})$ parameterizing $\gamma \in \mathcal{P}$, where $\delta_\gamma > 0$ denotes the smallest absolute value of any eigenvalue of the asymptotic operator of γ (see Definition 2.29).

The meaning of the energy bounds (2.20)–(2.21) is discussed further in §2.10.

Remark 2.21. Our statement of Proposition 2.20 is slightly more precise than that given by Hofer–Wysocki–Zehnder [HWZ02] in that they state it with an unspecified constant $\delta > 0$ in place of $\delta_\gamma > 0$. We emphasize, however, that knowing the correct constant is not needed for our work in this paper.

2.9 Gromov topology

We now define the Gromov topology on the moduli spaces $\overline{\mathcal{M}}(T)$, and we show that this topology is Hausdorff.

Recall that a *neighborhood* of a point x in a topological space X is a subset $N \subseteq X$ such that there exists an open set $U \subseteq X$ with $x \in U \subseteq N$. The *neighborhood filter* of x is the collection $\mathcal{N}_x \subseteq 2^X$ of all its neighborhoods. A *neighborhood base* at x is a cofinal subset $\mathcal{N}_x^\circ \subseteq \mathcal{N}_x$ (meaning, for every $N \in \mathcal{N}_x$ there exists $N' \in \mathcal{N}_x^\circ$ with $N' \subseteq N$). Given a neighborhood base \mathcal{N}_x° , we may recover \mathcal{N}_x as the collection of all subsets of X containing some element of \mathcal{N}_x° .

Lemma 2.22. *The system of neighborhood filters $\{\mathcal{N}_x \subseteq 2^X\}_{x \in X}$ at all the points of a topological space X satisfies the following properties:*

- (i) For every $N \in \mathcal{N}_x$, we have $x \in N$.
- (ii) For every $N \in \mathcal{N}_x$ and every $N' \supseteq N$, we have $N' \in \mathcal{N}_x$.
- (iii) For every $N, N' \in \mathcal{N}_x$, we have $N \cap N' \in \mathcal{N}_x$.
- (iv) For every $N \in \mathcal{N}_x$, there exists $N' \in \mathcal{N}_x$ such that $N \in \mathcal{N}_y$ for every $y \in N'$.

Conversely, any system of subsets $\{\mathcal{N}_x \subseteq 2^X\}_{x \in X}$ satisfying the above properties is the system of neighborhood filters for a unique topology on X . \square

We define the Gromov topology by specifying the neighborhood filter \mathcal{N}_x of any given pseudo-holomorphic building x and appealing to Lemma 2.22. More precisely, (1) we define the “ ε -neighborhood of x ” for any (sufficiently small) $\varepsilon > 0$, given certain choices of auxiliary data (metrics, etc.), (2) we define the neighborhood filter \mathcal{N}_x of x to be that resulting from declaring that the collection of all ε -neighborhoods (with respect to fixed auxiliary data) forms a neighborhood base \mathcal{N}_x° at x , and (3) we observe that \mathcal{N}_x is independent of the choice of auxiliary data.

Definition 2.23 (Gromov topology). The ε -neighborhood of a pseudo-holomorphic building $x = (\{C_v\}, \{p_{v,e}\}, \{u_v\}, \{\tilde{b}_e\}, \{m_e\})$ of type T is defined as follows. Let $\mathbf{g} \in G_{T'}$ be ε -close to the basepoint $0 \in G_{T'}$. Recall from §2.5 that \mathbf{g} determines a map $\pi : T \rightarrow T'$ together with a collection of (possibly only partially defined) maps

$$\pi_* : \hat{X}_v \rightarrow \hat{X}_{\pi(v)} \quad (2.23)$$

for $v \in V(T)$. Fix (independently of $\varepsilon > 0$) large compact subsets $\hat{X}_v^c \subseteq \hat{X}_v$. Consider the curves

$$C_{v'} := \bigsqcup_{\pi(v)=v'} C_v / \sim \quad (2.24)$$

for $v' \in V(T')$, where $p_{v_1,e} \sim p_{v_2,e}$ for edges $v_1 \xrightarrow{e} v_2$ with $\pi(v_1) = \pi(v_2) = v'$. These curves $C_{v'}$ are equipped with points $\{p_{v',e'}\}$ for edges $e' \in E(T')$ incident at v' , namely the images of the corresponding points $p_{v,e} \in C_v$ for $\pi(e) = e'$. Let $C'_{v'}$ be any curve obtained from $C_{v'}$ by the following operations:

- (i) An ε - C^∞ -small modification of the almost complex structure away from the ε -neighborhood of the points $\{p_{v',e'}\}$ and the nodes of $C_{v'}$.
- (ii) An arbitrary modification of the almost complex structure over the ε -neighborhood of the points $\{p_{v',e'}\}$ and the nodes of $C_{v'}$.
- (iii) Resolving the nodes of $C_{v'}$ created by the identifications $p_{v_1,e} \sim p_{v_2,e}$, and possibly resolving the other nodes of $C_{v'}$ (those coming from the nodes of C_v for $\pi(v) = v'$). (Resolving a node means that its ε -neighborhood $D^2 \vee D^2$ gets replaced with a cylinder $[0, 1] \times S^1$, equipped here with an arbitrary almost complex structure).

The locus referred to in (i) will be called the ε -thick part of $C'_{v'}$, and the locus referred to in (ii)–(iii) will be called the ε -thin part of $C'_{v'}$. The ε -thin part of $C'_{v'}$ is broken down further into the nodal ε -thin part (that coming from nodes of C_v) and the Reeb ε -thin part (that coming from the points $\{p_{v,e}\}$ of C_v , both those identified to form nodes in $C_{v'}$ and those which become the points $\{p_{v',e'}\}$). Now let $x' = (\{C'_{v'}\}, \{p_{v',e'}\}, \{u'_{v'}\}, \{\tilde{b}_{e'}\}, \{m_{e'}\})$ be a pseudo-holomorphic building of type T' , upon which we impose the following requirements. Most basically, we require that u' be ε - C^0 -close to u (measured with respect to fixed metrics on $\hat{X}_{v'}$ which are \mathbb{R} -invariant in the ends) in the following sense:

- (iv) Over the ε -thick part of $C'_{v'}$, we require $u'_{v'}$ to be ε - C^0 -close to $\pi_* \circ \bigsqcup_{\pi(v)=v'} u_v$.
- (v) Over the nodal ε -thin part of $C'_{v'}$, we require $u'_{v'}$ to map into the ε -ball around the image under $\pi_* \circ u_v$ of the corresponding node of C_v .

- (vi) Over the Reeb ε -thin part of $C'_{v'}$ coming from an edge $v_1 \xrightarrow{e} v_2$ with $\pi(v_1) = \pi(v_2) = v'$, we require $u'_{v'}$ to map into the ε -neighborhood of the trivial cylinder over γ_e inside the finite cylindrical region in $\hat{X}_{v'}$ in between $\pi_*(\hat{X}_{v_1}^c)$ and $\pi_*(\hat{X}_{v_2}^c)$ (recall the large compact subsets $\hat{X}_v^c \subseteq \hat{X}_v$ fixed above).
- (vii) Over the Reeb ε -thin part of $C'_{v'}$ near a positive (respectively, negative) puncture $p_{v',e'}$ associated to an edge $v_1 \xrightarrow{e} v_2$ not collapsed by π with $\pi(v_2) = v'$ (respectively, $\pi(v_1) = v'$), we require $u'_{v'}$ to map into the ε -neighborhood of the trivial cylinder over γ_e inside the half-infinite cylindrical region in $\hat{X}_{v'}$ above $\pi_*(\hat{X}_{v_2}^c)$ (respectively, below $\pi_*(\hat{X}_{v_1}^c)$).

Note that the *a priori* estimates of Hofer–Wysocki–Zehnder [HWZ02, Theorems 1.1, 1.2, and 1.3] recalled in Proposition 2.20 imply that (vi)–(vii) imply that over the Reeb ε -thin parts of $C'_{v'}$, the map $u'_{v'}$ actually converges very rapidly to a trivial cylinder (to apply Proposition 2.20, we should check that the integral of $d\lambda$ over these ε -thin parts is small, which can be seen by using Lemma 2.18 over the nearby ε -thick part and applying Stokes' theorem). It follows that our choice of fixed large compact subsets $\hat{X}_v^c \subseteq \hat{X}_v$ does not matter. This convergence to trivial cylinders also allows us to make sense of the following final requirements on the “discrete data” of x' :

- (viii) The asymptotic markers $\{\tilde{b}_{e'}\}$ and matching isomorphisms $\{m_{e'}\}$ of x' must agree with those descended from x .
- (ix) Over the glued edges $e \in E(T)$ (meaning those collapsed by π), the matching isomorphism induced by $u'_{v'}$ must agree with m_e .

The ε -neighborhood of x is defined as the set of all possible x' above. It is straightforward to check that the resulting neighborhood filter \mathcal{N}_x at x is independent of the choices involved. We leave it as an exercise to check that these neighborhood filters satisfy the conditions in Lemma 2.22 (only (iv) requires thought). We thus obtain a well-defined topology on each of the spaces $\overline{\mathcal{M}}(T)$, which we call the Gromov topology.

Remark 2.24. There are various other ways to define the curve $C'_{v'}$ and its identification with $C_{v'}$ away from their ε -thin parts. For example, one could add stabilizing marked points to $C_{v'}$ and then take $C'_{v'}$ to be a nearby fiber of the universal curve over Deligne–Mumford space. It is an exercise to check that this definition of $C'_{v'}$ gives rise to the same neighborhood filter.

Remark 2.25. A sequence (or net) x_i of pseudo-holomorphic buildings converges in the Gromov topology to a pseudo-holomorphic building x iff for every $\varepsilon > 0$, it holds that x_i lies in the ε -neighborhood of x for all sufficiently large i . Indeed, given any neighborhood base \mathcal{N}_x° at a point x of a topological space, a sequence (or net) x_i converges to x iff for every $N \in \mathcal{N}_x^\circ$, we have $x_i \in N$ for all sufficiently large i .

Lemma 2.26. *The Gromov topology on $\overline{\mathcal{M}}(T)$ is Hausdorff.*

Proof. Let $x^{(1)}$ and $x^{(2)}$ be two stable pseudo-holomorphic buildings such that for every $\varepsilon > 0$, there exists a pseudo-holomorphic building x^ε lying in the ε -neighborhood of both $x^{(1)}$ and $x^{(2)}$. Let $C^{(1)} := \bigsqcup_v C_v^{(1)} / \sim$ (where $p_{v_1,e} \sim p_{v_2,e}$ for edges $v_1 \xrightarrow{e} v_2$ as usual), and similarly define $C^{(2)}$ and C^ε .

Choose finite collections of points $P^{(1)} \subseteq C^{(1)}$ and $P^{(2)} \subseteq C^{(2)}$ (disjoint from the nodes and punctures) which stabilize each. By Remark 2.24, the curve C^ε may be equipped with marked points $P^{(1),\varepsilon}$ and $P^{(2),\varepsilon}$ (respectively) so that it is ε -close to $(C^{(1)}, P^{(1)})$ and $(C^{(2)}, P^{(2)})$ (respectively) in Deligne–Mumford space. We consider the curve $(C^\varepsilon, P^{(1),\varepsilon} \sqcup P^{(2),\varepsilon})$ obtained by adding both $P^{(1),\varepsilon}$ and $P^{(2),\varepsilon}$ as additional marked points (note that, without loss of generality, we may perturb either $P^{(1),\varepsilon}$ or $P^{(2),\varepsilon}$ so that they are disjoint).

Since Deligne–Mumford space is compact, we may assume that $(C^\varepsilon, P^{(1),\varepsilon} \sqcup P^{(2),\varepsilon})$ converges to some stable $(C, P^{(1)} \sqcup P^{(2)})$ as $\varepsilon \rightarrow 0$. Since forgetting marked points and stabilizing is continuous, we obtain canonical identifications

$$i_1 : (C, P^{(1)})^{\text{st}} \xrightarrow{\sim} (C^{(1)}, P^{(1)}), \quad (2.25)$$

$$i_2 : (C, P^{(2)})^{\text{st}} \xrightarrow{\sim} (C^{(2)}, P^{(2)}). \quad (2.26)$$

Any irreducible component $C_0 \subseteq C$ contracted by i_1 but not by i_2 would contradict stability of $x^{(2)}$. Thus every irreducible component $C_0 \subseteq C$ contracted by either i_1 or i_2 is contracted by both. We conclude that

$$i^{21} := i_2 \circ i_1^{-1} : C^{(1)} \xrightarrow{\sim} C^{(2)} \quad (2.27)$$

is an isomorphism and that $u^{(2)} \circ i^{21} = u^{(1)}$.

To conclude that i^{21} defines an isomorphism $x^{(1)} \rightarrow x^{(2)}$, it remains only to observe that the asymptotic markers and matching isomorphisms also agree as a consequence of (viii)–(ix) in Definition 2.23. \square

2.10 Compactness

We now recall the compactness results due to Bourgeois–Eliashberg–Hofer–Wysocki–Zehnder [BEHWZ03] in the context of the moduli spaces $\overline{\mathcal{M}}(T)$ which we have defined.

We begin by recalling from [BEHWZ03] the notion of *Hofer energy* first introduced in [Hof93]. Note that the usual notion of energy of a pseudo-holomorphic map (namely the integral of the pullback of the symplectic form) is not well-behaved when the target symplectic manifold is a symplectization \hat{Y} equipped with the symplectic form $\hat{\omega} := d\hat{\lambda} = d(e^s \lambda) = e^s(d\lambda + ds \wedge \lambda)$, as this energy is not invariant under \mathbb{R} -translation and is almost always infinite. We therefore remove the e^s factor, leaving $d\lambda + ds \wedge \lambda$, and we further treat each of these terms separately (and differently). The resulting notions of energy we now discuss are (beyond the case of symplectizations) only well-defined up to an overall multiplicative constant.

The ω -energy of a pseudo-holomorphic map $u : C \rightarrow \hat{Y}$ is the integral

$$E_\omega(u) := \int_C u^* \omega \quad \text{where } \omega := d\lambda. \quad (2.28)$$

To define the ω -energy of a pseudo-holomorphic map to a symplectic cobordism $u : C \rightarrow \hat{X}$, we integrate the closed 2-form ω defined by splicing together $\hat{\omega}$ with $\omega^\pm = d\lambda^\pm$ in the ends; more explicitly, in the positive end $\omega := d(\alpha(s)\lambda^+)$ for $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $\alpha(s) = e^s$ for $s \ll 0$, $\alpha'(s) = 0$ for $s \gg 0$, and $\alpha'(s) \geq 0$, and similarly in the negative end. To define the ω -energy of a pseudo-holomorphic map to an almost split symplectic cobordism $\hat{X}^{02,t}$ with

$t \rightarrow \infty$ as in case (IV), we integrate the closed 2-form $\omega^{02,t}$ defined as the descent of $\omega^{01} \sqcup \omega^{12}$ from $\hat{X}^{01} \sqcup \hat{X}^{12}$; over the identified regions these 2-forms agree up to a scaling factor which is fixed as $t \rightarrow \infty$ (and hence irrelevant).

The ω -energy of a pseudo-holomorphic building is simply the sum of the ω -energies of each of its components.

The integrand of the ω -energy is pointwise ≥ 0 for pseudo-holomorphic u , with equality iff $du = 0$ or u is tangent to the vertical distribution $\mathbb{R}\partial_s$ at a point of a symplectization region where ω is (up to a constant) $d\lambda$. If u is asymptotic to Reeb orbits at positive/negative infinity, then the ω -energy is the pairing between the relative homology class of u and the relative cohomology class of ω (both relative to the positive/negative asymptotics of u).

The λ -energy of a pseudo-holomorphic map $u : C \rightarrow \hat{Y}$ is the supremum

$$E_\lambda(u) := \sup_{\substack{\phi: \mathbb{R} \rightarrow [0,1] \\ \phi'(s) \geq 0 \\ \phi(s)=0 \ s \ll 0 \\ \phi(s)=1 \ s \gg 0}} \int_C u^*(d\phi \wedge \lambda). \quad (2.29)$$

The λ -energy of a pseudo-holomorphic map to a symplectic cobordism $u : C \rightarrow \hat{X}$ is the supremum of $\int_C u^*(d\phi \wedge \lambda^\pm)$ over all $\phi : \hat{X} \rightarrow [0, 1]$ which are constant away from the ends (1.8)–(1.9) and which in each of these ends are as above, namely they are functions of s satisfying $\phi'(s) \geq 0$ and $\phi(s) = 0$ for $s \ll 0$ and $\phi(s) = 1$ for $s \gg 0$. To define the λ -energy of a pseudo-holomorphic map to an almost split symplectic cobordism $\hat{X}^{02,t}$ for large t as in case (IV), we also allow ϕ to be non-constant over the middle finite symplectization region of $\hat{X}^{02,t}$ where the integrand is $d\phi \wedge \lambda^1$.

To define the λ -energy of a pseudo-holomorphic building of type T , we take the supremum (of the analogous sum of integrals) over all choices of $\{\phi_v : \hat{X}_v \rightarrow [0, 1]\}_v$ as follows. Each ϕ_v must be constant away from the symplectization regions, and over each symplectization region ϕ_v must be a function of s with $\phi'_v(s) \geq 0$. Finally, each ϕ_v must be constant in a neighborhood of positive/negative infinity; these constants must agree across interior edges of T and must equal 1 and 0 for input/output edges of T respectively.

The integrand of the λ -energy is pointwise ≥ 0 for pseudo-holomorphic u . An argument involving Stokes' theorem shows that any candidate value $\int_C u^*(d\phi \wedge \lambda)$ of the λ -energy is equal to (a constant times) any other candidate value, up to an error bounded above by (a constant times) the ω -energy.

The *Hofer energy* of a pseudo-holomorphic map or building is the sum of the ω -energy and the λ -energy:

$$E(u) := E_\omega(u) + E_\lambda(u). \quad (2.30)$$

Up to a multiplicative constant, the Hofer energy depends only on the asymptotics and homology class of the map/building u (for buildings, by this we mean only γ_{e^+} , $\{\gamma_{e^-}\}$, and $\#_v \beta_v$, not any of the intermediate Reeb orbits or the individual β_v); compare [BEHWZ03, Propositions 5.13 and 6.3]. Indeed, we have already seen above that the ω -energy depends only on the asymptotics and homology class of u . We have also seen above that the various candidate λ -energies differ by at most the ω -energy, so for the purposes of defining the Hofer energy, we may as well replace the λ -energy term with one specific such candidate value, such as the action of the positive asymptotic Reeb orbit or the sum of the actions of the

negative asymptotic Reeb orbits. We conclude that for pseudo-holomorphic buildings u of type T , we have

$$E(u) \asymp a(\gamma_{e^+}) + \sum_{e^-} a(\gamma_{e^-}) + \langle \omega, \#_v \beta_v \rangle \quad (2.31)$$

(equality up to a multiplicative constant). Finally, note that when we consider families of cobordisms, the last term on the right hand side may be trivially bounded above by $\sup_t \langle \omega^t, \#_v \beta_v \rangle$ in case (III) and by $\sup_t \langle \omega^{02,t}, \#_v \beta_v \rangle$ in case (IV) (note that the cohomology class of $\omega^{02,t}$ is independent of t for t sufficiently large, so this supremum is finite). Define $E(T)$ as the right hand side of (2.31) with $\langle \omega, \#_v \beta_v \rangle$ replaced with the relevant supremum in cases (III) and (IV), and note that $E(T) = E(T')$ for any morphism $T \rightarrow T'$.

Theorem 2.27 ([BEHWZ03, Theorems 10.1, 10.2, 10.3]). *For any fixed datum \mathcal{D} as in Setup I–IV, each moduli space $\overline{\mathcal{M}}(T)$ is compact, and for every $E < \infty$ there are at most finitely many (isomorphism classes of) $T \in \mathcal{S}$ with $\overline{\mathcal{M}}(T)$ non-empty and $E(T) \leq E$. In particular, given $T \in \mathcal{S}$ there are at most finitely many (isomorphism classes of) $T' \rightarrow T$ with $\overline{\mathcal{M}}(T')$ non-empty.*

The arguments of [BEHWZ03, §10] show that every net of pseudo-holomorphic buildings of Hofer energy $\leq E$ has a convergent subnet, which proves Theorem 2.27. Although [BEHWZ03, §10] is presented as a proof of sequential compactness, it applies essentially without change to yield the true compactness result we claim here. The fact that the moduli spaces considered in [BEHWZ03] differ slightly from those considered here (there is a surjective forgetful map from the moduli spaces in [BEHWZ03] to ours) also makes no serious difference to the argument.

2.11 Linearized operators

We now recall the relevant linearized operators associated to the pseudo-holomorphic curves that we consider.

Definition 2.28 (Linearized operators of Reeb orbits). For any $\gamma \in \mathcal{P}$ and a choice of parameterization $\tilde{\gamma} : S^1 \rightarrow Y$, denote by D_γ^N the operator $J\mathcal{L}_{R_\lambda}$ (where \mathcal{L} denotes the Lie derivative) acting on sections of $\tilde{\gamma}^* \xi$ over S^1 , and define $D_\gamma := D_\gamma^N \oplus i \cdot d$ for $i \cdot d$ acting on functions $S^1 \rightarrow \mathbb{C}$.

The operator D_γ^N defines a Fredholm map $H^s(S^1, \tilde{\gamma}^* \xi) \rightarrow H^{s-1}(S^1, \tilde{\gamma}^* \xi)$ for any $s \in \mathbb{R}$ (either by Fourier analysis or by appealing to general elliptic regularity results). By the analytic Fredholm theorem, the resolvent $(D_\gamma^N - zI)^{-1} : H^{s-1}(S^1, \tilde{\gamma}^* \xi) \rightarrow H^s(S^1, \tilde{\gamma}^* \xi)$ varies meromorphically in $z \in \mathbb{C}$, with poles corresponding to the eigenvalues $\sigma(D_\gamma^N)$ of D_γ^N . Since D_γ^N is formally self-adjoint, all these eigenvalues are real. The same discussion applies to $i \cdot d$ and to D_γ , and

$$\sigma(D_\gamma) = \sigma(D_\gamma^N) \cup \sigma(i \cdot d) = \sigma(D_\gamma^N) \cup \mathbb{Z}. \quad (2.32)$$

Clearly $0 \notin \sigma(D_\gamma^N)$ iff γ is non-degenerate (as is our standing assumption).

Definition 2.29. Denote by $\delta_\gamma > 0$ the smallest absolute value of any element of $\sigma(D_\gamma) \setminus \{0\}$.

Convention 2.30 (Choices of metric and connection on domains). For the purposes of defining function spaces, stating estimates, etc. involving a Riemann surface C equipped with punctures $\{p_e\}_e$, we use a choice of holomorphic cylindrical coordinates

$$[0, \infty) \times S^1 \rightarrow C \setminus p_e \quad (2.33)$$

near each p_e . We equip C with a Riemannian metric which equals $ds^2 + dt^2$ near p_e , and we equip TC with a j -linear connection for which ∂_s is parallel near p_e . Different choices of this data will result in uniformly commensurable norms and estimates, so the particular choice is not important.

Convention 2.31 (Choices of metric and connection on targets). For the purposes of defining function spaces, stating estimates, expressing linearized operators, etc. involving symplectizations \hat{Y} , we use any choice of \mathbb{R} -invariant Riemannian metric on \hat{Y} and any choice of connection on $T\hat{Y}$ which is \hat{J} -linear (meaning $\hat{J}\nabla_X Y = \nabla_X \hat{J}Y$, i.e. $\nabla \hat{J} = 0$) and which is pulled back from a connection on $T\hat{Y}|_{\{0\} \times Y} = TY \oplus \mathbb{R}$. On symplectic cobordisms, use metrics and connections which are of this form in the positive/negative ends. Different choices of metric and connection will always result in uniformly commensurable norms, so the particular choice is not important.

Convention 2.32 (Regularity classes of functions on nodal curves). A function on a nodal curve C of a certain regularity class (C^∞ , $W^{k,2}$, etc.) simply means a function on its normalization \tilde{C} (see Definition 2.17) of the given regularity class which agrees across each pair of points identified under $\tilde{C} \rightarrow C$ to form the nodes of C . We will only speak of functions on nodal C in regularity classes which embed into C^0 , so that the above makes sense. Recall, in particular, that $W^{k,2} \subseteq C^0$ when $k \geq 2$ [Ada75, Lemma 5.17].

We refer to Adams [Ada75] for basic properties and definitions of Sobolev spaces, including here only a brief discussion. Recall that for integers $k \geq 0$, a function is said to be of class $W^{k,2}$ iff all its derivatives (in the sense of distributions) of order $\leq k$ are in L^2 , and the $W^{k,2}$ -norm of such a function is defined as the sum of the L^2 -norms of these derivatives. A mollification argument [Ada75, Theorem 3.16] of Meyers–Serrin shows that any compactly supported $W^{k,2}$ function can be approximated in the $W^{k,2}$ -topology by a sequence of smooth functions supported in a neighborhood of the original support.

It is well-known that it is frequently useful to introduce “weights” into the definition of Sobolev spaces of functions on certain classes of non-compact manifolds. We now recall the weighted Sobolev spaces relevant to our setting of pseudo-holomorphic curves with cylindrical ends.

Remark 2.33. We frame our discussion for general k and δ so as to be precise as possible regarding the constraints on k and δ required for each step. Nevertheless, it may help orient the reader to remark that, for our intended purpose, it is sufficient to fix some sufficiently large value of k ($k \geq 4$ is sufficient) and to restrict our attention to $\delta > 0$ satisfying $\delta < 1$ and $\delta < \delta_\gamma$ for all Reeb orbits γ under consideration.

Definition 2.34 (Weighted Sobolev spaces $W^{k,2,\delta}$). For $k \geq 0$ and $\delta < 1$, we define weighted Sobolev spaces

$$W^{k,2,\delta}(C, u^*T\hat{X}), \quad (2.34)$$

$$W^{k,2,\delta}(C, u^*T\hat{X}_j \otimes_{\mathbb{C}} \Omega_C^{0,1}), \quad (2.35)$$

where $u : C \setminus \{p_e\}_e \rightarrow \hat{X}$ is a smooth map which in a neighborhood of each puncture p_e is C^∞ -convergent with weight $\delta' > \delta$ to a Reeb orbit $\gamma_e \in \mathcal{P}^\pm$ in the sense of Definition 2.6. For (2.34), we allow C to have nodes, though in this case we require $k \geq 2$ in accordance with Convention 2.32.

The $W^{k,2,\delta}$ -norm is defined as follows. Away from $\{p_e\}_e$, we use the usual $W^{k,2}$ -norm, and near a given p_e , the contribution to the norm squared is given by

$$\int_{[0,\infty) \times S^1} \sum_{j=0}^k |D^j f|^2 e^{2\delta s} ds dt \quad (2.36)$$

for any choice of local coordinates (2.33). Equivalently up to commensurability, one can set $\|f\|_{k,2,\delta} := \|\mu \cdot f\|_{k,2}$ for some smooth function μ which equals 1 away from the ends and which equals $e^{\delta|s|}$ in any end (multiplication by any such μ thus defines an isomorphism $W^{k,2,\delta} \rightarrow W^{k,2}$, allowing properties of $W^{k,2}$ to be lifted to $W^{k,2,\delta}$). Our hypotheses that $\delta < 1$ and that u is C^∞ -convergent with weight $\delta' > \delta$ imply that the $W^{k,2,\delta}$ -norm is independent up to commensurability of the choices involved in its definition.

Lemma 2.35. *Smooth functions of compact support are dense in $W^{k,2,\delta}$.*

Proof. By the Meyers–Serrin theorem [Ada75, Theorem 3.16] it is enough to show that functions of compact support are dense in $W^{k,2,\delta}$. It is enough to examine the model case of an end $[0, \infty) \times S^1$. Let $\varphi : \mathbb{R} \rightarrow [0, 1]$ be any smooth (cutoff) function satisfying $\varphi(s) = 1$ for $s \leq 1$ and $\varphi(s) = 0$ for $s \geq 2$. Now the sequence $\varphi(s/N)f(s, t)$ (or even $\varphi(s - N)f(s, t)$) converges to $f(s, t)$ in the $W^{k,2,\delta}$ topology as $N \rightarrow \infty$. \square

Next, we introduce a space $\tilde{W}^{k,2,\delta}(C, u^*T\hat{X})$ which is a modification of the space $W^{k,2,\delta}(C, u^*T\hat{X})$ taking into account the space of deformations of the almost complex structure of C equipped with the doubly marked points $\{p_e\}$. Recall from §2.6 that when C is smooth, this deformation theory is governed by $H^\bullet(C, TC(-2P))$ where $P := \sum_e p_e \subseteq C$. When C has nodes, this deformation theory (not taking into account resolutions of the nodes) is governed by $H^\bullet(\tilde{C}, T\tilde{C}(-2P - \tilde{N}))$, where $\tilde{N} \subseteq \tilde{C}$ denotes the inverse images of the nodes $N \subseteq C$ (recall the normalization \tilde{C} from Definition 2.17).

Definition 2.36 (Weighted Sobolev spaces $\tilde{W}^{k,2,\delta}$). Let k, δ , and u be as in Definition 2.34. We define a weighted Sobolev space

$$\tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) \quad (2.37)$$

as

$$W^{k,2,\delta}(C, u^*T\hat{X}) \oplus C_c^\infty(C \setminus \{p_e\} \cup N, \text{End}^{0,1}(TC)) \quad (2.38)$$

modulo the subspace

$$\left\{ Xu \oplus \mathcal{L}_{Xj} \left| \begin{array}{l} X \text{ holomorphic near } p_e \text{ and } \tilde{N} \\ X(p_e) = dX(p_e) = 0 \\ X(\tilde{n}) = 0 \text{ for } \tilde{n} \in \tilde{N} \end{array} \right. \right\}. \quad (2.39)$$

Since X is holomorphic and vanishes to second order at each p_e (so in local coordinates (2.33) it equals $f \cdot \partial_s$ for some holomorphic function f vanishing at p_e), a calculation shows that $Xu \in W^{k,2,\delta}(C, u^*T\hat{X})$ since $\delta < 1$ and u converges to a trivial cylinder with weight $\delta' > \delta$.

There is a unique commensurability class of norm on $\tilde{W}^{k,2,\delta}$ for which the natural map $W^{k,2,\delta} \rightarrow \tilde{W}^{k,2,\delta}$ is Fredholm, since this map has kernel $H^0(\tilde{C}, T\tilde{C}(-2P - \tilde{N}))$ and cokernel $H^1(\tilde{C}, T\tilde{C}(-2P - \tilde{N}))$, both of which are finite-dimensional.

To better understand the spaces $\tilde{W}^{k,2,\delta}$, we discuss a few special cases for smooth C .

If $H^1(C, TC(-2P)) = 0$, namely if all variations of C with its markings at $\{p_e\}$ are trivial up to gauge, then we have

$$\tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) = W^{k,2,\delta}(C, u^*T\hat{X}) / \ker\left(\mathbf{aut}(C, \{p_e\}_e) \rightarrow \bigoplus_e \mathfrak{gl}(T_{p_e}C)\right),$$

where $\ker(\mathbf{aut}(C, \{p_e\}_e) \rightarrow \bigoplus_e \mathfrak{gl}(T_{p_e}C)) = H^0(C, TC(-2P))$ denotes the (finite-dimensional!) Lie algebra of the group of automorphisms of C which fix each p_e and act as the identity on each $T_{p_e}C$. This Lie algebra of vector fields on C maps injectively (by stability) to $W^{k,2,\delta}(C, u^*T\hat{X})$ by differentiating u . In this case, $\tilde{W}^{k,2,\delta}$ is equipped with the natural quotient norm from the presentation above.

If $H^0(C, TC(-2P)) = 0$, namely if C equipped with its markings at $\{p_e\}$ has no nontrivial infinitesimal automorphisms, then we have

$$\tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) = W^{k,2,\delta}(C, u^*T\hat{X}) \oplus V, \quad (2.40)$$

where $V \subseteq C_c^\infty(C \setminus \{p_e\}_e, \text{End}^{0,1}(TC))$ is any subspace projecting isomorphically onto $H^1(C, TC(-2P))$ (which by §2.6 is the space of deformations of C equipped with the doubly marked points $\{p_e\}$). In this case, $\tilde{W}^{k,2,\delta}$ is equipped with the direct sum norm induced by (2.40) for any norm on V . This norm on $\tilde{W}^{k,2,\delta}$ is well-defined up to commensurability, since every other such subspace V' may be obtained as the image of $v \mapsto v + \mathcal{L}_{X(v)}j$ for some unique $X : V \rightarrow C^\infty(C, TC)$ with $X = 0$ and $dX = 0$ at p_e , and the two resulting spaces (2.40) are identified via $(\xi, v) \mapsto (\xi + X(v)u, v + \mathcal{L}_{X(v)}j)$.

To be even more concrete, suppose C is smooth of genus zero. If $\#\{p_e\} = 1$ (i.e. the case of the plane), then it falls into the first case above with a one-dimensional (over \mathbb{C}) space of infinitesimal automorphisms, corresponding to translations of \mathbb{C} (the automorphisms of \mathbb{C} acting trivially on the tangent space at infinity). If $\#\{p_e\} = 2$ (i.e. the case of the cylinder), then it falls into the second case above with a one-dimensional space of deformations of the almost complex structure. Namely, all cylinders are biholomorphic to $\mathbb{C} \setminus \{0\}$, the space of tangent space markings at 0 and ∞ is $(\mathbb{C}^\times)^2$ which is two-dimensional, and quotienting by $\text{Aut}(\mathbb{C} \setminus \{0\}) = \mathbb{C}^\times$ reduces this to the one-dimensional \mathbb{C}^\times . In fact, $\#\{p_e\} \geq 2$ all fall into the second case above.

Definition 2.37 (Linearized operators). Let k, δ , and u be as in Definition 2.34, and suppose $k \geq 1$. For any choice of \hat{J} -linear connection ∇ as in Convention 2.31, there is a linearized operator

$$D_u^\nabla : W^{k,2,\delta}(C, u^*T\hat{X}) \oplus C_c^\infty(C \setminus \{p_e\}_e, \text{End}^{0,1}(TC)) \rightarrow W^{k-1,2,\delta}(\tilde{C}, u^*T\hat{X}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}}^{0,1})$$

(recall the normalization \tilde{C} from Definition 2.17) expressing the first order change in $(du)^{0,1}$, measured with respect to the chosen connection ∇ , as u and j vary, namely

$$D_u^\nabla(\xi, j') := \frac{d}{d(\xi, j')} \Big|_{(\xi, j')=(0, j)} \text{PT}_{\exp_u \xi \rightarrow u}^\nabla \left[(d \exp_u \xi)_{j', j}^{0,1} \right] \quad (2.41)$$

where PT denotes parallel transport. When u is \hat{J} -holomorphic, the operator D_u^∇ is independent of the choice of connection and descends to a map

$$D_u : \tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) \rightarrow W^{k-1,2,\delta}(\tilde{C}, u^*T\hat{X}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}}^{0,1}). \quad (2.42)$$

For general u , we may still define such a map (2.42) by fixing a choice of connection ∇ and a choice of subspace of (2.38) which projects isomorphically onto $\tilde{W}^{k,2,\delta}$ and intersects $W^{k,2,\delta}$ in a subspace of finite codimension.

The operator

$$D_u^\nabla : W^{k,2,\delta}(C, u^*T\hat{X}) \rightarrow W^{k-1,2,\delta}(\tilde{C}, u^*T\hat{X}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}}^{0,1}) \quad (2.43)$$

is a real Cauchy–Riemann operator, in the sense that $D_u^\nabla(fs) = fD_u^\nabla(s) + (\bar{\partial}f)s$ for real-valued functions f . In particular, it is elliptic.

Proposition 2.38. *The linearized operator (2.42) is Fredholm provided $\pm\delta \notin \sigma(D_\gamma)$ for every orbit γ which u is positively/negatively asymptotic to.*

Proof. It is equivalent to show that (2.43) is Fredholm under the given hypothesis on δ . The main point of the proof of this (as compared with the standard Fredholm results for elliptic operators on compact manifolds) is to prove suitable elliptic estimates in the ends of C , which follow from the hypothesis $\delta \notin \sigma(D_\gamma)$. This argument can be viewed as a special case of a general result of Lockhart–McOwen [LM85], and proofs in the specific context of pseudo-holomorphic curves with cylindrical ends can be found in Salamon [Sal99, §2] and Wendl [Wen16, Lecture 4]. \square

Definition 2.39 (Regularity of pseudo-holomorphic buildings). Given a pseudo-holomorphic building $\{u_v : C_v \rightarrow \tilde{X}_v\}_{v \in V(T)}$, we consider the linearized operator

$$D_u : \bigoplus_{v \in V(T)} \tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) \rightarrow \bigoplus_{v \in V(T)} W^{k-1,2,\delta}(\tilde{C}, u^*T\hat{X}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}}^{0,1}) \quad (2.44)$$

with $\delta \in (0, 1)$ and $\delta < \delta_\gamma$ for every Reeb orbit $\gamma = \gamma_e$ for $e \in E(T)$. This operator is Fredholm by Proposition 2.38, and its kernel and cokernel are independent of the choice of k and δ by elliptic regularity.

A point in a moduli space $\overline{\mathcal{M}}(T)$ is called *regular* iff this linearized operator is surjective. A point is called *weakly regular* iff the linearized operator is surjective after adding variations in $t \in \mathfrak{s}(T)$ to its domain (this is only relevant in cases (III) and (IV)).

