

PARAMETRIC GROMOV WIDTH OF LIOUVILLE DOMAINS

FILIP BROĆIĆ AND DYLAN CANT

ABSTRACT. The classical Gromov width measures the largest symplectic ball embeddable into a symplectic manifold; inspired by the symplectic camel problem, we generalize this to ask how large a symplectic ball can be embedded as a family over a parameter space N . Given a smooth map $f : N \rightarrow \Omega$, where Ω is a symplectic manifold, we define the *parametric Gromov width* $\text{Gr}(f, \Omega)$ as the supremum of capacities $a > 0$ for which there exists a family of balls, parametrized by N , of capacity a whose centers trace out the map f . For Liouville domains Ω , we establish upper bounds on $\text{Gr}(f, \Omega)$ using the Floer cohomology persistence module associated to Ω . Specializing to fiberwise starshaped domains in the cotangent bundle T^*M , we derive computable bounds via filtered string topology. Specific examples of Ω – including disk cotangent bundles of thin ellipsoids, open books, and tori – demonstrate our bounds, and reveal constraints on parameterized symplectic embeddings beyond the classical Gromov width.

1. Introduction

The *symplectic camel theorem* (see [EG91, §3.4.B] and [Vit92, MT94]) produced the first example of a connected symplectic manifold for which the space of symplectic embeddings of a ball is disconnected. Let us explain the result in a slightly non-standard way: consider $W = \mathbb{R}/\mathbb{Z} \times \mathbb{R}^{2n-1}$, with coordinates $x_1, y_1, x_2, y_2, \dots, x_n, y_n$ and the standard symplectic structure. Define, for $n > 1$, the space:

$$X = (W \setminus \{x_1 = 0\}) \cup \{x_n^2 + y_n^2 \leq \pi^{-1}\epsilon\}.$$

In words, any loop in X with winding number 1 (relative the $\mathbb{R}/\mathbb{Z}$ factor) must pass through the “hole” of capacity ϵ . Then the classical camel theorem can be stated as:

Proposition 1. *If $a > \epsilon$, there does not exist a map $F : \mathbb{R}/\mathbb{Z} \times B(a) \rightarrow X$ such that:*

- (1) $t \mapsto F(t, 0)$ has winding number 1 relative the $\mathbb{R}/\mathbb{Z}$ -factor,
- (2) $z \mapsto F(t, z)$ is a symplectic embedding $B(a) \rightarrow X$ for each t ,

where $B(a)$ is the ball of capacity a .

The results in this paper provide a framework for detecting such phenomena in Liouville domains. Our methods are based on Floer cohomology, and recover Proposition 1.

Let Ω be a symplectic manifold. Denote by $B(a) \subset \mathbb{C}^n$ the compact ball of symplectic capacity a , and denote by $\mathfrak{B}(a, \Omega)$ the space of symplectic embeddings $B(a) \rightarrow \Omega$.

A smooth map $f : N \rightarrow \Omega$ is said to *lift* to $\mathfrak{B}(a, \Omega)$ provided there exists a smooth map $F : N \times B(a) \rightarrow \Omega$ such that the restriction $F(\eta, -) : B(a) \rightarrow \Omega$ is a symplectic embedding for each $\eta \in N$, and such that $f(\eta) = F(\eta, 0)$. It is necessary that $f^*T\Omega$ is symplectically trivializable for such F to exist.

To handle the case when $f^*T\Omega$ is not symplectically trivializable, we refine the notion of lifting. We say f lifts to $\mathfrak{B}(a, \Omega)/U(n)$ provided that:

- (1) N admits an open cover by sets U_α ,
- (2) there are maps $F_\alpha : U_\alpha \times B(a) \rightarrow \Omega$ such that $F_\alpha(\eta, 0) = f(\eta)$ for $\eta \in U_\alpha$, and which are embeddings on each $B(a)$ factor.
- (3) there are maps $g_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow U(n)$ such that:

$$F_\beta(\eta, z) = F_\alpha(\eta, g_{\alpha\beta}(\eta)z),$$

for $\eta \in U_\alpha \cap U_\beta$ and $z \in B(a)$.

These conditions implies that the $g_{\alpha\beta}$ form a cocycle, and hence define a unitary bundle over N . The derivatives of the F_α establish a symplectic isomorphism between this unitary bundle and $f^*T\Omega$.

The *parametric Gromov width* of f in Ω is defined by the formula:

- (1) $\text{Gr}(f, \Omega) := \sup\{a : f \text{ admits a lift to } \mathfrak{B}(a, \Omega)/U(n)\}.$

It is not hard to see that $\text{Gr}(f, \Omega)$ is invariant under homotopies of f . Moreover, a simple Moser-type argument proves $\text{Gr}(f, \Omega) > 0$ holds for all maps f . Furthermore, if $f^*T\Omega$ is trivializable, then f lifts to $\mathfrak{B}(a, \Omega)/U(n)$ if and only if it lifts to $\mathfrak{B}(a, \Omega)$.

One obvious choice of f is the inclusion of a point $f = [\text{pt}]$, in which case $\text{Gr}([\text{pt}], \Omega)$ is simply the classical Gromov width of Ω .

The goal of this paper is to provide Floer theoretic upper bounds on $\text{Gr}(f, \Omega)$ when Ω is the interior of a Liouville domain. Recall that a *Liouville domain* $(\bar{\Omega}, \omega = d\lambda)$ is a compact connected $2n$ -dimensional exact symplectic manifold with boundary $\partial\Omega$, such that the Liouville vector field Z defined by $Z \lrcorner d\lambda = \lambda$ is outwardly transverse to $\partial\Omega$.

Henceforth we reserve the symbol Ω for the interior of a Liouville domain.

The main examples of Ω we will consider are fiberwise starshaped domains in cotangent bundles, and the upper bounds we state below ultimately come from the relationship between string topology and the BV-algebra structure on Floer cohomology.

1.1. Examples. In this section we state some applications of our methods. The proofs are contained in §2.2.

1.1.1. Thin ellipsoids 1. Let Ω_a be the unit codisk bundle in T^*S^n associated to the metric obtained by embedding S^n into $\mathbb{R}^{n+1}$ as the level set:

$$\{x_0^2 + x_1^2 + \cdots + a^{-2}(x_{n-1}^2 + x_n^2) = 1\},$$

where $a \leq 1$ and $n \geq 2$. Let $[S^n]$ be the inclusion of the zero section $S^n \rightarrow \Omega$. We will show:

Theorem 2. $\text{Gr}([S^n], \Omega_a) = 2\pi a$. If $a \leq 1$, then:

$$\text{Gr}([\text{pt}], \Omega_a) \leq 4\pi a,$$

while, when $a = 1$, $\text{Gr}([S^n], \Omega_1) = \text{Gr}([\text{pt}], \Omega_1) = 2\pi$.

The fact that $\text{Gr}([\text{pt}], \Omega_1) = 2\pi$ is known; see [KS21, §6.3].

Remark. It can be shown, by comparison with the unit cotangent bundle of the cylinder $a^{-2}(x_{n-1}^2 + x_n^2) = 1$, that, as $a \rightarrow 0$, the ratio $\text{Gr}([\text{pt}], \Omega_a)/4\pi a$ converges to 1. Indeed, in dimension $n = 2$, it has been shown in [FRV23] that the Gromov width of Ω_a eventually equals $4\pi a$. This remark illustrates that the parametric Gromov width can be non-zero and strictly less than the usual Gromov width.

1.1.2. Thin ellipsoids 2. Let Ω_a be the unit codisk bundle in T^*S^n associated to the metric obtained by embedding S^n into $\mathbb{R}^{n+1}$ as the level set:

$$\{x_0^2 + x_1^2 + \cdots + a^{-2}(x_{n-3}^2 + x_{n-2}^2 + x_{n-1}^2 + x_n^2) = 1\},$$

with $n \geq 3$. In this case our methods give:

Theorem 3. $\text{Gr}([\text{pt}], \Omega_a) = \text{Gr}([S^n], \Omega_a) = 2\pi a$.

1.1.3. Open books with trivial monodromy. Let $(V, \partial V)$ be a compact and connected manifold with boundary and let:

$$M = (V \times \mathbb{R}/\mathbb{Z}) \cup (\partial V \times D(1))$$

be considered as a smooth open book with page $(V, \partial V)$ and trivial monodromy. Let $\mathcal{L}_+$ be the set of oriented loops of the form:

- (1) $\{v\} \times \mathbb{R}/\mathbb{Z}$, $v \in V$,
- (2) $\{v\} \times \partial D(r)$, $v \in \partial V$ and $r \leq 1$,

which form a singular foliation of M (the singularities occur along the binding, where the loops are constant). Similarly let $\mathcal{L}_-$ be the same set of loops but with the reverse orientation. For each loop $q \in \mathcal{L}_\pm$, pick a parametrization and define:

$$(2) \quad \ell_\Omega(q) = \int_0^1 \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\} dt.$$

This quantity is independent of the choice of parametrization, and should be considered as the length measured using Ω . Define:

$$(3) \quad E_\pm = \sup\{\ell_\Omega(q) : q \in \mathcal{L}_\pm\} \text{ and } e_\pm = \inf\{\ell_\Omega(q) : q \in \mathcal{L}_\pm\}.$$

Our methods give the upper bounds:

Theorem 4. Let $[M]$ be the inclusion of the zero section, and let $[V]$ be the inclusion of the page $V \times \{0\}$. Then:

- (1) $\text{Gr}([\text{pt}], \Omega) \leq E_+ + E_-$,
- (2) if $\partial V \neq \emptyset$ then $\text{Gr}([M], \Omega) \leq \min\{E_+, E_-\}$, and,
- (3) if $\partial V = \emptyset$, then $\text{Gr}([V], \Omega) \leq \min\{E_+ + e_-, E_- + e_+\}$.

There are examples with $\partial V \neq \emptyset$ for which both (1) and (2) are equalities, and examples with $\partial V = \emptyset$ for which both (1) and (3) are equalities.

1.1.4. Product with a torus. The case of $M = V \times \mathbb{R}/\mathbb{Z}$ where $\partial V = \emptyset$, one can use the non-contractibility of the orbits $\{x\} \times \mathbb{R}/\mathbb{Z}$ to bound the Gromov width without appealing to the structure of the BV-operator (we will present an argument which uses only classical displacement energy ideas in §1.4). However, our methods still give interesting bounds on the parametric Gromov width which do not seem to be accessible with the more classical methods of §1.4. Indeed, part (3) of Theorem 4 is already such a result. In this section we will state additional results for manifolds of the form $M = V \times T^d$, where V is a closed manifold.

Fix a fiberwise starshaped domain $\Omega \subset T^*(V \times T^d)$. In the following, the class $[V \times T^k]$ represents the inclusion $V \times T^k \rightarrow V \times T^d$, where $T^k \subset T^d$ is the subset of points of the form $(x_1, \dots, x_k, 0, \dots, 0)$. We also denote by:

$$\mathcal{L}_- = \{\text{loops of the form } t \in \mathbb{R}/\mathbb{Z} \mapsto (v, x_1, \dots, x_{d-1}, -t)\},$$

$$\mathcal{L}_+^k = \{\text{loops of the form } t \in \mathbb{R}/\mathbb{Z} \mapsto (v, 0, \dots, 0, x_{k+1}, \dots, x_{d-1}, t)\},$$

where $v \in V$, and where we require that $k < d$. Let E_- be the maximum ℓ_Ω -length of loops in $\mathcal{L}_-$ and E_+^k the maximum ℓ_Ω -length of loops in class $\mathcal{L}_+^k$, similarly to (3).

Theorem 5. *With the notation set in the preceding paragraph, we have:*

$$\text{Gr}([T^k], \Omega) \leq E_- + E_+^k,$$

where $[T^k]$ is represented by $x \mapsto (v_0, x_1, \dots, x_k, 0, \dots, 0)$ for some basepoint v_0 (the homotopy class is independent of v_0).

This result may seem abstruse; however, in the case when $V = \text{pt}$ and $k = 1$ we should note that it obstructs a loop of symplectic balls similarly to the classical camel theorem Proposition 1. In fact, this example will be used to prove Proposition 1 and the argument is given in §2.2.6.

1.1.5. Non-orientable surfaces. Let Σ be a compact non-orientable surface, and let $\Omega \subset T^*\Sigma$ be a fiberwise starshaped domain. Define:

$$\mathcal{L} = \{\text{set of loops } q : \mathbb{R}/\mathbb{Z} \rightarrow \Sigma \text{ such that } q^*T\Sigma \text{ is non-orientable}\},$$

and let:

$$E = \inf\{\ell_\Omega(q) + \ell_\Omega(\bar{q}) : q \in \mathcal{L}\},$$

where $\bar{q}$ denotes the loop q traversed in reverse. Then our methods bound the parametric Gromov width for the inclusion of the zero section $[\Sigma]$:

Theorem 6. *With the above notation, $\text{Gr}([\Sigma], \Omega) \leq E$.*

This theorem is interesting because it applies to surfaces with negative curvature, where it is generally hard to bound symplectic capacities. It is also interesting to ponder the role of non-orientability. For instance, since certain non-orientable surfaces Σ embed as Lagrangians in $\mathbb{C}^2$, the Hofer-Zehnder capacity of any disk cotangent bundle over such Σ is finite. However, it seems to be an open question whether the same fact is true for orientable surfaces of genus at least two (see, e.g., [Bim24, pp. 105]).

1.2. Floer cohomology persistence module. Let W be the completion of $\bar{\Omega}$ obtained by attaching the symplectization end $\partial\Omega \times [1, \infty)$ to $\bar{\Omega}$ in such a way that Z extends to $r\partial_r$, where r is the projection to $[1, \infty)$, and such that the extension of Z is a Liouville vector field.

Let $\mathcal{H}$ be the space of all Hamiltonian functions H such that $H = cr$ holds outside of a compact set, for some $c \in \mathbb{R}$.

Fix an almost complex structure J which is invariant under the flow by Z in the end. For a time-dependent family of Hamiltonian functions $H_t \in \mathcal{H}$, whose flow φ_t has a non-degenerate time-1 map, and for which the Floer cohomology chain complex $\text{CF}(H_t)$ is well-defined. Here the Floer complex is the $\mathbb{Z}/2\mathbb{Z}$ vector space generated by the 1-periodic orbits of φ_t , and the differential uses the almost complex structure J . Denote by $\text{HF}(H_t, J)$ the resulting homology.

Using continuation maps, the resulting homology group depends only on the *slope*, namely the average value of H_t/r in the end; we denote by V_c the homology group for a system with slope c . It is well-known that continuation maps endow $c \mapsto V_c$ with the structure of a persistence module, namely, a functor from $(\mathbb{R}, \leq) \rightarrow \text{Vect}(\mathbb{Z}/2\mathbb{Z})$. Precise details of this construction are recalled in §3.4.3. The colimit of V_c as $c \rightarrow \infty$ is the so-called *symplectic cohomology* $\text{SH}(W)$.

This persistence module V_c has three structures relevant to our paper:

- (1) the product structure $* : V_{c_1} \otimes V_{c_2} \rightarrow V_{c_1+c_2}$ induced by the Floer cohomology pair-of-pants operation,
- (2) the BV-operator $\Delta : V_c \rightarrow V_c$ induced by counting $\mathbb{R}/\mathbb{Z}$ -families of Floer cylinders,
- (3) the PSS morphism $\text{PSS} : H^*(W) \rightarrow V_c$ for $c > 0$;

we refer the reader to [AS10, Abo15] for background on these structures.

Our main result is the following:

Theorem 7. *Suppose that $\zeta_i \in V_{c_i}$, $i = 1, 2$, with $c_i > 0$, are such that:*

$$\text{PSS}(\beta) = \Delta(\zeta_1) * \zeta_2 \text{ holds in } V_{c_1+c_2}.$$

If a map $f : N \rightarrow \Omega$ has a non-zero mod 2 homological intersection number with the cohomology class $\beta \in H^(W)$, then $\text{Gr}(f, \Omega) \leq c_1 + c_2$.*

The idea in the proof of Theorem 7 is to define a sort of *evaluation map* $V_c \rightarrow \mathbb{Z}/2\mathbb{Z}$ using the family of ball embeddings $N \times B(a) \rightarrow \Omega$. The map is defined and non-trivial on $\text{PSS}(\beta)$ provided that the slope c is smaller than the capacity a . Arguing using a special action filtration defined using family Floer cohomology (in the sense of [Hut08]), we will show that the non-triviality of this map obstructs the existence of a solution to the equation appearing in Theorem 7. Thus, if the equation can be solved, the slope c must be larger than the capacity a , as desired.

The details of this argument are given in §3.6.

1.2.1. Relative Hofer-Zehnder capacities. In fact, the argument we give easily generalizes to the following statement concerning relative Hofer-Zehnder

capacities. Let us define a *Hofer-Zehnder admissible function* $K : \Omega \rightarrow \mathbb{R}$ with *depth* A and *well* U (a non-empty open set) to be a function such that:

- (1) K is compactly supported in Ω ,
- (2) $-A = \min K$, and the interior of $\{K = -A\}$ is U ,
- (3) X_K has no nonconstant 1-periodic orbits.

Then we have:

Theorem 8. *Suppose that the hypotheses of Theorem 7 are satisfied for slope $c > 0$ and map $f : N \rightarrow \Omega$. For each Hofer-Zehnder admissible function K whose well U contains the image of f , we have:*

$$A + \text{Gr}(f, U) \leq c,$$

where A is the depth of K . In particular, the so-called relative Hofer Zehnder capacity of $f(N) \subset \Omega$ is bounded from above by c .

This is proved in §3.6.9. Here we recall the relative Hofer Zehnder capacity of a pair $Y \subset X$ (where X is a symplectic manifold) is simply the deepest depth A one can achieve amongst Hofer-Zehnder admissible functions $K : X \rightarrow \mathbb{R}$ subject to the constraint that the well U contains Y .

1.2.2. Dilation classes. The equation appearing in Theorem 7 can be considered as a generalization of the *dilation class equation* introduced in [SS12]. Recall that a class ζ is called a *dilation class* if $\text{PSS}(1) = \Delta\zeta$. Thus our theorem implies:

Corollary 9. *The existence of a dilation class in the Floer cohomology group V_c bounds the usual Gromov width by c .*

This result is probably not so surprising to experts, as the slopes c for which dilation classes appear are already used in quantitative symplectic geometry; see, e.g., [Sei14, Zho21].

The result implied by Theorem 3 on the Gromov width on the cotangent bundle of the round S^3 then follows from the existence of a dilation class in the Floer cohomology group of the appropriate slope, and the above corollary. Indeed, in Theorem 3, our proof essentially shows that the dilation class equation $\text{PSS}(1) = \Delta\zeta$ can be solved in the appropriate group V_c .

It is interesting to compare the situation with our results on tori, as the cotangent bundles of tori (or more generally $K(\pi, 1)$ spaces) never have dilation classes. However, the more general equation in Theorem 7 does admit solutions (as we exploit in Theorem 5).

There are other classes of manifolds which are known to admit dilation classes. For instance, [SS12] show that the total space of certain Lefschetz fibrations admit dilation classes. To apply the above corollary one would need present the total space as a completion of a Liouville domain, and then estimate the precise slopes for which the dilation class equation can be solved in total spaces of Lefschetz fibrations; we do not analyze this question in this paper.

1.2.3. Subcritical handle attachment. Another interesting class of Liouville domains where the equation in Theorem 7 can be solved are domains Ω

obtained by a subcritical handle attachment to a Liouville domain Ω_0 . Recall that Ω is obtained by attaching a handle along an isotropic sphere S^{k-1} in the contact boundary $\partial\Omega_0$, and if $k < n$ we say the handle attachment is subcritical. For instance, attaching a 1-handle to a 4-dimensional Liouville domain is an example of a subcritical handle attachment. For more details on handle attachments see [Wei91, Cie02, Fau20].

If the attaching sphere is null-homologous in Ω_0 , then the *cocore disk* (which is a properly embedded open disk $D^{2n-k} \rightarrow \Omega$ which intersects the attaching disk in a single point) defines a cohomology class in Ω . It is known that:

Proposition 10. *The cohomology class of the cocore disk β satisfies:*

$$\text{PSS}(\beta) = 0$$

in V_c for c sufficiently large.

In particular, β satisfies the equation in Theorem 7.

Proof. This follows from [Cie02] (see also [Fau20, Theorem 1.3]) which proves the so-called Viterbo restriction map from [Vit99] is an isomorphism from the symplectic cohomology of W_0 to the symplectic cohomology of W (the completions of Ω_0 and Ω , respectively).

It is also known that the Viterbo restriction map (V.R.) commutes with PSS and the pullback map on cohomology associated to the inclusion $i : \Omega_0 \rightarrow \Omega$:

$$(4) \quad \begin{array}{ccc} H^*(W) & \xrightarrow{i^*} & H^*(W_0) \\ \downarrow \text{PSS} & & \downarrow \text{PSS} \\ \text{SH}(W) & \xrightarrow{\text{V.R.}} & \text{SH}(W_0), \end{array}$$

Since i^* takes β to 0, and V.R. is an isomorphism, it follows that $\text{PSS}(\beta) = 0$ in $\text{SH}(W)$, as desired. $\square$

Theorem 7 then bounds the parametric Gromov width of any map $f : N \rightarrow \Omega$ which has a non-zero intersection number with the cocore disk. This brings us close to the original camel problem Proposition 1; indeed, in one formulation of the camel problem, the relevant space is the domain obtained by attaching a 1-handle to a ball, see [MS17, Figure 1.3].

To apply our methods to obtain explicit bounds on the parametric Gromov width, one would need to set up a careful model for the handle attachment $\Omega_0 \rightarrow \Omega$, and determine at which slopes c the equation $\text{PSS}(\beta) = 0$ can be solved in V_c . We do not perform such an analysis in this paper. We will instead recover Proposition 1 using a string topological computation in §2.2.

1.3. Comparison between string topology and Floer cohomology. For our applications to fiberwise starshaped domains in cotangent bundles, we will use the well-established strategy of comparing string topology and Floer cohomology. The relationship between string topology and Floer cohomology of cotangent bundles is developed in [Vit99, AS06, SW06, AS10, Abo15].

Our approach is inspired by Shelukhin's proof of the Viterbo conjecture for certain manifolds in [She22a, She22b]. In particular, his approach in [She22a]

exploits the relationship between the string topology and symplectic cohomology for cotangent bundles of manifolds which are *string point-invertible*; the notion of string point-invertibility is defined in terms the string bracket which measures the failure of the BV operator Δ to satisfy the Leibniz rule.

Another inspiration is Irie's bound on the Hofer-Zehnder capacity of disk cotangent bundles of manifolds admitting circle actions with non-contractible fibers [Iri14]; his argument appeals to the aforementioned product structures.

Our framework for string topology is as follows. Let Ω be a fiberwise starshaped domain in T^*M . For $c > 0$ and $k = 0, 1, \dots$, introduce $Z_k(\Lambda_c)$ as the monoid of smooth maps $A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M$ where:

- (1) $\dim P = k$,
- (2) $\ell_\Omega(A(x, -)) < c$, for all $x \in P$, where ℓ_Ω is the length function defined in (2);

here $A_1 + A_2$ is the coproduct $(P_1 \sqcup P_2) \times \mathbb{R}/\mathbb{Z} \rightarrow M$.

A *cobordism* between $A_1, A_2 \in Z_k(\Lambda_c)$ is a cobordism Q between P_1, P_2 and a smooth map $C : Q \times \mathbb{R}/\mathbb{Z} \rightarrow M$ satisfying:

- (2') $\ell_\Omega(C(q, -)) < c$ for all $q \in Q$,

and such that C restricts to A_1, A_2 on the boundary. Define $H_k(\Lambda_c)$ be the quotient of $Z_k(\Lambda_c)$ by the cobordism relation. The monoid structure on $Z_k(\Lambda_c)$ induces on $H_k(\Lambda_c)$ the structure of a vector space over $\mathbb{Z}/2\mathbb{Z}$. Write $Z(\Lambda_c), H(\Lambda_c)$ for the direct sum of the graded pieces $Z_k(\Lambda_c), H_k(\Lambda_c)$, respectively. If $A \in Z(\Lambda_c)$, then we write its image $\alpha \in H(\Lambda_c)$ using the symbol $\alpha = [A]$. One thinks of $H(\Lambda_c)$ as a convenient proxy for the homology of the loop space.

The obvious inclusion morphisms $H(\Lambda_{c_1}) \rightarrow H(\Lambda_{c_2})$ endow $c \mapsto H(\Lambda_c)$ with the structure of a persistence module defined for $c > 0$.

As in the Floer cohomology case, this persistence module has three natural structures:

- (1) the *Chas-Sullivan product* $* : H(\Lambda_{c_1}) \otimes H(\Lambda_{c_2}) \rightarrow H(\Lambda_{c_1+c_2})$,
- (2) the *BV-operator* $\Delta : H(\Lambda_c) \rightarrow H(\Lambda_c)$,
- (3) an *inclusion of constant loops* morphism $\mathfrak{i} : H^*(W) \rightarrow H(\Lambda_c)$ for $c > 0$, which sends a degree d cohomology class to an element in $H_{n-d}(\Lambda_c)$.

The technical result relating string topology and the Floer cohomology persistence module is:

Theorem 11. *There is a morphism of persistence modules $\Theta_c : H(\Lambda_c) \rightarrow V_s$ such that:*

- (1) $\circ * (\Theta_{c_1}, \Theta_{c_2}) = \Theta_{c_1+c_2} \circ *$ holds on $H(\Lambda_{c_1}) \otimes H(\Lambda_{c_2})$,
- (2) $\Delta \circ \Theta_c = \Theta_c \circ \Delta$,
- (3) $\text{PSS} = \Theta_c \circ \mathfrak{i}$.

The proof of this theorem occupies §4. Most of the arguments are similar to those in [AS06, AS10, Abo15], with small modifications due to the fact we work with bordism classes of loops (rather than chains in the Morse

homology associated to an energy functional). However, working directly with bordism classes makes the existing proof that the product structures are identified difficult to implement (the existing proof is not a direct argument, and instead is based on a homological algebra argument). For this reason, we provide a new proof that the product structures are identified. Interestingly enough, this leads to an *adiabatic gluing problem*, similar to those considered in [FO97, Ekh07, OZ11]. This argument is given in §4.4.

Combining Theorem 7 and Theorem 11 yields the following upper bound on the relative Gromov width in terms of string topology:

Theorem 12. *Suppose there are classes $\alpha_i \in H(\Lambda_{c_i})$, $i = 1, 2$, with $c_i > 0$, such that:*

$$i(\beta) = \Delta(\alpha_1) * \alpha_2 \text{ holds in } H(\Lambda_{c_1+c_2}).$$

If a map $f : N \rightarrow \Omega$ has a non-zero mod 2 homological intersection number with the cohomology class $\beta \in H^(W)$, then $\text{Gr}(f, \Omega) \leq c_1 + c_2$. $\square$*

This corollary will be the result used in the demonstration of our examples. One of course has a similar corollary concerning the relative Hofer-Zehnder capacity by combining Theorem 11 with Theorem 8.

1.4. String topology in non-trivial free homotopy classes. In the case when there is sufficiently rich string topology in non-trivial free homotopy classes, one can establish bounds on the Gromov width without appealing to the BV-operator. For instance, a main result of [Iri14] uses the existence of two classes $\alpha_i \in H(\Lambda_{c_i})$, $i = 1, 2$, such that:

- (1) $\alpha_1 * \alpha_2 = i([M])$,
- (2) α_i lies in a non-trivial free homotopy class,

to prove the Hofer-Zehnder capacity of Ω is finite. Typically the way one obtains such classes α_1, α_2 is via $\mathbb{R}/\mathbb{Z}$ -actions whose orbits are non-contractible.

Such considerations yield the following variation of Theorem 7:

Theorem 13. *Suppose that $\zeta_i \in V_{c_i}(\Omega)$, $i = 1, 2$, with $c_i > 0$, are such that:*

$$\text{PSS}(\beta) = \zeta_1 * \zeta_2 \text{ holds in } V_{c_1+c_2},$$

and suppose that ζ_i can be represented by a cycle whose orbits are all non-contractible.

If a map $f : N \rightarrow \Omega$ has a non-zero mod 2 homological intersection number with the cohomology class $\beta \in H^(W)$, then $\text{Gr}(f, \Omega) \leq c_1 + c_2$.*

The proof is given in §3.6.10. There is also an obvious variation of Theorem 8 based on Theorem 13, whose statement we omit.

In certain cases where orbits are non-contractible, one can prove bounds on the Gromov width using more classical ideas of Hamiltonian displacement. Indeed, we have:

Proposition 14. *Suppose f_t is a smooth isotopy of a closed manifold M such that $f_0 = f_1 = \text{id}$ and such that the orbits $x \mapsto f_t(x)$ are non-contractible; denote by κ the free homotopy class containing these orbits. There exists a constant $\text{const}(f_t, \Omega)$ such that, for any covering space $M' \rightarrow M$ to which*

loops in κ do not lift, the following holds: if $K \subset \Omega$ is a compact set which does admit a lift to $K' \subset T^*M'$, then K' has Hofer displacement energy at most $\text{const}(f_t, \Omega)$.

Remark. This implies that the Gromov width of Ω is bounded by $\text{const}(f_t, \Omega)$. It also implies that Lagrangians which lift to T^*M' bound holomorphic disks with symplectic area at most $\text{const}(f_t, \Omega)$.

Proof. Let Φ_t the canonical lift of f_t , which lifts to an isotopy Φ'_t of T^*M' such that Φ'_1 is a canonical lift of a deck transformation of $M' \rightarrow M$. Then Φ'_1 displaces any set K' as in the statement. One can cut-off Φ'_t in a uniform way without changing how it acts on the inverse image of Ω . The cut-off can be constructed so that it has a finite Hofer length. The desired result follows. $\square$

1.5. Conventions.

1.5.1. On the cohomology of W . For convenience, we define $H^*(W)$ to be the group of smooth proper maps $C : S \rightarrow W$, where S is a smooth manifold, modulo proper cobordisms. This is graded by the codimension of F . This is a proxy for the cohomology of W . Prototypical examples are the fundamental class $\text{id} : W \rightarrow W$ and, in the case when $W = T^*M$, the inclusion of the fiber $T^*M_q \rightarrow T^*M$ for some q .

1.6. Acknowledgements. First of all, the authors wish to thank Egor Shelukhin for many interesting discussions and helpful comments. The first author is thankful to Kai Cieliebak for useful suggestions. The second author also wishes to thank Georgios Dimitroglou-Rizell, Tobias Ekholm, and Yin Li, for key discussions during his visit to Uppsala University. The authors also wish to thank useful discussions with Alberto Abbondandolo and Johanna Bimmermann after the first version of this text was posted. F.B. is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 517480394. D.C. is supported by the ANR project CoSy.

2. String topology

In this section, we continue the discussion of string topology from where the introduction left off.

2.1. Three structures on the string topology persistence module. In order to apply Theorem 12 it is necessary to explain the three structures $*$, Δ , $\mathbf{i}$.

2.1.1. The Chas-Sullivan product. The product $\alpha_1 * \alpha_2 \in H_{k_1+k_2-n}(\Lambda_{c_1+c_2})$ of two classes $\alpha_i \in H_{k_i}(\Lambda_{c_i})$, $i = 1, 2$, originally defined in [CS99], can be thought of as a combination of the intersection product and the Pontrjagin (concatenation) product.

It is defined as follows: write $\alpha_i = [A_i]$ where A_1, A_2 are in *general position*, which means that the *evaluation-at-zero* maps:

$$e_i : x \in P_i \mapsto A_i(x, 0),$$

$i = 0, 1$, are transverse to one-another. Such representatives always exist. Define P_3 to be the transverse fiber product of these two maps:

$$\begin{array}{ccc} P_3 & \longrightarrow & P_2 \\ \downarrow & & \downarrow e_2 \\ P_1 & \xrightarrow{e_1} & M, \end{array}$$

so that P_3 is a compact manifold of dimension $k_1 + k_2 - n$, where $n = \dim M$. Concretely P_3 is the submanifold of pairs $(x_1, x_2) \in P_1 \times P_2$ satisfying the incidence $e_1(x_1) = e_2(x_2)$. Define $A_3 : P_3 \times \mathbb{R}/\mathbb{Z} \rightarrow M$ by the formula:

$$A_3(x_1, x_2, t) := \begin{cases} A_1(x_1, \beta(2t)) & \text{for } t \in [0, 1/2], \\ A_2(x_2, \beta(2t - 1)) & \text{for } t \in [1/2, 1], \end{cases}$$

and extended by 1-periodicity; here β is a standard smooth cut-off function which equals 0 for $t \leq 0$ and equals 1 for $t \geq 1$. The cut-offs ensure A_3 is a smooth map. It is important to note that $A_3 \in Z(\Lambda_{c_1+c_2})$, because the length function is additive under concatenation and invariant under time reparametrizations.

Lemma 15. *The class $[A_3]$ in $H_{k_1+k_2-n}(\Lambda_{c_1+c_2})$ is independent of the choice of representatives A_1, A_2 in general position, and depends only on α_1, α_2 .*

Proof. This is a straightforward argument in differential topology, similar to the arguments in [Mil65b], and is left to the reader. $\square$

Observe also that $(x_3, t) \mapsto A_3(x_3, t + 1/2)$ represents the class of $\alpha_2 * \alpha_1$, and hence $\alpha_1 * \alpha_2 = \alpha_2 * \alpha_1$, since we can homotope from one to the other via the formula $(x_3, t, s) \mapsto A_3(x_3, t + s)$.

Similarly, one can prove that $*$ is an associative product, although we leave the details of this to the reader.

2.1.2. The BV-operator. Like the Chas-Sullivan product, this string topology operation is introduced in [CS99]. It is defined as follows; given $\alpha \in H_k(\Lambda_c)$, write $\alpha = [A]$, where $A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M$, and introduce:

$$\Delta(A) : (\mathbb{R}/\mathbb{Z} \times P) \times \mathbb{R}/\mathbb{Z} \rightarrow M \text{ given by } \Delta(A)(\theta, x)(t) = A(x, t - \theta).$$

Then $\Delta(A) \in Z_{k+1}(\Lambda_c)$ and we define $\Delta(\alpha) = [\Delta(A)]$. It is trivial to check that $\Delta(\alpha)$ is independent of the choice of representative A .

2.1.3. Inclusion of the constant loops. Let $\beta \in H^d(T^*M)$ be represented by a smooth proper map $C : S \rightarrow T^*M$ which is transverse to the zero section. The class $i(\beta) \in H_{n-d}(\Lambda_c)$ is the image of β under a sort of Thom isomorphism.

Define $f : P \rightarrow M$ via the fiber product:

$$\begin{array}{ccc} P & \longrightarrow & S \\ \downarrow f & & \downarrow C \\ M & \longrightarrow & T^*M, \end{array}$$

and define $A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M$ by $A(x, t) = f(x)$. Set:

$$\mathbf{i}(\beta) := [A] \in H_{n-d}(\Lambda_c).$$

It is straightforward to prove that $\mathbf{i}(\beta)$ in $H_{n-d}(\Lambda_c)$ is independent of the representative $C : S \rightarrow T^*M$ (since, by our proxy definition of $H^*(T^*M)$, any two representatives are properly cobordant).

The following discussion sheds light on the map $\mathbf{i}$, and will be used in §4.3 in the comparison between string topology and Floer cohomology; it shows that $\mathbf{i}$ is essentially a Thom isomorphism between cohomology of T^*M and homology of M .

Lemma 16. *Any class $C : S \rightarrow M$ is cohomologous to the class $C' : S' \rightarrow M$ determined by the fiber product:*

$$\begin{array}{ccc} S' & \xrightarrow{C'} & T^*M \\ \downarrow & & \downarrow \\ P & \xrightarrow{f} & M, \end{array}$$

where f is defined above.

Proof. Observe that C and C' have the same transverse intersection with the zero section. It follows from a straightforward surgery operation that $C + C'$ (the sum is taken in the cohomology cobordism group) is properly cobordant to a proper map $G : T \rightarrow T^*M$ whose image is disjoint from the zero section.

Then the map $T \times [0, \infty) \rightarrow T^*M$ given by $x, t \mapsto \rho_t(G(x, t))$, where ρ_t is the Liouville flow, is a proper cobordism from G to $\emptyset$ (it is proper because the image of G is disjoint from the zero section). Thus $[C + C'] = 0$, and hence the desired relation $[C] = [C']$ holds. $\square$

2.2. Proofs for §1.1. In this section, we provide the proofs of the Theorems stated in §1.1. The upper bounds in Theorem 2 will ultimately follow from Theorem 4. The proof of Theorem 4 is based on a geometric construction of classes solving the equation appearing in Theorem 12. The proof of Theorem 3 will follow from similar considerations, but with a different classes in the string topology of open books. We prove Theorem 5 and explain how to recover Proposition 1 in §2.2.6. Finally we prove Theorem 6 in §2.2.7.

2.2.1. Action classes. One recurring idea is the notion of an *action class*. Let $\zeta_t : M \rightarrow M$, $t \in \mathbb{R}/\mathbb{Z}$, be a circle action. The *action class* is the class:

$$A : M \times \mathbb{R}/\mathbb{Z} \rightarrow M \text{ given by } A(x, t) = \zeta_t(x).$$

More generally, for any bordism class $g : P \rightarrow M$, one defines:

$$A_{g,\pm} : P \times \mathbb{R}/\mathbb{Z} \rightarrow M \text{ given by } A_{g,\pm}(x, t) = \zeta_{\pm t}(g(x)).$$

It is clear that $[A_{g,\pm}]$ depends only on the bordism class of g in M .

Denote by $E_{g,\pm}$ the maximal ℓ_Ω -length of the orbits appearing in $A_{g,\pm}$.

Such classes behave well with respect to the Chas-Sullivan product.

Lemma 17. *For two transverse maps $g_i : P_i \rightarrow M$, $i = 1, 2$, one has:*

$$[A_{g_1,+}] * [A_{g_2,-}] = [g_1 \cap g_2] \text{ in } \Lambda_{E_{g_1,+} + E_{g_2,-}},$$

where $[g_1 \cap g_2]$ is the class of constant loops:

$$(x_1, x_2) \in P_3 \times \mathbb{R}/\mathbb{Z} \mapsto g_1(x_1) = g_2(x_2),$$

where P_3 is the fiber product of g_1 and g_2 .

Proof. The evaluation at 0 map of $A_{g_i,\pm}$ is g_i . Thus $A_{g_1,+} * A_{g_2,-}$ is represented by the map:

$$(5) \quad (x_1, x_2, t) \in P_3 \times \mathbb{R}/\mathbb{Z} \mapsto \begin{cases} \zeta_{\beta(2t)}(g_1(x_1)) & \text{for } t \in [0, 1/2], \\ \zeta_{-\beta(2t-1)}(g_2(x_2)) & \text{for } t \in [1/2, 1]. \end{cases}$$

We claim this class is homotopic to the class of $[g_1 \cap g_2]$ described in the statement of the lemma; furthermore, we claim the homotopy can be made inside of $\Lambda_{E_{g_1,+} + E_{g_2,-}}$. To show this, define $G : P_3 \times [0, 1] \times \mathbb{R}/\mathbb{Z} \rightarrow M$ by:

$$G(x_1, x_2, s, t) := \begin{cases} \zeta_{\beta(2st)}(g_1(x_1)) & \text{for } t \in [0, 1/2], \\ \zeta_{\beta(s) - \beta(2st-s)}(g_2(x_2)) & \text{for } t \in [1/2, 1]. \end{cases}$$

This defines a smooth homotopy, within $\Lambda_{E_{g_1,+} + E_{g_2,-}}$, between the class of constant loops $g_1 \cap g_2$, at $s = 0$, and (5), at $s = 1$. $\square$

Another lemma which will be used in the proof of Theorem 4 is:

Lemma 18. *For two transverse maps $g_i : P_i \rightarrow M$, $i = 1, 2$, then:*

$$[A_{g_1,\pm}] * [g_2] = [A_{g_1 \cap g_2, \pm}] \text{ in } \Lambda_{E_{g_1,\pm}},$$

where $[g_2]$ is considered as a class of constant loops, and $g_1 \cap g_2 : P_3 \rightarrow M$ is as in Lemma 17.

Proof. The proof is straightforward, and easier than Lemma 17. $\square$

Given a smooth map $g : P \rightarrow M$, denote by $\zeta g : \mathbb{R}/\mathbb{Z} \times P \rightarrow M$ the class defined by $\zeta g(\theta, x) = \zeta_\theta(g(x))$. The final result we require concerning action classes is:

Lemma 19. *Given a smooth map $g : P \rightarrow M$, it holds that:*

$$\Delta[A_{g,\pm}] = [A_{\zeta g, \pm}]$$

in $H(\Lambda_{E_{g,\pm}})$.

Proof. This is a straightforward calculation: by definition, $\Delta[A_{g,\pm}]$ is the class represented by $\mathbb{R}/\mathbb{Z} \times P \times \mathbb{R}/\mathbb{Z} \rightarrow M$ given by:

$$(\theta, x, t) \mapsto \zeta_{\pm(t-\theta)}(g(x)).$$

On the other hand, $A_{\zeta g, \pm}$ is given by:

$$(\theta, x, t) \mapsto \zeta_{\pm t + \theta}(g(x)).$$

In the case of $-$ sign, these classes are literally the same. In the case of $+$ sign, the classes differ by the diffeomorphism $\mathbb{R}/\mathbb{Z} \times P \mapsto \mathbb{R}/\mathbb{Z} \times P$ given by $(\theta, x) \mapsto (-\theta, x)$, and so represent the same class in cobordism. $\square$

2.2.2. Proof of Theorem 4. Let $M = (V \times \mathbb{R}/\mathbb{Z}) \cup (\partial V \times D(1))$ be an open book with trivial monodromy as in §1.1.3. There is an obvious circle action ζ which acts as:

$$\zeta_t(x, \theta) = \begin{cases} (x, \theta + t), & (x, \theta) \in V \times \mathbb{R}/\mathbb{Z}, \\ (x, r, \theta + t), & (x, r, \theta) \in \partial V \times D(1), \end{cases}$$

where we use polar coordinates on $D(1)$ factor near the binding.

