Gluing pseudoholomorphic curves along branched covered cylinders I

Michael Hutchings and Clifford Henry Taubes

Abstract

This paper and its sequel prove a generalization of the usual gluing theorem for two index 1 pseudoholomorphic curves u_+ and u_- in the symplectization of a contact 3-manifold. We assume that for each embedded Reeb orbit γ, the total multiplicity of the negative ends of u_+ at covers of γ agrees with the total multiplicity of the positive ends of u_- at covers of γ. However, unlike in the usual gluing story, here the individual multiplicities are allowed to differ. In this situation, one can often glue u_+ and u_- to an index 2 curve by inserting genus zero branched covers of $\mathbb{R}$-invariant cylinders between them. We establish a combinatorial formula for the signed count of such gluings. As an application, we deduce that the differential ∂ in embedded contact homology satisfies $\partial^2 = 0$.

This paper explains the more algebraic aspects of the story, and proves the above formulas using some analytical results from part II.

Contents

1 Statement of the gluing theorem

1.1 Pseudoholomorphic curves in symplectizations

Our gluing theorem concerns pseudoholomorphic curves in the symplectization of a contact 3-manifold. We now recall some mostly standard definitions and introduce some notation regarding such curves. The geometric setup here is essentially that of Hofer, Wysocki, and Zehnder [8], and the four-dimensional case of the setup used to define symplectic field theory [5].

Let Y be a closed oriented 3-manifold. Let λ be a contact form on Y, i.e. a 1-form λ such that $\lambda \wedge d\lambda > 0$. The associated contact structure is the 2-plane field $\xi := \operatorname{Ker}(\lambda)$. The contact form λ determines a vector field R on Y, called the Reeb vector field, which is characterized by $d\lambda(R, \cdot) = 0$ and $\lambda(R) = 1$. A *Reeb orbit* is a closed orbit of the flow R, i.e. a map $\gamma : \mathbb{R}/T\mathbb{Z} \to Y$ for some $T \in \mathbb{R}$, modulo reparametrization, such that $\partial_t \gamma(t) = R(\gamma(t))$. We do not require γ to be an embedding. Define the homology class of the Reeb orbit by $[\gamma] := \gamma_*[\mathbb{R}/T\mathbb{Z}] \in H_1(Y)$.

If γ is a Reeb orbit passing through a point $y \in Y$, then the linearized return map of the flow R along γ determines a symplectic linear map $P_\gamma : \xi_y \to \xi_y$. The eigenvalues of P_γ do not depend on y. We say that the Reeb orbit γ is *nondegenerate* if P_γ does not have 1 as an eigenvalue. We assume throughout that all Reeb orbits are nondegenerate; this condition holds for generic contact forms λ.

We now choose an almost complex structure J on the 4-manifold $\mathbb{R} \times Y$. We always assume that J is "admissible" in the following sense.

Definition 1.1. An almost complex structure J on $\mathbb{R} \times Y$ is *admissible* if:

- $J(\partial_s) = R$, where s denotes the $\mathbb{R}$ coordinate on $\mathbb{R} \times Y$.
- $J(\xi) = \xi$.
- J rotates ξ "positively" in the sense that $d\lambda(v, Jv) \geq 0$ for all $v \in \xi$.
- J is invariant under the $\mathbb{R}$ action on $\mathbb{R} \times Y$ that translates s.

A *J-holomorphic curve* in $\mathbb{R} \times Y$ is a triple (C, j, u) where C is a smooth surface, j is a complex structure on C, and $u : C \to \mathbb{R} \times Y$ is a smooth map such that $J \circ du = du \circ j$. The triple (C, j, u) is equivalent to the triple (C', j', u') iff there is a biholomorphic map $\phi : (C, j) \xrightarrow{\simeq} (C', j')$ such that $u' \circ \phi = u$. We always assume that the domain (C, j) is a punctured compact Riemann surface, possibly disconnected. We usually denote a J-holomorphic curve simply by u.

A *positive end* of u at a Reeb orbit γ is an end of C that can be parametrized by $(s, t) \in [R, \infty) \times S^1$ for some $R \in \mathbb{R}$, such that $u(s, t) = (s, y(s, t))$ and $\lim_{s \to \infty} y(s, \cdot)$ is a reparametrization of γ. A *negative end* of u at γ is defined analogously with $s \in (-\infty, -R]$.

For each embedded Reeb orbit γ, fix a point $y \in Y$ in the image of γ. If m is a positive integer, let γ^m denote the Reeb orbit that is an m-fold cover of γ. If u has an end at γ^m, then the intersection of this end with $\{s\} \times Y$ for $|s|$ large is an m-fold covering of γ via a normal bundle projection. An *asymptotic marking* of the end is an inverse image of y under this covering. This notion does not depend on the choice of normal bundle projection or on the choice of s with $|s|$ large. Note that $\mathbb{Z}/m$ acts freely and transitively on the set of asymptotic markings of an end of u at γ^m.

Definition 1.2. Let $\alpha = (\alpha_1, \ldots, \alpha_k)$ and $\beta = (\beta_1, \ldots, \beta_l)$ be ordered lists of Reeb orbits, possibly repeated. Define $\mathcal{M}^J(\alpha, \beta)$ to be the moduli space of J-holomorphic curves $u : C \to \mathbb{R} \times Y$ as above, such that u has ordered and asymptotically marked positive ends at $\alpha_1, \ldots, \alpha_k$, ordered and asymptotically marked negative ends at $\beta_1, \ldots, \beta_l$, and no other ends.

Note that $\mathcal{M}^J(\alpha, \beta) \neq \emptyset$ only if $\sum_{i=1}^k [\alpha_i] = \sum_{j=1}^l [\beta_j] \in H_1(Y)$. Also, since the $\mathbb{R}$ action on $\mathbb{R} \times Y$ preserves J, it induces an $\mathbb{R}$ action on $\mathcal{M}^J(\alpha, \beta)$.

Definition 1.3. If $u \in \mathcal{M}^J(\alpha, \beta)$, define the (Fredholm) *index*

$$\operatorname{ind}(u) := -\chi(C) + 2c_1(u^*\xi, \tau) + \sum_{i=1}^{k} \mathrm{CZ}_\tau(\alpha_i) - \sum_{j=1}^{l} \mathrm{CZ}_\tau(\beta_j). \tag{1.1}$$

The terms on the right hand side of equation (1.1) are defined as follows. First, C is the domain of u as above. Second, τ is a trivialization of ξ over the Reeb orbits α_i and β_j; it turns out that $\operatorname{ind}(u)$ does not depend on τ, although individual terms in its definition do. Next, $c_1(u^*\xi, \tau)$ denotes the relative first Chern class of the complex line bundle $u^*\xi$ over C with respect to the trivializations τ at the ends. This is defined by counting the zeroes of a generic section which at the ends is nonvanishing and constant with respect to the chosen trivializations. Finally, $\mathrm{CZ}_\tau(\gamma)$ denotes the Conley-Zehnder index of γ with respect to τ.

In the present setting where $\dim(Y) = 3$, this Conley-Zehnder index is described explicitly as follows. Let γ be an embedded Reeb orbit. Let τ be a trivialization of ξ over γ, and use τ to trivialize ξ over γ^m for each positive integer m. Our assumption that all Reeb orbits are nondegenerate implies that the linearized return map $P_{\gamma^m} = P_\gamma^m$ does not have 1 as an eigenvalue. Let λ, λ^{-1} denote the eigenvalues of P_γ. We say that γ is *positive hyperbolic* if $\lambda, \lambda^{-1} > 0$, *negative hyperbolic* if $\lambda, \lambda^{-1} < 0$, and *elliptic* if λ, λ^{-1} are on the unit circle. If γ is hyperbolic, then there is an integer n such that the linearized Reeb flow along γ rotates the eigenspaces by angle πn with respect to τ, and

$$\mathrm{CZ}_\tau(\gamma^m) = mn. \tag{1.2}$$

The integer n is even when γ is positive hyperbolic and odd when γ is negative hyperbolic. If γ is elliptic, then there is an irrational number θ, which we call the "monodromy angle", such that

$$\mathrm{CZ}_\tau(\gamma^m) = 2\lfloor m\theta \rfloor + 1. \tag{1.3}$$

Here τ is homotopic to a trivialization in which the linearized Reeb flow along γ rotates by angle $2\pi\theta$.

We say that $u \in \mathcal{M}^J(\alpha, \beta)$ is "not multiply covered" if u does not multiply cover any component of its image. We say that u is "unobstructed" if the linear deformation operator associated to u is surjective; then $\mathcal{M}^J(\alpha, \beta)$ is a manifold near u. The following proposition is the 3-dimensional case of a result proved in [4], using an index calculation from [16].

Proposition 1.4. *If J is generic, and if $u \in \mathcal{M}^J(\alpha, \beta)$ is not multiply covered, then u is unobstructed, so that $\mathcal{M}^J(\alpha, \beta)$ is a manifold near u. Moreover, this manifold has dimension* $\operatorname{ind}(u)$.

Assume henceforth that J is generic in this sense.

Following [3], one can "coherently" orient all the moduli spaces of non-multiply covered J-holomorphic curves by making one orientation choice for each Reeb orbit. (We use slightly different conventions from [3], and we make a canonical choice for each elliptic Reeb orbit; see [13, §9] for details.) Given coherent orientations, if M is a non-$\mathbb{R}$-invariant component of such a moduli space, we orient M using the $\mathbb{R}$ direction first. That is, if $u \in M$, if $v_1 \in T_uM$ denotes the derivative of the $\mathbb{R}$ action on M, and if $(v_1, \ldots, v_n)$ is an oriented basis for T_uM, then we declare that the projection of $(v_2, \ldots, v_n)$ is an oriented basis for $T_u(M/\mathbb{R})$. If u has index 1, then the above convention defines a sign, which we denote by $\epsilon(u) \in \{\pm 1\}$.

Remark 1.5. It follows from [3] that a system of coherent orientations behaves as follows under the diffeomorphisms between moduli spaces obtained by changing the orderings and asymptotic markings of the ends. If one switches the order of two ends, then this switches the orientation if and only if both ends are at positive hyperbolic Reeb orbits. (Here an even cover of a negative hyperbolic orbit is classified as positive hyperbolic.) If γ is an embedded Reeb orbit, and if one acts on the asymptotic marking of an end at γ^m by a generator of $\mathbb{Z}/m$, then this switches the orientation if and only if m is even and γ is negative hyperbolic.

1.2 Branched covered cylinders

To prepare for the statement of the gluing theorem, we now calculate the index of branched covers of $\mathbb{R}$-invariant cylinders.

Definition 1.6. If $a_1, \ldots, a_k$ and $b_1, \ldots, b_l$ are positive integers with

$$\sum_{i=1}^{k} a_i = \sum_{j=1}^{l} b_j = M$$

and if θ is an irrational number, define

$$\operatorname{ind}_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l) := 2\left(\sum_{i=1}^k \lceil a_i\theta\rceil - \sum_{j=1}^l \lfloor b_j\theta\rfloor - 1\right).$$

Note that $\operatorname{ind}_\theta \geq 0$, because $\sum_i \lceil a_i\theta\rceil \geq \lceil M\theta\rceil$ and $\sum_j \lfloor b_j\theta\rfloor \leq \lfloor M\theta\rfloor$.

Lemma 1.7. *Suppose $u \in \mathcal{M}^J(\alpha,\beta)$ is a branched cover of $\mathbb{R}\times\gamma$, where γ is an embedded Reeb orbit. Then $\operatorname{ind}(u) \geq 0$, with equality only if:*

(i) Each component of the domain C of u has genus 0.

(ii) If γ is hyperbolic, then the covering $u : C \to \mathbb{R}\times\gamma$ has no branch points.

Proof. Without loss of generality, C is connected; let g denote its genus. Write $\alpha = (\gamma^{a_1},\ldots,\gamma^{a_k})$ and $\beta = (\gamma^{b_1},\ldots,\gamma^{b_l})$. To calculate $\operatorname{ind}(u)$, choose a trivialization τ of $\gamma^*\xi$, and use this to trivialize ξ over all of the ends of u. Since τ extends to a trivialization of ξ over $\mathbb{R}\times\gamma$, it follows that $c_1(u^*\xi,\tau) = 0$. Thus

$$\operatorname{ind}(u) = -\chi(C) + \sum_{i=1}^k \mathrm{CZ}_\tau(\gamma^{a_i}) - \sum_{j=1}^l \mathrm{CZ}_\tau(\gamma^{b_j}). \tag{1.4}$$

If γ is hyperbolic, then by equation (1.2), the Conley-Zehnder index terms in equation (1.4) cancel, so $\operatorname{ind}(u) = -\chi(C) \geq 0$. If equality holds, then C is a cylinder, and by Riemann-Hurwitz there are no branch points.

Now suppose that γ is elliptic with monodromy angle θ with respect to τ. Then by equations (1.3) and (1.4),

$$\begin{aligned}\operatorname{ind}(u) &= (2g-2+k+l) + \sum_{i=1}^k (2\lceil a_i\theta\rceil - 1) - \sum_{j=1}^l (2\lfloor b_j\theta\rfloor + 1)\\ &= 2g + \operatorname{ind}_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l).\end{aligned} \tag{1.5}$$

Since $\operatorname{ind}_\theta \geq 0$, it follows that $\operatorname{ind}(u) \geq 0$, with equality only if $g = 0$. □

Recall that a *partition* of a nonnegative integer M is a list of positive integers $(a_1,\ldots,a_k)$ modulo reordering, possibly with repetitions, such that $\sum_{i=1}^k a_i = M$. In connection with the above index calculation, we now define a partial order on the set of partitions of M.

Definition 1.8. Fix θ irrational. We say that $(a_1,\ldots,a_k) \geq_\theta (b_1,\ldots,b_l)$ if whenever γ is an elliptic Reeb orbit with monodromy angle θ, there exists an index zero branched cover of $\mathbb{R}\times\gamma$ in $\mathcal{M}^J((\gamma^{a_1},\ldots,\gamma^{a_k}),(\gamma^{b_1},\ldots,\gamma^{b_l}))$. It is an exercise (which we will not need) to check that $\geq_\theta$ is a partial order.

1.3 Statement of the gluing problem

The following definition specifies the kinds of pairs of curves that we will be gluing.

Definition 1.9. A *gluing pair* is a pair of immersed J-holomorphic curves $u_+ \in \mathcal{M}^J(\alpha_+, \beta_+)$ and $u_- \in \mathcal{M}^J(\beta_-, \alpha_-)$ such that:

(a) $\operatorname{ind}(u_+) = \operatorname{ind}(u_-) = 1$.

(b) u_+ and u_- are not multiply covered, except that they may contain unbranched covers of $\mathbb{R}$-invariant cylinders.

(c) For each embedded Reeb orbit γ, the total covering multiplicity of Reeb orbits covering γ in the list β_+ is the same as the total for β_-. (In contrast, for the usual form of gluing one would assume that $\beta_+ = \beta_-$.)

(d) If γ is an elliptic embedded Reeb orbit with monodromy angle θ, let $a'_1, \ldots, a'_{k'}$ denote the covering multiplicities of the $\mathbb{R}$-invariant cylinders over γ in u_+, and let $b'_1, \ldots, b'_{l'}$ denote the corresponding multiplicities in u_-. Then under the partial order $\geq_\theta$ in Definition 1.8, the partition $(a'_1, \ldots, a'_{k'})$ is minimal, and the partition $(b'_1, \ldots, b'_{l'})$ is maximal.

Let (u_+, u_-) be a gluing pair. Our gluing theorem computes an integer $\#G(u_+, u_-)$ which, roughly speaking, is a signed count of ends of the index 2 part of the moduli space $\mathcal{M}^J(\alpha_+, \alpha_-)/\mathbb{R}$ that break into u_+ and u_- along with some index zero branched covers of $\mathbb{R}$-invariant cylinders between them. The precise definition of $\#G(u_+, u_-)$ is a bit technical and occupies the rest of this subsection. There is some subtlety here when u_+ or u_- contain covers of $\mathbb{R}$-invariant cylinders; in this case we will use condition (d) above in showing that $\#G(u_+, u_-)$ is well-defined.

To prepare for the definition of the count $\#G(u_+, u_-)$, we first define a set $\mathcal{G}_\delta(u_+, u_-)$ of index 2 curves in $\mathcal{M}^J(\alpha_+, \alpha_-)$ which, roughly speaking, are close to breaking in the above manner. For the following definition, choose an arbitrary product metric on $\mathbb{R} \times Y$.

Definition 1.10. For $\delta > 0$, define $\mathcal{C}_\delta(u_+, u_-)$ to be the set of immersed (except possibly for finitely many singular points) surfaces in $\mathbb{R} \times Y$ that can be decomposed as $C_- \cup C_0 \cup C_+$, such that the following hold:

- There is a real number R_-, and a section ψ_- of the normal bundle to u_- with $|\psi_-| < \delta$, such that C_- is the $s \mapsto s + R_-$ translate of the $s \leq 1/\delta$ part of the exponential map image of ψ_-.

- Likewise, there is a real number R_+, and a section ψ_+ of the normal bundle to u_+ with $|\psi_+| < \delta$, such that C_+ is the $s \mapsto s + R_+$ translate of the $s \geq -1/\delta$ part of the exponential map image of ψ_+.
- $R_+ - R_- > 2/\delta$.
- C_0 is contained in the union of the radius δ tubular neighborhoods of the cylinders $\mathbb{R} \times \gamma$, where γ ranges over the embedded Reeb orbits covered by orbits in $\beta_\pm$.
- $\partial C_0 = \partial C_- \sqcup \partial C_+$, where the positive boundary circles of C_- agree with the negative boundary circles of C_0, and the positive boundary circles of C_0 agree with the negative boundary circles of C_+.

Let $\mathcal{G}_\delta(u_+, u_-)$ denote the set of index 2 curves in $\mathcal{M}^J(\alpha_+, \alpha_-) \cap \mathcal{C}_\delta(u_+, u_-)$.

To see that this definition does what it is supposed to, we have:

Lemma 1.11. *Given a gluing pair (u_+, u_-), there exists $\delta_0 > 0$ with the following property. Let $\delta \in (0, \delta_0)$ and let $\{[u_n]\}_{n=1,2,\ldots}$ be a sequence in $\mathcal{G}_\delta(u_+, u_-)/\mathbb{R}$. Then there is a subsequence which converges in the sense of [2] either to a curve in $\mathcal{M}^J(\alpha_+, \alpha_-)/\mathbb{R}$, or to a broken curve in which the top level is u_+, the bottom level is u_-, and all intermediate levels are unions of index zero branched covers of $\mathbb{R}$-invariant cylinders.*

Proof. By the compactness theorem in [2], any sequence of index 2 curves in $\mathcal{M}^J(\alpha_+, \alpha_-)/\mathbb{R}$ has a subsequence which converges to some broken curve. Moreover, the indices of the levels of the broken curve sum to 2.

If the sequence is in $\mathcal{G}_\delta(u_+, u_-)/\mathbb{R}$ with $\delta > 0$ sufficiently small, then by Lemma 1.7 and the definition of $\mathcal{G}_\delta$, one of the following two scenarios occurs:

(i) One level of the broken curve contains the index 1 component of u_+, and some lower level contains the index 1 component of u_-.

(ii) Some level contains two index 1 components or one index 2 component.

Moreover, all other components of all levels are index zero branched covers of $\mathbb{R}$-invariant cylinders. By condition (d) in the definition of gluing pair, any covers of $\mathbb{R}$-invariant cylinders in the top and bottom levels of the broken curve must be unbranched. It follows that in case (i), the top level is u_+ and the bottom level is u_-, while in case (ii), there are no other levels. $\square$

Definition 1.12. Fix coherent orientations and generic J as in Proposition 1.4, and let (u_+, u_-) be a gluing pair. If $\delta \in (0, \delta_0)$, then by Lemma 1.11 one can choose an open set $U \subset \mathcal{M}^J(\alpha_+, \alpha_-)/\mathbb{R}$ such that:

- $\mathcal{G}_{\delta'}(u_+, u_-)/\mathbb{R} \subset U \subset \mathcal{G}_\delta(u_+, u_-)/\mathbb{R}$ for some $\delta' \in (0, \delta)$.
- The closure $\overline{U}$ has finitely many boundary points.

Define $\#G(u_+, u_-) \in \mathbb{Z}$ to be minus the signed count of boundary points of $\overline{U}$. By Lemma 1.11 , this does not depend on the choice of δ or U.

1.4 Statement of the main theorem

Let (u_+, u_-) be a gluing pair. The main result of this paper gives a combinatorial formula for $\#G(u_+, u_-)$. To state the formula, note first that by Lemma 1.7, if $\#G(u_+, u_-) \neq 0$ then for each hyperbolic Reeb orbit γ, the multiplicities of the negative ends of u_+ at covers of γ agree, up to reordering, with the multiplicities of the positive ends of u_- at covers of γ. When this is the case, assume that the orderings of the negative ends of u_+ and of the positive ends of u_- are such that for each positive hyperbolic orbit γ, the aforementioned multiplicities appear in the same order for u_+ and for u_-. With this ordering convention, the statement of the main theorem is as follows:

Theorem 1.13. *Fix coherent orientations. If J is generic and if (u_+, u_-) is a gluing pair, then*

$$\#G(u_+, u_-) = \epsilon(u_+)\epsilon(u_-) \prod_\gamma c_\gamma(u_+, u_-). \tag{1.6}$$

Here the product is over embedded Reeb orbits γ such that u_+ has a negative end at a cover of γ. The integer $c_\gamma(u_+, u_-)$, defined below, depends only on γ and on the multiplicities of the $\mathbb{R}$-invariant and non-$\mathbb{R}$-invariant negative ends of u_+ and positive ends of u_- at covers of γ.

To complete the statement of Theorem 1.13, we now define the "gluing coefficients" $c_\gamma(u_+, u_-)$ that appear in equation (1.6). We will use the following notation. Let $a_1, \ldots, a_k$ denote the multiplicities of the non-$\mathbb{R}$-invariant negative ends of u_+ at covers of γ (in some arbitrary order). Let $a'_1, \ldots, a'_{k'}$ denote the multiplicities of the $\mathbb{R}$-invariant components of u_+ at covers of γ. Likewise, let $b_1, \ldots, b_l$ denote the multiplicities of the non-$\mathbb{R}$-invariant positive ends of u_- at covers of γ, and let $b'_1, \ldots, b'_{l'}$ denote the multiplicities of the $\mathbb{R}$-invariant components of u_- at covers of γ. We will define

$$c_\gamma(u_+, u_-) := c_\gamma(a_1, \ldots, a_k; a'_1, \ldots, a'_{k'} \mid b_1, \ldots, b_l; b'_1, \ldots, b'_{l'}), \tag{1.7}$$

where the right hand side of (1.7) is defined below.

1.5 The gluing coefficients c_γ for hyperbolic γ

The gluing coefficient c_γ is relatively straightforward when γ is hyperbolic. In this case the gluing over γ does not involve any branch points, by Lemma 1.7. One just needs to match up negative ends of u_+ at covers of γ with positive ends of u_- at covers of γ with the same multiplicity. Also, when gluing two ends at the m-fold cover γ^m, there are m possibilities for matching up the sheets. The signs of these different matchings are related according to Remark 1.5. In many cases, the various possibilities all cancel out because of the orientations; while in the remaining cases, all possibilities have the same sign.

Definition 1.14. Suppose γ is hyperbolic. Then $c_\gamma = 0$ unless:

(a) The list of multiplicities $(a_1, \ldots, a_k, a'_1, \ldots, a'_{k'})$ is a permutation of the list $(b_1, \ldots, b_l, b'_1, \ldots, b'_{l'})$.

(b) If γ is positive hyperbolic, then the numbers $a_1, \ldots, a'_{k'}$ are distinct.

(c) If γ is negative hyperbolic, then the numbers $a_1, \ldots, a'_{k'}$ are all odd.

If (a), (b), and (c) hold, then for each positive integer m, let $r(m)$ denote the number of times that the number m appears in the list $a_1, \ldots, a'_{k'}$, and define

$$c_\gamma := \prod_{m=1}^{\infty} m^{r(m)} \cdot r(m)!$$

1.6 The gluing coefficients c_γ for elliptic γ

The interesting case of the gluing coefficient c_γ is when γ is elliptic with monodromy angle θ. Here the only relevant feature of γ is the irrational number θ, so we denote c_γ by c_θ. In this section we give a recursive definition of c_θ which is easy to compute with. An alternate definition of c_θ as a sum over forests, which is useful for proving certain symmetry properties of c_θ, is given in §4.

To simplify the notation, denote the arguments of the function c_θ by

$$S := (a_1, \ldots, a_k; a'_1, \ldots, a'_{k'} \mid b_1, \ldots, b_l; b'_1, \ldots, b'_{l'}). \tag{1.8}$$

When $k' = 0$ or $l' = 0$ we drop the corresponding semicolon from the notation (1.8). It is always assumed that

$$\sum_{i=1}^{k} a_i + \sum_{i=1}^{k'} a'_i = \sum_{j=1}^{l} b_j + \sum_{j=1}^{l'} b'_j. \tag{1.9}$$

Definition 1.15. If S as in (1.8) satisfies (1.9), define a positive integer

$$\kappa_\theta(S) := \sum_{i=1}^{k} \lceil a_i\theta \rceil + \sum_{i=1}^{k'} \lceil a_i'\theta \rceil - \sum_{j=1}^{l} \lfloor b_j\theta \rfloor - \sum_{j=1}^{l'} \lfloor b_j'\theta \rfloor . \tag{1.10}$$

The significance of $\kappa_\theta(S)$ is that by the calculation (1.5), any index zero branched cover of $\mathbb{R} \times \gamma$ with positive ends of multiplicities $a_1, \ldots, a'_{k'}$ and negative ends of multiplicities $b_1, \ldots, b'_{l'}$ must consist of $\kappa_\theta(S)$ genus zero components.

To define $c_\theta(S)$, we first reduce to the case where $\kappa_\theta(S) = 1$. We need to consider the different ways that the ends of a branched cover can be divided among $\kappa_\theta(S)$ different components.

Definition 1.16. A *θ-decomposition* of S is decomposition

$$\begin{aligned} \{1, \ldots, k\} &= I_1 \sqcup \cdots \sqcup I_{\kappa_\theta(S)}, \\ \{1, \ldots, k'\} &= I_1' \sqcup \cdots \sqcup I'_{\kappa_\theta(S)}, \\ \{1, \ldots, l\} &= J_1 \sqcup \cdots \sqcup J_{\kappa_\theta(S)}, \\ \{1, \ldots, l'\} &= J_1' \sqcup \cdots \sqcup J'_{\kappa_\theta(S)}, \end{aligned} \tag{1.11}$$

such that for each $\nu = 1, \ldots, \kappa_\theta(S)$, the sets I_ν, I_ν', J_ν, and J_ν' are not all empty, and

$$S_\nu := ((a_i \mid i \in I_\nu); (a_i' \mid i \in I_\nu') \mid (b_j \mid j \in J_\nu); (b_j' \mid j \in J_\nu')) \tag{1.12}$$

satisfies the sum condition (1.9). Note that since κ_θ is always positive, we must have $\kappa_\theta(S_\nu) = 1$ for each ν. Declare two θ-decompositions to be equivalent iff they differ by applying a permutation of the set $\{1, \ldots, \kappa_\theta(S)\}$ to the indexing on the right hand side of (1.11). We sometimes abuse notation and denote a θ-decomposition by $\{S_\nu\}$.

Lemma 1.17. *With the notation of* (1.8)*, a θ-decomposition of S exists if and only if* $(a_1, \ldots, a'_{k'}) \geq_\theta (b_1, \ldots, b'_{l'})$.

Proof. This follows directly from the above discussion. □

Definition 1.18. For any S as in (1.8) satisfying the sum condition (1.9), define

$$c_\theta(S) := \sum_{\substack{\text{equivalence classes of} \\ \theta\text{-decompositions of } S}} \prod_{\nu=1}^{\kappa_\theta(S)} c_\theta(S_\nu).$$

To complete this definition, the rest of this subsection defines $c_\theta(S)$ when $\kappa_\theta(S) = 1$. We first need to define some auxiliary functions.

Notation 1.19. If a and b are positive integers, define a positive integer

$$\delta_\theta(a,b) := b\lceil a\theta\rceil - a\lfloor b\theta\rfloor.$$

Definition 1.20. Given *ordered* lists of positive integers $a_1,\ldots,a_k$ and $b_1,\ldots,b_l$ with the same sum $\sum_{i=1}^k a_i = \sum_{j=1}^l b_j$, define a positive integer $f_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l)$ recursively as follows. To start the recursion, if $k=l=0$, then $f_\theta(|) := 1$.

For $k\ge 1$, the recursion involves a sum over subsets

$$I = \{i_1 < \cdots < i_q\} \subset \{1,\ldots,l\}$$

such that

$$\sum_{j=1}^{q-1} b_{i_j} < a_1 \le \sum_{j=1}^{q} b_{i_j}. \tag{1.13}$$

We also require that equality holds in (1.13) only when $k=1$. (This requirement is automatically satisfied in the case of interest where $\kappa_\theta = 1$.) The formula is now

$$f_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l) := \sum_I f_\theta(a_2,\ldots,a_k \mid b_I)\prod_{n=1}^{q}\delta_\theta\left(a_1 - \sum_{j=1}^{n-1} b_{i_j},\, b_{i_n}\right). \tag{1.14}$$

Here b_I denotes the arguments b_i for $i\notin I$, arranged in order, together with (when $k>1$) one additional argument equal to $\sum_{i=2}^k a_i - \sum_{i\notin I} b_i$, inserted in the position that b_{i_q} would occupy in the order.

Remark 1.21. If $\kappa_\theta = 1$, then f_θ is always a positive integer, because the sum (1.14) always has at least one term. This follows by induction, since one can find a subset I satisfying (1.13) by just taking $I=\{1,\ldots,q\}$, where q is the smallest integer such that $\sum_{j=1}^q b_j \ge a_1$.

The definition of $c_\theta(S)$ when $\kappa_\theta(S)=1$ is now divided into several cases depending on the value of $k'+l'$.

Definition 1.22. If $\kappa_\theta(S)=1$ and $k'=l'=0$, then $c_\theta(S)$ is defined as follows. Choose a reordering of the a_i's and b_j's so that

$$\frac{\lceil a_i\theta\rceil}{a_i} \le \frac{\lceil a_{i+1}\theta\rceil}{a_{i+1}}, \qquad \frac{\lfloor b_j\theta\rfloor}{b_j} \ge \frac{\lfloor b_{j+1}\theta\rfloor}{b_{j+1}}. \tag{1.15}$$

Then

$$c_\theta(S) := f_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l). \tag{1.16}$$

Remark 1.23. It is not obvious from Definition 1.22 that $c_\theta(S)$ is independent of the choice of reordering satisfying (1.15). This fact follows from the analysis used to prove Theorem 1.13, in particular Corollary 3.6 and Proposition 5.1. It can also be proved combinatorially, as described in Remark 4.7.

Another nonobvious property which follows from the analysis is the symmetry

$$c_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l) = c_{-\theta}(b_1, \ldots, b_l \mid a_1, \ldots, a_k). \tag{1.17}$$

We will give a combinatorial proof of this in §4.

Remark 1.24. The geometric significance of the condition (1.15) is that if a J-holomorphic curve u has ordered negative ends of multiplicities $a_1, \ldots, a_k$ at a given elliptic Reeb orbit of monodromy angle θ, then generically, the i^{th} negative end of u decays no faster than the $(i+1)^{st}$ negative end. Likewise, if u has ordered positive ends of multiplicities $b_1, \ldots, b_l$ at this Reeb orbit, then generically, the j^{th} positive end decays no faster than the $(j+1)^{st}$ positive end.

Definition 1.25. If $\kappa_\theta(S) = 1$ and $k' + l' = 1$, re-order the a_i's and b_j's in accordance with (1.15). If $k' = 1$, define

$$c_\theta(S) := f_\theta(a_1, \ldots, a_k, a'_1 \mid b_1, \ldots, b_l).$$

Likewise, if $l' = 1$ define

$$c_\theta(S) := f_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l, b'_1).$$

Definition 1.26. If $\kappa_\theta(S) = 1$ and $k' + l' \geq 2$, define $c_\theta(S) := 0$ unless $k = l = 0$ and $k' = l' = 1$, in which case define

$$c_\theta(;a'|;a') := a'.$$

1.7 Examples and applications

The following are some important examples of elliptic gluing coefficients $c_\theta(S)$ where $k' = l' = 0$.

Example 1.27. If $k = l = 1$, then $c_\theta(a \mid a) = a$. In the gluing theorem this corresponds to a situation with no branch points, and reflects the fact that there are a ways to match up the sheets of a negative end of u_+ and a positive end of u_- along an a-fold cover of Reeb orbit.

Example 1.28. Suppose $k = 1$ and $\theta \in (0, 1/a_1)$. (Elliptic orbits with θ close to zero arise when λ is a perturbation of a Morse-Bott contact form.)

Then $\kappa_\theta(a_1 \mid b_1, \ldots, b_l) = 1$, the sum in (1.14) has only one term, and we find that

$$c_\theta(a_1 \mid b_1, \ldots, b_l) = \prod_{j=1}^{l} b_j.$$

Example 1.29. In §7 we will use Theorem 1.13 to prove that the differential ∂ in embedded contact homology (ECH) satisfies $\partial^2 = 0$. For this purpose, one has to calculate c_θ in a certain special (but nontrivial) case as follows. Given an irrational number θ, for each nonnegative integer M there are two distinguished partitions of M, called the "incoming partition" and the "outgoing partition", and denoted here by $P_\theta^{\mathrm{in}}(M)$ and $P_\theta^{\mathrm{out}}(M)$ respectively, see §7.1. If u is a J-holomorphic curve contributing to the ECH differential, then the multiplicities of the negative (resp. positive) ends of u at each elliptic Reeb orbit are determined by the corresponding incoming (resp. outgoing) partitions. Thus a key part of the proof that $\partial^2 = 0$ is to show that if

$$P_\theta^{\mathrm{in}}(M) = (a_1, \ldots, a_k), \qquad P_\theta^{\mathrm{out}}(M) = (b_1, \ldots, b_l), \tag{1.18}$$

then

$$c_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l) = 1. \tag{1.19}$$

We will do so in Proposition 7.26. In fact, similar arguments show that the converse is also true, i.e. (1.19) implies (1.18). Thus the proof here that $\partial^2 = 0$ is quite delicate.

Remark 1.30. Symplectic field theory [5] defines a differential D on a supercommutative algebra over $\mathbb{Q}$ generated by all "good" (not necessarily embedded) Reeb orbits. The differential D counts points in some abstract perturbations of the compactified moduli spaces of index 1 J-holomorphic curves in $\mathbb{R}\times Y$. Even if one knows all J-holomorphic curves, it is a nontrivial problem to read off the differential D. Theorem 1.13 gives some constraints on the answer in our three-dimensional case (SFT is defined for contact manifolds of any dimension). To give the simplest example, let γ_1 be an embedded elliptic Reeb orbit with monodromy angle $\theta \in (0, 1/2)$, and let γ_2 denote the double cover of γ_1. Suppose that the only index 1 J-holomorphic curves modulo translation with ends at γ_1 or γ_2 are a curve $u_+ \in \mathcal{M}^J((\gamma_+), (\gamma_2))$ and a curve $u_- \in \mathcal{M}^J((\gamma_1, \gamma_1)), (\gamma_-))$, say with $\epsilon(u_\pm) = 1$. There is some contribution $D_+ \in \mathbb{Q}$ to the differential coefficient $\langle D\gamma_+, \gamma_1^2\rangle$ arising from broken curves consisting of u_+ together with a branched double cover of $\mathbb{R} \times \gamma_1$. Likewise there is some contribution $D_- \in \mathbb{Q}$ to the differential coefficient $\langle D\gamma_2, \gamma_-\rangle$ arising from broken curves consisting of a branched double cover of $\mathbb{R} \times \gamma_1$ together with u_-. Presumably D_+ and D_- may depend on the choice of abstract perturbations. However, since $D^2 = 0$, Theorem 1.13 requires that

$$D_+ + D_- = c_\theta(2 \mid 1, 1) = 1.$$

1.8 Overview of the proof of the main theorem

We now describe the proof of Theorem 1.13. To simplify notation, we will restrict attention to the special case where conditions (i) and (ii) below hold:

(i) There is an embedded elliptic Reeb orbit α such that all negative ends of u_+ and all positive ends of u_- are at covers of α.

Let θ denote the monodromy angle of α. Suppose that u_+ has negative ends at $\alpha^{a_1}, \ldots, \alpha^{a_{N_+}}$, and u_- has positive ends at $\alpha^{a_{-1}}, \ldots, \alpha^{a_{-N_-}}$. The second condition is then:

(ii) $\kappa_\theta(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-}) = 1$.

The strategy for gluing u_+ and u_- is as follows. Fix large constants $R >> r >> 0$. Let Σ be a connected genus zero branched cover of $\mathbb{R} \times \alpha$ which has positive ends of covering multiplicities $a_1, \ldots, a_{N_+}$ and negative ends of covering multiplicities $a_{-1}, \ldots, a_{-N_+}$, such that all ramification points have $|s| \leq R$. Form a "preglued" curve by using appropriate cutoff functions to patch the negative ends of the $s \mapsto s + R + r$ translate of u_+ to the positive ends of Σ, and the positive ends of the $s \mapsto s - R - r$ translate of u_- to the negative ends of Σ. Now try to perturb the preglued curve to a J-holomorphic curve, where near the ramification points of the branched cover we only perturb in the directions normal to $\mathbb{R} \times \alpha$.

It turns out that there is an "obstruction bundle" $\mathcal{O}$ over the moduli space $\mathcal{M}_R$ of branched covers Σ as above, and a section $\mathfrak{s} : \mathcal{M}_R \to \mathcal{O}$ of this bundle, such that the preglued curve determined by $\Sigma \in \mathcal{M}_R$ can be perturbed as above to a J-holomorphic curve if and only if $\mathfrak{s}(\Sigma) = 0$. In this way we will identify the count $\#G(u_+, u_-)$ in Theorem 1.13 with $\epsilon(u_+)\epsilon(u_-)$ times an appropriate count of the zeroes of $\mathfrak{s}$. The section $\mathfrak{s}$ is defined rather indirectly from the analysis, but there is an approximation $\mathfrak{s}_0$ to $\mathfrak{s}$ which is given by an explicit formula. We will see that if J is generic, then the sections $\mathfrak{s}_0$ and $\mathfrak{s}$ have the same count of zeroes, because one can deform $\mathfrak{s}$ to $\mathfrak{s}_0$ without any zeroes crossing the boundary of the moduli space $\mathcal{M}_R$. We will then use a detailed analysis of the obstruction bundle to count the zeroes of $\mathfrak{s}_0$ and recover the combinatorial gluing coefficient $c_\theta(u_+, u_-)$.