Lemma 2.40. *The trivial cylinder $\text{id} \times \tilde{\gamma} : \mathbb{R} \times S^1 \rightarrow \mathbb{R} \times Y$ over any Reeb orbit γ is regular. In fact, the associated linearized operator is an isomorphism.*

Proof. The linearized operator D may be decomposed into the tangential and normal deformation operators D^T and D^N . Precisely, there is a diagram whose rows are short exact sequences (we write $C = \mathbb{R} \times S^1$):

$$\begin{array}{ccccc} \tilde{W}^{k,2,\delta}(C, TC) & \xrightarrow{du^\circ} & \tilde{W}^{k,2,\delta}(C, \gamma^*TY \oplus \mathbb{R}\partial_s) & \longrightarrow & W^{k,2,\delta}(C, \gamma^*\xi) \\ \downarrow D^T & & \downarrow D & & \downarrow D^N \\ W^{k-1,2,\delta}(C, TC \otimes_{\mathbb{C}} \Omega_C^{0,1}) & \xrightarrow{du^\circ} & W^{k-1,2,\delta}(C, (\gamma^*TY \oplus \mathbb{R}\partial_s) \otimes_{\mathbb{C}} \Omega_C^{0,1}) & \rightarrow & W^{k-1,2,\delta}(C, \gamma^*\xi \otimes_{\mathbb{C}} \Omega_C^{0,1}). \end{array} \quad (2.45)$$

Note that the domain of D^T includes variations in complex structure on C , while the domain of D^N does not. It suffices to show that both D^T and D^N are isomorphisms.

To show that D^N is an isomorphism, write $D^N = \partial_s + D_\gamma$, where $D_\gamma = J\mathcal{L}_{R_\lambda}$ is the linearized operator for γ (Definition 2.28). It follows from the fact that δ is not an eigenvalue of D_γ that D^N is an isomorphism.

To show that D^T is an isomorphism, let us first observe that the complex

$$W^{k,2,\delta}(C, TC) \xrightarrow{D^T} W^{k-1,2,\delta}(C, TC \otimes_{\mathbb{C}} \Omega_C^{0,1}) \quad (2.46)$$

calculates $H^\bullet(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty]))$, that is we have

$$\ker D^T = H^0(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty])), \quad (2.47)$$

$$\text{coker } D^T = H^1(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty])), \quad (2.48)$$

where we identify C with $\mathbb{P}^1 \setminus \{0, \infty\}$.

To show (2.47), argue as follows. By elliptic regularity, $\ker D^T$ consists precisely of those holomorphic sections of TC over C which lie in $W^{k,2,\delta}$. Thus we must show that a holomorphic section of $T\mathbb{P}^1$ over $\mathbb{P}^1 \setminus \{0, \infty\}$ lies in $W^{k,2,\delta}$ iff it vanishes to order ≥ 2 at $0, \infty \in \mathbb{P}^1$. In fact, it suffices to show the seemingly much weaker statement that a holomorphic section of $T\mathbb{P}^1$ over $\mathbb{P}^1 \setminus \{0, \infty\}$ lying in $W^{k,2,\delta}$ must be meromorphic at $0, \infty \in \mathbb{P}^1$ (this implies the previous assertion by a straightforward calculation). So, let f be a holomorphic section lying in $W^{k,2,\delta}$, and let us show that it is meromorphic at the punctures. Since f lies in $W^{k,2,\delta}$, it can be integrated (with respect to a fixed non-vanishing area form on \mathbb{P}^1) against any smooth function on \mathbb{P}^1 (more precisely, smooth section of the real dual of $T_{\mathbb{P}^1}$) which vanishes to order N at 0 and ∞ , for some $N < \infty$ depending on k and δ . The space of such smooth functions is of finite codimension in the space of all smooth functions on \mathbb{P}^1 , so we may extend f arbitrarily to a distribution \tilde{f} on \mathbb{P}^1 . Since f is holomorphic on $\mathbb{P}^1 \setminus \{0, \infty\}$, the distribution $\bar{\partial}\tilde{f}$ is supported at $0, \infty \in \mathbb{P}^1$, so by Lemma 2.41 it is a finite sum of delta functions and their derivatives at $0, \infty \in \mathbb{P}^1$. It follows that $z^M \bar{\partial}\tilde{f} = 0$ for sufficiently large $M < \infty$, for any local coordinate z near 0 or ∞ . We have $\bar{\partial}(z^M \tilde{f}) = z^M \bar{\partial}\tilde{f} = 0$, so by elliptic regularity $z^M \tilde{f}$ is holomorphic at 0 and ∞ . It follows that f is meromorphic at $0, \infty \in \mathbb{P}^1$, which finishes the proof of (2.47). To show (2.48), apply the same reasoning to the formal adjoint to show that

$$(\text{coker } D^T)^\vee = \ker((D^T)^\vee) = H^0(\mathbb{P}^1, \Omega_{\mathbb{P}^1}^{\otimes 2}(2[0] + 2[\infty])) \quad (2.49)$$

and apply Serre duality.

It follows from (2.47)–(2.48) that D^T is an isomorphism. Indeed, $\ker D^T = H^0(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty])) = 0$, and by definition $\tilde{W}^{k,2,\delta}$ is $W^{k,2,\delta}$ direct sum a space which maps isomorphically to $H^1(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty])) = \text{coker } D^T$. \square

The following is an easy exercise (see [Hör69, Theorem 1.5.3]) and is frequently useful.

Lemma 2.41. *A distribution supported at a single point is a linear combination of the delta function at that point and its derivatives.* \square

2.12 Index of moduli spaces

We now define a notion of *index* and *virtual dimension* for objects of \mathcal{S} .

Definition 2.42 (Index $\mu(T)$). We define $\mu(T)$ to be the Fredholm index of the linearized operator (2.44) for non-nodal C_v .

Note that (2.44) makes sense for any smooth building $\{u_v\}_v$ for which u_v approach trivial cylinders sufficiently rapidly, and varies nicely in families of such $\{u_v\}_v$, hence $\mu(T)$ is well-defined.

Standard arguments allow one to express $\mu(T)$ in terms of the Conley–Zehnder indices of the Reeb orbits γ_{e^+} and $\{\gamma_{e_i^-}\}_i$ and the homology class of $\#_v \beta_v$ (see [EGH00, Proposition 1.7.1] or [BM04, Proposition 4]):

$$\mu(T) = [\text{CZ}_{\tau^+}(\gamma_{e^+}) + n - 3] - \sum_i [\text{CZ}_{\tau_i^-}(\gamma_{e_i^-}) + n - 3] + \langle 2c_1(T\hat{X}, \tau), \#_v \beta_v \rangle, \quad (2.50)$$

where $\tau = (\tau^+, \{\tau_i^-\}_i)$ denotes any collection of trivializations of ξ^\pm over the orbits γ , and $c_1(T\hat{X}, \tau)$ denotes the first Chern class relative to the boundary trivialization of $T\hat{X}$ obtained by summing τ with the tautological trivialization of $\mathbb{R}R_\lambda \oplus \mathbb{R}\partial_s$ (the expression on the right hand side is independent of the choice of τ). We may thus define the *homological grading* of a null-homologous Reeb orbit to be

$$|\gamma| := \text{CZ}_\tau(\gamma) + n - 3 + \langle 2c_1(\xi, \tau), \beta \rangle \in \mathbb{Z}/2c_1(\xi) \cdot H_2(Y) \quad (2.51)$$

for any trivialization τ and any null-homology β of γ . Analogously, there is a relative grading on the set of (monomials of) Reeb orbits in a fixed homology class.

The index satisfies $\mu(T) = \mu(T')$ for any morphism $T \rightarrow T'$ (this is evident from the formula in terms of Conley–Zehnder indices), and is additive under concatenations, that is $\mu(\#_i T_i) = \sum_i \mu(T_i)$ (trivial by definition).

Definition 2.43 (Virtual dimension $\text{vdim}(T)$). The *virtual dimension* of T is defined as

$$\text{vdim}(T) := \mu(T) - \#V_s(T) + \dim \mathfrak{s}(T), \quad (2.52)$$

recalling that $V_s(T) \subseteq V(T)$ denotes the set of symplectization vertices, i.e. those v for which $*^+(v) = *^-(v)$. This is the “expected dimension” of $\overline{\mathcal{M}}(T)$. For a morphism $T' \rightarrow T$, let $\text{codim}(T'/T) := \text{vdim } T - \text{vdim } T'$.

Lemma 2.44. *We have $\text{codim}(T'/T) \geq 0$ with equality iff $T' \rightarrow T$ is an isomorphism, and $\text{codim}(T'/T) > 1$ iff $T \rightarrow T'$ factors nontrivially.* \square

2.13 Orientations of moduli spaces

We now review the theory of orientations in contact homology. The general analytic methods used to orient moduli spaces of pseudo-holomorphic curves were introduced by Floer–Hofer [FH93] (see also Bourgeois–Mohnke [BM04]). The resulting algebraic structure relevant for contact homology was worked out by Eliashberg–Givental–Hofer [EGH00] (see also Bourgeois–Mohnke [BM04]).

We find it most convenient to encode orientations via the formalism of *orientation lines*, as introduced into the subject by Seidel. An orientation line is simply a free \mathbb{Z} -module of rank one, equipped with a $\mathbb{Z}/2$ -grading (i.e. declared to be either even or odd). For every Reeb orbit γ with basepoint b , we will define an orientation line $\mathfrak{o}_{\gamma,b}$. For every object $T \in \mathcal{S}$, we will define an orientation line \mathfrak{o}_T° , and we will show that there is a canonical isomorphism

$$\mathfrak{o}_T^\circ = \mathfrak{o}_{\gamma^+,b^+} \otimes \bigotimes_{\gamma^- \in \Gamma^-} \mathfrak{o}_{\gamma^-,b^-}. \quad (2.53)$$

The virtual orientation sheaf of $\overline{\mathcal{M}}(T)$ is canonically isomorphic to

$$\mathfrak{o}_T := \mathfrak{o}_T^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)}. \quad (2.54)$$

Recall that we are using the super tensor product \otimes (see §1.1).

Definition 2.45. Denote by $\mathfrak{o}_V := H_{\dim V}(V, V \setminus 0)$ with parity $\dim V$ the orientation line² of the vector space V . For a Fredholm map $A : E \rightarrow F$, define $\mathfrak{o}(A) = \mathfrak{o}_A := \mathfrak{o}_{\ker A} \otimes \mathfrak{o}_{\text{coker } A}^\vee$.

Definition 2.46 (Orientation lines $\mathfrak{o}_{\gamma,b}$ of Reeb orbits). Let $\gamma \in \mathcal{P} = \mathcal{P}(Y, \lambda)$, and fix a constant speed parameterization $\tilde{\gamma} : S^1 \rightarrow Y$ of γ (equivalently, fix a basepoint $b = \tilde{\gamma}(0) \in |\gamma|$). We consider the bundle $V := \tilde{\gamma}^*\xi \oplus \mathbb{C}$ over $[0, \infty) \times S^1 \subseteq \mathbb{C}$. The bundle V is equipped with a connection, namely the connection on $\tilde{\gamma}^*\xi$ induced by the Lie derivative \mathcal{L}_{R_λ} plus the trivial connection on \mathbb{C} , and this gives rise to a real Cauchy–Riemann operator $\bar{\partial}$ on V . Now extend the pair $(V, \bar{\partial})$ to all of \mathbb{C} , and define

$$\mathfrak{o}_{\gamma,b} := \mathfrak{o}(V, \bar{\partial}) := \mathfrak{o}(W^{k,2,\delta}(\mathbb{C}, V) \rightarrow W^{k-1,2,\delta}(\mathbb{C}, V \otimes_{\mathbb{C}} \Omega_{\mathbb{C}}^{0,1})) \quad (2.55)$$

for small $\delta > 0$. We show in Lemma 2.47 immediately below that $\mathfrak{o}_{\gamma,b}$ is well-defined.

Note that (because of the well-definedness of $\mathfrak{o}_{\gamma,b}$) any path between basepoints $b \rightarrow b'$ in the sense of §2.1 gives rise to an isomorphism $\mathfrak{o}_{\gamma,b} \rightarrow \mathfrak{o}_{\gamma,b'}$.

Lemma 2.47. *The orientation line $\mathfrak{o}_{\gamma,b}$ defined by (2.55) is independent of the choice of extension $(V, \bar{\partial})$ up to canonical isomorphism.*

(Refer to Solomon [Sol06, §2, Proposition 2.8] for a similar result.)

²One should distinguish the orientation line \mathfrak{o}_V from the determinant line $\det V := \wedge^{\dim V} V$, as there is no functorial isomorphism $\mathfrak{o}_V \otimes_{\mathbb{Z}} \mathbb{R} = \det V$ (rather only up to scaling by $\mathbb{R}_{>0}$).

Proof. The space of extensions of $(V, \bar{\partial})$ is homotopy equivalent to $\text{Maps}(S^2, BU(n))$ (noting that $\pi_1(BU(n)) = \pi_0(U(n)) = 0$ so V is trivial over S^1) and we have $\pi_i \text{Maps}(S^2, BU(n)) = \pi_{i+2}(BU(n)) = \pi_{i+1}(U(n))$; in particular $\pi_0 = \mathbb{Z}$ and $\pi_1 = 0$. By simple connectivity, the line $\mathfrak{o}_{\gamma,b}$ depends at most on the choice of connected component of $\text{Maps}(S^2, BU(n))$ (classified by relative Chern class).

To identify (canonically) the $\mathfrak{o}_{\gamma,b}$ from different relative Chern classes, argue as follows. Choose an extension $(V, \bar{\partial})$ for which $\bar{\partial}$ is complex linear in a neighborhood of the origin $0 \in \mathbb{C}$ (this is a contractible condition), so $(V, \bar{\partial})$ is a holomorphic vector bundle near the origin. Hence we may define $(V(-[0]), \bar{\partial})$ near the origin as the holomorphic vector bundle whose sheaf of holomorphic sections is the subsheaf of holomorphic sections of V which vanish at the origin, and we extend $(V(-[0]), \bar{\partial})$ to coincide with $(V, \bar{\partial})$ away from the origin. Thus $(V(-[0]), \bar{\partial})$ is another extension with relative Chern class one less than $(V, \bar{\partial})$, so it suffices to produce a canonical isomorphism $\mathfrak{o}(V, \bar{\partial}) = \mathfrak{o}(V(-[0]), \bar{\partial})$.

To produce this isomorphism, it suffices to construct the expected exact sequence

$$0 \rightarrow H_\delta^0(V([-0]), \bar{\partial}) \rightarrow H_\delta^0(V, \bar{\partial}) \rightarrow V_0 \rightarrow H_\delta^1(V([-0]), \bar{\partial}) \rightarrow H_\delta^1(V, \bar{\partial}) \rightarrow 0 \quad (2.56)$$

where $H_\delta^i(V, \bar{\partial})$ denotes the cohomology of the two term complex appearing in (2.55) and V_0 denotes the fiber of V over $0 \in \mathbb{C}$ (note that V_0 is a complex vector space and hence is canonically oriented). To produce this exact sequence, it suffices to apply the snake lemma to the diagram obtained from the inclusion of the complex calculating $H_\delta^\bullet(V(-[0]), \bar{\partial})$ into that calculating $H_\delta^\bullet(V, \bar{\partial})$. At the level of smooth sections, this is the diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & C^\infty(\mathbb{C}, V(-[0])) & \longrightarrow & C^\infty(\mathbb{C}, V) & \longrightarrow & J^{0,\infty}(V)_0 \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & C^\infty(\mathbb{C}, V(-[0]) \otimes \Omega_{\mathbb{C}}^{0,1}) & \longrightarrow & C^\infty(\mathbb{C}, V \otimes \Omega_{\mathbb{C}}^{0,1}) & \longrightarrow & J^{0,\infty}(V \otimes \Omega_{\mathbb{C}}^{0,1})_0 \longrightarrow 0, \end{array} \quad (2.57)$$

where $J^{0,\infty}(-)_0$ denotes the infinite anti-holomorphic jet space at $0 \in \mathbb{C}$, i.e. keeping track of all anti-holomorphic derivatives $(\frac{\partial}{\partial \bar{z}})^k$ at zero; note that the right vertical map is surjective with kernel V_0 . Now the usual diagram chase (in combination with some appeals to elliptic regularity) produces the desired exact sequence (2.56). In doing this diagram chase, it may be helpful to remark that, to define the connecting homomorphism, one may, for convenience, restrict consideration to lifts of elements of V_0 to $W^{k,2,\delta}(\mathbb{C}, V)$ which are holomorphic near zero; also note that by elliptic regularity, all elements of $H_\delta^1(\mathbb{C}, V)$ can be represented by smooth sections of $V \otimes \Omega_{\mathbb{C}}^{0,1}$ supported inside any given non-empty open subset of \mathbb{C} . \square

Definition 2.48 (Parity of Reeb orbits). The *parity* $|\gamma| \in \mathbb{Z}/2$ of $\gamma \in \mathcal{P}$ is the parity of \mathfrak{o}_γ .

By (2.50) we have $|\gamma| = \text{CZ}(\gamma) + n - 3$, i.e. the parity is simply the reduction to $\mathbb{Z}/2$ of the grading (2.51) (whenever the latter is defined). This formula for $|\gamma|$ in terms of the Conley–Zehnder index can also be expressed as $|\gamma| = \text{sign}(\det(I - A_\gamma)) \in \{\pm 1\} = \mathbb{Z}/2$ where A_γ denotes the linearized return map of γ acting on ξ_p (using the general property of the Conley–Zehnder index $(-1)^{\text{CZ}(\Psi)} = (-1)^{n-1} \text{sign}(\det(I - \Psi(1)))$ for $\Psi : [0, 1] \rightarrow \text{Sp}_{2n-2}(\mathbb{R})$). It thus follows that

$$|\gamma| = \#(\sigma(A_\gamma) \cap (0, 1)) \in \mathbb{Z}/2, \quad (2.58)$$

where $\sigma(\cdot)$ denotes the spectrum (recall that the spectrum of any matrix $A \in \mathrm{Sp}_{2n}(\mathbb{R})$ lies in $\mathbb{R}^\times \cup \{z \in \mathbb{C} : |z| = 1\}$ and that $1 \notin \sigma(A_\gamma)$ is equivalent to the non-degeneracy of γ). Note that by definition γ is non-degenerate iff $1 \notin \sigma(A_\gamma)$. It follows that the index of the k -fold multiple cover γ_k of γ is given by

$$|\gamma_k| = |\gamma| + (k+1)\#(\sigma(A_\gamma) \cap (-1, 0)) \in \mathbb{Z}/2. \quad (2.59)$$

Definition 2.49 (Good and bad Reeb orbits). There is an action of \mathbb{Z}/d_γ on $\mathfrak{o}_{\gamma,b}$ by functoriality, which just amounts to a homomorphism $\mathbb{Z}/d_\gamma \rightarrow \{\pm 1\}$ (independent of b since \mathbb{Z}/d_γ is abelian). The orbit γ is called *good* iff this homomorphism is trivial (and *bad* otherwise). For good γ , we thus have an orientation line \mathfrak{o}_γ independent of b up to canonical isomorphism.

The bad Reeb orbits are precisely the even multiple covers γ_{2k} of *simple* orbits γ with $\#(\sigma(A_\gamma) \cap (-1, 0))$ odd by [EGH00, Lemma 1.8.8, Remark 1.9.2]. To see this, it suffices to show that a generator of the \mathbb{Z}/k action on \mathfrak{o}_{γ_k} acts by $(-1)^{|\gamma_k| - |\gamma|}$ and use (2.59). This can be proven by pulling back the operator from (2.55) under $z \mapsto z^k$ and analyzing the representations of \mathbb{Z}/k occurring in the kernel and cokernel (see [BM04, Proof of Theorem 3]).

Definition 2.50 (Orientation lines \mathfrak{o}_T°). For any $T \in \mathcal{S}$, we define the orientation line \mathfrak{o}_T° to be the orientation line of the linearized operator (2.44) at any smooth building of type T .

We show in Lemma 2.51 immediately below that \mathfrak{o}_T° is well-defined (i.e. is independent of the choice of smooth building, up to canonical isomorphism). Note that it does not matter whether we define \mathfrak{o}_T° in terms of the linearized operator with domain $W^{k,2,\delta}$ or $\tilde{W}^{k,2,\delta}$, since their “difference” is a complex vector space and thus is canonically oriented. Note also that the parity of \mathfrak{o}_T° is just the index $\mu(T) \bmod 2$.

Lemma 2.51. *The orientation line of the linearized operator (2.44) at a smooth building of type T is canonically isomorphic to $\mathfrak{o}_{\gamma_{e^+}, b_{e^+}} \otimes \bigotimes_{e^-} \mathfrak{o}_{\gamma_{e^-}, b_{e^-}}^\vee$.*

Proof. We consider the linearized operator (2.44) direct sum the operators from (2.55) for each output Reeb orbit γ_{e^-} of T . The orientation line of this direct sum is thus $\mathfrak{o}_T^\circ \otimes \bigotimes_{e^-} \mathfrak{o}_{\gamma_{e^-}, b_{e^-}}$.

Now the Floer–Hofer kernel gluing operation [FH93] (see also Bourgeois–Mohnke [BM04]) provides a canonical isomorphism between the orientation line of this operator and the orientation line of the operator obtained by gluing the domains end to end in the specified way (note the use of the basepoints b_{e^-} for γ_{e^-} specified by T). Note that, for technical convenience, it is enough to do this gluing just for maps u_v which coincide with trivial cylinders near infinity.

This glued up operator is now exactly of the form from (2.55) for γ_{e^+} , so we deduce an isomorphism $\mathfrak{o}_T^\circ \otimes \bigotimes_{e^-} \mathfrak{o}_{\gamma_{e^-}, b_{e^-}} = \mathfrak{o}_{\gamma_{e^+}, b_{e^+}}$. \square

Remark 2.52. Gluing the linearized operators with domain $W^{k,2,\delta}$ results in an index increase of two, gluing the linearized operators with domain $\tilde{W}^{k,2,\delta}$ preserves the index, and gluing the linearized operators with domain $W^{k,2,\delta}$ plus decay to constants in the ends results in an index decrease of two. The Floer–Hofer kernel gluing operation can be done in any of these contexts; Bourgeois–Mohnke choose the last.

For any morphism $T \rightarrow T'$, there is a canonical isomorphism $\mathfrak{o}_T^\circ \rightarrow \mathfrak{o}_{T'}^\circ$, again by the Floer–Hofer kernel gluing operation. Associativity of this operation implies that this isomorphism respects the identifications from Lemma 2.51 and that it makes \mathfrak{o}° into a functor from \mathcal{S}_1 to the category of orientation lines and isomorphisms. For any concatenation $\{T_i\}$, there is a tautological identification $\mathfrak{o}_{\#_i T_i}^\circ = \bigotimes_i \mathfrak{o}_{T_i}^\circ$; associativity of the Floer–Hofer operation implies that this identification is compatible with Lemma 2.51.

Remark 2.53. The above reasoning, in particular the proof of Lemma 2.51, relies on the fact that the topology of the curves in question is particularly simple. To prove the analogous result in the SFT setting requires a more complicated argument; see Bourgeois–Mohnke [BM04, Proposition 8].

For any $T \in \mathcal{S}$, define

$$\mathfrak{o}_T := \mathfrak{o}_T^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)}, \quad (2.60)$$

where $\mathfrak{o}_{\mathfrak{s}(T)}$ denotes the global sections of the orientation sheaf of \mathfrak{s} considered as a real manifold. We shall see that the virtual orientation sheaf of $\overline{\mathcal{M}}(T)$ is canonically isomorphic to \mathfrak{o}_T . Note that the parity of \mathfrak{o}_T equals $\text{vdim}(T)$.

For any morphism $T' \rightarrow T$, there is a canonical identification

$$\mathfrak{o}_{T'} \otimes \mathfrak{o}_{G_{T'/T}} \otimes \mathfrak{o}_{\mathfrak{s}(T')}^\vee = \mathfrak{o}_T, \quad (2.61)$$

where $G_{T'/T}$ denotes the space of gluing parameters defined in §2.5. The appearance of $\mathfrak{o}_{\mathfrak{s}(T')}^\vee$ is explained by the fact that $G_{T'/T}$ includes all variations of t , rather than just those normal to $\mathfrak{s}(T')$.

For any concatenation $\{T_i\}$, there is a tautological identification

$$\mathfrak{o}_{\#_i T_i} = \bigotimes_i \mathfrak{o}_{T_i}. \quad (2.62)$$

2.14 Local structure of moduli spaces via local models G

We now state the precise sense in which the spaces $G_{T'/T}$ from §2.5 are local topological models for the regular loci in the moduli spaces $\overline{\mathcal{M}}(T)$. We also state the compatibility of this local topological structure with the natural maps on orientation lines discussed earlier. These statements are in essence a gluing theorem, and their proofs are given in §5.

We denote by $\overline{\mathcal{M}}(T)^{\text{reg}} \subseteq \overline{\mathcal{M}}(T)$ the locus of points which are regular and have trivial isotropy (we do not discuss the entire locus of regular points, possibly with isotropy, just for sake of simplicity, although it would not really present any additional difficulty to do so).

The (Banach space) implicit function theorem implies that $\mathcal{M}(T)^{\text{reg}}$ is a (smooth) manifold of dimension $\text{vdim } T$ over the locus without nodes. Recall that for any morphism $T \rightarrow T'$, we have $\mu(T) = \mu(T')$. Denote by $N(x)$ the set of nodes of the domain curve of a given point $x \in \overline{\mathcal{M}}(T)$.

Theorem 2.54 (Local structure of $\overline{\mathcal{M}}(T)^{\text{reg}}$). *Let $x_0 \in \overline{\mathcal{M}}(T)^{\text{reg}}$ be of type $T' \rightarrow T$. Then $\mu(T) - \#V_s(T') - 2\#N(x_0) \geq 0$ and there is a local homeomorphism*

$$(G_{T'/T} \times \mathbb{C}^{N(x_0)} \times \mathbb{R}^{\mu(T) - \#V_s(T') - 2\#N(x_0)}, (0, 0, 0)) \rightarrow (\overline{\mathcal{M}}(T), x_0) \quad (2.63)$$

whose image lands in $\overline{\mathcal{M}}(T)^{\text{reg}}$ and which commutes with the maps from both sides to $\mathcal{S}_{T'/T} \times \overline{\mathfrak{s}}(T)$ (as well as the stratifications by number of nodes).

Proof. See §§5.1–5.3. □

The (Banach space) implicit function theorem moreover gives a canonical identification

$$\mathfrak{o}_{\overline{\mathcal{M}}(T)^{\text{reg}}} = \mathfrak{o}_T. \quad (2.64)$$

More precisely, this identification is made over the locus in $\mathcal{M}(T)^{\text{reg}}$ without nodes, and has a unique continuous extension to all of $\overline{\mathcal{M}}(T)^{\text{reg}}$ by virtue of the local topological description in Theorem 2.54. It is also compatible with morphisms $T' \rightarrow T$ in the following precise sense.

Theorem 2.55 (Compatibility of the “analytic” and “geometric” maps on orientations). *The following diagram commutes:*

$$\begin{array}{ccc} \mathfrak{o}_{\overline{\mathcal{M}}(T')^{\text{reg}}} \otimes \mathfrak{o}_{G_{T'/T}} \otimes \mathfrak{o}_{\mathfrak{s}(T')^{\vee}} & \xrightarrow{(2.64)} & \mathfrak{o}_{T'} \otimes \mathfrak{o}_{G_{T'/T}} \otimes \mathfrak{o}_{\mathfrak{s}(T')^{\vee}} \\ \parallel (2.63) & & \parallel (2.61) \\ \mathfrak{o}_{\overline{\mathcal{M}}(T)^{\text{reg}}} & \xrightarrow{(2.64)} & \mathfrak{o}_T, \end{array} \quad (2.65)$$

where the left vertical map is the “geometric” map induced by the local topological structure of $\overline{\mathcal{M}}(T)^{\text{reg}}$ given in (2.63), and the right vertical map is the “analytic” map (2.61) defined earlier via the “kernel gluing” operation.

Proof. See §5.4. □

3 Implicit atlases

In this section, we define (topological) implicit atlases with cell-like stratification on the moduli spaces $\overline{\mathcal{M}}(T)$ stratified by \mathcal{S}/T . The notion of an implicit atlas with cell-like stratification was introduced in [Par16, §§3,6], and our constructions of implicit atlases follow the general procedure introduced in [Par16, §§1–2,9–10]. We will, however, give a self-contained treatment.

3.1 Implicit atlases with cell-like stratification

We review the notion of an implicit atlas with cell-like stratification from [Par16, Definitions 3.1.1, 3.2.1, 6.1.2, 6.1.6]. The notion we present here is, in fact, more general, as we must allow stratifications by categories rather than just posets (recall Definitions 2.14 and 2.15). Specifically, we consider here stratifications by categories \mathcal{T} satisfying the following properties:

- (i) \mathcal{T} has a final object.
- (ii) \mathcal{T} is finite, meaning $\#\mathcal{T} < \infty$ and $\#\text{Hom}(\mathfrak{t}_1, \mathfrak{t}_2) < \infty$ for $\mathfrak{t}_1, \mathfrak{t}_2 \in \mathcal{T}$.
- (iii) $\mathcal{T}_{\mathfrak{t}_/}$ is a poset for every $\mathfrak{t}_/ \in \mathcal{T}$.

It will be useful to consider such categories \mathcal{T} equipped with *dimension and orientation data*, which consists of:

- (i) A function $\dim : \mathcal{T} \rightarrow \mathbb{Z}$ such that for every morphism $f : \mathfrak{t} \rightarrow \mathfrak{t}'$, we have $\dim \mathfrak{t} \leq \dim \mathfrak{t}'$ with equality iff f is an isomorphism.
- (ii) For every $\mathfrak{t} \in \mathcal{T}$, an orientation line $\mathfrak{o}_{\mathfrak{t}}$ of parity $\dim \mathfrak{t}$ (more formally, this means a functor from the subcategory of isomorphisms in \mathcal{T} to the category of orientation lines).
- (iii) For every $\mathfrak{t} \rightarrow \mathfrak{t}'$ of codimension one (meaning $\dim \mathfrak{t}' = \dim \mathfrak{t} + 1$), an isomorphism $\mathfrak{o}_{\mathfrak{t}} \otimes \mathfrak{o}_{\mathbb{R}_{\geq 0}} \xrightarrow{\sim} \mathfrak{o}_{\mathfrak{t}'}$ (where $\mathfrak{o}_{\mathbb{R}_{\geq 0}} = \mathfrak{o}_{\mathbb{R}}$ is denoted as $\mathfrak{o}_{\mathbb{R}_{\geq 0}}$ to indicate that it represents the orientation line of directions transverse to the boundary of a manifold with boundary).

Definition 3.1. Let \mathcal{T} be a category as above equipped with dimension and orientation data. An *oriented cell-like stratification* is a stratification $X \rightarrow \mathcal{T}$ such that:

- (i) For every $\mathfrak{t} \in \mathcal{T}$, the closed stratum $X_{/\mathfrak{t}}$ is a topological manifold with boundary of dimension $\dim \mathfrak{t}$, whose interior is precisely the locally closed stratum $X_{\mathfrak{t}} \subseteq X_{/\mathfrak{t}}$.
- (ii) There are specified isomorphisms $\mathfrak{o}_{X_{/\mathfrak{t}}} = \mathfrak{o}_{\mathfrak{t}}$ such that for every $\mathfrak{t} \rightarrow \mathfrak{t}'$ of codimension one, the isomorphism $\mathfrak{o}_{X_{/\mathfrak{t}}} \otimes \mathfrak{o}_{\mathbb{R}_{\geq 0}} \xrightarrow{\sim} \mathfrak{o}_{X_{/\mathfrak{t}'}}$ induced by the local topology agrees with the specified isomorphism $\mathfrak{o}_{\mathfrak{t}} \otimes \mathfrak{o}_{\mathbb{R}_{\geq 0}} \xrightarrow{\sim} \mathfrak{o}_{\mathfrak{t}'}$.

Definition 3.2. Let X be a compact Hausdorff space with stratification by a category \mathcal{T} as above equipped with dimension and orientation data. An *oriented implicit atlas with cell-like stratification* on X is an index set A along with the following data:

- (i) (Covering groups) A finite group Γ_{α} for all $\alpha \in A$ (let $\Gamma_I := \prod_{\alpha \in I} \Gamma_{\alpha}$).
- (ii) (Obstruction spaces) A finitely generated $\mathbb{R}[\Gamma_{\alpha}]$ -module E_{α} for all $\alpha \in A$ (let $E_I := \bigoplus_{\alpha \in I} E_{\alpha}$).
- (iii) (Thickenings) A Hausdorff Γ_I -space X_I for all finite $I \subseteq A$, equipped with a Γ_I -invariant stratification $X_I \rightarrow \mathcal{T}$, and a homeomorphism $X = X_{\emptyset}$ respecting the stratification.
- (iv) (Kuranishi maps) A Γ_{α} -equivariant function $s_{\alpha} : X_I \rightarrow E_{\alpha}$ for all $\alpha \in I \subseteq A$ (for $I \subseteq J$, let $s_I : X_J \rightarrow E_I$ denote $\bigoplus_{\alpha \in I} s_{\alpha}$).
- (v) (Footprints) A Γ_I -invariant open set $U_{IJ} \subseteq X_I$ for all $I \subseteq J \subseteq A$.
- (vi) (Footprint maps) A Γ_J -equivariant function $\psi_{IJ} : (s_{J \setminus I}|X_J)^{-1}(0) \rightarrow U_{IJ}$ for all $I \subseteq J \subseteq A$.
- (vii) (Regular and trivial isotropy locus) A Γ_I -invariant subset $X_I^{\text{reg}} \subseteq X_I$ for all $I \subseteq A$.

which must satisfy the following ‘‘compatibility axioms’’:

- (i) $\psi_{IJ}\psi_{JK} = \psi_{IK}$ and $\psi_{II} = \text{id}$.
- (ii) $s_I\psi_{IJ} = s_I$.
- (iii) $U_{IJ_1} \cap U_{IJ_2} = U_{I, J_1 \cup J_2}$ and $U_{II} = X_I$.
- (iv) The restriction of the stratification on X_J to $(s_{J \setminus I}|X_J)^{-1}(0)$ is identified with the pullback of the stratification on X_I via ψ_{IJ} . These identifications must be compatible in triples $I \subseteq J \subseteq K$ (this condition can be nontrivial when \mathcal{T} is a category rather than a poset).
- (v) (Homeomorphism axiom) The footprint map ψ_{IJ} induces a homeomorphism $(s_{J \setminus I}|X_J)^{-1}(0)/\Gamma_{J \setminus I} \xrightarrow{\sim} U_{IJ}$.

and the following “transversality axioms”:

- (vi) $\psi_{IJ}^{-1}(X_I^{\text{reg}}) \subseteq X_J^{\text{reg}}$.
- (vii) $\Gamma_{J \setminus I}$ acts freely on $\psi_{IJ}^{-1}(X_I^{\text{reg}})$.
- (viii) (Openness axiom) $X_I^{\text{reg}} \subseteq X_I$ is open.
- (ix) (Submersion axiom) The stratification of X_I^{reg} is cell-like of dimension $\dim \mathfrak{t} + \dim E_I$ and oriented with respect to $\mathfrak{o}_{\mathfrak{t}} \otimes \mathfrak{o}_{E_I}$. Furthermore, near every point of $\psi_{IJ}^{-1}(X_I^{\text{reg}}) \subseteq X_J^{\text{reg}}$, the map $s_{J \setminus I} : X_J \rightarrow E_{J \setminus I}$ is locally modelled on the projection $X_I^{\text{reg}} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$ compatibly with oriented stratifications.
- (x) (Covering axiom) $X_{\emptyset} = \bigcup_{I \subseteq A} \psi_{\emptyset I}((s_I|X_I^{\text{reg}})^{-1}(0))$.

Definition 3.3. An *oriented implicit atlas with boundary* is an oriented implicit atlas with cell-like stratification by the two-element poset $\{\partial, \circ\}$ with $\partial < \circ$.

Remark 3.4. Given an oriented implicit atlas with cell-like stratification, one may obtain an oriented implicit atlas with boundary simply by collapsing all strata other than the maximal stratum \circ into a single stratum ∂ .

The *virtual dimension* d of a space equipped with an implicit atlas shall mean $\dim \mathfrak{t}^{\max}$, where $\mathfrak{t}^{\max} \in \mathcal{T}$ is the final object, and the *virtual orientation line* \mathfrak{o} of such a space shall mean $\mathfrak{o}_{\mathfrak{t}^{\max}}$.

3.2 Stratifications of G are cell-like

We now show that the natural stratifications $G_{T'} \rightarrow \mathcal{S}_{T'}$ are cell-like in the sense of Definition 3.1, where we equip $\mathcal{S}_{T'}$ with the dimension function

$$\mathcal{S}_{T'} \rightarrow \mathbb{Z} \tag{3.1}$$

$$(T \rightarrow T') \mapsto \text{codim}(T/T') + \dim \mathfrak{s}(T) \tag{3.2}$$

and orientation data

$$(T \rightarrow T') \mapsto (\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(T')} \otimes (\mathfrak{o}_{\mathbb{R}})^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T')}. \tag{3.3}$$

Lemma 3.5. *The stratification $(G_I)_{T'} \rightarrow (\mathcal{S}_I)_{T'}$ is cell-like.*

Proof. By inspection. □

Lemma 3.6. *The stratification $(G_{II})_{T'} \rightarrow (\mathcal{S}_{II})_{T'}$ is cell-like.*

Proof. Observe that for any $f : T \rightarrow T'$, we have

$$\begin{aligned} (G_{II})_{T//T'} &= \prod_{*(v')=00} (G_I^+)_{f^{-1}(v')/} \times \prod_{*(v')=01} (G_{II})_{f^{-1}(v')/} \times \prod_{*(v')=11} (G_I^-)_{f^{-1}(v')/}, \\ (\mathcal{S}_{II})_{T//T'} &= \prod_{*(v')=00} (\mathcal{S}_I^+)_{f^{-1}(v')/} \times \prod_{*(v')=01} (\mathcal{S}_{II})_{f^{-1}(v')/} \times \prod_{*(v')=11} (\mathcal{S}_I^-)_{f^{-1}(v')/}, \end{aligned}$$

compatibly with stratifications. Note that a product of cell-like stratifications is cell-like. Thus by induction (say, on the number of vertices of T), it suffices to show that $(G_{II})_{T'}$ is a topological manifold with boundary, whose interior coincides with the top stratum.