Theorem 4 will follow from Theorem 12 applied to the action classes associated to this circle action. In order to apply Theorem 12, we need to show that the relevant action classes lie in the image of the BV-operator, among other things.

Abbreviate $A_{\pm} = A_{\text{id}, \pm}$ and $E_{\pm} = E_{\text{id}, \pm}$, using the notation in §2.2.1. Then $E_{\pm}$ agree with the numbers denoted using the same symbols in the statement of Theorem 4.

Lemma 20. *There exist $[B_{\pm}] \in H(\Lambda_{E_{\pm}})$ such that $\Delta[B_{\pm}] = [A_{\pm}]$ in $H(\Lambda_{E_{\pm}})$.*

Proof. This is obvious in the case in $\partial V = \emptyset$ (one can simply appeal to Lemma 19). The argument is harder when $\partial V \neq \emptyset$.

The class $B_{\pm}$ is defined as a map $D(V) \times \mathbb{R}/\mathbb{Z} \rightarrow M$ where here the *double* $D(V)$ is the closed manifold obtained by gluing V to another copy of V , say $\bar{V}$, along the boundary using the identity map.

Then $A_{\pm}$ is defined on $M \times \mathbb{R}/\mathbb{Z}$, and $\Delta(B_{\pm})$ is defined on $\mathbb{R}/\mathbb{Z} \times D(V) \times \mathbb{R}/\mathbb{Z}$. There is a cobordism X between $\mathbb{R}/\mathbb{Z} \times D(V)$ and M , and there is map $C_{\pm} : Q \times \mathbb{R}/\mathbb{Z} \rightarrow M$ which extends $A_{\pm}, B_{\pm}$, and which remains in $\Lambda_{E_{\pm}}$; see Figure 1.

In fact, the construction of $B_{\pm}$ and $C_{\pm}$ follows from a more general argument which we present in §2.2.4. For this reason, we omit the details of the present proof. $\square$




FIGURE 1. The cobordism between M and $\mathbb{R}/\mathbb{Z} \times D(V)$ can be visualized as attaching a generalized handle to M . The simplest example is when $V = [0, 1]$ where $M \simeq S^2$ and $\mathbb{R}/\mathbb{Z} \times D(V) \simeq T^2$.

We now prove Theorem 4 using this lemma and the results in §2.2.1.

First of all, part (1) follows easily from Theorem 12, and the fact that:

$$\mathbf{i}(\text{PD}(T^*M)) = [A_+] * [A_-] = \Delta[B_+] * \Delta[B_-] \text{ in } H(\Lambda_{E_+ + E_-}),$$

where we have used Lemma 17 in the first equality, and the fact that $\mathbf{i}(\text{PD}(T^*M))$ is represented by the constant loops $(x, t) \in M \times \mathbb{R}/\mathbb{Z} \mapsto x$. The second equality follows from the previous lemma. By Theorem 12, this equation bounds the regular Gromov width from above by $E_+ + E_-$, since $f : \text{pt} \rightarrow \Omega$ is dual to the fundamental class $\text{PD}(T^*M)$.

To prove part (2), we use Lemma 18, with $g_1 = \text{id}$ and g_2 the inclusion of a point pt . We then have:

$$(6) \quad \Delta[B_{\pm}] * \mathbf{i}(T^*M_{\text{pt}}) = [A_{\pm}] * [g_2] = [A_{g_2, \pm}] \text{ in } H(E_{\pm}),$$

where we use that $\mathbf{i}(T^*M_{\text{pt}})$ is represented by a single constant loop based at pt .

The class $A_{g_2, \pm}$ is represented by a single loop going around the open book. Since $\partial V \neq \emptyset$, this single loop can be homotoped to a constant loop in the binding; moreover, this homotopy can be taken in $H(E_{\pm})$.

Since the class of $f : M \rightarrow M$ is dual to $\text{PD}(T^*M_{\text{pt}})$, Theorem 12 applied with equation (6) bounds the parametric Gromov width $\text{Gr}([M], \Omega)$ from above by $\min\{E_+, E_-\}$.

To prove part (3), we first appeal to Lemma 19 where the map $g : P \rightarrow M$ is the inclusion of a point pt ; this shows

$$\Delta[A_{g, \pm}] = [A_{\zeta g, \pm}].$$

Then we apply Lemma 17 to conclude:

$$\Delta[A_{g, \pm}] * \Delta[B_{\mp}] = \Delta[A_{\zeta g, \pm}] * [A_{\mp}] = [(\zeta g) \cap \text{id}] \text{ in } \Lambda_{E_{g, \pm} + E_{\mp}},$$

A moment's thought reveals that the class of constant loops $(\zeta g) \cap \text{id}$ (based at points which travel around a single orbit) equals $\mathbf{i}(\beta)$ where β is dual to the inclusion of $V \rightarrow M$. Thus we conclude from Theorem 12 that $\text{Gr}([V], \Omega)$ is bounded by $\min\{E_{g, +} + E_-, E_{g, -} + E_+\}$. Since we can take g to be an arbitrary point, we can make $E_{g, +} = e_+$ or $E_{g, -} = e_-$ (the constants appearing in Theorem 4). This gives the desired result for (3).

To complete the proof of Theorem 4, it remains for us to explain why the estimates are sharp in some cases.

One example with $\partial V \neq \emptyset$ where both (1) and (2) are sharp is $M = S^2$ with the ellipsoidal metric induced by $\{x_0^2 + a^{-2}(x_1^2 + x_2^2) = 1\} \subset \mathbb{R}^3$ for a sufficiently small; that the bound $\text{Gr}([S^2], \Omega_a) \leq 2\pi a$ is sharp is the content of Theorem 2, and further details are given in §2.2.3. That the bound $\text{Gr}([\text{pt}], \Omega_a) \leq 4\pi a$ is sharp is proved in [FRV23] for a small enough.

One example with $\partial V = \emptyset$ where both (1) and (3) are sharp is $M = T^2$ with the flat metric obtained by $T^2 = \mathbb{R}/a\mathbb{Z} \times \mathbb{R}\mathbb{Z}$ with $a \leq 1$. That the Gromov width of the unit codisk bundle Ω_a is $4\pi a$ is proved in [Bro25]; using canonical translations, one can transport this ball in a coherent manner to be based at all the points in $V = \{0\} \times \mathbb{R}/\mathbb{Z}$, proving that $\text{Gr}([V], \Omega_a)$ is also $4\pi a$, as desired (in general, for the codisk bundle Ω of a flat metric

on a torus, or, more generally, the codisk bundle of any Lie group with a bi-invariant metric, it holds that $\text{Gr}(\text{pt}, \Omega) = \text{Gr}(f, \Omega)$.

2.2.3. Proof of Theorem 2. As we will show momentarily, the upper bounds in Theorem 2 follow from Theorem 4. The lower bounds will follow from explicit constructions.

There are two metrics we will consider on S^n in the proof. Set:

$$S^n = \{(x_0, \dots, x_n) : x_0^2 + \dots + x_n^2 = 1\} \subset \mathbb{R}^{n+1}.$$

The first metric g_a we consider is the pullback of the Euclidean metric under the ellipsoid embedding $\varphi_{\text{ell}} : S^n \rightarrow \mathbb{R}^{n+1}$ given by:

$$\varphi_{\text{ell}}(x_0, \dots, x_{n-1}, x_n) = (x_0, \dots, ax_{n-1}, ax_n).$$

The unit codisk bundle of g_a is the domain Ω_a under consideration.

The second metric is the round metric g_a^{round} is the pullback of the Euclidean metric under the embedding $\varphi_{\text{round}} : S^n \rightarrow \mathbb{R}^{n+1}$ given by:

$$\varphi_{\text{round}}(x) = ax;$$

we denote by Ω_a^{round} the associated unit codisk bundle.

Observe that S^n is identified with the open book with page $V = D^{n-1}$, via the parametrization:

$$V \times \mathbb{R}/\mathbb{Z} \ni (x, \theta) \rightarrow (x, \sqrt{1 - \|x\|^2} \cos(2\pi\theta), \sqrt{1 - \|x\|^2} \sin(2\pi\theta)) \in S^n.$$

Since the domain Ω_a is fiberwise symmetric, we have $E_+ = E_-$. An easy computation shows that every loop in $\mathcal{L}_+$ has the length at most $2\pi a$ as measured using Ω_a ; this maximum is achieved for the loop $(0, \cos(2\pi\theta), \sin(2\pi\theta))$. Hence, from Theorem 4 we get

$$\text{Gr}([S^n], \Omega_a) \leq 2\pi a, \text{ and } \text{Gr}([\text{pt}], \Omega_a) \leq 4\pi a.$$

To obtain the equality $\text{Gr}([S^n], \Omega_a) = 2\pi a$, we will first show that $\Omega_a^{\text{round}} \subset \Omega_a$. Thus it follows that $\text{Gr}([S^n], \Omega_a^{\text{round}}) \leq \text{Gr}([S^n], \Omega_a)$. Then, we will show that $\text{Gr}([S^n], \Omega_a^{\text{round}}) = 2\pi a$.

It is a general fact that, if there is an inequality between Riemannian metric $g_0 \leq g_1$, then the unit codisk bundle determined by g_1 contains the unit codisk bundle determined by g_0 . An easy computation shows that $g_a^{\text{round}} \leq g_a$.

Now we argue that $\text{Gr}([S^n], \Omega_a^{\text{round}}) = 2\pi a$. It is known that for the class of the point we have $\text{Gr}([\text{pt}]; \Omega_a^{\text{round}}) = 2\pi a$, i.e., for every $\epsilon > 0$ there is a symplectic embedding $e : B^{2n}(2\pi a - \epsilon) \rightarrow \Omega_a^{\text{round}}$. We can further assume that $e(0) = (e_1, 0)$, where $e_1 = (1, 0, \dots, 0) \in S^n$. To estimate the parametric Gromov width from below, it is enough to find a family $A : S^n \rightarrow \text{Symp}(\Omega_a^{\text{round}})$ which satisfies $A(q)(e_1, 0) = (q, 0)$. Indeed, such a family, together with an embedding e induces $A_e : S^n \rightarrow \mathfrak{B}(2\pi a - \epsilon, \Omega_a)$ by

$$A_e(q) := A(q) \circ e \in \mathfrak{B}(2\pi a - \epsilon, \Omega_a).$$

One way to think about Ω_a^{round} is

$$\Omega_a^{\text{round}} \cong \{q + ip \in \mathbb{C}^{n+1} \mid \|q\| = 1, \langle q, p \rangle = 0, \|p\| \leq a\}.$$

Each $A \in U(n+1)$ maps Ω_a^{round} to itself, and it is a symplectomorphism. Hence, it is enough to find a family $A : S^n \rightarrow U(n+1)$ of unitary matrices such that $A(q)(e_1, 0) = (q, 0)$. This is equivalent to finding n families of vectors $z_i(q) \in \mathbb{C}^{n+1}$ such that $\langle q, z_1(q), \dots, z_n(q) \rangle$ is a unitary basis of $\mathbb{C}^{n+1}$ for every $q \in S^n$. Such a family of vectors can be found since the complexified tangent bundle $TS^n \otimes \mathbb{C}$ is trivial¹. This completes the proof of Theorem 2.

2.2.4. Diagonal actions on open books. If one has an $\mathbb{R}/\mathbb{Z}$ -action on the page V , we have the induced diagonal $\mathbb{R}/\mathbb{Z}$ action on the open book M . The action class we used in §2.2.2 is obtained from the trivial action on V .

The goal in this section is to prove that the action classes $A_{\pm} : M \times \mathbb{R}/\mathbb{Z} \rightarrow M$ lie in the image of the BV operator Δ , in $H(\Lambda_c)$ for specific c . To this end, it is useful to give a description of the open book M as a subset of $V \times \mathbb{C}$.

Let $(V, \partial V)$ be a smooth manifold with boundary, and let $f : V \rightarrow [0, 1]$ be a smooth function such that $f(1) = \partial V$ is a regular level set. Define:

$$\text{OB}(V) = \{(v, z) \in V \times \mathbb{C} : f(v) + |z|^2 = 1\}.$$

The bordism classes we consider are related to a circle action of the following form: suppose that $(t, v) \mapsto \zeta_t(v)$ is a circle action on V which preserves the level sets of f , and consider the resulting circle action:

$$(t, v, z) \mapsto (\zeta_t(v), e^{2\pi i t} z)$$

which acts on $\text{OB}(V)$.

First, we deform the construction of $\text{OB}(V)$ as follows. Define:

$$M = \{(v, z) \in V \times \mathbb{C} : f(v) + \rho(|z|)^2 = 1\},$$

where ρ is non-decreasing, $\rho(x) = 0$ for $x \leq 1$, $\rho(x) = 2$ for $x \geq 2$, and $\rho'(x) > 0$ whenever $\rho(x) \in (0, 1]$. Then:

Lemma 21. *The identity map $\text{id} : \text{OB}(V) \rightarrow \text{OB}(V)$ is cobordant to the smooth map $R : M \rightarrow \text{OB}(V)$ given by:*

$$R(v, z) = (v, g(z)z) \text{ where } g(z) = |z|^{-1}\rho(|z|).$$

which is well-defined; i.e., $f(v) + |g(z)z|^2 = 1$.

Proof. We define the cobordism explicitly:

$$Q = \{(v, z, s) \in V \times \mathbb{C} \times [0, 1] : ((1-s)|z| + s\rho(|z|))^2 + f(v) = 1\}.$$

A short computation shows that Q is cut transversally. Moreover, the boundary $\partial Q = \{s = 0, 1\}$ is identified with $\text{OB}(V) \sqcup M$ in the obvious way. Define a map:

$$S : Q \rightarrow \text{OB}(V)$$

by the formula:

$$S(v, z, s) = (v, g_s(z)z) \text{ where } g_s(z) = (1-s) + s|z|^{-1}\rho(|z|),$$

which is a smooth function. This provides the desired cobordism, since $S(v, z, 0)$ is identified with the identity map and $S(v, z, 1)$ with $R(v, z)$. $\square$

¹This is a well-known fact, a fairly simple way to see this is from the existence of a Lagrangian immersion $S^n = \{(x, y) \in \mathbb{R}^n \times \mathbb{R} \mid \|x\|^2 + y^2 = 1\} \mapsto (1 + iy)x \in \mathbb{C}^n$ (see [MS17, Example 13.2.4]).

Next, we consider the action class:

$$A : \mathbb{R}/\mathbb{Z} \times \text{OB}(V) \rightarrow \text{OB}(V) \text{ given by } (t, v, z) \mapsto (\zeta_t(v), e^{2\pi it} z).$$

Our goal is prove that $[A]$ lies in the image of Δ .

It follows easily from Lemma 21 that A is cobordant to the class:

$$B : \mathbb{R}/\mathbb{Z} \times M \rightarrow \text{OB}(V) \text{ given by } (t, v, z) \mapsto (\zeta_t(v), e^{2\pi it} g(z)z),$$

indeed, one can re-use the same cobordism Q , and A, B extend to the cobordism.

Remark. During the cobordism, the lengths of loops never exceeds the maximum length of loops in the class of A ; indeed, each loop appearing is exactly one of the loops appearing in the family A .

Thus, in order to prove that $[A]$ lies in the image of Δ , it is sufficient to prove that $[B]$ lies in the image of Δ .

The next step is to pick a smooth non-increasing cut-off function $\psi : \mathbb{R} \rightarrow [0, 2]$ such that $\psi(x) = 2$ for $x \leq 0$, $\psi(x) = 0$ for $x \geq 1$, and $\psi'(x) \neq 0$ if $\psi(x) \in (0, 1]$. Then we define:

$$N = \{(v, z) \in V \times \mathbb{C} : f(v) + \psi(|z|)^2 + \rho(|z|)^2 = 1\}.$$

It is important to note that $\psi(|z|)$ and $\rho(|z|)$ are supported in different regions (the former in $|z| \leq 1$ and the latter in $|z| \geq 1$). A quick computation shows that N is cut transversally.

Lemma 22. *The map $B : \mathbb{R}/\mathbb{Z} \times M \rightarrow \text{OB}(V)$ is cobordant to the map:*

$$C_\vartheta : \mathbb{R}/\mathbb{Z} \times N \rightarrow \text{OB}(V)$$

given by:

$$C_\vartheta(t, v, z) = \begin{cases} (\zeta_t(v), e^{2\pi i \vartheta} \psi(|z|)) & \text{if } |z| \leq 1, \\ (\zeta_t(v), e^{2\pi i t} g(|z|)z) & \text{if } |z| \geq 1, \end{cases}$$

where $g(|z|) = |z|^{-1} \rho(|z|)$.

Proof. It is clear from the construction that the map is smooth and agrees on the overlap (since $\psi(x)$ and $\rho(x)$ vanish to all orders when $x = 1$). Moreover, the map is well-defined and valued in $\text{OB}(V)$, as can be checked by a direct computation. It remains to prove the map is cobordant to B .

To see this, we define the cobordism explicitly:

$$X = \{(v, z, s) \in V \times \mathbb{C} \times [0, 1] : f(v) + s^2 \psi(|z|)^2 + \rho(|z|)^2 = 1\}.$$

It is not hard to see that this is cut transversally; when $s = 0$, it is already known to be transverse, and when $s \neq 0$, the derivative with respect to s is non-zero where $\psi(|z|) \neq 0$; elsewhere it is known to be transverse by the same argument used to prove M was cut transversally.

Now we define the map $T_\vartheta : \mathbb{R}/\mathbb{Z} \times X \rightarrow \text{OB}(V)$ by the formula:

$$T_\vartheta(t, v, z, s) = \begin{cases} (\zeta_t(v), e^{2\pi i \vartheta} s \psi(|z|)) & \text{if } |z| \leq 1, \\ (\zeta_t(v), e^{2\pi i t} g(|z|)z) & \text{if } |z| \geq 1, \end{cases}$$

where $g(|z|) = |z|^{-1} \rho(|z|)$. This map is easily seen to be a cobordism between B and C_ϑ , as desired. $\square$

Remark. Unlike the cobordism between A and B , the cobordism between B and C_ϑ introduces new loops; namely, it introduces the loops

$$(7) \quad t \mapsto (\zeta_t(v), e^{2\pi i \vartheta} r) \text{ where } f(v) = 1 - r^2$$

Thus, if c is greater than the maximum length of loops which appear in the class A , and in the family (7) (for any chosen ϑ), then the cobordism from A to C_ϑ occurs within $H(\Lambda_c)$. However, if the action on the page V is trivial, these loops are constant and therefore do not increase the length threshold; this observation is relevant to obtaining the stated bound in Lemma 20.

Our final result in this subsection is that $[C_\vartheta] = \Delta([D_\vartheta])$ for a family of loops D_ϑ . To show this, define:

$$P = \{(v, r) \in V \times (0, \infty) : f(v) + \psi(r)^2 + \rho(r)^2 = 1\} \subset N;$$

essentially since $\psi(0) = 2$ is larger than 1, it is easy to see that P is a closed manifold. Then we define: $D_\vartheta : \mathbb{R}/\mathbb{Z} \times P \rightarrow \text{OB}(V)$ by:

$$D_\vartheta(t, v, r) = \begin{cases} (\zeta_t(v), e^{2\pi i \vartheta} s\psi(r)) & \text{if } |z| \leq 1, \\ (\zeta_t(v), e^{2\pi i t} g(r)r) & \text{if } |z| \geq 1. \end{cases}$$

Lemma 23. *It holds that $\Delta([D_\vartheta]) = [C_\vartheta]$ within $H(\Lambda_c)$ provided $C_\vartheta \in Z(\Lambda_c)$.*

Proof. Observe that $\Delta(D_\vartheta)$ is represented by $D'_\vartheta : \mathbb{R}/\mathbb{Z} \times \mathbb{R}/\mathbb{Z} \times P \rightarrow \text{OB}(V)$ given by:

$$D'_\vartheta(t, \theta, v, r) = \begin{cases} (\zeta_{t-\theta}(v), e^{2\pi i \vartheta} s\psi(r)) & \text{if } |z| \leq 1, \\ (\zeta_{t-\theta}(v), e^{2\pi i(t-\theta)} g(r)r) & \text{if } |z| \geq 1. \end{cases}$$

Now consider the diffeomorphism $h : \mathbb{R}/\mathbb{Z} \times P \rightarrow N$ given by:

$$h(\theta, v, r) = (\zeta_{-\theta}(v), e^{-2\pi i \theta} r).$$

Under this diffeomorphism we have:

$$C_\vartheta(t, h(\theta, v, r)) = \begin{cases} (\zeta_{t-\theta}(v), e^{2\pi i \vartheta} \psi(r)) & \text{if } |z| \leq 1, \\ (\zeta_{t-\theta}(v), e^{2\pi i(t-\theta)} g(r)r) & \text{if } |z| \geq 1, \end{cases}$$

so $[C_\vartheta] = [D'_\vartheta] = \Delta([D_\vartheta])$. Note that two maps $\mathbb{R}/\mathbb{Z} \times P_i \rightarrow \text{OB}(V)$, $i = 0, 1$, differing by a diffeomorphism $P_0 \rightarrow P_1$ are cobordant in a trivial sense; Such a cobordism does not change any lengths of loops, so the cobordism from C_ϑ to $\Delta(D_\vartheta)$ occurs within $H(\Lambda_c)$. $\square$

2.2.5. Proof of Theorem 3. We begin with an auxiliary lemma about the Hopf flow. Consider the Hopf $\mathbb{R}/\mathbb{Z}$ -action on $S^3 \subset \mathbb{C}^2$ given by:

$$\zeta_t(z_1, z_2) = (e^{2\pi i t} z_1, e^{2\pi i t} z_2).$$

Then:

Lemma 24. *The Hopf action is homotopic to the trivial $\mathbb{R}/\mathbb{Z}$ -action through the loops of length $\leq 2\pi$.*

Proof. After conjugating the Hopf action with $\psi(z_1, z_2) = (z_1, \bar{z}_2)$ we get:

$$t \cdot (z_1, z_2) = (e^{2\pi it} z_1, e^{-2\pi it} z_2).$$

This $\mathbb{R}/\mathbb{Z}$ -action is the restriction of the $SU(2) \cong S^3$ action to the unit circle $\{(z, 0) : z \in S^1\} \subset S^3$. Consider a homotopy $h : \mathbb{R}/\mathbb{Z} \times [0, 1] \rightarrow S^3$ through loops of length $\leq 2\pi$ such that $h(t, 0) = (e^{2\pi it}, 0)$, $h(t, 1) = (1, 0)$. The homotopy h induces a homotopy:

$$H(z_1, z_2, t, s) = h(t, s) \cdot (z_1, z_2),$$

which satisfies the conclusion of the Lemma. $\square$

Similarly to §2.2.3 one gets that the radius a codisk bundle Ω_a^{round} of the round metric is a subset of Ω_a , where Ω_a is the unit codisk bundle determined by the embedding of:

$$\{x_0^2 + x_1^2 + \cdots + a^{-2}(x_{n-3}^2 + x_{n-2}^2 + x_{n-1}^2 + x_n^2) = 1\} \subset \mathbb{R}^{n+1}.$$

Since $Gr([S^n]; \Omega_a^{\text{round}}) = 2\pi a$ we get:

$$Gr([\text{pt}]; \Omega_a) \geq Gr([S^n]; \Omega_a) \geq 2\pi a.$$

To prove the upper bound, one appeals to the results of §2.2.4 with the diagonal $\mathbb{R}/\mathbb{Z}$ -action given by:

$$(8) \quad \zeta_t(x_0, \dots, x_{n-4}, z_1, z_2) = (x_0, \dots, x_{n-4}, e^{2\pi it} z_1, e^{2\pi it} z_2).$$

One sees S^n as an open book with page $V \subset \mathbb{D}^{n-3} \times \mathbb{C}$ given by:

$$V = \{x_0^2 + \cdots + x_{n-4}^2 + a^{-2}|z_1|^2 \leq 1\}.$$

Then we have:

$$S^n = \text{OB}(V) = \{(x, z_1, z_2) \in V \times \mathbb{C} \mid f(x, z_1) + a^{-2}|z_2|^2 = 1\},$$

where $f(x, z) = x_0^2 + \cdots + x_{n-4}^2 + a^{-2}|z|^2$. An $\mathbb{R}/\mathbb{Z}$ action on V is given by the rotation of the $\mathbb{C}$ coordinate, and it preserves the levels of f , hence from Lemma 23, the induced action class A is in the image of Δ ; in brief, there exists a class B such that $[A] = \Delta[B] \in H(\Lambda_{2\pi a})$. The stated length threshold of $2\pi a$ is obtained by inspection of the length thresholds in 2.2.4.

Finally, the contractibility of the Hopf flow on S^3 implies that the $\mathbb{R}/\mathbb{Z}$ -action (8) is homotopic to the trivial $\mathbb{R}/\mathbb{Z}$ action, hence $[A] = [S^n] \in H(\Lambda_{2\pi a})$. Then Theorem 12 applied to the class $\beta = [T^*S^n]$, $\alpha_1 = [B]$ (and $\alpha_2 = [S^n]$) produces $Gr([\text{pt}]; \Omega_a) \leq 2\pi a$, as desired.

2.2.6. Proof of Theorem 5 and Proposition 1. The proof of Theorem 5 follows very similar lines to the proofs of Theorem 4 in the case when $\partial V = \emptyset$. Indeed, one can define $B_-^d : V \times T^{d-1} \times \mathbb{R}/\mathbb{Z} \rightarrow V \times T^d$ by:

$$B_-(v, x, t) = (v, x_1, \dots, x_{d-1}, -t),$$

and, for $k < d$, define $B_+^k : V \times T^{d-k-1} \times \mathbb{R}/\mathbb{Z} \rightarrow V \times T^d$ by

$$B_+^k(x, t) = (v, 0, \dots, 0, x_1, \dots, x_{d-k-1}, t).$$

By Lemma 17 and 19, it holds that:

$$\Delta[B_-^d] * \Delta[B_+^k] = [V \times T^{d-k}] \text{ in } \Lambda_{E_- + E_+^k}$$

where $[V \times T^{d-k}]$ is the class of $(v, x) \mapsto (v, 0, \dots, 0, x_1, \dots, x_{d-k})$.

Since $[V \times T^{d-k}]$ has non-zero homological intersection number with the class $[T^k]$ represented by the map $f : T^k \rightarrow V \times T^d$ given by $x \mapsto (v_0, x_1, \dots, x_k)$, we conclude from Theorem 12 that $\text{Gr}([T^k], \Omega) \leq E_- + E_+^k$, as desired.

In the remainder of this subsection, we will explain how Theorem 5 can be used to recover the camel theorem (Proposition 1).

Recall the set-up of Proposition 1; we had defined the symplectic manifold:

$$X = (W \setminus \{x_1 = 0\}) \cup \{x_n^2 + y_n^2 \leq \pi^{-1}\epsilon\},$$

where $W = \mathbb{R}/\mathbb{Z} \times \mathbb{R}^{2n-1}$, for $n > 1$, and claimed that $\text{Gr}(f, X) \leq \epsilon$ provided that $f : \mathbb{R}/\mathbb{Z} \rightarrow X$ has non-trivial winding number.

We argue by contradiction: suppose there is a family of balls:

$$F : \mathbb{R}/\mathbb{Z} \times B(a) \rightarrow X$$

with $a > \epsilon$, such that $t \mapsto F(t, 0) = (t, 0, \dots, 0)$ (here we appeal to the fact that the winding number classifies the free homotopy class of a loop). It is convenient to fix some small parameter $\delta > 0$. We begin with an auxiliary lemma, which converts X into a space more compatible with Theorem 5.

Lemma 25. *There exists a symplectic isotopy $\psi_s : X \rightarrow W$ so that, abbreviating $y_i = y_i \circ \psi_1 \circ F(t, z)$ and $x_i = x_i \circ \psi_1 \circ F(t, z)$, we have:*

- (1) $\psi_1 \circ F(t, 0) = (t, 0, \dots, 0)$,
- (2) $x_i \in (-1/2, 1/2)$ for $i = 2, 3, \dots, n$,
- (3) $y_n > -\epsilon/2 - 2\delta$,
- (4) $x_1 = 0 \implies y_n < \epsilon/2 + \delta$.

Proof. Condition (1) is already satisfied. Condition (2) is easily satisfied using the transformation $(x_i, y_i) \mapsto (a^{-1}x_i, ay_i)$ where $a > \max 2|x_i|$, for each $i = 2, \dots, n$, where the maximum is taken over the image of the family of balls. Let us denote ρ_a the result of this transformation. It is convenient to also take $\pi a^2 > 4\epsilon$.

Now we focus on attaining (4). The transformation used to achieve (2) implies that:

$$x_1 \circ \rho_a(F(t, z)) = 0 \implies \rho_a(F(t, z)) \subset \{a^2 x_n^2 + a^{-2} y_n^2 \leq \pi^{-1}\epsilon\} = R.$$

Since the projection R_n of R to x_n, y_n -plane is an ellipse of area ϵ and is contained in $S = \{x_n \in (-1/2, 1/2)\} \subset \mathbb{R}^2$, one can find a symplectic isotopy φ_s of the strip S so φ_1 sends R_n into the rectangle:

$$(-1/2, 1/2) \times (-\epsilon/2 - \delta, \epsilon/2 + \delta),$$

since the rectangle has strictly larger area than the ellipse. One can arrange that $\varphi_1(0) = 0$. The isotopy φ_s extends as a product isotopy $\text{id} \times \varphi_s$ to all of W . Then $\varphi_1 \circ \rho_a \circ F$ satisfies (1), (2), (4). It remains only to ensure (3) holds.

For this step, we consider the isotopy generated by $H = f(x_1)x_n$, where $f(0) = 0$ and $f'(0) = 0$, and $f(x_1) = 1$ outside a small neighborhood U of $x_1 = 0$. The Hamiltonian vector field generated by this function is:

$$f(x_1)\partial_{y_n} + f'(x_1)x_n\partial_{y_1}$$

This vector field equals ∂_{y_n} outside of U , and equals 0 when $x_1 = 0$. We pick the neighborhood U small enough that:

$$\varphi_1 \circ \rho_a \circ x_1 \in U \implies \varphi_1 \circ \rho_a \circ y_n \geq -\epsilon/2 - 2\delta,$$

which can be achieved by a simple compactness argument, since we already know (4) holds.

The flow by X_H increases the y_n coordinates when $x_1 \notin U$; thus flowing long enough will therefore satisfy condition (3) everywhere. Denoting by η_1 the long time-flow by X_H , we set $\psi_1 = \eta_1 \circ \varphi_1 \circ \rho_a$ to complete the proof. $\square$

Henceforth replace F by $\psi_1 \circ F$.

Proof of Proposition 1. Introduce the domain $\Omega \subset T^*T^n$ determined by the conditions:

- (a) $p_n \geq -\epsilon/2 - \delta$,
- (b) $q_1 = 0 \implies p_n \leq \epsilon/2 + 2\delta$,

using canonical coordinates $(p_1, q_1, \dots, p_n, q_n)$. Define the map:

$$(x_1, y_1, \dots, x_n, y_n) \mapsto T^*T^n \text{ by } x_i = q_i \text{ and } y_i = -p_i.$$

The family of balls F projects to a family $F : \mathbb{R}/\mathbb{Z} \times B(a) \rightarrow T^*T^n$ which restricts to embeddings $z \mapsto F(t, z)$, since we have arranged that F is valued in the region $x_i \in (-1/2, 1/2)$ for $i = 2, \dots, n$. Moreover, by construction, F is actually a family of balls in Ω . Thus the existence of F implies $\text{Gr}(f, \Omega) \geq a$, where $q \circ f(t) = (t, 0, \dots, 0)$.

On the other hand, Theorem 5 with $k = 1$ and $d = n$ yields:

$$\text{Gr}(f, \Omega) \leq \epsilon + 3\delta,$$

as can be seen by approximating the domain Ω satisfying (a) and (b) from within by smooth fiberwise starshaped domains. Thus we conclude $a \leq \epsilon + 3\delta$. Since δ could be taken arbitrarily small, we have $a \leq \epsilon$, as desired. $\square$

2.2.7. Proof of Theorem 6. The proof is straightforward application of Theorem 12 using a simple string topology computation: if q is a loop such that $q^*T\Sigma$ is non-orientable, then $\Delta(q) * \Delta(\bar{q}) = \text{pt}$ in $H(\Lambda_E)$, where pt denotes the class represented by a single constant loop. Since $\text{pt} = i(T^*M_{\text{pt}})$, and T^*M_{pt} has non-zero homological intersection number with the zero section, the desired result $\text{Gr}([\Sigma], \Omega) \leq E$ follows. $\square$

3. The Floer cohomology persistence module

In this section we explain the aspects of our paper pertaining to the Floer cohomology persistence module, with the goal of proving Theorem 7.

3.1. Hamiltonian functions and isotopies.

3.1.1. Class of Hamiltonian functions. As in §1, throughout we fix a Liouville domain $\bar{\Omega}$, and denote by W its completion. The structure of W as a completion yields a distinguished “convex end” which is symplectomorphic to the positive half of the symplectization of $\partial\Omega$ with the contact structure $\ker(\lambda|_{\partial\Omega})$. This yields a function r which is one-homogeneous with respect to the Liouville flow in the convex end, and which satisfies $\partial\Omega = \{r = 1\}$. It is well-known that the Hamiltonian vector field X_r is equivariant with respect to the Liouville flow, and restricts to $\partial\Omega$ as the Reeb vector field for the contact form $\lambda|_{\partial\Omega}$.

Let us introduce the notation $\Omega(r_0) = \{r < r_0\}$, so $\Omega = \Omega(1)$.

As in §1.2, this distinguishes a class of Hamiltonian functions: define $\mathcal{H}$ to be those smooth functions $H : W \rightarrow \mathbb{R}$ such that $H = ar$ holds when $r \geq r_0$, for some r_0 (which depends on H).

3.1.2. Smooth families of Hamiltonians. One says that a family $H_\tau \in \mathcal{H}$, where τ is valued in a smooth manifold T (potentially with boundary and corners) is *smooth* provided:

- (1) the map $(\tau, w) \in T \times W \mapsto H_\tau(w)$ is smooth,
- (2) $H_\tau(w) = c_\tau r$ for $r \geq r_0$, for some r_0 which depends continuously on τ , and where the slope c_τ varies smoothly with τ .

3.2. Hamiltonian connections on surfaces. The Floer theory PDEs used in this text are based on the notion of a *Hamiltonian connection* as described in [MS12, §8]; see also [AAC23, §2.2.3] which works in a similar context to the present paper.

3.2.1. Domains. Let us agree that a *domain* is a compact smooth surface Σ with boundary $\partial\Sigma$ and the data of:

- (1) a complex structure j , namely a smooth section of $\text{End}(T\Sigma)$ whose square is -1 , and,
- (2) two collections $\Gamma_\pm$ of j -holomorphic embeddings of a closed disk $D(1) \rightarrow \Sigma$ whose images are mutually disjoint, and are also disjoint from $\partial\Sigma$.

A *smooth family of domains*, parametrized by $\sigma \in S$, is simply the parametric version of (1) and (2); more prosaically, this means that one considers the product $S \times \Sigma$, with projection π onto Σ , and picks:

- (1') a smooth section j of $\text{End}(\pi^*T\Sigma)$ whose square is -1 , and
- (2') two collections $\Gamma_\pm$ of smooth embeddings $S \times D(1) \rightarrow S \times \Sigma$ which sends $\{\sigma\} \times D(1)$ into $\{\sigma\} \times \Sigma$ and which satisfies (2) when restricted to the disk $\{\sigma\} \times D(1)$, for the appropriate complex structure.

A prototypical example of such a family is given by $\Sigma = \mathbb{CP}^1$, and S the product of k -copies of $T\mathbb{CP}^1$ with an appropriate closed set removed. To each collection of k non-zero tangent vectors, there exist k uniquely determined biholomorphisms of $\mathbb{CP}^1$ which send $(0, 1) \in T\mathbb{CP}^1$ to the chosen tangent vectors. These biholomorphisms take $D(1) \subset \mathbb{CP}^1$ onto k biholomorphically

embedded disks. There is a closed subset of S for which these k disks intersect; after removing this closed subset, one has a smooth family.

3.2.2. Remark on notation and terminology. If either j or $\Gamma_{\pm}$ is not germane to the discussion, we will suppress it from the notation, and simply refer to $S \times \Sigma$ as a family of domains.

It is also convenient to sometimes refer to $\Gamma_{\pm}$ as punctures, in which case we forget the holomorphic embedding and consider only the evaluation at the center of the disk. This produces a submanifold $\Gamma_{\pm} \subset S \times \Sigma$ of codimension 2 which is a disjoint union of sections of $S \times \Sigma \rightarrow S$; we call a connected component $\zeta \subset \Gamma_{\pm}$ a puncture.

Frequently it is expedient to define Σ as the already punctured surface: $\Sigma = \mathbb{R} \times \mathbb{R}/\mathbb{Z}$, $\Sigma = \mathbb{C}$, or $\Sigma = \mathbb{C} \setminus \{z_1, z_2\}$. In these cases, we should understand the underlying closed surface to be the Riemann sphere $\mathbb{CP}^1$.

At other times, we will refer to $\Gamma_{\pm}$ as cylindrical ends, in which case we precompose the holomorphic embedding of the disk $D(1) \rightarrow \Sigma$ with the conformal reparametrization $(s, t) \mapsto e^{\mp 2\pi(s+it)}$. Note that the ends in Γ_+ are modelled on $[0, \infty) \times \mathbb{R}/\mathbb{Z}$ while those in Γ_- are modelled on $(-\infty, 0] \times \mathbb{R}/\mathbb{Z}$. We will sometimes appeal to cylindrical coordinates near a puncture $\zeta \subset \Gamma_{\pm}$, and this is to be understood in the context of this remark.

3.2.3. Connection 1-forms. Let $(S \times \Sigma, \Gamma_{\pm})$ be a smooth family of domains. A connection 1-form is a singular 1-form $\mathfrak{a}$ on $S \times \Sigma \times W$ such that:

- (1) above a compact coordinate chart $(\sigma, z = x + iy)$ on $S \times \Sigma$ disjoint from the punctures, $\mathfrak{a} = H_{\sigma, x, y} dx + K_{\sigma, x, y} dy$, where H, K are smooth families in $\mathcal{H}$;
- (2) for each puncture $\zeta \subset \Gamma_{\pm}$, one has $\mathfrak{a} = H_{\zeta, t} dt$ near ζ , where $H_{\zeta, t}$ is a smooth family in $\mathcal{H}$ on $\mathbb{R}/\mathbb{Z}$; this uses the appropriate cylindrical coordinates near ζ . Moreover, the Hamiltonian isotopy generated by $H_{\zeta, t}$ has non-degenerate 1-periodic orbits.

One should think that $\mathfrak{a}$ has prescribed singularities near the punctures, with a requirement on the dynamics of the induced asymptotic Hamiltonian system. Note that, because the asymptotic Hamiltonian system commutes with the Liouville flow outside of a compact set, the requirement that the orbits are non-degenerate forces all of the orbits to remain in a compact set, and hence there are only finitely many orbits.

Remark. One slightly subtle requirement which plays a role in our proof of the maximum principle, Proposition 28, is the following: the connection one form should appear in the form $\mathfrak{a} = H_{\zeta, t} dt$, where $H_{\zeta, t}$ is non-degenerate, on cylindrical ends, *and on some number of finite length cylinders*, such that the domain obtained by removing the cylindrical ends and the finite length cylinders is contained in a fixed compact subset of Σ . This is because our maximum principle is based on confining Floer cylinders, and then bounding the distance of other points to a region where $\mathfrak{a} = H_{\zeta, t} dt$. This set-up is relevant when, for instance, gluing together two continuation cylinders.

3.2.4. Perturbation 1-forms. Connection 1-forms are used to define a certain PDE. It is well-understood that Floer theory relies on the moduli spaces of solutions to this PDE being transversally cut out. We introduce in this section a perturbation term $\mathfrak{p}$ to obtain the requisite transversality.

Given a smooth family of domains $(S \times \Sigma, \Gamma_{\pm})$, a perturbation 1-form is a smooth 1-form $\mathfrak{p}$ on $S \times \Sigma \times W$ such that:

- (1) above a compact coordinate chart $(\sigma, z = x + iy)$ on $S \times \Sigma$ disjoint from the punctures, $\mathfrak{p} = h_{\sigma,x,y}dx + k_{\sigma,x,y}dy$, where h, k are smooth and uniformly bounded in C^1 families of functions $W \rightarrow \mathbb{R}$, and,
- (2) $\mathfrak{p}$ vanishes above a neighborhood of the punctures $\Gamma_{\pm}$.

The boundedness of the functions appearing in $\mathfrak{p}$ is used in an essential way when establishing a priori estimates on the energy integrals in §3.3.5; see the comment at the end of §3.3.7.

3.2.5. Families of connection and perturbation 1-forms. Let $S \times \Sigma$ be a smooth family of domains. One can pull this back to $T \times S \times \Sigma$ for any manifold T to obtain a new family of domains. A connection or perturbation 1-form on $T \times S \times \Sigma \times W$ is, by definition, a smooth family of connection or perturbation 1-forms on $S \times \Sigma \times W$ parametrized by T .