Without conditions (i) and (ii) above, one also needs to keep track of the different Reeb orbits where gluing takes place, and also to consider disconnected branched covers of $\mathbb{R}$ cross an elliptic Reeb orbit. Since this does not involve any additional analysis, in an attempt to keep the notation manageable we will continue to assume (i) and (ii) below and in [13].

The harder analytic parts of the above proof are carried out in the sequel [13]. The present paper explains the more algebraic aspects and is organized as follows. In §2 we define the obstruction bundle over the moduli space of branched covers and discuss its basic properties. In §3 we define the section

$\mathfrak{s}_0$ of the obstruction bundle and quote results from [13] relating $\#G(u_+, u_-)$ to an appropriate count of zeroes of $\mathfrak{s}_0$. §4 is almost completely independent of the previous two sections, and discusses the combinatorics of the gluing coefficients c_θ in detail. §5 brings the analysis and the combinatorics together to count the zeroes of $\mathfrak{s}_0$, thereby completing the proof of Theorem 1.13 modulo the aforementioned results from [13]. §6 then ties up loose ends by proving some estimates on the obstruction bundle which were used in §2 and §5. Finally, §7 explains the application to embedded contact homology; this section is independent of §2–§6.

For some other obstruction bundle calculations in symplectic field theory, concerning index 1 branched covers of $\mathbb{R}$-invariant cylinders, see [6].

2 The obstruction bundle

As described in §1.8, the number of gluings in Theorem 1.13 is determined by counting zeroes of a certain section of an "obstruction bundle" $\mathcal{O}$ over a moduli space $\mathcal{M}$ of genus zero branched covers of $\mathbb{R} \times S^1$. We now define this bundle and discuss its basic properties. In §2.1 we define the moduli space $\mathcal{M}$, and in §2.2 we review the asympotic operator associated to a Reeb orbit. We then define the bundle $\mathcal{O}$ in §2.3; the fiber of $\mathcal{O}$ over a branched cover Σ is the dual of the cokernel of a certain operator D_Σ. In §2.4 we introduce some special elements of $\operatorname{Coker}(D_\Sigma)$; later we will study the section of $\mathcal{O}$ by evaluating it on these. In §2.5 we give some estimates on the behavior of a special cokernel element in terms of the combinatorics of the branched cover on which it is defined. Finally, §2.6 defines an orientation of $\mathcal{O}$, and §2.7 defines a useful compactification of $\mathcal{M}/\mathbb{R}$.

In this section fix positive integers $a_1, \ldots, a_{N_+}$ and $a_{-1}, \ldots, a_{-N_-}$ with

$$\sum_{i=1}^{N_+} a_i = \sum_{j=-1}^{-N_-} a_j = M. \tag{2.1}$$

Write $N := N_+ + N_-$, and to avoid trivialities assume that $N > 2$. Also, fix an irrational number θ, and assume that

$$\kappa_\theta(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-}) = 1. \tag{2.2}$$

Finally, fix an admissible almost complex structure J on $\mathbb{R} \times Y$ and an embedded Reeb orbit α. In §2.3–§2.6 we assume that α is elliptic with monodromy angle θ.

2.1 Branched covers and trees

The following basic definitions will be used throughout the paper.

Definition 2.1. Let $\mathcal{M} = \mathcal{M}(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ denote the moduli space of degree M branched covers $\pi : \Sigma \to \mathbb{R} \times S^1$ such that:

- Σ is connected and has genus zero.
- The positive ends of Σ are labeled by $1, \ldots, N_+$, and the negative ends of Σ are labeled by $-1, \ldots, -N_-$.
- The end of Σ labeled by i has covering multiplicity a_i.
- The ends are asymptotically marked. That is, an identification is chosen between the i^{th} positive end of Σ and $[R, \infty) \times (\mathbb{R}/2\pi a_i \mathbb{Z})$, respecting the projection to $\mathbb{R} \times (\mathbb{R}/2\pi\mathbb{Z})$. Likewise for the negative ends.

We declare $\pi : \Sigma \to \mathbb{R} \times S^1$ to be equivalent to $\pi' : \Sigma' \to \mathbb{R} \times S^1$ if there is a diffeomorphism $\phi : \Sigma \xrightarrow{\simeq} \Sigma'$ such that $\pi' \circ \phi = \pi$, and ϕ respects the labelings and asymptotic markings of the ends. We often abuse notation and denote an element of $\mathcal{M}$ by Σ.

Note that $\mathcal{M}$ is a finite-sheeted covering space of the space of meromorphic functions on $\mathbb{CP}^1$ with poles of order $a_1, \ldots, a_{N_+}$ and zeroes of order $a_{-1}, \ldots, a_{-N_-}$, modulo automorphisms of $\mathbb{CP}^1$. In particular, $\mathcal{M}$ is a complex manifold of dimension

$$\dim_{\mathbb{C}}(\mathcal{M}) = N - 2. \tag{2.3}$$

We now explain how to associate, to each branched cover $\Sigma \in \mathcal{M}$, a tree with certain additional structure. In this paper, a *tree* is a finite, connected, simply connected graph T, such that each vertex has degree either one (a *leaf*) or at least three (an *internal vertex*). We denote the set of internal vertices by $\dot{V}(T)$. The tree T is *trivalent* if every internal vertex has degree three.

For any two vertices v and w in a tree, let $P_{v,w}$ denote the unique (non-backtracking combinatorial) path from v to w. Given three distinct leaves i, j, and k, the triple intersection of the paths $P_{i,j}$, $P_{i,k}$, and $P_{j,k}$ consists of a single vertex, which we call the *central vertex* for i, j, and k.

Definition 2.2. An *oriented weighted tree* is a tree T such that:

- Each edge e has an orientation $\mathfrak{o}(e)$ and a positive integer weight $m(e)$, which we call the "multiplicity" of e.
- For each internal vertex, the sum of the multiplicities of the outgoing edges equals the sum of the multiplicities of the incoming edges.

In an oriented tree, we call a leaf *positive* if the incident edge points towards the leaf, and *negative* otherwise. An "upward" path will mean a positively oriented path, and a "downward" path will mean a negatively oriented path. A vertex v is a *splitting vertex* if it has at least two outgoing edges, and a *joining vertex* if it has at least two incoming edges.

Definition 2.3. Let $T(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ denote the set of oriented weighted trees such that:

- The positive leaves are labeled by $1, \ldots, N_+$, and the negative leaves are labeled by $-1, \ldots, -N_-$.
- The (edge incident to the) i^{th} leaf has multiplicity a_i.

Definition 2.4. Let $\mathcal{T} = \mathcal{T}(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ denote the set of oriented weighted trees $T \in T(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ such that:

- Each internal vertex v is labeled by a real number $\rho(v)$.
- If v and w are internal vertices, and if there is an oriented edge from v to w, then $\rho(v) < \rho(w)$.

Definition 2.5. Define a map

$$\tau : \mathcal{M} \longrightarrow \mathcal{T}$$

as follows. Given a branched cover $\pi : \Sigma \to \mathbb{R} \times S^1$ in $\mathcal{M}$, let ρ denote the composition $\Sigma \to \mathbb{R} \times S^1 \to \mathbb{R}$. Define two points in Σ to be equivalent if they are connected by a path on which ρ is constant. The quotient space of Σ by this equivalence relation is a one-dimensional CW complex $\tau(\Sigma)$, which is a tree with a continuous map

$$\rho : \tau(\Sigma) \longrightarrow \mathbb{R}. \tag{2.4}$$

In the tree $\tau(\Sigma)$, a vertex v of degree $d \geq 3$ corresponds to an equivalence class $R(v)$ in Σ containing ramification points with total ramification index $d - 2$. The complement $\Sigma \setminus \cup_v R(v)$ is a collection of cylinders, which correspond to the edges of $\tau(\Sigma)$. We orient the edges via the direction in which ρ increases, and define the multiplicity of an edge to be the covering multiplicity of the corresponding cylinder in Σ. An example is shown in Figure 1.

We next consider the extent to which the tree $T := \tau(\Sigma)$ determines the branched cover Σ. In the "generic" case when T is trivalent, each internal vertex corresponds to a unique ramification point in Σ, so there is a well-defined map

$$\phi_T : \tau^{-1}(T) \longrightarrow (S^1)^{\dot{V}(T)} \tag{2.5}$$

which sends a branched cover Σ to the S^1-coordinates of π of the ramification points. Let $E(T)$ denote the set of edges of T.

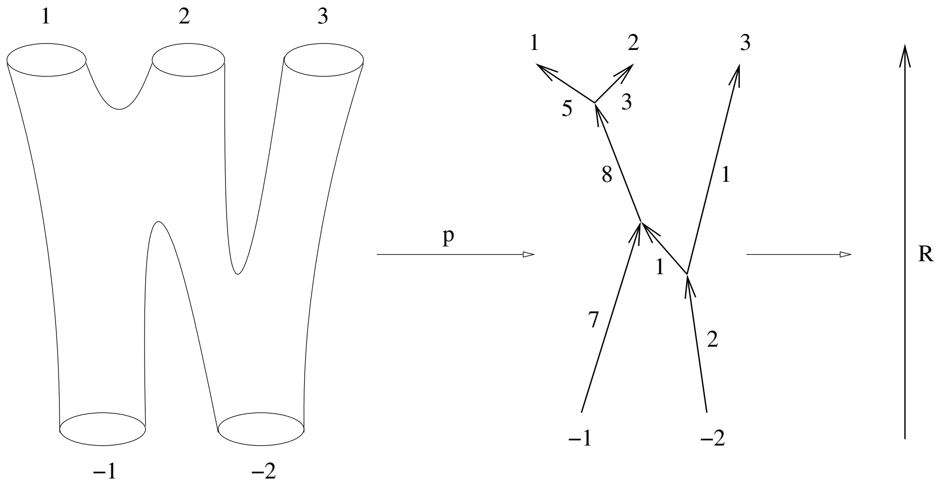

Figure 1: A branched cover $\Sigma \in \mathcal{M}(5,3,1 \mid 7,2)$, its associated oriented weighted tree $\tau(\Sigma)$, and their projections to $\mathbb{R}$. The end labels and edge weights are shown.

Lemma 2.6. *If T is trivalent, then the map ϕ_T in (2.5) is a covering of degree*

$$\deg(\phi_T) = \prod_{e \in E(T)} m(e). \tag{2.6}$$

Proof. Given a trivalent tree T and an element of S^1 for each vertex, a corresponding branched cover Σ is obtained by taking a pair of pants for each vertex and gluing them together as dictated by the internal edges of T. For each internal edge e there are $m(e)$ possible gluings, and for each external edge e there are $m(e)$ possible asymptotic markings of the corresponding end of Σ. □

The following subset of $\mathcal{M}$ will play a crucial role.

Definition 2.7. Given $R > 0$, define

$$\mathcal{M}_R = \mathcal{M}_R(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$$

to be the set of $\pi : \Sigma \to \mathbb{R} \times S^1$ in $\mathcal{M}$ such that if $x \in \Sigma$ is a ramification point and $\pi(x) = (s,t)$, then $|s| \leq R$. Let $\partial\mathcal{M}_R$ denote the set of $\Sigma \in \mathcal{M}_R$ having a ramification point with $|s| = R$.

Lemma 2.8. *$\mathcal{M}_R$ is compact*[1]*.*

[1]Note that we will use the assumption (2.2) here. Thanks to A. Cotton-Clay for pointing out a mistake in this regard in an earlier draft of this paper.

Proof. Let X denote the symmetric product $\mathrm{Sym}^{N-2}([-R,R]\times S^1)$. Consider the map $\phi:\mathcal{M}_R\to X$ which sends a branched cover $\pi:\Sigma\to\mathbb{R}\times S^1$ in $\mathcal{M}_R$ to the set of π-images of the ramification points in Σ, repeated according to their ramification indices. Note that the symmetric product X is compact, the map ϕ is continuous, and each point in X has only finitely many inverse images under ϕ. Furthermore, ϕ defines a covering space over each stratum in the symmetric product. Thus to prove that $\mathcal{M}_R$ is compact, it is enough to show that if $\eta:[0,1]\to X$ is a path that maps all of $[0,1)$ to the same stratum, then η has a lift to $\mathcal{M}_R$ starting at any given $\pi\in\phi^{-1}(\eta(0))$. The issue is to check that whenever two branch points in $[-R,R]\times S^1$ collide, the corresponding ramification points in Σ either do not interact or can be merged. More precisely, it is enough to show that if $x_0,x_1\in\Sigma$ are two distinct ramification points, and if γ is an embedded path in $[-R,R]\times S^1$ from $\pi(x_0)$ to $\pi(x_1)$, then γ has at most one lift to a path in Σ.

Suppose to the contrary that γ has two distinct lifts $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$. We can then make a new branched cover $\pi':\Sigma'\to\mathbb{R}\times S^1$ by cutting Σ along the paths $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$, and gluing each side of $\widetilde{\gamma}_1$ to the opposite side of $\widetilde{\gamma}_2$. This operation reduces the total ramification index by 2. Since the original branched cover Σ had genus zero, it follows by Riemann-Hurwitz that the new branched cover Σ' is disconnected. On the other hand, Σ' still has positive ends of multiplicities $a_1,\ldots,a_{N_+}$, and negative ends of multiplicities $a_{-1},\ldots,a_{-N_-}$. Hence there are decompositions $\{1,\ldots,N_+\}=I_1\sqcup I_2$ and $\{-1,\ldots,-N_-\}=J_1\sqcup J_2$ into proper subsets such that $\sum_{i\in I_1}a_i=\sum_{j\in J_1}a_j$. It follows that $\kappa_\theta(a_1,\ldots,a_{N_+}\mid a_{-1},\ldots,a_{-N_-})\geq 2$, contradicting the assumption (2.2). $\square$

2.2 The asymptotic operator

We now review the asymptotic operator associated to a Reeb orbit. This operator plays a fundamental role in the analysis.

Recall that we are fixing an embedded Reeb orbit α. By rescaling the t coordinate, we may assume that α is parametrized by $S^1:=\mathbb{R}/2\pi\mathbb{Z}$. The linearized Reeb flow on the contact planes along α defines a symplectic connection ∇^R on the 2-plane bundle $\alpha^*\xi$ over S^1.

Definition 2.9. Define the *asymptotic operator*

$$L:=L_\alpha:=J\nabla_t^R$$

acting on sections of $\alpha^*\xi$ over S^1. More generally, if m is a positive integer, let $L_m:=L_{\alpha^m}$ denote the pullback of L to $\mathbb{R}/2\pi m\mathbb{Z}$.

To describe the operator L_m more explicitly, fix a complex linear, symplectic trivialization of $\alpha^*\xi$. For $t\in\mathbb{R}$, let $\Psi(t):\mathbb{R}^2\to\mathbb{R}^2$ denote the linearized Reeb flow with respect to this trivialization, as the S^1 coordinate increases from 0 to t. Let J_0 denote the standard complex structure on $\mathbb{R}^2$. For $t \in S^1$, define a matrix $S(t)$ by writing the derivative of the linearized Reeb flow as

$$\frac{d\Psi(t)}{dt}\Psi(t)^{-1} =: J_0 S(t). \tag{2.7}$$

Then in the above trivialization,

$$L_m = J_0 \frac{d}{dt} + S(t)$$

acting on complex functions on $\mathbb{R}/2\pi m\mathbb{Z}$. Since the connection ∇^R on $\alpha^*\xi$ is symplectic, it follows from (2.7) that the matrix $J_0 S(t)$ is in the Lie algebra of the symplectic group, which means that the matrix $S(t)$ is symmetric. In particular, the operator L_m is self-adjoint.

Our standing assumption that all Reeb orbits are nondegenerate implies that $0 \notin \operatorname{Spec}(L_m)$. The reason is that if $L_m\gamma = 0$, then it follows from equation (2.7) that $\gamma(t) = \Psi(t)\gamma(0)$. Thus $\gamma(0)$ is an eigenvector of $\Psi(2\pi m)$ with eigenvalue 1, and if $\gamma \neq 0$ this contradicts the nondegeneracy of α^m.

To describe more spectral properties of the operator L_m, note that if γ is an eigenfunction with eigenvalue λ, then γ solves the ODE

$$\frac{d\gamma(t)}{dt} = J_0(S(t) - \lambda)\gamma(t). \tag{2.8}$$

It follows from (2.8) and the uniqueness of solutions to ODE's that if γ is nonzero, then it is nonvanishing. Then the loop $\gamma : \mathbb{R}/2\pi m\mathbb{Z} \to \mathbb{C}$ has a well-defined winding number around 0. We denote this winding number by $\eta(\gamma) \in \mathbb{Z}$.

Example 2.10. An important special case is where $S(t) = \theta$ for all t. In this case $\Phi(t) = e^{i\theta t}$, the operator L_m is complex linear, and the eigenfunctions are complex multiples of the functions $\gamma(t) = e^{i\eta t/m}$ for $\eta \in \mathbb{Z}$. Such an eigenfunction has eigenvalue $\theta - \eta/m$ and winding number η.

It follows by analytic perturbation theory that in the general case, eigenvalues are related to winding numbers as follows; for details see [9, §3].

Lemma 2.11. *(a) For each integer η, the sum of the eigenspaces whose nonzero eigenfunctions have winding number η is 2-dimensional.*

(b) If γ and γ' are eigenfunctions of L_m with eigenvalues $\lambda \leq \lambda'$, then $\eta(\gamma) \geq \eta(\gamma')$.

(c) Suppose the Reeb orbit α is elliptic with monodromy angle θ. If γ is an eigenfunction of L_m with eigenvalue λ, then

$$\lambda > 0 \Longleftrightarrow \eta(\lambda) < m\theta.$$

If η and m are integers with $m > 0$, let $E^-_{\eta/m} \le E^+_{\eta/m}$ denote the two eigenvalues of L_m, given by Lemma 2.11(a), whose associated eigenfunctions have winding number η.

Remark 2.12. These eigenvalues depend only on the rational number η/m, because if γ is an eigenfunction of L_m with winding number η, and if d is a positive integer, then the covering $\mathbb{R}/2\pi dm\mathbb{Z} \to \mathbb{R}/2\pi m\mathbb{Z}$ pulls back γ to an eigenfunction of L_{dm} with the same eigenvalue and with winding number $d\eta$. Also

$$\frac{\eta_1}{m_1} < \frac{\eta_2}{m_2} \Longrightarrow E^-_{\eta_1/m_1} > E^+_{\eta_2/m_2};$$

this follows by pulling back eigenfunctions to $\mathbb{R}/2\pi m_1 m_2\mathbb{Z}$ and applying Lemma 2.11(b). By Lemma 2.11(c), if α is elliptic with monodromy angle θ, then the largest negative eigenvalue of L_m is $E^+_{\lceil m\theta\rceil/m}$, while the smallest positive eigenvalue of L_m is $E^-_{\lfloor m\theta\rfloor/m}$.

2.3 The operator D_Σ and the obstruction bundle

We now introduce an operator D_Σ, which arises in connection with deformations of J-holomorphic curves given by branched covers of $\mathbb{R}\times\alpha$ in directions normal to $\mathbb{R}\times\alpha$.

For the analysis to follow, fix a Hermitian metric on each $\Sigma \in \mathcal{M}$ which varies smoothly over $\mathcal{M}$, which agrees with the pullback of the standard metric on $\mathbb{R}\times S^1$ at points in Σ with distance ≥ 1 from the ramification points, and which within distance 1 of a ramification point depends only on the local structure of the branched cover within distance $2N$. (Here the "distance" between two points $x, y \in \Sigma$ is defined to be the infimum, over all paths P in Σ from x to y, of the length of the projection of P to $\mathbb{R}\times S^1$.)

Now let $\pi : \Sigma \to \mathbb{R}\times S^1$ be a branched cover in $\mathcal{M}$. Let (s,t) denote the usual coordinates on $\mathbb{R}\times S^1$, and write $z := s + it$. Recall the notation L_m and $S(t)$ from §2.2.

Definition 2.13. Define a real linear operator

$$\begin{aligned} D_\Sigma : L^2_1(\Sigma,\mathbb{C}) &\longrightarrow L^2(T^{0,1}\Sigma), \\ D_\Sigma &:= \overline{\partial} + \frac{1}{2}\pi^*(S(t)d\overline{z}). \end{aligned} \tag{2.9}$$

Note that over an end of Σ of multiplicity m, identified with $[R,\infty)\times\mathbb{R}/2\pi m\mathbb{Z}$ or $(-\infty,-R]\times\mathbb{R}/2\pi m\mathbb{Z}$, we have

$$D_\Sigma f = (\partial_s + L_m)f \otimes \frac{d\overline{z}}{2}.$$

Since $0 \notin \operatorname{Spec}(L_m)$, it follows by standard arguments that the operator D_Σ is Fredholm.

Assume henceforth that the Reeb orbit α is elliptic with monodromy angle θ. Recall that θ is assumed to satisfy (2.2). The index of D_Σ is then given as follows:

Lemma 2.14. $\operatorname{ind}(D_\Sigma) = -\dim_{\mathbb{R}}(\mathcal{M})$.

Proof. By a standard index formula (cf. [16]),

$$\operatorname{ind}(D_\Sigma) = \chi(\Sigma) + \sum_{i=1}^{N_+} \mu(a_i) - \sum_{j=-1}^{-N_-} \mu(a_j).$$

Here $\mu(m)$, for a positive integer m, denotes the Maslov index of the path of symplectic matrices $\{\Psi(t) \mid t \in [0, 2\pi m]\}$, which is given explicitly by

$$\mu(m) = 2\lfloor m\theta \rfloor + 1.$$

Of course,

$$\chi(\Sigma) = 2 - N.$$

The lemma follows directly from the above three equations, together with (2.2) and (2.3). □

We now consider the cokernel of D_Σ. We can identify $\operatorname{Coker}(D_\Sigma)$ with the space of smooth $(0,1)$-forms σ on Σ that are in L^2 and annihilated by the formal adjoint D_Σ^* of D_Σ.

A nonzero cokernel element σ has the following asymptotic behavior. Over the i^{th} positive end of Σ, which the asymptotic marking identifies with $[R, \infty) \times \mathbb{R}/2\pi a_i \mathbb{Z}$, write $\sigma = \sigma_i(s,t) d\bar{z}$. The function σ_i satisfies the equation

$$(\partial_s - L_{a_i})\sigma_i = 0.$$

Since σ is in L^2, we can expand

$$\sigma_i(s,t) = \sum_{\lambda < 0} e^{\lambda s} \gamma_{i,\lambda}(t) \tag{2.10}$$

where the sum is over negative eigenvalues λ of L_{a_i}, and $\gamma_{i,\lambda}$ is a (possibly zero) eigenfunction with eigenvalue λ. Let λ_i denote the largest negative eigenvalue for which γ_{i,λ_i} is nonzero, and write $\gamma_i := \gamma_{i,\lambda_i}$. If $\kappa > 0$ is the difference between λ_i and the second largest negative eigenvalue, then by (2.10), there is an s-independent constant A such that

$$\left| \sigma_i(s,t) - e^{\lambda_i s} \gamma_i(t) \right| \le A e^{(\lambda_i - \kappa)s}.$$

It follows that when s is large, σ_i has no zeroes, and has winding number $\eta(\gamma_i)$ around the i^{th} positive end of Σ. We denote this winding number by $\eta_i^+(\sigma)$. Likewise, for $j \in \{-1, \ldots, -N_-\}$, one can expand σ on the j^{th} negative end of Σ as

$$\sigma_j(s,t) = \sum_{\lambda>0} e^{\lambda s}\gamma_{j,\lambda}(t), \tag{2.11}$$

and σ has a well-defined winding number $\eta_j^-(\sigma)$ around the j^{th} negative end. Lemma 2.11(c) then gives the winding bounds

$$\eta_i^+(\sigma) \geq \lceil a_i\theta \rceil, \qquad \eta_j^-(\sigma) \leq \lfloor a_j\theta \rfloor. \tag{2.12}$$

Given a nonzero cokernel element σ, let $Z(\sigma)$ denote the number of ends of Σ for which the inequalities (2.12) are strict.

Lemma 2.15. *(a) If $0 \neq \sigma \in \operatorname{Coker}(D_\Sigma)$, then the zeroes of σ are isolated and have negative multiplicity, and the algebraic count of zeroes is bounded by*

$$\#\sigma^{-1}(0) \geq \frac{\operatorname{ind}(D_\Sigma)}{2} + 1 + Z(\sigma).$$

(b) $\dim(\operatorname{Coker}(D_\Sigma)) = -\operatorname{ind}(D_\Sigma)$*, or equivalently* $\operatorname{Ker}(D_\Sigma) = \{0\}$.

Proof. (a) Since D_Σ^* is $\overline{\partial}^*$ plus a zeroth order term, the zeroes of σ are isolated and have negative multiplicity. For any $(0,1)$-form σ with finitely many zeroes, the algebraic count of zeroes is given by

$$\#\sigma^{-1}(0) = \chi(\Sigma) + \sum_{i=1}^{N_+} \eta_i^+(\sigma) - \sum_{j=-1}^{-N_-} \eta_j^-(\sigma). \tag{2.13}$$

If σ is a nonzero cokernel element, then putting the winding bounds (2.12) into (2.13) and using Lemma 2.14 proves part (a).

(b) If $\dim(\operatorname{Coker}(D_\Sigma)) > -\operatorname{ind}(D_\Sigma)$, then one can find a nonzero cokernel element σ with zeroes at $-\operatorname{ind}(D_\Sigma)/2$ given points in Σ. Since all zeroes of σ have negative multiplicity, this contradicts part (a). $\square$

Lemma 2.15(b) implies that the cokernels of the operators D_Σ for $\Sigma \in \mathcal{M}$ comprise a smooth real vector bundle over $\mathcal{M}$, which we denote by $\mathcal{O}^*$.

Definition 2.16. Define the *obstruction bundle* $\mathcal{O} \to \mathcal{M}$ to be the dual of the bundle of cokernels $\mathcal{O}^* \to \mathcal{M}$. Thus the fiber of $\mathcal{O}$ over Σ is

$$\mathcal{O}_\Sigma = \operatorname{Hom}(\operatorname{Coker}(D_\Sigma), \mathbb{R}).$$

By Lemmas 2.14 and 2.15(b), the rank of $\mathcal{O}$ equals the dimension of $\mathcal{M}$.

2.4 Special cokernel elements

We now introduce some special elements of $\mathrm{Coker}(D_\Sigma)$ which will play a key role in the obstruction bundle calculations in §5.

To define the special cokernel elements, for $\sigma \in \mathrm{Coker}(D_\Sigma)$ we need to consider the "leading terms" of the asymptotic expansions (2.10) and (2.11). Namely, to the i^{th} end of Σ we associate a real vector space $\mathcal{A}_i$ as follows. For $i \in \{1, \ldots, N_+\}$, define $\mathcal{A}_i$ to be the direct sum of those eigenspaces of L_{a_i} whose nonzero elements have winding number $\lceil a_i\theta \rceil$. By Lemma 2.11, $\mathcal{A}_i$ is two-dimensional, and consists of the eigenspace for the largest negative eigenvalue of L_{a_i}, together with the eigenspace for the second largest negative eigenvalue when the largest negative eigenvalue has multiplicity one. Likewise, for $j \in \{-1, \ldots, -N_-\}$, define $\mathcal{A}_j$ to be the direct sum of those eigenspaces of L_{a_j} whose nonzero elements have winding number $\lfloor a_j\theta \rfloor$.

Given $\sigma \in \mathrm{Coker}(D_\Sigma)$, for $i \in \{1, \ldots, N_+\}$, define $\Phi_i(\sigma) \in \mathcal{A}_i$ to be the sum of those eigenfunctions $\gamma_{i,\lambda}$ in the expansion (2.10) that have winding number $\lceil a_i\theta \rceil$. Likewise, for $j \in \{-1, \ldots, -N_-\}$, define $\Phi_j(\sigma) \in \mathcal{A}_j$ to be the sum of those eigenfunctions in the expansion (2.11) that have winding number $\lfloor a_j\theta \rfloor$.

We can now define the special cokernel elements:

Definition 2.17. If i, j, and k label distinct ends of Σ, define a subspace $V_{i,j,k}$ of $\mathrm{Coker}(D_\Sigma)$ by

$$V_{i,j,k} := \{\sigma \in \mathrm{Coker}(D_\Sigma) \mid \Phi_l(\sigma) = 0 \quad \forall l \notin \{i,j,k\}\}.$$

Lemma 2.18. *(a) Φ_i restricts to an isomorphism $V_{i,j,k} \xrightarrow{\simeq} \mathcal{A}_i$.*

(b) Every nonzero element of $V_{i,j,k}$ is nonvanishing.

Proof. By Lemmas 2.14 and 2.15(b), we know that

$$\dim(\mathrm{Coker}(D_\Sigma)) = 2(N-2). \tag{2.14}$$

It follows that $\dim(V_{i,j,k}) \geq 2$. On the other hand, Lemma 2.15(a) implies that if $0 \neq \sigma \in V_{i,j,k}$, then σ is nonvanishing and $\Phi_i(\sigma) \neq 0$. Assertions (a) and (b) follow. □

Consider now the oriented weighted tree $\tau(\Sigma)$ associated to the branched cover $\pi : \Sigma \to \mathbb{R} \times S^1$. A nonzero special cokernel element $\sigma \in V_{i,j,k}$, being nonvanishing, has a well-defined winding number around the cylinder in Σ corresponding to each edge e of the tree $\tau(\Sigma)$, which we denote by $\eta(\sigma, e)$. We now derive a useful formula for these winding numbers, which will play an essential role in the calculations in §5. First a preliminary lemma:

Lemma 2.19. *Let $T \in T(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$; and let v be an internal vertex of T with outgoing edges of multiplicities $m_1^+, \ldots, m_p^+$ and incoming edges of multiplicities $m_1^-, \ldots, m_q^-$. Then*

$$\sum_{l=1}^{p} \lceil m_l^+ \theta \rceil - \sum_{l=1}^{q} \lfloor m_l^- \theta \rfloor - 1 = 0. \tag{2.15}$$

Proof. Let $\mathrm{ind}_\theta(v)$ denote twice the left hand side of (2.15). By the sum condition on the weights, $\sum_{l=1}^{p} m_l^+$ and $\sum_{l=1}^{q} m_l^-$ are equal, say to r, so

$$\mathrm{ind}_\theta(v) \geq 2(\lceil r\theta \rceil - \lfloor r\theta \rfloor - 1) = 0.$$

On the other hand, a straightforward calculation shows that the quantity ind_θ from Definition 1.6 satisfies

$$\mathrm{ind}_\theta(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-}) = \sum_{v \in \dot{V}(T)} \mathrm{ind}_\theta(v).$$

Thanks to our assumption (2.2), the left hand side of the above equation is zero, and this completes the proof. □

Lemma 2.20. *Let $\sigma \in V_{i,j,k}$ be a nonzero special cokernel element, and let v denote the central vertex for i, j, and k.*

- *If e is an edge on one of the paths $P_{v,i}$, $P_{v,j}$, or $P_{v,k}$, then*

$$\eta(\sigma, e) = \begin{cases} \lceil m(e)\theta \rceil, & e \text{ points away from } v, \\ \lfloor m(e)\theta \rfloor, & e \text{ points towards } v. \end{cases} \tag{2.16}$$

- *If e is not on one of the paths $P_{v,i}$, $P_{v,j}$, or $P_{v,k}$, then*

$$\eta(\sigma, e) = \begin{cases} \lceil m(e)\theta \rceil + 1, & e \text{ points away from } v, \\ \lfloor m(e)\theta \rfloor - 1, & e \text{ points towards } v. \end{cases} \tag{2.17}$$

Proof. Suppose first that e is an external edge, incident to the l^{th} leaf. If $l \in \{i, j, k\}$, then by Lemma 2.18(a), we have $\eta(\sigma, e) = \lceil a_l \theta \rceil$ when $l > 0$ and $\eta(\sigma, e) = \lfloor a_l \theta \rfloor$ when $l > 0$. Thus equation (2.16) holds in this case. If $l \notin \{i, j, k\}$, then the definition of $V_{i,j,k}$ implies that $\eta(\sigma, e) \geq \lceil a_l \theta \rceil + 1$ when $l > 0$, and $\eta(\sigma, e) \leq \lfloor a_l \theta \rfloor - 1$ when $l < 0$. These inequalities must be equalities, or else equation (2.13) would give $\#\sigma^{-1}(0) > 0$, a contradiction. Thus equation (2.17) also holds in this case.

To prove that (2.16) and (2.17) hold for internal edges e, we will use downward induction on the distance (i.e. number of edges on the path) from e to the central vertex v.

To carry out the inductive step, let $w \neq v$ be an internal vertex with outgoing edges $e_1^+, \ldots, e_p^+$ and incoming edges $e_1^-, \ldots, e_q^-$. Since $w \neq v$, there is a unique edge incident to w which is closest to v. By symmetry, we may assume that this edge is incoming, say e_q^-. We may inductively assume that the winding numbers of σ around all other edges incident to w are given by (2.16) and (2.17).

Counting zeroes of σ as in (2.13) over a neighborhood in Σ of the circle $R(w)$ (see Definition 2.5), and using the assumption that σ is nonvanishing, shows that

$$\sum_{l=1}^{p} \eta(\sigma, e_l^+) - \sum_{l=1}^{q} \eta(\sigma, e_l^-) = p + q - 2. \tag{2.18}$$

Now e_q^- points away from v, all other incoming edges of w point towards v, and all outgoing edges of w point away from v. If e_q^- is on one of the paths $P_{v,i}$, $P_{v,j}$, or $P_{v,k}$, then so is exactly one other edge incident to the vertex w, whence by inductive hypothesis,

$$\sum_{l=1}^{p} \eta(\sigma, e_l^+) - \sum_{l=1}^{q-1} \eta(\sigma, e_l^-) = \sum_{l=1}^{p} \lceil m(e_l^+)\theta \rceil - \sum_{l=1}^{q-1} \lfloor m(e_l^-)\theta \rfloor + p + q - 2. \tag{2.19}$$

Combining this with equation (2.18) and Lemma 2.19 gives

$$\eta(\sigma, e_q^-) = \lceil m(e_q^-)\theta \rceil,$$

as desired. If e_q^- is not on one of the paths $P_{v,i}$, $P_{v,j}$, or $P_{v,k}$, then neither is any other edge adjacent to w, so a modification of equation (2.19) holds where we add 1 to the right hand side, giving

$$\eta(\sigma, e_q^-) = \lceil m(e_q^-)\theta \rceil + 1. \qquad \square$$

2.5 Estimates on nonvanishing cokernel elements

Given a nonvanishing cokernel element $\sigma \in \operatorname{Coker}(D_\Sigma)$, we now state an estimate on the relative sizes of the restrictions of σ to different parts of Σ, and some other related estimates, which will be used in §5 and [13].

To state the first estimate, recall that $\tau(\Sigma)$ denotes the oriented weighted tree associated to Σ, which is a one-dimensional CW complex with continuous maps $p : \Sigma \to \tau(\Sigma)$ and $\rho : \tau(\Sigma) \to \mathbb{R}$, such that $\rho \circ p$ equals the composition $\Sigma \xrightarrow{\pi} \mathbb{R} \times S^1 \to \mathbb{R}$. We give $\tau(\Sigma)$ the metric for which ρ restricts to an isometry on each edge.

Given $x, y \in \tau(\Sigma)$, let $P_{x,y}$ denote the path in $\tau(\Sigma)$ from x to y. By an "edge of $P_{x,y}$", we mean an edge e of $\tau(\Sigma)$ such that a positive length subset of e is on the path $P_{x,y}$. Let $P_{x,y}^+$ denote the set of edges of $P_{x,y}$ that are oriented in the direction pointing from x to y, and let $P_{x,y}^-$ denote the set of edges of $P_{x,y}$ whose orientation points from y to x. If e is an edge of $P_{x,y}$, let $\ell(e) > 0$ denote the length of the portion of e that is on the path $P_{x,y}$.

If e is an edge of $\tau(\Sigma)$, then in the notation of §2.2, define

$$E_\pm(\sigma, e) := E^\pm_{\eta(\sigma,e)/m(e)}.$$

Proposition 2.21. *There exists $r' > 0$ with the following property. Let $\Sigma \in \mathcal{M}$ and suppose $\sigma \in \operatorname{Coker}(D_\Sigma)$ is nonvanishing. Let $z, w \in \Sigma$ and define $x := p(z)$ and $y := p(w)$. Then*

$$\log|\sigma(w)| - \log|\sigma(z)| \le \sum_{e \in P^+_{x,y}} \ell(e) E_+(\sigma, e) - \sum_{e \in P^-_{x,y}} \ell(e) E_-(\sigma, e) + r'. \quad (2.20)$$

The proof of Proposition 2.21 is given in §6.3. The following is an important special case:

Corollary 2.22. *If $S(t) = \theta$, then in Proposition 2.21 we can replace* (2.20) *by*

$$\left| \log|\sigma(w)| - \log|\sigma(z)| - \left(\sum_{e \in P^+_{x,y}} - \sum_{e \in P^-_{x,y}} \right) \ell(e) \left(\theta - \frac{\eta(\sigma, e)}{m(e)} \right) \right| \le r'.$$

Proof. In general, by switching the role of z and w in (2.20) we get

$$\log|\sigma(w)| - \log|\sigma(z)| \ge \sum_{e \in P^+_{x,y}} \ell(e) E_-(\sigma, e) - \sum_{e \in P^-_{x,y}} \ell(e) E_+(\sigma, e) - r'. \quad (2.21)$$

The assumption $S(t) = \theta$ implies that

$$E^+_{\eta/m} = E^-_{\eta/m} = \theta - \frac{\eta}{m}.$$

Combining this with (2.20) and (2.21) proves the corollary. □

In the general case, another useful corollary of Proposition 2.21 is that a special cokernel element decays away from the central vertex, in the following sense:

Corollary 2.23. *There exist $c, \kappa > 0$ with the following property. Let $\Sigma \in \mathcal{M}$, let $\sigma \in V_{i,j,k}$ be a special cokernel element, let $z, w \in \Sigma$, and suppose that $p(z)$ is the central vertex for i, j, k. Then*

$$|\sigma(w)| \le c \exp(-\kappa \cdot \operatorname{dist}(p(w), p(z))) |\sigma(z)|. \quad (2.22)$$

Proof. Let $x := p(z)$ and $y := p(w)$. The winding number calculations in Lemma 2.20, and the relation between winding numbers and signs of eigenvalues in Lemma 2.11(c), show that $E_+(\sigma, e) < 0$ for each $e \in P^+_{x,y}$, and $E_-(\sigma, e) > 0$ for each $e \in P^-_{x,y}$. Now for our given α and a_i's, there are only finitely many possible values of $E_\pm(\sigma, e)$. We can then find $\kappa > 0$ such that $E_+(\sigma, e) < -\kappa$ for each $e \in P^+_{x,y}$ and $E_-(\sigma, e) > \kappa$ for each $e \in P^-_{x,y}$. Putting these inequalities into (2.20) proves (2.22). □

We now state one more estimate which we will need.