To show that $(G_{\text{II}})_{T/}$ is a topological manifold with boundary, we perform a change of variables $h = e^{-\mathfrak{s}} \in [0, 1)$. For convenience, we will allow $h \in [0, \infty)$ (this relaxation is certainly permitted for the present purpose). The relation $\mathfrak{g}_v = \mathfrak{g}_e + \mathfrak{g}_{v'}$ now becomes $h_v = h_e h_{v'}$. Under this relation $h_v = h_e h_{v'}$, observe that $h_v \in [0, \infty)$ and $h_e^2 - h_{v'}^2 \in (-\infty, \infty)$ determine $h_e \in [0, \infty)$ and $h_{v'} \in [0, \infty)$ uniquely, since

$$(h_e + i h_{v'})^2 = (h_e^2 - h_{v'}^2) + 2i h_v. \quad (3.4)$$

Thus if we perform another change of variables $q_e = h_e^2 - h_{v'}^2$ for $v \xrightarrow{e} v'$, then we have

$$(G_{\text{II}})_{T/} = \left\{ \begin{array}{ll} h_{v^{\max}} \in [0, \infty) & \text{if } *(v^{\max}) = 00 \\ q_e \in (-\infty, \infty) & \text{for } v \xrightarrow{e} v' \text{ with } *(v') = 00 \\ h_e \in [0, \infty) & \text{for } *(e) = 1 \end{array} \right\}. \quad (3.5)$$

This is clearly a topological manifold with boundary of dimension $\#V_s(T)$. Furthermore, the top stratum is the locus where $h_{v^{\max}} > 0$ and $h_e > 0$, which is clearly its interior.

Though not logically necessary due to the inductive reasoning above, let us remark that one can also express the stratification of $(G_{\text{II}})_{T/}$ by $(\mathcal{S}_{\text{II}})_{T/}$ concretely in terms of the coordinates (3.5) and thereby verify explicitly that it is cell-like. \square

Lemma 3.7. *The stratification $(G_{\text{III}})_{T/} \rightarrow (\mathcal{S}_{\text{III}})_{T/}$ is cell-like.*

Proof. Express the underlying forest of $T \in \mathcal{S}_{\text{III}}$ as the disjoint union of $T_i \in \mathcal{S}_{\text{II}}$, so we have

$$(G_{\text{III}})_{T/} = \prod_{i \in I} (G_{\text{II}})_{T_i/} \times \left\{ \begin{array}{ll} [0, 1) & \mathfrak{s}(T) = \{0\} \\ (0, 1) & \mathfrak{s}(T) = (0, 1) \\ (0, 1] & \mathfrak{s}(T) = \{1\}, \end{array} \right. \quad (3.6)$$

$$(\mathcal{S}_{\text{III}})_{T/} = \prod_{i \in I} (\mathcal{S}_{\text{II}})_{T_i/} \times \left\{ \begin{array}{ll} \{\{0\} < (0, 1)\} & \mathfrak{s}(T) = \{0\} \\ \{(0, 1)\} & \mathfrak{s}(T) = (0, 1) \\ \{(0, 1) > \{1\}\} & \mathfrak{s}(T) = \{1\}. \end{array} \right. \quad (3.7)$$

Now apply Lemma 3.6 and note that the product of two cell-like stratifications is again cell-like. \square

Lemma 3.8. *The stratification $(G_{\text{IV}})_{T/} \rightarrow (\mathcal{S}_{\text{IV}})_{T/}$ is cell-like.*

Proof. For $\mathfrak{s} \neq \{\infty\}$, this is just Lemma 3.7.

For $\mathfrak{s} = \{\infty\}$, argue as follows. For any given $T \rightarrow T'$, we may express $(G_{\text{IV}})_{T//T'} \rightarrow (\mathcal{S}_{\text{IV}})_{T//T'}$ as a product as in the first step of the proof of Lemma 3.6, by expressing T' as a concatenation of maximal T_i . Thus by induction, it suffices to show that $(G_{\text{IV}})_{T/}$ is a topological manifold with boundary whose interior is the stratum corresponding to the maximal contraction of T . To see this, we perform the same change of variables as in Lemma 3.6 to write

$$(G_{\text{IV}})_{T/} = \left\{ \begin{array}{ll} h_{v^{\max}} \in [0, \infty) & \text{if } *(v^{\max}) = 00 \\ q_e \in (-\infty, \infty) & \text{for } v \xrightarrow{e} v' \text{ with } *(v') = 00 \\ \tau \in [0, \infty) & \\ q_e \in (-\infty, \infty) & \text{for } v \xrightarrow{e} v' \text{ with } *(v') = 11 \\ h_e \in [0, \infty) & \text{for } *(e) = 2 \end{array} \right\}. \quad (3.8)$$

This is clearly a topological manifold of dimension $\#V_s(T) + 1$, and its maximal stratum is the locus where $\tau > 0$, $h_{v^{\max}} > 0$, and $h_e > 0$, which is clearly its interior. \square

3.3 Sets of thickening datums $A(T)$

We define sets of “thickening datums” $A(T)$ for the implicit atlases on $\overline{\mathcal{M}}(T)$. Roughly speaking, a thickening datum $\alpha \in A(T)$ is a collection of data $(r_\alpha, D_\alpha, E_\alpha, \nu_\alpha)$ from which we will construct a “thickened” version of a given moduli space by: (1) adding r_α marked points to the domain (constrained to lie on the codimension two submanifold \hat{D}_α), (2) adding an extra parameter e_α lying in the vector space E_α , and (3) adding an extra term $\nu_\alpha(e_\alpha)$ (which depends on the location of the r_α added points and the positive/negative ends $E^{\text{ext}}(T)$ in the domain) to the pseudo-holomorphic curve equation.

Recall from §2.7 that $\overline{\mathcal{M}}_{0,n}$ ($n \geq 3$) denotes the Deligne–Mumford moduli space of stable nodal Riemann surfaces of genus zero with n marked points labeled with $\{1, \dots, n\}$. We denote by $\overline{\mathcal{C}}_{0,n} \rightarrow \overline{\mathcal{M}}_{0,n}$ the universal family. Recall that $\overline{\mathcal{M}}_{0,n}$ is a compact smooth manifold. We usually prefer to label the marked points using a set other than $\{1, \dots, n\}$, so we will also use the notation $\overline{\mathcal{M}}_{0,n}$ and $\overline{\mathcal{C}}_{0,n}$ when n is a *finite set* ($\#n \geq 3$) used to label the marked points.

Definition 3.9 (Set of thickening datums A_I). A *thickening datum* α for $T \in \mathcal{S}_I$ consists of the following data:

- (i) $r_\alpha \geq 0$ an integer such that $r_\alpha + \#E^{\text{ext}}(T) \geq 3$.
- (ii) $n_\alpha \geq 0$ an integer (let $E_\alpha := \mathbb{R}^{n_\alpha}$) and an action on E_α of S_{r_α} (the group of permutations of $\{1, \dots, r_\alpha\}$).
- (iii) $D_\alpha \subseteq Y$ a compact codimension two submanifold with boundary. We let $\hat{D}_\alpha := \mathbb{R} \times D_\alpha \subseteq \hat{Y}$.
- (iv) $\nu_\alpha : E_\alpha \rightarrow C^\infty(\hat{Y} \times \overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}, T\hat{Y} \otimes_{\mathbb{R}} \Omega_{\overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}} / \overline{\mathcal{M}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}}^{0,1})^{\mathbb{R}}$ an S_{r_α} -equivariant linear map which vanishes in a neighborhood of the nodes and the $E^{\text{ext}}(T)$ -marked points of the fibers of $\overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}} \rightarrow \overline{\mathcal{M}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}$. The superscript \mathbb{R} indicates taking the subspace of \mathbb{R} -invariant sections (where \mathbb{R} acts on \hat{Y} by translation).

We denote by $A_I(T)$ the set of such thickening datums. Note that for any morphism $T \rightarrow T'$, there is a tautological identification $A_I(T) = A_I(T')$.

Definition 3.10 (Set of thickening datums A_{II}). A *thickening datum* α for $T \in \mathcal{S}_{II}$ consists of the following data:

- (i) r_α, E_α as in Definition 3.9(i), (ii).
- (ii) $D_\alpha^\pm \subseteq Y^\pm, \nu_\alpha^\pm : E_\alpha \rightarrow C^\infty(\hat{Y}^\pm \times \overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}, T\hat{Y}^\pm \otimes_{\mathbb{R}} \Omega_{\overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}} / \overline{\mathcal{M}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}}^{0,1})^{\mathbb{R}}$ as in Definition 3.9(iii), (iv).
- (iii) $\hat{D}_\alpha \subseteq \hat{X}$ a closed codimension two submanifold with boundary. We require that \hat{D}_α coincide (via (1.8)–(1.9)) with \hat{D}_α^\pm outside a compact subset of \hat{X} .
- (iv) $\nu_\alpha : E_\alpha \rightarrow C^\infty(\hat{X} \times \overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}, T\hat{X} \otimes_{\mathbb{R}} \Omega_{\overline{\mathcal{C}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}} / \overline{\mathcal{M}}_{0, E^{\text{ext}}(T) \cup \{1, \dots, r_\alpha\}}}^{0,1})^{\mathbb{R}}$ an S_{r_α} -equivariant linear map vanishing in a neighborhood of the nodes and $E^{\text{ext}}(T)$ -marked points. We require that ν_α coincide (via (1.8)–(1.9)) with ν_α^\pm outside a compact subset of \hat{X} .

We denote by $A_{II}(T)$ the set of such thickening datums. Note that for any morphism $T \rightarrow T'$, there is a tautological identification $A_{II}(T) = A_{II}(T')$.

Definition 3.11 (Set of thickening datums A_{III}). A *thickening datum* α for a connected non-empty $T \in \mathcal{S}_{\text{III}}$ is defined identically as a thickening datum for $T \in \mathcal{S}_{\text{II}}$. This makes sense since in Setup III, the identifications (1.8)–(1.9) are independent of t and the definition of a thickening datum does not make reference to λ^t or J^t . We denote by $A_{\text{III}}(T)$ the set of such thickening datums. Note that for any morphism $T \rightarrow T'$, there is a tautological identification $A_{\text{III}}(T) = A_{\text{III}}(T')$.

Definition 3.12 (Set of thickening datums A_{IV}). A *thickening datum* α for a connected non-empty $T \in \mathcal{S}_{\text{IV}}$ consists of the following data:

- (i) r_α, E_α as in Definition 3.9(i),(ii).
- (ii) $\{\hat{D}_\alpha^i \subseteq Y^i\}_{i=0,1,2}, \{\nu_\alpha^i : E_\alpha \rightarrow C^\infty(\hat{Y}^i \times \bar{\mathcal{C}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}}, T\hat{Y}^i \otimes_{\mathbb{R}} \Omega_{\bar{\mathcal{C}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}} / \bar{\mathcal{M}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}}}^{0,1})^{\mathbb{R}}\}_{i=0,1,2}$ as in Definition 3.9(iii),(iv).
- (iii) $\{\hat{D}_\alpha^{i,i+1} \subseteq \hat{X}^{i,i+1}\}_{i=0,1}$ as in Definition 3.10(iii).
- (iv) $\hat{D}_\alpha^{02,t} \subseteq \hat{X}^{02,t}$ a smoothly varying family of compact codimension two submanifolds with boundary for $t \in [0, \infty)$ as in Definition 3.10(iii) which coincides with the descent of $\hat{D}_\alpha^{01} \sqcup \hat{D}_\alpha^{12}$ for sufficiently large t . “Smoothly varying” means the projection to $[0, \infty)$ is a submersion (and remains a submersion when restricted to the boundary).
- (v) $\{\nu_\alpha^{i,i+1} : E_\alpha \rightarrow C^\infty(\hat{X}^{i,i+1} \times \bar{\mathcal{C}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}}, T\hat{X}^{i,i+1} \otimes_{\mathbb{R}} \Omega_{\bar{\mathcal{C}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}} / \bar{\mathcal{M}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}}}^{0,1})\}_{i=0,1}$ as in Definition 3.10(iv).
- (vi) $\nu_\alpha^{02,t} : E_\alpha \rightarrow C^\infty(\hat{X}^{02,t} \times \bar{\mathcal{C}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}}, T\hat{X}^{02,t}, \otimes_{\mathbb{R}} \Omega_{\bar{\mathcal{C}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}} / \bar{\mathcal{M}}_{0, \text{Ext}(T) \cup \{1, \dots, r_\alpha\}}}^{0,1})$ as in Definition 3.10(iv) coinciding with the descent of $\nu_\alpha^{01} \sqcup \nu_\alpha^{12}$ for sufficiently large t , and varying smoothly with t .

We denote by $A_{\text{IV}}(T)$ the set of such thickening datums. Note that for any morphism $T \rightarrow T'$, there is a tautological identification $A_{\text{IV}}(T) = A_{\text{IV}}(T')$.

3.4 Index sets $\bar{A}(T)$

We define the index sets $\bar{A}(T)$ of the implicit atlases on $\bar{\mathcal{M}}(T)$ as unions of copies of the sets of thickening datums $A(T)$. Precisely, we define

$$\bar{A}(T) := \bigsqcup_{T' \subseteq T} A(T'), \quad (3.9)$$

where the disjoint union is over all *subtrees* T' of T . What qualifies as a subtree depends on which case (I), (II), (III), (IV) we are in, and we discuss each case individually. Informally, in cases (I) and (II), a subtree of T is a tree T' which appears in some concatenation yielding T , however this interpretation breaks down somewhat in cases (III) and (IV) due to connectedness issues.

Let us first explain what we mean by a subtree of an ordinary directed tree T (or forest), as opposed to an object of one of the categories \mathcal{S} . Choosing some subset of the vertices $V(T') \subseteq V(T)$, we define $E(T') \subseteq E(T)$ as all edges incident to an element of $V(T')$ (at least one of its endpoints). The resulting graph T' will be called a subtree when it is connected and non-empty. Below, we will use the word “physical subtree” for this notion, to distinguish it from the notion of a subtree of an object of some \mathcal{S} . An example is illustrated in Figure 9.




Figure 9: A tree and its six physical subtrees.

Definition 3.13 (Index set \bar{A}_I). For $T \in \mathcal{S}_I$, we define

$$\bar{A}_I(T) := \bigsqcup_{T \supseteq T' \in \mathcal{S}_I} A_I(T'). \quad (3.10)$$

More precisely, there is exactly one subtree $T \supseteq T' \in \mathcal{S}_I$ for every physical subtree of the underlying tree of T . A physical subtree becomes an object of \mathcal{S}_I simply by restricting the vertex and edge decorations from T . Note that there is no canonical way to choose basepoints for the input/output edges of subtrees $T' \in \mathcal{S}_I$, however the disjoint union (3.10) remains well-defined since the subgroup of $\text{Aut}(T')$ given by paths between basepoints acts trivially on $A_I(T')$.

Definition 3.14 (Index set \bar{A}_{II}). For $T \in \mathcal{S}_{II}$, we define

$$\bar{A}_{II}(T) := \bigsqcup_{T \supseteq T' \in \mathcal{S}_I^+} A_I^+(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_I^-} A_I^-(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{II}} A_{II}(T'). \quad (3.11)$$

More precisely, physical subtrees of T give rise to terms in this disjoint union as follows:

- (i) $T \supseteq T' \in \mathcal{S}_{II}$ are those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$.
- (ii) $T \supseteq T' \in \mathcal{S}_I^+$ are those for which all edges and vertices have $* = 0$.
- (iii) $T \supseteq T' \in \mathcal{S}_I^-$ are those for which all edges and vertices have $* = 1$.

Note that those physical subtrees $T' \subseteq T$ for which all edges and vertices have $* = 0$ and which have no negative external edges can be regarded as objects of \mathcal{S}_I^+ or \mathcal{S}_{II} , and thus appear twice in (3.11).

Definition 3.15 (Index set \bar{A}_{III}). For $T \in \mathcal{S}_{III}$, we define

$$\begin{aligned} \bar{A}_{III}(T) := & \bigsqcup_{T \supseteq T' \in \mathcal{S}_I^+} A_I^+(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_I^-} A_I^-(T') \\ & \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{II}^{t=0}} A_{II}^{t=0}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{II}^{t=1}} A_{II}^{t=1}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{III}} A_{III}(T'). \end{aligned} \quad (3.12)$$

More precisely, physical subtrees of T give rise to terms in this disjoint union as follows:

- (i) $T \supseteq T' \in \mathcal{S}_{III}$ are those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$.
- (ii) $T \supseteq T' \in \mathcal{S}_{II}^{t=0}$ are those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{0\}$.
- (iii) $T \supseteq T' \in \mathcal{S}_{II}^{t=1}$ are those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{1\}$.
- (iv) $T \supseteq T' \in \mathcal{S}_I^+$ are those for which all edges and vertices have $* = 0$.
- (v) $T \supseteq T' \in \mathcal{S}_I^-$ are those for which all edges and vertices have $* = 1$.

It should be emphasized that subtrees $T \supseteq T' \in \mathcal{S}_{\text{II}}^{t=0}$ (respectively $T \supseteq T' \in \mathcal{S}_{\text{II}}^{t=1}$) are only present when $\mathfrak{s}(T) = \{0\}$ (respectively $\mathfrak{s}(T) = \{1\}$). It should also be emphasized that we consider only *subtrees* (required to be connected and non-empty), rather than subforests.

Definition 3.16 (Index set \bar{A}_{IV}). For $T \in \mathcal{S}_{\text{IV}}$, we define

$$\begin{aligned} \bar{A}_{\text{IV}}(T) := & \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{I}}^0} A_{\text{I}}^0(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{I}}^2} A_{\text{I}}^2(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{II}}^{02}} A_{\text{II}}^{02}(T') \\ & \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{II}}^{01}} A_{\text{II}}^{01}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{I}}^1} A_{\text{I}}^1(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{II}}^{12}} A_{\text{II}}^{12}(T') \\ & \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\text{IV}}} A_{\text{IV}}(T'). \end{aligned} \quad (3.13)$$

More precisely, physical subtrees of T give rise to terms in this disjoint union as follows:

- (i) $T \supseteq T' \in \mathcal{S}_{\text{IV}}$ are those with $*(e^+) = 0$ and $*(e^-) = 2$ for $e^\pm \in E^\pm(T')$.
- (ii) $T \supseteq T' \in \mathcal{S}_{\text{II}}^{01}$ are those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{\infty\}$.
- (iii) $T \supseteq T' \in \mathcal{S}_{\text{II}}^{12}$ are those with $*(e^+) = 1$ and $*(e^-) = 2$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{\infty\}$.
- (iv) $T \supseteq T' \in \mathcal{S}_{\text{II}}^{02}$ are those with $*(e^+) = 0$ and $*(e^-) = 2$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{0\}$.
- (v) $T \supseteq T' \in \mathcal{S}_{\text{I}}^0$ are those for which all edges and vertices have $*$ = 0.
- (vi) $T \supseteq T' \in \mathcal{S}_{\text{I}}^1$ are those for which all edges and vertices have $*$ = 1.
- (vii) $T \supseteq T' \in \mathcal{S}_{\text{I}}^2$ are those for which all edges and vertices have $*$ = 2.

A morphism $T \rightarrow T'$ induces a natural inclusion

$$\bar{A}(T') \hookrightarrow \bar{A}(T), \quad (3.14)$$

since given $T \rightarrow T'$, any subtree $T'' \subseteq T'$ pulls back to a subtree of T .

For any concatenation $\{T_i\}_i$, there is a natural inclusion

$$\bigsqcup_i \bar{A}(T_i) \hookrightarrow \bar{A}(\#_i T_i). \quad (3.15)$$

3.5 Thickened moduli spaces

We now define the thickened moduli spaces for the implicit atlases on the moduli spaces $\bar{\mathcal{M}}$.

Definition 3.17 (Moduli space $\mathcal{M}_{\text{I}}(T)_I$). Let $T \in \mathcal{S}_{\text{I}}$ and let $I \subseteq \bar{A}_{\text{I}}(T)$ be finite. An *I-thickened pseudo-holomorphic building of type T* consists of the following data:

- (i) Domains C_v and punctures $p_{v,e}$ as in Definition 2.7(i). For $\alpha \in I$, let $C_\alpha := \bigsqcup_{v \in T_\alpha} C_v / \sim$, where $T \supseteq T_\alpha \in \mathcal{S}_{\text{I}}$ denotes the subtree indexing the term in (3.10) containing $\alpha \in I$, and \sim identifies $p_{v,e} \sim p_{v',e}$ for interior edges $v \xrightarrow{e} v'$ of T_α (see Figure 10).
- (ii) Maps u_v , asymptotic markers \tilde{b}_e , and matching isomorphisms m_e as in Definition 2.7(ii),(iii),(iv),(v).
- (iii) For all $\alpha \in I$, we require that $(u|C_\alpha) \pitchfork \hat{D}_\alpha$ with exactly r_α intersections, which together with $\{p_{v,e}\}_{v,e}$ stabilize C_α . By $(u|C_\alpha) \pitchfork \hat{D}_\alpha$, we mean that $\gamma_e \cap D_\alpha = \emptyset$ for edges $e \in E(T_\alpha)$, $(u|C_\alpha)^{-1}(\partial \hat{D}_\alpha) = \emptyset$, $(u|C_\alpha)^{-1}(\hat{D}_\alpha)$ does not contain any node, and $(du)_p : T_p C_\alpha \rightarrow T_{u(p)} \hat{Y} / T_{u(p)} \hat{D}_\alpha$ is surjective for $p \in (u|C_\alpha)^{-1}(\hat{D}_\alpha)$.




Figure 10: A tree with its six subtrees, and a corresponding pseudo-holomorphic building with its corresponding subbuildings.

- (iv) $\{\phi_\alpha : C_\alpha \rightarrow \overline{\mathcal{C}}_{0, E^{\text{ext}}(T_\alpha) \cup \{1, \dots, r_\alpha\}}\}_{\alpha \in I}$, where each ϕ_α maps C_α isomorphically onto a fiber of $\overline{\mathcal{C}}_{0, E^{\text{ext}}(T_\alpha) \cup \{1, \dots, r_\alpha\}}$, where C_α is equipped with its given marked points $p_{v,e}$ for $e \in E^{\text{ext}}(T_\alpha)$ and any marking of $(u|C_\alpha)^{-1}(\hat{D}_\alpha)$ with $\{1, \dots, r_\alpha\}$. Note that under (iii) above, choosing ϕ_α is equivalent to choosing a marking of $(u|C_\alpha)^{-1}(\hat{D}_\alpha)$ with $\{1, \dots, r_\alpha\}$.
- (v) $\{e_\alpha \in E_\alpha\}_{\alpha \in I}$.
- (vi) We require that u satisfy

$$\left(du + \sum_{\alpha \in I} \nu_\alpha(e_\alpha)(\phi_\alpha(\cdot), u(\cdot)) \right)_j^{0,1} = 0. \quad (3.16)$$

Note that the term in $\sum_{\alpha \in I}$ corresponding to α makes sense only over C_α , and we define it to be zero elsewhere.

An isomorphism $(\{C_v\}, \{p_{v,e}\}, \{u_v\}, \{\tilde{b}_e\}, \{m_e\}, \{\phi_\alpha\}, \{e_\alpha\}) \rightarrow (\{C'_v\}, \{p'_{v,e}\}, \{u'_v\}, \{\tilde{b}'_e\}, \{m'_e\}, \{\phi'_\alpha\}, \{e'_\alpha\})$ between I -thickened pseudo-holomorphic buildings of type T is defined as in Definition 2.7, with the additional requirements that $e_\alpha = e'_\alpha$ and $\phi_{\alpha,v} = \phi'_{\alpha,v} \circ i_v$ for $v \in T_\alpha$. We denote by $\mathcal{M}_I(T)_I$ the set of isomorphism classes of stable I -thickened pseudo-holomorphic buildings of type T . Note the tautological action of $\text{Aut}(T)$ on $\mathcal{M}(T)$ by changing the marking.

Note that the sum over α in (3.16) is supported away from the punctures $p_{v,e} \in C_v$, and hence u_v is genuinely \hat{J} -holomorphic near $p_{v,e}$. Note also that (3.16) is equivalent to the assertion that the graph $(\text{id}, u_v) : C_v \rightarrow C_v \times \hat{Y}$ is pseudo-holomorphic for the almost complex structure on $C_v \times \hat{Y}$ given by

$$\left(\begin{array}{cc} j_{C_v} & 0 \\ (\sum_{\alpha \in I} \nu_\alpha(e_\alpha)(\phi_\alpha(\cdot), \cdot))^{0,1} & \hat{J} \end{array} \right). \quad (3.17)$$

Hence solutions to the I -thickened pseudo-holomorphic curve equation enjoy all of the nice local *a priori* estimates from §2.8. Specifically, Lemma 2.18 applied to the graph of u implies that C^0 -estimates imply C^∞ -estimates, and Lemma 2.19 and Proposition 2.20 apply to u

itself since the perturbation terms ν_α are supported away from the nodes and the long cylinders of the domain.

Definition 3.18 (Moduli space $\mathcal{M}_{\text{II}}(T)_I$). Let $T \in \mathcal{S}_{\text{II}}$ and let $I \subseteq \bar{A}_{\text{II}}(T)$. An I -thickened pseudo-holomorphic building of type T consists of the following data:

- (i) Domains C_v and punctures $p_{v,e}$ as in Definition 2.7(i).
- (ii) Maps u_v , asymptotic markers \tilde{b}_e , and matching isomorphisms m_e as in Definition 2.7(ii),(iii),(iv),(v) (the target of u_v is $\hat{X}^{*(v)}$).
- (iii) For all $\alpha \in I$, we require that $(u|_{C_\alpha}) \pitchfork \hat{D}_\alpha$ with exactly r_α intersections, which together with $\{p_{v,e}\}_{v,e}$ stabilize C_α . More precisely, we intersect u_v with the relevant $\hat{D}_\alpha \subseteq \hat{X}$ or $\hat{D}_\alpha^\pm \subseteq \hat{Y}^\pm$ depending on $*(v)$.
- (iv) $\{\phi_\alpha : C_\alpha \rightarrow \bar{\mathcal{C}}_{0, E^{\text{ext}}(T_\alpha) \cup \{1, \dots, r_\alpha\}}\}_{\alpha \in I}$ as in Definition 3.17(iv).
- (v) We require that u satisfy (3.16), noting to use ν_α or ν_α^\pm depending on $*(v)$.

An *isomorphism* between pseudo-holomorphic buildings of type T is defined as in Definition 3.17, except that there is a translation $s_v \in \mathbb{R}$ only if v is a symplectization vertex. We denote by $\mathcal{M}_{\text{II}}(T)_I$ the set of isomorphism classes of stable I -thickened pseudo-holomorphic buildings of type T .

Definition 3.19 (Moduli space $\mathcal{M}_{\text{III}}(T)_I$). Let $T \in \mathcal{S}_{\text{III}}$ and let $I \subseteq \bar{A}_{\text{III}}(T)$ be finite. We denote by $\mathcal{M}_{\text{III}}(T)_I$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable I -thickened pseudo-holomorphic buildings of type T .

Definition 3.20 (Moduli space $\mathcal{M}_{\text{IV}}(T)_I$). Let $T \in \mathcal{S}_{\text{IV}}$ and let $I \subseteq \bar{A}_{\text{IV}}(T)$ be finite. We denote by $\mathcal{M}_{\text{IV}}(T)_I$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable I -thickened pseudo-holomorphic buildings of type T .

Definition 3.21 (Moduli spaces $\bar{\mathcal{M}}_I$). For $T \in \mathcal{S}$ and $I \subseteq \bar{A}(T)$, we define

$$\bar{\mathcal{M}}(T)_I := \bigsqcup_{\substack{T' \rightarrow T \\ \bar{\mathcal{M}}(T') \neq \emptyset}} \mathcal{M}(T')_I / \text{Aut}(T'/T). \quad (3.18)$$

Each such set $\bar{\mathcal{M}}(T)_I$ has a natural Gromov topology which is Hausdorff.

The stratifications (2.5) are clearly defined on the thickened moduli spaces $\bar{\mathcal{M}}(T)_I$. The tautological functorial structure (2.3) (combined with (3.14)) and (2.4) (combined with (3.15)) also exists for the thickened moduli spaces.

Definition 3.22 (Regularity of I -thickened pseudo-holomorphic buildings). As in Definition 2.39, given an I -thickened pseudo-holomorphic building of type T , we consider the linearized operator

$$E_I \oplus \bigoplus_{v \in V(T)} \tilde{W}^{k,2,\delta}(C_v, u_v^* T \hat{X}_v) \rightarrow \bigoplus_{v \in V(T)} W^{k-1,2,\delta}(\tilde{C}_v, u_v^*(T \hat{X}_v)_{j_v}) \otimes_{\mathbb{C}} \Omega_{\tilde{C}_v}^{0,1} \quad (3.19)$$

(recall that $E_I := \bigoplus_{\alpha \in I} E_\alpha$). Note that when $r_\alpha > 0$ for some $\alpha \in I$, we must take $k \geq 2$ so that $W^{k,2} \subseteq C^0$ (see [Ada75, Lemma 5.17]) so that the first variation of ϕ_α under $W^{k,2}$ deformations of u is defined.

A point in a moduli space $\bar{\mathcal{M}}(T)_I$ is called *regular* iff this linearized operator is surjective. We denote by $\bar{\mathcal{M}}(T)_I^{\text{reg}} \subseteq \bar{\mathcal{M}}(T)_I$ the locus of points which are regular and which have trivial automorphism group.

3.6 Implicit atlas structure

We now define the rest of the implicit atlas structure for the implicit atlases $\bar{A}(T)$ on $\bar{\mathcal{M}}(T)$.

Recall that the covering groups Γ_α and the obstruction spaces E_α are already built into the definition of $\bar{A}(T)$ from §§3.3–3.4. The thickened moduli spaces $\bar{\mathcal{M}}(T)_I$ and their regular loci $\bar{\mathcal{M}}(T)_I^{\text{reg}} \subseteq \bar{\mathcal{M}}(T)_I$ were defined in §3.5. The Kuranishi maps $s_\alpha : \bar{\mathcal{M}}(T)_I \rightarrow E_\alpha$ are the tautological maps picking out e_α (Definition 3.17(v)), and the footprint maps $\psi_{IJ} : (s_{J \setminus I}|_{\bar{\mathcal{M}}(T)_J})^{-1}(0) \rightarrow \bar{\mathcal{M}}(T)_I$ are the tautological forgetful maps.

The footprint $U_{IJ} \subseteq \bar{\mathcal{M}}(T)_I$ is defined as the locus of buildings which satisfy the transversality condition Definition 3.17(iii) for all $\alpha \in J$, and it follows immediately from this definition that ψ_{IJ} induces a bijection $(s_{J \setminus I}|_{\bar{\mathcal{M}}(T)_J})^{-1}(0)/\Gamma_{J \setminus I} \rightarrow U_{IJ}$. Inspection of the definition of the Gromov topology shows that this map is a homeomorphism and that $U_{IJ} \subseteq \bar{\mathcal{M}}(T)_I$ is open.

Theorem 3.23. *The above data define oriented implicit atlases $\bar{A}(T)$ on $\bar{\mathcal{M}}(T)$.*

Proof. Of the axioms which have not already been discussed (and which are nontrivial), the covering axiom follows from Propositions 3.26–3.30 below, and the openness and submersion axioms follow from Theorems 3.31 and 3.32 below. \square

3.7 Stabilization of pseudo-holomorphic curves

We now verify the covering axiom for the implicit atlases $\bar{A}(T)$ on $\bar{\mathcal{M}}(T)$. Namely, we show that the moduli spaces $\bar{\mathcal{M}}(T)$ are covered by the regular loci in their thickenings $\bar{\mathcal{M}}(T)_I^{\text{reg}}$. The essential content is to show that at every point in each moduli space $\bar{\mathcal{M}}(T)$, there exists a codimension two submanifold (as in Definition 3.9(iii)) which stabilizes the domain (i.e. satisfies Definition 3.17(iii)).

Lemma 3.24. *Let $u : D^2 \rightarrow (X, J)$ be J -holomorphic (for an almost complex manifold (X, J)). Then either $du : T_p C \rightarrow T_{u(p)} X$ is injective for some $p \in D^2$ or³ u is constant.*

Proof. If du is non-injective, it must be zero by J -holomorphicity. \square

Lemma 3.25. *Let $u : D^2 \rightarrow (\hat{Y}, \hat{J})$ be \hat{J} -holomorphic (for (Y, λ, J) as in Setup I). Denote by $\pi_\xi : T\hat{Y} \rightarrow \xi$ the projection under the splitting $T\hat{Y} = \xi \oplus \mathbb{R}R_\lambda \oplus \mathbb{R}\partial_s$. Then either $\pi_\xi du : T_p C \rightarrow \xi_{u(p)}$ is injective for some $p \in D^2$ or⁴ u factors through $\text{id} \times \gamma : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \times Y$ for some Reeb trajectory $\gamma : \mathbb{R} \rightarrow Y$.*

Proof. If $\pi_\xi du$ is non-injective, it must be zero. If $\pi_\xi du$ vanishes identically, then du is everywhere tangent to the 2-dimensional foliation of \hat{Y} by $\mathbb{R}R_\lambda \oplus \mathbb{R}\partial_s$, and thus u factors through one of its leaves. \square

Recall that $U_{IJ} \subseteq \bar{\mathcal{M}}(T)_I$ denotes the locus of buildings satisfying the transversality condition in Definition 3.17(iii) for $\alpha \in J \setminus I$.

³In fact, it is a standard (but nontrivial) fact that the first alternative can be strengthened to state that the zeroes of du form a discrete set (see [MS04, Lemma 2.4.1]).

⁴In fact, it is a standard (but nontrivial) fact that the first alternative can be strengthened to state that the zeroes of $\pi_\xi du$ form a discrete set (see Hofer–Wysocki–Zehnder [HWZ95, Proposition 4.1]).

Proposition 3.26. *For every $x \in \overline{\mathcal{M}}_I(T)$, there exists $\alpha \in A_I(T)$ such that $x \in U_{\emptyset, \{\alpha\}}$ and $\psi_{\emptyset, \{\alpha\}}^{-1}(x) \subseteq \overline{\mathcal{M}}_I(T)_{\{\alpha\}}^{\text{reg}}$.*

Proof. The point x is an isomorphism class of stable pseudo-holomorphic building of type $T' \rightarrow T$.

We claim that for all $v \in V(T')$, either C_v is stable (i.e. the degree of v is ≥ 3) or $\pi_\xi du_v$ is injective somewhere on C_v . To see this, suppose that $\pi_\xi du_v \equiv 0$ and apply Lemma 3.25 to $u_v : C_v \rightarrow \hat{Y}$. If the resulting Reeb trajectory $\gamma : \mathbb{R} \rightarrow Y$ is not a closed orbit, then u_v factors through $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \times Y$, and consideration of the positive puncture of C_v leads to a contradiction. Thus u_v factors through a trivial cylinder $\mathbb{R} \times S^1 \rightarrow \mathbb{R} \times Y$ for some simple Reeb orbit $\gamma : S^1 \rightarrow Y$. The map $C_v \rightarrow \mathbb{R} \times S^1$ is holomorphic, and it must have ramification points, as otherwise the building x would be unstable. It now follows from Riemann–Hurwitz that C_v is stable. Thus the claim is valid.

Now using the claim, it follows from Sard’s theorem that there exists $D_\alpha \subseteq Y$ satisfying Definition 3.17(iii) for some $r_\alpha \geq 0$. Now to show the existence of E_α and ν_α so that x is regular for the thickening datum $\alpha = (D_\alpha, r_\alpha, E_\alpha, \nu_\alpha)$, it suffices to show that

$$\bigoplus_{v \in V(T')} C_c^\infty(\tilde{C}_v \setminus (\{p_{v,e}\}_e \cup \tilde{N}_v), u_v^* T\hat{Y}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}_v}^{0,1}) \quad (3.20)$$

surjects onto the (finite-dimensional) cokernel of the linearized operator at x (where $\tilde{N}_v \subseteq \tilde{C}_v$ denotes the pre-images of the nodes of C_v). Let ε be a continuous linear functional on

$$\bigoplus_{v \in V(T')} W^{k-1,2,\delta}(\tilde{C}_v, u_v^* T\hat{Y}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}_v}^{0,1}) \quad (3.21)$$

which vanishes both on (3.20) and on the image of the linearized operator; it suffices to show that $\varepsilon = 0$. By Lemma 2.35, the space

$$\bigoplus_{v \in V(T')} C_c^\infty(\tilde{C}_v \setminus \{p_{v,e}\}_e, u_v^* T\hat{Y}_j \otimes_{\mathbb{C}} \Omega_{\tilde{C}_v}^{0,1}) \quad (3.22)$$

is dense in (3.21), so it suffices to show that ε vanishes as a distribution. Since ε annihilates (3.20), it is supported over the finite set \tilde{N}_v . When $k = 2$, no nonzero distribution of finite support is continuous on $W^{k-1,2}$ (recall Lemma 2.41 and note that $W^{1,2} \not\subseteq C^0$ [Ada75, Example 5.26]). Since we did not use the fact that ε annihilates the image of the linearized operator, we have in fact shown that (3.20) is dense in (3.21). \square

Remark 3.27. There are two alternative ways to conclude the proof of Proposition 3.26. First, using the fact that ε annihilates the image of D , it follows that $D^* \varepsilon$ (D^* denotes the formal adjoint of D) is a linear combination of δ -functions supported over \tilde{N}_v . But such δ -functions live in H^{-2} (by Sobolev embedding $H^2 \subseteq C^0$ [Ada75, Lemma 5.17]) and the formal adjoint is elliptic, so this means $\varepsilon \in H^{-1}$, which contains no nonzero distributions supported at single points as argued above. Second, we could define the linearized operator for nodal curves using weights near the nodes (as we do in §5.2.1 to prove gluing) and simply cite Lemma 2.35 to conclude that (3.20) is dense in the target weighted Sobolev space for all $k \geq 2$.