In this sense, we may speak of homotopies of 1-forms by setting $T = [0, 1]$.

3.2.6. Connection associated to a 1-form. Let $\mathfrak{a}, \mathfrak{p}$ be connection and perturbation 1-forms on $S \times \Sigma \times W$. Define:

$$\mathfrak{H} = TW^{\perp\Omega} \text{ where } \Omega = \text{pr}_W^* \omega - d\mathfrak{a} - d\mathfrak{p}.$$

Then $\mathfrak{H}$ is an Ehresmann connection on $S \times \Sigma \times W \rightarrow S \times \Sigma$.

The coordinate distribution TS is contained in $\mathfrak{H}$. On the other hand, the horizontal lifts of $\partial_x, \partial_y \subset T\Sigma$ are given by:

$$\partial_x^{\mathfrak{H}} = \partial_x + X_{H+h} \text{ and } \partial_y^{\mathfrak{H}} = \partial_y + X_{K+k},$$

where $\mathfrak{a} + \mathfrak{p} = (H + h)dx + (K + k)dy$.

The pullback of $\mathfrak{H}$ to $\{\sigma\} \times \Sigma \times W \rightarrow \{\sigma\} \times \Sigma$ will be denoted by $\mathfrak{H}_{\sigma}$ in the sequel, and should be considered as a family of connections on $\Sigma \times W \rightarrow \Sigma$.

3.3. A general form of Floer's equation. In this section, we will explain a general form of Floer's equation which will specialize to the various cases used in the proofs of our main results.

3.3.1. Almost complex structures. Let $S \times \Sigma$ be a family of domains, with punctures $\Gamma_{\pm}$. An almost complex structure on $S \times \Sigma \times W$ is a section J of the pullback bundle $\text{pr}_W^* \text{End}(TW) \rightarrow S \times \Sigma \times W$ satisfying $J^2 = -1$.

It is convenient to think of this as a family $J_{\sigma,z}$ of almost complex structures on W parametrized by $(\sigma, z) \in S \times \Sigma$. We require that:

- (1) For (σ, z) near a puncture $\zeta \in \Gamma_{\pm}$, we have $J_{\sigma,z} = J_{\zeta}$ for an almost complex structure J_{ζ} on W .
- (2) J is ω -tame, i.e., $\omega(v, Jv)$ is a positive quadratic form for $v \in TW$.

- (3) Each $J_{\sigma,z}$ is invariant under the Liouville flow outside of $\Omega \subset W$.

3.3.2. Lagrangian boundary conditions. Suppose that the surface Σ has boundary $\partial\Sigma$. For the purposes of this text, a *Lagrangian boundary condition* is a smoothly varying family of Lagrangian submanifolds $L_{\sigma,z} \subset W$ where $(\sigma, z) \in S \times \partial\Sigma$. Here “smoothly varying” means that:

$$L = \{(\sigma, z, p) : p \in L_{\sigma,z}\} \subset S \times \partial\Sigma \times W$$

should be a smooth submanifold. A submanifold is required to be properly embedded and without boundary. We require two additional properties:

- (1) $L_{\sigma,z}$ is weakly exact, i.e., ω vanishes on disks with boundary on $L_{\sigma,z}$.
- (2) outside of Ω , the Liouville vector field is tangent to $L_{\sigma,z}$.

In this text, Lagrangian boundary conditions will only be used in the comparison between string topology of M and Floer cohomology of T^*M , and we will have that $L_{\sigma,z} = T^*M_{q(\sigma,z)}$ where $q(\sigma, z)$ depends smoothly on σ, z .

3.3.3. Floer’s equation for a Hamiltonian connection. The data required to formulate Floer’s equation is:

- (1) a family of domains $(S \times \Sigma, \Gamma_{\pm}, j)$ as in §3.2.1,
- (2) a connection 1-form $\mathfrak{a}$ and perturbation 1-form $\mathfrak{p}$ on $S \times \Sigma \times W$ as in §3.2.3 and §3.2.4,
- (3) an almost complex structure J on $S \times \Sigma \times W$ as in §3.3.1,
- (4) a Lagrangian boundary condition L for $S \times \Sigma \times W$ as in §3.3.2.

The data of (2) produces an Ehresmann connection $\mathfrak{H}$ for the fiber bundle $S \times \Sigma \times W \rightarrow S \times \Sigma$. For each $\sigma \in S$, we consider the Ehresmann connection $\mathfrak{H}_{\sigma}$ obtained by pullback to $\{\sigma\} \times \Sigma \times W \rightarrow \{\sigma\} \times \Sigma$. Then $\mathfrak{H}_{\sigma}$ is identified with $T\Sigma$ via the projection map, and we use (1) and (3) to define:

$$\tilde{J}_{\sigma,z,w} = \begin{bmatrix} J_{\sigma,z,w} & 0 \\ 0 & j_{\sigma,z} \end{bmatrix} \text{ on } TW_w \oplus (\mathfrak{H}_{\sigma})_{z,w} = T(\Sigma \times W)_{z,w}.$$

This is considered as a family of almost complex structures on $\Sigma \times W$.

A pair (σ, u) solves Floer’s equation for this data if:

$$\begin{cases} u : \Sigma \rightarrow W \text{ is a smooth map,} \\ \text{the section } z \mapsto (z, u(z)) \text{ is } (j_{\sigma}, \tilde{J}_{\sigma})\text{-holomorphic,} \\ u(z) \in L_{\sigma,z} \text{ for all } z \in \partial\Sigma. \end{cases}$$

In the sequel, we will say that such (σ, u) solves §3.3.3.

3.3.4. Regularity. We assume the reader is familiar with the standard elliptic regularity results, in particular, the following fact: if (σ_n, u_n) is a sequence of solutions to §3.3.3 and:

- (1) σ_n converges to $\sigma_{\infty} \in S$,
- (2) the image $u_n(\Sigma)$ remains in a compact set,
- (3) the first derivatives of $u_n : \Sigma \rightarrow W$ are bounded on compact subsets,

then a subsequence of u_n converges in the C_{loc}^∞ -topology to a smooth map u_∞ such that $(\sigma_\infty, u_\infty)$ solves §3.3.3.

This result is true because of the general theory of pseudo-holomorphic curves in $\Sigma \times W$, and because the submanifold $L_\sigma \subset \Sigma \times W$ in §3.3.2 is totally real. We refer the reader to [MS12, §B] for further details.

To establish (2) and (3) we rely on a priori estimates on the energy integral.

3.3.5. Energy integral. Let $u : \Sigma \rightarrow W$ be a smooth map, and let $\mathfrak{H}$ be an Ehresmann connection on $\Sigma \times W$. Consider the 2-form on Σ defined by:

$$\omega(\Pi_{\mathfrak{H}}dv(-), \Pi_{\mathfrak{H}}dv(-)),$$

where $v(z) = (z, u(z))$ and $\Pi_{\mathfrak{H}} : T(\Sigma \times W) \rightarrow TW$ is the projection whose kernel is $\mathfrak{H}$. The integral of this 2-form is called the *energy* of the map u relative the connection $\mathfrak{H}$.

If (σ, u) solves §3.3.3, and $\mathfrak{H} = \mathfrak{H}_\sigma$, then the above 2-form is non-negative, since J is assumed to be ω -tame. We say that (σ, u) has *finite energy* provided the energy of u relative $\mathfrak{H}_\sigma$ is finite. The moduli space of all finite energy solutions of 3.3.3 plays a central role in this text, and might be denoted:

$$\mathcal{M}(S \times \Sigma \times W, \Gamma_\pm, j, \mathfrak{a}, \mathfrak{p}, J, L);$$

however, we will typically suppress arguments from the notation and simply use the symbol $\mathcal{M}$ as the required data can be inferred from the context.

The finite energy condition implies the following results:

Proposition 26 (Asymptotic convergence). *If (σ, u) solves §3.3.3 and has finite energy, then, in cylindrical coordinates $z = s + it$ near the puncture ζ , u is asymptotic to a solution of the s -independent equation:*

$$\partial_s u + J_\zeta(u)(\partial_t u - X_{\zeta,t}(u)) = 0,$$

where $J_{\sigma,z} = J_\zeta$ and $\mathfrak{a} = H_{\zeta,t}dt$ near ζ . Moreover, this asymptotic solution satisfies $\partial_s u = 0$, and is thus tracing out an orbit of the vector field $X_{\zeta,t}$.

Proof. See, e.g., [Sal97]. □

Proposition 27 (Gradient bound). *If (σ_n, u_n) is a sequence of solutions of §3.3.3, with uniformly bounded energy, and σ_n converges, then u_n has bounded first derivatives with respect to a Riemannian metric g on W which is invariant under the Liouville flow in the end, and a metric on Σ which is cylindrical near the punctures.*

Proof. Because W is symplectically aspherical, and the Lagrangians appearing as boundary conditions are weakly exact, the stated result follows from standard bubbling analysis. The argument is simplified by the fact that the almost complex structures J are invariant under the Liouville flow in the end. For details, we refer the reader to, e.g., [BC24, AAC23]. □

Proposition 28 (Maximum principle). *If (σ_n, u_n) is a sequence of solutions of §3.3.3 with uniformly bounded energy, and σ_n converges, then $u_n(\Sigma)$ remains in a compact subset of W .*

Proof. This is proved in [BC24] by a soft argument (using the Liouville-invariance of J and the Hamiltonian systems in 3.1); see also [AAC23]. We note that the proof does use the assumption in §3.2.3 that the Hamiltonian isotopy generated by the asymptotic system $H_{\zeta,t}$ has no 1-periodic orbits outside of a compact set. $\square$

3.3.6. Compactness and Floer differential cylinders. By combining Propositions 27 and 28 with the regularity result in §3.3.4, one sees that, in order to conclude that a sequence (σ_n, u_n) of solutions to §3.3.3 has a convergent subsequence in C_{loc}^∞ , it suffices to ensure that σ_n has a convergent subsequence and that u_n has uniformly bounded energy. Thus the main consideration for ensuring compactness is a priori energy estimates.

Even if (σ_n, u_n) has a subsequence such that $u_n \rightarrow u_\infty$ in C_{loc}^∞ , it is not true in general that u_n will converge uniformly to u_∞ ; it is possible that solutions of the s -independent equation from Proposition 26 break-off at the punctures. This phenomenon is a cornerstone of Floer theory, and we will not explain it any further except via the illustration in Figure 2.




FIGURE 2. Compactness up-to-breaking of Floer differential cylinders at the punctures; a sequence of solutions u_n defined on the pair-of-pants surface converging on compact subsets to a limit u_∞ , with the breaking of two solutions v_0, v_1 of the s -independent equation at the punctures.

Because of the breaking phenomenon, the s -independent equation in Proposition 26 plays a central role in Floer theory; we refer to this equation as *Floer's differential equation*, and it depends on the choice of an almost complex structure J_ζ and a time-dependent Hamiltonian vector field $X_{\eta,t}$ generated by a smooth family $H_t \in \mathcal{H}$ satisfying $H_{t+1} = H_t$.

3.3.7. A priori energy estimates. Consider a connection one-form $\mathfrak{a}$ and perturbation one-form $\mathfrak{p}$ on $S \times \Sigma \times W$. This induces an Ehresmann connection $\mathfrak{H}$ on $S \times \Sigma \times W \rightarrow S \times \Sigma$. It is well-understood that the curvature of the connection $\mathfrak{H}$ plays a role in a priori energy estimates for the energy integral for solutions to §3.3.3; see, e.g., [AAC23, §2.3.3]. In fact, there is an identity for the energy integral involving the two-form $\mathfrak{r}$ on $S \times \Sigma \times W$ called the curvature potential characterized by the property that, above a local coordinate patch $z = x + iy$ on Σ where:

$$\mathfrak{a} = H_{\sigma,x,y}dx + K_{\sigma,x,y}dy \text{ and } \mathfrak{p} = h_{\sigma,x,y}dx + k_{\sigma,x,y}dy$$

we have:

$$\mathfrak{r} = (\partial_x(K_{\sigma,x,y} + k_{\sigma,x,y}) - \partial_y(H_{\sigma,x,y} + h_{\sigma,x,y}) + \omega(X_{\sigma,s,t}^{H+h}, X_{\sigma,x,y}^{K+k}))dx \wedge dy,$$

where $X_{\sigma,x,y}^{H+h}, X_{\sigma,x,y}^{K+k}$ are the domain dependent Hamiltonian vector fields tangent to the fibers of $S \times \Sigma \times W \rightarrow S \times \Sigma$.

We note that if $\mathbf{p} = 0$ and $\mathbf{a} = H_t dt$ then $\mathbf{r} = 0$, and so the curvature potential $\mathbf{r}$ vanishes above a neighbourhood of the punctures $\Gamma \subset S \times \Sigma$.

It is convenient to phrase the energy identity in terms of the amount of energy a solution has on a compact subdomain $\Sigma' \subset \Sigma$. A priori energy estimates for the full energy are then obtained by taking a limit over larger and larger subdomains. The energy identity we will use is:

Lemma 29. *If (σ, u) solves §3.3.3 and $\Sigma' \subset \Sigma$ is a compact subdomain with boundary containing the support of $\mathbf{p}_\sigma$, then the energy of $u|_{\Sigma'}$ is equal to:*

$$E(u|_{\Sigma'}) = \int_{\Sigma'} u^* \omega + v^* \mathbf{r}_\sigma - \int_{\partial \Sigma'} v^* \mathbf{a}_\sigma,$$

where $\mathbf{a}_\sigma, \mathbf{p}_\sigma, \mathbf{r}_\sigma$ denote the pullbacks to $\{\sigma\} \times \Sigma \times W$ and $v(z) = (z, u(z))$ is considered as a section of $\Sigma \times W \rightarrow \Sigma$.

Proof. This is a standard computation; see, e.g., [AAC23, Lemma 2.3]. $\square$

Let us say that the data $\mathbf{a}, \mathbf{p}$ has *curvature bounded from above* provided:

$$\sup \left\{ \int_{\Sigma} v^* \mathbf{r} : v \text{ is a smooth section } \Sigma \rightarrow \{\sigma\} \times \Sigma \times W \text{ for some } \sigma \right\} < \infty,$$

where $\mathbf{r}$ is the curvature potential associated to $\mathbf{a}, \mathbf{p}$. We denote the quantity on the left by $\text{const}(\mathbf{r})$. This leads to the following a priori energy estimate:

Lemma 30. *If $\mathbf{a}, \mathbf{p}$ has curvature bounded from above, then any finite energy solution (σ, u) of §3.3.3 satisfies:*

$$E(u) \leq \omega(u) + \int_0^1 \sum_{\Gamma_-} H_{\zeta,t}(\gamma_\zeta(t)) - \sum_{\Gamma_+} H_{\zeta,t}(\gamma_\zeta(t)) dt - \int_{\partial \Sigma} v^* \mathbf{a}_\sigma + \text{const}(\mathbf{r}),$$

where u is asymptotic to the orbit $\gamma_\zeta(t)$ at the puncture ζ , and $\mathbf{a} = H_{\zeta,t} dt$ holds near the puncture ζ . In particular, the energy of finite energy solutions is uniformly bounded provided:

- (1) the symplectic area $\omega(u)$ is bounded above for solutions (σ, u) ,
- (2) the 1-periodic orbits of $H_{\zeta,t}$ are contained in a compact set,
- (3) the integrals of $\mathbf{a}_\sigma$ over sections $v : \partial \Sigma \rightarrow W$ valued in the Lagrangian boundary conditions L_σ are bounded below as σ varies in S .

Proof. This follows from the energy identity; see [AAC23, Lemma 2.4]. $\square$

We briefly explain why these three conditions can be assumed to hold in the context of our paper. First note (2) holds by fiat; it is part of our assumptions from §3.2.3 on the connection one-form. Second, in all of the constructions involving a non-empty boundary $\partial \Sigma$, we ensure that $\mathbf{a}$ appears in the form $H_{\sigma,t} dt$ where t is a coordinate on the boundary, $H_{\sigma,t}$ is non-negative outside of a compact set, and moreover that $\min H_{\sigma,t}$ is uniformly bounded from below as σ varies in S . It then follows easily that the integrals of $v^* \mathbf{a}_\sigma$ will be bounded from below, i.e., (3) will be satisfied.

Morally speaking, condition (1) holds because we assume the symplectic form is exact. There is some subtlety when $\partial \Sigma \neq \emptyset$, and we will explain how

to deal with this in §4, when we first consider moduli spaces with boundary conditions.

Remark. In general, condition (1) cannot be established, and the usual way one deals with this is via the introduction of Novikov coefficients in the Floer cohomology groups.

To conclude this subsection, we make the following observation:

Lemma 31. *If $(\mathfrak{a}, 0)$ has curvature bounded above (i.e., $\mathfrak{p} = 0$), then so does the perturbed data $(\mathfrak{a}, \mathfrak{p})$.*

Proof. This is due to the fact that the functions appearing in $\mathfrak{p}$ are assumed to be uniformly bounded in C^1 ; see [AAC23, Lemma 2.2]. $\square$

3.3.8. Transversality. Each finite-energy solution (σ, u) determines a linearized operator:

$$D_{\sigma, u} : TS_{\sigma} \oplus W^{1,p}(u^*TW) \rightarrow L^p(\text{Hom}^{0,1}(T\Sigma, u^*TW)),$$

where $\text{Hom}^{0,1}$ is the bundle of (j_{σ}, J_{σ}) -antilinear homomorphisms. We always assume $p > 2$ when discussing the Sobolev space $W^{1,p}$.

The details of the construction of the linearized operator are well-known; see, e.g., [Sal97, Can22]. Let us comment that it is obtained by differentiating the local coordinate representations of the non-linear PDE with respect to local deformations of u (and σ); this yields a linear differential operator acting on the local deformations; these can be patched together to obtain the global linearized operator.

Restricting $D_{\sigma, u}$ to variations fixing σ yields a particular type of differential operator, namely, a Cauchy-Riemann operator with asymptotic conditions. Because we required that the asymptotic Hamiltonian systems are non-degenerate in §3.2.3, the linearized operator is Fredholm operator. If the cokernel of $D_{\sigma, u}$ is zero for each solution (σ, u) , then we say that transversality holds. In this case, the space of solutions has the structure of a smooth manifold, and its dimension is equal to the dimension of the kernel of $D_{\sigma, u}$. This dimension can be computed by a general formula for the Fredholm index of a Cauchy-Riemann operator with asymptotic conditions. This is all a quite standard part of Floer theory.

We will appeal to the following general transversality lemma:

Lemma 32. *Given a family of domains, almost complex structures, Lagrangian boundary conditions, and the connection one-form $\mathfrak{a}$ on $S \times \Sigma \times W$, transversality can be achieved for all solutions (σ, u) to the perturbed equation §3.3.3 by choosing $\mathfrak{p}$ to be a generic perturbation one-form.*

Proof. This is a standard application of the usual Sard-Smale argument; see [AAC23, §2.3.10] and the references therein. $\square$

As discussed in §3.3.5 and §3.3.6, a special role is played by solutions of the s -independent equation (Floer's differential equation):

$$(9) \quad \partial_s u + J_{\zeta}(u)(\partial_t u - X_{\zeta, t}(u)) = 0,$$

where u is defined on a cylinder and $J_\zeta, X_{\zeta,t}$ are the asymptotic almost complex structure and Hamiltonian vector field associated to a puncture ζ .

This equation also has a linearized operator, and it is also important to achieve transversality for this. However, the perturbation term can no longer be used, since we require $\mathfrak{p}$ to be supported away from the punctures. In this case, we achieve transversality by assuming the asymptotic vector fields $X_{\zeta,t}$ are sufficiently generic. As shown in [FHS95], transversality can be achieved by perturbing $X_{\zeta,t}$ in a small, compactly supported way, away from its non-degenerate orbits; see also [BC24, §4.1]. This requires us to slightly amend condition (2) from §3.2.3 by requiring that $X_{\zeta,t}$ is sufficiently generic (in addition to having non-degenerate 1-periodic orbits). Henceforth, we will assume asymptotic data is sufficiently generic to achieve transversality for Floer's differential equation.

3.4. Floer cohomology. In this section we define the Floer cohomology persistence module $c \in \mathbb{R} \mapsto V_c \in \text{Vect}(\mathbb{Z}/2\mathbb{Z})$, with its three additional structures: the product, the BV-operator, and the map $H^*(W) \rightarrow V_0$.

The reader who is comfortable with these structures should feel free to skip to §3.5.

3.4.1. The Floer complex. Let $H_t, t \in \mathbb{R}/\mathbb{Z}$, be a smooth family in $\mathcal{H}$, and suppose that the associated Hamiltonian vector field X_t has non-degenerate 1-periodic orbits. Let J be an almost complex structure on W which is ω -tame and Liouville invariant outside of Ω . Moreover, suppose that X_t is sufficiently generic that the space of finite energy solutions to Floer's differential equation is cut transversally, as explained in §3.3.8. Let us call such data (H_t, J) *admissible* for defining the Floer complex.

The requirements in §3.2 and §3.3 ensure that the asymptotics $(H_{\zeta,t}, J_\zeta)$ which appear in the general form of Floer's equation §3.3.3 are admissible for defining the Floer complex.

Given admissible data, the *Floer complex* $\text{CF}(H_t, J)$ is defined to be the $\mathbb{Z}/2\mathbb{Z}$ vector space generated by the 1-periodic orbits of X_t with the differential d given by counting the one-dimensional components of the moduli space of finite energy solutions to Floer's differential equation. The right asymptotic of the solution is considered as the input of the differential and the left asymptotic is considered as the output.

The homology of the complex is denoted $\text{HF}(H_t, J)$, and is called the Floer cohomology of the pair (H_t, J) .

The details of the definition of the Floer differential, and why it squares to zero, are well-known; see, e.g., [Flo89, HS95, Sal97]. The requisite regularity and compactness results follow from our general discussion in §3.3. One aspect of the argument we do not explain is the Floer theory gluing argument used to relate the coefficients appearing in $d \circ d$ to the two-dimensional components of the moduli space of Floer differential cylinders; we refer the reader to, e.g., [Sal97, §3.3] for an exposition of this gluing theory.

Let us comment on one slightly non-standard aspect of the systems that we consider. We do not assume that $H_t = cr$ holds in the end for a fixed slope c ; rather, our definition allows $H_t = c_t r$ where c_t varies with t . We define the *slope* of the family H_t to be the real number:

$$c(H_t) = \int_0^1 c_t dt$$

This number plays a special role in the persistence module: V_c is defined as a formal limit over all the Floer cohomology groups $\mathrm{HF}(H_t, J)$ with $c(H_t) = c$.

The slope $c(H_t)$ has a geometric interpretation: $X_t = c_t R$ holds in the region $r \geq r_0$ for r_0 large enough (how large depends on the smooth family $H_t \in \mathcal{H}$; see §3.1.1), where $R = X_r$ is the so-called Reeb vector field. Thus the isotopy obtained by integrating X_t agrees with the Reeb flow for time $c(H_t)$ in the region $r \geq r_0$, up to a time reparametrization.

$$\text{output} \left(\begin{array}{c} \partial_s u + J(u)(\partial_t u - X_t(u)) = 0 \end{array} \right) \text{input}$$

FIGURE 3. The Floer differential

3.4.2. Continuation maps. Continuation maps are chain maps which relate the Floer complexes for different choices of admissible data. Because we use Hamiltonian functions H_t with varying slope, a bit of care is needed to define continuation maps in the generality we will require.

We define *continuation data* to be a connection 1-form $\mathfrak{a}$ and almost complex structure J on $\Sigma \times W$, where $\Sigma = \mathbb{R} \times \mathbb{R}/\mathbb{Z}$, such that:

$$\mathfrak{a} = K_{s,t} ds + H_{s,t} dt,$$

where $H_{s,t}, K_{s,t}$ are smooth families in $\mathcal{H}$ satisfying:

- (1) $K_{s,t} = b_{s,t} r$ and $H_{s,t} = c_{s,t} r$ for $r \geq r_0$, for an r_0 independent of s, t ,
- (2) $\partial_s c_{s,t} \leq \partial_t b_{s,t}$,
- (3) $K_{s,t} = \partial_s H_{s,t} = 0$ for $|s| \geq s_0$,
- (4) $H_{s,t} = H_{\zeta_1,t}$ for $s \leq -s_0$ and $H_{s,t} = H_{\zeta_0,t}$ for $s \geq s_0$,
- (5) $J = J_{s,t}$ satisfies $J_{s,t} = J_{\zeta_1}$ for $s \leq -s_0$ and $J_{s,t} = J_{\zeta_0}$ for $s \geq s_0$,
- (6) $(H_{\zeta_0,t}, J_{\zeta_0})$ and $(H_{\zeta_1,t}, J_{\zeta_1})$ are admissible for defining CF.

Such data is called continuation data from $(H_{\zeta_0,t}, J_{\zeta_0})$ to $(H_{\zeta_1,t}, J_{\zeta_1})$, i.e., as in the Floer differential, we consider the right asymptotic to be the input.

A *homotopy* of continuation data is a connection one-form and almost complex structure on $\mathbb{R} \times \Sigma \times W$ such that the restriction to $\{\tau\} \times \Sigma \times W$ satisfies (1) through (6) for each τ , with fixed asymptotic data $(H_{\zeta_0,t}, J_{\zeta_0})$ and $(H_{\zeta_1,t}, J_{\zeta_1})$, and such that the numbers r_0, s_0 can be taken to be independent of τ, s, t . One says that the continuation data obtained by restricting to $\{0\} \times \Sigma \times W$ and $\{1\} \times \Sigma \times W$ are homotopic.

Lemma 33. *Given asymptotic data $(H_{\zeta_0,t}, J_{\zeta_0})$ and $(H_{\zeta_1,t}, J_{\zeta_1})$ satisfying (6), there exists continuation data between them if and only if the slope of $H_{\zeta_0,t}$ is at most the slope of $H_{\zeta_1,t}$. In this case, there is a homotopy between any two choices of continuation data.*

Proof. First suppose there exists continuation data. Introduce the averages:

$$c(s) = \int_0^1 c(s, t) dt.$$

Then one uses the requirement (2) to show:

$$\partial_s c(s) \leq \int_0^1 \partial_t b(s, t) dt = 0.$$

In particular, the slope at the left end is greater than the slope at the right end. This proves the “only if” part of the first assertion.

For the “if” part, we fix a standard cut-off function $\beta(s)$ and define:

$$\begin{aligned} H_{s,t} &= \beta(s)H_{\zeta_0,t} + (1 - \beta(s))H_{\zeta_1,t} \\ K_{s,t} &= \int_0^t \frac{\partial}{\partial s} H_{s,\tau} d\tau - t \frac{\partial}{\partial s} \int_0^1 H_{s,\tau} d\tau. \end{aligned}$$

It is straightforward to check that $H_{s,t}, K_{s,t}$ are 1-periodic in the t -variable, and that properties (1) through (4) are satisfied. It is important that the slope of $H_{\zeta_0,t}$ is less than the slope of $H_{\zeta_1,t}$ in order for (2) to be satisfied, as it ensures that:

$$\frac{\partial}{\partial s} \int_0^1 H_{s,\tau} d\tau \leq 0 \text{ on the region where } r \geq r_0.$$

One uses the contractibility of the space of almost complex structures to extend the asymptotic complex structures J_{ζ_0} and J_{ζ_1} such that (5) is satisfied.

Finally we prove that any two continuation data between the same asymptotic systems are homotopic. This is straightforward, as one can simply take a convex combination of two one-forms $\mathfrak{a}$ and easily verify the properties (1) through (4) (which are all preserved under convex combinations), and again appeal to the contractibility of the space of almost complex structures. This completes the proof. $\square$

Given continuation data $(\mathfrak{a}, J)$ between $(H_{\zeta_0,t}, J_{\zeta_0})$ and $(H_{\zeta_1,t}, J_{\zeta_1})$, one interprets the counts the rigid solutions of the moduli space of finite energy solutions to §3.3.3 with generic perturbation term $\mathfrak{p}$ as defining the coefficients in a linear map (called a *continuation map*):

$$\text{CF}(H_{\zeta_0,t}, J_{\zeta_0}) \rightarrow \text{CF}(H_{\zeta_1,t}, J_{\zeta_1})$$

Here *rigid* means the counting the zero-dimensional components of the moduli space, which is assumed to be cut transversally.

Let us comment on one aspect related to why the count of rigid elements is finite: condition (2) ensures that the connection 1-form $\mathfrak{a}$ has curvature bounded from above. Indeed, the curvature two-form was given by:

$$\mathfrak{r} = (\partial_s H_{s,t} - \partial_t K_{s,t} + \omega(X_{s,t}^K, X_{s,t}^H)) ds \wedge dt,$$

and since $X_{s,t}^K, X_{s,t}^H$ are proportional when $r \geq r_0$ (both point in the direction of the Reeb flow), we have:

$$\mathfrak{r} = (\partial_s c_{s,t} - \partial_t b_{s,t}) r ds \wedge dt$$

when $r \geq r_0$, and this is non-positive.

The details of the construction of the continuation map are standard; these maps form a basic ingredient in the Floer cohomology TQFT which has been carefully constructed in the open case by [Rit13]; see also [Sch95, HS95, Abo15].

Standard arguments show:

- (1) the continuation map is a chain map,
- (2) the chain homotopy class of the map is independent of the generic perturbation $\mathfrak{p}$ or the homotopy class of continuation data,
- (3) the composition of two continuation maps is a continuation map,
- (4) the continuation map $\mathrm{HF}(H_t, J) \rightarrow \mathrm{HF}(H_t, J)$ is the identity map.

We will not review these standard arguments, and refer the reader to, e.g., [Sch95, HS95, Rit13, Abo15]. Let us comment that (1), (2), and (3) are based on gluing arguments similar to those used to prove $d \circ d = 0$.

A standard consequence of (1) through (4) is, if $(H_{\zeta_0, t}, J_{\zeta_0})$ and $(H_{\zeta_1, t}, J_{\zeta_1})$ are admissible for defining CF, and $H_{\zeta_0, t}$ and $H_{\zeta_1, t}$ have the same slope, then their Floer homologies are canonically isomorphic. We will continue with this line of reasoning in the next subsection §3.4.3.

3.4.3. The Floer cohomology persistence module. Let $\mathcal{D}$ be the category whose objects are pairs (H_t, J) which are admissible for defining the Floer complex, and where there is a unique morphism $(H_{\zeta_0, t}, J_{\zeta_0}) \rightarrow (H_{\zeta_1, t}, J_{\zeta_1})$ if the slope of $H_{\zeta_0, t}$ is at most the slope of $H_{\zeta_1, t}$, and no morphisms otherwise. The assignment:

$$(H_t, J) \mapsto \mathrm{HF}(H_t, J),$$

together with the continuation map construction induces a functor:

$$\mathrm{HF} : \mathcal{D} \rightarrow \mathrm{Vect}(\mathbb{Z}/2\mathbb{Z}),$$

where $\mathrm{Vect}(\mathbb{Z}/2\mathbb{Z})$ is the category of vector spaces over the field $\mathbb{Z}/2\mathbb{Z}$. In a certain sense to made precise, this functor factorizes through the functor $\mathcal{D} \rightarrow (\mathbb{R}, \leq)$ which sends (H_t, J) to its slope.

Summarizing, we have a diagram of the form:

$$\begin{array}{ccc} & \mathcal{D} & \\ \text{slope} \swarrow & & \searrow \mathrm{HF} \\ (\mathbb{R}, \leq) & \xrightarrow{\quad V \quad} & \mathrm{Vect}(\mathbb{Z}/2\mathbb{Z}), \end{array}$$

and the claim is that there is a persistence module V which makes the diagram commute in the following sense:

Lemma 34. *There exists a functor $V : (\mathbb{R}, \leq) \rightarrow \mathrm{Vect}(\mathbb{Z}/2\mathbb{Z})$ equipped with a natural isomorphism:*

$$I : V \circ (\text{slope}) \rightarrow \mathrm{HF}$$

of functors defined on $\mathcal{D}$. Moreover, any two such (V, I) , (V', I') are isomorphic, in that there is a unique natural isomorphism $U : V \rightarrow V'$ making the

diagram:

$$\begin{array}{ccc} V \circ \text{slope} & \xrightarrow{U} & V' \circ \text{slope} \\ \downarrow I & & \downarrow I' \\ \text{HF} & \xrightarrow{\text{id}} & \text{HF} \end{array}$$

commute; the diagram is valued in the category of functors $\mathcal{D} \rightarrow \text{Vect}(\mathbb{Z}/2\mathbb{Z})$.

Proof. The argument is essentially abstract nonsense. For each c , consider the full subcategory $\mathcal{D}(c)$ of objects (H_t, J) where the slope of H_t is no less than c . The restriction of HF to this category is again a functor, and we define the category theory limit:

$$V_c = \lim_{\mathcal{D}(c)} \text{HF};$$

see, e.g., [Mac71, Alu09] for details on category theory limits.

Prosaically speaking, we define V_c as the inverse limit of the Floer cohomologies of Hamiltonian systems whose slopes are at least c . If c is not the period of a Reeb orbit of the Reeb flow associated to $\partial\Omega$, then $\mathcal{D}(c)$ has initial objects (H_t, J) , and then the limit is isomorphic to $\text{HF}(H_t, J)$; this produces the natural isomorphism I . If c is the period of a Reeb orbit, then $\mathcal{D}(c)$ does not have an initial object; in this case, it is possible that V_c is an infinite dimensional vector space; these outputs of the persistence module can be thought of as living in the “completion” of the category of finite dimensional vector spaces.

Proving that any other construction (V', I') is canonically isomorphic to our (V, I) is straightforward application of universal properties of limits, and other abstract nonsense. $\square$

Henceforth, we will refer to the functor $V : \mathbb{R} \rightarrow \text{Vect}(\mathbb{Z}/2\mathbb{Z})$ as the *Floer cohomology persistence module* associated to Ω .

Let us comment that, if the set of periods of the Reeb flow (i.e., its spectrum) is *discrete*, then V_c is finite dimensional for all c (although we will not use this fact). In the same vein, if $c \in \mathbb{R}$ is such that $(c, c + \epsilon)$ contains no points in the spectrum of the Reeb flow, then V_c is finite dimensional. As a special case, V_0 is finite dimensional.

3.4.4. The pair-of-pants product. In this section, we briefly recall the pair-of-pants product structure on Floer cohomology groups. The details of the construction will be important in the sequel, but for now, let us summarize the resulting algebraic structure.

Let $\mathcal{D}^3$ be the full subcategory of $\mathcal{D} \times \mathcal{D} \times \mathcal{D}$ consisting of data:

$$((H_{\zeta_0, t}, J_{\zeta_0}), (H_{\zeta_1, t}, J_{\zeta_1}), (H_{\zeta_\infty, t}, J_{\zeta_\infty})),$$

where $H_{\zeta_\infty, t}$ has a slope no less than the sum of the slopes of $H_{\zeta_0, t}$ and $H_{\zeta_1, t}$. There are two functors defined on $\mathcal{D}^3$:

$$T_1 = \text{HF}(H_{\zeta_0, t}, J_{\zeta_0}) \otimes \text{HF}(H_{\zeta_1, t}, J_{\zeta_1}),$$

$$T_2 = \text{HF}(H_{\zeta_\infty, t}, J_{\zeta_\infty}),$$

the “pair-of-pants product” is a natural transformation $*$: $T_1 \rightarrow T_2$.

Via suitable abstract nonsense, this induces a natural transformation:

$$* : V_{c_0} \otimes V_{c_1} \rightarrow V_{c_\infty}$$

between two functors defined on the subcategory of $\mathbb{R}^3$ (a partially ordered set) consisting of those objects (c_0, c_1, c_∞) where $c_\infty \geq c_0 + c_1$. The partial order on $\mathbb{R}^3$ is the one where $(c_0, c_1, c_\infty) \leq (c'_0, c'_1, c'_\infty)$ if and only if $c_i \leq c'_i$ for each $i = 0, 1, \infty$.

The construction of $*$ is formally similar to the construction of continuation maps: one defines a class of *pair-of-pants data* which consists of connection one-forms $\mathfrak{a}$ and almost complex structures J , and then uses perturbations $\mathfrak{p}$ to define a transversally cut-out moduli space whose rigid counts are packaged into a chain map. The resulting map on homology is independent of the perturbation $\mathfrak{p}$ and homotopy class of pair-of-pants data $(\mathfrak{a}, J)$. One shows the map on homology commutes with continuation maps (in a suitable sense) and therefore induces the aforementioned natural transformation $*$.

In the rest of this section, we describe the construction of the pair-of-pants product, with focus on the details relevant to the proof of Theorem 7.

We begin by defining pair-of-pants data, which is a connection one-form $\mathfrak{a}$ and almost complex structure on $\Sigma \times W$ where $\Sigma = \mathbb{C} \setminus \{0, 1\}$. We require the following properties:

- (1) $\mathfrak{a} = H_{\zeta_i, t} dt$ and $J = J_{\zeta_i}$ holds in standard cylindrical ends around the $i = 0, 1, \infty$ punctures, and $(H_{\zeta_i, t}, J_{\zeta_i})$ are admissible for defining the Floer complex
- (2) $\mathfrak{a} = f_{x,y} r dx + g_{x,y} r dy$ holds outside of $r \geq r_0$, where $f, g \in \mathbb{R}$ vary smoothly with x, y ,
- (3) $\partial_x g_{x,y} - \partial_y f_{x,y} \leq 0$.

Here a standard cylindrical end around $z = 0, 1$ is obtained by parametrizing a disk around z by $[0, \infty) \times \mathbb{R}/\mathbb{Z}$ using the exponential map, and a standard cylindrical end around $z = \infty$ is obtained by parametrizing the complement of a disk around 0 by $(-\infty, 0] \times \mathbb{R}/\mathbb{Z}$, again using the exponential map. One breaks the rotational symmetry by requiring that the line $t = 0$ is aligned with the positive real axis.

We say such data goes from $(H_{\zeta_0, t}, J_{\zeta_0})$, $(H_{\zeta_1, t}, J_{\zeta_1})$ to $(H_{\zeta_\infty, t}, J_{\zeta_\infty})$.

Similarly to the case of the continuation data in §3.4.2, one can speak of homotopies of pair-of-pants data. The analog of Lemma 33 is:

Lemma 35. *Given $(H_{\zeta_i, t}, J_{\zeta_i})$ for $i = 0, 1, \infty$, there exists pair-of-pants data if and only if the slope of $H_{\zeta_\infty, t}$ is no less than the sum of the slopes of $H_{\zeta_0, t}$ and $H_{\zeta_1, t}$. In this case, there is a unique homotopy class of pair-of-pants data.*

Proof. The key idea in the proof is to consider $\mathfrak{a} = f_{x,y} dx + g_{x,y} dy$ as a one-form on Σ . The integral of $d\mathfrak{a}$ (which is compactly supported) over Σ is the difference in slopes:

$$\int_{\Sigma} d\mathfrak{a} = c(H_{\zeta_0, t}) + c(H_{\zeta_1, t}) - c(H_{\zeta_\infty, t});$$

such an observation appears in, e.g., [Rit13]. This proves the “only if” part of the first assertion. The “if” part is a straightforward construction. Finally, the uniqueness of the homotopy class follows from the convexity of the space of pair-of-pants data (and the contractibility of the space of almost complex structures). $\square$

Given pair-of-pants data $(\mathfrak{a}, J)$, and a generic perturbation one-form $\mathfrak{p}$, one packages the counts of the rigid finite energy solutions to §3.3 into a map:

$$(10) \quad * : \mathrm{CF}(H_{\zeta_0,t}, J_{\zeta_0}) \otimes \mathrm{CF}(H_{\zeta_1,t}, J_{\zeta_1}) \rightarrow \mathrm{CF}(H_{\zeta_\infty,t}, J_{\zeta_\infty}).$$

The counts of rigid elements are finite because $\mathfrak{a}$ has curvature bounded above — this is a consequence of requirement (3), and also our assumption that ω is exact; see the discussion in §3.3.7.

As explained in [Sch95, Rit13], this map is a chain map, and its chain homotopy class is independent of the perturbation term $\mathfrak{p}$ and the homotopy class of the pair-of-pants data. The resulting map on homology is the product structure explained at the start of this section. The claimed naturality of the product operation follows from the fact that (10) commutes with continuation maps, up to chain homotopy; this standard fact is proved in, e.g., [Sch95, Rit13]; one indication of why this holds is that if one “glues” continuation data to pair-of-pants data, one obtains new pair-of-pants data. The detailed argument involves Floer theoretic gluing and we defer the precise argument to the aforementioned references.

We end this section by commenting on one detail which will be important in the sequel. The product operation is defined by counting finite energy solutions to §3.3, and the fact that solutions have non-negative energy integral implies an inequality involving the values of $\int_0^1 H_{\zeta_i,t}$ at the three asymptotics (see the energy estimate Lemma 30). For well-chosen pair-of-pants data, this inequality essentially says the product “respects action filtrations;” this will be used in a crucial way in the proof of Theorem 7 in §3.6.8. We will return to this point and discuss the relevant action filtrations in §3.6.7. For other results on the interaction between the product structure and action filtration, we refer the reader to [Sch00, EP03, KS21, AAC23].

3.4.5. The BV-operator on Floer cohomology. The BV-operator is a natural endomorphism of functors $\mathcal{D} \rightarrow \mathrm{Vect}(\mathbb{Z}/2\mathbb{Z})$:

$$\Delta : \mathrm{HF} \rightarrow \mathrm{HF};$$

abstract nonsense yields an endomorphism Δ of the persistence module V . The goal of this section is to briefly explain its construction. We refer the reader to [Abo15, pp. 326] for a detailed exposition.