Definition 2.24. Let $\Sigma \in \mathcal{M}$ and let $\sigma \in \operatorname{Coker}(D_\Sigma)$ be nonvanishing. Define a $(0,1)$-form $\Pi_W\sigma$ on those points $z \in \Sigma$ for which $p(z)$ is not a vertex of $\tau(\Sigma)$, as follows. Let x be a point in the CW-complex $\tau(\Sigma)$ which is in the interior of an edge e. Then $p^{-1}(x)$ is a circle in Σ; choose an identification $p^{-1}(x) \simeq \mathbb{R}/2\pi m(e)\mathbb{Z}$ commuting with the projections to $\mathbb{R} \times S^1$. Let $W(e)$ denote the sum of the eigenspaces of $L_{m(e)}$ whose eigenfunctions have winding number $\eta(\sigma, e)$. On $p^{-1}(x)$, use $d\overline{z}$ to identify $(0,1)$-forms with complex-valued functions, and define $\Pi_W\sigma$ to be the L^2-orthogonal projection of σ onto $W(e)$.

The following proposition, which is proved in §6.2, asserts roughly that away from the ramification points, σ is well approximated by $\Pi_W\sigma$.

Proposition 2.25. *Given $\varepsilon_0 > 0$, there exists $R > 1$ with the following property. Let $\Sigma \in \mathcal{M}$, let $\sigma \in \operatorname{Coker}(D_\Sigma)$ be nonvanishing, let $z \in \Sigma$, and suppose that $p(z)$ has distance at least R from all vertices in $\tau(\Sigma)$. Then*

$$|\sigma(z) - \Pi_W\sigma(z)| < \varepsilon_0 |\Pi_W\sigma(z)|. \tag{2.23}$$

2.6 Orientation of the obstruction bundle

We now specify an orientation of the obstruction bundle $\mathcal{O} \to \mathcal{M}$ associated to an elliptic Reeb orbit. This will be needed later to define various signs.

In the special case $S(t) = \theta$, the operators L_m and D_Σ are complex linear, and so the real vector bundle $\mathcal{O}^*$ has a canonical orientation, which determines an orientation of the dual real vector bundle $\mathcal{O}$.

To orient the obstruction bundle in the general case, note that for any elliptic Reeb orbit, we can deform J to a different admissible almost complex structure J' for which $S(t) = \phi(t)$ for some function $\phi : S^1 \to \mathbb{R}$. One can then deform the trivialization of $\alpha^*\xi$ to arrange that $\phi(t) = \theta$. This process gives rise to a continuous family of Fredholm operators D_Σ, cf. [7]. Since the space of admissible almost complex structures and trivializations of $\alpha^*\xi$ is contractible, we obtain a canonical bijection between orientations of the obstruction bundle for a general $S(t)$ and orientations of the obstruction bundle in the special case $S(t) = \theta$.

2.7 A compactification of $\mathcal{M}/\mathbb{R}$

As in §1.1, $\mathbb{R}$ acts on $\mathcal{M} = \mathcal{M}(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ by translating the $\mathbb{R}$ coordinate on $\mathbb{R} \times S^1$. Given $\Sigma \in \mathcal{M}$, let $[\Sigma]$ denote the equivalence class of Σ in $\mathcal{M}/\mathbb{R}$. We now define a compactification of $\mathcal{M}/\mathbb{R}$, which is slightly different from the symplectic field theory compactification in [2, 5]. This compactification will be used in the analysis in §6.2 and in [13].

Definition 2.26. An element of $\overline{\mathcal{M}/\mathbb{R}}$ is a tuple $(T; [\Sigma_{*1}], \ldots, [\Sigma_{*p}])$ where:

- T is an oriented weighted tree in $T(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ with p internal vertices, and with orderings of the edges and internal vertices such that the edge ordering restricts to the given orderings of the positive and negative leaves.
- Let $m_{j,1}, m_{j,2}, \ldots$ denote the multiplicities of the outgoing edges of the j^{th} internal vertex, and let $n_{j,1}, n_{j,2}, \ldots$ denote the multiplicities of the incoming edges of the j^{th} internal vertex, in their given order. Then

$$\Sigma_{*j} \in \mathcal{M}^{(j)} := \mathcal{M}(m_{j,1}, m_{j,2}, \ldots \mid n_{j,1}, n_{j,2}, \ldots).$$

Two such tuples are equivalent if they differ by the following operations:

- Reordering the edges and internal vertices.
- For an internal edge e from vertex j' to vertex j, acting on the asymptotic markings of the corresponding positive end of $\Sigma_{*j'}$ and negative end of Σ_{*j} by the same element of $\mathbb{Z}/m(e)$.

There is an inclusion $\mathcal{M}/\mathbb{R} \to \overline{\mathcal{M}/\mathbb{R}}$ sending $[\Sigma] \mapsto (T, [\Sigma])$, where T has only one internal vertex.

Definition 2.27. A sequence $\{[\pi_k : \Sigma_k \to \mathbb{R} \times S^1]\}_{k=1,2,\ldots}$ in $\mathcal{M}/\mathbb{R}$ converges to $(T; [\Sigma_{*1}], \ldots, [\Sigma_{*p}]) \in \overline{\mathcal{M}/\mathbb{R}}$ if for all k sufficiently large, there are disjoint closed subsets $\Sigma_{k1}, \ldots, \Sigma_{kp} \subset \Sigma_k$ such that:

(a) Each ramification point in Σ_k is contained in some Σ_{kj}.

(b) Each Σ_{kj} is a component of the π_k-inverse image of a cylinder in $\mathbb{R} \times S^1$, and the length of $\pi_k(\Sigma_{kj})$ goes to infinity as $k \to \infty$.

(c) Let s_{kj} denote the s coordinate of the central circle of $\pi_k(\Sigma_{kj})$. Then the function $\pi_k^* s - s_{kj}$ on the set of ramification points in Σ_{kj} has a k-independent upper bound.

(d) T is obtained from $\tau(\Sigma_k)$ by for each $j \in \{1, \ldots, p\}$ collapsing the vertices corresponding to the ramification points in Σ_{kj} to a single vertex.

(e) For each internal edge e of T, there is an identification $\Phi_{k,e}$ of the corresponding cylinder in Σ_k with an interval cross $\mathbb{R}/2\pi m(e)\mathbb{Z}$, commuting with the projections to $\mathbb{R}\times S^1$, such that the following holds. For each internal vertex j of T, let $\widehat{\Sigma}_{kj} \in \mathcal{M}^{(j)}$ be obtained by attaching half-infinite cylinders to the boundary circles of Σ_{kj}, and asymptotically marking the ends corresponding to internal edges using the identifications $\Phi_{k,e}$. Let $T_{-s_{kj}}$ denote the translation $s \mapsto s - s_{kj}$. Choose the representatives Σ_{*j} so that their ramification points are centered at $s = 0$. Then $\lim_{k\to\infty} T_{-s_{kj}}(\widehat{\Sigma}_{kj}) = \Sigma_{*j}$ in $\mathcal{M}^{(j)}$.

Lemma 2.28. *Any sequence in $\mathcal{M}/\mathbb{R}$ has a subsequence which converges in $\overline{\mathcal{M}/\mathbb{R}}$.*

Proof. Let $\{[\pi_k : \Sigma_k \to \mathbb{R}\times S^1]\}_{k=1,2,\ldots}$ be a sequence in $\mathcal{M}/\mathbb{R}$. The number of ramification points in Σ_k counted with multiplicity is $N-2$, which is independent of k. Hence we can pass to a subsequence so that there is an integer $p \in \{1,\ldots,N-2\}$, and for each k a partition of the ramification points in Σ_k into subsets $\{\Lambda_{kj}\}_{j=1,\ldots,p}$, such that:

- The diameter in Σ_k of Λ_{kj} has a k-independent upper bound.
- If $j \neq j'$, then the distance in Σ_k between Λ_{kj} and $\Lambda_{kj'}$ is greater than k.

We can then refine the sequence so that there are disjoint closed subsets $\Sigma_{k1},\ldots,\Sigma_{kp} \subset \Sigma_k$ satisfying conditions (a)–(c) in Definition 2.27.

To keep track of the combinatorics of the Σ_{kj}'s, for each k define an oriented weighted tree $T_k \in T(a_1,\ldots,a_{N_+} \mid a_{-1},\ldots,a_{-N_+})$ as follows. The internal vertices of T_k are labeled by $1,\ldots,p$ and identified with $\Sigma_{k1},\ldots,\Sigma_{kp}$. Edges correspond to components of $\Sigma_k \setminus \bigcup_k \Sigma_{kj}$, and the multiplicity of an edge is the degree of π_k on the corresponding cylinder. The internal (resp. external) vertices at the endpoints of an edge are determined by the boundary circles (resp. ends) of the associated cylinder. Edges are oriented in the increasing s direction.

Since there are only finitely many oriented weighted trees in the set $T(a_1,\ldots,a_{N_+} \mid a_{-1},\ldots,a_{-N_+})$, we can pass to a subsequence such that all the trees T_k are isomorphic, preserving the leaf labels, to a single oriented weighted tree T. Note that these isomorphisms $T_k \simeq T$ are canonical since a tree has no nontrivial automorphisms fixing the leaves. Choose orderings of the edges and internal vertices of T as in Definition 2.26, and use the vertex ordering to order $\Sigma_{k1},\ldots,\Sigma_{kp}$. We have now achieved requirement (d) in Definition 2.27.

Next, for each internal edge e of T, fix an identification $\Phi_{k,e}$ of the corresponding cylinder in Σ_k with an interval cross $\mathbb{R}/2\pi m(e)\mathbb{Z}$, commuting with the projections to $\mathbb{R}\times S^1$. Then $\widehat{\Sigma}_{kj}\in\mathcal{M}^{(j)}$ is defined. Moreover, since condition (c) in Definition 2.27 holds, there is a k-independent constant R such that each $T_{-s_{kj}}(\widehat{\Sigma}_{kj})$ is in $\mathcal{M}_R(m_{j,1},m_{j,2},\ldots\mid n_{j,1},n_{j,2},\ldots)$. By Lemma 2.19, we have

$$\kappa_\theta(m_{j,1},m_{j,2},\ldots\mid n_{j,1},n_{j,2},\ldots)=1.$$

Hence we can apply Lemma 2.8 to find a further subsequence satisfying condition (e) in Definition 2.27. □

3 The linearized section of the obstruction bundle

In this section, fix an admissible almost complex structure J on $\mathbb{R}\times Y$, an elliptic Reeb orbit α with monodromy angle θ, and

$$S=(a_1,\ldots,a_{\overline{N}_+};a_{\overline{N}_++1},\ldots,a_{N_+}\mid a_{-1},\ldots,a_{\overline{N}_-};a_{-(\overline{N}_-+1)},\ldots,a_{-N_-}) \tag{3.1}$$

satisfying the sum condition (1.9). Assume as usual that $N:=N_++N_->2$. Assume also that $\kappa_\theta(S)=1$. Finally, we make:

Assumption 3.1 (Partition minimality). Under the partial order $\geq_\theta$ in Definition 1.8, the partition $(a_{\overline{N}_++1},\ldots,a_{N_+})$ is minimal, and the partition $(a_{-(\overline{N}_-+1)},\ldots,a_{-N_-})$ is maximal.

By §2, there is an obstruction bundle $\mathcal{O}$ over the moduli space of branched covers $\mathcal{M}=\mathcal{M}(a_1,\ldots,a_{N_+}\mid a_{-1},\ldots,a_{-N_-})$ with $\operatorname{rank}(\mathcal{O})=\dim(\mathcal{M})$.

We will now define a section $\mathfrak{s}_0$ of $\mathcal{O}$ over $\mathcal{M}_R$ (see Definition 2.7). We will then quote results from [13] which relate an appropriate count of zeroes of $\mathfrak{s}_0$ to the count of gluings in Theorem 1.13.

3.1 Definition of the linearized section $\mathfrak{s}_0$

For $i=1,\ldots,\overline{N}_+$, let λ_i denote the largest negative eigenvalue of the operator L_{a_i}. Likewise, for $i=-1,\ldots,-\overline{N}_-$, let λ_i denote the smallest positive eigenvalue of the operator L_{a_i}. Let $\mathcal{B}_i$ denote the λ_i-eigenspace of L_{a_i}, and let $\mathcal{B}$ denote the direct sum of the $\mathcal{B}_i$'s for $i=1,\ldots,\overline{N}_+$ and $i=-1,\ldots,-\overline{N}_-$. Given $\Sigma\in\mathcal{M}$ and $\sigma\in\operatorname{Coker}(D_\Sigma)$, over the i^{th} end of Σ (identified with $[R,\infty)\times\mathbb{R}/2\pi a_i\mathbb{Z}$ or $(-\infty,-R]\times\mathbb{R}/2\pi a_i\mathbb{Z}$ via the asymptotic marking), write $\sigma=\sigma_i(s,t)d\overline{z}$.

Definition 3.2. Fix $\gamma=\{\gamma_i\}\in\mathcal{B}$ and real numbers $R,r>0$. Define the *linearized section* $\mathfrak{s}_0$ of the obstruction bundle $\mathcal{O}\to\mathcal{M}_R$ as follows. If $\Sigma\in\mathcal{M}_R$ and $\sigma\in\operatorname{Coker}(D_\Sigma)$, then

$$\mathfrak{s}_0(\Sigma)(\sigma):=\sum_{i=1}^{\overline{N}_+}\langle\gamma_i,\sigma_i(R+r,\cdot)\rangle-\sum_{i=-1}^{-\overline{N}_-}\langle\gamma_i,\sigma_i(-R-r,\cdot)\rangle\in\mathbb{R}. \tag{3.2}$$

Here the brackets denote the real inner product on $L^2(\mathbb{R}/2\pi a_i\mathbb{Z}, \mathbb{R}^2)$.

We will study the linearized section under the assumption that the eigenvectors γ are "admissible" in the following sense. If $i, j \in \{1, \ldots, \overline{N}_+\}$ and $\lceil a_i\theta \rceil / a_i = \lceil a_j\theta \rceil / a_j$, then we can identify $\mathcal{B}_i = \mathcal{B}_j$ using coverings as follows. Let $\phi := \lceil a_i\theta \rceil / a_i$, so that $\lambda_i = \lambda_j = E^+_\phi =: \lambda$. Write $\phi = \eta/m$ where $\eta, m \in \mathbb{Z}$ and $m > 0$ is as small as possible; then the integers a_i and a_j are both divisible by m. For any positive integer d, we can pull back an eigenfunction of L_m with eigenvalue λ to an eigenfunction of L_{md} with the same eigenvalue. In this way we identify the λ-eigenspaces of L_{md} for different d with each other. Likewise, if $i, j \in \{-1, \ldots, -\overline{N}_-\}$, then we can identify $\mathcal{B}_i = \mathcal{B}_j$ when $\lfloor a_i\theta \rfloor / a_i = \lfloor a_j\theta \rfloor / a_j$.

Note also that the cyclic group $\mathbb{Z}/a_i$ acts linearly on the eigenspace $\mathcal{B}_i$, via pullback from its action on $\mathbb{R}/2\pi a_i\mathbb{Z}$ by deck transformations of the covering map to $S^1 = \mathbb{R}/2\pi\mathbb{Z}$.

Definition 3.3. An element $\gamma = \{\gamma_i\} \in \mathcal{B}$ is *admissible* if:

(a) $\gamma_i \neq 0$ for all $i = 1, \ldots, \overline{N}_+$ and $i = -1, \ldots, -\overline{N}_-$.

(b) If $i, j \in \{1, \ldots, \overline{N}_+\}$ and $\lceil a_i\theta \rceil / a_i = \lceil a_j\theta \rceil / a_j$, or $i, j \in \{-1, \ldots, -\overline{N}_-\}$ and $\lfloor a_i\theta \rfloor / a_i = \lfloor a_j\theta \rfloor / a_j$, then for all $g_i \in \mathbb{Z}/a_i$ and $g_j \in \mathbb{Z}/a_j$ we have $g_i \cdot \gamma_i \neq g_j \cdot \gamma_j$ in $\mathcal{B}_i = \mathcal{B}_j$.

3.2 Counting zeroes of $\mathfrak{s}_0$

We now want to count zeroes of $\mathfrak{s}_0$, for which purpose we will use the following formalism.

Definition 3.4. Let ψ be a section of $\mathcal{O}$ over $\mathcal{M}_R$. Suppose that ψ is nonvanishing on $\partial\mathcal{M}_R$. Then define the *relative Euler class*

$$e(\mathcal{O} \to \mathcal{M}_R, \psi) \in \mathbb{Z}$$

as follows. Let ψ' be a section of $\mathcal{O}$ over $\mathcal{M}_R$ such that $\psi = \psi'$ on $\partial\mathcal{M}_R$, and all zeroes of ψ' are nondegenerate. Define $e(\mathcal{O} \to \mathcal{M}_R, \psi)$ to be the signed count of zeroes of ψ', where the signs are determined using the orientation of $\mathcal{M}_R$ as a complex manifold and the orientation of $\mathcal{O}$ defined in §2.6. By Lemma 2.8, this count is well-defined and depends only on $\psi|_{\partial\mathcal{M}_R}$. We usually denote this relative Euler class by $\#\psi^{-1}(0)$, even though the zeroes of ψ itself may be degenerate.

Note that if $\{\psi_t \mid t \in [0,1]\}$ is a homotopy of sections of $\mathcal{O} \to \mathcal{M}_R$ such that $\psi_t|_{\partial\mathcal{M}_R}$ is nonvanishing for all $t \in [0,1]$, then $\#\psi_0^{-1}(0) = \#\psi_1^{-1}(0)$.

The following lemma gives sufficient conditions for $\mathfrak{s}_0$ to be nonvanishing on $\partial\mathcal{M}_R$. To state it, let $\lambda := \min\{|\lambda_i|\}$ and $\Lambda := \max\{|\lambda_i|\}$.

Lemma 3.5. *Given admissible $\gamma = \{\gamma_i\} \in \mathcal{B}$, if r is sufficiently large with respect to γ, and if $R > 3\Lambda r/\lambda$, then $\mathfrak{s}_0$ has no zeroes on $\mathcal{M}_R \setminus \mathcal{M}_{R-r}$.*

Proof. This follows from [13, Prop. 8.2], as in the proof of [13, Cor. 8.6]. □

It follows that if γ, r, and R satisfy the hypotheses of Lemma 3.5, then the relative Euler class $\#\mathfrak{s}_0^{-1}(0)$ is defined. Moreover:

Corollary 3.6. *Under the assumptions of Lemma 3.5, the relative Euler class $\#\mathfrak{s}_0^{-1}(0)$ depends only on S and θ.*

Proof. As in §2.6, we can deform the almost complex structure J on $\mathbb{R} \times Y$ and the trivialization of $\alpha^*\xi$ so as to deform the operator D_Σ to a complex linear operator depending only on S and θ, in which the operators L_{a_i} are also complex linear. We can simultaneously deform the collection of eigenvectors γ while preserving the admissibility conditions. Once the operators L_{a_i} are complex linear, the set of admissible γ is connected, because now the conditions $\gamma_i = 0$ and $g_i \cdot \gamma_i = g_j \cdot \gamma_j$ have real codimension 2. Thus all of the different versions of the section $\mathfrak{s}_0$ are homotopic. For any given homotopy of the almost complex structure, the trivialization, and the admissible γ, if we fix r and R sufficiently large, then Lemma 3.5 will apply throughout the homotopy so that the count of zeroes does not change. Increasing r and R does not change the count of zeroes for the same reason. □

3.3 Zeroes of $\mathfrak{s}_0$ and gluings

The significance of $\#\mathfrak{s}_0^{-1}(0)$ for the gluing story is as follows. Let (u_+, u_-) be a gluing pair satisfying conditions (i) and (ii) in §1.8. Order the negative ends of u_+ and the positive ends of u_- such that the $\mathbb{R}$-invariant negative ends of u_+ are those labeled by $\overline{N}_+ + 1, \ldots, N_+$, and the $\mathbb{R}$-invariant positive ends of u_- are those labeled by $-(\overline{N}_- + 1), \ldots, -N_-$. The main result of [13] can be stated as follows:

Theorem 3.7. *[13, Thm. 1.1] Fix coherent orientations and generic J. If (u_+, u_-) is a gluing pair as above, then*

$$\#G(u_+, u_-) = \epsilon(u_+)\epsilon(u_-)\#\mathfrak{s}_0^{-1}(0),$$

where $\#\mathfrak{s}_0^{-1}(0)$ is defined as in Corollary 3.6.

A few words are in order concerning the proof of this theorem. As sketched in §1.8, we have

$$\#G(u_+, u_-) = \epsilon(u_+)\epsilon(u_-)\#\mathfrak{s}^{-1}(0),$$

where $\mathfrak{s}$ is a section of $\mathcal{O}$ over $\mathcal{M}_R$ arising from the gluing analysis. The idea is to relate the section $\mathfrak{s}$ to the linearized section $\mathfrak{s}_0$, where the latter is defined using certain eigenfunctions $\gamma \in \mathcal{B}$ that are determined by the asymptotics of the negative ends of u_+ and the positive ends of u_-. If J is generic, then the asymptotic eigenfunctions γ are admissible. In this case a generalization of Lemma 3.5 shows that the homotopy of sections $\mathfrak{s}_t := t\mathfrak{s} + (1-t)\mathfrak{s}_0$ has no zeroes on $\partial\mathcal{M}_R$ for $t \in [0,1]$, so that $\#\mathfrak{s}^{-1}(0) = \#\mathfrak{s}_0^{-1}(0)$.

3.4 Consequences of partition minimality

By Theorem 3.7, to prove Theorem 1.13 we just need to compute $\#\mathfrak{s}_0^{-1}(0)$. Before doing so, we note two basic facts about the numbers $\overline{N}_+$ and $\overline{N}_-$ in (3.1).

Lemma 3.8. *The assumption $\kappa_\theta(S) = 1$ and Assumption 3.1 imply that $\overline{N}_+ \geq N_+ - 1$ and $\overline{N}_- \geq N_- - 1$.*

Proof. By symmetry it is enough to show that $\overline{N}_+ \geq N_+ - 1$. Suppose to the contrary that $N_+ - \overline{N}_+ \geq 2$. Then we can construct a tree $T \in T(S)$ which has a splitting vertex v with one incoming edge and with outgoing edges incident to the positive leaves labeled by $N_+ - 1$ and N_+. Then by Lemma 2.19,

$$\lceil a_{N_+-1}\theta \rceil + \lceil a_{N_+}\theta \rceil = \lceil (a_{N_+-1} + a_{N_+})\theta \rceil .$$

It follows from this that

$$(a_{\overline{N}_++1}, \ldots, a_{N_+}) >_\theta (a_{\overline{N}_++1}, \ldots, a_{N_+-2}, a_{N_+-1} + a_{N_+}).$$

This contradicts the assertion in Assumption 3.1 that $(a_{\overline{N}_++1}, \ldots, a_{N_+})$ is minimal. □

Lemma 3.9. *If γ is admissible and $\mathfrak{s}_0(\Sigma) = 0$, then $\overline{N}_+ = N_+$ or $\overline{N}_- = N_-$.*

Proof. If not, then by Lemma 3.8, $\overline{N}_+ = N_+ - 1$ and $\overline{N}_- = N_- - 1$. Since we are assuming that $N > 2$, there is an end labeled by $i \notin \{N_+, -N_-\}$. Since γ is admissible, $\gamma_i \neq 0$. By Lemma 2.18(a), the projection $V_{i,N_+,-N_-} \to \mathcal{B}_i$ is surjective. Hence we can find a special cokernel element $\sigma \in V_{i,N_+,-N_-}$ such that the σ_i term in (3.2) is nonzero. Then $\mathfrak{s}_0(\Sigma)(\sigma) \neq 0$, because all other terms on the r.h.s. of (3.2) are zero. □

Corollary 3.10. *If $\overline{N}_+ \neq N_+$ and $\overline{N}_- \neq N_-$, then*

$$\#\mathfrak{s}_0^{-1}(0) = c_\theta(S). \tag{3.3}$$

Proof. The left side of (3.3) is zero by Lemma 3.9, while the right hand side of (3.3) is zero by Definition 1.26 and our assumption that $N > 2$. □

4 Combinatorics of the elliptic gluing coefficients

This section gives a diagrammatic reinterpretation of the combinatorial gluing coefficient $c_\theta(S)$ from §1.6. Recall that when $\kappa_\theta(S) = 1$, the integer $c_\theta(S)$ is determined by an auxiliary function f_θ, evaluated on a certain reordering of the list S. The main result of this section, Lemma 4.5 below, expresses $f_\theta(S)$ as a sum, over "admissible" trivalent trees with "edge pairings", of certain positive integer weights. (A straightforward extension of this interprets $c_\theta(S)$ when $\kappa_\theta(S) > 1$ as a sum over forests.) This alternate definition is lengthier but more symmetric, and leads to a proof that $c_\theta(S)$ satisfies the symmetry property (1.17). The combinatorics introduced here will be used in the computation of $\#\mathfrak{s}_0^{-1}(0)$ in §5.

For the rest of this section fix $S = (a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$, where $a_1, \ldots, a_{N_+}$ and $a_{-1}, \ldots, a_{-N_-}$ are positive integers satisfying (2.1).

4.1 Edge pairings, weights, and admissible trees

We begin with some combinatorial definitions.

Definition 4.1. Let T be a trivalent tree. An *edge pairing* P on T is an assignment, to each internal vertex v of T, of two distinct edges e_v^+ and e_v^- incident to v, such that the sets $\{e_v^+, e_v^-\}$ and $\{e_w^+, e_w^-\}$ are disjoint whenever v and w are adjacent (and hence whenever v and w are distinct) internal vertices.

Note that in a trivalent tree, the number of edges equals twice the number of internal vertices plus one. Hence for any edge pairing P on T, there is a distinguished edge e_0 which is not one of the edges e_v^+ or e_v^- for any internal vertex v.

Definition 4.2. If T is an oriented weighted trivalent tree, if P is an edge pairing on T, and if θ is an irrational number, define the *weight* $W_\theta(T, P) \in \mathbb{Z}^{>0}$ as follows. For an internal vertex v, define m_v^+ to be "the outward flow along the edge e_v^+", namely $m(e_v^+)$ if e_v^+ points outward from v, and $-m(e_v^+)$ if e_v^+ points inward toward v. Similarly define m_v^- to be "the inward flow along the edge e_v^-", namely $m(e_v^-)$ if e_v^- points inward toward v, and $-m(e_v^-)$ if e_v^- points outward from v. Then

$$W_\theta(T, P) := m(e_0) \prod_{v \in \dot{V}(T)} \Big(m_v^- \lceil m_v^+ \theta \rceil - m_v^+ \lfloor m_v^- \theta \rfloor \Big). \tag{4.1}$$

The above combinatorial notions arise naturally in the analysis, as explained in Remark 5.20.

Next, recall that $T(S)$ denotes the set of oriented weighted trees whose positive leaves are labeled by $1, \ldots, N_+$ and whose negative leaves are labeled by $-1, \ldots, -N_-$, such that the i^{th} leaf has multiplicity a_i.

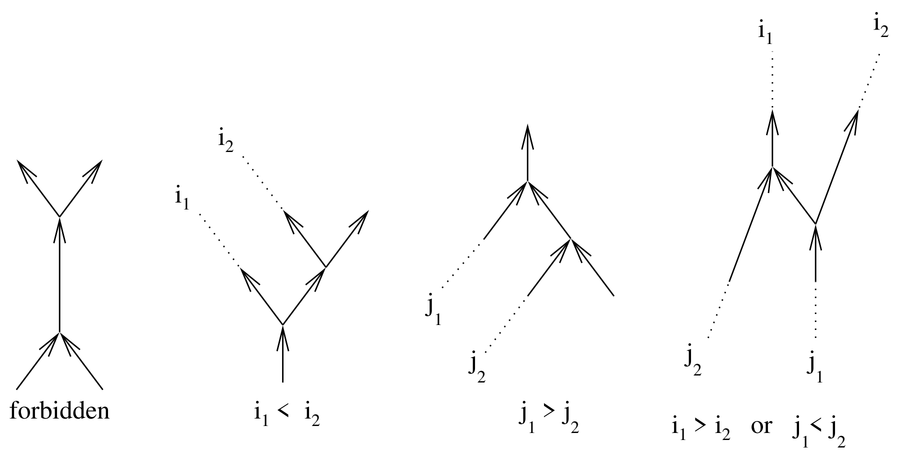

Figure 2: The rules for an admissible tree.

Definition 4.3. A trivalent tree $T \in T(S)$ is *admissible* if the following conditions hold (see Figure 2):

(a) No oriented edge starts at a joining vertex and ends at a splitting vertex.

(b) Let v be a splitting vertex with outgoing edges e_1 and e_2. Suppose there is an upward path starting along e_1 and ending at the positive leaf i_1. Suppose e_2 is incident to another splitting vertex w, from which there is an upward path leading to the positive leaf i_2. Then $i_1 < i_2$.

(c) Symmetrically to (b), let v be a joining vertex with incoming edges e_1 and e_2. Suppose there is a downward path starting along e_1 and ending at the negative leaf j_1. Suppose e_2 is incident to another joining vertex w, from which there is a downward path leading to the negative leaf j_2. Then $j_1 > j_2$.

(d) Let e be an edge from a splitting vertex w to a joining vertex v. Suppose there is an upward path from v to the positive leaf i_1 and a downward path from w to the negative leaf j_1. Suppose there is an upward path from w, not containing e, to the positive leaf i_2, and a downward path from v, not containing e, to the negative leaf j_2. Then $i_1 > i_2$ or $j_1 < j_2$ (or both).

Lemma 4.4. *If T is admissible, then there is a unique edge pairing P_T on T such that (see Figure 3):*

(i) If v is a splitting vertex, then e_v^- is the incoming edge; there is a unique upward path starting along e_v^+, say to the positive leaf i_1; and if i_2

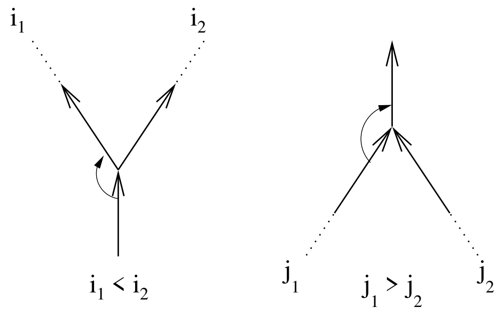

Figure 3: The canonical edge pairing for an admissible tree. The symbolism is that at each internal vertex v, the small arc starts at e_v^- and ends at e_v^+.

is a positive leaf reached by an upward path starting along the other outgoing edge, then $i_1 < i_2$*.*

(ii) Symmetrically to (i), if v *is a joining vertex, then* e_v^+ *is the outgoing edge; there is a unique downward path starting along* e_v^-*, say to the negative leaf* j_1*; and if* j_2 *is a negative leaf reached by a downward path starting along the other incoming edge, then* $j_1 > j_2$*.*

Proof. Let v be a splitting vertex. By admissibility condition (b), at least one outgoing edge of v is incident to a joining vertex or to a positive leaf. By admissibility condition (a), there is a unique upward path starting along such an edge. By condition (b) again, there is a unique such edge e_v^+ satisfying condition (i) above.

Symmetrically, if v is a joining vertex, then there is a unique incoming edge e_v^- satisfying condition (ii) above.

To see that P_T satisfies the disjointness condition in the definition of edge pairing, we must check that an edge e from an internal vertex v to an internal vertex w cannot be in both of the sets $\{e_v^+, e_v^-\}$ and $\{e_w^+, e_w^-\}$. There are four cases to consider, depending on whether the vertices v and w are joining or splitting. These four cases are precisely covered by admissibility conditions (a)–(d). $\square$

Let $\mathcal{A}(S)$ denote the set of admissible trees in $T(S)$. We can now interpret $f_\theta(S)$ as a sum over admissible trees of the weights associated to their canonical edge pairings:

Lemma 4.5.

$$f_\theta(S) = \sum_{T \in \mathcal{A}(S)} W_\theta(T, P_T).$$

The proof of this lemma is postponed to §4.2. We can now prove the symmetry property (1.17). By the definition of c_θ, it is enough to show:

Corollary 4.6. *The function f_θ satisfies the symmetry property*

$$f_\theta(a_1,\ldots,a_{N_+} \mid a_{-1},\ldots,a_{-N_-}) = f_{-\theta}(a_{-1},\ldots,a_{-N_-} \mid a_1,\ldots,a_{N_+}). \quad (4.2)$$

Proof. The definition of admissible tree is symmetric, in that reversing edge orientations defines a bijection

$$\phi : \mathcal{A}(a_1,\ldots,a_{N_+} \mid a_{-1},\ldots,a_{-N_-}) \longrightarrow \mathcal{A}(a_{-1},\ldots,a_{-N_-} \mid a_1,\ldots,a_{N_+}).$$

The canonical edge pairing is also symmetric, in that $P_{\phi(T)}$ is obtained from P_T by switching e_v^+ and e_v^- for all v. It now follows from equation (4.1), using the identity $-\lceil x\rceil = \lfloor -x\rfloor$, that

$$W_{-\theta}(\phi(T), P_{\phi(T)}) = W_\theta(T, P_T).$$

Equation (4.2) then follows from Lemma 4.5. $\square$

Remark 4.7. Assume that $N_+ > 1$ and $\kappa_\theta(S) = 1$. Let S' be obtained from switching a_1 and a_2:

$$S' := (a_2, a_1, a_3, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-}).$$

Let $\widehat{S}$ be obtained from S by adding the first two entries:

$$\widehat{S} := (a_1 + a_2, a_3, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-}).$$

One can then show that

$$f_\theta(S) - f_\theta(S') = (a_2\lceil a_1\theta\rceil - a_1\lceil a_2\theta\rceil)\cdot f_\theta(\widehat{S}). \quad (4.3)$$

Equation (4.3) implies that $c_\theta(S)$ does not depend on the choice of reordering satisfying (1.15), as discussed in Remark 1.23. To prove (4.3), one can first show that when $\kappa_\theta(S) = 1$, the function W_θ on trees with edge pairings satisfies a version of the IHX relation. (For the usual IHX relation see e.g. [1].) One can then expand both sides of equation (4.3) using Lemma 4.5 and show that the difference is a linear combination of IHX relations. We omit the details, because we will give a different proof that $c_\theta(S)$ is well-defined in §5.

4.2 Enumerating admissible trees

To prove Lemma 4.5, we now introduce a way to enumerate admissible trees, which is less symmetric, but more closely related to the definition of f_θ and to the obstruction bundle calculations in §5.

Definition 4.8. Let $\mathcal{E}(S)$ denote the set of N_+-tuples $(E_1, \ldots, E_{N_+})$ of nonempty subsets of $\{-1, \ldots, -N_-\}$, such that:

(i) For $i = 1, \ldots, N_+ - 1$, let $E_i' := E_i \setminus \{\min(E_i)\}$. Then

$$\{-1, \ldots, -N_-\} = E_{N_+} \sqcup \bigsqcup_{i=1}^{N_+-1} E_i'. \tag{4.4}$$

(ii) If $N_+ > 1$, then

$$\sum_{j \in E_1'} a_j < a_1 < \sum_{j \in E_1} a_j. \tag{4.5}$$

Also, let

$$\overline{S} := (a_2, \ldots, a_{N_+} \mid a_{-E_1})$$

where a_{-E_1} denotes the arguments a_j for $0 > j \notin E_1$ arranged in decreasing order, with one additional argument equal to $\sum_{j \in E_1} a_j - a_1$, inserted into the position that $\min(E_1)$ would occupy in the order. Let

$$\xi : \{-1, \ldots, -N_-\} \setminus E_1' \longrightarrow \{-1, \ldots, -(N_- - |E_1'|)\}$$

denote the order-preserving bijection. Then

$$\overline{E} := (\xi(E_2), \ldots, \xi(E_{N_+})) \in \mathcal{E}(\overline{S}). \tag{4.6}$$

Definition 4.9. We define a function

$$\phi : \mathcal{E}(S) \longrightarrow \mathcal{A}(S)$$

as follows. If $N_+ = 1$, then by (4.4) there is a unique element $E \in \mathcal{E}(S)$, given by $E_1 = \{-1, \ldots, -N_-\}$. In the tree $\phi(E)$, the path from the positive leaf to the negative leaf $-N_-$ goes through $N_- - 1$ trivalent joining vertices. These joining vertices are adjacent to the negative leaves $-1, \ldots, -(N_- - 1)$ in that order. See Figure 4.

If $N_+ > 1$, then given $(E_1, \ldots, E_{N_+}) \in \mathcal{E}$, construct the admissible tree $T = \phi(E_1, \ldots, E_{N_+})$ inductively as follows. We first define an oriented weighted tree T_1 with two positive leaves and with negative leaves indexed by E_1. To construct T_1, draw a downward path γ from the first positive leaf to the negative leaf indexed by $\min(E_1)$. This path will go through $|E_1'|$

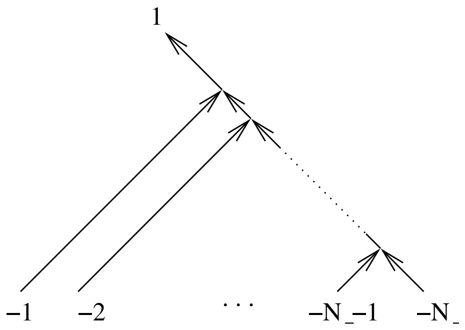

Figure 4: When $N_+ = 1$, there is a unique admissible tree.