Proposition 3.28. *For every $x \in \overline{\mathcal{M}}_{\text{II}}(T)$, there exists $\alpha \in A_{\text{II}}(T)$ such that $x \in U_{\emptyset, \{\alpha\}}$ and $\psi_{\emptyset, \{\alpha\}}^{-1}(x) \subseteq \overline{\mathcal{M}}_{\text{II}}(T)_{\{\alpha\}}^{\text{reg}}$.*

Proof. The point x is an isomorphism class of stable pseudo-holomorphic building of type $T' \rightarrow T$.

As in the proof of Proposition 3.26, for $v \in V(T')$ with $*(v) = 00$ or $*(v) = 11$, either C_v is stable or $\pi_\varepsilon du_v$ is injective somewhere on C_v . For $*(v) = 01$, every irreducible component of C_v is either stable or has a point where du_v injective by Lemma 3.24.

Now we can find a (compact) codimension two submanifold $\hat{D}_\alpha \subseteq \hat{X}$ such that $u \pitchfork \hat{D}_\alpha$ with intersections stabilizing every C_v with $*(v) = 01$. Now we consider the remaining unstable C_v ($*(v) = 00$ or $*(v) = 11$), and we choose codimension two submanifolds $D_\alpha^\pm \subseteq Y^\pm$ stabilizing these. We then cutoff \hat{D}_α^\pm near infinity in \hat{X} and add this to \hat{D}_α . Thus $u \pitchfork \hat{D}_\alpha$ with r_α intersections which stabilize C .

Now (E_α, ν_α) are constructed as in Proposition 3.26. □

Proposition 3.29. *For every $x \in \overline{\mathcal{M}}_{\text{III}}(T)$, there exist $\alpha_i \in A_{\text{III}}(T_i)$ such that $x \in U_{\emptyset, \{\alpha_i\}_i}$ and $\psi_{\emptyset, \{\alpha_i\}_i}^{-1}(x) \subseteq \overline{\mathcal{M}}_{\text{III}}(T)_{\{\alpha_i\}_i}^{\text{reg}}$ (writing $T = \bigsqcup_i T_i$ with T_i connected and non-empty).*

Proof. Apply Proposition 3.28 to each subbuilding of type T_i to get α_i . □

Proposition 3.30. *For every $x \in \overline{\mathcal{M}}_{\text{IV}}(T)$, there exist $\alpha_i \in A_{\text{IV}}(T_i)$ such that $x \in U_{\emptyset, \{\alpha_i\}_i}$ and $\psi_{\emptyset, \{\alpha_i\}_i}^{-1}(x) \subseteq \overline{\mathcal{M}}_{\text{IV}}(T)_{\{\alpha_i\}_i}^{\text{reg}}$ (writing $T = \bigsqcup_i T_i$ with T_i connected and non-empty).*

Proof. Essentially the same as the proof of Propositions 3.28–3.29. □

3.8 Local structure of thickened moduli spaces via local models G

We now state the precise sense in which the spaces $G_{T'/T}$ are local topological models for the regular loci in the thickened moduli spaces $\overline{\mathcal{M}}(T)$. We also state the compatibility of this local topological structure with the natural maps on orientation lines discussed earlier. These statements (which contain those from §2.14 as a special case) are in essence a gluing theorem, and their proofs are given in §5. They imply that the openness and submersion axioms hold for the implicit atlases we have defined, and they give canonical isomorphisms $\mathfrak{o}_{\overline{\mathcal{M}}(T)} = \mathfrak{o}_T$.

The (Banach space) implicit function theorem implies that $\mathcal{M}(T)_I^{\text{reg}}$ is a (smooth) manifold of dimension $\mu(T) - \#V_s(T) + \dim \mathfrak{s}(T) + \dim E_I$ over the locus without nodes.

Theorem 3.31 (Local structure of $\overline{\mathcal{M}}(T)_I^{\text{reg}}$). *Let $I \subseteq J \subseteq \bar{A}(T)$. Let $x_0 \in \overline{\mathcal{M}}(T)_J$ be of type $T' \rightarrow T$ and satisfy $s_{J \setminus I}(x_0) = 0$ and $\psi_{IJ}(x_0) \in \overline{\mathcal{M}}(T)_I^{\text{reg}}$. Then $\mu(T) - \#V_s(T') - 2\#N(x_0) + \dim E_I \geq 0$ and there is a local homeomorphism*

$$(G_{T'/T} \times E_{J \setminus I} \times \mathbb{C}^{N(x_0)} \times \mathbb{R}^{\mu(T) - \#V_s(T') - 2\#N(x_0) + \dim E_I}, (0, 0, 0)) \rightarrow (\overline{\mathcal{M}}(T)_J, x_0) \quad (3.23)$$

whose image lands in $\overline{\mathcal{M}}(T)_J^{\text{reg}}$ and which commutes with the maps from both sides to $\mathfrak{S}_{T'/T} \times \mathfrak{s}(T) \times E_{J \setminus I}$ (as well as the stratifications by number of nodes).

Proof. See §§5.1–5.3. □

The (Banach space) implicit function theorem moreover gives a canonical identification

$$\mathfrak{o}_{\overline{\mathcal{M}}(T)_I^{\text{reg}}} = \mathfrak{o}_T \otimes \mathfrak{o}_{E_I}. \quad (3.24)$$

More precisely, this identification is made over the locus in $\mathcal{M}(T)_I^{\text{reg}}$ without nodes, and has a unique continuous extension to all of $\overline{\mathcal{M}}(T)_I^{\text{reg}}$ by virtue of the local topological description in Theorem 3.31. This identification is easily seen to be compatible with $\psi_{I,J}$ and with concatenations. It is also compatible with morphisms $T' \rightarrow T$ in the following precise sense.

Theorem 3.32 (Compatibility of the “analytic” and “geometric” maps on orientations). *The following diagram commutes:*

$$\begin{array}{ccc} \mathfrak{o}_{\overline{\mathcal{M}}(T')_I^{\text{reg}}} \otimes \mathfrak{o}_{G_{T'/T}} \otimes \mathfrak{o}_{s(T')}^\vee & \xrightarrow{(3.24)} & \mathfrak{o}_{T'} \otimes \mathfrak{o}_{E_I} \otimes \mathfrak{o}_{G_{T'/T}} \otimes \mathfrak{o}_{s(T')}^\vee \\ \parallel (3.23) & & \parallel (2.61) \\ \mathfrak{o}_{\overline{\mathcal{M}}(T)_I^{\text{reg}}} & \xrightarrow{(3.24)} & \mathfrak{o}_T \otimes \mathfrak{o}_{E_I}, \end{array} \quad (3.25)$$

where the left vertical map is the “geometric” map induced by the local topological structure of $\overline{\mathcal{M}}(T)_I^{\text{reg}}$ coming from (3.23), and the right vertical map is the “analytic” map (2.61) defined earlier via the “kernel gluing” operation.

Proof. See §5.4. □

4 Virtual fundamental cycles

In this section, we prove Theorem 1.1 as stated in the introduction. Specifically, Theorem 1.1 follows by combining Definition 4.32 with Remark 4.15, Proposition 4.33, and Proposition 4.34.

4.1 Review of the VFC package

We begin with a review of the framework introduced in [Par16] for defining the virtual fundamental cycle of a space equipped with an implicit atlas. We state here all the results from [Par16] which we will be appealing to in the rest of this section. We use \mathbb{Q} coefficients throughout, and orientation lines are tacitly tensored up to \mathbb{Q} .

4.1.1 The basic formalism

Let X be a compact Hausdorff space equipped with an implicit atlas A with boundary in the sense of Definition 3.3 (recall also Remark 3.4), oriented with respect to \mathfrak{o} and of virtual dimension d (meaning these are the orientation line and virtual dimension of the top stratum). Assume for the moment that A is finite, and let

$$C_\bullet(E; A) := C_{\dim E_A + \bullet}(E_A, E_A \setminus 0; \mathfrak{o}_{E_A}^\vee)^{\Gamma_A} \quad (4.1)$$

following [Par16, Definition 4.2.4] (the superscript indicates taking Γ_A -invariants); note that there is a canonical isomorphism $H_\bullet(E; A) = \mathbb{Q}$ (concentrated in degree zero). Now for any (X, A) as above, [Par16, Definition 4.2.6] defines a cochain complex

$$C_{\text{vir}}^\bullet(X \text{ rel } \partial; A) \quad (4.2)$$

together with a canonical pushforward map

$$C_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) \rightarrow C_{-\bullet}(E; A), \quad (4.3)$$

and [Par16, Theorem 4.3.4] provides a canonical isomorphism

$$H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) = \check{H}^\bullet(X, j_! \mathfrak{o}), \quad (4.4)$$

where $j : X \setminus \partial X \hookrightarrow X$. The map (4.3) is the (chain level!) virtual fundamental cycle of X (taking cohomology and combining with (4.4) yields a map $\check{H}^d(X, j_! \mathfrak{o}) \rightarrow \mathbb{Q}$ which is, by definition, integration over the virtual fundamental class of X).

Remark 4.1. As an aside, let us provide some very brief motivation for this formalism of virtual fundamental cycles (the reader may refer to [Par16, §§1–2] for details). Suppose that for some $\alpha \in A$, the α -thickened space X_α is everywhere regular (namely $X_\alpha = X_\alpha^{\text{reg}}$), has no isotropy (meaning $\Gamma_\alpha = 1$), and has footprint all of X (that is $U_{\emptyset, \{\alpha\}} = X_\emptyset$). In this situation, we have $X = s_\alpha^{-1}(0)$ for $s_\alpha : X_\alpha \rightarrow E_\alpha$, and we obviously want to define $[X]^{\text{vir}} := [X_\alpha] \cap s_\alpha^* \tau_{E_\alpha}$, where τ_{E_α} denotes the Thom class of E_α . Another way of stating the equality $[X]^{\text{vir}} = [X_\alpha] \cap s_\alpha^* \tau_{E_\alpha}$ is to say that evaluation on $[X]^{\text{vir}}$ is given by the composition

$$\check{H}^\bullet(X) = H_{\dim X_\alpha - \bullet}(X_\alpha, X_\alpha \setminus X) \xrightarrow{(s_\alpha)^*} H_{\dim X_\alpha - \bullet}(E_\alpha, E_\alpha \setminus 0) = \mathbb{Q}[\dim E_\alpha - \dim X_\alpha] \quad (4.5)$$

(ignoring orientation lines for the moment), where the first equality is a form of Poincaré duality. In general, the complex $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$ is built out of relative chain groups of (roughly speaking) the regular loci in the thickened spaces X_I^{reg} , and (4.3)–(4.4) are a direct generalization of (4.5).

For an inclusion of implicit atlases $A \subseteq A'$ on the same space X , there are canonical quasi-isomorphisms

$$C_{\text{vir}}^\bullet(X \text{ rel } \partial; A) \xrightarrow{\sim} C_{\text{vir}}^\bullet(X \text{ rel } \partial; A'), \quad (4.6)$$

$$C_\bullet(E; A) \xrightarrow{\sim} C_\bullet(E; A'), \quad (4.7)$$

which compose as expected and are compatible with (4.3) (on the chain level), see [Par16, Definition 4.2.7].⁵ These maps allow us to remove our assumption that A is finite, namely by defining $C_\bullet(E; A)$ and $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$ for general A by taking the direct limit over the collection of finite subsets of A .

Lemma 4.2 ([Par16, Lemma 5.2.6]). *If $X = X^{\text{reg}}$, then the map $\check{H}^d(X; j_! \mathfrak{o}) \rightarrow \mathbb{Q}$ from (4.3) is evaluation on the ordinary fundamental class $[X] \in H_d(X, \partial X; \mathfrak{o}^\vee)$.*

⁵More precisely, the maps defined in [Par16] carry an extra factor of $\otimes C_\bullet(E; A' \setminus A)$ on the left. Let $[E_\alpha] \in C_0(E; \alpha)$ be the fundamental cycle obtained by pulling back some $[\mathbb{R}^n] \in C_n(\mathbb{R}^n, \mathbb{R}^n \setminus 0; \mathfrak{o}_{\mathbb{R}^n}^\vee)$ (fixed once and for all) under the specified identification $E_\alpha = \mathbb{R}^{n_\alpha}$ and averaging over Γ_α (alternatively, we could modify the definition of a thickening datum to include the data of a fundamental cycle $[E_\alpha] \in C_0(E; \alpha)$). Now (4.6)–(4.7) are defined by pre-composing the maps in [Par16] with $\otimes \bigotimes_{\alpha \in A' \setminus A} [E_\alpha]$.

4.1.2 Homotopy sheaves

Recall from [Par16, Definition A.2.5] that a presheaf of complexes \mathcal{F}^\bullet on the category of compact subsets of a locally compact Hausdorff space X is called a *homotopy sheaf* iff it satisfies the following three axioms:

- (hSh1) $\mathcal{F}^\bullet(\emptyset)$ is acyclic.
- (hSh2) $[\mathcal{F}^\bullet(K_1 \cup K_2) \rightarrow \mathcal{F}^{\bullet-1}(K_1) \oplus \mathcal{F}^{\bullet-1}(K_2) \rightarrow \mathcal{F}^{\bullet-2}(K_1 \cap K_2)]$ is acyclic for all $K_1, K_2 \subseteq X$.
- (hSh3) $\varinjlim_{\substack{K \subseteq U \\ U \text{ open}}} \mathcal{F}^\bullet(\overline{U}) \rightarrow \mathcal{F}^\bullet(K)$ is a quasi-isomorphism for all $K \subseteq X$.

A homotopy sheaf should be thought of as a complex of presheaves which calculates its own cohomology, in the sense that the natural map $H^\bullet \mathcal{F}^\bullet(K) \rightarrow \check{H}^\bullet(K, \mathcal{F}^\bullet)$ is an isomorphism [Par16, Proposition A.4.14].

Recall from [Par16, Definition A.5.1] that a homotopy sheaf \mathcal{F}^\bullet is called *pure* iff it satisfies the following two axioms:

$$\text{For all } p \in X \text{ and } i \neq 0, \text{ we have } H^i \mathcal{F}^\bullet(\{p\}) = 0. \quad (4.8)$$

$$\text{For all } p \in X, \text{ there exists a neighborhood } U \subseteq X \text{ of } p \text{ and an integer } i_0 \quad (4.9)$$

such that $H^i \mathcal{F}^\bullet(K) = 0$ for all $K \subseteq U$ and $i \leq i_0$.

From a conceptual standpoint, it is condition (4.8) which is most significant, and it should be thought of as saying that \mathcal{F}^\bullet is (stalkwise) quasi-isomorphic to a sheaf $H^0 \mathcal{F}^\bullet$ [Par16, Lemma A.5.3]. In particular, for any pure homotopy sheaf \mathcal{F}^\bullet , there is a canonical isomorphism

$$H^\bullet \mathcal{F}^\bullet(X) = \check{H}^\bullet(X; H^0 \mathcal{F}^\bullet) \quad (4.10)$$

by [Par16, Proposition A.5.4].

The complex $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$ extends naturally to a presheaf of complexes

$$K \mapsto C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)_K \quad (4.11)$$

on compact subsets $K \subseteq X$, with complex of global sections $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A) = C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)_X$.⁶ By [Par16, Proposition 4.3.3], this presheaf is a pure homotopy sheaf. In more detail, axiom (hSh1) is trivial, axiom (hSh2) follows from Mayer–Vietoris for singular chains, axiom (hSh3) follows from compactness of simplices, axiom (4.8) is a local homology calculation (essentially $H_{\dim M - \bullet}(M, M \setminus p) = \mathfrak{o}_{M,p}$), and axiom (4.9) follows from the finite-dimensionality of the thickened moduli spaces comprising the implicit atlas A .

The sheaf $K \mapsto H_{\text{vir}}^0(X \text{ rel } \partial; A)_K$ associated to (4.11) is canonically isomorphic to $j_! \mathfrak{o}$ on X . Furthermore, this isomorphism has a natural local description in terms of Poincaré duality [Par16, Lemma 4.3.2]. The isomorphism (4.4) is thus a special case of (4.10).

4.1.3 Stratifications

Let X be a compact Hausdorff space equipped with an implicit atlas A with oriented cell-like stratification in the sense of Definition 3.2. The implicit atlas on X induces implicit atlases on each closed stratum $X_{/t}$ for $t \in \mathcal{T}$, simply by defining $(X_{/t})_I := (X_I)_{/t}$. To understand

⁶In [Par16], the relatively inferior notation $C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)$ is used in place of $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)_K$.

how the virtual fundamental cycles of each of the strata $X_{/t}$ fit together, [Par16, Definition 6.2.2] defines

$$C_{\text{vir}}^\bullet(X, \mathcal{J}; A) := \bigoplus_{t \in \mathcal{J}} [C_{\text{vir}}^\bullet(X_{/t} \text{ rel } \partial; A) \otimes \mathfrak{o}_t^\vee \otimes \mathfrak{o}_{t^{\max}}]^{\text{Aut}(t)}, \quad (4.12)$$

(the superscript indicates taking $\text{Aut}(t)$ -invariants) equipped with the differential given by (the internal differential plus) pushing forward along all codimension one maps $t \rightarrow t'$ in \mathcal{J} (and contracting on the left with a chosen orientation of \mathbb{R}). There is now a canonical isomorphism

$$H_{\text{vir}}^\bullet(X; A) = \check{H}^\bullet(X; \mathfrak{o}_{t^{\max}}). \quad (4.13)$$

This is shown in [Par16, Propositions 6.2.3 and 4.3.3] for stratifications by posets, and the reasoning there applies without modification to the present setting of stratifications by categories in the sense of Definition 2.15.

To be more precise, the proof of the isomorphism (4.13) proceeds by considering the presheaf of complexes on X given by

$$K \mapsto C_{\text{vir}}^\bullet(X, \mathcal{J}; A)_K := \bigoplus_{t \in \mathcal{J}} [C_{\text{vir}}^\bullet(X_{/t} \text{ rel } \partial; A)_{K/t} \otimes \mathfrak{o}_t^\vee \otimes \mathfrak{o}_{t^{\max}}]^{\text{Aut}(t)}, \quad (4.14)$$

with differential as in (4.12). It is easy to see that $K \mapsto C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)_K$ being a homotopy sheaf implies the same for $K \mapsto C_{\text{vir}}^\bullet(X, \mathcal{J}; A)_K$ [Par16, Lemma A.2.11], and it is straightforward to argue that purity also passes. The sheaf $K \mapsto H_{\text{vir}}^0(X, \mathcal{J}; A)_K$ now has a filtration whose associated graded is the direct sum

$$K \mapsto \bigoplus_{t \in \mathcal{J}} [H_{\text{vir}}^0(X_{/t} \text{ rel } \partial; A)_{K/t} \otimes \mathfrak{o}_t^\vee \otimes \mathfrak{o}_{t^{\max}}]^{\text{Aut}(t)} \quad (4.15)$$

There is thus a natural stalkwise isomorphism of $K \mapsto H_{\text{vir}}^0(X, \mathcal{J}; A)_K$ with the constant sheaf $\mathfrak{o}_{t^{\max}}$ on X , and a local construction [Par16, Example 6.2.1] along with [Par16, Lemma 4.3.2] shows that it comes from an isomorphism of sheaves.

4.1.4 Products

For spaces X and Y with implicit atlases A and B , by [Par16, Definition 6.3.2] there is a canonical product map

$$C_{\text{vir}}^\bullet(X \text{ rel } \partial; A) \otimes C_{\text{vir}}^\bullet(Y \text{ rel } \partial; B) \rightarrow C_{\text{vir}}^\bullet(X \times Y \text{ rel } \partial; A \sqcup B), \quad (4.16)$$

compatible with (4.3) and (4.6)–(4.7). Under the isomorphisms (4.4), this map is simply the usual Künneth product map. This is not mentioned explicitly in [Par16], rather it follows from Lemma 4.3 below and recalling from [Par16, Definition 6.3.2] that the map (4.16) extends naturally to a family of maps

$$C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)_K \otimes C_{\text{vir}}^\bullet(Y \text{ rel } \partial; B)_{K'} \rightarrow C_{\text{vir}}^\bullet(X \times Y \text{ rel } \partial; A \sqcup B)_{K \times K'} \quad (4.17)$$

compatible with restriction.

Lemma 4.3. *Let $\mathcal{F}^\bullet, \mathcal{G}^\bullet, \mathcal{H}^\bullet$ be pure homotopy sheaves on spaces $X, Y, X \times Y$, respectively. Let $\mathcal{F}^\bullet(K) \otimes \mathcal{G}^\bullet(K') \rightarrow \mathcal{H}^\bullet(K \times K')$ be a collection of maps compatible with restriction. Then the following diagram commutes:*

$$\begin{array}{ccc}
H^\bullet \mathcal{F}^\bullet(X) \otimes H^\bullet \mathcal{G}^\bullet(Y) & \longrightarrow & H^\bullet \mathcal{H}^\bullet(X \times Y) \\
\downarrow (4.10) & & \downarrow (4.10) \\
\check{H}^\bullet(X; H^0 \mathcal{F}^\bullet) \otimes \check{H}^\bullet(Y; H^0 \mathcal{G}^\bullet) & \longrightarrow & \check{H}^\bullet(X \times Y; H^0 \mathcal{H}^\bullet),
\end{array} \tag{4.18}$$

where the bottom horizontal arrow is the usual Künneth cup product map.

Proof. We consider the following commutative diagram:

$$\begin{array}{ccccc}
H^\bullet \mathcal{F}^\bullet(X) \otimes H^\bullet \mathcal{G}^\bullet(Y) & \longrightarrow & H^\bullet [\mathcal{F}^\bullet(X) \otimes \mathcal{G}^\bullet(Y)] & \longrightarrow & H^\bullet \mathcal{H}^\bullet(X \times Y) \\
\downarrow \sim & & \downarrow & & \downarrow \sim \\
\check{H}^\bullet(X; \mathcal{F}^\bullet) \otimes \check{H}^\bullet(Y; \mathcal{G}^\bullet) & \longrightarrow & H^\bullet [\check{C}^\bullet(X; \mathcal{F}^\bullet) \otimes \check{C}^\bullet(Y; \mathcal{G}^\bullet)] & \longrightarrow & \check{H}^\bullet(X \times Y, \mathcal{H}^\bullet) \\
\uparrow \sim & & \uparrow & & \uparrow \sim \\
\check{H}^\bullet(X; \tau_{\leq 0} \mathcal{F}^\bullet) \otimes \check{H}^\bullet(Y; \tau_{\leq 0} \mathcal{G}^\bullet) & \rightarrow & H^\bullet [\check{C}^\bullet(X; \tau_{\leq 0} \mathcal{F}^\bullet) \otimes \check{C}^\bullet(Y; \tau_{\leq 0} \mathcal{G}^\bullet)] & \rightarrow & \check{H}^\bullet(X \times Y, \tau_{\leq 0} \mathcal{H}^\bullet) \\
\downarrow \sim & & \downarrow & & \downarrow \sim \\
\check{H}^\bullet(X; H^0 \mathcal{F}^\bullet) \otimes \check{H}^\bullet(Y; H^0 \mathcal{G}^\bullet) & \rightarrow & H^\bullet [\check{C}^\bullet(X; H^0 \mathcal{F}^\bullet) \otimes \check{C}^\bullet(Y; H^0 \mathcal{G}^\bullet)] & \rightarrow & \check{H}^\bullet(X \times Y, H^0 \mathcal{H}^\bullet).
\end{array} \tag{4.19}$$

Note that there is a natural map $\tau_{\leq 0} \mathcal{F}^\bullet \otimes \tau_{\leq 0} \mathcal{G}^\bullet \rightarrow \tau_{\leq 0}(\mathcal{F}^\bullet \otimes \mathcal{G}^\bullet) \rightarrow \tau_{\leq 0} \mathcal{H}^\bullet$, but no such natural map with $\tau_{\geq 0}$ in place of $\tau_{\leq 0}$. It is therefore important that the vertical columns above are of the form $\mathcal{F}^\bullet \leftarrow \tau_{\leq 0} \mathcal{F}^\bullet \rightarrow H^0 \mathcal{F}^\bullet$ rather than $\mathcal{F}^\bullet \rightarrow \tau_{\geq 0} \mathcal{F}^\bullet \leftarrow H^0 \mathcal{F}^\bullet$.

Now the far right and far left vertical maps above are all isomorphisms and define (4.10) (see the proof of [Par16, Proposition A.5.4]). The outer square is exactly the desired diagram, so we are done. \square

4.2 Sketch of the construction

We now sketch the construction of virtual fundamental cycles on the moduli spaces $\overline{\mathcal{M}}(T)$. The remainder of this section is then devoted to turning this sketch into an actual proof.

The virtual fundamental cycle of a space with an implicit atlas is represented by the pushforward map (4.3). Specializing this to the case at hand, the virtual fundamental class $[\overline{\mathcal{M}}(T)]^{\text{vir}}$ is represented by a canonical map

$$C_{\text{vir}}^{\text{vdim}(T)+\bullet}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \bar{A}(T)) \rightarrow C_{-\bullet}(E; \bar{A}(T)). \tag{4.20}$$

The domain of this map is quasi-isomorphic to the complex of Čech cochains $\check{C}^{\text{vdim}(T)+\bullet}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \mathfrak{o}_T)$ by (4.4), and the target is quasi-isomorphic to \mathbb{Q} . Thus, up to quasi-isomorphism, (4.20) can be regarded as a map

$$\check{C}^{\text{vdim}(T)+\bullet}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \mathfrak{o}_T) \rightarrow \mathbb{Q}, \tag{4.21}$$

which is nothing other than a cycle

$$[\overline{\mathcal{M}}(T)]^{\text{vir}} \in \overline{\mathcal{C}}_{\text{vdim}(T)}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \mathfrak{o}_T^\vee). \quad (4.22)$$

We note that the dual of Čech cochains \check{C}^\bullet is Steenrod chains $\overline{\mathcal{C}}_\bullet$ (see [Par16, §A.9] and the references therein), although this level of precision is not really relevant for the present sketch.

Now the functoriality of (4.20) with respect to morphisms and concatenations in \mathcal{S} implies that the virtual fundamental cycles (4.22) satisfy

$$\partial[\overline{\mathcal{M}}(T)]^{\text{vir}} = \sum_{\text{codim}(T'/T)=1} \frac{1}{|\text{Aut}(T'/T)|} [\overline{\mathcal{M}}(T')]^{\text{vir}}, \quad (4.23)$$

$$[\overline{\mathcal{M}}(\#_i T_i)]^{\text{vir}} = \frac{1}{|\text{Aut}(\{T_i\}_i/\#T_i)|} \prod_i [\overline{\mathcal{M}}(T_i)]^{\text{vir}}, \quad (4.24)$$

which clearly imply the desired master equations.

To turn this sketch into an actual construction, we must carry out the above arguments on the chain level, with all the necessary chain level functoriality with respect to morphisms and concatenations in \mathcal{S} . This is not completely trivial, since many of the chain maps we are given go in the wrong direction (basically, we need to “invert quasi-isomorphisms”). The rest of this section is devoted to performing the necessary algebraic manipulations.

4.3 \mathcal{S} -modules

We begin by introducing the notion of an \mathcal{S} -module. This notion formalizes the way in which the moduli spaces $\overline{\mathcal{M}}(T)$ fit together under morphisms and concatenations in \mathcal{S} (namely (2.3)–(2.4)) and will be used below to efficiently encode the identities (4.23)–(4.24) satisfied by the virtual fundamental cycles. In fact, many (or even most) of the objects introduced and studied earlier in this paper are \mathcal{S} -modules, as we explain in the examples which follow the definition. The notion of an \mathcal{S} -module plays a key organizational role in what follows.

Definition 4.4. An object $T \in \mathcal{S}$ will be called *effective* iff $\overline{\mathcal{M}}(T) \neq \emptyset$.

Note that (1) for any morphism $T \rightarrow T'$, if T is effective, then so is T' , and (2) for any concatenation $\{T_i\}_i$, every T_i is effective iff $\#_i T_i$ is effective. *For the remainder of this section, we will use \mathcal{S} to denote the full subcategory spanned by effective objects.*

Definition 4.5. An \mathcal{S}_I -module X_I valued in a symmetric monoidal category \mathcal{C}^\otimes consists of the following data:

- (i) A functor $X_I : \mathcal{S}_I \rightarrow \mathcal{C}$.
- (ii) For every concatenation of $\{T_i\}_i$ in \mathcal{S}_I , a morphism

$$\bigotimes_i X_I(T_i) \rightarrow X_I(\#_i T_i), \quad (4.25)$$

such that the following diagrams commute:

$$\begin{array}{ccccc}
\bigotimes_i X_I(T_i) & \xrightarrow{(4.25)} & X_I(\#_i T_i) & & \bigotimes_i X_I(\#_j T_{ij}) \\
\downarrow & & \downarrow & \nearrow_{\otimes_i (4.25)} & \searrow_{(4.25)} \\
\bigotimes_i X_I(T'_i) & \xrightarrow{(4.25)} & X_I(\#_i T'_i) & \xrightarrow{(4.25)} & X_I(\#_{ij} T_{ij}), \\
& & & \nearrow_{i,j} & \\
& & & \bigotimes_{i,j} X_I(T_{ij}) &
\end{array} \tag{4.26}$$

for any morphism of concatenations $\{T_i\}_i \rightarrow \{T'_i\}_i$ and any composition of concatenations, respectively.

A morphism of \mathcal{S}_I -modules is a natural transformation of functors compatible with (4.25).

An $\mathcal{S}_I^{\text{op}}$ -module is defined similarly, except that $X_I : \mathcal{S}_I^{\text{op}} \rightarrow \mathcal{C}$ and the vertical arrows in the leftmost diagram in (4.26) are reversed. Note that an $\mathcal{S}_I^{\text{op}}$ -module valued in \mathcal{C} is *not* the same thing as an \mathcal{S}_I -module valued in \mathcal{C}^{op} (this is analogous to the difference between a lax monoidal functor and an oplax monoidal functor).

Example 4.6. The functor $\overline{\mathcal{M}}_I$ is an \mathcal{S}_I -module (valued in the category of compact Hausdorff spaces, with the product symmetric monoidal structure).

Example 4.7. The functor \mathfrak{o}° is an \mathcal{S}_I -module (valued in the category of orientation lines and isomorphisms, with the super tensor product symmetric monoidal structure).

Example 4.8. The functor \bar{A}_I is an $\mathcal{S}_I^{\text{op}}$ -module (valued in the category of sets, with the disjoint union symmetric monoidal structure).

Example 4.9. The functor $T \mapsto (\mathcal{S}_I)_{/T}$ is an \mathcal{S}_I -module (valued in the category of categories, with the product symmetric monoidal structure).

Definition 4.10. An \mathcal{S}_{II} -module X_{II} valued in \mathcal{C}^\otimes consists of the following data:

- (i) An \mathcal{S}_I^+ -module X_I^+ valued in \mathcal{C}^\otimes .
- (ii) An \mathcal{S}_I^- -module X_I^- valued in \mathcal{C}^\otimes .
- (iii) A functor $X_{II} : \mathcal{S}_{II} \rightarrow \mathcal{C}$.
- (iv) For every concatenation $\{T_i\}_i$ in \mathcal{S}_{II} , a morphism

$$\bigotimes_{T_i \in \mathcal{S}_I^+} X_I^+(T_i) \otimes \bigotimes_{T_i \in \mathcal{S}_I^-} X_I^-(T_i) \otimes \bigotimes_{T_i \in \mathcal{S}_{II}} X_{II}(T_i) \rightarrow X_{II}(\#_i T_i), \tag{4.27}$$

such that the following diagrams commute:

$$\begin{array}{ccccc}
\bigotimes_i X_{I/II}(T_i) & \longrightarrow & X_{I/II}(\#_i T_i) & & \bigotimes_i X_{I/II}(\#_j T_{ij}) \\
\downarrow & & \downarrow & \nearrow & \searrow \\
\bigotimes_i X_{I/II}(T'_i) & \longrightarrow & X_{I/II}(\#_i T'_i) & \longrightarrow & X_{I/II}(\#_{ij} T_{ij}), \\
& & & \nearrow_{i,j} & \\
& & & \bigotimes_{i,j} X_{I/II}(T_{ij}) &
\end{array} \tag{4.28}$$

for any morphism or composition of concatenations, respectively.

A morphism of \mathcal{S}_{II} -modules consists of natural transformations of functors compatible with (4.25), (4.27).

It would perhaps be more proper to speak of “ $(\mathcal{S}_{\text{I}}^{\pm}, \mathcal{S}_{\text{II}})$ -modules $(X_{\text{I}}^{\pm}, X_{\text{II}})$ ”, though such notation rapidly becomes unwieldy. It will usually be clear from context what X_{I}^{\pm} are once we have specified X_{II} , though at times we will specify the pair $(X_{\text{I}}^{\pm}, X_{\text{II}})$ for sake of clarity.

Examples 4.6–4.9 all generalize: $(\overline{\mathcal{M}}_{\text{I}}^{\pm}, \overline{\mathcal{M}}_{\text{II}})$ is an $(\mathcal{S}_{\text{I}}^{\pm}, \mathcal{S}_{\text{II}})$ -module, etc.

Definition 4.11. An \mathcal{S}_{III} -module X_{III} valued in \mathcal{C}^{\otimes} consists of the following data:

- (i) An $\mathcal{S}_{\text{I}}^{+}$ -module X_{I}^{+} valued in \mathcal{C}^{\otimes} .
- (ii) An $\mathcal{S}_{\text{I}}^{-}$ -module X_{I}^{-} valued in \mathcal{C}^{\otimes} .
- (iii) An $(\mathcal{S}_{\text{I}}^{\pm}, \mathcal{S}_{\text{II}}^{t=0})$ -module $(X_{\text{I}}^{\pm}, X_{\text{II}}^{t=0})$ valued in \mathcal{C}^{\otimes} .
- (iv) An $(\mathcal{S}_{\text{I}}^{\pm}, \mathcal{S}_{\text{II}}^{t=1})$ -module $(X_{\text{I}}^{\pm}, X_{\text{II}}^{t=1})$ valued in \mathcal{C}^{\otimes} .
- (v) A functor $X_{\text{III}} : \mathcal{S}_{\text{III}} \rightarrow \mathcal{C}$.
- (vi) For every concatenation $\{T_i\}_i$ in \mathcal{S}_{III} , a morphism

$$\bigotimes_i X_{\text{I/II/III}}(T_i) \rightarrow X_{\text{III}}(\#T_i), \quad (4.29)$$

satisfying the natural compatibility conditions, as in Definition 4.10.

A morphism of \mathcal{S}_{III} -modules consists of natural transformations of functors compatible with (4.25), (4.27), (4.29).

Definition 4.12. An \mathcal{S}_{IV} -module X_{IV} valued in \mathcal{C}^{\otimes} consists of the following data:

- (i) \mathcal{S}_{I}^i -modules X_{I}^i valued in \mathcal{C}^{\otimes} for $0 \leq i \leq 2$.
- (ii) $(\mathcal{S}_{\text{I}}^{i,j}, \mathcal{S}_{\text{II}}^{ij})$ -modules $(X_{\text{I}}^{i,j}, X_{\text{II}}^{ij})$ valued in \mathcal{C}^{\otimes} for $0 \leq i < j \leq 2$.
- (iii) A functor $X_{\text{IV}} : \mathcal{S}_{\text{IV}} \rightarrow \mathcal{C}$.
- (iv) For every concatenation $\{T_i\}_i$ in \mathcal{S}_{IV} , a morphism

$$\bigotimes_i X_{\text{I/II/IV}}(T_i) \rightarrow X_{\text{IV}}(\#_i T_i), \quad (4.30)$$

satisfying the natural compatibility conditions, as in Definition 4.10.

A morphism of \mathcal{S}_{IV} -modules consists of natural transformations of functors compatible with (4.25), (4.27), (4.30).

4.4 Sketch of the construction (revisited)

We now revisit the sketch from §4.2 using the language of \mathcal{S} -modules.

Recall that the virtual fundamental cycles of the moduli spaces $\overline{\mathcal{M}}(T)$ come packaged via the maps (4.20). Now the key coherence properties of these cycles (4.23)–(4.24) shall be encoded in the fact that (4.20) is a map of \mathcal{S} -modules. Thus our first task is to define the \mathcal{S} -modules $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$ and $C_{\bullet}(E)$ along with the pushforward map

$$C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial) \rightarrow C_{\bullet}(E). \quad (4.31)$$

We then must relate $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$ to $\check{C}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial; \mathfrak{o})$ (with an appropriate \mathcal{S} -module structure), and we must relate $C_{\bullet}(E)$ to \mathbb{Q} (with the trivial \mathcal{S} -module structure). If all goes well, this will give virtual fundamental cycles (4.22) satisfying (4.23)–(4.24).

In reality, it is somewhat cumbersome to make sense out of $\check{C}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial; \mathfrak{o})$ as an \mathcal{S} -module, so we will take a shortcut by taking advantage of the fact that we are only interested in integrating the constant function 1 (and not more general cohomology classes) over the virtual fundamental cycles. We will define an \mathcal{S} -module $\mathbb{Q}[\mathcal{S}]$, which can be thought of as “the subcomplex of $\check{C}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial; \mathfrak{o})$ generated by the Poincaré duals of the closed strata”, and we will construct a corresponding map of \mathcal{S} -modules

$$\mathbb{Q}[\mathcal{S}] \rightarrow C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial). \quad (4.32)$$

Now combining this map with (4.31) and our understanding of $C_{\bullet}(E) \cong \mathbb{Q}$, we obtain a map of \mathcal{S} -modules

$$\mathbb{Q}[\mathcal{S}] \rightarrow \mathbb{Q}. \quad (4.33)$$

Morally speaking, this map is the intersection pairing between the closed strata of $\overline{\mathcal{M}}(T)$ and its virtual fundamental cycle. From any map (4.33), it is straightforward to read off virtual moduli counts satisfying the master equations.