We define *BV-data* to be a connection one-form $\mathfrak{a} = K_{\theta,s,t}ds + H_{\theta,s,t}dt$ and almost complex structure J on the family of domains $\mathbb{R}/\mathbb{Z} \times \Sigma \times W$, where $\Sigma = \mathbb{R} \times \mathbb{R}/\mathbb{Z}$ is the cylinder, satisfying:

- (1) $\mathfrak{a} = H_{t+\theta}dt$ for $s \leq -s_0$,
- (2) $\mathfrak{a} = H_tdt$ for $s \geq s_0$,
- (3) $H_{\theta,s,t} = c_{\theta,s,t}r$ and $K_{\theta,s,t} = b_{\theta,s,t}r$ for $r \geq r_0$,
- (4) $\partial_s c_{\theta,s,t} \leq \partial_t b_{\theta,s,t}$,

- (5) J is fixed ω -tame almost complex structure on W , which is Liouville equivariant outside of $\Omega(1)$,
- (6) (H_t, J) is admissible for defining the Floer complex,

for some positive r_0, s_0 , and where θ is the $\mathbb{R}/\mathbb{Z}$ coordinate on $\mathbb{R}/\mathbb{Z} \times \Sigma \times W$. As in §3.4.2 and §3.4.4, one can speak about homotopies of BV-data; such homotopies are required to satisfy the above properties with fixed H_t, J at all moments.

There is an analog of Lemma 33 and Lemma 35.

Lemma 36. *There exists BV-data for any data (H_t, J) which is admissible for defining CF. Moreover, there is a unique homotopy class of BV-data.*

Proof. The existence of at most one homotopy class follows from the convexity of the space of BV-data with fixed (H_t, J) . It remains only to prove there exists some BV-data.

The construction is rather simple; one defines:

$$H_{\theta,s,t} := (1 - \beta(s))H_{t+\theta} + \beta(s)H_t,$$

where $\beta(s)$ is a standard cut-off function, and then defines:

$$K_{\theta,s,t} := \int_0^t \partial_s H_{\theta,s,\tau} d\tau - t \int_0^1 \partial_s H_{\theta,s,\tau} d\tau = \int_0^t \partial_s H_{\theta,s,\tau} d\tau.$$

Note the similarity between this and the construction used in Lemma 33. The formula for $K_{\theta,s,t}$ simplifies because $\partial_s H_{\theta,s,t}$ has zero time average — the slope of $H_{\theta,s,t}$ is constant as s varies.

One easily verifies the enumerated conditions hold. Indeed, one verifies directly that $\partial_s H_{\theta,s,t} = \partial_t K_{\theta,s,t}$, which implies² (4). This completes the proof. $\square$

The way BV-data is used to define a map $\Delta : \text{CF}(H_t, J) \rightarrow \text{CF}(H_t, J)$ follows the same lines as §3.4.2 and §3.4.4. Briefly, one picks generic perturbation one-form $\mathfrak{p}$ on $\mathbb{R}/\mathbb{Z} \times \Sigma \times W$, and defines the map whose coefficients are the counts of rigid finite-energy solutions to §3.3.3.

Let us briefly explain how to interpret the count of solutions as coefficients in a matrix. If (θ, u) is a rigid finite energy solution, then $u(s, t) \rightarrow \gamma_{\text{out}}(t + \theta)$ as $s \rightarrow -\infty$ and $u(s, t) \rightarrow \gamma_{\text{in}}(t)$ as $s \rightarrow \infty$, where γ_{in} and γ_{out} are 1-periodic orbits of the system generated by H_t ; we consider u as contributing to the coefficient in the matrix with entry $(\gamma_{\text{out}}, \gamma_{\text{in}})$.

3.4.6. PSS and inclusion of the constant loops. Recall the full subcategory $\mathcal{D}(0)$ of objects $(H_t, J) \in \mathcal{D}$ where the slope of H_t is no less than 0. In this section we explain how to define the PSS morphism, which can be considered as a natural transformation:

$$\text{PSS} : H^*(W) \rightarrow \text{HF}|_{\mathcal{D}(0)};$$

²Let us observe that, for any BV data, $\partial_s c_{\theta,s,t} - \partial_t b_{\theta,s,t}$ has zero time-average, and hence must be constant. Thus we could replace (4) by the apparently stronger condition $\partial_s c_{\theta,s,t} = \partial_t b_{\theta,s,t}$ without any loss of generality.

in order to interpret PSS as a natural transformation, the domain $H^*(W)$ is considered as a constant functor on $\mathcal{D}(0)$. Abstract nonsense (i.e., taking limits) shows that PSS descends to a homomorphism:

$$\text{PSS} : H^*(W) \rightarrow V_0;$$

see §3.4.3.

The PSS construction goes back to [PSS96], and has been generalized to the setting of convex-at-infinity symplectic manifolds in [FS07, Rit13]. Typically one works with a Morse theory version of $H^*(W)$. In our framework, recall from §1.5.1 that we prefer to work with a proxy for singular cohomology, and instead define $H^*(W)$ to be the group of smooth proper maps $C : S \rightarrow W$ modulo proper cobordisms.

We define PSS-data to be a connection one-form $\mathfrak{a}$ and almost complex structure J on $\Sigma \times W$ where³ $\Sigma = \mathbb{C}$ which, in the cylindrical coordinates $z = e^{-2\pi(s+it)}$, satisfies:

- (1) $\mathfrak{a} = K_{s,t}ds + H_{s,t}dt$,
- (2) $\mathfrak{a} = 0$ for $s \geq s_0$,
- (3) $\mathfrak{a} = H_t dt$ for $s \leq -s_0$,
- (4) $K_{s,t} = b_{s,t}r$, $H_{s,t} = c_{s,t}r$ for $r \geq r_0$,
- (5) $\partial_s c_{s,t} \leq \partial_t b_{s,t}$,
- (6) J is fixed as in §3.4.5,
- (7) $(H_t, J) \in \mathcal{D}$.

Note that it follows from (2), (4), and (7) that $(H_t, J) \in \mathcal{D}(0)$. PSS-data is similar to continuation data from $(0, J)$ to (H_t, J) , the only difference being that $(0, J)$ is not actually admissible for defining the Floer complex, and so we consider the puncture at $s = \infty$ (namely $z = 0$) as a removable singularity.

For a cycle $C : S \rightarrow W$, we define $\text{PSS}(C) \in \text{CF}(H_t, J)$ as a cycle. Fix PSS data $(\mathfrak{a}, J)$, which we pull back to the family $S \times \Sigma \times W$. Fix a generic perturbation term $\mathfrak{p}$, and count the rigid finite energy solutions (σ, u) to §3.3 satisfying the incidence constraint:

$$u(0) = C(\sigma).$$

One counts the asymptotic orbits of u as a linear combination in $\text{CF}(H_t, J)$. The arguments from [PSS96, FS07, Rit13] adapt easily to the present case to show that this linear combination is a cycle in $\text{CF}(H_t, J)$. The cohomology class of this cycle is independent of $\mathfrak{p}$, the PSS-data chosen, and even the proper cobordism class of F . By construction, one sees that PSS is a linear map $H^*(W) \rightarrow \text{HF}(H_t, J)$. Finally, the standard Floer theory gluing arguments show that PSS is a natural homomorphism, i.e., it commutes with continuation maps.

This concludes the specification of the Floer theory framework used in the paper.

³More precisely, in the terminology of §3.2.1, $\Sigma = \mathbb{C}P^1$ with $\Gamma_- = \{\infty\}$, so the punctured surface is $\mathbb{C}$.

3.5. Evaluation maps and ball embeddings. Fix a map $f : N \rightarrow \Omega$, where N is a closed manifold. The goal of this section is to explain how a lift of f to $\mathfrak{B}(a, \Omega)/U(n)$, as in §1, is used to define an evaluation map $\mathfrak{e} : V_c \rightarrow \mathbb{Z}/2\mathbb{Z}$ for $0 < c < a$ which satisfies:

$$\mathfrak{e}(\text{PSS}(\beta)) = 1$$

provided that $\beta = [C]$ where $C : S \rightarrow W$ is a cycle with non-zero homological intersection number with f .

In particular, a consequence we will deduce by the end of this section is:

Proposition 37. *If f lifts to $\mathfrak{B}(a, \Omega)/U(n)$, and $\beta \in H^*(W)$ has non-zero homological intersection number with f , then $\text{PSS}(\beta) \neq 0$ in V_c for $c < a$.*

Remark. One can define a capacity by the number:

$$c(\beta, \Omega) = \inf\{c : \text{PSS}(\beta) = 0 \text{ in } V_c\};$$

Proposition 37 implies $\text{Gr}(f, \Omega) \leq c(\beta, \Omega)$. In the special case when $\beta = [W]$, we conclude the estimate that $\text{Gr}(f, \Omega) \leq c([W], \Omega)$; see, e.g., [BK22].

Without any true loss of generality, we assume that f lifts to $\mathfrak{B}(a', \Omega)/U(n)$ for some number $a' > a$. We will appeal to such an extension in some of the subsequent arguments.

3.5.1. Family of Hamiltonian functions associated to a family of balls. Introduce a convex function $\mu : \mathbb{R} \rightarrow \mathbb{R}$ such that:

- (1) $\mu(x) = x$ for $x \geq 1$,
- (2) $\mu(x) = 1/2$ for $x \leq 0$,

and introduce the autonomous Hamiltonian function:

$$G_{c,\delta} = c\delta\mu(\delta^{-1}(r-1)) + c$$

for a fixed non-negative constant c and a very small parameter $\delta > 0$. Then:

- (1) $G_{c,\delta}$ agrees with cr outside of a neighbourhood of Ω (indeed, the neighborhood is $\{r \geq 1 + \delta\}$),
- (2) $G_{c,\delta}$ equals the constant $c + c\delta/2$ on the domain Ω .

One should consider this Hamiltonian as the background Hamiltonian system.

Recall that the parametric Gromov width concerns maps $f : N \rightarrow \Omega$ together with:

- (1) an open cover U_α of N ,
- (2) smooth maps $g_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow U(n)$,
- (3) extensions $F_\alpha : U_\alpha \times B(a) \rightarrow \Omega$ satisfying $F_\beta(\eta, z) = F_\beta(\eta, g_{\alpha\beta}(\eta)z)$ for $\eta \in U_\alpha \cap U_\beta$.

We will use the extensions of F_α to define an η -parametric family of perturbations of the Hamiltonian $G_{c,\delta}$.

Let us denote by B_η the image of $z \mapsto F_\alpha(\eta, z)$, for any α such that $\eta \in U_\alpha$. The image of this ball, and the radial coordinate $|z|$, is independent of the choice of α .

Let $D_\eta : W \rightarrow \mathbb{R}$ be a smooth family of functions, parametrized by $\eta \in N$, with compact support in Ω and such that:

- (1) $\{D_\eta < 0\}$ is the interior of B_η ,
- (2) $D_\eta \circ B_\eta(z) = \pi|z|^2 - a$.

Such a family will be constructed in Lemma 38 below using slight extensions of B_η . Note that the constants are chosen so that D_η vanishes on the boundary of B_η .

The η -parametric family of perturbations of $G_{c,\delta}$ we will use in our construction is:

$$(11) \quad H_{c,\delta,\epsilon,\eta} = G_{c,\delta} + \epsilon D_\eta,$$

which depends on parameters $c, \delta, \epsilon > 0$, and $\eta \in N$. We will denote the Hamiltonian vector field of this function by $X_{c,\delta,\epsilon,\eta}$.

Remark. In the following lemma, we discuss the *action* of an orbit $\gamma(t)$ of a system generated by Hamiltonian function H_t , which is defined as:

$$\int H_t(\gamma(t))dt - \int \gamma^* \lambda,$$

as is typical in Floer theory in exact symplectic manifolds.

Remark. We will also refer to the period spectrum of the Reeb vector field X_r , which we denote by $\text{Per}(X_r)$.

Lemma 38. *The family D_η can be chosen so that, for $\epsilon < 1$ and c not a period of X_r , the contractible 1-periodic orbits of $X_{c,\delta,\epsilon,\eta}$ are of three types:*

- (1) *constant orbits in Ω lying outside B_η , with action at least $c + c\delta/2$,*
- (2) *a single constant orbit at the center of B_η , with action $c + c\delta/2 - \epsilon a$,*
- (3) *non-constant periodic orbits of X_r lying in the hypersurfaces defined by $c\mu'(\delta^{-1}(r_0 - 1)) = b \in \text{Per}(X_r)$, with action at least $c - b$.*

In particular, if $\epsilon a > b + c\delta/2$ for all $b \in \text{Per}(X_r) \cap [0, c]$, there are no non-constant finite energy Floer cylinders for $X_{c,\delta,\epsilon,\eta}$ whose negative asymptotic is in (2). If $a > c$, then we can achieve this conclusion for ϵ close enough to 1 and δ close enough to 0.

Proof. As mentioned in §3.5, we will appeal to a small extension $B(a) \subset B(a')$. Fix a smooth function $\rho : \mathbb{R} \rightarrow \mathbb{R}$ such that $\rho(x) = x$ for $x \leq a$, $\rho(x) \geq a$ for $x \in [a, a']$, and $\rho(x) = a$ for $x \geq a'$, and such that $|\rho'| \leq 1$ holds at all points. Then we define $D(z) = \rho(\pi|z|^2) - a$ which can be pushed forward using the ball embedding $B_\eta(a') \rightarrow \Omega$ to define the desired family of functions D_η .

Then it is easily seen that the only orbits inside Ω are the orbits of types (1) and (2) with the claimed actions. The only other orbits are orbits of $G_{c,\delta}$ appearing outside of Ω . A standard computation shows that these orbits are orbits lying in a fixed radius $r = r_0$ satisfying $c\mu'(\delta^{-1}(r_0 - 1)) = b \in \text{Per}(X_r)$.

In this case we use the fact that $\lambda(X_r) = r$ to compute their action as:

$$\text{action} = c\delta\mu(\delta^{-1}(r_0 - 1)) + c - c\mu'(\delta^{-1}(r_0 - 1))r_0.$$

Simplifying we see that:

$$\text{action} = c - b + cf(r_0 - 1)$$

where:

$$f(x) = -\mu'(\delta^{-1}x)x + \delta\mu(\delta^{-1}x).$$

Using the convexity of μ , one sees that $f'(x) \leq 0$, and since $f(x) = 0$ for $x \geq \delta$ (since $\mu(s) = s$ for $s \geq 1$), we conclude that $f(x) \geq 0$. This gives the desired result. $\square$

Remark. In the proof we implicitly appealed to the well-known fact that the action of the left end of a Floer cylinder is at least the action of the right end; this is a special case of the general energy identity proved in Lemma 29.




FIGURE 4. Graph of $\rho : \mathbb{R} \rightarrow \mathbb{R}$

3.5.2. Definition of the evaluation-map. We continue in the context of the previous subsection. Let $(H_t, J) \in \mathcal{D}$ where H_t has slope at most $a > 0$, where a is the capacity appearing in the lift of f to $\mathfrak{B}(a, \Omega)/U(n)$.

In this subsection, we explain how a family of continuation data from (H_t, J) to $(H_{c,\delta,\epsilon,\eta}, J_\eta)$ parametrized by $\eta \in N$ can be used to define an evaluation map provided $\epsilon a > c(1 + \delta/2)$ and where c is no less than the slope of H_t .

Let us therefore define *evaluation data* to be a connection one-form $\mathfrak{a}$ and almost complex structure J on the family on $N \times \Sigma \times W$ where $\Sigma = \mathbb{R} \times \mathbb{R}/\mathbb{Z}$, satisfying:

- (1) $\mathfrak{a} = K_{\eta,s,t}ds + H_{\eta,s,t}dt$,
- (2) $K_{s,t} = b_{\eta,s,t}r$ and $H_{s,t} = a_{\eta,s,t}r$ for $r \geq r_0$,
- (3) $\partial_s c_{\eta,s,t} \leq \partial_t b_{\eta,s,t}$,
- (4) $\mathfrak{a} = H_t dt$ for $s \geq s_0$,
- (5) $\mathfrak{a} = H_{c,\delta,\epsilon,\eta} dt$ for $s \leq -s_0$, where $\epsilon a > c(1 + \delta/2)$,
- (6) J is fixed, i.e., independent of η, s, t , for $s \geq s_0$,
- (7) $J = J_\eta$ for some fixed family J_η for $s \leq -s_0$.

We comment that such data deviates slightly from the framework established in §3.2 and §3.3, in that we allow the asymptotic data to depend on $\eta \in N$ at the left end. We will explain shortly why this does not cause any issues.

Introducing a generic perturbation one-form $\mathfrak{p}$, one can still consider the finite energy solutions to Floer's equation §3.3.3.

For data $(\mathfrak{a}, \mathfrak{p}, J)$, introduce $\mathcal{M}$ as the moduli space of finite energy solutions (η, u) such that the left asymptotic of u is the constant orbit located at the center of the ball B_η . Since this central orbit is non-degenerate for each η , the space of solutions $\mathcal{M}$ behaves as if the asymptotic data at the left end were independent of η , at least as far as the results in §3.3 are concerned.

Remark. Some care is needed to properly define the linearization framework when the asymptotic data changes; see, e.g., the proof of Lemma 40.

The key result about $\mathcal{M}$ is the following:

Lemma 39. *Suppose the perturbation term $\mathfrak{p}$ is such that $\mathcal{M}$ is cut transversally. If (η_n, u_n) is a sequence of rigid solutions $\mathcal{M}$, then (η_n, u_n) has a convergent subsequence in $\mathcal{M}$. If (η_n, u_n) is a sequence of solutions in the 1-dimensional component of $\mathcal{M}$, then (η_n, u_n) has a subsequence which either converges in $\mathcal{M}$, or which breaks into configuration of a rigid solution in $\mathcal{M}$ on the left and a Floer differential cylinder for (H_t, J) on the right (in the sense discussed in §3.3.6).*

Proof. This is mostly standard Floer theory; the only non-standard thing we need to check is that a sequence of solutions in the 1-dimensional component of $\mathcal{M}$ does not break into a configuration of a non-stationary Floer differential cylinder for $(H_{c,\delta,\epsilon,\eta}, J_\eta)$ on the left and a rigid solution in $\mathcal{M}$ on the right. However, such a breaking can be precluded by action considerations, since we have shown in Lemma 38 that there are no non-constant Floer differential cylinders for $(H_{c,\delta,\epsilon,\eta}, J_\eta)$ whose left asymptotic is the central orbit. $\square$

To define the map $\mathfrak{e} : \text{CF}(H_t, J) \rightarrow \mathbb{Z}/2$, we define:

$$\mathfrak{e}(\gamma) = \#\{u \in \mathcal{M} : \lim_{s \rightarrow \infty} u(s, t) = \gamma(t)\} \bmod 2.$$

Lemma 39 and standard Floer theory gluing arguments imply that $\mathfrak{e}$ is a chain map, i.e., $\mathfrak{e} \circ d = 0$.

Since the space of evaluation data is weakly contractible, for fixed input system (H_t, J) , the standard arguments show that the chain homotopy class of $\mathfrak{e}$ is independent of the choice of evaluation data. The resulting map on homology is denoted $\mathfrak{e} : \text{HF}(H_t, J) \rightarrow \mathbb{Z}/2$. Summarizing, we have constructed a map:

$$\mathfrak{e} : \text{HF}(H_t, J) \rightarrow \mathbb{Z}/2\mathbb{Z}$$

on the full subcategory $\mathcal{D}_{<a} \subset \mathcal{D}$ consisting of objects with slope at most a . Finally, the usual gluing arguments prove that $\mathfrak{e}$ is a natural transformation, where $\mathbb{Z}/2\mathbb{Z}$ is interpreted as a constant functor.

In the next section, we prove that $\mathfrak{e} \circ \text{PSS} : H^*(W) \rightarrow \mathbb{Z}/2\mathbb{Z}$ sends β to 1 provided that β has non-zero intersection with the cycle $\eta \mapsto F(\eta, 0)$.

3.5.3. Non-triviality of the evaluation map. We continue in the context of the previous two subsections. Let $\beta \in H^*(W)$ have non-zero homological intersection with the cycle $\eta \mapsto F(\eta, 0)$. Pick a representative $C : S \rightarrow W$ of β which is transverse to the cycle $\eta \mapsto F(\eta, 0)$, so that the set of pairs (σ, η) solving $C(\sigma) = F(\eta, 0)$ is a finite odd set of points.

Define *glued evaluation PSS data* to be a connection one-form $\mathfrak{a}$ and almost complex structure J on the family $N \times \mathbb{C} \times W$ such that, in the cylindrical coordinates $z = e^{-2\pi(s+it)}$,

- (1) $\mathfrak{a} = K_{\eta,s,t}ds + H_{\eta,s,t}dt$,
- (2) $K_{s,t} = b_{\eta,s,t}r$ and $H_{s,t} = c_{\eta,s,t}r$ for $r \geq r_0$,
- (3) $\partial_s c_{\eta,s,t} \leq \partial_t b_{\eta,s,t}$,

- (4) $\mathfrak{a} = 0$ for $s \geq s_0$,
- (5) $\mathfrak{a} = H_{c,\delta,\epsilon,\eta} dt$ for $s \leq -s_0$, where $\epsilon a > c(1 + \delta/2)$, $c \geq 0$,
- (6) $J = J_{+,\eta}$ for $s \leq -s_0$ and $J = J_{-,\eta}$ for $s \geq s_0$, for two fixed families $J_{-,\eta}$ and $J_{+,\eta}$.

Roughly speaking, glued evaluation PSS data is the data one obtains by gluing the evaluation data from §3.5.2 to the PSS-data from §3.4.6.

Similarly to §3.4.6 and §3.5.2, we pull back glued evaluation PSS data to the family $S \times N \times \mathbb{C} \times W$, and for a generic perturbation term $\mathfrak{p}$, we consider the moduli space $\mathcal{M}$ of finite energy solutions (σ, η, u) to §3.3.3 satisfying the incidence condition:

$$u(0) = C(\sigma),$$

and such that the asymptotic orbit of u is the constant orbit located at the center of the ball B_η .

Lemma 40. *We have the following:*

- (a) *all components of $\mathcal{M}$ are compact,*
- (b) *the count of points in the rigid component of $\mathcal{M}$ is independent of the choice of glued evaluation PSS data and perturbation one-form $\mathfrak{p}$,*
- (c) *the count of points in the rigid component of $\mathcal{M}$ equals $\mathfrak{e}(\text{PSS}(\beta))$,*
- (d) *the count of points in the rigid component of $\mathcal{M}$ is odd;*

all counts are taken mod 2. Thus $\mathfrak{e}(\text{PSS}(\beta)) = 1$, proving Proposition 37.

Proof. Part (a) follows from the same compactness-up-to-breaking argument used in the proof of Lemma 39. A sequence of solutions will, in general, have a subsequence which converges or breaks into a configuration of a non-stationary Floer differential cylinder and another solution in $\mathcal{M}$. There is only one end where the breaking can happen (the left end). Since we require that solutions in $\mathcal{M}$ are asymptotic to the central orbit at the left end, and the only Floer differential cylinders whose left end is the central orbit are constant (by Lemma 38), we conclude there can be no breaking.

Part (b) follows from a cobordism argument: because the space of data is path connected, the rigid counts for two choices of data are cobordant finite sets of points, and hence have the same cardinality mod 2; here we note that the parametric moduli space used in such a cobordism argument will be compact by the exact same argument used for (a).

Part (c) follows from Floer theoretic gluing. By picking the glued evaluation PSS data as a concatenation of evaluation data (§3.5.2) and PSS-data (§3.4.6), a standard gluing argument shows the count of rigid points in $\mathcal{M}$ equals the count obtained by applying the map $\mathfrak{e}$ to the cycle $\text{PSS}(C)$.

The non-standard part of the argument is establishing (d). For this part of the argument, we use an explicit choice of data, and directly show the count equals the (odd) count of pairs (σ, η) satisfying $C(\sigma) = F(\eta, 0)$.

Let us pick the data such that $c = 0$, so that δ becomes irrelevant, and:

$$H_{c,\delta,\epsilon,\eta} = \epsilon D_\eta.$$

Pick $\mathbf{a} = (1 - \beta(s))\epsilon D_\eta dt$, and pick $J = J_\eta$ (so $J_+ = J_-$), where J_η agrees with the standard almost complex structure when pulled back to the ball B_η . We will show that all solutions in $\mathcal{M}$ are regular with $\mathbf{p} = 0$, and therefore the count without perturbation equals the count with perturbation.

Claim. For ϵ sufficiently small, any solution (σ, η, u) with this specific data must be such that u is constant and equal to the center of the ball B_η .

This claim follows from a simple adiabatic compactness argument: restricting to a half-infinite cylinder $(-\infty, R] \times \mathbb{R}/\mathbb{Z}$ we can pull u back to a map valued in $B(a)$ satisfying:

$$(12) \quad \partial_s u + J(\partial_t u - \epsilon(1 - \beta(s))X(u)) = 0$$

where X is the Hamiltonian vector field for $\pi|z|^2$, and J is the standard almost complex structure. Either we can take $R = +\infty$, or we can take R to be maximal in which case $u(R, t)$ hits the boundary $\partial B(a)$ for some t .

We can take ϵ small enough so that $|\partial_t u|$ is everywhere less than $\sqrt{a/\pi}$, by a simple compactness argument.

Consider the center of mass $\xi(s) = \int_0^1 u(s, t) dt$, so $\xi : (-\infty, R] \rightarrow B(a)$ is a smooth curve. Using the linearity of the above equation (bearing in mind that $X(u)$ is actually a linear function of u) one concludes that:

$$\partial_s \xi(t) = \epsilon(1 - \beta(s))JX(\xi).$$

It is well-known that JX is a vector field which points radially inwards. Thus, since $\xi(s)$ converges to 0 as $s \rightarrow -\infty$, we must have that $\xi(s) = 0$ holds identically. It therefore follows that $u(s, t)$ has mean zero for $s \leq R$.

Thus we have $R = +\infty$, since u cannot touch the boundary $\partial B(a)$ because $|\partial_t u(R, t)|$ is smaller than the radius of $\partial B(a)$ and $u(R, -)$ has mean zero. Then one estimates the energy integral of u by:

$$\text{energy of } u = \int |\partial_s u|^2 ds dt \leq -a + \int \beta'(s)a = 0,$$

which implies u must be constant, and thus equal to the center of the ball.

Thus $\mathcal{M}$ is in bijection with the set of pairs (σ, η) satisfying $C(\sigma) = F(\eta, 0)$; the bijection is simply the projection map $(\sigma, \eta, u) \mapsto (\sigma, \eta)$. It remains only to prove that these solutions are in fact regular.

To analyze the regularity of these solutions, we need to open the “black box” of the linearization framework. Let us fix a solution (σ_0, η_0, u_0) , and use the coordinate system induced by the embedded ball B_{η_0} .

If (σ, η, u) is nearby (σ_0, η_0, u_0) , then it solves the equation provided that:

$$(13) \quad \begin{cases} u = w + f(\eta) \text{ for } w \in W^{1,p}(\Sigma) \\ \partial_s w + J_\eta(w + f(\eta))(\partial_t w - V_\eta(w)) = 0 \\ w(0) + f(\eta) - c(\sigma) = 0, \end{cases}$$

where $f(\eta)$, $c(\sigma) = C(\sigma)$, and $V_\eta(z) = \epsilon X_\eta(z + f(\eta))$ are represented in the coordinate system. Note that the maps c and f are transverse at their intersection (σ_0, η_0) .

Remark. When linearizing with a varying asymptotic, it is important to use such an auxiliary map w , as it is defined on a fixed Banach space of maps.

The second two equations in (13) can be considered as a non-linear map defined on a neighborhood of 0 between Banach manifolds:

$$W^{1,p} \times S \times N \rightarrow L^{1,p} \times \mathbb{R}^{2n}.$$

Differentiating this non-linear map gives the linearized operator.

If (σ', η', w') is a tangent vector at $(\sigma_0, \eta_0, 0)$, the linearized operator is:

$$\begin{bmatrix} \partial_s w' + J \partial_t w' - \epsilon J X(w') - J \partial V_\eta / \partial \eta(0) \eta' \\ w'(0) + df(\eta_0) \eta' - dc(\sigma_0) \sigma' \end{bmatrix} \in L^p \oplus \mathbb{R}^{2n}.$$

The fact that $V_\eta(0) = 0$ holds for all η implies $\partial V_\eta / \partial \eta(0) = 0$. Thus the first component is a Cauchy-Riemann operator $W^{1,p} \rightarrow L^p$. Because ϵX has a non-degenerate orbit at the origin, this operator is an isomorphism $W^{1,p} \rightarrow L^p$ (by standard Fredholm theory for Floer's equation). On the other hand, since f and c are transverse and have complementary dimensions, the second equation is an isomorphism $TS_{\sigma_0} \oplus TN_{\eta_0} \rightarrow \mathbb{R}^{2n}$. Thus the linearized operator is an isomorphism, and so the constant solution (σ_0, η_0, u_0) is regular. This completes the proof. $\square$

3.6. Family Floer cohomology. In this section we develop a version of family Floer cohomology following the scheme in [Hut08].

Let us briefly comment on the strategy. The identity $\mathfrak{e}(\text{PSS}(\beta)) = 1$ proved in §3.5 implies that any cycle $\sum \gamma_i$ in $\text{CF}(H_t, J)$ representing $\text{PSS}(\beta)$ has at least one orbit γ_i which appears as the right end of one of the cylinders used to define $\mathfrak{e}$; see Figure 5.

$$\text{center of } B_\eta \left(\overline{\partial_s u - Y_{\eta,s,t}(u) + J_{\eta,s,t}(u)(\partial_t u - X_{\eta,s,t}(u)) = 0} \right) \gamma_i$$

FIGURE 5. A consequence of $\mathfrak{e}(\sum \gamma_i) = 1$ is the existence of a cylinder (η, u) joining γ_i to the center of the ball B_η . In the figure Y, X are the Hamiltonian generators of K, H where $(\mathfrak{a} = K_{\eta,s,t} ds + H_{\eta,s,t} dt, J_{\eta,s,t})$ is evaluation data as in §3.5.2.

It is a general principle that the action of the left asymptotic (the center of the ball) is at least the action of the right asymptotic (the orbit γ_i), up to an error depending on the curvature of the Hamiltonian connection determined by $\mathfrak{a}$; see Lemma 30.

In Lemma 38 it was shown that the action of the center of the ball is equal to $(c + c\delta/2) - a$; thus if we can prove that:

- (1) the action of γ_i is greater than $c + c\delta/2 - a$, and
- (2) the error coming from the curvature can be made arbitrarily small,

then we have successfully obstructed the family of balls.

Item (1) will ultimately be a consequence of Theorem 7's hypothesis:

$$\text{PSS}(\beta) = \Delta(\zeta_1) * \zeta_2 \text{ holds in } V_{c_1+c_2},$$

where $\zeta_i \in V_{c_i}$, with $c_i > 0$, provided $c := c_1 + c_2 < a$.

However, because the cylinder in Figure 5 appears at an unknown parameter value $\eta \in N$, it is not possible to uniformly control the curvature because the input system (H_t, J) is independent of η .

For this reason, we instead work with a *family of input systems* $(H_{\eta,t}, J_\eta)$, and upgrade the map $\mathfrak{e}$ to be defined on a suitable family version of Floer cohomology. As we will show, with the family version, one can control the curvature and simultaneously achieve (1) and (2).

The family Floer cohomology essentially contains no new algebraic information (it is the regular Floer cohomology tensored with the cohomology of N); however, family Floer cohomology grants one access to new action filtrations on the Floer complex.

The rest of this section is dedicated to developing the necessary theory, and, in §3.6.8, completing the proof of Theorem 7.

3.6.1. Definition of family Floer cohomology. Let N be a compact manifold. By definition, *family Floer data* is:

- (1) a Morse-Smale pseudogradient P on N (e.g., a Morse-Smale gradient-like vector field in the terminology of [Mil65a]),
- (2) a smooth family $H_{\eta,t}$ in $\mathcal{H}$ parametrized by $N \times \mathbb{R}/\mathbb{Z}$,
- (3) a fixed almost complex structure J , which is, as usual, ω -tame and Liouville invariant in the convex end,

which satisfies:

- (4) for each zero η_0 of P , there is a neighborhood $\text{Op}(\eta_0)$ of η_0 such that $H_{\eta,t} = H_{\eta_0,t}$, and $(H_{\eta_0,t}, J)$ is admissible for defining $\text{CF}(H_{\eta_0,t}, J)$.
- (5) $H_{\eta,t} = cr$ for $r \geq r_0$.

Note. For simplicity, we do not allow $H_{\eta,t}$ to have a time-dependent slope, unlike the data for the non-family Floer complex. The reason for this is that we have no need to appeal to time-dependent slopes in the family Floer complex; on the other hand, time-dependent slopes will be used for the regular Floer complex in §4.

For each zero η_0 of P , associate the one-dimensional vector space $\mathbb{Z}/2\mathbb{Z}\eta_0$; for family Floer data $(P, H_{\eta,t}, J)$ define the *family Floer complex*:

$$\text{CFF}(P, H_{\eta,t}, J) = \bigoplus \text{CF}(H_{\eta_0,t}, J) \otimes \mathbb{Z}/2\mathbb{Z}\eta_0,$$

where the direct sum is over the zeros of P . Thus it makes sense to refer to a generator $\gamma \otimes \eta_0$ whenever $\gamma(t)$ is a 1-periodic orbit of $H_{\eta_0,t}$.

In the rest of this section we will explain how to define the differential d_{FF} on the family Floer complex. As part of the definition, we will need to require that certain auxiliary moduli spaces are cut transversally, and this will impose a genericity condition on the family $H_{\eta,t}$.

For each pair (η_1, η_0) of zeros of P , one can consider the (open) manifold $\mathcal{P}(\eta_1, \eta_0)$ of parametrized flow lines $\pi(s)$ from η_1 to η_0 for the negative pseudogradient $-P$; here η_1 is the left asymptotic and η_0 is the right asymptotic.

Fixing (η_1, η_0) , there is an induced connection and perturbation one-form, and almost complex structure, on the family $\mathcal{P}(\eta_1, \eta_0) \times \Sigma \times W$, where Σ is the cylinder; precisely, one defines the family of data by:

$$\mathbf{a}_{\pi,s,t} = H_{\pi(s),t} dt \quad J_{\pi,s,t} = J.$$

It is important to note that, since $H_{\eta,t} = cr$ for $r \geq r_0$, the curvature of $\mathbf{a}$ vanishes outside of a compact set.

One can consider the finite energy solutions to §3.3.3 for this family of data. Unwinding the definitions, one sees that a solution is a pair (π, u) solving:

$$\partial_s u + J(u)(\partial_t u - X_{\pi(s),t}(u)) = 0.$$

One defines $\mathcal{M}(\eta_1, \eta_0)$ to be the moduli space of such finite energy solutions.

Lemma 41. *For generic perturbation of $H_{\eta,t}$ away from the zeroes of P , all the moduli spaces $\mathcal{M}(\eta_1, \eta_0)$ are cut transversally.*

Proof. Note that, since $\pi(s)$ lies in the neighborhood $\text{Op}(\eta_0) \cup \text{Op}(\eta_1)$ outside of a compact set, $X_{\pi(s),t}$ is s -independent outside of a compact set.

If $\eta_1 \neq \eta_0$, then there is a non-empty interval where $X_{\pi(s),t}$ is actually s -dependent. In that case we can perturb $X_{\eta,t}$ generically away from the zeroes in such a way that $X_{\pi(s),t}$ is modified in a generic fashion only a compact part of the cylinder; such variations are sufficient to ensure transversality (compare with [Sei15, pp. 971]).

In the case when $\eta_1 = \eta_0$, then u solves the Floer differential equation for $(H_{\eta_0,t}, J)$, and so requirement (4) implies u is regular. In other words, the moduli space $\mathcal{M}(\eta_0, \eta_0)$ is simply the moduli space considered in §3.4.1 for the admissible data $(H_{\eta_0,t}, J)$. $\square$

We add one additional requirement to our family Floer data:

- (6) the data $H_{\eta,t}$ is chosen generically away from the zeros of P so that all moduli spaces $\mathcal{M}(\eta_1, \eta_0)$ are cut transversally.

Data $(P, H_{\eta,t}, J)$ which satisfies all of the requirements is said to be admissible for defining the family Floer complex.

The moduli space $\mathcal{M}(\eta_1, \eta_0)$ carries an $\mathbb{R}$ -action given by translation:

$$(\pi(s), u(s, t)) \mapsto (\pi(s + s_0), u(s + s_0, t)).$$

Finally, one defines the differential by the formula:

$$d_{\text{FF}}(\gamma_0 \otimes \eta_0) := \sum_{\gamma_1 \otimes \eta_1} \# \{ [\pi, u] \in \mathcal{M}(\eta_1, \eta_0) / \mathbb{R} : u \text{ joins } \gamma_1, \gamma_0 \} \cdot \gamma_1 \otimes \eta_1;$$

more precisely, we count the one-dimensional components of $\mathcal{M}(\eta_1, \eta_0)$ with the advertised asymptotics.

It is perhaps interesting to note that one can decompose d_{FF} into a sum:

$$d_{\text{FF}} = d_0 + d_1 + d_2 + \dots,$$

where d_i maps a generator $\gamma_0 \otimes \eta_0$ into the piece of CFF generated by terms $\gamma_1 \otimes \eta_1$ where $\text{Index}(\eta_1) = \text{Index}(\eta_0) + i$. A simple inspection proves that d_0 preserves the summand $\text{CF}(H_{\eta_0,t}, J) \otimes \mathbb{Z}/2\mathbb{Z}\eta_0$ and acts as $d \otimes \text{id}$ where d is the usual Floer differential.

The first key lemma of family Floer cohomology is:

Lemma 42. *The family Floer differential squares to zero: $d_{\text{FF}}^2 = 0$.*

Proof. This is outlined in [Hut08, Proposition 3.9] in the context of family Morse homology. We also refer the reader to [Sei15, pp.970] for a Floer cohomology set-up closer to the present context.

Briefly, one shows that the number of ends of the one-dimensional component of $\mathcal{M}(\eta_1, \eta_0)/\mathbb{R}$ equals the matrix entry for d_{FF}^2 of type:

$$\text{CF}(H_{\eta_0,t}, J) \otimes \mathbb{Z}/2\mathbb{Z}\eta_0 \rightarrow \text{CF}(H_{\eta_1,t}, J) \otimes \mathbb{Z}/2\mathbb{Z}\eta_1.$$

Since the number of ends is even, one concludes $d_{\text{FF}}^2 = 0$. We refer the reader who wishes for additional details to the proof of Lemma 43. $\square$

We denote by $\text{HFF}(P, H_{\eta,t}, J)$ the homology of $\text{CFF}(P, H_{\eta,t}, J)$ with respect to d_{FF} .

3.6.2. From Floer cohomology to family Floer cohomology. In this section we explain how to define a map $i : \text{HF}(H_t, J) \rightarrow \text{HFF}(P, H_{\eta,t}, J)$ whenever the slope $c \geq 0$ of $H_{\eta,t}$ equals the slope of H_t . This map will be referred to as the *comparison map*.

To keep things as simple as necessary, we assume that H_t satisfies $H_t = cr$ for $r \geq r_0$, i.e., we disallow time-dependent slopes in H_t , and that the same fixed almost complex structure is used for (H_t, J) and $(P, H_{\eta,t}, J)$.

We will construct on $\text{HFF}(P, H_{\eta,t}, J)$:

- (1) a product $\ast_{\text{FF}}$ in §3.6.3,
- (2) a BV operator Δ_{FF} in §3.6.5,
- (3) a map ϵ_{FF} associated to a family of ball embeddings in §3.6.6.

The comparison map i will be shown to respect these structures.

To define i , we follow the usual strategy of defining *comparison data*, and then counting the rigid elements in an associated moduli space.

Fix (H_t, J) and $(H_{\eta,t}, J, P)$ which are admissible for defining the Floer complex and family Floer complex. Define *comparison data* to be a connection one-form $\mathfrak{a}$ on the family $N \times \Sigma \times W$ where Σ is the cylinder, satisfying:

- (1) $\mathfrak{a} = H_{\eta,s,t} dt$,
- (2) $H_{\eta,s,t} = cr$ for $r \geq r_0$,
- (3) $H_{\eta,s,t} = H_t$ for $s \geq s_0$ and $H_{\eta,s,t} = H_{\eta,t}$ for $s \leq -s_0$.

One easily shows that the space of comparison data is convex and non-empty; to see it is non-empty, one can simply take the linear interpolation (bearing in mind that we assume that $H_{\eta,t}, H_t$ both equal cr for $r \geq r_0$).

As in the definition of d_{FF} , we do not directly consider the solutions of the moduli space associated to this family (indeed, this family has asymptotics which depend on η , which would require special treatment). Rather, we will pull back this data using flow lines of the pseudogradient P .

For each zero η_1 , introduce $\mathcal{P}(\eta_1)$ as the space of flow lines $\pi(s)$ of $-\beta(-s)P$ which converge to η_1 as $s \rightarrow -\infty$. Notice that $\pi(s)$ is constant for $s \geq 0$.

The space of points appearing as $\pi(0)$ for $\pi \in \mathcal{P}(\eta_1)$ is simply the unstable manifold of η_1 .