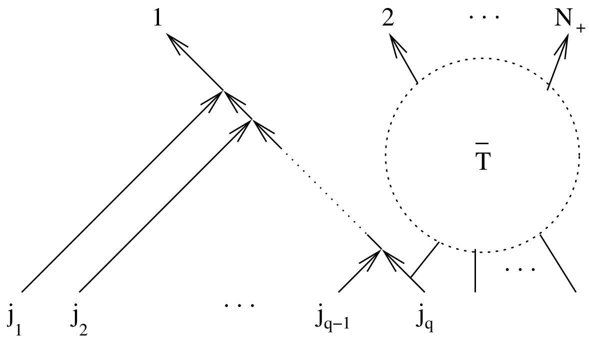

Figure 5: The inductive construction of an admissible tree when $N_+ > 1$. Here $E_1 = \{j_1 > \cdots > j_q\}$.

trivalent joining vertices. For each of these joining vertices, the downward edge not on γ is incident to a negative leaf. These negative leaves are those indexed by E_1', in decreasing order. Now add a trivalent splitting vertex below the lowest joining vertex; the new outgoing edge of this vertex is incident to the second positive leaf of T_1. The negative leaf of T_1 indexed by $j \in E_1$ has weight a_j. The first positive leaf has weight a_1 and the second positive leaf has weight $\sum_{j\in E_1} a_j - a_1$. The leaf weights of T_1 extend to unique weights on the internal edges satisfying the conservation condition at the vertices of T_1. Condition (4.5) insures that the edges of T_1 all have positive weight.

Next, by (4.6) and induction, $\overline{E}$ determines an oriented weighted tree $\overline{T} = \phi(\overline{E}) \in \mathcal{A}(\overline{S})$. To construct the tree T, glue the second positive leaf of T_1 to the negative leaf of $\overline{T}$ indexed by $\xi(\min(E_1))$. (See Figure 5.) It is straightforward to verify that the tree T is admissible.

The above notions are connected to the definition of f_θ as follows.

Lemma 4.10.

$$f_\theta(S) = \sum_{E\in\mathcal{E}(S)} W_\theta(\phi(E), P_{\phi(E)}). \tag{4.7}$$

Proof. Summary: Unraveling the definitions shows that if one expands the recursive formula (1.14) for $f_\theta(S)$, then one obtains a sum indexed by elements $E = (E_1, \ldots, E_{N_+}) \in \mathcal{E}(S)$, and the summand corresponding to E agrees with the summand on the right hand side of (4.7).

Details: If $N_+ = 1$ then (4.7) is immediate from the definitions since there is just one summand. Suppose now that $N_+ > 1$. Write $E_1 = \{j_1 > \cdots > j_q\}$. By equation (1.14) and induction on N_+ we can write

$$f_\theta(S) = \sum_{E_1} \sum_{\overline{E}\in\mathcal{E}(\overline{S})} W_\theta\left(\overline{T}, P_{\overline{T}}\right) \prod_{n=1}^{q} \delta_\theta\left(a_1 - \sum_{k=1}^{n-1} a_{j_k}\, ,\, a_{j_n}\right). \tag{4.8}$$

Here the sum is over E_1 satisfying condition (4.5). To clarify this, write $\overline{E} =: (\xi(E_2), \ldots, \xi(E_{N_+}))$ and $E := (E_1, E_2, \ldots, E_{N_+})$. By (4.6), $\overline{E} \in \mathcal{E}(\overline{S})$ is equivalent to $E \in \mathcal{E}(S)$. Thus (4.8) can be regarded as a sum over $E \in \mathcal{E}(S)$.

Now the tree T_1 in the definition of $\phi(E)$ is admissible, and with its canonical edge pairing has weight

$$W_\theta(T_1, P_{T_1}) = m(e_0) \prod_{n=1}^{q} \delta_\theta\left(a_1 - \sum_{k=1}^{n-1} a_{j_k}\, ,\, a_{j_n}\right). \tag{4.9}$$

Here e_0 denotes the unpaired edge in the canonical edge pairing P_{T_1}. Since e is also the edge along which T_1 and $\overline{T}$ are glued together, we have

$$W_\theta(\phi(E), P_{\phi(E)}) = \frac{W_\theta(T_1, P_{T_1})}{m(e_0)} W_\theta\left(\overline{T}, P_{\overline{T}}\right). \tag{4.10}$$

Regarding (4.8) as a sum over $E \in \mathcal{E}(S)$ and plugging in (4.9) and (4.10) gives (4.7). □

Lemma 4.5 now follows from Lemma 4.10 and:

Lemma 4.11. *The function $\phi : \mathcal{E}(S) \to \mathcal{A}(S)$ is a bijection.*

Proof. Observe that if $E = (E_1, \ldots, E_{N_+}) \in \mathcal{E}$, then E_i is the set of negative leaves of $\phi(E)$ that are accessible by downward paths starting at the i^{th} positive leaf. Thus it is enough to show the following: Given an admissible tree $T \in \mathcal{A}(S)$, let E_i denote the set of negative leaves of T that are accessible by downward paths starting at the i^{th} positive leaf. Then $(E_1, \ldots, E_{N_+}) \in \mathcal{E}$ and $T = \phi(E_1, \ldots, E_{N_+})$.

To begin the proof of this, let γ denote the downward path from the first positive leaf to the negative leaf indexed by $\min(E_1)$. By the definition of

E_1, for each $j \in E_1'$ there must be a joining vertex w_j on γ such that there is a downward path γ_j from w_j to the negative leaf indexed by j, where γ_j does not intersect γ except at w_j. By admissibility condition (c), the first edge on γ_j cannot be incident to a joining vertex (other than w_j). Then by admissibility condition (a), w_j is the only joining vertex on the path γ_j. Hence the vertices w_j are distinct. By condition (c) again, if $j > j'$ then w_j is above $w_{j'}$ on γ. If $N_+ = 1$ then we are done, so assume henceforth that $N_+ > 1$.

By the definition of E_1, the path γ cannot meet any joining vertex except for the w_j's. We have seen above that γ_j cannot meet any joining vertex except for w_j. By admissibility condition (d), the path γ_j cannot meet any splitting vertex. Hence the path γ_j consists of a single edge incident to w_j and the negative leaf indexed by j. Since the tree T is connected, the path γ must meet at least one splitting vertex. By admissibility condition (a), any such splitting vertex is below all of the joining vertices on γ. By admissibility condition (b), there is only one splitting vertex on γ, call it v_1.

Let e denote the outgoing edge of v_1 that is not on γ. Cut T along e to obtain two oriented weighted trees T_1 and $\overline{T}$, where T_1 contains γ and $\overline{T}$ does not. Order the positive and negative leaves of $\overline{T}$ with the orderings induced from those of T (where the negative leaf of $\overline{T}$ takes the place of $\min(E_1)$ in the ordering). Then $\overline{T}$ is admissible, so by induction,

$$(\xi(E_2), \ldots, \xi(E_{N_+})) \in \mathcal{E}(a_2, \ldots, a_{N_+} \mid a_{-E_1})$$

and $\overline{T} = \phi(\xi(E_2), \ldots, \xi(E_{N_+}))$. Clearly E_1' is disjoint from $E_2, \ldots, E_{N_+}$, so $(E_1, \ldots, E_{N_+}) \in \mathcal{E}$. (Condition (4.5) holds because the two outgoing edges of v_1 have positive multiplicity.) Now T_1 and $\overline{T}$ here are the same as in the definition of ϕ, so $T = \phi(E_1, \ldots, E_{N_+})$. $\square$

5 Counting zeroes of the linearized section

Continue with the assumptions from the first paragraph of §3. We will now show that the relative Euler class $\#\mathfrak{s}_0^{-1}(0)$ defined in §3.2 agrees with the combinatorial quantity $c_\theta(S)$ defined in §1.6, and also that the latter does not depend on the choice of ordering in Definition 1.22 or 1.25.

To start, by Corollary 3.10 we may assume that $\overline{N}_+ = N_+$ or $\overline{N}_- = N_-$. Next, observe that $\#\mathfrak{s}_0^{-1}(0)$ does not depend on the ordering of the a_i's for $i \in \{1, \ldots, \overline{N}_+\}$ or on the ordering of the a_i's for $i \in \{-1, \ldots, -\overline{N}_-\}$. Let us then choose these two orderings such that, as in (1.15), we have:

- If $0 < i < j \le \overline{N}_+$, then $\lceil a_i\theta \rceil / a_i \le \lceil a_j\theta \rceil / a_j$.
- If $0 > i > j \ge -\overline{N}_-$, then $\lfloor a_i\theta \rfloor / a_i \ge \lfloor a_j\theta \rfloor / a_j$.

Then by Definitions 1.22 and 1.25, to prove that $\#\mathfrak{s}_0^{-1}(0) = c_\theta(S)$ and that the latter is well-defined, it suffices to prove:

Proposition 5.1. *Under the above assumptions,*

$$\#\mathfrak{s}_0^{-1}(0) = f_\theta(S). \tag{5.1}$$

The rest of this section proves Proposition 5.1. As explained in §1.8, this will complete the proof of Theorem 1.13, assuming the results from [13] that are quoted in §3.

5.1 Setup for the proof of Proposition 5.1

To start, note that the statement of Proposition 5.1 is symmetric under switching positive ends with negative ends and replacing θ by $-\theta$. (The symmetry for $f_\theta(S)$ holds by Corollary 4.6, while the symmetry for $\#\mathfrak{s}_0^{-1}(0)$ is a straightforward consequence of the definitions.) By this symmetry, we can assume that

$$\overline{N}_- = N_-.$$

Now fix admissible $\gamma = \{\gamma_i\}$, r, and R for use in the definition of $\mathfrak{s}_0$. (We will be more particular about these choices later.) By Corollary 3.6, for the purposes of computing $\#\mathfrak{s}_0^{-1}(0)$ we may assume that

$$S(t) = \theta. \tag{5.2}$$

Then the operator D_Σ is complex linear, and $\operatorname{Coker}(D_\Sigma)$ is a complex vector space. Let $\mathcal{O}_\mathbb{C} \to \mathcal{M}_R$ denote the complex vector bundle whose fiber over Σ is $\operatorname{Hom}_\mathbb{C}(\operatorname{Coker}(D_\Sigma), \mathbb{C})$. There is a natural identification of real vector bundles

$$\mathcal{O} = \mathcal{O}_\mathbb{C}. \tag{5.3}$$

Under this identification, the section $\mathfrak{s}_0$ of $\mathcal{O}$ corresponds to a section $\mathfrak{s}_\mathbb{C}$ of $\mathcal{O}_\mathbb{C}$ defined by

$$\mathfrak{s}_\mathbb{C}(\Sigma)(\sigma) := \mathfrak{s}_0(\Sigma)(\sigma) + i\mathfrak{s}_0(\Sigma)(-i\sigma).$$

Equivalently, $\mathfrak{s}_\mathbb{C}$ is defined as in Definition 3.2, but with real inner products replaced by complex inner products. Under the identification (5.3), the orientation of $\mathcal{O}_\mathbb{C}$ as a complex vector bundle differs from the orientation of $\mathcal{O}$ defined in §2.6 by $(-1)^{\operatorname{rank}_\mathbb{C}(\mathcal{O}_\mathbb{C})}$. Thus

$$\#\mathfrak{s}_0^{-1}(0) = (-1)^N \#\mathfrak{s}_\mathbb{C}^{-1}(0). \tag{5.4}$$

We will now compute $\#\mathfrak{s}_\mathbb{C}^{-1}(0)$. To describe $\mathfrak{s}_\mathbb{C}$ more concretely, choose an isomorphism of each $\mathcal{B}_i$ with $\mathbb{C}$ as a complex vector space, and use these isomorphisms to regard the γ_i's as complex numbers. If $\Sigma \in \mathcal{M}_R$ and $\sigma \in$

$\mathrm{Coker}(D_\Sigma)$, then for each i in $\{1,\ldots,N_+\}$ or $\{-1,\ldots,-N_-\}$ labeling an end of Σ, the projection of $\sigma_i(\pm(R+r),\cdot)$ to $\mathcal{B}_i$ (where $\pm$ denotes the sign of i) also corresponds to a complex number, which we denote simply by σ_i. In this notation,

$$\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \sum_{i=1}^{N_+} \overline{\gamma_i}\sigma_i - \sum_{i=-1}^{N_-} \overline{\gamma_i}\sigma_i \in \mathbb{C}. \tag{5.5}$$

Here we interpret $\gamma_{N_+} = 0$ when $\overline{N}_+ \neq N_+$.

5.2 Outline of the argument

The relative Euler class $\#\mathfrak{s}_{\mathbb{C}}^{-1}(0)$ is determined by the restriction of $\mathfrak{s}_{\mathbb{C}}$ to the boundary of $\mathcal{M}_R$. The prototypical fact is that given a generic smooth function $f : D^2 \to \mathbb{C}$ which does not vanish on ∂D^2, the algebraic count of points $x \in D^2$ with $f(x) = 0$ is equal to a count of points $x \in \partial D^2$ with $f(x) > 0$. Roughly speaking, our strategy for computing $\#\mathfrak{s}_{\mathbb{C}}^{-1}(0)$ is to understand the relevant boundary behavior by induction on the dimension of boundary strata.

The precise procedure is as follows. We will choose a large constant $r_1 > 0$, and assume that the constant r in §3.2 is chosen sufficiently large that $r > Nr_1$. If k is a positive integer, define

$$R_k := R - kr_1.$$

Definition 5.2. For $k = 0,\ldots,N-2$, define $\mathcal{M}^k$ to be the set of $\Sigma \in \mathcal{M}_R$ such that:

- The tree $\tau(\Sigma)$ has trivalent vertices $v_1,\ldots,v_k$ with $\rho(v_i) \in \{\pm R_i\}$.
- All other vertices w of $\tau(\Sigma)$ have $\rho(w) \in [-R_{k+1}, R_{k+1}]$.

Let $\partial\mathcal{M}^k$ denote the set of $\Sigma \in \mathcal{M}^k$ such that there is at least one vertex w with $\rho(w) \in \{\pm R_{k+1}\}$.

Note that the interior of $\mathcal{M}^k$ is a smooth manifold of dimension $2N - 4 - k$. Also $\mathcal{M}^k$ is contained in the interior of $\partial\mathcal{M}^{k-1}$, and this inductively determines an orientation of the interior of $\mathcal{M}^k$. We will later define a smaller space $\mathcal{N}^k$ obtained by discarding certain components from $\mathcal{M}^k$. This will satisfy $\mathcal{N}^0 = \mathcal{M}^0$ and $\mathcal{N}^k \subset \partial\mathcal{N}^{k-1} := \mathcal{N}^{k-1} \cap \partial\mathcal{M}^{k-1}$. Next, for each $k = 1,\ldots,N-2$ and each component of $\mathcal{N}^{k-1}$, we will pick a suitable $\epsilon_k \in \{1,\ldots,N_+\} \cup \{-1,\ldots,-N_-\}$ labeling an end of the Σ's. Then to each component of $\mathcal{N}^{k-1}$ we will have associated k ends $\epsilon_1,\ldots,\epsilon_k$ (since $\mathcal{N}^{k-1} \subset \mathcal{N}^i$ for $i < k$), and these will be chosen to be distinct.

Definition 5.3. Let n_k denote the algebraic count of points $\Sigma \in \mathcal{N}^k$ such that there exist constants $\Lambda_1, \ldots, \Lambda_k > 0$ satisfying the *k-boundary equation*

$$\forall \sigma \in \operatorname{Coker}(D_\Sigma) : \quad \mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_1 \sigma_{\epsilon_1} + \cdots + \Lambda_k \sigma_{\epsilon_k}. \tag{5.6}$$

To be more precise, n_k is defined by making a small perturbation of $\mathfrak{s}_{\mathbb{C}}$ to a section $\mathfrak{s}'_{\mathbb{C}}$ on $\mathcal{N}^k$ so that all solutions to the $\mathfrak{s}'_{\mathbb{C}}$ analogue of (5.6) on $\mathcal{N}^k$ with $\Lambda_1, \ldots, \Lambda_k \geq 0$ have $\Lambda_1, \ldots, \Lambda_k > 0$ and are cut out transversely, and then counting these solutions with signs. We will now specify the sign convention and then explain why the count does not depend on the perturbation.

If $\epsilon_{k+1}, \ldots, \epsilon_N$ are the remaining ends in any order, then a complex basis for $\operatorname{Coker}(D_\Sigma)$ is given by $(\sigma^{(1)}, \ldots, \sigma^{(N-2)})$, where $\sigma^{(i)}$ denotes the special cokernel element satisfying

$$\sigma^{(i)} \in V_{\epsilon_i, \epsilon_{N-1}, \epsilon_N}, \qquad \sigma^{(i)}_{\epsilon_i} = 1. \tag{5.7}$$

The k-boundary equation (5.6) for the perturbed section $\mathfrak{s}'_{\mathbb{C}}$ is then equivalent to the open condition

$$\mathfrak{s}'_{\mathbb{C}}(\Sigma)(\sigma^{(i)}) \neq 0, \qquad i = 1, \ldots, k$$

together with the equations

$$\begin{aligned} \arg\left(\mathfrak{s}'_{\mathbb{C}}(\Sigma)\left(\sigma^{(i)}\right)\right) &= 0, \qquad i = 1, \ldots, k, \\ \mathfrak{s}'_{\mathbb{C}}(\Sigma)(\sigma^{(i)}) &= 0, \qquad i = k+1, \ldots, N-2. \end{aligned} \tag{5.8}$$

Writing the equations in this order determines the sign convention for n_k.

Now n_0 is well-defined and equal to the integer that we want to compute, namely

$$n_0 = \#\mathfrak{s}_{\mathbb{C}}^{-1}(0), \tag{5.9}$$

because Lemma 3.5 guarantees that all zeroes of $\mathfrak{s}_{\mathbb{C}}$ over $\mathcal{M}_R$ are in the interior of $\mathcal{N}^0$, and the sign conventions for counting agree. The following lemma provides an inductive strategy for computing n_0.

Lemma 5.4. *Suppose that the following hold for all $k = 1, \ldots, N-1$:*

- *(Ind1) If $\Sigma \in \partial \mathcal{N}^{k-1}$ solves the k-boundary equation with $\Lambda_1, \ldots, \Lambda_{k-1} > 0$ and $\Lambda_k \geq 0$, then $\Sigma \in \mathcal{N}^k$.*
- *(Ind2) If $\Sigma \in \mathcal{N}^{k-1}$ solves the k-boundary equation[2] with $\Lambda_1, \ldots, \Lambda_k \geq 0$, then $\Lambda_1, \ldots, \Lambda_{k-1} > 0$.*

[2] When $k = N-1$, statement (Ind1) is vacuously true since $\partial\mathcal{N}^{N-2} = \emptyset$, while statement (Ind2) is to be interpreted as saying that if $\Sigma \in \mathcal{N}^{N-2}$ satisfies the $(N-2)$-boundary equation with $\Lambda_1, \ldots, \Lambda_{N-2} \geq 0$, then $\Lambda_1, \ldots \Lambda_{N-2} > 0$.

Then for all $k = 1, \dots, N-2$:

(a) n_k is well-defined, independent of the small perturbation $\mathfrak{s}'_{\mathbb{C}}$ of $\mathfrak{s}_{\mathbb{C}}$.

(b) $n_k = (-1)^{k-1} n_{k-1}$.

Proof. First note that for all $k = 1, \dots, N-2$, by combining statement (Ind1) for k with statement (Ind2) for $k+1$, we have:

(Ind1′) If $\Sigma \in \partial\mathcal{N}^{k-1}$ solves the k-boundary equation with $\Lambda_1, \dots, \Lambda_{k-1} > 0$ and $\Lambda_k \geq 0$, then $\Sigma \in \mathcal{N}^k$ and $\Lambda_k > 0$.

(a) To see that n_k is well-defined, we need to show that if $\Sigma \in \mathcal{N}^k$ solves the k-boundary equation with $\Lambda_1, \dots, \Lambda_k \geq 0$, then (i) $\Lambda_1, \dots, \Lambda_k > 0$ and (ii) $\Sigma \notin \partial\mathcal{N}^k$. Assertion (i) follows from statement (Ind2) for $k+1$. Assertion (ii) then follows from statement (Ind1′) for $k+1$.

(b) Consider the set

$$Z := \{\Sigma \in \mathcal{N}^{k-1} \mid \Sigma \text{ solves the } k\text{-boundary equation with } \Lambda_1, \dots, \Lambda_k \geq 0\}.$$

Conditions (Ind2) and (Ind1′) assert that:

- Every $\Sigma \in Z$ has $\Lambda_1, \dots, \Lambda_{k-1} > 0$.
- If $\Sigma \in Z \cap \partial\mathcal{N}^{k-1}$, then $\Sigma \in \mathcal{N}^k$ (whence $\Sigma \in \operatorname{int}(\partial\mathcal{N}^{k-1})$) and $\Lambda_k > 0$.

It follows that we can choose the small perturbation $\mathfrak{s}'_{\mathbb{C}}$ of the section $\mathfrak{s}_{\mathbb{C}}$ over $\mathcal{N}^{k-1}$ to arrange not only that the points counted by n_{k-1} and n_k are cut out transversely, but also that the $\mathfrak{s}'_{\mathbb{C}}$ version of Z, call it Z', is a one-manifold with boundary

$$\partial Z' = (Z' \cap \mathcal{N}^k) \bigsqcup \{\Sigma \in Z' \mid \Lambda_k = 0\}.$$

Also Z' is compact by Lemma 2.8. Thus Z' is a cobordism between the set of solutions to the perturbed k-boundary equation on $\mathcal{N}^k$ with $\Lambda_1, \dots, \Lambda_k > 0$, and the set of solutions to the perturbed $(k-1)$-boundary equation on $\mathcal{N}^{k-1}$ with $\Lambda_1, \dots, \Lambda_{k-1} > 0$. After an orientation check it follows that $n_k = (-1)^{k-1} n_{k-1}$. □

We will see that if the ϵ_k's and $\mathcal{N}^k$'s and the various constants are chosen carefully, then points (Ind1) and (Ind2) hold for each k. We will then be reduced to the problem of computing n_{N-2}, i.e. counting solutions to the $(N-2)$-boundary equation on $\mathcal{N}^{N-2}$. Since every $\Sigma \in \mathcal{N}^{N-2}$ has a trivalent tree $\tau(\Sigma)$, it follows that $\mathcal{N}^{N-2}$ is a union of $(N-2)$-dimensional tori (cf. Lemma 2.6), and counting the solutions to the equations (5.8) will reduce to a determinant calculation.

5.3 Decay estimates

To work with the k-boundary equation (5.6), we need a preliminary discussion of the relative sizes of the different contributions to $\mathfrak{s}_{\mathbb{C}}(C)(\sigma)$ in equation (5.5). By Corollary 3.6, we can choose the γ_i's such that

$$\begin{aligned} |\gamma_i| &> r_2|\gamma_{i+1}|, \qquad i > 0, \\ |\gamma_j| &> r_2|\gamma_{j-1}|, \qquad j < 0, \end{aligned} \tag{5.10}$$

where $r_2 > 1$ is a large constant. We can also assume that

$$i \in \{1, \ldots, \overline{N}_+\} \cup \{-1, \ldots, -N_-\} \Longrightarrow \gamma_i \notin \mathbb{R}. \tag{5.11}$$

Let i, j, and k be distinct ends. Recall from (5.5) that for a special cokernel element $\sigma \in V_{i,j,k}$, we have $\sigma_l = 0$ for $l \notin \{i, j, k\}$, so only three terms contribute to $\mathfrak{s}_{\mathbb{C}}(C)(\sigma)$. Often one term dominates the other two, in the following sense. By (5.4) and (5.11), if K is sufficiently large, and if nonzero $\sigma \in V_{i,j,k}$ satisfy $|\overline{\gamma_i}\sigma_i| > K|\overline{\gamma_j}\sigma_j|, K|\overline{\gamma_k}\sigma_k|$, then

$$\mathfrak{s}_{\mathbb{C}}(C)(\sigma) \neq 0, \tag{5.12}$$
$$\arg(\mathfrak{s}_{\mathbb{C}}(C)(\sigma)) \neq \arg(\sigma_i). \tag{5.13}$$

Definition 5.5. Write
$$i \searrow j, k$$
if nonzero $\sigma \in V_{i,j,k}$ satisfy (5.12) and (5.13).

The following lemma will be used repeatedly.

Lemma 5.6. *If r is sufficiently large, if r_2 is sufficiently large with respect to r_1, and if R is sufficiently large with respect to all other choices, then the following holds: Let i be a positive end and let $j > j'$ be negative ends. Let v denote the central vertex for i, j, and j'. Then:*

(a) If the path $P_{v,i}$ stays above the level $\rho = R_N$ and if $i \leq \overline{N}_+$, then $i \searrow j, j'$.

(b) If the path $P_{v,j}$ stays below the level $\rho = -R_N$ (e.g. if the path $P_{j,j'}$ does), then $j \searrow i, j'$.

Proof. We begin with a key estimate. Let $\sigma \in V_{i,j,j'}$ be normalized so that $|\sigma(z)| = 1$ for some $z \in \Sigma$ with $p(z) = v$. Let $x_i \in \tau(\Sigma)$ denote the point on the edge corresponding to the i^{th} end for which $\rho(x_i) = R + r$. Let $x_j \in \tau(\Sigma)$ denote the point on the edge corresponding to the j^{th} end for which $\rho(x_j) = -(R + r)$, and define $x_{j'}$ likewise. By Lemma 2.20, Corollary 2.22, and Proposition 2.25, if r is sufficiently large then there is a constant r' such that

$$\left|\log|\sigma_i| + \sum_{e\in P^+_{v,x_i}} \ell(e)\left(\frac{\lceil m(e)\theta\rceil}{m(e)} - \theta\right) + \sum_{e\in P^-_{v,x_i}} \ell(e)\left(\theta - \frac{\lfloor m(e)\theta\rfloor}{m(e)}\right)\right| \le r'. \tag{5.14}$$

The estimate (5.14) also holds if i is replaced by j or j'.

(a) If the path $P_{v,i}$ stays above the level $\rho = R_N$, then the estimate (5.14) implies that there is a constant $\kappa > 0$ such that

$$\log|\sigma_i| \ge -r\left(\frac{\lceil a_i\theta\rceil}{a_i} - \theta\right) - \kappa r_1 - r'.$$

The analogue of (5.14) for j and j' implies that the constant κ can be chosen so that

$$\log|\sigma_j|, \log|\sigma_{j'}| \le -\kappa R + r'.$$

By the above two inequalities, assertion (a) holds provided that R is sufficiently large with respect to all of the other choices.

(b) If the path $P_{v,j}$ stays below the level $\rho = -R_N$, then (5.14) and its analogues for j and j' imply that there is a constant $\kappa > 0$ such that

$$\begin{aligned}
\log|\sigma_i| &\le -\kappa R + r',\\
\log|\sigma_j| &\ge -r\left(\theta - \frac{\lfloor a_j\theta\rfloor}{a_j}\right) - \kappa r_1 - r',\\
\log|\sigma_{j'}| &\le -r\left(\theta - \frac{\lfloor a_{j'}\theta\rfloor}{a_{j'}}\right) + r'.
\end{aligned}$$

Recall that our ordering convention gives $\lfloor a_j\theta\rfloor / a_j \ge \lfloor a_{j'}\theta\rfloor / a_{j'}$. So by (5.10), assertion (b) holds provided that r_2 is large enough with respect to r_1 and r', and R is large enough with respect to all other choices. $\square$

The obvious symmetric analogue of Lemma 5.6 with positive and negative ends switched also holds. Henceforth assume that the constants are chosen so that the conclusions of Lemma 5.6 and its symmetric analogue hold. (We will later need to choose r_1 large.)

5.4 Processing the positive ends

To begin the inductive process, we now define ϵ_k and $\mathcal{N}^k$ when $k < N_+$ and verify that the crucial properties (Ind1) and (Ind2) hold in this case.

When $k < N_+$, we choose ϵ_k to be the positive end labeled by k.

In the definition of $\mathcal{N}^k$ and below, we will use the following notation. If v is a vertex of a tree and e is an edge incident to v, let $A(v,e)$ denote the set of ends that are accessible via paths starting from v along the edge e. Also, if there is a unique downward path from v to a negative leaf, then we denote the corresponding negative leaf by v^-. We generally refer to the leaves of a tree $\tau(\Sigma)$ as "ends", and identify the ends with their labels in $\{1,\ldots,N_+\} \cup \{-1,\ldots,-N_-\}$.

Definition 5.7. For $k < N_+$, define $\mathcal{N}^k$ to be the set of $\Sigma \in \mathcal{M}^k$ that satisfy the following conditions for all $i = 1,\ldots,k$:

(a) v_i is a trivalent splitting vertex with $\rho(v_i) = -R_i$.

(b) v_i has an outgoing edge e_i^0 such that $\{i+1,\ldots,N_+\} \subset A(v_i, e_i^0)$.

(c) If j is a negative end and $j \notin A(v_i, e_i^0)$, then $j \geq v_i^-$.

Let e_i^+ denote the outgoing edge of v_i other than e_i^0, and let e_i^- denote the incoming edge of v_i.

To better the above definition, we now consider the following additional structure associated to elements of $\mathcal{N}^k$.

Definition 5.8. Let $\Sigma \in \mathcal{N}^k$ with $k < N_+$. For each $i = 0,\ldots,k$ define a tree τ_i^+ and a forest τ_i^- inductively as follows.

- $\tau_0^+ = \tau(\Sigma)$.
- For $i = 1,\ldots,k$, the tree τ_i^+ is obtained from the tree τ_{i-1}^+ by cutting along the edge e_i^0 and keeping the half that contains the positive ends $i+1,\ldots,N_+$.
- τ_i^- is the complement of τ_i^+ in $\tau(\Sigma)$.

Lemma 5.9. *Let $\Sigma \in \mathcal{N}^k$ with $k < N_+$. Then for each $i = 1,\ldots,k$:*

(a) The forest τ_i^- contains the vertices $v_1,\ldots,v_i$, and no other splitting vertices.

(b) The positive leaves of τ_i^- are the first i positive ends, together with, for each component of τ_i^-, a positive leaf where the component of τ_i^- is attached to τ_i^+.

(c) An upward path in $\tau(\Sigma)$ from e_i^+ to a positive end must terminate at the i^{th} end.

(d) If there is a downward path from the end i to the end j, then $j \geq v_i^-$.

Proof. Let s, p, and m denote the numbers of splitting vertices, positive leaves, and components respectively in τ_i^-. By construction, τ_i^- contains the splitting vertices $v_1, \ldots, v_i$, so $s \geq i$. Also, the positive leaves of τ_i^- consist of one positive leaf in each component where it attaches to τ_i^+, together with some subset of the first i positive ends. Thus $p \leq m + i$. But since τ_i^- contains no loops, we have $p \geq m + s$. Therefore $s = i$ and $p = i + m$, and these facts prove parts (a) and (b) respectively of the lemma.

To prove (c), note that the path under consideration stays in $\tau_i^- \setminus \tau_{i-1}^-$. We are then done by part (b).

To prove part (d), note that the downward path from i to j intersects the upward path from v_i to i. By part (c), the latter path does not contain e_i^0. It then follows that $j \notin A(v_i, e_i^0)$, so we are done by Definition 5.7(c). □

Lemma 5.10. *If $k < N_+$, then statements (Ind1) and (Ind2) hold.*

Proof. (Ind1) Suppose $\Sigma \in \partial\mathcal{N}^{k-1}$ satisfies the k-boundary equation (5.6) with $\Lambda_1, \ldots, \Lambda_k > 0$. We need to show that $\Sigma \in \mathcal{N}^k$. Since we already know that $\Sigma \in \mathcal{N}^{k-1}$, we just need to verify conditions (a)–(c) in Definition 5.7 for $i = k$, and we also need to check that $\Sigma \in \mathcal{M}^k$. We proceed in four steps.

Step 1. We first show that every vertex $v \in \tau(\Sigma)$ has $\rho(v) < R_N$.

Let v be a vertex with $\rho(v) \geq R_N$. We can assume that $\rho(v)$ is maximal. Suppose first that v is a splitting vertex. Then v has (at least) two outgoing edges incident to positive ends $i < i'$. Since $\Sigma \in \mathcal{N}^{k-1}$, Lemma 5.9 implies that $i, i' \geq k$. We can also find a downward path from v to a negative end j, so that v is the central vertex for i, i', and j. By the symmetric analogue of Lemma 5.6(b), we have $i \searrow i', j$. Let $0 \neq \sigma \in V_{i,i',j}$. If $i > k$, then the k-boundary equation asserts that $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$, contradicting (5.12). If $i = k$ then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_i \sigma_i$, contradicting (5.13).

The remaining possibility is that v is a joining vertex with one outgoing edge incident to a positive end i. Then similarly to the proof of Lemma 3.8, the partition minimality assumption implies that $\overline{N}_+ = N_+$. We can find downward paths from v, starting along distinct edges, to negative ends j and j'. Then v is the central vertex for i, j, and j'. Lemma 5.6(a) implies that $i \searrow j, j'$. Let $0 \neq \sigma \in V_{i,j,j'}$. If $i > k$, then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$, contradicting (5.12). If $i \leq k$, then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_i \sigma_i$, which contradicts (5.13).

Step 2. We now show that any vertex v with $\rho(v) = -R_k$ is a trivalent splitting vertex.

Suppose first that v has (at least) two incoming edges, and let $j_1 > j_2$ be negative ends reached by downward paths starting along these two edges. Lemma 5.6(b) then gives $j_1 \searrow j_2, N_+$. But if $0 \neq \sigma \in V_{j_1,j_2,N_+}$, then the k-boundary equation implies that $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$, which is a contradiction.

So v has only one incoming edge. In particular v is a splitting vertex, so v cannot be in the forest τ_{k-1}^- by Lemma 5.9(a). Thus any upward path starting at v stays in τ_{k-1}^+, and hence by Lemma 5.9(b) terminates at a positive end indexed by $k, \ldots, N_+$. If v has more than two outgoing edges, then at least two of these outgoing edges lead to positive ends $i_1, i_2 > k$. If $\sigma \in V_{i_1,i_2,v^-}$, then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$, while the symmetric analogue of Lemma 5.6(a) gives $v^- \searrow i_1, i_2$, which is a contradiction.

Step 3. Let v_k be a trivalent splitting vertex with $\rho(v_k) = -R_k$. We now show that v_k is unique and satisfies conditions (b) and (c) in Definition 5.7.

To prove (b), let e_k^0 and e_k^+ denote the outgoing edges of v_k. The sets $A(v_k, e_k^0)$ and $A(v_k, e_k^+)$ cannot both contain positive ends that are greater than k, or else we obtain a contradiction as in Step 2. So without loss of generality, $A(v_k, e_k^+)$ does not contain any positive ends indexed by $i > k$. Since $\Sigma \in \mathcal{N}^{k-1}$, the incoming edge e_k^- of v_k either comes out of the forest τ_{k-1}^- or is incident to a negative end. Hence $A(v_k, e_k^-)$ does not contain any positive ends indexed by $i > k$. Therefore all of the positive ends indexed by $i > k$ must be contained in $A(v_k, e_k^0)$.

To prove condition (c) in Definition 5.7, suppose $A(v_k, e_k^+) \cup A(v_k, e_k^-)$ contains a negative end indexed by j with $v_k^- > j$. Since $N_+ \in A(v_k, e_k^0)$, the central vertex for j, v_k^-, and N_+ is on the downward path from v_k to v_k^-. Lemma 5.6(b) then gives $v_k^- \searrow j, N_+$. But if $0 \neq \sigma \in V_{j, v_k^-, N_+}$ then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$, which is a contradiction.

To prove that v_k is unique, note that Step 2 and Lemma 5.9(a) imply that $v_k \in \tau_{k-1}^+$, so by Lemma 5.9(b) there is a unique upward path P starting along e_k^+, and the path P leads to the k^{th} positive end. Now suppose that w is another trivalent splitting vertex with $\rho(w) = -R_k$. Then w must also have an outgoing edge e such that $A(w, e)$ contains no positive ends indexed by $i > k$, there is a unique upward path P' starting along e, and the path P' leads to the k^{th} positive end. The two upward paths P and P' must intersect. By proceeding from w along P' to its intersection with P, and then backwarrds along P to v_k, we find that $N_+ \in A(w, e)$, which is a contradiction.

Step 4. To complete the proof that $\Sigma \in \mathcal{N}^k$, we must check that $\Sigma \in \mathcal{M}^k$, i.e. that any vertex w other than $v_1, \ldots, v_k$ satisfies $\rho(w) \in [-R_{k+1}, R_{k+1}]$.

We know from Steps 1–3 that $\rho(w) \in (-R_k, R_N)$. Suppose to get a contradiction that $\rho(w) \in (-R_k, -R_{k+1})$. We can assume that $\rho(w)$ is minimal. If w has more than one incoming edge, then we get a contradiction as in Step 2. So w has (at least) two outgoing edges. By Lemma 5.9(a) we know that w is in τ_k^+ (which is well-defined by Steps 1–3), so upward paths starting along these outgoing edges lead to positive ends with labels in $k+1, \ldots, N_+$. This again gives a contradiction as in Step 2.

(Ind2) Suppose that $\Sigma \in \mathcal{N}^{k-1}$ solves the k-boundary equation with $\Lambda_1, \ldots, \Lambda_k \geq 0$. We need to show that $\Lambda_1, \ldots, \Lambda_{k-1} \neq 0$. Given $i \in \{1, \ldots, k-1\}$, let $0 \neq \sigma \in V_{i,N_+,v_i^-}$. Observe that v_i is the central vertex for i, N_+, and v_i^-. Hence the symmetric analogue of Lemma 5.6(a) implies that $v_i^- \searrow i, N_+$, so $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) \neq 0$. However the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_i \sigma_i$, whence $\Lambda_i \neq 0$. □

5.5 Processing the negative ends

We now define ϵ_k and $\mathcal{N}^k$ for $k = N_+, \ldots, N-2$. Here ϵ_k depends on the component of $\mathcal{N}^{N_+-1}$. We will then prove that conditions (Ind1) and (Ind2) continue to hold.

Given $\Sigma \in \mathcal{N}^{N_+-1}$, for each $i = 1, \ldots, N_+$ let E_i denote the set of negative ends that are accessible by downward paths in the tree $\tau(\Sigma)$ starting from the i^{th} positive end. By Lemma 5.9(d), if $i < N_+$ then v_i^- is the smallest element of the set E_i. For $i = 1, \ldots, N_+$ define $E_i' := E_i \setminus \{v_i^-\}$, where we interpret $v_{N_+}^-$ to be the smallest element of the set E_{N_+}.