4.5 \mathcal{S} -modules \mathbb{Q} and $\mathbb{Q}[\mathcal{S}]$

We now introduce the two basic \mathcal{S} -modules \mathbb{Q} and $\mathbb{Q}[\mathcal{S}]$. We also make the elementary but crucial observation that a map of \mathcal{S} -modules $\mathbb{Q}[\mathcal{S}] \rightarrow \mathbb{Q}$ is nothing other than a collection of virtual moduli counts satisfying the relevant master equations.

Definition 4.13 (\mathcal{S} -module \mathbb{Q}). Denote by \mathbb{Q} the \mathcal{S} -module defined by $\mathbb{Q}(T) = \mathbb{Q}$ for all $T \in \mathcal{S}$, where the pushforward maps are the identity and the concatenation maps are multiplication.

Before defining $\mathbb{Q}[\mathcal{S}]$, we motivate it as follows. It is not hard to check that $\check{C}^{\bullet}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \mathfrak{o}_T)$ is quasi-isomorphic to the total complex

$$\left[\check{C}^{\bullet}(\overline{\mathcal{M}}(T); \mathfrak{o}_T) \longrightarrow \bigoplus_{\text{codim}(T'/T)=1} \check{C}^{\bullet-1}(\overline{\mathcal{M}}(T')/\text{Aut}(T'/T); \mathfrak{o}_{T'}) \longrightarrow \bigoplus_{\text{codim}(T''/T)=2} \check{C}^{\bullet-2}(\overline{\mathcal{M}}(T'')/\text{Aut}(T''/T); \mathfrak{o}_{T''}) \longrightarrow \dots \right]. \quad (4.34)$$

This chain model is convenient because it is manifestly an \mathcal{S} -module: a morphism $T' \rightarrow T$ clearly induces a map of complexes (4.34) of degree $\text{codim}(T'/T)$, while it is not so clear how to canonically define a corresponding map $\check{C}^{\bullet}(\overline{\mathcal{M}}(T') \text{ rel } \partial; \mathfrak{o}_{T'}) \rightarrow \check{C}^{\bullet+\text{codim}(T'/T)}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \mathfrak{o}_T)$. The complex $\mathbb{Q}[\mathcal{S}](T)$ which we now define should be thought of as the subcomplex of (4.34) spanned by the images of the constant sections

$$\mathfrak{o}_{T'} \rightarrow \check{C}^0(\overline{\mathcal{M}}(T'); \mathfrak{o}_{T'}) \rightarrow \check{C}^0(\overline{\mathcal{M}}(T')/\text{Aut}(T'/T); \mathfrak{o}_{T'}) \subseteq \check{C}^{\text{codim}(T'/T)}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \mathfrak{o}_T) \quad (4.35)$$

for $T' \rightarrow T$.

Definition 4.14 (\mathcal{S} -module $\mathbb{Q}[\mathcal{S}]$). For $T \in \mathcal{S}$, we define

$$\mathbb{Q}[\mathcal{S}](T) := \mathbb{Q}[\mathcal{S}/_T] = \bigoplus_{T' \rightarrow T} \mathfrak{o}_{T'}[\text{vdim}(T')] \quad (4.36)$$

(note that this definition makes sense since automorphisms of T' over T act trivially on $\mathfrak{o}_{T'}$; though the action is trivial, it may be conceptually helpful to think of $\mathfrak{o}_{T'}$ instead as its $\text{Aut}(T'/T)$ -coinvariants). We equip $\mathbb{Q}[\mathcal{S}](T)$ with the differential given by the sum over all codimension one maps $T'' \rightarrow T'$ (maps in $\mathcal{S}/_T$) of the boundary map $\mathfrak{o}_{T'} \rightarrow \mathfrak{o}_{T''}$ induced by (2.61) and pairing on the left with a chosen orientation of \mathbb{R} .

A morphism $T' \rightarrow T$ induces a map

$$\mathbb{Q}[\mathcal{S}](T') \rightarrow \mathbb{Q}[\mathcal{S}](T) \quad (4.37)$$

via pushforward under the functor $\mathcal{S}/_{T'} \rightarrow \mathcal{S}/_T$.

A concatenation $\{T_i\}_i$ induces a map

$$\bigotimes_i \mathbb{Q}[\mathcal{S}](T_i) \rightarrow \mathbb{Q}[\mathcal{S}](\#_i T_i) \quad (4.38)$$

by consideration of the isomorphism $\mathcal{S}/_{\#_i T_i} = \prod_i \mathcal{S}/_{T_i}$, covered by the tautological identifications $\mathfrak{o}_{\#_i T'_i} = \bigotimes_i \mathfrak{o}_{T'_i}$ multiplied by the factor $|\text{Aut}(\{T'_i\}_i/\#_i T'_i)| = |\text{Aut}(\{T_i\}_i/\#_i T_i)|$ (this factor is explained by it being the “degree” of the map $\prod_i \overline{\mathcal{M}}(T'_i) \rightarrow \overline{\mathcal{M}}(\#_i T'_i)$).

Thus $\mathbb{Q}[\mathcal{S}]$ is an \mathcal{S} -module (meaning $\mathbb{Q}[\mathcal{S}_I]$ is an \mathcal{S}_I -module, $(\mathbb{Q}[\mathcal{S}_I^\pm], \mathbb{Q}[\mathcal{S}_{II}])$ is an $(\mathcal{S}_I^\pm, \mathcal{S}_{II})$ -module, etc.).

Remark 4.15. It is easy to check that there is a natural bijection between morphisms of \mathcal{S} -modules $\mathbb{Q}[\mathcal{S}] \rightarrow \mathbb{Q}$ and collections of virtual moduli counts $\#\overline{\mathcal{M}}(T)^{\text{vir}} \in (\mathfrak{o}_T^\vee)^{\text{Aut}(T)}$ for $\text{vdim}(T) = 0$ satisfying

$$0 = \sum_{\text{codim}(T'/T)=1} \frac{1}{|\text{Aut}(T'/T)|} \#\overline{\mathcal{M}}(T')^{\text{vir}} \quad (4.39)$$

$$\#\overline{\mathcal{M}}(\#_i T_i)^{\text{vir}} = \frac{1}{|\text{Aut}(\{T_i\}_i/\#_i T_i)|} \prod_i \#\overline{\mathcal{M}}(T_i)^{\text{vir}} \quad (4.40)$$

(where $\overline{\mathcal{M}}(T)$ is interpreted as zero if $\text{vdim}(T) \neq 0$). Note that $\text{Aut}(T)$ -invariance forces $\#\overline{\mathcal{M}}(T)^{\text{vir}}$ to vanish if any of the input/output edges of T is labeled with a bad Reeb orbit.

4.6 \mathcal{S} -modules $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$ and $C_\bullet(E)$

We now define the two \mathcal{S} -modules $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$ and $C_\bullet(E)$.

Defining $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$ as an \mathcal{S} -module is nontrivial for the following reason. We would like to associate to $T \in \mathcal{S}$ the complex $C_{\text{vir}}^{\bullet+\text{vdim}(T)}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \bar{A}(T))$. However, a map $T' \rightarrow T$ does not induce a map on such complexes as desired, rather only a diagram

$$\begin{array}{ccc} C_{\text{vir}}^{\bullet+\text{vdim}(T')}(\overline{\mathcal{M}}(T') \text{ rel } \partial; \bar{A}(T)) & & \\ \downarrow \sim & \searrow & \\ C_{\text{vir}}^{\bullet+\text{vdim}(T')}(\overline{\mathcal{M}}(T') \text{ rel } \partial; \bar{A}(T')) & & C_{\text{vir}}^{\bullet+\text{vdim}(T)}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \bar{A}(T)). \end{array} \quad (4.41)$$

Fortunately, this diagram also suggests a solution. Namely, we instead associate to $T \in \mathcal{S}$ a suitable homotopy colimit of $C_{\text{vir}}^{\text{vdim}(T')+\bullet}(\overline{\mathcal{M}}(T') \text{ rel } \partial, \bar{A}(T'))$ over $T' \in \mathcal{S}_{/T}$.

Definition 4.16 (Homotopy diagram). Let \mathcal{T} be a finite category (meaning $\#|\mathcal{T}| < \infty$ and $\# \text{Hom}(\mathfrak{t}_1, \mathfrak{t}_2) < \infty$ for $\mathfrak{t}_1, \mathfrak{t}_2 \in \mathcal{T}$). A *homotopy diagram* over \mathcal{T} shall mean the following:

- (i) Let Δ^p denote the *simplex category* (so that a functor $\Delta^p \rightarrow \mathcal{T}$ is a chain of morphisms $\mathfrak{t}_0 \rightarrow \cdots \rightarrow \mathfrak{t}_p$). For every functor $\sigma : \Delta^p \rightarrow \mathcal{T}$, we specify a complex $A^\bullet(\sigma)$, and for every map of simplices $r : \Delta^q \rightarrow \Delta^p$, we specify a map $A^\bullet(\sigma) \rightarrow A^\bullet(\sigma \circ r)$. In other words, A^\bullet is a functor from the category whose objects are morphisms $\Delta^p \rightarrow \mathcal{T}$ and whose morphisms $(\Delta^p \rightarrow \mathcal{T}) \rightarrow (\Delta^q \rightarrow \mathcal{T})$ are factorizations $\Delta^q \rightarrow \Delta^p \rightarrow \mathcal{T}$.
- (ii) We require that if $r : \Delta^q \rightarrow \Delta^p$ satisfies $r(0) = 0$, then the induced map $A^\bullet(\sigma) \rightarrow A^\bullet(\sigma \circ r)$ is an isomorphism.

For example, an ordinary diagram B^\bullet over \mathcal{T} gives rise to a homotopy diagram by setting $A^\bullet(\sigma) := B^\bullet(\sigma(0))$.

Definition 4.17 (Homotopy colimit). Let \mathcal{T} be a finite category, and let A^\bullet be a homotopy diagram over \mathcal{T} . We define

$$\text{hocolim}_{\mathcal{T}} A^\bullet := \bigoplus_{p \geq 0} \bigoplus_{\mathfrak{t}_0 \rightarrow \cdots \rightarrow \mathfrak{t}_p} A^\bullet(\mathfrak{t}_0 \rightarrow \cdots \rightarrow \mathfrak{t}_p)_{\text{Aut}(\mathfrak{t}_0 \rightarrow \cdots \rightarrow \mathfrak{t}_p)} \otimes \mathfrak{o}_{\Delta^p}, \quad (4.42)$$

equipped with the differential which arises upon regarding the right hand side as chains on the nerve of \mathcal{T} (the subscript indicates taking coinvariants with respect to the natural action of $\text{Aut}(\mathfrak{t}_0 \rightarrow \cdots \rightarrow \mathfrak{t}_p)$).

Given a functor $f : \mathcal{T} \rightarrow \mathcal{T}'$ and a homotopy diagram A^\bullet over \mathcal{T}' , there is a natural pullback diagram $f^* A^\bullet$ over \mathcal{T} and a natural map

$$\text{hocolim}_{\mathcal{T}} f^* A^\bullet \rightarrow \text{hocolim}_{\mathcal{T}'} A^\bullet. \quad (4.43)$$

Given a finite collection A_i^\bullet of homotopy diagrams over \mathcal{T}_i , there is an Eilenberg–Zilber quasi-isomorphism

$$\bigotimes_i \text{hocolim}_{\mathcal{T}_i} A_i^\bullet \xrightarrow{\sim} \text{hocolim}_{\prod_i \mathcal{T}_i} \bigotimes_i A_i^\bullet, \quad (4.44)$$

corresponding to the standard simplicial subdivision of $\Delta^p \times \Delta^q$ into $\binom{p+q}{p}$ copies of Δ^{p+q} .

Lemma 4.18. *If \mathcal{T} has a final object \mathfrak{t}^{\max} , then the natural map*

$$A^\bullet(\mathfrak{t}^{\max}) \xrightarrow{\sim} \text{hocolim}_{\mathcal{T}} A^\bullet \quad (4.45)$$

is a quasi-isomorphism.

Proof. Filter the homotopy colimit by the number of \mathfrak{t}_i which are not isomorphic to \mathfrak{t}^{\max} . \square

Definition 4.19 (\mathcal{S} -module $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$). For $T \in \mathcal{S}$, we define

$$C_{\text{vir}}^\bullet(\overline{\mathcal{M}} \text{ rel } \partial)(T) := \text{hocolim}_{T_0 \rightarrow \cdots \rightarrow T_p \rightarrow T} C_{\text{vir}}^{\bullet-\text{codim}(T_0/T)}(\overline{\mathcal{M}}(T_0) \text{ rel } \partial, \bar{A}(T_p)). \quad (4.46)$$

This homotopy colimit is over the category $\mathcal{S}_{/T}$, and the structure maps of the homotopy diagram are those appearing in (4.41) (note that Definition 4.16(ii) is satisfied as the leftmost map in (4.41) is a quasi-isomorphism). The natural inclusion map

$$C_{\text{vir}}^{\bullet}(\overline{\mathcal{M}}(T) \text{ rel } \partial; \bar{A}(T)) \hookrightarrow C_{\text{vir}}^{\bullet}(\overline{\mathcal{M}} \text{ rel } \partial)(T) \quad (4.47)$$

is a quasi-isomorphism by Lemma 4.18.

Now $C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)$ naturally has the structure of an \mathcal{S} -module, as follows. A morphism $T \rightarrow T'$ induces a natural pushforward map (4.43) induced by the functor $\mathcal{S}_{/T} \rightarrow \mathcal{S}_{/T'}$, which is covered by a natural isomorphism of homotopy diagrams. Given a concatenation $\{T_i\}_i$, there is a natural Eilenberg–Zilber map (4.44) induced by the isomorphism $\mathcal{S}_{/\#_i T_i} = \prod_i \mathcal{S}_{/T_i}$, which is covered by a morphism of homotopy diagrams coming from the product maps (4.16).

Definition 4.20 (\mathcal{S} -module $C_{\bullet}(E)$). For $T \in \mathcal{S}$, we define

$$C_{\bullet}(E)(T) := \text{hocolim}_{T_0 \rightarrow \dots \rightarrow T_p \rightarrow T} C_{\bullet}(E; \bar{A}(T_p)). \quad (4.48)$$

Again by Lemma 4.18, we have

$$H_{\bullet}(E)(T) = H_{\bullet}(E; \bar{A}(T)) = \mathbb{Q}. \quad (4.49)$$

As in Definition 4.19, $C_{\bullet}(E)$ naturally has the structure of a \mathcal{S} -module. The isomorphism above is in fact an isomorphism of \mathcal{S} -modules $H_{\bullet}(E) = \mathbb{Q}$.

There is a canonical map of \mathcal{S} -modules

$$C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial) \rightarrow C_{-\bullet}(E) \quad (4.50)$$

induced by (4.3).

Definition 4.21 (\mathcal{S}^{op} -module $\text{Hom}_{\mathcal{S}_{/T}}(\mathbb{Q}[\mathcal{S}], C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial))$). For $T \in \mathcal{S}$, we consider

$$\text{Hom}_{\mathcal{S}_{/T}}(\mathbb{Q}[\mathcal{S}], C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)). \quad (4.51)$$

Namely, an element of this group is a natural transformation of functors from $\mathcal{S}_{/T}$ to the category of *graded \mathbb{Q} -vector spaces*, and we equip it with the usual differential $d \circ f - (-1)^{|f|} f \circ d$.

Now (4.51) is an \mathcal{S}^{op} -module as follows. A map $T \rightarrow T'$ clearly gives a restriction map from homomorphisms over $\mathcal{S}_{/T'}$ to homomorphisms over $\mathcal{S}_{/T}$. A concatenation $\{T_i\}_i$ induces the necessary map by virtue of the fact that $\mathcal{S}_{/\#_i T_i} = \prod_i \mathcal{S}_{/T_i}$ and the concatenation maps for $\mathbb{Q}[\mathcal{S}]$ are isomorphisms.

Definition 4.22 (\mathcal{S}^{op} -module $\check{H}^{\bullet}(\overline{\mathcal{M}}(T))$). We equip the functor $T \mapsto \check{H}^{\bullet}(\overline{\mathcal{M}}(T))$ with the structure of an \mathcal{S}^{op} -module as follows. A map $T \rightarrow T'$ clearly gives a pullback map $\check{H}^{\bullet}(\overline{\mathcal{M}}(T')) \rightarrow \check{H}^{\bullet}(\overline{\mathcal{M}}(T))$. A concatenation $\{T_i\}_i$ induces a map

$$\bigotimes_i \check{H}^{\bullet}(\overline{\mathcal{M}}(T_i)) \rightarrow \check{H}^{\bullet}\left(\prod_i \overline{\mathcal{M}}(T_i)\right) \rightarrow \check{H}^{\bullet}\left(\prod_i \overline{\mathcal{M}}(T_i)\right)^{\text{Aut}(\{T_i\}_i/\#_i T_i)} \xleftarrow{\sim} \check{H}^{\bullet}(\overline{\mathcal{M}}(\#_i T_i)), \quad (4.52)$$

where the middle map is averaging over the group action, and on the right we use (2.4) together with the general result that the pullback map $\check{H}^\bullet(X/G) \xrightarrow{\sim} \check{H}^\bullet(X)^G$ is an isomorphism for any compact Hausdorff space X with an action by a finite group G [Par16, Lemma A.4.9]. Commutativity of the first diagram in (4.26) follows from the fact that a morphism of concatenations $\{T_i\}_i \rightarrow \{T'_i\}_i$ induces a map of diagrams (4.52), noting that $\text{Aut}(\{T_i\}_i/\#_i T_i) = \text{Aut}(\{T'_i\}_i/\#_i T'_i)$. Commutativity of the second diagram in (4.26) follows from a diagram chase.

Lemma 4.23. *The cohomology of the \mathcal{S}^{op} -module from Definition 4.21 is canonically isomorphic to the \mathcal{S}^{op} -module from Definition 4.22.*

Proof. First, observe that

$$\text{Hom}_{\mathcal{S}/T}(\mathbb{Q}[\mathcal{S}], C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)) = \prod_{T' \rightarrow T} [\mathfrak{o}_{T'}^\vee \otimes C_{\text{vir}}^\bullet(\overline{\mathcal{M}} \text{ rel } \partial)(T')]^{\text{Aut}(T'/T)}.$$

Now

$$\prod_{T' \rightarrow T} [\mathfrak{o}_{T'}^\vee \otimes C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(T') \text{ rel } \partial; \bar{A}(T))]^{\text{Aut}(T'/T)} \quad (4.53)$$

includes quasi-isomorphically into the right hand side above, as the composition of the final object quasi-isomorphism (4.47) and the atlas enlargement quasi-isomorphism (4.6). Now (4.53) is exactly the complex

$$C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(T), \mathcal{S}/T; \bar{A}(T)) \otimes \mathfrak{o}_T^\vee, \quad (4.54)$$

where the left factor is as defined in (4.12). Appealing to (4.13), we deduce the desired isomorphism

$$H^\bullet \text{Hom}_{\mathcal{S}/T}(\mathbb{Q}[\mathcal{S}], C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial)) = \check{H}^\bullet(\overline{\mathcal{M}}(T)). \quad (4.55)$$

It remains to compare pullback and concatenation maps.

Compatibility of (4.55) with pullback maps is shown as follows. The pullback map on $\text{Hom}_{\mathcal{S}/T}(\mathbb{Q}[\mathcal{S}], C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial))$ for $T' \rightarrow T$ becomes, under the quasi-isomorphisms above, the natural restriction map

$$C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(T), \mathcal{S}/T; \bar{A}(T)) \rightarrow C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(T'), \mathcal{S}/T'; \bar{A}(T)) \xrightarrow{\sim} C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(T'), \mathcal{S}/T'; \bar{A}(T')). \quad (4.56)$$

Since the isomorphism (4.10) is natural, the map $\check{H}^\bullet(\overline{\mathcal{M}}(T)) \rightarrow \check{H}^\bullet(\overline{\mathcal{M}}(T'))$ induced by (4.56) is simply the map induced by the corresponding map of H^0 sheaves, which by definition is just restriction, as needed.

To show compatibility of (4.55) with concatenation maps, again under the quasi-isomorphisms above, the concatenation maps for $\text{Hom}_{\mathcal{S}/T}(\mathbb{Q}[\mathcal{S}], C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial))$ become the natural maps

$$\begin{aligned} \bigotimes_i C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(T_i), \mathcal{S}/T_i; \bar{A}(T_i)) &\xrightarrow{(4.16)} C_{\text{vir}}^\bullet\left(\prod_i \overline{\mathcal{M}}(T_i), \prod_i \mathcal{S}/T_i; \prod_i \bar{A}(T_i)\right) \\ &\rightarrow C_{\text{vir}}^\bullet(\overline{\mathcal{M}}(\#_i T_i), \mathcal{S}/\#_i T_i, \bar{A}(\#_i T_i)). \end{aligned} \quad (4.57)$$

Now appealing to Lemma 4.3, it is enough to observe that the induced maps on H^0 sheaves are the natural ones. \square

4.7 Cofibrant \mathcal{S} -modules

We would like to upgrade the isomorphism of \mathcal{S} -modules $H_\bullet(E) = \mathbb{Q}$ to a quasi-isomorphism of \mathcal{S} -modules $C_\bullet(E) \xrightarrow{\sim} \mathbb{Q}$. Unfortunately, the natural strategy to construct such a quasi-isomorphism, namely by induction on T , fails. As a substitute, we introduce the notion of a *cofibrant* \mathcal{S} -module, and we construct a cofibrant \mathcal{S} -module $C_\bullet^{\text{cof}}(E)$ with a natural quasi-isomorphism $C_\bullet^{\text{cof}}(E) \xrightarrow{\sim} C_\bullet(E)$. Cofibrancy is exactly the condition which ensures that the natural inductive construction of a quasi-isomorphism $C_\bullet^{\text{cof}}(E) \xrightarrow{\sim} \mathbb{Q}$ succeeds. The resulting diagram

$$C_\bullet(E) \xleftarrow{\sim} C_\bullet^{\text{cof}}(E) \xrightarrow{\sim} \mathbb{Q} \quad (4.58)$$

turns out to be enough for our purposes. Cofibrancy of the \mathcal{S} -module $\mathbb{Q}[\mathcal{S}]$ will also end up being important for the inductive construction of virtual fundamental cycles.

Definition 4.24 (Cofibrant \mathcal{S} -module). An \mathcal{S} -module X valued in chain complexes shall be called *cofibrant* iff it satisfies the following two properties:

- (i) For all concatenations $\{T_i\}_i$ in \mathcal{S} , the induced map

$$\left[\bigotimes_i X(T_i) \right]_{\text{Aut}(\{T_i\}_i/\#T_i)} \xrightarrow{\sim} X(\#_i T_i) \quad (4.59)$$

is an isomorphism (the subscript indicates taking coinvariants with respect to the natural action of $\text{Aut}(\{T_i\}_i/\#T_i)$). Note that this follows if (4.25)/(4.27)/(4.29)/(4.30) is itself an isomorphism.

- (ii) For maximal $T \in \mathcal{S}$, the map

$$\text{colim}_{\text{codim}(T'/T) \geq 1} X(T') \rightarrow X(T) \quad (4.60)$$

is injective. The left side denotes the colimit over the full subcategory of \mathcal{S}/T spanned by objects $T' \rightarrow T$ with $\text{codim}(T'/T) \geq 1$.

Note that for $(\mathcal{S}_I^\pm, \mathcal{S}_{II})$ -modules, the above conditions are imposed over each of \mathcal{S}_I^\pm , \mathcal{S}_{II} , and so on.

Lemma 2.5 will be used frequently below.

Lemma 4.25. *Injectivity of (4.60) for maximal T implies injectivity for all T .*

Proof. Fix $T \in \mathcal{S}$, and write $T = \#_i T_i$ for maximal T_i . Now we have $\mathcal{S}/T = \prod_i \mathcal{S}/T_i$. Consider the cubical diagram

$$\bigotimes_i \left[\text{colim}_{\text{codim}(T'_i/T_i) \geq 1} X(T'_i) \rightarrow X(T_i) \right]. \quad (4.61)$$

Now (4.60) for T is precisely the map to the top vertex of the cube (4.61) from the colimit over its remaining vertices. This map is clearly injective given that each map in (4.61) is injective. \square

Lemma 4.26. *The \mathcal{S} -module $\mathbb{Q}[\mathcal{S}]$ is cofibrant.*

Proof. The concatenation maps are isomorphisms by definition since $\mathcal{S}_{/\#_i T_i} = \prod_i \mathcal{S}_{/T_i}$. Let us now show that

$$\operatorname{colim}_{\operatorname{codim}(T'/T) \geq 1} \mathbb{Q}[\mathcal{S}](T') \xrightarrow{\sim} \mathbb{Q}[\mathcal{S}](T) \quad (4.62)$$

is an isomorphism onto the subspace of $\mathbb{Q}[\mathcal{S}](T)$ generated by those $T' \rightarrow T$ of codimension ≥ 1 (certainly this is sufficient).

Both sides of (4.62) are graded by $|\mathcal{S}_{/T}|$, so it suffices to fix $T' \rightarrow T$ of codimension ≥ 1 and show that (4.62) is an isomorphism on T' -graded pieces. The map (4.62) is certainly surjective onto the T' -graded piece $\mathfrak{o}_{T'} \subseteq \mathbb{Q}[\mathcal{S}](T)$. Moreover, there is a section from $\mathfrak{o}_{T'}$ back to the left side of (4.62) via the inclusion $\mathfrak{o}_{T'} \subseteq \mathbb{Q}[\mathcal{S}](T')$. Now it is enough to argue that this section is surjective (onto the T' -graded piece), but this is clear since a contribution of $\mathbb{Q}[\mathcal{S}](T'')$ to the T' -graded piece is the same thing as a factorization $T' \rightarrow T'' \rightarrow T$. \square

Definition 4.27 (Partial order on $|\mathcal{S}|$). For $T, T' \in \mathcal{S}$, let us write $T' \preceq T$ iff there is a morphism $\#_i T_i \rightarrow T$ with some T_i isomorphic to T' .

Since we are restricting to effective objects of \mathcal{S} , compactness of $\overline{\mathcal{M}}$ implies that the relation \preceq is *well-founded* (i.e. there is no infinite strictly decreasing sequence $T_1 \succ T_2 \succ \dots$). The relation \preceq is also a partial order (reflexivity and transitivity are obvious, and antisymmetry follows from well-foundedness). It follows that induction on $|\mathcal{S}|$ partially ordered by \preceq is justified.

Definition 4.28 (Cofibrant \mathcal{S} -module $C_{\bullet}^{\operatorname{cof}}(E)$). We now define a cofibrant \mathcal{S} -module $C_{\bullet}^{\operatorname{cof}}(E)$ together with a quasi-isomorphism

$$C_{\bullet}^{\operatorname{cof}}(E) \xrightarrow{\sim} C_{\bullet}(E) \quad (4.63)$$

which is surjective for maximal T . Furthermore, the action of the “paths between basepoints” subgroup of $\operatorname{Aut}(T)$ on $C_{\bullet}^{\operatorname{cof}}(E)(T)$ will be trivial for all T (as it is for $C_{\bullet}(E)(T)$).

We construct $C_{\bullet}^{\operatorname{cof}}(E)(T)$ by induction on T , partially ordered as in Definition 4.27. For T non-maximal, write $T = \#_i T_i$ with T_i maximal. The definition of cofibrancy both forces us to take $C_{\bullet}^{\operatorname{cof}}(E)(T) := \bigotimes_i C_{\bullet}^{\operatorname{cof}}(E)(T_i)$ and assures that the \mathcal{S} -module structure maps with target $C_{\bullet}^{\operatorname{cof}}(E)(T)$ exist and are unique. For T maximal, consider the diagram

$$\begin{array}{ccc} \operatorname{colim}_{\operatorname{codim}(T'/T) \geq 1} C_{\bullet}^{\operatorname{cof}}(E)(T') & \xrightarrow{\sim} & C_{\bullet}^{\operatorname{cof}}(E)(T) \\ \downarrow (4.63) & & \sim \downarrow (4.63) \\ \operatorname{colim}_{\operatorname{codim}(T'/T) \geq 1} C_{\bullet}(E)(T') & \longrightarrow & C_{\bullet}(E)(T). \end{array} \quad (4.64)$$

We define $C_{\bullet}^{\operatorname{cof}}(E)(T)$ to be the mapping cylinder of the composition of the two solid maps, which clearly fits into the diagram as desired. Now the top horizontal map defines the \mathcal{S} -module structure maps with target $C_{\bullet}^{\operatorname{cof}}(E)(T)$, and the commutativity of the diagram ensures that (4.63) is a map of \mathcal{S} -modules.

Lemma 4.29. *Let $\mathcal{T} \subseteq \mathcal{S}_{/T_a}$ be a full subcategory which is “downward closed” in the sense that $T' \rightarrow T$ and $T \in \mathcal{T}$ implies $T' \in \mathcal{T}$. Then the natural map*

$$C_\bullet N_\bullet \mathcal{T} \cong \operatorname{hocolim}_{T \in \mathcal{T}} C_\bullet^{\operatorname{cof}}(E)(T) \xrightarrow{\sim} \operatorname{colim}_{T \in \mathcal{T}} C_\bullet^{\operatorname{cof}}(E)(T) \quad (4.65)$$

is a quasi-isomorphism. In particular, the homology of the right side is supported in degrees ≥ 0 and in degree zero is generated by the images of $H_0^{\operatorname{cof}}(E)(T)$ for $T \in \mathcal{T}$.

Proof. We argue by induction on $\#\mathcal{T}$ (note that $|\mathcal{S}_{/T}|$ is finite by our restriction to effective objects), the case $\#\mathcal{T} = 0$ being trivial.

Pick a maximal object $T_0 \in \mathcal{T}$. If $\mathcal{T} = \mathcal{T}^{\leq T_0}$, then the desired conclusion follows from Lemma 4.18; otherwise argue as follows. Let us abbreviate $X := C_\bullet^{\operatorname{cof}}(E)$. We claim that there is a short exact sequence

$$0 \rightarrow \operatorname{colim}_{T \in \mathcal{T}^{< T_0}} X(T) \rightarrow \operatorname{colim}_{T \in \mathcal{T}^{\leq T_0}} X(T) \oplus \operatorname{colim}_{T \in \mathcal{T} \setminus \{T_0\}} X(T) \rightarrow \operatorname{colim}_{T \in \mathcal{T}} X(T) \rightarrow 0. \quad (4.66)$$

(Note that the expressions $\mathcal{T}^{< T_0}$, $\mathcal{T}^{\leq T_0}$, $\mathcal{T} \setminus \{T_0\}$ refer to *full subcategories* of \mathcal{T} , not any sort of over-categories.) Exactness on the right and in the middle follows from the universal property of colim . To show exactness on the left, we show the stronger statement that

$$\operatorname{colim}_{T \in \mathcal{T}^{< T_0}} X(T) \hookrightarrow \operatorname{colim}_{T \in \mathcal{T}^{\leq T_0}} X(T) \quad (4.67)$$

is injective. To see this, note that cofibrancy of X (and Lemma 4.25) and exactness of $\operatorname{Aut}_{\mathcal{T}}(T_0)$ -coinvariants imply that the map

$$\left(\operatorname{colim}_{T \in \mathcal{T}/T_0} X(T) \right)_{\operatorname{Aut}_{\mathcal{T}}(T_0)} \hookrightarrow X(T_0)_{\operatorname{Aut}_{\mathcal{T}}(T_0)} \quad (4.68)$$

is injective. The right hand sides of (4.68) and (4.67) coincide, and the left hand side of (4.68) surjects onto the left hand side of (4.67), so we conclude the desired injectivity of (4.67).

Now (4.66) induces a long exact sequence on homology. There is a similar long exact sequence with $\operatorname{hocolim}$ in place of colim , together with a map from the sequence of $\operatorname{hocolim}$'s to the sequence of colim 's. Now the five lemma and the induction hypothesis (together with the fact that $\mathcal{T}^{\leq T_0} \subsetneq \mathcal{T}$) give the desired result. \square

Lemma 4.30. *There exists a quasi-isomorphism of \mathcal{S} -modules*

$$C_\bullet^{\operatorname{cof}}(E) \xrightarrow{\sim} \mathbb{Q} \quad (4.69)$$

inducing the canonical isomorphism $H_\bullet^{\operatorname{cof}}(E) = H_\bullet(E) = \mathbb{Q}$ from (4.49).

Proof. We argue by induction on $T \in \mathcal{S}$, partially ordered as in Definition 4.27.

For T non-maximal, write $T = \#_i T_i$ for T_i maximal. Then cofibrancy of $C_\bullet^{\operatorname{cof}}(E)$ both forces us to take $p_*(T) := \bigotimes_i p_*(T_i)$ and assures that this choice is compatible with the maps defined thus far.

For maximal T , we would like to fill in the diagram

$$\begin{array}{ccc}
\operatorname{colim}_{\operatorname{codim}(T'/T) \geq 1} C_{\bullet}^{\operatorname{cof}}(E)(T') & \xrightarrow{\quad} & C_{\bullet}^{\operatorname{cof}}(E)(T) \\
\downarrow p_* & \swarrow p_* & \\
\mathbb{Q} & &
\end{array} \tag{4.70}$$

with a map $C_{\bullet}^{\operatorname{cof}}(E)(T) \rightarrow \mathbb{Q}$ in a particular chain homotopy class. The horizontal map is injective since $C_{\bullet}^{\operatorname{cof}}(E)$ is cofibrant; it follows that it is enough to show that (4.70) commutes up to chain homotopy. For chain complexes C_{\bullet} over \mathbb{Q} , we have an isomorphism $H_0 \operatorname{Hom}(C_{\bullet}, \mathbb{Q}) \xrightarrow{\sim} \operatorname{Hom}(H_0 C_{\bullet}, \mathbb{Q})$, so it suffices to show that (4.70) commutes on homology, which follows from Lemma 4.29. Finally, we may ensure p_* is $\operatorname{Aut}(T)$ -invariant by averaging (this is necessary for p_* to be a natural transformation of functors). \square

Lemma 4.31. *There exists a map of \mathcal{S} -modules*

$$\mathbb{Q}[\mathcal{S}] \rightarrow C_{\operatorname{vir}}^{\bullet+\operatorname{vdim}}(\overline{\mathcal{M}} \operatorname{rel} \partial) \tag{4.71}$$

with the following property. Such a map determines for all $T \in \mathcal{S}$ a cycle in (4.51), which by Lemma 4.23 determines an element of $\check{H}^0(\overline{\mathcal{M}}(T))$. We require these elements coincide with the class of the constant function $1 \in \check{H}^0(\overline{\mathcal{M}}(T))$ for all $T \in \mathcal{S}$.

Proof. We argue by induction on T .

For T non-maximal, write $T = \#_i T_i$ with T_i maximal; now cofibrancy of $\mathbb{Q}[\mathcal{S}]$ (Lemma 4.26) determines the map completely on $\mathbb{Q}[\mathcal{S}](T)$, and the map has the desired property by Lemma 4.23.

For maximal T , consider the following restriction map:

$$\prod_{T' \rightarrow T} [\mathfrak{o}_{T'}^{\vee} \otimes C_{\operatorname{vir}}^{\bullet}(\overline{\mathcal{M}} \operatorname{rel} \partial)(T')]^{\operatorname{Aut}(T'/T)} \rightarrow \prod_{\operatorname{codim}(T'/T) \geq 1} [\mathfrak{o}_{T'}^{\vee} \otimes C_{\operatorname{vir}}^{\bullet}(\overline{\mathcal{M}} \operatorname{rel} \partial)(T')]^{\operatorname{Aut}(T'/T)}. \tag{4.72}$$

The part of the map (4.71) defined thus far determines a cycle on the right, which we must lift to a cycle on the left in a particular cohomology class. Since (4.72) is surjective, it suffices to show that the desired cohomology class on the left is mapped to the given cohomology class on the right. Identifying (4.72) with the map on global sections of a corresponding map of pure homotopy sheaves as in Lemma 4.23, we see that on cohomology this map is simply the restriction map $\check{H}^{\bullet}(\overline{\mathcal{M}}(T)) \rightarrow \check{H}^{\bullet}(\partial \overline{\mathcal{M}}(T))$. It thus suffices to show that the cycle on the right equals the constant function $1 \in \check{H}^0(\partial \overline{\mathcal{M}}(T))$, which can be checked on each $\overline{\mathcal{M}}(T')$ for $\operatorname{codim}(T'/T) \geq 1$, where it holds by the induction hypothesis. \square

4.8 Sets Θ

We now conclude by defining the sets Θ which index all possible choices of the extra data necessary to fix coherent virtual fundamental cycles on the moduli spaces $\overline{\mathcal{M}}$.

Definition 4.32 (Sets $\Theta(\mathcal{D})$). Given a datum \mathcal{D} as in any of Setups I–IV, an element $\theta \in \Theta(\mathcal{D})$ consists of a commuting diagram⁷ of \mathcal{S} -modules

$$\begin{array}{ccccc} \mathbb{Q}[\mathcal{S}] & \xrightarrow{\tilde{w}_*} & C_{-\bullet}^{\text{cof}}(E) & \xrightarrow{p_*} & \mathbb{Q} \\ \downarrow w_* & & \sim \downarrow (4.63) & & \\ C_{\text{vir}}^{\bullet+\text{vdim}}(\overline{\mathcal{M}} \text{ rel } \partial) & \xrightarrow{(4.50)} & C_{-\bullet}(E) & & \end{array} \quad (4.73)$$

satisfying the following properties:

- We require that p_* induce the canonical isomorphism $H_{\bullet}^{\text{cof}}(E) = H_{\bullet}(E) = \mathbb{Q}$ from Definition 4.20.
- We require that w_* satisfy the conclusion of Lemma 4.31.