Comparison data induces a connection one-form $\mathfrak{a}$ on $\mathcal{P}(\eta_1) \times \Sigma \times W$ given by the formula:

$$\mathfrak{a}_{\pi,s,t} = \mathfrak{a}_{\pi(s),s,t},$$

and it is important to note that this family has constant asymptotics, namely:

- (1) $\mathfrak{a}_{\pi,s,t} = H_{\eta_1,t}$ for $s \leq -s_1(\pi)$,
- (2) $\mathfrak{a}_{\pi,s,t} = H_t$ for $s \geq s_0$.

The threshold $s_1(\pi)$ is not constant throughout the family (note that $\mathcal{P}(\eta_1)$ is an open manifold), but it is locally constant which is sufficient for the framework established in §3.2.3 and §3.3.

For a generic perturbation term $\mathfrak{p}$ on $\mathcal{P}(\eta_1) \times \Sigma \times W$, and using the almost complex structure J , one consider the moduli space $\mathcal{M}(\eta_1)$ of finite energy solutions (π, u) of §3.3.3. To ground the discussion with a concrete formula, here is the equation we are counting when $\mathfrak{p} = 0$:

$$\partial_s u + J(u)(\partial_t u - X_{\pi(s),s,t}(u)) = 0;$$

the perturbation $\mathfrak{p}$ simply changes the right hand side from 0 to a small vector field.

By counting the rigid elements in $\mathcal{M}(\eta_1)$ whose right asymptotic is equal to $\gamma \in \text{CF}(H_t, J)$, we obtain chains $i_{\eta_1}(\gamma) \in \text{CF}(H_{\eta_1,t}, J)$. The comparison map i is the sum over all zeros η_1 of P :

$$i(\gamma) := \sum i_{\eta_1}(\gamma) \otimes \eta_1.$$

The key lemma concerning this map is:

Lemma 43. *The map $i : \text{CF}(H_t, J) \rightarrow \text{CFF}(P, H_{\eta,t}, J)$ is a chain map with respect to d and d_{FF} . The chain homotopy class is independent of the generic perturbation $\mathfrak{p}$ or the precise comparison data used.*

Proof. This argument has no surprising parts. However, the moduli spaces we consider are less standard, and so we attempt to give a bit more detail than we have in previous arguments.

First we will prove that i is a chain map. The key is to consider the one-dimensional component $\mathcal{M}_1(\eta_1)$. It is convenient to focus on the one-dimensional components which contain solutions (π, u) where u has asymptotics γ_1, γ_0 . Let us refer to this 1-dimensional manifold by $\mathcal{M}_1(\eta_1, \gamma_1, \gamma_0)$.

As usual with parametric moduli spaces, this admits a smooth map to $\mathcal{P}(\eta_1)$ simply given by $(\pi, u) \mapsto \pi$. Because $\mathcal{P}(\eta_1)$ is an open manifold, there are two possible failures of compactness for a sequence $(\pi_n, u_n) \in \mathcal{M}_1(\eta_1, \gamma_1, \gamma_0)$:

- (1) (π_n, u_n) has no convergent subsequence, but π_n does;
- (2) π_n has no convergent subsequence.

Each non-compact component of $\mathcal{M}_1(\eta_1, \gamma_1, \gamma_0)$ has two non-compact ends, and each such end is either of type (1) or (2).

In the case of an end of type (1), the usual Floer compactness-up-to-breaking arguments imply that we can pass to a subsequence so that π_n converges to

π and u_n breaks into a configuration of a rigid element in $\mathcal{M}_0(\eta_1)$ connected to a non-stationary Floer differential cylinder for $(H_{\eta_1,t}, J)$ at the left end, or for (H_t, J) at the right end. By consideration of dimensions, the Floer differential cylinders which broke-off live in one-dimensional families, and hence are counted by the Floer differential. The gluing result complementary to this compactness-up-to-breaking result proves that the number of ends of type (1) equals the coefficient of $\gamma_1 \otimes \eta_1$ appearing in:

$$d_0(i(\gamma_0)) + i(d(\gamma_0)),$$

where we recall $d_{\text{FF}} = d_0 + d_1 + \dots$. Thus, to complete the argument, it suffices to prove that the number of ends of type (2) equals the coefficient of $\gamma_1 \otimes \eta_1$ appearing in:

$$d_1(i(\gamma_0)) + d_2(i(\gamma_0)) + \dots$$

Let us therefore focus on an end (π_n, u_n) of type (2). Because P is assumed to be a Morse-Smale pseudogradient, one can pass to a subsequence so that π_n breaks into a configuration in the product $(\pi_1, \pi_0) \in \mathcal{P}(\eta_1, \eta') \times \mathcal{P}(\eta')$ where the index of η' is strictly less than the index of η_1 ; see Figure 6.




FIGURE 6. Morse theoretic breaking of the flow lines into two pieces. The circle around η_0 signifies the open set $\text{Op}(\eta')$ where $H_{\eta,t} = H_{\eta',t}$.

As π_n breaks into (π_1, π_0) , the equation which u_n solves separates into two equations, in the following sense: $u_n(s + s_n, t)$ will solve the equation for d_{FF} on compact subsets, for any sequence $s_n \rightarrow \infty$, while the non-translated solution $(\pi_n, u_n(s, t))$ converges to a solution $(\pi_0, u_+) \in \mathcal{M}(\eta')$. By picking s_n correctly, the translated solution $(\pi_n(s + s_n), u_n(s + s_n))$ will converge to a solution (π_1, u_-) . By consideration of dimensions, (π_0, u_+) is a rigid element of $\mathcal{M}(\eta')$, and (π_1, u_-) is a rigid-up-to-translation element of the moduli space used to define d_k , where $k > 0$ is the index difference of η_1 and η' . Following similar gluing theory as in [Sei15, pp. 972], each such configuration actually appears as a non-compact end of type (2), and thereby one shows:

$$0 = i(d(\gamma_0)) + d_0(i(\gamma_0)) + d_1(i(\gamma_0)) + d_2(i(\gamma_0)) + \dots \text{ modulo } 2,$$

because each coefficient in the output is the count of the non-compact ends of a one-manifold. This completes the proof that i is a chain map.

The proof that the chain homotopy class is independent of the perturbation $\mathbf{p}$ or the comparison data follows similar lines, and we omit the details. $\square$

Remark. The appeal to gluing theory, while appearing non-standard, actually follows from a general parametric gluing result for continuation cylinders; this is because the equation which u solves near a breaking can be considered as a continuation cylinder for concatenated continuation data (varying in a

parameter space); see Figure 7. The arguments in [Sal97] can be employed in such a case.

$$\gamma_1 \left(\left(\partial_s u + J(u)(\partial_t u - X_{\eta',t}(u)) = 0 \right) \gamma_0 \right)$$

FIGURE 7. Solution near the breaking; on a large subcylinder (determined by $\pi(s) \in \text{Op}(\eta')$), the equation appears as Floer's equation.

3.6.3. The family pair-of-pants product. The product structure on family Floer cohomology is defined as a combination of the Morse cohomology product, defined using flow trees as in [Fuk97], and the pair-of-pants product from §3.4.4. For details on a different Floer theoretic product combining flow lines and pairs-of-pants, we refer the reader to [Sei15, §4.3].

First we introduce a framework for flow trees: having fixed a Morse-Smale pseudogradient P on the parameter space N , introduce time-dependent vector fields $P_{0,s}$, $P_{1,s}$, which are defined for $s \in [-1, \infty)$, vanish when $s \in [-1, 1]$, and agree with P when $s \in [2, \infty)$. Let $P_{\infty,s} = \beta(-s-1)P$.

A *flow tree* is a configuration $(\pi_0, \pi_1, \pi_\infty)$ where π_i is a flow line for $-P_{i,s}$, defined on $[-1, \infty)$ when $i = 0, 1$ and on $(-\infty, 1]$ when $i = \infty$, and such that $\pi_0(0) = \pi_1(0) = \pi_\infty(0) = \eta'$.

Remark. The fact that $\pi_i(s)$ is defined for $s \in [-1, 1]$ and is constant on this interval will be a convenience in some of the subsequent formulas.

A flow tree has asymptotic zeros $\eta_0, \eta_1, \eta_\infty$ of P at its non-compact ends; the space of flow lines with these asymptotics is denoted $\mathcal{T}(\eta_0, \eta_1, \eta_\infty)$.

Notice that the junction point η' lies in the intersection of (deformations of) the stable manifolds of η_1, η_0 and the unstable manifold of η_∞ . In particular, assuming these deformed stable and unstable manifolds are transverse, then the possible choices for η' form a (potentially open) manifold of dimension:

$$\dim \mathcal{T}(\eta_0, \eta_1, \eta_\infty) = \text{Index}(\eta_\infty) - \text{Index}(\eta_0) - \text{Index}(\eta_1).$$

This dimension is also the dimension of the space of flow trees, since the junction point determines the flow tree.

Next we explain how to set-up a family of connection one-forms on the pair-of-pants parametrized by the space of flow trees. To do this, define *family pair-of-pants data* to be:

- (1) a fixed almost complex structure J ,
- (2) families of Hamiltonian functions $H_{i,\eta,t}$, $i = 0, 1, \infty$ such that $(H_{i,\eta,t}, J)$ is admissible for defining the family Floer complex,
- (3) satisfying $H_{0,\eta,t} + H_{1,\eta,t} = H_{\infty,\eta,t}$.

For each flow tree $(\pi_0, \pi_1, \pi_\infty)$, one considers the families of Hamiltonians $H_{i,\pi_i(s),t}$ defined on the “branches” of the flow tree; at the junction point, one has three Hamiltonians related by $H_{0,\eta',t} + H_{1,\eta',t} = H_{\infty,\eta',t}$.

To lift this to the pair-of-pants surface, we follow an ad hoc recipe. Fix the pair-of-pants surface $\Sigma = \mathbb{C} \setminus \{0, 1\}$, as in §3.4.4, and consider the punctured




FIGURE 8. Illustration of a flow tree.

disks $D(1/3)^\times$, $1 + D(1/3)^\times$, and $\mathbb{C} \setminus D(2)$, as cylindrical ends, parametrized in the standard way so that the line $t = 0$ is aligned with the positive real axis. The cylindrical ends around $0, 1$ are parametrized by $[0, \infty) \times \mathbb{R}/\mathbb{Z}$ while the end around ∞ is parametrized by $(-\infty, 0] \times \mathbb{R}/\mathbb{Z}$.

Fix smooth functions $t_i : \Sigma \rightarrow \mathbb{R}/\mathbb{Z}$ such that:

$t_i = t$ in the i th cylindrical end and ∞ th cylindrical end,

and so that t_0 is constant in the 1 cylindrical end while t_1 is constant in the 0 cylindrical end. The differentials $\mathbf{m}_i = dt_i$ give two closed real-valued differential forms on Σ such that:

$\mathbf{m}_i = dt$ in the i th cylindrical end and ∞ th cylindrical end;

note that $\mathbf{m}_0$ vanishes in the 1 cylindrical end while $\mathbf{m}_1$ vanishes in the 0 cylindrical end.

One more piece of data needed to lift the equation to the pair-of-pants is a smooth map $b : \Sigma \rightarrow \mathbb{R}$ satisfying $b(z) = s(z)$ when z is in any of the three cylindrical ends. The values of $b(z)$ should be contained in $[-1, 1]$ on the complement of the cylindrical ends.

For each flow tree $\pi = (\pi_0, \pi_1, \pi_\infty)$ define:

$$\mathbf{a}_\pi = \begin{cases} H_{\infty, \pi_\infty(b(z)), t} dt & \text{in the } \infty\text{th cylindrical end,} \\ H_{0, \pi_0(b(z)), t_0(z)} \mathbf{m}_0 + H_{1, \pi_1(b(z)), t_1(z)} \mathbf{m}_1 & \text{otherwise.} \end{cases}$$

Notice that, for z outside of the cylindrical ends, $\pi_i(b(z)) = \eta'$. This explains our choice of having π_i defined on $[-1, 1]$ for all branches.

The key outcome of this construction is:

Lemma 44. *The family of connection one-forms on $\mathcal{T}(\eta_0, \eta_1, \eta_\infty) \times \Sigma \times W$ induced by $\mathbf{a}$ is smooth and has curvature bounded from above. Moreover:*

$$\mathbf{a} = H_{i, \pi_i(s), t} dt$$

holds in the i th cylindrical end.

Proof. To see that $\mathbf{a}$ has curvature bounded from above, one only needs to observe that, outside of a compact set one has $H_{i, \eta, t} = c_i r$, and therefore the curvature of $\mathbf{a}$ vanishes outside of a compact set because the forms $\mathbf{m}_i$ are closed.

The statement about the form of $\mathbf{a}$ in the cylindrical ends follows immediately from the construction of b , t_i , and $\mathbf{m}_i$. $\square$

The data of $\mathbf{a}$ on the family $\mathcal{T}(\eta_0, \eta_1, \eta_\infty) \times \Sigma \times W$ together with the fixed almost complex structure J and a generic perturbation term $\mathbf{p}$ leads to a moduli space $\mathcal{M}(\eta_0, \eta_1, \eta_\infty)$ of finite energy solutions to §3.3.3. Counting the rigid

elements asymptotic to orbits $\gamma_0, \gamma_1, \gamma_\infty$ gives a number $N_{\eta_0, \eta_1, \eta_\infty}(\gamma_0, \gamma_1, \gamma_\infty)$ in $\mathbb{Z}/2\mathbb{Z}$. Define:

$$(\gamma_0 \otimes \eta_0) *_{\text{FF}} (\gamma_1 \otimes \eta_1) := \sum N_{\eta_0, \eta_1, \eta_\infty}(\gamma_0, \gamma_1, \gamma_\infty)(\gamma_\infty \otimes \eta_\infty),$$

where the sum is over all η_∞ and γ_∞ . As expected:

Lemma 45. *The operation $*_{\text{FF}}$ is a chain map:*

$$*_{\text{FF}} : \text{CFF}(P, H_{0,\eta,t}, J) \otimes \text{CFF}(P, H_{1,\eta,t}, J) \rightarrow \text{CFF}(P, H_{\infty,\eta,t}, J);$$

the chain homotopy class of the map is independent of the choice of $\mathfrak{m}_i, t_i$, the perturbed vector fields $P_{i,s}$, and the perturbation one-form $\mathfrak{p}$.

Proof. This follows standard lines; to see that it is a chain map, one inspects the non-compact ends of the one-dimensional components of $\mathcal{M}(\eta_0, \eta_1, \eta_\infty)$ in a manner similar to the proof of Lemma 43. Let us note that a key part of the argument is understanding the failures of compactness in the moduli space of flow trees $\mathcal{T}(\eta_0, \eta_1, \eta_\infty)$; the same considerations used to prove the flow tree product on Morse cohomology is a chain map will be used here.

To see that the chain homotopy class is independent of the auxiliary choices, one needs to find a path between two choices, then set-up a parametric moduli space and apply the usual Floer theory arguments (see, e.g., [Abo15, pp. 314 and pp. 341]). We only comment on why one can find a path between two such choices: clearly one can find paths between the choices of $\mathfrak{P}_{i,s}$, b , and $\mathfrak{p}$ (simply by a linear interpolation). To find a path between the choices of t_0, t_1 and, say, t'_0, t'_1 , one can consider the circle-valued functions $t_i - t'_i$; by construction, these functions vanish in at least two-out-of-three cylindrical ends. Any such circle-valued function necessarily induces the zero map $\pi_1(\Sigma) \rightarrow \pi_1(\mathbb{R}/\mathbb{Z})$, and thus lifts to $\mathbb{R}$. The space of $\mathbb{R}$ -valued functions is convex, and hence the desired path can be taken to be a linear interpolation between the lifts. $\square$

As mentioned in §3.6.2, the map $i : \text{HF}(H_t, J) \rightarrow \text{HF}(P, H_{\eta,t}, J)$ is compatible with the product structures, in the following sense:

Lemma 46. *Given Hamiltonians $H_{i,t}$ and $H_{i,\eta,t}$, $i = 0, 1, \infty$, such that:*

- (1) *the slope of $H_{i,t}$ equals the slope of $H_{i,\eta,t}$*
- (2) *$(H_{i,t}, J)$ and $(P, H_{i,\eta,t}, J)$ are admissible for defining CF and CFF*
- (3) *$H_{\infty,t} = H_{0,t} + H_{1,t}$, and $H_{\infty,\eta,t} = H_{0,\eta,t} + H_{1,\eta,t}$,*

then we have an equality:

$$*_{\text{FF}} \circ (i \otimes i) = i \circ *,$$

of maps $\text{HF}(H_{0,t}, J) \otimes \text{HF}(H_{1,t}, J) \rightarrow \text{HFF}(P, H_{\infty,\eta,t}, J)$.

Proof. The argument has no surprises; one simply follows their nose. The strategy is as follows: define a parametric moduli space $\mathcal{M}_{\text{param}}$ admitting a map τ to $\mathbb{R}$. The one-dimensional components of $\mathcal{M}_{\text{param}}$ have non-compact ends of three types:

- (1) ends containing sequences with $\tau \rightarrow \infty$; these ends will be asymptotic to the configurations composing $*_{\text{FF}} \circ (i \otimes i)$;

- (2) ends containing sequences with $\tau \rightarrow -\infty$; these ends will be asymptotic to the configurations composing $i \circ *$;
- (3) ends which project under τ to precompact sets in $\mathbb{R}$; these ends will be asymptotic to chain homotopy terms.

The total count of ends is even, and one concludes an equation of the form:

$$(14) \quad *_\text{FF} \circ (i \otimes i) + i \circ * + d_\text{FF} K + K(d \otimes \text{id} + \text{id} \otimes d) = 0 \text{ mod } 2,$$

as maps on the chain complexes; each of the three summands corresponds to one type of ends. The desired result follows.

We now describe the construction of the parametric moduli space $\mathcal{M}_\text{param}$.

Let us define the space $\mathcal{F}$ of pairs (τ, π) where $\pi = (\pi_0, \pi_1, \pi_\infty)$ is a flow tree of the following type:

- (1) $\pi_i : [-1, \infty) \rightarrow N$ is a flow line for $-\beta(\tau - s)P_{i,s}$, $i = 0, 1$,
- (2) $\pi_\infty : (-\infty, 1] \rightarrow N$ is a flow line for $-\beta(\tau - s)P_{\infty,s}$,
- (3) $\pi_0(0) = \pi_1(0) = \pi_\infty(0) = \eta'$;

here β is the standard cut-off function. Notice that $\pi_i(s)$ is constant for $s \geq \tau$; consequently, each element $(\tau, \pi) \in \mathcal{F}$ is completely determined by the parameter value τ and the junction point η' which must be a point in the unstable manifold of $\eta_\infty = \lim_{s \rightarrow -\infty} \pi_\infty(s)$. In particular, if we let $\mathcal{F}(\eta_\infty)$ be the subset of flow trees with fixed asymptotic η_∞ , then $\mathcal{F}(\eta_\infty)$ is an open manifold diffeomorphic to $\mathbb{R} \times (\text{unstable manifold of } \eta_\infty)$.

Following a similar construction used in the definition of $*_\text{FF}$, we obtain a connection one form $\mathfrak{a}$ on the family $\mathcal{F}(\eta_\infty) \times \Sigma \times W$, as follows:

$$\mathfrak{a}_{\tau, \pi} := \begin{cases} H_{\infty, b(z) - \tau, \pi_\infty(b(z)), t} dt & \text{in the } \infty\text{th cylindrical end,} \\ H_{0, b(z) - \tau, \pi_0(b(z)), t_0(z)} \mathfrak{m}_0 + H_{1, b(z) - \tau, \pi_1(b(z)), t_1(z)} \mathfrak{m}_1 & \text{otherwise,} \end{cases}$$

where:

$$H_{i, s, \eta, t} = (1 - \beta(s))H_{i, \eta, t} + \beta(s)H_{i, t}.$$

As in the proof of Lemma 44, this $\mathfrak{a}$ has curvature bounded from above.

Morally, this $\mathfrak{a}_{\tau, \pi}$ is a sort of hybrid between a continuation from $H_{i, t}$ to $H_{i, s, \eta, t}$ and the family pair-of-pants product. The “continuation part” is in the region where $b(z) \approx \tau$, which occurs in the positive ends $i = 0, 1$ when τ is large and positive, and is in the negative end $i = \infty$ when τ is large and negative. As $\tau \rightarrow \infty$, this equation “breaks” into a concatenation of the equations defining the i -map at the 0 and 1 punctures and the family pair-of-pants product. When $\tau \rightarrow -\infty$, the equation breaks into a concatenation of the equation defined by:

$$\mathfrak{a}_{-\infty} = \begin{cases} H_{\infty, t} dt & \text{in the } \infty\text{th cylindrical end,} \\ H_{0, t_0(z)} \mathfrak{m}_0 + H_{1, t_1(z)} \mathfrak{m}_1 & \text{otherwise,} \end{cases}$$

(which is just the non-family pair-of-pants product and does not depend on any flow tree) and the equation defining the i -map.

The desired moduli space $\mathcal{M}_\text{param}$ is a union of components $\mathcal{M}_\text{param}(\eta_\infty)$; the ends of this component compose the terms in (14) with output in the summand $\text{CF}(H_{\infty, \eta_\infty, t}, J) \otimes \mathbb{Z}/2\mathbb{Z}\eta_\infty$.

This component $\mathcal{M}_{\text{param}}(\eta_\infty)$ is defined using the above $\mathfrak{a}$, the fixed almost complex structure J , and a perturbation term $\mathfrak{p}$ on $\mathcal{F}(\eta_\infty) \times \Sigma \times W$. One constructs $\mathfrak{p}$ “recursively,” in the following sense: if (τ, π) is close to breaking (either when $\tau \rightarrow \pm\infty$, or π approaches the boundary of the unstable manifold of η_∞), $\mathfrak{p}$ should be determined by perturbation terms chosen for the equations which appear in the breaking. Such recursive choices of perturbations are a standard ingredient in Floer theory (see, e.g., [Sei08, pp.109]).

The analysis of the non-compact ends of $\mathcal{M}_{\text{param}}(\eta_\infty)$ and the derivation of the desired chain-level equation (14) follows similar lines to the proof of Lemma 43, and we omit further details. $\square$

3.6.4. Special Hamiltonians associated to a family of ball embeddings. In this section, we will describe a particular choice of data $H'_{\eta,t}$, $\eta \in N$, associated to a lift of f to $\mathfrak{B}(a, \Omega)/U(n)$.

Recall from (11) in §3.5.1 the Hamiltonian function $H_{c,\delta,\epsilon,\eta}$ which has a minimum located at the center of the ball B_η . The parameters c, δ, ϵ are explained in §3.5. We assume that $\epsilon a > c(1 + \delta/2)$ as this ensures the evaluation map is defined; see §3.5.2.

Given a pseudogradient P on N , the family $\eta \mapsto H_{c,\delta,\epsilon,\eta}$ is not valid data for the family Floer complex, since it is probably not constant in neighborhoods of the zeros of P . We correct for this by modifying F as follows: simply precompose F using a smooth map $N \rightarrow N$ which is close to the identity in the C^0 distance (associated to a Riemannian metric g) and is constant on neighborhoods of the zeros of P . Provided the C^0 distance is small enough, the modified F is homotopic to the original F .

Fix an almost complex structure J . We define our family as:

$$H'_{\eta,t} = H_{c,\delta,\epsilon,\eta} + \kappa_{\eta,t},$$

where $\kappa_{\eta,t}$ vanishes in the ball B_η , is supported in $\Omega(1 + \delta)$, and is locally η -independent whenever η is in a neighborhood of the zeros of P . We require that $(P, H'_{\eta,t}, J)$ is admissible for defining the family Floer complex (this can be achieved if κ_t is chosen generically).

It will be important when considering action filtrations on $\text{CFF}(P, H'_{\eta,t}, J)$ to make the following quantity very small:

$$(15) \quad \max_{\eta \in N} \|\partial_\eta H'_{\eta,t}\|_g \times (\max g\text{-length of flow lines of } P).$$

This quantity can be made arbitrarily small by picking the pseudogradient P to have only very short flow lines;⁴ we note the size of $\partial_\eta H'_{\eta,t}$ is uniformly controlled during the construction (for each P one can choose a modification of F so that the η derivative of $H'_{c,\delta,\epsilon,\eta}$ is bounded independently of P).

⁴The construction of a pseudogradient with only very short flow lines is an exercise left for the reader. For a similar result see [EP22, Lemma 2.6].

3.6.5. The family BV-operator. The goal in this subsection is to construct the operator Δ_{FF} . At the end, we will analyze how Δ_{FF} acts on the special family $H'_{\eta,t}$ introduced in §3.6.4.

Let $(P, H_{\eta,t}, J)$ be admissible for defining the family Floer complex. The definition of Δ_{FF} is straightforward; we define a connection one-form $\mathfrak{a}$ on the family $\mathcal{P}(\eta_1, \eta_0) \times \mathbb{R}/\mathbb{Z} \times \Sigma \times W$, where Σ is the cylinder, by the formula:

$$\mathfrak{a}_{\pi, \theta, s, t} = [(1 - \beta(s))H_{\pi(s), t+\theta} + \beta(s)H_{\pi(s), t}]dt.$$

Using the almost complex structure J , and a generic perturbation one-form $\mathfrak{p}$, we have an associated moduli space $\mathcal{M}(\eta_1, \eta_0)$ for each pair η_1, η_0 . Counting the rigid elements produces a map $\Delta_{\text{FF}, \eta_0, \eta_1} : \text{CF}(H_{\eta_0, t}, J) \rightarrow \text{CF}(H_{\eta_1, t}, J)$, and we define:

$$\Delta_{\text{FF}}(\gamma_0 \otimes \eta_0) = \sum \Delta_{\text{FF}, \eta_0, \eta_1}(\gamma_0) \otimes \eta_1,$$

where the sum is over all zeros η_1 . Similar arguments to those in §3.6.2 and §3.6.3 prove that Δ_{FF} is a chain map $\text{CFF}(P, H_{\eta, t}, J) \rightarrow \text{CFF}(P, H_{\eta, t}, J)$, and the chain homotopy class is independent of the perturbation term $\mathfrak{p}$.

As with $*_{\text{FF}}$, one has $\Delta_{\text{FF}} \circ i = i \circ \Delta$ as maps $\text{HF}(H_t, J) \rightarrow \text{HFF}(P, H_{\eta, t}, J)$, provided that $H_t, H_{\eta, t}$ have the same slope c .

In the rest of this subsection, we will analyze how the map Δ_{FF} acts on the specific family $H'_{\eta, t}$ constructed in §3.6.4. We will show:

Proposition 47. *If the length of flow lines of P are sufficiently short, and the perturbation one-form $\mathfrak{p}$ and the perturbation $\kappa_{\eta, t}$ used in the definition of $H'_{\eta, t}$ are sufficiently small, then the following holds: any chain:*

$$\sum_{i=1}^k \gamma_i \otimes \eta_i$$

in the output of Δ_{FF} is such that γ_i , $i = 1, \dots, k$, is not the center of B_{η_i} .

Proof. First, we show that, if γ_0 is the center of the ball B_{η_0} , then $\Delta_{\text{FF}}(\gamma_0 \otimes \eta_0)$ does not contain any term $\gamma_i \otimes \eta_i$ where η_i is the center of the ball B_{η_i} . The key idea is to exploit the dimension of the moduli space $\mathcal{M}(\eta_1, \eta_0)$.

Suppose there exists a solution $(\pi, u) \in \mathcal{M}(\eta_1, \eta_0)$ which joins the center x_1 of the ball B_{η_1} at the left end to the center x_0 of the ball B_{η_0} at the right end. Pick a generic section $\mathfrak{s}$ of $\det_{\mathbb{C}} TW$ which is non-vanishing at x_0 and x_1 , and moreover is homotopic through non-vanishing sections to the standard trivialization of $\det_{\mathbb{C}} TW_{x_0}$ and $\det_{\mathbb{C}} TW_{x_1}$; this is trivial if x_0, x_1 are different points, and follows by an easy construction when $x_0 = x_1$, provided we assume that π is short enough that the balls $B_{\pi(s)}$ all contain the point $x_0 = x_1$.

As is well-known (see, e.g., [Can22]), the zero set $\mathfrak{s}^{-1}(0)$ is Poincaré dual to the first Chern class of W , and the dimension of $\mathcal{M}(\eta_1, \eta_0)$ near (π, u) is:

$$1 + \dim \mathcal{P}(\eta_1, \eta_0) + 2[u] \cdot \mathfrak{s}^{-1}(0);$$

the Conley-Zehnder indices do not appear because the orbits at the center of the ball have the same indices when they are computed using the homotopy

class of trivializations induced by $\mathfrak{s}$. Thus, if we can prove that u is null-homologous (bearing in mind that u is topologically a sphere), then (π, u) cannot lie in a rigid component of $\mathcal{M}(\eta_1, \eta_0)$; this gives the desired result. The rest of the proof is dedicated to showing that u must be null-homologous provided the perturbations are small enough, and the flow lines of the pseudogradient are short enough.

To prove that the sphere u is null-homologous we will argue that the diameter of each loop $t \mapsto u(s, t)$ is smaller than the injectivity radius, and hence u bounds a three-dimensional ball. To show this, we will analyze the equation, and estimate the energy of u in terms of $\mathfrak{p}, \kappa_t$ and the length of flow lines.

Unpacking the definitions, one sees that u solves:

$$\partial_s u + J(u)(\partial_t u - X_{\pi(s)}(u)) = V_{\pi, s, t}(u),$$

where $V_{\pi, s, t}$ is due to the perturbation one-form $\mathfrak{p}$ and perturbation term κ , and X_η is the Hamiltonian vector field for $H_{c, \delta, \epsilon, \eta}$.

By construction, $V_{\pi, s, t}$ is compactly supported in the cylinder, and can be taken to be as small as desired. The energy integral of u is equal to:

$$\int \frac{\partial H_{c, \delta, \epsilon, \pi(s)}}{\partial s}(u(s, t)) ds dt + \text{Error},$$

where the error term depends only on $\kappa_t, \mathfrak{p}$, and can be made as small as desired. We then estimate:

$$|\partial_s H_{c, \delta, \epsilon, \pi(s)}| \leq \max_{\eta} |\partial_\eta H_{c, \delta, \epsilon, \eta}|_g |\pi'(s)|_g.$$

Integrating this over the cylinder, one concludes the energy is bounded by the error plus the length of π times a uniform constant; see the discussion in §3.6.4. Since we assume the length of π is short, we can assume the energy of u is as small as desired.

The proof is finished by appealing to a compactness argument. Suppose we have a sequence of solutions u_n of the above form, with perturbation terms $\mathfrak{p}_n, \kappa_{n, \eta, t}$, lengths of flow lines of P_n , and, consequently, energies all tending to zero. Because ω is tamed by J , we obtain:

$$E_n = \int |\partial_t u_n - X_{\pi_n(s), t}(u_n)|^2 ds dt \rightarrow 0 \text{ as } n \rightarrow \infty,$$

as is well-known in estimates of the Floer theory energy integrals; see, e.g., [Sal97, pp. 12].

By standard bubbling analysis, we can assume $|\partial_s u_n|$ and $|\partial_t u_n|$ are uniformly bounded, say by $C > 0$.

Consider the loops $\gamma_{n, s}(t) = u_n(s, t)$. For any s , there must be a nearby point s' such that:

$$|s - s'| \leq E_n^{1/2} \text{ and } \int_{\mathbb{R}/\mathbb{Z}} |\partial_t \gamma_{n, s'}(t) - X_{\pi_n(s'), t}(\gamma_{n, s'}(t))|^2 dt \leq E_n^{1/2}.$$

By the gradient bound, we conclude that $\gamma_{n, s}(t)$ lies in the $CE_n^{1/2}$ neighborhood of $\gamma_{n, s'}(t)$. Thus it suffices to prove that, if γ_n is a sequence of loops

sampled from $u_n(s, t)$ such that:

$$\int_{\mathbb{R}/\mathbb{Z}} |\partial_t \gamma_n(t) - X_{\eta_n, t}(\gamma_n(t))|^2 dt \leq E_n^{1/2},$$

for some $\eta_n \in N$, then $\gamma_n(t)$ has a subsequence which converges uniformly to a point; this will imply all loops sampled from u have a small enough diameter (for n sufficiently large).

By standard bootstrapping for ODEs, similar to the argument in [BC24, §2.2.2], it follows that a subsequence $\gamma_n(t)$ converges in the C^1 topology to an orbit of $X_{\eta, t}$ for some $\eta \in N$. The orbits of $X_{\eta, t}$ are either constants, or have action uniformly far from the action of the center of the ball (see Lemma 38). Assume that γ_n does not converge to a point; it then follows that $\gamma_n(t)$ has action far from the action of the center of the ball as $n \rightarrow \infty$. Since $\gamma_n(t)$ was sampled from $u_n(s, t)$, and u_n joins two centers of balls, we conclude that u_n must have a minimum positive amount of energy (using the well-known principle that the energy integral governs the change in action; the presence of the perturbation terms in the energy integral will not ruin the application of this principle). This minimum amount of energy of u_n contradicts our assumption, and the proof is complete in this case.

The second thing to show is that $\Delta_{\text{FF}}(\gamma_0 \otimes \eta_0)$ does not contain any term $\gamma_i \otimes \eta_i$ where γ_i is the center of B_{η_i} provided γ_0 is *not* the center of B_{η_0} . This case is much easier (for sufficiently small perturbations and lengths of flow lines), since the action of the center has the lowest action among all orbits, and Δ_{FF} increases actions (up to an error which becomes as small as desired as the perturbations and lengths of flow lines tend to zero). $\square$

3.6.6. Evaluation map associated to a family of ball embeddings. In this section we construct a map $\mathfrak{e}_{\text{FF}} : \text{HFF}(P, H_{\eta, t}, J) \rightarrow \mathbb{Z}/2\mathbb{Z}$ using a lift of f to $\mathfrak{B}(a, \Omega)/U(n)$ in a similar way to §3.5.

We assume, as in §3.5, that $\epsilon a > c(1 + \delta/2)$, and that the family $H_{\eta, t}$ has slope at most c .

Consider the family $\mathcal{P}'(\eta_0)$ of flow lines π of $-\beta(s)P$ which are asymptotic at the positive end to the zero η_0 . Each $\pi \in \mathcal{P}'(\eta_0)$ is determined by its terminal point $\eta' = \pi(0)$, and the set of such terminal points is the stable manifold of η_0 .

Define a connection one-form $\mathfrak{a}$ on $\mathcal{P}'(\eta_0) \times \Sigma \times W$ where Σ is the cylinder by the equation:

$$\mathfrak{a}_{\pi, s, t} = (1 - \beta(s))H_{c, \delta, \epsilon, \pi(s)} dt + \beta(s)H_{\pi(s), t} dt;$$

in words, $\mathfrak{a}$ is a continuation data from $H_{\eta_0, t}$ to $H_{c, \delta, \epsilon, \eta'}$ where $\eta' = \pi(0)$. Note that, as in §3.5.2, this connection one-form has a varying asymptotic at the left end, and so some care is needed when considering the associated moduli space.

Fixing a generic perturbation term $\mathfrak{p}$ and an almost complex structure J , we consider the moduli space $\mathcal{M}(\eta_0)$ of solutions (π, u) whose left asymptotic is the central orbit of $H_{c, \delta, \epsilon, \pi(0)}$. This moduli space has similar compactness and regularity properties as if the asymptotics of $\mathfrak{a}$ were fixed, essentially

because the left asymptotic orbit is always non-degenerate (see the discussion in §3.5.2).

Define:

$$\mathfrak{e}_{\text{FF}}(\gamma_0 \otimes \eta_0) := \sum (\text{rigid elements in } \mathcal{M}(\eta_0) \text{ whose right asymptotic is } \gamma_0).$$

The main structural result is:

Lemma 48. *The map $\mathfrak{e}_{\text{FF}} : \text{CFF}(P, H_{\eta,t}, J) \rightarrow \mathbb{Z}/2\mathbb{Z}$ is a chain map, and the chain homotopy class is independent of the generic perturbation term $\mathfrak{p}$.*

Proof. The argument is the same as in §3.5.2, with the modifications needed to work with family Floer cohomology used in, e.g., §3.6.2. $\square$

We now show that $\mathfrak{e}_{\text{FF}}$ is compatible with the i -map.

Lemma 49. *Fix H_t so that (H_t, J) is admissible for defining the Floer complex, the same slope as $H_{\eta,t}$, so that the i -map $\text{CF}(H_t, J) \rightarrow \text{CFF}(P, H_{\eta,t}, J)$ is defined. Then:*

$$\mathfrak{e}_{\text{FF}} \circ i = \mathfrak{e}$$

as maps $\text{HF}(H_t, J) \rightarrow \mathbb{Z}/2\mathbb{Z}$.

Proof. The key is to consider the family $\mathcal{P}''$ of pairs (R, π) where π is a flow line of the vector field: $-\beta(s)\beta(R-s)P$. Note that each such π is locally constant outside of the interval $[0, R]$, and moreover π is determined by the point $\pi(0)$ which is an arbitrary point of N . Thus $\mathcal{P}''$ is diffeomorphic to the product $\mathbb{R} \times N$, via the map $(R, \pi) \mapsto (R, \pi(0))$.

Define a connection one-form $\mathfrak{a}$ on the family $\mathcal{P}'' \times \Sigma \times W$ by:

$$a_{\pi,s,t} = [(1 - \beta(s))H_{c,\delta,\epsilon,\pi(s)} + \beta(s)(\beta(R-s)H_{\pi(s),t} + (1 - \beta(R-s))H_t)]dt;$$

It is important to notice that, for $s \leq R-1$, $\mathfrak{a}_{\pi,s,t}$ agrees with the connection 1-form used to define $\mathfrak{e}_{\text{FF}}$, and on $s \geq 1$, $\mathfrak{a}_{\pi,s,t}$ agrees with the R -translated version of the connection 1-form used to define i .

Fixing a perturbation term $\mathfrak{p}$, we can therefore consider the one-dimensional component of the moduli space $\mathcal{M}$ consisting of triples (R, π, u) .

Claim: for each generic number R_0 , the fiber $\mathcal{M}(R_0)$ of triples (R_0, π, u) is a zero-dimensional manifold, and counting the points in $\mathcal{M}(R_0)$ defines a chain map $\text{CF}(H_t, J) \rightarrow \mathbb{Z}/2\mathbb{Z}$. The chain homotopy class of this map is independent of R_0 .

The claim is standard Floer theory, following similar arguments used in §3.5.2, and we omit further discussion.

The two crucial observations are that:

- (1) $\mathcal{M}(R_0)$ is precisely the moduli space used to define $\mathfrak{e}$, provided $R_0 \leq 0$,
- (2) as $R_0 \rightarrow +\infty$, $\mathcal{M}(R_0)$ “breaks” into configurations of rigid solutions $((\pi_-, u_-), (\pi_+, u_+))$, where $\pi_- \in \mathcal{P}'(\eta_0)$ and $\pi_+ \in \mathcal{P}(\eta_0)$, and where (π_-, u_-) contributes to $\mathfrak{e}_{\text{FF}}$ and (π_+, u_+) contributes to i .

In this manner, we conclude that $\mathfrak{e} = \mathfrak{e}_{\text{FF}} \circ i$, up to chain homotopy. One point which merits comment is that one should pick $\mathfrak{p}$ to be compatible with the breaking as $R_0 \rightarrow \infty$. This completes the proof. $\square$

We end this subsection with an analysis of how $\mathfrak{e}_{\text{FF}}$ acts on $\text{CFF}(P, H'_{\eta,t}, J)$ for the special family $H'_{\eta,t}$ constructed in §3.6.4.

Proposition 50. *If the length of flow lines of P are sufficiently short, and the perturbation one-form $\mathfrak{p}$ and the perturbation $\kappa_{\eta,t}$ used in the definition of $H'_{\eta,t}$ are sufficiently small, then the following holds:*

$$\mathfrak{e}_{\text{FF}}\left(\sum_{i=1}^k \gamma_i \otimes \eta_i\right) = 0$$

provided each γ_i is not the center of B_{η_i} . In the definition of $\mathfrak{e}_{\text{FF}}$ we use the same family of ball embeddings as is used in the family $H'_{\eta,t}$.

Proof. The argument is similar (and easier) than Proposition 47. One estimates that the energy of any solution u appearing in the moduli space used to define $\mathfrak{e}_{\text{FF}}$ is as small as desired; for this, it is important that the same family $H_{c,\delta,\epsilon,\eta}$ is used in the construction of $H'_{\eta,t}$ and in $\mathfrak{e}_{\text{FF}}$. Because the energy governs the action difference, any solution u joins asymptotic orbits γ_-, γ_+ where the action of γ_- is larger than the action of γ_+ up to an arbitrarily small error.

This error can be taken to be smaller than the distance between the action at the center of the ball and the action of any other orbit (see Lemma 38). In particular, since the center of the ball has the lowest action, we conclude that the only possible solutions contributing to $\mathfrak{e}_{\text{FF}}$ must be asymptotic at their positive end to $\gamma_0 \otimes \eta_0$ where γ_0 is the center of the ball B_{η_0} , as desired. $\square$

3.6.7. The family action filtration. In Propositions 47 and 50 we appealed to the actions of orbits in the context of family Floer cohomology, in the context when the flow lines of the pseudogradient P are short. In this subsection, we will formalize such considerations by introducing special action filtrations on the family Floer complex depending on a parameter $\varepsilon > 0$.

For $\varepsilon > 0$, define the ε -action of $\gamma_0 \otimes \eta_0 \in \text{CFF}(P, H_{\eta,t}, J)$ to be the number:

$$A_\varepsilon(P, H_{\eta,t}; \gamma_0 \otimes \eta_0) := \varepsilon \text{Index}(P; \eta_0) + \int H_{\eta_0,t}(\gamma_0(t)) dt - \int \gamma_0^* \lambda,$$

where, recall, λ is the Liouville form on W . One defines the ε -action of a chain in $\text{CFF}(P, H_{\eta,t}, J)$ to be the minimum action of a generator which appears in the chain.