Lemma 5.11. *For each $\Sigma \in \mathcal{N}^{N_+-1}$, the following hold:*

(a) $\{-1, \ldots, -N_-\} = E_{N_+} \sqcup \bigsqcup_{i=1}^{N_+-1} E_i'$.

(b) $v_{N_+}^- = -N_-$.

(c) $(E_1, \ldots, E_{N_+})$ *depends only on the component of* $\mathcal{N}^{N_+-1}$ *containing* Σ.

Proof. (a) Let j be a negative end; we need to show that there is a unique positive end i such that

(i) $j \in E_i'$ if $i < N_+$, and $j \in E_{N_+}$ if $i = N_+$.

Note that condition (i) is equivalent to

(ii) The path P from j to i is an upward path, and

> (*) for all $k \in \{1, \ldots, N_+ - 1\}$, if the path P meets the vertex v_k, then the path P continues along the edge e_k^0 in Definition 5.7.

The reason is that by Lemma 5.9(c), condition (ii) fails if and only if $j \notin E_i$ or there exists $k \in \{1, \ldots, N_+ - 1\}$ such that $j = v_k^-$ and $i = k$. But there is a unique upward path P starting at j and satisfying condition (*), because every vertex other than $v_1, \ldots, v_{N_+-1}$ is a joining vertex.

(b) We need to show that $-N_- \in E_{N_+}$. If not, then part (a) implies that $-N_- \in E_i'$ for some $i \in \{1, \ldots, N_+ - 1\}$. But then $-N_- > v_i^-$, which is impossible.

(c) The sets E_i can be characterized in terms of which ends are accessible from which edges incident to the vertices $v_1, \ldots, v_{N_+-1}$. The latter information depends only on the component of $\mathcal{N}^{N_+-1}$. □

Definition 5.12. For a given component of $\mathcal{N}^{N_+-1}$, define the sequence $\epsilon_{N_+}, \ldots, \epsilon_{N-2}$ by first listing the ends in E'_1 in decreasing order, then listing the ends in E'_2 in decreasing order, and so on up to E'_{N_+}.

Thus the two remaining ends are the positive end N_+ and the negative end $-N_-$; we denote these by ϵ_{N-1} and ϵ_N respectively.

Definition 5.13. For $k = N_+, \ldots, N-2$, define $\mathcal{N}^k$ to be the set of $\Sigma \in \mathcal{N}^{k-1} \cap \mathcal{M}^k$ such that:

(a) v_k is a trivalent joining vertex.

(b) For one of the incoming edges of v_k, call it e_k^-, there is a unique downward path starting along e_k^-, and this leads to the negative end ϵ_k.

(c) $\rho(v_k) = +R_k$.

Let e_k^0 denote the incoming edge of v_k other than e_k^-, and let e_k^+ denote the outgoing edge of v_k.

Lemma 5.14. *For $k = N_+, \ldots, N-1$, statements (Ind1) and (Ind2) hold.*

Proof. (Ind1) Recall that this is vacuous when $k = N-1$. Now given $k \in \{N_+, \ldots, N-2\}$, suppose $\Sigma \in \partial\mathcal{N}^{k-1}$ satisfies the k-boundary equation (5.6) with $\Lambda_1, \ldots, \Lambda_{k-1} > 0$ and $\Lambda_k \geq 0$. We need to show that $\Sigma \in \mathcal{N}^k$. We proceed in three steps.

Step 1. Let v_k be a vertex with $\rho(v_k) = \pm R_k$. We now show that v_k is unique and satisfies conditions (a) and (b) in Definition 5.13.

To start, the tree $\tau(\Sigma)$ contains at most $N_+ - 1$ splitting vertices, and these are accounted for by $v_1, \ldots, v_{N_+-1}$. Since $k \geq N_+$, it follows that v_k is a joining vertex with only one outgoing edge. Since v_k is above all of the splitting vertices, there is a unique upward path starting from v_k. Since $\Sigma \in \mathcal{N}^{k-1}$, a downward path starting at v_k cannot lead to an end in the set $\{\epsilon_1, \ldots, \epsilon_{k-1}\}$.

Now suppose that (a) or (b) fails. Then there are downward paths starting from v_k along distinct incoming edges, leading to negative ends $j > j'$ not in the set $\{\epsilon_1, \ldots, \epsilon_k\}$.

To get a contradiction, suppose first that $\rho(v_k) = -R_k$. In this case Lemma 5.6(b) gives $j \searrow j', N_+$. On the other hand, if $0 \neq \sigma \in V_{j,j',N_+}$ then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$, which contradicts (5.12).

Suppose next that $\rho(v_k) = +R_k$. Let i denote the positive end reached by the unique upward path starting from v_k. Let $0 \neq \sigma \in V_{i,j,j'}$. Note that v_k is the central vertex for i, j, and j'. If $i < N_+$, then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_i \sigma_i$, while Lemma 5.6(a) gives $i \searrow j, j'$. This contradicts (5.13). If $i = N_+$, then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = 0$. However the partition minimality assumption guarantees that $\overline{N}_+ = N_+$ here, so Lemma 5.6(a) applies again to give $i \searrow j, j'$, which contradicts (5.12). This completes the proof of (a) and (b).

To prove uniqueness of v_k, recall that there is a unique upward path from v_k, and there is a unique downward path starting along the incoming edge e_k^- of v_k. These paths lead respectively to the positive end i for which $\epsilon_k \in E_i'$, and to the negative end ϵ_k. If w is another vertex with these properties, then the downward paths meet at some vertex other than v_k or w (by uniqueness of these paths, since v_k and w are joining vertices). Then the downward paths and the upward paths together contain a loop in $\tau(\Sigma)$, which is a contradiction.

Step 2. We now show that v_k satisfies condition (c) in Definition 5.13. Suppose to the contrary that $\rho(v_k) = -R_k$. Choose a downward path from v_k starting along the incoming edge e_k^0; this leads to a negative end j with $\epsilon_k > j$. Then Lemma 5.6(b) gives $\epsilon_k \searrow j, N_+$. But if $0 \neq \sigma \in V_{N_+,\epsilon_k,j}$ then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_k \sigma_{\epsilon_k}$. If $\Lambda_k = 0$ then this contradicts (5.12), while if $\Lambda_k > 0$ then this contradicts (5.13).

Step 3. To complete the proof that $\Sigma \in \mathcal{N}^k$, we must now show that any vertex w other than $v_1, \ldots, v_k$ has $\rho(w) \in (-R_{k+1}, R_{k+1})$. The proof of this is essentially the same as the proof that v_k is unique.

(Ind2) Given $k \in \{N_+, \ldots, N-1\}$, suppose that $\Sigma \in \mathcal{N}^{k-1}$ solves the k-boundary equation with $\Lambda_1, \ldots, \Lambda_k \geq 0$. (When $k = N-1$, the hypothesis is that $\Sigma \in \mathcal{N}^{N-2}$ satisfies the $(N-2)$-boundary equation with $\Lambda_1, \ldots, \Lambda_{N-2} \geq 0$.) Let $j \in \{1, \ldots, k-1\}$; we must show that $\Lambda_j \neq 0$. There are two cases.

Case 1: $\epsilon_j \in E'_{N_+}$. Since $\Sigma \in \mathcal{N}^j$, Lemma 5.6(a) implies that $\epsilon_{N-1} \searrow \epsilon_j, \epsilon_N$, so if $0 \neq \sigma \in V_{\epsilon_j,\epsilon_{N-1},\epsilon_N}$ then $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) \neq 0$. But the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_j \sigma_{\epsilon_j}$ whence $\Lambda_j \neq 0$.

Case 2: $\epsilon_j \notin E'_{N_+}$. Then there exists $i \in \{1, \ldots, N_+ - 1\}$ such that either $j = i$ or $\epsilon_j \in E_i'$. Observe that v_i is the central vertex for j, v_i^-, and N_+. The path from v_i to v_i^- stays below the level $\rho = -R_N$, while the paths from v_i to j and N_+ go above the level $\rho = +R_N$. It then follows from the decay estimate (5.14) that $v_i^- \searrow j, N_+$. Let $0 \neq \sigma \in V_{j,v_i^-,N_+}$. There are now two subcases.

Case 2a: $v_i^- \notin \{\epsilon_1, \ldots, \epsilon_k\}$. Then the k-boundary equation asserts that $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_j \sigma_{\epsilon_j}$, which together with (5.12) implies that $\Lambda_j \neq 0$.

Case 2b: $v_i^- = \epsilon_l$ with $l \in \{1, \ldots, k\}$. If $\Lambda_j = 0$, then the k-boundary equation gives $\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \Lambda_l \sigma_{\epsilon_l}$. If $\Lambda_l = 0$ then this contradicts (5.12), while if $\Lambda_l > 0$ then this contradicts (5.13). $\square$

By Lemmas 5.10 and 5.14 and equations (5.4) and (5.9), we can apply Lemma 5.4 inductively to obtain

$$\#\mathfrak{s}_0^{-1}(0) = (-1)^{N_+ \frac{(N-2)(N-3)}{2}} \cdot n_{N-2}. \tag{5.15}$$

5.6 Rotation rates

To prepare to compute n_{N-2}, we now digress to consider the following question: Let $\Sigma \in \mathcal{N}^{N-2}$, let $\sigma \in V_{i,j,k}$ be a nonzero special cokernel element, and let v be a vertex of $\tau(\Sigma)$. Approximately how does the argument of σ_i/σ_j change as we rotate the corresponding branch point in the S^1 direction?

Definition 5.15. If $\Sigma \in \mathcal{M}$, if v is a vertex of $\tau(\Sigma)$, and if $\sigma \in V_{i,j,k}$ is nonzero, define the *rotation rate* $r(\sigma_i/\sigma_j, v) \in \mathbb{Q}$ as follows. Let e_i and e_j denote the edges of v that lead to the ends i and j respectively. Then

$$r(\sigma_i/\sigma_j, v) := \frac{\eta(\sigma, e_j)}{m(e_j)} - \frac{\eta(\sigma, e_i)}{m(e_i)}. \tag{5.16}$$

Recall that $\eta(\sigma, e)$ denotes the winding number of σ around e, which is computed by Lemma 2.20. Note that if v is not on the path $P_{i,j}$, then $r(\sigma_i/\sigma_j, v) = 0$.

The following lemma is a special case of Proposition 6.9, which is proved in §6.4.

Lemma 5.16. *For all $\varepsilon > 0$, if the constant r_1 in §5.2 is sufficiently large, then the following holds. Let $\Sigma \in \mathcal{N}^{N-2}$, let $0 \neq \sigma \in V_{i,j,k}$, and let v be a vertex of $\tau(\Sigma)$. If one rotates the corresponding branch point in the S^1 direction by angle $\varphi \in \mathbb{R}$, and if the resulting change in the argument of σ_i/σ_j is $r \in \mathbb{R}$, then*

$$|r - \varphi r(\sigma_i/\sigma_j, v)| < \varepsilon.$$

5.7 Beginning the computation of n_{N-2}

The integer n_{N-2} that we want to compute can be decomposed as a sum as follows. Recall that each component of $\mathcal{N}^{N-2}$ determines a data set $(E_1, \ldots, E_{N_+})$, where E_i denotes the set of negative ends that can be reached by downward paths starting at the i^{th} positive end. Since the outgoing edges incident to $v_1, \ldots, v_{N_+-1}$ have positive multiplicities, it follows from Lemma 5.11(a) that

$$E := (E_1, \ldots, E_{N_+}) \in \mathcal{E}(S).$$

(See §4.2 for the definition of $\mathcal{E}(S)$.) Given $E \in \mathcal{E}(S)$, let $\mathcal{N}(E)$ denote the corresponding union of components of $\mathcal{N}^{N-2}$. We can then write

$$n_{N-2} = \sum_{E \in \mathcal{E}(S)} n(E), \tag{5.17}$$

where $n(E)$ denotes the signed count of $\Sigma \in \mathcal{N}(E)$ solving the equations (5.8).

Observe that if $\Sigma \in \mathcal{N}(E)$, then the associated trivalent tree $\tau(\Sigma)$, with the function ρ forgotten, is exactly the tree $\phi(E)$ defined in §4.2.

Given $E \in \mathcal{E}(S)$, we now derive a formula for $n(E)$. To state the formula, let $k \in \{1, \ldots, N-2\}$, and consider a branched cover $\Sigma \in \mathcal{N}(E)$ and a nonzero special cokernel element

$$\sigma^{(k)} \in V_{\epsilon_k, \epsilon_{N-1}, \epsilon_N}.$$

Here, unlike in (5.7), we are not requiring $\sigma^{(k)}_{\epsilon_k} = 1$. Now there is a "dominant" end $d(k) \in \{\epsilon_k, \epsilon_{N-1}, \epsilon_N\}$ whose contribution to $\mathfrak{s}_{\mathbb{C}}(C)(\sigma^{(k)})$ is much larger than the other contributions, in the sense of Definition 5.5. That is:

Lemma 5.17. *Given $\Sigma \in \mathcal{N}^{N-2}$ and $k \in \{1, \ldots, N-2\}$, let v denote the central vertex for $\epsilon_k, \epsilon_{N-1}, \epsilon_N$. Then:*

(a) If $\rho(v) > 0$, then $\epsilon_{N-1} \searrow \epsilon_k, \epsilon_N$.

(b) If $\rho(v) < 0$, then $\epsilon_N \searrow \epsilon_k, \epsilon_{N-1}$.

Proof. By Lemma 5.11(b), there is a downward path P from $\epsilon_{N-1} = N_+$ to $\epsilon_N = N_-$. The central vertex v is somewhere on the path P. Suppose that $\rho(v) > 0$. The path P_{v,ϵ_k} must dip below the level $\rho = -R_N$, because all vertices with $\rho > 0$ are joining. It then follows as in the proof of Lemma 5.6(a) that $\epsilon_{N-1} \searrow \epsilon_k, \epsilon_N$. This proves assertion (a), and assertion (b) follows by a symmetric argument. □

Define $d(k) := \epsilon_{N-1}$ in case (a) above, and $d(k) := \epsilon_N$ in case (b).

Next, define a square matrix $A(E)$ over $\mathbb{Q}$ of size $N-2$ as follows. The rows of $A(E)$ correspond to the ends $\epsilon_1, \ldots, \epsilon_{N-2}$. The columns of $A(E)$ correspond to the vertices $v_1, \ldots, v_{N-2}$. The entries of $A(E)$ are defined by the rotation rates

$$A(E)_{k,l} := r\left(\sigma^{(k)}_{d(k)} / \sigma^{(k)}_{\epsilon_k}, v_l\right).$$

Let $\operatorname{Edge}(E)$ and $\operatorname{Vert}(E)$ denote the sets of edges and internal vertices respectively in the tree $\phi(E)$.

Lemma 5.18. *If r_1 is sufficiently large, then for each $E \in \mathcal{E}(S)$, we have*

$$n(E) = (-1)^{\frac{(N-2)(N-3)}{2} + (N_+ - 1)} \det(A(E)) \prod_{e \in \operatorname{Edge}(E)} m(e). \tag{5.18}$$

Proof. There is a natural action of $\mathbb{R}^{N-2}$ on $\mathcal{N}(E)$ that rotates the $N-2$ branch points in the S^1 direction at speed 2π. The kernel of this action is a nondegenerate lattice $\Lambda(E) \subset \mathbb{Z}^{N-2}$. In fact, the proof of Lemma 2.6 shows that $\Lambda(E)$ is the kernel of the homomorphism

$$\bigoplus_{\operatorname{Vert}(E)} \mathbb{Z} \longrightarrow \bigoplus_{e \in \operatorname{Edge}(E)} \mathbb{Z}/m(e)$$

that sends (the generator corresponding to) a vertex v to the sum of the outgoing edges of v minus the sum of the incoming edges of v. Thus we can identify

$$\mathcal{N}(E) \simeq \bigsqcup_{\pi_0\mathcal{N}(E)} \mathbb{R}^{N-2}/\Lambda(E). \tag{5.19}$$

By Lemma 2.6, we have

$$\det(\Lambda(E)) \cdot |\pi_0\mathcal{N}(E)| = \prod_{e\in\mathrm{Edge}(E)} m(e). \tag{5.20}$$

Now define a map

$$\begin{aligned} f : \mathcal{N}(E) &\longrightarrow (S^1)^{N-2}, \\ \Sigma &\longmapsto \left\{\arg\left(\sigma^{(k)}_{d(k)}/\sigma^{(k)}_{\epsilon_k}\right)\right\}_{k=1}^{N-2}. \end{aligned}$$

By the domination condition (5.13), the map f is homotopic to the map sending

$$\Sigma \longmapsto \left\{\arg\left(\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma^{(k)})/\sigma^{(k)}_{\epsilon_k}\right)\right\}_{k=1}^{N-2}.$$

Therefore the count of solutions to (5.8) on $\mathcal{N}(E)$ is given by

$$n(E) = \deg(f). \tag{5.21}$$

On the other hand, by Lemma 5.16, if r_1 is sufficiently large, then under the identification (5.19) the map f is homotopic to the linear map $A(E)$ on each component of $\mathcal{N}(E)$. Therefore

$$\deg(f) = (-1)^{\frac{(N-2)(N-3)}{2}+(N_+-1)} \cdot \det(\Lambda(E)) \cdot |\pi_0\mathcal{N}(E)| \cdot \det(A(E)). \tag{5.22}$$

Here the sign arises from the orientation convention for $\mathcal{N}^{N-2}$ in §5.2.

Combining equations (5.20), (5.21), and (5.22) proves the lemma. □

5.8 Calculating the determinant

Here is where we stand. By equations (5.15), (5.17), and (5.18), we have

$$\#\mathfrak{s}_0^{-1}(0) = (-1)^{N_-+1} \sum_{E\in\mathcal{E}(S)} \det(A(E)) \prod_{e\in\mathrm{Edge}(E)} m(e). \tag{5.23}$$

By equation (5.23) and Lemma 4.10, the following lemma will finish off the proof of Proposition 5.1. In the statement of this lemma, recall from §4.1 that $P_{\phi(E)}$ denotes the canonical edge pairing on the admissible tree $\phi(E)$, and $W_\theta(\phi(E), T_{\phi(E)})$ denotes the associated positive integer weight.

Lemma 5.19. *For each $E \in \mathcal{E}(S)$, we have*

$$\det(A(E)) \prod_{e \in \mathrm{Edge}(E)} m(e) = (-1)^{N_- - 1} W_\theta(\phi(E), P_{\phi(E)}).$$

Proof. It follows from the definitions that the canonical edge pairing $P_{\phi(E)}$ on $\phi(E)$ is given by $e^\pm_{v_k} = e^\pm_k$, where the edges $e^\pm_k$ are specified in Definitions 5.7 and 5.13.

Now define a matrix B as follows. Let A_l denote the l^{th} row of $A := A(E)$. Then the rows of B are given by the following prescription.

- If $i \in \{1, \ldots, N_+ - 1\}$, then

$$B_i := \begin{cases} A_i - A_k, & v_i^- = \epsilon_k \neq \epsilon_N, \\ A_i, & v_i^- = \epsilon_N. \end{cases}$$

- If $k \in \{N, \ldots, N-2\}$, then (cf. Lemma 5.11(a),(b))

$$B_k := \begin{cases} A_k - A_i, & \epsilon_k \in E'_i, \ i < N_+, \\ A_k, & \epsilon_k \in E'_{N_+}. \end{cases}$$

By equation (4.1), to prove the lemma it suffices to prove (i)–(iii) below:

(i) $\det(A) = \det(B)$.

(ii) B is lower triangular, for a suitable reordering of $\{1, \ldots, N-2\}$.

(iii) The l^{th} diagonal entry $B_{l,l}$ of B is given by

$$\frac{m(e_l^-)\left\lceil m(e_l^+)\theta \right\rceil - m(e_l^+)\left\lfloor m(e_l^-)\theta \right\rfloor}{m(e_l^+) m(e_l^-)} = \begin{cases} B_{l,l}, & l = 1, \ldots, N_+ - 1, \\ -B_{l,l}, & l = N_+, \ldots, N-2. \end{cases} \tag{5.24}$$

Proof of (i): The matrix B is obtained from A by performing the following row operations for $i = 1, \ldots, N_+ - 1$ in order:

- For each k such that $\epsilon_k \in E'_i$, subtract the i^{th} row from the k^{th} row.
- If $v_i^- = \epsilon_k$ with $k \neq N$, then subtract the k^{th} row (which has not yet been modified since $\epsilon_k \in E'_j$ for some $j > i$) from the i^{th} row.

Proof of (ii): We claim that the matrix B is lower triangular if one lists the numbers $1, \ldots, N-2$, which index the rows and columns of B, in the order

$$\{k \mid \epsilon_k \in E'_1\}, 1, \ldots, \{k \mid \epsilon_k \in E'_{N_+-1}\}, N_+ - 1, \{k \mid \epsilon_k \in E'_{N_+}\}. \tag{5.25}$$

Here the set $\{k \mid \epsilon_k \in E'_i\}$ is listed in increasing order of k for each i.

To prove lower triangularity, we first investigate the k^{th} row of B when $\epsilon_k \in E_i'$ and $i < N_+$. By the definition of A,

$$A_{i,l} = r(\sigma^{(i)}_{d(i)}/\sigma^{(i)}_i, v_l),$$
$$A_{k,l} = r(\sigma^{(k)}_{d(k)}/\sigma^{(k)}_{\epsilon_k}, v_l).$$

We now calculate these rotation rates using Definition 5.15 and Lemma 2.20. First note that the dominant ends $d(i)$ and $d(\epsilon_k)$ are equal. The reason is that the central vertex for $i, \epsilon_{N-1}, \epsilon_N$ is the same as the central vertex for $\epsilon_k, \epsilon_{N-1}, \epsilon_N$, because the paths from ϵ_k or i to $P_{\epsilon_{N-1},\epsilon_N}$ both pass through v_k.

More precisely, the path $P_{i,d(i)}$ passes first through the vertices v_j for $\epsilon_j \in E_i'$ in increasing order, then through the vertex v_i. If $v_i^- \neq \epsilon_N$, then at v_i the path $P_{i,d(i)}$ turns (at least temporarily) upward and passes through some additional internal vertices which are all in τ_i^+; otherwise the path $P_{i,d(i)}$ stays downward, and any additional internal vertices on this path are in τ_{i-1}^-. Likewise, the path $P_{\epsilon_k,d(k)}$ possibly first passes through some vertices in τ_{i-1}^-, then hits the vertex v_k, and then agrees with the rest of the path $P_{i,d(i)}$.

Since the central vertex for $i, \epsilon_{N-1}, \epsilon_N$ is the same as the central vertex for $\epsilon_k, \epsilon_{N-1}, \epsilon_N$, it follows by Definition 5.15 and Lemma 2.20 that $A_{i,l} = A_{k,l}$ whenever the paths $P_{i,d(i)}$ and $P_{\epsilon_k,d(\epsilon_k)}$ either both avoid v_l, or both pass through v_l along the same ordered pair of edges. By the above description of these two paths, this can fail only if $v_l \in \tau_{i-1}^-$, or $\epsilon_l \in E_i'$ with $l \leq k$. It follows that the k^{th} row of B has the required form for lower triangularity with respect to the ordering (5.25). Similar arguments show that all other rows of B have the required form.

Proof of (iii): We now prove equation (5.24) in several cases. In these calculations recall that e_l^0 denotes the edge of v_l that is neither e_l^+ nor e_l^-.

Suppose first that $k \in \{N_+, \ldots, N-2\}$. Then equation (5.24) for $l = k$ asserts that

$$B_{k,k} = -\frac{\lceil m(e_k^+)\theta \rceil}{m(e_k^+)} + \frac{\lfloor m(e_k^-)\theta \rfloor}{m(e_k^-)}. \tag{5.26}$$

If $\epsilon_k \in E_i'$ with $i < N_+$, then by the definition of A and Lemma 2.20, and using the descriptions of the paths $P_{i,d(i)}$ and $P_{\epsilon_k,d(k)}$ from the proof of part (ii), we obtain

$$A_{k,k} = \frac{\lfloor m(e_k^-)\theta \rfloor}{m(e_k^-)} - \frac{\lceil m(e_k^0)\theta \rceil}{m(e_k^0)},$$
$$A_{i,k} = \frac{\lceil m(e_k^+)\theta \rceil}{m(e_k^+)} - \frac{\lceil m(e_k^0)\theta \rceil}{m(e_k^0)}.$$

Subtracting these two equations gives (5.26). If $\epsilon_k \in E'_{N_+}$, then (5.26) holds since $B_{k,k} = A_{k,k}$ and $d(k) = \epsilon_{N-1}$.

Finally, suppose $i \in \{1, \ldots, N_+ - 1\}$. Then equation (5.24) for $l = i$ is

$$B_{i,i} = \frac{\lceil m(e_i^+)\theta \rceil}{m(e_i^+)} - \frac{\lfloor m(e_i^-)\theta \rfloor}{m(e_i^-)}. \tag{5.27}$$

By Lemma 2.20 and the definition of A, if $v_i^- = \epsilon_N$ then $A_{i,i}$ is given by the right hand side of (5.27). On the other hand, if $v_i^- = \epsilon_k \neq \epsilon_N$, then

$$A_{i,i} = \frac{\lceil m(e_i^+)\theta \rceil}{m(e_i^+)} - \frac{\lfloor m(e_i^0)\theta \rfloor}{m(e_i^0)}.$$

Also, since $\epsilon_k \in E'_j$ for some $j > i$, we have

$$A_{k,i} = \frac{\lfloor m(e_i^-)\theta \rfloor}{m(e_i^-)} - \frac{\lfloor m(e_i^0)\theta \rfloor}{m(e_i^0)}.$$

The above calculations imply (5.27). □

Remark 5.20. One might try to give a more direct proof of Proposition 5.1 as follows. If $\mathfrak{s}_{\mathbb{C}}(\Sigma) = 0$, then "generically" the tree $\tau(\Sigma)$ is trivalent, and given a nonzero special cokernel element $\sigma \in V_{i,j,k}$, in the equation

$$\mathfrak{s}_{\mathbb{C}}(\Sigma)(\sigma) = \pm\overline{\gamma_i}\sigma_i \pm \overline{\gamma_j}\sigma_j \pm \overline{\gamma_k}\sigma_k = 0,$$

one term is much smaller than the other two. The two larger terms specify two distinguished edges incident to the central vertex v for i, j, and k. One can check that these two edges depend only on v and define an edge pairing on $\tau(\Sigma)$, modulo the choice of which distinguished edge is e_v^+ and which is e_v^-. Moreover, similarly to the above calculations, the count of solutions with this tree and edge pairing is given by plus or minus the weight in Definition 4.2. Thus one finds that $\#\mathfrak{s}_{\mathbb{C}}^{-1}(0)$ is naturally given by a sum over certain trees with edge pairings of their corresponding weights. However the sum that arises is sometimes different than the sum over admissible trees in Lemma 4.5, and the combinatorics of this approach seems difficult.

6 Detailed analysis of the obstruction bundle

In this section, as in §2, fix positive integers $a_1, \ldots, a_{N_+}$ and $a_{-1}, \ldots, a_{-N_-}$ satisfying (2.1), fix an admissible almost complex structure J on $\mathbb{R} \times Y$, and fix an embedded elliptic Reeb orbit α with monodromy angle $\theta \in \mathbb{R} \setminus \mathbb{Q}$ satisfying (2.2). Let $\mathcal{M} := \mathcal{M}(a_1, \ldots, a_{N_+} \mid a_{-1}, \ldots, a_{-N_-})$ denote the moduli space of branched covers of the cylinder $\mathbb{R} \times S^1$ from Definition 2.1, and given $\Sigma \in \mathcal{M}$ recall the operator D_Σ defined in §2.3. As usual, identify an element of $\operatorname{Coker}(D_\Sigma)$ with a smooth, square integrable $(0,1)$ form σ on Σ satisfying $D_\Sigma^*\sigma = 0$, and away from the ramification points use $d\overline{z}$ to identify σ with a complex function.

In this section we give the previously deferred proof of Proposition 2.25, which describes the approximate behavior of nonvanishing cokernel elements away from the ramification points. We also prove a result on the approximate behavior of nonvanishing cokernel elements near isolated clusters of ramification points. The latter result is stated in §6.1, and the proofs of both results are given in §6.2. In §6.3 and §6.4 we use these results to give the previously deferred proofs of Proposition 2.21 and Lemma 5.16.

6.1 Isolated clusters of ramification points

We now state a result which asserts, roughly, that the behavior of a nonvanishing cokernel element near an isolated cluster of ramification points does not depend much on the nature of the distant ramification points.

We need the following preliminary definitions. Let $\pi : \Sigma \to \mathbb{R} \times S^1$ be a branched cover in $\mathcal{M}$. Recall from §2 that Σ determines a tree $\tau(\Sigma)$ with a metric and a map $p : \Sigma \to \tau(\Sigma)$.

Definition 6.1. A nonempty set Z of ramification points in Σ is a *cluster* if there is a connected set $B \subset \tau(\Sigma)$, such that a ramification point $z \in \Sigma$ is in Z if and only if $p(z) \in B$. In this case let Σ_Z denote the branched cover of $\mathbb{R} \times S^1$ obtained by attaching half-infinite cylinders to the boundary circles of $p^{-1}(B)$. The *diameter* of Z is the diameter of the set $p(Z)$ in $\tau(\Sigma)$. For a real number $R > 0$, the cluster Z is *R-isolated* if every vertex in $p(Z)$ has distance at least R from all vertices of $\tau(\Sigma)$ not in $p(Z)$.

Note that there is a canonical identification between a cluster of ramification points Z and the set of ramification points in Σ_Z. Also, if Z is R-isolated, then there is a canonical identification between the set of points in Σ within distance R of a ramification point in Z, and the set of points in Σ_Z within distance R of a ramification point in Σ_Z.

Definition 6.2. A $(0,1)$-form σ on Σ has *exponential growth* if there exists a constant c such that $|\sigma| \le c\exp(c|\pi^* s|)$ at every point in Σ. Define $\widetilde{\operatorname{Coker}}(D_\Sigma)$ to be the space of $(0,1)$-forms with exponential growth on Σ that are annihilated by D_Σ^*.

Proposition 6.3. *Given $r, \varepsilon_0 > 0$, there exists $R > 1/\varepsilon_0$ such that the following holds. Let $\Sigma \in \mathcal{M}$, let Z be an R-isolated cluster of ramification points in Σ of diameter $\le 2r$, and let $\sigma \in \operatorname{Coker}(D_\Sigma)$ be nonvanishing. Then there exists a nonvanishing $\sigma_Z \in \widetilde{\operatorname{Coker}}(D_{\Sigma_Z})$ such that $|\sigma - \sigma_Z| \le \varepsilon_0|\sigma|$ at all points in Σ within distance $1/\varepsilon_0$ of a ramification point in Z.*

To say more about the forms $\sigma_Z \in \widetilde{\mathrm{Coker}}(D_{\Sigma_Z})$ that can arise, first note the following basic lemma:

Lemma 6.4. *Fix a positive integer m and an integer η. Then there exists $\kappa > 0$ with the following property. Let σ be a complex function on $[0,\infty) \times \mathbb{R}/2\pi m\mathbb{Z}$ which is annihilated by $\partial_s - L_m$, is nonvanishing with winding number η, and has exponential growth. Then there is a normalized (L^2-norm 1) eigenfunction γ of L_m with eigenvalue E_γ and winding number η, and constants $\sigma_\gamma \neq 0$ and c_σ, such that*

$$|\sigma(s,t) - \sigma_\gamma \exp(E_\gamma s)\gamma(t)| \leq c_\sigma \exp((E_\gamma - \kappa)s)$$

for all (s,t) with $s \geq 0$.

Proof. This follows by writing $\sigma(s,t) = \sum_\gamma \sigma_\gamma \exp(E_\gamma s)\gamma(t)$, where the sum is over an orthonormal basis of $L^2(\mathbb{R}/2\pi m\mathbb{Z};\mathbb{R}^2)$ consisting of eigenfunctions γ of L_m with eigenvalues E_γ. □

Now suppose $\sigma \in \widetilde{\mathrm{Coker}}(D_\Sigma)$ is nonvanishing. As in §2.3, let η_i^+ and η_j^- denote the winding numbers of σ around the positive and negative ends of Σ. By equation (2.13), we must have

$$\sum_{i=1}^{N_+} \eta_i^+ - \sum_{j=-1}^{-N_-} \eta_j^- = N - 2. \tag{6.1}$$

Given integers η_i^+ and η_j^- satisfying (6.1), let $V(\eta_1^+,\ldots,\eta_{N_+}^+ \mid \eta_{-1}^-,\ldots,\eta_{-N_-}^-)$ denote the vector space of $\sigma \in \widetilde{\mathrm{Coker}}(D_\Sigma)$ such that either $\sigma = 0$, or σ is nonvanishing with winding numbers η_i^+ and η_j^-. Calculations similar to those in Lemmas 2.15 and 2.18, using Lemma 6.4, show that

$$\dim_{\mathbb{R}} V(\eta_1^+,\ldots,\eta_{N_+}^+ \mid \eta_{-1}^-,\ldots,\eta_{-N_-}^-) = 2. \tag{6.2}$$

In particular, there is a vector bundle

$$\mathcal{V}(\eta_1^+,\ldots,\eta_{N_+}^+ \mid \eta_{-1}^-,\ldots,\eta_{-N_-}^-) \longrightarrow \mathcal{M}, \tag{6.3}$$

whose fiber over $\Sigma \in \mathcal{M}$ is the vector space $V(\eta_1^+,\ldots,\eta_{N_+}^+ \mid \eta_{-1}^-,\ldots,\eta_{-N_-}^-)$ associated to Σ.

Remark 6.5. Also in connection with Lemma 6.4, one of the difficulties in proving Proposition 2.25 is that there is no a priori upper bound on the ratio $c_\sigma/|\sigma_\gamma|$. For example, fix $\theta \in (0,1)$ and take $S(t) = -\theta$ and $m = 1$. Let a be a nonnegative real number and take $\sigma = e^{-\theta s}\exp(ae^{-s+it})$ on $[0,\infty) \times S^1$. This σ is nonvanishing and square integrable with $\eta = 0$, so Lemma 6.4 gives

$\gamma(t) = 1\sqrt{2\pi}$. It is then easy to check that the smallest possible value of $c_\sigma/|\sigma_\gamma|$ limits to ∞ as $a \to \infty$.

The following even worse situation can occur over a compact cylinder. Again fix $\theta \in (0,1)$ and take $S(t) = -\theta$ and $m = 1$. Let a be a nonnegative real number and define $\sigma : \mathbb{R} \times S^1 \to \mathbb{C}$ by

$$\sigma(s,t) := e^{-\theta s} \exp\left(\frac{ia}{2}\left(e^{s-it} + e^{-s+it}\right)\right),$$

This is nonvanishing, is annihilated by $\partial_s - L$, and has winding number $\eta = 0$. Even so, there exists a such that $\sigma(0,\cdot)$ has no constant term in its Fourier series, so that $\Pi_W \sigma(0,\cdot) = 0$.

6.2 Proof of the approximation results

We now prove Propositions 2.25 and 6.3 together. If either of these propositions fails, then we can find constants $\varepsilon_0, r > 0$, and a sequence of pairs $\{(\Sigma_k, \sigma_k)\}_{k=1,2,\ldots}$ such that the conclusions of the propositions do not all hold for $(\Sigma, \sigma) = (\Sigma_k, \sigma_k)$ when $R < k$. Hence to prove Propositions 2.25 and 6.3, it is enough to prove the following statement:

- Consider a sequence $\{(\Sigma_k, \sigma_k)\}_{k=1,2,\ldots}$ where $\pi_k : \Sigma_k \to \mathbb{R} \times S^1$ is a branched cover in $\mathcal{M}$ and $\sigma_k \in \operatorname{Coker}(D_{\Sigma_k})$ is nonvanishing. Let r be given. Then we can pass to a subsequence (again indexed by $k = 1, 2, \ldots$) such that for all $\varepsilon_0 > 0$, there exists R such that the conclusions of Proposition 2.25 and 6.3 hold for $(\Sigma, \sigma) = (\Sigma_k, \sigma_k)$ whenever k is sufficiently large.

We now prove the above statement in 7 steps. The strategy is to pass to a subsequence with appropriate convergence properties, and then use estimates on the limit to produce R from ε_0.

Step 1. We begin by passing to a subsequence so that the sequence $\{\Sigma_k\}$ has certain convergence properties.

By Lemma 2.28, we can pass to a subsequence so that the sequence $\{[\Sigma_k]\}$ in $\mathcal{M}/\mathbb{R}$ converges, in the sense of Definition 2.27, to an element $(T; [\Sigma_{*1}], \ldots, [\Sigma_{*p}]) \in \overline{\mathcal{M}/\mathbb{R}}$. Fix Σ_{kj} and $\Phi_{k,e}$ as in Definition 2.27, and carry over the other notation from Definition 2.27. By passing to a further subsequence and increasing r if necessary, we may assume that:

- Σ_{kj} is a component of the π_k-inverse image of a subcylinder in $\mathbb{R} \times S^1$ of length $2k$.

- If $z \in \Sigma_{kj}$ is a ramification point, then $\pi_k(z)$ has distance $\leq r$ from the center of the subcylinder $\pi_k(\Sigma_{kj})$.

Since $\lim_{k\to\infty} T_{-s_{kj}}(\widehat{\Sigma}_{kj}) = \Sigma_{*j}$, we can, possibly after passing to a further subsequence, choose diffeomorphisms of the domains $\Psi_{kj} : \Sigma_{*j} \to \widehat{\Sigma}_{kj}$ such that:

- $T_{s_{kj}} \circ \pi_{*j} \circ \Psi_{kj}^{-1}$ agrees with the projection $\widehat{\Sigma}_{kj} \to \mathbb{R} \times S^1$ at all points in $\widehat{\Sigma}_{kj}$ that have distance ≥ 1 from the ramification points in $\widehat{\Sigma}_{kj}$.
- Let $\Psi_{kj}^{0,1}$ denote the composition of the pullback $\Psi_{kj}^* : T_{\mathbb{C}}^* \widehat{\Sigma}_{kj} \to T_{\mathbb{C}}^* \Sigma_{*j}$ with orthogonal projection $T_{\mathbb{C}}^* \Sigma_{*j} \to T^{0,1}\Sigma_{*j}$. Then the sequence of differential operators $\left\{\Psi_{kj}^{0,1} \circ D_{\widehat{\Sigma}_{kj}}^*\right\}_{k=1,2,\ldots}$ converges to $D_{\Sigma_{*j}}^*$.