Note that there are natural forgetful maps

$$\Theta_{\text{II}} \rightarrow \Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-, \quad (4.74)$$

$$\Theta_{\text{III}} \rightarrow \Theta_{\text{II}}^{t=0} \times_{\Theta_{\text{I}}^+ \times \Theta_{\text{I}}^-} \Theta_{\text{II}}^{t=1}, \quad (4.75)$$

$$\Theta_{\text{IV}} \rightarrow \Theta_{\text{II}}^{02} \times_{\Theta_{\text{I}}^0 \times \Theta_{\text{I}}^2} (\Theta_{\text{II}}^{01} \times_{\Theta_{\text{I}}^1} \Theta_{\text{II}}^{12}). \quad (4.76)$$

An element $\theta \in \Theta$ evidently gives rise to a morphism of \mathcal{S} -modules $p_* \circ \tilde{w}_* : \mathbb{Q}[\mathcal{S}] \rightarrow \mathbb{Q}$. Such a morphism corresponds to virtual moduli counts $\#\overline{\mathcal{M}}(T)_{\theta}^{\text{vir}}$ satisfying the relevant master equation by Remark 4.15.

Proposition 4.33. *If $\overline{\mathcal{M}}(T)$ is weakly regular and $\text{vdim}(T) = 0$, then $\overline{\mathcal{M}}(T) = \mathcal{M}(T)$ and $\#\overline{\mathcal{M}}(T)_{\theta}^{\text{vir}} = \#\mathcal{M}(T)$.*

Proof. There are no nontrivial (effective) $T' \rightarrow T$ for dimension reasons, so $\overline{\mathcal{M}}(T) = \mathcal{M}(T)$. Evaluating the diagram (4.73) at T and taking cohomology gives

$$\begin{array}{ccccc} \mathfrak{o}_T & \xrightarrow{\tilde{w}_*} & \mathbb{Q} & \xrightarrow{\text{id}} & \mathbb{Q} \\ \downarrow w_* & & \downarrow \text{id} & & \\ \check{H}^{\bullet}(\overline{\mathcal{M}}(T); \mathfrak{o}_T) & \xrightarrow{[\overline{\mathcal{M}}(T)]^{\text{vir}}} & \mathbb{Q} & & \end{array} \quad (4.77)$$

Unravelling the isomorphism in Lemma 4.31, we see that the left vertical map is just the tautological map to \check{H}^0 . The bottom horizontal map is by definition the virtual fundamental class $[\overline{\mathcal{M}}(T)]^{\text{vir}}$ from (4.3). Thus we would like to appeal to Lemma 4.2 to conclude that the virtual fundamental class coincides with the ordinary fundamental class.

For cases (I) and (II), our regularity assumption on $\overline{\mathcal{M}}(T)$ ensures that the hypothesis $\overline{\mathcal{M}}(T) = \overline{\mathcal{M}}(T)^{\text{reg}}$ of Lemma 4.2 is satisfied (one should observe here that for moduli spaces of virtual dimension zero, the regular locus automatically has trivial isotropy, due to our use of asymptotic markers and the fact that nodes appear only in codimension two). For cases (III) and (IV), there is a gap between our assumption that $\overline{\mathcal{M}}(T)$ is weakly regular

⁷Note that the categories \mathcal{S} are essentially small, so the collection of such diagrams forms a set.

and the hypothesis $\overline{\mathcal{M}}(T) = \overline{\mathcal{M}}(T)^{\text{reg}}$ of Lemma 4.2 (see Definition 2.39). It suffices, though, to observe that the implicit atlas on $\overline{\mathcal{M}}(T)$ remains an implicit atlas if we enlarge $\overline{\mathcal{M}}(T)_I^{\text{reg}}$ to include all points which are weakly regular and have trivial isotropy. The only thing to check is that the manifold/openness/submersion axioms still hold, but this is easy since the thickened moduli spaces have only the top stratum $\mathcal{M}(T)_I$ and hence can be described by the usual Banach manifold Fredholm setup. \square

Proposition 4.34. *The set Θ_I is non-empty (resp. (4.74)–(4.76) are surjective).*

Proof. Concretely, we must construct p_* , w_* , \tilde{w}_* as in Definition 4.32. We will in fact give a construction of such data by *induction* on $T \in \mathcal{S}$ with respect to the partial order from Definition 4.27. The inductive nature of our proof (in addition to being the only reasonable approach) implies surjectivity of (4.74)–(4.76) (as opposed to mere non-emptiness of Θ_{II} , Θ_{III} , Θ_{IV}).

The existence of p_* and w_* follow from Lemmas 4.30 and 4.31 respectively. Note that both are proved by induction on T , so their use here is permissible.

To show the existence of \tilde{w}_* , we argue by induction on T . Note that $\mathbb{Q}[\mathcal{S}]$ is cofibrant by Lemma 4.26. For non-maximal T , cofibrancy of $\mathbb{Q}[\mathcal{S}]$ determines $\tilde{w}_*(T)$ uniquely and implies that the diagram still commutes. For maximal T , we are faced with the following lifting problem:

$$\begin{array}{ccc}
 \text{colim}_{\text{codim}(T'/T) \geq 1} \mathbb{Q}[\mathcal{S}](T') & \xrightarrow{\tilde{w}_*} & C_{\bullet}^{\text{cof}}(E)(T) \\
 \downarrow & \nearrow \tilde{w}_* & \downarrow \sim (4.63) \\
 \mathbb{Q}[\mathcal{S}](T) & \xrightarrow{(4.50) \circ w_*} & C_{\bullet}(E)(T).
 \end{array} \tag{4.78}$$

The left vertical map is injective by Lemma 4.26, and the right vertical map is a surjective quasi-isomorphism by Definition 4.28. All four complexes are bounded below. Now Lemma 4.35 ensures that a lift exists, and we make it $\text{Aut}(T)$ -equivariant by averaging. \square

Lemma 4.35. *Consider a diagram of chain complexes bounded below over a ring R :*

$$\begin{array}{ccc}
 A_{\bullet} & \longrightarrow & X_{\bullet} \\
 \downarrow & \nearrow & \downarrow \sim \\
 B_{\bullet} & \longrightarrow & Y_{\bullet}
 \end{array} \tag{4.79}$$

where the right vertical map is a surjective quasi-isomorphism and the left vertical map is an injection whose cokernel is degreewise projective. Then there exists a lift as illustrated.

Proof. This is the fact that “cofibrations have the left lifting property with respect to acyclic fibrations” in the projective model structure on $\mathbf{Ch}_{\geq 0}(R)$. It can be proved by a straightforward diagram chase. \square

4.9 Symmetric monoidal structure on Θ

We now show that the sets Θ in cases (I) and (II) are naturally symmetric monoidal with respect to disjoint union of data \mathcal{D} . Concretely, this means we are relating the virtual

moduli counts associated to a disjoint union of contact manifolds or symplectic cobordisms with those associated to each piece (compare §1.7). In fact, to do this we will consider not Θ as defined in Definition 4.32 but rather an (equivalent) enlargement thereof (for which all of the results proven earlier remain valid), explained during the course of the proof.

Proposition 4.36. *There are functorial maps associated to finite disjoint unions (indexed by $i \in I$)*

$$\Theta_I(Y, \lambda, J) \leftarrow \prod_{i \in I} \Theta_I(Y_i, \lambda_i, J_i) \quad \text{for } (Y, \lambda, J) = \bigsqcup_{i \in I} (Y_i, \lambda_i, J_i) \quad (4.80)$$

$$\Theta_{II}(\hat{X}, \hat{\omega}, \hat{J}) \leftarrow \prod_{i \in I} \Theta_{II}(\hat{X}_i, \hat{\omega}_i, \hat{J}_i) \quad \text{for } (\hat{X}, \hat{\omega}, \hat{J}) = \bigsqcup_{i \in I} (\hat{X}_i, \hat{\omega}_i, \hat{J}_i) \quad (4.81)$$

preserving the virtual moduli counts and turning Θ_I and Θ_{II} into weak symmetric monoidal functors.

Proof. Note that (restricting to effective objects), we have

$$\mathcal{S}_I(Y, \lambda, J) = \bigsqcup_{i \in I} \mathcal{S}_I(Y_i, \lambda_i, J_i) \quad \text{for } (Y, \lambda, J) = \bigsqcup_{i \in I} (Y_i, \lambda_i, J_i), \quad (4.82)$$

$$\mathcal{S}_{II}(\hat{X}, \hat{\omega}, \hat{J}) = \bigsqcup_{i \in I} \mathcal{S}_{II}(\hat{X}_i, \hat{\omega}_i, \hat{J}_i) \quad \text{for } (\hat{X}, \hat{\omega}, \hat{J}) = \bigsqcup_{i \in I} (\hat{X}_i, \hat{\omega}_i, \hat{J}_i), \quad (4.83)$$

so an $\mathcal{S}_I(Y, \lambda, J)$ -module is the same as a tuple of $\mathcal{S}_I(Y_i, \lambda_i, J_i)$ -modules for $i \in I$ (and the same for \mathcal{S}_{II} -modules). There are also natural inclusions

$$A_I(Y, \lambda, J)(T) \hookrightarrow A_I(Y_i, \lambda_i, J_i)(T) \quad \text{for } T \in \mathcal{S}_I(Y_i, \lambda_i, J_i), \quad (4.84)$$

$$A_{II}(\hat{X}, \hat{\omega}, \hat{J})(T) \hookrightarrow A_{II}(\hat{X}_i, \hat{\omega}_i, \hat{J}_i)(T) \quad \text{for } T \in \mathcal{S}_{II}(\hat{X}_i, \hat{\omega}_i, \hat{J}_i). \quad (4.85)$$

Precisely, these maps preserve r , E , D , and extend λ as zero on $Y \setminus Y_i$ (resp. $\hat{X} \setminus \hat{X}_i$).

Now we modify the definition of Θ to include a choice of $\text{Aut}(T)$ -invariant subatlases $B(T) \subseteq A(T)$ on $\overline{\mathcal{M}}(T)$ for all maximal $T \in \mathcal{S}$ (e.g. in case (II), we choose $B_I^\pm(T) \subseteq A_I^\pm(T)$ and $B_{II}(T) \subseteq A_{II}(T)$). We define $\bar{B}(T)$ as in (3.10), and we use \bar{B} in place of \bar{A} in the definitions from §§4.6–4.7.

With this modification done, the map (4.80) is defined by taking the images of the sets $B_I(T)$ under (4.84) and using the “same” diagrams of \mathcal{S}_I -modules. The map (4.81) is defined similarly. It follows by definition that these maps are functorial and preserve the virtual moduli counts. \square

5 Gluing

This section is devoted to the proof of Theorems 3.31 and 3.32 (of which Theorems 2.54 and 2.55 are special cases). Namely, we prove that the regular loci in the thickened moduli spaces $\overline{\mathcal{M}}(T)_I$ admit the expected local topological description in terms of the spaces $G_{T'//T}$ from §2.5, and we verify that the natural “geometric” and “analytic” maps between orientation lines agree. We will give the argument for all cases (I), (II), (III), (IV) simultaneously.

The gluing techniques we use in this section are, to the best of our knowledge, standard in the field (perhaps the most similar published result is Hutchings–Taubes [HT07, HT09]). We learned these methods (while writing the gluing theorems in [Par16, §§B–C]) from Abouzaid [Abo12], McDuff–Salamon [MS04, MS94], and discussions with Mohammed Abouzaid, Tobias Ekholm, Helmut Hofer, and Rafe Mazzeo.

5.1 Gluing setup

Proof of Theorem 3.31. Fix $T \in \mathcal{S}$, $I \subseteq J \subseteq \bar{A}(T)$, and $x_0 \in \mathcal{M}(T)_J$ of type $T' \rightarrow T$ with $s_{J \setminus I}(x_0) = 0$ and $\psi_{I,J}(x_0) \in \mathcal{M}(T)_I^{\text{reg}}$. A neighborhood of $x_0 \in \overline{\mathcal{M}}(T)_J$ is stratified by $\mathcal{S}_{T'//T}$.

Our goal is to construct a local homeomorphism

$$(G_{T'//T} \times E_{J \setminus I} \times \mathbb{R}^{\mu(T') - \#V_s(T') + \dim E_I}, (0, 0, 0)) \rightarrow (\overline{\mathcal{M}}(T)_J, x_0) \quad (5.1)$$

which lands in $\overline{\mathcal{M}}(T)_J^{\text{reg}}$ and which commutes with the maps from both sides to $\mathcal{S}_{T'//T} \times \mathfrak{s}(T) \times E_{J \setminus I}$. We denote by $0 \in G_{T'//T}$ the basepoint corresponding to all gluing parameters equal to ∞ (i.e. corresponding to no gluing at all) and $t = t(x_0)$.

The basepoint x_0 corresponds to a map $u_0 : C_0 \rightarrow \hat{X}_0$ and an element $e_0 \in E_I \subseteq E_J$, along with “discrete data” consisting of asymptotic markers, matching isomorphisms, and markings $\phi_\alpha : (C_0)_\alpha \rightarrow \overline{\mathcal{C}}_{0, E^{\text{ext}}(T_\alpha) \cup \{1, \dots, r_\alpha\}}$ for $\alpha \in J$.

5.1.1 Domain stabilization via submanifolds $\hat{D}_{v,i}$ and points $q_{v,i}$

We first stabilize the domain C_0 by adding marked points $q_{v,i}$ where it intersects certain codimension two submanifolds $\hat{D}_{v,i}$, arguing much the same as we did in §3.7. Such $\hat{D}_{v,i}$ automatically stabilize the domains of all maps in a neighborhood of $x_0 \in \overline{\mathcal{M}}(T)_J$, and thus we need only consider stable domain curves in the main gluing argument which follows.

For every vertex $v \in V(T')$, we add marked points $q_{v,i} \in (C_0)_v$ and choose local codimension two submanifolds $\hat{D}_{v,i} \subseteq (\hat{X}_0)_v$ (required to be \mathbb{R} -invariant if v is a symplectization vertex) with $u_0(q_{v,i}) \in \hat{D}_{v,i}$ intersecting transversally, such that $(C_0)_v$ equipped with the marked points $\{p_{v,e}\}_e$ and $\{q_{v,i}\}_i$ is stable. These new marked points $q_{v,i}$ bear no logical relation to the marked points labeled by $\{1, \dots, r_\alpha\}$ for $\alpha \in J$ (in particular, there is no need to require that the $q_{v,i}$ be disjoint from these other marked points).

To show the existence of such points, it suffices to show that each unstable irreducible component of $(C_0)_v$ (equipped with just $\{p_{v,e}\}_e$ as marked points) has a point (and hence a non-empty open set) where du_0 (resp. $\pi_\xi du_0$ if v is a symplectization vertex) is injective. If $(C_0)_v \not\subseteq (C_0)_\alpha$ for all $\alpha \in I$, then $u_0|(C_0)_v$ is \hat{J}_v -holomorphic and the existence of such points follows from the arguments given in the proofs of Lemmas 3.26 and 3.28. For components $(C_0)_v \subseteq (C_0)_\alpha$ for some $\alpha \in J$, such points exist since $u_v|(C_0)_v$ satisfies Definition 3.17(iii).

5.1.2 Glued cobordisms \hat{X}_g and points q'_v

Given any gluing parameter $\mathbf{g} \in G_{T'//T}$, we may form the glued cobordism \hat{X}_g as follows (all statements involving a choice of $\mathbf{g} \in G_{T'//T}$ carry the (often tacit) assumption that \mathbf{g} lies in a sufficiently small neighborhood of 0). Note that \hat{X}_0 is equipped with cylindrical

coordinates (1.8)–(1.9) in each end. Now we truncate each positive (resp. negative) end $[0, \infty) \times Y_e$ (resp. $(-\infty, 0] \times Y_e$) to $[0, \mathbf{g}_e] \times Y_e$ (resp. $[-\mathbf{g}_e, 0] \times Y_e$) and identify truncated ends by translation by \mathbf{g}_e (if $\mathbf{g}_e = \infty$ we do nothing). By construction, $\hat{X}_{\mathbf{g}}$ is the target for pseudo-holomorphic buildings of type $T'_{\mathbf{g}}$ (denoting by $T' \rightarrow T'_{\mathbf{g}} \rightarrow T$ the image of \mathbf{g} under the map $G_{T'/T} \rightarrow \mathcal{S}_{T'/T}$).

For every symplectization vertex $v \in V_s(T')$, fix a section q'_v of the universal curve over a neighborhood of $C_0 \in \overline{\mathcal{M}}_{0, E^{\text{ext}}(T') \cup \{q_{v,i}\}_{v,i}}$ such that $q'_v(C_0) \in (C_0)_v$. Now over a neighborhood of $x_0 \in \overline{\mathcal{M}}(T)_J$, the sections q'_v determine points q'_v in the domain of each map. Over this neighborhood, we can identify the target with $\hat{X}_{\mathbf{g}}$ for a unique $\mathbf{g} \in G_{T'/T}$ by requiring that each q'_v be mapped to the corresponding “zero level” $0_v \subseteq \hat{X}_{\mathbf{g}}$ (i.e. the descent of $0_v := \{0\} \times Y_v \subseteq \hat{X}_0$).

5.1.3 Family of almost complex structures j_y on C_0

We now proceed to fix a family of almost complex structures on C_0 inducing a local biholomorphism onto the relevant moduli space of marked Riemann surfaces. Let $N \subseteq C_0$ denote the set of nodes (note that $\{p_{v,e}\}_{v,e}$ are not nodes), and let $N_v := N \cap (C_0)_v$.

We denote by $\overline{\mathcal{M}}_{0, n+m(2)} \rightarrow \overline{\mathcal{M}}_{0, n+m}$ the total space of the bundle over $\overline{\mathcal{M}}_{0, n+m}$ with fiber $\prod_{i=1}^m (T_{q_i} C \setminus 0)$ (thought of as the space of markings $\mathbb{C} \xrightarrow{\sim} T_{q_i} C$ at the last m marked points). As in §§2.6–2.7, the tangent space to the nodal stratum $\overline{\mathcal{M}}_{0, n+m(2)}^{\#\text{nodes}=r}$ at a given curve C with marked points P and doubly marked points Q is canonically isomorphic to $H^1(\tilde{C}, T\tilde{C}(-\tilde{P} - 2\tilde{Q} - \tilde{N}))$.

For every vertex $v \in V(T')$, fix a local biholomorphism

$$\mathcal{J}_v := \mathbb{C}^{2\#\{p_{v,e}\}_e + \#\{q_{v,i}\}_i - \#N_v - 3} \rightarrow \overline{\mathcal{M}}_{0, \{q_{v,i}\}_i + (\{p_{v,e}\}_e)_{(2)}}^{\#\text{nodes}=\#N_v} \quad (5.2)$$

sending 0 to $(C_0)_v$ equipped with its doubly marked points $\{p_{v,e}\}_e$ and marked points $\{q_{v,i}\}_i$. Further fix a smooth trivialization of the universal curve over these spaces (over a neighborhood of $(C_0)_v$, that is the origin in \mathcal{J}_v), which near the nodes and marked points of the fibers is holomorphic (on total spaces) and preserves the tangent space markings.

We thus obtain a description of the universal curve over \mathcal{J}_v as the curve $(C_0)_v$ with its doubly marked points $\{p_{v,e}\}_e$, its marked points $\{q_{v,i}\}_i$, and a modified almost complex structure j_y varying smoothly in $y \in \mathcal{J}_v$ and fixed in a neighborhood of $\{p_{v,e}\}_e$, $\{q_{v,i}\}_i$, and N_v . Since (5.2) is a local biholomorphism, the derivative

$$\frac{dj_y}{dy} : T_0 \mathcal{J}_v = \mathbb{C}^{2\#\{p_{v,e}\}_e + \#\{q_{v,i}\}_i - \#N_v - 3} \rightarrow C_c^\infty((C_0)_v \setminus (\{q_{v,i}\}_i \cup \{p_{v,e}\}_e \cup N_v), \text{End}^{0,1}(T(C_0)_v)) \quad (5.3)$$

induces an isomorphism between its domain and the deformation space of $(C_0)_v$ equipped with its doubly marked points $\{p_{v,e}\}_e$ and marked points $\{q_{v,i}\}_i$ as recalled above. We set

$$\mathcal{J} := \prod_{v \in V(T')} \mathcal{J}_v. \quad (5.4)$$

All statements involving a choice of $y \in \mathcal{J}_v$ or $y \in \mathcal{J}$ carry the (often tacit) assumption that y lies in a sufficiently small neighborhood of 0.

5.1.4 Cylindrical coordinates on C_0 and \hat{X}_0

We now fix positive (resp. negative) holomorphic cylindrical coordinates

$$[0, \infty) \times S^1 \rightarrow C_0 \quad (5.5)$$

$$(-\infty, 0] \times S^1 \rightarrow C_0 \quad (5.6)$$

near each positive (resp. negative) puncture. We also fix such cylindrical coordinates on either side of each node $n \in N$ (choosing which side is positive/negative arbitrarily).

We assume that with respect to these fixed cylindrical coordinates, we have

$$u_0(s, t) = (Ls, \tilde{\gamma}(t)) + o(1) \quad \text{as } |s| \rightarrow \infty \quad (5.7)$$

(i.e. the constant b in (2.1) vanishes) with respect to the cylindrical coordinates (1.8)–(1.9) on \hat{X}_0 . To ensure that we can achieve (5.7), we allow the possibility that (5.5)–(5.6) are only defined for $|s|$ sufficiently large. Recall from Proposition 2.20 due to Hofer–Wysocki–Zehnder [HWZ02, Theorems 1.1, 1.2, and 1.3] that the convergence (5.7) holds in C^∞ with error $O(e^{-\delta|s|})$ for every $\delta < \delta_\gamma$.

5.1.5 Glued curves $C_{\mathbf{g}}$ and points q_e''

Given any gluing parameter $\beta \in \mathbb{C}^{E^{\text{int}}(T')} \times \mathbb{C}^N$ (i.e. one for each interior edge of T' and one for each node $n \in N$), we may form the glued curve C_β as follows (all statements involving a choice of $\beta \in \mathbb{C}^{E^{\text{int}}(T')} \times \mathbb{C}^N$ carry the (often tacit) assumption that β lies in a sufficiently small neighborhood of 0). For each interior edge $v \xrightarrow{e} v'$ (or node $n \in N$), we truncate the positive (resp. negative) end $[0, \infty) \times S^1$ (resp. $(-\infty, 0] \times S^1$) to $[0, S] \times S^1$ (resp. $[-S, 0] \times S^1$) and identify them by $s = s' + S$ and $t = t' + \theta$ where $\beta = e^{-S-i\theta}$ (if $\beta = 0$ we do nothing). Note that the points $q_{v,i}$ and the complex structures j_y both descend naturally to C_β . This operation gives analytic families of curves over $\mathcal{J}_v \times \mathbb{C}^{N_v}$ and $\mathcal{J} \times \mathbb{C}^{E^{\text{int}}(T')} \times \mathbb{C}^N$ inducing local biholomorphisms

$$\mathcal{J}_v \times \mathbb{C}^{N_v} \rightarrow \overline{\mathcal{M}}_{0, \{q_{v,i}\}_i + (\{p_{v,e}\}_e)_{(2)}}, \quad (5.8)$$

$$\mathcal{J} \times \mathbb{C}^N \rightarrow \prod_{v \in V(T')} \overline{\mathcal{M}}_{0, \{q_{v,i}\}_i + (\{p_{v,e}\}_e)_{(2)}}. \quad (5.9)$$

The glued curve C_β is actually a bit too general for our purposes: in the current setting the gluing parameters at interior edges must come from the same gluing parameters $\mathbf{g} \in G_{T'/T}$ used to glue the target cobordism.

Namely, to any $\mathbf{g} \in G_{T'/T}$, we associate the gluing parameter $\beta = \beta(\mathbf{g}) \in \mathbb{C}^{E^{\text{int}}(T')}$ given by $S_e := L_e^{-1} \mathbf{g}_e$ and the unique θ_e corresponding to the matching isomorphism $S_{p_{v',e}} C_{v'} \rightarrow S_{p_{v,e}} C_v$. For $\mathbf{g} \in G_{T'/T} \times \mathbb{C}^N$, we thus have a glued curve $C_{\mathbf{g}}$. For e with $\mathbf{g}_e < \infty$, let $q_e'' \in C_{\mathbf{g}}$ denote the point in the middle of the corresponding neck, namely $s = \pm \frac{1}{2}S$ with respect to the coordinates (5.5)–(5.6) (the angular coordinate $t \in S^1$ of q_e'' is irrelevant as long as it is fixed). Now for any fixed $\mathbf{g} \in G_{T'/T}$, this construction gives a map

$$\mathcal{J} \times \mathbb{C}^N \rightarrow \prod_{v \in V(T'_{\mathbf{g}})} \overline{\mathcal{M}}_{0, \{q_{v',i}\}_{v' \mapsto v,i} + \{q_e''\}_{e \mapsto v} + (\{p_{v,e}\}_e)_{(2)}}. \quad (5.10)$$

This map is a local biholomorphism near $0 \in \mathcal{J} \times \mathbb{C}^N$ uniformly in \mathbf{g} , in the sense that there is a neighborhood of $0 \in \mathcal{J} \times \mathbb{C}^N$ over which (5.10) is a biholomorphism onto its image for all $\mathbf{g} \in G_{T'/T}$ near zero (we will prove this just below).

The purpose of including the points q_e'' as part of the structure of the glued curve is exactly to ensure that (5.10) is a local biholomorphism. Later, we will impose a codimension two condition at q_e'' on our glued maps, so that over a neighborhood of $x_0 \in \overline{\mathcal{M}}(T)_J$, the points q_e'' in the domain are determined uniquely, and hence the parameters $(y, \beta) \in \mathcal{J} \times \mathbb{C}^N$ are determined uniquely. We denote by $q_v' \in (C_{\mathbf{g}}, j_y)$ the value of the section q_v' at $(C_{\mathbf{g}}, j_y)$ for $\mathbf{g} \in G_{T'/T} \times \mathbb{C}^N$ and $y \in \mathcal{J}$ (note that this q_v' may *not* coincide with the descent of $q_v' \in C_0$, even for $\mathbf{g} = 0$).

To conclude, let us give the proof that (5.10) is a local biholomorphism, as claimed above. In fact, to make the proof work we impose the following additional compatibility condition between the cylindrical coordinates (5.5)–(5.6) and the smooth trivialization of the universal curve over $\overline{\mathcal{M}}_{0, \{q_{v,i}\}_i + (\{p_{v,e}\}_e)_{(2)}}^{\#\text{nodes}=\#N_v}$ near $(C_0)_v$ fixed in §5.1.3. The universal curve is pulled back under the forgetful map $\overline{\mathcal{M}}_{0, \{q_{v,i}\}_i + (\{p_{v,e}\}_e)_{(2)}} \rightarrow \overline{\mathcal{M}}_{0, \{q_{v,i}\}_i + \{p_{v,e}\}_e}$ which is a principal $(\mathbb{C} \setminus 0)^{\{p_{v,e}\}_e}$ -bundle (acting via changing the tangent space markings at $\{p_{v,e}\}_e$). We require that the induced action of $(\mathbb{C} \setminus 0)^{\{p_{v,e}\}_e}$ on the total space, when pulled back to $(C_0)_v$ under the chosen trivialization and then expressed in the cylindrical coordinates (5.5)–(5.6) near each $p_{v,e}$, must coincide with the tautological action of $\mathbb{C} \setminus 0$ on the standard end $[0, \infty) \times S^1$ (resp. $(-\infty, 0] \times S^1$). To construct such a smooth trivialization compatible with a given choice of cylindrical coordinates (5.5)–(5.6), first fix the trivialization over the image of a section of the forgetful map, then extend to a neighborhood of $\{p_{v,e}\}_e$ as dictated by the desired compatibility property, and then finally extend to the rest.

Given this compatibility property, the local biholomorphism property for (5.10) now follows easily from the standard local biholomorphism property for the usual ‘local gluing charts’ on Deligne–Mumford space. Namely, for any open subset $U \subseteq \overline{\mathcal{M}}_{0,n}^{\#\text{nodes}=r} \subseteq \overline{\mathcal{M}}_{0,n}$ (or $\overline{\mathcal{M}}_{0,n+m_{(2)}}$) and any choice of local holomorphic (on the total space) cylindrical coordinates

$$[0, \infty) \times S^1 \times \overline{\mathcal{M}}_{0,n}^{\#\text{nodes}=r} \rightarrow \overline{\mathcal{C}}_{0,n}^{\#\text{nodes}=r} \quad (5.11)$$

(defined over U) on each side of each of the r nodes, there is an induced holomorphic map

$$U \times \mathbb{C}^r \rightarrow \overline{\mathcal{M}}_{0,n} \quad (5.12)$$

given by the usual gluing construction (via $e^{-s_1-it_1} e^{-s_2-it_2} = \beta$). It is well-known that (5.12) is a local biholomorphism near U (indeed, it is enough to verify that for every curve C (a point of U) and every node $a \in C$, the map $\mathbb{C} \rightarrow \overline{\mathcal{M}}_{0,n}$ given by gluing C at a with respect to (5.11) with gluing parameter $\beta \in \mathbb{C}$ induces an isomorphism between $T_0\mathbb{C}$ and the tangent space to the local branch of $\overline{\mathcal{M}}_{0,n}^{\#\text{nodes}=r+1}$ corresponding to a , and this can be shown by a standard calculation). Now to deduce the local biholomorphism property for (5.10), it suffices to apply the local biholomorphism property for (5.12) in a setup which depends only on $T'_{\mathbf{g}}$ (and hence is uniform in \mathbf{g}).

5.1.6 Preglued maps $u_{\mathbf{g}}$

We now define a “preglued” map $u_{\mathbf{g}} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$ close to $u_0 : C_0 \rightarrow \hat{X}_0$. As we shall see later, this preglued map is very close to solving the relevant (thickened) pseudo-holomorphic curve equation. Our goal will then be to understand the true solutions near $u_{\mathbf{g}}$ and to show that this construction gives a local parameterization of the moduli space near u_0 .

Fix a smooth (cutoff) function $\chi : \mathbb{R} \rightarrow [0, 1]$ satisfying

$$\chi(x) = \begin{cases} 1 & x \leq 0 \\ 0 & x \geq 1. \end{cases} \quad (5.13)$$

Definition 5.1 (Flattening). For $\mathbf{g} \in G_{T'/T} \times \mathbb{C}^N$, we define the “flattened” map

$$u_{0|\mathbf{g}} : C_0 \rightarrow \hat{X}_0 \quad (5.14)$$

as follows. Away from the ends, $u_{0|\mathbf{g}}$ coincides with u_0 . Over a positive end asymptotic to a parameterized Reeb orbit $\tilde{\gamma}(t) := u_0(\infty, t)$, we define $u_{0|\mathbf{g}}$ by

$$u_{0|\mathbf{g}}(s, t) := \begin{cases} u_0(s, t) & s \leq \frac{1}{6}S \\ \exp_{(Ls, \tilde{\gamma}(t))} \left[\chi\left(s - \frac{1}{6}S\right) \cdot \exp_{(Ls, \tilde{\gamma}(t))}^{-1} u_0(s, t) \right] & \frac{1}{6}S \leq s \leq \frac{1}{6}S + 1 \\ (Ls, \tilde{\gamma}(t)) & \frac{1}{6}S + 1 \leq s. \end{cases}$$

Over a positive end at a node $n \in N$, we define $u_{0|\mathbf{g}}$ by

$$u_{0|\mathbf{g}}(s, t) := \begin{cases} u_0(s, t) & s \leq \frac{1}{6}S \\ \exp_{u_0(n)} \left[\chi\left(s - \frac{1}{6}S\right) \cdot \exp_{u_0(n)}^{-1} u_0(s, t) \right] & \frac{1}{6}S \leq s \leq \frac{1}{6}S + 1 \\ u_0(n) & \frac{1}{6}S + 1 \leq s. \end{cases}$$

An analogous definition applies over negative ends. Here $\exp : T\hat{X}_0 \rightarrow \hat{X}_0$ denotes any fixed exponential map (i.e. a smooth map defined in a neighborhood of the zero section satisfying $\exp(p, 0) = p$ and $d\exp(p, \cdot) = \text{id}_{T_p\hat{X}_0}$) which is \mathbb{R} -equivariant in any end.

Definition 5.2 (Pregluing). For $\mathbf{g} \in G_{T'/T} \times \mathbb{C}^N$, we define the “preglued” map

$$u_{\mathbf{g}} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}} \quad (5.15)$$

as the natural “descent” of $u_{0|\mathbf{g}}$ from C_0 to $C_{\mathbf{g}}$.

5.2 Gluing estimates

With the above setup understood, our aim is now to describe the “true solutions” close to the “approximate solution” $u_{\mathbf{g}} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$. This forms the core part of the gluing argument.

5.2.1 Weighted Sobolev norms

Our first step is simply to define the weighted Sobolev spaces $W^{k,2,\delta}$ relevant for the gluing argument and to specify norms on these spaces, up to commensurability uniform in \mathbf{g} near zero (uniformity in \mathbf{g} is crucial for the key gluing estimates to have any meaning or utility at all). The weighted Sobolev spaces we define here differ from those of §2.11 in that we put weights not just over the ends asymptotic to Reeb orbits but also near the nodes (this necessitates a comparison of linearized operators in the two contexts, which we do immediately after the definition of Sobolev norms below). We also put weights over the necks (which has meaning since we work up to commensurability uniform in \mathbf{g}).

Fix metrics and connections as in Conventions 2.30–2.31 on C_0 and \hat{X}_0 with respect to the cylindrical coordinates fixed above and which (for convenience) agree across the parts to be glued (thus descending to $C_{\mathbf{g}}$ and $\hat{X}_{\mathbf{g}}$). We suppose throughout that $\delta > 0$ satisfies $\delta < 1$ and $\delta < \delta_\gamma$ for every asymptotic Reeb orbit γ of u_0 . As in Remark 2.33, it is sufficient to fix any $k \geq 4$, though for the moment we only need to assume $k \geq 0$ ($k \geq 2$ if C_0 has nodes).

Remark 5.3. Although these choices of metrics, connections, k , and δ affect the constants appearing in the gluing estimates, they do *not* affect either the final gluing map or (the restriction to smooth functions/sections of) any of the intermediate maps we study and/or construct.

Away from the ends/neck of $C_{\mathbf{g}}$, we use the usual $W^{k,2}$ -norm. In an end asymptotic to a Reeb orbit, the contribution to the norm squared is

$$\int_{[0,\infty) \times S^1} \sum_{j=0}^k |D^j f|^2 e^{2\delta s} ds dt. \quad (5.16)$$

In an end asymptotic to a node $n \in N$, we distinguish two cases: for Sobolev spaces of $(0,1)$ -forms on $C_{\mathbf{g}}$, we use (5.16), and for spaces of functions we allow decay to a constant, i.e. the contribution to the norm squared is

$$|f(n)|^2 + \int_{[0,\infty) \times S^1} \sum_{j=0}^k |D^j [f - f(n)]|^2 e^{2\delta s} ds dt \quad (5.17)$$

The contribution of a neck is given as follows (first is for necks over Reeb orbits and necks over nodes for spaces of $(0,1)$ -forms; second is for necks over nodes for spaces of functions):

$$\int_{[0,S] \times S^1} \sum_{j=0}^k |D^j f|^2 e^{2\delta \min(s, S-s)} ds dt, \quad (5.18)$$

$$\int_{[0,S] \times S^1} \sum_{j=0}^k \left| D^j \left[f - \frac{1}{2\pi} \int_{S^1} f \left(\frac{1}{2}S, t \right) dt \right] \right|^2 e^{2\delta \min(s, S-s)} ds dt + \left| \frac{1}{2\pi} \int_{S^1} f \left(\frac{1}{2}S, t \right) dt \right|^2. \quad (5.19)$$

We use parallel transport to make sense of the differences $f - f(n)$, etc. To measure/differentiate $(0,1)$ -forms, we use the usual Riemannian metric on $[0,\infty) \times S^1$ (resp. $[0,S] \times S^1$) and usual flat connection on forms. Different choices of metrics and connections yield norms which are equivalent, uniformly in \mathbf{g} , for any fixed k and δ as above.

Remark 5.4. We emphasize that, although we refer to the weighted Sobolev norms in §2.11 informally as “without weights at the nodes”, they are *not* the same as the weighted Sobolev norms above specialized to $\delta = 0$ at the nodes.