Lemma 51. *The differential d_{FF} increases the ε -action provided:*

$$\max_{\eta,t,u} |\partial_\eta H_{\eta,t}(u)|_g dt \times \max\{g\text{-length of a flow line of } P\} < \varepsilon,$$

for some Riemannian metric g on N .

Proof. Recall that the differential counts $\mathbb{R}$ -families of finite-energy solutions (π, u) where $\pi \in \mathcal{P}(\eta_1, \eta_0)$ is a flow line of $-P$ and $u : \mathbb{R} \times \mathbb{R}/\mathbb{Z} \rightarrow W$ solves:

$$\partial_s u + J(u)(\partial_t u - X_{\pi(s),t}(u)) = 0,$$

where $X_{\eta,t}$ is the Hamiltonian vector field for $H_{\eta,t}$. The space of such solutions admits an $\mathbb{R}$ -action by translating π and u , and the differential counts the rigid elements in the quotient by this $\mathbb{R}$ -action.

Unpacking the definition from §3.3.5, one sees the energy of u is:

$$E(u) := \int \omega(\partial_s u, \partial_t u - X_{\pi(s),t}(u)) ds dt,$$

which is non-negative and can be computed as:

$$E(u) = \omega(u) + \int_0^1 H_{\eta_1,t}(\gamma_1(t)) dt - \int_0^1 H_{\eta_0,t}(\gamma_0(t)) dt + \int \partial_s H_{\pi(s),t}(u) ds dt,$$

where $\gamma_i(t)$ are the asymptotic orbits of u . If $\eta_1 = \eta_0$, then $\pi'(s) = 0$ and $\partial_s H_{\pi(s),t} = 0$. Otherwise, the last term can be estimated:

$$\int \partial_s H_{\pi(s),t}(u) ds dt \leq \max_{\eta,t,u} |\partial_\eta H_{\eta,t}(u)|_g \int_{\mathbb{R}} |\pi'(s)|_g ds \leq \varepsilon.$$

Using $\omega = d\lambda$, one obtains:

$$0 \leq A_\varepsilon(\gamma_1 \otimes \eta_1) - A_\varepsilon(\gamma_0 \otimes \eta_0);$$

this is proved in two cases: first, if $\eta_1 = \eta_0$, and, second, if $\eta_1 \neq \eta_0$, in which case their index difference is at least 1. $\square$

Thus, for ε satisfying the hypotheses of Lemma 51, $(\text{CFF}(P, H_t, J), d_{\text{FF}}, A_\varepsilon)$ is a (cohomologically) filtered complex, and so it makes sense to speak about the filtration level of a homology class (as the maximum action of all representative cycles).

Next we will show that A_ε is compatible with Δ_{FF} up to a bounded error.

Lemma 52. *For any homology class $\alpha \in \text{HFF}(P, H_{\eta,t}, J)$, it holds that:*

$$- \int \max_{\eta,u} |H_{\eta,t}(u) - H_{\eta,t+\theta}(u)| dt \leq A_\varepsilon(\Delta_{\text{FF}}(\alpha)) - A_\varepsilon(\alpha),$$

provided that ε satisfies the hypotheses of Lemma 51.

Remark. In particular, if $H_{\eta,t}$ is close to being autonomous (for each η), then Δ_{FF} is close to being action non-decreasing.

Proof. The argument is similar to Lemma 51. Recall that Δ_{FF} counts the rigid solutions (π, u) where $\pi \in \mathcal{P}(\eta_1, \eta_0)$ and u solves a perturbed version of the equation:

$$(16) \quad \partial_s u + J(u)(\partial_t u - X_{s,t}(u)) = 0$$

where $X_{s,t}$ is the Hamiltonian vector field for $(1 - \beta(s))H_{\pi(s),t+\theta} + \beta(s)H_{\pi(s),t}$. The energy of solutions for the unperturbed equation can be estimated as:

$$\begin{aligned} E(u) &\leq \int_0^1 \max_{\eta,u} |H_{\eta,t}(u) - H_{\eta,t+\theta}(u)| dt + \max_{u,\eta,t} |\partial_\eta H_{\eta,t}|_g \int_{\mathbb{R}} |\pi'(s)|_g \\ &\quad + \omega(u) + \int_0^1 H_{\eta_1,t+\theta}(\gamma_1(t+\theta)) dt - \int_0^1 H_{\eta_0,t}(\gamma_0(t)) dt. \end{aligned}$$

The proof uses the standard energy estimates for continuation map type equation like (16), similarly to Lemma 51.

If $\eta_1 = \eta_0$, then $\pi'(s) = 0$; otherwise the index difference is at least 1. In both cases, one rearranges to obtain the desired result.

In general, one needs to work with the perturbed equation $\mathfrak{p}$; however, since the chain homotopy class is independent of the perturbation, one can prove the estimate with small error terms due to the perturbation $\mathfrak{p}$, and taking a limit as $\mathfrak{p} \rightarrow 0$ will reduce to the above unperturbed analysis; see also the proof of Lemma 53 for further details on this step of the argument. $\square$

To conclude this section, we show that $*_{\text{FF}}$ respects the ε -action filtration, up to an error depending on ε .

Lemma 53. *Suppose that ε satisfies the hypotheses of Lemma 51. Then, for any two homology classes $\alpha_i \in \text{HFF}(P, H_{i,\eta,t}, J)$, $i = 0, 1$, we have:*

$$-(3 + \dim N)\varepsilon - C \max_{\eta, t_0, t_1} |\{H_{0,\eta,t_0}, H_{1,\eta,t_1}\}| \leq A_\varepsilon(\alpha_0 *_{\text{FF}} \alpha_1) - A_\varepsilon(\alpha_0) - A_\varepsilon(\alpha_1),$$

where C is a uniform constant depending only on the pair-of-pants surface; here $\alpha_0 *_{\text{FF}} \alpha_1 \in \text{HFF}(P, H_{\infty,\eta,t}, J)$ and $H_{\infty,\eta,t} = H_{0,\eta,t} + H_{1,\eta,t}$ is assumed to be admissible for defining the family Floer complex as required in the definition of $*_{\text{FF}}$, and $\{-, -\}$ is the Poisson bracket.

Proof. Recall from §3.6.3 that the product is defined by counting rigid solutions to the perturbed equation determined by the connection one-form:

$$\mathfrak{a} = \begin{cases} H_{i,\pi_i(s),t} dt & \text{in the } i \text{ cylindrical end, } i = 0, 1, \infty, \\ H_{0,\eta',t_0(z)} \mathfrak{m}_0 + H_{1,\eta',t_1(z)} \mathfrak{m}_1 & \text{otherwise,} \end{cases}$$

where $(\pi_0, \pi_1, \pi_\infty)$ is a flow tree, η' is the junction point of the flow tree, and the circle valued functions t_i , and their differentials $\mathfrak{m}_i = dt_i$ are as in §3.6.3. The branches of the flow tree are flow lines for perturbations of the pseudogradient.

Thus there are two relevant perturbations which play a role in the definition of the pair-of-pants product:

- (1) the usual perturbation term $\mathfrak{p}$,
- (2) the perturbations used to define the flow tree.

For each solution (π, u) to the family pair-of-pants equation, ideally, we would like to show that:

$$(17) \quad -(3 + \dim N)\varepsilon \leq A_\varepsilon(\gamma_\infty \otimes \eta_\infty) - A_\varepsilon(\gamma_0 \otimes \eta_0) - A_\varepsilon(\gamma_1 \otimes \eta_1),$$

where γ_i, η_i , $i = 0, 1, \infty$ are the asymptotics. We simplify our task as follows: it suffices to prove (17) when the perturbations (1) and (2) are turned off. To see why this is sufficient, recall that the chain homotopy class of the pair-of-pants is independent of the perturbation used. Thus we can take a sequence of solutions (π_n, u_n) solving the equation with perturbations (1) and (2) depending on n which tend to zero.

By the usual compactness theory for solutions to Floer's equation (see §3.3.6) and standard compactness results for ODEs, one concludes that (π_n, u_n) has a subsequence which converges to a configuration consisting of a central limit (π_∞, u_∞) solving the unperturbed pair-of-pants equation, together with some number of cylinders at the punctures solving the equation for d_{FF} ; see the illustration in Figure 9. Since we have already shown that d_{FF} increases the A_ε -action in Lemma 51, it is sufficient to prove (17) holds for solutions of

the unperturbed equation; this completes the explanation of why we can work only with the unperturbed equation.




FIGURE 9. Illustration of the pair-of-pants and limit (cylinders for d_{FF} break off) when perturbations are turned off

Referring to the notation in Figure 9, we claim the following:

$$\begin{aligned}
 (18) \quad & A_0(\gamma_\infty \otimes \eta_\infty) \geq A(H_{\infty, \eta', t}; \ell_\infty(t)) - \varepsilon \\
 & A(H_{0, \eta', t}; \ell_0(t)) \geq A_0(\gamma_0 \otimes \eta_0) - \varepsilon \\
 & A(H_{1, \eta', t}; \ell_1(t)) \geq A_0(\gamma_1 \otimes \eta_1) - \varepsilon \\
 & A(H_{\infty, \eta', t}; \ell_\infty(t)) \geq A(H_{0, \eta', t}; \ell_0(t)) + A(H_{1, \eta', t}; \ell_1(t)),
 \end{aligned}$$

where $A_0(\eta \otimes \gamma)$ is simply the action functional considered in §3.6.7 with ε set to zero, and:

$$A(H_{i, \eta', t}; \ell(t)) = \int H_{i, \eta', t}(\ell(t)) dt - \int \ell^* \lambda,$$

where η' is the junction point of the limit flow tree.

Combining everything, we conclude that:

$$A_0(\gamma_\infty \otimes \eta_\infty) + 3\varepsilon \geq A_0(\gamma_0 \otimes \eta_0) + A_0(\gamma_1 \otimes \eta_1).$$

Finally we observe that, along any unperturbed flow tree we have:

$$\text{Index}(\eta_\infty) \geq \max\{\text{Index}(\eta_1), \text{Index}(\eta_0)\},$$

and hence $\varepsilon \text{Index}(\eta_\infty) + \varepsilon \dim N \geq \varepsilon \text{Index}(\eta_0) + \varepsilon \text{Index}(\eta_1)$; this implies (17), as desired.

It remains only to verify (18). The first three lines of (18) follow from the same exact argument as Lemma 51. The rest of the proof is showing the last line; it will ultimately follow from a computation of the curvature of $\mathfrak{a}$ on the central region Σ_0 of the pair-of-pants shown in Figure 9.

By Lemma 29, it is sufficient to bound the intergral of the connection two-form $\mathfrak{r}$ associated to $\mathfrak{a}$ over the central region Σ_0 . Recall that:

$$\mathfrak{a} = H_{x,y} dx + K_{x,y} dy \implies \mathfrak{r} = (\partial_x K_{x,y} - \partial_y H_{x,y} + \omega(X_{x,y}, Y_{x,y})) dx \wedge dy,$$

where $X_{s,t}, Y_{s,t}$ are the Hamiltonian vector fields for $H_{s,t}, K_{s,t}$, respectively. In our case, above the central region, we have:

$$\mathfrak{a} = H_{0, \eta', t_0(z)} \mathfrak{m}_0 + H_{1, \eta', t_1(z)} \mathfrak{m}_1.$$

We claim that:

$$(19) \quad \mathfrak{r} = \omega(X_{0, \eta', t_0(z)}, X_{1, \eta', t_1(z)}) \mathfrak{m}_0 \wedge \mathfrak{m}_1.$$

To compute this, we first observe that $\partial_x K_{x,y} - \partial_y H_{x,y}$ is linear, and hence it suffices to prove it vanishes for each $H_{i,\eta',t_i(z)} \mathbf{m}_i$ separately. Write $t_i = f_i$, $H_{i,\eta',t_i(z)} = G_{i,f_i(z)}$ and $\mathbf{m}_i = df_i$. Then, in conformal coordinates $x + iy$,

$$G_{i,f_i(z)} df_i = G_{i,f_i(z)} \partial_x f_i dx + G_{i,f_i(z)} \partial_y f_i dy,$$

and a short computation shows the $\partial_x K_{x,y} - \partial_y H_{x,y}$ term vanishes. Thus the only term which contributes to $\mathfrak{r}$ is $\omega(X_{x,y}, Y_{x,y}) dx \wedge dy$, which expands to:

$$\omega(V_{0,f_0(z)} \partial_x f_0 + V_{1,f_1(z)} \partial_x f_1, V_{0,f_0(z)} \partial_y f_0 + V_{1,f_1(z)} \partial_y f_1) dx \wedge dy,$$

where $V_{i,f}$ is the Hamiltonian vector field of $G_{i,f}$. Simplifying, one obtains:

$$\omega(V_{0,f_0(z)}, V_{1,f_1(z)}) (\partial_x f_0 \partial_y f_1 - \partial_y f_0 \partial_x f_1) dx \wedge dy,$$

which equals (19), as desired. Finally, using the notation $v(z) = (z, u(z))$ from Lemma 29, we estimate:

$$\int v^* \mathfrak{r}_\sigma \leq \max_{\eta, t_0, t_1} |\{H_{0,\eta,t_0}, H_{1,\eta,t_1}\}| \int_{\Sigma_0} |\mathbf{m}_0 \wedge \mathbf{m}_1|,$$

which gives the desired result. $\square$

3.6.8. Proof of Theorem 7. Suppose that $H_{c,\delta,\epsilon,\eta}$ is the family as constructed in §3.5.1, using the lift of f to $\mathfrak{B}(a, \Omega)/U(n)$. The proof of Theorem 7 is an argument by contradiction: we assume $a > c$, and then derive a contradiction. As in the hypotheses of Theorem 7, suppose that there are classes $\zeta_i \in V_{c_i}$, $i = 1, 2$, where $c_i > 0$, such that:

$$\text{PSS}(\beta) = \Delta(\zeta_1) * \zeta_2 \text{ in } V_c,$$

where $c = c_1 + c_2$, and where β has non-zero homological intersection with f . Without loss of generality, we can suppose that $c_2 > 0$. Since $a > c$, we can use the evaluation map $\mathfrak{e}$ and §3.5.3 to conclude:

$$1 = \mathfrak{e}(\Delta(\zeta_1) * \zeta_2).$$

Let $\theta_i = c_i/c$, so that $\theta_1 + \theta_2 = 1$. Given a pseudogradient P on N , construct the family $H'_{\eta,t}$ as a small perturbation of $H_{c,\delta,\epsilon,\eta}$ as in §3.6.4. Let:

$$H_{i,\eta,t} = \theta_i H'_{\eta,t}, \text{ for } i = 1, 2,$$

Fix a small constant $\varepsilon > 0$, and pick the perturbation term used in the construction of $H'_{\eta,t}$ so that:

- (1) $H_{1,\eta,t}$, $H_{2,\eta,t}$, and $H_{1,\eta,t} + H_{2,\eta,t}$ are all admissible for defining the family Floer complex,
- (2) $\max_{\eta, t_0, t_1} |\{H'_{\eta,t_0}, H'_{\eta,t_1}\}| < \varepsilon/C$, where C is from Lemma 53.
- (3) the action of any non-central orbit appearing in $H_{1,\eta,t}$ is non-negative.

Lemma 38 implies that we can assume that:

- (4) the action of any orbit for $H_{2,\eta,t}$ is at least $\theta_2(c - a) - \varepsilon$.

By picking the pseudogradient P to have short enough flow lines, we can also suppose that:

- (5) $\max_{\eta,t} |\partial_\eta H'_{\eta,t}|_g \times \max\{g\text{-length of flow lines for } P\} < \varepsilon$.

- (6) the conclusion of Proposition 47 holds for $H_{1,\eta,t}$; here we note that $H_{1,\eta,t}$ is a perturbation of $\theta_1 H_{c,\delta,\epsilon,\eta} = H_{\theta_1 c, \delta, \theta_1 \epsilon, \eta}$, and Proposition 47 applies.

With this established, we claim:

Lemma 54. *For any classes $\zeta_i \in \text{HFF}(P, H_{i,\eta,t}, J)$, we have that ε satisfies the hypotheses of Lemma 51 for each $H_{i,\eta,t}$, so $\mathcal{A}_\varepsilon$ is valid cohomological filtration, and:*

$$\mathcal{A}_\varepsilon(\Delta_{\text{FF}}(\zeta_1) *_{\text{FF}} \zeta_2) \geq \theta_2(c - a) - (5 + \dim N)\varepsilon.$$

Proof. The first statement follows from (5). The next step is to use Proposition 47 and (6) to conclude that:

$$\Delta_{\text{FF}}(\zeta_1) \text{ is represented by cycles not containing any central orbit.}$$

In particular, using (3), we conclude:

$$\mathcal{A}_\varepsilon(\Delta_{\text{FF}}(\zeta_1)) \geq 0.$$

Use this, Lemma 53, and (2), to conclude:

$$\mathcal{A}_\varepsilon(\Delta_{\text{FF}}(\zeta_1) *_{\text{FF}} \zeta_2) \geq \mathcal{A}_\varepsilon(\Delta_{\text{FF}}(\zeta_1)) + \mathcal{A}_\varepsilon(\zeta_2) - (4 + \dim N)\varepsilon,$$

for $j = 1, \dots, k-1$. Then use (3) and (4) to conclude:

$$\mathcal{A}_\varepsilon(\Delta_{\text{FF}}(\zeta_1) *_{\text{FF}} \zeta_2) \geq \theta_2(c - a) - (5 + \dim N)\varepsilon,$$

as desired. $\square$

If the lengths of the flow lines are short enough, then Proposition 50 says $\mathfrak{e}_{\text{FF}} : \text{HFF}(P, H'_{\eta,t}, J) \rightarrow \mathbb{Z}/2\mathbb{Z}$ vanishes on classes which are represented by cycles which contain only non-central orbits. The discussion at the start of this subsection, together with the compatibility of $\mathfrak{e}_{\text{FF}}, \Delta_{\text{FF}}, *_{\text{FF}}$ and their non-family analogues implies:

$$\mathfrak{e}_{\text{FF}}(\Delta_{\text{FF}}(\zeta_1) *_{\text{FF}} \zeta_2) = 1.$$

The contradiction leading to the proof of Theorem 7 will therefore be completed provided we can prove:

$$(20) \quad \Delta_{\text{FF}}(\zeta_1) *_{\text{FF}} \zeta_2 \in \text{HFF}(P, H'_{\eta,t}, J)$$

is represented by a cycle which contains only non-central orbits. This fact follows from the action estimate in Lemma 54, provided ε is small enough. Indeed, Lemma 38 implies that the central orbit has action:

$$c + c\delta/2 - \epsilon a.$$

Picking ϵ close enough to 1, δ small enough, and picking ε small enough, we can use the fixed negative number $c - a$ to ensure:

$$\mathcal{A}_\varepsilon(\text{any central orbit}) \leq c + c\delta/2 - \epsilon a + \varepsilon \dim N < \theta_2(c - a) - (5 + \dim N)\varepsilon,$$

where we use $\theta_2 < 1$ (since $\theta_1 > 0$ and $\theta_1 + \theta_2 = 1$). It therefore follows that (20) must be represented by a cycle which does not contain central orbits. Thus $\mathfrak{e}_{\text{FF}}$ must vanish on it, providing the desired contradiction, proving Theorem 7. $\square$

3.6.9. Proof of Theorem 8. The argument is essentially exactly the same. The first difference is that one defines:

$$H_{c,\delta,\epsilon,\eta} = G_{c,\delta} + K + \epsilon D_\eta,$$

as a minor replacement of the functions defined in §3.5.1.

The second difference f is now assumed to lift to $\mathfrak{B}(a, U)/U(n)$, where U is the well of the Hofer-Zehnder admissible function K .

Everything else follows the same argument. This works because K is independent of $c, \delta, \epsilon, \eta$ (so many terms involving derivatives of $H_{c,\delta,\epsilon,\eta}$ with respect to its parameters vanish), and K only has constant orbits and U is contained in the set where $K = -A = \min K$ (this is used in Lemma 38 which estimates the actions). Because of the Hofer-Zehnder admissible function K , the action of the center of the ball is $c + c\delta/2 - A - \epsilon a$. Thereby one concludes the evaluation map is defined when $c < A + a$. $\square$

3.6.10. Proof of Theorem 13. The argument is very similar to the proof of Theorem 7, and in fact a bit easier. Since ζ_1 is represented by a cycle whose orbits are all non-contractible, ζ_1 is represented by a cycle which does not contain the central orbit. The proof then follows the same lines as §3.6.8. The details are left to the reader. $\square$

4. From string topology to Floer cohomology

The goal of this section is to prove Theorem 11 on the comparison between string topology and Floer cohomology. We specialize to the case when $W = T^*M$ and Ω is a fiberwise starshaped domain.

Much of the arguments in this section are straightforward modifications of the arguments of [AS10, Abo15] from their Morse theory approach to our bordism approach. The exception is the proof that the product structures are identified, as we instead give a seemingly novel adiabatic gluing argument.

Recall from §1.3 the monoid $Z(\Lambda_c)$ of smooth families of loops:

$$A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M,$$

such that P is a compact finite dimensional manifold, and $\ell_\Omega(A(x, -)) < c$ holds for all $x \in P$, where:

$$\ell_\Omega(q) = \int_0^1 \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\} dt$$

is the length of the loop q as measured by Ω . As in §1.3, this leads to the bordism group $H(\Lambda_c)$ as a quotient of $Z(\Lambda_c)$. Then $c \mapsto H(\Lambda_c)$ has the structure of a persistence module on the positive real line.

As explained in §2.1, there are three relevant structures on this persistence module:

- (1) the BV-operator $\Delta : H(\Lambda_c) \rightarrow H(\Lambda_c)$,
- (2) the inclusion of the constant loops $\mathfrak{i} : H^*(W) \rightarrow H(\Lambda_c)$, and,
- (3) the Chas-Sullivan product $*$: $H(\Lambda_{c_1}) \otimes H(\Lambda_{c_2}) \rightarrow H(\Lambda_{c_1+c_2})$.

In the first subsection §4.1, we will explain the definition of a morphism of persistence modules $\Theta : H(\Lambda_c) \rightarrow V_c$; afterwards, we show Θ intertwines $\Delta, \mathbf{i}, *$ with the analogous structures on the Floer cohomology persistence module.

4.1. Definition of the Θ -morphism. The definition of the Θ -morphism follows the strategy of [AS06, APS08, AS10, Abo15, CHO23, BCS25], and uses moduli spaces with moving Lagrangian boundary conditions similar to the ones considered in [Cie94, BC24].

4.1.1. Θ -data. Let (H_t, J) be a Hamiltonian system admissible for defining the Floer complex, and let $A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M$ be a family of loops in $Z(\Lambda_c)$. For such inputs (H_t, J) and A , we define Θ -data to be:

- (1) a Hamiltonian connection $\mathbf{a}$ on $P \times \Sigma \times W$ where Σ is the half-infinite cylinder $(-\infty, 0] \times \mathbb{R}/\mathbb{Z}$, so $\mathbf{a}_{x,s,t} = K_{x,s,t}ds + H_{x,s,t}$, where $K_{x,s,t} = 0$ and $H_{x,s,t} = H_t$ for $s \leq s_0$,
- (2) a smooth family of ω -tame and Liouville equivariant almost complex structures $J_{p,s,t}$ on $P \times \Sigma \times W$ so $J_{p,s,t} = J$ when $s \leq -s_0$,

which satisfies the following properties:

- (3) $H_{x,s,t} = c_{x,s,t}r$ and $K_{x,s,t} = b_{x,s,t}r$ for $r \geq r_0$,
- (4) $\partial_s c_{x,s,t} \leq \partial_t b_{x,s,t}$,
- (5) $c_{x,0,t} \geq \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)} \text{ where } q(t) = A(x, t)\}$, outside of $r \geq r_0$.

Similarly to §3.4.2, condition (4) is used when showing that the curvature of $\mathbf{a}$ is bounded from above, which is used in proving a priori energy estimates. Condition (5) is also used in the a priori energy estimate; see §4.1.2.

Let us note that (5) implies:

$$\int c_{x,0,t} dt \geq \ell_\Omega(q) \text{ for each } q = A(x, -),$$

and so it is necessary that the slope of H_t is at least the ℓ_Ω -length of all loops appearing in the family A .

Lemma 55. *If the slope of H_t is at least the c , where $A \in Z(\Lambda_c)$, then the space of Θ -data for (H_t, J) and A is weakly contractible.*

Proof. The argument is exactly as in Lemma 33; one can take convex combinations between any the connection one forms, and use the contractibility of the space of almost complex structures, to prove the space of Θ -data is either empty or weakly contractible. To prove it is non-empty, one picks a family of smooth functions $H_{x,0,t}$ so that $H_{x,0,t} = c_{x,0,t}r$ for $r \geq r_0$, and:

$$c_{x,0,t} \geq \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)} \text{ where } q(t) = A(x, t)\}.$$

Moreover, we can assume that:

$$(21) \quad \int c_{x,0,t} dt \leq c \leq \text{slope of } H_t \text{ for each } x \in P.$$

Similarly to the proof of Lemma 33, define:

$$\begin{cases} H_{x,s,t} = (1 - \beta(s+1))H_t + \beta(s+1)H_{x,0,t}, \\ K_{x,s,t} = \int_0^t \partial_s H_{x,s,\tau} d\tau - t \partial_s \int_0^1 H_{x,s,\tau} d\tau. \end{cases}$$

Then (21) implies condition (4). Condition (5) holds by construction, and the other properties are obvious. $\square$

4.1.2. A priori energy estimate. The Θ -map will be defined as follows: given Θ -data $(\mathfrak{a}, J)$, we will pick a generic perturbation one-form $\mathfrak{p}$ on $P \times \Sigma \times W$, which we assume vanishes above $s = 0$ and above $s \leq -s_0$. Then we will count the rigid solutions (x, u) to §3.3.3 satisfying the boundary conditions:

$$(22) \quad u(0, t) \in T^*M_{q(t)} \text{ where } q(t) = A(x, t).$$

In this subsection, we show such solutions satisfy a priori energy bound. Similar energy estimates are proved in [BC24, BCS25].

We will use the general energy estimate Lemma 30. First, (4) implies $\mathfrak{a}$ has curvature bounded from above. Therefore Lemma 31 implies $\mathfrak{a}, \mathfrak{p}$ also has curvature bounded from above. Thus, the a priori energy bound follows from:

Lemma 56. *Let $\mathfrak{a}, J$ be Θ -data, and let $\mathfrak{p}$ be a perturbation one-form as above. There is a uniform bound:*

$$\sup\{\omega(u) - \int_{\partial\Sigma} v^* \mathfrak{a}_x\} < \infty$$

where the supremum is over $v(z) = (z, u(z))$ where (x, u) is a finite energy solution of §3.3.3 for $(P \times \Sigma \times W, \mathfrak{a}, \mathfrak{p}, J)$ with boundary conditions (22).

Proof. Unpacking the definitions, one sees that:

$$\int_{\partial\Sigma} v^* \mathfrak{a}_x = \int_0^1 H_{x,0,t}(u(0, t)) dt,$$

and:

$$\omega(u) = \int_0^1 \lambda_{u(0,t)}(\partial_t u(0, t)) dt - \int \gamma^* \lambda,$$

where γ is the asymptotic orbit of u . Since there are only finitely many orbits, it is sufficient to bound:

$$I = \int_0^1 \lambda_{u(0,t)}(\partial_t u(0, t)) - H_{x,0,t}(u(0, t)) dt,$$

as can be verified by inspecting the terms in the general estimate Lemma 30.

As we are working with the cotangent bundle $W = T^*M$, we use the canonical Liouville form $\lambda = pdq$; substituting this, and using the boundary conditions for u , the above integral becomes:

$$I = \int_0^1 \langle p(t), q'(t) \rangle - H_{x,0,t}(u(0, t)) dt,$$

where $q(t) = A(x, t)$ and $u(0, t) = (p(t), q(t))$ considered as section of T^*M .

Write $r(t) = r(u(0, t))$. Then:

$$\langle p(t), q'(t) \rangle \leq \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\}r(t),$$

as can be proved by rescaling in the fiber direction until $p \in \partial\Omega$, and then using the fact the above estimate is invariant under rescaling in the fiber direction. Thus:

$$I \leq \int_0^1 \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\}r(t) - H_{x,0,t}(u(0, t))dt.$$

Because Θ -data satisfies property (5), the function:

$$F(t, w) = \max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\}r(w) - H_{x,0,t}(w)$$

is non-positive for $r(w) \geq r_0$. Thus F attains a finite maximum $F_{\max}$ over the space of t, w, x , independently of u . Then $I \leq F_{\max}$ holds independently of the solution (x, u) , completing the proof. $\square$

4.1.3. Definition of Θ . Let (H_t, J) be admissible for defining $\text{CF}(H_t, J)$, let $A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M$ be a family of loops in $Z(\Lambda_c)$, and let $(\mathfrak{a}, J)$ be Θ -data. We assume throughout this subsection that the slope of H_t is at least c .

For a generic perturbation term $\mathfrak{p}$ on $P \times \Sigma \times W$, consider the moduli space $\mathcal{M}$ of solutions to §3.3.3, satisfying the boundary conditions (22). Let $\mathcal{M}_0$ be the rigid component of $\mathcal{M}$.

We define:

$$\Theta(A) := \sum_{\gamma} \#\{(x, u) \in \mathcal{M}_0 : u \text{ is asymptotic to } \gamma\}\gamma.$$

By the above a priori energy estimate, this sum is finite. The key structural result is:

Lemma 57. *For each $A \in Z(\Lambda_c)$, $\Theta(A)$ is a cycle in the Floer complex. Moreover, the homology class of $\Theta(A)$ in $\text{HF}(H_t, J)$ is independent of the choice of Θ -data, perturbation one-form $\mathfrak{p}$, and is independent of the bordism class of A . The resulting map:*

$$\Theta : H(\Lambda_c) \rightarrow \text{HF}(H_t, J)$$

is a map of $\mathbb{Z}/2\mathbb{Z}$ -vector spaces.

Proof. To see that $\Theta(A)$ is a cycle, one considers the ends of the 1-dimensional component $\mathcal{M}_1 \subset \mathcal{M}$; one shows that the number of ends of the component of $\mathcal{M}_1$ asymptotic to γ equals the coefficient in front of γ appearing in the composition $d(\Theta(A))$. Thus $d(\Theta(A)) = 0$ as we count modulo two.

To see that the chain homotopy class of $\Theta(A)$ is independent of the Θ -data, or perturbation one-form $\mathfrak{p}$, one appeals to the usual parametric moduli space argument; see, e.g., the discussion at the end of §3.4.2.

A bit care is needed to prove the chain homotopy class of $\Theta(A)$ is independent of the bordism class of A . Nonetheless, the argument is still based on a parametric moduli space. Given a cobordism $C : Q \times \mathbb{R}/\mathbb{Z} \rightarrow M$ between A and A' , one can consider Θ -data for C ; this entails a connection one-form $\mathfrak{a}$ and almost complex structure on $Q \times \Sigma \times W$ satisfying properties (1) through (5), the only difference being that Q is a manifold with boundary rather than

a closed manifold. Invoking another generic perturbation term $\mathbf{p}$, this data leads to a parametric moduli space $\mathcal{M}^{\text{param}}$ of solutions (y, u) where $y \in Q$. We note that Θ -data for C exists since we assume the cobordism happens within $Z(\Lambda_c)$, and the slope of H_t is at least c , by assumption.

For generic perturbation term, the 1-dimensional component $\mathcal{M}_1^{\text{param}}$ has two boundary components:

$$\partial\mathcal{M}_1^{\text{param}} = \{(y, u) : y \in P \text{ or } y \in P'\},$$

where $P \sqcup P' = \partial Q$. These solutions on the boundary are rigid when y is restricted to variations tangent to ∂Q . A moment's reflection therefore reveals that:

$$\Theta(A) - \Theta(A') = \sum \#\{(y, u) \in \partial\mathcal{M}_1^{\text{param}} : u \text{ is asymptotic to } \gamma\}\gamma.$$

If $\mathcal{M}_1^{\text{param}}$ were compact, then we would be done, as the number of boundary points of a compact manifold with boundary is even. However, $\mathcal{M}_1^{\text{param}}$ can have non-compact components. Nonetheless, the non-compact ends of $\mathcal{M}_1^{\text{param}}$ can be understood via Floer breaking/gluing; the usual theory shows that:

$$d\left(\sum \#\{(y, u) \in \mathcal{M}_0^{\text{param}} : u \text{ is asymptotic to } \gamma\}\gamma\right) = dK = \sum N_{\gamma'}\gamma'$$

where $N_{\gamma'}$ is the number of non-compact ends of the component of $\mathcal{M}_1^{\text{param}}$ consisting of those (y, u) where u is asymptotic to γ' . Thus one proves:

$$\Theta(A) - \Theta(A') = dK,$$

as desired.

Finally, to see that Θ is a vector space map, it suffices to prove that Θ respects addition. Since addition in $H(\Lambda_c)$ is given by disjoint union, and we can pick the Θ -data independently for each component of the parameter space P , it follows easily that $\Theta(A + A') = \Theta(A) + \Theta(A')$. $\square$

4.1.4. Compatibility with continuation maps. To prove that Θ induces a map of persistence modules $H(\Lambda_c) \rightarrow V_c$, it is necessary to show that Θ is compatible with continuation maps between the Floer cohomologies.

Lemma 58. *If $\mathbf{c} : \text{HF}(H_{0,t}, J_0) \rightarrow \text{HF}(H_{1,t}, J_1)$ is a continuation map, as in §3.4.2, and the slope of $(H_{0,t}, J_0)$ at least c , then:*

$$\mathbf{c} \circ \Theta_0 = \Theta_1$$

as maps $H(\Lambda_c) \rightarrow \text{HF}(H_{1,t}, J_1)$, where $\Theta_i : H(\Lambda_c) \rightarrow \text{HF}(H_{i,t}, J_i)$ is the Θ -map constructed in §4.1.3.

Proof. This follows easily from the construction of $\mathbf{c}$ and Θ_i , and is left to the reader (see also the arguments in [AS10, Abo15]). $\square$

4.2. BV-operators. This subsection is concerned with showing that Θ is compatible with the BV-operator on $H(\Lambda_c)$ constructed in §2.1.2 and the BV-operator on $\text{HF}(H_t, J)$ constructed in §3.4.5. This is part of the content of Theorem 11.

Lemma 59. *If (H_t, J) is admissible for defining the Floer complex, and has slope at least c , then:*

$$\Delta \circ \Theta = \Theta \circ \Delta,$$

as maps $H(\Lambda_c) \rightarrow \text{HF}(H_t, J)$.

Proof. The argument has no surprises, and follows similar lines to the arguments in [AS10, Abo15]. We leave the details to the reader. $\square$

4.3. Inclusion of the constant loops and PSS. In this subsection, we argue that Θ intertwines the inclusion of constant loops map $\mathfrak{i} : H^*(W) \rightarrow H(\Lambda_c)$ with PSS : $H^*(W) \rightarrow \text{HF}(H_t, J)$. We assume that H_t, J has slope at least $c > 0$.

Lemma 60. *Suppose that $\beta \in H^*(W)$. Then $\text{PSS}(\beta) = \Theta(\mathfrak{i}(\beta))$ as homology classes in $\text{HF}(H_t, J)$, where $\mathfrak{i}(\beta)$ is considered as an element of $H(\Lambda_c)$.*

Proof. The first step in the proof is to appeal to the proof of Lemma 16; there it is shown that β is represented by the map $C : S \rightarrow T^*M$ in the fiber product diagram:

$$\begin{array}{ccc} S & \xrightarrow{C} & T^*M \\ \downarrow & & \downarrow \\ P & \xrightarrow{f} & M, \end{array}$$

where $f : P \rightarrow M$ is a smooth map and P is a compact manifold. One can think of C as the collection of cotangent fibers living over the map f .

As in Lemma 16, it holds that $\mathfrak{i}(\beta)$ is represented by:

$$A : P \times \mathbb{R}/\mathbb{Z} \rightarrow M \text{ given by } A(x, t) = f(x),$$

i.e., $\mathfrak{i}(\beta)$ is represented by a cycle of constant loops.

Let $(\mathfrak{a}^{\text{PSS}}, J)$ be PSS-data, so it is defined on $\mathbb{R} \times \mathbb{R}/\mathbb{Z} \times W$. Let us suppose that $\mathfrak{a}_{s,t}^{\text{PSS}} = 0$ for $s \geq s_0$. For each $R > s_0$, define:

$$\mathfrak{a}_{R,x,s,t} = \mathfrak{a}_{s+R,t}^{\text{PSS}};$$

note that $\mathfrak{a}_{R,x,s,t}$ is Θ -data for $A(x, t) = f(x)$, since $\mathfrak{a}_{R,x,0,t} = 0$, for each R .

One can think of $\mathfrak{a}_{R,x,s,t}$ as a connection one-form on the family:

$$(s_0, \infty) \times P \times (-\infty, 0] \times \mathbb{R}/\mathbb{Z},$$

and one can introduce a compatible perturbation one-form $\mathfrak{p}_{R,x,s,t}$ supported in the region where $s \leq -R$, which is sufficient to ensure transversality.

This leads to the parametric moduli space $\mathcal{M}$ of solutions (R, x, u) , satisfying the boundary conditions:

$$u(0, t) \in T^*M_{f(x)},$$

which admits a smooth map $(R, x, u) \mapsto R \in (s_0, \infty)$. Let us denote the fiber over $R_0 \in (s_0, \infty)$ by $\mathcal{M}(R_0)$. Since $\mathfrak{a}_{R_0,x,s,t}$ is Θ -data for each R_0 , the count of elements rigid elements of $\mathcal{M}(R_0)$ represents $\Theta(\mathfrak{i}(\beta))$.

The rigid elements of $\mathcal{M}(R_0)$ live in the one-dimensional component of $\mathcal{M}$. Let us note that $\mathcal{M}$ has three kinds of non-compact ends:

- (1) ends containing sequences (R_n, x_n, u_n) where $R_n \rightarrow s_0$; we will ignore these ends,
- (2) ends containing sequences (R_n, x_n, u_n) where R_n converges in (s_0, ∞) ; as is well-understood, these ends converge to configurations contributing to dK where K counts the rigid elements of $\mathcal{M}$; we will also ignore these ends,
- (3) ends containing sequences (R_n, x_n, u_n) where R_n converges to ∞ .

A bit of thought shows that, the count of ends of type (3), where the left asymptotic of u_n is γ , defines a coefficient N_γ , and $\sum N_\gamma \gamma$ represents $\Theta(\mathbf{i}(\beta))$; the argument is the same as the one given in [Can24, §5.3] — briefly, the count of rigid elements $\mathcal{M}(R_0)$, for R_0 large enough, equals the count of ends of type (3).

Thus it remains to show the count of ends of type (3) represents $\text{PSS}(\beta)$. To analyze this, consider the change of coordinates $w(s, t) = u(s - R, t)$, so that w is defined on $(-\infty, R] \times \mathbb{R}/\mathbb{Z}$. Then, along any sequence (R_n, x_n, u_n) , the shifted map (x_n, w_n) has a subsequence which converges to a solution (x, w) of the PSS-equation, where $\lim_{s \rightarrow \infty} w(s, t) = C(x)$. For this step to work properly, we should assume that $\mathbf{p}_{R, x, s-R, t}$ converges as $R \rightarrow \infty$ to a limiting perturbation term $\mathbf{p}_{x, s, t}$ compatible with the PSS equation.

On the other hand, for any solution (x, w) of the PSS equation, a standard gluing argument for solutions to Floer's equation with Lagrangian boundary conditions shows that each (x, w) arises as the limit of an end of type (3) in the above sense. Thus we conclude that $\text{PSS}(\beta) = \Theta(\mathbf{i}(\beta))$, as desired. $\square$

4.4. Product structures. In this final subsection, we explain why Θ is compatible with the product structures. Such a result was originally proved in [AS10], and is a cornerstone in the relationship between string topology and Floer cohomology.

Remark. Let us comment that there is another way in the literature which relates string topology with Floer cohomology, where one defines a map from $\text{HF}(H_t, J)$ to a suitable Morse homology of the free loop space. In this context, [Abo15, pp. 398] proves the product structures are identified, and the argument is simpler than the argument in [AS10] (and the argument we will explain in this section). That this direction of morphism provides simpler argument for the identification of ring structures was observed also in [AS12, pp. 500]. We should note that [CHO23, §5.4] show that the map from Floer cohomology to Morse homology considered by [Abo15] does respect filtrations (the argument relies on special choice of (H_t, J) and a novel action versus length estimate); moreover this argument can be applied to a suitable construction of the pair-of-pants product to show the product structures are preserved when the map is restricted to each filtration level.

Unfortunately, this direction of the morphism (going from Floer cohomology to string topology) does not seem to work well for our purposes. It does seem likely that, if we were to restrict to domains Ω which appear as the unit disk bundle associated to a Riemannian metric, then we could argue instead using Morse homology and using existing results in the literature, and reverse the direction of the Θ morphism, and obtain a proof (albeit a slightly convoluted

one) that the product structures are identified.⁵ We prefer to stick with the present direction of the Θ morphism, and give a direct proof.

4.4.1. *Set-up.* Throughout, we fix classes $\alpha_i \in H(\Lambda_{c_i})$, $i = 0, 1$.