The following notation will be used below. Choose a k-independent number $a > 2r+1$. Restrict attention to $k \geq a$. Let $V_{kj} \subset \Sigma_{kj}$ denote π_k^{-1} of the set of points with $s_{kj} - a < s < s_{kj} + a$. Let $U_{kj} \subset V_{kj}$ denote π_k^{-1} of the set of points with $s_{kj} - a + 1 < s < s_{kj} + a - 1$. Let π_{*j} denote the projection $\Sigma_{*j} \to \mathbb{R} \times S^1$, let $V_{*j} \subset \Sigma_{*j}$ denote π_{*j}^{-1} of the set of points with $|s| < a$, and let $U_{*j} \subset V_{*j}$ denote π_{*j}^{-1} of the set of points with $|s| < a - 1$.

Step 2. We now pass to a further subsequence so that the sequence $\{\sigma_k\}$ has certain convergence properties.

To start, normalize the σ_k's to have L^2 norm 1. Define θ_{kj} to be the L^2 norm of σ_k over V_{kj}.

Lemma 6.6. *For each j, there is a smooth $(0,1)$-form σ_{*j} on V_{*j} which is annihilated by $D_{\Sigma_{*j}}^*$ such that:*

(a) The sequence of $(0,1)$-forms $\left\{\Psi_{kj}^{0,1}\left(\theta_{kj}^{-1}\sigma_k|_{V_{kj}}\right)\right\}_{k=a,a+1,\ldots}$ has a subsequence that converges in the C^∞ topology[3] to σ_{*j}.

(b) σ_{*j} is nonvanishing.

Proof. A standard compactness argument using a priori elliptic estimates finds a subsequence of the sequence in (a) converging to a smooth $(0,1)$-form σ_{*j} on V_{*j} that is annihilated by $D_{\Sigma_{*j}}^*$. The $(0,1)$-form σ_{*j} is nonvanishing provided that it is not identically zero, because it is the C^0 limit of a sequence of nonvanishing $(0,1)$-forms, and any zero of σ_{*j} must have negative multiplicity. Thus it remains only to prove that σ_{*j} is not identically zero.

Suppose to the contrary that $\sigma_{*j} = 0$. By elliptic estimates, this assumption implies that for any neighborhood N of the boundary of V_{*j},

$$\lim_{k\to\infty} \int_{V_{*j}\setminus N} \left|\Psi_{kj}^{0,1}(\theta_{kj}^{-1}\sigma_k)\right|^2 = 0. \tag{6.4}$$

[3] Here and below, 'convergence in the C^∞ topology' means convergence in the C^n topology on any compact set for any integer n.

Now pass to a subsequence so that for each edge e of the tree T incident to the j^{th} internal vertex, the L^2 norm of $\Psi_{kj}^{0,1}(\sigma_k/\theta_{kj})$ over the component of $V_{*j} \setminus U_{*j}$ corresponding to e converges as $k \to \infty$ some $c_{j,e} \geq 0$. By (6.4) with $N = V_{*j} \setminus U_{*j}$, we must have $\sum_e c_{j,e} = 1$. Hence there is an edge e of the tree T adjacent to the j^{th} internal vertex with $c_{j,e} > 0$.

We will now show that if j is an internal vertex with $\sigma_{*j} = 0$, and if e is an edge of T incident to j with $c_{j,e} > 0$, then:

(i) e is an internal edge.

(ii) Let $j' \neq j$ denote the other internal vertex of T incident to e; then $\sigma_{*j'} = 0$.

(iii) If e' is an edge of T incident to j' with $c_{j',e'} > 0$, then $e \neq e'$.

By induction using (i) and (ii), we can find an infinite sequence of internal vertices $j_0 = j, j_1 = j', j_2, \ldots$ of T, and an infinite sequence of edges $e_0 = e, e_1 = e', e_2, \ldots$ such that $\sigma_{*j_i} = 0$; the edge e_i is incident to j_i and j_{i+1}; and $c_{j_i,e_i} > 0$. Then property (iii) implies $e_i \neq e_{i+1}$. Since T is a tree, this will give the desired contradiction.

Proof of (i): For each k, let $\mathcal{E}_k$ denote the component cylinder of $\Sigma_k \setminus \bigcup_{j'} U_{kj'}$ corresponding to e. Without loss of generality, $s \leq 1$ on $\mathcal{E}_k$, with $s = 1$ denoting the boundary circle of U_{kj} and $s = 0$ the boundary circle of V_{jk}. By (6.4) again,

$$\begin{aligned} \lim_{k\to\infty} \theta_{kj}^{-2} \int_{1/2}^{1} \|\sigma_k|_{s=\tau}\|^2 d\tau = 0, \\ \lim_{k\to\infty} \theta_{kj}^{-2} \int_{0}^{1/2} \|\sigma_k|_{s=\tau}\|^2 d\tau = c_{j,e}^2. \end{aligned} \tag{6.5}$$

To be more explicit, expand $\sigma_k|_{\mathcal{E}_k}$, regarded as a complex function, in terms of eigenfunctions of L_m as

$$\sigma_k|_{\mathcal{E}_k}(s,t) = \sum_\gamma \sigma_{k\gamma} \exp(E_\gamma s)\gamma(t). \tag{6.6}$$

It follows from (6.5) and (6.6) that for every real number Λ, there exists $c_\Lambda > 0$ such that for all $\varepsilon > 0$, if k is sufficiently large then

$$\sum_{E_\gamma > \Lambda} |\sigma_{k\gamma}|^2 \leq \varepsilon\theta_{kj}^2, \qquad \sum_{E_\gamma \leq \Lambda} |\sigma_{k\gamma}|^2 \geq c_\Lambda(1-\varepsilon)c_{j,e}^2\theta_{kj}^2. \tag{6.7}$$

Taking $\Lambda < 0$ in the right most inequality shows that $\mathcal{E}_k$ is compact. This is because if $\mathcal{E}_k$ is not compact, then square integrability of σ_k requires that $\sigma_{k\gamma} = 0$ when $E_\gamma < 0$.

Proof of (ii): Let $s_k < 0$ denote the value of s on the boundary circle of $\mathcal{E}_k$ in $V_{kj'}$. Let m denote the multiplicity of the edge e, and let $\Pi_{\le\Lambda}$ denote the orthogonal projection in $L^2(\mathbb{R}/2\pi m\mathbb{Z};\mathbb{R}^2)$ onto the span of the eigenfunctions of L_m with eigenvalue $\le \Lambda$. Likewise let $\Pi_{>\Lambda}$ denote the projection onto the sum of the eigenspaces with eigenvalues $> \Lambda$. It follows from (6.7) that if k is sufficiently large then

$$\frac{\|\Pi_{>\Lambda}\sigma_k|_{s=s_k}\|_2}{\|\Pi_{\le\Lambda}\sigma_k|_{s=s_k}\|_2} = \frac{\sum_{E_\gamma>\Lambda} |\sigma_{k\gamma}|^2 \exp(2E_\gamma s_k)}{\sum_{E_\gamma\le\Lambda} |\sigma_{k\gamma}|^2 \exp(2E_\gamma s_k)} \le \exp(\kappa_\Lambda s_k) \tag{6.8}$$

where κ_Λ is a positive constant. By the convergence in (a) to $\sigma_{*j'}$, it follows from (6.8) that $\sigma_{*j'} = 0$ on the boundary circle of $V_{*j'} \setminus U_{*j'}$ corresponding to the edge e, and hence on all of $V_{*j'}$, because Λ can be taken arbitrarily negative and $s_k \to -\infty$ as $k \to \infty$.

Proof of (iii): The L^2 norm of $\theta_{kj'}^{-1}\sigma_k$ over $\mathcal{E}_k \cap V_{kj'}$ must converge to zero, because otherwise the analogue of (6.7) for j', in which the inequalities on the eigenvalues are reversed, would contradict (6.8). $\square$

Now pass to a subsequence such that the convergence in Lemma 6.6 holds for each j. This convergence (or an argument independent of Lemma 6.6 using winding bounds) allows us to pass to a further subsequence such that for each edge e of the tree T, the winding number of σ_k around the component of $\Sigma_k \setminus \bigcup_k U_{kj}$ corresponding to e does not depend on k.

Step 3. We now show that σ_{*j} has an extension over Σ_{*j} with various nice properties. More properties of σ_{*j} will be established later in Lemma 6.8.

Lemma 6.7. *For each j, the $(0,1)$-form σ_{*j} extends to a smooth $(0,1)$-form σ_{*j} on Σ_{*j} which is annihilated by $D^*_{\Sigma_{*j}}$, and is such that:*

(a) Let e be an external edge of T incident to the j^{th} internal vertex, and let $\mathcal{E}_k$ denote the corresponding noncompact component of $\Sigma_k \setminus \bigcup_{j'} U_{kj'}$. On the corresponding component of $\Sigma_{*j} \setminus V_{*j}$, the $(0,1)$-form σ_{*j} is square integrable, and the limit in the C^∞ topology of the sequence $\{\Psi^{0,1}_{kj}(\theta_{kj}^{-1}\sigma_k|_{\mathcal{E}_k\cap\Sigma_{kj}})\}_{k=a,a+1,\ldots}$. In particular, σ_{*j} is nonvanishing here.

(b) σ_{*j} has exponential growth on all of Σ_{*j}.

Proof. (a) We extend σ_{*j} over the end in question as follows. Without loss of generality, $s \le -a+1$ on $\mathcal{E}_k$. Expand σ_k on $\mathcal{E}_k$ by the formula (6.6), and on the corresponding end of Σ_{*j} where $-a < s \le -a+1$ write

$$\sigma_{*j} = \sum_\gamma \sigma_{*j\gamma} \exp(E_\gamma s)\gamma(t). \tag{6.9}$$

By the convergence in Lemma 6.6(a) at $s=-a+1$, we have

$$\lim_{k\to\infty}\sum_\gamma \exp(2E_\gamma(-a+1))|\sigma_{*j\gamma}-\theta_{kj}^{-1}\sigma_{k\gamma}|^2=0. \tag{6.10}$$

If $E_\gamma<0$, then square integrability of σ_k implies that $\sigma_{k\gamma}=0$, and hence $\sigma_{*j\gamma}=0$ also by (6.10). Consequently (6.9) defines an extension of σ_{*j} over the component of $\Sigma_{*j}\setminus V_{*j}$ corresponding to $\mathcal{E}_k$ that has all of the required properties.

(b) We now extend each σ_{*j} over the rest of Σ_{*j}. Let e be an internal edge of T and let $\mathcal{E}_k$ denote the corresponding compact component of $\Sigma_k\setminus\bigcup_j U_{kj}$. Let j and j' denote the upper and lower vertices of e, and suppose without loss of generality that $s_k\le s\le -a+1$ on $\mathcal{E}_k$. Expand σ_k on $\mathcal{E}_k$ where $-a\le s\le -a+1$ as in (6.6), and expand σ_{*j} where $-a\le s\le -a+1$ as in (6.9). Meanwhile, expand $\sigma_{*j'}$ where $a-1\le s<a$ on the component of $\Sigma_{*j'}\setminus U_{*j'}$ that corresponds to $\mathcal{E}_k$ as in (6.9) but with j' replacing j. By the convergence in Lemma 6.6(a) at $s=-a+1$ and at $s=s_k$, we have

$$\lim_{k\to\infty}\sum_\gamma \exp(2E_\gamma(-a+1))|\sigma_{*j\gamma}-\theta_{kj}^{-1}\sigma_{k\gamma}|^2=0, \tag{6.11}$$

$$\lim_{k\to\infty}\sum_\gamma |\exp(E_\gamma(a-1))\sigma_{*j'}-\theta_{kj'}^{-1}\exp(E_\gamma s_k)\sigma_{k\gamma}|^2=0. \tag{6.12}$$

It follows that if $\sigma_{*j\gamma}\neq 0$ and $\sigma_{*j'\gamma'}\neq 0$ then $E_\gamma\ge E_{\gamma'}$, because otherwise

$$\frac{\sigma_{*j\gamma'}\sigma_{*j'\gamma}}{\sigma_{*j\gamma}\sigma_{*j'\gamma'}}=\lim_{k\to\infty}\exp((E_{\gamma'}-E_\gamma)(-s_k+a-1))$$

is infinite since $\lim_{k\to\infty}s_k=-\infty$.

We know from Lemma 6.6(b) that σ_{*j} and $\sigma_{*j'}$ are nonzero, so there exist γ,γ' with $\sigma_{*j\gamma}$ and $\sigma_{*j'\gamma'}$ nonzero. Hence there is a smallest eigenvalue E_+ such that $E_+=E_\gamma$ with $\sigma_{*j\gamma}\neq 0$, and a largest eigenvalue E_- such that $E_-=E_{\gamma'}$ with $\sigma_{*j'\gamma'}\neq 0$. Hence (6.9) defines an extension of σ_{*j} over the negative end of Σ_{*j} corresponding to e, and this extension is a smooth $(0,1)$-form with exponential growth annihilated by $D_{\Sigma_{*j}}$. Likewise, $\sigma_{*j'}$ extends over the positive end of $\Sigma_{*j'}$ corresponding to e as a smooth $(0,1)$-form with exponential growth annihilated by $D_{\Sigma_{*j'}}$. □

Step 4. We now show that for any $\varepsilon_0>0$, there exists R such that for all k, the conclusions of Proposition 2.25 hold for $(\Sigma,\sigma)=(\Sigma_k,\sigma_k)$ whenever $p(z)$ is on an external edge e of the tree $\tau(\Sigma)$.

The external edge e of $\tau(\Sigma)$ corresponds to an external edge of T which we also denote by e. Let m denote the multiplicity of e and let j denote the internal vertex of T incident to e. By symmetry, we may assume that the leaf incident to e is negative. Let $\mathcal{E}_k$ denote the corresponding noncompact component of $\Sigma_k\setminus\bigcup_{j'} U_{kj'}$. As usual, there is no loss of generality in assuming that $s \le -a+1$ on $\mathcal{E}_k$. Let $s_k^+ + 1 > -a+1$ denote the s value of the closest ramification point in Σ_k to the $s=-a+1$ circle in $\mathcal{E}_k$. The cylinder $\mathcal{E}_k$ then extends as $\overline{\mathcal{E}}_k = (-\infty, s_k^+ + 1) \times \mathbb{R}/2\pi m\mathbb{Z}$. The convergence of the sequence of branched covers $T_{-s_{kj}}(\Sigma_{kj})$ from Step 1 implies that the sequence $\{s_k^+\}$ converges to a number s_+ with $|s_+ + 1| \le r$.

Expand σ_k and σ_{*j} on $\overline{\mathcal{E}_k}$ as in (6.6) and (6.9). By the convergence in Lemma 6.6(a) at $s = s_+$, we have

$$\lim_{k\to\infty} \sum_{\gamma} \exp(2E_\gamma s_k^+) \left|\sigma_{*j\gamma} - \theta_{kj}^{-1}\sigma_{k\gamma}\right|^2 = 0. \tag{6.13}$$

Suppose that σ_k has winding number η on $\overline{\mathcal{E}_k}$ for all k. Let γ_+ and γ_- be orthonormal eigenfunctions of L_m with winding number η and eigenvalues $E_+ \ge E_-$. By Lemma 6.4, the following hold for each k:

- At least one of the coefficients $\sigma_{k\gamma_-}$, $\sigma_{k\gamma_+}$ is nonzero.
- If γ is an eigenfunction of L_m with $E_\gamma < E_-$, then $\sigma_{k\gamma} = 0$.

By Lemmas 6.6(a) and 6.7(a), the function σ_{*j} also has winding number η on $\overline{\mathcal{E}_k}$, so the above two properties also hold for the coefficients $\sigma_{*j\gamma}$.

It now follows from (6.13) that there is a k-independent number c with

$$\sum_{E_\gamma > E_+} \exp(2E_\gamma s_k^+)|\sigma_{k\gamma}|^2 < c\left(\exp(2E_- s_k^+)|\sigma_{k\gamma_-}|^2 + \exp(2E_+ s_k^+)|\sigma_{k\gamma_+}|^2\right). \tag{6.14}$$

Let $\kappa > 0$ denote the difference between E_+ and the next largest eigenvalue. Then it follows from (6.14) and elliptic regularity for the operator D_Σ^* that there is a k-independent constant c with

$$\begin{aligned}
&\left|\sigma_k(s,t) - \sigma_{k\gamma_-}\exp(E_- s)\gamma_-(t) - \sigma_{k\gamma_+}\exp(E_+ s)\gamma_+(t)\right| \\
&\qquad < c\cdot\exp(\kappa(s - s_k^+))\left|\sigma_{k\gamma_-}\exp(E_- s)\gamma_-(t) + \sigma_{k\gamma_+}\exp(E_+ s)\gamma_+(t)\right|
\end{aligned}$$

for all $(s,t) \in \overline{\mathcal{E}_k}$ with $s \le s_k^+ - 1$. Given $\varepsilon_0 > 0$, choose $R \ge 2$ sufficiently large that

$$c\exp(-\kappa(R-1)) < \varepsilon_0. \tag{6.15}$$

Then the conclusions of Proposition 2.25 follow when $(\Sigma,\sigma) = (\Sigma_k,\sigma_k)$ and the point $p(z)$ lies in the external edge e of $\tau(\Sigma)$.

Step 5. We now show that for any $\varepsilon_0 > 0$, there exists R such that for all k, the conclusions of Proposition 2.25 hold for $(\Sigma,\sigma) = (\Sigma_k,\sigma_k)$ whenever $p(z)$ is on an internal edge of the tree $\tau(\Sigma)$.

To start, we can assume that $R > 2r+1$. This ensures that we only have to consider z in a compact component $\mathcal{E}_k$ of $\Sigma_k \setminus \cup_{j''} U_{kj''}$ corresponding to an internal edge e of T. Let j' and j denote the lower and upper vertices respectively of e. Let $s_k^+ + 1 > -a+1$ denote the s value of the nearest ramification point in Σ_{kj} to the $s=-a+1$ circle in $\mathcal{E}_k$, and let $s_k^- - 1 < s_k$ denote the s value of the nearest critical point in $\Sigma_{kj'}$ to the $s = s_k$ circle in $\mathcal{E}_k$. Thus $\mathcal{E}_k$ extends to a cylinder $\overline{\mathcal{E}_k} \simeq (s_k^- - 1, s_k^+ + 1) \times \mathbb{R}/2\pi m(e)\mathbb{Z}$. As in Step 4, the sequence $\{s_k^+\}$ converges to a number s_+, while $\lim_{k\to\infty} s_k^- = -\infty$. Define E_+ and E_- as in the proof of Lemma 6.7(b), and let $\gamma_\pm$ be a normalized eigenfunction with eigenvalue $E_\pm$. We assume in what follows that if $E_+ = E_-$, then the corresponding eigenspace is one dimensional; the argument in the case when the dimension is two has no substantive differences.

Similarly to (6.14), there are k-independent numbers c_+ and c_- such that

$$\begin{aligned} \sum_{E_\gamma > E_+} \exp(2E_\gamma s_k^+)|\sigma_{k\gamma}|^2 &< c_+ \exp(2E_+ s_k^+)|\sigma_{k\gamma_+}|^2, \\ \sum_{E_\gamma < E_-} \exp(2E_\gamma s_k^-)|\sigma_{k\gamma}|^2 &< c_- \exp(2E_- s_k^-)|\sigma_{k\gamma_-}|^2. \end{aligned} \tag{6.16}$$

It follows that there are k-independent numbers $c, \kappa > 0$ such that

$$\begin{aligned} \left| \sum_{E_\gamma > E_+} \sigma_{k\gamma} \exp(E_\gamma s)\gamma(t) \right| &< c|\sigma_{k\gamma_+}| \exp(E_+ s) \exp(\kappa(s - s_k^+)), \\ \left| \sum_{E_\gamma < E_-} \sigma_{k\gamma} \exp(E_\gamma s)\gamma(t) \right| &< c|\sigma_{k\gamma_-}| \exp(E_- s) \exp(\kappa(s_k^- - s)) \end{aligned} \tag{6.17}$$

whenever $s_k^- \le s \le s_k^+$.

Suppose that σ_k has winding number η on $\mathcal{E}_k$ for all k. We now show that the eigenfunction γ_+ has winding number η. If γ is a normalized eigenfunction with $E_\gamma < E_+$, then since $\sigma_{*j\gamma_+} \neq 0$ and $\sigma_{*j\gamma} = 0$, it follows from (6.13) that

$$\lim_{k\to\infty} \frac{\sigma_{k\gamma}}{\sigma_{k\gamma_+}} = 0.$$

Combining this limit for $E_- \le E_\gamma < E_+$ with the inequalities (6.17), we deduce that for any $\varepsilon > 0$, if k is sufficiently large then

$$\frac{|\sigma_k(s,t) - \sigma_{k\gamma_+} \exp(E_+ s)\gamma_+(t)|}{|\sigma_{k\gamma_+}| \exp(E_+ s)} < c\exp(\kappa(s - s_k^+)) + \varepsilon \exp((E_- - E_+)s) \tag{6.18}$$

whenever $s_k^- \le s \le s_k^+$. By taking s sufficiently small and then taking ε sufficiently small (both of which we can do by taking k sufficiently large), we can make the right hand side of (6.18) less than $\min_t |\gamma_+(t)|$. Hence there exist k and s such that that γ_+ has the same winding number as $\sigma_k(s,\cdot)$, and of course the latter winding number is η.

Likewise, γ_- has winding number η. In particular, there are no eigenvalues between E_- and E_+. There are now two cases to consider regarding E_- and E_+.

Suppose first that $E_+ > E_-$. Then the inequalities (6.16) imply that

$$\begin{aligned} \left|\sigma_k(s,t) - \sigma_{k\gamma_+}\exp(E_+s)\gamma_+(t) - \sigma_{k\gamma_-}\exp(E_-s)\gamma_-(t)\right| < \qquad (6.19) \\ c\left|\exp(\kappa(s-s_k^+))\sigma_{k\gamma_+}\exp(E_+s)\gamma_+(t) + \exp(\kappa(s_k^- - s))\sigma_{k\gamma_-}\exp(E_-s)\gamma_-(t)\right| \end{aligned}$$

whenever $s_k^- + 1 \le s \le s_k^+ - 1$. Given $\varepsilon_0 > 0$, choose $R \ge 2$ sufficiently large that (6.15) holds. Then (6.19) implies the conclusions of Proposition 2.25 when $(\Sigma, \sigma) = (\Sigma_k, \sigma_k)$ and $p(z)$ is in the edge of $\tau(\Sigma)$ corresponding to e.

Suppose next that $E_+ = E_-$. Recall that we are assuming that the corresponding eigenspace is one dimensional, so that $\gamma_+ = \gamma_-$. Then (6.19) holds with the γ_- term on the left hand side deleted. So given $\varepsilon_0 > 0$, it is enough choose $R \ge 2$ sufficiently large that $c\exp(-\kappa(R-1)) < \varepsilon_0/2$.

This completes the proof of Proposition 2.25. □

Step 6. We now prove an addendum to Lemma 6.7.

Lemma 6.8. *For each j, and for each internal edge e of T incident to the j^{th} internal vertex, the following two points hold:*

(i) σ_{*j} is nonvanishing on the component of $\Sigma_{*j} \setminus V_{*j}$ corresponding to e.

(ii) The sequence $\{\Psi_{kj}^{0,1}(\sigma_k) - \theta_{kj}\sigma_{*j}\}_{k=a,a+1,\ldots}$ converges to zero in the C^∞ topology on the end in Σ_{*j} that corresponds to e.

Proof. Let m denote the multiplicity of e and write $\widetilde{S^1} := \mathbb{R}/2\pi m\mathbb{Z}$. Without loss of generality, the component $\mathcal{E}_k$ of $\Sigma_k \setminus \bigcup_{j'} U_{kj'}$ corresponding to e is identified with $[s_k, -a+1] \times \widetilde{S^1}$. On the corresponding end of Σ_{*j} where $-a \le s \le -a+1$, expand σ_{*j} as in (6.9). Recall from Step 3 that there is a smallest eigenvalue E_+ of L_m such that $E_+ = E_{\gamma_+}$ with $\sigma_{*j\gamma_+} \neq 0$; and in particular the expansion (6.9) is valid for all $s \le -a+1$. It follows that if $-s$ is large, then the winding number of $\sigma_{*j}(s,\cdot)$ around $\widetilde{S^1}$ equals the winding number of γ_+. By Step 5 and Lemma 6.6(a), the latter is the winding number of $\sigma_{*j}(s,\cdot)$ when $s > -a$. Since all zeroes of σ_{*j} have negative degree, we conclude that σ_{*j} is nonvanishing on $(-\infty, -a] \times \widetilde{S^1}$. This proves (i).

To prove (ii), it is enough to show that given $s \le -a$,

$$\lim_{k\to\infty} \sum_\gamma \exp(2E_\gamma s)|\theta_{kj}\sigma_{*j\gamma} - \sigma_{k\gamma}|^2 = 0. \qquad (6.20)$$

Here we have expanded σ_k on $\mathcal{E}_k$ as in (6.6). Since $\theta_{kj} \le 1$, it follows from the convergence in Lemma 6.6(a) that (ii) holds when $s = -a+1$. It is then enough to show that given $s \leq -a$,

$$\lim_{k\to\infty} \sum_{E_\gamma < E_+} \exp(2E_\gamma s)|\sigma_{k\gamma}|^2 = 0. \tag{6.21}$$

By (6.16), there is a k-independent constant c such that

$$\sum_{E_\gamma < E_+} \exp(2E_\gamma s)|\sigma_{k\gamma}|^2 < c \exp(2E_- s)|\sigma_{k\gamma_-}|^2 \tag{6.22}$$

whenever $s_k \leq s \leq -a+1$. For any given $s \leq -a$, if k is sufficiently large then $s_k \leq s$ so that (6.22) is applicable. The inequality (6.22) then implies (6.21) because $\lim_{k\to\infty} \sigma_{k\gamma_-} = 0$. □

Step 7. Let $\varepsilon_0 > 0$ be given; we now show that there exists R such that the conclusions of Proposition 6.3 hold for $(\Sigma, \sigma) = (\Sigma_k, \sigma_k)$ whenever k is sufficiently large.

In fact, we can take $R := r + 1$, where r was fixed in Step 1. To see why, let Z be a cluster of ramification points in Σ_k satisfying the assumptions of Proposition 6.3. Then Z contains all the ramification points in Σ_{kj} for some j, while our assumption that $k \geq a > 2r+1$ implies that Z contains no other ramification points. Thus $\Sigma_{kZ} = \widehat{\Sigma}_{kj}$. By Lemmas 6.6, 6.7, and 6.8, there is a nonvanishing $(0,1)$-form $\sigma_{*j} \in \widetilde{\operatorname{Coker}}(D_{\Sigma_{*j}})$ such that if k is sufficiently large, then

$$\left|\sigma_k - (\Psi_{kj}^{0,1})^{-1}(\theta_{kj}\sigma_{*j})\right| < \frac{\varepsilon_0}{2}|\sigma_k|$$

at all points in Σ_k within distance $1/\varepsilon_0$ of a ramification point in Z. By the conditions on Ψ_{kj}, and using the vector bundle structure on (6.3), if k is sufficiently large then we can also find a nonvanishing $(0,1)$-form $\sigma_{kZ} \in \widetilde{\operatorname{Coker}}(D_{\Sigma_{kZ}})$ such that

$$\left|(\Psi_{kj}^{0,1})^{-1}(\theta_{kj}\sigma_{*j}) - \sigma_{kZ}\right| < \frac{\varepsilon_0}{2}|\sigma_k|$$

at all points in Σ_k within distance $1/\varepsilon_0$ of a ramification point in Z. Combining the above two inequalities shows that the conclusions of Proposition 6.3 hold for $(\Sigma, \sigma) = (\Sigma_k, \sigma_k)$ whenever k is sufficiently large.

This completes the proof of Proposition 6.3. □

6.3 Proof of the relative size estimate

We now prove Proposition 2.21. The proof has four steps.

Step 1. We begin by using Proposition 2.25 to derive an estimate for the change in $|\sigma|$ along a cylinder away from the ramification points. Let

$\Sigma \in \mathcal{M}$, let $\sigma \in \operatorname{Coker}(D_\Sigma)$ be nonvanishing, and let e be an edge of the tree $\tau(\Sigma)$ of multiplicity m. Let $\mathcal{E}$ denote the cylinder in Σ corresponding to e, and identify $\mathcal{E}$ with an interval cross $\mathbb{R}/2\pi m\mathbb{Z}$ as usual. Then on $\mathcal{E}$ we can write

$$\Pi_W \sigma(s,t) = \exp(E_-(\sigma,e)s)\gamma_-(t) + \exp(E_+(\sigma,e)s)\gamma_+(t), \tag{6.23}$$

where $\gamma_\pm$ are orthogonal eigenfunctions of L_m with eigenvalues $E_\pm(\sigma,e)$, and at least one of $\gamma_\pm$ is nonzero. It follows from the above equation that

$$\log\|\Pi_W\sigma(s,\cdot)\| = \max\left\{E_-(\sigma,e)s + \log\|\gamma_-\|, E_+(\sigma,e)s + \log\|\gamma_+\|\right\} + \text{Error}$$

where $\|\cdot\|$ denotes the L^2 norm on $\mathbb{R}/2\pi m\mathbb{Z}$, and $0 \le \text{Error} \le \frac{\log 2}{2}$. Consequently, if $s' < s$ then

$$\log\frac{\|\Pi_W\sigma(s,\cdot)\|}{\|\Pi_W\sigma(s',\cdot)\|} \in \Big[E_-(\sigma,e)(s-s') - \log 2, E_+(\sigma,e)(s-s') + \log 2\Big]. \tag{6.24}$$

Next, observe that there are only finitely many possible values of the winding number $\eta(\sigma,e)$ when $\Sigma \in \mathcal{M}$, $\sigma \in \operatorname{Coker}(D_\Sigma)$ is nonvanishing, and e is an edge of $\tau(\Sigma)$. This follows from the winding bounds (2.12) together with equations (2.13) and (2.18). Hence there is a constant $\varepsilon_0 > 0$ such that for any $\Sigma \in \mathcal{M}$ and nonvanishing $\sigma \in \operatorname{Coker}(D_\Sigma)$, on the cylinder $\mathcal{E}$ in Σ corresponding to an edge e of $\tau(\Sigma)$, the inequality (2.23) implies that

$$\big|\log|\sigma(s,t)| - \log\|\Pi_W\sigma(s,\cdot)\|\big| < 1/\varepsilon_0. \tag{6.25}$$

Let R denote the constant provided by Proposition 2.25 for this ε_0. Then by the inequalities (6.24) and (6.25), we conclude that there is a constant $c > 0$ such that on a cylinder $\mathcal{E}$ corresponding to an edge e, if $s' < s$ have distance at least R from the endpoints of the corresponding interval, then

$$\log|\sigma(s,\cdot)| - \log|\sigma(s',\cdot)| \in \big[E_-(\sigma,e)(s-s') - c, E_+(\sigma,e)(s-s') + c.\big] \tag{6.26}$$

Step 2. We now inductively define certain constants d_k, r_k, r_k' for $k = 1,\ldots,N$, whose significance will become clear in subsequent steps.

To start, define $d_1 := 0$.

Next, supposing that d_k has been defined, we want to choose $r_k > R$ and r_k' with the following property: Let $\Sigma \in \mathcal{M}$ and let $\sigma \in \operatorname{Coker}(D_\Sigma)$ be nonvanishing. Let $B \subset \tau(\Sigma)$ be a compact connected set such that:

(i) The vertices in B correspond to an r_k-isolated cluster of ramification points with diameter $\le d_k$ and total ramification index k.

(ii) Each boundary point in B has distance exactly R from the nearest vertex in B.

Let $z_1, z_2 \in \Sigma$ with $p(z_1), p(z_2) \in B$. Then

$$\big|\log|\sigma(z_1)| - \log|\sigma(z_2)|\big| \le r'_k. \tag{6.27}$$

The existence of such r_k and r'_k follows by applying Proposition 6.3 with $r = d_k/2$ and $\varepsilon_0 < 1/\max(R, d_k)$, and using compactness of the projectivization of the vector bundle $\mathcal{V}(\cdots \mid \cdots)$ in (6.3) over $\mathcal{M}_r(\cdots \mid \cdots)$. This last compactness follows from Lemmas 2.8 and 2.19.

Finally, let $k > 1$, and suppose that d_i, r_i, r'_i have been defined for all $i < k$. Then

$$d_k := \max_{i+j=k,\ 0<i,j<k} (d_i + d_j + \max\{r_i, r_j\}). \tag{6.28}$$

Step 3. We claim now that for any $\Sigma \in \mathcal{M}$, the internal vertices in the tree $\tau(\Sigma)$ can be partitioned into disjoint subsets $V_1, \ldots, V_l$ such that for each $i = 1, \ldots, l$, the following two properties hold. Let k_i denote the total ramification index of the ramification points in Σ corresponding to vertices in V_i.

(a) The set V_i has diameter at most d_{k_i} in $\tau(\Sigma)$, and is contained in a connected set B which does not intersect any V_j with $i \neq j$.

(b) Let e be an edge incident to vertices in V_i and V_j with $i \neq j$. Then the length of e is greater than $\max\{r_{k_i}, r_{k_j}\}$.

We construct a partition satisfying (a) and (b) by induction as follows. Start with the partition into sets of cardinality one. Then (a) automatically holds. For the induction step, suppose we have a partition satisfying (a) but not (b). Then there exists an edge e incident to vertices in V_i and V_j with $i \neq j$ whose length is at most $\max\{r_{k_i}, r_{k_j}\}$. We now modify the partition by merging the subsets V_i and V_j into a single subset. Condition (a) still holds because

$$\operatorname{diam}(V_i \cup V_j) = \operatorname{diam}(V_i \cup V_j \cup e) \le d_{k_i} + d_{k_j} + \max\{r_{k_i}, r_{k_j}\} \le d_{k_i + k_j}$$

by (6.28). Since there are only finitely many vertices, repeating this step must eventually yield a partition satisfying both (a) and (b).

Step 4. We now complete the proof of Proposition 2.21. By conditions (a) and (b) in Step 3, we can find compact connected subsets $B_1, \ldots, B_l$ of $\tau(\Sigma)$, together containing all of the internal vertices, such that each B_i satisfies conditions (i) and (ii) in Step 2 with $B = B_i$ and $k = k_i$. Then to prove the estimate (2.20), divide the path $P_{x,y}$ into segments, each of which is either contained in one of the B_i's or outside the interiors of all of the B_i's. Use (6.27) to estimate the change in $\log|\sigma|$ along segments of the former type, and use (6.26) to estimate the change in $\log|\sigma|$ along segments of the latter type. $\square$

6.4 How moving a ramification point affects the cokernel

This subsection proves Proposition 6.9 below, which describes how the cokernel of D_Σ changes as one modifies Σ by moving a ramification point. Lemma 5.16 is a special case of Proposition 6.9.

To state Proposition 6.9, assume that $S(t) = \theta$, so that the operator D_Σ is $\mathbb{C}$-linear for each $\Sigma \in \mathcal{M}$. Fix integers η_i^+ for $i = 1, \ldots, N_+$ and η_j^- for $j = -1, \ldots, -N_-$ satisfying (6.1). Assume also that

$$\eta_i^+ \geq \lceil a_i\theta \rceil, \qquad \eta_j^- \leq \lfloor a_j\theta \rfloor. \tag{6.29}$$

Then the vector space $V := V(\eta_1^+, \ldots, \eta_{N_+}^+ \mid \eta_{-1}^-, \ldots, \eta_{-N_-}^-)$ from (6.2) is a complex linear subspace of $\operatorname{Coker}(D_\Sigma)$ with complex dimension 1.

Given $\Sigma \in \mathcal{M}$ and an edge e of the tree $\tau(\Sigma)$, let $\mathcal{E}$ denote the cylinder in Σ corresponding to e. Given $0 \neq \sigma \in V$, let $\eta(e) := \eta(\sigma, e)$ denote the winding number of σ around $\mathcal{E}$; by equation (2.18), this depends only on the numbers η_i^+ and η_j^-. Also recall the notation $W(e)$ from Definition 2.24; here we have $\dim_{\mathbb{C}} W(e) = 1$. If we choose an identification of $\mathcal{E}$ with an interval cross $\mathbb{R}/2\pi m(e)\mathbb{Z}$ commuting with the projections to $\mathbb{R} \times S^1$, then as in (6.23), on $\mathcal{E}$ we can write

$$\Pi_W \sigma(s,t) = \exp\left(\left(\theta - \frac{\eta(e)}{m(e)}\right) s\right) \sigma_e(t) \tag{6.30}$$

where $\sigma_e \in W(e)$ is given by

$$\sigma_e(t) = a_e \exp\left(\frac{\eta(e)}{m(e)} it\right) \tag{6.31}$$

for some $a_e \in \mathbb{C}^\times$.

Since $\dim_{\mathbb{C}} V = 1$, the eigenfunction σ_e determines σ, which in turn determines $\sigma_{e'}$ for any other edge e' of $\tau(\Sigma)$. Thus for every pair of edges e, e', the map sending σ_e to $\sigma_{e'}$ is an isomorphism

$$\Phi_{e,e'}(\Sigma) \in \operatorname{Hom}(W(e), W(e')) = W(e)^* \otimes W(e') \tag{6.32}$$

which depends only on the branched cover $\Sigma \in \mathcal{M}$ and on our fixed integers η_i^+ and η_j^-. We now want to study how $\Phi_{e,e'}$ changes as we rotate the ramification points in the t direction.

If v is an internal vertex, define a rational number $r(e, e'; v)$ as follows: Let f and f' denote the edges incident to v that lead from v to e and e' respectively, and define

$$r(e, e'; v) := \frac{\eta(f)}{m(f)} - \frac{\eta(f')}{m(f')}.$$

Proposition 6.9. *For all $\varepsilon > 0$ there exists $r > 0$ such that the following holds. Fix integers η_i^+ and η_j^- satisfying* (6.1) *and* (6.29). *Fix $\Sigma \in \mathcal{M}$ such that $\tau(\Sigma)$ is trivalent and each edge of $\tau(\Sigma)$ has length $\geq r$. Let v be an internal vertex of $\tau(\Sigma)$, and let Σ' be obtained from Σ by rotating the ramification point corresponding to v by angle $\varphi \in \mathbb{R}$ in the t direction. Then*

$$\Phi_{e,e'}(\Sigma') = (1 + O(\varepsilon)) \exp(i\varphi r(e, e'; v))\Phi_{e,e'}(\Sigma). \tag{6.33}$$

Here and below, '$O(\varepsilon)$' denotes a complex number z with $|z| < \varepsilon$.