We now provide the promised comparison of linearized operators defined using the weighted Sobolev norms above (with weights near the nodes) and using the weighted Sobolev norms from §2.11 (without weights near the nodes). That is, we argue that both linearized operators are Fredholm, with the same kernel and cokernel. Let us temporarily write $W^{k,2,\delta,\delta}$ for the weighted Sobolev spaces defined above and $W^{k,2,\delta}$ for the weighted Sobolev spaces from §2.11. We suppose that k is large enough so that the linearized operators are defined (see Definitions 2.37 and 3.22). We have already seen in Proposition 2.38 that the linearized operator acting on $W^{k,2,\delta}$ is Fredholm for such δ . The same argument applies to the linearized operator acting on $W^{k,2,\delta,\delta}$, using the fact that $\delta \in (0, 1)$ is not an eigenvalue of the asymptotic operators at the nodes (all of which have spectrum \mathbb{Z}). There are inclusions $W^{k,2,\delta} \subseteq W^{k,2,\delta,\delta}$ (since $\delta < 1$) which are dense by Lemma 2.35. We get induced maps from the kernel and cokernel of the linearized operator acting on $W^{k,2,\delta}$ to those of the operator acting on $W^{k,2,\delta,\delta}$. The map on kernels is obviously injective, and the map on cokernels is obviously surjective. To show surjectivity on kernels and injectivity on cokernels, it suffices to show that if $D\xi = \eta$ with $\xi \in W^{k,2,\delta,\delta}$ and $\eta \in W^{k-1,2,\delta}$, then $\xi \in W^{k,2,\delta}$. Now as distributions, we have $D\xi = \eta + \varepsilon$, for some distribution ε supported over the inverse images of the nodes $\tilde{N}_v \subseteq \tilde{C}_v$. By elliptic regularity, it suffices to show that $\varepsilon = 0$. But now for any smooth test function φ supported over a small neighborhood of \tilde{N}_v , we have

$$\langle \varepsilon, \varphi \rangle = \langle D\xi - \eta, \varphi \rangle = \langle \xi, D^*\varphi \rangle - \langle \eta, \varphi \rangle. \quad (5.20)$$

The first term is bounded by a constant times $\|\varphi\|_{1,1}$ since D^* (the formal adjoint) is first order and $\xi \in W^{k,2,\delta,\delta} \subseteq C^0$. The second term is bounded by a constant times $\|\varphi\|_2$ since $\eta \in W^{k-1,2,\delta} \subseteq L^2$. Now ε is supported at a finite number of points, so must be a linear combination of δ -functions and their derivatives (see Lemma 2.41), but these do not satisfy such bounds (recall that $W^{1,1} \not\subseteq C^0$ since we are in two dimensions [Ada75, Example 5.25]).

5.2.2 Nonlinear Fredholm setup for fixed \mathbf{g}

We now formulate precisely what we mean by “solutions close to $u_{\mathbf{g}} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$ ”. What we mean is “small zeroes of $\mathcal{F}_{\mathbf{g}}$ ”, where $\mathcal{F}_{\mathbf{g}}$ is the map

$$\begin{aligned} \mathcal{F}_{\mathbf{g}} : W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T\hat{X}_{\mathbf{g}}) &\oplus \mathcal{J} \oplus E_J \\ &\xrightarrow[\pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda} \xi(q''_e)=0]{\xi(q_{v,i}) \in T\hat{D}_{v,i}} W^{k-1,2,\delta}(\tilde{C}_{\mathbf{g}}, u_{\mathbf{g}}^*(T\hat{X}_{\mathbf{g}})_{\hat{J}_{\mathbf{g}t}} \otimes_{\mathbb{C}} \Omega_{C_{\mathbf{g}}}^{0,1}) \oplus \mathbb{R}^{V_s(T')} \end{aligned} \quad (5.21)$$

defined by

$$\begin{aligned} \mathcal{F}_{\mathbf{g}}(\xi, y, e) := & \left[(\text{PT}_{\exp_{u_{\mathbf{g}}}}^{\mathbf{g}t} \xi \rightarrow u_{\mathbf{g}}} \otimes I_y) \left(d(\exp_{u_{\mathbf{g}}} \xi) + \sum_{\alpha \in J} \nu_{\alpha} ((e_0 + e)_{\alpha}) (\phi_{\alpha}^{\mathbf{g}, \xi}, \exp_{u_{\mathbf{g}}} \xi) \right)_{j_y, \hat{J}_{\mathbf{g}t}} \right]^{0,1} \\ & \oplus \bigoplus_{v \in V_s(T')} \pi_{\mathbb{R}}(\exp_{u_{\mathbf{g}}} \xi)(q'_v(y)). \end{aligned} \quad (5.22)$$

We explain the notation. We denote by $\exp : T\hat{X}_0 \rightarrow \hat{X}_0$ a fixed exponential map which is \mathbb{R} -equivariant in ends and over symplectizations and agrees across the parts to be glued, thus descending to $\hat{X}_{\mathbf{g}}$. We fix a \hat{J}_0 -linear connection on $T\hat{X}_0$ which is \mathbb{R} -equivariant in ends and over symplectizations and agrees across the parts to be glued, thus descending to $\hat{X}_{\mathbf{g}}$. We denote by $\text{PT}^{\mathbf{g}t}$ parallel transport with respect to the $\hat{J}_{\mathbf{g}t}$ -linear part of this fixed connection. The map $I_y : \Omega_{C_{\mathbf{g},j_y}}^{0,1} \rightarrow \Omega_{C_{\mathbf{g},j_0}}^{0,1}$ denotes the composition $\Omega_{C_{\mathbf{g},j_y}}^{0,1} \rightarrow \Omega_{C_{\mathbf{g}}}^1 \otimes_{\mathbb{R}} \mathbb{C} \rightarrow \Omega_{C_{\mathbf{g},j_0}}^{0,1}$ (which is \mathbb{C} -linear).

The map $\mathcal{F}_{\mathbf{g}}$ is defined over the ball of some fixed radius $c'_{k,\delta} > 0$ uniformly in \mathbf{g} near zero, for any $k \geq 3$. The constraint $k \geq 3$ ensures that $W^{k,2} \subseteq C^1$ (see [Ada75, Lemma 5.17]), which is needed so that $\phi_{\alpha}^{\mathbf{g},\xi}$ is defined for $\|\xi\|_{k,2,\delta}$ small.

5.2.3 Estimate for $\|\mathcal{F}_{\mathbf{g}}(0)\|$

We now show that $\mathcal{F}_{\mathbf{g}}(0)$ is very small (i.e. the preglued map $u_{\mathbf{g}} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$ is very close to being a true solution).

Lemma 5.5. *We have*

$$\|\mathcal{F}_{\mathbf{g}}(0)\|_{k-1,2,\delta} \rightarrow 0 \quad \text{as } \mathbf{g} \rightarrow 0 \quad (5.23)$$

for all $k \geq 1$.

Proof. Away from the necks and ends, the 1-form part of $\mathcal{F}_{\mathbf{g}}(0)$ is only nonzero because of using $\hat{J}_{\mathbf{g}t}$ in place of \hat{J}_0 and because of using $\phi_{\alpha}^{\mathbf{g},0}$ in place of $\phi_{\alpha}^{0,0} = \phi_{\alpha}$. This difference clearly goes to zero as $\mathbf{g} \rightarrow 0$.

Over Reeb ends, the 1-form part of $\mathcal{F}_{\mathbf{g}}(0)$ is identically zero. Over Reeb necks, the 1-form part of $\mathcal{F}_{\mathbf{g}}(0)$ is supported near $\frac{1}{6}S$ and $\frac{5}{6}S$, and the desired estimate follows from the exponential convergence of (5.7) and the fact that $\delta < \delta_{\gamma}$.

The same applies to nodal ends/neck, except that in addition there is a term coming from using $\hat{J}_{\mathbf{g}t}$ in place of \hat{J}_0 . This again is bounded as desired since $\delta < 1$.

The $\mathbb{R}^{V_s(T')}$ part of $\mathcal{F}_{\mathbf{g}}(0)$ satisfies the desired estimate since $q'_v \in (C_{\mathbf{g}}, j_0)$ approaches the descent of $q'_v \in (C_0, j_0)$ as $\mathbf{g} \rightarrow 0$. \square

5.2.4 Regularity of the map $\mathcal{F}_{\mathbf{g}}$ (quadratic estimate)

We now prove a “quadratic estimate” quantifying the regularity of $\mathcal{F}_{\mathbf{g}}$ near zero, i.e. giving a uniform upper bound on certain of its derivatives near zero. This estimate is used when we apply the (Banach space) inverse function theorem to understand $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ near zero. We begin with a general discussion of the smoothness properties of $\mathcal{F}_{\mathbf{g}}$ before specializing to the specific estimate we need.

The first term in $\mathcal{F}_{\mathbf{g}}$ (the usual pseudo-holomorphic curve equation) is smooth and local. The second term (the “thickening” terms ν_{α}) is non-local; its only non-smoothness comes from the association $\xi \mapsto \phi_{\alpha}^{\mathbf{g},\xi}$. It thus is C^{ℓ} as long as the function which assigns to ξ the set $(\exp_{u_{\mathbf{g}}} \xi)^{-1}(\hat{D}_{\alpha})$ is C^{ℓ} . By the inverse function theorem, this is the case whenever $W^{k,2} \subseteq C^{\ell}$, which in turn holds whenever $k \geq \ell + 2$ [Ada75, Lemma 5.17]. The third term is also C^{ℓ} whenever $W^{k,2} \subseteq C^{\ell}$ (these both come down to the fact that the evaluation map $W^{k,2}(C, X) \times C \rightarrow X$ is of class C^{ℓ} whenever $W^{k,2} \subseteq C^{\ell}$).

The specific sort of estimate we need to apply the inverse function theorem to understand $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is a “quadratic estimate” quantifying, uniformly in \mathbf{g} , the fact that the derivative of $\mathcal{F}_{\mathbf{g}}$ is Lipschitz. It is thus sufficient to assume that $k \geq 4$, so that the above discussion shows that $\mathcal{F}_{\mathbf{g}}$ is of class C^2 .

Proposition 5.6. *For $\|\zeta\|_{k,2,\delta}, \|\xi\|_{k,2,\delta} \leq c'_{k,\delta}$, we have*

$$\|\mathcal{F}'_{\mathbf{g}}(0, \xi) - \mathcal{F}'_{\mathbf{g}}(\zeta, \xi)\|_{k-1,2,\delta} \leq c_{k,\delta} \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta} \quad (5.24)$$

for constants $c_{k,\delta} < \infty$ and $c'_{k,\delta} > 0$ uniformly in \mathbf{g} near 0, for all $k \geq 4$.

Proof. We have already observed above that $\mathcal{F}_{\mathbf{g}}$ is of class C^2 for $k \geq 4$, which implies that (5.24) holds for some $c_{k,\delta} < \infty$ and $c'_{k,\delta} > 0$ possibly depending on \mathbf{g} . Our present task is thus to make the above reasoning precise enough to extract bounds which are uniform in \mathbf{g} . In a word, the point of the proof is simply that the constants $c_{k,\delta} < \infty$ and $c'_{k,\delta} > 0$ in question depend only on the $W^{k,\infty}$ -norms of $u_{\mathbf{g}}$ and the various operations comprising $\mathcal{F}_{\mathbf{g}}$ (in particular, the total area of $C_{\mathbf{g}}$ is irrelevant).

Choose a covering $\{U_{\alpha}\}_{\alpha}$ of $C_{\mathbf{g}}$ of locally bounded geometry by open sets U_{α} of bounded diameter (e.g. over ends/necks, we can use the covering by sets of the form $[s, s+2] \times S^1$ for integers s). The local geometry of the cover can be bounded uniformly in \mathbf{g} , and thus we can write any $\xi \in W^{k,2,\delta}$ as a sum $\xi = \sum_{\alpha} \xi_{\alpha}$ where ξ_{α} is supported inside U_{α} and

$$\sum_{\alpha} \|\xi_{\alpha}\|_{k,2,\delta} \leq c_{k,\delta} \|\xi\|_{k,2,\delta} \quad (5.25)$$

for some constant $c_{k,\delta} < \infty$ uniform in \mathbf{g} . Hence it is enough to show the desired estimate (5.24) for ξ whose support has bounded diameter. We know from our earlier reasoning that the estimate (5.24) holds for all ξ , and it is clear that the resulting constants $c_{k,\delta} < \infty$ and $c'_{k,\delta} > 0$ depend only on the derivatives of $u_{\mathbf{g}}$ over the support of ξ (which we now assume has bounded diameter, and thus bounded geometry). The derivatives of $u_{\mathbf{g}}$ are bounded uniformly in \mathbf{g} , so we conclude that the constants $c_{k,\delta} < \infty$ and $c'_{k,\delta} > 0$ are as well. \square

Note that integrating (5.24) from ξ_1 to ξ_2 gives

$$\|D_{\mathbf{g}}(\xi_1 - \xi_2) - (\mathcal{F}_{\mathbf{g}}\xi_1 - \mathcal{F}_{\mathbf{g}}\xi_2)\|_{k-1,2,\delta} \leq c_{k,\delta} \|\xi_1 - \xi_2\|_{k,2,\delta} \max(\|\xi_1\|_{k,2,\delta}, \|\xi_2\|_{k,2,\delta}) \quad (5.26)$$

for $\|\xi_1\|_{k,2,\delta}, \|\xi_2\|_{k,2,\delta} \leq c'_{k,\delta}$.

5.2.5 Bounded right inverses and kernel gluing I: relating D_0 and $D_{\mathbf{g}}$

The final step in understanding $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is to construct a sufficiently nice bounded right inverse $Q_{\mathbf{g}}$ for $D_{\mathbf{g}} := \mathcal{F}'_{\mathbf{g}}(0, \cdot)$. In particular, we will need to show that $\|Q_{\mathbf{g}}\|$ is bounded uniformly for \mathbf{g} near 0, and that $\text{im } Q_{\mathbf{g}}$ “varies continuously” (in a sense which we will make precise) as \mathbf{g} varies. We will also construct a natural “kernel gluing” isomorphism $\ker D_0 \rightarrow \ker D_{\mathbf{g}}$.

To study the linearized operator $D_{\mathbf{g}}$, and in particular to construct $Q_{\mathbf{g}}$, we consider the following diagram, which allows us to relate $D_{\mathbf{g}}$ to D_0 .

$$\begin{array}{ccc}
W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T\hat{X}_{\mathbf{g}})_{\substack{\xi(q_{v,i}) \in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda} \xi(q_e'')=0}} \oplus \mathcal{J} \oplus E_J & \xrightarrow{D_{\mathbf{g}}} & W^{k-1,2,\delta}(\tilde{C}_{\mathbf{g}}, u_{\mathbf{g}}^*(T\hat{X}_{\mathbf{g}})_{j_{\mathbf{g}t}} \otimes_{\mathbb{C}} \Omega_{\tilde{C}_{\mathbf{g}}}^{0,1}) \oplus \mathbb{R}^{V_s(T')} \\
\uparrow \text{calib} & & \parallel \\
W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T\hat{X}_{\mathbf{g}})_{\xi(q_{v,i}) \in T\hat{D}_{v,i}} \oplus \mathcal{J} \oplus E_J & \xrightarrow{D_{\mathbf{g}}} & W^{k-1,2,\delta}(\tilde{C}_{\mathbf{g}}, u_{\mathbf{g}}^*(T\hat{X}_{\mathbf{g}})_{j_{\mathbf{g}t}} \otimes_{\mathbb{C}} \Omega_{\tilde{C}_{\mathbf{g}}}^{0,1}) \oplus \mathbb{R}^{V_s(T')} \\
\uparrow \text{glue} & & \downarrow \text{break} \\
W^{k,2,\delta}(C_0, u_{0|\mathbf{g}}^* T\hat{X}_0)_{\xi(q_{v,i}) \in T\hat{D}_{v,i}} \oplus \mathcal{J} \oplus E_J & \xrightarrow{D_{0|\mathbf{g}}} & W^{k-1,2,\delta}(\tilde{C}_0, u_{0|\mathbf{g}}^*(T\hat{X}_0)_{j_{\mathbf{g}t}} \otimes_{\mathbb{C}} \Omega_{\tilde{C}_0}^{0,1}) \oplus \mathbb{R}^{V_s(T')} \\
\uparrow \text{PT} & & \uparrow \text{PToid}^{1,0} \\
W^{k,2,\delta}(C_0, u_0^* T\hat{X}_0)_{\xi(q_{v,i}) \in T\hat{D}_{v,i}} \oplus \mathcal{J} \oplus E_J & \xrightarrow{D_0} & W^{k-1,2,\delta}(\tilde{C}_0, u_0^*(T\hat{X}_0)_{j_0} \otimes_{\mathbb{C}} \Omega_{\tilde{C}_0}^{0,1}) \oplus \mathbb{R}^{V_s(T')}
\end{array} \tag{5.27}$$

The maps in this diagram are defined as follows.

The horizontal maps are the linearized operators at the maps u_0 , $u_{0|\mathbf{g}}$, $u_{\mathbf{g}}$.

The maps PT are parallel transport with respect to the fixed connection on $T\hat{X}_0$.

The map $\text{id}^{1,0} : (T\hat{X}_0)_{j_0} \rightarrow (T\hat{X}_0)_{j_{\mathbf{g}t}}$ is the $(1,0)$ -component of the identity map.

Let us define the break map from (5.27). Fix a smooth function $\bar{\chi} : \mathbb{R} \rightarrow [0, 1]$ such that

$$\bar{\chi}(x) = \begin{cases} 1 & x \leq -1 \\ 0 & x \geq +1 \end{cases} \quad \bar{\chi}(x) + \bar{\chi}(-x) = 1. \tag{5.28}$$

Now $\text{break}(\eta)$ is simply η except over the ends of C_0 , where we define it to be

$$\text{break}(\eta)(s, t) := \begin{cases} \eta(s, t) & s \leq \frac{1}{2}S - 1 \\ \bar{\chi}(s - \frac{1}{2}S) \cdot \eta(s, t) & \frac{1}{2}S - 1 \leq s \leq \frac{1}{2}S + 1 \\ 0 & \frac{1}{2}S + 1 \leq s. \end{cases} \tag{5.29}$$

Thus the “trace” of $\text{break}(\eta)$ from C_0 to $C_{\mathbf{g}}$ (adding along fibers) is precisely η .

Let us define the glue map from (5.27). The map glue acts only on the vector field component (it acts identically on the other components). Away from the necks, we set $\text{glue}(\xi) := \xi$, and in any particular neck $[0, S] \times S^1 \subseteq C_{\mathbf{g}}$, we define (respectively for necks near Reeb orbits and necks near nodes $n \in N$)

$$\text{glue}(\xi)(s, t) := \begin{cases} \xi(s, t) & s \leq \frac{1}{3}S - 1 \\ \chi(s - \frac{2}{3}S)\xi(s, t) + \chi(\frac{2}{3}S - s')\xi(s', t') & \frac{1}{3}S - 1 \leq s \leq \frac{2}{3}S + 1 \\ \xi(s', t') & \frac{2}{3}S + 1 \leq s, \end{cases} \\
\text{glue}(\xi)(s, t) := \begin{cases} \xi(s, t) & s \leq \frac{1}{3}S - 1 \\ \xi(n) + \chi(s - \frac{2}{3}S)[\xi(s, t) - \xi(n)] \\ \quad + \chi(\frac{2}{3}S - s')[\xi(s', t') - \xi(n)] & \frac{1}{3}S - 1 \leq s \leq \frac{2}{3}S + 1 \\ \xi(s', t') & \frac{2}{3}S + 1 \leq s, \end{cases}$$

(noting the corresponding ends $(s, t) \in [0, \infty) \times S^1 \subseteq C_0$ and $(s', t') \in (-\infty, 0] \times S^1 \subseteq C_0$, glued via $s = s' + S$ and $t = t' + \theta$).

Let us define the calib map from (5.27). For every edge $e \in E^{\text{int}}(T')$ with $\mathbf{g}_e < \infty$, we consider the vector field $C_{\mathbf{g}} \rightarrow u_{\mathbf{g}}^* T \hat{X}_{\mathbf{g}}$ given in this neck by

$$\chi(s - \frac{2}{3}S)\chi(\frac{2}{3}S - s') \cdot \partial_s u_{\mathbf{g}}. \quad (5.30)$$

We denote by \mathbf{X} the \mathbb{C} -span of these vector fields. Now we have

$$W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T \hat{X}_{\mathbf{g}})_{\xi(q_{v,i}) \in T \hat{D}_{v,i}} = W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T \hat{X}_{\mathbf{g}})_{\substack{\xi(q_{v,i}) \in T \hat{D}_{v,i} \\ \pi_{\mathbb{R} \partial_s \oplus \mathbb{R} R_\lambda} \xi(q_e'')=0}} \oplus \mathbf{X}, \quad (5.31)$$

and the map calib is simply the associated projection onto the first factor.

This completes the definition of the maps in (5.27).

Lemma 5.7. *The maps in (5.27) are all bounded uniformly in \mathbf{g} .*

Proof. To show uniform boundedness of the linearized operators $D_0, D_{0|\mathbf{g}}, D_{\mathbf{g}}$, it is enough to cover the domain curve C_0 or $C_{\mathbf{g}}$ with small open sets of bounded geometry and to prove a uniform bound on $D_0, D_{0|\mathbf{g}}, D_{\mathbf{g}}$ restricted to sections supported in each such small open set (compare the proof of Proposition 5.6). For sake of concreteness, we can take these “small open sets of bounded geometry” to be of the form $[n, n+2] \times S^1$ in the ends/necks (note that the total number of these open sets does not matter). The maps $u_0, u_{0|\mathbf{g}}, u_{\mathbf{g}}$ are very well behaved in the ends/necks, uniformly in \mathbf{g} , which is enough. More succinctly, the point is simply that the norms of the linearized operators $D_0, D_{0|\mathbf{g}}, D_{\mathbf{g}}$ depend only on the $W^{k,\infty}$ -norms of their “coefficients” and of $u_0, u_{0|\mathbf{g}}, u_{\mathbf{g}}$ (so, in particular, the total area of $C_{\mathbf{g}}$ is irrelevant).

The maps break and glue are uniformly bounded because the derivatives of χ and $\bar{\chi}$ are bounded and the weights agree (up to a constant factor) at the relevant corresponding points of the domains.

To show that the map calib is uniformly bounded, it is enough to show that the projection onto the second factor in (5.31) is uniformly bounded. This projection is uniformly bounded since the norm of (5.30) is uniformly commensurable with its value at $q_e'' \in [0, S] \times S^1$ (which holds because $\delta > 0$ and q_e'' is given the maximum weight).

The maps PT and $\text{id}^{1,0}$ are bounded uniformly by the same argument used for the linearized operators $D_0, D_{0|\mathbf{g}}, D_{\mathbf{g}}$. \square

5.2.6 Bounded right inverses and kernel gluing II: estimates

The diagram (5.27) does not commute, but is very close to commuting for \mathbf{g} close to zero, as the following estimates make precise.

Lemma 5.8. *We have the following estimates*

$$\|\text{PT} \circ D_0 - D_{0|\mathbf{g}} \circ \text{PT}\| \rightarrow 0 \quad (5.32)$$

$$\|(D_{\mathbf{g}} \circ \text{glue})(\xi) - \eta\| = o(1) \cdot \|\xi\| \quad \text{for } \text{break}(\eta) = D_{0|\mathbf{g}}\xi \quad (5.33)$$

$$\|D_{\mathbf{g}} \circ \text{calib} - D_{\mathbf{g}}\| \rightarrow 0 \quad (5.34)$$

as $\mathbf{g} \rightarrow 0$, for any fixed $k \geq 2$.

Proof. To prove (5.32), argue as follows. The first difference between the two operators is over the $[\frac{1}{6}S, \infty) \times S^1$ subset of some ends. In this region, both are linear differential operators, which we may write in local coordinates (s, t) on C_0 and exponential coordinates on the target near the asymptotic orbit or point. The desired bound then follows from the exponential convergence of (5.7) (near $n \in N$, observe that smoothness of u_0 implies decay of all derivatives like $e^{-|s|}$ in cylindrical coordinates). The second difference between the two operators is \hat{J}_0 vs $\hat{J}_{\mathbf{g}_t}$, and this is also bounded as desired, since $\hat{J}_{\mathbf{g}_t} \rightarrow \hat{J}_0$ in C^∞ .

To prove (5.33), argue as follows. Away from ends/necks, the difference is only nonzero due to using $\phi_{\alpha}^{\mathbf{g}, \xi}$ in place of $\phi_{\alpha}^{0, \xi}$, and it is straightforward to see that this is bounded as desired. In the ends/necks, the difference is only nonzero over the $([\frac{1}{3}S - 1, \frac{1}{3}S] \cup [\frac{2}{3}S, \frac{2}{3}S + 1]) \times S^1$ subsets of each neck. By symmetry, we discuss only the $[\frac{2}{3}S, \frac{2}{3}S + 1] \times S^1$ part, where it equals $D_{0|\mathbf{g}}(\chi(s - \frac{2}{3}S)\xi(s, t))$. Now we note that $D_{0|\mathbf{g}}(\chi(s - \frac{2}{3}S)\xi(s, t))$ has $W^{k-1, 2, \delta}(\tilde{C}_0)$ -norm bounded by $\|\xi\|_{k, 2, \delta}$. But we are interested in the $W^{k-1, 2, \delta}(\tilde{C}_{\mathbf{g}})$ -norm, where the weight is smaller by a factor of $e^{-\frac{1}{3}\delta S}$, giving the desired estimate since $\delta > 0$.

To prove (5.34), argue as follows. It suffices to show that $\|D_{\mathbf{g}}(X(\xi))\| = o(1)\|\xi\|$, where $X(\xi) \in \mathbf{X}$ denotes $\xi - \text{calib}(\xi)$, i.e. the second projection in (5.31). Note that the norm of $\xi(q_e'')$ (which determines $X(\xi)$) is bounded by a constant times $e^{-\frac{1}{2}\delta S}\|\xi\|_{k, 2, \delta}$ (i.e. $\|\xi\|_{k, 2, \delta}$ divided by the weight in the middle of the neck). Now $D_{\mathbf{g}}(X(\xi))$ is only nonzero over the $([\frac{1}{3}S - 1, \frac{1}{3}S] \cup [\frac{2}{3}S, \frac{2}{3}S + 1]) \times S^1$ subsets of each neck, where the weight is $e^{\frac{1}{3}\delta S}$. Its norm is thus bounded by $e^{(\frac{1}{3} - \frac{1}{2})\delta S}\|\xi\|_{k, 2, \delta}$, giving the desired result since $\delta > 0$. \square

5.2.7 Bounded right inverses and kernel gluing III: goal

Recall that by assumption, D_0 is surjective and the natural projection $\ker D_0 \rightarrow E_{J \setminus I}$ is surjective; indeed, this is what it means for $\psi_{IJ}(x_0)$ to lie in $\mathcal{M}(T')_I^{\text{reg}}$. Let Q_0 denote any bounded right inverse for D_0 , meaning $D_0 Q_0 = \mathbf{1}$. Then we have a direct sum decomposition

$$W^{k, 2, \delta}(C_0, u_0^* T \hat{X}_0)_{\xi(q_{v,i}) \in T \hat{D}_{v,i}} \oplus \mathcal{J} \oplus E_J = \ker D_0 \oplus \text{im } Q_0. \quad (5.35)$$

In fact, choosing a bounded right inverse Q_0 is equivalent to choosing a closed complement $\text{im } Q_0$ of $\ker D_0$. The classical Banach space implicit function theorem (taking as input Q_0 and the quadratic estimate (5.24)) then implies that the map from $\mathcal{F}_0^{-1}(0)$ to $\ker D_0$ by projection along $\text{im } Q_0$ is a local diffeomorphism near zero.

Our goal is to generalize this setup to \mathbf{g} in a neighborhood of zero (using (5.27) and Lemma 5.8). Namely, we will construct a right inverse $Q_{\mathbf{g}}$ for $D_{\mathbf{g}}$ (equivalently, we will choose a complement $\text{im } Q_{\mathbf{g}}$ for $\ker D_{\mathbf{g}}$), so we that have a direct sum decomposition

$$W^{k, 2, \delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T \hat{X}_{\mathbf{g}})_{\substack{\xi(q_{v,i}) \in T \hat{D}_{v,i} \\ \pi_{\mathbb{R} \partial_s \oplus \mathbb{R} R_\lambda} \xi(q_e'') = 0}} \oplus \mathcal{J} \oplus E_J = \ker D_{\mathbf{g}} \oplus \text{im } Q_{\mathbf{g}}. \quad (5.36)$$

The same implicit function theorem argument applies as long as $\|Q_{\mathbf{g}}\|$ is bounded uniformly for \mathbf{g} near zero. Note also that uniform boundedness of $Q_{\mathbf{g}}$ implies that both projections in (5.36) are uniformly bounded (since they are given by $\mathbf{1} - Q_{\mathbf{g}} D_{\mathbf{g}}$ and $Q_{\mathbf{g}} D_{\mathbf{g}}$ respectively).

Now to ensure that the individual parameterizations of $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ by $K_{\mathbf{g}}$ near zero fit together continuously as \mathbf{g} varies, we also need to show that the direct sum decomposition (5.36) is

“continuous in \mathbf{g} ” in some sense. Let us now describe more precisely the sense we mean. For some points $w_i \in C_0 \setminus (\{p_{v,e}\}_{v,e} \cup N)$, consider the linear functional

$$L_0 : W^{k,2,\delta}(C_0, u_0^* T \hat{X}_0)_{\xi(q_{v,i}) \in T \hat{D}_{v,i}} \oplus \mathcal{J} \oplus E_J \rightarrow \left(\bigoplus_i T_{u_0(w_i)} \hat{X}_0 \oplus \mathcal{J} \oplus E_J \right) / B$$

for some subspace B projecting trivially onto $E_{J \setminus I}$. Fix B so that $L_0|_{\ker D_0}$ is an isomorphism; this is possible since $\ker D_0 \rightarrow E_{J \setminus I}$ is surjective. Since $L_0|_{\ker D_0}$ is an isomorphism, we have a direct sum decomposition

$$W^{k,2,\delta}(C_0, u_0^* T \hat{X}_0)_{\xi(q_{v,i}) \in T \hat{D}_{v,i}} \oplus \mathcal{J} \oplus E_J = \ker D_0 \oplus \ker L_0. \quad (5.37)$$

Now denote by

$$L_{\mathbf{g}} : W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T \hat{X}_{\mathbf{g}})_{\substack{\xi(q_{v,i}) \in T \hat{D}_{v,i} \\ \pi_{\mathbb{R} \partial_s \oplus \mathbb{R} R_\lambda} \xi(q_e'')=0}} \oplus \mathcal{J} \oplus E_J \rightarrow \left(\bigoplus_i T_{u_{\mathbf{g}}(w_i)} \hat{X}_{\mathbf{g}} \oplus \mathcal{J} \oplus E_J \right) / B$$

the “same” linear functional, where $w_i \in C_{\mathbf{g}}$ denote the descents of $w_i \in C_0$, so that there is a natural identification $T_{u_{\mathbf{g}}(w_i)} \hat{X}_{\mathbf{g}} = T_{u_0(w_i)} \hat{X}_0$. We will show that $L_{\mathbf{g}}|_{\ker D_{\mathbf{g}}}$ is still an isomorphism, and hence there is a direct sum decomposition

$$W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T \hat{X}_{\mathbf{g}})_{\substack{\xi(q_{v,i}) \in T \hat{D}_{v,i} \\ \pi_{\mathbb{R} \partial_s \oplus \mathbb{R} R_\lambda} \xi(q_e'')=0}} \oplus \mathcal{J} \oplus E_J = \ker D_{\mathbf{g}} \oplus \ker L_{\mathbf{g}}. \quad (5.38)$$

We will construct $Q_{\mathbf{g}}$ with $\text{im } Q_{\mathbf{g}} = \ker L_{\mathbf{g}}$, i.e. the direct sum decompositions (5.36) and (5.38) coincide. We will also define natural “kernel gluing” isomorphisms $\ker D_0 \xrightarrow{\sim} \ker D_{\mathbf{g}}$ which agree with $L_{\mathbf{g}}^{-1} \circ L_0$.

5.2.8 Bounded right inverses and kernel gluing IV: construction

We now construct the right inverses $Q_{\mathbf{g}}$ and the kernel gluing isomorphisms $\ker D_0 \xrightarrow{\sim} \ker D_{\mathbf{g}}$ satisfying the desired properties discussed above.

We first recall the following general construction, which allows one to upgrade an “approximate right inverse” into a (true) right inverse.

Definition 5.9. Let $D : X \rightarrow Y$ be a bounded linear map between Banach spaces, and let $T : Y \rightarrow X$ be an *approximate right inverse*, meaning that $\|\mathbf{1} - DT\| < 1$. Then there is a (necessarily unique) associated right inverse $Q : Y \rightarrow X$ with the same image $\text{im } Q = \text{im } T$, namely $Q := T(DT)^{-1}$, where $DT : Y \rightarrow Y$ is invertible by the geometric series $\sum_{k=0}^{\infty} (\mathbf{1} - DT)^k$. Moreover, we have (trivially) that $\|Q\| \leq \|T\| (1 - \|\mathbf{1} - DT\|)^{-1}$.

To define the right inverse $Q_{\mathbf{g}}$, first define an approximate right inverse $T_{\mathbf{g}}$ of $D_{\mathbf{g}}$ as the following composition of maps in (5.27):

$$T_{\mathbf{g}} := \text{calib} \circ \text{glue} \circ \text{PT} \circ Q_0 \circ \text{PT} \circ \text{id}^{1,0} \circ \text{break}, \quad (5.39)$$

where Q_0 denotes the fixed right inverse of D_0 defined by the property that $\text{im } Q_0 = \ker L_0$. A consequence of the estimates (5.32)–(5.34) (expressing the fact that (5.27) almost commutes)

is that $\|\mathbf{1} - D_{\mathbf{g}}T_{\mathbf{g}}\| \rightarrow 0$ as $\mathbf{g} \rightarrow 0$ (see [Par16, Lemma B.7.6]). Let $Q_{\mathbf{g}}$ denote the associated true right inverse, which is uniformly bounded for \mathbf{g} near zero (since all the maps in (5.27) are uniformly bounded). Note that $L_{\mathbf{g}} \circ \text{calib} \circ \text{glue} \circ \text{PT} = L_0$ by inspection, so $\text{im } Q_0 = \ker L_0$ implies that $\text{im } Q_{\mathbf{g}} = \text{im } T_{\mathbf{g}} \subseteq \ker L_{\mathbf{g}}$.

We define the kernel gluing isomorphism $\ker D_0 \xrightarrow{\sim} \ker D_{\mathbf{g}}$ as the composition

$$(\mathbf{1} - Q_{\mathbf{g}}D_{\mathbf{g}}) \circ \text{calib} \circ \text{glue} \circ \text{PT} : \ker D_0 \rightarrow \ker D_{\mathbf{g}}. \quad (5.40)$$

Note that $L_{\mathbf{g}} \circ (\mathbf{1} - Q_{\mathbf{g}}D_{\mathbf{g}}) \circ \text{calib} \circ \text{glue} \circ \text{PT} = L_0$ by inspection, and hence (5.40) is injective. Now we have by definition that $\text{ind } D_0 = \mu(T') - \#V_s(T') - \#N + \dim E_J$ and $\text{ind } D_{\mathbf{g}} = \mu(T'_{\mathbf{g}}) - \#V_s(T'_{\mathbf{g}}) - \#N + \dim E_J$, where $T' \rightarrow T'_{\mathbf{g}}$ denotes the image of \mathbf{g} under the map $G_{T'/T} \rightarrow \mathcal{S}_{T'/T}$. These indices coincide as remarked below Definition 2.42, so (5.40) is an isomorphism as both D_0 and $D_{\mathbf{g}}$ are surjective. Since (5.40) is an isomorphism, so is $L_{\mathbf{g}}|_{\ker D_{\mathbf{g}}}$, and it thus follows that the inclusion $\text{im } Q_{\mathbf{g}} \subseteq \ker L_{\mathbf{g}}$ is in fact an equality $\text{im } Q_{\mathbf{g}} = \ker L_{\mathbf{g}}$.

Remark 5.10. It is possible to prove that (5.40) is surjective directly at the cost of proving a few more estimates (this is similar to [FH93, Proposition 9]). This thus gives an *a priori* proof that $\mu(T) = \mu(T')$ for $T' \rightarrow T$.

Let us sketch the argument. Let $\ell \in \ker D_{\mathbf{g}}$. The *a priori* estimate of exponential convergence to a trivial cylinder due to Hofer–Wysocki–Zehnder [HWZ02, Theorems 1.1, 1.2, and 1.3] (restated here as Proposition 2.20) has an easier linear analogue which says that in any neck, ℓ decays rapidly to a constant vector field tangent to the trivial cylinder; moreover, this constant vanishes in Reeb necks since $\pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda} \ell(q_e'') = 0$. It follows that we can apply an “ungluing” operation to produce a κ of commensurable norm $\|\kappa\| \asymp \|\ell\|$ with $\|D_0\kappa\| = o(1) \cdot \|\kappa\|$ and $L_0\kappa = L_{\mathbf{g}}\ell$. Now we have $\|\ell - (\text{calib} \circ \text{glue} \circ \text{PT})(\kappa)\| = o(1) \cdot \|\ell\|$ by explicit calculation, and it follows that the image of $(\mathbf{1} - Q_0D_0)\kappa \in \ker D_0$ under (5.40) is within distance $o(1) \cdot \|\ell\|$ of ℓ . Since this holds for all $\ell \in \ker D_{\mathbf{g}}$, it follows that (5.40) is surjective.

5.3 Gluing map

We now define the gluing map and show that it is a local homeomorphism. This is the “endgame” of the gluing argument, where we deduce the desired results from the technical work performed above.

5.3.1 Definition of the gluing map

We first recall (following our sketch in §5.2.7) how our work above implies that $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is a manifold near zero and that projection along $\text{im } Q_{\mathbf{g}}$ provides a diffeomorphism between it and $\ker D_{\mathbf{g}}$ near zero.

We have fixed a right inverse $Q_{\mathbf{g}}$ for $D_{\mathbf{g}} = \mathcal{F}'_{\mathbf{g}}(0, \cdot)$ with $\|Q_{\mathbf{g}}\|$ bounded uniformly for \mathbf{g} near zero. Now it follows from (5.24) that $Q_{\mathbf{g}}$ is an approximate right inverse to $\mathcal{F}'_{\mathbf{g}}$ over the ball of some radius $c'_{k,\delta} > 0$ (uniform in \mathbf{g}). Hence over this ball of radius $c'_{k,\delta} > 0$, the operator $\mathcal{F}'_{\mathbf{g}}$ is surjective, i.e. $\mathcal{F}_{\mathbf{g}}$ is transverse to zero. By the Banach space implicit function theorem, it thus follows that $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is a C^ℓ -submanifold (for $k \geq \ell + 2$) which is transverse to $\text{im } Q_{\mathbf{g}}$.