First, pick two Hamiltonian system $H_{i,t}$, $i = 0, 1$, defined on $[0, 1]$, and an almost complex structure J so that:

- (1) $H_{i,t} = c_i(t)r$ holds outside of $r \geq r_0$,
- (2) $H_{i,t} = 0$ unless $|t - 1/2| < 1/4$,
- (3) $(H_{i,t}, J)$ are admissible for defining CF, when $H_{i,t}$ is extended to $\mathbb{R}/\mathbb{Z}$ by 1-periodicity,
- (4) the slope of $H_{i,t}$ is at least c_i .

Define:

$$H_{\infty,t} := \begin{cases} 2H_{0,2t} & \text{for } t \in [0, 1/2], \\ 2H_{1,2t-1} & \text{for } t \in [1/2, 1], \end{cases}$$

and extend $H_{\infty,t}$ to $\mathbb{R}/\mathbb{Z}$ by 1-periodicity. Finally, suppose that:

- (5) $(H_{\infty,t}, J)$ is admissible for defining CF.

Since the slope of $H_{\infty,t}$ is the sum of the slopes of $H_{0,t}$ and $H_{1,t}$, there is a pair-of-pants map:

$$* : \text{HF}(H_{0,t}, J) \otimes \text{HF}(H_{1,t}, J) \rightarrow \text{HF}(H_{\infty,t}, J).$$

Second, pick representatives $A_i : P \times \mathbb{R}/\mathbb{Z} \rightarrow M$ for α_i so that:

- (1) $x \in P_i \mapsto A_i(x, 0)$, $i = 0, 1$, are mutually transverse,
- (2) $\max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\} \leq c_i(t)$ for each $q(t) = A_i(x, t)$.

the second condition can be achieved by a x -dependent smooth family of time-reparametrizations, since the integral of $c_i(t)$ is strictly larger than $\ell_{\Omega}(q(t))$. As a consequence, it follows that $A_i(x, 0) = A_i(x, t)$ unless $|t - 1/2| < 1/4$.

Define P_{∞} to be the fiber product of the transverse maps in (1), and define:

$$A_{\infty} : P_{\infty} \times \mathbb{R}/\mathbb{Z} \rightarrow M \text{ given by } A_{\infty}((x_0, x_1), t) = \begin{cases} A_0(x_0, 2t) & \text{for } t \in [0, 1/2], \\ A_1(x_1, 2t) & \text{for } t \in [1/2, 1], \end{cases}$$

which we extend from $[0, 1]$ to all of $\mathbb{R}$ by 1-periodicity. Then $A_{\infty} \in Z(\Lambda_{c_0+c_1})$, and $[A_{\infty}] = \alpha_0 * \alpha_1$.

The goal in this section is to prove:

$$(23) \quad \Theta(A_{\infty}) = \Theta(A_0) * \Theta(A_1),$$

as cycles in $\text{HF}(H_{\infty,t}, J)$.

⁵Note that the length function ℓ_{Ω} is the same as the length function associated to the fiberwise convex hull of Ω . In particular, there is the potential of a Morse theoretical argument, using the energy functional for (irreversible) Finsler metrics.

4.4.2. *An auxiliary map.* Before explaining why (23) holds, we will define an auxiliary cycle in $\text{CF}(H_{\infty,t}, J)$ which we will show represents both $\Theta(A_{\infty})$ and $\Theta(A_0) * \Theta(A_1)$.

Define a connection one-form $\mathfrak{a}$ on $\mathbb{C} \setminus \{0\}$ by the formula:

$$\mathfrak{a}_{s,t} = H_{\infty,t} dt$$

in cylindrical coordinates $z = e^{-2\pi(s+it)}$. Note that the connection one-form $\mathfrak{a}$ vanishes outside of the segments contained between rays $\mathbb{R}_+ i e^{\pm\pi i/4}$ (where its values are determined by $H_{1,t}$) and $\mathbb{R}_- i e^{\pm\pi i/4}$ (where its values are determined by $H_{0,t}$), as shown in Figure 10.




FIGURE 10. Illustration of the connection one-form $\mathfrak{a}$ on the domain $\mathbb{C} \setminus \{0\}$.

Let us say that two smoothly embedded closed disks D_0, D_1 are *shadowing* provided:

- (1) $D_0 \subset \{z = x + iy : y < 0\}$ and $D_1 \subset \{z = x + iy : y > 0\}$,
- (2) in coordinates $z = e^{-2\pi(s+it)}$, each ray with fixed t coordinate in $[1/8, 3/8]$ intersects ∂D_0 in points $s_-^0(t) < s_+^0(t)$, where $s_{\pm}$ are smooth functions on $[1/8, 3/8]$, and the radial vector ∂_s is transverse to ∂D_0 when $t \in [1/8, 3/8]$
- (3) in coordinates $z = e^{-2\pi(s+it)}$, each ray with fixed t coordinate in $[5/8, 7/8]$ intersects ∂D_1 in points $s_-^1(t) < s_+^1(t)$, where $s_{\pm}$ are smooth functions on $[5/8, 7/8]$, and the radial vector ∂_s is transverse to ∂D_1 when $t \in [5/8, 7/8]$

A *homotopy* of shadowing disks is an isotopy of embedded disks $D_{0,\tau}, D_{1,\tau}$ which satisfy the above properties for all τ .

The conditions admittedly look a bit strange, but it is perhaps best understood by comparing Figure 11 with Figure 10.




FIGURE 11. Shadowing disks block the rays from hitting the origin.

Given shadowing disks $D = (D_0, D_1)$, define a connection one-form $\mathfrak{a}^D$ on $\Sigma^D := \mathbb{C} \setminus \text{Interior}(D_0 \cup D_1)$

by the formula:

$$\mathfrak{a}_z^D = \begin{cases} 0 & \text{if the ray joining } 0 \text{ and } z \text{ misses the interior of } D_0 \cup D_1, \\ \mathfrak{a}_z & \text{otherwise;} \end{cases}$$

referring to Figure 11, the connection one-form $\mathfrak{a}^D$ is supported in the shaded region of Σ^D . It is straightforward to show that $\mathfrak{a}^D$ is smooth on Σ^D .

Shadowing disks $D = (D_0, D_1)$ also determine moving Lagrangian boundary conditions for $\partial\Sigma^D$. Referring to the notation $s_{\pm}^i(t)$ in (2) and (3), we define Lagrangians $L_{x_0, x_1, z}$ on $P_0 \times P_1 \times \partial\Sigma^D$ by the rule:

$$L_{x_0, x_1, z}^D := \begin{cases} T^*M_{A_0(x_0, 2t)} & \text{if } z = e^{-2\pi(s_-^0(t) + it)} \text{ for } t \in [1/8, 3/8], \\ T^*M_{A_1(x_1, 2t)} & \text{if } z = e^{-2\pi(s_-^1(t) + it)} \text{ for } t \in [5/8, 7/8], \\ T^*M_{A_0(x_0, 0)} & \text{if } z \in \partial D_0 \text{ and is not captured by previous rule,} \\ T^*M_{A_1(x_1, 0)} & \text{if } z \in \partial D_1 \text{ and is not captured by previous rule.} \end{cases}$$

In words, the Lagrangian boundary conditions are determined by the two loops $A_i(x_i, -)$, $i = 0, 1$, which are traversed along the “outer segment” determined by $s = s_-^i(t)$, and are parametrized using the angular coordinate associated to $z = e^{-2\pi(s + it)}$. Outside of the outer segments, the Lagrangian boundary conditions remain fixed at $T^*M_{A_i(x_i, 0)}$, the fiber over the basepoint of the loop.

This set-up leads to data $(\mathfrak{a}^D, L^D)$ on the family $P_0 \times P_1 \times \Sigma^D \times W$. For an almost complex structure J on the same family, so that $J_{x_0, x_1, z} = J$ is fixed outside of a sufficiently large disk, and a perturbation term $\mathfrak{p}$ on the same family, there is an associated moduli space $\mathcal{M}(\mathfrak{a}^D, \mathfrak{p}, J, L^D)$ of finite energy solutions (x_0, x_1, u) satisfying boundary conditions:

$$u(z) \in L_{x_0, x_1, z} \text{ for } z \in \partial\Sigma^D,$$

and which solve the general form of Floer’s equation described in §3.3.3.

Lemma 61. *There is an a priori energy bound on solutions:*

$$(x_0, x_1, u) \in \mathcal{M}(\mathfrak{a}^D, \mathfrak{p}, J, L^D)$$

which is independent of D and $\mathfrak{p}$, provided $\mathfrak{p}$ is sufficiently small.

Proof. The argument is similar to that used for Lemma 56, in that it suffices to bound the integrals:

$$I = \int_{\partial\Sigma^D} u^* \lambda - v^* \mathfrak{a}^D$$

independently of the solution; here $v(z) = (z, u(z))$ is as in Lemma 30. The integrand vanishes on points $z \in \partial D$ which do not lie in the outer segments $s = s_-^i(t)$ described above; thus it suffices to bound the integral over the outer segments. We parametrize the outer segments by the angular coordinate t in

$z_0(t) = e^{-2\pi(s_0^-(t)+it)}$ and $z_1(t) = e^{-2\pi(s_+^1(t)+it)}$. The integral decomposes as $I = I_0 + I_1$ where:

$$I_0 = \int_{1/8}^{3/8} \langle p(t), q'(t) \rangle - 2H_{0,2t}(u(z(t)))dt, \text{ where } q(t) = A_0(x_0, 2t).$$

and similarly for I_1 , with $[1/8, 3/8]$ replaced by $[5/8, 7/8]$. Reparametrizing the integral by $\tau = 2t$ for $i = 0$ and $\tau = 2t - 1$ for $i = 1$, so $\tau \in [1/4, 3/4]$, and using the same estimate as in Lemma 56 together with property (2) in the choice of $A_i(x_i, t)$, one concludes that each I_i is uniformly bounded independently of the solution. This completes the proof. $\square$

Counting the rigid finite energy solutions (x_0, x_1, u) in $\mathcal{M}(\mathfrak{a}^D, \mathfrak{p}, J, L^D)$ for generic $\mathfrak{p}$, where the asymptotic orbit of u equals γ (at the negative end, as seen in the coordinates $z = e^{-2\pi(s+it)}$) gives a coefficient N_γ . Define:

$$\Pi(D) := \sum N_\gamma \gamma \in \text{CF}(H_{\infty,t}, J).$$

Then:

Lemma 62. *The chain $\Pi(D)$ is a cycle for each choice $\mathfrak{p}, D, J$, provided $\mathfrak{p}$ is generic. The homology class of $\Pi(D)$ is independent of $\mathfrak{p}, J$, and is also independent of the homotopy class of D in the space of shadowing disks.*

Proof. The argument is similar to many other arguments in this paper, in particular, Lemma 57. We omit the proof. $\square$

To eliminate the apparent dependency on the homotopy class of shadowing disks, we select a distinguished homotopy class. Let us fix $D^{\text{std}} = (D_0^{\text{std}}, D_1^{\text{std}})$ to be *standard circular shadowing disks*, defined to be:

$$D_0^{\text{std}} = -i + D(r) \text{ and } D_1^{\text{std}} = i + D(r)$$

where $D(r)$ is the disk of radius $r \in (2^{-1/2}, 1)$; see Figure 12. Elementary geometry proves these are indeed shadowing disks, and the homotopy class is independent of the choice of r .




FIGURE 12. Standard circular shadowing disks

The two results whose proofs occupy the rest of the paper are:

Lemma 63. *If D^{std} comprises standard circular shadowing disks, then:*

$$\Pi(D^{\text{std}}) = \Theta(A_0) * \Theta(A_1),$$

as elements of $\text{HF}(H_{\infty,t}, J)$.

Lemma 64. *If D^{std} comprises standard circular shadowing disks, then:*

$$\Pi(D^{\text{std}}) = \Theta(A_\infty),$$

as elements of $\text{HF}(H_{\infty,t}, J)$.

Combining these, we conclude $\Theta(A_\infty) = \Theta(A_0) * \Theta(A_1)$, completing the proof that Θ is compatible with the product structures (since A_∞ represents the Chas-Sullivan product $\alpha_0 * \alpha_1$). Together with the results in §4.2 and §4.3, this completes the proof of Theorem 11.

Relatively speaking, Lemma 63 is easy and Lemma 64 is hard. We prove Lemma 63 in the next subsection §4.4.3, and the proof of Lemma 64 occupies the remaining parts of the paper.

4.4.3. Proof of Lemma 63. The argument is fairly standard, and we only explain the main construction, omitting the Floer theory details.

For $r \in (2^{-1/2}, 1)$, let $D_0(r) = -i + D(r)$ and $D_1(r) = i + D(r)$, and let $D(r) = (D_0(r), D_1(r))$; these form standard circular shadowing disks.

Parametrize the punctured disk $D_i(r)^\times$ as a positive cylindrical end using cylindrical coordinates $s + it$. Then, $\mathfrak{a}^{D(r)}$ extends to a collar $s \in [0, \epsilon]$ for some small ϵ , using the same formula (as slightly shrunk disks remain shadowing disks). Let us abbreviate:

$$\mathfrak{a}_i = \mathfrak{a}^{D_i(r)}|_{s \in [0, \epsilon]},$$

which takes the form:

$$\mathfrak{a}_i = K_{i,s,t}ds + G_{i,s,t}dt,$$

and a moment's thought reveals that $K_{i,s,t} = b_i(s, t)r$ and $G_{i,s,t} = c_i(s, t)r$ for $r \geq r_0$. Moreover, $b_i(s, t)ds + c_i(s, t)dt$ is a closed one-form, and:

$$\int_0^1 c_i(s, t)dt = \text{slope of } H_{i,t}.$$

Moreover, $\mathfrak{a}_i$ has zero curvature, since $\mathfrak{a}^{D_i(r)}$ has zero curvature. The goal is now to extend $\mathfrak{a}_i$ to all of $[0, \infty) \times \mathbb{R}/\mathbb{Z}$. Define:

$$\begin{aligned} G'_{i,s,t} &= \beta(1 - s/\epsilon)G_{i,s,t} + \beta(s/\epsilon)H_{i,t} \\ K'_{i,s,t} &= \int_0^t \partial_s G'_{i,s,\tau} d\tau - t \int_0^1 \partial_s G'_{i,s,\tau} d\tau. \end{aligned}$$

Note that $G'_{i,s,t} = G_{i,s,t}$ in a neighborhood of $s = 0$, by construction of the standard cut-off function β . Since $\mathfrak{a}_i$ has zero curvature, it holds that:

$$\partial_s G_{i,s,t} = \partial_t K_{i,s,t},$$

where we use the fact that $\mathfrak{a}_i(V_1), \mathfrak{a}_i(V_2)$ are Poisson-commuting for any two tangent vectors V_1, V_2 based at the same point (this holds since $\mathfrak{a}_i$ is obtained by restricting $\mathfrak{a}^{D(r)}$, which has this property). Thus it holds that $K'_{i,s,t} = K_{i,s,t}$ also holds in a neighborhood of $s = 0$. Thus:

$$\mathfrak{a}'_i = K'_{i,s,t}ds + G'_{i,s,t}dt$$

is a valid extension of $\mathfrak{a}_i$ from $[0, \epsilon] \times \mathbb{R}/\mathbb{Z}$ to $[0, \infty) \times \mathbb{R}/\mathbb{Z}$, provided we shrink ϵ . Moreover, as in §3.4.2, $\mathfrak{a}'_i$ has curvature bounded from above.

Denote by:

$$\mathfrak{a}' = \begin{cases} \mathfrak{a}'_0 & \text{in } -i + D(r)^\times, \\ \mathfrak{a}'_1 & \text{in } i + D(r)^\times, \\ \mathfrak{a}^{D(r)} & \text{elsewhere,} \end{cases}$$

which is a smooth connection one-form on $\mathbb{C} \setminus \{i, -i\}$. By construction, $\mathfrak{a}'$ agrees with $H_{\infty,t}dt$ on the cylindrical end around $z = \infty$, (since it agrees with $\mathfrak{a}^{D(r)}$), and agrees with $H_{0,t}dt$, $H_{1,t}dt$ in the cylindrical ends around the punctures $-i, i$. A bit of thought (and standard arguments) show that appropriately counting the rigid finite-energy solutions of the equation determined by $(\mathfrak{a}', J, \mathfrak{p})$, for some perturbation term $\mathfrak{p}$, defines a chain-level representation of the pair-of-pants operation:

$$* : \text{CF}(H_{0,t}, J) \otimes \text{CF}(H_{1,t}, J) \rightarrow \text{CF}(H_{\infty,t}, J).$$

Next we explain how to deform the equation defining $\Pi(D)$ to the equation defining $* \circ (\Theta(A_0), \Theta(A_1))$. The process is illustrated in Figure 13.




FIGURE 13. 1-parameter family used to prove Lemma 63; as the parameter R increases, the boundary components of the surface contract onto the points $\pm i$. Solutions in the associated parametric moduli space converge to configurations representing $\Theta(A_1) * \Theta(A_2)$

For each $R \in [0, \infty)$, define $\mathfrak{a}^R$ to be the restriction of $\mathfrak{a}'$ to the region obtained by removing from the disks $D_0(r)$ and $D_1(r)$ the cylindrical ends defined by $s > R$. Then $\mathfrak{a}^0$ agrees with $\mathfrak{a}^{D(r)}$, and $\mathfrak{a}^R$ “converges” to $\mathfrak{a}'$ as $R \rightarrow \infty$. Let us denote by $\Sigma(R)$ the surface with boundary obtained from this removal, so $\Sigma(0) = \Sigma^D$.

For each R and $x_i \in P_i$, there are moving Lagrangian boundary conditions for $(\mathfrak{a}^R, \Sigma(R))$ defined as follows: pick a reparametrization $\rho_{R,x_i} : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ so that:

$$\max\{\langle p, q'(t) \rangle : p \in \Omega \cap T^*M_{q(t)}\} < c_i(R, t) \text{ where } q(t) = A_i(x_i, \rho_{R,x_i}(t)),$$

which induces boundary conditions by requiring that $u(R, t) \in T^*M_{q(t)}$ for each solution (R, x_0, x_1, u) . The functions ρ_{R,x_i} should be chosen so that:

- (1) the boundary conditions for $R = 0$ agree with the previously defined boundary conditions for $\mathfrak{a}^D, \Sigma^D$,
- (2) $\rho_{R,x_i}(t) = t$ for R sufficiently large.

Because the set of reparametrizations achieving the above estimate is contractible (the condition is a convex condition $\partial_t \rho_{R,x_i}(t)$, since the slope of $H_{i,t}$ is large enough), it is possible to pick such ρ_{R,x_i} .

The data: $\mathfrak{a}^R, \Sigma(R), J$, the above (R, x_0, x_1) -dependent moving boundary conditions, and a generic perturbation term $\mathfrak{p}$ on the family, leads to a moduli space $\mathcal{M}$ of solutions (R, x_0, x_1, u) . The 1-dimensional component of $\mathcal{M}$ maps to $[0, \infty)$ via $(R, x_0, x_1, u) \rightarrow R$, and generic fibers $\mathcal{M}(R_0)$ defined by $R = R_0$ are zero-dimensional manifolds. By standard arguments, similar to those used in §2.1.3, the count of:

- (1) elements in $\mathcal{M}(0)$ and,
- (2) non-compact ends of $\mathcal{M}$ containing sequences $(R_n, x_{0,n}, x_{1,n}, u_n)$ with $R_n \rightarrow \infty$,

define homologous cycles in $\text{CF}(H_{\infty,t}, J)$. The count of (1) is exactly the count defining $\Pi(D)$, while the count of (2) represents $\Theta(A_1) * \Theta(A_2)$, as can be shown using standard Floer theory gluing arguments. We note that in this last part of the argument, one should pick the perturbation term $\mathfrak{p}$ so that it converges in an appropriate sense as $R \rightarrow \infty$. $\square$

4.4.4. Geometric set-up for Lemma 64. The main idea is to deform the shadowing disks to be rectangles with rounded corners, with a very thin neck between the disks, as shown on the left of Figure 14. Roughly speaking, as the neck gets thinner, the solutions approximate those used to define $\Theta(A_\infty)$ (in fact, the neck will degenerate to a flow line connecting two marked points on the boundary of the limiting curve). The presence of this flow line indicates that the gluing/compactness theory we will employ is similar to the adiabatic gluing/compactness theory of [FO97, Ekh07, OZ11, EENS13, CC23].

We now describe the surface and various regions inside the domain in more detail. The surfaces obtained by removing the approximately rectangular shadowing disks from $\mathbb{C}$ have a long neck region. As the space between the shadowing disks shrinks, the modulus of the neck grows. This neck region is shown in Figure 14.

Define the domain $\Sigma(R)$, for $R \geq R_0$, by replacing the inner neck of modulus $2R_0$ with a neck of modulus $2R$. Note that $\Sigma(R)$ is defined abstractly and is not defined as a subset of $\mathbb{C}$. One can still speak about the various regions inside of $\Sigma(R)$; i.e., it still makes sense to speak about $\mathcal{R}_1, \mathcal{R}_2$, and the necks.

In the limit $R \rightarrow \infty$, the surface $\Sigma(R)$ converges (in the moduli space of domains) to a surface $\Sigma(\infty)$ which is conformally equivalent to a disk with one interior puncture and two boundary punctures. The local model around the boundary punctures is illustrated in Figure 15. The limit surface $\Sigma(\infty)$ has two ends (rather than one neck).

There is an obvious connection one-form $\mathfrak{a}$ on $\Sigma(R_0)$, described by the shadowing disk construction. It has the property that $\mathfrak{a}$ vanishes on the neck region. When we define $\Sigma(R)$ by removing the inner neck and gluing in a longer neck, we obtain a connection one-form $\mathfrak{a}$ by requiring that it vanishes on the neck and agrees with the previously defined $\mathfrak{a}$ outside the neck.

The Lagrangian boundary conditions are defined for parameter R_0 by the shadowing disk construction. For $R \geq R_0$, one defines the Lagrangian boundary conditions to match those for R_0 away from the neck, and require solutions to map the boundary of the neck to cotangent fibers $T^*M_{A(x_0,0)}$ and $T^*M_{A(x_1,0)}$ (as in the boundary conditions for R_0).




FIGURE 14. (Left) Shadowing disks defining the domain $\Sigma(R_0)$; the region $\mathcal{R}_1$ is the region between the rays where the connection one-form $\mathfrak{a}_{R_0}$ is supported; the region $\mathcal{R}_2$ is the complement of a large disk centered on the origin. (Right) The neck region of $\Sigma(R_0)$ is the neck with modulus $2R_0$; the *total neck* is the neck with modulus $2\ell + 2R_0$.




FIGURE 15. The end region on the limiting disk around the left end; the picture around the right end is a reflected version.

Let us pick perturbation one-forms $\mathfrak{p}_R$ depending on R with the property that $\mathfrak{p}_R$ is supported in the region $\mathcal{R}_1 \cap \mathcal{R}_2^c$; in particular, the perturbation term vanishes near the neck, and the equation appears as Floer's equation for $(H_{\infty,t}, J)$ in the cylindrical end $\mathcal{R}_2$, using coordinates $z = e^{-2\pi(s+it)}$.

Once a complex structure is chosen (see §4.4.5), the construction then gives a parametric moduli space $\mathcal{M}$ of solutions (R, x_0, x_1, u) where (x_0, x_1, u) solves the equation for the domain $\Sigma(R)$. Let us denote by $\mathcal{M}$ the one-dimensional component, and by $\mathcal{M}_0$ the rigid component; we will not have occasion to consider higher dimensional components. We denote by $\mathcal{M}(R)$ the fiber of $\mathcal{M}$ of solutions (x_0, x_1, u) , so $\mathcal{M}(R)$ consists of rigid solutions for generic R .

As a special case, we define $\mathcal{M}(\infty)$ to be the moduli of rigid finite energy solutions on the limiting surface $\Sigma(\infty)$, where (x_0, x_1) lies in the fiber product P_∞ ; the finite energy condition implies any solution converges exponentially to a removable singularity lying on $T^*M_{A(x_0,0)} = T^*M_{A(x_1,0)}$ at the punctures.

We will require that $\mathfrak{p}_R$ is equal to $\mathfrak{p}_\infty$ for R sufficiently large, where $\mathfrak{p}_\infty$ is a perturbation term used to achieve transversality for $\mathcal{M}(\infty)$. As part of the

gluing analysis, we will show that such $\mathfrak{p}_R$ also achieve transversality for the parametric moduli space $\mathcal{M}$.

In §4.4.6 we will explain how to pick the parameter ℓ ; essentially it needs to be chosen large enough that the neck/ends are mapped into a small neighborhood of the cotangent fiber $T^*M_{A(x_0,0)}$.

At this stage, we claim the following:

Proposition 65. *The count of elements $(x_0, x_1, u) \in \mathcal{M}(\infty)$, weighted by the asymptotic orbit at the cylindrical end, defines a cycle in $\text{CF}(H_{\infty,t}, J)$ which represents $\Theta(A_\infty)$.*

Proof. The surface $\Sigma(\infty)$ is conformally equivalent to a closed disk $D(1)$ with two boundary punctures (say $-1, 1$) and one interior puncture 0 . The equation and boundary conditions are what is used to define Θ , with the sole exception that the cylindrical end at 0 is not the standard cylindrical end around 0 (the biholomorphism between $\Sigma(\infty)$ and $D(1)$ does not respect the cylindrical coordinates). However, the cycle one obtains is stable under small changes in the data, and so one can correct the cylindrical end very close to zero, without changing the resulting cycle. In this fashion, one proves the count obtained from $\mathcal{M}(\infty)$ is exactly the one used to define $\Theta(A_\infty)$. $\square$

Moreover, standard Floer theoretic arguments show:

Proposition 66. *The count of elements $(R_0, x_0, x_1, u) \in \mathcal{M}(R_0)$, weighted by the asymptotic orbit at the cylindrical end, and the count of non-compact ends of $\mathcal{M}$, also weighted by the asymptotic orbits, define homologous cycles in $\text{CF}(H_{\infty,t}, J)$. Both of these cycles represent $\Pi(D)$.*

Proof. The argument is standard and similar to other arguments in this paper. We only comment that the chain homotopy between the two cycles is given by counting the rigid elements in the parametric moduli space $\mathcal{M}_0$. $\square$

Thus the rest of the proof is dedicated to proving the cycle obtained by counting solutions in $\mathcal{M}(\infty)$ also represents the cycle $\Pi(D)$, i.e., represents the cycle obtained by counting the non-compact ends of $\mathcal{M}$. This step of the proof is what involves the delicate gluing argument. The argument is quite technical, and we have written it with an expert audience in mind; we assume the reader is already familiar with gluing theory in Floer theory, and we focus on the details needed to convince an expert in the validity of the approach.

4.4.5. Choice of almost complex structure. To simplify the argument as much as possible, we will make a particular choice of almost complex structure. The construction uses Riemannian geometry. For a given Riemannian metric g on TM (and hence on T^*M), and a smooth function $f : [0, \infty) \rightarrow (0, \infty)$, we let $J_{g,f}$ be the almost complex structure so that:

$$J_{g,f} = \begin{bmatrix} 0 & -f(\rho_g)g_* \\ (f(\rho_g)g_*)^{-1} & 0 \end{bmatrix} \text{ with respect to splitting } \text{Ver} \oplus \text{Hor}_g,$$

where $\rho_g(p) = g(p, p)$ is the fiberwise radius-squared function, and g_* is the duality isomorphism between TM and T^*M . The horizontal distribution Hor_g is defined using the Levi-Civita connection.

The key result about such almost complex structures is that:

Lemma 67. *Let $(\Sigma, \partial\Sigma)$ be a (not necessarily compact) Riemann surface with boundary. Let $f_z, z \in \Sigma$, be a family of functions $[0, \infty) \rightarrow (0, \infty)$, yielding a domain dependent family of almost complex structures J_{g, f_z} . For any J_{g, f_z} -holomorphic map $u : (\Sigma, \partial\Sigma) \rightarrow T^*M$, whose boundary is mapped onto cotangent fibers, the function $\rho_g(u)$ does not attain any local maxima.*

Proof. Write $u(s, t) = (p(s, t), q(s, t))$ where $q(s, t)$ is the projection to the base, and $p(s, t)$ is considered as section of u^*T^*M . The computations in [CC23, §2] imply that:

$$\nabla_s p = f_z(\rho_g)g_*\partial_t q \text{ and } \nabla_t p = -f_z(\rho_g)g_*\partial_s q.$$

A standard computation gives:

$$\Delta\rho_g = \Delta g(p, p) \geq 2g(\nabla_s \nabla_s p, p) + 2g(\nabla_t \nabla_t p, p),$$

and using the above holomorphic curve equation, and the well-known fact $\nabla_s(g_*\partial_t q) = \nabla_t(g_*\partial_s q)$ (see [CC23, Lemma 2.4]), we conclude:

$$\Delta\rho_g \geq (\partial_s \rho_g)f_z(\rho_g)^{-1}\partial_s(f_z(\rho_g)) - (\partial_t \rho_g)f_z(\rho_g)^{-1}\partial_t(f_z(\rho_g)).$$

If $\partial\Sigma = \emptyset$, the desired result then follows from the maximum principle [GT98, Chapter 3]. In general, observe that $d\rho_g$ vanishes on directions orthogonal to the boundary, and hence doubles to a C^2 function; the maximum principle can then be applied to this doubling. $\square$

Having established this, we construct a family $J_{x_0, z}$ of almost complex structures on the family $P_0 \times \mathbb{C}$, as follows. Pick a family of Riemannian metrics g_{x_0} so that g_{x_0} is flat in a neighborhood of $A_0(x_0, 0)$. Then pick the almost complex structures $J_{x_0, z}$ so that:

- (1) in $\mathcal{R}_2$, $J_{x_0, z} = J$ is fixed,
- (2) inside of a disk slightly smaller than the one bounding $\mathcal{R}_2$, it holds that $J_{x_0, z} = J_{g_{x_0}, f_z}$,
- (3) f_z is locally constant inside the total neck (where it equals f_0) and on a neighborhood of $\mathcal{R}_1$ (where it equals f_{z_1} for some $z_1 \in \mathcal{R}_1$).
- (4) $f_0(\rho) = 1$ holds for $\rho \leq r_1$,
- (5) $f_z(\rho) = \rho^{1/2}$ for $\rho \geq r_1 + 1$ holds at all points z .

Note that because $f_z(\rho) = \rho^{1/2}$ holds outside of a compact set, $J_{x_0, z}$ is Liouville equivariant outside of a compact set.

The construction is admittedly a bit ad hoc, but the benefits of the construction is that:

Lemma 68. *There is a constant C independent of R_0 and f_z for $z \notin \mathcal{R}_1$ (but depending on f_{z_1}) so that any finite energy solution (x_0, x_1, u) in the moduli space $\mathcal{M}(R)$, where $R \in [R_0, \infty]$, satisfies $\rho_{x_0}(u) \leq C$. In other words, fixing f_{z_1} , one can alter the choice of f_z for $z \notin \mathcal{R}_1$ so the above conclusion holds.*

Proof. We consider the two regions $\mathcal{R}_1, \mathcal{R}_2 \subset \Sigma(R)$ illustrated in Figure 14. By Lemma 67, the supremum of $\rho_{x_0}(u)$ is equal to the supremum of $\rho_{x_0}(u)$ on the region $\mathcal{R}_1 \cup \mathcal{R}_2$.

By the earlier maximum principle Proposition 28, the supremum of $\rho_{x_0}(u)$ over $\mathcal{R}_2$ depends only on $(H_{\infty,t}, J)$, and the apriori energy bound which is independent of the choice of family of almost complex structures and R .

Any point in $\mathcal{R}_1$ can be joined to $\mathcal{R}_2$ by a radial path. Standard bubbling analysis proves the derivatives along the path are bounded by a constant depending only on the Hamiltonian connection and f_{z_1} (the radial path remains inside $\mathcal{R}_1$). Therefore the supremum of $\rho_{x_0}(u)$ over $\mathcal{R}_1$ is also bounded independently of f_z for $z \notin \mathcal{R}_1$, and R_0 . This proves the statement. $\square$

Therefore we can pick r_1 larger than C , because C is independent of r_1 .

The upshot of this construction is the following:

Corollary 69. *Any solution (x_0, x_1, u) in $\mathcal{M}(R)$, $R \in [R_0, \infty]$, is J_0 -holomorphic on the intersection of the total neck and $u^{-1}(B(1))$ if $B(1)$ is a coordinate ball around $A(x_0, 0)$ where g_{x_0} is flat; here J_0 is the standard almost complex structure in canonical coordinates p, q .*

The standard almost complex structure is the one which satisfies $J_0 \partial_{p_i} = \partial_{q_i}$, $i = 1, \dots, n$. In other words, we can assume without loss of generality that $J_{x_0,z}$ is standard on canonical coordinates on the coordinate ball $B(1)$, and if z is in the total neck.

4.4.6. On the choice of ℓ . In this subsection, we prove that ℓ can be chosen large enough that certain properties hold for solutions in $\mathcal{M}(R)$ for $R \geq R_0$.

Lemma 70. *For any $\epsilon > 0$, the parameter ℓ can be chosen large enough that: for all solutions $(R, x_0, x_1, u) \in \mathcal{M}(R)$, u maps the neck (or ends if $R = \infty$) into the preimage of $B(\epsilon)$, where $B(\epsilon) \subset B(1)$ is contained in the coordinate ball around $A(x_0, 0)$. Recall the neck/ends exclude the pieces with modulus ℓ .*

Proof. This is a simple compactness argument, and is left to the reader. $\square$

The parameter ϵ will be chosen as follows. For each of the finitely many points (x_0, x_1) in $P_\infty \subset P_0 \times P_1$ so that there is some solution $(x_0, x_1, u) \in \mathcal{M}(\infty)$, consider the non-linear map:

$$(24) \quad (y_0, y_1) \in U \subset P_0 \times P_1 \mapsto A_1(y_1, 0) - A_0(y_0, 0)$$

defined on the open neighborhood of points (y_0, y_1) where $A_0(y_0, 0), A_1(y_1, 0)$ lie in the coordinate ball $B(1)$ (this difference vector is computed using the coordinate chart).

By the transversality of $A_0(-, 0), A_1(-, 0)$, this map is a submersion at (x_0, x_1) , and the kernel of the differential is $TP_{\infty, (x_0, x_1)}$. Pick a smooth open disk V passing through (x_0, x_1) so that $TV_{(x_0, x_1)}$ and $TP_{\infty, (x_0, x_1)}$ are complementary spaces. Then the restriction of (24) to V is a diffeomorphism between small neighborhoods of $(x_0, x_1) \in V$ and $0 \subset \mathbb{R}^n$. We assume that ϵ is small enough that $B(\epsilon)$ is contained in this small neighborhood; thus, for each vector q in $B(\epsilon)$, there is $(y_0, y_1) \in V$ so that:

$$(1) A_1(y_1, 0) - A_0(y_0, 0) = q.$$

Moreover, for appropriate metric on $P_0 \times P_1$ it follows that:

$$(2) \text{dist}((y_0, y_1), (x_0, x_1)) \leq |q|.$$

For the rest of argument, we fix ϵ, ℓ , and V , without further modifications so that the conclusion of the lemma holds.

4.4.7. Linearization framework for solutions of $\mathcal{M}(\infty)$. Let (x_0, x_1, u) be a solution in $\mathcal{M}(\infty)$. The goal in this section is to describe the linearization framework. As usual, the space of maps nearby u is modelled on an appropriate Banach space completion of sections of u^*TW . Since the domain of u is a punctured domain, we will use an appropriately weighted Sobolev space. For use as a weight, fix $\delta \in (0, \pi)$.

First we define $W^{1,p,\delta}$, $p > 2$, to be the Sobolev space of sections of u^*TW which are tangent to the Lagrangian boundary conditions with an exponential weight in the ends. On the end regions, there is a canonical trivialization of $u^*TW \simeq \mathbb{C}^n$ induced by the canonical coordinates associated to the coordinate ball $B(1)$ around $A(x_0, 0) = A(x_1, 0)$, and the linear deformations ξ are required to point in the real space $\mathbb{R}^n$.

There are two ends: the left end is identified with $[0, \infty) \times [0, 1]$ and the right end with $(-\infty, 0] \times [0, 1]$. In these ends, we require that:

$$e^{\pm \delta s} \xi(s, t) \text{ is in } W^{1,p},$$

where s -coordinate on the end is such that the line $s = 0$ corresponds to the innermost boundary of the strip of modulus ℓ .

Of course, if one only uses variations in $W^{1,p,\delta}$, then one could not change the position of the removable singularity of u at the two punctures. Thus it is necessary to stabilize to allow this point to vary. We therefore pick two linear maps $\Phi_{\pm}$ from $\mathbb{R}^n$ to the space of smooth variations of u so that:

- (1) $\Phi_{-}(v) = \beta(s)v$, $\Phi_{+}(v) = 0$ holds in the left end,
- (2) $\Phi_{-}(v) = 0$, $\Phi_{+}(v) = \beta(-s)v$ holds in the right end,

where β cuts off in the cut-off region shown in Figure 14.

This gives a family of variations:

$$(v_{-}, v_{+}, \xi) \in \mathbb{R}^n \oplus \mathbb{R}^n \oplus W^{1,p,\delta} \mapsto \Phi_{-}(v_{-}) + \Phi_{+}(v_{+}) + \xi,$$

which is sufficient to capture all nearby maps with the same boundary conditions. For further information on the use of stablized exponentially weighted Sobolev spaces in the context of maps with Lagrangian boundary conditions with boundary punctures, we refer the reader to [BC07].

Because $\mathcal{M}(\infty)$ is a parametric moduli space of triples (x_0, x_1, u) , where $(x_0, x_1) \in P_{\infty}$, there are also variations which move the point (x_0, x_1) . To account for this, we introduce a linear map Ψ taking $w = (w_0, w_1) \in TP_{\infty}$ into the space of smooth variations of u so that:

- (3) the projection of $\Psi(w)$ at $z \in \partial\Sigma(\infty)$ onto the zero section matches the variation of $A_i(x_i, t(z))$ in the direction of $w_i \in TP_{i,x_i}$; here the

- t coordinate is determined by the relation that $u(z) \in T^*M_{A_i(x_i, t(z))}$ and is described in greater detail in the shadowing disks construction;
- (4) $\Psi(w)$ is constant and imaginary in the ends (i.e., projects to a constant imaginary vector).

The total space of variations of (x_0, x_1, u) is then identified with:

$$\mathbb{R}^n \oplus \mathbb{R}^n \oplus TP_\infty \oplus W^{1,p,\delta}$$

and each tuple (v_-, v_+, w, ξ) produces a variation of u given by:

$$\Phi_-(v_-) + \Phi_+(v_+) + \Psi(w) + \xi.$$

By the usual procedure (for concreteness, we use a Riemannian exponential map), one can differentiate the differential equation in the directions of these these variations. This produces a linear differential operator $D_{x_0, x_1, u}$.

By using a Riemannian exponential map for a metric which agrees with the standard metric in the canonical coordinate system above the coordinate ball $B(1)$, it is arranged that:

$$(25) \quad D_{x_0, x_1, u}(\Phi_-(v_-) + \Phi_+(v_+) + \Psi(w) + \xi) = \bar{\partial}\xi \text{ on the ends,}$$

where $\bar{\partial} = \partial_s + i\partial_t$, using the aforementioned trivialization. Then it is clear that the left hand side of (25) is valued in the exponentially weighted L^p space.

To obtain this simple formula for the linearized operator on the ends, one argues that the nearby map associated to the variation (v_-, v_+, w, ξ) equals:

$$u + \Phi_-(v_-) + \Phi_+(v_+) + \Psi(w) + \xi,$$

on the ends (in the canonical coordinate system) provided v_-, v_+, w, ξ are not too big (this is possible since u lies above $B(\epsilon)$, and so there is lots of room before one leaves $B(1)$, the domain of the coordinates).

Because (x_0, x_1, u) is supposed to be rigid, the Fredholm index of (25) as a map $\mathbb{R}^n \oplus \mathbb{R}^n \oplus TP_\infty \oplus W^{1,p,\delta} \rightarrow L^{p,\delta}$ is zero. By the standard Sard-Smale argument, as in [MS12], the generic perturbation term $\mathbf{p}_\infty$ can be chosen so the operator is an isomorphism for all solutions (x_0, x_1, u) . In particular, it has a bounded inverse. This will be used in §4.4.9.

4.4.8. Pregluing. The strategy is now to:

- (1) For each rigid solution $(x_0, x_1, u) \in \mathcal{M}(\infty)$, construct preglued solutions which are supposed to approximate solutions in $\mathcal{M}(R)$ for R sufficiently large. Prove they approximately solve the equation, as measured with an appropriately weighted Sobolev space norm.
- (2) Prove that each preglued solution approximates exactly one of the non-compact ends of $\mathcal{M}$ containing sequences (R_n, x_0^n, x_1^n, u_n) with $R_n \rightarrow \infty$. This involves analyzing a linearized operator for each preglued solution (gluing).
- (3) Prove that genuine solutions in $\mathcal{M}(R)$ are close to the preglued solutions as $R \rightarrow \infty$. Together with (2), this establishes a bijection between the non-compact ends and solutions in $\mathcal{M}(\infty)$.

This is a standard strategy for Floer gluing, and the precise details are quite close to the “adiabatic” gluing results cited in §4.4.4. In this subsection, we are concerned with step (1) of the strategy.