Proof. It follows from the definitions that $\Phi_{e,e''}(\Sigma) = \Phi_{e',e''}(\Sigma)\circ\Phi_{e,e'}(\Sigma)$ and $r(e, e''; v) = r(e, e'; v) + r(e', e''; v)$. Hence by induction, it suffices to prove the lemma when e and e' are both incident to the same vertex w. We do so in three steps.

Step 1. As in Definition 6.1, let $\widehat{\Sigma} := \Sigma_{\{w\}}$ denote the thrice-punctured sphere obtained by attaching cylindrical ends to a neighborhood in Σ of the component of the constant s locus corresponding to w. Let $\widehat{V}$ denote the space of $\widehat{\sigma} \in \widetilde{\operatorname{Coker}}(D_{\widehat{\Sigma}})$ such that if $\widehat{\sigma} \neq 0$, then for each edge e of $\tau(\Sigma)$ incident to w, $\widehat{\sigma}$ has winding number $\eta(e)$ around the corresponding cylinder in $\widehat{\Sigma}$. By equation (2.18), the winding numbers in the definition of $\widehat{V}$ satisfy the appropriate version of equation (6.1), so that $\dim_{\mathbb{C}} \widehat{V} = 1$. Thus there is a well-defined element

$$\Phi_{e,e'}(\widehat{\Sigma}) \in W(e)^* \otimes W(e')$$

as in (6.32). Define $\widehat{\Sigma}'$, $\widehat{V}'$, and $\Phi_{e,e'}(\widehat{\Sigma}')$ analogously from Σ'. Propositions 2.25 and 6.3 imply that for any $\varepsilon > 0$, if r is sufficiently large then

$$\begin{aligned} \Phi_{e,e'}(\Sigma) &= (1 + O(\varepsilon))\Phi_{e,e'}(\widehat{\Sigma}), \\ \Phi_{e,e'}(\Sigma') &= (1 + O(\varepsilon))\Phi_{e,e'}(\widehat{\Sigma}'). \end{aligned} \tag{6.34}$$

Step 2. We now prove the lemma when $v \neq w$. Here it follows from the definition that $r(e, e'; v) = 0$. On the other hand, $\Phi_{e,e'}(\widehat{\Sigma}) = \Phi_{e,e'}(\widehat{\Sigma}')$, because in passing from Σ to $\widehat{\Sigma}$ or from Σ' to $\widehat{\Sigma}'$, the location of the ramification point corresponding to w is forgotten. Thus (6.33) follows from (6.34).

Step 3. We now prove the lemma when $v = w$. Here

$$r(e, e'; v) = \frac{\eta(e)}{m(e)} - \frac{\eta(e')}{m(e')}. \tag{6.35}$$

Observe that there is an isomorphism $\widehat{\Sigma} \to \widehat{\Sigma}'$ covering the automorphism of $\mathbb{R} \times S^1$ that sends $(s, t) \mapsto (s, t + \varphi)$. Given an element $\widehat{\sigma} \in \widehat{V}$, we can push it forward via this isomorphism to obtain an element $\widehat{\sigma}' \in \widehat{V}'$. It follows from (6.30) and (6.31) that

$$\widehat{\sigma}'_e = \exp\left(-i\varphi\frac{\eta(e)}{m(e)}t\right)\widehat{\sigma}_e,$$

and likewise for e'. Therefore

$$\Phi_{e,e'}(\widehat{\Sigma}') = \exp\left(i\varphi\left(\frac{\eta(e)}{m(e)} - \frac{\eta(e')}{m(e')}\right)t\right)\Phi_{e,e'}(\widehat{\Sigma}). \tag{6.36}$$

We are now done by (6.34), (6.35), and (6.36). □

7 Application to embedded contact homology

As in §1.1, let Y be a closed oriented 3-manifold with a contact form λ whose Reeb orbits are nondegenerate, and let J be an admissible almost complex structure on $\mathbb{R} \times Y$. Out of these data one can define the embedded contact homology (ECH), which is the homology of a chain complex whose differential ∂ counts certain (mostly) embedded J-holomorphic curves in $\mathbb{R} \times Y$. The significance of ECH is that as explained in [12, §1.1], it is conjecturally isomorphic to versions of the Ozsváth-Szabó and Seiberg-Witten Floer homologies defined in [15, 14]. However, most of the foundations of ECH have not yet been established.

In this section, we apply the gluing formula of Theorem 1.13 in a special (but nontrivial) case to prove that the ECH differential ∂ satisfies $\partial^2 = 0$. Essentially the same argument shows that the differential in the periodic Floer homology of mapping tori [11] also has square zero.

After some combinatorial preliminaries in §7.1, the definition of the ECH differential ∂ is reviewed in §7.2. The proof that $\partial^2 = 0$ is given in §7.3, using a gluing coefficient calculation which is carried out in §7.4 and §7.5. This section uses only §1 (if one accepts the statement of Theorem 1.13), and is not used elsewhere in the paper.

7.1 Incoming and outgoing partitions

To prove that $\partial^2 = 0$, we need to know the multiplicities of the ends of the curves that are counted by ∂. The description of these multiplicities in §7.2 requires the following preliminary combinatorial definitions.

Fix an irrational number θ. For each nonnegative integer M, we now define two distinguished partitions of M, called the "incoming partition" and the "outgoing partition", and denoted here by $P_\theta^{\mathrm{in}}(M)$ and $P_\theta^{\mathrm{out}}(M)$ respectively.

Definition 7.1. [10, §4] Define the *incoming partition* $P_\theta^{\text{in}}(M)$ as follows. Let S_θ denote the set of all positive integers a such that

$$\frac{\lceil a\theta \rceil}{a} < \frac{\lceil a'\theta \rceil}{a'}, \qquad \forall a' \in \{1, \ldots, a-1\}.$$

Let $a := \max(S_\theta \cap \{1, \ldots, M\})$. Define $P_\theta^{\text{in}}(0) = \emptyset$, and inductively define[4]

$$P_\theta^{\text{in}}(M) := (a) \cup P_\theta^{\text{in}}(M-a).$$

Define the *outgoing partition*

$$P_\theta^{\text{out}}(M) := P_{-\theta}^{\text{in}}(M).$$

We will make frequent use of the following alternate description of the incoming and outgoing partitions.

Definition 7.2. Let $\Lambda_\theta^{\text{in}}(M)$ denote the lowest convex polygonal path in the plane that starts at $(0,0)$, ends at $(M, \lceil M\theta \rceil)$, stays above the line $y = \theta x$, and has corners at lattice points. That is, the boundary of the convex hull of the set of lattice points $(x,y) \in \mathbb{Z}^2$ such that $0 \le x \le M$ and $y \ge \theta x$ consists of the ray $(x = 0, y \ge 0)$, the path $\Lambda_\theta^{\text{in}}(M)$, and the ray $(x = M, y \ge \lceil M\theta \rceil)$.

Lemma 7.3. *The integers in the incoming partition $P_\theta^{\text{in}}(M)$ are the horizontal displacements of the segments of the path $\Lambda_\theta^{\text{in}}(M)$ between lattice points.*

Proof. Let $(x_1, y_1) \in \mathbb{Z}^2$ denote the first lattice point on the path $\Lambda_\theta^{\text{in}}(M)$, after the initial endpoint $(0,0)$. Since $\Lambda_\theta^{\text{in}}(M)$ is convex, there are no lattice points in the open region bounded by the lines $y = \theta x$, $y = (y_1/x_1)x$, and $x = M$. Hence

$$\frac{\lceil x_1\theta \rceil}{x_1} = \frac{y_1}{x_1} \le \frac{\lceil x\theta \rceil}{x}, \quad \forall x = 1, \ldots, M. \tag{7.1}$$

Also, since the vector $(x_1, y_1) \in \mathbb{Z}^2$ is indivisible, equality can hold in (7.1) only when $x \ge x_1$. It follows that $x_1 = \max(S_\theta \cap \{1, \ldots, M\})$. Moreover, the rest of the path $\Lambda_\theta^{\text{in}}(M)$ is the translation of the path $\Lambda_\theta^{\text{in}}(M - x_1)$ by (x_1, y_1). The lemma follows by induction. □

Likewise, let $\Lambda_\theta^{\text{out}}(M)$ denote the highest concave polygonal path in the plane which starts at $(0,0)$, ends at $(M, \lfloor M\theta \rfloor)$, stays below the line $y = \theta x$,

[4] If $P = (a_1, \ldots, a_k)$ is a partition of M and $Q = (b_1, \ldots, b_l)$ is a partition of N, define a partition $P \cup Q$ of $M + N$ by

$$P \cup Q := (a_1, \ldots, a_k, b_1, \ldots, b_l).$$

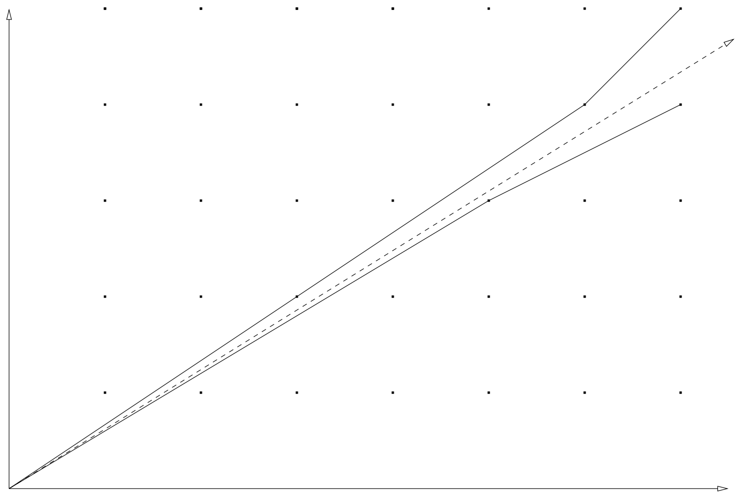

Figure 6: Here $3/5 < \theta < 2/3$, the dashed line is $y = \theta x$, the upper path is $\Lambda_\theta^{\mathrm{in}}(7)$, and the lower path is $\Lambda_\theta^{\mathrm{out}}(7)$. In particular, $P_\theta^{\mathrm{in}}(7) = (3,3,1)$ and $P_\theta^{\mathrm{out}}(7) = (5,2)$.

and has corners at lattice points. Then by Lemma 7.3, the integers in the outgoing partition $P_\theta^{\mathrm{out}}(M)$ are the horizontal displacements of the segments of the path $\Lambda_\theta^{\mathrm{out}}(M)$ between lattice points. An example is shown in Figure 6.

The following basic facts about the incoming and outgoing partitions will be needed later.

Lemma 7.4. *(a) If $P_\theta^{\mathrm{in}}(M) = (a_1, \ldots, a_k)$, then $\sum_{i=1}^k \lceil a_i\theta \rceil = \lceil M\theta \rceil$.*

(b) If $P_\theta^{\mathrm{out}}(M) = (b_1, \ldots, b_l)$ then $\sum_{j=1}^l \lfloor b_j\theta \rfloor = \lfloor M\theta \rfloor$.

Proof. This is an immediate consequence of the descriptions of $P_\theta^{\mathrm{in}}(M)$ and $P_\theta^{\mathrm{out}}(M)$ in terms of the paths $\Lambda_\theta^{\mathrm{in}}(M)$ and $\Lambda_\theta^{\mathrm{out}}(M)$. $\square$

Lemma 7.5. *Under the partial order $\geq_\theta$ on the set of partitions of M (see Definition 1.8), $P_\theta^{\mathrm{in}}(M)$ is maximal and $P_\theta^{\mathrm{out}}(M)$ is minimal.*

Proof. It is an exercise using either Definition 1.8 or Lemma 1.17 to show that $P \geq_\theta Q$ if and only if one can get from P to Q by a sequence of the following operations:

- Replace a_1, a_2 by $a_1 + a_2$ where $\lceil a_1\theta \rceil + \lceil a_2\theta \rceil = \lceil (a_1 + a_2)\theta \rceil$.
- Replace $a_1 + a_2$ by a_1, a_2 where $\lceil (a_1 + a_2)\theta \rceil = \lceil a_1\theta \rceil + \lceil a_2\theta \rceil - 1$.

Now to prove the lemma, by symmetry it is enough to show that $P_\theta^{\mathrm{in}}(M)$ is maximal. Suppose to the contrary that $Q \geq_\theta P_\theta^{\mathrm{in}}(M)$ and $Q \neq P_\theta^{\mathrm{in}}(M)$. Then at least one of the following situations occurs:

(i) $P_\theta^{\mathrm{in}}(M)$ contains m_1 and m_2 with $\lceil (m_1+m_2)\theta\rceil = \lceil m_1\theta\rceil + \lceil m_2\theta\rceil - 1$.

(ii) $P_\theta^{\mathrm{in}}(M)$ contains m_1+m_2 where $\lceil m_1\theta\rceil + \lceil m_2\theta\rceil = \lceil (m_1+m_2)\theta\rceil$.

In case (i), write $P_\theta^{\mathrm{in}}(M) = (m_1, m_2, \ldots, m_k)$. Then by Lemma 7.4(a),

$$\lceil (m_1+m_2)\theta\rceil + \sum_{i=3}^{k} \lceil m_i\theta\rceil = \sum_{i=1}^{k}\lceil m_i\theta\rceil - 1 = \lfloor M\theta\rfloor .$$

But this is impossible, since the left side is greater than $M\theta$ and the right side is smaller than $M\theta$.

In case (ii), since $m_1+m_2\in S_\theta$, by Lemma 7.3 the path $\Lambda_\theta^{\mathrm{in}}(m_1+m_2)$ is just a line segment from the origin to the point $(m_1+m_2, \lceil (m_1+m_2)\theta\rceil)$, and this line segment has no lattice points in its interior. By the definition of $\Lambda_\theta^{\mathrm{in}}(m_1+m_2)$, the slope of this line segment must be strictly less than that of the vectors $(m_1, \lceil m_1\theta\rceil)$ and $(m_2, \lceil m_2\theta\rceil)$. This contradicts (ii). □

Definition 7.6. The *standard ordering convention* for the incoming and outgoing partitions is to write $P_\theta^{\mathrm{in}}(M) = (a_1,\ldots,a_k)$ and $P_\theta^{\mathrm{out}}(M) = (b_1,\ldots,b_l)$ where

$$a_i \geq a_{i+1}, \qquad b_j \geq b_{j+1}.$$

Note that by Definition 7.1, this is equivalent to

$$\frac{\lceil a_i\theta\rceil}{a_i} \leq \frac{\lceil a_{i+1}\theta\rceil}{a_{i+1}}, \qquad \frac{\lfloor b_j\theta\rfloor}{b_j} \geq \frac{\lfloor b_{j+1}\theta\rfloor}{b_{j+1}}. \tag{7.2}$$

7.2 The ECH differential ∂

We now briefly review the definition of the differential ∂ in embedded contact homology, in preparation for showing that $\partial^2 = 0$.

A J-holomorphic curve may have several ends at covers of an embedded Reeb orbit γ, with various covering multiplicities. The ECH chain complex only keep tracks of the sum of these multiplicities. For this purpose we make the following definitions.

Definition 7.7. An *orbit set* is a finite set of pairs $\alpha = \{(\alpha_i, m_i)\}$ where the α_i's are distinct embedded Reeb orbits and the m_i's are positive integers[5]. Define $[\alpha] := \sum_i m_i[\alpha_i] \in H_1(Y)$. The orbit set α is *admissible* if $m_i = 1$ whenever α_i is hyperbolic. If $\beta = \{(\beta_j, n_j)\}$ is another orbit set with $[\alpha] = [\beta]$, define $H_2(Y,\alpha,\beta)$ to be the set of relative homology classes of 2-chains Z in Y with

$$\partial Z = \sum_i m_i\alpha_i - \sum_j n_j\beta_j.$$

[5]This is different from the notation in Definition 1.2, where α and β are ordered lists of Reeb orbits which might be multiply covered.

Thus $H_2(Y,\alpha,\beta)$ is an affine space over $H_2(Y)$.

Definition 7.8. If $\alpha = \{(\alpha_i, m_i)\}$ and $\beta = \{(\beta_j, n_j)\}$ are orbit sets with $[\alpha] = [\beta]$, let $\mathcal{M}^J(\alpha,\beta)$ denote the moduli space of J-holomorphic curves u with positive ends at covers of α_i with total multiplicity m_i, negative ends at covers of β_j with total multiplicity n_j, and no other ends. In contrast to Definition 1.2, the ends of u are not ordered or asymptotically marked. Note that the projection of each $u \in \mathcal{M}^J(\alpha,\beta)$ to Y has a well-defined relative homology class $[u] \in H_2(Y,\alpha,\beta)$. For $Z \in H_2(Y,\alpha,\beta)$ we then define

$$\mathcal{M}^J(\alpha,\beta,Z) := \{u \in \mathcal{M}^J(\alpha,\beta) \mid [u] = Z\}.$$

Definition 7.9. Given a homology class $\Gamma \in H_1(Y)$, the ECH chain complex $C_*(Y,\lambda;\Gamma)$ is a free $\mathbb{Z}$-module with one generator for each admissible orbit set α with $[\alpha] = \Gamma$.

To fix the signs in the differential below, for each admissible orbit set we need to choose an ordering of its positive hyperbolic orbits[6]. To simplify the discussion below, let us do this by fixing some ordering of all the embedded positive hyperbolic Reeb orbits in Y.

The relative index on this chain complex is defined as follows. (This should be contrasted with Definition 1.3.)

Definition 7.10. (cf. [10, §1]) If α and β are orbit sets with $[\alpha] = [\beta]$, and if $Z \in H_2(Y,\alpha,\beta)$, define the *ECH index*

$$I(\alpha,\beta,Z) := c_1(\xi|_Z,\tau) + Q_\tau(Z) + \sum_i \sum_{k=1}^{m_i} \mathrm{CZ}_\tau(\alpha_i^k) - \sum_j \sum_{k=1}^{n_j} \mathrm{CZ}_\tau(\beta_j^k).$$

Here τ is a trivialization of ξ over the α_i's and β_j's. As in §1.1, c_1 denotes the relative first Chern class of ξ over a surface representing Z, and $\mathrm{CZ}_\tau(\gamma^k)$ denotes the Conley-Zehnder index of the k^{th} iterate of γ. Also, Q_τ is the "relative self-intersection pairing" defined in [10, §2]. If $u \in \mathcal{M}^J(\alpha,\beta,Z)$, write $I(u) := I(\alpha,\beta,Z)$.

It is shown in [10] that I depends only on the orbit sets α and β and on the relative homology class Z. Also, I is additive in the following sense: if γ is another orbit set with $[\beta] = [\gamma]$ and if $W \in H_2(Y,\beta,\gamma)$, then there is a well-defined relative homology class $Z + W \in H_2(Y,\alpha,\gamma)$, and we have

$$I(\alpha,\gamma,Z+W) = I(\alpha,\beta,Z) + I(\beta,\gamma,W).$$

[6] Alternately one can define the chain complex to be generated by admissible orbit sets in which the positive hyperbolic orbits are ordered, modulo the relation that reordering the positive hyperbolic orbits in a generator multiplies the generator by the sign of the reordering permutation.

The key nontrivial property of the ECH index I is that it gives an upper bound on the Fredholm index ind from Definition 1.3. Moreover, curves that realize this upper bound are highly restricted. To give the precise statements, we need the following definitions.

Definition 7.11. (cf. [10, §4]) If γ is an embedded Reeb orbit and M is a positive integer, define two partitions of M, the *incoming partition* $P_\gamma^{\mathrm{in}}(M)$ and the *outgoing partition* $P_\gamma^{\mathrm{out}}(M)$, as follows.

- If γ is positive hyperbolic, then

$$P_\gamma^{\mathrm{in}}(M) := P_\gamma^{\mathrm{out}}(M) := (1, \ldots, 1). \tag{7.3}$$

- If γ is negative hyperbolic, then

$$P_\gamma^{\mathrm{in}}(M) := P_\gamma^{\mathrm{out}}(M) := \begin{cases} (2, \ldots, 2), & \text{if } M \text{ is even}, \\ (2, \ldots, 2, 1), & \text{if } M \text{ is odd}. \end{cases} \tag{7.4}$$

- If γ is elliptic with monodromy angle θ, then (see §7.1)

$$P_\gamma^{\mathrm{in}}(M) := P_\theta^{\mathrm{in}}(M), \qquad P_\gamma^{\mathrm{out}}(M) := P_\theta^{\mathrm{out}}(M).$$

The *standard ordering convention* for $P_\gamma^{\mathrm{in}}(M)$ or $P_\gamma^{\mathrm{out}}(M)$ is to list the entries in nonincreasing order.

Notation 7.12. Any J-holomorphic curve $u \in \mathcal{M}^J(\alpha, \beta)$ can be uniquely written as $u = u_0 \cup u_1$, where u_0 and u_1 are unions of components of u, each component of u_0 maps to an $\mathbb{R}$-invariant cylinder, and no component of u_1 does. Given an embedded Reeb orbit γ, let n_γ denote the total multiplicity of covers of $\mathbb{R} \times \gamma$ in u_0. Let m_γ^+ denote the total multiplicity of all positive ends of u at covers of γ, and let P_γ^+ denote the partition of $m_\gamma^+ - n_\gamma$ consisting of the multiplicities of the positive ends of u_1 at covers of γ. Define m_γ^- and P_γ^- analogously for the negative ends.

Definition 7.13. $u = u_0 \cup u_1 \in \mathcal{M}^J(\alpha, \beta)$ is *admissible* if:

(a) u_1 is embedded and does not intersect u_0.

(b) For each embedded Reeb orbit γ, under the standard ordering convention:

  - P_γ^+ is an initial segment of $P_\gamma^{\mathrm{out}}(m_\gamma^+)$.
  - P_γ^- is an initial segment of $P_\gamma^{\mathrm{in}}(m_\gamma^-)$.

We can now state the key index inequality.

Proposition 7.14. *Let $u = u_0 \cup u_1 \in \mathcal{M}^J(\alpha, \beta)$ and suppose that u_1 is not multiply covered. Then:*

(a) $\operatorname{ind}(u_1) \leq I(u_1) - 2\delta(u_1)$, *with equality only if for each embedded Reeb orbit γ:*

- $P_\gamma^+ = P_\gamma^{\mathrm{out}}(m_\gamma^+ - n_\gamma)$.
- $P_\gamma^- = P_\gamma^{\mathrm{in}}(m_\gamma^- - n_\gamma)$.

(b) $I(u_1) \leq I(u) - 2\#(u_0 \cap u_1)$, *with equality only if the following hold for each embedded Reeb orbit γ, under the standard ordering convention:*

- $P_\gamma^{\mathrm{out}}(m_\gamma^+ - n_\gamma)$ *is an initial segment of* $P_\gamma^{\mathrm{out}}(m_\gamma^+)$.
- $P_\gamma^{\mathrm{in}}(m_\gamma^- - n_\gamma)$ *is an initial segment of* $P_\gamma^{\mathrm{in}}(m_\gamma^-)$.

Here $\delta(u_1)$ is a count of the singularities of u_1 with positive integer weights; in particular $\delta(u_1) \geq 0$, with equality if and only if u_1 is embedded. Also $\#(u_0 \cap u_1)$ is the algebraic intersection number; by intersection positivity, each intersection point counts positively.

Proof. Part (a) is proved in [10, Eq. (18) and Prop. 6.1], and part (b) is proved in [10, Prop. 7.1], except for two issues. First, these results are proved in [10] in a slightly different setting where Y is a mapping torus and an analytical simplifying assumption ("local linearity") is made. The asymptotic analysis needed to transfer these results to the present setting is carried out in [17]. Second, the necessary condition for equality in part (b) is different from the one given in [10, Prop. 7.1]. However these two conditions are equivalent by Lemma 7.28(a),(d) below. □

We can now classify the curves with small ECH index for generic J.

Proposition 7.15. *Suppose that J is generic and $u = u_0 \cup u_1 \in \mathcal{M}^J(\alpha, \beta)$. Then:*

(a) $I(u) \geq 0$.

(b) If $I(u) = 0$, *then* $u_1 = \emptyset$.

(c) If $I(u) = 1$, *then u is admissible and* $\operatorname{ind}(u_1) = 1$.

(d) If $I(u) = 2$ *and α and β are admissible, then u is admissible and* $\operatorname{ind}(u_1) = 2$.

Proof. (This is based on [10, Lem. 9.5] with simplifications from [12, Cor. 11.5].) The image of u_1 is the union of k irreducible components $v_1, \ldots, v_k$, covered by u with positive integer multiplicities $d_1, \ldots, d_k$. Since J is generic, $\operatorname{ind}(v_i) \geq 1$ for each i.

Let u_1' be the union of d_i translates of v_i for $i = 1, \ldots, k$. Then $\operatorname{ind}(u_1') = \sum_{i=1}^k d_i \operatorname{ind}(v_i)$ by definition, and $I(u_1') = I(u_1)$ since u_1' and u_1 go between the same orbit sets and have the same relative homology class. So by Proposition 7.14(a) applied to u_1' and Proposition 7.14(b) applied to $u_0 \cup u_1'$, we obtain

$$\sum_{i=1}^{k} d_i \operatorname{ind}(v_i) \le I(u) - 2\delta(u_1') - 2\#(u_0 \cap u_1), \tag{7.5}$$

with equality only if condition (b) in Definition 7.13 holds. Parts (a)–(c) follow immediately from (7.5).

To prove part (d), note that if $I(u) = 2$ then $k > 0$, because a union of $\mathbb{R}$-invariant cylinders has $I = 0$. Furthermore the left hand side of (7.5) must equal 2, because by [10, Lem. 9.4], if α and β are admissible then $\operatorname{ind}(u)$ and $I(u)$ have the same parity. Now there are three possibilities: (i) $k = 2$ and $d_1 = d_2 = 1$; (ii) $k = 1$ and $d_1 = 1$; (iii) $k = 1$ and $d_1 = 2$. In cases (i) and (ii) we are done.

To complete the proof we now rule out case (iii). In this case we must have $\operatorname{ind}(v_1) = 1$. However, since α and β are admissible, and since $d_1 > 1$, all Reeb orbits in α and β are elliptic. Since elliptic orbits have odd Conley-Zehnder index, it follows from the definition of ind and the formula for the Euler characteristic of a surface that $\operatorname{ind}(v_1)$ is even, a contradiction. □

The differential ∂ in ECH counts $I = 1$ curves in $\mathcal{M}^J(\alpha, \beta)/\mathbb{R}$ where α and β are admissible orbit sets. Such curves may contain multiple covers of the $\mathbb{R}$-invariant cylinder $\mathbb{R} \times \gamma$ when γ is an elliptic embedded Reeb orbit. The differential ∂ only keeps track of the total multiplicity of such coverings for each γ. We now give the precise definition of ∂, in notation which will be convenient for the proof that $\partial^2 = 0$.

Definition 7.16. Let α and β be orbit sets. Define $\mathcal{M}_1^J(\alpha, \beta, Z)$ to be the set of curves $u \in \mathcal{M}^J(\alpha, \beta, Z)$ such that if γ is an elliptic Reeb orbit, then u does not contain $\mathbb{R} \times \gamma$ or any cover thereof.

Notation 7.17. If α and β are orbit sets, define a "product" orbit set $\alpha\beta$ by adding the multiplicities of all embedded Reeb orbits involved. (The index and differential are not well-behaved with respect to this "multiplication".) Write $\alpha|\beta$ if β is divisible by α in this sense, in which case denote the quotient by β/α. Call an orbit set "elliptic" if all of its Reeb orbits are elliptic.

Definition 7.18. Given a generic J and a system of coherent orientations, define the ECH differential

$$\partial : C_*(Y, \lambda, \Gamma) \longrightarrow C_{*-1}(Y, \lambda, \Gamma)$$

as follows. If α and β are admissible orbit sets with $[\alpha] = [\beta] = \Gamma$, then the coefficient of β in $\partial\alpha$ is

$$\langle\partial\alpha, \beta\rangle := \sum_{\substack{Z\in H_2(Y,\alpha,\beta)\\ I(\alpha,\beta,Z)=1}} \left(\sum_{\substack{\gamma \text{ elliptic orbit set}\\ \gamma|\alpha,\beta}} \# \frac{\mathcal{M}_1^J(\alpha/\gamma, \beta/\gamma, Z - [\mathbb{R}\times\gamma])}{\mathbb{R}} \right) \cdot \beta$$

Here the symbol '#' indicates the signed count.

To see why $\langle\partial\alpha, \beta\rangle$ is well-defined, first note that the set being counted is zero-dimensional, because by Proposition 7.15(c), if $I(\alpha, \beta, Z) = 1$ and

$$u_1 \in \mathcal{M}_1^J(\alpha/\gamma, \beta/\gamma, Z - [\mathbb{R}\times\gamma])$$

then $\operatorname{ind}(u_1) = 1$. By Remark 1.5, the sign $\epsilon(u_1)$ is well-defined, because admissibility of α and β ensures that u_1 does not have an end at a double cover of a negative hyperbolic orbit or more than one end at a positive hyperbolic orbit, and an ordering of all positive hyperbolic orbits has been chosen. Finiteness of the count results from the following compactness lemma.

Lemma 7.19. *If α and β are (not necessarily admissible) orbit sets and J is generic, then the set*

$$\bigsqcup_{\substack{Z\in H_2(Y,\alpha,\beta)\\ I(\alpha,\beta,Z)=1}} \mathcal{M}_1^J(\alpha, \beta, Z)/\mathbb{R}$$

is compact (and therefore finite).

Proof. (Cf. [10, §9].) Let $\{u_n\}$ be a sequence of curves in $\mathcal{M}_1^J(\alpha,\beta)/\mathbb{R}$ with $I(u_n) = 1$. By Stokes' theorem, the "energy" of u_n, namely the integral of $u_n^* d\lambda$ over the domain of u_n, is

$$\int_{u_n} d\lambda = \int_\alpha \lambda - \int_\beta \lambda,$$

which does not depend on n. So by Gromov compactness as in [10, Lem. 9.8], we can pass to a subsequence so that $\{u_n\}$ converges in the sense of [2] to a (possibly) broken curve with $\operatorname{ind} = 1$.

By Proposition 7.15(a) and the additivity of the ECH index, one level of the broken curve has $I = 1$ and all other levels have $I = 0$. By Proposition 7.15(b), Lemma 1.7, and the additivity of ind, the $I = 0$ levels also have $\operatorname{ind} = 0$. Then the top level cannot have $I = 0$, or else by Lemmas 1.7 and 7.5 it would have $\operatorname{ind} \geq 1$. Likewise the bottom level cannot have $I = 0$. Hence there is only one level.

The limiting curve cannot contain a cover of $\mathbb{R}\times\gamma$ with γ elliptic, because the u_n's contain no such covers, and any J-holomorphic curve in the same moduli space component as a cover of $\mathbb{R}\times\gamma$ is itself a cover of $\mathbb{R}\times\gamma$ because it has energy zero. $\square$

7.3 Proof that $\partial^2 = 0$

Theorem 7.20. *If J is generic, then the ECH differential ∂ satisfies $\partial^2 = 0$.*

The proof of Theorem 7.20 follows the standard strategy of analyzing ends of moduli spaces of $I = 2$ curves, and consists of a compactness argument and a gluing argument. The following are the kinds of pairs of curves that we will need to glue.

Definition 7.21. Let α_+ and α_- be admissible orbit sets. An *ECH gluing pair* is a pair of curves $u_+ \in \mathcal{M}^J(\alpha_+, \beta)$ and $u_- \in \mathcal{M}^J(\beta, \alpha_-)$ such that:

(a) $I(u_+) = I(u_-) = 1$.

(b) For each embedded elliptic Reeb orbit γ:

  (i) All covers of $\mathbb{R} \times \gamma$ in u_+ and u_- are unbranched.

  (ii) If u_+ (resp. u_-) contains covers of $\mathbb{R}\times\gamma$ with total multiplicity n_γ^+ (resp. n_γ^-), then the individual multiplicities comprise the outgoing partition $P_\gamma^{\text{out}}(n_\gamma^+)$ (resp. the incoming partition $P_\gamma^{\text{in}}(n_\gamma^-)$).

  (iii) u_+ and u_- do not both contain covers of $\mathbb{R} \times \gamma$.

To glue ECH gluing pairs, we will apply Theorem 1.13, for which purpose we will need the following calculation of gluing coefficients. If P is a partition in which the positive integer n appears $r(n)$ times, define

$$P! := \prod_{n=1}^{\infty} n^{r(n)} \cdot r(n)!.$$

In particular, if P is the (empty) partition of 0, then $P! = 1$.

Proposition 7.22. *Given integers $0 \le M_+, M_- \le M$, define*

$$S := (P_\theta^{\text{in}}(M_+); P_\theta^{\text{out}}(M - M_+) \mid P_\theta^{\text{out}}(M_-); P_\theta^{\text{in}}(M - M_-)).$$

Suppose that under the standard ordering convention,

$$\begin{array}{l} P_\theta^{\text{in}}(M_+) \text{ is an initial segment of } P_\theta^{\text{in}}(M), \text{ and} \\ P_\theta^{\text{out}}(M_-) \text{ is an initial segment of } P_\theta^{\text{out}}(M). \end{array} \tag{7.6}$$

Then:

(a) If $M_-, M_+ < M$, then for every θ-decomposition of S, see Definition 1.16, there exists ν with $I_\nu = J_\nu = \emptyset$ and $|I'_\nu| = |J'_\nu| = 1$.

(b) If $M_- = M$ then $c_\theta(S) = P_\theta^{\text{out}}(M - M_+)!$.

(c) If $M_+ = M$ then $c_\theta(S) = P_\theta^{\mathrm{in}}(M - M_-)!$.

The proof of this proposition is deferred to §7.4 and §7.5. We can now carry out the compactness part of the proof that $\partial^2 = 0$ and see how ECH gluing pairs arise.

Lemma 7.23. *Assume that J is generic and let α_+ and α_- be admissible orbit sets. Let $\{u_n\}$ be a sequence of curves in $\mathcal{M}_1^J(\alpha_+, \alpha_-)/\mathbb{R}$ such that $I(u_n) = 2$. Then after passing to a subsequence, $\{u_n\}$ converges in the sense of [2] either to a curve in $\mathcal{M}_1^J(\alpha_+, \alpha_-)/\mathbb{R}$, or to a broken curve $(u_+, \tau_1, \ldots, \tau_k, u_-)$ for some $k \geq 0$, such that each τ_i maps to a union of $\mathbb{R}$-invariant cylinders and (u_+, u_-) is an ECH gluing pair.*

Proof. As in the proof of Lemma 7.19, we can pass to a subsequence so that $\{u_n\}$ converges to a (possibly) broken curve with ind $= 2$, in which each level has $I \geq 0$, and the ECH indices of the levels sum to 2. The top level must have $I > 0$; otherwise, since α_+ is admissible, by Proposition 7.15(b) and Lemma 7.5 it would have ind ≥ 2, contradicting additivity of ind for the broken curve. Likewise the bottom level has $I > 0$.

Suppose there are at least two levels. Then it follows that the limiting broken curve has the form $(u_+, \tau_1, \ldots, \tau_k, u_-)$ where $I(u_+) = I(u_-) = 1$ and $I(\tau_i) = 0$ for all i. By Proposition 7.15(a), each τ_i maps to a union of $\mathbb{R}$-invariant cylinders.

To complete the proof we must verify condition (b) in the definition of ECH gluing pair. Let γ be an embedded elliptic Reeb orbit. By Proposition 7.15(c),(d), the u_n's and $u_\pm$ are admissible. It then follows from Definition 7.1 (cf. Lemma 7.28(a) below) that the multiplicities of the positive ends of u_+ (resp. negative ends of u_-) at covers of γ must comprise the outgoing (resp. incoming) partition of n_γ^+ (resp. n_γ^-). Assertion (i) now follows from Lemma 7.5 and additivity of ind as before. Assertion (ii) then follows from the above description of the multiplicities of the positive ends of u_+ and negative ends of u_-.

To prove assertion (iii), let θ denote the monodromy angle of γ, and let m_γ denote the total multiplicity of the negative ends of u_+ at covers of γ. We can glue $\tau_1, \ldots, \tau_k$ to obtain an index zero branched cover $\pi : \Sigma \to \mathbb{R} \times \gamma$, where each positive end of Σ is paired with a negative end of u_+, and each negative end of Σ is paired with a positive end of u_-. The multiplicities of the ends of the components of Σ determine a θ-decomposition of

$$S := (P_\theta^{\mathrm{in}}(m_\gamma - n_\gamma^+); P_\theta^{\mathrm{out}}(n_\gamma^+) \mid P_\theta^{\mathrm{out}}(m_\gamma - n_\gamma^-); P_\theta^{\mathrm{in}}(n_\gamma^-)).$$

If assertion (iii) is false, then Proposition 7.22(a) implies that Σ has a cylinder component which is attached to $\mathbb{R}$-invariant cylinders in u_+ and u_-. This contradicts the fact that the u_n's have no components mapping to $\mathbb{R} \times \gamma$. □

We now apply Theorem 1.13 to deduce the gluing lemma that will be needed in the proof that $\partial^2 = 0$. Note that by Lemma 7.5, an ECH gluing pair becomes a gluing pair as in Definition 1.9 after orderings and asymptotic markings of the ends of $u_\pm$ are chosen.

Lemma 7.24. *Assume J is generic and let (u_+, u_-) be an ECH gluing pair. If orderings and asymptotic markings of the ends of $u_\pm$ are chosen, then:*

(a) If β is not admissible then $\#G(u_+, u_-) = 0$.

(b) If β is admissible then

$$\#G(u_+, u_-) = \epsilon(u_+)\epsilon(u_-) \prod_{\gamma \text{ elliptic embedded Reeb orbit}} P^{\text{out}}_\gamma(n^+_\gamma)! P^{\text{in}}_\gamma(n^-_\gamma)!. \tag{7.7}$$

Proof. For each embedded Reeb orbit γ, let m_γ denote the total multiplicity of negative ends of u_+ at covers of γ. By Theorem 1.13, it is enough to show:

(c) If γ is hyperbolic and $m_\gamma = 1$ then $c_\gamma(u_+, u_-) = 1$.

(d) If γ is hyperbolic and $m_\gamma > 1$ then $c_\gamma(u_+, u_-) = 0$.