Let us now show that map $\ker D_{\mathbf{g}} \rightarrow \mathcal{F}_{\mathbf{g}}^{-1}(0)$ given by projection along $\text{im } Q_{\mathbf{g}}$ (is well-defined and) is a diffeomorphism near zero. The key point is that the map $\mathbf{1} - Q_{\mathbf{g}}\mathcal{F}_{\mathbf{g}}$ is a contraction mapping when restricted to any slice $(\xi + \text{im } Q_{\mathbf{g}}) \cap B(c'_{k,\delta})$ with ξ sufficiently small in terms of $c'_{k,\delta}$. This follows from (5.26) and (5.23), which moreover imply that the contraction constant approaches zero (uniformly in \mathbf{g}) as $c'_{k,\delta} \rightarrow 0$. This gives the desired result, and moreover shows that the projection along $\text{im } Q_{\mathbf{g}}$ to $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is given (over the whole ball $B(c'_{k,\delta})$) by the limit of the Newton–Picard iteration $\xi \mapsto \xi - Q_{\mathbf{g}}\mathcal{F}_{\mathbf{g}}\xi$.

We can now define the gluing map, by precomposing the above local diffeomorphisms $\ker D_{\mathbf{g}} \rightarrow \mathcal{F}_{\mathbf{g}}^{-1}(0)$ with the kernel gluing isomorphisms $\ker D_0 \xrightarrow{\sim} \ker D_{\mathbf{g}}$ and letting \mathbf{g} vary. In other words, the gluing map

$$\left(G_{T'/T} \times \mathbb{C}^N \times \ker D_0, (0, 0) \right) \rightarrow \left(\overline{\mathcal{M}}(T)_J, x_0 \right) \quad (5.41)$$

sends (\mathbf{g}, κ) to the map $\exp_{u_{\mathbf{g}}} \kappa_{\mathbf{g}}^{\infty} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$, where $\kappa_{\mathbf{g}}^{\infty} \in \mathcal{F}_{\mathbf{g}}^{-1}(0)$ is the unique intersection point $\mathcal{F}_{\mathbf{g}}^{-1}(0) \cap (((\mathbf{1} - Q_{\mathbf{g}}D_{\mathbf{g}}) \circ \text{calib} \circ \text{glue} \circ \text{PT})(\kappa) + \text{im } Q_{\mathbf{g}})$ in $B(c'_{k,\delta})$. The discrete data for $\exp_{u_{\mathbf{g}}} \xi : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$ is naturally inherited from that for $u_0 : C_0 \rightarrow \hat{X}_0$. Since $\mathcal{F}_{\mathbf{g}}$ is transverse to zero at $\kappa_{\mathbf{g}}^{\infty}$, it follows that the image of the gluing map is contained in $\overline{\mathcal{M}}(T)_J^{\text{reg}}$.

The gluing map evidently commutes with the maps from both sides to $\mathcal{S}_{T'/T} \times \overline{\mathfrak{s}}(T) \times E_{J \setminus I}$ (recall that $\text{im } Q_{\mathbf{g}} = \ker L_{\mathbf{g}}$ projects trivially onto $E_{J \setminus I}$ by definition).

Let us also note here that the inequality $\mu(T') - \#V_s(T') - 2\#N + \dim E_I \geq 0$ follows from the fact that D_0 is surjective, $\ker D_0 \rightarrow E_{J \setminus I}$ is surjective, and $\text{ind } D_0 = \mu(T') - \#V_s(T') - 2\#N + \dim E_J$.

5.3.2 Properties of the gluing map

We now show that the gluing map is a local homeomorphism. More precisely, we show that the gluing map is continuous and that it is a local bijection, for which the *a priori* estimates recalled in §2.8 will be crucial. We then appeal to some point set topology to see that the gluing map is a local homeomorphism (though continuity of the inverse could also be proven directly).

Lemma 5.11. *The gluing map is continuous.*

Proof. Recall that $\kappa_{\mathbf{g}}^{\infty}$ may be described via the Newton–Picard iteration as follows. Namely, $\kappa_{\mathbf{g}}^{\infty} = \lim_{i \rightarrow \infty} \kappa_{\mathbf{g}}^i$, where

$$\kappa_{\mathbf{g}}^{i+1} = \kappa_{\mathbf{g}}^i - Q_{\mathbf{g}}\mathcal{F}_{\mathbf{g}}\kappa_{\mathbf{g}}^i, \quad (5.42)$$

$$\kappa_{\mathbf{g}}^0 = (\text{calib} \circ \text{glue} \circ \text{PT})(\kappa). \quad (5.43)$$

Note that there is no $\mathbf{1} - Q_{\mathbf{g}}D_{\mathbf{g}}$ in the definition of $\kappa_{\mathbf{g}}^0$ (this is ok since $Q_{\mathbf{g}}D_{\mathbf{g}}\kappa_{\mathbf{g}}^0 \in \text{im } Q_{\mathbf{g}}$).

Now suppose $(\mathbf{g}_i, \kappa_i) \rightarrow (\mathbf{g}, \kappa)$ (a convergent net), and let us show that the net $\exp_{u_{\mathbf{g}_i}} (\kappa_i)_{\mathbf{g}_i}^{\infty} : C_{\mathbf{g}_i} \rightarrow \hat{X}_{\mathbf{g}_i}$ approaches $\exp_{u_{\mathbf{g}}} \kappa_{\mathbf{g}}^{\infty} : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}$ in the Gromov topology.

First, we claim that $\|(\kappa_i)_{\mathbf{g}_i}^{\infty} - \kappa_{\mathbf{g}_i}^{\infty}\|_{k,2,\delta} \rightarrow 0$. By uniform convergence of the Newton–Picard iteration, it suffices to show that $\|(\kappa_i)_{\mathbf{g}_i}^n - \kappa_{\mathbf{g}_i}^n\|_{k,2,\delta} \rightarrow 0$ for all n . The case $n = 0$ follows from uniform boundedness of $\text{calib} \circ \text{glue} \circ \text{PT}$. The desired claim then follows by

induction on n using (5.24). Now the claim implies that $\|(\kappa_i)_{\mathbf{g}_i}^\infty - \kappa_{\mathbf{g}_i}^\infty\|_\infty \rightarrow 0$, and thus it suffices to show that

$$\exp_{u_{\mathbf{g}_i}} \kappa_{\mathbf{g}_i}^\infty : C_{\mathbf{g}_i} \rightarrow \hat{X}_{\mathbf{g}_i} \text{ approaches } \exp_{u_{\mathbf{g}}} \kappa_{\mathbf{g}}^\infty : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}. \quad (5.44)$$

Define $(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0$ by (as the notation suggests) pregluing $\kappa_{\mathbf{g}}^\infty$ from $C_{\mathbf{g}}$ to $C_{\mathbf{g}_i}$ as follows. In any neck of $C_{\mathbf{g}_i}$ corresponding to a pair of ends of $C_{\mathbf{g}}$, we preglue via $\text{calib} \circ \text{glue} \circ \text{PT}$ as before (this operation is local to the ends/neck). In any neck of $C_{\mathbf{g}_i}$ corresponding to a neck of $C_{\mathbf{g}}$, we simply use parallel transport and a nice diffeomorphism between the two necks (say, converging to the identity map in the C^∞ topology as $\mathbf{g}_i \rightarrow \mathbf{g}$). We may assume without loss of generality that there are no pairs of ends of $C_{\mathbf{g}_i}$ corresponding to a neck of $C_{\mathbf{g}}$.

Now we claim that

$$\|\mathcal{F}_{\mathbf{g}_i}((\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0)\|_{k-1,2,\delta} \rightarrow 0. \quad (5.45)$$

Away from the necks/ends, the 1-form part of $\mathcal{F}_{\mathbf{g}_i}((\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0)$ is nonzero only because of using $\hat{J}_{(\mathbf{g}_i)_t}$ in place of $\hat{J}_{\mathbf{g}_i}$ and because of using $\phi_\alpha^{\mathbf{g}_i, \xi}$ in place of $\phi_\alpha^{\mathbf{g}, \xi}$. We have $\hat{J}_{(\mathbf{g}_i)_t} \rightarrow \hat{J}_{\mathbf{g}_i}$ and $\phi_\alpha^{\mathbf{g}_i, \xi} \rightarrow \phi_\alpha^{\mathbf{g}, \xi}$, so the desired estimate follows since $(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0 = \kappa_{\mathbf{g}}^\infty$ away from the ends/necks. Over the Reeb ends of $C_{\mathbf{g}_i}$, the 1-form part vanishes. Over the Reeb necks of $C_{\mathbf{g}_i}$ corresponding to necks of $C_{\mathbf{g}}$, the 1-form part approaches zero. Over the Reeb necks of $C_{\mathbf{g}_i}$ corresponding to pairs of ends of $C_{\mathbf{g}}$, the estimate follows from the exponential decay of $\kappa_{\mathbf{g}}^\infty$ and u_0 and the fact that $\delta < 1$ and $\delta < \delta_\gamma$ for all asymptotic Reeb orbits γ . The same applies to nodal necks, in addition considering the convergence $\hat{J}_{(\mathbf{g}_i)_t} \rightarrow \hat{J}_{\mathbf{g}_t}$. The $\mathbb{R}^{V_s(T')}$ part clearly approaches zero. This proves (5.45).

Now we consider the Newton–Picard iteration starting at $(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0$, with limit $(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^\infty \in \mathcal{F}_{\mathbf{g}_i}^{-1}(0)$. By uniform contraction of the iteration and (5.45), we conclude that $\|(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0 - (\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^\infty\|_{k,2,\delta} \rightarrow 0$. It thus follows that

$$\exp_{u_{\mathbf{g}_i}} (\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^\infty : C_{\mathbf{g}_i} \rightarrow \hat{X}_{\mathbf{g}_i} \text{ approaches } \exp_{u_{\mathbf{g}}} \kappa_{\mathbf{g}}^\infty : C_{\mathbf{g}} \rightarrow \hat{X}_{\mathbf{g}}. \quad (5.46)$$

Now we claim that $(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^\infty = \kappa_{\mathbf{g}_i}^\infty$, which is clearly sufficient to conclude the proof.

By construction, we have $L_{\mathbf{g}_i}((\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^\infty) = L_{\mathbf{g}_i}((\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^0) = L_{\mathbf{g}}(\kappa_{\mathbf{g}}^\infty) = L_{\mathbf{g}}(\kappa_{\mathbf{g}}^0) = L_0(\kappa)$ and similarly $L_{\mathbf{g}_i}(\kappa_{\mathbf{g}_i}^\infty) = L_{\mathbf{g}_i}(\kappa_{\mathbf{g}_i}^0) = L_0(\kappa)$. Thus $(\kappa_{\mathbf{g}}^\infty)_{\mathbf{g}_i}^\infty$ and $\kappa_{\mathbf{g}_i}^\infty$ differ by an element of $\ker L_{\mathbf{g}_i} = \text{im } Q_{\mathbf{g}_i}$, which is enough. \square

Lemma 5.12. *The gluing map (5.41) is a local bijection. That is, for every sufficiently small neighborhood U of the basepoint in the domain, there exists an open neighborhood V of the basepoint in the target such that every $v \in V$ has a unique inverse image $u \in U$.*

Proof. Let $x \in \overline{\mathcal{M}}(T)_J$, and denote the corresponding map by $u : C \rightarrow \hat{X}$. We assume that x is sufficiently close to x_0 , and we will show that x has a unique inverse image under the gluing map (5.41).

Concretely, x close to x_0 in the Gromov topology means the following. We may identify C with (C_β, j_w) , for arbitrarily small gluing parameters $\beta \in \mathbb{C}^{E^{\text{int}}(T')} \times \mathbb{C}^N$ and an almost complex structure j_w which agrees with j_0 except over a compact set (away from the ends/necks) where it is arbitrarily C^∞ close to j_0 . Furthermore, the map $u : (C_\beta, j_w) \rightarrow \hat{X}$ is arbitrarily C^0 -close to u_0 away from the Reeb ends/necks (which by Lemma 2.18 implies arbitrarily C^∞

close, at least away from the nodes of C_0 which are resolved in C ; the behavior of u in these neighborhoods is controlled by Lemma 2.19). We will discuss the situation over the Reeb ends/necks shortly.

Now observe that there are unique points $q_{v,i} \in C$ close to $q_{v,i} \in C_0$ where u intersects $D_{v,i}$ transversally. These points give rise to unique points $q'_v \in C$ according to the sections chosen in §5.1.2. Now regarding $u(q'_v) \in \hat{X}$ as lying on the “zero section” determines (uniquely) a gluing parameter $\mathbf{g} \in G_{T'/T}$ and an isomorphism $\hat{X} = \hat{X}_{\mathbf{g}}$. Thus x corresponds to a map $u : (C_\beta, j_w) \rightarrow \hat{X}_{\mathbf{g}}$.

Now in any Reeb end $[0, \infty) \times S^1 \subseteq C_\beta$ or Reeb neck $[0, S] \times S^1 \subseteq C_\beta$, we apply [HWZ02, Theorems 1.1, 1.2, and 1.3] (restated here as Proposition 2.20) to see that u decays exponentially to a trivial cylinder. For unglued edges, we combine this with the fact that u is C^∞ -close to u_0 over an arbitrarily compact set (which we can choose to go very deep into the end) to conclude that the tangent space marking of C_β at $\{p_{v,e}\}$ induced by the cylindrical coordinates on $\hat{X}_{\mathbf{g}}$ is arbitrarily close to the tangent space marking descended from C_0 . The same reasoning implies that the gluing parameters $\beta \in \mathbb{C}^{E^{\text{int}}(T')}$ (for the glued edges) are given by $L_e^{-1}g_e + o(1)$ and $\theta_e + o(1)$ (i.e. are very close to those coming from \mathbf{g}), and we recover the point $q''_e \in C_\beta$ as the inverse image of $\exp_{u_0|_{\mathbf{g}}(q''_e)}(\ker \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda})$. From these estimates, we conclude that x corresponds to a map $u : (C_{\mathbf{g}}, j_y) \rightarrow \hat{X}_{\mathbf{g}}$ for $\mathbf{g} \in G_{T'/T} \times \mathbb{C}^N$ which is arbitrarily C^∞ -close to $u_{\mathbf{g}}$ (including over the ends/necks, with respect to the cylindrical coordinates) respecting the tangent markings at $p_{v,e}$ and sending q''_e to $\exp_{u_0|_{\mathbf{g}}(q''_e)}(\ker \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda})$. Here \mathbf{g} and y can be assumed arbitrarily small, and we use the fact that (5.10) is a local diffeomorphism uniformly in $\mathbf{g} \in G_{T'/T}$.

The injectivity radius of our fixed exponential map on $\hat{X}_{\mathbf{g}}$ is bounded below uniformly, so x corresponds uniquely to some pair $(u = \exp_{u_{\mathbf{g}}} \xi : (C_{\mathbf{g}}, j_y) \rightarrow \hat{X}_{\mathbf{g}}, e)$, where

$$(\xi, y, e - e_0) \in W^{k,2,\delta}(C_{\mathbf{g}}, u_{\mathbf{g}}^* T\hat{X}_{\mathbf{g}})_{\substack{\xi(q_{v,i}) \in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda} \xi(q''_e) = 0}} \oplus \mathcal{J} \oplus E_J \quad (5.47)$$

has arbitrarily small norm due to the *a priori* estimates on u from Lemmas 2.18 and 2.19 and Proposition 2.20. Now we observed in the gluing construction that for any fixed sufficiently small \mathbf{g} , the gluing map gives a local diffeomorphism between $\ker D_0$ and $\mathcal{F}_{\mathbf{g}}^{-1}(0)$, over a ball of size uniformly bounded below. Thus our map is uniquely in the image of the gluing map. \square

Since the target of the gluing map is Hausdorff and the domain locally compact Hausdorff, it follows from continuity (Lemma 5.11) and bijectivity in a small neighborhood (Lemma 5.12) that the gluing map (5.41) is in fact a local homeomorphism, thus completing the proof of Theorem 3.31. \square

5.4 Orientations

We now prove the compatibility of the geometric and analytic maps on orientation lines, namely that (3.25) commutes. We rely heavily on the gluing construction in §§5.1–5.3.

Proof of Theorem 3.32. The gluing map (5.41) (in the case $I = J$) allows us to describe the left vertical “geometric” map in (3.25) (i.e. the map induced by the topological structure of

$\overline{\mathcal{M}}(T)_I^{\text{reg}}$) near the basepoint x_0 as follows. Note that we may assume (for convenience) that $N = \emptyset$, since this locus is dense. Recall that there is a canonical identification

$$\mathfrak{o}_{\ker D_0} = \mathfrak{o}_{T'}^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I}. \quad (5.48)$$

Now consider sufficiently small (\mathbf{g}, κ) , where $\mathbf{g} \in G_{T'/T}$ lies in the top stratum ($T' \rightarrow T$). Now $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is a submanifold with tangent space $\ker \mathcal{F}'_{\mathbf{g}}(\kappa_{\mathbf{g}}^\infty, \cdot)$, and there is a canonical identification

$$\mathfrak{o}_{\ker \mathcal{F}'_{\mathbf{g}}(\kappa_{\mathbf{g}}^\infty, \cdot)} = \mathfrak{o}_T^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I}. \quad (5.49)$$

The gluing map is differentiable with respect to κ since $\mathcal{F}_{\mathbf{g}}^{-1}(0)$ is a submanifold transverse to $\text{im } Q_{\mathbf{g}}$. Its derivative is clearly given by the composition of $\ker D_0 \rightarrow \ker D_{\mathbf{g}}$ and the map $\ker D_{\mathbf{g}} \rightarrow \ker \mathcal{F}'_{\mathbf{g}}(\kappa_{\mathbf{g}}^\infty, \cdot)$ given by projecting off $\text{im } Q_{\mathbf{g}}$. This map $\ker D_0 \rightarrow \ker \mathcal{F}'_{\mathbf{g}}(\kappa_{\mathbf{g}}^\infty, \cdot)$ thus gives the “geometric” map $\mathfrak{o}_{T'}^\circ \rightarrow \mathfrak{o}_T^\circ$ when combined with the isomorphisms above.

Now the right inverse $Q_{\mathbf{g}}$ to $D_{\mathbf{g}} = \mathcal{F}'_{\mathbf{g}}(0, \cdot)$ is an approximate right inverse to $\mathcal{F}'_{\mathbf{g}}(\xi, \cdot)$ for all $\xi \in B(c'_{k,\delta})$ by (5.24). Hence the kernel $\ker \mathcal{F}'_{\mathbf{g}}(\xi, \cdot)$ forms a vector bundle over $B(c'_{k,\delta})$ which is canonically oriented by $\mathfrak{o}_T^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I}$, and the map $\ker D_0 \rightarrow \ker \mathcal{F}'_{\mathbf{g}}(\xi, \cdot)$ makes sense for all such ξ . Thus the geometric map on orientations is also given by the simpler map at $\xi = 0$

$$\ker D_0 \xrightarrow{(1-Q_{\mathbf{g}}D_{\mathbf{g}}) \circ \text{calib} \circ \text{glue} \circ \text{PT}} \ker D_{\mathbf{g}} \quad (5.50)$$

combined with the canonical identifications

$$\mathfrak{o}_{\ker D_0} = \mathfrak{o}_{T'}^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I}, \quad (5.51)$$

$$\mathfrak{o}_{\ker D_{\mathbf{g}}} = \mathfrak{o}_T^\circ \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I}. \quad (5.52)$$

Now this map (5.50) is precisely the sort of kernel pregluing map which defines the “analytic” map on orientations.

Strictly speaking, the analytic map $\mathfrak{o}_{T'}^\circ \rightarrow \mathfrak{o}_T^\circ$ is defined using a slightly different linearized operator (no \mathcal{J} , E_I , or point conditions), but this is only a “finite-dimensional” difference (note also that \mathcal{J} is canonically oriented since it is a complex vector space). It is thus straightforward to relate them and see that they give rise to the same analytic map on orientations. \square

References

- [Abo12] Mohammed Abouzaid. Framed bordism and Lagrangian embeddings of exotic spheres. *Ann. of Math. (2)*, 175(1):71–185, 2012. 78
- [Ada75] Robert A. Adams. *Sobolev spaces*. Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London, 1975. Pure and Applied Mathematics, Vol. 65. 35, 36, 54, 56, 84, 85
- [AM12] Miguel Abreu and Leonardo Macarini. Contact homology of good toric contact manifolds. *Compos. Math.*, 148(1):304–334, 2012. 12

- [BEHWZ03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder. Compactness results in symplectic field theory. *Geom. Topol.*, 7:799–888, 2003. 1, 2, 4, 5, 6, 13, 32, 33, 34
- [BEM15] Matthew Strom Borman, Yakov Eliashberg, and Emmy Murphy. Existence and classification of overtwisted contact structures in all dimensions. *Acta Math.*, 215(2):281–361, 2015. 12
- [BH16] Erkao Bao and Ko Honda. Semi-global Kuranishi charts and the definition of contact homology. arXiv:1512.00580v3:1–90, 2016. 2
- [BH18] Erkao Bao and Ko Honda. Definition of cylindrical contact homology in dimension three. *J. Topol.*, 11(4):1002–1053, 2018. 2
- [BM04] Frédéric Bourgeois and Klaus Mohnke. Coherent orientations in symplectic field theory. *Math. Z.*, 248(1):123–146, 2004. 40, 41, 43, 44
- [Bou02] Frederic Bourgeois. *A Morse-Bott approach to contact homology*. ProQuest LLC, Ann Arbor, MI, 2002. Thesis (Ph.D.)–Stanford University. 15
- [Bou03] Frédéric Bourgeois. A Morse-Bott approach to contact homology. In *Symplectic and contact topology: interactions and perspectives (Toronto, ON/Montreal, QC, 2001)*, volume 35 of *Fields Inst. Commun.*, pages 55–77. Amer. Math. Soc., Providence, RI, 2003. 15
- [Bou06] Frédéric Bourgeois. Contact homology and homotopy groups of the space of contact structures. *Math. Res. Lett.*, 13(1):71–85, 2006. 14
- [Bou09] Frédéric Bourgeois. A survey of contact homology. In *New perspectives and challenges in symplectic field theory*, volume 49 of *CRM Proc. Lecture Notes*, pages 45–71. Amer. Math. Soc., Providence, RI, 2009. 13
- [BvK10] Frédéric Bourgeois and Otto van Koert. Contact homology of left-handed stabilizations and plumbing of open books. *Commun. Contemp. Math.*, 12(2):223–263, 2010. 12
- [CE12] Kai Cieliebak and Yakov Eliashberg. *From Stein to Weinstein and back*, volume 59 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2012. Symplectic geometry of affine complex manifolds. 13
- [CL09] Kai Cieliebak and Janko Latschev. The role of string topology in symplectic field theory. In *New perspectives and challenges in symplectic field theory*, volume 49 of *CRM Proc. Lecture Notes*, pages 113–146. Amer. Math. Soc., Providence, RI, 2009. 15
- [CMP19] Roger Casals, Emmy Murphy, and Francisco Presas. Geometric criteria for overtwistedness. *J. Amer. Math. Soc.*, 32(2):563–604, 2019. 12

- [DM69] P. Deligne and D. Mumford. The irreducibility of the space of curves of given genus. *Inst. Hautes Études Sci. Publ. Math.*, (36):75–109, 1969. 27
- [EES05] Tobias Ekholm, John Etnyre, and Michael Sullivan. The contact homology of Legendrian submanifolds in \mathbb{R}^{2n+1} . *J. Differential Geom.*, 71(2):177–305, 2005. 2
- [EGH00] Y. Eliashberg, A. Givental, and H. Hofer. Introduction to symplectic field theory. *Geom. Funct. Anal.*, (Special Volume, Part II):560–673, 2000. GAFA 2000 (Tel Aviv, 1999). 1, 2, 3, 12, 15, 40, 41, 43
- [EHK16] Tobias Ekholm, Ko Honda, and Tamás Kálmán. Legendrian knots and exact Lagrangian cobordisms. *J. Eur. Math. Soc. (JEMS)*, 18(11):2627–2689, 2016. 15
- [EKP06] Yakov Eliashberg, Sang Seon Kim, and Leonid Polterovich. Geometry of contact transformations and domains: orderability versus squeezing. *Geom. Topol.*, 10:1635–1747, 2006. 2
- [Eli89] Y. Eliashberg. Classification of overtwisted contact structures on 3-manifolds. *Invent. Math.*, 98(3):623–637, 1989. 12
- [Eli98] Yakov Eliashberg. Invariants in contact topology. In *Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998)*, number Extra Vol. II, pages 327–338, 1998. 1, 3, 12
- [Eli07] Yakov Eliashberg. Symplectic field theory and its applications. In *International Congress of Mathematicians. Vol. I*, pages 217–246. Eur. Math. Soc., Zürich, 2007. 3
- [EO17] Tobias Ekholm and Alexandru Oancea. Symplectic and contact differential graded algebras. *Geom. Topol.*, 21(4):2161–2230, 2017. 15
- [FH93] A. Floer and H. Hofer. Coherent orientations for periodic orbit problems in symplectic geometry. *Math. Z.*, 212(1):13–38, 1993. 41, 43, 91
- [FH18a] Joel W. Fish and Helmut Hofer. Lectures on Polyfolds and Symplectic Field Theory. arXiv:1808.07147v1:1–134, 2018. 2
- [FH18b] Joel W. Fish and Helmut Hofer. Polyfold and SFT Notes I: A Primer on Polyfolds and Construction Tools. arXiv:1806.07025v2:1–68, 2018. 2
- [FH18c] Joel W. Fish and Helmut Hofer. Polyfold and SFT Notes II: Local-Local M-Polyfold Constructions. arXiv:1808.04939v1:1–123, 2018. 2
- [FO99] Kenji Fukaya and Kaoru Ono. Arnold conjecture and Gromov-Witten invariant. *Topology*, 38(5):933–1048, 1999. 2

- [FOOO09a] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian intersection Floer theory: anomaly and obstruction. Part I*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI, 2009. 2
- [FOOO09b] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian intersection Floer theory: anomaly and obstruction. Part II*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI, 2009. 2
- [FOOO12] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Technical details on Kuranishi structure and virtual fundamental chain. arXiv:1209.4410v1:1–257, 2012. 2
- [FOOO15] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Kuranishi structure, Pseudo-holomorphic curve, and Virtual fundamental chain: Part 1. arXiv:1503.07631v1:1–203, 2015. 2
- [FOOO16a] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Exponential decay estimates and smoothness of the moduli space of pseudoholomorphic curves. arXiv:1603.07026v1:1–111, 2016. 2
- [FOOO16b] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Shrinking good coordinate systems associated to Kuranishi structures. *J. Symplectic Geom.*, 14(4):1295–1310, 2016. 2
- [FOOO17] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Kuranishi structure, Pseudo-holomorphic curve, and virtual fundamental chain: Part 2. arXiv:1704.01848v1:1–277, 2017. 2
- [FOOO18a] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Construction of Kuranishi structures on the moduli spaces of pseudo holomorphic disks: I. In *Surveys in differential geometry 2017. Celebrating the 50th anniversary of the Journal of Differential Geometry*, volume 22 of *Surv. Differ. Geom.*, pages 133–190. Int. Press, Somerville, MA, 2018. 2
- [FOOO18b] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. Construction of Kuranishi structures on the moduli spaces of pseudo holomorphic disks: II. arXiv:1808.06106v1:1–47, 2018. 2
- [Fuk17] Kenji Fukaya. Lie groupoid, deformation of unstable curve, and construction of equivariant Kuranishi charts. arXiv:1701.02840v1:1–86, 2017. 2
- [GM15] W. D. Gillam and Samouil Molcho. Log differentiable spaces and manifolds with corners. arXiv:1507.06752v1:1–128, 2015. 25
- [Gro85] M. Gromov. Pseudoholomorphic curves in symplectic manifolds. *Invent. Math.*, 82(2):307–347, 1985. 1, 28

- [HN16] Michael Hutchings and Jo Nelson. Cylindrical contact homology for dynamically convex contact forms in three dimensions. *J. Symplectic Geom.*, 14(4):983–1012, 2016. 2
- [Hof93] H. Hofer. Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three. *Invent. Math.*, 114(3):515–563, 1993. 1, 32
- [Hör69] Lars Hörmander. *Linear partial differential operators*. Third revised printing. Die Grundlehren der mathematischen Wissenschaften, Band 116. Springer-Verlag New York Inc., New York, 1969. 40
- [HT07] Michael Hutchings and Clifford Henry Taubes. Gluing pseudoholomorphic curves along branched covered cylinders. I. *J. Symplectic Geom.*, 5(1):43–137, 2007. 2, 78
- [HT09] Michael Hutchings and Clifford Henry Taubes. Gluing pseudoholomorphic curves along branched covered cylinders. II. *J. Symplectic Geom.*, 7(1):29–133, 2009. 2, 78
- [Hut02] Michael Hutchings. An index inequality for embedded pseudoholomorphic curves in symplectizations. *J. Eur. Math. Soc. (JEMS)*, 4(4):313–361, 2002. 2
- [Hut09] Michael Hutchings. The embedded contact homology index revisited. In *New perspectives and challenges in symplectic field theory*, volume 49 of *CRM Proc. Lecture Notes*, pages 263–297. Amer. Math. Soc., Providence, RI, 2009. 2
- [HWZ95] H. Hofer, K. Wysocki, and E. Zehnder. Properties of pseudo-holomorphic curves in symplectisations. II. Embedding controls and algebraic invariants. *Geom. Funct. Anal.*, 5(2):270–328, 1995. 2, 55
- [HWZ96] H. Hofer, K. Wysocki, and E. Zehnder. Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics. *Ann. Inst. H. Poincaré Anal. Non Linéaire*, 13(3):337–379, 1996. 2
- [HWZ98] H. Hofer, K. Wysocki, and E. Zehnder. Correction to: “Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics” [Ann. Inst. H. Poincaré Anal. Non Linéaire **13** (1996), no. 3, 337–379; MR1395676 (97e:58029)]. *Ann. Inst. H. Poincaré Anal. Non Linéaire*, 15(4):535–538, 1998. 2
- [HWZ99] H. Hofer, K. Wysocki, and E. Zehnder. Properties of pseudoholomorphic curves in symplectizations. III. Fredholm theory. In *Topics in nonlinear analysis*, volume 35 of *Progr. Nonlinear Differential Equations Appl.*, pages 381–475. Birkhäuser, Basel, 1999. 2

- [HWZ02] H. Hofer, K. Wysocki, and E. Zehnder. Finite energy cylinders of small area. *Ergodic Theory Dynam. Systems*, 22(5):1451–1486, 2002. 2, 21, 29, 31, 80, 91, 94
- [HWZ07] H. Hofer, K. Wysocki, and E. Zehnder. A general Fredholm theory. I. A splicing-based differential geometry. *J. Eur. Math. Soc. (JEMS)*, 9(4):841–876, 2007. 2
- [HWZ09a] Helmut Hofer, Kris Wysocki, and Eduard Zehnder. A general Fredholm theory. III. Fredholm functors and polyfolds. *Geom. Topol.*, 13(4):2279–2387, 2009. 2
- [HWZ09b] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder. A general Fredholm theory. II. Implicit function theorems. *Geom. Funct. Anal.*, 19(1):206–293, 2009. 2
- [HWZ10a] H. Hofer, K. Wysocki, and E. Zehnder. Integration theory on the zero sets of polyfold Fredholm sections. *Math. Ann.*, 346(1):139–198, 2010. 2
- [HWZ10b] Helmut Hofer, Kris Wysocki, and Eduard Zehnder. sc-smoothness, retractions and new models for smooth spaces. *Discrete Contin. Dyn. Syst.*, 28(2):665–788, 2010. 2
- [HWZ14a] Helmut Hofer, Kris Wysocki, and Eduard Zehnder. Applications of polyfold theory I: The polyfolds of Gromov–Witten theory. arXiv:1107.2097v3:i–vi,1–271, 2014. 2
- [HWZ14b] Helmut Hofer, Kris Wysocki, and Eduard Zehnder. Polyfold and Fredholm Theory I: Basic Theory in M-Polyfolds. arXiv:1407.3185v1:1–246, 2014. 2
- [HWZ17] Helmut Hofer, Kris Wysocki, and Eduard Zehnder. Polyfold and Fredholm Theory. arXiv:1707.08941v1:i–xviii,1–698, 2017. 2
- [Ish18] Suguru Ishikawa. Construction of general symplectic field theory. arXiv:1807.09455v3:1–320, 2018. 2
- [Joy12] Dominic Joyce. *D-manifolds and d-orbifolds: a theory of derived differential geometry*. Book Manuscript, 2012. 2
- [Joy14] Dominic Joyce. An introduction to d-manifolds and derived differential geometry. In *Moduli spaces*, volume 411 of *London Math. Soc. Lecture Note Ser.*, pages 230–281. Cambridge Univ. Press, Cambridge, 2014. 2
- [Joy15] Dominic Joyce. A new definition of Kuranishi space. arXiv:1409.6908v3:1–193, 2015. 2
- [Joy16] Dominic Joyce. A generalization of manifolds with corners. *Adv. Math.*, 299:760–862, 2016. 25
- [KM15] Chris Kottke and Richard B. Melrose. Generalized blow-up of corners and fiber products. *Trans. Amer. Math. Soc.*, 367(1):651–705, 2015. 25

- [LM85] Robert B. Lockhart and Robert C. McOwen. Elliptic differential operators on noncompact manifolds. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)*, 12(3):409–447, 1985. 38
- [LT98a] Jun Li and Gang Tian. Virtual moduli cycles and Gromov-Witten invariants of general symplectic manifolds. In *Topics in symplectic 4-manifolds (Irvine, CA, 1996)*, First Int. Press Lect. Ser., I, pages 47–83. Int. Press, Cambridge, MA, 1998. 2
- [LT98b] Gang Liu and Gang Tian. Floer homology and Arnold conjecture. *J. Differential Geom.*, 49(1):1–74, 1998. 2
- [LW11] Janko Latschev and Chris Wendl. Algebraic torsion in contact manifolds. *Geom. Funct. Anal.*, 21(5):1144–1195, 2011. With an appendix by Michael Hutchings. 13, 15
- [McD19] Dusa McDuff. Constructing the virtual fundamental class of a Kuranishi atlas. *Algebr. Geom. Topol.*, 19(1):151–238, 2019. 2
- [MNW13] Patrick Massot, Klaus Niederkrüger, and Chris Wendl. Weak and strong fillability of higher dimensional contact manifolds. *Invent. Math.*, 192(2):287–373, 2013. 13
- [MS94] Dusa McDuff and Dietmar Salamon. *J-holomorphic curves and quantum cohomology*, volume 6 of *University Lecture Series*. American Mathematical Society, Providence, RI, 1994. 78
- [MS04] Dusa McDuff and Dietmar Salamon. *J-holomorphic curves and symplectic topology*, volume 52 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2004. 28, 55, 78
- [Mul94] Marie-Paule Muller. Gromov’s Schwarz lemma as an estimate of the gradient for holomorphic curves. In *Holomorphic curves in symplectic geometry*, volume 117 of *Progr. Math.*, pages 217–231. Birkhäuser, Basel, 1994. 28
- [MW17a] Dusa McDuff and Katrin Wehrheim. Smooth Kuranishi atlases with isotropy. *Geom. Topol.*, 21(5):2725–2809, 2017. 2
- [MW17b] Dusa McDuff and Katrin Wehrheim. The topology of Kuranishi atlases. *Proc. Lond. Math. Soc. (3)*, 115(2):221–292, 2017. 2
- [MW18] Dusa McDuff and Katrin Wehrheim. The fundamental class of smooth Kuranishi atlases with trivial isotropy. *J. Topol. Anal.*, 10(1):71–243, 2018. 2
- [Nie06] Klaus Niederkrüger. The plastikstufe—a generalization of the overtwisted disk to higher dimensions. *Algebr. Geom. Topol.*, 6:2473–2508, 2006. 13
- [NW11] Klaus Niederkrüger and Chris Wendl. Weak symplectic fillings and holomorphic curves. *Ann. Sci. Éc. Norm. Supér. (4)*, 44(5):801–853, 2011. 13

- [Par16] John Pardon. An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves. *Geom. Topol.*, 20(2):779–1034, 2016. 1, 2, 28, 45, 58, 59, 60, 61, 62, 63, 70, 78, 91
- [Rua99] Yongbin Ruan. Virtual neighborhoods and pseudo-holomorphic curves. In *Proceedings of 6th Gökova Geometry-Topology Conference*, volume 23, pages 161–231, 1999. 2
- [Sal99] Dietmar Salamon. Lectures on Floer homology. In *Symplectic geometry and topology (Park City, UT, 1997)*, volume 7 of *IAS/Park City Math. Ser.*, pages 143–229. Amer. Math. Soc., Providence, RI, 1999. 38
- [Sol06] Jake P. Solomon. *Intersection theory on the moduli space of holomorphic curves with Lagrangian boundary conditions*. ProQuest LLC, Ann Arbor, MI, arXiv:math/0606429, 2006. Thesis (Ph.D.)—Massachusetts Institute of Technology. 41
- [Ust99] Ilya Ustilovsky. Infinitely many contact structures on S^{4m+1} . *Internat. Math. Res. Notices*, (14):781–791, 1999. 12
- [Wen14] Chris Wendl. Lectures on Holomorphic Curves in Symplectic and Contact Geometry. arXiv:1011.1690v2:i-x,1–184, 2014. 28
- [Wen16] Chris Wendl. Lectures on Symplectic Field Theory. arXiv:1612.01009v2:i-x,1–333, 2016. 38
- [Yau06] Mei-Lin Yau. Vanishing of the contact homology of overtwisted contact 3-manifolds. *Bull. Inst. Math. Acad. Sin. (N.S.)*, 1(2):211–229, 2006. With an appendix by Yakov Eliashberg. 12