First of all, each solution $(x_0, x_1, u) \in \mathcal{M}(\infty)$ determines two points p_-, p_+ in the cotangent fiber $T^*M_{A_0(x_0,0)} = T^*M_{A_1(x_1,0)} \simeq \mathbb{R}^n$ — the identification with $\mathbb{R}^n$ uses the canonical coordinates above the ball $B(1)$. The point p_- is the asymptotic at the left end, and p_+ is the asymptotic at the right end. These two points vary smoothly with the solution, and there are two constants $C_\pm$ so that:

- (1) on the left end, $|u(s, t) - p_-| \leq C_- e^{-\pi s}$, where $s \in [0, \infty)$,
- (2) on the right end, $|u(s, t) - p_+| \leq C_+ e^{\pi s}$, where $s \in (-\infty, 0]$.

This follows from the removable singularity theorem for J_0 -holomorphic maps.

The argument will require us splitting the neck of modulus $2R$ into specific regions; this is illustrated in Figure 16.




FIGURE 16. Decomposing $R = s + 1 + \tau$. The shaded regions are the *cut-off regions*.

For each (x_0, x_1, u) , $q_0 \in B(\epsilon/2)$, define:

$$\begin{cases} N_\tau : [-\tau - 1, \tau + 1] \times [0, 1] \rightarrow \mathbb{R}^n \times \mathbb{R}^n, \\ N_\tau(z) := iq_0 + 2\tau^{-1}((\tau - z)p_- + (\tau + z)p_+), \end{cases}$$

where $p_\pm$ are considered as real-vectors (in $\mathbb{R}^n \times \{0\}$). Let us observe that, the imaginary part of $N_\tau(z)$ equals $q_0 + 2\tau^{-1}(p_+ - p_-)t$, and so, provided τ is sufficiently large, the size of the imaginary part can be bounded by ϵ , and hence N_τ can be considered as a holomorphic map in $\mathbb{R}^n \times B(1) \subset T^*M$.

Let $y_0, y_1 \in V$ be such that:

- (3) $A_1(y_1, 0) - A_0(y_0, 0) = 2\tau^{-1}(p_+ - p_-)$, and
- (4) pick $q_0 = A_0(y_0, 0)$.

Then N_τ takes boundary conditions for $t = 0, 1$ on $T^*M_{A_0(y_0,0)}$ and $T^*M_{A_1(y_1,0)}$.

The idea is to glue this holomorphic neck to the existing solution (x_0, x_1, u) to form a preglued solution.

Since V is contained in a small neighborhood of (x_0, x_1) , we can use the Riemannian exponential map to define a variation $\Gamma(y_0, y_1)$ of u which pushes the boundary conditions from those for (x_0, x_1) to those for (y_0, y_1) , in a controlled way: we suppose that the C^1 size of $\Gamma(y_0, y_1)$ is controlled by the distance $\text{dist}((y_0, y_1), (x_0, x_1))$. Moreover, we suppose that Γ is supported away from the region $\mathcal{R}_2$. We denote this variation by $u + \Gamma(y_0, y_1)$, where

the “+” symbol is to be interpreted using the Riemannian exponential map. Since the size of $\Gamma(y_0, y_1)$ is controlled by $\mathfrak{r}^{-1}|p_+ - p_-|$, we can assume that $u + \Gamma(y_0, y_1)$ maps the neck into $\mathbb{R}^n \times B(2\epsilon)$ provided $\mathfrak{r}$ is large enough.

Finally, we define the preglued solution $\text{PG}_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u)$ as a piecewise function gluing together $u + \Gamma(y_0, y_1)$ and $N_{\mathfrak{r}}$, as illustrated in Figure 17.




FIGURE 17. Preglued solution; compare with Figure 14. On the shaded cut-off regions (which are mapped into $B(2\epsilon)$), one should interpolate between the solutions using standard cut-off functions.

In the left cut-off region, parametrized by $s, t \in [0, 1]^2$, the interpolation is given by:

$$\beta(s)N_{\mathfrak{r}}(z) + (1 - \beta(s))(u + \Gamma(y_0, y_1)),$$

and a similar (reflected) formula is used in the right cut-off region.

The decomposition of R into $\mathfrak{s} + 1 + \mathfrak{r}$ affects how this preglued solution is constructed, so different choices of $\mathfrak{s}, \mathfrak{r}$ yield different preglued solutions.

It will be important to measure sizes of variations of $\text{PG}_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u)$ using a weighted Sobolev norm. We introduce the weight:

$$\mathfrak{w} = \min\{e^{\delta(R+s)}, e^{\delta(-s+R)}\}$$

supported on the neck of length $2R$ parametrized so $s = 0$ is the middle of the neck (see Figure 14 for illustration of the region). Then we define $L^{p, \mathfrak{w}}$ as the set of L^p_{loc} sections η which are L^p integrable on the asymptotic cylindrical end, and so that $\mathfrak{w}\eta$ is L^p integrable on the neck; $W^{1, p, \mathfrak{w}}$ is defined analogously.

We then have:

Lemma 71. *If $F(y_0, y_1, w)$ is the non-linear map encoding the PDE for solutions of $\mathcal{M}(R)$, i.e., in local coordinates:*

$$F(w) = \partial_s w + J_{x_0, z}(w) \partial_t w - a(w),$$

then the $L^{p, \mathfrak{w}}$ size is bounded by:

$$\|F(y_0, y_1, \text{PG}_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u))\|_{L^{p, \mathfrak{w}}} \leq (C_+ + C_-)e^{-(\pi-\delta)\mathfrak{s}} + C_{\mathfrak{s}}\mathfrak{r}^{-1},$$

where $C_{\pm}$ are defined in (1) and (2), and $C_{\mathfrak{s}}$ is independent of $x_0, x_1, u, \delta, \mathfrak{r}$, but depends on $\mathfrak{s}$.

Proof. Since u solves the equation, $F(y_0, y_1, u + \Gamma(y_0, y_1))$ will approximately solve the equation up to an error whose C^0 size is controlled by

$$\mathfrak{r}^{-1}|p_+ - p_-|,$$

since that size governs the C^0 size of $\Gamma(y_0, y_1)$.

The integral of this error over the region disjoint from the neck is then bounded by $C_1 \mathfrak{r}^{-1}$. Moreover, the integral of this error over the necks of length $\mathfrak{s} + 1$ produces errors of size:

$$\mathfrak{r}^{-1} |p_+ - p_-| (\mathfrak{s} + 1)^{1/p} e^{\delta(\mathfrak{s}+1)} \leq C_2 \mathfrak{r}^{-1}.$$

In total, these errors can be bounded by $C_3 \mathfrak{r}^{-1}$.

There is an additional error arising due to the interpolation (we focus only on the left end as the right end follows from a reflected estimate):

$$\beta(s) N_{\mathfrak{r}}(z) + (1 - \beta(s))(u(z) + \Gamma(y_0, y_1)),$$

Since $N_{\mathfrak{r}}(z)$ and $u(z)$ both solve the equation (which is the standard $\bar{\partial}$ equation in the cut-off region), this produces an error of C^0 size:

$$|u(z) + \Gamma(y_0, y_1) - N_{\mathfrak{r}}(z)| + C_4 \mathfrak{r}^{-1},$$

where $C_4 \mathfrak{r}^{-1}$ is due to the fact that $\Gamma(y_0, y_1)$ has C^1 size bounded by $\mathfrak{r}^{-1}$. We use that $\beta'(s)$ is approximately 1 to avoid introducing another constant.

Then we estimate $|u(z) + \Gamma(y_0, y_1) - N_{\mathfrak{r}}(z)|$ in the region $s \in [-\mathfrak{r} - 1, -\mathfrak{r}]$ by:

$$|u(z) - p_-| + |p_- - N_{\mathfrak{r}}(z)| + |\Gamma(y_0, y_1)| \leq C_- e^{-\pi s} + C_5 \mathfrak{r}^{-1}.$$

The bound on the first and the last term is clear from the construction. Let us expand a bit why is $|p_- - N_{\mathfrak{r}}(z)|$ bounded in terms of $\mathfrak{r}^{-1}$. First note that $|iq_0|$ is bounded in terms of $\mathfrak{r}^{-1}$ since the distance between y_0 and x_0 is controlled by $\mathfrak{r}^{-1}$, and $A_0(\cdot, 0)$ is a smoot map, hence $q_0 - 0 = A_0(y_0, 0) - A_0(x_0, 0)$ is controlled by $\mathfrak{r}^{-1}$ as well. Secondly, since $s \in [-\mathfrak{r} - 1, -\mathfrak{r}]$, we have that $|\mathfrak{r} + z|$ is bounded by $\sqrt{2}$ and $|p_- - 2\mathfrak{r}^{-1}((\mathfrak{r} - z)p_- + (\mathfrak{r} + z)p_+)| = |2\mathfrak{r}^{-1}(\mathfrak{r} + z)(p_+ - p_-)|$.

The contribution to the $L^{p, \mathfrak{w}}$ size can be estimated by:

$$e^{+\delta s} C_- e^{-\pi s} + e^{+\delta s} C_5 \mathfrak{r}^{-1}.$$

Combining with the earlier estimate, we conclude:

$$\|F(y_0, y_1, \text{PG}_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u))\|_{L^{p, \mathfrak{w}}} \leq (C_- + C_+) e^{-(\pi - \delta)s} + C_5 \mathfrak{r}^{-1},$$

where $C_5 = (C_3 + 2e^{+\delta s}(C_4 + C_5))$, where we take into account the reflected estimate at the right end. This completes the proof. $\square$

In the rest of the argument, one should imagine that first $\mathfrak{s}$ is chosen large enough, and then, $\mathfrak{r}$ is chosen in terms of $\mathfrak{s}$.

4.4.9. Linearization framework for the preglued solutions. In this section we are concerned with step (2) of the gluing argument outlined above. The ideas are similar to those in §4.4.7; we will linearize the equation at the preglued solution producing a linearized operator. Then we will show this linearized operator is uniformly surjective as $R \rightarrow \infty$, provided the parameters $\mathfrak{s}, \mathfrak{r}$ are chosen sufficiently large.

Let us abbreviate $(y_0, y_1, u_{\mathfrak{s}, \mathfrak{r}}) = \text{PG}_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u)$. Using the Riemannian metric g_{x_0} as in §4.4.7, one can associate to each variation of $u_{\mathfrak{s}, \mathfrak{r}}$ a nearby map, and by the usual procedure of differentiating the equation, obtain a linearized operator $D_{y_0, y_1, u_{\mathfrak{s}, \mathfrak{r}}}$. As in §4.4.7, $D_{y_0, y_1, u_{\mathfrak{s}, \mathfrak{r}}}$ agrees with $\bar{\partial}$ on the neck region.

The space of variations of $(y_0, y_1, u_{s,\mathfrak{r}})$ we will work with is:

$$\mathbb{R}^n \oplus TP_0 \oplus TP_1 \oplus W^{1,p,\mathfrak{w}},$$

where:

- (1) TP_0, TP_1 account for variations w_0, w_1 based at y_0, y_1 ; these potentially move the boundary conditions,
- (2) $\mathbb{R}^n$ is a space of variations v which are real-valued and constant on the neck, similarly to the space of variations $\mathbb{R}^n \oplus \mathbb{R}^n$ used in §4.4.7,
- (3) $W^{1,p,\mathfrak{w}}$ are variations of u_R which are tangent to the boundary conditions and lie in the weighted Sobolev space.

It is convenient to split the short-exact sequence:

$$0 \rightarrow TP_\infty \rightarrow TP_0 \oplus TP_1 \rightarrow \mathbb{R}^n \rightarrow 0,$$

using the slice V , where the right map is the difference of the derivatives $y_i \mapsto A_i(y_i, 0)$; see §4.4.6 for the definition of V . With this splitting chosen, the space of variations is identified with:

$$\mathbb{R}^n \oplus TP_\infty \oplus \mathbb{R}^n \oplus W^{1,p,\mathfrak{w}}.$$

We will reuse the notation from §4.4.7 as much as possible. First, for each $v \in \mathbb{R}^n$, let $\Phi(v)$ be a function which is constant and equal to $v \in \mathbb{R}^n \times \{0\}$ in the neck region and “agrees” with $\Phi_-(v)$ and $\Phi_+(v)$ outside of the neck region. Here “agrees” should be interpreted up to the identification of variations of u with variations of u_R outside of the neck region (which appeals to a parallel transport map).

Second, for each $w \in TP_\infty$, let $\Psi(w)$ be a variation which is constant and imaginary in the neck region, and “agrees” with the old $\Psi(w)$ outside of the neck region.

Next, for $\mu_0, \mu_1 \in \mathbb{R}^n$, let $k(\mu_0, \mu_1) : [-\mathfrak{r} - 1, \mathfrak{r} + 1] \times [0, 1] \rightarrow \mathbb{C}^n$ be given by:

$$k(\mu_0, \mu_1) = 2\mathfrak{r}^{-1}i\mu_0 + 2\mathfrak{r}^{-1}((\mathfrak{r} - z)\mu_0 + (\mathfrak{r} + z)\mu_1).$$

This has boundary conditions on $\mathbb{R}^n \times 2\mathfrak{r}^{-1}\mu_0$ and $\mathbb{R}^n \times 2\mathfrak{r}^{-1}\mu_1$.

For $\mathfrak{r}$ large enough, $(2\mathfrak{r}^{-1}\mu_0, 2\mathfrak{r}^{-1}\mu_1)$ lies in the image of the derivatives of $(A_0(-, 0), A_1(-, 0))$ restricted to the slice V . Such pairs are uniquely determined by the difference vector $\mu = \mu_1 - \mu_0 \in \mathbb{R}^n$, and given such a pair we can find a variation $\gamma(\mu_0, \mu_1)$ by differentiating $\Gamma(y_0, y_1)$ in the direction of a vector tangent to V . Then $\gamma(\mu_0, \mu_1)$ has C^1 size controlled by $2\mathfrak{r}^{-1}|\mu_1 - \mu_0|$, and, by construction, $k(\mu_0, \mu_1)$ and $\gamma(\mu_0, \mu_1)$ have the same imaginary parts along the boundary. Then we define $K(\mu)$ to be $k(\mu_0, \mu_1)$ on the neck of modulus $2\mathfrak{r}$, the linear interpolation between $k(\mu_0, \mu_1)$ and $\gamma(\mu_0, \mu_1)$ on the the cut-off region (using the cut-off functions $\beta(s), \beta(-s)$), and $\gamma(\mu_0, \mu_1)$ everywhere else.

It suffices to say that the linearized operator takes the form:

$$(26) \quad (v, w, \mu, \xi) \mapsto D_{y_0, y_1, u_{s,\mathfrak{r}}}(\Phi(v) + \Psi(w) + K(\mu) + \xi) \in L^{p,\mathfrak{w}}.$$

We will now argue that this linearized operator is uniformly surjective.

Lemma 72. *For any $\eta \in L^{p,\mathfrak{w}}$, there exists (v, w, μ, ξ) so that:*

$$D_{y_0, y_1, u_{\mathfrak{s}, \mathfrak{r}}}(\Phi(v) + \Psi(w) + K(\mu) + \xi) = \eta,$$

and:

$$|v| + |w| + |\mu| + \|\xi\|_{W^{1,p,\mathfrak{w}}} \leq C_{\text{lin}} \|\eta\|,$$

where C_{lin} is independent of $\eta, \rho, \mathfrak{s}, \mathfrak{r}$ or the original solution (x_0, x_1, u) , provided $\mathfrak{r}$ is large enough.

Proof. Divide the neck $[-\mathfrak{r}, \mathfrak{r}]$ into two regions $[-\mathfrak{r}, 0] \cup [0, \mathfrak{r}]$ and introduce two cut-off functions:

$$\begin{aligned} (1) \quad f_-(s) &= \beta(1 - s/\mathfrak{r}), \\ (2) \quad f_+(s) &= \beta(1 + s/\mathfrak{r}), \end{aligned}$$

The key is that f_- is 1 on the left region, and cuts off in the right region, while f_+ cuts off in the left region, and is 1 on the right region. Another key is that the derivative of $f_{\pm}$ is $O(\mathfrak{r}^{-1})$.

We define two maps. First: $\mathfrak{R} : L^{p,\mathfrak{w}} \rightarrow L^{p,\delta}$ as an extension by zero map; this is illustrated in Figure 18. Modulo the small change in sizes due to parallel transport (taking variations of $u_{\mathfrak{s}, \mathfrak{r}}$ to u), this function is norm preserving by virtue of how the weights are defined. Second we define $\mathfrak{L} : L^{p,\delta} \rightarrow L^{p,\mathfrak{w}}$ using the cut-off functions f_-, f_+ as shown in Figure 19. The $\mathfrak{L}$ map is almost norm preserving (due to how the weights are defined), and only increases norms slightly.

Let us observe that $\mathfrak{L} \circ \mathfrak{R} = \text{id}$, and moreover $\mathfrak{L}$ maps $W^{1,p,\delta}$ into $W^{1,p,\mathfrak{w}}$ in an approximately norm preserving way.

We argue as follows: given $\eta \in L^{p,\mathfrak{w}}$, we can find v_-, v_+, w, ξ so that:

$$D_{x_0, x_1, u}(\Phi_-(v_-) + \Phi_+(v_+) + \Psi(w) + \xi) = \mathfrak{R}(\eta),$$

where $|v_-| + |v_+| + |w| + \|\xi\|_{W^{1,p,\delta}} \leq C_{\infty} \|\eta\|_{L^{p,\mathfrak{w}}}$, using the assumed surjectivity of the linearized operator for the genuine solution (x_0, x_1, u) .

Pick μ and v so that:

$$(27) \quad v + \mu_0 = v_- \text{ and } v + \mu_1 = v_+,$$

so that $\mu = v_+ - v_-$. We then claim that:

$$D_{y_0, y_1, u_{\mathfrak{s}, \mathfrak{r}}}(\Phi(v) + \Psi(w) + K(\mu) + \mathfrak{L}(\xi)) = \eta + O(\mathfrak{r}^{-1}) \|\eta\|.$$

Indeed, $K(\mu) + \Phi(v)$ is close to $\Phi_-(v_-)$ and $\Phi_+(v_+)$ on the cut-off regions, and the difference between $\bar{\partial}(K(\mu) + \Phi(v))$ and $\bar{\partial}(\Phi_{\pm}(v_{\pm}))$ is a small error of order $\mathfrak{r}^{-1}|\mu|$, using equation (27) and the fact $\gamma(\mu_0, \mu_1)$ has C^1 size controlled by $\mathfrak{r}^{-1}|\mu|$.

On the inner neck, the only contribution is $\bar{\partial}\mathfrak{L}(\xi)$ which is approximately $\eta = \mathfrak{L}(\mathfrak{R}(\eta))$ up to an error $O(\mathfrak{r}^{-1}) \|\eta\|$, due to derivatives of f_+, f_- .

The only other errors are due to the parallel transport, but this is of order $O(\mathfrak{r}^{-1}) \|\eta\|$, since the distance of curves we need to parallel transport along is $O(\mathfrak{r}^{-1})$ (noting that parallel transport acts identically in the neck).

Thus we can solve the equation up to an error of $O(\mathfrak{r}^{-1}) \|\eta\|$. Provided $O(\mathfrak{r}^{-1}) < 1/2$, we can then solve the equation by an iterative process, as is

common in solving linear equations in Banach spaces. The result gives an inverse image whose norm is uniformly bounded in terms of the input norm $\|\eta\|$ by some constant C_{lin} . $\square$




FIGURE 18. Extension by zero map $\mathfrak{R}$; this goes from deformations of the preglued map u_R to deformations of the original map u .




FIGURE 19. Partition of unity map $\mathfrak{L}$; this goes from deformations of u to deformations of u_R .

As a consequence of this result, and the previous result Lemma 71, one concludes by the usual application of the inverse function theorem in Banach spaces that there are indeed genuine solutions of $\mathcal{M}(R)$ close to the preglued solutions, provided the parameter $\mathfrak{s}$ is large, and $\mathfrak{r}$ is chosen much larger than $\mathfrak{s}$, so that the failure of the preglued solution to solve the equation is small enough (see Lemma 71).

Moreover, these genuine solutions are cut transversally, because the linearization at the preglued solutions is uniformly surjective. The solutions are rigid in $\mathcal{M}(R)$ as can be deduced by the index formula together with the above uniform surjectivity. Let us denote this genuine solution by $G_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u)$. This rigidity leads to the following conclusion:

Lemma 73. *Let $\mathfrak{s}_0 \geq 0$ be a constant, and let $P(\mathfrak{s})$ be such that $\mathfrak{r} \geq P(\mathfrak{s})$ and $\mathfrak{s} \geq \mathfrak{s}_0$ implies that the gluing argument applied to $PG_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u)$ converges to a genuine solution $G_{\mathfrak{s}, \mathfrak{r}}(x_0, x_1, u)$ (this requires $\mathfrak{r}$ to be large enough for the linearized operator to be uniformly surjective, and then the error in Lemma 71 to be small enough). Now suppose that $\mathfrak{s}(\tau), \mathfrak{r}(\tau)$ varies continuously and so that:*

- (1) $\mathfrak{s}(\tau) \geq \mathfrak{s}_0$ and $\mathfrak{r}(\tau) \geq P(\mathfrak{s}(\tau))$ for all τ ,
- (2) $\mathfrak{s}(\tau) + \mathfrak{r}(\tau) = \mathfrak{s}(0) + \mathfrak{r}(0)$,

then $G_{\mathfrak{s}(\tau), \mathfrak{r}(\tau)}(x_0, x_1, u) = G_{\mathfrak{s}(0), \mathfrak{r}(0)}(x_0, x_1, u)$. In particular, the glued solution in $\mathcal{M}(R)$ is $G_{\mathfrak{s}, R-\mathfrak{s}}$, provided $\mathfrak{s} \geq \mathfrak{s}_0$ and $R \geq \mathfrak{s} + P(\mathfrak{s})$.

Proof. This follows from the rigidity of the solutions, and the fact the pregluing construction is continuous with respect to variations of $\mathfrak{s}$ and $\mathfrak{r}$. $\square$

This completes part (2) of the gluing argument. In the next and final subsection, we will complete step (3) of the gluing argument.

4.4.10. Compactness for ends of $\mathcal{M}$. Suppose that (R_n, x_0^n, x_1^n, u_n) is a sequence of genuine solutions in $\mathcal{M}$, with $R_n \rightarrow \infty$. To complete the gluing argument, we need to show that (after passing to a subsequence) (R_n, x_0^n, x_1^n, u_n) eventually equals the glued solution $G_{\mathfrak{s}, \mathfrak{r}_n}(x_0, x_1, u)$, where $\mathfrak{r}_n = R_n - \mathfrak{s}$, and (x_0, x_1, u) is an appropriate Gromov limit of (x_0^n, x_1^n, u_n) .

First, pick the subsequence so that (x_0^n, x_1^n) converge to a limit (x_0, x_1) . For later use, let us suppose the subsequence is such that x_0^n is in the ball $B(\epsilon)$ around x_0 . After passing to a further subsequence, the restriction of u_n to the neck $[-R_n, R_n]$ converges on compact subsets to a limiting holomorphic strip with boundary on cotangent fibers $T^*M_{A_0(x_0, 0)}$ and $T^*M_{A_1(x_1, 0)}$. Since this limiting holomorphic strip has bounded energy (by Fatou's lemma), it must be that $T^*M_{A_0(x_0, 0)} = T^*M_{A_1(x_1, 0)}$, otherwise no such finite energy holomorphic strip exists. Thus we conclude $A_0(x_0, 0) = A_1(x_1, 0)$, so $(x_0, x_1) \in P_\infty$.

By passing to yet a further sequence, standard elliptic regularity and Floer theory compactness results imply that u_n converges on compact subsets of the limiting domain $\Sigma(\infty)$ to a smooth map u — here we use the identification of the surfaces with sufficiently deep ends/necks removed, so that any compact subset of $\Sigma(\infty)$ eventually is identified with a compact subset of $\Sigma(R_n)$.

Let $(\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n) = G_{\mathfrak{s}, \mathfrak{r}_n}(x_0, x_1, u)$. It is sufficient to prove there is a variation of $(\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n)$ of small norm which equals (x_0^n, x_1^n, u_n) after applying the Riemannian exponential map, as then we can appeal to the rigidity of the glued solution to conclude $(\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n) = (x_0^n, x_1^n, u_n)$.

It is clear from the construction that $(\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n)$ also converges to (x_0, x_1, u) as $n \rightarrow \infty$, in the same manner that (x_0^n, x_1^n, u_n) does.

In particular, the restrictions of $\bar{u}_n, u_n$ to $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$ are both holomorphic strips whose endpoints are eventually within the $C_\pm e^{-\pi \mathfrak{s}}$ neighborhoods of $p_\pm$, since the limit (x_0, x_1, u) has its endpoints within this neighborhood, as explained in §4.4.8.

Let us abbreviate by q_0^n, q_1^n and $\bar{q}_0^n, \bar{q}_1^n$ the imaginary parts of the boundary components of u_n and $\bar{u}_n$. Then we estimate:

Lemma 74. *It holds that:*

$$\mathfrak{r}_n |(q_1^n - q_0^n) - (\bar{q}_1^n - \bar{q}_0^n)| \leq (C_- + C_+) e^{-\pi \mathfrak{s}}.$$

Proof. Let $U = u_n - \bar{u}_n$, restricted to $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$. Let $q(s, t), p(s, t)$ be the coordinate projections of U to $\mathbb{R}^n \times B(1)$. Then:

$$\partial_s p = \partial_t q \implies 2\mathfrak{r}_n ((q_1^n - q_0^n) - (\bar{q}_1^n - \bar{q}_0^n)) = \int_0^1 p(\mathfrak{r}_n, t) - p(-\mathfrak{r}_n, t) dt$$

Using that U has its $s = \pm \mathfrak{r}_n$ boundary component in the $2C_\pm e^{-\pi \mathfrak{s}}$ neighborhood of 0, we conclude the desired result. $\square$

Next, we recall from §4.4.9 that there is a variation $K(\mu_n)$ associated to $\mu_n \in \mathbb{R}^n$ whose restriction to $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$ equals:

$$2\mathfrak{r}_n^{-1}i\mu_{n,0} + 2\mathfrak{r}_n^{-1}(\mu_{n,0}(z + \mathfrak{r}_n) + \mu_{n,1}(\mathfrak{r}_n - z)),$$

and which has a C^1 distance controlled by $\mathfrak{r}_n^{-1}(\mu_{n,1} - \mu_{n,0})$ outside of the neck $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$. Let us “add” this variation to $(\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n)$ (using the Riemannian exponential map) so that:

$$2\mathfrak{r}_n^{-1}(\mu_{n,1} - \mu_{n,0}) = q_1^n - q_0^n - (\bar{q}_1^n - \bar{q}_0^n).$$

Then it holds that:

$$u_n - (\bar{u}_n + K(\mu_n))$$

has boundary conditions on a single cotangent fiber. Moreover, Lemma 74 $|\mu_n|$ is controlled by $2^{-1}(C_- + C_+)e^{-\pi\mathfrak{s}}$. It then follows that:

$$\|u_n - (\bar{u}_n + K(\mu_n))\|_{W^{1,p,\mathfrak{w}}} \leq O(\mathfrak{r}_n^{-1}e^{\delta\mathfrak{s}}) + O(\mathfrak{e}^{\delta\mathfrak{s}}\text{dist}_{\mathfrak{s}}(u_n, \bar{u}_n))$$

where $\text{dist}_{\mathfrak{s}}$ is the C^1 distance computed on the complement of $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$. To see this, one uses:

- (1) the aforementioned bound of the C^1 size $K(\mu_n)$ on the complement of $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$,
- (2) the fact the $W^{1,p,\mathfrak{w}}$ distance on the complement of $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$ is bounded by $e^{\delta\mathfrak{s}}$ times the usual C^1 distance, and
- (3) $u_n - (\bar{u}_n + K(\mu))$ is holomorphic on $[-\mathfrak{r}_n, \mathfrak{r}_n] \times [0, 1]$, with boundary on a single cotangent fiber, and therefore the $W^{1,p,\mathfrak{w}}$ size on this neck is bounded by the C^1 size at the endpoints.

Thus we conclude that:

$$u_n = \bar{u}_n + K(\mu_n) + \xi_n$$

where $\|\xi_n\|_{W^{1,p,\mathfrak{w}}} = O(\mathfrak{r}_n^{-1}e^{\delta\mathfrak{s}}) + O(\mathfrak{e}^{\delta\mathfrak{s}}\text{dist}_{\mathfrak{s}}(u_n, \bar{u}_n))$. In particular, by taking n large enough, we can make $\|\xi_n\|_{W^{1,p,\mathfrak{w}}} + |\mu_n|$ as small as desired. In particular, (x_0, x_1, u_n) enters arbitrarily small neighborhoods of $(\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n)$ in the correct topology for the uniform surjectivity of the linearized operator to be applied. By the rigidity of the glued solution, it follows that $(x_0, x_1, u_n) = (\bar{x}_0^n, \bar{x}_1^n, \bar{u}_n)$, as desired.

This completes the proof that the count of non-compact ends of $\mathcal{M}$ equals the count of rigid solutions of $\mathcal{M}(\infty)$. $\square$

References

- [AAC23] H. Alizadeh, M. S. Atallah, and D. Cant, *Lagrangian intersections and the spectral norm in convex-at-infinity symplectic manifolds*, arXiv:2312.14752, 12 2023, Accepted to Math. Z., doi:10.48550/arXiv.2312.14752.
- [Abo15] M. Abouzaid, *Symplectic cohomology and Viterbo’s theorem*, Free Loop Spaces in Geometry and Topology, European Mathematical Society, 2015, pp. 271–486, doi:10.4171/153.
- [Alu09] P. Aluffi, *Algebra: Chapter 0*, Graduate Studies in Mathematics, vol. 104, AMS, 2009, doi:10.1090/gsm/104.
- [APS08] A. Abbondandolo, A. Portaluri, and M. Schwarz, *The homology of path spaces and Floer homology with conormal boundary conditions*, J. Fixed Point Theory Appl. **4** (2008), 263–293, doi:10.1007/S11784-008-0097-Y.

- [AS06] A. Abbondandolo and M. Schwarz, *On the Floer homology of cotangent bundles*, Comm. Pure Appl. Math. **59** (2006), 254–316, doi:10.1002/cpa.20090.
- [AS10] ———, *Floer homology of cotangent bundles and the loop product*, Geom. Topol. **14** (2010), 1569–1722, doi:10.2140/GT.2010.14.1569.
- [AS12] ———, *On product structures in Floer homology of cotangent bundles*, Global Differential Geometry, Springer Berlin Heidelberg, 2012, pp. 491–521.
- [BC07] P. Biran and O. Cornea, *Quantum structures for lagrangian submanifolds*, arXiv:0708.4221, 2007, doi:10.48550/arXiv.0708.4221.
- [BC24] F. Broćić and D. Cant, *Bordism classes of loops and Floer’s equation in cotangent bundles*, J. Fixed Point Theory Appl. **26** (2024), 1–29, doi:10.1007/s11784-024-01114-x.
- [BCS25] F. Broćić, D. Cant, and E. Shelukhin, *The chord conjecture for conormal bundles*, Math. Ann. (2025), 63, doi:10.1007/s00208-025-03171-0.
- [Bim24] J. Bimmermann, *Hofer–Zehnder capacity of magnetic disc tangent bundles over constant curvature surfaces*, Arch. Math. **123** (2024), 103–111, doi:10.1007/s00013-024-02003-y.
- [BK22] G. Benedetti and J. Kang, *Relative Hofer–Zehnder capacity and positive symplectic homology*, J. Fixed Point Theory Appl. **24** (2022), no. 44, 1–32, doi:10.1007/s11784-022-00963-8.
- [Bro25] F. Broćić, *Riemannian distance and symplectic embeddings in cotangent bundle*, Comm. Cont. Math. **27** (2025), no. 03, 1–28, doi:10.1142/S021919972450024X.
- [Can22] D. Cant, *A dimension formula for relative symplectic field theory*, Stanford University PhD Thesis. Available at: <https://dylancant.ca/thesis.pdf>, 7 2022.
- [Can24] ———, *Remarks on eternal classes in symplectic cohomology*, arXiv:2410.03914, 10 2024, Submitted, doi:10.48550/arXiv.2410.03914.
- [CC23] D. Cant and D. Chen, *Adiabatic compactness for holomorphic curves with boundary on nearby Lagrangians*, arXiv:2302.13391, 02 2023, 77 pages. Accepted to Kyoto J. Math., doi:10.48550/arXiv.2302.13391.
- [CHO23] K. Cieliebak, N. Hingston, and A. Oancea, *Loop coproduct in Morse and Floer homology*, J. Fixed Point Theory Appl. **25** (2023), no. 59, 1–84, doi:10.1007/s11784-023-01061-z.
- [Cie94] K. Cieliebak, *Pseudo-holomorphic curves and periodic orbits on cotangent bundles*, J. Math. Pure Appl. **9** (1994), no. 73, 251–278.
- [Cie02] ———, *Handle attaching in symplectic homology and the chord conjecture*, J. European Math. Soc. **4** (2002), no. 2, 115–142, doi:10.1007/s100970100036.
- [CS99] M. Chas and D. Sullivan, *String topology*, arXiv:math/9911159v1, 11 1999, doi:10.48550/arXiv.math/9911159.
- [EENS13] T. Ekholm, J. Etnyre, L. Ng, and M. G. Sullivan, *Knot contact homology*, Geom. Topol. **17** (2013), 975–1112, doi:10.2140/gt.2013.17.975.
- [EG91] Y. Eliashberg and M. Gromov, *Convex symplectic manifolds*, Several complex variables and complex geometry, part 2, Proc. Sympos. Pure Math., vol. 52, AMS, 1991, pp. 135–162, doi:10.1090/pspum/052.2.
- [Ekh07] T. Ekholm, *Morse flow trees and Legendrian contact homology in 1-jet spaces*, Geom. Topol. **11** (2007), 1083–1224, doi:10.2140/gt.2007.11.1083.
- [EP03] M. Entov and L. Polterovich, *Calabi quasimorphism and quantum homology*, Int. Math. Res. Not. **2003** (2003), 1635–1676, doi:10.1155/S1073792803210011.
- [EP22] Y. Eliashberg and D. Pancholi, *Honda–Huang’s work on contact convexity revisited*, arXiv:2207.07185, 2022, doi:10.48550/arXiv.2207.07185.
- [Fau20] A. Fauck, *On manifolds with infinitely many fillable contact structures*, Int. J. Math. **31** (2020), no. 13, 1–71, doi:10.1142/S0129167X20501086.
- [FHS95] A. Floer, H. Hofer, and D. Salamon, *Transversality in elliptic Morse theory for the symplectic action*, Duke Math. J. **80** (1995), no. 1, 251–292, doi:10.1215/S0012-7094-95-08010-7.
- [Flo89] A. Floer, *Symplectic fixed points and holomorphic spheres*, Comm. Math. Phys. **120** (1989), 575–611, doi:10.1007/BF01260388.

- [FO97] K. Fukaya and Y-G. Oh, *Zero-loop open strings in the cotangent bundle and Morse homotopy*, Asian J. Math. **1** (1997), no. 1, 96–180, doi:10.4310/AJM.1997.V1.N1.A5.
- [FRV23] B. Ferreira, V. G. B. Ramos, and A. Vicente, *Gromov width of the disk cotangent bundle of spheres of revolution*, arXiv:2301.08528, 2023, doi:10.48550/arXiv.2301.08528.
- [FS07] U. Frauenfelder and F. Schlenk, *Hamiltonian dynamics on convex symplectic manifolds*, Isr. J. Math. **159** (2007), 1–56, doi:10.1007/S11856-007-0037-3.
- [Fuk97] K. Fukaya, *Morse homotopy and its quantization*, Geometric topology: 1993 Georgia international topology conference (William H Kazez, ed.), Studies in Advanced Mathematics, vol. 2, American Mathematical Society and International Press of Boston, 1997, pp. 409–440, doi:10.1090/amsip/002.1.
- [GT98] D. Gilbarg and N. S. Trudinger, *Elliptic partial differential equations of second order*, revised 3rd printing ed., Grundlehren der mathematischen Wissenschaften, vol. 224, Springer, 1998, doi:10.1007/978-3-642-61798-0.
- [HS95] H. Hofer and D. Salamon, *Floer homology and Novikov rings*, The Floer memorial volume, Progr. Math., vol. 133, Birkhäuser, 1995, pp. 483–524, doi:10.1007/978-3-0348-9217-9.
- [Hut08] M. Hutchings, *Floer homology of families I*, Alg. Geom. Topol. **8** (2008), 435–492, doi:10.2140/agt.2008.8.435.
- [Iri14] K. Irie, *Hofer-Zehnder capacity of unit disk cotangent bundles and the loop product*, J. Eur. Math. Soc. **16** (2014), 2477–2497, doi:10.4171/JEMS/491.
- [KS21] A. Kislev and E. Shelukhin, *Bounds on spectral norms and barcodes*, Geometry & Topology **25** (2021), 3257–3350, doi:10.2140/gt.2021.25.3257.
- [Mac71] S. Mac Lane, *Categories for the working mathematician*, 2nd edition ed., Graduate texts in mathematics, vol. 5, Springer, 1971, doi:10.1007/978-1-4757-4721-8.
- [Mil65a] J. Milnor, *Lectures on the h-cobordism theorem*, Princeton University Press, 1965, doi:10.1515/9781400878055.
- [Mil65b] ———, *Topology from the differentiable viewpoint*, The University Press of Virginia, 1965, doi:10.2307/2314613.
- [MS12] D. McDuff and D. Salamon, *j-holomorphic curves and symplectic topology*, 2nd edition ed., American Mathematical Society, Colloquium Publications, 2012, doi:10.1090/coll/052.
- [MS17] ———, *Introduction to symplectic topology*, 3rd ed., Oxford University Press, 2017, doi:10.1093/oso/9780198794899.001.0001.
- [MT94] D. McDuff and L. Traynor, *The 4-dimensional symplectic camel and related results*, Symplectic Geometry (D. Salamon, ed.), London Mathematical Society Lecture Note Series, Cambridge University Press, 1994, pp. 169–182, doi:10.1017/CBO9780511526343.010.
- [OZ11] Y-G. Oh and K. Zhu, *Floer trajectories with immersed nodes and scale-dependent gluing*, J. Symplectic Geom. **9** (2011), no. 4, 483–636, doi:10.4310/JSG.2011.v9.n4.a4.
- [PSS96] S. Piunikhin, D. Salamon, and M. Schwarz, *Symplectic Floer-Donaldson theory and quantum cohomology*, Contact and symplectic geometry, Publications of the Newton Institute, vol. 8, Cambridge University Press, 1996, pp. 171–200.
- [Rit13] A. Ritter, *Topological quantum field theory structure on symplectic cohomology*, J. Topol. **6** (2013), 391–489, doi:10.1112/jtopol/jts038.
- [Sal97] D. Salamon, *Lectures on Floer homology*, IAS/Park City Graduate Summer School on Symplectic Geometry and Topology, 1997, doi:10.1090/pcms/007/05.
- [Sch95] M. Schwarz, *Cohomology operations from S^1 -cobordisms in Floer homology*, <https://www.math.uni-leipzig.de/~schwarz/diss.pdf>, 1995, PhD Dissertation.
- [Sch00] ———, *On the action spectrum for closed symplectically aspherical manifolds*, Pacific J. Math. **193** (2000), 419–461, doi:10.2140/PJM.2000.193.419.
- [Sei08] P. Seidel, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society, 2008, doi:10.4171/063.
- [Sei14] ———, *Disjoinable Lagrangian spheres and dilations*, Invent. Math. **197** (2014), 299–359, doi:10.1007/s00222-013-0484-x.

- [Sei15] ———, *The equivariant pair-of-pants product in fixed point Floer cohomology*, Geom. Funct. Anal. **25** (2015), 942–1007, doi:10.1007/s00039-015-0331-x.
- [She22a] E. Shelukhin, *Symplectic cohomology and a conjecture of Viterbo*, Geom. Funct. Anal. **32** (2022), 1514–1543, doi:10.1007/s00039-022-00619-2.
- [She22b] ———, *Viterbo conjecture for Zoll symmetric spaces*, Inventiones mathematicae **230** (2022), no. 1, 321–373, doi:10.1007/s00222-022-01124-x.
- [SS12] P. Seidel and J. P. Solomon, *Symplectic cohomology and q -intersection numbers*, Geom. Funct. Anal. **22** (2012), 443–477, doi:10.1007/s00039-012-0159-6.
- [SW06] D. Salamon and J. Weber, *Floer homology and the heat flow*, Geom. Funct. Anal. **16** (2006), 1050–1138, doi:10.1007/s00039-006-0577-4.
- [Vit92] C. Viterbo, *Symplectic topology as the geometry of generating functions*, Math. Ann. **292** (1992), 685–710, doi:10.1007/BF01444643.
- [Vit99] ———, *Functors and computations in Floer homology with applications I*, Geom. Funct. Anal. **9** (1999), 985–1033, doi:10.1007/s000390050106.
- [Wei91] A. Weinstein, *Contact surgery and symplectic handlebodies*, Hokkaido Mathematical Journal **20** (1991), no. 2, 241 – 251, doi:10.14492/hokmj/1381413841.
- [Zho21] Z. Zhou, *Symplectic fillings of asymptotically dynamically convex manifolds i* , Journal of Topology **14** (2021), no. 1, 112–182, doi:10.1112/topo.12177.

Email address: `filipvbrocic@gmail.com`

LEHRSTUHL FÜR ANALYSIS UND GEOMETRIE, UNIVERSITÄT AUGSBURG, UNIVERSITÄTSSTRASSE 14, 86159 AUGSBURG, GERMANY

Email address: `dylan@dylancant.ca`

INSTITUT DE MATHÉMATIQUE D’ORSAY, UNIVERSITÉ PARIS-SACLAY, BÂTIMENT 307, RUE MICHEL MAGAT, F-91405 ORSAY CEDEX, FRANCE