(e) If γ is elliptic then $c_\gamma(u_+, u_-) = P^{\text{out}}_\gamma(n^+_\gamma)! P^{\text{in}}_\gamma(n^-_\gamma)!$.

Assertion (c) follows immediately from Definition 1.14.

To prove (d), suppose γ is hyperbolic and $m_\gamma > 1$. Recall from Proposition 7.15(c) that u_+ and u_- are admissible. If γ is positive hyperbolic, this means that all ends of u_+ and u_- at (covers of) γ have multiplicity 1. It then follows immediately from Definition 1.14(b) that $c_\gamma(u_+, u_-) = 0$. If γ is negative hyperbolic, then admissibility implies that u_+ has at least one negative end at a double cover of γ. Then $c_\gamma(u_+, u_-) = 0$ by Definition 1.14(c).

Assertion (e) follows immediately from Proposition 7.22(b),(c), thanks to condition (b) in the definition of ECH gluing pair, and the admissibility of u_+ and u_-. □

Proof of Theorem 7.20. Let α_+ and α_- be admissible orbit sets. We will prove that $\langle \partial^2 \alpha_+, \alpha_- \rangle = 0$ in two steps.

Step 1. We first show that

$$\sum_{\beta \text{ admissible}} \sum_{\substack{\gamma_+ | \beta, \alpha_+ \\ \gamma_- | \beta, \alpha_-}} \sum_{Z_+, Z_-} \# \frac{\mathcal{M}^J_1(\alpha_+/\gamma_+, \beta/\gamma_+, Z_+ - [\mathbb{R} \times \gamma_+])}{\mathbb{R}} \cdot \# \frac{\mathcal{M}^J_1(\beta/\gamma_-, \alpha_-/\gamma_-, Z_- - [\mathbb{R} \times \gamma_-])}{\mathbb{R}} = 0. \tag{7.8}$$

Here γ_+ and γ_- are elliptic orbit sets with no common factor, while $Z_+ \in H_2(Y, \alpha_+, \beta)$ and $Z_- \in H_2(Y, \beta, \alpha_-)$ satisfy $I(\alpha_+, \beta, Z_+) = I(\beta, \alpha_-, Z_-) = 1$.

To prove (7.8), we study the ends of the one-dimensional manifold

$$\mathcal{M} := \coprod_{\substack{Z \in H_2(Y, \alpha_+, \alpha_-) \\ I(\alpha_+, \alpha_-, Z) = 2}} \frac{\mathcal{M}_1^J(\alpha_+, \alpha_-, Z)}{\mathbb{R}}.$$

If (u_+, u_-) is an ECH gluing pair, in which $u_+ \in \mathcal{M}^J(\alpha_+, \beta)$ and $u_- \in \mathcal{M}^J(\beta, \alpha_-)$ for some orbit set β, let $V(u_+, u_-) \subset \mathcal{M}_1^J(\alpha_+, \alpha_-)/\mathbb{R}$ be an open set like the open set U in Definition 1.12, but where the curves do not have asymptotic markings or orderings of the ends. Define

$$\overline{\mathcal{M}} := \mathcal{M} \setminus \coprod_{(u_+, u_-)} V(u_+, u_-).$$

By Lemma 7.23, $\overline{\mathcal{M}}$ is compact. Thus the signed count of boundary points is

$$0 = \#\partial\overline{\mathcal{M}} = \sum_{(u_+, u_-)} -\#\partial\overline{V(u_+, u_-)}.$$

To understand this sum, let $v_\pm$ denote the $\operatorname{ind} = 1$ component of $u_\pm$. Then

$$\begin{aligned} v_+ &\in \mathcal{M}_1^J(\alpha_+/\gamma_+, \beta/\gamma_+, Z_+ - [\mathbb{R} \times \gamma_+])/\mathbb{R}, \\ v_- &\in \mathcal{M}_1^J(\beta/\gamma_-, \alpha_-/\gamma_-, Z_- - [\mathbb{R} \times \gamma_-])/\mathbb{R}, \end{aligned} \tag{7.9}$$

where $\gamma_\pm$ and $Z_\pm$ are as above. Thus

$$\#\partial\overline{\mathcal{M}} = \sum_{\beta} \sum_{\substack{\gamma_+ | \alpha_+, \beta \\ \gamma_- | \beta, \alpha_-}} \sum_{Z_+, Z_-} \sum_{v_+, v_- \text{ as in (7.9)}} -\#\partial\overline{V(u_+, u_-)}. \tag{7.10}$$

By Lemma 7.24,

$$-\#\partial\overline{V(u_+, u_-)} = \begin{cases} 0, & \text{if } \beta \text{ is not admissible}, \\ \epsilon(v_+)\epsilon(v_-), & \text{if } \beta \text{ is admissible}. \end{cases} \tag{7.11}$$

Let us clarify the signs and factorials here. First, the signs $\epsilon(v_+)$ and $\epsilon(v_-)$ are well-defined when β is admissible. If β is not admissible, then to apply Lemma 7.24 one needs to choose some orderings and asymptotic markings of the ends of $u_\pm$. However, $-\#\partial\overline{V(u_+, u_-)}$ is defined independently of this choice. Second, the factorials in (7.7) have disappeared in (7.11), because the count $\#G(u_+, u_-) = -\#\partial\overline{U}$ distinguishes curves in $\partial\overline{U}$ that have different asymptotic markings and orderings of the ends but represent the same element of $\partial\overline{V(u_+, u_-)}$. More precisely, given $v \in \partial\overline{V(u_+, u_-)}$, the corresponding curves $u \in \partial\overline{U}$ differ from each other by the following operations:

- changing the asymptotic marking of a positive (resp. negative) end of u that corresponds to an $\mathbb{R}$-invariant component of u_+ (resp. u_-).
- switching the ordering of two positive (resp. negative) ends of u that correspond to identical $\mathbb{R}$-invariant components of u_+ (resp. u_-).

Since γ_+ and γ_- have no common factor, it follows that $\#(\partial\overline{U})$ equals $\#\partial\overline{V(u_+,u_-)}$ times the product of factorials in equation (7.7).

Since $\#\partial\overline{\mathcal{M}} = 0$, equations (7.10) and (7.11) imply (7.8).

Step 2. By definition, the coefficient of α_- in $\partial^2\alpha_+$ is given by

$$\begin{aligned}\langle\partial^2\alpha_+,\alpha_-\rangle &= \sum_{\beta \text{ admissible}} \langle\partial\alpha_+,\beta\rangle\langle\partial\beta,\alpha_-\rangle \\ &= \sum_{\beta \text{ admiss.}} \sum_{\substack{\gamma_+|\beta,\alpha_+\\ \gamma_-|\beta,\alpha_-}} \#\frac{\mathcal{M}_1^J(\alpha_+/\gamma_+,\beta/\gamma_+)}{\mathbb{R}}\cdot\#\frac{\mathcal{M}_1^J(\beta/\gamma_-,\alpha_-/\gamma_-)}{\mathbb{R}}.\end{aligned}$$

In the second line, γ_+ and γ_- are elliptic orbit sets. We are also implicitly summing over relative homology classes with $I = 1$, which are suppressed here in order to simplify the notation. To process the above sum, let γ_0 denote the greatest common divisor of γ_+ and γ_-. Then after dividing γ_+ and γ_- by γ_0, the above sum becomes

$$\sum_{\beta \text{ admissible}} \sum_{\substack{\gamma_0|\alpha_+,\beta,\alpha_-\\ \gamma_0\gamma_+|\beta,\alpha_+\\ \gamma_0\gamma_-|\beta,\alpha_-}} \#\frac{\mathcal{M}_1^J(\alpha_+/\gamma_0\gamma_+,\beta/\gamma_0\gamma_+)}{\mathbb{R}}\cdot\#\frac{\mathcal{M}_1^J(\beta/\gamma_0\gamma_-,\alpha_-/\gamma_0\gamma_-)}{\mathbb{R}}.$$

Here γ_0, γ_+, and γ_- are elliptic orbit sets such that γ_+ and γ_- have no common factor. Now we can sum over γ_0 first and divide β by γ_0 to obtain

$$\sum_{\gamma_0|\alpha_+,\alpha_-} \sum_{\beta \text{ admiss.}} \sum_{\substack{\gamma_+|\beta,\alpha_+/\gamma_0\\ \gamma_-|\beta,\alpha_-/\gamma_0}} \#\frac{\mathcal{M}_1^J(\alpha_+/\gamma_0\gamma_+,\beta/\gamma_+)}{\mathbb{R}}\cdot\#\frac{\mathcal{M}_1^J(\beta/\gamma_-,\alpha_-/\gamma_0\gamma_-)}{\mathbb{R}}.$$

Again, γ_0, γ_+, and γ_- are elliptic orbit sets such that γ_+ and γ_- have no common factor. For each γ_0, by equation (7.8) applied to $\alpha_\pm/\gamma_0$, the above sum over β equals zero. This completes the proof that $\partial^2 = 0$. □

Remark 7.25. There is also a "twisted" version of ECH, with coefficients in the group ring over $H_2(Y)$ (or a quotient thereof), which keeps track of the relative homology classes of the J-holomorphic curves, see [12, §11.2]. The same argument with a bit more notation shows that $\partial^2 = 0$ for the twisted chain complex as well.

7.4 Calculation of ECH gluing coefficients, first half

To prepare for the proof of Proposition 7.22, we now establish a special case:

Proposition 7.26. *For any irrational number θ and positive integer M,*

$$c_\theta(P_\theta^{\mathrm{in}}(M) \mid P_\theta^{\mathrm{out}}(M)) = 1.$$

The proof of Proposition 7.26 uses induction on M. The key to carrying out the induction is the following lemma.

Lemma 7.27. *Write $P_\theta^{\mathrm{in}}(M) = (a_1, \ldots, a_k)$ and $P_\theta^{\mathrm{out}}(M) = (b_1, \ldots, b_l)$, with the standard ordering convention* (7.2)*. Then:*

(a) There is a unique subset $I = \{i_1 < \cdots < i_m\} \subset \{1, \ldots, l\}$ such that

$$\sum_{j=1}^{m-1} b_{i_j} < a_1 \le \sum_{j=1}^{m} b_{i_j}, \tag{7.12}$$

and moreover $I = \{1, \ldots, m\}$ for some m.

(b) With m as above, if $1 \le n \le m$, then

$$\delta_\theta\left(a_1 - \sum_{j=1}^{n-1} b_j \,;\, b_n\right) = 1. \tag{7.13}$$

(c) $P_\theta^{\mathrm{in}}(M - a_1) = (a_2, \ldots, a_k)$ with the standard ordering convention.

(d) Let $\overline{b} := \sum_{j=1}^{m} b_j - a_1$. Then with the standard ordering convention,

$$P_\theta^{\mathrm{out}}(M - a_1) = \left(\overline{b}, b_{m+1}, \ldots, b_l\right). \tag{7.14}$$

Proof. We begin with some preliminary remarks. Note that with the ordering convention (7.2), the lattice points on the path $\Lambda_\theta^{\mathrm{in}}(M)$ are the points

$$\sum_{i=1}^{n} \begin{pmatrix} a_i \\ \lceil a_i\theta \rceil \end{pmatrix}, \quad n = 0, \ldots, k,$$

while the lattice points on the path $\Lambda_\theta^{\mathrm{out}}(M)$ are

$$\sum_{j=1}^{n} \begin{pmatrix} b_j \\ \lfloor b_j\theta \rfloor \end{pmatrix}, \quad n = 0, \ldots, l.$$

Let $\Delta_\theta(M)$ denote the open region in the plane consisting of points (x, y) that (i) have $0 \le x \le M$, (ii) are strictly below the path $\Lambda_\theta^{\mathrm{in}}(M)$, and

(iii) are strictly above the path $\Lambda_\theta^{\mathrm{out}}(M)$. A key observation, which we will use repeatedly below, is that by construction *the region $\Delta_\theta(M)$ contains no lattice points*. Note also that by Lemma 7.4, we have

$$\kappa_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l) = 1. \tag{7.15}$$

Proof of (a): It will suffice to show that for each $n = 1, 2, \ldots$,

$$\sum_{j=1}^{n-1} b_j < a_1 \Longrightarrow M - b_n < a_1. \tag{7.16}$$

To see that (7.16) suffices, suppose that I satisfies (7.12), and let n be the smallest positive integer that is not in I. Suppose to get a contradiction that I contains an integer larger than n. Then the first inequality in (7.12) implies that $\sum_{j=1}^{n-1} b_j < a_1$. So by (7.16) we have $M - b_n < a_1$. Since $n \notin I$, the second inequality in (7.12) is then impossible.

To prove (7.16), suppose to the contrary that

$$\sum_{j=1}^{n-1} b_j < a_1, \qquad a_1 \le M - b_n. \tag{7.17}$$

Consider the lattice point in the plane

$$\begin{pmatrix} x \\ y \end{pmatrix} := \begin{pmatrix} a_1 \\ \lceil a_1\theta \rceil \end{pmatrix} + \begin{pmatrix} b_n \\ \lfloor b_n\theta \rfloor \end{pmatrix}.$$

To get a contradiction we will show that $(x, y) \in \Delta_\theta(M)$.

(i) To start, the second inequality in (7.17) implies that $x \le M$.

(ii) Next, (x, y) is strictly below the path $\Lambda_\theta^{\mathrm{in}}(M)$, because the vector $(a_1, \lceil a_1\theta \rceil)$ is on the path $\Lambda_\theta^{\mathrm{in}}(M)$, while the vector $(b_n, \lfloor b_n\theta \rfloor)$ points to the right and has slope less than that of all subsequent edges on the path $\Lambda_\theta^{\mathrm{in}}(M)$. Indeed, $(b_n, \lfloor b_n\theta \rfloor)$ has slope less than θ, while all of the edges in the path $\Lambda_\theta^{\mathrm{in}}(M)$ have slope greater than θ.

(iii) To see that (x, y) is strictly above the path $\Lambda_\theta^{\mathrm{out}}(M)$, rewrite (x, y) as a sum of two vectors as follows:

$$\begin{pmatrix} x \\ y \end{pmatrix} = \sum_{j=1}^{n} \begin{pmatrix} b_j \\ \lfloor b_j\theta \rfloor \end{pmatrix} + \left(\begin{pmatrix} a_1 \\ \lceil a_1\theta \rceil \end{pmatrix} - \sum_{j=1}^{n-1} \begin{pmatrix} b_j \\ \lfloor b_j\theta \rfloor \end{pmatrix} \right).$$

Then the first vector is on the path $\Lambda_\theta^{\mathrm{out}}(M)$, while the second vector points to the right (by the first inequality in (7.17)), and has slope greater than that of all subsequent edges in the path $\Lambda_\theta^{\mathrm{out}}(M)$ (because it has slope greater than $\lceil a_1\theta \rceil / a_1 > \theta$).

Proof of (b): For $n = 1, \ldots, m$, let T_n denote the triangle with vertices

$$\sum_{j=1}^{n-1} \begin{pmatrix} b_j \\ \lceil b_j\theta \rceil \end{pmatrix}, \quad \sum_{j=1}^{n} \begin{pmatrix} b_j \\ \lceil b_j\theta \rceil \end{pmatrix}, \quad \begin{pmatrix} a_1 \\ \lceil a_1\theta \rceil \end{pmatrix}.$$

Then the interior of T_n is in $\Delta_\theta(M)$, and hence contains no lattice points, and the interiors of the edges of T_n also contain no lattice points, by the definition of the incoming and outgoing partitions. It follows that T_n has area $1/2$, i.e.

$$\det \begin{pmatrix} b_n & a_1 - \sum_{j=1}^{n-1} b_j \\ \lfloor b_n\theta \rfloor & \lceil a_1\theta \rceil - \sum_{j=1}^{n-1} \lfloor b_j\theta \rfloor \end{pmatrix} = 1. \tag{7.18}$$

Next, in the notation of Definition 1.6, we have

$$\begin{aligned} \operatorname{ind}_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l) = \operatorname{ind}_\theta &\left(a_1 \mid b_1, \ldots, b_{n-1}, a_1 - \sum_{j=1}^{n-1} b_j \right) + \\ &+ \operatorname{ind}_\theta \left(a_1 - \sum_{j=1}^{n-1} b_j, a_2, \ldots, a_k \mid b_n, \ldots, b_l \right). \end{aligned}$$

We know by Lemma 7.4 that the left side of this equation equals zero, and the two terms on the right are nonnegative. In particular, the first term on the right must equal zero, so

$$\lceil a_1\theta \rceil - \sum_{j=1}^{n-1} \lfloor b_j\theta \rfloor = \left\lceil \left(a_1 - \sum_{j=1}^{n-1} b_j \right) \theta \right\rceil. \tag{7.19}$$

Equations (7.18) and (7.19) imply equation (7.13).

Proof of (c): This follows immediately from the definition of the incoming partition and the ordering convention (7.2).

Proof (d): We begin with some preliminary calculations. Suppose that $k > 1$. Then $\overline{b} > 0$ by (7.15). Next observe that

$$\begin{aligned} \operatorname{ind}_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l) = \operatorname{ind}_\theta &\left(a_2, \ldots, a_k \mid \overline{b}, b_{m+1} \ldots, b_l \right) + \\ &+ \operatorname{ind}_\theta \left(a_1, \overline{b} \mid b_1, \ldots, b_m \right). \end{aligned}$$

Similarly to (7.19), the second term on the right must vanish and so

$$\left\lfloor \overline{b}\theta \right\rfloor = \sum_{j=1}^{m} \lfloor b_j\theta \rfloor - \lceil a_1\theta \rceil. \tag{7.20}$$

We now prove (7.14) up to reordering. Consider the polygonal path Λ whose initial vertex is $(a_1, \lceil a_1\theta \rceil)$, and whose subsequent vertices are the sums $\sum_{j=1}^{n}(b_j, \lfloor b_j\theta \rfloor)$ for $n = m, \ldots, l$. Note that the interior of the initial edge of Λ contains no lattice points, because it is inside the region $\Delta_\theta(M)$. It then suffices to show that Λ is the path $\Lambda_\theta^{\mathrm{out}}(M - a_1)$ translated by $(a_1, \lceil a_1\theta \rceil)$.

To prove this, first note that by equation (7.20), the first edge of Λ has slope $\lfloor \overline{b}\theta \rfloor / \overline{b} < \theta$, and hence all edges of Λ have slope less than θ. Also, by (7.21) the path Λ is concave. Second, the final endpoint of the path Λ is

$$\begin{pmatrix} M \\ \lfloor M\theta \rfloor \end{pmatrix} = \begin{pmatrix} a_1 \\ \lceil a_1\theta \rceil \end{pmatrix} + \begin{pmatrix} M - a_1 \\ \lfloor (M - a_1)\theta \rfloor \end{pmatrix},$$

by Lemma 7.4(a) applied to M and $M - a_1$ with the help of part (c). Third, there are no lattice points above the path Λ and below the translate by $(a_1, \lceil a_1\theta \rceil)$ of the line $y = \theta x$, because any such lattice point would lie in the region $\Delta_\theta(M)$. This completes the proof of (7.14) up to reordering.

To show that (7.14) respects the standard ordering convention, it is enough to show that if $m < l$ then

$$\frac{\lfloor \overline{b}\theta \rfloor}{\overline{b}} \geq \frac{\lfloor b_{m+1}\theta \rfloor}{b_{m+1}}. \tag{7.21}$$

If $m < l$ and equation (7.21) fails, consider the lattice point

$$\begin{pmatrix} x \\ y \end{pmatrix} := \begin{pmatrix} a_1 \\ \lceil a_1\theta \rceil \end{pmatrix} + \begin{pmatrix} b_{m+1} \\ \lfloor b_{m+1}\theta \rfloor \end{pmatrix}. \tag{7.22}$$

To get a contradiction, we will show that $(x, y) \in \Delta_\theta(M)$. (i) First observe that $x = a_1 + b_{m+1} < \sum_{j=1}^{m+1} b_j \leq M$. (ii) As in the proof of (a), it follows from (7.22) that (x, y) is strictly below the path $\Lambda_\theta^{\mathrm{in}}(M)$. (iii) By (7.20),

$$\begin{pmatrix} x \\ y \end{pmatrix} = \sum_{j=1}^{m+1} \begin{pmatrix} b_j \\ \lfloor b_j\theta \rfloor \end{pmatrix} - \begin{pmatrix} \overline{b} \\ \lfloor \overline{b}\theta \rfloor \end{pmatrix}. \tag{7.23}$$

By our assumption that (7.21) fails, the vector $(\overline{b}, \lfloor \overline{b}\theta \rfloor)$ has strictly smaller slope than the vectors $(b_j, \lfloor b_j \rfloor)$ for $j = 1, \ldots, m + 1$, so (x, y) is strictly above the path $\Lambda_\theta^{\mathrm{out}}(M)$. □

Proof of Proposition 7.26. By equation (7.15) and the definition of c_θ, we have

$$c_\theta(P_\theta^{\mathrm{in}}(M) \mid P_\theta^{\mathrm{out}}(M)) = f_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l). \tag{7.24}$$

By Lemma 7.27(a) and the definition of f_θ, we have

$$\begin{aligned} f_\theta(a_1, \ldots, a_k \mid b_1, \ldots, b_l) = f_\theta\left(a_2, \ldots, a_k \mid \overline{b}, b_{m+1}, \ldots, b_l\right) \cdot \\ \cdot \prod_{n=1}^{m} \delta_\theta\left(a_1 - \sum_{j=1}^{n-1} b_j\,,\, b_n\right). \end{aligned}$$

Then by Lemma 7.27(b),

$$f_\theta(a_1,\ldots,a_k \mid b_1,\ldots,b_l) = f_\theta\left(a_2,\ldots,a_k \mid \overline{b}, b_{m+1},\ldots,b_l\right). \tag{7.25}$$

By Lemma 7.27(c),(d),

$$f_\theta\left(a_2,\ldots,a_k \mid \overline{b}, b_{m+1},\ldots,b_l\right) = c_\theta(P_\theta^{\mathrm{in}}(M-a_1) \mid P_\theta^{\mathrm{out}}(M-a_1)). \tag{7.26}$$

Proposition 7.26 follows from (7.24), (7.25), and (7.26) by induction on k. $\square$

7.5 Calculation of ECH gluing coefficients, second half

We now prove Proposition 7.22. We begin by clarifying the hypothesis (7.6) in the statement of the proposition. If Λ_1 and Λ_2 are two paths in the plane, let $\Lambda_1\Lambda_2$ denote the concatenated path that first traverses Λ_1 and then traverses the appropriate translate of Λ_2.

Lemma 7.28. *For $0 \le M' \le M$, the following are equivalent:*

(a) $P_\theta^{\mathrm{in}}(M'+n) = P_\theta^{\mathrm{in}}(M') \cup P_\theta^{\mathrm{in}}(n)$ for all $n = 1,\ldots,M-M'$.

(b) $\lceil (M'+n)\theta \rceil = \lceil M'\theta \rceil + \lceil n\theta \rceil$ for all $n = 1,\ldots,M-M'$.

(c) $\Lambda_\theta^{\mathrm{in}}(M) = \Lambda_\theta^{\mathrm{in}}(M')\Lambda_\theta^{\mathrm{in}}(M-M')$.

(d) Under the standard ordering convention, $P_\theta^{\mathrm{in}}(M')$ is an initial segment of $P_\theta^{\mathrm{in}}(M)$.

Proof. (a) $\Rightarrow$ (b): For a given n, if $P_\theta^{\mathrm{in}}(M'+n) = P_\theta^{\mathrm{in}}(M') \cup P_\theta^{\mathrm{in}}(n)$, then it follows from Lemma 7.3 that the edge vectors in the path $\Lambda_\theta^{\mathrm{in}}(M'+n)$ are the same as the edge vectors in the path $\Lambda_\theta^{\mathrm{in}}(M)\Lambda_\theta^{\mathrm{in}}(n)$, possibly in a different order. Therefore these two paths have the same endpoints, so

$$\lceil (M'+n)\theta \rceil = \lceil M'\theta \rceil + \lceil n\theta \rceil.$$

(b) $\Leftrightarrow$ (c): Observe that (b) is equivalent to:

(b$'$) There are no lattice points above the line $y = \theta x$ and below the line $y - \lceil M'\theta \rceil = \theta(x - M')$ with $M' \le x \le M$.

By the interpretation of $\Lambda_\theta^{\mathrm{in}}(M)$ as the boundary of a convex hull, condition (c) is equivalent to the following two conditions: (i) there are no lattice points below $\Lambda_\theta^{\mathrm{in}}(M')\Lambda_\theta^{\mathrm{in}}(M-M')$ and above the line $y = \theta x$ with $0 \le x \le M$, and (ii) the edges in the path $\Lambda_\theta^{\mathrm{in}}(M')\Lambda_\theta^{\mathrm{in}}(M-M')$ have monotonically increasing slope. By the definition of $\Lambda_\theta^{\mathrm{in}}(M')$ and $\Lambda_\theta^{\mathrm{in}}(M-M')$, condition (b$'$) is equivalent to condition (i). But condition (i) implies condition (ii). To see this, note that to prove (ii), it is enough to show that slope of the last edge in the path $\Lambda_\theta^{\mathrm{in}}(M')$ does not exceed the slope of the first edge in the path $\Lambda_\theta^{\mathrm{in}}(M-M')$. If this fails, then the fourth vertex of the parallelogram on these two edges is a lattice point of the type ruled out by (i).

(b) $\Rightarrow$ (a): Since (b) implies (c), it follows by replacing M with $M'+n$ that (b) also implies $\Lambda_\theta^{\mathrm{in}}(M'+n) = \Lambda_\theta^{\mathrm{in}}(M')\Lambda_\theta^{\mathrm{in}}(n)$ for all $n = 1,\ldots,M-M'$. By Lemma 7.3, this implies (a).

(c) $\Leftrightarrow$ (d): By Lemma 7.3 and the convexity of $\Lambda_\theta^{\mathrm{in}}(M)$, condition (d) holds if and only if $\Lambda_\theta^{\mathrm{in}}(M')$ is an initial subpath of $\Lambda_\theta^{\mathrm{in}}(M)$. But if the latter holds, then the rest of the path $\Lambda_\theta^{\mathrm{in}}(M)$ is by definition $\Lambda_\theta^{\mathrm{in}}(M-M')$. $\square$

Proof of Proposition 7.22. We will use induction and Lemma 7.27. By symmetry, we can assume that $M_+ \le M_-$. (Otherwise we can replace θ by $-\theta$ and positive ends by negative ends.) By Proposition 7.26, we may further assume that $M_+ < M$. We now proceed in four steps.

Step 1. We begin with some setup and preliminary calculations. Write

$$S = (a_1,\ldots,a_k; a_1',\ldots,a_{k'}' \mid b_1,\ldots,b_l; b_1',\ldots,b_{l'}').$$

Order the a_i's and b_j's according to the standard convention (7.2), and order the a_i''s and b_j''s so that

$$\frac{\lfloor a_i'\theta\rfloor}{a_i'} \ge \frac{\lfloor a_{i+1}'\theta\rfloor}{a_{i+1}'}, \qquad \frac{\lceil b_j'\theta\rceil}{b_j'} \le \frac{\lceil b_{j+1}'\theta\rceil}{b_{j+1}'}. \tag{7.27}$$

For future reference we now compute $\kappa_\theta(S)$. By Lemma 7.4,

$$\kappa_\theta(S) = \lceil M_+\theta\rceil + \lfloor (M-M_+)\theta\rfloor + k' - \lfloor M_-\theta\rfloor - \lceil (M-M_-)\theta\rceil + l'.$$

Since $M_+ < M_-$, by the hypothesis (7.6) and Lemma 7.28(b) this becomes

$$\kappa_\theta(S) = \begin{cases} k', & M_- = M, \\ k'+l'-1, & M_- < M. \end{cases} \tag{7.28}$$

Let m denote the smallest integer such that $\sum_{j=1}^m b_j \ge a_1$. Observe that we must have a strict inequality $\sum_{j=1}^m b_j > a_1$. The reason is that since $M_+ < M$, the hypothesis (7.6) implies that (a_1) is a proper subpartition of $P_\theta^{\mathrm{in}}(M)$, while $(b_1,\ldots,b_m)$ is a subpartition of $P_\theta^{\mathrm{out}}(M)$. If these two subpartitions had the same size, then it would follow that $\kappa_\theta(P_\theta^{\mathrm{in}}(M) \mid P_\theta^{\mathrm{out}}(M)) \ge 2$, contradicting Lemma 7.4.

Next define $\overline{b} := \sum_{j=1}^m b_j - a_1$ and

$$\overline{S} := (a_2,\ldots,a_k; P^{\mathrm{out}}(M-M_+) \mid \overline{b}, b_{m+1},\ldots,b_l; P_\theta^{\mathrm{in}}(M-M_-)).$$

By Lemma 7.27(c),(d),

$$\overline{S} = (P^{\mathrm{in}}(M_+ - a_1); P^{\mathrm{out}}(M - M_+) \mid P^{\mathrm{out}}(M_- - a_1); P^{\mathrm{in}}(M - M_-)).$$

Moreover, the hypothesis (7.6) still holds when (M, M_+, M_-) are replaced by $(M - a_1, M_+ - a_1, M_- - a_1)$. The strategy of the induction will be to deduce the conclusions of the proposition for S from those for $\overline{S}$.

Step 2. We now show that if $J \subset \{1, \ldots, k\}$ and $J' \subset \{1, \ldots, l'\}$ satisfy

$$\sum_{j \in J} b_j + \sum_{j \in J'} b'_j \geq a_1, \tag{7.29}$$

then $\{1, \ldots, m\} \subset J$.

To prove this, first note that by the hypothesis (7.6) and Lemma 7.28(a),

$$P_\theta^{\mathrm{out}}(M) = (b_1, \ldots, b_l) \cup P_\theta^{\mathrm{out}}(M - M_-).$$

It follows that the b_j's for $j \in J$, together with $P_\theta^{\mathrm{out}}(M - M_-)$, comprise a subpartition of $P_\theta^{\mathrm{out}}(M)$. By (7.29), the sum of the numbers in this subpartition is

$$\sum_{j \in J} b_j + (M - M_-) \geq \sum_{j \in J} b_j + \sum_{j \in J'} b'_j \geq a_1. \tag{7.30}$$

By Lemma 7.27(a), this subpartition must contain the minimal initial segment of $P_\theta^{\mathrm{out}}(M)$ whose sum is at least a_1. By (7.6), this initial segment is $(b_1, \ldots, b_m)$.

Step 3. We claim now that if $\{S_\nu\}$ is a θ-decomposition of S (see Definition 1.16), reordered so that $1 \in I_1$ if $k > 0$, then it must have the following properties:

(i) If $k > 0$, then $I_1 = \{1, \ldots, k\}$; $I'_1 = \{i\}$ for some i with $a'_i = a'_1$; $J_1 = \{1, \ldots, q\}$ for some q; and $J'_1 = \emptyset$.

(ii) For all $\nu > 1$ (and also for $\nu = 1$ if $k = 0$), we have

$$|I_\nu| + |I'_\nu| = |J_\nu| + |J'_\nu| = 1.$$

(iii) If $M_- < M$, then there exists ν such that $I_\nu = J_\nu = \emptyset$.

We prove this claim by induction on k.

(Base case.) Suppose that $k = 0$ and let $\{S_\nu\}$ be a θ-decomposition of S. Since $k = 0$, the set I_ν is empty for each ν. Since ν runs from 1 to $\kappa_\theta(S)$, and since I'_ν is nonempty for each ν by the sum condition (1.9), it follows that $\kappa_\theta(S) \leq k'$. We then deduce from equation (7.28) that $l' \leq 1$ and $\kappa_\theta(S) = k'$, so $|I'_\nu| = 1$ for each ν.

By the hypothesis (7.6) and Lemma 7.28(c), we have

$$\Lambda_\theta^{\text{out}}(M) = \Lambda_\theta^{\text{out}}(M_-)\Lambda_\theta^{\text{out}}(M - M_-).$$

Therefore $l \le k'$, and $b_j = a'_j$ for all $j = 1, \ldots, l$. Recall from §7.1 that the ordering convention (7.27) implies that $a'_i \ge a'_{i+1}$ for all i. Now consider the ν for which $1 \in J_\nu$. Since $|I'_\nu| = 1$, by the sum condition (1.9) we must have $I'_\nu = \{i\}$ where $a'_i = a'_1$, and therefore $J_\nu = \{1\}$ and $J'_\nu = \emptyset$. Continuing by induction, the θ-decomposition can be reordered so that $J_\nu = \{\nu\}$ and $J'_\nu = \emptyset$ for $\nu = 1, \ldots, l$.

If $l' = 0$, then we have described all of $S_1, \ldots, S_\nu$. If $l' = 1$, then the description of $\{S_\nu\}$ is completed by noting that under the above reordering, $J_{l+1} = \emptyset$ and $J'_{l+1} = \{1\}$. Now points (i)–(iii) follow immediately from the above description of $\{S_\nu\}$.

(Induction step.) Suppose $k > 0$ and assume that the claim holds for $k - 1$. To carry out the induction we will relate θ-decompositions of S to θ-decompositions of $\overline{S}$. By equation (7.28), $\kappa_\theta(S) = \kappa_\theta(\overline{S})$. Thus we can identify a θ-decomposition of $\overline{S}$ with a decomposition

$$\begin{aligned}
\{2, \ldots, k\} &= \overline{I}_1 \sqcup \cdots \sqcup \overline{I}_{\kappa_\theta(S)},\\
\{1, \ldots, k'\} &= \overline{I}'_1 \sqcup \cdots \sqcup \overline{I}'_{\kappa_\theta(S)},\\
\{m, \ldots, l\} &= \overline{J}_1 \sqcup \cdots \sqcup \overline{J}_{\kappa_\theta(S)},\\
\{1, \ldots, l'\} &= \overline{J}'_1 \sqcup \cdots \sqcup \overline{J}'_{\kappa_\theta(S)},
\end{aligned}$$

such that for each $\nu = 1, \ldots, \kappa_\theta(S)$, the data set $\overline{S}_\nu$ satisfies the sum condition (1.9). Here $\overline{S}_\nu$ is defined as in (1.12), but with b_m replaced by $\overline{b}$.

Given a θ-decomposition $\{\overline{S}_\nu\}$ of $\overline{S}$, reorder the θ-decomposition so that $m \in \overline{J_1}$. We can then define a θ-decomposition $\{S_\nu\}$ of S by setting

$$I_1 := \{1\} \cup \overline{I}_1, \qquad I'_1 := \overline{I}'_1, \qquad J_1 := \{1, \ldots, m-1\} \cup \overline{J}_1, \qquad J'_1 := \overline{J}'_1$$

and leaving the components of the θ-decomposition for $\nu = 2, \ldots, \kappa_\theta(S)$ unchanged. It follows from Step 2 that every θ-decomposition of S is obtained this way from a θ-decomposition of $\overline{S}$. Points (i)–(iii) for θ-decompositions of S then follow from points (i)–(iii) for θ-decompositions of $\overline{S}$. Note that Lemma 7.27(d) gurarantees that when $k = 1$, the unique element i of I'_1 will satisfy $a'_i = a'_1$.

Step 4. We now complete the proof of the proposition. Part (a) is an immediate consequence of points (i)–(iii) from Step 3. We now prove part (b) by induction on k. (Part (c) then follows by symmetry.)

If $k = 0$ then

$$S = (; P_\theta^{\text{out}}(M) \mid P_\theta^{\text{out}}(M)).$$

In this case $\kappa_\theta(S) = k' = l$, and a θ-decomposition of S is equivalent to a permutation of $P_\theta^{\mathrm{out}}(M)$ that preserves the sizes of the elements. So it follows immediately from the definition of c_θ that $c_\theta(S) = P_\theta^{\mathrm{out}}(M)!$ as desired.

If $k > 0$, then as in the proof of Proposition 7.26, it follows from assertion (i) of Step 3 and Lemma 7.27(c),(d) that

$$c_\theta(S) = c_\theta(\overline{S}) \cdot \prod_{n=1}^{m} \delta_\theta \left(a_1 - \sum_{j=1}^{n-1} b_j \, , \, b_n \right) .$$

By Lemma 7.27(b), this becomes $c_\theta(S) = c_\theta(\overline{S})$. We are now done by induction. □

References

[1] D. Bar-Natan, *On the Vassiliev knot invariants*, Topology **34** (1995), 423–472.

[2] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888.

[3] F. Bourgeois and K. Mohnke, *Coherent orientations in symplectic field theory*, Math. Z. **248** (2004), 123–146.

[4] D. Dragnev, *Fredholm theory and transversality for noncompact pseudoholomorphic maps in symplectizations*, Comm. Pure Appl. Math **57** (2004), 726–763.

[5] Y. Eliashberg, A. Givental, and H. Hofer, *Introduction to symplectic field theory*, Geom. Funct. Anal. (2000), 560–673.

[6] O. Fabert, *Counting trivial curves in rational symplectic field theory*, in preparation.

[7] A. Floer and H. Hofer, *Coherent orientations for periodic orbit problems in symplectic geometry*, Math. Z. **212** (1993), 13–38.

[8] H. Hofer, *Holomorphic curves and dynamics in dimension three*, Symplectic geometry and topology (Park City, UT, 1997), 35–101, IAS/Park City Math. Ser. **7**, AMS, 1999.

[9] H. Hofer, K. Wysocki, E. Zehnder, *Properties of pseudo-holomorphic curves in symplectizations. II. Embedding controls and algebraic invariants*, Geom. Funct. Anal. **5** (1995), 270–328.

[10] M. Hutchings, *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. **4** (2002), 313–361.

[11] M. Hutchings and M. Sullivan, *The periodic Floer homology of a Dehn twist*, Algebr. Geom. Topol. **5** (2005), 301–354.

[12] M. Hutchings and M. Sullivan, *Rounding corners of polygons and the embedded contact homology of T^3*, Geometry and Topology **10** (2006), 169–266.

[13] M. Hutchings and C. H. Taubes, *Gluing pseudoholomorphic curves along branched covered cylinders II*, arxiv:0705.2074.

[14] P. Kronheimer and T. Mrowka, *Monopoles and three-manifolds*, book in preparation.

[15] P. Ozsváth and Z. Szabó, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. **159** (2004), 1027–1158.

[16] M. Schwarz, *Cohomology operations from S^1 cobordisms in Floer homology*, ETH Zürich PhD thesis, 1995.

[17] R. Siefring, *The relative asymptotic behavior of pseudoholomorphic half-cylinders*, math.SG/0702356.